



Configuration Guide for Red Hat[®] Linux[®] Host Attachment

Hitachi Virtual Storage Platform
Hitachi Universal Storage Platform V/VM

FASTFIND LINKS

[Document Organization](#)

[Product Version](#)

[Getting Help](#)

[Contents](#)

Copyright © 2010 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 photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd. (hereinafter referred to as "Hitachi") and Hitachi Data Systems Corporation (hereinafter referred to as "Hitachi Data Systems").

Hitachi Data Systems reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.

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

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

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

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

Microsoft product screen shot(s) reprinted with permission from Microsoft Corporation.



Contents

Preface	v
Intended Audience	vi
Product Version.....	vi
Document Revision Level	vi
Source Documents for this Revision	vii
Changes in this Revision	vii
Referenced Documents.....	vii
Document Organization	viii
Document Conventions.....	ix
Convention for Storage Capacity Values	x
Accessing Product Documentation	x
Getting Help	xi
Comments.....	xi
Introduction	1-1
About the Hitachi RAID Storage Systems.....	1-2
Device Types	1-3
Installation and Configuration Roadmap.....	1-7
Installing the Storage System	2-1
Requirements	2-2
Preparing for the Storage System Installation.....	2-3
Hardware Installation Considerations.....	2-3
LUN Manager Software Installation	2-4
Setting the Host Mode	2-4
Setting the Host Mode Options	2-5
Configuring the Fibre-Channel Ports.....	2-6
Port Address Considerations for Fabric Environments.....	2-7
Loop ID Conflicts	2-7

Connecting the Storage System to the Red Hat Linux Host.....	2-8
Configuring the Host Fibre-Channel HBA(s).....	2-9
Verifying New Device Recognition	2-10
Configuring the New Disk Devices.....	3-1
Setting the Number of Logical Units	3-2
Partitioning the Devices	3-3
Creating, Mounting, and Verifying the File Systems	3-4
Creating the File Systems	3-4
Creating the Mount Directories.....	3-4
Mounting the New File Systems.....	3-4
Verifying the File Systems.....	3-4
Setting the Auto-Mount Parameters.....	3-5
Failover and SNMP Operation	4-1
Host Failover.....	4-2
Path Failover.....	4-2
SNMP Remote System Management.....	4-3
Troubleshooting	5-1
General Troubleshooting.....	5-2
Calling the Hitachi Data Systems Support Center.....	5-3
Note on Using Veritas Cluster Server.....	A-1
Acronyms and Abbreviations	



Preface

This document describes and provides instructions for installing and configuring the devices on the Hitachi RAID storage systems for operations in a Red Hat® Linux® environment. The Hitachi RAID storage system models include the Hitachi Virtual Storage Platform (VSP) and the Hitachi Universal Storage Platform V and Hitachi Universal Storage Platform VM (USP V/VM).

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

This preface includes the following information:

- [Intended Audience](#)
- [Product Version](#)
- [Document Revision Level](#)
- [Source Documents for this Revision](#)
- [Changes in this Revision](#)
- [Referenced Documents](#)
- [Document Organization](#)
- [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 involved in installing, configuring, and operating the Hitachi RAID storage systems.

Readers of this document should meet the following requirements:

- You should have a background in data processing and understand RAID storage systems and their basic functions.
- You should be familiar with the Hitachi RAID storage system(s), and you should have read the *User and Reference Guide* for the storage system.
- You should be familiar with the Storage Navigator software for the Hitachi RAID storage system(s), and you should have read the *Storage Navigator User's Guide*.
- You should be familiar with the Red Hat Linux operating system and the hardware hosting the Red Hat Linux system.
- You should be familiar with the hardware used to attach the Hitachi RAID storage system to the Red Hat Linux host, including fibre-channel cabling, host bus adapters (HBAs), switches, and hubs.

Product Version

This document revision applies to the following microcode levels:

- Hitachi Virtual Storage Platform microcode 70-01-0x or later.
- Hitachi Universal Storage Platform V/VM microcode 60-01-3x or later.

Document Revision Level

Revision	Date	Description
MK-96RD640-P	February 2007	Initial Release
MK-96RD640-00	May 2007	Initial Release, supersedes and replaces MK-96RD640-P
MK-96RD640-01	September 2007	Revision 01, supersedes and replaces MK-96RD640-00
MK-96RD640-02	February 2010	Revision 02, supersedes and replaces MK-96RD640-01
MK-96RD640-03	October 2010	Revision 03, supersedes and replaces MK-96RD640-02

Source Documents for this Revision

- MK-96RD640-03a-1_RSDreview

Changes in this Revision

- Added the Hitachi Virtual Storage Platform storage system.
- Added information about the 8-Gbps fibre-channel interface ([Table 2-1](#)).
- Corrected the queue depth value from 1024 to 2048 ([Table 2-7](#)).
- Added a link to the Hitachi Data Systems interoperability site for specific information about supported OS versions, HBAs, drivers, hubs, and switches ([Table 2-1](#)).

Referenced Documents

Hitachi Virtual Storage Platform documentation:

- *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 documentation:

- *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
- *Cross-OS File Exchange User's Guide*, MK-96RD647
- *Hitachi Dynamic Link Manager for Red Hat Linux User's Guide*, MK-92DLM113

Red Hat Linux documentation

Document Organization

The following table provides an overview of the contents and organization of this document. Click the [chapter title](#) in the left column to go to that chapter. The first page of each chapter provides links to the sections in that chapter.

Chapter	Description
Chapter 1, Introduction	Provides a brief overview of the Hitachi RAID storage systems, supported device types, and an installation roadmap.
Chapter 2, Installing the Storage System	Provides instructions for installing and connecting the Hitachi RAID storage system to a Red Hat Linux host.
Chapter 3, Configuring the New Disk Devices	Provides instructions for configuring the new devices for use.
Chapter 4, Failover and SNMP Operation	Describes how to configure the Hitachi RAID storage system for failover and SNMP.
Chapter 5, Troubleshooting	Provides information for identifying and resolving problems.
Appendix A, Note on Using Veritas Cluster Server	Provides information about adding reserve keys for LUs to increase disk capacity.

Document Conventions





The terms “Virtual Storage Platform” and “VSP” refer to all models of the Hitachi Virtual Storage Platform storage system, unless otherwise noted.

The terms “Universal Storage Platform V/VM” and “USP V/VM” refer to all models of the Universal Storage Platform V/VM, unless otherwise noted.

This document uses the following typographic conventions:

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click OK .
<i>Italic</i>	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: copy <i>source-file target-file</i> Note: Angled brackets (< >) are also used to indicate variables.
screen/code	Indicates text that is displayed on screen or entered by the user. Example: # <code>pairdisplay -g oradb</code>
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: # <code>pairdisplay -g <group></code> Note: Italic font is also used to indicate variables.
[] 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.

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

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

Convention for Storage Capacity Values

Physical storage capacity values (e.g., 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 storage capacity values (e.g., logical device capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
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	1,024 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

The user documentation for the Hitachi RAID storage systems is available on the Hitachi Data Systems Portal: <https://hdssupport.hds.com>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting Help

The Hitachi Data Systems customer support staff is available 24 hours a day, seven days a week. If you need technical support, log on to the Hitachi Data Systems Portal for contact information: <https://hdssupport.hds.com>

Comments

Please send us your comments on this document: doc.comments@hds.com
Include the document title, number, and revision, and refer to specific section(s) and paragraph(s) whenever possible.

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

Introduction

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

- [About the Hitachi RAID Storage Systems](#)
- [Device Types](#)
- [Installation and Configuration Roadmap](#)

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™ software, application-centric storage management and logical partitioning, and simplified and unified data replication across heterogeneous storage systems. These storage systems are an integral part of the Services Oriented Storage Solutions architecture from Hitachi Data Systems, providing 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)
- Non-disruptive data migration
- Volume management across heterogeneous storage arrays
- Thin provisioning
- Security services (immutability, logging, auditing, data shredding)
- Data de-duplication
- 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 Red Hat Linux, UNIX platforms, Windows, VMware, 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 configured with OPEN-V logical units (LUs) and are compatible with most fibre-channel (FC) host bus adapters (HBAs). Users can perform additional LU configuration activities using the LUN Manager, Virtual LVI/LUN (VLL), and LUN Expansion (LUSE) features provided by the Storage Navigator software, which is the primary user interface for the storage systems.

For further information on storage solutions and the Hitachi RAID storage systems, please contact your Hitachi Data Systems account team.

Device Types

[Table 1-1](#) describes the types of logical devices (volumes) that can be installed and configured for operation with the Hitachi RAID storage systems on a Red Hat Linux operating system. [Table 1-2](#) lists the specifications for devices supported by the Hitachi RAID storage systems. Logical devices are defined to the host as SCSI disk devices, even though the interface is fibre channel. For information about configuring devices other than OPEN-V, contact your Hitachi Data Systems representative.

The sector size for the devices is 512 bytes.

Table 1-1 Logical Devices Supported by the Hitachi RAID Storage Systems

Device Type	Description
OPEN-V Devices	OPEN-V logical units (LUs) are disk devices (VLL-based volumes) that do not have a predefined size.
OPEN-x Devices	OPEN-x logical units (LUs) (e.g., OPEN-3, OPEN-9) are disk devices of predefined sizes. The Hitachi RAID storage systems support OPEN-3, OPEN-8, OPEN-9, OPEN-E, and OPEN-L, devices. For the latest information on usage of these device types, contact your Hitachi Data Systems account team.
LUSE Devices (OPEN-x*n)	LUSE devices are combined LUs that can be from 2 to 36 times larger than standard OPEN-x LUs. Using LUN Expansion (LUSE) remote console software, you can configure these custom-size devices. LUSE devices are designated as OPEN-x*n, where x is the LU type (e.g., OPEN-9*n) and $2 < n < 36$). For example, a LUSE device created from 10 OPEN-3 LUs is designated as an OPEN-3*10 disk device. This lets the host combine logical devices and access the data stored on the Hitachi RAID storage system using fewer LU numbers.
VLL Devices (OPEN-x VLL)	VLL devices are custom-size LUs that are smaller than standard OPEN-x LUs. Using Virtual LVI/LUN remote console software, you can configure VLL devices by "slicing" a single LU into several smaller LUs that best fit your application needs to improve host access to frequently used files. The product name for the OPEN-x VLL devices is OPEN-x-CVS (CVS stands for custom volume size). The OPEN-L LU type does not support Virtual LVI/LUN.
VLL LUSE Devices (OPEN-x*n VLL)	VLL LUSE devices combine Virtual LVI/LUN devices (instead of standard OPEN-x LUs) into LUSE devices. Use the Virtual LVI/LUN feature to create custom-size devices, then use 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 a0 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)	<p>The Hitachi Cross-OS File Exchange (FX) software allows you to share data across mainframe, UNIX, and PC server platforms using special multiplatform volumes. The VLL feature can be applied to FX devices for maximum flexibility in volume size. For more information about FX, see the <i>Cross-OS File Exchange User's Guide</i>, or contact your Hitachi Data Systems account team.</p> <p>FX devices are not SCSI disk devices, and must be installed and accessed as raw devices. UNIX/PC server hosts must use FX to access the FX devices as raw devices (no file system, no mount operation).</p> <p>The 3390-3B devices are write-protected from UNIX/PC server access. The Hitachi RAID storage system rejects all UNIX/PC server write operations (including fibre-channel adapters) for 3390-3B devices.</p> <p>Multiplatform devices are not write-protected for UNIX/PC server access. Do not execute any write operation by the fibre-channel adapters on these devices. Do not create a partition or file system on these devices. This will overwrite any data on the FX device and prevent the FX software from accessing the device.</p>

Table 1-2 Device Specifications

Device Type	Category (Note 1)	Product Name (Note 2)	# of Blocks (512 B/blk)	# of Cylinders	# of Heads	# of Sectors per Track	Capacity (MB) (Note 3)
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-V	SCSI disk	OPEN-V	125827200 max Note 4	Note 5	15	128	Note 6
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-V*n	SCSI disk	OPEN-L*n	Note 4	Note 5	15	128	Note 6
OPEN-3 VLL	SCSI disk	OPEN-3-CVS	Note 4	Note 5	15	96	Note 6
OPEN-8 VLL	SCSI disk	OPEN-8-CVS	Note 4	Note 5	15	96	Note 6
OPEN-9 VLL	SCSI disk	OPEN-9-CVS	Note 4	Note 5	15	96	Note 6
OPEN-E VLL	SCSI disk	OPEN-E-CVS	Note 4	Note 5	15	96	Note 6
OPEN-V VLL	SCSI disk	OPEN-V	Note 4	Note 5	15	128	Note 6
OPEN-3*n VLL	SCSI disk	OPEN-3*n-CVS	Note 4	Note 5	15	96	Note 6
OPEN-8*n VLL	SCSI disk	OPEN-8*n-CVS	Note 4	Note 5	15	96	Note 6
OPEN-9*n VLL	SCSI disk	OPEN-9*n-CVS	Note 4	Note 5	15	96	Note 6
OPEN-E*n VLL	SCSI disk	OPEN-E*n-CVS	Note 4	Note 5	15	96	Note 6
OPEN-V*n VLL	SCSI disk	OPEN-V*n	Note 4	Note 5	15	128	Note 6
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 4	Note 5	15	116	Note 6
3390-3B VLL	FXmto	3390-3B-CVS	Note 4	Note 5	15	116	Note 6
3390-3C VLL	FXotm	OP-C-3390-3C-CVS	Note 4	Note 5	15	116	Note 6
FX OPEN-3 VLL	FXoto	OPEN-3-CVS	Note 4	Note 5	15	96	Note 6

Note 1: The category of a device (SCSI disk or FX) determines its volume usage. [Table 1-3](#) shows the volume usage for SCSI disk devices and FX devices. The SCSI disk devices (OPEN-x, VLL, LUSE, and VLL LUSE) are usually formatted with file systems for Red Hat Linux operations. The FX devices (3390-3A/B/C, and OPEN-x-FXoto) must be installed as raw devices and can only be accessed using the FX software. Do not partition or create a file system on any device used for FX operations.

Table 1-3 Volume Usage for Device Categories

Category	Device Type	Volume Usage
SCSI Disk	OPEN-x, OPEN-x VLL, OPEN-x*n LUSE, OPEN-x*n VLL LUSE	File System or Raw Device (e.g., some applications use raw devices)
FX	3390-3A/B/C 3390-3A/B/C VLL OPEN-x for FXoto, OPEN-x VLL for FXoto	Raw Device

Note 2: The command device (used for Command Control Interface (CCI) operations) is distinguished by **-CM** on the product name (e.g., OPEN-3-CM, OPEN-3-CVS-CM). The product name for VLL devices is OPEN-x-CVS, where CVS = custom volume size.

Note 3: This capacity is the maximum size which can be entered using the **lvcreate** command. The device capacity can sometimes be changed by the BIOS or host bus adapter. Also, different capacities may be due to variations such as 1 MB = 1000² or 1024² bytes.

Note 4: The number of blocks for a VLL 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 2) × (15 heads) × (96 sectors per track) = 76320

Note 5: The number of data cylinders for a Virtual LVI/LUN volume is calculated as follows ($\uparrow \dots \uparrow$ 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 = $\uparrow (\text{capacity (MB)} \times 1024/720) \uparrow$

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

$$\begin{aligned} \text{\# of cylinders} &= \uparrow 37 \times 1024/720 \uparrow = \uparrow 52.62 \uparrow \\ &= 53 \text{ cylinders} \end{aligned}$$

- Number of data cylinders for an OPEN-V VLL volume =
of cylinders = $\uparrow (\text{capacity (MB) specified by user}) \times 16/15 \uparrow$

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

$$\text{\# of cylinders} = \uparrow 50 \times 16/15 \uparrow = \uparrow 53.33 \uparrow = 54 \text{ cylinders}$$

- Number of data cylinders for a VLL LUSE volume (except for OPEN-V) =
of cylinders = $\uparrow (\text{capacity (MB)} \times 1024/720) \uparrow \times n$

Example: For OPEN-3 VLL LUSE volume with capacity = 37 MB and $n = 4$:

$$\text{\# of cylinders} = \uparrow 37 \times 1024/720 \uparrow \times 4 = \uparrow 52.62 \uparrow \times 4 = 53 \times 4 = 212$$

- Number of data cylinders for an OPEN-V VLL LUSE volume =
of cylinders = $\uparrow (\text{capacity (MB) specified by user}) \times 16/15 \uparrow \times n$

Example: For OPEN-V VLL LUSE volume with capacity = 50 MB and $n = 4$:

$$\text{\# of cylinders} = \uparrow 50 \times 16/15 \uparrow \times 4 = \uparrow 53.33 \uparrow \times 4 = 54 \times 4 = 216$$

- Number of data cylinders for a 3390-3A/C =
of cylinders = (number of cylinders) + 9

- Number of data cylinders for a 3390-3B VLL volume =
of cylinders = (number of cylinders) + 7

S1 = maximum **lvcreate** size value for VLL, LUSE, and VLL LUSE devices. Calculate the maximum size value (in MB) as follows: $S1 = (\text{PE Size}) \times (\text{Free PE})$. **Note:** Do not exceed the maximum **lvcreate** size value of 128 GB.

Note 6: 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. The user specifies the volume size using the Virtual LVI/LUN software.

Installation and Configuration Roadmap

The steps in [Table 1-4](#) outline the general process you follow to install and configure the Hitachi RAID storage system on a Red Hat Linux operating system.

Table 1-4 Installation and Configuration Roadmap

	Task
1.	Verify that the system on which you are installing the Hitachi RAID storage system meets the minimum requirements for this release.
2.	Prepare the Hitachi RAID storage system for the installation.
3.	Connect the Hitachi RAID storage system to a Red Hat Linux host.
4.	Configure the fibre-channel HBAs for the installation.
5.	Verify recognition of the new devices.
6.	Set the number of logical units.
7.	Partition the disk devices.
8.	Create file systems and mount directories, mount and verify the file systems, and set and verify auto-mount parameters.

Installing the Storage System

This chapter describes how to install the Hitachi RAID storage system on a Red Hat Linux operating system:

- [Requirements](#)
- [Preparing for the Storage System Installation](#)
- [Configuring the Fibre-Channel Ports](#)
- [Connecting the Storage System to the Red Hat Linux Host](#)
- [Configuring the Host Fibre-Channel HBA\(s\)](#)
- [Verifying New Device Recognition](#)

Requirements

[Table 2-1](#) lists and describes the requirements for installing the Hitachi RAID storage system on the HP-UX operating system.

Table 2-1 Requirements

Item	Requirements
Hitachi RAID storage system	<p>The availability of features and devices depends on the level of microcode installed on the Hitachi RAID storage system.</p> <p>Use LUN Manager software to configure the fibre-channel ports.</p>
Red Hat Linux AS/ES operating system	<p>Please refer to the Hitachi Data Systems interoperability site for specific support information for the Red Hat Linux operating system: http://www.hds.com/products/interoperability</p> <p>Root (superuser) login access to the host system.</p>
Red Hat Linux server	<p>Refer to the Red Hat Linux user documentation for server hardware and configuration requirements.</p>
Fibre-channel HBAs	<p>The Hitachi RAID storage systems support fibre-channel HBAs equipped as follows:</p> <ul style="list-style-type: none"> ▪ 8-Gbps fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors ▪ 4 Gbps fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors. ▪ 2 Gbps fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors. ▪ 1 Gbps fibre-channel interface, including shortwave non-OFC optical interface and multimode optical cables with SC connectors. <p>If a switch or HBA with a 1Gbps transfer rate is used, configure the device to use a fixed 1Gbps setting instead of Auto Negotiation. Otherwise, it may prevent a connection from being established.</p> <p>However, the transfer speed of CHF port cannot be set as 1 Gbps when the CHF is 8US/8UFC/16UFC. Therefore 1 Gbps HBA and switch cannot be connected.</p> <p>Do not connect OFC-type fibre-channel interfaces to the Hitachi RAID storage system. For information about supported fibre-channel HBAs, optical cables, hubs, and fabric switches, contact your Hitachi Data Systems account team.</p> <p>For information about supported HBAs, drivers, hubs, and switches, see the Hitachi Data Systems interoperability site: http://www.hds.com/products/interoperability</p>
Fibre-channel utilities and tools	<p>Refer to the documentation for your fibre-channel HBA for information about installing the utilities and tools for your adapter.</p>
Fibre-channel drivers	<p>Do not install/load the driver(s) yet. When instructed in this guide to install the drives for your fibre-channel HBA, refer to the documentation for your adapter.</p>

Preparing for the Storage System Installation

The following sections describe preinstallation considerations to follow before installing the Hitachi RAID storage system.

Hardware Installation Considerations

The Hitachi Data Systems representative performs the hardware installation by following the precautions and procedures in the Maintenance Manual.

Hardware installation activities include:

- Assembling all hardware and cabling
- Installing and formatting the logical devices (LDEVs). Be sure to obtain the desired LDEV configuration information from the user, including the desired number of OPEN-x, LUSE, VLL, VLL LUSE, and multiplatform (FX) devices.
- Installing the fibre-channel HBAs and cabling. The total fibre cable length attached to each fibre-channel adapter must not exceed 500 meters (1,640 feet).
 - Do not connect any OFC-type connectors to the storage system.
 - Do not connect/disconnect fibre-channel cabling that is being actively used for I/O. This can cause the Red Hat Linux system to hang.
 - Always confirm that the devices on the fibre cable are offline before connecting/disconnecting the fibre cable.
- Configuring the fibre port topology. The fibre topology parameters for each fibre-channel port depend on the type of device to which the port is connected, and the type of port. Determine the topology parameters supported by the device, and set your topology accordingly (see [Configuring the Fibre-Channel Ports](#)).

Before starting the installation, check all specifications to ensure proper installation and configuration.

LUN Manager Software Installation

The LUN Manager software on Storage Navigator is used to configure the fibre-channel ports. The user or Hitachi Data Systems representative installs the LUN Manager software. For instructions on installing LUN Manager, see the *Storage Navigator User's Guide*.

Setting the Host Mode

The Hitachi RAID storage system has host modes that the storage administrator must set for all new installations (newly connected ports) to Red Hat Linux hosts. The required host mode for Red Hat Linux is **00**. Do not select a host mode other than **00** for Red Hat Linux.

Use LUN Manager to set the host mode for each newly connected port. For instructions, see the *LUN Manager User's Guide* for the USP V/VM storage system or the *Provisioning Guide for Open Systems* for the VSP storage system.



Caution: Changing host modes on a Hitachi RAID storage system that is already installed and configured is disruptive and requires the server to be rebooted.

Setting the Host Mode Options

When each new host group is added, the storage administrator must be sure that the host mode options (HMOs) are set for all host groups connected to Red Hat Linux hosts. Use LUN Manager to set the HMOs. For instructions, see the *LUN Manager User's Guide* for the USP V/VM storage system or the *Provisioning Guide for Open Systems* for the VSP storage system.

The host mode options for Red Hat Linux are **2** and **13**. Select host mode option **2** when the conditions in [Table 2-2](#) are met. Alternatively, common host mode option **13** can be selected (see [Table 2-3](#)).



Caution: Changing host mode options on a Hitachi RAID storage system that is already installed and configured is disruptive and requires the server to be rebooted.

Table 2-2 Host Mode Option for Red Hat Linux

No.	Host Mode Option	Select if the Following Conditions are Met	Remarks
2	Veritas Database Edition™/Advanced Cluster	If you are using either: <ul style="list-style-type: none">▪ Veritas Database Edition™/Advanced Cluster for Real Application Clusters, or▪ Veritas Cluster Server™ 4.0 or later (I/O fencing function).	Mandatory. Do not apply this option to Sun™ Cluster.

Table 2-3 Common Host Mode Option

No.	Host Mode Option	Select if the Following Conditions are Met	Remarks
13	SIM report at link failure	When you want SIM notification when the number of link failures detected between ports exceeds the threshold.	Optional

Configuring the Fibre-Channel Ports

Use LUN Manager to configure the fibre-channel ports with the appropriate fibre parameters. You select the appropriate settings for each port based on the device to which the port is connected. Determine the topology parameters supported by the device, and set your topology accordingly.

The Hitachi RAID storage system supports up to 2048 logical units per fibre-channel port (512 per host group). Check your fibre-channel adapter documentation and your Linux system documentation to determine the total number of devices that can be supported.

[Table 2-4](#) explains the settings for defining port parameters. For instructions, see the *LUN Manager User's Guide* for the USP V/VM storage system or the *Provisioning Guide for Open Systems* for the VSP storage system.

Table 2-4 Fibre Parameter Settings

Fabric	Connection	Provides
Enable	FC-AL	FL-port (fabric port)
Enable	Point-to-Point	F-port (fabric port)
Disable	FC-AL	NL-port (private arbitrated loop)
Disable	Point-to-Point	<i>Not supported</i>



Note:

- If you plan to connect different types of servers to the Hitachi RAID storage system via the same fabric switch, use the zoning function of the fabric switch.
 - Contact Hitachi Data Systems for information about port topology configurations supported by HBA/switch combinations. Not all switches support F-port connection.
-

Port Address Considerations for Fabric Environments

In fabric environments, port addresses are assigned automatically by fabric switch port number and are not controlled by the port settings. In arbitrated loop environments, the port addresses are set by entering an AL-PA (arbitrated-loop physical address, or loop ID).

[Table 2-5](#) shows the available AL-PA values ranging from **01** to **EF**. Fibre-channel protocol uses the AL-PAs to communicate on the fibre-channel link, but the software driver of the platform host adapter translates the AL-PA value assigned to the port to a SCSI TID.

Table 2-5 Available AL-PA Values

EF	CD	B2	98	72	55	3A	25
E8	CC	B1	97	71	54	39	23
E4	CB	AE	90	6E	53	36	1F
E2	CA	AD	8F	6D	52	35	1E
E1	C9	AC	88	6C	51	34	1D
E0	C7	AB	84	6B	4E	33	1B
DC	C6	AA	82	6A	4D	32	18
DA	C5	A9	81	69	4C	31	17
D9	C3	A7	80	67	4B	2E	10
D6	BC	A6	7C	66	4A	2D	0F
D5	BA	A5	7A	65	49	2C	08
D4	B9	A3	79	63	47	2B	04
D3	B6	9F	76	5C	46	2A	02
D2	B5	9E	75	5A	45	29	01
D1	B4	9D	74	59	43	27	
CE	B3	9B	73	56	3C	26	

Loop ID Conflicts

The Red Hat Linux operating system assigns port addresses from lowest (**01**) to highest (**EF**). To avoid loop ID conflict, assign the port addresses from highest to lowest (i.e., starting at **EF**). The AL-PAs should be unique for each device on the loop to avoid conflicts. Do not use more than one port address with the same TID in same loop (e.g., addresses **EF** and **CD** both have TID 0, see [Table 2-5](#)).

Connecting the Storage System to the Red Hat Linux Host

After you prepare the hardware, software, and fibre-channel HBA(s), connect the Hitachi RAID storage system to the Red Hat Linux system.

[Table 2-6](#) summarizes the steps for connecting Hitachi RAID storage system to the Red Hat Linux system host. Some steps are performed by the Hitachi Data Systems representative, while others are performed by the user.

Table 2-6 Connecting the Storage System to the Red Hat Linux Host

	Activity	Performed by	Description
1.	Verify storage system installation	Hitachi Data Systems representative	Confirm that the status of the fibre-channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the Red Hat Linux system	User	Power off the Red Hat Linux system before connecting the Hitachi RAID storage system. <ul style="list-style-type: none">▪ Shut down the Red Hat Linux system.▪ When shutdown is complete, power off the Red Hat Linux display.▪ Power off all peripheral devices except for the Hitachi RAID storage system.▪ Power off the host system. You are now ready to connect the Hitachi RAID storage system.
3.	Connect the Hitachi RAID storage system	Hitachi Data Systems representative	Install fibre-channel cables between the storage system and the Red Hat Linux system. Follow all precautions and procedures in the Maintenance Manual. Check all specifications to ensure proper installation and configuration.
4.	Power on the Red Hat Linux system	User	Power on the Red Hat Linux system after connecting the Hitachi storage system: <ul style="list-style-type: none">▪ Power on the Red Hat Linux system display.▪ Power on all peripheral devices. The Hitachi RAID storage system should be on, and the fibre-channel ports should be configured. If the fibre ports are configured after the Linux system is powered on, restart the system to have the new devices recognized.▪ Confirm the ready status of all peripheral devices, including the Hitachi RAID storage system.▪ Power on the Red Hat Linux system.
5	Boot the Red Hat Linux system		

Configuring the Host Fibre-Channel HBA(s)

Configure the fibre-channel HBA(s) connected to the Hitachi RAID storage system. The HBAs have many configuration options. This section provides the following minimum requirements for configuring host fibre-channel adapters for operation with the Hitachi RAID storage system. For more information, refer to the documentation for your fibre-channel HBA(s).

- The queue depth requirements for the devices on the Hitachi RAID storage system are specified in [Table 2-7](#). 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.
- The BIOS might need to be disabled to prevent the system from trying to boot from the Hitachi RAID storage system.

Several other parameters (e.g., FC, fabric) may also need to be set. Please refer to the user documentation that came with your HBA to determine whether other options are required to meet your operational requirements.



Note: Use the same settings and device parameters for all devices on the Hitachi RAID storage system.

Table 2-7 Queue Depth Requirements

Parameter	Required Value
IOCB Allocation (Queue depth) per LU	≤ 32 per LU
IOCB Allocation (Queue depth) per port (MAXTAGS)	≤ 2048 per port

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 2-1](#)).
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 2-8](#) 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 2-2](#)).

```
# 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 2-1 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 2-2 Example of Verifying Device Files

Table 2-8 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 Disk Devices

This chapter 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 operating system. Check your fibre-channel adapter documentation and Linux system documentation to ascertain the total number of devices that can be supported.

2. To set the Emulex Driver, as shown in [Figure 3-3](#), add the following line to the `/etc/modules.conf` file:

```
Alias scsi_hostadapter lpfcd
```

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 3-1](#), 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 3-2](#).

5. Reboot the system.

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

← *Comment out this line.*
← *Add this line.*

Figure 3-1 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
initrd /boot/initrd-2.4.x.scsiluns.img
```

← *This line is commented out.*
← *Add this line.*

Figure 3-2 Example of Setting the Number of LUs (GRUB)

```
Alias scsi_hostadapter lpfcd
```

← *Add this to /etc/modules.conf*

Figure 3-3 Example of Setting the Emulex Driver

Partitioning the Devices

After the setting the number of logical units, use the following procedure to set 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.



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

5. Repeat steps 1 through 4 for each new 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 /USP-LU00
```

Mounting the New File Systems

Use the **mount** command to mount each new file system (see example in [Figure 3-4](#)). 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 3-4](#).

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

Figure 3-4 Example of Mounting the New Devices

Verifying the File Systems

After mounting the file systems, verify the file systems (see the example in [Figure 3-5](#)).

```
# 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%    /USP-LU00
#
```

Figure 3-5 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 3-6](#)).

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

Figure 3-6 Example of Setting Auto-Mount Parameters

Failover and SNMP Operation

The Hitachi RAID storage systems support industry-standard products and functions that provide host and/or application failover, I/O path failover, and logical volume management (LVM). The Hitachi RAID storage systems also support the industry-standard simple network management protocol (SNMP) for remote storage system management from the UNIX/PC server host. SNMP is used to transport management information between the storage system and the SNMP manager on the host. The SNMP agent sends status information to the host(s) when requested by the host or when a significant event occurs.

This chapter describes how failover and SNMP operations are supported on the Hitachi RAID storage system:

- [Host Failover](#)
- [Path Failover](#)
- [SNMP Remote System Management](#)



Note: The user is responsible for configuring the failover and SNMP management software on the UNIX/PC server host. For assistance with failover and/or SNMP configuration on the host, refer to the user documentation, or contact the vendor's technical support.

Host Failover

The Hitachi RAID storage systems support the Veritas Cluster Server and host failover products for the Red Hat Linux operating system. The user must be sure to configure the host failover software and any other high-availability (HA) software as needed to recognize and operate with the newly attached devices.

For assistance with Veritas Cluster Server operations, refer to the Veritas user documentation, see [Note on Using Veritas Cluster Server](#), or contact Symantec technical support. For assistance with specific configuration issues related to the Hitachi RAID storage system, please contact your Hitachi Data Systems representative.

Path Failover

The Hitachi RAID storage systems support the Hitachi HiCommand Dynamic Link Manager (HDLM) and Veritas Volume Manager for the Red Hat Linux operating system. For further information, see the *Hitachi Dynamic Link Manager™ for Red Hat® Linux® User's Guide*. For assistance with Veritas Volume Manager operations, refer to the Veritas user documentation or contact Symantec technical support.

SNMP Remote System Management

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 Red Hat Linux server host via a notebook PC (see [Figure 4-1](#)). 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 Red Hat Linux server. Notification of error conditions is made in real time, providing the Red Hat Linux server user with the same level of monitoring and support available to the mainframe user. The SIM reporting via SNMP enables the user to monitor the Hitachi RAID storage system from the Red Hat Linux server 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.



Note: The user is responsible for configuring the SNMP manager on the Red Hat Linux server host. For assistance with SNMP manager configuration on the Red Hat Linux server host, refer to the user documentation, or contact the vendor's technical support.




Figure 4-1 SNMP Environment

Troubleshooting

This chapter provides troubleshooting information for Red Hat Linux host attachment and instructions for calling technical support.

- [General Troubleshooting](#)
- [Calling the Hitachi Data Systems Support Center](#)

General Troubleshooting

[Table 5-1](#) lists potential error conditions that may occur during storage system installation and provides instructions for resolving each condition. If you cannot resolve an error condition, please contact your Hitachi Data Systems representative for help, or call the Hitachi Data Systems Support Center for assistance.

For troubleshooting information on the Hitachi RAID storage system, see the User and Reference Guide for the storage system (e.g., *Hitachi Virtual Storage Platform User and Reference Guide*).

For troubleshooting information on Hitachi Storage Navigator, see the Storage Navigator User's Guide for the storage system (e.g., *Hitachi Virtual Storage Platform Storage Navigator User Guide*).

For information on errors messages displayed by Storage Navigator, see the Storage Navigator Messages document for the storage system (e.g., *Hitachi Virtual Storage Platform Storage Navigator Messages*).

Table 5-1 Troubleshooting

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 /etc/fstab file is correct.

Calling the Hitachi Data Systems Support Center

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

- The circumstances surrounding the error or failure.
- The exact content of any error messages displayed on the host system(s).
- The exact content of any error messages displayed by Storage Navigator.
- The Storage Navigator configuration information (use the FD 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. If you need technical support, log on to the Hitachi Data Systems Portal for contact information: <https://hdssupport.hds.com>

Note on Using Veritas Cluster Server

By issuing a SCSI-3 Persistent Reserve command for a Hitachi RAID storage system, the 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 group. 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 one storage system is 128 for a port. 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 128, key registration as well as operations such as installing an LU to the disk group fail. To avoid failure of reserve key registration, the number of reserve keys needs to be kept below 128. 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 A-1 Adding Reserve Keys for LUs to Increase Disk Capacity



Acronyms and Abbreviations

AL	arbitrated loop
AL-PA	arbitrated loop physical address
blk	block
CVS	custom volume size
FC	fibre-channel
FCP	fibre-channel protocol
FX	Hitachi Cross-OS File Exchange
GB	gigabytes
Gbps	gigabits per second
GRUB	Grand Unified Bootloader
HBA	host bus adapter
HDLM	Hitachi Dynamic Link Manager
I/O	input/output
LU	logical unit
LUN	logical unit, logical unit number
LUSE	LUN Expansion
LVI	logical volume image
LVM	Logical Volume Manager
MB	megabytes
MPE	maximum number of physical extents
OFC	open fibre control
PA	physical address
PC	personal computer
PP	physical partition
RAID	redundant array of independent disks
SCSI	small computer system interface
SIM	service information message
SNMP	simple network management protocol
TCO	total cost of ownership

TID	target ID
USP V/VM	Hitachi Universal Storage Platform V/VM
VLL	Virtual LVI/LUN
VSP	Hitachi Virtual Storage Platform
WWN	worldwide name

Hitachi Data Systems

Corporate Headquarters

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

Asia Pacific and Americas

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

Europe Headquarters

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



MK-96RD640-03