IBM TotalStorage DS300 and DS400 Best Practices Guide

- SAN management options
- Recommended configurations
- Software and hardware installation details

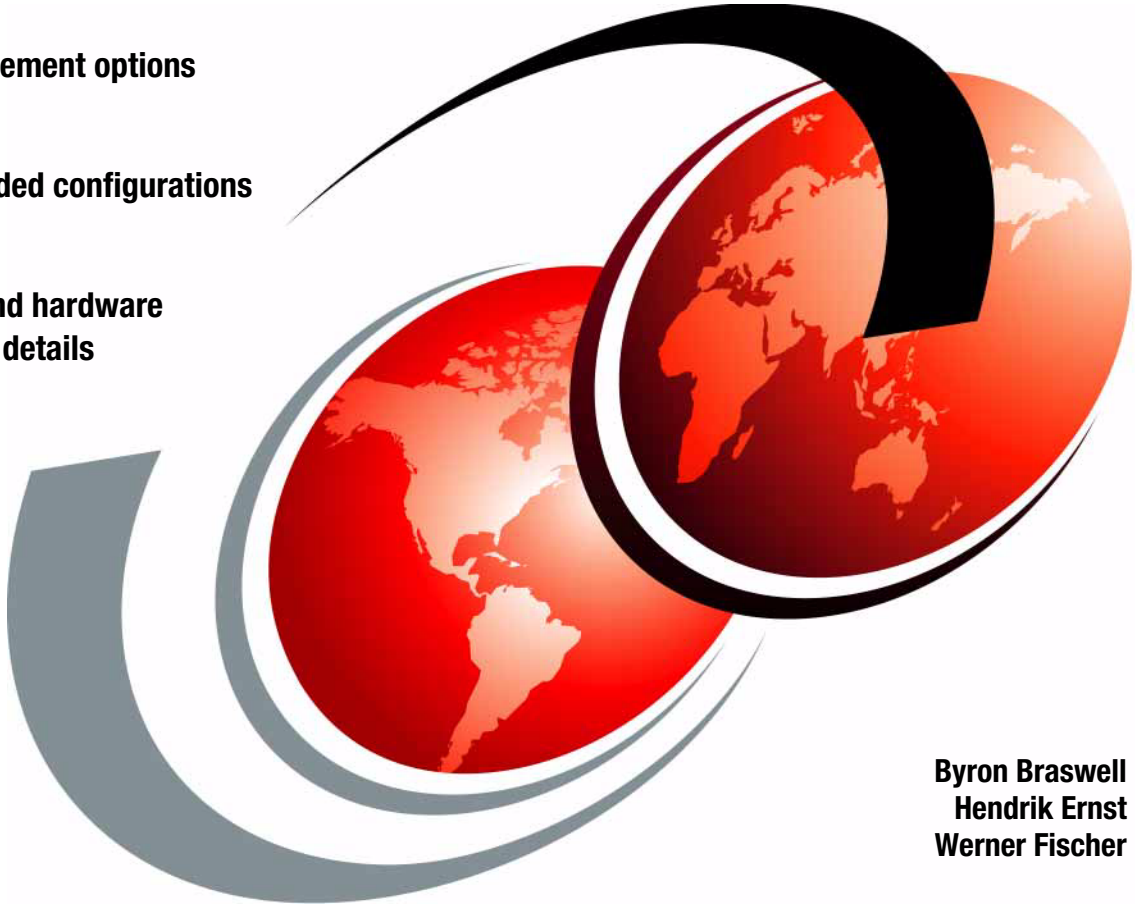

Byron Braswell
Hendrik Ernst
Werner Fischer

ibm.com/redbooks

Redbooks

International Technical Support Organization

IBM TotalStorage DS300 and DS400 Best Practices Guide

February 2006

SG24-7121-00

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

First Edition (February 2006)

This edition applies to Version 8.20 of ServeRAID Manager, and TotalStorage DS300 and DS400 Controller firmware V7.01.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

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

@server®	Enterprise Storage Server®	Notes®
eServer™	Enterprise Systems Connection Architecture®	Redbooks™
xSeries®	ESCON®	Redbooks (logo) ™
z/OS®	FICON®	RACF®
zSeries®	HACMP™	S/390®
AFS®	IBM®	ServerProven®
AIX®	Lotus Notes®	ServeRAID™
BladeCenter®	Lotus®	Tivoli®
DB2 Universal Database™	Micro Channel®	TotalStorage®
DB2®	Netfinity®	
DFS™		

The following terms are trademarks of other companies:

eXchange, IPX, Java, JVM, ONC, RSM, Sun, Ultra, VSM, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Preface

This IBM® Redbook represents a compilation of best practices for configuring the IBM TotalStorage® DS300 (iSCSI) and DS400 (Fibre). These workgroup storage subsystems for IBM @server® xSeries® and BladeCenter® servers are designed to provide an exceptional solution for workgroup storage applications such as file, print and Web serving, as well as, remote boot storage for diskless servers.

The redbook provides a conceptual framework for understanding the DS300 and DS400 in a Storage Area Network and includes recommendations, hints, and tips for the physical installation, cabling, and zoning. It shows:

- Supported configurations
- Proper use of the software and hardware initiators
- Array and logical drive configurations to support maximum throughput
- DSM and MPIO usage
- Understanding failover and failback
- Best practices with installation (for example, FC switches, BladeCenters, Ethernet switches)
- Single to dual upgrade

This book is intended for IBM technical professionals, Business Partners, and customers responsible for the planning, deployment, and maintenance of IBM TotalStorage DS300/DS400 products.

Part 1 is a technical overview of Storage Area Networks, iSCSI, the DS300 and DS400 product family, and TotalStorage management considerations.

Part 2 covers the IBM TotalStorage DS300 Subsystem sample configurations and installation details.

Part 3 covers the IBM TotalStorage DS400 Subsystem sample configurations and installation details.

Sample of callout boxes used in this redbook.

Throughout this redbook, callout boxes are used to illustrate example CLI commands used to perform installation, configuration and maintenance functions.

This redbook applies to Version 8.20 of ServeRAID™ Manager, and TotalStorage DS300 and DS400 Controller firmware V7.01.

The team that wrote this redbook

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

The team - Werner, Byron, Hendrik

Byron Braswell is a Networking Professional at the International Technical Support Organization, Raleigh Center. He received a bachelor's degree in Physics and a master's degree in Computer Sciences from Texas A&M University. He writes extensively in the areas of networking, host integration, and personal computer software. Before joining the ITSO four years ago, Byron worked in IBM Learning Services Development in networking education development.

Hendrik Ernst has worked for Computer Service GmbH in Erfurt since 1998 performing Netfinity® presales support. In 1999, he joined the German Country Postsales Support Team, specializing in the Netfinity Server, IBM @server xSeries, Microsoft® Cluster Services and storage solutions.

Werner Fischer is an IT Specialist in IBM Global Services, Integrated Technology Services in Austria. He has three years of experience in the high availability field. He has worked at IBM for two years, including one year at the EMEA Storage ATS (Advanced Technical Support) in Mainz, Germany. His areas of expertise include planning and implementation of Linux® high availability clusters, SAN disk and tape solutions, and hierarchical storage management environments. Werner holds a master's degree in Computer and

Media Security from the University of Applied Sciences of Upper Austria in Hagenberg where he now also teaches as Assistant Lecturer.

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

Rufus Credle
Margaret Ticknor
Joshua Young
Jeanne Tucker
Tamikia Barrow
David Watts
International Technical Support Organization, Raleigh Center

Shawn Andrews
Khalid Ansari
Roger Bullard
John Fasano
Darryl Gardner
Paul Goetz
Thomas Newsom
Trung V. Nguyen
Scott Partington
Rajesh Sukhramani
David Tareen
Tan T. Trinh
IBM Research Triangle Park

Silvio Erdenberger
Ullrich Hofmann
Sandro Thierbach
Ruediger Zoellner
IBM Germany

Joe Earhart
Cisco Systems

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks™ to be as helpful as possible. Send us your comments about this 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 email to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HZ8 Building 662
  P.O. Box 12195
  Research Triangle Park, NC 27709-2195

Part 1

Technical overview

In Part 1 we give an introduction to Storage Area Networks (SAN) and the TotalStorage DS300 and DS400 products.

1

Storage Area Networks (SAN) overview

What is a SAN? One definition of a SAN is a high-speed network, comparable to a LAN, that allows the establishment of direct connections between storage devices and processors (servers) centralized to the extent supported by the distance of the supported network (Fibre Channel or iSCSI). The SAN can be viewed as an extension to the storage bus concept that enables storage devices and servers to be interconnected using similar elements as in Local Area Networks (LANs) and Wide Area Networks (WANs): routers, hubs, switches, directors and gateways. A SAN can be shared between servers and/or dedicated to one server. It can be local or can be extended over geographical distances.

This chapter describes Storage Area Networks (SAN), what they are and what benefits they provide in the IT environment.

1.1 Storage Area Network benefits

SANs create new methods of attaching storage to servers. These new methods can enable great improvements in both availability and performance. Today's SANs are used to connect shared storage arrays and tape libraries to multiple servers, and are used by clustered servers for failover. They can interconnect mainframe disk or tape to mainframe servers where the SAN switch devices allow the intermixing of open systems (Windows®, AIX®, and so on) and mainframe traffic.

A SAN can be used to bypass traditional network bottlenecks. It facilitates direct, high speed data transfers between servers and storage devices, potentially in any of the following three ways:

- Server to storage: This is the traditional model of interaction with storage devices. The advantage is that the same storage device may be accessed serially or concurrently by multiple servers.
- Server to server: A SAN may be used for high-speed, high-volume communications between servers.
- Storage to storage: This outboard data movement capability enables data to be moved without server intervention, thereby freeing up server processor cycles for other activities like application processing. Examples include a disk device backing up its data to a tape device without server intervention, or remote device mirroring across the SAN.

Why would you want one? SANs allow applications that move data to perform better, for example, by having the data sent directly from source to target device without any server intervention. SANs also enable new network architectures where multiple hosts access multiple storage devices connected to the same network. Using a SAN can potentially offer the following benefits:

- Improvements to application availability: Storage is independent of applications and accessible through multiple data paths for better reliability, availability and serviceability.
- Higher application performance: Storage processing off-loaded from servers and moved onto a separate network.
- Centralized and consolidated storage: Simpler management, scalability, flexibility, and availability.
- Data transfer and vaulting to remote sites: Remote copy of data enabled for disaster protection and against malicious attack.
- Simplified centralized management: Single image of storage media simplifies management.Add text here (Body0).

1.2 Storage Area Network defined

The Storage Network Industry Association (SNIA) defines SAN as a network whose primary purpose is the transfer of data between computer systems and storage elements. A SAN consists of a communication infrastructure, which provides physical connections; and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. The term SAN is usually (but not necessarily) identified with block I/O services rather than file access services.

A SAN can also be a storage system consisting of storage elements, storage devices, computer systems, and/or appliances, plus all control software, communicating over a network.

Note: The SNIA definition specifically does not identify the term SAN with Fibre Channel technology. When the term SAN is used in connection with Fibre Channel technology, use of a qualified phrase such as *Fibre Channel SAN* is encouraged. According to this definition, an Ethernet-based network whose primary purpose is to provide access to storage elements would be considered a SAN. SANs are sometimes also used for system interconnection in clusters.

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

The SAN can be viewed as an extension to the storage bus concept that enables storage devices and servers to be interconnected using similar elements as in Local Area Networks (LANs) and Wide Area Networks (WANs): routers, hubs, switches, directors and gateways. A SAN can be shared between servers and/or dedicated to one server. It can be local or can be extended over geographical distances.

Figure 1-1 on page 6 shows a tiered overview of a Storage Area Network connecting multiple servers to multiple storage system.

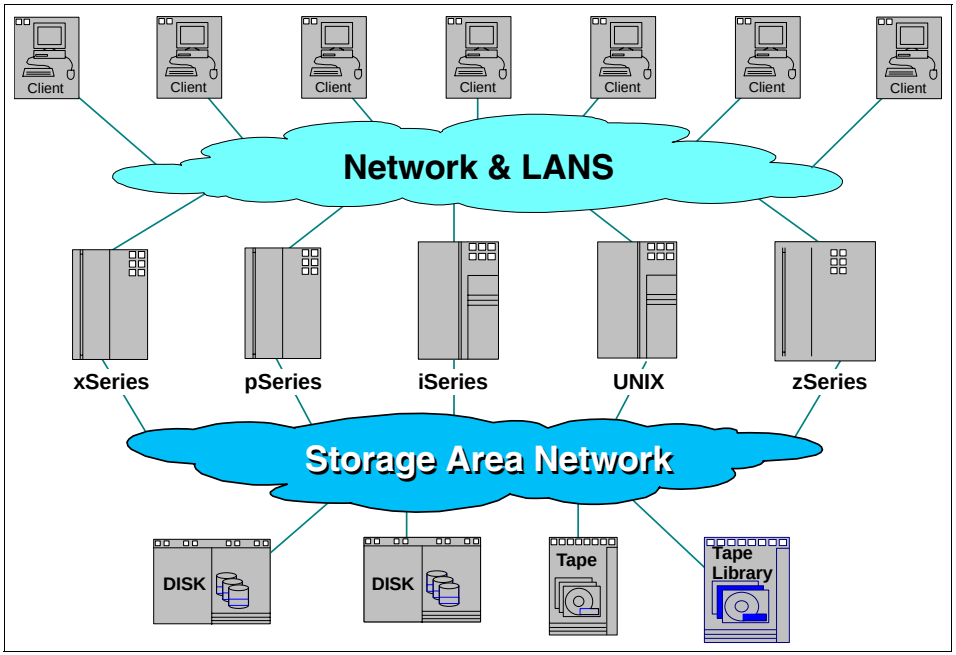

Figure 1-1 A Storage Area Network

SANs create new methods of attaching storage to servers. These new methods can enable great improvements in both availability and performance. Today's SANs are used to connect shared storage arrays and tape libraries to multiple servers, and are used by clustered servers for failover. They can interconnect mainframe disk or tape to mainframe servers where the SAN devices allow the intermixing of open systems (such as Windows, AIX) and mainframe traffic.

1.3 SAN versus Network Attached Storage (NAS)

Network Attached Storage (NAS) is basically a LAN-attached file server that serves files using a network protocol such as Network File System (NFS). NAS is a term used to refer to storage elements that connect to a network and provide file access services to computer systems. An NAS storage element consists of an engine that implements the file services (using access protocols such as NFS or Common Internet File System (CIFS), and one or more devices, on which data is stored. NAS elements may be attached to any type of network. From a SAN perspective, a SAN-attached NAS engine is treated just like any other server.

Figure 1-2 on page 7 illustrates the difference between a SAN and NAS.

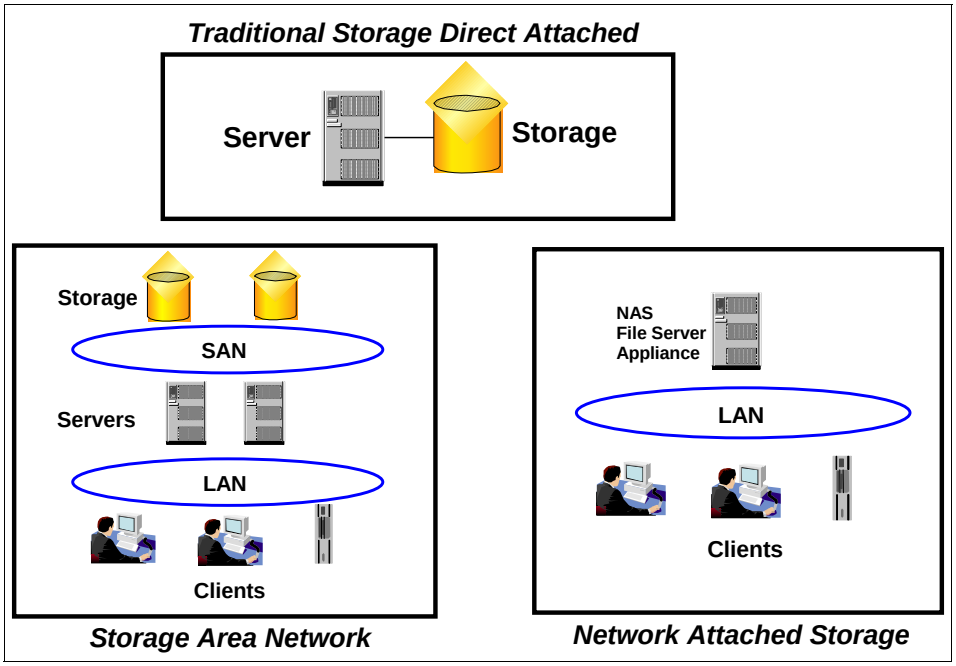

Figure 1-2 SAN versus NAS

1.4 SAN components

Fibre Channel is the predominant architecture upon which most SAN implementations are built, with FICON® as the standard protocol for z/OS® systems, and FCP as the standard protocol for open systems. The SAN components described in the following sections are Fibre Channel-based, and are shown in Figure 1-3 on page 8.

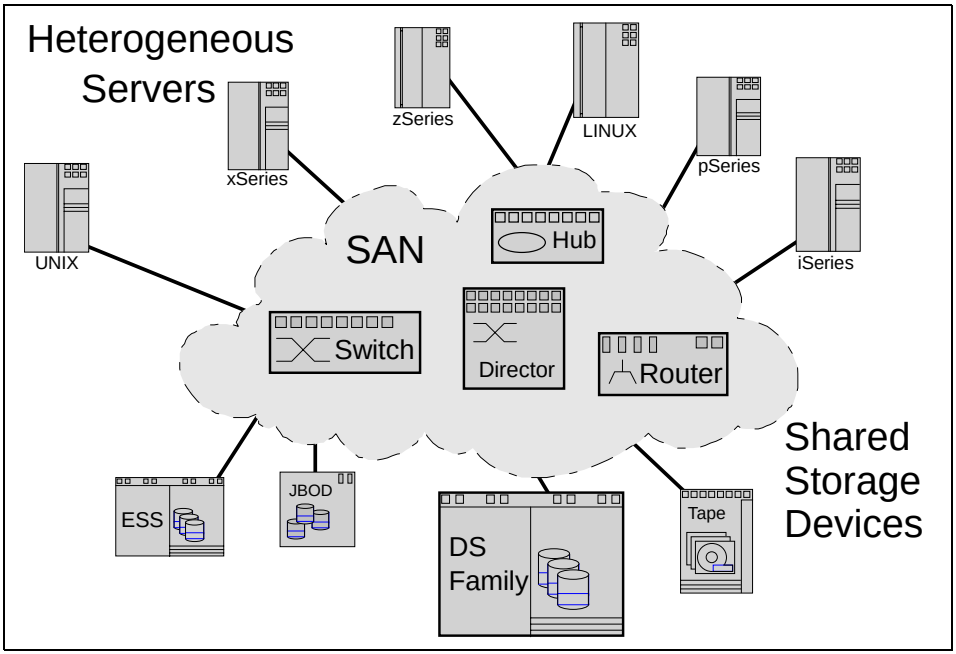

Figure 1-3 SAN components

1.4.1 SAN servers

The server infrastructure is the underlying reason for all SAN solutions. This infrastructure includes a mix of server platforms such as Windows, UNIX® (and its various flavors), and z/OS. With initiatives such as server consolidation and e-business, the need for SANs will increase, making the importance of storage in the network greater.

1.4.2 SAN storage

The storage infrastructure is the foundation on which information relies, and therefore must support a company's business objectives and business model. In this environment simply deploying more and faster storage devices is not enough. A SAN infrastructure provides enhanced network availability, data accessibility, and system manageability. It is important to remember that a good SAN begins with a good design. This is not only a maxim, but must be a philosophy when we design or implement a SAN.

The SAN liberates the storage device so it is not on a particular server bus, and attaches it directly to the network. In other words, storage is externalized and can be functionally distributed across the organization. The SAN also enables the

centralization of storage devices and the clustering of servers, which has the potential to make for easier and less expensive centralized administration that lowers the total cost of ownership (TCO).

1.4.3 SAN interconnects

The first element that must be considered in any SAN implementation is the connectivity of storage and server components typically using Fibre Channel. SANs, like LANs, interconnect the storage interfaces together into many network configurations and across long distances.

Much of the terminology used for SAN has its origins in IP network terminology. In some cases, the industry and IBM use different terms that mean the same thing, and in some cases, mean different things.

1.5 Additional information

For more information, see the redbook *Introduction to Storage Area Networks*, SG24-5470.

2

iSCSI overview

Internet SCSI (iSCSI) is defined in IETF RFC 3720. It allows SCSI block I/O protocols (commands, sequences and attributes) to be sent over a network using the TCP/IP protocol. This is similar to the way SCSI commands are already mapped to Fibre Channel, parallel SCSI, and SSA media. Do not confuse this with the SCSI cabling transport mechanism. We are talking about protocols.

This chapter provides a high level overview of iSCSI.

2.1 What Internet SCSI (iSCSI) is

Internet SCSI (iSCSI) is a transport protocol that carries SCSI commands from an initiator to a target. It is a data storage networking protocol that transports standard Small Computer System Interface (SCSI) block I/O protocol requests (commands, sequences and attributes) over the standard Transmission Control Protocol/Internet Protocol (TCP/IP) networking technology. SCSI data and commands are encapsulated in TCP/IP packets. iSCSI enables the implementation of IP-based Storage Area Networks (SANs), enabling customers to use the same networking technologies—from the box level to the Internet—for both storage and data networks. Taking this approach enables users to get full access to IP management and routing, and, as it uses TCP/IP, iSCSI is also well suited to run over almost any physical network. By eliminating the need for a second network technology just for storage, iSCSI can lower the costs of deploying networked storage and increase its potential market. ISCSI is a native IP interconnect that wraps SCSI data and commands in TCP/IP packets. The receiving device takes the command out of the IP packet and passes it to the SCSI controller, which forwards the request to the storage device. Once the data is retrieved, they are again wrapped in an IP packet and returned to the requesting device.

The main advantages of iSCSI are, when compared to a Fibre Channel installation, it is a low cost implementation, and there are no distance limitations.

There are some drawbacks to iSCSI. An IP network is based on design considerations different from those of storage concepts. The TCP/IP protocol is software-based and geared towards unsolicited packets, whereas storage protocols are hardware based and use solicited packets. Latency, introduced when putting storage on the same route as other network traffic, and security have been concerns for iSCSI. Additional latency is introduced because of the greater amount of processing to be done within the iSCSI chip to support TCP.

2.2 iSCSI components

iSCSI consists of several parts to provide a total solution. Figure 2-1 on page 13 shows the different pieces that may appear in an iSCSI environment.

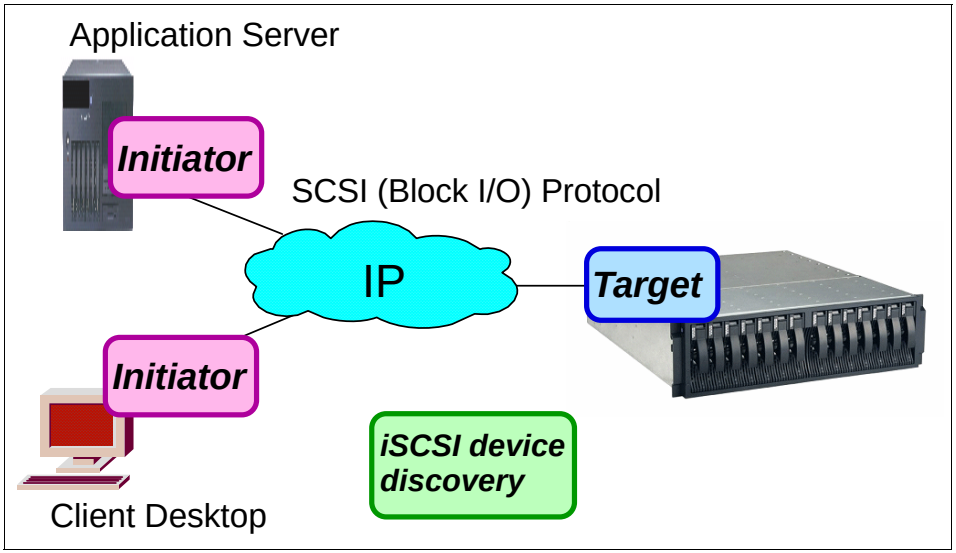

Figure 2-1 iSCSI components

- **iSCSI Initiator:** This consists of either software drivers or hardware adapters that initiate the SCSI request over IP to a target server. iSCSI initiators are available in software or hardware varieties. These drivers coexist with existing SCSI devices without disruption.
- **iSCSI Target:** This software manages the SCSI over IP requests on the target device.
- **iSCSI device discovery:** iSCSI device discovery software sits outside the datastream providing TCP/IP fabric services to aid device discovery, storage policy, logical unit number (LUN) virtualization (partitioning, mapping, volume management). See “iSCSI discovery” on page 14 for more information.

All iSCSI sessions are started by a client (application server or client desktop) requesting access to a target device.

2.2.1 iSCSI details

Internet SCSI (iSCSI) is defined in IETF RFC 3720. It allows SCSI block I/O protocols (commands, sequences and attributes) to be sent over a network using the TCP/IP protocol. This is similar to the way SCSI commands are already mapped to Fibre Channel, parallel SCSI, and SSA media. Do not confuse this with the SCSI cabling transport mechanism. We are talking about protocols.

We have discussed in general terms what iSCSI is supposed to be and its components. In an later section, we will compare it to Network Attached Storage

(NAS). The following sections take a little closer look at some technical aspects of iSCSI.

iSCSI technical introduction

As mentioned previously, iSCSI is the mapping of a SCSI remote procedure invocation model on top of the TCP protocol. A SCSI message is encapsulated inside the TCP frame and is called an “iSCSI protocol data unit”, or iSCSI PDU. iSCSI protocol is based on version three of the SCSI (SAM-2) protocol standard.

How does iSCSI keep track?

iSCSI uses multiple number techniques to keep track of commands and data.

- Command numbering is session wide and provides the means for ordered commands.
- Status numbering is per connection and aids recovery from connection failures.
- Data sequencing is per command and detects missing data packets.

Fields in the iSCSI PDU contain numbering information. With unidirectional traffic, special NOP-message PDUs may be utilized to synchronize communications between client and server. The status number is started after login. Data sequencing is required since multiple client-server paths may be used.

iSCSI discovery

Discovery allows an initiator to find the target(s) to which it has access. This requires a minimum of user configuration. Several methods of discovery may be used:

A list of targets at the initiator

An administrator can define the iSCSI targets to the host system initiator. This process allows the administrator to specify the iSCSI target node name and IP address:port to the host system initiator or its host bus adapter (HBA). See 2.6.3, “Hardware initiator” on page 26 for more information about HBAs. iSCSI HBAs should support an administrator defining this information. This type of discovery is useful in small installations.

Queries to known iSCSI servers

An iSCSI initiator can probe its environment and, when a possible iSCSI target is found, start a **`discovery session`** with the target by issuing a **`SendTargets`** command. The target can reply to a **`SendTargets`** command by returning a list of all iSCSI target nodes it knows about.

Queries to an Internet Storage Name Server (iSNS)

The Internet Storage Name Server permits iSCSI targets to register with a central point. The administrator can set up discovery domains so that when a host iSCSI initiator queries the central control point for the locations of iSCSI storage controllers, only the authorized controllers are reported. The iSNS server can be located by one of the following techniques:

- iSCSI initiators multicasting to the iSNS server
- setting the iSNS server IP address in the DHCP server
- setting the iSNS server IP address in the SLP server (see below)
- setting the iSNS server IP address in the iSCSI initiator or target

Service Location Protocol (SLP)

The Service Location Protocol can be used to locate iSCSI target devices. SLP operates with three agents:

- User agent (UA) - works on the client (iSCSI initiator) to help establish contact with a service (iSCSI target). It does this be retrieving information from service agents (SA) or directory agents (DA).
- Service agent (SA) - runs on the iSCSI target device to advertise the service and its capabilities.
- Directory agent (DA) - collects service advertisements from the iSCSI targets.

Starting iSCSI (iSCSI login)

The purpose of login is to establish a TCP connection for iSCSI to use. This process authenticates the iSCSI end stations. The session parameters are set and a security protocol association is created. A TCP connection is mapped to an iSCSI session. iSCSI connections use layer 5 (session layer) of the OSI seven layer reference model. See Figure 2-2.

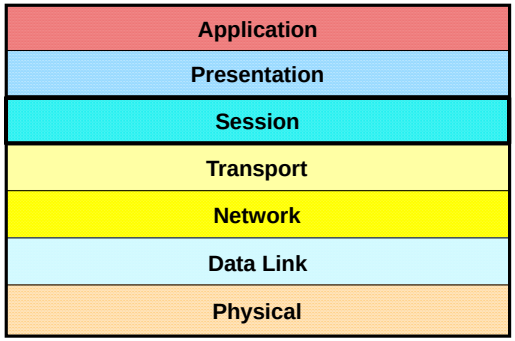

Figure 2-2 Seven layers of the OSI reference model

iSCSI full feature phase

After a successful login process, normal SCSI commands, data and messages can flow. Data and status for a given command must be sent over the same TCP connection. This is called connection allegiance and is based on port and IP number. Ordered command sets may use a different connection.

iSCSI data flows

Outgoing data (initiator to target) can be either solicited or unsolicited. Solicited data must be sent in response to Ready to Transfer PDUs (R2T) from the target. An initiator can send unsolicited data as part of a command or in a separate PDU. All subsequent data PDUs have to be solicited.

Figure 2-3 shows the iSCSI data packet frame.

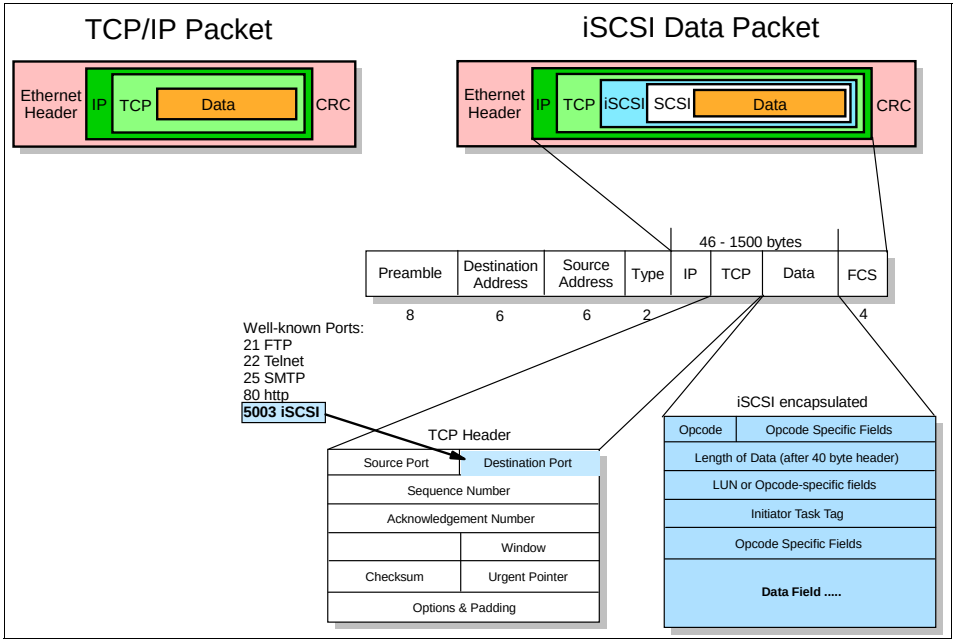

Figure 2-3 Data encapsulation of iSCSI frame

iSCSI connection termination

Connection terminations are exceptional events. Graceful connection shutdowns occur when there are no outstanding tasks with allegiances to the connection. If there are outstanding tasks, or tasks that have not yet sent status information, recovery actions may be needed.

Connection termination is also required as a prelude to recovery. Connection termination can avoid old PDUs from previous connections. (Logout is sent by the initiator.) Targets request the initiator to issue a Logout through an Asynchronous Event PDU.

iSCSI naming

Every initiator and target node must have a single name defined for it. This name is used as part of all sessions established between initiator and target nodes. An iSCSI initiator or target can have one of two types of names:

- iqn (iSCSI qualified name) - its format is iqn plus a date field plus the reverse DNS name with unique qualifiers. For example: `iqn.2005-08.com.ibm.ral.itso.testred`
- eui (enterprise unique identifier) - its format is *eui* plus 16 hex digits. For example: `eui.acdc15882005bdef`

Both types are meant to be long lived as well as unique in the world and independent of physical location. Both have a central naming authority that can ensure their uniqueness.

The DS300 uses iqn names.

The iSCSI name is used as a part of an iSCSI address.

iSCSI addressing

The iSCSI address has the format of:

```
<IP address>[:<port>]/<iSCSI name>
```

The IP address can be either IPv4, IPv6, or the fully qualified domain name. The <port> is optional; it specifies the TCP port that the target is listening for connections on. If it is not used, the well-known iSCSI port (3260) is assumed. The <iSCSI name> is the iqn or eui name of the device. It is optional.

The iSCSI address specifies a single path to an iSCSI target. The iSCSI address is primarily used during discovery. It ensures that the SCSI PDU gets delivered to the correct agent and does NOT affect any SCSI protocol specific addresses.

iSCSI message synchronization

iSCSI PDUs are of varying length, but TCP does not have a mechanism to identify these message boundaries within the TCP layer. iSCSI places message length information in the iSCSI header to delineate the end of the current message and the start of the next message. A potential problem still exists if iSCSI PDUs are delivered out of sequence or are missing.

There are several schemes that can be used to delineate the starting points of various messages. One of the easiest to implement is by using fixed interval markers. A special marker frame with a 32-bit pointer to the next iSCSI PDU start is used.

iSCSI synch and steering

An optional layer can be inserted between the iSCSI and TCP layers called “synch and steering”. This layer retains the ending address for every delivered iSCSI PDU. It includes the minimum header information as part of its encapsulation process. The information is used to deliver PDU contents to a final address and aid with recovery operations if TCP frames are lost during transmission.

2.3 iSCSI SAN and Network Attached Storage (NAS)

In this section, we give a brief comparison of iSCSI SAN and Network Attached Storage. Figure 2-4 illustrates a comparison between iSCSI, Fibre Channel, and NAS.

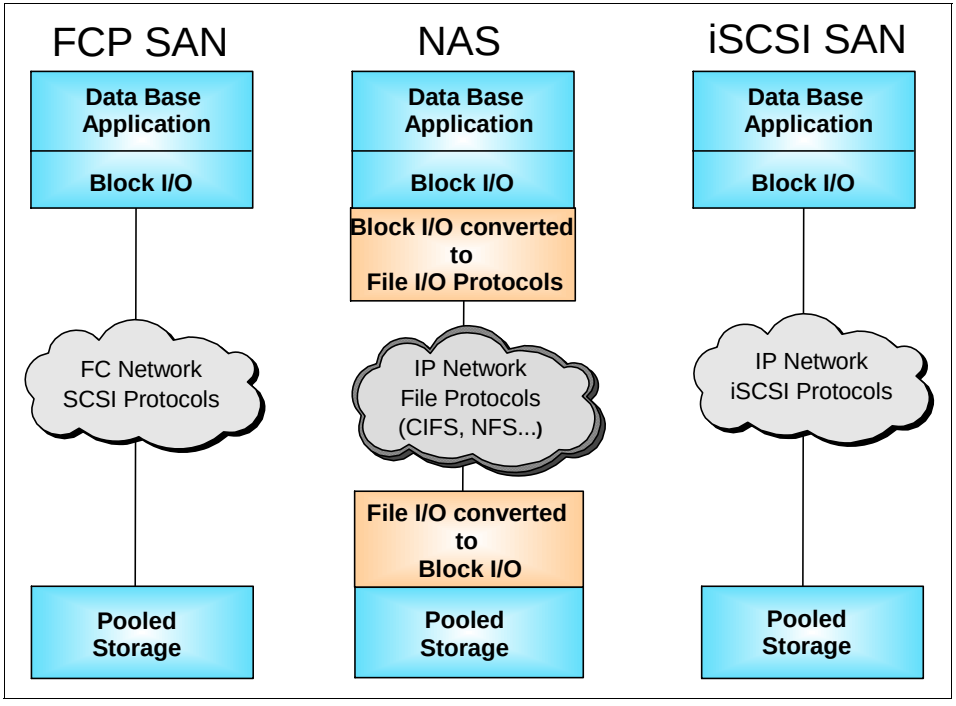

Figure 2-4 SAN/NAS comparison

2.3.1 Network Attached Storage (NAS)

Storage devices which optimize the concept of file sharing across the network have come to be known as NAS. NAS solutions utilize the mature Ethernet IP network technology of the LAN. Data is sent to and from NAS devices over the LAN using TCP/IP protocol. By making storage devices LAN addressable, the storage is freed from its direct attachment to a specific server, and any-to-any connectivity is facilitated using the LAN fabric. In principle, any user running any operating system can access files on the remote storage device. This is done by means of a common network access protocol, for example, Network File System (NFS) for UNIX servers, and Common Internet File System (CIFS) for Windows servers. In addition, a task, such as back-up to tape, can be performed across the LAN, using software like Tivoli® Storage Manager (TSM), enabling sharing of expensive hardware resources, such as automated tape libraries, between multiple servers. A storage device cannot just attach to a LAN. It needs intelligence to manage the transfer and the organization of data on the device. The intelligence is provided by a dedicated server to which the common storage is attached. It is important to understand this concept. NAS comprises a server, an operating system, plus storage which is shared across the network by many other servers and clients. So an NAS is a *device*, rather than a *network infrastructure*, and shared storage is attached to the NAS server.

One of the key differences of an NAS disk device, compared to Direct Attached Storage or other network storage solutions, such as FC SAN or iSCSI SAN, is that all I/O operations use file level I/O protocols. File I/O is a high level type of request that, in essence, specifies only the file to be accessed, but does not directly address the storage device. This is done later by other operating system functions in the remote NAS appliance. A file I/O specifies the file. It also indicates an offset into the file. For instance, the I/O may specify “Go to byte ‘1000’ in the file, (as though the file was a set of contiguous bytes), and read the next 256 bytes beginning at that position”. Unlike block I/O, there is no awareness of a disk volume or disk sectors in a file I/O request. Inside the NAS appliance the operating system keeps tracks of where files are located on disk. The OS issues a block I/O request to the disks to fulfill the file I/O read and write requests it receives.

In summary, the network access methods, NFS, CIFS, and NetWare, can only handle File I/O requests to the remote file system. This is located in the operating system of the NAS device. I/O requests are packaged by the initiator into TCP/IP protocols to move across the IP network. The remote NAS file system converts the request to block I/O and reads or writes the data to the NAS disk storage. To return data to the requesting client application the NAS appliance software re-packages the data in TCP/IP protocols to move it back across the network. A database application which is accessing a remote file located on an NAS device,

by default, is configured to run with File System I/O. It cannot utilize a “raw I/O” to achieve improved performance.

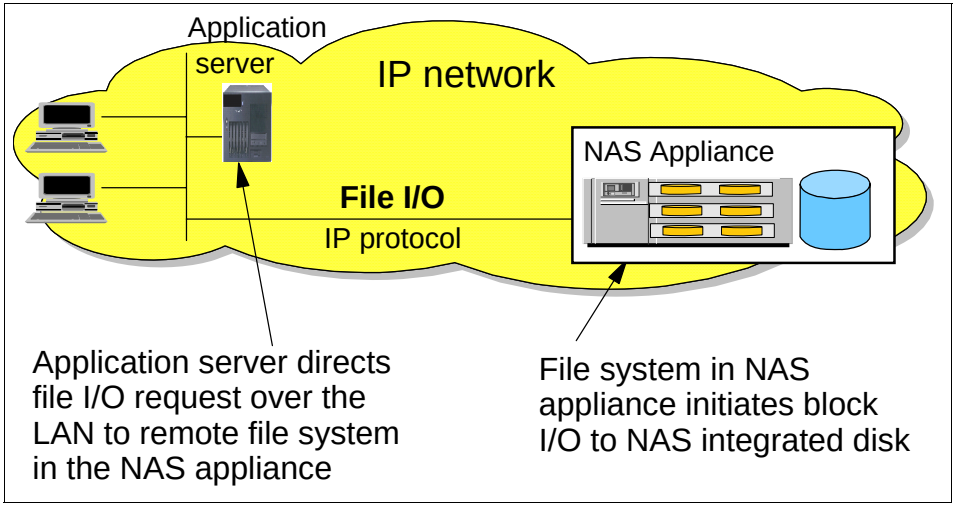

Figure 2-5 NAS uses file I/O

NAS offers a number of benefits, which address some of the limitations of directly attached storage devices and overcome some of the complexities associated with SANs.

- **Resource pooling**: An NAS appliance enables disk storage capacity to be consolidated and pooled on a shared network resource, at great distance from the clients and servers which will share it. Therefore, an NAS device can be configured as one or more file systems, each residing on specified disk volumes. All users accessing the same file system are assigned space within it on demand. This contrasts with individual DAS storage, when some users may have too little storage, and others may have too much. Consolidation of files onto a centralized NAS device can minimize the need to have multiple copies of files spread across distributed clients. Therefore, overall hardware costs can be reduced. NAS pooling can reduce the need physically to reassign capacity among users. The results can be lower overall costs through better utilization of the storage, lower management costs, increased flexibility, and increased control.
- **Exploits existing infrastructure**: Because NAS utilizes the existing LAN infrastructure, there are minimal costs of implementation. Introducing a new network infrastructure, such as a Fibre Channel SAN, can incur significant hardware costs. In addition new skills must be acquired, and a project of any size will need careful planning and monitoring to bring to completion.

- **Simple to implement**: Because NAS devices attach to mature, standard LAN infrastructures, and have standard LAN addresses, they are, typically, extremely easy to install, operate and administer. This plug and play operation results in low risk, ease of use, and fewer operator errors, so it contributes to a lower cost of ownership.

2.3.2 iSCSI

As mentioned in 2.1, “What Internet SCSI (iSCSI) is” on page 12, iSCSI is a network transport protocol for SCSI that operates on top of TCP. iSCSI encapsulates SCSI protocols into a TCP/IP frame, so that storage Controllers can be attached to IP networks.

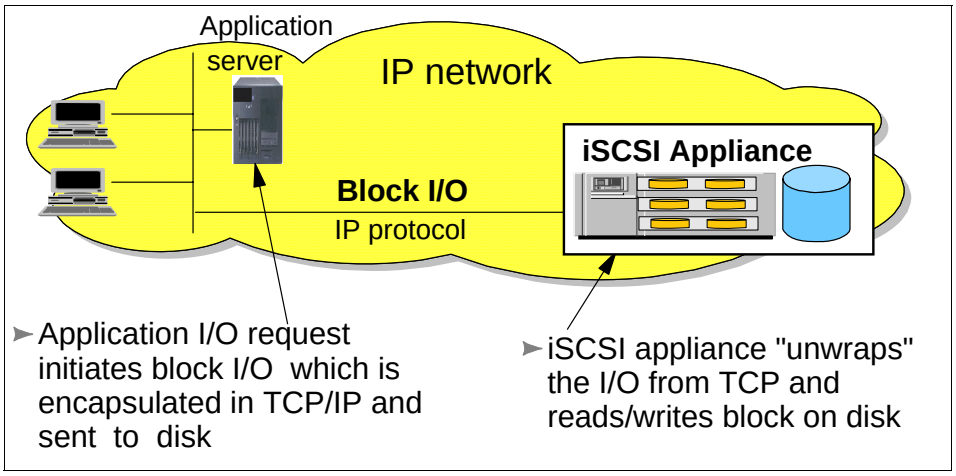

Figure 2-6 iSCSI uses block I/O

2.4 iSCSI benefits

The following topics discuss a few of the benefits that come with iSCSI storage connections.

2.4.1 Connectivity

iSCSI can be used for direct attached storage (DAS) or SAN connections. iSCSI capable devices could be placed on an existing LAN (shared with other applications) in a similar way to NAS devices. Also, iSCSI capable devices can be attached to a LAN which is dedicated to storage I/O (in other words an IP SAN), or even on a LAN connected to only one processor (like a DAS). These options are shown in Figure 2-7 on page 22.

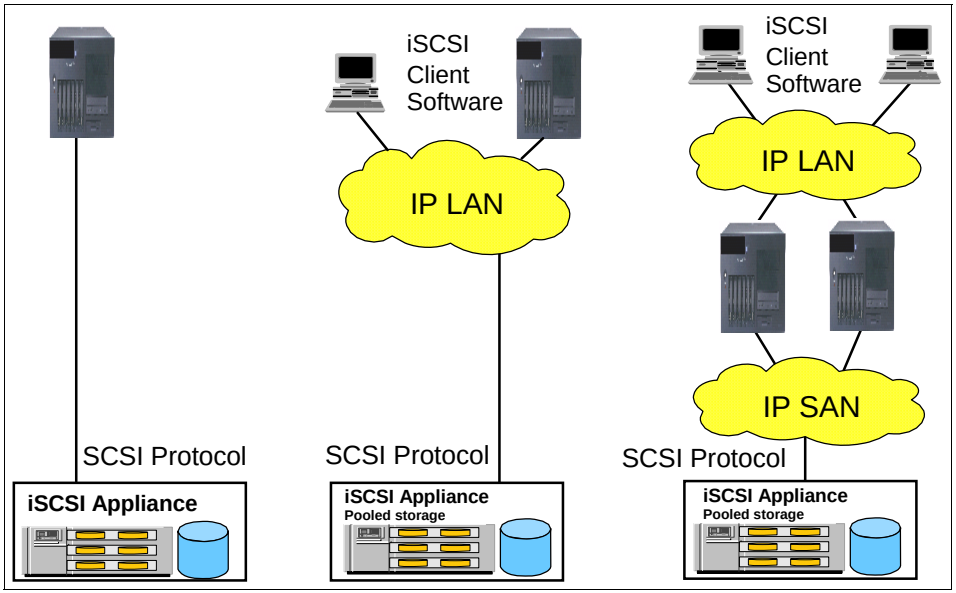

Figure 2-7 iSCSI appliance connection options

2.4.2 Extended distance

IP networks offer the capability easily to extend beyond the confines of a LAN, to include Metropolitan and Wide Area Networks (MANs and WANs). This gives greater flexibility, and at far less cost and complexity, compared to the inter connection of Fibre Channel SANs over wide areas.

2.4.3 Media and network attachments

iSCSI and NAS devices both attach to IP networks. This is attractive compared to Fibre Channel, because of the widespread use of IP networks. IP networks are already in place in most organizations and are supported by existing IT skills. TCP/IP-based networks can potentially support much longer distances than can pure Fibre Channel SANs, promising the possibility of scaling IP SANs to Storage Wide Area Networks (SWAN). Purchasing of attachment technologies is simplified, as they are the same as for LAN attachments.

2.4.4 Backup

Backup of data on the IP Storage is the same as for any direct attach storage; that is, via any method that supports SCSI-attached volumes. A backup application running on an external server, including the one hosting the iSCSI initiator code, will control the backup of data that is physically stored on the iSCSI appliance.

An NAS appliance, because it “hides” disk volumes from its clients, and often includes specialized backup facilities, may be easier to install and manage.

2.4.5 Management

iSCSI is managed like any direct-attach SCSI device. iSCSI-connected disk volumes are visible to attached processors. Compared to Fibre Channel SANs, iSCSI benefits from using IP networks for which there are established network management tools and people skills. Such tools enable network administrators to coordinate provision of bandwidth among users and applications, traffic management, and overall network operations. Training in new networking skills is minimized.

2.5 iSCSI planning considerations

Since the iSCSI appliance attaches to the existing Ethernet network, implementation is very similar to an NAS device. In this way, NAS and iSCSI are very similar. However, when you make the comparisons between iSCSI, Fibre Channel and NAS, it really comes down to Block I/O versus File I/O. Both iSCSI and Fibre Channel use Block I/O to transport data, whereas NAS uses File I/O. The determining factor might ultimately be what kind of application is to be used. Will the application work better in a Block I/O environment, or File I/O environment? Is the application database or file sharing oriented? How much does performance play in this application?

All File I/Os result at a lower layer in Block I/Os. This means that iSCSI can also support File I/O applications. However, if “visibility” and “sharing” of files are needed, File I/O applications are better supported by NAS than by iSCSI devices. Actual device sharing can only happen at the Block I/O level. This requires a “share-aware” file system. Windows NTFS does not support this.

Fibre channel SAN offers better performance, but is more expensive and requires a higher skill set to implement. iSCSI and NAS offer better pricing and skills may already be in place to implement them. However, both Fibre Channel and iSCSI offer the performance benefit of Block I/O. These are all considerations that must be taken into account when making the decision to add storage to the current environment.

2.5.1 Network congestion

In the case of NAS, and backup over IP networks, congestion may cause variable performance for other applications. This also remains a concern for iSCSI implementations. This can be alleviated with high speed Gigabit Ethernet, and largely overcome or masked by higher speed network transmission, such as 10 GBps.

2.5.2 Performance

A performance comparison is difficult to generalize because there are so many variables. That said, Fibre Channel at 100 MBps (1 Gigabit/second) is generally more efficient for I/O traffic than TCP/IP over Ethernet at equivalent bandwidth. iSCSI performs better than NAS (when both are on Ethernet) due to reduced protocol overhead. This is because it handles SCSI directly, rather than translating between file-I/O protocols and SCSI.

TCP/IP is a software intensive network design which requires significant processing overhead. This overhead can consume a substantial proportion of available processor cycles when handling Ethernet connections. This is a drawback for performance intensive storage applications.

A performance consideration, therefore, is the impact of the software protocol stack on processor utilization. Fibre Channel SANs support SCSI commands mapped directly to Fibre Channel media, and processor overhead for this mapping is low. In iSCSI, handling of the TCP/IP protocol requires processor cycles at both ends.

2.6 iSCSI initiators

As stated previously, iSCSI initiators consist of either software drivers or hardware that initiate the SCSI requests over the IP network to a target device. See 3.7, “iSCSI software and hardware” on page 58 for information specific to the DS300. There are three ways that an iSCSI initiator can be implemented in hardware and software. Each of the three approaches differs in price, performance and capabilities.

2.6.1 Software initiator

Many operating systems have iSCSI software initiators. An iSCSI software initiator is an iSCSI driver that works with the OS TCP/IP stack, network drivers, and NICs to provide iSCSI connectivity to other iSCSI devices via the IP network. Since an iSCSI software initiator depends on the OS IP stack, if the IP stack fails or is taken down by a user, access to the remote disk is lost. As a result, software initiators are not ideal for usage to boot the server up. See Figure 2-8.

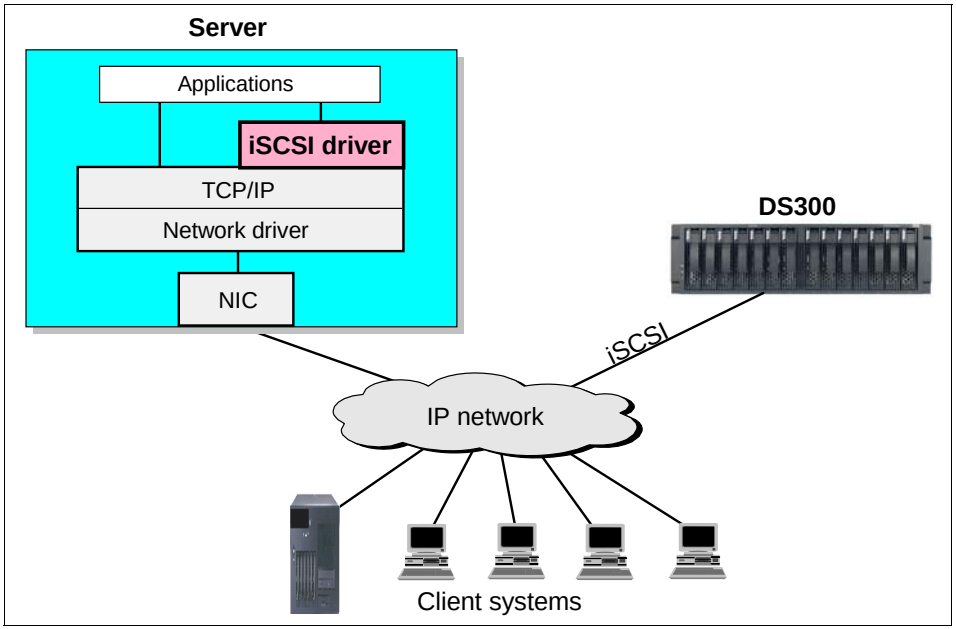

Figure 2-8 iSCSI software initiator

An iSCSI software initiator is an inexpensive alternative to an iSCSI HBA. It is used to provide iSCSI support via software instead of a separate iSCSI HBA.

An iSCSI software initiator adds additional workload to the server processor, TCP/IP stack and network driver to handle the additional SCSI protocol and data traffic. Since the host CPU is responsible for processing TCP/IP requests, iSCSI can suffer performance degradation - especially in high traffic settings. This performance limitation is especially dramatic when compared with Fibre Channel, which does not have TCP/IP overhead. However, iSCSI software initiators are a potential fit for casual demands for storage access.

One method to address the performance problem is to increase the speed of your host processor. Another method is to use a TCP/IP Offload Engine.

2.6.2 TCP/IP offload engine (TOE)

When using a TCP/IP offload engine, an iSCSI software initiator is still required, but the TCP/IP processing along with the network driver are handled on a special NIC. TOE is a hardware-based solution that removes the burden of IP processing from the CPU on a server and moves it down onto the NIC. Data is written directly to the NIC and it handles the IP processing necessary to transmit and receive on the network. This implementation is faster than the previously mentioned iSCSI software implementation since the host CPU does not have to process the iSCSI IP packets. It has the added benefit that the TOE functionality may potentially be used for all network traffic, not merely the iSCSI portions of it.

Note: A TOE may not allow you to run your standard TCP/IP connections off the NIC unless the vendor has provided some type of filter driver to intercept the standard TCP/IP requests.

2.6.3 Hardware initiator

The iSCSI initiator or driver can be implemented in a hardware adapter card, rather than in the host software. This can be done by using an iSCSI host bus adapter (HBA). The iSCSI processing is offloaded to the hardware adapter card instead of processing the iSCSI protocol in the host software. iSCSI TCP/IP processing is also offloaded to the TOE on the HBA. With both TCP and iSCSI processing on the adapter card, high-speed transport of block data with minimal CPU overhead is possible. This is the most expensive of the three iSCSI initiator options as it require the purchase of an iSCSI HBA, but it is the most capable and the best performer. All of the SCSI block-processing and TOE functions are integrated into the HBA. This frees the host CPU from having to do any of the iSCSI processing.

3

Components overview

This chapter describes the components of the IBM TotalStorage DS300 and DS400 disk subsystems along with the associated adapters, switches, etc. that are required to complete an operational SAN environment.

- TotalStorage DS300 disk subsystem
- TotalStorage DS400 disk subsystem
- EXP400 expansion unit
- iSCSI components
- Fibre Channel components
- Advanced TotalStorage DS300 and DS400 features

3.1 Overview

The IBM TotalStorage DS300 and DS400 are entry-level, low cost workgroup storage subsystems for xSeries and BladeCenter servers. The DS300 and DS400 offer a solution for workgroup storage applications, such as file, print and Web serving, as well as remote boot storage for diskless servers.

The main difference between the two products is that the DS300 is designed to use iSCSI to connect over an existing Ethernet infrastructure, whereas the DS400 is capable of using Fibre Channel connectors. However, a dedicated Ethernet infrastructure can also be used if desired for the DS300 to isolate the iSCSI SAN traffic. The DS400 is also capable of supporting more disk capacity than the DS300, through the use of EXP400 expansion units and both support Flashcopy.

The modular design of the DS300 and DS400 allows you to easily perform upgrades and additions to help meet the growing needs of your application environment—scaling to support increasing volumes of data flexibly and affordable. In addition, advanced software features on the DS300 and DS400 -including access control lists and online array expansion are designed to give administrators the power to configure storage according to changing usage needs. These features also help enable you to share storage across multiple application servers.

Both units are designed to deliver advanced functionality for business continuance and disaster recovery at an entry-level price. Using a rack mountable 3U enclosure with 14 SCSI drives and redundant design, hot-swap power and cooling modules, the DS300/400 features RAID reliability and high availability software to help maintain operations. The enclosure supports either one or two controllers for high-availability configurations. Certain models of the DS300 controller are designed to support up to three 1 Gbps Ethernet ports for management and redundant data paths so that—even if a line or a controller fails—multiple paths between servers and storage can help maintain a connection.

The DS300/400 also offers optional advanced features such as space efficient IBM Flashcopy solutions, which are designed to help reduce storage backup windows and improve storage utilization for Microsoft Exchange, Microsoft SQL Server 2000 and Lotus® Notes® applications.

The DS300 and DS400 are designed to offer excellent storage management and configuration capabilities through IBM ServeRAID Manager (the same software used on the xSeries servers for RAID management), helping enable clients to manage multiple IBM TotalStorage DS300 and DS400 subsystems, as well as direct-attached internal ServeRAID controllers, all through a single, easy-to-use

management console with an extensive online help library. Configuration wizards simplify initial storage setup. IBM ServeRAID Manager is designed to support most common operating environments for workgroup applications, including Microsoft Windows, Novell NetWare and Linux platforms.

3.2 Specifications

On both the DS300 and DS400, each controller runs on an Intel® 733 Mhz xScale 80200 Processor. It supports dual 100 Mhz PCI-X, DDR200 DDRAM memory, 32 MB compact flash and Gb ethernet. The second DDRAM memory DIMM is dedicated to the RAID chipset and not accessible by the xScale processor.

Each controller also contains a dual port 82545 Ethernet chipset. Port 1 of the chipset is used for management purposes and can auto detect 10, 100 or 1Gb transfer speeds. Port 1 is listed in the management software as interface0. Port 2 is used for the internal controller communication channel. It is forced to 1Gb and cannot be changed.

I/O data is processed by an Adaptec 7942 chipset which contains its own processor, memory and XOR. It supports RAID levels 0, 1, 10, 5 and 50. It also monitors the battery status of the card. A failed battery status will generate an alert. Cache battery backup is a Li Ion battery that can hold data for at least 72 hours. A DS300 controller has a single Adaptec 7942 chipset while a DS400 controller has two.

There are two separate SCSI buses on the midplane. Each bus is driven by a separate channel of the Adaptec 7942 chipset. Channel 0 contains drives 9 through 15, while channel 1 contains drives 0 through 5 and 8.

Table 3-1 Specifications of IBM TotalStorage DS300 and DS400 disk subsystems

Product	**DS300**	**DS400**
Machine/Models	1701-1RL, 1701-1RS, 1701-2RD	1700-1RS, 1700-2RD
Platform support	Windows 2000, Windows 2003, Linux	Windows 2000, Windows 2003, Linux
Host connectivity	1-2 1GB iSCSI per controller	Two 2GB Fibre Channel per controller
SAN support	Switched Ethernet	Direct, FC-AL, Switched Fabric
Availability features	Fault Tolerant, RAID, Redundant Hotswap Power, Hotswap drives	Fault Tolerant, RAID, Redundant Hotswap Power, Hotswap drives

Product	DS300	DS400
Controller	Single or dual active 1 GB iSCSI RAID Controllers	Single or dual active 2 GB FC RAID Controllers
RAID support	0, 1, 5, 10, 50	0, 1, 5, 10, 50
Capacity (max based on 300 GB Ultra320 HDDs)	min 36 GB, max 4.2 TB 14 HDD max	min 36 GB, max 12 TB with 2 EXP400 Expansion Units 40 HDD max
Drive interface	Ultra320 SCSI	Ultra320 SCSI
Drive support	36 GB, 73 GB, 146 GB, 300 GB 10,000 RPM Disk Drives; 36 GB, 73 GB, 146 GB 15,000 RPM Disk Drives	36 GB, 73 GB, 146 GB, 300 GB 10,000 RPM Disk Drives; 36 GB, 73 GB, 146 GB 15,000 RPM Disk Drives
Logical drives	512 max per subsystem	512 max per subsystem
Logical drives per array (SRM)	64 created using IBM ServeRAID Manager console	64 created using IBM ServeRAID Manager console
Logical drives per array (CLI)[a]	512 created using command line interface	512 created using command line interface
Max array size	2 terabytes	2 terabytes
RAID 0, 1, or 5	14 drives max per array	16 drives max per array
RAID 10 or 50	14 drives max per array	40 drives max per array
SCSI buses	2 internal	2 internal, 2 external
LUNs	32 max per initiator	32 max per initiator
Certifications	Microsoft Windows MSCS	Microsoft Windows MSCS

a. With 512 logical drives in one array, no more arrays can be created because the maximum number of logical drives per subsystem is reached.

3.3 IBM TotalStorage DS300 disk subsystem

Leveraging standard Ethernet infrastructure via iSCSI and low cost SCSI drives, the DS300 is designed to offer a simple, affordable Storage Area Network (SAN) solution that can be used in direct attached or network-attached environments. For even greater flexibility, the DS300 is designed to allow simultaneous support of a range of operating systems for xSeries and BladeCenter servers—in either single or dual-controller configurations.

iSCSI adapters in the servers are used to communicate over an IP network to the DS300 disc subsystem. An alternative to iSCSI hardware adapters in the servers is to use iSCSI software initiators instead. These have the advantage of reduced costs in the server. However, they have the disadvantage of less security and increased server CPU utilization.

The IBM TotalStorage DS300 is available in 3 models as shown in Table 3-2.

Table 3-2 DS300 models

	1701-1RL	**1701-1RS**	**1701-2RD**
iSCSI RAID controller	1	1	2
Controller upgrade	No	Yes	N/A
Controller system memory	512MB max	512MB, 1GB max	512MB, 1GB max
RAID cache (per controller)	none	256MB	256MB
Battery backup for cache	No	Yes	Yes
Flashcopy copy services	4 volumes	256 Flashcopy target logical drives with 1GB system memory	256 Flashcopy target logical drives with 1GB system memory
Gbps Ethernet ports/controller	1	3	3
Power supply w/fan unit	1	2	2
Disk bays	7, 14 max[a]	14	14
LED indicators	8	12	12

a. More than 7 hard disk drives require an optional power supply.

3.3.1 DS300 Model 1701-1RL

Figure 3-1 DS300 Model 1701-1RL - front view

The DS300 storage subsystem 1701-1RL contains one iSCSI RAID controller with only one 1 Gbps Ethernet connection. This connection is used for both iSCSI data traffic and management communication. The controller on the 1701-1RL contains only one 82545 Ethernet chipset. Therefore, the 1701-1RL controller cannot be upgraded to a different DS300 model.

The default configuration for DS300 Model 1701-1RL includes support for up to 7 drives. A total of 14 drives can be installed if the front filler panel is removed.

Important: Support for more than 7 drives requires an additional power supply with fan unit be installed to provide proper temperature control for the additional hard drives.

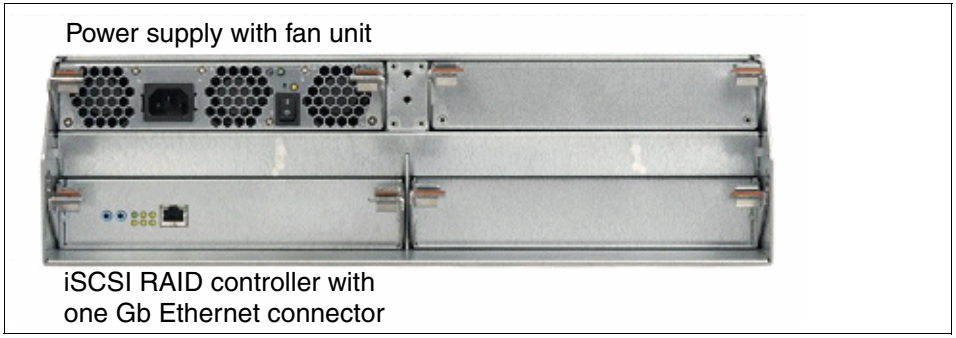

Figure 3-2 DS300 Model 1701-1RL - rear view

Attention: It is theoretically possible to upgrade the DS300 Model 1RL to a dual controller model. The already installed controller would have to be replaced (because FRUs do not come with the cache memory and battery backup cache), and a second controller would have to be installed along with a second power supply.

However, this upgrade option is not offered and not supported. In addition, this upgrade would cost more than purchasing a dual controller model!

3.3.2 DS300 Model 1701-1RS

Figure 3-3 DS300 Models 1701-1RS and 1701-2RD - front view

The DS300 storage subsystem 1701-1RS contains one 1GB iSCSI RAID controller with three 1 Gbps Ethernet connectors. One connector is used for management control, while the remaining two are used for data communication.

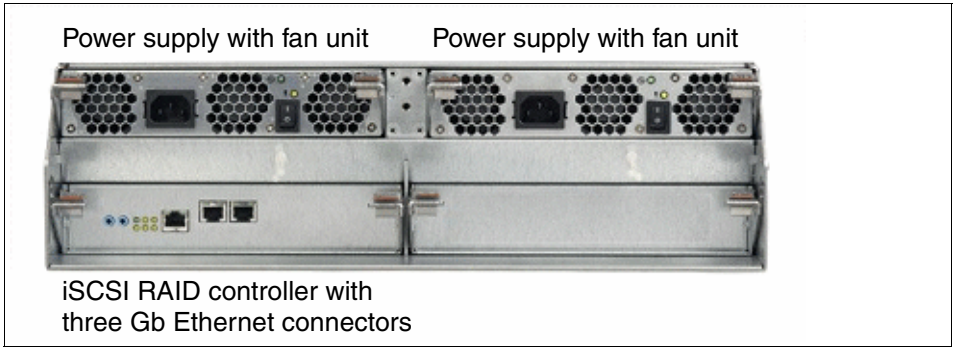

Figure 3-4 DS300 Model 1701-1RS - rear view

3.3.3 DS300 Model 1701-2RD

The DS300 storage subsystem 1701-1RD contains two 1GB iSCSI RAID controllers. Each controller contains three 1 Gbps Ethernet connectors. One connector is used for management control, while the remaining two are used for data communication.

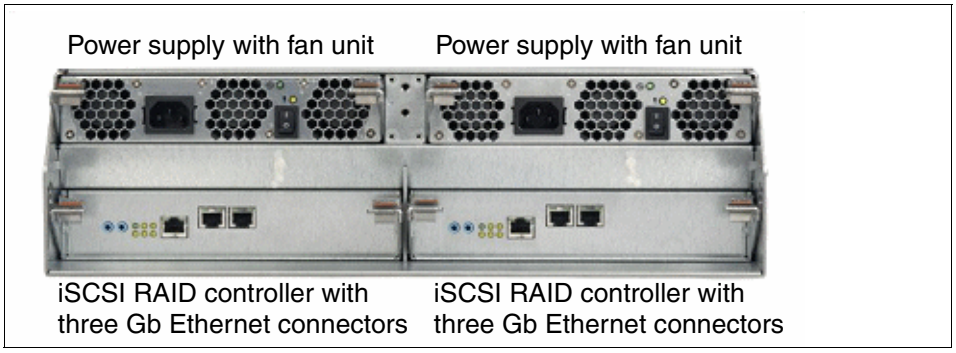

Figure 3-5 DS300 Model 1701-2RD - rear view

3.3.4 DS300 controls and LEDs

The following sections describe the various controls and LEDs that are located on the DS300 front and rear panels.

Front view

Figure 3-6 shows the LEDs on the front of the DS300 storage subsystem.

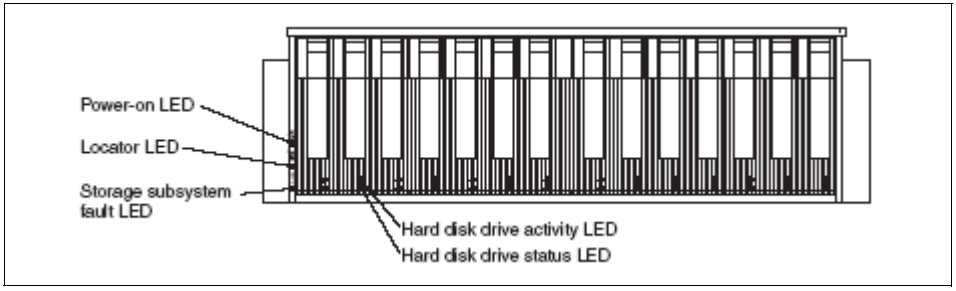

Figure 3-6 DS300 front LEDs

Table 3-3 DS300 front panel LEDs

DS300 problem indicators	Action
Green Power LED is off. When this LED is off, it indicates that DC power is not present, or that the power supply or the LED itself has failed. A power LED is also on the power supply.	Check Power Supply LEDs. Refer to power supply LEDs table. Note: To remove all electrical power from the storage subsystem, you must disconnect the power cord from the electrical outlet.
Blue Locator LED is lit. When this LED is lit, it has been remotely lit by the ServeRAID Manager program (running on the system that is the management station for the storage subsystem), to aid in visually locating the storage subsystem. In a dual controller system, when this LED is lit, it can possibly mean that one controller is down.	None required Note: The LED can be turned off by selecting Identify Enclosure from the ServeRAID Manager program. Telnet to the administration port of the surviving controller and enable the alternate controller with the following command: controller peer enable
Amber Storage Subsystem Fault LED is lit solidly. If the fault LED is lit continuously (not flashing), there is a problem with the storage subsystem. Use the ServeRAID Manager program to diagnose and repair the problem. One or more of the following conditions might exist: ► Power supply fault ► Fan fault ► Overtemperature fault ► RAID controller module fault ► Battery fault ► Cache DIMM fault ► XScale DIMM fault ► Disk drive fault ► Turned ON through SES control page	1. Check the LEDs for all components on enclosure (disk drives, power supplies, controllers) and follow the solution instructions provided for the respective illuminated LED. 2. If no other fault LED is on, replace the controller: 3. If replacing the controller does not resolve the issue, replace the chassis/backplane assembly.
Amber Storage Subsystem Fault LED is lit flashing. When this LED is lit flashing, it indicates that an illegal hardware configuration is detected.	Verify that the RAID controller and power supply are both installed in the slot A side. Note: This should only occur with the DS300 type 1701 model 1RL.
Green Hard disk drive activity LED is off. No drive activity is being detected.	Start your application and verify that the activity LED flashes as the drive is accessed.

DS300 problem indicators	Action
Amber Hard Disk Drive Status LED is lit solidly. Each hard disk drive has a status LED. When this LED is lit solidly, it indicates that the drive has failed. Note: When this LED is flashing slowly (one flash per second), it indicates that the drive is being rebuilt. When the LED is flashing rapidly (three flashes per second), it indicates that the RAID controller is identifying the drive.	For a single drive failure, replace the failed drive. For a multiple drive failure, attempt to determine the order in which the drives failed and then follow the appropriate recovery process for their RAID level.

Rear view

The following sections detail the LEDs, controls and connectors for the power supply and controller at the rear of the storage subsystem.

Power supply

Figure 3-7 shows the LEDs, controls and connectors on the power supply.

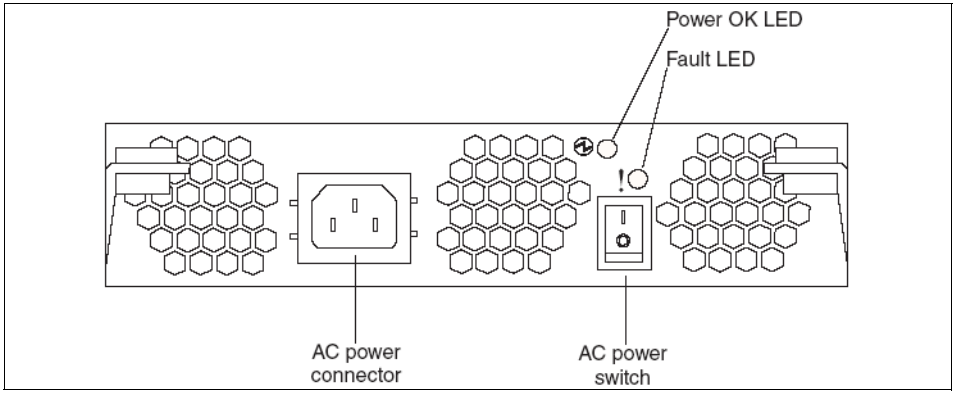

Figure 3-7 DS300 power supply LEDs, controls and connectors

- AC power connector - Connect the power cord to this connector.
- AC power on/off switch - Use this switch to turn the power supply on and off.
- Power OK LED (green) - When this LED is lit, it indicates that the power supply is turned on.
- Fault LED (amber) - When this LED is lit, there is a fault with either the power supply or a fan.

Table 3-4 Power supply LEDs

Power supply problem Indicator	Action
Green Power Supply Power LED is off.	Verify that the power cord is plugged in and that the storage subsystem power is on.
Amber Fault LED is lit solidly. The amber Fault LED will be ON solid when the enclosure services module detects a fault, such as one or more of the following conditions: ► DC output fault ► AC input fault ► Fan fault ► Turned ON through SES control page Note: The amber Fault LED flashes when the power/cooling module locate feature is selected.	Replace the power supply.

iSCSI RAID controller for DS300 model 1701-1RL

Figure 3-8 shows the LEDs, controls and connectors for the model 1701-1RL DS300 RAID controller with one Ethernet connector.

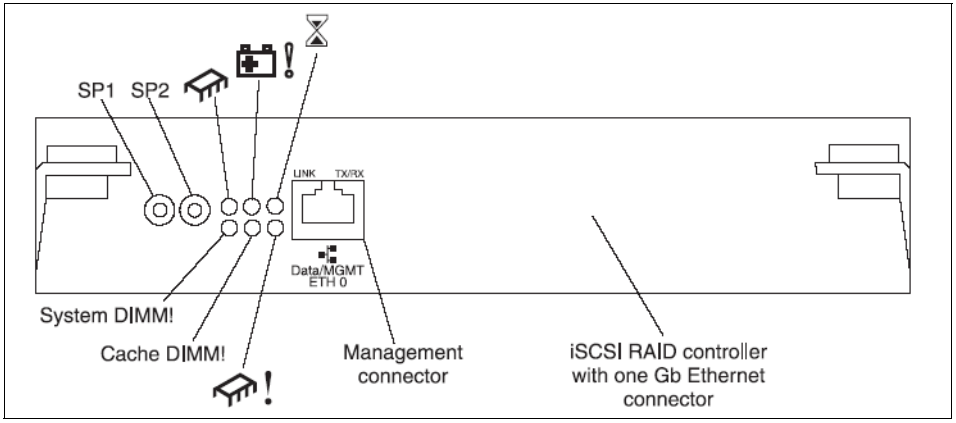

Figure 3-8 DS300 model 1701-1RL controller LEDs, controls and connections

- ► Serial connectors (SP1 and SP2) - The serial connectors are for diagnostic purposes only.
- ► Data and management iSCSI connector - The data and management connector is for ServeRAID Manager function and iSCSI data traffic.

The 1701-1RL controller has only one Ethernet port. This port is used for both management and data/traffic flow.

iSCSI RAID controller for DS300 models 1701-1RS and 1701-2RD

Figure 3-9 shows the LEDs, controls and connectors for the model 1701-1RS and 1701-2RD DS300 RAID controller with three Ethernet connectors.

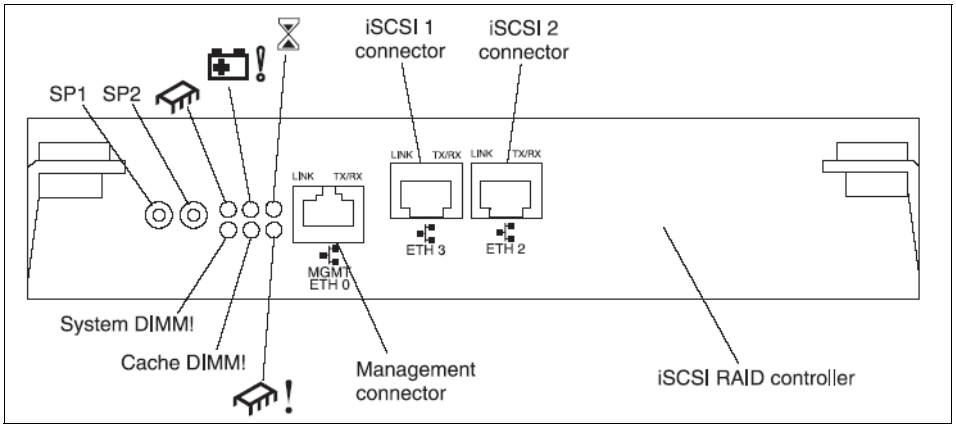

Figure 3-9 DS300 models 1701-1RS and 2RD controller LEDs, controls and connections

- Serial connectors (SP1 and SP2) - The serial connectors are for diagnostic purposes only.
- Management connector - The management connector is for ServeRAID Manager function.
- ETH 3 connector - The ETH 3 connector is for iSCSI data transfer.
- ETH 2 connector - The ETH 2 connector is for iSCSI data transfer.

The three network ports on the DS300 model 1701-1RS and 1701-2RD controllers are labeled ETH0, ETH2 and ETH3. ETH0 is used only to access the controller to make configuration changes and administer the unit. ETH2 and ETH3 are iSCSI interfaces used to pass data/traffic to the initiator. There is a fourth Ethernet port on the DS300 controller named ETH1. It is used for backchannel communication to keep the two controllers synchronized.

Table 3-5 iSCSI RAID controller LEDs

LED symbol	Controller problem indicator	Action
	Green Controller Ready LED is off.	If the Controller Ready LED is off, one or more of the RAID controller fault LEDs will be lit. Follow the recovery procedure for the Fault LED that is lit.

LED symbol	Controller problem indicator	Action
	Amber Battery Fault LED is lit. When this LED is lit, it indicates that the battery cannot sustain the RAID controller memory in case of power loss. This might be caused by any of the following conditions: ► The battery is removed. ► There is a battery-charger circuit failure. ► The battery temperature is too high. ► The battery voltage is out of range. ► The battery charge current is out of range.	1. Determine whether any other amber LEDs are lit and follow the appropriate recovery process for that component. 2. Remove the controller and replace the battery. Then reinsert the controller. Note: This action is not applicable to the type 1701 (model 1RL) because it does not have a battery
	Amber Cache Dirty LED is lit. When this LED is lit, it indicates that there is data in the cache that has not been written to a disk or synchronized to the redundant controller cache.	The Cache Dirty LED will go on and off at varying rates and durations during normal operation. No action is required when this LED goes on. If the Cache Dirty LED lit solidly, it indicates an improper controller shutdown during the controller initialization. Restart the controller from ServeRAID Manager to flush the cache. The Cache Dirty LED will turn off before the Controller Ready LED turns on and the server logins start. Attention: Do not turn off power to the enclosure or remove the controller from the enclosure when the Cache Dirty LED is illuminated. Doing so could cause loss of any data in the cache that is not mirrored. Notes: 1. This is not a FRU. 2. This action is not applicable to the DS300 type 1701 (model 1RL because it does not support a cache.

LED symbol	Controller problem indicator	Action
System DIMM!	Amber XScale (Processor) DIMM Fault LED is lit. When this LED is lit, it indicates that the XScale memory has failed. This is the result of one the following conditions: ► Cannot read the DIMM configuration data ► Nonrecoverable ECC failure	Replace the controller.
Cache DIMM!	Amber Cache DIMM Fault LED is lit. When this LED is lit, it indicates that the cache memory has failed. This is based on one of the following conditions: ► Cannot read the DIMM configuration data \ ► Nonrecoverable ECC failure	► Replace the DIMM 256 MB ► If the DIMM replacement does not fix the problem, replace the controller.
	Amber Controller Not Ready LED is lit. When this LED is lit, it indicates that the controller is not ready.	This LED might be lit because the controller did not finish POST. Contact your IBM support representative to attempt to recover the controller.
ETH0 Link	Green GbE Link Status LED is off. When this LED is lit, it indicates that the Gb Ethernet link is operational.	If this LED is off, perform the following actions: ► Check the status on the Ethernet switch. ► Replace the Ethernet cable. ► Switch the Ethernet cable to another port on the same switch in the same network segment.
ETH0 TX/RX	Green GbE Activity LED is off. This LED will flicker continuously between on and off. When this LED is lit, it indicates that data packets are being transmitted or received.	If this LED is off, then no data is being sent or received. Establish a connection and start passing data.

LED symbol	Controller problem indicator	Action
ETH3 Link	Green GbE Link Status LED is off. When this LED is lit, it indicates that the Gb Ethernet link is operational.	If this LED is off, perform the following actions: ► Check the status on the Ethernet switch. ► Replace the Ethernet cable. ► Switch the Ethernet cable to another port on the same switch in the same network segment.
ETH3 TX/RX	Green GbE Activity LED is off. This LED will flicker continuously between on and off. When this LED is lit, it indicates that data packets are being transmitted or received.	If this LED is off, then no data is being sent or received. Establish a connection and start passing data.
ETH2 Link	Green GbE Link Status LED is off. When this LED is lit, it indicates that the Gb Ethernet link is operational.	If this LED is off, perform the following actions: ► Check the status on the Ethernet switch. ► Replace the Ethernet cable. Switch the Ethernet cable to another port on the same switch in the same network segment
ETH2 TX/RX	Green GbE Activity LED is off. This LED will flicker continuously between on and off. When this LED is lit, it indicates that data packets are being transmitted or received.	If this LED is off, then no data is being sent or received. Establish a connection and start passing data

3.3.5 DS300 failover

Important: The description of failover on the DS300 subsystem was accurate at the time this book was written. However, failover processing and options may change in future releases.

The DS300 failover mechanism is IP failover. IP failover is not a true path failover mechanism. It only allows you to have controller redundancy. However, this failover method is independent of the iSCSI initiator used, so it can support both software based and hardware based initiators (refer to 3.7, “iSCSI software and hardware” on page 58).

Logical drives defined on a DS300 controller are visible only on the iSCSI ports of the controller that owns the array. The alternate controller in a dual controller DS300 does not see arrays that are defined on the other controller. Nor does the alternate controller respond with information such as passive paths as is done in the DS400.

Each DS300 controller has an Eth2 and Eth3 port for iSCSI traffic (see Figure 3-5 on page 34). Each of these ports must be assigned to a separate network segment. The Eth2 ports for both controller A and B must be assigned a different IP address in the same network segment, while the Eth3 ports for both controller A and B must be assigned a different IP address in another network segment.

The DS300 supports four modes for IP failover:

remote	the IP interfaces will failover between Eth2 → Eth2 and Eth3 → Eth3 between the controllers
none	the IP interfaces will not failover between controllers
local	the IP interfaces will failover from Eth2 → Eth3 or Eth3 → Eth2 within the same controller only
both	the IP interfaces will first failover between the two ports on the same controller and then to the remote interface for that port on the alternate controller. For example, Eth2[A] → Eth3[A] → Eth3[B] → Eth2[B]

For example, assume the DS300 is configured to support remote failover mode. The controller failover is initiated when a controller goes down or when a link to a controller goes down (a host bus adapter link failure does not cause a failover). The IP addresses of both Eth2 and Eth3 interfaces, and the arrays owned by the failing controller, are failed over to the alternate controller. If Eth2 port on controller A fails, sessions are failed over to Eth2 port on Controller B. The IP address on Eth2 port on controller A is migrated to Eth2 port on controller B. After failover occurs, the sessions are reconnected to the same IP address (but on Eth2 port on controller B, with a different MAC address) and I/O continues. Once the interface has failed over, and the IP address is online on Eth2 port on controller B, a Gratuitous ARP is sent to invalidate the current MAC address. A second Gratuitous ARP is sent to the switch to update the MAC address forwarding table for the new MAC bound to that IP address.

It is important to allow enough time for all the failover processing to occur. After the IP address has been moved to the new port (about 15 seconds), failover of the logical drives is sequential. It takes about 45-60 seconds for a failover of a single logical drive.

The DS300 failover mode can be set only with the CLI using the following command:

```
controller ipfailover <mode>
```

The default failover mode is **remote**.

Important: Use of the **both** failover mode is not recommended because it would require routing interfaces.

3.4 IBM TotalStorage DS400 disk subsystem

Unlike the DS300, the IBM TotalStorage DS400 storage subsystem offers 2 Gbps Fibre Channel (FC) host fabric ports per controller and can therefore be connected to Fibre Channel switches.

Up to 14 ultra 320 drives are supported in a DS400. A total of 40 ultra 320 drives can be archived by attaching up to two EXP400 enclosures. Each EXP400 enclosure can hold up to 13 drives.

The number of drives in an array depends on the RAID level chosen to be used. For RAID 0, 1 and 5 arrays up to 16 drives can be used to build a logical drive. With a RAID level 10 or 50 up to 40 drives are possible.

The DS400 supports asynchronous logical unit access (ALUA), which means that only one controller at a time can access the physical disk drives. This is often referred to as active/passive configuration.

The DS400 storage subsystem is available on two models as shown in Table 3-6.

Table 3-6 DS400 models

	1700-1RS	1700-2RD
Fibre channel RAID controller	1	2
Controller upgrade	Yes	N/A
Controller system memory	512MB, 1GB max	512MB, 1GB max
RAID cache (per controller)	256MB	256MB
Write cache enabled	Yes	No
Battery backup for cache	Yes	Yes
Flashcopy copy services	256 Flashcopy target logical drives with 1GB system memory	256 Flashcopy target logical drives with 1GB system memory

	1700-1RS	1700-2RD
Gbps Ethernet port	1	1
Fibre channel ports/controller	2	2
Power supply w/fan unit	2	2
Disk bays	14	14
LED indicators	8	12

The Gbps Ethernet connector on each controller is used for management.

3.4.1 Model 1700-1RS with one Fibre Channel RAID controller

The DS400 storage subsystem 1700-1RS contains one 2GB Fibre Channel RAID controller. The controller contains one 1 Gbps Ethernet connector for management control, and two 2 Gbps Fibre Channel connectors for data communication.

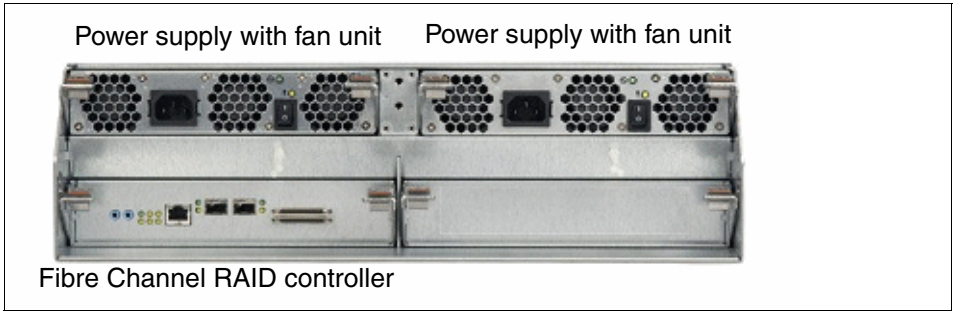

Figure 3-10 DS400 Model 1700-1RS - rear view

3.4.2 Model 1700-2RD with two Fibre Channel RAID controllers

The DS400 storage subsystem 1700-1RS contains two 2GB Fibre Channel RAID controllers. Each controller contains one 1 Gbps Ethernet connector for management control, and two 2 Gbps Fibre Channel connectors for data communication.

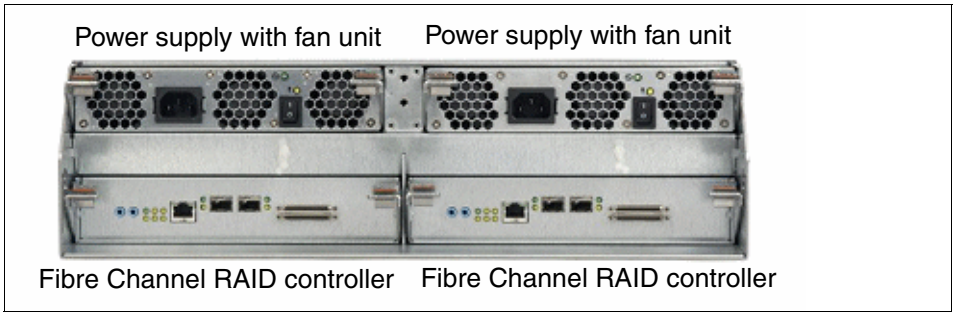

Figure 3-11 DS400 Model 1700-2RD - rear view

3.4.3 DS400 controls and LEDs

The following sections describe the various controls and LEDs that are located on the DS400 front and rear panels.

Front view

Figure 3-12 shows the LEDs on the front of the DS400 storage subsystem.

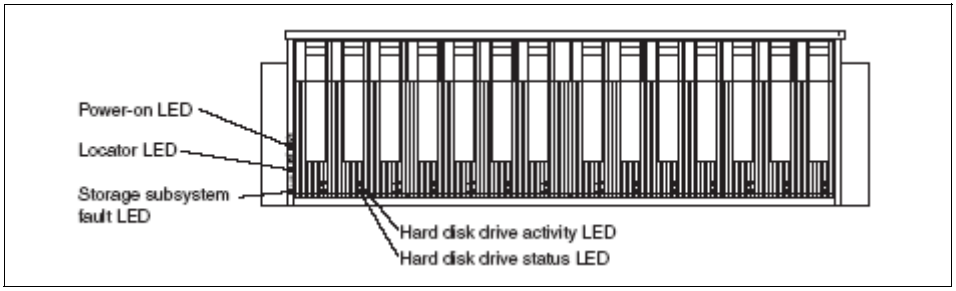

Figure 3-12 DS400 front LEDs

Table 3-7 DS400 front panel LEDs

DS400 problem indicators	Action
Green Power LED is off. When this LED is off, it indicates that DC power is not present, or that the power supply or the LED itself has failed. A power LED is also on the power supply.	Check Power Supply LEDs. Refer to power supply LEDs table. Note: To remove all electrical power from the storage subsystem, you must disconnect the power cord from the electrical outlet.

DS400 problem indicators	Action
Blue Locator LED is lit. When this LED is lit, it has been remotely lit by the ServeRAID Manager program (running on the system that is the management station for the storage subsystem), to aid in visually locating the storage subsystem. In a dual controller system, when this LED is lit, it can possibly mean that one controller is down.	None required Note: The LED can be turned off by selecting Identify Enclosure from the ServeRAID Manager program. Telnet to the administration port of the surviving controller and enable the alternate controller with the following command: controller peer enable
Amber Storage Subsystem Fault LED is lit solidly. If the fault LED is lit continuously (not flashing), there is a problem with the storage subsystem. Use the ServeRAID Manager program to diagnose and repair the problem. One or more of the following conditions might exist: ► Power supply fault ► Fan fault ► Overtemperature fault ► RAID controller module fault ► Battery fault ► Cache DIMM fault ► XScale DIMM fault ► Disk drive fault ► Turned ON through SES control page	1. Check the LEDs for all components on enclosure (disk drives, power supplies, controllers) and follow the solution instructions provided for the respective illuminated LED. 2. If no other fault LED is on, replace the controller: 3. If replacing the controller does not resolve the issue, replace the chassis/backplane assembly.
Green Hard disk drive activity LED is off. No drive activity is being detected.	Start your application and verify that the activity LED flashes as the drive is accessed.
Amber Hard Disk Drive Status LED is lit solidly. Each hard disk drive has a status LED. When this LED is lit solidly, it indicates that the drive has failed. Note: When this LED is flashing slowly (one flash per second), it indicates that the drive is being rebuilt. When the LED is flashing rapidly (three flashes per second), it indicates that the RAID controller is identifying the drive.	For a single drive failure, replace the failed drive. For a multiple drive failure, attempt to determine the order in which the drives failed and then follow the appropriate recovery process for their RAID level.

Rear view

The following sections detail the LEDs, controls and connectors for the power supply and controller at the rear of the storage subsystem.

Power supply

Figure 3-13 shows the LEDs, controls and connectors on the power supply.

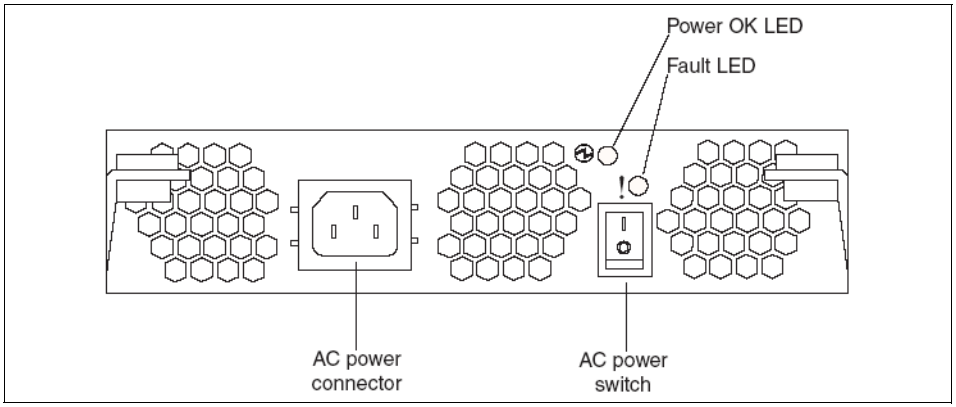

Figure 3-13 DS400 power supply LEDs, controls and connectors

- AC power connector - Connect the power cord to this connector.
- AC power on/off switch - Use this switch to turn the power supply on and off.
- Power OK LED (green) - When this LED is lit, it indicates that the power supply is turned on.
- Fault LED (amber) - When this LED is lit, there is a fault with either the power supply or a fan.

Table 3-8 Power supply LEDs

Power supply problem Indicator	Action
Green Power Supply Power LED is off.	Verify that the power cord is plugged in and that the storage subsystem power is on.

Power supply problem Indicator	Action
Amber Fault LED is lit solidly. The amber Fault LED will be ON solid when the enclosure services module detects a fault, such as one or more of the following conditions: ► DC output fault ► AC input fault ► Fan fault ► Turned ON through SES control page Note: The amber Fault LED flashes when the power/cooling module locate feature is selected.	Replace the power supply.

Fibre Channel RAID controller for DS400

Figure 3-14 shows the LEDs, controls and connectors for the DS400 Fibre Channel RAID controller.

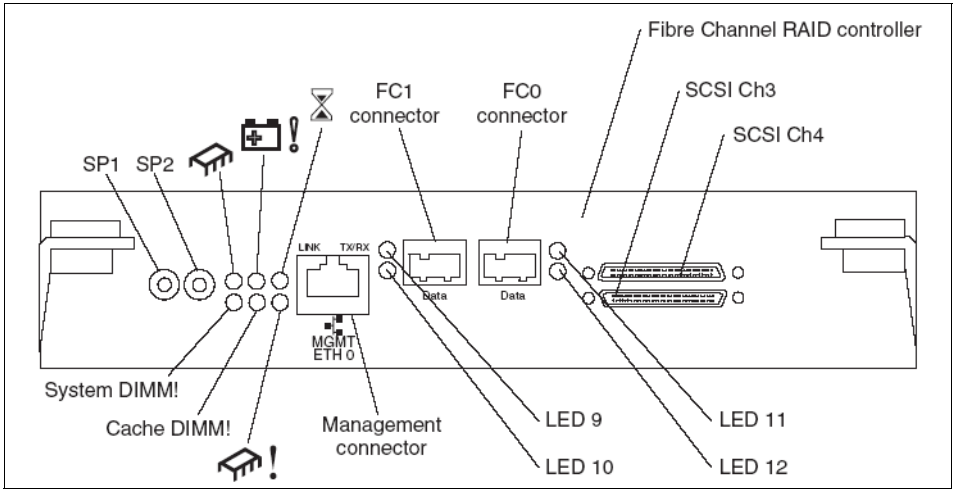

Figure 3-14 DS400 controller LEDs, controls and connections

- Serial connectors (SP1 and SP2) - The serial connectors are for diagnostic purposes only.
- Management connector - The management connector is for ServeRAID Manager function.
- FC 1 connector - The FC 1 connector is for Fibre Channel data transfer.
- FC 0 connector - The FC 0 connector is for Fibre Channel data transfer.
- SCSI channel 4 - SCSI channel 4 can be connected to an IBM EXP400 Storage Expansion Unit.

- SCSI channel 3 - SCSI channel 3 can be connected to an IBM EXP400 Storage Expansion Unit.

The Fibre Channel RAID controller has two status LEDs on each of the two Fibre Channel host connectors.

Table 3-9 Fibre Channel RAID controller LEDs

LED symbol	Controller problem indicator	Action
	Green Controller Ready LED is off.	If the Controller Ready LED is off, one or more of the RAID controller fault LEDs will be lit. Follow the recovery procedure for the Fault LED that is lit.
	Amber Battery Fault LED is lit. When this LED is lit, it indicates that the battery cannot sustain the RAID controller memory in case of power loss. This might be caused by any of the following conditions: ► The battery is removed. ► There is a battery-charger circuit failure. ► The battery temperature is too high. ► The battery voltage is out of range. ► The battery charge current is out of range.	1. Determine whether any other amber LEDs are lit and follow the appropriate recovery process for that component. 2. Remove the controller and replace the battery. Then reinsert the controller. Note: This action is not applicable to the type 1701 (model 1RL) because it does not have a battery
	Amber Cache Dirty LED is lit. When this LED is lit, it indicates that there is data in the cache that has not been written to a disk or synchronized to the redundant controller cache.	The Cache Dirty LED will go on and off at varying rates and durations during normal operation. No action is required when this LED goes on. If the Cache Dirty LED lit solidly, it indicates an improper controller shutdown during the controller initialization. Restart the controller from ServeRAID Manager to flush the cache. The Cache Dirty LED will turn off before the Controller Ready LED turns on and the server logins start. Attention: Do not turn off power to the enclosure or remove the controller from the enclosure when the Cache Dirty LED is illuminated. Doing so could cause loss of any data in the cache that is not mirrored.

LED symbol	Controller problem indicator	Action
System DIMM!	Amber XScale (Processor) DIMM Fault LED is lit. When this LED is lit, it indicates that the XScale memory has failed. This is the result of one the following conditions: ► Cannot read the DIMM configuration data ► Nonrecoverable ECC failure	Replace the controller.
Cache DIMM!	Amber Cache DIMM Fault LED is lit. When this LED is lit, it indicates that the cache memory has failed. This is based on one of the following conditions: ► Cannot read the DIMM configuration data \ ► Nonrecoverable ECC failure	► Replace the DIMM 256 MB ► If the DIMM replacement does not fix the problem, replace the controller.
	Amber Controller Not Ready LED is lit. When this LED is lit, it indicates that the controller is not ready.	This LED might be lit because the controller did not finish POST. Contact your IBM support representative to attempt to recover the controller.
ETH0 Link	Green GbE Link Status LED is off. When this LED is lit, it indicates that the Gb Ethernet link is operational.	If this LED is off, perform the following actions: ► Check the status on the Ethernet switch. ► Replace the Ethernet cable. ► Switch the Ethernet cable to another port on the same switch in the same network segment.
ETH0 TX/RX	Green GbE Activity LED is off. This LED will flicker continuously between on and off. When this LED is lit, it indicates that data packets are being transmitted or received.	If this LED is off, then no data is being sent or received. Establish a connection and start passing data.

Table 3-10 Fibre Channel port status LEDs

Green LEDs		Amber LEDs		Indication
FC1 (LED 9)	**FC 0 (LED 11)**	**FC 1 (LED 10)**	**FC 0 (LED12)**	
On		On		Power on

Green LEDs	Amber LEDs	Indication
On	Off	Online (normal operation)
Off	On	Link established (transitory condition)
Off	Flashing twice per second	Loss of synchronization
Green and amber LEDs alternate flashing twice per second		FC chip fault
Off	Flashing once per second	Beacon

Each of the two Fibre Channel ports (FC0 and FC1) has a unique world wide port name (WWPN). The management port (ETH0) is used only to access the controller to make configuration changes and administer the unit.

3.4.4 DS400 failover

All logical drives are seen on all 4 FC ports on a dual controller on DS400. However, logical drives can only be active on two of those ports at a time. Logical drives are configured with a “preferred path” and an “alternate path”. A DS400 controller will initiate a failover from the preferred to the alternate path any time the Fibre Channel link goes down on a Fibre Channel port, or the alternate controller goes offline for some reason. When this failover occurs, the logical drives will go into a transition state as the array ownership changes. The Device Specific Module (DSM) will monitor the logical unit numbers (LUNs) until they change into either a passive state or active state. Once this occurs, the DSM will issue a command to select the currently active path.

DSM

The Device Specific Module is a major part of the DS400 failover mechanism. A DSM is the “personality” module that is loaded by the Microsoft MPIO driver to handle customer failover calls. It is provided by a third party vendor (Qlogic for the DS400) to implement the actual device moves. A DSM is provided to differentiate each storage device. Multiple DSMs may be loaded for different storage device types, but only one instance of MPIO is required. Each DSM should claim only its own storage device. The DSM provided with the DS400 only works with the DS400.

A DSM will function only if the Microsoft MPIO driver is installed as well. A DSM can be thought of as a low level driver talking directly to the hardware. The MPIO driver is the failover mechanism that builds and manages the path definitions.

3.5 EXP400 expansion unit

Up to two EXP400 expansion units can be attached to an IBM DS400 subsystem for a total of 40 hard disk drives.

The EXP400 Storage Expansion Unit supports up to 14 LVD (low voltage differential) Ultra320 SCSI disk drives which transfer data at 320MB/s. The EXP400 supports SCSI drives on a single or dual bus. It supports a cable distance of up to 20 meters.

The EXP400 delivers fast, high-volume data transfer, retrieval, and storage functions across multiple drives. The modular, redundant disk drives, power supply with fan units, and SCSI Bus Expander and Enclosure Services Module (ESM) use hot-swap technology for easy replacement without turning off the expansion unit.

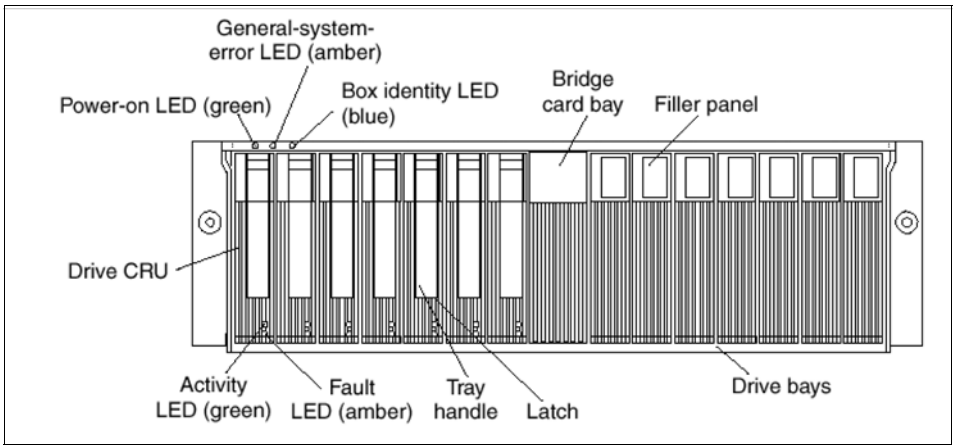

Figure 3-15 EXP400 front view

The EXP400 accepts only slim-high Ultra320 SCSI drives. These drives must be in the same Ultra320 tray that is supported in the IBM Ultra320 servers. Ultra320 SCSI drives up to 300GB's in size can be installed in the EXP400 which give it the ability to support up to 4TB's of storage capacity. The unit supports either 10k RPM or 15k RPM SCSI disk drives. To activate the Blue Box Identity LED, ServeRAID Manager must be used. The amber general fault LED is activated by any fault that occurs in the enclosure.

Figure 3-16 on page 53 displays a rear view of the EXP400 and the position of the two Enclosure Services Module (ESM) boards. An ESM provides the SCSI interface to the hard disk drives and monitors the overall status of the expansion unit. Each EXP400 supports up to two ESMs.

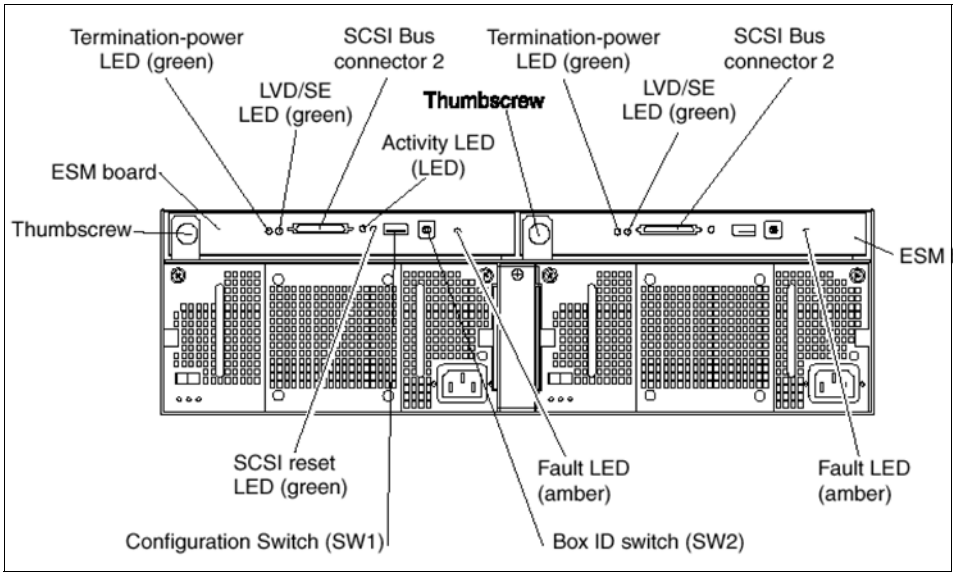

Figure 3-16 EXP400 back view

The EXP400 Expansion unit is available in 3 models:

1. 1733-1RU - comes with two 500-watt North America AC power supply with fan units, one Enclosure Services Module (the second ESM is optional), a filler panel to cover the empty ESM bay, and 14 drive filler panels.
2. 1733-1RX - same as 1733-1RU but with EMEA AC power.
3. 1733-2RX - comes with two -48 volt DC power supply with built-in fan units, two ESMs, and 14 drive filler panels.

The IBM DS400 storage controller has two external SCSI connectors (see Figure 3-14 on page 48) which can connect two optional EXP400 expansion units allowing up to 7.8 TB (assuming 300 GB hard drives) of additional data.

Note: SCSI ID 6 (slot 7) must be left free in each EXP400 enclosure for SCSI controller use. This means that only 13 drives can be installed in each enclosure.

When an EXP400 is attached to a DS400 with dual controllers, both EXP400 ESM boards must be connected to the DS400. It is important that the EXP400 be connected to SCSI Expansion Channel 3 on both DS400 controllers, or to SCSI Expansion Channel 4 on both DS400 controllers.

The IBM DS400 can be used in a split configuration with a single controller DS400 model 1700-1RS. See the EXP400 hardware maintenance manual for jumper settings for split configuration of the EXP400.

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-53436

Refer to the following documents for more information about the EXP400 Storage Expansion Units:

- *Hardware Maintenance Manual and Troubleshooting Guide*
- *Installation Guide*
- *User's Guide*

All of these documents are available at the following URL:

http://www.ibm.com/pc/support/site.wss/quickPath.do?quickPathEntry=1733

See 4.10, "Install premium feature license key" on page 167 for installing the feature key.

3.6 Premium feature license keys

Certain IBM DS300 and DS400 premium features require an additional *license key* to activate the feature. A premium feature allows use of additional functions in the IBM TotalStorage DS300 and DS400. These additional features are:

- more than 4 Flashcopys

  With this feature, up to 256 Flashcopys can be supported for a single controller. This feature requires purchase of a 1 Gbyte system memory upgrade for the controller.

- support for the EXP400 expansion unit with an IBM DS400 subsystem

  The EXP400 expansion unit is supported on single and dual controller IBM DS400 subsystems. To be able to see the drives in the EXP400 enclosure, or the enclosure itself, a license key is required.

When a premium feature is purchased, a booklet with a *purchase validation code* is included. This booklet also contains documentation on the procedure to follow to retrieve the *license key*. Figure 3-17 on page 55 shows the booklet included with the Flashcopy premium feature. The booklet included for an EXP400 license is the same except for the EXP400 label.

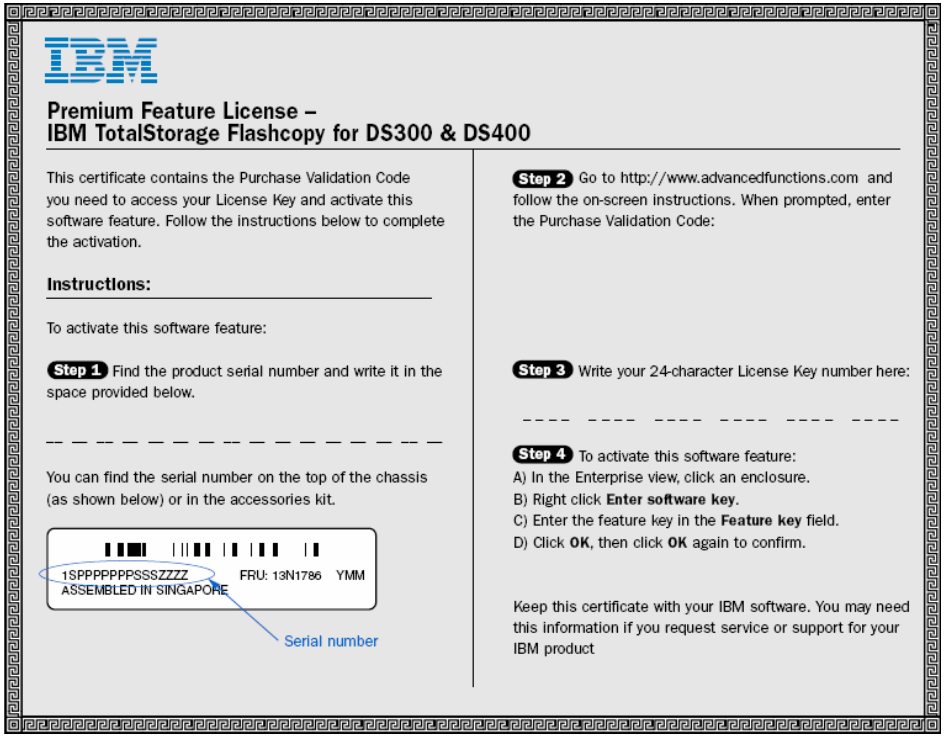

Figure 3-17 Booklet to retrieve a Flashcopy premium license key

To retrieve the license key, some specific information about the DS300 or DS400 onto which the feature is to be installed is required:

- Purchase Validation Code (included with the booklet supplied with the feature)
- Product serial number of the DS300 or DS400. This data can be found on the rear side of the DS300 or DS400.
- 24-character controller license key. See Figure 4-147 on page 167.

Connect to the following Web site to retrieve the feature license key:

`http://www.advancedfunctions.com`

The page shown in Figure 3-18 on page 56 is displayed.

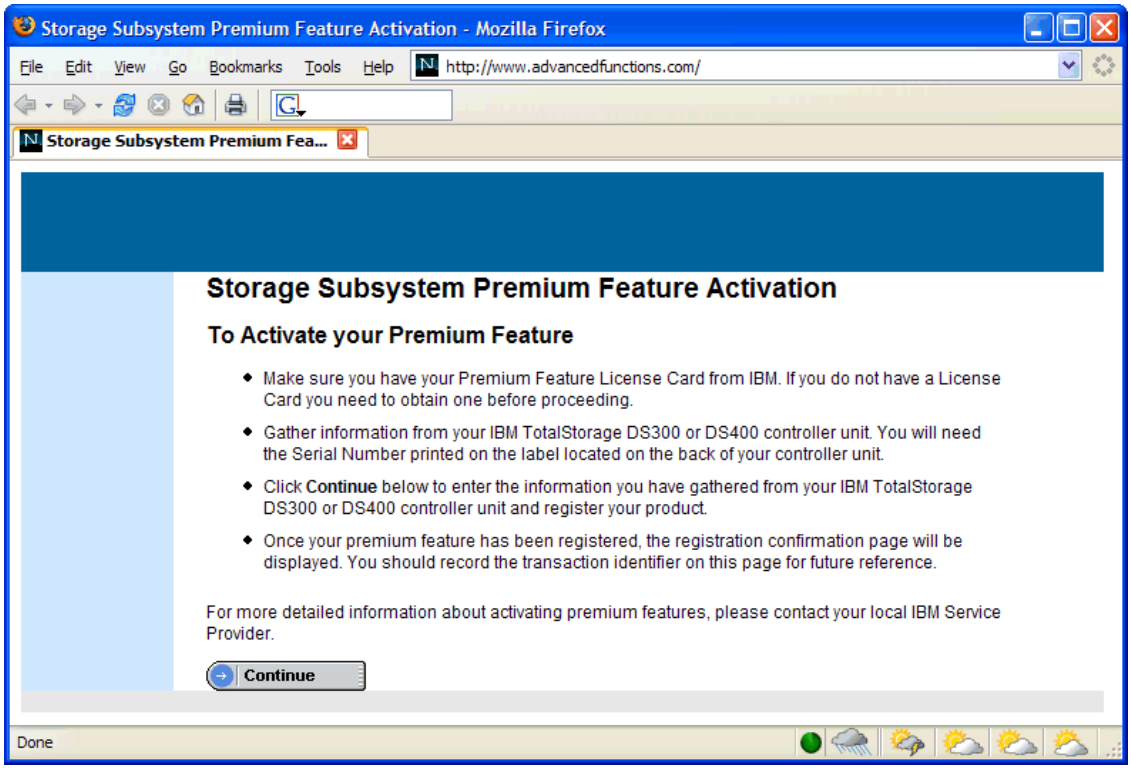

Figure 3-18 Premium feature activation - part 1

The Premium Feature License Card described in Figure 3-18 is referred to as the Purchase Validation Code in this example and on Figure 3-17 on page 55. Press **Continue** to proceed with the license key retrieval process. The page shown in Figure 3-19 on page 57 is displayed.

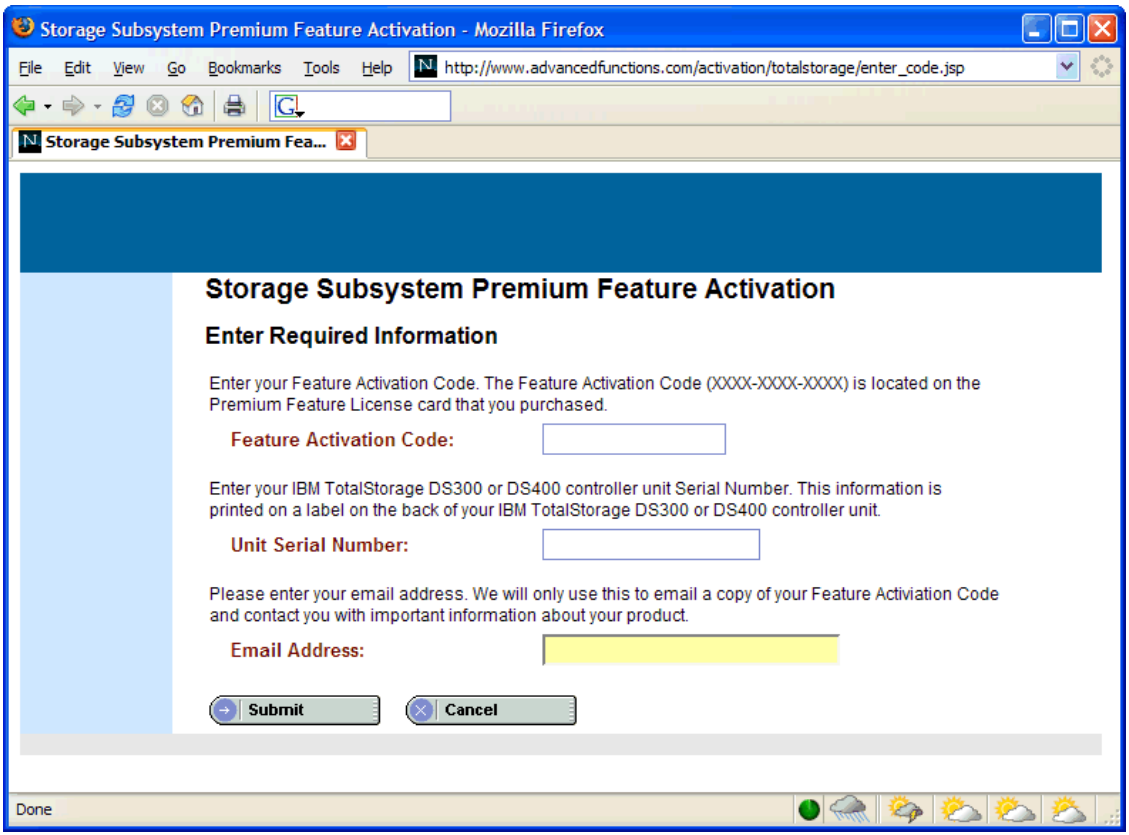

Figure 3-19 Premium feature activation - part 2

Enter the Purchase Validation Code/Feature Activation Code into the **Feature Activation Code** field. Enter the DS300/DS400 serial number into the **Unit Serial Number** field. Enter an email address to send the license key to. Press **Submit**.

Note: Be very careful to select the data from the correct DS300 or DS400. License keys cannot be transferred from one DS300 or DS400 to another!

See 4.10, “Install premium feature license key” on page 167 for information about how to enter the premium feature license key into the DS300/DS400 controller.

3.7 iSCSI software and hardware

The following sections describe software and hardware that support and implements iSCSI transport between xSeries and BladeCenter servers and the TotalStorage DS300. See 2.6, “iSCSI initiators” on page 24 for more information about iSCSI initiators.

3.7.1 Software initiators

Software initiators are available for Windows 2000, Windows 2003, and Linux using both the 2.4 and 2.6 kernel levels.

The Microsoft iSCSI Software initiator is available in two versions: V1.06 and V2.0. V2.0 has been tested only with DS300/DS400 firmware level V6.21 and higher.

Cisco Systems Inc. Linux software initiator 3.4.3 and 3.4.4 are also supported.

Note: At the time this book was written, iSCSI software initiators for both Windows and Linux were available from Cisco Systems, Inc. Consult with Cisco Systems to verify availability in you location.

3.7.2 Hardware initiators

Hardware initiators are available, but less prevalent than software initiators.

IBM iSCSI Server Adapter 73P3601

The IBM iSCSI Server Adapter enables block-oriented storage data to be carried by Ethernet IP networks. It is a hardware-based iSCSI Initiator for IBM @server xSeries servers to connect to iSCSI target devices, such as the IBM TotalStorage DS300 external storage enclosure. It provides remote boot capability. This option reduces CPU workloads by moving all protocol I/O processing off the host.

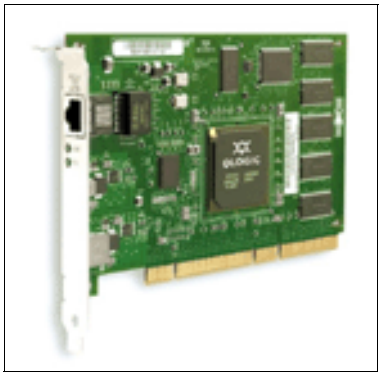

Figure 3-20 IBM iSCSI Server Adapter

- Works only at 1Gbps, full duplex. 10/100 speeds are not supported.
- Single-port PCI-X 64-bit/133MHz copper adapter over CAT 5E or CAT 6 cable
- Supports both 3.3v and 5v slots
- Half-high card with a standard size bracket
- Based on QLogic QLA4010C HBA with ISP4010 iSCSI/TOE chipset
- Full hardware TCP/IP Offload Engine (TOE) and iSCSI Initiator
- Supports iSCSI boot
- For all IBM @server xSeries IA32 systems.
- All driver software is provided by QLogic.
- Supports Windows 2000, Windows 2003, Linux Red Hat, SUSE Linux

The Windows 2000 and 2003 drivers and related software are included with the card. The drivers are digitally signed.

For the drivers and related software for Linux RHEL and SUSE SLES, go to the following Web site.

`http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-57073`

QLogic iSCSI TOE Expansion Card for IBM eServer BladeCenter 26K6487

The QLogic iSCSI Expansion Card for IBM eServer™ BladeCenter option is a hardware initiator that provides iSCSI communication from the blade server (HS20, HS40, and LS20) to an iSCSI storage device (for example, the DS300. This card offloads the iSCSI processing load from the CPU.

- Full iSCSI initiator function from HS20, HS40, and LS20 to iSCSI storage devices
- Enables blade servers to run diskless in a non-Fibre Channel SAN environment

- Full TOE for storage traffic only function (TCP/IP Offload Engine) to reduce CPU processing
- Dual port card utilizing QLogic QMC4052 enabling iSCSI connectivity through module bays 3 and 4
- Standard BladeCenter expansion card form factor (full-sized expansion card)

One of the following Ethernet Switch modules must be in chassis bays 3 and 4:

- 48P7054[1] IBM Gb Ethernet Switch (DLink)
- 13N0658 IBM Gb Ethernet Switch (DLink)
- 90P3776 IBM GbE Switch Module (Intel)
- 13N2281 Cisco Copper
- 26K6547 Cisco Fiber
- 26K6530 Nortel L2/L3 Copper
- 26K6531 Nortel L2/L3 Fiber
- 73P9057 Nortel L2-7
- 02R9080 Optical Passthrough Module (OPM)
- 73P6100 Copper Passthrough Module (CPM)

For the drivers and related software to support this card, go to the following Web site.

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-60280

3.8 Fibre Channel hardware

Fibre Channel technology is outlined in the *SCSI-3 Fibre Channel Protocol* (SCSI-FCP) standard. Fibre Channel is a high-speed data transport technology used for mass storage and networking.

Using a Fibre Channel Arbitrated loop (FC-AL), 126 Fibre Channel devices can be supported, compared to 15 small computer system interface (SCSI) devices with Ultra™ SCSI.

The following sections describe IBM hardware that supports and implements Fibre Channel transport between xSeries and BladeCenter servers and the TotalStorage DS400.

3.8.1 IBM TotalStorage DS4000 FC2-133 Host Bus Adapter 24P0960

The IBM TotalStorage DS4000 FC2-133 Host Bus Adapter is a 2 GB high-performance, direct memory access (DMA), bus master, Fibre Channel host

[1] Option part number

adapter designed for high-end systems. The function and performance are derived from the ISP2312 chip, making this Host Bus Adapter a leading-edge host adapter.

The DS4000 Adapter uses a multi-mode shortwave optical interface for distances up to 500 meters when operating at 1 Gbps and 300 meters when operating at 2 Gbps. It supports data transfer rates up to 200 MB per second half-duplex and 400 MB per second full-duplex on optical interfaces.

Attention: Multi-mode optical cabling is available in either a 50 micron or 62.5 micron core with a cladding of 125 microns. The FC2-133 Adapter requires use of 50 micron optical cabling.

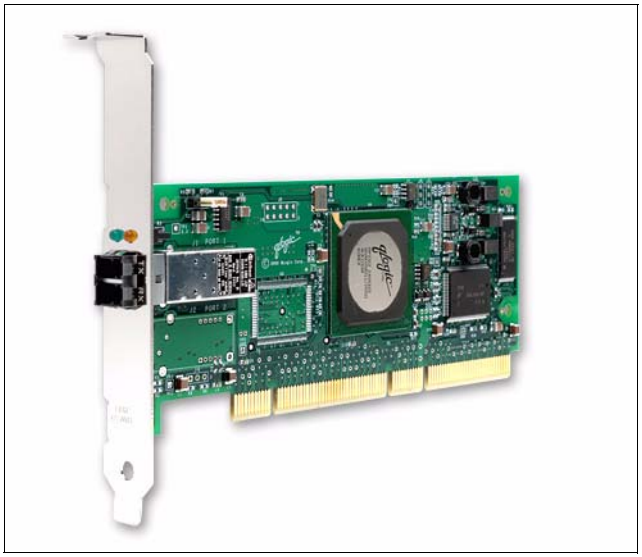

Figure 3-21 DS4000 FC2-133 Host Bus Adapter

The ISP2312 chip combines a reduced instruction set computer (RISC) processor, a Fibre Channel protocol manager (FPM) with one 2 GB Fibre Channel transceiver, and a peripheral component interconnect (PCI) or peripheral component interconnect-extended (PCI-X) local bus interface in a single-chip solution. The IBM TotalStorage DS4000 133 Host Bus Adapter supports all Fibre Channel (FC) peripheral devices that support private-loop direct attach (PLDA) and fabric-loop attach (FLA).

The DS4000 FC2-133 Adapter has the following features:

- Based on the QLogic QLA2340 Fibre Channel HBA
- Compliance with Intel PCI Local Bus version 2.2 specification

- Compliance with peripheral component interconnect-extended (PCI-X) addendum, revision 1.0 to the Intel PCI Local Bus version 2.2 specification
- Compliance with Third Generation Fibre Channel Physical and Signaling Interface (PC-PH-3), revision 9.2
- Compliance with Fibre Channel Arbitrated Loop (FC-AL-2) standard
- Compliance with U.S. and international safety and emissions standards
- Support for direct memory access (DMA)
- Support for bus master
- Fast!UTIL basic input/output system (BIOS) utility program to customize the configuration parameters on the FC2-133 Adapter and attached drives
- Support for Fibre Channel protocol SCSI (FCP-SCSI) and Fibre Channel Internet protocol (FCP-IP)
- Support for point-to-point fabric connection (F-PORT FABRIC LOGIN)
- Support for Fibre Channel service (Classes 2 and 3)

3.8.2 IBM TotalStorage DS4000 FC2-133 Dual Port Host Bus Adapter

The IBM TotalStorage DS4000 FC2-133 Dual Port Host Bus Adapter shares the same specifications as the The IBM TotalStorage DS4000 FC2-133 Host Bus Adapter except that it supports two Fibre Channel port attachments instead of one. It is based on the Based on the Qlogic QLA2342 Fibre Channel HBA.

3.8.3 IBM TotalStorage SMB Host Bus Adapter (QLA200) 13N1873

This host bus adapter is based on the Qlogic QLA200. It includes a single fibre channel 2 Gbps port supporting 2Gbps operations. This HBA is similar to the IBM TotalStorage DS4000 FC2-133 host bus adapter (QLA2340) with a reduced feature set. ASIC is the same. The driver for the IBM TotalStorage DS4000 host bus adapter and the IBM TotalStorage SMB host bus adapter are the same. The same failover driver can be used also for both adapters. Most of the functionality is reduced by firmware but also hardware is different. This adapter uses less memory then the QLA2340.

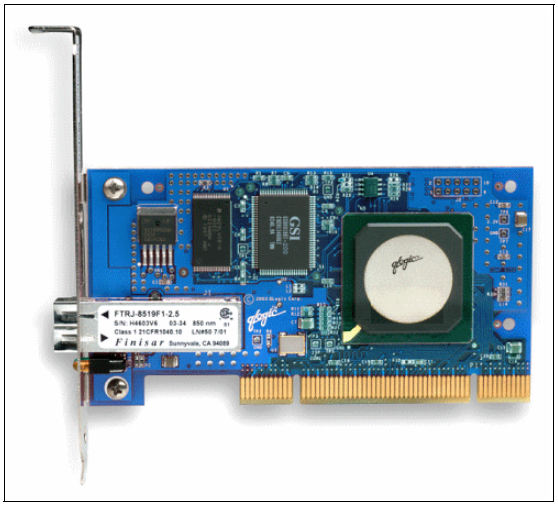

Figure 3-22 TotalStorage SMB Host Bus Adapter

This adapter has a limited functionality compared to a TotalStorage DS4000 host bus adapter. This limited functionality makes the adapter more price effective. The following table provides an introduction to the features of the IBM TotalStorage SMB HBA.

Table 3-11 Qlogic HBA comparison

Feature	IBM TotalStorage DS4000 HBA	IBM TotalStorage SMB HBA (QLA200)
Number of boot devices	4	2
Number of LUNs per target	256	16
64bit PCI support	Yes	No
FC speed auto negotiation	Yes	2Gb only
Load balancing	Yes	No
LUN masking	Yes	No

The QLA200 adapter allows up to two WWPN/LUN to be selected as boot devices.

3.8.4 IBM TotalStorage SAN Switch 2005-H08 and 2005-H16

The IBM H08 and H16 SAN Switches are well-suited to address small and medium business customer fibre channel requirements for infrastructure

simplification and improved business continuity. The H08 comes standard with eight fibre channel ports, while the H16 comes standard with sixteen fibre channel ports.

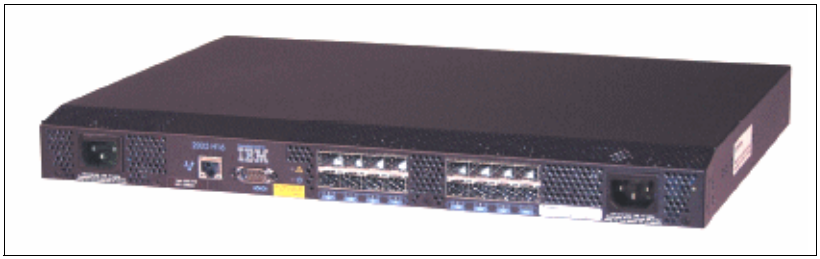

Figure 3-23 SAN Switch 2005-H16

They are designed to be fully interoperable with other products in the IBM TotalStorage SAN Switch family. You can configure a wide range of scalable solutions that help address SMB demands for affordability, simple to manage and integrated heterogeneous open server SANs.

The switches come with four shortwave fibre optic transceivers, and a single power supply. The support one connection to one other IBM SAN Switch in a two-switch fabric, as well as support for optional full-fabric activation.

Both switches support:

- Entry-fabric switch with advanced zoning and ability to attach to two IBM TotalStorage SAN Switch fabric for Microsoft Windows and UNIX server clustering
- Upgradable to full-enterprise capabilities for use in large core-to-edge fabrics
- Offers 8-port or 16-port Fibre Channel switching with wide range of Small Form-Factor Pluggable (SFP) transceiver features
- Improved availability capabilities include hot code activation and single field replaceable unit, FRU design
- Provides up to 2 Gigabit/sec. (Gbps) non-blocking performance. With Inter-Switch Link (ISL) Trunking, four links can be combined with an aggregate speed up to 8 Gbps
- Offers Advanced Security feature with extensive, policy based security capabilities

3.9 Advanced features

The following features are available for both DS300 and DS400 systems:

- Access Control Lists – ability to select which servers can access each storage partition or target (LUN mapping/masking). Access Control Lists allow the user to assign LUNs to servers and designate read/write or modify rights in order to improve security of data.
- Online logical drive expansion
- Online RAID level migration

> **Note:** Single drives and RAID 50 arrays cannot be migrated. In addition, migration to a RAID 50 array from any other RAID level is not supported.

See 4.8.6, “Increase logical drive capacity” on page 152 for information about how to increase logical drive capacity, and 4.8.7, “Expand or migrate arrays” on page 155 for information about how to migrate arrays.

3.9.1 Boot from SAN

IBM TotalStorage DS400 and DS300 allow you to boot from the SAN rather than from local disks on individual servers, that can enable you to maximize consolidation of IT resources and minimize equipment cost. Boot from SAN is a remote boot technology where the source of the boot disk is on the SAN. The server communicates with the SAN through host bus adapters (HBA). The HBA BIOS contains the instructions that enable the server to find the boot disk on the SAN. Booting from the SAN provides you a rapid Disaster Recovery as all the boot information and production data is stored on a local or remote SAN environment.

The DS300 operates on iSCSi and it supports booting from SAN using iSCSI interconnect to the SAN, provided iSCSI HBAs are used to enable the boot process.

> **Important:** Boot from SAN is not supported using an iSCSI software initiator. No sharing of boot images is allowed, as Windows servers cannot currently share a boot image. Each server requires its own dedicated LUN to boot.
>
> There is no path failover support for either the DS300 or the DS400 when booting from SAN.

See 9.1, “Boot Microsoft Windows Server 2003 from DS400 with IBM SMB Host Bus Adapter” on page 448 for a sample scenario.

Refer to the Microsoft Web site for more information and implementation on booting from SAN:

http://www.microsoft.com/windowsserversystem/wss2003/techinfo/plandeploy/BootfromSANinWindows.mspx

3.10 More information

For the latest information about IBM's TotalStorage offerings, refer to:

http://www-1.ibm.com/servers/storage/index.html

- The following link gives an A-Z listing of all IBM TotalStorage products:

  http://www-1.ibm.com/servers/storage/product/index.html

4

TotalStorage DS300 and DS400 management

This chapter reviews the two techniques that can be used to manage the TotalStorage DS300 and DS400 controllers:

- Command Line Interface
- IBM ServeRAID Manager

Examples of using these tools to create and modify SAN configurations, along with performing firmware upgrades, are covered.

This chapter covers the following topics:

- 4.8, “IBM TotalStorage DS300/DS400 array and logical drive management” on page 113
- 4.9, “Administer Management Station user accounts” on page 162
- 4.10, “Install premium feature license key” on page 167
- 4.11, “Flashcopy or SnapShot backup” on page 170
- 4.12, “Update TotalStorage DS300 and DS400 firmware” on page 179
- 4.13, “IBM ServeRAID Manager removal on Microsoft Windows” on page 190
- 4.14, “IBM ServeRAID Manager removal on Linux” on page 191
- 4.15, “Basic information gathering in case of trouble” on page 191

4.1 Management overview

You can use the IBM ServeRAID Manager or the command line interface (CLI) to manage a TotalStorage DS300 or DS400.

The command line interface has the greatest flexibility and can perform all management tasks. The command line interface is also a faster way to manage the TotalStorage DS300 or DS400 although it requires some knowledge of the command set. A description of the command line interface commands can be found in the problem determination guides for the DS300 or DS400. You may download them from

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-56887

In addition, the command line interface reference is included in Appendix A, “Command line interface reference” on page 565.

To use the command line interface, open a Telnet session to the management port of controller A or B of the TotalStorage DS300 or DS400. It is not required to connect to both controllers when using the command line interface.

The IBM ServeRAID Manager is a graphical utility to manage locally attached and network attached storage. Some management tasks cannot be performed from within the ServeRAID manager console. The command line interface is required instead. One example is the Flashcopy feature. This can only be managed from the command line interface.

An advantage of the IBM ServeRAID Manager is the ability to send notifications when problems with the managed storage occurs. Those notifications can be sent as operating system pop-ups, email or SNMP. This alerting can be configured from the ServeRAID manager console.

4.2 TotalStorage DS300 and DS400 command line interface (CLI)

password operator

The command line interface (CLI) is accessed via a Telnet session to one of the management ports of the controllers installed in the DS300/DS400. By default you have access as operator without the need for a password. You can get information about the health of your storage subsystem. A password can be assigned to the operator account to limit the access to persons who have the need for access.

Entering the command **administrator** gives you administrative rights to manage the storage subsystem. You need a password to authenticate. By default is the TotalStorage DS300 and DS400 shipped with the administrator password **PASSW0RD** (with a zero instead a O, all capital letters).

Note: Some firmware versions have a setup wizard which sets the password of the administrator to **passw0rd**.

Help can be retrieved in the command line interface by entering the command **help**. A list with the available command in the current section are provided. The example shows all command available in interface.

```
DS300-1[A] (interface)# help

manage network interfaces

failback               fail back interfaces to their original controller
info                   show information for some or all interfaces
list                   list all interfaces
manage                 manage an individual interface
statistics             show throughput information for some or all
interfaces


end                    leave this context and return to main mode
quit                   leave the CLI
help/?                 display this output

DS300-1[A] (interface)#
```

Figure 4-1 Help for interface

To get detailed help to a specific command, enter the command followed by the option help. The example shows help for the command **interface manage**.

```
DS300-1[A] (interface)# manage help

manage an individual interface

Usage:
manage <interface> ip <IP address> | <IP/mask>
                   netmask <netmask>
                   broadcast <broadcast>
                   dhcp enable | disable
                   info
                   failback
                   statistics
                   up
                   down
                   mtu <mtu>
                   rename <name>
                   speed 10 | 100 | auto

DS300-1[A] (interface)#
```

Figure 4-2 Help for interface manage

It is not required to enter a command completely as long as the beginning part of a command is unique. Pressing the tab key will complete the command after a the first part was entered. When nothing was entered a list with available commands or options is shown.

The U320 SCSI drives which can be installed in a TotalStorage DS300 or DS400 are not addressed by channel and SCSI ID. A unique fixed number is used instead. The table below gives an overview of these address numbers and the position in the enclosures. Enclosure 1 is always the TotalStorage DS300 or DS400. Enclosure 2 and 3 are the two EXP400 which might be attached to a TotalStorage DS400.

Different naming conventions are used to name a single drive and the position in the enclosure. ServeRAID Manager use channel and SCSI ID. The command line interface uses the enclosure number and the slot number beginning at 0. Documentation specially for the EXP400 uses slot 1 to 14. The following three tables will give an overview of the different names. It should be useful when working extensively with the CLI or analyzing the logs.

Table 4-1 Enclosure 1 — TotalStorage DS300 or DS400

Address number	Enclosure	SCSI Channel	Slot CLI [a]	SCSI ID	Front bezel number
65536	1	2	0	0	0
65792	1	2	1	1	1
66048	1	2	2	2	2
66304	1	2	3	3	3
66560	1	2	4	4	4
66816	1	2	5	5	5
67584	1	2	6	8	8
2304	1	1	7	9	9
2560	1	1	8	10	10
2816	1	1	9	11	11
3072	1	1	10	12	12
3328	1	1	11	13	13
3584	1	1	12	14	14
3840	1	1	13	15	15

a. Slot number as shown in the CLI

Table 4-2 Enclosure 2 — EXP400

Address number	Enclosure	SCSI Channel	Slot CLI	SCSI ID	Slot Number
136608	2	3	0	0	1
136864	2	3	1	1	2
137120	2	3	2	2	3
137376	2	3	3	3	4
137632	2	3	4	4	5
137888	2	3	5	5	6
DO NOT USE			6	6	7

Address number	Enclosure	SCSI Channel	Slot CLI	SCSI ID	Slot Number
138656	2	3	7	8	8
138912	2	3	8	9	9
139168	2	3	9	10	10
139424	2	3	10	11	11
139680	2	3	11	12	12
139936	2	3	12	13	13
140192	2	3	13	14	14

Table 4-3 Enclosure 3 — EXP400

Address number	Enclosure	SCSI Channel	Slot CLI	SCSI ID	Slot Number
196608	3	4	0	0	1
196864	3	4	1	1	2
197120	3	4	2	2	3
197376	3	4	3	3	4
197632	3	4	4	4	5
197888	3	4	5	5	6
DO NOT USE			6	6	7
198656	3	4	7	8	8
198912	3	4	8	9	9
199168	3	4	9	10	10
199424	3	4	10	11	11
199680	3	4	11	12	12
199936	3	4	12	13	13
200192	3	4	13	14	14

The CLI is described in detail in the DS300/DS400 Problem Determination MAP. For your convenience the CLI is included in Appendix A, “Command line interface reference” on page 565.

4.3 IBM ServeRAID Manager

IBM ServeRAID Manager is a graphical tool to manage, administer and configure onboard RAID controllers in IBM @server xSeries servers like LSI1030, IBM ServeRAID controller, Adaptec hostRAID controllers and network attached storage enclosures like IBM TotalStorage DS300 and DS400. These RAID controllers can be installed in a local or remote system. IBM ServeRAID Manager consists of multiple components:

- ServeRAID Manager console
- ServeRAID Manager agent
- Flashcopy agent
- Command line utilities

ServeRAID Manager console is the graphical component.

ServeRAID Manager agent is used to manage local and PCI adapter based RAID controllers. It is possible to manage them also over an existing IP network.

Flashcopy agent service comes with an additional command line interface, SSTOOL.EXE. This service is used to configure flashcopy and snapshots for locally attached storage. This service and command line interface cannot be used with TotalStorage DS300 and DS400.

The IBM ServeRAID Manager can be installed with a few command line utilities. Theses command line utilities are:

- ARCCONF.EXE - used to manage the IBM ServeRAID 8i RAID adapter.
- HRCONF.EXE - used to manage IBM ServeRAID 7e integrated RAID controllers based on Adaptec Host Raid controllers in the operating system. DOS uses two different utilities for SCSI and SATA based controllers: RAIDSEL.EXE and ACUICHSV.EXE.
- USER_TOOL.EXE - used to create, delete and modify Management Station user accounts. This utility is installed along with the Management Station service.
- SSTOOL.EXE - used to establish SnapShots from logical volumes. This tool is used only for supported locally attached storage.

HRCONF.EXE and ARCCONF.EXE are part of the command line tools.

4.3.1 Management Station

The Management Station service is used to talk via IP to a TotalStorage DS300 and DS400. It converts the output of the TotalStorage DS300 or DS400 to

something which can be understood by the ServeRAID Manager console. The ServeRAID Manager console converts this output again to present this in a readable format for the user like the event log. Management Station software is available from the Adaptec Web site at the following URL:

http://www.adaptec.com/

Important: As of 1 December, 2005, the ServeRAID Manager CD will no longer include the Management Station software. Management Station software will be a separate download from the Adaptec Web site.

Management Station, Flashcopy agent and ServeRAID Manager agent run as services on the system where the ServeRAID Manager is installed.

4.3.2 ServeRAID Manager requirements

To manage a TotalStorage DS300 or DS400, ServeRAID Manager requires two components:

- Management Station service
- Client station or console

The Management Station service is a new feature released together with the TotalStorage DS300 and DS400. The Management Station is required for network attached storage and communicates over an IP network connection with the DS300 and DS400 subsystems. Previous ServeRAID Manager could only manage direct attached storage.

Management Station system requirements

To successfully install and run the Management Station, you need:

- PC with an Intel-compatible 500 MHz processor (IA32, AMD32 or AMD64)
- a minimum of 256 MB RAM
- 35 MB hard disk drive space
- Microsoft Windows Advanced Server 2003, Windows XP, or Windows 2000

Client station system requirements

To successfully install and run an IBM ServeRAID Manager client station, you need:

- PC with an Intel-compatible 500 MHz processor (IA32, AMD32 or AMD64)
- a minimum of 256 MB RAM
- 35 MB hard disk drive space

- 256 color video mode

On client station systems, the ServeRAID Manager is supported on a broad range of operating systems, such as Windows, Linux, NetWare, UnixWare and OpenServer.

Note: When installing the ServeRAID Manager on client stations, refer to the OS-specific Readme.txt files on the ServeRAID Manager Application CD for a list of supported operating systems and other operating system specific information.

4.4 Initial network setup of TotalStorage DS300 and DS400

This section explains the setup of the network addresses of the TotalStorage DS300 and DS400 using the command line interface. IP address setup for the management interface cannot be performed using the IBM ServeRAID Manager. The TotalStorage DS300 and DS400 are shipped with default IP addresses. Table 4-4 contains the default addresses for the management network ports of a TotalStorage DS300 and DS400. The iSCSI ports of the DS300 do not have IP addresses configured.

Note: Please note the slightly unusual network mask.

Table 4-4 TotalStorage DS300 and DS400 management port default network addresses

Controller	Address
Controller A - Management port	192.168.70.123/255.255.0.0
Controller B - Management port	192.168.70.124/255.255.0.0

Note: When installing *multiple* TotalStorage DS300, DS400 or IBM BladeCenters, be sure not to run the unconfigured products at the same time. IP address conflicts will occur.

The link speed of the management ports can be set for every controller using the ServeRAID Manager. Global settings such as default gateway, DNS server, domain and host name can be set for the enclosure using the ServeRAID Manager.

This following sections set the IP addresses as shown in Table 4-5 on page 76.

Table 4-5 TotalStorage DS300 and DS400 network addresses

DSx00	Controller	Port	IP	Netmask	Broadcast
DS300	A	ETH0	192.168.1.230	255.255.225.224	192.168.1.255
DS300	A	ETH2	192.168.1.61	255.255.255.192	192.168.1.63
DS300	A	ETH3	192.168.1.125	255.255.255.192	192.168.1.127
DS300	B	ETH0	192.168.1.231	255.255.225.224	192.168.1.255
DS300	B	ETH2	192.168.1.62	255.255.255.192	192.168.1.63
DS300	B	ETH3	192.168.1.126	255.255.255.192	192.168.1.127
DS400	A	ETH0	192.168.1.226	255.255.225.224	192.168.1.255
DS400	B	ETH1	192.168.1.227	255.255.225.192	192.168.1.255

4.4.1 Setup of management ports of TotalStorage DS300 and DS400

Use a Telnet client to connect to the default IP address of the management port of the TotalStorage DS300 or DS400. The default IP address of the management port of controller A is 192.168.70.123/255.255.0.0. and 192.168.70.124/255.255.0.0 of controller B. It is possible that only one controller is installed in the TotalStorage DS300 or DS400. A crosslink cable can be used to connect a Thinkpad to the management port of controller A to perform the initial setup.

Note: Open a Telnet session to only one of the available management ports.

After establishing the connection to the management port with a Telnet client, the controller presents a setup wizard to perform the initial setup. This is only done when the controller is in factory default state. The first line of the output shows that the system is in a factory default state. The last line asks for a host name of the TotalStorage DS300. This is a unique value across both controllers. After answering all the questions and replying “yes” to Finish, the changes are applied immediately. See Figure 4-3 on page 77. This results in a loss of the Telnet connect. A reconnect to the new assigned IP address is required. You should run the save command to make your changes static and write them to flash.

```
:::::: Factory Default State ::::::

----[ IBM DS300: Command Line Interface ]---------------------------------

               Copyright:  IBM (c)2003-2004

                Firmware:  IBM TotalStorage DS Series (Build 0750)
                           [27/07/2005][6.21]

  Peer controller status:  Ready


           Typing 'help':  provides an overview of the commands available to
                           configure the controller.

                 TAB key:  use the TAB key as a means of getting assistance
                           in completing any commands that you issue.


Hostname [DS300]: DS300-1
Domain name [ibm.com]: rivers.local
SAN name []: RIVERSSAN
System Date [18:59:31@09/08/05]: 18:55:00@09/08/05
Timezone [GMT-5]: GMT-6
New administrator password? [No]:
New operator password? [Yes]: No
Default gateway [None]: 192.168.1.225
DNS servers [None]: 172.16.0.2 172.16.0.3
Use DHCP? [No]:
IP address [192.168.70.123]: 192.168.1.230
Netmask [255.255.0.0]: 255.255.255.224
Broadcast Address [192.168.255.255]: 192.168.1.255
Peer IP address [192.168.70.124]: 192.168.1.231
Peer Netmask [255.255.0.0]: 255.255.255.224
Peer Broadcast Address [192.168.255.255]: 192.168.1.255
Finish? [Yes]:
Please wait...
```

Figure 4-3 Setup wizard

The setup wizard can be started again by running system setup.

```
DS300-1[A]# setup
.
.
.
```

Figure 4-4 Starting the setup wizard

In the setup wizard, the command **back** can be used to go back to the previous question. Command **abort** will cancel the wizard.

To bring the system back into a default state see 4.4.4, “Reset network and account setting to factory default” on page 91.

If a question has a predefined or default answer, the answer appears in square brackets [ ]. Some options take an arbitrary string, in which case the commands back and abort cannot be used. In those circumstances, enclose the answer in double-quotes " " so that the answer will be accepted even if it is the same as a command.

Hostname — Supply the name by which the system is to be known.The special BACK command here repeats this prompt.

Domain name — Supply the domain in which the system belongs.

SAN name — Supply the Storage Area Network in which the system belongs.

system Date — Supply the system time and date in the format 24 hours:minutes:seconds@day/month/year.

timezone — Supply the system time zone. Use capital letters for the time zone. for example GMT-6 or GMT+1

New administrator password — Define whether an administrator password is required. If a password is required, you are prompted for that password, which must be supplied twice. Press **Ctrl** to leave the password unchanged. Press Return (or Enter) twice to require no administrator password.

New operator password — Define whether an operator password is required.If a password is required, you are prompted for that password, which must be supplied twice. Press **Ctrl** to leave the password unchanged. Press **Return** (or **Enter**) twice to require no administrator password.

Default gateway — Enter the name or IP number of a router to which packets for destinations that do not have routes defined for them in

	the system should be sent. Enter no to clear the default gateway.
DNS servers	Enter the IP numbers of up to 3 systems that will act as DNS servers. Enter no to clear the list of DNS servers. Only 2 DNS servers can be configured by the ServeRAID manager. The third server will not be visible or changeable from ServeRAID Manager console.
Use DHCP	Reply yes or no to determine whether DHCP should be used to discover the system's network values. If you enter yes, the wizard takes you to the Finish question. If you enter no, the wizard continues with questions about the systems management ethernet interface and the peer interface in case of a dual controller DS300 or DS400.
IP address	Enter the IP number in the conventional dotted quad format as shown in this example: 192.193.194.195. This will set the IP address of the controller where the Telnet session is connected to.
Netmask	Define the size of the netmask of the management interface. A netmask can be represented in a number of ways. The examples shown here all refer to the same netmask: /23 The number of bits set in the network part. 255.255.254.0 The set bits displayed in the same way as an IP number 0xfffffe00 The mask show as a hex number
Broadcast Address	Enter the address to be used for broadcasts for the management interface where the Telnet session is connected to.
Peer IP address	Enter the IP number in the conventional dotted quad format as shown in this example: 192.193.194.195. This will set the IP address of the alternate controller.
Peer Netmask	Define the size of the netmask of the management interface of the alternate controller.
Peer Broadcast Address	Enter the address to be used for broadcasts management interface of the alternate controller.
Finish?	Enter yes to apply the values you have entered, or no to return to the first question and change the values you have entered.

Note: If you changed the IP address, you must Telnet to the new IP address to reestablish the Telnet session.

The setup wizard is the same for a TotalStorage DS300 and DS400. Only the available management ports will be defined. When no alternate controller is available, the questions for peer IP, subnet mask and gateway will not be shown.

Note: The iSCSI interfaces for TotalStorage DS300 1701-1RS and 1701-2RD must be set with commands from the CLI or from the ServeRAID manager. See 4.4.2, “Setting the IP addresses of the iSCSI ports from the command line interface” on page 80 or 4.4.3, “Setting iSCSI port IP addresses from ServeRAID Manager” on page 87 to define the iSCSI interface network addresses.

The TotalStorage DS300 1701-1RL has only one ethernet port. This port is shared for management purposes and iSCSI traffic. The network settings for this port are defined with the setup wizard. It is not possible to set this IP address from an ServeRAID Manager console.

4.4.2 Setting the IP addresses of the iSCSI ports from the command line interface

After the management port interfaces are configured by the setup wizard, the network addresses of the iSCSI interfaces can be configured from the command line interface. it is not required to connect to both controllers. The IP settings can be performed from connection to only one controller. the controller is specified by a letter A or B behind the network interface. For example the ETH2 of controller B is ETH2[B].

1. Re-establish a Telnet session to one of the management ports of the TotalStorage DS300. Logon with administrative rights. To do this enter administrator and press enter. Default password of the administrator is PASSW0RD with a zero instead the letter O. Depending on the firmware level on the TotalStorage DS300 controller the password might be lower case or upper case. When a password for the operator is applied, as this is the default during the execution of the setup wizard, this password has to be entered. See Figure 4-5 on page 81.

```
----[ IBM DS300: Command Line Interface ]---------------------------------

               Copyright:  IBM (c)2003-2004

                Firmware:  IBM TotalStorage DS Series (Build 0750)
                           [27/07/2005][6.21]

  Peer controller status:  Ready


           Typing 'help':  provides an overview of the commands available to
                           configure the controller.

                 TAB key:  use the TAB key as a means of getting assistance
                           in completing any commands that you issue.



DS300-1[A]> administrator

Administrator level password: ********

DS300-1[A]#
```

Figure 4-5 Telnet logon

2. Get the current network settings of the first iSCSI interface on controller A by entering the command **interface manage eth2[A] info**. There is no IP address defined for this interface. See Figure 4-6 on page 82.

```
DS300-1[A]# interface manage eth2[A] info


----[ Interface 'eth2[A]' ]-----------------------------------------------

          Controller:  [A] local
            Services:  iscsi
              Status:  Enabled
         MAC address:  00:00:D1:DE:1B:70
          DHCP DISABLED
          IP address:
             Netmask:
           Broadcast:
                 MTU:  1500
               Speed:  AUTO


DS300-1[A]#
```

Figure 4-6 iSCSI interface information

3. Assign the IP address of the first iSCSI interface. The interface is called ETH2. A letter A or B in square brackets specifies the controller to which the interface belongs. Enter the command **interface manage eth2[A] ip** `192.168.1.61` and press **Enter**.

> **Note:** Broadcast address and netmask do not fit right now. This will be changed in the next steps.

```
DS300-1[A]# interface manage eth2[A] ip 192.168.1.61
Please wait.

Interface values are now:
          IP address:  192.168.1.61
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.0


DS300-1[A]#
```

Figure 4-7 Set iSCSI IP interface address

4. Change the subnet mask of the interface ETH2[A] by entering **interface manage eth2[A] netmask** `255.255.255.192` and press **Enter**.

```
DS300-1[A]# interface manage eth2[A] netmask 255.255.255.192
Please wait.

Interface values are now:
          IP address:  192.168.1.61
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.192


DS300-1[A]#
```

Figure 4-8 Set iSCSI interface subnet mask

5. Change the broadcast address of the interface ETH2[A] by entering **interface manage eth2[A] broadcast** `192.168.1.63` and press **Enter**.

```
DS300-1[A]# interface manage eth2[A] broadcast 192.168.1.63
Please wait.

Interface values are now:
          IP address:  192.168.1.61
   Broadcast address:  192.168.1.63
             Netmask:  255.255.255.192


DS300-1[A]#
```

Figure 4-9 Set iSCSI interface broadcast address

6. Set the IP address, subnet mask and broadcast address for the first iSCSI interface of controller B ETH2[B].

```
DS300-1[A]# interface manage eth2[B] ip 192.168.1.62
Please wait.

Interface values are now:
          IP address:  192.168.1.62
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.0


DS300-1[A]# interface manage eth2[B] netmask 255.255.255.192
Please wait.

Interface values are now:
          IP address:  192.168.1.62
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.192


DS300-1[A]# interface manage eth2[B] broadcast 192.168.1.63
Please wait.

Interface values are now:
          IP address:  192.168.1.62
   Broadcast address:  192.168.1.63
             Netmask:  255.255.255.192


DS300-1[A]#
```

Figure 4-10 Configure ETH2[B}

7. Set the IP address, subnet mask and broadcast address for the second iSCSI interface of controller A ETH3[A].

```
DS300-1[A]# interface manage eth3[A] ip 192.168.1.125
Please wait.

Interface values are now:
          IP address:  192.168.1.125
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.0


DS300-1[A]# interface manage eth3[A] netmask 255.255.255.192
Please wait.

Interface values are now:
          IP address:  192.168.1.125
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.192


DS300-1[A]# interface manage eth3[A] broadcast 192.168.1.127
Please wait.

Interface values are now:
          IP address:  192.168.1.125
   Broadcast address:  192.168.1.127
             Netmask:  255.255.255.192


DS300-1[A]#
```

Figure 4-11 Configure ETH3[A]

8. Set the IP address, subnet mask and broadcast address for the second iSCSI interface of controller B ETH3[B].

```
DS300-1[A]# interface manage eth3[B] ip 192.168.1.126
Please wait.

Interface values are now:
          IP address:  192.168.1.126
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.0


DS300-1[A]# interface manage eth3[B] netmask 255.255.255.192
Please wait.

Interface values are now:
          IP address:  192.168.1.126
   Broadcast address:  192.168.1.255
             Netmask:  255.255.255.192


DS300-1[A]# interface manage eth3[B] broadcast 192.168.1.127
Please wait.

Interface values are now:
          IP address:  192.168.1.126
   Broadcast address:  192.168.1.127
             Netmask:  255.255.255.192


DS300-1[A]#
```

Figure 4-12 Configure ETH3[B]

The network setup is now completed. The TotalStorage DS300 can now be managed by the IBM ServeRAID Manager. The agent must now be added to the Management Station. See section “Configure and run IBM ServeRAID Manager” on page 97.

To change the network settings for the management ports, use ETH0 on single controller models and ETH0[A] and ETH0[B] for dual controller models with the commands shown above.

Note: Changes of the network configuration of a TotalStorage DS300 or DS400 need to be saved to be available after a reboot. Run the command **save** in the CLI to save the configuration. This is required on only one controller.

4.4.3 Setting iSCSI port IP addresses from ServeRAID Manager

After the management port interfaces are configured with the setup wizard, the IBM ServeRAID manager can be used to define the network addresses of the iSCSI interfaces. To be able to do this, a few steps must be done first.

1. Install the IBM ServeRAID Manager on a host running Microsoft Windows. See 4.5, “IBM ServeRAID Manager installation on Microsoft Windows” on page 91
2. Start the IBM ServeRAID Manager console. See 4.7.1, “Start IBM ServeRAID Manager” on page 98.
3. Add a Management Station to the ServeRAID Manager console. See 4.7.2, “Add a Management Station to a ServeRAID Manager” on page 98.
4. Add an agent to the IBM ServeRAID Manager Management Station.

Follow these steps to configure the iSCSI interfaces on both controllers:

1. Right-click controller A to open the context menu. Select from the context menu **Configure ethernet port** → **ETH 2 (iSCSI)**.

interface manage eth2[A] info

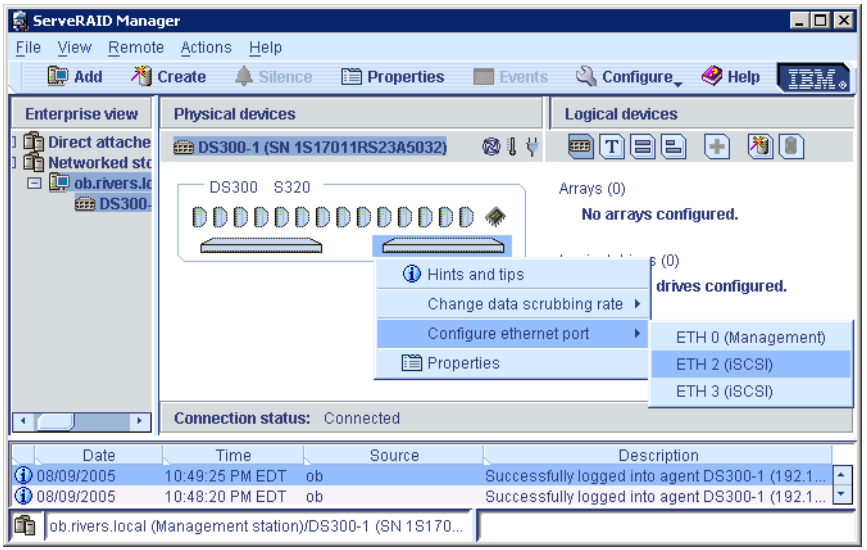

Figure 4-13 Configure ethernet port ETH2 controller A

2. Enter the IP address `192.168.1.61`, subnet mask `255.255.255.192` and broadcast address `192.168.1.63` into the dialog. Click **OK**. The settings become immediately effective.

interface manage eth2[A] ip 192.168.1.61

interface manage eth2[A] netmask 255.255.255.192

interface manage eth2[A] broadcast 192.168.1.63

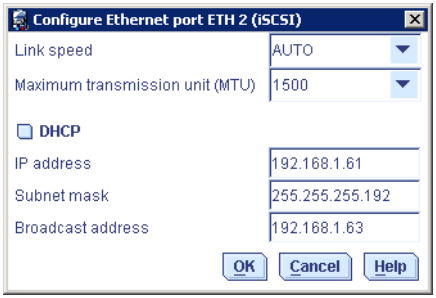

Figure 4-14 Ethernet port ETH2[A] settings

3. Right-click controller A to open the context menu. Select from the context menu **Configure ethernet port → ETH 3 (iSCSI)**.

interface manage eth3[A] info

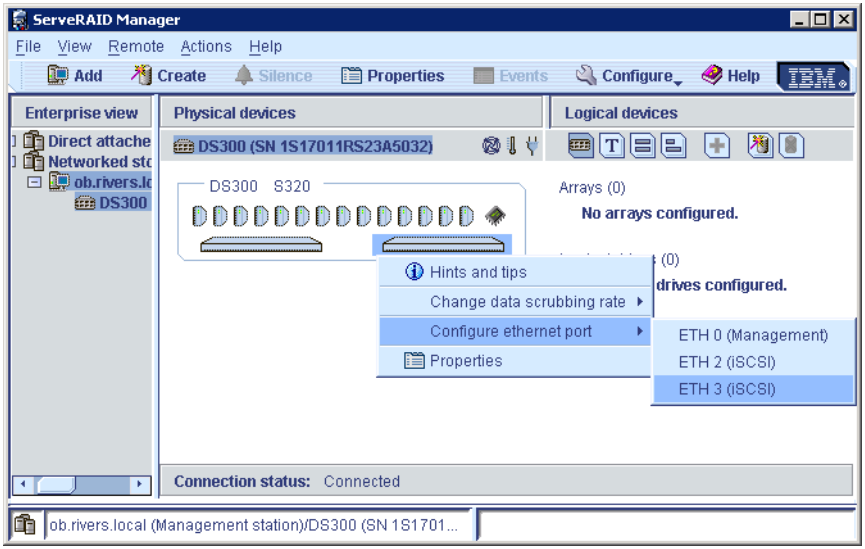

Figure 4-15 Configure ethernet port ETH3 controller A

4. Enter the IP address 192.168.1.125, subnet mask 255.255.255.192 and broadcast address 192.168.1.127 into the dialog. Click **OK**. The settings become immediately effective.

interface manage eth3[A] ip 192.168.1.125

interface manage eth3[A] netmask 255.255.255.192

interface manage eth3[A] broadcast 192.168.1.127

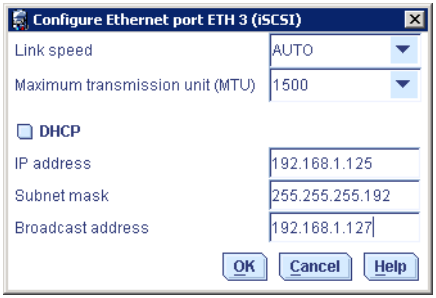

Figure 4-16 Ethernet port ETH3[A] settings

5. Right-click controller A to open the context menu. Select from the context menu **Configure ethernet port → ETH 2 (iSCSI)**.

interface manage eth2[B] info

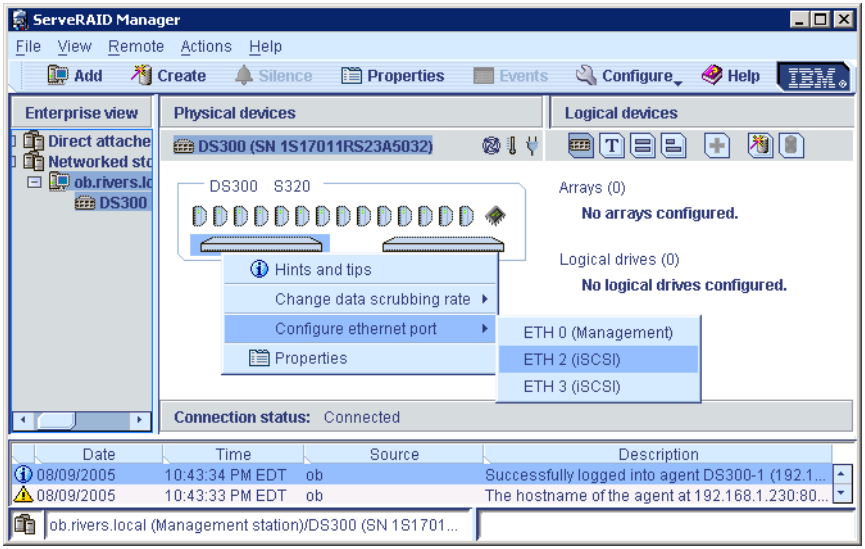

Figure 4-17 Configure ethernet port ETH2 controller B

6. Enter the IP address 192.168.1.62, subnet mask 255.255.255.192 and broadcast address 192.168.1.63 into the dialog. Click **OK**. The settings become immediately effective.

interface manage eth2[B] ip 192.168.1.62

interface manage eth2[B] netmask 255.255.255.192

interface manage eth2[B] broadcast 192.168.1.63

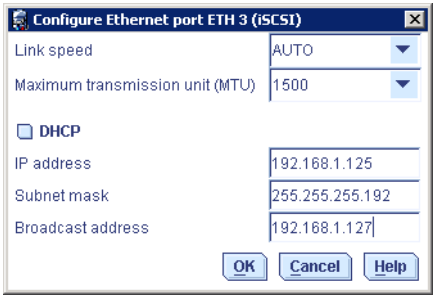

Figure 4-18 Ethernet port ETH2[B] settings

7. Right-click controller B to open the context menu. Select from the context menu **Configure ethernet port → ETH 3 (iSCSI)**.

interface manage eth3[B] info

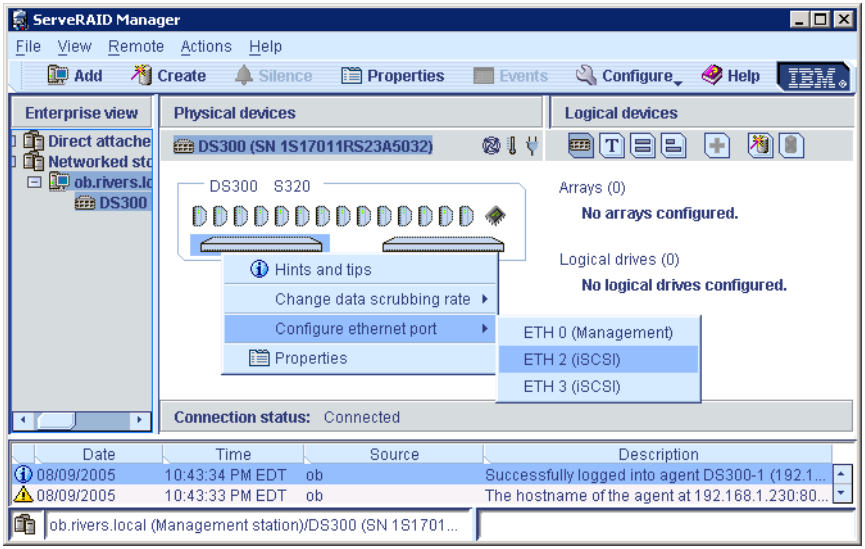

Figure 4-19 Configure ethernet port ETH3 controller B

8. Enter the IP address 192.168.1.126, subnet mask 255.255.255.192 and broadcast address 192.168.1.127 into the dialog. Click **OK**. The settings become effective immediately.

interface manage eth3[B] ip 192.168.1.126

interface manage eth3[B] netmask 255.255.255.192

interface manage eth3[B] broadcast 192.168.1.127

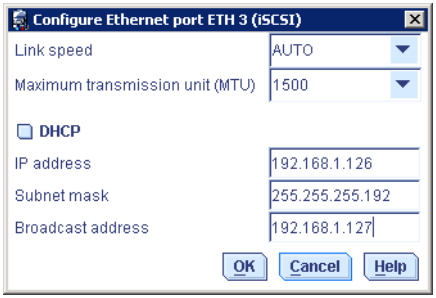

Figure 4-20 Ethernet port ETH3[B] settings

4.4.4 Reset network and account setting to factory default

To reset a TotalStorage DS300 or DS400 to factory default state, run the **system defaults** command. This command will reset all network settings, feature keys, and the passwords of the operator and administrator account. It will not delete arrays, logical drives, initiators or access control lists. These need to be deleted manually with the ServeRAID Manager console or the command line interface. A reboot is performed automatically. Connect to the default management IP and start network setup again.

```
DS300-1[A]# system defaults
                WARNING!!!
All configured values will be lost!
Restore system to factory defaults? [No]: yes


The system is about to reboot
because it has been reset to the factory default state

DS300-1[A]#
```

Figure 4-21 Reset factory default

4.5 IBM ServeRAID Manager installation on Microsoft Windows

To install IBM ServeRAID Manager, any previous installed version of the IBM ServeRAID Manager should be uninstalled first. For more details about uninstalling IBM ServeRAID Manager see “IBM ServeRAID Manager removal on Microsoft Windows” on page 190.

1. Insert the installation CD and wait for the Autorun executable to start the installation. If this does not occur, browse the CD and click **Autorun**.

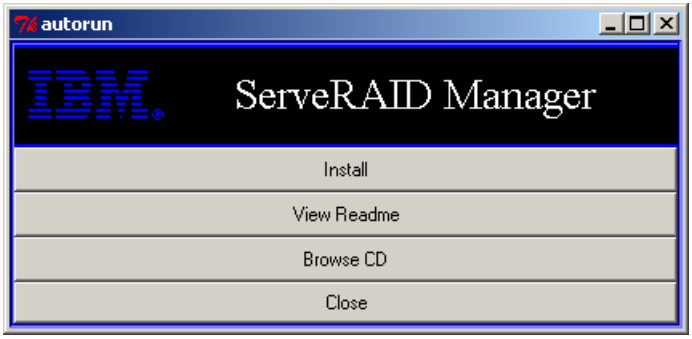

Figure 4-22 Autorun window

2. In the autorun windows click the **Install** button.
3. When the installation wizard opens, click **Next**.

Figure 4-23 InstallShield Wizard

4. Read and accept the terms of the license agreement, then click **Next**.

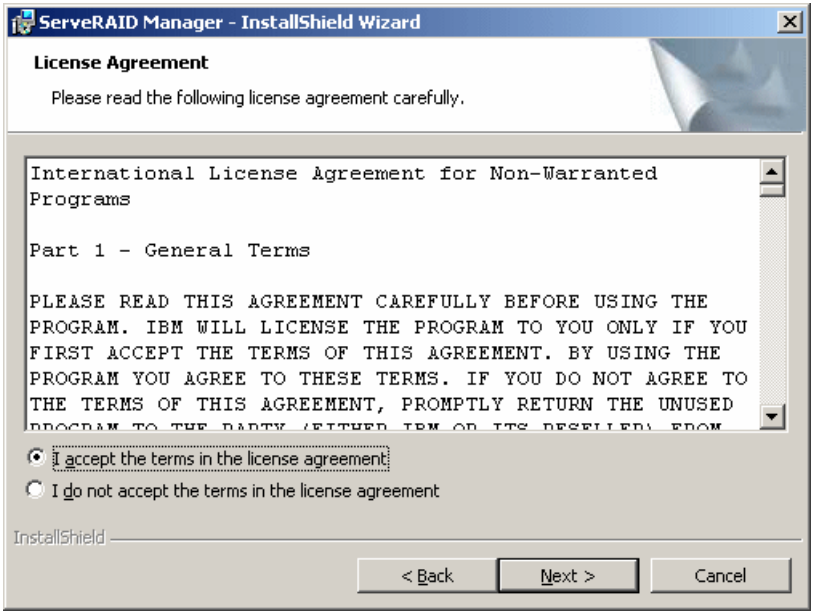

Figure 4-24 License Agreement

5. Click **Next** to accept the default installation setup.

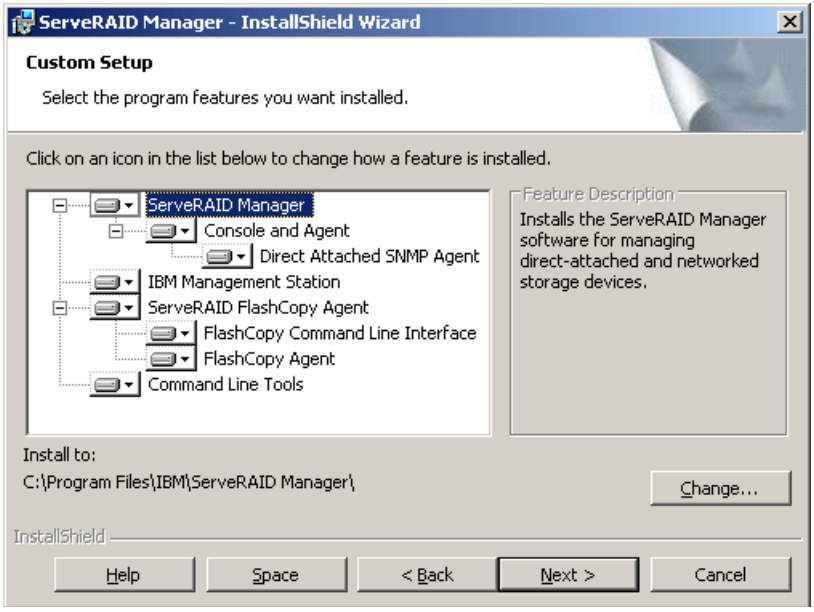

Figure 4-25 Custom setup

6. Alternatively, use the dropdown boxes to select the ServeRAID components that you want to install, then click **Next**.
7. This dialog allows you to create a user account for the ServeRAID Management Station. This is the user name and password that is used to logon onto the Management Station. By default, the user will be named `admin`. You must specify a password to proceed with installation. Click **Next**.

Note: If the account details created during installation are forgotten, new accounts can be created by the USER_TOOL.EXE. See 4.9.2, “Administer Management Station user accounts with USER_TOOL” on page 163 about more details.

More accounts can be added later by configuring the Management Station. See 4.9, “Administer Management Station user accounts” on page 162.

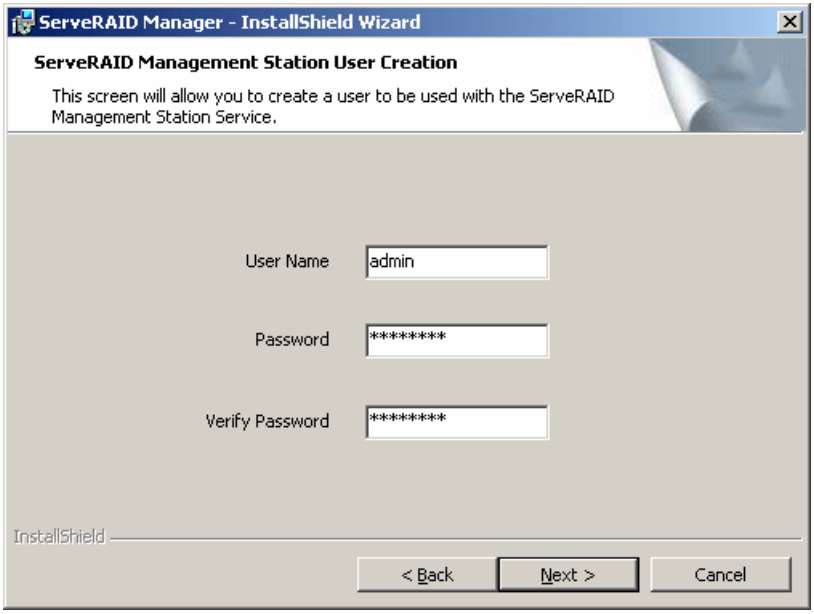

Figure 4-26 ServeRAID Management Station Account

Important: As of 1 December, 2005, the ServeRAID Manager CD will no longer include the Management Station software. Management Station software will be a separate download from the Adaptec Web site and will require a separate installation.

8. Click **Install**.

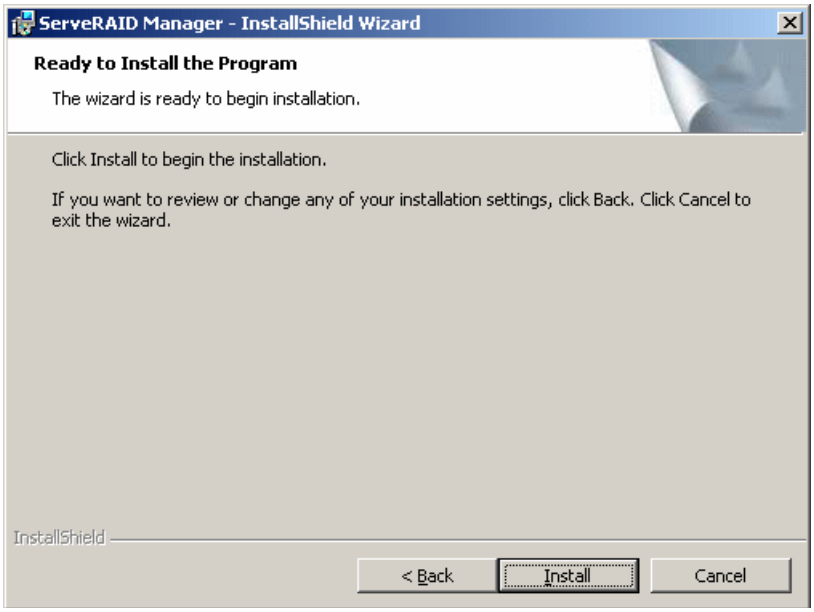

Figure 4-27 InstallShield Wizard — Ready to install

9. The installation wizard installs the software. This may take a few minutes to complete.
10. When the installation is complete, click **Finish**.

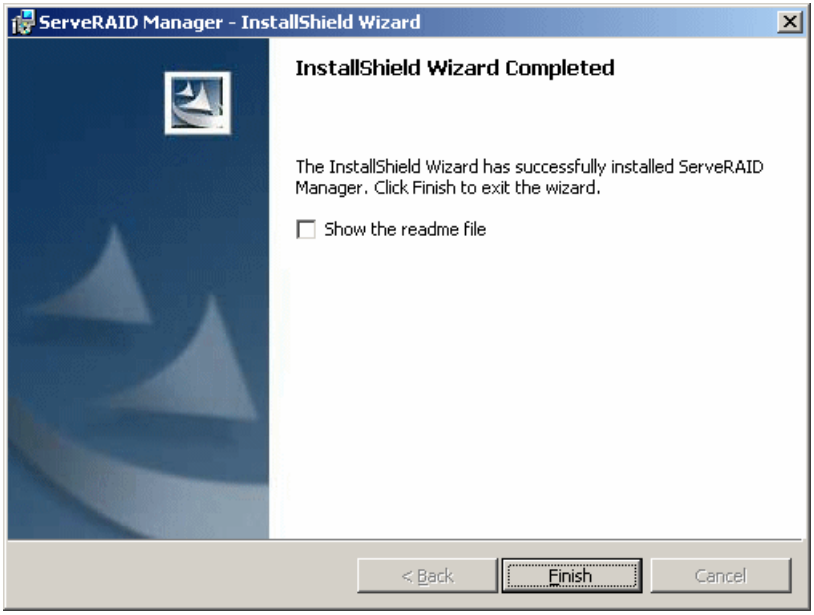

Figure 4-28 InstallShield Wizard completed

Previous versions of the IBM ServeRAID Manager have required a restart. This restart is no longer required.

The readme file is stored in the IBM ServeRAID Manager installation directory and can be read at any time. It contains the fixes of this version, new features, and troubleshooting tips.

Note: The readme might not be installed when an unattended installation is chosen. See the install.pdf in the book directory of the IBM ServeRAID Manager Application CD about details of performing unattended installation of the ServeRAID Manager.

4.6 IBM ServeRAID Manager installation on Linux

To install IBM ServeRAID Manager on Linux, follow these steps:

1. Insert the ServeRAID Manager CD and mount the CD if this is not done automatically.
2. Change to the folder **/linux/manager** on the CD.

3. Install ServeRAID Manager with the **rpm -ihv** command as shown in Figure 4-29.

```
danube:/media/dvd/linux/manager # ls
.  ..  RaidMan-8.20.i386.rpm
danube:/media/dvd/linux/manager # rpm -ihv RaidMan-8.20.i386.rpm
Preparing...                ########################################
[100%]

IBM
Version 8.20

   1:RaidMan                ########################################
[100%]
starting IBM ServeRAID Manager agent ...

raid_agent               0:off  1:off  2:off  3:on   4:off  5:on   6:off

Installation completed successfully.
The application can be started by running: /usr/RaidMan/RaidMan.sh
danube:/media/dvd/linux/manager #
```

Figure 4-29 ServeRAID Manager installation using the rpm command

4.7 Configure and run IBM ServeRAID Manager

To be able to manage a DS300/DS400, you need a ServeRAID Manager Management Station. This Management Station is used to talk via IP to the network attached storage enclosures. This storage enclosures are called agents.

Note: Only a Microsoft Windows operating system can be used to run the Management Station.

Only one Management Station can be used to manage a specific TotalStorage DS300 or DS400. However, multiple Management Stations may be used to manage different TotalStorage DS300 or DS400 subsystems. In addition, multiple TotalStorage DS300 or DS400 subsystems can be managed by a single Management Station. Before an agent can be managed by an other Management Station, the agent needs to be removed first from its current Management Station. For information about removing a agent from a Management Station see “Remove an agent from a Management Station” on page 104.

The ServeRAID Management Station service is an optional component of the IBM ServeRAID Manager. This component can be installed together with the IBM ServeRAID Manager. In the ServeRAID manager console, the local or multiple remote Management Stations can be added. Thus an installed ServeRAID Manager is required. Usable ServeRAID Manager versions are 7.2x.xx or higher. At time of writing the current version is 8.20.00.

Note: ServeRAID Manager version 8.00.xx cannot be used to manage a TotalStorage DS300 or DS400.

See 4.5, “IBM ServeRAID Manager installation on Microsoft Windows” on page 91 about the installation of the IBM ServeRAID Manager.

Start the ServeRAID Manager to add a Management Station and agent.

4.7.1 Start IBM ServeRAID Manager

Start the IBM ServeRAID Manager by clicking **Start** → **Programs** → **ServeRAID Manager** → **ServeRAID Manager** on a Microsoft Windows operating system or run **/usr/RaidMan/RaidMan.sh** on Linux.

The IBM ServeRAID Manager console opens and shows a window like Figure 4-30 on page 99.

Note: An X server is required to run the ServeRAID Manager console under Linux.

4.7.2 Add a Management Station to a ServeRAID Manager

A ServeRAID Management Station is required to manage a TotalStorage DS300 or DS400. The Management Station service can run on the local machine or on a remote machine. With ServeRAID Manager software available prior to 1 December, 2005, this Management Station software is installed together with the ServeRAID Manager console.

Important: As of 1 December, 2005, the ServeRAID Manager CD will no longer include the Management Station software. Management Station software will be a separate download from the Adaptec Web site and a separate installation.

The following steps explain the adding of a Management Station.

Note: Currently the Management Station service runs only under Windows. To manage a TotalStorage DS300 or DS400 from a Linux ServeRAID Manager console, you need to add a remote Management Station which runs on a Microsoft Windows operating system.

1. In the Enterprise view, right-click the icon for the Networked storage and select **Add management station**, as shown below.

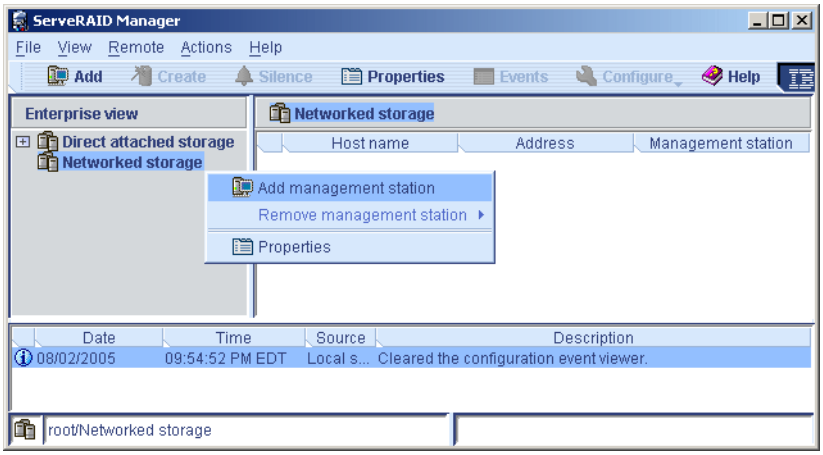

Figure 4-30 Add Management Station

2. The Add Management Station window opens.
3. In the Type field, select **Management Station**.
4. Enter the host name or TCP/IP address of the Management Station system. In case you have no separate Management Station, and the Management Station is running on the same system, use `localhost` or `127.0.0.1`.
5. Enter the Management Station user name and password, as they were defined during the ServeRAID Manager installation. (See Step 7 on page 94.)
6. Select **Save user name/password**. When the password is not saved, it needs to be entered every time the ServeRAID manager is started and access to the TotalStorage DS300 or DS400 is requested.
7. Click **Connect**.

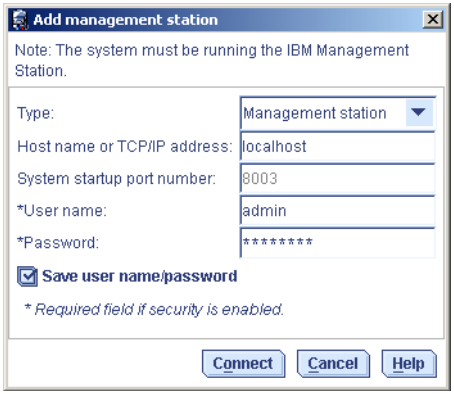

Figure 4-31 Add Management Station

The Management Station `localhost` is now added.

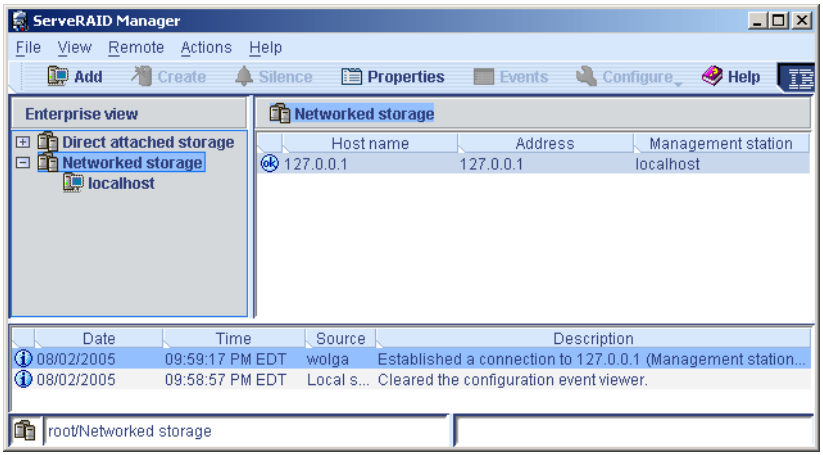

Figure 4-32 Management Station

In case you have added a remote Management Station which has already added agents, they will be usable after adding the remote Management Station. See 4.7.2, “Add a Management Station to a ServeRAID Manager” on page 98 for more details about adding a Management Station.

Multiple remote Management Stations can be added as shown in Figure 4-33 on page 101.

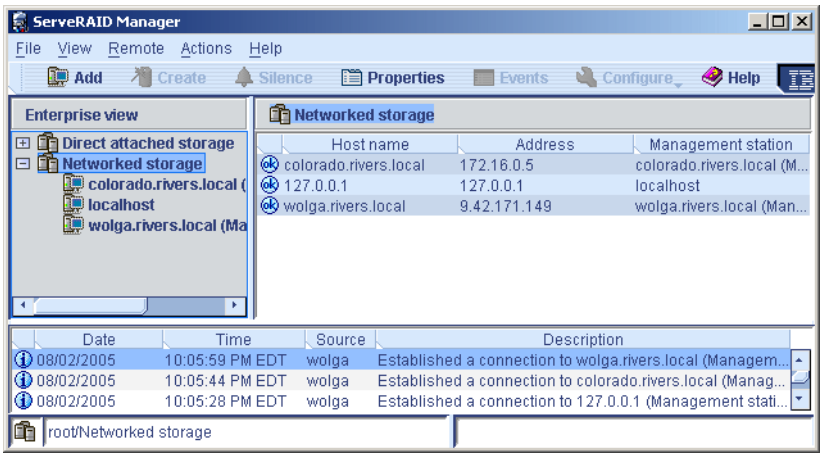

Figure 4-33 Multiple Management Stations

4.7.3 Remove Management Station from the IBM ServeRAID Manager console

Follow the steps below to remove a Management Station from the ServeRAID manager console.

1. Right-click **Networked storage**. In the context menu click **Remove management station → 127.0.0.1 (Management station)**.

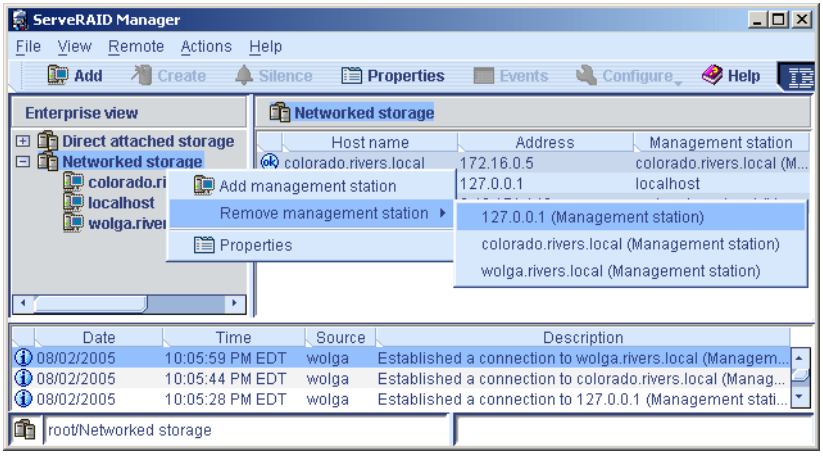

Figure 4-34 Remove Management Station from ServeRAID Manager console

2. Confirm the removal of the Management Station by clicking **Yes**.

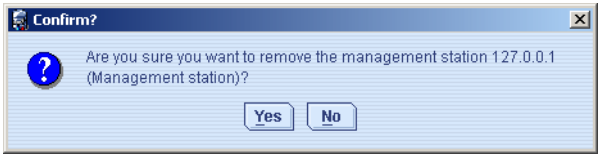

Figure 4-35 Confirm removal of the Management Station

3. Management Station is removed.

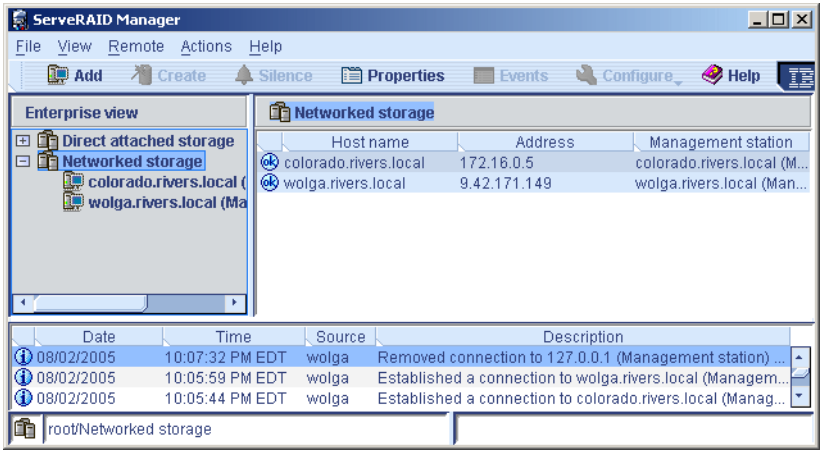

Figure 4-36 Removed Management Station

4.7.4 Add an agent to a Management Station

The agent represents a IBM TotalStorage DS300 or DS400. After adding the agent to a Management Station you will be able to manage a TotalStorage DS300 or DS400.

Note: An agent can only be managed by one Management Station. Before a agent can be managed from an other Management Station, the agent must be removed from the original Management Station. See 4.7.5, “Remove an agent from a Management Station” on page 104 for the details of removing a agent from a Management Station.

Perform the following steps to add an agent to a Management Station:

1. In the Enterprise view, right-click the Management Station in the **Networked storage** tree and select **Add agent**.

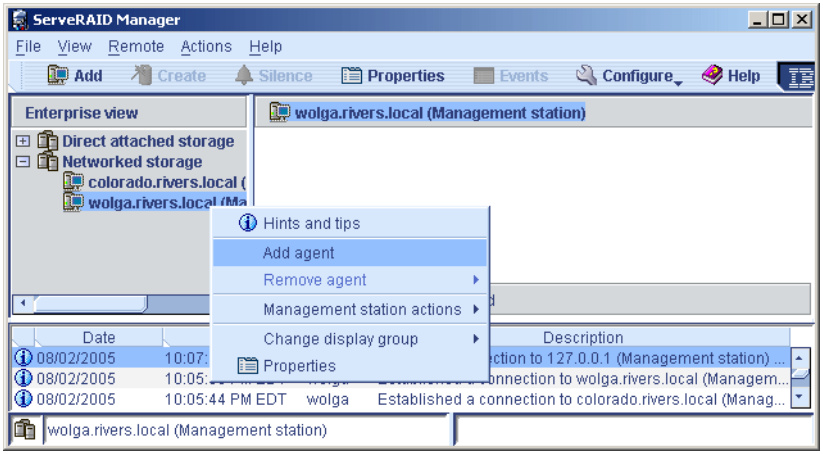

Figure 4-37 Select Add Agent

2. Type in the host name or TCP/IP address of one of the storage subsystem's management port. One IP address is enough in case you own a dual controller model.
3. Enter the administrator password of the storage subsystem.

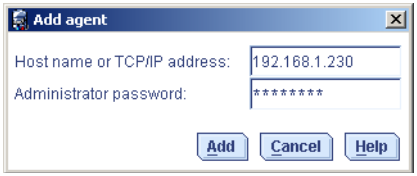

Figure 4-38 Add Agent

4. Click **Add**.
5. The Agent is added to the Management Station.

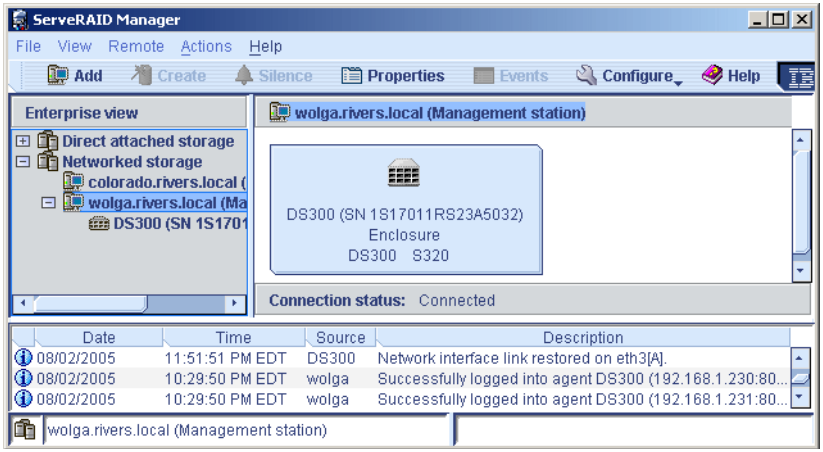

Figure 4-39 Added Agent

Multiple agents can be added to the same Management Station.

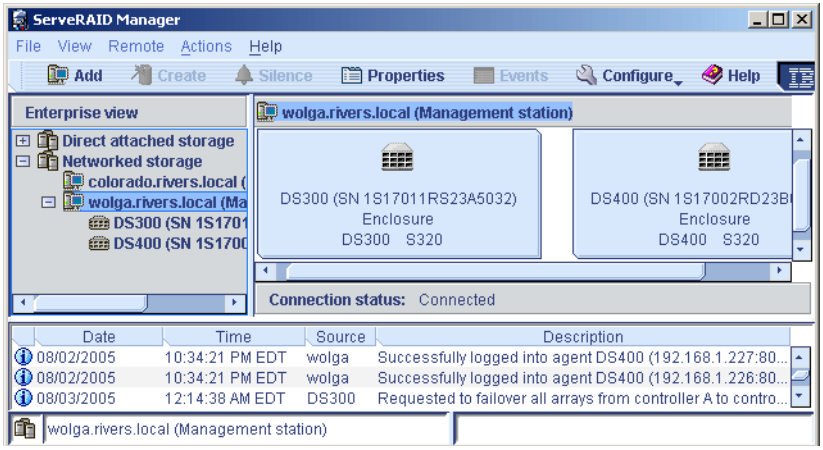

Figure 4-40 Management Station with Agents

4.7.5 Remove an agent from a Management Station

To be able to manage an agent from another Management Station or because the agent no longer exists, the agent must be removed from a Management Station. Follow the outlined steps below to remove an agent from a Management Station.

1. Right-click the Management Station which contains the agent. In the context menu select **Remove agent** → **<hostname>**.

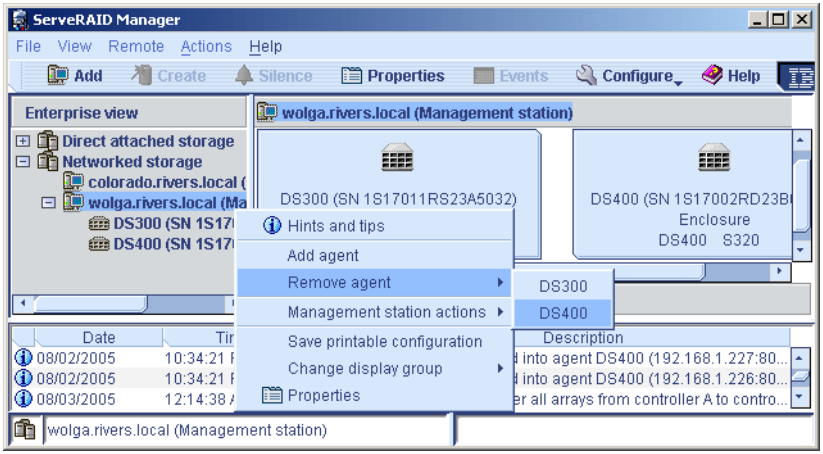

Figure 4-41 Remove agent from Management Station

2. Confirm the agent removal by clicking **Yes**.

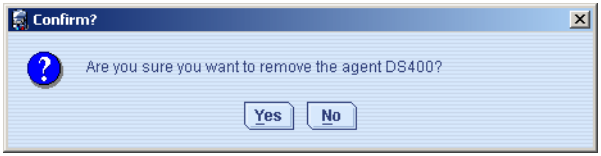

Figure 4-42 Agent removal confirmation

3. The Management Station no longer contains the agent. The agent might be added now to another Management Station.

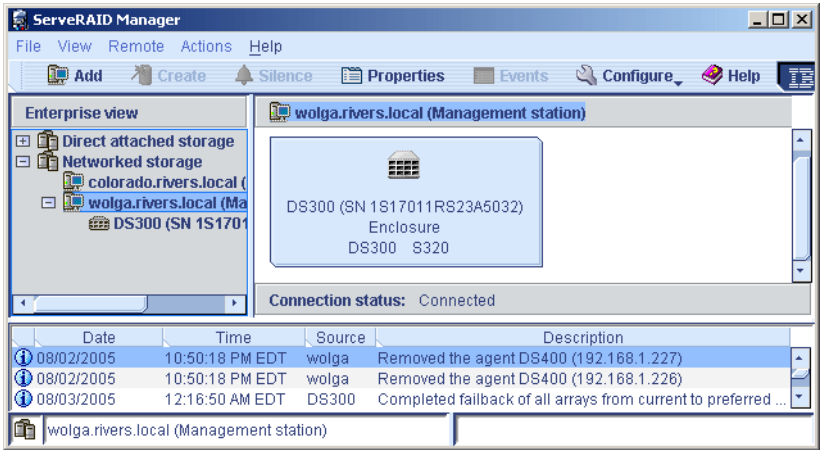

Figure 4-43 Removed agent

4.7.6 Display groups

Display groups might be used to group IBM ServeRAID Manager Management Stations together. This provides a management overview about the health of the Management Stations and agents in this group. After an agent has got a problem, the display group gets an attention mark to inform you that interactions are required. Use the property window of the display group to get a overview about all agents in the display group.

Note: Display groups are not available when a boot from a IBM ServeRAID Manager Support CD is performed or when the IBM ServeRAID Manager is used as IBM Director plug-in. IBM Director 5.10 may be able to support management of DS300 and DS400 controller.

Create a new display group

To create a new display group perform the following steps.

1. Select a Management Station within the enterprise view with a right-click. In the context menu click **Change display group** → **New group**.

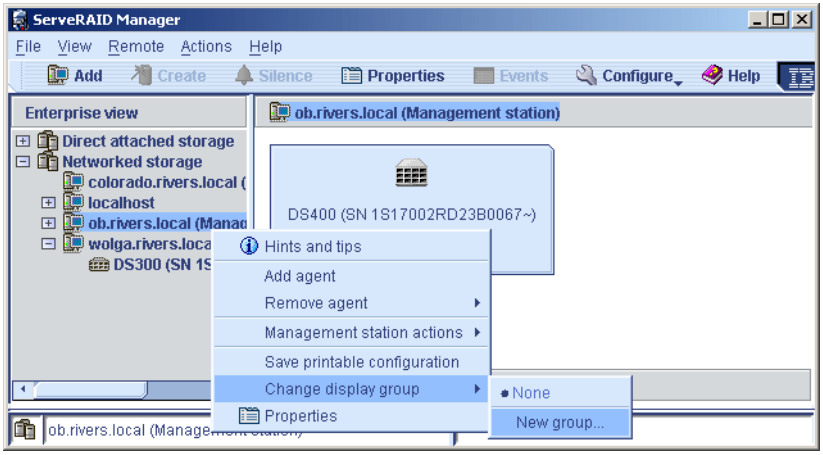

Figure 4-44 Menu change display group

2. Enter in the opened group name window a name for the display group. Click the **OK** button to create the display group.

Figure 4-45 Group name

3. In the enterprise view, the selected Management Station will be shown now under the new display group `Asia`.

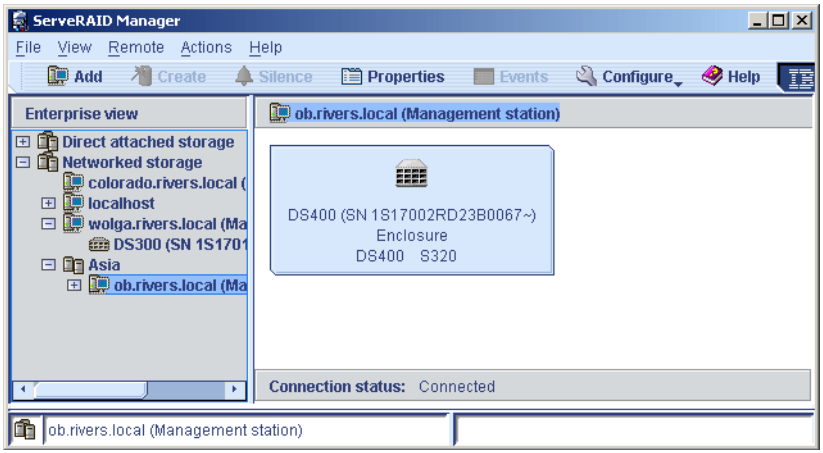

Figure 4-46 Added display group ASIA

Add a Management Station to a existing display group

To add a Management Station to a existing display group a display group must exist. See “Create a new display group” on page 106 about creating new display groups.

To add a Management Station to a existing display group perform the following steps:

1. Select a Management Station under **Networked storage** or from a other display group with right-click. In the context menu select **Change display group** → **Name of existing display group**. All existing display groups are listed in the menu. The figure below use as example the display group `Asia`.

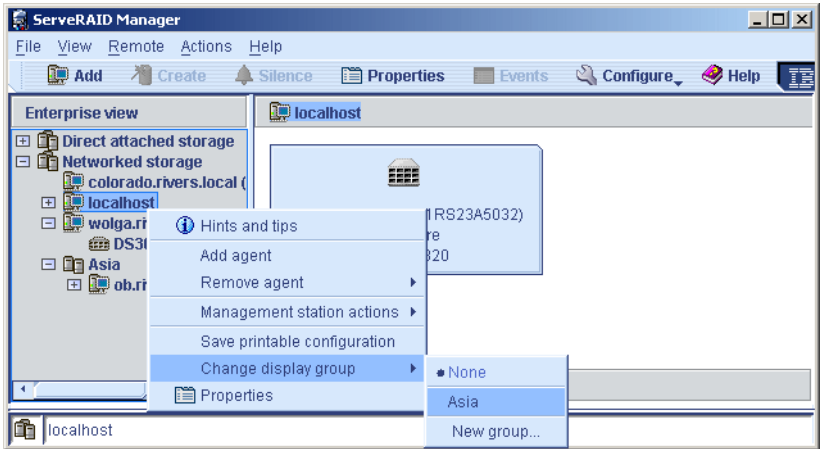

Figure 4-47 Add Management Station to existing display group

2. The Management Station `localhost` is now added to the display group `Asia`.

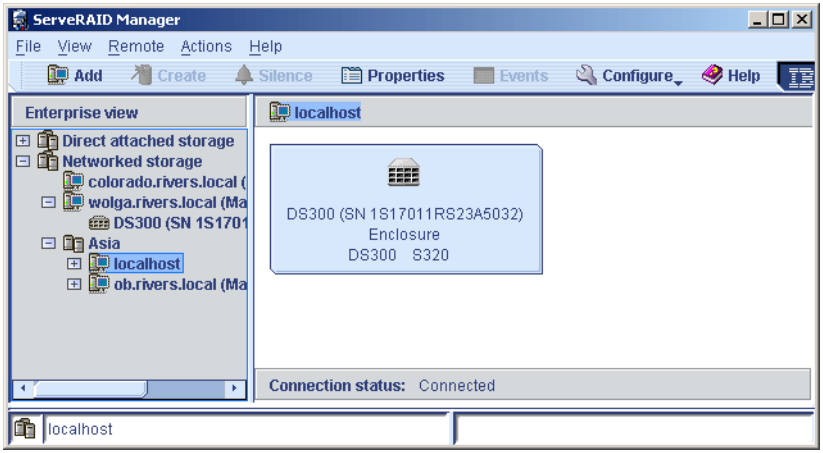

Figure 4-48 Added Management Station to existing display group

The next figure shows three display groups, two with one Management Station and one group with two Management Stations.

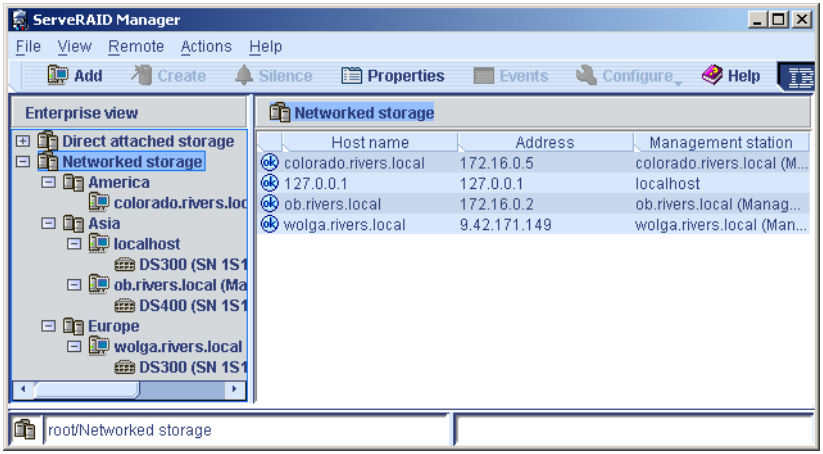

Figure 4-49 Multiple Display groups

Rename a display group

To rename a display group perform the following steps:

1. To rename a display group right-click the display group where the name needs to be changed. In the context menu click **Rename display group**.

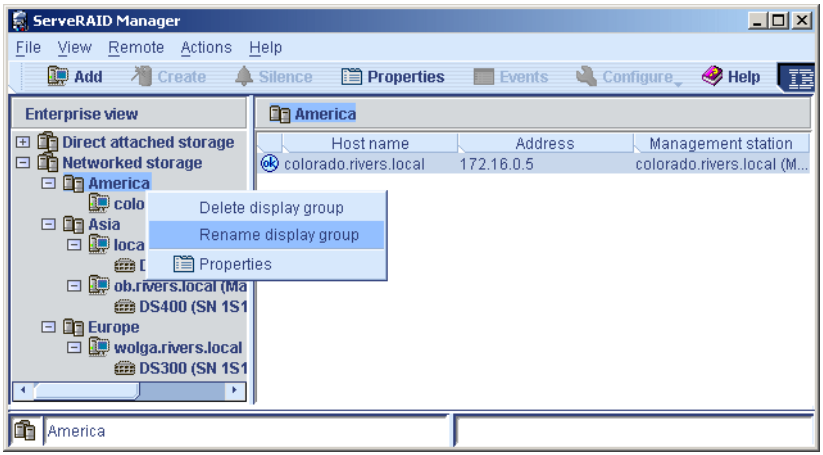

Figure 4-50 Rename

2. A group name dialog opens and shows the current display group name. Edit the current name or replace it with a new name. Then click **OK**.

Figure 4-51 Rename display group to local

3. The Group name window closes and you see the renamed display group.

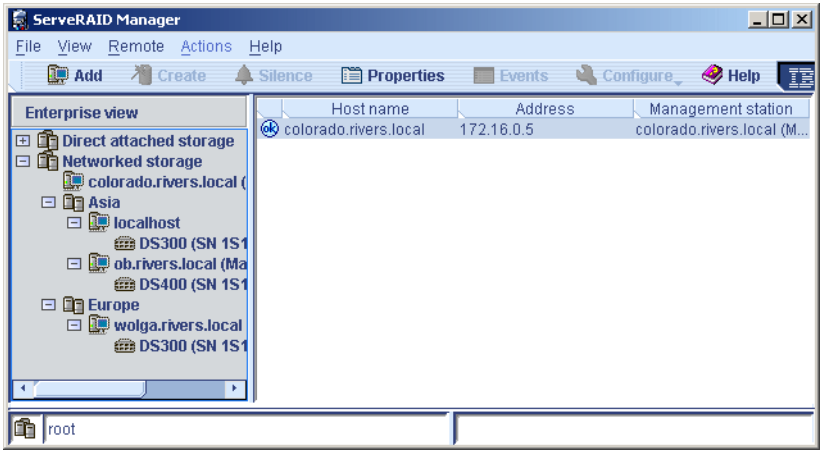

Figure 4-52 Renamed display group

Delete a display group

Display groups no longer required can be deleted. The Management Stations which where added to the given display group will appear under **Networked storage** as they where after adding by default.

To delete a not longer used display group perform the following steps:

1. Right-click the display group which is not longer required. Select **Delete display group** from the context menu.

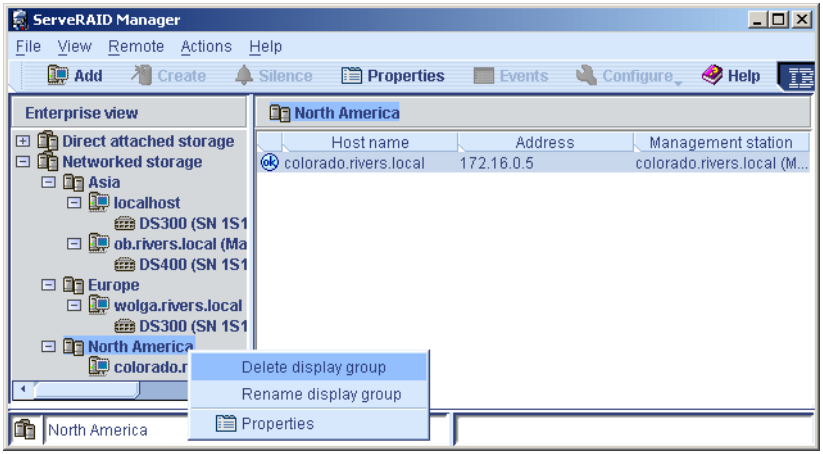

Figure 4-53 Delete a display group

2. The display group is removed and the Management Station is added under **Networked storage**.

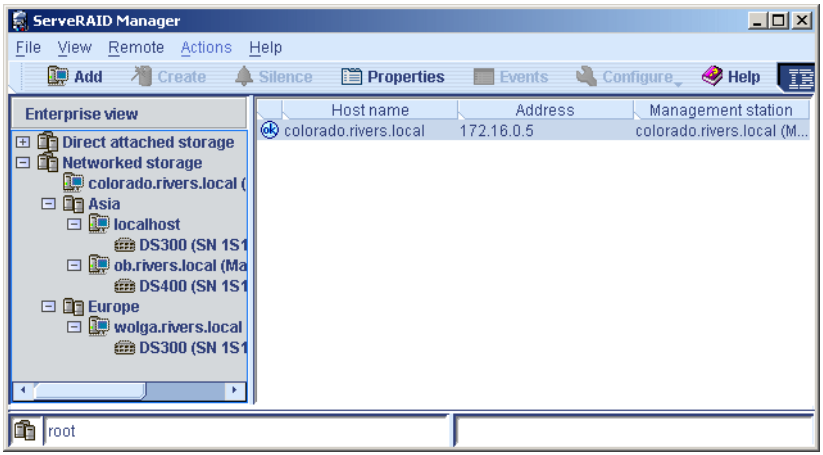

Figure 4-54 Deleted display group

Display group properties

Display group properties gives you a overview about the total health of the Management Stations added to the display group.

1. To see the overview right-click a display group and select **Properties**.

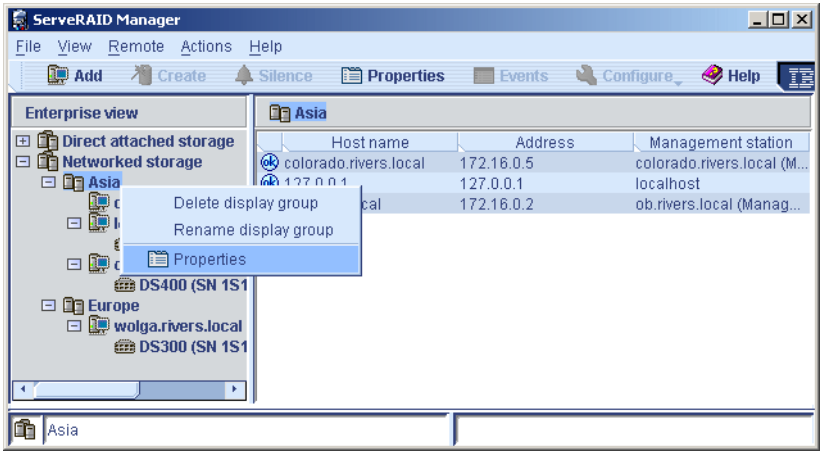

Figure 4-55 Menu display group properties

2. The properties window opens an a overview is shown. Click the close button in the top right corner to close the properties window.

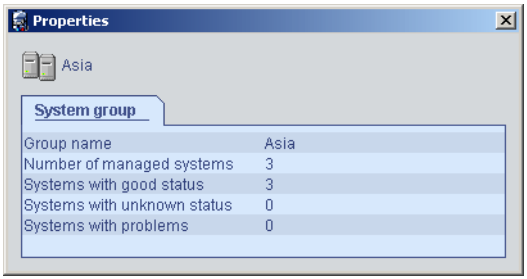

Figure 4-56 Display group properties

4.7.7 Replicate IBM ServeRAID Manager settings to another system

You can configure ServeRAID Manager settings on other servers exactly as they are configured on one server. To replicate the ServeRAID Manager tree, notification list, and security list, do the following:

1. Install the ServeRAID Manager on one server.
2. Start the ServeRAID Manager. Using the 'Add remote system' action, define the servers for your tree.
3. Open the Notification Manager. Using the 'Add system' action, define the notification list.
4. Open the Security Manager. Using the 'Add user' action, define the security list.

5. Exit the ServeRAID Manager.
6. Copy the following files onto a diskette from the directory where the ServeRAID Manager is installed:

RaidMSys.ser	to replicate the tree
RaidNLst.ser	to replicate the notification list
RaidSLst.ser	to replicate the security list
RaidSMTP.ser	to replicate the SMTP e-mail notification list
RaidJob.ser	to replicate the jobs in the Task Scheduler

7. Install the ServeRAID Manager on the other servers.
8. Copy the files from the diskette into the directory where the ServeRAID Manager is installed on the other servers.

4.8 IBM TotalStorage DS300/DS400 array and logical drive management

This section explains how to create arrays and logical drives. Authentication and access control settings are explained as well. The configuration of authentication settings and access control is slightly different between DS300 and DS400. See 4.8.1, “Create arrays, logical drives, and assign IQNs on the TotalStorage DS300” on page 114 about the array and logical drive management for a DS300 and 4.8.2, “Create arrays, logical drives, and access control on the TotalStorage DS400” on page 131. We use a sample configuration which is described later in one of the sample scenarios for the TotalStorage DS300 and DS400.

DS300 and DS400 are able to operate arrays with RAID level 0,1,5,10 and 50. Arrays with a RAID level x0 are able to handle more drives than arrays with a RAID level x. An array with RAID level 0,1 or 5 can have up to 16 drives and an array with RAID level 10 or 50 up to 40 drives. This is archived by using multiple arrays of the base RAID level 1 or 5 and span a RAID 0 across all. The result is called a spanned array.

Note: Only 14 drives can be installed in a TotalStorage DS300.

A TotalStorage DS400 can have only 14 drives. To use more drives, up to two EXP400 expansion units, each with up to 13 drives, can be attached to a TotalStorage DS400. A feature key is required to enable the external SCSI channels to attach the EXP400s.

Do not install drives in bay 6 of the EXP400. This SCSI ID is required for the SCSI chip in the alternate controller.

See 4.8.1, “Create arrays, logical drives, and assign IQNs on the TotalStorage DS300” on page 114 about the creation of arrays and logical drives in a TotalStorage DS300. 4.8.2, “Create arrays, logical drives, and access control on the TotalStorage DS400” on page 131 will explain the same for a TotalStorage DS400.

4.8.1 Create arrays, logical drives, and assign IQNs on the TotalStorage DS300

To explain the creation of arrays, logical drives, access control and authentication settings, a sample configuration is used. This configuration contains the following items:

- First array with two drives as RAID 1 and two logical drives with a size of `10GB`. The logical drives are named `COLORADO-BOOT-W` and `COLORADO-LOG`.
- Second array with 6 drives as RAID 50 with one logical drive. The logical drive size will be `150 GB` and the name is `COLORADO-Mail`.
- Third array with 4 drives as RAID 10 with one logical drive. The logical drive has a size of `50GB` and the name `COLORADO-DB`.
- One HS20 blade with QLogic iSCSI Expansion Card for IBM @server BladeCenter for access to the logical drives. Host name is `COLORADO`.
- Host `COLORADO` has access to the created logical drives only over one path.
- One global hotspare will be defined during creation of arrays.

Follow the steps below to create arrays, logical drives and user accounts and define initiators allowed to access the logical drives:

1. Start IBM ServeRAID Manager (SRM) and connect to a DS300.

Note: Press F5 to get a refreshed view. It is recommended to refresh the view from time to time. The view is not automatically refreshed in many conditions! For example, when the command line interface is used together with the ServeRAID manager console, some changes might be missing until the view is refreshed.

2. Right-click the TotalStorage DS300 enclosure in the enterprise view. Select **Configure enclosure** from the content menu. Or select **Actions** → **Configure enclosure** or click the **Create logical** drive button in the **Logical devices** column. This allows you to manage access control list, authentication and logical drives and arrays. This can be used later as well to perform changes to the existing configuration.

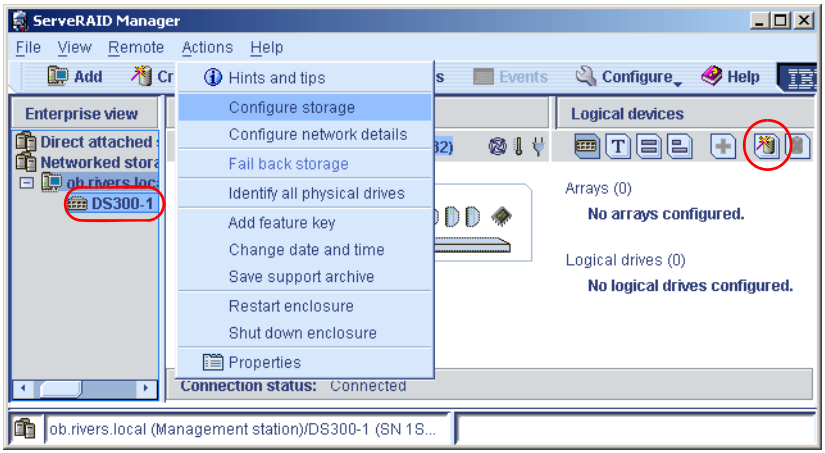

Figure 4-57 SRM with no configured logical drives or arrays

3. The window shown in Figure 4-58 on page 116 permits all access control lists or authentication settings to be managed. Selecting a logical drive and choosing the **Configure access control list** from the context menu allows only the modification of this access control list. The same is true for the authentication settings. This is called **Configure target information**. When the TotalStorage DS300 has no array or logical drive configuration applied, only a new logical drive can be created. All other options will be greyed out. For our example, verify that **Create logical drive** is selected as shown in Figure 4-58 on page 116, and click **Next**.

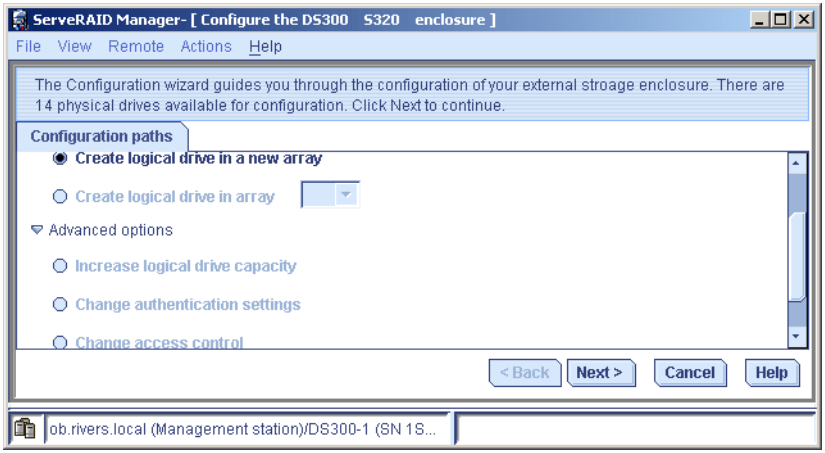

Figure 4-58 Configuration options

4. Hold **Ctrl** and click the two drives which will build the array. Select drives with ID 0 and 1. These two drives will build the first array which will be a RAID 1 array. See Figure 4-59.

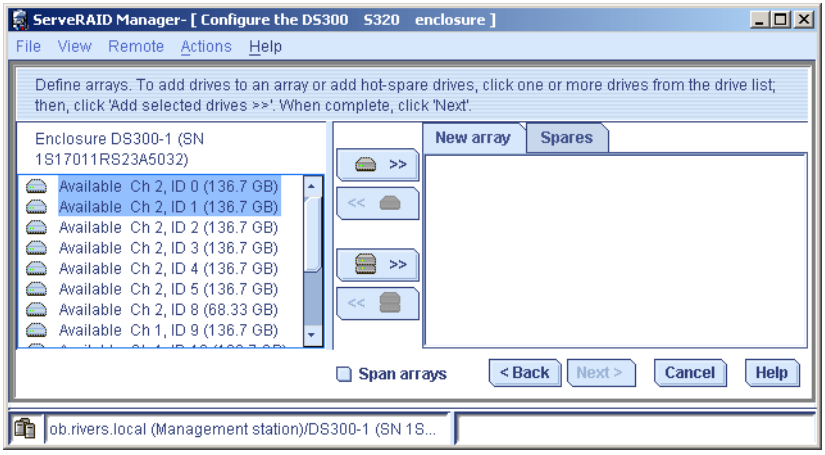

Figure 4-59 Select drives for new Array-1

5. Click the **Add selected drives to array** button. This is the button with the single drive symbol. The selected drives will be assigned to the new array on the right site. The name of **New array** will change to **Array-1** and a new tab with name **New array** is created. Figure 4-60 on page 117.

array create
`Array-1` **RAID1**
`65536 65792`

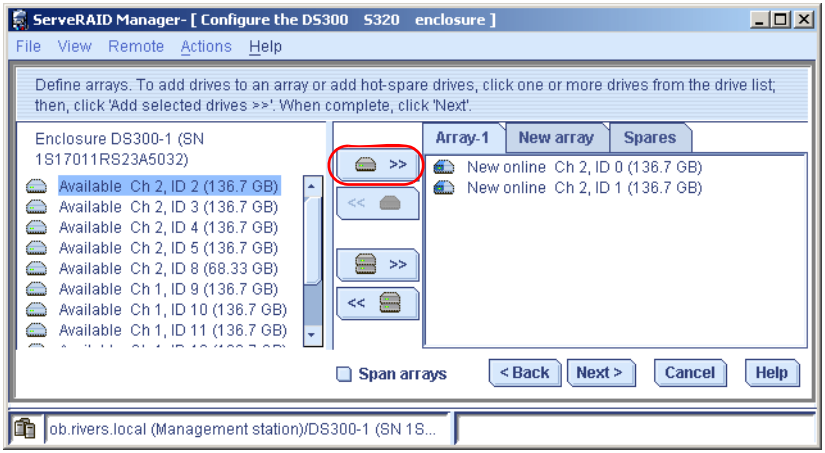

Figure 4-60 Array-1 with two assigned drives

6. Click the newly created tab **New array**. Select drives with the ID `10`, `11` and `12` and assign them to **New array**. This array will be used together with the array created in the next step to build a spanned array with a RAID level 50. Click the **Add selected drives to array** button. See Figure 4-61.

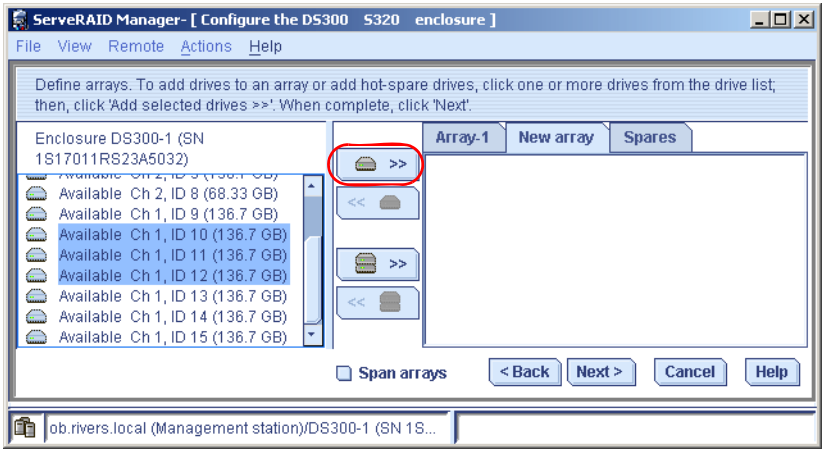

Figure 4-61 Selected drives for the new Array-2

7. The selected drives will be added to the new array and the array is named **Array-2**. A new tab named **New array** is created. Figure 4-62 on page 118.

array create
S-array-1
RAID50 128 2560
2816 3072 3328
3584 3840

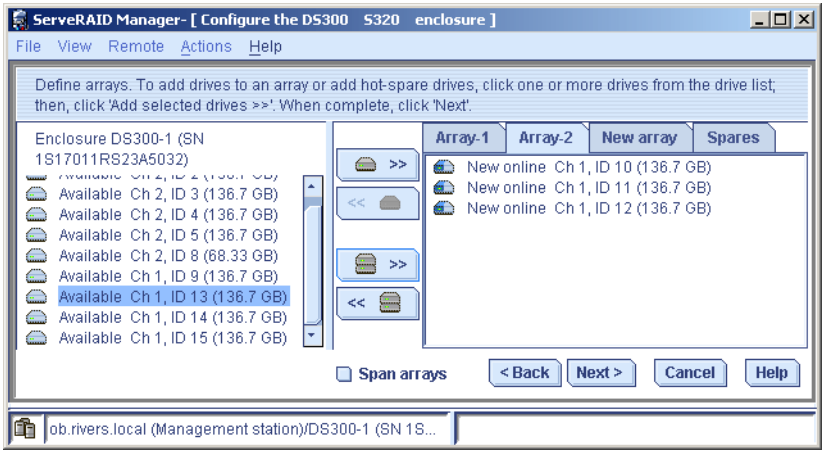

Figure 4-62 Array-2 with three drives

8. Select the **New array** tab and repeat step 6 on page 117 to create a new array with drives ID `13`, `14 and 15`. See Figure 4-63.

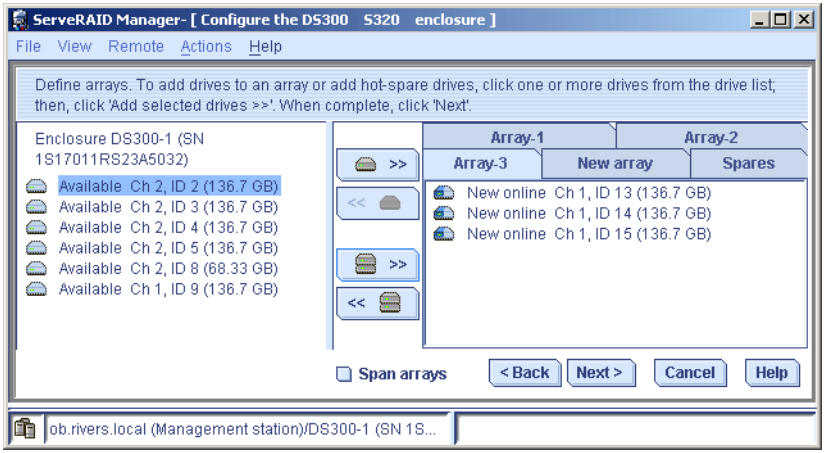

Figure 4-63 Array-3 with three drives

9. Create two more arrays with only two drives in each. Use the drives with ID `2` and `3` in `Array-4` and drives with ID `4` and `5` in `Array-5`. This arrays will be used together to build a spanned array with a RAID level 10. Refer to the previous steps on how to add drives. The final results are shown in Figure 4-64 on page 119.

array create
S-array-2
RAID10 256
66048 66304
66560 66816

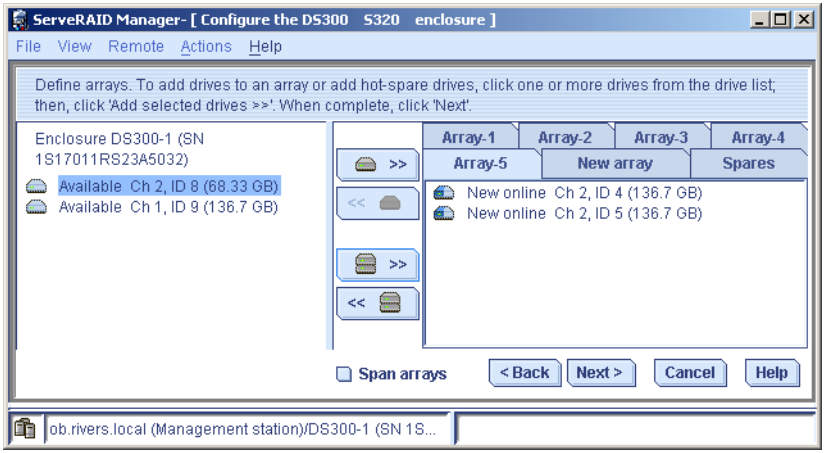

Figure 4-64 Array-5 with two drives

10. Click the **Spares** tab to create a global hotspare. See Figure 4-65.

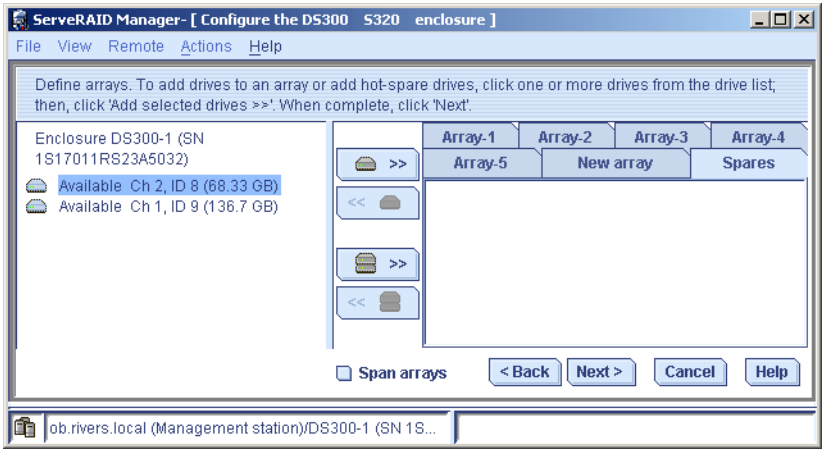

Figure 4-65 Spares tab without hotspare drives

11. Select the drive with ID 9 from the enclosure and define this drive as hotspare drive by clicking the **Add selected drives to spare** button. See Figure 4-66 on page 120.

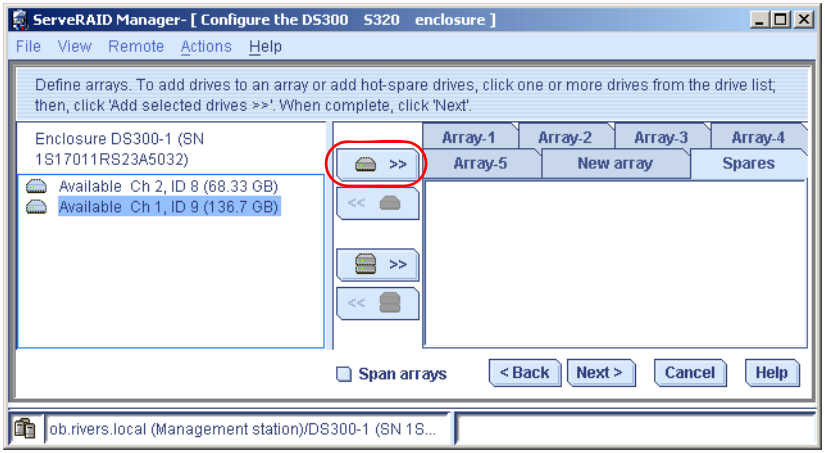

Figure 4-66 Spares tab

12. The selected drive will be moved to the **Spares** tab. This will be a global hotspare drive. See Figure 4-67.

> **Note:** Select a global hotspare drive that has the same capacity as the largest drives installed in the controller. Not all arrays and logical drives will be protected when drives with a smaller size are defined as hotspare. When small and large capacity drives are defined as global hotspares, the largest one is always selected first to rebuild an array with a disabled drive.

array spare add
2304

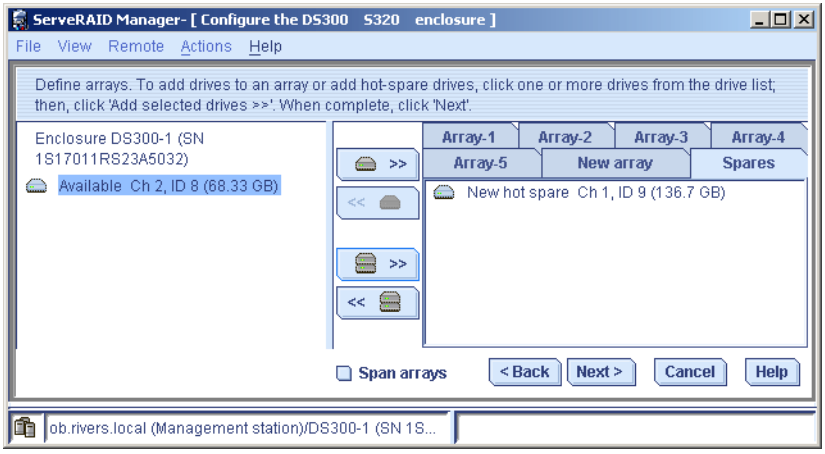

Figure 4-67 Spares tab with spare drive

13. Click the **Span arrays** check box shown in Figure 4-68 to create spanned arrays. Without checking this checkbox, spanned arrays cannot be created. Click **Next**.

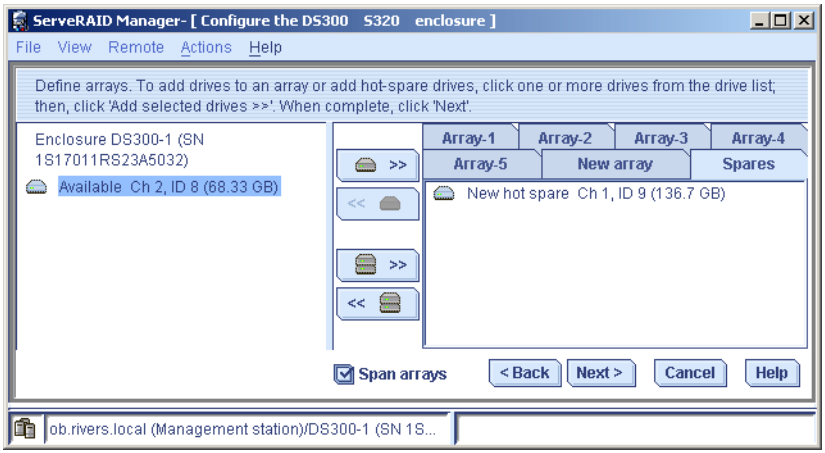

Figure 4-68 Checked Span arrays

14. The window shown in Figure 4-69 is used to create spanned arrays. Select the base arrays which will be used to build a spanned array. Select **Array-2** and **Array-3** and assign them to the **New spanned array**. The **New spanned array** tab will be renamed to **Spanned-array-1** and a new **New spanned array** tab is created. Create another spanned array with Array-4 and Array-5. See Figure 4-69. Click **Next**.

array create
S-array-1
RAID50 128 2560
2816 3072 3328
3584 3840

array create
S-array-2
RAID10 256
66048 66304
66560 66816

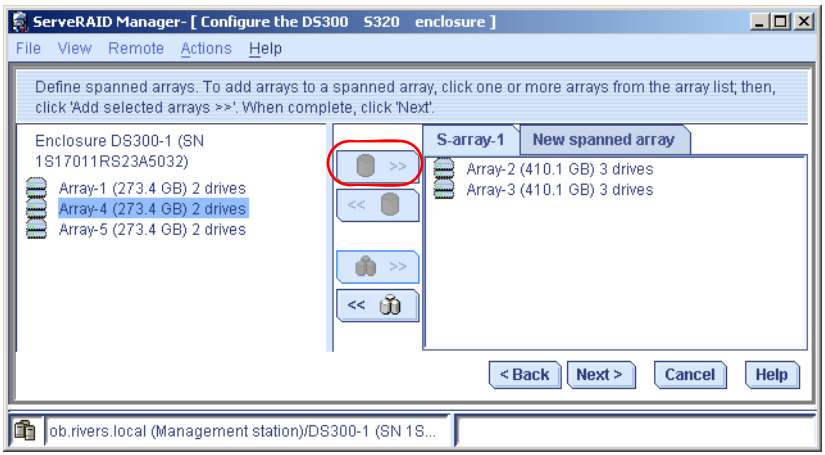

Figure 4-69 S-array-1 with base arrays

15. The window shown in Figure 4-70 allows you to create logical drives in all defined arrays. The first logical drive in each array is already defined with the maximum logical drive size. Enter a meaningful logical drive name. This logical drive name will be part of the IQN of the device and is visible on the iSCSI initiator. Select the RAID level. This RAID level is determined by the number of drives in the array. For a spanned array, the number of drives in the base array determines if a RAID 10 or 50 will be defined. Specify how large the logical drive will be. Choose the preferred controller.

    For our example, use a logical drive name `COLORADO-BOOT-W`. Specify RAID level `1` and a size of `10GB` for this logical drive. The logical drive and this array will be owned by controller A.

logical create
`COLORADO-BOOT-W`
`Array-1 10G`
restricted

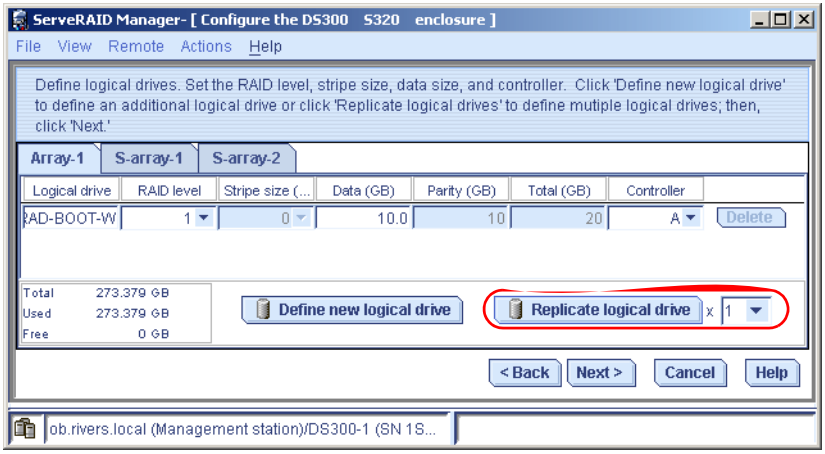

Figure 4-70 Logical drive COLORADO-BOOT-W

16. In Figure 4-70, the **Define new logical drive button** creates a new logical drive. The size will be the same as the remaining free space.

    Select the logical drive COLORADO-BOOT-W and click the **Replicate logical drive** button. A new logical drive with the same size is created as shown in Figure 4-71 on page 123. The dropdown box with the number 1 specifies how often the logical drive is replicated. Up to 63 logical drives can be replicated, depending on the remaining free space after creating each replicated logical drive.

    Assign the new replicated logical drive a name `COLORADO-LOG`.

logical create
COLORADO-LOG
Array-1 10G
restricted

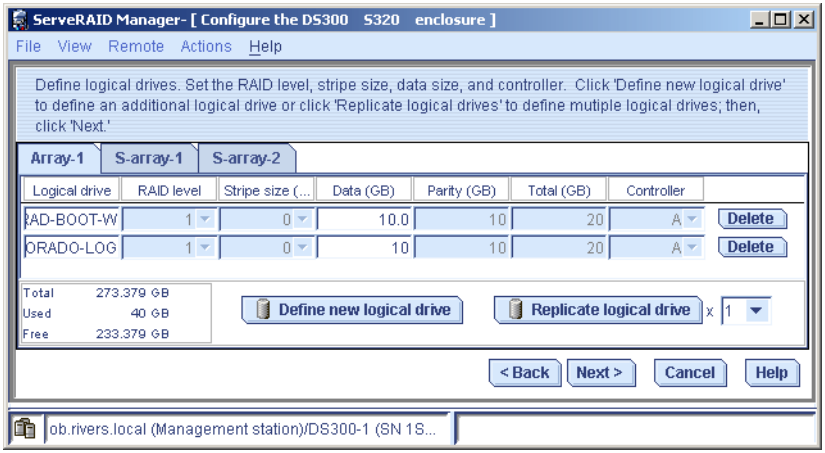

Figure 4-71 Array-1 with configured logical drives

17. As shown in Figure 4-72, select **S-array-1**. Assign a meaningful logical drive name. The RAID level is determined automatically by the used base arrays. For logical drives in RAID 5, RAID 10 and RAID 50 arrays the stripe size can be modified. Different stripe sizes are available. 256 is the default value. Smallest stripe size is 16K and largest stripe size is 1MB. Define the size of the logical drive and the preferred controller.

    Enter the logical drive name COLORADO-MAIL, specify a strip size of 128K and a size of 150GB. Change the owner of the logical drive to controller B.

logical create
COLORADO-MAIL
S-Array-1 150G
restricted

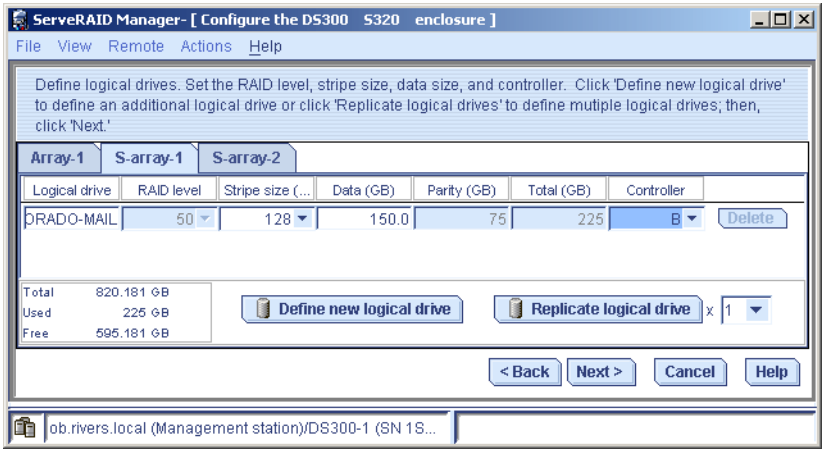

Figure 4-72 Logical drive COLORADO-MAIL

18. Select **S-array-2** as shown in Figure 4-73 on page 124. Assign a meaningful logical drive name. The RAID level determined automatically by the used

base arrays. For logical drives in RAID 5, RAID 10 and RAID 50 arrays the stripe size can be modified. different strip sizes are available. 256 is the default value. Smallest stripe size is 16K and largest stripe size is 1MB. Define the size of the logical drive and the preferred controller.

Enter the logical drive name COLORADO-DB and a size of 50GB. Change the owner of the logical drive to controller B. Click **Next.**

logical create
COLORADO-DB
S-array-2 50**G**
restricted

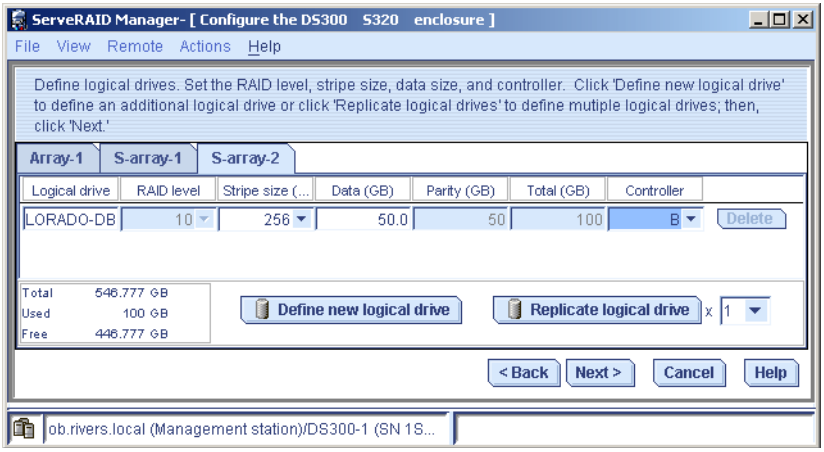

Figure 4-73 Logical drive COLORADO-DB

19. Select the logical drive COLORADO-DB from the enclosure. Change the **Authentication type** from **None** to **CHAP** for this logical drive. See Figure 4-74 on page 125.

Note: If a CHAP name and secret is specified during the configuration of the iSCSI host bus adapter in the host adapter settings, then CHAP needs to be selected for the boot logical drive and a user with the same name and the secret must be specified. See Figure 7-15 on page 236.

logical manage
COLORADO-DB
authentication
CHAP

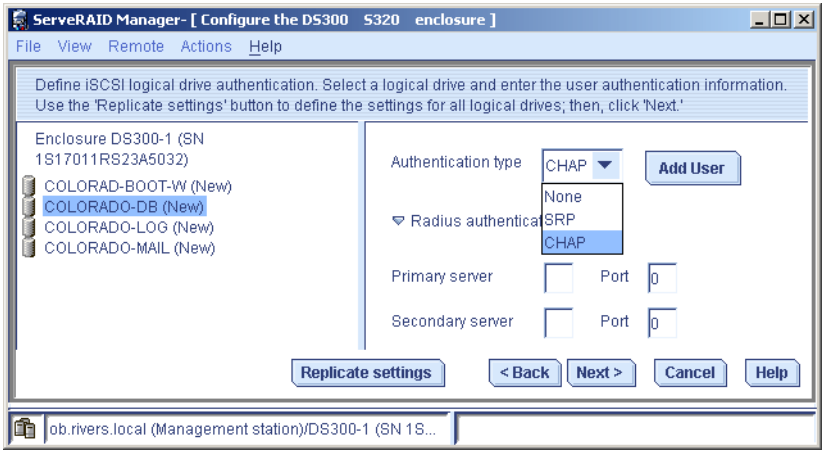

Figure 4-74 Authentication type of logical drive COLORADO-DB

20. Click the **Add user** button. Specify as **User name** the IQN of the HBA Port used to access logical drives from this TotalStorage DS300. Enter a password with at least 12 characters and confirm it in the **Confirm Password** field. The IQN `iqn.2000-04.com.qlogic:qmc4052.zj1ve454f01y.2` of the second port was used as user name and the password `a1234567890b`. See Figure 4-75.

    Click **Add** to add the new user account.

authentication add
`iqn.2000-04.com.qlogic:qmc4052.zj1ve454f01y.2`

Enter the password twice

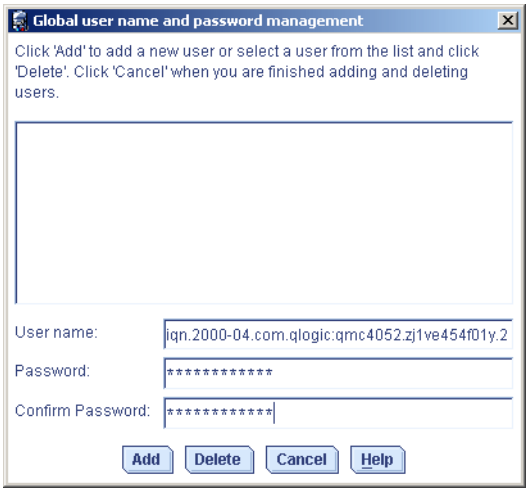

Figure 4-75 Global user and password management

21. Click **Cancel** after you have created all required accounts. See Figure 4-76 on page 126.

Figure 4-76 Global user and password management

22. Click **Next** in Figure 4-77.

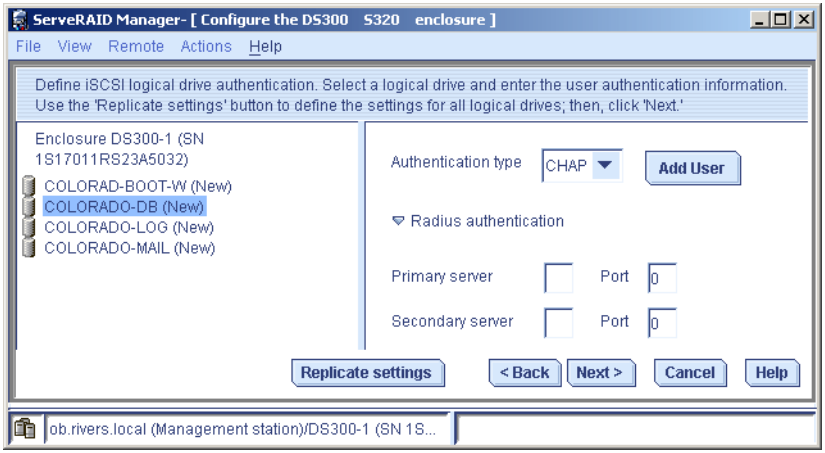

Figure 4-77 Authentication type of logical drive COLORADO-DB

23. Define the IQNs of the iSCSI initiators which will access the DS300. An alias is used to assign more meaningful names to the IQNs. Click **Add** to create a IQN. See Figure 4-78 on page 127.

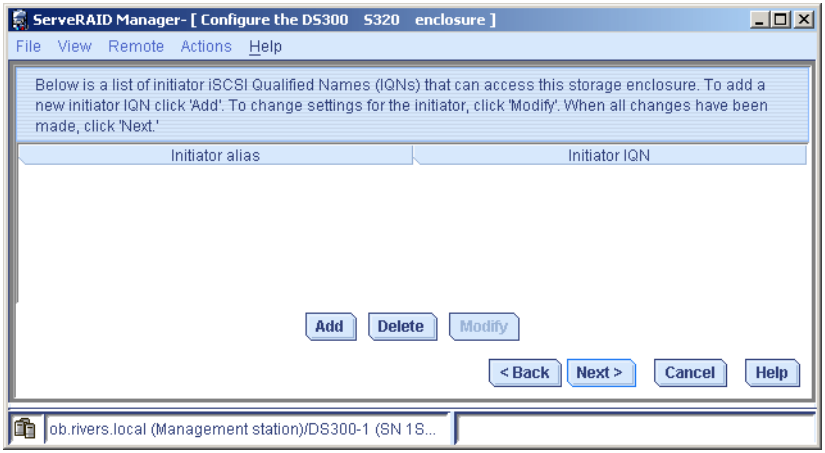

Figure 4-78 iSCSI initiator management

24. Add the IQN of an iSCSI initiator into the Initiator IQN field. Assign a IQN alias. Click **OK** to add the IQN to the configuration. Add all required IQN's here.

> **Note:** iSCSI initiators are not detectable like world wide names. They must be specified manually.

Define the IQN `iqn.2000-04.com.qlogic:qmc4052.zj1ve454f01y.1` of the first adapter port. The alias is `COLORADO-P0`. See Figure 4-79. Click **OK**.

initiator create
`COLORADO-P0`
`iqn.2000-04.com`
`.qlogic:qmc4052`
`.zj1ve454f01y.1`

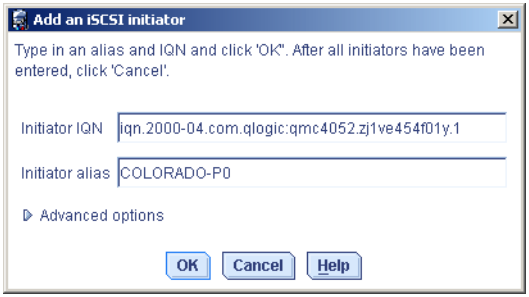

Figure 4-79 IQN of the first port of the iSCSI HBA

25. The IQN of the second port is `iqn.2000-04.com.qlogic:qmc4052.zj1ve454f01y.2` and the alias is `COLORADO-P1`. Enter these values and click **OK**. See Figure 4-80 on page 128.

initiator create
COLORADO-P1
iqn.2000-04.com
.qlogic:qmc4052
.zj1ve454f01y.2

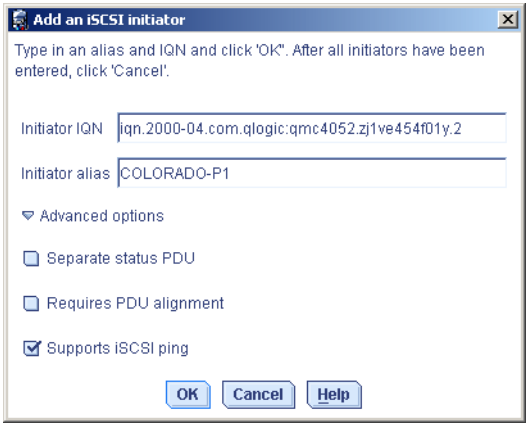

Figure 4-80 IQN of the second port of the iSCSI HBA

26. A new and empty **Add iSCSI initiator** dialog opens. Click **Cancel** to leave this dialog.

27. Click **Next** as shown in Figure 4-81.

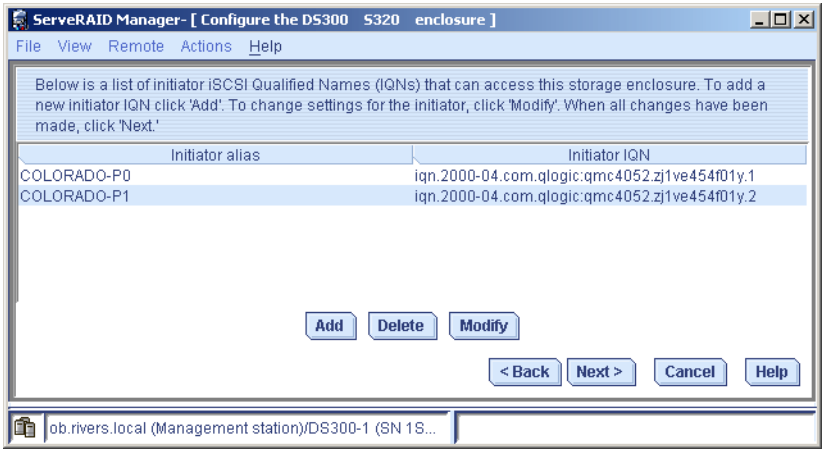

Figure 4-81 Defined iSCSI initiators and aliases

28. Assign the logical drive to the alias of an IQN. This will restrict the access to the selected logical drive to an iSCSI initiator with this IQN. When CHAP authentication is specified for this logical drive, a username/password pair is required together with the correct IQN to access this logical drive. A logical drive assigned to the Unrestricted tab can be accessed by any iSCSI initiator. When a logical drive is assigned to more than one IQN, a warning in red is shown explaining to how many IQNs the logical drive is assigned.

Select the logical drives COLORADO-BOOT and COLORADO-LOG and assign them to the initiator alias COLORADO-P0. The other port is assigned logical drives later. See Figure 4-82.

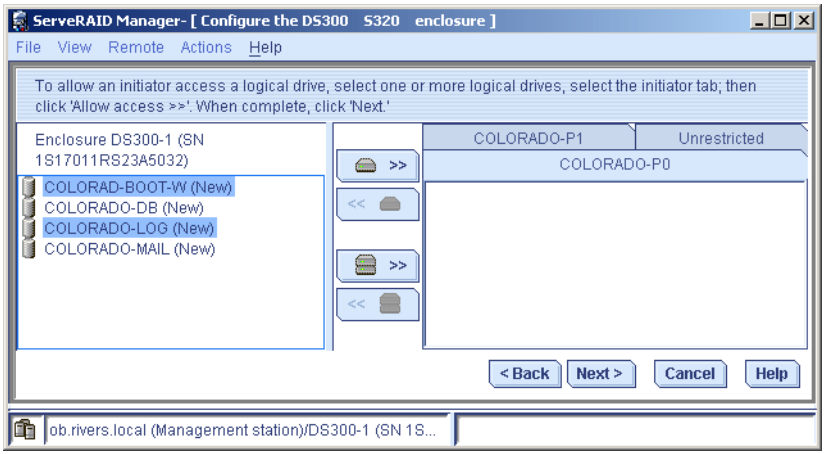

Figure 4-82 Assign logical drives to iSCSI initiator aliases

29. The assigned logical drives will be removed from the enclosure list and added under the initiator alias. When a logical drive is assigned more than once to an initiator, they are shown in red with the number of assigned initiators. Click **Next**. See Figure 4-83.

```
logical manage
COLORADO-BOOT-W
allow
COLORADO-P0

logical manage
COLORADO-LOG
allow
COLORADO-P0
```

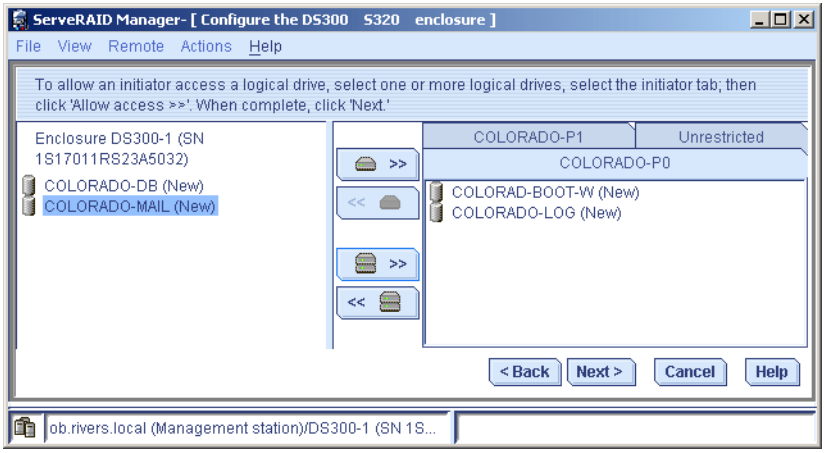

Figure 4-83 Assigned logical drives to an iSCSI initiator alias

30. Verify setup and click **Apply**. If modifications or corrections are required, use the **Back** button and correct the configuration. All logical drives are protected

by at least one spare drive. Two logical drives will not be accessible because no initiator was assigned. See Figure 4-84.

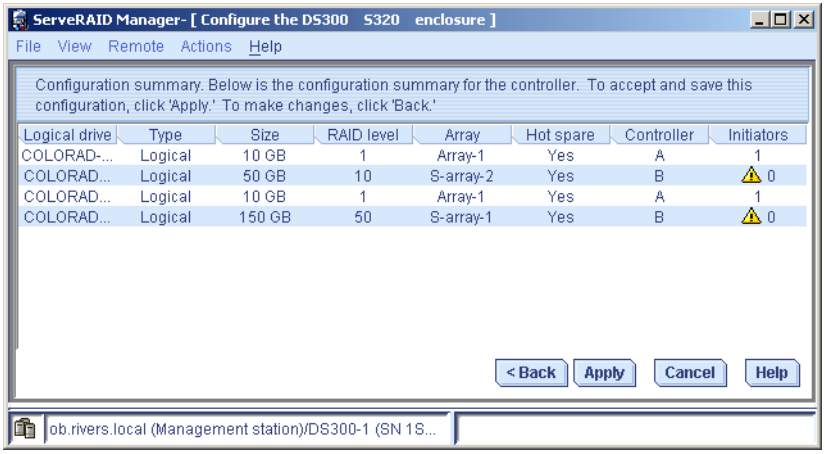

Figure 4-84 Configuration summary

31. Confirm the warning message to apply the configuration by clicking **Yes**.

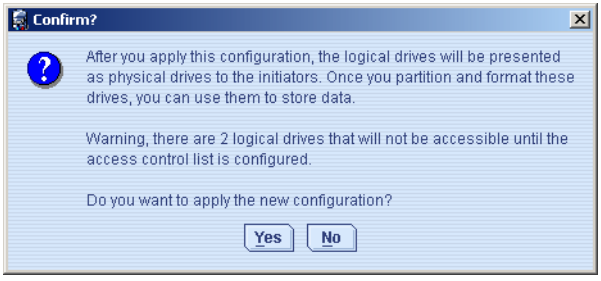

Figure 4-85 Confirm configuration

It will take some time to apply the configuration. The more complex the setup is, the more time is required to apply the configuration. See Figure 4-85.

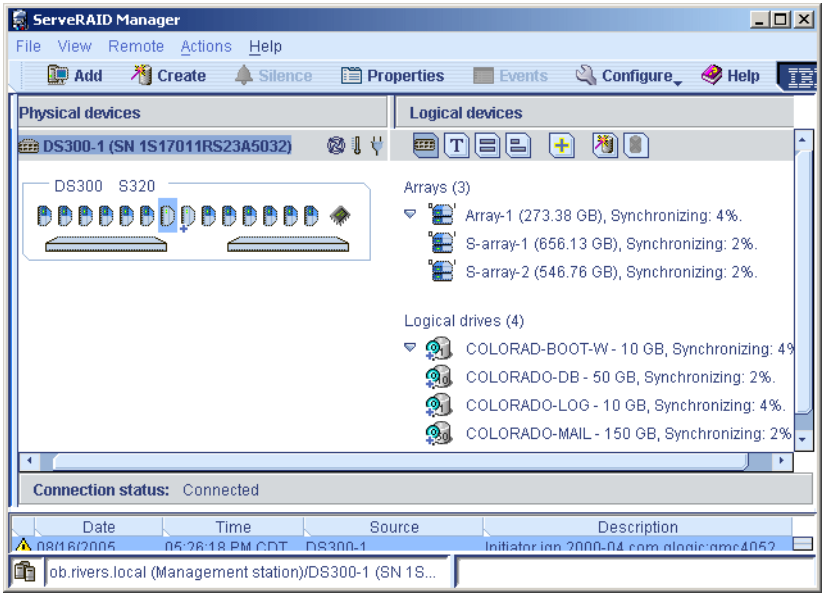

Figure 4-86 Applied TotalStorage DS300 configuration

4.8.2 Create arrays, logical drives, and access control on the TotalStorage DS400

To explain the creation of arrays, logical drives, access control and authentication settings, a sample configuration is used. This configuration contains the following items:

- First array with two drives as RAID 1 and one logical drive with a size of `10GB`. The logical drive will be named `WOLGA-BOOT`.
- Second array with six drives as RAID 50 with one logical drive. Size and name of the logical drive will be `100 GB` and `WOLGA-Media`.
- Third array with four drives as RAID 10 with four logical drives. The logical drives are a `15GB` large logical drive `WOLGA-DB`, a `11GB` logical drive `WOLGA-LOG` and two logical drives `Device-1` and `2` with a size of `11GB`.
- One host `WOLGA` with two fibre channel HBA's get access to the created logical drives. over redundant paths.
- Define one global hotspare during creation of arrays.

Follow the steps below to create arrays, logical drives and define initiators allowed to access the logical drives:

1. Start IBM ServeRAID Manager (SRM) and connect to a DS400.

Note: Press **F5** to get a refreshed view. It is recommended to refresh the view from time to time. The view is not automatically refreshed in many conditions!

2. Right-click the TotalStorage DS400 enclosure in the enterprise view. Select **Configure storage** from the content menu or click the **Create logical** drive button in the **Logical devices** column. This function is also accessible from menu **Action** → **Configure storage**.

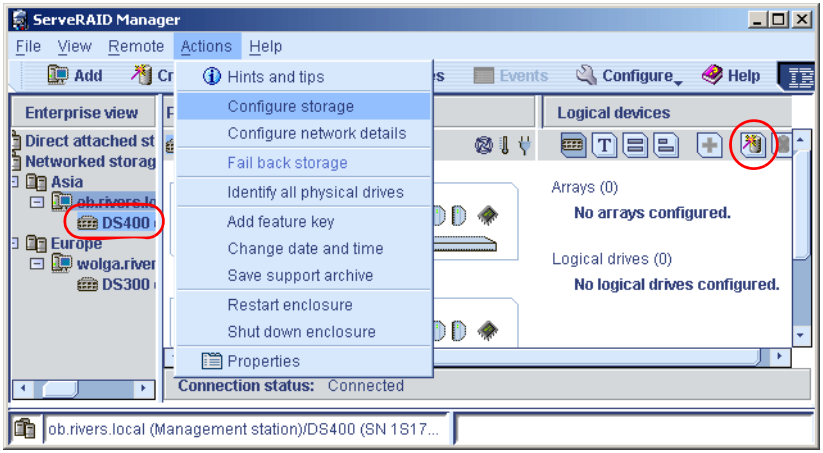

Figure 4-87 SRM with no configured logical drives or arrays

3. Verify that **Create logical drive** is selected as shown in Figure 4-88 on page 133 and click **Next**.

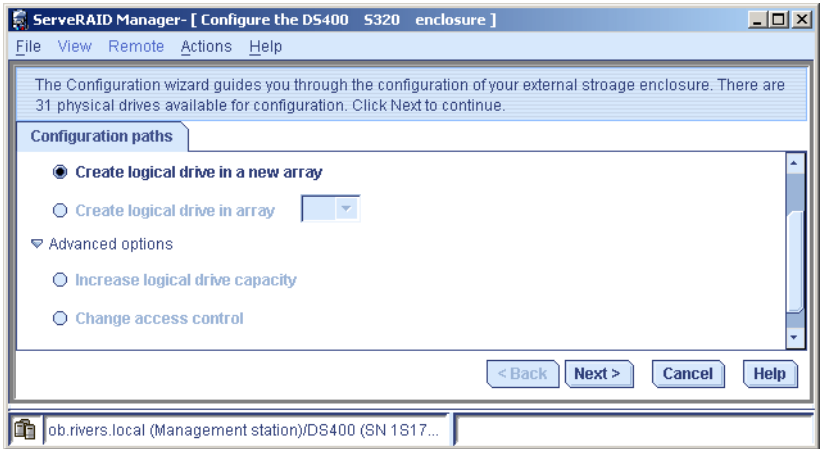

Figure 4-88 Configuration options

4. Hold **Ctrl** and click the two drives which will build the array (the first two drives selected to build the first array). Then click the **Add selected drives to array** button. See Figure 4-89.

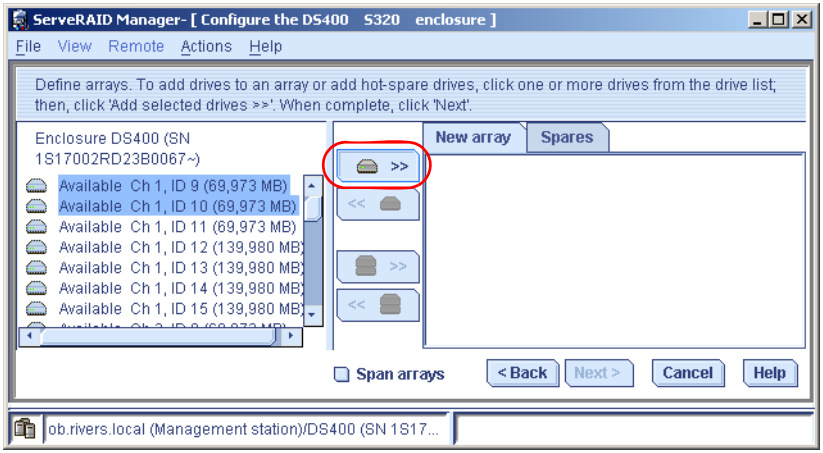

Figure 4-89 Select drives for the new array

5. The selected drives will move from the left panel to the right panel. The tab named **New array** will change to **Array-1** and a new tab with name **New array** is created. See Figure 4-90 on page 134.

array create
`Array-1` **RAID1**
`2304 2560`

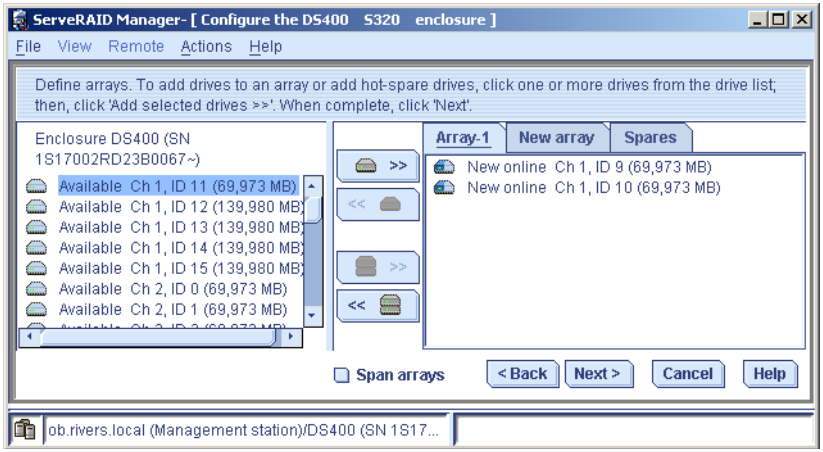

Figure 4-90 Array-1 with two drives

6. Click the newly created **New array** tab. Select 3 drives and assign them to the **New array**. This array will be used together with the array created in the next step to build a spanned array with a RAID level of 50. Then click the **Add selected drives to array** button. See Figure 4-91.

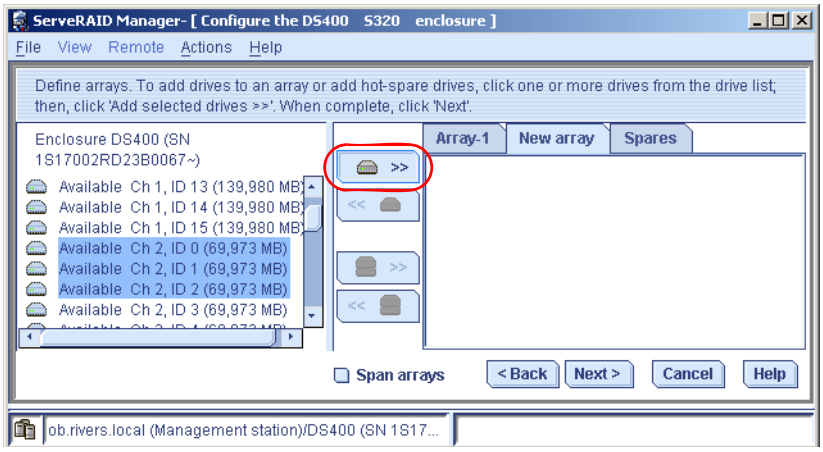

Figure 4-91 Selected drives for the new array

7. The selected drives will be added to the new array and the array is named **Array-2**. A new tab named **New array** is created. See Figure 4-92 on page 135.

array create
S-array-1
RAID50 128
65536 65792
66048 66304
66560 66816

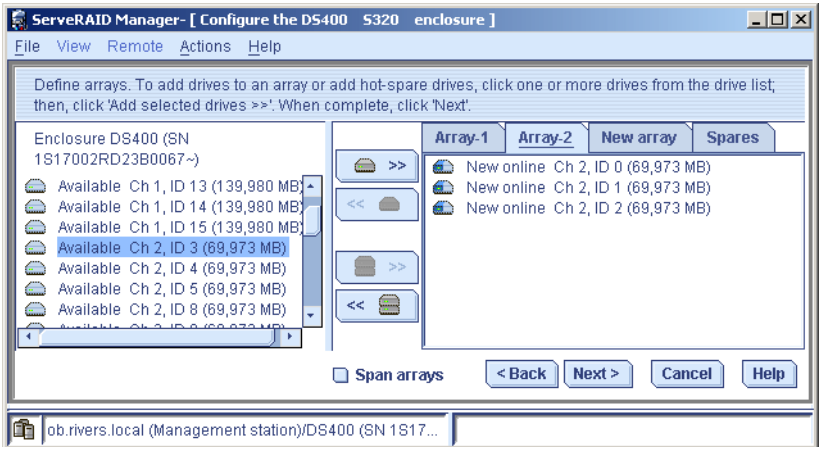

Figure 4-92 Array-3 with three drives

8. Select the **New array** tab and repeat step 6 on page 134 to add three drives with the IDs 3, 4 and 5 into a new array.
9. Create two more arrays with only two drives in each. Use drives with ID 12 and 13 for the first array and ID 14 and 15 for the second array. Refer to the previous steps on how to add drives. The final results are displayed in Figure 4-93.

array create
S-army-2
RAID10 256 3072
3328 3584 3840

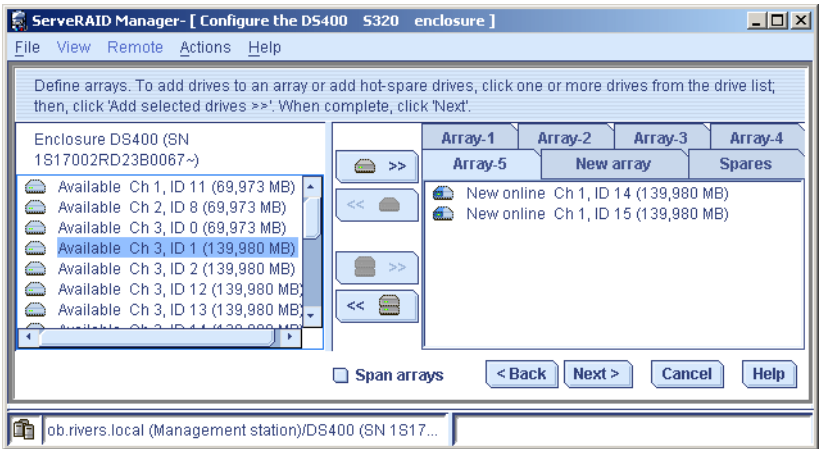

Figure 4-93 Array-5 with two drives

10. Click the **Spares** tab to create global hotspare drives. See Figure 4-94 on page 136.

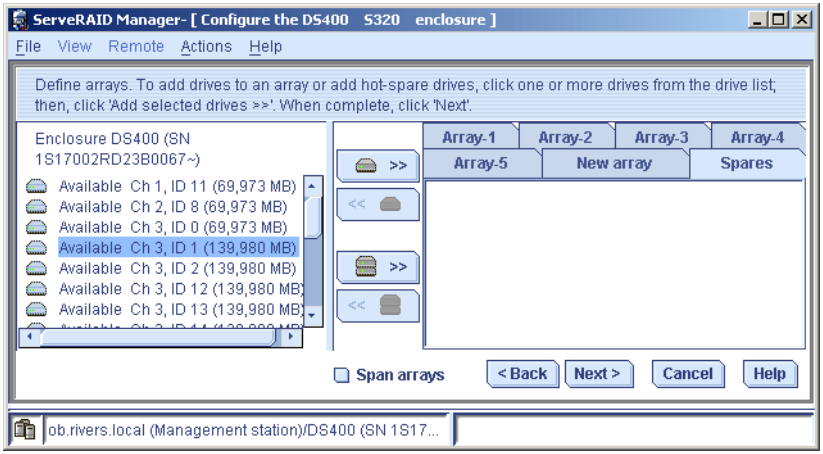

Figure 4-94 Spares tab

11. Select the physical drives which will be the hotspares. Then click the **Add selected drives to array** button. See Figure 4-95.

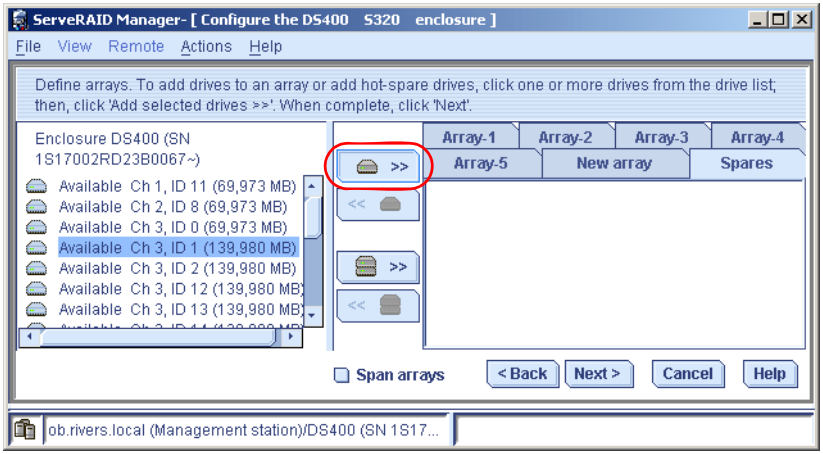

Figure 4-95 Spares tab without spare drive

12. The selected drives will be moved to the **Spares** tab. This will be the global hotspare drive. See Figure 4-96 on page 137.

Note: Select a global hotspare drive that has the same capacity as the largest drives installed in the controller. Not all arrays and logical drives will be protected when drives with a smaller size are defined as hotspare. When small and large capacity drives are defined as global hotspares, the largest one is always selected first to rebuild an array with a disabled drive.

array spare add
196864

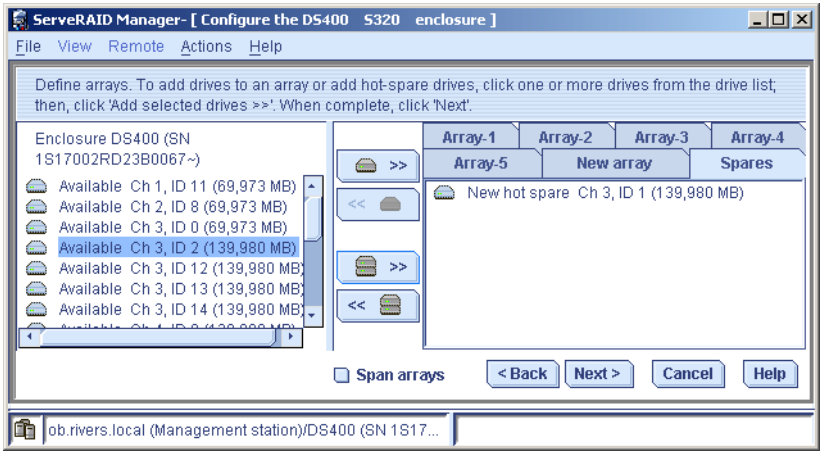

Figure 4-96 Spares tab with spare drive

13. Click the **Span arrays** check box shown in Figure 4-97 to create spanned arrays. Without checking this checkbox, spanned arrays cannot be created. Click **Next**.

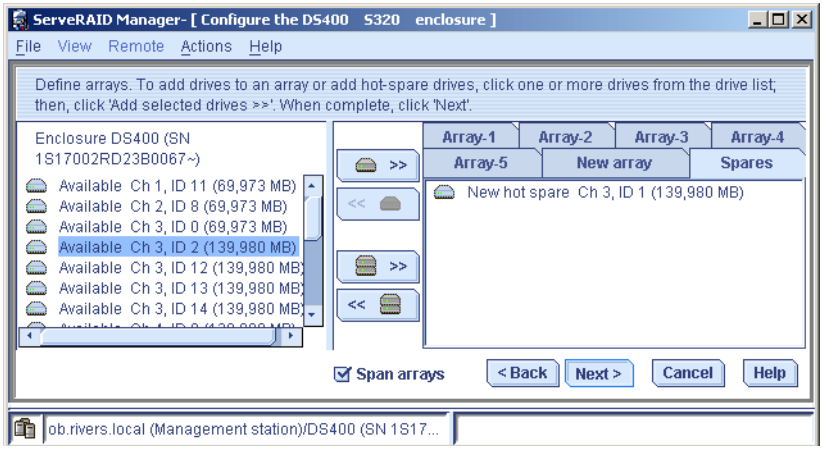

Figure 4-97 Checked Span arrays

14. The window shown in Figure 4-98 is used to create spanned arrays. Select the base arrays which will be used to build a spanned array. Select **Array-2** and **Array-3** and assign them to the **New spanned array**. The **New spanned array** tab will be renamed to **S-array-1** and a new **New spanned array** tab is created. Create another spanned array with Array-4 and Array-5. See Figure 4-98. Click **Next**.

array create
`Array-1` **RAID1**
`2304 2560`

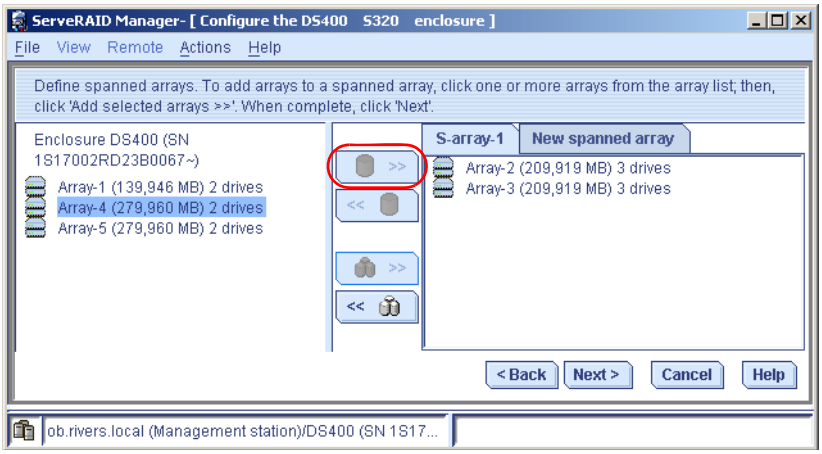

Figure 4-98 Spanned arrays with drives

15. The window shown in Figure 4-99 on page 139 allows you to create logical drives in all defined arrays. The first logical drive in each array is already defined with the maximum logical drive size. Enter a meaningful device name. Select the RAID level. This RAID level is determined by the number of drives in the array. For a spanned array, the number of drives in the base array determines if a RAID 10 or 50 will be defined. Specify how large the logical drive will be. Choose the preferred controller.

    For our example, use a logical drive name `VOLGA-BOOT`. Specify RAID level `1` and a size of `10GB` for this logical drive. The logical drive and this array will be owned by controller A.

logical create
WOLGA-BOOT
Array-1 10240**M**

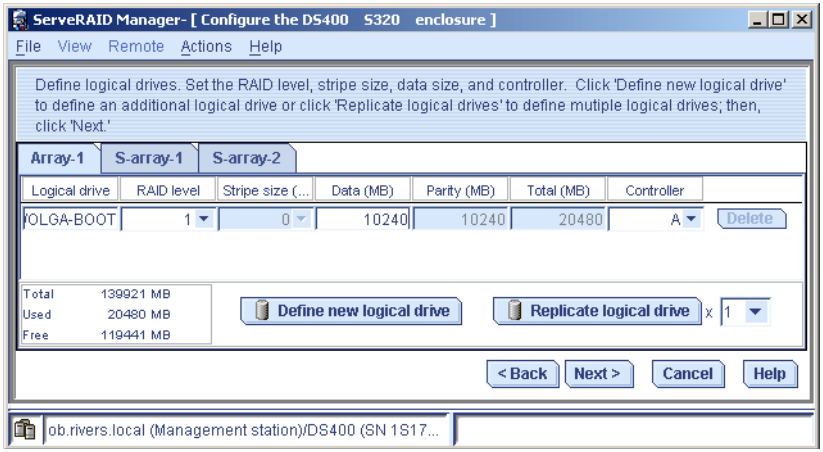

Figure 4-99 Create logical drives

16. In Figure 4-99, select **S-array-1**. Assign a meaningful device name for S-array-1. Figure 4-100 shows a logical drive named `WOLGA-Media` with a size of `100GB`. The RAID level is automatically determined by the base arrays. Base arrays with more than 2 drives become a RAID 50 and arrays with at least and not more than two drives become a RAID 10. For logical drives in RAID 10 and RAID 50 arrays, the stripe size can be modified. Different strip sizes are available. 256KB is the default value. Define the size of the logical drive and the preferred controller. Select **S-array-2.**

logical create
WOLGA-MEDIA
S-array-1
102400**M**

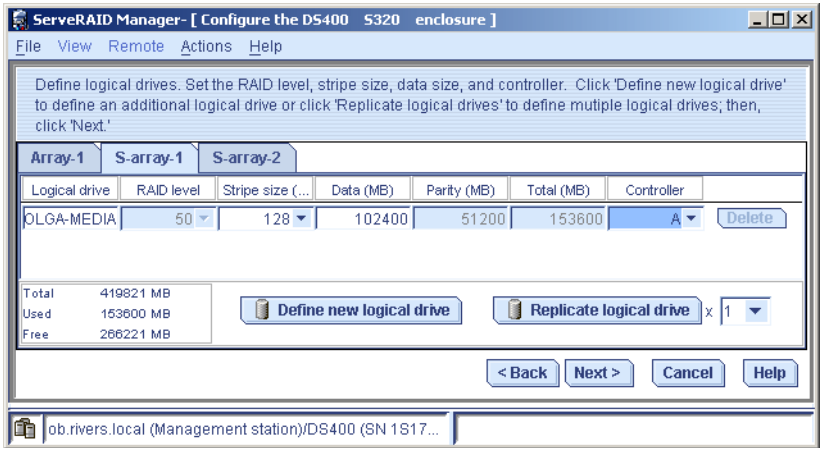

Figure 4-100 RAID level

17. For S-array-2, define the name of the logical drive, stripe size, size and controller. The window shown in Figure 4-101 on page 140 shows the first

logical drive `WOLGA-DB` with a size of `15GB` owned by controller `B`. Click the **Define new logical drive** button to create the second logical drive in **S-array-2**.

logical create
`WOLGA-DB`
`S-array-2`
`15360`**M**

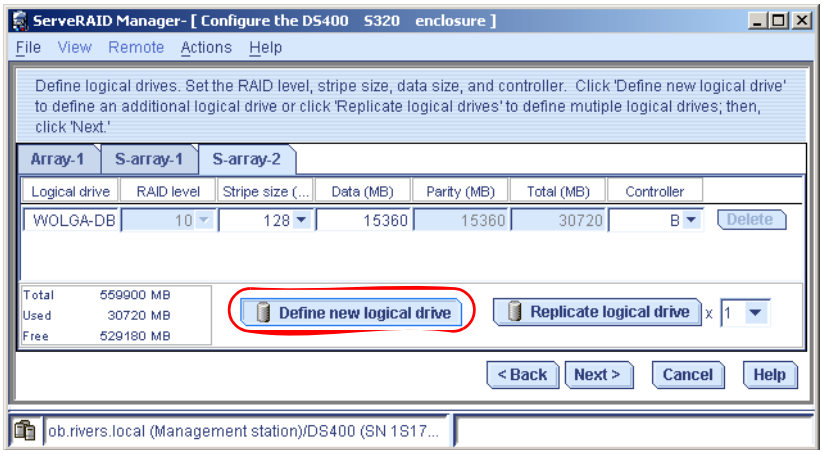

Figure 4-101 First logical drive for S-array-2

18. A new logical drive with name Device-1 and the maximum available space is created. Change the logical drive name to `WOLGA-LOG` and specify a size of `11GB`. See Figure 4-102.

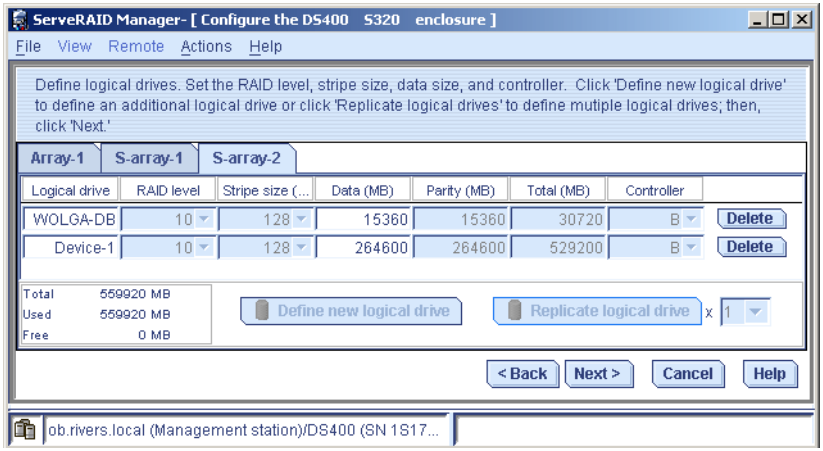

Figure 4-102 New defined logical drive

Note: Stripe size and controller are array specific values and cannot be set to different values in the same array.

19. When several logical drives are required, and the array has enough free space, up to 63 additional logical drives can be created. To replicate a logical drive, place the cursor on the line of the logical drive that is to be replicated. Select the number of replicas, and click the button **Replicate logical drive.** See Figure 4-103.

logical create
WOLGA-LOG
S-array-2
11264**M**

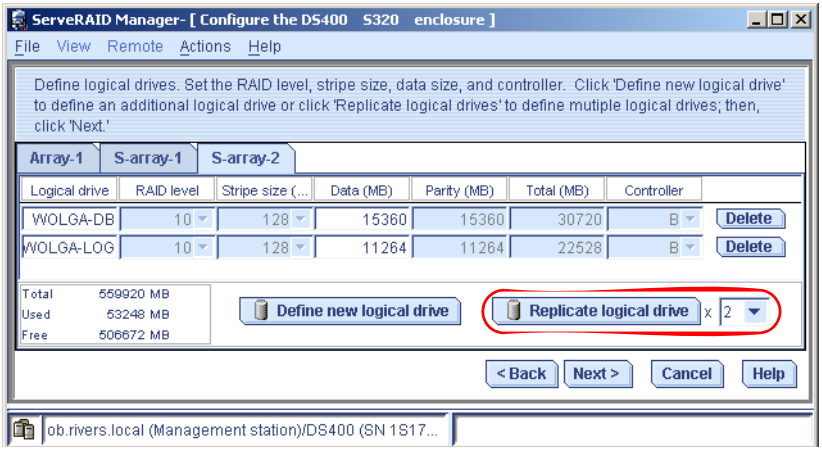

Figure 4-103 Replicate a logical drive

20. The requested number of logical drives will be created with default name of Device-x. See Figure 4-104. Click **Next** after all required logical drives are defined.

logical create
Device-1
S-array-2
11264**M**

logical create
Device-1
S-array-2
11264**M**

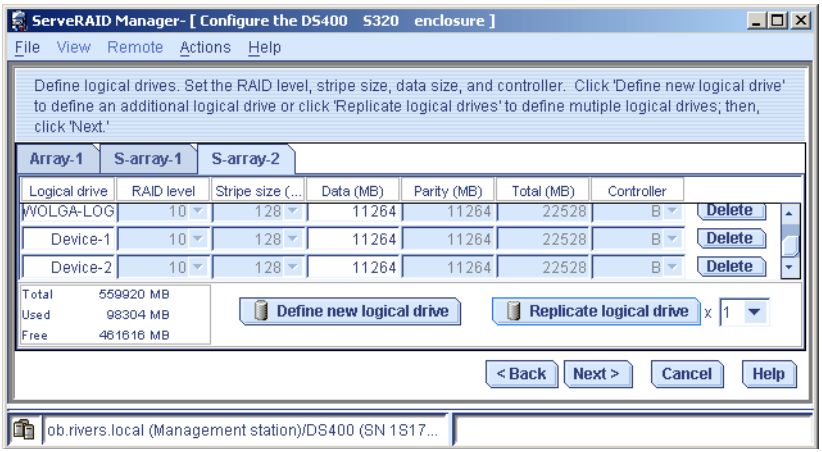

Figure 4-104 New replicated logical drives

21. The dialog shown in Figure 4-105 on page 142 is used to discover new initiators, add or manage existing initiators. Click the **Discover new initiators** button.

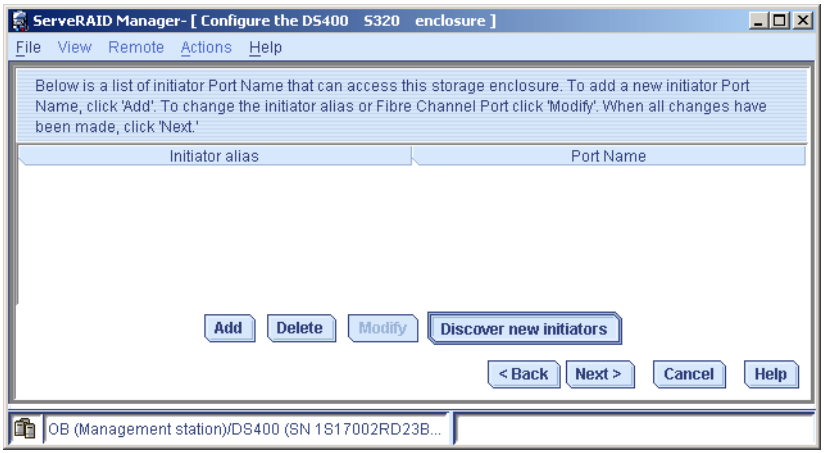

Figure 4-105 Discover new initiators

22. When new initiators are found, a dialog box opens with the discovered initiators. See Figure 4-106. Select a world wide port name (WWPN) from the **Port Name** drop down list. The default initiator alias is discovered-x. Assign a meaningful initiator alias to the initiator and click **OK**.

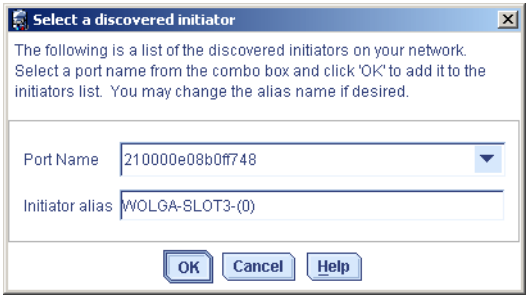

Figure 4-106 Discovered initiator

23. Repeat steps 21 on page 142 and 22 until all initiators are defined. An error message informs when no more initiators are available. See Figure 4-107.

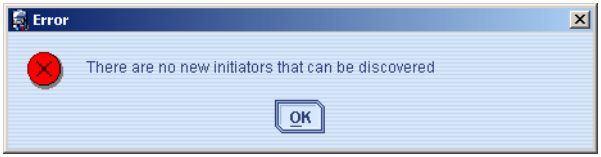

Figure 4-107 No new initiators can be discovered

24. With the **Modify** button, an initiator alias can be modified but not the initiator itself. Select the first listed initiator and click **Modify**. See Figure 4-108.

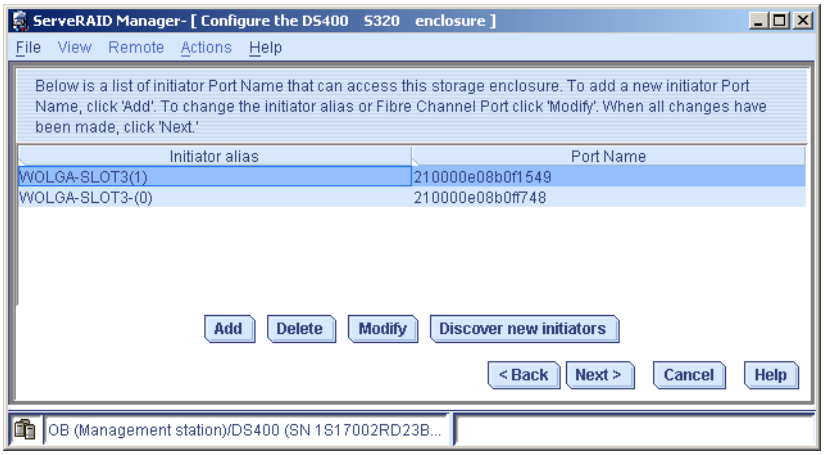

Figure 4-108 List of defined initiators

25. Modify the initiator alias and click **OK**.

> **Note:** The alias can have a length of 16 characters when modified with the ServeRAID Manager console. The CLI supports longer aliases. They will be shown correctly in the ServeRAID Manager console.

```
initiator
manage
WOLGA-SLOT3-(1)
rename
WOLGA-SLOT4-(1)
```

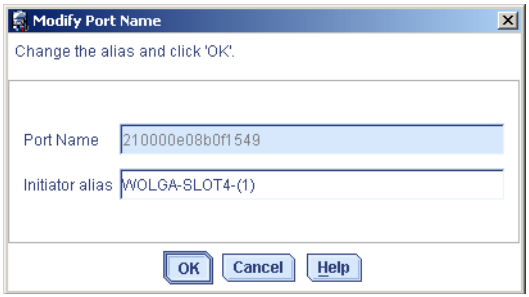

Figure 4-109 Rename the initiator alias

26. In Figure 4-110 on page 144, the **Add** button permits definition of initiators manually. Initiators can be deleted with **Delete**. Click **Next**.

> **Note:** When an initiator is added manually, the world wide name might be copied into the initiator field. The field has a length of 23 characters to allow a copy of WWN with colons or other separator characters. However, they need to be removed before the initiator is created.

initiator create
WOLGA-SLOT3-(3)
210000e08b0ff748

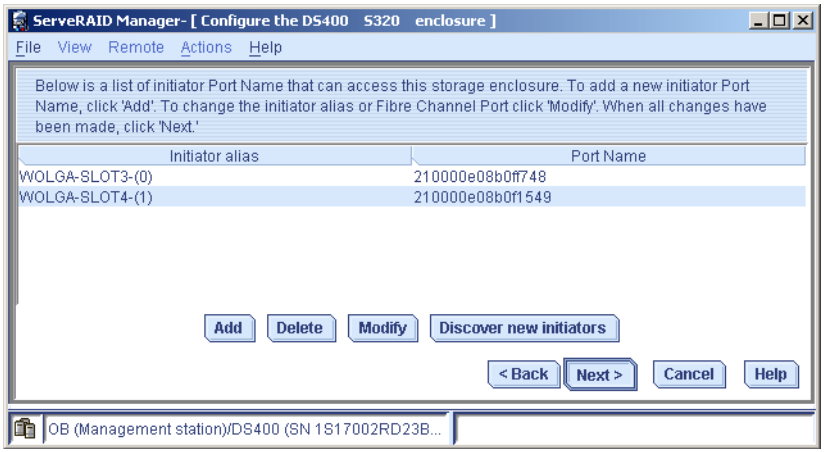

Figure 4-110 List of defined initiators

27. In Figure 4-111 on page 145, assign the logical drives to the initiator alias WOLGA-SLOT3-(0). The logical drives will be added under the initiator alias and removed from the enclosure window. Select the initiator alias WOLGA-SLOT4-(1). All logical drives are shown in the enclosure. Add the same logical drives to the selected initiator alias. A warning appears when a logical drive is assigned to more than one initiator. This should be required for systems with multiple HBAs installed and connected to a TotalStorage DS400 or for clustered systems. Click **Next**.

```
initiator manage
WOLGA-SLOT3-(0)
add WOLGA-BOOT 0

initiator manage
WOLGA-SLOT3-(0)
add WOLGA-DB 1

initiator manage
WOLGA-SLOT3-(0)
add WOLGA-LOG 2

initiator manage
WOLGA-SLOT3-(0)
add WOLGA-MEDIA
3

initiator manage
WOLGA-SLOT4-(1)
add WOLGA-BOOT 0

initiator manage
WOLGA-SLOT4-(1)
add WOLGA-DB 1

initiator manage
WOLGA-SLOT4-(1)
add WOLGA-LOG 2

initiator manage
WOLGA-SLOT4-(1)
add WOLGA-MEDIA
3
```

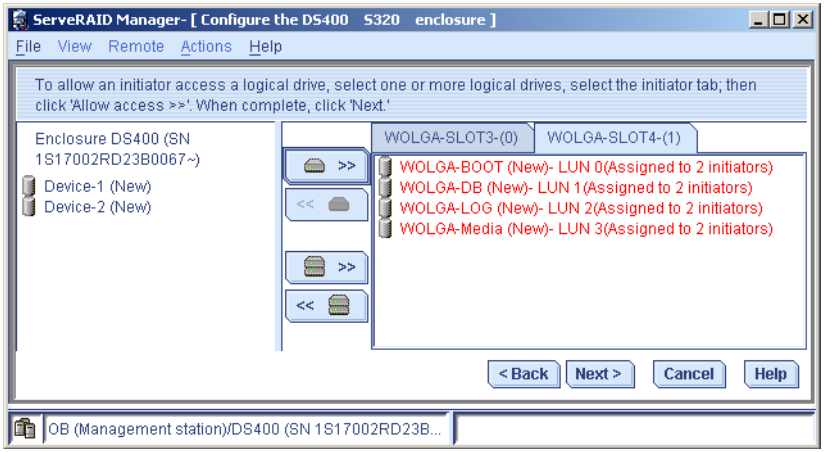

Figure 4-111 Assigned logical drives to initiators

28. When the order of assigning the logical drives to multiple initiators is not the same, a warning is shown after clicking **Next**. Usually there is no need to assign the same logical drive to more than one initiator. This only makes sense when multiple host bus adapters are installed in a server, or cluster nodes are attached to the TotalStorage DS300 or DS400. The logical drives need to be removed from the initiators and assigned in the same order to all initiators. See Figure 4-112.

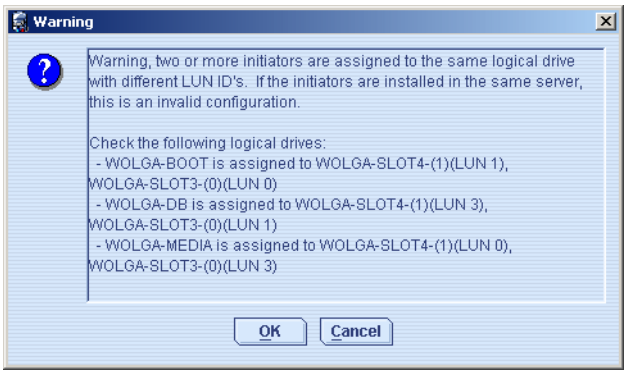

Figure 4-112 Warning about different LUN IDs assigned to multiple initiators

29. Verify setup and click **Apply**. In case something is not as desired, use the **Back** button and correct the configuration. The two drives which are not assigned to an initiator are not accessible by a host, and a message indicates this. See Figure 4-113 on page 146. Click **Apply**.

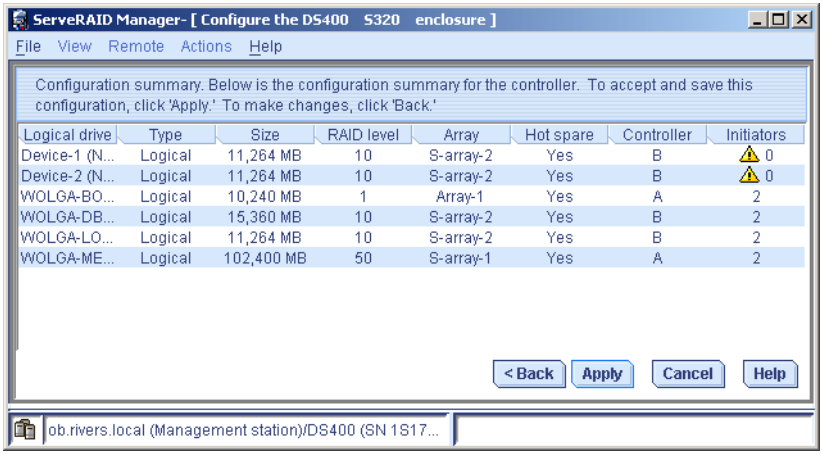

Figure 4-113 Configuration summary

30. Confirm the configuration and check for warnings. In this example, two logical drives are not accessible after applying the configuration. Apply the configuration by clicking **Yes**. Figure 4-114.

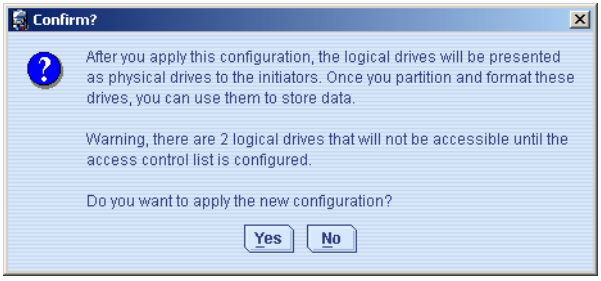

Figure 4-114 Confirm configuration

31. It will take some time to apply the configuration. The more complex the setup is, it will require more time to apply.

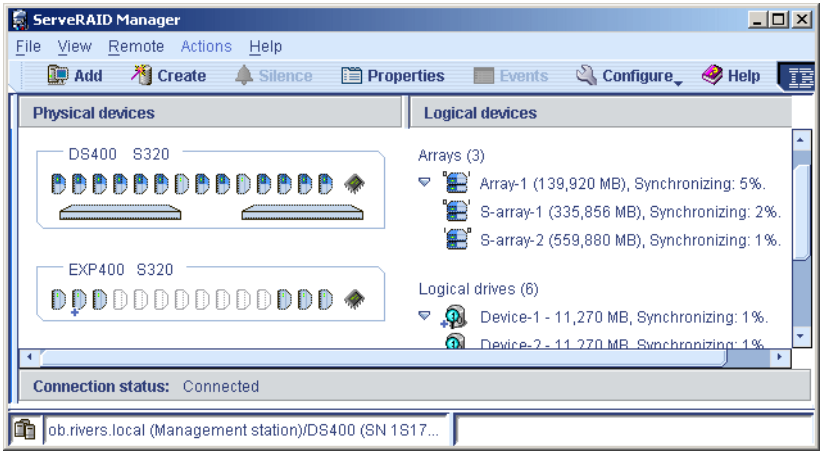

Figure 4-115 Applied configuration

4.8.3 Delete arrays on Total Storage DS300 and DS400

Arrays can be deleted. When an array is deleted, all logical drives of this array are deleted as well.

Note: Working with the CLI and deleting arrays requires that all logical drives be deleted first. See callout boxes for example commands.

1. Arrays are deleted by selecting the **array** → **right-click** → **Delete array**. All logical drives defined in the array also deleted. Initiators or user accounts are not deleted.

logical destroy
`Device-3`

array destroy
`Array-2`

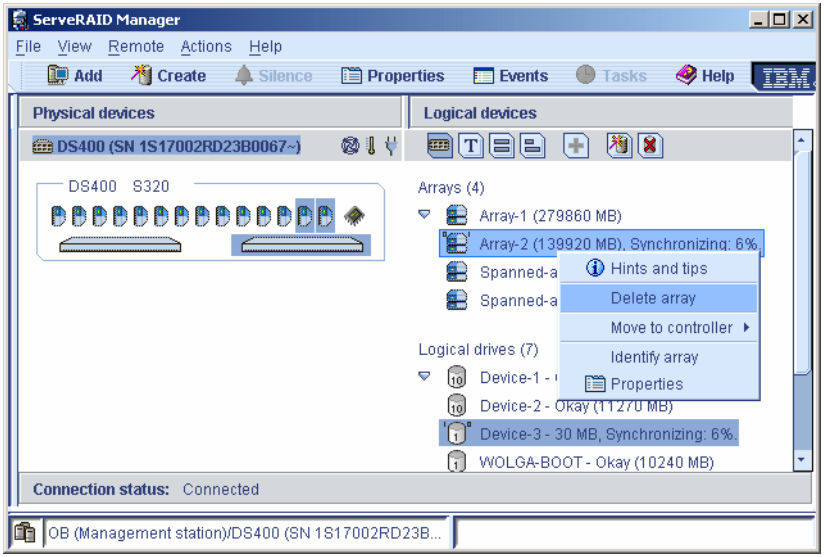

Figure 4-116 Select array to delete

2. Confirm the warning that all data of the logical drives in this array will be deleted by clicking **Yes**.

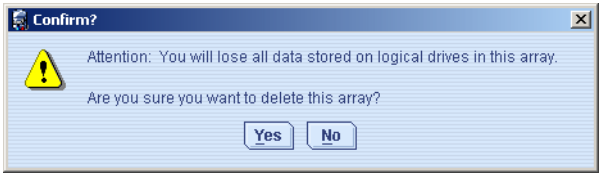

Figure 4-117 Array — data loss warning

3. After a while the change is applied and the array with the logical drive does not longer exist.

Note: Flashcopy target logical drives in this array must be deleted first.

4.8.4 Delete logical drives on TotalStorage DS300 and DS400

This section will explain how to delete logical drives. Follow the steps below:

1. Logical drives are deleted by selecting the logical drive with a **right-click** → **Delete logical drive**. The Delete logical drive button or the context menu might be used as well. Initiators or user accounts are not deleted when a logical drive is deleted.

logical destroy
`Device-1`

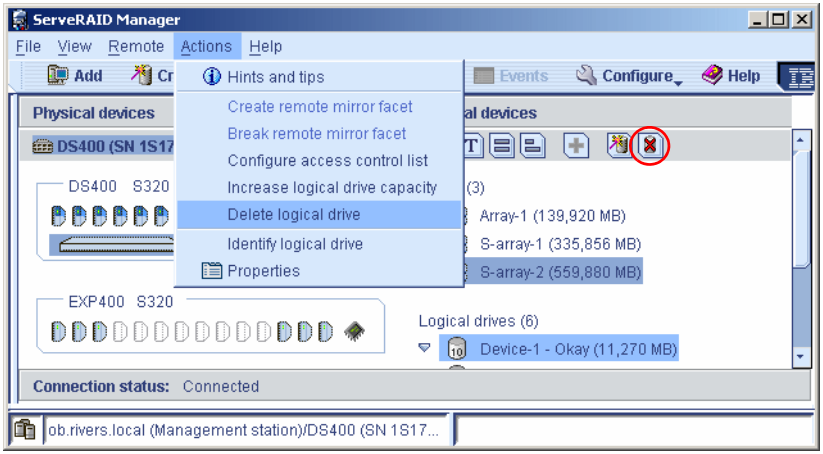

Figure 4-118 Select logical drive to delete

2. Confirm the warning that all data of the logical drive will be deleted by clicking **Yes**.

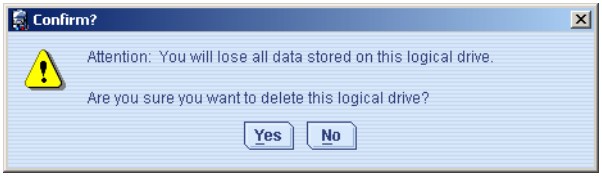

Figure 4-119 Logical drive — data loss warning

3. After a while the change is applied and the logical drive does not longer exist.

Note: Flashcopy target logical drives of this logical drive must be deleted first.

4.8.5 Create hotspare drives

This section will explain how to create a hotspare drive. Hotspares are drives which will be used to rebuild an array when one drive out of this array has failed. An array with a failed drive is in a critical condition and any further disk failure may make the array and all logical drives in this array inoperable. The chosen RAID level decides how many drives may fail until the array will be inoperable.

Two different versions of hotspares might be configured.

- Assigned hotspare drives
- Global hotspare drives

The difference between this two hotspare types are the arrays for which a hotspare will be used. Assigned hotspare drives will only be used for the array they are assigned for. Global hotspare drives will be used for all critical arrays. When both types of hotspares are configured, the assigned hotspares will be used first and than the global hotspares.

Multiple hotspares of either type can be assigned.

Note: Select a global hotspare drive that has the same capacity as the largest drives installed in the controller. Not all arrays and logical drives will be protected when drives with a smaller size are defined as hotspare. When small and large capacity drives are defined as global hotspares, the largest one is always selected first to rebuild an array with a disabled drive.

Create assigned hotspare drive

To create an assigned hotspare follow the steps below:

1. Right-click **drive 14** → **Create assigned hot-spare drive for** → **Spanned-array-1.**

array manage
Spanned-array-1
spare add 3584

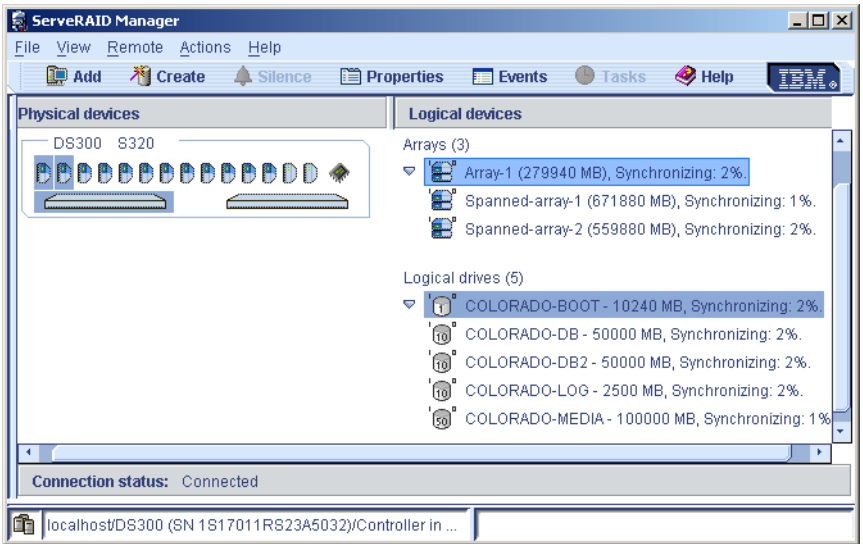

Figure 4-120 SRM - Create array specific hotspare

2. The selected dive 14 will get a small plus to symbolize that this drive will be used as a hotspare drive. Select the drive (by clicking) to see to which array it was assigned. The array and all logical drive will be highlighted as well.

Create a global hotspare drive

To create a global hotspare drive proceed with the step below.

1. To create a global hotspare select the drive which has to be chosen for a hotspare and click the **Create global hot-spare drive** button.

array spare add
3840

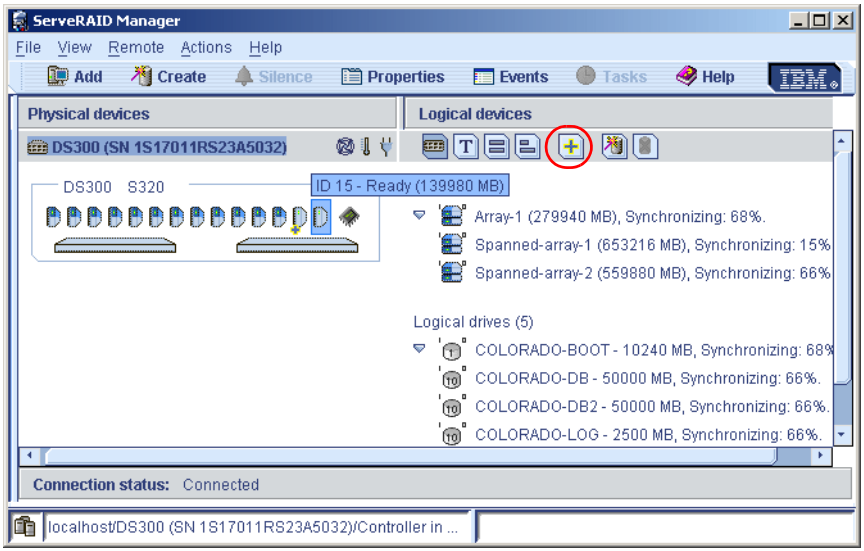

Figure 4-121 SRM - Create global hotspare

2. The properties of the hotspare drives will show the differences of those drives. By selecting the hotspare drive all protected logical drives will be marked.

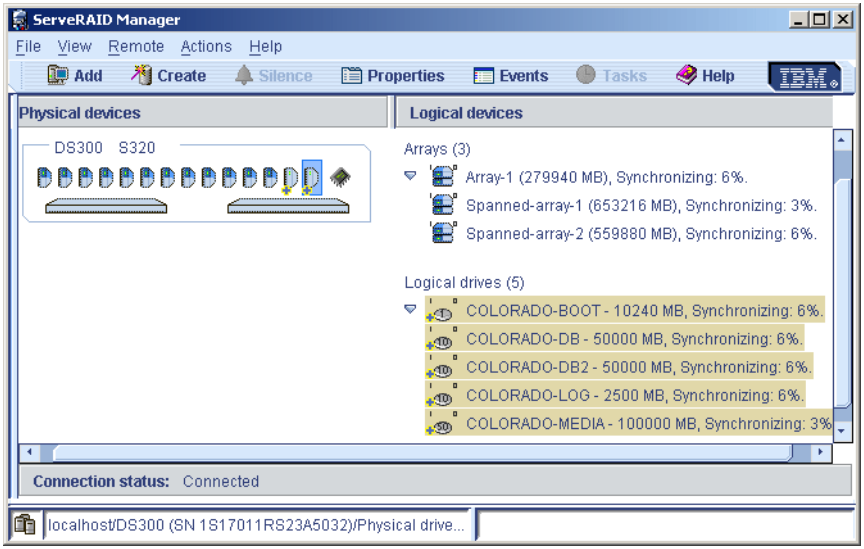

Figure 4-122 SRM - Logical drives protected by global hotspare

Now is the creation of the logical drives finished and they can be used by the hosts.

4.8.6 Increase logical drive capacity

The logical drive size can be expanded to a larger size when the array of the logical drive has free space. When no space is left in the array the size of the array or the RAID level of the array can be modified. Prerequisite for this functionality is firmware release 6.21. See 4.8.7, “Expand or migrate arrays” on page 155 for more details about increasing the space of an array.

To expand the size of an logical drive follow the steps below:

1. Use the context menu or the Action menu to increase the logical drive capacity. The following dialogs allow you to increase the capacity of the selected logical drive only. By using the Create logical drive button the next dialog has also the option to increase logical drive capacity. This option allows to modify all logical drives and into only one. The next steps use the Create logical drive button to start the expansion of the logical drive. The figure shows the current size of `10GB` of the logical drive `WOLGA-BOOT`.

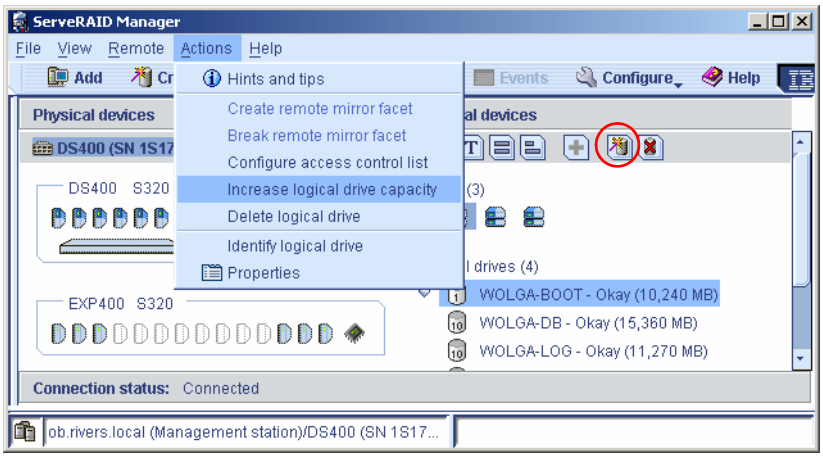

Figure 4-123 Menu Action → Increase logical drive capacity

2. Select the **Increase logical drive capacity** and click **Next**.

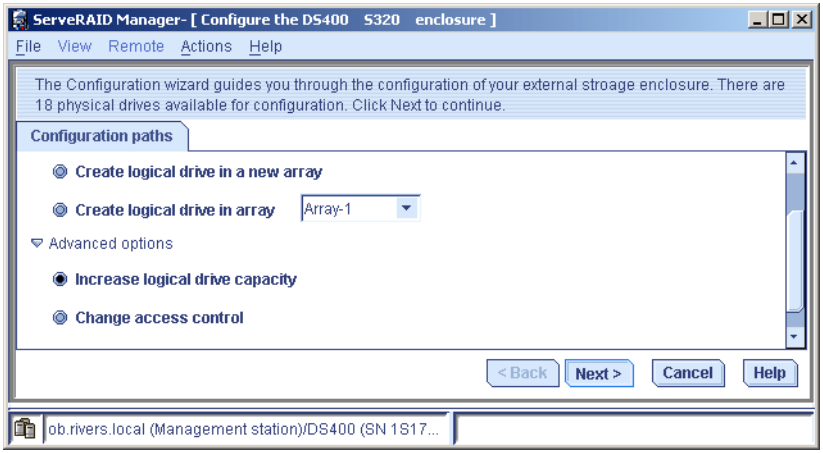

Figure 4-124 Select an option

3. The capacity of the logical drive `WOLGA-BOOT` after expansion has to be `20GB`. Additional space of `10GB` is required. Enter the amount of space the logical drive will be extended on in the **Additional (MB)** field. Then click **Next**.

> **Note:** The capacity unit depends on the display preferences and might be changed in **File** → **Preferences** register tab **Display option**.

logical manage WOLGA-BOOT **growby** 10240**M**

or :

logical manage WOLGA-BOOT **growto** 20480**M**

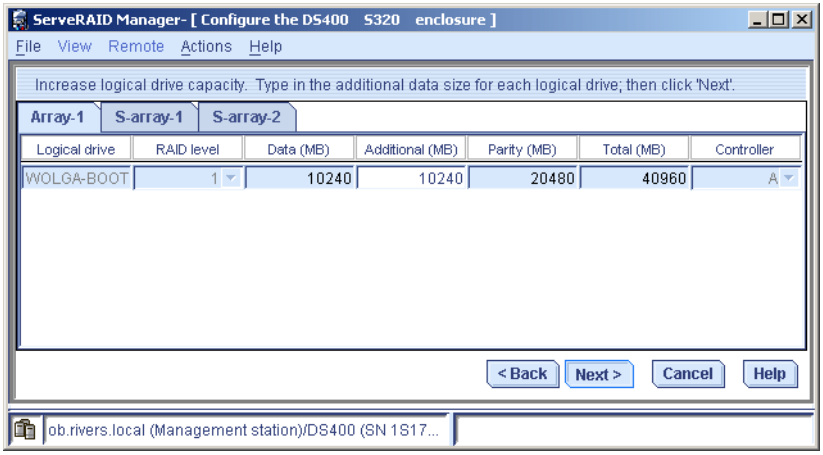

Figure 4-125 Enter the amount of space the logical drive is extended

4. Click the **Apply**.

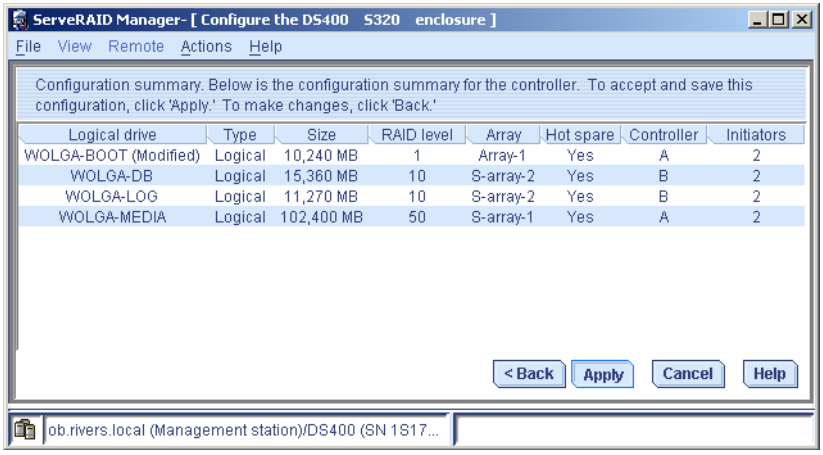

Figure 4-126 Configuration summary

5. Confirm the configuration by clicking **Yes**.

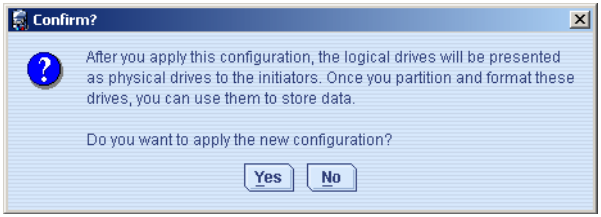

Figure 4-127 confirm configuration

The logical drive is now twice as large as before.

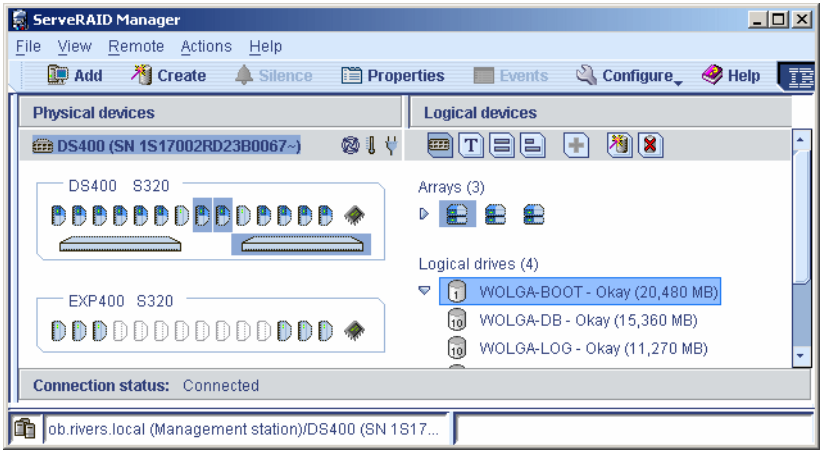

Figure 4-128 Extended logical drive WOLGA-BOOT

4.8.7 Expand or migrate arrays

With the release of firmware level 6.21 the size of an array can be extended or the RAID level can be migrated to an other RAID level. To use this function mostly unused drives are required. The array cannot be smaller after the operation than before. As long as a synchronize is running the The menu entry **Expand or migrate** does not exist.

Follow the steps below to migrate the Array-3 from RAID level 5 to RAID level 10 with 6 drives.

1. Select the array which need to be expanded or migrated. Use the menu **Action** → **Expand or migrate** or the context menu of the array `Array-2` **Expand or migrate** to migrate the Array-2 from RAID level 5 to RAID level 10 with 6 drives instead of 3.

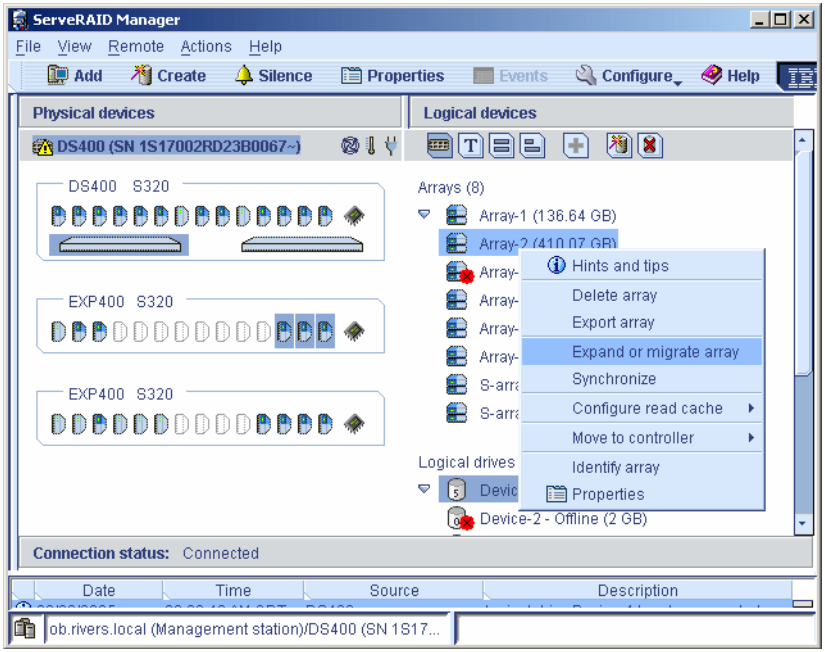

Figure 4-129 Expand or migrate array

2. Select **RAID 10** from the **Advanced settings** section and click **Next**.

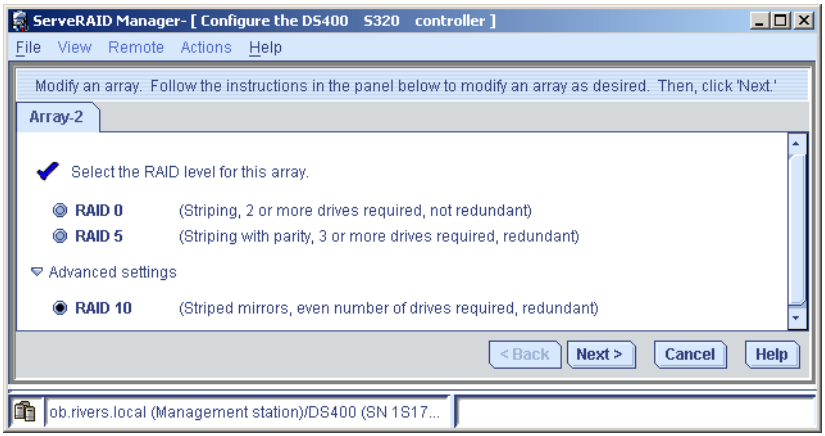

Figure 4-130 Choose RAID level for the migration

3. This screen provides help for the migration process. It gives an overview about the expected result and lets define the drives which will be used in the migrated array. On the right side of the blue arrow is information provided about the required additional drives for the migration process. This test is

updated with the selection of additional drives. The advanced settings section below the blue arrow provides an forever of the values of the migrated array.

Sub logical drives are needed to migrate to a RAID 10 array. This option defines how many RAID 1 arrays are used to build the RAID 10 array.

The physical view contains all drives in the enclosures. The drives currently in the array defined are already selected.

Note: When only new drives for the migrated array are selected and the original drives are deselected, ServeRAID Manager will show those drives as Alien drives until the migration is finished.

Proceed with the next step.

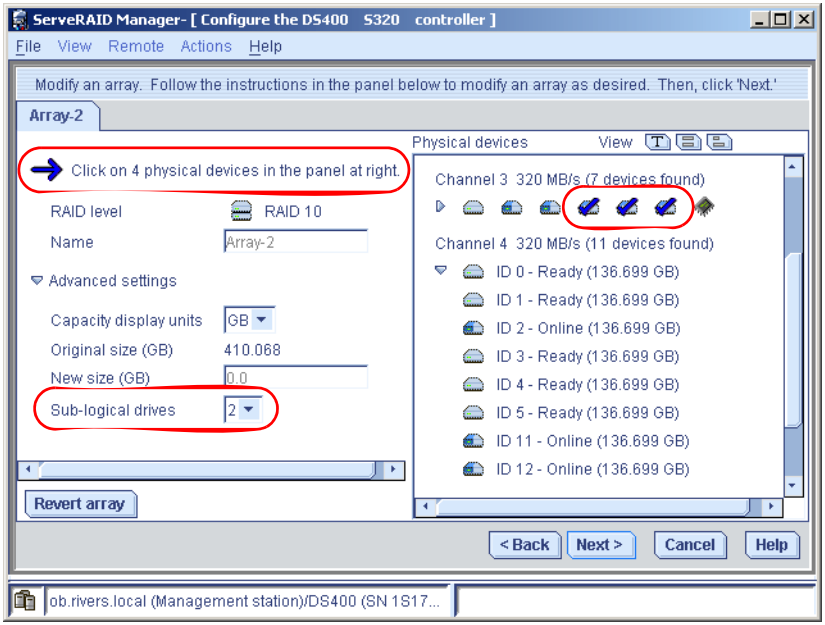

Figure 4-131 Physical drives view

4. Set sub logical drives to 3. Select three additional drives. Then click **Next**.

Note: The new size value will be 0 until a valid configuration is selected.

array migrate
`Array-2` **raid10 64**
197376 197632
197888

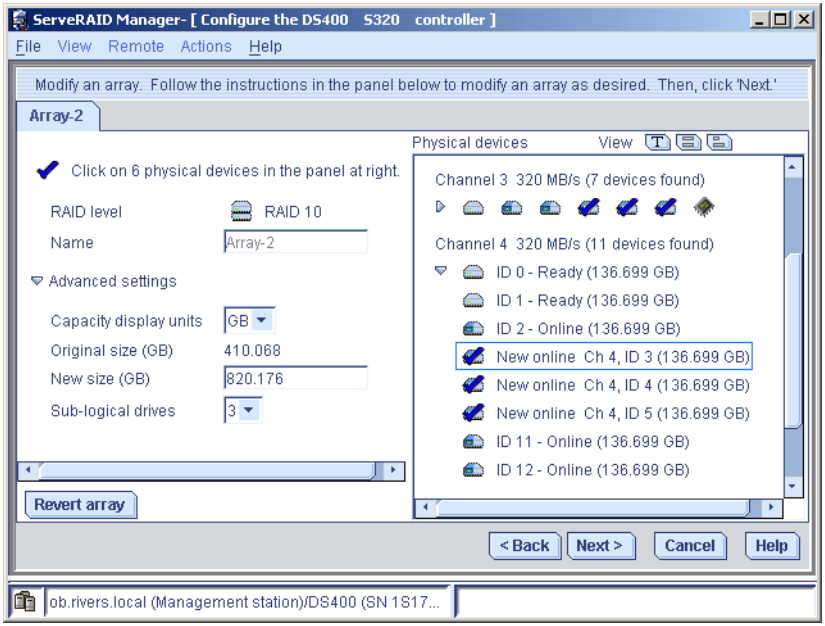

Figure 4-132 Select new drives for the migrated array

5. Click **Apply**.

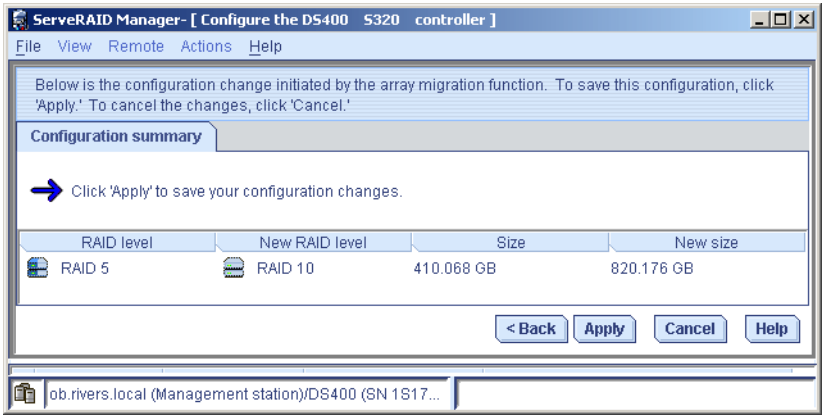

Figure 4-133 Configuration summary

6. Confirm the configuration by clicking **Yes**.

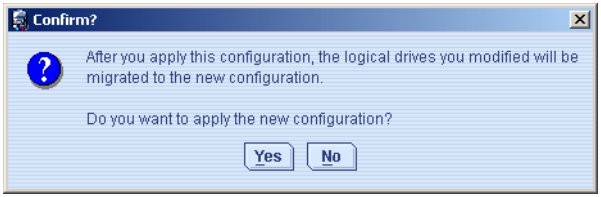

Figure 4-134 Confirm the configuration

7. Array migration or Expansion is in progress.

> **Note:** All logical drives in the array will be offline for the period of the migration or expansion. The logical drives become online after finishing the migration or expansion process.
>
> When Original drives are deselected, those drives are shown as alien drives until the migration is finished.
>
> There is no information in the ServeRAID manager console available that a expansion or migration is in progress. Use the CLI to check the status of the logical drives to determine when a problem has occurred or if process is still running.

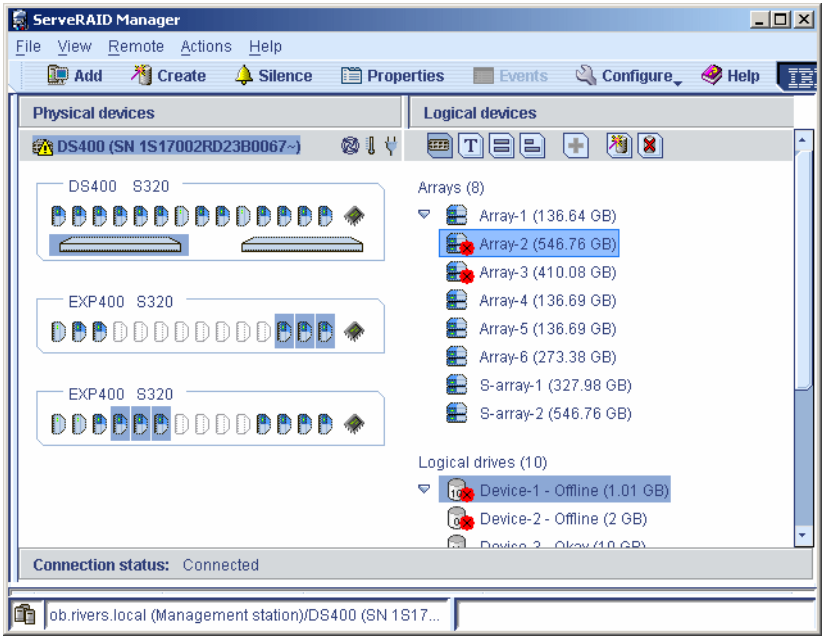

Figure 4-135 Array migration in progress

8. Verify the status of the logical drive from the CLI.

array manage
array-2 **info**

```
DS400[A]# array manage array-2 info

----[ array 'Array-2' ]-----------------------------------------------

               Status:  Migrating  (0% complete, task rate high)
                 Type:  RAID 10
           Controller:  [B] remote
  Preferred Controller:  [B] remote
         Stripe Depth:  64 KB
           Read Cache:  Enabled
          Write Cache:  Disabled

             Capacity:  273.38GB       (573,317,120 512 byte blocks)
            Available:  272.37GB       (100%)

----[ Group '0' ]-----------------------------------------------------

 ID     Device                                Capacity       Status
-------------------------------------------------------------------------
134144  IBM-ESXS Disk (enclosure 2 slot 11)  136.73GB       online
134400  IBM-ESXS Disk (enclosure 2 slot 12)  136.73GB       online

----[ Group '1' ]-----------------------------------------------------

 ID     Device                                Capacity       Status
-------------------------------------------------------------------------
134656  IBM-ESXS Disk (enclosure 2 slot 13)  136.73GB       online
197376  IBM-ESXS Disk (enclosure 3 slot 3)   136.73GB       online

----[ Group '2' ]-----------------------------------------------------

 ID     Device                                Capacity       Status
-------------------------------------------------------------------------
197632  IBM-ESXS Disk (enclosure 3 slot 4)   136.73GB       online
197888  IBM-ESXS Disk (enclosure 3 slot 5)   136.73GB       online
-------------------------------------------------------------------------

logical                          Used   (% of array)
-------------------------------------------------------------------------
Device-1                      1030.00MB   (0.37%)
-------------------------------------------------------------------------


DS400[A]#
```

Figure 4-136 RAID 10 array information

4.9 Administer Management Station user accounts

The Management Station has its own user management. Operators and administrators can be defined. Operator has the permission to look at items. Administrator can modify items. The operator level can only by created by a command line utility USER_TOOL. This utility is located in the installation directory of the IBM ServeRAID Manager.

Note: Management Station and USER_TOOL are only available in the Microsoft Windows operating system version of the IBM ServeRAID Manager.

Section 4.9.1, “Administer Management Station user accounts with ServeRAID Manager console” on page 162 explains the user management with the ServeRAID Manager console.

4.9.2, “Administer Management Station user accounts with USER_TOOL” on page 163 explains the account management from the command line.

4.9.1 Administer Management Station user accounts with ServeRAID Manager console

1. To manage the user accounts of the Management Station right-click the Management Station **Management station actions** → **Configure**.

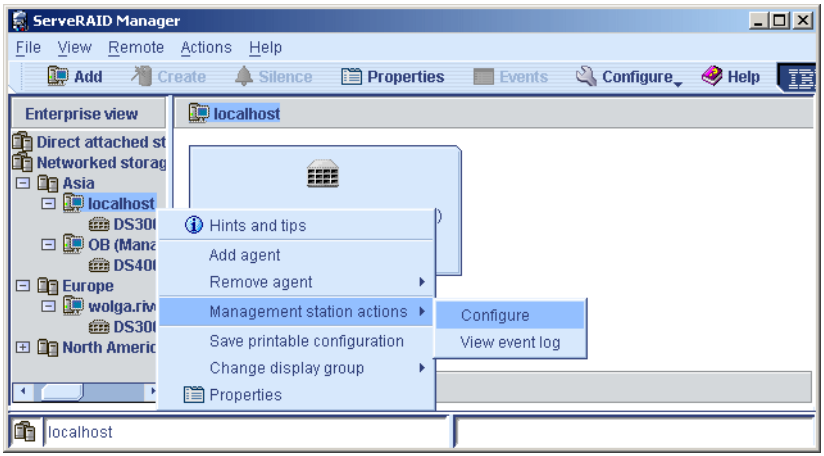

Figure 4-137 Configure Management Station

2. The configuration dialog of the Management Station opens. The listed user is the user created during installation of the ServeRAID Manager.

The user can be deleted by selecting the user and pressing the **Delete user** button.

Note: The last user cannot be deleted.

A new password can be assigned by selecting the user and clicking the **Modify user** button.

Note: The authority level of new or existing account cannot be changed from the ServeRAID Manager console.

To add a user click the **Add user** button. User name and password can be specified. All new created users are administrators.

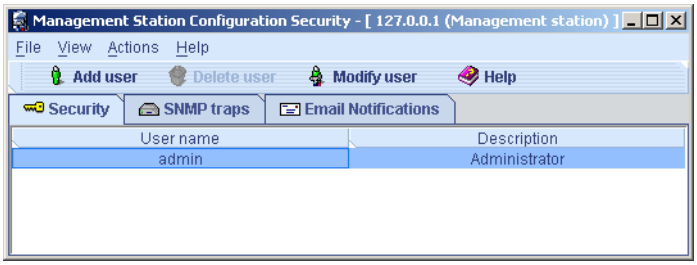

Figure 4-138 Security configuration of the Management Station

4.9.2 Administer Management Station user accounts with USER_TOOL

The installation directory of the IBM ServeRAID Manager contains a command line utility called USER_TOOL.EXE. This utility allows you to create user accounts for the ServeRAID Manager Management Station service on the command line.

Note: Do not create accounts with an authority level other than 3. This is the administrator level. Right now only accounts with an authority level of 3 are working.

Run the utility without specifying any options and a help text is shown.

```
C:\Program Files\IBM\ServeRAID Manager>user_tool.exe
Usage:
user_tool.exe list
user_tool.exe add user password level
user_tool.exe remove user
user_tool.exe change user level

C:\Program Files\IBM\ServeRAID Manager>user_tool.exe list
```

Figure 4-139 User_tool usage

In addition to the functionality of the ServeRAID manager console, this utility can specify the authority level of existing or new accounts.

To add a new account `administrator` with administrative rights, enter following command

```
user_tool add administrator passw0rd 3
```

The 3 stands for an account with administrative rights.

.

```
C:\Program Files\IBM\ServeRAID Manager>user_tool.exe add administrator
passw0rd 3
['user_tool.exe', 'add', 'administrator', 'passw0rd', '3']

C:\Program Files\IBM\ServeRAID Manager>
```

Figure 4-140 Create account with administrative authority

To add a new account `operator` with operator rights enter following command **user_tool add** `operator passw0rd` **1**.

```
C:\Program Files\IBM\ServeRAID Manager>user_tool.exe add operator passw0rd 1
['user_tool.exe', 'add', 'operator', 'passw0rd', '1']

C:\Program Files\IBM\ServeRAID Manager>
```

Figure 4-141 Create account with operator authority

A list of all defined accounts are retrieved with the list option.

Note: ServeRAID Management Station service may need a restart to reflect the modifications.

Stop and start the ServeRAID Management Station service from the services or from command line:

```
C:\>net stop "ServeRAID ManagementStation Service"
The ServeRAID Management Station Service service is stopping.
The ServeRAID Management Station Service service was stopped successfully.


C:\>net start "ServeRAID Management Station Service"

The ServeRAID Management Station Service service was started successfully.


C:\>
```

Figure 4-142 Restart Management Station service

List all available accounts with the **list** option.

```
C:\Program Files\IBM\ServeRAID Manager>user_tool.exe list
['user_tool.exe', 'list']
admin 3
administrator 3
operator 1

C:\Program Files\IBM\ServeRAID Manager>
```

Figure 4-143 List Management Station accounts

Delete accounts by using the option remove. The example shows the deletion of the user `administrator`.

Note: The command line utility **USER_TOOL.EXE** will reflect deletions of accounts immediately. ServeRAID Manager console requires a restart of **ServeRAID Management Station service**.

```
C:\Program Files\IBM\ServeRAID Manager>user_tool.exe remove administrator
['user_tool.exe', 'remove', 'administrator']

C:\Program Files\IBM\ServeRAID Manager>user_tool.exe list
['user_tool.exe', 'list']
admin 3
operator 1

C:\Program Files\IBM\ServeRAID Manager>
```

Figure 4-144 Delete Management Station account

The created accounts after restarting the ServeRAID Management Station service will look like this. The description shows the authority level of the existing accounts.

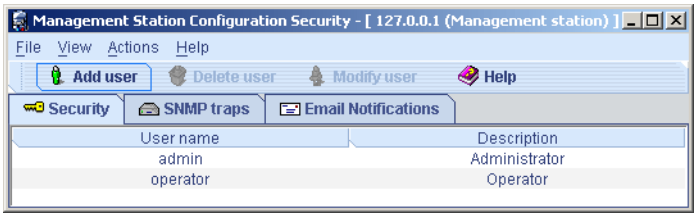

Figure 4-145 Management Station user accounts

Change the authority level by using the option change. The example below will change the authority level for the user `operator` from operator (1) to administrator (3).

```
C:\Program Files\IBM\ServeRAID Manager>user_tool.exe change operator 3
['user_tool.exe', 'change', 'operator', '3']

C:\Program Files\IBM\ServeRAID Manager>user_tool.exe list
['user_tool.exe', 'list']
admin 3
operator 3

C:\Program Files\IBM\ServeRAID Manager>
```

Figure 4-146 Change user authority level

4.10 Install premium feature license key

A premium feature allows the use of additional functions of the IBM TotalStorage DS300 and DS400. These additional premium features are, for example, support for 5 to 256 Flashcopys, and the EXP400 expansion unit attachment. When one of these premium features is purchased, a booklet with a *purchase validation code* is delivered. This booklet also contains documentation on the procedure to follow to retrieve the feature *license key*. See 3.6, “Premium feature license keys” on page 54.

When a premium feature license key is retrieved, it must be installed into the DS300/DS400 controller. This section describes the installation of a premium feature license key using the EXP400 expansion unit premium feature as an example.

1. Using the command line interface, the command **show license** displays the license information for the IBM DS300 or DS400. It also provides information about available additional features. In the example shown in Figure 4-147, the external channels are not enabled (the EXP400 expansion unit uses external channels) and the default number of 4 snapshots (Flashcopys) are allowed.

```
DS400[A]> show license
------------------------------------------------------------
                 Output for 'license'
------------------------------------------------------------

----[ License key ]-------------------------------------------

          ABCD-EFGH-IJKL-MNOP-TUVW-XYZ4

Enabled Features:
        Number of snapshots:  4
  Enabled external channels:  0
        Maximum drive count:  14

DS400[A]>
```

Figure 4-147 Output of the command show license

2. After the EXP400 license key has been retrieved following the procedure described in 3.6, “Premium feature license keys” on page 54, select the DS400 where you want to install the key. Open menu **Action** → **Add feature key**.

Note: Be very careful to select the correct DS300 or DS400. License keys cannot be transferred from one DS300 or DS400 to another!

license set
XXXX-XXXX-XXXX-XX
XX-XXXX-XXXX

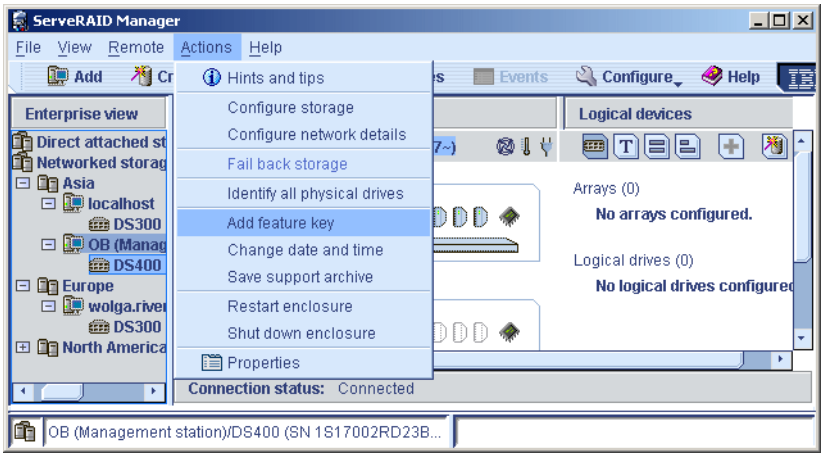

Figure 4-148 Menu actions → Add feature key

3. Enter the retrieved license key for this enclosure and click **OK**.

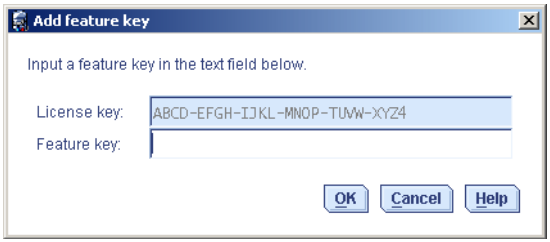

Figure 4-149 Dialog add feature key

4. A message dialog opens that the license key has been entered successfully. Click **OK**.
5. The output of the command line interface command **show license** now shows the enabled external channels.

```
DS400-1[A]# show license

------------------------------------------------------------
                Output for 'license'
------------------------------------------------------------

----[ License key ]--------------------------------------------

          ABCD-EFGH-IJKL-MNOP-TUVW-XYZ4

Enabled Features:
        Number of snapshots:  4
  Enabled external channels:  2
        Maximum drive count:  1024

DS400-1[A]#
```

Figure 4-150 Show license information

6. A reboot might be required to activate the new features.

The DS300/DS400 will not activate the premium feature if the license key has been entered improperly. If the license key is lost for any reason, a case must be opened with IBM support.

4.10.1 Loosing the premium feature license key

The premium feature license key might be lost under some circumstances. They are:

- Controller replacement in a single controller TotalStorage DS300 or DS400
- Midplane replacement in a dual controller TotalStorage DS300 or DS400
- Replacement of both controllers in a dual controller Totalstorage DS300 or DS400

If one of these circumstances occurs and the premium feature key is lost, contact IBM technical support to get a new feature key. A proof of purchase is required! Support for the retrieval of a replacement feature key may be available at the following Web site at a future date:

http://www.advancedfunctions.com

To find the technical support services in a specific country go to the following URL:

http://www.ibm.com/planetwide/

4.11 Flashcopy or SnapShot backup

Flashcopy and SnapShot are different names for the same feature. IBM uses the term SnapShot to refer to the Flashcopy feature. Other vendors may use the term SnapShot. The command line interface uses also the term SnapShot. ServeRAID Manager console uses Flashcopy.

A Flashcopy backup creates a point-in-time backup of logical drive data. Flashcopy backup might be used for drive cloning or running a test system with productive data without the need of stopping production. The greatest advantage of Flashcopy is the reduction of the backup window for an large application like an database. Backup of the files of the application requires that the application is stopped or has at least closed all open files. The data are then in a consistent state and a backup can be performed. Depending on the file size and number several hours or days might be needed to backup everything. During this time the application is not available for productive usage. This is usually not acceptably by a customer. Flashcopy may help in such a situation.

After stopping the application and getting a consistent state of the data an Flashcopy of the logical drives used by this application can be created. This will create a new logical drive with the date in this consistent state. This will only takes a few minutes. The application be started again and the new created logical drive can be used to perform a regular backup

To create a Flashcopy logical drive, an array with at least one logical drive must exist. The array needs space for the Flashcopy target logical drive.

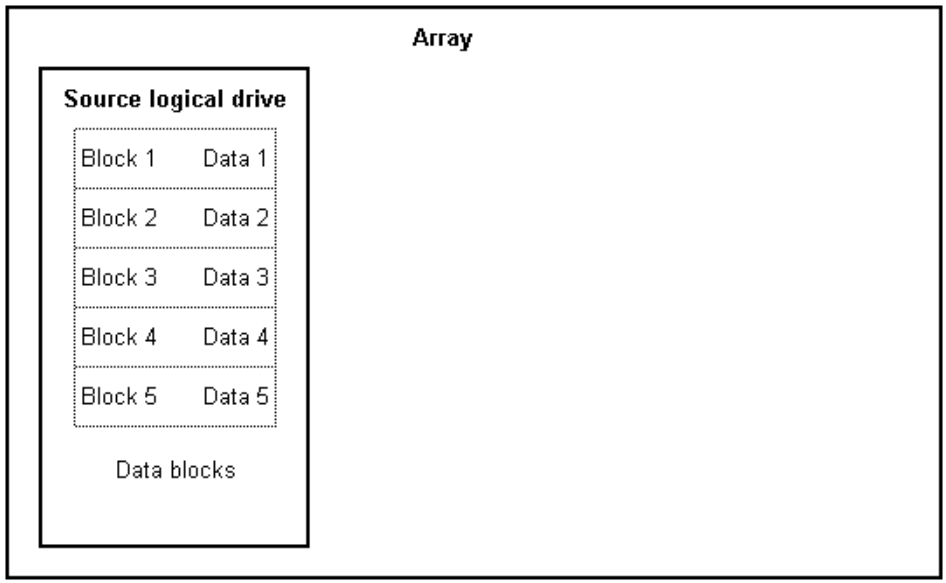

Figure 4-151 Array with logical drive and data

A Flashcopy of an existing logical drive is created by using the snapshot command in the command line interface. The process of creating the Flashcopy includes the creation of a new logical drive in the same array than the source logical drive. The target logical drive will have 10% of the capacity of the source logical drive. It contains two data regions. One region is used for changed data on the source logical drive and the other region is used for a lookup table. The lookup table maintains a relationship between the data blocks of the source and target logical drive. There are no data block in the target logical drive at this time.

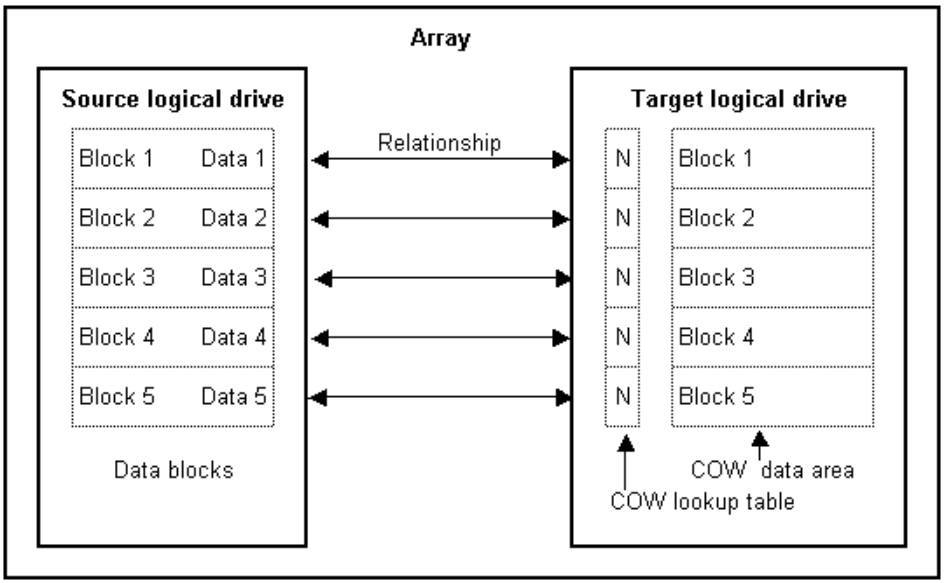

Figure 4-152 Array with Flashcopy source and target logical drive

The Flashcopy target logical drive can now be assigned to an other host than the source logical drive is assigned to. This is important because the date which can be read from both logical drives is identical. This means that the disk signature written by a Microsoft Windows operating system is also identical. This will cause problems when the Flashcopy target logical drive is assigned to the same host.

To read data from the Flashcopy target the controller is looking into the Copy-On-Write lookup table. This table explains if the requested block of data is already in the COW data area or must be read from the source logical drive. After the creation of the Flashcopy target logical drive, no data block are in the COW area and all blocks are read from the source logical drive.

Before a block with new data on the source logical drive is overwritten a background task in the storage enclosure will copy this block with the old data to the Flashcopy target logical drive. The lookup table is modified to reflect that the data block is now on the Flashcopy target logical drive. The relationship to the data block on the source logical drive is not longer required and removed. Future changes of this data block on the source logical drive will not be copied to the target logical drive. This kind of Flashcopy is also called Copy-On-Write (COW) because the data copy occurs as soon as a write to a data block is performed.

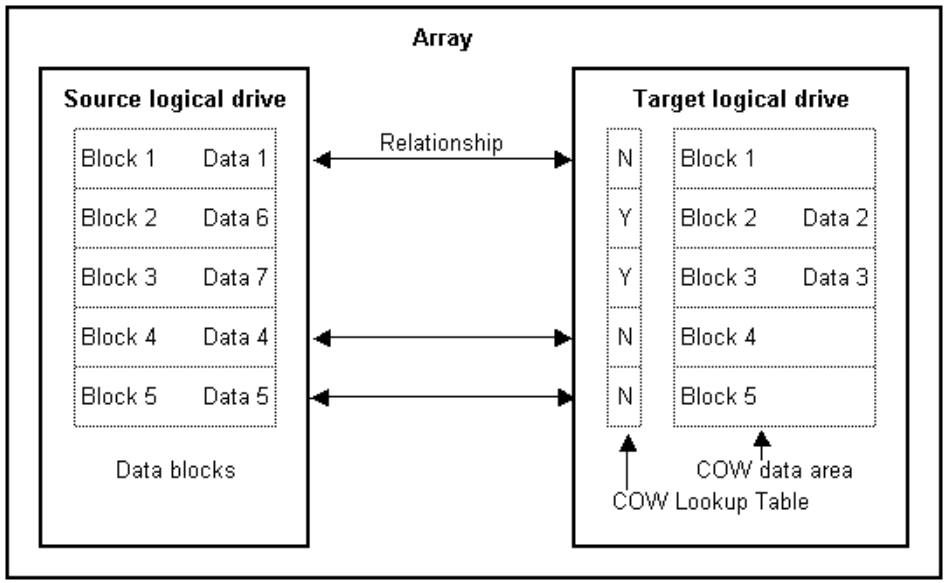

Figure 4-153 Array with Flashcopy source and target logical drive - Changed data on source logical drive

When the Flashcopy target logical drive was in a read/write mode assigned to an host system the data on the target logical drive can be changed. Reading of data is the same under a read/write conditions. When data is not in the COW area it will be taken from the source logical drive. Writing the data in the COW area will just change this data. When the data was read from the source logical drive and is written to the Flashcopy target logical drive the COW lookup table will be modified and the relationship between the data blocks dissolved.

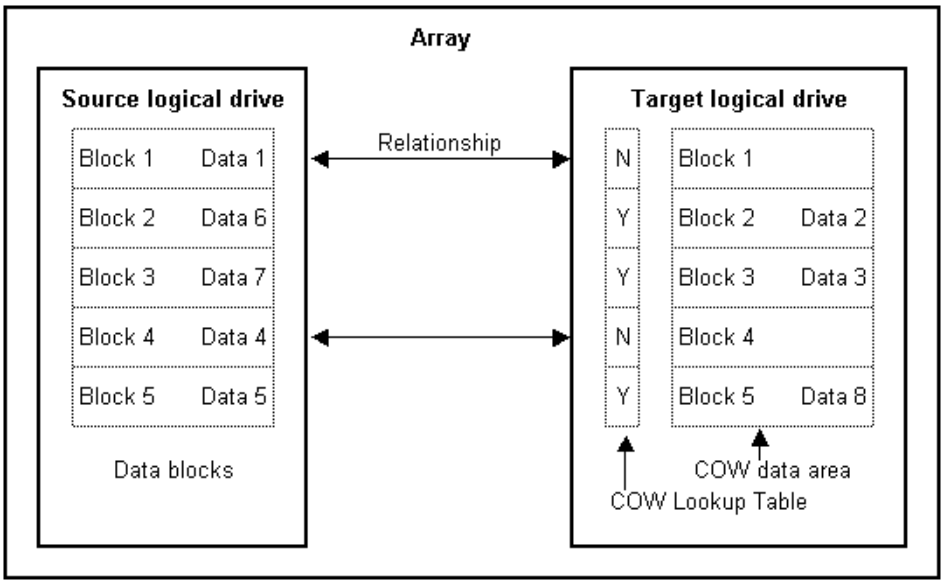

Figure 4-154 Array with Flashcopy source and target logical drive - Changed data on target logical drive

As more data is changed in the time of an backup or by writing on the target logical drive as more free space in the array is required to increase the size of the target logical drive. The target logical drive can grow up to the same size as the source logical drive. When there is not enough space for the increase of the target logical drive is available than becomes the Flashcopy invalid and must be deleted.

Multiple Flashcopys can be created from a single logical drive. The Flashcopys will build a Flashcopy group. The COW Lookup table and the COW data area shared between Flashcopys of the same Flashcopy group. Every Flashcopy group requires 2MB system memory. The firmware allows up to 4 Flashcopys per enclosure. When more Flashcopys are required, a feature key must be purchased to upgrade the possible four Flashcopys to 256. Depending on the usage, more system memory is also required. IBM offers a system memory upgrade to 1GB for the TotalStorage DS400 and the TotalStorage DS300 model 1701-1RS and 1701-2RD.

Note: With the TotalStorage DS300 1701-1RL, only 4 Flashcopys can be used. The system memory upgrade and the Flashcopy upgrade to 256 Flashcopys are not supported.

To create Flashcopy or SnapShot from a logical drive perform the following steps:

1. Open a Telnet session to on of the management ports of the TotalStorage DS300 or DS300.
2. The **show license** command informs about the maximum number of SnapShots that can be used.

```
DS300-1[A]# show license

------------------------------------------------------------
                Output for 'license'
------------------------------------------------------------

----[ License key ]-------------------------------------------

          ABCD-EFGH-IJKL-MNOP-TUVW-XYZ4

Enabled Features:
        Number of snapshots:  4
  Enabled external channels:  0
        Maximum drive count:  14

DS300-1[A]#
```

Figure 4-155 Show license information

3. Create a snapshot from the logical drive DANUBE-MAIL and call the SnapShot target logical drive DANUBE-MAIL-S1. The example below does not define an initiator which will access the SnapShot target logical drive. This can be done later or by adding the initiator aliases behind the command.

Note: It is important that your data is in a consistent state and no I/O is sent to the Flashcopy source logical drive. This can be archived by shutting down the system. This will disconnect the initiators from the logical drive. Other options are possible depending on the usage of the logical drive and the used application. Stopping I/O to a logical drive from which it is booted, for example, is not possible from the operating system.

```
DS300-1[A]# snapshot create DANUBE-MAIL-S1 DANUBE-MAIL readonly restricted

WARNING - all initiators MUST be disconnected to

prevent loss of data - proceed?
[No]: yes

Please wait....

Created the snapshot called DANUBE-MAIL-S1
Warning - this logical drive cannot be accessed until an ACL is defined for
it


DS300-1[A]#
```

Figure 4-156 Create snapshot

4. From the ServeRAID Manager console, the Flashcopy logical drives can be seen, deleted, access control and authentication be managed like every other logical drive as well. See 4.8, “IBM TotalStorage DS300/DS400 array and logical drive management” on page 113.
5. ServeRAID Manager console shows the created Flashcopy target logical drives with its own symbol to differentiate between regular logical drives and the Flashcopy target logical drives. Figure 4-157 on page 177 shows two Flashcopy target logical drives `COLORADO-DB-2005.08.19.001` and `COLORADO-DB-2005.08.19.002`.

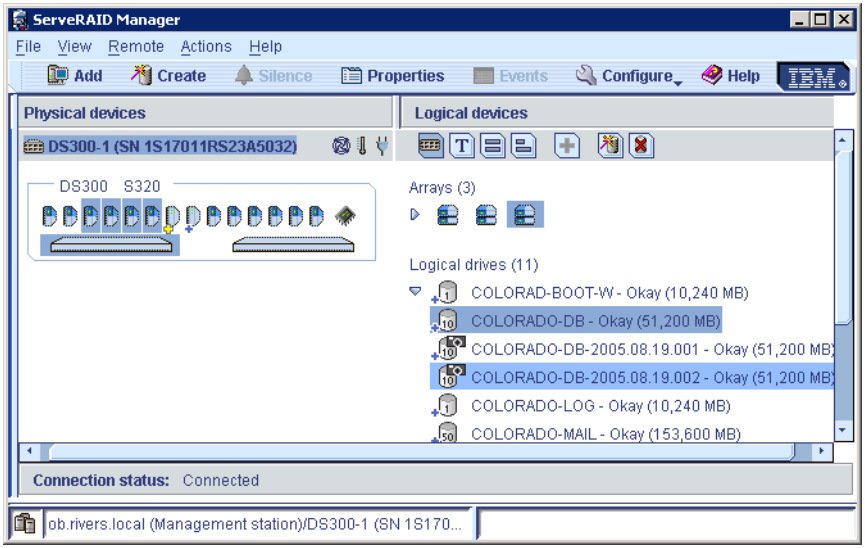

Figure 4-157 ServeRAID manager console with Flashcopy target logical drives

6. From the property windows of the source logical drive the date when the Flashcopy was taken can be found. Select the source Flashcopy logical drive and use the menu **Actions** → **Properties** or the context menu to open the property window. The corresponding Flashcopy target logical drives will be highlighted.

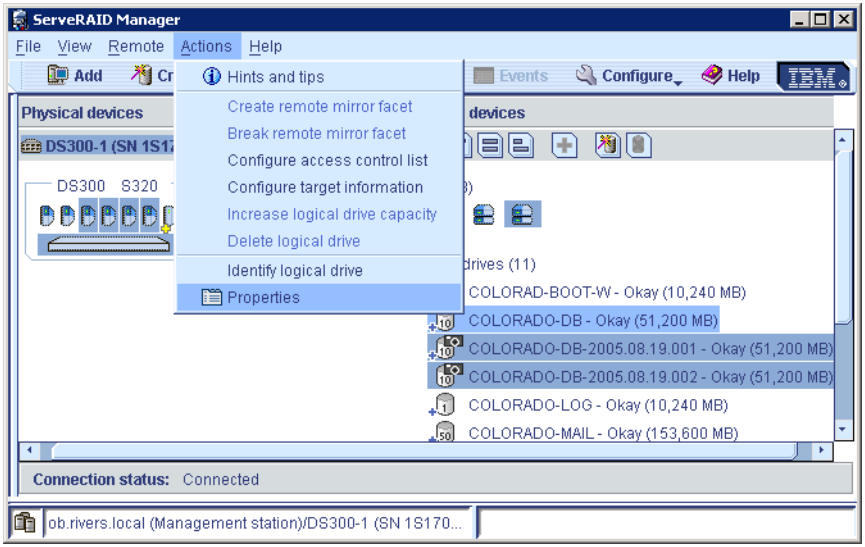

Figure 4-158 ServeRAID Manager menu action → properties

7. On the **Device** tab the Flashcopy type Parent can be seen.

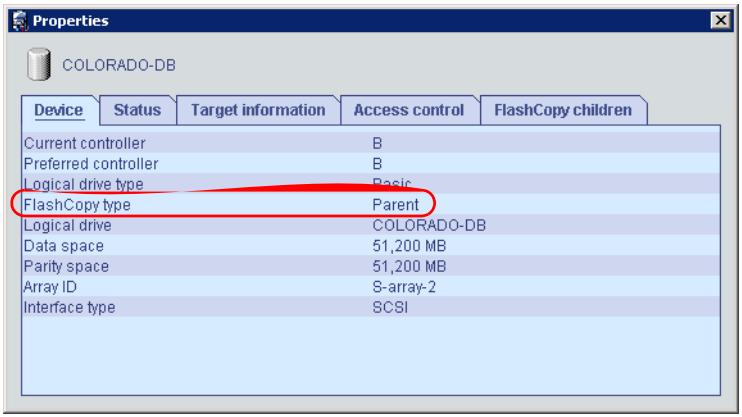

Figure 4-159 Flashcopy source logical drive — Device properties

8. The tab **Flashcopy children** exists only for Flashcopy source logical drives. This tab lists the dates when the Flashcopy was taken. It also shows the **Data Space** and the **Virtual space**. Data space is the current space required for this Flashcopy in the array. Virtual space shows the maximum size to which the Data space can grow up. In worst case this amount of free space is in the array required. This depends on the amount of changed data in the source or target logical drive. The figure shows also the second Flashcopy does not require data space. This is because the COW area and COW lookup table are shared between those two Flashcopys.

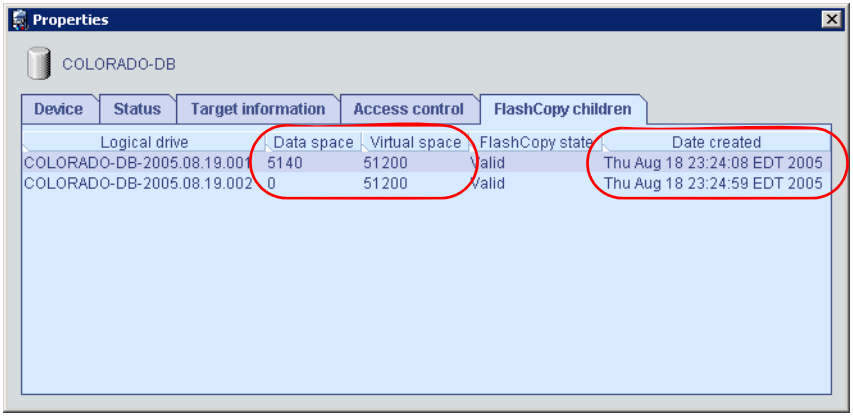

Figure 4-160 Flashcopy children properties of a Flashcopy source logical drive

Note: The value of the data space can only be found in the ServeRAID manager console. Command line interface does not provide this value.

4.12 Update TotalStorage DS300 and DS400 firmware

Two different firmware images can be stored on a DS300 and DS400. During boot of the controller, one of the images can be selected by pressing 1 or 2 for the first or second image. One firmware image is active and used to run the controller and the other is passive. Only the passive firmware image can be overwritten during a upgrade.

Note: The firmware update can be performed during normal operation. There is no danger because the update process does not restart the DS300 or DS400.

After the update is performed, choose a safe time for a reboot of the DS300 or DS400. After the reboot, the DS300 and DS400 will load the new firmware image because the update process sets the next boot image to the updated image number.

Firmware of the DS300 and DS400 can be updated with the ServeRAID Manager or over the CLI.

4.12.1 Update DS300 and DS400 firmware with CLI

TotalStorage DS300 or DS400 firmware update can be performed by using the command line interface. Three protocols for transferring the firmware file to the TotalStorage DS300 or DS400:

- FTP protocol
- HTTP Protocol
- 1k Xmodem protocol (not recommended)

The update of the firmware is performed one controller after the other.

Note: Do not perform the firmware update to both controller at the same time!

After upload and flash of one controller is finished, the command line interface informs you of the required reboot. Do not update the alternate controller before reboot. If you update both controllers before reboot, it is possible that you could cause the wrong configuration to be imported on the next reboot. You should only

update one controller and let the auto-synch function update the second controller.

Up to three reboots might be required to load a different firmware image. All three updates require a reboot and are performed after the switch to the different firmware image.

To perform a firmware update using the command line interface follow the outline steps below:

1. Open Hyperterminal and create a new connection. Enter a name for the new connection and click **OK**.

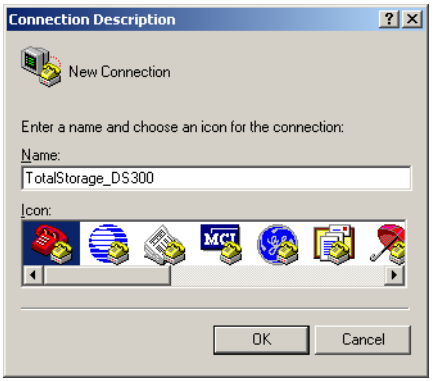

Figure 4-161 Hyperterminal — New connection

2. Select TCP/IP (WInsock) from Connect using:

Figure 4-162 Hyperterminal — Connect to

3. Enter the IP address of the management port of controller A. Click **OK**.

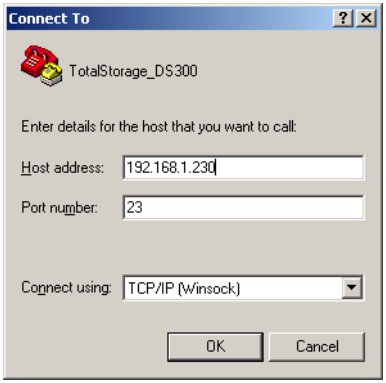

Figure 4-163 Hyperterminal — Connect to 192.168.1.230

4. Logon as administrator by entering **administrator**. Enter the administrator password and press enter. When a password for the operator account was defined than this password needs to be entered fist to logon as operator. Verify the current firmware version and image number with **image info**. Image 1 is version 6.21 and image 2 is version 6.04. The image which will be used for booting is image 1, version 6.21. Current image is also 6.21. The firmware update will overwrite image 2, version 6.04.

```
----[ IBM DS300: Command Line Interface ]---------------------------------

               Copyright:  IBM (c)2003-2004

                Firmware:  IBM TotalStorage DS Series (Build 0750)
                           [27/07/2005][6.21]

  Peer controller status:  Ready


           Typing 'help':  provides an overview of the commands available to
                           configure the controller.

                 TAB key:  use the TAB key as a means of getting assistance
                           in completing any commands that you issue.



DS300-1[A]> administrator

Administrator level password: ********

DS300-1[A]# image info

----[ System firmware images ]-------------------------------------------

           Image [1]:  [IBM TotalStorage DS Series (Build 0750)]
                       [27/07/2005][6.21]

           Image [2]:  [IBM TotalStorage DS Series (Build 0710)]
                       [23/06/2005][6.04]

       Current image:  [1]  [IBM TotalStorage DS Series (Build 0750)]

     Next boot image:  [1]  [IBM TotalStorage DS Series (Build 0750)]

DS300-1[A]#
```

Figure 4-164 Command line interface — Logon and boot image information

5. Start the firmware update with the command **image upload 1kxmodem**. Confirm with **yes** that the inactive firmware page will be overwritten. When the first C is shown as in Figure 4-165 on page 183 click menu Transfer → Send File from Hyperterminal. See Figure 4-166 on page 183

Note: Image upload `ftp://192.168.1.253/__DELETE/0750.upgrade` can be used to perform a firmware file upload via ftp. For HTTP replace the protocol ftp with http. A ftp or http server must be accessibly by the controller of the TotalStorage DS300 or DS400 and the firmware file must be stored there.

```
DS300-1[A]# image upload 1kxmodem
Overwrite [IBM TotalStorage DS Series (Build 0710)]? [No]: yes
CCCCC
```

Figure 4-165 Command line interface — 1K Xmodem firmware upload

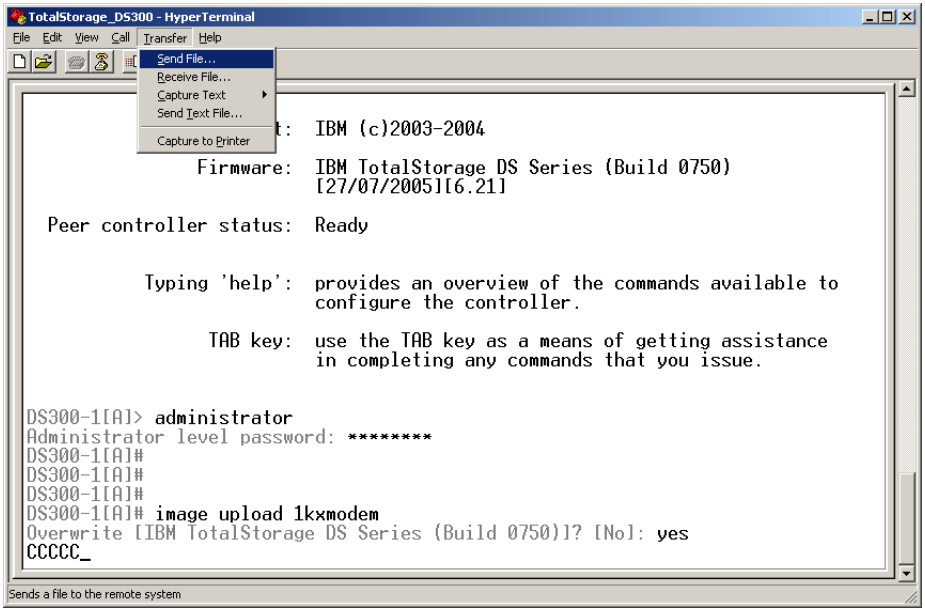

Figure 4-166 Hyperterminal — Transfer → Send file

6. Select the firmware file and the protocol 1K Xmodem. Click Send.

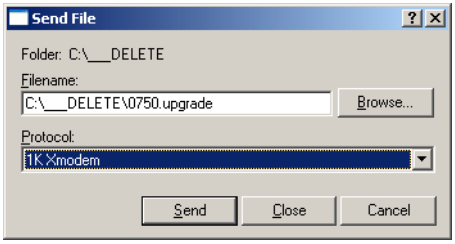

Figure 4-167 Hyperterminal — Send file

7. The firmware file will be transferred to the controller. This dialog will close after successful transmission or because to many errors.

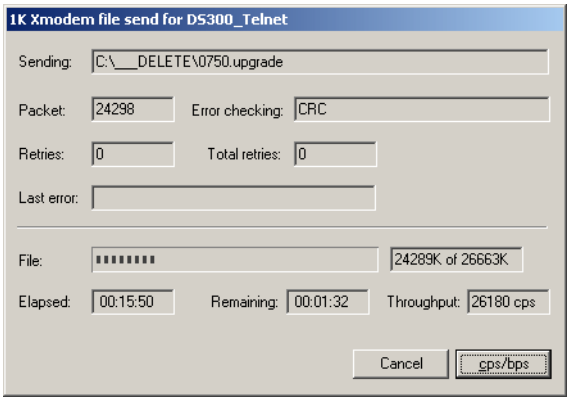

Figure 4-168 Hyperterminal — transfer firmware data

8. After the firmware file was successful transferred the command line interface informs about the next steps. The firmware is written to the controller.

```
 27302294 bytes read...
RedBoot already at version 05.05.05, skipping...
Writing kernel
Writing filesystem
  27302294 bytes read...
Image successfully uploaded.

Use the 'system reboot' command to run this new firmware image



DS300-1[A]#
```

Figure 4-169 Command line interface — successful image flash

9. When the firmware was successfully flashed and a alternate controller is available, repeat the previous steps with controller B.
10. Verify the new boot image with **boot info**. The second image was overwritten with the newer firmware version 6.21. The next boot image was set to 2.

```
DS300-1[A]# image info

----[ System firmware images ]------------------------------------------

           Image [1]:  [IBM TotalStorage DS Series (Build 0750)]
                       [27/07/2005][6.21]

           Image [2]:  [IBM TotalStorage DS Series (Build 0750)]
                       [27/07/2005][6.21]

       Current image:  [1]  [IBM TotalStorage DS Series (Build 0750)]

     Next boot image:  [2]  [IBM TotalStorage DS Series (Build 0750)]

DS300-1[A]#
```

Figure 4-170 Command line interface — boot image information

11. Enter **system reboot** to reboot the TotalStorage DS300 or DS400. Both controllers will reboot. The flashed firmware image will be used for booting after this reboot.

```
DS300-1[A]# system reboot

Reboot the system? [No]: yes


The system is about to reboot
```

Figure 4-171 Command line interface — Reboot TotalStorage DS300 or DS400

The next section explains how to perform the same firmware update wit the IBM ServeRAID Manager console.

4.12.2 Update DS300 and DS400 firmware with ServeRAID Manager

With ServeRAID Manager, multiple enclosures can be updated at the same time. This is a new feature introduced with the ServeRAID Manager 8.20 the first time.

The following steps give a detailed explanation of the firmware upgrade with the ServeRAID Manager console.

1. Start the ServeRAID Manager console and connect to all Management Stations and agents. Right-click **Networked storage** and select **Update controller images** from the context menu.

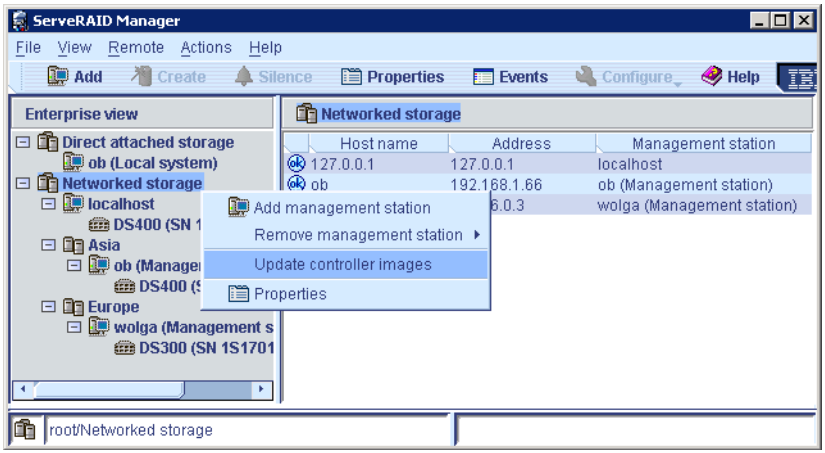

Figure 4-172 ServeRAID Manager console — Update controller images

2. The IBM ROM update wizard starts and introduces the process of upgrading the firmware. Click **Next**.

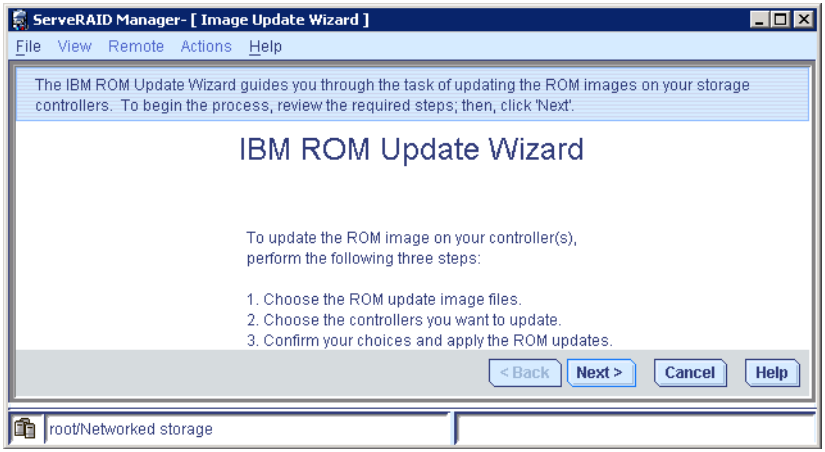

Figure 4-173 IBM ROM Update Wizard

3. Enter the path and the filename of the upgrade or use the add button to select the upgrade file. Click **Next**.

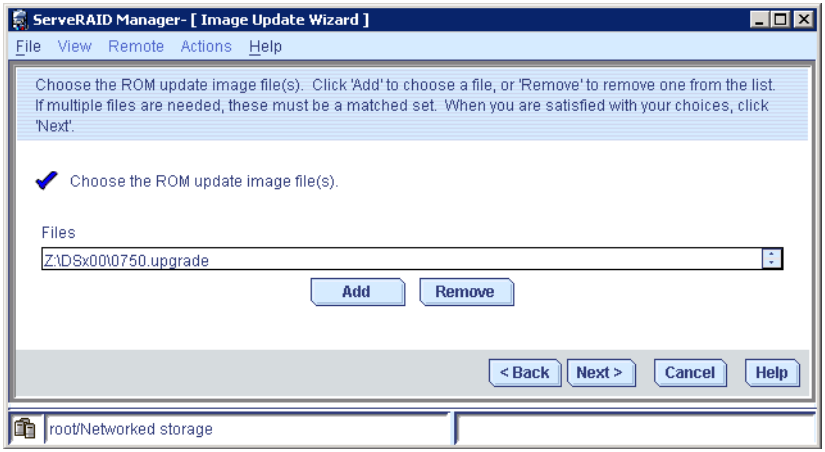

Figure 4-174 Select upgrade file

4. Deselect all those enclosures which should not be updated at this time. All enclosures which are on the same level are greyed out and cannot be selected or deselected. Click **Next**.

image upload
ftp://192.168.1.1/0750.upgrade

or:

image upload
http://192.168.1.1/0750.upgrade

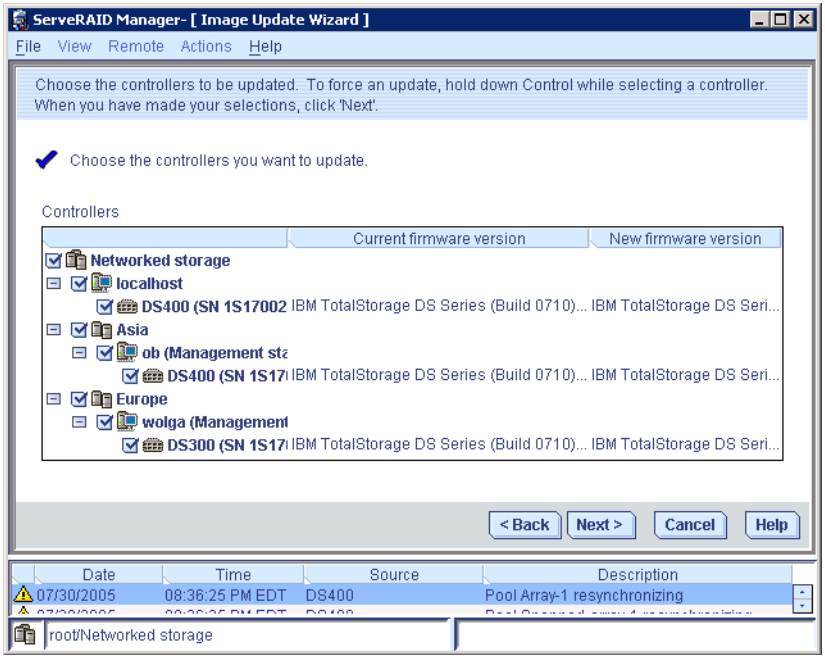

Figure 4-175 Select enclosures for update

5. Confirm the selected enclosures for update by clicking **Apply**.

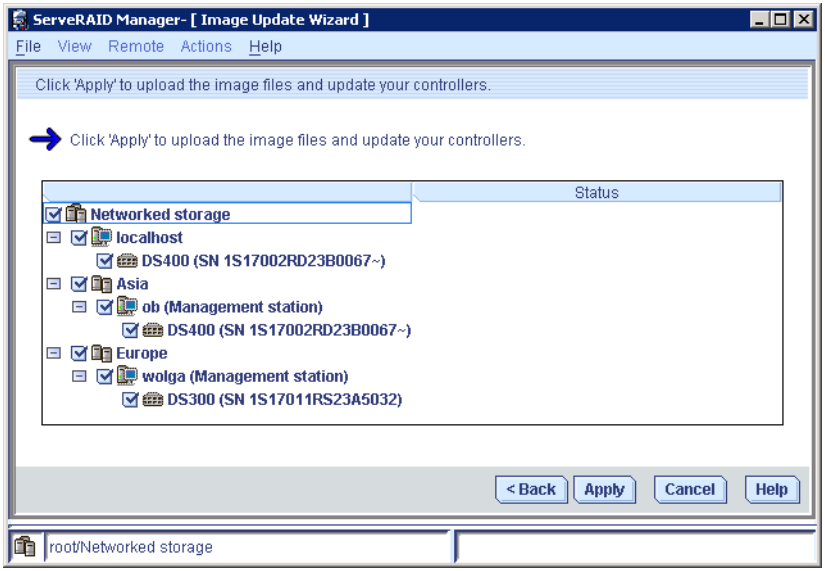

Figure 4-176 Confirm selected enclosures for update

6. Last chance to cancel the update. Click **Yes** to start the update on the selected enclosures.

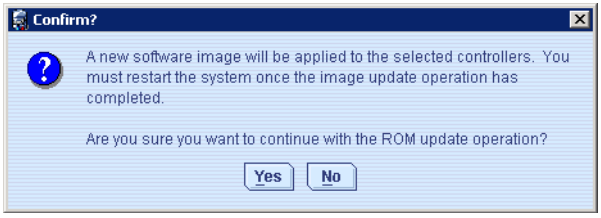

Figure 4-177 Confirmation question

7. First step of the upgrade is the upload of the image file to all selected Management Stations.

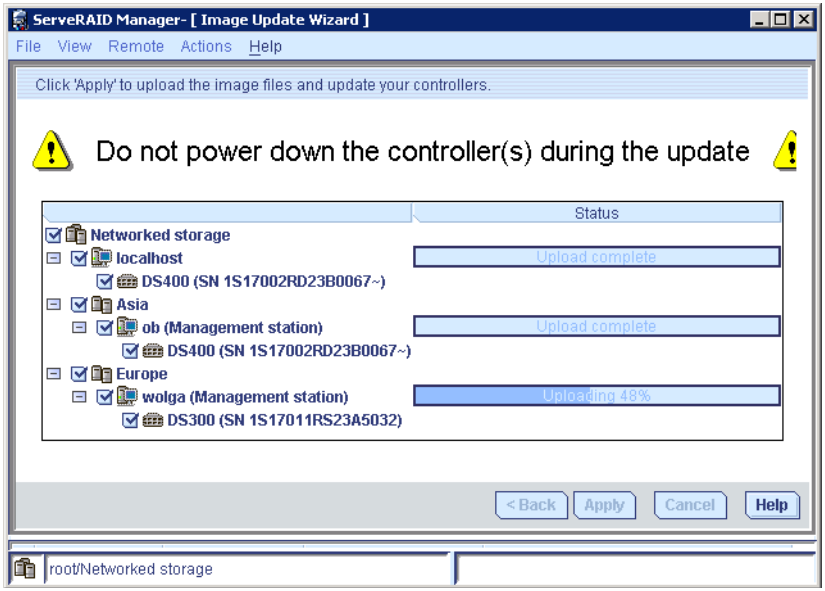

Figure 4-178 Upload firmware to Management Station

8. Second step is the download of the image file to the controller of the selected enclosures. After verifying the uploaded file the controller will be flashed.

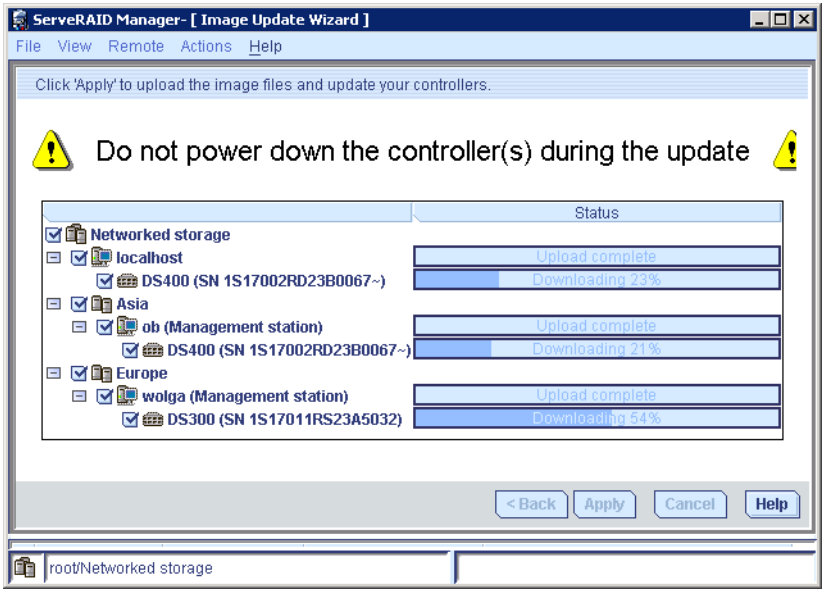

Figure 4-179 Download firmware to enclosures

9. A message will tell when all enclosures are updated and a reboot is required. Perform a reboot during a period of inactivity or off-hour to activate the new firmware code.

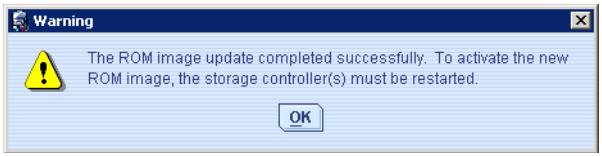

Figure 4-180 Update finished, reboot required

4.13 IBM ServeRAID Manager removal on Microsoft Windows

To uninstall a already installed IBM ServeRAID Manager perform the following steps:

1. Go to **Start** → **Settings** → **Control Panel** → **Add or Remove Programs**.
2. Select **ServeRAID Manager** in **Change or Remove Programs**.
3. Click the **Remove** button.

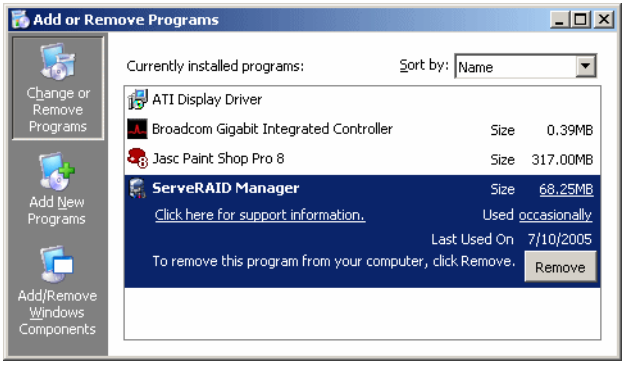

Figure 4-181 Add and Remove Programs

4. Click **Yes** to confirm that you want to remove the IBM ServeRAID Manager.

Figure 4-182 Question

5. Click **Yes** to reboot the system.

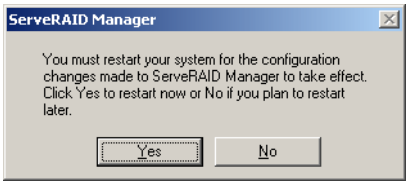

Figure 4-183 Reboot Confirmation

Note: Older versions of the ServeRAID Manager require a reboot. Newer versions do not. Reboot the system when a reboot is requested.

4.14 IBM ServeRAID Manager removal on Linux

To remove ServeRAID Manager on a Linux server it is only necessary to execute the **`rpm -e RaidMan`** command as shown in Figure 4-184.

```
danube:~ # rpm -e RaidMan
stopping IBM ServeRAID Manager agent ...

raid_agent               0:off  1:off  2:off  3:off  4:off  5:off  6:off

Uninstall completed successfully.
danube:~ #
```

Figure 4-184 ServeRAID Manager removal using the rpm command

4.15 Basic information gathering in case of trouble

To get help from the IBM technical support center, certain information about the current status of your IBM Total Storage DS300 or DS400 is required. This

section will help you to collect information about the current configuration and the current status. The basic information gathering consists of the following steps:

4.15.1, “Gather information about the system with eGatherer and Dynamic System Analysis” on page 192

4.15.2, “LED status of the Total Storage DS300 and DS400 and EXP400” on page 196

4.15.3, “Network cabling diagram” on page 196

4.15.4, “IBM ServeRAID Manager support archive” on page 196

4.15.5, “IBM ServeRAID Manager log files” on page 203

4.15.6, “Log files of the iSCSI SANsurfer Management software” on page 204

4.15.7, “Log files of the IBM FAStT Management Suite Java (MSJ)” on page 204

4.15.1 Gather information about the system with eGatherer and Dynamic System Analysis

eGatherer and Dynamic System Analysis (DSA) are utilities to gather information about the host operating system. DSA is supported only on Microsoft Windows, but eGatherer is supported as well on Linux. Both utilities can be freely downloaded from the IBM Support site.

eGatherer

Download the eGatherer from the listed URL below and store the file on the host attached to a IBM TotalStorage DS300 or IBM TotalStorage DS400 and the system used as Management Station.

Link for eGatherer:

`http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-4R5VKC`

Execute eGatherer with administrative rights direct on the server. In Microsoft Windows do not use a Terminal Server Session to execute the eGatherer utility.

A file *.eg2 will be created in the same directory from which you start eGatherer.

Dynamic System Analysis

Dynamic System Analysis (DSA) is available in two different versions. First version, the DSA Portable Edition, can be executed on a system without installing the utility. This is useful for a single time usage. The second version, the DSA

Installable Edition, needs to be installed first before it can be used. A start menu entry is created to run the utility. For more details about DSA and the download see the URL below. This book assumes that the DSA Portable Edition is used.

Link of DSA:

http://www-1.ibm.com/servers/eserver/xseries/systems_management/dsa.html

DSA can be executed with options. The order of the options are important. First are all options for the installer specified. To use options for the DSA itself, specify /a followed by those options. A few important options are:

/s accept the license, option is for the installer

/a option for the installer to pass all remaining options to DSA

/v DSA option to create HTML output files

More about the options are documented in the readme file of DSA which can be downloaded together with the appropriate DSA edition.

Download the DSA portable edition and store the file on all systems attached to the IBM TotalStorage DS300 and IBM TotalStorage DS300. Perform the following steps to gather information about the system:

1. Click Start → Run. In the Run window select the DSA executable.

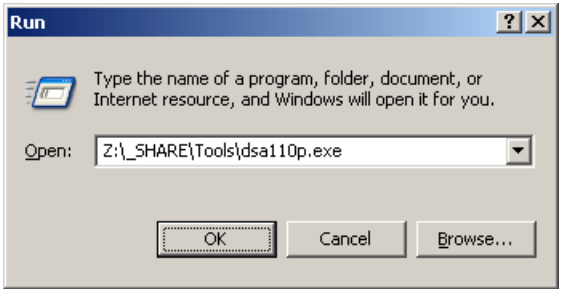

Figure 4-185 Run — DSA installer

2. Add the option /a /v behind the executable. then click **OK**.

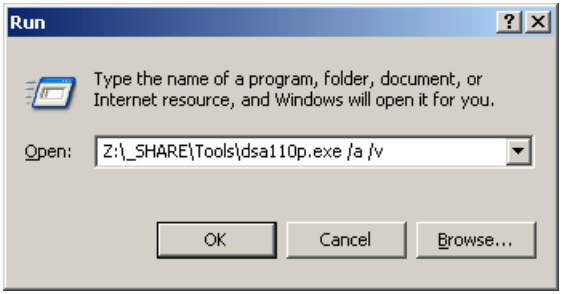

Figure 4-186 Run — DSA installer with options

3. The InstallShield Wizard opens. Click **Next**.

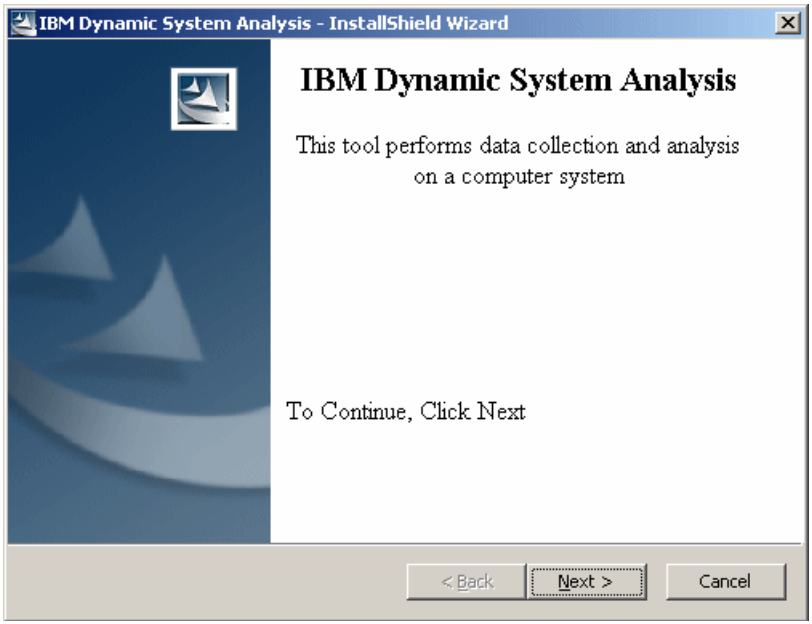

Figure 4-187 DSA InstallShield Wizard

4. Accept the DSA license agreement and click **Next**.

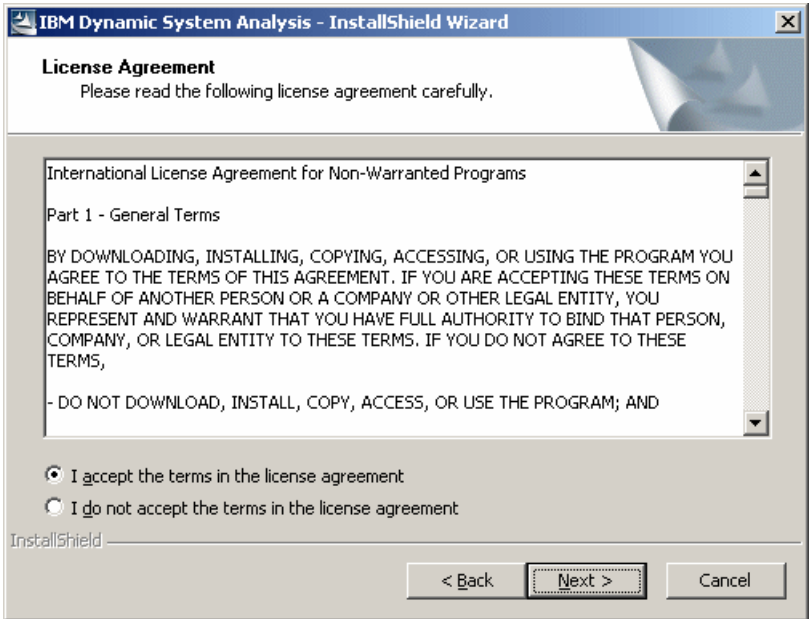

Figure 4-188 DSA license agreement

5. An command line window opens and DSA is executed. This windows will close as soon as DSA execution is finished.

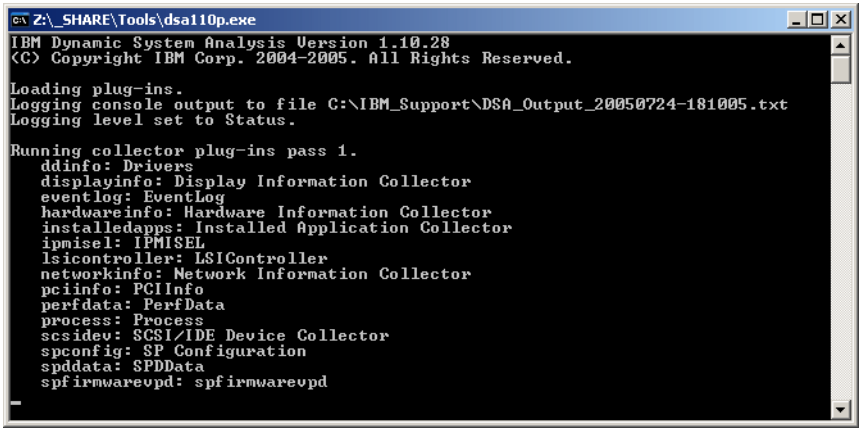

Figure 4-189 DSA

6. By default DSA saves the output in C:\IBM_Support. The output is a compressed xml file. in a subdirectory with the same name than the *.xml.gz file are the html files with all gathered information as well. Open the index.html file to see the gathered information.

4.15.2 LED status of the Total Storage DS300 and DS400 and EXP400

The TotalStorage DS300 and DS400 has on front and rear site LEDs to help identifying problems. It is important to know the current status of those LEDs. 3.3.4, “DS300 controls and LEDs” on page 34 describes the available LEDs for a TotalStorage DS300. 3.4.3, “DS400 controls and LEDs” on page 45 describes the LEDs for a TotalStorage DS400.

A description of error LEDs and actions to solve problems are described in the hardware maintenance manual of the TotalStorage DS300 or DS400. The documents can be downloaded from:

http://www.ibm.com/support/docview.wss?rs=1109&uid=psg1MIGR-58619

4.15.3 Network cabling diagram

Make a diagram of the current cabling of the hosts, TotalStorage DS300 or DS400 and all switches routers involved in this setup. Some samples are shown in Chapter 6, “TotalStorage DS300 sample configurations” on page 211 and Chapter 8, “TotalStorage DS400 example configurations” on page 433.

4.15.4 IBM ServeRAID Manager support archive

The Support Archive contains much of the important information about the storage product. Part of the array is the controller diagnostic file, printable configuration, the Management Station event log and the ServeRAID Manager error log. It can be generated for local attached storage as well as for network attached storage.

1. Create the support archive by right-clicking on the controller or enclosure in the **Enterprise view** and selecting **Save support archive** from the context menu.

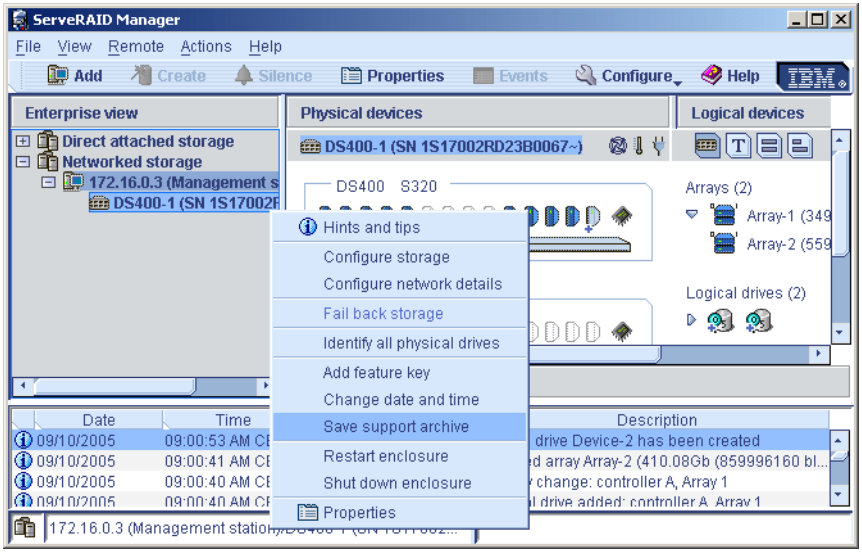

Figure 4-190 Save support archive

2. Save the support archive local. by default the host name of the TotalStorage DS300 or DS400 is used as file name.

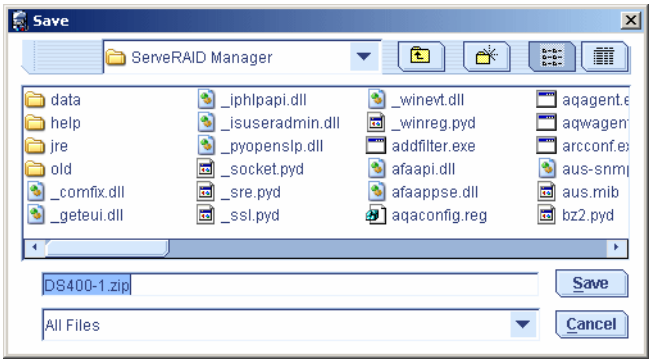

Figure 4-191 Save file

3. When problems exist with creating the support archive, information can be retrieved manually. Perform the following steps to get the information when the support archive cannot be generated:

   a. Save IBM ServeRAID Manager events

   Collect the console event log by clicking on the **Events** button in the toolbar. Save the events into a file. By default the file name is events.txt and it is saved in the installation directory of the IBM ServeRAID Manager.

This log file contains events from all Management Stations added to the ServeRAID Manager.

To save the events for a particular Management Station, select the Management Station and use the menu actions and save the events from there. Please note that by default the filename and location are the same.

b. Save printable configuration information

Save the printable configuration information for every Management Station. In the figure below there are three Management Stations added. This action has to be performed for every station. The printable configuration is stored in a file with the name RaidExt1.log in a subdirectory with the name of the Management Station in the IBM ServeRAID Manager installation directory.

c. Save controller diagnostic file

Use Hyperterminal to connect via TCOP/IP (Winsock) to Controller A.

> **Note:** A serial cable is required for this connection. An IBM Service representative must be onsite when this procedure is performed.

- Enter administrative mode.
- Enter **diagnostic dump 1xmodem** and press enter

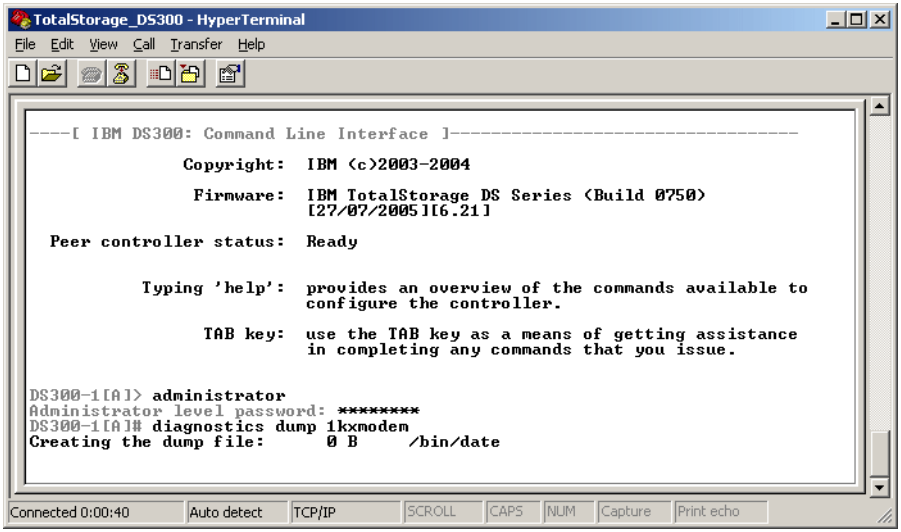

Figure 4-192 Hyperterminal — Logon and diagnostic dump

- After the message “Creating the dump file: 0kB” appears click the Hyperterminal menu **Transfer** → **Receive File**

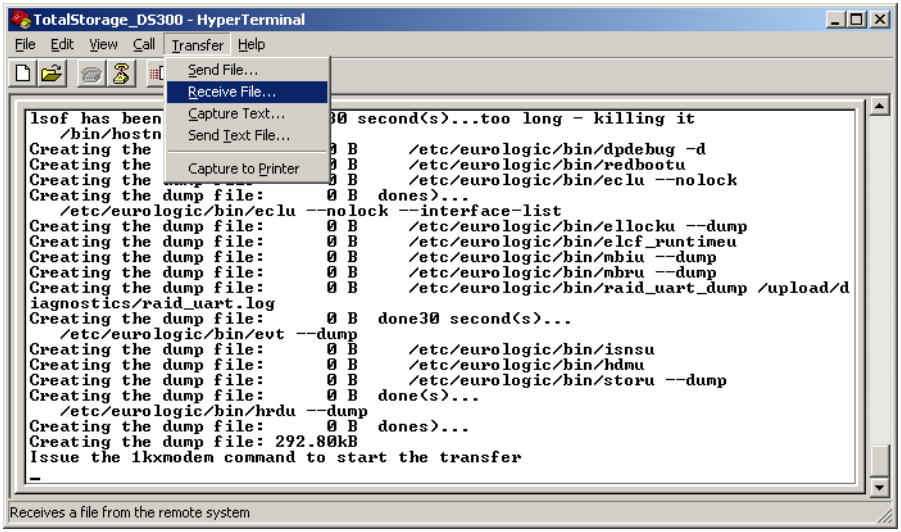

Figure 4-193 Hyperterminal — Menu transfer → Receive file

- Select the directory where the received file will be saved and choose the **1K Xmodem** protocol. Click **Receive**.

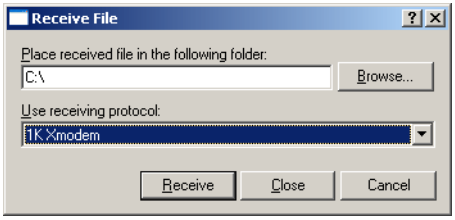

Figure 4-194 Hyperterminal — Chose protocol and location

- Enter the file name under which the received file will be saved. Click **OK**.

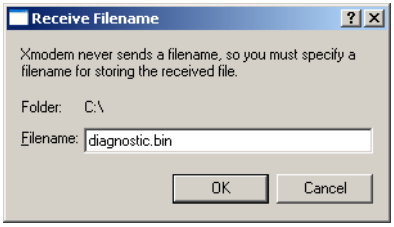

Figure 4-195 Hyperterminal — Enter file name

- The file will be downloaded from the controller and saved in the specified location. This window will close after the download is done.

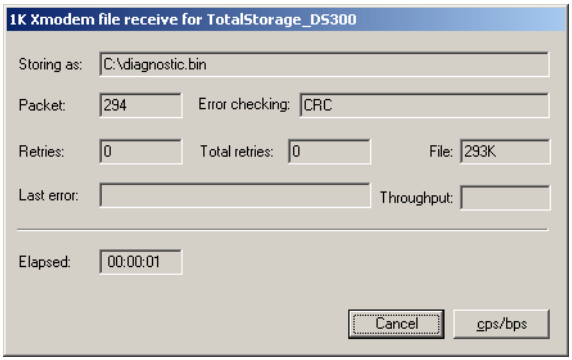

Figure 4-196 Hyperterminal — Receive file

- The command line interface informs you of the successful dump transmission.

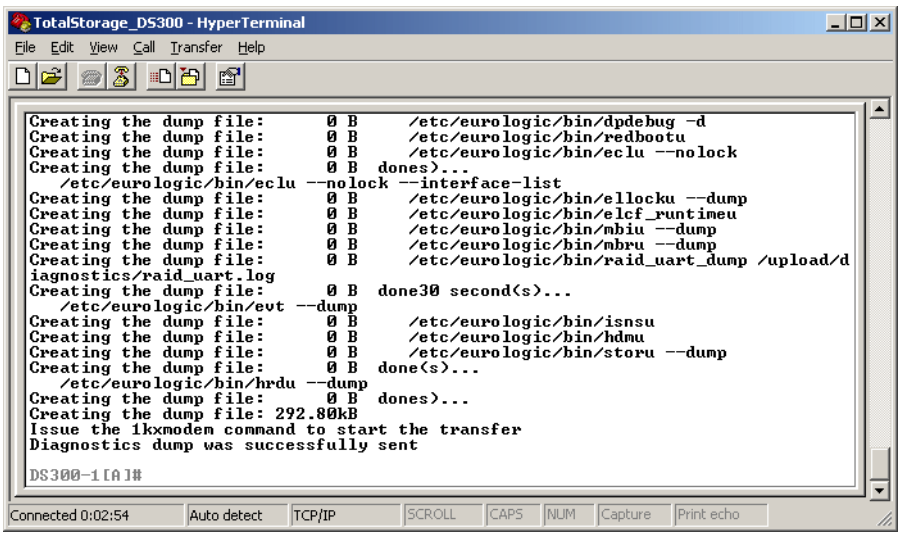

Figure 4-197 Hyperterminal — Dump successful sent

- Repeat this actions on the Controller B in case you have a dual controller TotalStorage DS300 or DS400.

d. CLI command output

Use Windows HyperTerminal or PUTTY or any other Telnet client which is capable of capturing the screen output into a file.

i. Use Hyperterminal to connect via TCOP/IP (Winsock) to Controller A.

ii. Logon as administrator.

iii. Click the Hyperterminal menu **Transfer** → **Capture Text**.

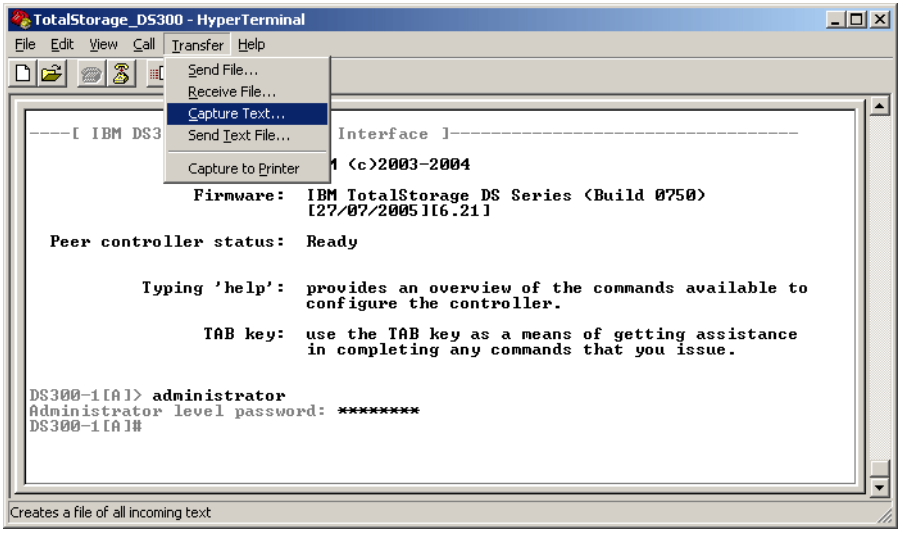

Figure 4-198 Hyperterminal — Menu Transfer → Capture text

iv. Enter path and file name to save the captured information. Click **Start**.

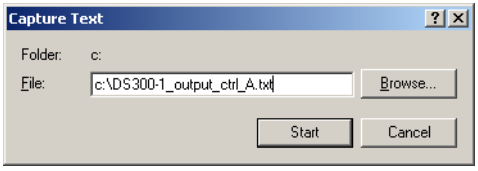

Figure 4-199 Hyperterminal — Enter path and file name

v. Run the command **show all**.

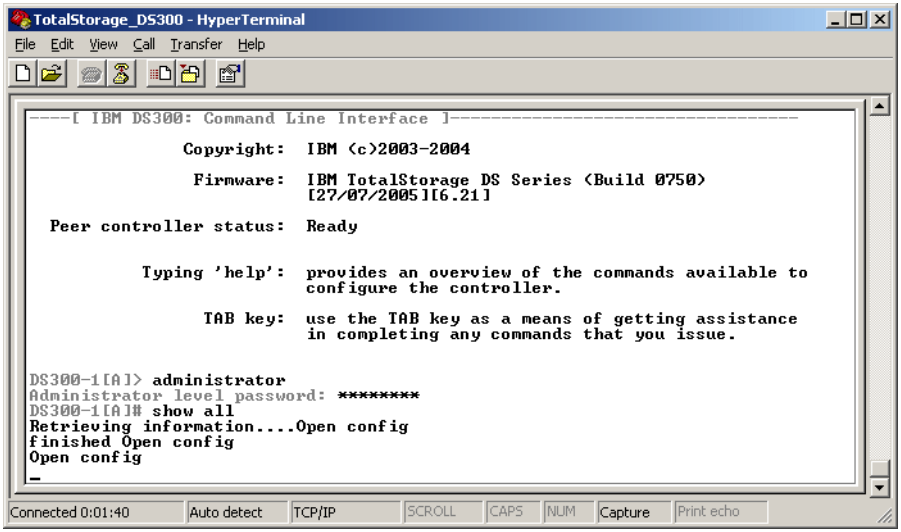

Figure 4-200 Hyperterminal — show all

vi. The output will be several pages long. When a page is filled, the output stops. The last line shows the amount of output. Press **Enter** to scroll to the next page. When the end is reached, Press **q** to quit the show all command.

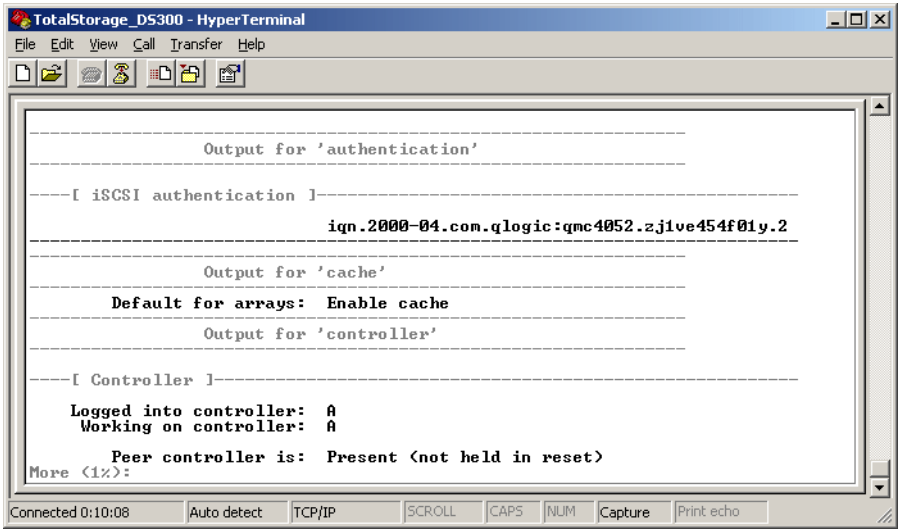

Figure 4-201 Hyperterminal — show all output

vii. When the output cannot be collected at once, try the separate sub commands of show. Run one command after the other. for example show array or show system. This is the list of command for a TotalStorage DS300. Enter show and press TAB to get a list of sub commands.

```
DS300-1[A]# show
all             fru             logical         snmp
array           image           network         system
authentication  initiator       route
cache           interface       scope
device          isns            service
events          license         snapshot
DS300-1[A]# show
```

Figure 4-202 Sub commands of the show command

viii. Run the command **events report all**. See the command show all for details. Please note that this can be a very long output.

ix. Repeat these commands on the alternate controller.

e. Save Management Station log files

The Management Station maintains a own set of log files. This log files are located in the temp directory of the windows installation directory. Windows installation directory is mostly `Windows\Temp` or `Winnt\Temp`. The log files are called:

- mgmtservice.log
- mgmtservice.log.old
- miniwinagent.log

4.15.5 IBM ServeRAID Manager log files

The IBM ServeRAID Manager maintains a few log files:

RaidEvt.log	Contains the information reported in the ServeRAID Manager event viewer for all local and remote systems.
RaidEvtA.log	Contains the information reported in the ServeRAID Manager event viewer for the local system.
RaidNot.log	Contains the information reported in the Notification Manager event viewer.
RaidSec.log	Contains the information reported in the Security Manager event viewer.

RaidErr.log	Contains Java™ messages generated by the ServeRAID Manager.
RaidErrA.log	Contains Java messages generated by the ServeRAID Manager agent.

Information written to these files are appended to the existing files to maintain a history. However, if any of the files reach a size of 200,000 bytes, the file is copied to a new file with the extension .OLD and the original (that is the LOG file) is deleted and recreated. If a .OLD file already exists, the existing .OLD file is destroyed.

Not all of those files might exist. IBM ServeRAID Manager creates the log files as soon as the first event for a specific log file appears.

4.15.6 Log files of the iSCSI SANsurfer Management software

In case of problems with a DS300, you may use the iSCSI SANsurfer Management software. This software saves a few files in the installation directory which are helpful in case of trouble.

In Microsoft Windows those files are:

```
C:\Program Files\QLogic Coproration\SANsurfer\iscsi_alarms.txt
C:\Program Files\QLogic Coproration\SANsurfer\iscsi_hists.txt
C:\Program Files\QLogic Coproration\SANsurfer\alarms.txt
C:\Program Files\QLogic Coproration\SANsurfer\events.txt
C:\Program Files\QLogic Coproration\SANsurfer\hosts.txt
```

4.15.7 Log files of the IBM FAStT Management Suite Java (MSJ)

In case of problems with a TotalStorage DS400, you may use the IBM FAStT Management Suite Java. This software saves a few files in the installation directory which are helpful in case of trouble. The newer utility is the SANsurfer utility. This tool has the capability of managing fibre channel hot bus adapter as well.

In Microsoft Windows those files are:

```
C:\Program Files\IBM FAStT MSJ\alarms.txt
C:\Program Files\IBM FAStT MSJ\events.txt
C:\Program Files\IBM FAStT MSJ\hosts.txt
```

5

Performance and configuration considerations

This chapter provides some notes about performance considerations which may be useful when configuring IBM DS300/DS400 subsystems.

5.1 Performance considerations

SAN storage products such as the IBM TotalStorage DS300 and DS400 subsystems are designed to support server workloads. Since servers are designed to handle multiple tasks in parallel, a DS300/DS400 is optimized to perform well with multi-threaded workloads. Using a single-threaded utility such as Copy or Xcopy does not take full advantage of the DS300/DS400 capabilities, and does not show how a server that can support multi-threaded workloads will really perform.

Similarly, dividing drives up into multiple smaller drive arrays, instead of using just one large drive array, improves DS300/DS400 performance.

TotalStorage DS300 and DS400 performance improves when both controllers are utilized. During performance evaluation, some evaluate the performance by applying work to only one of the two controllers. In general, performance can be improved by almost 2 times when work is applied to arrays on both controllers instead of just one controller.

TotalStorage DS300 and DS400 performance can be improved by dividing the drives into multiple smaller drive arrays, instead of using only one large drive array. When using two controllers, performance can be improved by utilizing two arrays per controller instead of one array per controller.

5.2 Configuration considerations

- An array can have up to 16 drives on a DS400 when RAID level 0,1 or 5 is used. On a DS300, up to 14 drives are supported.
- A spanned array on a DS400 can have a total of 40 drives.
- Logical drive settings, such as RAID level and stripe size, are defined for an array, and are the same for all logical drives in the same array.
- The maximum array size is 2TB.
- Up to 512 logical drives can be configured in a single array. Only 64 of those logical drives can be configured by using the command line interface of the TotalStorage DS300 or DS400.

network ntp
172.16.0.9

- A network time protocol server (NTP) helps to have correct time on the TotalStorage DS300 or DS400. All time stamps in the log files and messages are correct and no manual intervention is required to keep them synchronized. A time server can only be configured by using the command line interface.

- Only one default gateway can be defined. Additional routes can be set only from CLI.

With RAID technology, data is striped across a group of physical drives. This data distribution scheme complements the way the operating system requests data.

The granularity at which data is stored on one physical drive of the logical drive, before subsequent data is stored on the next physical drive of the logical drive, is called the stripe-unit size. The collection of stripe-units, from the first drive of the logical drive to the last drive of the logical drive, is called a stripe.

You can set the stripe-unit size to 16, 32, 64, 128, 256 (the default), 512 or 1024 KB. You can maximize the performance of your TotalStorage DS300 or DS400 by setting the stripe-unit size to a value that is close to the size of the system I/O requests. For example, performance in transaction-based environments, which typically involve large blocks of data, might be optimal when the stripe-unit size is set to 64 KB or 128 KB. However, performance in file and print environments, which typically involve multiple small blocks of data, might be optimal when the stripe-unit size is set to 16 KB.

Part 2

TotalStorage DS300

In Part 2 we discuss planning, installation, and configuration considerations for the IBM TotalStorage DS300 subsystem.

6

TotalStorage DS300 sample configurations

This chapter gives an overview of the sample DS300 SAN configurations that are used in later chapters to explain, in detail, TotalStorage DS300 installation and configuration.

This chapter contains these topics:

6.1 DS300 single controller sample configurations

In this section we describe sample configurations for the DS300 single controller models.

6.1.1 DS300 single controller switched or routed attachment (1701-1RL)

The DS300 model 1701-1RL has one GBit Ethernet port which is used for management and iSCSI traffic. All DS300 models have in common that they support one default gateway. This default gateway is normally used for the management network. As the DS300 model 1701-1RL uses one interface for management and iSCSI traffic, the default gateway can also be used for iSCSI traffic. For the other models, it is necessary to configure additional routes to enable routed access to the iSCSI ports.

Figure 6-1 shows a routed attachment of a single server. Additional servers can also be connected to the gigabit switch, or somewhere else in the network behind different routers.

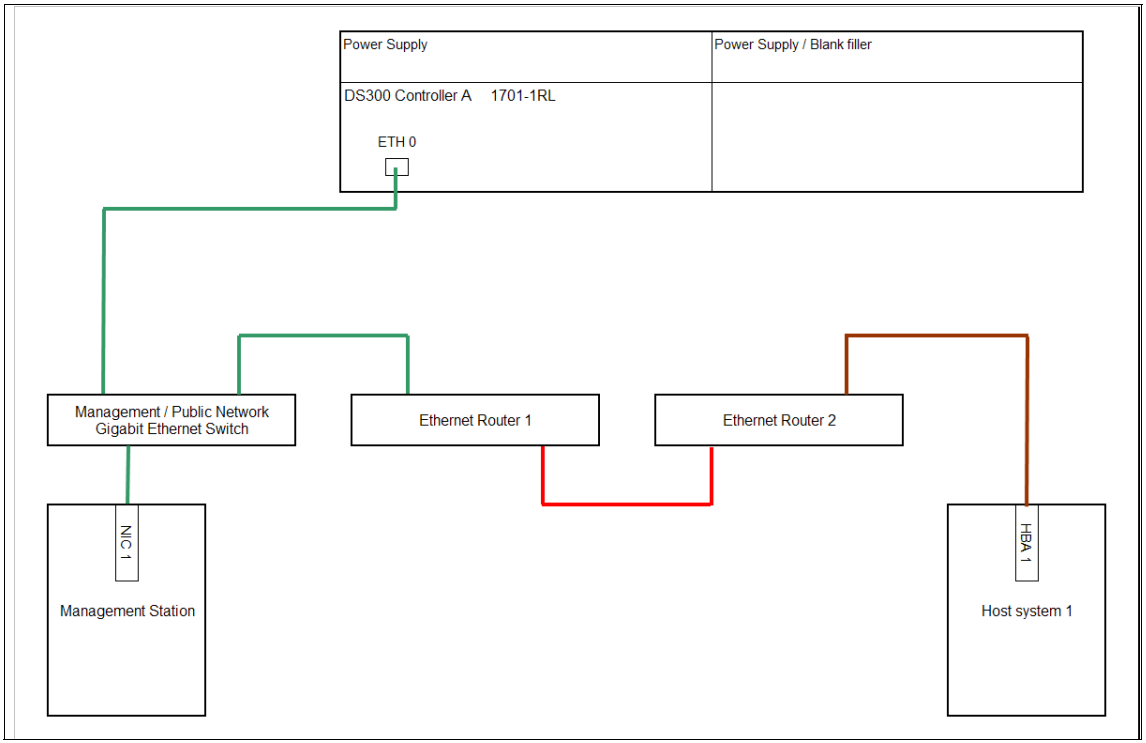

Figure 6-1 DS300 single controller routed attachment (model 1701-1RL)

6.1.2 DS300 single controller switched single path attachment

A server in this configuration has one initiator. There is only one path from each server to the DS300, but these paths go through a switch. Hence it is possible to also connect an Internet Storage Name Server (iSNS) to the networks. It is also possible to connect additional servers to the switches which can then connect to the DS300.

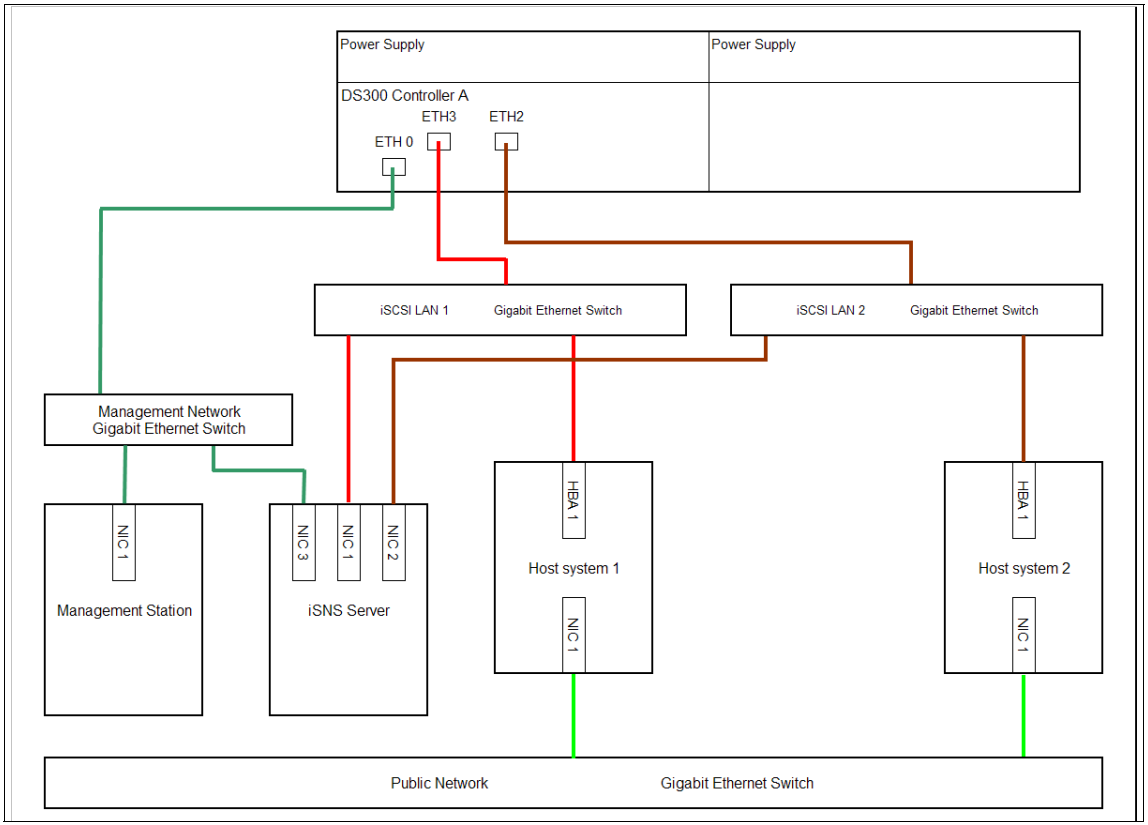

Figure 6-2 DS300 single controller switched single path attachment

Table 6-1 gives an overview of the usage of different types of initiators in this configuration.

Table 6-1 DS300 single controller switched single path attachment initiator remarks

OS type	Initiator[a]
Windows 2000 and 2003	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter
Windows 2000 and 2003	MS software initiator 1.06
Windows 2000 and 2003	MS software initiator 2.0
Linux	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter
Red Hat 3	iSCSI software initiator V3.6.3, V3.6.4

OS type	Initiator[a]
SLES 8	iSCSI software initiator V3.6.3, V3.6.4
Red Hat 4	embedded iSCSI software initiator
SLES 9	embedded iSCSI software initiator

a. Check driver readme files for detailed version numbers and service pack levels.

6.1.3 DS300 single controller switched dual path attachment

A server in this configuration has two interfaces (two hardware initiators or one software initiator using two network interfaces). Each of the interfaces is connected to a different subnet. Depending on the type of the initiator, it is possible to use multipathing via operating system specific multipathing mechanisms or multiple connections per session (MCS). Only with the software initiators is multipathing possible. Even when multipathing is possible, it is not always supported. Check support pages and firmware readme files for current support status. With hardware initiators like IBM iSCSI Server Adapter or the Qlogic iSCSI Expansion Card for IBM IBM @server BladeCenter, multipathing is not possible because at the time this redbook was written, no multipathing driver was available.

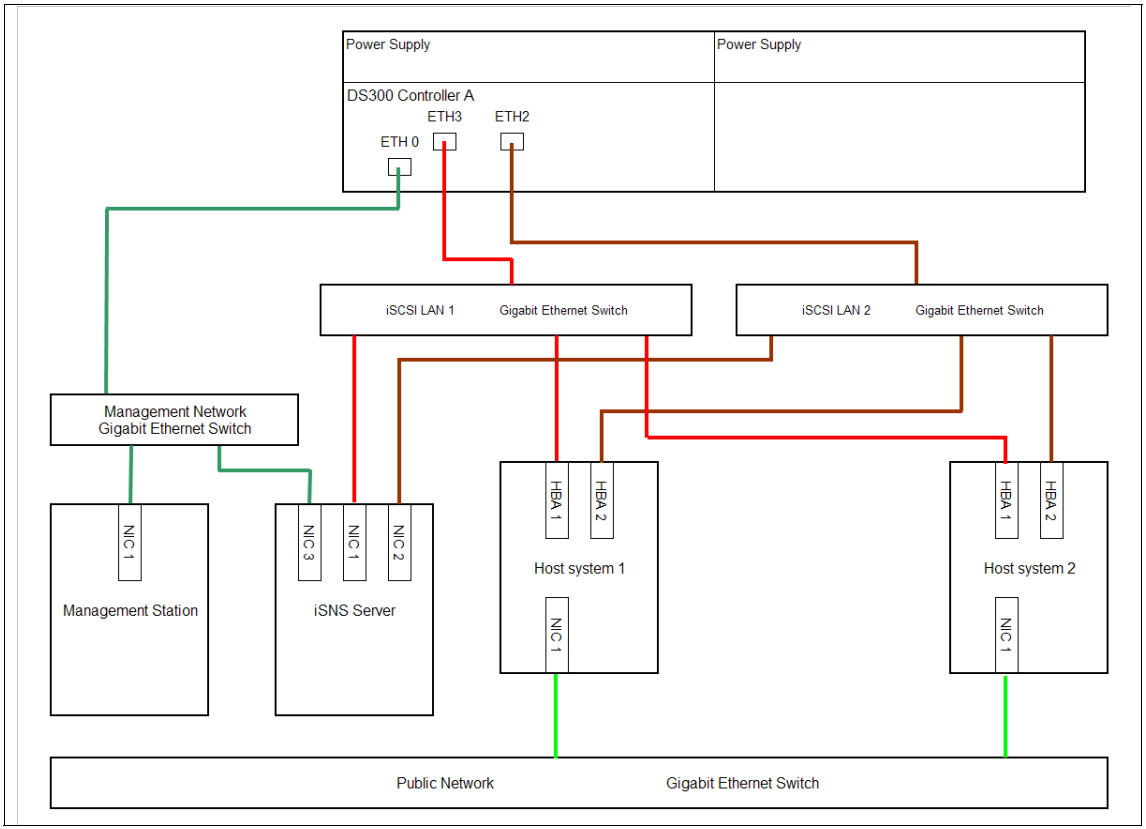

Figure 6-3 DS300 single controller switched dual path attachment

Table 6-2 gives an overview of the usage of different types of initiators in this configuration.

Table 6-2 DS300 single controller switched dual path attachment initiator remarks

OS type	Initiator[a]	Remarks
Windows 2000 and 2003	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter	spread logical drives across the HBAs, no multipathing
Windows 2000 and 2003	MS software initiator 1.06	spread logical drives across the HBAs, no multipathing
Windows 2000 and 2003	MS software initiator 2.0	spread logical drives across the HBAs, no multipathing

OS type	Initiator[a]	Remarks
Linux	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter	spread logical drives across the HBAs, no multipathing
Red Hat 3	iSCSI software initiator V3.6.3, V3.6.4	multipathing using device mapper or MCS possible, but not supported by IBM
SLES 8	iSCSI software initiator V3.6.3, V3.6.4	multipathing using device mapper or MCS possible, but not supported by IBM
Red Hat 4	embedded iSCSI software initiator	multipathing using device mapper or MCS possible, but not supported by IBM
SLES 9	embedded iSCSI software initiator	multipathing using device mapper or MCS possible, but not supported by IBM

a. Check driver readme files for detailed version numbers and service pack levels.

6.2 DS300 dual controller sample configurations

In this section we describe sample configurations for the DS300 dual controller models.

6.2.1 DS300 dual controller single switch attachment

A server in this configuration has two interfaces (two hardware initiators or one software initiator using two network interfaces teamed into one). Each of the interfaces is connected to the same subnet.

With software initiators, outbound traffic is always sent out over the first port of multiple ethernet NICs in the same subnet. However, inbound traffic can be received over either port. Some operating systems are not able to send outbound responses over the interface where inbound traffic was received. This can be addressed by using teaming software. Ethernet interfaces for software initiators must be teamed together using teaming software in Windows or ethernet channel bonding in Linux. You must insure that the iSCSI session is established over the virtual adapter and not the physical adapters.

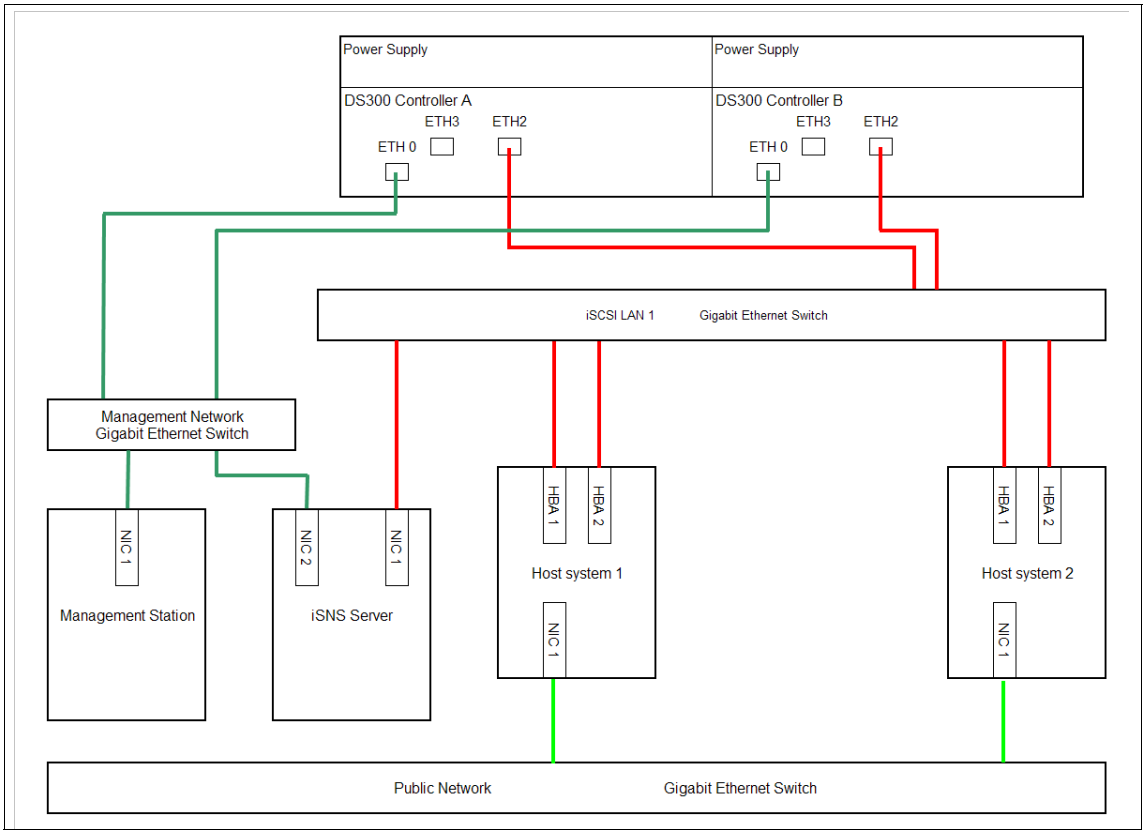

Figure 6-4 DS300 dual controller switched dual path attachment with one switch

Table 6-4 on page 220 gives an overview of the usage of different types of initiators in this configuration.

Table 6-3 DS300 dual controller switched dual path attachment with one switch remarks

OS type[a]	Initiator	Remarks
Windows 2000 and 2003	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter	place HBAs in same subnet, spread logical drives across the HBAs, no multipathing
Windows 2000 and 2003	MS software initiator 1.06	use NIC teaming software (failover mode)
Windows 2000 and 2003	MS software initiator 2.0	use NIC teaming software (failover mode)

OS type[a]	Initiator	Remarks
Linux	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter	place HBAs in same subnet, spread logical drives across the HBAs, no multipathing
Red Hat 3	iSCSI software initiator V3.6.3, V3.6.4	use ethernet channel bonding (failover mode)
SLES 8	iSCSI software initiator V3.6.3, V3.6.4	use ethernet channel bonding (failover mode)
Red Hat 4	embedded iSCSI software initiator	use ethernet channel bonding (failover mode)
SLES 9	embedded iSCSI software initiator	use ethernet channel bonding (failover mode)

a. Check driver readme files for detailed version numbers and service pack levels.

6.2.2 DS300 dual controller switched dual path attachment

A server in this configuration has two interfaces (two hardware initiators or one software initiator using two network interfaces). Each of the interfaces is connected to a different subnet. Depending on the type of the initiator, it is possible to use multipathing via operating system specific multipathing mechanisms or multiple connections per session (MCS). Only with the software initiators is multipathing possible. Even when multipathing is possible, it is not always supported. Check support pages and firmware readme files about current support status. With hardware initiators like IBM iSCSI Server Adapter or the Qlogic iSCSI Expansion Card for IBM IBM @server BladeCenter, multipathing is not possible because at the time this redbook was written, no multipathing driver was available.

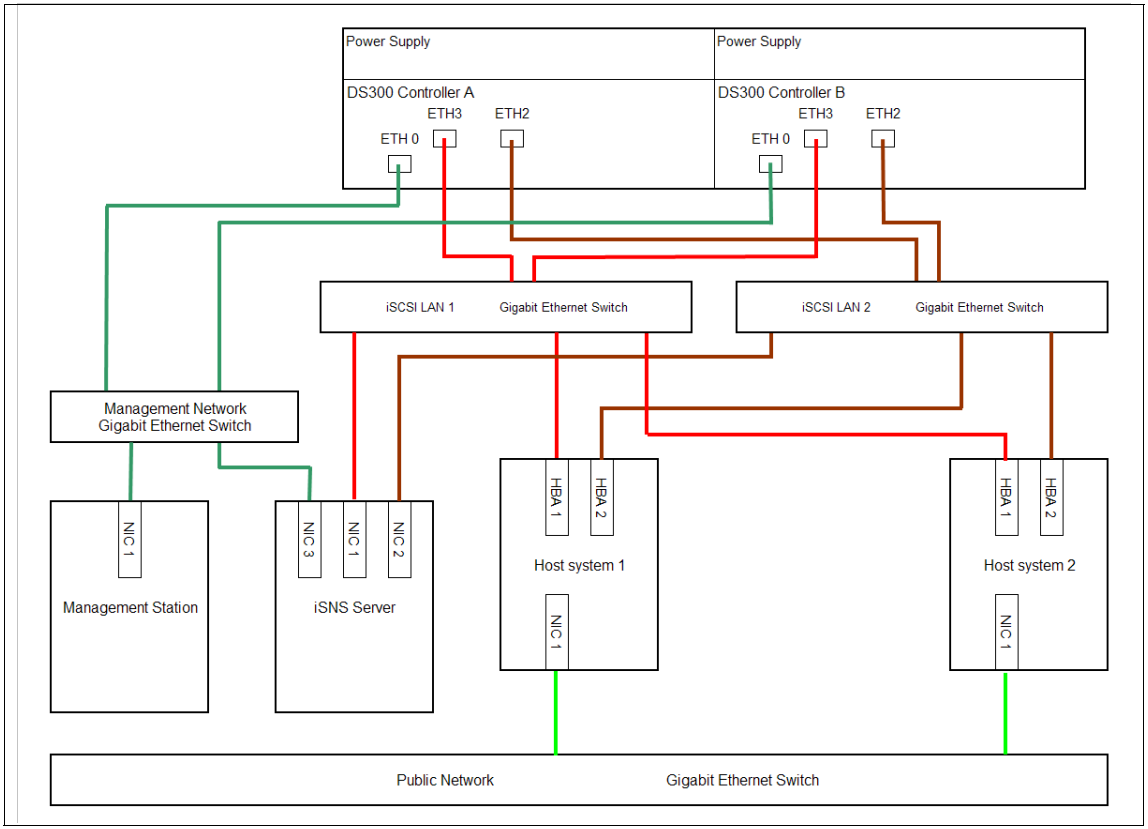

Figure 6-5 DS300 dual controller switched dual path attachment

Table 6-4 gives an overview about the usage of different types of initiators in this configuration.

Table 6-4 DS300 dual controller switched dual path attachment initiator remarks

OS type	Initiator[a]	Remarks
Windows 2000 and 2003	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter	spread logical drives across the HBAs, no multipathing
Windows 2000 and 2003	MS software initiator 1.06	spread logical drives across the HBAs, no multipathing
Windows 2000 and 2003	MS software initiator 2.0	spread logical drives across the HBAs, no multipathing

OS type	Initiator[a]	Remarks
Linux	IBM iSCSI Server Adapter or QLogic iSCSI Expansion Card for BladeCenter	spread logical drives across the HBAs, no multipathing
Red Hat 3	iSCSI software initiator V3.6.3, V3.6.4	multipathing using device mapper or MCS possible, but not supported by IBM
SLES 8	iSCSI software initiator V3.6.3, V3.6.4	multipathing using device mapper or MCS possible, but not supported by IBM
Red Hat 4	embedded iSCSI software initiator	multipathing using device mapper or MCS possible, but not supported by IBM
SLES 9	embedded iSCSI software initiator	multipathing using device mapper or MCS possible, but not supported by IBM

a. Check driver readme files for detailed version numbers and service pack levels.

7

TotalStorage DS300 installation

This chapter gives detailed step-by-step installation and configuration examples for the TotalStorage DS300 SAN controller.

In this chapter, these topics are discussed:

Note: At the time this book was written, RHEL4 had no iSCSI software initiator support. There is a statement about this in the Red Hat Enterprise Linux AS 4 Release Notes: The initial release of Red Hat Enterprise Linux 4 does not include iSCSI software initiator or target support. Support for iSCSI is being evaluated for addition in a future update to Red Hat Enterprise Linux 4.

DS300 LAB Setup

We use four IP networks in our lab setup. These networks are:

- Public LAN: used for file sharing, domain logon in Windows network
- Management LAN: used for management of DS300, BladeCenter and RSA adapters
- iSCSI LAN 1 and 2: used for iSCSI traffic

Table 7-1 Different networks in our DS300 LAB setup

Network	Network address	Subnet mask	Gateway	available IP addresses
Public LAN	172.16.0.0	255.255.192.0	172.16.0.1	172.16.0.2 - 172.16.63.254
Mgmt. LAN	192.168.1.224	255.255.255.224	192.168.1.225	192.168.1.226 - 192.168.1.254
iSCSI LAN 1	192.168.1.0	255.255.255.192	192.168.1.1	192.168.1.2 - 192.168.1.63
iSCSI LAN 2	192.168.1.64.	255.255.255.192	192.168.1.65	192.168.1.66 - 192.168.1.127

Both of the two DS300 controllers are connected to three networks as shown in Table 7-2.

Table 7-2 DS300 IP configuration

Controller	Management LAN IP (ETH0)	iSCSI LAN 1 IP (ETH2)	iSCSI LAN 2 IP (ETH3)
Controller A	192.168.1.230 255.255.255.224	192.168.1.61 255.255.255.192	192.168.1.125 255.255.255.192
Controller B	192.168.1.231 255.255.255.224	192.168.1.62 255.255.255.192	192.168.1.126 255.255.255.192

7.1 Boot from a DS300 with QLogic iSCSI Expansion Card for IBM eServer BladeCenter

This scenario will explain in detail the setup of an IBM @server Blade HS20 with QLogic iSCSI Expansion Card for IBM @server BladeCenter. There are no hard drives installed in the blade. One of the ports of the dual port iSCSI card will be used to boot the blade. The second port will be configured after installation of the operating system to get access to more logical drives of the TotalStorage DS300. Used operating system will be Microsoft Windows Server 2003 Enterprise Edition. Microsoft's iSCSI simple naming service will be used to discover iSCSI targets in the SAN.

- Planing and considerations
- SAN Setup
- DS300 Setup
- host bus adapter setup
- Microsoft Windows Server 2003 Enterprise Edition installation
- Driver and service pack installation

7.1.1 Planning and considerations

A host bus adapter failover driver did not exist at the time this book was written. The logical drives will be accessible only over one path. This means that logical drives should be distributed across controllers and host bus adapters.

The failover policy of the TotalStorage DS300 will be "remote". This is the default setting and allows a failover of arrays and iSCSI interface IP addresses to the alternate controller. A controller failover will not result in a inaccessible boot device. However, an iSCSI host bus adapter failure, at least the port used for booting, will bring the system down.

The BIOS setup of the Qlogic iSCSI Expansion Card host bus adapter does not allow for the configuration of additional logical drives. This needs to be done in the running operating system with the SANsurfer utility.

For every slot in the blade center, 4 ports can be routed out of the BladeCenter. Port one and two are used for the onboard dual port ethernet adapter. One port is routed to the module bay 1 and the other port is routed to the module bay 2. The add-in cards like the Qlogic iSCSI Expansion Card card are also dual port adapters. One port of the add-in card is routed to the module bay 3 and the other port of the add-in card is routed to module bay 4.

The BladeCenter module bays 1 and 2 are populated with D-Link gigabit ethernet switch modules. One of two ports of the onboard network adapter is used to access the public LAN. the other port is not used. Module bay 3 and 4 are populated with Nortel layer 2-3 gigabit ethernet switches. Those switches are used to connect to the two iSCSI networks. Depending on the port type of the external switches, pass through modules can be used as well. IBM offers a copper pass through and an optical pass through module. The copper pass through module can only operate at gigabit speed and can handle only ethernet. More flexibility in terms of the possible protocols is provided by the optical pass through module. The big advantage of the optical pass through module is the possibility that different add on cards can be installed in the blades of the same BladeCenter. For example a fibre channel host bus adapter can be installed in one blade and a gigabit ethernet card in a other blade. The pass through module provides a single cable for every port installed in a blade. This cable needs to be connected to an external gigabit ethernet switch or a fibre channel switch. More cabling is required than using switch modules in the module bays of the BladeCenter, but no configuration of the pass through modules is required.

Note: When using switch modules in the module bays 3 and 4, only adapter cards of the same type as the switch can be installed in the blades. For example, when a gigabit ethernet switch is installed in bays 3 and 4, then only gigabit adapter cards can be installed in the blades.

Verify the current support status before starting implementation. There might be changes in the supported devices.

7.1.2 SAN and network setup

The following diagram shows the cabling of the components. It follows the example configuration as explained in 6.2.2, “DS300 dual controller switched dual path attachment” on page 219.

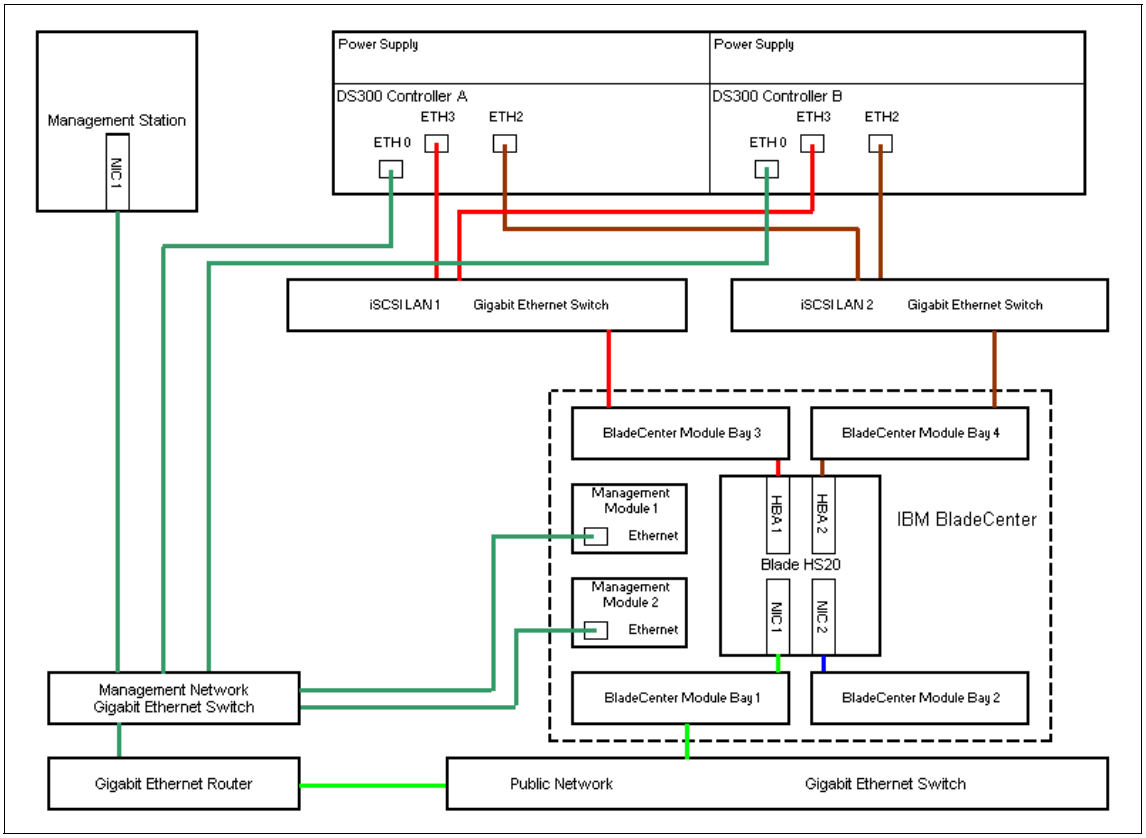

Figure 7-1 Configuration setup

The QLogic iSCSI Expansion Card is installed in an HS20 blade.

No special configuration of the D-Link gigabit ethernet switch modules in module bay 1 and 2 or the Nortel layer 2/3 gigabit ethernet switch modules in bay 3 and 4 is required. The switches are in factory default state. External ports where enabled to get the connect to the external Cisco 3550 series gigabit ethernet switches for iSCSI traffic and public LAN access.

The BladeCenter has two management modules installed and all four module bays are populated with gigabit ethernet switch modules. That is the reason that the management modules are connected into the management network which has no direct connect to either of the networks used for the switch modules in the blade center. When the management modules are direct connected into one of the networks used with the switch modules, the switches can see MAC addresses from devices twice. One time they can see it over the management

module connection and one time over the external ports of the switches. This will cause network problems.

The setup uses a few names which will be found in the various figures. They are:

- System name is COLORADO.
- The logical drive used to boot from is named COLORADO-BOOT-W.

The TotalStorage DS300 network ports were configured as described in 4.4.1, “Setup of management ports of TotalStorage DS300 and DS400” on page 76 and 4.4.2, “Setting the IP addresses of the iSCSI ports from the command line interface” on page 80.

7.1.3 TotalStorage DS300 setup

For this installation, one logical drive COLORADO-BOOT-W was created to boot from. The array of this logical drive is owned by controller A. An additional logical drive with name COLORADO-LOG was created to have an additional logical drive during the HBA setup. This logical drive is in its own array and also owned by Controller A. 4.8.1, “Create arrays, logical drives, and assign IQNs on the TotalStorage DS300” on page 114 describes how to create arrays and logical drives. The second HBA has at this time no access to logical drives.

Table 7-3 Array and logical drive configuration of the TotalStorage DS300

Array	Drives	Owned by	Raid Level	Logical Drive	Size
Array1	2	A	1	COLORADO-BOOT-W	10GB
Array2	6	A	50	COLORADO-LOG	10GB
Array3	4	B	10	COLORADO-MAIL	150GB

7.1.4 iSCSI host bus adapter setup and installation preparation

The following steps are done to configure the QLogic iSCSI Expansion Card for IBM @server BladeCenter in a IBM @server Blade HS20. The expansion card is installed and uses the module bays 3 and 4 to connect outside of the BladeCenter.

1. Update BIOS and firmware of the HBA to the latest level supported by IBM for the DS300. The BIOS can be downloaded from IBM’s support Web page:

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-60280
   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

See Appendix C, “How to get updates of code and documentation” on page 623. At time of writing the supported BIOS level was 1.47. The readme file of this update contains the details how the update must be performed.

2. During POST, the BIOS of the QLogic iSCSI Expansion Card for IBM @server BladeCenter shows a message **Press Ctrl+Q for Fast!UTIL**. Press **CTRL**+**Q** to enter the setup.

Note: ALT+Q works as well.

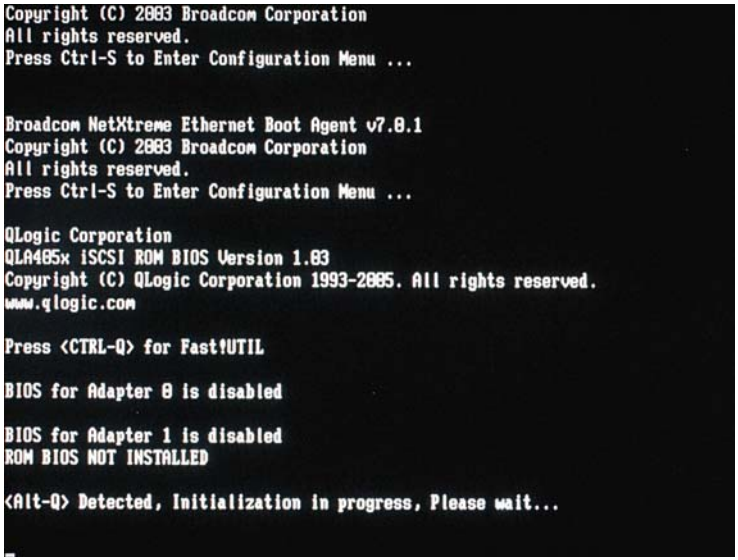

Figure 7-2 POST — QLogic iSCSI Expansion Card for IBM @server BladeCenter *BIOS message*

3. Select the second HBA and press **Enter**.

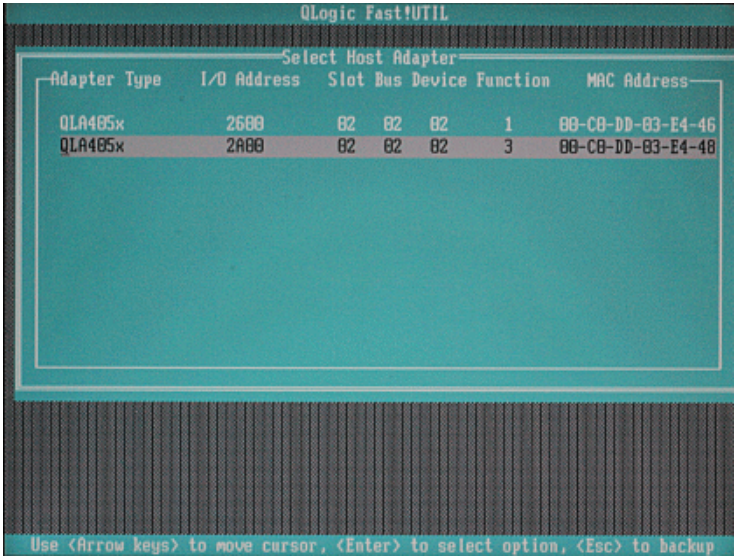

Figure 7-3 Select host adapter

4. Select **Configuration Settings** and press **Enter**.

Figure 7-4 Fast!Util options — Configuration settings

5. Select **Host Adapter Settings** from the **Configuration Settings** and press **Enter**.

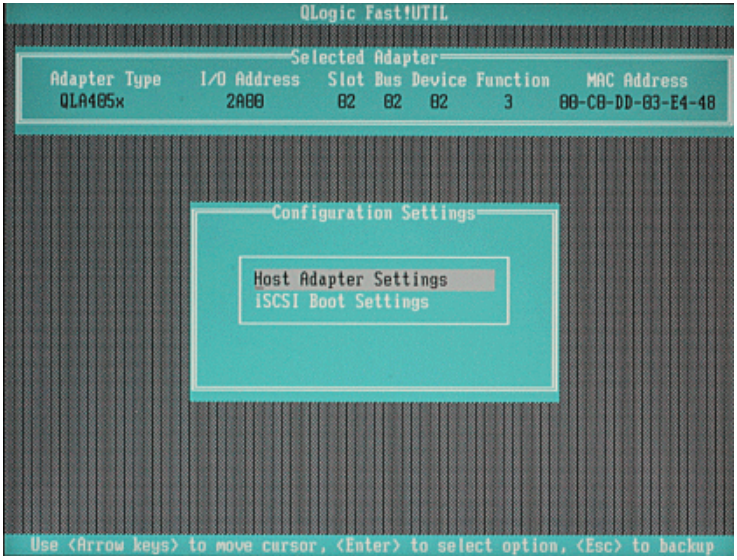

Figure 7-5 Configuration settings — Host adapter settings

6. Verify that the Adapter BIOS is disabled. Disable it if it is enabled. Select **Host Adapter BIOS: Enabled** and press **Enter**. The setting will change to **Disabled**. Then Press **ESC** to return to the **Configuration Settings**.

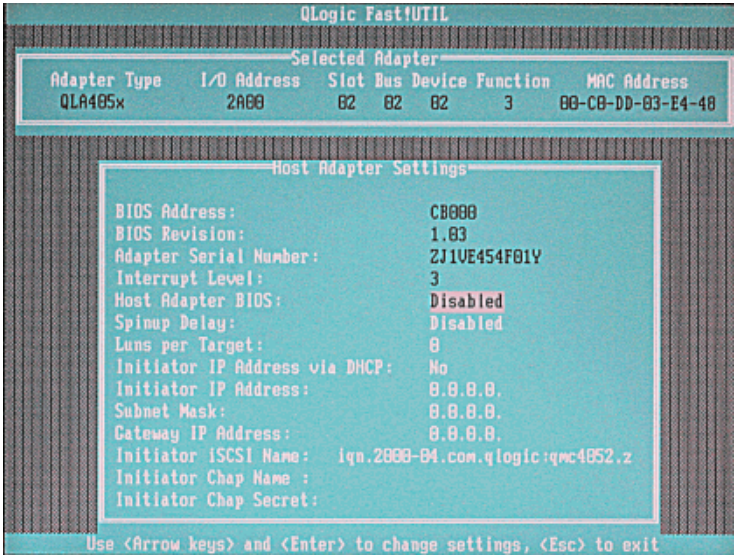

Figure 7-6 Host adapter settings

7. Select **iSCSI Boot Settings** and press **Enter**.

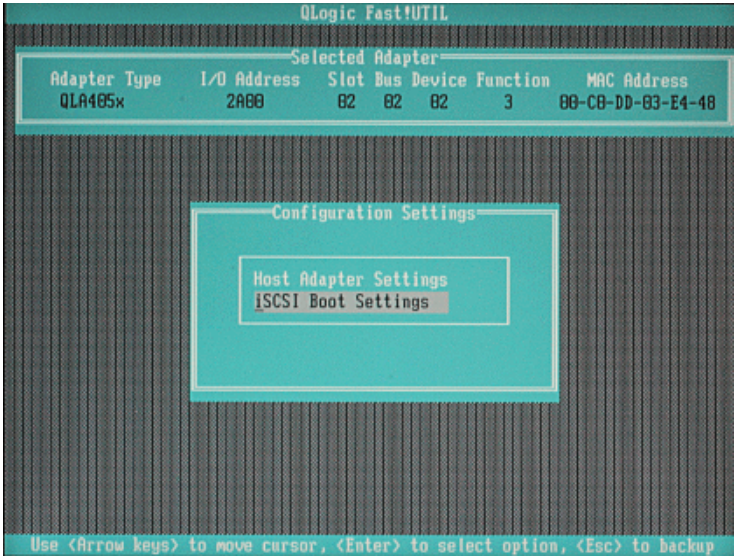

Figure 7-7 Configuration settings — iSCSI boot settings

8. Verify that the iSCSI Boot is disabled. When iSCSI Boot is enabled, select it and press **Enter** to change the value from enabled to disabled. An IP address is not set here. This will be done later from the operating system.

   Press **ESC** to leave **iSCSI Boot Settings** and return to **Configuration Settings**.

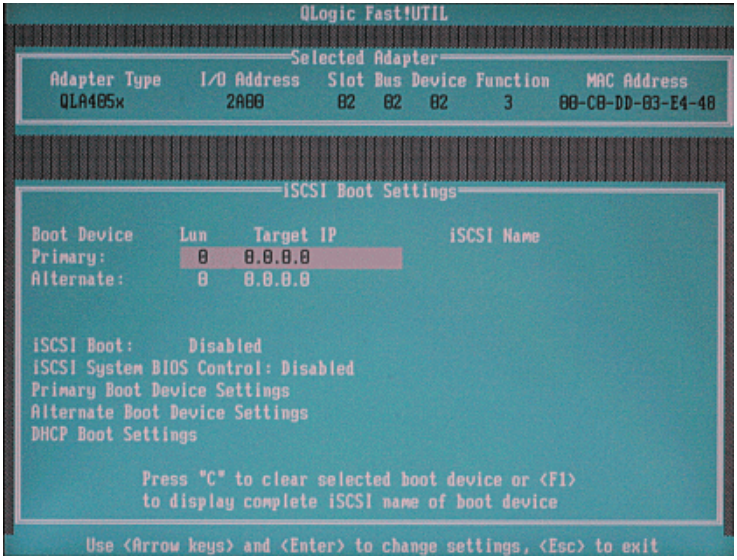

Figure 7-8 iSCSI boot settings

9. Press **ESC** to leave the **Configuration Settings**.

Figure 7-9 Configuration settings

10. When changes are made in **Host Adapter Settings** or **iSCSI Boot Settings,** these changes can be saved now. Select **Save changes** and press **Enter**. If no changes made in the setting, this dialog will not presented.

Figure 7-10 Save modifications

11. Select **Host Adapter** and press **Enter**.

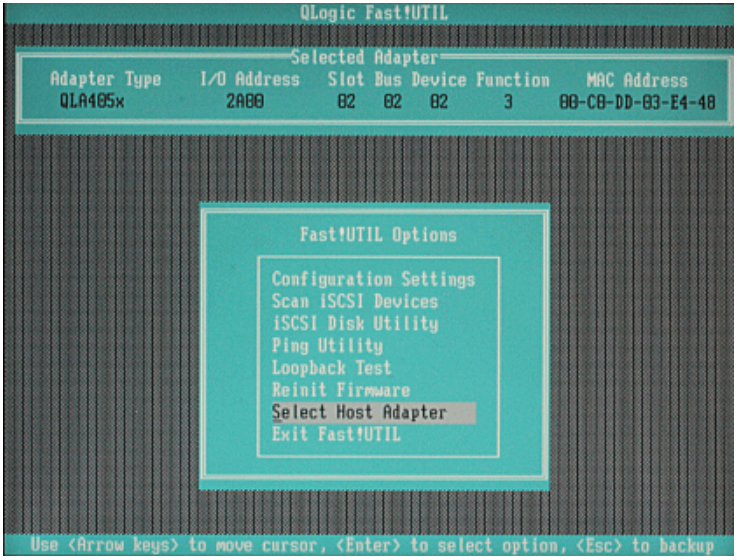

Figure 7-11 Fast!UTIL options — Select host adapter

12. Select the first host adapter and press **Enter**.

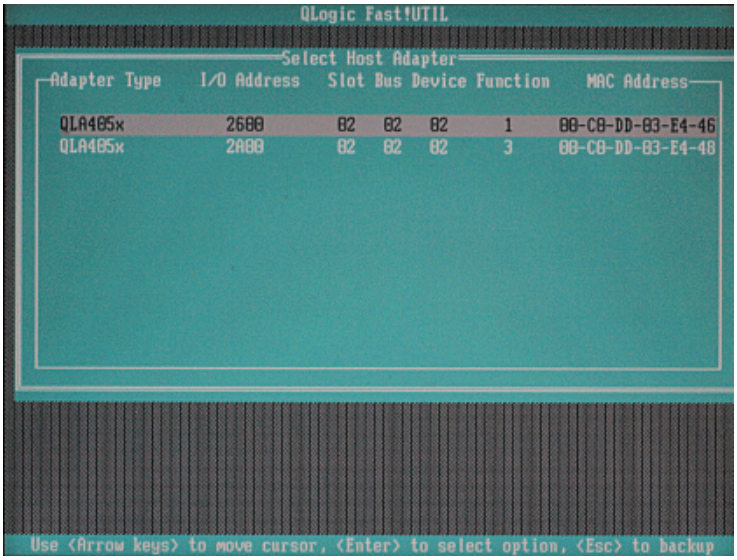

Figure 7-12 Select host adapter

13. Select **Configuration Settings** and press **Enter**.

Figure 7-13 Fast!UTIL options — Configuration Settings

14. Select **Host Adapter Settings** and press **Enter**.

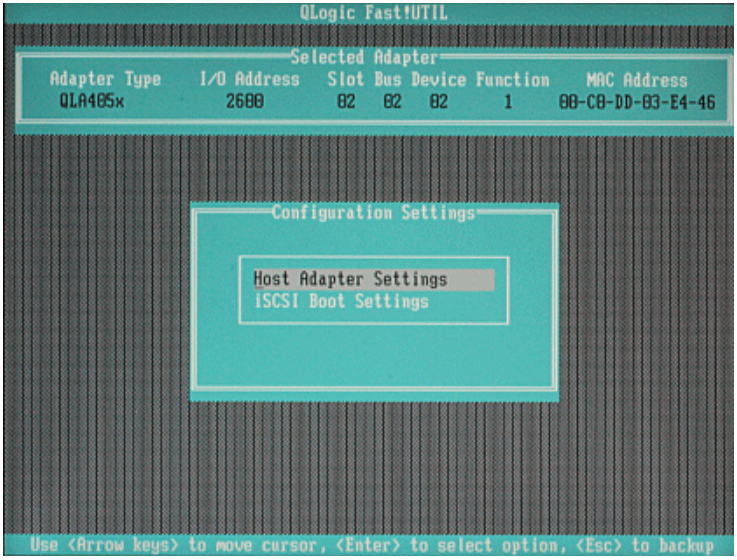

Figure 7-14 Configuration settings — Host adapter settings

15. Select **Host Adapter BIOS** and enable it by pressing **Enter**. The value changes from Disabled to Enabled. Then select **Initiator IP Address** and press **Enter**.

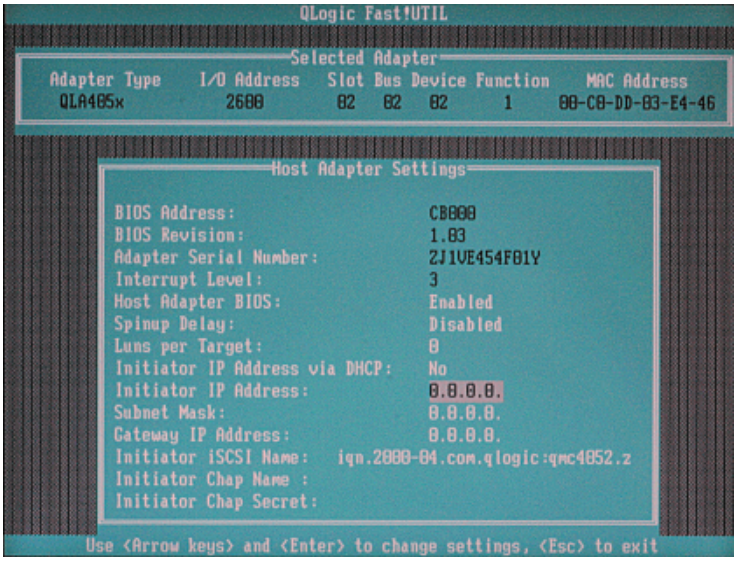

Figure 7-15 Host adapter settings

16. Enter the new IP address `192.168.1.5` of the first HBA and press **Enter**.

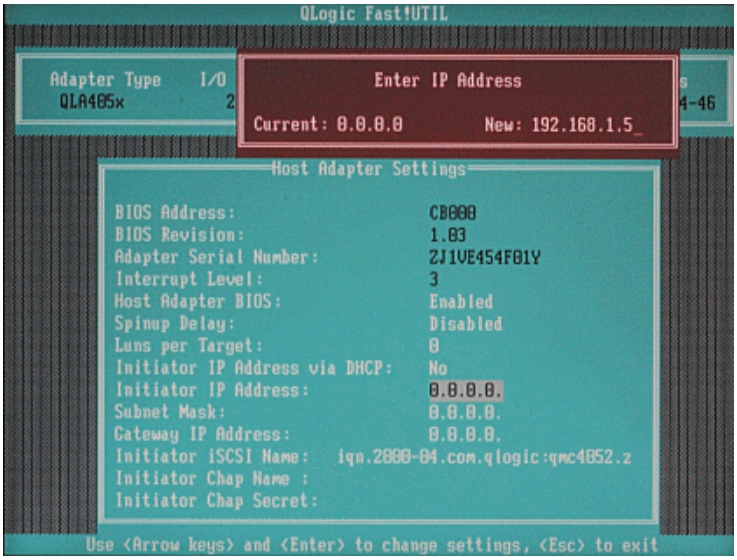

Figure 7-16 Enter HBA IP Address

17. The IP address was added into the dialog. Select **Subnet Mask** and press **Enter**.

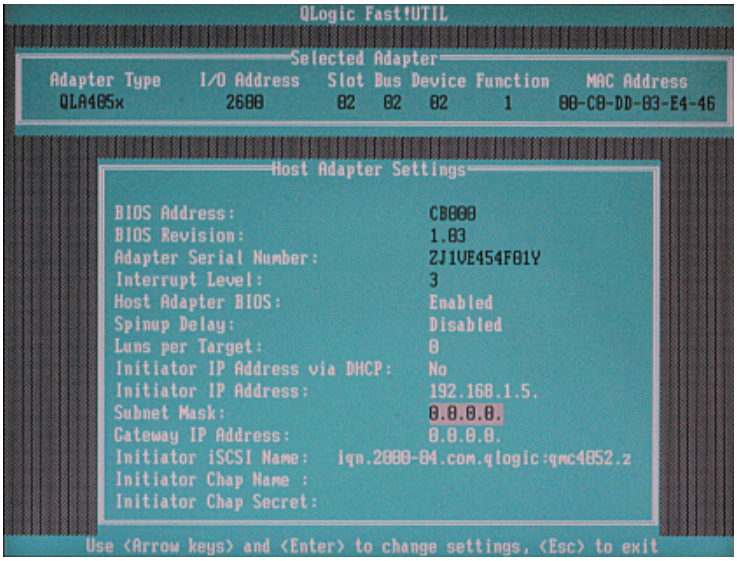

Figure 7-17 Subnet mask

18. Enter the subnet mask `255.255.255.192` and press **Enter**.

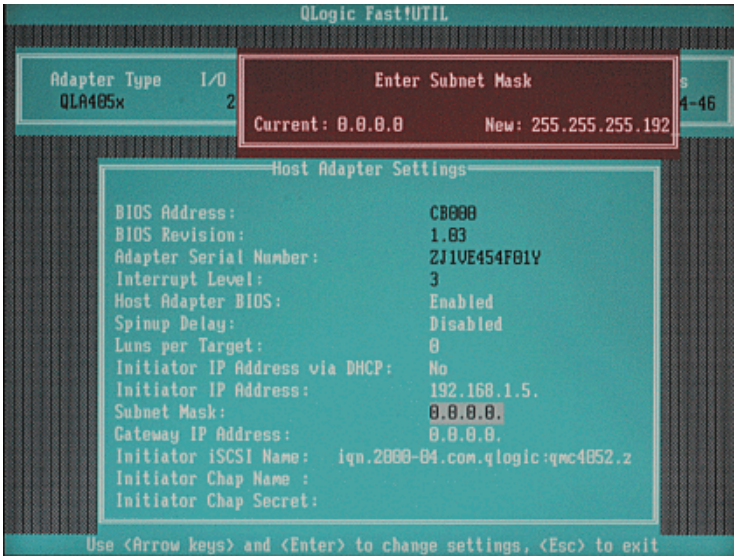

Figure 7-18 Enter subnet mask

19. Select the **Gateway IP Address** and press **Enter**.

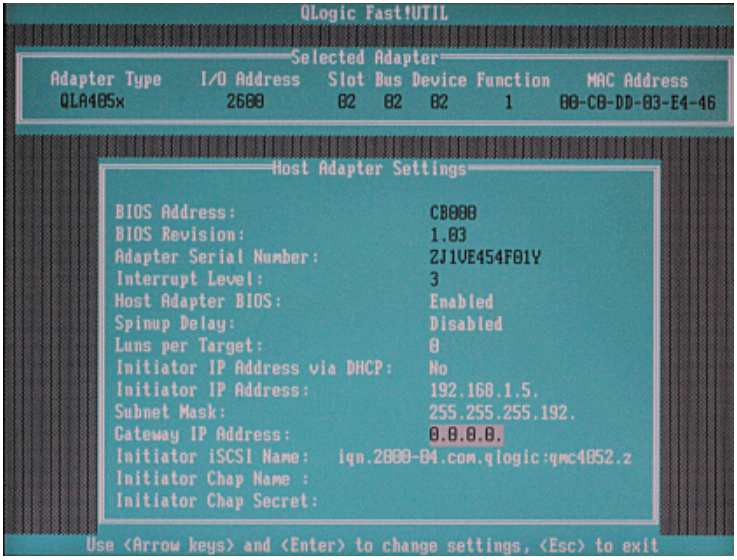

Figure 7-19 Select gateway IP address

20. Enter the gateway **IP** address 192.168.1.65 and press **Enter**.

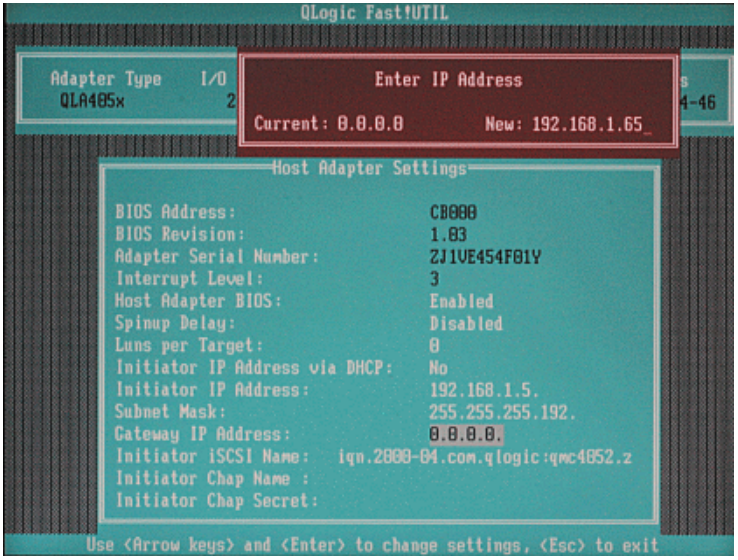

Figure 7-20 Enter gateway address

21. Select the **Initiator iSCSI Name** and press **Enter**.

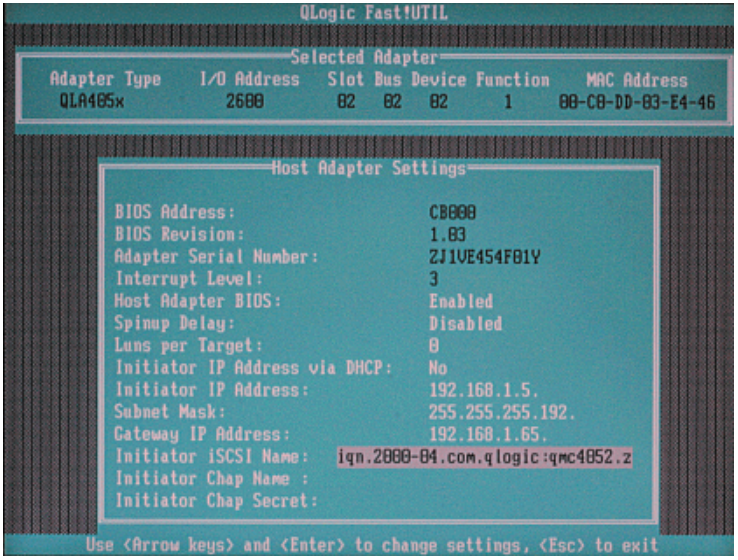

Figure 7-21 Select IQN

22. The full iSCSI qualified name is shown. Make note of the iSCSI qualified name because it will be required later to configure the access control list of the TotalStorage DS300. Press **ESC** to return to **Adapter Settings**.

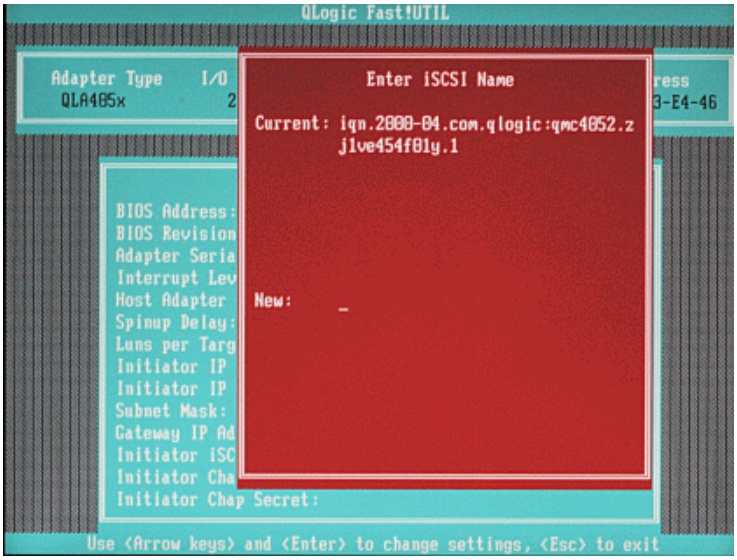

Figure 7-22 Full IQN of the selected HBA

23. Press ESC to return to Fast!UTIL options and save modifications.

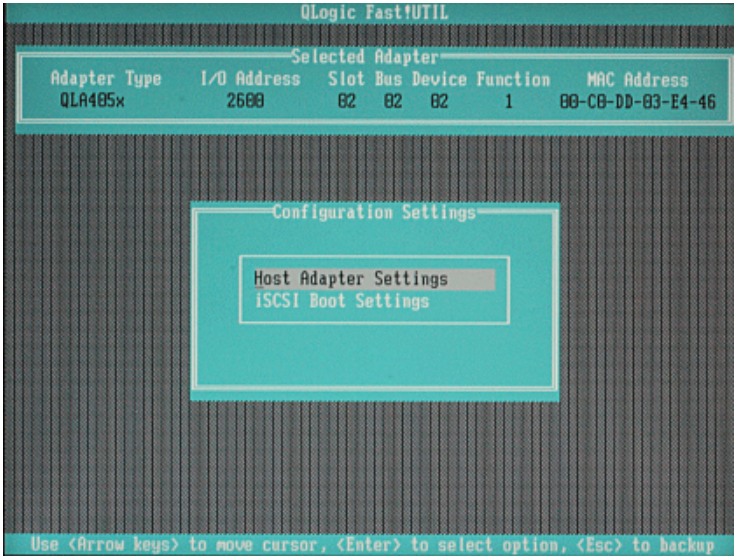

Figure 7-23 Configuration settings

24. Save the defined IP address by selecting **Save change**s and press **Enter**. This is important before defining iSCSI boot settings. The adapter first needs to connect to the storage subsystem before a LUN to boot from can be selected. Thus requires defined and working network settings.

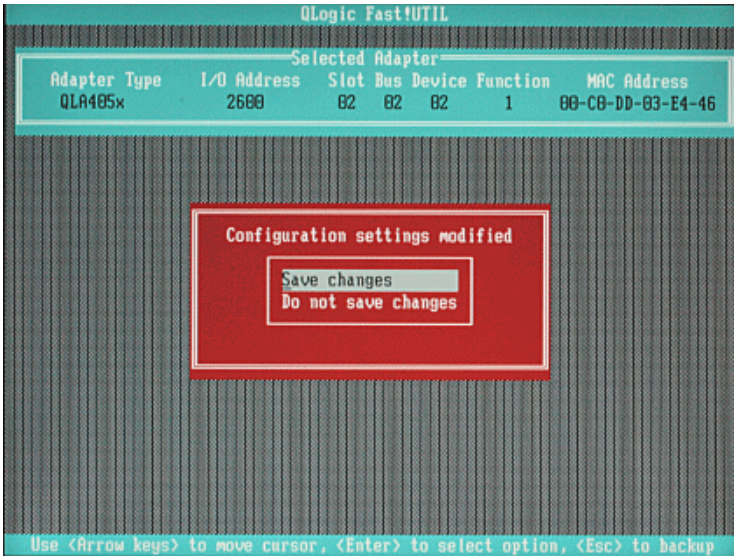

Figure 7-24 Save modifications

25. Select **Configuration Settings** and press **Enter**.

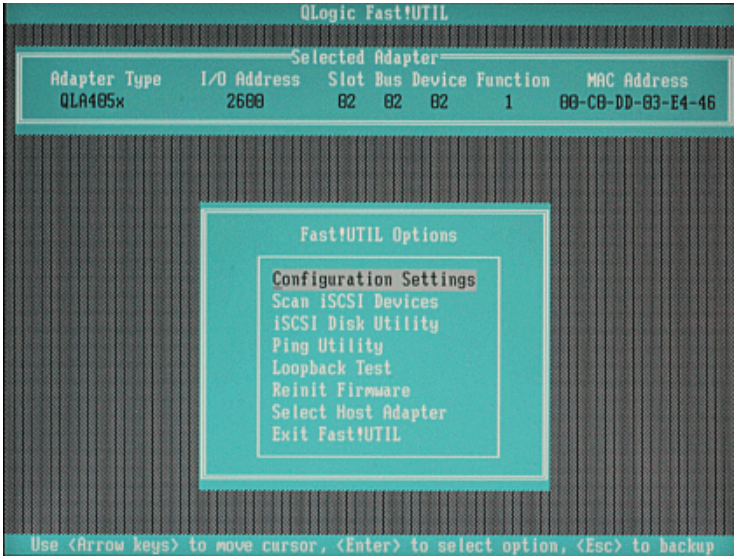

Figure 7-25 Fast!UTIL options — Configuration settings

26. Select **iSCSI Boot Settings** and press **Enter**.

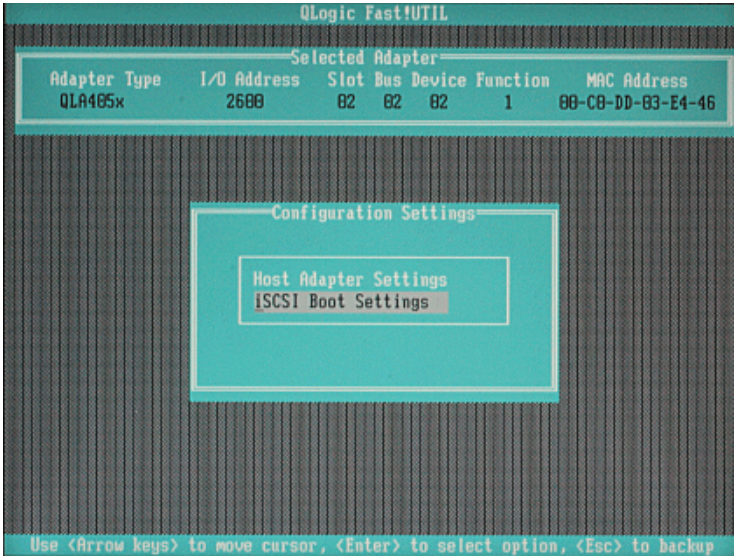

Figure 7-26 Configuration Settings — iSCSI Boot Settings

27. On Figure 7-27 on page 243, select **iSCSI Boot: Disabled**, and press **Enter** to enable iSCSI boot. Select **Primary Boot Device** and press **Enter**. The

adapter will scan for available iSCSI targets in local area network. It takes a moment to scan through all possible targets.

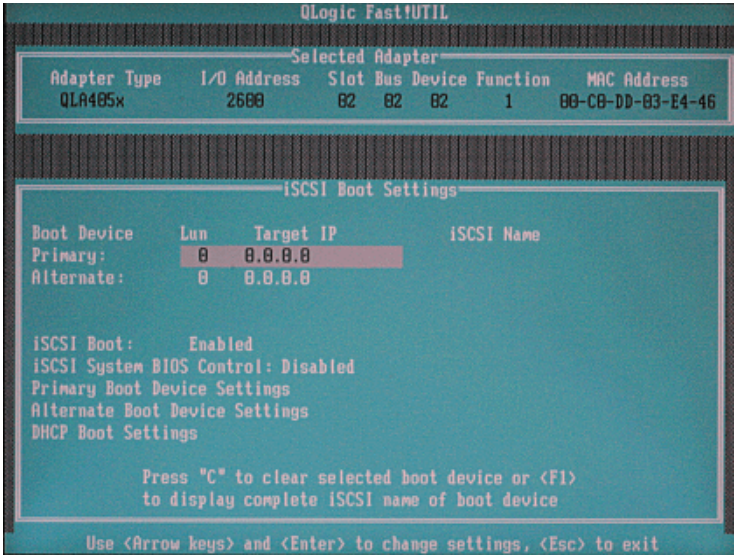

Figure 7-27 select primary boot device

28. A list with available targets is shown. See Figure 7-28 on page 244. Select the target to boot from and press **Enter**. You may scroll down the list to see more targets. They may not to be shown sequentially in the list.

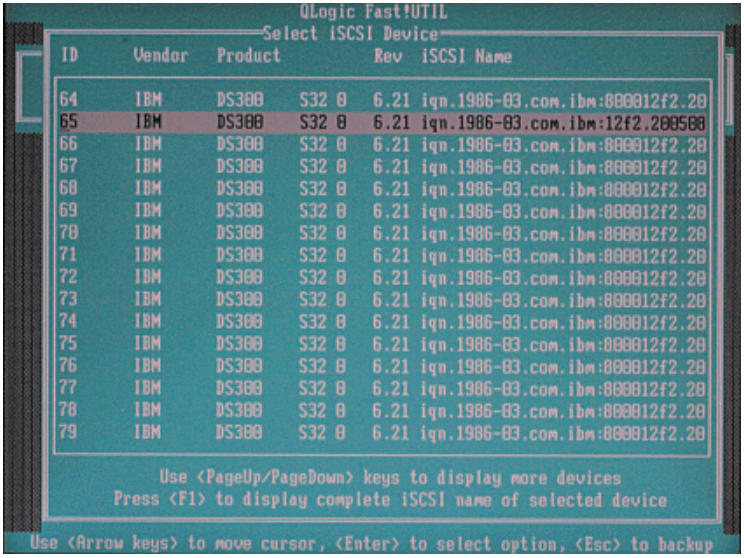

Figure 7-28 iSCSI targets

29. The selected target will be added as primary boot device. Select **Primary Boot Device Settings** to verify the full iSCSI qualified name of the selected boot device. Then press **Enter**.

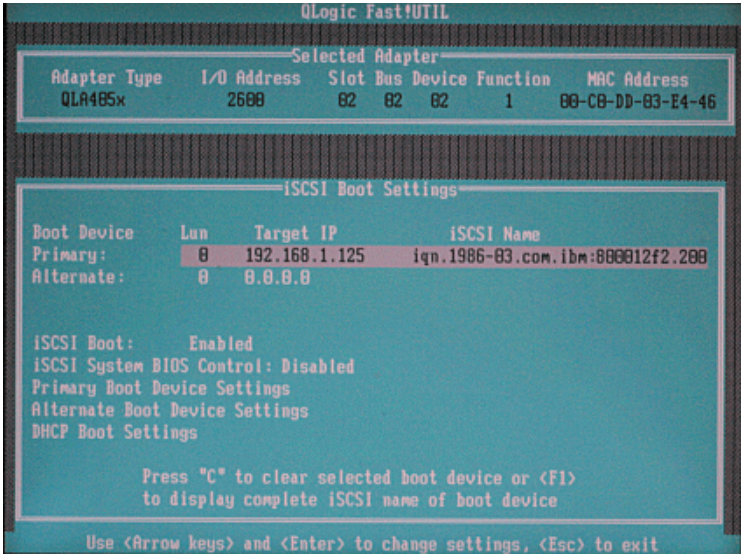

Figure 7-29 iSCSI boot settings — primary boot device settings

30. The name of the logical drive will be part of the IQN.

> **Note:** When a logical drive is renamed in the CLI of the TotalStorage DS300 the IQN will not change.

Press **ESC** to return to the **iSCSI Boot Settings**.

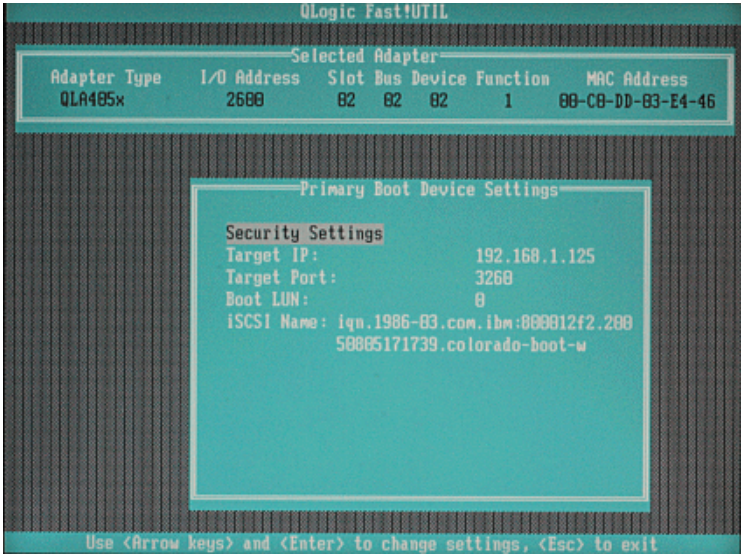

Figure 7-30 Detailed primary boot device settings

31. Press **ESC** to return to **Configuration Settings**.

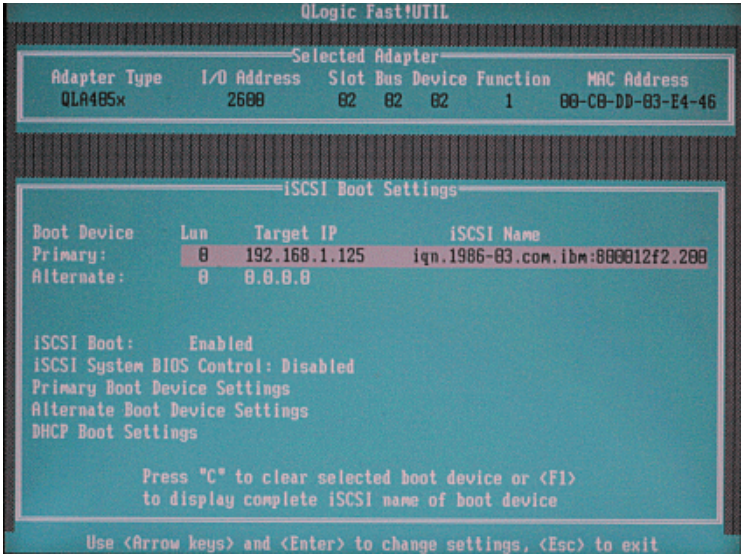

Figure 7-31 iSCSI Boot settings

32. Select **Save changes** and press **Enter**.

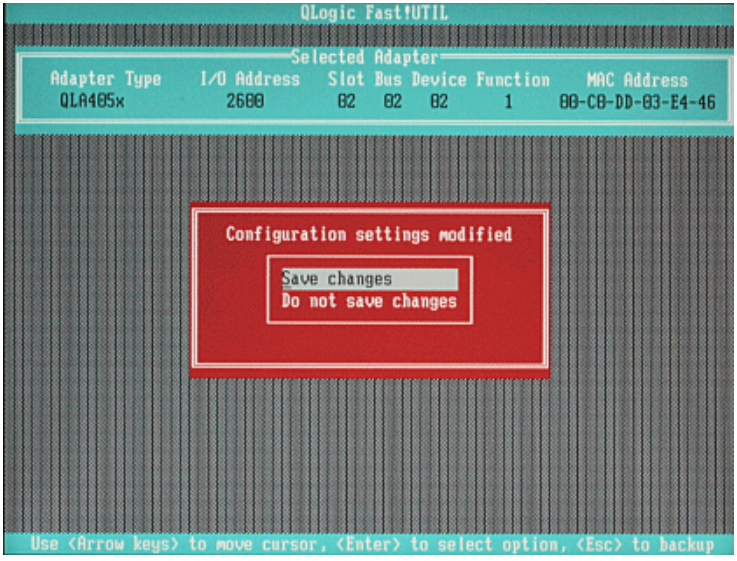

Figure 7-32 Save mollifications

33. Select Exit Fast!UTIL and press **Enter**.

Figure 7-33 Fast!Util options — Exit Fast!UTIL

34. Do not reboot the system at this time.

7.1.5 Create QLogic iSCSI Expansion Card for IBM eServer BladeCenter driver disk

The driver disk is required because the driver is not included on the Microsoft Windows Server 2003 Enterprise Edition CD.

1. Download the driver for the QLogic iSCSI Expansion Card for IBM @server BladeCenter from IBM's support Web site.

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

   There are different driver packages available. Use the SCSI port driver for Microsoft Windows 2000. With Windows Server 2003, use the storport driver. The packages also available with NDIS drivers. These drivers can be used to make the iSCSI host bus adapter appear as a network adapter available for the operating system. With these NDIS drivers, no TOE functions are used.

2. When using an SCSI port driver package without NDIS driver, unzip the driver archive and store all files in the root of a 3.5" diskette. When using an archive containing the SCSI port and NDIS driver, copy all files in the StorMiniport\2.1.1.5 directory into the root of a floppy diskette.

This diskette will be used as driver disk to install Microsoft Windows Server 2003.

```
Folder PATH listing
Volume serial number is 00530064 F08D:85A9
Z:.
├───Networking
│   └───2.0.0.0
└───StorMiniport
    └───2.1.1.5
```

Figure 7-34 iSCSI HBA driver archive with storport and NDIS driver

3. Proceed with the installation of Microsoft Windows Server 2003 as described in the next section.

7.1.6 Microsoft Windows Server 2003 Enterprise Edition installation

The installation of Windows 2003 is straightforward. The driver of the QLogic iSCSI Expansion Card for IBM @server BladeCenter is not included on the Microsoft Windows Server 2003 Enterprise Edition CD. This requires you to create a driver diskette first. See 7.1.5, “Create QLogic iSCSI Expansion Card for IBM eServer BladeCenter driver disk” on page 247 for information about the creation of the driver disk.

1. The system is still in the Host bus adapter BIOS setup. Select **Reboot System** and press **Enter** to leave the setup and restart the system now.

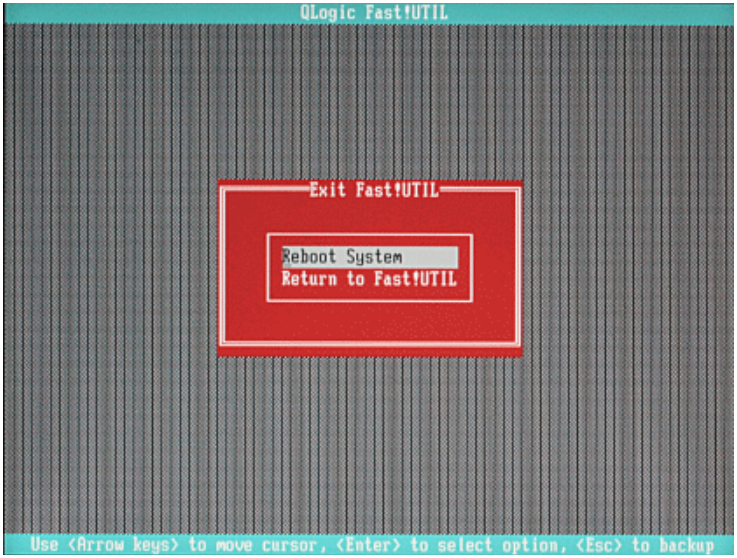

Figure 7-35 Exit Fast!UTIL

2. Insert the installation CD into the CD-ROM drive and boot from this CD. After the message Setup is inspecting your computer was shown, the black screen will switch to blue and the bottom line will give you the option to press **F6** to install third party SCSI or RAID drives which are required at boot time. Press **F6** to install the driver for the QLogic iSCSI Expansion Card for IBM @server BladeCenter.

Figure 7-36 Windows setup — Install third party SCSI or RAID drivers

3. Press **S** to specify additional SCSI or RAID controller driver.

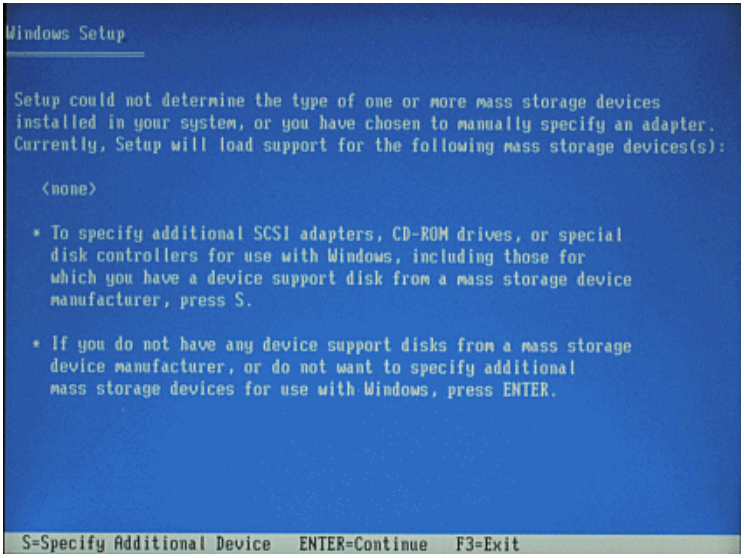

Figure 7-37 Specify additional SCSI adapter

4. Insert the driver disk created in 7.1.5, “Create QLogic iSCSI Expansion Card for IBM eServer BladeCenter driver disk” on page 247 into the floppy drive and press **Enter**.

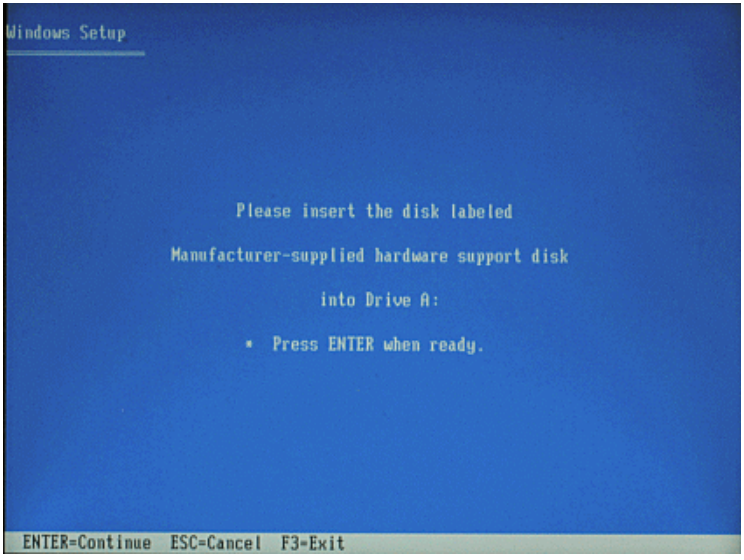

Figure 7-38 Windows setup — insert driver disk

5. The driver for the QLogic iSCSI Expansion Card for IBM @server BladeCenter will be found. Press **Enter** to select this driver.

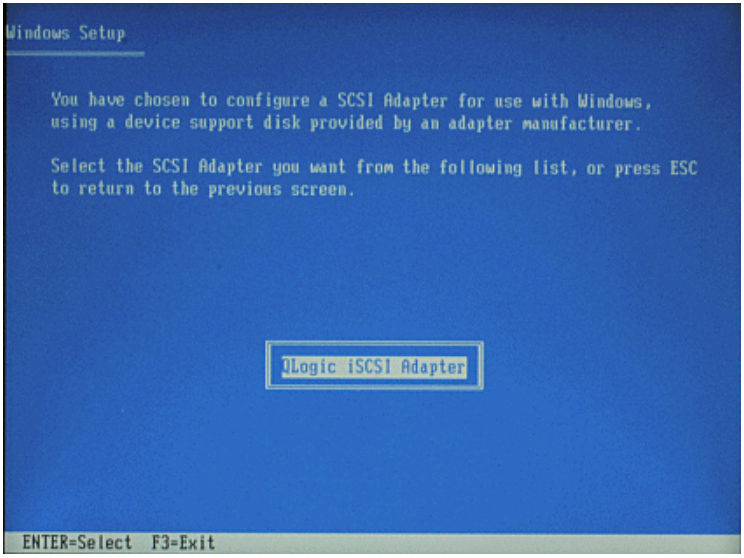

Figure 7-39 Windows setup — QLogic iSCSI Adapter

6. The driver for the QLogic iSCSI Expansion Card for IBM @server BladeCenter host bus adapter will be loaded and Windows setup returns to the screen to add more third party SCSI and RAID driver. Press **Enter** to proceed with the installation.

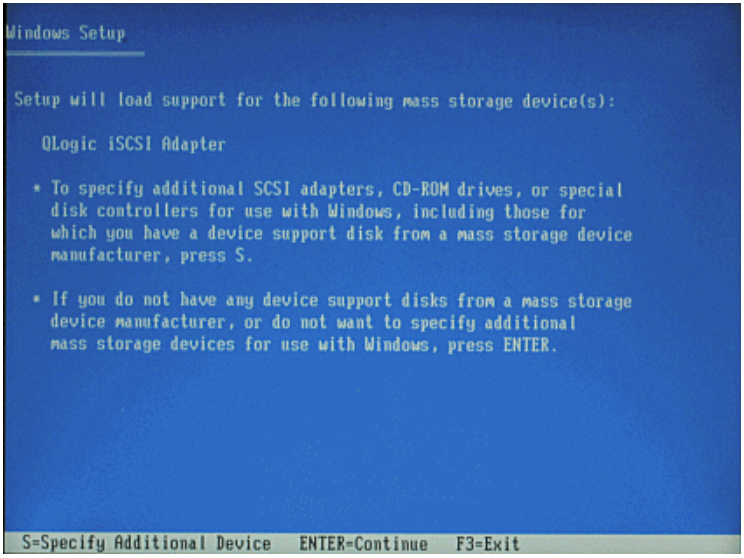

Figure 7-40 Windows setup — Supported mass storage drivers

7. For the remaining installation follow the instructions on the screen. There is no difference from the installation to a local attached drive. The planned disk to boot from is shown as last disk in the list.

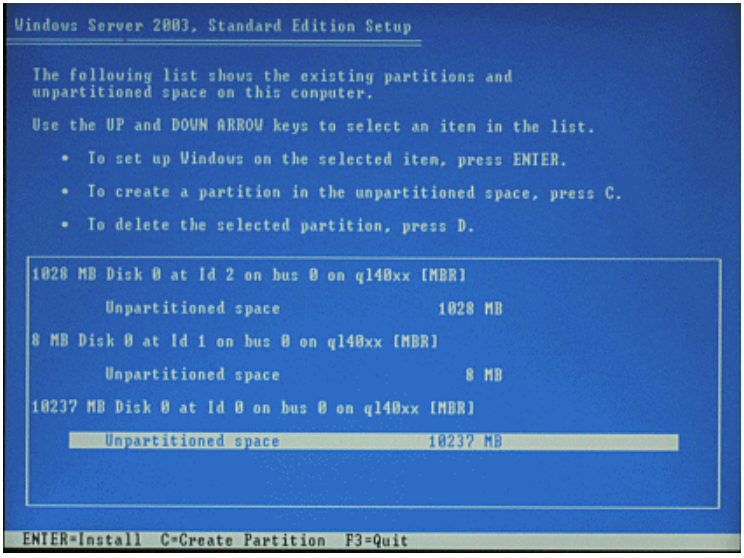

Figure 7-41 Detected disks by Windows setup

7.1.7 Driver and service pack installation and finishing the setup

Install the Broadcom network driver for the adapter. Broadcom NetXtreme Gigabit Ethernet #2 is connected to the port number 1. This port is also shared by the baseboard management controller (BMC). This port might be connected into a management network. Broadcom NetXtreme Gigabit Ethernet is routed to port number 2. This is the port close to the power supplies.

There are operating system install instructions for all IBM @serverxSeries servers on the support pages. See Appendix C, “How to get updates of code and documentation” on page 623.

1. Install service pack1 for Microsoft Windows Server 2003. Reboot the server.

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

2. Install remaining drivers not available during operating system installation:
   a. Install the Intel chipset drivers.
   b. Install ActivePCI.
   c. All unknown devices are now removed.
   d. Install IPMI driver. This driver is used for the BMC.
3. Reboot the server.

7.1.8 Install the SANsurfer utility

Install the SANsurfer utility to manage the iSCSI host bus adapter. The SANsurfer utility is a client server application. On hosts with the installed host bus adapters, at least the agent must be installed. The GUI can be installed on the same host or can be installed on a management station. It is also possible to install the GUI on the host and a management station. The latest release of the SANsurfer utility is able to manage fibre channel host bus adapter, iSCSI host bus adapter and the QLogic SAN switch for the IBM BladeCenter. The installation allows to select the components separately. Finest granularity can be achieved by the custom setup. The default setup installs everything.

1. Download the latest version from the following URL:

   http://www.ibm.com/support/docview.wss?rs=1109&uid=psg1MIGR-58619

2. Run the installer file.
3. Click **Next**.

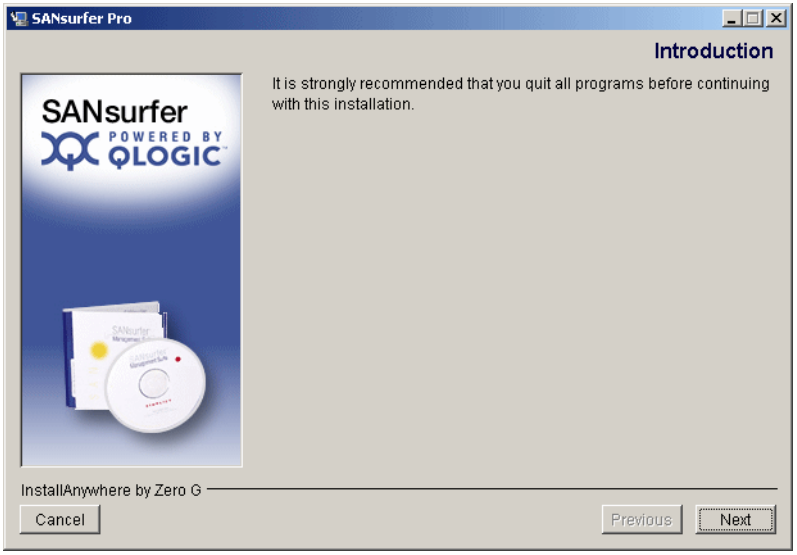

Figure 7-42 SANsurfer installation — Introduction

4. Read the important information and click **Next**.

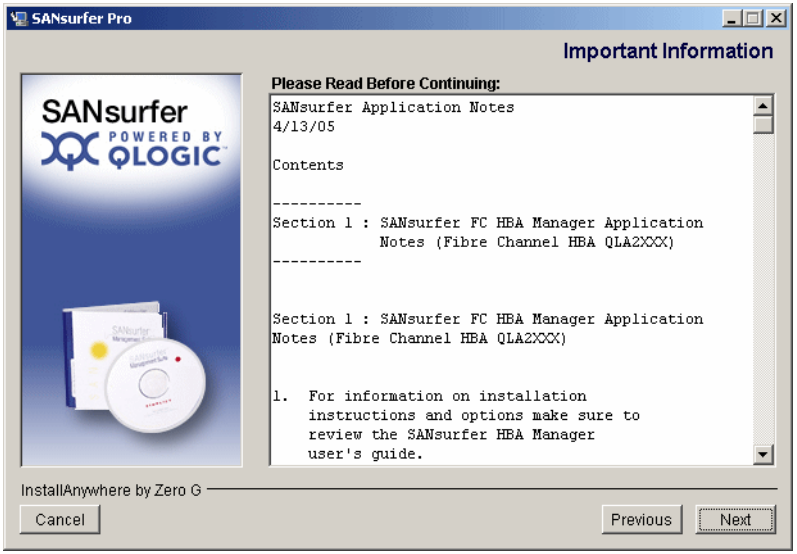

Figure 7-43 SANsurfer installation — Important information

5. Choose one of the setup options. By default all agents and GUIs are installed. Only the iSCSI agent and GUI is required for this setup. The SAN Surfer

iSCSI HBA Manager can be installed on a management station to manage the HBAs remote. Click **Next** to install all GUIs and agents.

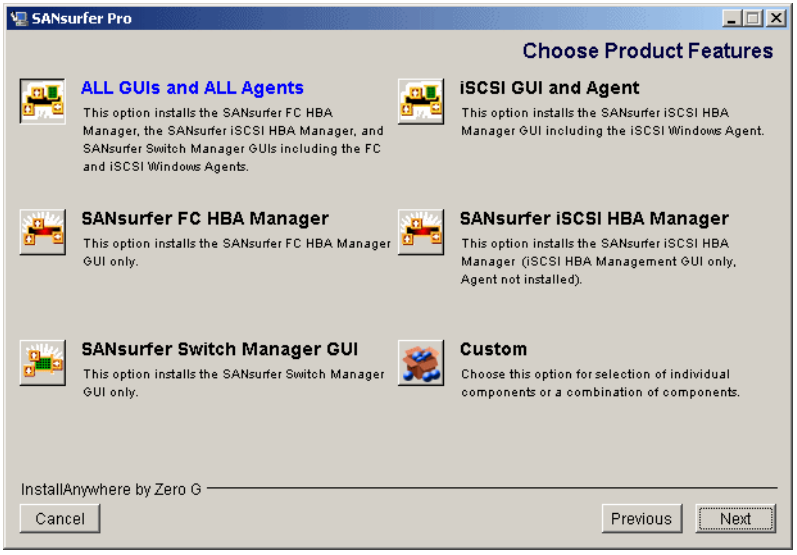

Figure 7-44 SANsurfer installation — Choose product features

6. Choose the folder where the SANsurfer will be installed. Click **Next**.

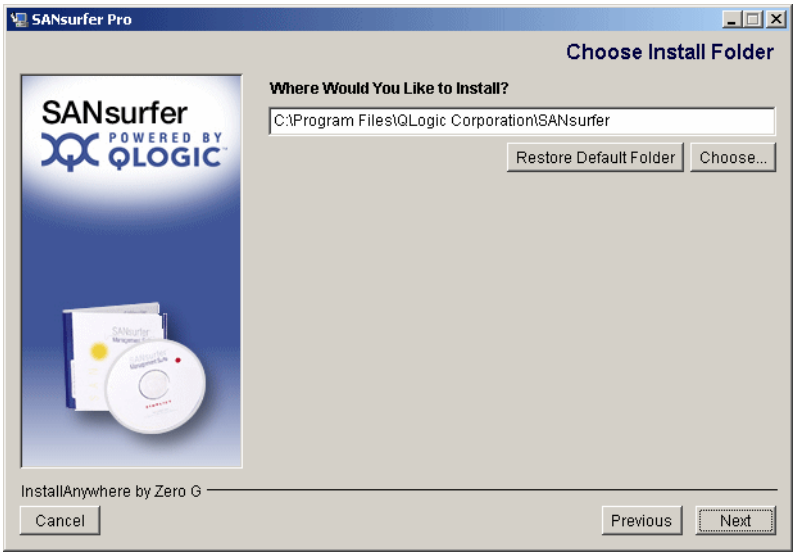

Figure 7-45 SANsurfer installation — Choose install folder

7. Select the profile where the shortcuts to SANsurfer will be created. Select All Users and click **Next**.

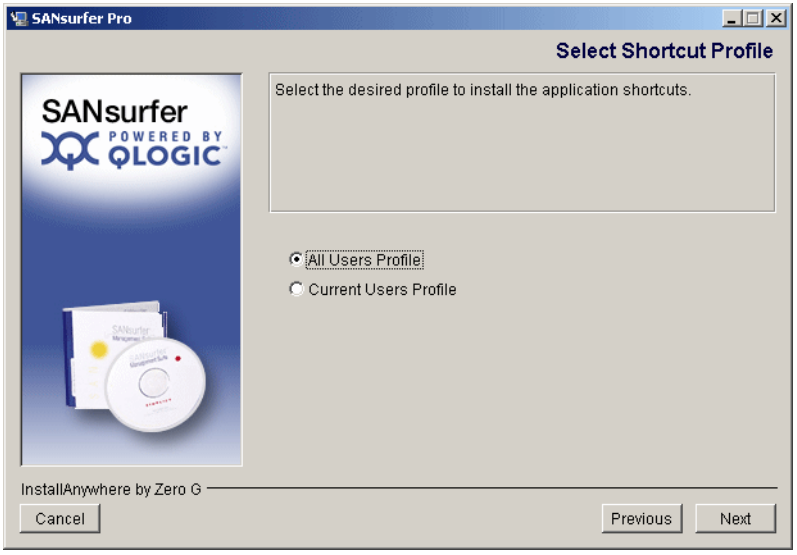

Figure 7-46 SANsurfer installation — Select shortcut profile

8. Click Next. A icon on the desktop will be created in the selected profile.

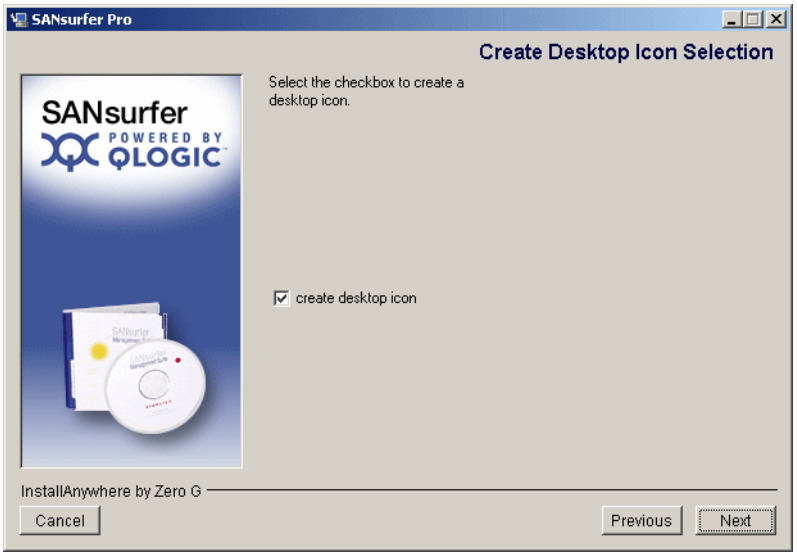

Figure 7-47 SANsurfer installation — Create desktop icon selection

9. A summary of the chosen options are shown. To install the application click **Install**.

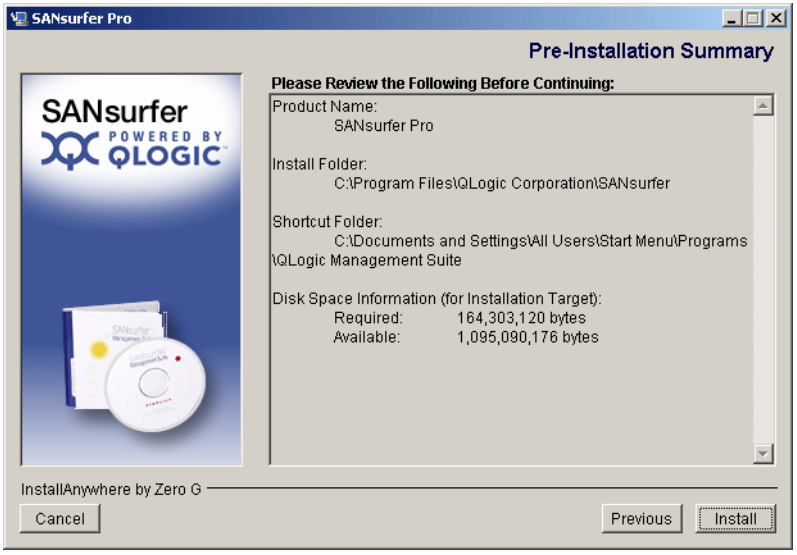

Figure 7-48 SANsurfer installation — Pre installation summary

10. After installation is finished, a window is shown to define if QLogic HBA failover can be configured. Do not activate the failover configuration and click Next.

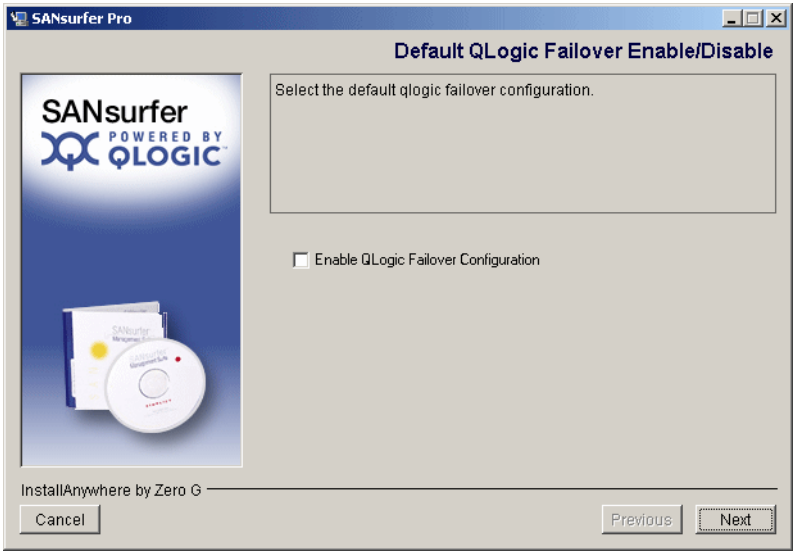

Figure 7-49 SANsurfer installation — Qlogic failover

11. Click **Done**. The installation wizard will close.

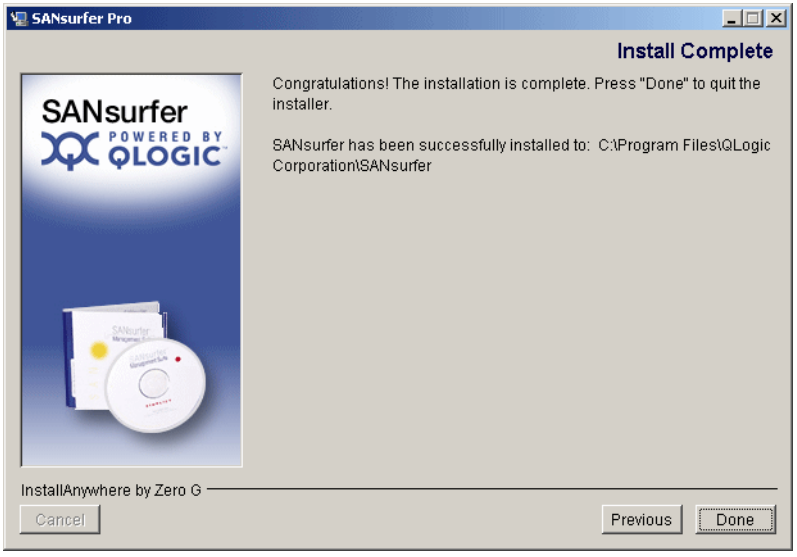

Figure 7-50 SANsurfer installation — Install complete

7.1.9 Use SANsurfer and verify current configuration

The installation of the SANsurfer utility is complete. On the desktop is now a SANsurfer icon. Use this icon to start SANsurfer and check the current HBA configuration.

1. After SANsurfer has started an empty window is shown. Depending on the chosen option during installation the left site of the window may have less tabs than shown below. SANsurfer will look like the figure below when all agents and GUIs are installed. Select the **iSCSI HBA** tab.

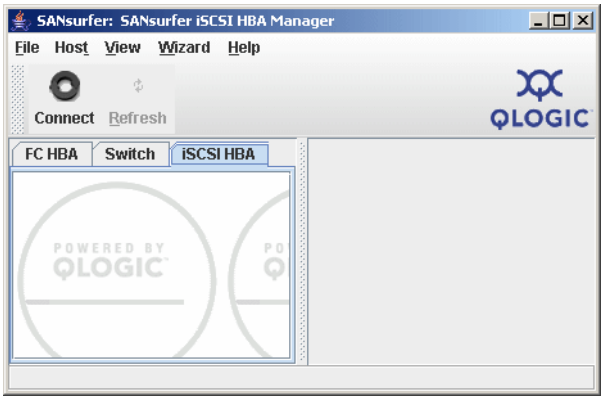

Figure 7-51 SANsurfer

2. Click **Host** → **Connect**.

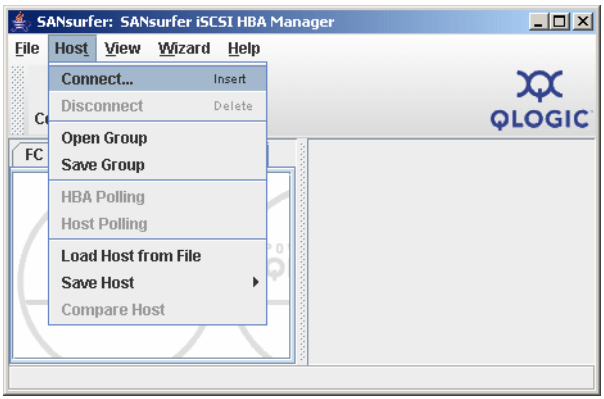

Figure 7-52 SANsurfer — Host → connect

3. Localhost is chosen by default. An IP address of a remote system where the agent is installed can be specified here as well. Click **Connect**.

Figure 7-53 SANsurfer — Connect

4. Open **HBA 0 QMC4052** by double-clicking on this host bus adapter. The two ports of this adapter are shown. Double-click the port 0. All targets seen by this port are listed. double-clicking the target shows the LUN of this target. and the first port (Port 0) the devices seen by this port are shown. Every logical drive configured for this host is seen as a separate device. The green color of the device indicates that there is an active connection from the HBA to the target. Blue is an unknown state or no active session and red is a failed session or the device is disappeared. Similar status is also seen for the port and the HBA. SANsurfer help explains the different status very well. See SANsurfer iSCSI HBA Manager Main Windows in the Getting started section.

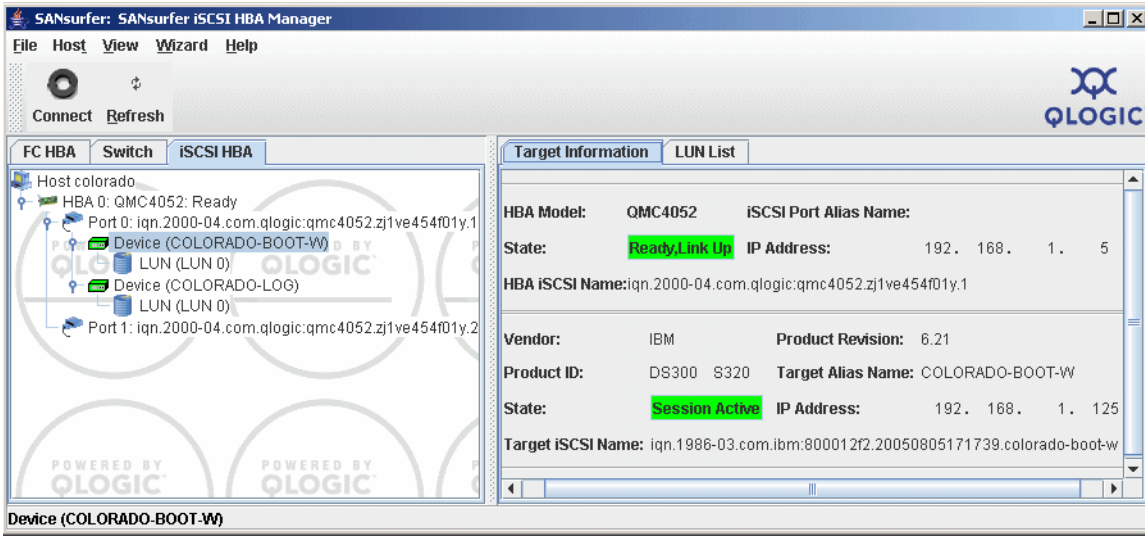

Figure 7-54 SANsurfer — Connected to host COLORADO

The second port does not see any devices because there are no devices configured for this port. This will be done in the next steps.

7.1.10 Configure iSCSI host bus adapter firmware settings

There are two settings which need to be changed in the iSCSI host bus adapter firmware. These settings are the same for the IBM iSCSI Server Adapter and the QLogic iSCSI Expansion Card for IBM @server BladeCenter and are required for the IP failover of the DS300 controllers. These are:

- ARP Redirect — set to enable
- Connection keep alive time-out — set to 180 (was 30)

To change these firmware settings perform the following steps:

1. From the SANsurfer utility select on the **iSCSI HBA** tab the first port of the iSCSI host bus adapter. Click the **Port Options** tab. On the Port options tab is a **Firmware** tab. Click the firmware tab. Click the **Open** button to edit **Configured Firmware Values**.

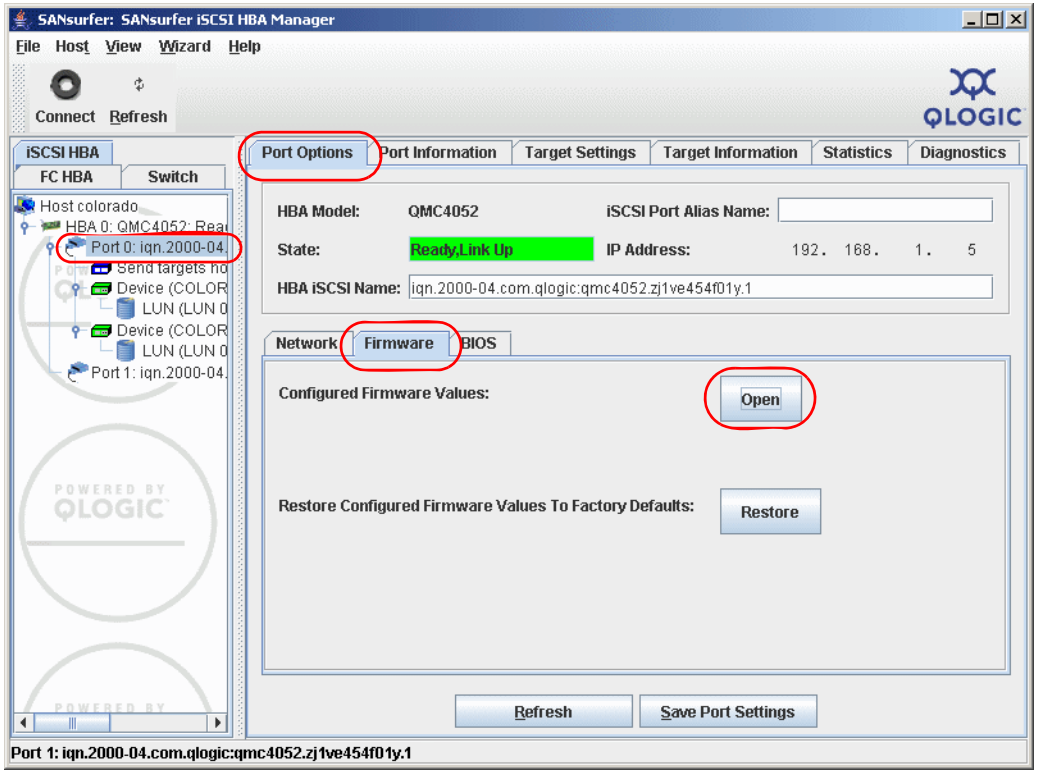

Figure 7-55 Select the port of the iSCSI adapter

2. A list with the current firmware settings is shown. The two settings which needs to be changed are highlighted in the figure. Click **Edit**.

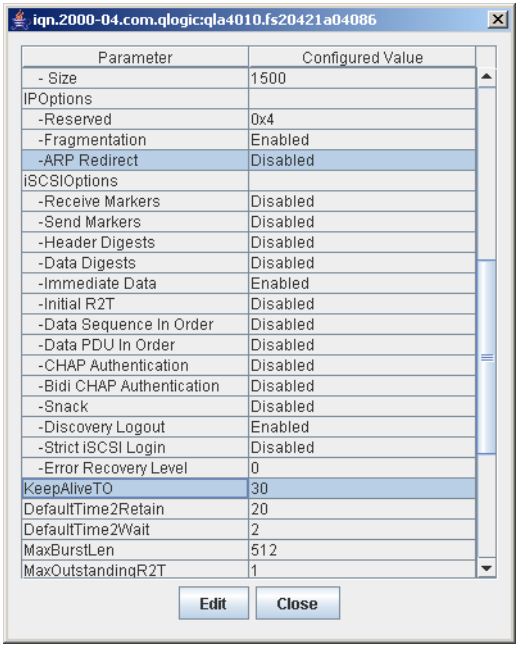

Figure 7-56 iSCSI adapter firmware settings

3. Scroll to the right site. Change the **Connection Keep Alive Timeout** from 30 to 180. Enable the **ARP Redirect** and click **OK**.

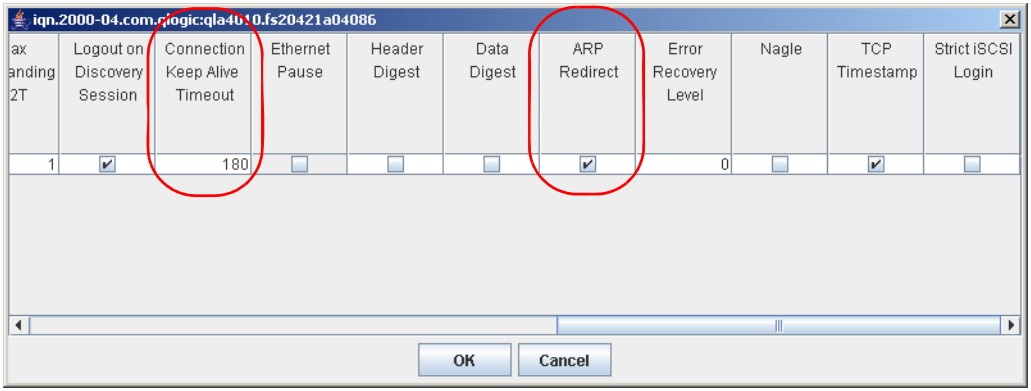

Figure 7-57 iSCSI adapter firmware settings

4. Click **OK**.

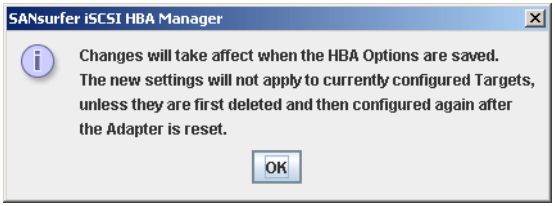

Figure 7-58 Change notification

5. The changed firmware settings are listed in Figure 7-59, but are not yet active. Click **Close**.

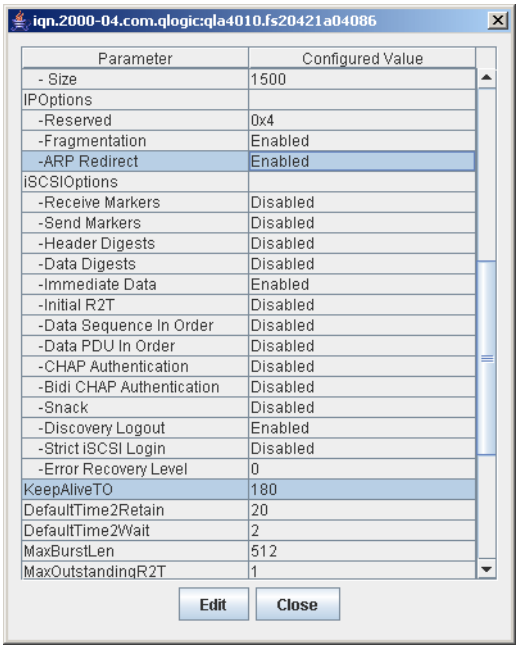

Figure 7-59 Modified firmware settings

6. Click **Save Port Settings** as shown in Figure 7-60 on page 264 to apply the firmware setting changes.

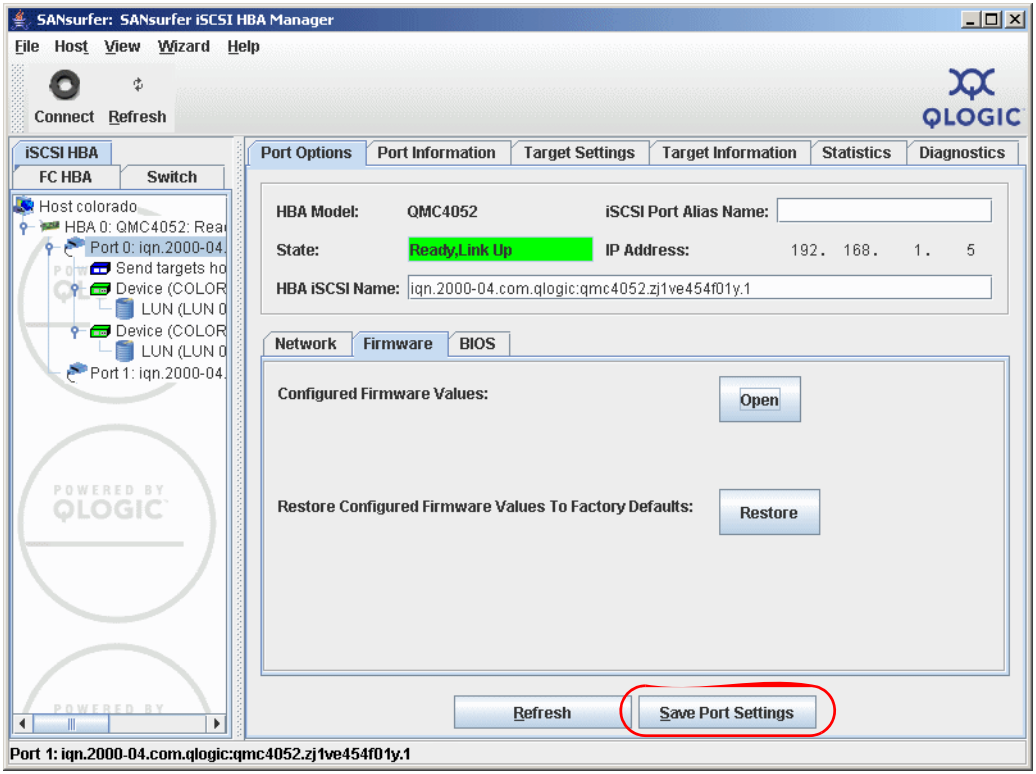

Figure 7-60 Save port settings

7. Click **Yes**.

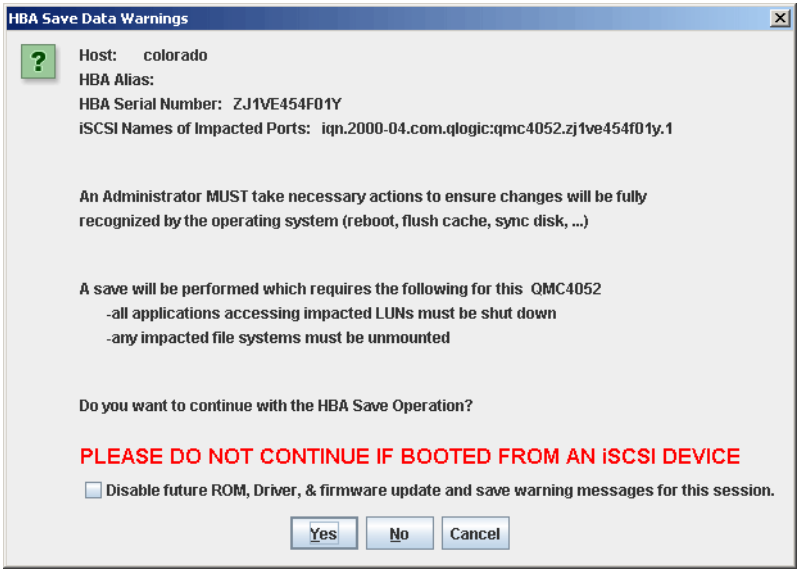

Figure 7-61 Save data warning

8. Enter the password and click **OK**. The default password is **config**.

Figure 7-62 Change notification

9. Repeat this steps for the second port of the host bus adapter.

7.1.11 Configure the second port of the iSCSI host bus adapter

The second port requires an IP address. This can be done during the POST BIOS menu of the adapter as it was shown in 7.1.4, “iSCSI host bus adapter setup and installation preparation” on page 228. SANsurfer provides more options to configure the host bus adapter than the POST BIOS menu of the host bus adapter. Two logical drives of the TotalStorage DS300 will be added to the host COLORADO: one logical drive without authentication and one with.

1. Start SANsurfer and connect to the host.
2. Configure the network address of the second port of the iSCSI host bus adapter installed in the HS20 blade

a. Select the second port of the HBA. The second port is port 1. From the right column select the **Port Options** tab. This port does not have an IP address configured.

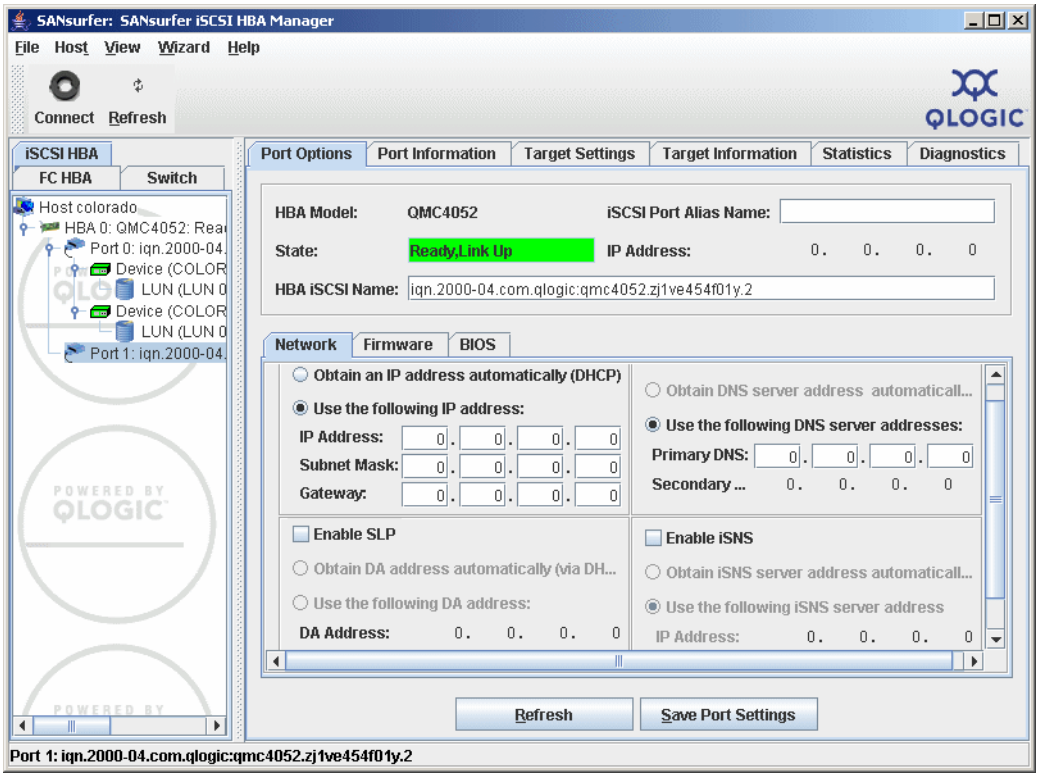

Figure 7-63 SANsurfer — Port1 - port options

b. Enter in the IP address field the IP address 192.168.1.69. Enter the Subnetmask 255.255.255.192 in the Subnet Mask field. Then click **Save Port Settings**.

Note: You may see a BSoD when the network settings are applied to the adapter. Applying network addresses to the adapter is resetting this adapter and this interrupts the access of the operating system to the C drive. It is safer to perform this configuration from the adapter BIOS.

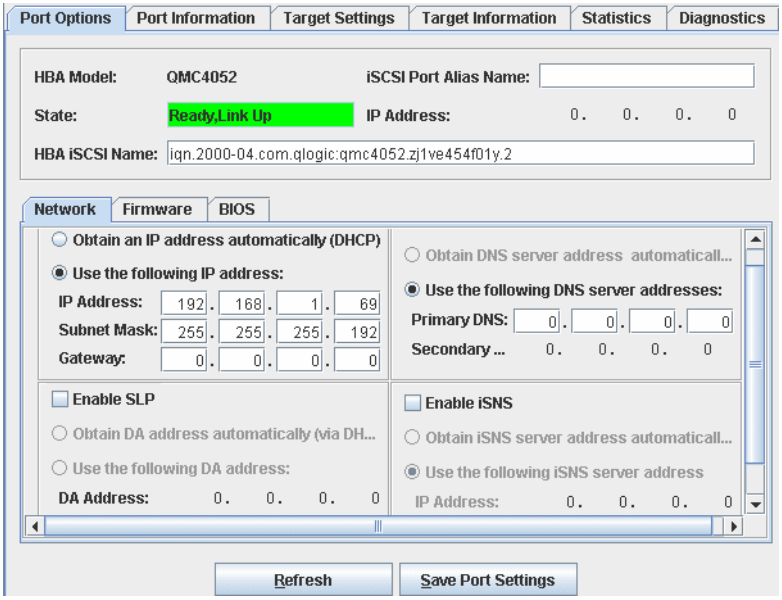

Figure 7-64 SANsurfer — Port1 - network settings

c. The Diagnostic tab allows to send a ping to an iSCSI target to verify the settings and the network connect.

3. Configure the access control list for the logical drive `COLORADO-MAIL`.

a. Use the ServeRAID manager to configure the access control list. Add the logical drive COLORADO-MAIL to the initiator alias COLORADO-P1.

logical manage
`COLORADO-MAIL`
allow
`COLORADO-P1`

4. Add a new iSCSI target to the host without initiator authentication.

a. Select the **Target Settings** tab. No targets are currently shown. Click the green plus sign to add a iSCSI target.

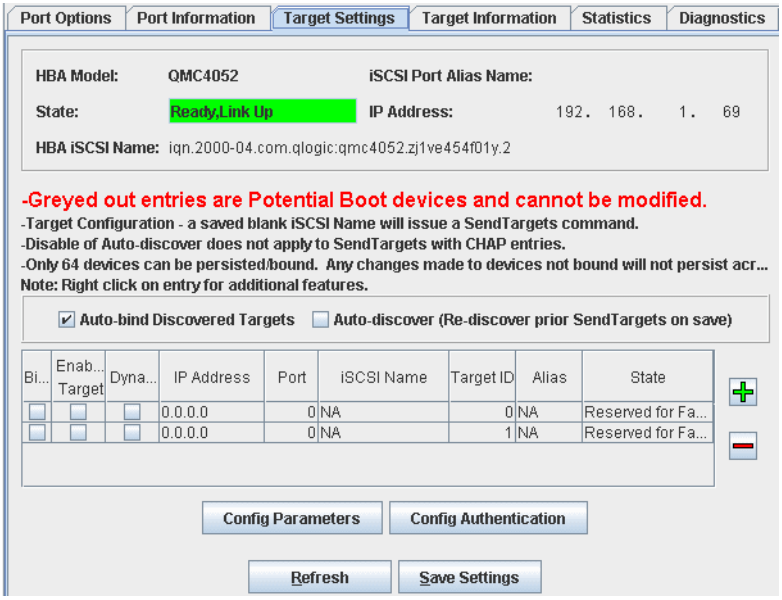

Figure 7-65 SANsurfer — Port1 - Target settings

b. Enter the iSCSI target IP address `192.168.1.126`. In this case it is the IP address of ETH3 of controller B. The HBA is connected to ETH3 on both controllers and the array of the logical drive `COLORADO-MAIL` is owned by controller B. Then click **OK**.

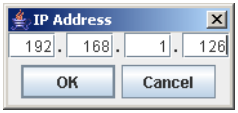

Figure 7-66 SANsurfer — iSCSI target IP address

c. The target is added to the target list. This connection is not yet saved. Click **Save Settings**.

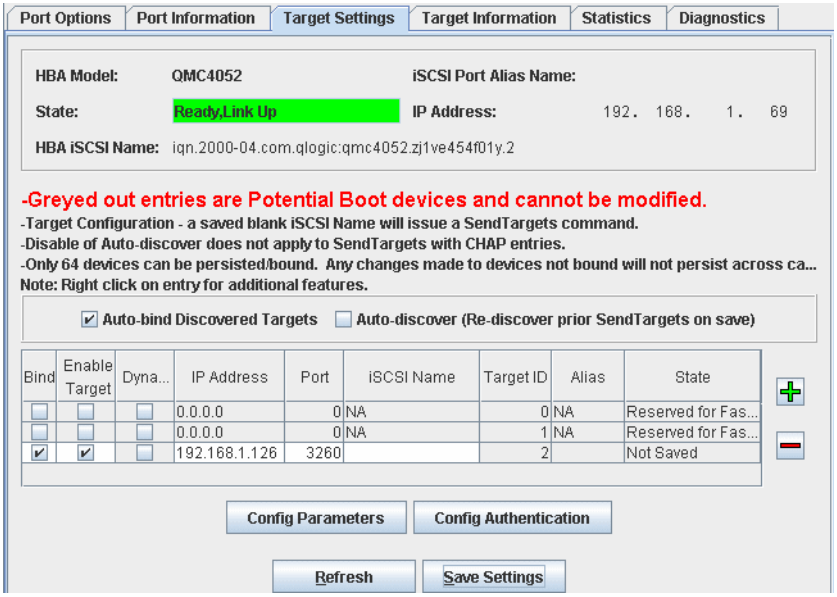

Figure 7-67 SANsurfer — iSCSI target list

d. A warning appears that the settings should not be saved when this device is used for booting. In this case the other port is used. Click **Yes**.

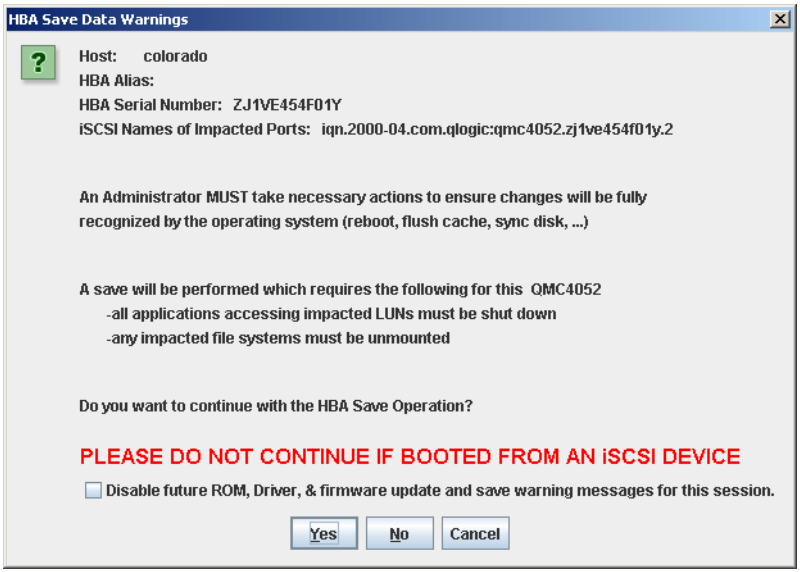

Figure 7-68 SANsurfer — Save settings warning

e. After refreshing the adapter configuration, a new target is seen under the HBA port 1. It is the blue Send targets host. Click the Target information tab. Then open the target list box. Only target ID 2 is shown. This is the Send targets host target. This target is used to retrieve a list of targets available on the TotalStorage DS300.

> **Note:** The Send targets host target will not show an IQN or session active. It will be always show as an unknown target or not active session.

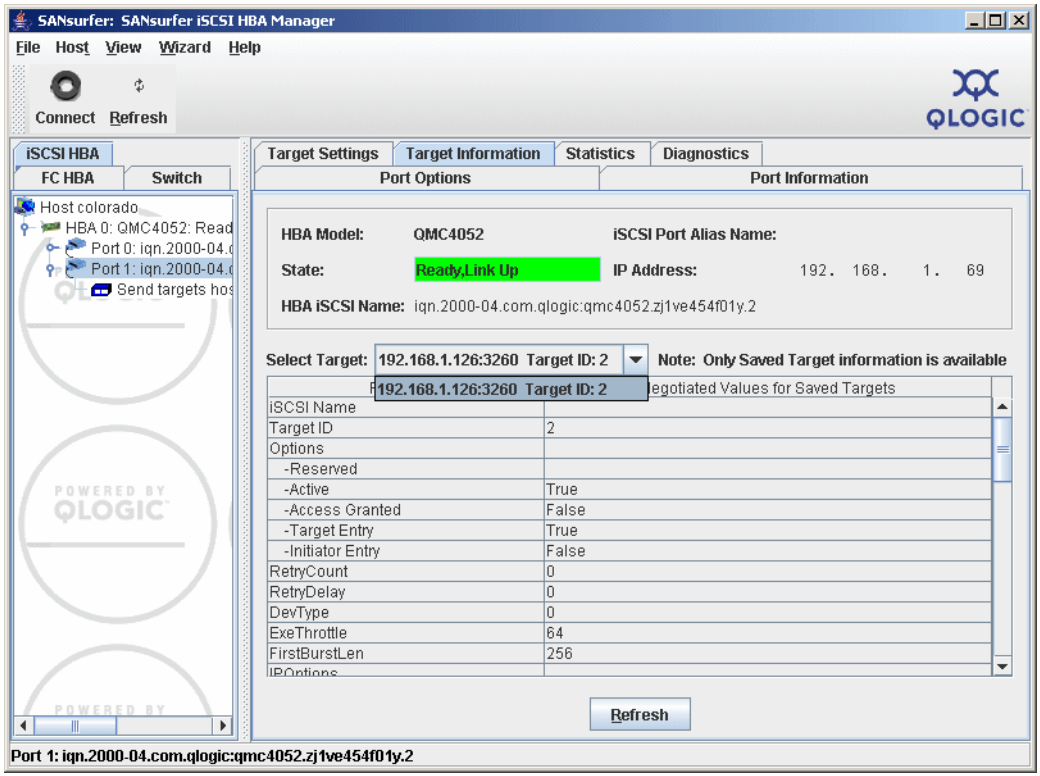

Figure 7-69 SANsurfer — Target information

f. After a short time SANsurfer asks for an refresh of the configuration information. Click **Yes** to refresh the configuration. The new target COLORADO-MAIL will be found.

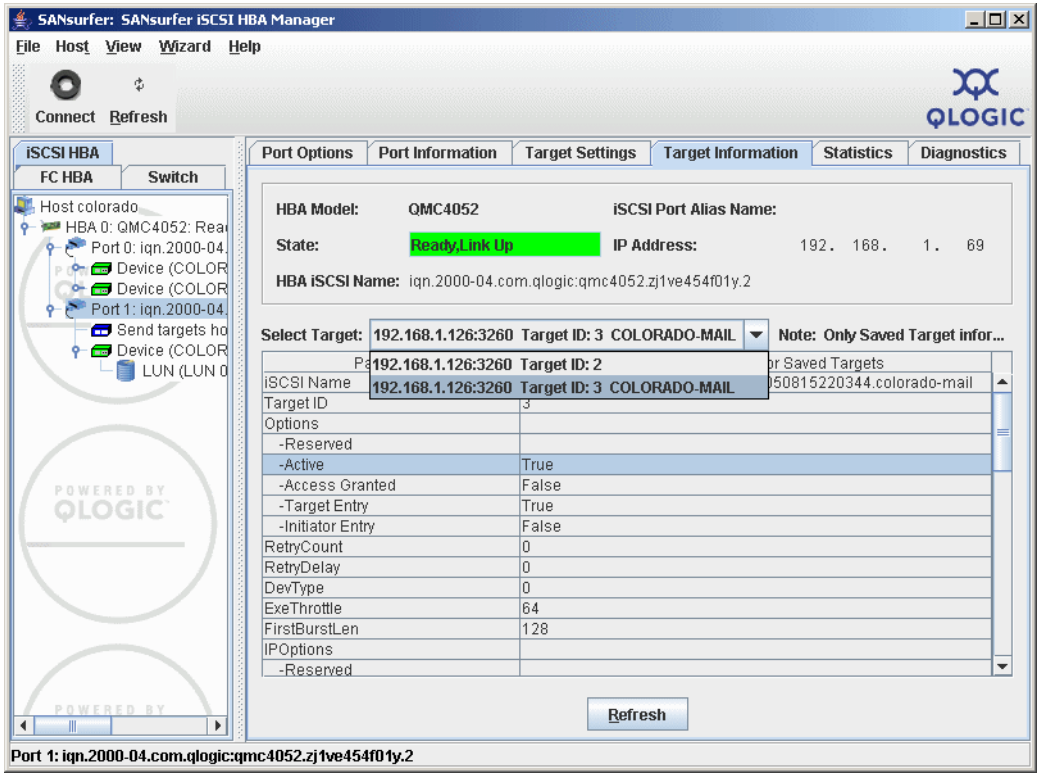

Figure 7-70 SANsurfer — Target information

After a hardware rescan or reboot, the logical drive is seen by the operating system.

The next steps will add a logical drive which can be accessed only with the proper credentials.

1. Create a user and modify the access control list.
2. Add logical drive COLORADO-DB to the host. Use initiator authentication.
   a. Select the port 1 of the iSCSI HBA and from the right column the **Target Settings** tab. No targets are currently shown.Click the green plus sign to add a iSCSI target.

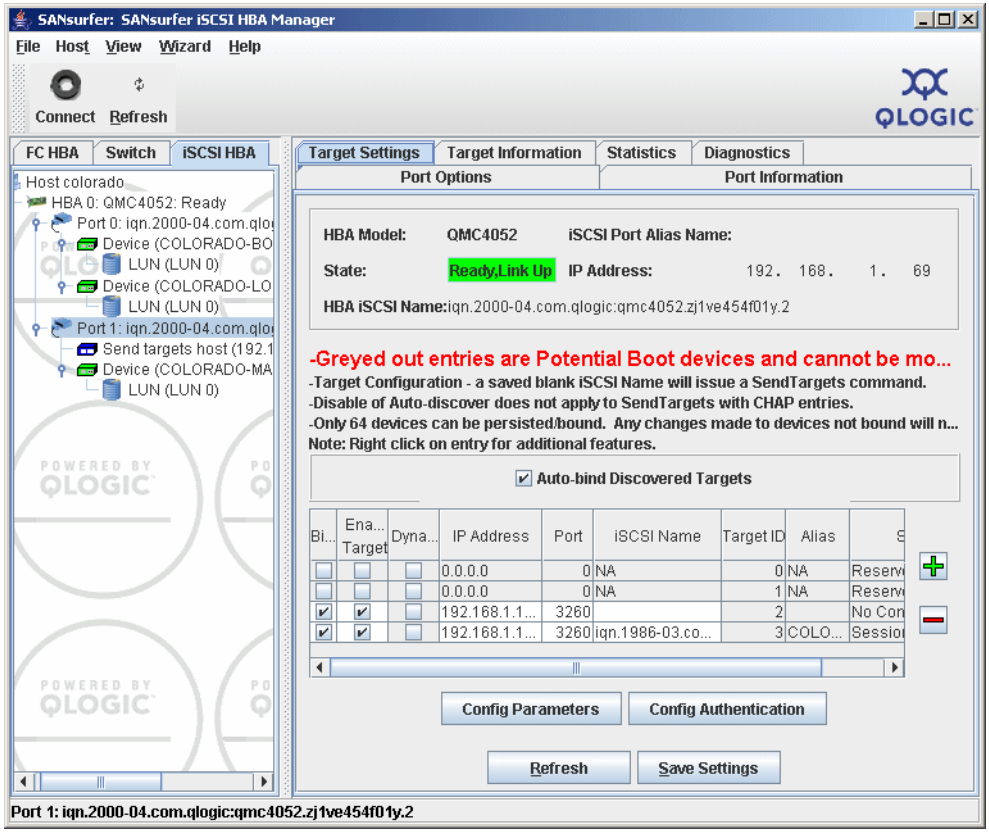

Figure 7-71 SANsurfer — Port1 - Target settings

b. Enter the iSCSI target IP address `192.168.1.126`. In this case it is the IP address of ETH3 of controller B. The HBA is connected to ETH3 on both controllers and the array of the logical drive `COLORADO-MAIL` is owned by controller B. Then click **OK**.

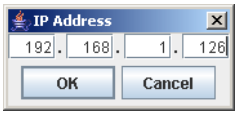

Figure 7-72 SANsurfer — iSCSI target IP address

c. The target is added to the target list. This connection is not yet saved. Click **Configure Authentication**.

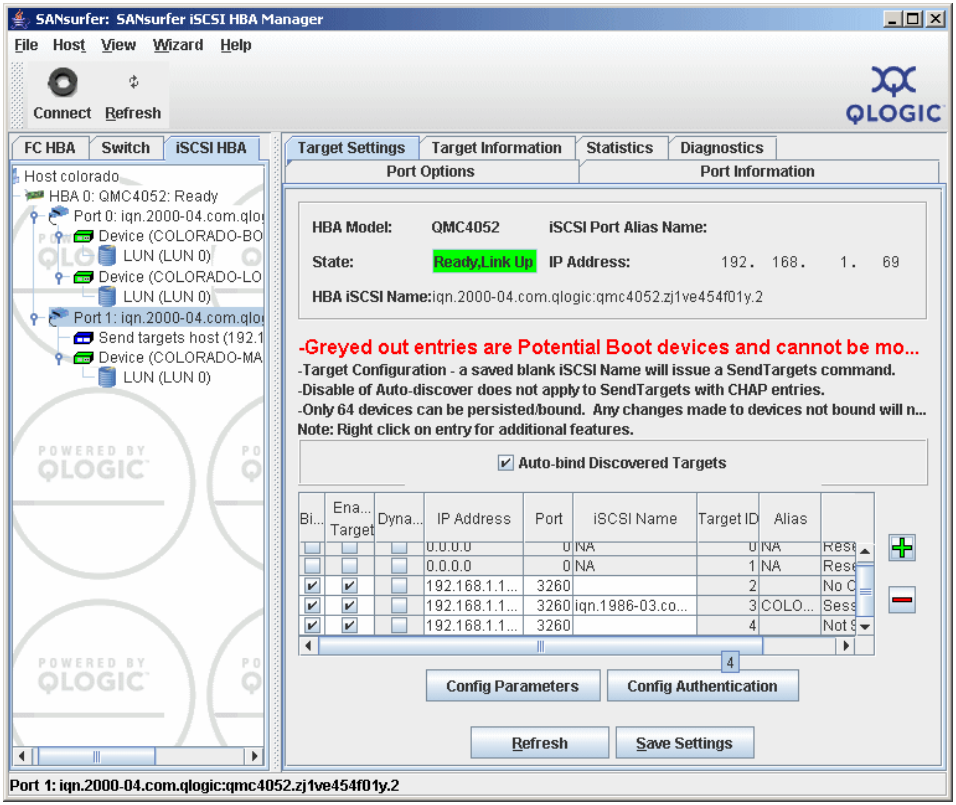

Figure 7-73 SANsurfer — iSCSI target list

d. Enter the password to change the authentication settings. The default password is config.

Figure 7-74 SANsurfer — Enter configuration password

e. Click the green plus sign to create a new CHAP entry. Enter the name of the initiator. Here was the IQN `iqn.2000-04.com.qlogic:qmc4052.zj1ve454f01y.2` of the port 1 chosen. In the initiator secret field enter at least 12 character long password.The password is `a1234567890b`. Then select the last target from the target list. It is the target with TID4. Select the created CHAP name and secret from the **CHAP name/secret** drop down list. Click **OK**.

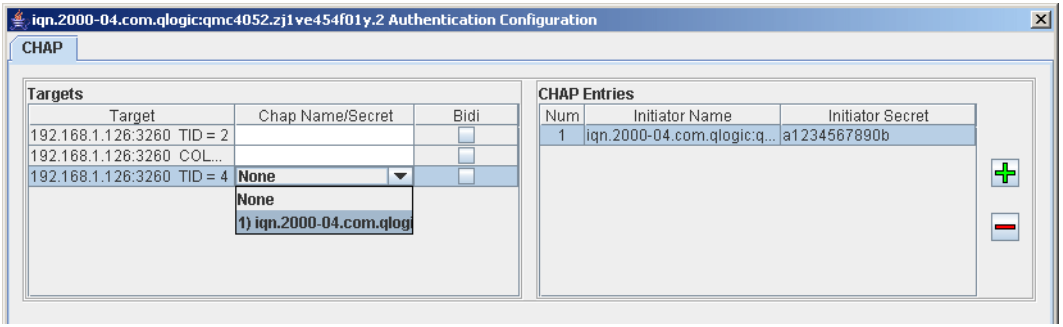

Figure 7-75 SANsurfer — CHAP name and secrets configuration

f. Click **Save Settings** to save all the target modifications.

g. A warning appears that the settings should not be saved when this device is used for booting. In this case the other port is used. Click **Yes**.

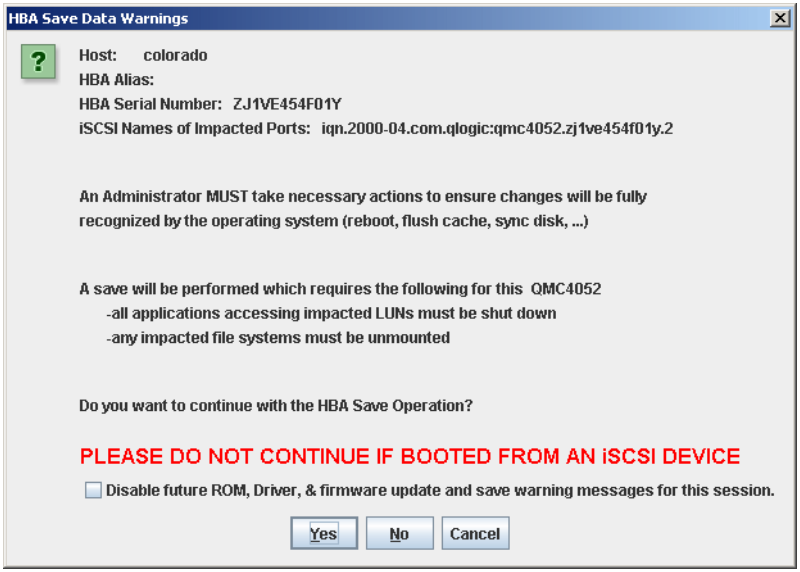

Figure 7-76 SANsurfer — Save settings warning

h. Enter the configuration password. The default password is config.

Figure 7-77 SANsurfer — Configuration Password

i. After refreshing the adapter configuration a new target is seen under the HBA port 1. Click **Yes** to refresh the configuration.

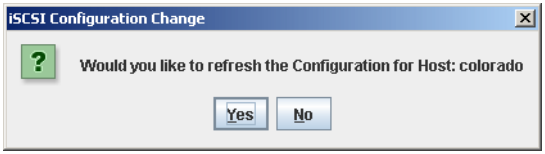

Figure 7-78 SANsurfer — Refresh configuration

j. The configuration shows now the new added iSCSI target `COLORADO-DB`.

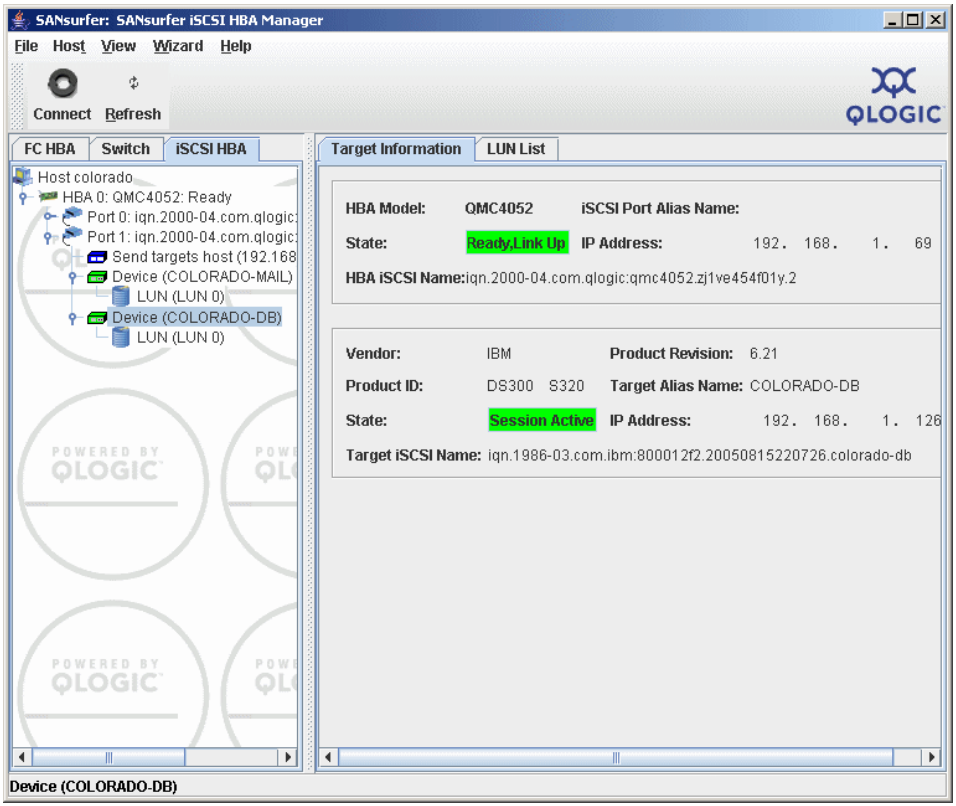

Figure 7-79 SANsurfer — New iSCSI target COLORADO-DB configured

After a hardware rescan or reboot the logical drive is seen by the operating system.

More logical drives may be added now. The setup is now completed and the testing of the setup can start.

7.2 Qlogic iSCSI Expansion Card for IBM eServer BladeCenter in Microsoft Windows Server 2003

The QLogic iSCSI Expansion Card for IBM @server BladeCenter can be used with an IBM HS20, LS20 and HS40 blade in a IBM @server BladeCenter. At time of writing only Microsoft Windows 2000 and Windows Server 2003 are supported. Please check IBM ServerProven® and driver availability when other operating systems are required.

ServerProven:

http://www.ibm.com/servers/eserver/serverproven/compat/us/

IBM Personal computing support:

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

Finish the operating system installation of the IBM Blade Server. Shutdown the the blade and install the QLogic iSCSI Expansion Card for IBM @server BladeCenter adapter in the blade.

After finishing the hardware installation of the iSCSI adapter, power on the blade and start the operating system. After logon you will be presented with a few Found New Hardware wizards. Before you start with updating the drivers for the iSCSI adapter, have a look at the Device Manager. In the Device Manager, four devices are listed in the **Other devices** section. These four devices are the iSCSI adapters.

In the Device Manager under Unknown Devices are listed two ethernet controllers and two network controllers. The four devices are 4 functions of the same PCI device. First ethernet controller is function 0 and second ethernet controller is function 2. Network controllers have the function 1 and 3 assigned. To use the iSCSI adapter like a regular network adapter this device can be used to access the iSCSI adapter from the operating system as a NIC. The network controller are used as SCSI devices and used for the iSCSI traffic.

There are different drivers available for this adapter:

- SCSI Miniport driver for Windows 2000 and Windows Server 2003
- SCSI Storport driver for Windows Server 2003
- An NDIS driver for Windows 2000 and Windows Server 2003

You should use the storport driver for Windows Server 2003 and the miniport driver for Windows 2000.

An infrastructure service host is used to provide the iSCSI simple naming services. This host is a multihomed host, connected to the iSCSI networks and to

the management network. Microsoft Internet Storage Name Server 3.0 was chosen for the iSNS server. The iSNS server needs to be defined on the DS300 with the command line interface. IBM ServeRAID Manager does not provide the capability to define this setting. The iSNS Server and the ServeRAID Manager management station service might be run on the same system.

The following diagram shows the cabling of the components. It follows the example configuration as explained in 6.2.2, “DS300 dual controller switched dual path attachment” on page 219.

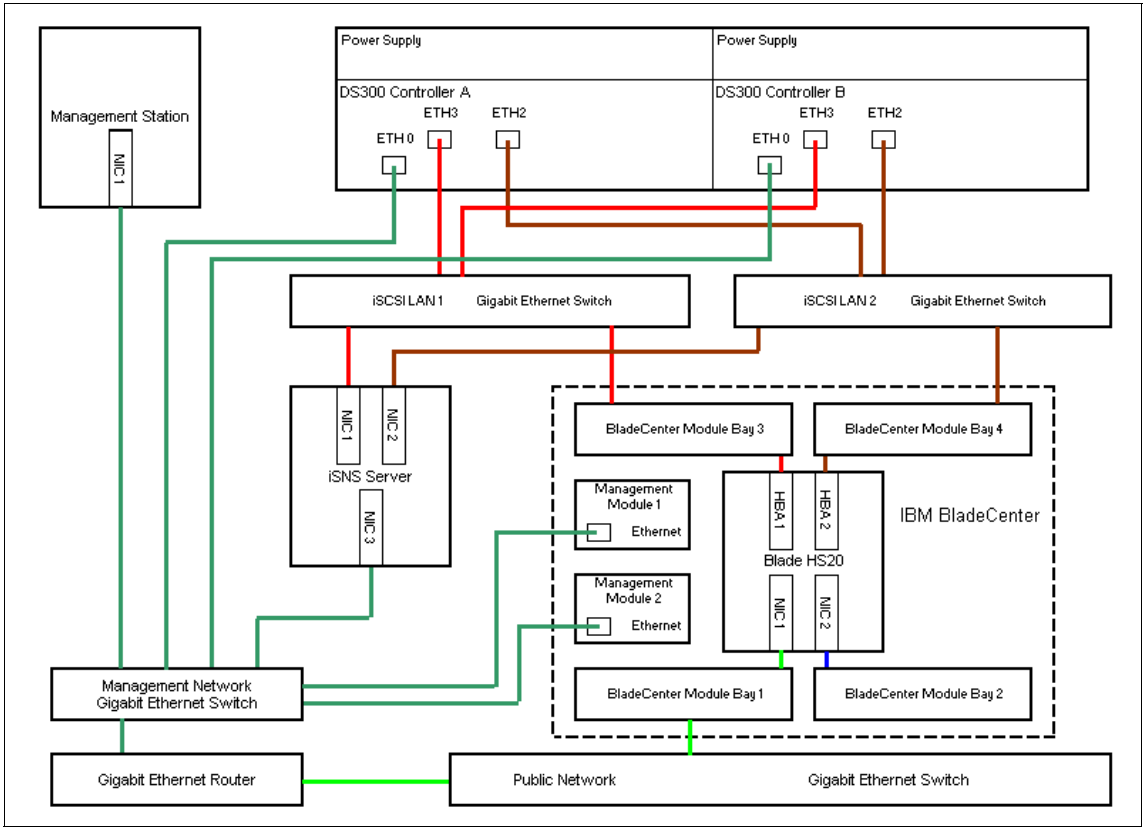

Figure 7-80 Configuration setup

7.2.1 Install the device drivers for the QLogic iSCSI Expansion Card for IBM eServer BladeCenter

For the card there will appear four devices in the **Other devices** tree of the Device Manager:

- Two ethernet controller (function 0 and 2)
  - These devices refer to the gigabit ethernet adapter functionality of the QLogic iSCSI Expansion Card.
  - After driver installation they will appear in the **Network adapters** tree of the device manager.
- Two network controller (function 1 and 3)
  - These devices refer to the iSCSI adapter functionality of the QLogic iSCSI Expansion Card.
  - After driver installation it will appear in the **SCSI and RAID controllers** tree of the device manager.

The drivers can be found via the following URL:

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-60280

To install the drivers for the QLogic iSCSI Expansion Card follow these steps:

1. Open the **Device Manager**. Right-click the first **Ethernet Controller** in the **Other devices** tree. Click **Update Driver...**.

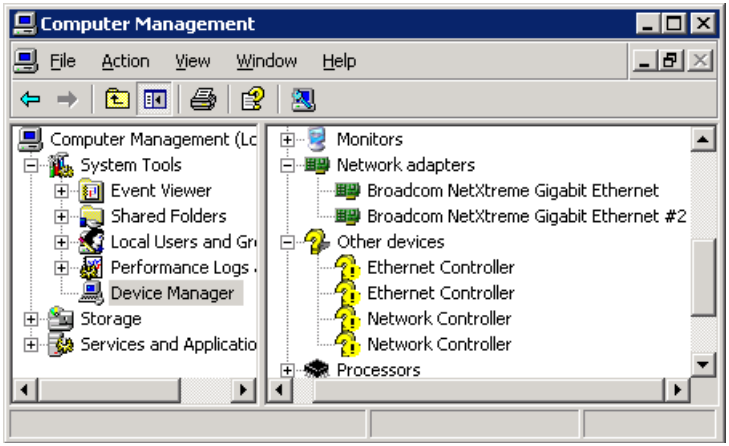

Figure 7-81 Windows device manager with unknown devices (BladeCenter iSCSI HBA)

2. In the **Welcome to the Found New Hardware Wizard** select **No, not at this time** and click **Next**.

Figure 7-82 Found new hardware wizard

3. Select **Install from a list or specific location (Advanced)** and click **Next**.

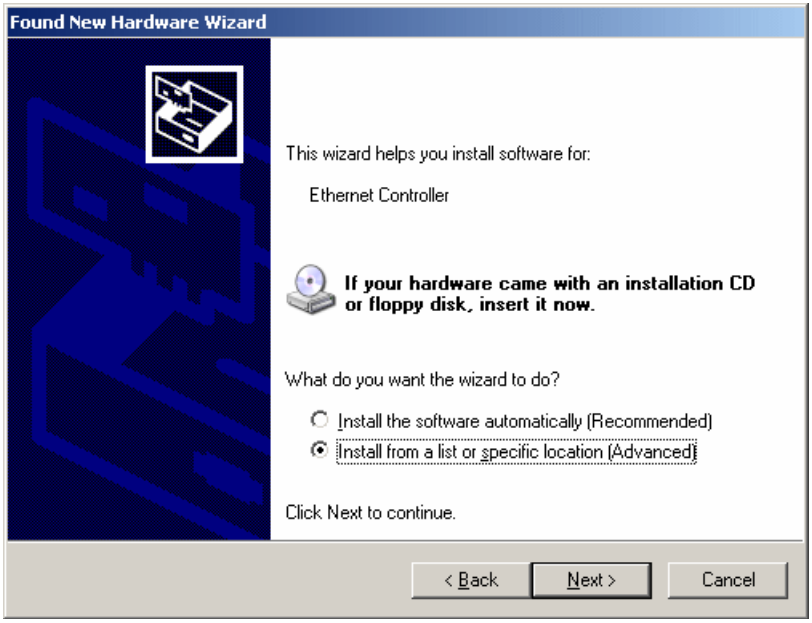

Figure 7-83 Install from a list

4. Browse to the folder where you have downloaded the NDIS network driver for the card. Click **Next**.

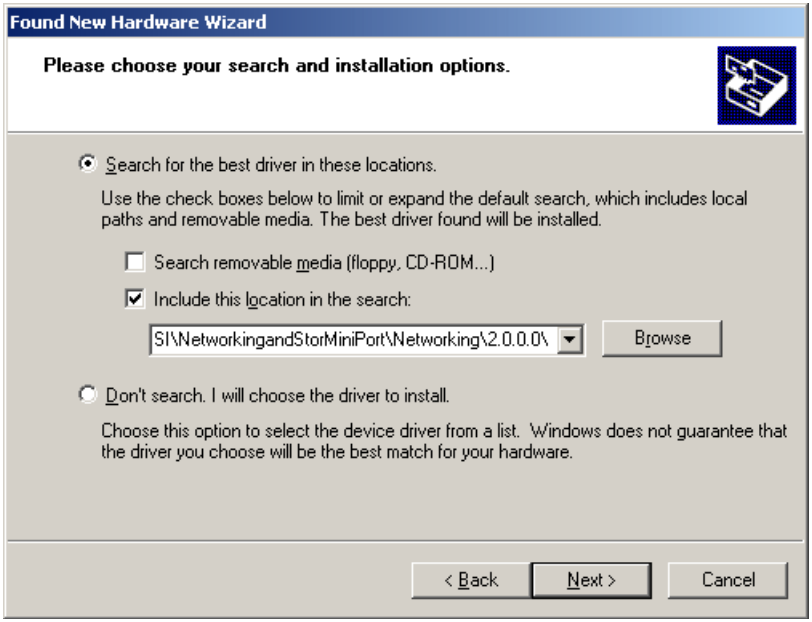

Figure 7-84 Search location

5. After the driver installation finishes, click **Finish** to close the Hardware Update Wizard.

Figure 7-85 Hardware update completed

6. In the **Network adapters** tree of the device manager appears a **QLogic QLA4010 PCI 1Gb Ethernet Adapter**. Redo steps 1-5 for the remaining ethernet controller in the **Other devices** tree.

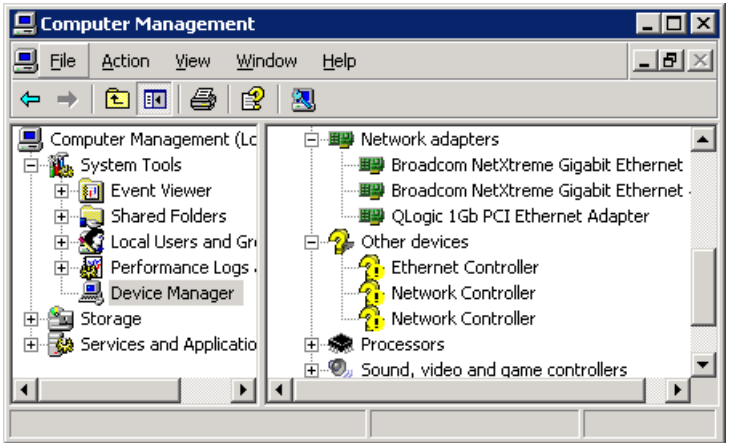

Figure 7-86 Windows device manager with unknown devices (network controller)

7. Right-click the first **Network Controller** in the **Other devices** tree. Click **Update Driver...**.

8. In the **Hardware Update Wizard** choose **No, not this time** and click **Next**.

Figure 7-87 Found new hardware wizard

9. Select **Install from a list or specific location (advanced)** and click **Next**.

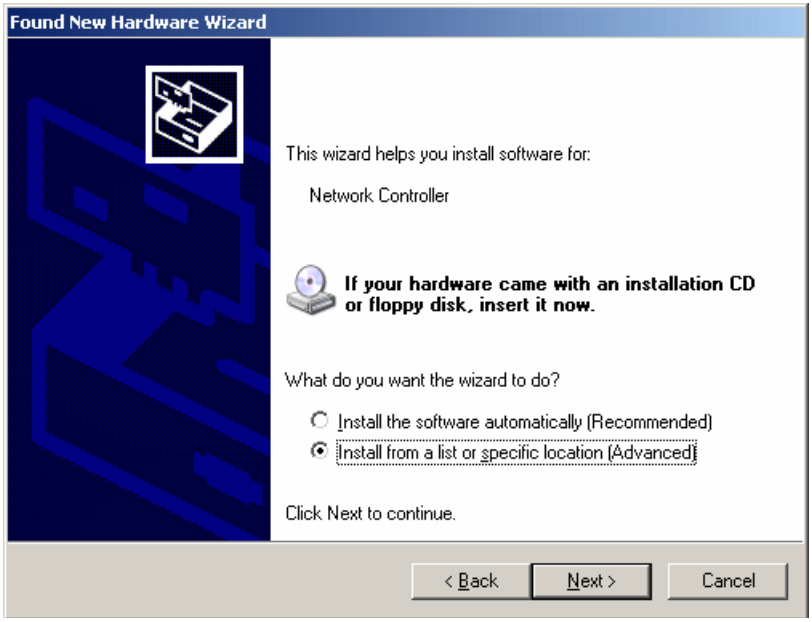

Figure 7-88 Install from a list

10. Deselect the **Search removable media option,** select **Include this location in the search** and specify the directory were you have stored the device driver for the iSCSI adapter. Click the **Next** button to start the installation.

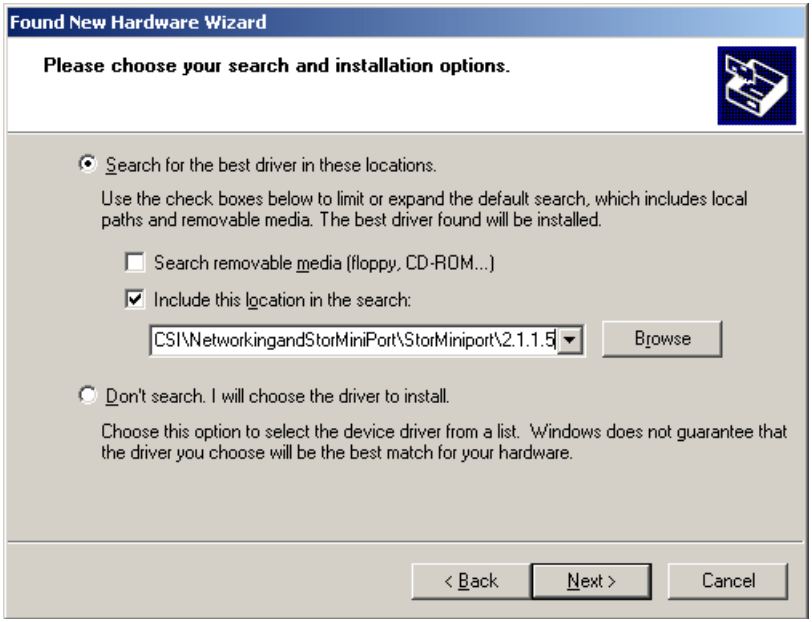

Figure 7-89 Search location

The driver files will be copied to the correction location.

11. Click the **Finish** button to complete the installation if the first iSCSI adapter.

Figure 7-90 Hardware update completed

When you check now the Windows device manager, you will find a new SCSI adapter device **QLogic iSCSI Adapter** in the **SCSI and RAID controllers** section and one unknown device less in the **Other devices** section.

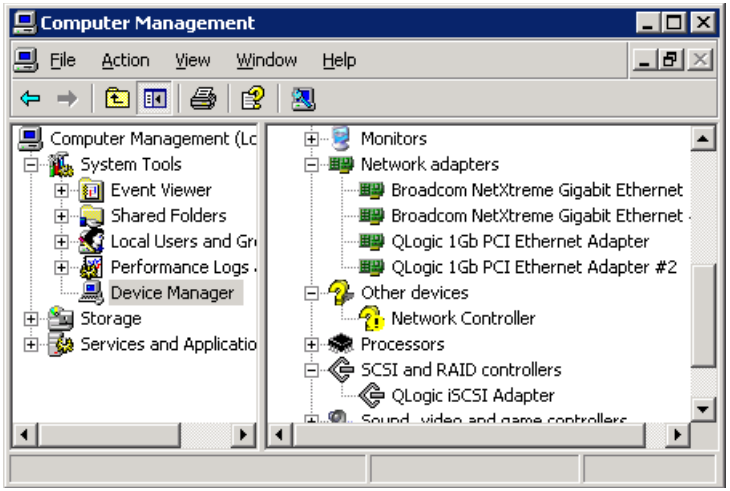

Figure 7-91 Windows device manager with unknown devices (network controller)

12. After finishing the driver installation for the first network controller of the iSCSI card, the **Welcome to the Found New Hardware Wizard** will start again**.** Repeat the steps 6 to 9 to install the driver for the second network controller of the iSCSI card.

You have now completed the installation of the device driver for the iSCSI card. The next step describes the installation of the management software for the iSCSI card.

7.2.2 Install the Qlogic iSCSI Expansion Card for IBM eServer BladeCenter management software

To manage the QLogic iSCSI Expansion Card for IBM @server BladeCenter, a software package called iSCSI SANsurfer Management is used. Management of the iSCSI adapter includes configuration and updating the adapter. This management software consists of a agent installed on the system where the iSCSI adapter is installed and a client or graphical user interface (GUI). The GUI can be installed on the same system or on a system used for management purposes. Multiple agents might be managed by a single client or GUI. Agent and GUI are currently only for Microsoft Windows available. This document uses the version 4.01.03 of the iSCSI SANsurfer Management software to demonstrate installation and configuration. To download the management software use this URL:

`http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-60280`

The following steps explain the installation and configuration of the iSCSI SANsurfer Management software.

1. Run the installer of the management software by executing the downloaded file.

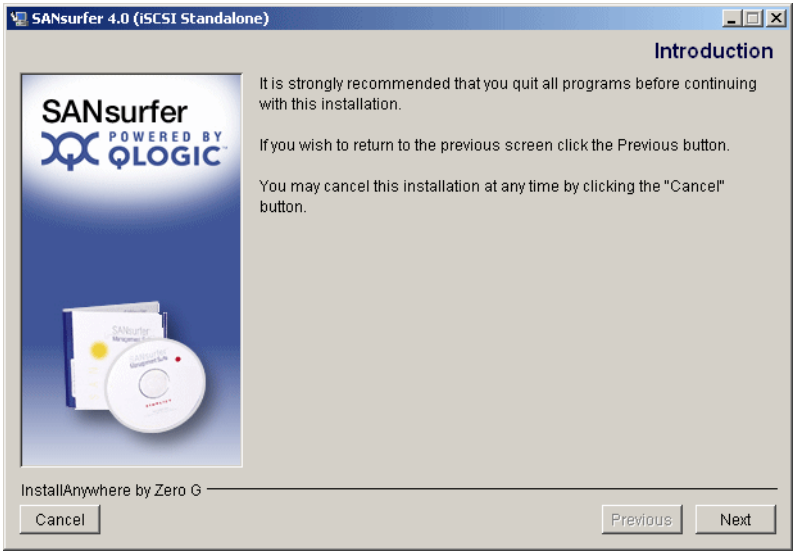

Figure 7-92 SANsurfer — introduction

2. The introduction screen gives you some advice about the installation. Click the **Next** button to proceed.

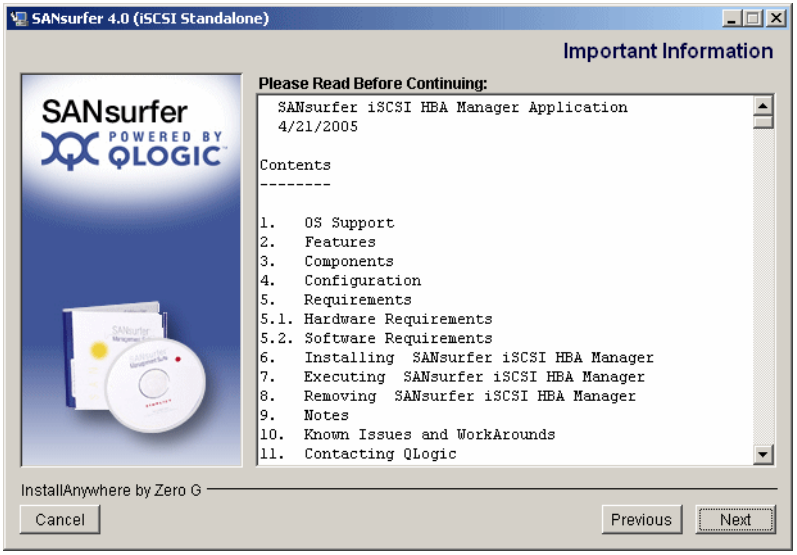

Figure 7-93 SANsurfer — Important Information

3. Read the important information about the management application and then click **Next**.

Figure 7-94 SANsurfer — Choose Product Features

4. Select the options you want to install. The installer allows to install GUI and agent alone or together. Depending on you planning select the option. The example below installs GUI and agent together on the same system. Click the **Next** button.

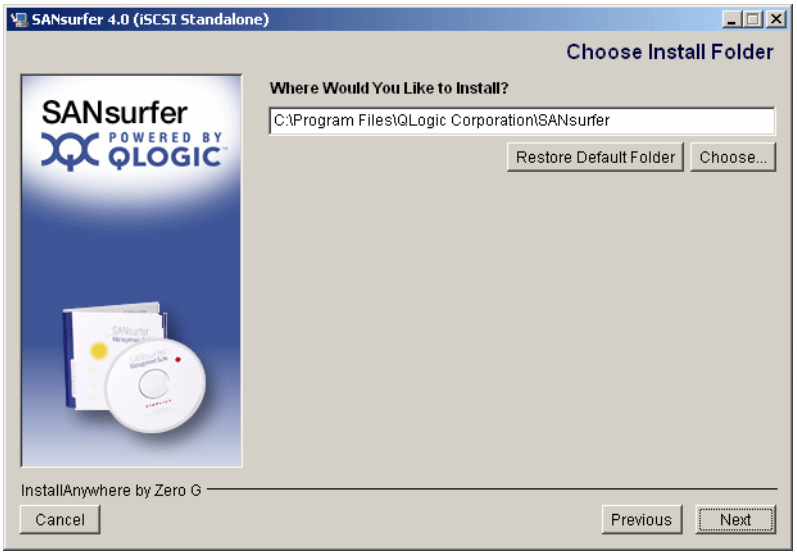

Figure 7-95 SANsurfer — Choose Install Folder

5. Accept the recommended installation location with the **Next** button.

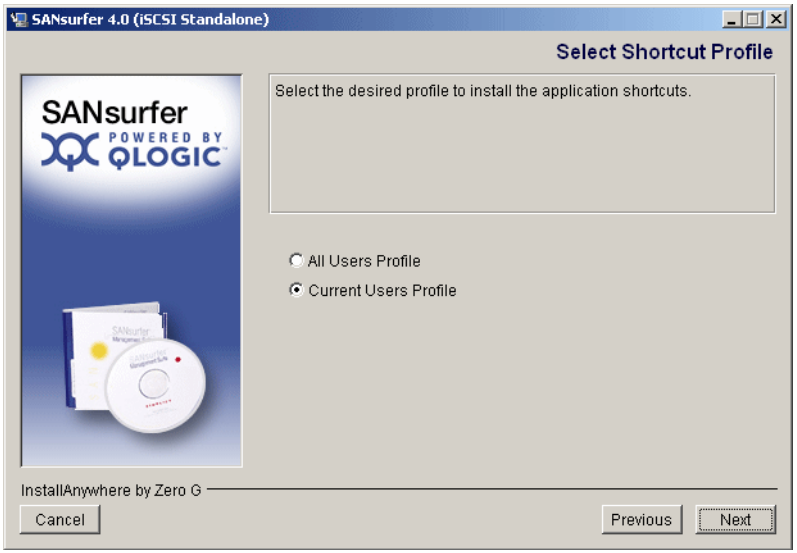

Figure 7-96 SANsurfer — Select Shortcut Profile

6. Menu entries and desktop shortcuts are placed by default into the current users profile. Select **All User Profiles** and click **Next**.

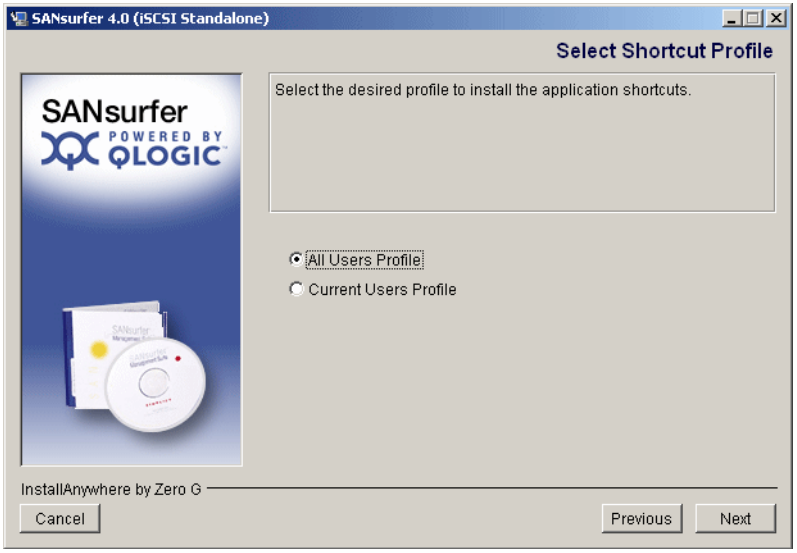

Figure 7-97 SANsurfer — Select Shortcut Profile

7. Click Next to let the installer create a desktop icon.

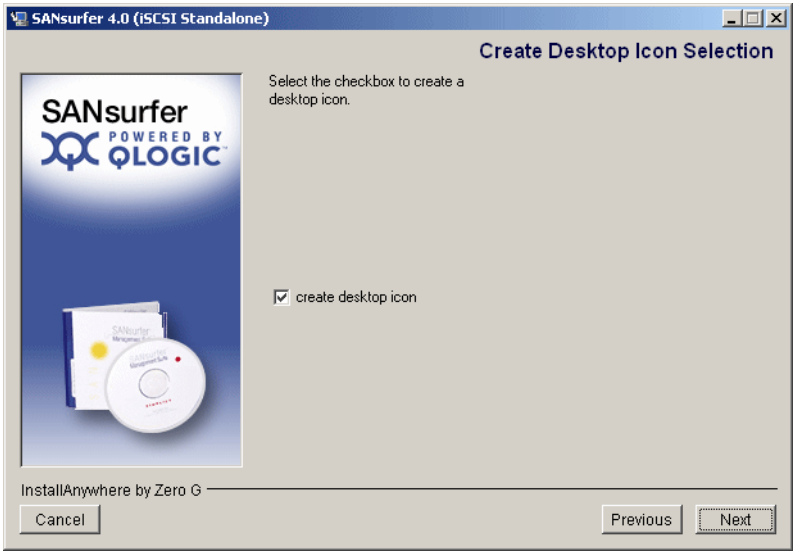

Figure 7-98 SANsurfer — Create Desktop icon Selection

8. Click **Install** in the **Pre-Installation Summary**.

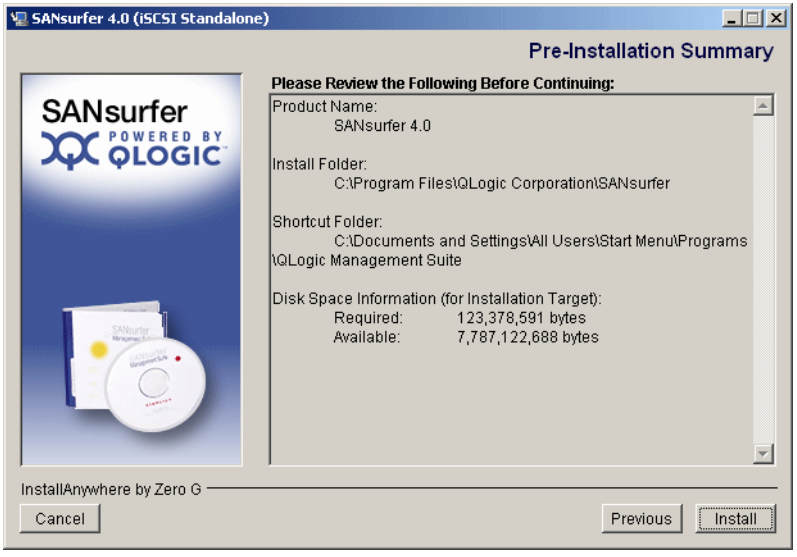

Figure 7-99 SANsurfer — Pre-installation Summary

All files will be installed now.

9. Click **Done** to close the installation application.

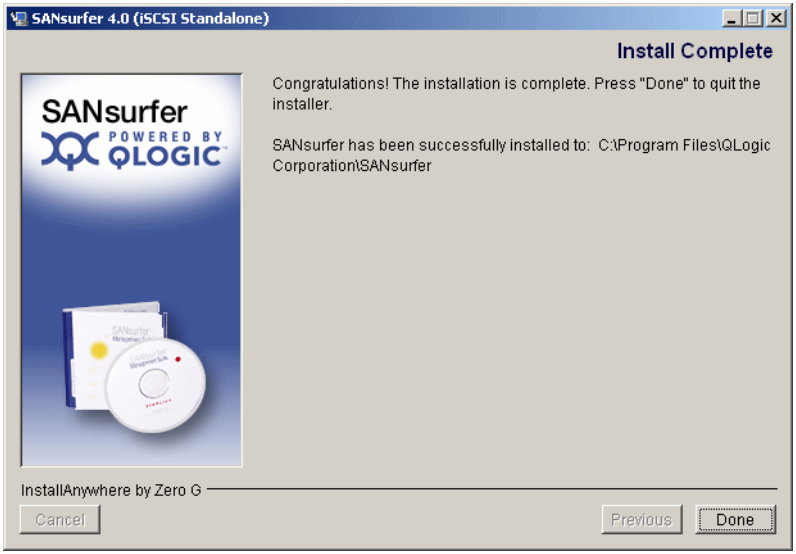

Figure 7-100 SANsurfer — Install Complete

The iSCSI SANsurfer Management Software for Windows is now installed and ready for use. The next section explains the basic setup of the Qlogic iSCSI Expansion Card for IBM @server BladeCenter.

7.2.3 Setup of the QLogic iSCSI Expansion Card for IBM eServer BladeCenter

The QLogic iSCSI Expansion Card for IBM @server BladeCenter can be configured in the **CTRL+Q** BIOS during system POST or with the iSCSI SANsurfer Management Software for Windows. This section explains the setup using the management software. The same options are available in the CTRL+Q BIOS Setup as well as in the management software. There are minor differences between those two possibilities of configuration. The update of BIOS, firmware and ROM of the iSCSI adapter can be only performed in the management software. Formatting or verifying of attached iSCSI drives are done in the CTRL+Q BIOS Setup.

Because of the redundant design of an IBM BladeCenter daughter cards like the Qlogic iSCSI Expansion Card for IBM @server BladeCenter are dual port adapter. The next sections will explain different ways to configure the adapter by using the two available ports.

- Configuring the IP Address and a target by using the configuration wizard. The target will be configured manually.

- The second port will be configured without the usage of the wizard and with help of an iSNS service.

Setup of Port 0 with the wizard

1. Start the iSCSI SANsurfer Management Software by using the **SANsurfer** desktop icon or use the start menu by clicking **Start** → **Programs** → **QLogic Management Suite** → **SANsurfer**. In SANsurfer GUI click the **Connect** button.

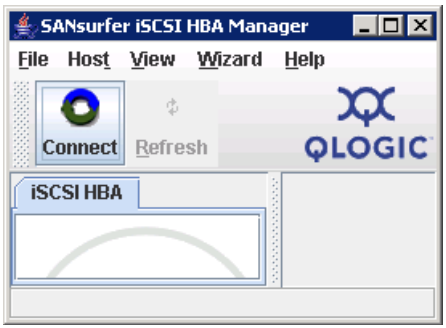

Figure 7-101 SANsurfer Manager — unconnected

2. In the **Connect to Host** window select `localhost` or enter the IP address of the host where the agent was installed. CLick **Connect**.

Figure 7-102 Connect to host window

3. SANsurfer connects to the specified host and presents a list of installed HBAs. The left site of the window contains host information.

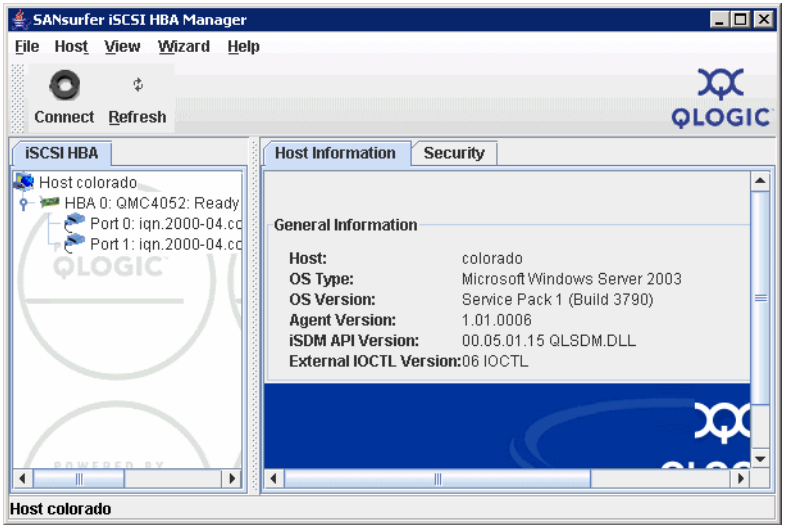

Figure 7-103 iSCSI HBA and host information

4. Select in HBA list a HBA and verify that firmware and BIOS level is up to date.

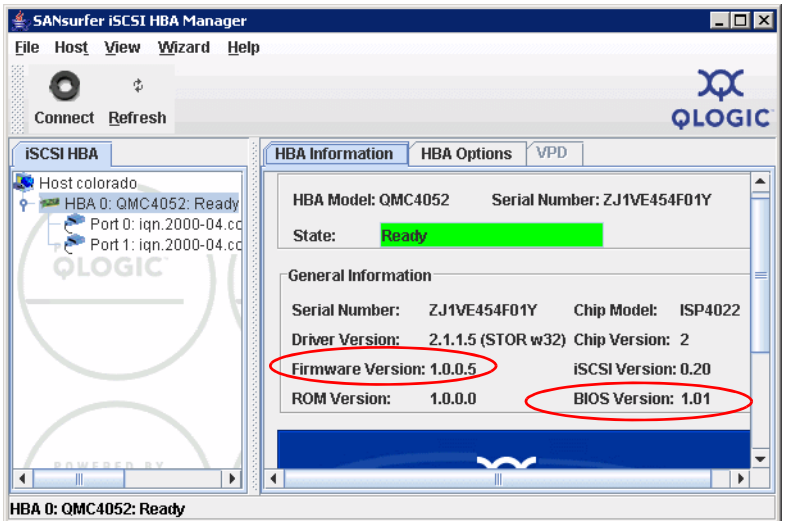

Figure 7-104 HBA information

5. Select the register card HBA Options when you have to perform a BIOS or firmware update. Buttons for updating the firmware and the BIOS are available.

6. A new adapter does not have a configuration on it. It is required to define the IP network settings first before the adapter can be used to access iSCSI targets. Click the menu **Wizard** → **General Configuration Wizard...** or press **Ctrl+G**.

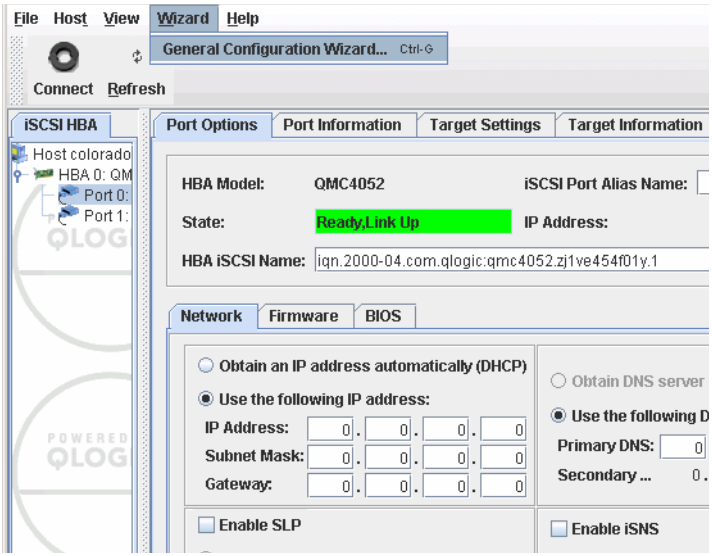

Figure 7-105 Port options

7. The wizard needs to know which HBA and port has to be configured. Select a port and click **Next**.

Figure 7-106 Wizard — HBA port selection

8. Verify by inspecting the HBA port data that the right HBA port was chosen. Click **Next**.

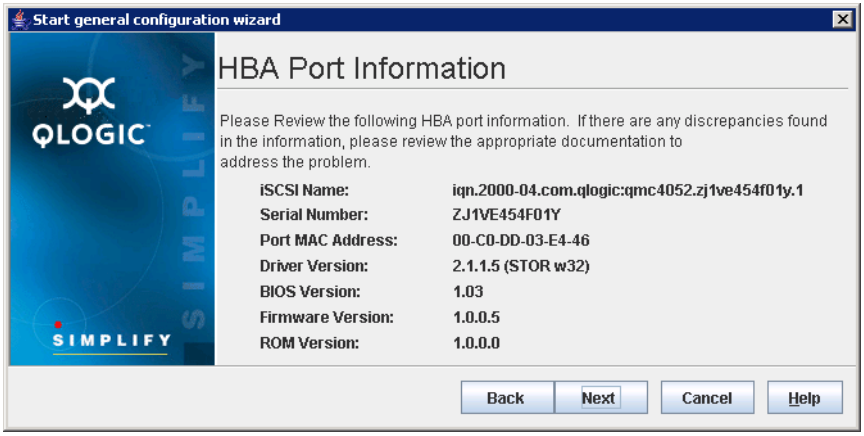

Figure 7-107 Wizard — HBA port information

9. Select **Use the following IP address** and set the IP address, subnet mask and gateway for this iSCSI HBA port. Select **Use the following DNS server addres**s and specify a DNS server. Click **Next**.

> **Note:** A DNS server is required when iSCSI targets will be discover by using the service location protocol (SLP).

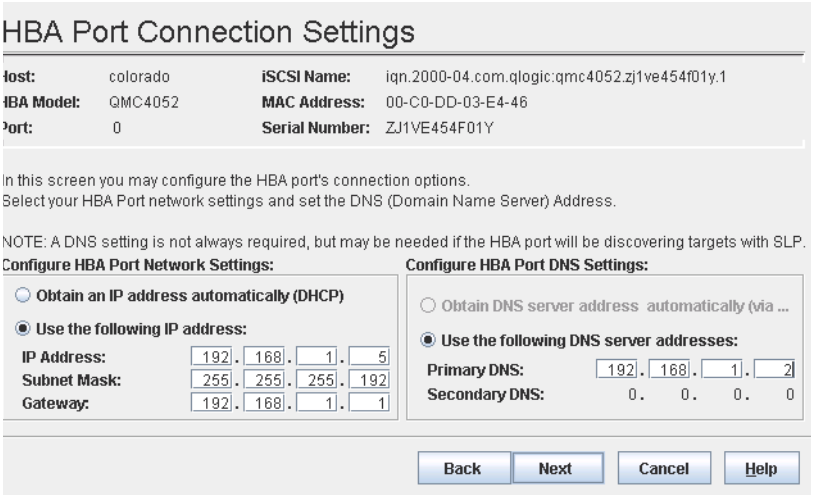

Figure 7-108 Wizard — HBA port connection settings

10. Click the green plus sign on the right site of the HBA port target configuration window. This button ads a new line in the target list and will ask for an IP address.

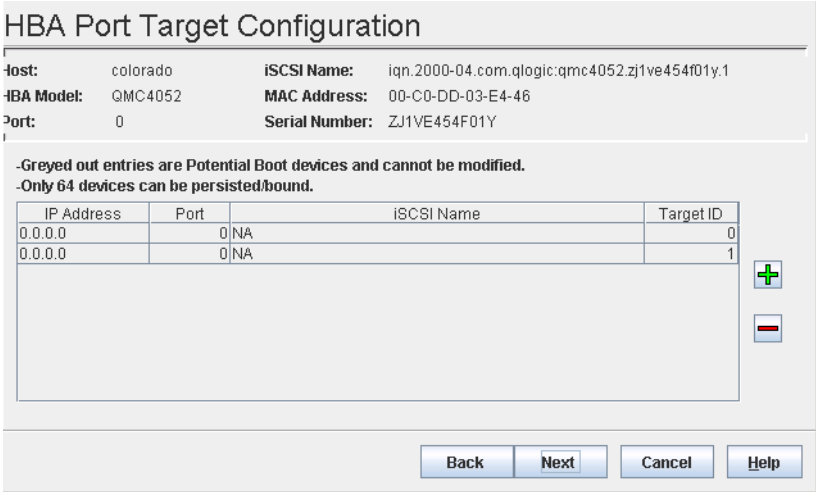

Figure 7-109 Wizard — HBA port target configuration

11..Specify the IP address of the iSCSI controller of an DS300. Click **OK**.

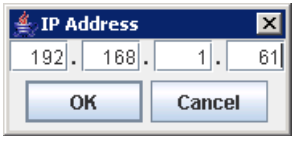

Figure 7-110 Wizard — iSCSI Device IP address

12. The specified iSCSI target will be added to the target list. Click **Next**.

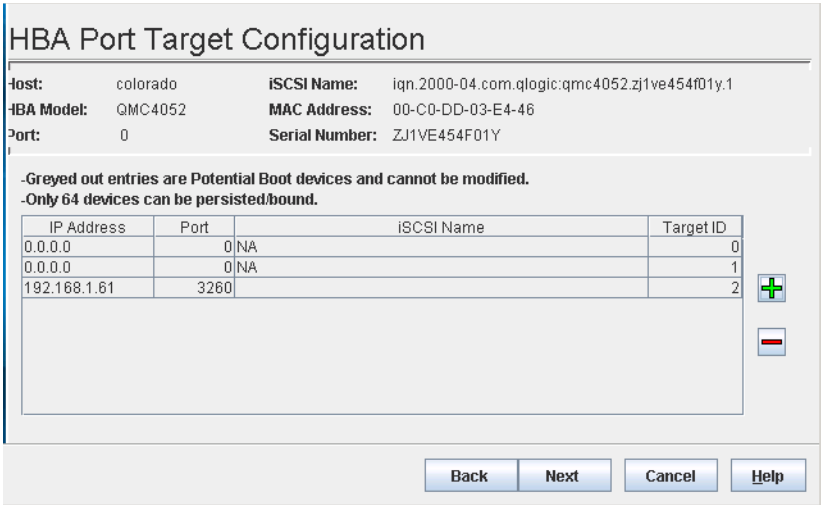

Figure 7-111 Wizard — new iSCSI target

13. Confirm the changes mad in the wizard by clicking **Next**. The configuration might be saved in a text file by pressing the **Save Configuration to File** button.

14. To apply the made configuration changes a password is required. By default is the password set to `config`. Enter the password and click **Next**.

Figure 7-112 Wizard - Security check dialog

15. The configuration will be now applied. This takes a short moment and the wizard will show a message that configuration update has successfully applied. Confirm this message by clicking **Next**.

16. Do not configure the second HBA port with the wizard. Click **Cancel** to finish the wizard.

17. Select target settings of a configured HBA port. Uncheck the **Auto-bind Discovered Targets** option.

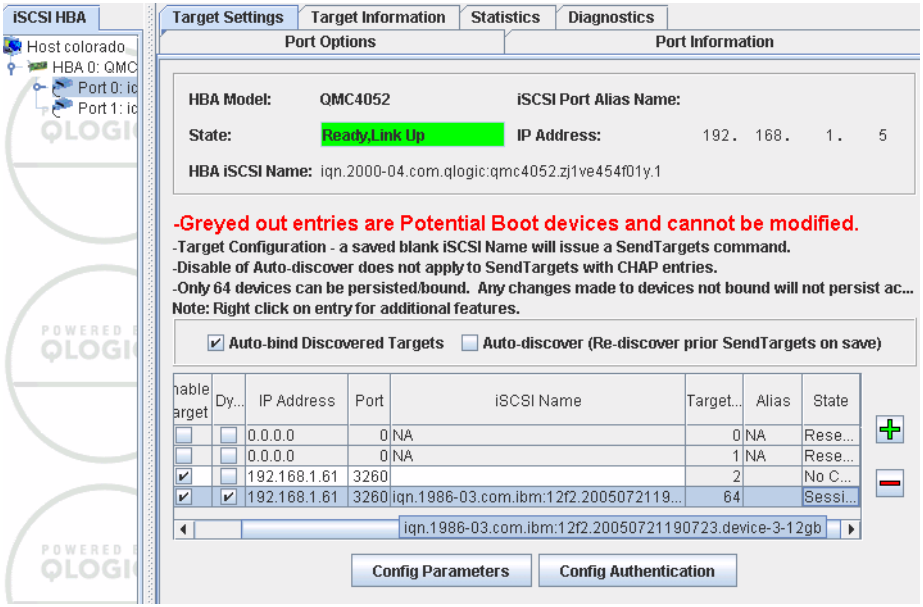

Figure 7-113 Uncheck Auto-bind

18. The last step is the authentication configuration. See 7.2.4, “Configure authentication” on page 300 for further actions.

Setup Port 1 manually and with iSNS

This procedure assumes that the host bus adapter address has already been configured (DHCP or static address), and you have manually specified a DNS or acquired one using DHCP.

1. Select Port 1. On the **Port Options** register card specify the IP address, subnet mask and gateway.
2. Enter a DNS server IP.
3. Select the check box **Enable iSNS** and specify the IP address of the iSNS server.
4. Click **Save Port Settings**.

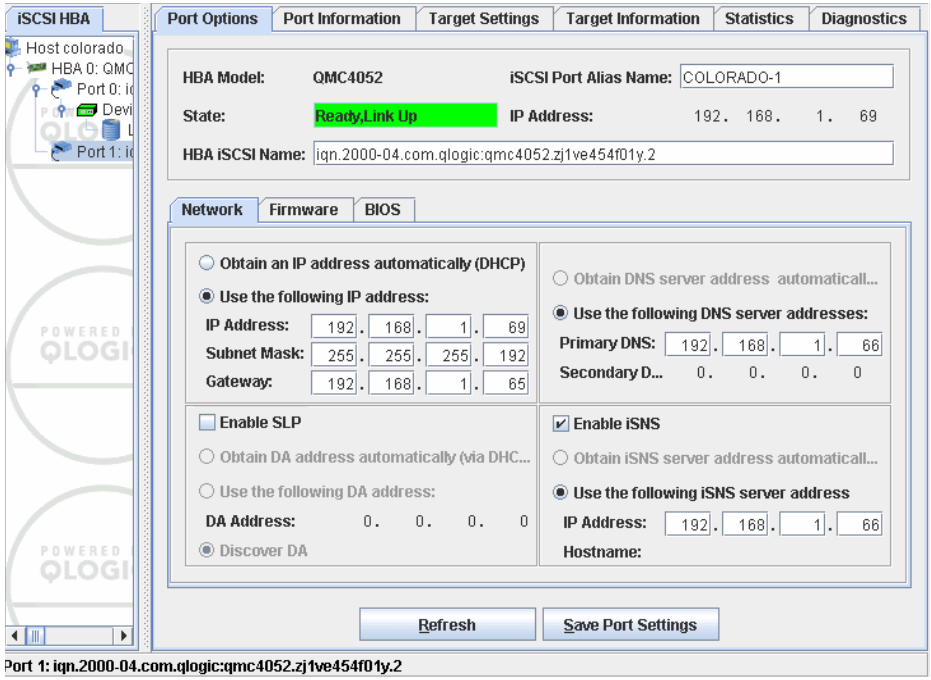

Figure 7-114 iSCSI HBA network settings

5. Select the register card **Target Settings** and select the **Auto discover** option. Save the settings by clicking **Save settings**. To save settings enter the password.The default password is `config`. Click the **Refresh** button. This allows you to get the a list of iSCSI targets from the iSNS server.

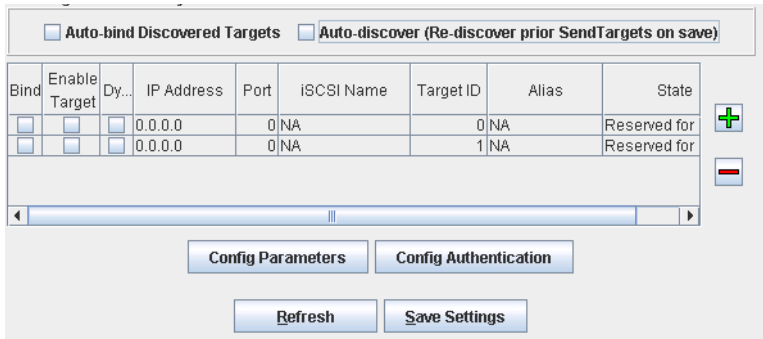

Figure 7-115 iSCSI target list

6. After a moment, the target list is updated with iSCSI targets from the iSNS server updated iSCSI target list.

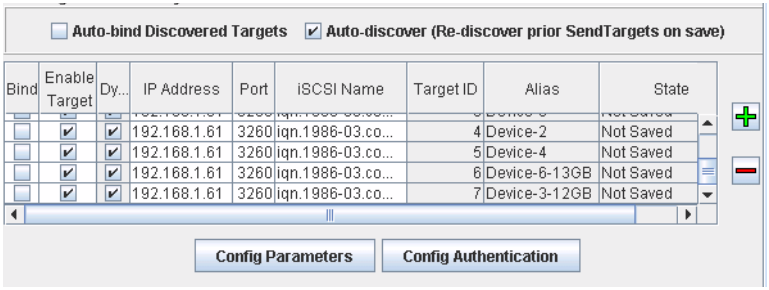

Figure 7-116 Updated iSCSI target list

7. The last step is the authentication configuration. See 7.2.4, “Configure authentication” on page 300 for further actions.

7.2.4 Configure authentication

1. From the **Target Settings** register card look for those iSCSI targets which will be attached to the host. Check the **Bind** check box for each target. Click the **Config Authentication** button.

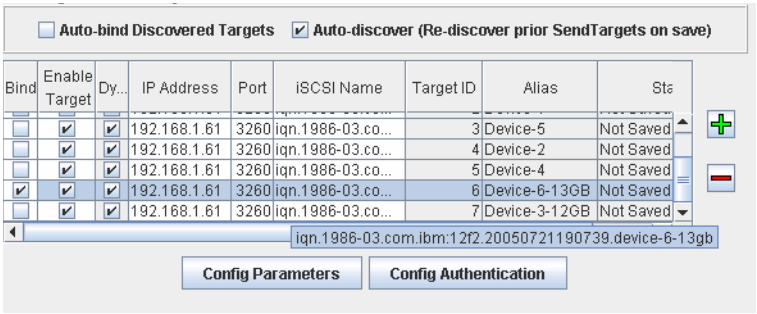

Figure 7-117 Bind

2. Enter the administration password. The default password is config.
3. Click **Yes**.

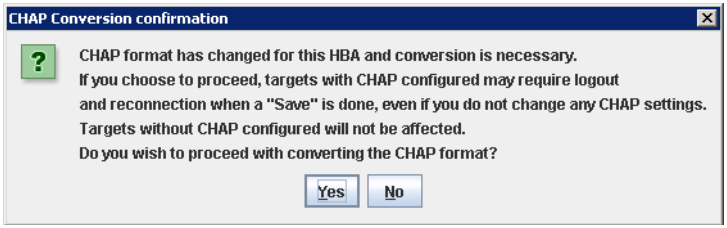

Figure 7-118 CHAP conversion confirmation

4. The CHAP register card is shown.

   This register card consists of 3 areas. A targets area, a CHAP entries area and a target table. Target table contains iSCSI target names and secrets to authenticate the target at the host. This is used together with the initiator authentication at the target. Initiator authentication can be used alone or bidirectional with the target authentication. The CHAP Entry area contains a list of user names and passwords to authenticate the initiator at the iSCSI target. In the Targets area a username and password pair is assigned to a iSCSI target.

5. Click the green plus sign in the CHAP entry section so create a new username password pair. Use as username the IQN. THe password needs to be at least 12 characters long

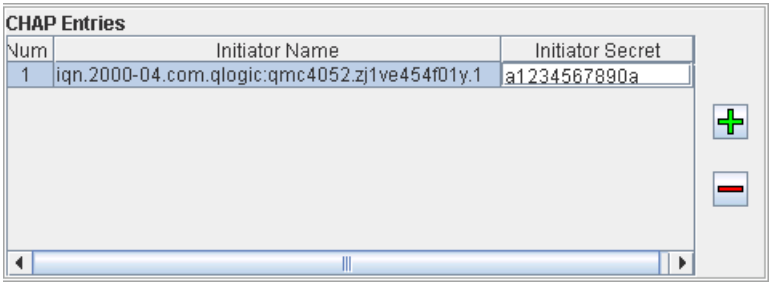

Figure 7-119 CHAP Username Password pair

6. Assign in the target area the userid password pair to a iSCSI target by selecting it from the drop down list.

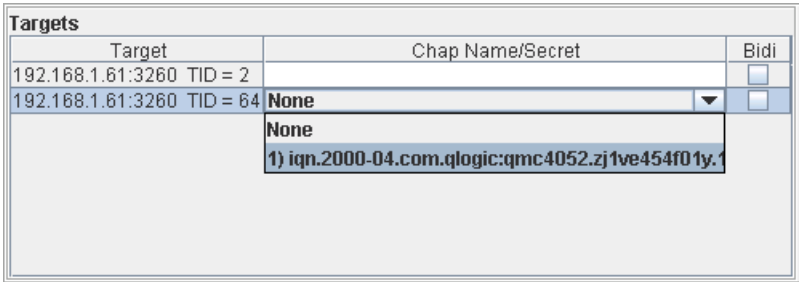

Figure 7-120 Select a userid/password pair

7. Click **OK** to finish.

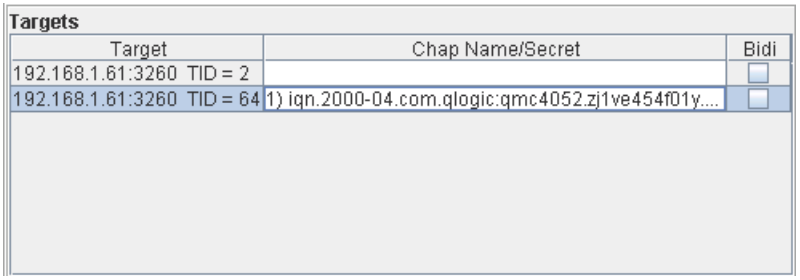

Figure 7-121 Finished

8. The state of the attached iSCSI target is now Session Active and the disk is accessible by the operating system.

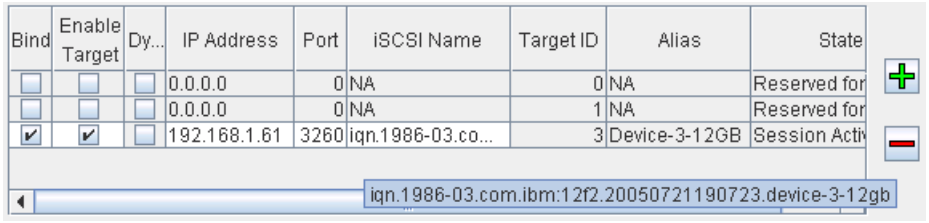

Figure 7-122 Session Active

7.3 QLogic iSCSI Expansion Card for IBM eServer BladeCenter in Linux

7.3.1 Driver installation

Follow these steps to install the driver on a Linux 2.6 kernel. These steps do not apply to a Linux 2.4 kernel.

1. Download the current driver.
2. Extract the files from download file using **`tar xzf <DRIVERFILE>`**
3. Change to the directory containing the source files. The directory name is the version number of the driver.
4. Build the driver using **`./extras/build.sh`**.
5. Install the driver using **`./extras/buil.sd install`**.
6. To ensure that the driver works reboot the server and verify the driver version using **`cat /proc/scsi/qla4xxx/<SCSI-HOST-NUMBER>`**.

Figure 7-123 on page 304 shows steps 2 - 5 in our test setup using SLES9 SP2.

```
colorado:~/qla4xxx # tar xzf qla4xxx-5.00.03-2.tar.gz
colorado:~/qla4xxx # cd 5.00.03/
colorado:~/qla4xxx/5.00.03 # ./extras/build.sh
qla4xxx -- Building the qla4xxx driver...
make: Entering directory `/usr/src/linux-2.6.5-7.191'
  CC [M]  /root/qla4xxx/5.00.03/ql4_os.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_init.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_mbx.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_iocb.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_isr.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_isns.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_nvram.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_dbg.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_cfg.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_cfgln.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_fo.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_foio.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_foioctl.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_foln.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_xioct.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_inioct.o
  CC [M]  /root/qla4xxx/5.00.03/ql4_32ioctl.o
  LD [M]  /root/qla4xxx/5.00.03/qla4xxx.o
  Building modules, stage 2.
  MODPOST
  CC      /root/qla4xxx/5.00.03/qla4xxx.mod.o
  LD [M]  /root/qla4xxx/5.00.03/qla4xxx.ko
make: Leaving directory `/usr/src/linux-2.6.5-7.191'
colorado:~/qla4xxx/5.00.03 #
```

Figure 7-123 Test setup

7.3.2 iSCSI SANsurfer management software installation

To install the SANsurfer management software for Linux, follow these steps:

1. Open a shell in your graphical environment. Extract the installation file and start the installation executable as shown in Figure 7-124 on page 305.

```
danube:~/SANsurfer # tar xzvf
iSCSI_SANsurfer_4_01_00_linux_x86_package.tgz.gz
Readme_SANsurfer_iSCSI_HBA_Mgr.txt
ReleaseNotes_SANsurfer_iSCSI_HBA_Mgr.pdf
iSCSI_SANsurfer_4_01_00_linux_x86.bin
danube:~/SANsurfer # ./iSCSI_SANsurfer_4_01_00_linux_x86.bin
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...

Launching installer...
```

Figure 7-124 SANsurfer Linux installation step 1

2. The SANsurfer graphical installer appears. The introduction screen gives you some advice about the installation. Click **Next** to proceed.

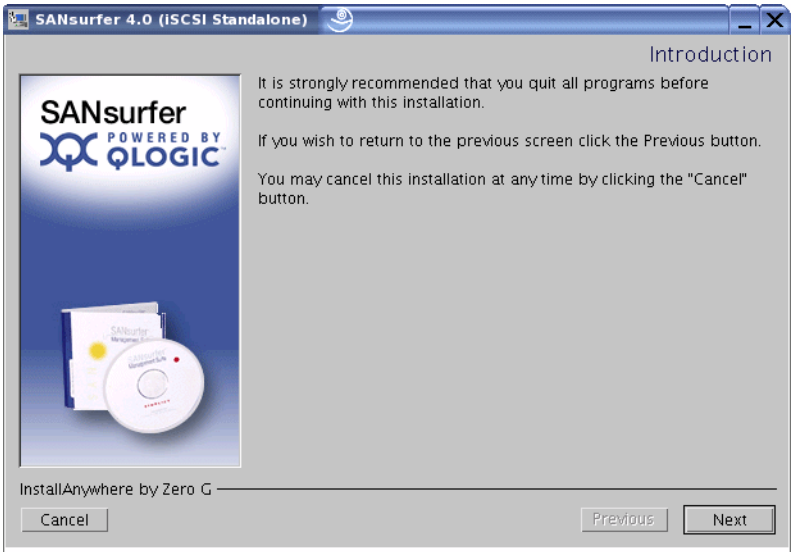

Figure 7-125 SANsurfer Linux installation step 2

3. Read the important information about the management application and then click **Next**.

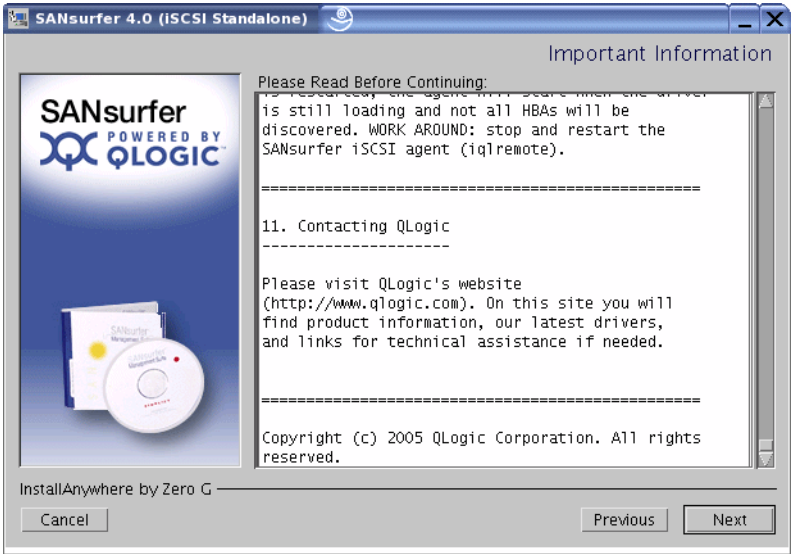

Figure 7-126 SANsurfer Linux installation step 3

4. Select the options you want to install. The installer allows to install the GUI and the agent alone or together. Depending on your needs select the appropriate option. The example below installs the GUI and the agent together on the same system. Click the **Next** button.

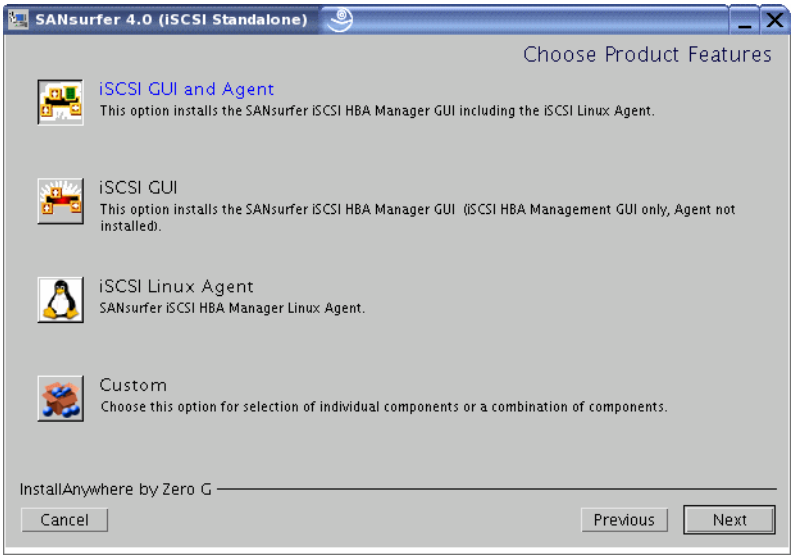

Figure 7-127 SANsurfer Linux installation step 4

5. You can specify a installation destination. Click **Next** to continue.

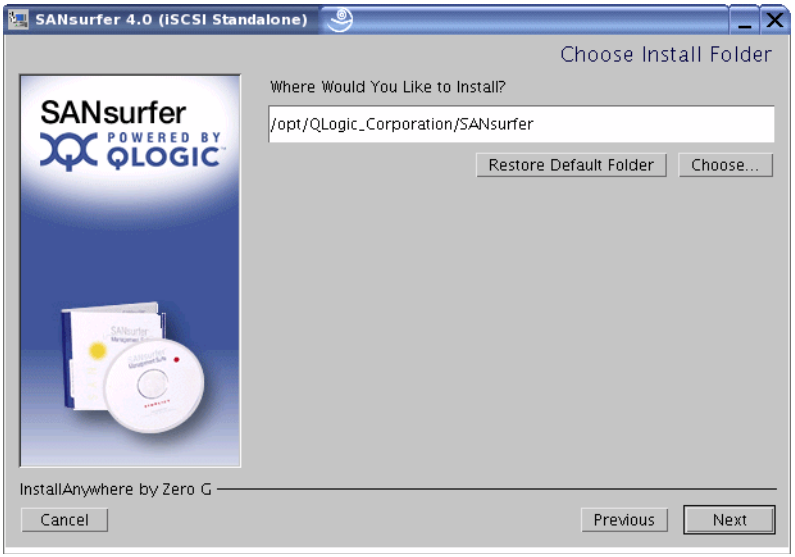

Figure 7-128 SANsurfer Linux installation step 5

6. Click **Install** in the pre-installation summary to start the installation.

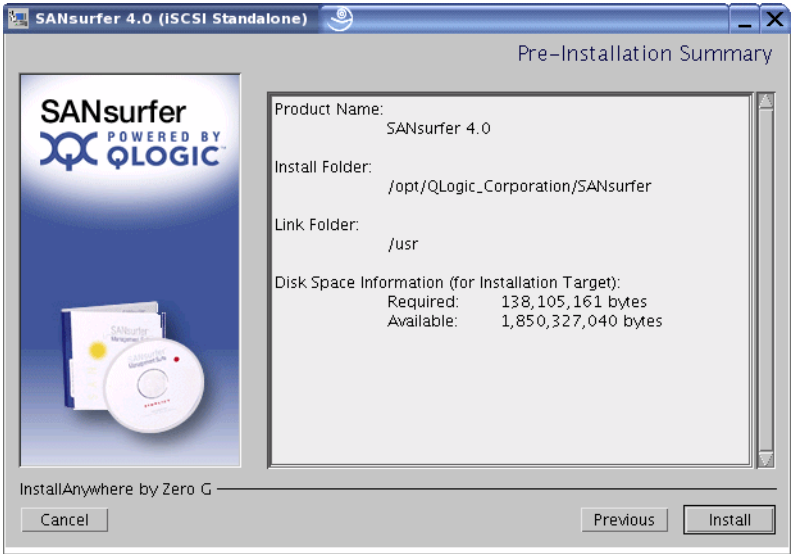

Figure 7-129 SANsurfer Linux installation step 6

7. Click **Done** to close the installation application.

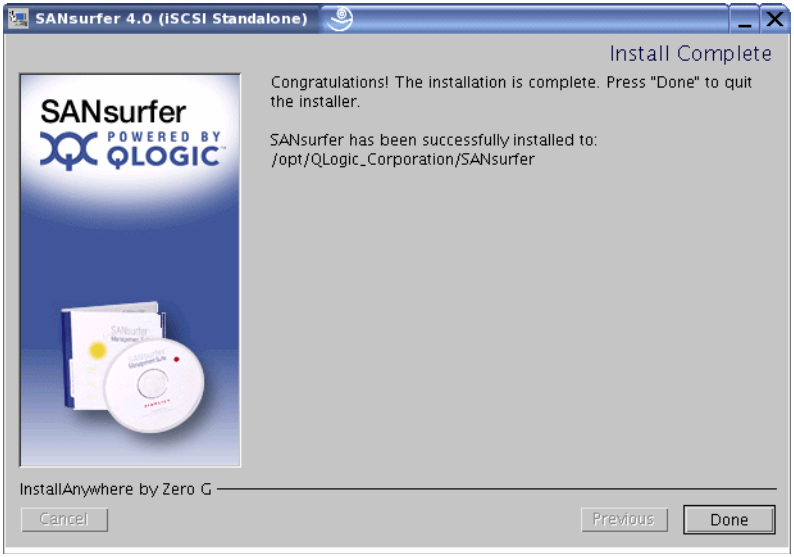

Figure 7-130 SANsurfer Linux installation step 7

7.3.3 Configuration of the QLogic iSCSI Expansion Card for IBM eServer BladeCenter

The Qlogic iSCSI Expansion Card for IBM @server BladeCenter can be configured in its BIOS during system POST (by pressing **CTRL+Q**) or with the iSCSI SANsurfer Management Software.

The configuration is almost the same as described in 7.2.3, “Setup of the QLogic iSCSI Expansion Card for IBM eServer BladeCenter” on page 291.

Figure 7-131 on page 309 shows SANsurfer on a system with two installed IBM iSCSI Server Adapter cards.

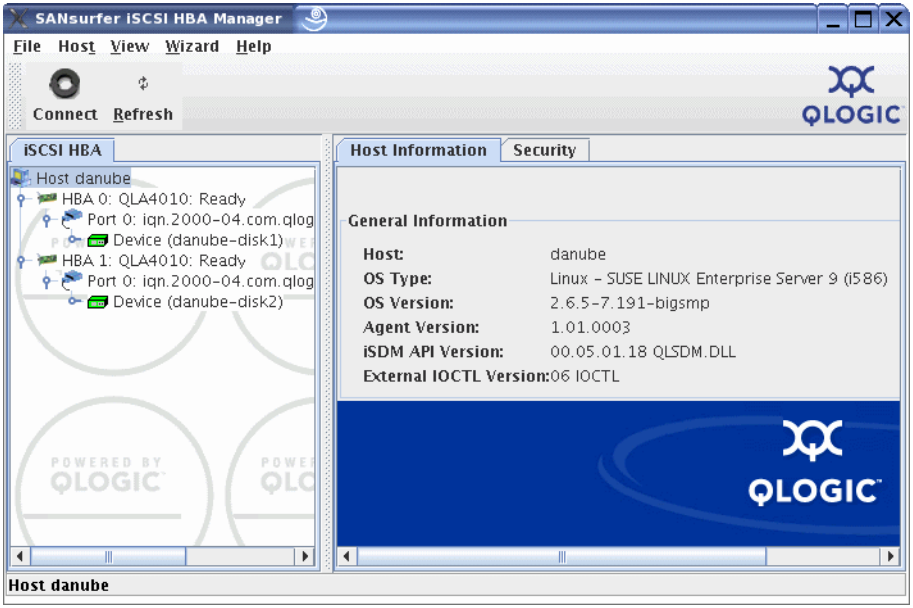

Figure 7-131 SANsurfer application in Linux

Note: When using two cards in a server, configure the cards to access different logical drives to do a load sharing. Currently the driver does not support multipathing to a logical drive using multiple HBAs. We tested device mapper and the multipath-tools for multipathing, but experienced problems as the readsector0 checker reported failed paths online again after some seconds.

7.4 IBM iSCSI Server Adapter in Windows Server 2003

7.4.1 Planning, considerations and network setup

Two IBM iSCSI Server adapter are installed in an IBM @serverxSeries 445 in slot 5 and 6. As operating system a Microsoft Windows Server 2003 Enterprise Edition with service pack 1 was chosen. Even with two iSCSI server adapters no redundant paths are usable because there is right now no multipath driver for the TotalStorage DS300 available. In case of an HBA failure all logical drives assigned to the failed initiator can be manually migrated to the other initiator. In worst case a reboot might be required. All drives should come up with the same drive letter as they had as they where accessed over the failed host bus adapter. The two iSCSI adapters where used to distribute the required logical drives over

the two controllers and the 4 available iSCSI interface ports. The network configuration looks like a redundant configuration but it isn't. The configuration allows only a controller failover. This prevents a for situations where a path between switch and controller or a controller itself fails. Every logical drive will have only one initiator assigned.

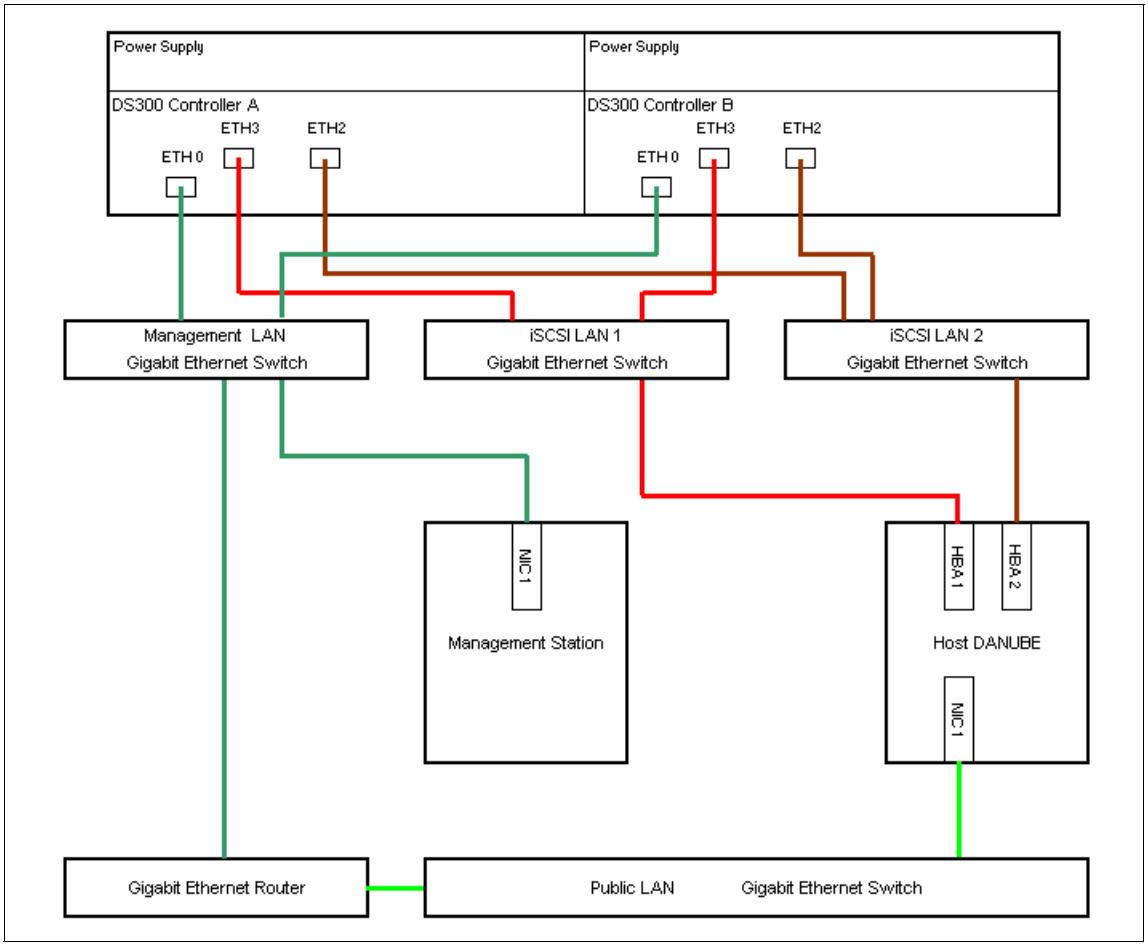

Figure 7-132 Configuration setup

7.4.2 Microsoft Windows Server 2003 Enterprise Edition installation

The operating system was installed as described in:

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-51935

After installation is complete install Service pack 1 and all remaining drivers See 7.4.3, “IBM iSCSI Server Adapter driver installation” on page 311 about installing drivers for the IBM iSCSI Server Adapter.

For operating system installation instructions see the IBM support Web site. Appendix C, “How to get updates of code and documentation” on page 623 describes how to find operating system instructions for other IBM @serverxSeries server.

7.4.3 IBM iSCSI Server Adapter driver installation

For each physically installed IBM iSCSI Server Adapter there will appear two devices in the **Other devices** tree of the device manager:

- Ethernet Controller (function 0)
  - This device refers to the gigabit ethernet adapter functionality of the IBM iSCSI Server Adapter.
  - After driver installation it will appear in the **Network adapters** tree of the device manager.
- Network Controller (function 1) - this refers to the iSCSI
  - This device refers to the iSCSI adapter functionality of the IBM iSCSI Server Adapter.
  - After driver installation it will appear in the **SCSI and RAID controllers** tree of the device manager.

The driver for the IBM iSCSI Server adapter is not included on the Microsoft Windows product CD. The drivers can be found via the following URL:

`http://www.ibm.com/servers/storage/support/disk/ds300/downloading.html`

To install the drivers for the IBM iSCSI Server Adapter follow these steps:

1. Open the **Device Manager**. Depending on the number of physically installed IBM iSCSI Server Adapters there are multiple devices in the **Other devices** tree. Right-click the first **Ethernet Controller** in the **Other devices** tree. Click **Update Driver...**.

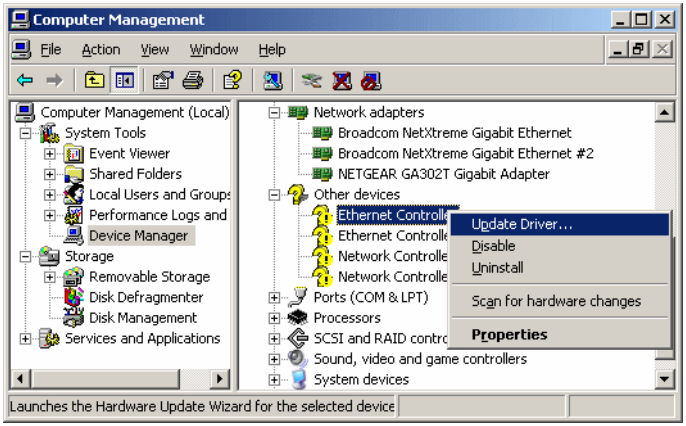

Figure 7-133 IBM iSCSI Server Adapter driver installation step 1

2. In the **Hardware Update Wizard** choose **No, not this time** and click **Next**.

Figure 7-134 IBM iSCSI Server Adapter driver installation step 2

3. Select **Install from a list or specific location (Advanced)** and click **Next**.

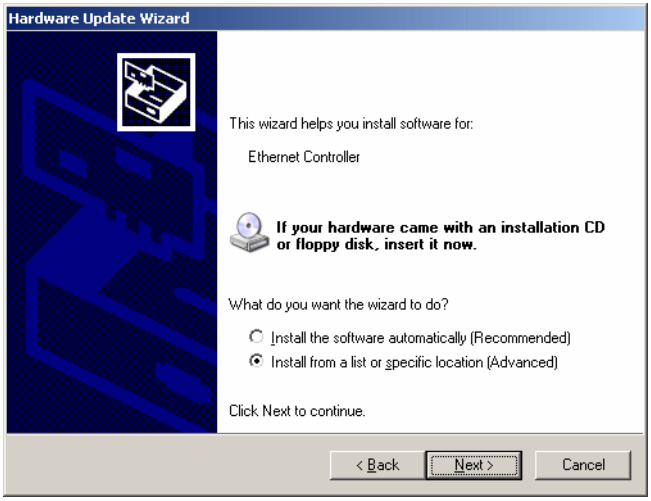

Figure 7-135 IBM iSCSI Server Adapter driver installation step 3

4. Browse to the folder where you have downloaded the NDIS network driver for the card. Click **Next**.

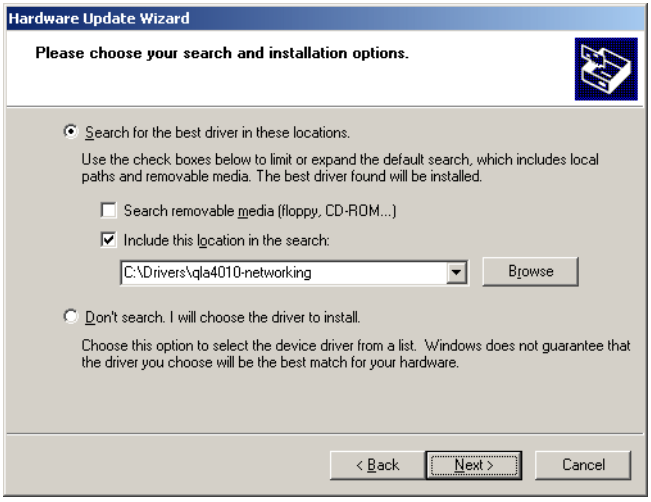

Figure 7-136 IBM iSCSI Server Adapter driver installation step 4

5. After the driver installation finishes, click **Finish** to close the Hardware Update Wizard.

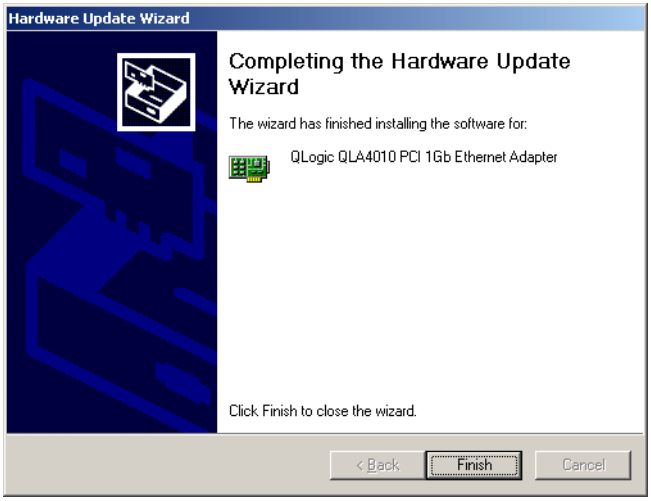

Figure 7-137 IBM iSCSI Server Adapter driver installation step 5

6. In the **Network adapters** tree of the device manager appears a **QLogic QLA4010 PCI 1Gb Ethernet Adapter**. Redo steps 1-5 for the remaining ethernet controllers in the **Other devices** tree.

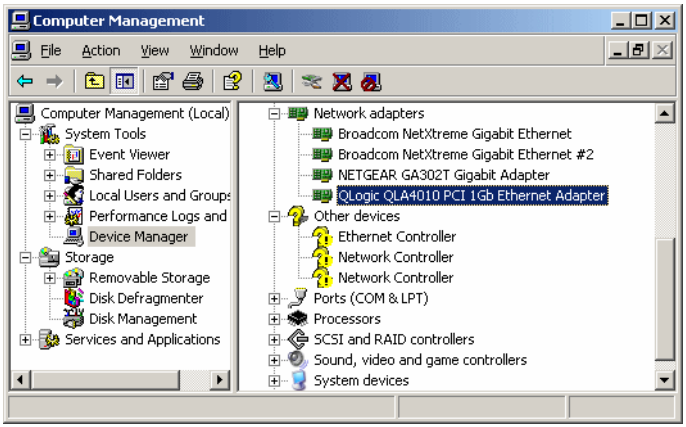

Figure 7-138 IBM iSCSI Server Adapter driver installation step 6

Note: When using the NDIS driver for the iSCSI host bus adapter do not assign the same IP address to the network interface seen by the operating system then than the host bus adapter has got with SANsurfer or in the BIOS Menu setup. There is no off load engine used when the iSCSI server adapter is used together with the NDIS driver.

7. Right-click the first **Network Controller** in the **Other devices** tree. Click **Update Driver...**.

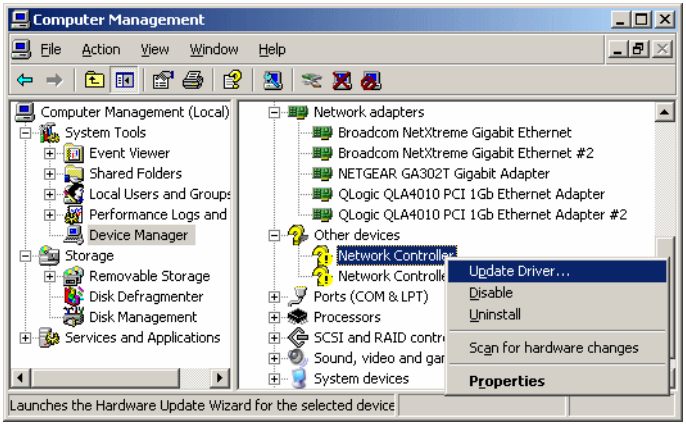

Figure 7-139 IBM iSCSI Server Adapter driver installation step 7

8. In the **Hardware Update Wizard** choose **No, not this time** and click **Next**.

Figure 7-140 IBM iSCSI Server Adapter driver installation step 8

9. Select **Install from a list or specific location (Advanced)** and click **Next**.

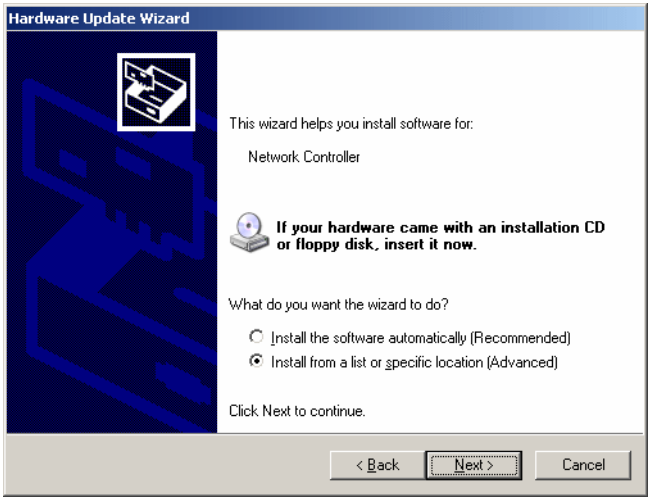

Figure 7-141 IBM iSCSI Server Adapter driver installation step 9

10. Browse to the folder where you have downloaded the storport driver for the card. Click **Next**.

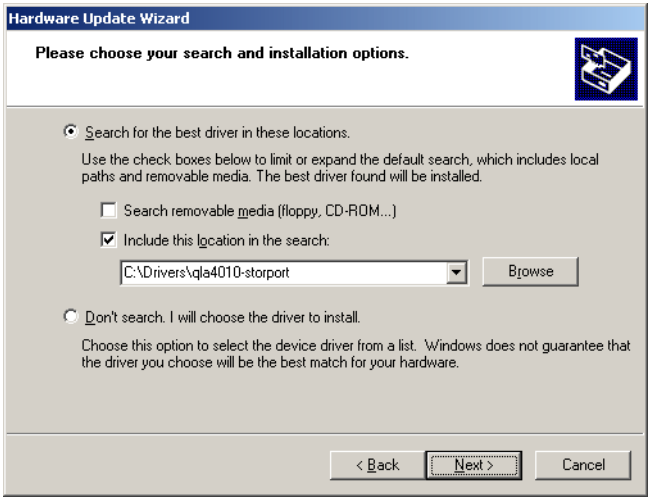

Figure 7-142 IBM iSCSI Server Adapter driver installation step 10

11. After the driver installation finishes, click **Finish** to close the Hardware Update Wizard.

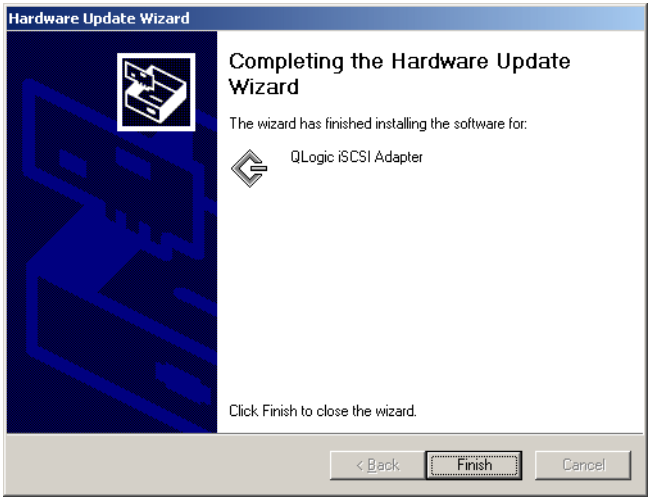

Figure 7-143 IBM iSCSI Server Adapter driver installation step 11

12. In the **SCSI and RAID controllers** tree of the device manager appears a **QLogic iSCSI Adapter**. Repeat steps 7-11 for the remaining network controllers in the **Other devices** tree.

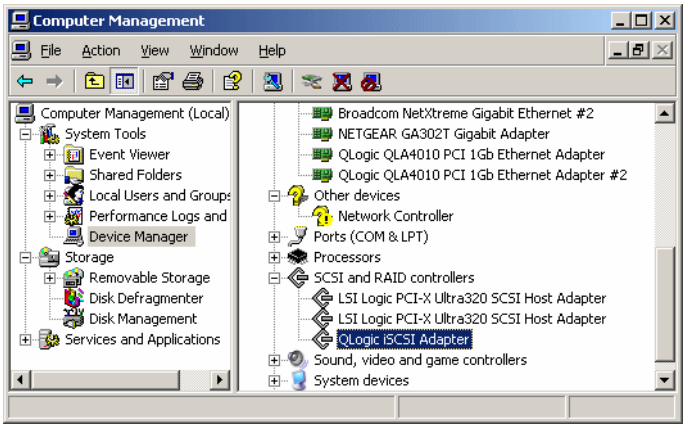

Figure 7-144 IBM iSCSI Server Adapter driver installation step 12

7.4.4 iSCSI SANsurfer management software installation

Refer to 7.2.2, “Install the Qlogic iSCSI Expansion Card for IBM eServer BladeCenter management software” on page 286. The installation of SANsurfer for the IBM iSCSI Server Adapter is the same as for the Qlogic iSCSI Expansion Card for IBM @server BladeCenter.

7.4.5 Configuration of the IBM iSCSI Server Adapter Card

The IBM iSCSI Server Adapter Card can be configured in its BIOS during system POST (by pressing **CTRL+Q**) or with the iSCSI SANsurfer Management Software.

The configuration is almost the same as described in 7.2.3, “Setup of the QLogic iSCSI Expansion Card for IBM eServer BladeCenter” on page 291. The only difference is that the IBM iSCSI Server Adapter Card is a single port card. Therefore, a single port is displayed in the SANsurfer iSCSI HBA manager application for each installed HBA.

Note: When using two cards in a server, configure the cards to access different logical drives to do load sharing. Currently, the driver does not support multipathing to a logical drive using multiple HBAs.

The final configuration looks as shown in figure Figure 7-145 on page 319. HBA 0 accesses the logical drives `DANUBE-DISK1` and `DANUBE-LOG` over the IP address 192.168.1.61. This is ETH2 of controller A. Those two logical drives are in the same array which is owned by controller A. Only HBA 0 has access to those logical drives.

HBA1 has access to the logical drives `DANUBE-MAIL` and `DANUBE-DB` over 192.168.1.126. This is the IP address of ETH3 of controller B. Both logical drives are in their own array. Both arrays are owned by controller B.

Logical drive `DANUBE-PRINT` is also accessed by HBA1. The array this logical drive is in is owned by controller A. HBA1 gets access over the IP address `192.168.1.125`. This is the IP of ETH3 of controller A.

Note: Please remember to configure the ARP redirect and Keep alive timeout as described in 7.1.10, “Configure iSCSI host bus adapter firmware settings” on page 261.

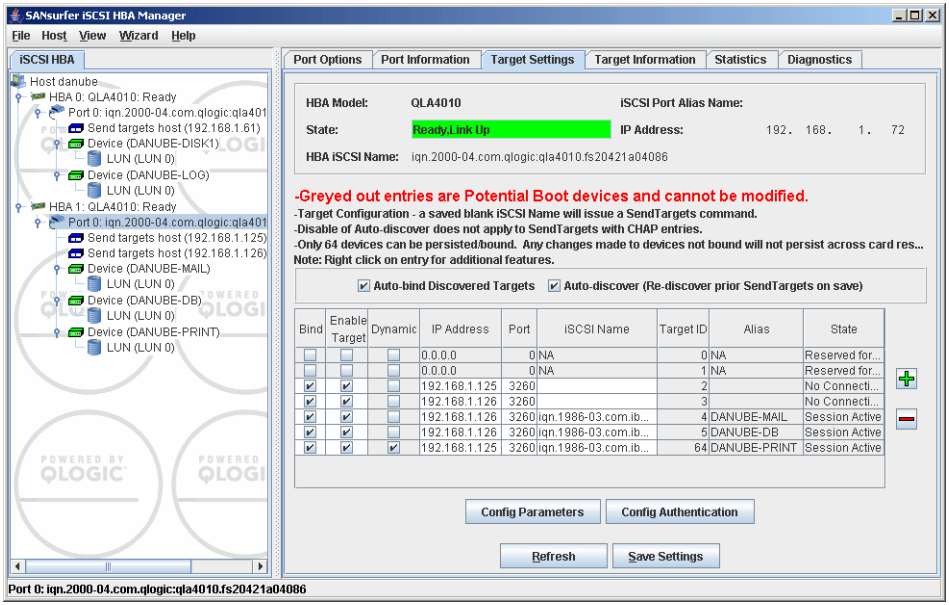

Figure 7-145 iSCSI host bus adapter configuration

7.5 IBM iSCSI Server Adapter in Microsoft Cluster Server

This section describes the installation of a Microsoft Cluster Server with two nodes and shared storage from a IBM TotalStorage DS300. Each cluster node is equipped with two IBM iSCSI Server adapter. This two adapter are currently not used for host bus adapter failover. The logical drives are accessed over one host bus adapter or the other but not both.

For security not only the HBA has to authenticate with CHAP on the TotalStorage DS300 but also the TotalStorage DS300 has to authenticate against the initatiors.This is called bidirectional authentication.

Microsoft Windows Server 2003 Enterprise Editions is used as operating system because this includes the cluster service.

7.5.1 Planning and considerations

The nodes will be equipped with two iSCSI server adapters. This permits the ability to distribute the load over two adapters. When a host bus adapter failover driver becomes available, it can be installed and the assignment of the logical

drives to the initiators can be extended. Arrays and logical drives are distributed over the two controllers of the TotalStorage DS300 and the two iSCSI interfaces on each port. This will be achieved by assigning the logical drives to different host bus adapters. The table below will explain this in detail:

Table 7-4 Assign logical drives to host bus adapters

Controller A
Array: ELBE-Quorum
Logical drive: ELBE-Quorum (500MB)
SEINE - HBA0
TIBER - HBA0
Logical drive: ELBE-MSDTC (510MB)
SEINE - HBA1
TIBER - HBA1
Array: ELBE-Mail
Logical drive: ELBE-Echange (8000MB)
SEINE - HBA1
TIBER - HBA1
Controller B
Array: ELBE-Domino
Logical drive: ELBE-Domino (10000MB)
SEINE - HBA0
TIBER - HBA0
Array: ELBE-DB
Logical drive: ELBE-SQL (9000MB)
SEINE - HBA1
TIBER - HBA1

Figure 7-146 on page 321 gives an overview of the cabling in this cluster setup.

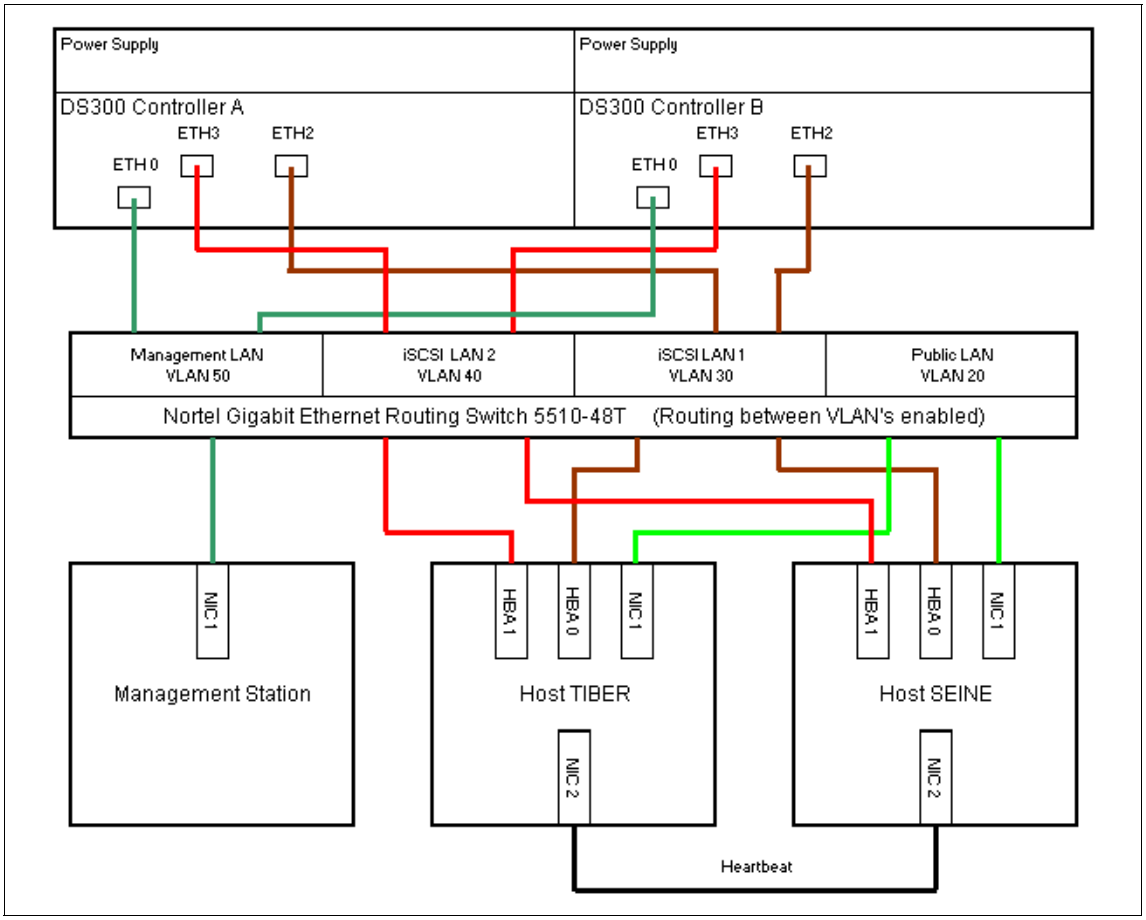

Figure 7-146 Cabling diagram

Table 7-5 lists the network settings of the TotalStorage DS300. See Chapter 6, “TotalStorage DS300 sample configurations” on page 211 for the basic network setup of a TotalStorage DS300.

Table 7-5 IBM TotalStorage DS300 network addresses

Usage	Address
DS300 - ETH0 Connected to	Management LAN
DS300 Controller A - ETH0 - IP	192.168.1.236
DS300 Controller B - ETH0 - IP	192.168.1.237
ETH0 - Subnet Mask	255.255.255.192

Usage	Address
ETH0 - Broadcast Address	192.168.1.255
Management LAN Gateway	192.168.1.254
DS300 - ETH2 Connected to	iSCSI LAN 1
DS300 Controller A - ETH2 - IP	192.168.1.13
DS300 Controller B - ETH2 - IP	192.168.1.14
ETH2 - Subnet Mask	255.255.255.192
ETH2 - Broadcast Address	192.168.1.63
iSCSI LAN 1 Gateway	192.168.1.62
DS300 - ETH3 Connected to	iSCSI LAN 2
DS300 Controller A - ETH3 - IP	192.168.1.77
DS300 Controller B - ETH3 - IP	192.168.1.78
ETH3 - Subnet Mask	255.255.255.192
ETH3 - Broadcast Address	192.168.1.127
iSCSI LAN 2 Gateway	192.168.1.126

Table 7-6 lists the network settings used on cluster node `SEINE`.

Table 7-6 Cluster node 1 network addresses

Usage	Address
Host name	SEINE
HBA0 - Port 0 connected to	iSCSI LAN 1
HBA0 - Port 0 IP Address	192.168.1.9
HBA0 - Port 0 Subnet Mask	255.255.255.192
HBA0 - Port 0 Gateway	192.168.1.62
HBA1 - Port 0 connected to	iSCSI LAN 2
HBA1 - Port 0 IP Address	192.168.1.73
HBA1 - Port 0 Subnet Mask	255.255.255.192
HBA1 - Port 0 Gateway	192.168.1.126

Usage	Address
NIC0 - IP Address	172.16.100.141
NIC0 - Subnet Mask	255.255.0.0
NIC1 - IP Address	10.0.0.1
NIC1 - Subnet Mask	255.0.0.0
DNS Server	172.16.100.10

Table 7-7 lists the network settings used on cluster node TIBER.

Table 7-7 Cluster node 2 network addresses

Usage	Address
Host name	TIBER
HBA0 - Port 0 connected to	iSCSI LAN 1
HBA0 - Port 0 IP Address	192.168.1.11
HBA0 - Port 0 Subnet Mask	255.255.255.192
HBA0 - Port 0 Gateway	192.168.1.62
HBA1 - Port 0 connected to	iSCSI LAN 2
HBA1 - Port 0 IP Address	192.168.1.75
HBA1 - Port 0 Subnet Mask	255.255.255.192
HBA1 - Port 0 Gateway	192.168.1.126
NIC0 - IP Address	172.16.100.142
NIC0 - Subnet Mask	255.255.0.0
NIC1 - IP Address	10.0.0.2
NIC1 - Subnet Mask	255.0.0.0
DNS Server	172.16.100.10

7.5.2 Install the server operating system

Install Microsoft Windows Server 2003 Enterprise Edition as described in the operating system installation guide for the chosen server. The matrix for the operating system installation guides can be found under

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-4QLNTQ

After installation of the operating system, apply all updates such as hardware updates, driver and service packs. Join the domain. The next section explains the driver installation of the IBM iSCSI Server Adapter.

7.5.3 Install drivers for the IBM iSCSI Server Adapter

Download the latest supported driver for the IBM iSCSI Server Adapter for a connection to the TotalStorage DS300 from IBM support site at:

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

Unpack the driver and be sure that the driver is supported for connecting to a IBM TotalStorage DS300.

Install the driver as described in the following steps:

1. Unpack the driver into a local directory.

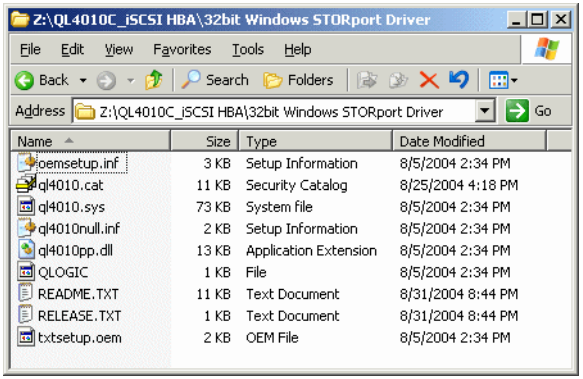

Figure 7-147 IBM iSCSI Server Adapter driver files

2. Open the Device Manager. In the Device Manager are four unknown devices. One Ethernet Controller and one Network Controller are different functions of the same device. The Location field in the properties window of those unknown devices can be used to identify those devices. The Ethernet Controller is function 0 and the network controller is function 1. After the installation of a ethernet network adapter driver for the function 0 device, a new ethernet network card can be used. The function 1 device will be used as a SCSI controller. Right-click the first Network Controller. In the context menu click **Update Driver...**.

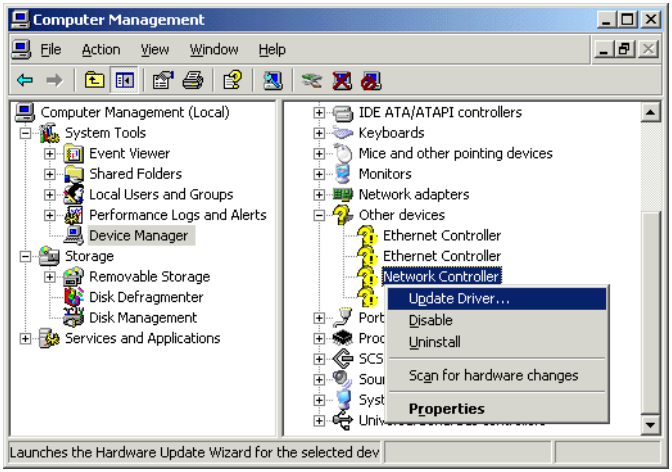

Figure 7-148 iSCSI Server Adapter — Network controller driver installation

3. Select **No, not this time** and click **Next**.

Figure 7-149 Hardware update wizard — Welcome

4. Select **Install from a list or specific location** and click **Next**.

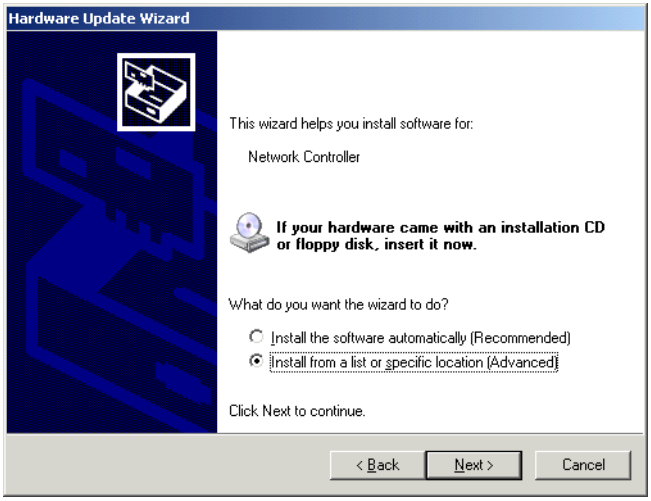

Figure 7-150 Hardware update wizard — Choose software installation method

5. Uncheck **Search removable media**. Check **Include this location in the search** and click **Browse**. In the **Browse For Folder** dialog choose the folder where the driver was unpacked into. Click **Next**.

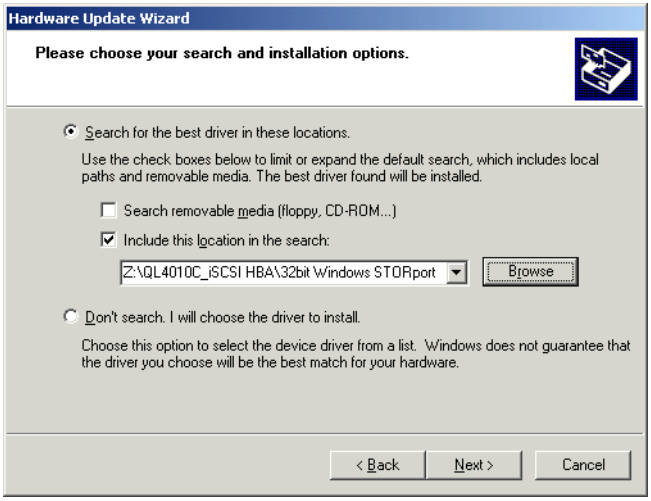

Figure 7-151 Hardware update wizard — Choose search options

6. In the specified location will be a driver searched for the Network Controller. The driver will be found and installed. Click **Finish**.

Figure 7-152 Hardware update wizard — Completing the wizard

7. The Network Controller will disappear under other devices and show up in the SCSI and RAID controllers section.

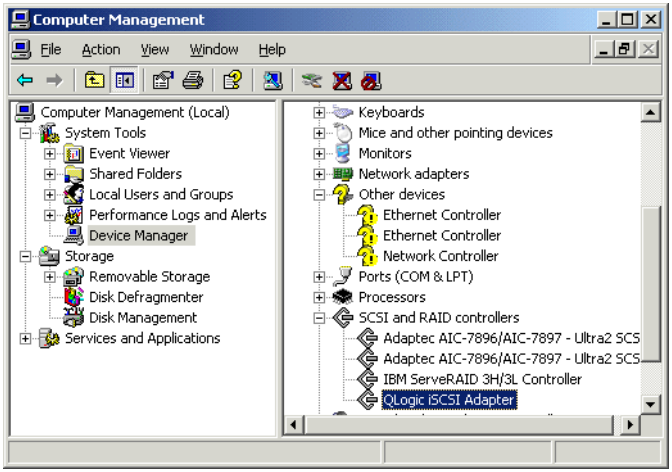

Figure 7-153 Device Manager with installed QLogic iSCSI Adapter

8. Repeat step 2 on page 324 till step 7 for the second Network Controller in the Device Manager under other devices.
9. Both network controllers are installed. Install now the driver the Ethernet Controller. Depending on the planned usage two different drivers can be installed. The storport driver package contains a dummy driver for the ethernet controller in case the network interface of this iSCSI host bus

adapter is not required to be usable within the operating system. An NDIS driver can be installed to use the iSCSI host bus adapter like any other network card. This driver is included in a combined package of storport and NDIS driver.

Note: When using the iSCSI HBA in the operating system additionally as a regular network adapter, assign a different IP address from the same IP range as it was done in the adapter BIOS or with SANsurfer Pro.

Install the dummy driver.

10. Right-click the Ethernet controller device in other devices. In the context menu click **Update Driver ...** .

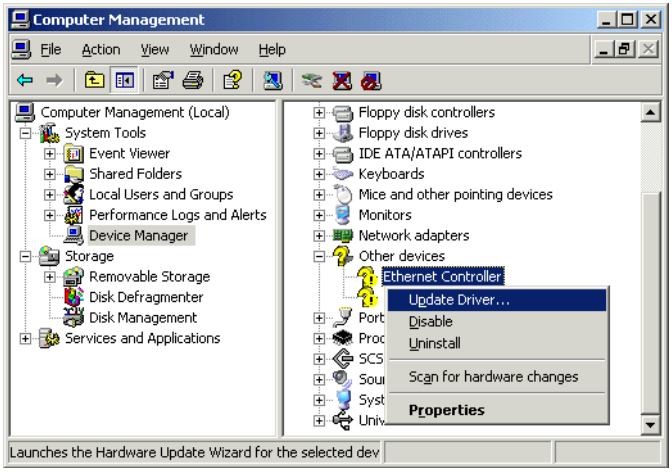

Figure 7-154 iSCSI Server Adapter — Network controller driver installation

11. Select **No, not this time** and click **Next**.

Figure 7-155 Hardware update wizard — Welcome

12. Select **Install from a list or specific location** and click **Next**.

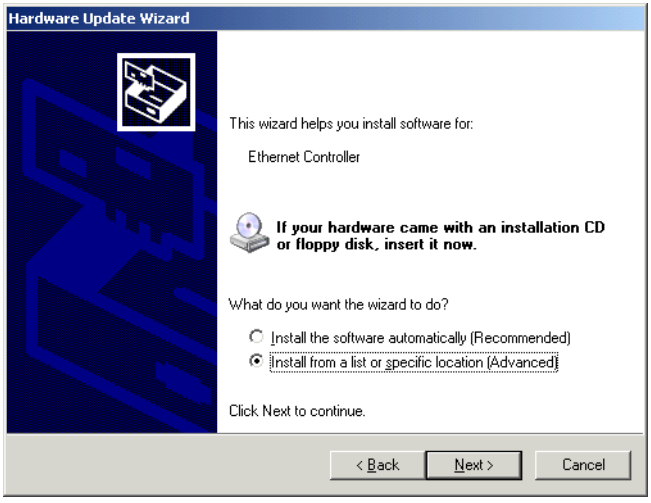

Figure 7-156 Hardware update wizard — Choose software installation method

13. Uncheck **Search removable media**. Check **Include this location in the search** and click **Browse**. In the **Browse For Folder** dialog choose the folder where the driver was unpacked into. Click **Next**.

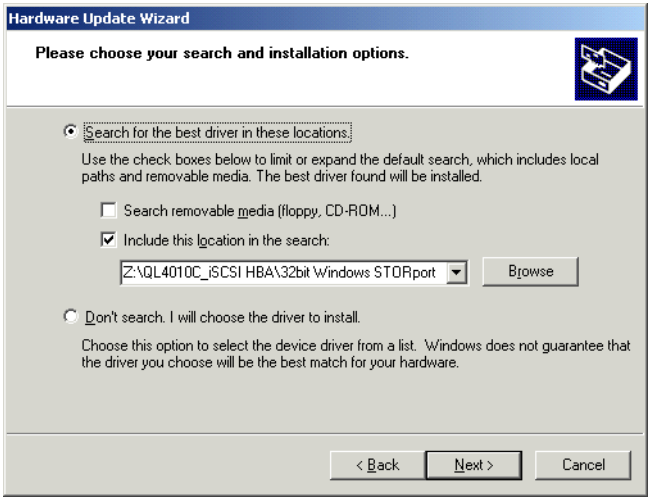

Figure 7-157 Hardware update wizard — Choose search options

14. Click **Finish**.

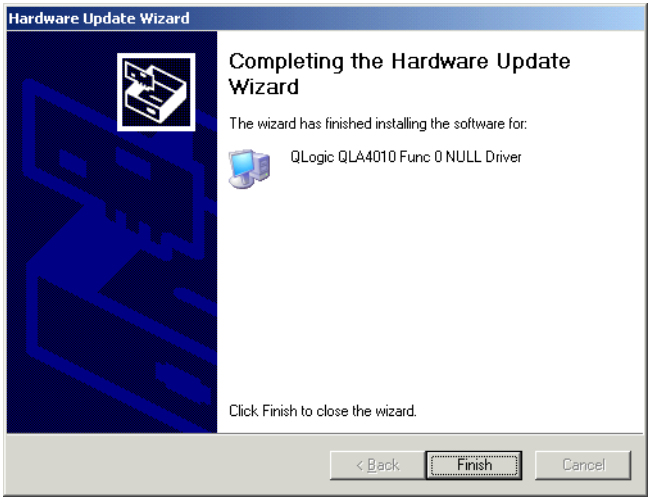

Figure 7-158 Hardware update wizard — Completing the wizard

15. Under System devices appears now a QLogic NULL driver device.

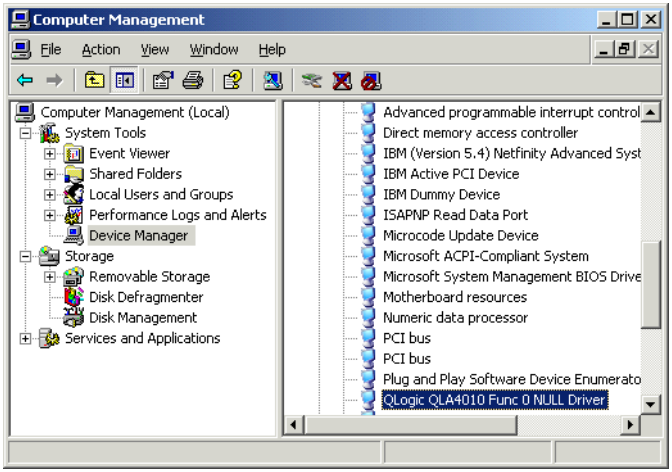

Figure 7-159 Device Manager with QLogic NULL Driver

16. Repeat step 10 on page 328 till step 15 on page 330 to install the NULL driver for the remaining Ethernet Controller.

> **Note:** In case the network functionality of the iSCSI host bus adapter is required, the NULL driver can be upgraded with the NDIS driver.

The drivers for the iSCSI host bus adapters are now installed.

7.5.4 Install SANsurfer Pro

SANsurfer Pro is a management application for the iSCSI host bus adapter. This software follows a client server model. Only the iSCSI agent required is required on both cluster nodes. A management station needs the graphical user interface of the SANsurfer Pro.

1. Download the SANsurfer Pro from the following URL:

   `http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58619`

2. Install on both cluster nodes TIBER and SEINE the SANsurfer Pro agent. Start the setup program.

3. Gather the IQNs of the host bus adapters.

7.5.5 Configure arrays and logical drives on the TotalStorage DS300

This section describes the configuration using the command line interface of the TotalStorage DS300. See 7.5.1, “Planning and considerations” on page 319 for details of the configuration.

1. Open a Telnet session to one of the controllers.
2. Logon as administrator.
3. Create the initiator definition by issuing the command **`initiator create`**.

```
initiator create TIBER-HBA0 iqn.2000-04.com.qlogic:qla4010.fs20430a05865
initiator create TIBER-HBA1 iqn.2000-04.com.qlogic:qla4010.fs20432a07396
initiator create SEINE-HBA0 iqn.2000-04.com.qlogic:qla4010.fs20417a03127
initiator create SEINE-HBA1 iqn.2000-04.com.qlogic:qla4010.fs20417a02573
```

Figure 7-160 Define initiators on the TotalStorage DS300

4. Create a user account with that the initiator can authenticate on the TotalStorage DS300.

   Note: The IQN is the default username used by the Microsoft Software Initiator. it is possible to change this user.

   After entering this command enter the password twice.

```
authentication add elbe
```

Figure 7-161 Create a user account on the TotalStorage DS300

5. Create the arrays. The parameter remote specifies that the array is owned by the alternate controller.

```
array create ELBE-Quorum raid1 65536 65792
array create ELBE-Exchange raid1 3584 3840
array create ELBE-DB remote Raid1 2304 2560
array create ELBE-Domino remote Raid1 3072 3328
```

Figure 7-162 Create arrays

6. Create the logical drives in the arrays.

```
logical create ELBE-Quorum ELBE-Quorum 500MB restricted SEINE-HBA0 TIBER-HBA0
logical create ELBE-MSDTC ELBE-Quorum 510MB restricted SEINE-HBA1 TIBER-HBA1
logical create ELBE-Domino ELBE-Domino 10000MB restricted SEINE-HBA0 TIBER-HBA0
logical create ELBE-Exchange ELBE-Exchange 8000MB restricted SEINE-HBA1 TIBER-HBA1
logical create ELBE-DB ELBE-DB 9000MB restricted SEINE-HBA1 TIBER-HBA1
```

Figure 7-163 Create logical drives

7. Modify the authentication type from none to CHAP for each logical drive.

```
logical manage ELBE-Quorum authentication CHAP
logical manage ELBE-MSDTC authentication CHAP
logical manage ELBE-Domino authentication CHAP
logical manage ELBE-Exchange authentication CHAP
logical manage ELBE-DB authentication CHAP
```

Figure 7-164 Modify authentication type

8. Define the username used for the bidirectional authentication. This user name will be used by the TotalStorage DS300 to authenticate against the initiator. There is only one account per enclosure. It is not possible to define this account within the IBM ServeRAID Manager console. Within the iSCSI Server Adapter or the QLogic iSCSI Expansion Card for IBM BladeCenter BIOS can be defined only one enclosure account. SANsurfer Pro has the ability to manage multiple enclosure accounts.

```
logical chap user ds300-4
```

Figure 7-165 Create enclosure account

9. Define the enclosure password.

```
logical chap password
```

Figure 7-166 Define the enclosure password

The configuration of the TotalStorage DS300 is now compete.

7.5.6 Configure the IBM iSCSI Server Adapter

This section describes the configuration of the iSCSI Server adapter.

1. Start the SANsurfer Pro on the Management station and connect to host SEINE. When you start the SANsurfer the first time it might be required to

specify the IP address or the host name of the host system running the SANsurfer iSCSI agent.

Figure 7-167 SANsurfer — Connect to host

2. Connect also to host `TIBER`.
3. Both cluster nodes are connected now.

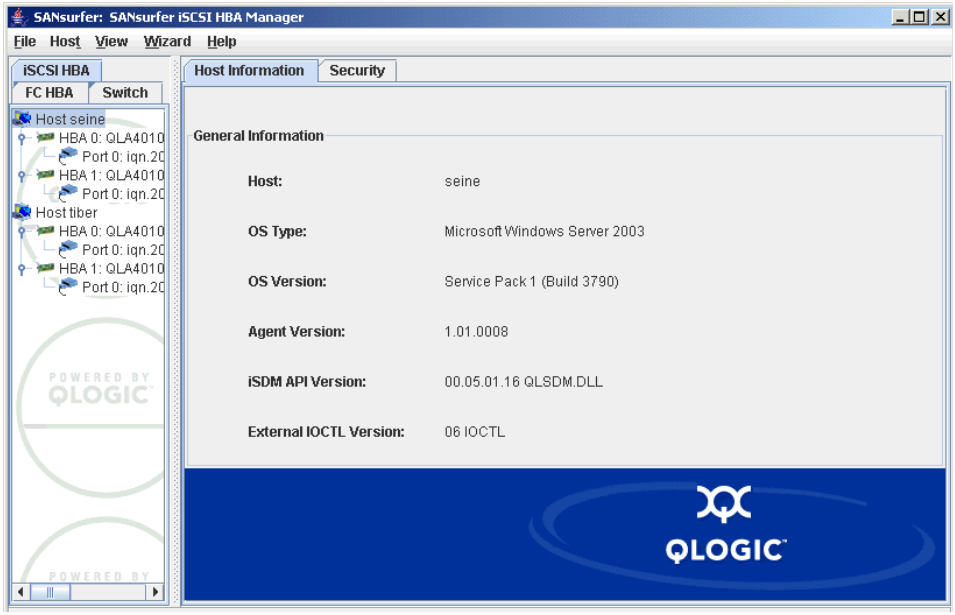

Figure 7-168 SANsurfer — Connected cluster nodes

4. Configure the network setting of all four host bus adapters.
5. Select **Port 0** of **HBA0** of host `SEINE`. Select register card **Port Options** and then **Network**. Enter the IP address `192.168.1.9`, subnet mask `255.255.255.192` and gateway `192.168.1.62` of the first host bus adapter in

the appropriate fields. Specify the primary DNS Server 172.16.100.10. Click **Save Port Settings**.

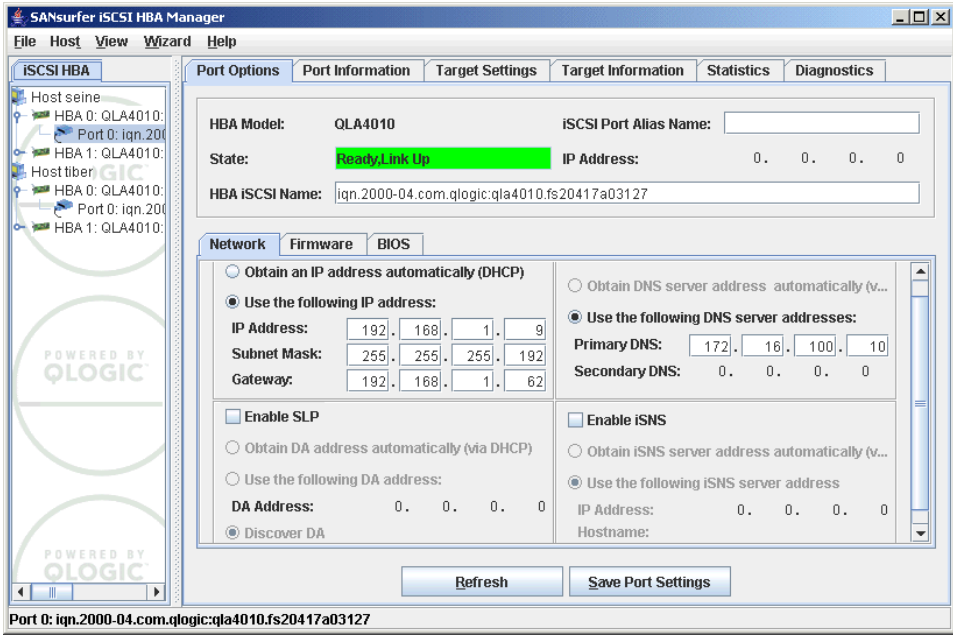

Figure 7-169 SANsurfer — Network settings of the iSCSI Server Adapter

6. Check the **Disable future warning messages for this session** checkbox and click **Yes**.

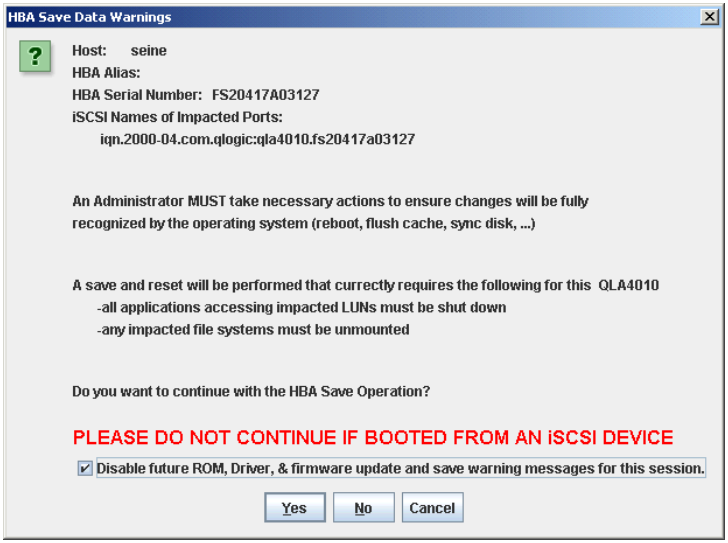

Figure 7-170 SANsurfer — Warnings

7. Enter the SANsurfer configuration password. The default password is config. Click **OK**.

Figure 7-171 SANsurfer — Security check

8. Click **Yes** to refresh the current configuration.

Figure 7-172 SANsurfer — iSCSI Configuration change

9. The network settings are now defined.

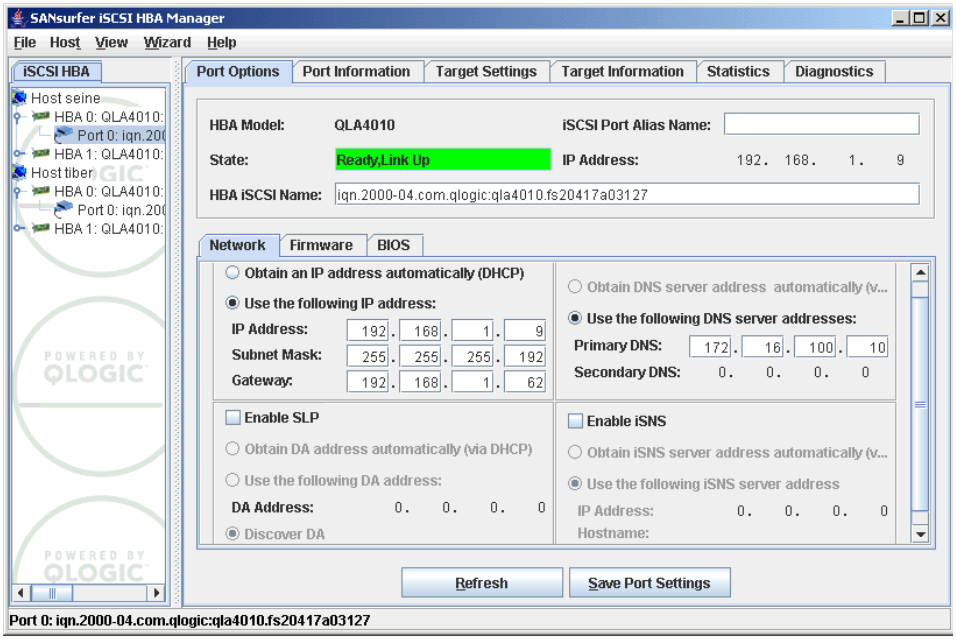

Figure 7-173 SANsurfer — Defined network settings

10. Repeat step 5 on page 334 to step 9 on page 336 for each remaining host bus adapters in both cluster nodes. See table Table 7-6 on page 322 and Table 7-7 on page 323 for network address details.

11. After the network configuration for each host bus adapter is done, the firmware settings **ARP Redirect** and **Connection Keep Alive timeout** must to be changed before LUNs can be bound to the adapter.

> **Note:** In case LUNs are already bound to an adapter, this binding needs to be deleted and reestablished to use the modified firmware settings.

Select port 0 of HBA0 of host `SEINE`. Then click register card **Port Options** and **Firmware**. Click the button **Open** to see configured firmware values.

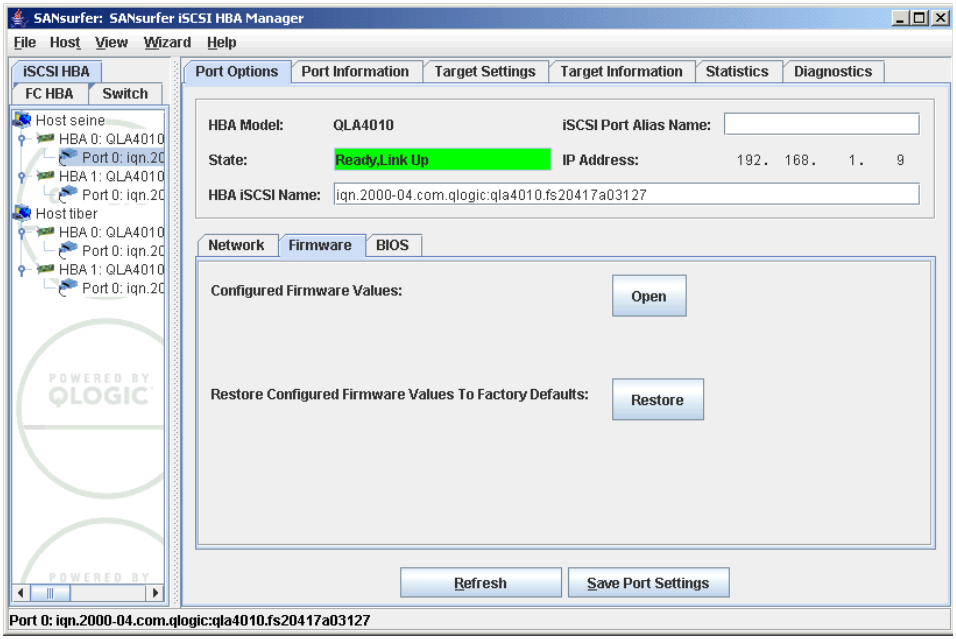

Figure 7-174 SANsurfer — HBA Port options

12. Click the Edit button to modify firmware values.

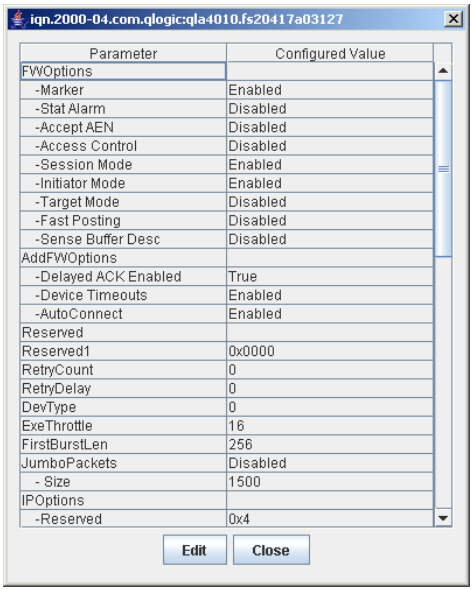

Figure 7-175 SANsurfer — Firmware values

13. Scroll to the right end and enable **ARP Redirect**. Then change **Connection Keep Alive Timeout** from 30 to 180. Click the **OK** button.

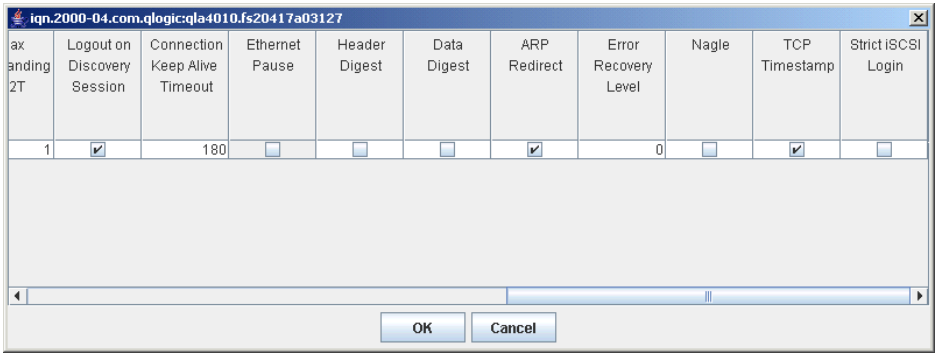

Figure 7-176 SANsurfer — Modified firmware values

14. Click **OK**.

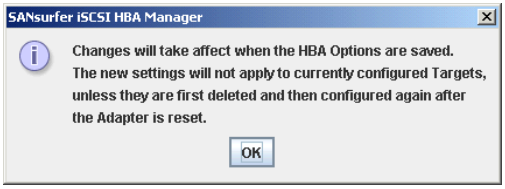

Figure 7-177 SANsurfer — Information

15. Click **Close**.

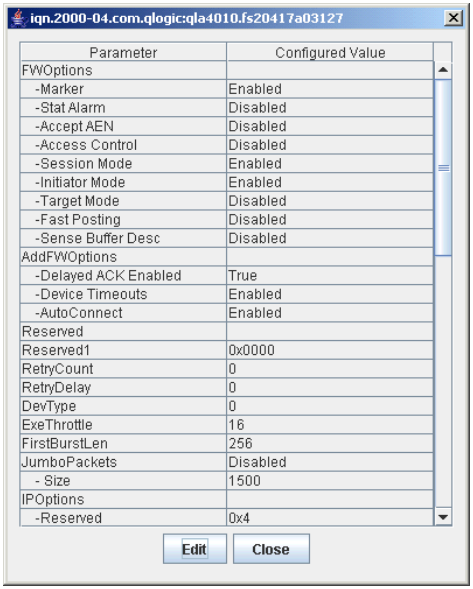

Figure 7-178 SANsurfer — Firmware values

16. Click **Save Port Settings** to save the modified firmware values.

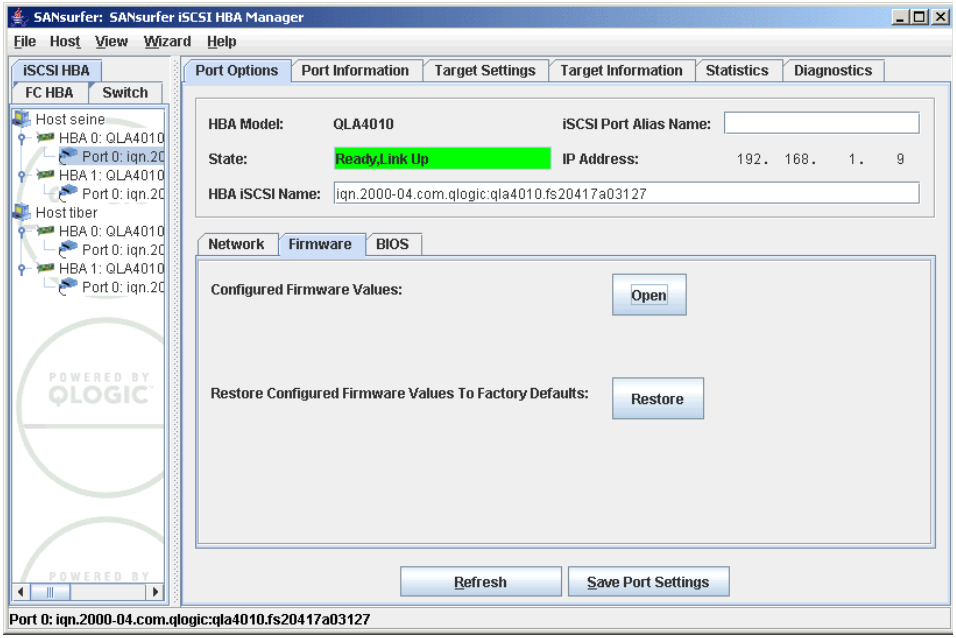

Figure 7-179 SANsurfer — Port Options

17. Click **Yes**.

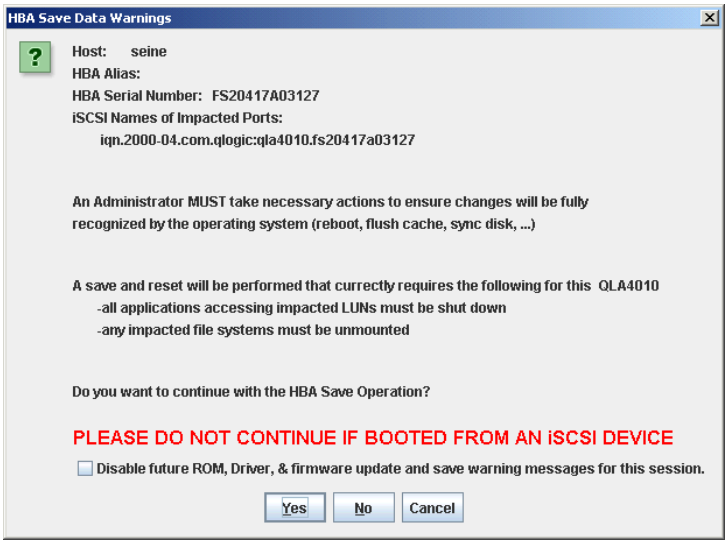

Figure 7-180 SANsurfer — Warning

18. Enter the configuration password and click **OK**.

Figure 7-181 SANsurfer — Security check

19. The configuration changes will be saved and a configuration refresh is requested by SANsurfer. Click **Yes**.

Figure 7-182 SANsurfer — iSCSI configuration change

20. Perform step 11 on page 337 to step 19 for each remaining host bus adapter in both cluster nodes.
21. Next steps will configure the iSCSI targets accessed by each host bus adapter.

22. Select **Port 0** of the first HBA of host `SEINE`. Select the register card **Target Settings**. In the target list click the green plus button to add a target portal IP.

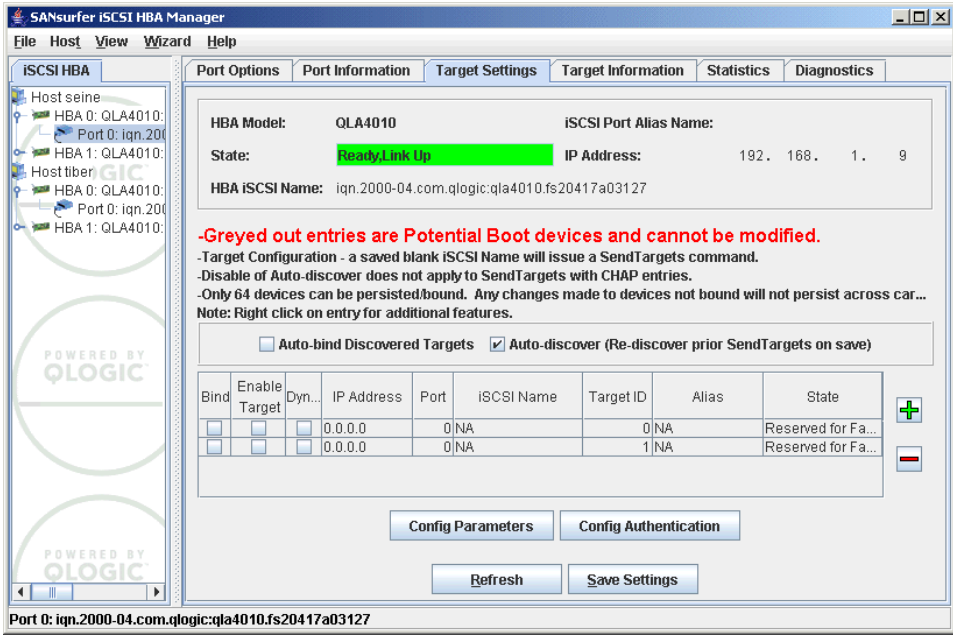

Figure 7-183 SANsurfer — Target settings

23. Enter the IP address of controller A's first iSCSI interface `192.168.1.7`.

Figure 7-184 SANsurfer — iSCSI Target IP address

24. The target list contains now an not yet saved entry for the iSCSI target 192.168.1.7. When no iSCSI qualified name is specified than this entry is used to perform a discovery of available LUNs on this target. A list of all LUNs is requested by issuing the sendtargets command. This automatic discovery is only performed when the Auto-discover checkbox is checked and the configuration is saved. Click **Save Settings**.

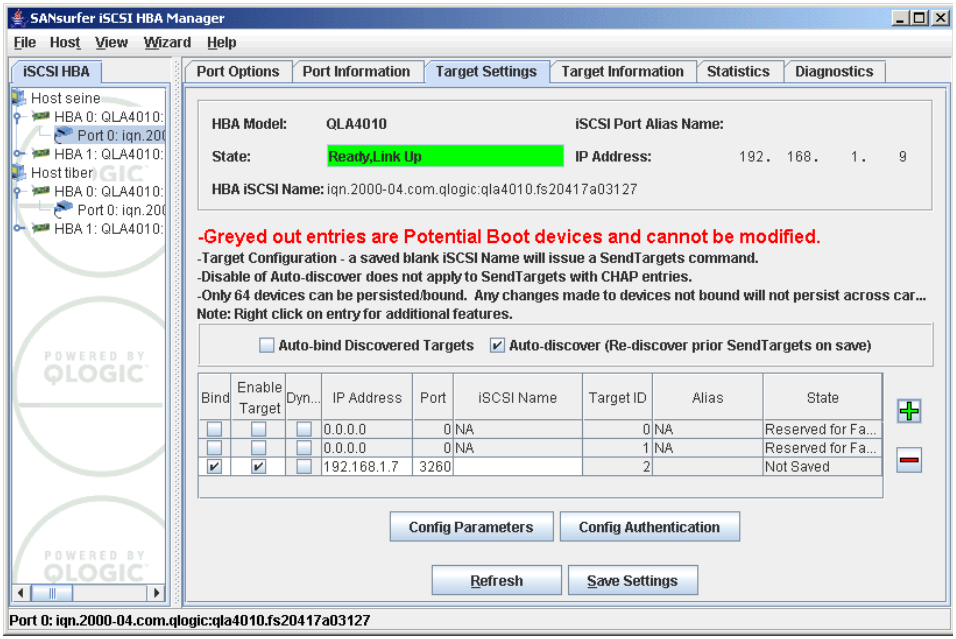

Figure 7-185 SANsurfer — iSCSI target portal

25. Enter the configuration password and click **OK**.

Figure 7-186 SANsurfer — Security check

26. The Configuration of the HBA will be refreshed by the SANsurfer utility. Click **Yes** to accept the refresh.

Figure 7-187 SANsurfer — iSCSI configuration change

27. The configured logical drive ELBE-Quorum was found but session is failed because no CHAP password was specified to access this iSCSI target. Click **Configure Authentication** to add the required authentication settings. Note

the target ID 64. Authentication configuration for this target will be configured in the next steps.

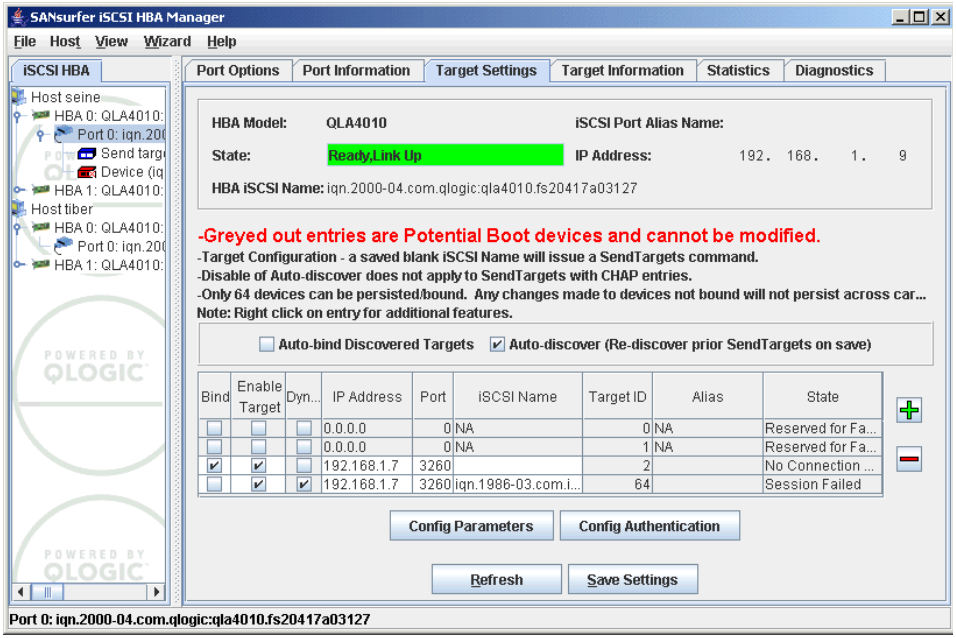

Figure 7-188 SANsurfer — Target with failed session

28. Enter the configuration password. This needs to be done because the authentication configuration contains plain text passwords.

Figure 7-189 SANsurfer — Security check

29. Define the username and password pair which will be used to authenticate the host bus adapter against the iSCSI target. Click the green plus button in the **CHAP Entries** list to add a empty line in this list.

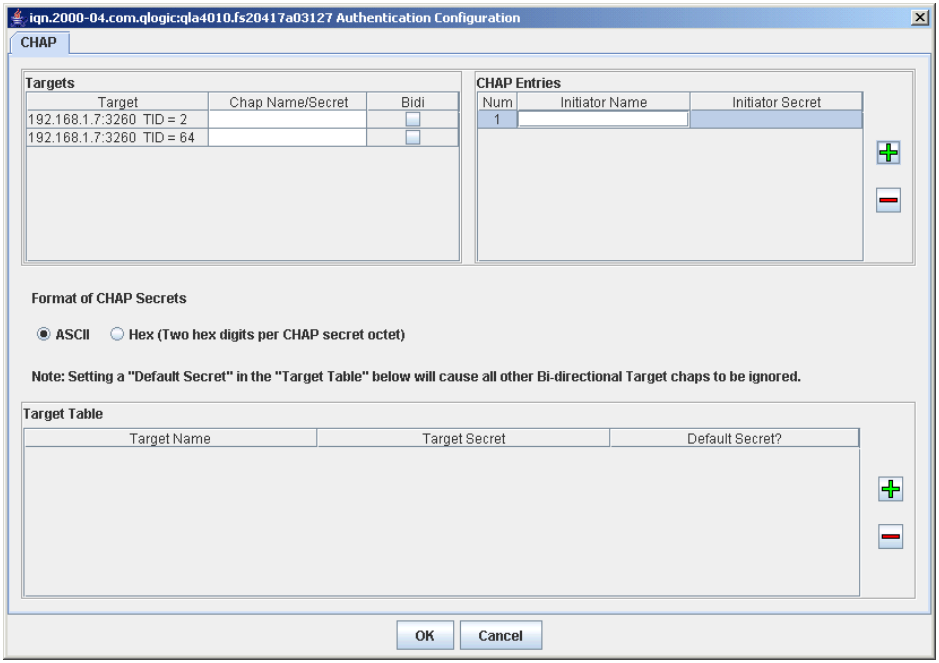

Figure 7-190 SANsurfer — Authentication configuration

30. Enter the username and password in the CHAP entries list in the new added line. The username can be any name or the iSCSI qualified name of an adapter. The password needs to be at least twelve characters long. Figure 7-191 on page 346 shows the username `elbe` and the password `a1234567890a` in the CHAP entries list.

    The specified CHAP entry can now be selected for the target in the target list. Select the CHAP Name/secret only for the target with the target ID 64.

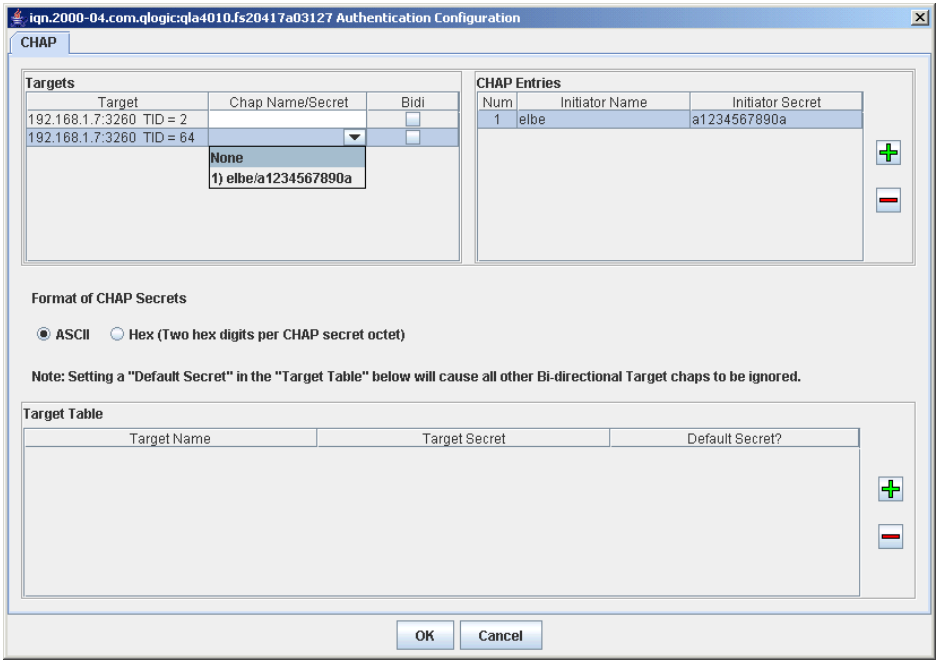

Figure 7-191 SANsurfer — Authentication configuration

31. Check the Bidi box for target 192.168.1.7. This enables the authentication of the storage subsystem against the initiator. Then click the green plus button in the targets table to add an empty line in this table.

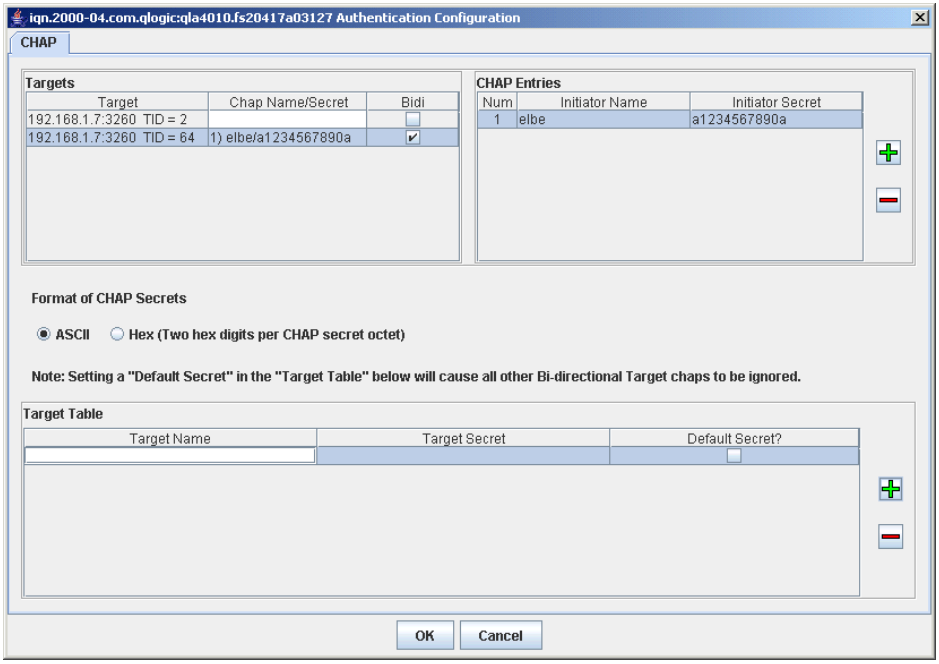

Figure 7-192 SANsurfer — Authentication Configuration

32. Enter the username in the field **Target Name**. Add the used password in the field **Target Secret**. Then click **OK**.

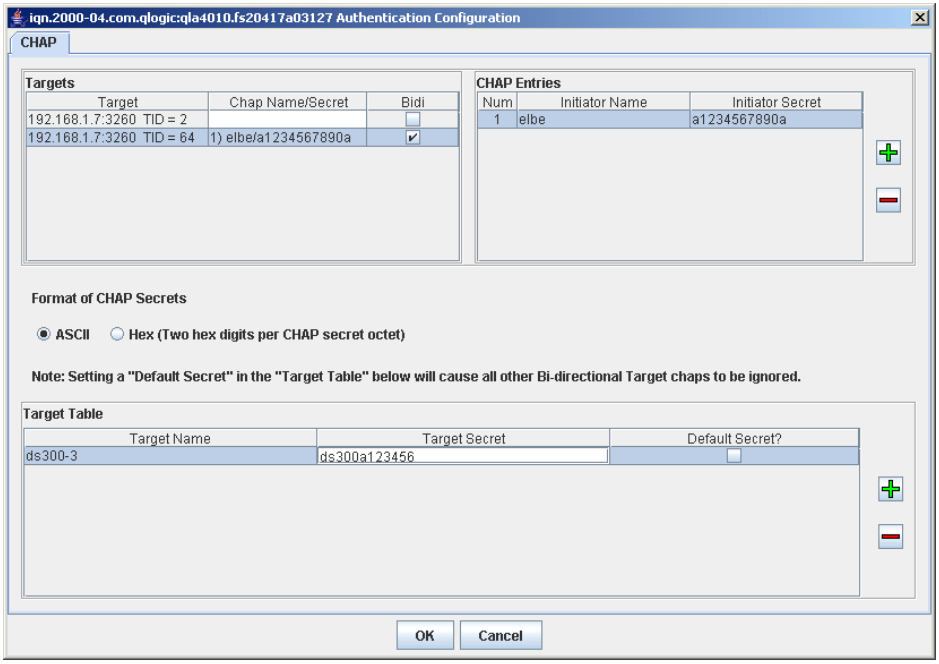

Figure 7-193 SANsurfer — Authentication Configuration

33. Check the **Bind** checkbox in the line of target ID 64 and then save the configuration by clicking the **Save Settings** button.

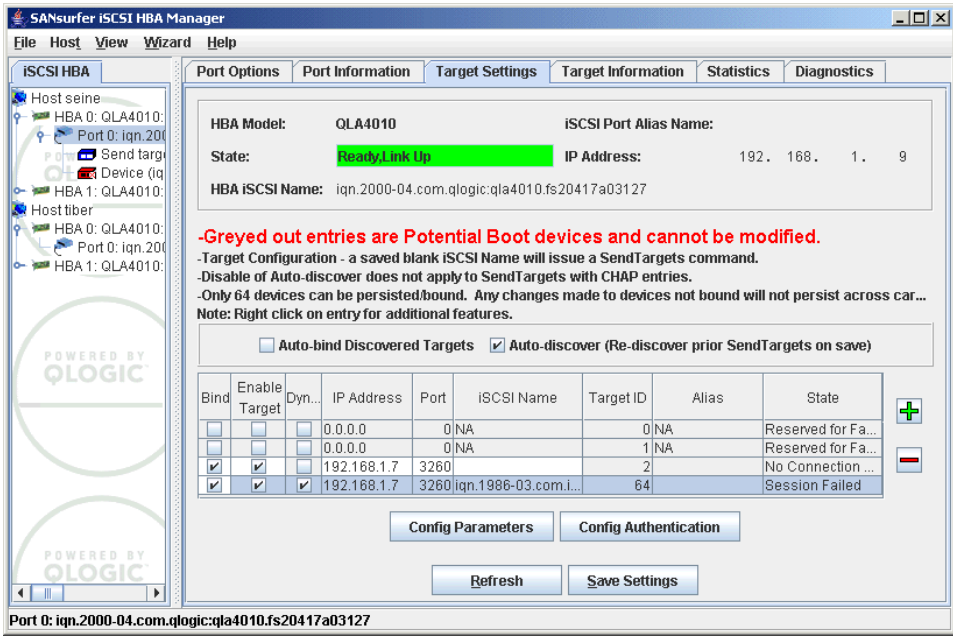

Figure 7-194 SANsurfer — Target settings

34. Enter the configuration password and click **OK**.

Figure 7-195 SANsurfer — Security check

35. The configuration changes will be saved and a configuration refresh is requested by SANsurfer. Click **Yes**.

Figure 7-196 SANsurfer — iSCSI configuration change

36. After the configuration refresh the target ID 64 was changed to ID 3 and the session to ELBE-Quorum is active.

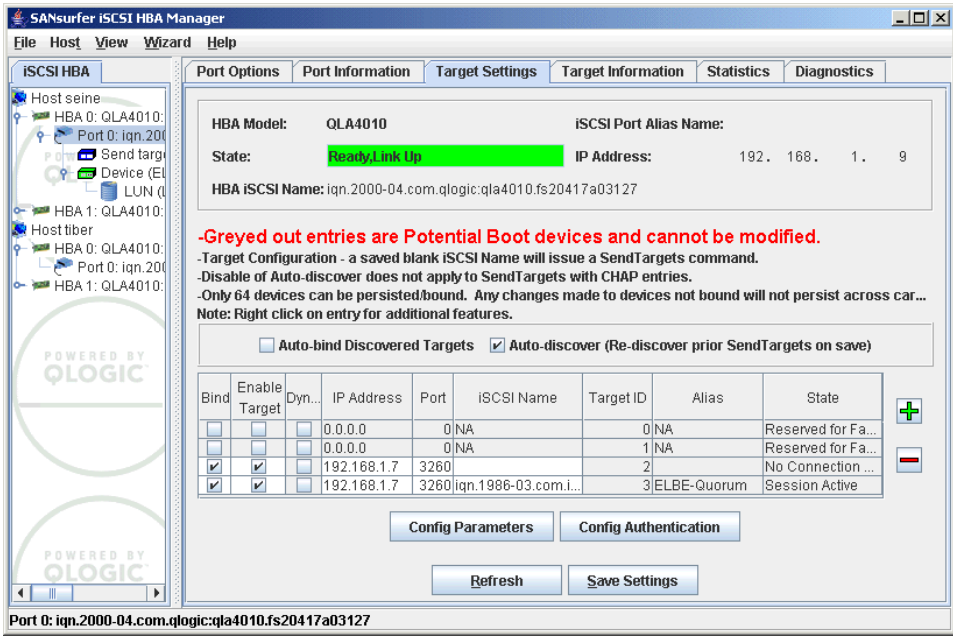

Figure 7-197 SANsurfer — Target settings

Note: When the CHAP name and secret are specified for the send target host entry, the settings of the CHAP name and secret and the bidirectional authentication are used also for all automatically discovered targets.

37. Repeat step 22 on page 342 to step 36 on page 349 for the first iSCSI interface of controller B.

38. The configuration looks like window shown in Figure 7-198 on page 351.

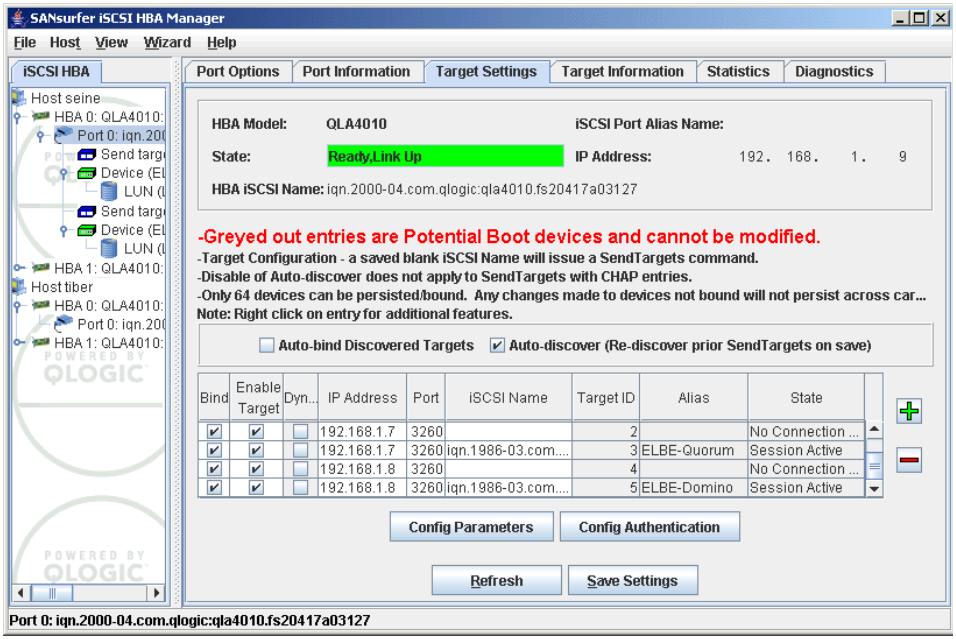

Figure 7-198 SANsurfer — first HBA of host SEINE with two iSCSI LUNs

39. Repeat step 22 on page 342 to step 38 on page 350 for the remaining host bus adapter in the cluster nodes `SEINE` and `TIBER`.

40. The final configuration of all iSCSI targets looks like this. All targets are defined for both hosts and the cluster service can be installed in the next step.

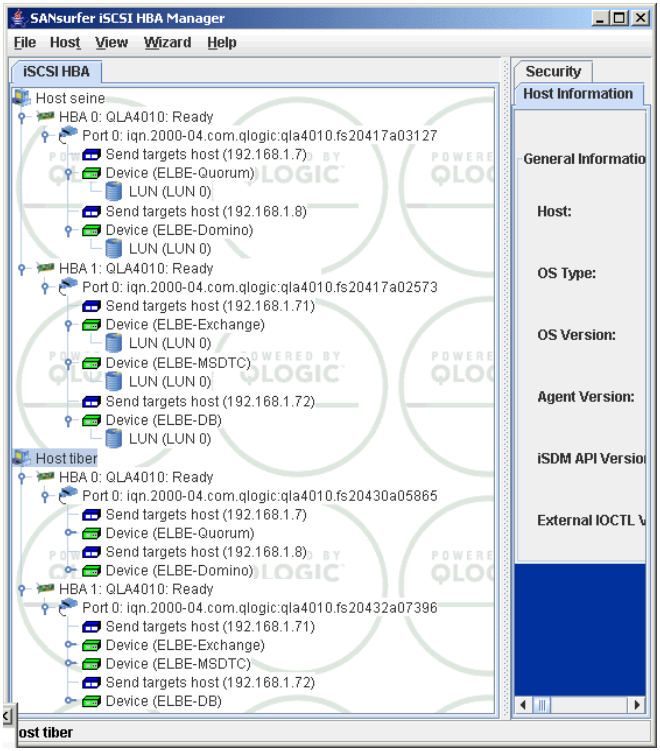

Figure 7-199 SANsurfer — Complete configuration

Note: The Send target host entries might be deleted now. Also the auto discovery function might be disabled. But even when a new iSCSI target is assigned to the initiators defined on the TotalStorage DS300, those targets will not be automatically accessible by the host. The CHAP authentication and the not defined CHAP name and secret for the send target host entries will prevent this.

The next step is the partitioning and formatting of the drives.

7.5.7 Initialize, partition and format drives

Before the cluster service can be installed, all shared disks needs to be initialized, partitioned and formatted with a NTFS file system. Drive letters must be assigned when no mount points are planned to be used. Perform this tasks only on node SEINE.

1. Open the **Computer Management**, select **Storage** and click **Disk Management**.

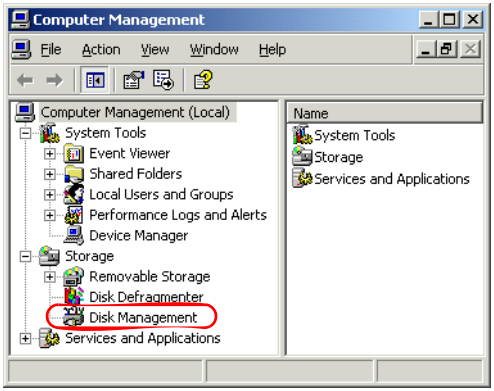

Figure 7-200 Computer Management (Local)

2. The **Initialize and Convert Disk Wizard** starts. Click **Next**.

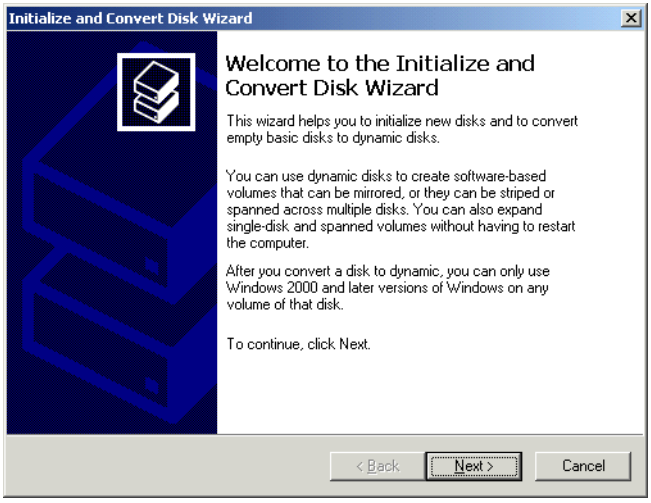

Figure 7-201 Initialize and Convert Disk Wizard

3. Select the disks to initialize them. Click **Next**.

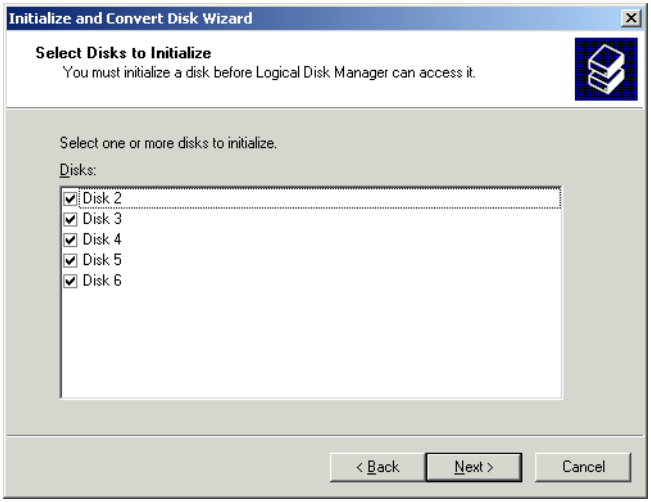

Figure 7-202 Initialize and Convert Disk Wizard — Select disks to initialize

4. Do not select disks in the convert dialog. Click Next.

Note: Dynamic disks are not supported in a cluster environment or with a Microsoft Software Initiator.

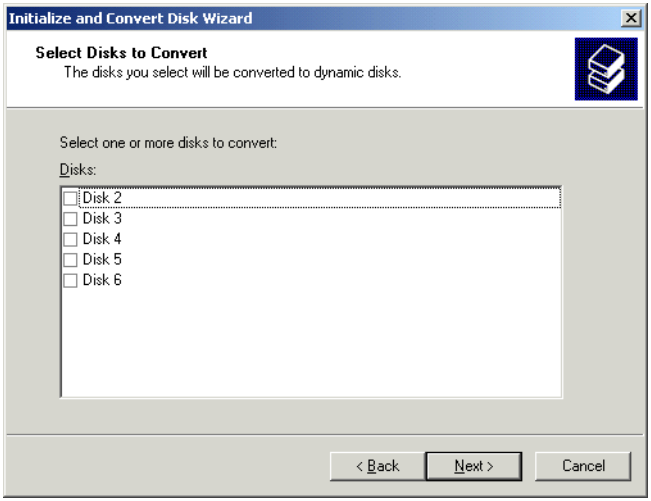

Figure 7-203 Initialize and Convert Disk Wizard — Select disks to convert

5. Verify the chosen settings and click **Finish**.

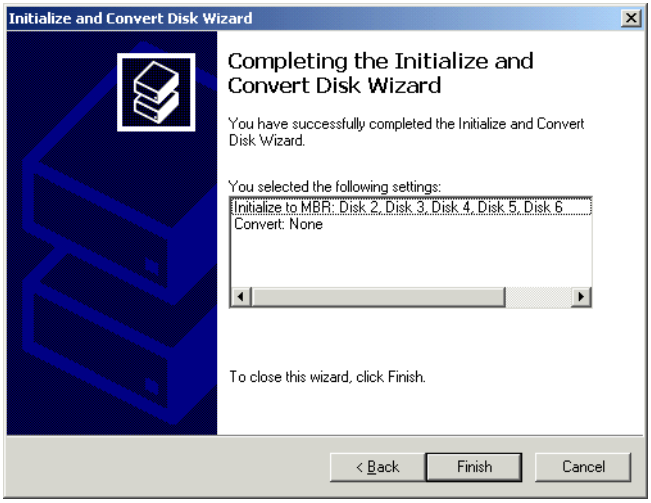

Figure 7-204 Initialize and Convert Disk Wizard — Verify settings

6. All drives are ready to be partitioned and formatted. Right-click the first drive. From the context menu select New Partition.

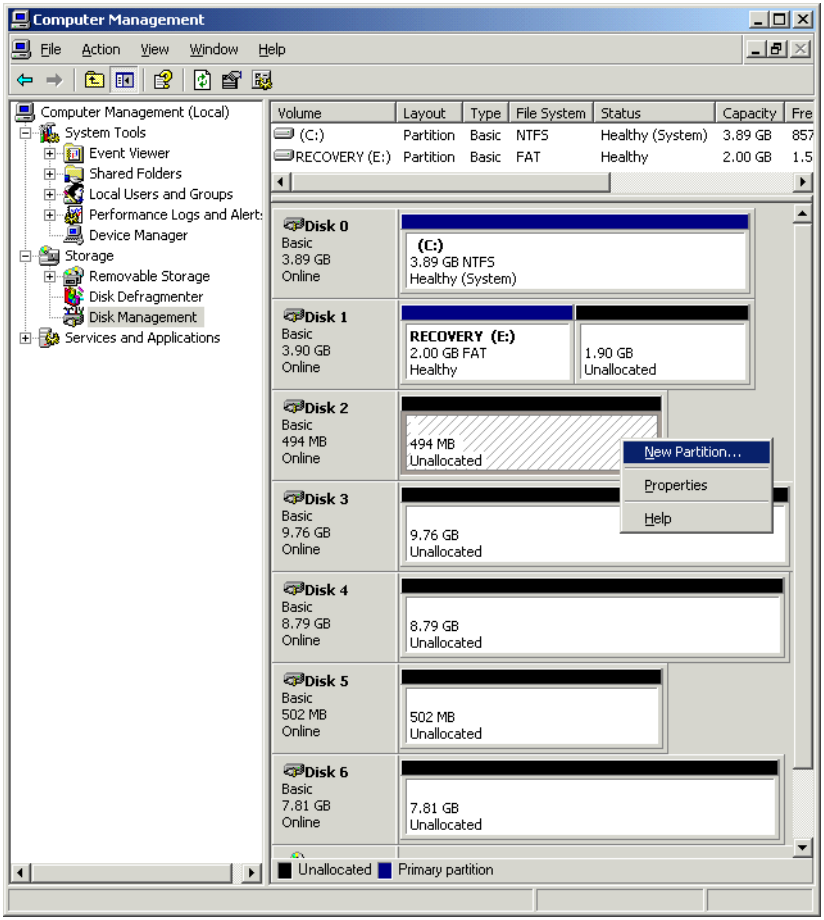

Figure 7-205 Disk Management — Initialized drives

7. The New partition wizard starts. Click **Next**.

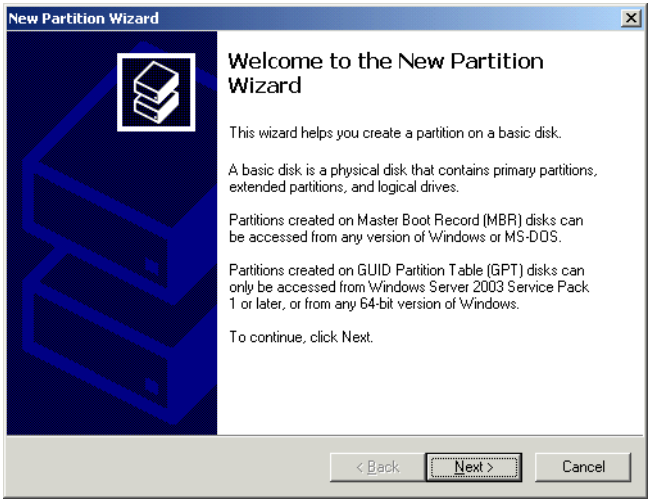

Figure 7-206 New Partition Wizard

8. Select **Primary partition** and click **Next**.

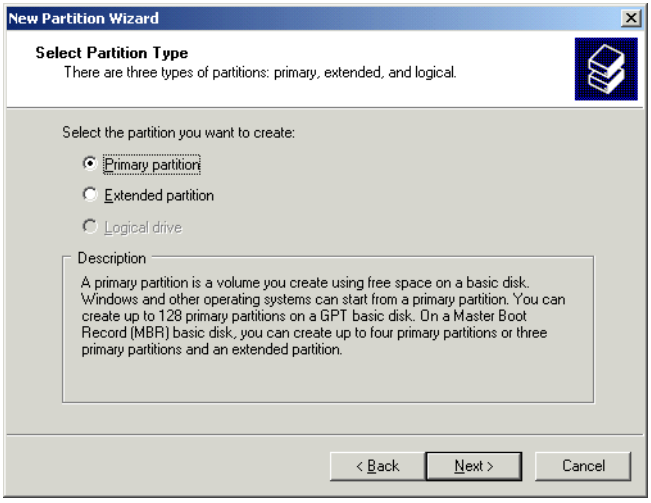

Figure 7-207 New Partition Wizard — Partition type

9. Choose the partition size and click **Next**.

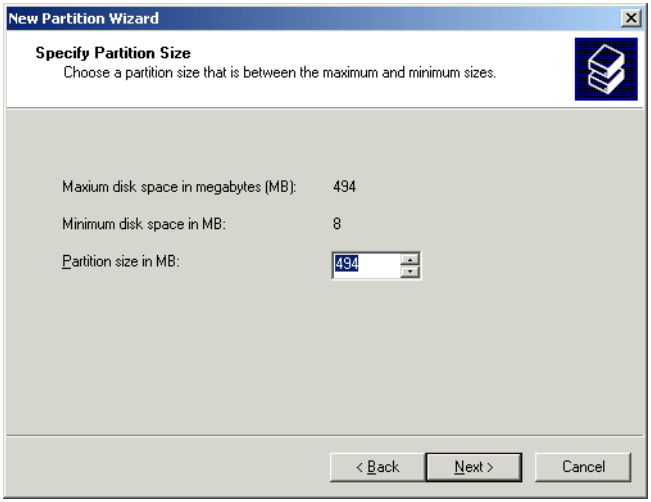

Figure 7-208 New Partition Wizard — Partition size

10. Select the drive letter from the drop down list which should be assigned to this partition.

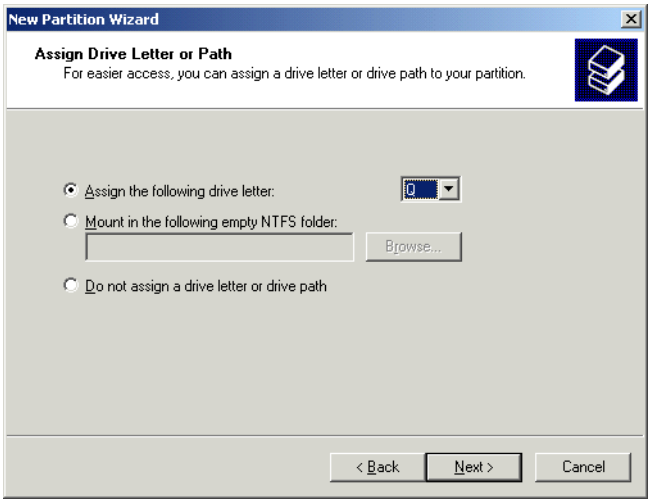

Figure 7-209 New Partition Wizard — Assign drive letter

11. Choose NTFS as file system. This is a requirement by the cluster server. Check quick format and enter a descriptive volume label. The volume label should contain the drive letter assigned to this partition. This helps in a clustered environment to verify with one view on the other node or nodes the assignment of the drives. Click **Next**.

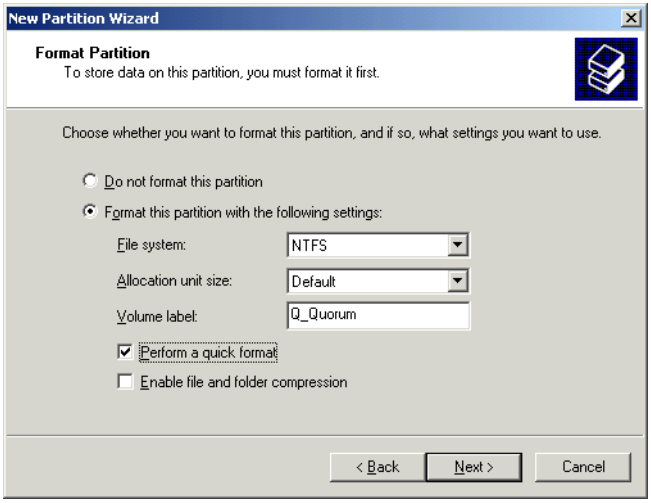

Figure 7-210 New Partition Wizard — Format partition

12. Verify the settings and confirm them by clicking **Finish**.

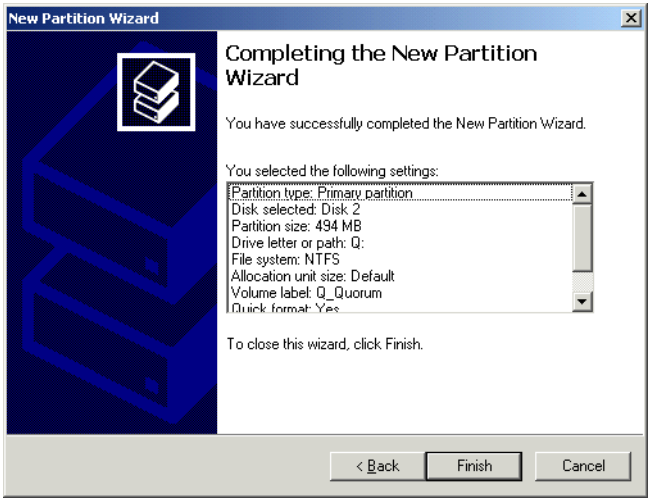

Figure 7-211 New Partition Wizard — Verify settings

13. The drive is now partitioned and formatted.

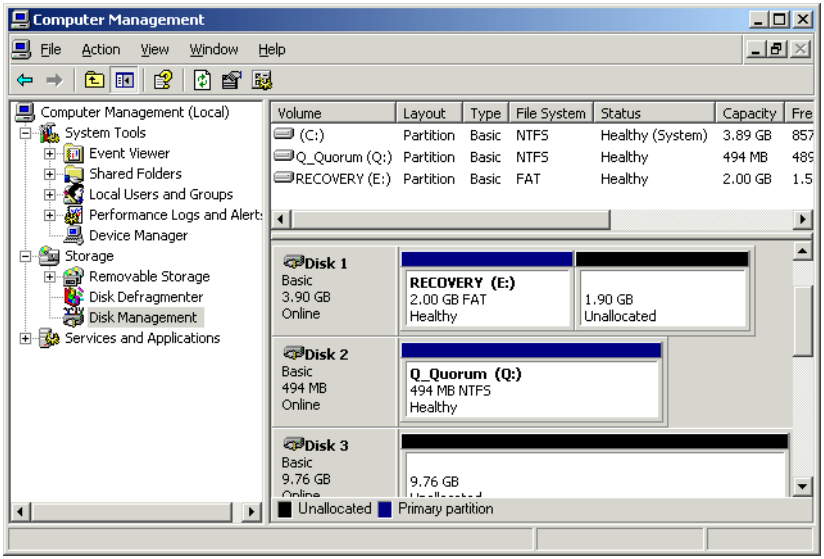

Figure 7-212 Disk Management — Partitioned and formatted quorum drive

14. Repeat step 6 on page 355 to step 13 on page 359 for the remaining unpartitioned and unformatted drives.

All drives are now ready to be used in the cluster. The next step is the installation of the cLuster service.

7.5.8 Installation of the Cluster Service

This section explains the installation of the cluster service. Install the cluster service on the first node of the cluster where also all drives are partitioned, formatted and configured with drive letters.

1. Start the Cluster Administrator. Click **Start**->**Run**. Enter cluadmin in the open field and click **OK**.

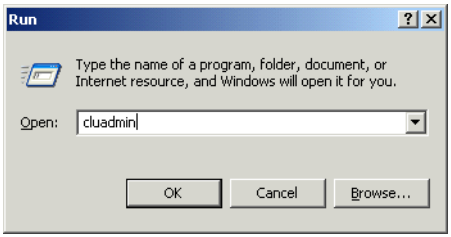

Figure 7-213 Open cluster administrator

2. Select **Create new cluster** and click **OK**.

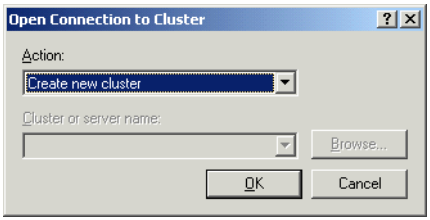

Figure 7-214 Cluster Administrator — Action create new cluster

3. The New Server Cluster Wizard starts. Click **Next**.

Figure 7-215 New Server Cluster Wizard

4. Choose domain and enter the cluster name. This is the name under which the cluster itself is accessible in the network. Click **Next**.

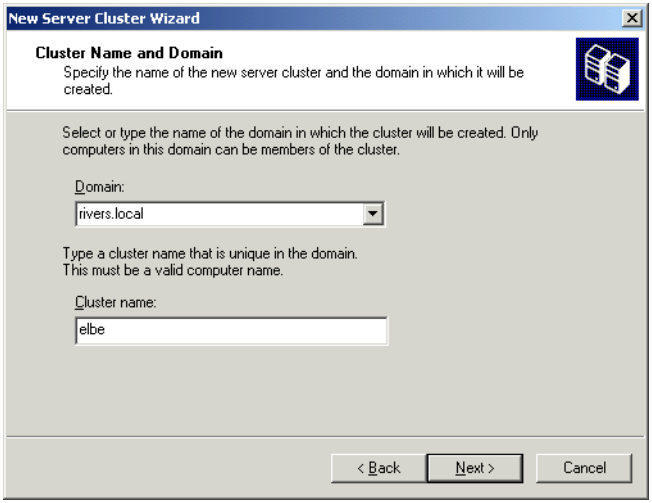

Figure 7-216 New Server Cluster Wizard — Cluster name and domain

5. Enter the host name of the first cluster node. Under Advanced can be a minimal test chosen. This helps to install the cluster when some tests fail and the cluster setup means that this is critical and the setup is stopped.

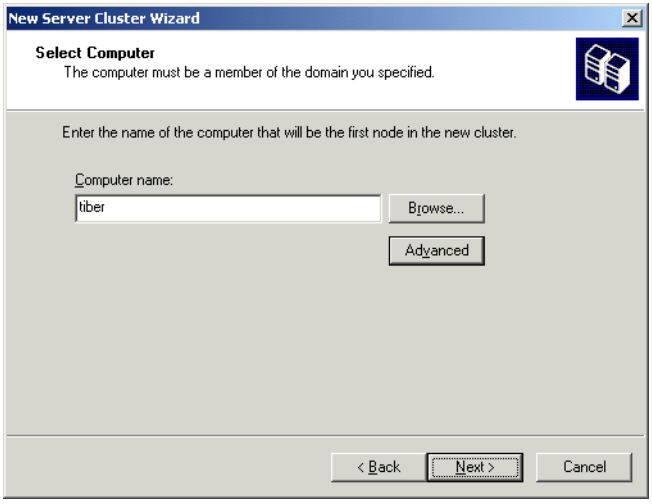

Figure 7-217 New Server Cluster Wizard — Select node

6. The cluster setup performs some test to verify that the hardware and software setup is able to operate a cluster. The warnings in this screen indicating that drives C: and E: are not manageable by the cluster service. This is correct,

because there are two partitions on the drive where the operating system is installed. Click **Next**.

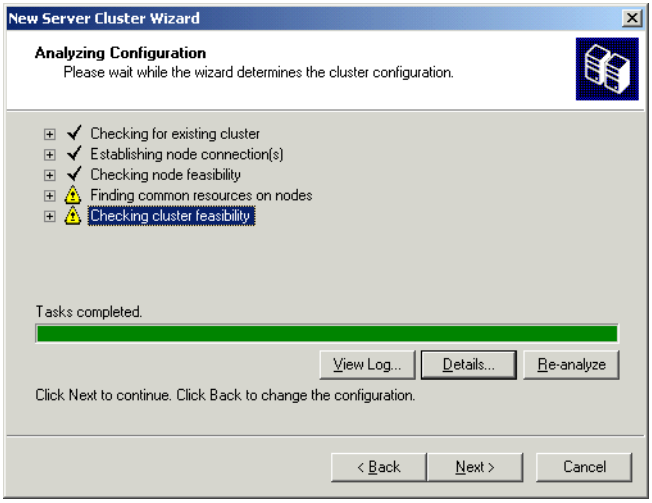

Figure 7-218 New Server Cluster Wizard — Analyzing configuration

7. Enter the IP address of the server cluster. Click **Next**.

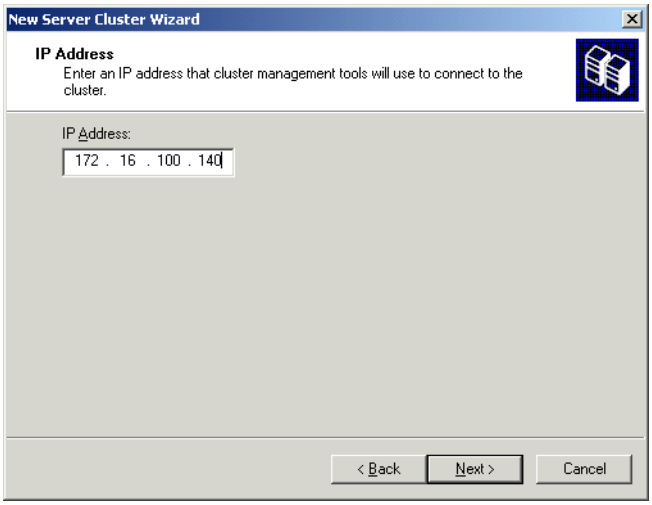

Figure 7-219 New Server Cluster Wizard — Cluster IP address

8. Specify a domain account which will be used as service account for this cluster. During the setup this account will be configured with all required rights. Click **Next**.

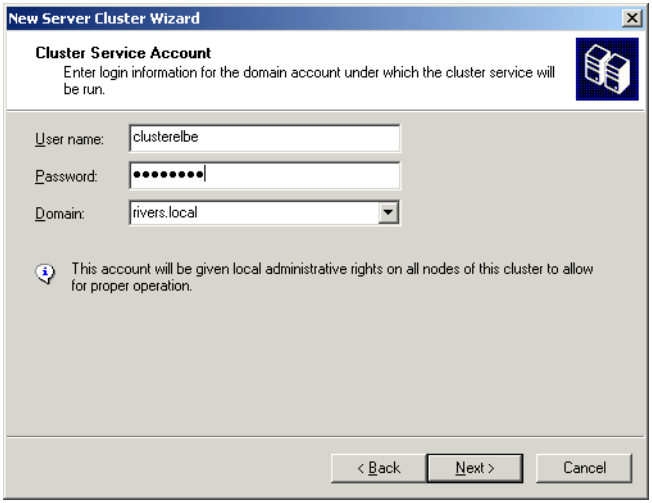

Figure 7-220 New Server Cluster Wizard — Cluster service account

9. Click **Quorum** to define the quorum drive.

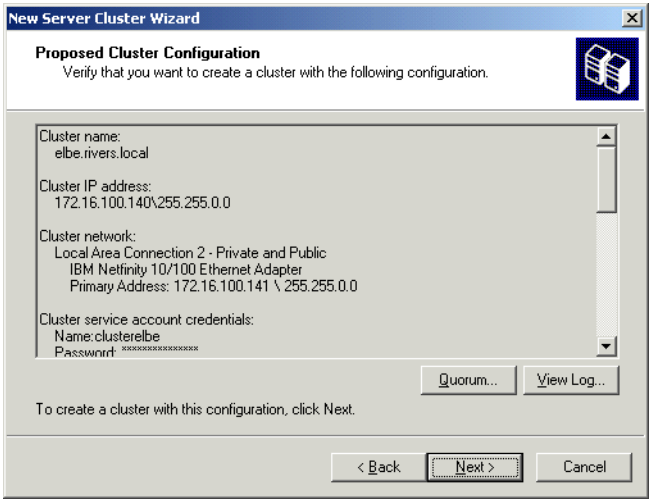

Figure 7-221 New Server Cluster Wizard — Proposed cluster configuration

10. Select the quorum drive from the list of available shared drives. Click **OK**.

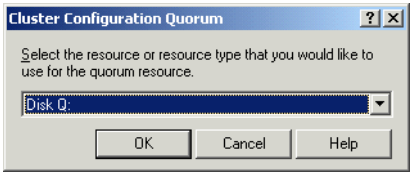

Figure 7-222 New Server Cluster Wizard — Cluster configuration quorum

11. Verify the settings and click **Next**.

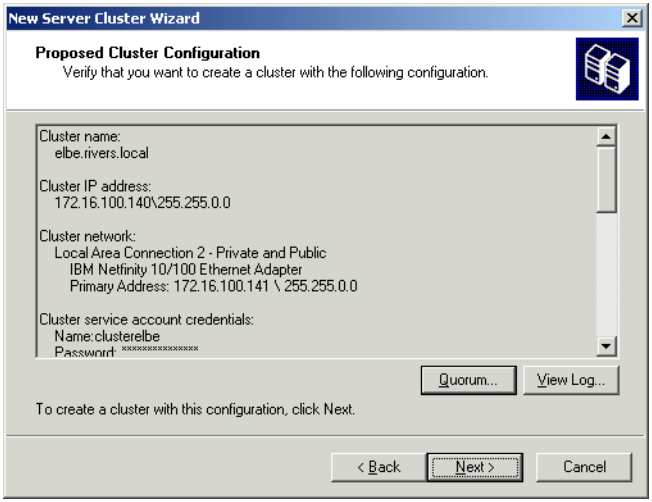

Figure 7-223 New Server Cluster Wizard — Proposed cluster configuration

12. Click **Next** after the cluster was successfully created.

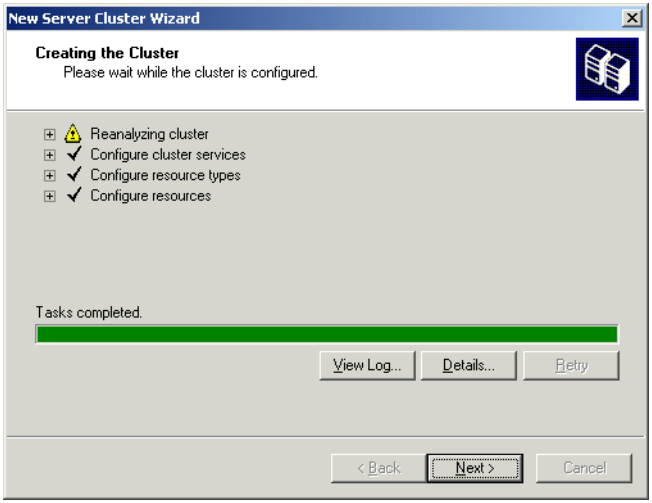

Figure 7-224 New Server Cluster Wizard — Creating the cluster

13. Verify the log file by clicking **View Log**. Click **Finish** to close the New Server Cluster Wizard.

Figure 7-225 New Server Cluster Wizard — Finish

14. The cluster administrator connects to the newly created cluster.

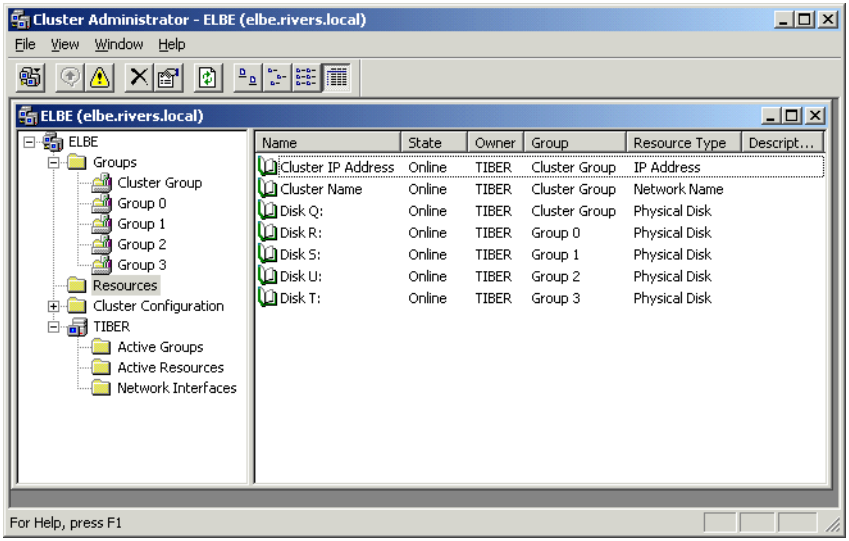

Figure 7-226 Cluster administrator connected to cluster ELBE

The cluster service now running on the first node. The cluster service must now be installed on the remaining nodes.

1. Open menu **File** → **Open Connection**.

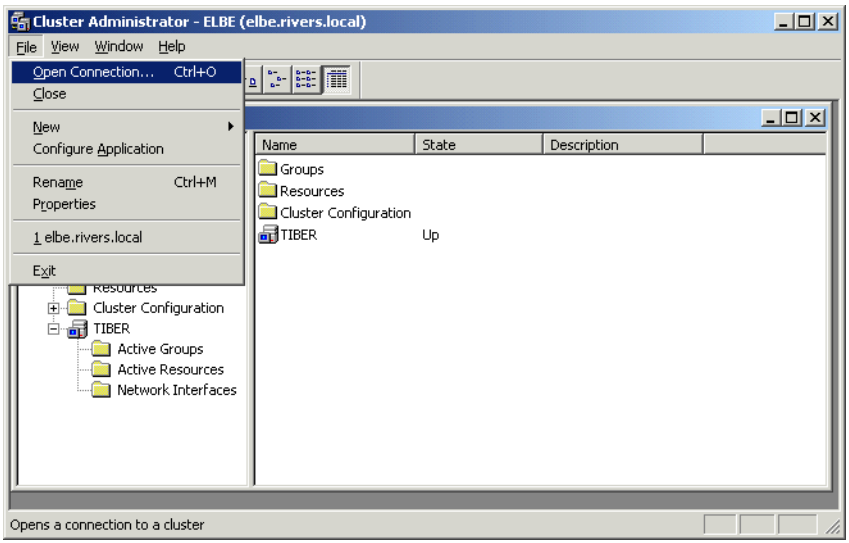

Figure 7-227 Cluster administrator — Menu file

2. Select **Add nodes to cluster** and click **OK**.

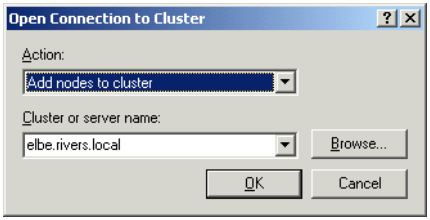

Figure 7-228 Cluster Administrator — Action add nodes to cluster

3. The Add Nodes Wizard starts. Click **Next**.

Figure 7-229 Add Nodes Wizard

4. Click **Browse** to browse for computers in the domain.

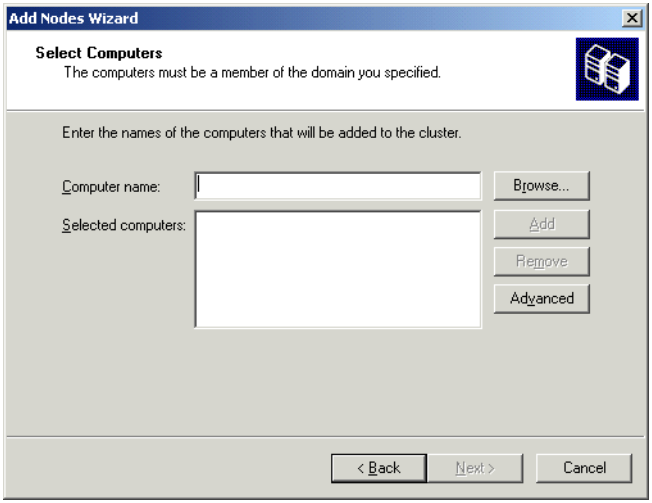

Figure 7-230 Add Nodes Wizard — Select computers

5. Click **Advanced**.

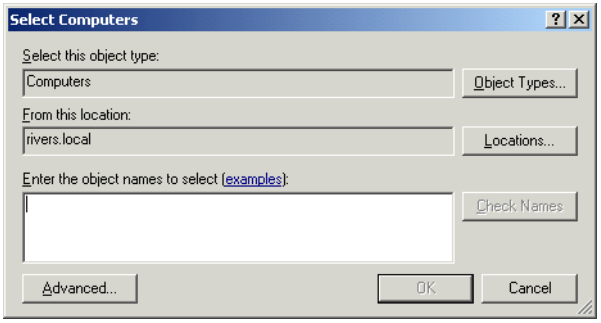

Figure 7-231 Select computers

6. Click **Find now**. This provides a list of all computers in the domain. A search term can be added to the Name field to limit the results returned by **Find Now**.

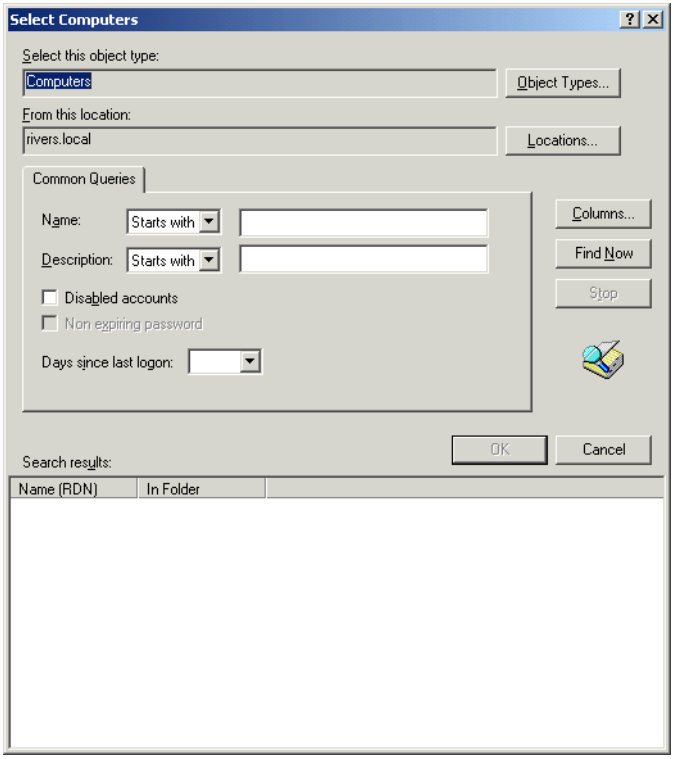

Figure 7-232 Select computers

7. Select those computers which will be added to the cluster as new nodes. Click **OK**.

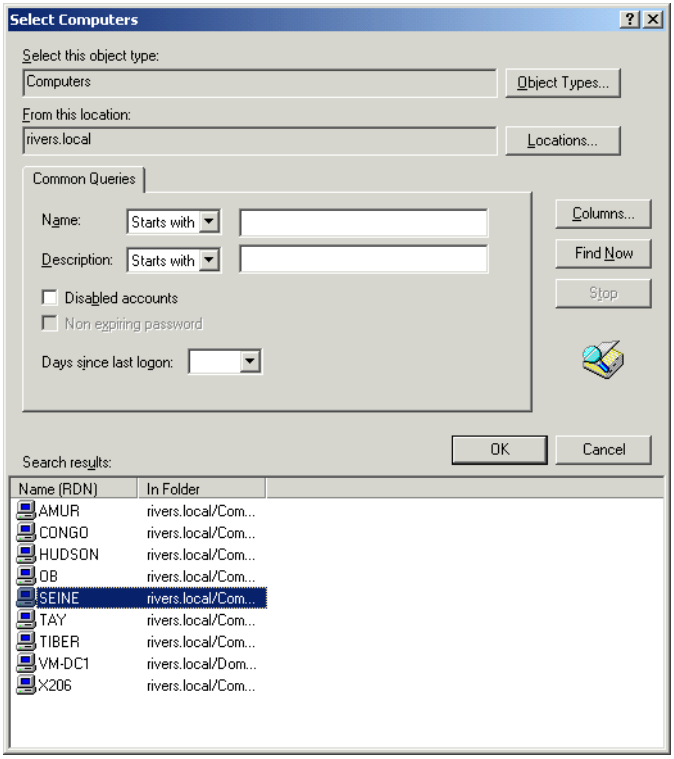

Figure 7-233 Select computers

8. The selected computers are added to the list of selected computers. Click **OK**.

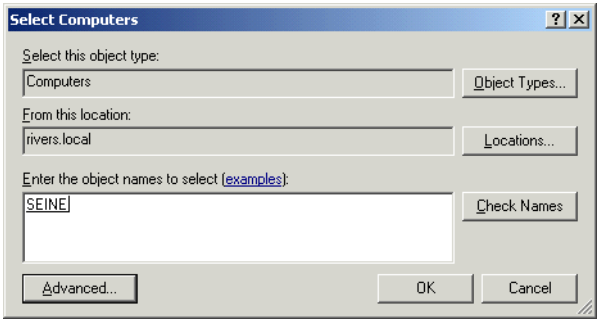

Figure 7-234 Select computers

9. Click **Add**.

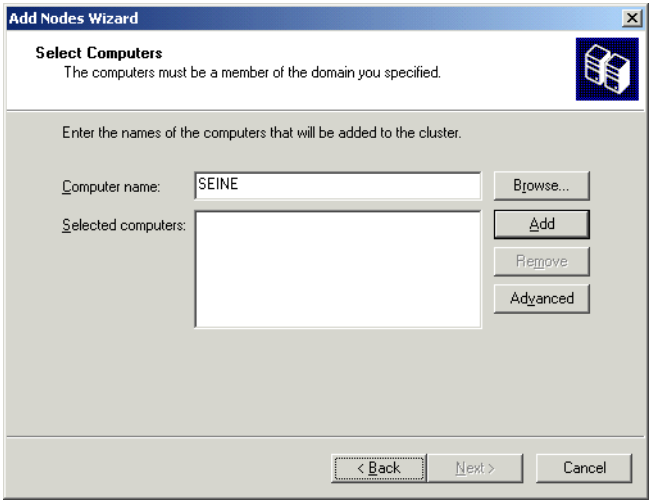

Figure 7-235 Add Nodes wizard — Add computer as new nodes

10. Click **Next**.

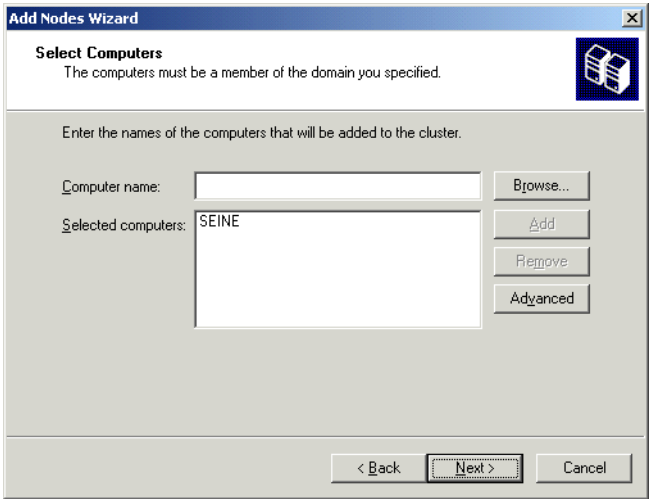

Figure 7-236 Add Nodes wizard — Added computer as new nodes

11. The Add Nodes Wizard analyzes the configuration. The shown warnings are the local C: and E: drives which are not managed by the cluster. After the successful analysis click **Next**.

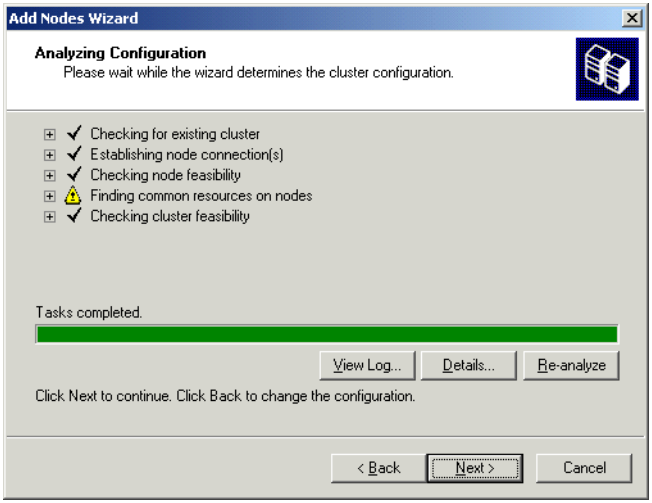

Figure 7-237 Add Nodes Wizard — Analyzing configuration

12. Enter the password of the cluster service account. The same account was specified during the installation of the first node. Click **Next**.

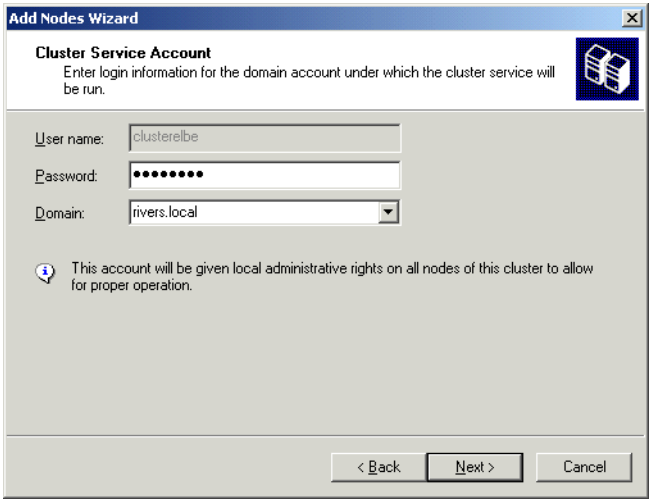

Figure 7-238 Add Nodes Wizard — Cluster service account

13. Verify settings and click **Next**.

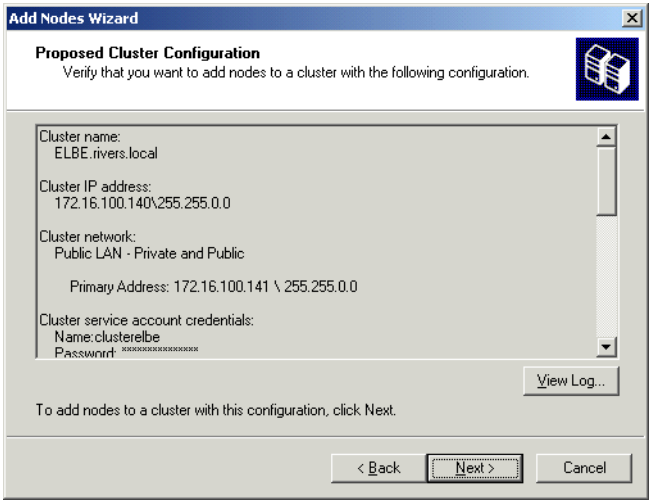

Figure 7-239 Add Nodes Wizard — proposed cluster configuration

14. After all specified nodes successfully are added to the cluster click **Next**.

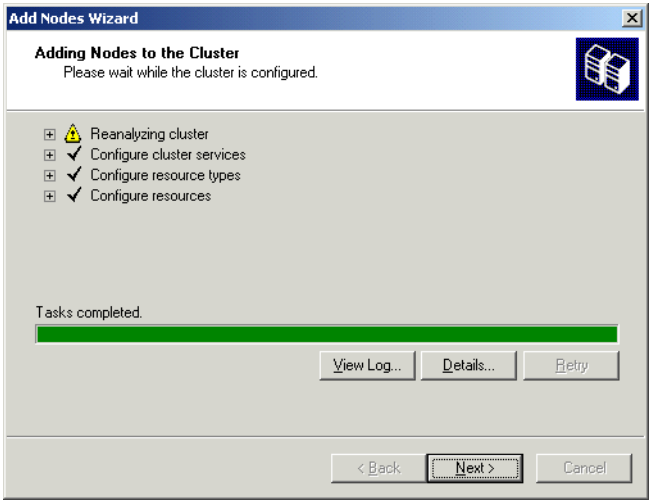

Figure 7-240 Add Nodes Wizard — Adding nodes to the cluster

15. Verify the log for any problems that might have occurred and click **Finish**.

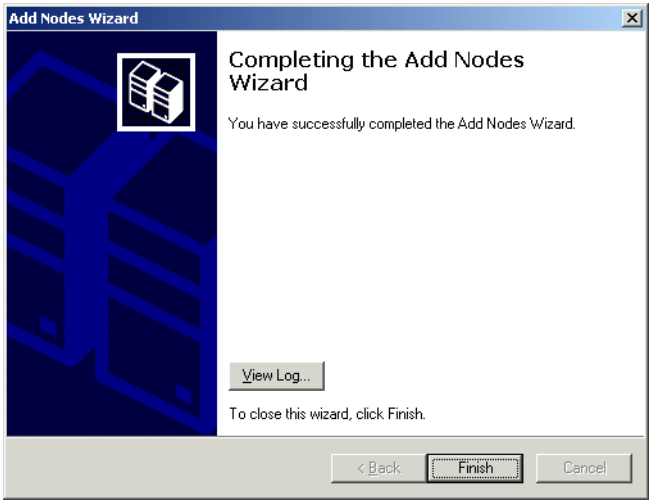

Figure 7-241 Add Nodes Wizard — Finished

16. The cluster administrator shows the added nodes.

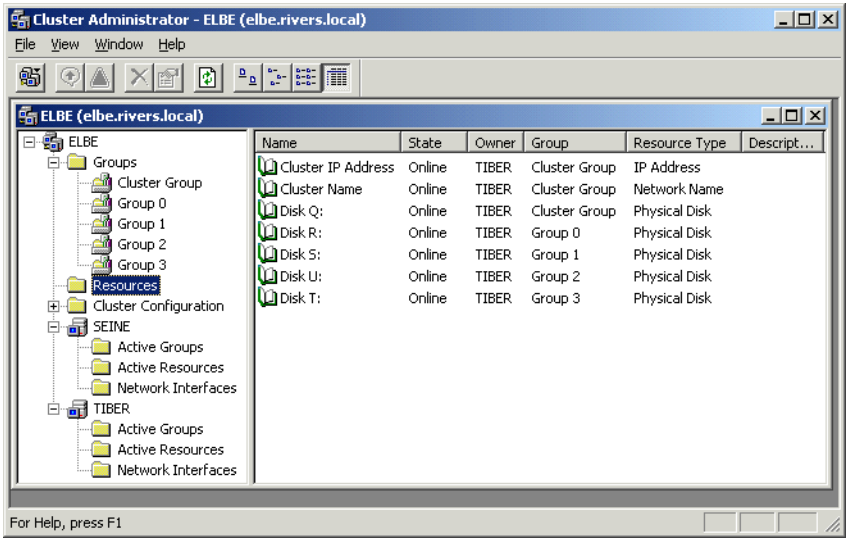

Figure 7-242 Cluster administrator with new added nodes

The Cluster service is now on all nodes configured. The last step is the configuration of the network priority and the failover tests.

1. Right-click the cluster name. From the context menu select **Properties**.

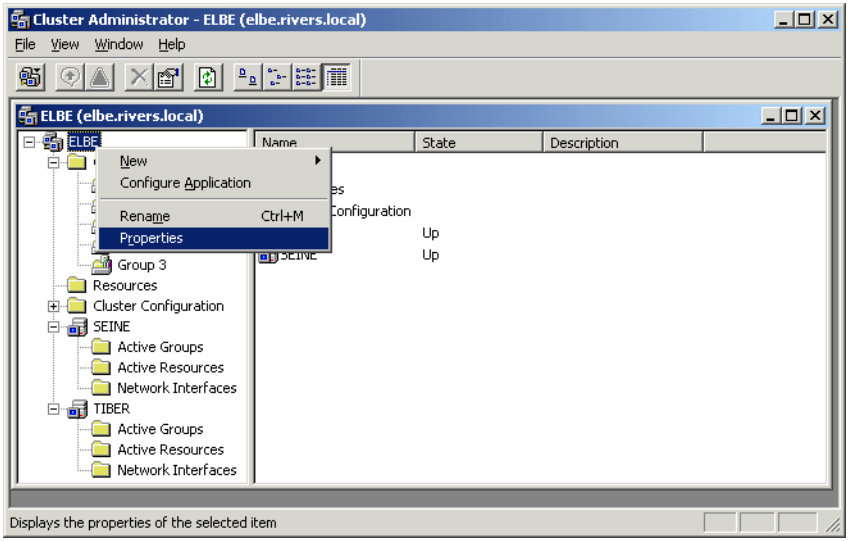

Figure 7-243 Cluster administrator — Menu cluster properties

2. Select register card Network priority.

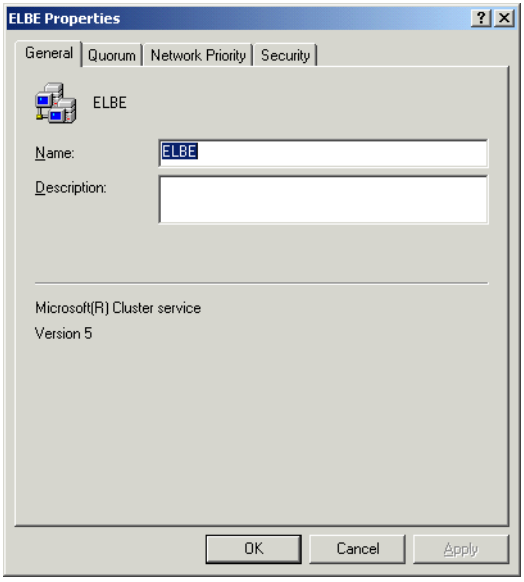

Figure 7-244 General cluster properties

3. Select the Heartbeat network and click **Properties**.

Note: The shown network names are the names used for the network connections and can differ from cluster to cluster.

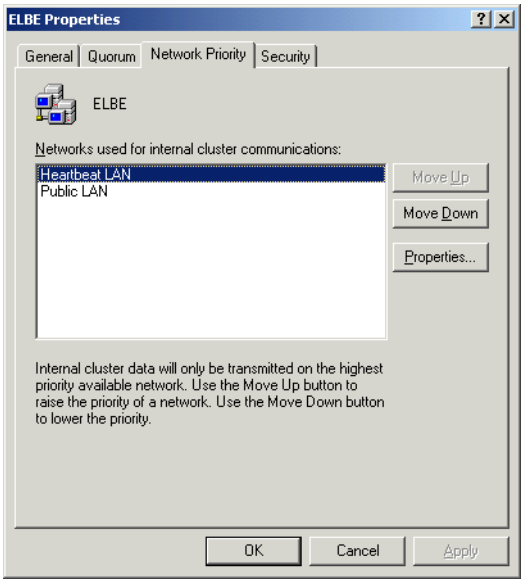

Figure 7-245 Cluster network priority properties

4. Change the role of the heartbeat network from **All communications** to **Internal cluster communications only**. Click **OK**.

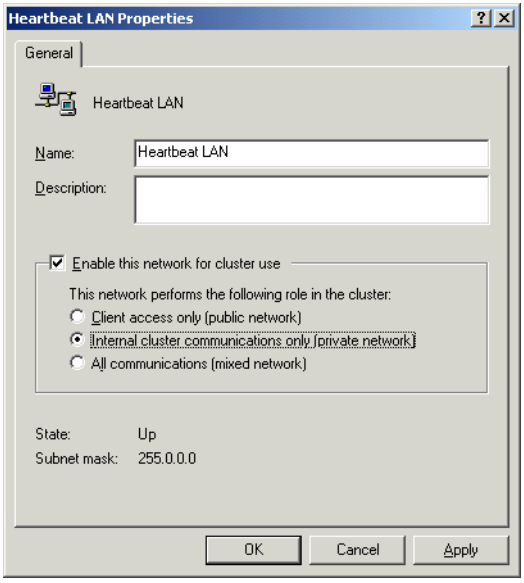

Figure 7-246 Heartbeat properties

5. Click OK to close the cluster properties and return to the cluster administrator.

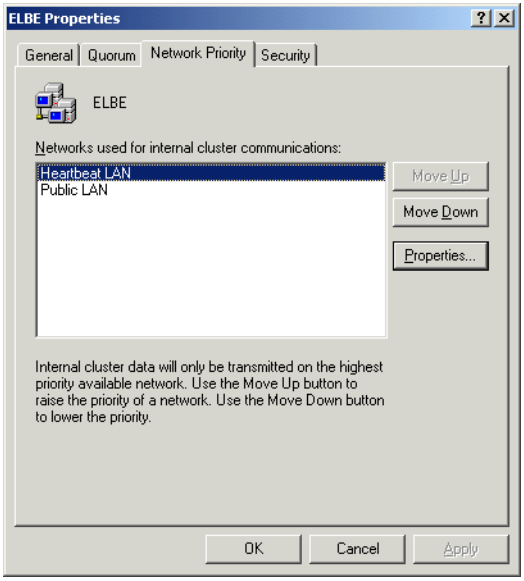

Figure 7-247 Cluster network priority properties

7.6 IBM iSCSI Server Adapter in SLES9

The scenario shows the setup of an IBM x445 equipped with two IBM iSCSI Server Adapters connected via switches to a DS300. SLES9 SP2 is used in this setup.

7.6.1 Planning and considerations

In this setup two iSCSI Server Adapters are used. Each adapter connects to its own array on the DS300. At the time of writing this book, it was not possible to implement a multipathing over two iSCSI Server Adapter.

7.6.2 SAN and network setup

Figure 7-248 shows the SAN and network setup.

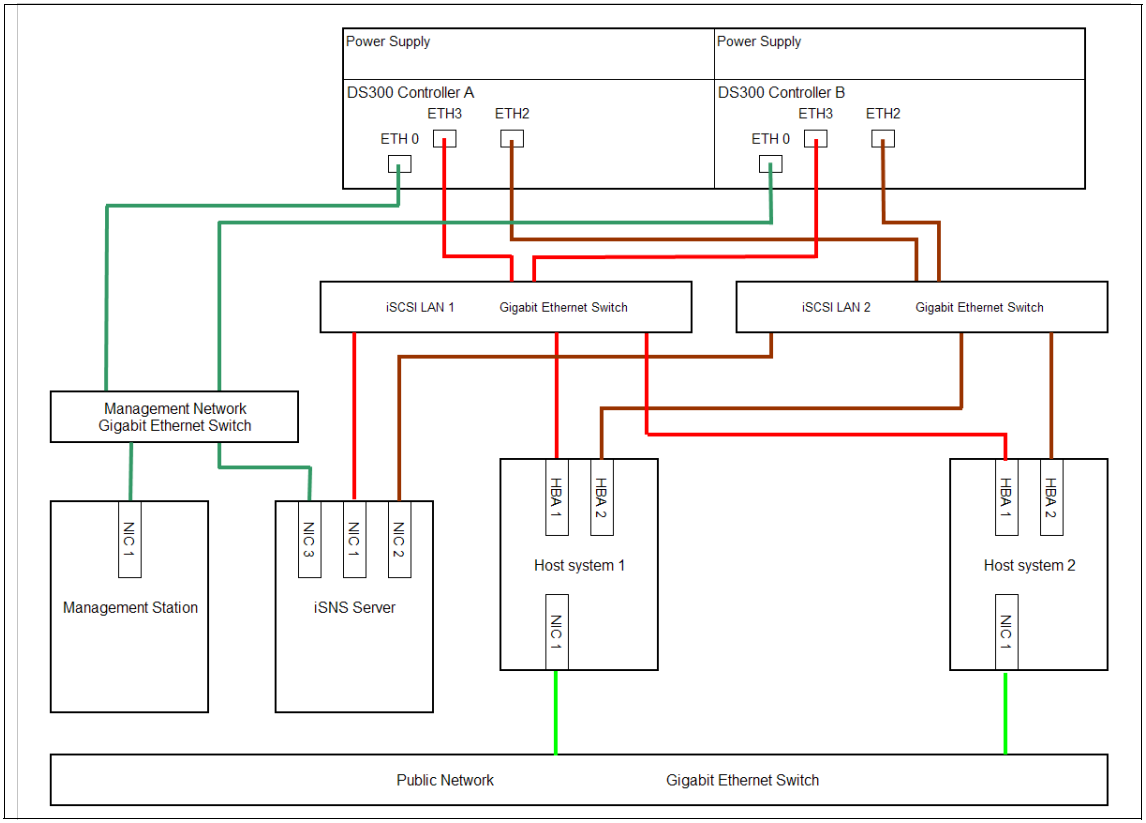

Figure 7-248 iSCSI network configuration with SLES9

7.6.3 IBM iSCSI Server Adapter driver installation

Before using the IBM iSCSI Server Adapter in SLES9, download the latest driver from the DS300 support page:

http://www.ibm.com/servers/storage/support/disk/ds300/

Note: For this setup we used a beta version of the qla4xxx driver as the initial version of Kernel 2.6 drivers for the IBM iSCSI Server Adapter were not available at the time of writing this book. The version number of the released qla4xxx driver for Kernel 2.6 will differ.

Follow these steps to install the driver:

1. Download the current driver from the DS300 support page.
2. Extract the files from download file using **`tar xzf <DRIVERFILE>`**
3. Change to the directory containing the source files. The directory name is the version number of the driver.
4. Build the driver using **`./extras/build.sh`**.
5. Install the driver using **`./extras/buil.sd install`**.
6. To ensure that the driver works reboot the server and verify the driver version using **`cat /proc/scsi/qla4xxx/<SCSI-HOST-NUMBER>`**.

Figure 7-249 on page 381 shows steps 2 - 5 in our test setup using SLES9 SP2.

```
danube:~/qla4xxx # tar xzf qla4xxx-5.00.03b19test3.tar.gz
danube:~/qla4xxx # cd 5.00.03b19test3/
danube:~/qla4xxx/5.00.03b19test3 # ./extras/build.sh
qla4xxx -- Building the qla4xxx driver...
make: Entering directory `/usr/src/linux-2.6.5-7.191'
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_os.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_init.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_mbx.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_iocb.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_isr.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_isns.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_nvram.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_dbg.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_cfg.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_cfgln.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_fo.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_foio.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_foioctl.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_foln.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_xioct.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_inioct.o
  CC [M]  /root/qla4xxx/5.00.03b19test3/ql4_32ioctl.o
  LD [M]  /root/qla4xxx/5.00.03b19test3/qla4xxx.o
  Building modules, stage 2.
  MODPOST
  CC      /root/qla4xxx/5.00.03b19test3/qla4xxx.mod.o
  LD [M]  /root/qla4xxx/5.00.03b19test3/qla4xxx.ko
make: Leaving directory `/usr/src/linux-2.6.5-7.191'
danube:~/qla4xxx/5.00.03b19test3 # ./extras/build.sh install
qla4xxx -- Building the qla4xxx driver...
make: Entering directory `/usr/src/linux-2.6.5-7.191'
  Building modules, stage 2.
  MODPOST
make: Leaving directory `/usr/src/linux-2.6.5-7.191'
qla4xxx -- Installing the qla4xxx modules to
/lib/modules/2.6.5-7.191-bigsmp/kernel/drivers/scsi/qla4xxx/...
danube:~/qla4xxx/5.00.03b19test3 #
```

Figure 7-249 Driver installation for IBM iSCSI Server Adapter in SLES9 SP2

The driver is now ready for use. To configure LUNs it is necessary to use the iSCSI SANsurfer management software.

7.6.4 iSCSI SANsurfer management software installation

To install the SANsurfer management software for Linux, follow these steps:

1. Open a shell in your graphical environment. Extract the installation file and start the installation executable as shown in Figure 7-250.

```
danube:~/SANsurfer # tar xzvf
iSCSI_SANsurfer_4_01_00_linux_x86_package.tgz.gz
Readme_SANsurfer_iSCSI_HBA_Mgr.txt
ReleaseNotes_SANsurfer_iSCSI_HBA_Mgr.pdf
iSCSI_SANsurfer_4_01_00_linux_x86.bin
danube:~/SANsurfer # ./iSCSI_SANsurfer_4_01_00_linux_x86.bin
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...

Launching installer...
```

Figure 7-250 SANsurfer Linux installation step 1

2. The SANsurfer graphical installer appears. The introduction screen gives you some advice about the installation. Click **Next** to proceed.

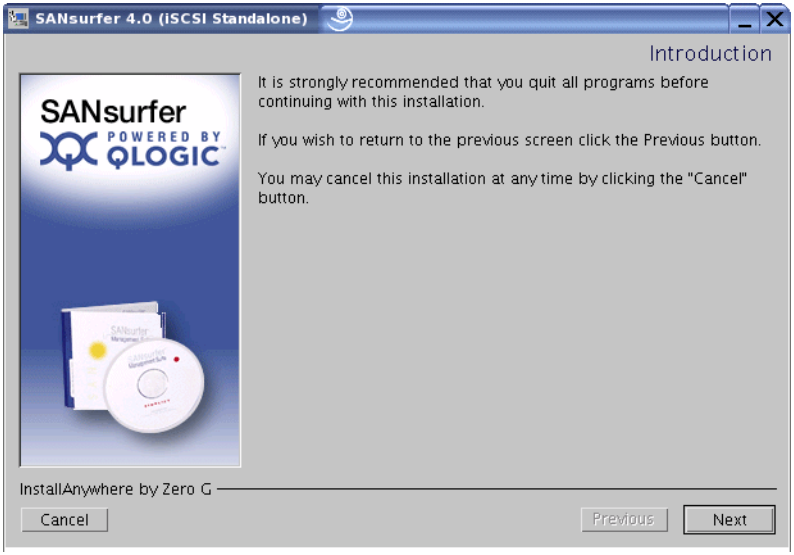

Figure 7-251 SANsurfer Linux installation step 2

3. Read the important information about the management application and then click **Next**.

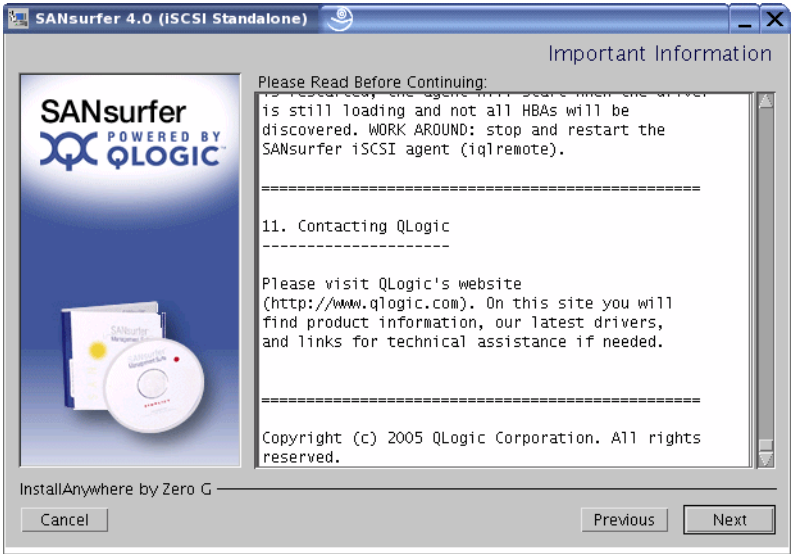

Figure 7-252 SANsurfer Linux installation step 3

4. Select the options you want to install. The installer allows to install the GUI and the agent alone or together. Depending on your needs select the appropriate option. The example below installs the GUI and the agent together on the same system. Click the **Next** button.

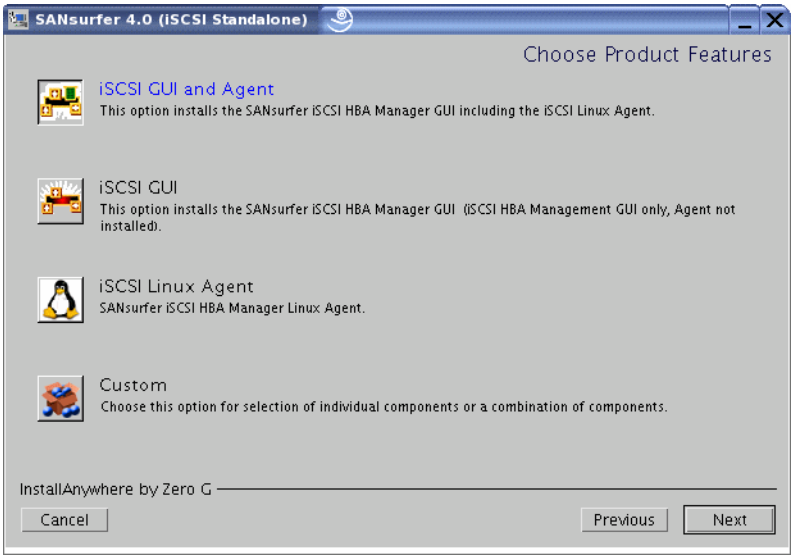

Figure 7-253 SANsurfer Linux installation step 4

5. You can specify a installation destination. Click **Next** to continue.

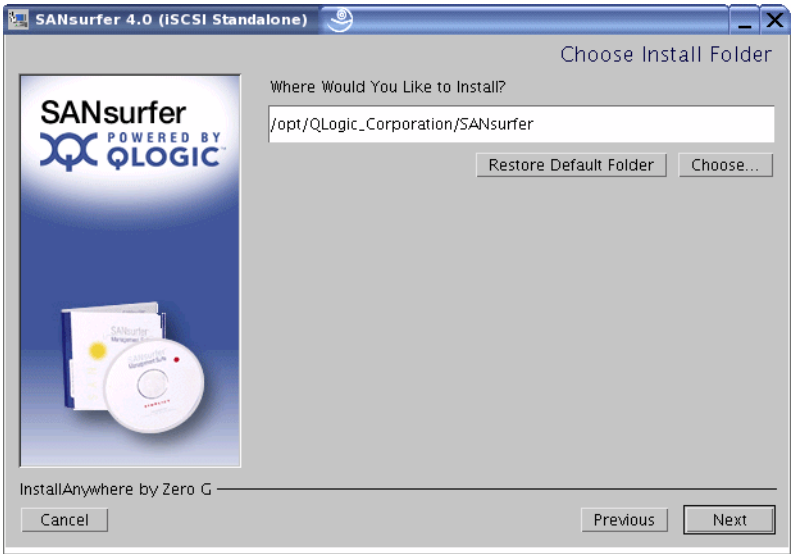

Figure 7-254 SANsurfer Linux installation step 5

6. Click **Install** in the pre-installation summary to start the installation.

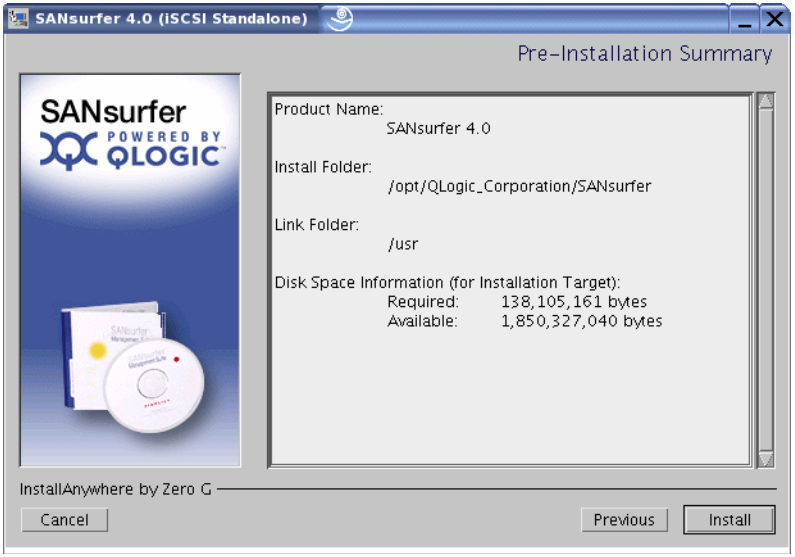

Figure 7-255 SANsurfer Linux installation step 6

7. Click **Done** to close the installation application.

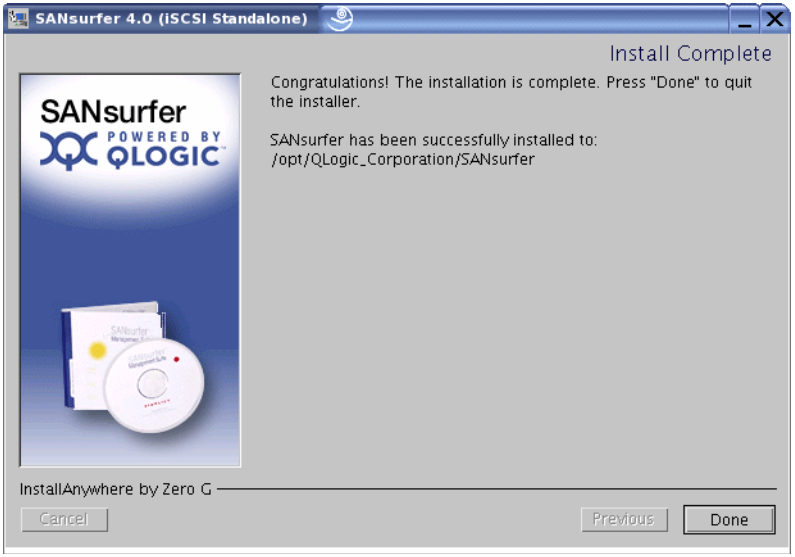

Figure 7-256 SANsurfer Linux installation step 7

7.6.5 Configuration of the IBM iSCSI Server Adapter Card

The IBM iSCSI Server Adapter Card can be configured in its BIOS during system POST (by pressing **CTRL+Q**) or with the iSCSI SANsurfer Management Software.

The configuration is almost the same as described in 7.2.3, “Setup of the QLogic iSCSI Expansion Card for IBM eServer BladeCenter” on page 291. The only difference is that the IBM iSCSI Server Adapter Card is a single port card. So in the SANsurfer iSCSI HBA manager application appears a single port for each installed HBA.

Figure 7-257 on page 386 shows SANsurfer on a system with two installed IBM iSCSI Server Adapter cards.

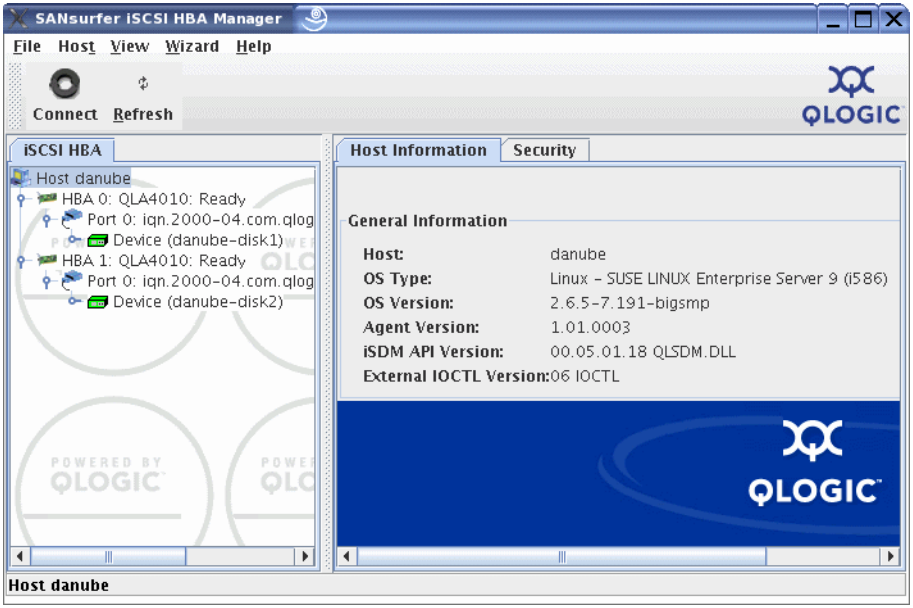

Figure 7-257 SANsurfer application in Linux

Note: When using two cards in a server, configure the cards to access different logical drives to do a load sharing. Currently the driver does not support multipathing to a logical drive using multiple HBAs. We tested device mapper and the multipath-tools for multipathing, but experienced problems as the readsector0 checker reported failed paths online again after some seconds.

7.7 Microsoft iSCSI software initiator 1.06

In this section the installation and configuration of the Microsoft iSCSI software initiator 1.06 on Windows 2003 Enterprise Edition Service Pack 1 is described. This section is included as the Microsoft iSCSI software initiator 2.0 was only around for a short time at the time of writing and there might be environments where the Microsoft iSCSI software initiator 1.06 is already in use. Refer to chapter 7.8, “Microsoft iSCSI software initiator 2.0 in Windows 2003” on page 398 for the newer Microsoft iSCSI software initiator.

7.7.1 Planning and considerations

The Microsoft iSCSI software initiator 1.06 includes MPIO support but does not contain the required Microsoft MPIO binaries. The Microsoft iSCSI software initiator 1.06 itself can only be used for single path configurations to an iSCSI target. Microsoft iSCSI software initiator 2.0 has support for both MPIO and iSCSI multiple connections per session for multipathing.

7.7.2 Installation of Microsoft iSCSI software initiator 1.06

Follow these steps to install the Microsoft iSCSI software initiator 1.06:

1. Download and execute the installation package 1.06-initiator-x86fre.msi. On the welcome screen click **Next** to continue.

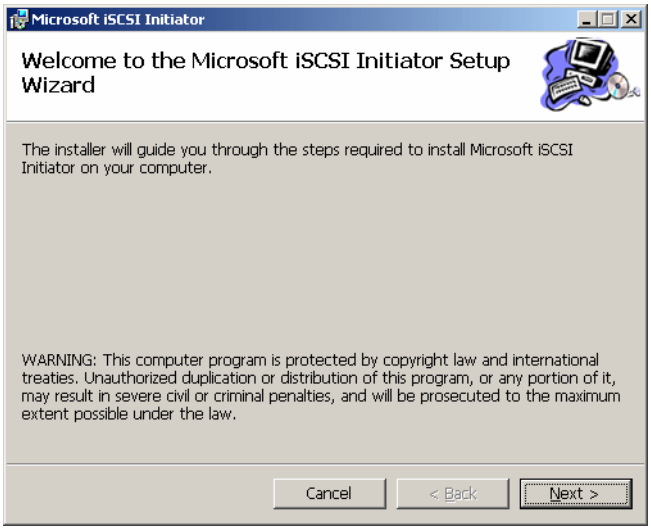

Figure 7-258 Microsoft iSCSI software initiator 1.06 installation step 1

2. Select the installation folder and choose **Everyone** if appropriate for you before clicking **Next**.

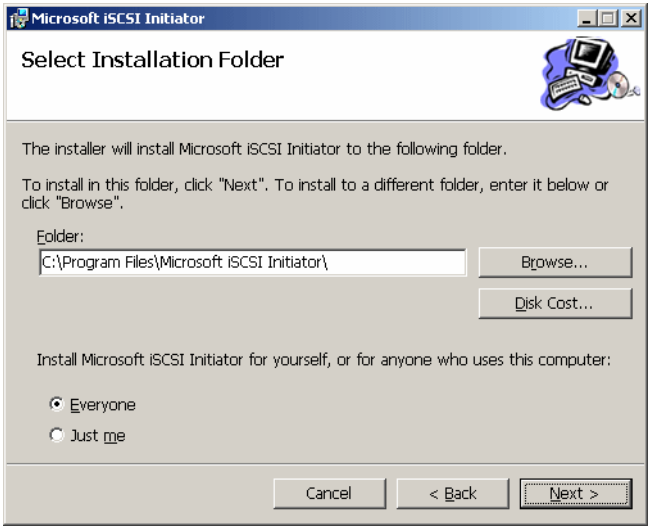

Figure 7-259 Microsoft iSCSI software initiator 1.06 installation step 2

3. Click **Next** to start the installation.

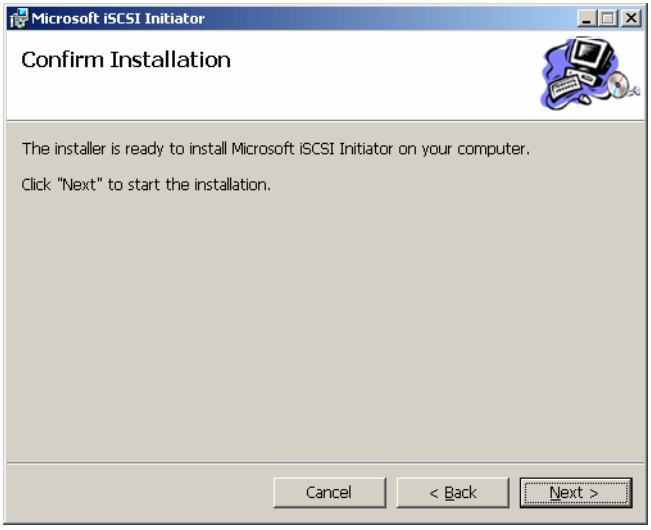

Figure 7-260 Microsoft iSCSI software initiator 1.06 installation step 3

4. Click **I Agree** and **Next** to continue.

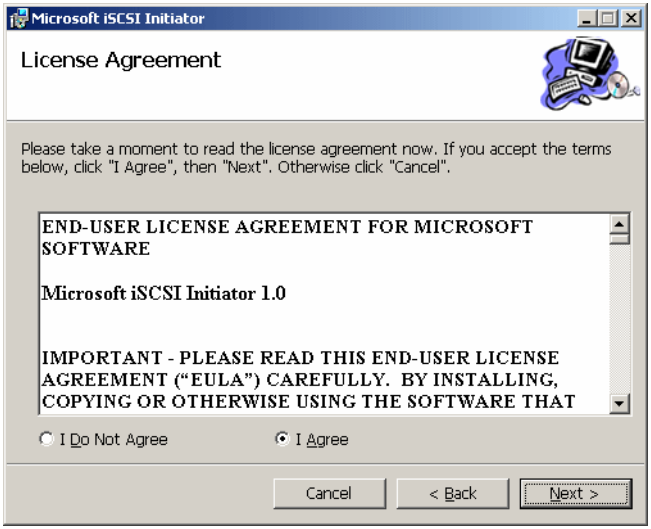

Figure 7-261 Microsoft iSCSI software initiator 1.06 installation step 4

5. Ensure that **Install Complete iSCSI Initiator** is selected and click **OK** to continue.

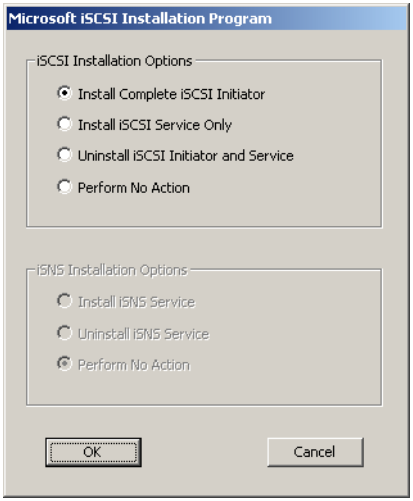

Figure 7-262 Microsoft iSCSI software initiator 1.06 installation step 5

6. Click **Agree** to continue.

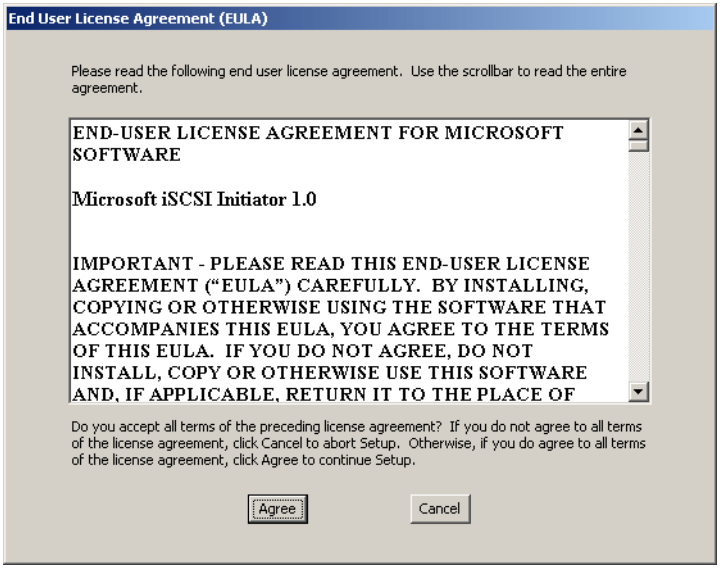

Figure 7-263 Microsoft iSCSI software initiator 1.06 installation step 6

7. A message that the iSCSI initiator was installed successfully and has been started appears. Click **OK** to close the message window.

Figure 7-264 Microsoft iSCSI software initiator 1.06 installation step 7

8. There are some hints displayed in the following window, that help to find the problems if the installation did not finish successfully in step 7. Click **Next**.

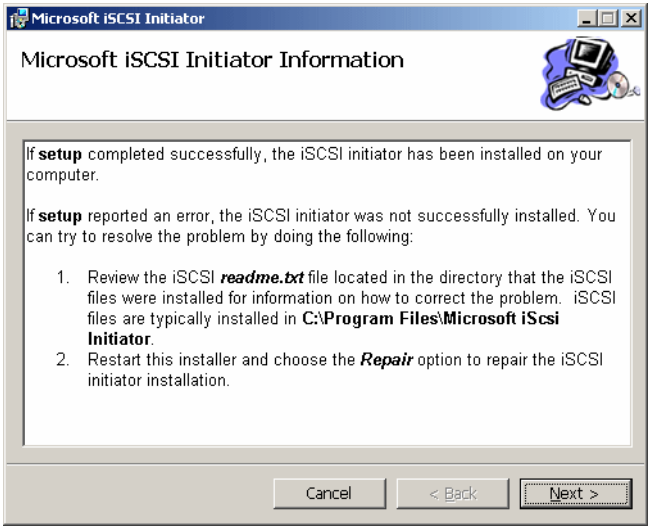

Figure 7-265 Microsoft iSCSI software initiator 1.06 installation step 8

9. Click **Close** to exit the installation program.

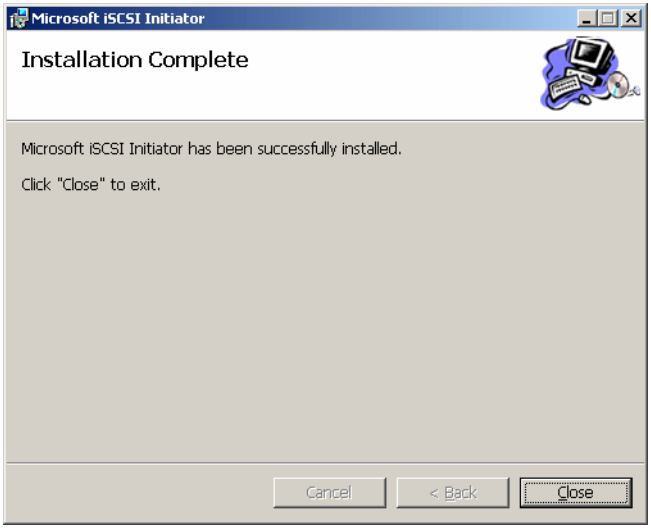

Figure 7-266 Microsoft iSCSI software initiator 1.06 installation step 9

7.7.3 Configuration of Microsoft iSCSI software initiator 1.06

In this section we describe the configuration of the Microsoft iSCSI software initiator 1.06. Before using Microsoft iSCSI software initiator in a production

environment, it is recommend to read the *Microsoft iSCSI Software Initiator Users Guide* that can be downloaded as uguide.doc with the software initiator.

Open the iSCSI Initiator properties dialog by starting **iSCSI Initiator** in the control panel. Execute the following steps to do the configuration:

1. Click the **Initiator Settings** tab to view the initiator node name. The logical drives on the DS300 that are used must allow access to the logical drives from this initiator.

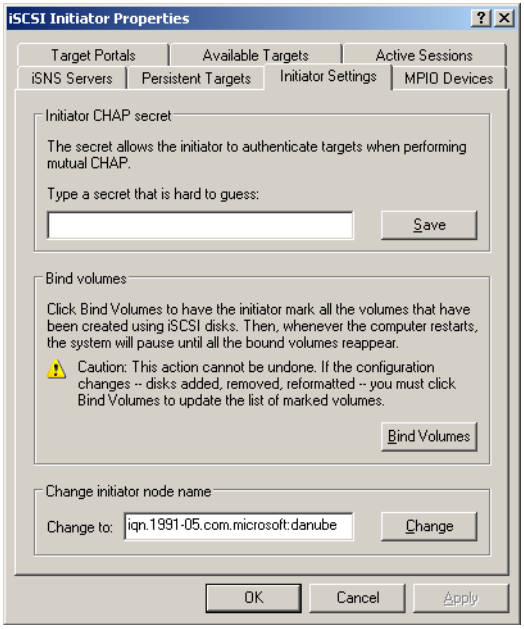

Figure 7-267 Microsoft iSCSI software initiator 1.06 configuration step 1

2. Device discovery can be done whether by entering target portals in the **Target Portals** tab or by entering iSNS server IPs in the **iSNS Servers** tab. We show how to configure device discovery by providing IPs of target portals. Click the **Target Portals** tab.

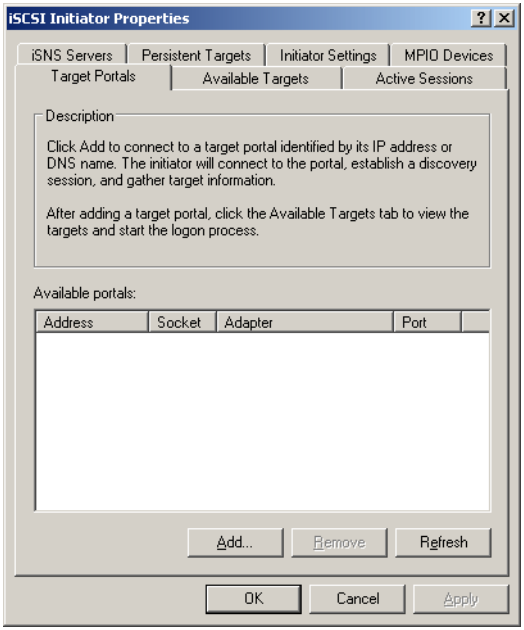

Figure 7-268 Microsoft iSCSI software initiator 1.06 configuration step 2

3. Click **Add...** and enter the IP address of the target portal.

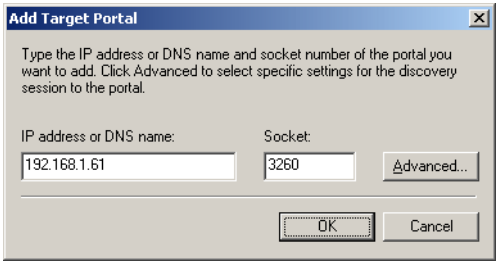

Figure 7-269 Microsoft iSCSI software initiator 1.06 configuration step 3

4. The target portal now appears in the list of available portals.

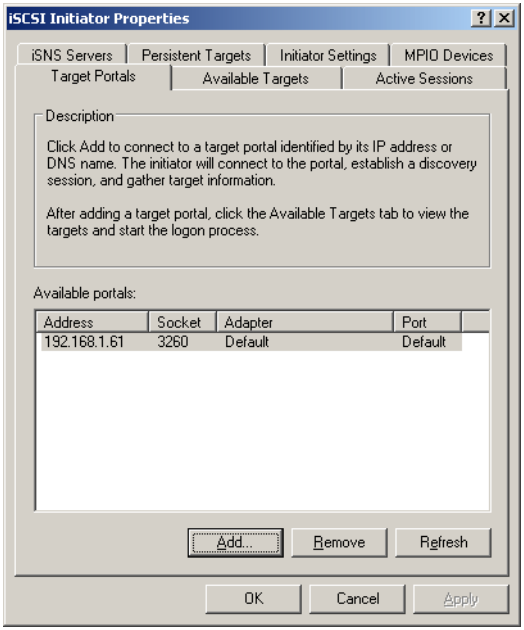

Figure 7-270 Microsoft iSCSI software initiator 1.06 configuration step 4

5. Click the **Available Targets** tab. Select the first target you want to use and click Log **On...**.

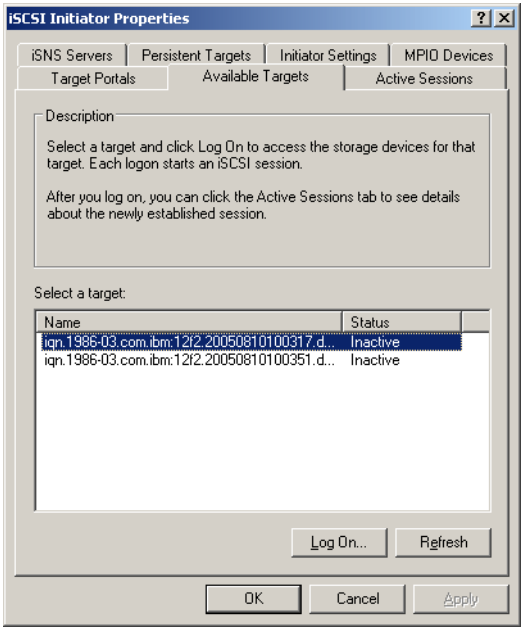

Figure 7-271 Microsoft iSCSI software initiator 1.06 configuration step 5

6. Enable the checkbox **Automatically restore this connection when the system boots** and click **Advanced...**.

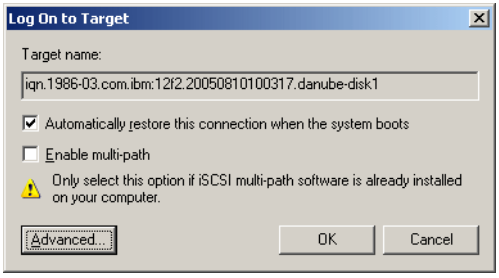

Figure 7-272 Microsoft iSCSI software initiator 1.06 configuration step 6

7. Select **Microsoft iSCSI Initiator** in the **Local adapter** pull-down. Provide the logon information if you are using CHAP authentication. Click **OK** to the **Advanced Settings** dialog.

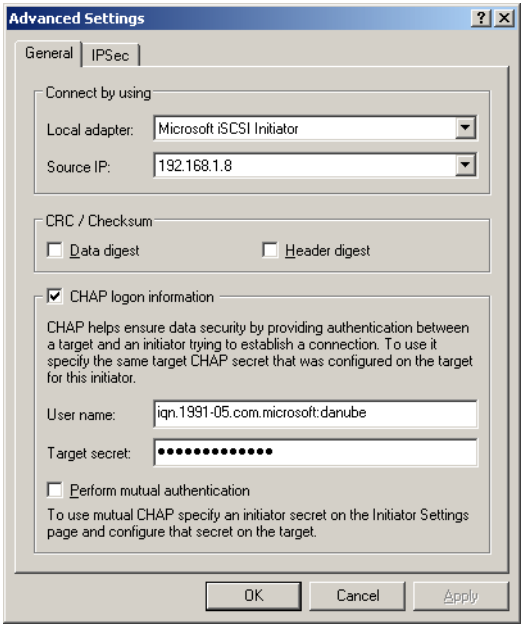

Figure 7-273 Microsoft iSCSI software initiator 1.06 configuration step 7

8. Click **OK** in the **Log On to Target** dialog. Redo steps 5-7 for each target you want to connect to. The status of the targets in the **Available Targets** tab of the **iSCSI Initiator Properties** dialog is now **Connected**.

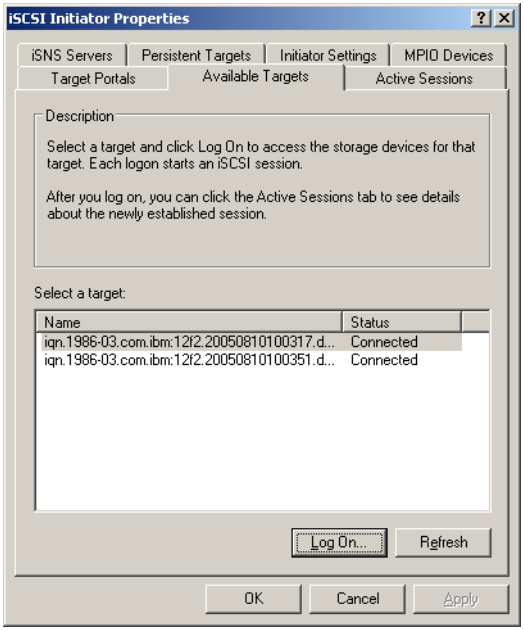

Figure 7-274 Microsoft iSCSI software initiator 1.06 configuration step 8

9. The active sessions to the connected targets are shown in the **Active Sessions** tab.

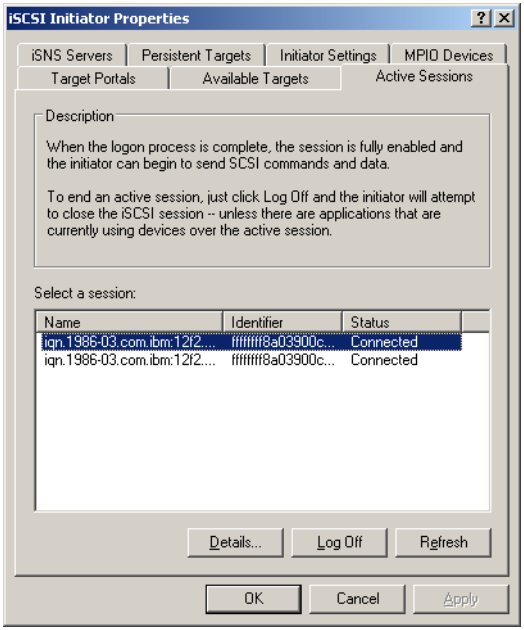

Figure 7-275 Microsoft iSCSI software initiator 1.06 configuration step 9

7.8 Microsoft iSCSI software initiator 2.0 in Windows 2003

In this section, the installation and configuration of the Microsoft iSCSI software initiator 2.0 on Windows 2003 Enterprise Edition Service Pack 1 is described.

7.8.1 Planning and considerations

In this example, MPIO is used for multipath. Note that the Microsoft iSCSI software initiator 2.0 allows to use MPIO or iSCSI multiple connections per session for multipathing.

7.8.2 SAN and network setup

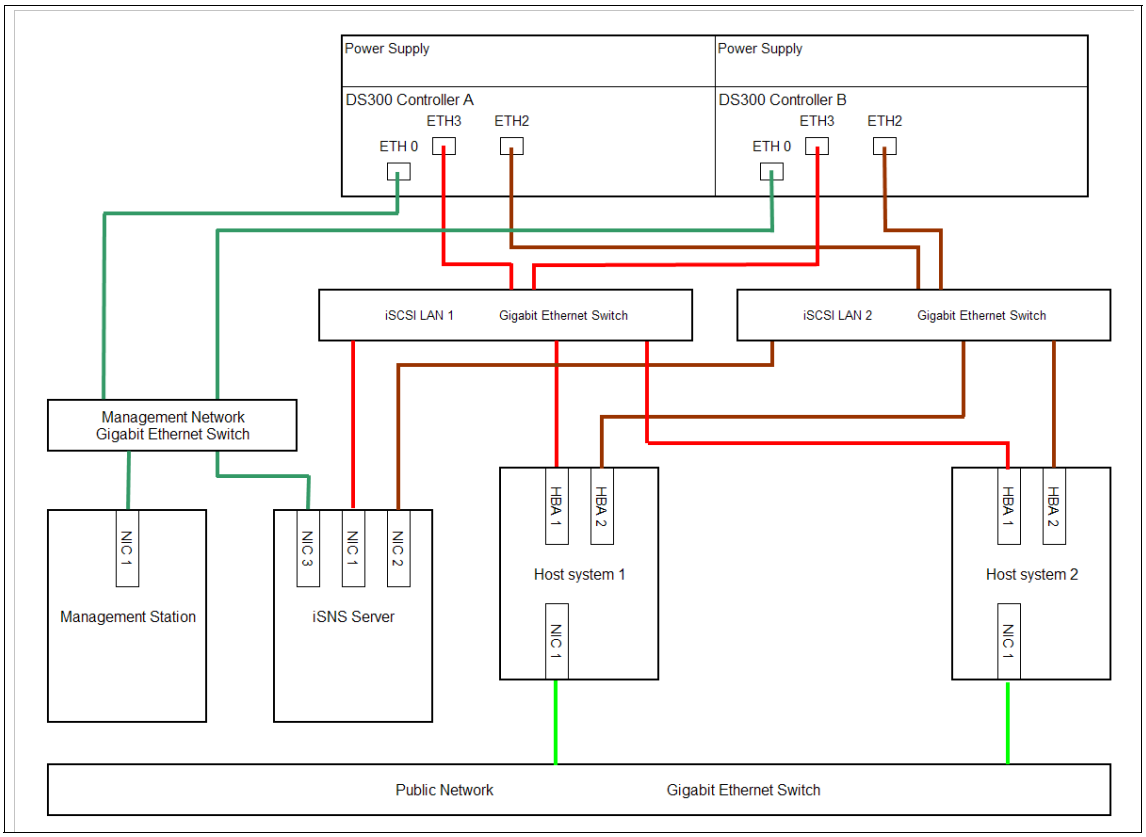

Figure 7-276 Network configuration

7.8.3 Installation of Microsoft iSCSI software initiator 2.0

Follow these steps to install the Microsoft iSCSI software initiator 2.0:

1. Download and execute the installation binary iSCSI-2.0-x86fre.exe. The software update installation wizard opens. Click **Next** to continue.

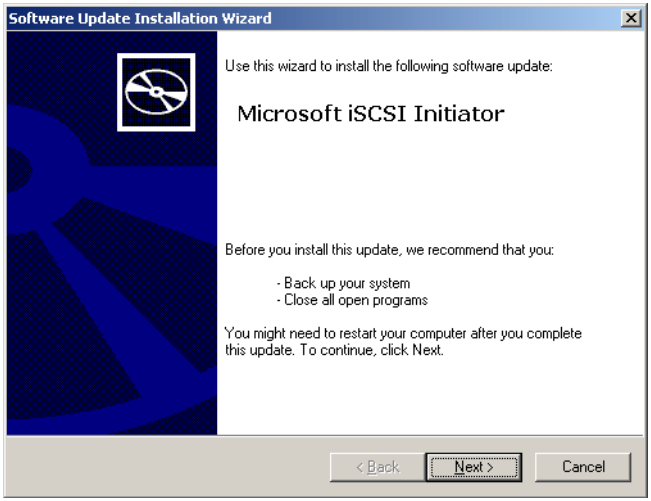

Figure 7-277 Microsoft iSCSI software initiator 2.0 installation Step 1

2. Activate the **Microsoft MPIO Multipathing Support for iSCSI** checkbox and click **Next**.

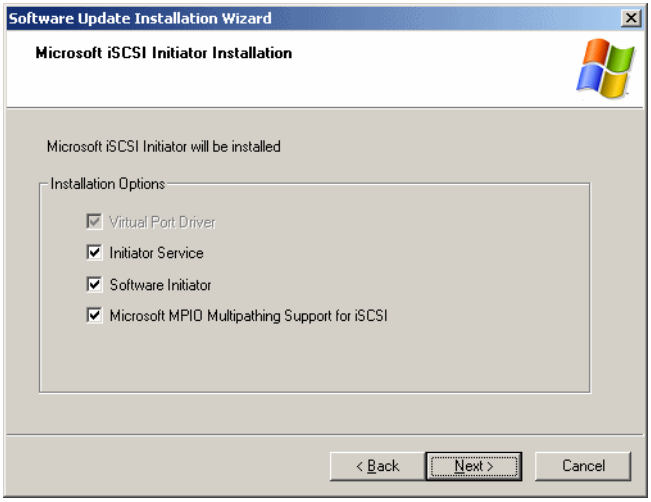

Figure 7-278 Microsoft iSCSI software initiator 2.0 installation step 2

3. Select the radio button **I agree** and click **Next**.

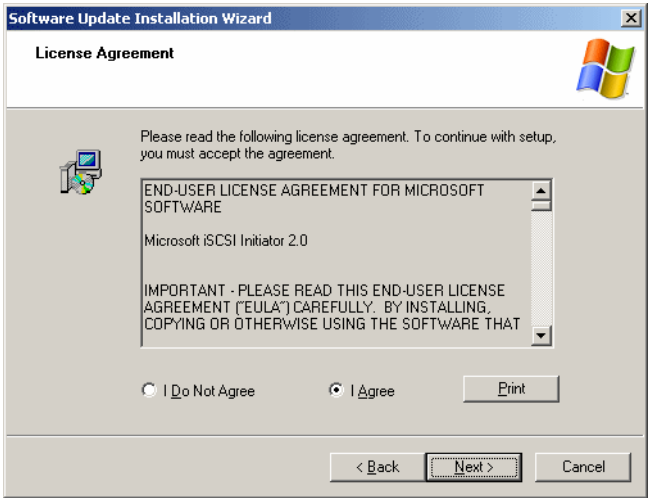

Figure 7-279 Microsoft iSCSI software initiator 2.0 installation step 3

4. Click **Finish** to close the installation wizard. The computer will be restarted.

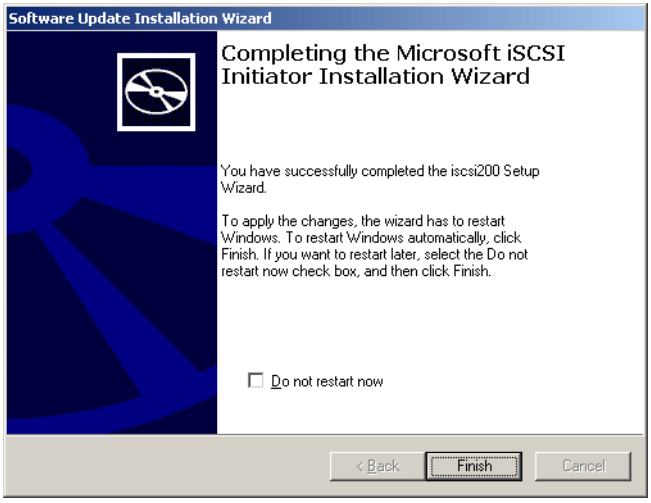

Figure 7-280 Microsoft iSCSI software initiator 2.0 installation step 4

7.8.4 Configuration of Microsoft iSCSI software initiator 2.0

In this section we describe the configuration of the Microsoft iSCSI software initiator 2.0. Before using Microsoft iSCSI software initiator in a production environment, it is recommend to read the *Microsoft iSCSI Software Initiator Users Guide* that can be downloaded as uguide.doc with the software initiator.

Open the iSCSI Initiator properties dialog by starting **iSCSI Initiator** in the control panel. Execute the following steps to do the configuration:

1. The initiator node name is shown in the **General** tab. The logical drives on the DS300 that are used must allow access to the logical drives from this initiator.

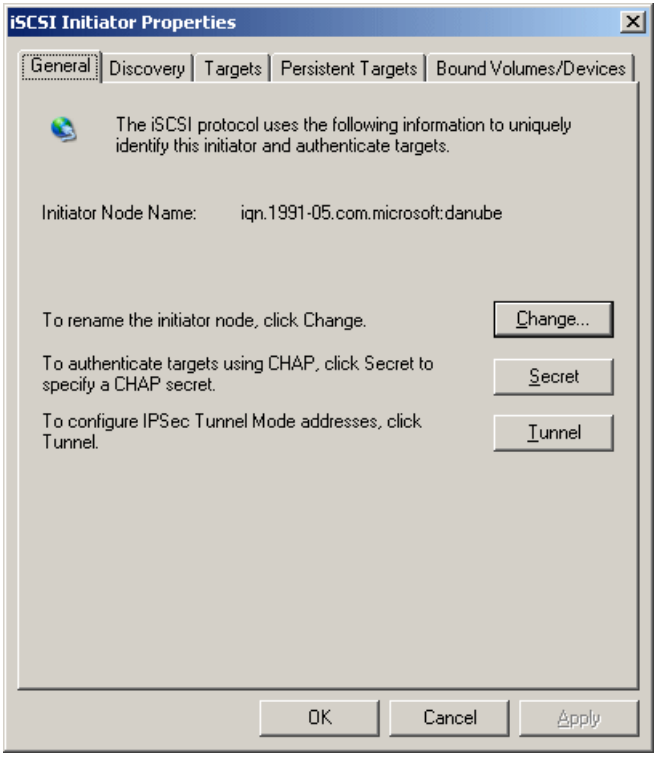

Figure 7-281 Microsoft iSCSI software initiator 2.0 configuration step 1

2. In the **Discovery** tab it is possible to search for targets by adding IP addresses of target portals or by specifying an iSNS server. Click **Add** in the appropriate area to add a target portal or an iSNS server. This example shows the configuration by using an iSNS server.

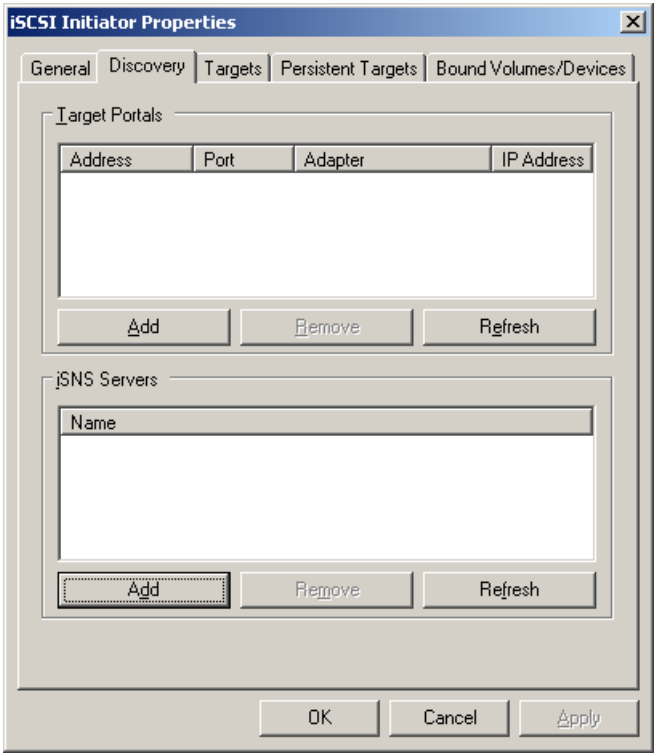

Figure 7-282 Microsoft iSCSI software initiator 2.0 configuration step 2

3. Enter the IP address of the iSNS server and click **OK**.

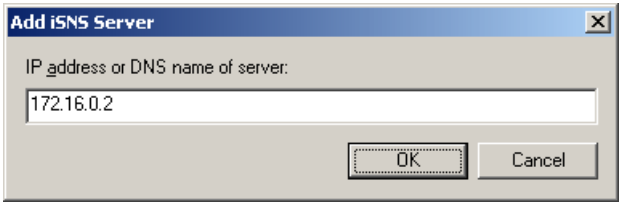

Figure 7-283 Microsoft iSCSI software initiator 2.0 configuration step 3

4. Switch to the **Targets** tab to display all targets known by the iSNS server. Select the target you want to connect to and click **Log On...**.

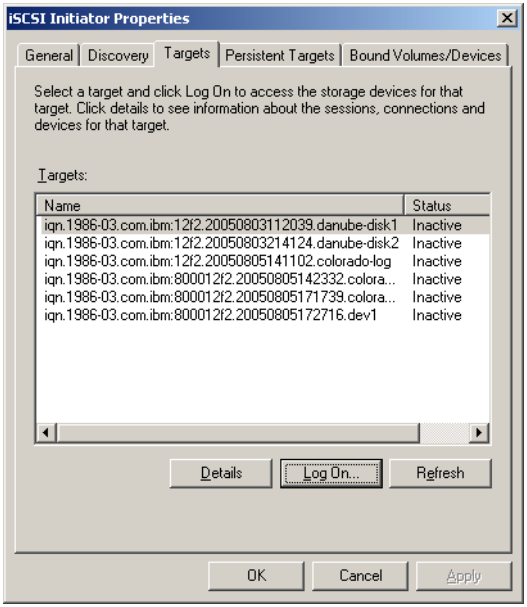

Figure 7-284 Microsoft iSCSI software initiator 2.0 configuration step 4

5. Activate **Automatically restore this connection when the system boots** and **Enable multi-path**. Click **Advanced** to configure the details.

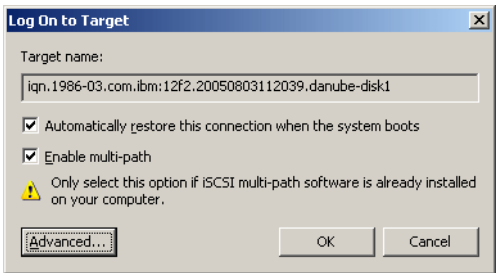

Figure 7-285 Microsoft iSCSI software initiator 2.0 configuration step 5

6. In the **General** tab of the **Advanced Settings** dialog choose the appropriate values for the *first* path in the connect by using area. Ensure that the source IP and the target portal are both in the same subnet. In our example the source IP and the target portal are in iSCSI LAN 1 (see Table 7-1 on page 224). If you are using CHAP authentication, activate the checkbox for CHAP and enter username and target secret. Click **OK** to close the advanced settings dialog. Then click **OK** to close the log on to target dialog.

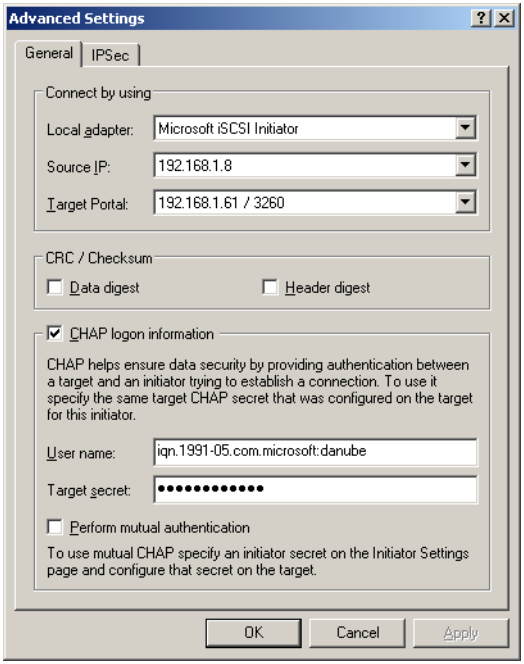

Figure 7-286 Microsoft iSCSI software initiator 2.0 configuration step 6

7. The **Targets** tab in the **iSCSI Initiator Properties** dialog indicates now a status of connected for the configured target. If you have multiple ports to the target, ensure that the connected target is highlighted and click **Log On...** again.

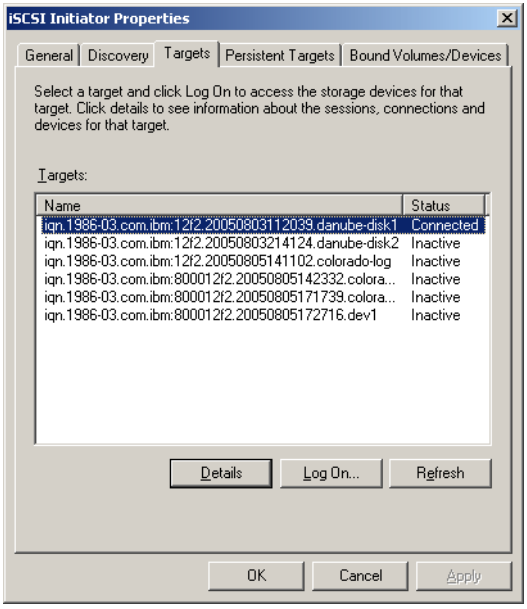

Figure 7-287 Microsoft iSCSI software initiator 2.0 configuration step 7

8. Activate **Automatically restore this connection when the system boots** and **Enable multi-path**. Click **Advanced** to configure the details.

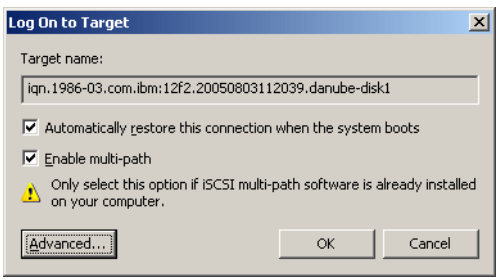

Figure 7-288 Microsoft iSCSI software initiator 2.0 configuration step 8

9. In the **General** tab of the **Advanced Settings** dialog choose the appropriate values for the *second* path in the connect by using area. Ensure that the source IP and the target portal are both in the same subnet. In our example the source IP and the target portal are in iSCSI LAN 2 (see Table 7-1 on page 224). If you are using CHAP authentication, activate the checkbox for CHAP and enter username and target secret. Click **OK** to close the advanced settings dialog. Then click **OK** to close the log on to target dialog.

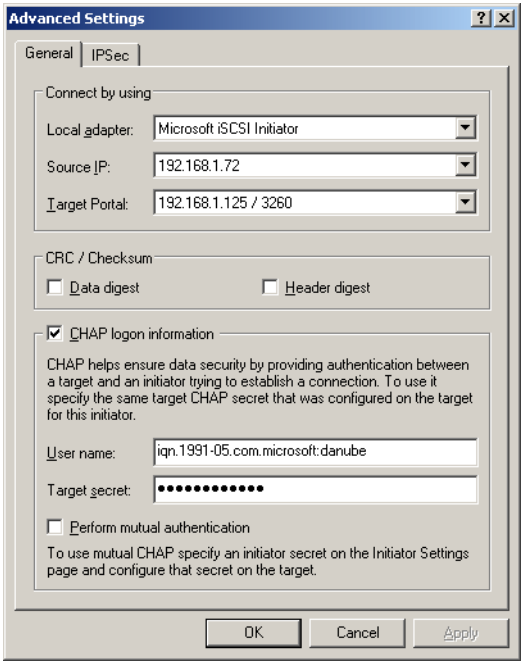

Figure 7-289 Microsoft iSCSI software initiator 2.0 configuration step 9

10. The status of the drive is still connected in the **Targets** tab of the **iSCSI Initiator Properties** dialog. Ensure that the drive is marked and click **Details**.

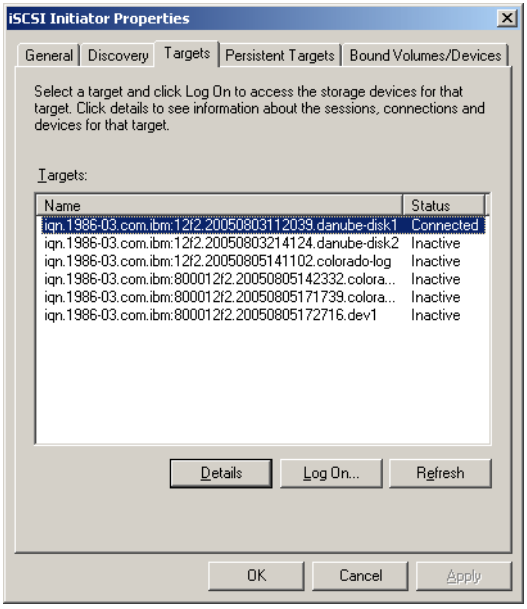

Figure 7-290 Microsoft iSCSI software initiator 2.0 configuration step 10

11. The two sessions that are connected to the target are shown. Click the **Devices** tab.

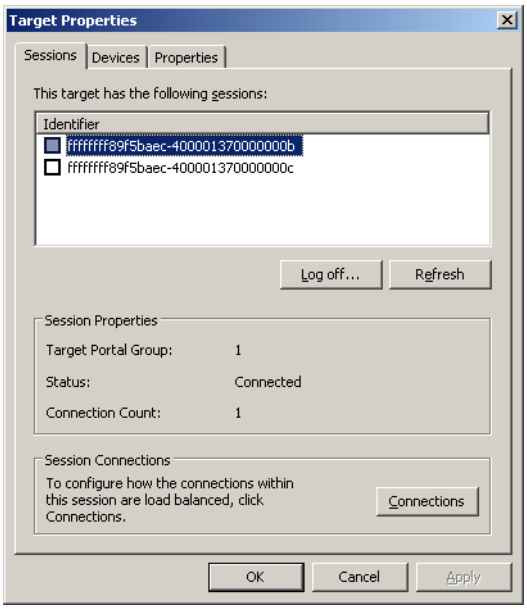

Figure 7-291 Microsoft iSCSI software initiator 2.0 configuration step 11

12. Click **Advanced** to open the device details dialog.

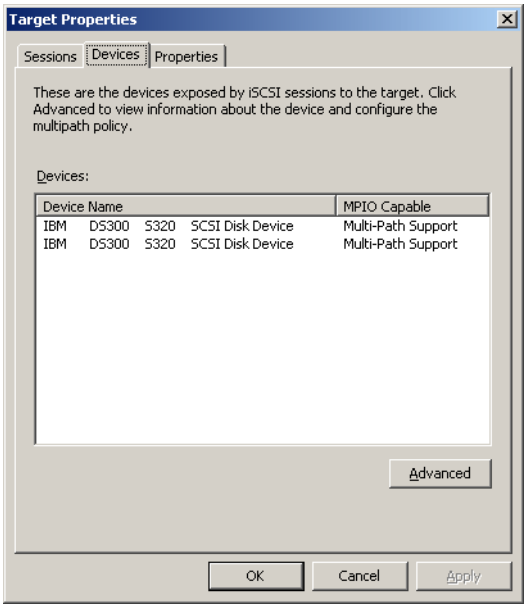

Figure 7-292 Microsoft iSCSI software initiator 2.0 configuration step 12

13. Click the **MPIO** tab.

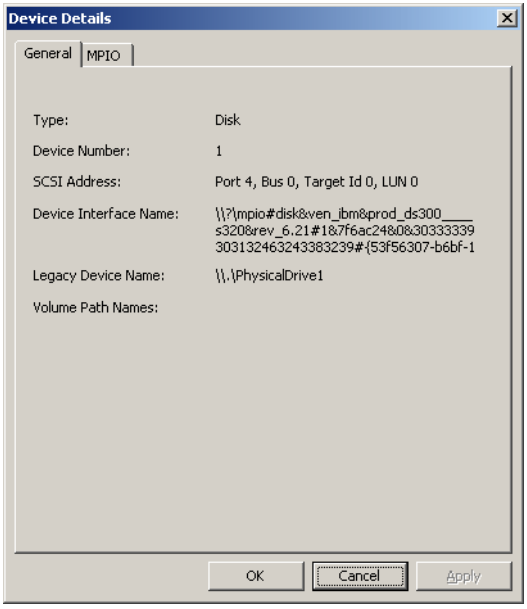

Figure 7-293 Microsoft iSCSI software initiator 2.0 configuration step 13

14. The load balancing policy and the paths to the device are shown. Highlight the active path and click **Details**.

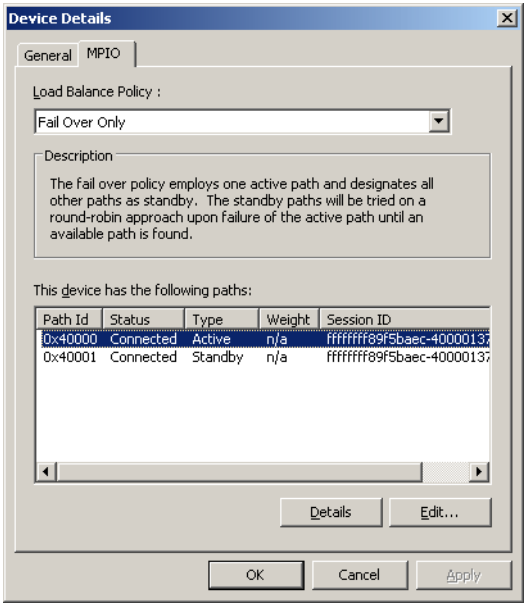

Figure 7-294 Microsoft iSCSI software initiator 2.0 configuration step 14

15. The MPIO path details are shown. Verify the source portal and the target portal of the connection. Click **OK** to close the **MPIO Path Details** dialog. Highlight the standby path, click **Details** again to verify now the settings for the standby path.

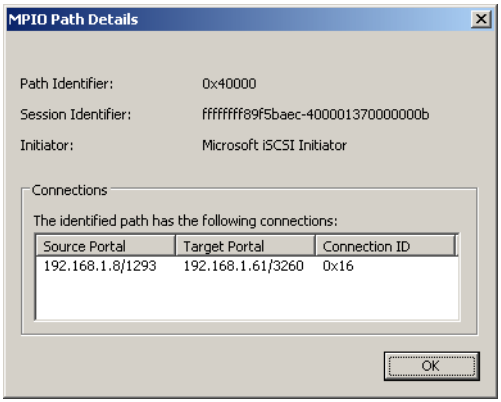

Figure 7-295 Microsoft iSCSI software initiator 2.0 configuration step 15

16. Close all open dialogs to return to the main **iSCSI Initiator Properties** dialog. Execute steps 4-16 again for other logical drives you want to connect to.

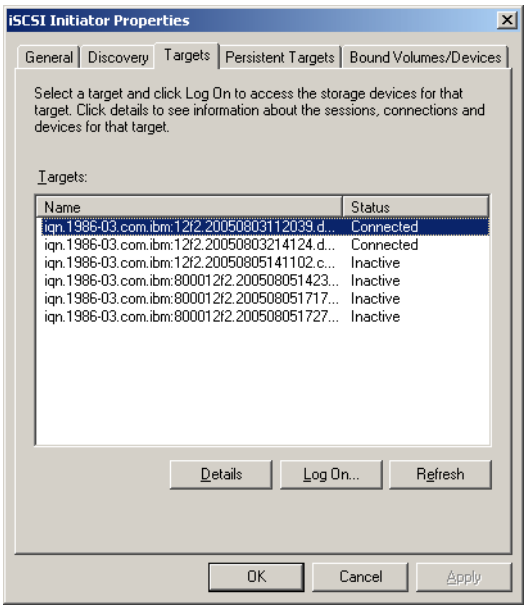

Figure 7-296 Microsoft iSCSI software initiator 2.0 configuration step 16

17. Click the **Persistent Targets** tab to verify that there is an entry for every MPIO path in the list.

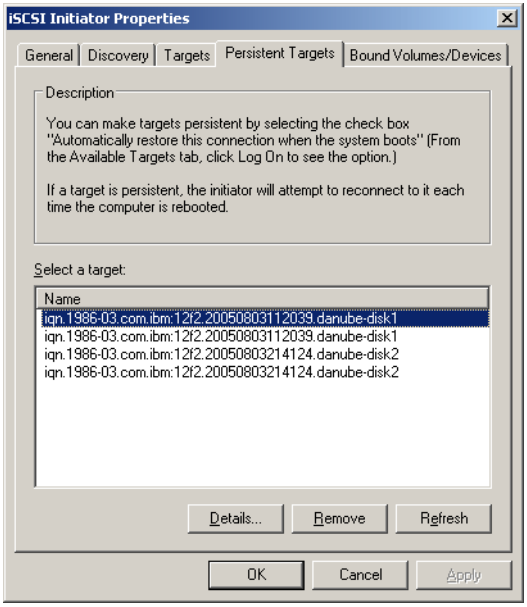

Figure 7-297 Microsoft iSCSI software initiator 2.0 configuration step 17

18. Open the computer management and click **Disk Management**.

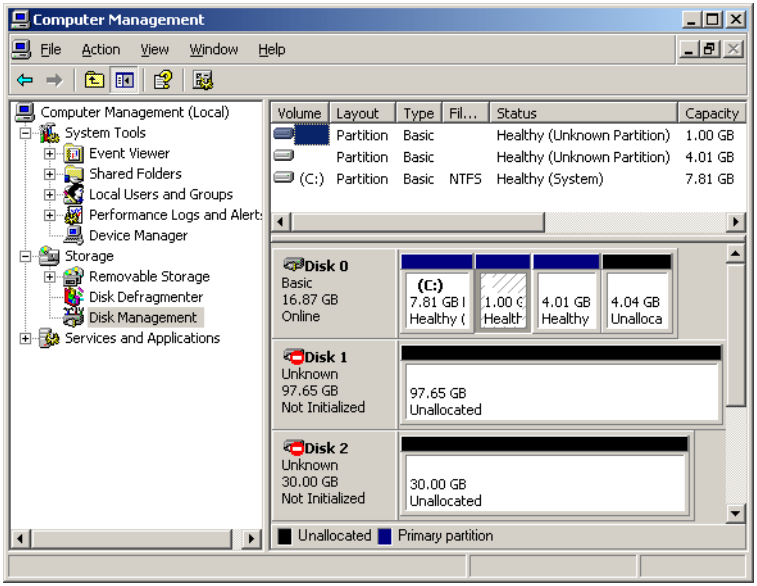

Figure 7-298 Microsoft iSCSI software initiator 2.0 configuration step 18

19. The **Initialize and Convert Disk** wizard starts up automatically. Click **Next** to continue.

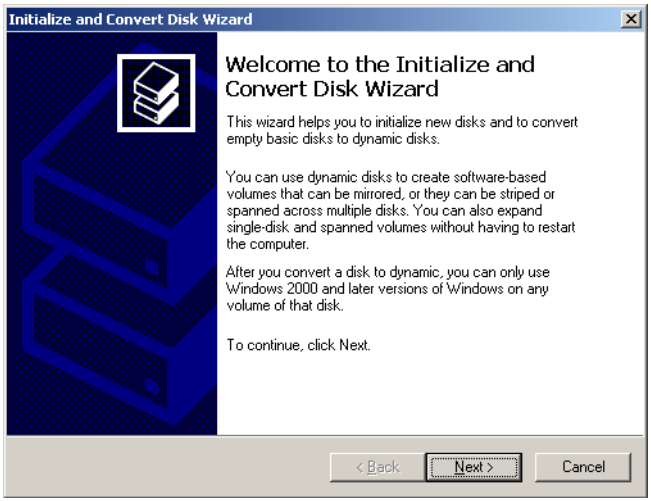

Figure 7-299 Microsoft iSCSI software initiator 2.0 configuration step 19

20. Ensure that the new disks are selected and click **Next**.

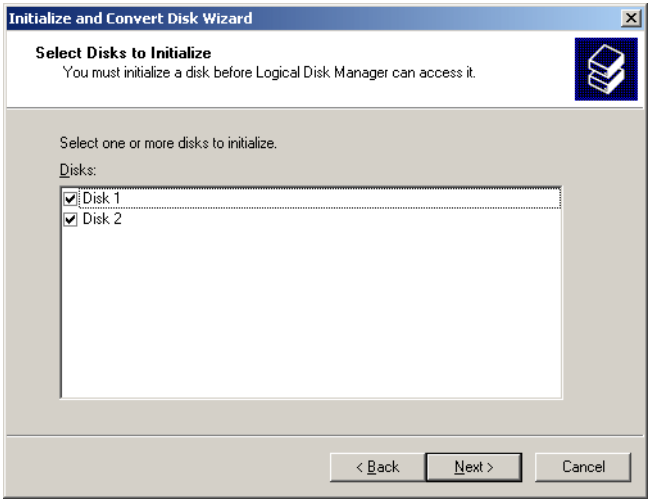

Figure 7-300 Microsoft iSCSI software initiator 2.0 configuration step 20

21. According to the users guide shipped with the Microsoft iSCSI software initiator (uguide.doc), dynamic disks are not supported for volumes made available through the Microsoft software initiator. Ensure to keep the disks unselected before you click **Next**.

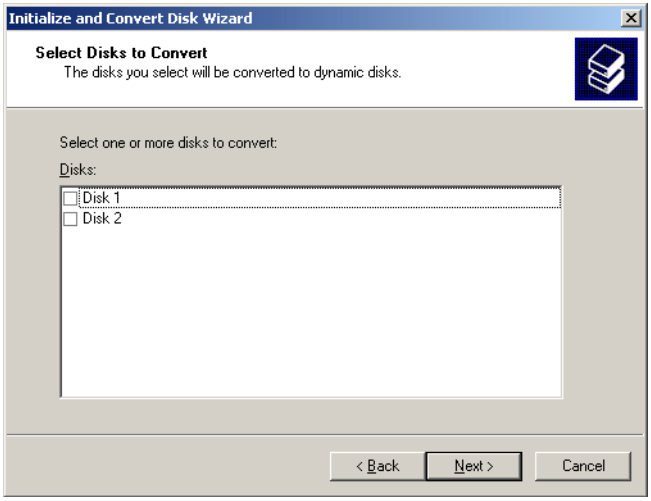

Figure 7-301 Microsoft iSCSI software initiator 2.0 configuration step 21

22. Click **Finish** to close the wizard.

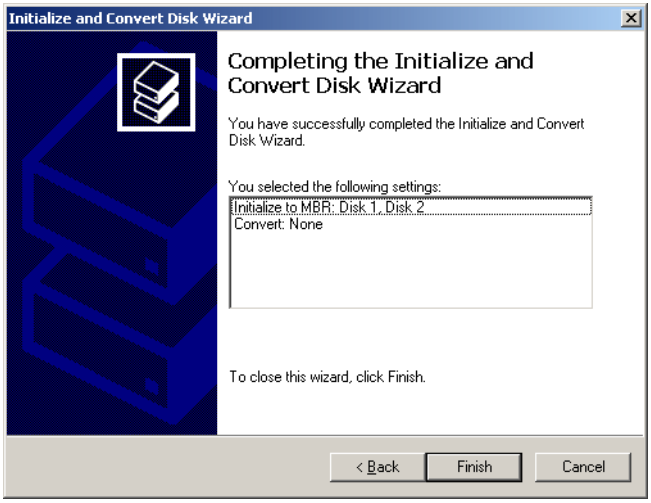

Figure 7-302 Microsoft iSCSI software initiator 2.0 configuration step 22

23. **Disk Management** shows now the new disks online.

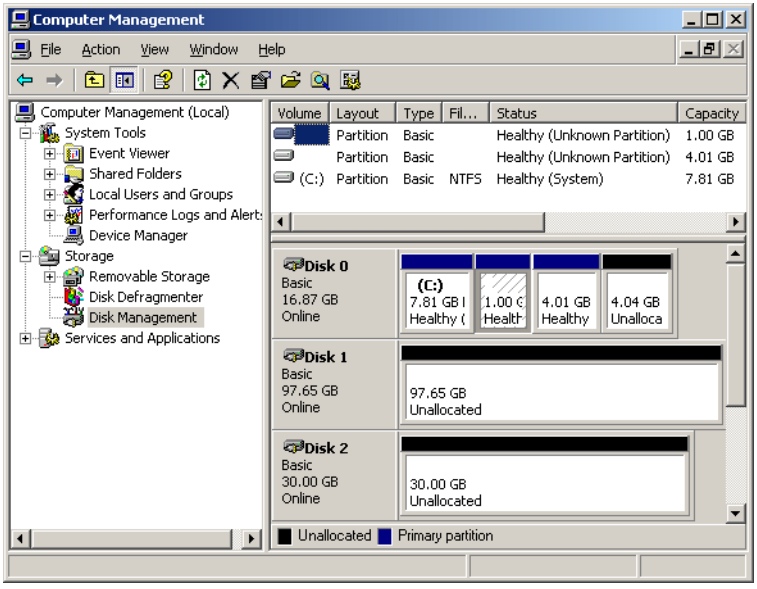

Figure 7-303 Microsoft iSCSI software initiator 2.0 configuration step 23

24. To partition and format the new drives, right-click the drive, click **New partition...** and follow the wizard. After partitioning and formatting the drives they appear as healthy drives.

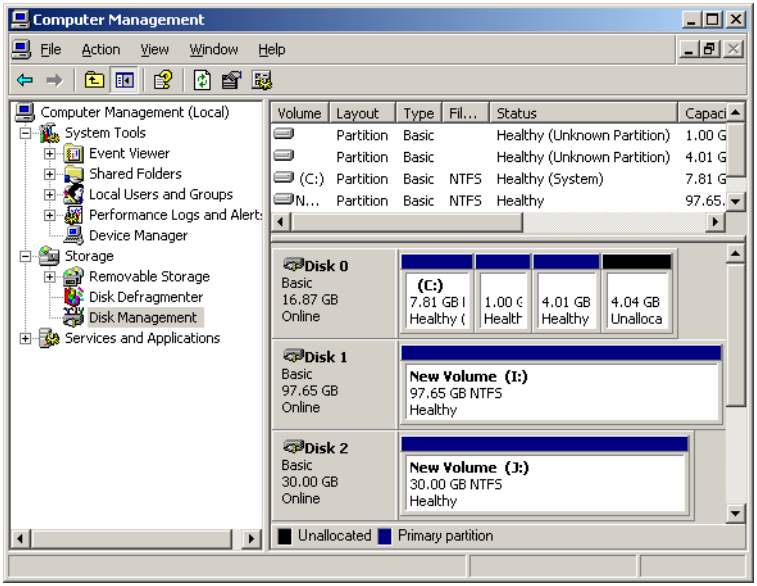

Figure 7-304 Microsoft iSCSI software initiator 2.0 configuration step 24

25. To ensure that the drives are available on startup they have to be persistently bound. This is done in the **Bound Volumes/Devices** tab in the iSCSI initiator properties dialog. Click **Add** to bind a volume.

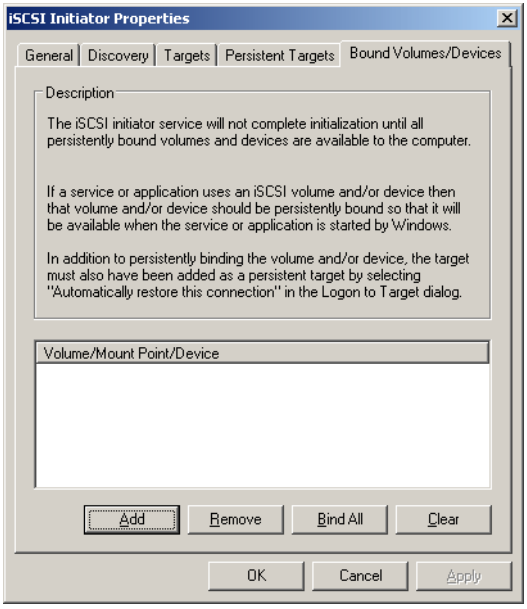

Figure 7-305 Microsoft iSCSI software initiator 2.0 configuration step 25

26. Enter the drive letter and click **OK**. Repeat these steps for every new logical drive.

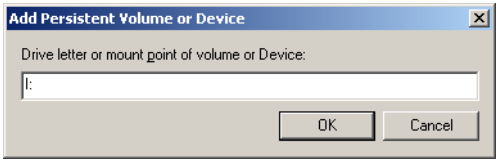

Figure 7-306 Microsoft iSCSI software initiator 2.0 configuration step 26

27. Click **OK** to close the iSCSI initiator properties dialog.

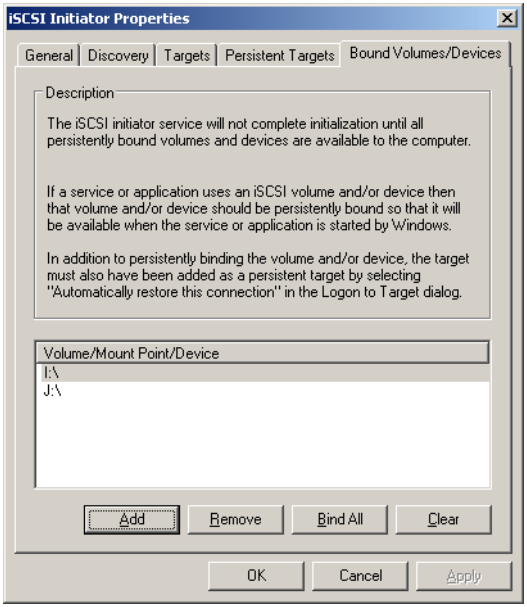

Figure 7-307 Microsoft iSCSI software initiator 2.0 configuration step 27

7.9 Linux iSCSI software initiator in SLES9

This description about the iSCSI software initiator configuration is based on SLES9 SP2. We describe both a single path and a multiple path setup.

7.9.1 Single path IO setup

The linux-iscsi RPM package must be installed before using the Linux iSCSI software initiator. The package is shipped with SLES9 and can be installed for example during installation of SLES9 SP2, or later on using YaST2. It is also possible to use the **rpm** command line utility to install the RPM package. In Example 7-1 the installation of linux-iscsi from the SP2 CD 1 with the **rpm** command line utility is shown. During installation of linux-iscsi the file /etc/initiatorname.iscsi containing the initiator name is created. After installation the initiator name can be queried with the **iscsi-iname** command.

Example 7-1 RPM installation of linux-iscsi using SLES9 SP2 CD1

```
danube:/media/dvd/suse/i586 # rpm -Uhv linux-iscsi-4.0.1-88.18.i586.rpm
Preparing...                ########################################### [100%]
   1:linux-iscsi            ########################################### [100%]
danube:/media/dvd/suse/i586 # cat /etc/initiatorname.iscsi
```

```
## DO NOT EDIT OR REMOVE THIS FILE!
## If you remove this file, the iSCSI daemon will not start.
## If you change the InitiatorName, existing access control lists
## may reject this initiator.  The InitiatorName must be unique
## for each iSCSI initiator.  Do NOT duplicate iSCSI InitiatorNames.
InitiatorName=iqn.1987-05.com.cisco:01.b411fc34c6b
danube:/media/dvd/suse/i586 # iscsi-iname
iqn.1987-05.com.cisco:01.d23d679b4b8
danube:/media/dvd/suse/i586 #
```

Before starting linux-iscsi, it is necessary to do the configuration in /etc/iscsi.conf. This file already contains many possible configuration parameters and some description. Detailed information about the configuration file is also available in the manpage of iscsi.conf. An initial configuration is shown in Example 7-2. This configuration can be used to ensure that the setup works, but should not be used for production use as it lacks any security configuration.

Example 7-2 initial /etc/iscsi.conf for single path IO

```
DiscoveryAddress=192.168.1.61
```

With this configuration, linux-iscsi can be started as shown in Example 7-3.

Example 7-3 starting linux-iscsi

```
danube:~ # /etc/init.d/iscsi start
Starting iSCSI: iscsi iscsid fsck/mount                              done
danube:~ #
```

Linux-iscsi logs interesting information to /var/log/messages. Example 7-4 shows logging information during startup of linux-iscsi.

Example 7-4 Log information in /var/log/messages after starting linux-iscsi with no multipath

```
Aug  3 11:21:51 danube kernel: iSCSI: 4.0.188.21 ( 21-May-2004) built for Linux
2.6.5-7.191-bigsmp
Aug  3 11:21:51 danube kernel: iSCSI: will translate deferred sense to current sense on disk
command responses
Aug  3 11:21:51 danube kernel: iSCSI: control device major number 254
Aug  3 11:21:51 danube kernel: scsi2 : SFNet iSCSI driver
Aug  3 11:21:51 danube kernel: iSCSI:detected HBA host #2
Aug  3 11:21:51 danube iscsid[7149]: version 4.0.188.18 ( 21-May-2004)
Aug  3 11:21:51 danube kernel: iSCSI: bus 0 target 2 =
iqn.1986-03.com.ibm:12f2.20050803112039.danube-disk1
Aug  3 11:21:51 danube kernel: iSCSI: bus 0 target 2 portal 0 = address 192.168.1.61 port 3260
group 1
Aug  3 11:21:51 danube kernel: iSCSI: bus 0 target 2 portal 1 = address 192.168.1.125 port 3260
group 1
```

```
Aug  3 11:21:51 danube kernel: iSCSI: starting timer thread at 847068
Aug  3 11:21:51 danube kernel: iSCSI: bus 0 target 2 trying to establish session to portal 0,
address 192.168.1.61 port 3260 group 1
Aug  3 11:21:51 danube kernel: iSCSI: bus 0 target 2 established session #1 to portal 0,
address 192.168.1.61 port 3260 group 1, alias danube-disk1
Aug  3 11:21:56 danube kernel:   Vendor: IBM       Model: DS300    S320     Rev: 6.21
Aug  3 11:21:56 danube kernel:   Type:   Direct-Access                     ANSI SCSI revision:
04
Aug  3 11:21:56 danube kernel: iSCSI: session iscsi_bus 0 target id 2 recv_cmd cdb 0x0, status
0x2, response 0x0, senselen 18, key 06, ASC/ASCQ 29/01, itt 6 to (2 0 2 0), danube-disk1
Aug  3 11:21:56 danube kernel: iSCSI: Sense 70000600 0000000a 00000000 29010000 0000
Aug  3 11:21:56 danube kernel: SCSI device sdb: 204800000 512-byte hdwr sectors (104858 MB)
Aug  3 11:21:56 danube kernel: SCSI device sdb: drive cache: write through
Aug  3 11:21:56 danube kernel:  sdb: unknown partition table
Aug  3 11:21:56 danube kernel: Attached scsi disk sdb at scsi2, channel 0, id 2, lun 0
Aug  3 11:21:56 danube kernel: Attached scsi generic sg2 at scsi2, channel 0, id 2, lun 0,
type 0
Aug  3 11:21:57 danube /etc/hotplug/block.agent[7178]: new block device /block/sdb
```

Note: The SCSI host number of iSCSI devices will raise each time linux-iscsi is restarted. In the example above, the disk sdb is at scsi2, channel 0, id 0, lun 0. A restart of linux-iscsi via its initscript will lead to a SCSI address with SCSI host scsi3. After each reboot, the counting starts again (in our example starting with scsi2).

The command **iscsi-ls** can be used to display details about connected iscsi volumes as shown in Example 7-5.

Example 7-5 output of iscsi-ls command with no multipath

```
danube:~ # iscsi-ls
*******************************************************************************
        Cisco iSCSI Driver Version ... 4.0.188.18 ( 21-May-2004 )
*******************************************************************************
TARGET NAME             : iqn.1986-03.com.ibm:12f2.20050803112039.danube-disk1
TARGET ALIAS            : danube-disk1
HOST NO                 : 2
BUS NO                  : 0
TARGET ID               : 2
TARGET ADDRESS          : 192.168.1.61:3260
SESSION STATUS          : ESTABLISHED AT Wed Aug  3 11:21:52 2005
NO. OF PORTALS          : 2
PORTAL ADDRESS 1        : 192.168.1.61:3260,1
PORTAL ADDRESS 2        : 192.168.1.125:3260,1
SESSION ID              : ISID 00023d000001 TSID 0e
*******************************************************************************
danube:~ #
```

In the example a single iSCSI volume is present. This volume appears also at the end of /proc/scsi/scsi as shown in Example 7-6.

Example 7-6 contents of /proc/scsi/scsi with no multipath

```
danube:~ # cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: VIRTUAL DISK  IM Rev: 1998
  Type:   Direct-Access                    ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 08 Lun: 00
  Vendor: IBM      Model: 25P3495a S320  1 Rev: 1
  Type:   Processor                        ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 02 Lun: 00
  Vendor: IBM      Model: DS300    S320    Rev: 6.21
  Type:   Direct-Access                    ANSI SCSI revision: 04
danube:~ #
```

Important: Because of variations in process scheduling and different network delays, iSCSI volumes may be mapped to different SCSI device names every time linux-iscsi is started. Ensure to use some form of persistent device naming (like udev) if you use more than a single iSCSI LUN. Take also into account that the SCSI host number changes for iSCSI devices when linux-iscsi is restarted.

The device is now ready for use. It can be partitioned and formatted.

7.9.2 Multipath IO setup

In this section we describe how to setup a multipath configuration with the device mapper multipath module of the Linux kernel, and the multipath-tools user-space package. SLES9 SP2 or higher is necessary for this configuration.

Important: We show that device mapper and the multipath-tools can be used for multipathing of iSCSI targets. We tested this configuration successfully in our lab setup and show the possibilities of device mapper. However, this configuration is not supported by IBM.

Before going further using device mapper multipath IO (DM MPIO), it is strongly recommend to read through Novell's paper *How to Setup/Use Multipathing on SLES*, which is online at the following URL:

http://portal.suse.com/sdb/en/2005/04/sles_multipathing.html

It is necessary to have the following packages installed to use DM MPIO:

- device-mapper (in our setup we use device-mapper-1.01.01-1.2, which is shipped with SP2 - check for the current version)
- multipath-tools (in our setup we use multipath-tools-0.4.4-0.22, which is shipped with SP2 - check for the current version)

To get a working DM MPIO configuration for linux-iscsi, the configuration file /etc/iscsi.conf needs additional parameters as shown in Example 7-7. Again, the shown configuration can be used to ensure that the setup works, but should not be used for production use as it lacks any security configuration.

Example 7-7 initial /etc/iscsi.conf for multipath IO

```
DiscoveryAddress=192.168.1.61
Multipath=portal
ConnFailTimeout=10
```

With Multipath=portal, linux-iscsi establishes iSCSI sessions to the target using each of the network portals in the target. Each session is bound to a different SCSI address in /proc/scsi/scsi. This will lead to two block devices (sdb and sdc) for the single logical drive in our example.

ConnFailTimeout=10 is necessary to fail SCSI commands after 10 Seconds. The value depends on your needs and on your network setup. The default value for ConnFailTimeout is 0, which means no timeout. The timeout is necessary to notify device mapper about a failed path.

Important: ConnFailTimeout=10 ensures a quick failover in case a network path from the server to the switch fails. If a controller failover in the DS300 happens, ConnFailTimeout=10 will lead to both paths failing as the controller failover takes longer than 10 seconds (see 3.3.5, “DS300 failover” on page 41). In this example we handle this with the queue_if_no_path setting in /etc/multipath.conf.

When starting linux-iscsi with this configuration, the log information changes compared to the single path IO configuration as shown in Figure 7-308 on page 423.

```
Aug  3 13:15:58 danube iscsid[6989]: version 4.0.188.18 ( 21-May-2004)
Aug  3 13:15:58 danube kernel: iSCSI: bus 0 target 2 =
iqn.1986-03.com.ibm:12f2.20050803112039.danube-disk1
Aug  3 13:15:58 danube kernel: iSCSI: bus 0 target 2 portal 0 = address 192.168.1.61 port 3260 group 1
Aug  3 13:15:58 danube kernel: iSCSI: starting timer thread at 268361
Aug  3 13:15:58 danube kernel: iSCSI: bus 0 target 2 trying to establish session to portal 0, address
192.168.1.61 port 3260 group 1
Aug  3 13:15:58 danube kernel: iSCSI: bus 1 target 2 =
iqn.1986-03.com.ibm:12f2.20050803112039.danube-disk1
Aug  3 13:15:58 danube kernel: iSCSI: bus 1 target 2 portal 0 = address 192.168.1.125 port 3260 group 1
Aug  3 13:15:58 danube kernel: iSCSI: bus 1 target 2 trying to establish session to portal 0, address
192.168.1.125 port 3260 group 1
Aug  3 13:15:58 danube kernel: iSCSI: bus 0 target 2 established session #1 to portal 0, address
192.168.1.61 port 3260 group 1, alias danube-disk1
Aug  3 13:15:58 danube kernel: iSCSI: bus 1 target 2 established session #1 to portal 0, address
192.168.1.125 port 3260 group 1, alias danube-disk1
Aug  3 13:16:03 danube kernel:   Vendor: IBM       Model: DS300    S320     Rev: 6.21
Aug  3 13:16:03 danube kernel:   Type:   Direct-Access                      ANSI SCSI revision: 04
Aug  3 13:16:03 danube kernel: iSCSI: session iscsi_bus 0 target id 2 recv_cmd cdb 0x0, status 0x2,
response 0x0, senselen 18, key 06, ASC/ASCQ 29/01, itt 6 to (2 0 2 0), danube-disk1
Aug  3 13:16:03 danube kernel: iSCSI: Sense 70000600 0000000a 00000000 29010000 0000
Aug  3 13:16:03 danube kernel: SCSI device sdb: 204800000 512-byte hdwr sectors (104858 MB)
Aug  3 13:16:03 danube kernel: SCSI device sdb: drive cache: write through
Aug  3 13:16:03 danube kernel:  sdb: unknown partition table
Aug  3 13:16:03 danube kernel: Attached scsi disk sdb at scsi2, channel 0, id 2, lun 0
Aug  3 13:16:03 danube kernel: Attached scsi generic sg2 at scsi2, channel 0, id 2, lun 0,  type 0
Aug  3 13:16:03 danube kernel:   Vendor: IBM       Model: DS300    S320     Rev: 6.21
Aug  3 13:16:03 danube kernel:   Type:   Direct-Access                      ANSI SCSI revision: 04
Aug  3 13:16:03 danube kernel: iSCSI: session iscsi_bus 1 target id 2 recv_cmd cdb 0x0, status 0x2,
response 0x0, senselen 18, key 06, ASC/ASCQ 29/01, itt 6 to (2 1 2 0), danube-disk1
Aug  3 13:16:03 danube kernel: iSCSI: Sense 70000600 0000000a 00000000 29010000 0000
Aug  3 13:16:03 danube kernel: SCSI device sdc: 204800000 512-byte hdwr sectors (104858 MB)
Aug  3 13:16:03 danube kernel: SCSI device sdc: drive cache: write through
Aug  3 13:16:03 danube kernel:  sdc: unknown partition table
Aug  3 13:16:03 danube kernel: Attached scsi disk sdc at scsi2, channel 1, id 2, lun 0
Aug  3 13:16:03 danube kernel: Attached scsi generic sg3 at scsi2, channel 1, id 2, lun 0,  type 0
Aug  3 13:16:03 danube /etc/hotplug/block.agent[7021]: new block device /block/sdb
Aug  3 13:16:03 danube /etc/hotplug/block.agent[7044]: new block device /block/sdc
```

Figure 7-308 Log information in /var/log/messages after starting linux-iscsi with Multipath=portal

Note: Notice that with the configuration of Multipath=portal in /etc/iscsi.conf we have now two block devices (sdb and sdc) visible for the single logical drive danube-disk1.

The **iscsi-ls** command shows now the logical drive with the alias danube-disk1 twice as shown in Example 7-8 on page 424.

Example 7-8 output of iscsi-ls command with Multipath in iscsi.conf

```
danube:~ # iscsi-ls
*******************************************************************************
        Cisco iSCSI Driver Version ... 4.0.188.18 ( 21-May-2004 )
*******************************************************************************
TARGET NAME             : iqn.1986-03.com.ibm:12f2.20050803112039.danube-disk1
TARGET ALIAS            : danube-disk1
HOST NO                 : 2
BUS NO                  : 0
TARGET ID               : 2
TARGET ADDRESS          : 192.168.1.61:3260
SESSION STATUS          : ESTABLISHED AT Wed Aug  3 13:15:58 2005
NO. OF PORTALS          : 1
PORTAL ADDRESS 1        : 192.168.1.61:3260,1
SESSION ID              : ISID 00023d000001 TSID 1c
*******************************************************************************
TARGET NAME             : iqn.1986-03.com.ibm:12f2.20050803112039.danube-disk1
TARGET ALIAS            : danube-disk1
HOST NO                 : 2
BUS NO                  : 1
TARGET ID               : 2
TARGET ADDRESS          : 192.168.1.125:3260
SESSION STATUS          : ESTABLISHED AT Wed Aug  3 13:15:58 2005
NO. OF PORTALS          : 1
PORTAL ADDRESS 1        : 192.168.1.125:3260,1
SESSION ID              : ISID 00023d000002 TSID 1d
*******************************************************************************
danube:~ #
```

This disk also appears twice in /proc/scsi/scsi as shown in Example 7-9.

Example 7-9 contents of /proc/scsi/scsi with Multipath=portal in iscsi.conf

```
danube:~ # cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: VIRTUAL DISK  IM Rev: 1998
  Type:   Direct-Access                    ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 08 Lun: 00
  Vendor: IBM      Model: 25P3495a S320  1 Rev: 1
  Type:   Processor                        ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 02 Lun: 00
  Vendor: IBM      Model: DS300    S320    Rev: 6.21
  Type:   Direct-Access                    ANSI SCSI revision: 04
Host: scsi2 Channel: 01 Id: 02 Lun: 00
  Vendor: IBM      Model: DS300    S320    Rev: 6.21
  Type:   Direct-Access                    ANSI SCSI revision: 04
```

```
danube:~ #
```

Before using the multipath-tools it is advantageous to create the configuration file /etc/multipath.conf. There is an example configuration file located in /usr/share/doc/packages/multipath-tools/multipath.conf.annotated for reference. As we are using ConnFailTimeout=10 in /etc/iscsi.conf we configure queue_if_no_path in /etc/multipath.conf as shown in Example 7-10. With this configuration, DM MPIO queues all IO in case of errors leading loss of all paths, and never propagates errors upwards.

We also added sda to the devnode_blacklist as this internal drive (sda) in the server should not be scanned for DM MPIO.

Example 7-10 /etc/multipath.conf

```
defaults {
        default_features "1 queue_if_no_path"
}

#
# name    : devnode_blacklist
# scope   : multipath & multipathd
# desc    : list of device names to discard as not multipath candidates
# default : cciss, fd, hd, md, dm, sr, scd, st, ram, raw, loop
#
devnode_blacklist {
        devnode cciss
        devnode fd
        devnode hd
        devnode md
        devnode dm
        devnode sr
        devnode scd
        devnode st
        devnode ram
        devnode raw
        devnode loop
        devnode sda
}
```

It is now necessary to create device targets for the multipath device. This is done by executing the following command:

/etc/init.d/boot.multipath start

Now the multipath daemon can be started:

/etc/init.d/multipathd start

Multipath devices now show up under /dev/disk/by-name/. With **multipath -l** it is possible to query the current path states.

Example 7-11 multipath device information and mounting

```
danube:~ # ls -l /dev/disk/by-name/
total 0
drwxr-xr-x  2 root root  88 Aug  3 11:33 .
drwxr-xr-x  5 root root 120 Jul 26 20:21 ..
lrwxrwxrwx  1 root root  10 Aug  3 11:33 20003390012f2c829 -> ../../dm-0
danube:~ # multipath -l
20003390012f2c829
[size=97 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [active][best]
  \_ 3:0:2:0 sdb  8:16    [active]
\_ round-robin 0 [enabled]
  \_ 3:1:2:0 sdc  8:32    [active]

danube:~ #
```

A filesystem can now be created on the device. Then the filesystem can be mounted.

Example 7-12 creating filesystem

```
danube:~ # mkfs.reiserfs /dev/disk/by-name/20003390012f2c829
mkfs.reiserfs 3.6.19 (2003 www.namesys.com)
[...]
ReiserFS is successfully created on /dev/disk/by-name/20003390012f2c829.
danube:~ # mkdir /mnt/danube-disk1
danube:~ # mount /dev/disk/by-name/20003390012f2c829 /mnt/danube-disk1/
danube:~ # df -h /mnt/danube-disk1/
Filesystem            Size  Used Avail Use% Mounted on
/dev/dm-0              98G   33M   98G   1% /mnt/danube-disk1
danube:~ #
```

Important: The setup uses DM MPIO with the iSCSI software initiator linux-iscsi-4.0.188.18 and the IP failover mechanism of the DS300. We tested it successfully at the time this book was written to show the possibilities of DM MPIO. If you rely on a fully supported configuration, do not use this setup as it is not supported by IBM! Also note that failover processing and options may be enhanced in future releases of DS300 firmware (see 3.3.5, “DS300 failover” on page 41). So the shown DM MPIO setup might not work with other failover mechanisms than IP failover.

To show how the DM MPIO setup works, we documented the following two test scenarios:

- Failover when network path from eth0 to switch fails
- “Failover in case of a controller failover in the DS300” on page 429

Failover when network path from eth0 to switch fails

In this test, the network cable from eth0 to the switch is unplugged. IO is created during the test with the following command:

```
dd if=/dev/urandom of=/mnt/danube-disk1/10GBfile bs=1M count=10000
```

Example 7-13 shows logging information in /var/log/messages during the outage of eth0.

Example 7-13 /var/log/messages during outage of eth0

```
Aug  3 13:23:20 danube kernel: tg3: eth0: Link is down.
Aug  3 13:23:29 danube kernel: iSCSI: 10 second timeout expired for session iscsi bus 0 target
id 2, rx 710007, ping 715006, now 720007
Aug  3 13:23:29 danube kernel: iSCSI: session iscsi bus 0 target id 2 iscsi_xmit_task -
xmit_data failed to send 65584 bytes,rc 54928
[...]
Aug  3 13:23:29 danube kernel: iSCSI: session iscsi bus 0 target id 2 to danube-disk1 dropped
Aug  3 13:23:29 danube kernel: iSCSI: bus 0 target 2 trying to establish session to portal 0,
address 192.168.1.61 port 3260 group 1
Aug  3 13:23:34 danube kernel: iSCSI: Requested ioctl not found
Aug  3 13:23:39 danube kernel: iSCSI: session iscsi bus 0 target id 2 replacement timed after
10 seconds, drop 720007, now 730007, failing all commands
Aug  3 13:23:39 danube multipathd: 8:16: readsector0 checker reports path is down
Aug  3 13:23:39 danube multipathd: checker failed path 8:16 in map 20003390012f2c829
Aug  3 13:23:39 danube multipathd: devmap event (2) on 20003390012f2c829
Aug  3 13:23:39 danube kernel: Device sdb not ready.
Aug  3 13:23:39 danube kernel: end_request: I/O error, dev sdb, sector 577864
Aug  3 13:23:39 danube kernel: device-mapper: dm-multipath: 8:16 (#577872): IO error - error:
-5 - bi_rw: 9 - bi_flags: 10 - bi_error: 00000000
Aug  3 13:23:39 danube kernel: device-mapper: dm-multipath: Failing path 8:16
Aug  3 13:23:39 danube kernel: device-mapper: dm-multipath: 8:16 (#577864): Requeued sector as
#1
Aug  3 13:23:39 danube kernel: device-mapper: dm-multipath: 8:16 (#577880): IO error - error:
-5 - bi_rw: 9 - bi_flags: 10 - bi_error: 00000000
Aug  3 13:23:39 danube kernel: device-mapper: dm-multipath: 8:16 (#577872): Requeued sector as
#2
[...]
```

The **multipath -l** command shows the new situation. The remaining path via eth1 is used to continue IO operations.

Example 7-14 multipath -l output after eth0 outage

```
danube:~ # multipath -l
```

```
20003390012f2c829
[size=97 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [enabled]
  \_ 2:0:2:0 sdb  8:16    [failed]
\_ round-robin 0 [active][best]
  \_ 2:1:2:0 sdc  8:32    [active]

danube:~ #
```

After the failed connection from eth0 to the switch comes up again, the available path is recognized by multipathd.

Example 7-15 /var/log/messages as eth0 comes up again

```
Aug  3 13:30:31 danube kernel: tg3: eth0: Link is up at 1000 Mbps, full duplex.
[...]
Aug  3 13:31:12 danube kernel: iSCSI: bus 0 target 2 trying to establish session to portal 0,
address 192.168.1.61 port 3260 group 1
Aug  3 13:31:12 danube kernel: iSCSI: bus 0 target 2 established session #2 to portal 0,
address 192.168.1.61 port 3260 group1, alias danube-disk1
Aug  3 13:31:16 danube kernel: iSCSI: session iscsi_bus 0 target id 2 recv_cmd cdb 0x28, status
0x2, response 0x0, senselen 18, key 06, ASC/ASCQ 29/01, itt 3761 to (2 0 2 0), danube-disk1
Aug  3 13:31:16 danube kernel: iSCSI: Sense 70000600 0000000a 00000000 29010000 0000
Aug  3 13:31:16 danube kernel: iSCSI: session iscsi bus 0 recv_cmd itt 3761, to (2 0 2 0), cdb
0x28,  U underflow, received 0, residual 512, expected 512
Aug  3 13:31:16 danube multipathd: 8:16: readsector0 checker reports path is up
Aug  3 13:31:16 danube multipathd: 8:16: reinstated
Aug  3 13:31:16 danube multipathd: devmap event (3) on 20003390012f2c829
```

The path via eth0 appears now as active (instead of failed) in the **`multipath -l`** output.

Example 7-16 multipath -l output after eth0 comes up again

```
danube:~ # multipath -l
20003390012f2c829
[size=97 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [enabled]
  \_ 2:0:2:0 sdb  8:16    [active]
\_ round-robin 0 [active][best]
  \_ 2:1:2:0 sdc  8:32    [active]

danube:~ #
```

Failover in case of a controller failover in the DS300

In this test, the network cable from eth3 of controller A of the DS300 is unplugged. This leads to a controller failover in the DS300. IO is again created during the test with the following command:

```
dd if=/dev/urandom of=/mnt/danube-disk1/10GBfile bs=1M count=10000
```

Example 7-17 shows some logging information as the controller failover starts.

Example 7-17 /var/log/messages as controller failover starts

```
Aug  3 13:39:56 danube kernel: iSCSI: target requests logout within 5 seconds
for session to danube-disk1
Aug  3 13:39:56 danube iscsid[7009]: socket 1 closed by target
Aug  3 13:39:56 danube iscsid[7009]: discovery session to  192.168.1.61:3260
failed to recv a PDU response, reconnecting
Aug  3 13:39:56 danube iscsid[7009]: cannot make connection to
192.168.1.61:3260: Connection refused
Aug  3 13:39:56 danube iscsid[7009]: connect to 192.168.1.61 failed
```

Because of the ConnFailTimeout=10 setting in /etc/iscsi.conf both paths fail as the controller failover takes about 45 - 60 seconds in the test.

Example 7-18 multipath -l during the controller failover

```
danube:~ # multipath -l
20003390012f2c829
[size=97 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [enabled][best]
  \_ 2:0:2:0 sdb  8:16    [failed]
\_ round-robin 0 [enabled]
  \_ 2:1:2:0 sdc  8:32    [failed]

danube:~ #
```

The queue_if_no_path setting /etc/multipath.conf ensures that IO will be queued in this situation.

Example 7-19 /var/log/messages as controller failover finishes

```
Aug  3 13:40:36 danube kernel: iSCSI: bus 1 target 2 trying to establish
session to portal 0, address 192.168.1.125 port 3260group 1
Aug  3 13:40:36 danube kernel: iSCSI: bus 1 target 2 established session #3 to
portal 0, address 192.168.1.125 port 3260 group 1, alias danube-disk1
Aug  3 13:40:36 danube kernel: iSCSI: session iscsi_bus 1 target id 2 recv_cmd
cdb 0x28, status 0x2, response 0x0, senselen 18, key 06, ASC/ASCQ 29/01, itt
18716 to (2 1 2 0), danube-disk1
```

```
Aug  3 13:40:36 danube kernel: iSCSI: Sense 70000600 0000000a 00000000 29010000
0000
Aug  3 13:40:36 danube kernel: iSCSI: session iscsi bus 1 recv_cmd itt 18716,
to (2 1 2 0), cdb 0x28,  U underflow, received 0, residual 512, expected 512
Aug  3 13:40:36 danube multipathd: 8:32: readsector0 checker reports path is up
Aug  3 13:40:36 danube multipathd: 8:32: reinstated
Aug  3 13:40:36 danube multipathd: devmap event (8) on 20003390012f2c829
Aug  3 13:40:40 danube kernel: iSCSI: Requested ioctl not found
Aug  3 13:40:41 danube kernel: iSCSI: bus 0 target 2 trying to establish
session to portal 0, address 192.168.1.61 port 3260 group 1
Aug  3 13:40:41 danube kernel: iSCSI: bus 0 target 2 established session #3 to
portal 0, address 192.168.1.61 port 3260 group1, alias danube-disk1
Aug  3 13:40:41 danube kernel: iSCSI: session iscsi_bus 0 target id 2 recv_cmd
cdb 0x28, status 0x2, response 0x0, senselen 18, key 06, ASC/ASCQ 29/01, itt
5818 to (2 0 2 0), danube-disk1
Aug  3 13:40:41 danube kernel: iSCSI: Sense 70000600 0000000a 00000000 29010000
0000
Aug  3 13:40:41 danube kernel: iSCSI: session iscsi bus 0 recv_cmd itt 5818, to
(2 0 2 0), cdb 0x28,  U underflow, received 0, residual 512, expected 512
Aug  3 13:40:41 danube multipathd: 8:16: readsector0 checker reports path is up
Aug  3 13:40:41 danube multipathd: 8:16: reinstated
Aug  3 13:40:41 danube multipathd: devmap event (9) on 20003390012f2c829
```

Now both paths are up again.

Note: Failing iSCSI paths may cause many log entries. When testing your environment watch the size of /var/log/messages in different failure scenarios. Placing /var/log on a separate filesystem helps to prevent a filled root filesystem in case that logfiles grow too fast.

7.10 Linux iSCSI software initiator in RHEL4

At the time of writing this book, RHEL4 had no iSCSI software initiator support. There is a statement about that in the Red Hat Enterprise Linux AS 4 Release Notes: The initial release of Red Hat Enterprise Linux 4 does not include iSCSI software initiator or target support. Support for iSCSI is being evaluated for addition in a future update to Red Hat Enterprise Linux 4.

Part 3

TotalStorage DS400

In Part 3 we discuss planning, installation, and configuration considerations for the IBM TotalStorage DS400 storage subsystem.

8

TotalStorage DS400 example configurations

This chapter gives an overview of the sample DS400 SAN configurations that are used in later chapters to explain, in detail, TotalStorage DS400 installation and configuration.

This chapter contains these topics:

8.1 SAN zoning considerations

Fibre Channel switches allow the configuration of zoning. Zones are used to create sets of devices which should be able to communicate with each other through a fabric. A device is either an initiator such as an HBA or a target such as a disk controller or tape drive.

With single HBA zoning, each zone contains exactly one initiator (HBA), but can contain multiple targets. If an HBA should connect to both disk and tape devices, we suggest you create two zones for this HBA (one containing the HBA and the disk devices, and the other one containing the HBA and the tape devices). This leads to at least as many zones as HBAs exist in the SAN. But in this way, each zone contains only a limited number of members.

Whenever a device is added to or removed from a fabric, a so-called *Registered State Change Notification* (RSCN) occurs. Without zoning, all devices attached to a fabric receive this RSCN. Each device then queries the nameserver of the fabric to get updated information about the membership in the fabric. This also occurs if the notified device does not care about the device which caused the RSCN.

In a SAN, initiators only communicate with targets. Initiators do not communicate with each other, as targets also do not communicate with each other. So initiators have no use for RSCNs caused by other initiators. The same applies for targets.

Single HBA zoning helps to reduce the impact of adding or removing devices in a SAN only to those members which are in the same zone. It fences off different operating systems to reduce challenges related to operating system and device driver interoperability. It also can minimize unwanted interactions between devices and facilitate fault isolation.

More information about zoning can be found for example in the whitepaper *Zoning implementation strategies for Brocade SAN switches*, available from:

http://www.brocade.com/san/white_papers/pdf/Zoning_Imp_WP_00.pdf

8.2 DS400 single controller example configurations

In this section we describe two example configurations for the DS400 single controller model.

8.2.1 DS400 single controller direct attachment

The DS400 single controller model has two Fibre Channel ports. Thus it allows the direct attachment of up to two host systems, each having a single path to the DS400.

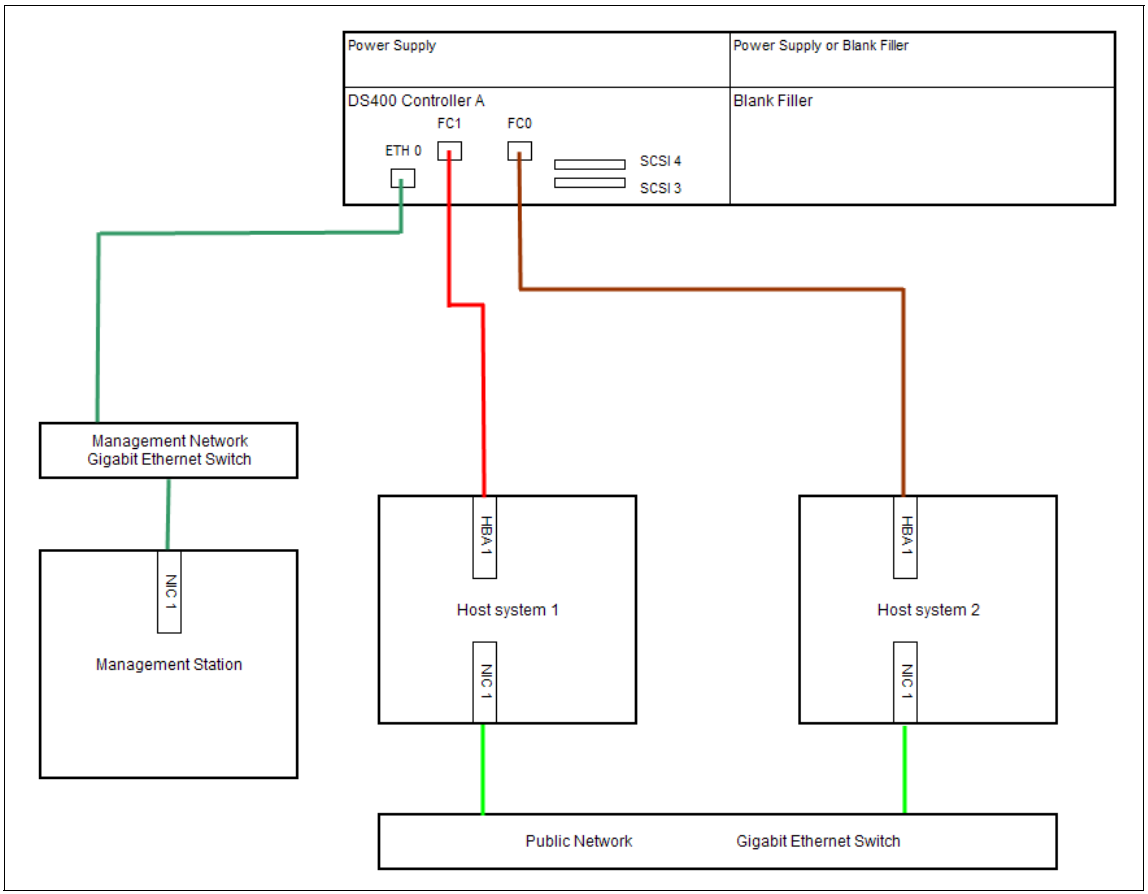

Figure 8-1 DS400 single controller direct attachment

8.2.2 DS400 single controller fabric attachment

In case that more than two host systems should be attached to a DS400 single control model, a SAN switch is necessary. Figure 8-2 on page 436 shows a configuration with a SAN switch which allows additional host systems to be connected to the DS400.

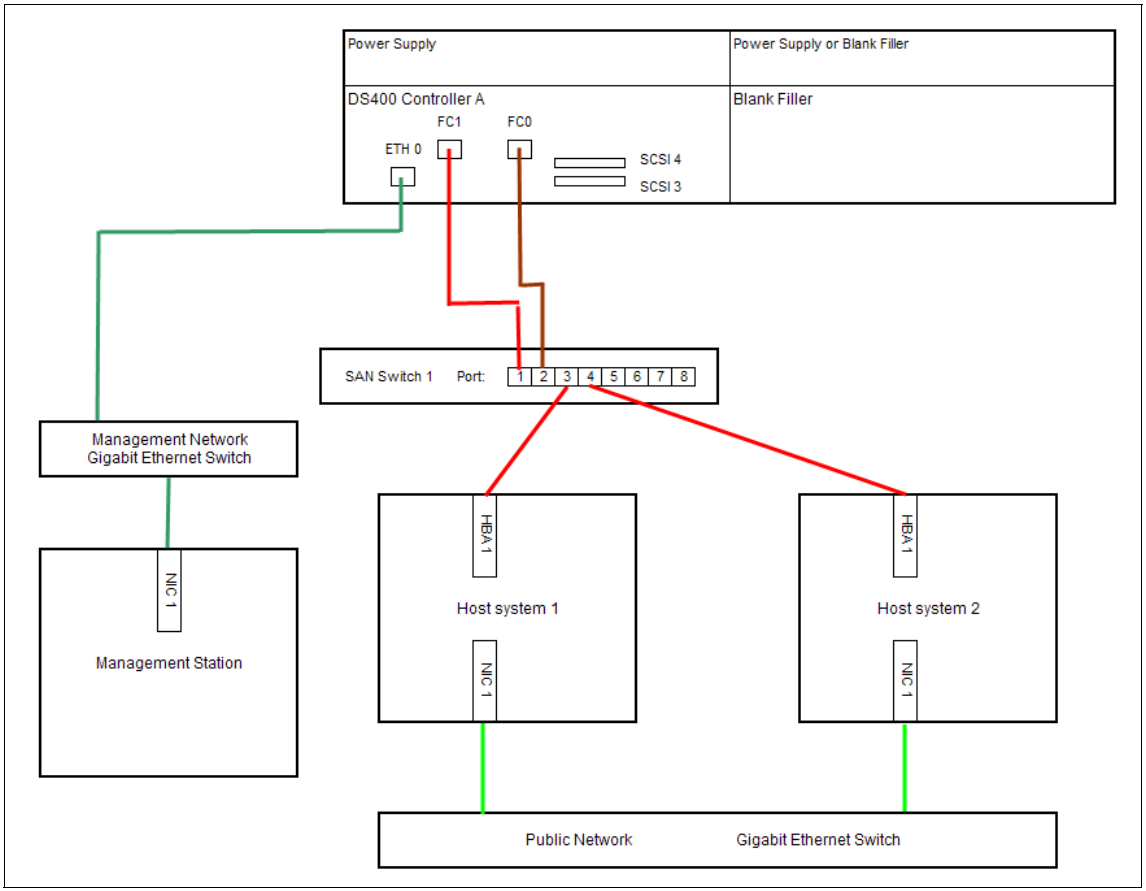

Figure 8-2 DS400 single controller fabric attachment

Table 8-1 shows a example zoning configuration for DS400 single controller fabric attachment.

Table 8-1 Zoning configuration fabric attachment with single paths

Zone	Initiator	Target
Zone1	Host1_HBA1	DS400_ControllerA_FC1
Zone2	Host2_HBA1	DS400_ControllerA_FC0

8.3 DS400 dual controller example configurations

In this section we describe example configurations for the DS400 dual controller model.

The DS400 dual controller model allows the configuration of a fully redundant path from host systems to the logical drives, preventing single points of failure. We recommend to use two HBAs in each attached server to take advantage of this possibility.

Important: Controller failover within the DS400 occurs on an array basis, not on a logical drive basis. If a host system accesses a logical drive, and loses its path to the controller currently managing the array where the logical drive resides, the whole array will be switched to the other controller. Other servers also accessing logical drives within this array are also forced to switch to the other controller.

It is best practice to assign logical drives within an individual array only to a single host system. This eliminates the impact of such failovers to other host systems.

8.3.1 DS400 dual controller direct attachment single path

In small environments it is possible to directly attach up to four host systems, each with one host bus adapter, with a single path to the TotalStorage DS400. No multipath driver is required. Most operating systems should have no problems with this configuration.

There are some limitations which may result in a host system no longer being able to access the logical drives of the TotalStorage DS400. This can happen after a controller failover. In case of a controller failure, all arrays with the logical drives will be failed over to the alternate controller. The hosts attached to the original controller will not longer be able than to access the logical drives because there exists only one path. This is simple configuration that leaves very little room for failover and recovery and should be avoided.

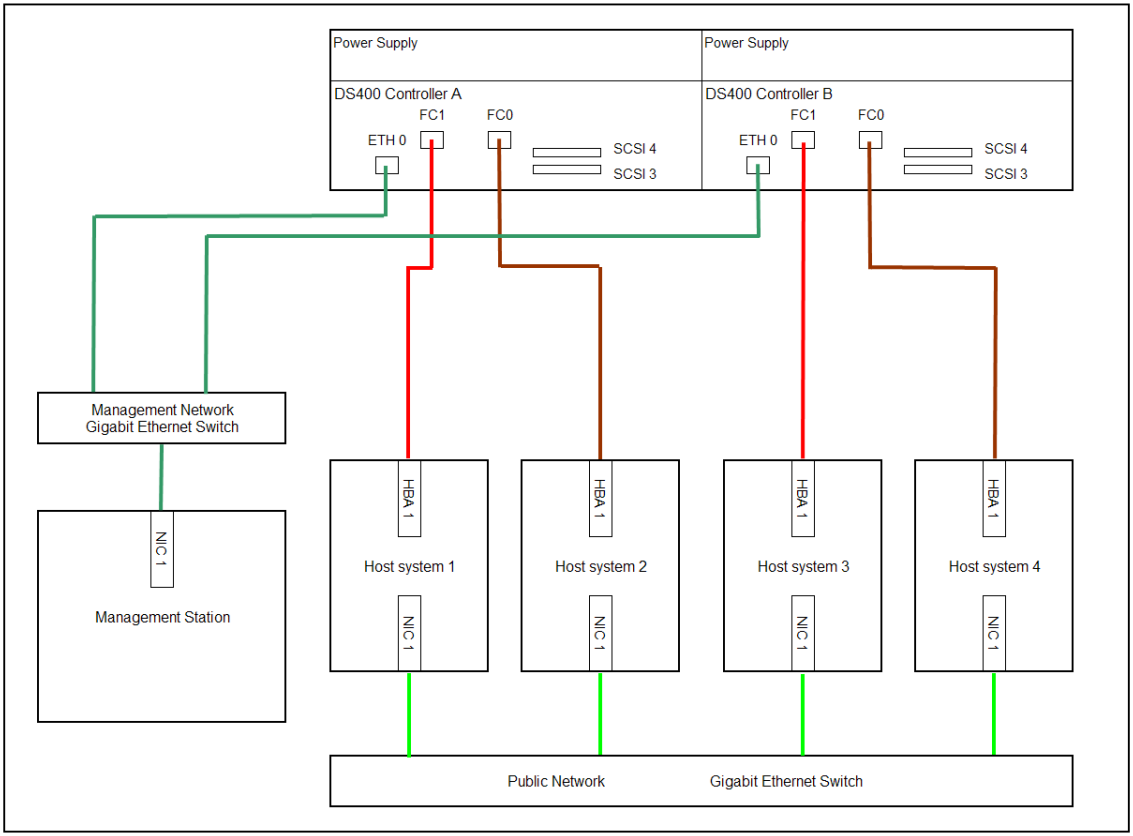

Figure 8-3 DS400 single path attachment

8.3.2 DS400 dual controller direct attachment dual path

In small environments it is possible to directly attach up to two host systems, each with two HBAs, with two paths to the DS400. It is important that both hosts be attached to both controllers to be protected against a controller failure. A failover driver is required for this situation. The logical must be assigned to both initiators.

When no failover driver is available for the host operating system, each logical drive must be assigned only to a single initiator. This is a configuration with two single paths.

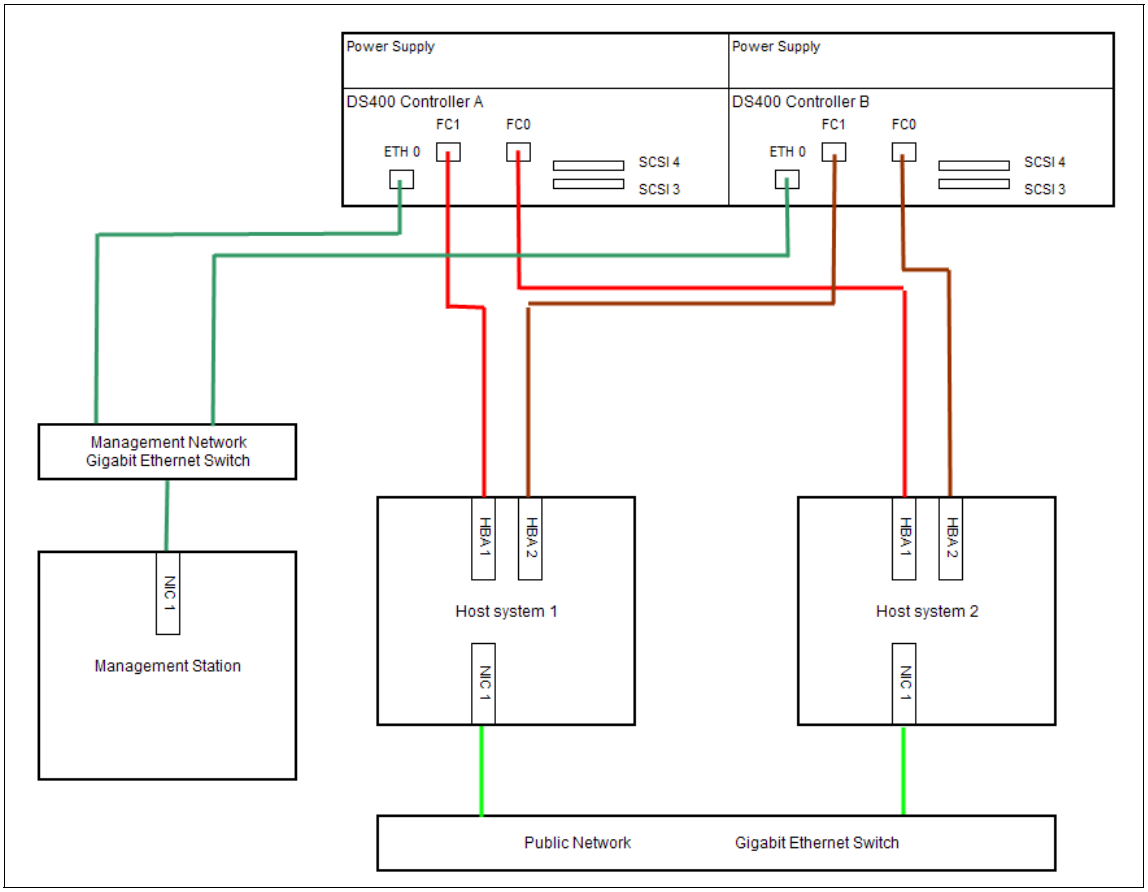

Figure 8-4 DS400 dual controller direct attachment

Important: Assign logical drives within an individual array only to a single host system in this configuration. Otherwise the access of a second host system with to a logical drive may be interrupted, if the first host system forces a controller failover for the array.

8.3.3 DS400 dual controller fabric attachment

In case that more than two host systems should be attached with multiple paths to a DS400 dual controller model, SAN switches are necessary. To prevent single points of failure, two SAN switches or SAN directors should be used.

We describe three example configurations of a DS400 dual controller fabric attachment:

- Fabric attachment with dual paths
- Fabric attachment with quad paths through interswitch link
- Fabric attachment with quad paths without interswitch link

Fabric attachment with dual paths

A very straightforward configuration of a fabric attachment with dual paths is shown in Figure 8-5 on page 441. Each host system has two paths to the DS400. One path is from the first HBA of the host system to controller A of the DS400. The other path is from the second HBA of the host system to controller B of the DS400.

An outage of an HBA or its connection to the SAN switch leads to a controller failover for an array if the active path ran over this HBA. The other configurations with quad paths cope with such an outage and can use a second path to the same controller, preventing a controller failover.

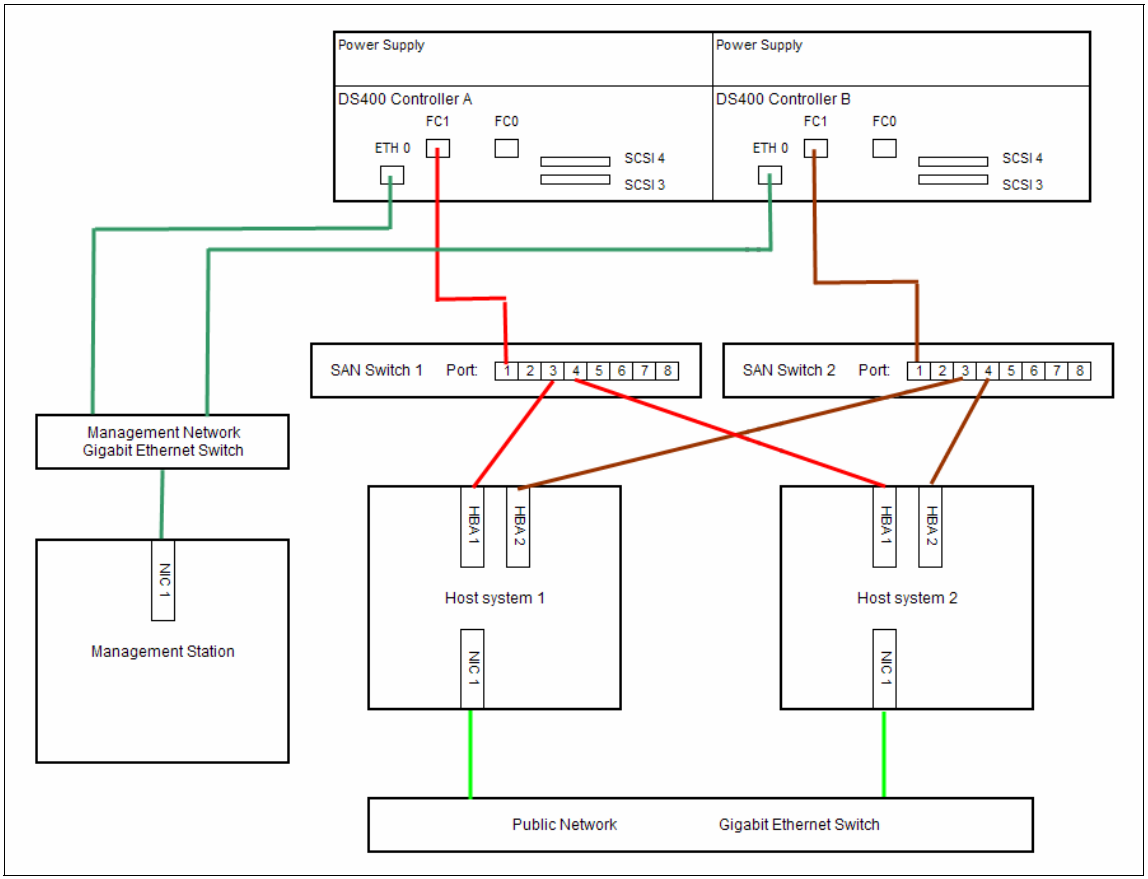

Figure 8-5 DS400 dual controller fabric attachment with dual paths

Table 8-2 shows a example zoning configuration for DS400 dual controller fabric attachment with dual paths.

Table 8-2 Zoning configuration fabric attachment with dual paths

Zone	Initiator	Target
Zone1 on switch1	Host1_HBA1	DS400_ControllerA_FC1
Zone1 on switch2	Host1_HBA2	DS400_ControllerB_FC1
Zone2 on switch1	Host2_HBA1	DS400_ControllerA_FC1
Zone2 on switch2	Host2_HBA2	DS400_ControllerB_FC1

Important: Assign logical drives within an individual array only to a single host system in this configuration. Otherwise the access of a second host system to a logical drive may be interrupted, if the first host system forces a controller failover for the array.

Fabric attachment with quad paths through interswitch link

In this configuration the two switches build a single fabric through the interswitch link (ISL) as shown in Figure 8-6 on page 443. They share their zone configuration.

Attention: Connecting two switches to a single fabric through an interswitch link raises the impact of configuration failures. Zone configuration data is automatically updated to all switches which are part of the fabric. With two separated fabrics, hosts may still have connectivity to the DS400 even if one fabric fails because of incorrect configuration. Therefore we recommend to use the configuration described in “Fabric attachment with quad paths without interswitch link” on page 444 instead of this one, but we discuss this configuration to give an overview about possible implementations.

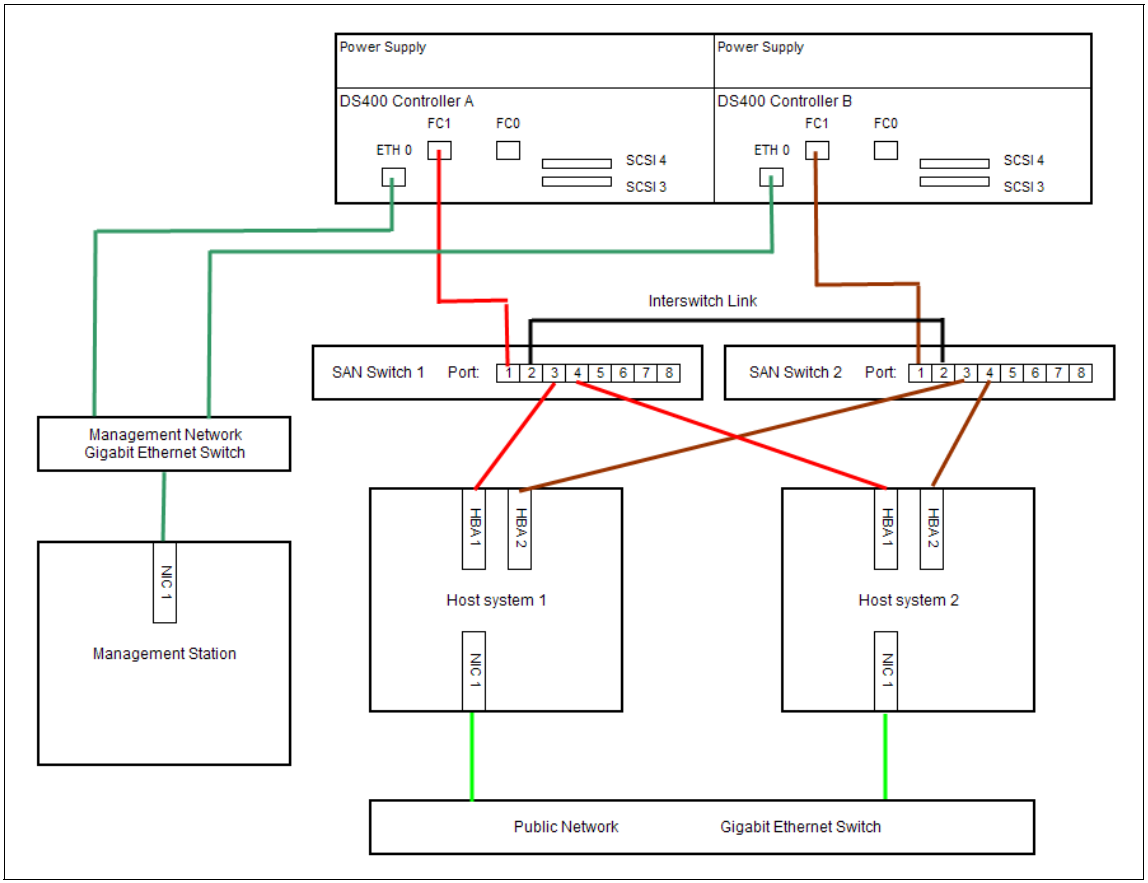

Figure 8-6 DS400 dual controller fabric attachment with quad paths with ISL

Table 8-3 shows a example zoning configuration for DS400 dual controller fabric attachment with quad paths through ISL.

Table 8-3 Zoning configuration fabric attachment with quad paths through ISL

Zone	Initiator	Targets
Zone1	Host1_HBA1	DS400_ControllerA_FC1 DS400_ControllerB_FC1
Zone2	Host1_HBA2	DS400_ControllerA_FC1 DS400_ControllerB_FC1
Zone3	Host2_HBA1	DS400_ControllerA_FC1 DS400_ControllerB_FC1

Zone	Initiator	Targets
Zone4	Host2_HBA2	DS400_ControllerA_FC1 DS400_ControllerB_FC1

Note: Although in this configuration a controller failover for an array is less often necessary compared to “Fabric attachment with dual paths” on page 440, it is recommend that logical drives within an individual array be assigned only to a single host system.

Fabric attachment with quad paths without interswitch link

This configuration provides four paths from a host system to the DS400. In the event of an outage of a HBA or its connection to the SAN switch the host system can use a second path to controller in the DS400, preventing a failover for the array.

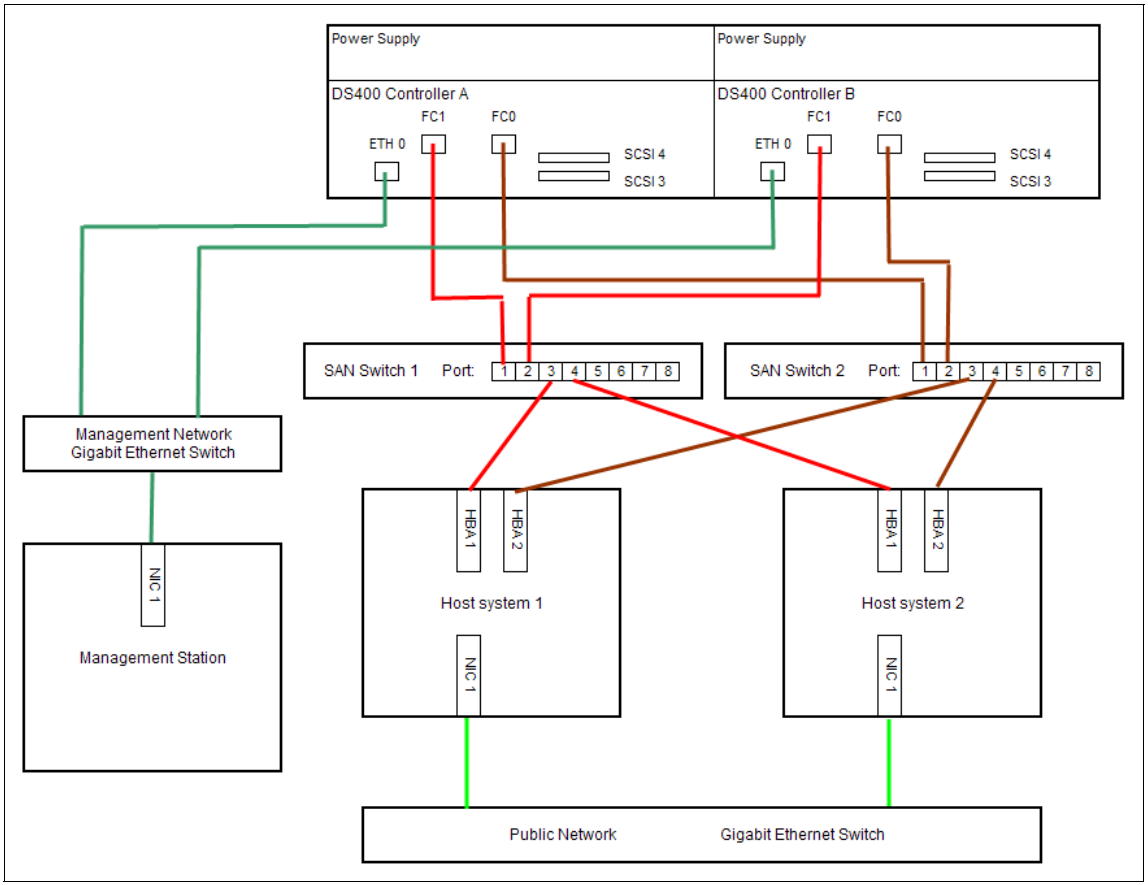

Figure 8-7 DS400 dual controller fabric attachment with quad paths without ISL

Table 8-4 shows a example zoning configuration for DS400 dual controller fabric attachment with quad paths without ISL.

Table 8-4 Zoning configuration fabric attachment with quad paths without ISL

Zone	Initiator	Targets
Zone1 on switch1	Host1_HBA1	DS400_ControllerA_FC1 DS400_ControllerB_FC1
Zone1 on switch2	Host1_HBA2	DS400_ControllerA_FC0 DS400_ControllerB_FC0
Zone2 on switch1	Host2_HBA1	DS400_ControllerA_FC1 DS400_ControllerB_FC1

Zone	Initiator	Targets
Zone2 on switch2	Host2_HBA2	DS400_ControllerA_FC0 DS400_ControllerB_FC0

Note: Although in this configuration a controller failover for an array is less often necessary compared to “Fabric attachment with dual paths” on page 440, it is recommend that logical drives within an individual array be assigned only to a single host system.

9

TotalStorage DS400 installation

This chapter provides some examples and describes the setup in more detail. The following scenarios are described:

- 9.1, “Boot Microsoft Windows Server 2003 from DS400 with IBM SMB Host Bus Adapter” on page 448
- 9.2, “Boot Microsoft Windows Server 2003 from DS400 with the IBM DS4000 FC2-133 Host Bus Adapter” on page 501

9.1 Boot Microsoft Windows Server 2003 from DS400 with IBM SMB Host Bus Adapter

The setup of an IBM @server xSeries 236 with two IBM SMB Host Bus Adapters connected with multiple paths to a DS400 will be explained in detail. The IBM SMB Host Bus Adapter is used to boot the system from the DS400.

9.1.1 Planning and considerations

The IBM SMB host bus adapter has a feature which allows you to boot from a LUN 0 of any target the adapter finds during the POST. This allows you to boot an IBM @serverxSeries 236 without specifying a dedicated boot target and LUN. The advantage is a simpler setup and the ability that the server will boot no matter which controller owns the logical drive to boot from. This is archived by using the IBM SMB Host Bus Adapter and the cabling of the TotalStorage DS400 as shown in “Fabric attachment with quad paths without interswitch link” on page 444.

When a problem occurs which might prevent the server from booting, the troubleshooting is much more complicated than using local boot. There is no operating system to check what went wrong.

The setup consists of redundant paths to the storage. This results in seeing the same disk multiple times. This is a kind of shared storage and this cannot be handled by Microsoft Windows. A setup change for the period of operating system installation is required. The steps are described in this example when they are needed.

All management interfaces of the various components are connected into their own management network. This is good practice and allows access restrictions to those persons who need access and restricts it from those who do not. When using an IBM BladeCenter server to connect to a TotalStorage DS400 the separate management network is a requirement to prevent network problems.

9.1.2 SAN and network setup

The following diagram shows the cabling of the components. It follows the example configuration in “Fabric attachment with quad paths without interswitch link” on page 444.

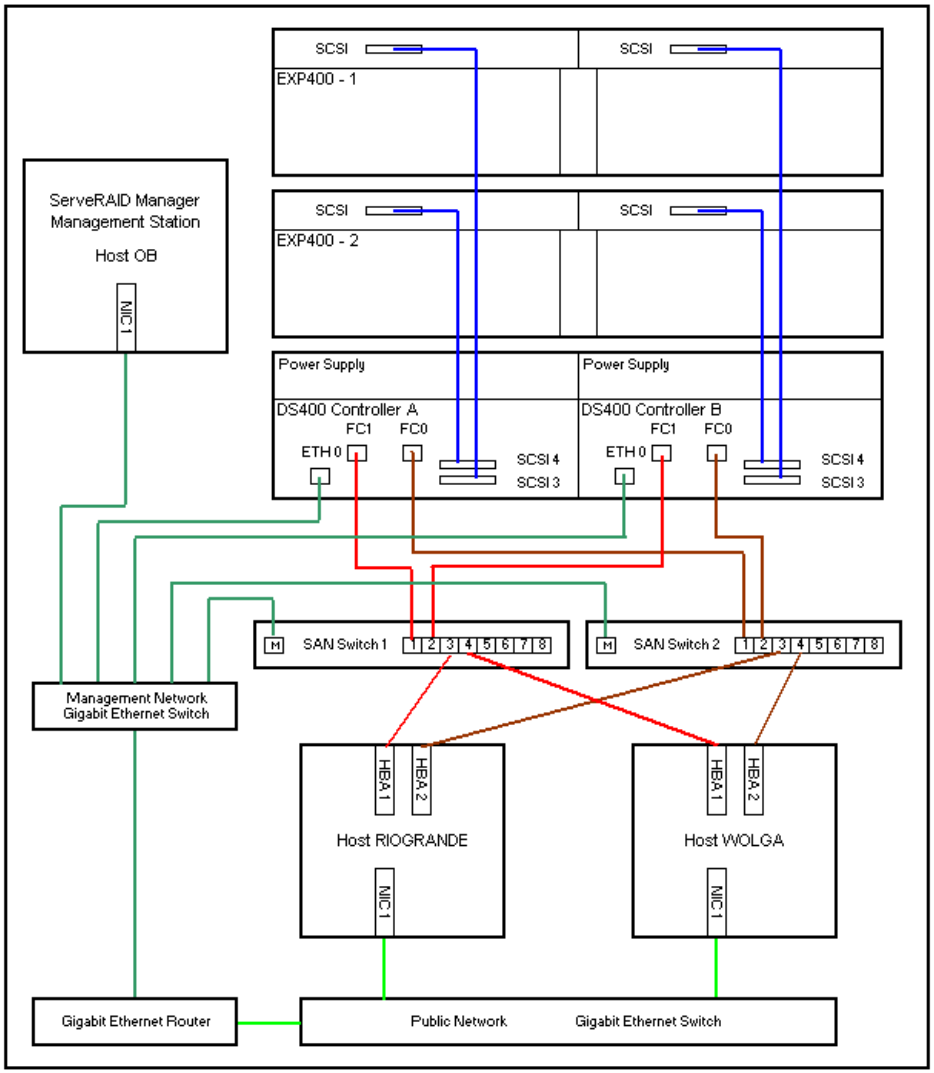

Figure 9-1 Configuration setup

The setup has used a few names which will be found in the various screenshots. They are:

- System name is `RIOGRANDE`.
- First host bus adapter (HBA1), IBM SMB Host Bus Adapter, is installed in slot 3 and is detected by IBM FAStT MSJ as port 0. The world wide name has got the alias `RIOGRANDE-S3-(0)` assigned. This HBA has access to FC0 from controller A and B. Port FC0 as described on the rear site of the TotalStorage

DS400 and has the name A0 in the CLI for controller A and B0 for controller B.

- Second Host bus adapter (HBA2), an IBM SMB Host Bus Adapter, is installed in slot 4 and is detected by IBM FAStT MSJ as port 1. The world wide name has got the alias RIOGRANDE-S4-(1) assigned. This HBA has access to FC1 from controller A and B. Port FC1 as described on the rear site of the TotalStorage DS400 and has the name A1 in the CLI for controller A and B1 for controller B.
- The logical drive used to boot from is named RIOGRANDE-BOOT.

Each switch has a zone for each connected HBA configured. Soft zoning was used. Every world wide name has got an alias. The alias was used during zone configuration. SAN switch 1 has got configured a zone Z_RIOGRANDE_HBA1 with the members A_RIOGRANDE_HBA1, A_DS400_1_A_FC0 and A_DS400_1_B_FC0. SAN switch 2 has got configured a zone Z_RIOGRANDE_HBA2 with the members A_RIOGRANDE_HBA2, A_DS400_1_A_FC1 and A_DS400_1_B_FC1. The zones where added to a configuration and the configuration was activated.

Note: It is important for this description that the array of the logical drive RIOGRANDE-BOOT is owned by controller A. When the alternate controller owns the array of the logical drive, the port names needs to be changed to reflect that the access goes over a port of controller B and not A.

9.1.3 TotalStorage DS400 setup

For the installation, only one logical drive was created. The world wide names where manually added into the initiator configuration and this logical drive was added to both initiators. At time this redbook was written, no multipathing driver was available. This driver can be installed after the operating system installation. To prevent trouble with multiple disk drives with the same signatures, the allowed ports for an initiator to access a logical drive where modified. The result is that only a single logical drive is seen during installation time and no recabling or change in the zoning configuration is required. By default, an initiator is allowed to access a logical drive over all for fibre channel ports. The CLI supports disabling and enabling every port for a given initiator. To get more details on managing a TotalStorage DS400 see Chapter 4, “TotalStorage DS300 and DS400 management” on page 67. 4.8.2, “Create arrays, logical drives, and access control on the TotalStorage DS400” on page 131 explains in detail the creation of logical drives and array. The steps below explain the tasks performed after the array and logical drive creation.

Note: The described tasks can only be performed from the command line interface (CLI).

1. Use a Telnet to connect to one of the management IPs of the DS400. This book uses as example controller A. Logon with administrative rights.

```
----[ IBM DS400: Command Line Interface ]---------------------------------

               Copyright:  IBM (c)2003-2004

                Firmware:  IBM TotalStorage DS Series (Build 0743)
                           [20/07/2005][6.21]

  Peer controller status:  Ready


           Typing 'help':  provides an overview of the commands available to
                           configure the controller.

                 TAB key:  use the TAB key as a means of getting assistance
                           in completing any commands that you issue.


DS400[A]> administrator
Administrator level password: ********
DS400[A]#
```

Figure 9-2 Logon

2. Get the list of initiators and aliases defined during setup of the DS400.

```
DS400[A]# initiator list

Fibre Channel Initiators

      210000e08b0f1549        [WOLGA-SLOT4-(1)]
      210000e08b0ff748        [WOLGA-SLOT3-(0)]
      210000e08b170d75        [RIOGRANDE-S4-(1)]
      210000e08b17d774        [RIOGRANDE-S3-(0)]

DS400[A]#
```

Figure 9-3 Initiator list

3. Disable the access of the second HBA with alias `RIOGRANDE-S4-(1)` to any port.

```
DS400[A]# initiator manage RIOGRANDE-S4-(1) deny all
Access denied to initiator RIOGRANDE-S4-(1) on all ports

DS400[A]#
```

Figure 9-4 Disable initiator access

4. To be able to install onto the logical drive on the DS400, one port must allow access for an initiator. The port used is A0. A1, B0 and B1 will be deactivated for the first HBA.

```
DS400[A]# initiator manage RIOGRANDE-S3-(0) deny A1 B0 B1
Access denied to initiator RIOGRANDE-S3-(0) on ports A1 B0 and B1

DS400[A]#
```

Figure 9-5 Disable initiator access for certain ports

5. Verify that only one port is enabled for only one HBA. Only port A0 allows access for the first HBA.

```
DS400[A]# initiator manage RIOGRANDE-S3-(0) info

----[ Initiators ]-----------------------------------------------------


----[ Initiator 'RIOGRANDE-S3-(0)' ]-----------------------------------

                WWN:  210000e08b17d774
              Alias:  RIOGRANDE-S3-(0)
                 ID:  4
        Port access:
                      Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------
                      YES            NO             NO             NO


       LUN Mappings:
                LUN                 logical          [Capacity]
-------------------------------------------------------------------------

                0                   RIOGRANDE-BOOT   [ 10.00GB]
                1                   RIOGRANDE-BOOT2  [9000.00MB]

DS400[A]#
```

Figure 9-6 Initiator information

6. No access on any port is allowed for the second HBA.

```
DS400[A]# initiator manage RIOGRANDE-S4-(1) info

----[ Initiators ]-----------------------------------------------------


----[ Initiator 'RIOGRANDE-S4-(1)' ]-----------------------------------

                 WWN:  210000e08b170d75
               Alias:  RIOGRANDE-S4-(1)
                  ID:  3
         Port access:
                        Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------------
                        NO             NO             NO             NO


        LUN Mappings:
                 LUN                 logical          [Capacity]
-------------------------------------------------------------------------------

                 0                   RIOGRANDE-BOOT   [ 10.00GB]
                 1                   RIOGRANDE-BOOT2  [9000.00MB]

DS400[A]#
```

Figure 9-7 Initiator information

9.1.4 Host bus adapter setup and installation preparation

The IBM SMB Host Bus Adapter has limited functionality compared with a IBM DS4000 host bus adapter. Please see “IBM TotalStorage SMB Host Bus Adapter (QLA200) 13N1873” on page 62 for more details.

The following steps are done to configure the IBM SMB Host Bus Adapters in the x236.

1. Update BIOS and firmware of the HBA to a supported level by IBM. During our setup, the BIOS level 1.23 was used.

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

2. Reset the HBA to adapter defaults.

   a. During POST, the BIOS of the IBM SMB Host Bus Adapter displays a message. Press **Ctrl+Q** to access Setup. Figure 9-8 on page 456.

Note: ALT+Q works as well.

Figure 9-8 QLA200 BIOS Message

b. Select the second HBA from the HBA list.

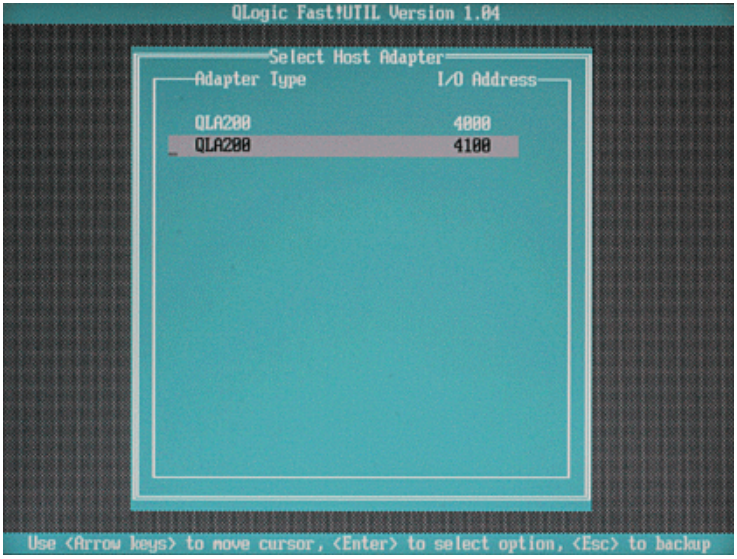

Figure 9-9 Select host adapters

c. Select **Configuration Settings** and press **Enter**.

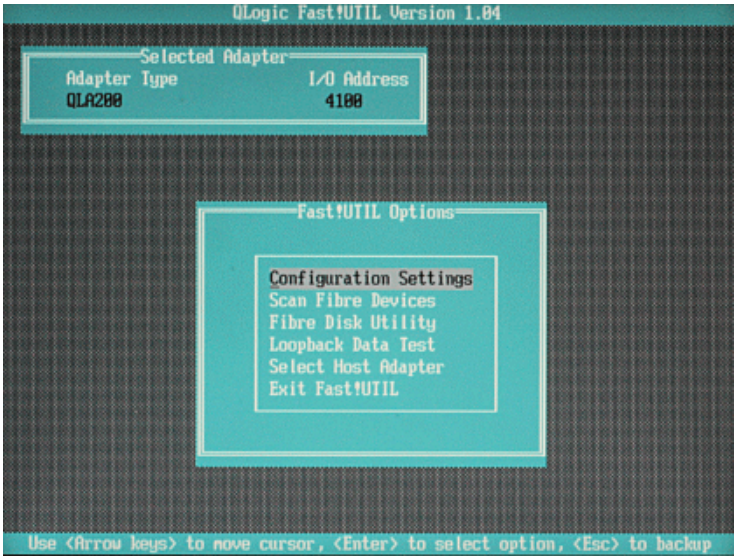

Figure 9-10 Fast!UTIL options — Configuration settings

d. Select **Restore Default Settings** and press **Enter**.

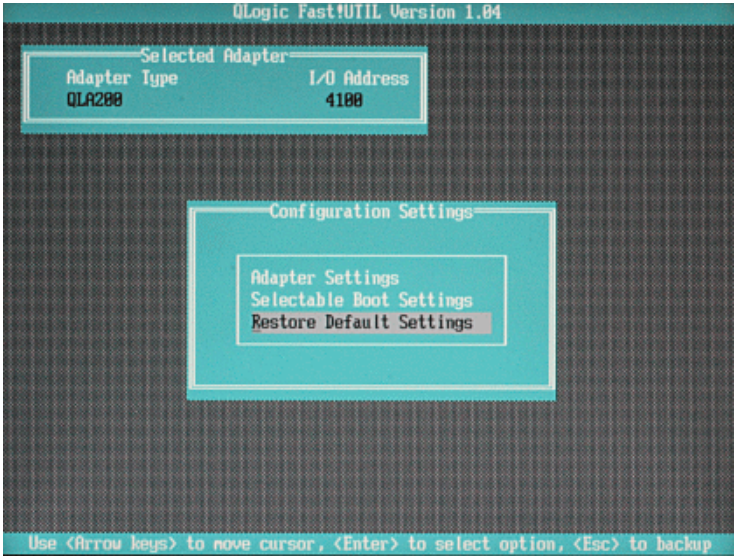

Figure 9-11 Configuration settings — Restore default settings

e. The adapter settings will be restored to the default settings.Press any key to return to **Configuration Settings** dialog.

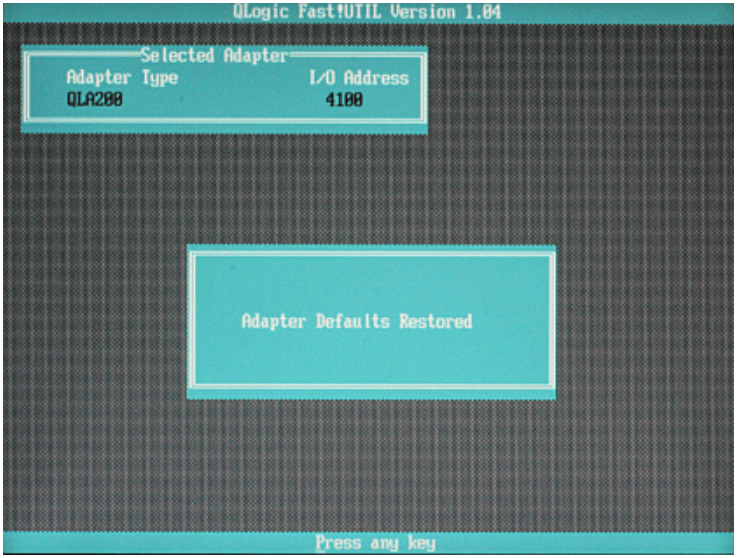

Figure 9-12 Select host adapters

f. Press **ESC** key to leave **Configuration Settings**.

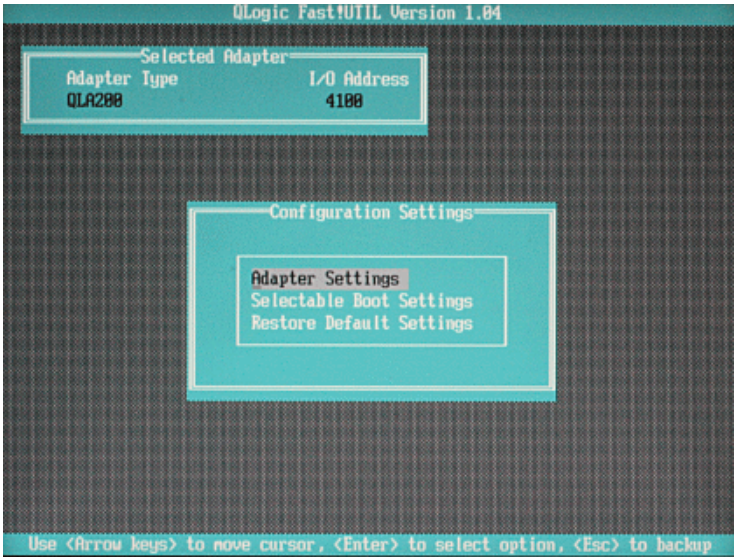

Figure 9-13 Select host adapters

g. The configuration was modified and needs to be saved. Select **Save changes** and press **Enter**.

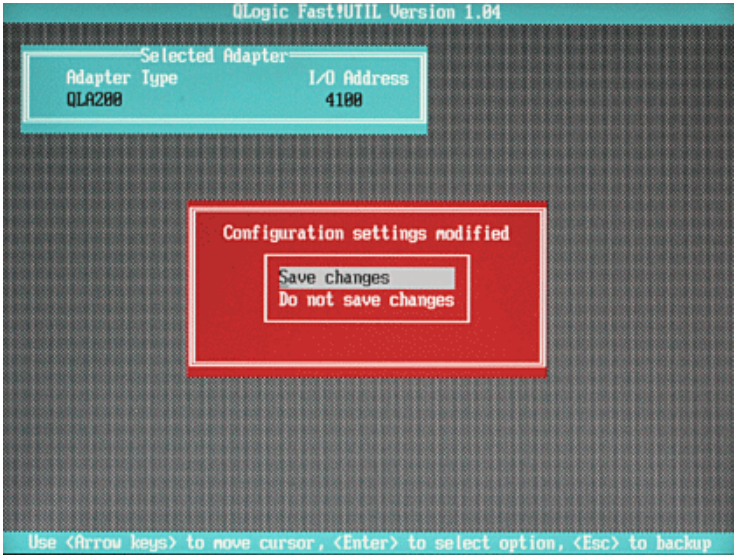

Figure 9-14 Restored adapter defaults

h. Select **Host Adapter** and press **Enter**.

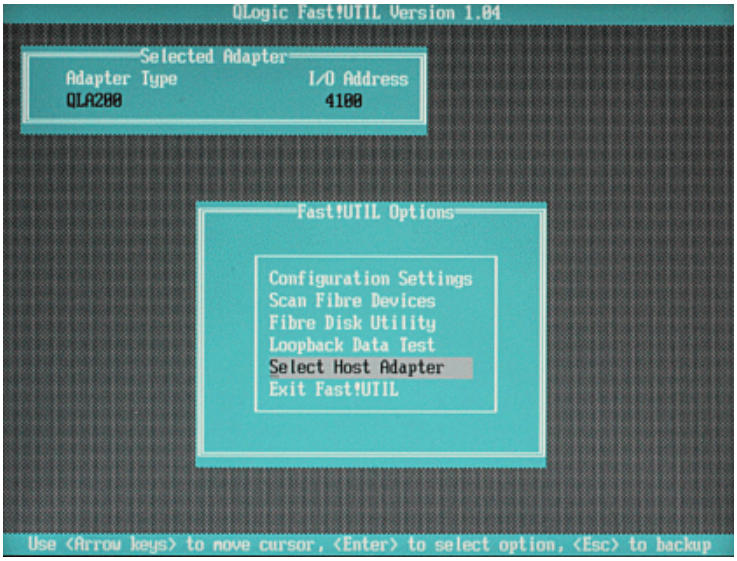

Figure 9-15 Fast!UTIL options — Select host adapter

i. Select the first HBA and press **Enter**.

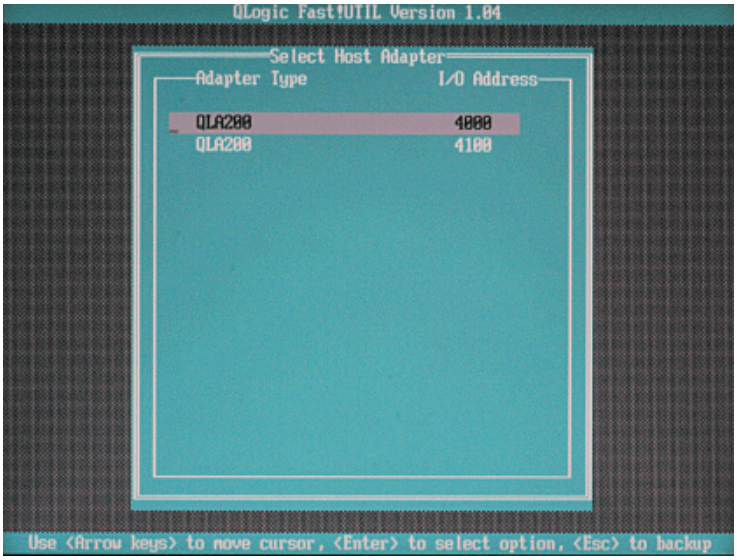

Figure 9-16 Select host adapters

j. Select **Configuration Settings** and press **Enter**.

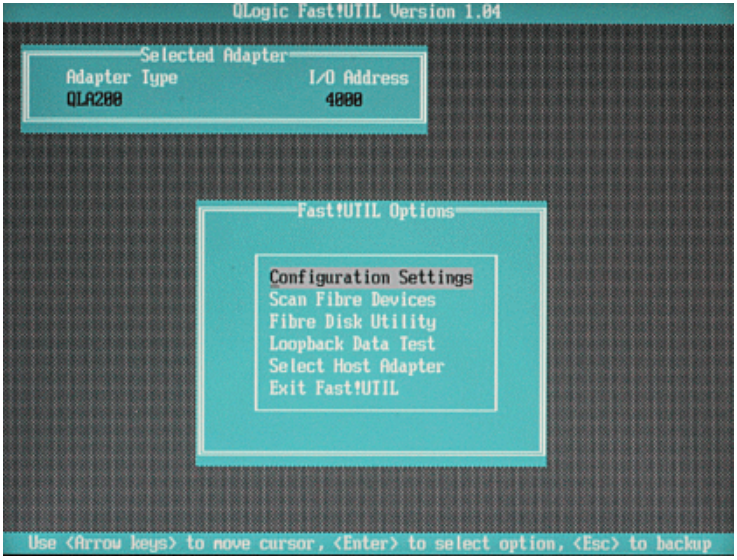

Figure 9-17 Fast!UTIL options — Configuration settings

k. Select **Restore Default Settings** and press **Enter**.

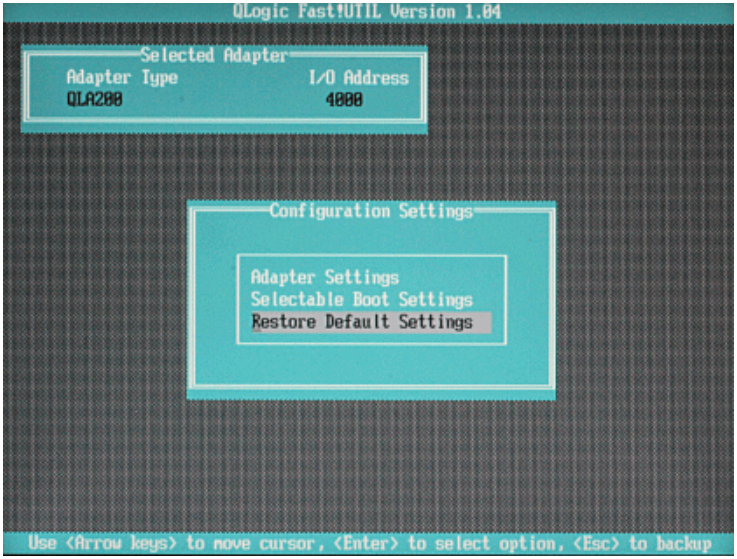

Figure 9-18 Configuration settings — Restore default settings

l. The default settings of the adapter will be applied. Press any **key** to return to **Configuration Settings**.

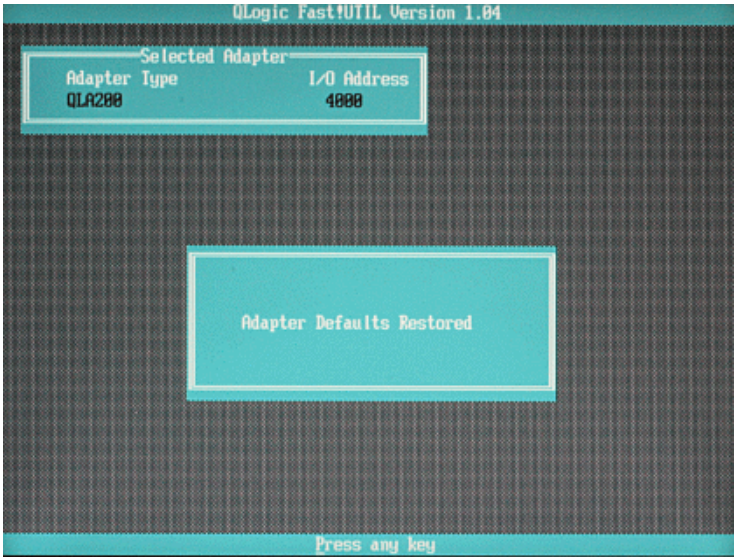

Figure 9-19 Restore default settings

3. Activate the adapter BIOS and define the boot settings now.

a. Select **Adapter Settings** and press **Enter**.

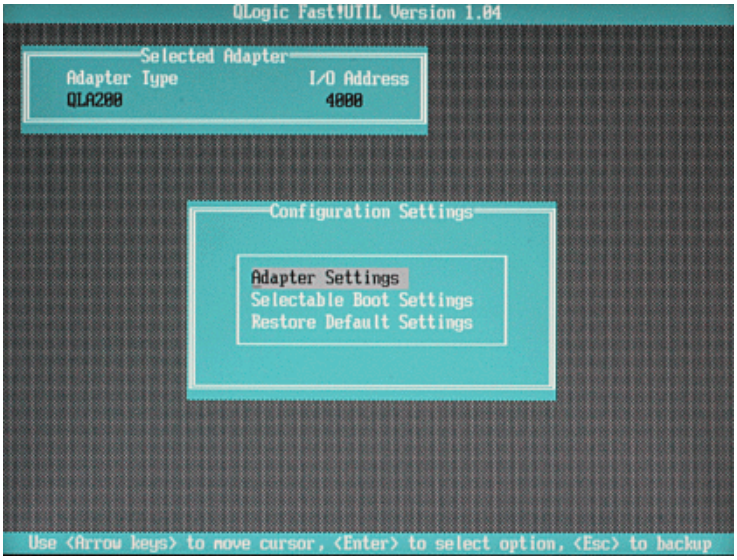

Figure 9-20 Configuration Settings — Adapter settings

b. Enable the host adapter BIOS by selecting **Host Adapter BIOS: disabled** and press **Enter**. The value will change from disabled to enabled.Then Press **ESC** to return to **Configuration Settings** dialog.

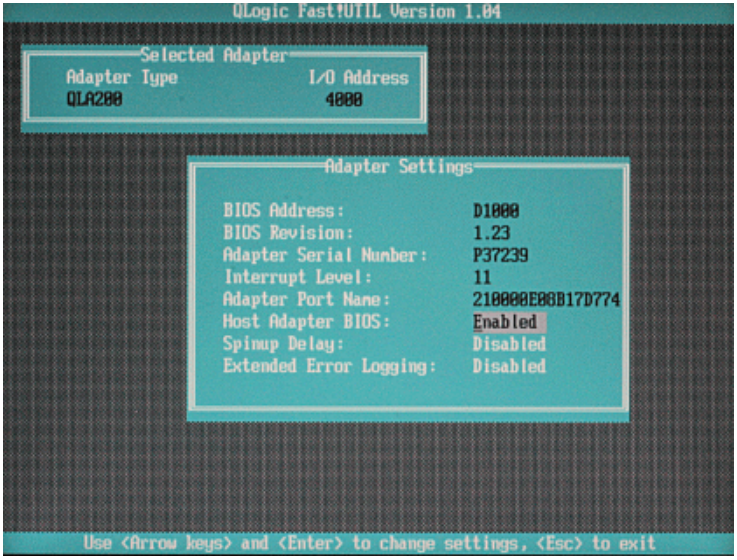

Figure 9-21 Enable host adapter BIOS

c. Select **Selectable Boot Settings** and press **Enter**.

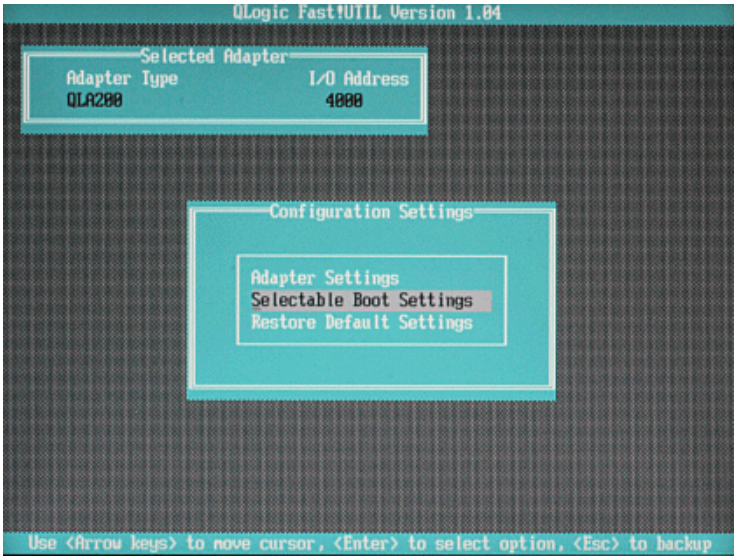

Figure 9-22 Configuration settings — Selectable boot settings

d. Select **Selectable Boot: Disabled** and press **Enter**. The value will change from Disabled to Enabled.

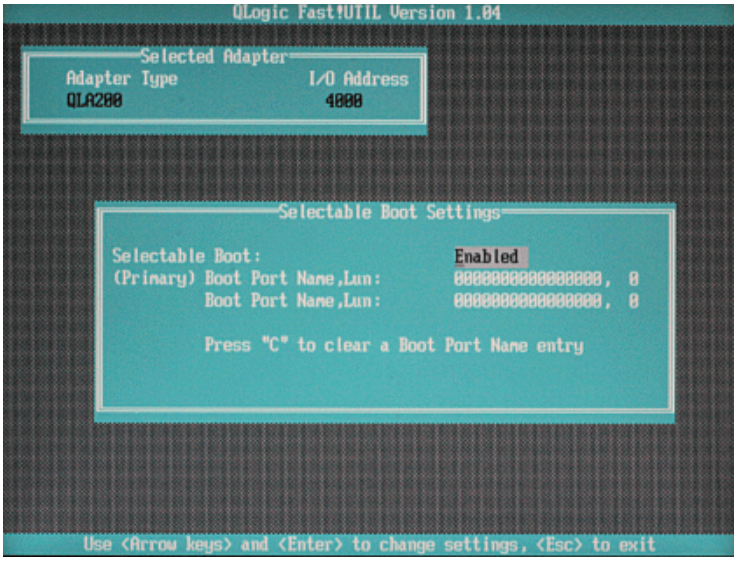

Figure 9-23 Enable selectable boot

e. Depending on the planning it is possible not to choose a specific target and LUN to boot from The adapter has the functionality to decide from which LUN0 the system will boot. As long as only one storage subsystem is attached to this system there are no problems. With multiple storage systems attached multiple different LUN0 may exist. This can lead to boot problems when a LUN0 from the wrong storage system is chosen to boot from. The next steps explain how to select s specific target and device to boot from.

f. Select (Primary) Boot Port Name and LUN and press **Enter**.

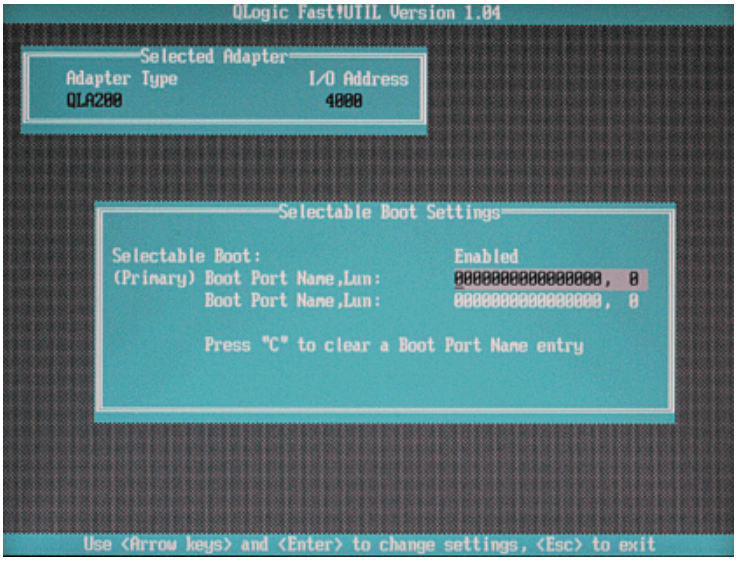

Figure 9-24 Primary Boot Port and LUN

g. Select the target which owns the logical drive to boot from. Than press **Enter**. In this case the logical drive is owned by controller A. The initiator is allowed only to access port FC0 of controller A. The third and fourth digit characterize the port on the Total Storage DS400. 00 represents FC0 on controller A, 01 represents FC1 on controller A, 02 represents FC0 on controller B and 03 represents FC1 on controller B. Other ports of the TotalStorage DS400 do not provide in the current configuration LUNs.

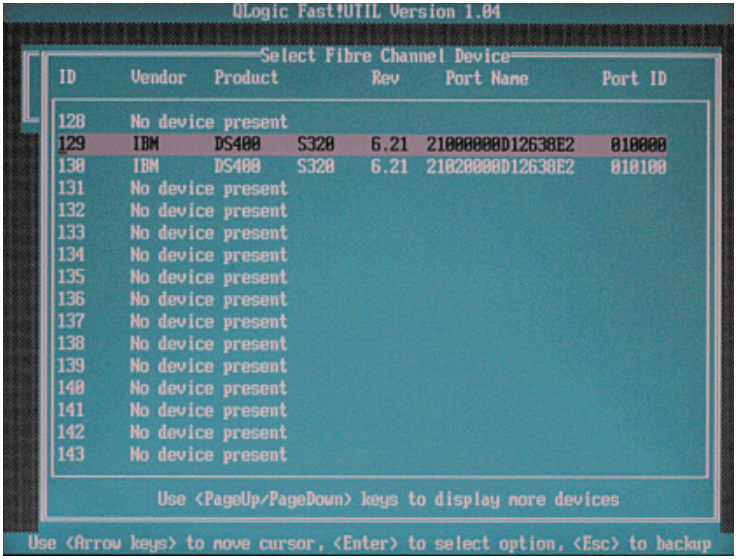

Figure 9-25 Select target

h. Select the LUN to boot from. Than press **Enter**. The logical drive to boot from was assigned to the initiator as LUN 0 on the TotalStorage DS400.

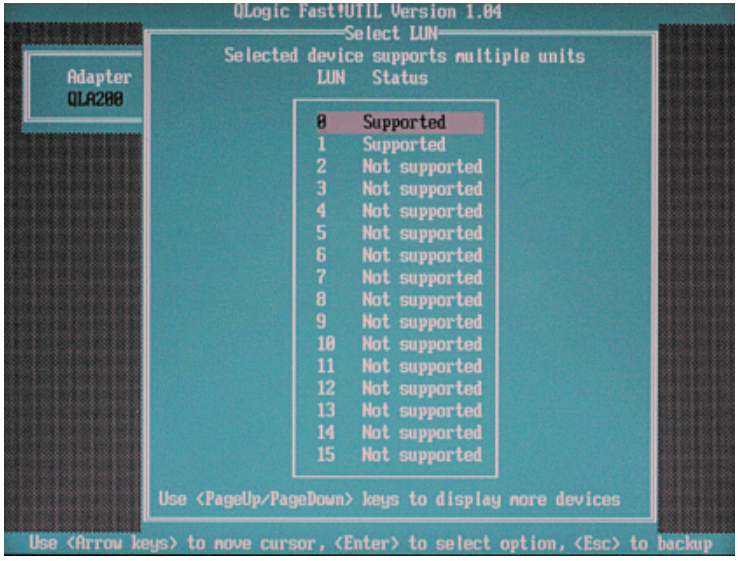

Figure 9-26 Select LUN to boot from

i. Press **ESC** to leave the **Selectable Boot Settings** and return to **Configuration Settings**.

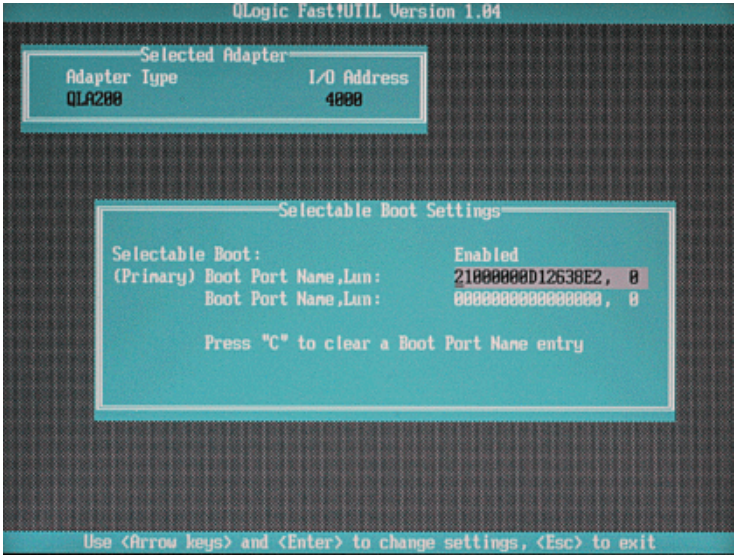

Figure 9-27 Selected boot target and LUN

j. Press **ESC** to leave **Configuration Settings**.

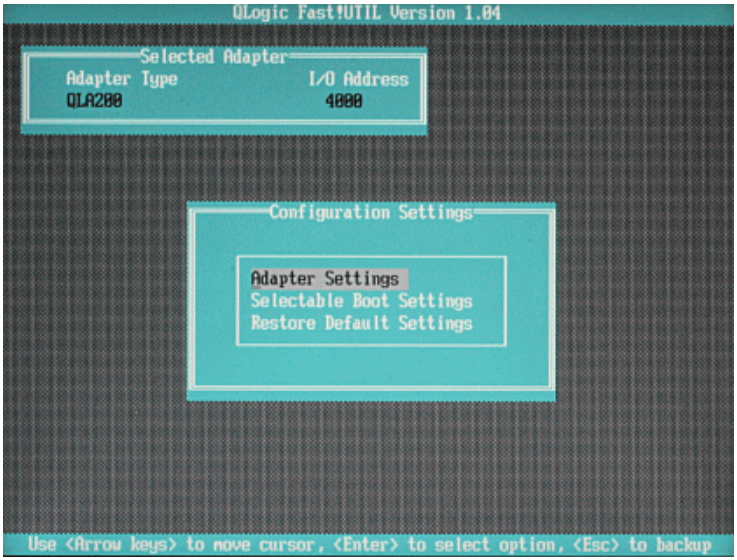

Figure 9-28 Configuration Settings

k. Select **Save changes** and press Enter. This will save all modified settings.

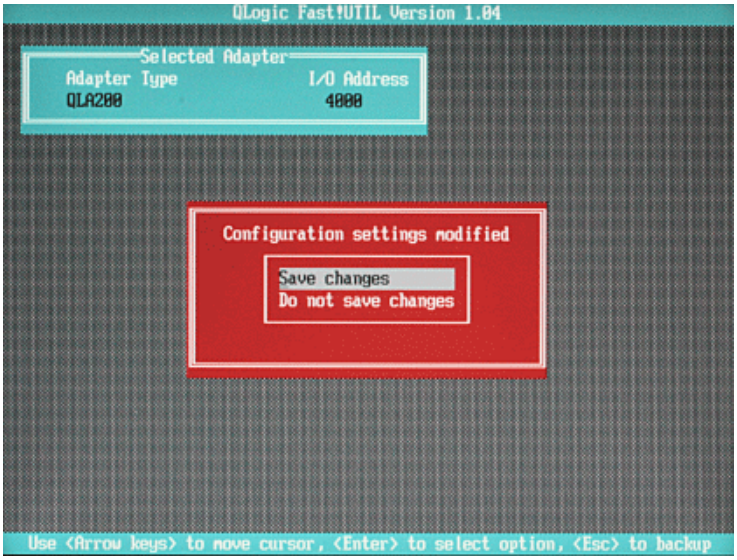

Figure 9-29 Save modifications

l. Select **Exit Fast!UTIL** and press **Enter**.

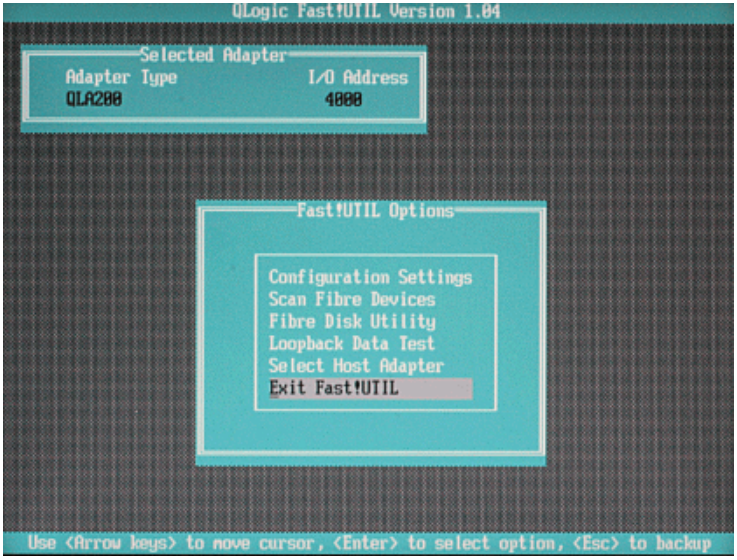

Figure 9-30 Fast!UTIL options — Exit Fast!UTIL

4. Do not reboot the system at this time. It will be rebooted in a later step.

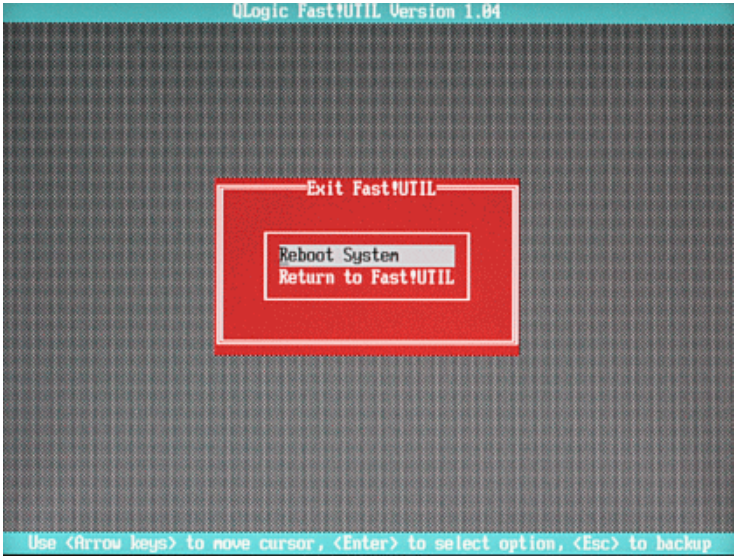

Figure 9-31 Reboot system

The HBAs are now ready for the installation of Microsoft Windows Server 2003 Enterprise Edition.

9.1.5 HBA driver disk creation

After the IBM SMB HBA is configured to use a LUN as boot device, create a driver diskette for the installation of Microsoft Windows Server 2003. To create the driver diskette, download the driver for the IBM SMB Host Bus Adapter. Use the software and driver matrix of the DS400 to get the right driver for the host bus adapter in combination with the TotalStorage DS400.

Note: IBM DS4000 Host Bus Adapter and IBM SMB Host Bus Adapter use separate drivers, but they use the same driver set. BIOS and firmware are different.

Unzip the driver archive and store all files in the root of a 3.5” diskette. This diskette will be required to install Microsoft Windows Server 2003.

Proceed with the installation of Microsoft Windows Server 2003 as described in the next section.

9.1.6 Microsoft Windows Server 2003 Enterprise Edition installation

The installation of Windows 2003 is straightforward. The driver of the IBM SMB Host Bus Adapter is not included on the Microsoft Windows Server 2003 Enterprise Edition CD. This requires you to create a driver diskette first. See 9.1.5, “HBA driver disk creation” on page 468 for information the creation of the driver diskette.

1. The system is still in the host bus adapter BIOS setup. Select **Reboot System** and press **Enter**.

Figure 9-32 Reboot system

2. Insert the installation CD into the CD-ROM drive and boot from this CD. After the message Setup is inspecting your computer was shown the black screen will switch to blue and the bottom line give you the option to press F6 to install third party SCSI or RAID drives which are required at boot time. Press **F6** to install the driver for the IBM SMB Host Bus Adapter.

Figure 9-33 Windows setup — Install third party SCSI or RAID drivers

3. Insert the created driver disk and press **Enter**.

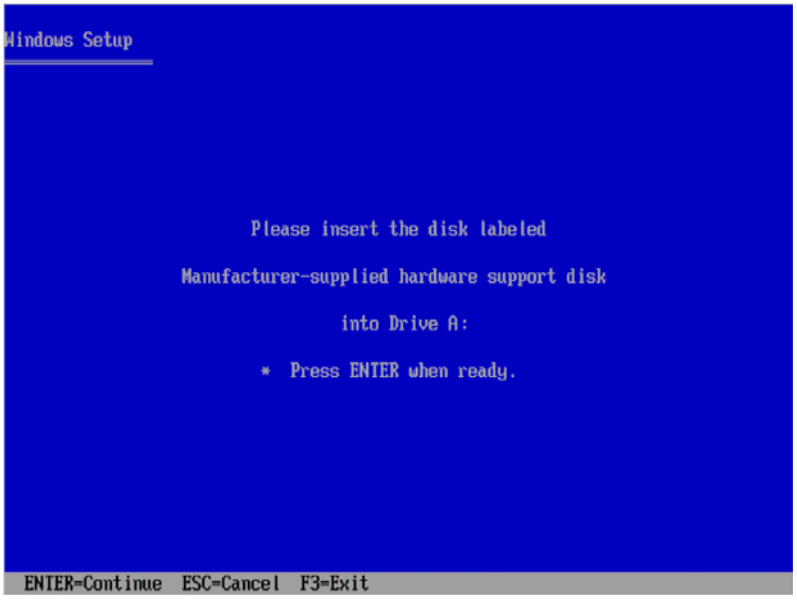

Figure 9-34 Windows setup — insert driver disk

4. The driver for the QLogic Fibre Channel Adapter will be found. Press **Enter** to select this driver.

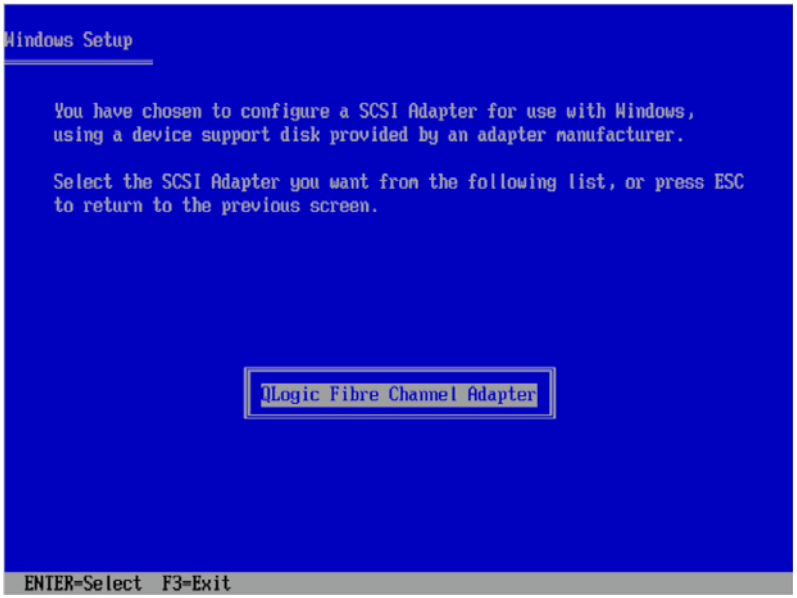

Figure 9-35 Windows setup — QLogic Fibre Channel Adapter

5. The driver is newer that the driver included on the CD. Press **S** to use the newer driver from the disk.

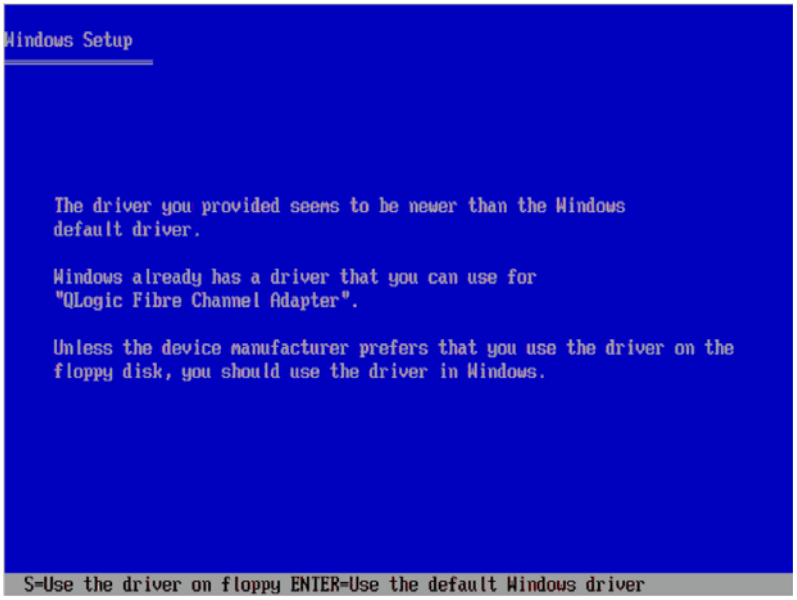

Figure 9-36 Windows setup — newer driver found

6. The driver for the IBM SMB Host Bus Adapter will be loaded and Windows setup returns to the screen to add more third party SCSI and RAID driver. Press Enter to proceed with the installation. For the remaining installation follow the instructions on the screen. There is now difference from the installation to a local attached drive.

Figure 9-37 Windows setup — Supported mass storage drivers

9.1.7 Driver and service pack installation and finishing the setup

Install the Broadcom network driver for the adapter. Broadcom NetXtreme Gigabit Ethernet #2 is connected to the port number 1. This port is also shared by the baseboard management controller (BMC). This port might be connected into a management network. Broadcom NetXtreme Gigabit Ethernet is routed to port number 2. This is the port close to the power supplies.

There are operating system install instructions for all IBM @serverxSeries server on the support pages. See Appendix C, “How to get updates of code and documentation” on page 623.

1. Install the service pack1 for Microsoft Windows Server 2003. Reboot the server.

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

2. Install remaining drivers not available during operating system installation:
   a. Install the Intel chipset drivers.
   b. Install ActivePCI.
   c. All unknown devices are now removed.
   d. Install IPMI driver. This driver is used for the BMC.

3. Install the IBM FAStT Management Suite Java (MSJ) to manage the IBM SMB Host Bus Adapter.

   a. Download the latest version of the FAStT MSJ from the IBM support Web page and run the installer file.

      http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56707

      Click **Run** to execute the installer.

Figure 9-38 Open file security warning

   b. Click Next.

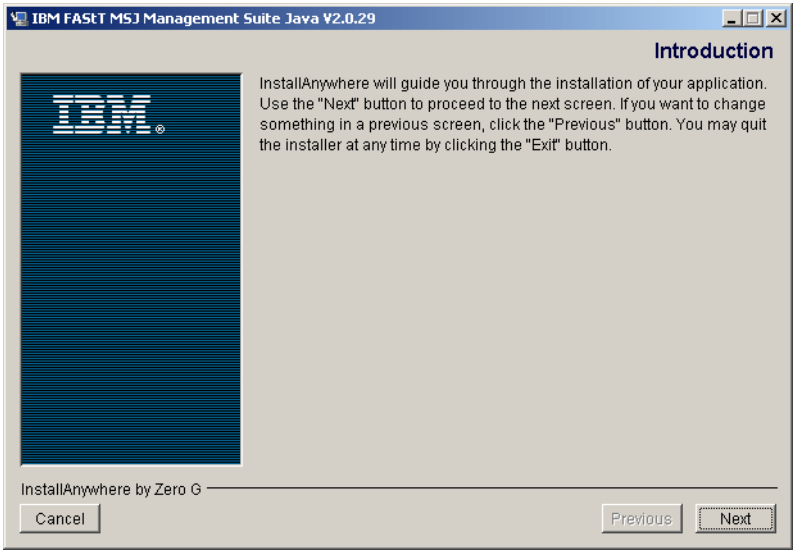

Figure 9-39 MSJ installation — Introduction

   c. Accept the license agreement and click **Next**.

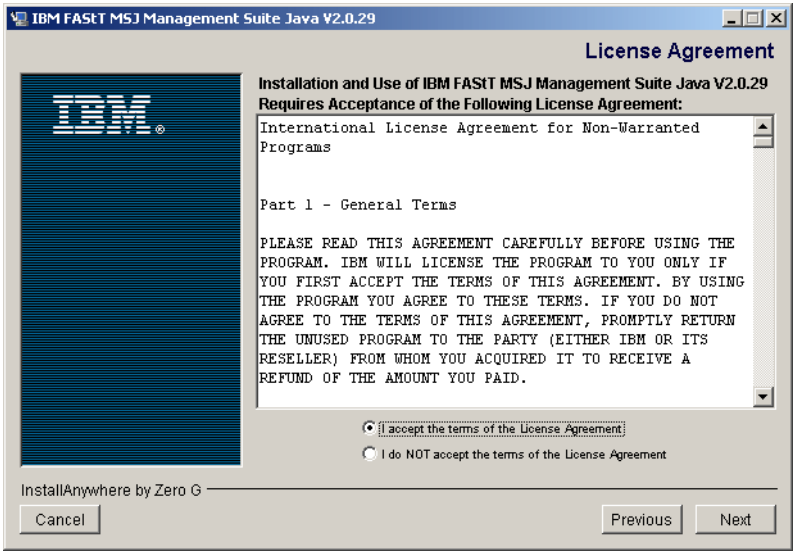

Figure 9-40 MSJ installation — license agreement

d. Read important information and click Next.

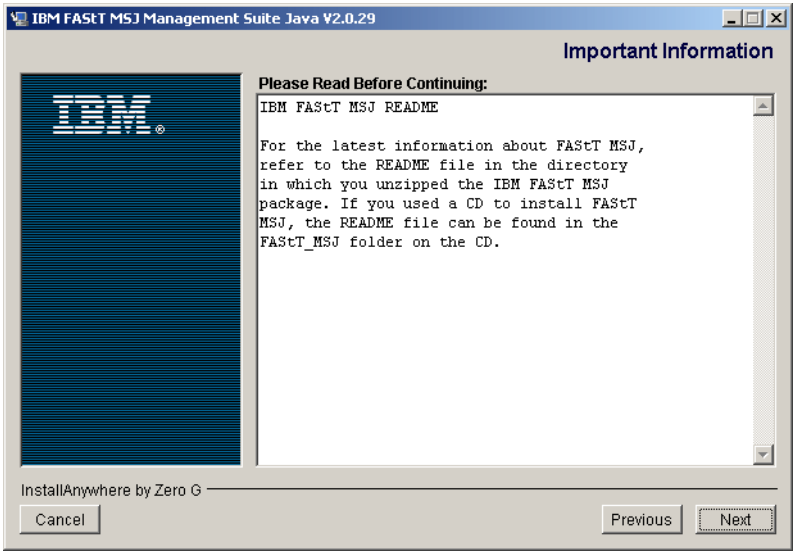

Figure 9-41 MSJ installation — important information

e. Select **GUI and NT agent** and click **Next**. For a single system the MSJ GUI and agent should be installed. The agent must be installed on the

system with the host bus adapters. The GUI is required only once in a network but can be installed on multiple systems.

Figure 9-42 MSJ installation — Choose product features

f. Specify the installation directory and click **Next**.

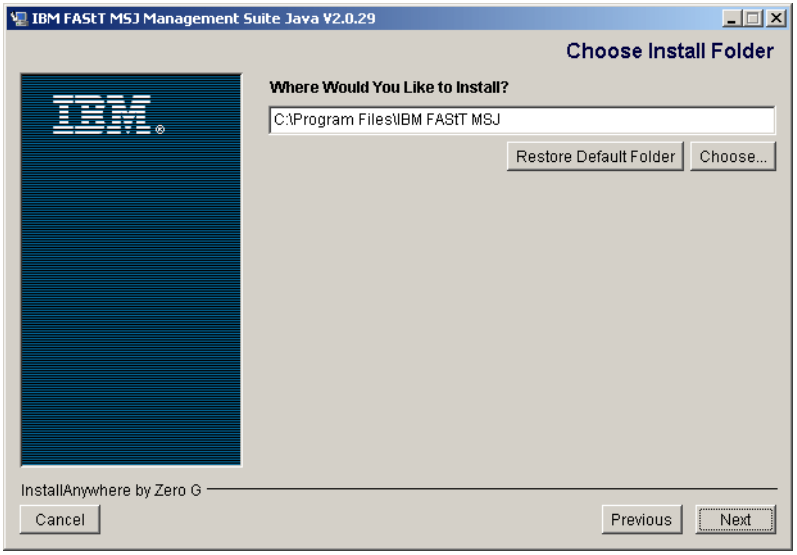

Figure 9-43 MSJ installation — Choose install directory

g. Specify the location of the shortcuts and click **Next**.

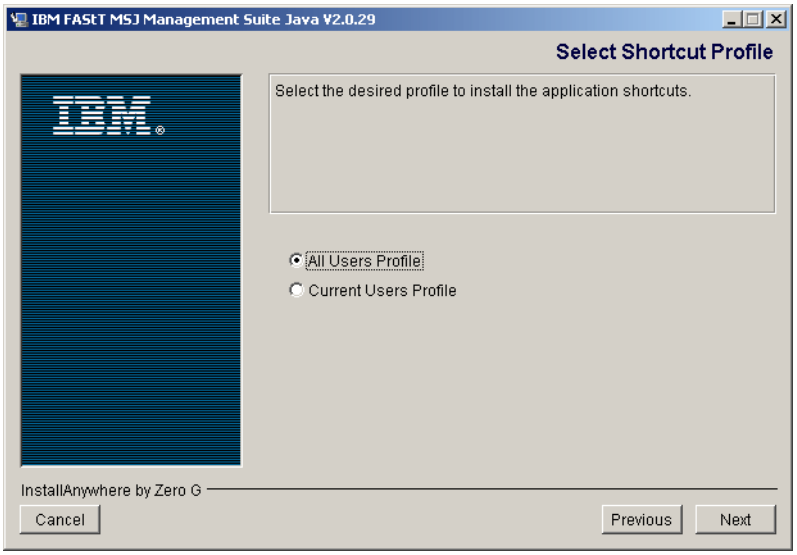

Figure 9-44 MSJ installation — select shortcut profile

h. Choose if a desktop icon is required and click **Install**.

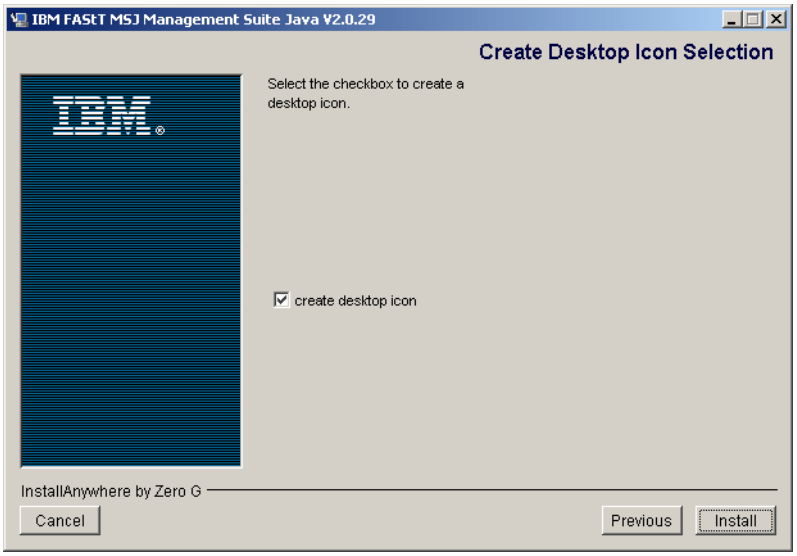

Figure 9-45 MSJ installation — create desktop icon

i. Files will be copied. When installation is complete click **Done**.

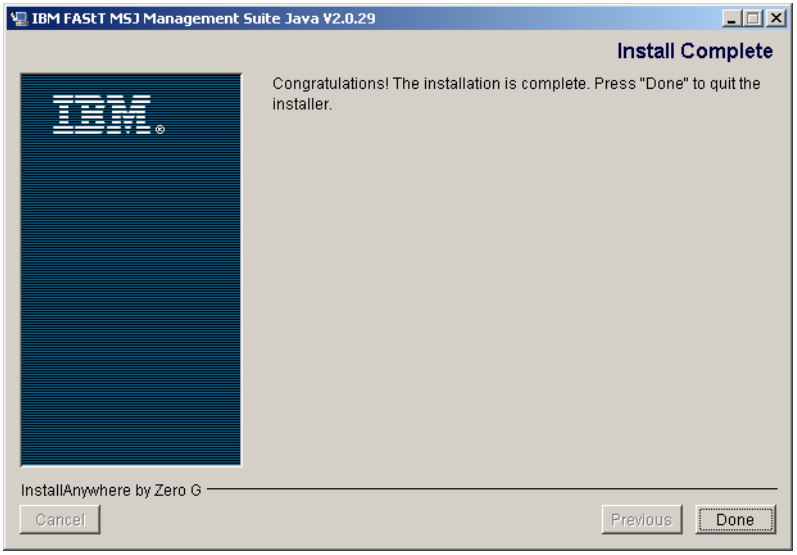

Figure 9-46 MSJ installation — Done

4. Install the QLogic Device Specific Module (DSM) with the multipath input/output (MPIO) driver.

   QLogic provides a special Device Specific Module for Microsoft multipath input/output driver. This DSM is specifically designed to work with the IBM TotalStorage DS400. See 3.4.4, “DS400 failover” on page 51 for failover details.

   For more information about MPIO, see Microsoft whitepaper “Highly Available Storage: Multipathing and the Microsoft MPIO Driver Architecture”:

   http://www.microsoft.com/windowsserver2003/technologies/storage/mpio/default.mspx

   Frequently Asked Questions: Microsoft Multipath I/O:

   http://www.microsoft.com/WindowsServer2003/technologies/storage/mpio/faq.mspx

   To install the QLogic Device Specific Module together with the Multipath IO driver perform the following steps.

   a. Download and extract the driver archive. See DS400 software and driver matrix for the correct versions of the HBA driver and the DSM and MPIO.

      http://www.ibm.com/support/docview.wss?rs=1110&uid=psg1MIGR-58619

   b. Run the install.bat file included in the driver archive.

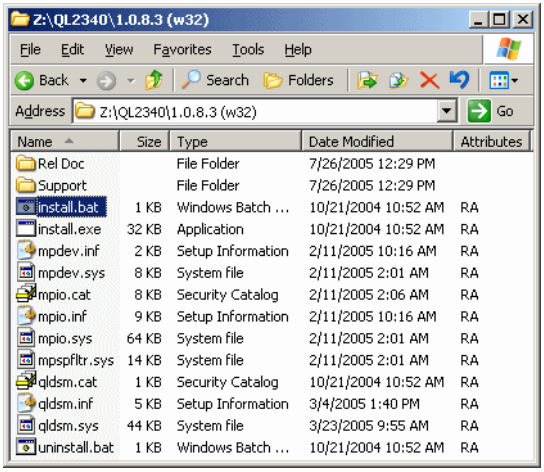

Figure 9-47 Content of the DSM and MPI O driver archive

c. Confirm the open file security warning by clicking **Run**.

Figure 9-48 Open file security warning

d. A command line box opens and displays the progress of the drive installation.

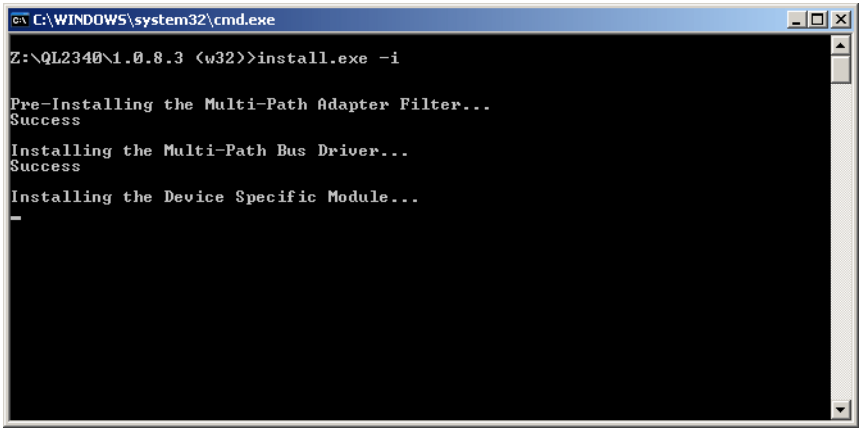

Figure 9-49 Driver installation output

e. The DSM driver is currently not yet digitally signed. Confirm the warning to install the unsigned driver.

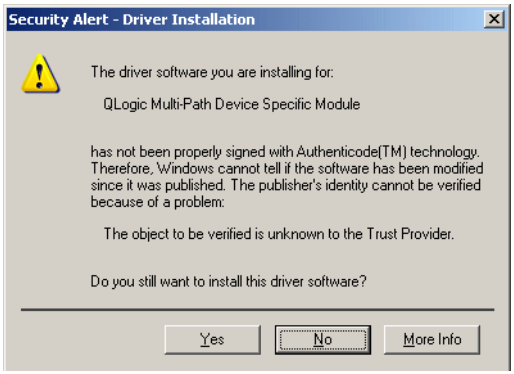

Figure 9-50 Unsigned driver warning

f. After the DSM module was installed the command line window will close. In the device manager a new device in the **SCSI and raid controller** section is displayed.

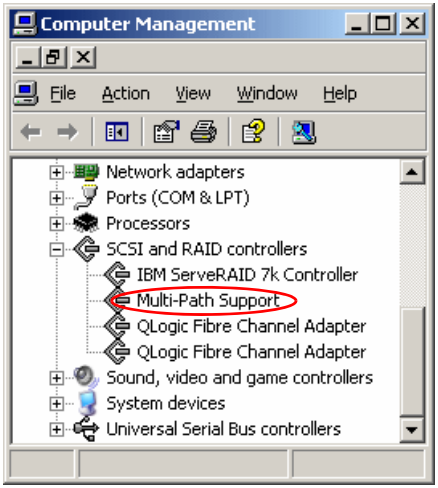

Figure 9-51 Device Manager after installing DSM and MPIO

5. Reboot the server.

9.1.8 Enable initiator access to all TotalStorage DS400 ports

Microsoft Windows Server 2003 is installed along with the multipath drivers. There are multiple paths available but no LUNs are accessible here. This step will activate the ports of the TotalStorage DS400 for the initiators to gain access over all ports. The access was restricted in 9.1.3, “TotalStorage DS400 setup” on page 451.

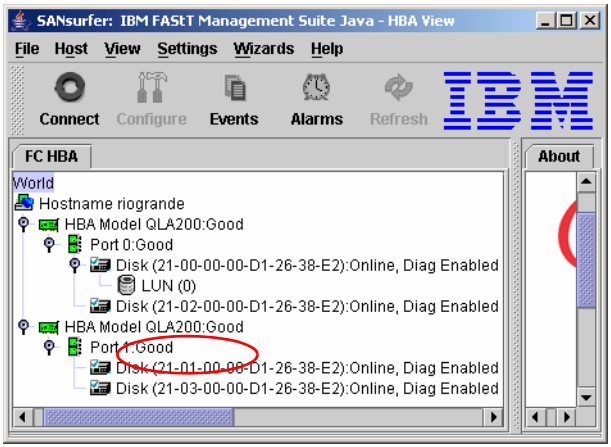

Figure 9-52 Assigned LUNs to the HBAs

1. Open a Telnet connection to one of the controllers in the DS400 and logon as administrator. Enter **Telnet** 192.168.1.226. To logon with administrative rights enter **administrator**.

```
----[ IBM DS400: Command Line Interface ]---------------------------------

               Copyright:  IBM (c)2003-2004

                Firmware:  IBM TotalStorage DS Series (Build 0743)
                           [20/07/2005][6.21]

  Peer controller status:  Ready


           Typing 'help':  provides an overview of the commands available to
                           configure the controller.

                 TAB key:  use the TAB key as a means of getting assistance
                           in completing any commands that you issue.


DS400[A]> administrator
Administrator level password: ********
DS400[A]#
```

Figure 9-53 Logon

2. Get a list of the available initiators by entering **initiator list**.

```
DS400[A]# initiator list

Fibre Channel Initiators

      210000e08b0f1549        [WOLGA-SLOT4-(1)]
      210000e08b0ff748        [WOLGA-SLOT3-(0)]
      210000e08b170d75        [RIOGRANDE-S4-(1)]
      210000e08b17d774        [RIOGRANDE-S3-(0)]

DS400[A]#
```

Figure 9-54 Initiator list

3. Look for the initiators of RIOGRANDE and check port access for every singly initiator by entering **initiator manage** RIOGRANDE-S4-(1) **info.** All ports are disabled for this initiator.

```
DS400[A]# initiator manage RIOGRANDE-S4-(1) info

----[ Initiators ]-----------------------------------------------------------


----[ Initiator 'RIOGRANDE-S4-(1)' ]-----------------------------------------

                 WWN:  210000e08b170d75
               Alias:  RIOGRANDE-S4-(1)
                  ID:  3
         Port access:
                       Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------------
                       NO             NO             NO             NO


        LUN Mappings:
                 LUN                 logical         [Capacity]
-------------------------------------------------------------------------------

                 0                   RIOGRANDE-BOOT  [ 10.00GB]

DS400[A]#
```

Figure 9-55 Initiator information

4. Enable the ports by entering the command **initiator manage** RIOGRANDE-S4-(1) **allow all**. This will activate the ports for this initiator. The port access will be set to yes. Verify this with the command entered in step 3 on page 482.

```
DS400[A]# initiator manage RIOGRANDE-S4-(1) allow all
Access allowed to initiator RIOGRANDE-S4-(1) on all ports

DS400[A]#
```

Figure 9-56 Enable initiator access for all ports

5. The other initiator of the system RIOGRANDE had one port enabled to install the operating system onto a logical drive from a DS400. The remaining three ports must be enabled as well. Check the ports with **initiator manage** RIOGRANDE-S3-(0) **info** and enable the remaining ports with **initiator manage** RIOGRANDE-S3-(0) **allow A1 B0 B1**.

```
DS400[A]# initiator manage RIOGRANDE-S3-(0) info

----[ Initiators ]-----------------------------------------------------


----[ Initiator 'RIOGRANDE-S3-(0)' ]-----------------------------------

                WWN:  210000e08b17d774
              Alias:  RIOGRANDE-S3-(0)
                 ID:  4
        Port access:
                      Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------
                      YES            NO             NO             NO


       LUN Mappings:
                LUN                 logical         [Capacity]
-------------------------------------------------------------------------

                0                   RIOGRANDE-BOOT  [ 10.00GB]

DS400[A]# initiator manage  RIOGRANDE-S3-(0) allow A1 B0 B1
Access allowed to initiator RIOGRANDE-S3-(0) on ports A1 B0 and B1

DS400[A]#
```

Figure 9-57 Initiator information

6. Reboot the server and verify after reboot that you see 4 drives in the device manager and one multipath device.

   IBM FAStT MSJ shows 4 targets and every target has one LUN. This is correct.

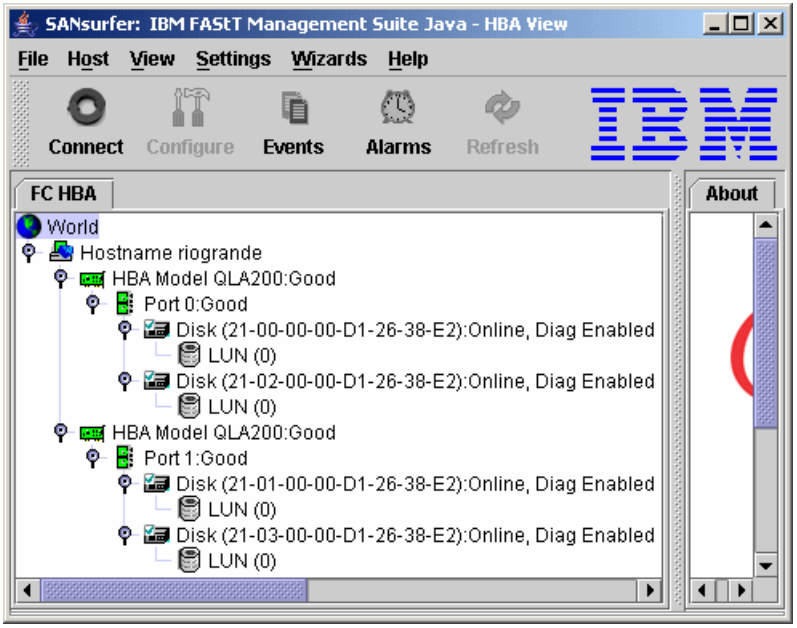

Figure 9-58 FAStT MSJ with 4 LUNs.

Windows device manager shows now 4 drives and one multipath device.

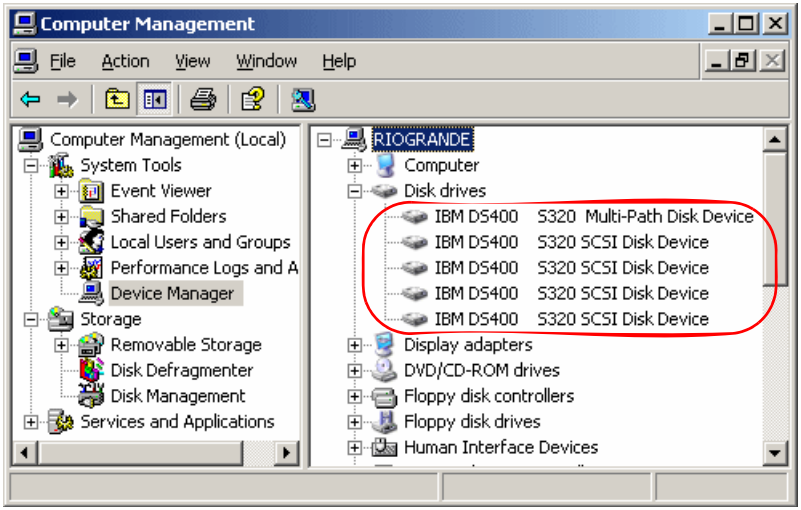

Figure 9-59 Device manager with 4 drives and one multipath device

Target and LUN ID can be retrieved from the disk device property. Those values are also found in the properties of the multipath device.

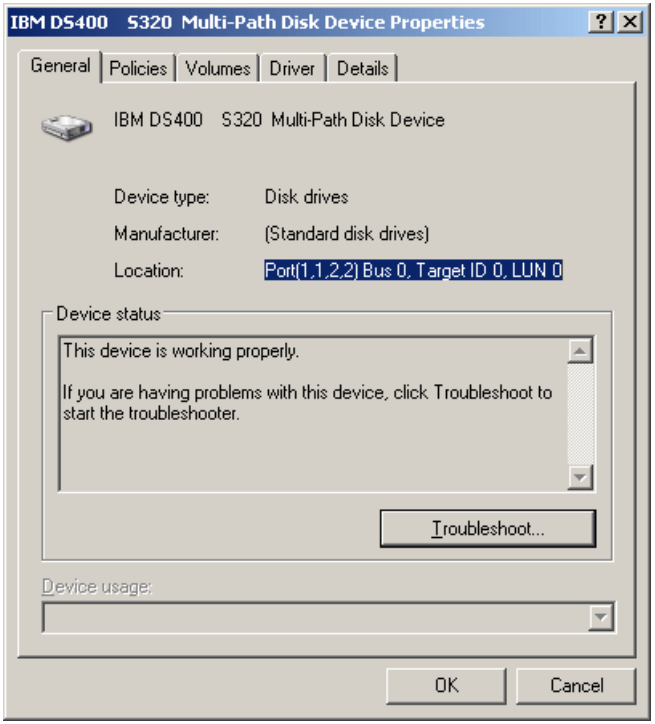

Figure 9-60 Multipath disk device properties

9.1.9 Finish IBM SMB Host Bus Adapter BIOS setup

During the planning phase when it is decided where the boot targets need to be specified, the following steps are required to complete the setup. If no boot targets need to be specified in the host bus adapter BIOS, only the adapter BIOS and the selectable boot need to be enabled. In any case, the HBA must be enabled in the system BIOS.

The initial setup of the host bus adapters has only one boot device initially defined. Now, the remaining three devices will be defined. This is required to boot the server no matter which controller owns the array of the logical drive or when one HBA fails.

Reboot the server and enter the BIOS setup of the host bus adapter by pressing **CTRL+Q**.

During the initial setup of the host bus adapter, only one boot device was defined. To insure that the system boot can take place no matter which controller owns the array of the boot logical drive, the alternate controller must be defined as a

boot target. The appropriate LUN must be selected as well. This setting needs to be defined also for the second host bus adapter to allow booting the system in case the first HBA fails. Not under all circumstances will this be the case. There is no failover driver available, therefore you have to configure the second adapter if there is a failure.

1. Reboot the system. During POST press **CTRL+Q** when the adapter displays the message “Press <CTRL+Q> for Fast!UTIL” to enter the host bus adapter BIOS setup.

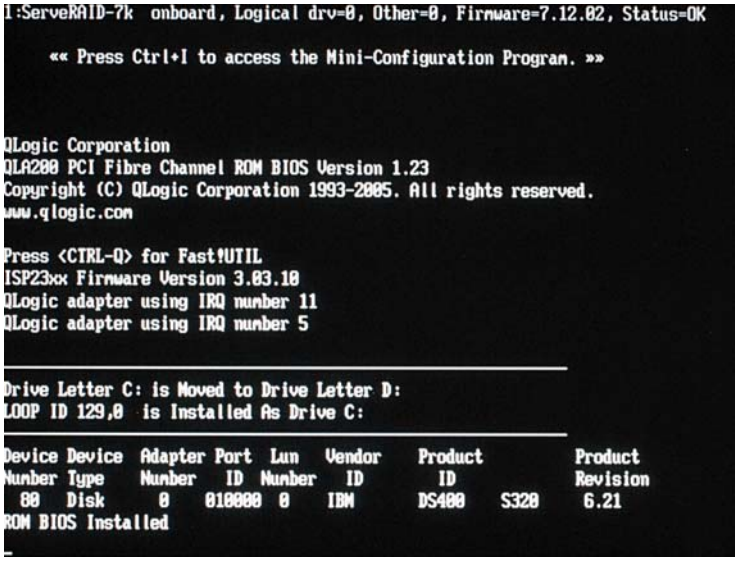

Figure 9-61 POST messages of the host bus adapter

2. Select the first HBA and press **Enter**. See Figure 9-62 on page 488.

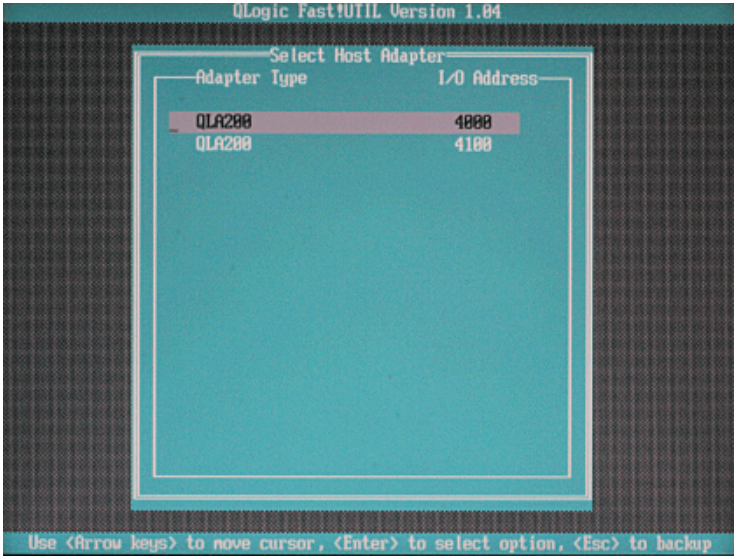

Figure 9-62 Select host adapter

3. Select **Configuration Settings** and press **Enter**.

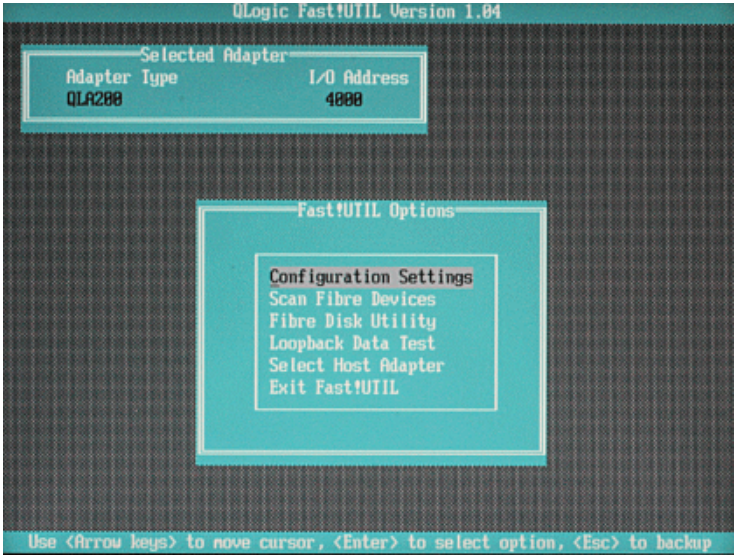

Figure 9-63 Fast!UTIL options — Configuration settings

4. Select **Selectable Boot Settings** as shown in Figure 9-64 on page 489 and press **Enter**.

Figure 9-64 Configuration settings — Selectable boot settings

5. Select the **Boot Port name, LUN** after the primary entry and press **Enter**.

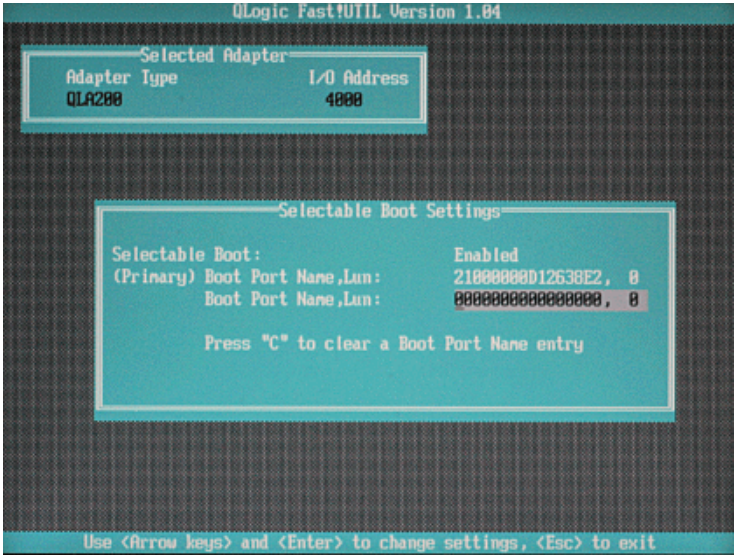

Figure 9-65 Second boot port name and LUN

6. Select the FC0 interface of controller B. From the world wide name the interface can be detected. Digit 3 and 4 of the world wide name is 02.

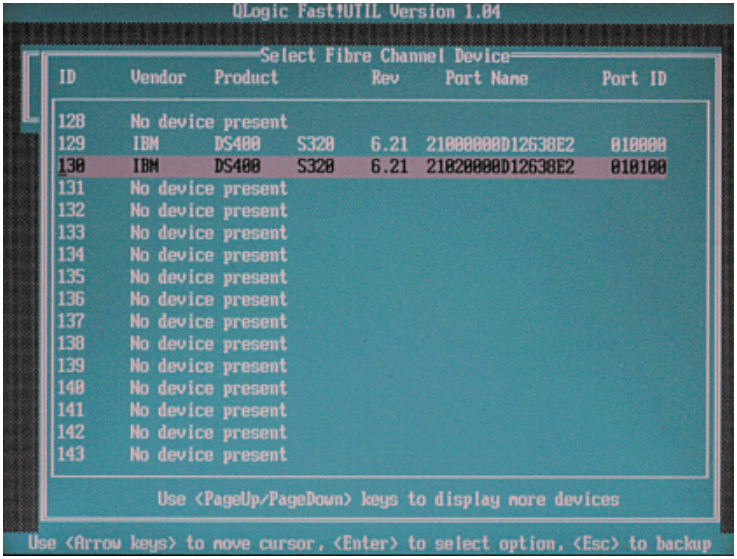

Figure 9-66 Select boot target

7. A list of all LUNs of this target is presented. Select LUN 0 and press **Enter**.

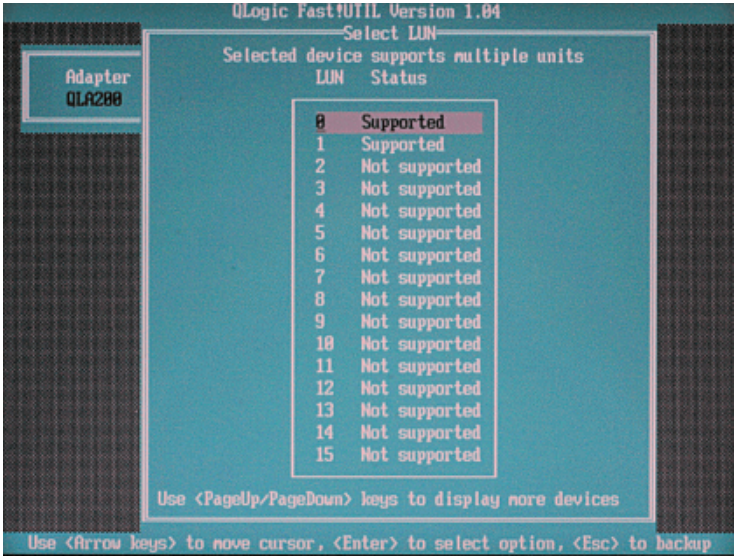

Figure 9-67 Select boot LUN

8. The target and the LUN is added to the boot device list. Press **ESC** to return to **Adapter Settings**.

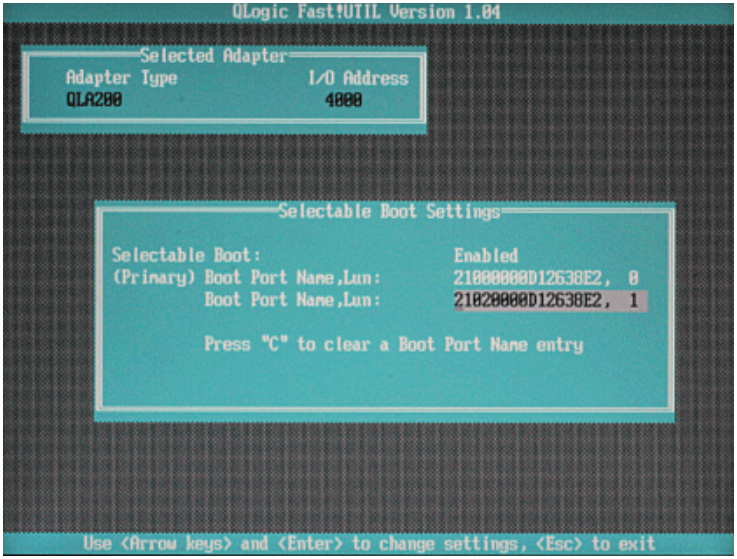

Figure 9-68 Defined boot targets and LUNs

9. Press **ESC** to leave **Adapter Settings** and return to **Configuration Settings**.

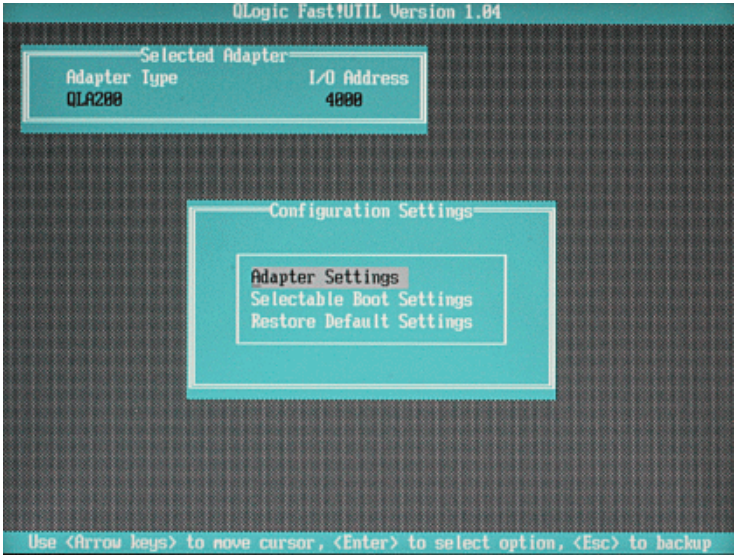

Figure 9-69 Configuration settings

10. Save the modifications by select **Save changes** and press **Enter**.

Figure 9-70 Save modifications

11. Select **Host Adapter** and press **Enter**.

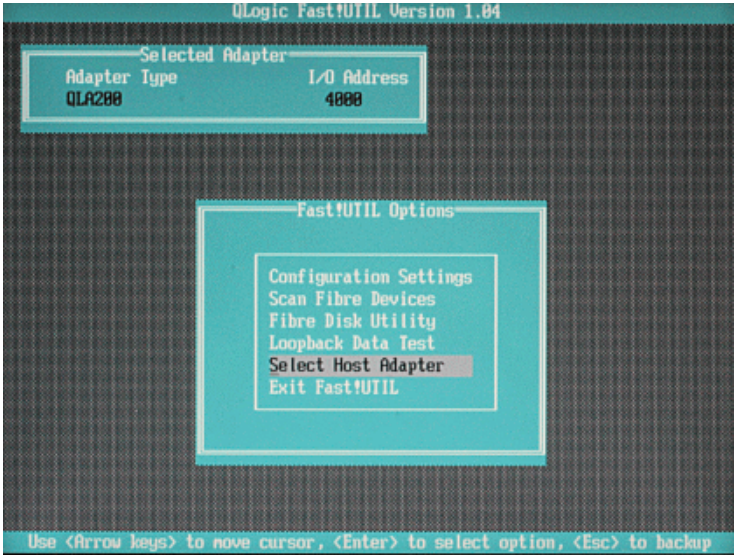

Figure 9-71 Fast!UTIL options — Select host adapter

12. Select the second HBA and press **Enter**.

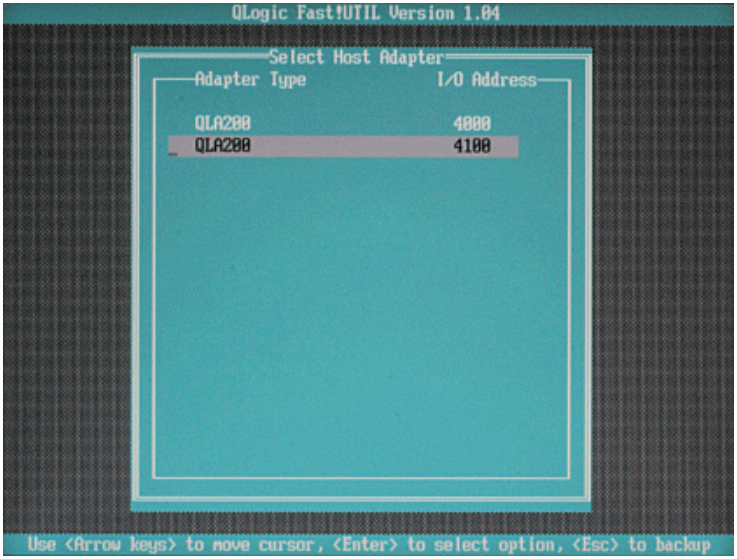

Figure 9-72 Select host adapter

13.. Select **Configuration Settings** and press **Enter**.

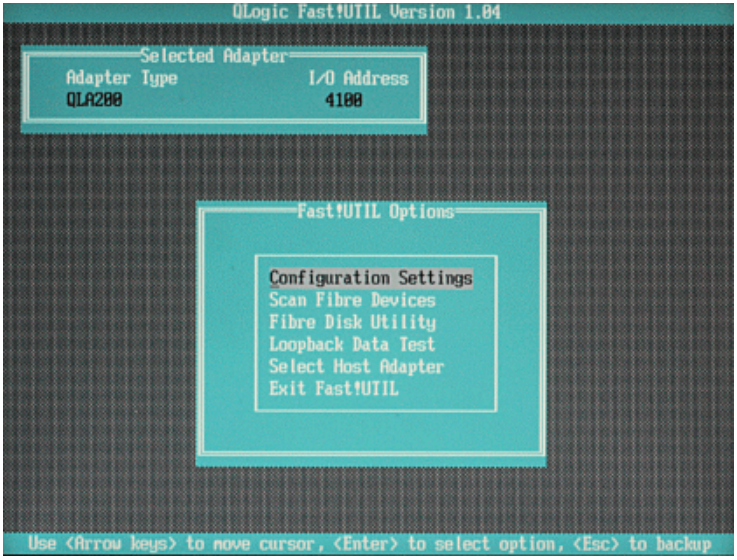

Figure 9-73 Fast!UTIL options — Configuration settings

14.Select **Adapter Settings** and press **Enter**.

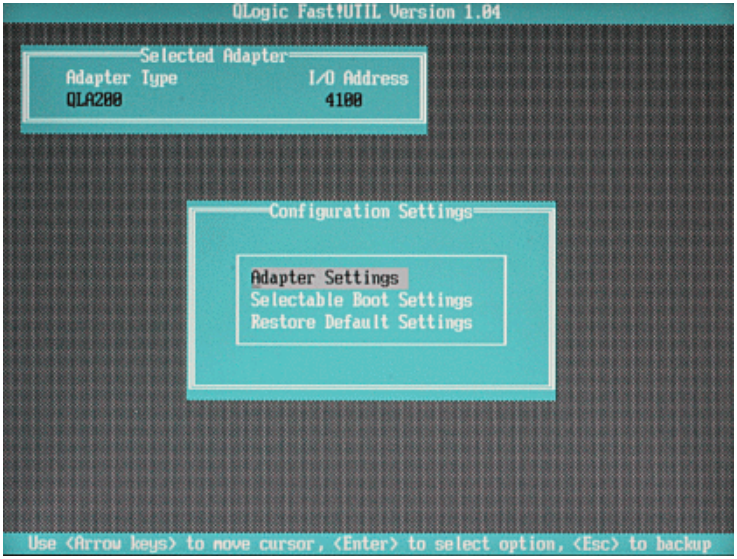

Figure 9-74 Configuration settings — Adapter settings

15. Select **Host Adapter BIOS disabled** and press Enter. The value will change from Disabled to Enabled. Then Press **ESC** o return to **Configuration Settings**.

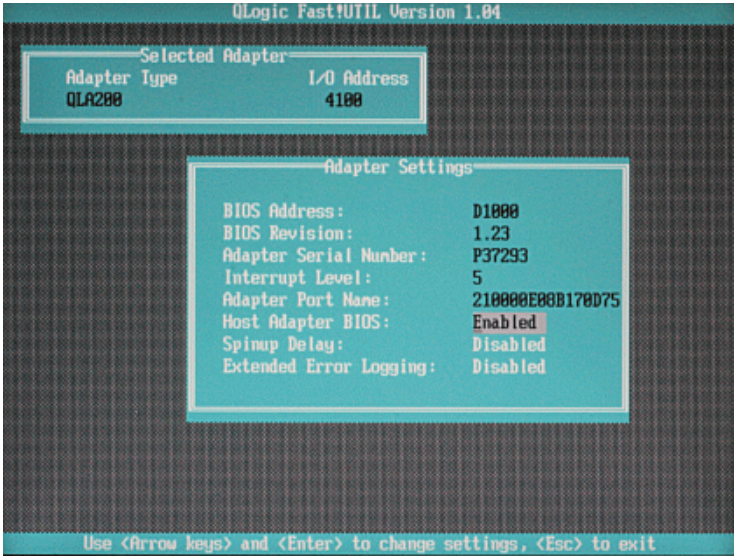

Figure 9-75 Enable host adapter BIOS

16. Select **Selectable Boot Settings** and press **Enter**.

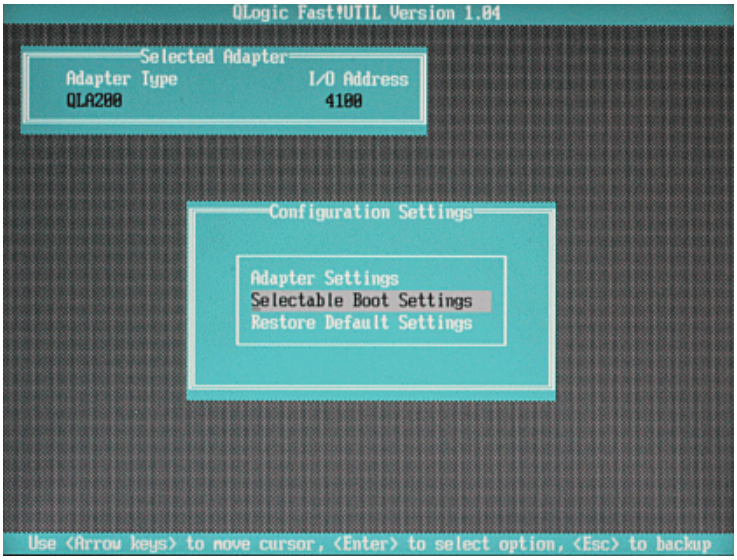

Figure 9-76 Configuration settings — Selectable boot settings

17. Select **Selectable Boot** and press **Enter**. The value changes from Disabled to Enabled.

Figure 9-77 Enable selectable boot

18. Select the **(Primary) Boot Port name, LUN** and press **Enter**.

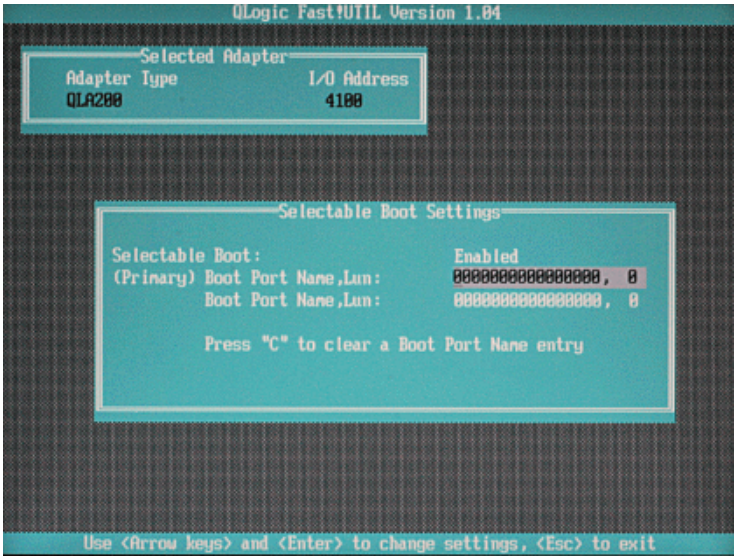

Figure 9-78 primary boot port name and LUN

19. Select the FC1 interface of controller A. From the world wide name the interface can be detected. Digit 3 and 4 of the world wide name is 01.

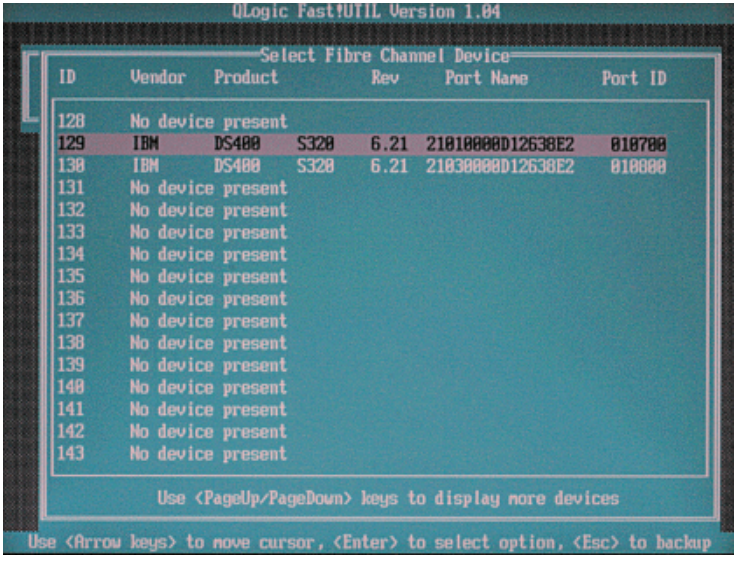

Figure 9-79 Boot target FC1 (01) controller A

20. A list of all luns of this target presented. Select LUN 0 and press Enter.

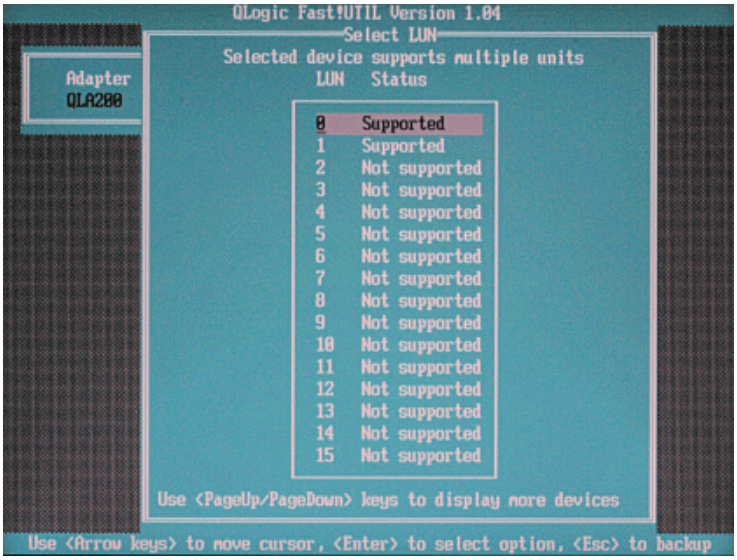

Figure 9-80 Boot LUNs

21. Select the **Boot Port name, LUN** after the primary entry and press **Enter**.

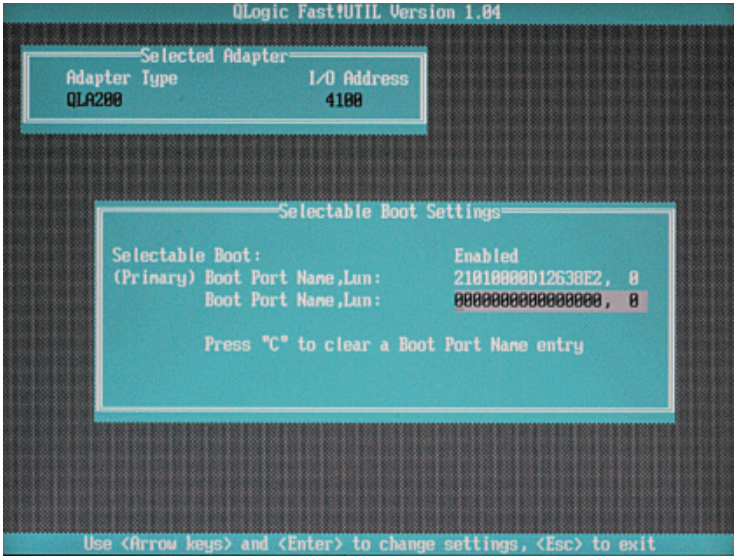

Figure 9-81 Second boot port name and LUN

22. Select the FC1 interface of controller B. From the world wide name the interface can be detected. Digit 3 and 4 of the world wide name is 03.

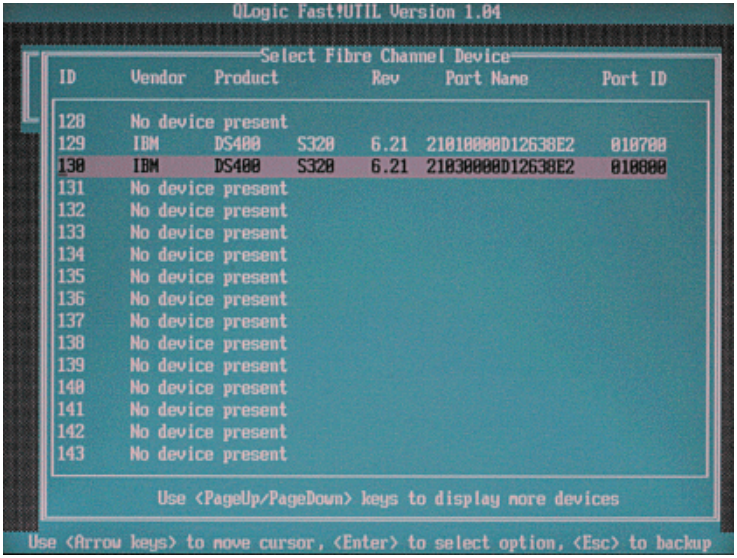

Figure 9-82 Boot Target

A list of all LUNs of this target presented. Select LUN 0 and press Enter

.

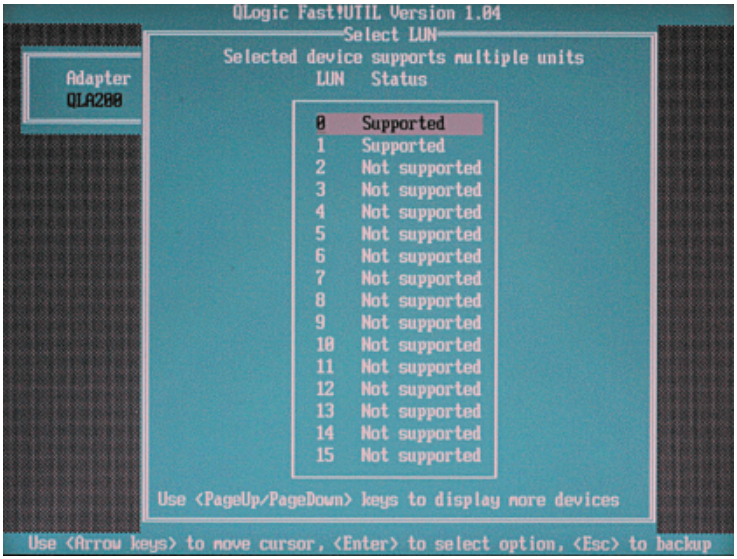

Figure 9-83 Boot LUNs

23. The target and the LUN is added to the boot device list. Press **ESC** to return to **Adapter Settings**.

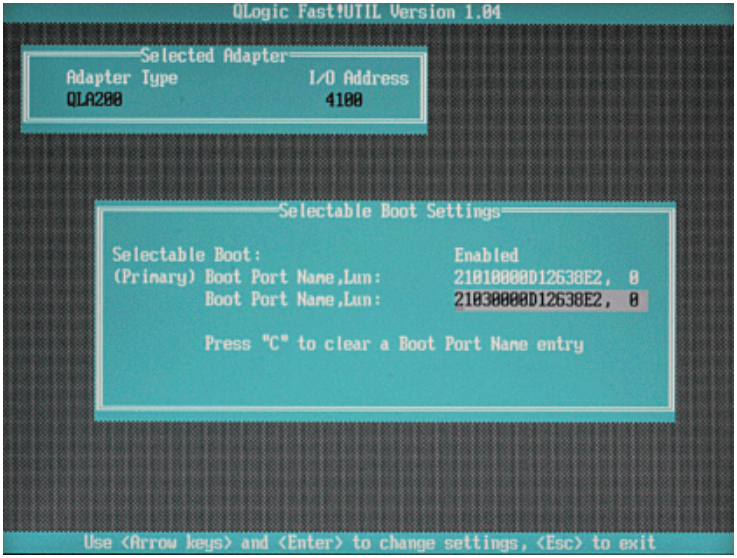

Figure 9-84 Defined boot targets and LUNs

24. Press **ESC** to leave **Adapter Settings** and return to **Configuration Settings**.

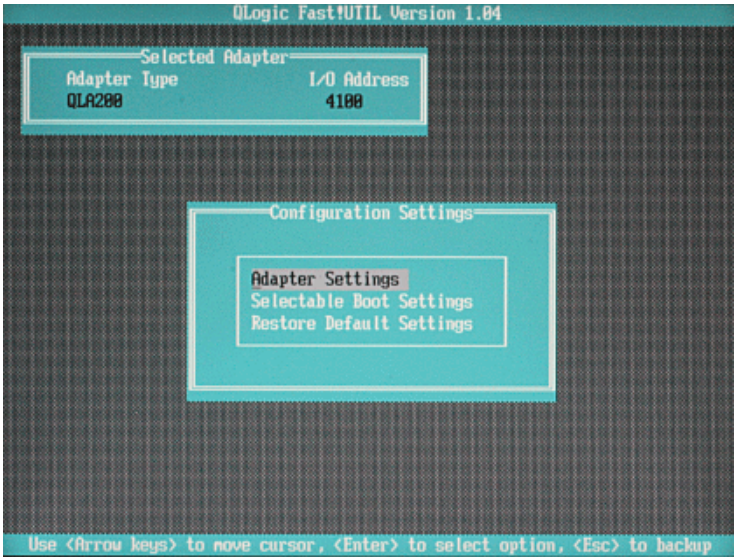

Figure 9-85 Configuration settings — Adapter settings

25. Save the modifications by selecting **Save changes** as shown in Figure 9-86 and press **Enter**.

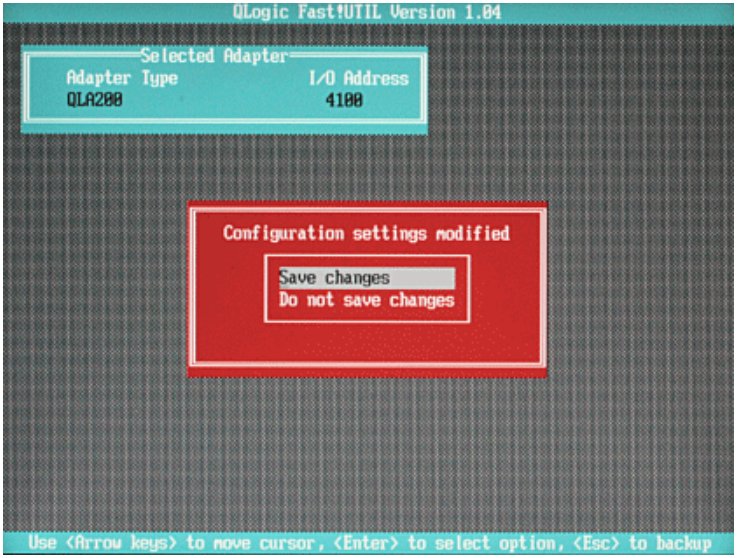

Figure 9-86 Save modifications

26. Select **Exit Fast!UTIL** and press **Enter** as shown in Figure 9-87.

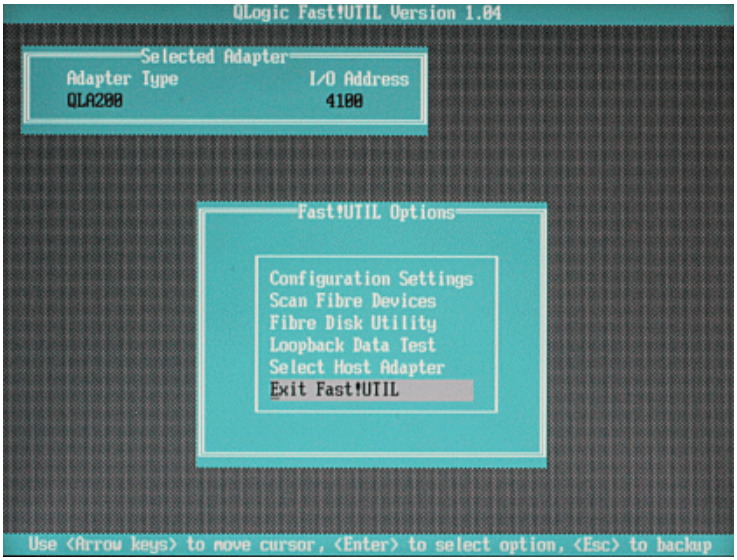

Figure 9-87 Fast!UTIL options —Exit Fast!Util

27. Select Reboot System end press Enter.

Figure 9-88 Reboot system

More logical drives my be added now. The setup is completed and the testing of the setup can start.

9.2 Boot Microsoft Windows Server 2003 from DS400 with the IBM DS4000 FC2-133 Host Bus Adapter

The setup of an IBM @serverxSeries 236 with two IBM DS4000 FC2-133 host bus adapters connected with multiple paths to a TotalStorage DS400 will be explained in detail. The IBM @serverxSeries 236 will use the TotalStorage DS4000 FC2-133 host bus adapter to boot from the TotalStorage DS400.

- 9.2.7, “Driver and service pack installation and finishing the setup” on page 533
- 9.2.8, “Enable initiator access to all TotalStorage DS400 ports” on page 541
- 9.2.9, “Finish the IBM TotalStorage DS4000 FC2-133 Host Bus Adapter BIOS setup” on page 548

9.2.1 Planning and considerations

The DS4000 FC2-133 host bus adapter does not search for an accessible LUN to boot from on an attached target like the IBM SMB Host Bus Adapter does. This requires you to specify the boot LUN in the BIOS menu of the host bus adapter. Up to four boot LUNs may be specified. Due to the fact that the host bus adapter does not has any failover capabilities, there are situations where the boot of the server will fail. There is no 100% redundancy during the boot time! Such a situation can only be solved with manual intervention.

When a problem occurs which will prevent the server to boot from the troubleshooting is much more complicated than using local boot. There is no operating system to check what went wrong.

The setup consists of redundant paths to the storage. This results in seeing the same disk multiple times. This is a kind of shared storage and this cannot be handled by Microsoft Windows An setup change for the period of operating system installation is required. The steps are described in this example when thy are needed.

All management interfaces of the various components are connected into its now management network this good practice and allows access restrictions to persons who need access and who not. When using IBM blade server to connect to a TotalStorage DS400 the separate management network is a requirement to prevent network problems.

9.2.2 SAN and network setup

The following diagram shows the cabling of the components. It follows the example configuration as explained in “Fabric attachment with quad paths without interswitch link” on page 444.

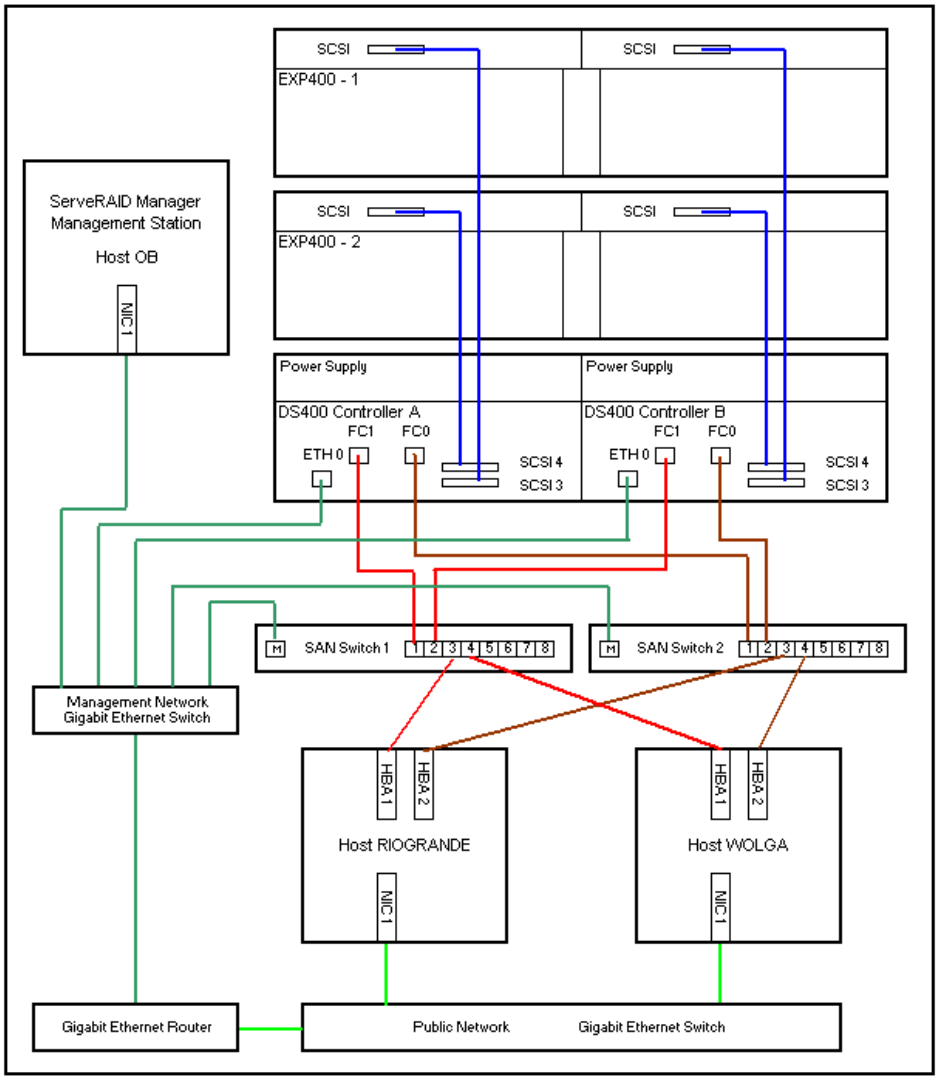

Figure 9-89 Configuration setup

The setup has used a few names which will be found in the various screenshots.

They are:

- System name is WOLGA.
- First HBA (HBA1), a IBM TotalStorage DS4000 FC2-133 host bus adapter, is installed in slot 3 and is detected by IBM FAStT MSJ as port 0. The world wide name has got the alias WOLGA-SLOT3-(0) assigned. This HBA has

access to FC0 from controller A and B. Port FC0, as described on the rear site of the TotalStorage DS400, has the name A0 in the CLI for controller A and B0 for controller B.

- Second HBA (HBA2) is also a IBM TotalStorage DS4000 FC2-133 host bus adapter and is installed in slot 4. IBM FAStT MSJ detects this HBA as port 1. The world wide name has got the alias WOLGA-SLOT4-(1) assigned. This HBA has access to FC1 from controller A and B. Port FC1, as described on the rear site of the TotalStorage DS400, has the name A1 in the CLI for controller A and B1 for controller B.
- The logical drive used to boot from is named `WOLGA-BOOT`.

Each switch has a a zone for each connected HBA configured. Soft zoning was used. Every world wide name has got an alias. The alias was used during zone configuration. SAN switch 1 has got configured a zone `Z_WOLGA_HBA1` with the members `A_WOLGA_HBA1`, `A_DS400_1_A_FC0` and `A_DS400_1_B_FC0`. SAN switch 2 has got configured a zone `Z_WOLGA_HBA2` with the members `A_WOLGA_HBA2`, `A_DS400_1_A_FC1` and `A_DS400_1_B_FC1`. The zones where added to a configuration and the configuration was activated.

9.2.3 TotalStorage DS400 setup

For this installation, five logical drive were created but only one is required to boot from. 4.8.2, “Create arrays, logical drives, and access control on the TotalStorage DS400” on page 131 describes how to create arrays and logical drives. The boot logical drive was put in its own array. following RAID configuration was applied to the TotalStorage DS400:

Table 9-1 Array and logical drive configuration of the TotalStorage DS400

Array	Drives	Owned by	Raid Level	Logical Drive	Size
Array1	2	A	1	WOLGA-BOOT	10GB
Array2	3	B	5	WOLGA-TEMP	900MB
S-array-1	6	A	50	WOLGA-MEDIA	100GB
S-array-2	4	B	10	WOLGA-DB	15GB
				WOLGA-LOG	11GB

Note: It is important for this description that the `Array-1` of the logical drive `WOLGA-BOOT` is owned by controller A. When the alternate controller owns the array of the logical drive the port names needs to be changed to reflect that the access goes over a port of controller B and not A.

Generally it is also possible to use controller B to boot from.

Access control was also defined during the creation of the array and logical drives. Both initiators have access to all five logical drives. The setup of the TotalStorage DS400 and the SAN switches are done.

The current setup allows to see each defined logical drive four times. During installation of Microsoft Windows all drives seen separately because Windows has no knowledge about multiple paths to the same disk. There is no multipathing driver or software running at this time. The result is a kind of shared storage. The multipath driver or software can be installed after the installation was success full.

Disconnecting one HBA dose not solve the issue with seeing disks multiple times because one HBA has access to controller A and controller B and the logical drive is seen over both controllers. Three ports of the TotalStorage DS400 can be disconnected together with one HBA so that every defined logical drive on the TotalStorage DS400 is seen only one by the host. This is very unpractical in productive environments.

Reconfigure the zoning on the SAN switches might be used, but this affects also all systems connected to the switch and I/O might be interrupted.

With help of the CLI the access of an initiator to a port of the TotalStorage DS400 can be restricted. Per default a initiator can access all four ports of the TotalStorage DS400. This has also the advantage the target will be seen but no logical drives. For problem determination and verification that the cabling is correct this steps need not to be reverted. The steps below explain the tasks to restrict the available logical drives seen by the host during installation of Windows Server 2003 Enterprise Edition.

Note: The described tasks can only be performed from the command line interface (CLI).

1. Use a Telnet to connect to one of the management IPs of the DS400. The examples below showing controller A. Logon with administrative rights.

```
----[ IBM DS400: Command Line Interface ]---------------------------------

               Copyright:  IBM (c)2003-2004

                Firmware:  IBM TotalStorage DS Series (Build 0743)
                           [20/07/2005][6.21]

  Peer controller status:  Ready


           Typing 'help':  provides an overview of the commands available to
                           configure the controller.

                 TAB key:  use the TAB key as a means of getting assistance
                           in completing any commands that you issue.


DS400[A]> administrator
Administrator level password: ********
DS400[A]#
```

Figure 9-90 Login as administrator

2. Get the list of initiators and aliases defined during setup of the DS400.

```
DS400[A]# initiator list

Fibre Channel Initiators

      210000e08b0f1549        [WOLGA-SLOT4-(1)]
      210000e08b0ff748        [WOLGA-SLOT3-(0)]
      210000e08b170d75        [RIOGRANDE-S4-(1)]
      210000e08b17d774        [RIOGRANDE-S3-(0)]

DS400[A]#
```

Figure 9-91 List initiators

3. Disable the access of the second HBA with alias `WOLGA-SLOT4-(1)` to any port.

```
DS400[A]# initiator manage RIOGRANDE-S4-(1) deny all
Access denied to initiator RIOGRANDE-S4-(1) on all ports

DS400[A]#
```

Figure 9-92 Deny access to an initiator

4. To be able to install onto the logical drive on the DS400 at least one port must allow access for an initiator. The port used is A0. **A1**, **B0** and **B1** will be deactivated for the first HBA.

> **Note:** The array of the logical drive to boot from needs to be owned by controller A. First HBA needs access to the first port(FC0) of the Controller A. Outside of the TotalStorage DS400 the ports are named FC0 and FC1 on each controller. CLI use A0 for FC0 and A1 for FC1 of controller A and B0 for FC0 and B1 for FC1 of controller B.
>
> Basically every port and HBA can be used as long as the host can see every logical drive only once during installation of Microsoft Windows.

```
DS400[A]# initiator manage RIOGRANDE-S3-(0) deny A1 B0 B1
Access denied to initiator RIOGRANDE-S3-(0) on ports A1 B0 and B1

DS400[A]#
```

Figure 9-93 Deny access on ports A1, B0 and B1

5. Verify that only one port is enabled for only one HBA. Only port A0 allows access for the first HBA.

The access of the initiators to the ports needs to be enabled after finishing the installation of the operating system and the multipath driver.

```
DS400[A]# initiator manage WOLGA-SLOT3-(0) info

----[ Initiators ]-------------------------------------------------------


----[ Initiator 'WOLGA-SLOT3-(0)' ]---------------------------------------

                WWN:  210000e08b0ff748
              Alias:  WOLGA-SLOT3-(0)
                 ID:  0
        Port access:
                      Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------
                      YES            NO             NO             NO


       LUN Mappings:
                LUN                 logical          [Capacity]
-------------------------------------------------------------------------

                0                   WOLGA-BOOT       [ 30.00GB]
                1                   WOLGA-DB         [ 15.00GB]
                2                   WOLGA-LOG        [ 11.01GB]
                3                   WOLGA-MEDIA      [100.00GB]
                4                   WOLGA-TEMP       [900.00MB]

DS400[A]#
```

Figure 9-94 Verify port A0 access

6. No access on any port is allowed for the second HBA.

```
DS400[A]# initiator manage WOLGA-SLOT4-(1) info

----[ Initiators ]-----------------------------------------------------


----[ Initiator 'WOLGA-SLOT4-(1)' ]-------------------------------------

                WWN:  210000e08b0f1549
              Alias:  WOLGA-SLOT4-(1)
                 ID:  1
        Port access:
                      Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------
                      NO             NO             NO             NO


       LUN Mappings:
                LUN                 logical          [Capacity]
-------------------------------------------------------------------------

                0                   WOLGA-BOOT       [ 30.00GB]
                1                   WOLGA-DB         [ 15.00GB]
                2                   WOLGA-LOG        [ 11.01GB]
                3                   WOLGA-MEDIA      [100.00GB]
                4                   WOLGA-TEMP       [900.00MB]

DS400[A]#
```

Figure 9-95 Verify no port has access

9.2.4 Host bus adapter setup and installation preparation

The following steps are done to configure the IBM TotalStorage DS4000 FC2-133 host bus adapter in the IBM @server x236. The adapters are installed in PCI slots 3 and 4.

1. Update BIOS and firmware of the HBA to the latest level supported by IBM for the DS400. The BIOS can be downloaded from the IBM support Web page

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

   See Appendix C, “How to get updates of code and documentation” on page 623. At time of writing the supported BIOS level is 1.47. The readme file of this update contains the details how the update need to be performed.

2. Reset the HBA to adapter Defaults.

a. During POST the BIOS of the TotalStorage DS4000 FC2-133 HBA shows a message **Press Ctrl+Q for Fast!UTIL**. Press **CTRL**+**Q** to enter the setup.

Note: ALT+Q works as well.

Figure 9-96 POST — TotalStorage DS4000 FC2-133 adapter BIOS message

b. Select the second HBA and press Enter.

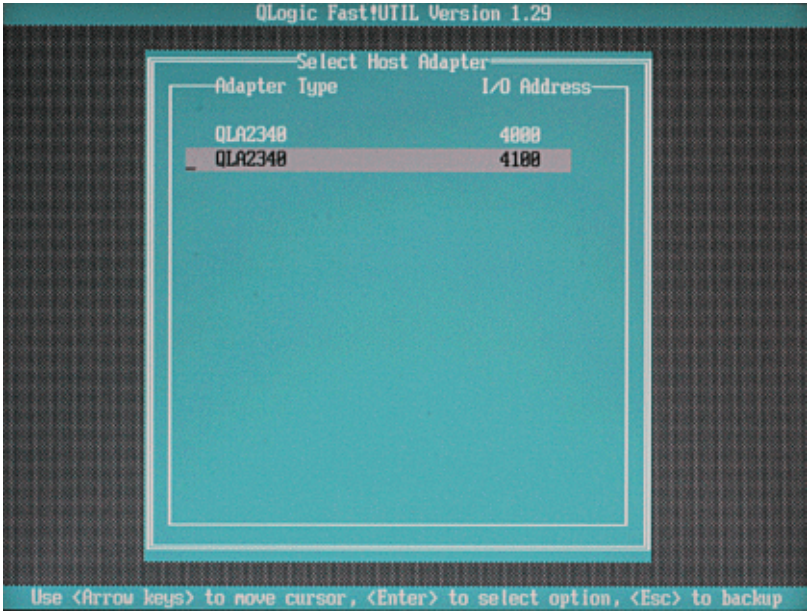

Figure 9-97 Select host adapters

c. Select **Configuration Settings** and press **Enter**.

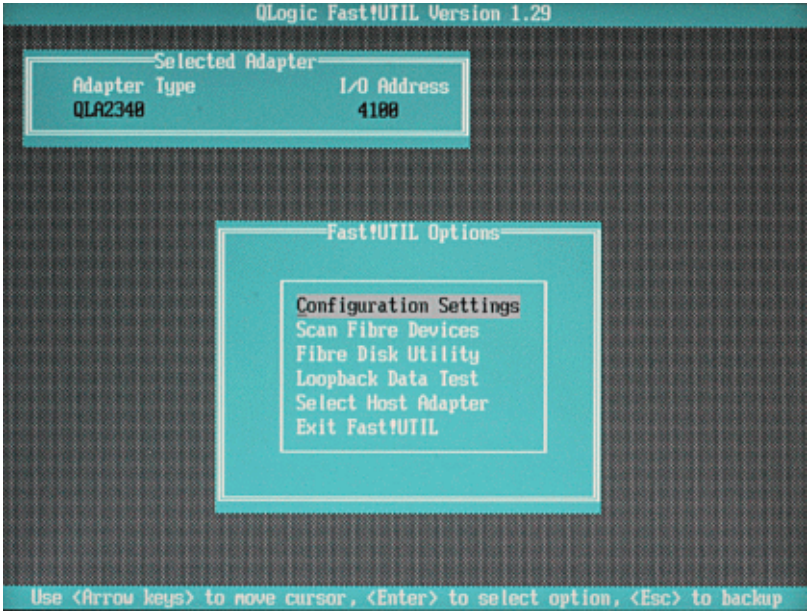

Figure 9-98 Fast!Util options — Configuration settings

d. Select **Restore Default Settings** and press **Enter**.

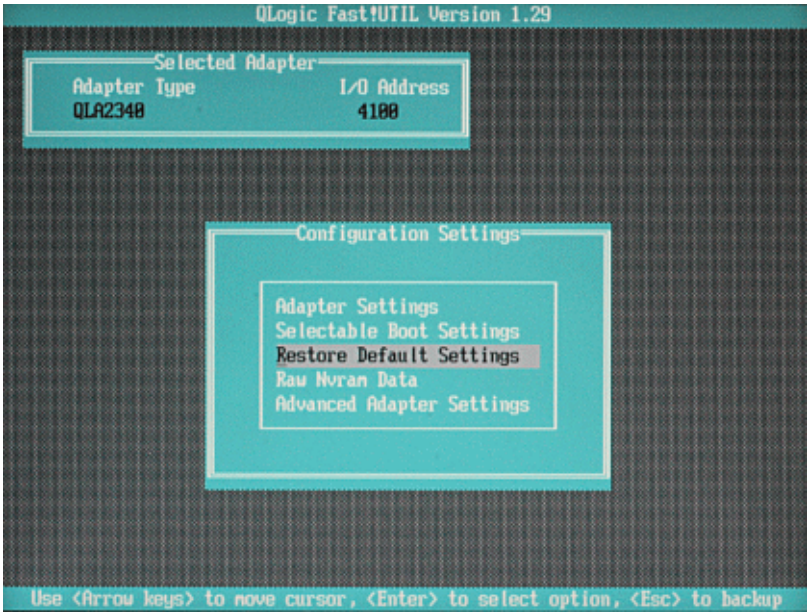

Figure 9-99 Configuration settings — Restore default settings

e. The default setting of the adapter will be applied. Press **any** key to return to **Configuration Settings**.

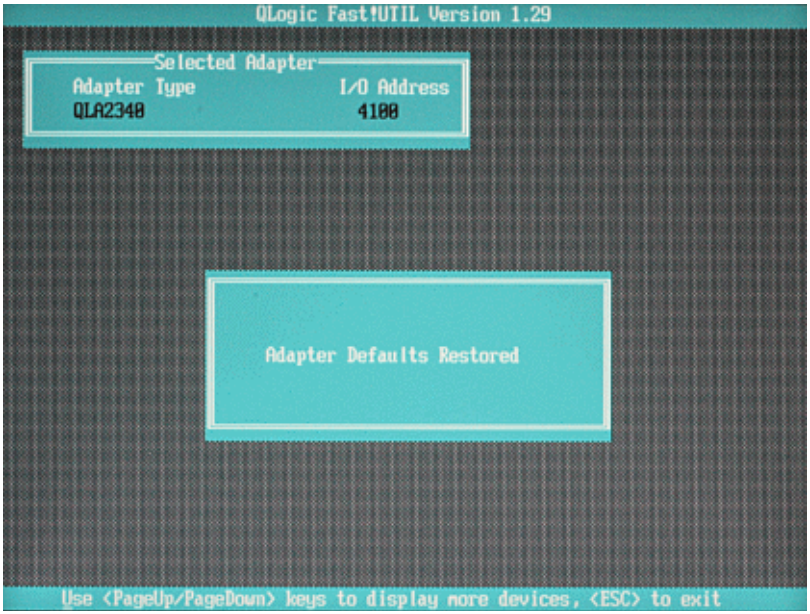

Figure 9-100 Adapter defaults restored

f. Press **ESC** to return to **Configuration Settings**.

g. Press **ESC** again to leave **Configuration Settings** and return to **Fast!UTIL Options**.

h. Save the modifications. by pressing **Enter**.

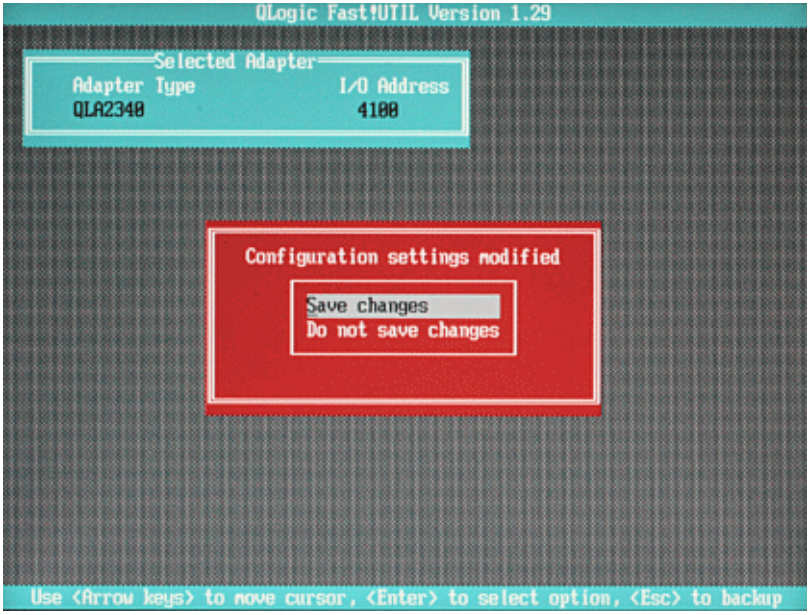

Figure 9-101 Save modifications

i. After all modifications are saved in NVRAM, select **Select Host Adapter** from the **Fast!UTIL Options** and press **Enter**.

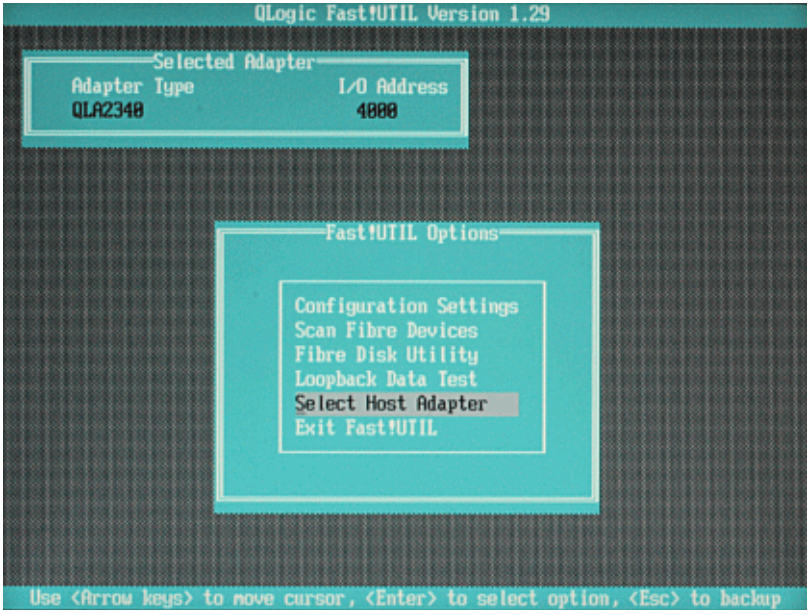

Figure 9-102 Fast!UTIL options — Select host adapter

j. Select the first HBA and press **Enter**.

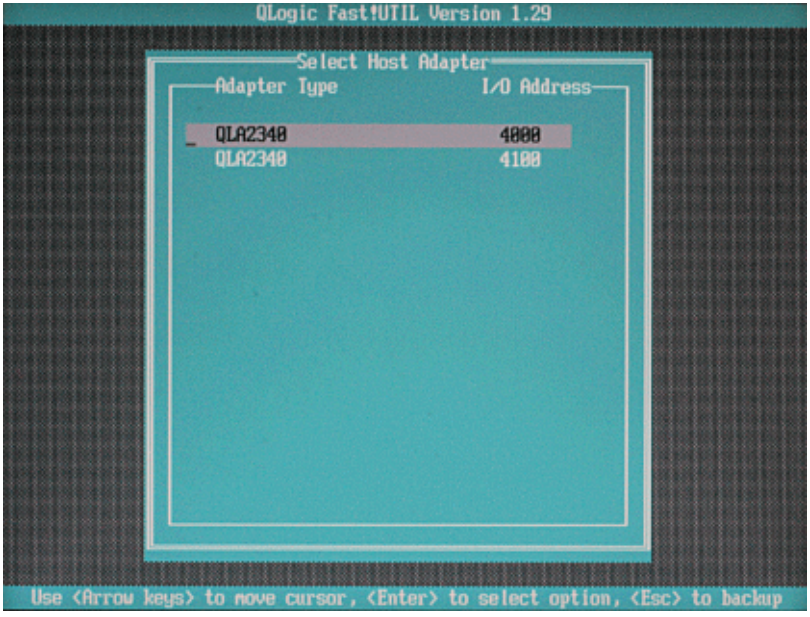

Figure 9-103 Select host adapters

k. Reset the first HBA to factory default. Perform Step c on page 511 to Step f on page 513.

3. Verify that the first adapter can see one port of each controller.

   a. Select **Scan Fibre Devices** from **Configuration Settings** and press **Enter**.

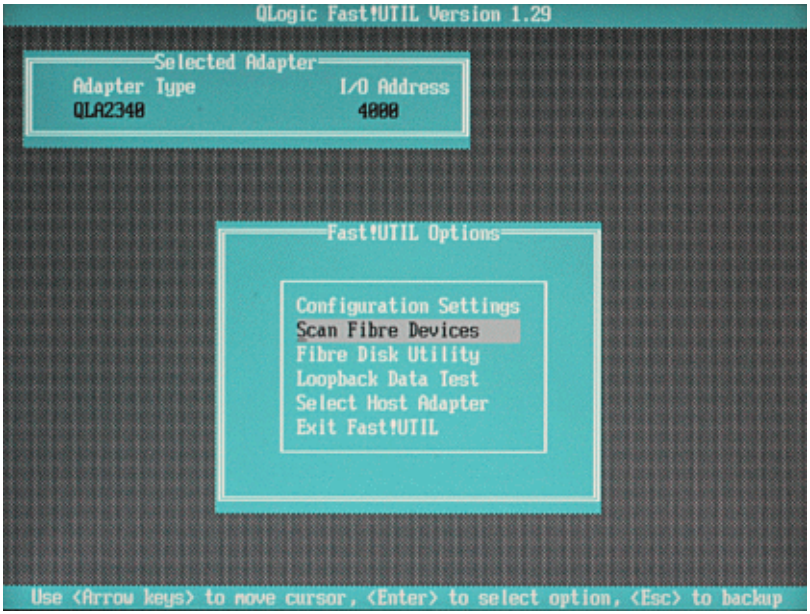

Figure 9-104 Fast!UTIL Options

   b. A list with all detected devices will be shown. The list needs two entries for a connected TotalStorage DS400. One for one port of controller A and one for one port of controller B. The port name can be used to identify the TotalStorage DS400. The last 8 digits of the port name are identifying the TotalStorage DS400 across both controllers and ports. The third and fourth digit identifies the port to which the HBA has a connect. The screen shows that the first HBA sees ID 129 which is the Controller A port FC0 (A0) and ID 130 which is the Controller B port FC0 (B0). This screen will only show the target devices but no LUNs.

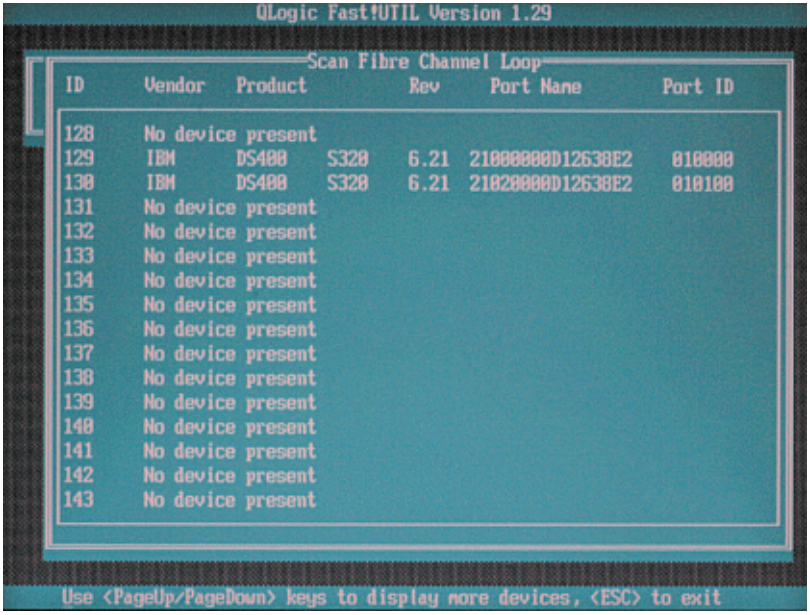

Figure 9-105 Fibre channel device list

c. Press **ESC** to return to the **Fast!UTIL Options**.

4. Enable the adapter BIOS for the first HBA.

   a. Select Configuration Settings from the **Fast!UTIL Options** and press **Enter**.

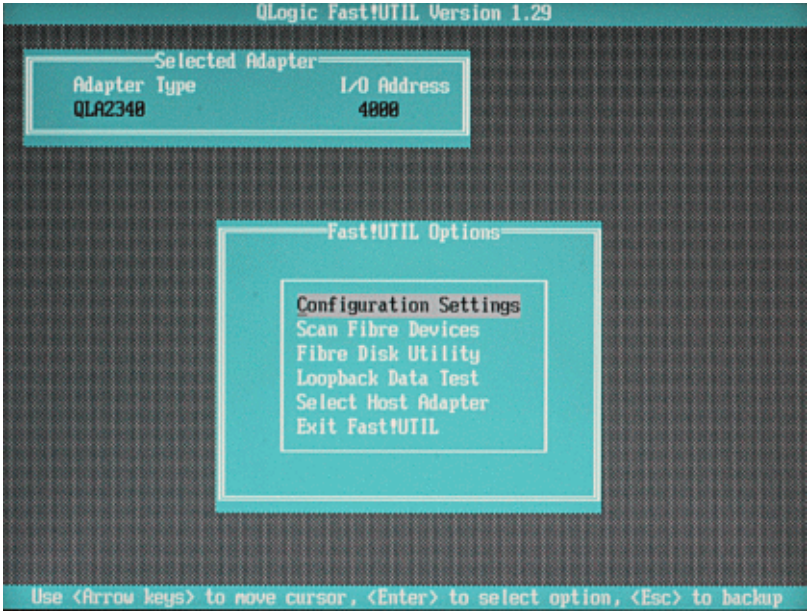

Figure 9-106 Fast!UTIL Options — Configuration Settings

b. Select **Adapter Settings** from the **Configuration Settings** and press **Enter**.

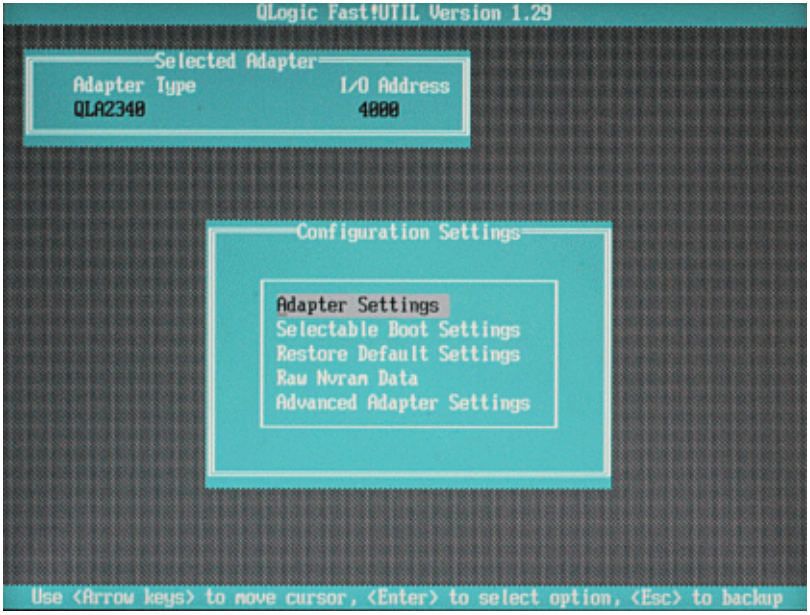

Figure 9-107 Configuration settings — Adapter Settings

c. From the **Adapter Settings** select **Host Adapter BIOS: Disabled** and press **Enter**. The setting will change to **Enabled**. Then Press **ESC** to return to the **Configuration Settings**.

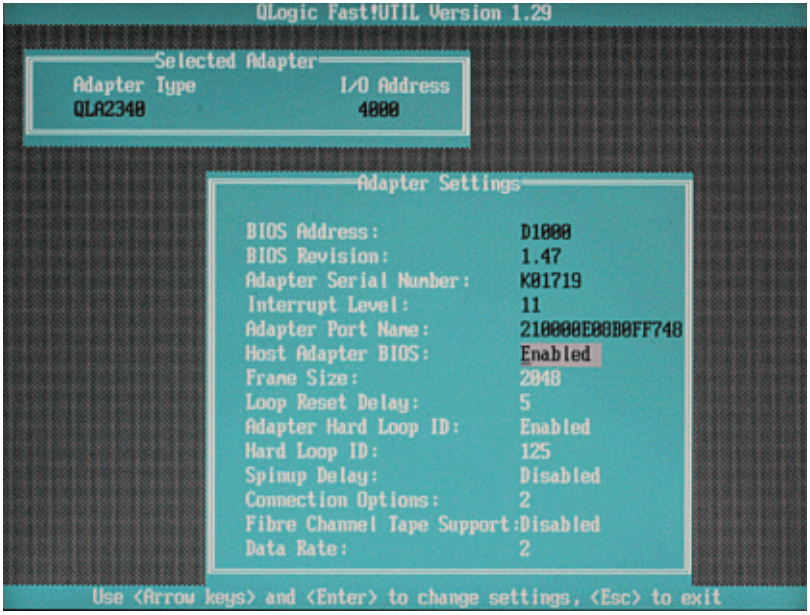

Figure 9-108 Adapter Settings

5. Define the BOOT LUNs for this first HBA.
   a. Select **Selectable Boot Settings** from the **Configuration Settings** and press **Enter**.

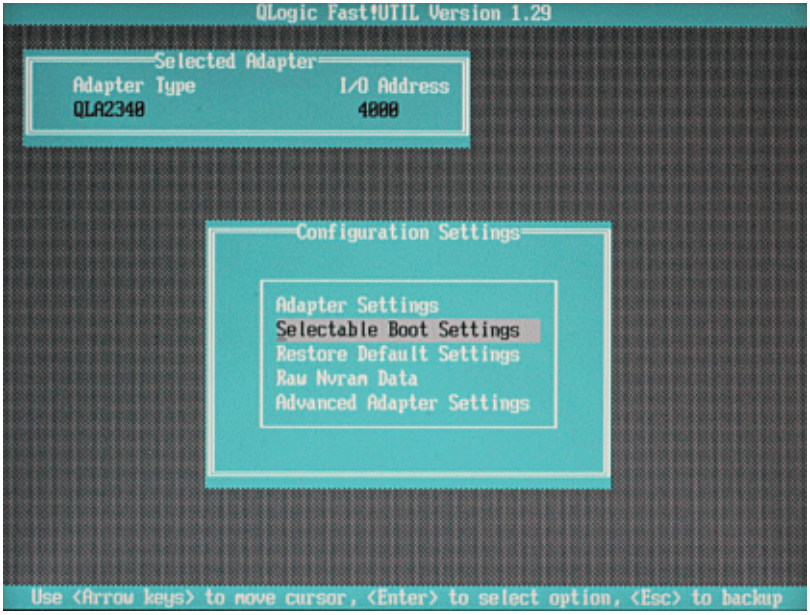

Figure 9-109 Configuration settings — Selectable boot settings

b. Enable the **Selectable Boot** by selecting this option and press **Enter**. The value will change from Disabled to Enabled.

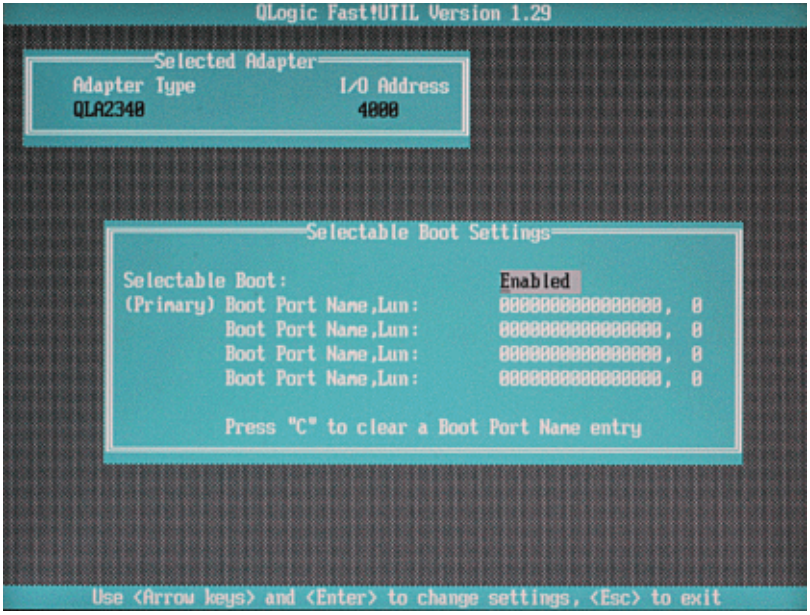

Figure 9-110 Selectable boot settings

c. Select **(Primary) Boot Port Name, Lun** and press **Enter**.

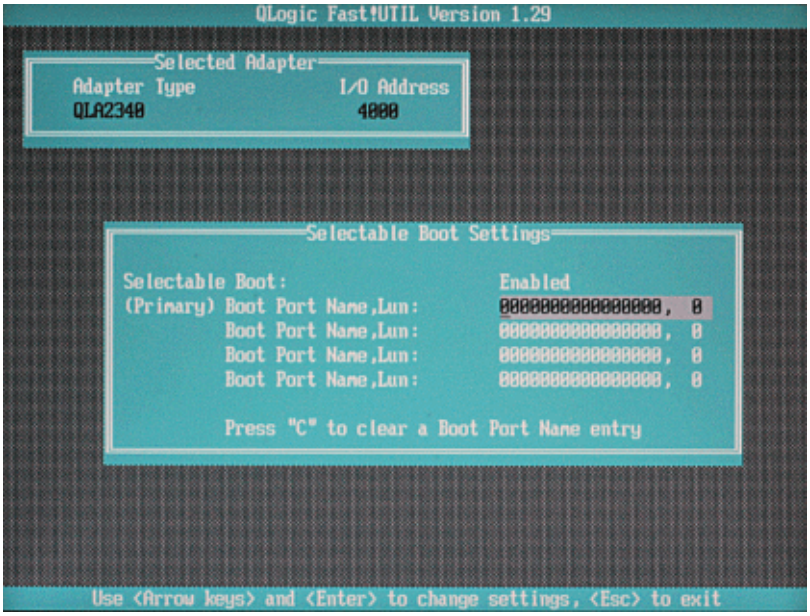

Figure 9-111 Selectable boot settings — Primary boot port

d. Select now the target which owns the logical drive to boot from. In this case the logical drive is owned by controller A and the initiator `WOLGA-SLOT3-(0)` (HBA1) is allowed to access the port FC0 of controller A. Press **Enter**. See 9.2.3, “TotalStorage DS400 setup” on page 504.

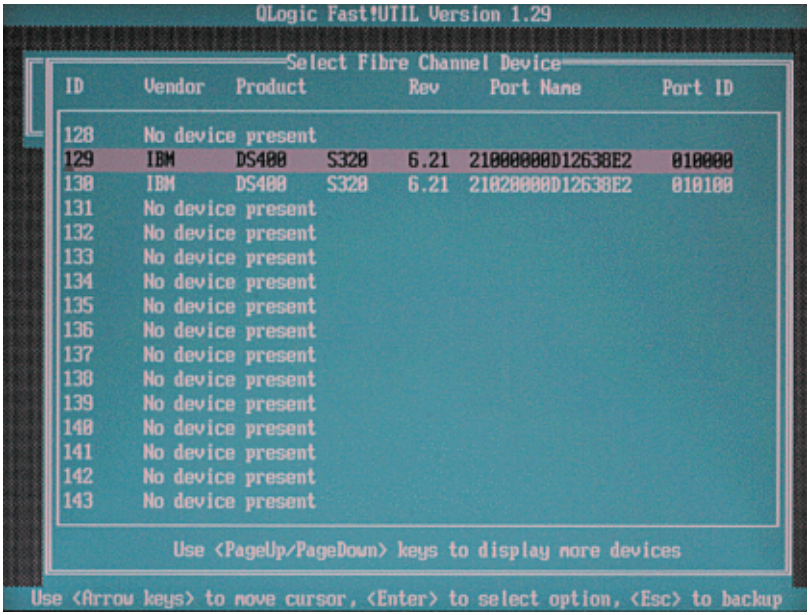

Figure 9-112 Fibre channel targets

When you select the second target in the list an error message is received that the selected device is not a disk device. Please remember that the HBA is not allowed to access this target at the moment. This is required to have only a logical drive only over a single path accessible during the installation of Microsoft Windows Server 2003. This is a limitation by the operating system. The same is the case for the second HBA.

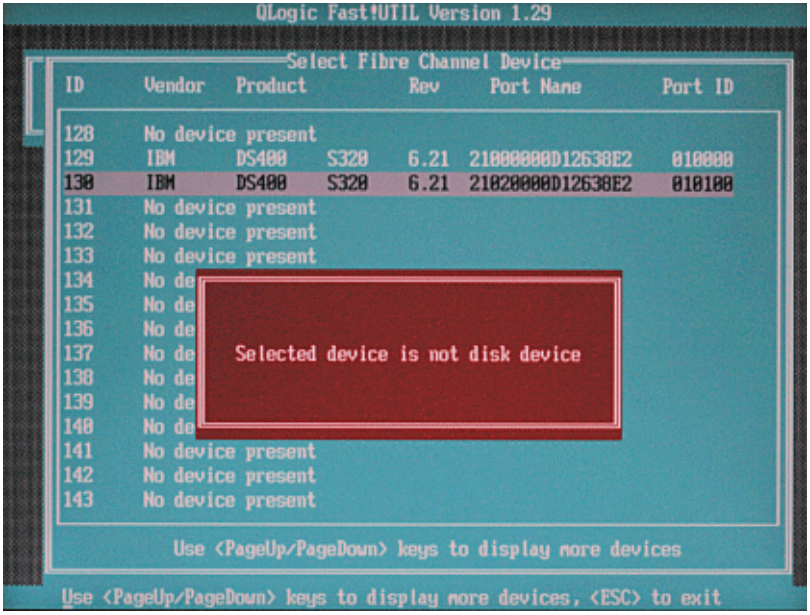

Figure 9-113 Selected device is not a disk device

e. Select from the list of LUNs the LUN from which the server will boot. The logical drive `WOLGA-BOOT` was as first logical drive assigned to the initiators. The DS400 will present the configured logical drives in that order to the host as they where assigned to the initiator. Select the `LUN 0` end press **Enter**.

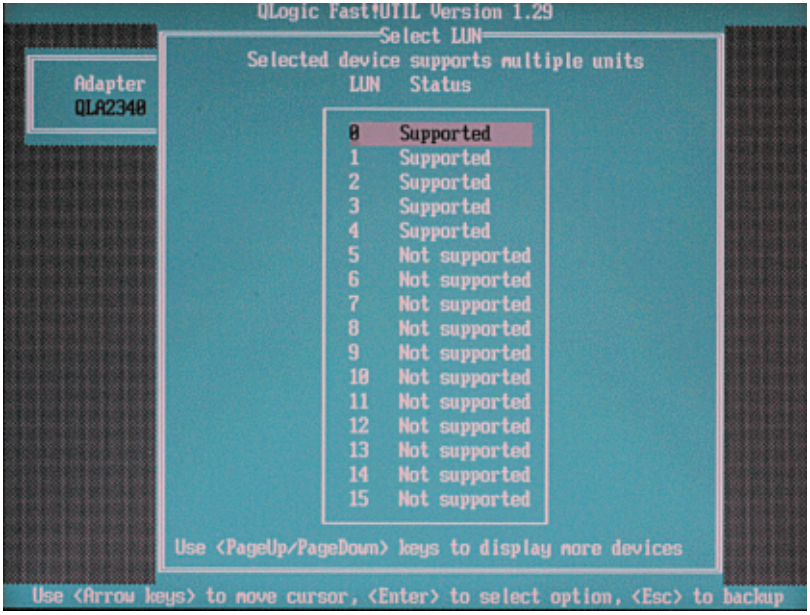

Figure 9-114 Available LUN s of the selected fibre channel target

f. The world wide port name the target and the LUN are shown as values. Press **ESC** to return to **Configuration Settings**.

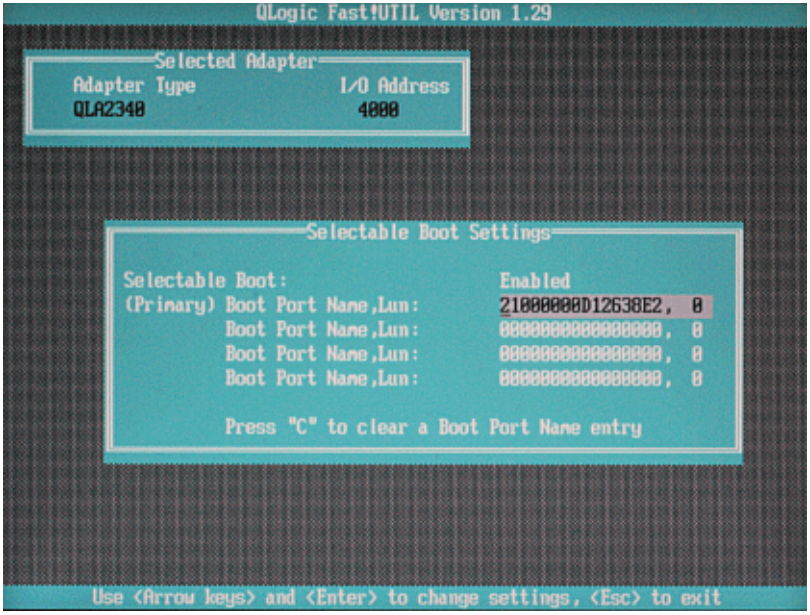

Figure 9-115 Selected boot device

g. Press **ESC** to leave the Configuration Settings dialog. All configuration modifications will be saved now. Select Save changes and press **Enter** to save the modifications.

Figure 9-116 Save modifications

h. The Fast!Util options dialog is shown. Select **Exit Fast!UTIL** and press **Enter** to leave the BIOS setup.

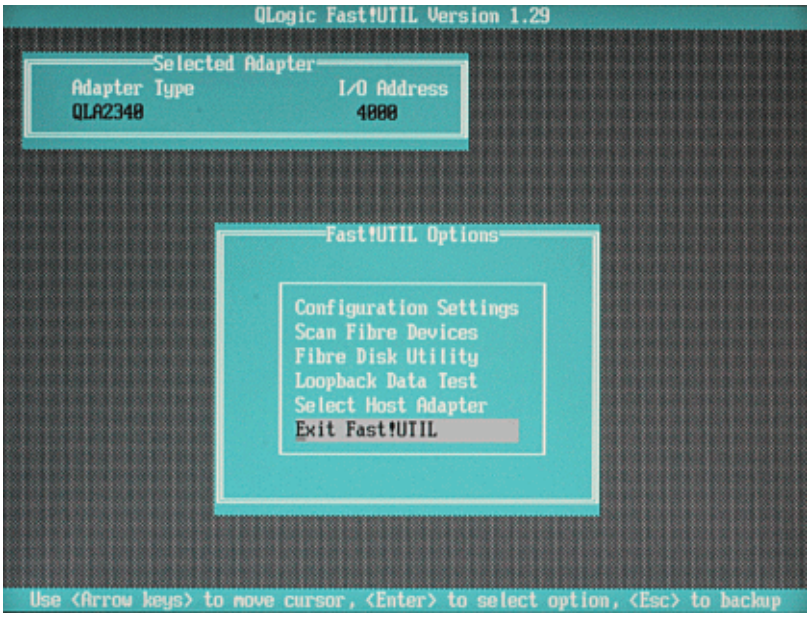

Figure 9-117 Modifications

i. Do not reboot the system until the driver for the HBA is ready.

6. The HBAs are now configured for the Installation of Microsoft Windows Server 2003. Each logical drive of the TotalStorage DS400 is seen by the host only once, even with multiple paths.

9.2.5 Create IBM DS4000 Host Bus Adapter driver disk

The driver disk is required because the driver included on the Microsoft Windows Server 2003 Enterprise Edition CD is out of date.

1. Create the driver disk for the TotalStorage DS4000 FC2-133 host bus adapter. Download the driver for the TotalStorage DS4000 FC2-133 host bus adapter from the QLogic Web site. Use the software and driver matrix of the DS400 to get the right driver for the host bus adapter in combination with the TotalStorage DS400.

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

   Unzip the driver archive and store all files in the root of a 3.5” diskette. This diskette will be required to install Microsoft Windows Server 2003.

2. Proceed with the installation of Microsoft Windows Server 2003 as described in the next section.

9.2.6 Microsoft Windows Server 2003 Enterprise Edition installation

The installation of Windows 2003 is straight forward as already known. The driver of the TotalStorage DS4000 FC2-133 host bus adapter is included on the Microsoft Windows Server 2003 Enterprise Edition CD. However, this driver is to old and may not detect disks. Thus requires to create a driver diskette first. See chapter Create IBM DS4000 Host Bus Adapter driver disk about the creation of the driver disk.

1. The system is still in the Host bus adapter BIOS setup. Select **Reboot System** and press **Enter** to leave the setup and restart the system now.

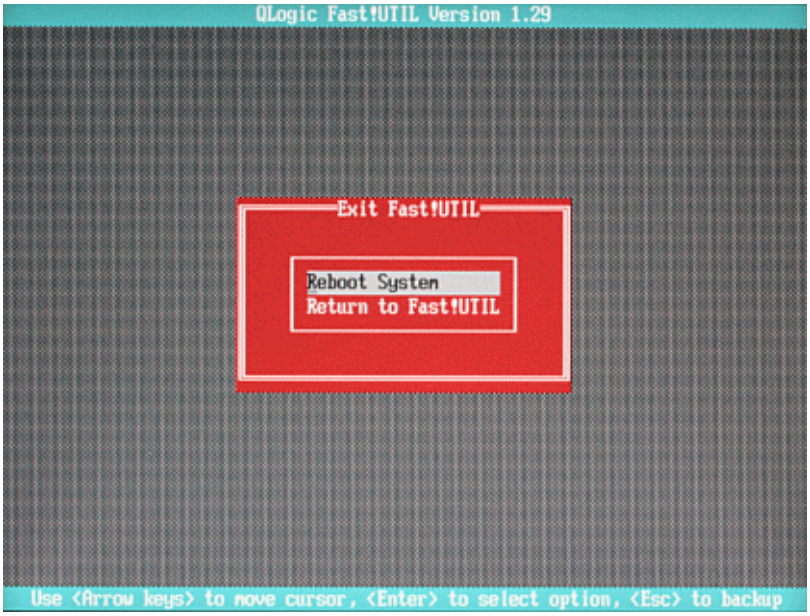

Figure 9-118 Exit Fast!UTIL img_8420.gif 500DPI

2. Insert the installation CD into the CD-ROM drive and boot from this CD. After the message Setup is inspecting your computer was shown the black screen will switch to blue and the bottom line give you the option to press F6 to install third party SCSI or RAID drives which are required at boot time. Press **F6** to install the driver for the IBM TotalStorage DS4000 FC2-133 host bus adapter.

Figure 9-119 Windows setup — install third party SCSI or RAID drivers

3. Insert the driver disk created in “Create IBM DS4000 Host Bus Adapter driver disk” on page 528 and press **Enter**.

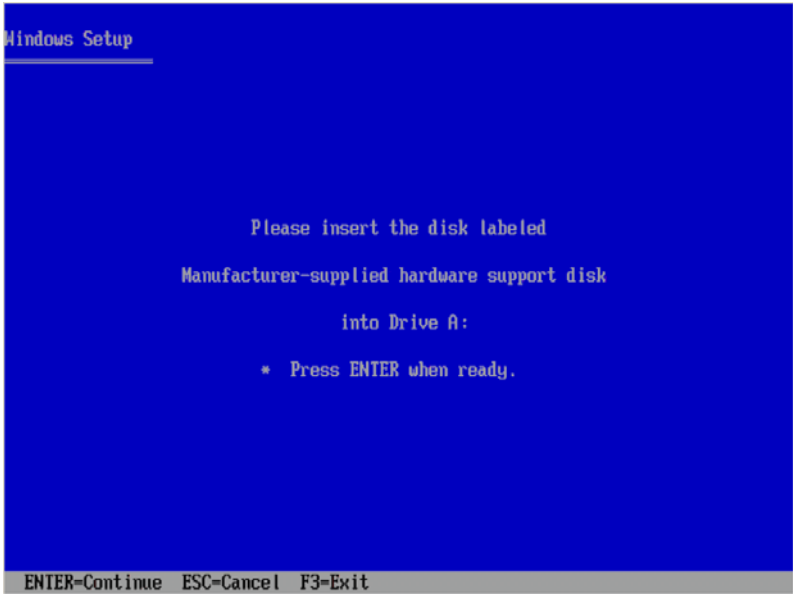

Figure 9-120 Windows setup — insert driver disk

4. The driver for the QLogic Fibre Channel Adapter will be found. Press **Enter** to select this driver.

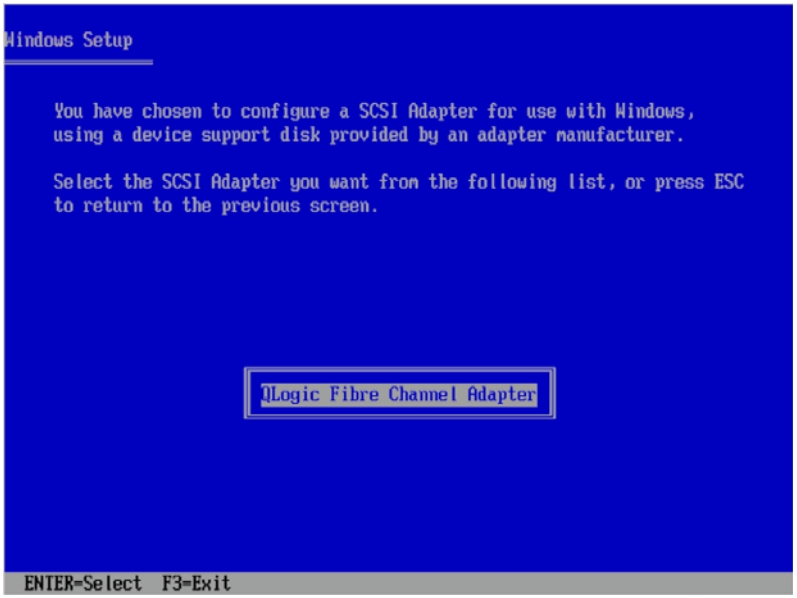

Figure 9-121 Windows setup — QLogic Fibre Channel Adapter

5. The driver is newer than the driver included on the CD. Press **S** to use the newer driver from the disk.

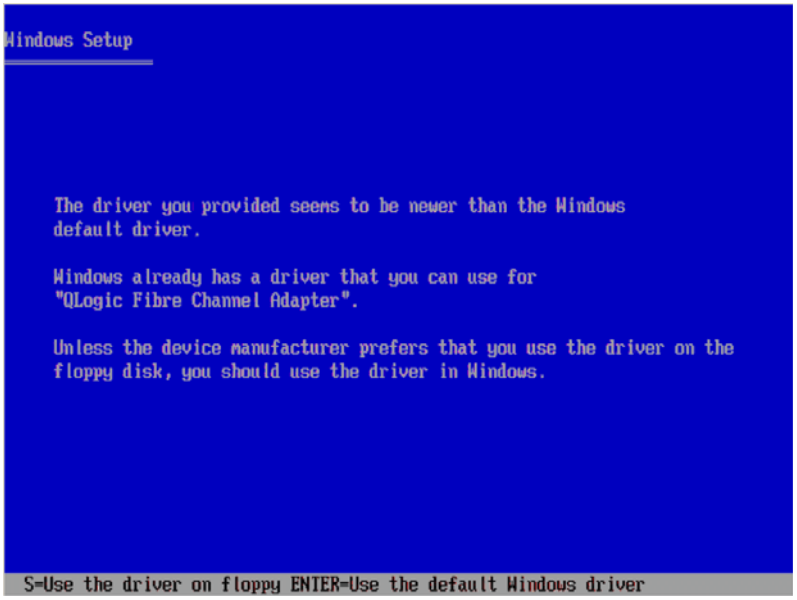

Figure 9-122 Windows setup — newer driver found

6. The driver for the DS4000 FC2-133 host bus adapter will be loaded and Windows setup returns to the screen to add more third party SCSI and RAID driver. Press Enter to proceed with the installation.

Figure 9-123 Windows setup — Supported mass storage drivers

7. For the remaining installation follow the instructions on the screen. There is no difference from the installation to a local attached drive.

9.2.7 Driver and service pack installation and finishing the setup

Install the Broadcom network driver for the adapter. Broadcom NetXtreme Gigabit Ethernet #2 is connected to the port number 1. This port is also shared by the baseboard management controller (BMC). This port might be connected into a management network. Broadcom NetXtreme Gigabit Ethernet is routed to port number 2. This is the port close to the power supplies.

There are operating system install instructions for all IBM IBM @serverxSeries server on the support pages. See Appendix C, “How to get updates of code and documentation” on page 623.

8. Install the service pack1 for Microsoft Windows Server 2003. Reboot the server.

   http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

9. Install remaining drivers not available during operating system installation:

   a. Install the Intel chipset drivers.

   b. Install ActivePCI.

c. All unknown devices are now removed.

d. Install IPMI driver. This driver is used for the BMC.

10. Install the IBM FAStT MSJ to manage the DS4000 FC2-133 host bus adapter.

a. Download the latest version of the FAStT MSJ from IBM's support Web page and run the installer file.

http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56707

Click **Run** to execute the installer.

Figure 9-124 open file security warning

b. Click Next.

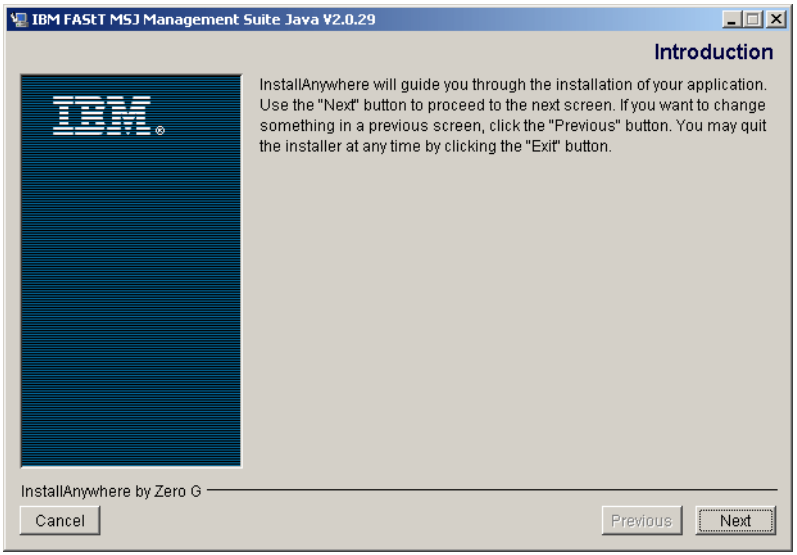

Figure 9-125 MSJ installation — Introduction

c. Accept the license agreement and click **Next**.

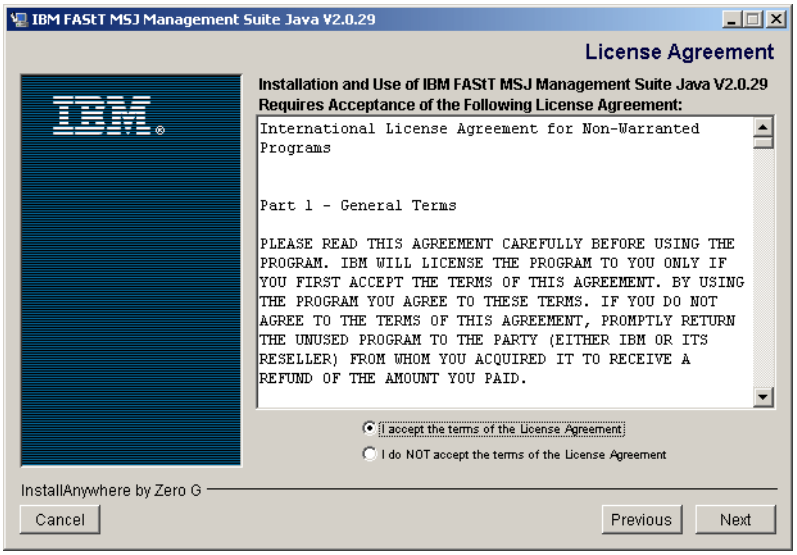

Figure 9-126 MSJ installation — license agreement

d. Read important information and click Next.

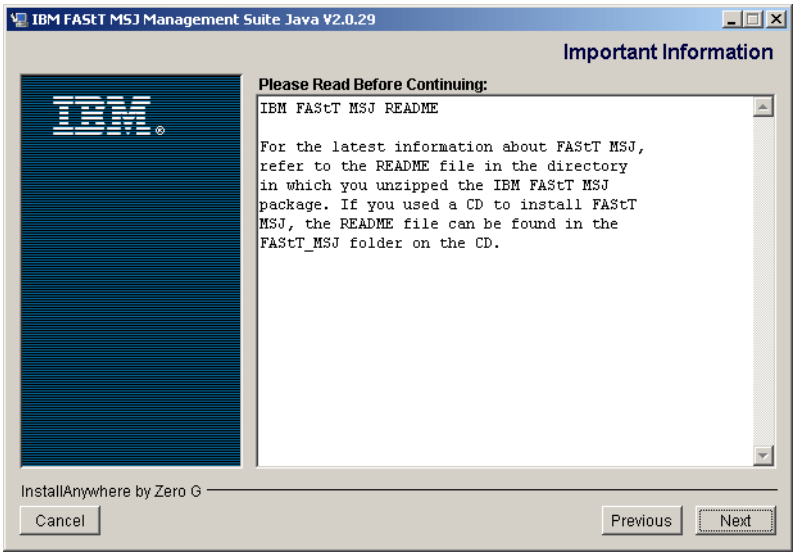

Figure 9-127 MSJ installation — important information

e. Select **GUI and NT agent** and click **Next**. For a single system the MSJ GUI and agent should be installed. The agent need to be installed on the system with the host bus adapters. The GUI is only once required in a network but might be installed on multiple systems.

Figure 9-128 MSJ installation — Choose product features

f. Specify the installation directory and click **Next**.

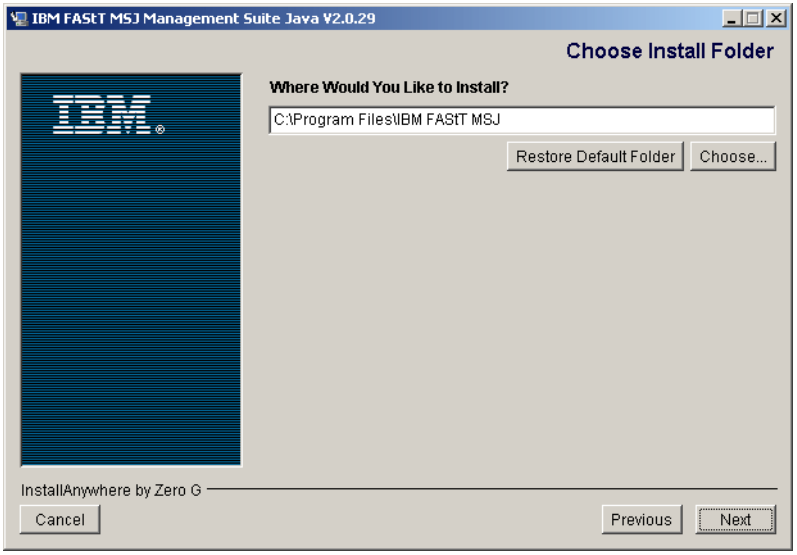

Figure 9-129 MSJ installation — Choose install directory

g. Specify the location of the shortcuts and click Next.

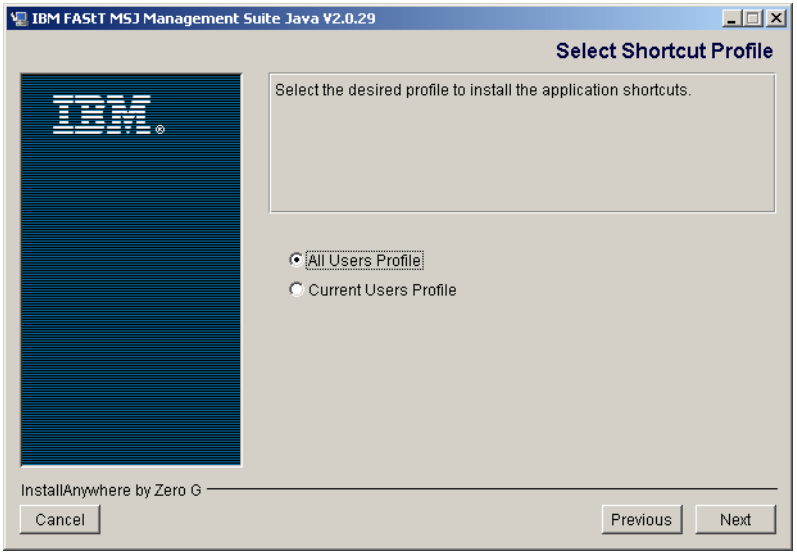

Figure 9-130 MSJ installation — select shortcut profile

h. Choose if a desktop icon is required and click **Install**.

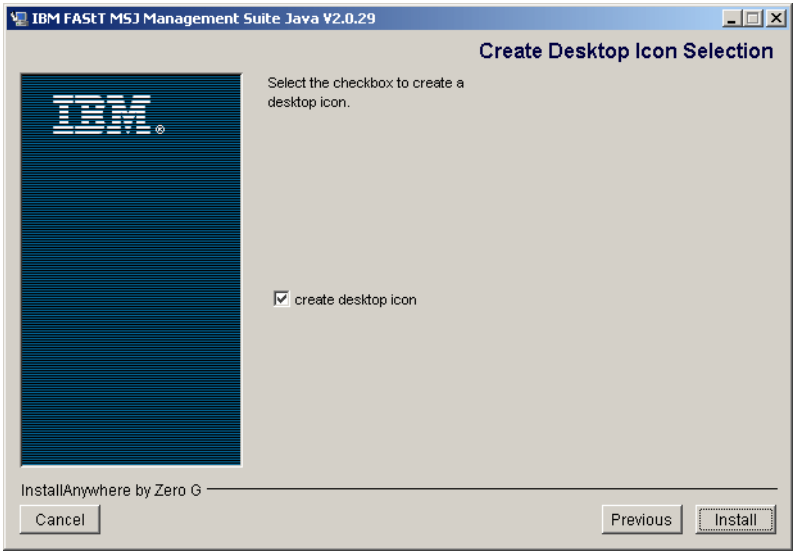

Figure 9-131 MSJ installation — create desktop icon

i. Files will be copied. When installation is complete click **Done**.

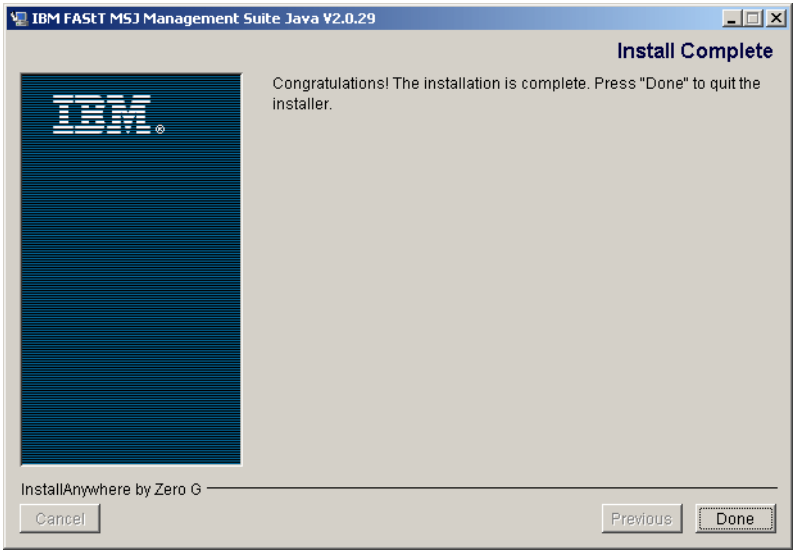

Figure 9-132 MSJ installation — Done

11. Install the QLogic Device Specific Module (DSM) with the multipath input/output (MPIO) driver.

QLogic provides a special Device Specific Module for Microsoft multipath input/output driver. This DSM is specifically designed to work with the IBM TotalStorage DS400. See 3.4.4, “DS400 failover” on page 51 for failover details.

For more information about MPIO, see Microsoft whitepaper “Highly Available Storage: Multipathing and the Microsoft MPIO Driver Architecture”:

http://www.microsoft.com/windowsserver2003/technologies/storage/mpio/default.mspx

Frequently Asked Questions: Microsoft Multipath I/O:

http://www.microsoft.com/WindowsServer2003/technologies/storage/mpio/faq.mspx

To install the QLogic Device Specific Module together with the Multipath IO driver perform the following steps.

a. Download and extract the driver archive. See DS400 software and driver matrix for the correct versions of the HBA driver and the DSM and MPIO.

   http://www.ibm.com/support/docview.wss?rs=1110&uid=psg1MIGR-58619

b. Run the install.bat file included in the driver archive.

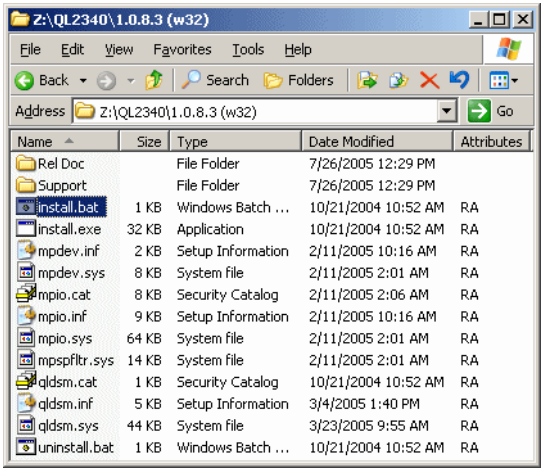

Figure 9-133 Content of the DSM and MPI O driver archive

c. Confirm the open file security warning by clicking **Run**.

Figure 9-134 Open file security warning

d. A command line box opens and displays the progress of the drive installation.

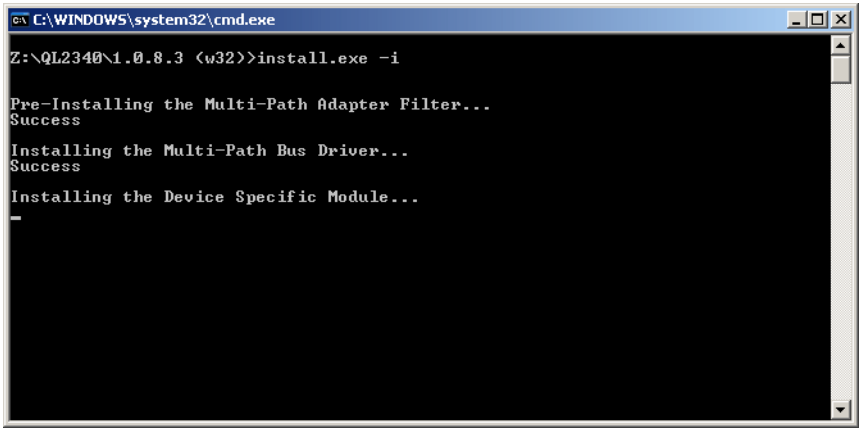

Figure 9-135 Driver installation output

e. The DSM driver is currently not yet digitally signed. Confirm the warning to install the unsigned driver.

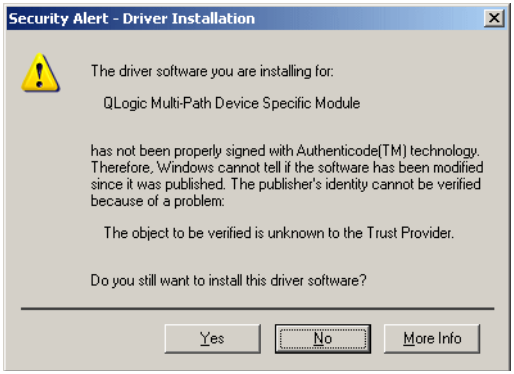

Figure 9-136 Unsigned driver warning

f. After the DSM module was installed the command line window will close. In the device manager a new device in the **SCSI and raid controller** section is displayed.

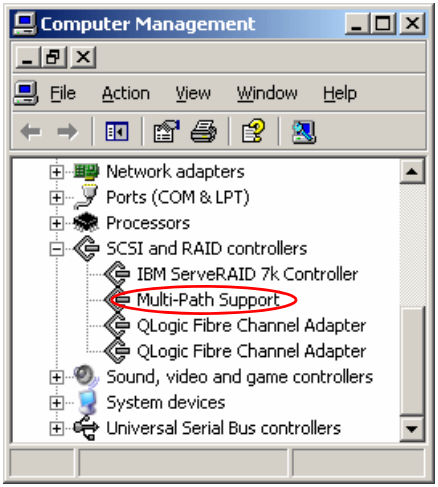

Figure 9-137 Device Manager after installing DSM and MPIO

g. Reboot the server.

9.2.8 Enable initiator access to all TotalStorage DS400 ports

Microsoft Windows Server 2003 is installed and the multipath driver as well. There are multiple paths available but no LUNs are accessible here. This step will activate the ports of the TotalStorage DS400 for the initiators to gain access over

all ports. The access was restricted in 9.2.3, “TotalStorage DS400 setup” on page 504.

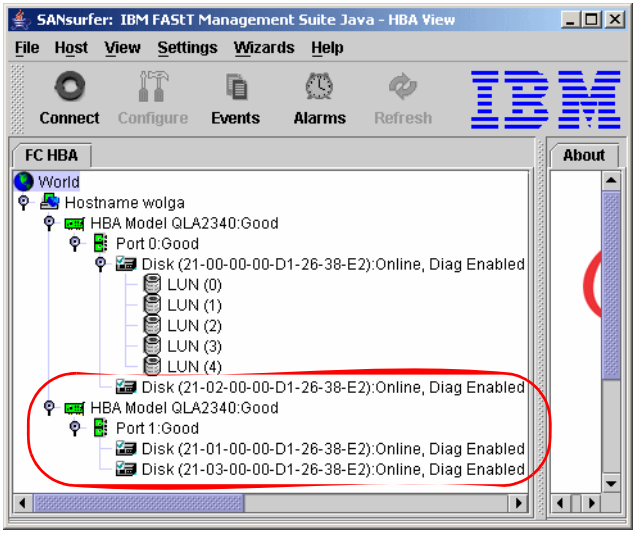

Figure 9-138 Assigned LUNs to the HBAs

Follow the steps below to activate access to all ports:

1. Open a Telnet connection to one of the controllers in the DS400 and logon as administrator. To logon with administrative rights enter **administrator**.

```
----[ IBM DS400: Command Line Interface ]---------------------------------

               Copyright:  IBM (c)2003-2004

                Firmware:  IBM TotalStorage DS Series (Build 0743)
                           [20/07/2005][6.21]

  Peer controller status:  Ready


           Typing 'help':  provides an overview of the commands available to
                           configure the controller.

                 TAB key:  use the TAB key as a means of getting assistance
                           in completing any commands that you issue.


DS400[A]> administrator
Administrator level password: ********
DS400[A]#
```

Figure 9-139 Logon as administrator

2. Get a list of the available initiators by entering **initiator list**.

```
DS400[A]# initiator list

Fibre Channel Initiators

      210000e08b0f1549        [WOLGA-SLOT4-(1)]
      210000e08b0ff748        [WOLGA-SLOT3-(0)]
      210000e08b170d75        [RIOGRANDE-S4-(1)]
      210000e08b17d774        [RIOGRANDE-S3-(0)]

DS400[A]#
```

Figure 9-140 List initiators

3. Look for the initiators of WOLGA and check port access for every singly initiator by entering **initiator manage** WOLGA-SLOT4-(1) **info.** All ports are disabled for this initiator.

```
DS400[A]# initiator manage WOLGA-SLOT4-(1) info

----[ Initiators ]-----------------------------------------------------


----[ Initiator 'WOLGA-SLOT4-(1)' ]-------------------------------------

                WWN:  210000e08b0f1549
              Alias:  WOLGA-SLOT4-(1)
                 ID:  1
        Port access:
                      Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------
                      NO             NO             NO             NO


       LUN Mappings:
                LUN                 logical          [Capacity]
-------------------------------------------------------------------------

                0                   WOLGA-BOOT       [ 30.00GB]
                1                   WOLGA-DB         [ 15.00GB]
                2                   WOLGA-LOG        [ 11.01GB]
                3                   WOLGA-MEDIA      [100.00GB]
                4                   WOLGA-TEMP       [900.00MB]

DS400[A]#
```

Figure 9-141 Verify all ports are disabled

4. Enable the ports by entering the command **initiator manage** WOLGA-SLOT4-(1) **allow all**. This will activate the ports for this initiator. The port access will be set to yes. Verify this with the command entered in step 3 on page 482.

```
DS400[A]# initiator manage WOLGA-SLOT4-(1) allow all
Access allowed to initiator WOLGA-SLOT4-(1) on all ports

DS400[A]#
```

Figure 9-142 Enable all ports

5. The other initiator WOLGA-SLOT3-(0) of the system WOLGA had one port enabled to install the operating system onto a logical drive from a TotalStorage DS400. The remaining three ports must be enabled as well. Check the ports with **initiator manage** WOLGA-SLOT3-(0) **info.**

```
DS400[A]# initiator manage WOLGA-SLOT3-(0) info

----[ Initiators ]-----------------------------------------------------


----[ Initiator 'WOLGA-SLOT3-(0)' ]-------------------------------------

                WWN:  210000e08b0ff748
              Alias:  WOLGA-SLOT3-(0)
                 ID:  4
        Port access:
                      Port A0        Port A1        Port B0        Port B1
-------------------------------------------------------------------------
                      YES            NO             NO             NO


       LUN Mappings:
                LUN                 logical          [Capacity]
-------------------------------------------------------------------------

                0                   WOLGA-BOOT       [ 30.00GB]
                1                   WOLGA-DB         [ 15.00GB]
                2                   WOLGA-LOG        [ 11.01GB]
                3                   WOLGA-MEDIA      [100.00GB]
                4                   WOLGA-TEMP       [900.00MB]

DS400[A]#
```

Figure 9-143 Initiator information

6. Enable the remaining ports with **initiator manage** `WOLGA-SLOT3-(0)` **allow A1 B0 B1**.

```
DS400[A]# initiator manage WOLGA-SLOT3-(0) allow all
Access allowed to initiator WOLGA-SLOT3-(0) on A1, B0 and B1

DS400[A]#
```

Figure 9-144 Enable initiator access to all ports

7. Reboot the server and verify after reboot that you see twenty drives in the device manager and five multipath device.

   IBM FAStT MSJ shows four targets and every target has five LUNs. This is as expected.

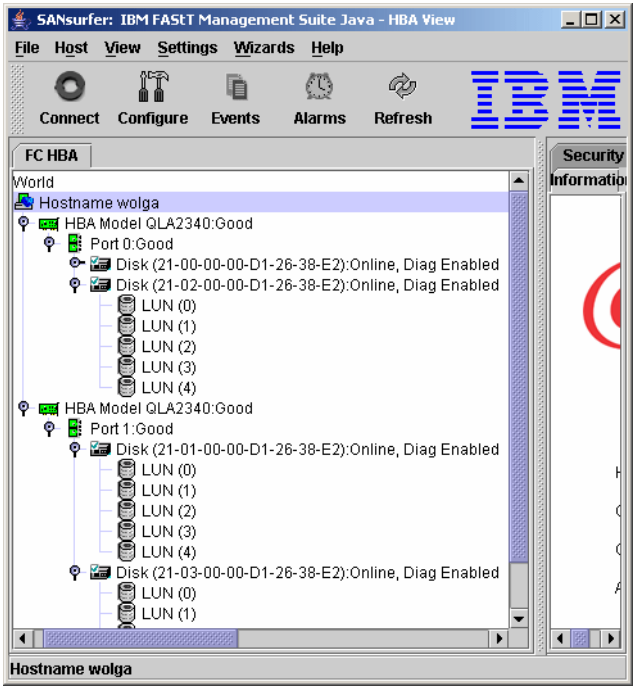

Figure 9-145 FAStT MSJ with 4 LUNs

Windows device manager shows now 20 **IBM DS400 S320 SCSI Disk Device**, five logical drives over 4 paths are 20 drives and for each logical drive defined on the DS400 one **IBM DS400 S320 Multi-Path Disk Device**.

Figure 9-146 Device manager with 5 drives and one multipath device

Target and LUN ID can be retrieved from the disk device property. Those values are also found in the properties of the multipath device.

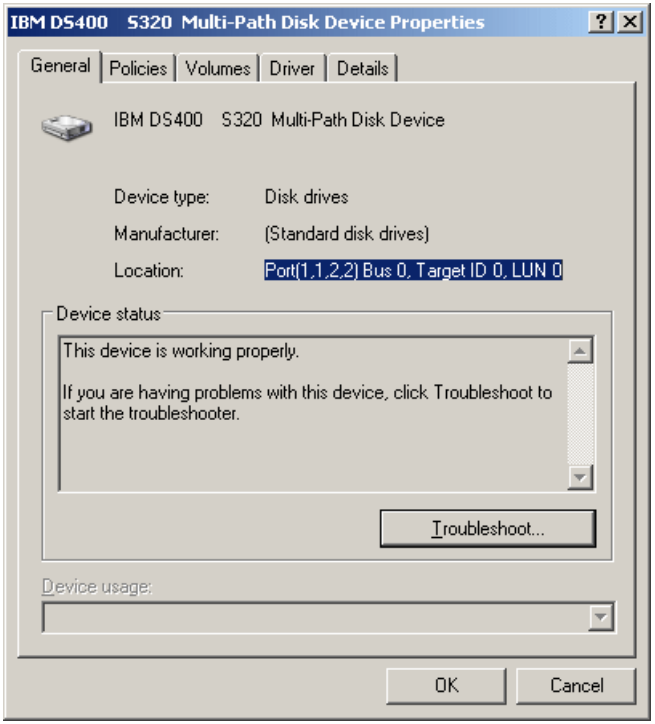

Figure 9-147 Multipath disk device properties

9.2.9 Finish the IBM TotalStorage DS4000 FC2-133 Host Bus Adapter BIOS setup

During the initial setup of the host bus adapter, only one boot device was defined.

To insure that the system boot can take place no matter which controller owns the array of the boot logical drive, the alternate controller must be defined as a boot target. The appropriate LUN must be selected as well. This settings needs to be defined also for the second host bus adapter to allow booting the system in case the first HBA fails. This will not be the case under all circumstances. There is no failover driver available, so it depends on the kind of problem if the system can use the second host bus adapter to boot from.

1. Reboot the system. During POST press **CTRL+Q** when the adapter displays the message “Press <CTRL+Q> for Fast!UTIL” to enter the host bus adapter BIOS setup.

Figure 9-148 POST messages of the host bus adapter

2. Select the first HBA and press **Enter**.

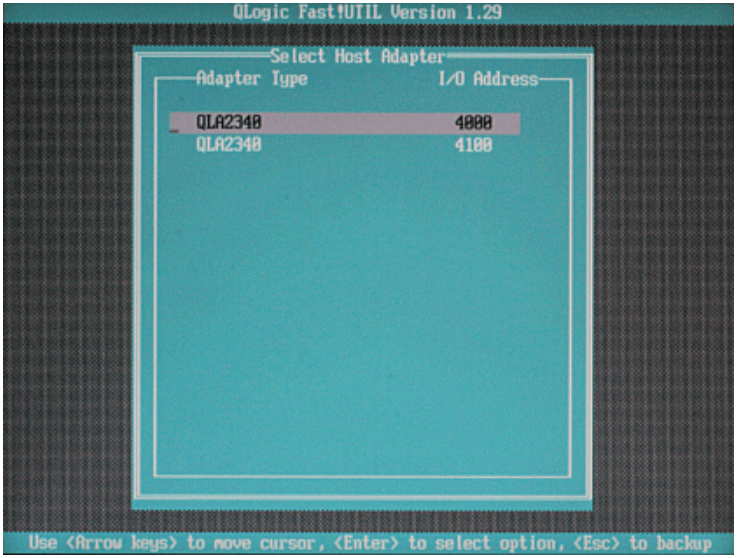

Figure 9-149 Select host adapter

3. Select **Configuration Settings** and press **Enter**.

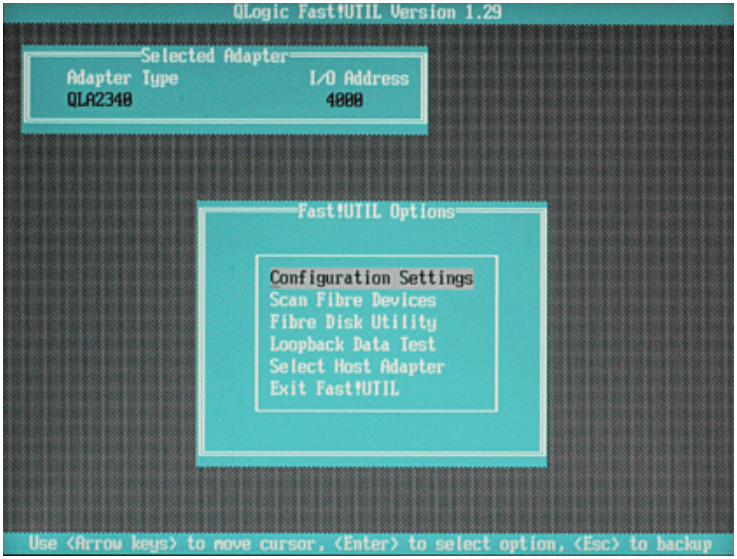

Figure 9-150 Fast!UTIL options — Configuration settings

4. Select **Selectable Boot Settings** and press **Enter**.

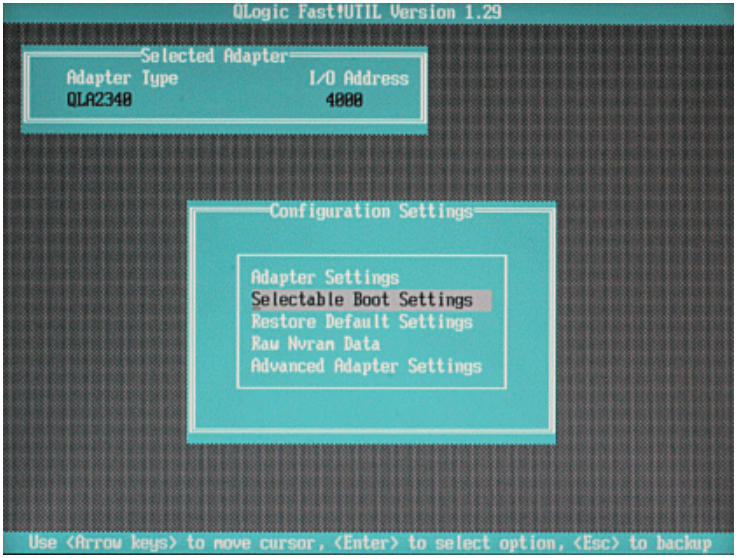

Figure 9-151 Configuration settings — Selectable boot settings

5. Select the **Boot Port name, LUN** after the primary entry and press **Enter**.

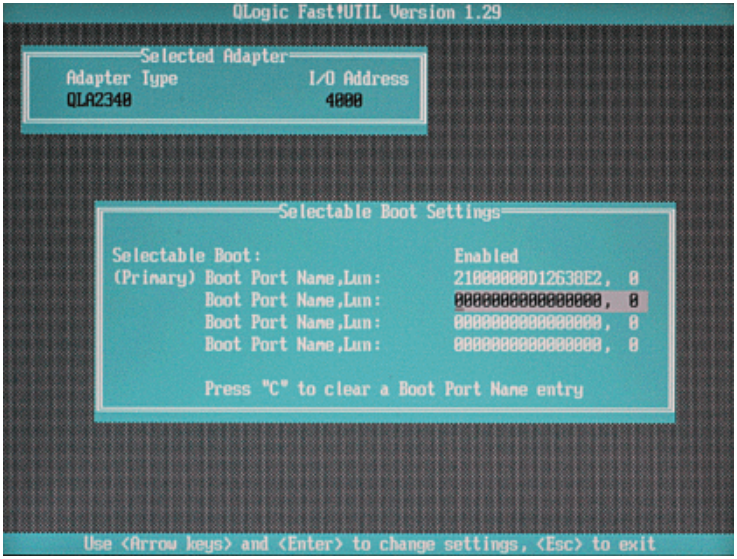

Figure 9-152 Second boot port name and LUN

6. Select the FC0 interface of controller B. From the world wide name the interface can be detected. Digit 3 and 4 of the world wide name is 02.

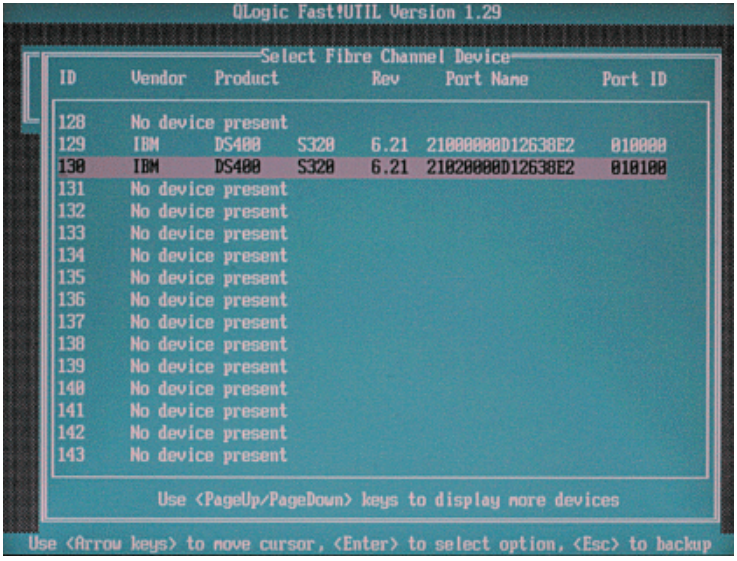

Figure 9-153 Select boot target

7. A list of all LUNs of this target presented. Select LUN 0 and press **Enter**.

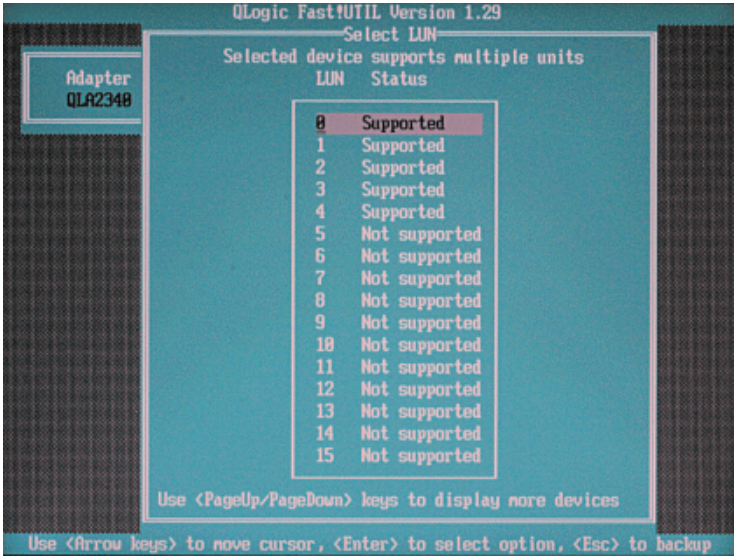

Figure 9-154 Select boot LUN

8. The target and the LUN is added to the boot device list. Press **ESC** to return to **Adapter Settings**.

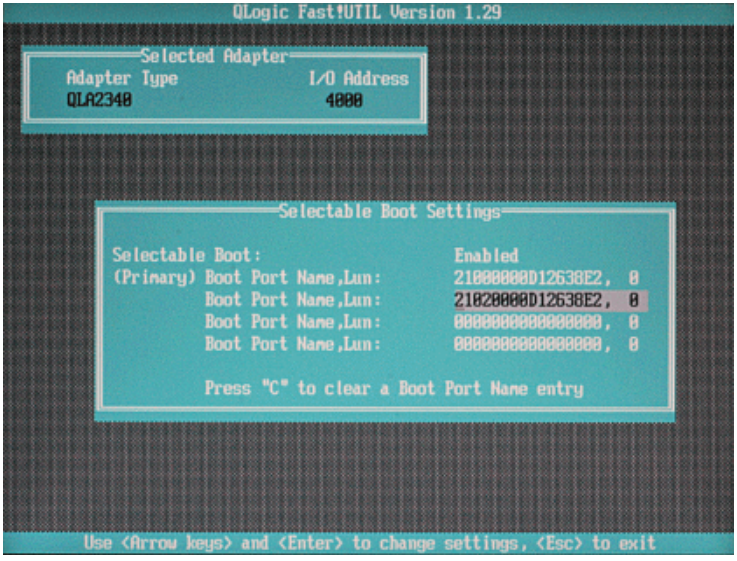

Figure 9-155 Defined boot targets and LUNs

9. Press **ESC** to leave **Adapter Settings** and return to **Configuration Settings**.

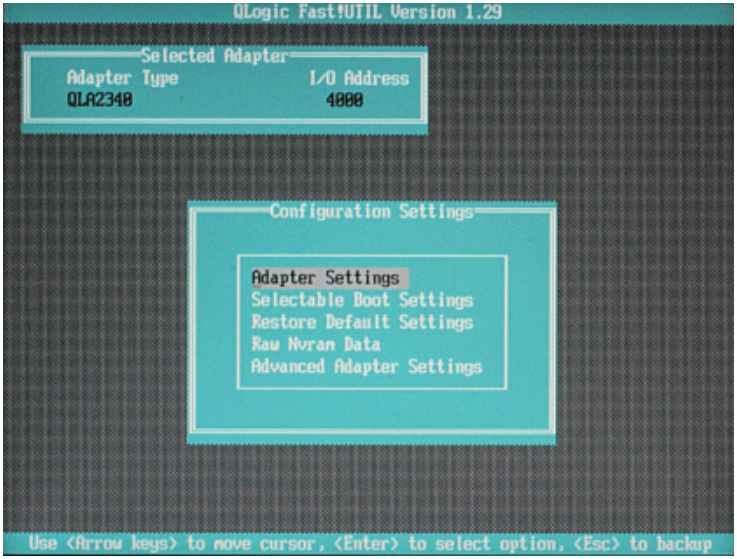

Figure 9-156 Select adapter settings

10. Save the modifications by select **Save changes** and press **Enter**.

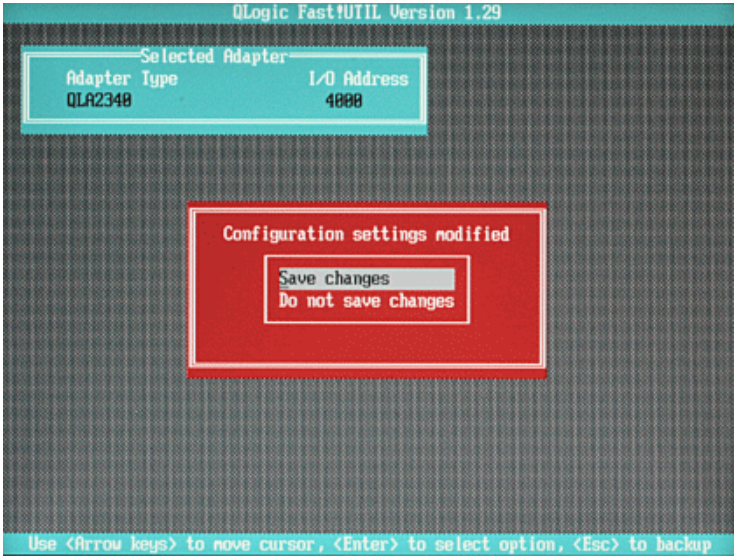

Figure 9-157 Save modifications

11. Select **Host Adapter** and press **Enter**.

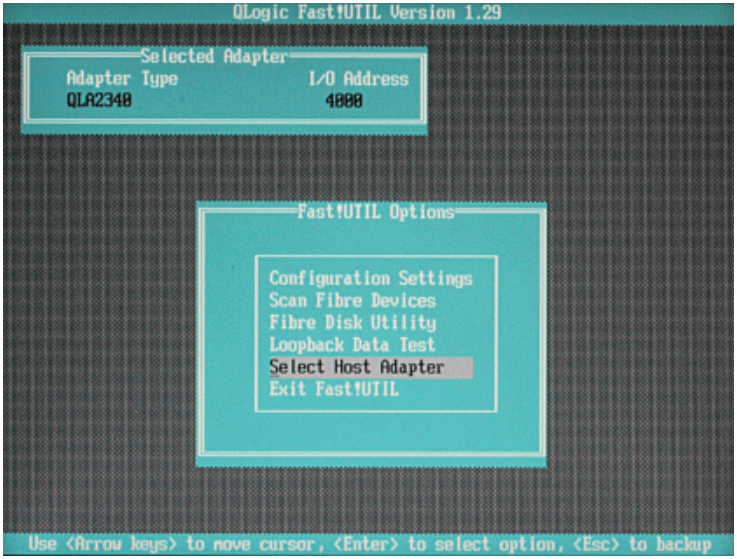

Figure 9-158 Fast!UTIL options — Select host adapter

12.Select the second HBA and press **Enter**.

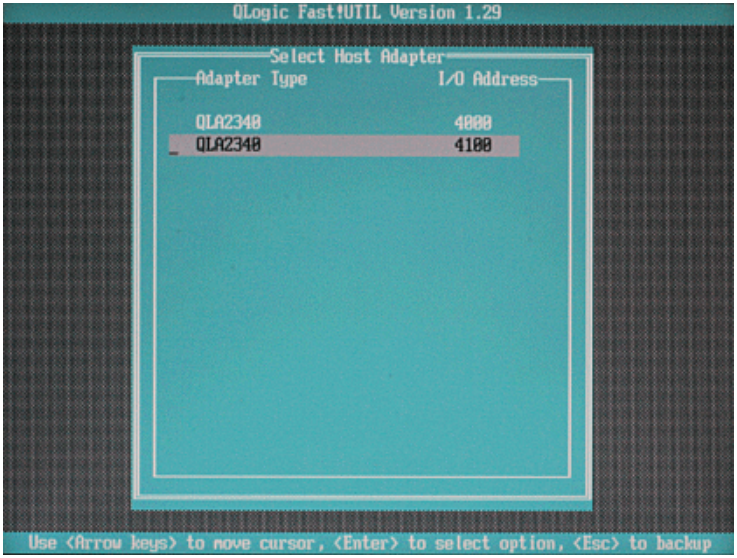

Figure 9-159 Select host adapter

13.. Select **Configuration Settings** and press **Enter**.

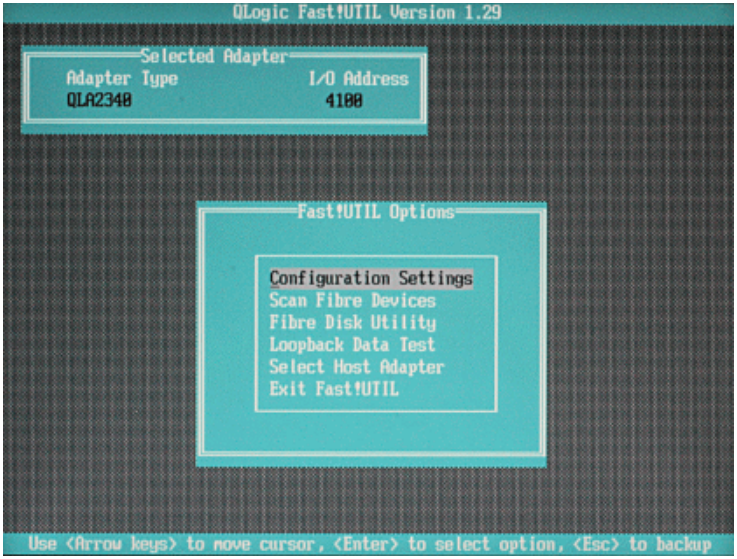

Figure 9-160 Fast!UTIL options — Configuration settings

14. Select **Adapter Settings** and press **Enter**.

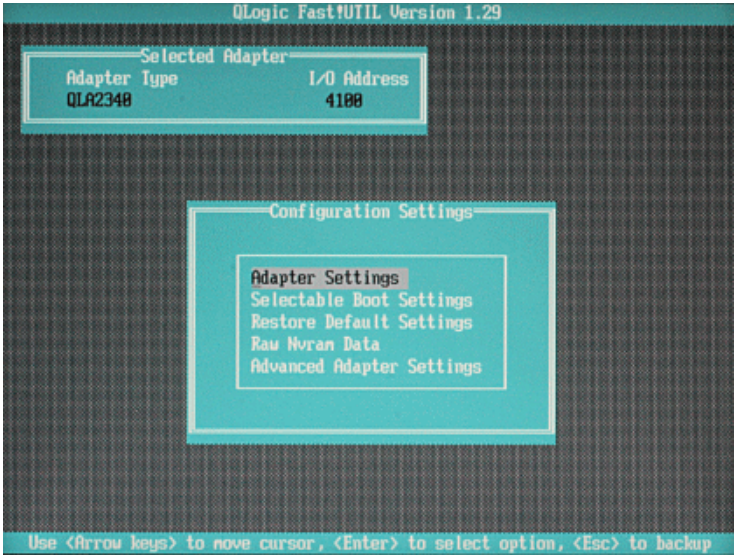

Figure 9-161 Configuration settings — Adapter settings

15. Select **Host Adapter BIOS disabled** and press **Enter**. The value will change from Disabled to Enabled. Then Press **ESC** o return to **Configuration Settings**.

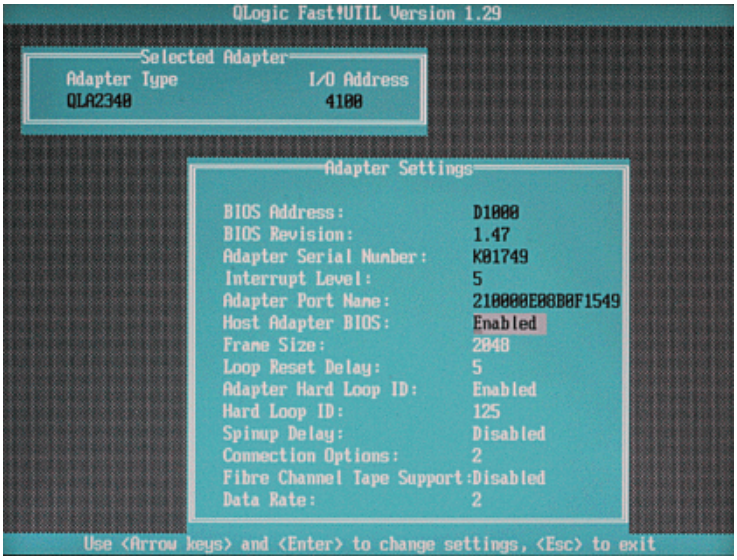

Figure 9-162 Enable host adapter BIOS

16. Select **Selectable Boot Settings** and press **Enter**.

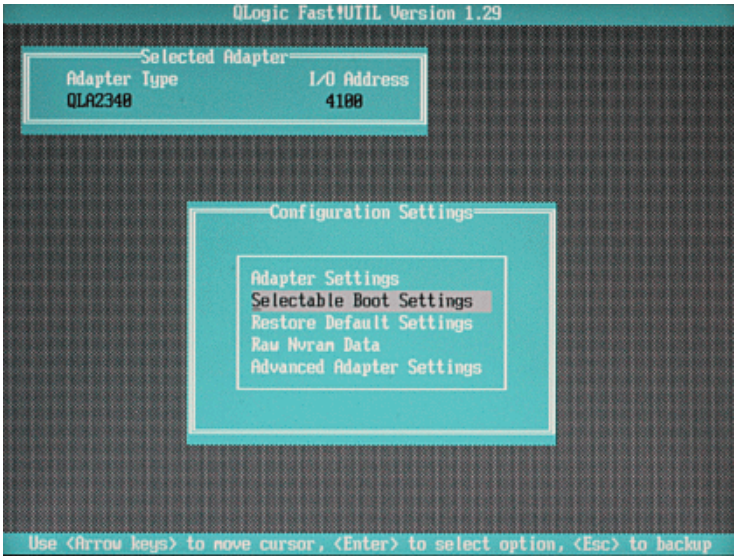

Figure 9-163 Configuration settings — Selectable boot settings

17. Select **Selectable Boot** and press **Enter**. The value changes from Disable to Enabled.

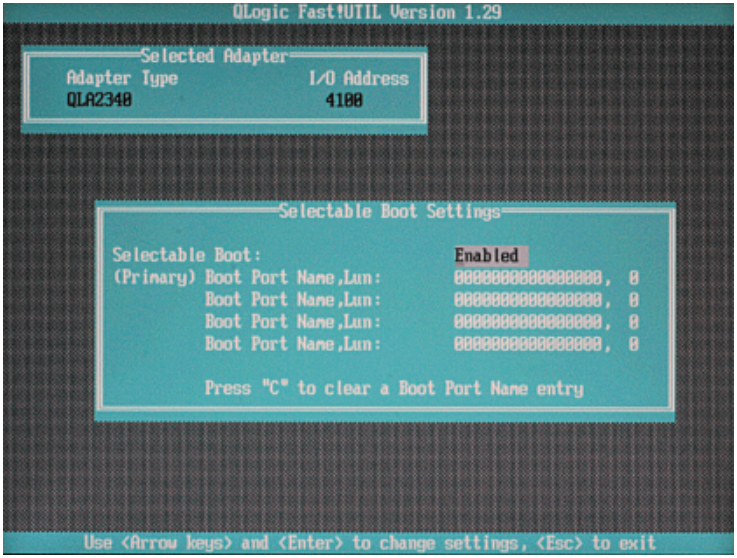

Figure 9-164 Enable selectable boot

18. Select the **(Primary) Boot Port name, LUN** and **Enter**.

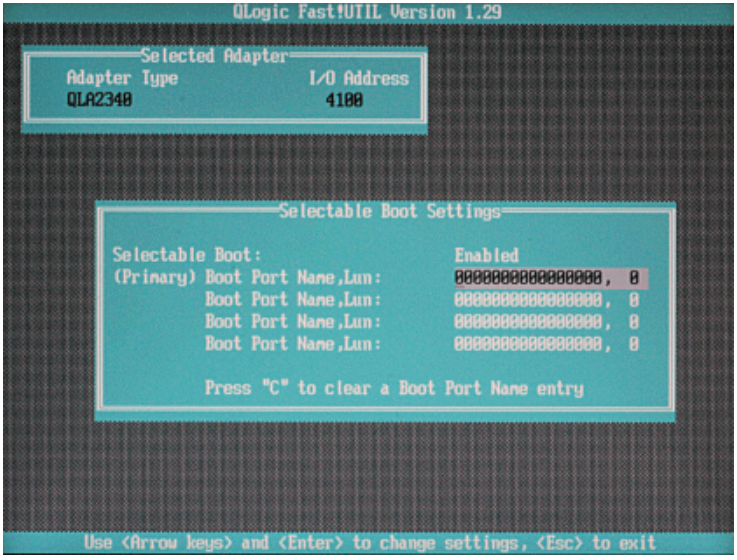

Figure 9-165 primary boot port name and LUN

19. Select the FC1 interface of controller A. From the world wide name the interface can be detected. Digit 3 and 4 of the world wide name is 01.

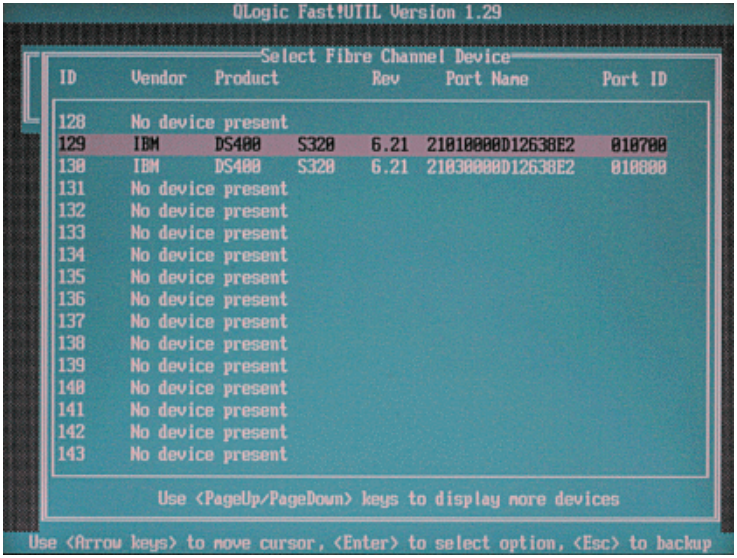

Figure 9-166 Boot target FC1 (01) controller A

20. A list of all LUNs of this target presented. Select LUN 0 and press **Enter**.

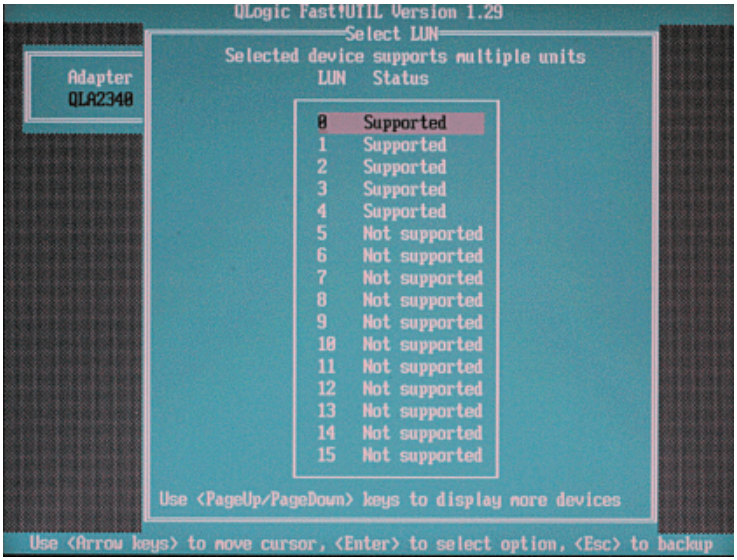

Figure 9-167 Boot LUNs

21. Select the **Boot Port name, LUN** after the primary entry and press **Enter**.

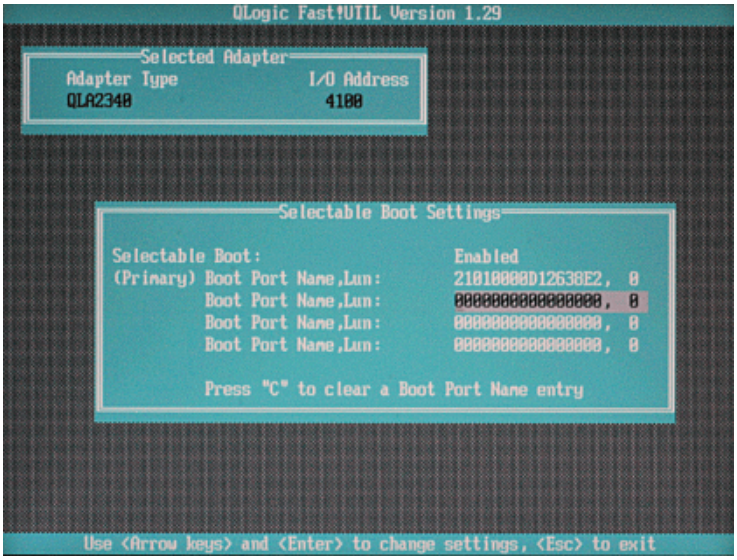

Figure 9-168 Second boot port name and LUN

22. Select the FC1 interface of controller B. From the world wide name the interface can be detected. Digit 3 and 4 of the world wide name is 03.

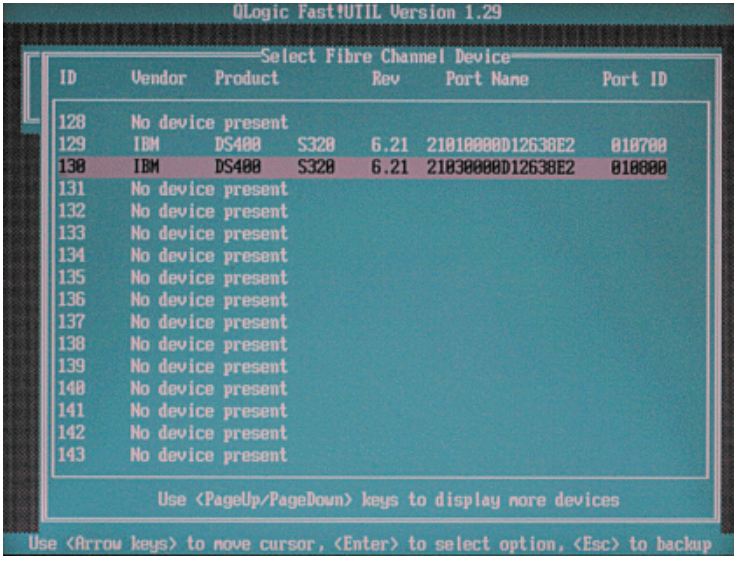

Figure 9-169 Boot Target

A list of all LUNs of this target presented. Select LUN 0 and press **Enter**.

.

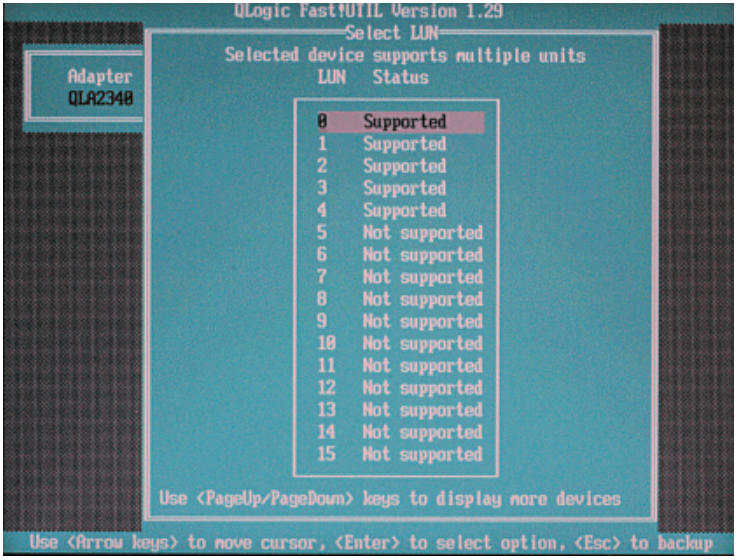

Figure 9-170 Boot LUNs

23. The target and the LUN is added to the boot device list. Press **ESC** to return to **Adapter Settings**.

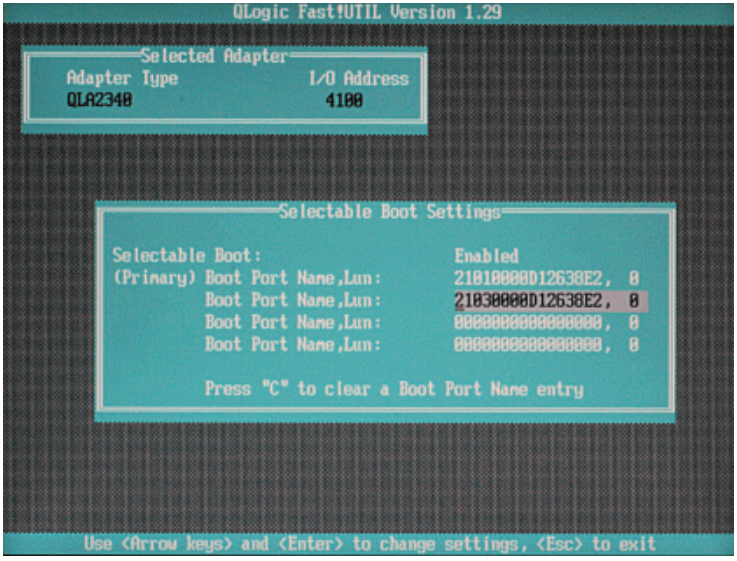

Figure 9-171 Defined boot targets and LUNs

24. Press **ESC** to leave **Adapter Settings** and return to **Configuration Settings**.

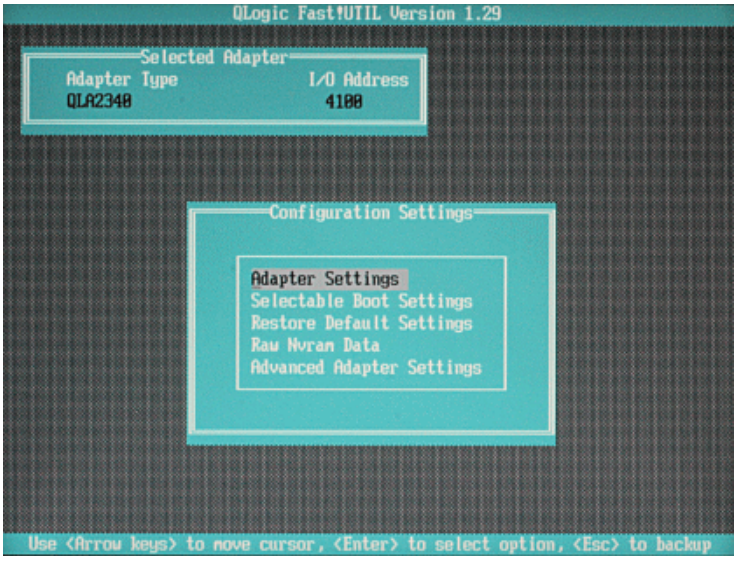

Figure 9-172 Configuration settings — Adapter settings

25. Save the modifications by select **Save changes** and press **Enter**.

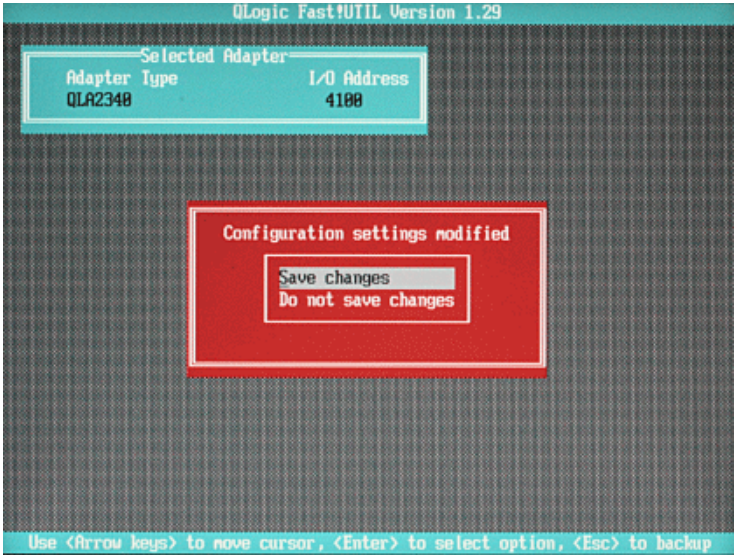

Figure 9-173 Save modifications

26. Select Exit Fast!UTIL and press **Enter**.

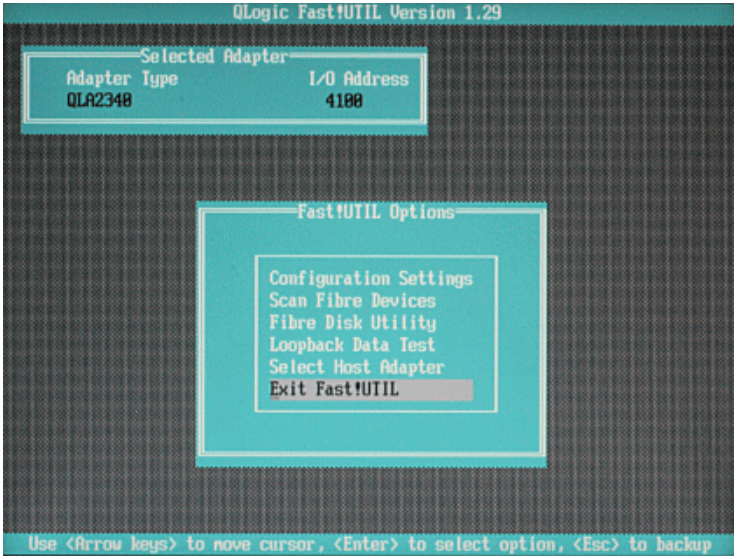

Figure 9-174 Fast!UTIL options —Exit Fast!Util

27. Select **Reboot System** and press **Enter**.

Figure 9-175 Reboot system

More logical drives may be added now. The setup is completed and the testing of the setup can start.

Part 4

Appendixes

A

Command line interface reference

You can manage the DS300 and DS400 via the ServeRAID GUI, or via the command line interface (CLI) utility. This is done by connecting to the controller using Telnet port 23, or the serial connection on the back of the controller.

The CLI is a Linux driven utility.

This appendix provides a list of the CLI commands available to configure and control the DS300/DS400 controllers.

When using the CLI, if at any point you do not know the rest of a command, you can press <TAB> and the controller will provide the options.

Important: Be extremely careful when using the Command Line Interface. Some commands can destroy arrays and logical drives as well as change the network settings. This could prevent you from communicating with the enclosure.

Note: The CLI prompt displays as DS400[x] or DS300[x] depending on the machine type. The [x] defines the controller (either A or B).

A.1 Overview of CLI

The CLI software is a terminal-based command line interface that allows you to monitor and manage the IBM TotalStorage external storage subsystem. iSCSI targets and FC volumes that will appear as local disk storage on your server can be created and managed. Comprehensive event logging features allow you to monitor the entire.

A.1.1 Accessing the CLI

You can access CLI two ways. Each method presents CLI automatically:

- Attach a VT100 terminal emulation utility (for example, hyperterminal) to the serial port of the storage subsystem.

  Use these terminal emulation program settings to access the CLI:

  - Null Modem Cable
  - Baud Rate:115,200
  - Data Bits: 8
  - Parity: None
  - Stop Bits: 1
  - Flow Control: Off

- Connect to the storage subsystem through a Telnet session (once the initial network values have been set.)

  Telnet to the IP address that is assigned to the Management port on the storage subsystem.

Note: The factory default IP address for the management port for Controller A is 192.168.70.123 (ETH0[A]). The management port for Controller B defaults to 192.168.70.124 (ETH0[B]).

A.1.1.1 Access Levels and Passwords

The CLI has two levels of access: Operator or Administrator. From the Operator level you can monitor the current status of the system, but you cannot make any configuration changes. To set the Operator password type:

```
password operator <Enter>
new password ***** <Enter>
confirm new password ***** <Enter>
```

Note: You must be in Administrator level to set passwords.

From the Administrator level, you have full access to all the features of the software and can make changes to the configuration of the system.

Note: The default administrator password is “passw0rd”. (Note the zero instead of the letter O.)

To set the Administrator password type:

```
Password administrator <Enter>
new password ***** <Enter>
confirm new password ***** <Enter>
```

To switch from operator to administrator level, type `administrator` and enter the password. To switch from administrator to operator level, type `operator`.

A.1.2 CLI commands

This section provides an overview of what CLI commands are and how they are used.

A *command* is like an area of functionality (for example, array) followed by an action (for example, create), and possibly also followed by a list of parameters (for example, disk names).

Parameters may be in a fixed or variable order. When parameters are in a fixed order, the CLI uses their position on the line to determine what they signify. When parameters are in a variable order, they are preceded by a keyword that the CLI uses to determine what they signify.

The *manage* commands are distinctive, because they take as a parameter an item to be acted upon (for instance, a logical drive or a disk) and apply to it other commands with their own sets of parameters. For example:

```
port manage port topology topology
```

where `topology` may be considered a separate command within the `manage` command.

The commands available to you depend on the level at which you are working (operator or administrator) and in which context the CLI is running.

When you are in a context other than the global context, you don't have to enter the part of the command that defines the context (or, in some cases, the object being worked on).

Many commands take lists, which are terminated by the end of the line or by another keyword.

A.1.2.1 Character Case and White Space

The CLI is not case-sensitive. Commands and parameters can be typed using any combination of upper and lower case letters 1. All IQNs, iSCSI Aliases, and initiator names must be entered in lower-case letters.

Commands and parameters are separated with one or more spaces and/or tab characters (white space). White space at the start and end of a command line is ignored.

If a white space character has to be entered as part of a parameter, enclose the parameter in a pair of quotes (" "). Everything between the quotes, including any enclosed white space, is used as the parameter (the quotes are not used as part of the parameter).

A.1.2.2 Command Completion and Abbreviation

You can request a command completion at any time by pressing the Tab key.

If, after considering what has been typed on the line up to the tab, there is only one word that can legally be entered, that word is added to the end of the line. In cases where several words can be entered, these are listed and the line redisplayed. If these words have a common root, it is added to the line. If there are no words that can be legally added, the line does not change.

Sometimes, the complete list of options cannot be printed, such as with a range of numbers, or a class of objects. In this case, a description enclosed in < > is shown (e.g. <name>).

Sometimes, when a list of numbers is given, a description will be added in angle brackets to make clear what is being asked for (e.g., <stripedepth>, as shown in the example below):

```
dual-ip3[B]# array create a2 raid5 <TAB>
1024  16   32   64
128   256  512  <stripedepth>
```

Most commands and parameters can be abbreviated to their shortest unique abbreviation. For example, if you are in the Operator level and want to change to the Administrator level, all you have to type is ad as there are no other commands starting with "ad".

A.1.2.3 Command Format in This Guide

In this Guide, a value shown in italics describes the type of the value. Otherwise, the string itself is the value. For example, *int* means any integer, whereas `int` means the string `int`.

If a parameter is in [ ], it is optional. If a command's parameters are all optional, at least one must be supplied unless the complete parameters list is also enclosed in [ ].

A set of values joined with a | means that only one of those values may be given.

A.1.2.4 Command Line Editing

Commands and parameters can be edited as you type by using common keystrokes. The next Table lists the editing keystrokes that can be used.

Keystroke	Function
Ctrl-A	Jump to the start of the line
Ctrl-B	Move cursor back one character (left)
Ctrl-C	Escapes and terminates lengthy tasks
Ctrl-D	Delete the character at, or in front of, the cursor
Ctrl-E	Jump to the end of the current line
Ctrl-F	Moves cursor forward one character
Ctrl-H; Backspace	Erases a mistake when entering a command
Ctrl-K	Delete from the cursor forward to the end of the line
Ctrl-N; down-arrow	Display the next command in the command history
Ctrl-P; up-arrow	Display the previous command in the command history
Ctrl-S	Search for a command in the command history
Ctrl-U	Delete complete line
Esc B	Move cursor back one word
Esc D	Deletes from cursor to end of the word
Esc F	Moves cursor forward one word
Tab	Complete the command or display alternatives if there is more than one possible command

A.1.2.5 Help

Command completion can be requested at any time by pressing the Tab key. This gives a list of the words that may be added to the command line, taking into account what has already been typed.

You can request help at any time by typing help or ? anywhere in a line. The depth of help information provided (e.g., valid commands or parameters, or all available commands) depends on where you are in the CLI.

A.1.2.6 Viewing Long Reports

When a command generates a report (including help messages) that is more than one page long, the first page is displayed. Other parts of the report may be viewed by using the keystrokes shown in the next Table.

Keystroke	Function
B	Previous page
D or J	Down to next line
E or >	Advance to end of report
H or ?	Display this list
N or F or <Enter> or <Space>	Next page
P or <	Back to first page
Q	Quit report
R or <F5>	Redraw current page
U or K	Up to previous line

A.2 CLI parameter glossary

This section provides definitions for parameters used in the rest of this appendix to describe the commands and options of the CLI.

absolute_date

This defines an absolute time and date in the range 1970 to 2036. Unless it is explicitly mentioned, the time is relative to the current time zone.

The absolute_date parameter can include date and time, date only, or time only. When no time is entered, the system assumes a time of 00:00:00. When no date is entered, the system assumes the current date.

The date has three fields, day, month, year, which can be expressed in any of the common formats, with a `/' character separating each field:

- Month may be a number, an abbreviated name of 3 characters, or the full name.
- Year must consist of 2 characters.

If you use numbers for both the day and the month, a day number greater than 12 is an error. If you use a format that is unambiguous (that is, uses the name of the month rather than a number representing that month), you can express the date in any order (day/month/year, month/day/year).

Time is expressed as 00:00:00 (hours:minutes:seconds). The seconds field is optional.

Note: If a time is given with no seconds then 0 seconds is assumed.

Time may be supplied either before or after the date, separated from it by a single '@ ' character, as shown in this example:

```
06:00:00@25/07/04
```

Note: If a date is given without a time, the time is assumed to be 00:00:00. If a time is given without a date then the current date is assumed. If a time is given with no seconds then 0 seconds is assumed.

bytes

The number of bytes may be an integer or a real number followed by a multiplier letter (case is ignored) to multiply the given number by powers of 1024. The multipliers are:

k	1024
M	10242
G	10243
T	10244

ip_host

A network node that can be defined using either its IP address or host name. If a host name is used, it is resolved to its IP address to ensure that the name is valid.

ip_mask

An IP number and a netmask. This parameter can take two forms:

- ip_numb/mask

  where /mask is a netmask size (see netmask)

- ip_numb netmask

  where ip_numb defines the IP address and netmask defines the netmask (see netmask)

ip_numb

An IP number in the conventional dotted quad format, as shown in this example:

```
192.193.194.195
```

level

Some commands take an event log level (also known as a severity). These levels are listed here from most severe (highest) to least severe (lowest).

- Fatal
- Serious
- Error
- Warning
- Information
- Audit

*_lists

Some parameters are described as lists, which are supplied by entering the list of elements separated by spaces. Normally, the list is the last thing on a line; when it's not, the next parameter must be introduced by a fixed label to distinguish it from being part of the list.

netmask

A netmask splits an IP number into a network part and a host part. A netmask consists of a bit mask that starts with a set of bits at 1 followed by a set of bits at 0. It may not have 1s and 0s mixed up.

A netmask can be represented two ways. The examples shown here refer to the same netmask:

/23	The number of bits set in the network part.
255.255.254.0	The set bits displayed in the same way as an IP number

percentage

This is a number followed by a '% ' sign. Percentages allow both integers and numbers with decimal points.

relative_period

This defines a period of time that is usually added to or subtracted from a time or date.

Note: Application of the relative_period parameter is described in the relevant command section of this appendix.

The relative_period is expressed as a number followed by an upper or lower case letter which indicates the units to which the number refers. The letters are:

S	Seconds
M	Minutes
H	Hours
D	Days
W	Weeks

For example:

90m (or 90M) expresses 1 hour 30 minutes.

3h (or 3H) expresses 3 hours

Only a single value can be entered.

When the relative period is to be subtracted, the sign of the relative_period parameter is ignored. When the relative period is to be added, the sign of the relative_period parameter is used; however, an error is generated if it is illegal to have a negative relative period for that command.

A.3 Setup Wizard

The Setup Wizard runs if the system is in factory default state when the CLI is started. You can also run it with the setup command. The Setup Wizard helps you enter the information required to connect the system to a local network.

The Setup Wizard steps you through a set of questions. In addition to typing in the answers to the questions, you can also use three commands, as shown in this Table:

Command	Function
Abort	Quit the Setup Wizard without changing anything
Back	Go to previous question
Help	Describe the current options

If a question has a predefined or default answer, the answer appears in square brackets [ ]. Some options take an arbitrary string, in which case the commands back and abort cannot be used. In those circumstances, enclose the answer in double-quotes " " so that the answer will be accepted even if it is the same as a command.

hostname

Supply the name by which the system is to be known.

The special BACK command here repeats this prompt.

domain name

Supply the domain in which the system belongs.

SAN name

Supply the Storage Area Network in which the system belongs.

timezone

Supply the system time zone.

system Date

Supply the system date. The format is as for the absolute_date (see absolute date).

define administrator password

Define whether an administrator password is required. If a password is required, you are prompted for that password, which must be supplied twice.

Press Ctrl to leave the password unchanged.

Press Return (or Enter) twice to require no administrator password.

define operator password

Define whether an operator password is required. (See define administrator password for instructions on entering this information.)

dhcp

Reply yes or no to determine whether DHCP should be used to discover the system's network values.

If you enter yes, the wizard takes you to the Save new configuration question (see save new configuration). If you enter no, the wizard continues with questions about the system's main ethernet interface.

IP address

Enter the IP number to use in the format described in ip_numb.

Alternatively, the netmask (see netmask) may be defined by entering its width after the IP number. (See the interface manage commands for more information.)

netmask

Define the size of the netmask. Any of the common formats are accepted, as described in netmask.

Note: If you entered the netmask as part of the previous IP address, this question is skipped.

broadcast address

Enter the address to be used for broadcasts. You can use any of the forms described for the `interface manage interface broadcast` command.

default gateway

Enter the name or IP number of a router to which packets for destinations that do not have routes defined for them in the system should be sent. Enter None to clear the default gateway.

DNS servers

Enter the IP numbers of up to 3 systems that will act as DNS servers. Enter None to clear the list of DNS servers.

save new configuration

When the Setup Wizard has finished, the system is configured but the new configuration is not saved immediately in non-volatile memory. Enter yes to save the configuration immediately, or no to defer saving the configuration.

The `save` command saves the configuration.

Finish?

Enter `yes` to apply the values you have entered, or `no` to return to the first question and change the values you have entered.

A.4 Administrator command

Use this command to enter administrator level. If a password has been defined, you are prompted for that password.

This command is a global command and is always available.

Note: The factory default password is `passw0rd` (note the number zero not the letter O). For security reason you should change the PW.

A.5 Array commands

These commands control device arrays. They are available after the global `array` command, or as global commands by preceding them by `array`.

Command	Parameters	Levels
create	name [remote] raid_level [options] device_list	administrator
migrate	[array] raid_level [stripe_depth] device_list	
Destroy	unused_or_alien_array_list (requires confirmation)	administrator
export	name_list (requires confirmation)	administrator
failback		administrator
Identify	[start | stop] name_list	all
import	foreign_array_list	administrator
info	[name_list]	all

Command	Parameters	Levels
list	[local | foreign | virtualized | alien | remote]	all
manage	name (See Array manage array commands)	all
move	local | remote name_list	administrator
Rebuild	RAID_array device_list	administrator
Unused		all
Verify	start | stop	administrator
Verify	status	administrator
Verify	rate high | low | medium	Administrator

create *name* [remote] *raid_level* [options] *device_list*

Create a new array called `name` with no RAID and enter a list of devices to add to the new array. The `name` must not be the name of any other array.

create name raid0 stripe-depth device_list

Create a RAID 0 array with the specified stripe depth and devices. You must specify at least two devices.

create name raid1 device_list

Create a RAID 1 array with the two specified devices. You must specify exactly two devices.

create name raid5 stripe-depth device_list

Create a RAID 5 array with the specified stripe depth and devices. You must specify at least three devices.

create name raid10 stripe-depth device_list

Create a RAID 10 array with the specified stripe depth and devices. Devices are automatically grouped into pairs, and therefore must always be an even number.

create name raid50 stripe-depth device_list group num_group

Create a RAID 50 array with the specified stripe depth and devices. You must specify at least six devices and arrange them into groups. The number of groups is specified by the group parameter.

migrate name raid0 stripe-depth device_list

Migrate a RAID 0 array with the specified stripe depth and devices. You must specify at least two devices.

migrate name raid5 stripe-depth device_list

Migrate a RAID 5 array with the specified stripe depth and devices. You must specify at least three devices.

migrate name raid10 stripe-depth device_list

Migrate a RAID 10 array with the specified stripe depth and devices. Devices are automatically grouped into pairs, and therefore must always be an even number.

destroy unused_or_alien_array_list

Destroy the listed arrays. When the arrays are destroyed, the disks and RAID groups associated with them become available for use.

Note: Arrays assigned to logical drives cannot be destroyed. Foreign arrays must be imported before they can be destroyed.

An `alien array` is an array with an incomplete or corrupted signature that cannot be used.

export name_list

Export the arrays in `name_list`. When these arrays are exported, the disks in those arrays are no longer available.

If all of the disks in a array are transferred to another system, they can be used in that system after you complete an `array import`. (See import name_list.)

array failback

Return all storage to its original controller, if it has been moved (failed over). The original controller, from which the array was removed, must be functional again before you use this command.

identify [start | stop] name_list

Identify the devices that make up the arrays in `name_list`, or all of the disks that are assigned to any array if `name_list` is not supplied.

Disks are identified by flashing LEDs; flashing continues for 30 seconds or until an `array identify stop` or `device identify stop` command is given for that disk.

import foreign_array

Import a set of disks that have been assigned to an array in this or another unit and subsequently exported using the `array export` command.

Note: All of the disks that make up the array must be present for this command to be successful.

info [name_list]

Show information about arrays. If a `name_list` list is not supplied, all of the arrays are included. Specify arrays by providing a `name_list` list.

For a single array, this command is the same as the `array manage name info` command.

list [local | foreign | virtualized | alien | remote]

List all of the known arrays. (Minimal information about each is also included.) If no parameters are entered, all arrays are listed. You can also limit the list to only foreign, or alien, or only local arrays.

move local | remote array_list

Change the ownership of an array to either the remote or local controller. If the `remote` keyword is specified, only arrays on the local controller may be specified. If the `local` keyword is specified, only arrays on the remote controller may be specified.

rebuild name device_list

Rebuild an array that has had devices replaced. The new set of devices must be specified in its entirety.

unused

List all of arrays that are not associated with a logical volume.

A.5.1 Array verify commands

These commands control array verifications.

Command	Parameters	Levels
verify	start | stop	administrator
status		administrator

Command	Parameters	Levels
rate	high | low | medium	administrator

array verify start | stop

Begins a parity check and corrects parity information where deficient. All arrays are checked.

array verify status

Indicates whether a verification check is in progress, and reports the speed at which it is running.

array verify rate high | low | medium

Sets the speed of a running verification check.

A.5.2 Array manage array commands

These commands control a specific array. They are available after the global `array manage array` command, or as global commands by preceding them with `array manage array`, or from the array context by preceding them by *manage array*.

Command	Parameters	Levels
cache	read enable | disable	administrator
identify	[start | stop]	administrator
info		all
spare	add unused_device_list | list | remove device_list	all
task abort		administrator
task rate	high | medium | low	administrator
verify		administrator

identify [start | stop]

Blink the LEDs of the disks which make up this array.

cache read enable | disable
Either enable or disable the hardware cache for this array. The global Cache command settings take precedence over that of the arrays'.

info
List details about the array, such as its type and which devices it includes.

spare add unused_device_list | list | remove device_list
Add one or more spare devices, list all spares, or remove the specified spares.

task abort
Aborts whatever task is running on the array.

task rate
Sets the speed of whatever task is running on the specified array (building, rebuilding, resynchronizing, or verifying).

verify
Verifies the specified array.

A.5.3 Array spare commands

These commands control the management of global hotfix devices.

Note: Use the array manage array spare commands to manage hotfix devices for individual arrays. (See cache boolean.)

Command	Parameters	Levels
add	device_list	administrator
list		all
remove	device_list	administrator

add device_list
Add the specified devices to be used as global hotfix devices.

list
List the global hotfix devices.

remove device_list

Remove the list of devices from the global hotfix devices.

A.6 Authentication commands (iSCSI only)

Use the authentication commands to manage user access to the system's targets.

You can use these commands after the authentication command with no parameters, or as global commands by preceding them with `authentication`.

Command	Parameters	Levels
add	name	administrator
info		all
list		all
password	name	administrator
remove	name_list	administrator

add name

Add a new user called name. Enter the password when prompted.

info

List information about authentication options.

list

List the known users.

password name

Change the password for the specified user name. Enter the password when prompted.

remove name_list

Remove the listed users.

A.7 Cache commands

These commands manage the caching policy of arrays.

Command	Parameters	Levels
info		all
policy	array | logical boolean	administrator

info

Output information about caching policies.

policy array [enable | disable]

Set the caching policy for arrays. This setting overrides the cache settings of the arrays.

A.8 Controller commands

These commands output device, array, and interface information, and manage the settings of the current controller on a dual-controller machine.

Command	Parameters	Levels
info		All
set	controller	Administrator
Ipfailover *iSCSI only	local | remote | both | none	Administrator
peer	enable | disable	Administrator

info

Display information about the current controller, as shown in this example:

Example: A-1 Controller info command

```
----[ Controller ]--------------------
Logged into controller:  B
Working on controller:  B
Peer controller is:  Present (not held in reset)
Peer controller status:  Ready
```

```
iSCSI IP failover policy:  remote
```

'Logged into controller' means the controller that the Telnet or serial session was started on.

The 'Working on' controller can be changed with the `controller set` command.

controller set controller

Set the current controller. The controller value can be A or B. The current Telnet or serial session will remain logged into the controller it started with.

ipfailover local | remote | both | none

(iSCSI only) Set the policy for the failover of iSCSI interfaces, using these options:

`local` - if an interface fails, an interface on the same controller will be chosen if possible (ETH2 → ETH3 or ETH3 → ETH2).

`remote` - if an interface fails, an interface on the other controller will be used (ETH2 → ETH2 and ETH3 → ETH3)

`both` - if an interface fails, attempt to failover to an interface on the same controller (ETH2 → ETH3 or ETH3 → ETH2). If an interface on the same controller cannot be used, an interface on the other controller will be used (ETH2 → ETH2 and ETH3 → ETH3).

`none` - no failover of iSCSI interfaces will be attempted.

Note: The default setting is Remote failover option.

peer enable | disable

Specify `enable` to bring up the other controller.

Specify `disable` to take down (halt) the other controller. This takes a while, so you must keep looking at 'controller info' to check on the status of the other controller.

A.9 Device commands

These commands control device discovery. You can use these commands after a device command with no parameters, or as global commands by preceding them by `device`.

Command	Parameters	Levels
identify	[device_list] [start | stop]	all
info[	device_list]	all
initialize	[local_device_list]	all
list	[foreign] [spare]	all
manage	device (See Device manage device commands)	
unused		all

identify [device_list] [start | stop]

Identify one or more devices by blinking their LED.

The default action is to start the LED blinking and return to the prompt. The LED will stop blinking after about 30 seconds or when a `device identify stop` or `array identify stop` command is given for that device.

To identify all of the disks in an array, use the `array identify` command.

info [device_list]

Output information about some or all devices. (You can specify which devices by entering a parameter for `device_list`.)

This command is identical to the device manage device info command (see Device manage device commands) when only one device is specified.

initialize [local_device_list]

Remove all array membership information for the local devices in `local_device_list`.

list [foreign] [spare]

List the known physical and/or spare devices. All devices are listed if no parameters are specified.

unused

List any devices which have not been allocated to an array volume.

Device manage device commands

These commands control devices that are found using device discovery. You can use these commands after a `device manage device` command with no parameters, or as global commands by preceding them with `device manage device`.

Command	Parameters	Levels
info		all
identify	[start | stop]	all

info

List details about this device.

identify [start | stop]

Identify the specified device. (See identify [device_list] [start | stop] for a description of options.)

A.10 Diagnostics commands

These commands run diagnostics tests. They are available after a `diagnostics` command with no parameters, or as global commands by preceding them with `diagnostics`.

Command	Parameters	Levels
loopback *FC only	port continual | iterations [stop]	administrator
dump	xmodem | 1kxmodem	administrator
ping	host [timeout time] [count number] [interface interface]	administrator
tracert	host	administrator

loopback FC port iterations | [continual] [stop]

Perform a loopback test on the specified FC port.

The `iterations` parameter specifies the number of times this test is performed.

Use the `continual` option to have the test performed continuously.

The `stop` option will cause the test to be halted immediately an error is detected.

Press **Esc** to abort the test. Press **Return** (or **Enter**) while the test is running to output ongoing results.

dump [1kxmodem] [xmodem]

Send a diagnostics dump. The `diagnostic.bin` file is created and sent to the host.

This file is the same file that is generated from the ServeRAID manager when "Save Support Archive" is selected from the controller pop-up menu.

To save the file, do the following from a CLI prompt using the left most Serial Port:

```
DS300[A] (diagnostics)# dump 1kxmodem
Creating the dump file: 145.10kB
```

Issue the `1kxmodem` command to start the transfer → at this time select "receive" from the Terminal application menu (Hyperterminal).

```
Diagnostics dump was successfully sent
```

From a Telnet session:

```
DS400 (diag)# dump xmodem
```

→ press **Enter**

```
Creating the dump file:  0kB
```

→ at this time select "receive" from the Telnet File menu

```
Creating the dump file:      0 B  Diagnostics dump was successfully sent
```

ping [timeout time] [count number] host [interface]

Check the connectivity from a port to the specified network node. Each ping attempt result is reported immediately. A line is output for each ping attempt.

The optional `timeout` option defines the number of seconds to try each ping attempt (default 1 second). The `time` parameter is the number of seconds.

The optional `count` parameter defines the number of times to attempt the ping if an attempt fails (default 1)

The command returns immediately if a ping attempt succeeds; otherwise, the delay is approximately the product of `time` and `number`.

tracert host

Trace the route to a host on the network. You can terminate this command with Ctrl-C.

A.11 End command

If you are in a context, use this command to leave that context and return to the previously active context.

Note: This command is not valid from the global context.

A.12 Events commands

These commands control reporting events. You can use these commands after the events global command with no parameters, or as global commands by preceding them with events.

Command	Parameters	Levels
clear		administrator
inactive	all | parameters (See Events inactive Commands.)	administrator
info		all
report	all | parameters (See Events report commands.)	

clear

Clear the event log.

info

Display these values associated with the event log:

- Size of the event log
- List of actions for each level
- Current status which means the highest active severity level

A.12.1 Events inactive commands

These commands are used to remove the active state from the selected log entries. The commands can be executed as global commands by preceding

them by `events inactive` or after the events global commands by preceding them by inactive.

Except for the `all` command, any number of these commands can be combined on a line to define the set of events that will be selected. Each command restricts the number of events that are selected.

The selected events will be listed and you will be asked for confirmation before they are set inactive. At least one command must be supplied for a report to be generated.

Command	Parameters	Levels
all	Cannot be combined with other parameters	administrator
from	absolute_date | relative_period	administrator
level	See list of valid levels.	administrator
originator	name	administrator
to	absolute_date | relative_period	administrator

all

Select every event. It must be the only command on the line.

```
from absolute_date | relative_period
```

Only report events which were generated at or after the time and date specified by the supplied value. All the reports from the specified time and date to the present will be reported unless the for or to commands further restrict the reported events to a range of dates and times.

The `relative_period` is the period before the current time and date.

level

See list of valid levels.

originator name

Only select events that were originated by name.

```
to absolute_date | relative_period
```

Restrict the events that are selected to those which occurred before the date and time given by the value. All reports from the oldest stored event up to the

supplied time and date will be reported unless the from command is also supplied to define the earliest event to report.

The `relative_period` is back from the current time and date.

A.12.2 Events report commands

These commands are used to generate event log reports. These commands can be executed as global commands by preceding them by events report, or when in the events context by preceding them by report.

Command	Parameters	Levels
active	boolean	all
all		all
for	relative_period	all
from	absolute_date | relative_period	all
level	level | [level]-[level]	all
long		all
originator	name	all
to	absolute_date | relative_period	all

These commands take exactly the same options as the Events Inactive commands. (See Events inactive commands.) In addition, these commands are permitted:

```
active boolean
```

Only include active events in the report if the boolean value is `enable`; or only include events that are not active if the boolean value is `disable`.

```
long
```

Display the long text associated with each event. The associated short text is always reported.

A.13 FRU commands

These commands manage Field Replaceable Units (FRUs).

Command	Parameters	Levels
info		all

info
Output information about Field Replaceable Units.

A.14 Help | ? [all] | [global] | [help | ?] | [keys]

Without parameters, this command prints all commands that are available in the current context. Global commands are not included unless you are in the global context.

Use the `global` parameter to print all of the current context and global commands.

Use the `keys` parameter to list all of the key sequences which can be used when typing in commands.

Use the `help` or `?` parameter to display help about the help command.

Use the `all` parameter to list all the information described in this section.

A.15 Image commands

Image commands manage the images used to boot the system. These commands are available after the `image` global command, or can be run as global commands by preceding them by `image`.

Command	Parameters	Levels
boot	image number	administrator
info		all
list		all
upload	1kxmodem | url	administrator

boot image_number
Set the numbered image as the image to use the next time the unit is rebooted.

Note: If the image_name of the available images is not unique, the image_number must be used.

info

Display information about the current images. This is the same as the `show images` command.

----[System firmware images]----------	
Image 1 name:	SWR-00031-01{Z.AA}
Build date:	11/11/02
Image 2 name:	SWR-00031-01{2.01}
Build date:	29/10/02
Current image:	[1] SWR-00031-01{Z.AA}
Image to use at the next boot:	[1] SWR-00031-01{Z.AA}

list

List the names and numbers of the available images. The list consists of lines in this format:

[1]	SWR-00031--1{Z.AA}

where the number in the square brackets [ ] is the image number (in this example, 1), and the text is the image's name (in this example, SWR-00031{Z.A}).

upload 1kxmodem

Upload a new image to be used the next time the unit is rebooted. If the upload is successful the new image will be used the next time the unit is booted and the current image is set as the alternative.

Issue the `1kxmodem` command to start the transfer

→ select **Send** from the terminal application menu (Hyperterminal).

upload url

Upload a new image to be used the next time the unit is rebooted. If the upload is successful the new image will be used the next time the unit is booted and the current image is set as the alternative.

The `url` option gives the URL which the unit will access to download the new image. The url can take any of the common formats:

- http:// host:port/pathname
- http:// host/pathname
- ftp:// host/pathname

A.16 Initiator commands (iSCSI)

Note: For information about initiator commands for FC, see Initiator Commands (FC).

These commands are used to manage remote iSCSI initiators. They are available after the `initiator` global command, or as global commands by preceding them by `initiator.`

Note: Remote iSCSI initiators are also implicitly created by other commands

Command	Parameters	Levels
create	Name iqn [options]	administrator
destroy	initiator_list	administrator
info	[initiator_list]	all
list		all
local		all
manage	initiator (See initiator manage initiator commands)	

create name iqn [align boolean] [piggyback boolean] [ping boolean]

Create a new remote iscsi initiator with the local name where `iqn` is the IQN of that initiator.

The `align`, `piggyback` and `ping` parameters set up initial values for the initiator's attributes. By default, the `align` parameter is enabled, the `piggyback` and `ping` parameters are disabled. For more information, (see initiator manage initiator commands).

destroy initiator_list

Destroy the remote iscsi initiators in the `initiator_list` list. This removes all references to these initiators from any targets or other entities which refer to them.

info [name_list]

Output information about all specified remote iSCSI initiators

Including only one initiator in the `name_list` list produces the same results as using the `iscsi initiators manage initiator info` command.

list

List all of the known remote iSCSI initiators.

local

Display IQN of local system's initiator.

A.16.1 initiator manage initiator commands

These commands manage remote iSCSI initiators. These commands are available after the global `initiator manage initiator` command, from the global context by preceding them by `initiator initiator`, or from the iSCSI initiators context by preceding them with `manage initiator`.

Command	Parameters	Levels
noalign	enable | disable	Administrator
info		All
piggyback	enable | disable	Administrator
ping	enable | disable	Administrator

align enable | disable

Change whether Protocol Data Units (PDUs) should be aligned on Ethernet frames

info
Output information about this remote iSCSI initiator, including which other entities (such as targets) refer to it.

piggyback enable | disable
Change the piggyback SCSI status of this remote iSCSI initiator.

ping enable | disable
Change whether to ping periodically to test the connection.

A.17 Initiator commands (FC)

These commands are available after the `initiator` global command and are used to manage remote FC initiators. They can also be run as global commands by preceding them by `initiator`.

Note: Remote FC initiators are also implicitly created by other commands.

Command	Parameters	Levels
create	name WWN	administrator
destroy	initiator_list	administrator
info	[initiator_list]	all
list		all
manage	initiator (See initiator manage initiator commands)	
mode	lunmapped | auto	administrator

create name WWN
Create a new remote FC initiator with the local name where WWN is the port WWN of that initiator.

destroy initiator_list
Destroy the remote FC initiators in the initiator_list list. This removes all references to these initiators from any volumes or other entities which refer to them.

info [name_list]

Output information about all specified remote iSCSI initiators

Including only one initiator in the `name_list` list produces the same results as using the `iscsi initiators manage initiator info` command (see initiator manage initiator commands).

list

List all of the known remote FC initiators.

A.17.1 initiator manage initiator commands

These commands manage remote FC initiators. These commands are available after the global `initiator manage initiator` command, from the global context by preceding them by `initiator initiator`, or from the iSCSI initiators context by preceding them with `manage initiator`.

Command	Parameters	Levels
add	logical LUN	administrator
allow	FC_Port | all	administrator
deny	FC_port | all	administrator
info		all
move	old_LUN new_LUN	administrator
remove	LUN	administrator
rename	new_name	administrator

The maximum number of supported LUN mappings per FC initiator is 256. LUNs are numbered 0-255.

add logical LUN

Map a volume to appear as `LUN LUN` to the specified initiator.

allow FC_port | all

Allow access to the specified initiator on the FC port `FC_port`, or on all ports

deny FC_port | all

Deny access to the specified initiator on the FC port `FC_port`, or on all ports

info
Output information about the specified remote FC initiator.

move old_LUN new_LUN
Change the mapping of a volume to a new LUN.

remove LUN
Remove the mapping that contains `LUN LUN`.

rename old_name new_name
Give a FC initiator a new user label.

A.18 Interface commands

These commands are available after the `interface` global command has been executed and deal with interface options in general. They can also be run as global commands by preceding them by `interface`.

Command	Parameters	Levels
failback	[failed-over_interface_list]	all
info	[interface_list]	all
list		all
manage	interface (See interface manage commands)	
statistics	[interface_list]	all

Interfaces are specified by their hardware names (such as eth0), unless you have renamed them using the `interface rename` command (see rename name).

failback [failed-over_interface_list]
If the IP addresses of any iSCSI interfaces have been transferred (failed over) to another interface because of problems, this command will restore them to their original interface (assuming the problems have been fixed).

info [interface_list]

Output details about the specified interfaces including statistics. If no `interface_list` parameter is supplied, the details about all of the interfaces are output.

See the interface manage commands for more information.

list

Output a list of the known interfaces and their state.

statistics [interface_list]

Output statistics about the specified interfaces. If no `interface_list` parameter is supplied the statistics for all of the interfaces are output.

See statistics for more information about what is displayed by this command.

interface manage commands

Use commands to manage network interfaces. These commands are available after the `interface manage interface` global command, and can also be run as global commands by preceding them by `interface manage interface`, or from the interface context by preceding them with `manage interface`.

Note: If DHCP is enabled for the interface, these commands may store values that will be used when DHCP is disabled. They do not override the values supplied by DHCP.

Command	Parameters	Levels
broadcast 2	broadcast	administrator
dhcp	enable | disable	administrator
down		administrator
failback		administrator
info		all
ip 1	ip_address [ip_mask]	administrator
mtu	mtu	administrator
netmask 1	netmask	administrator
rename	name	administrator

Command	Parameters	Levels
speed	10 | 100 | auto	administrator
statistics		all
Up		administrator

broadcast broadcast

Set the broadcast address. An explicit value may be given in the conventional way, for example 192.112.28.255. Alternatively, the special values listed in the Table above may be entered. These will look at the current address and netmask and deduce the value to use. The value `high` sets all of the bits in the host part and the value `low` clears them.

This command may be combined with the `ip` and `netmask` commands on a single line. The `ip` command describes the default value for this option if just the IP number is changed

If DHCP is enabled this command changes the stored value which will be used if DHCP is disabled. It does not change the actual value of the interface if DHCP is enabled otherwise it resets the interface.

info

Show the details about the interface. This includes:

- IP number
- Netmask
- Broadcast address
- MAC address
- Whether DHCP is enabled
- Whether it is configured (up)
- The interface's speed

If DHCP is enabled the output will show the active values for the interface plus the values which will be applied if DHCP is disabled.

dhcp enable | disable

Either enable or disable DHCP depending on the value of the `boolean` parameter.

ip ip_address [ip_mask]

Specify the IP number for the interface.

Optionally, the netmask may be set at the same time by adding the `ip_mask` value, which defines the number of bits that make up the network. Alternatively, the netmask can be defined by adding the `netmask` command to the same line.

This command may be combined with the `broadcast` and `netmask` commands on a single line. If they are not given, default values will be deduced from the address supplied.

If DHCP is enabled, this command changes the stored value which will be used if DHCP is disabled, but otherwise does not change the actual value of the interface. If DHCP is disabled, this command resets the interface.

mtu mtu

Specify the size of the Maximum Transmission Unit (MTU) for the interface. Press **Tab** to see a list of acceptable size values.

netmask netmask

Set the interface's netmask to be the value supplied.

This command may be combined with the `ip` and `netmask` commands on a single line. The `ip` command describes the default value for this option if just the IP number is changed.

If DHCP is enabled, this command changes the stored value which will be used if DHCP is disabled, but otherwise does not change the actual value of the interface. If DHCP is disabled, this command resets the interface.

rename name

Specify a new name for the interface.

speed 10 | 100 | auto

For devices that can auto-negotiate their speed, clamp their speed to a specified value.

A value of `auto` will re-establish auto-negotiation if the device is capable of doing that. The `full` option is not valid with `auto`.

statistics

Display these statistics for the open interface:

- Number of collisions
- Number of transmitted packets with breakdown of failures
- Number of received packets with breakdown of failures

A.19 iSNS commands (iSCSI only)

These commands manage iSNS configuration.

Command	Parameters	Levels
info		all
server	clear | isns_server	administrator

info

Output information about the iSNS server.

server hostname | IP_address | clear

Set or clear the iSNS server.

A.20 License commands

These commands manage setting and displaying the license key.

Command	Parameters	Levels
info		all
set	[xxxx- xxxx- xxxx- xxxx- xxxx- xxxx]	administrator

info

Output the license key, detailing enabled and disabled features.

set

Set a new license key.

A.21 Logical commands (iSCSI)

These commands control the iSCSI logical drives. They are available after the global `logical` command. They may also be used as global commands by preceding them by `logical`.

If the logical drive is a mirror set or snapshot, you can use these commands after the global `mirror` and `snapshot` commands, or by preceding them by `mirror` or `snapshot`.

Command	Parameters	Levels
chap user	[clear | chap user] | info | password	administrator
create	logical_drive_name available_array available | size restricted | unrestricted	administrator
destroy	logical_list	administrator
down	logical_drive_up	administrator
info	[logical_list]	all
list		all
manage	logical (see logical manage logical commands)	
up	logical_drive_down	administrator

chap user [clear | chap user] | info | password

Manage Challenge Handshake Authentication Protocol (CHAP) details for logical drives. Enter `chap user` to display the current username for CHAP authentication. Enter `chap user username` to set a new username. Enter `chap password` to prompt for a new password.

create logical_drive_name available_array available | size restricted | unrestricted

Create a new logical drive with a user name of `logical`. The IQN name that the initiator uses to access this logical drive is automatically generated. The character set used for a logical drive must be the character set which is valid for an IQN.

The logical drive takes its space from the named `array`. Size is entered in bytes, as a percentage of the total size of the pool. If the `size` is not specified, all of the available space in that pool is used.

The Access Control must be explicitly defined using the `restricted` or `unrestricted` option. If you enter restricted, you must also enter a list of initiators which will be used to define the initial ACL.

destroy logical_list
Destroy the named logical drives.

down logical_drive_up
Take down the named logical drives.

info [logical_list]
Show information about all of the logical drives if a `logical_list` list is not supplied, or just the named logical drives if it is supplied. For single logical drives, this command is the same as the `logical logical info` command.

up logical_list_down
Bring up the named logical drives

list
List the logical drives known to the system.

A.21.1 Logical manage logical commands

These commands manipulate the attributes of the specified logical drive. They are available after the `logical manage logical` command, or can also be run as global commands by preceding them by `logical manage logical` or by `manage logical` from the logical context.

Command	Parameters	Levels
allow	[ro] initiator_list	administrator
authentication	SRP | CHAP| none	administrator
deny	initiator_in_ACL	administrator
down		administrator
growby	size	administrator
growto	size	administrator
info		All
interface	all | interface_list	administrator
rename	name	administrator
restricted		administrator

Command	Parameters	Levels
scope	scope_list	administrator
sessions	parameters	administrator
size	bytes	administrator
unrestricted		administrator

allow [ro] initiator_list
Enter initiators to be added to the list of initiators that can access this logical drive

authentication SRP | CHAP | none
Specify the authentication to use for this logical drive.

deny initiator_in_ACL
Enter initiators to be removed from the list of initiators that can access this logical drive.

down
Take down this logical drive.

growby size
Grow the logical drive by the amount specified.

growto size
Grow the logical drive to the amount specified.

info
Output information about this logical drive.

interface all | interface_list
Redefine the list of interfaces that can be used to access this logical drive. The keyword all selects `all` of the interfaces.

Note: Not all interfaces can be used for logical drives. If you select an interface that is unavailable for use, access is denied although no error is generated.

rename name
Enter a new name (iSCSI alias) for the specified logical drive.

restricted
Activate the logical drive's ACL.

Activating the ACL restricts access to the logical drive to only those initiators defined by the `logical manage logical allow` command.

scope scope_list
Redefine a new list of scopes to which the logical drive should belong. A logical drive is always in the DEFAULT scope.

sessions info | list | reset session_ids
Manage active iSCSI sessions. The `info` parameter shows detail information for iSCSI sessions; `list` displays a list of sessions; `reset` shuts down the specified sessions, or all sessions if none are specified.

unrestricted
Disable the logical drive's ACL. Disabling the ACL allows unrestricted access to this logical drive.

A.22 Logical commands (FC)

These commands control the FC logical drives. They are available after the global `logical` command, or as global commands by preceding them by `logical`.

Command	Parameters	Levels
clear	logical_drive	administrator
create	logical_drive available_array available | size	administrator
destroy	Logical	administrator
info	[logical_list]	all
list		all
manage	logical (see logical manage logical commands)	

clear logical_drives
Clear the ACL on a list of logical drives.

logical_drive available_array available | size

Create a new logical drive with a user name of `name`. The logical drive takes its space from the named `array`. If the `available` option is entered, all of the available space in that array is used. If the `size` option is used, the `size` parameter is the size in bytes.

destroy logical_list

Destroy the named logical drives.

info [logical_list]

Show information about all of the logical drives if the `logical_list` list is not supplied, or just the named logical drives if it is supplied. For a single logical drive, this is the same as the `logical target info` command.

list

List the logical drives known to the system.

A.22.1 Logical manage logical commands

These commands manipulate the attributes of the specified logical drive. They are available after the `logical manage logical` command, or can also be run as global commands by preceding them by `logical manage logical` or by just manage `logical` from the logical context.

Command	Parameters	Levels
growby	size	administrator
growto	size	administrator
info	all	rename
new_name	administrator	

growby size

Grow the logical drive by the amount specified by size.

growto size

Grow the logical drive to the amount specified by size.

info

Output information about this logical drive.

rename logical

Enter a new name for this logical drive.

A.23 Network commands

These commands control aspects of the unit's identity on the network. They are available after the `network` command has been typed. They can also be run as global commands by preceding them with `network`.

Command	Parameters	Levels
dns	ip_list | clear	administrator
domain	name	administrator
gateway	ip_host | clear	administrator
hostname	name	administrator
info		all
ntp	ip_host | clear	administrator
sanname	name	administrator
smtp	host host port port | clear	administrator

dns ip_list | clear

Define a new list of DNS servers where `ip_list` is a list of up to three IP numbers of name servers to use. It replaces the previous list. If a `disable` value is given instead of IP numbers, the list of DNS servers is cleared.

domain name

Set the DNS domain name to be `name`.

This command may be combined with the `hostname` command on the same line to configure the network details in one command. The order of the commands does not matter.

gateway ip_host | clear

Set the gateway network node (or default route) for the unit.

The parameter may be the name or IP number of a router to use as the default gateway or a `disable` value if there is no default gateway.

hostname name

Set the host name of this unit. When added to the `domain` it gives the fully qualified network name.

This command may be combined with the `domain` command on the same line to configure the network details in one command. The order of the commands does not matter.

info

Do a dump of all network information.

ntp ip_host | clear

Set the NTP server to be `ip_host`. For no NTP server a valid disable value should be given. If the `ip_host` name is a valid string for a `disable` value enclose the `ip_host` string in "" characters.

sanname name

Set the name of the Storage Area Network.

This command may be combined with the `domain` command on the same line to configure the network details in one command. The order of the commands does not matter.

host host port port | clear

Set the SMTP server to use. Optionally, the port can be specified. If the port is not specified, the value is not changed; if the port is given as 0, the default port number of 25 is used.

If a `disable` value is supplied instead of a host, the SMTP server name will be cleared.

A.24 Operator command

Use this command to go from the current level to operator level. If a password has been defined for operator level, enter it when prompted.

A.25 Password operator | administrator command

Use this command to define a new password to enter the operator or administrator level, or remove the password requirement.

This command prompts for the password twice. If you enter no password and press **Return** (or **Enter**) twice, the password requirement is removed for that level.

To leave the existing password unchanged, type `CTRL-C` at either one of the password prompts.

The old password is not displayed after this command but does indicate whether the password has been changed or not.

A.26 Port commands (FC only)

These commands control FC ports. They are available after the global `port` command, or as global commands by preceding them with `port`.

Command	Parameters	Levels
identify	Port iterations continual | iterations	administrator
info	[ports]	all
list		all
manage	(See Port manage port commands)	

Note: FC ports are numbered A0, A1, B0, B1, where A and B refer to the controller they are on.

identify port iterations continual | iterations

Flash the LEDs of the specified port. If `continual` is entered, the LEDs remain flashing until you press **Esc**. If `continual` is not entered, the `iterations` value specifies the number of flashes.

info [port_list]

Show port information. Include all ports by not entering a `port_list` list. Include only the named ports by entering a `port_list` list.

For a single port, this command is the same as the `port manage port info` command.

list

List the FC ports known to the system.

A.26.1 Port manage port commands

These commands manipulate the attributes of the specified port. They are available after the `port manage port` command, or as global commands by preceding them with `port manage port` or `manage port` from the port context.

Command	Parameters	Levels
alpa	AL_PA enable | disable	administrator
info		all
speed	1g | 2g | auto	administrator
topology	topology	administrator

alpa AL_PA | enable | disable

Configure the hard AL_PA (Arbitrated Loop Port Address) of the specified port, or enable or disable its use. The change will take effect after a system reboot.

info

Get information about the managed port.

topology topology

Configure the topology of the specified port.

The permissible values are `auto`, `f_port`, `fl_port`, and `loop`.

The change will take effect after a system reboot.

A.27 Quit command

Exit the CLI.

A.28 Route commands

These commands control static routes. They are available after the global `route` command. They may also be used as global commands by preceding them by `route`.

Command	Parameters	Levels
add	parameters	administrator

Command	Parameters	Levels
delete	parameters	administrator
info		All

add ip/mask gateway

Add a static route that goes to the network defined by the combination of `destination` and `netmask` via the given gateway. gateway may be an IP number or a resolvable machine name (as described for the `ip_host` parameter type). The netmask must be specified 192.193.194.0/24 format and the destination must be an IP number (as described for `ip_numb`). If `netmask` is a separate parameter, `destination` may be a resolvable machine name (as described for `ip_host`).

An `interface` that is up and that does not use DHCP to get its IP number may be used instead of `gateway`. If you use such an interface, its IP number is used as the gateway address. Because the route is not associated with the interface, changing the interface's IP number doesn't change the route.

delete ip/mask

Remove the existing static route defined by `destination` and `netmask`. If there is only one existing route that matches `destination` and `netmask`, `gateway` is not required; but if it is supplied it must match the static route. If there are several static routes that match destination and netmask, `gateway` must be supplied to distinguish which static route to remove.

info

Display the current set of static routes plus the default gateway (if there is one). The default gateway is defined by a network command.

A.29 Save command

Use this command to ensure that changes have been written to flash. Changes are normally written at a time chosen by the underlying libraries.

A.30 Scope commands (iSCSI only)

These commands manage scopes. They are available after the global `scope` command. They may also be used as global commands by preceding them by `scope`.

Command	Parameters	Levels
create	name [logical_drive_list]	administrator
destroy	name_list	administrator
info	[name_list]	all
list		all
manage	name (See Scope Manage Scope commands)	all

create name [logical_drive_list]

Create a new scope called `name` which must not be the same as any other scopes. Optionally, add one or more targets to it immediately with the `logical_drive_list` list.

destroy name_list

Destroy the named scopes in the `name_list` list.

info [name_list]

Show information about all of the scopes if `name_list` is not supplied, or the scopes named in the `name_list` list if it is supplied. This is the same as the scope manage name info command if `name_list` consists of a single scope.

list

List all of the known scopes.

A.30.1 Scope Manage Scope commands

These commands control the logical drives in a specific scope. They are available after the global `scope manage scope` command. They may also be used as global commands by preceding them by `scope manage scope` or from the scope context by preceding them by `manage scope`.

Command	Parameters	Levels
add	logical_drive_list	administrator
info		All
remove	logical_drive_list	administrator

add logical_drive _list

Add the list of logical drives to the current scope.

info

Output information about the scope and its logical drives.

remove logical_drive _list

Remove the logical drives in the logical_drive_list list from the scope.

A.31 Service commands

This command produces details about the services running on the system. It is available after the global `service` command or may be used as a global command by preceding it with `service`.

Command	Parameters	Levels
info		All

info

Output details about the services. The information output will include:

- Name
- Status

A.32 Show option command

Use this command to display information about various parts of the system. These are the valid values for this option:

all	Use this value to ask for all the other options in this list
arrays	The same as array_info
authentication (iSCSI only)	The same as authentication list
cache	The same as cache info
devices	The same as device list
events	The same as events info

fru	The same as fru info
images	The same as image info
initiators	The same as initiators list
interfaces	The same as interface info
isns	The same as isns info
license	The same as license info
logicals	The same as logical info
networks	The same as network info
ports (FC only)	The same as port info
routes	The same as route info
scopes (iSCSI only)	The same as scope info
service	The same as service info
snapshot	The same as snapshot info
snmp	The same as snmp info
system	The same as system info

A.33 Snapshot commands (iSCSI)

These commands manage snapshot logical drives. They are available after the global `snapshot` command, or as global commands by preceding them with `snapshot`.

A snapshot is a special logical drive, so many of these commands are identical to the same `logical` command. See Logical Commands (iSCSI) for more information.

Command	Parameters	Levels
create	parameters	administrator
destroy	snapshot_child	administrator
down	snapshot_child	administrator
info	[snapshot_child]	all

Command	Parameters	Levels
list	[parent]	all
manage	snapshot_child (See snapshot manage snapshot_child commands)	
rollback	snapshot_child	administrator
up	snapshot_child	administrator

Note: Only logical drives which have been created by the snapshot create command can be managed by these commands.

create name snapshot_parent [readonly] unrestricted | restricted [initiator_list]

Create a new snapshot of `snapshot_parent` to be called name using array.

The snapshot may be specified as restricted to a limited number of initiators (which you can specify in this command or specify later using the `snapshot manage allow` command) or it can be created as unrestricted. You can also specify access to be read-only.

The `create` command creates the logical drive, so logical must be a name not used by any other logical drive (of any type) before.

destroy snapshot_child

Destroys the snapshot of logical drives given in the `snapshot_child` list and frees the space from each facet's array.

info [snapshot_child]

Show information about all of the snapshot logical drives or the snapshot logical drives in the `snapshot_child` list.

Each snapshot logical drive will output the information shown by the `logical info` command.

list [parent]

List all snapshot logical drives. If the `parent` filter is supplied, only snapshot's parents will be listed.

rollback snapshot_child

Rolls back the snapshots in the `snapshot_child` list.

A.33.1 snapshot manage snapshot_child commands

These commands manipulate the attributes of the specified snapshot logical drive. They are available after the `snapshot manage snapshot_child` command. They can also be run as global commands by preceding them with `snapshot manage snapshot_child` or by just manage `snapshot_child` from the snapshot context.

Many of these commands are actually `logical manage logical` commands that can be applied to snapshot logical drives as well as to non-snapshot logical drives. See logical manage logical commands for detailed information.

Command	Parameters	Levels
allow	[ro] initiator_list	administrator
authentication	srp | chap | none	administrator
deny	initiator_in_ACL	administrator
growby	size	administrator
growto	size	administrator
info		all
interface	all | interface_list	administrator
rename	new_name	administrator
restricted		administrator
scope	scope_list	administrator
sessions	info | list | reset initiator	administrator

A.34 Snapshot commands (FC)

These commands manage snapshot logical drives. They are available after the global `snapshot` command, or as global commands by preceding them with `snapshot`.

A snapshot is a special logical drive, so many of these commands are identical to the same `logical` command. See Logical Commands (FC) for more information.

Command	Parameters	Levels
create	parameters	administrator
destroy	snapshot	administrator
info	[snapshot]	all
list	[parent]	all
manage	logical (See snapshot manage logical commands)	
rollback	snapshot_child	administrator

Note: Only logical drives which have been created by the snapshot create command can be managed by these commands.

create name logical [readonly]

Create a new snapshot of `logical` to be called `name` using `pool`. The `create` command creates the logical drive so `logical` must be a name not used by any other logical drive (of any type) before.

The snapshot will be created using the same pool as logical, so enough space must be remaining in that pool.

destroy snapshot

Destroys the snapshot of logical drives given in the `snapshot` list and frees the space from each facet's array.

info [snapshot]

Show information about all of the snapshot logical drives or the snapshot logical drives in the `snapshot` list.

Each snapshot logical drive will output the information shown by the `logical info` command.

list [parent]

List all snapshots. If the `parent` filter is supplied, only snapshots' parents will be listed.

rollback snapshot_child

Rolls back the snapshots given in the snapshot_child list.

A.34.1 snapshot manage logical commands

These commands manipulate the attributes of the specified snapshot logical drive. They are available after the `snapshot manage logical` command. They can also be run as global commands by preceding them with `snapshot manage logical` or by just `manage logical` from the snapshot context.

Many of these commands are actually `logical manage logical` commands that can be applied to snapshot logical drives as well as to non-snapshot logical drives. See logical manage logical commands for detailed information.

Command	Parameters	Levels
growby	size	administrator
growto	size	administrator
info		all
rename	new_name	administrator

A.35 SNMP commands

The SNMP command should not be used. SNMP should be configured using ServeRAID Manager on the Management Station.

A.36 System commands

These commands set various aspects of the system. They are available after the `system` command has been typed. They can also be run as global commands by preceding them with `system`.

Command	Parameters	Levels
date	absolute_date | relative_period	administrator
defaults		administrator
halt		administrator
info		All
keyboard	layout	administrator
reboot		administrator

Command	Parameters	Levels
status		All
timezone	TZ	administrator

date absolute_date | relative_period

Redefine the time and/or date. The `relative_period` is relative to current time and date and is used to adjust the date and time.

This command can be called with no parameters in Operator mode.

defaults

Reset the system to its factory default state.

halt

Halt the system.

info

Display this information about the operating system:

- BIOS version and date
- Spitfire software version and date
- Hardware serial number
- Active boot image
- Uptime since last boot
- Current date, time and time zone

reboot

Halt the system and boot it again.

status

Display the current status of the unit. This indicates the highest active severity level.

timezone TZ

Change the timezone. A parameter starting with a ' + ' or ' - ' sign is the number of hours in front or behind Greenwich Mean Time (GMT); a parameter starting with a letter is the standard abbreviation for the requested time zone.

This command can be called with no parameters in Operator mode.

B

Network design considerations

The network examples used in this redbook use IP subnetting. In the example network used for iSCSI, LAN 1 is subnet zero and the management network is the all-ones subnet. These two subnets are not differentiable from a full class C network. The subnet mask for the iSCSI LAN 1 subnet is identical to that of a full class C network. The broadcast address for the management LAN is also identical to a full class C network. This can lead to problems in some network installations with multiple routers.

More details are described on Cisco's Web site at:

http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080093f18.shtml

A careful network design is required to prevent problems in your environment.

C

How to get updates of code and documentation

This appendix describes how to retrieve new firmware, drivers and guides from the IBM Support Web page.

IBM Web support site

1. Open a Web browser and enter the URL:

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

2. A redirect to another URL (which is not so easily remembered) occurs. There are two options to navigate to the required information:
   - First option is by entering the IBM machine type, a 4 digit number found on most machines on the front, in the **Search technical support** panel.
   - Second option is to use select **Downloads and drivers**. in the **Support topics** panel. See Figure C-1 on page 625.

Figure C-1 Support downloads page

3. The Downloads and drivers page opens. Select the required product from the **Category** drop down box. The example shown in Figure C-2 on page 626 uses Storage to outline the usage of this page.

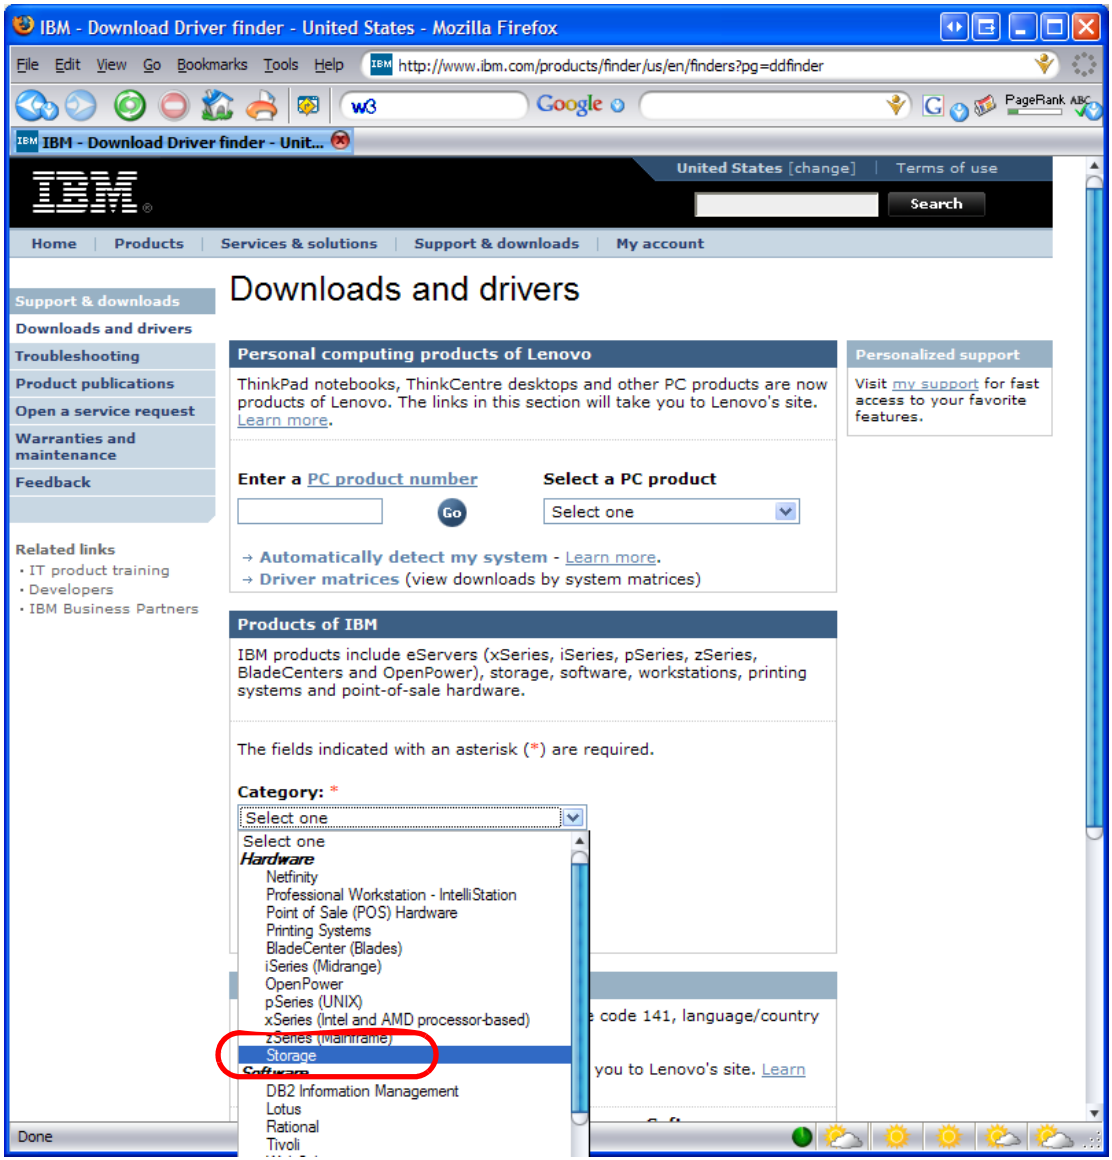

Figure C-2 Select a category

4. The page refreshes and allows the selection of the specific storage product. Select DS300 or DS400 from the **Entry-level disk systems** heading as shown in Figure C-3 on page 627.

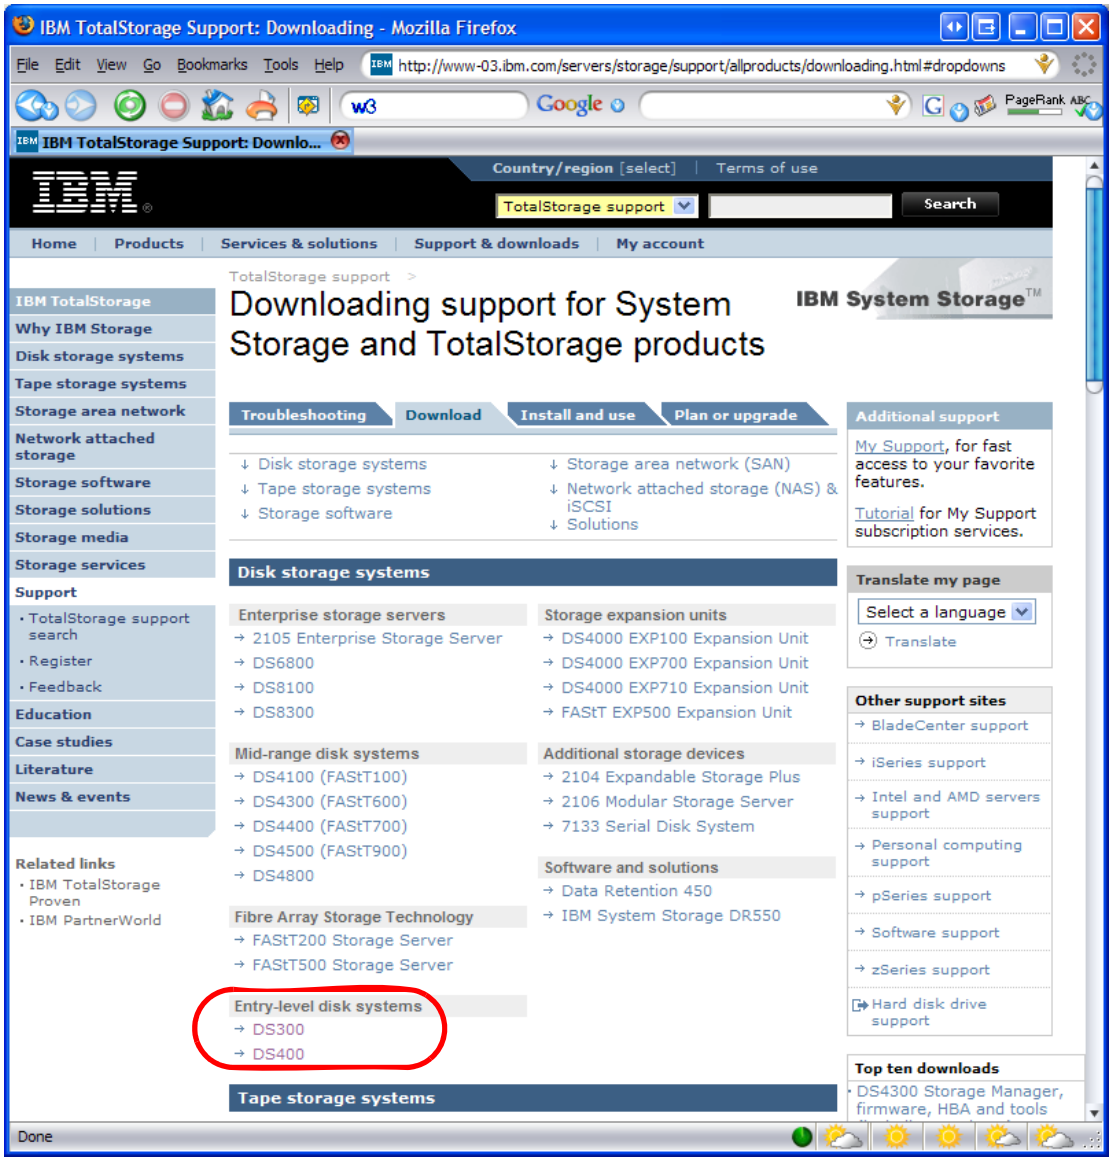

Figure C-3 Select DS300 or DS400 download support

The next page supports driver downloads, publications, and firmware updates for the DS300 and DS400 single and dual controllers.

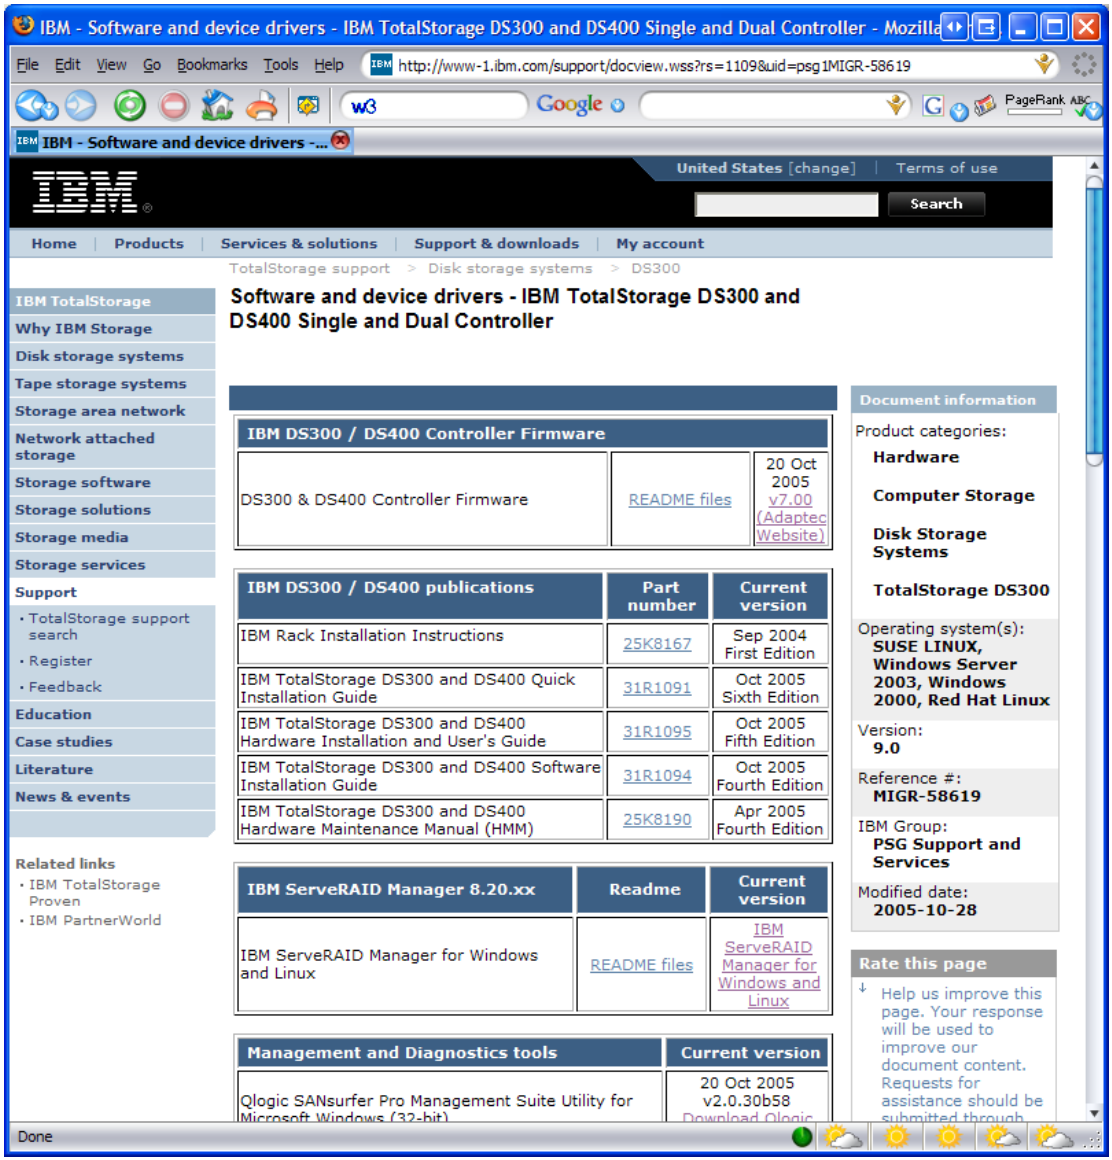

Figure 9-176 Software and device drivers

IBM TotalStorage DS300 and DS400 problem determination guides:

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-56887

TotalStorage DS300 and DS400 software and device drivers:

http://www.ibm.com/pc/support/site.wss/document.do?sitestyle=ibm&lndocid=MIGR-58619

Glossary

A

Agent A software entity that runs on endpoints and provides management capability for other hardware or software. An example is an SNMP agent. An agent has the ability to spawn other processes.

AL See arbitrated loop.

Allocated storage The space that is allocated to volumes, but not assigned.

Allocation The entire process of obtaining a volume and unit of external storage, and setting aside space on that storage for a data set.

Arbitrated loop A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate. See also Fibre Channel Arbitrated Loop and loop topology.

Arbitration The process of selecting one respondent from a collection of several candidates that request service concurrently.

Address Resolution Protocol (ARP) ARP is a layer 3 protocol and is required to allow a sending station to gather address information used in forming a layer 2 frame complete with destination and source MAC addresses. Every router maintains a table listing IP addresses and respective hardware addresses (e.g. MAC addresses) of devices that exist on the network. This table is called an ARP cache and is referenced by the router when it is looking up a hardware address of a device for which it knows the IP address and needs to forward a datagram to it. If no hardware address is found in the ARP cache then an ARP broadcast is sent on to the adjoining media (ARP only applies to the connecting wire). This broadcast is read by every station including the destination station. The destination station sends back an ARP reply with its hardware address so that the IP datagram can now be forwarded to it by the router.

Array An arrangement of related disk drive modules that have been assigned to a group.

B

Backup A copy of computer data that is used to recreate data that has been lost, mislaid, corrupted, or erased. The act of creating a copy of computer data that can be used to recreate data that has been lost, mislaid, corrupted or erased.

Bandwidth A measure of the data transfer rate of a transmission channel.

BCP See Business Continuity planning

Bridge (1) A component used to attach more than one I/O unit to a port. (2) A data communications device that connects two or more networks and forwards packets between them. The bridge may use similar or dissimilar media and signaling systems. It operates at the data link level of the OSI model. Bridges read and filter data packets and frames. Facilitates communication with LANs, SANs, and networks with dissimilar protocols.

Bridge/Router A device that can provide the functions of a bridge, router or both concurrently. A bridge/router can route one or more protocols, such as TCP/IP, and bridge all other traffic.

Business Continuity Planning An enterprise wide planning process which creates detailed procedures to be used in the case of a disaster. Business Continuity Plans take into account processes, people, facilities, systems, and external elements.

C

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

CIFS See Common Internet File system.

CIM See Common Information Model.

CIM Object Manager The CIMOM is the core component to the implementation of the CIM specification. The CIMOM manages the CIM schema, instantiation, communication, and operation of the physical Providers that represent the CIM classes stored within the namespace of the local host.

CIMOM See CIM Object Manager.

Client A function that requests services from a server, and makes them available to the user. A term used in an environment to identify a machine that uses the resources of the network.

Client authentication The verification of a client in secure communications where the identity of a server or browser (client) with whom you want to communicate is discovered. A sender's authenticity is demonstrated by the digital certificate issued to the sender.

Client-server relationship Any process that provides resources to other processes on a network is a server. Any process that employs these resources is a client. A machine can run client and server processes at the same time.

Cluster A type of parallel or distributed system that consists of a collection of interconnected whole computers and is used as a single, unified computing resource.

Common Information Model An object oriented description of the entities and relationships in a business' management environment maintained by the Distributed Management Task Force.

Common Internet File System CIFS provides an open cross-platform mechanism for client systems to request file services from server systems over a network. It is based on the SMB protocol widely used by PCs and workstations running a wide variety of operating systems.

Console A user interface to a server.

core (1) In an optical cable, the central region of an optical fiber through which light is transmitted. (2) In an optical cable, the central region of an optical fiber that has an index of refraction greater than the surrounding cladding material.

CWDM See Wave Division Multiplexing.

D

DASD Acronym for Direct Access Storage Device. This term is common in the z/OS environment to designate a disk or z/OS volume.

Data Sharing A SAN solution in which files on a storage device are shared between multiple hosts.

DATABASE 2 (DB2®) A relational database management system. DB2 Universal Database™ is the relational database management system that is Web-enabled with Java support.

Device driver A program that enables a computer to communicate with a specific device, for example, a disk drive.

Disaster Recovery Planning Functions as a logical subset to the BCP process. IT operation's teams manage the DRP process to ensure continuity of operations in the event of a wide variety of disaster scenarios

Disk group A set of disk drives that have been configured into one or more logical unit numbers. This term is used with RAID devices.

Disk Mirroring A fault-tolerant technique that writes data simultaneously to two hard disks using the same hard disk controller.

Disk Pooling A SAN solution in which disk storage resources are pooled across multiple hosts rather than be dedicated to a specific host.

DRP See Disaster Recovery Planning

duplex Pertaining to communication in which data or control information can be sent and received at the same time, from the same node. Contrast with *half duplex.*

DWDM See Wave Division Multiplexing.

E

E_Port An E_Port is used as an inter-switch expansion port to connect to the E_Port of another Fibre Channel switch, in order to build a larger switched fabric.

Enterprise network A geographically dispersed network under the backing of one organization.

ESCON® Enterprise Systems Connection Architecture®. A zSeries® 900 and S/390® computer peripheral interface. The I/O interface utilizes S/390 logical protocols over a serial interface that configures attached units to a communication fabric.

Event In the Tivoli environment, any significant change in the state of a system resource, network resource, or network application. An event can be generated for a problem, for the resolution of a problem, or for the successful completion of a task. Examples of events are: the normal starting and s ping of a process, the abnormal termination of a process, and the malfunctioning of a server.

F

F_Port An F_Port is a fabric port that is not loop capable. It is used to connect a N_Port to a switch.

Fabric The Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is often used to describe a more complex network utilizing hubs, switches, and gateways.

FC See Fibre Channel.

FCIP See Fibre Channel over IP.

FCP See Fibre Channel protocol.

FCS See Fibre Channel standard.

Fiber optic The medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

Fibre Channel A technology for transmitting data between computer devices at a data rate of up to 2 Gbps. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

Fibre Channel Arbitrated Loop A reference to the FC-AL standard, a shared gigabit media for up to 127 nodes, one of which can be attached to a switch fabric. See also arbitrated loop and loop topology.

Fibre Channel over IP Fibre Channel over IP is defined as a tunneling protocol for connecting geographically distributed Fibre Channel SANs transparently over IP networks.

Fibre Channel Protocol the mapping of SCSI-3 operations to Fibre Channel.

Fibre Channel Standard An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. Refer to ANSI X3.230-199x.

FICON An I/O interface based on the Fibre Channel architecture. In this new interface, the ESCON protocols have been mapped to the FC-4 layer, that is, the Upper Level Protocol layer, of the Fibre Channel Architecture. It is used in the S/390 and z/Series environments.

File system An individual file system on a host. This is the smallest unit that can monitor and extend. Policy values defined at this level override those that might be defined at higher levels.

FL_Port An F_Port is a fabric port that is loop capable. It is used to connect NL_Ports to the switch in a loop configuration.

field replaceable unit (FRU) An assembly that is replaced in its entirety when any one of its required components fails.

Full-Duplex A mode of communications allowing simultaneous transmission and reception of frames.

G

Gateway In the SAN environment, a gateway connects two or more different remote SANs with each other. A gateway can also be a server on which a gateway component runs.

Gigabit Interface Converter (GBIC) a transceiver to convert the internal electrical communication transport to fibre gigabit transport. For 1 Gbps fibre.

Gigabit Link Module (GLM) - a generic Fibre Channel transceiver unit that integrates the key functions necessary for installation of a Fibre channel media interface on most systems.

Gratuitous ARP Most hosts on a network will send out a Gratuitous ARP when they are initializing their IP stack. This Gratuitous ARP is an ARP request for their own IP address and is used to check for a duplicate IP address. If there is a duplicate IP address, then the stack does not complete initialization. See Address Resolution Protocol.

H

half duplex In data communication, pertaining to transmission in only one direction at a time. Contrast with *duplex*.

Hardware zoning Hardware zoning is based on physical ports. The members of a zone are physical ports on the fabric switch. It can be implemented in the following configurations: one to one, one to many, and many to many.

HBA See host bus adapter.

Host Any system that has at least one Internet address associated with it. A host with multiple network interfaces can have multiple Internet addresses associated with it. This is also referred to as a server.

Host bus adapter A Fibre Channel HBA connection that allows a workstation to attach to the SAN network. The function of an HBA is to convert the parallel electrical signals from the host bus into a serial signal to pass to the SAN.

Hub A Fibre Channel device that connects up to 126 nodes into a logical loop. All connected nodes share the bandwidth of this one logical loop. Hubs automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

I

I/O group A group containing two SVC nodes defined by the configuration process. The nodes in the I/O group provide access to the vDisks in the I/O group.

ICAT IBM Common Information Model [CIM] Agent Technology.

iFCP See Internet Fibre Channel Protocol.

Internet Fibre Channel Protocol The Internet Fibre Channel Protocol specification defines iFCP as a gateway-to-gateway protocol for the implementation of a Fibre Channel fabric in which TCP/IP switching and routing elements replace Fibre Channel components.

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

Internet SCSI Internet SCSI encapsulates SCSI commands into TCP packets; therefore enabling the transport of I/O block data over IP networks.

IP Internet protocol.

iSCSI See Internet SCSI.

Isochronous Transmission Data transmission which supports network-wide timing requirements. A typical application for isochronous transmission is a broadcast environment which needs information to be delivered at a predictable time.

J

Java A programming language that enables application developers to create object-oriented programs that are very secure, portable across different machine and operating system platforms, and dynamic enough to allow expendability.

Java runtime environment The underlying, invisible system on your computer that runs applets the browser passes to it.

Java Virtual Machine The execution environment within which Java programs run. The Java virtual machine is described by the Java Machine Specification which is published by Sun™ Microsystems. Because the Tivoli Kernel Services is based on Java, nearly all ORB and component functions execute in a Java virtual machine.

JBOD Just a Bunch Of Disks. A disk group configured without the disk redundancy of the RAID arrangement. When configured as JBOD, each disk in the disk group is a rank in itself.

JRE See Java runtime environment.

JVM™ See Java Virtual Machine.

L

Latency A measurement of the time it takes to send a frame between two locations.

Logical unit number (LUN) a unique identifier used on a SCSI bus that enables it to differentiate between up to eight separate devices (each of which is a logical unit). This number provides a volume identifier that is unique among all storage servers. The LUN is synonymous with a physical disk drive or a SCSI device. For disk subsystems such as the IBM Enterprise Storage Server®, a LUN is a logical disk drive. This is a unit of storage on the SAN which is available for assignment or unassignment to a host server.

Loop topology In a loop topology, the available bandwidth is shared with all the nodes connected to the loop. If a node fails or is not powered on, the loop is out of operation. This can be corrected using a hub. A hub opens the loop when a new node is connected and closes it when a node disconnects. See also Fibre Channel Arbitrated Loop and arbitrated loop.

LUN See logical unit number.

LUN assignment criteria The combination of a set of LUN types, a minimum size, and a maximum size used for selecting a LUN for automatic assignment.

LUN masking This allows or blocks access to the storage devices on the SAN. Intelligent disk subsystems like the IBM Enterprise Storage Server provide this kind of masking.

M

MAN See Metropolitan Area Network.

Managed node A managed node is a computer, a storage system, a gateway, a media device such as a switch or hub, a control instrument, a software product such as an operating system or an accounting package, or a machine on a factory floor, such as a robot.

Managed object A variable of a managed node. This variable contains one piece of information about the node. Each node can have several objects.

Managed resource A physical element to be managed.

Management Agent A process that exchanges a managed node's information with a management station.

Management Information Base A logical database residing in the managed system which defines a set of MIB objects. A MIB is considered a logical database because actual data is not stored in it, but rather provides a view of the data that can be accessed on a managed system.

Management Station A host system that runs the management software.

MDM See Multiple Device Manager.

MDS See Metadata Server Engine.

Metadata Server Engine Within SAN File System, an MDS is the physical hardware on which a Metadata server and an Administrative server run.

Metropolitan Area Network A network that connects nodes distributed over a metropolitan (city-wide) area as opposed to a local area (campus) or wide area (national or global).

MIB See Management Information Base.

MIB object A MIB object is a unit of managed information that specifically describes an aspect of a system. Examples are CPU utilization, software name, hardware type, and so on. A collection of related MIB objects is defined as a MIB.

Mirroring The process of writing data to two separate physical devices simultaneously

Multiple Device Manager The Multiple Device Manager is software that has been designed to allow administrators to manage Storage Area Networks (SANs) and storage from a single console.

N

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

NAS See Network attached storage.

Network attached storage (NAS) Network Attached Storage - a term used to describe a technology where an integrated storage system is attached to a messaging network that uses common communications protocols, such as TCP/IP. An NAS device is attached to a TCP/IP- based network (LAN or WAN), and accessed using CIFS and NFS - specialized I/O protocols for file access and file sharing.

Network topology A physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

NL_Port node loop port A node port that supports arbitrated loop devices.

O

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

P

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

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

Port An endpoint for communication between applications, generally referring to a logical connection. A port provides queues for sending and receiving data. Each port has a port number for identification. When the port number is combined with an Internet address, it is called a socket address.

Port zoning In Fibre Channel environments, port zoning is the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones. See also LUN masking and subsystem masking.

POST In HTTP, a parameter on the METHOD attribute of the FORM tag that specifies that a browser will send form data to a server in an HTTP transaction separate from that of the associated URL.

PPRC See Peer-to-Peer remote copy.

Protocol The set of rules governing the operation of functional units of a communication system if communication is to take place. Protocols can determine low-level details of machine-to-machine interfaces, such as the order in which bits from a byte are sent. They can also determine high-level exchanges between application programs, such as file transfer.

R

RAID Redundant array of inexpensive or independent disks. A method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

RAID 0 Level 0 RAID support - Striping, no redundancy.

RAID 1 Level 1 RAID support - mirroring, complete redundancy.

RAID 5 Level 5 RAID support, Striping with parity.

RDAC See Redundant Disk Array Controller.

Redundant Disk Array Controller Controller failover facility provided for some operating systems with the FAStT product line.

Remote volume mirroring A hardware based remote copy option that provides a synchronous volume copy with the FAStT product line.

Router (1) A device that can decide which of several paths network traffic will follow based on some optimal metric. Routers forward packets from one network to another based on network-layer information. (2) A dedicated computer hardware and/or software package which manages the connection between two or more networks. See also *Bridge, Bridge/Router*.

RVM See Remote volume mirroring.

S

SAN See Storage Area Network.

SAN agent A software program that communicates with the manager and controls the subagents. This component is largely platform independent. See also subagent.

SAN File System SAN file systems allow computers attached via a SAN to share data. They typically separate the actual file data from the metadata, using the LAN path to serve the metadata, and the SAN path for the file data.

SAN Integration Server The SIS is a pre-packaged system comprising an SVC, backend storage, SAN and Ethernet Switches, and a master controller assembled and pre-configured in a rack.

SAN Volume Controller The SVC is a SAN appliance designed for attachment to a variety of host computer systems, which carries out block level virtualization of disk storage.

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

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

SCSI-FCP A standard that defines the protocol used to transfer Small Computer System Interface (SCSI) commands over the transport physical layer of the fibre-channel interface. This standard is published by ANSI as X3.269-1996.

SDD See subsystem device driver.

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

Server A program running on a mainframe, workstation, or file server that provides shared services. This is also referred to as a host.

Shared storage Storage within a storage facility that is configured such that multiple homogeneous or divergent hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. You need to design the programs to handle the effects of concurrent access.

Simple Network Management Protocol A protocol designed to give a user the capability to remotely manage a computer network by polling and setting terminal values and monitoring network events.

SIS See SAN Integration Server.

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

Small Form Factor Pluggable media (SF) a transceiver to convert the internal electrical communication transport to fibre gigabit transport. For 2 Gbps fibre.

SNMP See Simple Network Management Protocol.

SNMP agent An implementation of a network management application which is resident on a managed system. Each node that is to be monitored or managed by an SNMP manager in a TCP/IP network, must have an SNMP agent resident. The agent receives requests to either retrieve or modify management information by referencing MIB objects. MIB objects are referenced by the agent whenever a valid request from an SNMP manager is received.

SNMP manager A managing system that executes a managing application or suite of applications. These applications depend on MIB objects for information that resides on the managed system.

SNMP trap A message that is originated by an agent application to alert a managing application of the occurrence of an event.

Software zoning Is implemented within the Simple Name Server (SNS) running inside the fabric switch. When using software zoning, the members of the zone can be defined with: node WWN, port WWN, or physical port number. Usually the zoning software also allows you to create symbolic names for the zone members and for the zones themselves.

SQL Structured Query Language.

SSA See Serial Storage Architecture.

Storage administrator A person in the data processing center who is responsible for defining, implementing, and maintaining storage management policies.

Storage Area Network A managed, high-speed network that enables any-to-any interconnection of heterogeneous servers and storage systems.

Striping A method for achieving higher bandwidth using multiple N_Ports in parallel to transmit a single information unit across multiple levels.

Subagent A software component of SAN products which provides the actual remote query and control function, such as gathering host information and communicating with other components. This component is platform dependent. See also SAN agent.

Subsystem Device Driver The Subsystem Device Driver is a multipathing software designed especially to use with the IBM Enterprise Storage Server 2105. It cannot be used with any other storage servers or storage devices.

Subsystem masking The support provided by intelligent disk storage subsystems like the Enterprise Storage Server. See also LUN masking and port zoning.

SVC See SAN Volume Controller.

Switch A component with multiple entry and exit points or ports that provide dynamic connection between any two of these points.

Switch topology A switch allows multiple concurrent connections between nodes. There can be two types of switches, circuit switches and frame switches. Circuit switches establish a dedicated connection between two nodes. Frame switches route frames between nodes and establish the connection only when needed. A switch can handle all protocols.

T

TCP See Transmission Control Protocol.

TCP/IP Transmission Control Protocol/Internet Protocol - a set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

Topology An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, arbitrated loop, and switched fabric are all Fibre Channel topologies.

Transmission Control Protocol A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for Internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

U

UNIX A highly portable operating system originally developed by Bell Laboratories that features multiprogramming in a multiuser environment. UNIX is implemented in the C language. UNIX was originally developed for use on minicomputers but has been adapted on mainframes and microcomputers. It is especially suitable for multiprocessor, graphics, and vector-processing systems.

V

vDisk See Virtual Disk.

Virtual Disk An SVC device that appears to host systems attached to the SAN as a SCSI disk. Each vDisk is associated with exactly one I/O group.

Virtualization An abstraction of storage where the representation of a storage unit to the operating system and applications on a server is divorced from the actual physical storage where the information is contained.

W

WAN Wide Area Network.

Wave Division Multiplexing WDM allows the simultaneous transmission of a number of data streams over the same physical fiber, each using a different optical wavelength. WDM receives incoming optical signals from many sources (Fibre Channel, IP, ESCON, FICON) which it converts to electrical signals, it then assigns them a specific wavelength (or lambdas) of light and retransmits them on that wavelength. This method relies on the large number of wavelengths available within the light spectrum. Coarse WDM (CWDM) and Dense WDM (DWDM) are based on the same methodology as WDM enabling more data streams over the same physical fiber.

WDM See Wave Division Multiplexing.

World Wide Name A unique number assigned to Fibre Channel devices (including hosts and adapter ports) - analogous to a MAC address on a network card.

WWN See World Wide Name.

Z

Zoning In Fibre Channel environments, zoning allows for finer segmentation of the switched fabric. Zoning can be used to instigate a barrier between different environments. Ports that are members of a zone can communicate with each other but are isolated from ports in other zones. Zoning can be implemented in two ways: hardware zoning and software zoning.

Abbreviations and acronyms

ACL	Access Control List
AD	Microsoft Active Directory
AFS®	Andrew File System
AIX	Advanced Interactive eXecutive
ALUA	asynchronous logical unit access
ANSI	American National Standards Institute
API	Application Programming Interface
APPC	Advanced Program to Program Communication
ARP	Address Resolution Protocol
ASCII	American National Standard Code for Information Interchange
ASIC	application-specific integrated circuit
ATM	Asynchronous Transfer Mode
BCP	Business Continuity Planning
BDC	Backup Domain Controller
BIND	Berkeley Internet Name Domain
BIOS	basic input/output system
BOS	Base Operating System
BMC	baseboard management controller
BSD	Berkeley Software Distribution
CA	Certification Authorities
CDE	Common Desktop Environment
CDMF	Commercial Data Masking Facility
CDS	Cell Directory Service
CERT	Computer Emergency Response Team
CGI	Common Gateway Interface
CHAP	Challenge Handshake Authentication Protocol
CLI	command line interface
CIFS	Common Internet File System
CIM	Common Information Model
CIMOM	CIM Object Manager
COW	Copy-On-Write
CPM	Copper Passthrough Module
C-SPOC	Cluster Single Point Of Control
CWDM	Coarse Wave Division Multiplexing
DA	Directory agent
DASD	Direct Access Storage Device
DBM	DataBase Management
DCE	Distributed Computing Environment
DCOM	Distributed Component Object Model
DDE	Dynamic Data Exchange
DDNS	Dynamic Domain Name System
DEN	Directory Enabled Network
DES	Data Encryption Standard
DFS™	Distributed File System
DHCP	Dynamic Host Configuration Protocol
DLC	Data Link Control
DLL	Dynamic Load Library
DMA	direct memory access
DNS	Domain Name System
DRP	Disaster Recovery Planning
DSA	Dynamic System Analysis
DSM	Device Specific Module
DWDM	Dense Wave Division Multiplexing
EFS	Encrypting File Systems
EMS	Event Management Services
ERD	Emergency Repair Disk
ERP	Enterprise Resources Planning
ESC	escape character
ESCON	Enterprise System Connection
ESM	Enclosure Services Module
ESS	Enterprise Storage Server

ETL Enterprise Tape Library
FAT File Allocation Table
FC Fibre Channel
FC-AL Fibre Channel Arbitrated Loop
FCIP Fibre Channel over IP
FCP Fibre Channel Protocol
FCS Fibre Channel Standard
FDDI Fiber Distributed Data Interface
FEC Fast EtherChannel technology
FICON fibre-channel connection
FIFO First In/First Out
FLA fabric-loop attach
FQDN Fully Qualified Domain Name
FRU field-replaceable unit
FSF File Storage Facility
FTD Fault-Tolerant Disk
FTP File Transfer Protocol
GBIC Gigabit Interface Converter
GDA Global Directory Agent
GDS Global Directory Service
GID Group Identifier
GL Graphics Library
GLM gigabit link module
GUI Graphical User Interface
HA High Availability
HACMP™ High Availability Cluster Multiprocessing
HAL Hardware Abstraction Layer
HBA Host Bus Adapter
HCL Hardware Compatibility List
HSM Hierarchical Storage Management
HTML Hypertext Markup Language
HTTP Hypertext Transfer Protocol
IBM International Business Machines Corporation
ICAT IBM Common Information Model (CIM) Agent Technology
IDE Integrated Drive Electronics
IEEE Institute of Electrical and Electronic Engineers
IETF Internet Engineering Task Force
IFCP Internet Fibre Channel Protocol
IGMP Internet Group Management Protocol
IGS IBM Global Services
IIS Internet Information Server
IKE Internet Key Exchange
IMAP Internet Message Access Protocol
IP Internet Protocol
IQN iSCSI qualified name
IPC Interprocess Communication
IPL Initial Program Load
IPsec Internet Protocol Security
IPX™ Internetwork Packet eXchange™
ISA Industry Standard Architecture
iSCSI Internet SCSI
ISDN Integrated Services Digital Network
ISL interswitch link
iSNS Internet Storage Name Service
ISO International Standards Organization
ISS Interactive Session Support
ISV Independent Software Vendor
ITSO International Technical Support Organization
JBOD Just a Bunch Of Disks
JFS Journaled File System
JRE Java Runtime Environment
JVM Java Virtual Machine
LAN Local Area Network
LCN Logical Cluster Number
LDAP Lightweight Directory Access Protocol
LED light-emitting diode

LFS	Log File Service (Windows NT®)
LFT	Low Function Terminal
LOS	Layered Operating System
LPAR	Logical Partition
LPC	Local Procedure Call
LPP	Licensed Program Product
LSA	Local Security Authority
LUID	Login User Identifier
LUN	Logical Unit Number
LVCB	Logical Volume Control Block
LVDD	Logical Volume Device Driver
LVM	Logical Volume Manager
MAC	Medium Access Control
MAN	Metropolitan Area Network
MBR	Master Boot Record
MBps	Megabytes per second
MCA	Micro Channel® Architecture
MCS	Multiple Connections per Session
MDC	Meta Data Controller
MDM	Multiple Device Manager
MDS	Metadata Server Engine
MFT	Master File Table
MIB	Management Information Base
MIPS	Million Instructions Per Second
MMC	Microsoft Management Console
MOCL	Managed Object Class Library
MPIO	multipath I/O
MPTN	Multi-protocol Transport Network
MSCS	Microsoft Cluster Server
MSJ	Management Suite Java
MSS	Modular Storage Server
MTU	Maximum Transmission Unit
MWC	Mirror Write Consistency
NAS	Network Attached Storage
NBC	Network Buffer Cache
NBPI	Number of Bytes per I-node
NCP	NetWare Core Protocol
NCS	Network Computing System
NCSC	National Computer Security Center
NDIS	Network Device Interface Specification
NDMP	Network Data Management Protocol
NDS	NetWare Directory Service
NFS	Network File System
NetBEUI	NetBIOS Extended User Interface
NetDDE	Network Dynamic Data Exchange
NETID	Network Identifier
NFS	Network File System
NIC	Network Information Center
NIM	Network Installation Manager
NIS	Network Information System
NNS	Novell Network Services
NTFS	NT File System
NTLDR	NT Loader
NTLM	NT LAN Manager
NTP	Network Time Protocol
NVRAM	Non-Volatile Random Access Memory
ODBC	Open Database Connectivity
ODM	Object Data Manager
OLTP	OnLine Transaction Processing
OMG	Object Management Group
ONC™	Open Network Computing
OPM	Optical Passthrough Module
OS	Operating System
OSF	Open Software Foundation
OSI	open systems interconnection
PAP	Password Authentication Protocol
PB	Petabytes
PCI	Peripheral Component Interconnect
PCI-X	Peripheral Component Interconnect-X
PCMCIA	Personal Computer Memory Card International Association
PDC	Primary Domain Controller

PDF Portable Document Format
PDS Partitioned data set
PDSE Partitioned data set extended
PDT Performance Diagnostic Tool
PDU Protocol Data Unit
PEX PHIGS Extension to X
PHIGS Programmer's Hierarchical Interactive Graphics System
PID Process Identification Number
PIN Personal Identification Number
PLDA private-loop direct attach
PMTU Path Maximum Transfer Unit
POSIX Portable Operating System Interface for Computer Environment
POST Power-On Self Test
PP Physical Partition
PPP Point-to-Point Protocol
PPRC Peer to Peer Remote Copy
PPTP Point-to-Point Tunneling Protocol
PSM Persistent Storage Manager
PSSP Parallel System Support Program
PV Physical Volume
PVID Physical Volume Identifier
QoS Quality of Service
RACF® Resource Access Control Facility
RAID Redundant Array of Independent Disks
RAM random access memory
RAS Remote Access Service
RDAC Redundant Disk Array Controller
RDBMS Relational Database Management System
RFC Request for Comments
RGID Real Group Identifier
RISC Reduced Instruction Set Computer
RMC Resource Monitoring and Control
ROM read-only memory
ROS Read-Only Storage
RPC Remote Procedure Call
RRIP Rock Ridge Internet Protocol
RSA Rivest-Shamir-Adleman algorithm
RSCN Registered State Change Notification
RSCT Reliable Scalable Cluster Technology
RSM™ Removable Storage Management
RSVP Resource Reservation Protocol
RVM Remote Volume Mirroring
SA Service agent
SAK Secure Attention Key
SAM Security Account Manager
SAN Storage Area Network
SASL Simple Authentication and Security Layer
SCSI Small Computer System Interface
SDD Subsystem Device Driver
SDK Software Developer's Kit
SFP Small form-factor pluggable
SID Security Identifier
SIS SAN Integration Server
SLIP Serial Line Internet Protocol
SMB Server Message Block system-managed buffering
SMIT System Management Interface Tool
SMP Symmetric Multiprocessor
SMS Systems Management Server
SNA Systems Network Architecture
SNIA Storage Network Industry Association
SNMP Simple Network Management Protocol
SP System Parallel
SPX Sequenced Packet eXchange
SQL Structured Query Language
SRM Storage Resource Manager
SSA Serial Storage Architecture
SSL Secure Sockets Layer

SVC	SAN Volume Controller
SWAN	Storage Wide Area Networks
TCB	Trusted Computing Base
TCO	Total cost of ownership
TCP/IP	Transmission Control Protocol/Internet Protocol
TDI	Transport Data Interface
TDP	Tivoli Data Protection
TLS	Transport Layer Security
TOE	TCP/IP Offload Engine
TOS	Type of Service
TSM	Tivoli Storage Manager
UA	User agent
UCS	Universal Code Set
UDB	Universal Database
UDF	Universal Disk Format
UDP	User Datagram Protocol
UFS	UNIX File System
UID	User Identifier
UMS	Ultimedia Services
UNC	Universal Naming Convention
UPS	Uninterruptible Power Supply
URL	Universal Resource Locator
USB	Universal Serial Bus
UTC	Universal Time Coordinated
UUCP	UNIX to UNIX Communication Protocol
UUID	Universally Unique Identifier
VAX	Virtual Address eXtension
VCN	Virtual Cluster Name
VFS	Virtual File System
VG	Volume Group
VGDA	Volume Group Descriptor Area
VGID	Volume Group Identifier
VGSA	Volume Group Status Area
VIPA	Virtual IP Address
VMM	Virtual Memory Manager
VPD	Vital Product Data
VPN	Virtual Private Network
VSM™	Virtual System Management
VTS	Virtual Tape Server
W3C	World Wide Web Consortium
WAN	Wide Area Network
WDM	Wavelength Division Multiplexing
WINS	Windows Internet Name Service
WLM	Workload Manager
WWN	World Wide Name
WWPN	world wide port number
WWW	World Wide Web
XCMF	X/Open Common Management Framework
XDM	X Display Manager
XNS	XEROX Network Systems

Related publications

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

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 648. Note that some of the documents referenced here may be available in softcopy only.

- *The IBM TotalStorage Solutions Handbook,* SG24-5250
- *Using iSCSI Solutions’ Planning and Implementation*, SG24-6291
- *DS4000 Best Practices and Performance Tuning Guide*, SG24-6363
- *Introduction to Storage Area Networks*, SG24-5470
- *Designing an IBM Storage Area Network*, SG24-5758
- *IBM SAN Survival Guide*, SG24-6143
- *IBM SAN Survival Guide Featuring the Cisco Portfolio*, SG24-9000
- *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149
- *IBM SAN Survival Guide Featuring the CNT Portfolio*, SG24-6150

Other publications

These publications are also relevant as further information sources:

- *iSCSI The Universal Storage Connection,* by John L. Hufferd. Addison-Wesley Professional, First Edition, November 2002. ISBN 020178419X

Online resources

These Web sites and URLs are also relevant as further information sources:

- IBM TotalStorage Web Site

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

- IBM Redbooks

  http://www.redbooks.ibm.com/
- IBM TotalStorage Productivity Center

  http://www-1.ibm.com/servers/storage/software/center/index.html
- IBM Storage networking products

  http://www.storage.ibm.com/snetwork/index.html

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

C

D

E

F

G

L

M

N

O

P

Q

R

S

T

U

V

W

Z

IBM TotalStorage DS300 and DS400 Best Practices Guide

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM TotalStorage DS300 and DS400 Best Practices Guide

SAN management options

Recommended configurations

Software and hardware installation details

This IBM Redbook represents a compilation of best practices for configuring the IBM TotalStorage DS300 (iSCSI) and DS400 (Fibre). These workgroup storage subsystems for IBM @server xSeries and BladeCenter servers are designed to provide an exceptional solution for workgroup storage applications such as file, print and Web serving, as well as, remote boot storage for diskless servers.

The redbook provides a conceptual framework for understanding the DS300 and DS400 in a Storage Area Network and includes recommendations, hints, and tips for the physical installation, cabling, and zoning. It shows:

- Supported configurations
- Proper use of the software and hardware initiators
- Array and Logical drive configurations to support maximum throughput
- DSM and MPIO usage
- Understanding failover and failback
- Best practices with installation (i.e., FC switches, BladeCenters, Ethernet switches)
- Single to dual upgrade

This book is intended for IBM technical professionals, Business Partners, and customers responsible for the planning, deployment, and maintenance of IBM TotalStorage DS300/DS400 products.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7121-00 ISBN 0738496316