IBM BladeCenter iSCSI SAN Solution

- Include failover support utilizing the IBM N3700 and NetApp FAS270
- QLogic iSCSI Expansion Card using the Microsoft iSCSI software initiator
- High Availability using Nortel L2/3 and L2-7 Switch Modules

Rufus Credle
Khalid Ansari
Scott Lorditch
P Chris Mojica

Redpaper

ibm.com/redbooks

International Technical Support Organization

IBM BladeCenter iSCSI SAN Solution

August 2006

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

First Edition (August 2006)

This edition applies to IBM BladeCenter, NetAPP FAS270 and N-Series N3700 storage system.

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®	xSeries®	System x™
@server®	AIX®	System Storage™
Redbooks (logo) ™	BladeCenter®	TotalStorage®
eServer™	IBM®	
iSeries™	Redbooks™	

The following terms are trademarks of other companies:

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

NetApp, the Network Appliance logo, the bolt design, are registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

QLogic, the QLogic logo SANbox, SAN Pro and SAN Express are registered trademarks of QLogic Corporation.

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

Preface

This IBM® Redpaper provides the instructions to configure an IBM BladeCenter® Boot from iSCSI NAS from the IBM N-Series N3700 storage system. Our Boot from the iSCSI NAS test environment included the use of the Nortel and Cisco Switch Modules for the IBM BladeCenter. We also discuss High Availability utilizing the Nortel Networks L2-3 and L2-7 Switch Modules for the IBM BladeCenter.

The team that wrote this Redpaper

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

Rufus Credle is a Certified Consulting IT Specialist at the ITSO, Raleigh Center. In his role as Project Leader, he conducts residencies and develops IBM Redbooks™ about network operating systems, ERP solutions, voice technology, high availability and clustering solutions, Web application servers, pervasive computing, and IBM and OEM e-business applications, all running IBM System x™ and IBM BladeCenter. Rufus's various positions during his IBM career have included assignments in administration and asset management, systems engineering, sales and marketing, and IT services. He holds a BS degree in business management from Saint Augustine's College. Rufus has been employed at IBM for 26 years.

Khalid Ansari is the Technical Team Lead for the IBM Blade Infrastructure Solution Center team in Research Triangle Park, North Carolina. His responsibilities include assisting BladeCenter pre-sale customers world-wide with new proof of concepts and pilot testing. He has developed numerous technical documents and best practices procedures for BladeCenter complex technology solutions. He was a participant in developing the IBM Storage Networking Solutions V1 Certification. Khalid started with IBM in August 1998 as an ATM Networking Specialist in Level 2 support and later worked in SAN Solutions support.

Scott Lorditch is a Sales Network Architect for Blade Network Technologies, Inc. (formerly the Blade Switching Server business unit of Nortel Networks). He develops designs and proposals for customers and potential customers of Blade Network Technologies, Inc. GbESM products for the IBM BladeCenter, including overall network architecture assessments. He has also developed several training and lab sessions for IBM technical and sales personnel and has provided field feedback to the product team. His background before working for Blade Network Technologies, Inc., includes almost 20 years working on networking, including electronic securities transfer projects for a major bank based in New York City, as a Senior Network Architect for a multi-national soft drink company, and as Product Manager for managed hosting services for a large telecommunications provider. He holds a BS in Operations Research with specialization in Computer Science from Cornell University.

P Chris Mojica is a Staff Systems Engineer for QLogic supporting various IBM xSeries® and iSeries™ groups with pre-sales and post-sales responsibilities. He joined QLogic in 2001 and has supported multi-generational OEM product life spans and published numerous whitepapers and design guides focusing on subjects from chip level design to boot from SAN. Prior to coming to QLogic Chris had successful careers in the networking and broadband fields as a Network Systems Consultant with Lucent Network Care division, formerly International Network Services, and as a Field Applications Engineer with Xpeed Corporation. Chris began his professional career co-oping with IBM Network Hardware Division.

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

Tamikia Barrows, Jeanne Tucker, Margaret Ticknor
International Technical Support Organization, Raleigh Center

Kenny Speer, Senior Engineer, SAN Interoperability
Network Appliance RTP

Kalman Z. Meth, Ph.D. and Julian Satran
IBM Haifa Labs

Madhav Ponamgi, Computational Scientist, Advanced Technical Services Life Sciences/Healthcare
IBM Furlong PA

Fabiano Matassa, ITA
IBM United Kingdom

David Tareen, Product Marketing Manager - Storage
IBM Research Triangle Park

Shawn Andrews, e-Server Storage Development
IBM Research Triangle Park

Robert Ray, IBM System x Systems Engineer
IBM Cincinnati

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

1

iSCSI implementation with the IBM BladeCenter

This chapter discusses those items to be considered when implementing an IBM BladeCenter Boot from the iSCSI SAN solution and the products we used for our test environment.

1.1 Introduction

There are two main ways to implement iSCSI in general, and specifically on Blade servers within a BladeCenter chassis. The two options are to use hardware initiators or software initiators. It is generally not possible to use both on the same server at the same time, but iSCSI will not work without one of these two options.

1.2 Why iSCSI versus Fibre Channel

The iSCSI protocol is a transport for SCSI over TCP/IP. Until recently, standard IP protocol infrastructure (for example, Ethernet) could not provide the necessary high bandwidth and low-latency needed for storage access. Special communications infrastructure, mainly Fibre Channel running FCP (SCSI over Fibre Channel), was developed to allow for Storage Area Networks (SANs). With the recent advances in Ethernet technology, it is now practical (from a performance perspective) to access storage devices over an IP network. 1 Gigabit Ethernet is now widely available and is competitive with 1 and 2 Gigabit Fibre Channel. 10 Gigabit Ethernet is readily coming available. Similar to FCP, iSCSI allows storage to be accessed over a Storage Area Network, allowing shared access to storage. A major advantage of iSCSI over FCP is that iSCSI can run over standard, off-the-shelf Ethernet network components. A network that incorporates iSCSI SANs need use only a single kind of network infrastructure (Ethernet) for both data and storage traffic, while use of FCP requires a separate type of infrastructure (Fibre Channel) and administration for the storage. Furthermore, iSCSI (TCP) based SANs can extend over arbitrary distances, and are not subject to distance limitations that currently limit FCP.

Since iSCSI is designed to run on an IP network, it can take advantage of existing features and tools that were already developed for IP networks. The very use of TCP utilizes TCP's features of guaranteed delivery of data and congestion control. IPSec can be leveraged to provide security of an iSCSI SAN, while a new security mechanism may have to be developed for the Fibre Channel. Service Location Protocol (SLP) can be used by iSCSI to discover iSCSI entities in the network. Thus, in addition to iSCSI running on standard, cheaper, off-the-shelf hardware, iSCSI also benefits from using existing, standard IP-based tools and services.

Note: Fibre Channel security is available. It is defined by the T11 organization (that defines all the Fibre Channel specs). It is the FC-SP spec and can be found at the following Web site:

http://www.t11.org/ftp/t11/pub/fc/sp/06-157v0.pdf

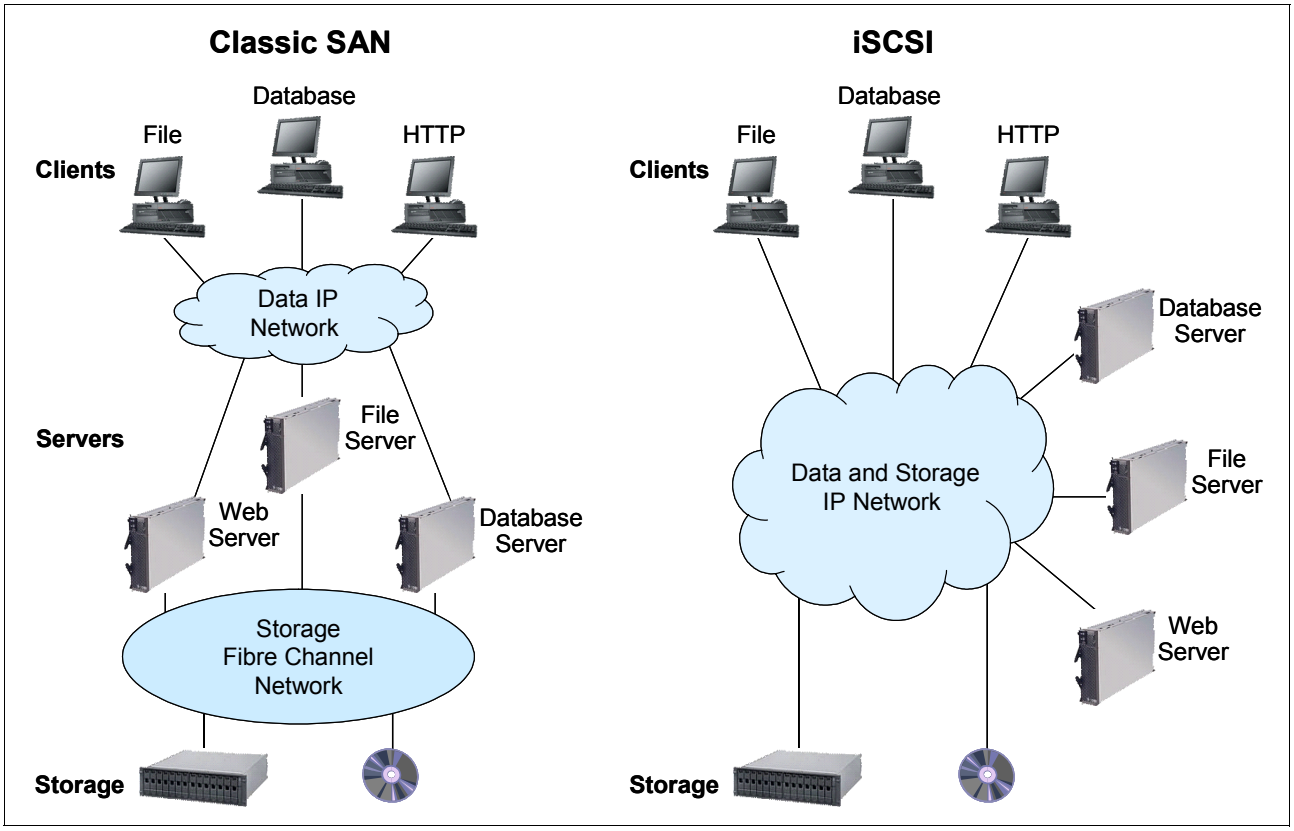

Figure 1-1 Classic SAN and an iSCSI SAN topology

Note: Where the Boot from iSCSI or Fibre Channel SAN would serve and satisfy the requirements of several client environments, there are some engineering and scientific applications that require the use of the local disk for better performance and may not be feasible for a Boot from iSCSI or Fibre Channel SAN solution. For example:

- Nastran (linear finite element analysis codes for engineering)
- Dytran (non-linear finite element analysis)
- Marc (another non-linear code)
- Fluent (fluid dynamics)
- Gaussian (computation chemistry)
- Amber (computational chemistry)
- GCG (computational chemistry)

Therefore, you should know your applications and where best they should be run.

1.3 TOE benefits

We know that the processing of TCP packets from an Ethernet connection consumes many CPU resources, and iSCSI protocol only adds another layer of processing. With the number of packets and their corresponding interrupts required for iSCSI, the software iSCSI packet processing can burden the host system with 50–65% CPU usage. Depending upon the signaling options used, high CPU usage may even render some host applications unusable.

Therefore, it is important that efficient iSCSI systems depend on a hardware TCP Offload Engine (TOE) to handle the transportation protocols of iSCSI. A TOE network interface card

(NIC) is a special interface card specifically designed for interfacing a server to the IP-SAN and will offload iSCSI as well and TCP/IP encapsulation from the server CPUs. A hardware TOE implements the entire standard TCP and iSCSI protocol stacks into the hardware. This approach completely offloads the iSCSI protocol from the primary CPU, processing storage communications efficiently and enabling applications to run faster and more reliably. By using the TCP Offload Engine, a single system can run multiple initiators for improved throughput.

In our Boot from iSCSI example, we use the QLogic iSCSI Expansion Card for IBM BladeCenter. This iSCSI Expansion Card option is a hardware initiator that provides iSCSI communication from the blade server to an iSCSI storage device. It delivers full TCP/IP Offload Engine (TOE) functionality to reduce CPU processing. For more information view the Web link:

http://www-306.ibm.com/common/ssi/rep_ca/4/897/ENUS105-194/ENUS105-194.PDF

1.4 Software initiators

A configuration that uses software initiators includes the following:

- Microsoft® iSCSI software initiator or equivalent. The Microsoft software is used in the examples that follow.
- One or two Ethernet switch modules, preferably the use of two Ethernet switch modules for high availability.
- There is no daughter card required on the server blades to implement a connection to an iSCSI SAN environment. However, it is possible to use an Ethernet daughter card and deploy Ethernet switch modules in I/O module bays 3 and 4 in addition to the Ethernet switch modules for data networking in bays 1 and 2. For our example in Chapter 2, “Boot from iSCSI SAN using iSCSI HBA and initiator with failover support” on page 9, it was necessary that we use a TOE network interface card (QLogic HBA) and Microsoft software initiator to successfully implement *a boot from iSCSI and failover solution*.

Note: If you are using BladeCenter bays 3 and 4, you will need an iSCSI daughter card or an Ethernet daughter card installed on the blade. If not, your blade will not have connectivity to the storage subsystems via the switch modules in bays 3 and 4.

Announcement: In April 2006, Microsoft announced that it will supply select partners with specialized code that allows them to toggle the boot order, making a boot from iSCSI SAN possible and supported. Two of Microsoft's partners are IBM and emBoot, with its WinBoot/i iSCSI boot-from-SAN software initiator.

For more information about Microsoft support for iSCSI visit the following Web site:

http://www.microsoft.com/windowsserver2003/technologies/storage/iscsi/msfiSCSI.mspx

1.5 Hardware initiators

A configuration that uses hardware initiators includes the following:

- An iSCSI daughter card for each blade, which will access the storage array, and the associated drivers and other software. The examples that follow later in this Redpaper used the QMC4052 card from QLogic.

- One or two Ethernet switch modules, preferably the use of two Ethernet switch modules for high availability. Note that these switch modules can only be installed in I/O module bays 3 and 4. As above, the examples that follow in this Redpaper were tested using Nortel Networks L2/3 ESM and the Cisco Systems Intelligent GbESM for IBM BladeCenter.
- The BladeCenter Management Module will detect if there is a conflict between the daughter cards installed on the server blades and the switch modules installed in I/O bays 3 and 4. It will disable power to the switch modules or connectivity to the blades to prevent damage to any of the parts inserted in the chassis.

1.6 Choosing between hardware and software initiators

The key benefits of using hardware initiators are:

- Their performance will be noticeably faster.
- They will not interfere with data networking traffic if the network topology is designed to segregate storage traffic such as by using a different set of switches to connect servers to iSCSI storage.
- The traffic that passes through them will not load the server's CPU to the same extent that would be the case if the storage traffic passed through the standard IP stack.
- It is possible to implement *iSCSI boot from SAN* with hardware initiators.

> **Note:** At the time of the writing of this Redpaper, Microsoft announced that it will make its software initiator capable of booting from iSCSI SAN possible and supported.

The key benefits of using software initiators are:

- The cost of the hardware (daughter cards) is avoided.
- It is possible to use one set of switches for both data and storage networking, avoiding the cost of additional switches but possibly impacting performance.
- It is possible to access other networked storage devices such as NAS, NFS, or other file servers using the same network interfaces as are used for iSCSI.

1.7 QLogic iSCSI Expansion Card for IBM BladeCenter

The iSCSI Expansion Card option (p/n 32R1923) is a hardware initiator that provides iSCSI communication from the blade server (HS20, HS40, JS20, and LS20) to an iSCSI storage device (target).

With the iSCSI Expansion Card, you gain:

- Full hardware-based iSCSI initiator and TOE for storage traffic only function (true TOE adapter as well as full iSCSI offload HBA, allows host processor hand off storage processing so it can worry about host application processing)
- Blades enabled to run diskless in a non-Fibre Channel SAN environment (Boot from iSCSI)
- Dual port card utilizing QLogic QMC4052 enabling iSCSI connectivity through switch module bays 3 and 4

  - Since HBA is presented as a true SCSI storage adapter (hiding the iSCSI interface) application interoperability is guaranteed.
  - Supports NIC and iSCSI functionality on both channels simultaneously.
- Easy-to-use GUI and CLI to manage the HBA
- Standard BladeCenter expansion card form factor (full-sized expansion card)
- NEBS-compliant for use in NGN/Telco environments
- Fully interoperable with DS300, N3700, N5200, N5500, IBM SDLT/LTO, IBM TapeSLX
- Fully interoperable with RHEL3 and 4, SLES 8 and 9, Windows® 2000/2003, and (AIX® 5.2 and 5.3 (JS20 and JS21 only))

The portfolio now includes iSCSI-SCSI over Ethernet. iSCSI technology combines Ethernet, SCSI, and TCP/IP to deliver a non-Fibre Channel SAN using stable and familiar standards. Ultimately, iSCSI is a viable alternative to local IDE and local SCSI drives, allowing customers to run servers in a diskless/stateless environment.

For more information see the announcement at the following site:

http://w3-3.ibm.com/sales/ssi/rep_ca/7/897/ENUSC05-027/ENUSC05-027.PDF

1.8 Ethernet switch modules

The supported Ethernet switch modules for the IBM BladeCenter and suggested for the Boot from iSCSI solution referred to in this chapter are:

- Nortel Networks Layer 2/3 Fiber Gigabit Ethernet Switch Module for IBM BladeCenter
- Nortel Networks Layer 2/3 Copper Gigabit Ethernet Switch Module for IBM BladeCenter
- Nortel Networks Layer 2-7 Gigabit Ethernet Switch Module for IBM BladeCenter
- Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module for IBM BladeCenter
- Cisco Systems Intelligent Gigabit Ethernet Switch Module for IBM BladeCenter
- Server Connectivity Module for IBM BladeCenter

More information about each switch module can be viewed at the following Web site:

http://www-03.ibm.com/systems/bladecenter/switch/switch_ethernet_overview.html

Attention: As of February 2006, the Nortel switch modules were acquired by Blade Network Technologies:

http://www.bladenetwork.net/

1.9 Storage

In our Boot from iSCSI SAN documented in this chapter, we only tested the use of the IBM System Storage™ N3700 and NetApp 270. To gain more knowledge regarding the IBM System Storage N series IP SAN solution, please visit the following Web site and view the data sheet:

http://www-03.ibm.com/servers/storage/network/software/iscsi/

1.9.1 IBM System Storage N3700

IBM System Storage N3700 is designed to offer you fast data access with extremely low maintenance requirements for a highly capable data storage solution. The N3700 filer integrates storage and storage processing into a single unit, facilitating affordable network

deployments. These advanced storage systems leverage a proven storage architecture and offer standard IBM System Storage N series elements, including integrated I/O, high availability via clustering, and Fibre Channel disk drives. IBM System Storage N3700 models are designed to integrate easily into existing IT environments to deliver unified storage for organizations with NAS, iSCSI, or combined environments, making enterprise-level storage a realistic goal for company sites regardless of size or staffing. For more information visit the following Web site:

http://www-03.ibm.com/servers/storage/network/n3700/

Note: It is worth mentioning that IBM System Storage N5000 and NetApp 3000 series will support this iSCSI SAN environment since they basically operate from the same ONTAP kernel and have the same functions.

More information about the N5000 can be found at:

http://www-03.ibm.com/servers/storage/network/n5000/

More information about the NetApp FAS3000 can be found at:

http://www.netapp.com/products/filer/fas3000_ds.html

1.9.2 NetApp FAS270

The NetApp FAS270 is a part of the NetApp FAS200 series systems. Based on an innovative hardware design that shrinks the traditional *filer head* to a form factor that fits within a single storage shelf, the FAS200 series provides all of the software functionality enabled by Data ONTAP and is completely compatible with and can be upgraded to the full range of NetApp systems.

The NetAPP FAS200 series is ideally suited for large, distributed enterprises with remote office and branch office storage requirements. The FAS270 is a midrange system that offers an entry-level Fibre Channel SAN solution while providing strong price/performance for NAS and iSCSI infrastructures. For more information visit the following Web site:

http://www.netapp.com/products/filer/fas200_ds.html

1.9.3 IBM TotalStorage DS300

Although this Redpaper includes the Boot from iSCSI SAN example utilizing the N-Series N3700 and NetApp FAS270 in an IBM BladeCenter environment, we are compelled to direct your attention to the Boot from iSCSI solution for the IBM TotalStorage® DS300. This information can be found in section 7.1 of the IBM Redbook, *IBM TotalStorage DS300 and DS400 Best Practices Guide*, SG24-7121, located at the following Web site:

http://www.redbooks.ibm.com/redbooks/pdfs/sg247121.pdf

2

Boot from iSCSI SAN using iSCSI HBA and initiator with failover support

This chapter illustrates the iSCSI SAN boot for HS20 Blade server for Windows 2003 Server SP1 operating system. The iSCSI SAN boot uses the QLogic iSCSI Expansion Card (hardware initiator) in conjunction with the MPIO component of the Microsoft iSCSI software initiator service for failover support.

2.1 Important prerequisites

Listed here are some important prerequisites you must adhere to prior to setting up this environment:

- From the BladeCenter MM → IO Modules, ensure that the external ports for the Ethernet switch modules in bays 3 and 4 are set to Enabled state.
- From the BladeCenter MM → IO Modules, ensure that the *management over external ports* for all the Ethernet switch modules in bays 1–4 is set to Disabled state.
- Upgrade the switch module to the current release. The latest switch module firmware can be downloaded from the following URL:

  `http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54597`
- Configure the internal switch port of the blade server used to implement the iSCSI SAN Boot environment and the N3700 storage subsystem ports in the same VLAN.
- Update the HS20 blade BIOS and Baseboard Management Controller (BMC) to the latest level.
- Update the BIOS and firmware on the hardware initiator, for example, QMC4052, to the latest level.
- Download the latest device driver for the QMC4052 HBA and create a device driver diskette to be used while installing the operating system.
- Assign static IP addresses on the N3700 storage subsystem and the QMC4052 HBA.
- Verify IP communication between the host and storage subsystem from the QMC4052 BIOS Utility.
- The Windows 2003 Server SP1 CD image must be installed. Otherwise it will bluescreen when it gets to the point of where the GUI screen initializes and it is installing devices.

2.2 Configuration overview

In this section we provide an overview of the Boot from iSCSI SAN in a single topology (Figure 2-1). For example, the host has access to the boot LUN via the switch module in bay 3. The second path will be activiated once the failover driver is installed, as shown in 2.3.6, “Multipath topology configuration” on page 31.

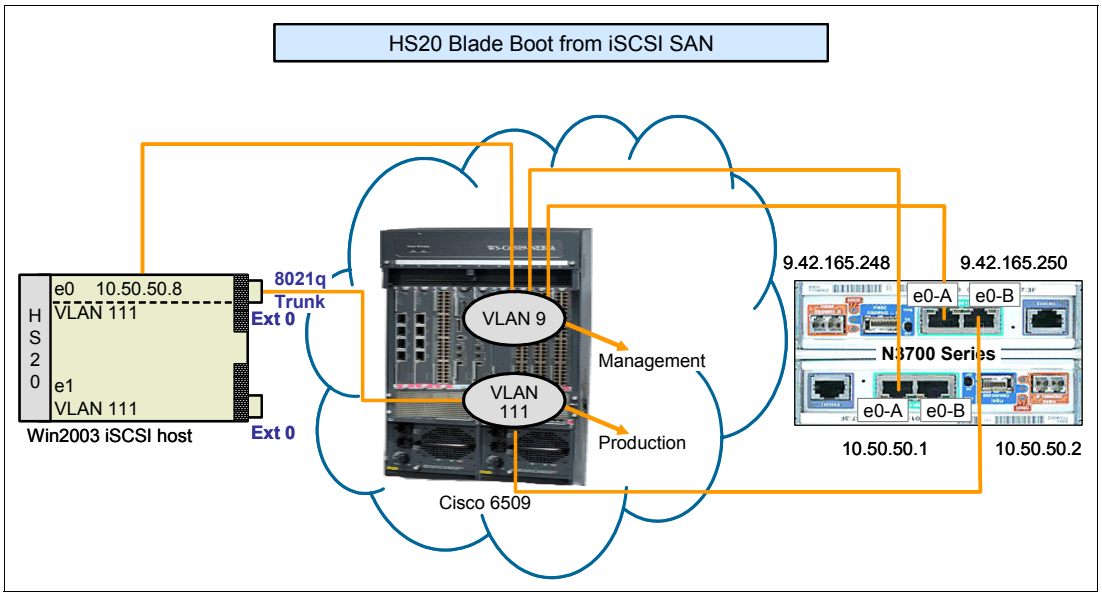

Figure 2-1 Boot from iSCSI SAN - single topology

2.2.1 Host configuration

Here we begin with the necessary requirements to configure your host system. They are:

1. Install the QLogic iSCSI Expansion Card (HBA-hardware initiator) on the HS20 blade server, if not already installed.
2. Select the CDROM as the first boot device from the BIOS.
3. Configure the parameters on the first port of the QMC4052 hardware initiator from the QLogic Fast!UTIL menu.

> **Note:** In order to modify the HBA configuration settings, select the value and press Enter.

Host adapter settings

Listed here are the host adapter settings to be updated or modified:

- iSCSI IQN

> **Important:** Initiator IQN Name
>
> The same Internet Qualified Name (IQN) name applies to all iSCSI HBA ports on the host and the Microsoft iSCSI software initiator in the system. All iSCSI HBAs and software initiators configured on the same machine must share the same IQN name.

Attention: If it is predetermined that the Microsoft MPIO driver packaged with the MS iSCSI initiator service will be used to achieve failover in a iSCSI SAN Boot environment, then it is critical that the MS IQN naming format be used. The Microsoft MPIO driver (bundled with MS iSCSI intiator service) installer utility automatically without warning or prompting overwrites the IQN name of the system as well as the iSCSI expansion card. This means that in an iSCSI SAN boot environment, the host will lose access to the boot LUN (where the OS resides) during the MPIO driver install process, thus causing the system to crash. Worse, the system cannot be recovered even after the boot path is restored by redefining the host and LUN mapping on the storage. Thus re-installing the OS is the only option available at this point.

It is very critical that the MS IQN naming format be used from the beginning for the SAN Boot environment. The format is very simple:

`iqn.1991-05.com.microsoft:{Windows Server Name}`

For the environments that uses Active Directory® server, the server name will be the full FQDN name. For example, servername.domain.com and thus the IQN name will be iqn.1991-05.com.microsoft:servername.domain.com.

- IP address
- Subnet mask
- Gateway (optional)

iSCSI boot settings

The iSCSI boot settings are:

- Adapter Boot Mode
- Primary Boot Device Setting

  Configure the IP address of the primary target device.

2.2.2 N3700 storage configuration

The necessary requirements to configure your storage system are:

1. Create the boot logical unit number (LUN).
2. Define the I/O group.
   - Host type of Windows
   - IQN of the iSCSI initiator
3. Map the LUN to the I/O group.

Note: We highly recommend that you configure your storage system prior to completing the setup of the HBA.

2.2.3 Network configuration

In this section we discuss the items that you should consider regarding your network environment:

1. For a successful install of the Windows 2003 Server SP1 operating system over the iSCSI SAN attached boot disk, it is *critical* that there exist a single and unique path from the initiator to the target device. Additionally, the OS should only see the boot LUN during the OS install process. Any additional LUNs can be added after the OS and MPIO drivers are installed.

We highly recommended that the host and target ports reside in the same VLAN.

2. Once the network configuration is complete, validate the IP connectivity from the host to the target or vice versa by pinging the IP address of the device interfaces.

For our network test environment, we used both the Nortel Networks Layer 2/3 Copper Gigabit Ethernet Switch Module for IBM BladeCenter and the Cisco Systems Intelligent Gigabit Ethernet Switch Module for IBM BladeCenter.

2.3 Boot from iSCSI with failover support implementation

The steps in this section provide you with specific instructions to implement a successful Boot from iSCSI SAN environment with failover support.

Management Module configuration

In this section we begin our setup by configuring the BladeCenter Management Module:

1. Access the Management Module using the Web browser.

   a. Log in with admin privileges to the Management Module using the browser (Inetnet Explorer or Firefox). From the BladeCenter Management Module window, click **I/O Module Tasks** → **Admin/Power/Restart**. You will see a window similar to Figure 2-2.

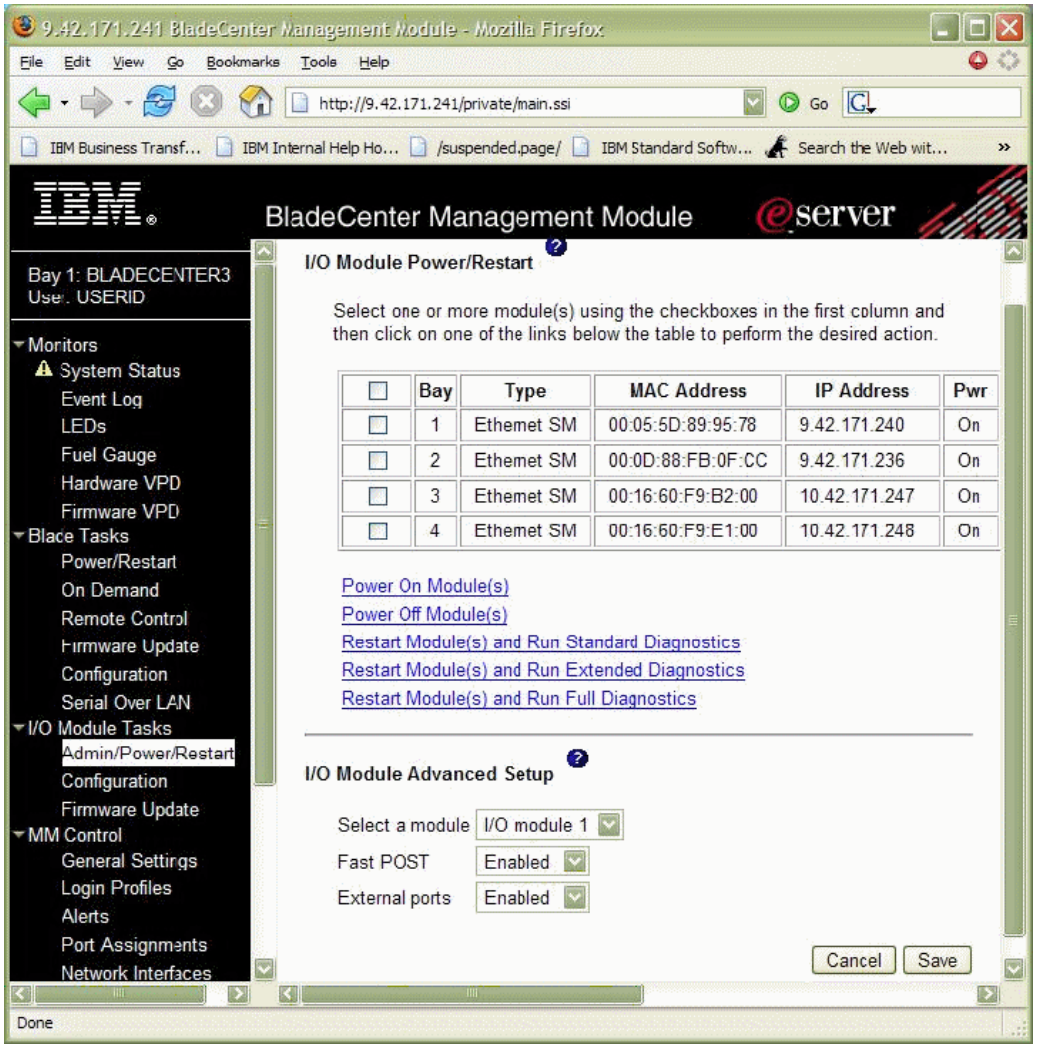

Figure 2-2 BladeCenter Management Module - Admin/Power/Restart window

b. Verify and confirm that the *external ports* for I/O modules are set as enabled. From the BladeCenter Management Module window, click **I/O Module Tasks** → **Configuration**, as shown in Figure 2-3.

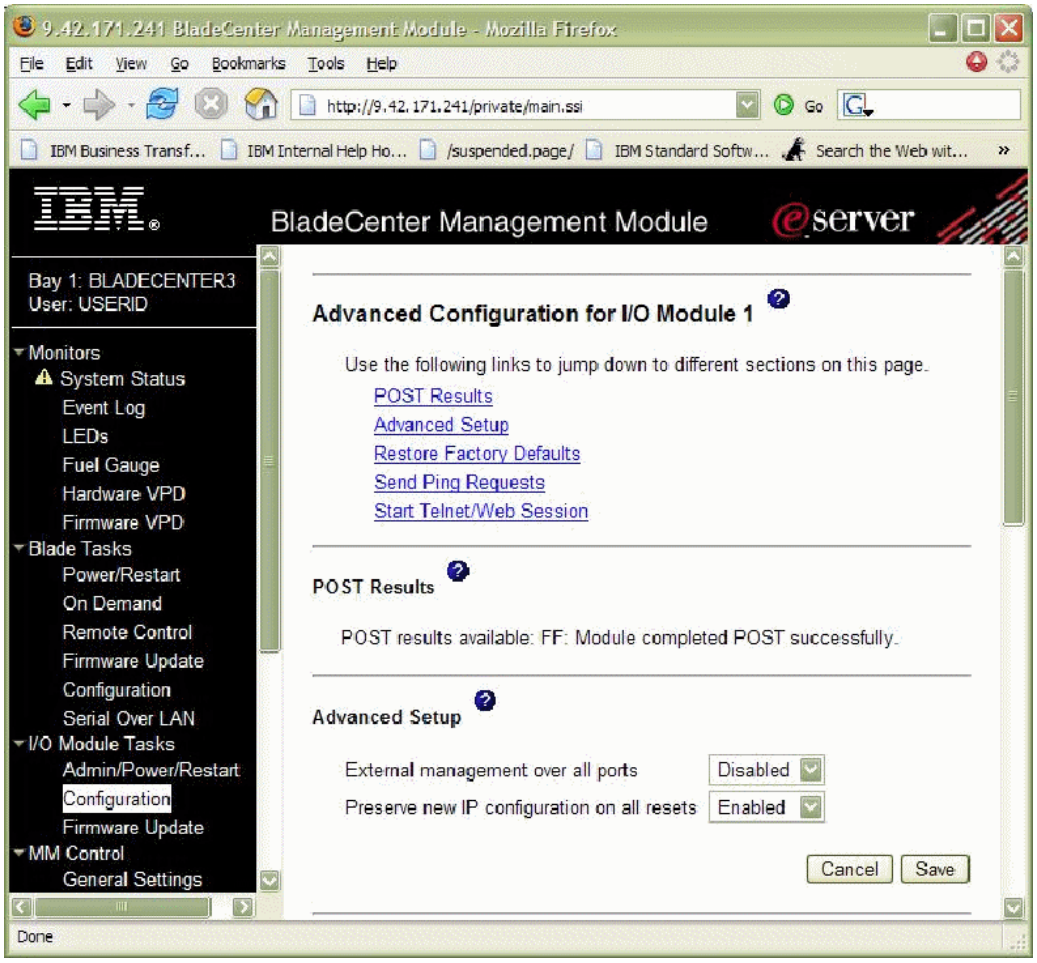

Figure 2-3 BladeCenter Management Module - configuration window

c. Verify and confirm that the *External management over all ports* is disabled on the I/O module external ports.

2.3.1 Network configuration

In this section we provide specific information pertaining to our network environment. The following network configuration is performed on the Cisco IGESM for IBM BladeCenter:

1. Access the BladeCenter Ethernet switch module in bay 3 and configure the two internal network interfaces on the blade in VLAN 111, as shown in Example 2-1. In Example 2-1, the HS20 blade is in slot 8, so internal port number 8 is configured in VLAN 111.

Example 2-1 Network configuration 1

```
Ethernet Switch Module in Bay 3
   interface GigabitEthernet0/8
   description blade8
   switchport access vlan 111
   switchport mode access
   spanning-tree portfast
```

```
Ethernet Switch Module in Bay 4
   interface GigabitEthernet0/8
   description blade8
   switchport access vlan 111
   switchport mode access
   spanning-tree portfast
```

2. Configure the external port on the BladeCenter Ethernet switch module in bay 3 to allow traffic from VLAN 111, as shown Example 2-2.

Example 2-2 Network configuration 2

```
!
   interface GigabitEthernet0/17
    description extern1
    switchport trunk native vlan 100
    switchport trunk allowed vlan 2-4094
    switchport mode trunk
   spanning-tree portfast trunk
```

3. Configure the Ethernet interfaces on storage controllers A and B as in VLAN 111, as shown in Example 2-3.

Example 2-3 Network configuration 3

```
   interface GigabitEthernet0/5
   switchport access vlan 111
   spanning-tree portfast
   switchport mode access
   !
   interface GigabitEthernet0/6
   switchport access vlan 111
   spanning-tree portfast
   switchport mode access
!
```

The corresponding Nortel network configurations for the examples above can be viewed in Appendix C, “Corresponding switch module configurations” on page 73.

2.3.2 N3700 storage configuration

This section exclusively covers the configuration components necessary to implement the iSCSI SAN boot environment. At this point it is assumed the N3700 is set up and accessible from the host. For setup details, refer to the system administration guide for the N3700 storage system.

1. Access the N3700 storage subsystem active controller console via a serial cable connected to the serial port on the active controller. In our example, we used a desktop with a serial connection to the storage controller using HyperTerminal.
2. Issue a cluster failover status command to ensure that the cluster is *up*, as shown in Figure 2-4.

```
NetApp> cf status
Cluster enabled, netapp2 is up.
```

Figure 2-4 Command-line execution for status check of cluster failover

Note: In Figure 2-4, you will noticed that the NetApp command-line software serves as an interface to the N-Series N3700 product.

Create the boot LUN

Create a boot LUN of size 10g of type Windows to be mapped to a Windows 2003 host ,as shown in Figure 2-5.

```
NetApp> lun create -s 10g -t windows /vol/testvol/win2003bootlun1

lun create: created a LUN of size: 10.0g (10742215680)
```

Figure 2-5 Command line execution for creating a boot LUN

Create an I/O group

From the N3700 console prompt, create an I/O group to define the host using the iSCSI WWN and map the LUN to the host, as shown in Figure 2-6.

```
igroup create -i -t windows winboothost iqn.1991-95.com.microsoft:win2k3-iscsi-b1
```

Figure 2-6 Command-line execution for creating an I/O group

LUN mapping

Map the LUN to the I/O group created in the preceding step, as shown in Figure 2-7.

```
NetApp> lun map /vol/testvol/win2003bootlun1 winboothost 0
```

Figure 2-7 Command-line execution for LUN mapping

Verification

Validate the LUN, I/O group, and LUN mapping configuration performed in the preceding steps, as shown in Figure 2-8 and Figure 2-20 on page 25.

```
NetAPP> lun show

/vol/testvol/testlun1.lun 200.0g (214778511360) (r/w, online, mapped)

/vol/testvol/win2003bootlun1  10.0g (10742215680) (r/w, online, mapped)

/vol/testvola/linux/linuxlun_a.1  10.0g (10737418240) (r/w, online)
```

Figure 2-8 Command-line execution for verification

2.3.3 Host configuration - QLogic iSCSI Expansion Card for IBM BladeCenter

This section illustrates how to configure the QLogic iSCSI Expansion Card for IBM BladeCenter (QMC4052 HBA) for iSCSI boot and assumes that an iSCSI target has not previously been configured on the HBA.

Before you begin:

1. Before entering the HBA BIOS to configure the HBA, verify that the HBA is directly connected to the storage device using a BladeCenter pass-through module or BladeCenter switch.
2. Ensure that the latest BIOS is installed on the HBA.
3. Ensure that the storage device and switch (if used) are both powered up and completely initialized before proceeding further.
4. If using a switch, verify that the external ports are enabled and that all devices are visible to all other devices on the switch.
5. If VLANs are enabled, the HBA and storage device must be on the same VLAN. Refer to the switch vendor documentation for details on switch configuration.

To configure the HBA, perform the following steps:

1. Power on the blade server and when you are prompted to launch the *QLogic Corporation* QMC4052 BIOS, as shown in Figure 2-9, press the Ctrl+Q keys.

```
QLogic Corporation
QMC405x iSCSI ROM BIOS Version 1.04
Copyright (C) QLogic Corporation 1993-2005. All rights reserved.
www.qlogic.com

Press <CTRL-Q> for Fast!UTIL

<CTRL-Q> Detected, Initialization in progress, Please wait...

ISP4022 Firmware Version 2.00.00.07
```

Figure 2-9 QLogic Corporation Fast!UTIL window

Note: If you use SANsurfer to update the BIOS and the system is rebooted, the card is identified as QLA405X instead of QMC405x.

2. After pressing the Ctrl+Q keys, you will enter the QLogic iSCSI Expansion Card for IBM BladeCenter hardware initiator BIOS utility. You will see a window similar to Figure 2-10.

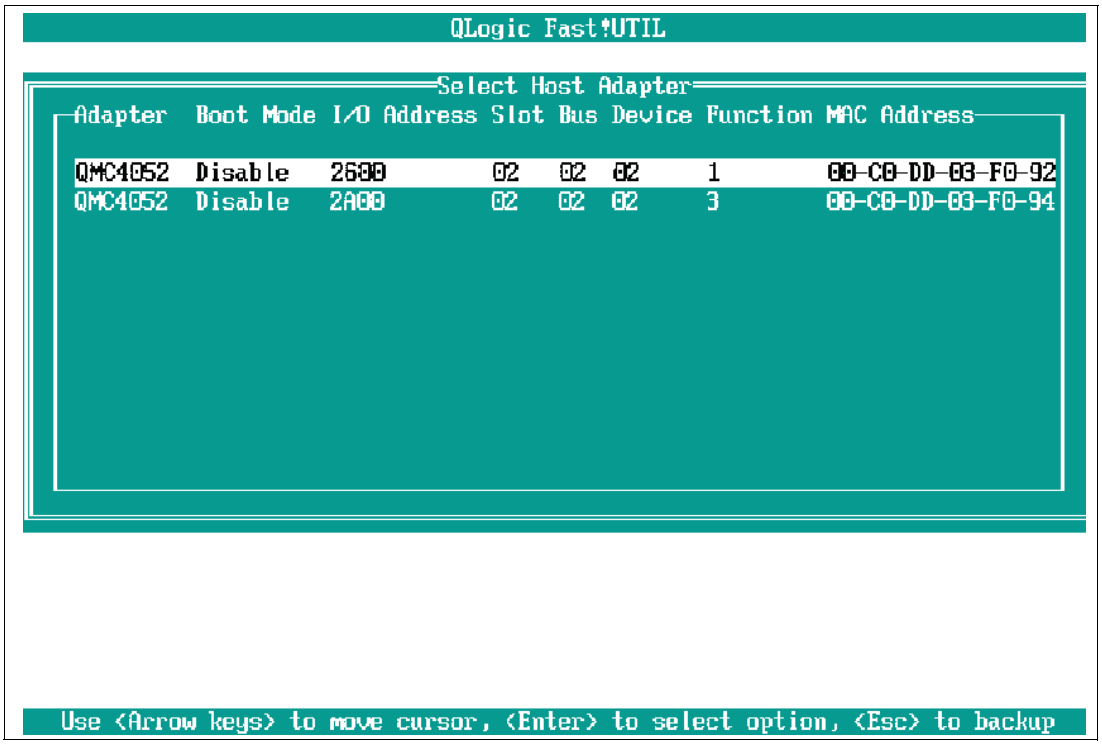

Figure 2-10 QLogic Fast!UTIL menu window

3. From the QLogic Fast!UTIL menu, select the first host bus adapter port (port 0) and press Enter. You will see a window similar to Figure 2-11.

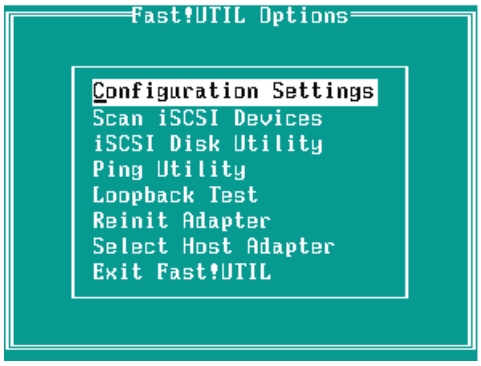

Figure 2-11 Fast!UTIL Options menu window

4. From the Fast!UTIL Options window, select **Configuration Settings** and press Enter. Next select **Host Adapter Settings** and press Enter to input the HBA configuration parameters, which must be specified before any targets can be configured. You will see a window similar to Figure 2-12.

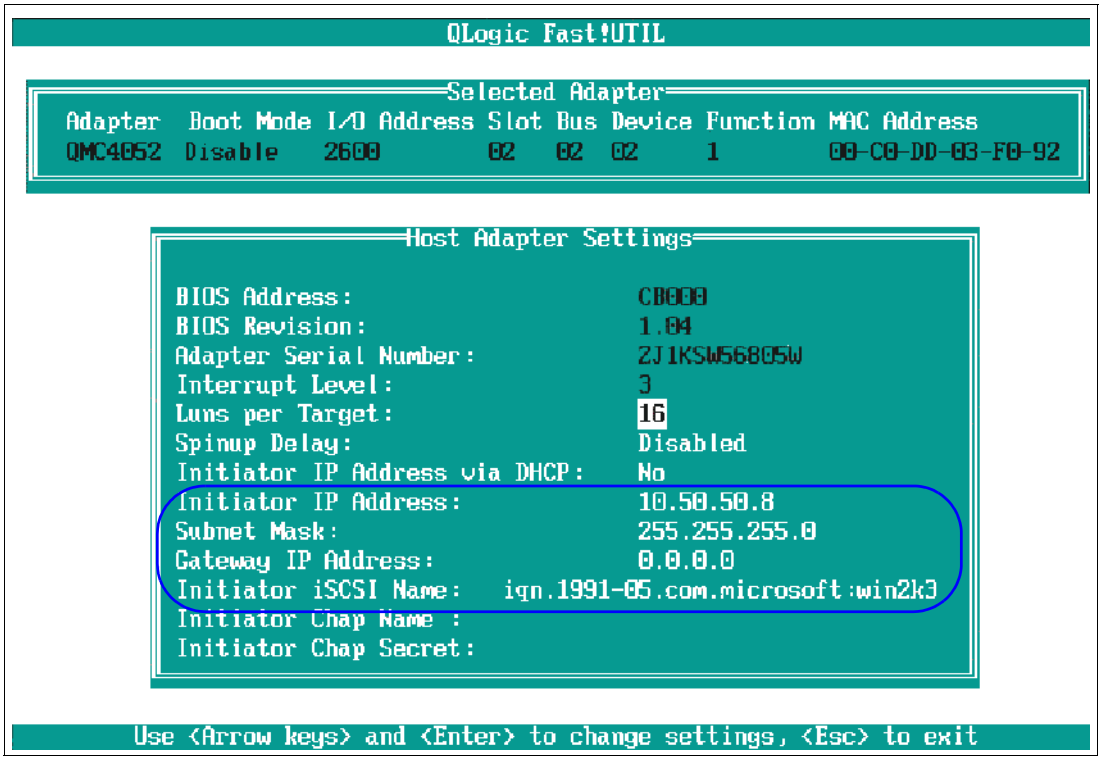

Figure 2-12 Host Adapter Settings window

Specify the Initiator IP Address, Subnet Mask, and Gateway IP Address if manually configuring these parameters. If using DHCP to obtain these parameters from a DHCP server, you need only enable the Initiator IP Address via DHCP.

Specify the Initiator iSCSI Name. A default name is provided but you have the option to manually change it. The name must be in the form of a correctly formatted RFC 3720 IQN (iSCSI Qualified name).

Optionally, if using Chap authentication with the target, specify the Initiator Chap Name and Initiator Chap Secret on this window.

5. After changing each setting and pressing Enter, press Esc to exit. You will see a window similar to Figure 2-13. In our example, we only changed the IP address, subnet mask, and iSCSI name.

Note: In order to modify the HBA configuration settings, select the value, modify the setting, and press Enter to save the changes.

Notice the format of the iSCSI name of the initiator (Figure 2-12 on page 20).

Figure 2-13 iSCSI Boot Setting window

6. From the Configuration Settings menu, select **iSCSI Boot Settings**. You will see a window similar to Figure 2-15.

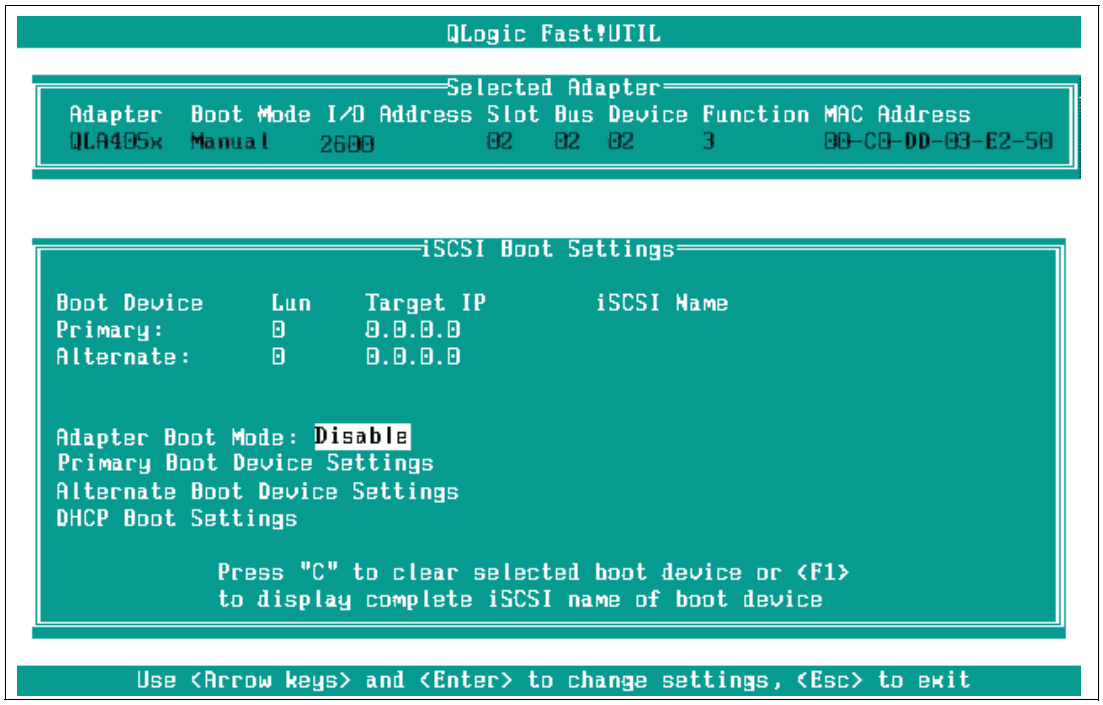

Figure 2-14 iSCSI Boot Settings window

7. From the iSCSI Boot Settings menu, select **Adapter Boot Mode** and set it to Manual to enable the BIOS on this port. In order to modify the iSCSI boot settings, select the value, make the change, and press Enter.

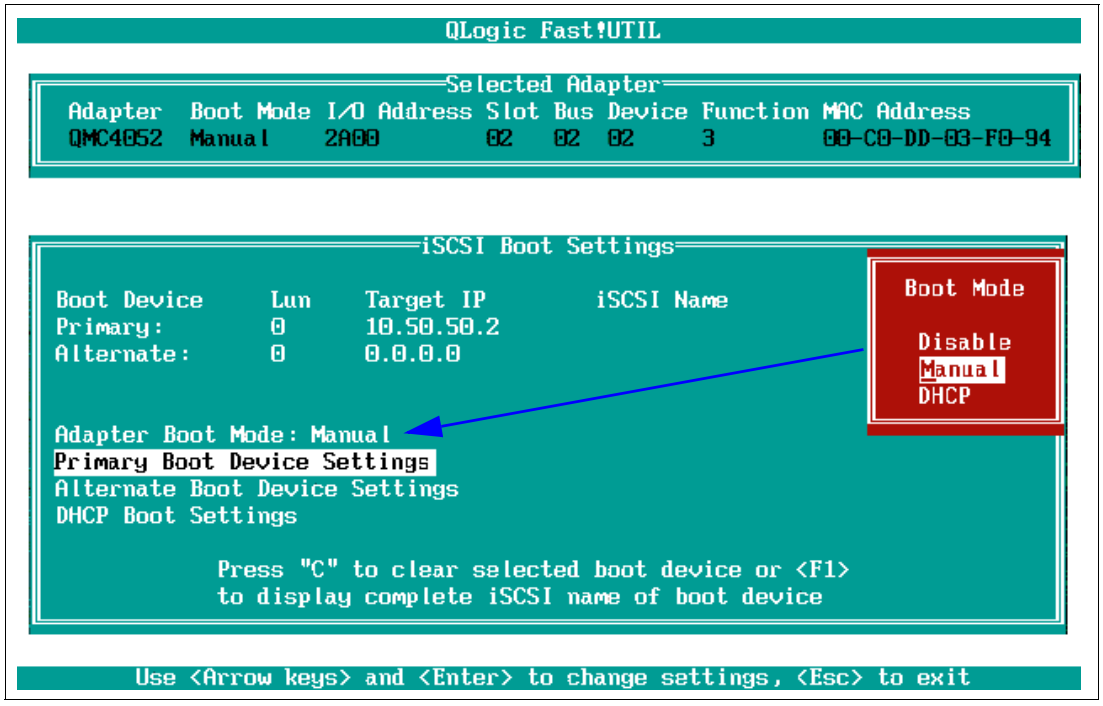

Figure 2-15 iSCSI Boot Settings window

8. Also, as in Figure 2-15 on page 22, select **Primary Boot Device Settings**. You will see a window similar to Figure 2-16.

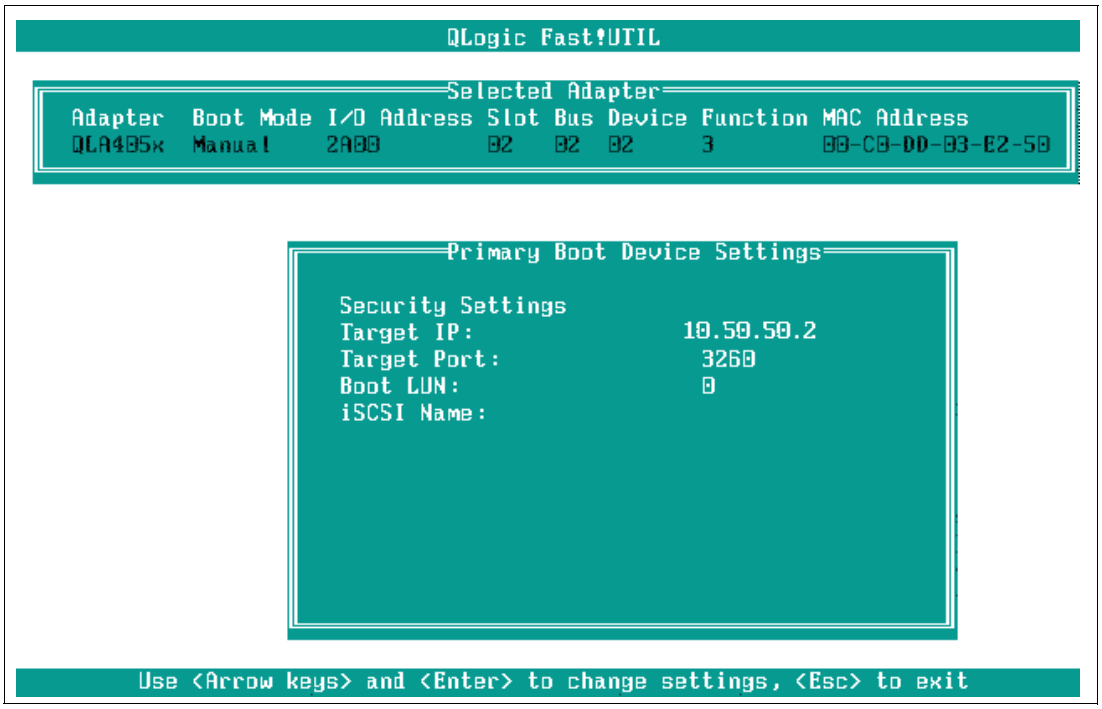

Figure 2-16 Primary Boot Device Settings window

9. Specify the IP address and Boot LUN number of the iSCSI target (Figure 2-6 on page 17). There is no need to change the default target port or provide the iSCSI name. After the IP address is configured, press Enter to change and save settings. Press Esc to exit. You will see a window similar to Figure 2-17.

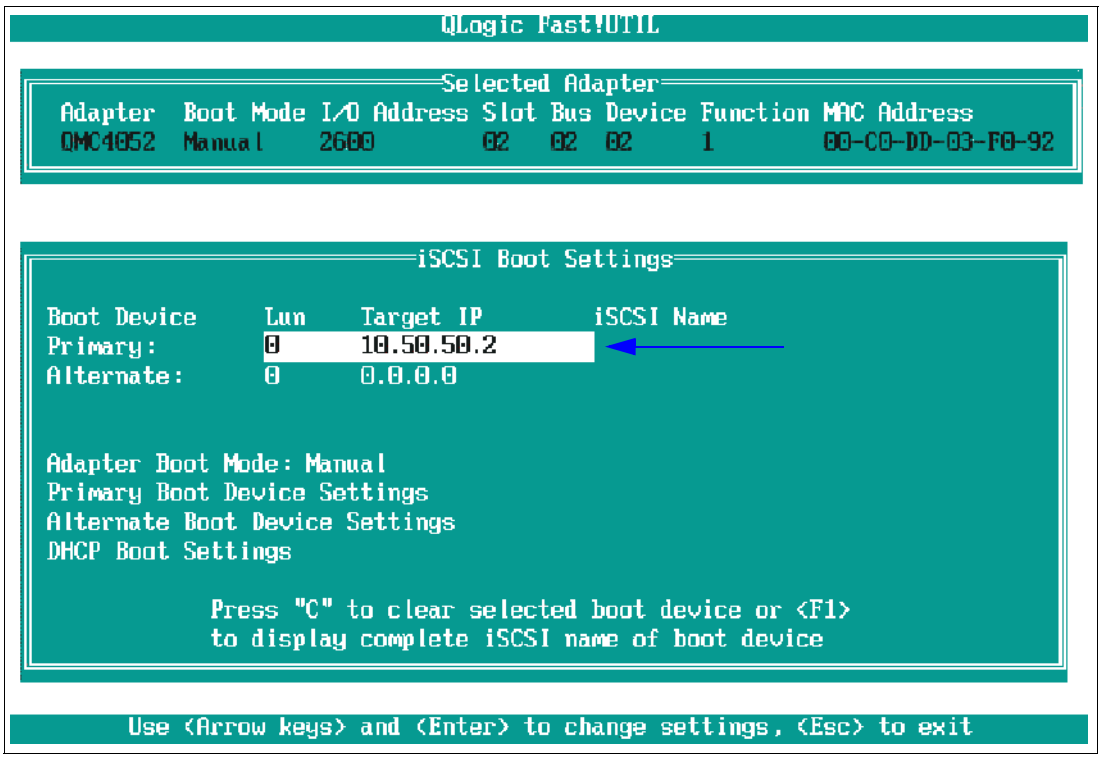

Figure 2-17 iSCSI Boot Settings window

10. To allow the initiator to query the target's IP for the target's IQN name, the target's information and initiator IP address must first be saved to the HBA and the firmware reinitialized. To perform this necessary step, press Esc twice and when prompted select **Save changes**.

Figure 2-18 Fast!UTIL Options menu

11. Select **Reinit Adapter** and press Enter. You will see a status window indicating `Reiniting Firmware... please wait`. See Figure 2-19.

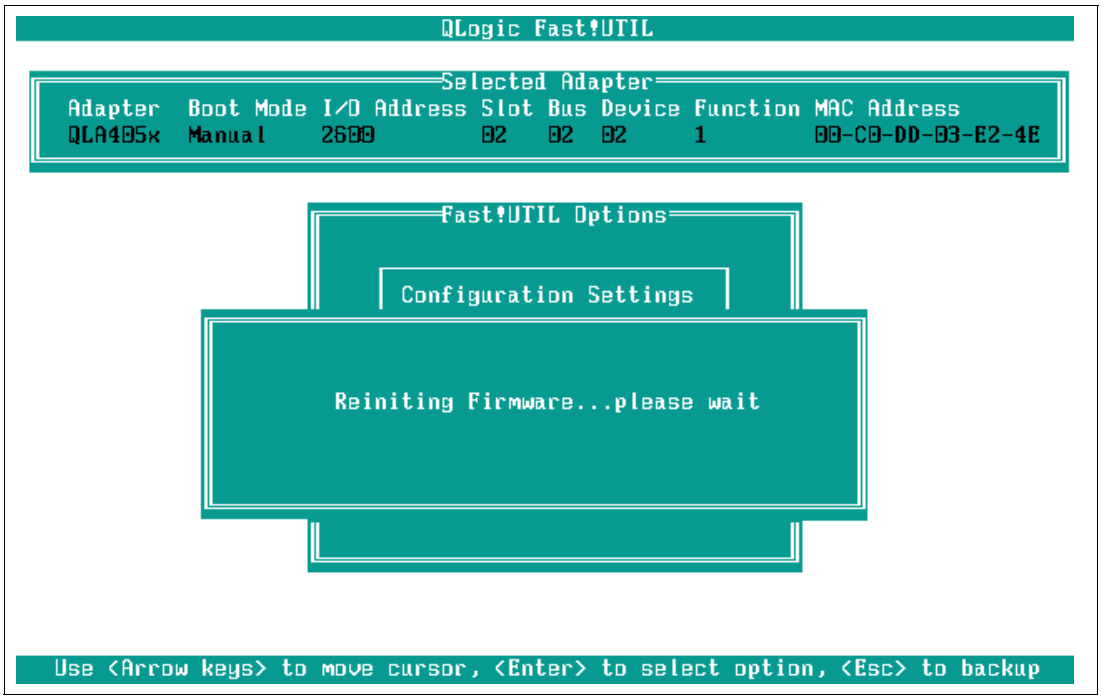

Figure 2-19 Reiniting status window

12. Return to the Configuration Settings window and select **iSCSI Boot Settings** again. Select the primary device that now has the target IP address and LUN you configured.
13. Upon successful firmware re-initialization of the initiator port, the host should successfully log in to the storage device.
14. Verify host login from storage (Figure 2-20).

```
NetApp> igroup show

winboothost (iSCSI) (ostype: windows):iqn.1991-95.com.microsoft:win2k3-iscsi-b1 (logged
in on: e0a)

linux_group (iSCSI) (ostype: linux):iqn.2000-04.com.qlogic:blade11b (not logged in)
```

Figure 2-20 Command-line execution for verification

15. Select **Scan iSCSI Devices** to discover the target device. You will see a window similar to Figure 2-21.

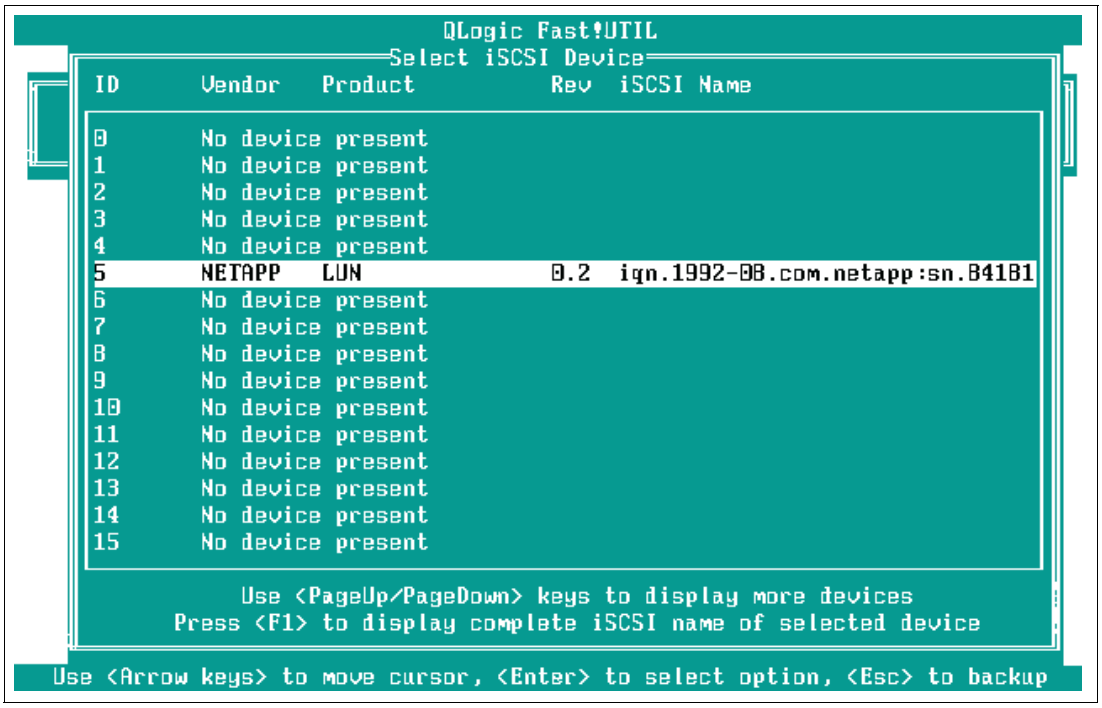

Figure 2-21 Select iSCSI Device window

16. Now select the correct iSCSI boot device and press Enter.
17. To save changes to the HBA, press Esc twice. When prompted, select **Save** changes. The HBA is now configured to allow for iSCSI boot.
18. The operating system installation can now begin as though installing to a local disk. For a trouble-free installation, ensure that there is only one path available to the operating system during the installation.

At this point, the QLogic BIOS configuration process on the host (iSCSI hardware initiator) is complete.

This completes the host and storage configuration process necessary for the primary path. Now proceed to install the Windows 2003 Server (Standard or Enterprise edition) operating system.

2.3.4 Operating system installation

Prior to installing the Windows 2003 Server SP1 operating system, you should perform the following instructions:

1. Download the QMC4052 device driver. The QLogic iSCSI Expansion Card device driver for Windows is available at:

   http://support.qlogic.com/support/oem_detail_all.asp?oemid=369

2. Download the following device driver: 32 bit STOR Miniport Storage Only.
3. Extract the driver to diskette.

4. During the loading of the Windows 2003 operating system, select **F6 option to install the iSCSI device driver**. You will see a window similar to Figure 2-22.

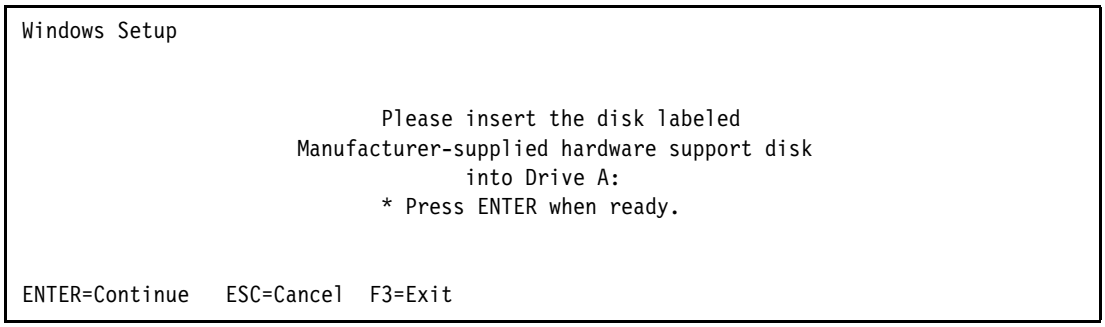

Figure 2-22 Windows Setup window

5. Insert the iSCSI device driver and press Enter. You will see a window similar to Figure 2-23 identifying the QLogic iSCSI Adapter option.

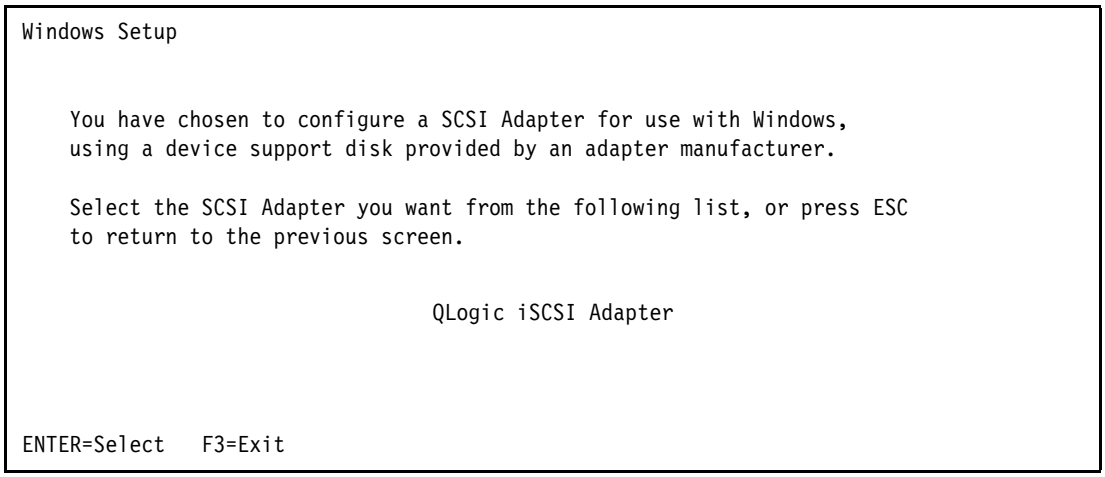

Figure 2-23 Windows Setup window

6. Press Enter to select the QLogic iSCSI Adapter. Complete the Windows 2003 installation process and proceed to the next section.

Note: If you have problems installing Windows 2003 Server and the required QLogic iSCSI Adapter drivers, make sure that you are using a Windows 2003 Server SP1 CD image.

2.3.5 Enable failover support for the boot disk

After the successful installation of your Windows 2003 Server, you will enable failover support for your boot disk. This is done by incorporating the use of multipaths. In order to achieve multipath and failover support for the boot LUN, we must install the Microsoft MPIO driver packaged with the Microsoft iSCSI Initiator Service.

Note: If you have questions regarding the Microsoft Multipath I/O, please visit the following Web site:

http://www.microsoft.com/WindowsServer2003/technologies/storage/mpio/faq.mspx

The Microsoft iSCSI Initiator Service package consists of four components:

- iSCSI port driver (iscsiprt): This is always checked and cannot be unchecked. All configurations require the port driver and thus it is always installed.
- Initiator service: This is the usermode iSCSI Initiator Service and is required for all iSCSI Software Initiator installations using iSCSI HBA or the iSCSI Software Initiator.
- Software initiator: This is the kernel mode iSCSI software initiator driver and is used to connect to iSCSI devices via the Windows TCP/IP stack using NICs. If this option is selected then the Initiator Service option is also selected automatically.
- Microsoft MPIO multipathing support for iSCSI: This installs the core MS MPIO files and the Microsoft iSCSI Device Specific Module (DSM). This will enable the MS iSCSI software initiator and HBA to perform session-based multipathing to a target that supports multiple sessions to a target.

Perform the following steps to configure failover support for the boot disk:

1. Download the Microsoft iSCSI Software Initiator Version 2.01 (build 1748):

   http://www.microsoft.com/downloads/details.aspx?FamilyID=12cb3c1a-15d6-4585-b385-befd1319f825&DisplayLang=en

2. Install the Microsoft iSCSI Initiator Service by selecting the downloaded executable file shown in Figure 2-24.

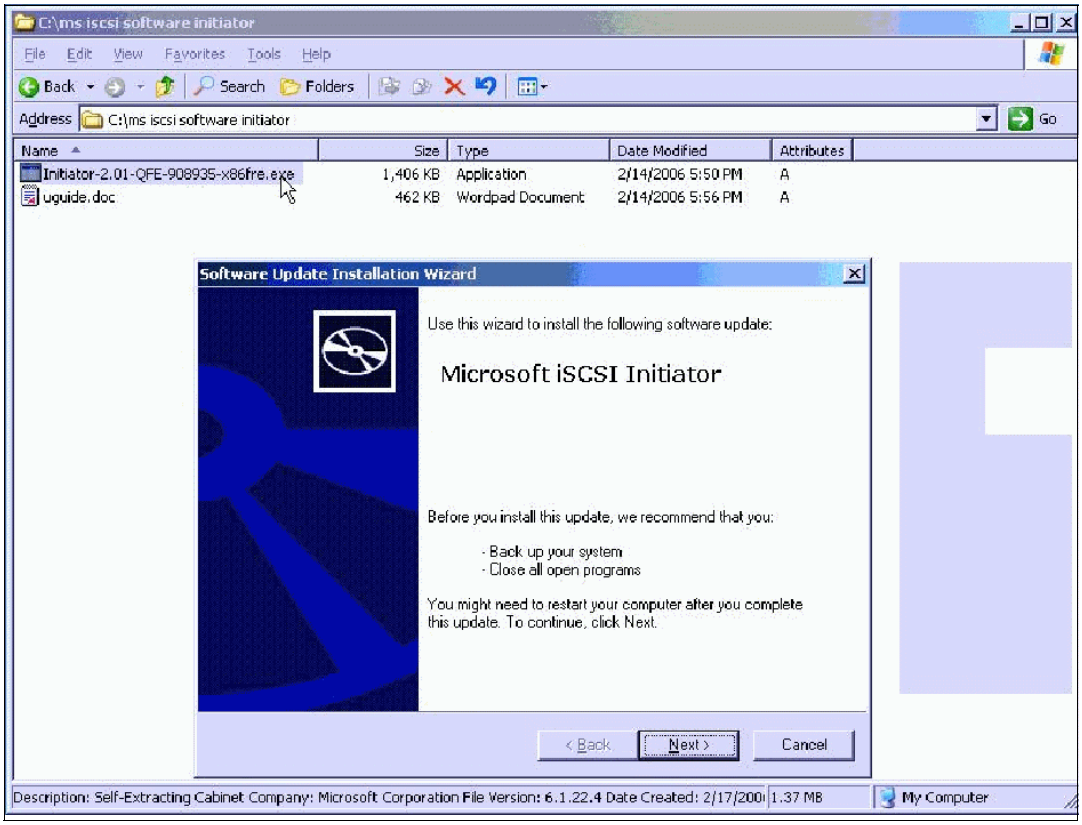

Figure 2-24 Microsoft iSCSI Initiator installation window

3. Click **Next** to proceed. You will see a window similar to Figure 2-25.

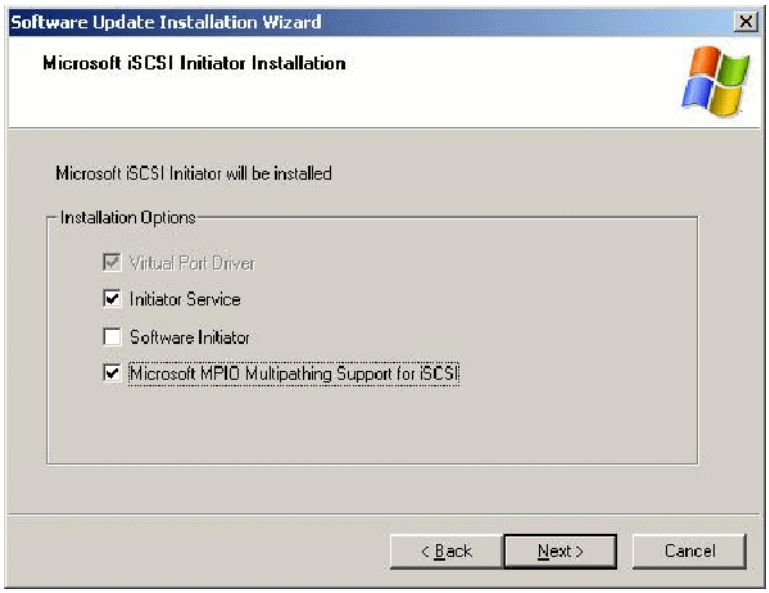

Figure 2-25 Microsoft iSCSI Initiator Installation window

4. The following installation options should be checked:
   - Virtual Port Driver
   - Initiator Service
   - Microsoft MPIO Multipathing Support for iSCSI

Note: Notice that the Software Initiator option is unchecked because the QLogic iSCSI Expansion Card for IBM BladeCenter hardware initiator (HBA) is installed on this system.

5. Click **Next** to proceed. When the installation is complete, you will see a window similar to Figure 2-26.

Figure 2-26 Software Update Installation Wizard completion window

6. Ensure that the Do not restart now check box is selected because additional paths need to be added from the host to the target device prior to restarting the system. Click **Finish**.

2.3.6 Multipath topology configuration

Figure 2-27 illustrates dual paths from each QLogic iSCSI Expansion Card (hardware initiator) port to the two storage controllers (also known as the N3700 CPU modules) for high availability and failover support.

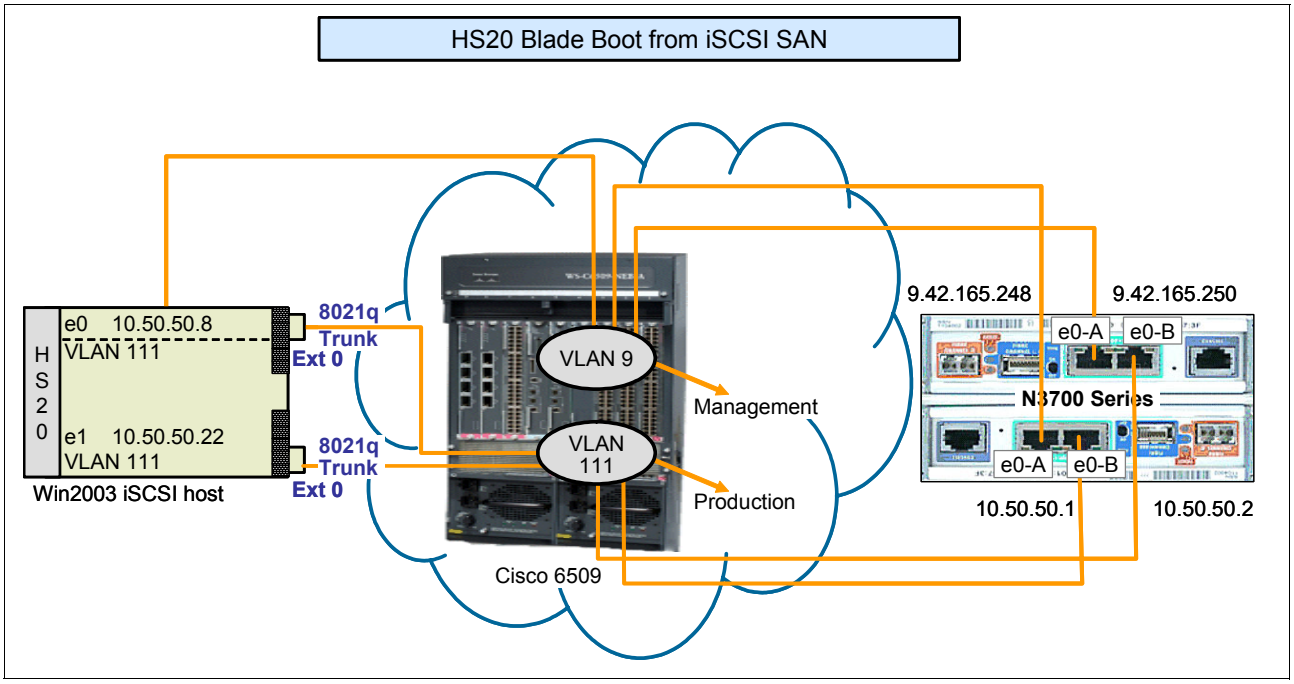

Figure 2-27 Multipath topology

Network configuration for multipath topology

In order achieve failover capability, configure the external port on the BladeCenter Ethernet switch module in bay 4 to allow traffic from VLAN 111, as shown Example 2-4.

Example 2-4 Example network configuration

```
!
interface GigabitEthernet0/17
description extern1
switchport trunk native vlan 100
switchport trunk allowed vlan 2-4094
switchport mode trunk
```

Perform the following steps to set up the multipath environment:

1. At this point, restart the host. During reboot, press Ctrl+Q to interrupt the boot process and access the QLogic Fast!UTIL shown in Figure 2-28.

```
QLogic Corporation
QLA405x iSCSI ROM BIOS Version 1.04
Copyright (C) QLogic Corporation 1993-2005. All rights reserved.
www.qlogic.com

Press <CTRL-Q> for Fast!UTIL

<CTRL-Q> Detected, Initialization in progress, Please wait...

ISP4022 Firmware Version 2.00.00.07
```

Figure 2-28 QLogic Corporation Fast!UTIL window

2. In this multipath topology, each host adapter *port* has access to one Ethernet interface on N3700 storage controllers A and B.

Important: Notice in the topology diagram (Figure 2-27) that one Ethernet interface on each controller is dedicated into VLAN 9, thus prohibiting access from the iSCSI initiator to devices in VLAN 9. Whereas the second Ethernet interface is mapped into production VLAN 111 for iSCSI host access. This design is required for the failover to work correctly.

The following steps illustrate the multipath configuration procedure.

Note: Ensure that the network configuration from initiator to the storage is complete, meaning that each QMC4052 host adapter *port* has access to the Ethernet interfaces on N3700 controllers A and B. Refer to 2.3.1, “Network configuration” on page 15. Also see Appendix B, “N-Series N3700 and NetApp FAS270 CPU Module” on page 71, for a better view of the N3700 controllers.

a. Enable the path from QLogic iSCSI Expansion Card (HBA) port 0 to access the Ethernet interface on N3700 controller B. From the QLogic Fast!UTIL menu, configure the Alternate Boot Device IP address by selecting the **QLogic HBA host adapter 0** → **Configuration Settings** → **iSCSI Boot Settings** → **Alternate Boot Device Settings**, as shown in Figure 2-29.

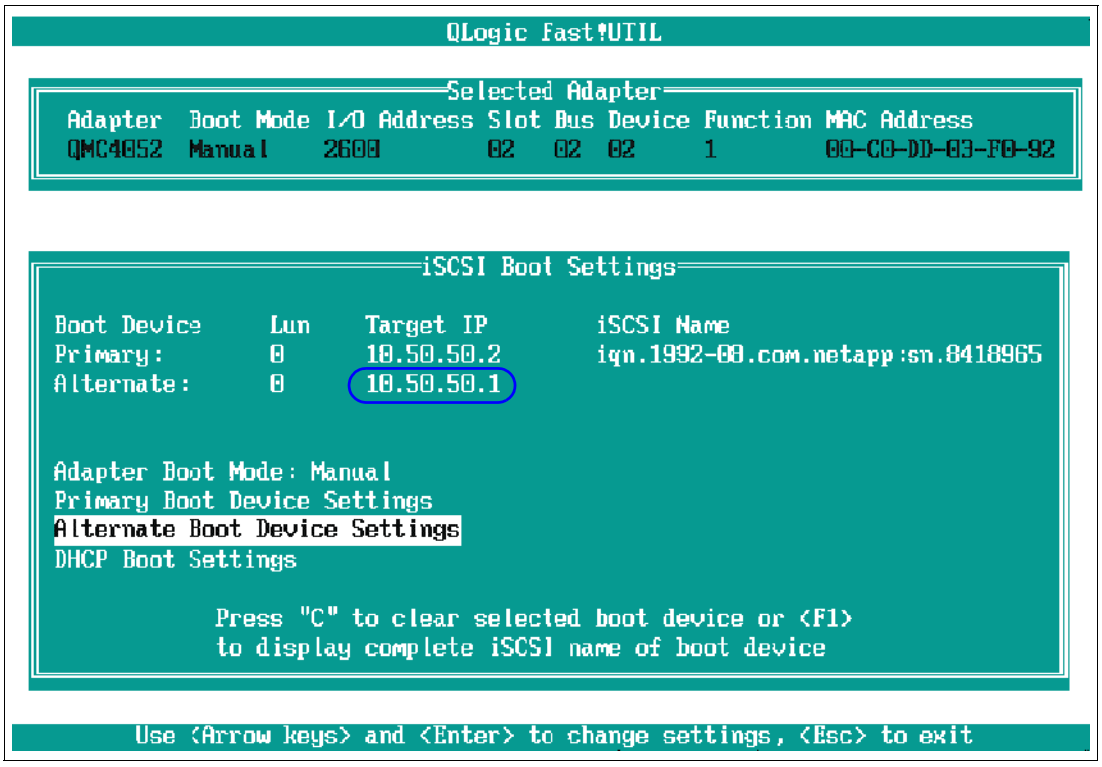

Figure 2-29 QLogic Fast!UTIL window

Note: In order to modify the HBA configuration settings, select the value and press Enter.

b. Now configure host adapter *port 1* on QMC4052 to access the N3700 controllers A and B. This requires performing the following sequence of steps:

i. Select the HBA host adapter port 1 as shown in Figure 2-30.

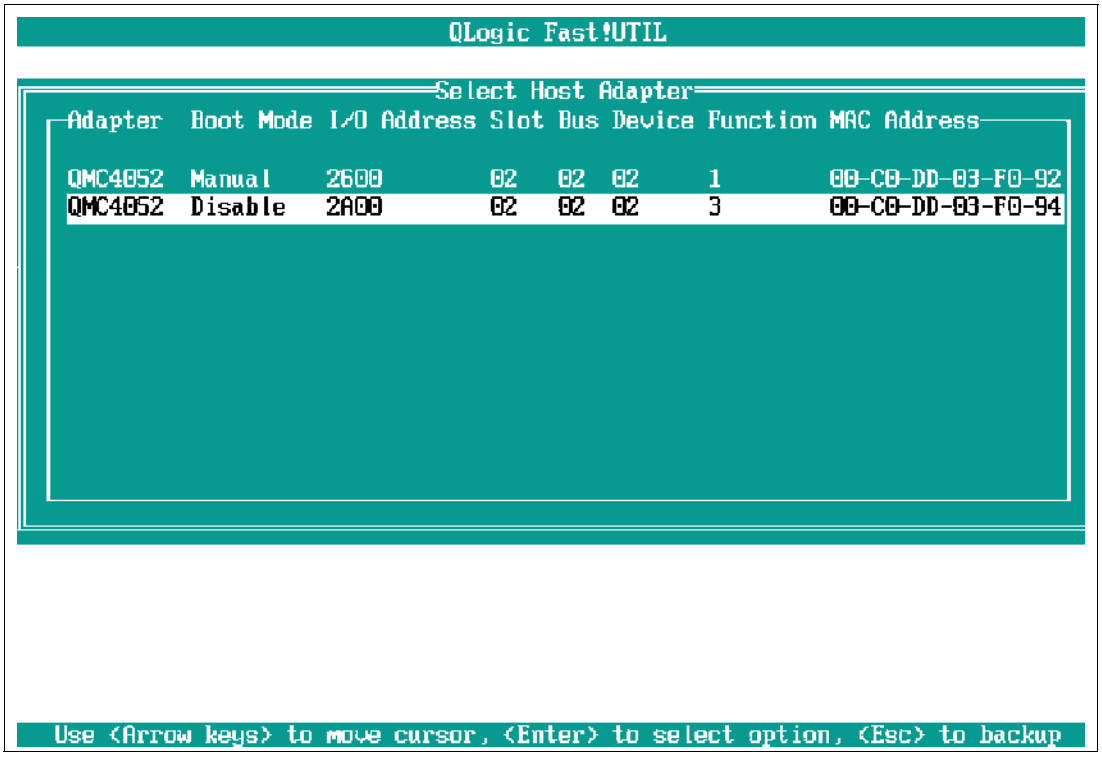

Figure 2-30 QLogic Fast!UTIL window

ii. Select **Configuration Settings** → **Host Adapter Settings**. Configure the IP address on QMC4052, as shown in Figure 2-31.

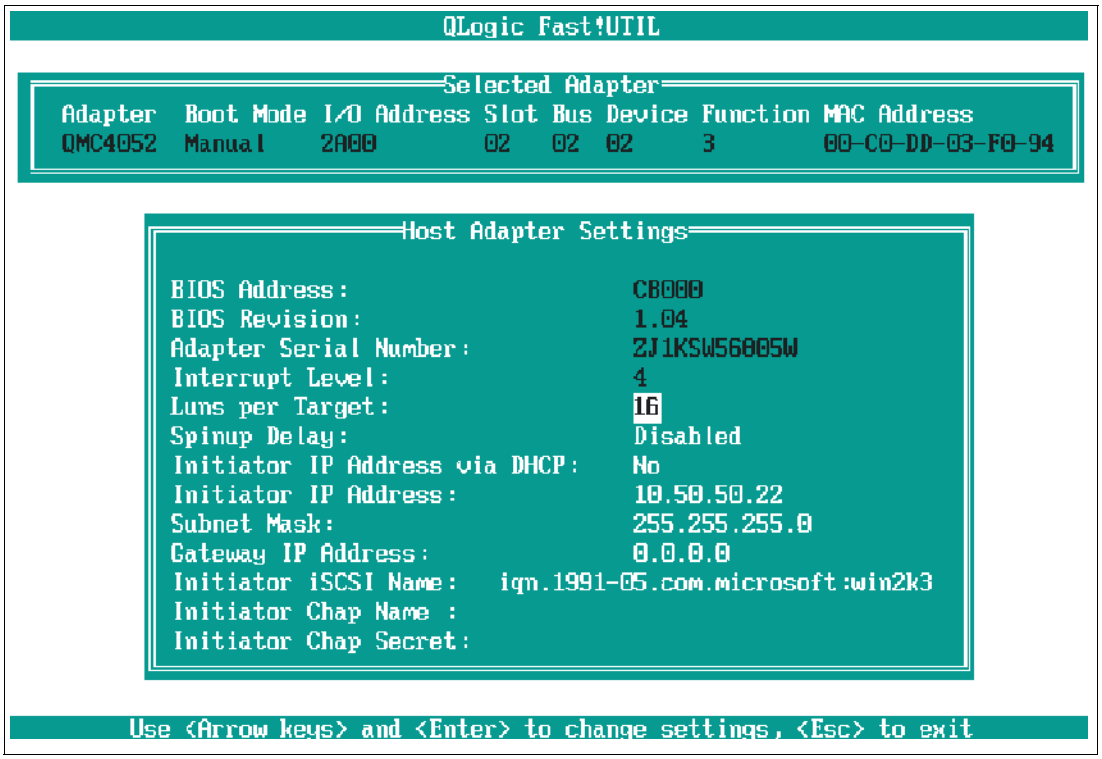

Figure 2-31 Host Adapter Settings window

iii. Configure the Primary boot device IP address (Figure 2-32).

iv. Configure the Alternate boot device IP address.

v. Set Adapter Boot Mode to Manual.

See 2.3.8, “Failover verification” on page 42 to validate the failover configuration from the QLogic iSCSI Expansion Card BIOS.

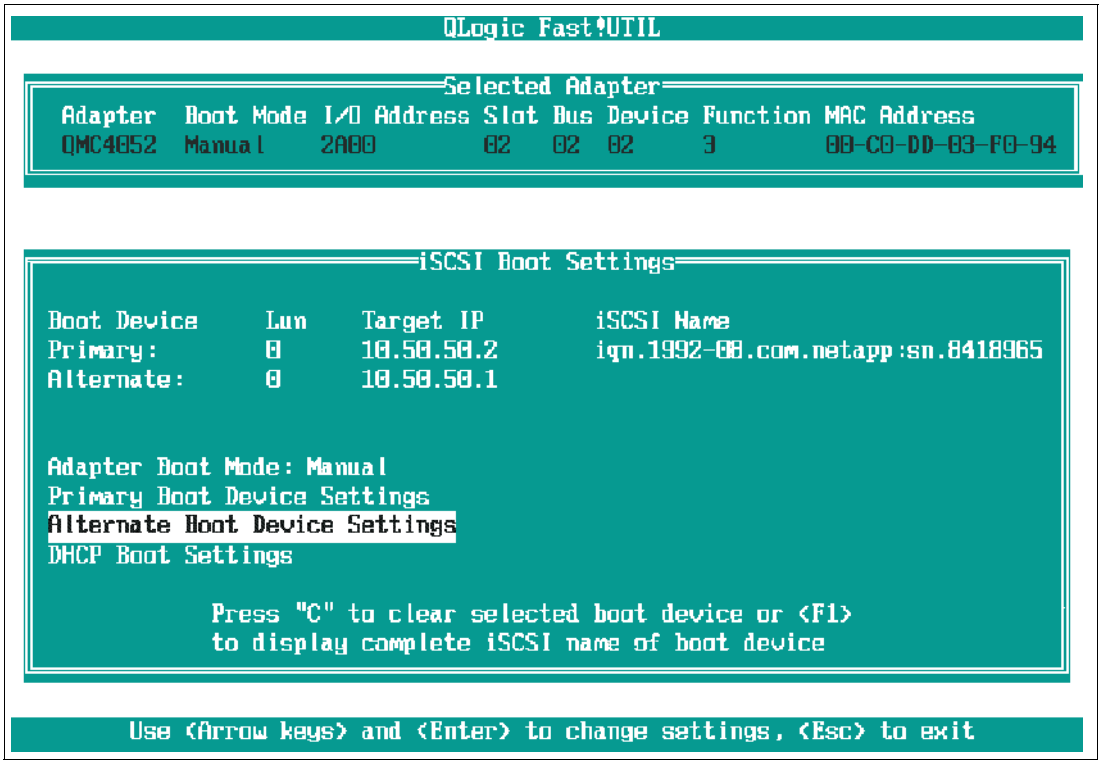

Figure 2-32 iSCSI Boot Settings window

Upon configuration completion, select Esc to save and exit the QLogic Fast!UTIL.

2.3.7 Verify end-to-end network connectivity

The network connectivity from the initiator to the target can be verified by pinging the IP addresses of Ethernet interfaces on N3700 controllers A and B from the QLogic Fast!UTIL menu for both QMC4052 host adapter ports, as shown in the following examples.

From the QLogic Fast!UTIL menu, select **Ping Utility** for each QMC host adapter, as shown in Figure 2-33.

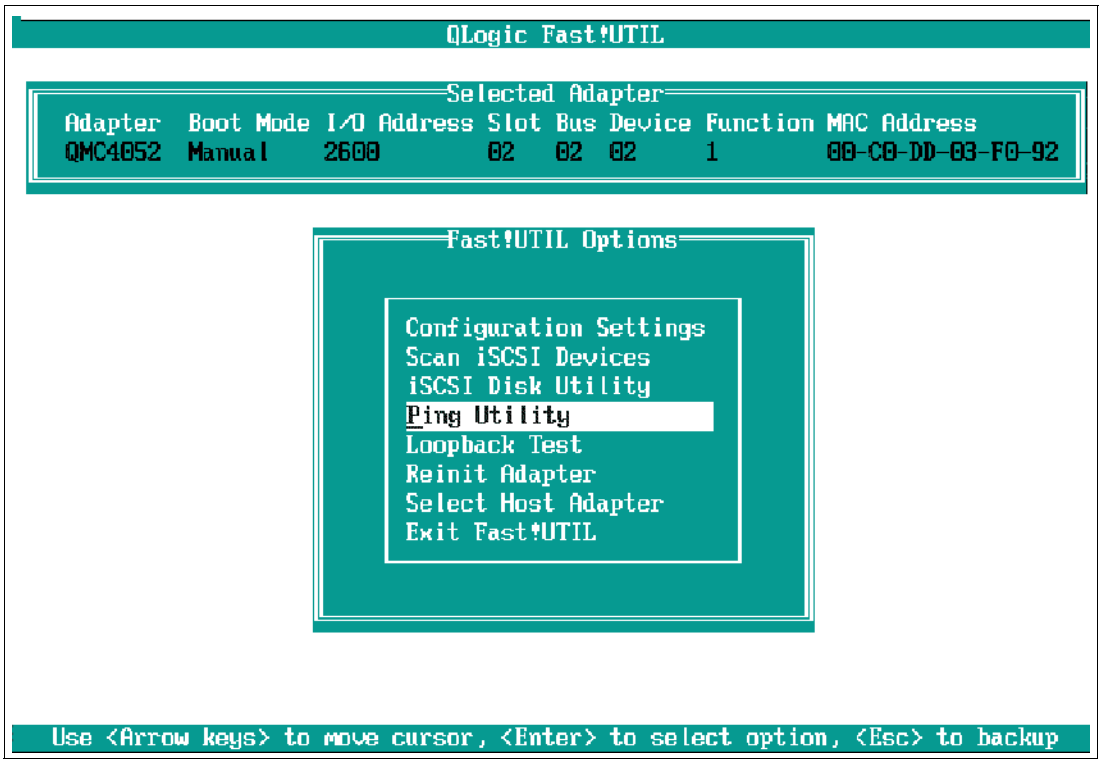

Figure 2-33 Fast!UTIL Options window

Figure 2-34 displays a successful ping response from the QMC4052 host adapter to the N3700 Storage Controller A at IP address 10.50.50.2 (e0-B).

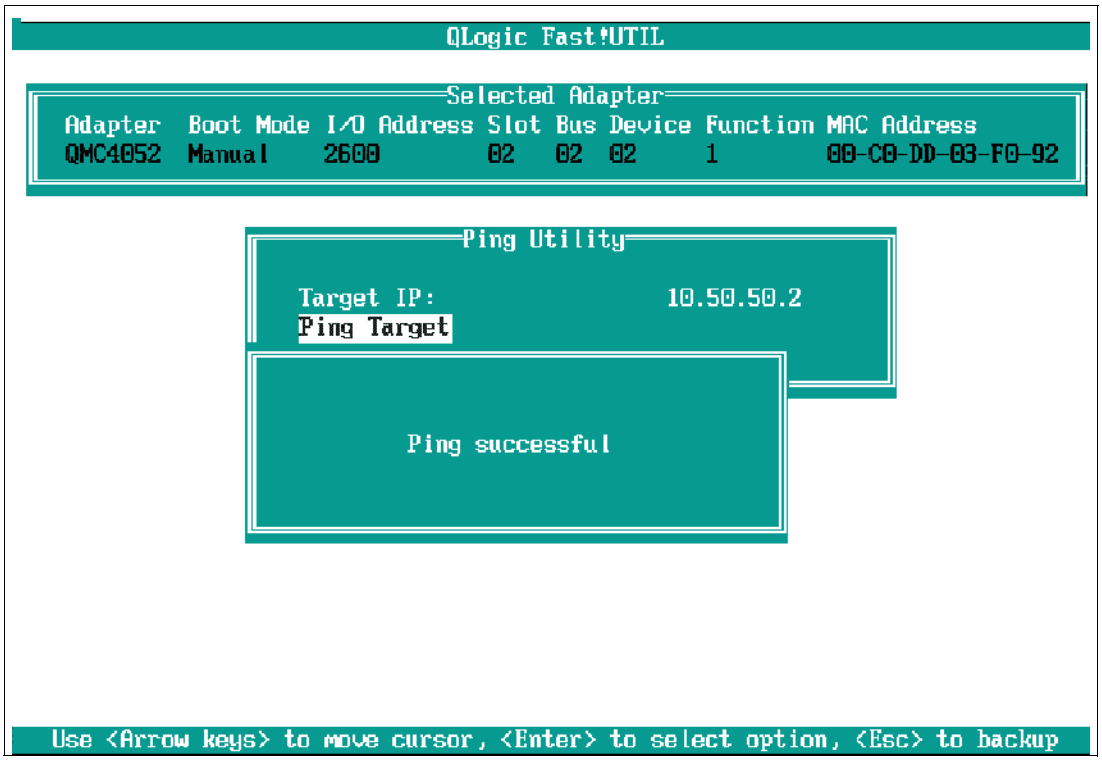

Figure 2-34 Ping Utility window

Figure 2-35 displays a successful response from the QMC4052 host adapter to the N3700 Storage Controller A at IP address 10.50.50.1 (e0-A).

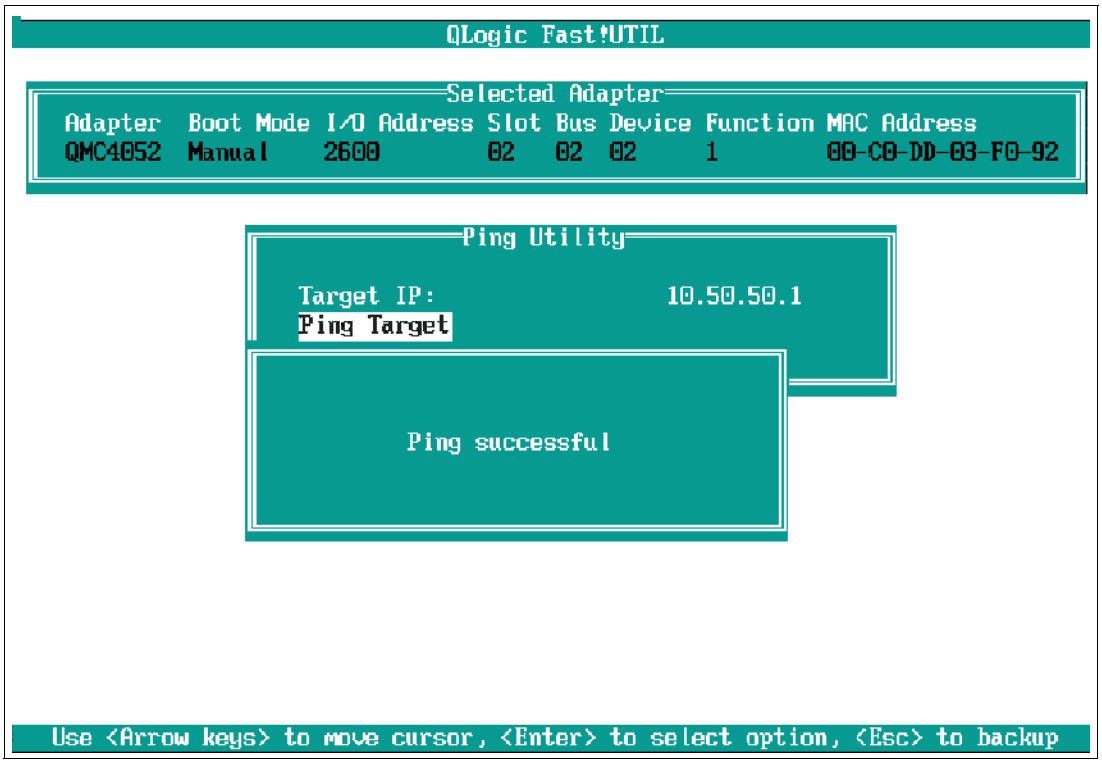

Figure 2-35 Ping Utility window

Figure 2-36 displays a successful response from the QMC4052 host adapter to the N3700 Storage Controller B at IP address 10.50.50.2 (e0-B).

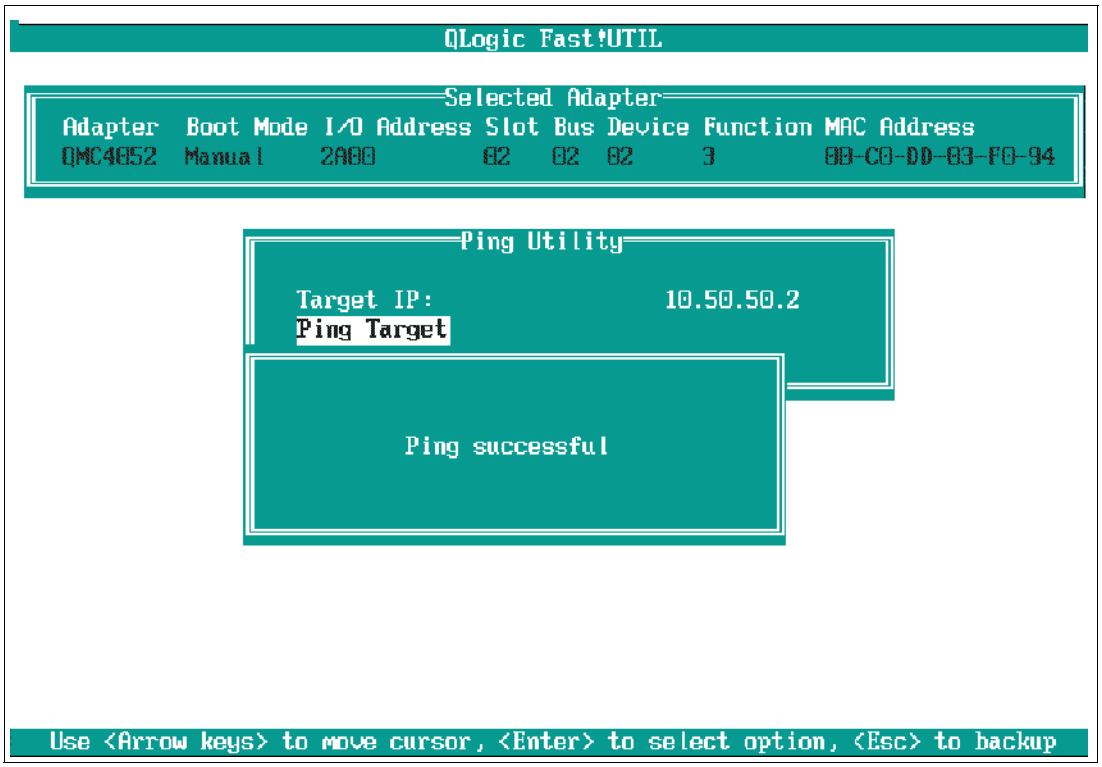

Figure 2-36 Ping Utility window

Figure 2-37 example displays a successful response from the QMC4052 host adapter to the N3700 Storage Controller B at IP address 10.50.50.1(e0-A).

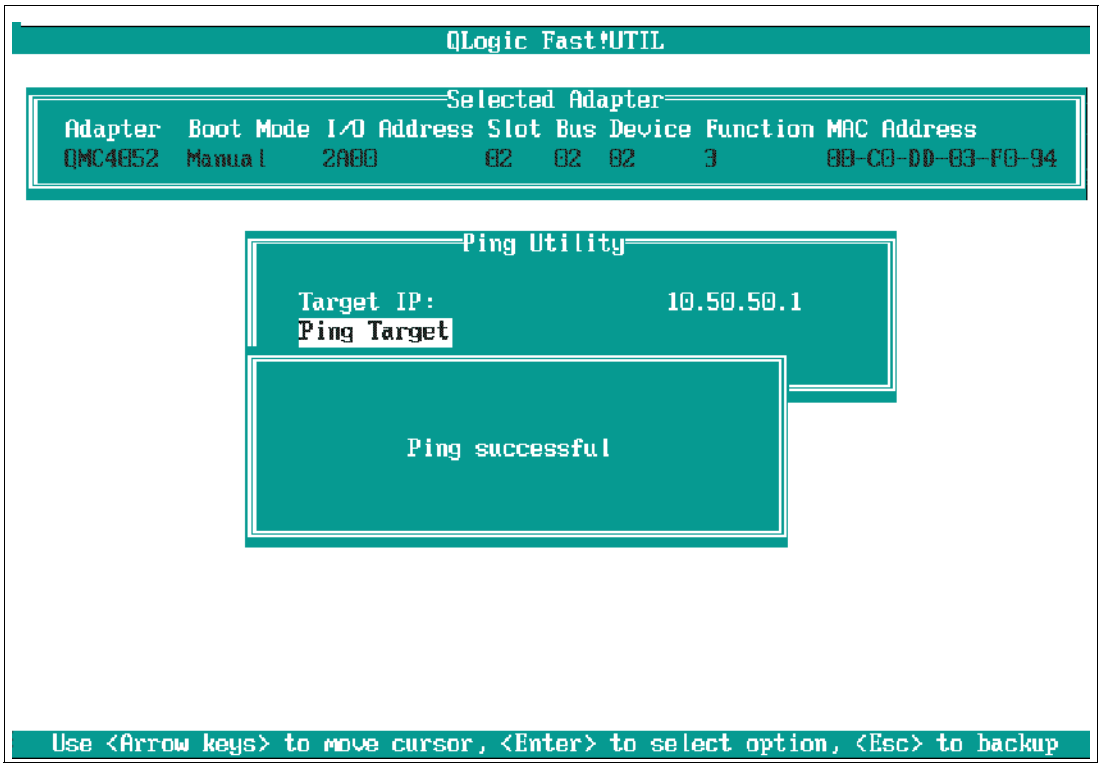

Figure 2-37 Ping Utility window

Verify network connectivity

Notice that all the device mac-addresses are learned in VLAN 111 as expected on the network core switch (Cisco 6509) to which the HS20 blade and N3700 storage devices are connected. The mac-addresses were learned dynamically by the switch in VLAN 111 as a result of the successful pings in the preceding examples. See Figure 2-38.

Note: The command **`show mac-address-table vlan 111`** is not applicable on a different vendor's switch.

```
#show mac-address-table vlan 111
          Mac Address Table
-------------------------------------------
Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
 111    0011.211b.55c0    STATIC      CPU
 111    0100.0ccc.cccc    STATIC      CPU
 111    0100.0ccc.cccd    STATIC      CPU
 111    0100.0cdd.dddd    STATIC      CPU
 111    000f.904b.a8fe    DYNAMIC     Gi0/17
 111    00a0.9802.14e7    DYNAMIC     Gi0/17      <-- Storage Port MAC Address
 111    00a0.9802.14ed    DYNAMIC     Gi0/17
 111    00c0.dd03.f092    DYNAMIC     Gi0/8       <-- Host Port MAC Address
 111    00c0.dd03.f094    DYNAMIC     Gi0/17
Total Mac Addresses for this criterion: 9
```

Figure 2-38 Show MAC address table for vlan 111 execution

2.3.8 Failover verification

The following failover tests are performed from the host side by disabling the active path from the host to the network. This failover test is performed by starting a file copy job from the CDROM to the *C drive* (for example, boot LUN from the operating system and also by booting the host via the secondary path while the primary path is in failed state.

In this section we test our failover connections to ensure that failover is working correctly:

1. From your Windows desktop, open the Windows Device Manager menu by right-clicking **My Computer** → **Manage**, as shown in Figure 2-39.

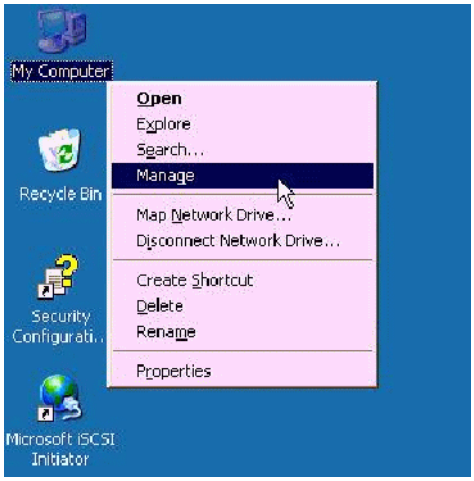

Figure 2-39 Windows Device Manager selection window

2. Once the Computer Management window is displayed, verify redundancy (dual paths) at the disk and controller levels, as shown in Figure 2-40.

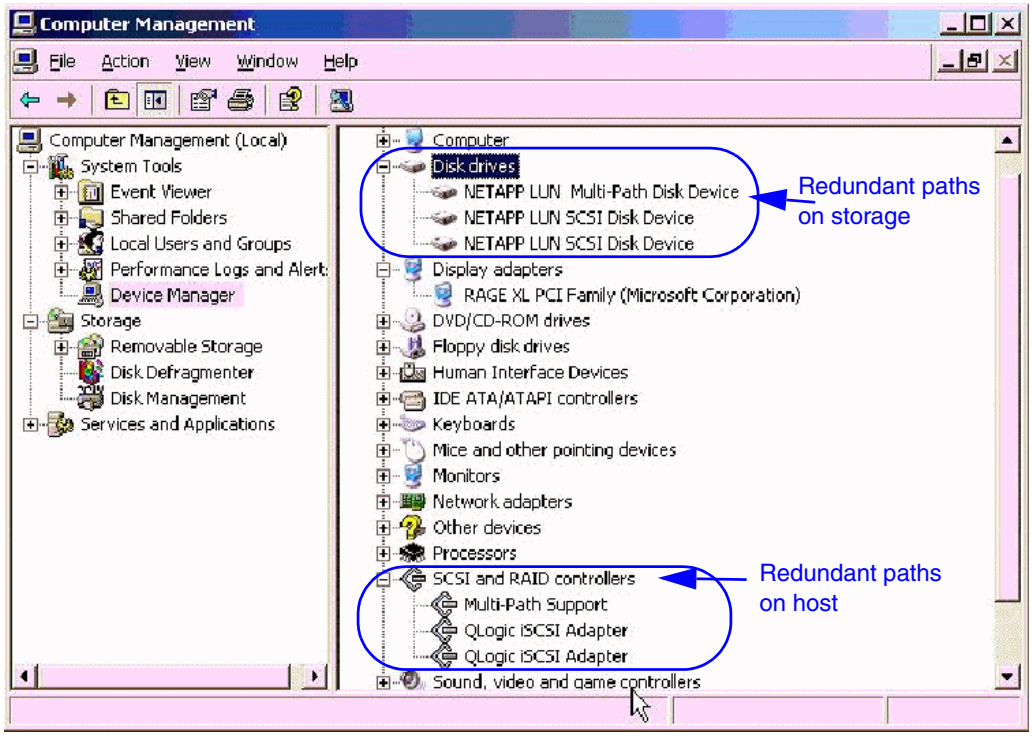

Figure 2-40 Computer Management window

3. Now start I/O to the primary path by copying files from CDROM to the boot disk *C drive* and manually inject the failure by disabling the host controller or storage host adapter on the switch or pulling the Ethernet cable from the primary storage controller. The I/O should successfully resume once a new path from the host to the storage is established.

4. The failover of the boot LUN can also be verified by rebooting the host and confirming that the initiator finds the boot LUN via the secondary path and successfully initializes.
5. From the Windows Device Manager view, there will be a single path from the host to the storage, as shown in Figure 2-41.

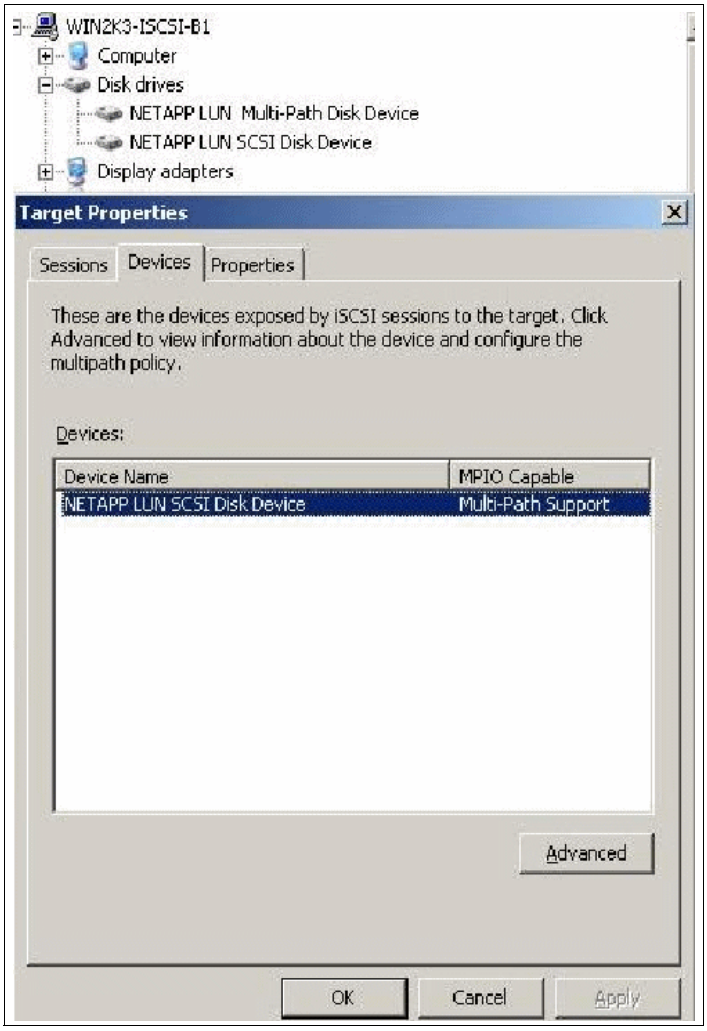

Figure 2-41 Target Properties window

This completes the failover verification process and iSCSI SAN Boot for Windows 2003 SP1 on the HS20 blade server with N3700 solution.

3

iSCSI SAN High Availability

This chapter demonstrates the implementation of High Availability whether you are using a hardware initiator or software initiator solution on an IBM BladeCenter HS20. The configuration used in our testing is described in the following sections.

3.1 Topology, naming, and addressing conventions used in these examples

Figure 3-1 on page 47 is an example of the iSCSI SAN topology based on our lab configuration using the hardware initiator.

Our lab configuration used for testing included the following:

- In a standard chassis, blades 1 and 2 had hardware initiators with the associated QLogic HBA card. Note that in normal practice, either hardware or software initiators would have been chosen for the entire configuration. Our test bed included both, so that we could test both options.
- Four Nortel L2/3 switches were used in the chassis. Bays 1 and 2 were for data traffic and bays 3 and 4 were for traffic to and from the iSCSI hardware initiators.
- The iSCSI storage array used was the IBM N-Series N3700 (utilizes the same operating system as the NetApp FAS270). The storage array had two controllers, each of which has two Gigabit Ethernet ports. The names assigned were *netapp* and *netapp2*. The two controllers were configured in cluster failover mode.
- The two Nortel L2/3 switches in bays 3 and 4 are each connected to both of the controllers on the storage array.

Note: The example used in this section works with any switch. However, the use of the Nortel configurations in this chapter will not work with a Cisco device because of the Nortel command-line format. Also, the Cisco device is unable to route the traffic between the blades and the storage device, so some of the HA attributes of the Nortel design will not be available. In the design tested, the 10.10.10.x server ports can reach the 10.10.20.x storage ports. However, the Cisco device is unable to do this without inserting an external router between the switches and the storage device.

The topology of the test network is shown in Figure 3-1.

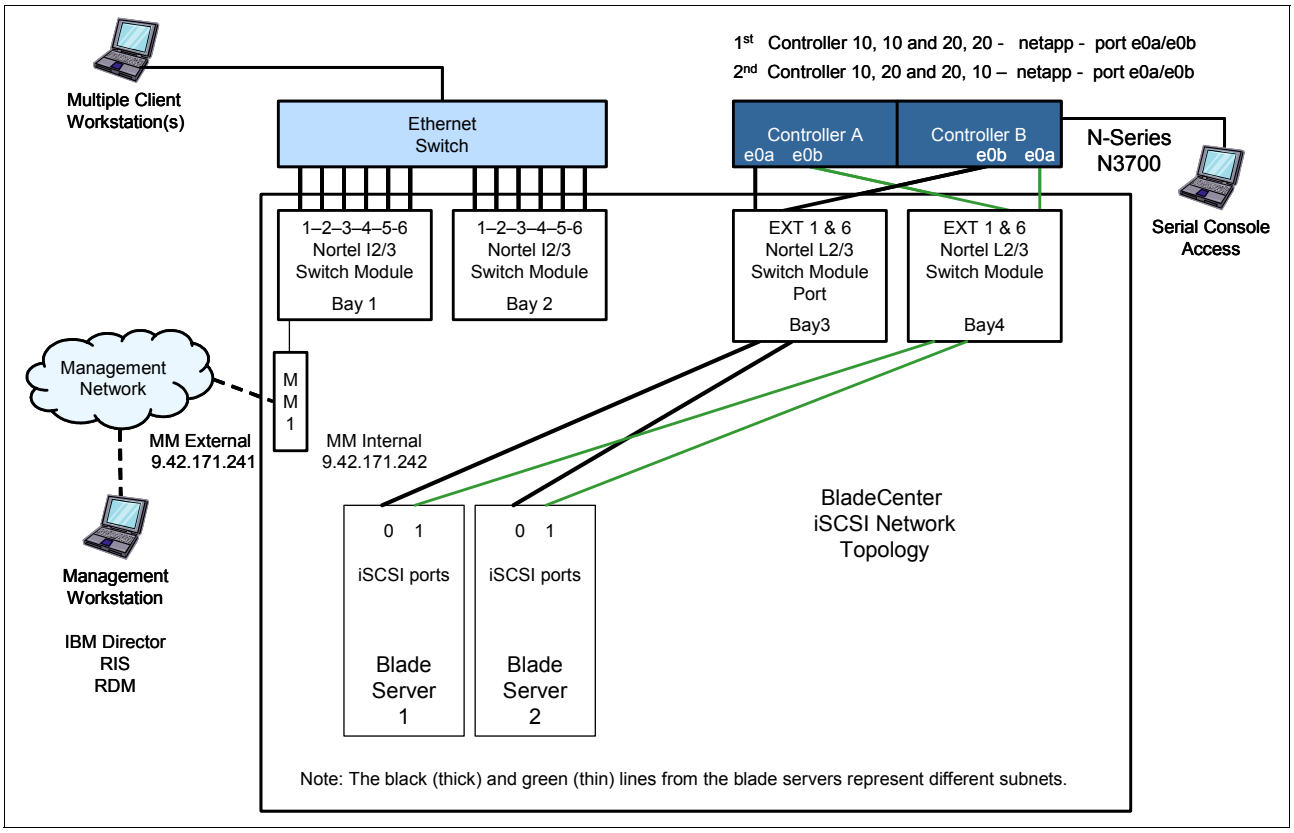

Figure 3-1 HS20 Blade Server 1 and 2 and N3700 Storage Subsystem using iSCSI hardware initiator

In Figure 3-1 the connections between the Ethernet switch modules and the N3700 controllers are shown as direct connections, but this is permissible and commonplace for intermediate Ethernet switches to be used. And therefore following this design, different VLANs would be required for the 10.10.10.x subnet and the 10.10.20.x subnet.

The iqn names and IP addresses for blade servers 1 and 2 are shown in Table 3-1.

Table 3-1 HBA iqn names and IP addresses

Server	HBA#1 address	name	HBA#2 address	name
Blade 1	10.10.10.14	iqn.2000-04.qlogic:bc1s1p0	10.10.20.24	iqn.2000-04.qlogic:bc1s1p1
Blade 2	10.10.10.32	iqn.2000-04.qlogic:port 0	10.10.20.32.	iqn.2000-04.qlogic:port1

The connections from the Ethernet switches to the e0b ports are shown in green because they are on the same subnet as the other green lines. The associated IP addresses for the storage controllers and Ethernet switches are listed in Table 3-2 and Table 3-3 on page 48.

Table 3-2 Storage controllers port, IP addresses, and node names

Storage array - controllers			
	e0a address	e0b address	node name
Controller A	10.10.10.10	10.10.20.10	netapp

Storage array - controllers			
Controller B	10.10.10.20	10.10.20.20	netapp2

Table 3-3 Ethernet switches IP addresses

Switches		
Bay 3	10.10.10.1	10.10.20.1
Bay 4	10.10.10.2	10.10.20.2

Figure 3-2 is an example of the iSCSI SAN topology based on our lab configuration using the software initiator.

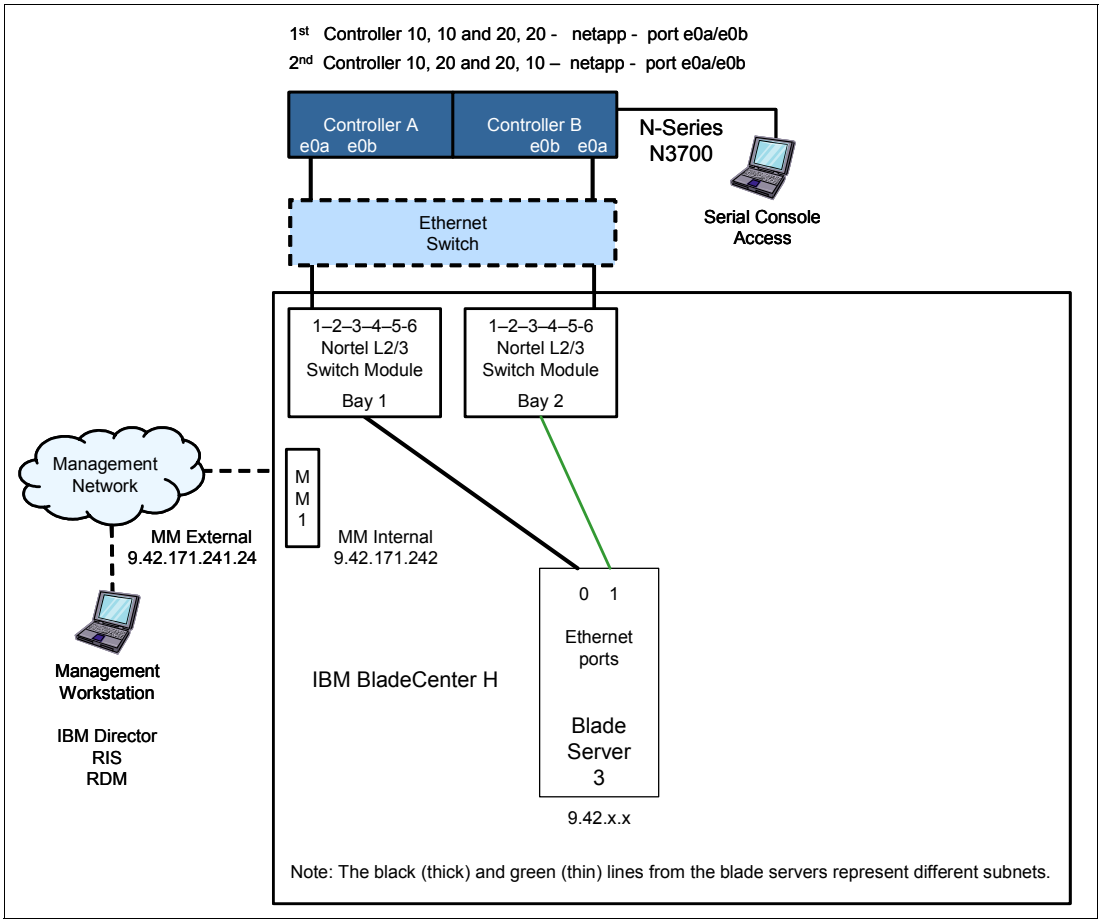

Figure 3-2 HS20 Blade Server 3 and N3700 Storage Subsystem using software initiator

Our lab configuration used for testing included the following:

- In a standard chassis, blade 3 (an HS20) used the Microsoft software initiator.
- Two Nortel L2/3 switches were used in the chassis (bays 1 and 2) for traffic generated by the software initiator.
- The iSCSI storage array used was the IBM N-Series N3700 (utilizes the same operating system as the NetApp FAS270). The storage array had two controllers, each of which has two Gigabit Ethernet ports. The names assigned were *netapp* and *netapp2*.

Shown in Figure 3-2 on page 48, the software initiator on Blade Server 3 communicates to the storage array via the Ethernet switch modules in bay 1 and 2, using the regular Broadcom NICs on the planar board in the blade. This means that NIC teaming can and probably should be used so that the server has one address that applies across the two NICs. We used NIC teaming in our testing and the server used an address in the 9.42.x.x subnet. (This address was assigned by DHCP in our testing but would probably have been statically assigned in a real configuration).

In Figure 3-2 on page 48, an intermediate switch is shown between the switch modules in bays 1 and 2 and the storage device. This is optional and not required. The use of intermediate switches allows servers on multiple BladeCenter chassis to access LUNs on the same storage array.

Table 3-4 HS20 Ethernet controller

Server	Ethernet ports	IP address
Blade 3	0, 1	9.42.x.x

The IP addresses for the storage controllers and Ethernet switches are listed in Table 3-5.

Table 3-5 Storage controllers port, IP addresses, and node names

Storage array - controllers			
	e0a address	e0b address	node name
Controller A	10.10.10.10	10.10.20.10	netapp
Controller B	10.10.10.20	10.10.20.20	netapp2

3.2 NetApp/IBM N3700 configurations

This section shows the steps to configure the N3700 for iSCSI access. However, storage configuration (defining, initializing, and formatting LUNs, and similar steps) are not covered here. It is assumed that these steps were performed previous to the steps in the following sections.

3.2.1 IP addressing

The **ifconfig** command (Figure 3-3) is used to configure each port of each of controller. Each port has its own IP address and no two ports on the same controller should be on the same subnet. Also, the partner parameter should be used to enable the controllers to take over for each other in the event of an outage.

```
ifconfig -a

e0a: flags=848043<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
 inet 10.10.20.10 netmask 0xffffff00 broadcast 10.10.20.255
 partner inet 10.10.20.20 (not in use)
 ether 00:a0:98:01:80:34 (auto-1000t-fd-up) flowcontrol full

e0b: flags=848043<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
 inet 10.10.10.20 netmask 0xffffff00 broadcast 10.10.10.255
 partner inet 10.10.10.10 (not in use)
 ether 00:a0:98:01:80:35 (auto-1000t-fd-up) flowcontrol full

lo: flags=1948049<UP,LOOPBACK,RUNNING,MULTICAST,TCPCKSUM> mtu 8160
 inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1
 ether b8:bc:cf:37:04:00 (VIA Provider)
```

Figure 3-3 ifconfig command window

3.2.2 Initiators and Initiator groups

The initiators (clients) must be specifically allowed to access each LUN that the server blade is going to use. This is done with the **igroup** command. Note that the QLogic hardware initiators have a different name (Figure 3-4) for each of the two ports on the HBA card and both of these initiator names should be authorized to access the LUN if you need a High Availability configuration. The Microsoft software initiator uses one nodename no matter how many physical NICs are available on the server. The initiator names used in our testing are as follows (names including iqn.2000-04.com.qlogic are hardware initiators; names including iqn.1999-05.com.microsoft are software initiators).

```
netapp2> iscsi initiator show

Initiators connected:

  TSIH  TPGroup  Initiator
   141     200   iqn.2000-04.com.qlogic:bc1s1p0 / 40:0f:21:1f:4f:00
   142     200   iqn.2000-04.com.qlogic:bc1s1p1 / 40:0f:21:1f:4f:00
   258     200   iqn.2000-04.com.qlogic:port0 / 40:0f:21:00:00:00
   270     200   iqn.1991-05.com.microsoft:bc3srv3 / 40:01:37:00:00:00
```

Figure 3-4 iscsi initiator show window

3.2.3 Storage array node name

The N-Series N3700 (and NetApp FAS270) array has a default nodename of the form iqn.2000-04.com.qlogic:sn.<serial number> (Figure 3-5). This can be changed but we did not change it in our testing. Each of the two controllers has a different serial number, so the storage array has two nodenames. The two controllers each have a configured *alias*. The ones configured in our testing were netapp and netapp2.

```
netapp2> iscsi nodename

iSCSI target nodename: iqn.1992-08.com.netapp:sn.84178161
```

Figure 3-5 iscsi nodename window

3.3 QLogic iSCSI HBA and software configuration

In this section we briefly describe the install of the iSCSI HBA and then the install of the SANsurfer software used to manage the HBA.

3.3.1 Hardware and software installation

The following steps are needed to install the SANsurfer software on a Windows blade server.

First, install the iSCSI HBA card in the bracket towards the rear of the blade. Note that if there are any server blades with non-Ethernet HBA cards installed that are in the chassis and powered on, the switch modules in bays 3 and 4 will not power up. This includes Fiber Channel, Myrinet, Infiniband, and any others that may be developed.

Second, install SANsurfer from the installation CD and then reboot the server. Additional steps are required to install BIOS support for iSCSI Boot from SAN. Those steps are detailed

in Chapter 2, “Boot from iSCSI SAN using iSCSI HBA and initiator with failover support” on page 9.

The iSCSI GUI and Agent should both be installed, as shown in Figure 3-6.

Figure 3-6 SANsurfer 4.0 (iSCSI Standalone) window

The SANsurfer software should be installed for all users unless you are signed on with an administrator or equivalent ID and wish to restrict the software to be used only on that ID (Figure 3-7).

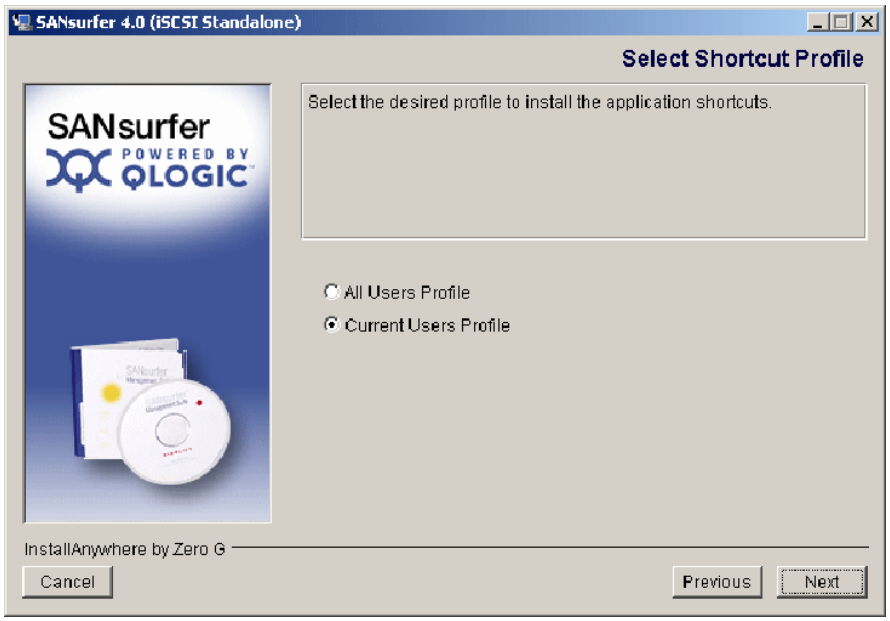

Figure 3-7 SANsurfer - Select Shortcut Profile window

After the software installation is complete the HBA installed on the blade should be visible (see Figure 3-8).

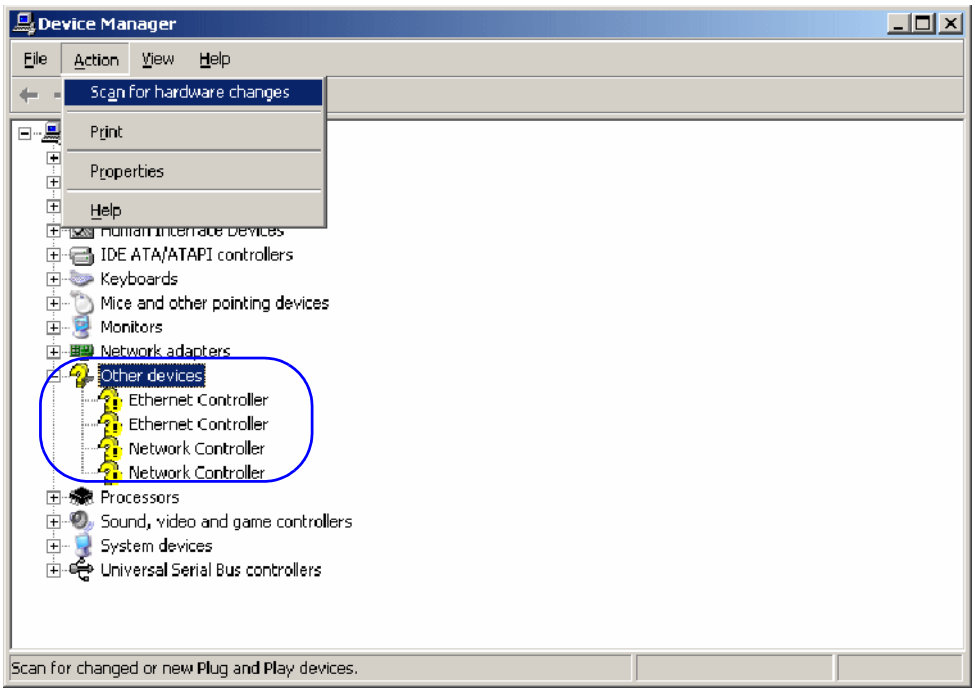

Figure 3-8 Scan for hardware changes window

The scan for hardware changes will prompt you for the location of the drivers for the newly discovered devices. You should specify the location, as shown in Figure 3-9 and Figure 3-10 on page 54.

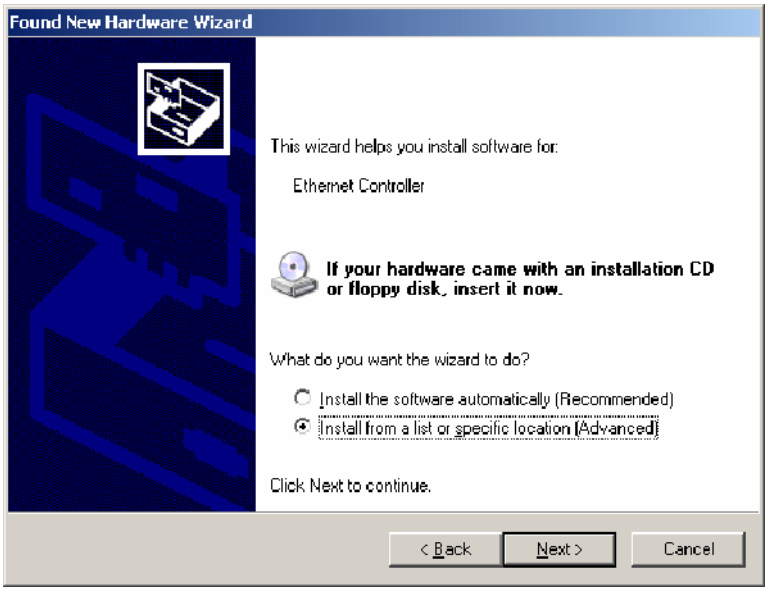

Figure 3-9 Found New Hardware Wizard window

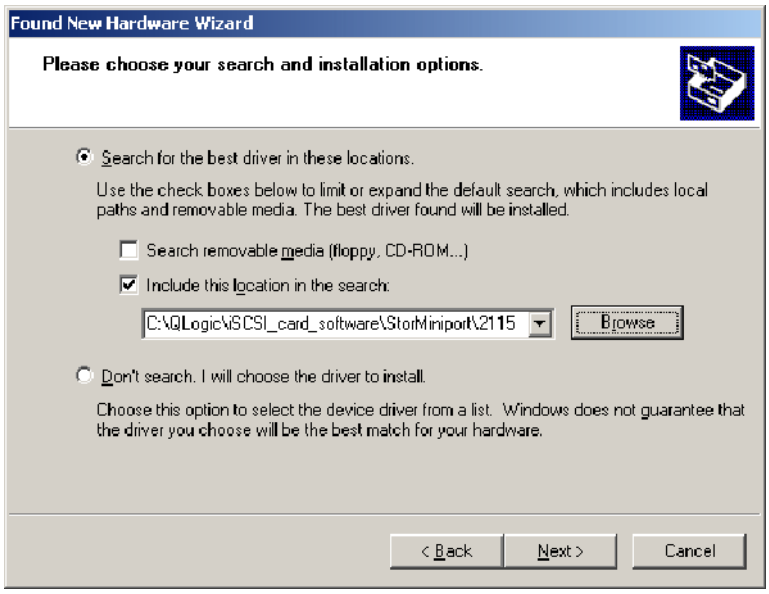

Figure 3-10 Found New Hardware Wizard window

After the drivers have successfully been installed, the controllers will be visible in the Device Manager of the server (Figure 3-11).

Figure 3-11 Device Manager window

3.3.2 SANsurfer configuration

The next step is to configure the iSCSI card and associated software using the SANsurfer GUI. The first step is to connect to the server being configured. In our testing, we ran the GUI on the server so we could connect to localhost. If you run the GUI on a management station or elsewhere, specify the host name or IP address of the server to be configured (see Figure 3-12).

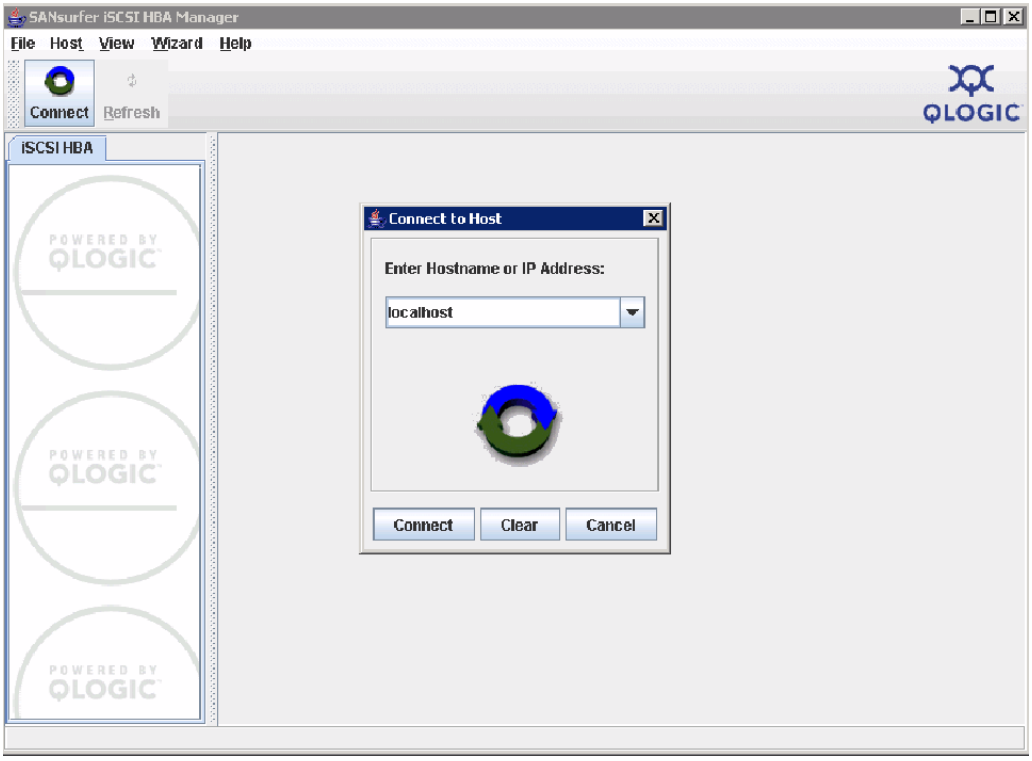

Figure 3-12 SANsurfer iSCSI HBA Manager window

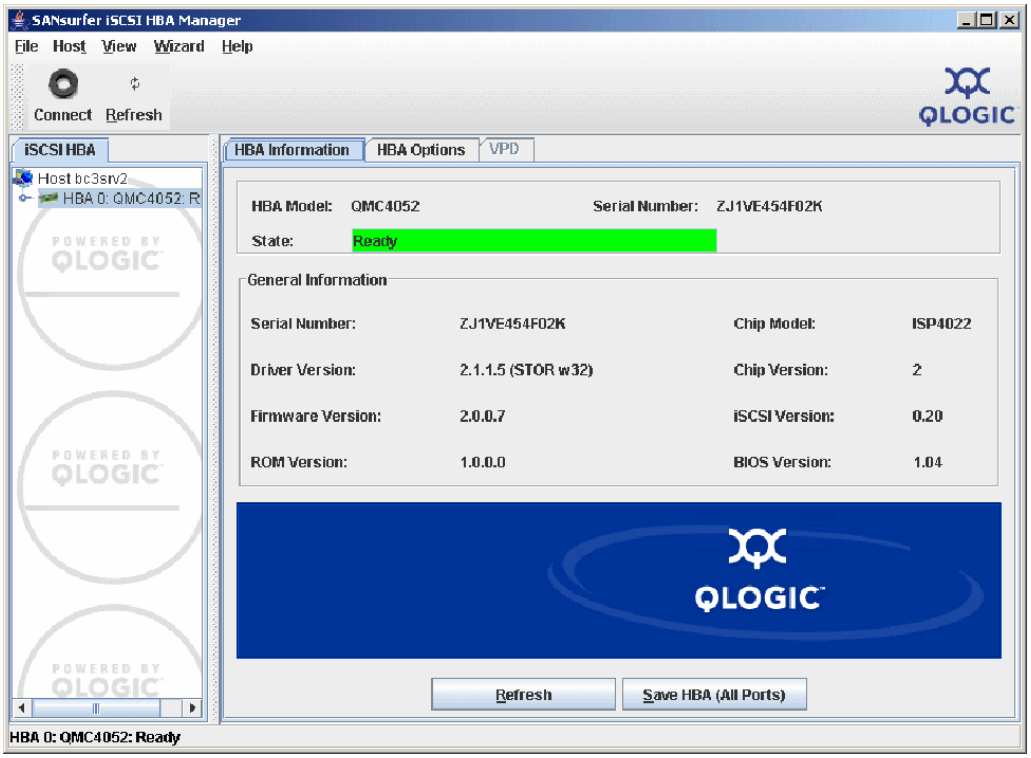

Figure 3-13 SANsurfer iSCSI HBA Manager window

After connecting to the server, you have the ability to update the firmware, ROM, and BIOS on the daughter card itself (see Figure 3-13).

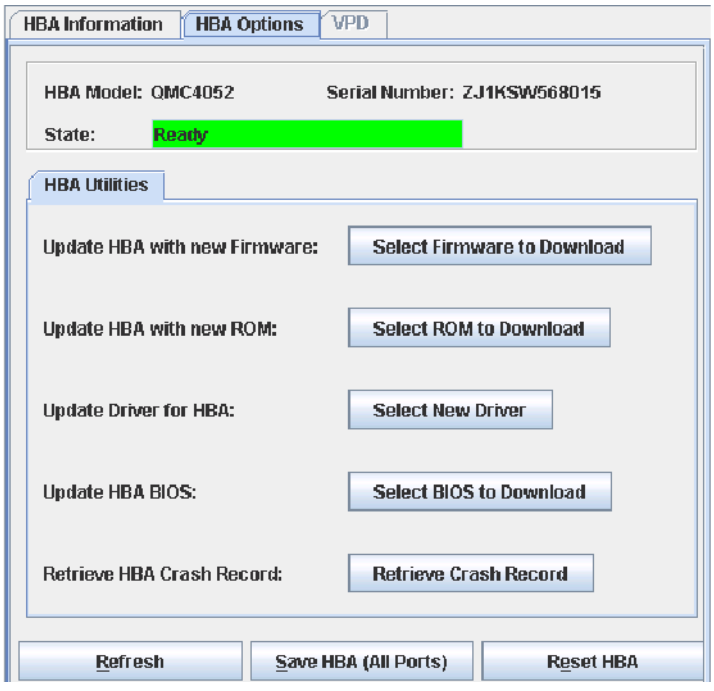

Figure 3-14 HBA Options window

3.3.3 Network configuration

SANsurfer presents the HBA card on the blade servers as a single HBA with two ports. Each port has a distinct initiator nodename. Each port also has its own IP address. The two ports were configured in our testing with addresses on different subnets because the ports on the controllers on the N-Series N3700 array are strongly recommended to be in different subnets. Addresses are assigned through the Port Options page in SANsurfer, which is shown below. This step must be performed once for each port.

In Figure 3-15 we assign the IP address of 10.10.10.14 and the initiator name (iSCSI Name). In a production environment with a more complex network topology, a default gateway would be required as well.

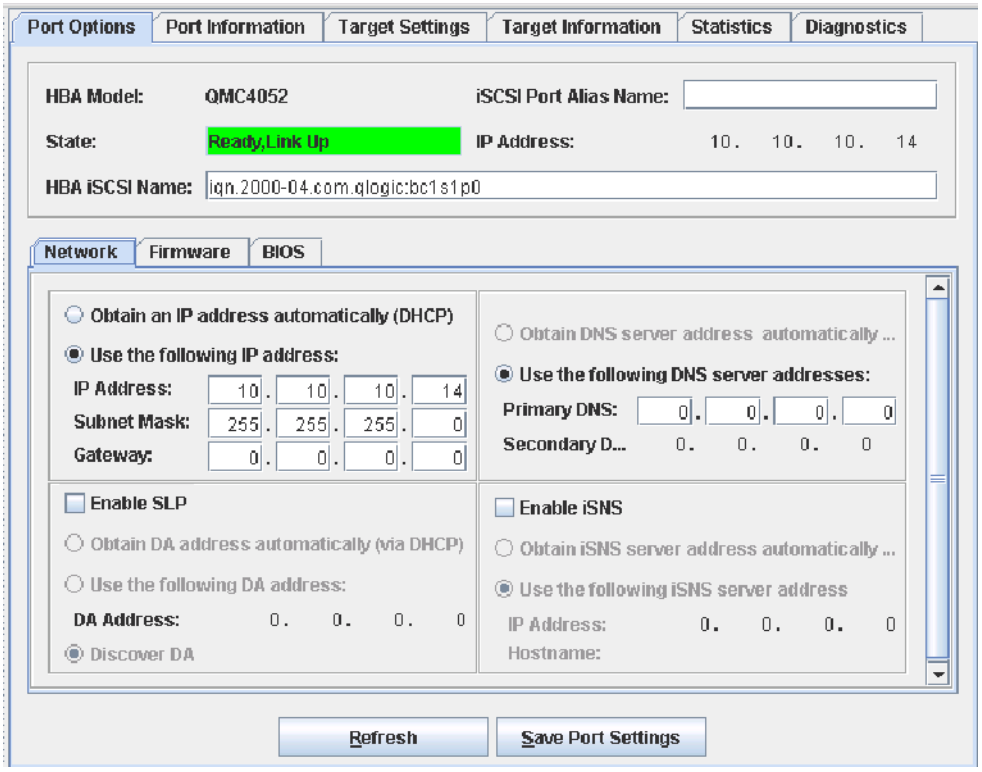

Figure 3-15 Port Options tab window

Note that this card cannot be seen by or used by the TCP/IP stack included in the server's operating system. Running the **ipconfig** command on a Windows server would not show these ports or their addresses (see Figure 3-16).

```
Microsoft Windows [ Version 5.2.3790]
<C> Copyright 1985-2003 Microsoft Corp.
C:\Documents and Settings\Administrator\ipconfig
Windows IP Configuration

Ethernet adapter ITSO 2:
   Connection-specific DNS Suffix .: itso.ral.ibm.com
   IP Address. . . . . . . . . . . . : 9.42.171.164
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway ..... . . . . . : 9.42.171.3

Ethernet adapter ITSO 1:
   Connection-specific DNS Suffix .: itso.ral.ibm.com
   IP Address. . . . . . . . . . . . : 9.42.171.201
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway ..... . . . . . : 9.42.171.3

C:\Documents and Settings\Administrator>_
```

Figure 3-16 Command-line interface window

3.3.4 Specification of targets

The initiators on the server must be configured with the targets that represent the storage devices to which it is to attach itself. Targets can be found by explicitly entering their fully qualified names (which begin with iqn) or by entering their IP addresses. If IP addresses are entered, a discovery process takes place. The initiators will contact the configured IP address and learn the associated name as well as learning other IP addresses through which this resource can be reached.

Note: The discovery process described above makes it difficult to use the virtual addressing functionality of the Nortel L2-7 switch with iSCSI. The initiator will start the discovery process with the virtual address, if so configured, and will then learn all of the real addresses associated with the resources at that address and their fully qualified names. Actual iSCSI connections will then be opened with the resource via one or more of the real addresses, negating most of the value of load balancing.

In Figure 3-17 we filled in those entries that contain only an IP address. The other entries that also contain iSCSI names (Figure 3-18 on page 60) were discovered by the software after it contacted the storage array at the addresses we specified.

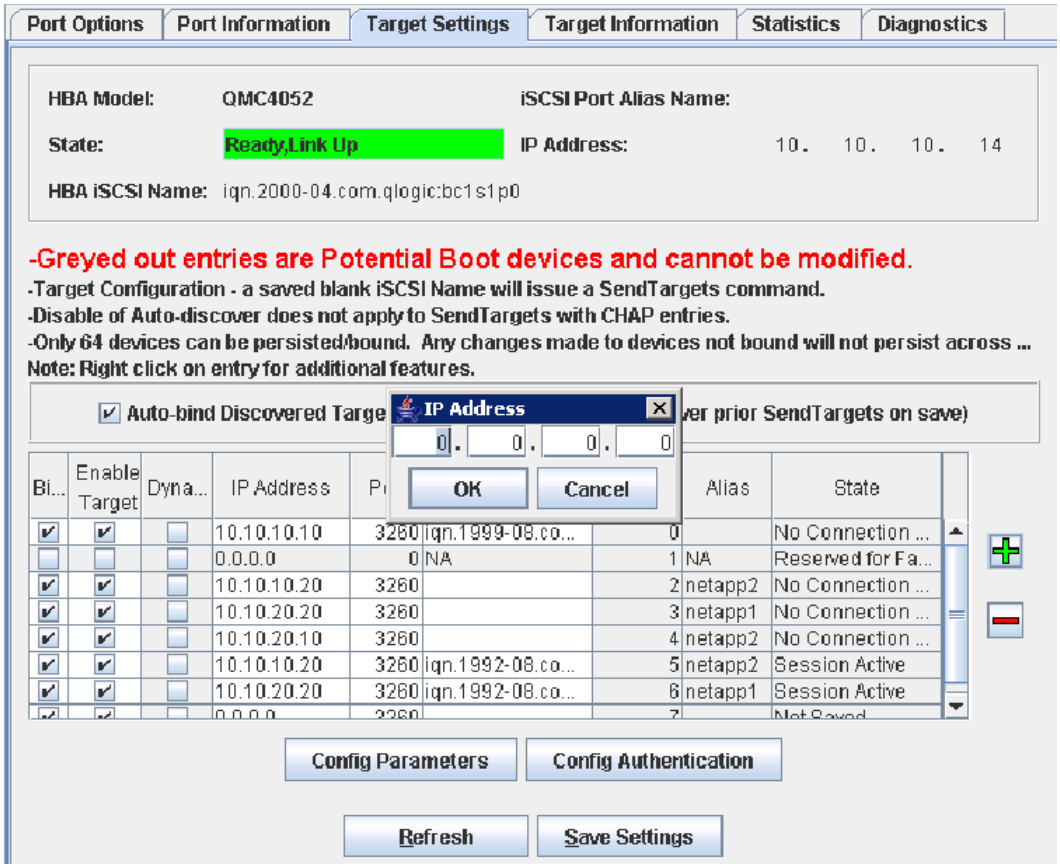

Figure 3-17 Target Setting tab window

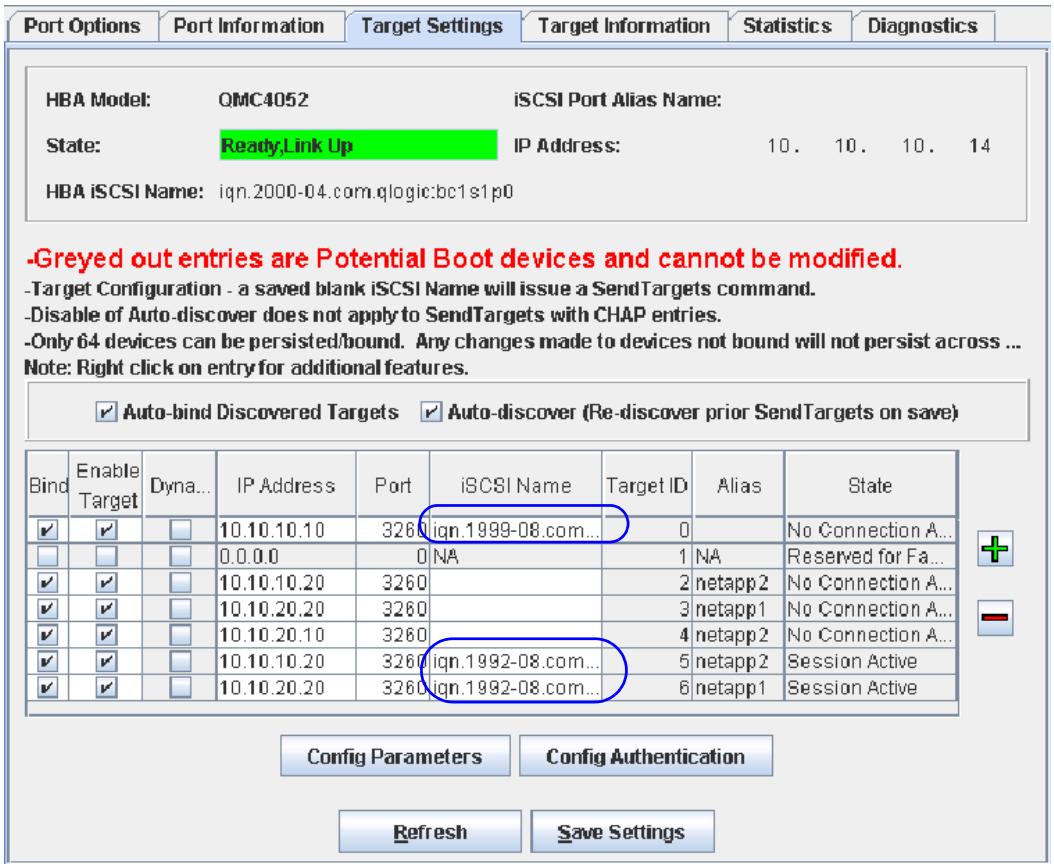

Figure 3-18 Target Setting tab window

It is possible to verify connectivity to the targets' IP addresses with a PING, as shown in Figure 3-19. The PING command issued from a command prompt will not verify connectivity between the HBA and the storage devices.

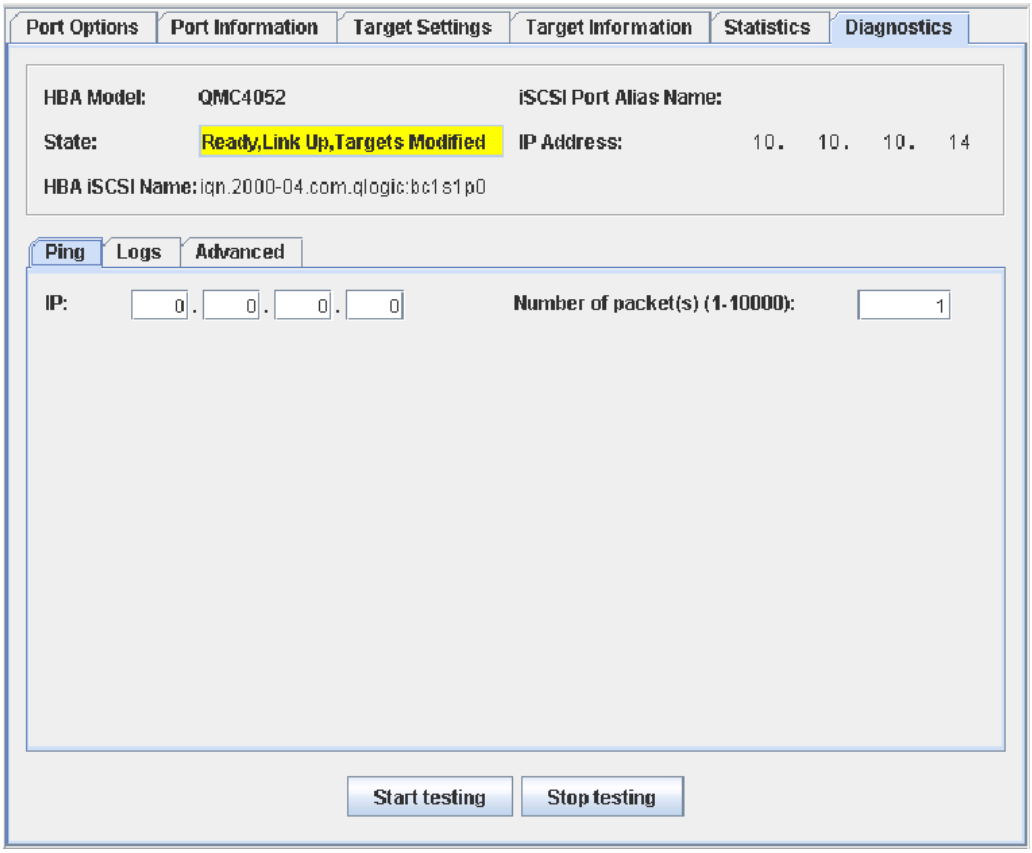

Figure 3-19 Diagnostics tab window

Once the HBA successfully connects to a LUN on the storage array, you can expand the target's icon in the left-hand portion of the SANsurfer window and obtain information about the LUN and the status of the connection to it (see Figure 3-20).

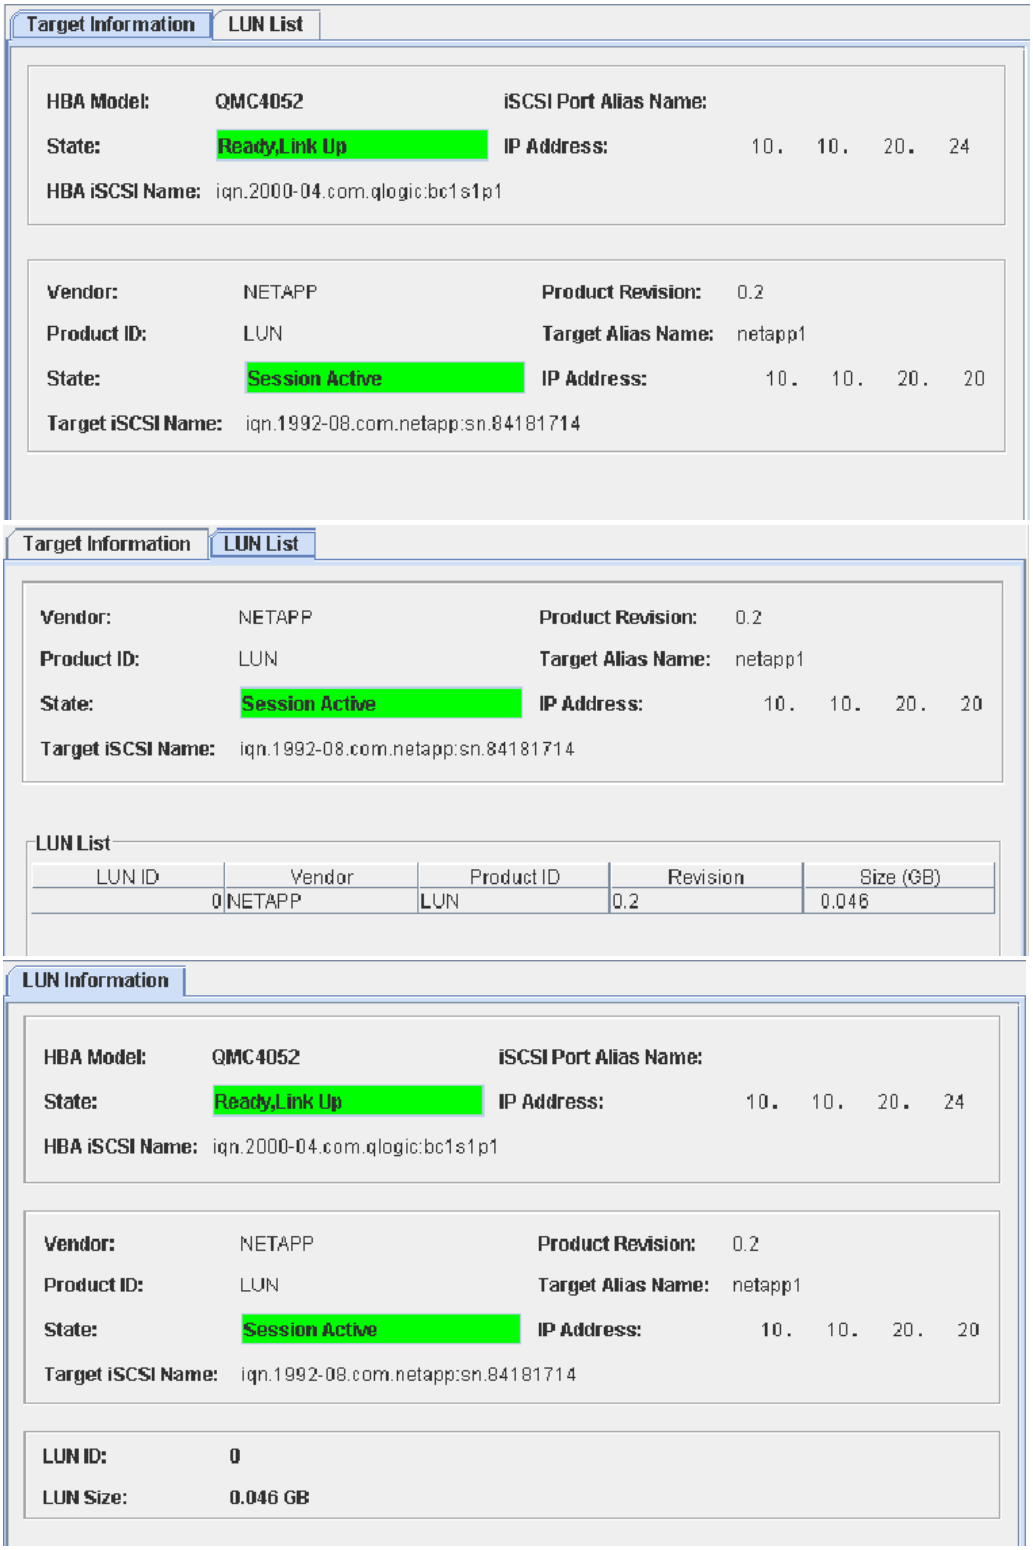

Figure 3-20 Target Information tab, LUN List tab, LUN Information windows

With the LUN connected you can bring the disk online from the disk manager. The first time this is done, you will need to first assign a drive letter to the disk. The iSCSI disk is disk 1.

Click **Start** → **Control Panel** → **Administrative Tools** → **Computer Management** → **Disk Management** to access the disk manager. Right-click **Disk 1** to assign a drive letter and to bring the drive online.

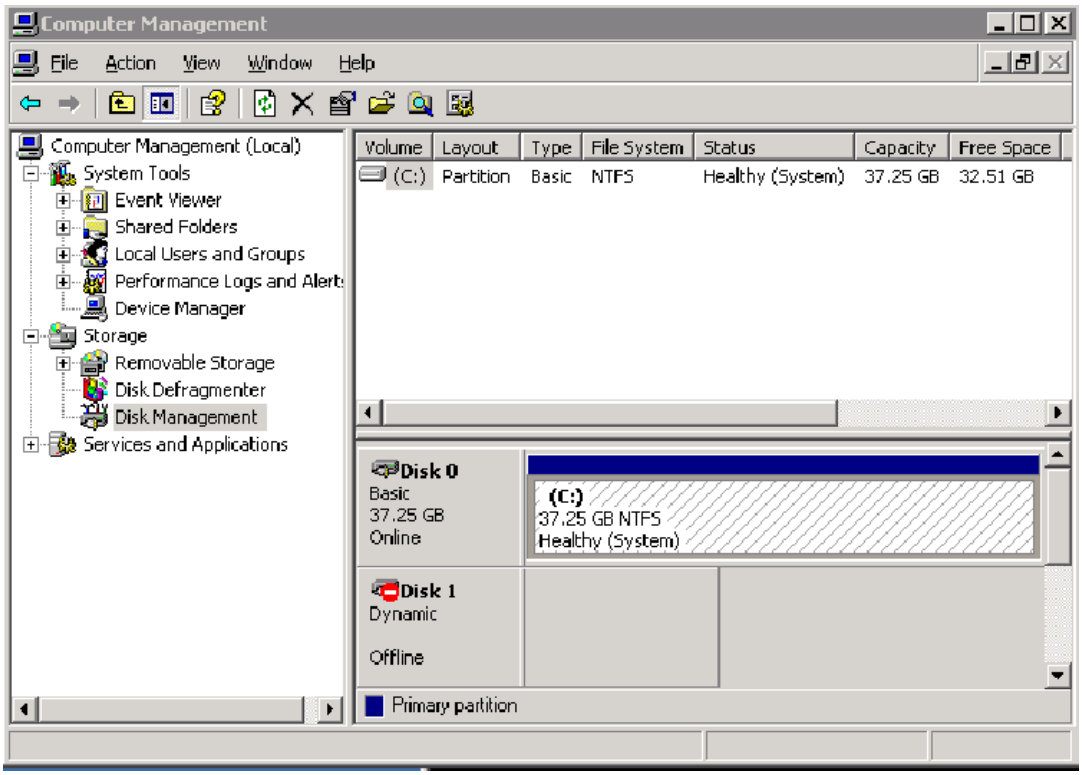

Figure 3-21 Computer Management window

Once the disk is online, it will appear under Disk Management, as in Figure 3-22. The *at risk* and *error* indications appear because the disk was not reachable to be brought online when the server first booted up.

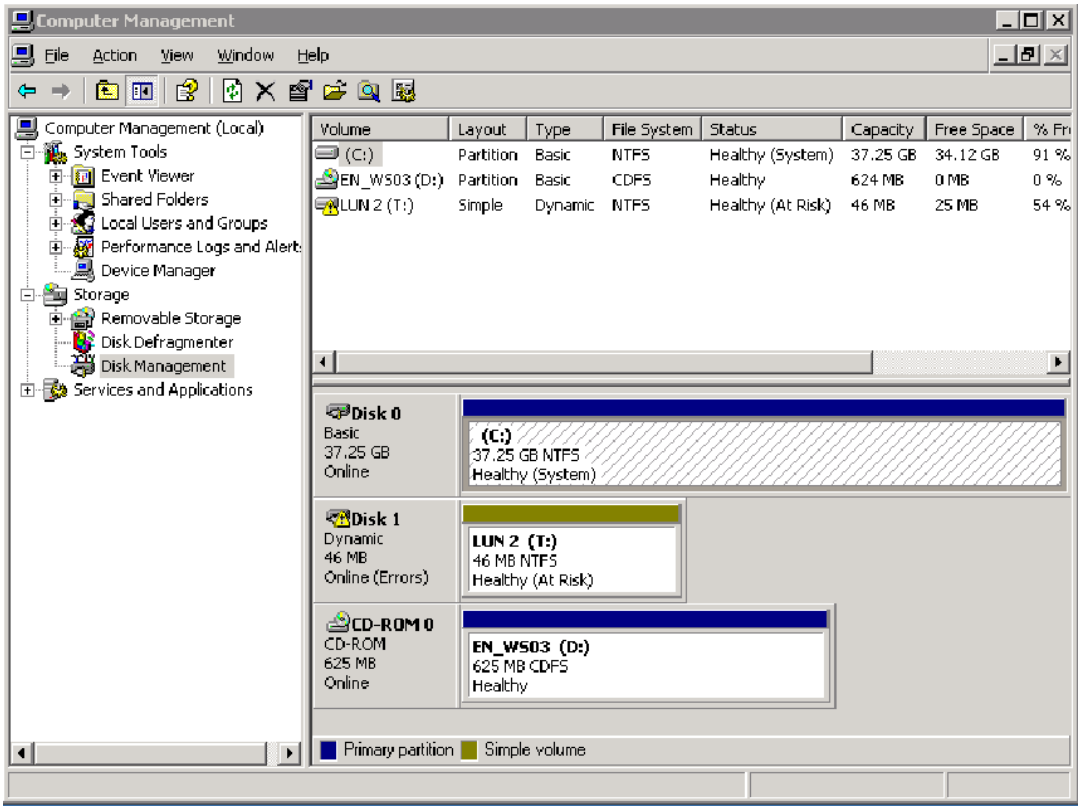

Figure 3-22 Computer Management window

3.4 Nortel switch configurations

The Nortel switches in this test were configured to allow either switch module access to any of the four ports on the Network Appliance array. This was done by configuring the switches with interface addresses on each of the subnets used for storage resources. Better containment of broadcasts could be achieved by assigning each subnet to its own VLAN, but in that case the ports would not be all interchangeable, and would not all be reachable from any port on any switch.

The actual configuration tested is shown in Example 3-1 and an alternate configuration with the VLAN assignment is shown in Example 3-2 on page 65. The alternate configuration relies entirely on the High Availability functions included in the iSCSI standards and does not use the switches' ability to find an available path to the destination. Further, in the alternate configuration the two Nortel switches are not cross-connected to each other (as is done via port EXT4 in the tested configuration).

Example 3-1 Switch 1 (iSCSI data on VLAN 1)

```
# Switch 1: iSCSI data on VLANs 1, subnets 10.10.10.x and 10.10.20.x, cross connected to
switch 2
# Switch 2 identical except as shown
```

```
script start "Layer 2-3 Gigabit Ethernet Switch Module for IBM eServer BladeCenter" 5
/**** DO NOT EDIT THIS LINE!
/* Configuration dump taken 20:59:24 Sat Jan  3, 2070
/* Version 1.2.2,  Base MAC address 00:16:60:f9:e1:00
/c/port EXT5
   pvid 2
/c/l2/vlan 1
   def INT1 INT2 INT3 INT4 INT5 INT6 INT7 INT8 INT9 INT10 INT11 INT12 INT13 INT14 EXT1 EXT2
EXT3 EXT4 EXT5 EXT6
/c/l3/if 1
   ena
   addr 10.10.10.1 /* 10.10.10.2 on switch 2 */
   mask 255.255.255.0
   broad 10.10.10.255
/c/l3/if 2
   ena
   addr 10.10.20.1  /* 10.10.20.2 on switch 2 */
   mask 255.255.255.0
   broad 10.10.20.255
/c/l3/if 9
   ena
   addr 9.42.171.50  /* management interface via port EXT5*/
   mask 255.255.255.0
   broad 9.42.171.255
   vlan 2
/c/l3/gw 1
   ena
   addr 9.42.171.3
/cfg/acl/acl 10/tcpudp/sport 3260 0xffff
/cfg/acl/acl 10/action setcos 5
/cfg/acl/acl 10/stats ena
/cfg/acl/acl 20/tcpudp/dport 3260 0xffff
/cfg/acl/acl 20/action setcos 5
/cfg/acl/acl 20/stats ena
/cfg/port INT1/acl/add acl 10
/cfg/port INT1/acl/add acl 20
/cfg/port INT2/acl/add acl 10
/cfg/port INT2/acl/add acl 20
/cfg/port EXT1/acl/add acl 10
/cfg/port EXT1/acl/add acl 20
/cfg/port EXT4/acl/add acl 10
/cfg/port EXT4/acl/add acl 20
/cfg/port EXT6/acl/add acl 10
/cfg/port EXT6/acl/add acl 20
/
script end  /**** DO NOT EDIT THIS LINE!

>> Configuration#
```

Example 3-2 Switch 1 (iSCSI data on VLAN 10)

```
# Switch 1 - iSCSI data on VLAN 10, subnet 10.10.10.x, ports 1 and 6 connected to different
controllers on storage

script start "Layer 2-3 Gigabit Ethernet Switch Module for IBM eServer BladeCenter" 5
/**** DO NOT EDIT THIS LINE!
/* Configuration dump taken 20:59:24 Sat Jan  3, 2070
/* Version 1.2.2,  Base MAC address 00:16:60:f9:e1:00
/c/port EXT5
   pvid 2
```

```
/c/l2/vlan 1
   def INT3 INT4 INT5 INT6 INT7 INT8 INT9 INT10 INT11 INT12 INT13 INT14 EXT2 EXT3 EXT5
/c/port INT1
   pvid 10
/c/port INT2
   pvid 10
/c/port EXT1
   pvid 10
/c/port ext6
   pvid 10
/c/l2/vlan 10
   def INT1 INT2 EXT1 EXT6
/c/l3/if 1
   ena
   addr 10.10.10.2
   mask 255.255.255.0
   broad 10.10.10.255
   vlan 10
/c/l3/if 9       /* management interface via EXT5 */
   ena
   addr 9.42.171.57
   mask 255.255.255.0
   broad 9.42.171.255
   vlan 2
/c/l3/gw 1
   ena
   addr 9.42.171.3
/cfg/acl/acl 10/tcpudp/sport 3260 0xffff
/cfg/acl/acl 10/action setcos 5
/cfg/acl/acl 10/stats ena
/cfg/acl/acl 20/tcpudp/dport 3260 0xffff
/cfg/acl/acl 20/action setcos 5
/cfg/acl/acl 20/stats ena
/cfg/port INT1/acl/add acl 10
/cfg/port INT1/acl/add acl 20
/cfg/port INT2/acl/add acl 10
/cfg/port INT2/acl/add acl 20
/cfg/port EXT1/acl/add acl 10
/cfg/port EXT6/acl/add acl 10
/cfg/port EXT6/acl/add acl 20
/
script end  /**** DO NOT EDIT THIS LINE!

# Switch 2 - iSCSI data on VLAN 20, subnet 10.10.20.x, ports 1 and 6 connected to different
controllers on storage

script start "Layer 2-3 Gigabit Ethernet Switch Module for IBM eServer BladeCenter" 5
/**** DO NOT EDIT THIS LINE!
/* Configuration dump taken 20:59:24 Sat Jan  3, 2070
/* Version 1.2.2,  Base MAC address 00:16:60:f9:e1:00
/c/port EXT5
   pvid 2
/c/l2/vlan 1
   def INT3 INT4 INT5 INT6 INT7 INT8 INT9 INT10 INT11 INT12 INT13 INT14 EXT2 EXT3 EXT5
/c/port INT1
   pvid 20
/c/port INT2
   pvid 20
/c/port EXT1
   pvid 20
```

```
/c/port ext6
   pvid 20
/c/l2/vlan 20
   def INT1 INT2 EXT1 EXT6
/c/l3/if 1
   ena
   addr 10.10.20.2
   mask 255.255.255.0
   broad 10.10.10.255
   vlan 20
/c/l3/if 9       /* management interface via EXT5 */
   ena
   addr 9.42.171.50
   mask 255.255.255.0
   broad 9.42.171.255
   vlan 2
/c/l3/gw 1
   ena
   addr 9.42.171.3
/cfg/acl/acl 10/tcpudp/sport 3260 0xffff
/cfg/acl/acl 10/action setcos 5
/cfg/acl/acl 10/stats ena
/cfg/acl/acl 20/tcpudp/dport 3260 0xffff
/cfg/acl/acl 20/action setcos 5
/cfg/acl/acl 20/stats ena
/cfg/port INT1/acl/add acl 10
/cfg/port INT1/acl/add acl 20
/cfg/port INT2/acl/add acl 10
/cfg/port INT2/acl/add acl 20
/cfg/port EXT1/acl/add acl 10
/cfg/port EXT6/acl/add acl 10
/cfg/port EXT6/acl/add acl 20
/
script end  /**** DO NOT EDIT THIS LINE!

>> Configuration#
```

The switch configuration also includes Quality of Service (QoS) commands to give iSCSI traffic higher priority than other traffic. The QoS commands are included in both Example 3-1 on page 64 and Example 3-2 on page 65. They use the DSCP functionality of the switch to set the priority field in the IP header of all traffic to and from port 3260, which is the default port for iSCSI. This form of prioritization gives a higher weight to iSCSI packets, and packets are sent out from the switch using a weighted round-robin algorithm.

The example shown is not the only way to prioritize iSCSI traffic. It is possible to give iSCSI unconditional priority over other traffic, meaning that if there is an iSCSI packet waiting it will be sent before any other packets are considered. This can be done by setting the DSCP for iSCSI to 6. Alternatively, it is possible to assign iSCSI traffic a weight less than that associated with DSCP 5.

Note: If iSCSI traffic is segregated from normal data traffic, then the use of QoS adds little value. If iSCSI traffic and normal data traffic will use the same switches, connections, and so on, then QoS may be useful. However, QoS is quite complex and a thorough analysis of all of the traffic flowing to and from the BladeCenter is strongly recommended.

3.5 High availability options and configurations

In this section we discuss some of the options and configurations for HA.

3.5.1 HA from storage

The iSCSI storage array we tested has HA capabilities of its own. They are described below.

The N-Series N3700 used for testing was configured with two network controllers, each of which has two Gigabit Ethernet ports. This provides High Availability in the following ways:

- The controllers can back each other up using cluster failover. Cluster failover is a software feature of the N-Series N3700 that is specifically licensed. It allows one controller to take over the LUNs, initiators, and other definitions of a partner including its IP addresses. Cluster failover is enabled with the **`cf enable`** command. The ports can specify their partner using the partner operand of the **`ipconfig`** command.
- The two ports within a cluster can back each other up. Each port will have its own IP address, but the node discovery process will enable an initiator (client) to learn all of the addresses that can be used to reach a particular target portal. In the event that one such address becomes unreachable the initiator can re-establish the connection using one of the other valid addresses. In the testing we performed, the two ports on each controller were attached to two different switches in the BladeCenter chassis.

3.5.2 HA from switches

The Nortel switch modules used in our testing added to the HA capabilities of the N-Series N3700 in the following ways: Layer 3 configuration:

- The N-Series N3700 issues a diagnostic message when the two ports on a controller are configured with IP addresses on the same subnet. In our testing, we configured the ports on two different subnets (10.10.10.x and 10.10.20.x).
- The QLogic initiator is configured with two ports, one of which is attached to each controller. One of the ports is configured with a 10.10.10.x address and the other with a 10.10.20.x.
- Because the Nortel switch is configured for routing, either port on a particular server can reach either port on either controller on the storage array. This allows the environment to survive the following types of outages:
  - Failure or removal of one of the network controllers on the storage array
  - Failure or removal of one of the switch modules
  - Removal of a cable causing a loss of connectivity between the switch modules and the storage array.

Notes: Neither the QLogic hardware initiators with SANsurfer nor the N-Series N3700 support port aggregation (also known as Multi Link Trunking, 802.3ad, or Etherchannel).

The QLogic hardware does not support NIC teaming such as that provided by the BASP driver on the on-board NIC chips on the HS20. The two ports on the server (one for bay 3 and one for bay 4) will each have their own IP address.

A

Install local hard disk for high-load environments

We strongly recommend that you install the local hard disk for high-load environments. The operating system or applications may experience errors and become unstable due to latency with accessing the page file. Thus, having pagefile on a local disk ensures reliable access to the pagefile. Specifically, we elect to use the local disk for swap/RAS.

Refer to the following Microsoft KB (Knowledge base) articles on Boot from SAN issues:

- Microsoft SAN Support

  http://www.microsoft.com/windowsserversystem/storage/sansupport.mspx#top
- Support for Booting from a SAN

  http://support.microsoft.com/default.aspx?scid=kb;en-us;305547
- Support for Multiple Clusters Attached to the Same SAN Device

  http://support.microsoft.com/default.aspx?scid=kb;en-us;304415

B

N-Series N3700 and NetApp FAS270 CPU Module

Figure 3-23 has been inserted here as an illustration to allow you a better view of the ports to be configured on the N-Series N3700 in 2.3.6, “Multipath topology configuration” on page 31.

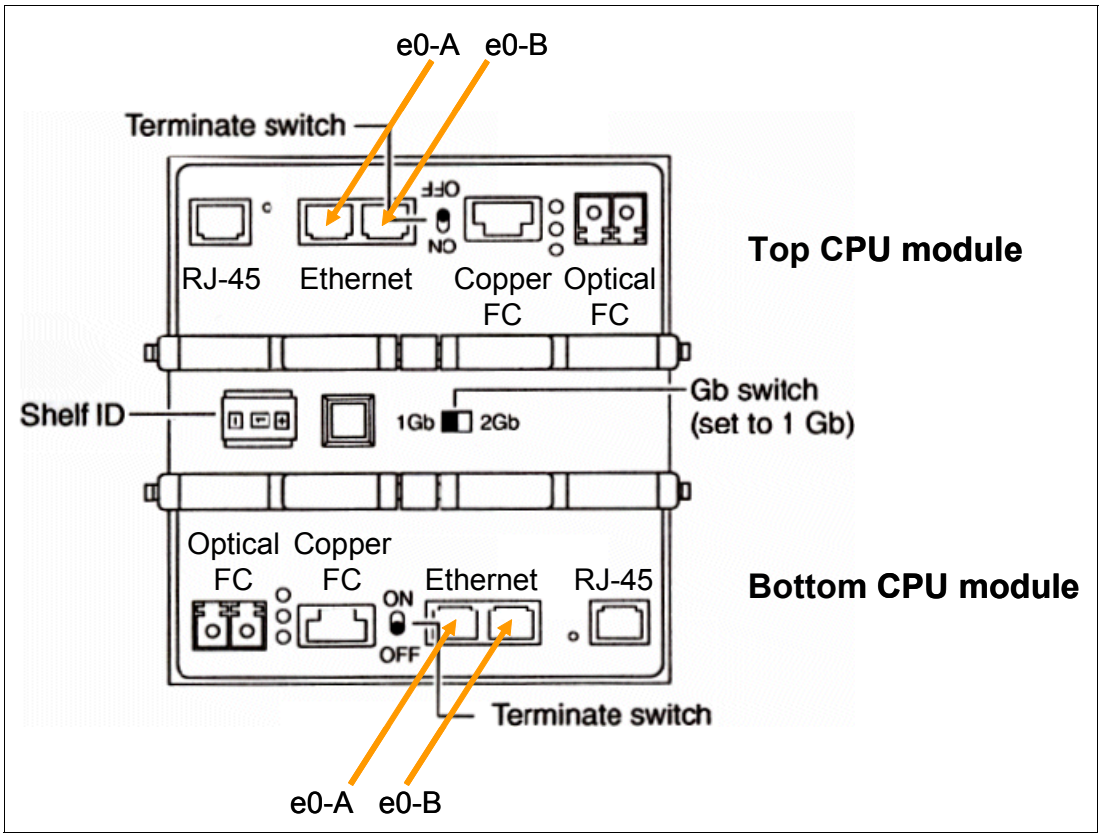

Figure 3-23 CPU module (controller) of N3700 dual-controller storage system

C

Corresponding switch module configurations

The configurations listed in this appendix correspond to the Cisco switch configuration in 2.3.1, “Network configuration” on page 15.

Example C-1 corresponds to Example 2-1 on page 15.

Example: C-1 Network configuration 1

```
! Switches in Bay 3 and 4 use the configuration below
! Set blade 8 to use VLAN 111
/c/l2/vlan 111/ena
/c/port INT8/pvid 111
! Note that the below is optional and will disable Serial-over-LAN
/c/port INT8/tag d
! Note that spanning tree is disabled by default on internal ports
```

Example C-2 corresponds to Example 2-2 on page 16.

Example: C-2 Network configuration 2

```
! Switches in bay 3 and 4 can use identical configuration here as well if both are
connected to an upstream switch
/c/port EXT1/tag e
/c/port EXT1/pvid 100
/c/port EXT1/upfast enable
/c/l2/vlan 100/ena
/c/l2/vlan 100/add EXT1
/c/l2/vlan 111/add EXT1
! note that the below is optional and only needed if two uplinks ports (EXT1 and another)
are to be bound together with LACP
! the other EXT port would also be configured as shown in this example
! The N3700 does not support LACP so an intermediate switch would be required.
/c/l2/lacp/port EXT1/mode active
/c/l2/lacp/port EXT1/adminkey 1
```

The following note corresponds to Example 2-3 on page 16.

Note: If a Cisco 6500 is used as an upstream switch located between the BladeCenter and the storage device, the config in Example 2-3 on page 16 can be used unchanged regardless of which switches are in the BladeCenter chassis. If the storage is connected directly to the BladeCenter switches, then configurations like those in 3.4, “Nortel switch configurations” on page 64, can be used.

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 76. Note that some of the documents referenced here may be available in softcopy only.

- *Using iSCSI Solutions' Planning and Implementation*, SG24-6291
- *IP Storage Networking: IBM NAS and iSCSI Solutions*, SG24-6240
- *IBM TotalStorage DS300 and DS400 Best Practices Guide*, SG24-7121

Online resources

These Web sites and URLs are also relevant as further information sources:

- QLogic iSCSI Expansion Card for IBM BladeCenter

  http://www-306.ibm.com/common/ssi/rep_ca/4/897/ENUS105-194/ENUS105-194.PDF
- BladeCenter Ethernet Switch Modules

  http://www-03.ibm.com/systems/bladecenter/switch/switch_ethernet_overview.html
- BladeCenter switch module firmware

  http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54597
- Blade Network Technologies, Inc.

  http://www.bladenetwork.net/
- N-Series N3700

  http://www-03.ibm.com/servers/storage/network/n3700/
- NetApp products

  http://www.netapp.com/products/filer/fas200_ds.html
- QLogic iSCSI Expansion Card device driver

  http://support.qlogic.com/support/oem_detail_all.asp?oemid=369
- 32 bit STOR Miniport Storage Only

  http://support.qlogic.com/support/Eula_1.asp?path=http://download.qlogic.com/drivers/29651/StorMiniport.zip
- Microsoft Multipath I/O

  http://www.microsoft.com/WindowsServer2003/technologies/storage/mpio/faq.mspx
- Microsoft iSCSI Software Initiator Version 2.01

  http://www.microsoft.com/downloads/details.aspx?FamilyID=12cb3c1a-15d6-4585-b385-befd1319f825&DisplayLang=en
- Fibre Channel security

  http://www.t11.org/ftp/t11/pub/fc/sp/06-157v0.pdf

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

IBM BladeCenter iSCSI SAN Solution

Redpaper

Include failover support utilizing the IBM N3700 and NetApp FAS270

QLogic iSCSI Expansion Card using the Microsoft iSCSI software initiator

High Availability using Nortel L2/3 and L2-7 Switch Modules

This IBM Redpaper provides instructions to configure an IBM BladeCenter Boot from iSCSI SAN from the IBM N-Series N3700 storage system. Our Boot from iSCSI SAN test environment includes the use of the Nortel and Cisco Switch Modules for the IBM BladeCenter. We also discuss High Availability utilizing the Nortel Networks L2-3 and L2-7 Switch Modules for the IBM BladeCenter.

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