HITACHI
Inspire the Next

Open-Systems Host Attachment Guide

Hitachi Virtual Storage Platform G200, G400, G600, G800
Hitachi Virtual Storage Platform F400, F600, F800
Hitachi Virtual Storage Platform G1000, G1500
Hitachi Virtual Storage Platform F1500
Hitachi Virtual Storage Platform
Hitachi Unified Storage VM

Hitachi Data Systems

MK-90RD7037-09
October 2016

© 2014, 2016 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd., or Hitachi Data Systems Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Data Systems Corporation at https://support.hds.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi Data Systems Corporation agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Data Systems Corporation.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

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

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Contents

Preface ix
Intended audience x
Product version x
Release notes x
Document revision level xi
Changes in this revision xi
Referenced documents xii
Document conventions xiii
Convention for storage capacity values xv
Accessing product documentation xv
Getting help xvi
Comments xvi

Overview of host attachment 1-1
About the Hitachi RAID storage systems 1-2
Device types 1-3
Host queue depth 1-6
Host attachment workflow 1-7

Preparing for host attachment 2-1
Installation and configuration requirements 2-2
Installing the Hitachi RAID storage system 2-4
Configuring the Hitachi RAID storage system 2-5
Setting the system option modes 2-5
Configuring the ports 2-6
Setting the host modes and host mode options 2-7
Installing and configuring the host 2-8
Installing the host OS software 2-8
Installing the LVM software 2-8
Installing the failover software 2-9

Installing the SNMP software 2-10
Installing and configuring the host adapters 2-11
Connecting the Hitachi RAID storage system to the host 2-13
Configuring the new hosts and new LU paths 2-14

AIX® configuration and attachment 3-1

Hitachi storage system configuration for AIX® operations 3-2
Verifying new device recognition 3-3
Configuring the new devices 3-5
Changing the default device parameters 3-5
Changing device parameters from the AIX® command line 3-6
Assigning new devices to volume groups and setting partition sizes 3-7
Creating, mounting, and verifying file systems 3-10
Creating the file systems 3-10
Mounting and verifying file systems 3-14
Using the Object Data Manager with Hitachi RAID storage 3-15
Overview of ODM 3-15
ODM advantages and cautions 3-16
Using ODM 3-17
Discovering new devices 3-17
Deleting devices 3-17
Queue depth and read/write timeout values 3-17
Online device installation 3-18
Online LUSE configuration 3-19
Creating and mounting the file systems 3-19
Expanding the logical volume 3-21
Expanding the file system (up to 3 GB) 3-22
Increasing the file system (up to 40 GB) 3-22
Troubleshooting for AIX® host attachment 3-24

HP-UX configuration and attachment 4-1

Hitachi storage system configuration for HP-UX operations 4-2
Configuring the new devices 4-3
Verifying new device recognition 4-4
Verifying device files and the driver 4-7
Creating device files 4-8
Partitioning disk devices 4-10
Creating physical volumes 4-11
Creating volume groups 4-12
Creating logical volumes 4-15
Creating file systems 4-17
Setting device parameters 4-18
Setting the IO time-out parameter 4-18

Setting the queue depth parameter 4-19
Creating and registering the queue depth start-up script 4-21
Creating mount directories 4-25
Mounting and verifying file systems 4-26
Setting and verifying auto-mount parameters 4-27
Online device installation 4-28
Troubleshooting for HP-UX host attachment 4-29

Red Hat Linux configuration and attachment 5-1

Hitachi storage system configuration for Red Hat Linux operations 5-2
Device Mapper (DM) Multipath configuration 5-3
Verifying new device recognition 5-4
Configuring the new devices 5-6
Setting the number of logical units 5-6
Partitioning the devices 5-7
Creating, mounting, and verifying the file systems 5-8
Creating the file systems 5-8
Creating the mount directories 5-8
Mounting the new file systems 5-8
Verifying the file systems 5-8
Setting the auto-mount parameters 5-9
Troubleshooting for Red Hat Linux host attachment 5-10

Solaris configuration and attachment 6-1

Hitachi storage system configuration for Solaris operations 6-2
FCA configuration for Solaris 6-3
Verifying the FCA installation 6-3
Setting the disk and device parameters 6-4
Configuring the new devices 6-6
Setting and recognizing the LUs 6-6
Verifying recognition of new devices 6-9
Partitioning and labeling the new devices 6-10
Creating and mounting the file systems 6-26
Creating the file systems 6-26
Creating and verifying the mount directories 6-27
Mounting and verifying the file systems 6-28
Setting and verifying the auto-mount parameters 6-30
ALUA support on Solaris 6-32
Troubleshooting for Solaris host attachment 6-33
Verbose mode troubleshooting 6-33
Online device installation 6-35
Sun fibre-channel host bus adapter installation 6-35
Using MPxIO path failover software 6-36

SUSE Linux configuration and attachment 7-1
Hitachi storage system configuration for SUSE Linux operations 7-2
Device Mapper (DM) Multipath configuration..... 7-3
Verifying new device recognition..... 7-4
Configuring the new devices..... 7-6
Setting the number of logical units 7-6
Partitioning the devices 7-7
Creating, mounting, and verifying file systems..... 7-8
Creating file systems 7-8
Creating mount directories 7-8
Mounting new file systems 7-8
Verifying file systems..... 7-8
Setting auto-mount parameters 7-9
Troubleshooting for SUSE Linux host attachment 7-10

VMware configuration and attachment 8-1
Hitachi storage system configuration for VMware operations 8-2
VMware host configuration for Hitachi RAID storage..... 8-3
SAN configuration 8-3
VMware vSphere API operations..... 8-3
VMware ESX Server and VirtualCenter compatibility 8-3
Installing and configuring VMware..... 8-4
Creating and managing VMware infrastructure components..... 8-4
FCA configuration for VMware 8-6
Settings for QLogic adapters 8-6
Settings for Emulex adapters 8-7
Configuring the new devices..... 8-8
Creating the VMFS datastores 8-8
Adding a hard disk to a virtual machine 8-9
Troubleshooting for VMware host attachment 8-10

Windows configuration and attachment..... 9-1
Hitachi storage system configuration for Windows operations..... 9-2
Verifying the disk and device parameters..... 9-3
Verifying the disk I/O timeout value (TOV)..... 9-3
Verifying the queue depth 9-4
Verifying new device recognition..... 9-6
Configuring the new disk devices 9-8
Writing the signatures 9-8
Creating and formatting the partitions 9-11
Verifying file system operations..... 9-15
Verifying auto-mount 9-17

Changing the enable write caching option 9-18
Creating an online LUSE volume 9-19
Enabling MultiPath IO (MPIO) 9-23
Troubleshooting for Windows host attachment 9-29

XenServer configuration and attachment 10-1

Hitachi storage system configuration for XenServer operations 10-2
Recognizing the new devices 10-3
Creating storage repositories 10-4
Configuring the new storage devices for host use 10-6
Troubleshooting for XenServer host attachment 10-7

General troubleshooting 11-1

General troubleshooting 11-2
Contacting Hitachi Data Systems technical support 11-3

SCSI TID Maps for FC adapters A-1

Note on using Veritas Cluster Server B-1

Disk parameters for Hitachi disk types C-1

Parameter values for OPEN-*x* disk types C-1
Parameter values for VLL disk types C-3
Parameter values for LUSE disk types C-4
Parameter values for VLL LUSE disk types C-5
Parameter values for OPEN-8 disk types C-6

Host modes and host mode options D-1

Host modes and host mode options for USP V/VM D-1
Host modes and host mode options for VSP D-8
Host modes and host mode options for VSP G1000, G1500 and VSP F1500 ..... D-12
Host modes and host mode options for HUS VM D-17
Host modes and host mode options for VSP Gx00 and Fx00 ..... D-21

Acronyms and abbreviations

Preface

This document describes and provides instructions for installing and configuring the storage devices on the Hitachi RAID storage systems for attachment to open-systems hosts. The Hitachi RAID storage systems include the following models:

- Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP Gx00)
- Hitachi Virtual Storage Platform F400, F600, F800 (VSP Fx00)
- Hitachi Virtual Storage Platform G1000, G1500 (VSP G1000, G1500)
- Hitachi Virtual Storage Platform F1500 (VSP F1500)
- Hitachi Virtual Storage Platform (VSP)
- Hitachi Unified Storage VM (HUS VM)
- Hitachi Universal Storage Platform V (USP V)
- Hitachi Universal Storage Platform VM (USP VM)

Please read this document carefully to understand how to use this product, and maintain a copy for reference purposes.

☐ Intended audience
☐ Product version
☐ Release notes
☐ Document revision level
☐ Changes in this revision
☐ Referenced documents
☐ Document conventions
☐ Convention for storage capacity values
☐ Accessing product documentation
☐ Getting help
☐ Comments

Intended audience

This document is intended for system administrators, Hitachi Data Systems representatives, and authorized service providers who are install, configure, and operate the Hitachi RAID storage systems.

Readers of this document should be familiar with the following:

- Data processing and RAID storage systems and their basic functions.
- The Hitachi RAID storage system and the *Hardware Guide* or *User and Reference Guide* for the storage system.
- The management software for the storage system (for example, Hitachi Command Suite, Hitachi Device Manager - Storage Navigator) and the applicable user manual (for example, *Hitachi Command Suite User Guide*, *System Administrator Guide*, *Storage Navigator User Guide*).
- The host operating system (OS), the hardware hosting the system, and the hardware used to attach the Hitachi RAID storage system to the host, including fibre-channel cabling, host adapters, switches, and hubs.

Product version

This document revision applies to the following versions:

- SVOS 7.0
- VSP Gx00 and VSP Fx00: 83-04-0*x* or later
- VSP G1000, G1500 and VSP F1500: 80-05-0*x* or later
- HUS VM: 73-01-0*x* or later
- VSP: 70-01-0*x* or later
- USP V/VM: 60-05-0*x* or later

Release notes

The release notes for these products are available on Hitachi Data Systems Support Connect: https://knowledge.hds.com/Documents. Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document.

Document revision level

Revision	Date	Description
MK-90RD7037-00	Apr 2014	Initial release. ***Note:*** This document replaces the following documents: *Configuration Guide for HP-UX Host Attachment*, MK-96RD638 *Configuration Guide for IBM® AIX® Host Attachment*, MK-96RD636 *Configuration Guide for Red Hat Linux Host Attachment*, MK-96RD640 *Configuration Guide for Solaris Host Attachment*, MK-96RD632 *Configuration Guide for SUSE Linux Host Attachment*, MK-96RD650 *Configuration Guide for VMware Host Attachment*, MK-98RD6716 *Configuration Guide for Windows Host Attachment*, MK-96RD639 *Configuration Guide for XenServer Host Attachment*, MK-90RD6766
MK-90RD7037-01	Oct 2014	Supersedes and replaces revision 0
MK-90RD7037-02	Feb 2015	Supersedes and replaces revision 1
MK-90RD7037-03	Apr 2015	Supersedes and replaces revision 2
MK-90RD7037-04	Aug 2015	Supersedes and replaces revision 3
MK-90RD7037-05	Nov 2015	Supersedes and replaces revision 4
MK-90RD7037-06	Dec 2015	Supersedes and replaces revision 5
MK-90RD7037-07	Feb 2016	Supersedes and replaces revision 6
MK-90RD7037-08	May 2016	Supersedes and replaces revision 7
MK-90RD7037-09	Oct 2016	Supersedes and replaces revision 8

Changes in this revision

- Added the Hitachi Virtual Storage Platform G1500 and Hitachi Virtual Storage Platform F1500 storage systems.
- Updated the host mode option (HMO) information for VSP Gx00, VSP Fx00, VSP G1000, G1500, and VSP F1500:
  - Updated the description of HMO 88.
  - Added HMO 105 for Task Set Full response.
- Updated the SAN configuration information for VMware vSphere (SAN configuration).
- Added information about ALUA support on Solaris hosts (ALUA support on Solaris).
- Corrected references to the host mode and host mode option tables in Appendix D.

Referenced documents

Hitachi Virtual Storage Platform Gx00 and Fx00 documents:

- *Provisioning Guide for Hitachi VSP Gx00 and Fx00 Models*, MK-94HM8014
- *System Administrator Guide*, MK-94HM8016
- *VSP G200 Hardware Reference Guide*, MK-94HM8020
- *VSP G200 Hardware Installation and Reference Guide*, FE-94HM8020
- *VSP G400, G600 Hardware Reference Guide*, MK-94HM8022
- *VSP G400, G600 Hardware Installation and Reference Guide*, FE-94HM8022
- *VSP G800 Hardware Reference Guide*, MK-94HM8026
- *VSP G800 Hardware Installation and Reference Guide*, FE-94HM8026
- *VSP F400, F600 Hardware Reference Guide*, MK-94HM8045
- *VSP F400, F600 Hardware Installation and Reference Guide*, FE-94HM8045
- *VSP F800 Hardware Reference Guide*, MK-94HM8046
- *VSP F800 Hardware Installation and Reference Guide*, FE-94HM8046

Hitachi Virtual Storage Platform G1000, G1500, F1500 documents:

- *Provisioning Guide for Open Systems*, MK-92RD8014
- *Hardware Guide*, MK-92RD8007

Hitachi Unified Storage VM documents:

- *User and Reference Guide*, MK-92HM7005
- *Provisioning Guide*, MK-92HM7012
- *Storage Navigator User Guide*, MK-92HM7016
- *Storage Navigator Messages*, MK-92HM7017

Hitachi Virtual Storage Platform documents:

- *Provisioning Guide for Open Systems*, MK-90RD7022
- *Storage Navigator User Guide*, MK-90RD7027
- *Storage Navigator Messages*, MK-90RD7028
- *User and Reference Guide*, MK-90RD7042

Hitachi Universal Storage Platform V/VM documents:

- *Storage Navigator Messages*, MK-96RD613
- *LUN Manager User's Guide*, MK-96RD615
- *LUN Expansion (LUSE) User's Guide*, MK-96RD616
- *Storage Navigator User's Guide*, MK-96RD621
- *Virtual LVI/LUN and Volume Shredder User's Guide*, MK-96RD630
- *User and Reference Guide*, MK-96RD635

Hitachi Command Suite documents:

- *Hitachi Command Suite User Guide*, MK-90HC172
- *Hitachi Command Suite Administrator Guide*, MK-90HC175
- *Hitachi Command Suite Messages Guide*, MK-90HC178

Hitachi Dynamic Link Manager documents:

- *HDLM for AIX® User Guide*, MK-92DLM111
- *HDLM for HP-UX User Guide*, MK-92DLM112
- *HDLM for Linux User Guide*, MK-92DLM113
- *HDLM for Solaris User Guide*, MK-92DLM114
- *HDLM for Windows User Guide*, MK-92DLM129
- *HDLM for VMware User Guide*, MK-92DLM130

Document conventions

This document uses the following terminology conventions:

Convention	Description
Hitachi RAID storage system, storage system	Refers to all configurations and models of the Hitachi RAID storage systems, unless otherwise noted.
Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP G200, G400, G600, G800) VSP Gx00 models	Refers to all models of the Hitachi Virtual Storage Platform G200, G400, G600, G800 storage systems, unless otherwise noted.
Hitachi Virtual Storage Platform F400, F600, F800 (VSP F400, F600, F800) VSP Fx00 models	Refers to all models of the Hitachi Virtual Storage Platform F400, F600, F800 storage systems, unless otherwise noted.

This document uses the following typographic conventions:

Convention	Description
Bold	▪ Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click **OK**. ▪ Indicates emphasized words in list items.
Italic	▪ Indicates a document title or emphasized words in text. ▪ Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: `pairdisplay -g group` (For exceptions to this convention for variables, see the entry for angle brackets.)
`Monospace`	Indicates text that is displayed on screen or entered by the user. Example: `pairdisplay -g oradb`
< > angled brackets	Indicates variables in the following scenarios: ▪ Variables are not clearly separated from the surrounding text or from other variables. Example: `Status-<report-name><file-version>.csv` ▪ Variables in headings

Convention	Description
[ ] 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.
↓*value*↓ floor floor(*value*)	Floor function (round down *value* to the next integer)
↓*value*↓ ceiling ceiling(*value*)	Ceiling function (round up *value* to the next integer)
_ (underlined text)	Default value

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to important or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions or consequences (for example, disruptive operations).
	WARNING	Warns the user of severe conditions or consequences (for example, destructive operations).

Convention for storage capacity values

Physical capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 KB	1,000 (10^3) bytes
1 MB	1,000 KB or $1{,}000^2$ bytes
1 GB	1,000 MB or $1{,}000^3$ bytes
1 TB	1,000 GB or $1{,}000^4$ bytes
1 PB	1,000 TB or $1{,}000^5$ bytes
1 EB	1,000 PB or $1{,}000^6$ bytes

Logical capacity values (for example, logical device capacity, cache capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Open-systems: ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1{,}024^2$ bytes
1 GB	1,024 MB or $1{,}024^3$ bytes
1 TB	1024 GB or $1{,}024^4$ bytes
1 PB	1,024 TB or $1{,}024^5$ bytes
1 EB	1,024 PB or $1{,}024^6$ bytes

Accessing product documentation

Product documentation is available on Hitachi Data Systems Support Connect: https://knowledge.hds.com/Documents. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

Hitachi Data Systems Support Connect is the destination for technical support of products and solutions sold by Hitachi Data Systems. To contact technical support, log on to Hitachi Data Systems Support Connect for contact information: https://support.hds.com/en_us/contact-us.html.

Hitachi Data Systems Community is a new global online community for HDS customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. Join the conversation today! Go to http://community.hds.com, register, and complete your profile.

Comments

Please send us your comments on this document: doc.comments@hds.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Data Systems Corporation.

Thank you!

1

Overview of host attachment

This chapter provides an overview of the Hitachi RAID storage systems and open-systems host attachment:

- ☐ About the Hitachi RAID storage systems
- ☐ Device types
- ☐ Host queue depth
- ☐ Host attachment workflow

About the Hitachi RAID storage systems

The Hitachi RAID storage systems offer a wide range of storage and data services, including thin provisioning with Hitachi Dynamic Provisioning, application-centric storage management and logical partitioning, and simplified and unified data replication across heterogeneous storage systems. These storage systems provide the foundation for matching application requirements to different classes of storage and delivering critical services such as:

- Business continuity services
- Content management services (search, indexing)
- Nondisruptive data migration
- Volume management across heterogeneous storage arrays
- Thin provisioning
- Security services (immutability, logging, auditing, encryption, shredding)
- Data deduplication
- I/O load balancing
- Data classification
- File management services

The Hitachi RAID storage systems provide heterogeneous connectivity to support multiple concurrent attachment to a variety of host operating systems, including UNIX, Windows, VMware, Linux, and mainframe servers, enabling massive consolidation and storage aggregation across disparate platforms. The storage systems can operate with multi-host applications and host clusters, and are designed to handle very large databases as well as data warehousing and data mining applications that store and retrieve terabytes of data. The Hitachi RAID storage systems are compatible with most fibre-channel host bus adapters (HBAs), iSCSI adapters, and FC-over-ethernet (FCoE) converged network adapters (CNAs).

Hitachi RAID storage system models

This document applies to the following Hitachi RAID storage systems:

- Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP Gx00)
- Hitachi Virtual Storage Platform F400, F600, F800 (VSP Fx00)
- Hitachi Unified Storage VM (HUS VM)
- Hitachi Virtual Storage Platform G1000, G1500 (VSP G1000, G1500)
- Hitachi Virtual Storage Platform F1500 (VSP F1500)
- Hitachi Virtual Storage Platform (VSP)
- Hitachi Universal Storage Platform V/VM (USP V/VM)

For information about availability of Virtual Storage Platform F1500, contact your Hitachi Data Systems representative.

Device types

Table 1-1 lists and describes the types of logical devices (LDEVs) on the Hitachi RAID storage systems that can be configured and used by open-systems hosts. The LDEVs on the Hitachi RAID storage systems are defined to the host as SCSI disk devices, even though the interface is fibre channel or iSCSI. For information about configuring LDEVs other than OPEN-V, contact your Hitachi Data Systems representative.

Table 1-1 Logical devices provided by the Hitachi RAID storage systems

Device type	Description
OPEN-V devices	SCSI disk devices (VLL-based volumes) that do not have a predefined size.
OPEN-*x* devices	SCSI disk devices of predefined sizes: ▪ OPEN-3 (2.3 GB) ▪ OPEN-8 (6.8 GB) ▪ OPEN-9 (6.9 GB) ▪ OPEN-E (13.5 GB) ▪ OPEN-L (33 GB) For information on the use of these devices, contact your Hitachi Data Systems account team.
VLL devices (OPEN-*x* VLL)	Custom-size LUs that are configured by "slicing" a single LU into two or more smaller LUs to improve host access to frequently used files. VLL devices are configured using the Virtual LVI/LUN (VLL) feature. The product name for OPEN-*x* VLL devices is OPEN-*x*-CVS, in which CVS stands for *custom volume size*. OPEN-L devices do not support VLL.
LUSE devices (OPEN-*x***n*)	Combined LUs composed of multiple OPEN-*x* devices. LUSE devices are configured using the LUN Expansion (LUSE) feature. A LUSE device can be from 2 to 36 times larger than a fixed-size OPEN-*x* LU. LUSE devices are designated as OPEN-*x***n*, where *x* is the LU type and 2< *n* < 36. For example, a LUSE device created by combining 10 OPEN-3 LUs is designated as an OPEN-3*10 device. LUSE lets the host access the data stored on the Hitachi RAID storage system using fewer LU numbers. ***Note:*** LUSE devices are not supported on the VSP G1000, VSP G1500, VSP F1500, VSP Gx00, or VSP Fx00 storage systems.
VLL LUSE devices (OPEN-*x***n* VLL)	Combined LUs composed of multiple VLL devices. VLL LUSE devices are configured first using the Virtual LVI/LUN feature to create custom-size devices and then using the LUSE feature to combine the VLL devices. You can combine from 2 to 36 VLL devices into one VLL LUSE device. For example, an OPEN-3 LUSE volume created from 10 OPEN-3 VLL volumes is designated as an OPEN-3*10 VLL device (product name OPEN-3*10-CVS).
FX devices (3390-3A/B/C, OPEN-*x*-FXoto)	The Hitachi Cross-OS File Exchange (FX) feature allows you to share data across mainframe and open-systems platforms using special multiplatform volumes called *FX devices*. FX devices are installed and accessed as raw devices (not SCSI disk devices). Windows hosts must use FX to access the FX devices as raw devices (no file system, no mount operation). The 3390-3B devices are write-protected from Windows host access. The Hitachi RAID storage system rejects all Windows host write operations (including FC adapters) for 3390-3B devices. The 3390-3A/C and OPEN-*x*-FXoto devices are not write-protected for Windows host access. Do not execute any write operations on these devices. Do not create a partition or file system on these devices. This will overwrite data on the FX device and prevent the Cross-OS File Exchange software from accessing the device. The VLL feature can be applied to FX devices for maximum flexibility in volume size. For more information about Hitachi Cross-OS File Exchange, see the *Hitachi Cross-OS File Exchange User Guide*, or contact your Hitachi Data Systems account team.

Table 1-2 lists the specifications for the LDEVs on the Hitachi RAID storage systems. The sector size for the devices is 512 bytes.

Table 1-2 Device specifications

Device type (Note 1)	Category (Note 2)	Product name (Note 3)	# of blocks (512 B/blk)	# of cylinders	# of heads	# of sectors per track	Capacity (MB) (Note 4)
OPEN-3	SCSI disk	OPEN-3	4806720	3338	15	96	2347
OPEN-8	SCSI disk	OPEN-8	14351040	9966	15	96	7007
OPEN-9	SCSI disk	OPEN-9	14423040	10016	15	96	7042
OPEN-E	SCSI disk	OPEN-E	28452960	19759	15	96	13893
OPEN-L	SCSI disk	OPEN-L	71192160	49439	15	96	34761
OPEN-3*n	SCSI disk	OPEN-3*n	4806720*n	3338*n	15	96	2347*n
OPEN-8*n	SCSI disk	OPEN-8*n	14351040*n	9966*n	15	96	7007*n
OPEN-9*n	SCSI disk	OPEN-9*n	14423040*n	10016*n	15	96	7042*n
OPEN-E*n	SCSI disk	OPEN-E*n	28452960*n	19759*n	15	96	13893*n
OPEN-L*n	SCSI disk	OPEN-L*n	71192160*n	49439*n	15	96	34761*n
OPEN-3 VLL	SCSI disk	OPEN-3-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-8 VLL	SCSI disk	OPEN-8-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-9 VLL	SCSI disk	OPEN-9-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-E VLL	SCSI disk	OPEN-E-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-V VLL	SCSI disk	OPEN-V	***Note 5***	***Note 6***	15	128	***Note 7***
OPEN-3*n VLL	SCSI disk	OPEN-3*n-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-8*n VLL	SCSI disk	OPEN-8*n-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-9*n VLL	SCSI disk	OPEN-9*n-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-E*n VLL	SCSI disk	OPEN-E*n-CVS	***Note 5***	***Note 6***	15	96	***Note 7***
OPEN-V*n VLL	SCSI disk	OPEN-V*n	***Note 5***	***Note 6***	15	128	***Note 7***
3390-3A	FX otm/mto	3390-3A	5820300	3345	15	116	2844
3390-3B	FXmto	3390-3B	5816820	3343	15	116	2844
3390-3C	FXotm	OP-C-3390-3C	5820300	3345	15	116	2844
FX OPEN-3	FXoto	OPEN-3	4806720	3338	15	96	2347
3390-3A VLL	FX otm/mto	3390-3A-CVS	***Note 5***	***Note 6***	15	116	***Note 7***
3390-3B VLL	FXmto	3390-3B-CVS	***Note 5***	***Note 6***	15	116	***Note 7***
3390-3C VLL	FXotm	OP-C-3390-3C-CVS	***Note 5***	***Note 6***	15	116	***Note 7***
FX OPEN-3 VLL	FXoto	OPEN-3-CVS	***Note 5***	***Note 6***	15	96	***Note 7***

Notes:

1. The availability of specific device types depends on the storage system model and the level of microcode installed on the storage system.

2. The category of a device (SCSI disk or Cross-OS File Exchange) determines its volume usage. SCSI disk devices (for example, OPEN-V) are usually formatted with file systems but can also be used as raw devices (for example, some applications use raw devices).

3. The product name for Virtual LVI/LUN devices is OPEN-*x*-CVS, where CVS = custom volume size. The command device (used for Command Control Interface operations) is distinguished by **–CM** on the product name (for example, OPEN-V-CM).

4. This capacity is the maximum size that can be entered. The device capacity can sometimes be changed by the BIOS or host adapter. Also, different capacities may be due to variations such as 1 MB = 1000^2 bytes or 1024^2 bytes.

5. The number of blocks for a Virtual LVI/LUN volume is calculated as follows:

# of blocks = (# of data cylinders) × (# of heads) × (# of sectors per track)

The number of sectors per track is 128 for OPEN-V and 96 for the other emulation types.

Example: For an OPEN-3 VLL volume with capacity = 37 MB:

# of blocks = (53 cylinders – see Note 3) × (15 heads) × (96 sectors per track) = 76320

6. The number of data cylinders for a Virtual LVI/LUN volume is calculated as follows (↑**...**↑ means that the value should be rounded up to the next integer):

- Number of data cylinders for OPEN-*x* VLL volume (except for OPEN-V) =
  # of cylinders = ↑ (capacity (MB) × 1024/720 ↑

  Example: For OPEN-3 VLL volume with capacity = 37 MB:

  # of cylinders = ↑37 × 1024/720↑ = ↑52.62↑ = 53 cylinders
- Number of data cylinders for an OPEN-V VLL volume =
  # of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑

  Example: For OPEN-V VLL volume with capacity = 50 MB:

  # of cylinders = ↑50 × 16/15↑ = ↑53.33↑ = 54 cylinders

7. The size of an OPEN-*x* VLL volume is specified by capacity in MB, not number of cylinders. The size of an OPEN-V VLL volume can be specified by capacity in MB or number of cylinders.

Host queue depth

Each operating system chapter in this document describes the specific configuration files and file format syntax required to configure the queue depth settings on your Hitachi RAID storage systems. The requirements for host queue depth depend on the Hitachi RAID storage system model.

- **USP V/VM (and earlier).** The Universal Storage Platform V/VM requires that the host queue depth (or max tag count) be set appropriately due to the queue depth limits of 32 per LUN and 2048 per port. This is because each MP in the USP V/VM can process a maximum of 4096 I/Os, and each MP manages two ports.
- **VSP, HUS VM, VSP G1000, VSP G1500, VSP F1500, VSP Gx00, VSP Fx00.** Due to their advanced architecture, the I/O limit per MP in these storage systems has increased substantially. However, while the technical limit to queue depth is much higher, the appropriate queue depth settings for each operational environment must be carefully researched and determined.

  To ensure smooth processing at the ports and best average performance, the recommended queue depth setting (max tag count) for these storage systems is 2048 per port and 32 per LDEV. Other queue depth settings, higher or lower than these recommended values, can provide improved performance for certain workload conditions.



Caution: Higher queue depth settings (greater than 2048 per port) can impact host response times, so caution must be exercised in modifying the recommended queue depth settings.

Host attachment workflow

1. Install the new Hitachi RAID storage system, or install the new physical storage devices on the existing Hitachi RAID storage system. This task is performed by the Hitachi Data Systems representative. See Installing the Hitachi RAID storage system.
2. Configure the Hitachi RAID storage system for host attachment. This task is performed by the Hitachi Data Systems representative and the user. See Configuring the Hitachi RAID storage system.
3. Configure the host for connection to the Hitachi RAID storage system, including host OS, middleware, and SNMP. This task is performed by the user. See Installing and configuring the host.
4. Install and configure the FC adapters for connection to the Hitachi RAID storage system. This task is performed by the user. See Installing and configuring the host adapters.
5. Connect the Hitachi RAID storage system to the host. This task is performed by the Hitachi Data Systems representative and the user. See Connecting the Hitachi RAID storage system to the host.
6. Configure the newly attached hosts and LU paths. This task is performed by the user. See Configuring the new hosts and new LU paths.
7. Configure the new storage devices for use on the host. This task is performed by the user. See the following chapters:
   - AIX® configuration and attachment
   - HP-UX configuration and attachment
   - Red Hat Linux configuration and attachment
   - Solaris configuration and attachment
   - SUSE Linux configuration and attachment
   - VMware configuration and attachment
   - Windows configuration and attachment
   - XenServer configuration and attachment

2

Preparing for host attachment

This chapter describes how to install and configure the Hitachi RAID storage system, host, and host adapters in preparation for host attachment.

- ☐ Installation and configuration requirements
- ☐ Installing the Hitachi RAID storage system
- ☐ Configuring the Hitachi RAID storage system
- ☐ Installing and configuring the host
- ☐ Installing and configuring the host adapters
- ☐ Connecting the Hitachi RAID storage system to the host
- ☐ Configuring the new hosts and new LU paths

Installation and configuration requirements

Table 2-1 lists the requirements for installing and configuring the Hitachi RAID storage system for attachment to an open-systems host server.

Table 2-1 Installation and configuration requirements

Item	Requirements
Hitachi RAID storage system	▪ The availability of features and devices depends on the Hitachi RAID storage system model and the level of microcode installed on the storage system. ▪ The Hitachi Storage Navigator software must be installed and operational. For details, see the *System Administrator Guide* or the *Storage Navigator User Guide* for the storage system. ▪ The Hitachi LUN Manager feature must be enabled. For details, see the *System Administrator Guide* or the *Storage Navigator User Guide* for the storage system.
Host server hardware	▪ Review the hardware requirements for attaching new storage to the host server. For details, see the user documentation for the host server. ▪ For details about supported host server hardware, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html
Hardware for host attachment	▪ For details about supported hardware for host attachment (optical cables, hubs, switches, and so on), see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html
Host operating system	▪ This document covers the following host platforms. Check the Hitachi Data Systems interoperability site for the latest information about host OS support. – AIX – HP-UX – Red Hat Linux – Solaris – SUSE Linux – VMware – Windows – XenServer ▪ Verify that the OS version, architecture, relevant patches, and maintenance levels are supported by the Hitachi RAID storage system. For details about supported OS versions, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html ▪ Verify that the host meets the latest system and software requirements for attaching new storage. For details, see the host OS user documentation. ▪ Verify that you have the host OS software installation media. ▪ Verify that you have root/administrator login access to the host system.

Item	Requirements
Host adapters (HBAs and CNAs)	▪ **HBAs:** The Hitachi RAID storage systems support FC HBAs equipped as follows: – 8-Gbps FC interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors. – 4-Gbps FC interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors. – 2-Gbps FC interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors. – 1-Gbps FC interface, including shortwave non-OFC optical interface and multimode optical cables with SC connectors. For OM3 fiber and 200-MB/s data transfer rate, the total cable length attached to each FC HBA must not exceed 500 meters (1,640 feet). Do not connect any OFC type connectors to the Hitachi RAID storage system. ▪ **iSCSI HBAs:** The VSP G1000, VSP G1500, VSP F1500, VSP Gx00, and VSP Fx00 storage systems support iSCSI adapters, with the following iSCSI SAN requirements: – 10 Gigabit Ethernet switch – 10 Gb NIC or HBA card in each host computer – 10 Gb iSCSI initiator – LC-LC optical cables VSP G1000, VSP G1500, VSP F1500: – Minimum microcode level: 80-03-3*x* For details, see the *Hardware Guide* for your storage system model. ▪ **CNAs:** The VSP G1000, VSP G1500, VSP F1500, and VSP storage systems support FCoE converged network adapters (CNAs) equipped as follows: – 10 Gbps fibre-channel over Ethernet interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors. For OM3 fiber and 10-Gb/s transfer rate, the total cable length attached to each CNA must not exceed 300 meters (984 feet). The diskless storage system model (no internal drives) does not support the FCoE option. VSP G1000, VSP G1500, VSP F1500: – Minimum microcode level: 80-02-0x – Host OS: Red Hat Enterprise Linux, VMware, Windows VSP: – Host OS: VMware, Windows ▪ For details about installing the adapter and using the utilities and tools for the adapter, see the user documentation for the adapter. ▪ For details about supported host adapters and drivers, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html
Storage area network (SAN)	▪ A SAN may be required to connect the Hitachi RAID storage system to the host. For details about supported switches, topology, and firmware versions for SAN configurations, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html

Installing the Hitachi RAID storage system

The Hitachi RAID storage systems come with all hardware and cabling required for installation. The Hitachi Data Systems representative follows the instructions and precautions in the Maintenance Manual for the storage system when installing the product. The installation tasks include:

- Checking all specifications to ensure proper installation and configuration.
- Installing and assembling all hardware and cabling.
- Verifying that the Storage Navigator software is installed and ready for use. For details, see the *System Administrator Guide* for the storage system (for VSP, HUS VM, and USP V/VM, see the *Storage Navigator User's Guide*).
- Installing and formatting the logical devices (LDEVs). The user provides the desired parity group and LDEV configuration information to the Hitachi Data Systems representative. For details, see the *Provisioning Guide* for the storage system (for USP V/VM see the manuals for LUN Manager, LUN Expansion, and Virtual LVI/LUN).

Configuring the Hitachi RAID storage system

Complete the following tasks to configure the Hitachi RAID storage system for attachment to the host server:

- ☐ Setting the system option modes
- ☐ Configuring the ports
- ☐ Setting the host modes and host mode options

Setting the system option modes

To provide greater flexibility, the Hitachi RAID storage systems have additional operational parameters called *system option modes* (SOMs) that allow you to tailor the storage system to your unique operating requirements. The SOMs are set on the service processor by the Hitachi Data Systems representative.

To set and manage the SOMs

1. Review the list of SOMs in the hardware guide for your storage system:
   - *VSP G200 Hardware Reference Guide*, MK-94HM8020
   - *VSP G400, G600 Hardware Reference Guide*, MK-94HM8022
   - *VSP G800 Hardware Reference Guide*, MK-94HM8026
   - *VSP F400, F600 Hardware Reference Guide*, MK-94HM8045
   - *VSP F800 Hardware Reference Guide*, MK-94HM8046
   - *VSP G1000, G1500 and VSP F1500 Hardware Guide*, MK-92RD8007
   - *VSP User and Reference Guide*, MK-90RD7042
   - *HUS VM Block Module Hardware User Guide*, MK-92HM7005
   - *USP V/VM User and Reference Guide*, MK-96RD635
   - *USP/NSC User and Reference Guide*, MK-94RD231
2. Work with your Hitachi Data Systems team to ensure that the appropriate SOMs for your operational environment are set on your storage system.
3. Check each new revision of the hardware guide for SOM changes that may apply to your operational environment, and contact your Hitachi Data Systems representative as needed.

Configuring the ports

Before the storage system is connected to the host, you must configure the ports on the Hitachi RAID storage system. Select the appropriate settings for each port based on the device to which the port is connected. The settings include attribute, security, speed, address, fabric, and connection type. For the latest information about port topology configurations supported by OS versions and adapter/switch combinations, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html

For details on configuring the ports, see the *Provisioning Guide* for the storage system (or the *LUN Manager User's Guide* for the USP V/VM).



Note:

- If you plan to use LUN security, enable the security setting now before the port is attached to the host. If you enable LUN security on a port when host I/O is in progress, I/Os will be rejected with a security guard after LUN security is enabled.
- If you plan to connect different types of servers to the RAID storage system via the same fabric switch, use the zoning function of the fabric switch.

Setting the host modes and host mode options

Before the storage system is connected to the hosts, you must configure the host groups or iSCSI targets for the new hosts and set the host mode and host mode options (HMOs) for each host group/iSCSI target. When you connect multiple hosts of different platforms to a single port, you must group hosts connected to the storage system by host groups/iSCSI targets that are segregated by platform. For example, if VMware, Windows, and Solaris hosts will be connected to a single port, you must create a host group/iSCSI target for each platform and set the host mode and HMOs for each host group/iSCSI target. When the storage system is connected to the hosts, you will register the hosts in the appropriate host groups/iSCSI targets.

While a host group can include more than one WWN, it is recommended that you create one host group for each host adapter and name the host group the same as the nickname for the adapter. Creating one host group per host adapter provides flexibility and is the only supported configuration when booting hosts from a SAN.

For instructions on setting the host modes and HMOs, see the *Provisioning Guide* for the storage system (or the *LUN Manager User's Guide* for the USP V/VM). For lists of the host modes and HMOs for the storage systems, see Host modes and host mode options. ***Important:*** There are differences in host mode and HMO support among the Hitachi storage system models, so it is important that you refer to the lists for your specific storage system model.



WARNING:

- Changing host modes or HMOs on a Hitachi RAID storage system that is already installed and attached to the host is disruptive and requires the host server to be rebooted.
- Before setting any HMO, review its functionality carefully to determine whether it can be used for your configuration and environment. If you have any questions or concerns, contact your Hitachi Data Systems representative.

Installing and configuring the host

This section describes general host configuration tasks that must be performed before the Hitachi RAID storage system is attached to the host server.

- ☐ Installing the host OS software
- ☐ Installing the LVM software
- ☐ Installing the failover software
- ☐ Installing the SNMP software



Note: The user is responsible for configuring the host system as needed for the new storage devices.

- For assistance with host configuration, see the user documentation for the product or contact the vendor's technical support.
- For assistance with specific configuration issues related to the Hitachi RAID storage system, contact your Hitachi Data Systems representative.

Installing the host OS software

The host operating system (OS) software must be loaded, configured, and operational before the Hitachi RAID storage system is attached.

1. Verify that the OS version, architecture, relevant patches, and maintenance levels are supported by the Hitachi RAID storage system. For details about supported OS versions, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html
2. Verify that the host meets the latest system and software requirements for attaching new storage. For details, see the user documentation for the OS.
3. Verify that you have the host OS software installation media.
4. Verify that you have root/administrator login access to the host system.

Installing the LVM software

The Hitachi RAID storage systems support industry-standard products and functions that provide logical volume management (LVM). You must configure the LVM products on the host servers to recognize and operate with the new storage devices before the new storage is attached. For assistance with LVM operations, see the user documentation for the LVM software or contact the vendor's technical support.

Installing the failover software

The Hitachi RAID storage systems support industry-standard products and functions that provide host, application, and path failover. You should configure the failover products to recognize and operate with the new storage devices before the new storage is attached.

- Supported host and application failover products include High Availability Cluster Multi-Processing (HACMP), Veritas Cluster Server, Sun Cluster, Microsoft Cluster Server (MSCS), and MC/ServiceGuard.
- Supported path failover products include Hitachi Dynamic Link Manager (HDLM), Veritas Volume Manager, DM Multipath, XenCenter dynamic multipathing, and HP-UX alternate link path failover.

For assistance with failover operations, see the user documentation for the failover product or contact the vendor's technical support.

For details about HDLM, see the HDLM User's Guide for the host platform (for example, *Hitachi Dynamic Link Manager User's Guide for Windows*), or contact your Hitachi Data Systems representative.



Note: Failover products may not provide a complete disaster recovery or backup solution and are not a replacement for standard disaster recovery planning and backup/recovery.

Installing the SNMP software

The Hitachi RAID storage systems support the industry-standard simple network management protocol (SNMP) for remote storage system management from the host servers. You must configure the SNMP software on the host before the new storage is attached. For assistance with SNMP configuration on the host, see the SNMP user documentation or contact the vendor's technical support.

SNMP is a part of the TCP/IP protocol suite that supports maintenance functions for storage and communication devices. The Hitachi RAID storage systems use SNMP to transfer status and management commands to the SNMP Manager on the host (see Figure 2-1). When the SNMP manager requests status information or when a service information message (SIM) occurs, the SNMP agent on the storage system notifies the SNMP manager on the host. Notification of error conditions is made in real time, enabling you to monitor the storage system from the open-systems host.

When a SIM occurs, the SNMP agent initiates trap operations, which alert the SNMP manager of the SIM condition. The SNMP manager receives the SIM traps from the SNMP agent and can request information from the SNMP agent at any time.

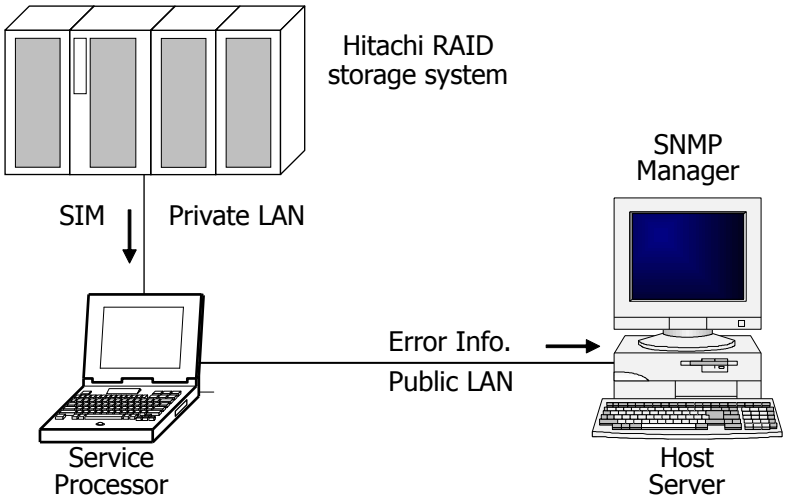


Figure 2-1 SNMP Environment

Installing and configuring the host adapters

The host adapters must be installed on the host before the Hitachi RAID storage system is attached. You also need to discover and write down the WWNs of the adapters to be connected to the storage system.

iSCSI (VSP Gx00, VSP Fx00, VSP G1000, VSP G1500, VSP F1500): Follow the instructions in your vendor documentation for preparing your hosts, HBAs, CNAs, NICs, and iSCSI initiators for use with the storage system. For iSCSI specifications and requirements, see the *Hardware Guide* for your storage system model.



Note: The user is responsible for installing and configuring the adapters as needed for the new storage devices.

- For assistance with host adapter configuration, see the user documentation for the adapter or contact the vendor's technical support.
- For assistance with specific configuration issues related to the Hitachi RAID storage system, contact your Hitachi Data Systems representative.

To install the host adapters:

1. **Verify interoperability.** Verify that the host adapters are supported by the Hitachi RAID storage system. For details, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html
2. **Install and verify the adapters.** Install the host adapters on the host server, and verify that the adapters are functioning properly. For details about installing the adapter and using the utilities for the adapter, see the user documentation for the adapter.

   Note:

   - Do not connect OFC-type FC interfaces to the Hitachi storage system.
   - If a switch or adapter with a 1-Gbps transfer rate is used, configure the device to use a fixed 1-Gbps setting instead of Auto Negotiation. Otherwise, it may prevent a connection from being established. However, the transfer speed of CHF port cannot be set as 1 Gbps when the CHF model type is 8US/8UFC/16UFC. Therefore 1-Gbps adapter and switch cannot be connected.
3. **Configure the adapter.** Use the setup utilities to configure the adapters to be connected to the Hitachi RAID storage system. The adapters have many configuration options. The minimum requirements for configuring the adapters for operation with the Hitachi RAID storage system are:
   - **I/O timeout value (TOV).** The disk I/O timeout value (TOV) requirement for the Hitachi storage system is **60 seconds** (0×3c hex).
   - **Queue depth.** The queue depth requirements for the Hitachi storage system devices are listed below. You can adjust the queue depth for the devices later as needed (within the specified range) to optimize the I/O performance of the devices. For details, see Host queue depth.

Parameter	Recommended value for VSP G1000, VSP G1500, VSP F1500, VSP Gx00, VSP Fx00, HUS VM, VSP	Required value for USP V/VM
Queue depth per LU	32 per LU	≤ 32 per LU
Queue depth per port	2048 per port	≤ 2048 per port

- **BIOS.** The BIOS may need to be disabled to prevent the system from trying to boot from the storage system.

Use the same settings and device parameters for all devices on the Hitachi RAID storage system. Several other parameters (for example, FC fabric) may also need to be set. Refer to the user documentation for the host adapter to determine whether other options are required to meet your operational requirements.

4. **Record the WWNs of the adapters.** Find and write down the WWN of each host adapter. You will need to enter these WWNs when you configure the new hosts on your storage system.

   For details about finding the WWN of an adapter, see the user documentation for the adapter. The method for finding the WWN varies depending on the adapter type, host platform, and topology. You can use the adapter utility (for example, the LightPulse Utility for Emulex), or the host OS (for example, the **`dmesg |grep Fibre`** command in Solaris), or the fabric switch connected to the host (for example, an AIX® host).

Connecting the Hitachi RAID storage system to the host

After the Hitachi RAID storage system and host have been configured, the Hitachi RAID storage system can be physically connected to the host system. Some of the steps in this procedure are performed by the Hitachi Data Systems representative, and some are performed by the user.



Note: The Hitachi Data Systems representative must use the ***Maintenance Manual*** for the storage system during all installation activities. Follow all precautions and procedures in the ***Maintenance Manual***, and always check all specifications to ensure proper installation and configuration.

To connect the Hitachi RAID storage system to the host system:

1. **Verify the storage system installation.** The Hitachi Data Systems representative verifies the configuration and operational status of the Hitachi RAID storage system ports, LDEVs, and paths.
2. **Shut down and power off the host.** The user shuts down and powers off the host. The power must be off when the FC/FCoE/iSCSI cables are connected.
3. **Connect the Hitachi RAID storage system to the host system.** The Hitachi Data Systems representative connects the cables between the Hitachi RAID storage system and the host or switch. Verify the ready status of the storage system and peripherals.
4. **Power on and boot the host system.** The user powers on and boots the host system after the storage system has been connected:
   - Power on the host system display.
   - Power on all peripheral devices. The Hitachi RAID storage system must be on, and the ports and modes must be configured before the host is powered on. If the ports are configured after the host is powered on, the host may need to be restarted to recognize the new settings.
   - Confirm the ready status of all peripheral devices, including the Hitachi RAID storage system.
   - Power on and boot the host system.

Configuring the new hosts and new LU paths

After discovering the WWNs of the host adapters and connecting the storage system to the host, you need to configure the new hosts and new LU paths on the Hitachi RAID storage system.

FC: To configure the newly attached hosts and LUs:

1. **Add new hosts.** Before you can configure LU paths, you must register the new hosts in host groups/iSCSI targets. For details, see the *Provisioning Guide* for the storage system (*LUN Manager User's Guide* for USP V/VM).

   When registering hosts in multiple host groups, set the security switch (LUN security) to enabled, and then specify the WWN of the host adapter.
2. **Configure LU paths.** Configure the LU paths for the newly attached storage devices, including defining primary LU paths and alternate LU paths and setting the UUID. For details, see the *Provisioning Guide* for the storage system (*LUN Manager User's Guide* for USP V/VM).
3. **Set fibre-channel authentication.** Set fibre-channel authentication as needed on host groups, ports, and fabric switches of the storage system. For details, see the *Provisioning Guide* for the storage system (*LUN Manager User's Guide* for USP V/VM).

iSCSI (VSP Gx00, VSP Fx00, VSP G1000, VSP G1500, VSP F1500): For details about iSCSI network configuration (for example, registering hosts in iSCSI targets, adding CHAP users, defining LU paths), see the *Provisioning Guide* for the storage system.

After configuring the newly attached hosts and LUs, you are ready to configure the new storage devices for use on the host system. For details, see the following chapters:

- AIX® configuration and attachment
- HP-UX configuration and attachment
- Red Hat Linux configuration and attachment
- Solaris configuration and attachment
- SUSE Linux configuration and attachment
- VMware configuration and attachment
- Windows configuration and attachment
- XenServer configuration and attachment

3

AIX® configuration and attachment

This chapter describes how to configure and manage the new Hitachi disk devices on an AIX® host:

- ☐ Hitachi storage system configuration for AIX® operations
- ☐ Verifying new device recognition
- ☐ Configuring the new devices
- ☐ Using the Object Data Manager with Hitachi RAID storage
- ☐ Online device installation
- ☐ Online LUSE configuration
- ☐ Troubleshooting for AIX® host attachment



Note: Configuration of the devices should be performed by the AIX® system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Hitachi storage system configuration for AIX® operations

The storage system must be fully configured before being attached to the AIX® host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for AIX® operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)
- Cross-OS File Exchange (FX) (3390-3A/B/C, OPEN-*x*-FXoto)

Host mode. The required host mode for AIX® is **0F**. Do not select a host mode other than **0F** for IBM AIX. For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Verifying new device recognition

The first step after attaching to the AIX® host is to verify that the host system recognizes the new devices. The host system automatically creates a device file for each new device recognized.

The devices should be installed and formatted with the fibre ports configured before the host system is powered on. Enter the **`cfgmgr`** command to check 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 (see Figure 3-1):

   `lsdev -C -c disk`
3. Verify that the host system recognizes all new disk devices, including OPEN-*x*, LUSE, VLL, VLL LUSE, and FX devices. The devices are listed by device file name.
4. Record the following device data for each new device: device file name, bus number, TID, LUN, and device type. Table 3-1 shows a sample worksheet for recording the device data. You need this information in order to change the device parameters.



Note: When you create the FX volume definition file (`datasetmount.dat`), provide the device file names for the FX devices. For example, if hdisk3 is a 3390-3B FX device, the entry for this volume in the FX volume definition file is: `\\.\PHYSICALDRIVE3 XXXXXX 3390-3B` (where `XXXXXX` is the VOLSER)

```
# lsdev -C -c disk                                          ← Display device data.
hdisk0  Available 10-68-00-0,0 16 Bit SCSI Disk Drive
hdisk1  Available 00-01-00-2,0  Hitachi Disk Array (Fibre)  ← New device.
hdisk2  Available 00-01-00-2,1  Hitachi Disk Array (Fibre)  ← New device.
   ↖ Device file name = hdiskx.
 :
#
```

This example shows the following information:

- The device hdisk1 is TID=2, LUN=0 on bus 1.
- The device hdisk2 is TID=2, LUN=1 on bus 1.

Figure 3-1 Verifying new device recognition on AIX host

Table 3-1 Device data table for AIX

Device File Name	Bus No.	TID	LUN	Device Type	Alternate Paths	
hdisk1					TID:____ LUN:____	TID:____ LUN:____
hdisk2					TID:____ LUN:____	TID:____ LUN:____
hdisk3					TID:____ LUN:____	TID:____ LUN:____
hdisk4					TID:____ LUN:____	TID:____ LUN:____
hdisk5					TID:____ LUN:____	TID:____ LUN:____
hdisk6					TID:____ LUN:____	TID:____ LUN:____
hdisk7					TID:____ LUN:____	TID:____ LUN:____
hdisk8					TID:____ LUN:____	TID:____ LUN:____
hdisk9					TID:____ LUN:____	TID:____ LUN:____
and so on…						

Configuring the new devices

This section describes how to configure the new disk devices on an AIX® host:

- Changing the default device parameters
- Assigning new devices to volume groups and setting partition sizes
- Creating, mounting, and verifying file systems

Changing the default device parameters

After the Hitachi storage system is installed and connected and the device files have been created, the AIX® system sets the device parameters to the system default values. If necessary, you can change the read/write time-out, queue type, and queue depth parameters for each new device using the System Management Information Tool (SMIT) or the AIX® command line (see Changing device parameters from the AIX® command line).



Note: When you set parameters for the FX devices and SCSI disk devices, use the same settings and device parameters for all storage system devices.



Note: If you installed the ODM update, skip this section and go to Assigning new devices to volume groups and setting partition sizes.

Table 3-2 specifies the read/write time-out and queue type requirements for the devices. Table 3-3 specifies the queue depth requirements for the devices. To optimize the I/O performance of the devices, you can adjust the queue depth for the devices later within the specified range. For details, see Host queue depth.

Table 3-2 Read/write time-out and queue type requirements

Parameter Name	Default Value	Requirement
Read/write time-out	30	60
Queue type	none	simple

Table 3-3 Queue depth

Parameter	Recommended value for VSP G1000, VSP G1500, VSP F1500, VSP Gx00, VSP Fx00, HUS VM, VSP	Required value for USP V/VM
Queue depth per LU	32 per LU	≤ 32
Queue depth per port (MAXTAGS)	2048 per port	≤ 2048 per port

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 (see Figure 3-3).

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

```
cfgmr
rmdev -l hdisk$i
chdev -l hdisk$i -a reserve_policy=no_reserve -a queue_depth=x -a algorithm=round_robin
mkdev -l hdisk$i
```

   Note: *x* is used to indicate the desired queue depth within the limits specified in Table 3-3.

3. Repeat steps 1 and 2 for each new device.

4. Type the following command to verify that the parameters for all devices were changed (see Figure 3-2):

```
lsattr -E -1 hdiskx
```

```
#lsattr -E -1 hdisk1
scsi_id         0xef                                SCSI ID
lun_id          0x0                                 LUN ID
location                                            Location Label
ww_name         0x500490e802757500                  FC World Wide Name for this LUN
pvid            000432871c6bbceb0000000000000000    Physical volume identifier
queue_depth     8                                   Queue DEPTH
q_type          simple                              Queuing TYPE
q_err           yes                                 Use QERR bit
clr_q           no                                  Device CLEARS its Queue on error
rw_timeout      60                                  READ/WRITE time out value
start_timeout   60                                  START unit time out value
reassign_to     120                                 REASSIGN time out value
```

Figure 3-2 Verifying the device parameters using the lsattr -E -1 hdiskx command

```
#lscfg -vl hdisk1
DEVICE            LOCATION          DESCRIPTION
hdisk1            20-58-01          Other FC SCSI Disk Drive
Manufacturer.................HITACHI
Machine Type and Model......OPEN-3              Type of device emulation
ROS Level and ID............30313130
Serial Number...............04007575      Type of System and serial number (hex)
Device Specific.(Z0)........000002026300003A
Device Specific.(Z1)........0200 1A              LCU (02) LDEV (00) and port (1A)
Device Specific.(Z2)........
```

Figure 3-3 Verifying the device parameters using the lscfg -vl hdisk1 command

Assigning new devices to volume groups and setting partition sizes

After you change the device parameters, assign the new SCSI disk devices to new or existing volume groups and set the partition size using SMIT. If SMIT is not installed, see the IBM® AIX® user guide for instructions on assigning new devices to volume groups using AIX® commands.

- Table 3-4 specifies the partition sizes for standard LUs.
- Table 3-5 specifies the partition sizes for VLL LUSE devices.
- Table 3-6 specifies the partition sizes for LUSE devices (OPEN-x*n).



Note: **Do not assign the FX devices (for example, 3390-3A/B/C) to volume groups.** If you are configuring storage devices for databases that use a "raw" partition, do not assign those devices to volume groups.

To assign the SCSI disk devices to volume groups and set the partition size:

1. At the AIX® command line prompt, type the following command to start SMIT and open the System Management panel: `smit`
2. Select **System Storage Management (Physical & Logical Storage)** to open the System Storage Management panel.
3. Select **Logical Volume Manager** to open the Logical Volume Manager panel.
4. Select **Volume Groups** to open the Volume Group panel.
5. Select **Add a Volume Group** to open the Add a Volume Group panel.
6. Using the Add a Volume Group panel (see Figure 3-4), you can assign one or more devices (physical volumes) to a new volume group and set the physical partition size:
   a. Place the cursor in the **VOLUME GROUP name** entry field. Enter the name of the new volume group (for example, **VSPvg0**). A volume group can contain multiple hdisk devices, depending on the application.
   b. Place the cursor in the **Physical partition SIZE in megabytes** field, and press the **F4** key. When the size menu appears, select the correct partition size for the devices.
   c. Place the cursor in the **PHYSICAL VOLUME names** entry field. Enter the device file names for the desired devices (for example, hdisk1), or press **F4** and select the device file names from the list.
   d. Place the cursor in the **Activate volume group AUTOMATICALLY** entry field.
   e. Type **yes** to activate the volume group automatically at system restart, or type **no** if you are using a High Availability Cluster Multi-Processing (HACMP) product.
7. Press the **Enter** key.
8. When the confirmation panel opens, select **Yes** to assign the specified devices to the specified volume group with the specified partition size.

9. When the Command Status panel opens, wait for OK to appear on the Command Status line (this response ensures that the devices have been assigned to a volume group).
10. To continue creating volume groups, press F3 until the Add a Volume Group panel opens.
11. Repeat steps 2 through 10 until all new disk devices are assigned to a volume group.

```
                                  Add a Volume Group
Type or select values in entry fields.
Press Enter AFTER making all desired changes.
                                      [Entry Fields]
  VOLUME GROUP name                   [VSPvg0]              ← Enter volume group.
  Physical partition SIZE in megabytes  4                   ← Enter partition size.
  PHYSICAL VOLUME names               [hdisk1]              ← Enter device file names.
  Activate volume group AUTOMATICALLY    yes                ← Enter no for HACMP.
    at system restart
  Volume Group MAJOR NUMBER              []
  Create VG Concurrent Capable?
  Auto-varyon in Concurrent Mode?
F1=Help          F2=Refresh         F3=Cancel          F4=List
F5=Reset         F6=Command         F7=Edit            F8=Image
F9=Shell         F10=Exit           Enter=Do
```

Figure 3-4 Assigning Devices to Volume Groups and Setting the Partition Size

Table 3-4 Partition sizes for standard LUs

Device Type	Partition Size
OPEN-3	4
OPEN-8	8
OPEN-9	8
OPEN-E	16
OPEN-L	64
OPEN-V	256 (default size)

Table 3-5 Partition sizes for VLL LUSE devices

Device Type	LU Size (MB)	Partition Size (MB)
OPEN-*x**n VLL	35-1800	2
	1801-2300	4
	2301-7000	8
	7001-16200	16
	13201-32400	32
	32401-64800	64
	64801-126000	128
	126001–259200	256
	259201–518400	512
	518401 and higher	1024

Table 3-6 Partition sizes for LUSE devices

Device Type	LUSE Configuration	Partition Size (MB)
OPEN-3	OPEN-3	4
	OPEN-3*2-OPEN-3*3	8
	OPEN-3*4-OPEN-3*6	16
	OPEN-3*7-OPEN-3*13	32
	OPEN-3*14-OPEN-3*27	64
	OPEN-3*28-OPEN-3*36	128
OPEN-8	OPEN-8	8
	OPEN-8*2	16
	OPEN-8*3-OPEN-8*4	32
	OPEN-8*5-OPEN-8*9	64
	OPEN-8*10-OPEN-8*18	128
	OPEN-8*19-OPEN-8*36	256
OPEN-9	OPEN-9	8
	OPEN-9*2	16
	OPEN-9*3-OPEN-9*4	32
	OPEN-9*5-OPEN-9*9	64
	OPEN-9*10-OPEN-9*18	128
	OPEN-9*19-OPEN-9*36	256
OPEN-E	OPEN-E	16
	OPEN-E*2	32
	OPEN-E*3,OPEN-E*4	64
	OPEN-E*5-OPEN-E*9	128
	OPEN-E*10-OPEN-E*18	256
OPEN-L	OPEN-L	64
	OPEN-L*2-OPEN-L*3	128
	OPEN-L*4-OPEN-L*7	256
OPEN-V	OPEN-V is a VLL-based volume	

Creating, mounting, and verifying file systems

After you assign SCSI disk devices to volume groups and set the partition sizes, you can create and verify the file systems for the new SCSI disk devices.

- ☐ Creating the file systems
- ☐ Mounting and verifying file systems

Note: Do not create file systems or mount directories for the FX devices (for example, 3390-3A). These devices are accessed as raw devices and do not require any further configuration after being partitioned and labeled.

Creating the file systems

To create the file systems for the newly installed SCSI disk devices:

1. At the AIX® command line prompt, type the following command to start SMIT and open the System Management panel: **`smit`**

Note: If SMIT is not installed, see the IBM® AIX® user guide for instructions on creating file systems using AIX® commands.

2. Select **System Storage Management (Physical & Logical Storage).** The System Storage panel opens.
3. Select **File Systems**. The File Systems panel opens.
4. Select **Add/Change/Show/Delete File Systems**. The Add/Change panel opens.
5. Select **Journaled File Systems**. The Journaled File System panel opens.
6. Select **Add a Standard Journaled File System**. The Volume Group Name panel opens.
7. Move the cursor to the selected volume group, then press **Enter**.
8. Select the desired value, then press **Enter** (see Figure 3-5).
9. In the **SIZE of file system** field, enter the desired file system size (see Table 3-7).
10. In the **Mount Point** field, enter the desired mount point name (for example, **/VSP_VG00**). Record the mount point name and file system size for use later in the configuration process.
11. In the **Mount AUTOMATICALLY** field, type **`yes`** to auto-mount the file systems.

Note: If you are using a HACMP product, do not set the file systems to auto-mount.

12. In the **Number of bytes per inode** field, enter the correct value for the selected device (see Table 3-8, Table 3-9, and Table 3-10).

13. Be sure that the file system size, mount point name, auto-mount options, and number of bytes per inode are correct. Press **Enter** to create the Journaled File System.
14. The Command Status panel appears. To be sure the Journaled File System has been created, wait for **OK** to appear on the Command Status line (see Figure 3-6).
15. Repeat steps 2 through 14 for each Journaled File System that you want to create. To continue creating Journaled File Systems press the **F3** key until you return to the Add a Journaled File System panel.
16. To exit SMIT, press **F10**.

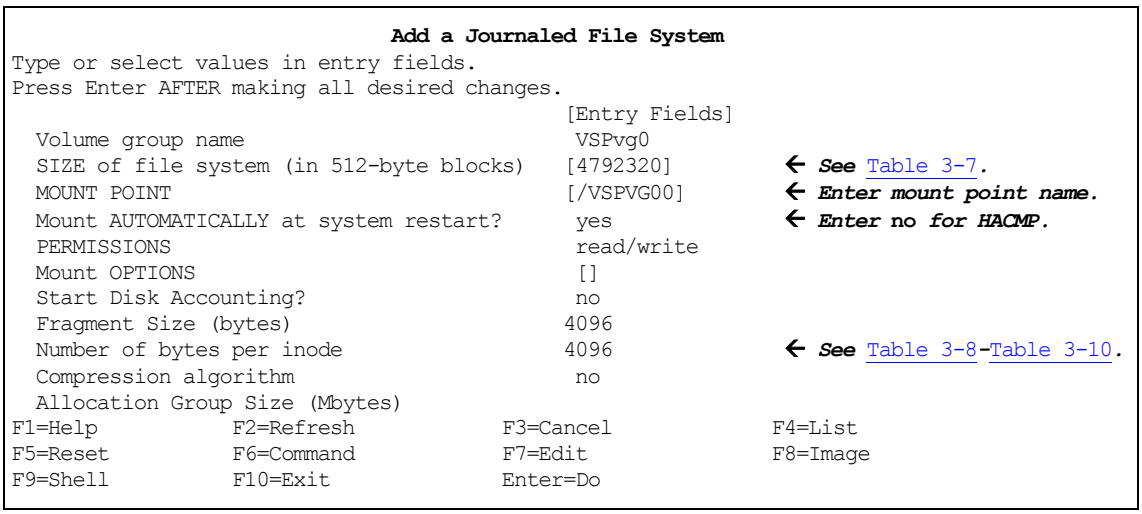


Figure 3-5 Adding a Journaled File System Using SMIT

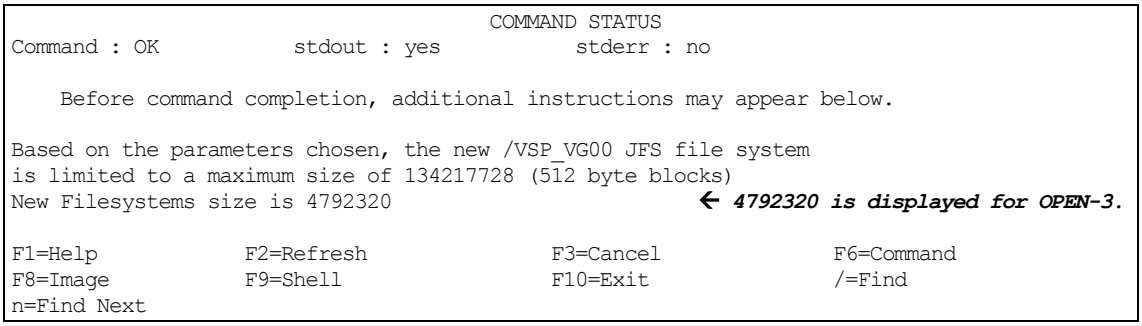


Figure 3-6 Verifying creation of Journaled File System

Table 3-7 Journaled File System size

Device Type	LU Product Name	Capacity (in 512-Byte Blocks)	Maximum File System Size *(see Note 1)* (in 512-Byte Blocks)
Standard LU	OPEN-3	4806720	4792320
	OPEN-8	14351040	14319616
	OPEN-9	14423040	14401536
	OPEN-E	28452960	28409856
	OPEN-L	71192160	71041024
	OPEN-V	Max.125827200	*Max.125566976*
	OPEN-*x**n	See Table 1-2.	(see ***Note 2***)
LUSE Device	OPEN-*x**n VLL	See Table 1-2.	(see ***Note 2***)
VLL LUSE Device	OPEN-*x**n VLL	See Table 1-2.	(see ***Note 2***)

Note 1: When determining SIZE of file system at **Add a Journaled File System**, AIX® already uses an unspecified amount of disk space. You must determine the remaining space available for physical partitions.

Note 2: Calculate the file system size for these devices as follows:

1. Display the number of free physical partitions (FREE PPs) and physical partition size (PP SIZE) by entering the following command (see Figure 3-7): **lsvg**
2. Calculate the maximum size of the file system as follows:
   (FREE PPs – 1) × (PP SIZE) × 2048

Figure 3-7 shows an example for OPEN-3*20 LUSE:
Maximum file system size = (733 – 1) × (64) × 2048 = 95944704

```
# lsvg VSPvg0
VOLUMEGROUP:   VSPvg0              VG IDENTIFIER:   0083665612e98521
VG STATE:      active               PP SIZE:        64 megabyte(s)
VG PERMISSION: read/write           TOTAL PPs:      733 (46912 megabytes)
MAX LVs:       256                  FREE PPs:       733 (46912 megabytes)
LVs:           0                    USED PPS:       0 (0 megabytes)
OPEN LVs:      0                    QUORUM:         2
TOTAL PVs:     1                    VG DESCRIPTORS: 2
STALE PVs;     0                    STALE PPs       0
ACTIVE PVs     1                    AUTO ON:        yes
Concurrent:    Non-Capable          Auto-Concurrent: Disabled
VG Mode:       Non-Concurrent
```

Figure 3-7 Determining the maximum file system size

Table 3-8 Number of bytes per inode for LUSE devices

Device type	LU product name	Number of bytes per inode
OPEN-3	OPEN-3, OPEN-3*2-OPEN-3*28	4096
	OPEN-3*29-OPEN-3*36	8192
OPEN-8	OPEN-8, OPEN-8*2-OPEN-8*9	4096
	OPEN-8*10-OPEN-8*18	8192
	OPEN-8*19-OPEN-8*36	16384
OPEN-9	OPEN-9, OPEN-9*2-OPEN-9*9	4096
	OPEN-9*10-OPEN-9*18	8192
	OPEN-9*19-OPEN-9*36	16384
OPEN-E	OPEN-E, OPEN-E*2-OPEN-E*4	4096
	OPEN-E*5-OPEN-E*9	8192
	OPEN-E*10-OPEN-E*18	16384
OPEN-L	OPEN-L	4096
	OPEN-L*2-OPEN-L*3	8192
	OPEN-L*4-OPEN-L*7	16384
OPEN-V	See Table 3-10	

Table 3-9 Number of bytes per inode for VLL

Device type	LU product name	Number of bytes per inode
OPEN-*x* VLL	OPEN-3 VLL, OPEN-8 VLL, OPEN-9 VLL, OPEN-E VLL, OPEN-V VLL	4096

Table 3-10 Number of bytes per inode for VLL LUSE

Device type	LU size in megabytes	Number of bytes per inode
OPEN-*x**n VLL	35-64800	4096
	64801-126000	8192
	126001 and higher	16384

Mounting and verifying file systems

After you create the Journaled File Systems, mount the file systems and verify that the file systems were created correctly and are functioning properly.

To mount and verify the file systems:

1. At the AIX® command line prompt, type the following command:

   mount ***<mount_point_name>*** (for example, **mount/VSP_VG00**)
2. Repeat step 1 for each new file system.
3. Use the **df** command to verify the size of the file systems you created.



Note: The file system capacity is listed in 512-byte blocks by default. To list capacity in 1024-byte blocks, use the **df -k** command.

4. Verify that the new devices and file systems are fully operational by performing some basic operations (for example, file creation, copying, deletion) on each device (see Figure 3-8).
5. Restart the system and verify that the file systems have successfully auto-mounted by using the **mount** or **df** command to display all mounted file systems (see Figure 3-9). Any file systems that were not auto-mounted can be set to auto-mount using SMIT.



Note: If you are using a HACMP™ product, do not set the file systems to auto-mount.

```
# cd /VSPVG00      ← Go to mount point.
# cp /smit.log /VSPVG00/smit.log.back1                              ← Copy file.
# ls -l VSPVG00                                                   ← Verify file copy.
-rw-rw-rw-   1   root  system      375982 Nov 30 17:25 smit.log.back1
# cp smit.log.back1 smit.log.back2                                  ← Copy file again.
# ls -l                                                           ← Verify copy again.
-rw-rw-rw-   1   root  system      375982 Nov 30 17:25 smit.log.back1
-rw-rw-rw-   1   root  system      375982 Nov 30 17:28 smit.log.back2
# rm smit.log.back1                                               ← Remove test file.
# rm smit.log.back2                                               ← Remove test file.
```

Figure 3-8 Verifying the auto-mounted file systems

```
# df                                                              ← List mounted file systems.
File system  512-blocks      free  %Used  Iused  %Iused  Mounted on
/dev/hd4           8192      3176   61%   652     31%   /
/dev/hd2        1024000    551448   46%  6997      5%   /usr
/dev/hd9var        8192      5512   32%    66      6%   /var
/dev/hd3          24576     11608   52%    38      0%   /tmp
/dev/hd1           8192      7840    4%    17      1%   /home
/dev/lv00       4792320   4602128    4%    16      1%   /VSPVG00        ← OPEN-3 device.
/dev/lv01       4792320   4602128    4%    16      1%   /VSPVG01        ← OPEN-3 device.
/dev/lv02      14401536  13949392    4%    16      1%   /VSPVG02        ← OPEN-9 device.
```

Figure 3-9 Final file system verification

Using the Object Data Manager with Hitachi RAID storage

This section describes the IBM® AIX® Object Data Manager (ODM) and its relationship with the Hitachi RAID storage system:

- ☐ Overview of ODM
- ☐ ODM advantages and cautions
- ☐ Using ODM

Overview of ODM

The ODM is a repository of system information that 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® also provides a set of commands to manipulate the ODM and procedures to package ODM updates. For details, see the following references:

- **Device Configuration Database**
  - http://www16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/kernextc/device_config_db_over.htm
- **Device Configuration System**
  - http://www16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/kernextc/device_config_subsys.htm#a4d56110chri

- **List of ODM commands and subroutines**
  - http://publib.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/genprogc/odm_cmds_subrs.htm
  - http://publib.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/genprogc/odm.htm (Chapter 17. Object Data Manager)
  - http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/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

ODM advantages and cautions

Advantages

The Hitachi Data Systems ODM updates enable the AIX® system to recognize Hitachi disk devices and set the proper attributes. If the attributes for queue type, queue depth, and read/write timeout are not the same for all Hitachi devices, disk errors can be logged both on the storage system 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 default values recommended by the manufacturer. For Hitachi disk devices, the default queue depth is **2** (with a range of 1-32) and the default read/write timeout value is **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.

For details about AIX ODM for Hitachi storage, see the Hitachi Data Systems Technical Upload Facility (TUF) site:
https://tuf.hds.com/instructions/servers/AIXODMUpdates.php

Cautions

Since the Hitachi ODM update changes attributes, it is possible that you may experience problems if you share ports on the Hitachi RAID storage system with multiple AIX® servers at different ODM update levels (for example, one AIX® host at 5.4.0.0 and one AIX® host at 5.4.0.4). Contact your Hitachi Data Systems representative for more information on restrictions when sharing ports.

Using ODM

This section describes how to use ODM with Hitachi storage:

- ☐ Discovering new devices
- ☐ Deleting devices
- ☐ Queue depth and read/write timeout values

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, then 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 (for example, 2105 or LVD SCSI disk drive). This customized definition is persistent and will remain 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 (for example, the WWN) changes. The definitions in the ODM are persistent and remain until a system administrator removes them.

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 60.
- The default value for queue depth is 2.

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. The disk devices on the Hitachi RAID storage system should be defined as **Hitachi Disk Array (Fibre)**. See Table 3-3 for queue depth requirements for the Hitachi RAID disk devices.

Online device installation

After initial installation and configuration of the Hitachi RAID storage system, additional devices can be installed or removed online without having to restart the AIX® system. After online installation, the device parameters for new volumes must be changed to match the LUs defined under the same fibre-channel port (see Changing the default device parameters). This procedure should be performed by the system administrator (that is, super-user).



Note: For additional instructions about online installation and reinstallation of LUs, see the Maintenance Manual for the storage system.

To install or uninstall a device online without having to restart the system:

1. Log on to the AIX® system as **root**.
2. At the AIX® command line prompt, type the following command to start SMIT and open the System Management panel: **`smit`**



Note: If SMIT is not installed, see the IBM® AIX® user guide for instructions on assigning new devices to volume groups using AIX® commands.

3. Select **Devices** to open the Devices panel.
4. Select **Install/Configure Devices Added After IPL** to open the Install/Configure Devices Added After IPL panel.
5. Select **INPUT device/directory** for software, then press **Enter**. The AIX® system scans the buses for new devices.
6. To verify that the new device is installed, type the following command:

   `lsdev -C -c disk`



Note: See Verifying new device recognition for complete instructions. Record the device file names for the new devices.

Configure the new devices for AIX® operations as described in Configuring the new devices and Using the Object Data Manager with Hitachi RAID storage.

Online LUSE configuration

Online LUSE is LU Expansion that is performed after mounting (2GB => 5GB). Before you begin, verify that the size of corresponding LUN in the storage system can be expanded online. Online LUSE involves the following steps:

- ☐ Creating and mounting the file systems
  - Unmounting the file system
  - Varying off the volume group
  - Expanding the size of LU from the Hitachi RAID storage system
  - Varying on the volume group
  - Changing the volume group
  - Mounting the file system
- ☐ Expanding the logical volume
- ☐ Expanding the file system (up to 3 GB)
- ☐ Increasing the file system (up to 40 GB)



Note:

- There is no unmount during this process.
- Online LUSE is available for AIX® 5.2 and later.

Creating and mounting the file systems

1. Type the following command to unmount all file systems in the affected volume group:

   #umount /mnt/h00

2. Type the following command to vary off the volume group:

   #varyoff vg_fc00

3. Expand the size of LU from the Hitachi RAID storage system.
4. Vary on the volume group:

```
#varyonvg vg_fc00
0516-1434 varyonvg: Following physical volumes appear to be grown in size
Run chvg command to activate the new space.
hdisk1
```

5. Change the volume group:

```
#chvg  -g vg_fc00
0516-1224 chvg: WARNING, once this operation is completed, volume group vg_fc00
cannot be imported into AIX 510 or lower versions. Continue (y/n) ?

y
0516-1164 chvg: Volume group vg_fc04 changed. With given characteristics vg_fc00
can include up to 16 physical volumes with 2032 physical partitions each.
```

6. Type the following command to mount all file systems unmounted in step 1:

   #mount /mnt/h00

7. Type the **df-k** command as follows:

```
# df -k
/dev/lv00         2097152   2031276    4%       17     1% /mnt/h00
```

8. Type the **lsvg vg_fc00** command:

```
# lsvg vg_fc00
VOLUME GROUP:   vg_fc00                  VG IDENTIFIER:
0007d6dc00004c00000000f3305f5d36
VG STATE:       active                   PP SIZE:        128 megabyte(s)
VG PERMISSION:  read/write               TOTAL PPs:      543 (69504 megabytes)
MAX LVs:        256                      FREE PPs:       526 (67328 megabytes)
LVs:            2                        USED PPs:       17 (2176 megabytes)
OPEN LVs:       2                        QUORUM:         2
TOTAL PVs:      1                        VG DESCRIPTORS: 2
STALE PVs:      0                        STALE PPs:      0
ACTIVE PVs:     1                        AUTO ON:        yes
MAX PPs per PV: 1016                     MAX PVs:        32
LTG size:       128 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:      no                       BB POLICY:      relocatable
```

9. Type the **lslv lv00** command:

```
# lslv lv00
LOGICAL VOLUME:     lv00                   VOLUME GROUP:   vg_fc00
LV IDENTIFIER:      0007d6dc00004c00000000f3305f5d36.2 PERMISSION:     read/write
VG STATE:           active/complete        LV STATE:       opened/syncd
TYPE:               jfs                    WRITE VERIFY:   off
MAX LPs:            512                    PP SIZE:        128 megabyte(s)
COPIES:             1                      SCHED POLICY:   parallel
LPs:                16                     PPs:            16
STALE PPs:          0                      BB POLICY:      relocatable
INTER-POLICY:       minimum                RELOCATABLE:    yes
INTRA-POLICY:       middle                 UPPER BOUND:    32
MOUNT POINT:        /mnt/h00               LABEL:          /mnt/h00
```

Expanding the logical volume

1. Type the **`extendlv lv00 400`** command:

```
# extendlv lv00 400

# lsvg vg_fc00
VOLUME GROUP:   vg_fc00                  VG IDENTIFIER:
0007d6dc00004c00000000f3305f5d36
VG STATE:       active                   PP SIZE:        128 megabyte(s)
VG PERMISSION:  read/write               TOTAL PPs:      543 (69504 megabytes)
MAX LVs:        256                      FREE PPs:       126 (16128 megabytes)
LVs:            2                        USED PPs:       417 (53376 megabytes)
OPEN LVs:       2                        QUORUM:         2
TOTAL PVs:      1                        VG DESCRIPTORS: 2
STALE PVs:      0                        STALE PPs:      0
ACTIVE PVs:     1                        AUTO ON:        yes
MAX PPs per PV: 1016                     MAX PVs:        32
LTG size:       128 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:      no                       BB POLICY:      relocatable
```

2. Type the **`lslv lv00`** command:

```
# lslv lv00
LOGICAL VOLUME:     lv00                    VOLUME GROUP:   vg_fc00
LV IDENTIFIER:      0007d6dc00004c00000000f3305f5d36.2 PERMISSION:     read/write
VG STATE:           active/complete         LV STATE:       opened/syncd
TYPE:               jfs                     WRITE VERIFY:   off
MAX LPs:            512                     PP SIZE:        128 megabyte(s)
COPIES:             1                       SCHED POLICY:   parallel
LPs:                416                     PPs:            416
STALE PPs:          0                       BB POLICY:      relocatable
INTER-POLICY:       minimum                 RELOCATABLE:    yes
INTRA-POLICY:       middle                  UPPER BOUND:    32
MOUNT POINT:        /mnt/h00                LABEL:          /mnt/h00
MIRROR WRITE CONSISTENCY: on/ACTIVE
EACH LP COPY ON A SEPARATE PV ?: yes
Serialize IO ?:     NO
```

Note:

- To determine the parameters for LUSE expansion, see Table 3-5 (Partition Sizes for VLL LUSE Devices), Table 3-6 (Partition Sizes for LUSE Devices), and Table 3-8 (Number of Bytes per inode for LUSE Devices).
- To correspond to the capacity per emulation type, physical partitions such as PPs, LPs, and inodes will need to be adjusted. They cannot be set with the OS default value.
- The number of bytes per inode cannot be changed with online LUSE

Expanding the file system (up to 3 GB)

1. Type the **chfs** command to change the size of the file system to 10485760:

   # chfs -a size=+3G /mnt/h00

2. Type the **df-k** command:

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            32768     18496   44%     1474     9% /
/dev/hd2           851968     33396   97%    24029    12% /usr
/dev/hd9var         32768      4712   86%      436     6% /var
/dev/hd3            32768     31620    4%       47     1% /tmp
/dev/hd1            32768     29936    9%       97     2% /home
/proc                   -         -    -        -     -  /proc
/dev/hd10opt        32768     24108   27%      395     5% /opt
/dev/lv00         5242880   5078268    4%       17     1% /mnt/h00
```

Increasing the file system (up to 40 GB)

1. Type the **chfs** command to change the file system size to 31457280:

   # chfs -a size=+10G /mnt/h00

2. Type the **df-k** command:

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            32768     18496   44%     1474     9% /
/dev/hd2           851968     33396   97%    24029    12% /usr
/dev/hd9var         32768      4584   87%      436     6% /var
/dev/hd3            32768     31620    4%       47     1% /tmp
/dev/hd1            32768     29936    9%       97     2% /home
/proc                   -         -    -        -     -  /proc
/dev/hd10opt        32768     24108   27%      395     5% /opt
/dev/lv00        15728640  15234908    4%       17     1% /mnt/h00
```

3. Type the **lsvg vg_fc00** command:

```
# lsvg vg_fc00
VOLUME GROUP:   vg_fc00                  VG IDENTIFIER:
0007d6dc00004c00000000f3305f5d36
VG STATE:       active                   PP SIZE:        128 megabyte(s)
VG PERMISSION:  read/write               TOTAL PPs:      543 (69504 megabytes)
MAX LVs:        256                      FREE PPs:       126 (16128 megabytes)
LVs:            2                        USED PPs:       417 (53376 megabytes)
OPEN LVs:       2                        QUORUM:         2
TOTAL PVs:      1                        VG DESCRIPTORS: 2
STALE PVs:      0                        STALE PPs:      0
ACTIVE PVs:     1                        AUTO ON:        yes
MAX PPs per PV: 1016                     MAX PVs:        32
LTG size:       128 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:      no                       BB POLICY:      relocatable
```

4. Type the **chfs** command to change the size of the file system to 94371840:

   # chfs -a size=+30G /mnt/h00

5. Type the **`lsvg vg_fc00`** command:

```
# lsvg vg_fc00
VOLUME GROUP:   vg_fc00                  VG IDENTIFIER:
0007d6dc00004c00000000f3305f5d36
VG STATE:       active                   PP SIZE:        128 megabyte(s)
VG PERMISSION:  read/write               TOTAL PPs:      543 (69504 megabytes)
MAX LVs:        256                      FREE PPs:       126 (16128 megabytes)
LVs:            2                        USED PPs:       417 (53376 megabytes)
OPEN LVs:       2                        QUORUM:         2
TOTAL PVs:      1                        VG DESCRIPTORS: 2
STALE PVs:      0                        STALE PPs:      0
ACTIVE PVs:     1                        AUTO ON:        yes
MAX PPs per PV: 1016                     MAX PVs:        32
LTG size:       128 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:      no                       BB POLICY:      relocatable
#
```

6. Type the **`lslv lv00`** command:

```
# lslv lv00
LOGICAL VOLUME:     lv00                   VOLUME GROUP:   vg_fc00
LV IDENTIFIER:      0007d6dc00004c00000000f3305f5d36.2 PERMISSION:     read/write
VG STATE:           active/complete        LV STATE:       opened/syncd
TYPE:               jfs                    WRITE VERIFY:   off
MAX LPs:            512                    PP SIZE:        128 megabyte(s)
COPIES:             1                      SCHED POLICY:   parallel
LPs:                416                    PPs:            416
STALE PPs:          0                      BB POLICY:      relocatable
INTER-POLICY:       minimum                RELOCATABLE:    yes
INTRA-POLICY:       middle                 UPPER BOUND:    32
MOUNT POINT:        /mnt/h00               LABEL:          /mnt/h00
MIRROR WRITE CONSISTENCY: on/ACTIVE
EACH LP COPY ON A SEPARATE PV ?: yes
Serialize IO ?:     NO
```

7. Type the **`df-k`** command to increase the volume size to 47 GB and fully expand the file system size:

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            32768     18496   44%     1474     9% /
/dev/hd2           851968     33396   97%    24029    12% /usr
/dev/hd9var         32768      4584   87%      436     6% /var
/dev/hd3            32768     31620    4%       47     1% /tmp
/dev/hd1            32768     29936    9%       97     2% /home
/proc                   -         -    -        -     -  /proc
/dev/hd10opt        32768     24108   27%      395     5% /opt
/dev/lv00        47185920  45704828    4%       17     1% /mnt/h00
```

Troubleshooting for AIX® host attachment

Table 3-11 lists potential error conditions that might occur during storage system installation on an AIX® host and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 3-11 Troubleshooting for AIX® host attachment

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Make sure that the READY indicator lights on the storage system are ON. Run `cfgmgr` to recheck the fibre channel for new devices. Make sure that LUSE devices are not intermixed with normal LUs or with FX devices on the same fibre-channel port. Verify that LUNs are configured properly for each TID.
The file system is not mounted after rebooting.	Make sure that the system was restarted properly. Verify that the values listed under Journaled File System are correct.
If a new path is added while an existing path is in I/O processing in alternate path configuration, the status of the added path becomes offline.	Run an online operation on the offline path with the Alternate Path software. For details, see the user documentation for the Alternate Path software.

4

HP-UX configuration and attachment

This chapter describes how to configure and manage the new Hitachi disk devices on an HP-UX host:

- Hitachi storage system configuration for HP-UX operations
- Configuring the new devices
- Online device installation
- Troubleshooting for HP-UX host attachment

Hitachi storage system configuration for HP-UX operations

The storage system must be fully configured before being attached to the HP-UX host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for HP-UX operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)
- Cross-OS File Exchange (FX) (3390-3A/B/C, OPEN-*x*-FXoto)

Host mode. The required host mode for HP-UX is **03**. Do not select a host mode other than **03** for HP-UX. For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Configuring the new devices

This section describes how to configure the new disk devices on an HP-UX host:

- Verifying new device recognition
- Verifying device files and the driver
- Partitioning disk devices
- Creating file systems
- Setting device parameters
- Creating mount directories
- Mounting and verifying file systems
- Setting and verifying auto-mount parameters



Note: Configuration of the devices should be performed by the HP-UX system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Verifying new device recognition

The first step in configuring the new 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.

The devices should be installed and formatted with the ports configured before the host system is powered on. Type the **cfgmgr** command to force the system to check the buses for new devices.

To verify new device recognition:

1. Login to the HP-UX system as **root** as shown in Figure 4-1.
2. Use the **ioscan -f** command to display the device data. Verify that the system recognizes the newly installed devices (see Figure 4-2). If desired, use the -**C disk** command option (**ioscan -fnC disk**) to limit the output to disk devices only.



Notes:

- If UNKNOWN appears as the Class type, the HP-UX system may not be configured properly. Refer to the HPE documentation or contact HPE technical support.
- If information for unused devices remains in the system, get the system administrator's permission to renew the device information. To renew the device information, delete the **/etc/ioconfig** and **/stand/ioconfig** files (**rm** command), reboot the server, and then issue the **ioinit -c** command. Now issue the **ioscan -f** command to recognize the logical devices again.

3. Make a blank table (see Table 4-1) for recording the device data. The table must have nine columns for the following data: bus number, bus instance number, disk number, H/W path, driver, device type, target ID, LUN, and device file name. You will need three more columns for entering the major and minor numbers later.
4. Enter the device data for each device (disk devices and raw/FX devices) in your table including the device file name. The device file name has the following structure:

   File name = c***X***t***Y***d***Z***, where ***X*** = bus instance #, ***Y*** = target ID, ***Z*** = LUN.

   The "c" stands for controller, the "t" stands for target ID, and the "d" stands for device. The SCSI target IDs are hexadecimal (0 through F) and the LUN is decimal (0 through 7).

5. Verify that the SCSI TIDs correspond to the assigned port address for all connected ports (see SCSI TID Maps for FC adapters). If so, the logical devices are recognized properly. If not:
   a. Check the AL-PA for each port using the LUN Manager software. If the same port address is set for multiple ports on the same loop (AL with HUB), all port addresses except one changed to another value, and the relationship between AL-PA and TID does not correspond to the mapping in SCSI TID Maps for FC adapters. Set a different address for each port, reboot the server, and then verify new device recognition again.
   b. If unused device information remains, the TID-to-AL-PA mapping will not correspond to the mapping in SCSI TID Maps for FC adapters. Renew the device information (see step 2 for instructions) and then verify new device recognition again.

```
The system is ready.

GenericSysName [HP Release B.11.0] (see /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 4-1 Logging in as root

```
# 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                fcT1   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 OPEN-9
                       /dev/dsk/c2t6d0   /dev/rdsk/c2t6d0
disk       4   8/12.8.8.255.0.6.1  sdisk  CLAIMED DEVICE      HITACHI OPEN-9
                       /dev/dsk/c2t6d1   /dev/rdsk/c2t6d1
disk       5   8/12.8.8.255.0.8.0  sdisk  CLAIMED DEVICE      HITACHI 3390*3B
                       /dev/dsk/c2t8d0   /dev/rdsk/c2t8d0
:
#
```

This sample screen shows the following new devices recognized:

- HITACHI OPEN-9 device: bus no. = 8/12, bus instance = 2, target ID = 6, LUN = 0, driver = sdisk
- HITACHI OPEN-9 device: bus no. = 8/12, bus instance = 2, target ID = 6, LUN = 1, driver = sdisk
- HITACHI 3390-3B device: bus no. = 8/12, bus instance = 2, target ID = 8, LUN = 0, driver = sdisk

Figure 4-2 Verifying new device recognition

Table 4-1 Sample device data table for HP-UX (showing data for devices in Figure 4-2)

Bus No.	Instance (XX)	Disk No.	H/W Path	Driver	Device Type	TID (Y)	LUN (Z)	Device File	Major # for Char. Files	Major # for Block Files	Minor # 0xXX Y Z 00
8/12	02	3	8/12.8.8.255.0.6.0	sdisk	OPEN-9	6	0	c2t6d0			
8/12	02	4	8/12.8.8.255.0.6.1	sdisk	OPEN-9	6	1	c2t6d1			
8/12	02	5	8/12.8.8.255.0.8.0	sdisk	3390-3B	8	0	c2t8d0			

Verifying device files and the driver

The device files for all new devices (SCSI disk and raw/FX) should be created automatically during system startup. Each device should have a block-type device file in the **/dev/dsk** directory and a character-type device file in the **/dev/rdsk** directory. The SCSI disk devices must have both device files. Raw/FX devices only require the character-type device file.



Note: Some HP-compatible systems do not create the device files automatically. If the device files were not created automatically, follow the instructions in Creating device files to create the device files manually.

To verify that the device files for the new devices were successfully created:

1. Display the block-type device files in the **/dev/dsk** directory using the **ll** command (equivalent to **ls -l**) with the output piped to **more** (see Figure 4-3). Verify that there is one block-type device file for each device.
2. Use your completed device data table (see Creating device files and Table 4-2) to verify that the block-type device file name for each device is correct.
3. Display the character-type device files in the **/dev/rdsk** directory using the **ll** command with the output piped to **more** (see Figure 4-4). Verify that there is one character-type device file for each new device.
4. Use your completed device data table (see Creating device files and Table 4-2) to verify that the character-type device file name for each device is correct.
5. After verifying the block-type and character-type device files, verify the HP-UX driver for the storage system using the **ioscan -fn** command (see Figure 4-5).

```
# ll /dev/dsk | more                                    ← Check block-type files.
total 0
brw-r-----  1 bin   sys    28 0x000000 Oct  4 11:01 c0t0d0
brw-r-----  1 bin   sys    28 0x006000 Dec  6 15:08 c0t6d0
brw-r-----  1 bin   sys    28 0x006100 Dec  6 15:08 c0t6d1    ← Block-type device file.
     Bus instance # = 0, SCSI target ID = 6, LUN = 1 ↗
```

Figure 4-3 Verifying block-type device files

```
# ll /dev/rdsk | more     ← Check character-type files.
total 0
crw-r-----  1 bin   sys   177 0x000000 Oct  4 11:01 c0t0d0
crw-r-----  1 bin   sys   177 0x006000 Dec  6 15:08 c0t6d0
crw-r-----  1 bin   sys   177 0x006100 Dec  6 15:08 c0t6d1  ← Character-type device file.
     Bus instance # = 0, SCSI target ID = 6, LUN = 1 ↗
```

Figure 4-4 Verifying character-type device files

```
# 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              fcT1   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 OPEN-9
                       /dev/dsk/c2t6d0   /dev/rdsk/c2t6d0
disk       4   8/12.8.8.255.0.6.1  sdisk  CLAIMED DEVICE    HITACHI OPEN-9
                       /dev/dsk/c2t6d1   /dev/rdsk/c2t6d1
disk       5   8/12.8.8.255.0.8.0  sdisk  CLAIMED DEVICE    HITACHI 3390*3B
                        /dev/dsk/c2t8d0   /dev/rdsk/c2t8d0
:
#
```

Figure 4-5 Verifying the HP-UX driver

Creating device files

If the device files were not created automatically when the HP-UX system was restarted, issue the **insf -e** command in the **/dev** directory (see Figure 4-6) to instruct the HP-UX system to create the device files. After executing this command, repeat the procedure in Verifying new device recognition to verify new device recognition and the device files and driver.

```
# cd /dev
# insf -e
insf: Installing special files for mux2 instance 0 address 8/0/0
      :          :         :        :
      :          :         :        :
#
```

Figure 4-6 Issuing a command to create the device files

If the device files for the new devices cannot be created automatically, use the **mknod** command to create the device files manually:

1. Obtain your Device Data table on which you recorded the data for the new devices (see Table 4-2). You should have the following information for all new devices:
   - Bus number
   - Bus instance number
   - Disk number
   - Driver
   - Device type
   - Target ID
   - LUN
2. Build the device file name for each device, and enter the device file names into your table. Example:

   File name = c***X***t***Y***d***Z***, where ***X*** = bus instance #, ***Y*** = target ID, ***Z*** = LUN.

3. Build the minor number for each device, and enter the minor numbers into your table. Example:

   0x**XXYZ**00, where **XX** = bus instance #, **Y** = SCSI target ID, and **Z** = LUN.

4. Display the driver information for the system using the **lsdev** command (see Figure 4-7).

5. Enter the major numbers for the drivers into your table. You should now have all required device and driver information in the table (see Table 4-2).

6. Create the device files for all new devices (SCSI disk and raw/FX devices) using the **mknod** command (see Figure 4-8). Be sure to create the block-type device files in the **/dev/dsk** directory and the character-type device files in the **/dev/rdsk** directory.

   The character-type device file is required for volumes used as raw devices (for example, 3390-3A). The block-type device file is not required for raw devices.

   If you need to delete a device file, use the **rm -i** command.

```
# lsdev                                          ← Display driver information.
  Character     Block     Driver     Class
      :           :         :         :
     188          31       sdisk      disk
#
```

This sample screen shows the following system information for the "sdisk" device driver:
Major number of driver sdisk for character-type files: **188**
Major number of driver sdisk for block-type files: **31**

Figure 4-7 Displaying driver information

Table 4-2 Completed device data table (sample)

Bus No.	Instance (XX)	Disk No.	H/W Path	Driver	Device Type	TID (Y)	LUN (Z)	Device File	Minor # 0xXXYZ00	Major # – Char. Files	Major # – Block Files
8/12	02	3	8/12.8.8.255.0.6.0	sdisk	OPEN-9	6	0	c2t6d0	0x026000	188	31
8/12	02	4	8/12.8.8.255.0.6.1	sdisk	OPEN-9	6	1	c2t6d1	0x026100	188	31
8/12	02	5	8/12.8.8.255.0.8.0	sdisk	3390-3B	8	0	c2t8d0	0x028000	188	31

```
# cd /dev/dsk                                    ← Go to /dev/dsk directory.
# mknod /dev/dsk/c2t6d0 b 31 0x026000            ← Create block-type file.
          ↖File name  ↖ b = block-type, 31 = major #, 0x026000 = minor #
 :
# mknod /dev/dsk/c2t8d0 b 31 0x028000
          ↖ File name  ↖ b = block-type, 31 = major #, 0x028000 = minor #
```

```
# cd /dev/rdsk                                   ← Go to /dev/rdsk directory.
# mknod /dev/rdsk/c2t6d0 c 188 0x026000          ← Create character-type file.
          ↖ File name   ↖ c = character-type, 177 = major #, 0x026000 = minor #
 :
# mknod /dev/rdsk/c2t8d0 c 188 0x028000
          ↖ File name   ↖ c = character-type, 188 = major #, 0x028000 = minor #
```

Figure 4-8 Creating block-type & character-type device files manually

Partitioning disk devices

The HP-UX system uses the Logical Volume Manager (LVM) to manage the disk devices on all peripheral storage devices including the Hitachi RAID storage system. Under LVM disk management, a volume group consisting of multiple disks is formed, and then the volume group is divided into logical partitions and managed as a logical volume. These procedures should be executed for all device files corresponding to the new Hitachi SCSI disk devices.



WARNING: Do not partition the raw/FX devices (for example, 3390-3A/B/C). These volumes are not managed by LVM and do not need any further configuration after their character-type device files have been created and verified.

To partition the new SCSI disk devices for LVM operation:

- Create a physical volume for each new SCSI disk device (see Creating physical volumes).
- Create new volume groups as desired (see Creating volume groups). To increase the maximum volume groups (**maxvgs**) setting.
- Create a logical volume for each new SCSI disk device (see Creating logical volumes).

This section provides general instructions and basic examples for partitioning the Hitachi SCSI devices for LVM operations using UNIX commands. These instructions do not explicitly cover all LVM configuration issues. For more information about LVM configuration, see the appropriate user documentation or contact HPE technical support.



Note: If desired, the HP-UX System Administrator Manager (SAM) can be used instead of UNIX commands to configure the SCSI disk devices.

Creating physical volumes

The first step in partitioning the new devices is to create a physical volume for each new disk device. Once the physical volumes have been created, you will be able to assign these new physical volumes to new or existing volume groups for management by LVM.



Note: Do not create physical volumes for raw/FX devices (for example, 3390-3A/B/C).

To create the physical volumes for the new disk devices:

1. Use the **pvcreate** command to create the physical volume with the character-type device file as the argument (see Figure 4-9). Specify the **/dev/rdsk** directory for the character file. You can only create one physical volume at a time.



WARNING: Do not use the **-f** (force) option with the **pvcreate** command. This option creates a new physical volume forcibly and overwrites the existing volume.

2. Repeat step 1 for each new disk device on the Hitachi RAID storage system.

```
# pvcreate /dev/rdsk/c2t6d0                          ← Create physical volume.
                    ↖ Character-type file for disk device.
Physical volume "/dev/rdsk/c2t6d0" has been successfully created.
# pvcreate /dev/rdsk/c2t6d1
Physical volume "/dev/rdsk/c2t6d1" has been successfully created.
:
```

Figure 4-9 Creating physical volumes

Creating volume groups

After the physical volumes for the disk devices have been created, you can begin creating new volume groups for the new physical volumes as needed. If desired, you can also add any of the new physical volumes on the Hitachi RAID storage system to existing volume groups using the **vgextend** command. The physical volumes, which make up one volume group, can be located in the same disk system or in different disk systems.

Notes:

- Do not assign the raw/FX devices (for example, OPEN-*x*-FXoto) to volume groups.
- You may need to modify the HP-UX system kernel configuration (**maxvgs** setting) to allow more volume groups to be created (see Online device installation).

To create a volume group:

1. Use the **ls** command to display the existing volume groups (see Figure 4-10).
2. Use the **mkdir** command to create the directory for the new volume group (see Figure 4-11). Choose a name for the new volume group that is different than all other group names. Do not use an existing volume group name.

   If you need to delete a directory, use the **rmdir** command (for example, **rmdir /dev/vgnn**).
3. Use the **ls** command to verify the new directory (see Figure 4-11).
4. Use the **ll** command to verify the minor numbers for existing group files with the output piped to **grep** to display only the files containing "group" (see Figure 4-12).
5. Choose a minor number for the new group file in sequential order (that is, when existing volume groups are vg00-vg05 and next group name is vg06, use minor number 06 for the vg06 group file). Do not to duplicate any minor numbers.

   The minor numbers are hexadecimal (for example, the tenth minor number is 0x0a0000, not 0x100000).
6. Use the **mknod** command to create the group file for the new directory (see Figure 4-13). Specify the correct volume group name, major number, and minor number. The major number for all group files is **64**.

   If you need to delete a group file, use the **rm -r** command to delete the group file and the directory at the same time (for example, **rm -r /dev/vgnn**), and start again at step 2.
7. Repeat steps 5 and 6 for each new volume group.

8. Use the **vgcreate** command to create the volume group (see Figure 4-14).

   To allocate more than one physical volume to the new volume group, add the other physical volumes separated by a space (for example, **vgcreate /dev/vg06 /dev/dsk/c0t6d0 /dev/dsk/c0t6d1**).

   For LUSE volumes with more than 17 OPEN-8/9 LDEVs or more than 7043 MB (OPEN 8/9*n-CVS), use the **-s** and **-e** physical extent (PE) parameters of **vgcreate** (see Figure 4-14).

   Table 4-3 lists the PE and maximum PE (MPE) parameters for the LUSE devices on the Hitachi RAID storage system.

   If you need to delete a volume group, use the **vgremove** command (for example, **vgremove /dev/vgnn**). If the **vgremove** command does not work because the volume group is not active, use the **vgexport** command (for example, **vgexport /dev/vgnn**).

9. Use the **vgdisplay** command to verify that the volume group was created correctly (see Figure 4-15). The **-v** option displays the detailed volume group information.

```
# ls /dev                                          ← Display existing volume group names.
vg00
 :
vg05
#
```

Figure 4-10 Displaying existing volume group names

```
# mkdir /dev/vg06                                  ← Make directory for new volume group.
# ls /dev                                          ← Verify directory for new volume group.
vg00
 :
vg06
#
```

Figure 4-11 Creating and verifying a directory for the new volume group

```
# ll /dev/vg* | grep group                         ← Display existing group files.
crw-rw-rw-    1 root     root     64 0x000000 Nov 7 08:13 group
                                    ↖ Minor number of existing group file = 00
:
#
```

Figure 4-12 Displaying minor numbers for existing group files

```
# mknod /dev/vg06/group c 64 0x060000              ← Create new group file.
                ↖        ↖  ↖ Group name = vg06, major number of group file = 64,
                              Minor number of new group file = 06
:
#
```

Figure 4-13 Creating group file for new volume group

```
# vgcreate /dev/vg06 /dev/dsk/c2t6d0                    ← Create new volume group.
                ↖ Vol group name  ↖ Device file name
Volume group "/dev/vg06" has been successfully created.
Volume group configuration for /dev/vg06 has been saved in /etc/1vmconf/vg06.cof.
# vgcreate -s 8 -e 15845 /dev/vg09 /dev/dsk/c2t7d0      ← Example for LUSE with n=18.
           ↖PE Size  ↖ Max Physical Extent Size (MPE)
Volume group "/dev/vg09" has been successfully created.
Volume Group configuration for /dev/vg09 has been saved in /etc/lvmconf/vg09.cof
```

Figure 4-14 Creating new volume group

```
# vgdisplay /dev/vg06                                   ← Verify new volume group.
--- Volume groups ---
VG Name                     /dev/vg06
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      0
Open LV                     0
Max PV                      16
Cur PV                      1
Act PV                      1
Max PE per PV               1016                        ← Verify MPE for LUSE devices.
VGDA                        2
PE Size (Mbytes)            4                           ← Verify PE for LUSE devices.
Total PE                    586
Alloc PE                    0
Free PE                     586
Total PVG                   0
```

Figure 4-15 Verifying new volume group

Table 4-3 PE and MPE parameters for LUSE devices

Device type		Physical Extent Size (PE)	Max Number of Physical Extents (MPE)
OPEN-3/8/9/E OPEN-3*n (n= 2 to 36) OPEN-3-CVS OPEN-3*n-CVS (n = 2 to 36)		default	default
OPEN-8/9*n	n = 2 to 17	default	default
	n = 18	8	15845
OPEN-E*n	n = 2 to 9	default	default
OPEN-L*n	n=2 to 3	default	default
OPEN-8/9/E-CVS, OPEN-V		default	default
OPEN-8/9/E*n-CVS, OPEN-V*n (n = 2 to 36)	70-119731(MB) × N1	8	default
	119732- (MB) × N1	8	N2

N1 = [Virtual LVI/LUN volume capacity (in MB)] × n
N2 = ↑ N1 / PE ↑ (↑↑ means round up to next integer.)

Example: Volume capacity is 6000 MB for OPEN-9*22-CVS volume:
N1 = 6000 × 22 = 132000
N2 = ↑ 132000/8 ↑ = 16500

Creating logical volumes

After you create the new volume groups, create the logical volumes for each new disk device on the Hitachi RAID storage system.

Note: Do not create logical volumes for raw/FX devices (for example, 3390-3A/B/C).

To create the logical volumes:

1. Use the **lvcreate -L** command to create the logical volume, and specify the volume size and volume group for the new logical volume (see Figure 4-16).

   The HP-UX system assigns the logical volume numbers automatically (lvol1, lvol2, lvol3, …). Use the capacity values specified in Table 1-1 for the size parameter (for example, OPEN-3 = 2344, OPEN-V = 61432 in maximum size). To calculate S1 for VLL, LUSE, and VLL LUSE volumes:

   Use the **vgdisplay** command to display the physical extent size (**PE Size**) and usable number of physical extents (**Free PE**) for the volume (see Figure 4-17). Calculate the maximum size value (in MB) as follows:

   S1 = (PE Size) × (Free PE)

2. Use the **lvdisplay** command to verify that the logical volume was created correctly (see Figure 4-18). If desired, wait until all logical volumes have been created, then use the * wildcard character with the **lvdisplay** command to verify all volumes at one time by (for example, **lvdisplay /dev/vg06/lvol***).

3. Repeat steps 1 and 2 for each logical volume to be created. You can only create one logical volume at a time, but you can verify more than one logical volume at a time.

   If you need to delete a logical volume, use the **lvremove** command (for example, **lvremove /dev/vgnn/lvolx**).

   If you need to increase the size of an existing logical volume, use the **lvextend** command (for example, **lvextend -L** size **/dev/vg**nn**/lvol**x).

   If you need to decrease the size of an existing logical volume, use the **lvreduce** command (for example, **lvreduce -L** size **/dev/vg**nn**/lvol**x).

```
# lvcreate -L 2344 /dev/vg06                    ← Create new logical volume.
                  ↖ Size of volume = 2344 MB (OPEN-3)
Logical volume "/dev/vg06/lvol1" has been successfully created with character device
"/dev/vg06/rlvol1".
Logical volume "/dev/vg06/lvol1" has been successfully extended.
Volume Group configuration for /dev/vg06 has been saved in /etc/lvmconf/vg06.cof.
```

Figure 4-16 Creating a logical volume

```
# vgdisplay /dev/vg01

--- Volume groups ---
VG Name             /dev/vg01
VG Write Access     read/write
VG Status           available
Max LV              255
Cur LV              0
Open LV             0
Max PV              16
Cur PV              1
Act PV              1
Max PE per PV       1016
VGDA                2
PE Size (Mbytes)    4                    ← Physical extent size.
Total PE            586
Alloc PE            0
Free PE             586                  ← Number of physical extents.
Total PVG           0
```

This example shows the following information for /dev/vg01:
Physical extent size = 4
Usable number of physical extents = 586

Therefore, maximum size value = 4 × 586 = 2344

Figure 4-17 Calculating volume size for VLL, LUSE, and VLL LUSE devices

```
# lvdisplay /dev/vg06/lvol1                                  ← Verify new logical volume.
--- Logical volume ---
LV Name                 /dev/vg06/lvol1
VG Name                 /dev/vg06
LV Permission           read/write
LV Status               available/syncd
Mirror copies           0
Consistency Recovery    MWC
Schedule                parallel
LV Size (Mbytes)        2344        (7040 for OPEN-9)      ← 2344 = 586 x 4 = OPEN-3
Current LE              586         (1760 for OPEN-9)      ← LE = logical extent
Allocated PE            586         (1760 for OPEN-9)      ← PE = physical extent
Stripes                 0
Stripe Size (Kbytes)    0
Bad block               on
Allocation              strict
```

Figure 4-18 Verifying a logical volume

Creating file systems

After you create logical volumes, you are ready to create the file system for each new logical volume on the Hitachi RAID storage system. The default file system type for HP-UX version 11i is vxfs.



Note: Do not create file systems for the raw/FX devices (for example, 3390-3A/B/C).

To create the file system on a new logical volume:

1. Use the **newfs** command to create the file system with the logical volume as the argument.
   - Figure 4-19 shows an example of creating the file system for an OPEN-3 volume.
   - Figure 4-20 shows an example of creating the file system for an OPEN-9 volume.
   - Figure 4-21 shows examples of specifying the file system type (vxfs) with the **newfs** command.
2. Repeat step 1 for each new logical volume on the storage system.

```
# newfs /dev/vg06/rlvol1                                    ← Create file system.
newfs: /etc/default/fs is used for determining the file system type
mkfs (vxfs): Warning -272 sector(s) in the last cylinder are not allocated.
mkfs (vxfs): /dev/vg06/rlvol1 - 2400256 sectors in 3847 cylinders of 16 tracks,
2457.9MB in 241 cyl groups (16 c/g, 10.22Mb/g, 1600 i/g)
Super block backups (for fsck -b) at:
    16,  10040,  20064,  30038,  40112,  50136,  60160,  70184,  80208,  90232,
  ...
2396176
#
```

Figure 4-19 Creating a file system (default file system, OPEN-3 shown)

```
# newfs /dev/vg06/rlvol1                                    ← Create file system.
newfs: / etc/default/fs is used for determining the file system type
mkfs (vxfs): ...
 :
7188496, 7198520, 7208544
#
```

Figure 4-20 Creating a file system (default file system, OPEN-9 shown)

```
# newfs -F vxfs /dev/vg06/rlvol1                            ← Specify file system type.
 :
# newfs -F vxfs /dev/vg06/rlvol2
```

Figure 4-21 Specifying file system type

Setting device parameters

When device files are created, the HP-UX system sets the IO time-out parameter to its default value of 20 seconds and the queue depth parameter to its default value of either 2 or 8. You must change these values for all new disk devices on the Hitachi RAID storage system. For details about queue depth, see Host queue depth.



Note: Do not change the device parameters for raw/FX devices (for example, 3390-3A/B/C).

Setting the IO time-out parameter

The IO time-out parameter for the disk devices on the Hitachi RAID storage system must be set to **60 seconds**. To change the IO time-out parameter:

1. Use the **pvdisplay** command to verify the current IO time-out value (see Figure 4-22).
2. Use the **pvchange -t** command to change the IO time-out value to 60 (see Figure 4-23).
3. Use the **pvdisplay** command to verify that the new IO time-out value is 60 seconds (see Figure 4-24).
4. Repeat steps 1 through 3 for each new disk device on the storage system.

```
# pvdisplay /dev/dsk/c0t6d0                          ← Checking current IO time-out value.
--- Physical volumes ---
PV Name                 /dev/dsk/c0t6d0
VG Name                 /dev/vg06
PV Status               available
Allocatable             yes
VGDA                    2
Cur LV                  1
PE Size (Mbytes)        4
Total PE                586          ← This value is 586 for OPEN-3 and 1760 for OPEN-9.
Free PE                 0
Allocated PE            586          ← This value is 586 for OPEN-3 and 1760 for OPEN-9.
Stale PE                0
IO Timeout (Seconds)    default                              ← Default IO time-out value.
```

Figure 4-22 Checking current IO time-out value

```
# pvchange -t 60 /dev/dsk/c0t6d0                     ← Change IO time-out value.
Physical volume "/dev/dsk/c0t6d0" has been successfully changed.
Volume Group configuration for /dev/vg06 has been saved in /etc/lvmconf/vg06.cof
```

Figure 4-23 Changing IO time-out value

```
# pvdisplay /dev/dsk/c0t6d0                    ← Verify new IO time-out value.
--- Physical volumes ---
PV Name                /dev/dsk/c0t6d0
VG Name                /dev/vg06
PV Status              available
 :
Stale PE               0
IO Timeout (Seconds)   60                       ← New IO time-out value.
```

Figure 4-24 Verifying new IO time-out value

Setting the queue depth parameter

The HP-UX system automatically sets the queue depth to a default value of 2 or 8, depending on the installed HPE options and drivers. The queue depth for the Hitachi disk devices must be set as specified Table 4-4. For details about queue depth, see Host queue depth.

Using the **scsictl** command, you can view and change the queue depth parameter for each device one volume at a time. However, the queue depth is reset to the default value the next time the system restarts. Therefore, you must create and register a start-up script to set the queue depth for the disk devices each time the system restarts (see Creating and Registering the Queue Depth Start-Up Script).



Note: Do not set the queue depth for the raw/FX devices (for example, 3390-3A/B/C).

Table 4-4 Queue depth for HP-UX

Parameter	Recommended value for VSP G1000, VSP G1500, VSP F1500, VSP Gx00, VSP Fx00, HUS VM, VSP	Required value for USP V/VM
Queue depth per LU	32 per LU	≤ 8
Queue depth per port	2048 per port	≤ 2048 per port

To set the queue depth parameter for the new Hitachi devices:

1. If you cannot shut down and restart the system at this time, use the **scsictl** command to set the queue depth for each new device (see Figure 4-25). The **scsictl** commands to set queue depth should be registered as HP-UX start-up script for future reboot.

2. Check the **/sbin/init.d** and **/sbin/rc1.d** directories to see whether the script name **queue** is already used (link name **Sxxxqueue** or **Kxxxqueue**) (see Figure 4-26). Choose a unique name for the start-up script as follows:
   a. If there is no script named queue and no link file named **Sxxxqueue** or **Kxxxqueue**, use the name queue for the new script and go to step 3.
   b. If the script queue and the link file **Sxxxqueue** or **Kxxxqueue** exist and the script is used to set the queue depth for other previously installed Hitachi RAID storage systems, check the script file to see whether the queue depth is set to the desired number (per Table 4-4) and add a line for each new disk device. If necessary, restart the HP-UX system to set the queue depth for the new volumes.
   c. If the script queue and the link file **Sxxxqueue** or **Kxxxqueue** already exist and the script is not used for setting the queue depth for the Hitachi RAID storage system, use another name for the new queue-depth script for the storage system (for example, **hitachi_q**) and go to step 3.



Note: If the link **Sxxxqueue** and/or **Kxxxqueue** exists, but there is no script file named **queue**, delete the link files, use the name **queue** for the new script, and go to step 3.

3. Choose a unique 3-digit number for the link name. This number cannot be used in any other links. The link name is derived as follows: **S** stands for "start up script," **K** stands for "kill script," the three-digit number is unique to each link, and the script file name follows the three-digit number (for example, **S890queue** or **S890hitachi_q**).
4. Create and register the new start-up script for the Hitachi RAID storage system (see Creating and registering the queue depth start-up script for an example).
5. Shut down and restart the HP-UX system, so the new start-up script sets the queue depth for the disk devices to the specified value (per Table 4-4).
6. After restarting the system or setting the queue depths manually, use the **scsictl** command to verify the queue depth for each Hitachi disk device (see Figure 4-27).

```
# /usr/sbin/scsictl -m queue_depth=8 -a /dev/rdsk/c0t6d0  ←Set queue depth per Table 4-4.
                                                          ↖ Character-type device file
# /usr/sbin/scsictl -m queue_depth=8 -a /dev/rdsk/c0t6d1
# /usr/sbin/scsictl -m queue_depth=8 -a /dev/rdsk/c0t6d2
# /usr/sbin/scsictl -m queue_depth=8 -a /dev/rdsk/c0t6d3
 :
 :
# /usr/sbin/scsictl -m queue_depth=8 -a /dev/rdsk/c0t8d0
```

Figure 4-25 Changing queue depth

```
# ls /sbin/init.d                                          ← Check for QUEUE.
OspfMib     clean_ex    dfs         hpether     names       nis.server  savecore    swconfig
SnmpHpunix  clean_tmps  diagnostic  iforls      ncs         pd          sendmail    syncer
 :
clean_adm   ddfa        hparray     mrouted     nis.client  rwhod       swcluster   xntpd
# ls /sbin/rc1.d                                  ← Check for SxxxQUEUE and KxxxQUEUE.
K230audio       K340xntpd   K420dfs         K475rarpd       K630named       S420set_date
K240auditing    K356vjed    K430dce         K480rdpd        K660net         S440savecore
K250envd        K358egcd    K435OspfMib     K490gated       K700nettl       S500swap_start
K258diagnostic  K360kks     K435SnmpHpunix  K500inetd       K770ptydaemon   S520syncer
K270cron        K370vt      K435SnmpMib2    K510mrouted     K780syslogd
K278pd          K380xfs     K440SnmpMaster  K570nfs.client  K900swagentd
K280lp          K390rbootd  K450ddfa        K580nis.client  S100localmount
K290hparray     K400iforls  K460sendmail    K590nis.server  S320hostname
K300acct        K410ncs     K470rwhod       K600nfs.core    S400set_prvgrp
```

Figure 4-26 Checking existing script names and link names

```
# /usr/sbin/scsictl -a /dev/rdsk/c0t6d0                    ← Verify new queue depth.
                                  ↖ Character-type device file
immediate_report = 0; queue_depth = 8                      ← Queue depth = 8.
:
:
# /usr/sbin/scsictl -a /dev/rdsk/c0t8d0                    ← Verify new queue depth.
immediate_report = 0; queue_depth = 8                      ← Queue depth = 8.
```

Figure 4-27 Verifying queue depth

Creating and registering the queue depth start-up script

The **queue** (or **hitachi_q**) start-up script sets the queue depth to 2 for all new volumes (SCSI disk devices) on the Hitachi RAID storage system each time the HP-UX system restarts. If the **queue** script exists for a previously installed Hitachi RAID storage system, check the script file to verify that the queue depth value is set to the desired value (see Table 4-4), and add a line for each new volume (see Figure 4-28). If the script does not exist, create and register the script as shown in Figure 4-28. You can use the UNIX **vi** editor or other text editor to create or edit the script.



Note: For questions about creating and registering the start-up script, refer to the UNIX and HPE user documentation, or ask your Hitachi Data Systems representative for assistance.

```
# cp /sbin/init.d/template /sbin/init.d/queue      ← Copy start-up script template file.
# vi /sbin/init.d/queue                              ← Edit script file as shown below.
-------------------------file(/sbin/init.d/queue)----------------------
# !/sbin/sh
#
# @(#) $Revision: 78.1 $
#
# NOTE:    This script is not configurable!  Any changes made to this
#          script will be overwritten when you upgrade to the next
#          release of HP-UX.
#
# WARNING: Changing this script in any way may lead to a system that
#          is unbootable.  Do not modify this script.
#
# <Insert comment about your script here>
#
# Allowed exit values:
#  0 = success; causes "OK" to show up in checklist.
#  1 = failure; causes "FAIL" to show up in checklist.
#  2 = skip; causes "N/A" to show up in the checklist.
#           Use this value if execution of this script is overridden
#     by the use of a control variable, or if this script is not
#     appropriate to execute for some other reason.
#       3 = reboot; causes the system to be rebooted after execution.
# Input and output:
#       stdin is redirected from /dev/null
#       stdout and stderr are redirected to the /etc/rc.log file
#       during checklist mode, or to the console in raw mode.

PATH=/usr/sbin:/usr/bin:/sbin
export PATH

# NOTE: If your script executes in run state 0 or state 1,then /usr
#      might not be available.  Do not attempt to access commands or
#      files in /usr unless your script executes in run state 2 or
#      greater.  Other file systems typically not mounted until run
#      state 2 include /var and /opt.

rval=0

# Check the exit value of a command run by this script.  If non-zero,
# the exit code is echoed to the log file and the return value of this
# script is set to indicate failure.
set_return()  {
       x=$?
       if [ $x -ne 0 ]; then
              echo "EXIT CODE: $x"
              rval=1 # script FAILed
       fi
}
```

Figure 4-28 Example start-up script with changes for Hitachi devices (continues on the following pages)

```
# Kill the named process(es).
# $1=<search pattern for your process>

killproc()  {
       pid='ps -el | awk '( )$NF ~ /'"$1"'/) && ($4 !=mypid) && ($5 !=
mypid)  ){ print $4 }' mypid=$$ '
       if [ "X$pid" != "X" ]; then
              if kill "$pid"; then
                     echo "$1 stopped"
              else
                     rval=1
                     echo "Unable to stop $1"
              fi
       fi
}

case $1 in
'start_msg')
        # Emit a _short_ message relating to running this script with
        # the "start" argument; this message appears as part of the
        # checklist.
        echo "Setting the queue value"                          ← Edit text here.
        ;;
'stop_msg')
       # Emit a _short_ message relating to running this script with
       # the "stop" argument; this message appears as part of the
       # checklist.
       echo "Stopping the <specific> system"
       ;;

'start')
-------------------------------------------------------------------
       # source the system configuration variables                 |
       if [ -f /etc/rc.config ] ; then                           | ← Delete these lines.
              . /etc/rc.config                                   |
       else                                                      |
              echo "ERROR: /etc/rc.config defaults file MISSING" |
       fi                                                        |
                                                                 |
       # Check to see if this script is allowed to run...        |
       if [ "$CONTROL_VARIABLE" != 1 ]; then                     |
              rval=2                                             |
       else                                                      |
                                                                 |
       # Execute the commands to stop your system                |
       :                                                         |
       fi                                                        |
-------------------------------------------------------------------

       /usr/sbin/scsictl -m queue_depth=8 /dev/rdsk/c0t6d0       ← Add one line for each
       /usr/sbin/scsictl -m queue_depth=8 /dev/rdsk/c0t6d1       ←      new disk device.
       /usr/sbin/scsictl -m queue_depth=8 /dev/rdsk/c0t8d0       ←
        :
       ;;
'stop')
```

Figure 4-28 Example start-up script with changes for Hitachi devices (continued)

```
        # source the system configuration variables
        if [ -f /etc/rc.config ] ; then
                . /etc/rc.config
        else
                echo "ERROR: /etc/rc.config defaults file MISSING"
        fi

        # Check to see if this script is allowed to run...
        if [ "$CONTROL_VARIABLE" != 1 ]; then
                rval=2
        else
         :
        # Execute the commands to stop your system

        fi
        ;;
*)
        echo "usage: $0 {start|stop|start_msg|stop_msg}"
        rval=1
        ;;
esac
exit $rval
-------------------------end of file(/sbin/init.d/queue)---------------------

# ls /sbin/rc1.d                                                  ← Check link names.
K230audio       K340xntpd   K420dfs         K475rarpd       K630named       S420set_date
K240auditing    K356vjed    K430dce         K480rdpd        K660net         S440savecore
K250envd        K358egcd    K435OspfMib     K490gated       K700nettl       S500swap_start
K258diagnostic  K360kks     K435SnmpHpunix  K500inetd       K770ptydaemon   S520syncer
K270cron        K370vt      K435SnmpMib2    K510mrouted     K780syslogd
K278pd          K380xfs     K440SnmpMaster  K570nfs.client  K900swagentd
K280lp          K390rbootd  K450ddfa        K580nis.client  S100localmount
K290hparray     K400iforls  K460sendmail    K590nis.server  S320hostname
K300acct        K410ncs     K470rwhod       K600nfs.core    S400set_prvgrp

# ln -s /sbin/init.d/queue /sbin/rc1.d/S890queue                  ← Create link file.
                                          ↖ Be sure this file name does not already exist.
```

Figure 4-28 Example start-up script with changes for Hitachi devices (continued)

Creating mount directories

After you create the file systems and set the device parameters, create the mount directory for each volume. Choose a unique name for each mount directory that identifies the logical volume.

To create the mount directories:

1. Use the **mkdir** command to create the mount directory with the new mount directory name as the argument (see Figure 4-29).
2. Use the **ls -x** command to verify the new mount directory (see Figure 4-29).
3. Repeat steps 1 and 2 for each new device on the Hitachi RAID storage system.

   If you need to delete a mount directory, use the **rmdir** command.

```
# mkdir /VSP-LU00                                    ← Create new mount directory.
# ls -x                                              ← Verify new mount directory.
VSP-LU00    bin     dev       device    etc       export
floppy      home    hstsboof  kadb      kernel    lib
#
```

Figure 4-29 Creating and verifying a mount directory

Mounting and verifying file systems

After you create the mount directories, mount the file system for each new logical volume and verify the file systems.

To mount and verify the file systems:

1. Use the **mount** command to mount the file system for the volume (see Figure 4-30).
2. Repeat step 1 for each new logical volume on the Hitachi RAID storage system.
3. Use the **bdf** command to verify that the file systems are correct (see Figure 4-31). Be sure the capacity (listed under **Kbytes**) is correct for each device.
4. Perform basic UNIX operations, such as file creation, copying, and deletion, on each logical device to be sure the new devices on the Hitachi RAID storage system are fully operational (see Figure 4-32).
5. If you want to unmount a file system after it has been mounted and verified, use the **umount** command (for example, **umount /VSP-LU00**).

```
# mount /dev/vg06/lvol1 /VSP-LU00                    ← Mount file system.
        ↖ Block-type lvol name  ↖ Mount directory name
#
```

Figure 4-30 Mounting a file system

```
# bdf                                                ← Verify file systems.
Filesystem          Kbytes    used   avail  %used  Mounted on
/dev/vg00/lvol1      59797   59364       0  100%   /
 :
/dev/vg06/lvol1    2348177       9 2113350    0%   /VSP-LU00          ← OPEN-3
/dev/vg07/lvol1    2348177       9 2113350    0%   /VSP-LU01          ← OPEN-3
/dev/vg08/lvol1    7052764       9 6347478    0%   /VSP-LU02          ← OPEN-9
```

Figure 4-31 Verifying file systems

```
# mount /dev/vg06/lvol1 /VSP-LU00                                  ← Mount LUN.
# cd /VSP-LU00                                                ← Go to LUN mount directory.
# cp /bin/vi /VSP-LU00/vi.back1                               ← Copy any file to LUN.
# ll                                                          ← Verify file copy.
drwxr-xr-t   2 root     root        8192 Mar 15 11:35  lost+found
-rwxr-xr-x   1 root     sys       217088 Mar 15 11:41  vi.back1
# cp vi.back1 vi.back2                                        ← Copy file again.
# ll                                                          ← Verify second file copy.
drwxr-xr-t   2 root     root        8192 Mar 15 11:35  lost+found
-rwxr-xr-x   1 root     sys       217088 Mar 15 11:41  vi.back1
-rwxr-xr-t   1 root     sys       217088 Mar 15 11:52  vi.back2
# rm vi.back1                                                 ← Delete first test file.
# rm vi.back2                                                 ← Delete second test file.
```

Figure 4-32 Final verification of a file system for one volume

Setting and verifying auto-mount parameters

The final step in configuring the Hitachi RAID storage system volumes for LVM operations is to set up and verify the auto-mount parameters for each new volume. The **/etc/fstab** file contains the auto-mount parameters for the logical volumes. If you do not plan to auto-mount the new devices, you can skip this section.

To set and verify the auto-mount parameters:

1. Edit the **/etc/fstab** file to add a line for each new volume (SCSI disk device) on the Hitachi RAID storage system (see Figure 4-33). Table 4-5 shows the auto-mount parameters.
2. After you finish editing the **/etc/fstab** file, reboot the HP-UX system. If you cannot reboot at this time, issue the **mount -a** command.
3. Use the **bdf** command to verify the device file systems again (see Figure 4-31).

```
# cp -ip /etc/fstab /etc/fstab.standard                    ← Make backup before editing.
# vi /etc/fstab                                              ← Edit the file (vi shown).
/dev/vg00/lvol1  /            vxfs    rw         0    1    # root
/dev/vg00/lvol2  swap         ignore sw          0    0    # primary swap
 :
/dev/vg06/lvol1  /VSP-LU00    vxfs    defaults   0    2    # VSP-LU00
/dev/vg06/lvol2  /VSP-LU01    vxfs    defaults   0    2    # VSP-LU01
      ①           ②           ③       ④        ⑤    ⑥       ⑦          ← See Table 4-5.
```

Figure 4-33 Setting auto-mount parameters

Table 4-5 Auto-mount parameters

Parameter #	Name	Enter:
①	Device to mount	Block-type device file name
②	Mount point	Mount directory name
③	File system	Type of file system (for example, vxfs)
④	Mount options	Usually “defaults”
⑤	Enhance	“0”
⑥	File system check (fsck pass)	Order for performing file system checks
⑦	Comment	Any comment statement

Online device installation

After initial installation and configuration of the Hitachi RAID storage system, additional devices can be installed or de-installed online without having to restart the HP-UX system. This procedure should be performed by the system administrator (that is, super-user).

Use the normal disruptive device configuration procedure in the following cases:

- **Fibre:** If a new fibre-channel connection is being installed. New fibre-channel connections can only be installed when the host system is powered off. New devices under existing fibre-channel ports can be installed and configured nondisruptively.
- **Maxvgs:** If the **maxvgs** parameter needs to be changed. The procedure for changing the **maxvgs** value in the system kernel requires a system reboot.

To perform online device installation and configuration:

1. Verify that the new devices on the Hitachi RAID storage system are ready to be configured. The Hitachi Data Systems representative should have completed hardware installation and verified the normal status of the new devices (see Installing the Hitachi RAID storage system).
2. Be sure that you are logged in as **root**.
3. Enter the **insf -e** command to perform online device recognition. The **insf -e** command creates device files for the new devices on the existing fibre busses (see Creating device files).
4. Configure the new disk devices for HP-UX operations described in HP-UX configuration and attachment. For raw/FX devices, you only need to verify the device files and driver. Do not partition or create a file system on any raw/FX device.
5. Configure the application failover, path failover (that is, **vgextend**), and/or SNMP software on the HP-UX system as needed to recognize the new disk devices. For additional information about online installation and reinstallation of LUs, see the Maintenance Manual for the storage system.

Troubleshooting for HP-UX host attachment

Table 3-11 lists potential error conditions that might occur during storage system installation on an HP-UX host and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 4-6 Troubleshooting for HP-UX host attachment

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Make sure that the READY indicator lights on the storage system are ON. Make sure that the FC cables are correctly installed and firmly connected. Make sure that LUSE devices are not intermixed with normal LUs on the same fibre-channel port. Verify that LUNs are configured properly for each TID. Run **sr-probe** to recheck the fibre channel for new devices.
A physical volume cannot be created (PVCREATE command).	Ensure the Hitachi RAID storage system devices are properly formatted. Ensure the character-type device file exists. Ensure the correct character-type device file name is used with **pvcreate**.
A volume group cannot be created (VGCREATE command).	Ensure the directory for the new volume group exists. Ensure the control file exists. Ensure the correct major # (64) and minor # are used with **mknod**. Ensure the block-type file exists and is entered correctly with **vgcreate**. Ensure the physical volume is not already allocated to another volume group.
A logical volume cannot be created (LVCREATE command).	Ensure the specified capacity is not greater than 4096 MB. Ensure the capacity of the volume group is not less than the capacity of the partitioned logical volume.
File system cannot be created (**newfs**).	Ensure the character-type device file is entered correctly with **newfs**.
The file system is not mounted after rebooting.	Ensure the system was restarted properly. Ensure the auto-mount information in the /**etc/fstab** file is correct.
The HP-UX system does not reboot properly after hard shutdown.	If the HP-UX system is powered off without executing the shutdown process, wait **three minutes** before restarting the HP-UX system. This allows the Hitachi RAID storage system internal time-out process to purge all queued commands so that the storage system is available (not busy) during system startup. If the HP-UX system is restarted too soon, the Hitachi RAID storage system will continue trying to process the queued commands and the HP-UX system will not reboot successfully.

5

Red Hat Linux configuration and attachment

This chapter describes how to configure the new Hitachi disk devices on a Red Hat Linux host:

- ☐ Hitachi storage system configuration for Red Hat Linux operations
- ☐ Device Mapper (DM) Multipath
- ☐ Verifying new device recognition
- ☐ Configuring the new devices
- ☐ Troubleshooting for Red Hat Linux host attachment



Note: Configuration of the devices should be performed by the Linux system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Hitachi storage system configuration for Red Hat Linux operations

The storage system must be fully configured before being attached to the Red Hat Linux host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for Red Hat Linux operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)
- Cross-OS File Exchange (FX) (3390-3A/B/C, OPEN-*x*-FXoto)

Host mode. The required host mode for Red Hat Linux is **00**. Do not select a host mode other than **00** for Red Hat Linux. For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Veritas Cluster Server: See Note on using Veritas Cluster Server for important information about using Veritas Cluster Server.

Device Mapper (DM) Multipath configuration

The Hitachi RAID storage systems support DM Multipath operations for Red Hat Enterprise Linux (RHEL) version 5.4 X64 or X32 or later.



Note: Refer to the native multipath configuration documentation from the OS vendor for important information about required settings and parameters for DM Multipath operations, including but not limited to:

- Disabling the HBA failover function
- Installing the `kpartx` utility
- Creating the multipath device with the `multipath` command
- Editing the `/etc/modprobe.conf` file
- Editing the `/etc/multipath.conf` file
- Configuring LVM
- Configuring raw devices
- Creating partitions with DM Multipath

Verifying new device recognition

The final step before configuring the new 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 (see Figure 5-1).
2. Record the device file name for each new device. You will need this information when you partition the devices (see Verifying new device recognition). See Table 5-1 for a sample SCSI path worksheet.
3. The device files are created under the **/dev** directory. Verify that a device file was created for each new disk device (see Figure 5-2).

```
# dmesg | more
:
:
scsi0 : Qlogic QLA2200 PCI to Fibre Channel Host Adapter: 0 device 14 irq 11
       Firmware version:  1.17.26, Driver version  2.11 Beta

scsi : 1 host.
  Vendor: HITACHI        Model: OPEN-3            Rev: 0111
  Type:   Direct-Access                       ANSI SCSI revision: 02
Detected scsi disk sda at scsi0, channel 0, id 0, lun 0
         ↖ Device file name of this disk = /dev/sda   ↖ Logical unit number

  Vendor: HITACHI        Model: OPEN-9            Rev: 0111
  Type:   Direct-Access                       ANSI SCSI revision: 02
Detected scsi disk sdb at scsi0, channel 0, id 0, lun 1
:
:
```

In this example, the HITACHI OPEN-3 device (TID 0, LUN 0) and the HITACHI OPEN-9 device (TID 0, LUN 1) are recognized by the Red Hat Linux server.

Figure 5-1 Example of verifying new device recognition

```
# ls -l /dev | more
:
brw-rw----   1 root    disk    8,  0 May  6  1998 sda      ← Device file = sda
```

Figure 5-2 Example of verifying device files

Table 5-1 Sample SCSI path worksheet

LDEV (CU:LDEV)	Device Type	LUSE (*n)	VLL (MB)	Device File Name	Path	Alternate Path
0:00					TID:_____ LUN:_____	TID:_____ LUN:_____
0:01					TID:_____ LUN:_____	TID:_____ LUN:_____
0:02					TID:_____ LUN:_____	TID:_____ LUN:_____
0:03					TID:_____ LUN:_____	TID:_____ LUN:_____
0:04					TID:_____ LUN:_____	TID:_____ LUN:_____
0:05					TID:_____ LUN:_____	TID:_____ LUN:_____
0:06					TID:_____ LUN:_____	TID:_____ LUN:_____
0:07					TID:_____ LUN:_____	TID:_____ LUN:_____
0:08					TID:_____ LUN:_____	TID:_____ LUN:_____
0:09					TID:_____ LUN:_____	TID:_____ LUN:_____
0:0A					TID:_____ LUN:_____	TID:_____ LUN:_____
0:0B					TID:_____ LUN:_____	TID:_____ LUN:_____
0:0C					TID:_____ LUN:_____	TID:_____ LUN:_____
0:0D					TID:_____ LUN:_____	TID:_____ LUN:_____
0:0E					TID:_____ LUN:_____	TID:_____ LUN:_____
0:0F					TID:_____ LUN:_____	TID:_____ LUN:_____

Configuring the new devices

This section describes how to configure the new disk devices on the Red Hat Linux system host:

- ☐ Setting the number of logical units
- ☐ Partitioning the devices
- ☐ Creating, mounting, and verifying the file systems

Setting the number of logical units

To set the number of LUs:

1. Edit the **/etc/modules.conf** file to add the following line:

   options scsi_mod max_scsi_luns=*xx*

   where ***xx*** is the maximum number of LUs supported by your Linux OS. Check your host adapter documentation and Linux system documentation to determine the total number of devices that can be supported.
2. To set the Emulex Driver, add the following line to the **/etc/modules.conf** file, as shown in Figure 5-3:

   Alias scsi_hostadapter lpfcdd
3. To activate the above modification, make an image file for booting.

   Example: **# mkinitrd /boot/initrd-2.4.x.scsiluns.img 'uname -r'**
4. Use one of the following methods to change the setting of Bootloader:
   a. LILO used as Bootloader. Edit the **lilo.conf** file as shown in Figure 5-4, then issue the **lilo** command to activate the **lilo.conf** setting with selecting the label. Example: **# lilo**
   b. Grand Unified Bootloader (GRUB) is used as Bootloader. Edit the **/boot/grub/grub.conf** file as shown in Figure 5-5.
5. Reboot the system.

```
Alias scsi_hostadapter lpfcdd                    ← Add this to /etc/modules.conf.
```

Figure 5-3 Example of setting the Emulex driver

```
image=/boot/vmlinuz-qla2x00
       label=Linux-qla2x00
       append="max_scsi_luns=16"
       # initrd=/boot/initrd-2.4.x.img                    ← Comment out this line.
       initrd=/boot/initrd-2.4.x.scsiluns.img                  ← Add this line.
 root=/dev/sda7
 read-only
#sbin/lilo
```

Figure 5-4 Example of setting the number of LUs (LILO)

```
kernel /boot/vmlinuz-2.4.x ro root=/dev/hda1
# initrd /boot/initrd-2.4.x.img                          ← Comment out this line.
initrd /boot/initrd-2.4.x.scsiluns.img                        ← Add this line.
```

Figure 5-5 Example of setting the number of LUs (GRUB)

Partitioning the devices

After the setting the number of logical units, you need to create the partitions on the new disk devices.



Note: For important information about required settings and parameters for DM Multipath operations, refer to the native multipath configuration documentation from the OS vendor.

To create the partitions on the new 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.



Tip: Other useful commands include **d** to remove partitions and **q** to stop a change.

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

Creating, mounting, and verifying the file systems

Creating the file systems

After you partition the devices, create the file systems. Be sure the file system are appropriate for the primary and/or extended partition for each logical unit.

To create the file system, issue the **mkfs** command:

```
# mkfs /dev/sda1
```

where **/dev/sda1** is device file of primary partition number 1.

Creating the mount directories

To create the mount directories, issue the **mkdir** command:

```
# mkdir /VSP-LU00
```

Mounting the new file systems

Use the **mount** command to mount each new file system (see example in Figure 5-6). The first parameter of the **mount** command is the device file name (**/dev/sda1**), and the second parameter is the mount directory, as shown in Figure 5-6.

```
# mount /dev/sda1 /VSP-LU00
        ↖ Device file name  ↖ Mount directory name
#
```

Figure 5-6 Example of mounting the new devices

Verifying the file systems

After mounting the file systems, verify the file systems (see the example in Figure 5-7).

```
# df -h
Filesystem            Size      Used     Avail    Used%    Mounted on
/dev/sda1             1.8G      890M      866M     51%      /
/dev/sdb1             1.9G      1.0G      803M     57%      /usr
/dev/sdc1             2.2G       13k      2.1G      0%      /VSP-LU00
#
```

Figure 5-7 Example of verifying the file system

Setting the auto-mount parameters

To set the auto-mount parameters, edit the **/etc/fstab** file (see the example in Figure 5-8).

```
# cp -ip /etc/fstab /etc/fstab.standard                    ← Make a backup of /etc/fstab.
# vi /etc/fstab                                                ← Edit /etc/fstab.
:
/dev/sda1          /VSP-LU00 ext2     defaults    0    2          ← Add new device.
```

Figure 5-8 Example of setting the auto-mount parameters

Troubleshooting for Red Hat Linux host attachment

Table 5-2 lists potential error conditions that might occur during storage system installation on a Red Hat Linux host and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 5-2 Troubleshooting for Red Hat Linux host attachment

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Be sure that the READY indicator lights on the Hitachi RAID storage system are ON. Be sure that the LUNs are properly configured. The LUNs for each target ID must start at 0 and continue sequentially without skipping any numbers.
The file system cannot be created.	Be sure that the device name is entered correctly with **mkfs**. Be sure that the LU is properly connected and partitioned.
The file system is not mounted after rebooting.	Be sure that the system was restarted properly. Be sure that the auto-mount information in the **/etx/fstab** file is correct.

6

Solaris configuration and attachment

This chapter describes how to configure the new Hitachi disk devices on a Solaris host:

- ☐ Hitachi storage system configuration for Solaris operations
- ☐ FCA configuration for Solaris
- ☐ Configuring the new devices
- ☐ ALUA support on Solaris
- ☐ Troubleshooting for Solaris host attachment
- ☐ Online device installation
- ☐ Using MPxIO path failover software



Note: Configuration of the devices should be performed by the Solaris system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Hitachi storage system configuration for Solaris operations

The storage system must be fully configured before being attached to the Solaris host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for Solaris operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)
- Cross-OS File Exchange (FX) (3390-3A/B/C, OPEN-*x*-FXoto)

Host mode. The required host mode for Solaris is **09**. Do not select a host mode other than **09** for Solaris. For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).



Note: You must set **HOST MODE=09** before installing Sun Cluster, or the Quorum Device will not be assigned to the Hitachi RAID storage system.

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Veritas Cluster Server: See Note on using Veritas Cluster Server for important information about using Veritas Cluster Server.

FCA configuration for Solaris

This section describes how to configure the fibre-channel adapters (FCAs) that will be attached to the Solaris host.

- Verifying the FCA installation
- Setting the disk and device parameters

Verifying the FCA installation

Before configuring the fibre-channel HBAs, verify the HBA installation and recognition of the fibre-channel HBA and driver.

1. Log in to the Solaris system as **root**, and confirm that all existing devices are powered on and properly connected to the Solaris system.
2. Display the host configuration using the **dmesg** command (see Figure 6-1). The fibre information (underlined in the following example) includes the recognition of the **fibre channel adapter**, **SCSI bus characteristics**, **world wide name**, and **FCA driver**. Ensure the host recognizes these four classes. If this information is not displayed or if error messages are displayed, the host environment may not be configured properly.

```
# dmesg

Nov  9 23:14
ems, Inc.
mem = 65536K (0x4000000)
avail mem = 60129280
Ethernet address = 8:0:20:92:32:48
root nexus = Sun Ultra 1 SBus (UltraSPARC 167MHz)
sbus0 at root: UPA 0x1f 0x0 ...
espdma0 at sbus0: SBus0 slot 0xe offset 0x8400000
esp0:   esp-options=0x46
esp0 at espdma0: SBus0 slot 0xe offset 0x8800000 Onboard device sparc9 ipl 4
sd0 at esp0: target 0 lun 0
sd0 is /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@0,0
        <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
sd6 at esp0: target 6 lun 0
sd6 is /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@6,0
fca0: JNI Fibre Channel Adapter (1062 MB/sec), model FC                    <- Verify that
fca0:   SBus 1: IRQ 4: FCODE Version 11.0.9 [1a6384]: SCSI ID 125: AL PA 01 <- these items
fca0:   Fibre Channel WWN: 100000e0690000d5                                <- are listed.
fca0:   FCA Driver Version 2.2.HIT.03, Oct 09, 1999 Solaris 2.5, 2.6

fca0:   All Rights Reserved.
fca0:   < Total IOPB space used: 1125824 bytes >
fca0:   < Total DMA space used: 565277 bytes >
root on /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@0,0:a fstype ufs
zs0 at sbus0: SBus0 slot 0xf offset 0x1100000 Onboard device sparc9 ipl 12
zs0 is /sbus@1f,0/zs@f,1100000
zs1 at sbus0: SBus0 slot 0xf offset 0x1000000 Onboard device sparc9 ipl 12
zs1 is /sbus@1f,0/zs@f,1000000
keyboard is </sbus@1f,0/zs@f,1000000> major <29> minor <2>
mouse is </sbus@1f,0/zs@f,1000000:b> major <29> minor <3>
stdin is </sbus@1f,0/zs@f,1000000> major <29> minor <2>
 . . . . . .
```

Figure 6-1 Displaying the fibre device information (Jaycor FC-1063)

Setting the disk and device parameters

The queue depth (**max_throttle**, **max_pending** for Solaris ZFS) for the Hitachi RAID storage system devices must be set as specified in Table 6-1. You can adjust the queue depth for the devices later as needed (within the specified range) to optimize the I/O performance. For details about queue depth, see Host queue depth.

Table 6-1 Queue depth requirements for Solaris

Parameter	Recommended value for VSP G1000, VSP G1500, VSP F1500, VSP Gx00, VSP Fx00, HUS VM, VSP	Requirements for USP V/VM
Queue depth	32 per LU 2048 per port	queue_depth ≤ 32 (# of LUs) × (queue_depth) ≤ 2048 For USP V/VM, it is recommended that queue_depth be specified between 8 and 16 per LU.

Caution: Inappropriate settings, including max_pending/throttle and number of LUNs per ZFS pool, can significantly impact the SAN environment (for example, C3 discards). If you have any questions or concerns, contact your Hitachi Data Systems representative for important information about these settings.

The required I/O time-out value (TOV) for Hitachi RAID storage system devices is 60 seconds (default TOV=60). If the I/O TOV has been changed from the default, change it back to 60 seconds by editing the **sd_io_time** or **ssd_io_time** parameter in the **/etc/system** file.

Several other parameters (for example, FC fibre support) may also need to be set. See the user documentation for the HBA to determine whether other options are required to meet your operational requirements.

Use the same settings and device parameters for all Hitachi RAID storage system devices. For fibre-channel, the settings in the system file apply to the entire system, not to just the HBAs.

To set the queue depth and I/O TOV:

1. Make a backup of the **/etc/system** file:
   ```
   cp /etc/system /etc/system.old
   ```
2. Edit the **/etc/system** file.
3. To set the TOV, add the following to the **/etc/system** file (see Figure 6-2): `set sd:sd_io_time=0x3c`

   For Sun generic HBA: `set ssd:ssd_io_time=0x3c`
4. To set the queue depth, add the following to the **/etc/system** file (see Figure 6-3): `set sd:sd_max_throttle=x` (for **x** see Table 6-1)

   For Sun generic HBA: `set ssd:ssd_max_throttle=x`

For Solaris ZFS: **`set zfs:zfs_vdev_max_pending=x`**

5. Save your changes, and exit the text editor.
6. Shut down and reboot to apply the I/O TOV setting.

```
*ident  "@(#)system     1.18    97/06/27 SMI" /* SVR4 1.5 */
*
* SYSTEM SPECIFICATION FILE
*
:
*       To set a variable named 'debug' in the module named 'test_module'
*
*               set test_module:debug=0x13
                set sd:sd_io_time=0x3c                  ← Add this line to /etc/system
                set ssd:ssd_io_time=0x3c                ← Add this line to /etc/system
                                                                (for Sun generic HBA)
```

Figure 6-2 Setting the I/O TOV

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

               set sd:sd_max_throttle=8                 ← Add this line to /etc/system
               set ssd:ssd_max_throttle=8               ← Add this line to /etc/system
                                                                (for Sun HBA)
               set vdev:vdev_max_pending=8              ← Add this line to /etc/system
                                                              (for Solaris ZFS)
```

Figure 6-3 Setting the queue depth

Configuring the new devices

This chapter describes how to configure the new disk devices that you attached to the Solaris system:

- ☐ Setting and recognizing the LUs
- ☐ Verifying recognition of new devices
- ☐ Partitioning and labeling the new devices
- ☐ Creating and mounting the file systems

Setting and recognizing the LUs

Once the Hitachi RAID storage system is installed and connected, set and recognize the new LUs by adding the logical devices to the **sd.conf** file (**/kernel/drv/sd.conf**). The **sd.conf** file includes the SCSI TID and LUN for all LDEVs connected to the Solaris system. After editing the **sd.conf** file, you will halt the system and reboot.

To set and recognize LUs:

1. Log in as root, and make a backup copy of the **/kernel/drv/sd.conf file**:

   ```
   cp -ip /kernel/drv/sd.conf /kernel/drv/sd.conf.standard
   ```

2. Edit the **/kernel/drv/sd.conf file** as shown in Figure 6-4. Be sure to make an entry (SCSI TID and LUN) for each new device being added to the Solaris system.

   If the LUs have already been added to the **sd.conf** file, verify each new LU.

3. Exit the vi editor by entering the command:

   ESC + `:wq`

4. Halt the Solaris system:

   ```
   halt
   ```

5. Reboot the Solaris system:

   ```
   boot -r
   ```

6. Log in to the system as root, and verify that the system recognizes the Hitachi RAID storage system (see Figure 6-5):

   ```
   dmesg | more
   ```

7. Verify that the vendor name, product name, and number of blocks match the values shown in Figure 6-5.

```
# cp -ip /kernel/drv/sd.conf /kernel/drv/sd/conf/standard      ← Make backup of file.
#
# vi /kernel/drv/sd.conf                                       ← Edit the file (vi shown).
#ident  "@(#)sd.conf    1.8     93/05/03 SMI"
name="sd" class="scsi"                                         ← The SCSI class type name
        target=0 lun=0;                                          is used because the SCSI
                                                                 driver is used for fibre
name="sd" class="scsi"                                           channel.
        target=1 lun=0;

name="sd" class="scsi"
        target=2 lun=0;

name="sd" class="scsi"                                         ← Add this information for
        target=2 lun=1;                                          all new target IDs
                                                                 and LUNs.
name="sd" class="scsi"
        target=3 lun=0;

name="sd" class="scsi"
        target=4 lun=0;
#
# halt                                                         ← Enter halt.
Jan 11 10:10:09 sunss20 halt:halted by root
Jan 11 10:10:09 sunss20 syslogd:going down on signal 15
Syncing file systems... done
Halted

Program terminated
Type help for more information
OK

volume management starting.
The system is ready.

host console login: root                                       ← Log in as root.
Password:                                                      ← Password is not displayed.
Oct 11 15:28:13 host login: ROOT LOGIN /dev/console
Last login:Tue Oct 11 15:25:12 on console
Sun Microsystems inc.  SunOS 5.5       Generic September 1993
#
#
#
```

Figure 6-4 Setting and recognizing LUs

```
# dmesg | more
  :
sbus0 at root: UPA 0x1f 0x0 ...
fas0:  rev 2.2 FEPS chip

SUNW,fas0 at sbus0: SBus0 slot 0xe offset 0x8800000 and slot 0xe offset 0x8810000 Onboard
device sparc9 ipl 4
SUNW,fas0 is /sbus@1f,0/SUNW,fas@e,8800000
sd0 at SUNW,fas0: target 0 lun 0
sd0 is /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
     <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
sd6 at SUNW,fas0: target 6 lun 0
sd6 is /sbus@1f,0/SUNW,fas@e,8800000/sd@6,0
WARNING: fca0: fmle: sc1: 000e0000  sc2: 00000000
fca0: JNI Fibre Channel Adapter (1062 MB/sec), model FC
fca0:  SBus 1 / IRQ 4 / FCODE Version 10 [20148b] / SCSI ID 125 / AL_PA 0x1
fca0:  Fibre Channel WWN: 100000e0690002b7
fca0:  FCA Driver Version 2.1+, June 24, 1998 Solaris 2.5, 2.6
fca0:  All Rights Reserved.
fca0:  < Total IOPB space used: 1100624 bytes >
fca0:  < Total DMA space used: 532644 bytes >
fca0: <HITACHI :OPEN-3         :5235> target 2 (alpa 0xe4) lun 0 online
sd192 at fca: target 2 lun 0
                      ↖ LUN = 0
                ↖ target ID = 2
sd192 is /sbus@1f,0/fca@1,0/sd@2,0

WARNING: /sbus@1f,0/fca@1,0/sd@2,0 (sd192)
       corrupt label - wrong magic number                    ← Not yet labeled.
        Vendor 'HITACHI', product 'OPEN-3', 4806720 512 byte blocks
                ↖ Vendor name              ↖ Number of blocks
                                  ↖ Product name
fca0: <HITACHI :OPEN-3          :5235> target 2 (alpa 0xdc) lun 2 online
sd193 at fca: target 2 lun 1           (LUN=1, target ID=2)
sd193 is /sbus@1f,0/fca@1,0/sd@2,1
WARNING: /sbus@1f,0/fca@1,0/sd@2,1 (sd193)
        corrupt label - wrong magic number
        Vendor 'HITACHI', product 'OPEN-3', 4806720 512 byte blocks
fca0: <HITACHI :OPEN-9          :5235> target 6 (alpa 0xdc) lun 0 online
sd.. at fca: target  lun 0   (LUN=0, target ID=6)
sd.. is /sbus@1f,0/fca@1,0/sd@4,0
WARNING: /sbus@1f,0/fca@1,0/sd@4,0 (sd..)
        corrupt label - wrong magic number                   ← Not yet labeled.
        Vendor 'HITACHI', product 'OPEN-9', 14423040 512 byte blocks
sd.. at fca: target 6 lun 0                                  ← Verify target ID.
        corrupt label - wrong magic number                   ← Not yet labeled.
        Vendor 'HITACHI', product 'OPEN-9', 14423040 512 byte blocks
sd.. is /sbus@1f,0/fca@1,0/sd@5,0
WARNING: /sbus@1f,0/fca@1,0/sd@5,0 (sd..)
        corrupt label - wrong magic number                   ← Not yet labeled.
        Vendor 'HITACHI', product '3390-3B', 5822040 512 byte blocks
sd.. is /sbus@1f,0/fca@1,0/sd@6,0
WARNING: /sbus@1f,0/fca@1,0/sd@6,0 (sd..)
        corrupt label - wrong magic number                   ← Not yet labeled.
        Vendor 'HITACHI', product '3390-3A', 5825520 512 byte blocks
sd.. is /sbus@1f,0/fca@1,0/sd@8,0
```

Figure 6-5 Fibre device recognition



Note: If the FX volumes (for example, 3390-3A/B/C) are customized, their block number may be lower than the number displayed in this example.

Verifying recognition of new devices

After system start-up, log in as root and use the **`dmesg | more`** command to verify that the Solaris system recognizes the Hitachi storage system. Confirm that the displayed vendor names, product names, and number of blocks match the values in Figure 6-6. If the results are different than the intended system configuration, the path definition or fibre cabling might be wrong.

Note: When the Solaris system accesses the multiplatform devices, the message "Request sense couldn't get sense data" may be displayed. You can disregard this message.

```
# dmesg | more
  :
sbus0 at root: UPA 0x1f 0x0 ...
fas0:  rev 2.2 FEPS chip

SUNW,fas0 at sbus0: SBus0 slot 0xe offset 0x8800000 and slot 0xe offset 0x8810000 Onboard device
sparc9 ipl 4
SUNW,fas0 is /sbus@1f,0/SUNW,fas@e,8800000
sd0 at SUNW,fas0: target 0 lun 0
sd0 is /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
     <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
sd6 at SUNW,fas0: target 6 lun 0
sd6 is /sbus@1f,0/SUNW,fas@e,8800000/sd@6,0
WARNING: fca0: fmle: sc1: 000e0000  sc2: 00000000
fca0: JNI Fibre Channel Adapter (1062 MB/sec), model FC
fca0: SBus 1 / IRQ 4 / FCODE Version 10 [20148b] / SCSI ID 125 / AL_PA 0x1
fca0: Fibre Channel WWN: 100000e0690002b7
fca0: FCA Driver Version 2.1+, June 24, 1998 Solaris 2.5, 2.6
fca0: All Rights Reserved.
fca0: < Total IOPB space used: 1100624 bytes >
fca0: < Total DMA space used: 532644 bytes >
fca0: <HITACHI :OPEN-3          :5235> target 2 (alpa 0xe4) lun 0 online
sd192 at fca: target 2 lun 0
                     ↖      ↖ LUN = 0
                       target ID = 2
sd192 is /sbus@1f,0/fca@1,0/sd@2,0
WARNING: /sbus@1f,0/fca@1,0/sd@2,0 (sd192)
        corrupt label - wrong magic number                          ← Not yet labeled.
        Vendor 'HITACHI', product 'OPEN-3',    4806720 512 byte blocks
                  ↖ Vendor name    ↖ Product name   ↖ Number of blocks
fca0: <HITACHI :OPEN-3          :5235> target 2 (alpa 0xdc) lun 1 online
sd193 at fca: target 2 lun 1           (LUN=1, target ID=2)
sd193 is /sbus@1f,0/fca@1,0/sd@2,1
WARNING: /sbus@1f,0/fca@1,0/sd@2,1 (sd193)
        corrupt label - wrong magic number
        Vendor 'HITACHI', product 'OPEN-3', 4806720 512 byte blocks
```

This example shows two new disks on fca@1: target ID is 2, LUNs are 0 and 1, vendor name is "HITACHI", product name is "OPEN-3", and number of blocks is 4806720. LUNs 0 and 1 are assigned as device names sd192 and sd193, respectively. Details for other disks:

- vendor name "HITACHI", product name "OPEN-9" and 14423040 512-byte blocks
- vendor name "HITACHI", product name "3390-3B" and 5822040 512-byte blocks
- vendor name "HITACHI", product name "3390-3A" and 5825520 512-byte blocks

Figure 6-6 Verifying new devices

Partitioning and labeling the new devices

After the Solaris system recognizes the new devices, partition and label the devices. All new devices, including all SCSI disk devices and FX devices, must be partitioned and labeled using the **format** utility (see ***WARNING*** below).

- Each SCSI disk device (for example, OPEN-x) can have more than one partition.
- Each FX device (for example, 3390-3A) must have one partition of fixed size.

The disk partitioning and labeling procedure involves the following tasks:

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

A good way to partition and label the disks is to partition and label all devices of one type (for example, OPEN-3), then all devices of the next type (for example, OPEN-9), and so on until you partition and label all new devices. You will enter this information into the Solaris system during the disk partitioning and labeling procedure.



WARNING: Be extremely careful when using the Solaris **format** utility. Do not use any **`format`** commands not described in this document. The **format** utility is designed for Sun disks. Some **`format`** commands are not compatible with the Hitachi RAID storage system and can overwrite the data on the disk. The Hitachi RAID storage system will not respond to the **`format`** command (devices are formatted using the SVP), and will not report any defect data in response to the **`defect`** command.

To partition and label the new devices/disks:

1. Enter **`format`** at the root prompt to start the **format** utility (see Figure 6-7).
   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). For troubleshooting information see Troubleshooting for Solaris host attachment.
   b. Write down the character-type device file names (for example, c1t2d0) for all of the new 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. Remember the device type of this device (for example, OPEN-3).

3. When prompted to label the disk, 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 in Table 1-2 (vendor name + product name, for example, HITACHI OPEN-3).
5. If the disk type for the selected device is already defined, enter the number for that disk type and skip to step 7.

Note:

- Do not use HITACHI-OPEN-*x*-0315, HITACHI-3390-3A/B-0315. These disk types are created automatically by the Solaris system and cannot be used for the Hitachi RAID storage system devices.
- LU capacity must be less than 1 TB. In case of selecting other type, the disk type parameters described below cannot be set for an LU larger than 32,767 data cylinders.

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 6-8.
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 6-9 and Table 6-2 through Table 6-9.
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.

Note: This step does not apply to the multiplatform devices (for example, 3390-3A/B/C), because these devices can only have one partition of fixed size.

13. After setting the partitions for the selected device, enter **label** at the **partition>** prompt, and enter **y** to label the device (see Figure 6-10).

Note: The Solaris system displays the following warnings when an FX device (for example, 3390-3A/B/C) is labeled. You can ignore these warnings.

```
Warning: error warning VTOC.
Warning: no backup labels.
Label failed.
```

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. Verify that 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 device type is defined (for example, HITACHI OPEN-3), you can label all devices of that same type without having to enter the parameters (skipping steps 6 and 7). For this reason, you may want to label the devices by type (for example, labeling all OPEN-3 devices, then all OPEN-9 devices, and so on) 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

c1t2d0: configured with capacity of 2.29GB   (OPEN-3)  ← These devices are not yet labeled.
c1t2d1: configured with capacity of 2.29GB   (OPEN-3)  ←
c2t4d0: configured with capacity of 6.88GB   (OPEN-9)  ←
c2t5d0: configured with capacity of 2.77GB   (3390-3B) ←
c2t6d0: configured with capacity of 2.78GB   (3390-3A) ←
  ↖ These character-type device file names are used later to create the file systems.

AVAILABLE DISK SELECTIONS:

       0. c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>                ← Already labeled.
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0
       1. c0t3d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>                ←Already labeled.
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@3,0
       2. c1t2d0 <HITACHI-OPEN-3-52-34                                ← Not yet labeled:
             ↖ LUN          ↖ Product version
           -       ↖ Vendor  ↖ Product ID
          –   ↖ Target Id
            ↖ Logical Controller ID
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@2,0    OPEN-3, TID=2, LUN=0
       3. c1t2d1 <HITACHI-OPEN-3-52-34                                    ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@2,1    OPEN-3, TID=2, LUN=1
       4. c1t4d0 <HITACHI-OPEN-9-52-34                                    ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@4,0    OPEN-9, TID=4, LUN=0
       5. c1t5d0 <HITACHI-3390-3B-52-34                                  ←Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@5,0   3390-3B, TID=5, LUN=0
       6. c1t6d0 <HITACHI-3390-3A-52-34                                   ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@6,0   3390-3A, TID=6, LUN=0
Specify disk (enter its number): 2                                        ← Select device.
selecting c1t2d0
[disk formatted]
Disk not labeled.  Label it now ? n                                       ← Enter "n" for no.
:
#
```

Figure 6-7 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
             :
        14. SUN2.1G
        15. HITACHI-OPEN-3-0315                   ← Do not select this disk type.
        16. other                                   (Note 3)
Specify disk type (enter its number):16           ← Enter number for "other" to define.
Enter number of data cylinders:3336               ← Enter value from Table 6-2 (Note 1)
Enter number of alternate cylinders[2]:2          ← Enter value from Table 6-2
Enter number of physical cylinders[3338]:           (press Enter for default)
Enter number of heads:15                          ← Enter value from Table 6-3
Enter number of physical sectors/track[defaults]:   (press Enter for default)
Enter rpm of drive [3600]:10000                   ← Enter value from Table 6-2 (Note 2)
Enter format time[defaults]:                        (press Enter for default)
Enter cylinder skew[defaults]:                      (press Enter for default)
Enter track skew[defaults]:                         (press Enter for default)
Enter track per zone[defaults]:                     (press Enter for default)
Enter alternate tracks[defaults]:                   (press Enter for default)
Enter alternate sectors[defaults]:                  (press Enter for default)
Enter cache control[defaults]:                      (press Enter for default)
Enter prefetch threshold[defaults]:                 (press Enter for default)
Enter minimum prefetch[defaults]:                   (press Enter for default)
Enter maximum prefetch[defaults]:                   (press Enter for default)
Enter disk type name(remember quotes):"HITACHI OPEN-3"   ← Enter name from Table 1-2.
selecting c1t2d0
[disk formatted]
No defined partition tables.
Disk not labeled.  Label it now ? n               ← Enter "n" for no.
format>
```

Figure 6-8 Defining and setting the disk type

Figure notes:

1. The number of cylinders for the 3390-3B is 3346, and the Hitachi RAID storage system returns '3346 cylinder' to the Mode Sense command, and '5822040 blocks' (Maximum LBA 5822039) to the Read capacity command. When 3390-3B is not labeled yet, Solaris displays 3344 data cylinders and 2 alternate cylinders. When 3390-3B is labeled by the Solaris format type subcommand, use 3340 for data cylinder and 2 for alternate cylinder. This is similar to the 3390-3B VLL.
2. The Hitachi RAID storage system reports the RPM of the physical disk drive in response to the type subcommand parameter.
3. It is also possible to follow the procedure using type => "0. Auto configure" => label the drive without calculating detail values like as Cylinder, Header, Blocks/Tracks.

4. Setting host mode 16 affects the geometry parameter reported by the Hitachi RAID storage system (see Table 6-2) as follows:

- Setting host mode option 16 to ON increases the number of cylinders by 4 and reduces the number of blocks per track by ¼.
- Setting host mode option 16 to OFF lowers the number of cylinders by ¼ and increases the number of blocks per track by 4. Therefore, if you use host mode option 16, please account for these differences. For example, if you change the host mode option 16 from OFF to ON, you may want to make either of the following changes in the Format Menu:

  - Increase the number of block setting per track by ¼ and the number of heads by 4.

  - Increase the number of blocks per track to ¼, the number of cylinders by 2, and the number of heads by 2.

  If the number of cylinders entered exceeds 65,533, the total LU block number equals or is less than 65,533. Use the Format Menu to specify the numbers of cylinders, heads, and blocks per track.

```
format> disk

AVAILABLE DISK SELECTIONS
       0. c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0
       1. c0t3d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@3,0
       2. c1t2d0 <HITACHI OPEN-3 cyl 3336 alt 2 hd 15 sec 96> ...already labeled
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,0
       3. c1t2d1 <HITACHI-OPEN-3-52-34 .....>                  ...not yet labeled
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,1
       4. c1t4d0 <HITACHI-OPEN-9-52-34 .....>                 ...not yet labeled
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@4,0
       5. c1t5d0 <HITACHI-3390-3B-52-34 ....>                 ...not yet labeled
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@5,0
       6. c1t6d0 <HITACHI-3390-3A-52-34 ....>                 ...not yet labeled
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@6,0

Specify disk (enter its number): 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
       volume    - set 8-character volume name
       quit
format> type                                          ← Enter type.

AVAILABLE DRIVE TYPES
       0.Auto configure
          :
       13. SUN1.3G
       14. SUN2.1G
       15. HITACHI-OPEN-3-52-34
       16. HITACHI OPEN-3
       17. other
Specify disk type (enter its number):16               ← Enter the number for
:                                                       the desired drive type.
```

Figure 6-9 Setting the partitions (continues on the next page)

```
selecting c0t2d0
[disk formatted]
No defined partition tables.
Disk not labeled.  Label it now ? n                    ← Enter n for no.
format>
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  unassigned    wm        0 -           0          (0/0/0)

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 [0b, 0c, 0.00mb]:3336c            ← Enter size (Table 6-2).
partition> print                                       ← Display partition table.
:                                                                  (see Note, below)
Current partition table (unnamed)
Part        Tag    Flag      Cylinders     Size              Blocks
  0        root    wm        0 -           0          (0/0/0)         0
  1        swap    wm        0 -           0          (0/0/0)         0
  2      backup    wu        0 - 3335      2.29 GB    (3336/0/0)   4803840
  3  unassigned    wu        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      336 - 3335      204 GB     (2970/0/0)   4276800
  7  unassigned    wm        0 -           0          (0/0/0)         0
```

Figure 6-9 Setting the partitions (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.
  * (see Note, below)
partition> quit                                               ← Return to format.
format> disk                                                  ← Display disks.

AVAILABLE DISK SELECTIONS
        0. c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
           /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0
        1. c0t3d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
           /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@3,0
        2. c1t2d0 <HITACHI OPEN-3 cyl 3336 alt 2 hd 15 sec 96>      ← Verify disk label.
                                                      ↖ Track size.
                                                   ↖ Number of heads.
                                                ↖ Number of alternate cylinders.
                                             ↖ Number of data cylinders.
                                          ↖ Disk type name.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,0
        3. c1t2d1 <HITACHI-OPEN-3-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,1
        4. c1t4d0 <HITACHI-OPEN-9-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@4,0
        5. c1t5d0 <HITACHI-3390-3B-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@5,0
        6. c1t6d0 <HITACHI-3390-3A-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@6,0

Specify disk (enter its number): 3                  ← Enter number for next disk to label,
                                                      or press Ctrl-d to quit.
```

Figure 6-10 Labeling the disk and verifying the disk label



Note: The Solaris system displays the following warnings when an FX device (for example, 3390-3A) is labeled. You can ignore these warnings:

```
Warning: error warning VTOC.
Warning: no backup labels. Label failed.
```

Table 6-2 Device geometry parameters

Device Type	# of Data Cylinders	# of Alternate Cylinders	RPM	Partition Size (sample)
OPEN-3	3336	2	10,000	3336c
OPEN-8	9964	2	10,000	9964c
OPEN-9	10014	2	10,000	10014c
OPEN-E	19757	2	10,000	19757c
OPEN-L	19013	2	10,000	19013c
OPEN-3*n	N1*	2	10,000	N4*
OPEN-8*n	N26*	2	10,000	N29*
OPEN-9*n	N5*	2	10,000	N8*
OPEN-E*n	N30*	2	10,000	N33*
OPEN-L*n	N34	2	10,000	N37
OPEN-x VLL	See Table 1-2	2	10,000	See Table 1-2
OPEN-3*n VLL	N22*	2	10,000	N25*
OPEN-8*n VLL	N22*	2	10,000	N25*
OPEN-9*n VLL	N22*	2	10,000	N25*
OPEN-E*n VLL	N22*	2	10,000	N25*
OPEN-V*n VLL	N22*	2	10,000	N25*
3390-3A	3346	2	10,000	3346c
3390-3B	3340	2	10,000	3340c
3390-3C	3346	2	10,000	3346c
FX OPEN-3	3336	2	10,000	3336c
3390-3A VLL	See Table 1-2	2	10,000	See Table 1-2
3390-3B VLL	See Table 1-2	2	10,000	See Table 1-2
3390-3C VLL	See Table 1-2	2	10,000	See Table 1-2
FX OPEN-3 VLL	See Table 1-2	2	10,000	See Table 1-2
Note: For the values indicated by Nxx (for example, N15, N22), see Table 6-3 through Table 6-9.				

Table 6-3 Geometry parameters for OPEN-3*n LUSE devices

n	Data Cylinders-N1 Partition Size-N4	Heads-N2	Blocks/ Track-N3	Usable Blocks (N1+2)*N2*N3	Provided Blocks =3338*15*96*n	Diff.
2	6674	15	96	9613440	9613440	0
3	10012	15	96	14420160	14420160	0
4	13350	15	96	19226880	19226880	0
5	16688	15	96	24033600	24033600	0
6	20026	15	96	28840320	28840320	0
7	23364	15	96	33647040	33647040	0
8	26702	15	96	38453760	38453760	0
9	30040	15	96	43260480	43260480	0
10	16688	30	96	48067200	48067200	0
11	20026	33	80	52873920	52873920	0
12	20026	30	96	57680640	57680640	0
13	20026	39	80	62487360	62487360	0
14	23364	30	96	67294080	67294080	0
15	16688	45	96	72100800	72100800	0
16	26702	30	96	76907520	76907520	0
17	30040	34	80	81714240	81714240	0
18	30040	30	96	86520960	86520960	0
19	30040	38	80	91327680	91327680	0
20	16688	60	96	96134400	96134400	0
21	23364	45	96	100941120	100941120	0
22	30040	55	64	105747840	105747840	0
23	30040	46	80	110554560	110554560	0
24	20026	60	96	115361280	115361280	0
25	16688	45	160	120168000	120168000	0
26	20026	39	160	124974720	124974720	0
27	30040	45	96	129781440	129781440	0
28	23364	60	96	134588160	134588160	0
29	30040	58	80	139394880	139394880	0
30	16688	45	192	144201600	144201600	0
31	30040	62	80	149008320	149008320	0
32	26702	60	96	153815040	153815040	0
33	30040	55	96	158621760	158621760	0
34	30040	64	85	163428480	163428480	0

n	Data Cylinders-N1 Partition Size-N4	Heads-N2	Blocks/ Track-N3	Usable Blocks (N1+2)*N2*N3	Provided Blocks =3338*15*96*n	Diff.
35	30040	56	100	168235200	168235200	0
36	30040	60	96	173041920	173041920	0

Notes:

N1,**N2**,**N3**: Use value in Table 6-2.

N4: Use same value as N1. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 6674c for OPEN-3*2).

Table 6-4 Geometry parameters for OPEN-8*n LUSE devices

n	Data Cylinders-N26 Partition Size-N29	Heads -N27	Blocks/ Track-N28	Usable Blocks (N26+2)*N27*N28	Provided Blocks =9966*15*96*n	Diff.
2	19930	15	96	28702080	28702080	0
3	29896	15	96	43053120	43053120	0
4	29896	20	96	57404160	57404160	0
5	29896	25	96	71755200	71755200	0
6	29896	30	96	86106240	86106240	0
7	29896	35	96	100457280	100457280	0
8	29896	40	96	114808320	114808320	0
9	29896	45	96	129159360	129159360	0
10	29896	50	96	143510400	143510400	0
11	29896	55	96	157861440	157861440	0
12	29896	60	96	172212480	172212480	0
13	29896	52	120	186563520	186563520	0
14	29896	56	120	200914560	200914560	0
15	29896	60	120	215265600	215265600	0
16	29896	64	120	229616640	229616640	0
17	29896	34	240	243967680	243967680	0
18	29896	36	240	258318720	258318720	0
19	29896	38	240	272669760	272669760	0
20	29896	40	240	287020800	287020800	0
21	29896	42	240	301371840	301371840	0
22	29896	44	240	315722880	315722880	0
23	29896	46	240	330073920	330073920	0
24	29896	48	240	344424960	344424960	0
25	29896	50	240	358776000	358776000	0
26	29896	52	240	373127040	373127040	0
27	29896	54	240	387478080	387478080	0
28	29896	56	240	401829120	401829120	0
29	29896	58	240	416180160	416180160	0

n	Data Cylinders-N26 Partition Size-N29	Heads -N27	Blocks/ Track-N28	Usable Blocks (N26+2)*N27*N28	Provided Blocks =9966*15*96*n	Diff.
30	29896	60	240	430531200	430531200	0
31	29896	62	240	444882240	444882240	0
32	29896	64	240	459233280	459233280	0
33	32614	60	242	473584320	473584320	0
34	29896	64	255	487935360	487935360	0
35	30655	64	256	502284288	502286400	2112
36	31531	64	256	516636672	516637440	768

Notes:

N26,**N27**,**N28** : Use values in Table 1-2.

N29: Use same value as N26. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 19930c for OPEN-8*2).

Note: Data cylinders must be less than or equal to **32767**, heads must be less than or equal to **64**, blocks per track must be less than or equal to **256** when these values are specified as parameters of Solaris format type subcommand. The whole data blocks of OPEN-3*2 ~ OPEN-3*36 can be used by above parameters.

Table 6-5 Geometry parameters for OPEN-9*n LUSE devices

n	Data Cylinders-N5 Partition Size-N8	Heads -N6	Blocks/ Track-N7	Usable Blocks (N5+2)*N6*N7	Provided Blocks =10016*15*96*n	Diff.
2	20030	15	96	28846080	28846080	0
3	30046	15	96	43269120	43269120	0
4	30046	20	96	57692160	57692160	0
5	30046	25	96	72115200	72115200	0
6	30046	30	96	86538240	86538240	0
7	30046	35	96	100961280	100961280	0
8	30046	40	96	115384320	115384320	0
9	30046	45	96	129807360	129807360	0
10	30046	50	96	144230400	144230400	0
11	30046	55	96	158653440	158653440	0
12	30046	60	96	173076480	173076480	0
13	30046	52	120	187499520	187499520	0
14	30046	56	120	201922560	201922560	0
15	30046	60	120	216345600	216345600	0
16	30046	64	120	230768640	230768640	0
17	30046	34	240	245191680	245191680	0
18	30046	36	240	259614720	259614720	0

n	Data Cylinders-N5 Partition Size-N8	Heads -N6	Blocks/ Track-N7	Usable Blocks (N5+2)*N6*N7	Provided Blocks =10016*15*96*n	Diff.
19	30046	38	240	274037760	274037760	0
20	30046	40	240	288460800	288460800	0
21	30046	42	240	302883840	302883840	0
22	30046	44	240	317306880	317306880	0
23	30046	46	240	331729920	331729920	0
24	30046	48	240	346152960	346152960	0
25	30046	50	240	360576000	360576000	0
26	30046	52	240	374999040	374999040	0
27	30046	54	240	389422080	389422080	0
28	30046	56	240	403845120	403845120	0
29	30046	58	240	418268160	418268160	0
30	30046	60	240	432691200	432691200	0
31	30046	62	240	447114240	447114240	0
32	30046	64	240	461537280	461537280	0
33	30985	64	240	475960320	475960320	0
34	31924	64	240	490383360	490383360	0
35	31298	63	256	504806400	504806400	0
36	31689	64	256	519225344	519229440	4096

Notes:

N5,**N6**,**N7**: Use value in Table 6-2 and Table 6-3.

N8: Use same value as N5. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 20030c for OPEN-9*2).

Table 6-6 Geometry parameters for OPEN-E*n LUSE devices

n	Data Cylinders-N30 Partition Size-N33	Heads-N31	Blocks/ Track-N32	Usable Blocks (N30+2)*N31*N32	Provided Blocks =9966*15*96*n	Diff.
2	19757	30	96	56905920	56905920	0
3	19757	45	96	85358880	85358880	0
4	19757	60	96	113811840	113811840	0
5	19757	30	240	142264800	142264800	0
6	19757	45	192	170717760	170717760	0
7	19757	60	168	199170720	199170720	0
8	19757	60	192	227623680	227623680	0
9	19757	60	216	256076640	256076640	0
10	19757	60	240	284529600	284529600	0
11	27166	60	192	312975360	312982560	7200
12	29636	60	192	341429760	341435520	5760
13	32106	60	192	369884160	369888480	4320
14	27660	60	240	398332800	398341440	8640
15	29636	60	240	426787200	426794400	7200
16	31612	60	240	455241600	455247360	5760
17	31612	60	255	483694200	483700320	6120
18	31257	64	256	512147456	512153280	5824

Notes:

N30,**N31**,**N32**: Use value in Table 6-2.

N33: Use same value as N30. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 19757c for OPEN-E*2).



Note: Data cylinders must be less than or equal to **32767**, heads must be less than or equal to 64, blocks per track must be less than or equal to 256 when these values are specified as parameters of Solaris format type subcommand. The whole data blocks of OPEN-E*2~OPEN-E*10 can be used by above parameters. About OPEN-E*11~OPEN-E*18, some blocks must become unusable.

Table 6-7 Geometry parameters for OPEN-L*n LUSE devices

n	Data Cylinders-N34 Partition Size-N37	Heads-N35	Blocks/ Track-N36	Usable Blocks (N34+2)*N35*N36	Provided Blocks =49439*15*96*n	Diff.
2	19013	64	117	142384320	142384320	0
3	30422	36	195	213576480	213576480	0
4	30422	45	208	284768640	284768640	0
5	30422	60	195	355960800	355960800	0
6	30422	60	234	427152960	427152960	0
7	30897	63	256	498339072	498345120	6048

Notes:

N34, **N35**, **N36**: Use value in Table 6-2.

N37: Use same value as N34. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 19013c for OPEN-L*2).



Note: Data cylinders must be less than or equal to **32767**, heads must be less than or equal to 64, blocks per track must be less than or equal to 256 when these values are specified as parameters of Solaris format type subcommand. The whole data blocks of OPEN-L*2~OPEN-L*6 can be used by above parameters. About OPEN-L*7, some blocks must become unusable.

Table 6-8 Geometry parameters for OPEN-*x**n VLL-LUSE devices (example)

Data Cylinders-N22 Partition Size-N25	Heads-N23	Blocks/ Track-N24	Usable Blocks (N22+2)*N23*N24	Provided Blocks-N21	Diff.
98	15	96	144000	35MB×2 volumes ↑35×1024/720↑×2=100 100×15×96=144000	0
2590	15	96	3732480	50MB×36 volumes ↑50×1024/720↑×36=2592 2592×15×96=3732480	0
284	15	96	411840	100MB×2 volumes ↑100×1024/720↑×2=286 286×15×96=411840	0
5694	15	96	8202240	500MB×8 volumes ↑500×1024/720↑×8=5696 5696×15×96=8202240	0
22758	30	96	65548800	2000MB×2 volumes ↑2000×1024/720↑×16=45520 45520×15×96=65548800	0
27455	40	188	206476640	2800MB×36 volumes ↑2800×1024/720↑×36=143388 143388×15×96=206478720	2080

Notes:

N21 # of blocks of LUSE composed by VLL volumes are calculated by:

N21 = N20 × (# of heads) × (# of sectors per track).

N22: N20 – 2, Use total cylinder – 2.

N23, N24: Use value in Table 6-2 and Table 6-3.

N25: Use same value as N22.

Table 6-9 Geometry parameters for OPEN-V*n VLL-LUSE devices (example)

Data Cylinders-N22 Partition Size-N25	Heads -N23	Blocks/ Track-N24	Usable Blocks (N22+2)*N23*N24	Provided Blocks-N21	Diff.
48	15	128	92160	45 MB volumes ↑ 45 × 16/15 ↑ = 48 48 × 15 × 128 = 92160	0
27305	30	128	104858880	50 GB volumes ↑ 50 × 1024 × 16/15 ↑ = 54614 54614 × 15 × 128 = 104858880	0
10921	150	128	209721600	10 GB × 10 volumes ↑ 10 × 1024 × 16/15 ↑ × 10 = 109230 109230 × 15 × 128 = 209721600	0
32767	100	128	419443200	20 GB × 10 volumes ↑20 × 1024 × 16/15↑ × 10 = 218460 218460 × 15 × 128 = 419443200	0

Notes:

N21 # of blocks of LUSE composed by VLL volumes are calculated by:

N21 = N20 × (# of heads) × (# of sectors per track).

N22: N20 – 2, Use total cylinder – 2.

N23, N24: Use value in Table 6-2 and Table 6-3.

N25: Use same value as N22.

- For OPEN-V, because the capacity is not fixed, AutoConfig is recommended for the geometry setting.
- The geometry setting is also available manually as needed. In this case, the OPEN-V geometry of X GB can be calculated according to the equations in ***Example 1*** and ***Example 2*** using the values of N22, N23, and N24:

Example 1:

N22(Cyl)×N23(Head)×N24(Block/Trk)×512(Byte) ≤ X GB (=×1024×1024×1024 Byte) is as follows:

```
16000(Cyl)×256(Head)×256(Block)×512(Byte)=536870912000Byte=500GB
32000(Cyl)×128(Head)×256(Block)×512(Byte)=536870912000Byte=500GB
```

Example 2 (a variation of ***Example 1***):

22(Cyl)×N23(Head)×N24(Block/Trk) ≤ X GB (=×1024×1024×1024 Byte) / 512 (Byte) = Usable Blocks is as follows:

```
15000(Cyl)×256(Head)×256(Block)×512(Byte)=536870912000Byte=468.75GB<500GB
```

Creating and mounting the file systems

After you partition and label all new disks, you can create and mount the file systems for the SCSI disk devices.

- ☐ Creating the file systems
- ☐ Creating and verifying the mount directories
- ☐ Mounting and verifying the file systems
- ☐ Setting and verifying the auto-mount parameters



Note: Do not create file systems or mount directories for the FX devices (for example, 3390-3A). These devices are accessed as raw devices and do not require any further configuration after being partitioned and labeled.

Creating the file systems

To create the file systems for the newly installed SCSI disk devices:

1. Create the file system using the **`newfs -C <maxcontig>`** command (see Figure 6-11).
   a. Use 6 or one of the following multiples of 6 as the **maxcontig** value for all SCSI disk devices on the Hitachi RAID storage system: 12, 18, 24, or 30. If 6 is used, the Solaris system will access 48 KB as a unit (6 × 8 KB), which matches the track size of the OPEN-x devices. These **maxcontig** values (6, 12, 18, 24, 30) optimize the I/O performance by keeping the I/O data range on one track. The **maxcontig** value that you choose depends on your applications, and you can always change the **maxcontig** parameter to a different value at any time.
   b. Use the character-type device file as the argument. For example:

      /dev/rdsk/c1t2d0s0
2. When the confirmation appears, verify that the device file name is correct. If so, enter **`y`** for yes. If not, enter **`n`** for no, and then repeat step (1) using the correct device file name.
3. Repeat steps (1) and (2) for each new SCSI disk device on the storage system. Be sure to use the same **maxcontig** value for all Hitachi RAID storage system devices.

```
# newfs -C 6 /dev/rdsk/c1t2d0s0                          ← Create file system.
newfs:construct a new file system /dev/rdsk/c1t2d0s0:(y/n) y   ← Verify correct device.
/dev/rdsk/c1t2d0s0: 4803840 sectors in 3336 cylinders of 15 tracks, 96 sectors
    2345.6MB in 209 cyl groups (16 c/g, 11.25MB/g, 5440 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
 32, 23168, 46304, 69440, 92576, 115712, 138848, 161984, 185120, 208256,
  :
4747616, 4770752, 4792352,
# newfs -C 6 /dev/rdsk/c1t2d1s0                          ← Create file system on next disk
                                                           using the same maxcontig value.
```

Figure 6-11 Creating the file systems

Creating and verifying the mount directories

After you create the file systems, create and verify the mount directories for the new SCSI disk devices. Each logical partition requires a unique mount directory, and the mount directory name should identify the logical volume and the partition.

To create the mount directories for the newly installed SCSI disk devices:

1. Go to the root directory (see Figure 6-12).
2. Use the **mkdir** command to create the mount directory.

   To delete a mount directory, use the **rmdir** command (for example, **rmdir /VSP_LU00**).
3. Choose a name for the mount directory that identifies both the logical volume and the partition. For example, to create a mount directory named VSP_LU00, enter:

   mkdir /VSP_LU00
4. Use the **ls -x** command to verify the new mount directory.
5. Repeat steps 2 and 3 for each logical partition on each new SCSI disk device.

```
# cd                                                    ← Go to the root directory.
# pwd                                                   ← Display current directory.
/
# mkdir /VSP_LU00                                       ← Create new mount directory.
# ls -x                                                 ← Verify new mount directory.
VSP_LU00  bin    dev      device    etc      export correctly
floppy    home   hstsboof kadb      kernel   lib
#
```

Figure 6-12 Creating and verifying a mount directory

Mounting and verifying the file systems

After you create the mount directories, mount and verify the file systems for the new SCSI disk devices. The file system for each logical partition should be mounted and verified to ensure that all new logical units are fully operational.

To mount and verify the file systems for the new devices (see Figure 6-13):

1. Mount the file system using the mount command. Be sure to use the correct block-type device file name and mount directory for the device/partition. For example, to mount the file **/dev/dsk/c1t2d0s0** with the mount directory **/VSP_LU00**, enter:

   ```
   mount /dev/dsk/c1t2d0s0 /VSP_LU00
   ```

   To unmount a file system, use the `umount` command (for example, `umount /VSP_LU00`).



Note: If you already set the auto-mount parameters (see Setting and verifying the auto-mount parameters), you do not need to specify the block-type device file, only the mount directory.

2. Repeat step 1 for each partition of each newly installed SCSI disk device.
3. Display the mounted devices using the `df -k` command, and verify that all new SCSI disk devices are displayed correctly. OPEN-x devices will display as OPEN-3, OPEN-9, OPEN-E, OPEN-L devices.
4. As a final verification, perform some basic UNIX operations (for example, file creation, copying, and deletion) on each logical unit to ensure the new file systems are fully operational.

```
# mount /dev/dsk/c1t2d0s0 /VSP_LU00                          ← Mount file system.
                ↖ Block-type device file name
# mount /dev/dsk/c1t2d1s0 /VSP_LU01                          ← Mount next file system.
                            ↖ Mount directory name
# mount /dev/dsk/c1t2d2s0 /VSP_LU02                          ← Mount next file system.
 :
 :
#
# df -k                                                      ← Display file systems.
File system         Kbytes    used    avail  capacity   Mounted on
/dev/dsk/c0t3d0s0    28775   27706        0    100%     /
/dev/dsk/c0t3d0s6   269191  234897     7384     97%     /usr
/proc                    0       0        0      0%     /proc
fd                       0       0        0      0%     /dev/fd
/dev/dsk/c0t3d0s4s   57567   29515    22302     57%     /var
swap                142204      20   142184      0%     /tmp
/dev/dsk/c0t3d0s7   462119  206000   209909     50%     /export/home
/dev/dsk/c0t3d0s5    47975   42059     1126     97%     /opt
/dev/dsk/c1t2d0s0  2256436       9  2030787      0%     /VSP_LU00
                                                         ↖ Verify file systems.
/dev/dsk/c1t2d1s0  2256436       9  2030787      0%     /VSP_LU01
                                                         ↖ OPEN-3 device.
/dev/dsk/c1t2d2s0  6774358       9  6548709      0%     /VSP_LU02
                                                         ↖ OPEN-9 device.
:
# mount /dev/dsk/c1t2d0s0 /VSP_LU00                          ← Mount file system.
# cd /VSP_LU00                                               ← Go to mount directory.
# cp /bin/vi /VSP_LU00/vi.back1                              ← Copy a file.
# ls -l                                                      ← Verify the file copy.
drwxr-xr-t   2 root     root       8192 Mar 15 11:35  lost+found
-rwxr-xr-x   1 root     sys     2617344 Mar 15 11:41  vi.back1
# cp vi.back1 vi.back2                                       ← Copy file again.
# ls -l                                                      ← Verify file copy again.
drwxr-xr-t   2 root     root       8192 Mar 15 11:35  lost+found
-rwxr-xr-x   1 root     sys     2617344 Mar 15 11:41  vi.back1
-rwxr-xr-t   1 root     sys     2617344 Mar 15 11:52  vi.back2
# rm vi.back1                                                ← Remove test files.
# rm vi.back2                                                ← Remove test files.
```

Figure 6-13 Mounting and verifying the file system

Setting and verifying the auto-mount parameters

You can add any or all of the new SCSI disk devices to the **/etc/vfstab** file to specify the auto-mount parameters for each device. Once a device is added to this file, you can mount the device without having to specify its block-type device file name (for example, **mount /VSP_LU00**), since the **/etc/vfstab** file associates the device with its mount directory.

To set the auto-mount parameters for the desired devices (see Figure 6-14):

1. Make a backup copy of the **/etc/vfstab** file:

   cp /etc/vfstab /etc/vfstab.standard

2. Edit the **/etc/**vfstab file to add one line for each device to be auto-mounted. Table 6-10 shows the auto-mount parameters. If you make a mistake while editing, exit the **vi** editor without saving the file, and then begin editing again.
3. Reboot the Solaris system after you are finished editing the **/etc/vfstab** file.
4. Use the **df -k** command to display the mounted devices and verify that the desired devices were auto-mounted.

```
# cp -ip /etc/vfstab /etc/vfstab.standard                    ← Make backup before editing.
# vi /etc/vfstab                                             ← Edit the file.
#device            device              mount       FS      fsck  mount    mount
#to mount          to fsck             point       type    pass  at boot  options
  ①                ②                   ③           ④       ⑤     ⑥        ⑦←See Table 6-10


/proc              -                   /proc       procfs  -     no       -
fd                 -                   /dev/fd     fd      -     no       -
swap               -                   /tmp        tmpfs   -     yes      -
/dev/dsk/c0t3d0s0  /dev/rdsk/c0t3d0s0  /           ufs     1     no       -
/dev/dsk/c0t3d0s6  /dev/rdsk/c0t3d0s6  /usr        ufs     2     no       -
/dev/dsk/c0t3d0s7  /dev/rdsk/c0t3d0s7  /export     ufs     3     yes      -
/dev/dsk/c0t3d0s5  /dev/rdsk/c0t3d0s5  /opt        ufs     4     yes      -
/dev/dsk/c0t3d0s1  -                   -           swapfs  -     no       -
/dev/dsk/c1t2d0s0  /dev/rdsk/c1t2d0s0  /VSP_LU00   ufs     5     yes      -   ← Add one line
/dev/dsk/c1t2d1s0  /dev/rdsk/c1t2d1s0  /VSP_LU01   ufs     5     yes      -     for each LUN.
```

Figure 6-14 Setting the auto-mount parameters

Table 6-10 Auto-mount parameters

Parameter #	Name	Enter:
①	Device to mount	Block-type device file name
②	Device to fsck	Character-type device file name
③	Mount point	Mount directory name
④	FS type	File system type (for example, ufs)
⑤	Fsck pass	Order for performing file system checks
⑥	Mount at boot	Yes = auto-mounted at boot/mountall No = not auto-mounted at boot/mountall
⑦	Mount options	Desired mount options: - no options (typical) -ro read-only access (for example, for 3390-3B devices)

ALUA support on Solaris

ALUA is supported on Solaris systems with the following two restrictions:

- You must modify the `/etc/driver/drv/scsi_vhci.conf` file (shown below) and then reboot the system.
- You must have ALUA enabled for all LUNs.



Note:

- LUNs that do not have ALUA enabled are not visible after you enable ALUA support. If you plan to migrate non-ALUA LUNs, perform the migration first and then enable ALUA after the migration.
- This information applies to Solaris versions 10.*x* and 11.*x*.

Define the following in `/etc/driver/drv/scsi_vhci.conf`:

```
scsi-vhci-failover-override =
"HITACHI OPEN", "f_tpgs";
and "scsi_vhci_f_sym_hds" and "scsi_vhci_f_tpgs" defined in ddi-forceload.
ddi-forceload =
"misc/scsi_vhci/scsi_vhci_f_asym_sun",
"misc/scsi_vhci/scsi_vhci_f_asym_lsi",
"misc/scsi_vhci/scsi_vhci_f_asym_emc",
"misc/scsi_vhci/scsi_vhci_f_sym_emc",
"misc/scsi_vhci/scsi_vhci_f_sym_hds",
"misc/scsi_vhci/scsi_vhci_f_sym",
# "misc/scsi_vhci/scsi_vhci_f_sym_enc",
# "misc/scsi_vhci/scsi_vhci_f_tpgs_tape",
# "misc/scsi_vhci/scsi_vhci_f_tape",
"misc/scsi_vhci/scsi_vhci_f_tpgs";
```

Troubleshooting for Solaris host attachment

Table 6-11 lists potential error conditions that might occur during storage system installation on a Solaris host and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 6-11 Troubleshooting for Solaris host attachment

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Ensure the READY indicator lights on the storage system are ON. Ensure the fibre-channel cables are correctly installed and firmly connected. Run **dmesg** to recheck the fibre buses for new devices. Verify the contents of **/kernel/drv/sd.conf** file.
File system cannot be created (**newfs** command)	Ensure the character-type device file is specified for **newfs** command. Verify that logical unit is correctly labeled by UNIX **format** command.
The file system is not mounted after rebooting.	Ensure the system was restarted properly. Ensure the file system attributes are correct. Ensure the **/etc/vfstab** file is correctly edited.
The Solaris system does not reboot properly after hard shutdown.	If the Solaris system is powered off without executing the shutdown process, wait three minutes before restarting the Solaris system. This allows the storage system's internal time-out process to purge all queued commands so that the storage system is available (not busy) during system startup. If the Solaris system is restarted too soon, the storage system will continue trying to process the queued commands, and the Solaris system will not reboot successfully.
The Hitachi RAID storage system performed a self-reboot because the system was busy or it logged a panic message.	Reboot the Solaris system.
The Hitachi RAID storage system responds Not Ready, or displays Not Ready and timed itself out.	Contact Hitachi Data Systems technical support.
The system detects a parity error.	Ensure the HBA is installed properly. Reboot the Solaris system.

Verbose mode troubleshooting

One way to troubleshoot Solaris operations involves the "verbose" mode for the HBA configuration file. This section provides examples of error messages that may occur. A possible debugging method is to select the device and turn on verbose mode, then attempt the boot process again. Verbose error messages provide information that help isolate the problem.

To turn on the verbose flag, use the commands shown in Figure 6-15. Figure 6-16 shows examples of error messages.

```
ok  " /sbus/fca" select-dev
ok   true to fca-verbose
ok   boot fcadisk
```

Figure 6-15 Turning on the verbose flag

```
Error message:
Cannot Assemble drivers for /sbus@1f,0/fcaw@1,0/sd@0,0:a
Cannot Mount root on  /sbus@1f,0/fcaw@1,0/sd@0,0:a
Problem:
The process of copying the OS to the fibre channels was not complete, or the drive
specified on the boot command is not the same as the one the OS was constructed on.
```

```
Error message:
Can't open boot device
Problem:
The wwn specified with the set-bootn0-wwn does not correspond to the wwn of the device.
Could also be a cable problem - the adapter cannot initialize.
```

```
Error message:
The file just loaded does not appear to be bootable
Problem:
The bootblk was not installed on the target.
```

```
Error message:
mount: /dev/dsk/c0t0d0s0 - not of this fs type
Problem:
At this point the process hangs. This happens if the /etc/vfstab
File has not been updated on the fibrechannel boot drive to reflect the new target.
```

```
Error message:
Get PortID request rejected by nameserver
Problem:
The wwn of the target is not correct. Select the adapter and perform set-bootn0-wwn. If
this is correct, check the switch to see that target is properly connected.
```

```
Error message:
Can't read disk label
Problem:
The selected target is not a Solaris filesystem.
Error message:
Nport init failed -
Problem:
Card is connected to an arbitrated loop device, but wants to initialize as an NPORT. The
bootn0-wwn property has probably been set to a valid WWN.
```

```
Error message:
Panic dump not saved
Problem:
After the system is successfully booted to Solaris from the fibrechannel and a panic occurs
the panic does not get saved to the swap device.
This can be the result not properly defined the swap partition.
Use the format command to view the slices on the fibre channel drive.
Take the partition option, then the print option.
The swap partition should look something like this:
      1      swap      wm      68-459      298.36MB      (402/0/0)      611040
Sizes and cylinders will probably be different on your system. Make sure that the flag is
wm and that the sizes are defined (not 0). Then use the label option from partition to
write the label to the drive. After this the panic should be saved to the swap partition.
If the partition needs to be changed chose the partition option, and enter 1 to select
slice 1.
```

Figure 6-16 Examples of error messages

Online device installation

After initial installation and configuration of the Hitachi RAID storage system, additional devices can be installed or de-installed online without having to restart the Solaris system. After online installation, the device parameters for new volumes must be changed to match the LUs defined under the same fibre-channel port (see Verifying recognition of new devices). This procedure should be performed by the system administrator (that is, super-user).



Note: For additional instructions about online installation and deinstallation of LUs, see the Maintenance Manual.

Sun fibre-channel host bus adapter installation

To perform online installation of the Sun fibre-channel HBA:

1. Set up the Solaris server:
   - Confirm that the Sun fibre-channel HBAs are installed.
   - Confirm that Sun StorEdge SAN Foundation Software version 4.2 or later is installed.
2. Set up the Hitachi RAID storage system:
   - Ensure the latest microcode is loaded. Non-disruptive version-up requires alternate path.
   - Install the front-end directors and LDEVs, and connect fibre cable if necessary.
   - Execute online LU installation from the service processor (SVP) or the Storage Navigator software.
   - Verify the SCSI path configuration.
3. Execute the **`Format`** command. Solaris will recognize the new volumes.
4. If new volumes are not recognized, the following operation is not needed. Refer to the Solaris documentation as needed.
   - Disconnect and reconnect the fibre cable connected to the paths on which you are adding LUs.
   - Use the following command to display available paths to the HBAs:

     `luxadm –e port`
   - With the path from the output, issue the following command:

     `luxadm –e forcelip path`
   - Use the following command to display devices:

     `cfgadm –al`
   - Bring fabric devices back onto the system.
   - Execute the **`Format`** command.

Using MPxIO path failover software

The Hitachi RAID storage systems are compatible with the Solaris Operating Environment Multi-path I/O (MPxIO) multi-pathing driver that offers hardware transparency and multi-pathing capabilities. MPxIO is fully integrated within the Solaris operating system (beginning with Solaris 8) and enables I/O devices to be accessed through multiple host controller interfaces from a single instance of the I/O device.

MPxIO enables you to more effectively to represent and manage devices that are accessible through multiple I/O controller interfaces within a single instance of the Solaris operating system. The MPxIO architecture:

- Helps protect against I/O outages due to I/O controller failures. Should one I/O controller fail, MPxIO automatically switches to an alternate controller.
- Increases I/O performance by load balancing across multiple I/O channels.

For the Hitachi RAID storage system to work with MPxIO:

1. Configure the Hitachi RAID storage system to use host mode **09** (see Setting the host modes and host mode options).
2. Modify the configuration file **/kernel/drv/scsi_vhci.conf** to enable MPxIO to manage the path failover:

   mpxio-disable="no";



Note: You do not have to edit **/kernel/drv/sd.conf**.

3. Connect the Hitachi RAID storage system to the Solaris system.
4. Reboot the server.
5. After reboot, login to the system and issue the following command:

   cfgadm -la

   The following information appears:

```
bigc2 > cfgadm -la
Ap_Id                  Type         Receptacle   Occupant     Condition
ac0:bank0              memory       connected    configured   ok
ac0:bank1              memory       connected    configured   ok
ac1:bank0              memory       connected    configured   ok
ac1:bank1              memory       connected    configured   ok
ac2:bank0              memory       connected    configured   ok
ac2:bank1              memory       connected    configured   ok
ac3:bank0              memory       connected    configured   ok
ac3:bank1              memory       connected    configured   ok
c0                     scsi-bus     connected    configured   unknown
c0::dsk/c0t2d0         disk         connected    configured   unknown
c0::dsk/c0t3d0         disk         connected    configured   unknown
c0::dsk/c0t6d0         CD-ROM       connected    configured   unknown
c0::rmt/0              tape         connected    configured   unknown
c5                     fc-fabric    connected    configured   unknown
c5::20000001730037eb   unavailable  connected    unconfigured failed
c5::200000017380a45b   unknown      connected    unconfigured unknown
```

```
c5::210000e08b042791   unknown      connected    unconfigured unknown
c5::210000e08b049755   unknown      connected    unconfigured unknown
c5::210100e08b276f6d   unknown      connected    unconfigured unknown
c5::500060e8029eb604   disk         connected    configured   unknown
c5::50060e80034e5a05   disk         connected    configured   unknown
c5::50060e8004272f01   disk         connected    configured   unknown
c6                     fc-fabric    connected    configured   unknown
c6::200000017300380d   unavailable  connected    unconfigured failed
c6::200000017300a45b   unknown      connected    unconfigured unknown
c6::210000e08b076f6d   unknown      connected    unconfigured unknown
c6::210100e08b242791   unknown      connected    unconfigured unknown
c6::500060e8029eb614   disk         connected    unconfigured unknown
c6::50060e80034e5a15   disk         connected    unconfigured unknown
c6::50060e8004272f11   disk         connected    configured   unknown
```

6. Check for the target not configured (in red). Then issue the following command to see the unconfigured LUNs:

```
cfgadm -c configure c6::500060e8029eb614   c6::50060e80034e5a15
```

7

SUSE Linux configuration and attachment

This chapter describes how to configure the new Hitachi disk devices on a SUSE Linux host:

- ☐ Hitachi storage system configuration for SUSE Linux operations
- ☐ Device Mapper (DM) Multipath configuration
- ☐ Verifying new device recognition
- ☐ Configuring the new devices
- ☐ Troubleshooting for SUSE Linux host attachment



Note: Configuration of the devices should be performed by the Linux system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Hitachi storage system configuration for SUSE Linux operations

The storage system must be fully configured before being attached to the SUSE Linux host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for SUSE Linux operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)

Host mode. The required host mode for SUSE Linux is **00**. Do not select a host mode other than **00** for IBM AIX. For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Device Mapper (DM) Multipath configuration

The Hitachi RAID storage systems support DM Multipath operations.



Note: Refer to the native multipath configuration documentation from the OS vendor for important information about required settings and parameters for DM Multipath operations, including but not limited to:

- Disabling the HBA failover function
- Installing `kpartx` utility
- Creating the multipath device with the `multipath` command
- Editing the `/etc/modprobe.conf` file
- Editing the `/etc/multipath.conf` file
- Configuring LVM
- Configuring raw devices
- Creating partitions with DM Multipath

Verifying new device recognition

The final step before configuring the new 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. Display the devices using the **dmesg** command (see example in Figure 7-1). In this example, the HITACHI OPEN-V device (TID 0, LUN 0) and the HITACHI OPEN-V device (TID 0, LUN 1) are recognized by the SUSE Linux server.
2. Record the device file name for each new device. You will need this information when you partition the devices (see Partitioning the devices). See Table 7-1 for a sample SCSI path worksheet.
3. The device files are created under the **/dev** directory. Verify that a device file was created for each new disk device (see Figure 7-2).

```
# dmesg | more
:
:
scsi0 : Qlogic QLA2200 PCI to Fibre Channel Host Adapter: 0 device 14 irq 11
       Firmware version:  1.17.26, Driver version  2.11 Beta

scsi : 1 host.
  Vendor: HITACHI        Model: OPEN-V            Rev: 0111
  Type:   Direct-Access                      ANSI SCSI revision: 02
Detected scsi disk sda at scsi0, channel 0, id 0, lun 0
          ↖ Device file name of this disk = /dev/sda   ↖ Logical unit number

  Vendor: HITACHI        Model: OPEN-V            Rev: 0111
  Type:   Direct-Access                      ANSI SCSI revision: 02
Detected scsi disk sdb at scsi0, channel 0, id 0, lun 1
:
:
```

Figure 7-1 Verifying new device recognition

```
# ls -l /dev | more
:
brw-rw----   1 root    disk    8,  0 May  6  1998 sda      ← Device file = sda
```

Figure 7-2 Verifying device files

Table 7-1 Sample SCSI path worksheet

LDEV (CU:LDEV)	Device Type	LUSE (*n)	VLL (MB)	Device File Name	Path	Alternate Path
0:00					TID:____ LUN:____	TID:____ LUN:____
0:01					TID:____ LUN:____	TID:____ LUN:____
0:02					TID:____ LUN:____	TID:____ LUN:____
0:03					TID:____ LUN:____	TID:____ LUN:____
0:04					TID:____ LUN:____	TID:____ LUN:____
0:05					TID:____ LUN:____	TID:____ LUN:____
0:06					TID:____ LUN:____	TID:____ LUN:____
0:07					TID:____ LUN:____	TID:____ LUN:____
0:08					TID:____ LUN:____	TID:____ LUN:____
0:09					TID:____ LUN:____	TID:____ LUN:____
0:0A					TID:____ LUN:____	TID:____ LUN:____
0:0B					TID:____ LUN:____	TID:____ LUN:____
0:0C					TID:____ LUN:____	TID:____ LUN:____
0:0D					TID:____ LUN:____	TID:____ LUN:____
0:0E					TID:____ LUN:____	TID:____ LUN:____
0:0F					TID:____ LUN:____	TID:____ LUN:____

Configuring the new devices

This section describes how to configure the newly attached disk devices:

- ☐ Setting the number of logical units
- ☐ Partitioning the devices
- ☐ Creating, mounting, and verifying file systems

Setting the number of logical units

To set the number of LUs:

1. Edit the `/etc/modules.conf` file to add the following line, for example:

   `options scsi_mod max_scsi_luns=xx`

   xx = the maximum number of LUs supported by your Linux operating system. Check your adapter documentation and your Linux documentation to determine the total number of devices that can be supported.

2. To set the Emulex driver, add the following line to the `/etc/modules.conf` file:

   `Alias scsi_hostadapter lpfcdd`

3. To activate the above modification, make an image file for booting, for example:

   `# mkinitrd /boot/initrd-2.4.x.scsiluns.img 'uname –r'`

4. To change the setting of Bootloader, use one of the following methods (see Figure 7-3 and Figure 7-4):

   a. **LILO used as Bootloader.** You need to edit the `lilo.conf` file and then execute the **lilo** command to activate the `lilo.conf` setting with selecting the label. For example: `# lilo`

   b. **GRUB (Grand Unified Bootloader) is used as Bootloader.** You need to edit the `/boot/grub/grub.conf` file.

5. Reboot the system.

```
image=/boot/vmlinuz-qla2x00
       label=Linux-qla2x00
       append="max_scsi_luns=16"
       initrd=/boot/initrd-2.4.x.img
       root=/dev/sda7
 read-only
#sbin/lilo
```

Figure 7-3 Setting the number of LUs (LILO)

```
Initrd_modules = "lpfcdd"                    ← Add "lpfcdd" in  /etc/rc.config
```

Figure 7-4 Setting the Emulex driver module to load with Ramdisk

Partitioning the devices

After the setting the number of logical units, you can set the partitions.



Note: For important information about creating partitions with DM Multipath, refer to the native multipath configuration documentation from the OS vendor.

To partition the new disk devices:

1. Enter `fdisk/dev/<device_name>` (for 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 as an alternative, you can make one extended partition. The extended partition can be divided into a maximum of 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.

   Other commands that you might want to use include:

   - To remove partitions, select **d**.
   - To stop a change, select **q**.
5. Repeat the above steps for each new disk device.

Creating, mounting, and verifying file systems

Creating file systems

After you have partitioned the devices, you can create the file systems, making sure that they are appropriate for the primary and/or extended partition for each logical unit.

To create the file system, execute the **mkfs** command:

`# mkfs /dev/sda1` (where **/dev/sda1** is device file of primary partition number 1.)

Creating mount directories

To create the mount directories, execute the **mkdir** command:

```
# mkdir /VSP-LU00
```

Mounting new file systems

Use the **mount** command to mount each new file system (see example in Figure 7-5). The first parameter of the **mount** command is the device file name (**/dev/sda1**), and the second parameter is the mount directory.

```
# mount /dev/sda1 /VSP-LU00
        ↖ Device file name  ↖ Mount directory name
#
```

Figure 7-5 Mounting new devices

Verifying file systems

After mounting the file systems, you should verify the file systems (see example in Figure 7-6).

```
# df -h
Filesystem            Size      Used     Avail    Used%    Mounted on
/dev/sda1             1.8G      890M      866M     51%      /
/dev/sdb1             1.9G      1.0G      803M     57%      /usr
/dev/sdc1             2.2G       13k      2.1G      0%      /VSP-LU00
#
```

Figure 7-6 Verifying the file systems

Setting auto-mount parameters

To set the auto-mount parameters, edit the **/etc/fstab** file (see example in Figure 7-7).

```
# cp -ip /etc/fstab /etc/fstab.standard                    ← Make a backup of /etc/fstab.
# vi /etc/fstab                                                  ← Edit /etc/fstab.
:
/dev/sda1          /VSP-LU00 ext2     defaults    0   2          ← Add new device.
```

Figure 7-7 Setting the auto-mount parameters

Troubleshooting for SUSE Linux host attachment

Table 7-2 lists potential error conditions that may occur during installation of new storage and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 7-2 Troubleshooting for SUSE Linux host attachment

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Be sure that the READY indicator lights on the Hitachi RAID storage system are ON. Be sure that the LUNs are properly configured. The LUNs for each target ID must start at 0 and continue sequentially without skipping any numbers.
The file system cannot be created.	Be sure that the device name is entered correctly with **mkfs**. Be sure that the LU is properly connected and partitioned.
The file system is not mounted after rebooting.	Be sure that the system was restarted properly. Be sure that the auto-mount information in the **/etx/fstab** file is correct.

8

VMware configuration and attachment

This chapter describes how to configure the new Hitachi disk devices on a VMware host:

- ☐ Hitachi storage system configuration for VMware operations
- ☐ VMware host configuration for Hitachi RAID storage
- ☐ FCA configuration for VMware
- ☐ Configuring the new devices
- ☐ Troubleshooting for VMware host attachment



Note: Configuration of the devices should be performed by the VMware system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Hitachi storage system configuration for VMware operations

The storage system must be fully configured before being attached to the VMware host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for VMware operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)

Host mode. Table 8-1 lists and describes the required host modes for VMware host attachment. You must use either host mode **01** or host mode **21**. For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Note: For VMware, host groups are created per VMware cluster or per ESX host on the ports on each storage cluster that the VMware cluster or ESX hosts can access.

Table 8-1 Host modes for VMware operations

Host Mode	Description
01[VMware]	If you use host mode 01[VMware], you will not be able to create a LUSE volume using a volume to which an LU path has already been defined. Before performing a LUSE operation on a volume with a path defined from a VMware host, make sure that the host mode is 21[VMware Extension].
21[VMware Extension]	Use host mode 21 if you plan to create LUSE volumes.

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

VMware host configuration for Hitachi RAID storage

This section provides reference information to help you implement VMware software with the Hitachi RAID storage systems:

- ☐ SAN configuration
- ☐ VMware vSphere API operations
- ☐ VMware ESX Server and VirtualCenter compatibility
- ☐ Installing and configuring VMware
- ☐ Creating and managing VMware infrastructure components

SAN configuration

For vSphere 5.5 and earlier, a SAN is required to connect the Hitachi RAID storage system to the VMware ESX Server host as VMware does not support FC-AL and direct-connect connections to storage systems.

For vSphere 6.x, direct-connect connection to the storage system is supported with 8G FC connections and FC-AL. For information about setting up storage arrays for VMware ESX Server, see the VMware user documentation.

For details about supported switches, topology, and firmware versions for SAN configurations, see the Hitachi Data Systems interoperability site: https://support.hds.com/en_us/interoperability.html

VMware vSphere API operations

The Hitachi RAID storage systems support the VMware vSphere API for Array Integration (VAAI). VAAI enables the offload of specific storage operations from the VMware ESX host to the Hitachi RAID storage system for improved performance and efficiency. These APIs, available in VMware vSphere 4.1 and later, provide integration with the advanced features and capabilities of the Hitachi RAID storage systems such as thin provisioning, dynamic tiering, and storage virtualization. For details, see the following sites:

- www.hds.com/go/vmware/
- http://www.vmware.com/products/vsphere/features/storage-api.html

VMware ESX 4.1 or later is required for VAAI operations.

VMware ESX Server and VirtualCenter compatibility

VMware recommends that you install VirtualCenter with the ESX Server software. VirtualCenter lets you provision virtual machines and monitor performance of physical servers and virtual machines, monitor performance and utilization of physical servers and the virtual machines they are running, and export VirtualCenter data to HTML and Excel formats for integration with other reporting tools.

Make sure that your VMware ESX server and VirtualCenter versions are compatible. For details, refer to your VMware Release Notes and the VMware website at www.vmware.com.

Installing and configuring VMware

You must verify that your server, I/O, storage, guest operating system, management agent, and backup software are all compatible before you install and configure VMware.

Consult the following documents for information about VMware ESX Server installation, configuration, and compatibility:

- **Installing and Configuring VMware ESX Server:** Refer to the VMware documentation when installing and configuring VMware ESX Server. Follow the configuration steps for licensing, networking, and security.
- **Upgrading an ESX Server and VirtualCenter Environment:** Refer to the VMware documentation when upgrading an ESX Server and VirtualCenter environment.

Creating and managing VMware infrastructure components

After VMware ESX Server installation has been completed, including all major components of the VMware Infrastructure, you can perform the following tasks to manage your VMware infrastructure components:

- **Use the VI client** to manage your ESX Server hosts either as a group through VirtualCenter or individually by connecting directly to the host.
- **Set up a datacenter** to bring one or more ESX Server hosts under VirtualCenter management, create virtual machines, and determine how you want to organize virtual machines and manage resources.
- **Create a Virtual Machine** manually, from templates, or by cloning existing virtual machines.
- **Configure permissions and roles for users** to allocate access to VirtualCenter, its administrative functions, and its resources.
- **Use resource pools** to partition available CPU and memory resources hierarchically.
- **Configure network connections** to ensure that virtual machine traffic does not share a network adapter with the service console for security purposes.
- **Install a guest operating system** in a virtual machine.
- **Manage virtual machines** to learn how to power them on and off.
- **Monitor the status of your virtual infrastructure** using tasks and events.
- **Schedule automated tasks** to perform actions at designated times.

- **Configure alarm notification messages to be sent** when selected events occur to or on hosts or virtual machines.

FCA configuration for VMware

The fibre-channel adapters (FCAs) on the VMware host must be fully configured before being attached to the Hitachi RAID storage system, as described in Installing and configuring the host adapters. This section provides recommended settings for QLogic and Emulex host adapters for Hitachi RAID storage attached to a VMware host.

- Settings for QLogic adapters
- Settings for Emulex adapters

Settings for QLogic adapters

Table 8-2 lists the recommended QLogic adapter settings for Hitachi RAID storage attached to a VMware host. Use the setup utility for the adapter to set the required options for your operational environment. For details and instructions, see the user documentation for the adapter.

For the latest information about QLogic adapters and Hitachi RAID storage systems, see the QLogic interoperability matrix for Hitachi Data Systems storage: http://www.qlogic.com/Interoperability/SANInteroperability/Pages/home.aspx?vendor=HitachiDataSystems

Table 8-2 Settings for QLogic adapters on VMware hosts

Parameter	Setting
Host Adapter BIOS	Disabled
Number of LUNs per target	Determined by the number of LUNs in your configuration. Multiple LUN support is typically for RAID arrays that use LUNs to map drives. The default is 8. If you do not need multiple LUN support, set the number of LUNs to 0.
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes
Connection option	Point-to-point only

Settings for Emulex adapters

Table 8-3 lists the recommended Emulex adapter settings for Hitachi RAID storage attached to a VMware host. Use the setup utility for the adapter to set the required options for your operational environment. For details and instructions, see the user documentation for the adapter.

For the latest information about Emulex adapters and Hitachi RAID storage systems, see Emulex interoperability matrix for Hitachi Data Systems storage: http://www.emulex.com/interoperability/results/matrix-action/Interop/by-partner/?tx_elxinterop_interop%5Bpartner%5D=Hitachi%20Data%20Systems&tx_elxinterop_interop%5Bsegment%5D=Storage&cHash=0c0fd579327662cb4144494f046b41dc

Table 8-3 Settings for Emulex adapters on VMware hosts

Parameter	Setting
Host Adapter BIOS	Disabled
Topology	Fabric Point-to-Point

Configuring the new devices

This section provides information about configuring the new storage devices on the Hitachi RAID storage system for operation with the VMware host.

☐ Creating the VMFS datastores

☐ Adding a hard disk to a virtual machine

Creating the VMFS datastores

Use the software on the VMware host (for example, vSphere Client) to create the VMFS datastores on the new storage devices in the Hitachi RAID storage system. Make sure to create only one VMFS datastore for each storage device. For details about configuring new storage devices (for example, supported file and block sizes), see the VMware user documentation.

Use the following settings when creating a VMFS datastore on a Hitachi RAID storage device:

- LUN properties
  - Path policy: Round robin.
  - Preference: Preferred. Always route traffic over this port when possible.
  - State: Enabled. Make this path available for load balancing and failover.
- VMFS properties
  - Storage type: disk/LUN
  - Maximum file size: 256 GB, block size 1 MB
  - Capacity: Maximum capacity



TIP: You do not need to create the VMFS datastores again on other hosts that may need access to the new storage devices. Use the storage refresh and rescan operations to update the datastore lists and storage information on the other hosts.

Adding a hard disk to a virtual machine

Use the following settings when adding a hard disk to a virtual machine for Hitachi RAID storage devices:

- When creating a new virtual disk:
  - Disk capacity (can be changed later)
  - Location: on the same datastore as the virtual machine files, or specify a datastore
- When adding an existing virtual disk: browse for the disk file path.
- When adding a mapped SAN LUN:
  - Datastore: Virtual Machine
  - Compatibility mode: physical
  - Store LUN mapping file on the same datastore as the virtual machine files
- Virtual device node: Select a node that is local to the virtual machine.
- Virtual disk mode options: Independent mode (persistent or nonpersistent)

Troubleshooting for VMware host attachment

Table 8-4 lists potential error conditions that may occur during installation of new storage and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 8-4 Troubleshooting for VMware host attachment

Error condition	Recommended action
Virtual Machine adapter does not see Lun8 and greater.	Verify cabling, storage LUN, switch and storage security and LUN masking. Verify that the Disk.MaxLUN parameter in the Advance Settings (VMware Management Interface) is set to more than 7.
Guest OS virtual machine booting up but not installing the OS.	It is possible that there is an existing corrupted vmdk file (due to an incomplete installation). Delete the vmdk file from the File Manager and remove it from the Guest OS. Add a new device for the Guest OS and recreate a new vmdk image file.
Cannot add Meta Data File for raw device.	The Meta Data File for the raw device may have existed. Selected the existing Meta Data File or delete the old Meta Data File and create a new one.
Guest OS virtual machine boots up, but does not install the operating system.	There may be a corrupt vmdk file (usually because of previous incomplete installation). Delete the vmdk file from the File Manager and remove it from the Guest OS. Add a new device for the Guest OS and recreate a new vmdk image file.
Cannot add Meta Data File for raw device.	The Meta Data File for the raw device may have existed. Select the existing Meta Data File or delete the old Meta Data File and create a new one.
Volume label is not successful.	Limit the number of characters to 30.
Cannot delete a VMFS file.	It is possible that there is an active swap file on the same extended partition. Manually turn off the swap device (using `vmkfstools` command) from the service console and try again. Relocate the swap file to another disk.
Guest OS cannot communicate with the server or outside network.	Make sure a virtual switch is created and bound to a connected network adapter.
vmkfstools -s does not add LUN online.	Delete the LUN. Select and add another LUN and retry the process again. Repeat the command or perform the Rescan SAN function in the Storage Management of the VMware Management Interface and display again.
Service console discovers online LUN addition, but the Disks and LUNs do not.	Rescan SAN and refresh.
VMware ESX Server crashes while booting up.	Check for the error message on the screen. It could be because of mixing different types of adapters in the server.

9

Windows configuration and attachment

This chapter describes how to configure the new Hitachi disk devices on a Microsoft® Windows® host:

- ☐ Hitachi storage system configuration for Windows operations
- ☐ Verifying the disk and device parameters
- ☐ Verifying new device recognition
- ☐ Configuring the new disk devices
- ☐ Creating an online LUSE volume
- ☐ Enabling MultiPath IO (MPIO)
- ☐ Troubleshooting for Windows host attachment



WARNING: Changes made to the Registry without the direct assistance of Hitachi Data Systems may jeopardize the proper operation of your Windows system and are the sole responsibility of the user.



Note: Configuration of the devices should be performed by the Windows system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Hitachi storage system configuration for Windows operations

The storage system must be fully configured before being attached to the Windows host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for Windows operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)
- Cross-OS File Exchange (FX) (3390-3A/B/C, OPEN-*x*-FXoto)

Host mode. The following table lists the required host modes for Windows host attachment. You must use either host mode **0C** or host mode **2C**. Do not select a host mode other than **0C** or **2C** for Windows. Either setting is required to support MSCS failover and to recognize more than eight LUs.

Host Mode	Description
0C[Windows]	If you use host mode 0C, you will not be able to create a LUSE volume using a volume to which an LU path has already been defined. Before performing a LUSE operation on an LDEV with a path defined from a Windows host, make sure that the host mode is 2C (Windows Extension).
2C[Windows Extension]	Use host mode **2C Windows Extension** if you plan to create LUSE volumes. If you plan to create a LUSE volume using a volume to which an LU path has already been defined, you must use host mode **2C**.

For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Verifying the disk and device parameters

Before you configure the new disk devices, verify the disk I/O timeout value, queue depth, and other required parameters such as fabric support. If you need to change any settings, reboot the Windows system, and use the setup utility for the adapter to change the settings.

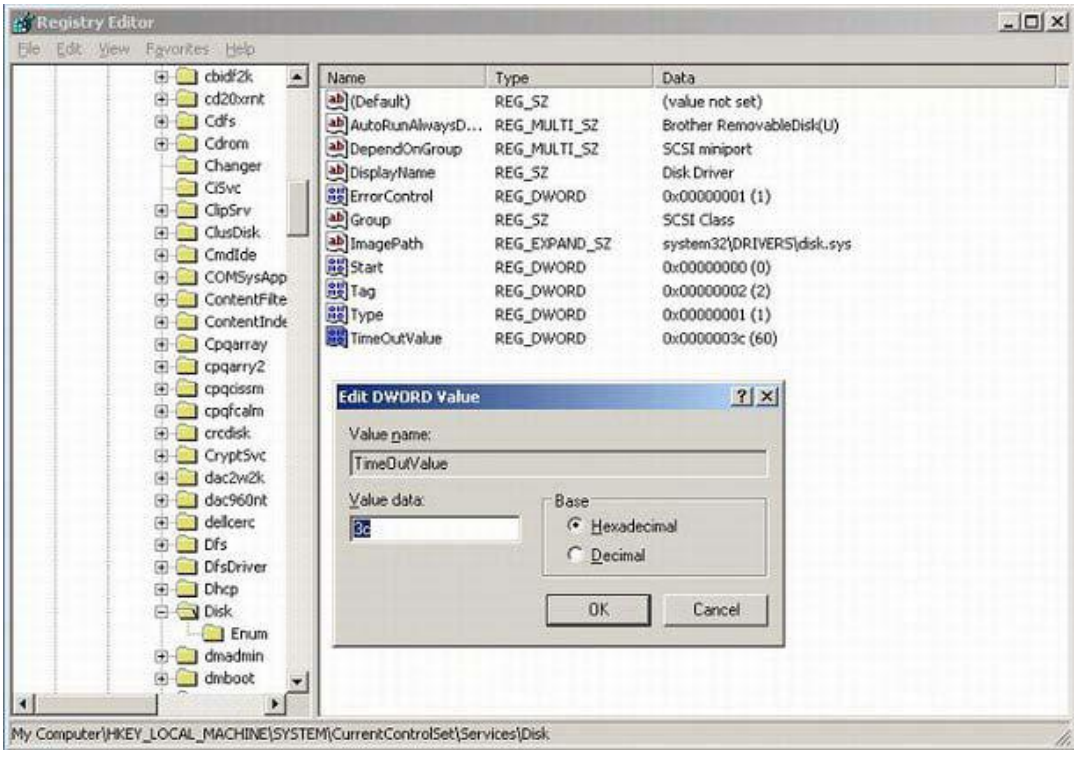
Verifying the disk I/O timeout value (TOV)

The disk I/O TOV parameter, which applies to all SCSI disk devices attached to the Windows system, must be set to 60 seconds. The default setting is hexadecimal 0x3c (decimal 60).

WARNING: The following procedure is intended for the system administrator with the assistance of the Hitachi Data Systems representative. **Use the Registry Editor with extreme caution.** Do not change the system registry without the direct assistance of Hitachi Data Systems. For information and instructions about the registry, refer to the online help for the Registry Editor.

To verify the disk I/O TOV using the Registry Editor:

1. Start the Windows Registry Editor: click **Start**, click **Run**, and enter **regedt32** in the Run dialog box.
2. Go to **HKEY_LOCAL_MACHINE** → **SYSTEM** → **CurrentControlSet** → **Services** → **Disk** to display the disk parameters.
3. Verify that the **TimeOutValue** disk parameter is set to 60 seconds (0x3c), as shown below.



4. Verify other required settings for your operational environment (for example, FC fabric support). Refer to the user documentation for the adapter as needed.
5. Exit the Registry Editor.
6. If you need to change any settings, reboot the Windows system, and use the setup utility for the adapter to change the settings. If you are not able to change the settings using the setup utility, ask your Hitachi Data Systems representative for assistance.

Verifying the queue depth

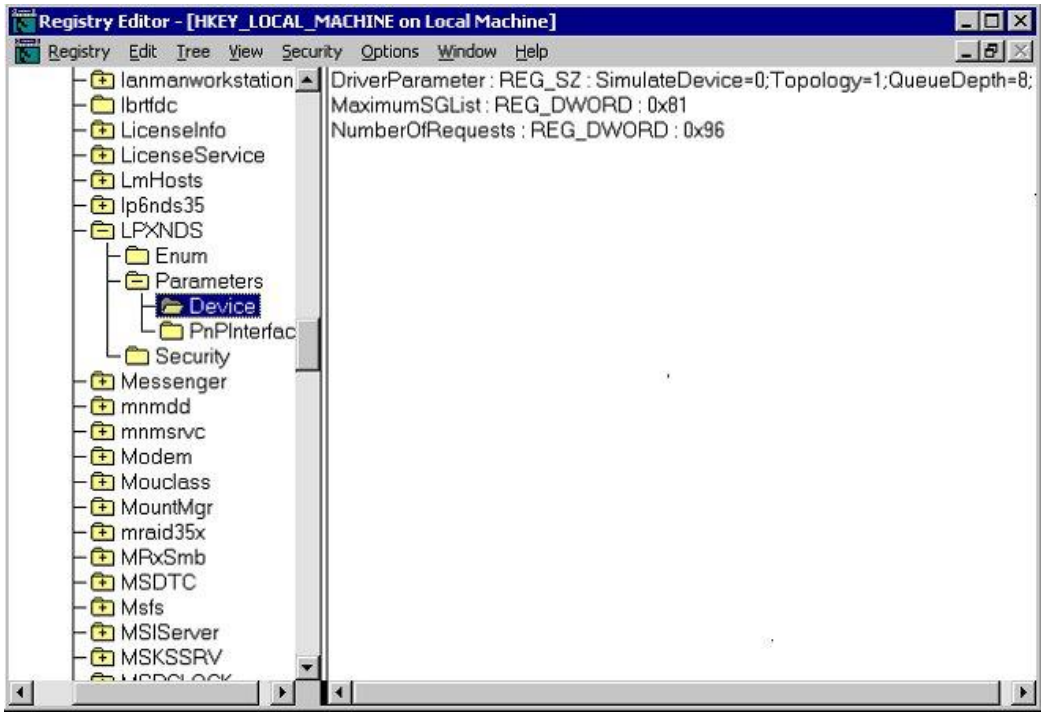
The following sample instructions describe how to verify the queue depth for a QLogic HBA using the Registry Editor.



WARNING: The following procedure is intended for the system administrator with the assistance of the Hitachi Data Systems representative. **Use the Registry Editor with extreme caution.** Do not change the system registry without the direct assistance of Hitachi Data Systems. For information and instructions about the registry, refer to the online help for the Registry Editor.

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

1. Start the Windows Registry Editor: click **Start**, click **Run**, and enter **regedt32** in the Run dialog box.
2. Go to **HKEY_LOCAL_MACHINE** → **SYSTEM** → **CurrentControlSet** → **Services** → **ql2200** (or **2300)** → **Parameters** → **Device** to display the device parameters for the QLogic HBA.



3. Verify that the queue depth value in **DriverParameter** meets the requirements for the Hitachi storage system. For details about queue depth, see Host queue depth.

Parameter	Recommended value for VSP G1000, VSP G1500, VSP F1500, VSP Gx00, VSP Fx00, HUS VM, VSP	Required Value for USP V/VM
IOCB Allocation (queue depth) per LU	32	≤ 32 per LU
IOCB Allocation (queue depth) per port (MAXTAGS)	2048	≤ 2048 per port

4. If connected to a fabric switch, make sure **FabricSupported=1** appears in **DriverParameter**.
5. Verify other required settings for your environment (for example, support for more than eight LUNs per target ID). Refer to the HBA documentation as needed.
6. Make sure the device parameters are the same for all devices on the Hitachi RAID storage system.
7. Exit the Registry Editor.
8. If you need to change any settings, reboot the Windows system, and use the HBA setup utility to change the settings. If you are not able to change the settings using the HBA utility, ask your Hitachi Data Systems representative for assistance.

Verifying new device recognition

When the adapter connected to the storage system shows the new devices (see Figure 9-1), pause the screen and record the disk number for each new device on your SCSI Device worksheet (see Table 9-1). You will need this information when you write signatures on the devices (see Writing the signatures).

Disk number assignments

The Windows system assigns the disk numbers sequentially starting with the local disks and then by adapter, and by TID/LUN. If the Hitachi RAID storage system 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 Hitachi RAID storage system 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 Hitachi RAID storage system is attached to the second adapter, the disk numbers for the storage system will start at 41.

Note: When disk devices are added to or removed from the Windows system, the disk numbers are reassigned automatically. For the FX devices, be sure to update your FX volume definition file (**datasetmount.dat**) with the new disk numbers.

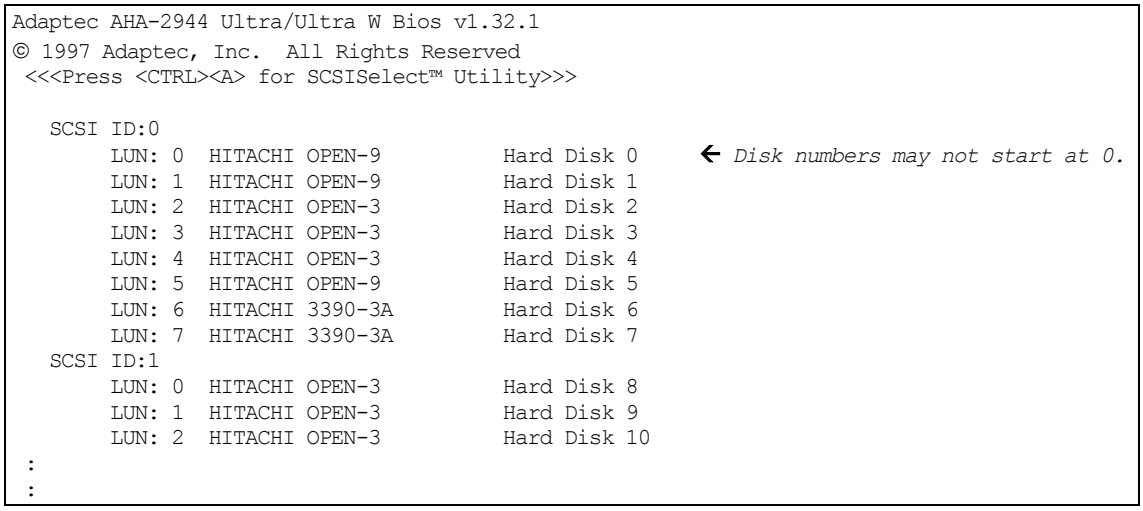


Figure 9-1 Recording the disk numbers for the new devices

Table 9-1 Sample SCSI device information worksheet

LDEV (CU:LDEV)	LU Type	VLL (MB)	Device Number	Bus Number	Path 1	Alternate Paths	
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:____
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…							

Configuring the new disk devices

This section describes how to configure the new disk devices on the Windows host.

- ☐ Writing the signatures
- ☐ Creating and formatting the partitions
- ☐ Verifying file system operations
- ☐ Verifying auto-mount
- ☐ Changing the enable write caching option

Notes:

- Do not create partitions on the FX devices. If the FX devices will be used in the MSCS environment, you must write a signature on each FX device. If not, do not write a signature.
- For information about the FC AL-PA to SCSI TID mapping, see SCSI TID Maps for FC adapters.
- Online LUSE expansion: data migration is not needed for OPEN-V (required for other LU types). A host reboot is not required for Windows. For more information, contact your Hitachi Data Systems representative.

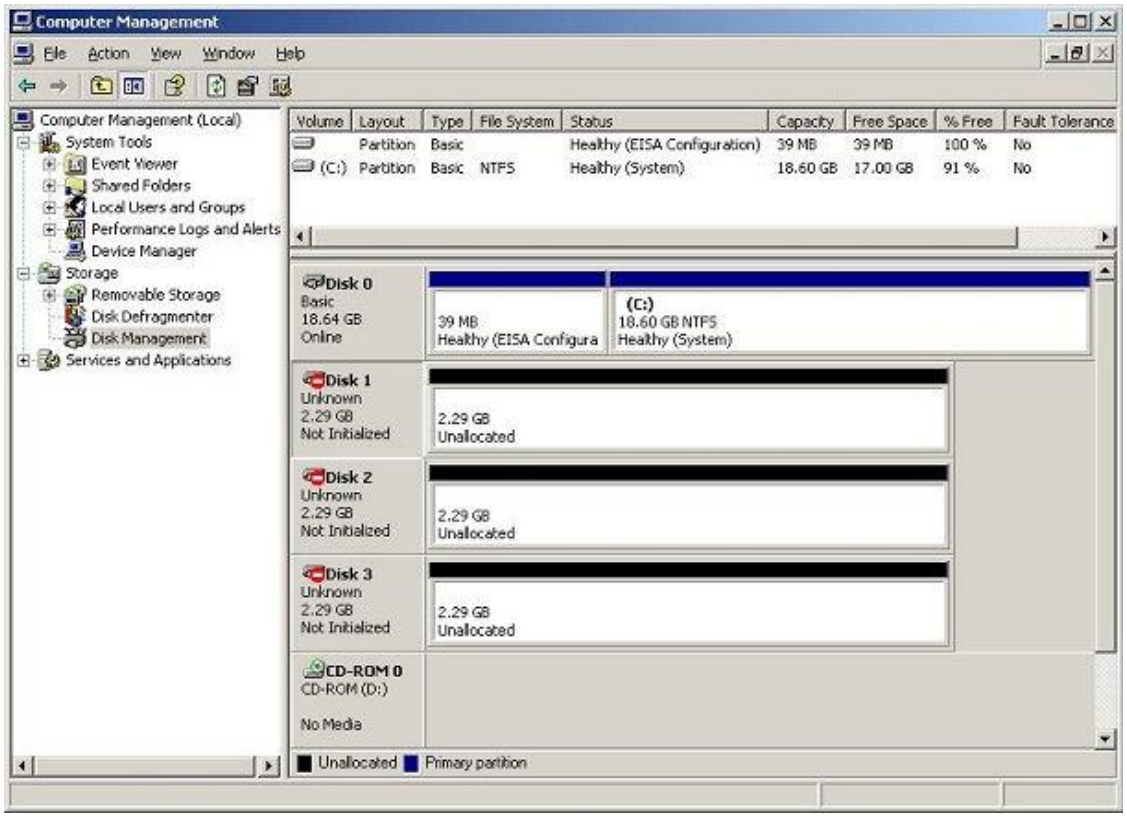
Writing the signatures

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

Note: Microsoft Windows assigns disk numbers sequentially, starting with the local disks and then by adapter, and by TID/LUN. If the Hitachi RAID storage system 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 Hitachi RAID storage system 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 Hitachi RAID storage system is attached to the second adapter, the disk numbers for the Hitachi RAID storage system start at 41.

To write the signatures on the new disk devices (see Figure 9-2):

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. The Disk Management also notifies you if any disks were removed.



Note: In the example in this figure, disk 0 is the local disk, disk 1 is an OPEN-3 device, disk 2 is an OPEN-3 device, and disk 3 is an OPEN-3 device.

Figure 9-2 Disk Management window showing new devices

3. The Disk Management displays each new device by disk number and asks if you want to write a signature on the disk (see Figure 9-3). You may only write a signature once on each device. Refer to your completed SCSI Path Worksheet (see Table 9-1) to verify the device type for each disk number.
   - For all SCSI disk devices, click **OK** to write a signature.
   - For FX devices without MSCS, click **No**.
   - For FX devices with MSCS, click **Yes** and observe this warning:



WARNING: After a signature has been written on an FX device, there is no way to distinguish the FX device from a SCSI disk device. Use extreme caution to not accidentally partition and format an FX device. This will overwrite any data on the FX device and prevent the FX software from accessing the device.

4. After you write or decline to write a signature on each new device, the Disk Management window displays the devices by disk number (see Figure 9-2). The total capacity and free space is displayed for each disk device with a signature. **Configuration information not available** indicates no signature. For directions on creating partitions on the new SCSI disk devices, see Creating and formatting the partitions.

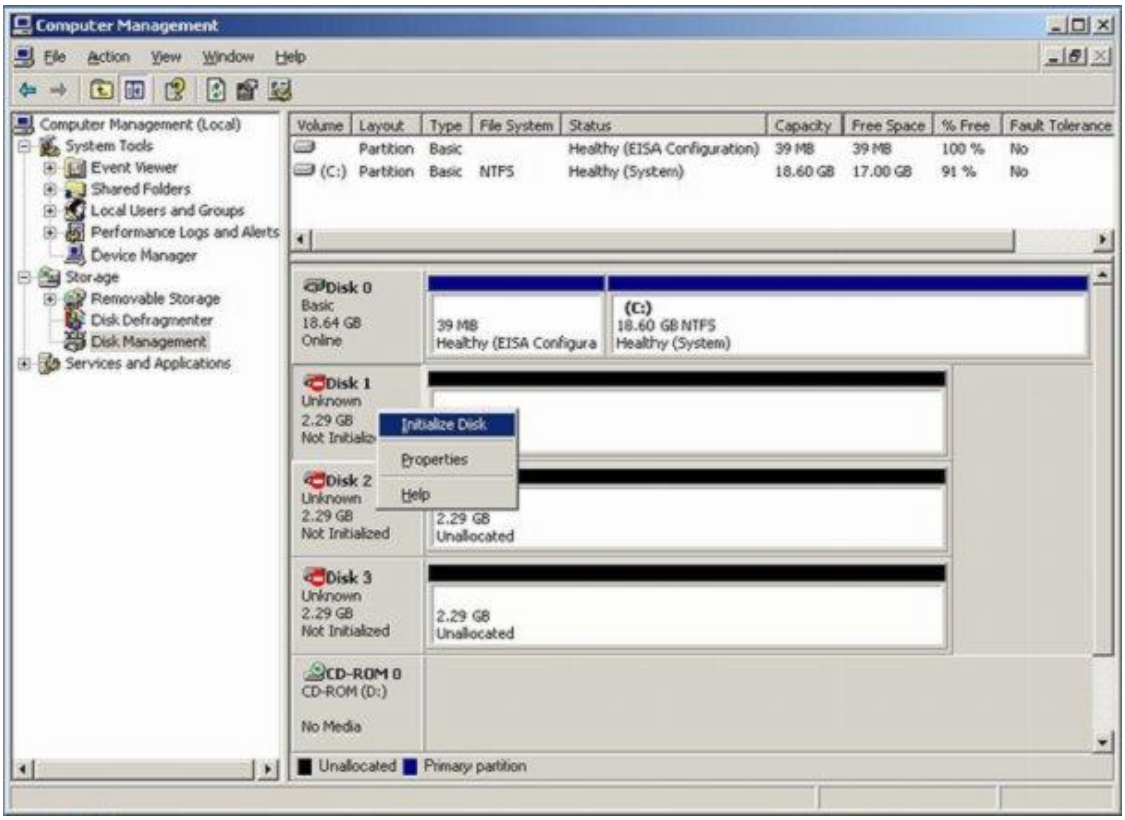


Figure 9-3 Writing the signatures

Creating and formatting the partitions

After writing signatures on the new devices, you can create and format the partitions on the new disk devices. Use your completed SCSI Device Worksheet (see Table 9-1) to verify disk numbers and device types.

Dynamic Disk is supported with no restrictions for the Hitachi RAID storage system connected to the Windows operating system. For more information, refer to the Microsoft Windows online help.

Note: Do not partition or create a file system on a device that will be used as a raw device. All FX devices are raw devices.

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

1. On the Disk Management window, select the unallocated area for the SCSI disk you want to partition, click the **Action** menu, and then click **Create Partition** to launch the New Partition Wizard.
2. When the Select Partition Type dialog box appears (see Figure 9-4), select the desired type of partition and click **Next**.

Note: The Hitachi RAID storage systems do not support Stripe Set Volume with parity.

3. When the Specify Partition Size dialog box appears (see Figure 9-5), specify the desired partition size. If the size is greater than 1024 MB, you will be asked to confirm the new partition. Click **Next**.
4. When the Assign Drive Letter or Path dialog box appears (see Figure 9-6), select a drive letter or path, or specify no drive letter or drive path. Click **Next**.
5. When the Format Partition dialog box appears (see Figure 9-7), click **Format this partition with the following settings** and select the following options:
   - **File System:** Select **NTFS** (enables the Windows system to write to the disk).
   - **Allocation unit size: Default**. Do not change this entry.
   - **Volume label:** Enter a volume label, or leave 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.
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 dialog box 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 (see Figure 9-8). Click **Finish** to close the New Partition Wizard.

8. Verify that the Disk Management window shows the correct file system (NTFS) for the formatted partition (see Figure 9-9). The word **Healthy** indicates that the partition has been created and formatted successfully.
9. Repeat steps 1-8 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.

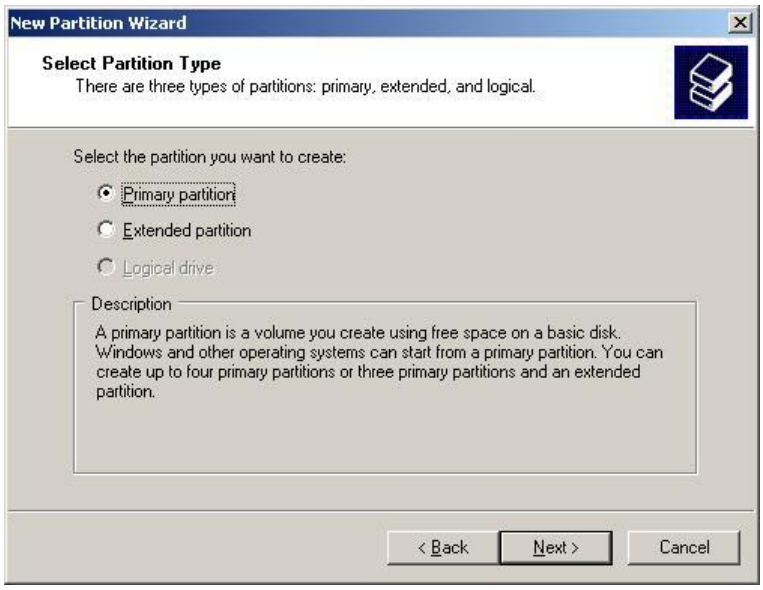


Figure 9-4 New Partition Wizard

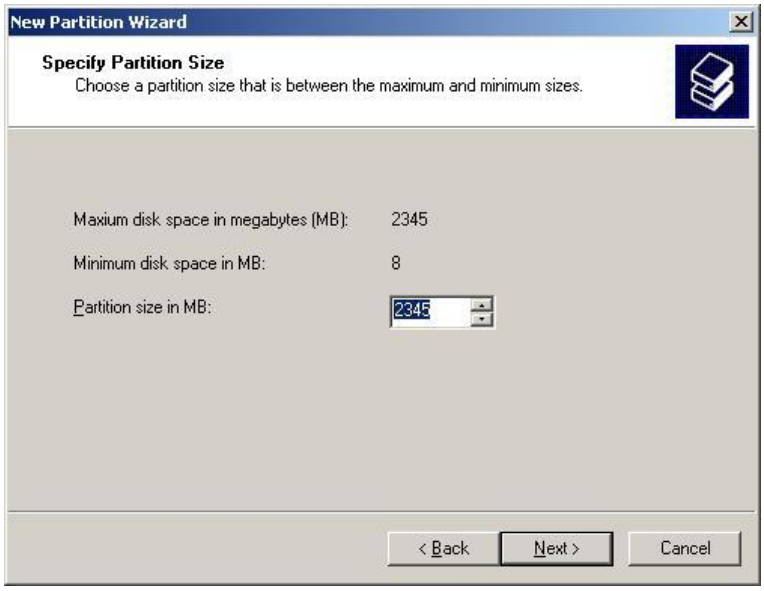


Figure 9-5 Specifying the partition size

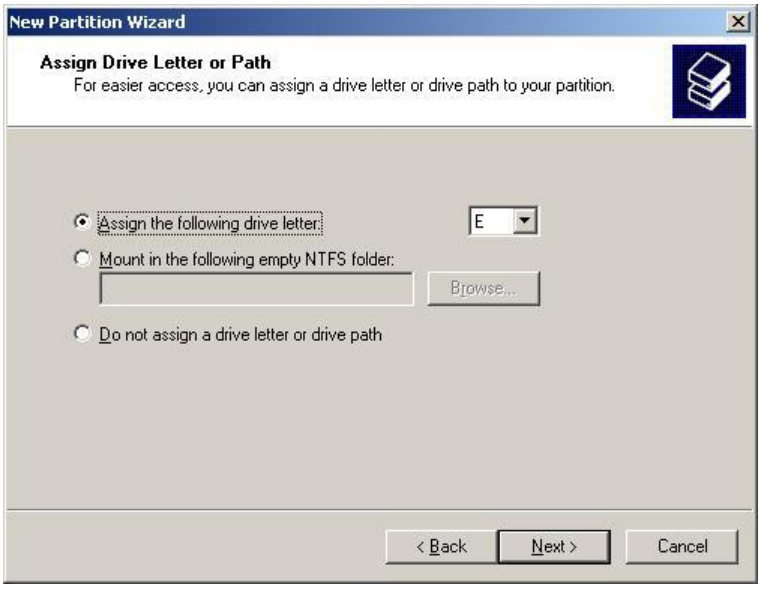


Figure 9-6 Assigning the drive letter or path

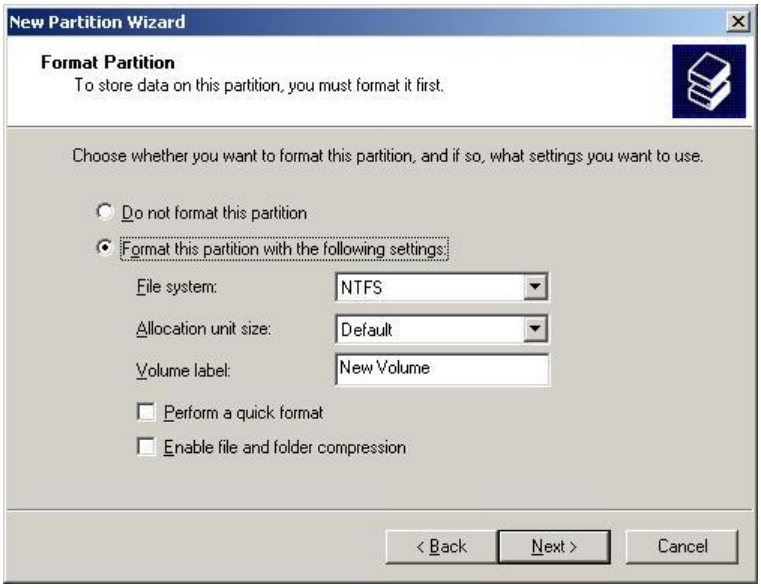


Figure 9-7 Formatting the partition



Figure 9-8 Confirmation of successful formatting

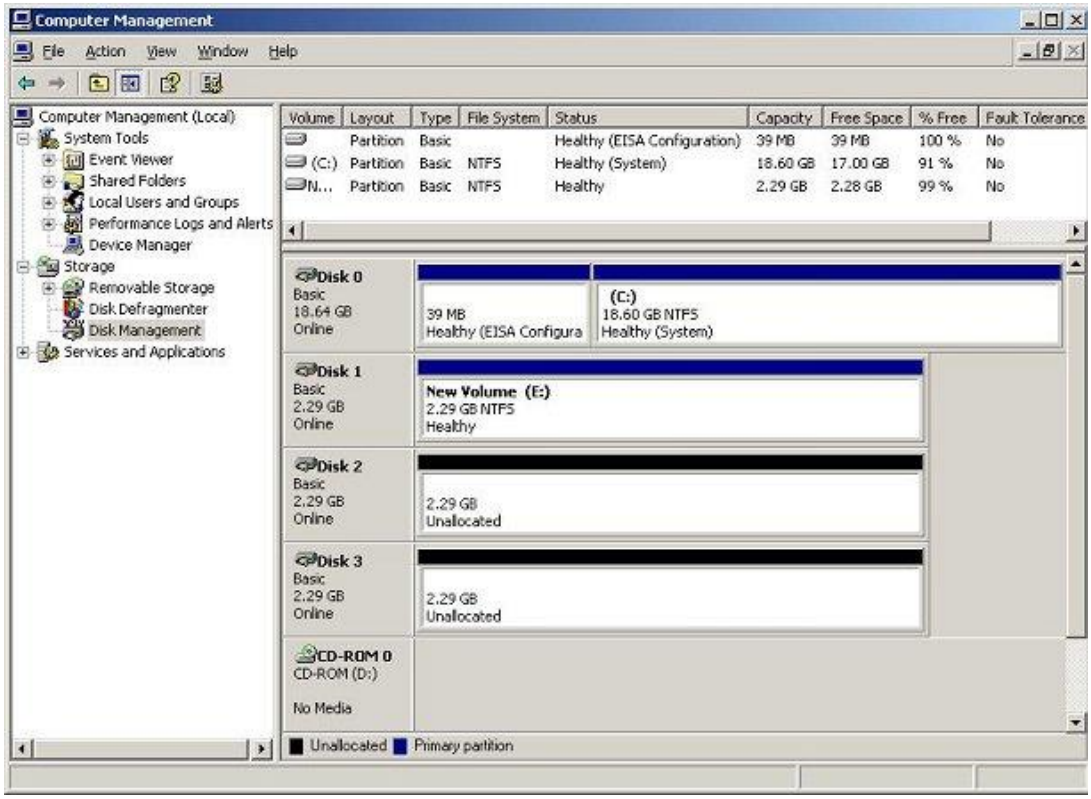


Figure 9-9 Verifying the formatted partition

Verifying file system operations

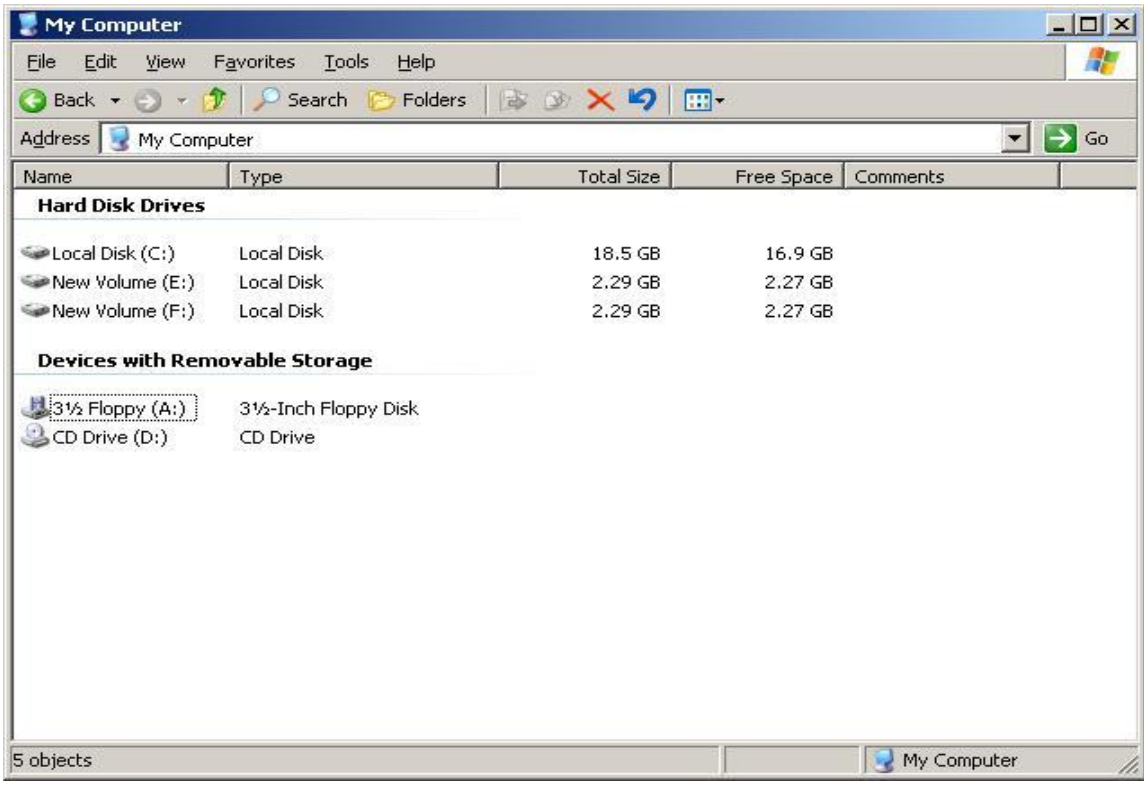
After you create and format the partitions, verify that the file system is operating properly on each new SCSI disk device. The file system enables the Windows host to access the devices. You can verify file system operation easily by copying a file onto each new device. If the file is copied successfully, this verifies that the file system is operating properly and that Windows can access the new device.



Note: Do not perform this procedure for FX and other raw devices. Instead, use the FX File Conversion Utility (FCU) or File Access Library (FAL) to access the FX devices.

To verify file system operations for the new SCSI disk devices:

1. From the Windows desktop, double-click **My Computer** to display all connected devices. All newly partitioned disks appear in this window (see Figure 9-10).
2. Select the device you want to verify, then display its Properties using either of the following methods:
   - On the **File** menu, click **Properties**.
   - Right-click and select **Properties**.
3. On the Properties dialog box (see Figure 9-11), verify that the following properties are correct:
   - Label (optional)
   - Type
   - Capacity
   - File system
4. Copy a small file to the new device.
5. Display the contents of the new device to be sure the copy operation completed successfully (see Figure 9-12). The copied file should appear with the correct file size. If desired, compare the copied file with the original file to verify no differences.
6. Delete the copied file from the new device, and verify the file was deleted successfully.
7. Repeat steps 2 through 6 for each new SCSI disk device.



Note: In the example above, (E:) and (F:) are the new devices.

Figure 9-10 Displaying the connected devices

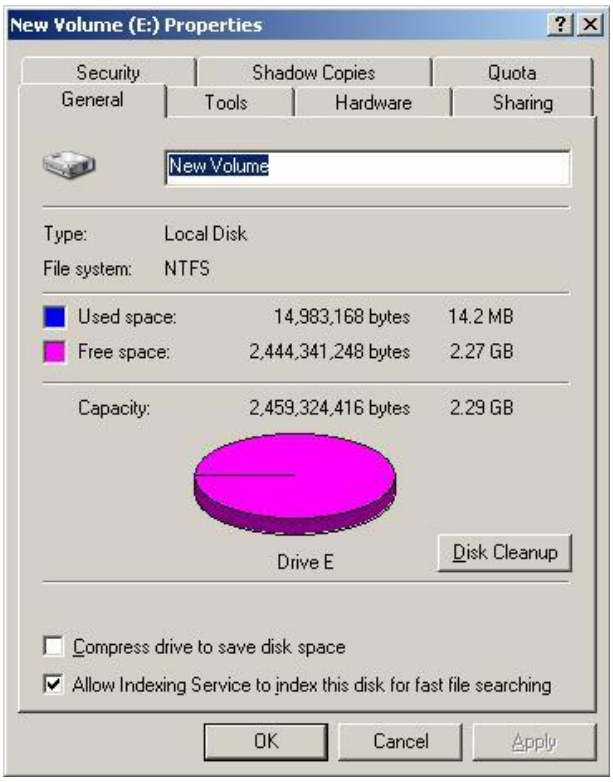


Figure 9-11 Verifying the new device properties

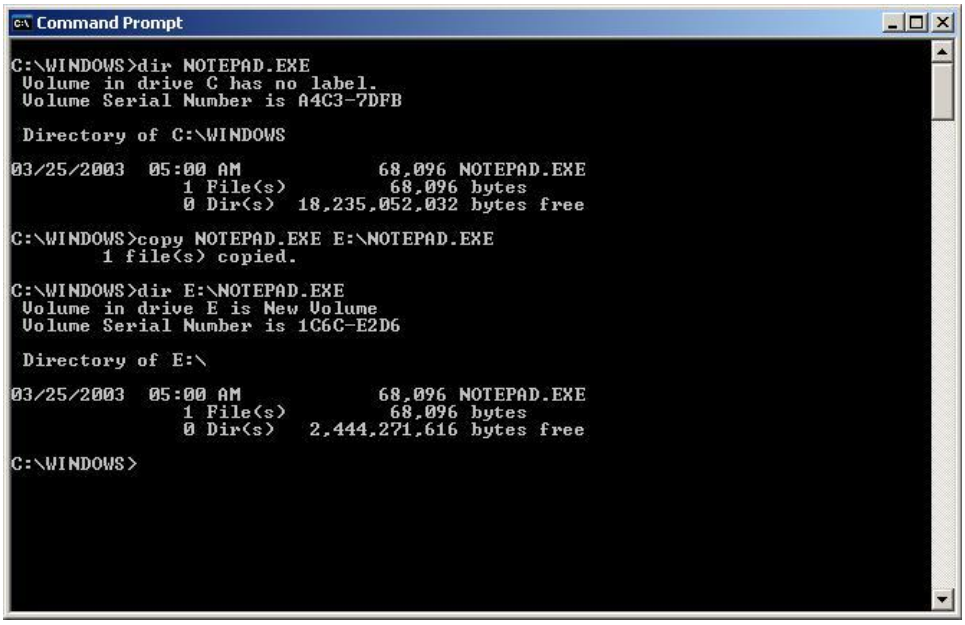


Figure 9-12 Verifying the file copy operation

Verifying auto-mount

The last step in configuring the new devices is to verify that all new devices are mounted automatically at system boot-up.

To verify auto-mount of the new devices:

1. Shut down and then restart the Windows system.
2. Open **My Computer** and verify that all new SCSI disk devices are displayed.
3. Verify that the Windows host can access each new device by repeating the procedure in Verifying file system operations:
   a. Verify the device properties for each new device (see Figure 9-11).
   b. Copy a file to each new device to be sure the devices are working properly (see Figure 9-12).

Changing the enable write caching option

The Enable Write Cache option has no effect on the cache algorithm when used with HDS storage systems and is not related to any internal Windows server caching. Microsoft and Hitachi Data Systems both recommend that you enable this option because it will provide a small improvement to Microsoft error reporting.

To enable or disable the setting **Enable write caching on the disk**:

1. Right-click **My Computer**.
2. Click **Manage**.
3. Click **Device Manager**.
4. Click the plus sign (+) next to **Disk Drives**. A list of all the disk drives appears.
5. Double-click the first HDS system disk drive.
6. Click the **Policies** or **Disk Properties** tab.
7. If **Enable write caching on the disk** is enabled, a check mark appears next to it. To disable this option, clear the check mark (see Figure 9-13).

   If the **Enable Write Cache** option is grayed-out, this option is disabled.
8. Repeat this procedure for all additional HDS system disks.

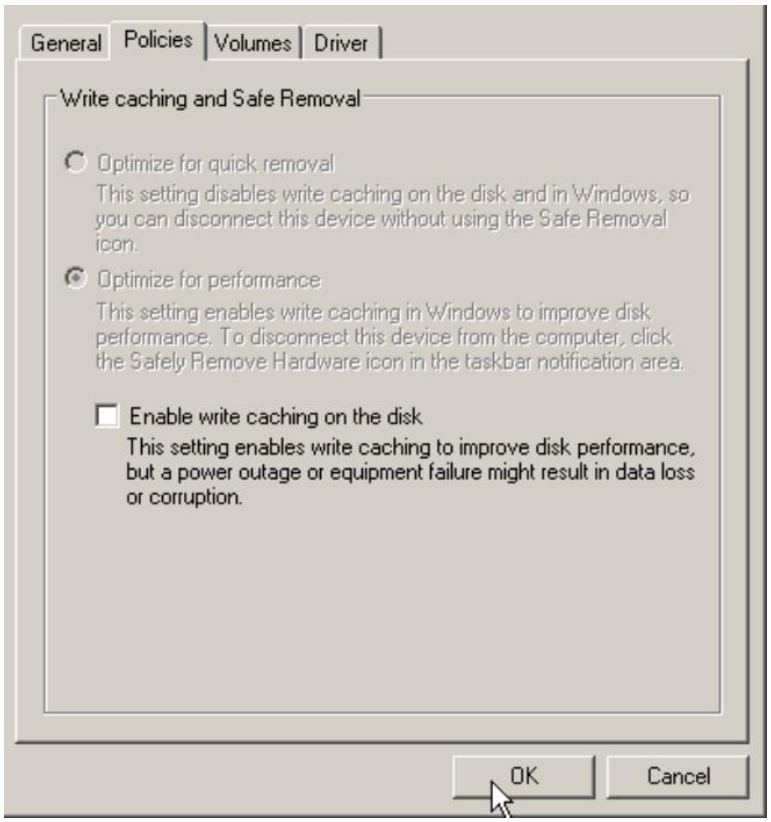


Figure 9-13 Example of disabling Enable write caching on the disk

Creating an online LUSE volume

This section explains how to safely expand a LUSE volume in an online Windows operating system.

Note:

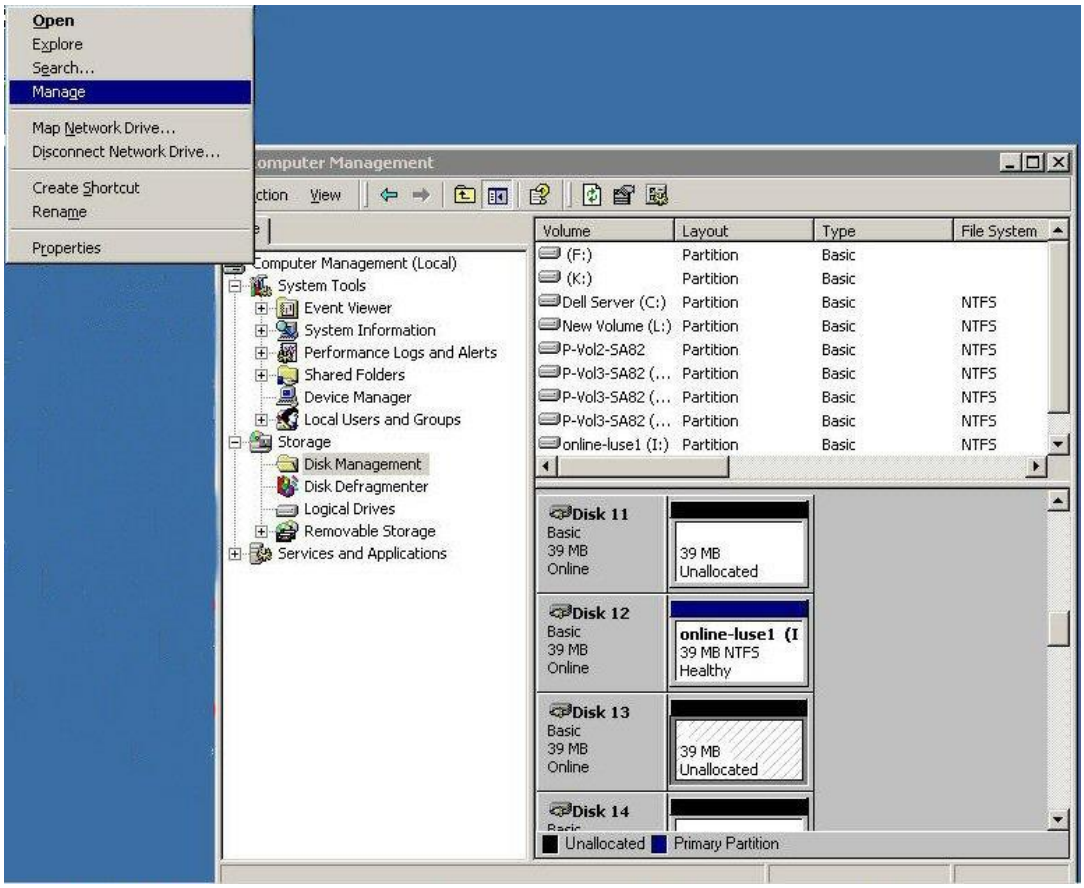
- It is recommended that you stop all I/O activity before you perform an online LUSE expansion.
- Data migration is not needed for OPEN-V (required for other LU types). A host reboot is not required for Windows. For more information, see your Hitachi Data Systems representative.

The following information applies to the instructions below:

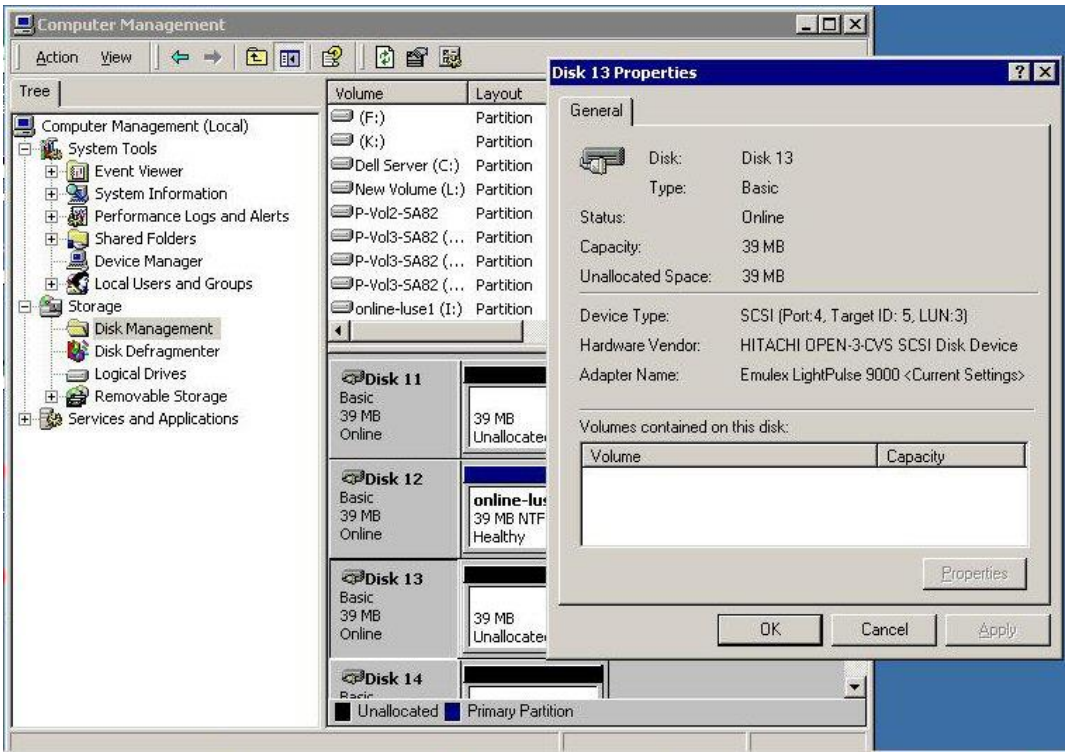
LDEV # = 0:32
Mount point = i
capacity = 40 MB

To expand a LUSE volume in an online Windows operating system:

1. On the Windows host, confirm that Disk I is mounted and Disk 12 (the disk to be expanded) is on this system: open Windows **Computer Management**, expand **Storage**, and select **Disk Management**.



2. View the disk properties (right click on the disk and select **Properties**) to get detailed information. In this example, details for Disk 13 are displayed.

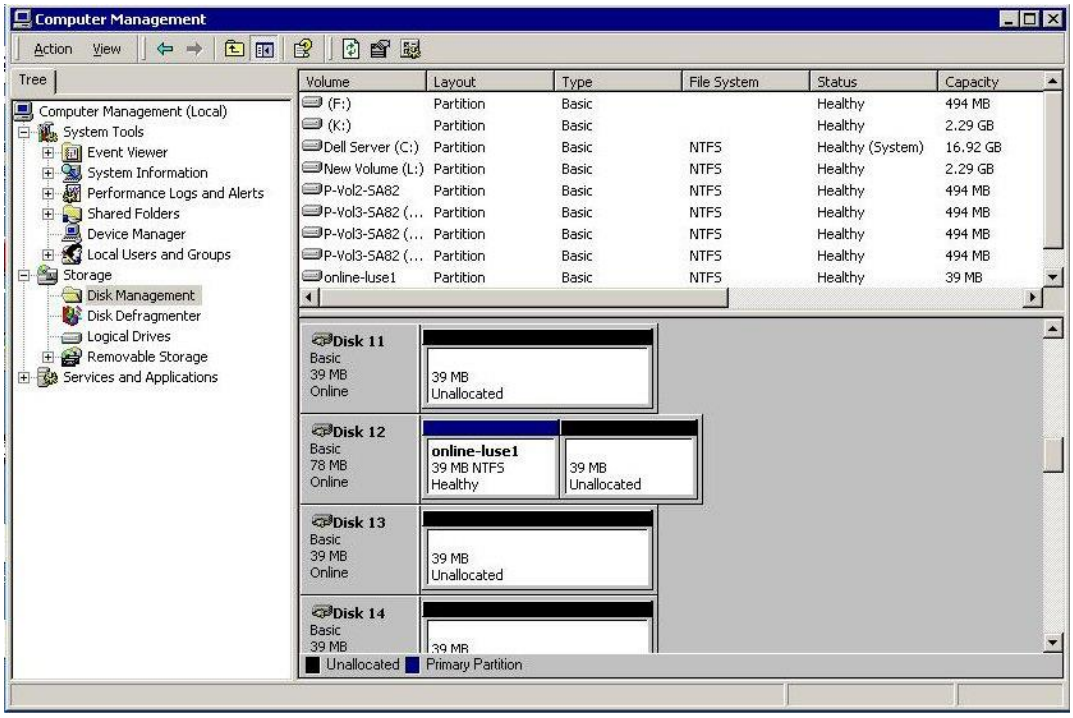


3. Create a LUSE volume. For instructions, see the *Provisioning Guide* for the storage system (or the *LUSE User's Guide* for USP V/VM).

   After creating the LUSE volume, you can configure the Windows host to recognize the expanded LDEV (for example, using DISKPART).

4. Return to the Windows **Computer Management** application, and refresh the display: select **Action** from the Menu bar, and then select **Rescan**.

   When this is done, the mounted volume I:\ (disk 12) is expanded from 40 MB to 80 MB, but the newly added disk is not yet formatted. You must now combine the new partition (for example, using DISKPART).



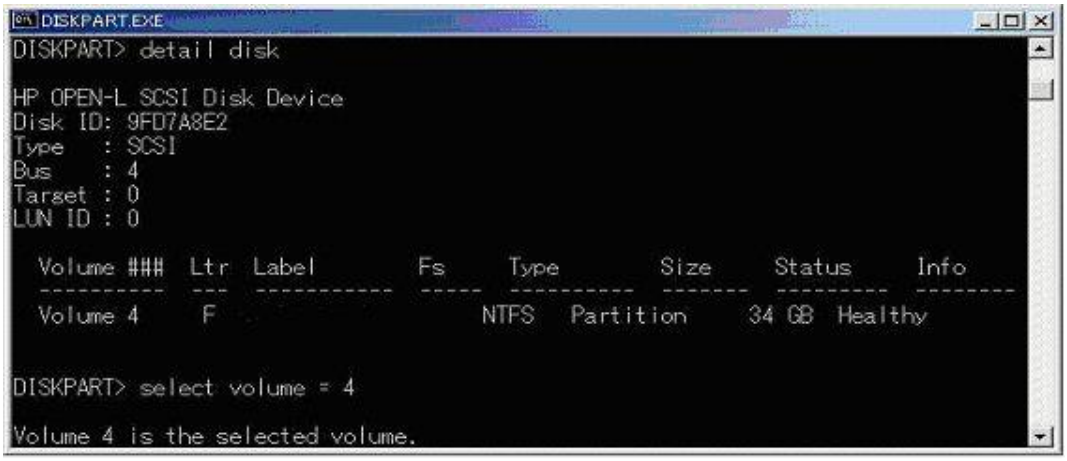
Note: Before using DISKPART, please read all applicable instructions.

5. At a command prompt, enter `Diskpart`, and press **Enter**.
6. At the `DISKPART>` prompt enter `list disk`, and press **Enter** to display the list of disks.
7. When you have identified the disk to be expanded (Disk 1 in this example), enter `select disk=1` (for this example), and press **Enter**. Disk 1 is now the selected disk on which the operations will be performed.

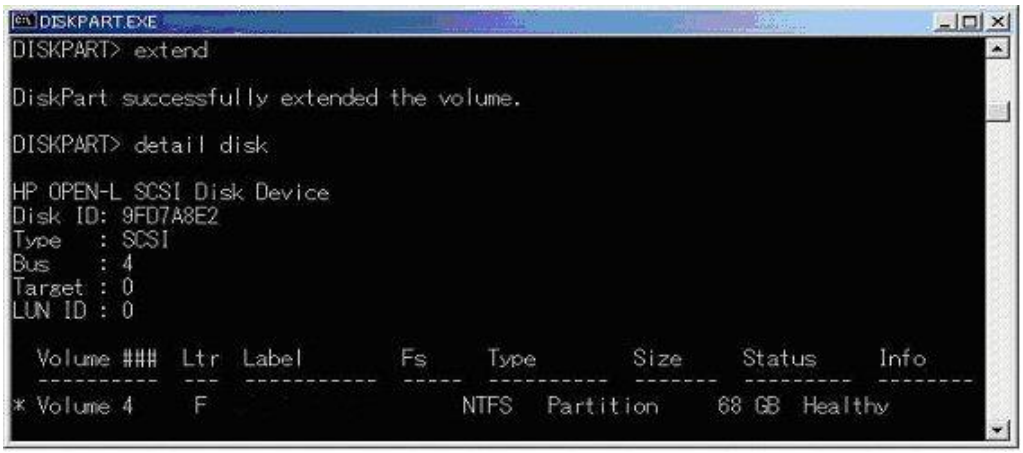


8. At the `DISKPART>` prompt enter `detail disk`, and press **Enter** to display the disk details.

9. Select the volume to be used. For this example, enter `select volume = 4`, and press **Enter**.



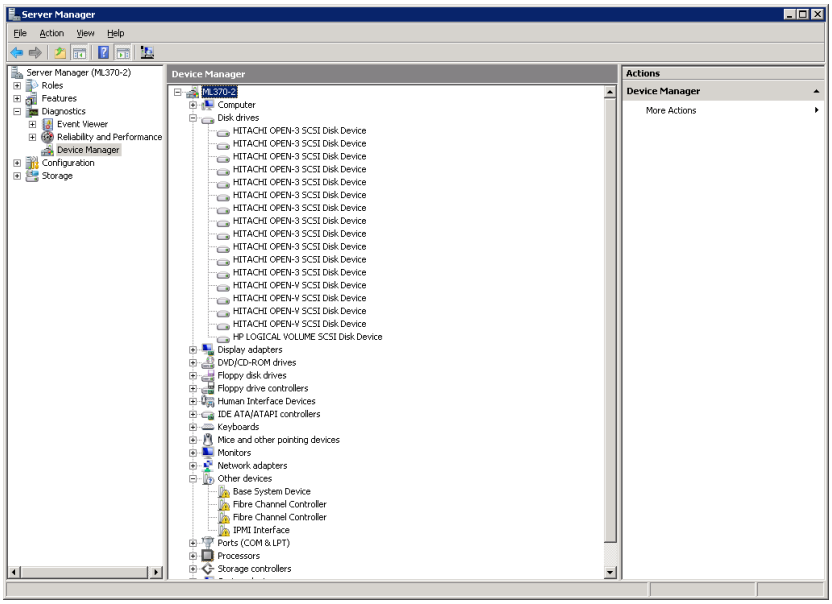
10. At the `DISKPART>` prompt, enter `extend`, and press **Enter** to combine the available volumes for the selected disk into a single partition.
11. Enter `detail disk` at the `DISKPART>` prompt, and press **Enter** to verify that the size is 68G.



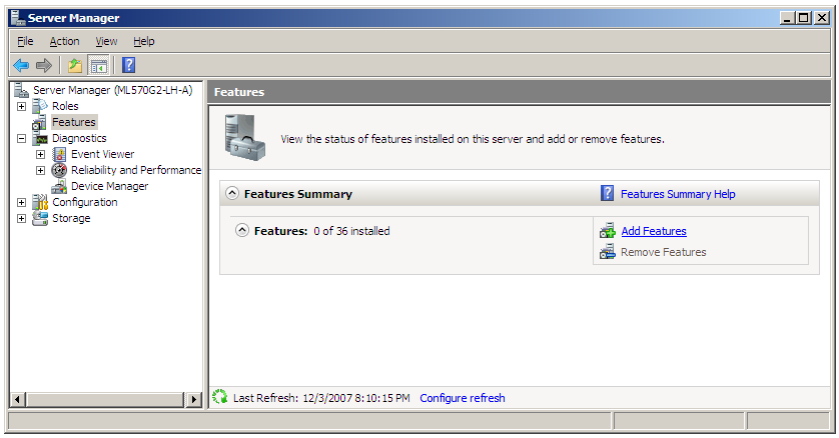
Enabling MultiPath IO (MPIO)

To enable and configure the MultiPath IO (Input/Output) feature of the Windows Server Manager for the Hitachi storage systems:

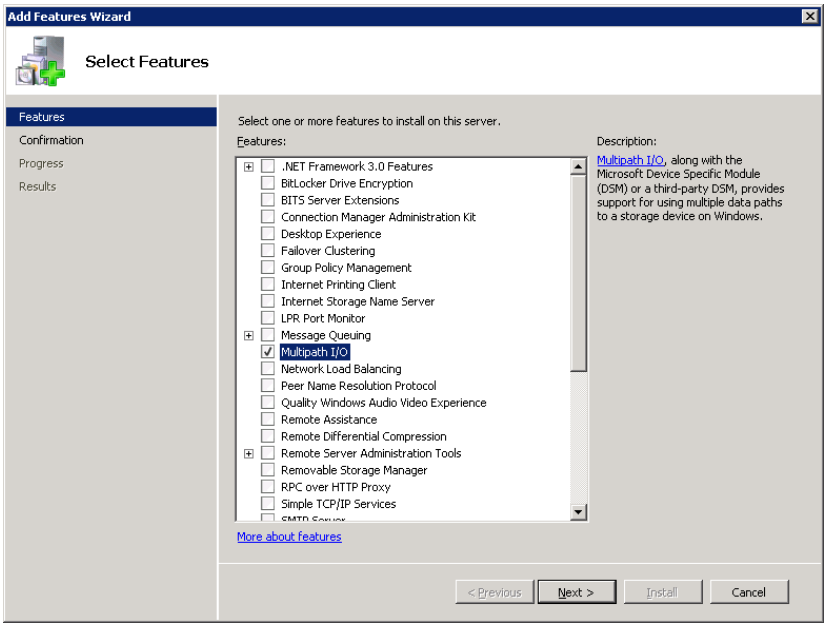
1. Launch Server Manager, and open the Administrator Tools menu.
2. Select Diagnostics, and then open Device Manager window and verify that HITACHI OPEN-x SCSI Disk Device is displayed as having n LDEV x 2 paths=2n devices.



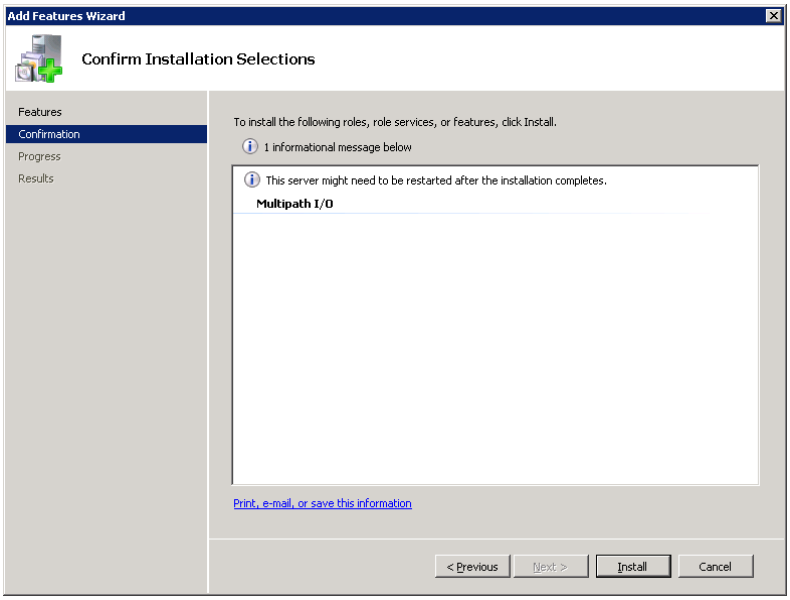
3. From Server Manager, select **Features** and click **Add Features**.



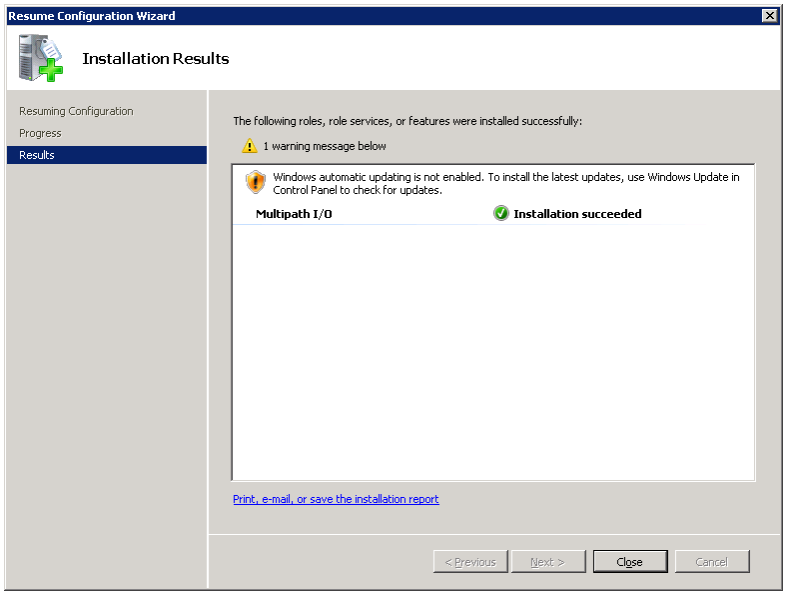
4. In the Select Features window select “Multipath I/O” and Click “Next” If the Cluster option is selected, “Multipath I/O” and “Failover Clustering” must be selected.



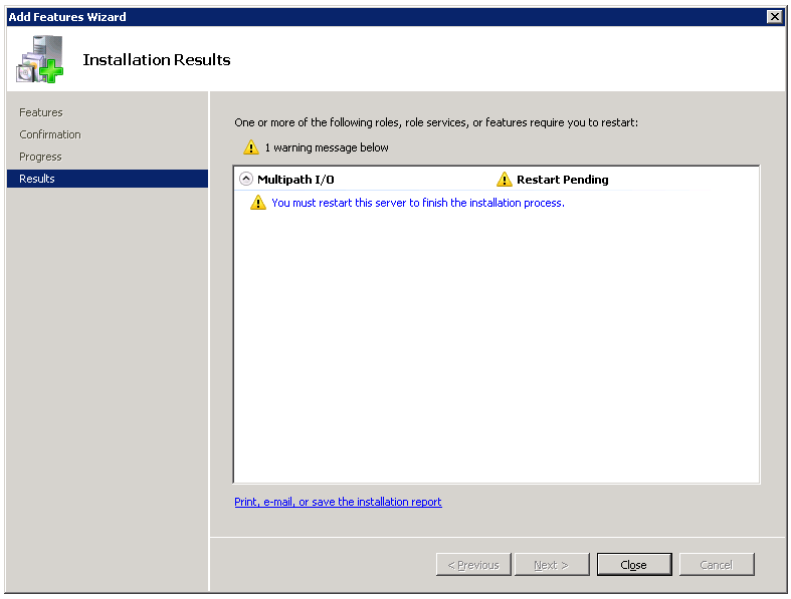
5. Confirm the installed content (Mutlipath I/O) and Click Install to start the installation.



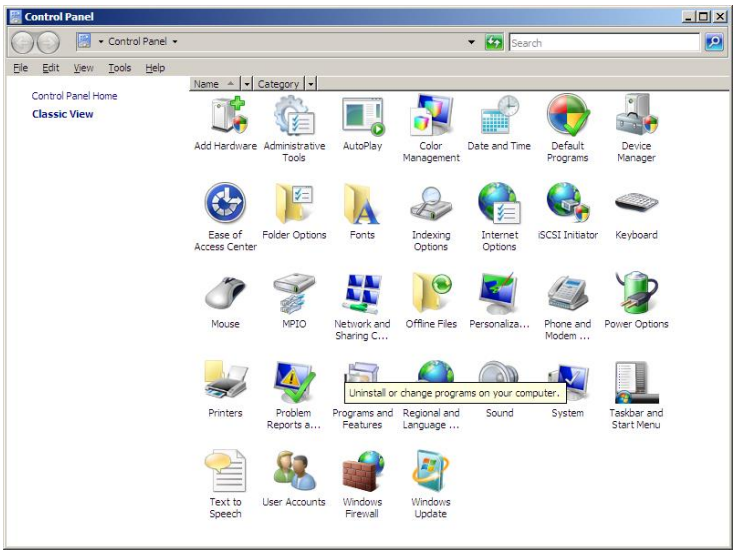
6. When the Installation Results window appears, review and confirm (if successful) by clicking **Close**.



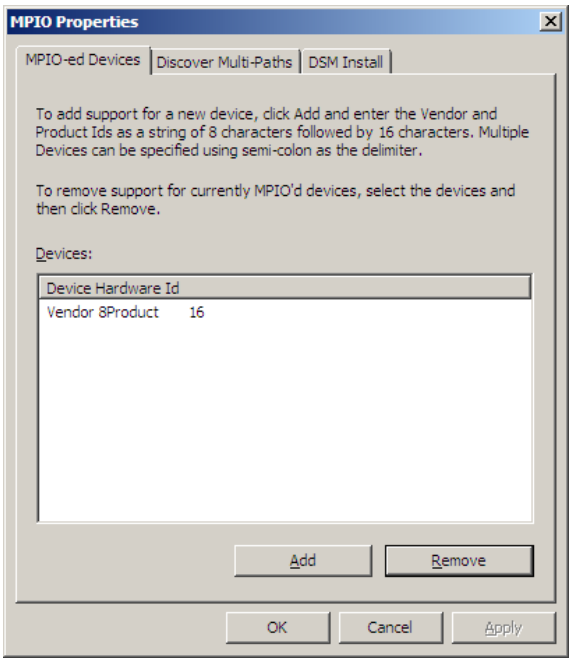
Note: If the system notice shown below appears, restart the server.



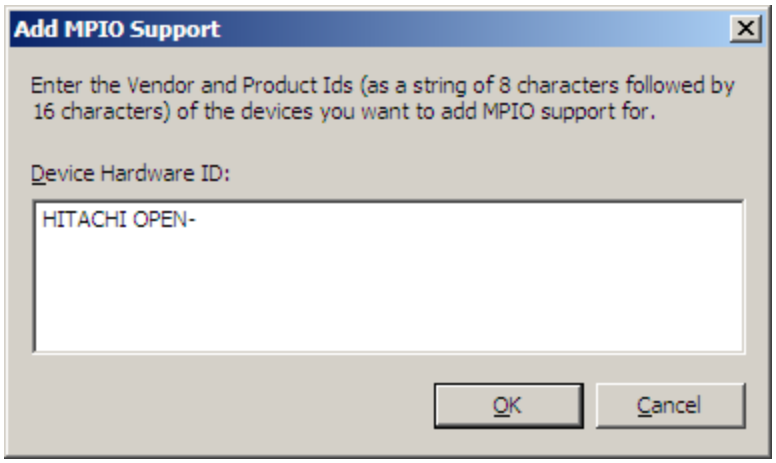
7. To launch MPIO, select **Start**, then from the Control Panel, double-click the MPIO icon.



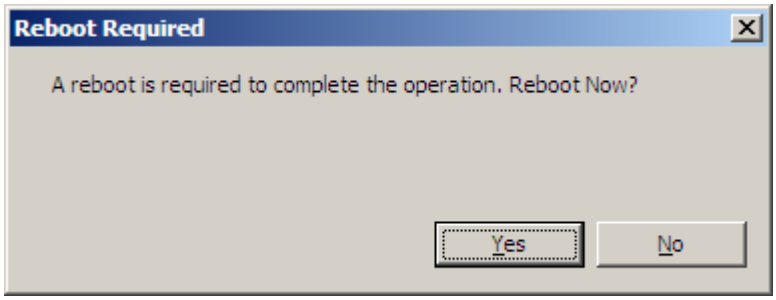
8. On the MPIO Properties window, select the **MPIO-ed Devices** tab, select the device to add, and click **Add**.



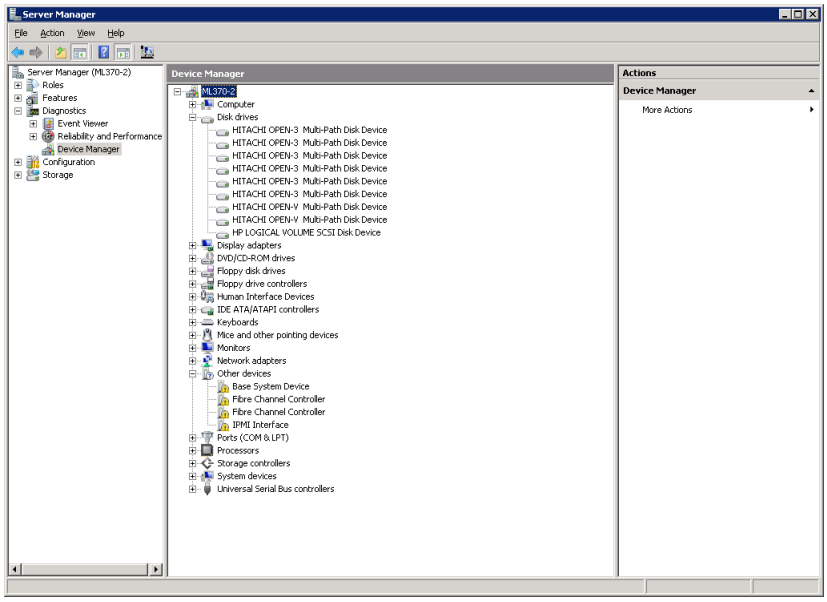
9. When the Add MPIO Support window opens, enter **HITACHI OPEN-**, and click **OK**.



10. When the Reboot Required message appears, click **Yes**.

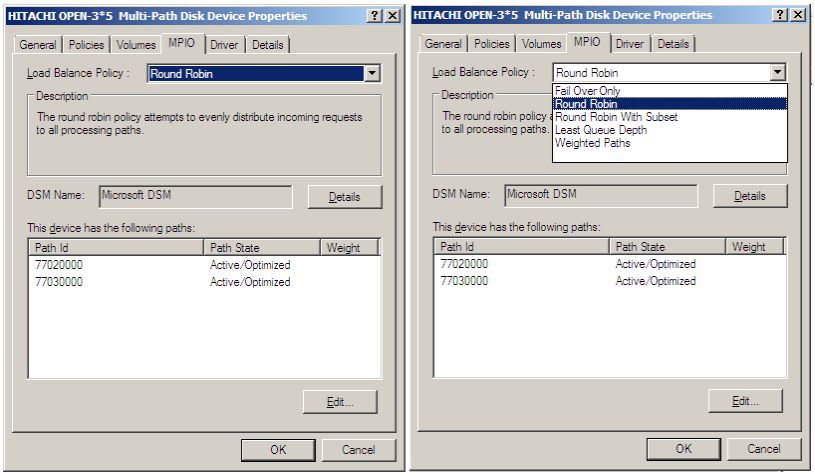


11. After the reboot, go to Server Manager, select Diagnostics and in the Device Manager window, and verify that "HITACHI OPEN-x Multi-Path Disk Device" is displayed correctly.



12. To set the Balance Policy, select the device and right-click to access its properties window. Select **Round Robin** for each LU. This policy setting is selectable on a per device basis.

    This completes enabling and configuring the MPIO feature.



Troubleshooting for Windows host attachment

Table 9-2 lists potential error conditions that may occur during installation of new storage and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 9-2 Troubleshooting for Windows host attachment

Error Condition	Recommended Action
The devices are not recognized by the system.	Be sure the READY indicator lights on the storage system are ON. Be sure the fibre cables are correctly installed and firmly connected.
The Windows system does not reboot properly after hard shutdown.	If the Windows system is powered off unexpectedly (without the normal shutdown process), wait three minutes before restarting the Windows system. This lets the storage system's internal time-out process to purge all queued commands so the storage system is available (not busy) during system startup. If the Windows system is restarted too soon, the storage system tries to process the queued commands and the Windows system will not reboot successfully.

10

XenServer configuration and attachment

This chapter describes how to configure the new Hitachi disk devices on a XenServer host:

- ☐ Hitachi storage system configuration for XenServer operations
- ☐ Recognizing the new devices
- ☐ Creating storage repositories
- ☐ Configuring the new storage devices for host use
- ☐ Troubleshooting for XenServer host attachment



Note: Configuration of the devices should be performed by the XenServer system administrator. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact your Hitachi Data Systems representative.

Hitachi storage system configuration for XenServer operations

The storage system must be fully configured before being attached to the XenServer host, as described in Configuring the Hitachi RAID storage system.

Devices types. The following devices types are supported for XenServer operations. For details, see Device types.

- OPEN-V
- OPEN-3/8/9/E/L
- LUSE (OPEN-*x**n)
- VLL (OPEN-*x* VLL)
- VLL LUSE (OPEN-*x**n VLL)

Host mode. The required host mode for XenServer is **00**. Do not select a host mode other than **00** for XenServer. For a complete list of host modes for the storage system, see Host modes and host mode options. For instructions on setting the host modes, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Host mode options. You may also need to set host mode options (HMOs) to meet your operational requirements. For a complete list of HMOs for the storage system, see Host modes and host mode options. For instructions on setting the HMOs, see the *Provisioning Guide* for the storage system (for USP V/VM see the *LUN Manager User's Guide*).

Recognizing the new devices

Once the Hitachi RAID storage system has been installed and connected, you are ready to recognize and configure the new storage devices on the Hitachi RAID storage system. The devices on the Hitachi RAID storage system do not require any special procedures and are configured in the same way as any new (HBA-attached) SCSI disk devices. You can use the XenCenter software or the XenServer CLI (**sr-probe** command) to recognize and configure the new storage devices. For details and instructions, see the XenServer user documentation.

Figure 10-1 shows the XenCenter New Storage wizard for configuring new storage. Under **Virtual disk storage** select **Hardware HBA** for the new devices on the Hitachi RAID storage system.

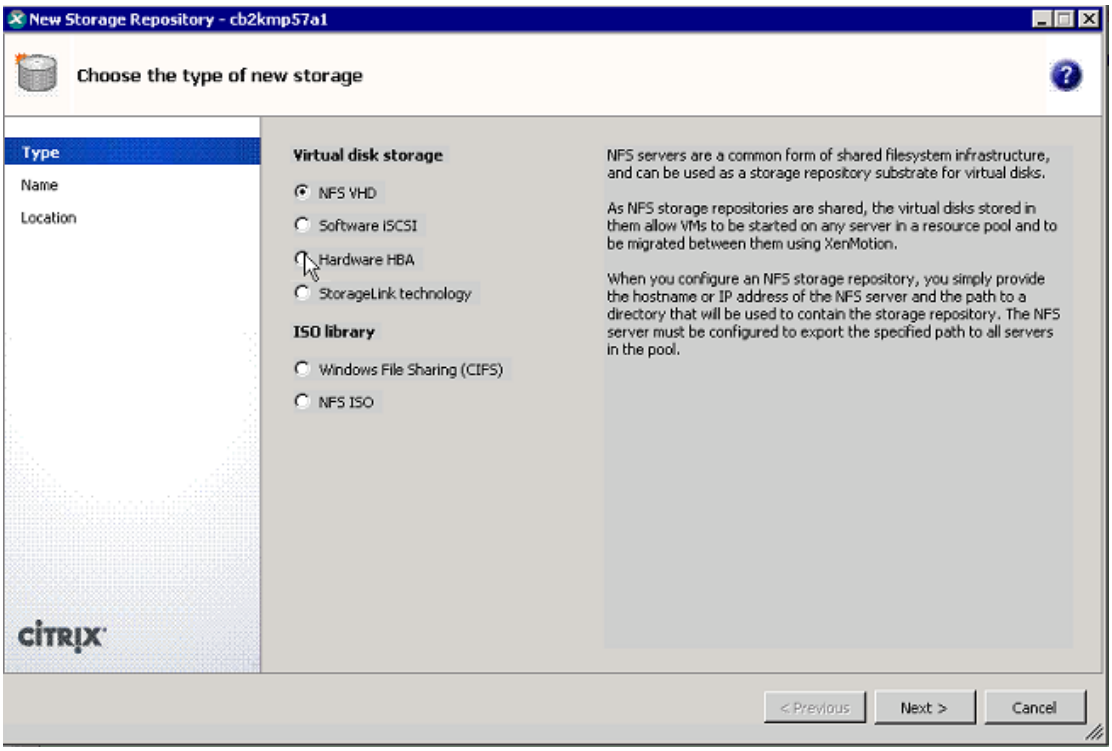


Figure 10-1 Recognizing the new storage devices

The new storage devices are recognized by the XenServer host as new scsi disk devices that are symlinked under the directory **/dev/disk/by_id** using the unique **scsi_ids**. To display the **scsi_id** for a specific device, use the **sginfo** command with the device path, for example:

```
sginfo /dev/disk/by_id/ {scsi_id}
```

Creating storage repositories

After recognizing the new disk devices, you can create storage repositories (SRs) for the new storage. Figure 10-2 shows the creation of an SR using the XenCenter software. Figure 10-3 shows the device status (**OK**, **Connected**) and multipathing status (**2 of 2 paths active**) of a new SR (called **new lun**) for a device on a Hitachi RAID storage system.

For details about SRs and instructions for creating and managing SRs, see the XenServer user documentation.

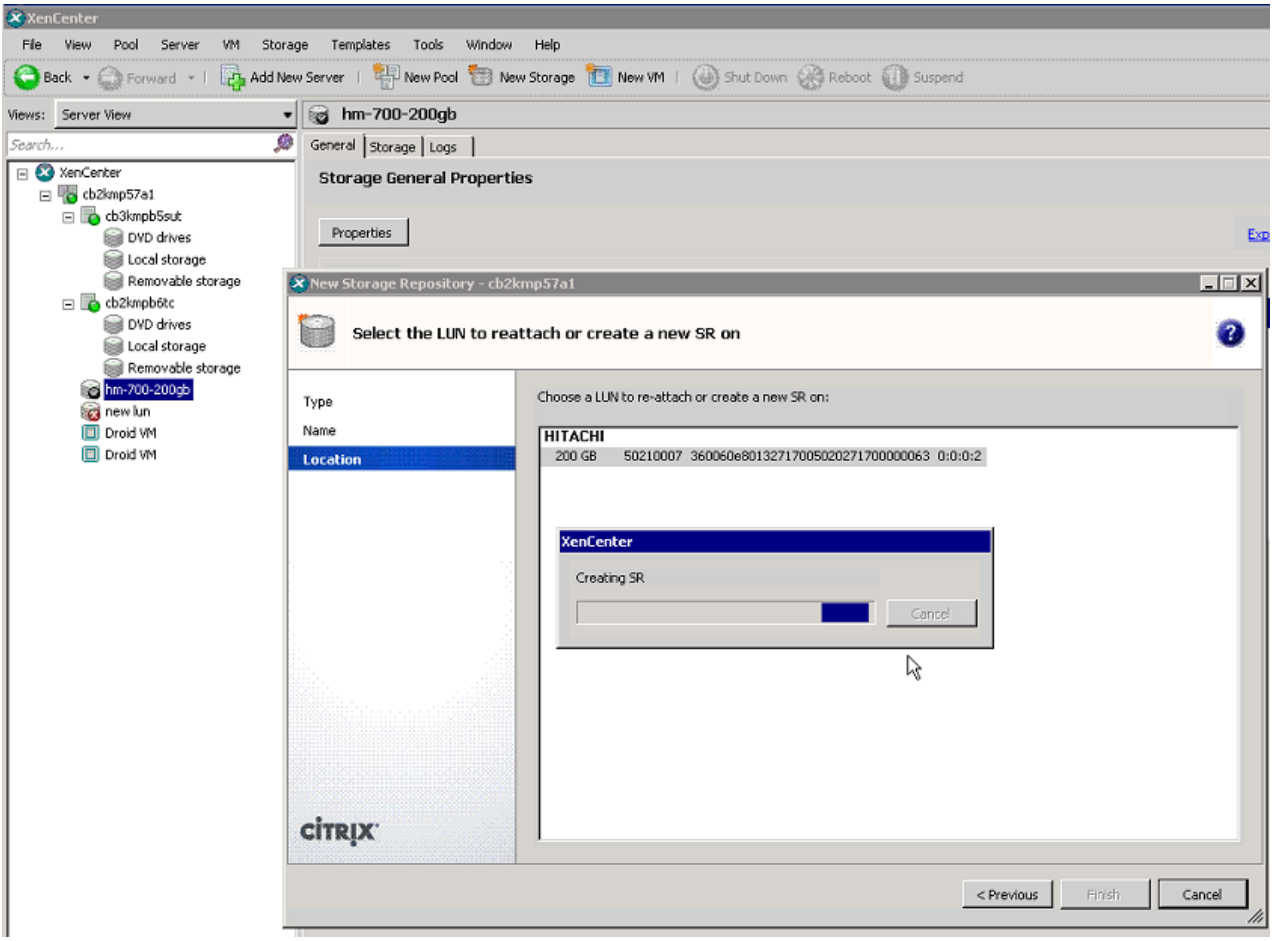


Figure 10-2 Creating a new storage repository

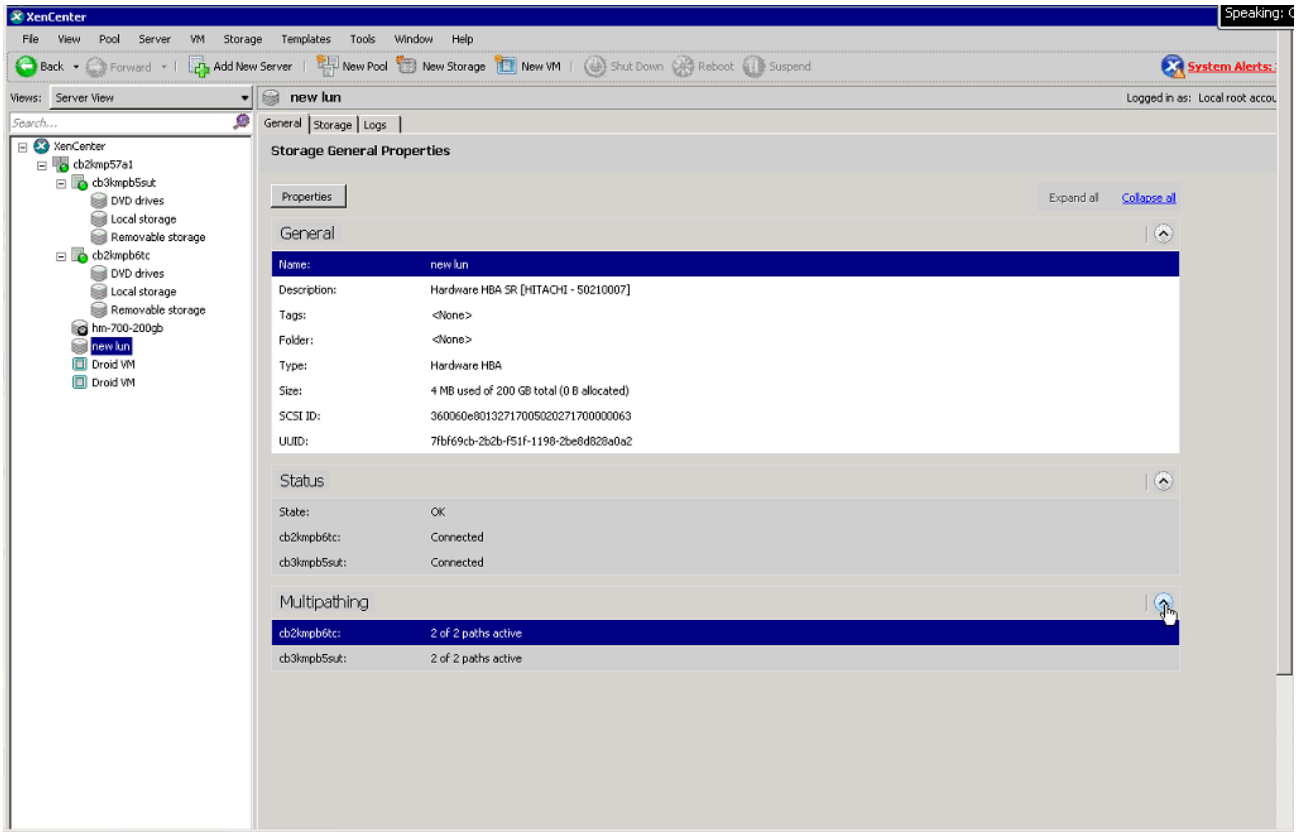


Figure 10-3 Verifying new device status

Configuring the new storage devices for host use

After the SRs have been created and the status of the new SRs has been verified, you can configure the new storage devices for use by the Citrix XenServer host, for example, adding virtual disks (vdisks) and dynamic LUNs.

For details and instructions for configuring and managing fibre-channel attached storage devices, see the Citrix XenServer user documentation.

Troubleshooting for XenServer host attachment

Table 10-1 lists potential error conditions that might occur during storage system installation on a XenServer host and provides instructions for resolving the conditions. If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Table 10-1 Troubleshooting for XenServer host attachment

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Be sure the READY indicator lights on the storage system are ON. Run **sr-probe** to recheck the fibre channel for new devices. Be sure LUSE devices are not intermixed with normal LUs on the same fibre-channel port. Verify that LUNs are configured properly for each TID.

11

General troubleshooting

This chapter provides general troubleshooting information and instructions for contacting Hitachi Data Systems technical support.

- ☐ General troubleshooting
- ☐ Contacting Hitachi Data Systems technical support

General troubleshooting

For general troubleshooting information, see the following documentation:

- For troubleshooting information for the Hitachi RAID storage system, see the User and Reference Guide for the storage system (for example, *Hitachi Virtual Storage Platform User and Reference Guide*).
- For troubleshooting information for the Hitachi Command Suite software, see the *Hitachi Command Suite Administrator Guide*.
- For troubleshooting information for the Storage Navigator software, see the *Hitachi Storage Navigator User Guide* for the storage system.
- For information about error messages displayed by Hitachi Command Suite, see the *Hitachi Command Suite Messages Guide*.
- For information about error messages displayed by Storage Navigator, see the *Storage Navigator Messages* document for the storage system.

If you cannot resolve an error condition, contact your Hitachi Data Systems representative, or contact Hitachi Data Systems technical support for assistance. For instructions, see Contacting Hitachi Data Systems technical support.

Contacting Hitachi Data Systems technical support

If you need to contact Hitachi Data Systems technical support, please 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 systems.
- The exact content of any error messages displayed by the Hitachi Command Suite software.
- The exact content of any error messages displayed by the Storage Navigator software.
- The Storage Navigator configuration information (use the Dump Tool).
- The service information messages (SIMs), including reference codes and severity levels, displayed by Storage Navigator.

The Hitachi Data Systems customer support staff is available 24 hours a day, seven days a week. To contact technical support, log on to Hitachi Data Systems Support Connect for contact information: https://support.hds.com/en_us/contact-us.html

SCSI TID Maps for FC adapters

When an arbitrated loop (AL) is established or re-established, the port addresses are assigned automatically to prevent duplicate target IDs (TID). When using the SCSI over fibre-channel protocol (FCP), TIDs are no longer needed. SCSI is a bus-oriented protocol requiring each device to have a unique address since all commands go to all devices.

For fibre channel, the AL-PA is used instead of the TID to direct packets to the desired destination. Unlike traditional SCSI, once control of the loop is acquired, a point-to-point connection is established from the initiator to the target. To enable transparent use of FCP, the host operating system "maps" a TID to each AL-PA.

Table A-1 and Table A-2 identify the fixed mappings between the bus/TID/LUN addresses assigned by the host OS and the fibre-channel native addresses (AL_PA/SEL_ID) for fibre-channel adapters. There are two potential mappings depending on the value of the ScanDown registry parameter:

- For ScanDown = 0 (default) see Table A-1.
- For ScanDown = 1 see Table A-2.



Note: When Hitachi RAID storage system devices and other types of devices are connected in the same arbitrated loop, the mappings defined in Table A-1 and Table A-2 cannot be guaranteed.

Table A-1 SCSI TID map (ScanDown=0)

Bus #	TID	LUN	AL_PA	SEL_ID
0	0-31	0-7	NONE	NONE
1	0	0-7	0x01	0x7D
	1	0-7	0x02	0x7C
	2	0-7	0x04	0x7B
	3	0-7	0x08	0x7A
	4	0-7	0x0F	0x79
	5	0-7	0x10	0x78
	6	0-7	0x17	0x77
	7	0-7	0x18	0x76
	8	0-7	0x1B	0x75
	9	0-7	0x1D	0x74
	10	0-7	0x1E	0x73
	11	0-7	0x1F	0x72
	12	0-7	0x23	0x71
	13	0-7	0x25	0x70
	14	0-7	0x26	0x6F
	15	0-7	0x27	0x6E
	16	0-7	0x29	0x6D
	17	0-7	0x2A	0x6C
	18	0-7	0x2B	0x6B
	19	0-7	0x2C	0x6A
	20	0-7	0x2D	0x69
	21	0-7	0x2E	0x68
	22	0-7	0x31	0x67
	23	0-7	0x32	0x66
	24	0-7	0x33	0x65
	25	0-7	0x34	0x64
	26	0-7	0x35	0x63
	27	0-7	0x36	0x62
	28	0-7	0x39	0x61
	29	0-7	0x3A	0x60
	30	0-7	0x3C	0x5F
	31	0-7	NONE	NONE

Table A-1 SCSI TID map (ScanDown=0) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
2	0	0-7	0x43	0x5E
	1	0-7	0x45	0x5D
	2	0-7	0x46	0x5C
	3	0-7	0x47	0x5B
	4	0-7	0x49	0x5A
	5	0-7	0x4A	0x59
	6	0-7	0x4B	0x58
	7	0-7	0x4C	0x57
	8	0-7	0x4D	0x56
	9	0-7	0x4E	0x55
	10	0-7	0x51	0x54
	11	0-7	0x52	0x53
	12	0-7	0x53	0x52
	13	0-7	0x54	0x51
	14	0-7	0x55	0x50
	15	0-7	0x56	0x4F
	16	0-7	0x59	0x4E
	17	0-7	0x5A	0x4D
	18	0-7	0x5C	0x4C
	19	0-7	0x63	0x4B
	20	0-7	0x65	0x4A
	21	0-7	0x66	0x49
	22	0-7	0x67	0x48
	23	0-7	0x69	0x47
	24	0-7	0x6A	0x46
	25	0-7	0x6B	0x45
	26	0-7	0x6C	0x44
	27	0-7	0x6D	0x43
	28	0-7	0x6E	0x42
	29	0-7	0x71	0x41
	30	0-7	0x72	0x40
	31	0-7	NONE	NONE

Table A-1 SCSI TID map (ScanDown=0) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
3	0	0-7	0x73	0x3F
	1	0-7	0x74	0x3E
	2	0-7	0x75	0x3D
	3	0-7	0x76	0x3C
	4	0-7	0x79	0x3B
	5	0-7	0x7A	0x3A
	6	0-7	0x7C	0x39
	7	0-7	0x80	0x38
	8	0-7	0x81	0x37
	9	0-7	0x82	0x36
	10	0-7	0x84	0x35
	11	0-7	0x88	0x34
	12	0-7	0x8F	0x33
	13	0-7	0x90	0x32
	14	0-7	0x97	0x31
	15	0-7	0x98	0x30
	16	0-7	0x9B	0x2F
	17	0-7	0x9D	0x2E
	18	0-7	0x9E	0x2D
	19	0-7	0x9F	0x2C
	20	0-7	0xA3	0x2B
	21	0-7	0xA5	0x2A
	22	0-7	0xA6	0x29
	23	0-7	0xA7	0x28
	24	0-7	0xA9	0x27
	25	0-7	0xAA	0x26
	26	0-7	0xAB	0x25
	27	0-7	0xAC	0x24
	28	0-7	0xAD	0x23
	29	0-7	0xAE	0x22
	30	0-7	0xB1	0x21
	31	0-7	NONE	NONE

Table A-1 SCSI TID map (ScanDown=0) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
4	0	0-7	0xB2	0x20
	1	0-7	0xB3	0x1F
	2	0-7	0xB4	0x1E
	3	0-7	0xB5	0x1D
	4	0-7	0xB6	0x1C
	5	0-7	0xB9	0x1B
	6	0-7	0xBA	0x1A
	7	0-7	0xBC	0x19
	8	0-7	0xC3	0x18
	9	0-7	0xC5	0x17
	10	0-7	0xC6	0x16
	11	0-7	0xC7	0x15
	12	0-7	0xC9	0x14
	13	0-7	0xCA	0x13
	14	0-7	0xCB	0x12
	15	0-7	0xCC	0x11
	16	0-7	0xCD	0x10
	17	0-7	0xCE	0x0F
	18	0-7	0xD1	0x0E
	19	0-7	0xD2	0x0D
	20	0-7	0xD3	0x0C
	21	0-7	0xD4	0x0B
	22	0-7	0xD5	0x0A
	23	0-7	0xD6	0x09
	24	0-7	0xD9	0x08
	25	0-7	0xDA	0x07
	26	0-7	0xDC	0x06
	27	0-7	0xE0	0x05
	28	0-7	0xE1	0x04
	29	0-7	0xE2	0x03
	30	0-7	0xE4	0x02
	31	0-7	NONE	NONE

Table A-1 SCSI TID map (ScanDown=0) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
5	0	0-7	0xE8	0x01
	1	0-7	0xEF	0x00
	2	0-7	NONE	NONE
	3	0-7	NONE	NONE
	4	0-7	NONE	NONE
	5	0-7	NONE	NONE
	6	0-7	NONE	NONE
	7	0-7	NONE	NONE
	8	0-7	NONE	NONE
	9	0-7	NONE	NONE
	10	0-7	NONE	NONE
	11	0-7	NONE	NONE
	12	0-7	NONE	NONE
	13	0-7	NONE	NONE
	14	0-7	NONE	NONE
	15	0-7	NONE	NONE
	16	0-7	NONE	NONE
	17	0-7	NONE	NONE
	18	0-7	NONE	NONE
	19	0-7	NONE	NONE
	20	0-7	NONE	NONE
	21	0-7	NONE	NONE
	22	0-7	NONE	NONE
	23	0-7	NONE	NONE
	24	0-7	NONE	NONE
	25	0-7	NONE	NONE
	26	0-7	NONE	NONE
	27	0-7	NONE	NONE
	28	0-7	NONE	NONE
	29	0-7	NONE	NONE
	30	0-7	NONE	NONE
	31	0-7	NONE	NONE

Table A-2 SCSI TID map (ScanDown=1)

Bus #	TID	LUN	AL_PA	SEL_ID
0	0-31	0-7	NONE	NONE
1	0	0-7	0xEF	0x00
	1	0-7	0xE8	0x01
	2	0-7	0xE4	0x02
	3	0-7	0xE2	0x03
	4	0-7	0xE1	0x04
	5	0-7	0xE0	0x05
	6	0-7	0xDC	0x06
	7	0-7	0xDA	0x07
	8	0-7	0xD9	0x08
	9	0-7	0xD6	0x09
	10	0-7	0xD5	0x0A
	11	0-7	0xD4	0x0B
	12	0-7	0xD3	0x0C
	13	0-7	0xD2	0x0D
	14	0-7	0xD1	0x0E
	15	0-7	0xCE	0x0F
	16	0-7	0xCD	0x10
	17	0-7	0xCC	0x11
	18	0-7	0xCB	0x12
	19	0-7	0xCA	0x13
	20	0-7	0xC9	0x14
	21	0-7	0xC7	0x15
	22	0-7	0xC6	0x16
	23	0-7	0xC5	0x17
	24	0-7	0xC3	0x18
	25	0-7	0xBC	0x19
	26	0-7	0xBA	0x1A
	27	0-7	0xB9	0x1B
	28	0-7	0xB6	0x1C
	29	0-7	0xB5	0x1D
	30	0-7	0xB4	0x1E
	31	0-7	NONE	NONE

Table A-2 SCSI TID map (ScanDown=1) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
2	0	0-7	0xB3	0x1F
	1	0-7	0xB2	0x20
	2	0-7	0xB1	0x21
	3	0-7	0xAE	0x22
	4	0-7	0xAD	0x23
	5	0-7	0xAC	0x24
	6	0-7	0xAB	0x25
	7	0-7	0xAA	0x26
	8	0-7	0xA9	0x27
	9	0-7	0xA7	0x28
	10	0-7	0xA6	0x29
	11	0-7	0xA5	0x2A
	12	0-7	0xA3	0x2B
	13	0-7	0x9F	0x2C
	14	0-7	0x9E	0x2D
	15	0-7	0x9D	0x2E
	16	0-7	0x9B	0x2F
	17	0-7	0x98	0x30
	18	0-7	0x97	0x31
	19	0-7	0x90	0x32
	20	0-7	0x8F	0x33
	21	0-7	0x88	0x34
	22	0-7	0x84	0x35
	23	0-7	0x82	0x36
	24	0-7	0x81	0x37
	25	0-7	0x80	0x38
	26	0-7	0x7C	0x39
	27	0-7	0x7A	0x3A
	28	0-7	0x79	0x3B
	29	0-7	0x76	0x3C
	30	0-7	0x75	0x3D
	31	0-7	NONE	NONE

Table A-2 SCSI TID map (ScanDown=1) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
3	0	0-7	0x74	0x3E
	1	0-7	0x73	0x3F
	2	0-7	0x72	0x40
	3	0-7	0x71	0x41
	4	0-7	0x6E	0x42
	5	0-7	0x6D	0x43
	6	0-7	0x6C	0x44
	7	0-7	0x6B	0x45
	8	0-7	0x6A	0x46
	9	0-7	0x69	0x47
	10	0-7	0x67	0x48
	11	0-7	0x66	0x49
	12	0-7	0x65	0x4A
	13	0-7	0x63	0x4B
	14	0-7	0x5C	0x4C
	15	0-7	0x5A	0x4D
	16	0-7	0x59	0x4E
	17	0-7	0x56	0x4F
	18	0-7	0x55	0x50
	19	0-7	0x54	0x51
	20	0-7	0x53	0x52
	21	0-7	0x52	0x53
	22	0-7	0x51	0x54
	23	0-7	0x4E	0x55
	24	0-7	0x4D	0x56
	25	0-7	0x4C	0x57
	26	0-7	0x4B	0x58
	27	0-7	0x4A	0x59
	28	0-7	0x49	0x5A
	29	0-7	0x47	0x5B
	30	0-7	0x46	0x5C
	31	0-7	NONE	NONE

Table A-2 SCSI TID map (ScanDown=1) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
4	0	0-7	0x45	0x5D
	1	0-7	0x43	0x5E
	2	0-7	0x3C	0x5F
	3	0-7	0x3A	0x60
	4	0-7	0x39	0x61
	5	0-7	0x36	0x62
	6	0-7	0x35	0x63
	7	0-7	0x34	0x64
	8	0-7	0x33	0x65
	9	0-7	0x32	0x66
	10	0-7	0x31	0x67
	11	0-7	0x2E	0x68
	12	0-7	0x2D	0x69
	13	0-7	0x2C	0x6A
	14	0-7	0x2B	0x6B
	15	0-7	0x2A	0x6C
	16	0-7	0x29	0x6D
	17	0-7	0x27	0x6E
	18	0-7	0x26	0x6F
	19	0-7	0x25	0x70
	20	0-7	0x23	0x71
	21	0-7	0x1F	0x72
	22	0-7	0x1E	0x73
	23	0-7	0x1D	0x74
	24	0-7	0x1B	0x75
	25	0-7	0x18	0x76
	26	0-7	0x17	0x77
	27	0-7	0x10	0x78
	28	0-7	0x0F	0x79
	29	0-7	0x08	0x7A
	30	0-7	0x04	0x7B
	31	0-7	NONE	NONE

Table A-2 SCSI TID map (ScanDown=1) (continued)

Bus #	TID	LUN	AL_PA	SEL_ID
5	0	0-7	0x02	0x7C
	1	0-7	0x01	0x7D
	2	0-7	NONE	NONE
	3	0-7	NONE	NONE
	4	0-7	NONE	NONE
	5	0-7	NONE	NONE
	6	0-7	NONE	NONE
	7	0-7	NONE	NONE
	8	0-7	NONE	NONE
	9	0-7	NONE	NONE
	10	0-7	NONE	NONE
	11	0-7	NONE	NONE
	12	0-7	NONE	NONE
	13	0-7	NONE	NONE
	14	0-7	NONE	NONE
	15	0-7	NONE	NONE
	16	0-7	NONE	NONE
	17	0-7	NONE	NONE
	18	0-7	NONE	NONE
	19	0-7	NONE	NONE
	20	0-7	NONE	NONE
	21	0-7	NONE	NONE
	22	0-7	NONE	NONE
	23	0-7	NONE	NONE
	24	0-7	NONE	NONE
	25	0-7	NONE	NONE
	26	0-7	NONE	NONE
	27	0-7	NONE	NONE
	28	0-7	NONE	NONE
	29	0-7	NONE	NONE
	30	0-7	NONE	NONE
	31	0-7	NONE	NONE

B

Note on using Veritas Cluster Server

By issuing a SCSI-3 Persistent Reserve command for a Hitachi RAID storage system, Veritas Cluster Server (VCS) provides the I/O fencing function that can prevent data corruption from occurring if the cluster communication stops. Each node of VCS registers reserve keys to the storage system, which enables these nodes to share a disk to which the reserve key is registered.

Each node of VCS registers the reserve key when importing a disk groups. One node registers the identical reserve key for all paths of all disks (LU) in the disk group. The reserve key contains a unique value for each disk group and a value to distinguish nodes.

Key format: `<Node # + disk group-unique information>`

Example: `APGR0000`, `APGR0001`, `BPGR0000`, and so on

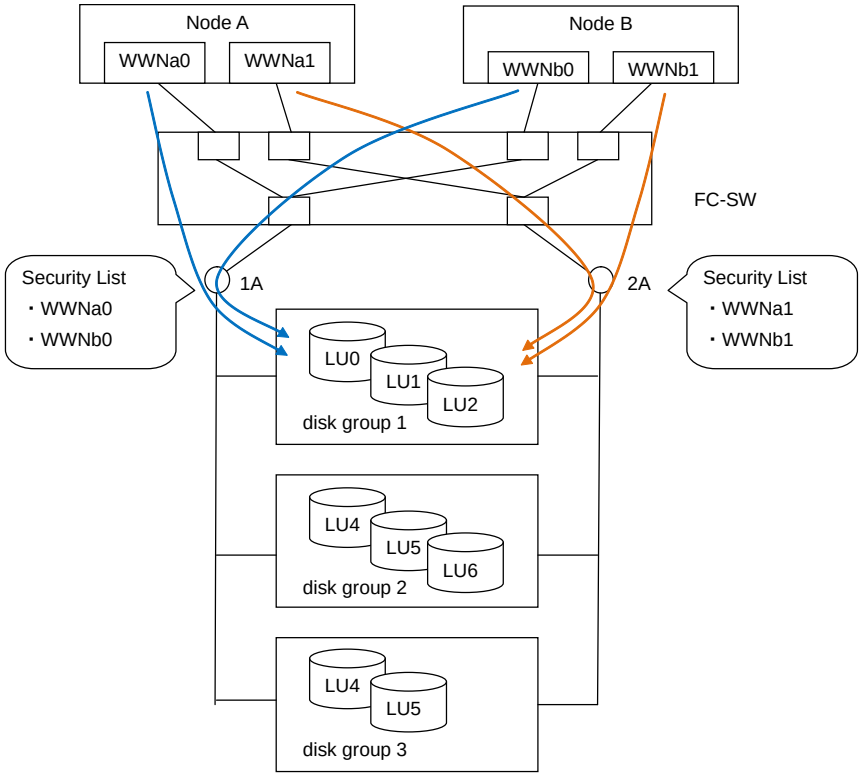
When the Hitachi RAID storage system receives a request to register the reserve key, the reserve key and Port WWN of node are recorded on a key registration table of each port of storage system where the registration request is received. The number of reserve keys that can be registered to a storage system is 2,048 per port on VSP G1000, VSP G1500, VSP F1500, VSP Gx00, and VSP Fx00. For VSP and HUS VM, you must set HMO 61 ON to increase the maximum number of reserve keys per port from 128 to 2,048. For USP V/VM the maximum number of reserve keys per port is 128. The storage system confirms duplication of registration by a combination of the node Port WWN and reserve key. Therefore, the number of entries of the registration table does not increase even though any request for registering duplicated reserve keys is accepted.

Calculation formula for the number of used entries of key registration table:

`(number of nodes) × (number of Port WWN of node) × (number of disk groups)`

When the number of registered reserve keys exceeds the upper limit of 2,048, key registration as well as operations such as installing an LU to the disk group fails. To avoid failure of reserve key registration, the number of reserve keys needs to be kept below 2,048. For this, restrictions such as imposing a limit on the number of nodes or on the number of server ports using LUN security function or maintaining the number of disk groups appropriate are necessary.

Example: When adding an LU to increase disk capacity, do not add the number of disk groups, but add an LU to the current disk group.



Key registration table for Port-1A		
Entry	**Reserve Key**	**WWN**
0	APGR0001	WWNa0
1	APGR0002	WWNa0
2	APGR0003	WWNa0
3	BPGR0001	WWNb0
4	BPGR0002	WWNb0
5	BPGR0003	WWNb0
6	-	-
:	:	:
127	-	-

Key registration table for Port-2A		
Entry	**Reserve Key**	**WWN**
0	APGR0001	WWNa1
1	APGR0002	WWNa1
2	APGR0003	WWNa1
3	BPGR0001	WWNb1
4	BPGR0002	WWNb1
5	BPGR0003	WWNb1
6	-	-
:	:	:
127	-	-

Figure B-1 Adding Reserve Keys for LUs to Increase Disk Capacity

For AIX® systems: The persistent reservation of a logical unit (LU) may not be canceled due to some reason when multiple hosts share a volume group rather than making up a cluster configuration.

C

Disk parameters for Hitachi disk types

The following tables list the disk parameters for the Hitachi SCSI disk devices. For information about configuring devices other than OPEN-V, contact your Hitachi Data Systems representative.

- ☐ Parameter values for OPEN-x disk types
- ☐ Parameter values for VLL disk types
- ☐ Parameter values for LUSE disk types
- ☐ Parameter values for VLL LUSE disk types
- ☐ Parameter values for OPEN-8 disk types

Parameter values for OPEN-*x* disk types

Parameter		Disk Type			
		OPEN-3	OPEN-9	OPEN-E	OPEN-L
Ty	Disk category	winchester	winchester	winchester	winchester
Dt	Control type	SCSI	SCSI	SCSI	SCSI
Ns	sectors/tracks	96	96	96	96
nt	tracks/cylinder	15	15	15	15
nc	Number of all cylinders	3338	10016	19759	19759
rm	Number of rotations of the disk	6300	6300	6300	6300
oa	a partition offset (Starting block in a partition)	Set optionally	Set optionally	Set optionally	Set optionally
ob	b partition offset (Starting block in b partition)	Set optionally	Set optionally	Set optionally	Set optionally
oc	c partition offset (Starting block in c partition)	0	0	0	0
od	d partition offset (Starting block in d partition)	Set optionally	Set optionally	Set optionally	Set optionally
oe	e partition offset (Starting block in e partition)	Set optionally	Set optionally	Set optionally	Set optionally
of	f partition offset (Starting block in f partition)	Set optionally	Set optionally	Set optionally	Set optionally

Parameter		Disk Type			
		OPEN-3	OPEN-9	OPEN-E	OPEN-L
og	g partition offset (Starting block in g partition)	Set optionally	Set optionally	Set optionally	Set optionally
oh	h partition offset (Starting block in h partition)	Set optionally	Set optionally	Set optionally	Set optionally
pa	a partition size	Set optionally	Set optionally	Set optionally	Set optionally
pb	b partition size	Set optionally	Set optionally	Set optionally	Set optionally
pc	c partition size	4806720	14423040	28452960	28452960
pd	d partition size	Set optionally	Set optionally	Set optionally	Set optionally
pe	e partition size	Set optionally	Set optionally	Set optionally	Set optionally
pf	f partition size	Set optionally	Set optionally	Set optionally	Set optionally
pg	g partition size	Set optionally	Set optionally	Set optionally	Set optionally
ph	h partition size	Set optionally	Set optionally	Set optionally	Set optionally
ba	a partition block size	8192	8192	8192	8192
bb	b partition block size	8192	8192	8192	8192
bc	c partition block size	8192	8192	8192	8192
bd	d partition block size	8192	8192	8192	8192
be	e partition block size	8192	8192	8192	8192
bf	f partition block size	8192	8192	8192	8192
bg	g partition block size	8192	8192	8192	8192
bh	h partition block size	8192	8192	8192	8192
fa	a partition fragment size	1024	1024	1024	1024
fb	b partition fragment size	1024	1024	1024	1024
fc	c partition fragment size	1024	1024	1024	1024
fd	d partition fragment size	1024	1024	1024	1024
fe	e partition fragment size	1024	1024	1024	1024
ff	f partition fragment size	1024	1024	1024	1024
fg	g partition fragment size	1024	1024	1024	1024
fh	h partition fragment size	1024	1024	1024	1024

Parameter values for VLL disk types

Parameter		Disk Type		
		OPEN-3 VLL	OPEN-9 VLL	OPEN-E VLL
ty	Disk category	winchester	winchester	winchester
dt	Control type	SCSI	SCSI	SCSI
ns	sectors/tracks	96	96	96
nt	tracks/cylinder	15	15	15
nc	Number of all cylinders	Depends on CV configuration	Depends on CV configuration	Depends on CV configuration
rm	Number of rotations of the disk	6300	6300	6300
oa	a partition offset (Starting block in a partition)	Set optionally	Set optionally	Set optionally
ob	b partition offset (Starting block in b partition)	Set optionally	Set optionally	Set optionally
oc	c partition offset (Starting block in c partition)	0	0	0
od	d partition offset (Starting block in d partition)	Set optionally	Set optionally	Set optionally
oe	e partition offset (Starting block in e partition)	Set optionally	Set optionally	Set optionally
of	f partition offset (Starting block in f partition)	Set optionally	Set optionally	Set optionally
og	g partition offset (Starting block in g partition)	Set optionally	Set optionally	Set optionally
oh	h partition offset (Starting block in h partition)	Set optionally	Set optionally	Set optionally
pa	a partition size	Set optionally	Set optionally	Set optionally
pb	b partition size	Set optionally	Set optionally	Set optionally
pc	c partition size	Depends on CV configuration	Depends on CV configuration	Depends on CV configuration
pd	d partition size	Set optionally	Set optionally	Set optionally
pe	e partition size	Set optionally	Set optionally	Set optionally
pf	f partition size	Set optionally	Set optionally	Set optionally
pg	g partition size	Set optionally	Set optionally	Set optionally
ph	h partition size	Set optionally	Set optionally	Set optionally
ba	a partition block size	8192	8192	8192
bb	b partition block size	8192	8192	8192
bc	c partition block size	8192	8192	8192
bd	d partition block size	8192	8192	8192
be	e partition block size	8192	8192	8192
bf	f partition block size	8192	8192	8192
bg	g partition block size	8192	8192	8192
bh	h partition block size	8192	8192	8192
fa	a partition fragment size	1024	1024	1024
fb	b partition fragment size	1024	1024	1024
fc	c partition fragment size	1024	1024	1024
fd	d partition fragment size	1024	1024	1024
fe	e partition fragment size	1024	1024	1024
ff	f partition fragment size	1024	1024	1024
fg	g partition fragment size	1024	1024	1024
fh	h partition fragment size	1024	1024	1024

Parameter values for LUSE disk types

Parameter		Disk Type			
		OPEN-3*n (n = 2 to 36)	OPEN-9*n (n = 2 to 36)	OPEN-E*n (n = 2 to 36)	OPEN-L*n (n = 2 to 12)
ty	Disk category	winchester	winchester	winchester	winchester
dt	Control type	SCSI	SCSI	SCSI	SCSI
ns	sectors/tracks	96	96	96	96
nt	tracks/cylinder	15	15	15	15
nc	Number of all cylinders	3338*n	Depends on CV configuration	19759*n	19759*n
rm	Number of rotations of the disk	6300	6300	6300	6300
oa	a partition offset (Starting block in a partition)	Set optionally	Set optionally	Set optionally	Set optionally
ob	b partition offset (Starting block in b partition)	Set optionally	Set optionally	Set optionally	Set optionally
oc	c partition offset (Starting block in c partition)	0	0	0	0
od	d partition offset (Starting block in d partition)	Set optionally	Set optionally	Set optionally	Set optionally
oe	e partition offset (Starting block in e partition)	Set optionally	Set optionally	Set optionally	Set optionally
of	f partition offset (Starting block in f partition)	Set optionally	Set optionally	Set optionally	Set optionally
og	g partition offset (Starting block in g partition)	Set optionally	Set optionally	Set optionally	Set optionally
oh	h partition offset (Starting block in h partition)	Set optionally	Set optionally	Set optionally	Set optionally
pa	a partition size	Set optionally	Set optionally	Set optionally	Set optionally
pb	b partition size	Set optionally	Set optionally	Set optionally	Set optionally
pc	c partition size	4806720*n	Depends on CV configuration	28452960*n	28452960*n
pd	d partition size	Set optionally	Set optionally	Set optionally	Set optionally
pe	e partition size	Set optionally	Set optionally	Set optionally	Set optionally
pf	f partition size	Set optionally	Set optionally	Set optionally	Set optionally
pg	g partition size	Set optionally	Set optionally	Set optionally	Set optionally
ph	h partition size	Set optionally	Set optionally	Set optionally	Set optionally
ba	a partition block size	8192	8192	8192	8192
bb	b partition block size	8192	8192	8192	8192
bc	c partition block size	8192	8192	8192	8192
bd	d partition block size	8192	8192	8192	8192
be	e partition block size	8192	8192	8192	8192
bf	f partition block size	8192	8192	8192	8192
bg	g partition block size	8192	8192	8192	8192
bh	h partition block size	8192	8192	8192	8192
fa	a partition fragment size	1024	1024	1024	1024
fb	b partition fragment size	1024	1024	1024	1024
fc	c partition fragment size	1024	1024	1024	1024
fd	d partition fragment size	1024	1024	1024	1024
fe	e partition fragment size	1024	1024	1024	1024
ff	f partition fragment size	1024	1024	1024	1024
fg	g partition fragment size	1024	1024	1024	1024
fh	h partition fragment size	1024	1024	1024	1024

Parameter values for VLL LUSE disk types

Parameter		Disk Type		
		OPEN-3 VLL*n (n = 2 to 36)	OPEN-9 VLL*n (n = 2 to 36)	OPEN-E VLL*n (n = 2 to 36)
ty	winchester	winchester	winchester	winchester
dt	SCSI	SCSI	SCSI	SCSI
ns	96	96	96	116
nt	15	15	15	15
nc	Depends on CV configuration[3]	19759	10016*n	Depends on CV configuration
rm	6300	6300	6300	6300
oa	Set optionally	Set optionally	Set optionally	Set optionally
ob	Set optionally	Set optionally	Set optionally	Set optionally
oc	0	0	0	0
od	Set optionally	Set optionally	Set optionally	Set optionally
oe	Set optionally	Set optionally	Set optionally	Set optionally
of	Set optionally	Set optionally	Set optionally	Set optionally
og	Set optionally	Set optionally	Set optionally	Set optionally
oh	Set optionally	Set optionally	Set optionally	Set optionally
pa	Set optionally[2]	Set optionally	Set optionally	Set optionally
pb	Set optionally	Set optionally	Set optionally	Set optionally
pc	Depends on CV configuration[3]	28452960	14423040*n	Depends on CV configuration
pd	Set optionally	Set optionally	Set optionally	Set optionally
pe	Set optionally	Set optionally	Set optionally	Set optionally
pf	Set optionally	Set optionally	Set optionally	Set optionally
pg	Set optionally	Set optionally	Set optionally	Set optionally
ph	Set optionally	Set optionally	Set optionally	Set optionally
ba	8192	8192	8192	8192
bb	8192	8192	8192	8192
bc	8192	8192	8192	8192
bd	8192	8192	8192	8192
be	8192	8192	8192	8192
bf	8192	8192	8192	8192
bg	8192	8192	8192	8192
bh	8192	8192	8192	8192
fa	1024	1024	1024	1024
fb	1024	1024	1024	1024
fc	1024	1024	1024	1024
fd	1024	1024	1024	1024
fe	1024	1024	1024	1024
ff	1024	1024	1024	1024
fg	1024	1024	1024	1024
fh	1024	1024	1024	1024

Parameter values for OPEN-8 disk types

Parameter		Disk Type			
		OPEN-8	OPEN-8*n (n = 2 to 36)	OPEN-8 VIR	OPEN-8*n VIR (n = 2 to 36)
ty	Disk category	winchester	winchester	winchester	winchester
dt	Control type	SCSI	SCSI	SCSI	SCSI
ns	sectors/tracks	96	96	96	116
nt	tracks/cylinder	15	15	15	15
nc	Number of all cylinders	9966	9966*n	Depends on CV configuration	Depends on CV configuration
rm	Number of rotations of the disk	6300	6300	6300	6300
oa	a partition offset (Starting block in a partition)	Set optionally	Set optionally	Set optionally	Set optionally
ob	b partition offset (Starting block in b partition)	Set optionally	Set optionally	Set optionally	Set optionally
oc	c partition offset (Starting block in c partition)	0	0	0	0
od	d partition offset (Starting block in d partition)	Set optionally	Set optionally	Set optionally	Set optionally
oe	e partition offset (Starting block in e partition)	Set optionally	Set optionally	Set optionally	Set optionally
of	f partition offset (Starting block in f partition)	Set optionally	Set optionally	Set optionally	Set optionally
og	g partition offset (Starting block in g partition)	Set optionally	Set optionally	Set optionally	Set optionally
oh	h partition offset (Starting block in h partition)	Set optionally	Set optionally	Set optionally	Set optionally
pa	a partition size	Set optionally	Set optionally	Set optionally	Set optionally
pb	b partition size	Set optionally	Set optionally	Set optionally	Set optionally
pc	c partition size	14351040	14351040*n	Depends on CV configuration	Depends on CV configuration
pd	d partition size	Set optionally	Set optionally	Set optionally	Set optionally
pe	e partition size	Set optionally	Set optionally	Set optionally	Set optionally
pf	f partition size	Set optionally	Set optionally	Set optionally	Set optionally
pg	g partition size	Set optionally	Set optionally	Set optionally	Set optionally
ph	h partition size	Set optionally	Set optionally	Set optionally	Set optionally
ba	a partition block size	8192	8192	8192	8192
bb	b partition block size	8192	8192	8192	8192
bc	c partition block size	8192	8192	8192	8192
bd	d partition block size	8192	8192	8192	8192
be	e partition block size	8192	8192	8192	8192
bf	f partition block size	8192	8192	8192	8192
bg	g partition block size	8192	8192	8192	8192
bh	h partition block size	8192	8192	8192	8192
fa	a partition fragment size	1024	1024	1024	1024
fb	b partition fragment size	1024	1024	1024	1024
fc	c partition fragment size	1024	1024	1024	1024
fd	d partition fragment size	1024	1024	1024	1024
fe	e partition fragment size	1024	1024	1024	1024
ff	f partition fragment size	1024	1024	1024	1024
fg	g partition fragment size	1024	1024	1024	1024
fh	h partition fragment size	1024	1024	1024	1024

Host modes and host mode options

This appendix lists the host modes and host mode options (HMOs) for the Hitachi storage systems. Refer to the section below for your storage system model, as the host modes and HMOs are different for each storage system (for example, new HMOs 80, 81, 82, and 83 for VSP Gx00 and Fx00).

- Host modes and host mode options for USP V/VM
- Host modes and host mode options for VSP
- Host modes and host mode options for VSP G1000, G1500 and VSP F1500
- Host modes and host mode options for HUS VM
- Host modes and host mode options for VSP Gx00 and Fx00

Host modes and host mode options for USP V/VM

Table D-1 Host Modes for USP V/VM

Host mode	When to select this mode
00 Standard	When registering Red Hat Linux server hosts or IRIX server hosts in the host group.
01 VMware	When registering VMware server hosts in the host group (see ***Notes***).
03 HP	When registering HP-UX server hosts in the host group.
05 OpenVMS	When registering OpenVMS server hosts in the host group.
07 Tru64	When registering Tru64 server hosts in the host group.
09 Solaris	When registering Solaris server hosts in the host group.
0A NetWare	When registering NetWare server hosts in the host group.
0C Windows	When registering Windows server hosts in the host group (see ***Notes***).
0F AIX	When registering AIX server hosts in the host group
21 VMware Extension	When registering VMware server hosts in the host group (see ***Notes***).
2C Windows Extension	When registering Windows server hosts in the host group (see ***Notes***).

Host mode	When to select this mode
4C UVM	When registering another USP V/VM storage system in the host group for mapping by using Universal Volume Manager. If this mode is used when the USP V/VM is being used as external storage of another USP V/VM storage system, the data of the MF-VOL in the USP V/VM storage system can be transferred. Refer to emulation types below for the MF-VOL. The data of the MF-VOL cannot be transferred when the storage systems are connected with the host mode other than "4C UVM", and a message requiring formatting appears after the mapping. In this case, cancel the message requiring formatting, and set the host mode to "4C UVM" when you want to transfer data. The following device types can be transferred: 3390-3A, 3380-KA, 3380-3A, 3390-9A, 3390-LA.

Notes:

- If Windows server hosts are registered in a host group, ensure that the host mode of the host group is 0C Windows or 2C Windows Extension.

  If the host mode of a host group is 0C Windows and an LU path is defined between the host group and a logical volume, the logical volume cannot be combined with other logical volumes to form a LUSE volume (that is, an expanded LU).

  If the host mode of a host group is 2C Windows Extension and an LU path is defined between the host group and a logical volume, the logical volume can be combined with other logical volumes to form a LUSE volume (that is, an expanded LU). If you plan to expand LUs by using LUSE in the future, set the host mode 2C Windows Extension. For detailed information about LUSE, see the *LUN Expansion User's Guide*.
- If VMware server hosts are registered in a host group, ensure that the host mode of the host group is 01 VMware or 21 VMware Extension.

  If the host mode of a host group is 01 VMware and an LU path is defined between the host group and a logical volume, the logical volume cannot be combined with other logical volumes to form a LUSE volume (that is, an expanded LU).

  If the host mode of a host group is 21 VMware Extension and an LU path is defined between the host group and a logical volume, the logical volume can be combined with other logical volumes to form a LUSE volume (that is, an expanded LU). If you plan to expand LUs by using LUSE in the future, set the host mode 21 VMware Extension. For detailed information about LUSE, see the *LUN Expansion User's Guide*.
- If you plan to expand LUs by using LUSE in case of Windows virtual host on VMware recognizing LU by Raw Device Mapping (RDM) method, set the host mode 2C Windows Extension. If the host mode 2C Windows Extension is not set, change the host mode to 2C. Before changing the host mode, back up the LUSE volume. After changing the mode, restore the LUSE volume. For detailed information about LUSE, see the *LUN Expansion User's Guide*.
- Besides the host modes mentioned above, the **Host Mode** list displays the **Reserve** host modes. Please do not select any **Reserve** host mode without assistance from technical support.

Table D-2 Host Mode Options for USP V/VM

No.	Function	Description
2	VERITAS DBC+RAC	When VERITAS Database Edition/Advanced Cluster for Real Application Clusters is used. When VERITAS Cluster Server 4.0 or later (I/O fencing function) is used. When Oracle RAC Cluster Ready Services is used. Anything using I/O fencing.
6	TPRLO (Third-party process layout)	Use when all the following conditions are satisfied: • The host mode 0C Windows or 2C Windows Extension is used • The Emulex host bus adapter is used • The mini-port driver is used • TPRLO=2 is specified for the mini-port driver parameter of the host bus adapter

No.	Function	Description
7	Automatic recognition function of LUN	Use when all the following conditions are satisfied: • The host mode **00 Standard** or **09 Solaris** is used. • SUN StorEdge SAN Foundation Software Version 4.2 or later is used • You want to automate recognition of increase and decrease of devices when genuine SUN HBA is connected.
12	No display for ghost LUN	Use when all the following conditions are satisfied: • The host mode **03 HP** is used. • You want to suppress creation of device files for devices to which paths are not defined.
13	SIM report at link failure	Use when you want to be informed by SIM (service information message) that the number of link failures detected between ports exceeds the threshold. ***Caution:*** Configure this HMO only when requested to do so.
14	HP TruCluster with TrueCopy function	Use when all the following conditions are satisfied: • The host mode **07 Tru64** is used. • You want to use TruCluster to set a cluster to each of P-VOL and S-VOL for TrueCopy or Universal Replicator.
15	HACMP	Use when all the following conditions are satisfied: • The host mode **0F AIX** is used. • HACMP 5.1 Version 5.1.0.4 or later, HACMP4.5 Version 4.5.0.13 or later, or HACMP5.2 or later is used.
22	Veritas Cluster Server	Use when all the following conditions are satisfied: • The host mode **0F AIX** is used. • Veritas Cluster Server is used. **Note:** Before setting HMO 22, ask your Hitachi Data Systems representative for assistance.
33	Set/Report Device Identifier enable	Use when all the following conditions are satisfied: • Host mode **03 HP** or **05 OpenVMS** is used. Set the UUID when you set HMO 33 and host mode 05 OpenVMS is used. • You want to enable commands to assign a nickname of the device. • You want to set UUID to identify a logical volume from the host.
39	A target reset	Resets a job and returns UA to all initiators connected to the host group where Target Reset has occurred. **ON**: • Job reset range: Performs a reset to the jobs of all he initiators connected to the host group where Target Reset has occurred. • UA set range: Returns UA to all the initiators connected to the host group where Target Reset has occurred. **OFF** (default): • Job reset range: Performs a reset to the jobs of the initiator that has issued Target Reset. • UA set range: Returns UA to the initiator that has issued Target Reset. ***Note:*** This HMO is used in the SVC environment, and the job reset range and UA set range must be controlled per host group when Target Reset has been received.
40	DP-Vol expansion	Notifies the host OS through SCSI protocol that DP-VOL capacity has been expanded. The host operating system must accept this notification and adjust to the increase in DP-VOL capacity. If the host operating system is one that does not react to the notification by automatically adjusting to the capacity change, then the host must be manually commanded to recognize the change.

No.	Function	Description
41	Prioritized device recognition command	Gives priority to starting Inquiry/Report LUN issued from the host where this HMO is set. **ON:** Inquiry/ Report LUN is started by priority. **OFF (default):** The operation is the same as before.
42	Prevent "OHUB PCI retry"	When CHA PCI is accessed from MP, the behavior when the status is busy differs depending on the mode status as follows. **ON:** The PCI retry is not returned, and the PCI bus is occupied. **OFF (default):** The PCI retry is returned. ***Note:*** When IBM Z10 Linux is connected, set this mode to ON. In other cases, set the mode to OFF.
43	Queue Full Response	When Queue Full occurs, this HMO is used to return Queue Full to the host. **ON:** When Queue Full occurs, Queue Full is always returned to the host. **OFF (default):** When Queue Full occurs with Host Mode HP-UX, Busy is returned to the host. ***Note:*** Set this HMO to ON when HP-UX 11.x or later is connected. However, if the setting of queue depth on the host is made based on the configuration guide, the mode setting is not necessary since Queue Full/ Busy will not occur.
48	HAM S-VOL Read	By setting this HMO to ON, in normal operation, the pair status of S-VOL is not changed to SSWS even when Read commands exceeding the threshold (1,000/6 min) are issued while a specific application is used. **ON:** The pair status of S-VOL is not changed to SSWS if Read commands exceeding the threshold are issued. **OFF (default):** The pair status of S-VOL is changed to SSWS if Read commands exceeding the threshold are issued. ***Notes:*** 1. Set this HMO to ON for the host group if the transition of the pair status to SSWS is not desired in the case that an application, which issues Read commands (*1) exceeding the threshold (1,000/6 min) to S-VOL, is used in HAM environment. (*1: Currently, the vxdisksetup command of Solaris VxVM serves.) 2. Even when a failure occurs in P-VOL, if this option is set to ON, which means that the pair status of S-VOL is not changed to SSWS (*2), the response time of Read command to the S-VOL whose pair status remains as Pair takes several msecs. On the other hand, if the option is set to OFF, the response time of Read command to the S-VOL is recovered to be equal to that to P-VOL by judging that an error occurs in the P-VOL when Read commands exceeding the threshold are issued. (*2: Until the S-VOL receives a Write command, the pair status of S-VOL is not changed to SSWS.)

No.	Function	Description
49	BB Credit Set Up Option 1	Set this HMO when you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 50. This HMO determines the BB_Credit value. (HMO#49: Low_bit). **ON:** The storage system operates with BB_Credit value of 80 or 255. ***Caution:*** Set this HMO to ON only for the 8US package. **OFF (default):** The storage system operates with BB_Credit value of 40 or 128. HMOs 50/49: BB_Credit value is determined by 2 bits of the HMOs: 00: Existing mode (BB_Credit value = 40) 01: BB_Credit value = 80 10: BB_Credit value = 128 11: BB_Credit value = 255 ***Notes:*** 1. Apply this HMO when the following two conditions are met: • Data frame transfer in long distance connection exceeds the BB_Credit value. • System option mode (SOM) 769 is set to OFF (retry operation is enabled at TC/UR path creation). 2. When HMO 49 is set to ON, SSB log of link down is output on the MCU (M-DKC). 3. This HMO functions only when both the MCU (M-DKC) and RCU (R-DKC) have the microcode that supports this function. 4. This HMO is applied only to Initiator-Port. This function is applicable only when the 8US PCB is used on the MCU/RCU. 5. If this HMO is used, FC point-to-point setting is required. 6. If you need to remove the 8US PCB, set HMO 49 to OFF first, and then remove the PCB. 7. If HMO 49 is set to ON while SOM 769 is ON, path creation may fail after automatic port switching. 8. Make sure to set HMO 49 to ON or OFF after the pair is suspended or when the load is light. 9. The RCU Target that is connected to the MCU on which HMO 49 is ON cannot be used for UR. 10. This function is intended for use in long-distance data transfer. If HMO 49 is set to ON with distance of 0 km, data transfer errors may occur on RCU side.

No.	Function	Description
50	BB Credit Set Up Option 2	Set this HMO when you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 49. This HMO determines the BB_Credit value. (HMO#50: High_bit). **ON:** The storage system operates with BB_Credit value of 128 or 255. ***Caution:*** Set this HMO to ON only for the 8US package. **OFF (default):** The storage system operates with BB_Credit value of 40 or 80. HMOs 50/49: BB_Credit value is dectermined by 2 bits of the HMOs: 00: Existing mode (BB_Credit value = 40) 01: BB_Credit value = 80 10: BB_Credit value = 128 11: BB_Credit value = 255 ***Notes:*** 1. Apply this HMO when the following two conditions are met: • Data frame transfer in long distance connection exceeds the BB_Credit value. • System option mode (SOM) 769 is set to OFF (retry operation is enabled at TC/UR path creation). 2. When HMO 50 is set to ON, SSB log of link down is output on the MCU (M-DKC). 3. This HMO functions only when both the MCU and RCU have the microcode that supports this function. 4. The HMO setting is only applied to Initiator-Port. This function is only applicable when the 8US PCB is used on RCU/MCU. 5. If this HMO is used, Point-to-Point setting is necessary. 6. When removing 8US PCB, the operation must be executed after setting this HMO to OFF. 7. If this HMO is set to ON while SOM 769 is ON, path creation may fail after automatic port switching. 8. Make sure to set this HMO from OFF to ON or from ON to OFF after the pair is suspended or when the load is low. 9. The RCU Target that is connected to the MCU on which this HMO is ON cannot be used for UR. 10. This function is intended for use in long-distance data transfer. If this HMO is set to ON with distance of 0 km, data transfer errors may occur on RCU side.

No.	Function	Description
51	Round Trip Set Up Option	Set this HMO if you want to adjust the response time of the host I/O, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. This HMO selects the operation condition of TrueCopy. **ON:** TrueCopy operates in the performance improvement logic. When a WRITE command is issued, FCP_CMD/FCP_DATA is continuously issued while XFER_RDY issued from RCU side is prevented. ***Caution:*** Set this HMO to ON only for the 8US package. **OFF (default):** TrueCopy operates in the existing logic. ***Notes:*** 1. This HMO is applied when the following two conditions are met: • Data frame transfer in long distance connection exceeds the BB_Credit value. • System option mode (SOM) 769 is set to OFF (retry operation is enabled at TC/UR path creation). 2. When this HMO is set to ON, SSB log of link down is output on the MCU (M-DKC). 3. This HMO functions only when both the MCU and RCU have the microcode that supports this function. 4. The HMO setting is only applied to Initiator-Port. This function is only applicable when the 8US PCB is used on RCU/MCU. 5. If this option is used, Point-to-Point setting is necessary. 6. When removing 8US PCB, the operation must be executed after setting this HMO to OFF. 7. If this HMO is set to ON while SOM 769 is ON, path creation may fail after automatic port switching. 8. Make sure to set this HMO from OFF to ON or from ON to OFF after the pair is suspended or when the load is low. 9. When this HMO is set to ON using USP V/VM as the MCU and VSP as the RCU, the USP V/VM microode must be 60-07-63-00/00 or later (within 60-07-6*x* range) or 60-08-06-00/00 or later. 10. Path attribute change (Initiator Port → RCU-Target Port, RCU-Target Port → Initiator Port) together with Hyperswap is enabled after HMO 51 is set to ON. If HMO 51 is already set to ON on both paths, HMO 51 continues to be applied on the paths even after execution of Hyperswap.

No.	Function	Description
54	Enable XCOPY command on VMWare ESX server	Enables the XCOPY command. **ON**: The XCOPY command can be used. **OFF** (default): When the XCOPY command is received, Check Condition is returned as an unsupported command (0x05/0x2000). Also used in combination with system option mode (SOM) 808 to set the ANSI version of Standard Inquiry (microcode 60-08-07 or later): • HMO 54: ON SOM 808: ON **4** is returned as the ANSI version of Standard Inquiry. • HMO 54: ON SOM 808: OFF **2** is returned as the ANSI version of Standard Inquiry. • HMO 54: OFF SOM 808: ON or OFF **2** is returned as the ANSI version of Standard Inquiry. ***Notes:*** 1. Set this HMO to ON only when VMWare ESXi (version 5.0 or later) is connected and the VAAI function is used. 2. If this HMO is not applied, the VMWare support function, Cloning file blocks, cannot be used. 3. When the Block Zero function is used in the ESXi 5 environment with RAID600 (60-08-07/00 and later), make sure to set HMO 54 and SOM 808 to ON.
57	Conversion of sense code/key	Converts the sense code/key that is returned when an S-VOL is accessed. Apply this HMO when the sense code/key response needs to be converted when an old data volume of an HAM pair is accessed. **ON:** Sense code/key 05/2500 (LDEV blockage) converted from 0b/c0000 is returned when SSB=B8A0 is output. **OFF (default):** Sense code/key 0b/c0000 is returned when SSB=B8A0 is output.

Host modes and host mode options for VSP

Table D-3 Host Modes for VSP

Host mode	When to select this mode
00 Standard	When registering Red Hat Linux server hosts or IRIX server hosts in the host group.
01 VMware	When registering VMware server hosts in the host group (see ***Notes***).
03 HP	When registering HP-UX server hosts in the host group.
05 OpenVMS	When registering OpenVMS server hosts in the host group.
07 Tru64	When registering Tru64 server hosts in the host group.
09 Solaris	When registering Solaris server hosts in the host group.
0A NetWare	When registering NetWare server hosts in the host group.
0C Windows	When registering Windows server hosts in the host group (see ***Notes***).
0F AIX	When registering AIX server hosts in the host group
21 VMware Extension	When registering VMware server hosts in the host group (see ***Notes***).

Host mode	When to select this mode
2C Windows Extension	When registering Windows server hosts in the host group (see ***Notes***).
4C UVM	When registering another VSP storage system in the host group for mapping by using Universal Volume Manager. If this mode is used when the VSP is being used as external storage of another VSP storage system, the data of the MF-VOL in the VSP storage system can be transferred. Refer to emulation types below for the MF-VOL. The data of the MF-VOL cannot be transferred when the storage systems are connected with the host mode other than "4C UVM", and a message requiring formatting appears after the mapping. In this case, cancel the message requiring formatting, and set the host mode to "4C UVM" when you want to transfer data. The following device types can be transferred: 3390-3A, 3380-KA, 3380-3A, 3390-9A, 3390-LA.

Notes:

- If Windows server hosts are registered in a host group, ensure that the host mode of the host group is 0C Windows or 2C Windows Extension.

  If the host mode of a host group is 0C Windows and an LU path is defined between the host group and a logical volume, the logical volume cannot be combined with other logical volumes to form a LUSE volume (that is, an expanded LU).

  If the host mode of a host group is 2C Windows Extension and an LU path is defined between the host group and a logical volume, the logical volume can be combined with other logical volumes to form a LUSE volume (that is, an expanded LU). If you plan to expand LUs by using LUSE in the future, set the host mode 2C Windows Extension.
- If VMware server hosts are registered in a host group, ensure that the host mode of the host group is 01 VMware or 21 VMware Extension.

  If the host mode of a host group is 01 VMware and an LU path is defined between the host group and a logical volume, the logical volume cannot be combined with other logical volumes to form a LUSE volume (that is, an expanded LU).

  If the host mode of a host group is 21 VMware Extension and an LU path is defined between the host group and a logical volume, the logical volume can be combined with other logical volumes to form a LUSE volume (that is, an expanded LU). If you plan to expand LUs by using LUSE in the future, set the host mode 21 VMware Extension.
- If you plan to expand LUs by using LUSE in case of Windows virtual host on VMware recognizing LU by Raw Device Mapping (RDM) method, set the host mode 2C Windows Extension. If the host mode 2C Windows Extension is not set, change the host mode to 2C. Before changing the host mode, back up the LUSE volume. After changing the mode, restore the LUSE volume.
- Besides the host modes mentioned above, the **Host Mode** list displays the **Reserve** host modes. Please do not select any **Reserve** host mode without assistance from technical support.

Table D-4 Host Modes Options for VSP

No.	Function	When to select this option
2	VERITAS Database Edition / Advanced Cluster	Use when VERITAS Database Edition/Advanced Cluster for Real Application Clusters or VERITAS Cluster Server 4.0 or later (I/O fencing function) is used.
6	TPRLO (Third-party process layout)	Use when all the following conditions are satisfied: • The host mode **0C Windows** or **2C Windows Extension** is used • The Emulex host bus adapter is used • The mini-port driver is used • TPRLO=2 is specified for the mini-port driver parameter of the host bus adapter

No.	Function	When to select this option
7	Automatic recognition function of LUN	Use when all the following conditions are satisfied: • The host mode **00 Standard** or **09 Solaris** is used. • SUN StorEdge SAN Foundation Software Version 4.2 or later is used • You want to automate recognition of increase and decrease of devices when genuine SUN HBA is connected. Enable this HMO when connected to Hitachi NAS Platform version 12.5 or later. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
12	No display for ghost LUN	Use when all the following conditions are satisfied: • The host mode **03 HP** is used. • You want to suppress creation of device files for devices to which paths are not defined.
13	SIM report at link failure[1]	Use when you want to be informed by SIM (service information message) that the number of link failures detected between ports exceeds the threshold.
14	HP TruCluster with TrueCopy function	Use when all the following conditions are satisfied: • The host mode **07 Tru64** is used. • You want to use TruCluster to set a cluster to each of P-VOL and S-VOL for TrueCopy or Universal Replicator.
15	HACMP	Use when all the following conditions are satisfied: • The host mode **0F AIX** is used. • HACMP 5.1 Version 5.1.0.4 or later, HACMP 4.5 Version 4.5.0.13 or later, or HACMP 5.2 or later is used.
22	Veritas Cluster Server	When Veritas Cluster Server is used.
33	Set/Report Device Identifier enable	Use when all the following conditions are satisfied: • Host mode **03 HP** or **05 OpenVMS**[2] is used. Set the UUID when you set HMO 33 and host mode 05 OpenVMS is used. • You want to enable commands to assign a nickname of the device. • You want to set UUID to identify a logical volume from the host.
39	Change the nexus specified in the SCSI Target Reset	When you want to control the following ranges per host group when receiving Target Reset: • Range of job resetting. • Range of UAs (Unit Attentions) defined.
40	V-VOL expansion	When all of the following conditions are satisfied: • The host mode 0C Windows or 2C Windows Extension is used • You want to automate recognition of the DP-VOL capacity after increasing the DP-VOL capacity.
41	Prioritized device recognition command	When you want to execute commands to recognize the device preferentially.
42	Prevent "OHUB PCI retry"	When IBM Z10 Linux is used.
43	Queue Full Response	When the command queue is full in the VSP storage system connecting with the HP-UX host, and if you want to respond Queue Full, instead of Busy, from the storage system to the host.
48	HAM S-VOL Read	When you do not want to generate the failover from MCU to RCU, and when the applications that issue the Read commands more than the threshold to S-VOL of the pair made with High Availability Manager are performed.

No.	Function	When to select this option
49	BB Credit Set Up Option1[3]	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 50.
50	BB Credit Set Up Option2[3]	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 49.
51	Round Trip Set Up Option[3]	If you want to adjust the response time of the host I/O, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 65. **Note:** Set host mode option 51 for both ports on MCU and RCU.
52	HAM and Cluster software for SCSI-2 Reserve	When a cluster software using the SCSI-2 reserve is used in the High Availability Manager environment.
54	(VAAI) Support Option for the EXTENDED COPY command	When the VAAI (vStorage API for Array Integration) function of VMware ESX/ESXi 4.1 is used.
57	HAM response change	When you use **0C Windows**, **2C Windows Extension**, **01 VMware**, or **21 VMware Extention** as the host mode in the High Availability Manager environment.
60	LUN0 Change Guard	When HP-UX 11.31 is used, and when you want to prevent adding or deleting of LUN0.
61	Expanded Persistent Reserve Key	Increases Reservation Keys from 128 to 2,048.
63	(VAAI) Support Option for vStorage APIs based on T10 standards	When you connect the storage system to VMware ESXi 5.0 and use the VAAI function for T10.
65	Round Trip extended set up option[3]	If you want to adjust the response time of the host I/O when you use the host mode option 51 and the host connects the TrueCopy pair. For example, when the configuration using the maximum number of processor blades is used. Use the combination of this host mode option and the host mode option 51.
67	Change of the ED_TOV value	When the OPEN fibre channel port configuration applies to following: • The topology is the Fibre Channel direct connection. • The port type is **Target** or **RCU Target**.
68	Support Page Reclamation for Linux	When using the Page Reclamation function from the environment which is being connected to the Linux host. When connected to Hitachi NAS Platform to allow the NAS host to retrieve thin provisioning information and support the SCSI UNMAP command. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
69	Online LUSE expansion	When you want the host to be notified of expansion of LUSE volume capacity.
71	Change the Unit Attention for Blocked Pool-VOLs	When you want to change the unit attention (UA) from **NOT READY** to **MEDIUM ERROR** during the pool-VOLs blockade.

No.	Function	When to select this option
72	AIX GPFS Support	When using General Parallel File System (GPFS) in the VSP storage system connecting to the AIX host.
73	Support Option for WS2012	When using the following functions provided by Windows Server 2012 (WS2012) from an environment which is being connected to the WS2012: • Thin Provisioning function • Offload Data Transfer (ODX) function

Notes:

1. Configure these host mode options only when requested to do so.
2. Set the UUID when you set host mode option 33 and host mode 05 OpenVMS is used.
3. Host mode options 49, 50, 51, and 65 are enabled only for the 8UFC/16UFC package.

Host modes and host mode options for VSP G1000, G1500 and VSP F1500

Table D-5 Host Modes for VSP G1000, G1500 and VSP F1500

Host mode	When to select this mode
00 Standard	When registering Red Hat Linux server hosts or IRIX server hosts in the host group.
01 VMware	When registering VMware server hosts in the host group.[1]
03 HP	When registering HP-UX server hosts in the host group.
05 OpenVMS	When registering OpenVMS server hosts in the host group.
07 Tru64	When registering Tru64 server hosts in the host group.
09 Solaris	When registering Solaris server hosts in the host group.
0A NetWare	When registering NetWare server hosts in the host group.
0C Windows	When registering Windows server hosts in the host group.[2]
0F AIX	When registering AIX server hosts in the host group
21 VMware Extension	When registering VMware server hosts in the host group. If the virtual host on VMware recognizes LUs by the Raw Device Mapping (RDM) method, set the host mode related to OS of the virtual host.
2C Windows Extension	When registering Windows server hosts in the host group.

Notes:

1. There are no functional differences between host mode 01 and 21. When you first connect a host, it is recommended that you set host mode 21.
2. There are no functional differences between host mode 0C and 2C. When you first connect a host, it is recommended that you set host mode 2C.

Table D-6 Host Modes Options for VSP G1000, G1500 and VSP F1500

No.	Function	When to select this option
2	VERITAS Database Edition / Advanced Cluster	Use when VERITAS Database Edition/Advanced Cluster for Real Application Clusters or VERITAS Cluster Server 4.0 or later (I/O fencing function) is used.
6	TPRLO (Third-party process layout)	Use when all the following conditions are satisfied: • The host mode **0C Windows** or **2C Windows Extension** is used • The Emulex host bus adapter is used • The mini-port driver is used • TPRLO=2 is specified for the mini-port driver parameter of the host bus adapter
7	Automatic recognition function of LUN	Use when all the following conditions are satisfied: • The host mode **00 Standard** or **09 Solaris** is used. • SUN StorEdge SAN Foundation Software Version 4.2 or later is used • You want to automate recognition of increase and decrease of devices when genuine SUN HBA is connected. Enable this HMO when connected to Hitachi NAS Platform version 12.5 or later. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
12	No display for ghost LUN	Use when all the following conditions are satisfied: • The host mode **03 HP** is used. • You want to suppress creation of device files for devices to which paths are not defined.
13	SIM report at link failure[1]	Use when you want to be informed by SIM (service information message) that the number of link failures detected between ports exceeds the threshold.
14	HP TruCluster with TrueCopy function	Use when all the following conditions are satisfied: • The host mode **07 Tru64** is used. • You want to use TruCluster to set a cluster to each of P-VOL and S-VOL for TrueCopy or Universal Replicator.
15	HACMP	Use when all the following conditions are satisfied: • The host mode **0F AIX** is used. • HACMP 5.1 Version 5.1.0.4 or later, HACMP 4.5 Version 4.5.0.13 or later, or HACMP 5.2 or later is used.
22	Veritas Cluster Server	When Veritas Cluster Server is used.
25	Support SPC-3 behavior on Persistent Reservation	When running the PERSISTENT RESERVE OUT (Service Action = REISTER AND IGNORE EXISTING KEY) command, if there is no reserved key, response is changed depending on the option setting as follows. Mode 25 = ON: Good Status (SPC-3 response) is returned without any processing. Mode 25 = OFF (default): Reservation Conflict (SPC-2 response) is returned. **Notes:** 1. The option is applied if the following are used while there is no reserved key to delete when running the PERSISTENT RESERVE OUT command: - Windows Server Failover Clustering (WSFC) - Microsoft Failover Cluster (MSFC) - Symantec Cluster Server (Previously named as Veritas Cluster Server (VCS)) 2. Depending on host type, the response when the option is set to OFF is expected.

No.	Function	When to select this option
33	Set/Report Device Identifier enable	Use when all the following conditions are satisfied: • Host mode **03 HP** or **05 OpenVMS**[2] is used. Set the UUID when you set HMO 33 and host mode 05 OpenVMS is used. • You want to enable commands to assign a nickname of the device. • You want to set UUID to identify a logical volume from the host.
39	Change the nexus specified in the SCSI Target Reset	When you want to control the following ranges per host group when receiving Target Reset: • Range of job resetting. • Range of UAs (Unit Attentions) defined.
40	V-VOL expansion	When all of the following conditions are satisfied: • The host mode 0C Windows or 2C Windows Extension is used • You want to automate recognition of the DP-VOL capacity after increasing the DP-VOL capacity.
41	Prioritized device recognition command	When you want to execute commands to recognize the device preferentially.
43	Queue Full Response	When the command queue is full in the VSP storage system connecting with the HP-UX host, and if you want to respond Queue Full, instead of Busy, from the storage system to the host.
49	BB Credit Set Up Option1	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 50.
50	BB Credit Set Up Option2	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 49.
51	Round Trip Set Up Option	If you want to adjust the response time of the host I/O, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 65. **Notes:** 1. Set HMO 51 for both ports on the local and remote storage systems. 2. This HMO does not support channel packages for 8FC16 and 16FE10. If these channel packages are used, do not set the host mode option 51.
54	(VAAI) Support Option for the EXTENDED COPY command	When the VAAI (vStorage API for Array Integration) function of VMware ESX/ESXi 4.1 is used.
60	LUN0 Change Guard	When HP-UX 11.31 is used, and when you want to prevent adding or deleting of LUN0.
63	(VAAI) Support Option for vStorage APIs based on T10 standards	When you connect the storage system to VMware ESXi 5.0 and use the VAAI function for T10.
67	Change of the ED_TOV value	When the OPEN fibre channel port configuration applies to following: • The topology is the Fibre Channel direct connection. • The port type is **Target** or **RCU Target**.

No.	Function	When to select this option
68	Support Page Reclamation for Linux	When using the Page Reclamation function from the environment which is being connected to the Linux host. When connected to Hitachi NAS Platform to allow the NAS host to retrieve thin provisioning information and support the SCSI UNMAP command. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
71	Change the Unit Attention for Blocked Pool-VOLs	When you want to change the unit attention (UA) from **NOT READY** to **MEDIUM ERROR** during the pool-VOLs blockade.
72	AIX GPFS Support	When using General Parallel File System (GPFS) in the VSP G1000, VSP G1500, or VSP F1500 storage system connecting to the AIX host.
73	Support Option for WS2012	When using the following functions provided by Windows Server 2012 (WS2012) from an environment which is being connected to the WS2012: • Thin Provisioning function • Offload Data Transfer (ODX) function Microcode: DKCMAIN 80-01-22-00/00 and later.
78	Non-preferred path option	When all of following conditions are satisfied: • Global-active device is used in the configuration with the data centers (Metro configuration). • Hitachi Dynamic Link Manager is used as the alternative path software. • The host group is on the non-optimized path of Hitachi Dynamic Link Manager. • The performance deterioration of I/O responses can be avoided without I/O using the non-optimized path of Hitachi Dynamic Link Manager. Microcode: DKCMAIN 80-01-42-00/00 and later.
80	Multi Text OFF	When using an iSCSI interface and the storage system is connected to a host OS that does not support the Multi Text function. For instance, connecting the storage system and an RHEL5.0 host that does not support the Multi Text function. Microcode: DKCMAIN 80-03-31-00/00 and later.
81	NOP-In Suppress Mode	In an iSCSI connection environment, the delay replying of the Delayed Acknowledgment function that is located on the upper layer is restrained by sending NOP-IN of executing of sense commands such as Inquiry, Test unit ready, or Mode sense. Select this option when connecting the storage system and a host to which it is not necessary to send the NOP-IN. For example: • When connecting the storage system and the Open Enterprise Server of Novell Co., Ltd. • When connecting the storage system and winBoot/i of emBoot Co., Ltd. Microcode: DKCMAIN 80-03-31-00/00 and later.
82	Discovery CHAP Mode	Select this option when the CHAP authentication is performed at the time of the discovery login in the iSCSI connection environment. For example: When the CHAP authentication is performed at the time of the discovery login in the iSCSI environment of the VMware host and the storage system. Microcode: DKCMAIN 80-03-31-00/00 and later.

No.	Function	When to select this option
83	Report iSCSI Full Portal List Mode	When configuring alternate paths in the environment of connecting the VMware host and storage system: If waiting for the reply of the target information from the HMO 83 enabled ports other than ports of discovery login, select this host mode option. When both of the following are satisfied, select this host mode option: • Configuring alternate paths in the environment of connecting the VMware host and the storage system • Waiting for replying of the target information from the ports other than ports of discovery login Microcode: DKCMAIN 80-03-31-00/00 and later.
88	Nondisruptive migration with HP-UX hosts	When converging multiple host-target ports used in the migration source storage system on the migration target storage system, and enable LUN path definition from a host group belonging to a virtual storage machine to an LDEV defined in a different virtual storage machine. Mode 88 = ON: LUN path definition is enabled. Mode 88 = OFF (default): LUN path definition is disabled. Microcode: DKCMAIN 80-02-01-00/01 and later. **Notes:** 1. Apply this option when all the following conditions are met. • You are using the nondisruptive migration function to migrate volumes in multiple old storage models that use the same server. • You need to reducte the number of Target ports used on the migration target storage system. • The host is an HP-UX server. 2. Applying this option to a server other than HP-UX can cause the following: • Path addition from the server to the migration target storage system might fail. • Display of devices that the server recognizes might be invalid. 3. If a LUN path is defined to an LDEV defined in a virtual storage machine different from the one to which the host group belongs, this option cannot be set to OFF.
96	Change the nexus specified in the SCSI Logical Unit Reset	When you want to control the following ranges per host group when receiving LU Reset: • Range of job resetting. • Range of UAs (Unit Attentions) defined.
97	Proprietary ANCHOR command support	This option is used to support the Proprietary ANCHOR command (operation code=0xC1). HMO 97 = ON: The Proprietary ANCHOR command is supported. HMO 97 = OFF (default): The Proprietary ANCHOR command is not supported. Microcode: DKCMAIN 80-03-31-00/00 and later. **Notes:** (1) The option is applied when using the Proprietary ANCHOR command in the HNAS environment. (2) The option is used only in the HNAS environment. Any other environments than HNAS do not issue the Proprietary ANCHOR command. (3) When the option is set to ON, make sure that SOM 1079 (a system option mode to disable Proprietary ANCHOR command) is set to OFF. SOM 1079 is used to disable the Proprietary ANCHOR command so as to enable microcode downgrade from a version that supports the Proprietary ANCHOR command to a version that does not support the command. If SOM 1097 is set to ON, the Proprietary ANCHOR command cannot be run even when HMO 97 is set to ON.

No.	Function	When to select this option
102	(GAD) Standard Inquiry Expansion for HCS	When all of the following conditions are satisfied: • The OS of the host is Windows or AIX, and the MPIO function is used. • Global-active device (GAD) or nondisruptive migration (NDM) is used. • Hitachi Device Manager (HDvM) agent is used.
105	Task Set Full response in the event of I/O overload	When all of following conditions are satisfied: • The host mode 0C Windows or 2C Windows Extension is used. • You want to return Task Set Full response from the storage system to the host when an overload of I/Os occurs on the storage system.
Notes: 1. Configure these host mode options only when requested to do so. 2. Set the UUID when you set host mode option 33 and host mode 05 OpenVMS is used.		

Host modes and host mode options for HUS VM

Table D-7 Host Modes for HUS VM

Host mode	When to select this mode
00 Standard	When registering Red Hat Linux server hosts or IRIX server hosts in the host group.
01 VMware	When registering VMware server hosts in the host group (see ***Notes***).
03 HP	When registering HP-UX server hosts in the host group.
05 OpenVMS	When registering OpenVMS server hosts in the host group.
07 Tru64	When registering Tru64 server hosts in the host group.
09 Solaris	When registering Solaris server hosts in the host group.
0A NetWare	When registering NetWare server hosts in the host group.
0C Windows	When registering Windows server hosts in the host group (see ***Notes***).
0F AIX	When registering AIX server hosts in the host group
21 VMware Extension	When registering VMware server hosts in the host group (see ***Notes***).
2C Windows Extension	When registering Windows server hosts in the host group (see ***Notes***).

Notes:

- If Windows server hosts are registered in a host group, ensure that the host mode of the host group is 0C Windows or 2C Windows Extension.

  If the host mode of a host group is 0C Windows and an LU path is defined between the host group and a logical volume, the logical volume cannot be combined with other logical volumes to form a LUSE volume (that is, an expanded LU).

  If the host mode of a host group is 2C Windows Extension and an LU path is defined between the host group and a logical volume, the logical volume can be combined with other logical volumes to form a LUSE volume (that is, an expanded LU). If you plan to expand LUs by using LUSE in the future, set the host mode 2C Windows Extension.
- If VMware server hosts are registered in a host group, ensure that the host mode of the host group is 01 VMware or 21 VMware Extension.

  If the host mode of a host group is 01 VMware and an LU path is defined between the host group and a logical volume, the logical volume cannot be combined with other logical volumes to form a LUSE volume (that is, an expanded LU).

  If the host mode of a host group is 21 VMware Extension and an LU path is defined between the host group and a logical volume, the logical volume can be combined with other logical volumes to form a LUSE volume (that is, an expanded LU). If you plan to expand LUs by using LUSE in the future, set the host mode 21 VMware Extension.

- If you plan to expand LUs by using LUSE in case of Windows virtual host on VMware recognizing LU by Raw Device Mapping (RDM) method, set the host mode 2C Windows Extension. If the host mode 2C Windows Extension is not set, change the host mode to 2C. Before changing the host mode, back up the LUSE volume. After changing the mode, restore the LUSE volume.

Table D-8 Host Mode Options for HUS VM

No.	Function	When to select this option
2	VERITAS Database Edition / Advanced Cluster	Use when VERITAS Database Edition/Advanced Cluster for Real Application Clusters or VERITAS Cluster Server 4.0 or later (I/O fencing function) is used.
6	TPRLO (Third-party process layout)	Use when all the following conditions are satisfied: • The host mode **0C Windows** or **2C Windows Extension** is used • The Emulex host bus adapter is used • The mini-port driver is used • TPRLO=2 is specified for the mini-port driver parameter of the host bus adapter
7	Automatic recognition function of LUN	Use when all the following conditions are satisfied: • The host mode **00 Standard** or **09 Solaris** is used. • SUN StorEdge SAN Foundation Software Version 4.2 or higher is used • You want to automate recognition of increase and decrease of devices when genuine SUN HBA is connected. Enable this HMO when connected to Hitachi NAS Platform version 12.5 or later. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
12	No display for ghost LUN	Use when all the following conditions are satisfied: • The host mode **03 HP** is used. • You want to suppress creation of device files for devices to which paths are not defined.
13	SIM report at link failure[1]	Use when you want to be informed by SIM (service information message) that the number of link failures detected between ports exceeds the threshold.
14	HP TruCluster with TrueCopy function	Use when all the following conditions are satisfied: • The host mode **07 Tru64** is used. • You want to use TruCluster to set a cluster to each of P-VOL and S-VOL for TrueCopy or Universal Replicator.
15	HACMP	Use when all the following conditions are satisfied: • The host mode **0F AIX** is used. • HACMP 5.1 Version 5.1.0.4 or later, HACMP 4.5 Version 4.5.0.13 or later, or HACMP 5.2 or later is used.
22	Veritas Cluster Server	When Veritas Cluster Server is used.
33	Set/Report Device Identifier enable	Use when all the following conditions are satisfied: • Host mode **03 HP** or **05 OpenVMS**[2] is used. Set the UUID when you set HMO 33 and host mode 05 OpenVMS is used. • You want to enable commands to assign a nickname of the device. • You want to set UUID to identify a logical volume from the host.
39	Change the nexus specified in the SCSI Target Reset	When you want to control the following ranges per host group when receiving Target Reset: • Range of job resetting. • Range of UAs (Unit Attentions) defined.

No.	Function	When to select this option
40	V-VOL expansion	When all of the following conditions are satisfied: • The host mode 0C Windows or 2C Windows Extension is used • You want to automate recognition of the DP-VOL capacity after increasing the DP-VOL capacity.
41	Prioritized device recognition command	When you want to execute commands to recognize the device preferentially.
42	Prevent "OHUB PCI retry"	When IBM Z10 Linux is used.
43	Queue Full Response	When the command queue is full in the HUS VM storage system connecting with the HP-UX host, and if you want to respond Queue Full, instead of Busy, from the storage system to the host.
48	HAM S-VOL Read	When you do not want to generate the failover from MCU to RCU, and when the applications that issue the Read commands more than the threshold to S-VOL of the pair made with High Availability Manager are performed.
49	BB Credit Set Up Option1[3]	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 50.
50	BB Credit Set Up Option2[3]	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 49.
51	Round Trip Set Up Option[3]	If you want to adjust the response time of the host I/O, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 65. **Note:** Set host mode option 51 for both ports on MCU and RCU.
52	HAM and Cluster software for SCSI-2 Reserve	When a cluster software using the SCSI-2 reserve is used in the High Availability Manager environment.
54	(VAAI) Support Option for the EXTENDED COPY command	When the VAAI (vStorage API for Array Integration) function of VMware ESX/ESXi 4.1 is used.
57	HAM response change	When you use **0C Windows**, **2C Windows Extension**, **01 VMware**, or **21 VMware Extention** as the host mode in the High Availability Manager environment.
60	LUN0 Change Guard	When HP-UX 11.31 is used, and when you want to prevent adding or deleting of LUN0.
61	Expanded Persistent Reserve Key	Increases Reservation Keys from 128 to 2,048.
63	(VAAI) Support Option for vStorage APIs based on T10 standards	When you connect the storage system to VMware ESXi 5.0 and use the VAAI function for T10.
67	Change of the ED_TOV value	When the OPEN fibre channel port configuration applies to following: • The topology is the Fibre Channel direct connection. • The port type is **Target** or **RCU Target**.

No.	Function	When to select this option
68	Support Page Reclamation for Linux	When using the Page Reclamation function from the environment which is being connected to the Linux host. When connected to Hitachi NAS Platform to allow the NAS host to retrieve thin provisioning information and support the SCSI UNMAP command. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
69	Online LUSE expansion	When you want the host to be notified of expansion of LUSE volume capacity.
71	Change the Unit Attention for Blocked Pool-VOLs	When you want to change the unit attention (UA) from **NOT READY** to **MEDIUM ERROR** during the pool-VOLs blockade.
72	AIX GPFS Support	When using General Parallel File System (GPFS) in the HUS VM storage system connecting to the AIX host.
73	Support Option for WS2012	When using the following functions provided by Windows Server 2012 (WS2012) from an environment which is being connected to the WS2012: • Thin Provisioning function • Offload Data Transfer (ODX) function

Notes:

1. Configure these host mode options only when requested to do so.
2. Set the UUID when you set host mode option 33 and host mode 05 OpenVMS is used.
3. Host mode options 49, 50, and 51 are enabled only for the HF8G package.

Host modes and host mode options for VSP Gx00 and Fx00

Table D-9 Host Modes for VSP Gx00 and VSP Fx00 models

Host mode	When to select this mode
00 Standard	When registering Red Hat Linux server hosts or IRIX server hosts in the host group.
01 VMware	When registering VMware server hosts in the host group.[1]
03 HP	When registering HP-UX server hosts in the host group.
05 OpenVMS	When registering OpenVMS server hosts in the host group.
07 Tru64	When registering Tru64 server hosts in the host group.
09 Solaris	When registering Solaris server hosts in the host group.
0A NetWare	When registering NetWare server hosts in the host group.
0C Windows	When registering Windows server hosts in the host group.[2]
0F AIX	When registering AIX server hosts in the host group
21 VMware Extension	When registering VMware server hosts. If the virtual host on VMware recognizes LUs by the Raw Device Mapping (RDM) method, set the host mode related to OS of the virtual host. Example: If a LUN/LDEV is formatted as VMFS (where virtual machines and their VMDK's usually reside), it should be set with HMO 21. However, if a LUN/LDEV is formatted as a specific file system format (for example, NTFS) and has application requirements to be presented directly to a virtual machine as an RDM, it should be set to the HMO specific to that OS/filesystem (e.g., such as HMO 2C for Windows). A common example of VM's with this mix is: C: drive – OS VMDK on VMFS D: drive – RDM for application data In this example, 2 different Host Groups should be created for a single host with different HMO and LUN's assigned.
2C Windows Extension	When registering Windows server hosts in the host group.

Notes:

1. There are no functional differences between host mode 01 and 21. When you first connect a host, it is recommended that you set host mode 21.
2. There are no functional differences between host mode 0C and 2C. When you first connect a host, it is recommended that you set host mode 2C.

Table D-10 Host Modes for VSP Gx00 and VSP Fx00 models

No.	Function	When to select this option
2	VERITAS Database Edition / Advanced Cluster	Use when VERITAS Database Edition/Advanced Cluster for Real Application Clusters or VERITAS Cluster Server 4.0 or later (I/O fencing function) is used.
6	TPRLO (Third-party process layout)	Use when all the following conditions are satisfied: • The host mode **0C Windows** or **2C Windows Extension** is used • The Emulex host bus adapter is used • The mini-port driver is used • TPRLO=2 is specified for the mini-port driver parameter of the host bus adapter

No.	Function	When to select this option
7	Automatic recognition function of LUN	Use when all the following conditions are satisfied: • The host mode **00 Standard** or **09 Solaris** is used. • SUN StorEdge SAN Foundation Software Version 4.2 or higher is used • You want to automate recognition of increase and decrease of devices when genuine SUN HBA is connected. Enable this HMO when connected to Hitachi NAS Platform (unified or external) version 12.5 or later. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
12	No display for ghost LUN	Use when all the following conditions are satisfied: • The host mode **03 HP** is used. • You want to suppress creation of device files for devices to which paths are not defined.
13	SIM report at link failure[1]	Use when you want to be informed by SIM (service information message) that the number of link failures detected between ports exceeds the threshold.
14	HP TruCluster with TrueCopy function	Use when all the following conditions are satisfied: • The host mode **07 Tru64** is used. • You want to use TruCluster to set a cluster to each of P-VOL and S-VOL for TrueCopy or Universal Replicator.
15	HACMP	Use when all the following conditions are satisfied: • The host mode **0F AIX** is used. • HACMP 5.1 Version 5.1.0.4 or later, HACMP 4.5 Version 4.5.0.13 or later, or HACMP 5.2 or later is used.
22	Veritas Cluster Server	When Veritas Cluster Server is used.
25	Support SPC-3 behavior on Persistent Reservation	Select this option when one of the following conditions is satisfied: ▪ Using Windows Server Failover Clustering (WSFC) ▪ Using Microsoft Failover Cluster (MSFC) ▪ Using Symantec Cluster Server, also known as Veritas Cluster Server (VCS) ▪ Using a configuration other than above with the `PERSISTENT RESERVE OUT` (Service Action=REGISTER AND IGNORE EXISTING KEY) command, change the status response from Reservation-Conflict to Good-Status when there is no registered key to be deleted. Microcode: DKCMAIN earlier than 83-02-01-20/00
33	Set/Report Device Identifier enable	Use when all the following conditions are satisfied: • Host mode **03 HP** or **05 OpenVMS**[2] is used. Set the UUID when you set HMO 33 and host mode 05 OpenVMS is used. • You want to enable commands to assign a nickname of the device. • You want to set UUID to identify a logical volume from the host.
39	Change the nexus specified in the SCSI Target Reset	When you want to control the following ranges per host group when receiving Target Reset: • Range of job resetting. • Range of UAs (Unit Attentions) defined.
40	V-VOL expansion	When all of the following conditions are satisfied: • The host mode 0C Windows or 2C Windows Extension is used • You want to automate recognition of the DP-VOL capacity after increasing the DP-VOL capacity.

No.	Function	When to select this option
41	Prioritized device recognition command	When you want to execute commands to recognize the device preferentially.
43	Queue Full Response	When the command queue is full in the VSP Gx00 or VSP Fx00 storage system connecting with the HP-UX host, and if you want to respond Queue Full, instead of Busy, from the storage system to the host.
49	BB Credit Set Up Option1	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy or GAD pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 50.
50	BB Credit Set Up Option2	When you want to adjust the number of buffer-to-buffer credits (BBCs) to control the transfer data size by the fibre channel, for example when the distance between MCU and RCU of the TrueCopy pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 49.
51	Round Trip Set Up Option	If you want to adjust the response time of the host I/O, for example when the distance between MCU and RCU of the TrueCopy or GAD pair is long (approximately 100 kilometers) and the Point-to-Point topology is used. Use the combination of this host mode option and the host mode option 65. **Note:** Set HMO 51 for ports on the remote site of the TrueCopy pair or the global-active device pair.
54	(VAAI) Support Option for the EXTENDED COPY command	When the VAAI (vStorage API for Array Integration) function of VMware ESX/ESXi 4.1 is used.
60	LUN0 Change Guard	When HP-UX 11.31 is used, and when you want to prevent adding or deleting of LUN0.
63	(VAAI) Support Option for vStorage APIs based on T10 standards	When you connect the storage system to VMware ESXi 5.0 and use the VAAI function for T10.
67	Change of the ED_TOV value	When the OPEN fibre channel port configuration applies to following: • The topology is the Fibre Channel direct connection. • The port type is **Target** or **RCU Target**.
68	Support Page Reclamation for Linux	When using the Page Reclamation function from the environment which is being connected to the Linux host. When connected to Hitachi NAS Platform (unified or external) to allow the NAS host to retrieve thin provisioning information and support the SCSI UNMAP command. ***Note:*** When you use Hitachi Storage Advisor (HSA) to allocate storage to the NAS modules on ports 1A and 2A, this HMO is enabled automatically.
71	Change the Unit Attention for Blocked Pool-VOLs	When you want to change the unit attention (UA) from **NOT READY** to **MEDIUM ERROR** during the pool-VOLs blockade.
72	AIX GPFS Support	When using General Parallel File System (GPFS) in the VSP Gx00 or VSP Fx00 storage system connecting to the AIX host.
73	Support Option for WS2012	When using the following functions provided by Windows Server 2012 (WS2012) from an environment which is being connected to the WS2012: • Thin Provisioning function • Offload Data Transfer (ODX) function

No.	Function	When to select this option
78	The non-preferred path option	When all of following conditions are satisfied: • Global-active device is used in the configuration with the data centers (Metro configuration). • Hitachi Dynamic Link Manager is used as the alternative path software. • The host group is on the non-optimized path of Hitachi Dynamic Link Manager. • The performance deterioration of I/O responses can be avoided without I/O using the non-optimized path of Hitachi Dynamic Link Manager. Microcode: DKCMAIN 83-01-21-20/00 and later.
80	Multi Text OFF	When using an iSCSI interface and the storage system is connected to a host OS that does not support the Multi Text function. For instance, connecting the storage system and an RHEL5.0 host that does not support the Multi Text function.
81	NOP-In Suppress Mode	In an iSCSI environment, the delay replying of the Delayed Acknowledgment function that is located on the upper layer is restrained by sending NOP-IN of executing of sense commands such as Inquiry, Test unit ready, or Mode sense. Select this option when connecting the storage system and a host to which it is not necessary to send the NOP-IN. For example: • When connecting the storage system and the Open Enterprise Server of Novell. • When connecting the storage system and winBoot/i of emBoot Co., Ltd.
82	Discovery CHAP Mode	Select this option when the CHAP authentication is performed at the time of the discovery login in the iSCSI connection environment. For example: When the CHAP authentication is performed at the time of the discovery login in the iSCSI environment of the VMware host and the storage system.
83	Report iSCSI Full Portal List Mode	When configuring alternate paths in the environment of connecting the VMware host and storage system: If waiting for the reply of the target information from the HMO 83 enabled ports other than ports of discovery login, select this host mode option. When both of the following are satisfied, select this host mode option: • Configuring alternate paths in the environment of connecting the VMware host and the storage system • Waiting for replying of the target information from the ports other than ports of discovery login
88	Nondisruptive migration with HP-UX hosts	When converging multiple host-target ports used in the migration source storage system on the migration target storage system, and enable LUN path definition from a host group belonging to a virtual storage machine to an LDEV defined in a different virtual storage machine. Mode 88 = ON: LUN path definition is enabled. Mode 88 = OFF (default): LUN path definition is disabled. Microcode: DKCMAIN 80-02-01-00/01 and later. **Notes:** 1. Apply this option when all the following conditions are met. • You are using the nondisruptive migration function to migrate volumes in multiple old storage models that use the same server. • You need to reducte the number of Target ports used on the migration target storage system. • The host is an HP-UX server. 2. Applying this option to a server other than HP-UX can cause the following: • Path addition from the server to the migration target storage system might fail. • Display of devices that the server recognizes might be invalid. 3. If a LUN path is defined to an LDEV defined in a virtual storage machine different from the one to which the host group belongs, this option cannot be set to OFF.

No.	Function	When to select this option
96	Change the nexus specified in the SCSI Logical Unit Reset	Select this option when you want to control the following ranges per host group when receiving LU Reset: • Range of job resetting. • Range of UAs (Unit Attentions) defined.
97	Proprietary ANCHOR command support	Select this option when connecting to Hitachi NAS Platform. Microcode: DKCMAIN 83-02-01-20/00 and later.
100	Hitachi HBA (Fabric Emulation Mode) Connection Option[1]	Select this option when connecting the 8 Gbps channel port (in the storage system) and the adapter of BladeSymphony/HA8000 Hitachi Gigabit Fibre Channel by using the Fabric Emulation mode.
102	(GAD) Standard Inquiry Expansion for Hitachi Command Suite	Select this option when all of the following conditions are satisfied: • The OS of the host is Windows or AIX, and the MPIO function is used. • GAD (global-active device) or nondisruptive migration (NDM) is used. • Hitachi Device Manager (HDvM) agent is used.
105	Task Set Full response in the event of I/O overload	When all of following conditions are satisfied: • The host mode 0C Windows or 2C Windows Extension is used. • You want to return Task Set Full response from the storage system to the host when an overload of I/Os occurs on the storage system.

Notes:

1. Configure these host mode options only when requested to do so.
2. Set the UUID when you set host mode option 33 and host mode 05 OpenVMS is used.

Acronyms and abbreviations

AL	arbitrated loop
AL-PA	arbitrated loop physical address
ALUA	asymmetric logical unit access
blk	block
CNA	converged network adapter
CVS	custom volume size
dev	device
FC	fibre-channel
FCP	fibre-channel protocol
GAD	global-active device
GB	gigabyte
Gbps	gigabits per second
HA	high availability
HBA	host bus adapter
HDS	Hitachi Data Systems
HMO	host mode option
HUS VM	Hitachi Unified Storage VM
I/O	input/output
IPL	initial program load
KB	kilobyte
LCU	logical control unit
LDEV	logical device
LU	logical unit
LUN	logical unit number
LUSE	LUN Expansion
LV	logical volume
LVI	logical volume image
LVM	Logical Volume Manager, logical volume management
MB	megabyte

OFC	open fibre control
OS	operating system
PA	physical address
PB	petabyte
PC	personal computer
RAID	redundant array of inexpensive disks
RHEL	Red Hat Enterprise Linux
rw	read-write
SIM	service information message
SNMP	simple network management protocol
SOM	system option mode
SR	storage repository
TB	terabyte
TID	target ID
TUF	Technical Upload Facility
USP V	Hitachi Universal Storage Platform V
USP VM	Hitachi Universal Storage Platform VM
VLL	Virtual LVI/LUN
VSP	Hitachi Virtual Storage Platform
VSP Fx00	Hitachi Virtual Storage Platform F400, F600, F800
VSP Gx00	Hitachi Virtual Storage Platform G200, G400, G600, G800
WWN	worldwide name

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific
+852 3189 7900
hds.marketing.apac@hds.com

Contact Us
www.hds.com/en-us/contact.html

Hitachi Data Systems

MK-90RD7037-09