

Hitachi Modular Host Installation Guide for Fibre Channel

FASTFIND LINKS

[Document Organization](#)

[Getting Help](#)

[Table of Contents](#)

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd. and Hitachi Data Systems Corporation (hereinafter referred to as "Hitachi").

Hitachi, Ltd. and Hitachi Data Systems reserve the right to make changes to this document at any time without notice and assume no responsibility for its use. Hitachi, Ltd. and Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements.

All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

Notice: Hitachi Data Systems products and services can be ordered only under the terms and conditions of Hitachi Data Systems' applicable agreement(s). The use of Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Hitachi is a registered trademark of Hitachi, Ltd. in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi in the United States and other countries.

VMware and ESX server are trademarks of VMware, Inc. IBM is a registered trademark of International Business Machines.

All other trademarks, service marks, and company names are properties of their respective owners.



Table of Contents

Preface	vii
Document Revision Level	viii
Changes in this Release	viii
Intended Audience	viii
Document Organization	ix
Related Web Sites	x
Document Conventions	xi
Safety and Warnings	xi
Typographic Conventions	xii
Convention for Storage Capacity Values	xii
Getting Help	xiii
Support Contact Information	xiii
Hitachi Data Systems Support Web site	xiii
Comments	xiii
 1 System Configuration Prerequisites	 1-1
Planning Your Configuration	1-2
Installation Prerequisites	1-2
Bootting from a SAN-Attached Disk	1-3
Vendor High Availability (HA) Cluster Configurations	1-4
HA Multipath Configurations	1-4
Upgrading Firmware	1-4
 2 Microsoft Windows	 2-1
Preparing the Host Server	2-2
Host Bus Adapter Guidelines	2-2
Operating System Guidelines	2-2
Connecting to the Array	2-2
Disk Number Assignments	2-4
Configuring the Host Fibre Channel HBAs	2-5
Installing Multi-path I/O (MPIO)	2-6

What is Multipathing?	2-6
Installation Prerequisites	2-6
Environmental Prerequisites	2-6
Hardware Prerequisites	2-6
Microsoft MPIO for Windows 2000 and 2003	2-7
Installing and Configuring MPIO on Windows 2008	2-7
Installation Instructions for Windows 2008	2-7
Configuring Native MPIO for Your Hitachi Storage System	2-8
Verifying Disk and Device Parameters	2-11
Verifying the Device Parameters	2-11
Configuring New Disk Devices	2-12
Writing Signatures	2-12
Creating and Formatting the Partitions	2-15
Instructions for Windows 2003	2-15
Instructions for Windows 2008	2-22
3 Solaris	3-1
Preparing the Host Server	3-2
Host Bus Adapter Guidelines	3-2
Operating System Guidelines	3-2
Multipath HA Configurations with MPxIO	3-3
Setting the Disk and Device Parameters	3-4
Configuring Fibre Channel on the Host	3-5
Verifying and Discovering LUNs	3-6
Setting and Discovering LUNs	3-6
Recognizing LUNs	3-8
Partitioning and Labeling the New Devices	3-9
4 VMware	4-1
Preparing the Host Server	4-2
Configuring the Fibre Channel HBAs	4-2
Setting the Queue Depth Parameter	4-3
Setting the Queue Depth Parameter on QLogic HBAs	4-3
Setting the Queue Depth Parameter on Emulex HBAs	4-4
Upgrading the Firmware Online	4-4
Configuring Host Groups for VMware	4-5
Logging in to Navigator 2	4-5
Reviewing RAID Groups	4-7
Creating or Editing Host Groups for VMware	4-8
Creating a VMFS and Attaching a Raw Device	4-12
Creating a Virtual Machine File System	4-12
Attaching a Raw Device	4-20

5	Red Hat Enterprise Linux	5-1
	Connecting to a Red Hat Enterprise Linux Host	5-2
	Configuring the Host Fibre Channel HBAs	5-3
	Verifying New Device Recognition	5-4
	High Availability (HA) Cluster Configurations	5-5
	HA Multipath Configurations	5-5
	Verifying and Discovering LUNs	5-6
	Partitioning the Devices	5-6
6	SuSE Linux Enterprise Server	6-1
	Connecting to an SLES Host	6-2
	Configuring the Host Fibre Channel HBAs	6-3
	Verifying New Devices	6-3
	Verifying and Discovering LUNs	6-5
	Partitioning the Devices	6-5
	High Availability (HA) Cluster Configurations	6-5
	HA Multipath Configurations	6-5
7	Asianux	7-1
	Preparing the Host Server	7-2
	Configuring the Host Fibre Channel HBAs	7-3
	High Availability (HA) Cluster Configurations	7-4
	HA Multipath Configurations	7-4
	Verifying and Discovering LUNs	7-4
	Partitioning the Devices	7-4
	Verifying New Device Recognition	7-5
8	Apple OS X	8-1
	Connecting the Host Server	8-2
	Configuring the Host Fibre Channel HBAs	8-3
	Verifying and Discovering LUNs	8-4
	Connecting the Array and Server	8-4
	Verifying Recognition of the Array	8-4
9	IBM AIX	9-1
	Preparing the Host Server	9-2
	Setting the Disk and Device Parameters	9-3
	Changing Device Parameters using SMIT	9-4
	Changing Device Parameters from the AIX Command Line	9-5
	Configuring the New Disk Devices	9-6
	Verifying New Device Recognition	9-6
	Using AIX with Object Data Manager (ODM)	9-7

ODM Overview	9-7
Installing the ODM 5.0.0.1 Fibre Update Package.	9-8
Installing the ODM Fibre Update Package for 5.0.0.4	9-10
Uninstalling the ODM Fibre Update Package	9-12
Installing the IBM AIX MPIO ODM_6000I Update	9-13
Using ODM.	9-14
Discovering New Devices.	9-14
Deleting Devices.	9-14
Queue Depth and Read/Write Timeout Values	9-15
ODM Advantages and Cautions	9-15
Advantages	9-15
Sharing Ports	9-15
ODM References.	9-16
Device Configuration Database.	9-16
Device Configuration System	9-16
List of ODM Commands and Subroutines.	9-16
IBM Redbook	9-16
10 HP-UX.	10-1
Preparing the Host Server	10-2
Host Bus Adapter Guidelines	10-2
Operating System Guidelines	10-2
Connecting to the Array	10-2
Setting Disk and Device Parameters	10-4
Verifying and Discovering LUNs.	10-4
11 Troubleshooting	11-1
Potential Error Conditions.	11-2
Calling the Hitachi Data Systems Support Center	11-2
Glossary	
Index	



Preface

This document provides guidelines on how to connect a Hitachi Modular Fibre Channel array to the following host operating systems:

- Microsoft Windows 2000 (Service Pack 4), Windows 2003 (Service Pack1), and Windows 2008
- Solaris
- VMware
- Red Hat Enterprise Linux
- SuSE Linux Enterprise (SLES)
- Asianux
- Apple/Macintosh
- IBM AIX
- HP-UX

This guide also provides information on where to find resources to ensure proper functioning between the array and the host. Please read this document carefully to understand how to implement your operating system with the Hitachi array using Fibre Channel. It is recommended that you maintain a copy for reference purposes.



NOTE: This guide assumes that the array is installed and functioning correctly. If it is not, install the system before reading this guide. Refer to your array user's guide and quick installation guide for more information.

Notice: The use of a Hitachi Modular array and all Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Document Revision Level

This section provides a history of the revision changes to this document.

Revision	Date	Description
MK-97DF8189-00	December 2008	Initial Release
MK-97DF8189-01	December 2008	Supersedes and replaces MK-97DF8189-00
MK-97DF8189-02	February 2009	Supersedes and replaces MK-97DF8189-01
MK-97DF8189-03	April 2009	Supersedes and replaces MK-97DF8189-02
MK-97DF8189-04	August 2009	Supersedes and replaces MK-97DF8189-03
MK-97DF8189-05	November 2009	Supersedes and replaces MK-97DF8189-04
MK-97DF8189-06	January 2010	Supersedes and replaces MK-97DF8189-05
MK-97DF8189-07	April 2010	Supersedes and replaces MK-97DF8189-06
MK-97DF8189-08	May 2010	Supersedes and replaces MK-97DF8189-07
MK-97DF8189-09	August 2010	Supersedes and replaces MK-97DF8189-08

Changes in this Release

- Revised [Figure 5-1 on page 5-4](#) and [Figure 5-2 on page 5-4](#).

Intended Audience

This document is intended for system administrators, Hitachi Data Systems representatives, Authorized Service Providers, and modular storage customers who are involved in verifying that the LUNs connected to the host server are discovered by the array and can be used by the operating system and other applications. It also assumes the user is familiar with the following:

- Your Hitachi array and the LUN configurations described in supporting Hitachi documents.
- Your server hardware and operating system.
- Your server's Fibre Channel host bus adapters (HBAs).
- UNIX® file system, system commands, and utilities (if applicable).
- Direct-access storage device (DASD) systems and their basic functions (if applicable).

For the latest information about operating systems supported by your storage system, refer to the interoperability matrix at:

<http://www.hds.com/products/interoperability/>

Document Organization

The following table provides an overview of the contents and organization of this document. Click the [chapter title](#) in the first column to go to that chapter. The first page of every chapter or appendix contains a brief list of the contents of that section of the manual, with links to the pages where the information is located.

Chapter/Appendix Title	Description
Chapter 1, System Configuration Prerequisites	Describes the prerequisites for setting up an array to work with supported hosts.
Chapter 2, Microsoft Windows	Describes how to connect a Windows host to the array.
Chapter 3, Solaris	Describes how to connect a Solaris host to the array.
Chapter 4, VMware	Describes how to connect a VMware host to the array.
Chapter 5, Red Hat Enterprise Linux	Describes how to connect a Red Hat Enterprise Linux host to the array.
Chapter 6, SuSE Linux Enterprise Server	Describes how to connect an SUSE Linux Enterprise Server host to the array.
Chapter 7, Asianux	Describes how to connect an Asianux host to the array.
Chapter 8, Apple OS X	Describes how to connect an Apple Macintosh host server running Apple OS X to the array.
Chapter 9, IBM AIX	Describes how to connect an IBM AIX host to the array.
Chapter 10, HP-UX	Describes how to connect an HP-UX host to the array.
Chapter 11, Troubleshooting	Contains troubleshooting information you can use in the unlikely event you encounter a problem with your array.
Glossary	Defines special terms and acronyms used in this document.
Index	Provides links to specific information in this manual.

Related Web Sites

Before you install Fibre Channel host bus adapters (HBAs), drivers, and BIOSes, verify they are supported by Hitachi Data Systems and the OS vendor by referring to the following Web sites:

- For devices supported by your array, check the interoperability matrix at <http://www.hds.com/products/interoperability/>.
- For information about Fibre Channel HBAs, check the HBA vendor Web sites.
- Operating systems-related issues, check the Web site for your operating system:
 - For Microsoft Windows:
<http://www.microsoft.com/en/us/default.aspx>
 - For Solaris:
<http://www.sun.com/>
 - For VMware:
<http://www.vmware.com/>
 - For Red Hat Enterprise Linux:
<http://www.redhat.com>
 - For SuSE Linux Enterprise Server:
<http://www.novell.com/>
 - For Asianux:
<http://www.asianux.com/asianux.do>
 - For Apple:
<http://www.apple.com>
 - For IBM AIX:
<http://www.ibm.com/us/>
 - For HP-UX:
<http://www.hp.com>












NOTE: Hitachi Data Systems is not responsible for the availability of third-party web sites mentioned in this document. Hitachi Data Systems does not endorse and is not responsible or liable for any content advertising, products, or other materials that are available on or through such sites or resources. Hitachi Data Systems will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content goods, or services that are available on or through such sites or resources.

Document Conventions

This document uses the following conventions to draw your attention to certain information.

Safety and Warnings

This document also uses the following symbols to draw your attention to certain information.

Symbol	Meaning	Description
	Tip	Tips provide helpful information, guidelines, or suggestions for performing tasks more effectively.
	Note	Notes emphasize or supplement important points of the main text.
	Caution	Cautions indicate that failure to take a specified action could result in damage to the software or hardware.
	WARNING	Warnings indicate that failure to take a specified action could result in loss of data or serious damage to the hardware.
	DANGER	The Danger symbol warns users of possible injury or death if instructions are not followed.
	ELECTRIC SHOCK HAZARD!	This symbol warns users of electric shock hazard. Failure to take appropriate precautions such as not opening or touching hazardous areas of the equipment could result in injury or death.
	Electrostatic Sensitive	The ESD symbol warns users that the equipment is sensitive to electrostatic discharge (ESD) and could be damaged if users do not take appropriate precautions such as using a grounded wrist strap when touching or handling the equipment.
	Burn Hazard	HOT SURFACE! Turn off power and allow to cool before touching.
	Sharp Edges or Corners	WARNING! Sharp edges or corners. Avoid touching or wear gloves

Typographic Conventions

The following typographic conventions are used in this document.

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click OK.
<i>Italic</i>	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: copy source-file target-file.
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: # pairdisplay -g <group>
screen/code	Indicates text that is displayed on screen or entered by the user. Example: # pairdisplay -g oradb
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.
underline	Indicates the default value. Example: [a b]

Convention for Storage Capacity Values

Storage capacity values for hard disk drives (HDDs) in Hitachi Data Systems' storage products are calculated based on the following values:

- 1 KB = 1,000 bytes
- 1 MB = 1,000² bytes
- 1 GB = 1,000³ bytes
- 1 TB = 1,000⁴ bytes

For further information on Hitachi Data Systems products and services, please contact your Hitachi Data Systems account team, or visit Hitachi Data Systems online at <http://www.hds.com>.

Getting Help

Support Contact Information

If you purchased this product from an authorized Hitachi Data Systems reseller, contact that reseller for support. For the name of your nearest Hitachi Data Systems authorized reseller, refer to the Hitachi Data Systems support website for locations and contact information.

To contact the Hitachi Data Systems Support Center, please visit the Hitachi Data Systems website for current telephone numbers and other contact information. <https://extranet.hds.com/http://aim.hds.com/portal/dt>

Please provide at least the following information about the problem:

- Product name, model number, part number (if applicable) and serial number
- System configuration, including names of optional features installed, host connections, and storage configuration such as RAID groups and LUNs
- Operating system name and revision or service pack number
- Exact content of any error message(s) displayed on the host system(s)
- The circumstances surrounding the error or failure
- A detailed description of the problem and what has been done

Hitachi Data Systems Support Web site

The following pages on the Hitachi Data Systems support Web site contain other further help and contact information:

Home Page: <https://extranet.hds.com/http://aim.hds.com/portal/dt>

Comments

Please send us your comments on this document:

doc.comments@hds.com. Include the document title, number, and revision, and refer to specific section(s) and paragraph(s) whenever possible.

Thank you! (All comments become the property of Hitachi Data Systems Corporation.)

System Configuration Prerequisites

This chapter provides an overview of the tasks that need to be completed before you connect the host to a Hitachi array.

This chapter covers the following key topic:

- ❑ [Planning Your Configuration](#)
- ❑ [Installation Prerequisites](#)
- ❑ [Bootting from a SAN-Attached Disk](#)
- ❑ [Vendor High Availability \(HA\) Cluster Configurations](#)
- ❑ [Upgrading Firmware](#)

Planning Your Configuration

There are several ways to configure your Hitachi array. Therefore, before you connect the host server to the array, we recommend you plan your configuration. For more configuration information, refer to your array's user's guide.

Before implementing any configuration, decide whether the following redundancy is useful and appropriate for your needs. Redundancy, such as Fibre Channel LUN multipathing, on the I/O path between the host and the LUNs on the array is created using multi-path software. Multipathing eliminates host system outages due to single-path I/O failures, thereby increasing system reliability.

Installation Prerequisites

Table 1-1 summarizes required installation and configuration tasks for the host and array to work properly. It includes references where you can find detailed information for each task. Verify the tasks outlined below have been completed before connecting the host to the array.

Table 1-1: Installation Tasks

Task	Reference
Install all array hardware and cabling. To avoid damage to the array or array components due to electrostatic discharge, wear an anti-static wrist strap when handling the array. Connect the clip to an unpainted part of the array chassis frame to safely channel any static electricity generated by your body to ground. If no wrist strap is available, ground yourself by touching an unpainted part of the array chassis frame.	Your array comes with all the hardware and cabling required for installation. Refer to your array's user's guide.
Upgrade to the latest array firmware if necessary.	Refer to the following: <ul style="list-style-type: none">• Latest Release Notes for your array• Your array's user's guide• Storage Navigator Modular 2 (Navigator 2) online help.
Identify and configure the topology in which the array will be used.	Refer to your array's user's guide.
Set up arrays and create LUNs	Use Navigator 2 and refer to the online help.
Set the host group operation parameters.	Use Navigator 2 and refer to the online help.

Table 1-1: Installation Tasks (Continued)

Task	Reference
Host Bus Adapters (HBAs)	One or more supported HBAs with the latest supported BIOS and driver are required. Verify that the HBAs, drivers, and BIOSes are the latest versions supported by the array, and are functioning properly. <ul style="list-style-type: none">• To check the latest supported versions, refer to the Hitachi interoperability matrix: http://www.hds.com/products/interoperability/.• For information about HBAs supported by your operating system, refer to the HBA vendor and operating system Web sites.
Install the HBA in the host server.	For installation information, check the documentation and Web site of your HBA vendor. Be sure the HBA is supported by your array (refer to the Hitachi interoperability matrix listed above).
Set operating system parameters.	Refer to the operating system vendor's documentation or Web site for parameter settings.
For high availability, check whether the multiple-port function is set.	Refer to the interoperability matrix listed above.

Booting from a SAN-Attached Disk

To boot from a SAN-attached disk, verify that the configurations on your array and your HBA support boot from SAN.

To verify boot from SAN support on your system:

- Contact your Hitachi Data Systems account representative to determine the supported configurations when booting from your array LUN.
- Contact your HBA vendor or check the vendor's Web site to verify your HBA card will support a boot-from-SAN configuration (for example, OS maintenance level, HBA model, HBA BIOS level, etc.).
- Check your operating system's site for available documentation on booting from SAN and how to install the operating system on external LUNs (direct or in a fabric).

Vendor High Availability (HA) Cluster Configurations

For vendor cluster software supported by Hitachi Data Systems:

- Consult the vendor documentation for installation and configuration details.
- For Windows 2003 or 2008 operating systems, the Hitachi array requires no unique cluster-related settings.
- For required array parameter/mode settings related to supported vendor clusters, refer to your array's user's guide and the Navigator 2 online help.

HA Multipath Configurations

Hitachi modular arrays supports various HA multipathing software products for supported operating systems. Contact your Hitachi Data Systems account representative for the latest information on supported software products, or refer to the Hitachi Data Systems interoperability support matrix:

<http://www.hds.com/products/interoperability/>

Multipathing software should be installed and configured before connecting to the array.

Upgrading Firmware

If you upgrade fault-tolerant firmware in a heavy loaded environment, I/O may time out.

Microsoft Windows

This chapter discusses guidelines on how to prepare the following host servers for connection to the array and verify that the host server can connect to the array:

- Microsoft Windows 2000 (Service Packs 3/4)
- Microsoft Windows 2003 and 2003 Server (Service Pack 1)
- Microsoft Windows 2008
- Microsoft Windows XP (Service Pack 2)

This chapter covers the following key topics:

- ❑ [Preparing the Host Server](#)
- ❑ [Connecting to the Array](#)
- ❑ [Disk Number Assignments](#)
- ❑ [Configuring the Host Fibre Channel HBAs](#)
- ❑ [Installing Multi-path I/O \(MPIO\)](#)
- ❑ [Verifying Disk and Device Parameters](#)
- ❑ [Configuring New Disk Devices](#)

Preparing the Host Server

This section provides guidelines to consider when preparing your host server.

Host Bus Adapter Guidelines

One or more host bus adapters (HBAs) are required to connect your host server to the array. Before installing and connecting HBAs, verify that the HBAs are working properly and have the latest drivers and BIOSes supported by your operating system and the array.

For information about:

- The latest HBAs supported by the array, refer to the Hitachi interoperability matrix at:
<http://www.hds.com/products/interoperability/>.
- Installing and configuring your HBAs, refer to the HBA documentation and the HBA vendor's Web site.
- HBAs supported by your operating system, refer to the Microsoft Windows Web site:
<http://www.microsoft.com/>.

Operating System Guidelines

Verify that the OS version, architecture, relevant patches, and maintenance levels you intend to use are supported by Hitachi Data Systems. For information about supported versions, refer to the Hitachi interoperability matrix.

Connecting to the Array

After you prepare the array hardware and software and Fibre Channel HBA(s), connect the array to the Windows 2003 or 2008 system. The array comes with all the hardware and cabling required for connection to the host system(s).

[Table 2-1 on page 2-3](#) lists the tasks to follow to prepare the host server. To prepare the server, you need one or more supported HBAs with the latest supported BIOS and driver. Verify the HBAs, drivers, and BIOSes are the latest supported versions by Hitachi Data Systems and are functioning properly.

Table 2-1: Steps for Connecting to a Windows Host

Step	Activity	Description
1.	Verify the array installation.	Confirm that the status of the Fibre Channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the Windows system.	<ul style="list-style-type: none">• Shut down the Windows system.• Power off all peripheral devices, except for the array.• Power off the host system.• You are now ready to connect the array.
3.	Connect the array.	<ul style="list-style-type: none">• Install Fibre Channel cables between the array and the Windows system.• Follow all precautions and procedures in your array's user's guide.• Check all specifications to ensure proper installation and configuration.
4.	Define LUN security if connected to a fabric.	If the array is connected to a fabric: <ul style="list-style-type: none">• Control access to the array before you power on Windows to avoid failures on other systems connected to the same fabric.• Use any method to ensure that Windows sees only the LUs it owns.
5.	Power on the Windows system.	<ul style="list-style-type: none">• Power on the Windows system display.• Power on all peripheral devices. The array should be on, the Fibre Channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the Windows system is powered on, restart the system to have the new devices recognized.• Confirm the ready status of all peripheral devices, including the array.• Power on the Windows system.
6.	Verify new device recognition and record the disk numbers.	When the adapter connected to the array shows the new devices, pause the screen and record the disk number for each new device on your SCSI Device worksheet (Table 2-2 on page 2-4).
7.	Boot the Windows system.	

Disk Number Assignments

The Windows 2003 and 2008 operating systems assign disk numbers sequentially, starting with the local disks and then by adapter, and by TID/LUN.

- If the array is attached to the first adapter (displayed first during system start-up), the disk numbers for the new devices will start at 1 (the local disk is 0).
- If the array is not attached to the first adapter, the disk numbers for the new devices will start at the next available disk number. For example, if 40 disks are attached to the first adapter (disks 1-40) and the array is attached to the second adapter, the disk numbers for the array will start at 41.



NOTE: When disk devices are added to or removed from the Windows 2003 or 2008 system, the disk numbers are reassigned automatically.

Table 2-2: Sample SCSI Device Information Worksheet

LDEV (CU:LDEV)	LU Type	VLL (MB)	Device Number	Bus Number	Path 1	Alternate Path(s)	
0:00					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:01					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:02					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:03					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:04					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:05					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:06					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:07					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:08					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:09					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:0a					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:0b					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:0c					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:0d					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____

Table 2-2: Sample SCSI Device Information Worksheet (Continued)

LDEV (CU:LDEV)	LU Type	VLL (MB)	Device Number	Bus Number	Path 1	Alternate Path(s)	
0:0e					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:0f					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
0:10					TID: _____ LUN: _____	TID: _____ LUN: _____	TID: _____ LUN: _____
and so on...					TID: _____ LUN: _____		

Configuring the Host Fibre Channel HBAs

Configure the Fibre Channel HBA(s) connected to the array. HBAs have many configuration options. This section provides the minimum requirements for configuring host Fibre Channel HBAs for operation with the array. For more information, refer to the documentation for your Fibre Channel HBA(s).

- The queue depth requirements for array devices are specified in [Table 2-3](#).
- You may need to disable the BIOS to prevent the system from trying to boot from the array.
- Several other parameters (for example, Fibre Channel and/or fabric options) may also need to be set. Please refer to the documentation for your HBA to determine whether other options are required to meet your operational requirements.
- Use the same settings and device parameters for all array devices.

If your HBA does not have a setup utility, or if your HBA setup utility does not provide access to the required parameters, use the Windows 2003 or 2008 Registry Editor to set the required parameters. Ask your Hitachi Data Systems representative for assistance in setting these parameters. See [Verifying the Device Parameters on page 2-11](#) for instructions about configuring the HBA settings using the Registry Editor.

Table 2-3: Queue Depth Requirements for the Array

Parameter	Required Value for Array
Queue depth per LU	≤ 32 ports
Queue depth per port (MAXTAGS)	≤ 512 per port

Installing Multi-path I/O (MPIO)

This section provides information about installing and using Multipath I/O (MPIO). MPIO is a feature that provides support for using multiple data paths to a storage device. Multipathing increases the availability of storage resources by providing path failover from a server or cluster to a storage subsystem.

You must install MPIO on a server if it will access a logical unit number (LUN) through multiple Fibre Channel HBAs.

What is Multipathing?

Multipathing gives systems the ability to use more than one read/write path to a storage device using redundant physical path components — adapters, cables, and switches — between the server and storage device. Multipathing is key in keeping mission-critical data continuously available in the enterprise environment, where the goal for many organizations is continuous availability.

Because multipathing allows two or more data paths to be used simultaneously for read/write operations, a failure with one or more components still allows applications to access their data. In addition to providing fault tolerance, multipathing also serves to redistribute the read/write load among multiple paths between the server and storage, helping to remove bottlenecks and balance workloads.

Installation Prerequisites

Observe the following prerequisite before installing MPIO.

Environmental Prerequisites

- The host computer(s) must be connected to the Hitachi storage system using dual-redundant data paths, either end-to-end or via routing devices such as network switches. The driver-level MPIO then manages the redundant connections.
- All Fibre Channel HBAs and their associated drivers should be installed to the host computer(s) before you install MPIO.

Hardware Prerequisites

- To create path redundancy in a storage network or applications using other host-side links, apply at least two single-ported HBAs or two dual-ported Fibre Channel HBAs on the host computer(s). To optimize performance, place the Fibre Channel HBAs on different buses in the server to distribute the workload over the server's PCI bus architecture.
- Before installing and using MPIO, RAID volumes must be created and properly associated with host ID/LUNs to use the MPIO functions.

Microsoft MPIO for Windows 2000 and 2003

For information about MPIO for Fibre Channel arrays on Windows 2000/2003 servers, please refer to the Microsoft Web site:

<http://www.microsoft.com/en/us/default.aspx>

Installing and Configuring MPIO on Windows 2008

On Windows 2008, MPIO is available as a standard installable feature. The following procedure provides general guidelines for installing MPIO on a Windows 2008 operating system. For more information, refer to <http://go.microsoft.com/fwlink/?LinkId=81020>.

Installation Instructions for Windows 2008

Before proceeding with the installation, observe the following guidelines:

- If you will be using MPIO, refer to <http://technet.microsoft.com/en-us/library/cc725907.aspx> and your host server documentation for information about implementing MPIO.
- If you will enable access to a LUN from a cluster, be sure that Failover Clustering is installed on each server in the cluster; otherwise, data loss can occur. For more information, refer to <http://go.microsoft.com/fwlink/?LinkId=86168>.
- You must have Administrator privileges on the computer to install MPIO.

To install MPIO on a Windows 2008 operating system:

1. In the Server Manager console tree, click the **Features** node.
2. In the Features pane, under **Features Summary**, click **Add Features**.
3. In the Add Features Wizard, select the **Multipath I/O** check box and click **Next**.
4. Follow the steps in the Add Features Wizard to complete the installation.
5. After the wizard is completed, click the **Close** button and verify that MPIO is installed:
 - a. Click the Windows **Start** button, point to **Administrative Tools**, and click **Server Manager**. The Server Manager window opens.
 - b. In the left pane, click the **Features** category if it is not selected.
 - c. Confirm that **Multipath I/O** appears under the Features group (see [Figure 2-1 on page 2-8](#)). You may need to expand the group to see all installed features.
 - d. Close or minimize the Server Manager window.

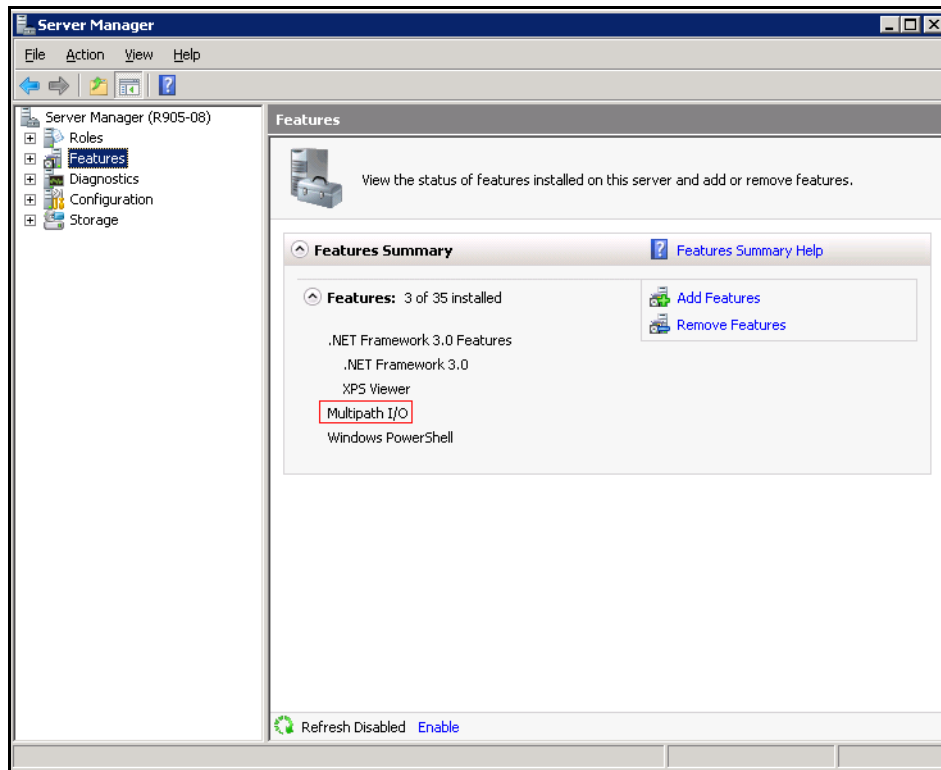


Figure 2-1: Server Manager Window with Multipath I/O Displayed

Configuring Native MPIO for Your Hitachi Storage System

After MPIO is installed, use the following procedure to add the device hardware ID for your Hitachi storage system. This procedure requires you to reboot your system. Optionally, you can also remove the default MPIO hardware ID, but this requires a second system reboot.

1. Click the Windows **Start** button.
2. Point to **Administrative Tools** and click **MPIO**. The **MPIO-ed Devices** tab of the MPIO Properties dialog box appears. In this panel, the **Device Hardware Id** field shows the IDs of the MPIO hardware devices installed. The first time this dialog box appears, the default string in [Figure 2-2 on page 2-9](#) is displayed to indicate that the vendor ID is limited to 8 characters and the product ID is limited to 16 characters.
3. Optional: To remove the default MPIO hardware ID, perform the following steps; otherwise, proceed to step 4 on the next page.
 - a. Select the default string in the **Device Hardware Id** field and click the **Remove** button to delete the string. The Reboot Required message in [Figure 2-3 on page 2-9](#) appears.

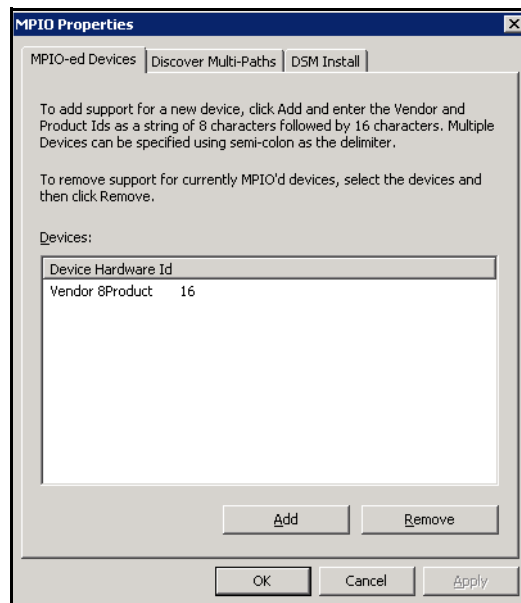


Figure 2-2: MPIO Properties Dialog Box




Figure 2-3: Reboot Required Message

- b. Click **Yes** to reboot your host.
- c. After your host reboots, click the Windows **Start** button, point to **Administrative Tools**, and click **MPIO** to display the **MPIO-ed Devices** tab of the MPIO Properties dialog box again.
- d. Continue with step 4.
4. Click the **Add** button. The Add MPIO Support dialog box appears (see [Figure 2-4](#)).

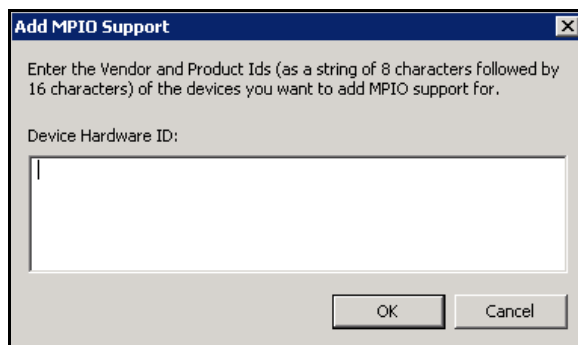


Figure 2-4: Add MPIO Support Dialog Box

5. In the **Device Hardware ID** field, type: **HITACHI DF600F**

6. Click the **OK** button. The Reboot Required message in [Figure 2-5](#) appears.




Figure 2-5: Reboot Required Message

7. Click **Yes** to reboot your host.
8. After the host reboots, confirm that the device hardware ID was added to the MPIO configuration:
 - a. Click the Windows **Start** button.
 - b. Point to **Administrative Tools**.
 - c. Click **MPIO**. The **MPIO-ed Devices** tab of the MPIO Properties dialog box appears. In this panel, verify that the **Device Hardware Id** field shows the vendor and product IDs of your Hitachi MPIO hardware (as well as any other MPIO hardware devices that may be installed). See [Figure 2-6](#).

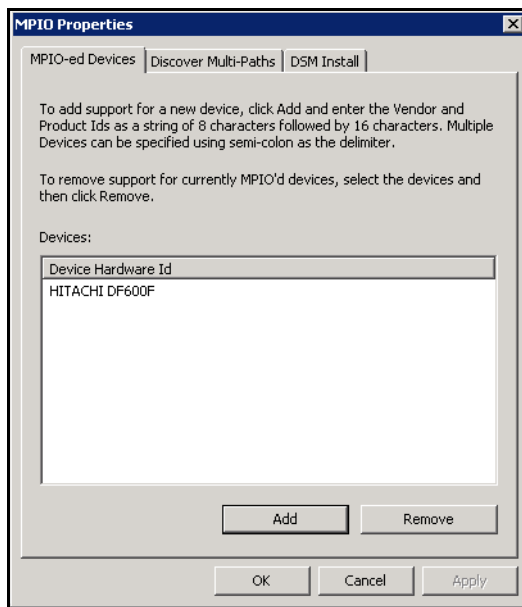


Figure 2-6: MPIO Properties Dialog Box with Hitachi Vendor and Product IDs

This completes the procedure for configuring MPIO on Windows 2008 operating systems for use with your Hitachi storage system.

Verifying Disk and Device Parameters

After you configure the fibre-channel ports during boot-up, verify the required disk and device parameters using the Windows 2003 or 2008 Registry. As part of this task, you should:

- Verify the queue depth.
- Verify other required parameters such as FC fabric support and link down timeout.

Verifying the Device Parameters

The queue depth parameter for the array devices must be set as specified in [Table 2-3 on page 2-5](#). In addition, verify all other required settings for your operational environment (for example, FC fabric support) and be sure device parameters are the same for all array devices.

This section provides sample instructions for the Emulex and QLogic fibre-channel HBAs. For other HBAs, refer to the user documentation for the HBAs.



NOTE: For Emulex HBAs, you can also use Emulex utilities to set driver parameters (refer to <http://www.emulex.com>).

To verify the queue depth and other device parameters using the Registry Editor:

1. Start the Windows 2003 or 2008 Registry Editor.
2. For each Emulex Fibre Channel HBA:
 - a. For the SCSI Port Driver, Queue Depth (Maximum Queue Depth) is set using Emulex Configuration Tool as shown in [Figure 2-7 on page 2-12](#). Queue depth requirements are listed in [Table 2-3 on page 2-5](#).
 - b. For the SCSI Port Driver, the Link Timer value should be kept at its default value of 30.
3. For each QLogic Fibre Channel HBA:
 - a. Go to **HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Services > ql2300 > Parameters > Device** to display the device parameters for the QLogic fibre-channel HBA.
 - b. If connected to a fabric switch, add **FabricSupported=1** to the **DriverParameter**.
 - c. Verify all other required settings for your environment (e.g., support for more than eight LUNs per array ID). Refer to the HBA documentation as needed.
4. To change any HBA settings, reboot the Windows 2003 or 2008 system, and use the HBA setup utility. If you do not want to reboot, edit the registry carefully (Note on previous page about using the Registry Editor).
5. Save your changes and exit the Registry Editor.




Figure 2-7: Emulex Configuration Tool Panel

Configuring New Disk Devices

This section describes how to prepare the host servers for connection to the array.



NOTE: The array logical devices are defined to the host as SCSI disk devices, even though the interface is Fibre Channel.

Writing Signatures

The first step when configuring new devices is to write a signature on each device using the Windows 2003 or 2008 Disk Management. You must write a signature on each SCSI disk device to enable the Windows 2003 or 2008 system to vary the device online. For MSCS environments, you must also write signatures on the raw devices. The 32-bit signature identifies the disk to the Windows 2003 or 2008 system. If the disk's TID and/or LUN is changed, or if the disk is moved to a different controller, the Disk Management and Windows 2003 or 2008 fault-tolerant driver will continue to recognize it.



NOTE: Microsoft Windows 2003 and 2008 assign disk numbers sequentially, starting with the local disks and then by adapter, and by TID/LUN. If the array is attached to the first adapter (displayed first during system start-up), the disk numbers for the new devices start at 1 (the local disk is 0). If the array is not attached to the first adapter, the disk numbers for the new devices start at the next available disk number. For example, if 40 disks are attached to the first adapter (disks 1-40) and the array is attached to the second adapter, the disk numbers for the array start at 41.

To write the signatures on the new disk devices:

1. Click the **Start** button, point to **Programs**, point to **Administrative Tools (Computer Management)**, and click **Disk Management** to start the Disk Manager. Initialization takes a few seconds.
 2. When the Disk Management notifies you that one or more disks have been added, click **OK** to allow the system configuration to be updated.
-



NOTE: If you removed any disks, the Disk Management will notify you at this time.

3. The Disk Management displays each new device by disk number and asks if you want to write a signature on the disk ([Figure 2-10 on page 2-15](#)). You may only write a signature once on each device.
 - For all SCSI disk devices, click **OK** to write a signature.
4. After you write or decline to write a signature on each new device, switching to offline enables the Disk Management main panel to display the devices by disk number ([Figure 2-8](#) and [Figure 2-9 on page 2-14](#)). The total capacity and free space is displayed for each disk device with a signature. Configuration information not available indicates no signature. Proceed to [Creating and Formatting the Partitions on page 2-15](#) to create partitions on the new SCSI disk devices.




Figure 2-8: Disk Management Panel Showing New Devices




Figure 2-9: Switching to Offline




Figure 2-10: Writing the Signatures

Creating and Formatting the Partitions

After write signatures on the new devices, create and format the partitions on the new SCSI disk devices

Dynamic Disk is supported with no restrictions for the array connected to the Windows 2003 or 2008 operating system. Using the Dynamic Disk on a Windows Server 2008 operating system may require you to reduce the logical unit capacity. For more information, refer to the Microsoft Windows online help.



NOTE: Do not partition or create a file system on a device which will be used as a raw device (e.g., some database applications use raw devices).

Instructions for Windows 2003

To create and format partitions on the new SCSI disk devices using Windows 2003:

1. On the Disk Management main panel, select the unallocated area for the SCSI disk you want to partition, click the **Action** menu, and click **Create Partition** to launch the New Partition Wizard.
2. When the Select Partition Type panel appears (Figure 2-11), select the desired type of partition and click **Next**.



NOTE: The array does not support Stripe Set Volume with parity.

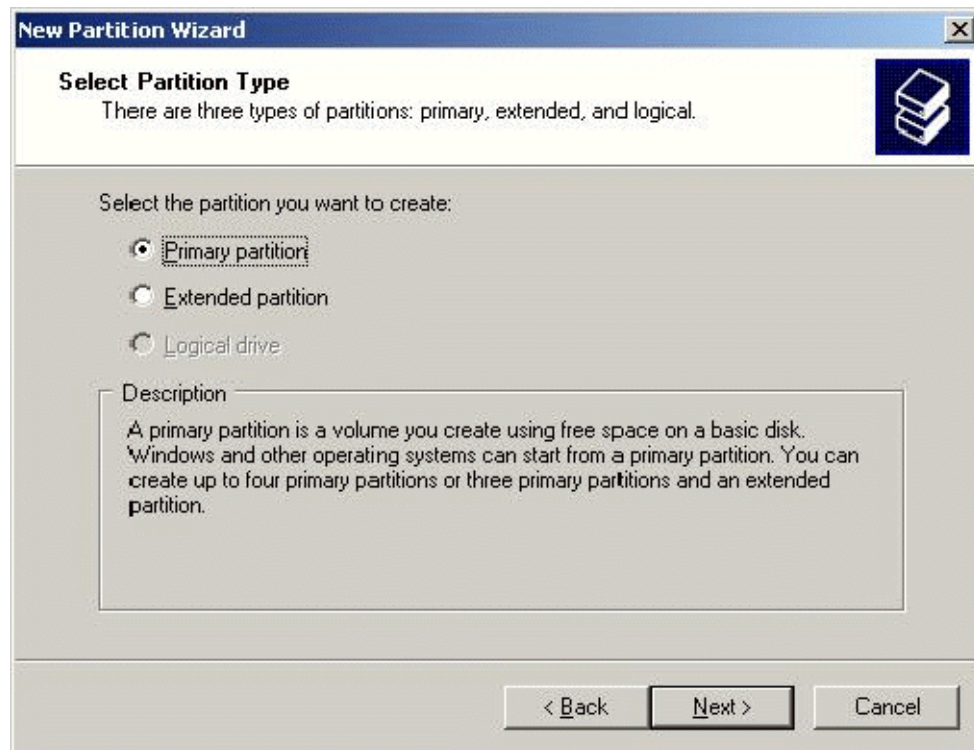


Figure 2-11: New Partition Wizard Panel (Windows 2003)

3. When the Specify Partition Size panel appears ([Figure 2-12](#)), specify the desired partition size. If the size is greater than 1024 MB, you will be asked to confirm the new partition. Click **Next**.

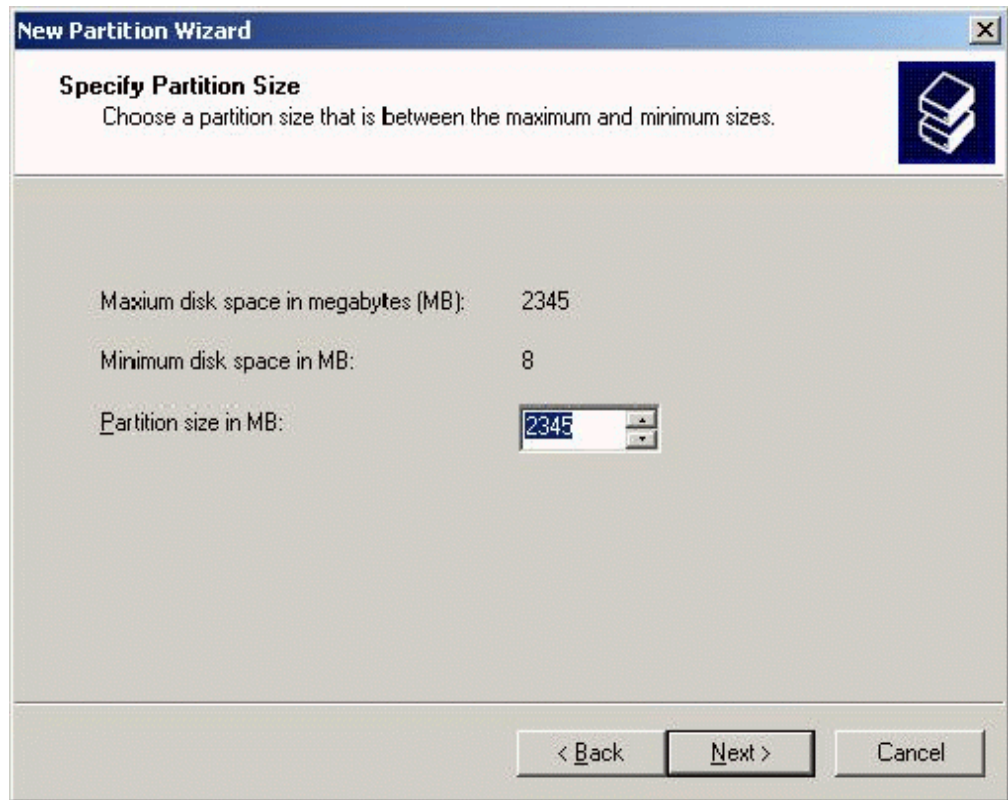


Figure 2-12: Specifying the Partition Size

4. When the Assign Drive Letter or Path panel appears ([Figure 2-13 on page 2-18](#)), select a drive letter or path, or specify no drive letter or drive path. Click **Next**.

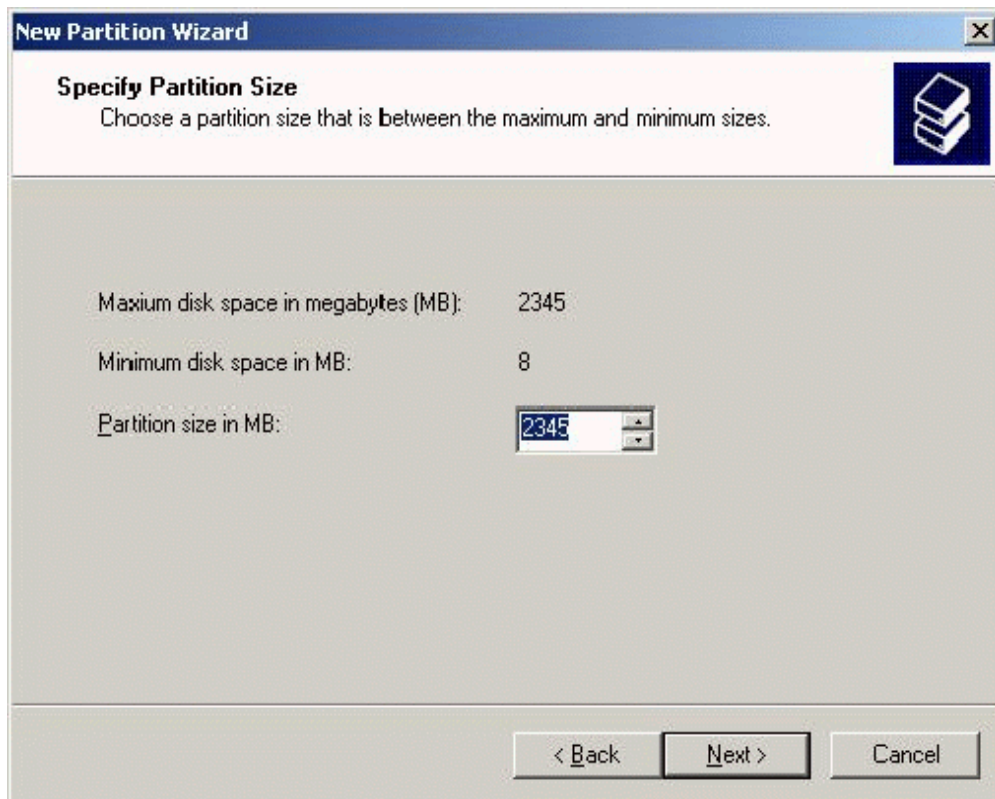


Figure 2-13: Assigning the Drive Letter or Path

5. When the Format Partition panel appears ([Figure 2-14 on page 2-19](#)), click **Format this partition with the following settings** and select the following options:
 - File System: Select **NTFS** (enables the Windows 2003 or 2008 system to write to the disk).
 - Allocation unit size: **Default**. Do not change this entry.
 - Volume label: Enter a volume label, or leave this field blank for no label.
 - Format Options: Select **Perform a Quick Format** to decrease the time required to format the partition. Select **Enable file and folder compression only** if you want to enable compression.

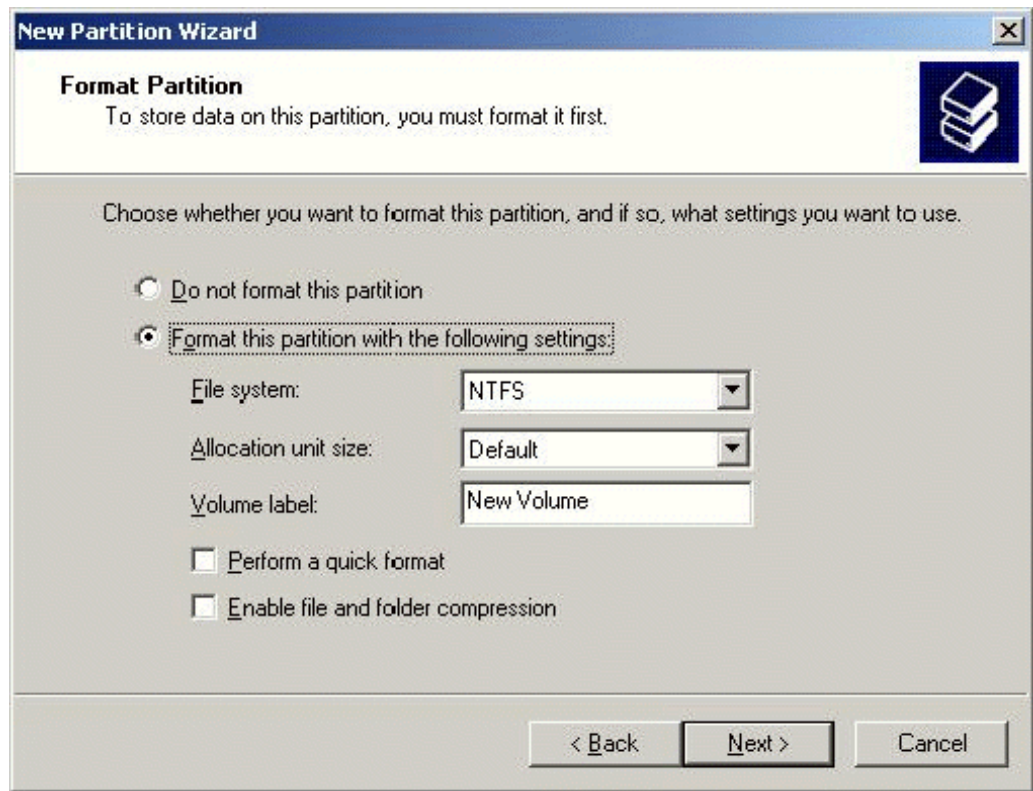


Figure 2-14: Formatting the Partition

6. Select **Next** to format the partition as specified. When the format warning appears (this new format will erase all existing data on disk), click **OK** to continue. The Format panel shows the progress of the format partition operation.
7. When the format operation is complete, click **OK**. The New Partition Wizard displays the new partition ([Figure 2-15 on page 2-20](#)). Click **Finish** to close the New Partition Wizard.

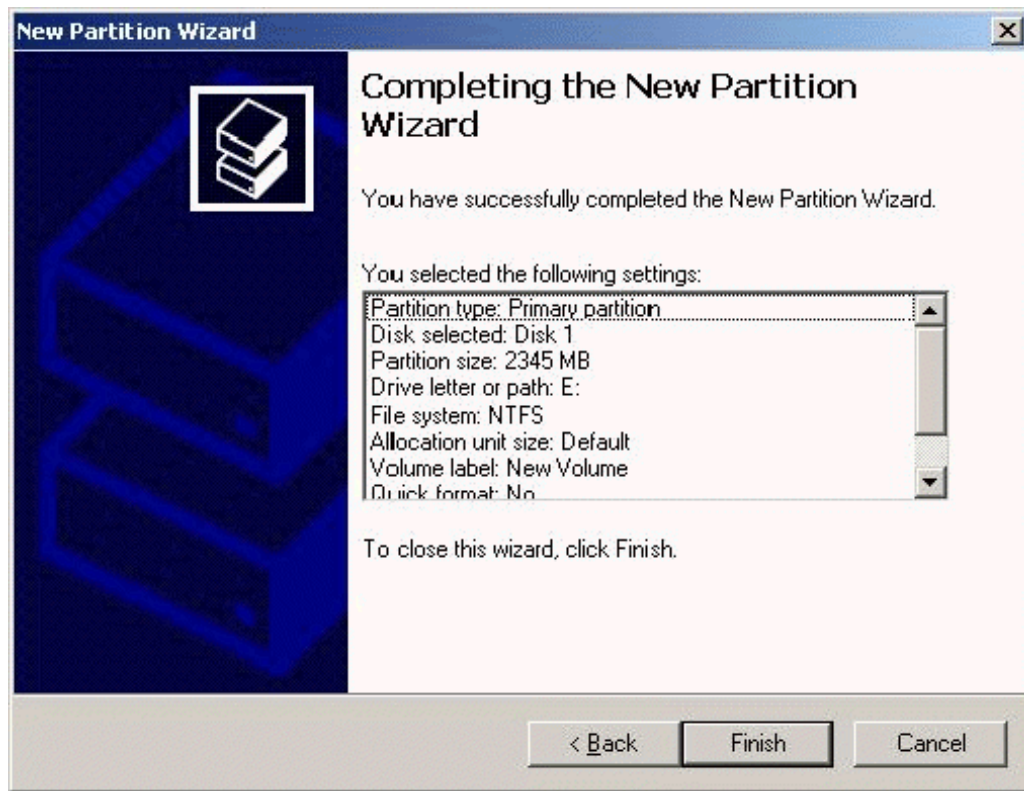


Figure 2-15: Successful Formatting Confirmation Panel

8. Verify that the Disk Management main panel shows the correct file system (NTFS) for the formatted partition ([Figure 2-16 on page 2-21](#)). The word **Healthy** indicates that the partition has been created and formatted successfully.




Figure 2-16: Verifying the Formatted Partition

9. Repeat steps 1 through 8 for each new SCSI disk device. When you finish creating and formatting partitions, exit the Disk Management. When the disk configuration change message appears, click **Yes** to save your changes.



NOTE: Be sure to make your new Emergency Repair Disk using RDISK.EXE.

Instructions for Windows 2008

To create and format partitions on the new SCSI disk devices using Windows 2008:

1. On the Disk Management main panel, select the unallocated area for the SCSI disk you want to partition.
2. On the **Action** menu, click **Create Partition** to launch the New Simple Volume Wizard (Figure 2-17).




Figure 2-17: Starting the New Simple Volume Wizard

3. Click **Next**.
4. When the Select Partition Type panel appears, select the desired type of partition and click **Next**.



NOTE: The array does not support Stripe Set Volume with parity.

5. When the Specify Volume Size panel appears (Figure 2-18 on page 2-23), specify the desired volume size. If the size is greater than 1024 MB, you will be asked to confirm the new volume. Click **Next**.

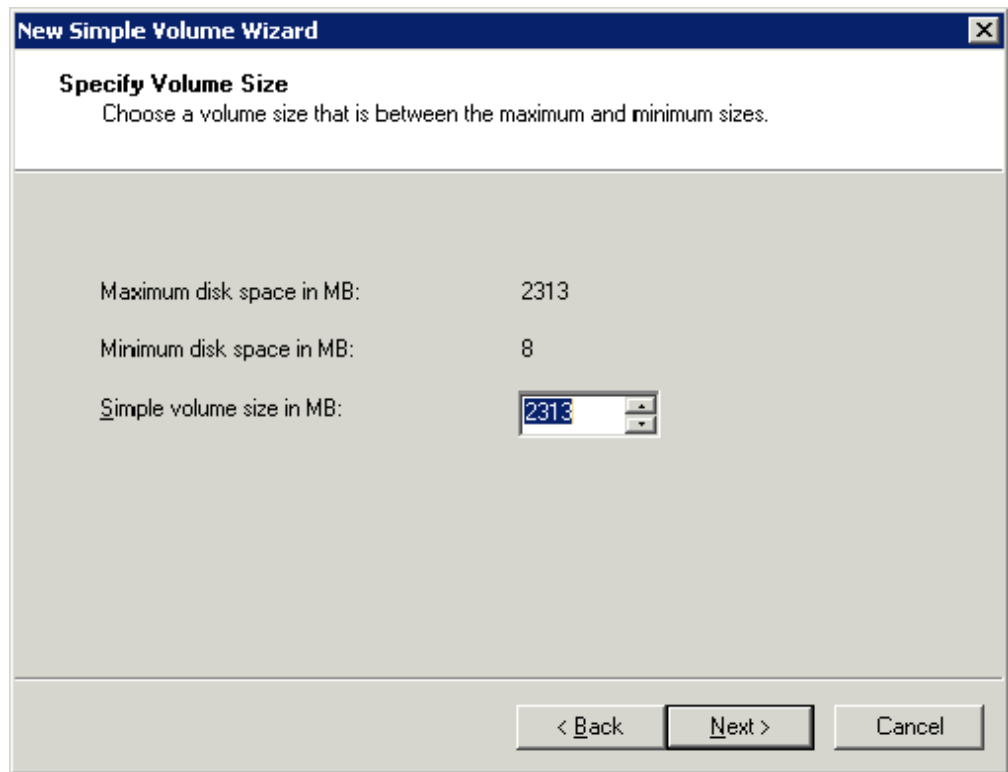


Figure 2-18: Specify Volume Size Panel

6. When the Assign Drive Letter or Path panel appears ([Figure 2-19](#)), select a drive letter or path, or specify no drive letter or drive path. Click **Next**.

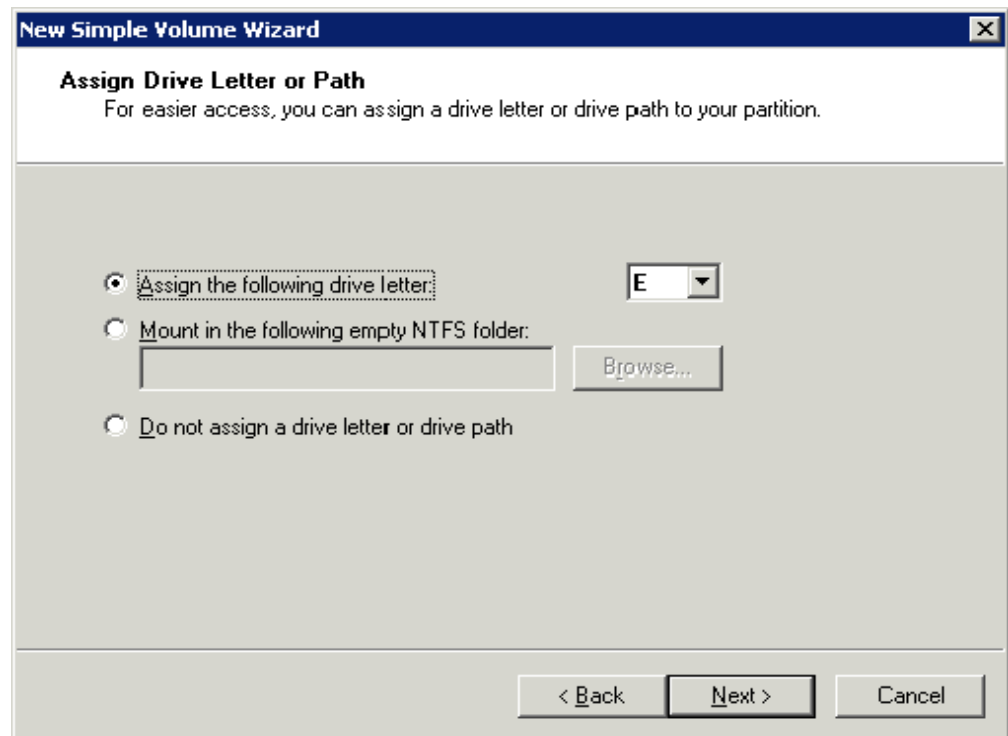


Figure 2-19: Assign Drive Letter or Path Panel

7. When the Format Partition panel appears ([Figure 2-20](#)), click Format this partition with the following settings and select the following options:
 - **File System:** Select **NTFS** (enables the Windows 2008 system to write to the disk).
 - **Allocation unit size:** **Default**. Do not change this entry.
 - **Volume label:** Enter a volume label, or leave this field blank for no label.
 - **Format Options:** Select **Perform a Quick Format** to decrease the time required to format the partition. Select **Enable file and folder compression** if you want to enable compression.




Figure 2-20: Format Partition Panel

8. Click **Next** to format the partition as specified. When the format warning appears (this new format will erase all existing data on disk), click **OK** to continue. The Format panel shows the progress of the format partition operation.
9. When the format operation finishes, click **OK**. The New Partition Wizard displays the new partition ([Figure 2-21 on page 2-25](#)). Click **Finish** to close the New Partition Wizard.

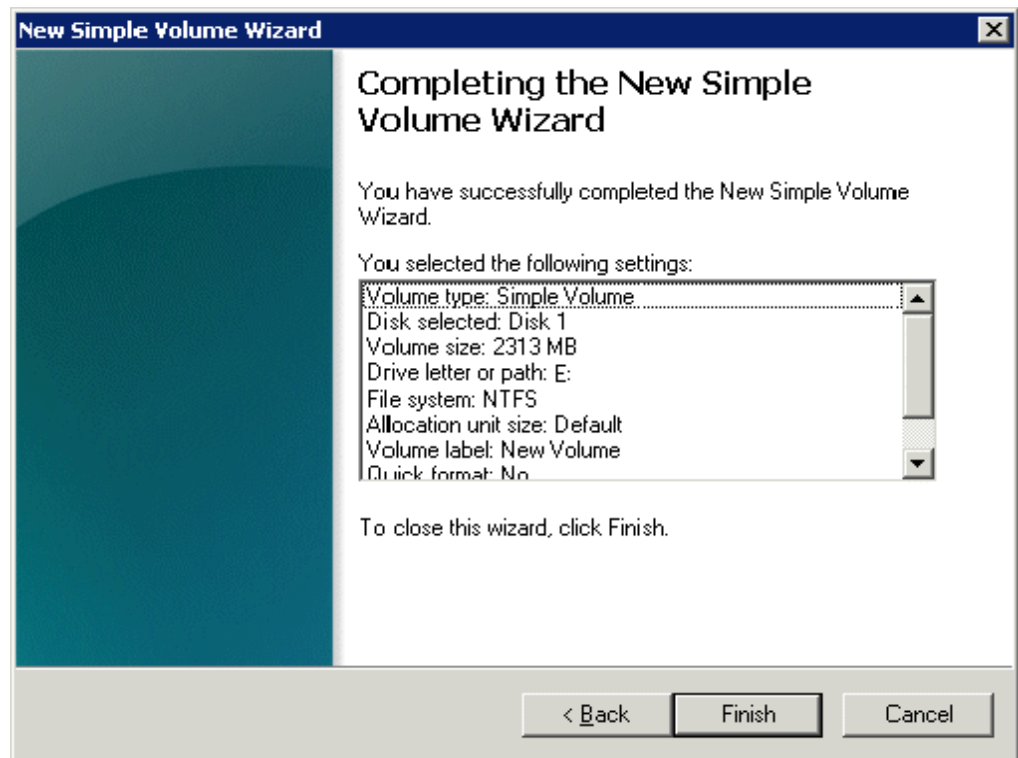


Figure 2-21: Successful Formatting Confirmation Panel

10. Verify that the Disk Management main panel shows the correct file system (NTFS) for the formatted partition ([Figure 2-22](#)). The word **Healthy** indicates that the partition has been created and formatted successfully.




Figure 2-22: Verifying the Formatted Partition

11. Repeat steps 1 through 10 for each new SCSI disk device. When you finish creating and formatting partitions, exit the Disk Management. When the disk configuration change message appears, click **Yes** to save your changes.



NOTE: Make your new Emergency Repair Disk using RDISK.EXE.

Solaris

This chapter discusses guidelines on how to prepare a Solaris host server for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Preparing the Host Server](#)
- ❑ [Setting the Disk and Device Parameters](#)
- ❑ [Configuring Fibre Channel on the Host](#)
- ❑ [Verifying and Discovering LUNs](#)

Preparing the Host Server

This section provides guidelines to consider when preparing your host server for connection to the array.

Host Bus Adapter Guidelines

One or more host bus adapters (HBAs) are required to connect your host server to the array. Before installing and connecting HBAs, verify that the HBAs are working properly and have the latest drivers and BIOSes supported by your operating system and the array.

For information about:

- The latest HBAs supported by the array, refer to the Hitachi interoperability matrix at:
<http://www.hds.com/products/interoperability/>
- Installing and configuring your HBAs, refer to the HBA documentation and the HBA vendor's Web site.
- HBAs supported by your operating system, refer to the Sun Web site:

<http://www.sun.com/>

Operating System Guidelines

Verify that the operating system version, architecture, relevant patches, and maintenance levels you intend to use are supported by Hitachi Data Systems. For information about supported versions, refer to the Hitachi interoperability matrix.

Multipath HA Configurations with MPxIO

MPxIO provides multiple paths to the same LUN, allowing failover to another path. If you use MPxIO for Sun Solaris on a SPARC station running Solaris v10 OS U5 or higher, you can download a patch from the following URL that will enable the use of both array controllers in a multipath HA configuration:

<http://sunsolve.sun.com/search/document.do?assetkey=1-21-138308-02-1>

To use this patch, your server and array system must meet the requirements in [Table 3-1](#) and [Table 3-2](#).

Table 3-1: MPxIO Settings for Hitachi Modular Storage Systems

Parameters		Settings
Storage Navigator Modular 2 Settings	Simple Settings	Platform = Solaris Middleware = none
	Host Connection Mode 1	Standard Mode
	Host Connection Mode 2	None

Table 3-2: MPxIO Host Settings

Activity	Description
Solaris 8, Solaris 9, and Solaris 10 U4 or earlier	Patch = not necessary
	scsi_vhci.conf = edit as below: <ul style="list-style-type: none">• load-balance = round-robin• auto-failback = enable• device-type-scsi-options-list = HITACHI DF600F• symmetric-option• symmetric-option = 0x1000000
Solaris 10 U5	Patch = one of the following patches is required, based on the Solaris version installed on the host: <ul style="list-style-type: none">• SPARC = patch 138308-02• x64/x86 = patch 138309-02
	scsi_vhci.conf = use default settings
Solaris 10 U6	Patch = not necessary
	scsi_vhci.conf = use default settings

Setting the Disk and Device Parameters

The queue depth parameter (max_throttle) for the array devices must be set according to one of the options specified in [Table 3-3](#).

Table 3-3: Max Throttle (Queue Depth) Requirements

Queue Depth Option	Requirements
Option 1	32 commands per LUN
Option 2	512 commands per port



NOTE: You can adjust the queue depth for the array devices later as needed (within the specified range) to optimize the I/O performance of the array devices.

Several other parameters (e.g., Fibre Channel support) may also need to be set. Please refer to the user documentation for your HBA to determine whether other options are required to meet your requirements.



NOTE: Use the same settings and device parameters for all array devices. For Fibre Channel, the settings in the system file apply to the entire system, not to just the HBA(s).

To set the queue depth for the array devices:

1. Make a backup of the /etc/system file: `cp /etc/system /etc/system.old`
2. Edit the /etc/system file.
3. Add the following to the /etc/system file
set sd:sd_max_throttle = x (for **x** see [Figure 3-1](#))
For Sun generic HBA: **set ssd:ssd_max_throttle = x**
4. Save your changes, and exit the text editor.
5. Shutdown and reboot to apply the Max Throttle setting.

```
:
*   To set a variable named 'debug' in the module named 'test_module'
*
*       set test_module:debug = 0x13

set sd:sd_max_throttle = 8
set ssd:ssd_max_throttle = 8
```

Add this line to /etc/system
Add this line to /etc/system
(for Sun generic Host Bus Adapter)

Figure 3-1: Setting the Max Throttle (Queue Depth)

Configuring Fibre Channel on the Host

After you prepare the array hardware and software and Fibre Channel HBA(s), connect the array to the Sun system. The array comes with all the hardware and cabling required for connection to the host system(s).

Table 3-4 summarizes the steps for connecting the array to the Sun host.

Table 3-4: Steps for Connecting to a Solaris Host

Step	Activity	Description
1.	Verify the array installation.	Confirm that the status of the Fibre Channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the Solaris system.	<ul style="list-style-type: none">• Shut down the Solaris system.• When shutdown is complete, power off the Solaris display.• Power off all peripheral devices, except for the array.• Power off the host system.• You are now ready to connect the array.
3.	Connect the array.	<ul style="list-style-type: none">• Install Fibre Channel cables between the array and the Solaris system.• Follow all precautions and procedures in your array's user's guide.• Check all specifications to ensure proper installation and configuration (see Table 3-1 on page 3-3).
4.	Power on the Solaris system.	<ul style="list-style-type: none">• Power on the Solaris system display.• Power on all peripheral devices. The array should be on, the Fibre Channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the Solaris system is powered on, restart the system to have the new devices recognized.• Confirm the ready status of all peripheral devices, including the array.• Power on the Solaris system.
5.	Boot the Solaris system.	When the OK prompt appears, boot the system using the boot -r command to rebuild the devices. Using boot by itself will not build the devices on the newly installed array.

Verifying and Discovering LUNs

This section describes how to prepare the host servers for connection to the array.

Setting and Discovering LUNs

After the array is installed and connected, halt and reboot the system to have the LUs recognized following the reboot.

To set and recognize LUs:

1. Halt the Sun system:

halt

2. Reboot the Sun system:

boot -r

3. Log in to the system as root and verify that the system recognizes the array (see [Figure 3-2](#)):

dmesg | more

4. Verify that the vendor name, product name, and number of blocks match the values in [Figure 3-2](#).

```
# dmesg

Jan  8 18:50:24 Solaris9 genunix: [ID 540533 kern.notice] ^MSunOS Release 5.9
Version Generic_117171-07 64-bit
Jan  8 18:50:24 Solaris9 genunix: [ID 943905 kern.notice] Copyright 1983-2003 Sun
Microsystems, Inc. All rights reserved.

Jan  8 19:33:22 Solaris9 qlc: [ID 630585 kern.info] NOTICE: Qlogic qlc(0): Loop
ONLINE
Jan  8 19:33:23 Solaris9 scsi: [ID 799468 kern.info] ssd0 at fp0: name
w50060e8010200497,3, bus address ef
Jan  8 19:33:23 Solaris9 genunix: [ID 936769 kern.info] ssd0 is /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,3
Jan  8 19:33:23 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,3 (ssd0):
Jan  8 19:33:23 Solaris9 Corrupt label; wrong magic number
Jan  8 19:33:23 Solaris9 scsi: [ID 799468 kern.info] ssd1 at fp0: name
w50060e8010200497,2, bus address ef
Jan  8 19:33:23 Solaris9 genunix: [ID 936769 kern.info] ssd1 is /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,2
Jan  8 19:33:23 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,2 (ssd1):
Jan  8 19:33:23 Solaris9 Corrupt label; wrong magic number
Jan  8 19:33:23 Solaris9 scsi: [ID 799468 kern.info] ssd2 at fp0: name
w50060e8010200497,1, bus address ef
Jan  8 19:33:23 Solaris9 genunix: [ID 936769 kern.info] ssd2 is /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,1
Jan  8 19:33:23 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,1 (ssd2):
Jan  8 19:33:23 Solaris9 Corrupt label; wrong magic number
Jan  8 19:33:23 Solaris9 scsi: [ID 799468 kern.info] ssd3 at fp0: name
w50060e8010200497,0, bus address ef
Jan  8 19:33:23 Solaris9 genunix: [ID 936769 kern.info] ssd3 is /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,0
Jan  8 19:33:23 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@le,600000/
SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,0 (ssd3):
Jan  8 19:33:23 Solaris9 Corrupt label; wrong magic number
```

Figure 3-2: Recognizing Devices in Solaris

Recognizing LUNs

After system start-up, log in as root and use the **dmesg | more** command to verify that the Sun system recognizes the array. Confirm that the displayed vendor names, product names, and number of blocks match the values in [Figure 3-3](#). If the results are different than the intended system configuration, the Navigator 2 LUN configuration or fiber cabling might be wrong.



NOTE: When the Sun system accesses the multiplatform devices, the message “Request sense couldn't get sense data” may appear. If this message appears, disregard it.

```
# dmesg | more
:
Jan  8 19:41:41 Solaris9 qlc: [ID 139792 kern.info] Qlogic qlc(0) FCA Driver v20070212-2.19
Jan  8 19:41:41 Solaris9 qlc: [ID 657001 kern.info] Qlogic qlc(0) WWPN=210000e08b121f37 :
WWNN=200000e08b121f37
Jan  8 19:41:57 Solaris9 qlc: [ID 630585 kern.info] NOTICE: Qlogic qlc(0): Loop OFFLINE
Jan  8 19:41:57 Solaris9 qlc: [ID 694252 kern.info] NOTICE: qlc(0): Firmware version 3.3.24

Jan  8 19:43:02 Solaris9 scsi: [ID 799468 kern.info] ssd0 at fp0: name w50060e8010200497,3, bus
address ef
Jan  8 19:43:02 Solaris9 genunix: [ID 936769 kern.info] ssd0 is /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,3
Jan  8 19:43:02 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,3 (ssd0):
Jan  8 19:43:02 Solaris9 Corrupt label; wrong magic number
Jan  8 19:43:02 Solaris9 scsi: [ID 799468 kern.info] ssd1 at fp0: name w50060e8010200497,2, bus
address ef
Jan  8 19:43:02 Solaris9 genunix: [ID 936769 kern.info] ssd1 is /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,2
Jan  8 19:43:02 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,2 (ssd1):
Jan  8 19:43:02 Solaris9 Corrupt label; wrong magic number
Jan  8 19:43:02 Solaris9 scsi: [ID 799468 kern.info] ssd2 at fp0: name w50060e8010200497,1, bus
address ef
Jan  8 19:43:02 Solaris9 genunix: [ID 936769 kern.info] ssd2 is /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,1
Jan  8 19:43:02 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,1 (ssd2):
Jan  8 19:43:02 Solaris9 Corrupt label; wrong magic number
Jan  8 19:43:02 Solaris9 scsi: [ID 799468 kern.info] ssd3 at fp0: name w50060e8010200497,0, bus
address ef
Jan  8 19:43:02 Solaris9 genunix: [ID 936769 kern.info] ssd3 is /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,0
Jan  8 19:43:02 Solaris9 scsi: [ID 107833 kern.warning] WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/
ssd@w50060e8010200497,0 (ssd3):
Jan  8 19:43:02 Solaris9 Corrupt label; wrong magic number
```

Figure 3-3: Verifying New Devices

Partitioning and Labeling the New Devices

After the Sun System recognizes the new devices, partition and label the devices. All new devices, including all SCSI disk devices, must be partitioned and labeled using the Sun format utility (the WARNING below). Each SCSI disk device can have more than one partition.

Disk partitioning and labeling procedure involves the following steps:

1. Defining and setting the disk type.
2. Setting the partition(s).
3. Labeling the disk (required for devices to be managed by HDLM).
4. Verifying the disk label.

You enter this information into the Sun system when performing the disk partitioning and labeling procedure.



WARNING! Be extremely careful when using the Sun format utility. Do not use any format commands not described in this document. The Sun format utility is designed for Sun disks. Some format commands are not compatible with the array and can overwrite the data on the disk. The array will not respond to the format command (disks are formatted using a notebook PC), and will not report any defect data in response to the defect command.

To partition and label the devices/disks on the newly installed array:

1. Enter format at the root prompt to start the format utility ([Figure 3-4 on page 3-11](#)).
 - a. Verify that all new devices are displayed. If not, exit the format utility (**quit** or Ctrl d), and then be sure the SCSI/fibre-to-LDEV paths were defined for all devices and that all new devices were added to the driver configuration file).
 - b. Write the character-type device file names (e.g., c1t2d0) for all of the new array devices. You will need this information later to create the file systems.
2. When prompted to specify the disk, enter the number from the list for the device to be partitioned and labeled.
3. When prompted to label the disk, perform one of the following steps:
 - Enter **y** for “yes” and enter the desired label. Devices that will be managed by HDLM require a label.
 - If you are sure that the device will not need a label, you can enter **n** for “no.”
4. When the format menu appears, enter **type** to display the disk types. The disk types are listed (vendor name + product name, e.g., HITACHI DF600F).
5. If the disk type for the selected device is already defined, enter the number for that disk type and skip to step 7.
6. If the disk type for the selected device is not already defined, enter the number for other to define a new disk type.

7. Enter the disk type parameters for the selected device using the data provided above. Be sure to enter the parameters exactly as shown in [Figure 3-5 on page 3-12](#).
8. When prompted to label the disk, enter **n** for "no."
9. When the format menu appears, enter **partition** to display the partition menu.
10. Enter the desired partition number and the partition parameters in [Figure 3-6 on page 3-13](#).
11. At the **partition>** prompt, enter **print** to display the current partition table.
12. Repeat steps 9 and 10 as needed to set the desired partitions for the selected device.
13. After setting the partitions for the selected device, enter **label** at the partition> prompt and enter **y** to label the device ([Figure 3-7 on page 3-14](#)).
14. Enter **quit** to exit the **partition** utility and return to the format utility.
15. At the **format>** prompt, enter **disk** to display the available disks. Be sure the disk you just labeled is displayed with the proper disk type name and parameters.
16. Repeat steps 2 through 15 for each new device to be partitioned and labeled. After a array disk type is defined (e.g., HITACHI DF600F), you can label all devices of that same type without having to enter the parameters (skipping step 6). For this reason, you may want to label the devices until all new devices have been partitioned and labeled.
17. When you finish partitioning and labeling the disks and verifying the disk labels, exit the **format** utility by entering **quit** or Ctrl-d.

```

# format Start format utility
Searching for disks...done

c4t0d0: configured with capacity of 975.00MB These devices are not yet labeled
c4t0d1: configured with capacity of 1.96GB
c4t0d2: configured with capacity of 2.95GB
c4t0d3: configured with capacity of 3.96GB

AVAILABLE DISK SELECTIONS:
  0. clt1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@1c,600000/scsi@2/sd@1,0
  1. c4t0d0 <HITACHI-DF600F-0000 cyl 52 alt 2 hd 50 sec 768> D Not yet labeled:
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,0
  2. c4t0d1 <HITACHI-DF600F-0000 cyl 107 alt 2 hd 50 sec 768> D Not yet labeled:
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,1
  3. c4t0d2 <HITACHI-DF600F-0000 cyl 161 alt 2 hd 50 sec 768> D Not yet labeled:
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,2
  4. c4t0d3 <HITACHI-DF600F-0000 cyl 216 alt 2 hd 50 sec 768> D Not yet labeled:
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,3
Specify disk (enter its number): Jan  8 19:48:22 Solaris9 Corrupt label;
wrong magic number
Jan  8 19:48:22 Solaris9 scsi: WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w500
60e8010200497,1 (ssd2):
Jan  8 19:48:22 Solaris9 Corrupt label; wrong magic number
Jan  8 19:48:22 Solaris9 scsi: WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w500
60e8010200497,1 (ssd2):
Jan  8 19:48:22 Solaris9 Corrupt label; wrong magic number
Jan  8 19:48:22 Solaris9 scsi: WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w500
60e8010200497,0 (ssd3):
Jan  8 19:48:22 Solaris9 Corrupt label; wrong magic number
Jan  8 19:48:22 Solaris9 scsi: WARNING: /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w500
60e8010200497,0 (ssd3):
Jan  8 19:48:22 Solaris9 Corrupt label; wrong magic number
1
selecting c4t0d0
[disk formatted]
Disk not labeled. Label it now? n Enter "n" for no

#

```

Figure 3-4: Verifying New Devices for Disk Partitioning

```

FORMAT MENU:
    disk      - select a disk
    type      - select (define) a disk type
    partition - select (define) a partition table
    current   - describe the current disk
    format    - format and analyze the disk
    repair    - repair a defective sector
    label     - write label to the disk
    analyze   - surface analysis
    defect    - defect list management
    backup    - search for backup labels
    verify    - read and display labels
    save      - save new disk/partition definitions
    inquiry   - show vendor, product and revision
    volume    - set 8-character volume name
    quit

# format> type Enter type
:
AVAILABLE DRIVE TYPES
    0. Auto configure
?
    22. HITACHI-DF600F-0000 Do not select this disk type
    23. other
Specify disk type (enter its number)[19]: 23 Enter number for "other" to define
Enter number of data cylinders: 52 Enter value from Figure 3-4("cyl")
Enter number of alternate cylinders[2]: 2 Enter value from Figure 3-4("alt")
Enter number of physical cylinders[54]: 54 (press Enter for default)
Enter number of heads: 50 Enter value from Figure 3-4("hd")
Enter physical number of heads[default]:
Enter number of data sectors/track: 768 Enter value from Figure 3-4("sec")
Enter number of physical sectors/track[default]: (press Enter for default)
Enter rpm of drive[3600]: 10000 Enter value for SAS drive
                        7200 Enter value for SATA drive
Enter format time[default]: (press Enter for default)
Enter cylinder skew[default]: (press Enter for default)
Enter track skew[default]: (press Enter for default)
Enter tracks per zone[default]: (press Enter for default)
Enter alternate tracks[default]: (press Enter for default)
Enter alternate sectors[default]: (press Enter for default)
Enter cache control[default]: (press Enter for default)
Enter prefetch threshold[default]: (press Enter for default)
Enter minimum prefetch[default]: (press Enter for default)
Enter maximum prefetch[default]: (press Enter for default)
Enter disk type name (remember quotes): "HITACHI DF600F" Enter name
selecting c4t0d0
[disk formatted]
Disk not labeled. Label it now? n Enter "n" for no
format>

```

Figure 3-5: Defining and Setting the Disk Type


```

format> disk

AVAILABLE DISK SELECTIONS:
  0. c1t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@1c,600000/scsi@2/sd@1,0
  1. c4t0d0 <HITACHI-DF600F-0000 cyl 52 alt 2 hd 50 sec 768>
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,0
  2. c4t0d1 <HITACHI-DF600F-0000 cyl 107 alt 2 hd 50 sec 768>
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,1
  3. c4t0d2 <HITACHI-DF600F-0000 cyl 161 alt 2 hd 50 sec 768>
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,2
  4. c4t0d3 <HITACHI-DF600F-0000 cyl 216 alt 2 hd 50 sec 768>
    /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,3
Specify disk (enter its number)[2]: 3

FORMAT MENU:
  disk      - select a disk
  type      - select (define) a disk type
  partition - select (define) a partition table
  current   - describe the current disk
  format    - format and analyze the disk
  repair    - repair a defective sector
  label     - write label to the disk
  analyze   - surface analysis
  defect    - defect list management
  backup    - search for backup labels
  verify    - read and display labels
  save      - save new disk/partition definitions
  inquiry   - show vendor, product and revision
  volname   - set 8-character volume name
  !<cmd>    - execute <cmd>, then return
  quit

format> type Enter type
?
AVAILABLE DRIVE TYPES:
  0. Auto configure
  ?
    22. HITACHI-DF600F-0000
    23. HITACHI DF600F
    ???24. other
Specify disk type (enter its number)[19]: 23 Enter the desired lun size
selecting c4t0d1
[disk formatted]
Disk not labeled. Label it now? n Enter n for no
format>

```

Figure 3-6: Setting the Partition(s) (continues on the next page)

```

FORMAT MENU:
    disk      - select a disk
    type      - select (define) a disk type
    partition - select (define) a partition table
    current   - describe the current disk
    format    - format and analyze the disk
    repair    - repair a defective sector
    label     - write label to the disk
    analyze   - surface analysis
    defect    - defect list management
    backup    - search for backup labels
    verify    - read and display labels
    save      - save new disk/partition definitions
    inquiry   - show vendor, product and revision
    volname   - set 8-character volume name
    <cmd>     - execute <cmd>, then return
    quit

format> partition                                Display partition menu

PARTITION MENU
    0      - change '0' partition
    1      - change '1' partition
    2      - change '2' partition
    3      - change '3' partition
    4      - change '4' partition
    5      - change '5' partition
    6      - change '6' partition
    7      - change '7' partition
    select - select a predefined table
    modify - modify a predefined partition table
    name   - name the current table
    print  - display the current table
    label  - write partition map and label to the disk
    quit

partition> 0                                       Select partition number

Part      Tag      Flag      Cylinders      Size      Blocks
  0        root      wm         0 - 1        37.50MB    (2/0/0)    76800

Enter partition id tag [root]:                       Press enter for default
Enter partition permission flags [wm]:                Press enter for default
Enter new starting cyl [0]:                          Press enter for default
Enter partition size [76800b, 2c, 1e, 37.50mb, 0.04gb]: 51c Enter size (Figure 3-2)
partition> print                                    Display partition table
:
Current partition table (unnamed):
Total disk cylinders available: 52 + 2 (reserved cylinders)

Part      Tag      Flag      Cylinders      Size      Blocks
  0        root      wm         0 - 50        956.25MB    (51/0/0) 1958400
  1        swap      wu          2 - 5         75.00MB     (4/0/0)  153600
  2        backup     wu         0 - 51        975.00MB    (52/0/0) 1996800
  3 unassigned      wm          0              0           (0/0/0)    0
  4 unassigned      wm          0              0           (0/0/0)    0
  5 unassigned      wm          0              0           (0/0/0)    0
  6         usr      wm         6 - 51        862.50MB    (46/0/0) 1766400
  7 unassigned      wm          0              0           (0/0/0)    0

```

Figure 3-7: Setting the Partition(s) (continued)

```

PARTITION MENU
  0      - change '0' partition
  1      - change '1' partition
  2      - change '2' partition
  3      - change '3' partition
  4      - change '4' partition
  5      - change '5' partition
  6      - change '6' partition
  7      - change '7' partition
select - select a predefined table
modify - modify a predefined partition table
name    - name the current table
print   - display the current table
label   - write partition map and label to the disk
quit

partition> label                                Label the disk
Ready to label disk, continue? Y                Enter Y for yes
*
partition> quit                                    Return to format
format> disk                                       Display disks

AVAILABLE DISK SELECTIONS
  0. clt1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
      /pci@1c,600000/scsi@2/sd@1,0
  1. c4t0d0 <HITACHI DF600F cyl 52 alt 2 hd 50 sec 768>
>      Verify disk label

                                     Track size
                                     Number of heads
                                     Number of alternate cylinders
                                     Number of data cylinders
                                     Disk type name
      /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,0
  2. c4t0d1 <HITACHI-DF600F-0000 cyl 107 alt 2 hd 50 sec 768> Not yet labeled
      /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,1
  3. c4t0d2 <HITACHI-DF600F-0000 cyl 161 alt 2 hd 50 sec 768> Not yet labeled
      /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,2
  4. c4t0d3 <HITACHI-DF600F-0000 cyl 216 alt 2 hd 50 sec 768> Not yet labeled
      /pci@1e,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e8010200497,3
Specify disk (enter its number): 2      Enter number for next disk to label,
                                     or press Ctrl-d to quit

```

Figure 3-8: Labeling the Disk and Verifying the Disk Label

VMware

This chapter discusses guidelines on how to prepare a VMware host server for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Preparing the Host Server](#)
- ❑ [Configuring the Fibre Channel HBAs](#)
- ❑ [Setting the Queue Depth Parameter](#)
- ❑ [Upgrading the Firmware Online](#)
- ❑ [Configuring Host Groups for VMware](#)
- ❑ [Creating a VMFS and Attaching a Raw Device](#)

Preparing the Host Server

Table 4-1 lists guidelines and tasks you need to follow to prepare the host server for connection to the array.

Table 4-1: Host Server Preparation Guidelines

Item	Task
HBAs	One or more supported host bus adapters (HBAs) with the latest supported BIOS and driver are required. Verify that the HBAs, drivers, and BIOSes are the latest supported versions by the array, and are functioning properly. <ul style="list-style-type: none">To check the latest supported versions, refer to the Hitachi interoperability matrix: http://www.hds.com/products/interoperability/.For information about HBAs supported by your operating system, refer to the VMware Web site: http://www.vmware.com/.
Install the HBA and in the host server.	For installation information, check the documentation and Web site of your HBA vendor. Be sure the HBA is supported by your array (refer to the Hitachi interoperability matrix).
VMware ESX Server Operating System	Verify that planned OS version, architecture, relevant patches, and maintenance levels are supported by HDS. Refer to the Hitachi Data Systems interoperability matrix for information on supported versions.

Configuring the Fibre Channel HBAs

After connecting the array and rebooting the VMware server, configure the Fibre Channel HBAs connected to the array. For information about configuring your HBA(s), refer to the:

- Documentation that came with your HBA
- HBA vendor's Web site
- Hitachi interoperability matrix.
- Server Configuration Guide on the VMware Web site:

http://www.vmware.com/pdf/vi3_301_201_server_config.pdf

The user documentation for your HBA should describe whether other options are required to meet your operational requirements.

After configuring the HBA, you may have to reset the VMware server to have the change take effect before restarting the VMware server.

Setting the Queue Depth Parameter

You may need to change the queue depth value on the server. If the number is small, I/O performance can deteriorate. The array reports a queue full status when the queue depth exceeds an allowable limit. The system may not operate correctly when the queue is full and a large value is set. Set an appropriate number according to your configuration. If necessary, set a queue depth number for each server. Refer to the documentation for your HBA before setting a value.

Guidelines for settings:

- 32 commands per LUN
- 512 commands per port
- 30 or more for device timeout value on the array LU.

The following sections provide general guidelines for setting queue depth on QLogic and Emulex HBAs. For detailed information, please refer to the documentation for your HBA and the HBA vendor's Web site.

Setting the Queue Depth Parameter on QLogic HBAs

To set the queue depth parameter on QLogic HBAs:

1. Issue the following command to identify the HBA's driver name:

```
# vmkload_mode -l | grep qla
```

Sample output: qla2300_707_vmw

2. Issue the following command lines, with the following considerations:
 - Substitute the <driver_name> parameter with the name output in step 1.
 - Substitute the "nn" parameter with the queue depth value calculated above.
 - The last line reboots the server.

```
# esxcfg-module -s "ql2xmaxqdepth=nn" <driver_name>
```

```
# esxcfg-boot -b
```

```
# reboot
```

Setting the Queue Depth Parameter on Emulex HBAs

To set the queue depth parameter on Emulex HBAs:

1. Issue the following command to identify the HBA's driver name:

```
# vmkload_mode -l | grep lpfcdd
```

Sample output: lpfcdd_7xx
2. Issue one of the following command lines, with the following considerations:
 - Substitute the <driver_name> parameter below with the name output in step 1.
 - Substitute the "nn" parameter with the queue depth value calculated above.
 - a. If your server has one Emulex HBA, issue the following command:

```
# esxcfg-module -s "lpfc0_lun_queue_depth=nn"  
  <driver_name>
```
 - b. If your server has two Emulex HBAs, issue the following command:

```
# esxcfg-module -s "lpfc0_lun_queue_depth=nn  
  lpfc1_lun_queue_depth=nn" <driver_name>
```
3. Issue the following command lines. The last line reboots the server.

```
# esxcfg-boot -b  
# reboot
```

Upgrading the Firmware Online

I/O may be suspended for up to 30 seconds before and after the firmware upgrade procedure.

Configuring Host Groups for VMware

This section describes how to use Storage Navigator Modular 2 (Navigator 2) to configure host groups on an Hitachi array for VMware. As part of this procedure, you will review the RAID groups defined for a selected array.

This procedure assumes that the Navigator 2 software has been installed. If it is not installed, refer to your array's user's guide for installation instructions. Then return to this chapter.

Logging in to Navigator 2

To log in to Navigator 2, use the following procedure.

1. Start your browser and enter the following Uniform Resource Locator in the browser's address bar:

<http://<IP address>:23015/StorageNavigatorModular/Login>

where <IP address> is the IP address of the Navigator 2 host.

2. At the login window, enter the user ID and password (default user ID is **system** and default case-sensitive password is **manager**). For security, each typed password character appears as a bullet (•).

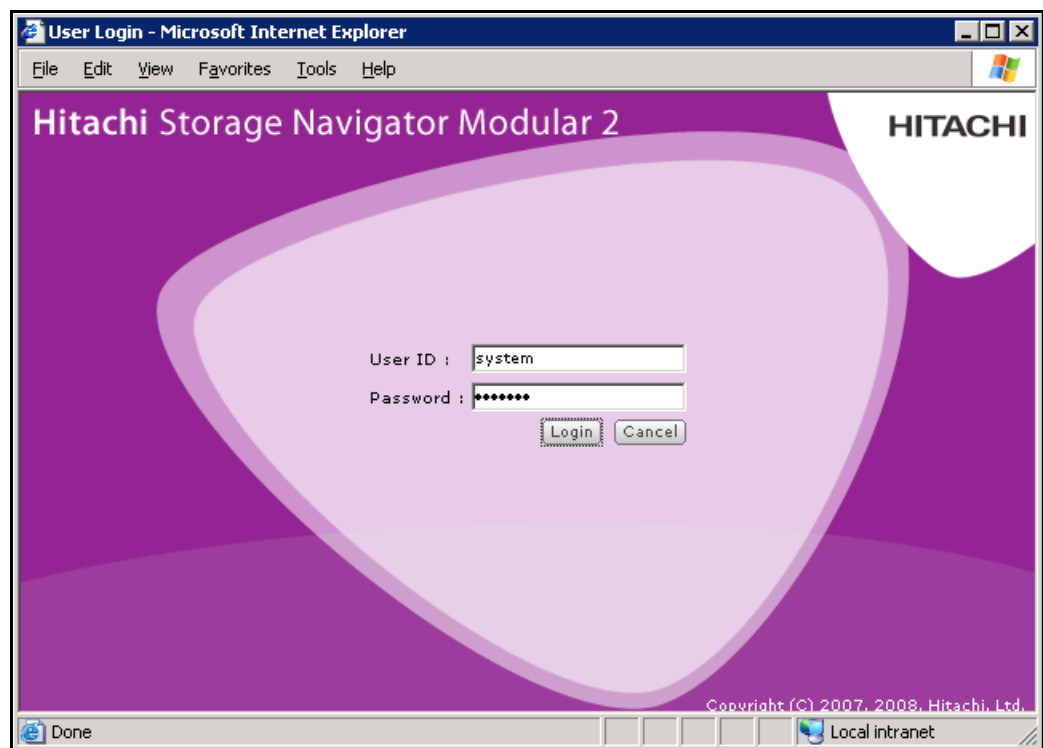


Figure 4-1: Navigator 2 Login Window

3. Click the **Login** button. An Arrays List window similar to the one in [Figure 4-2 on page 4-6](#) appears.

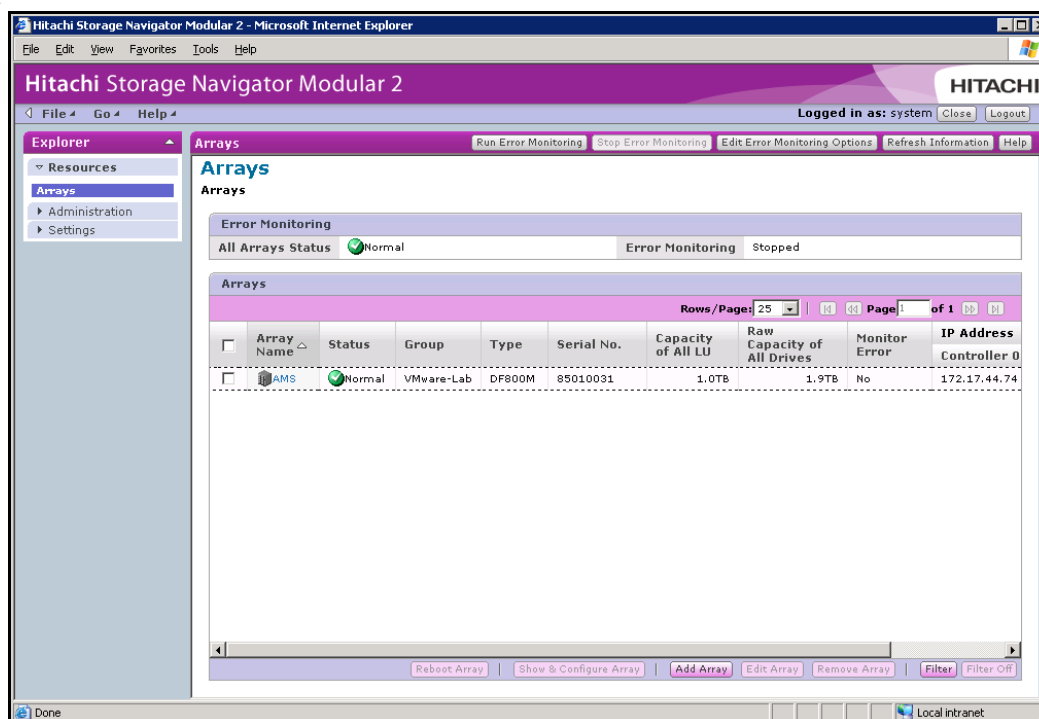


Figure 4-2: Example of an Arrays List Window

4. In the left column, check the checkbox next to the array you want to configure.
5. Click **Show & Configure Array**. An Array Properties window similar to the one in [Figure 4-3](#) appears for the selected array.




Figure 4-3: Example of an Array Properties Window

Reviewing RAID Groups

Use the following procedure to review the RAID groups associated with the array you selected.

1. In the **Arrays** pane in the center of the Array Properties window, click **Groups > RAID Groups**. A RAID Groups screen similar to the one in [Figure 4-4](#) appears, listing the available RAID groups on the selected array.

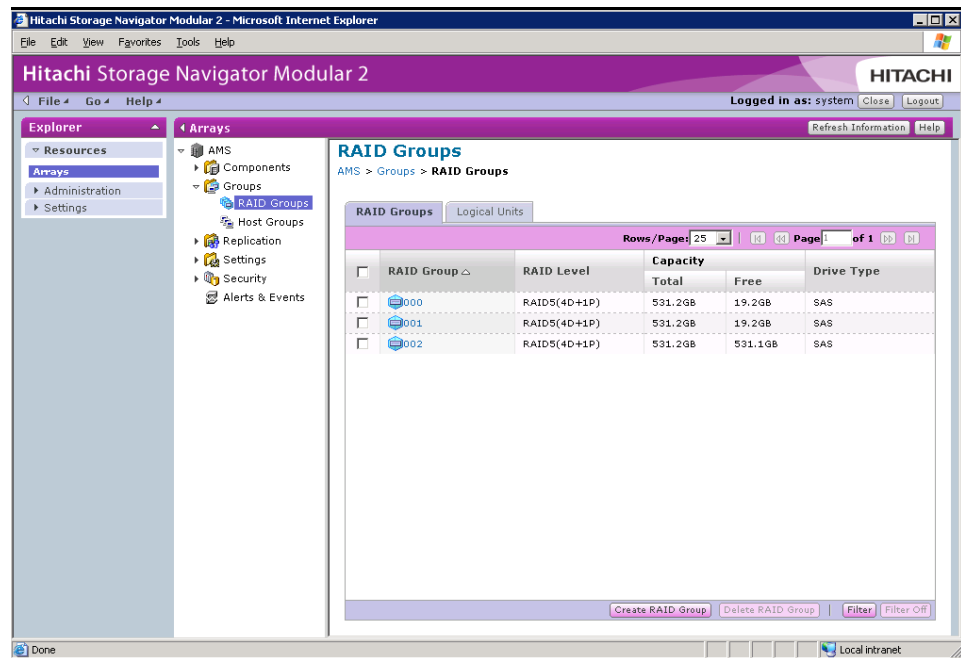


Figure 4-4: Example of RAID Groups Screen

2. Review the RAID groups, if any, that have been created for the array. If necessary, use the **Create RAID Group** button to create more RAID groups. For more information, refer to the Navigator 2 online help.
3. Click the **Logical Units** tab ([Figure 4-5 on page 4-8](#)). Then review the logical units, if any, that have been created. If necessary, use the **Create LU** button to create more logical units. For more information, refer to the Navigator 2 online help.

2. Create or edit a host group for VMware.

- To create a host group, click **Create Host Group**. A Create Host Group page appears (Figure 4-7).
- To edit an existing host group, check the checkbox in the left column for the host group you want to edit and click **Edit Host Group**. An Edit Host Group page appears. This page is similar to the Create Host Group page, but shows the host group settings defined in an earlier Navigator 2 session.

The screenshot shows the 'Create Host Group' page in the HSNM2 web interface. The page is titled 'Create Host Group' and has a 'Host Group Property' section. In this section, there are fields for 'Host Group No.' (001), 'Name' (ESX0_0_HBA0), and 'Options' (Platform: not specified, Middleware: not specified). There is also an 'Available Ports' section with a list of ports (0A, 0B, 0C, 0D) and a 'Port' dropdown menu. Below this, there are two tables: 'Detected WWNs' and 'Selected WWNs'. The 'Detected WWNs' table is empty. The 'Selected WWNs' table has one row with columns 'Nickname', 'Port Name', and 'Port'. The 'Nickname' is 'ESX0_0_HBA0', the 'Port Name' is '210000E08B817360', and the 'Port' is '0A'. There are 'Add' and 'Remove' buttons between the tables. At the bottom, there are 'OK' and 'Cancel' buttons.

Figure 4-7: Example of Create Host Group Page

3. To create or edit the host group information shown:
- a. In the **Host Group No** box, enter or modify the host group number.
 - b. In the **Name** box, enter or modify the host group name.



TIP: While host groups can be created with more than one WWN, best practices dictate that you create one host group for each HBA and name the host group the same as the nickname for the HBA. While either of these methods is supported, creating one host group per HBA is more flexible and the only supported configuration when booting hosts from a SAN.

- c. Under **Available Ports**, check the available port(s) on which the host group will be created.

4. In the **WWNs** tab at the bottom of the page, confirm that **Select From List** is selected. Then select a World Wide Name (WWN) in the **Detected WWNs** box on the left and click **Add >** to move the selected WWN to the **Selected WWNs** box on the right. To add more WWNs to the **Selected WWNs** box, repeat this step.

NOTE: You can also use the **Enter WWNs Manually** selection to add WWNs manually by specifying a World Wide Port Name (WWPN). For more information, refer to the Navigator 2 online help.

Create Host Group - Microsoft Internet Explorer

HSNM2

HITACHI

Create Host Group

Host Group Property

Enter the information for the host group to be created.

* Host Group No.: 001
From 1 to 127 .

* Create to:

Available Ports

* Name: ESK0_0_HBA0
32 characters or less (alphanumeric characters, ' ', '@', '\$', '%', '&', '!', '+', '-', '.', ':', ';', '=', '<div>Options: Platform: not specified
Middleware: not specified

WWNs

Logical Units

Options

Check H-LUN and available LUN, and push add button when assign logical unit to host group.
Check selected LUN and push remove button when disassign logical unit from host group.

LU Mapping:

H-LUNs

Rows / Pages: 100 of 21

H-LUN
<input type="radio"/> 0048
<input type="radio"/> 0049
<input checked="" type="radio"/> 0050
<input type="radio"/> 0051
<input type="radio"/> 0052
<input type="radio"/> 0053

Available Logical Units

Rows / Pages: 25 of 1

LUN	Capacity	RAID Group	RAID Lev
<input checked="" type="checkbox"/> 0301 50.0MB 002 RAID5(4D			
<input checked="" type="checkbox"/> 0302 50.0MB 002 RAID5(4D			

Assigned Logical Units

Rows / Pages: 25 of 1

H-LUN	LUN	Capacity	RAID Group
<input type="checkbox"/> 0001 0001 256.0GB 000			
<input type="checkbox"/> 0002 0002 256.0GB 000			
<input type="checkbox"/> 0010 0101 256.0GB 001			
<input type="checkbox"/> 0011 0102 256.0GB 001			

* Required field

OK Cancel

Done

Local intranet

6. In the **H-LUNs** table, select a host LUN. Then, in the **Available Logical Units** table below it, check one or more logical units you want to assign to this host group. Click the **Add >** button to move the selected logical unit(s) to the **Assigned Logical Units** table on the right.

7. Click the **Options** tab. A tab similar to the one in [Figure 4-9](#) appears. The list on the left shows the available (selectable) LUNs that have been created.

Create Host Group

Enter the information for the host group to be created.

* Host Group No.: 001
From 1 to 127.

* Name: ESX0_0_HBA0
32 characters or less (alphanumeric characters, !, #, \$, %, &, ' , * , + , - , . , : , ; , < , = , > , [, \ ,] , ^ , _ , ` , { , | , } , ~ , or space).

Options:
Platform: VMware
Middleware: not specified

* Create to: Available Ports
☐ Port
☒ 0A
☐ 0B
☐ 0C
☐ 0D

WWNs Logical Units **Options**

Select or check options for host group.
Mode setting is cleared and set again automatically when platform or middleware are settled.

Platform: VMware Middleware: not specified

Mode Settings:

Common Setting: Standard Mode

Additional Setting:

Enable HP-UX Mode:	<input type="checkbox"/> Yes
Enable PSUE Read Reject Mode:	<input type="checkbox"/> Yes
Enable Mode Parameters Changed Notification Mode:	<input type="checkbox"/> Yes
Enable NACA Mode:	<input type="checkbox"/> Yes
Enable Task Management Isolation Mode:	<input type="checkbox"/> Yes
Enable Unique Reserve Mode 1:	<input type="checkbox"/> Yes
Enable Port-ID Conversion Mode:	<input type="checkbox"/> Yes
Enable Tru Cluster Mode:	<input type="checkbox"/> Yes
Enable Product Serial Response Mode:	<input type="checkbox"/> Yes
Enable Same Node Name Mode:	<input type="checkbox"/> Yes
Enable CCHS Mode:	<input type="checkbox"/> Yes
Enable Inquiry Serial Number Conversion Mode:	<input type="checkbox"/> Yes

* Required field

OK Cancel

Figure 4-9: Example of the Options Tab

8. Click the **Platform** drop-down list and select **VMware** ([Figure 4-9](#)). Keep **Middleware** set to **not specified**. In addition, for VMware, standard mode settings are used, eliminating the need to check any **Additional Settings**.
9. Click **OK**.
10. If desired, use the **WWNs** tab in the Host Groups screen to view the WWPNS mapped to host groups ([Figure 4-10 on page 4-12](#)).

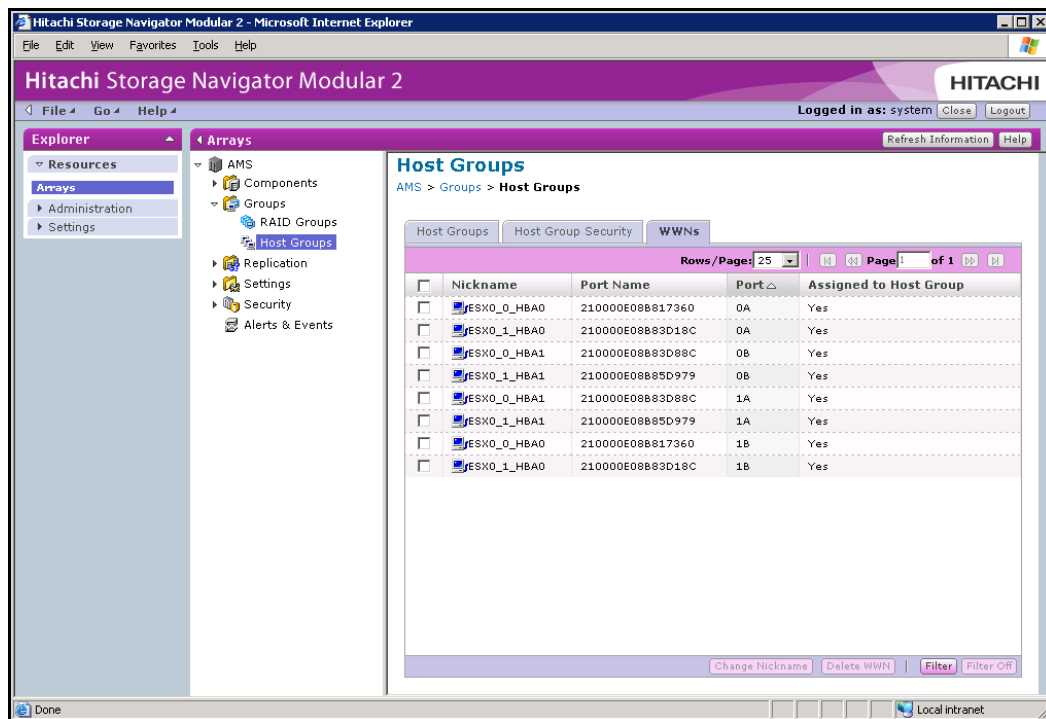


Figure 4-10: Example of WWPNS Mapped to Host Groups

This completes the procedure for configuring host groups for VMware.

Creating a VMFS and Attaching a Raw Device

This section describes how to create a virtual machine file system (VMFS) and attach a raw device.

Creating a Virtual Machine File System

The following procedure describes how to create a VMFS.

1. Log on to your VMware ESX host or to your virtual center using VMware Infrastructure Client. The VMware Infrastructure page appears.
2. With the VMware Infrastructure page displayed:
 - a. In the left pane, select a VMware host.
 - b. Click the **Configuration** tab in the VMware Infrastructure page.
 - c. Under **Hardware** on the left side of the page, click **Storage Adapters**. A page similar to the one in [Figure 4-11 on page 4-13](#) appears.



NOTE: If your **Configuration** tab is not completely populated, do not be concerned. Performing the next step will populate the tab with the LUNs in your system.

3. Click **Rescan** in the top-right area above **Storage Adapters** to scan the Fibre Channel HBAs in your system for LUNs.

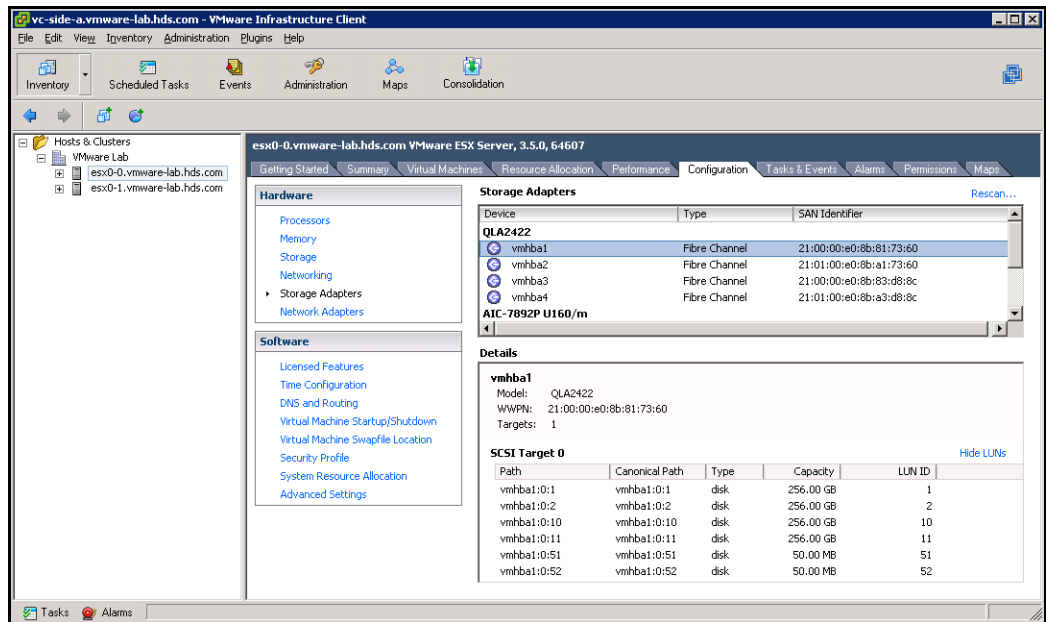


Figure 4-11: Example of Initial VMware Infrastructure Client Page

4. Select a Fibre Channel HBA under **Storage Adapters**. Then right-click a LUN in the lower part of the screen (under **SCSI Target 0** in the figure above) and click **Properties**. A LUN Properties dialog box similar to the one in [Figure 4-12](#) appears.

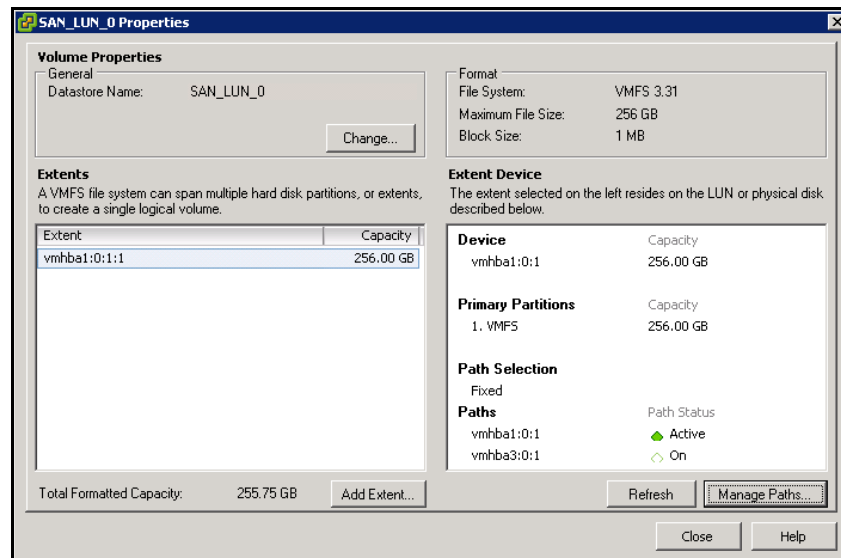


Figure 4-12: Example of LUN Properties Dialog Box

5. In the LUN Properties dialog box, click **Manage Paths**. A Manage Paths dialog box similar to the one in [Figure 4-13 on page 4-14](#) appears. This dialog box shows the paths defined for the selected LUN. In the example, below, one path is active while the other path is on.




Figure 4-13: Example of Manage Paths Dialog Box

6. In the **Policy** area, click the **Change** button. The Manage Paths - Selection Policy dialog box appears.
7. The default policy is **Most Recently Used**. The recommended path policy is **Fixed** as shown in Figure 4-14.



NOTE: Please apply VMware ESX 3.5 Patch ESX350-200804401-BG, which also sets the default path policy to **Fixed** for your array.

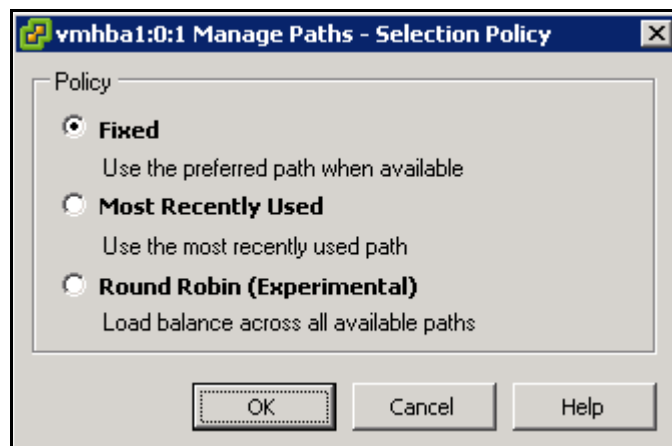


Figure 4-14: Manage Paths - Selection Policy Dialog Box

8. Click **OK** to exit the Manage Paths - Selection Policy dialog box.
9. In the Manage Paths dialog box, click the **Change** button in the lower right side of the dialog box. The Change Path State dialog box appears (Figure 4-15 on page 4-15).

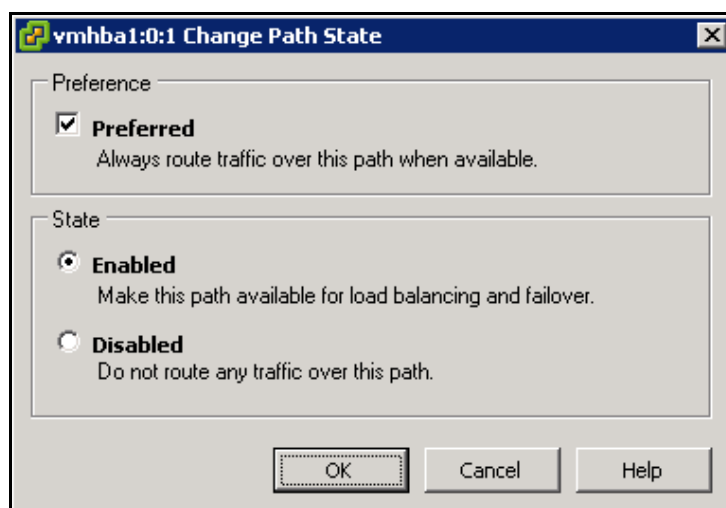


Figure 4-15: Change Path State Dialog Box

10. Under **Preference**, confirm that **Preferred** is checked. Under **State**, confirm that **Enabled** is selected. If these are not configured this way, do so now.
11. Click **OK** to exit the Change Path State dialog box.
12. Click **OK** to exit the Managed Paths dialog box.
13. Click **Close** to exit the LUN Properties dialog box.
14. From the VMware Infrastructure page (Figure 4-16), select a VMware host in the left pane and click **Storage** under **Hardware**.

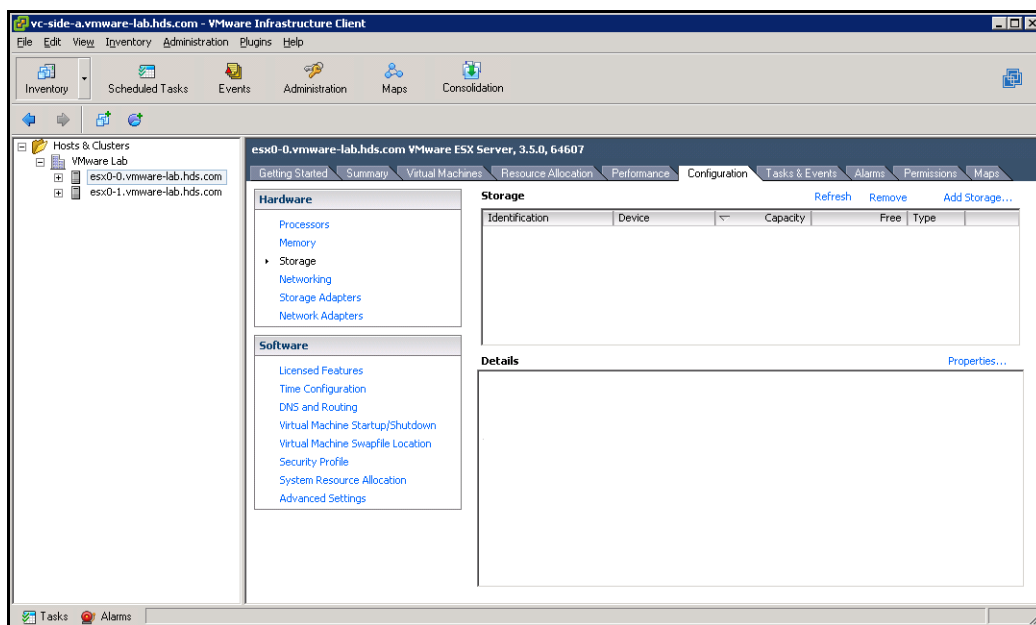


Figure 4-16: Initial VMware Infrastructure Client Page with Storage Selected

15. Click **Add Storage** in the top right the top-right area above **Storage**. The Add Storage Wizard starts, with the Select Storage Type screen displayed (Figure 4-17 on page 4-16).

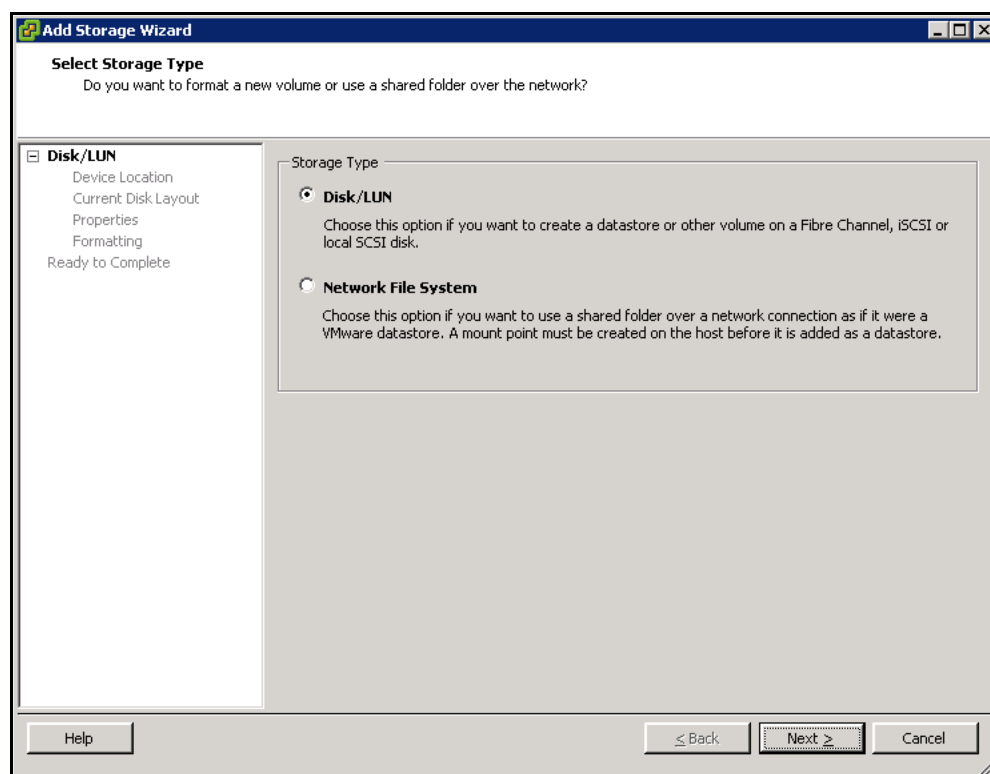


Figure 4-17: Select Storage Type Screen (Add Storage Wizard)

16. Under **Storage Type**, select **Disk/LUN** (if it is not already selected, then click **Next**. The Select Disk /LUN screen appears (Figure 4-18).

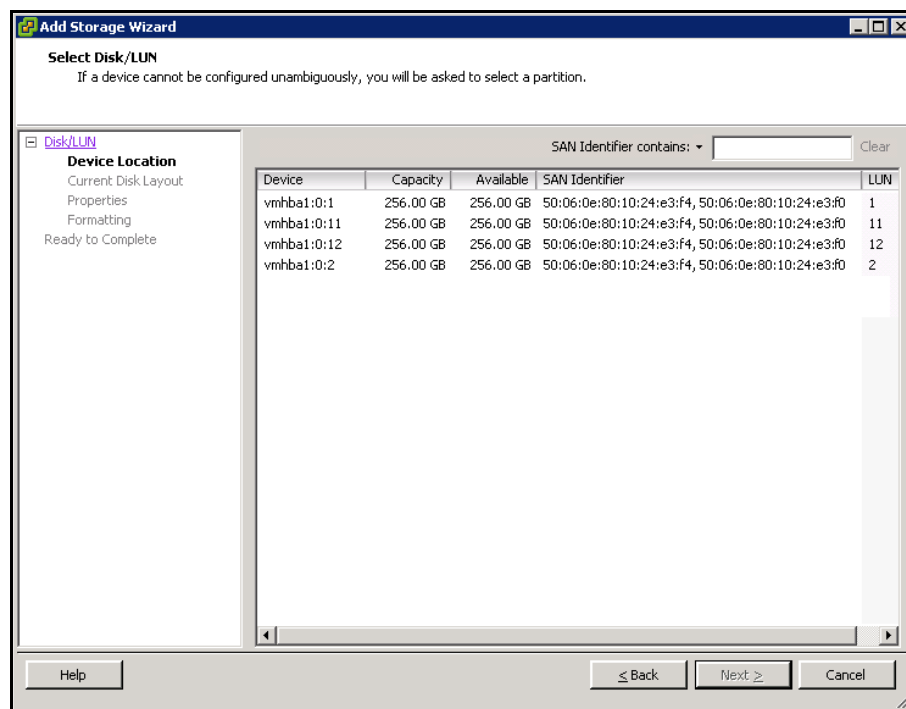


Figure 4-18: Select Disk/LUN Screen (Add Storage Wizard)

17. Select the LUN on which you want to create the VMFS and click **Next**.
The Current Disk Layout screen appears (Figure 4-19).




Figure 4-19: Current Disk Layout Screen (Add Storage Wizard)

18. Click **Next**. The Disk/LUN Properties screen appears (Figure 4-20).

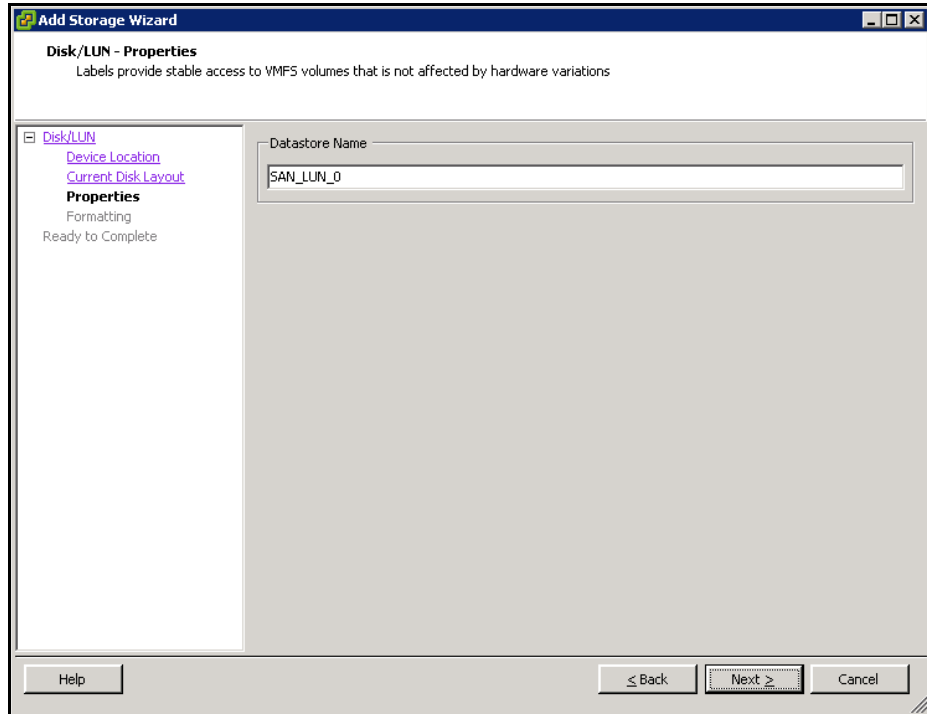


Figure 4-20: Disk/LUN Properties Screen (Add Storage Wizard)

19. Under **Datastore Name**, enter a name for the VMFS volume and click **Next**. The Disk LUN Formatting screen in [Figure 4-21](#) appears.

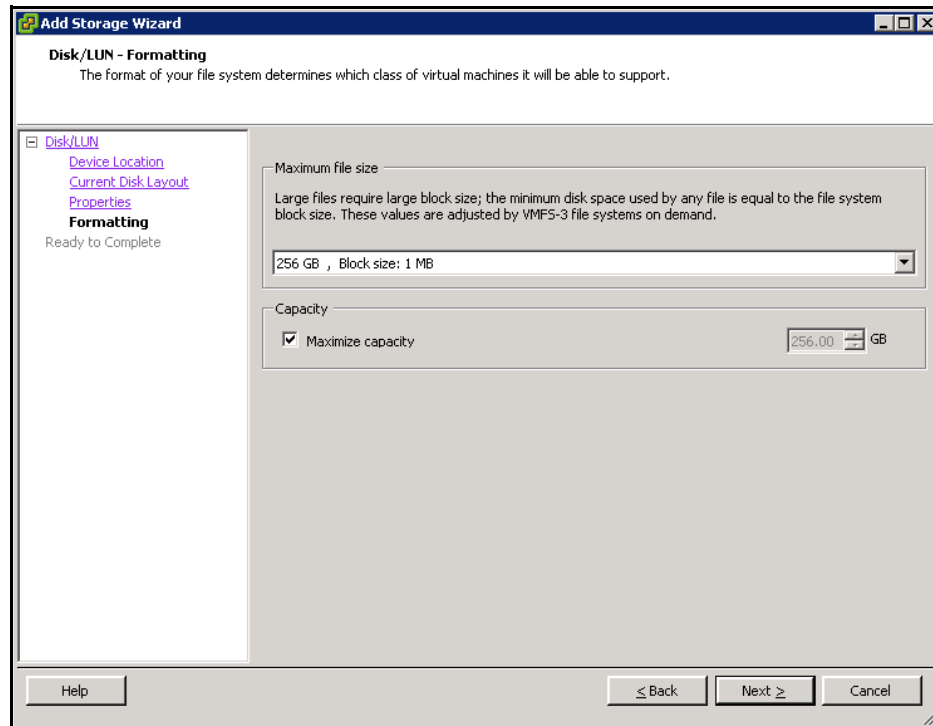


Figure 4-21: Disk/LUN Formatting Screen (Add Storage Wizard)

20. Accept the default selection in the Disk LUN Formatting screen and confirm that **Maximize capacity** under **Capacity** is checked.
21. Click **Next**. The Ready to Complete screen appears ([Figure 4-22 on page 4-19](#)).

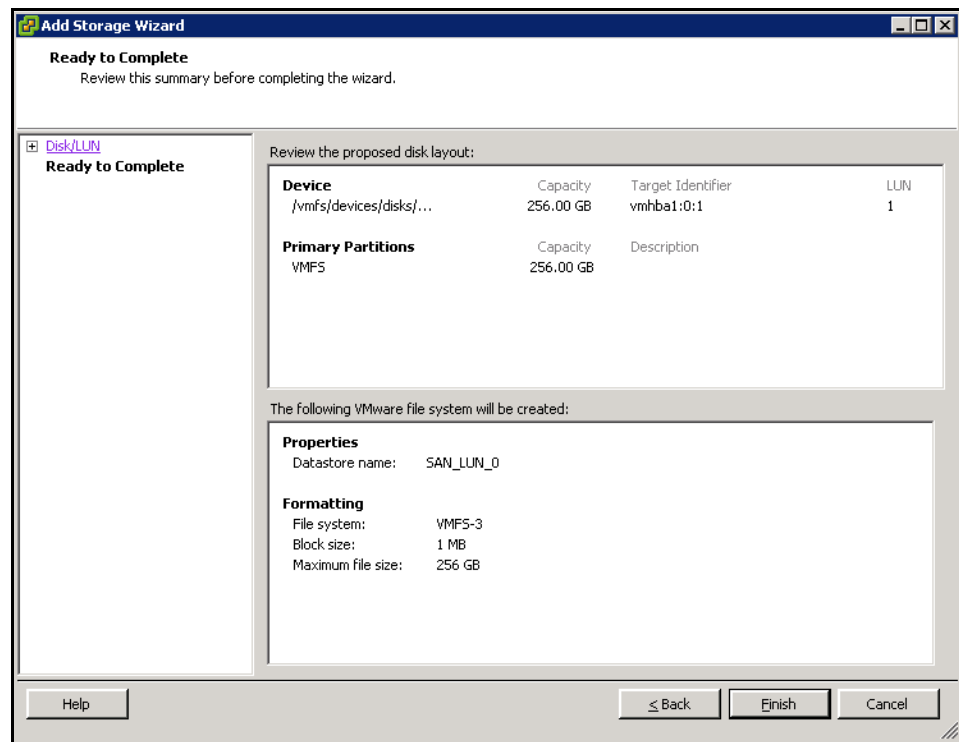


Figure 4-22: Ready to Complete Screen (Add Storage Wizard)

22. Review the summary information displayed. If you need to change a setting, click the **Back** button to return to the appropriate screen, change the setting, and click **Next** until you return to the Ready to Complete screen.
23. After confirming that the settings are correct, click **Finish**. The VMware Infrastructure page appears, with your LUN/VMFS settings displayed (Figure 4-23).

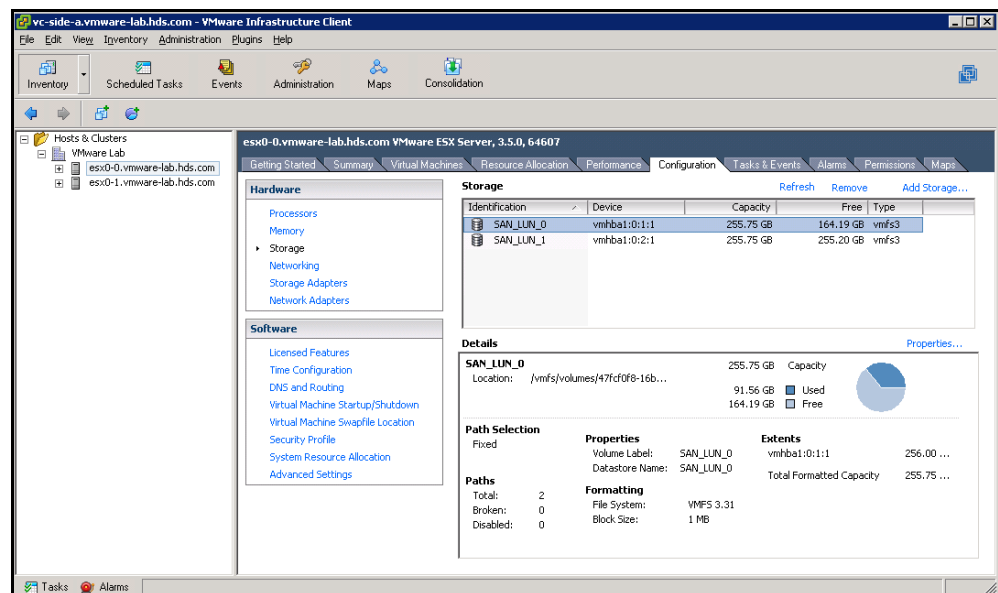


Figure 4-23: Example of VMware Infrastructure Page

24. If needed, repeat steps 15 through 23 to create additional VMFS's.



TIP: On any additional hosts that will have access to the LUNs, you do not have to create the VMFS's again. Instead, on the other host, click **Refresh** above **Storage** on the VMware Infrastructure page ([Figure 4-23](#)).

This completes the procedure for creating a VMFS. Proceed to the section below for instructions on attaching a raw device.

Attaching a Raw Device

The following procedure describes how to attach a raw device, such as a command device, to a virtual machine.

1. From the VirtualCenter client, right-click a virtual machine and select **Edit Settings**. The Virtual Machine Properties dialog box appears ([Figure 4-24](#)).

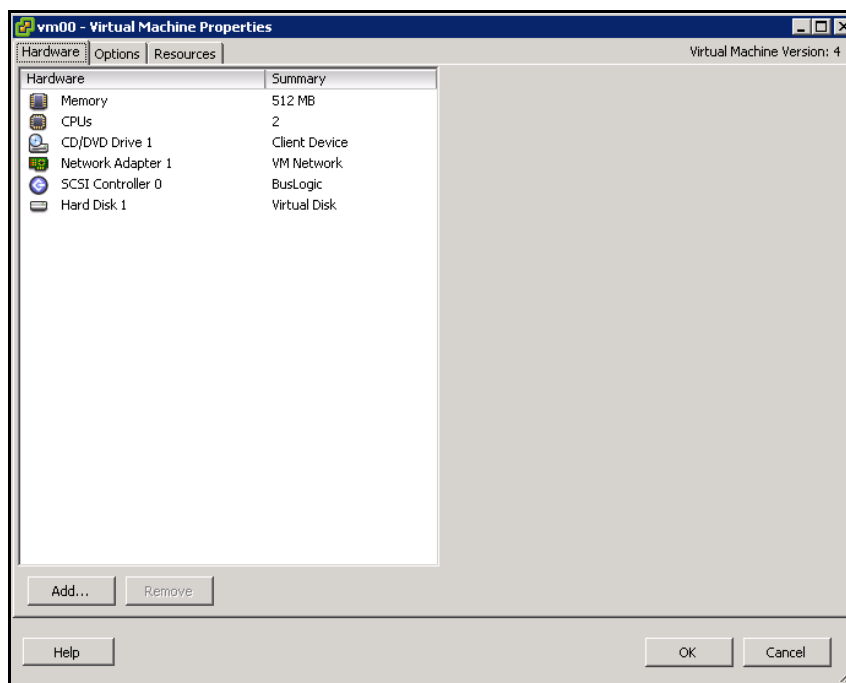


Figure 4-24: Virtual Machine Properties Dialog Box

2. Select **Add**. The Add Hardware Device wizard starts, with the Select Device Type screen displayed ([Figure 4-25 on page 4-21](#)).

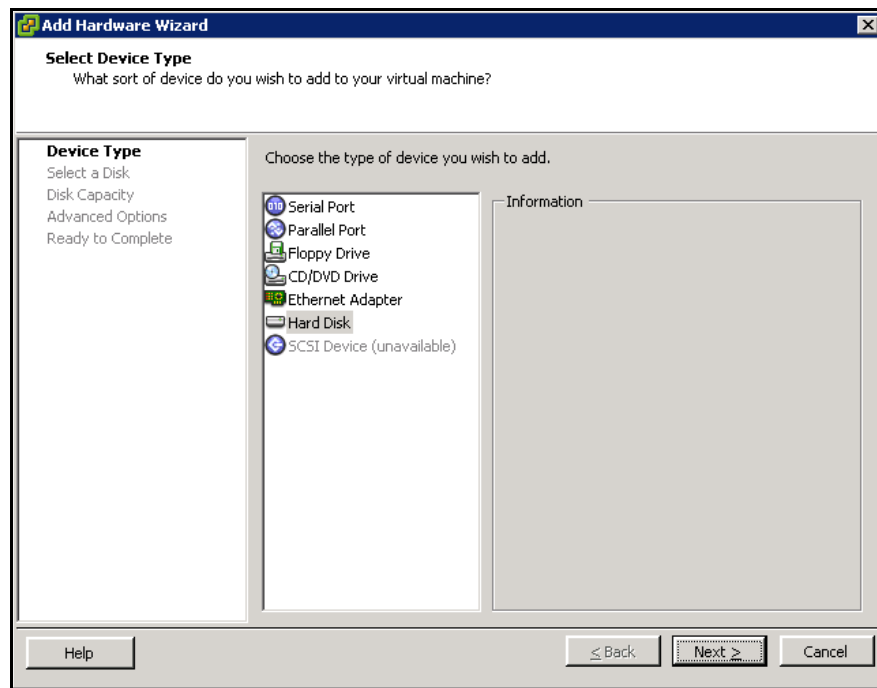


Figure 4-25: Select Device Type Screen (Add Hardware Wizard)

3. Select **Hard Disk** and click **Next**. The Select a Disk Screen appears.
4. Select **Raw Device Mappings**, as shown in [Figure 4-26](#).

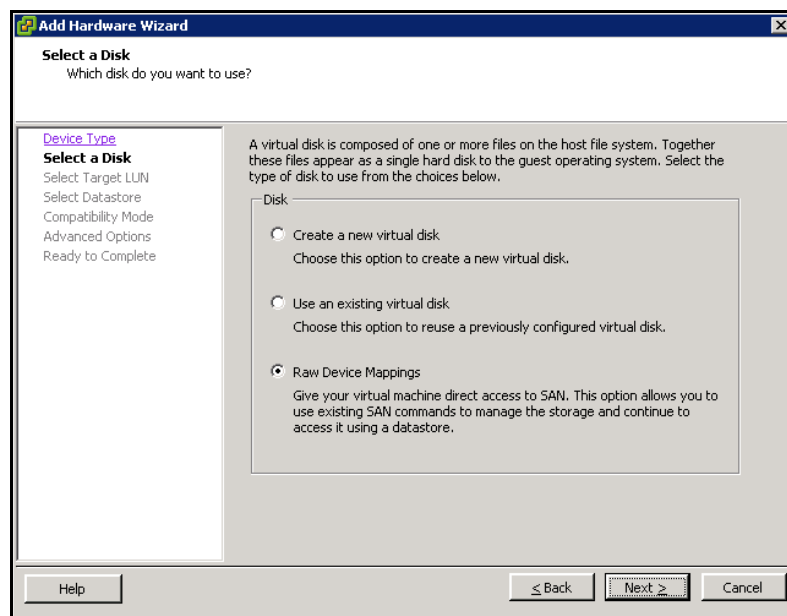


Figure 4-26: Select a Disk Screen (Add Hardware Wizard)

5. Click **Next**. The Select and Configure a Raw LUN screen appears ([Figure 4-27 on page 4-22](#)).

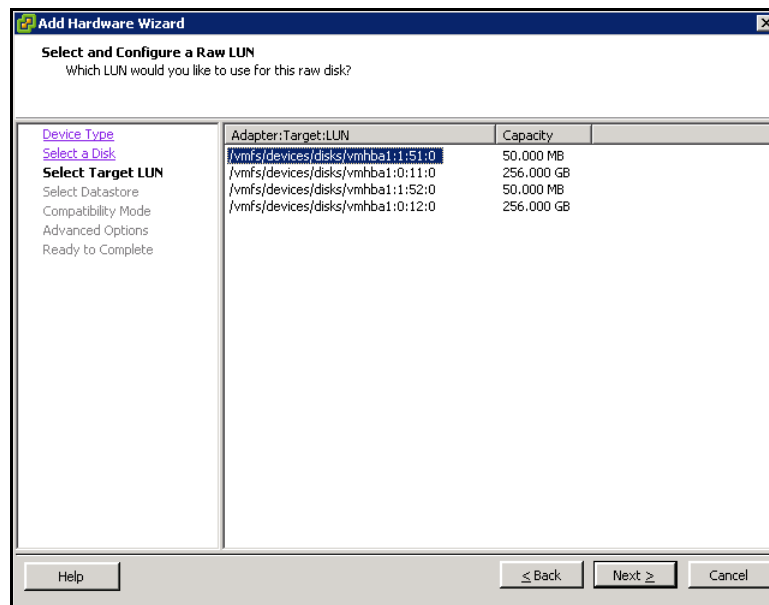


Figure 4-27: Select and Configure a Raw LUN Screen (Add Hardware Wizard)

6. Select the LUN you want to use for the raw device (Command Device) and click **Next**. The Select a Datastore screen appears (Figure 4-28).




Figure 4-28: Select a Datastore Screen (Add Hardware Wizard)

7. Confirm that **Store with Virtual Machine** is selected. If it is not, select it.
8. Click **Next**. The Select Compatibility Mode screen appears (Figure 4-29 on page 4-23).

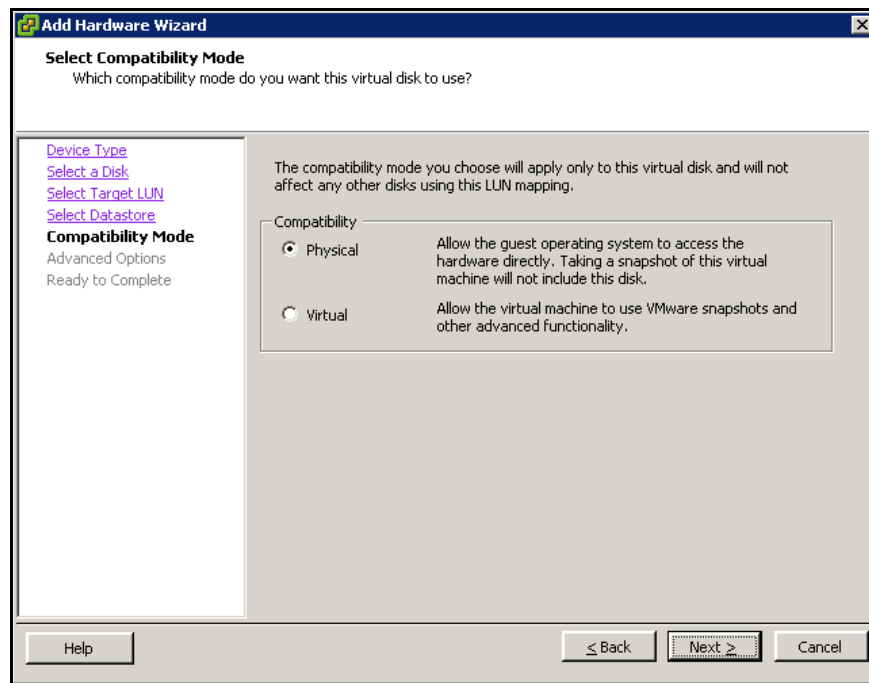


Figure 4-29: Select Compatibility Mode (Add Hardware Wizard)

9. Under **Compatibility**, confirm that **Physical** is selected. If it is not selected, select it.
10. Click **Next**. The Specify Advanced Options screen appears (Figure 4-30).

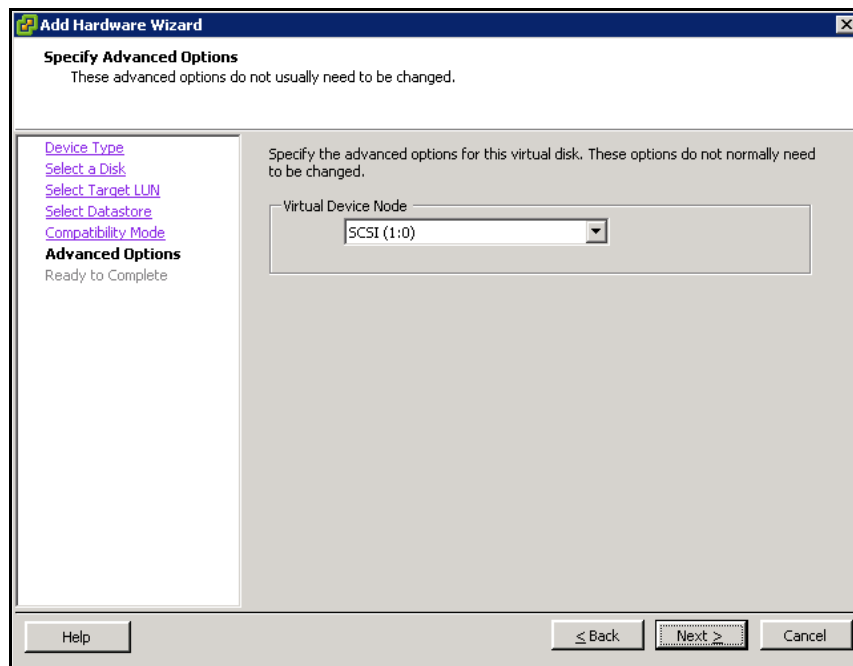


Figure 4-30: Specify Advanced Options Screen (Add Hardware Wizard)

11. Under **Virtual Device Node**, select a virtual device node local to the Virtual Machine.

12. Click **Next**. The Ready to Complete Screen appears (Figure 4-31).

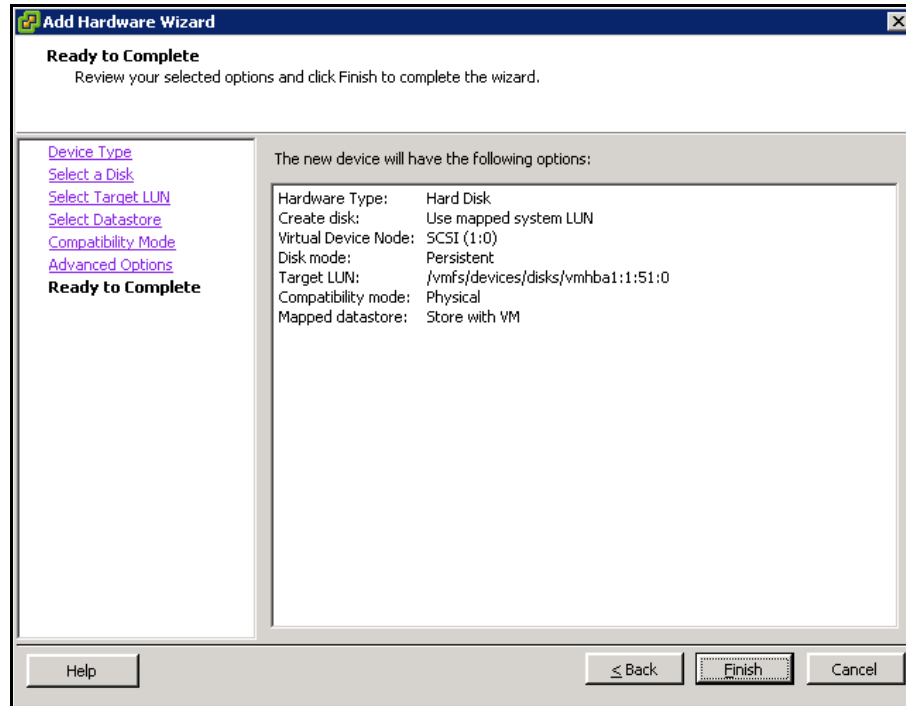


Figure 4-31: Ready to Complete Screen (Add Hardware Wizard)

13. Review the summary information displayed. If you need to change a setting, click the **Back** button to return to the appropriate screen, change the setting, and click **Next** until you return to the Ready to Complete screen.
14. After confirming that the settings are correct, click **Finish**. The Virtual Machine Properties dialog box appears (Figure 4-32).

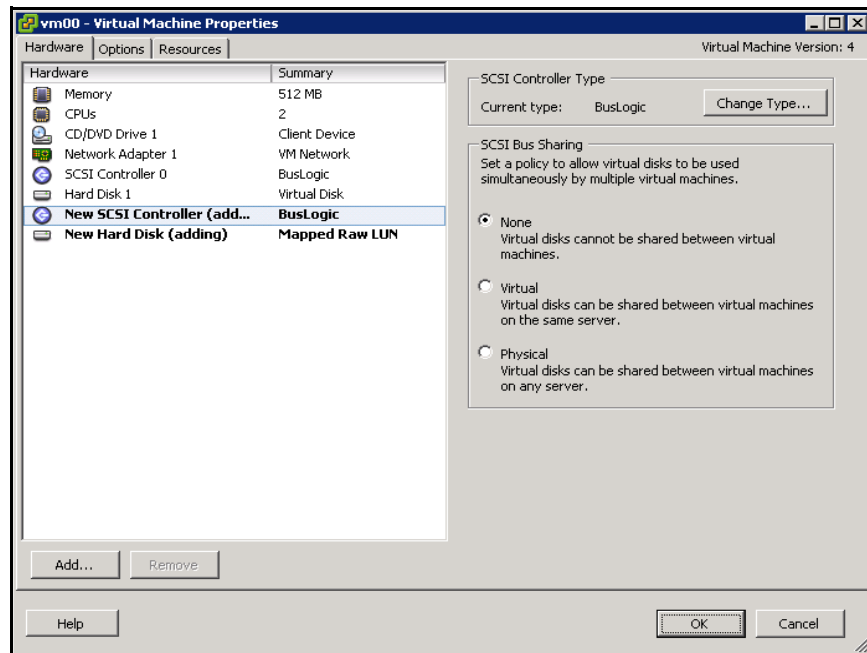


Figure 4-32: Virtual Machine Properties Dialog Box

15. Click **OK** to exit the Virtual Machine Properties dialog box.
16. Right-click the same virtual machine and select **Edit Settings**. The Virtual Machine Properties dialog box appears, with the new SCSI controller and hard disk you added (Figure 4-33).

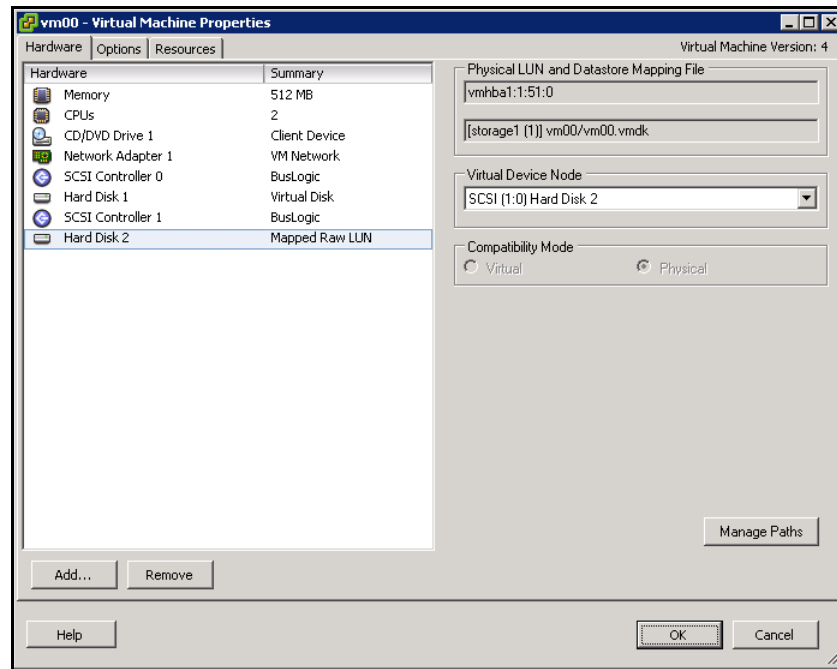


Figure 4-33: Example of Virtual Properties Dialog Box with New SCSI Controller and Hard Disk Added

17. Click **Cancel** to exit the Virtual Properties dialog box.
- This completes the procedure for adding a raw device.

Red Hat Enterprise Linux

This chapter discusses guidelines on how to prepare a Red Hat Enterprise Linux host server for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Connecting to a Red Hat Enterprise Linux Host](#)
- ❑ [Configuring the Host Fibre Channel HBAs](#)
- ❑ [Verifying New Device Recognition](#)
- ❑ [High Availability \(HA\) Cluster Configurations](#)
- ❑ [HA Multipath Configurations](#)
- ❑ [Verifying and Discovering LUNs](#)

Connecting to a Red Hat Enterprise Linux Host

After you prepare array hardware and software and Fibre Channel HBA(s), connect the array to the Red Hat Enterprise Linux system. The array comes with all the hardware and cabling required for connection to the host system(s).

Table 5-1 lists the tasks to follow to connect to the host server. To prepare the server, you need one or more supported HBAs with the latest supported BIOS and driver. Verify the HBAs, drivers, and BIOSes are the latest supported versions by Hitachi Data Systems and are functioning properly.

- To check the latest supported versions, refer to the Hitachi interoperability matrix at <http://www.hds.com/products/interoperability/>.
- For information about your HBAs, refer to the HBA vendor's Web site.
- For information about HBAs supported by your operating system, refer to the HP Web site: <http://www.redhat.com>.

Table 5-1: Steps for Connecting to a Red Hat Enterprise Linux Host

Step	Activity	Description
1.	Verify the array installation.	Confirm that the status of the Fibre Channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the Red Hat Enterprise Linux system.	<ul style="list-style-type: none">• Shut down the Red Hat Enterprise Linux system.• When shutdown is complete, power off the Red Hat Enterprise Linux display.• Power off all peripheral devices, except for the array.• Power off the host system.• You are now ready to connect the array.
3.	Connect the array.	<ul style="list-style-type: none">• Install Fibre Channel cables between the array and the Red Hat Enterprise Linux system.• Follow all precautions and procedures in your array's user's guide.• Check all specifications to ensure proper installation and configuration.
4.	Power on and boot the Red Hat Enterprise Linux system.	<ul style="list-style-type: none">• Power on the Red Hat Enterprise Linux system display.• Power on all peripheral devices. The array should be on, the Fibre Channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the Red Hat Enterprise Linux system is powered on, restart the system to have the new devices recognized.• Confirm the ready status of all peripheral devices, including the array.• Power on the Red Hat Enterprise Linux system.• Boot the Red Hat Enterprise Linux system.

Configuring the Host Fibre Channel HBAs

Configure the Fibre Channel HBA(s) connected to the array. HBAs have many configuration options. This section provides the minimum requirements for configuring host Fibre Channel HBAs for operation with the array. For more information, refer to the documentation for your Fibre Channel HBAs.

- Queue depth requirements for array devices are specified in [Table 5-2](#).
- You may need to disable the BIOS to prevent the system from trying to boot from the array.
- Several other parameters (e.g., Fibre Channel and/or fabric options) may also need to be set. Please refer to the documentation for your HBA to determine whether other options are required to meet your operational requirements.
- Use the same settings and device parameters for all array devices.

Table 5-2: Queue Depth Requirements for the Array

Parameter	Required Value for Array
Queue depth per LU	≤ 32 ports
Queue depth per port (MAXTAGS)	≤ 512 per port

To optimize I/O performance, you can adjust the queue depth for the array devices later as needed, so long as you keep within the specified range.

Verifying New Device Recognition

The final step before configuring the new array disk devices is to verify that the host system recognizes the new devices. The host system automatically creates a device file for each new device recognized.

To verify new device recognition:

1. Use the **dmesg** command to display the devices (Figure 5-1).
2. Record the device file name for each new array device. You will need this information when you partition the devices.
3. The device files are created under the `/dev` directory. Verify that a device file was created for each new array disk device (Figure 5-2).

```
# dmesg | more
:
:
scsil : ioc1: LSI53C1030, FwRev=01032700h, Ports=1, MaxQ=255, IRQ=209
      Emulex LightPulse Fibre Channel SCSI driver 8.1.10.3
      Copyright(c) 2004-2006 Emulex. All rights reserved.

scsi : 1 host.
Vendor: HITACHI      Model: DF600F      Rev: 0000
Type:   Direct-Access      ANSI SCSI revision: 04
sd 0:0:0:0: Attached scsi disk sda
      ^      scsi ID0/channel 0/ID 0/lun 0      Device file name = /dev/sda

Vendor: HITACHI      Model: DF600F      Rev: 0000
Type:   Direct-Access      ANSI SCSI revision: 04
sd 0:0:0:1: Attached scsi disk sdb
:
:
```

Figure 5-1: Example of Verifying New Device Recognition



NOTE: In Figure 5-1, the HITACHI DF600F device (TID 0, LUN 0) and the HITACHI DF600F device (TID 0, LUN 1) are recognized by the Red Hat Enterprise Linux server.

```
# ls -l /dev | more
:
brw-rw---- 1 root  disk  8, 0 Oct 30 2007 sda      Device file name = sda
```

Figure 5-2: Example of Verifying Device Files

High Availability (HA) Cluster Configurations

Hitachi Modular storage systems are compatible with various cluster software applications. For more information about:

- Compatible cluster software applications, refer to the Hitachi Data Systems interoperability matrix at <http://www.hds.com/products/interoperability/>.
- Installing and configuring the cluster software, refer to the documentation provided by the cluster software vendor.

HA Multipath Configurations

Hitachi Modular storage systems support various HA multipathing software products for the Red Hat Enterprise Linux operating system. Refer to the Hitachi Data Systems interoperability matrix at <http://www.hds.com/products/interoperability/> for currently supported HA software applications. Then consult the documentation provided by the HA multipathing vendor for information about installing, configuring, operating, and best practices when using the software with Active/Active storage systems like the Hitachi Modular storage systems.

If the Hitachi Data Systems interoperability matrix show that Red Hat Enterprise Linux's bundled multipathing software, referred to as "Device Mapper," is supported for the intended Red Hat operating system version / update level and you want to use this bundled Multipath solution, refer to the appropriate Red Hat Device Mapper documentation for proper installation, configuration, and operation. Some Device Mapper release level documentation can be obtained at the following link:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.4/html/DM_Multipath/queueifnopath_issues.html

To ensure Active/Active I/O activity by the host Linux I/O to the Hitachi Modular storage system, confirm that the following minimum parameters are set in the file `/etc/multipath.conf`:

- Vendor: Hitachi
- Product: DF600F
- Path_grouping_policy: Multibus

Install and configure the multipathing software on the server before connecting the server to the Hitachi Modular storage system.

Verifying and Discovering LUNs

This section describes how to configure the new array disk devices that you attached to the Red Hat Enterprise Linux system host in the previous sections.

Partitioning the Devices

After the setting the number of logical units, use the following procedure to set the partitions on the array disk devices:

1. Enter **fdisk/dev/<device_name>**
Example: **fdisk/dev/sda**
where **dev/sda** is the device file name
2. Select **p** to display the present partitions.
3. Select **n** to make a new partition. You can make up to four primary partitions (1-4) or one extended partition. The extended partition can be organized into 11 logical partitions, which can be assigned partition numbers from 5 to 15.
4. Select **w** to write the partition information to disk and complete the fdisk command.



NOTE: Other commands that you might want to use include **d** to remove partitions or **q** to stop a change.

5. Repeat steps 1 through 4 for each new array disk device.

SuSE Linux Enterprise Server

This chapter discusses guidelines on how to prepare a SuSE Linux Enterprise (SLES) server for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Connecting to an SLES Host](#)
- ❑ [Configuring the Host Fibre Channel HBAs](#)
- ❑ [Verifying New Devices](#)
- ❑ [Verifying and Discovering LUNs](#)
- ❑ [High Availability \(HA\) Cluster Configurations](#)
- ❑ [HA Multipath Configurations](#)

Connecting to an SLES Host

After you prepare the array hardware and software and fibre-channel HBA(s), connect the array to the SLES system. The array comes with all the hardware and cabling required for connection to the host system(s).

Table 6-1 lists the tasks to follow to prepare the host server. To prepare the server, you need one or more supported HBAs with the latest supported BIOS and driver. Verify the HBAs, drivers, and BIOSes are the latest supported versions by Hitachi Data Systems and are functioning properly.

- To check the latest supported versions, refer to the Hitachi interoperability matrix at <http://www.hds.com/products/interoperability/>.
- For information about your HBAs, refer to the HBA vendor's Web site.
- For information about HBAs supported by your operating system, refer to the Novell Web site: <http://www.novell.com/linux/>.

Table 6-1: Steps for Connecting to an SLES Host

Step	Activity	Description
1.	Verify the array installation.	Confirm that the status of the Fibre Channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the SLES system.	<ul style="list-style-type: none">• Shut down the SLES system.• When shutdown is complete, power off the SLES display.• Power off all peripheral devices, except for the array.• Power off the host system.• You are now ready to connect the array.
3.	Connect the array.	<ul style="list-style-type: none">• Install Fibre Channel cables between the array and the SLES system.• Follow all precautions and procedures in your array's user's guide.• Check all specifications to ensure proper installation and configuration.
4.	Power on and boot the SLES system.	<ul style="list-style-type: none">• Power on the SLES system display.• Power on all peripheral devices. The array should be on, the Fibre Channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the SLES system is powered on, restart the system to have the new devices recognized.• Confirm the ready status of all peripheral devices, including the array.• Power on the SLES system.• Boot the SLES system.

Configuring the Host Fibre Channel HBAs

Configure the Fibre Channel HBA(s) connected to the array. HBAs have many configuration options. This section provides the minimum requirements for configuring host Fibre Channel HBAs for operation with the array. For more information, refer to the documentation for your fibre-channel HBAs.

- Queue depth requirements for array devices are specified in [Table 6-2](#).
- You may need to disable the BIOS to prevent the system from trying to boot from the array.
- Several other parameters (e.g., Fibre Channel and/or fabric options) may also need to be set. Please refer to the documentation for your HBA to determine whether other options are required to meet your operational requirements.
- Use the same settings and device parameters for all array devices.

Table 6-2: Queue Depth Requirements for the Array

Parameter	Required Value for Array
Queue depth per LU	≤ 32 ports
Queue depth per port (MAXTAGS)	≤ 512 per port



NOTE: To optimize I/O performance, you can adjust the queue depth for the array devices later as needed, so long as you keep within the specified range.

Verifying New Devices

The final step before configuring the new array disk devices is to verify that the host system recognizes the new devices. The host system automatically creates a device file for each new device recognized.

To verify new device recognition:

1. Use the **dmesg** command to display the devices ([Figure 6-1](#)).
2. Record the device file name for each new device. You will need this information when you partition the devices.
3. The device files are created under the `/dev` directory. Verify that a device file was created for each new disk device ([Figure 6-2 on page 6-4](#)).

```
# dmesg | more
:
:
scsil : ioc1: LSI53C1030, FwRev=01032700h, Ports=1, MaxQ=255, IRQ=209
      Emulex LightPulse Fibre Channel SCSI driver 8.1.10.3
      Copyright(c) 2004-2006 Emulex. All rights reserved.

scsi : 1 host.
Vendor: HITACHI      Model: DF600F      Rev: 0000
Type:   Direct-Access      ANSI SCSI revision: 04
sd 0:0:0:0: Attached scsi disk sda
      scsi ID0/channel 0/ID 0/lun 0      Device file name of this disk = /dev/sda

Vendor: HITACHI      Model: DF600F      Rev: 0000
Type:   Direct-Access      ANSI SCSI revision: 04
sd 0:0:0:1: Attached scsi disk sdb
:
:
```

Figure 6-1: Example of Verifying New Device Recognition



NOTE: In [Figure 6-1](#), the HITACHI DF600F device (TID 0, LUN 0) and the HITACHI DF600F device (TID 0, LUN 1) are recognized by the SLES server.

```
# ls -l /dev | more
:
brw-rw---- 1 root  disk   8,  0 Oct 30 2007 sda      Device file = sda
```

Figure 6-2: Example of Verifying Device Files

Verifying and Discovering LUNs

This section describes how to configure the new array disk devices that you attached to the SUSE Linux Enterprise system host in the previous sections.

Partitioning the Devices

After the setting the number of logical units, use the following procedure to set the partitions on the array disk devices:

1. Enter **fdisk/dev/<device_name>**
Example: **fdisk/dev/sda**
where **dev/sda** is the device file name
2. Select **p** to display the present partitions.
3. Select **n** to make a new partition. You can make up to four primary partitions (1-4) or one extended partition. The extended partition can be organized into 11 logical partitions, which can be assigned partition numbers from 5 to 15.
4. Select **w** to write the partition information to disk and complete the fdisk command.



NOTE: Other commands that you might want to use include **d** to remove partitions or **q** to stop a change.

5. Repeat steps 1 through 4 for each new array disk device.

High Availability (HA) Cluster Configurations

The Hitachi Modular storage system is compatible with various cluster software applications. For more information about:

- Compatible cluster software applications, refer to the Hitachi Data Systems interoperability matrix at <http://www.hds.com/products/interoperability/>.
- Installing and configuring the cluster software, refer to the documentation provided by the cluster software vendor.

HA Multipath Configurations

The Hitachi Modular storage system supports various HA multipathing software products for the SuSE Linux operating system. Refer to the Hitachi Data Systems interoperability matrix at <http://www.hds.com/products/interoperability/> for currently supported HA software applications. Then consult the documentation provided by the HA multipathing vendor for information about installing, configuring, operating, and best practices when using the software with Active/Active storage systems like the Hitachi Modular storage systems.

If the Hitachi Data Systems interoperability matrix show that SuSE Linux's bundled multipathing software, referred to as "Device Mapper," is supported for the intended SuSE operating system version /update level and you want

to use this bundled Multipath solution, refer to the appropriate SuSE Device Mapper documentation for proper installation, configuration, and operation. Some Device Mapper release level documentation can be obtained at the following link:

<http://www.novell.com/>

To ensure Active/Active I/O activity by the host Linux I/O to the Hitachi Modular storage system, confirm that the following minimum parameters are set in the file `/etc/multipath.conf`:

- Vendor: Hitachi
- Product: DF600F
- Path_grouping _policy: Multibus

Install and configure the multipathing software on the server before connecting the server to the Hitachi Modular storage system.

Asianux

This chapter discusses guidelines on how to prepare an Asianux host server for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Preparing the Host Server](#)
- ❑ [Configuring the Host Fibre Channel HBAs](#)
- ❑ [High Availability \(HA\) Cluster Configurations](#)
- ❑ [HA Multipath Configurations](#)
- ❑ [Verifying and Discovering LUNs](#)
- ❑ [Verifying New Device Recognition](#)

Preparing the Host Server

Table 7-1 lists the tasks to follow to prepare the host server. To prepare the server, you need one or more supported HBAs with the latest supported BIOS and driver. Verify the HBAs, drivers, and BIOSes are the latest supported versions by Hitachi Data Systems and are functioning properly.

- To check the latest supported versions, refer to the Hitachi interoperability matrix at <http://www.hds.com/products/interoperability/>.
- For information about your HBAs, refer to the HBA vendor's Web site.
- For information about HBAs supported by your operating system, refer to the HP Web site:
<http://www.asianlinux.net>.

Table 7-1: Steps for Connecting to an Asianux Host

Step	Activity	Description
1.	Verify the array installation.	Confirm that the status of the Fibre Channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the Asianux system.	<ul style="list-style-type: none">• Shut down the Asianux system.• When shutdown is complete, power off the Asianux display.• Power off all peripheral devices, except for the array.• Power off the host system.• You are now ready to connect the array.
3.	Connect the array.	<ul style="list-style-type: none">• Install Fibre Channel cables between the array and the Asianux system.• Follow all precautions and procedures in your array's user's guide.• Check all specifications to ensure proper installation and configuration.
4.	Power on and boot the Asianux system.	<ul style="list-style-type: none">• Power on the Asianux system display.• Power on all peripheral devices. The array should be on, the Fibre Channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the Asianux system is powered on, restart the system to have the new devices recognized.• Confirm the ready status of all peripheral devices, including the array.• Power on the Asianux system.• Boot the Asianux system.

Configuring the Host Fibre Channel HBAs

Configure the Fibre Channel HBA(s) connected to the array. HBAs have many configuration options. This section provides the minimum requirements for configuring host Fibre Channel HBAs for operation with the array. For more information, refer to the documentation for your Fibre Channel HBAs.

- Queue depth requirements for array devices are specified in [Table 7-2](#).
- You may need to disable the BIOS to prevent the system from trying to boot from the array.
- Several other parameters (e.g., Fibre Channel and/or fabric options) may also need to be set. Please refer to the documentation for your HBA to determine whether other options are required to meet your operational requirements.
- Use the same settings and device parameters for all array devices.

Table 7-2: Queue Depth Requirements for the Array

Parameter	Required Value for Array
Queue depth per LU	≤ 32 ports
Queue depth per port (MAXTAGS)	≤ 512 per port



NOTE: To optimize I/O performance, you can adjust the queue depth for array devices later as needed, so long as you keep within the specified range.

High Availability (HA) Cluster Configurations

The Hitachi Modular storage system is compatible with various cluster software applications. For more information about:

- Compatible cluster software applications, refer to the Hitachi Data Systems interoperability matrix at <http://www.hds.com/products/interoperability/>.
- Installing and configuring the cluster software, refer to the documentation provided by the cluster software vendor.

HA Multipath Configurations

The Hitachi Modular storage system supports various HA multipathing software products for the Asianux Linux operating system. Refer to the Hitachi Data Systems interoperability matrix at <http://www.hds.com/products/interoperability/> for currently supported HA software applications. Then consult the documentation provided by the HA multipathing vendor for information about installing, configuring, operating, and best practices when using the software with Active/Active storage systems like the Hitachi Modular storage systems.

If the Hitachi Data Systems interoperability matrix show that Asianux's bundled multipathing software, referred to as "Device Mapper," is supported for the intended Asianux operating system version /update level and you want to use this bundled Multipath solution, refer to the appropriate Asianux Device Mapper documentation for proper installation, configuration, and operation. Some Device Mapper release level documentation can be obtained at the following link:

<http://www.asianux.com/asianux.do>

To ensure Active/Active I/O activity by the host Linux I/O to the Hitachi Modular storage system, confirm that the following minimum parameters are set in the file `/etc/multipath.conf`:

- Vendor: Hitachi
- Product: DF600F
- Path_grouping_policy: Multibus

Install and configure the multipathing software on the server before connecting the server to the Hitachi Modular storage system.

Verifying and Discovering LUNs

This section describes how to configure the new array disk devices that you attached to the Asianux system host.

Partitioning the Devices

After the setting the number of logical units, use the following procedure to set the partitions on array disk devices:

1. Enter **`fdisk/dev/<device_name>`**

Example: **fdisk/dev/sda**

where **dev/sda** is the device file name

2. Select **p** to display the present partitions.
3. Select **n** to make a new partition. You can make up to four primary partitions (1-4) or one extended partition. The extended partition can be organized into 11 logical partitions, which can be assigned partition numbers from 5 to 15.
4. Select **w** to write the partition information to disk and complete the fdisk command.



NOTE: Other commands that you might want to use include **d** to remove partitions or **q** to stop a change.

5. Repeat steps 1 through 4 for each new array disk device.

Verifying New Device Recognition

The final step before configuring the new array disk devices is to verify that the host system recognizes the new devices. The host system automatically creates a device file for each new device recognized.

To verify new device recognition:

1. Use the **dmesg** command to display the devices ([Figure 7-1](#)).



NOTE: In the example in [Figure 7-1](#), the HITACHI DF600F device (TID 0, LUN 0) and the HITACHI DF600F device (TID 0, LUN 1) are recognized by the Asianux server.

2. Record the device file name for each new array device. You will need this information when you partition the devices.
3. The device files are created under the **/dev** directory. Verify that a device file was created for each new array disk device ([Figure 7-2](#)).

```
# dmesg | more
:
: QLogic Fibre Channel HBA Driver
  qla2400 0000:06:01.0: Found an ISP2422, irq 225, iobase 0xffffffff0000022000
  qla2400 0000:06:01.0: LOOP UP detected (4 Gbps).
  qla2400 0000:06:01.0: Topology - (Loop), Host Loop address 0x1
  QLogic Fibre Channel HBA Driver: 8.01.04-d7
  QLogic QLA2462 - PCI-X 2.0 to 4Gb FC, Dual Channel
  ISP2422: PCI-X Mode 1 (100 MHz) @ 0000:06:01.0 hdma+, host#=0, fw=4.00.18 [IP]

  Vendor: HITACHI   Model: DF600F           Rev: 0000
  Type:   Direct-Access           ANSI SCSI revision: 04
  qla2400 0000:06:01.0: scsi(0:0:0:32): Enabled tagged queuing, queue depth 32.
                                scsi ID0/channel 0/ID 0/lun 32

  SCSI device sda: 8388608 512-byte hdwr sectors (4295 MB)
  SCSI device sda: drive cache: write through
  SCSI device sda: 8388608 512-byte hdwr sectors (4295 MB)
  SCSI device sda: drive cache: write through
  sda: unknown partition table
  Attached scsi disk sda at scsi0, channel 0, id 0, lun 32
                                scsi ID0/channel 0/ID 0/lun 32

  Vendor: HITACHI   Model: DF600F           Rev: 0000
  Type:   Direct-Access           ANSI SCSI revision: 04
  qla2400 0000:06:01.0: scsi(0:0:0:33): Enabled tagged queuing, queue depth 32.
                                scsi ID0/channel 0/ID 0/lun 33
:
```

Figure 7-1: Example of Verifying New Device Recognition

```
# ls -l /dev | more
:
brw-rw----  1 root   disk   8,  0 Oct 30  2007 sda           Device file = sda
```

Figure 7-2: Example of Verifying Device Files

Apple OS X

This chapter discusses guidelines on how to prepare the Apple Macintosh host servers running OS X for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Connecting the Host Server](#)
- ❑ [Configuring the Host Fibre Channel HBAs](#)
- ❑ [Verifying and Discovering LUNs](#)

Connecting the Host Server

The array comes with all the hardware and cabling required for connection to the host system(s).

Table 8-1 lists the tasks to follow to prepare the host server. To prepare the server, you need one or more supported HBAs with the latest supported BIOS and driver. Verify the HBAs, drivers, and BIOSes are the latest supported versions by Hitachi Data Systems and are functioning properly.

- To check the latest supported versions, refer to the Hitachi interoperability matrix at <http://www.hds.com/products/interoperability/>.
- For information about your HBAs, refer to the HBA vendor's Web site.
- For information about HBAs supported by your operating system, refer to the Apple Web site:
<http://www.apple.com>.

Table 8-1: Steps for Connecting to an Apple Macintosh OS X Host

Step	Activity	Description
1.	Verify the array installation.	Confirm that the status of the Fibre Channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the Apple Macintosh OS X system.	<ul style="list-style-type: none">• Shut down the Apple Macintosh OS X system.• Power off all peripheral devices, except for the array.• Power off the host system.• You are now ready to connect the array.
3.	Connect the array.	<ul style="list-style-type: none">• Install Fibre Channel cables between the array and the Apple Macintosh OS X system.• Follow all precautions and procedures in your array's user's guide.• Check all specifications to ensure proper installation and configuration.
4.	Define LUN security if connected to a fabric	<ul style="list-style-type: none">• If the array is connected to a fabric, control access to the array before you power on the Apple Macintosh OS X host to avoid failures on other systems connected to the same fabric.• Use any method to ensure that Apple Macintosh OS X host sees only the LUs it owns.
5.	Power on and boot the Apple Macintosh OS X system.	<ul style="list-style-type: none">• Power on the Apple Macintosh OS X system display.• Power on all peripheral devices. The array should be on, the Fibre Channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the Apple Macintosh OS X system is powered on, restart the system to have the new devices recognized.• Confirm the ready status of all peripheral devices, including the array.• Power on the Apple Macintosh OS X system.

Table 8-1: Steps for Connecting to an Apple Macintosh OS X Host (Continued)

Step	Activity	Description
6.	Verify new device recognition, and record the disk numbers	When the HBA connected to the array shows the new devices, record the number of disks. You will need this information.
7.	Boot the Apple Macintosh OS X system	

Configuring the Host Fibre Channel HBAs

Use the following procedure to configure the HBAs using the Apple Macintosh host running OS X.

1. Select **System preferences** from the Apple menu.
2. Select **Fibre Channel** from **System references**. The Fibre Channel panel shown in [Figure 8-1](#) appears.




Figure 8-1: Fibre Channel Panel

3. Click the key at the lower left side of the Fibre Channel panel to unlock the settings.
4. On the left side of the panel, select the port to which you want to connect.
5. On the right side of the panel, select the port speed and port topology. You can select the **Auto setting** for both, but doing so may prevent a connection from being made. We recommend you set the port speed to match the port speed of the switch used in your configuration. Then set the port topology to **Point to Point** when you connect the switch.

Verifying and Discovering LUNs

This section describes how to configure the new array disk devices that you attached to the Apple Macintosh OS X host in the previous section.

Connecting the Array and Server

To connect the array and the Apple Macintosh OS X host:

1. Shut down the server.
2. After checking that the server shut down, turn off power to the server and all peripherals.
3. Connect the array to the server using a Fibre Channel cable.
4. Turn on power to the array and verify that it is ready.
5. Turn on power to all the peripherals, check that they are ready.
6. Turn on power to the server.
7. Log in to the Apple Macintosh OS X operating system.

Verifying Recognition of the Array

After connecting the array and the Apple Macintosh OS X host, use the following procedure to confirm that the host recognizes the array.

1. Select **Utilities** from the **Apple** menu (Figure 8-2).



Figure 8-2: Utilities Icon

2. Open the Disk Utility window from **Utilities** (Figure 8-3).



Figure 8-3: Disk Utility Window

3. Open **Disk Utility** and confirm that the device **HITACHI DF600F** is recognized ([Figure 8-4](#)).




Figure 8-4: Opening the Disk Utility

4. In the left pane ([Figure 8-5](#)), select the disk you want to use (for example, the disk you want to partition).




Figure 8-5: Example of Selecting a Drive

IBM AIX

This chapter discusses guidelines on how to prepare an IBM AIX host server for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Preparing the Host Server](#)
- ❑ [Setting the Disk and Device Parameters](#)
- ❑ [Changing Device Parameters using SMIT](#)
- ❑ [Changing Device Parameters from the AIX Command Line](#)
- ❑ [Configuring the New Disk Devices](#)
- ❑ [Using AIX with Object Data Manager \(ODM\)](#)

Preparing the Host Server

Table 9-1 lists the tasks to follow to prepare the host server. To prepare the server, you need one or more supported host bus adapters (HBAs) with the latest supported BIOS and driver. Verify the HBAs, drivers, and BIOSes are the latest supported versions by Hitachi Data Systems and are functioning properly.

- To check the latest supported versions, refer to the Hitachi interoperability matrix at <http://www.hds.com/products/interoperability/>.
- For information about your HBAs, refer to the HBA vendor's Web site.
- For information about HBAs supported by your operating system, refer to the HP Web site: <http://www.ibm.com/us/>.

Table 9-1: Steps for Connecting to an IBM AIX Host

Step	Activity	Description
1.	Verify the array installation.	Confirm that the status of the Fibre Channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the AIX system.	<ul style="list-style-type: none">• Shut down the IBM AIX system.• When shutdown is complete, power off the IBM AIX display.• Power off all peripheral devices, except for the array.• Power off the host system.• You are now ready to connect the array.
3.	Connect the array.	<ul style="list-style-type: none">• Install Fibre Channel cables between the array and the IBM AIX system.• Follow all precautions and procedures in your array's user's guide.• Check all specifications to ensure proper installation and configuration.
4.	Power on the IBM AIX system.	<ul style="list-style-type: none">• Power on the IBM AIX system display.• Power on all peripheral devices. The array should be on, the Fibre Channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the IBM AIX system is powered on, restart the system to have the new devices recognized.• Confirm the ready status of all peripheral devices, including the array.• Power on the IBM AIX system.

Setting the Disk and Device Parameters

After the array is installed and connected, and device files are created, the IBM AIX system sets device parameters to system default values. Using the AIX System Management Interface Tool (SMIT) or the AIX command line, you can change the queue type and queue depth parameters for each new array device. Both methods are described in the following sections.



NOTE: When you set parameters for Fibre Channel disk devices, use the same settings and device parameters for all array devices.

Table 9-2 lists the queue type requirements for array devices. Table 9-3 lists the queue depth requirements for array devices.

Table 9-2: Queue Type Requirements for Modular Arrays

Parameter Name	Default Value	Required Value for Array
Queue type	None	Simple

Table 9-3: Queue Depth Requirements for Modular Arrays

Parameter Name	Required Value for Array
Queue depth per LU	1 to 32 commands per LUN
Queue depth per port	512 commands per port



NOTE: To optimize the I/O performance of the devices, you can adjust the queue depth for array devices within the specified range later as needed.

Changing Device Parameters using SMIT

To change the device parameters using SMIT:

1. At the AIX command line prompt, type the following command on the command line to start SMIT and open the System Management panel:
smit
2. On the SMIT System Management panel, select **Devices**. The Devices panel appears.
3. Select **Fixed Disk**. The Fixed Disk panel appears.
4. Select **Change/Show Characteristics of a Disk**. The Disk panel appears.
5. Select the desired device from the **Disk** menu. The Change/Show Characteristics of a Disk panel appears.
6. Enter the desired queue depth ([Table 9-3 on page 9-3](#)), queue type (simple). Press Enter to complete the parameter changes.
7. Repeat steps 5 and 6 for each new device on the array.
8. Type the command in [Figure 9-1](#) to verify that the parameters for all devices were changed:

Change/Show Characteristics of a Disk

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[MORE...11]

[Entry Fields]

FC Node Name

Physical volume IDENTIFIER

ASSIGN physical volume identifier

Queue DEPTH

Queuing TYPE

Use QERR bit

Device CLEARS its Queue on error

READ/WRITE time out value

Maximum LUN inquired on device

START UNIT time out value

REASSIGN time out value

Maximum TRANSFER Size

Apply change to DATABASE only

[BOTTOM]

none

no

[1] <- See Table 9-3

[simple] <- Enter **simple** here.

[yes]

[no]

[30]

[]

[60]

[120]

[0x40000]

no

F1=Help

F2=Refresh

F3=Cancel

F4=List

Esc+5=Reset

Esc+6=Command

Esc+7=Edit

Esc+8=Image

Esc+9=Shell

Esc+0=Exit

Enter=Do

Figure 9-1: Changing Device Parameters Using SMIT

Changing Device Parameters from the AIX Command Line

To change the device parameters from the AIX command line:

1. Type the following command at the AIX command line prompt to display the parameters for the specified device:

lsattr -E -l hdiskx



NOTE: 'hdiskx' is the device file name (for example, hdisk2). You can also use the **lscfg -vl hdiskx** command ([Figure 9-2](#)).

2. Type the following command to change the device parameters:

chdev -l hdiskx -a rw_timeout='30' -a q_type='simple' -a queue_depth='x'

```
# chdev -l hdisk4 -a rw_timeout=30
hdisk4 changed
EP8000-521# lsattr -El hdisk4
PR_key_value      none                Reserve Key                True
clr_q             no                 Device CLEARS its Queue on error True
location          Location Label      True
lun_id            0x100000000000000 Logical Unit Number ID     False
max_transfer      0x40000            Maximum TRANSFER Size      True
node_name         FC Node Name       False
pvid              00cle29aec8fa056000000000000000000000 Physical Volume ID        False
q_err             yes                Use QERR bit               False
q_type            simple             Queue TYPE                  True
queue_depth       16                Queue DEPTH                 True
reassign_to       120               REASSIGN time out          True
reserve_policy    single_path        Reserve Policy              True
rw_timeout        30                READ/WRITE time out        True
scsi_id           0xef              SCSI ID                     False
start_timeout     60                START UNIT time out        True
www_name          0x50060e8010400780 FC World Wide Name         False
```

Figure 9-2: Changing Timeout Device Parameters Using SMIT

```
# chdev -l hdisk4 -a queue_depth=16
hdisk4 changed
```

Figure 9-3: Changing Queue_depth Device Parameters Using SMIT



NOTE: **x** is used to indicate the desired queue depth within the limits specified in [Table 9-3 on page 9-3](#).

3. Repeat steps 1 and 2 for each new array device.
4. Type the following command to verify that the parameters for all devices were changed (see [Figure 9-4 on page 9-6](#) and [Figure 9-5 on page 9-6](#)).

#lsattr -E -l hdisk4			
clr_q	no	Device CLEARS its Queue on error	True
location		Location Label	True
lun_id	0x10000000000000	Logical Unit Number ID	False
max_transfer	0x40000	Maximum TRANSFER Size	True
node_name		FC Node Name	False
pvid	00cle29aec8fa056000000000000000	Physical volume identifier	False
q_err	yes	Use QERR bit	True
q_type	simple	Queuing TYPE	True
queue_depth	1	Queue DEPTH	True
reassign_to	120	REASSIGN time out value	True
rw_timeout	30	READ/WRITE time out value	True
scsi_id	0xef	SCSI ID	False
start_timeout	60	START unit time out value	True
ww_name	0x50060e8010400780	FC World Wide Name	False

**Figure 9-4: Verifying the Device Parameters
Using the lsattr -E -l hdiskx Command**

```

lscfg -vl hdisk4
  hdisk4  U787A.001.DPM0GZG-P1-C1-T1-W50060E8010400780-L1000000000000  Other FC
SCSI Disk Drive
  Manufacturer.....HITACHI
  Machine Type and Model.....DF600F
  Part Number.....P
  ROS Level and ID.....30303030
  Serial Number.....83000120
  EC Level.....
  FRU Number.....0120
  Device Specific.(Z0).....00000432B3001102
  Device Specific.(Z1).....0001

```

**Figure 9-5: Verifying the Device Parameters
Using the lscfg -vl hdisk1 Command**

Configuring the New Disk Devices

This section describes how to configure the array disk devices that you attached to the IBM AIX system host in the previous sections.

Verifying New Device Recognition

This section provides guidelines on how to connect the array to the host.

Before you connect to the system, verify the steps in [Table 9-1 on page 9-2](#) are completed.

Hitachi Data Systems recommends that you install and format the devices with the fibre ports configured before the host system is powered on. Enter the **cfgmgr** command to force the system to check the buses for new devices.

To verify new device recognition:

1. Log in to the host system as root.
2. Display the system device data by entering the following command ([Figure 9-6 on page 9-7](#) or [Figure 9-7 on page 9-7](#)):

lsdev -C -c disk

```
# lsdev -C -c disk                                     Display device data
hdisk0 Available 08-08-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 08-08-00-5,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 08-08-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk3 Available 04-08          Other FC SCSI Disk Drive      New device
hdisk4 Available 04-08          Other FC SCSI Disk Drive      New device
hdisk5 Available 04-08          Other FC SCSI Disk Drive      New device
hdisk6 Available 04-08          Other FC SCSI Disk Drive      New device
Device file name = hdiskx
:
#
```

Figure 9-6: Verifying New Device Recognition without Object Data Manager (ODM)

```
# lsdev -Cc disk                                     Display device data
hdisk0 Available 08-08-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 08-08-00-5,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 08-08-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk3 Available 04-08          Hitachi Disk Array (Fibre)    New device
hdisk4 Available 04-08          Hitachi Disk Array (Fibre)    New device
hdisk5 Available 04-08          Hitachi Disk Array (Fibre)    New device
hdisk6 Available 04-08          Hitachi Disk Array (Fibre)    New device
Device file name = hdiskx
```

Figure 9-7: Verifying New Device Recognition with ODM

Using AIX with Object Data Manager (ODM)

This section describes the AIX Object Data Manager (ODM) and its relationship with the array.

ODM Overview

The ODM is a repository of system information, which includes the basic components of object classes and characteristics. Information is stored and maintained as objects with associated characteristics.

System data managed by ODM includes:

- Device configuration information
- Display information for SMIT (menus, selectors, and dialogs)
- Vital product data for installation and update procedures
- Communications configuration information
- System resource information

IBM provides a predefined set of devices (PdDv), and attributes (PdAt). Hitachi Data Systems has added its own device definitions to the ODM, based on classes defined as objects with associated characteristics. This allows you to add devices that are recognized when the system boots or when the configuration manager command (cfgmgr) is executed. These devices have their own set of predefined attributes, which allows you to customize device definitions easily and automatically, thereby minimizing the amount of work required to define a device.

IBM provides a set of commands to manipulate the ODM and procedures to package ODM updates. For more information, see [ODM References on page 9-16](#).

Installing the ODM 5.0.0.1 Fibre Update Package

The ODM update filesets include the ODM parameters for all Hitachi disks. The current base installation file contains two install images:

- 5.0.0.1I AIX Support for Hitachi Disk Arrays Base Installation
- 5.0.0.0I HACMP Support for Hitachi Disk Arrays

Both filesets are required when using concurrent resources with High Availability Multi-Cluster Processing (HACMP). HACMP must be installed prior to installing the HACMPTM support fileset.

The following four updates can be installed the same time you perform the base installation:

- 5.0.0.4U — update for device support
- 5.0.52.1U — update for AIX 5.2 and above
- 5.0.52.2U — update for HP/HDS co-existence
- 5.0.52.3U — update for HMDE devices



NOTE: ODM updates can be found on the Product Documentation Library (PDL) CDs that ship with the hardware. The ODM update for Hitachi Data Systems LUN definitions is only for releases of Hitachi Dynamic Link Manager (HDLM). Previous versions of the ODM must be uninstalled before installing version 5.0.0.1 because this is a base level fileset.

To install the Hitachi Data Systems ODM Fibre updates for AIX:

1. Type the following command to copy the fileset to an AIX directory (/usr/sys/inst.images)

```
# cp HDS_FC_ODM2.bff /usr/sys/inst.images
```

2. Type the following commands to create a new .toc file:

```
#cd /usr/sys/inst.images
```

```
#inutoc .
```

3. Use the command **installp** to install the AIX support fileset:

```
#smitty install
```

4. Select **5.0.0.1 AIX Support for Hitachi Disk Arrays**.
5. Select **Install and Update Software**.
6. Select **Install and Update from LATEST Available Software**.
7. When prompted for an input device, enter the directory where the fileset was copied.
8. Press F4 (or Esc+4) to list the available installation packages.
9. Press F7 (or Esc+7) to select **5.0.0.1 AIX Support for Hitachi Disk Arrays**, then select **Enter**.
10. Press Enter to start the installation process.

11. If necessary, remove any Hitachi Data Systems devices already defined as **Other FC SCSI**.

- To remove one device, type the command:
#rmdev -dl hdisk(n)
- To remove all devices on a path, type the command:
#rmdev -dl fscsi(n) -R

12. Reboot AIX.

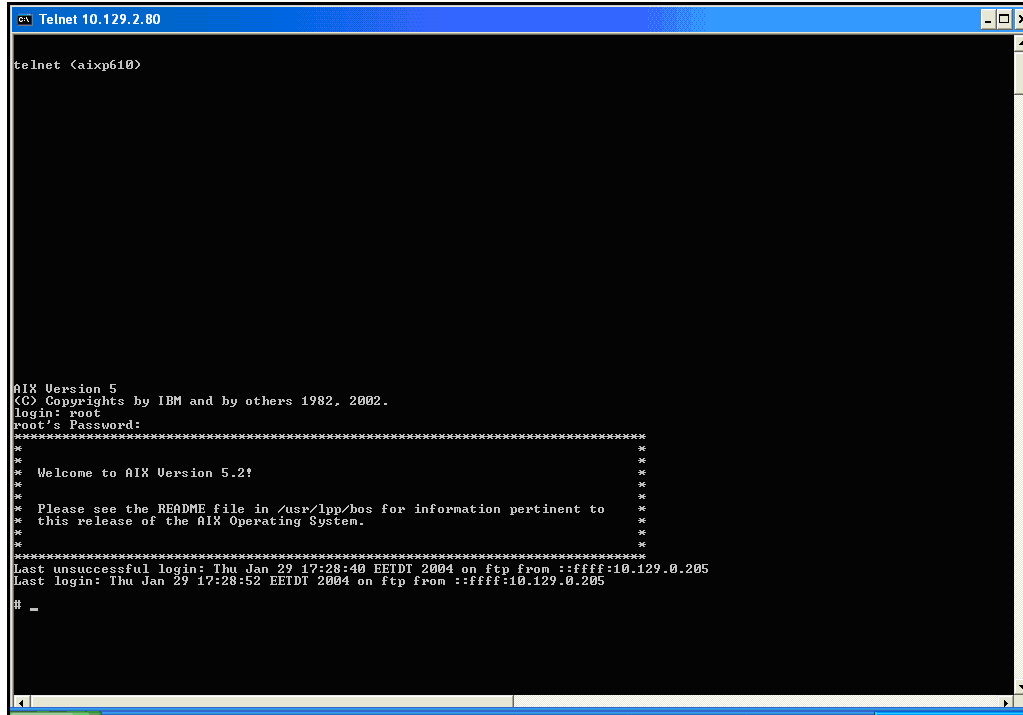
13. Confirm that the parameters have the following values:

```
# lsattr -El hdisk4
PR_key_value      none
Reserve Key              True
clr_q              no
Device CLEARS its Queue on error True
location              Location Label
True
lun_id             0x40000000000000
Logical Unit Number ID      False
max_transfer       0x40000
Maximum TRANSFER Size      True
node_name          0x50060e8004280212
FC Node Name              False
pvid               0007bcda3efc5c400000000000000000
Physical Volume ID          False
q_err              yes
Use QERR bit              False
q_type             simple
Queue TYPE              True
queue_depth        2
Queue DEPTH              True
reassign_to        120
REASSIGN time out      True
reserve_policy      single_path
Reserve Policy              True
rw_timeout         30
READ/WRITE time out      True
scsi_id            0x13200
SCSI ID                  False
start_timeout       60
START UNIT time out      True
ww_name            0x50060e8004280212
FC World Wide Name      False
```

Installing the ODM Fibre Update Package for 5.0.0.4

The following procedure describes how to install the ODM Fibre Update Package for 5.0.0.4. You can only install the update for 5.0.0.4 if the base fileset is 5.0.0.1.


1. Download and unzip the file ODM 5.0.0.4U.
2. Enter IBM AIX.




```
Telnet 10.129.2.80
telnet <aixp610>

AIX Version 5
(C) Copyrights by IBM and by others 1982, 2002.
login: root
root's Password:
*****
* Welcome to AIX Version 5.2! *
* Please see the README file in /usr/lpp/bos for information pertinent to *
* this release of the AIX Operating System. *
* *
*****
Last unsuccessful login: Thu Jan 29 17:28:40 EETDT 2004 on ftp from ::ffff:10.129.0.205
Last login: Thu Jan 29 17:28:52 EETDT 2004 on ftp from ::ffff:10.129.0.205
# _
```


3. Type the following command to copy the fileset to the AIX directory /usr/sys/inst.images:
cp HDS_FC_ODMUPD.bff /usr/sys/inst.images
4. Type the following commands to create a new .toc file:
#cd /usr/sys/inst.images
#inutoc .
5. Use **installp** to install the AIX support fileset:
#smitty install
6. Select **5.0.0.4 AIX Support for Hitachi Disk Arrays**.
7. Select **Install and Update Software**.



8. Select **Install and Update from LATEST Available Software**.



9. When prompted for an input device, enter the directory where the fileset was copied (for example, /usr/sys/inst.images).



10. Press F4 (or Esc+4) to list the available installation packages.
11. Press F7 (or Esc+7) to select **5.0.0.4 AIX Support for Hitachi Disk Arrays**, then select **Enter**.



NOTE: You can install any additional available Hitachi ODM updates at this time. Check with Hitachi Data Systems for the most current Hitachi ODM updates available.

12. Press Enter to start the installation process.
13. Reboot AIX.

Uninstalling the ODM Fibre Update Package

To uninstall the Hitachi Data Systems ODM Fibre updates for AIX:

1. Use the **rmdev** command to remove all Hitachi LUNs from the system:
 - To remove a single hdisk, type the command:
rmdev -dl hdisk(n)
 - To remove all hdisks on a path, type the command:
rmdev -dl fscsi(n) -R
2. From the command line, type:
smitty install
3. Select **Software Maintenance and Utilities**.
4. Select **Remove Installed Software**.
5. Select F4 (or Esc+4) for **SOFTWARE name** to list all installed software.
6. Search (/) for **Hitachi**.
7. Select the package to remove pressing F7 (or Esc+7) and then Enter.

8. Use either the Tab key or F4 (or Esc+4) to change **PREVIEW only** to **No**.
9. Select **Enter** to remove the package.
10. Select **Enter** to confirm the warning message.
11. Reboot AIX.

Installing the IBM AIX MPIO ODM_6000I Update

The following procedure describes how to install the IBM AIX MPIO ODM_6000I update for the Modular storage systems. IBM AIX MPIO ODM_6000I is the only ODM update that defaults the queue depth settings to 8.

1. Download the MPIO ODM_6000I update.
2. Type the following command to navigate to the following AIX directory where the update was downloaded: /usr/sys/inst.images
#cd /usr/sys/inst.images
3. Use the following command to list the contents and confirm that the file HDS_FC_ODMUPD_6000I exists:
#ls
4. Use the following command to create a .toc file:
#inutoc .
5. Use the following command to start SMIT to install the update
#smitty install
6. When prompted, select **Install and Update Software**.
7. Select **Install and Update from ALL Available Software**.
8. When prompted for an input device, enter the following directory:
/usr/sys/inst.images
9. Press F4 (or Esc+4) to list the available installation packages.
10. Select **SOFTWARE to Install** and press F7 (or Esc+7).
11. Press Enter.
12. Press Enter again. A warning message asks whether you are sure you want to proceed.
13. Press Enter to start the installation.
14. Issue the **lsdev** command to display disks:

```
# lsdev -Cc disk
hdisk0 Available 10-60-00-10,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 10-60-00-12,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 10-68-01      Other FC SCSI Disk Drive
hdisk3 Available 10-68-01      Other FC SCSI Disk Drive
```

15. Reboot AIX.
16. After AIX is running, issue the **lsdev** command again to check disks:

```
# lsdev -Cc disk
hdisk0 Available 10-60-00-10,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 10-60-00-12,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 10-68-01      Hitachi MPIO Disk AMS.SMS.WMS (Fibre)
hdisk3 Available 10-68-01      Hitachi MPIO Disk AMS.SMS.WMS (Fibre)
```



NOTE: To confirm that the ODM is working, verify that **Hitachi MPIO Disk AMS.SMS.WMS (Fibre)** replaces **MPIO Other FC SCSI Disk Drive** (as shown in the example above).

17. Use the following commands to check the ODM:

```
# lslpp -l |grep -i hitachi
devices.fcp.disk.Hitachi.modular.mpio.rte 6.0.0.0 COMMITTED AIX Support
for Hitachi Disk
```

This completes the ODM installation.

Using ODM

Discovering New Devices

When the system boots and a new device is discovered, the system checks the ODM for a device definition that matches the new device. For a disk device, this is based on the SCSI inquiry command. If a match is found, a customized definition (CuDv and CuAt) is built for that device using the default attributes for that device definition. The new device then has the description based in the ODM for that device. This customized definition is persistent and remains until the device is removed from the system.

An active device will have an “available” status and is ready for use. A device that was available, but has been physically removed from the system will have a “defined” status and cannot be used.

Deleting Devices

A device's definition remains until it is removed using the **rmdev** command. Some device attributes (such as physical volume identifier, SCSI ID, or Target ID) are unique to a device and remain until the device is removed using the **rmdev** command.

A device definition remains in the ODM when an attribute (e.g., the WWN) changes. The definitions in the ODM are persistent and remain until a system administrator removes them.

Default Device Definitions for “Other FC SCSI Disk Drive”

Devices that are predefined in the ODM will be uniquely identified and have a predefined set of default attributes. Disk devices that do not have a predefined definition are considered generic devices by AIX and are defined as “Other FC SCSI Disk Drive.”

Queue Depth and Read/Write Timeout Values

The default IBM read/write timeout and queue depth values are different from the recommended and required values for Hitachi disk devices. For Hitachi disk devices:

- The required value for read/write timeout is 30.
- The default queue depth value is 2/8 depending on the ODM update installed (see [Table 9-3 on page 9-3](#) for queue depth and port specifications).

If AIX defines a device as "Other FC SCSI Disk Drive," the queue depth setting for that device is ignored, which can have a negative impact on performance.

Hitachi disk devices must be defined as "Hitachi," not as "Other," and the queue depth should be at least 2.

ODM Advantages and Cautions

Advantages

Hitachi Data Systems ODM updates lets AIX recognize Hitachi disk devices and set the proper attributes. If the attributes for queue type, queue depth, and read/write timeout for all Hitachi disk devices are not the same, disk errors can be logged on both the disk array and in the AIX error log.

If the Hitachi ODM update is installed and a device is discovered, a match will be found in the ODM and the attributes will be set to the defaults recommended by the manufacturer. For Hitachi disk devices, the default queue depth will be set to 2 or 8, depending on the ODM update installed (with a range of 1-32), and the read/write timeout value will be set to 60. If the Hitachi ODM update is not installed, a system administrator will be required to run a **chdev** (change device) command for every device on the system to change the default attributes.

Sharing Ports

The Hitachi ODM update changes attributes. This can cause problems if you share ports and use or do not use host storage domains on Modular storage systems connected to multiple AIX servers running different levels of ODM updates. An example of AIX servers running different ODM update levels is where one AIX host runs v5.0.0.0 and another AIX host runs v5.0.0.1 and above. (AIX v5.0.0.0 is obsolete and has different mode data settings, which can cause problems when sharing a storage port without using HSDs.) For this reason, be sure to use the same ODM update with both controllers in the Modular storage systems.

ODM References

For more information about the configuration database and ODM, refer to the following references:

Device Configuration Database

http://www16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprgpd/kernextc/device_config_db_over.htm

Device Configuration System

http://www16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprgpd/kernextc/device_config_subsys.htm#a4d56110chri

List of ODM Commands and Subroutines

http://publib.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprgpd/genprogc/odm_cmds_subrs.htm

http://publib.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprgpd/genprogc/odm.htm (Chapter 17. Object Data Manager)

http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprgpd/genprogc/pkging_sw4_install.htm (Chapter 20. Packaging Software for Installation)

IBM Redbook

Certification Study Guide-pSeries® - AIX System Support

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246199.pdf>

HP-UX

This chapter discusses guidelines on how to prepare an HP-UX host server for connection to the array and verify that the host server can connect to the array.

This chapter covers the following key topics:

- ❑ [Preparing the Host Server](#)
- ❑ [Connecting to the Array](#)
- ❑ [Setting Disk and Device Parameters](#)
- ❑ [Verifying and Discovering LUNs](#)

Preparing the Host Server

This section provides guidelines to consider when preparing your host server.

Host Bus Adapter Guidelines

One or more host bus adapters (HBAs) are required to connect your host server to the array. Before installing and connecting HBAs, verify that the HBAs are working properly and have the latest drivers and BIOSes supported by your operating system and the array.

For information about:

- The latest HBAs supported by the array, refer to the Hitachi interoperability matrix at <http://www.hds.com/products/interoperability/>.
- Installing and configuring your HBAs, refer to the HBA documentation and the HBA vendor's Web site.
- HBAs supported by your operating system, refer to the HP Web site at <http://www.hp.com/>.

Operating System Guidelines

Verify that the OS version, architecture, relevant patches, and maintenance levels you intend to use are supported by Hitachi Data Systems. For information about supported versions, refer to the Hitachi interoperability matrix.

Connecting to the Array

This section provides guidelines on how to connect the array to the host.

Before you connect to the system:

1. Verify the items described under [Preparing the Host Server](#), above, were completed.
2. Log in to the host system as root, and make sure that all existing devices are powered on and are properly connected to the host system.
3. Verify the host recognizes the following four classes:
 - Fibre channel adapter name
 - Hardware (H/W) path
 - Software (S/W) state
 - Fibre channel driver

If this information does not appear or if error messages are displayed, the host environment may not be configured properly.

After configuring Fibre Channel ports on the array, verify that the HBAs are installed properly using the **ioscan f** command (Figure 10-1)

```
# ioscan -f
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
=====						
bc	0		root	CLAIMED	BUS_NEXUS	
bc	1	8	bc	CLAIMED	BUS_NEXUS	Pseudo Bus Converter
bc	2	8/0	bc	CLAIMED	BUS_NEXUS	Bus Converter
tty	0	8/0.0	mux2	CLAIMED	INTERFACE	MUX
ext_bus	0	8/4	c720	CLAIMED	INTERFACE	GSC add on Fast/Wide SCSI INTERFACE
target	0	8/4.5	tgt	CLAIMED	DEVICE	
disk	0	8/4.5.0	sdisk	CLAIMED	DEVICE	SEGATE ST321S1W
target	1	8/4.7	tgt	CLAIMED	DEVICE	
ctl	0	8/4.7.0	sctl	CLAIMED	DEVICE	Initiator
fc	0	8/12	fcTl	CLAIMED	INTERFACE	HP Fibre Channel Mass Storage Adapter
lan	1	8/12.5	fcTl_cntl	CLAIMED	INTERFACE	HP Fibre Channel Mass Storage Cntl
fc	0	8/12.8	fc	CLAIMED	INTERFACE	FCP Protocol Adapter
ext_bus	2	8/12.8.0.255.0	fcpdev	CLAIMED	INTERFACE	FCP Device Interface
target	7	8/12.8.0.255.0.6	tgt	CLAIMED	DEVICE	
target	8	8/12.8.0.255.0.8	tgt	CLAIMED	DEVICE	
target	9	8/12.8.0.255.0.9	tgt	CLAIMED	DEVICE	
#						

Figure 10-1: Verifying Host Fibre Channel Adapter Installation

4. Connect the array to the HP-UX host.



NOTE: For information about the HBA-specific text displayed on the screen, refer to the MAN pages and user documentation for the HBA.

5. Use Storage Navigator Modular 2 to configure the array as follows (for more information, refer to the Storage Navigator Modular 2 online help):
 - Platform = **HP-UX**
 - Mode Settings:
 - Common Setting = **Standard Mode**
 - Additional Setting = **HP-UX Mode** and **PSUE Read Reject Mode**



NOTE: HP-UX supports following multi-pathing software:

- PVLink(HP-UX 11.i V1.0/V2.0)
- Native Multi-path(HP-UX 11.i V3.0)
- -VxVM

Setting Disk and Device Parameters

To set or change the timeout value to 60 seconds for a disk in HP-UX, run the following commands:

```
#for i in `ls /dev/dsk/*`  
>do  
>pvchange -t 60 $i  
>done  
:
```

An **IO_timeout** value of zero (0) causes the system to use the default value supplied by the device driver associated with the physical device.

IO_timeout is used by the device driver to determine the amount of time to wait for disk transactions to complete before concluding that an I/O request cannot be completed (and the device is offline or unavailable).

Verifying and Discovering LUNs

The first step in configuring the array devices for HP-UX operations is to verify that the HP system recognizes the new devices on the array. The devices must be installed and formatted and the Fibre Channel ports defined before the host system is powered on. If the fibre ports are configured after the HP system is powered on, halt and restart the system to allow the system to recognize the new devices.

To verify that the HP system recognizes new array devices:

1. Login to HP system as root ([Figure 10-2 on page 10-5](#)).
2. Display device data using the **ioscan -f** command, and verify that the system recognizes newly installed devices on the array ([Figure 10-3](#) and [Figure 10-4 on page 10-6](#) and [Figure 10-5 on page 10-6](#) — these figures are used with HP-UX 11.31 only). If desired, use the **-C disk** command option (**ioscan -fnC disk**) to limit output to disk devices only.



NOTE: If **UNKNOWN** appears as the Class type, the HP system may not be configured properly. Refer to HP documentation or contact HP technical support. If information for unused devices remains in the system, obtain system administrator permission to renew the device information. To renew device information, delete the `/etc/ioconfig` and `/stand/ioconfig` files (**rm** command), reboot server, and issue the **ioinit -c** command followed by the **ioscan -f** command to recognize logical devices again. You can also use the command **rmsf -H <hardware address of unused devices>** to remove unused device files.

3. Enter device data for each array device in your table, including device file name. Device file names have the following structure:

cXtYdZ except persistent device file

In this filename structure:

- **c** = controller
- **X** = bus instance #
- **t** = target ID

- **Y** = target ID
- **d** = device
- **Z** = LUN



NOTE: The SCSI target IDs are in hexadecimal (0 through F) notation and the LUN is in decimal (0 through 7) notation.

```
The system is ready.

GenericSysName [HP Release B.11.23] (/etc/issue)
Console Login: root                               Log in as root
Password:                                           Enter password (not displayed)
Please wait...checking for disk quotas
(c)Copyright 1983-1995 Hewlett-Packard Co., All Rights Reserved.
:
#
```

Figure 10-2: Logging In

Figure 10-3 shows the following new array devices recognized:

- HITACHI DF600F device: bus number = 8/12, bus instance = 8, target ID = 6, LUN = 0, driver = sdisk.
- HITACHI DF600F device: bus number = 8/12, bus instance = 8, target ID = 6, LUN = 1, driver = sdisk.

```
# ioscan -fn
Class      I    H/W Path      Driver      S/W State H/W Type  Description
=====
bc          0                      root        CLAIMED   BUS_NEXUS
bc          1    8                      bc          CLAIMED   BUS_NEXUS Bus Converter
fc          0    8/12                  fcTl        CLAIMED   INTERFACE HP Fibre Channel Mass Storage
fcp         0    8/12.8                fcp         CLAIMED   INTERFACE FCP Protocol Adapter
ext_bus     2    8/12.8.0.255.0        fcpdev      CLAIMED   INTERFACE FCP Device Interface
target      7    8/12.8.0.255.0.6      tgt         CLAIMED   DEVICE
disk        3    8/12.8.8.255.0.6.0    sdisk       CLAIMED   DEVICE      HITACHI DF600F
              /dev/dsk/c2t6d0      /dev/rdisk/c2t6d0
disk        4    8/12.8.8.255.0.6.1    sdisk       CLAIMED   DEVICE      HITACHI DF600F
              /dev/dsk/c2t6d1      /dev/rdisk/c2t6d1
#
```

Figure 10-3: Verifying New Device Recognition (HP-UX 11.23)

```
# ioscan -nkf
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
=====						
slot	0	0/2/1	pci_slot	CLAIMED	SLOT	PCI Slot
fc	0	0/2/1/0	fcd	CLAIMED	INTERFACE	HP A6826-60001 2Gb
			/dev/fcd0			
fc	7	0/2/1/0.1fcd_fc	fcd_fc	CLAIMED	INTERFACE	FCP Domain
ext_bus	32	0/2/1/0.1.0.0.0	fcd_vbus	CLAIMED	INTERFACE	FCP Array Interface
target	1	0/2/1/0.1.0.0.0.0	tgt	CLAIMED	DEVICE	
disk	306	0/2/1/0.1.0.0.0.0.0	sdisk	CLAIMED	DEVICE	HITACHI DF600F
			/dev/dsk/c32t0d0		/dev/rdisk/c32t0d0	
slot	5	0/6/1	pci_slot	CLAIMED	SLOT	PCI Slot
fc	6	0/6/1/0	fcd	CLAIMED	INTERFACE	HP A6826-60001 2Gb
			/dev/fcd6			
fc	3	0/6/1/0.8	fcd_fc	CLAIMED	INTERFACE	FCP Protocol Adapter
ext_bus	11	0/6/1/0.8.0.2.0	fcd_vbus	CLAIMED	INTERFACE	FCP Array Interface
target	11	0/6/1/0.8.0.2.0.0	tgt	CLAIMED	DEVICE	
disk	33	0/6/1/0.8.0.2.0.0.0	sdisk	CLAIMED	DEVICE	HITACHI DF600F
			/dev/dsk/c11t0d0		/dev/rdisk/c11t0d0	

**Figure 10-4: Verifying New Device Recognition
(HP-UX 11.31, Legacy DSF format)**

```
# ioscan -nfkN
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
=====						
fc	0	0/2/1/0	fcd	CLAIMED	INTERFACE	HP A6826-60001 2Gb
			/dev/fcd0			
tgtpath	4	0/2/1/0.0x50060e8010227240			estp	CLAIMED
TGT_PATH		fibre_channel target served by fcd driver				
lunpath	4	0/2/1/0.0x50060e8010227240.0x0			eslpt	CLAIMED
LUN_PATH		LUN path for ctl19				
lunpath	5	0/2/1/0.0x50060e8010227240.0x4000000000000000			eslpt	CLAIMED
LUN_PATH		LUN path for disk322				
fc	6	0/6/1/0	fcd	CLAIMED	INTERFACE	HP A6826-60001 2Gb
			/dev/fcd6			
tgtpath	7	0/6/1/0.0x50060e8010227242			estp	CLAIMED
TGT_PATH		fibre_channel target served by fcd driver				
lunpath	34	0/6/1/0.0x50060e8010227242.0x0			eslpt	CLAIMED
LUN_PATH		LUN path for ctl19				
lunpath	35	0/6/1/0.0x50060e8010227242.0x4000000000000000			eslpt	CLAIMED
LUN_PATH		LUN path for disk322				
esvroot	0	64000/0xfa00	esvroot	CLAIMED	VIRTBUS	Escsi virtual root
ctl	19	64000/0xfa00/0x18			esctl	CLAIMED
DEVICE		HITACHI DF600F				
			/dev/pt/pt19			
disk	322	64000/0xfa00/0x19			esdisk	CLAIMED
DEVICE		HITACHI DF600F				
			/dev/disk/disk322		/dev/rdisk/disk322	

**Figure 10-5: Verifying New Device Recognition
(HP-UX 11.31, for Persistent format)**

```
# ioscan -m dsf /dev/rdisk/c32t0d0
Persistent DSF      Legacy DSF(s)
=====
/dev/rdisk/disk322  /dev/rdisk/c32t0d0
```

**Figure 10-6: Verifying New Device Recognition for Multipath
(Legacy DSF Format)**

```
# ioscan -m dsf /dev/rdisk/disk322
Persistent DSF          Legacy DSF(s)
=====
/dev/rdisk/disk322      /dev/rdisk/c32t0d0
                        /dev/rdisk/c11t0d0
```

Figure 10-7: Verifying New Device Recognition for Multipath for Persistent Format

Troubleshooting

This chapter provides troubleshooting information for your array and instructions for calling technical support.

This chapter covers the following key topics:

- ❑ [Potential Error Conditions](#)
- ❑ [Calling the Hitachi Data Systems Support Center](#)

Potential Error Conditions

s lists potential error conditions that can occur when configuring the array with a host. It also provides instructions for resolving each condition. If you are unable to resolve an error condition, please contact your Hitachi Data Systems representative or VAR for help, or log on to the Hitachi Customer Support Center:

<https://extranet.hds.com/http://aim.hds.com/portal/dt>

Table 11-1: Error Conditions and Recommended Actions

Error Condition	Recommended Action
The array LUN devices are not recognized by the system.	<ul style="list-style-type: none">• Verify that the READY LEDs on the array are working properly.• Check that the Fibre Channel interface cables are correctly installed and firmly connected.• Recheck the Fibre Channel buses for new devices.<ul style="list-style-type: none">- For Windows, use Disk Management.- For Solaris, use the <code>format</code> command.- For Red Hat Enterprise Linux, Asianux, SLES, Apple, IBM, and HP-UX, use the <code>lsdev</code> command.
The array's red Alarm LED is ON.	Please contact the Hitachi Data Systems Support Center.

Calling the Hitachi Data Systems Support Center

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any error messages displayed on the host system(s).
- The exact content of any error messages displayed by Storage Navigator Modular 2.
- The Storage Navigator Modular 2 configuration information.

The Hitachi Data Systems customer support staff is available 24 hours/day, seven days a week. If you need technical support, please call:

- United States: (800) 446-0744
- Outside the United States: (858) 547-4526



Glossary

This glossary provides definitions of general storage networking terms as well as specific terms related to the technology that supports your array. Click the letter of the glossary section to display that page.

A

array

A set of hard disks mounted in a single enclosure and grouped logically together to function as one contiguous storage space.

B

BIOS

Basic Input Output System, built-in software code that determines the functions that a computing device can perform without accessing programs from a disk.

Bps

Bits per second, the standard measure of data transmission speeds.

C

capacity

The amount of information (usually expressed in megabytes) that can be stored on a disk drive. It is the measure of the potential contents of a device; the volume it can contain or hold. In communications, capacity refers to the maximum possible data transfer rate of a communications channel under ideal conditions.

channel

The link between the central processor and the peripherals. A channel can be the physical cabling that connects the nodes on a network, an electronic signal traveling over a pathway, or a sub-channel in a carrier frequency.

CLI

See command line interface.

cluster

A group of disk sectors. The operating system assigns a unique number to each cluster and then keeps track of files according to which clusters they use.

cluster capacity

The total amount of disk space in a cluster, excluding the space required for system overhead and the operating system. Cluster capacity is the amount of space available for all archive data, including original file data, metadata, and redundant data.

Glossary–2

command line interface (CLI)

A method of interacting with an operating system or software using a command line interpreter. With Hitachi's Storage Navigator Modular Command Line Interface, CLI is used to interact with and manage Hitachi storage and replication systems.

D

data volume

A volume that stores database information. Other files, such as index files and data dictionaries, store administrative information (metadata).

direct access storage device (DASD) fast write (DFW)

An attribute of record caching (while DASD Fast Write Access is a function of record caching) in which a specified record ID is placed in the cache and nonvolatile storage when a file-type macro is issued. If the cache is not available or the nonvolatile storage is not available, the record is written directly to the DASD surface.

disk array

An enterprise storage system containing multiple disk drives. Also referred to as "disk array device" or "disk storage system."

E

Ethernet

Local area networking technology, based on transmission packets that go between physical ports over different types of electrical and optical media.

F

fabric

The hardware that connects workstations and servers to storage devices in a SAN. The SAN fabric enables any-server-to-any-storage device connectivity through the use of Fibre Channel switching technology.

failover

The automatic substitution of a functionally equivalent system component for a failed one. The term failover is most often applied to intelligent controllers connected to the same storage devices and host computers. If one of the controllers fails, failover occurs, and the survivor takes over its I/O load.

FC

See Fibre Channel.

FCP

Fibre-Channel Protocol.

Fibre Channel

A gigabit-speed network technology primarily used for storage networking.

firmware

Software embedded into a storage device. It may also be referred to as Microcode.

G

Gbps

Gigabit per second.

GUI

Graphical user interface.

H

HBA

See Host Bus Adapter

High Availability (HA) software

An application designed for use during a primary host or disk failure. The software switches the failed host to a standby host (fail-over). High Availability software must be installed on the primary and secondary hosts.

Host Bus Adapter

An I/O adapter that connects a host I/O bus to the memory system of a computer.

I

I/O

Input/output.

Initiator

A system component that originates an I/O command over an I/O bus or network, such as an I/O adapters or network interface cards.

L

LAN

Local Area Network, a computer network that spans a relatively small area, such as a single building or group of buildings.

logical

Describes a user's view of the way data or systems are organized. The opposite of logical is physical, which refers to the real organization of a system. A logical description of a file is that it is a quantity of data collected together in one place. The file appears this way to users. Physically, the elements of the file could live in segments across a disk.

logical unit

See logical unit number.

logical unit number (LUN)

An address for an individual disk drive, and by extension, the disk device itself. Used in the SCSI protocol as a way to differentiate individual disk drives within a common SCSI target device, like a disk array. LUNs are normally not entire disk drives but virtual partitions (or volumes) of a RAID set.

logical volume

An area on a disk consisting of device files that are logically integrated using a volume manager.

LU

Logical unit.

LUN

See logical unit number.

LUN Manager

This storage feature is operated through Storage Navigator Modular 2 software and manages access paths among host and logical units for each port in your array.

M

microcode

The lowest-level instructions directly controlling a microprocessor. Microcode is generally hardwired and cannot be modified. It is also referred to as firmware embedded in a storage subsystem.

Middleware

Software that connects two otherwise separate applications. For example, a middleware product can be used to link a database system to a Web server. Using forms, users request data from the database; then, based on the user's requests and profile, the Web server returns dynamic Web pages to the user.

N

node

In networks, a node is a processing location. A node can be a computer or other device, such as a printer. Every node has a unique network address.

Q

queue depth

When a host queues successive commands to the array before execution of a previous command can complete, the number of times successive commands are issued is called queue depth. When two or more hosts are connected to a port of an array, the number of queue commands for the port is increased because the host issues commands to each array separately.

S

SAN

Storage Area Network, a network of shared storage devices that contain disks for storing data.

SCSI

Small Computer System Interface, a parallel interface standard that provides faster data transmission rates than standard serial and parallel ports.

SNM2

See Storage Navigator Modular 2.

SNMP

Simple Network Management Protocol, a protocol used to facilitate monitoring and management of clusters through an external interface. SNMP sends notifications to IP addresses whenever certain types of events occur.

Storage Navigator Modular 2

A multi-featured scalable storage management application that is used to configure and manage the storage functions of Hitachi arrays. Also referred to as "Navigator 2".

T

target port

A port-type which differs from an "Initiator Port" or "Remote Control Unit Target Port". The target is used without configuration of Fibre Remote Copy. It allows LOGIN of host computers, but does not allow LOGIN of MCUs.

V

volume

A disk array object that most closely resembles a physical disk from the operating environment's viewpoint. The basic unit of storage as seen from the host.

W

World Wide Name (WWN)

A unique identifier for an open systems host. It consists of a 64-bit physical address (the IEEE 48-bit format with a 12-bit extension and a 4-bit prefix). The WWN is essential for defining the SANTinel™ parameters because it determines whether the open systems host is to be allowed or denied access to a specified logical unit or a group of logical units.

A	B	C	D	E	F	G	H	I	M	N	Q	S	T	W
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Index

A

Array

connecting to

Apple [8-2](#)

Asianux [7-2](#)

HP-UX [10-2](#)

IBM [8-4](#)

Red Hat Enterprise Linux [5-2](#)

SLES [6-2](#)

Solaris [3-2](#)

VMware [4-2](#)

Windows [2-2](#)

verifying recognition of (Apple) [8-4](#)

Attaching a raw device (VMware) [4-20](#)

B

Boot from SAN [1-3](#)

C

Cluster configurations [1-4](#)

Asianux [7-4](#)

SuSE Linux [6-5](#)

Configuring host groups (VMware) [4-5](#)

Creating a Virtual Machine File System
(VMware) [4-12](#)

D

Devices

and disk parameters

HP-UX [10-4](#)

IBM [9-3](#)

Solaris [3-4](#)

changing parameters

from the AIX command line [9-5](#)

using SMIT [9-4](#)

configuring new

Asianux [7-4](#)

Windows [2-12](#)

discovering using ODM [9-14](#)

partitioning

Asianux [7-4](#)

Red Hat Enterprise Linux [5-6](#)

SLES [6-5](#)

Windows [2-15](#)

recognition

Asianux [7-5](#)

Red Hat Enterprise Linux [5-4](#)

SLES [6-3](#)

Disk number assignments (Windows) [2-4](#)

E

Emulex HBAs, setting queue depth

VMware [4-4](#)

Windows [2-11](#)

G

Guidelines

host bus adapters

HP-UX [10-2](#)

Solaris [3-2](#)

Windows [2-2](#)

operating system

HP-UX [10-2](#)

Solaris [3-2](#)

Windows [2-2](#)

H

High availability

cluster configurations for Asianux [7-4](#)

cluster configurations for SuSE Linux [6-5](#)

multipath configurations for Asianux [7-4](#)

multipath configurations for SuSE Linux [6-5](#)

High availability configurations

cluster [1-4](#)

multipath [1-4](#)

Host bus adapters

configuring

Apple [8-3](#)

Asianux [7-3](#)

Red Hat Enterprise Linux [5-3](#)

SLES [6-3](#)

- Windows [2-5](#)
- guidelines
 - HP-UX [10-2](#)
 - Solaris [3-2](#)
 - Windows [2-2](#)
- IBM [9-2](#)
- setting queue depth (VMware) [4-3](#)
- Host groups (VMware)
 - configuring [4-5](#)
 - creating and editing [4-8](#)
- Host server
 - connecting
 - Apple [8-2](#)
 - Asianux [7-2](#)
 - IBM [9-2](#)
 - Red Hat Enterprise Linux [5-2](#)
 - SLES [6-2](#)
 - Solaris [3-5](#)
 - Windows [2-3](#)
 - preparing
 - Apple [8-2](#)
 - Asianux [7-2](#)
 - IBM [9-2](#)
 - Red Hat Enterprise Linux [5-2](#)
 - SLES [6-2](#)
 - VMware [4-2](#)
 - Windows [2-2](#)

I

Installation prerequisites and tasks [1-2](#)

L

Logging in to Storage Navigator Modular 2 (VMware) [4-5](#)

M

Multipath configurations [1-4](#)

- Asianux [7-4](#)
- SuSE Linux [6-5](#)

Multi-path I/O (MPIO) (Windows) [2-6](#)

O

Object Data Manager (IBM)

- device
 - deleting [9-14](#)
 - discovering new [9-14](#)
- installing the 5.0.0.1 fibre update package [9-8, 9-13](#)
- installing the 5.0.0.4 fibre update package [9-10](#)
- overview [9-7](#)
- queue depth and read/write timeout values [9-15](#)
- references [9-16](#)
- uninstalling the fibre update package [9-12](#)

Operating system guidelines

- HP-UX [10-2](#)
- Solaris [3-2](#)

Windows [2-2](#)

P

Partitioning devices

- Asianux [7-4](#)
- Red Hat Enterprise Linux [5-6](#)
- SLES [6-5](#)
- Windows [2-15](#)

Potential error conditions [11-2](#)

Product support [11-2](#)

Q

QLogic HBAs, setting queue depth

- VMware [4-3](#)
- Windows [2-11](#)

Queue depth

- parameters
 - IBM [9-4, 9-5, 9-15](#)
 - Solaris [3-4](#)
 - VMware [4-4](#)
 - Windows [2-11](#)
- requirements
 - Asianux [7-3](#)
 - IBM [9-3](#)
 - Red Hat Enterprise Linux [5-3](#)
 - SLES [6-3](#)
 - Solaris [3-4](#)
 - Windows [2-5](#)

R

RAID groups in array (VMware) [4-5](#)

Raw devices

- VMware [4-20](#)

Read/write timeout values, ODM [9-15](#)

Recognizing new devices

- Asianux [7-5](#)
- Red Hat Enterprise Linux [5-4](#)
- SLES [6-3](#)

References for ODM [9-16](#)

S

Signatures (Windows) [2-12](#)

SMIT, using to change device parameters (IBM) [9-4](#)

Storage Navigator Modular 2 (VMware) [4-5](#)

Support Center [11-2](#)

V

Verifying host recognition of array (Apple) [8-4](#)

Verifying new devices

- Asianux [7-5](#)
- Red Hat Enterprise Linux [5-4](#)
- SLES [6-3](#)

Virtual Machine File System (VMware) [4-12](#)

W

Writing signatures (Windows) [2-12](#)

Index-2

Hitachi Data Systems

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com



MK-08DF8189-09