HITACHI
Inspire the Next

Hitachi Universal Storage Platform V/VM

Configuration Guide for SGI® IRIX® Host Attachment

FASTFIND LINKS

Hitachi Data Systems

MK-96RD651-01

Contents

Preface

Welcome to the Configuration Guide for SGI® IRIX® Host Attachment. This document provides information and instructions for installing, configuring, and operating the Hitachi Universal Storage Platform V storage system (USP V) in an SGI IRIX environment.

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:

Note: The Storage Navigator windows shown in this document were captured on a Windows® system with the Internet Explorer web browser. The Storage Navigator windows may display differently on other operating systems and browsers. Please refer to the *Storage Navigator User's Guide* for information on other supported operating systems and browsers.

Notice: The use of the Hitachi USP V storage system and all other Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

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 Universal Storage Platform V storage system in an SGI IRIX host environment. To use this document, you should have a working knowledge of the following:

- The Hitachi Universal Storage Platform V storage system
- The SGI IRIX operating system and the UNIX® file system, system commands, and utilities
- The hardware hosting the SGI IRIX system
- The hardware used to attach the Hitachi Universal Storage Platform V to the SGI IRIX host, including fibre-channel cabling, host bus adapters (HBAs), switches, and hubs

Product Version

This document revision applies to Universal Storage Platform V microcode 60-01-3*x* and higher.

Document Revision Level

Revision	Date	Description
MK-96RD651-P	May 2007	Preliminary Release
MK-96RD651-00	September 2007	Initial Release, supersedes and replaces MK-96RD651-P
MK-96RD651-01	November 2007	MK-96RD651-01, supersedes and replaces MK-96RD651-00

Changes in this Revision

- Added support for USP VM.
- In Table 2-2, indicated that host mode option 13 is a common host mode option.

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
Introduction	Provides a brief overview of the Hitachi USP V, supported device types, and an installation roadmap.
Installing the USP V	Provides instructions for installing and connecting the USP V to an SGI IRIX host.
Configuring the New Disk Devices	Provides instructions for configuring the new devices on the USP V for use.
Failover and SNMP Operation	Describes how to configure the USP V for failover and SNMP.
Troubleshooting	Provides information for identifying and resolving problems.
SCSI TID Maps for Fibre-Channel Adapters	Describes SCSI TID Maps for fibre-channel adapters.
Online Device Installation	Provides instructions for online installation of new devices.
Acronyms and Abbreviations	Defines the acronyms and abbreviations used in this document.

Referenced Documents

- *Hitachi Universal Storage Platform V Storage Navigator User's Guide*, MK-96RD621)
- *Hitachi Universal Storage Platform V LUN Manager User's Guide*, MK-96RD615
- *Hitachi Universal Storage Platform V LUN Expansion (LUSE) and Virtual LVI/LUN User's Guide*, MK-96RD630
- *Hitachi Universal Storage Platform V Cross-OS File Exchange User's Guide*, MK-96RD647

Note: For more information about SGI IRIX, refer to the IRIX user documentation or contact SGI customer support services.

Document Conventions

The terms "Universal Storage Platform V" and "USP V" refer to all models of the Universal Storage Platform V, 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, and labels.
Italic	Indicates a variable.
`screen/code`	Indicates text displayed or entered on screen or at the command prompt.
boldface screen font	Information you must enter is in boldface screen font.
< italic screen >	Variables appear in italic screen font between angle brackets.
[ ]	Default responses to system prompts are in square brackets.
#	A **pound sign** at the beginning of a line indicates an operating system command line prompt.

This document uses icons to draw your attention to certain information. In order of how critical the information is to your system, these items will be marked as note, tip, caution, WARNING, or DANGER.

Icon	Meaning	Description
	Note	Notes call attention to important and/or additional information.
	Tip	Tips provide helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Cautions notify the user of adverse conditions and/or consequences (e.g., disruptive operations).
	WARNING	Warnings notify the user of severe conditions and/or consequences (e.g., destructive operations).
	DANGER	Dangers provide information about how to avoid physical injury to yourself and others.
	ELECTRIC SHOCK HAZARD!	This symbol warns the user of the possibility of electric shock. Failure to take appropriate precautions (e.g., do not touch) could result in serious injury.
	ESD Sensitive	This symbol indicates that the hardware is sensitive to electrostatic discharge (ESD). Failure to take appropriate precautions (e.g., grounded wriststrap) could result in damage to the hardware.

Icon	Meaning	Description
	HOT SURFACE!	This symbol indicates that the surface of the hardware is hot. Failure to take appropriate precautions (e.g., turn off power and allow to cool before touching) could result in injury.
	WARNING! Sharp edges or corners.	This symbol indicates that the hardware has sharp edges and/or corners. Failure to take appropriate precautions (e.g., avoid touching or wear gloves) could result in injury.

Convention for Storage Capacity Values

The RAID storage system calculates physical storage capacity values (e.g., disk drive capacity) based on the following values:

1 KB = 1,000 bytes
1 MB = $1{,}000^2$ bytes
1 GB = $1{,}000^3$ bytes
1 TB = $1{,}000^4$ bytes
1 PB = $1{,}000^5$ bytes

The RAID storage system calculates logical storage capacity values (e.g., logical device capacity) based on the following values:

1 KB = 1,024 bytes
1 MB = $1{,}024^2$ bytes
1 GB = $1{,}024^3$ bytes
1 TB = $1{,}024^4$ bytes
1 PB = $1{,}024^5$ bytes
1 block = 512 bytes

Getting Help

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

- The circumstances surrounding the error or failure.
- The exact content of any error message(s) displayed on the host system(s).
- The data in the CCI error log file and trace data (all files in HORCM_LOG directory).
- The remote service information messages (R-SIMs), including reference codes and severity levels, displayed by Storage Navigator and/or HiCommand Device Manager.

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

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

Comments

Please send us your comments on this document. Make sure to include the document title, number, and revision. Please refer to specific section(s) and paragraph(s) whenever possible.

- **E-mail:** doc.comments@hds.com
- **Fax:** 858-695-1186
- **Mail:**
  Technical Writing, M/S 35-10
  Hitachi Data Systems
  10277 Scripps Ranch Blvd.
  San Diego, CA 92131

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

1

Introduction

This chapter provides an overview of the USP V from Hitachi Data Systems:

- ☐ About the Hitachi Universal Storage Platform V
- ☐ Device Types
- ☐ Installation and Configuration Roadmap

About the Hitachi Universal Storage Platform V

The USP V from Hitachi Data Systems is a comprehensive and fully integrated universal storage services platform that builds on the proven technical superiority and features of Hitachi Data Systems' USP disk arrays. As an evolution of the USP solutions, the USP V combines refinements to existing technology with new functionality to deliver higher system performance, increased scalability, and greater levels of availability.

The USP V provides heterogeneous data replication and copy services between any platform at any distance, with higher performance, increased data protection, and more efficient use of network bandwidth than previously available. The result is reduced costs, simplified management, and improved data protection and disaster recovery capability. The following list summarizes the USP V's key benefits and the advantages they deliver.

- **Massive multi-protocol consolidation**. The USP V provides an integrated platform for consolidating and centralizing data, while providing access to hundreds or thousands of servers with differing connectivity requirements. USP V support for fibre channel, ESCON, and FICON - coupled with logical partitioning, virtual storage ports, and centralized management with HiCommand Software - provides the most comprehensive and flexible platform for massive consolidation of storage systems with multi-protocol requirements, while delivering the software tools to monitor, manage, and optimize the environment to ensure application quality of service.
- **Simplified storage management**. Embedded second-generation virtualization features allow multiple storage systems to be aggregated (or consolidated) into a single storage pool to reduce storage management complexities and costs. Customers can now manage the entire infrastructure from a single pane of glass, or have multiple storage administrators manage particular sets of storage resources.
- **Tiered storage foundation for data lifecycle management**. The USP V reduces costs and improves efficiency by allowing heterogeneous tiered storage to be managed from a single console and moved non-disruptively between storage tiers as the value of the data or the usage model changes. With the USP V, a tiered storage infrastructure can be built internally or in combination with externally attached storage, including FICON for IBM z/OS mainframe operating systems.
- **Investment protection**. USP V customers can enjoy a storage technology refresh while protecting their current storage system investment by attaching their legacy storage to the USP V externally and using powerful software tools such as unique virtualization, storage partitioning, and universal replication with their legacy storage. Customers can non-disruptively migrate their data to the USP V, or any other platform within the USP V's external storage pool (including FICON for IBM z/OS customers), at the appropriate time.

- **Reduced Total Cost of Ownership (TCO) with improved data availability**. A new thin provisioning software feature improves data availability while reducing storage TCO by allowing customers to allocate "virtual" disk storage based on their anticipated future needs, with less physical disk than initially required. Additional physical disks can be purchased later and installed transparently without interrupting application services.
- **Increased Return on Investment (ROI)**. The USP V increases ROI by providing investment protection and longer life for legacy storage, superior storage and server utilization, improved application quality of service, lower software licenses/costs, and simplified and centralized storage management.
- **Reduced mainframe storage costs**. The USP V lets mainframe customers exploit SATA storage either internal to the USP V or externally attached to improve data protection, cost-effectively store larger amounts DFSMS Migration Level 1 data and offload costly DFSMS mainframe processing cycles, or deploy a long-term and secure data archive solution for regulatory compliance. The USP V's virtual private storage machines are the ideal compliment to mainframe LPARs for providing workload isolation and optimum quality of service.
- **Application data mobility solutions**. The USP V is complemented by a rich suite of software solutions for managing storage across an entire heterogeneous storage pool:
  - HiCommand Tiered Storage Manager - lets customers continually adjust their storage infrastructure to meet their changing business needs.
  - HiCommand Tuning Manager and HiCommand Device Manager – let customers continually monitor and forecast their storage infrastructure needs.
  - Hitachi ShadowImage In-System Replication and Hitachi Universal Replication - deliver proven scalable data-protection for meeting the most stringent business continuance and compliance requirements of the world's largest financial institutions.
  - Hitachi Universal Replication and Hitachi ShadowImage now support more volumes and replication pairs than before. Performance has also been improved, with support added for more consistency groups spread across four USP V systems located at each site within a 2- or 3-datacenter configuration.

Note: The availability of the Hitachi USP V features and functions depends on the level of microcode installed.

Device Types

Table 1-1 describes the types of logical devices (volumes) that can be installed and configured for operation with the USPV on an SGI IRIX operating system.

Table 1-2 lists the specifications for devices supported by the USP V. For information about configuring devices other than OPEN+V, contact your Hitachi Data Systems representative.

Note: Logical devices are defined to the host as SCSI disk devices, even though the interface is fibre channel.

Table 1-1 Types of Logical Devices Supported by the USP V

Device Type	Description
OPEN-x Devices	OPEN-*x* logical units (LUs) (e.g., OPEN-3, OPEN-9) are disk devices of predefined sizes. USP V supports OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-L, and OPEN-V devices (OPEN-V is a VLL-based volume and has no standard size). For the latest information about supported LU types, please 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 USP V using fewer LU numbers. For more information about LUSE, see the *Hitachi Universal Storage Platform V LUN Expansion (LUSE) and Virtual LVI/LUN User's Guide* (MK-96RD630).
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. For more information about the Virtual LVI/LUN feature, see the *Hitachi Universal Storage Platform V LUN Expansion (LUSE) and Virtual LVI/LUN User's Guide* (MK-96RD630). 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).

Device Type	Description
FX Devices (3390-3A/B/C, OPEN-x-FXoto)	Using USP V's Cross-OS File Exchange (FX) feature, you can 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 *Hitachi Cross-OS File Exchange User's Guide* (MK-96RD647), or contact your Hitachi Data Systems account team. 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). The 3390-3B devices are write-protected from UNIX/PC server access. The USP V rejects all UNIX/PC server write operations (including fibre-channel adapters) for 3390-3B devices. 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.

Table 1-2 USP V Device Specifications (continues on the next page)

Device Type *(Note 1)*	Category *(see Note 2)*	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Data Cylinders	# of Heads	# of Sectors per Track	Capacity MB *(Note 3)*
OPEN-3	SCSI Disk	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
OPEN-8	SCSI Disk	HITACHI	OPEN-8	14351040	512	9966	15	96	7007
OPEN-9	SCSI Disk	HITACHI	OPEN-9	14423040	512	10016	15	96	7042
OPEN-E	SCSI disk	HITACHI	OPEN-E	28452960	512	19759	15	96	13893
OPEN-L	SCSI disk	HITACHI	OPEN-L	71192160	512	49439	15	96	34761
OPEN-V	SCSI disk	HITACHI	OPEN-V	125827200 max ***Note 4***	512	***Note 5***	15	128	***Note 6***
OPEN-3*n	SCSI Disk	HITACHI	OPEN-3*n	4806720*n	512	3338*n	15	96	2347*n
OPEN-8*n	SCSI Disk	HITACHI	OPEN-8*n	14351040*n	512	9966*n	15	96	7007*n
OPEN-9*n	SCSI Disk	HITACHI	OPEN-9*n	14423040*n	512	10016*n	15	96	7042*n
OPEN-E*n	SCSI disk	HITACHI	OPEN-E*n	28452960*n	512	19759*n	15	96	13893*n
OPEN-L*n	SCSI disk	HITACHI	OPEN-L*n	71192160*n	512	49439*n	15	96	34761*n
OPEN-V*n	SCSI disk	HITACHI	OPEN-V*n	125827200 max ***Note 4***	512	***Note 5***	15	128	***Note 6***
OPEN-3 VLL	SCSI Disk	HITACHI	OPEN-3-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-8 VLL	SCSI Disk	HITACHI	OPEN-8-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-9 VLL	SCSI Disk	HITACHI	OPEN-9-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-E VLL	SCSI disk	HITACHI	OPEN-E-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-V VLL	SCSI disk	HITACHI	OPEN-V	***Note 4***	512	***Note 5***	15	128	***Note 6***
OPEN-3*n VLL	SCSI Disk	HITACHI	OPEN-3*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-8*n VLL	SCSI Disk	HITACHI	OPEN-8*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
OPEN-9*n VLL	SCSI Disk	HITACHI	OPEN-9*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***

Table 1-2 USP V Device Specifications (continued)

Device Type *(Note 1)*	Category *(see Note 2)*	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Data Cylinders	# of Heads	# of Sectors per Track	Capacity MB *(Note 3)*
OPEN-V*n VLL	SCSI disk	HITACHI	OPEN-V*n-CVS	***Note 4***	512	***Note 5***	15	128	***Note 6***
OPEN-E*n VLL	SCSI disk	HITACHI	OPEN-E*n-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***
3390-3A	FXotm/ mto	HITACHI	3390-3A	5820300	512	3345	15	116	2844
3390-3B	FXmto	HITACHI	3390-3B	5816820	512	3343	15	116	2842
3390-3C	FXotm	HITACHI	OP-C-3390-3C	5820300	512	3345	15	116	2844
FX OPEN-3	FXoto	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
3390-3A VLL	FXotm/ mto	HITACHI	3390-3A-CVS	***Note 4***	512	***Note 5***	15	116	***Note 6***
3390-3B VLL	FXmto	HITACHI	3390-3B-CVS	***Note 4***	512	***Note 5***	15	116	***Note 6***
3390-3C VLL	FXotm	HITACHI	3390-3C-CVS	***Note 4***	512	***Note 5***	15	116	***Note 6***
FX OPEN-3 VLL	FXoto	HITACHI	OPEN-3-CVS	***Note 4***	512	***Note 5***	15	96	***Note 6***

Note 1: The availability of a specific USP V device type depends on the level of microcode installed on the USP V system.

Note 2: 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, VLL LUSE) are usually formatted with file systems for SGI IRIX operations. The FX devices (3390-3A/B/C, OPEN-*x*-FXoto) must be installed as raw devices and can only accessed using FX. Do not create a partition or 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 (SCSI disk devices can also be used as raw devices. For example, some database applications use raw devices.	File System*
FX	3390-3A/B/C 3390-3A/B/C VLL OPEN-x for FXoto, OPEN-x VLL for FXoto	Raw Device

* The SCSI disk devices can also be used as raw devices (e.g., some database applications use raw devices).

Note 3: The BIOS or HBA can sometimes change the device capacity. These device capacities are calculated based on 1 MB = 1024^2 bytes rather than 1000^2 bytes.

Note 4: 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

Note 5: 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):

The number of data cylinders for an OPEN-*x* VLL volume = (except for OPEN-V):

# of cylinders = ↑ (capacity (MB) × 1024/720 ↑

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

# of cylinders = ↑37 × 1024/720↑ = ↑52.62↑
= 53 cylinders

The number of data cylinders for an OPEN-V VLL volume =

# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑

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

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

The number of data cylinders for a VLL LUSE volume = (except for OPEN-V):

# of cylinders = ↑ (capacity (MB) × 1024/720 ↑ × n

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

# of cylinders = ↑37 × 1024/720↑ × 4 = ↑52.62↑ × 4 = 53 × 4 = 212

The number of data cylinders for an OPEN-V VLL LUSE volume =

# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑ × n

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

# of cylinders = ↑50 × 16/15↑ × 4 = ↑53.33↑× 4 = 54 × 4 = 216

The number of data cylinders for a 3390-3A/C =

# of cylinders = (number of cylinders) + 9

The number of data cylinders for a 3390-3B VLL volume =

# of cylinders = (number of cylinders) + 7

Note 6: The size of an OPEN-3/8/9/E/L 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. For more information, please see *Hitachi Universal Storage Platform V LUN Expansion and Virtual LVI/LUN User's Guide* (MK-96RD630).

Installation and Configuration Roadmap

The steps in Table 1-4 and Figure 1-1 outline the general process you follow to install and configure the USP V on an SGI IRIX operating system.

Table 1-4 Installation and Configuration Roadmap

	Task
1.	Verify that the system on which you are installing the USP V meets the minimum requirements for this release.
2.	Prepare the USP V for the installation.
3.	Prepare the fibre-channel HBAs for the installation.
4.	Connect the USP V to an SGI IRIX host.
5.	Verify recognition of the new devices.
6.	Partition the disk devices and enable command tag queuing.
7.	Create the file systems and mount directories, and mount the devices,.
8.	Verify the file systems and set the auto-mount parameters.

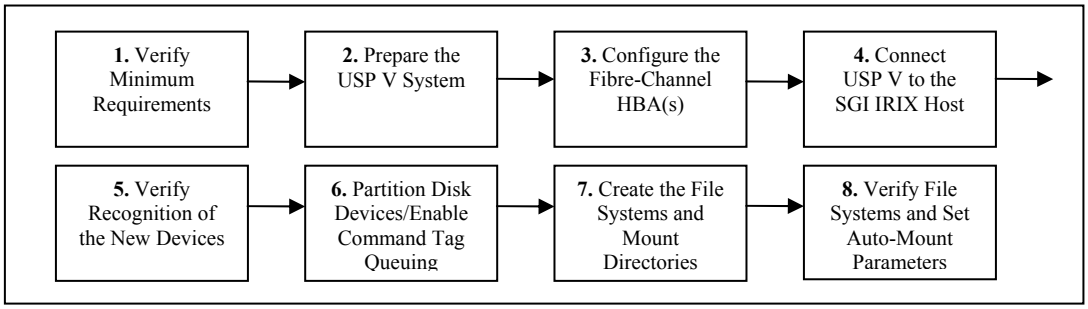

Figure 1-1 Installation and Configuration Roadmap

2

Installing the USP V

This chapter describes how to install the USP V on an SGI IRIX operating system:

- ☐ Installation and Configuration Checklist
- ☐ Preparing for the USP V Installation
- ☐ Configuring the Host Fibre-Channel HBA(s)
- ☐ Connecting the USP V System to the SGI IRIX Host
- ☐ Verifying New Device Recognition

Installation and Configuration Checklist

Table 2-1 is a checklist for installing and configuring the USP V on an SGI IRIX operating system. The two columns on the left indicate requirements to be met, tasks to be completed, and information to be obtained before you perform the installation. Use the column on the right to confirm that you have the required minimum requirements and completed the task. You can also write any notes about a requirement.

Table 2-1 Installation and Configuration Checklist

Installation Requirements	Description	Requirements Confirmed?
Hitachi USP V system, all-open or multiplatform configuration	The availability of features and devices depends on the level of microcode installed on the USP V. Use LUN Manager software to configure the fibre-channel ports. If this feature is not installed, contact your Hitachi Data Systems account team.	
SGI IRIX system hardware	▪ SGI O2 ▪ Octane® ▪ Onyx2® ▪ Challenge® ▪ Origin® 200 Please contact your Hitachi Data Systems account team for further information on server hardware requirements.	
SGI IRIX operating system, version 6.5.*x* (6.5.16 and 6.5.24 are under evaluation)	Root login access to the SGI IRIX system is required. Please contact Silicon Graphics® to make sure that the most current OS patches are installed on the SGI systems(s). For the latest information about SGI IRIX version support, contact your Hitachi Data Systems account team.	
Fibre-channel HBAs	The USP V supports fibre-channel HBAs equipped as follows: ▪ 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. 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. Do not connect OFC-type fibre-channel interfaces to the USP V. For information about supported fibre-channel HBAs, optical cables, hubs, and fabric switches, contact your Hitachi Data Systems account team.	
Fibre-channel utilities and tools	Refer to the documentation for your fibre-channel HBA for information about installing the utilities and tools for your adapter.	

Fibre-channel drivers	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.	

Preparing for the USP V Installation

The following sections describe pre-installation considerations to follow before installing the USP V.

Hardware Installation Considerations

The Hitachi Data Systems representative performs the hardware installation by following the precautions and procedures in the USP V Maintenance Manual. The USP V comes with all the hardware and cabling required for the installation.

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 USP V system.
  - Do not connect/disconnect fibre-channel cabling that is being actively used for I/O. This can cause the SGI IRIX 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 USP V 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 USP V Fibre-Channel Ports).

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

LUN Manager Software Installation

LUN Manager software is used to configure the USP V fibre-channel ports. The user or Hitachi Data Systems representative installs the LUN Manager software. For instructions about installing the LUN Manager software, see the *Hitachi Universal Storage Platform V Storage Navigator User's Guide* (MK-96RD621). For more information about LUN Manager, see the *Hitachi Universal Storage Platform V LUN Manager User's Guide* (MK-96RD615).

Note: If the LUN Manager feature is not installed, the Hitachi Data Systems representative can configure fibre-channel ports for you. Contact your Hitachi Data Systems account team.

Setting the Host Mode

The USP V has host modes that the storage administrator must set for all new installations (newly connected ports) to IRIX hosts.

- The required host mode for SGI IRIX is: **00**

Do not select a host mode other than **00** for SGI IRIX.

You can use either the LUN Manager software (Storage Navigator) or the HiCommand® Device Manager software to set the host mode. Figure 2-1 shows the Add New Host Group panel from LUN Manager. For more information about using LUN Manager, see the *Hitachi Universal Storage Platform V LUN Manager User's Guide* (MK-96RD615).

WARNING: Changing host modes on USP V systems that are already installed and configured is disruptive and requires the server to be rebooted.

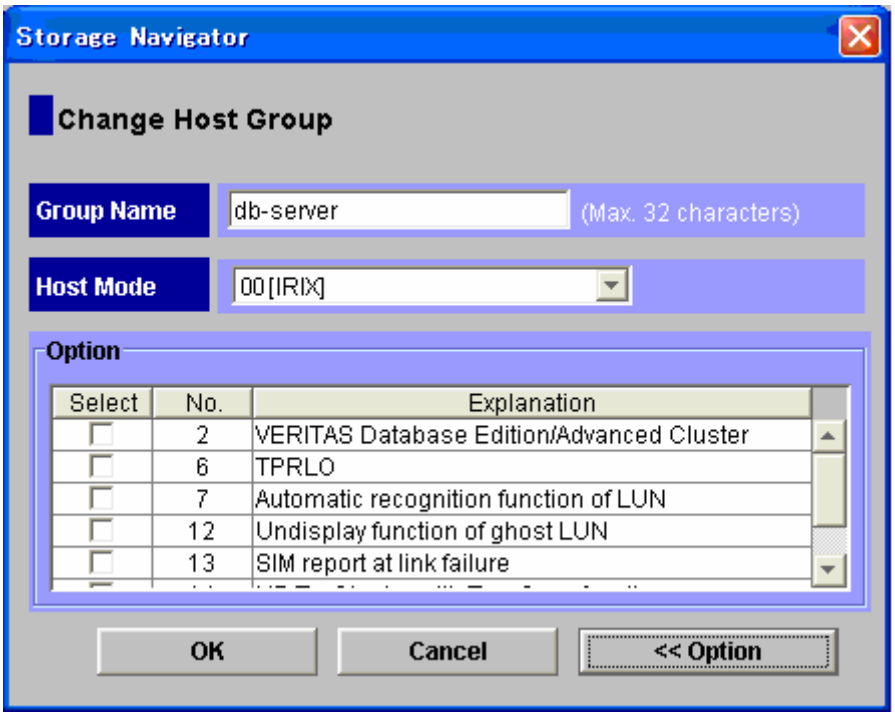

Figure 2-1 Example of Setting the Host Mode

Setting the Host Mode Options

When each new host group is added, the storage administrator must be sure that all host mode options are set for all host groups connected to SGI IRIX SGI IRIX hosts.

Select the common host mode option **13** when the conditions in Table 2-2 are met.

WARNING: Changing host mode options on USP V systems that are already installed and configured is disruptive and requires the server to be rebooted.

Table 2-2 Common Host Mode Options for the USP V

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

You can use either the LUN Manager software (Storage Navigator) or the HiCommand® Device Manager software to set the host mode options. For more information about using LUN Manager, see the *Hitachi Universal Storage Platform V LUN Manager User's Guide* (MK-96RD615).

Configuring the USP V Fibre-Channel Ports

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

Figure 2-2 shows the LUN Manager panel for defining port parameters, and Table 2-3 explains the panel settings. For more information, see the *Hitachi Universal Storage Platform V LUN Manager User's Guide* (MK-96RD615).

Table 2-3 Fibre Parameter Settings for the USP V

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	*Not supported*

Note: The USP V supports up to 2048 LUs per fibre-channel port.

Note: If you plan to connect different types of servers to the USP V via the same fabric switch, use the zoning function of the fabric switch.

Note: Contact Hitachi Data Systems for information about port topology configurations supported by HBA/switch combinations. Not all switches support F-port connection.

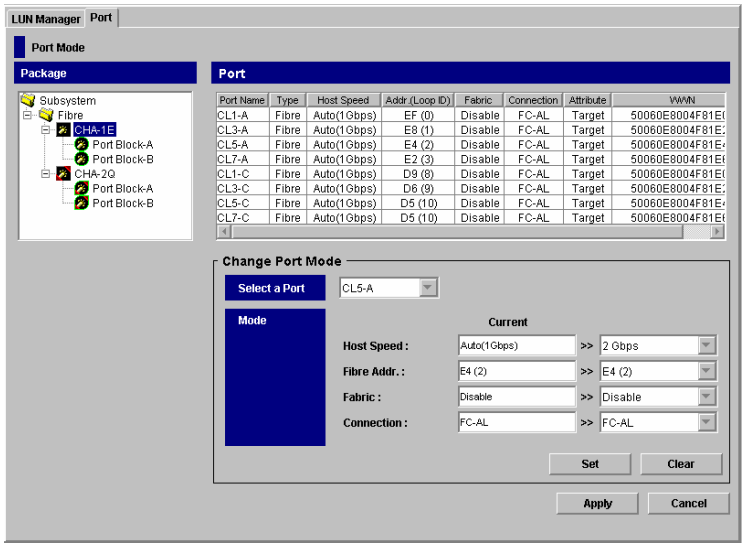

Figure 2-2 Setting the Fibre-Channel Port Parameters

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 USP V port settings. In arbitrated loop environments, the port addresses are set by entering an AL-PA (arbitrated-loop physical address, or loop ID).

Table 2-4 shows the available USP V 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 USP V port to a SCSI TID. See Appendix A for a description of the AL-PA-to-TID translation.

Table 2-4 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 SGI IRIX 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 Appendix A for the TID-to-AL-PA mapping).

Configuring the Host Fibre-Channel HBA(s)

Configure the fibre-channel HBA(s) connected to the USP V. The HBAs have many configuration options. For more information, refer to the documentation for your fibre-channel HBA(s).

Connecting the USP V System to the SGI IRIX Host

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

Table 2-5 summarizes the steps for connecting the USP V to the SGI system host. Some steps are performed by the Hitachi Data Systems representative, while others are performed by the user.

Table 2-5 Steps for Connecting USP V to an SGI IRIX Host

	Activity	Performed by	Description
1.	Verify USP V installation	Hitachi Data Systems representative	Confirm that the status of the fibre-channel HBA(s) and LDEVs is NORMAL.
2.	Shut down the SGI system	User	Power off the SGI system before connecting the USP V. ▪ Shut down the SGI system. ▪ When shutdown is complete, power off the SGI IRIX display. ▪ Power off all peripheral devices except for the USP V system. ▪ Power off the host system. You are now ready to connect the USP V system.
3.	Connect the USP V system	Hitachi Data Systems representative	Install fibre-channel cables between the USP V and the SGI system. Follow all precautions and procedures in the USP V Maintenance Manual. Check all specifications to ensure proper installation and configuration.
4.	Power on the SGI system	User	Power on the SGI system after connecting the USP V system: ▪ Power on the SGI IRIX system display. ▪ Power on all peripheral devices. The USP V should be on, the fibre-channel ports should be configured, and the driver configuration file and system configuration file should be edited. If the fibre ports are configured or configuration files edited after the SGI system is powered on, restart the system to have the new devices recognized. ▪ Confirm the ready status of all peripheral devices, including the USP V. ▪ Power on the SGI system.

Verifying New Device Recognition

The final step before configuring the new USP V 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.

Hitachi Data Systems recommends that the devices should be installed and formatted with the fibre ports configured before the host system is powered on. Enter the **`cfgmgr`** command to force the system to check the buses for new devices.

To verify new device recognition:

1. Log in to the host system as **root**.
2. Display the peripheral device information using the **hinv** command (see Figure 2-3 and Figure 2-4)
3. Verify that the system recognizes all new disk devices, including OPEN-*x*, LUSE, VLL, VLL LUSE, and FX devices. The devices are listed by device file name.

Note: LUN 0 is implied when no LU number is listed.

```
# hinv                                                    ← Display device info.
4 250 MHZ IP27 Processors
CPU: MIPS R10000 Processor Chip Revision: 3.4
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
Main memory size: 2048 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 4 Mbytes
Integral SCSI controller 2: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 0: Version QL1040B (rev. 2), single ended
  Disk drive: unit 1 on SCSI controller 0
  CDROM: unit 6 on SCSI controller 0
Integral SCSI controller 1: Version QL1040B (rev. 2), single ended
Integral SCSI controller 3: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 4: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 5: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 6: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 8: Version Fibre Channel QL2200
  Disk drive: unit 0 on SCSI controller 8            ← TID=0, LUN = 0, SCSI controller 8.
  Disk drive: unit 0, lun 1 on SCSI controller 8     □ TID=0, LUN = 1, SCSI controller 8.
  Disk drive: unit 0, lun 2 on SCSI controller 8     □ TID=0, LUN = 2, SCSI controller 8.
  Disk drive: unit 0, lun 3 on SCSI controller 8
  Disk drive: unit 0, lun 4 on SCSI controller 8
  Disk drive: unit 0, lun 5 on SCSI controller 8
  Disk drive: unit 0, lun 6 on SCSI controller 8
  Disk drive: unit 0, lun 7 on SCSI controller 8
  Disk drive: unit 0, lun 8 on SCSI controller 8
  Disk drive: unit 0, lun 9 on SCSI controller 8
  Disk drive: unit 0, lun 10 on SCSI controller 8
  Disk drive: unit 0, lun 11 on SCSI controller 8
  Disk drive: unit 0, lun 12 on SCSI controller 8
  Disk drive: unit 0, lun 13 on SCSI controller 8
  Disk drive: unit 0, lun 14 on SCSI controller 8
  Disk drive: unit 0, lun 15 on SCSI controller 8
Integral SCSI controller 7: Version Fibre Channel AIC-1160, revision 2
IOC3 serial port: tty1
IOC3 serial port: tty2
Integral Fast Ethernet: ef0, version 1, module 1, slot io1, pci 2
Origin FIBRE CHANNEL board, module 1 slot 8: Revision 4
Origin BASEIO board, module 1 slot 1: Revision 4
Origin FIBRE CHANNEL board, module 1 slot 4: Revision 4
Origin FIBRE CHANNEL board, module 1 slot 7: Revision 4
Origin PCI XIO board, module 1 slot 2: Revision 4
IOC3 external interrupts: 1#
```

Figure 2-3 Verifying New Device Recognition

```
# hinv                                                    ← Display device info.
4 250 MHZ IP27 Processors
CPU: MIPS R10000 Processor Chip Revision: 3.4
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
Main memory size: 2048 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 4 Mbytes
Integral SCSI controller 2: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 0: Version QL1040B (rev. 2), single ended
  Disk drive: unit 1 on SCSI controller 0
  CDROM: unit 6 on SCSI controller 0
Integral SCSI controller 1: Version QL1040B (rev. 2), single ended
Integral SCSI controller 3: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 4: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 5: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 8: Version Fibre Channel QL2200
  Fabric Disk: node 50000e10ff809999 port 0 lun 0 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 1 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 2 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 3 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 4 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 5 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 6 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 7 on SCSI controller 8
  Fabric Disk: node 50000e10ff809999 port 0 lun 8 on SCSI controller 8
Integral SCSI controller 6: Version Fibre Channel AIC-1160, revision 2
Integral SCSI controller 7: Version Fibre Channel AIC-1160, revision 2
IOC3 serial port: tty1
IOC3 serial port: tty2
Integral Fast Ethernet: ef0, version 1, module 1, slot io1, pci 2
Origin FIBRE CHANNEL board, module 1 slot 7: Revision 4
Origin BASEIO board, module 1 slot 1: Revision 4
Origin FIBRE CHANNEL board, module 1 slot 4: Revision 4
Origin PCI XIO board, module 1 slot 2: Revision 4
Origin FIBRE CHANNEL board, module 1 slot 3: Revision 4
IOC3 external interrupts: 1
```

Figure 2-4 Verifying New Device Recognition (with Fabric)

Device Files and WWNs

The SGI IRIX system creates device files for new devices automatically during server startup. For LUN 0, the device name may not specify the LU number. The **rdsk** devices use a raw interface, while the **dsk** devices use a block interface. The **port#** and **nodename** indicate the worldwide name (WWN) and device port numbers (fibre-channel disks have two ports). The **vh** and **vol** devices are only in the **rdsk** directory and are normally used only for ioctl and raw access. The format for device file names in SGI IRIX is:

- FC-AL (and parallel SCSI)

  ```
  /dev/rdsk/dkscontroller#ddrive#{spartition#|vh|vol}
  /dev/rdsk/dkscontroller#ddrive#llun#{spartition#|vh|vol}
  /dev/dsk/dkscontroller#ddrive#spartition#
  /dev/dsk/dkscontroller#ddrive#llun#spartition#
  ```

  ***Example:* /dev/rdsk/dks8d0l8s0**

- Fibre-channel fabric

  ```
  /dev/rdsk/nodename/lunlun#{spartition#|vh|vol}/ccontroller#pport#
  /dev/dsk/nodename/lunlun#spartition#/ccontroller#pport#
  ```

  ***Example:* /dev/rdsk/50000e10ff809999/lun1vol/c8p50000e10ff809999**

Note: When the Brocade SilkWorm® 2800 fabric switch is used, the WWN information is displayed by the **nsShow** command (see Figure 2-5). The **PortName** (column 4) is the WWN and the **NodeName** is the device port number.

```
switch:admin> nsShow
    The Local Name Server has 7 entries {
Type Pid    COS     PortName                NodeName            TTL(sec)
*N   011200; 2,3;10:00:00:60:69:00:ab:ba;10:00:00:60:69:00:ab:ba; 60
     FC4s: FCIP
N    021200; 2,3;10:00:00:60:69:00:03:19;30:00:00:60:69:00:03:19; na
     FC4s: FCIP
N    021300; 3;10:00:00:60:69:00:02:d6;20:00:00:60:69:00:02:d6; na
NL   0214e2; 3;21:00:00:fa:ce:00:21:1e;20:00:00:fa:ce:00:21:1e; na
     FC4s: FCP [STOREX  RS2999FCPH3     MT09]
NL   0214e4; 3;21:00:00:fa:ce:00:21:e1;20:00:00:fa:ce:00:21:e1; na
     FC4s: FCP [STOREX  RS2999FCPH3     CD09]
NL   0214e8; 3;21:00:00:fa:ce:04:83:c9;20:00:00:fa:ce:04:83:c9; na
     FC4s: FCP [STOREX  RS2999FCPH3     NS09]
NL   0214ef; 3;21:00:00:ad:bc:04:6f:70;20:00:00:ad:bc:04:6f:70; na
     FC4s: FCP [STOREX  RS2999FCPH3     JB09]
     }
```

Figure 2-5 Displaying the WWN (NodeName) on Brocade Fabric Switch

Note: PortName = WWN; NodeName = device port number.

3

Configuring the New Disk Devices

This chapter describes how to configure the new USP V disk devices that you attached to the SGI IRIX system host in the previous chapter:

- ☐ Partitioning the Disk Devices
- ☐ Enabling Command Tag Queuing
- ☐ Creating the File Systems
- ☐ Creating the Mount Directories and Mounting the Devices
- ☐ Verifying the File Systems
- ☐ Setting the Auto-Mount Parameters

Note: The USP V logical devices are defined to the host as SCSI disk devices, even though the interface is fibre channel.

For information about configuring USP V for failover and SNMP, see Chapter 4.

For information about SCSI TID maps for fibre-channel adapters, see Appendix A.

For information about fibre port addressing (AL-PA to SCSI TID mapping) for SGI IRIX systems, see Appendix A.

For information about configuring newly installed USP V devices without rebooting the IRIX system, see Appendix B.

Partitioning the Disk Devices

After new device recognition has been verified, partition the new SCSI disk devices using the **fx** utility (see Figure 3-1 and Figure 3-2). After setting the partitions for a device, verify the partitions using the **prtvtoc** command (see Figure 3-3).

Available partitions: The IRIX system controls disk devices using partitions. One LU can be divided into a maximum of sixteen partitions (primary partition 0 through 15). The maximum capacity per partition is not limited. Partition #8 (vh) and partition #9 are reserved and are used for storing disk management information. Partition #10 is also reserved. Therefore, the number of available partitions per device is thirteen (0-7 and 11-15).

WARNING: Do not partition or label a disk device that will be accessed as a raw device (e.g. some database applications use raw devices).

Note: Do not change partitions 8, 9, or 10. Set the partitions (except partition 10) so that no partition is set on top of another partition. Because partition 8 usually uses parts between 0 and 6, set that base of first setting partition is 6.

```
# fx -x    ← Start the fx utility.
fx version 6.5, Jul 11, 1999
fx: "device-name" = (dksc)                                  ← Press Enter to label a disk.
fx: ctlr# = (0) 8                                           ← Enter the SCSI controller #.
fx: drive# = (1) 0                                                    ← Enter the TID.
fx: lun# = (0) 15                                                   ← Enter the LUN #.
...opening dksc(8,0,15)
...drive selftest...OK
Scsi drive type == Hitachi      OPEN-3          5244

----- please choose one (? For help, .. to quit this menu)-----
[exi]t             [d]ebug/           [l]abel/           [a]uto
[b]adblock/        [exe]rcise/        [r]epartition/
fx> r                                                    ← Enter r for partition menu.

----- partitions-----
part  type        blocks            Megabytes   (base+size)
  0: xfs      266240 + 2048000      130 + 1000
  1: xfs     2052096 + 2048000     1002 + 1000
  8: volhdr        0 + 4096            0 + 2
 10: volume        0 + 4806720         0 + 2347

capacity is 4806720 blocks
----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive           [o]ptiondrive         [e]xpert
[u]srrootdrive        [re]size
fx/repartition> e                                        ← Enter e for expert menu.

Warning: you will need to re-install all software and restore user data
from backups after changing the partition layout.  Changing partitions
will cause all data on the drive to be lost.  Be sure you have the drive
backed up if it contains any user data.  Continue? y              ← Enter y to continue.
Enter .. when done
```

Figure 3-1 Partitioning and Labeling the Disk Devices (without Fabric) (continues on the next page)

```
fx/repartition/expert: change partition = (0) 0                ← Enter partition number.
before:  type xfs        block  266240,       130 MB
                         len:  4540416 blks, 2217 MB
fx/repartition/expert: partition type = (xfs) xfs               ← Enter partition type.
fx/repartition/expert: base in megabytes = (130) 2
fx/repartition/expert: size in megabytes (max 2347) = (2217) 1024
after:  type xfs        block    4096,         2 MB
                         len:  2097152 blks, 1024 MB
fx/repartition/expert: change partition = (1) 1                ← Enter partition number.
before:  type xfs        block    4096,         2 MB
                         len:   262144 blks,  128 MB
fx/repartition/expert: partition type = (xfs) xfs               ← Enter partition type.
fx/repartition/expert: base in megabytes = (2) 1026
fx/repartition/expert: size in megabytes (max 1323) = (128) 1024
 after:  type xfs        block 2097152,      1024 MB
                         len:  2097152 blks, 1024 MB
fx/repartition/expert: change partition = (2) 2                ← Enter partition number.
before:  type xfs        block       0,         0 MB
                         len:        0 blks,    0 MB
fx/repartition/expert: partition type = (xfs) xfs               ← Enter partition type.
fx/repartition/expert: base in megabytes = (0) 2050
fx/repartition/expert: size in megabytes (max 299) = (0) 297
 after:  type xfs        block 4194304,      2048 MB
                         len:   608256 blks,  297 MB
:
:
----- partitions-----
part  type        blocks            Megabytes    (base+size)
  0: xfs           0 + 2097152        2 + 1024
  1: xfs     2097152 + 2097152     1026 + 1024
  3: xfs     4194304 + 608256      2050 + 297
  8: volhdr         0 + 4096                                    0 + 2
 10: volume          0 + 4806656              0 + 2347

capacity is 4806656 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive           [o]ptiondrive         [e]xpert
[u]srrootdrive        [re]size
fx/repartition> ..                                          ← Enter ".." to quit menu.
----- please choose one (? for help, .. to quit this menu)-----
[exi]t             [d]ebug/           [l]abel/           [a]uto
[b]adblock/        [exe]rcise/        [r]epartition/
fx>exi                                                      ← Exit the fx utility.
#
```

Figure 3-1 Partitioning and Labeling the Disk Devices (without fabric) (continued)

```
# fx -x -d /dev/rdsk/50000e10ff809999/lun2vol/c8p50000e10ff809999← Enter device file.
fx version 6.5, Jan 11, 2000
...opening /dev/rdsk/50000e10ff809999/lun2vol/c8p50000e10ff809999
...drive selftest...OK
Scsi drive type == HITACHI      OPEN-3          5245

----- please choose one (? for help, .. to quit this menu)-----
[exi]t             [d]ebug/           [l]abel/           [a]uto
[b]adblock/        [exe]rcise/        [r]epartition/
fx> r                                                  ← Enter r for partition menu.
```

Figure 3-2 Partitioning and Labeling a Device Connected via Fabric Switch

```
# prtvtoc /dev/dsk/dks8d0l15s0   ← Verifying partition 0 of LUN 15, TID 0, controller 8.
```

Figure 3-3 Verifying a Partition

Enabling Command Tag Queuing

Command tag queuing (CTQ) must be enabled to optimize the performance of the USP V devices. Since CTQ is disabled by default in IRIX systems, you need to enable it and set the queue depth for each USP V logical device using the **fx** utility. Table 3-1 lists the queue depth requirements for the USP V devices.

Table 3-1 Queue Depth Requirements for the USP V Devices

Parameter	Required Value
Queue depth per LU	≤ 32
Queue depth per port	≤ 1024

Note: You can adjust the queue depth for the USP V devices later as needed (within the specified range) to optimize the I/O performance of the USP V devices. If I/O response time will be long, you must adjust queue depth parameter. You can check the response time using the **sar** command (see Figure 3-4) and set queue depth within 10 seconds.

```
sgi 1# sar -d 1 10                                          ← Input sar command.
IRIX64 sgi 6.5-ALPHA-1276737220 09080737 IP27    10/17/00

11:01:02     device %busy  avque  r+w/s  blks/s    w/s wblks/s  avwait  avserv  ← Response
11:01:03     dks0d1     0    0.0    0.0       0    0.0       0     0.0     0.0  time (msec)
             dks0d6     0    0.0    0.0       0    0.0       0     0.0     0.0
           dks14d0    100    4.0    2.0     133    2.0     133  1745.0   290.0
         dks14d0l1    100    5.5    3.9     500    1.0     125  1117.5   285.0
         dks14d0l2    100    4.5    2.0     250    0.0       0   955.0   510.0
         dks14d0l3    100    5.2    3.9    2219    1.0       2  1735.0   255.0
         dks14d0l4    100    5.7    2.9     252    1.0       2   963.3   276.7
         dks14d0l5    100    5.0    2.0     250    0.0       0  1585.0   370.0
```

Figure 3-4 Checking the I/O Response Time Using the Sar Command

To enable CTQ and set the queue depth for the USP V devices (see Figure 3-5):

1. Start the **fx** disk utility and select the desired device to configure.
2. Once the device is selected and the **fx>** prompt reappears, enter **/label/set/para** to set the command tag queuing and queue depth options.
3. When prompted, enter **enable** to enable CTQ and enter the desired CTQ depth (e.g., 32).
4. When prompted, enter **yes** to modify the drive parameters as specified.
5. Exit the **fx** utility and enter **yes** to write out (save) the changes to the drive parameters.
6. Repeat steps 1 through 5 for each new USP V disk device.

```
# fx -x "dksc(8,0,15)"                                    ← Start fx and enter USP device.

fx version 6.5, Jan 11, 2000
...opening dksc(8,0,15)
...drive selftest...OK
fx: Warning:  invalid label from disk driver, ignored
Scsi drive type == HITACHI      OPEN-3-CVS      5244
...creating default bootinfo
...created default partitions, use /repartition menu to change
...creating default volume directory

----- please choose one (? for help, .. to quit this menu)-----
[exi]t             [d]ebug/           [l]abel/           [a]uto
[b]adblock/        [exe]rcise/        [r]epartition/
fx> /label/set/param                                   ← Set the device parameters.

fx/label/set/parameters: Error correction = (enabled)
fx/label/set/parameters: Data transfer on error = (enabled)
fx/label/set/parameters: Report recovered errors = (enabled)
fx/label/set/parameters: Delay for error recovery = (enabled)
fx/label/set/parameters: Err retry count = (0)
fx/label/set/parameters: Transfer of bad data blocks = (disabled)
fx/label/set/parameters: Auto bad block reallocation (write) = (enabled)
fx/label/set/parameters: Auto bad block reallocation (read) = (enabled)
fx/label/set/parameters: Read ahead caching = (enabled)
fx/label/set/parameters: Write buffering = (enabled)
fx/label/set/parameters: Drive disable prefetch = (0)
fx/label/set/parameters: Drive minimum prefetch = (0)
fx/label/set/parameters: Drive maximum prefetch = (0)
fx/label/set/parameters: Drive prefetch ceiling = (0)
fx/label/set/parameters: Enable CTQ = (disabled) enable              ← Enter "enable".
fx/label/set/parameters: CTQ depth = (2) 32                ← Enter desired queue depth.
fx/label/set/parameters: Read buffer ratio = (0/256)
fx/label/set/parameters: Write buffer ratio = (0/256)
 * * * * * W A R N I N G * * * * *
about to modify drive parameters on disk dksc(8,0,15)! ok? yes          ← Enter "yes".

----- please choose one (? for help, .. to quit this menu)-----
[exi]t             [d]ebug/           [l]abel/           [a]uto
[b]adblock/        [exe]rcise/        [r]epartition/
fx> exi                                                ← Exit the fx utility.

label info has changed for disk dksc(8,0,15). write out changes? (yes) yes← Enter "yes".
```

Figure 3-5 Enabling Command Tag Queuing (CTQ) and Setting the Queue Depth

Creating the File Systems

After partitioning and enabling CTQ for the new devices, you can create the file systems on the new SCSI disk devices. The standard file system is ESF™ and the extended file system is XFS®. The EFS™ file system creates one file system of 2 GB or less on a single device without the extended logical volume manager (XLV). The XFS® file system creates a 64-bit file system capable of scaling to handle extremely large files and file systems. The file system created is application-dependent. Make sure to select the correct file system for your operational setup.

Note: Do not create a file system on partition 8 or 10.

WARNING: Do not create a file system on a disk device that will be accessed as a raw device (e.g. some database applications use raw devices).

EFS™ File System

To create an EFS™ file system:

1. Use the **mkfs** command to create an EFS™ file system. For example, to create an EFS™ file system for controller 8, drive (TID) 0, logical unit 15, partition 0, enter:

   mkfs /dev/rdsk/dks8d0l15s0

Note: For fabric-connected devices, use the fabric device file name (see section 2.5.1) (e.g., /dev/rdsk/50000e10ff809999/lun2sl/c8p50000e10ff809999).

2. Repeat step 1 for each device partition on which you want to create an EFS™ file system.

XFS® File System

To create an XFS® file system:

1. Use the **mkfs** command to create an XFS® file system. For example, to create an XFS® file system for controller 8, drive (TID) 0, logical unit 15, partition 0, enter:

   mkfs -t xfs -d name=/dev/rdsk/dks8d0l15s0

Note: For fabric-connected devices, use the fabric device file name (see section 2.5.1) (e.g., /dev/rdsk/50000e10ff809999/lun2sl/c8p50000e10ff809999).

2. Repeat step 1 for each device partition on which you want to create an XFS® file system.

Creating the Mount Directories and Mounting the Devices

After you have created the file systems for the new USP V SCSI disk devices, you can create the mount directories and mount the new devices. Make sure to choose a unique directory name which identifies the logical volume.

To create the mount directories and mount the new SCSI disk devices:

1. Create the desired new mount directories using the **mkdir** command. For example, to create a mount directory for logical unit 0 on the USP V, enter:

   mkdir /USP_LU00

2. Mount all new USP V devices using the **mount** command. For example, to mount partition 0 of LUN 15, drive (TID) 0, controller 8, enter:

   mount /dev/dsk/dks8d0l15s0 /USP_LU00

Verifying the File Systems

Verify the file systems for the new USP V disk devices using the **df** command (see Figure 3-6). Make sure that the capacity value (**kB**) for each device is correct.

```
# df -k    ← List file systems.
File system             Type   Kbytes     use    avail  %use  Mounted on
/dev/root               xfs    969857  414702   555155   43%   /
/dev/dsk/dks8d0l15s0    xfs   1048576      13  1048563    0%   /USP_LU0     ← New device.
:
:
#
```

Figure 3-6 Verifying the New File Systems

Setting the Auto-Mount Parameters

The final step in configuring the USP V devices for SGI IRIX operations is to set the auto-mount parameters for the new devices. For each device to be auto-mounted, you will add the device to the system auto-mount table (**/etc/fstab** file). If you do not plan to auto-mount any of the USP V devices, you can skip this section.

To add new devices to the system auto-mount table:

1. First make a backup copy of the mount table: **cp /etc/fstab /etc/fstab.backup**
2. Add each desired new device to the mount table as shown in Figure 3-7. Table 3-2 describes the auto-mount parameters.

```
# cp /etc/fstab  /etc/fstab.backup                                  ← Make backup.
# vi /etc/fstab                                                   ← Edit mount table.
/dev/root            /             xfs  rw,raw=/dev/rroot  0  0
/dev/dsk/dks8d0l15s0 /DKC310_LU00  xfs  rw,noquota         0  1    ← Enter new device.
          ①              ②         ③      ④              ⑤ ⑥      ← See Table 3-2.
:
```

Figure 3-7 Setting the Auto-Mount Parameters

Table 3-2 Auto-Mount Parameters

No.	Description
①	Device to mount (device file name)
②	Mount point (mount directory)
③	File system (FS) type
④	Mount options (usually [rw,noquota])
⑤	Enhance – enter 0 for USP V devices
⑥	fsck pass – order in which FS checks are to be performed

4

Configuring the New Disk Devices

The USP V supports industry-standard products and functions that provide host and/or application failover, I/O path failover, and logical volume management (LVM). The USP V also supports 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 USP V 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 USP V storage system. The topics covered in this chapter are:

- ☐ Host/Application 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/Application Failover

The USP V supports the IRIS® FailSafe® and SGI Advanced Cluster Environment (ACE) software products for the SGI IRIX operating system. The user must make sure to configure the host failover software and any other high-availability (HA) software as needed to recognize and operate with the newly attached USP V devices.

For assistance with IRIS FailSafe and/or SGI ACE operations, refer to the user documentation or contact SGI 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 USP V uses SNMP to transfer status and management commands to the SNMP Manager on the UNIX/PC 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 USP V notifies the SNMP manager on the UNIX/PC server. Notification of USP V error conditions is made in real time, providing the UNIX/PC 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 USP V from the UNIX/PC server host.

When a SIM occurs, the USP V 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 SGI IRIX server host. For assistance with SNMP manager configuration on the SGI IRIX server host, refer to the user documentation, or contact the vendor's technical support.

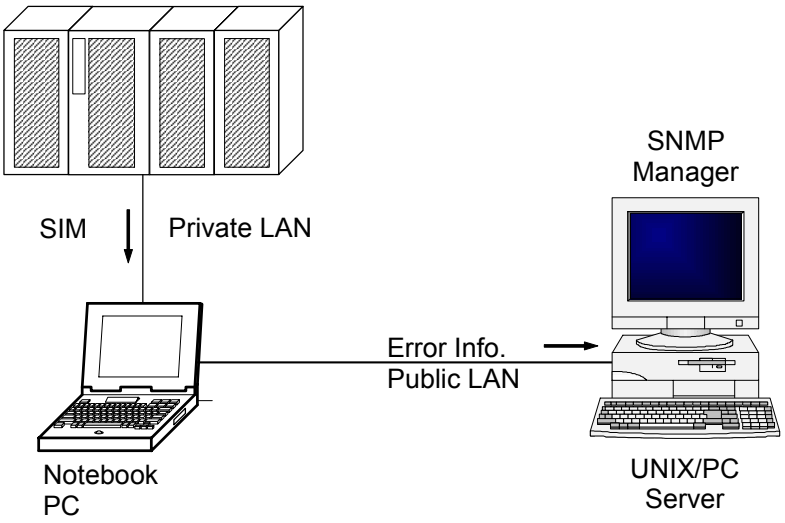

Figure 4-1 Example of a USP V SNMP Environment

5

Troubleshooting

This chapter provides information to help you identify and resolve problems in the unlikely event you encounter a problem with the USP V. The topics covered in this chapter are:

- ☐ Troubleshooting Error Conditions
- ☐ Calling the Hitachi Data Systems Support Center

Note: For additional troubleshooting information, refer to the *Hitachi Universal Storage Platform V User and Reference Guide* (MK-96RD635).

Troubleshooting Error Conditions

Table 5-1 lists potential error conditions that may occur during the USP V 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.

Table 5-1 Troubleshooting

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Make sure that the READY indicator lights on the USP V system are ON. Make sure that the fibre cables are correctly installed and firmly connected. Make sure that the LUNs are properly configured. The LUNs for each target ID must start at 0 and continue sequentially without skipping any numbers. Make sure that the SCSI IDs (unit IDs) on each bus are unique. Do not connect two devices with the same SCSI ID on the same bus.
The SGI system does not reboot properly after hard shutdown	If the SGI system is powered off without executing the shutdown process, wait three minutes before restarting the SGI system. This allows the USP V internal time-out process to purge all queued commands so that the USP V is available (not busy) during system startup. If the SGI system is restarted too soon, the USP V will continue trying to process the queued commands and the SGI system will not reboot successfully.

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. Include the circumstances surrounding the error or failure, the exact content of any messages displayed and the severity levels and reference codes of the R-SIMs on the R-SIM panel.

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

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

SCSI TID Maps for Fibre-Channel Adapters

When an arbitrated loop (AL) is established or re-established, the port addresses are assigned automatically to prevent duplicate TIDs. With the SCSI over fibre-channel protocol (FCP) there is no longer a need for target IDs in the traditional sense. 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 initiator to target. To enable transparent use of FCP, the SGI IRIX system "maps" a TID to each AL-PA.

The host maps SCSI protocol to fibre-channel protocol and detects and accesses fibre-connected devices using device files (/dev/dsk/c*t*d* and /dev/rdsk/c*t*d*) in the same way as for SCSI-connected devices. The device files for fibre-connected devices are configured in a different way than SCSI-connected devices because fibre supports 126 addresses per path while SCSI supports 16 TIDs per path.

Table A-1 identifies the fixed mappings between the TID (drive) values assigned by the IRIX system and the FC native addresses (AL_PA/SEL_ID) for FC adapters. The controller number (the **dks** value in /dev/dsk/dks*d*l*s*) depends on the server configuration and a different value is assigned per each column of Table A-1.

Note: The mapping defined in Table A-1 cannot be guaranteed under the following conditions:

When USP V devices and other types of devices are connected in the same loop.

When information for unused devices remains in server system.

When multiple ports participate in the same arbitrated loop.

Table A-1 Filter Port Addressing

AL-PA	T value	AL-PA	T value	AL-PA	T value	AL-PA	T value
EF	0	CD	16	B2	32	98	48
E8	1	CC	17	B1	33	97	49
E4	2	CB	18	AE	34	90	50
E2	3	CA	19	AD	35	8F	51
E1	4	C9	20	AC	36	88	52
E0	5	C7	21	AB	37	84	53
DC	6	C6	22	AA	38	82	54
DA	7	C5	23	A9	39	81	55
D9	8	C3	24	A7	40	80	56
D6	9	BC	25	A6	41	7C	57
D5	10	BA	26	A5	42	7A	58
D4	11	B9	27	A3	43	79	59
D3	12	B6	28	9F	44	76	60
D2	13	B5	29	9E	45	75	61
D1	14	B4	30	9D	46	74	62
CE	15	B3	31	9B	47	73	63
72	64	55	80	3A	96	23	112
71	65	54	81	39	97	23	113
6E	66	53	82	36	98	1F	114
6D	67	52	83	35	99	1E	115
6C	68	51	84	34	100	1D	116
6B	69	4E	85	33	101	1B	117
6A	70	4D	86	32	102	18	118
69	71	4C	87	31	103	17	119
67	72	4B	88	2E	104	10	120
66	73	4A	89	2D	105	0F	121
65	74	49	90	2C	106	08	122
63	75	47	91	2B	107	04	123
5C	76	46	92	2A	108	02	124
5A	77	45	93	29	109	01	125
59	78	43	94	27	110		
56	79	3C	95	26	111		

B

Online Device Installation

This appendix provides instructions for online installation of new devices. After initial installation and configuration of the USP V system, additional devices can be installed or de-installed online without having to restart the SGI system. These procedures are to be performed after logging in as a **super-user**.

Figure B-1 shows the two commands that must be executed to cause the SGI IRIX system to recognize newly added targets. Use the **scsiha –p #** command (# is the controller number) to scan the existing controller for new devices. Use the **ioconfig –f /hw** command to add the new device information to the hardware graph.

```
# scsiha   -p #                    ← Probe controller, where # is the controller number.
# ioconfig -f /hw                  ← Update the hardware graph with new device information.
```

Figure B-1 Recognizing New Devices Installed Online

Acronyms and Abbreviations

ACE	Advanced Cluster Environment
AL	arbitrated loop
AL-PA	arbitrated loop physical address
CTQ	command tag queuing
CVS	custom volume size
CXFS™	clustered version of XFS® file system
EFS™	IRIX standard file system
ESCON®	Enterprise System Connection (IBM trademark for optical channels)
ExSA™	Extended Serial Adapter™
FC	fibre-channel
FCA	fibre-channel adapter
FC-AL	fibre-channel arbitrated loop
FCP	fibre-channel protocol
FX	Hitachi Cross-OS File Exchange
fx	IRIX disk utility
GB	gigabyte
Gbps	gigabits per second
I/O, IO	input/output
kB	kilobytes
LDEV	logical device
LU	logical unit
LUN	logical unit number, logical unit
LUSE	LUN Expansion
LVI	Logical Volume Image
LVM	Logical Volume Manager, logical volume management
MB	megabytes
OFC	open fibre control
PA	physical address
PC	personal computer system
P-P	point-to-point

RAID	redundant array of independent disks
SCSI	small computer system interface
SGI	Silicon Graphics, Incorporated
SIM	service information message
SNMP	simple network management protocol
SVP	service processor
TID	target ID
USP	Universal Storage Platform
VLL	Virtual LVI/LUN
WWN	worldwide name
XFS®	IRIX extended file system
XLV	extended logical volume manager

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

Hitachi Data Systems

MK-96RD651-01