



**Hitachi Freedom Storage™
Lightning 9900™
HP-UX® Configuration Guide**

© 2000 Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

Notice: 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 Data Systems Corporation.

Hitachi Data Systems reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements. All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

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.

Trademarks

Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., and the Hitachi Data Systems design mark is a trademark and service mark of Hitachi, Ltd.

Extended Serial Adapter, ExSA, Hitachi Freedom Storage, Hitachi LUN Manager, and Lightning 9900 are trademarks of Hitachi Data Systems Corporation.

HP and HP-UX are registered trademarks of Hewlett-Packard Company (HP).

ESCON and S/390 are registered trademarks of International Business Machines Corporation (IBM).

UNIX is a registered trademark of X/Open Company Limited in the United States and other countries and is licensed exclusively through X/Open Company Limited.

All other brand or product names are or may be trademarks or service marks of and are used to identify products or services of their respective owners.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Please contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

| Revision | Date | Description |
|--------------|---------------|--|
| MK-90RD016-P | May 2000 | Preliminary Release |
| MK-90RD016-0 | July 2000 | Initial Release |
| MK-90RD016-1 | November 2000 | Revision 1, supersedes and replaces MK-90RD016-0 |

Source Document Revision Level

The following source documents were used to produce this 9900 configuration guide:

- *Hitachi DKC310 Disk Subsystem SCSI/Fibre Installation Manual, HP9000/800 Series, HP-UX 10.10, 10.20, 11.0*, revision 3.
- *Lightning 9900™ LUN Manager User's Guide*, MK-90RD006-1.
- ECN TW-90RD002 (Technical Writing ECN).

Preface

The *Hitachi Freedom Storage™ Lightning 9900™ HP-UX® Configuration Guide* describes and provides instructions for configuring the devices on the Hitachi Lightning 9900™ disk array subsystem for operation with the HP-UX® operating system. This configuration guide assumes that:

- the user has a background in data processing and understands direct-access storage device subsystems and their basic functions,
- the user is familiar with the Hitachi Lightning 9900™ array subsystem,
- the user is familiar with the HP-UX® operating system, the HP® 9000 system, and the fibre-channel adapters, and
- the user is familiar with the UNIX® file system, system commands, and utilities.

Note: The term “9900” refers to the entire Hitachi Lightning 9900™ subsystem family, unless otherwise noted. Please refer to the *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide* (MK-90RD008) for further information on the 9900 disk array subsystems.

Note: For further information on the HP-UX® operating system, please consult the HP-UX® user documentation, or contact HP® technical support.

Contents

Chapter 1 Overview of 9900 HP-UX® Configuration

| | | |
|-----|--|---|
| 1.1 | 9900 HP-UX® Configuration..... | 1 |
| 1.2 | Hitachi Lightning 9900™ Subsystem..... | 1 |
| 1.3 | Device Types and Configuration Procedures..... | 2 |

Chapter 2 Preparing for New Device Configuration

| | | |
|-------|--|----|
| 2.1 | Configuration Requirements..... | 7 |
| 2.2 | Installing the 9900 Subsystem..... | 8 |
| 2.3 | Preparing to Connect the 9900..... | 9 |
| 2.3.1 | Setting the Host Mode for the 9900 Ports..... | 9 |
| 2.3.2 | Configuring the 9900 Fibre-Channel Ports..... | 10 |
| 2.3.3 | Verifying the Host Fibre-Channel Adapter Installation..... | 12 |
| 2.3.4 | Verifying System Diagnostics..... | 12 |
| 2.4 | Connecting the 9900 Subsystem to the HP® System..... | 13 |

Chapter 3 Configuring the 9900 Devices

| | | |
|-------|---|----|
| 3.1 | Verifying New Device Recognition..... | 16 |
| 3.2 | Verifying the Device Files and Driver..... | 18 |
| 3.2.1 | Creating the Device Files..... | 20 |
| 3.3 | Partitioning the Disk Devices..... | 22 |
| 3.3.1 | Creating the Physical Volumes..... | 23 |
| 3.3.2 | Creating Volume Groups..... | 24 |
| 3.3.3 | Creating the Logical Volumes..... | 28 |
| 3.4 | Creating the File Systems..... | 30 |
| 3.5 | Setting the Device Parameters..... | 31 |
| 3.5.1 | Setting the IO Time-Out Parameter..... | 31 |
| 3.5.2 | Setting the Queue Depth Parameter..... | 32 |
| 3.5.3 | Creating and Registering the Queue Depth Start-Up Script..... | 34 |
| 3.6 | Creating the Mount Directories..... | 37 |
| 3.7 | Mounting and Verifying the File Systems..... | 38 |
| 3.8 | Setting and Verifying the Auto-Mount Parameters..... | 39 |

Chapter 4 Failover and SNMP Configuration

| | | |
|-----|---------------------------------------|----|
| 4.1 | Application Failover..... | 41 |
| 4.2 | I/O Path Failover..... | 41 |
| 4.3 | SNMP Remote Subsystem Management..... | 42 |

Chapter 5 Troubleshooting

| | | |
|-----|--|----|
| 5.1 | Troubleshooting..... | 43 |
| 5.2 | Calling the Hitachi Data Systems Support Center..... | 44 |

| | | |
|------------|---------------------------------|----|
| Appendix A | Acronyms and Abbreviations..... | 45 |
|------------|---------------------------------|----|

| | |
|---|----|
| Appendix B SCSI TID Map for Fibre-Channel Adapters | 47 |
| Appendix C Configuring the Boot Disk on the 9900 | 49 |
| Appendix D Reference Information for SAM | |
| D.1 Configuring the 9900 Devices Using SAM | 54 |
| D.2 Setting the Maximum Number of Volume Groups Using SAM | 54 |
| Appendix E Online Device Installation | 57 |

List of Figures

| | | |
|-------------|--|----|
| Figure 2.1 | Setting the Host Mode | 9 |
| Figure 2.2 | Fibre Parameter Panel (from the LUN Manager Software) | 11 |
| Figure 2.3 | Verifying Host Fibre-Channel Adapter Installation | 12 |
| Figure 2.4 | Shutting Down and Powering Off the HP® System | 13 |
| Figure 3.1 | Log In as Root | 17 |
| Figure 3.2 | Verify New Device Recognition | 17 |
| Figure 3.3 | Verifying the Block-Type Device Files | 18 |
| Figure 3.4 | Verifying the Character-Type Device Files | 18 |
| Figure 3.5 | Verifying the HP-UX® Driver | 19 |
| Figure 3.6 | Creating Device Files Automatically | 20 |
| Figure 3.7 | Displaying the Driver Information | 21 |
| Figure 3.8 | Creating the Block-Type and Character-Type Device Files Manually | 21 |
| Figure 3.9 | Creating Physical Volumes | 23 |
| Figure 3.10 | Displaying the Existing Volume Group Names | 25 |
| Figure 3.11 | Creating and Verifying a Directory for the New Volume Group | 25 |
| Figure 3.12 | Displaying the Minor Numbers for the Existing Group Files | 25 |
| Figure 3.13 | Creating the Group File for the New Volume Group | 25 |
| Figure 3.14 | Creating the New Volume Group | 25 |
| Figure 3.15 | Verifying the New Volume Group | 26 |
| Figure 3.16 | Creating a Logical Volume | 28 |
| Figure 3.17 | Calculating the Volume Size for CVS, LUSE, and CVS-LUSE Devices | 29 |
| Figure 3.18 | Verifying a Logical Volume | 29 |
| Figure 3.19 | Creating a File System (default file system, OPEN-3 shown) | 30 |
| Figure 3.20 | Creating a File System (default file system, OPEN-9 shown) | 30 |
| Figure 3.21 | Specifying the File System Type | 30 |
| Figure 3.22 | Checking the Current IO Time-Out Value | 31 |
| Figure 3.23 | Changing the IO Time-Out Value | 31 |
| Figure 3.24 | Verifying the New IO Time-Out Value | 31 |
| Figure 3.25 | Changing the Queue Depth for the 9900 Disk Devices | 33 |
| Figure 3.26 | Checking the Existing Script Names and Link Names | 33 |
| Figure 3.27 | Verifying the Queue Depth for the 9900 Disk Devices | 33 |
| Figure 3.28 | Creating and Verifying a Mount Directory | 37 |

| | | |
|-------------|--|----|
| Figure 3.29 | Mounting a File System | 38 |
| Figure 3.30 | Verifying the File Systems | 38 |
| Figure 3.31 | Final Verification of the File System for One Volume | 38 |
| Figure 3.32 | Setting the Auto-Mount Parameters | 39 |
| Figure 4.1 | 9900 SNMP Environment | 42 |
| Figure C.1 | Checking the Status and Path of LU X and LU 0 | 50 |
| Figure C.2 | Confirming the Firmware (BIOS) Version | 50 |
| Figure C.3 | Setting the New Boot Disk | 51 |
| Figure C.4 | Confirming Boot from 9900 LU X | 51 |

List of Tables

| | | |
|-----------|--|-------|
| Table 1.1 | 9900 Device Specifications | 3-4 |
| Table 1.2 | Volume Usage for Device Categories | 5 |
| Table 2.1 | Fibre Parameter Settings on the 9900 Remote Console PC | 10 |
| Table 2.2 | Available AL-PA Values | 11 |
| Table 3.1 | Device Data Table (Sample) | 17 |
| Table 3.2 | Completed Device Data Table (Sample) | 21 |
| Table 3.3 | PE and MPE Parameters for LUSE Devices | 26-27 |
| Table 3.4 | Queue Depth Requirements for the 9900 Devices | 32 |
| Table 3.5 | Auto-Mount Parameters | 39 |
| Table 5.1 | Troubleshooting | 43 |
| Table B.1 | AL-PA to SCSI TID Mapping (t value) for HP-UX® Systems | 48 |

Chapter 1 Overview of 9900 HP-UX® Configuration

1.1 9900 HP-UX® Configuration

This document describes the requirements and procedures for connecting the 9900 subsystem to an HP-UX® system and configuring the new 9900 devices for operation with the HP-UX® operating system. The Hitachi Data Systems representative performs the physical installation of the 9900 subsystem. The user prepares for 9900 subsystem installation, and then configures the new 9900 devices with assistance as needed from the Hitachi Data Systems representative.

Configuration of the 9900 disk devices for HP-UX® operations includes:

- Verifying new device recognition (see section 3.1),
- Verifying the device files and driver (see section 3.2),
- Partitioning the disk devices (see section 3.3),
- Creating the file systems for the new volumes (see section 3.4),
- Setting the device parameters for the disk devices (see section 3.5),
- Creating the mount directories (see section 3.6),
- Mounting and verifying the file systems (see section 3.7), and
- Setting and verifying the auto-mount parameters (see section 3.8).

Note on the term "SCSI disk": The 9900 logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

1.2 Hitachi Lightning 9900™ Subsystem

The Hitachi Lightning 9900™ RAID subsystem supports concurrent attachment to multiple UNIX®-based and PC-server platforms. Please contact your Hitachi Data Systems account team for the latest information on platform support. The 9900 subsystem provides continuous data availability, high-speed response, scaleable connectivity, and expandable capacity for PC server and open-system storage. The 9900 subsystem can operate with multihost applications and host clusters, and is designed to handle very large databases as well as data warehousing and data mining applications that store and retrieve terabytes of data.

The Hitachi Lightning 9900™ subsystem can be configured with fibre-channel ports and/or Extended Serial Adapter™ (ExSA™) ports (compatible with ESCON® protocol) to provide connectivity with S/390® mainframe hosts as well as open-system hosts. For further information on the 9900 subsystem, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide* (MK-90RD008), or contact your Hitachi Data Systems account team.

1.3 Device Types and Configuration Procedures

The 9900 subsystem allows the following types of logical devices (also called volumes) to be installed and configured for operation with the HP-UX[®] operating system. Table 1.1 lists the device specifications for the 9900 devices. Table 1.2 shows the volume usage (i.e., file system or raw device) for the 9900 devices.

OPEN-x Devices. The OPEN-x logical units (LUs) (e.g., OPEN-3, OPEN-9) are disk devices of predefined sizes. The 9900 subsystem currently supports OPEN-3, OPEN-8, OPEN-9, OPEN-K, and OPEN-E devices. Please contact your Hitachi Data Systems account team for the latest information on supported LU types.

LUSE Devices (OPEN-x*n). The LUSE devices are combined LUs which can be from 2 to 36 times larger than standard OPEN-x LUs. The Logical Unit Size Expansion (LUSE) feature of the 9900 subsystem enables you to 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 $2 \leq n \leq 36$. For example, a LUSE device created from ten OPEN-3 LUs would be designated as an OPEN-3*10 disk device. This capability enables the server host to combine logical devices and access the data stored on the 9900 subsystem using fewer LU numbers (LUNs). For further information on the LUSE feature, please refer to the *Hitachi Freedom Storage[™] Lightning 9900[™] LUN Manager User's Guide* (MK-90RD006).

CVS Devices (OPEN-x CVS). The CVS devices are custom-size LUs which are smaller than standard OPEN-x LUs. The Custom Volume Size (CVS) feature of the 9900 subsystem (also called Virtual LUN and Virtual LVI) enables you to configure CVS devices. The CVS capability enables you to “slice up” a single LU into several smaller LUs to best fit the application needs and improve host access to frequently used files. For further information on the CVS feature, please refer to the *Hitachi Freedom Storage[™] Lightning 9900[™] Virtual LVI/LUN User's Guide* (MK-90RD005).

CVS LUSE Devices (OPEN-x*n CVS). The CVS LUSE devices combine CVS devices (instead of standard OPEN-x LUs) into LUSE devices. The CVS feature is used first to create custom-size devices, and then the LUSE feature is used to combine (concatenate) these CVS devices. The user can combine from 2 to 36 CVS devices into one CVS LUSE device. For example, an OPEN-3 LUSE volume that was created from ten OPEN-3 CVS volumes would be designated as an OPEN-3*10 CVS device.

Configuration of the 9900 disk devices for HP-UX[®] operations includes:

- Verifying new device recognition (see section 3.1),
- Verifying the device files and driver (see section 3.2),
- Partitioning the disk devices (see section 3.3),
- Creating the file systems for the new volumes (see section 3.4),
- Setting the device parameters for the disk devices (see section 3.5),
- Creating the mount directories (see section 3.6),
- Mounting and verifying the file systems (see section 3.7), and
- Setting and verifying the auto-mount parameters (see section 3.8).

HMDE Devices (3390-3A/B/C, 3380-KA/B/C, OPEN-x-HMDEoto). The Hitachi Multiplatform Data Exchange (HMDE) feature of the 9900 subsystem enables user data to be shared across S/390®, UNIX®, and PC server platforms using special multiplatform volumes. The CVS feature can also be applied to HMDE devices for maximum flexibility in volume size. For further information on HMDE, please refer to the *Hitachi Multiplatform Data Exchange User's Guide* (MK-90RD020), or contact your Hitachi Data Systems account team.

The HMDE devices must be installed and accessed as raw devices. UNIX®/PC server hosts must use HMDE to access the HMDE devices as raw devices (i.e., no disk partition, no file system, no mount operation).

Note: The 3390-3B and 3380-KB devices are write-protected from UNIX®/PC server access. The 9900 subsystem will reject all UNIX®/PC server write operations (including fibre-channel adapters) for the 3390-3B and 3380-KB devices.

WARNING: The 3390-3A/C, 3380-KA/C, and OPEN-x-HMDEoto devices are *not* write-protected for UNIX®/PC server access. Do not execute any write operations on these devices. Do not create a partition or file system on these devices. This will overwrite any data on the HMDE device and also prevent the HMDE software from accessing the device.

Configuration of the 9900 HMDE (raw) devices for HP-UX® operations includes:

- Verifying new device recognition (see section 3.1), and
- Verifying the device files and driver (see section 3.2).

Table 1.1 9900 Device Specifications (continues on the next page)

| Device Type (Note 1) | Category (Note 2) | Vendor Name | Product Name (Note 3) | # of Blocks (512-byte blk) | Sector Size (bytes) | # of Data Cylinders | # of Heads | # of Sectors per Track | Capacity MB (Note 4) |
|-------------------------|----------------------|----------------|--------------------------|-------------------------------|------------------------|------------------------|---------------|---------------------------|-------------------------|
| OPEN-3 | SCSI disk | HITACHI | OPEN-3 | 4806720 | 512 | 3338 | 15 | 96 | 2344 |
| OPEN-9 | SCSI disk | HITACHI | OPEN-9 | 14423040 | 512 | 10016 | 15 | 96 | 7040 |
| OPEN-K | SCSI disk | HITACHI | OPEN-K | 3661920 | 512 | 2543 | 15 | 96 | 1740 |
| OPEN-8 | SCSI disk | HITACHI | OPEN-8 | 14351040 | 512 | 9966 | 15 | 96 | 7004 |
| OPEN-E | SCSI disk | HITACHI | OPEN-E | 28452960 | 512 | 19759 | 15 | 96 | 13888 |
| OPEN-3*n | SCSI disk | HITACHI | OPEN-3*n | 4806720*n | 512 | 3338*n | 15 | 96 | S1 |
| OPEN-9*n | SCSI disk | HITACHI | OPEN-9*n | 14423040*n | 512 | 10016*n | 15 | 96 | S1 |
| OPEN-K*n | SCSI disk | HITACHI | OPEN-K*n | 3661920*n | 512 | 2543*n | 15 | 96 | S1 |
| OPEN-8*n | SCSI disk | HITACHI | OPEN-8*n | 14351040*n | 512 | 9966*n | 15 | 96 | S1 |
| OPEN-E*n | SCSI disk | HITACHI | OPEN-E*n | 28452960*n | 512 | 19759*n | 15 | 96 | S1 |
| OPEN-3 CVS | SCSI disk | HITACHI | OPEN-3-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-9 CVS | SCSI disk | HITACHI | OPEN-9-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-K CVS | SCSI disk | HITACHI | OPEN-K-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-8 CVS | SCSI disk | HITACHI | OPEN-8-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-E CVS | SCSI disk | HITACHI | OPEN-E-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |

Table 1.1 9900 Device Specifications (continued)

| Device Type (Note 1) | Category (Note 2) | Vendor Name | Product Name (Note 3) | # of Blocks (512-byte blk) | Sector Size (bytes) | # of Data Cylinders | # of Heads | # of Sectors per Track | Capacity MB (Note 4) |
|-------------------------|----------------------|----------------|--------------------------|-------------------------------|------------------------|------------------------|---------------|---------------------------|-------------------------|
| OPEN-3*n CVS | SCSI disk | HITACHI | OPEN-3*n-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-9*n CVS | SCSI disk | HITACHI | OPEN-9*n-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-K*n CVS | SCSI disk | HITACHI | OPEN-K*n-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-8*n CVS | SCSI disk | HITACHI | OPEN-8*n-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| OPEN-E*n CVS | SCSI disk | HITACHI | OPEN-E*n-CVS | Note 5 | 512 | Note 6 | 15 | 96 | S1 |
| 3390-3A | HMDE otm/mt0 | HITACHI | 3390-3A | 5820300 | 512 | 3345 | 15 | 116 | -- |
| 3380-KA | HMDE otm/mt0 | HITACHI | 3380-KA | 3833280 | 512 | 2662 | 15 | 96 | -- |
| 3390-3B | HMDEmt0 | HITACHI | 3390-3B | 5816820 | 512 | 3343 | 15 | 116 | -- |
| 3380-KB | HMDEmt0 | HITACHI | 3380-KB | 3830400 | 512 | 2660 | 15 | 96 | -- |
| 3390-3C | HMDEotm | HITACHI | OP-C-3390-3C | 5820300 | 512 | 3345 | 15 | 116 | -- |
| 3380-KC | HMDEotm | HITACHI | OP-C-3380-KC | 3833280 | 512 | 2662 | 15 | 96 | -- |
| HMDE OPEN-3 | HMDEoto | HITACHI | OPEN-3 | 4806720 | 512 | 3338 | 15 | 96 | -- |
| 3390-3A CVS | HMDE otm/mt0 | HITACHI | 3390-3A-CVS | Note 5 | 512 | Note 6 | 15 | 116 | -- |
| 3380-KA CVS | HMDE otm/mt0 | HITACHI | 3380-KA-CVS | Note 5 | 512 | Note 6 | 15 | 96 | -- |
| 3390-3B CVS | HMDEmt0 | HITACHI | 3390-3B-CVS | Note 5 | 512 | Note 6 | 15 | 116 | -- |
| 3380-KB CVS | HMDEmt0 | HITACHI | 3380-KB-CVS | Note 5 | 512 | Note 6 | 15 | 96 | -- |
| 3390-3C CVS | HMDEotm | HITACHI | OP-C-3390-3C- CVS | Note 5 | 512 | Note 6 | 15 | 116 | -- |
| 3380-KC CVS | HMDEotm | HITACHI | OP-C-3380-KC- CVS | Note 5 | 512 | Note 6 | 15 | 96 | -- |
| HMDE OPEN-3 CVS | HMDEoto | HITACHI | OPEN-3-CVS | Note 5 | 512 | Note 6 | 15 | 96 | -- |

Note 1: The availability of a specific 9900 device type depends on the level of microcode installed on the 9900 subsystem.

Note 2: The category of a device (SCSI disk or HMDE) determines its volume usage. Table 1.2 shows the volume usage for SCSI disk devices and HMDE devices. The SCSI disk devices (OPEN-x, CVS, LUSE, CVS LUSE) are usually formatted with file systems for HP-UX® operations. The HMDE devices (3390-3A/B/C, 3380-KA/B/C, OPEN-x-HMDEoto) must be installed as raw devices and can only be accessed using the HMDE software. Do not partition or create a file system on any device used for HMDE operations.

Note on the term "SCSI disk": The 9900 logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

Table 1.2 Volume Usage for Device Categories

| Category | Device Type | Volume Usage |
|-----------|--|--------------|
| SCSI Disk | OPEN-x, OPEN-x CVS, OPEN-x*n LUSE, OPEN-x*n CVS LUSE | File System* |
| HMDE | 3390-3A/B/C, 3380-KA/B/C 3390-3A/B/C CVS, 3380-KA/B/C CVS OPEN-x for HMDEoto, OPEN-x CVS for HMDEoto | Raw Device |

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

Note 3: The 9900 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).

Note 4: This capacity is the maximum size which can be entered using the **lvcreate** command.

Note 5: The number of blocks for a CVS volume is calculated as follows:

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

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

of blocks = (53 cylinders—see Note 5) × (15 heads) × (96 sectors per track) = 76320

Note 6: The number of data cylinders for a CVS 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 CVS volume =
of cylinders = ↑ (capacity (MB) specified on Remote Console PC) × 1024/720 ↑
 Example: For an OPEN-3 CVS volume with capacity = 37 MB:
 # of cylinders = ↑37 × 1024/720↑ = ↑52.62↑ (rounded up to next integer) = 53 cylinders
- The number of data cylinders for a CVS LUSE volume =
of cylinders = ↑ (capacity (MB) specified on Remote Console PC) × 1024/720 ↑ × n
 Example: For an OPEN-3 CVS 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 a 3390-3A/C or 3380-KA/C CVS volume =
of cylinders = (number of cylinders specified on Remote Console PC) + 9
- The number of data cylinders for a 3390-3B or 3380-KB CVS volume =
of cylinders = (number of cylinders specified on Remote Console PC) + 7

S1 = maximum **lvcreate** size value for CVS, LUSE, and CVS LUSE devices. Calculate the maximum size value (in MB) as follows: **S1 = (PE Size) × (Free PE)**

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

The requirements for 9900 HP-UX[®] configuration are:

- Hitachi Lightning 9900[™] subsystem, all-open or multiplatform configuration:
 - The LUN Manager software on the 9900 Remote Console PC is used to configure the fibre-channel (FC) ports. If remote LUN Manager is not installed, please contact your Hitachi Data Systems account team for information on LUN configuration services.

Note: The availability of 9900 features and devices depends on the level of microcode installed on the 9900 subsystem.

- HP[®] 9000/800 series system. **Note:** Connectability of the 9900 subsystem to model E35 of the HP[®] 9000/800 Series has been verified. For assistance with other HP[®] models, please contact the Hitachi Data Systems Support Center.
- HP-UX[®] operating system, version 10.20 or 11.0. **Important:** Please contact HP[®] technical support to make sure the most current OS patches are installed on the HP[®] system(s).
 - **root** log-in (superuser) access to the HP[®] system is required.

Note: Hitachi Data Systems plans to support future releases of the HP-UX[®] operating system. For the latest information on HP-UX[®] version support, please contact your Hitachi Data Systems account team.

- Fibre-channel adapters. Make sure to install all utilities, tools, and drivers that come with the adapter(s). For information on driver requirements for the adapters, please refer to the user documentation for the adapter or contact the vendor.
 - The 9900 subsystem supports full-speed (100 MB/s) fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface, and multimode optical cables with SC connectors. Do not connect any OFC-type fibre-channel interface to the 9900 subsystem. For information on supported FC adapters (FCAs), optical cables, hubs, and fabric switches, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center (see section 5.2).

Note: If you plan to connect different types of servers to the 9900 subsystem via the same fabric switch, you must use the **zoning** function of the fabric switch.

2.2 Installing the 9900 Subsystem

The 9900 subsystem comes with all hardware and cabling required for installation. Installation of the 9900 subsystem involves the following activities:

1. **Hardware installation.** The Hitachi Data Systems representative performs hardware installation as specified in the 9900 maintenance manual. Follow all precautions and procedures in the 9900 maintenance manual. Check all specifications to ensure proper installation and configuration. Hardware installation includes:
 - Assembling all hardware and cabling.
 - Loading the latest microcode and service processor (SVP) updates for full fibre-channel support.
 - Installing and formatting the logical devices (LDEVs) using the SVP. Make sure to get the desired LDEV configuration information from the user, including the desired number of OPEN-x, LUSE, CVS, CVS LUSE, and multiplatform (HMDE) devices.

Note: Although the 9900 subsystem supports up to 256 LUs per port, connectivity with HP[®] systems is limited to eight (8) LUs per port.

- Installing the fibre-channel adapters 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 connector to the 9900 subsystem. Do not connect/disconnect fibre-channel cabling that is being actively used for I/O. This can cause the HP-UX[®] system to hang. Always confirm that the devices on the fibre cable are offline before connecting/disconnecting the fibre cable.

9900 FC Port: The fibre topology parameters for each 9900 fibre-channel port depend on the type of device to which the 9900 port is connected. Determine the topology parameters supported by the device, and set your topology accordingly (see section 2.3.2). The type of 9900 port is also important.

Note: The Hitachi Data Systems representative must use the 9900 Maintenance Manual during all installation activities. Follow all precautions and procedures in the maintenance manual, and always check all specifications to ensure proper installation and configuration.

2. **Remote Console PC and LUN Manager installation.** The user or Hitachi Data Systems representative can perform this activity. You will use the LUN Manager software on the 9900 Remote Console PC to configure the 9900 fibre-channel ports. For instructions on installing the Remote Console PC and LUN Manager remote console software, please refer to the *9900 Remote Console User's Guide* (MK-90RD003) and the *9900 LUN Manager User's Guide* (MK-90RD006) respectively.

Note: If the remote LUN Manager feature is not installed, the Hitachi Data Systems representative can configure FC ports for you using the SVP of the subsystem. Please contact your Hitachi Data Systems account team for further information on fibre-channel configuration services.

2.3 Preparing to Connect the 9900

Before the 9900 is connected to the HP® system, you must perform the following tasks:

- Set the host mode for the 9900 fibre-channel port(s) (see section 2.3.1),
- Configure the 9900 fibre-channel ports (see section 2.3.2),
- Verify the host fibre-channel adapter installation (see section 2.3.3), and
- Verify the system diagnostics setting (see section 2.3.4).

You will use the LUN Manager remote console software to set the host modes for and configure the 9900 fibre ports. For instructions on using the LUN Manager software, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006). After completing these steps, you will shut down the HP® system, connect the 9900 subsystem, and then restart the HP® system (see section 2.4).

Note: If the remote LUN Manager feature is not installed, please contact your Hitachi Data Systems account team for information on fibre-channel configuration services.

2.3.1 Setting the Host Mode for the 9900 Ports

The 9900 ports have special modes which must be set for the connected operating system. Use the LUN Manager remote console software to set the host mode for each port (see Figure 2.1). The required host mode setting for 9900 HP-UX® operations is **08**.

Note: If you plan to connect different types of servers to the 9900 subsystem via the same fabric switch, you must use the **zoning** function of the fabric switch.

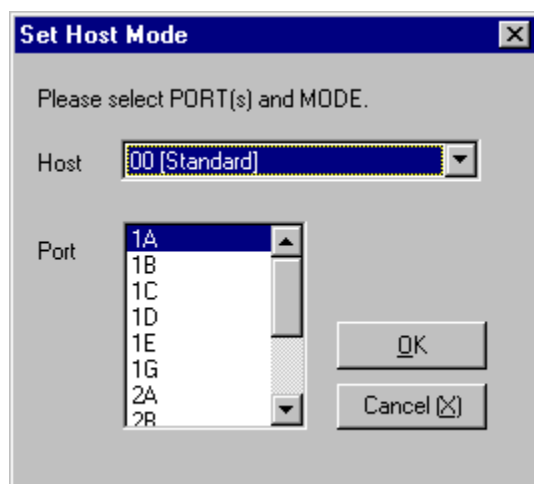


Figure 2.1 Setting the Host Mode

2.3.2 Configuring the 9900 Fibre-Channel Ports

You also need to configure the 9900 FC ports to define the fibre parameters (see Figure 2.2 and Table 2.1) and port addresses (see Table 2.2). You will use the LUN Manager remote console software to configure the 9900 FC ports. For instructions on using LUN Manager, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

Note: The 9900 subsystem supports up to 256 LUs per fibre-channel port, but connectivity with HP® systems is limited to eight (8) LUs per port.

Fibre topology. Figure 2.2 shows the Fibre Parameter panel (part of the LUN Manager software), and Table 2.1 explains the settings on this panel. You will select the appropriate settings for each 9900 FC 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 type of 9900 port is also important. **Note:** If you plan to connect different types of servers to the 9900 via the same fabric switch, you must use the **zoning** function of the fabric switch.

Port address. In fabric environments, the port addresses are assigned automatically by fabric switch port number and are not controlled by the 9900 port settings. In arbitrated loop environments, the port addresses are set by entering an AL-PA (arbitrated-loop physical address, or loop ID). Table 2.2 shows the available 9900 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 9900 port to a SCSI target ID (TID). See Appendix B for a description of the AL-PA-to-TID translation.

Note on loop ID conflict: The HP® 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, refer to Appendix B for the TID-to-AL-PA mapping).

Table 2.1 Fibre Parameter Settings on the 9900 Remote Console PC

| Fabric Parameter | Connection Parameter | Provides: |
|------------------|----------------------|-----------------------------------|
| ON | FC-AL | Not supported |
| ON | Point-to-Point | Not supported |
| OFF | FC-AL | AL-port (private arbitrated loop) |
| OFF | Point-to-Point | Not supported |

Fibre Parameter

Port Name : 1A

Port Type : Fibre Channel

Please select new Parameter

Fibre Address

☒ AL-PA EF -> EF

☐ Loop ID 0 -> 0

Fabric

OFF -> OFF

Connection

FC_AL -> FC_AL

OK Cancel

Figure 2.2 Fibre Parameter Panel (from the LUN Manager Software)

Table 2.2 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 | |

2.3.3 Verifying the Host Fibre-Channel Adapter Installation

After configuring the fibre-channel ports on the 9900 subsystem, you need to verify that the host fibre-channel adapters are installed properly using the **ioscan -f** command (see Figure 2.3). Please verify that the underlined items are displayed as a result of **ioscan -f** command. If these items are not displayed, the host environment may not be installed properly. In this case, check the host adapter installation (hardware and driver installation) or the host configuration.

| # ioscan -f | | | | | | |
|-------------|---|------------------|-----------|-----------|-----------|--|
| Class | I | H/W Path | Driver | S/W State | H/W Type | Description |
| ===== | | | | | | |
| bc | 0 | | root | CLAIMED | BUS_NEXUS | |
| bc | 1 | 8 | bc | CLAIMED | BUS_NEXUS | Pseudo Bus Converter |
| bc | 2 | 8/0 | bc | CLAIMED | BUS_NEXUS | Bus Converter |
| tty | 0 | 8/0.0 | mux2 | CLAIMED | INTERFACE | MUX |
| ext_bus | 0 | 8/4 | c720 | CLAIMED | INTERFACE | GSC add on Fast/Wide SCSI INTERFACE |
| target | 0 | 8/4.5 | tgt | CLAIMED | DEVICE | |
| disk | 0 | 8/4.5.0 | sdisk | CLAIMED | DEVICE | SEagate ST321S1W |
| target | 1 | 8/4.7 | tgt | CLAIMED | DEVICE | |
| ctl | 0 | 8/4.7.0 | sct1 | CLAIMED | DEVICE | Initiator |
| fc | 0 | 8/12 | fcT1 | CLAIMED | INTERFACE | <u>HP Fibre Channel Mass Storage Adapter</u> |
| lan | 1 | 8/12.5 | fcT1 cnt1 | CLAIMED | INTERFACE | <u>HP Fibre Channel Mass Storage Cntl</u> |
| fcp | 0 | 8/12.8 | fcp | CLAIMED | INTERFACE | <u>FCP Protocol Adapter</u> |
| ext_bus | 2 | 8/12.8.0.255.0 | fcpdev | CLAIMED | INTERFACE | <u>FCP Device Interface</u> |
| target | 7 | 8/12.8.0.255.0.6 | tgt | CLAIMED | DEVICE | |
| target | 8 | 8/12.8.0.255.0.8 | tgt | CLAIMED | DEVICE | |
| target | 9 | 8/12.8.0.255.0.9 | tgt | CLAIMED | DEVICE | |
| # | | | | | | |

Figure 2.3 Verifying Host Fibre-Channel Adapter Installation

2.3.4 Verifying System Diagnostics

Before connecting the 9900 subsystem to an HP® 9000 N4000 enterprise server system, verify the following parameter in the `/etc/rc.config.d/diagnostic` file:

- DIAGNOSTICS = 0

If DIAGNOSTICS = 1, then change the value to 0, and reboot the HP® server.

Note: Setting the DIAGNOSTICS parameter to zero (0) disables error logging for FC. If FC problems occur, you might need to re-enable the DIAGNOSTICS parameter.

2.4 Connecting the 9900 Subsystem to the HP® System

After you have configured the 9900 fibre ports and verified the host fibre-channel adapter installation, you are ready to connect the 9900 subsystem to the HP-UX® system. The 9900 subsystem comes with all hardware and cabling required for connection to the host system(s).

Note: If you want to set up the HP® system to boot from the 9900, please see Appendix C.

To connect the 9900 subsystem:

1. **Verify subsystem installation.** The Hitachi Data Systems representative verifies the fibre-port address configuration and the status of the FC adapters and LDEVs (normal).
2. **Shut down the HP-UX® system.** The user should perform this activity. You must shut down and power off the HP-UX® system before connecting the 9900:
 - a) Shut down the HP® system (see Figure 2.4).
 - b) When shutdown is complete, power off the HP® display.
 - c) Power off all peripheral devices except for the 9900 subsystem.
 - d) Power off the HP® system. You are now ready to connect the 9900 subsystem.
3. **Connect the 9900 to the HP® system.** The Hitachi Data Systems representative installs the fibre-channel cables between the 9900 and the HP-UX® system. **Note:** The Hitachi Data Systems representative must use the 9900 maintenance manual during all installation activities. Follow all precautions and procedures in the maintenance manual, and always check all specifications to ensure proper installation and configuration.
4. **Power on the HP® system.** The user or Hitachi Data Systems representative can perform this activity. To power on the HP® system after connecting the 9900 subsystem:
 - a) Power on the HP® system display.
 - b) Power on all peripheral devices. The 9900 subsystem should already be on, the host modes should already be set, and the fibre-channel ports should already be configured. If the host modes or fibre ports are configured after the HP® system is powered on, the system must be restarted to recognize the new devices.
 - c) Confirm the ready status of all devices.
 - d) Power on the HP® system. You are now ready to configure the new 9900 devices.

```
# cd
# shutdown -h now
:
System has halted
OK to turn off power or reset system
UNLESS "WAIT for UPS to turn off power" message was printed above
```

← Shut down the HP® system.

← Power off the system.

Figure 2.4 Shutting Down and Powering Off the HP® System

Chapter 3 Configuring the 9900 Devices

Configuration of the 9900 disk devices is performed by the user and requires root (superuser) access to the HP-UX[®] system. The host modes for the 9900 fibre ports must already be set, the 9900 fibre ports must already be configured, and the host adapter(s) must already be verified (see section 2.3). If any of these steps is performed after the HP[®] system is powered on, you must halt and restart the system before configuring the new devices.

The activities involved in 9900 device configuration are:

- Verifying new device recognition (see section 3.1),
- Verifying the device files and driver (see section 3.2),
- Partitioning the disk devices (see section 3.3),
- Creating the file systems for the new volumes (see section 3.4),
- Setting the device parameters for the disk devices (see section 3.5),
- Creating the mount directories (see section 3.6),
- Mounting and verifying the file systems (see section 3.7), and
- Setting and verifying the auto-mount parameters (see section 3.8).

AL-PA to SCSI TID mapping. For information on the fibre-channel AL-PA to SCSI TID mapping, please refer to Appendix B.

Booting from the 9900. For information on configuring the HP-UX[®] 11.0 boot disk on the 9900 subsystem, please refer to Appendix C.

HP[®] System Administrator Manager (SAM). This chapter provides instructions for configuring the new devices using UNIX[®] commands. For information on using the HP[®] System Administrator Manager (SAM) instead of UNIX[®] commands to configure the 9900 SCSI disk devices, please refer to Appendix D.

Online device installation. For information on configuring newly installed 9900 devices without rebooting the HP[®] system, please refer to Appendix E.

Note on the term "SCSI disk": The 9900 logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

3.1 Verifying New Device Recognition

The first step in configuring the 9900 devices for HP-UX[®] operations is to verify that the HP[®] system recognizes the new devices on the 9900 subsystem. The devices must be installed and formatted and the fibre-channel ports defined before the host system is powered on. If the fibre ports are configured after the HP[®] system is powered on, the user must halt and restart the system to allow the system to recognize the new devices.

To verify that the HP[®] system recognizes the new 9900 devices:

1. Login to the HP[®] system as **root** as shown in Figure 3.1.
2. Display the device data using the **ioscan -f** command, and verify that the system recognizes the newly installed devices on the 9900 subsystem (see Figure 3.2). If desired, use the **-C disk** command option (**ioscan -fnC disk**) to limit the output to disk devices only.

Note: If UNKNOWN is displayed as the Class type, the HP[®] system may not be configured properly. Please refer to the HP[®] documentation or contact HP[®] technical support.

Note: If information for unused devices remains in the system, get the system administrator's permission to renew the device information. To renew the device information, delete the **/etc/ioconfig** and **/stand/ioconfig** files (**rm** command), reboot the server, and then issue the **ioinit -c** command. Now issue the **ioscan -f** command to recognize the logical devices again.

3. Make a blank table (see Table 3.1) for recording the 9900 device data. The table must have seven columns for the following data: bus number, bus instance number, disk number, H/W path, driver, device type, target ID, LUN, and device file name. You will enter the major and minor numbers later.
4. Enter the device data for each 9900 device (disk devices and multiplatform devices) in your table, including the device file name. The device file name has the following structure:
File name = **cXtYdZ**, where **X** = bus instance #, **Y** = target ID, **Z** = LUN.
The "c" stands for controller, the "t" stands for target ID, and the "d" stands for device.
The SCSI target IDs are hexadecimal (0 through F), and the LUN is decimal (0 through 7).
5. Verify that the SCSI TIDs correspond to the assigned port address for all connected ports (see mapping in Appendix B). If so, the logical devices are recognized properly. If not:
 - a) Check the AL-PA for each port using the LUN Manager software. If the same port address is set for multiple ports on the same loop (AL with HUB), all port addresses except one changed to another value, and the relationship between AL-PA and TID does not correspond to the mapping given in Appendix B. Set a different address for each port, reboot the server, and then verify new device recognition again.
 - b) If unused device information remains, the TID-to-AL-PA mapping will not correspond to the mapping given in Appendix B. Renew the device information (see step (2) for instructions), and then verify new device recognition again.

```

The system is ready.

GenericSysName [HP Release B.10.20] (see /etc/issue)
Console Login: root
Password:
Please wait...checking for disk quotas
(c)Copyright 1983-1995 Hewlett-Packard Co., All Rights Reserved.
:
#

```

Figure 3.1 Log In as Root

```

# ioscand -fn
Class      I   H/W Path      Driver      S/W State H/W Type  Description
=====
bc         0           root        CLAIMED    BUS_NEXUS
bc         1   8           bc          CLAIMED    BUS_NEXUS  Bus Converter
fc         0   8/12          fcT1        CLAIMED    INTERFACE  HP Fibre Channel Mass Storage
fcp        0   8/12.8        fcp         CLAIMED    INTERFACE  FCP Protocol Adapter
ext_bus    2   8/12.8.0.255.0 fcpdev      CLAIMED    INTERFACE  FCP Device Interface
target     7   8/12.8.0.255.0.6 tgt          CLAIMED    DEVICE
disk       3   8/12.8.8.255.0.6.0 sdisk       CLAIMED    DEVICE      HITACHI OPEN-9
                        /dev/dsk/c2t6d0 /dev/rdisk/c2t6d0
disk       4   8/12.8.8.255.0.6.1 sdisk       CLAIMED    DEVICE      HITACHI OPEN-9
                        /dev/dsk/c2t6d1 /dev/rdisk/c2t6d1
disk       5   8/12.8.8.255.0.8.0 sdisk       CLAIMED    DEVICE      HITACHI 3390*3B
                        /dev/dsk/c2t8d0 /dev/rdisk/c2t8d0
:
#

```

This sample screen shows the following new 9900 devices recognized:

HITACHI OPEN-9 device: bus number = 8/12, bus instance = 2, target ID = 6, LUN = 0, driver = sdisk.

HITACHI OPEN-9 device: bus number = 8/12, bus instance = 2, target ID = 6, LUN = 1, driver = sdisk.

HITACHI 3390-3B device: bus number = 8/12, bus instance = 2, target ID = 8, LUN = 0, driver = sdisk.

Figure 3.2 Verify New Device Recognition

Table 3.1 Device Data Table (Sample)

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

Note: This table shows the device data for the devices displayed in Figure 3.2.

3.2 Verifying the Device Files and Driver

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

Note: Some HP-compatible systems do not create the device files automatically. If the device files were not created automatically, follow the instructions in section 3.2.1 to create the device files manually.

To verify that the device files for the 9900 devices were successfully created:

1. Display the block-type device files in the **/dev/dsk** directory using the **ll** command (equivalent to **ls -l**) with the output piped to **more** (see Figure 3.3). Verify that there is one block-type device file for each 9900 device.
2. Use your completed device data table (refer to section 3.2.1 and Table 3.2) to verify that the block-type device file name for each device is correct.
3. Display the character-type device files in the **/dev/rdsd** directory using the **ll** command with the output piped to **more** (see Figure 3.4). Verify that there is one character-type device file for each 9900 device.
4. Use your completed device data table (refer to section 3.2.1 and Table 3.2) to verify that the character-type device file name for each device is correct.
5. After verifying the block-type and character-type device files, verify the HP-UX[®] driver for the 9900 subsystem using the **ioscan -fn** command (see Figure 3.5).

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

← Check block-type files.

← Block-type device file.

Figure 3.3 Verifying the Block-Type Device Files

```
# ll /dev/rdsd | more
total 0
crw-r----- 1 bin sys 177 0x000000 Oct 4 11:01 c0t0d0
crw-r----- 1 bin sys 177 0x006000 Dec 6 15:08 c0t6d0
crw-r----- 1 bin sys 177 0x006100 Dec 6 15:08 c0t6d1
Bus instance # = 0, SCSI target ID = 6, LUN = 1 ↗
```

← Check character-type files.

← Character-type device file.

Figure 3.4 Verifying the Character-Type Device Files

```

# ioscan -fn
Class      I   H/W Path          Driver S/W State H/W Type  Description
=====
bc          0                   root   CLAIMED   BUS_NEXUS
bc          1   8                   bc     CLAIMED   BUS_NEXUS Bus Converter
fc          0   8/12                fcT1    CLAIMED   INTERFACE HP Fibre Channel Mass Storage
fcp         0   8/12.8              fcp     CLAIMED   INTERFACE FCP Protocol Adapter
ext_bus     2   8/12.8.0.255.0      fcpdev  CLAIMED   INTERFACE FCP Device Interface
target      7   8/12.8.0.255.0.6    tgt     CLAIMED   DEVICE
disk        3   8/12.8.8.255.0.6.0  sdisk  CLAIMED   DEVICE      HITACHI OPEN-9
                        /dev/dsk/c2t6d0    /dev/rdisk/c2t6d0
disk        4   8/12.8.8.255.0.6.1  sdisk  CLAIMED   DEVICE      HITACHI OPEN-9
                        /dev/dsk/c2t6d1    /dev/rdisk/c2t6d1
disk        5   8/12.8.8.255.0.8.0  sdisk  CLAIMED   DEVICE      HITACHI 3390*3B
                        /dev/dsk/c2t8d0    /dev/rdisk/c2t8d0
:
#

```

Figure 3.5 Verifying the HP-UX® Driver

3.2.1 Creating the Device Files

If the device files were not created automatically when the HP® system was restarted, issue the **insf -e** command in the **/dev** directory (see Figure 3.6) to instruct the HP® system to create the device files. After executing this command, repeat the procedures in sections 3.1 and 3.2 to verify new device recognition and the device files and driver.

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

Figure 3.6 Creating Device Files Automatically

If the device files for the new 9900 devices cannot be created automatically, you must create the device files manually using the **mknod** command as follows:

1. Get your Device Data table on which you recorded the data for the new 9900 devices (Table 3.2 in section 3.2.1). You should already have the following information for all new devices: bus number, bus instance number, disk number, driver, device type, target ID, and LUN.

2. If you have not already done so, construct the device file name for each device, and enter the device file names into your table.

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

3. Construct the minor number for each device, and enter the minor numbers into your table. The minor number is constructed as follows:

0xXXYZ00, where **XX** = bus instance number, **Y** = SCSI target ID, and **Z** = LUN.

4. Display the driver information for the system using the **lsdev** command (see Figure 3.7).
5. Enter the major numbers for the device drivers into your table. You should now have all required device and driver information in your table (see Table 3.2).
6. Create the device files for all 9900 devices (SCSI disk and HMDE devices) using the **mknod** command (see Figure 3.8). Make sure to create the block-type device files in the **/dev/dsk** directory and the character-type device files in the **/dev/rdsk** directory.

Note: The character-type device file **is required** for volumes used as raw devices (e.g., 3390-3A/B/C). The block-type device file is not required for volumes used as raw devices.

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

```
# lsdev
Character      Block      Driver      Class
:             :          :          :
188          31      sdisk      disk
⚡ Major # for character-type files = 188, major # for block-type files = 31
⚡ Driver for HITACHI devices = sdisk
#
```

← Display driver information.

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

Figure 3.7 Displaying the Driver Information

Table 3.2 Completed Device Data Table (Sample)

| Bus No. | Instance (XX) | Disk No. | H/W Path | Driver | Device Type | TID (Y) | LUN (Z) | Device File | Minor # 0xXYZ00 | Major # for Char. Files | Major # for Block Files |
|---------|---------------|----------|--------------------|--------|-------------|---------|---------|-------------|-----------------|-------------------------|-------------------------|
| 8/12 | 02 | 3 | 8/12.8.8.255.0.6.0 | sdisk | OPEN-9 | 6 | 0 | c2t6d0 | 0x026000 | 188 | 31 |
| 8/12 | 02 | 4 | 8/12.8.8.255.0.6.1 | sdisk | OPEN-9 | 6 | 1 | c2t6d1 | 0x026100 | 188 | 31 |
| 8/12 | 02 | 5 | 8/12.8.8.255.0.8.0 | sdisk | 3390-3B | 8 | 0 | c2t8d0 | 0x028000 | 188 | 31 |

```
# cd /dev/dsk
# mknod /dev/dsk/c2t6d0 b 31 0x026000
⚡ File name ⚡ b = block-type, 31 = major #, 0x026000 = minor #
:
# mknod /dev/dsk/c2t8d0 b 31 0x028000
⚡ File name ⚡ b = block-type, 31 = major #, 0x028000 = minor #

# cd /dev/rdisk
# mknod /dev/rdisk/c2t6d0 c 188 0x026000
⚡ File name ⚡ c = character-type, 177 = major #, 0x026000 = minor #
:
# mknod /dev/rdisk/c2t8d0 c 188 0x028000
⚡ File name ⚡ c = character-type, 188 = major #, 0x028000 = minor #
```

← Go to /dev/dsk directory.
← Create block-type file.

← Go to /dev/rdisk directory.
← Create character-type file.

Figure 3.8 Creating the Block-Type and Character-Type Device Files Manually

3.3 Partitioning the Disk Devices

The HP-UX[®] system uses the Logical Volume Manager (LVM) to manage the disk devices on all peripheral storage devices, including the Hitachi Lightning 9900[™] subsystem. Under LVM disk management, a volume group consisting of multiple disks is formed, and then the volume group is divided into logical partitions and managed as a logical volume. These procedures should be executed for all device files corresponding to the new 9900 SCSI disk devices.

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

To partition the 9900 SCSI disk devices for LVM operations, you need to:

- Create a physical volume for each new SCSI disk device (see section 3.3.1).
- Create new volume groups as desired (see section 3.3.2). If you need to increase the maximum volume groups (**maxvgs**) setting, please see section D.2 for instructions.
- Create a logical volume for each new SCSI disk device (see section 3.3.3).

This section provides general instructions and basic examples for partitioning the 9900 SCSI devices for LVM operations using UNIX[®] commands. These instructions do not explicitly cover all LVM configuration issues. For further information on LVM configuration, please refer to the appropriate user documentation, or contact HP[®] technical support.

Note: If desired, the HP[®] System Administrator Manager (SAM) can be used instead of UNIX[®] commands to configure the 9900 SCSI disk devices. See Appendix D for information on using SAM to configure the 9900 devices for LVM operations.

3.3.1 Creating the Physical Volumes

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

Note: Do not create physical volumes for the HMDE devices (e.g., 3390-3A/B/C).

To create the physical volumes for the new 9900 disk devices:

1. Create the physical volume using the **pvcreate** command with the character-type device file as the argument (see Figure 3.9). Make sure to specify the **/dev/rdisk** directory for the character file. You can only create one physical volume at a time.

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

2. Repeat step (1) for each new disk device on the 9900 subsystem.

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

Figure 3.9 Creating Physical Volumes

3.3.2 Creating Volume Groups

After the physical volumes for the disk devices have been created, you can begin creating new volume groups for the new physical volumes as needed. If desired, you can also add any of the new physical volumes on the 9900 subsystem to existing volume groups using the **vgextend** command. The physical volumes which make up one volume group can be located in the same disk subsystem or in different disk subsystems.

Note: You may need to modify the HP® system kernel configuration (**maxvgs** setting) to allow more volume groups to be created. Please see section D.2 in Appendix D for instructions.

Note: Do not assign the HMDE devices (e.g., OPEN-x-HMDEoto) to volume groups

To create a volume group:

1. Display the existing volume groups using the **ls** command (see Figure 3.10).
2. Create the directory for the new volume group using the **mkdir** command (see Figure 3.11). Choose a name for the new volume group which is different than all other group names. Be careful not to use an existing volume group name.
Note: If you need to delete a directory, use the **rmdir** command (e.g., `rmdir /dev/vgnn`).
3. Verify the new directory using the **ls** command (see Figure 3.11).
4. Verify the minor numbers for existing group files using the **ll** command with the output piped to **grep** to display only the files containing “group” (see Figure 3.12).
5. Choose a minor number for the new group file in sequential order (i.e., when existing volume groups are vg00-vg05 and next group name is vg06, use minor number 06 for the vg06 group file). Be careful not to duplicate any minor numbers. Remember that the minor numbers are hexadecimal (e.g., the tenth minor number is 0x0a0000, not 0x100000).
6. Create the group file for the new directory using the **mknod** command (see Figure 3.13). Make sure to specify the correct volume group name, major number, and minor number. The major number for all group files is **64**.
Note: If you need to delete a group file, use the **rm -r** command to delete the group file and the directory at the same time (e.g., `rm -r /dev/vgnn`), and start again at step (2).
7. Repeat steps (5) and (6) for each new volume group.
8. Create the volume group using the **vgcreate** command (see Figure 3.14).

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

For LUSE volumes which contain more than 17 OPEN-8/9 LDEVs or more than 7043 MB (OPEN-8/9*n-CVS), use the **-s** and **-e** physical extent (PE) parameters of the **vgcreate** command (see Figure 3.14). Table 3.3 lists the PE and maximum PE (MPE) parameters for the 9900 LUSE devices.

Note: If you need to delete a volume group, use the **vgremove** command (e.g., `vgremove /dev/vgnn`). If the **vgremove** command does not work because the volume group is not active, use the **vgexport** command (e.g., `vgexport /dev/vgnn`).

9. Verify that the volume group was created correctly using the **vgdisplay** command (see Figure 3.15). The **-v** option displays the detailed volume group information.

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

Figure 3.10 Displaying the Existing Volume Group Names

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

Figure 3.11 Creating and Verifying a Directory for the New Volume Group

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

Figure 3.12 Displaying the Minor Numbers for the Existing Group Files

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

Figure 3.13 Creating the Group File for the New Volume Group

```
# vgcreate /dev/vg06 /dev/dsk/c2t6d0    ← Create new volume group.
      ↖ Vol group name ↖ Device file name

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

Figure 3.14 Creating the New Volume Group

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

Figure 3.15 Verifying the New Volume Group

Table 3.3 PE and MPE Parameters for LUSE Devices

| Device Type | | Physical Extent Size (PE) | Max Number of Physical Extents (MPE) |
|--|-------------|---------------------------|--------------------------------------|
| OPEN-3/8/9/K/E OPEN-3/K*n (n = 2 to 36) OPEN-3/K-CVS OPEN-3/K*n-CVS (n = 2 to 36) | | default | default |
| OPEN-8/9*n | n = 2 to 17 | default | default |
| | n = 18 | 8 | 15845 |
| | n = 19 | 8 | 16725 |
| | n = 20 | 8 | 17606 |
| | n = 21 | 8 | 18486 |
| | n = 22 | 8 | 19366 |
| | n = 23 | 8 | 20247 |
| | n = 24 | 8 | 21127 |
| | n = 25 | 8 | 22007 |
| | n = 26 | 8 | 22888 |
| | n = 27 | 8 | 23768 |
| | n = 28 | 8 | 24648 |
| | n = 29 | 8 | 25529 |
| | n = 30 | 8 | 26409 |
| | n = 31 | 8 | 27289 |
| | n = 32 | 8 | 28170 |
| | n = 33 | 8 | 29050 |
| | n = 34 | 8 | 29930 |
| | n = 35 | 8 | 30810 |
| | n = 36 | 8 | 31691 |
| OPEN-E*n | n = 2 to 9 | default | default |
| | n = 10 | 8 | 17366 |
| | n = 11 | 8 | 19102 |
| | n = 12 | 8 | 20839 |

Table 3.3 PE and MPE Parameters for LUSE Devices (continued)

| Device Type | | Physical Extent Size (PE) | Max Number of Physical Extents (MPE) |
|-----------------------------------|--------------------|---------------------------|--------------------------------------|
| OPEN-E*n (cont.) | n = 13 | 8 | 22576 |
| | n = 14 | 8 | 24312 |
| | n = 15 | 8 | 26049 |
| | n = 16 | 8 | 27786 |
| | n = 17 | 8 | 29522 |
| | n = 18 | 8 | 31259 |
| | n = 19 | 8 | 32995 |
| | n = 20 | 8 | 34732 |
| | n = 21 | 8 | 36469 |
| | n = 22 | 8 | 38205 |
| | n = 23 | 8 | 39942 |
| | n = 24 | 8 | 41679 |
| | n = 25 | 8 | 43415 |
| | n = 26 | 8 | 45152 |
| | n = 27 | 8 | 46889 |
| | n = 28 | 8 | 48625 |
| | n = 29 | 8 | 50362 |
| | n = 30 | 8 | 52098 |
| | n = 31 | 8 | 53835 |
| | n = 32 | 8 | 55572 |
| | n = 33 | 8 | 57308 |
| | n = 34 | 8 | 59045 |
| | n = 35 | 8 | 60782 |
| | n = 36 | 8 | 62518 |
| OPEN-8/9/E-CVS | | default | default |
| OPEN-8/9/E*n-CVS (n = 2 to 36) | 70-119731(MB) × N1 | 8 | default |
| | 119732- (MB) × N1 | 8 | N2 |

N1 = [CVS volume capacity (in MB) from Remote Console PC] × n

N2 = $\uparrow N1 / PE \uparrow$ ($\uparrow \uparrow$ means round up to next integer.)

Example: CVS volume capacity is 6000 MB for OPEN-9*22-CVS volume:

N1 = $6000 \times 22 = 132000$

N2 = $\uparrow 132000 / 8 \uparrow = 16500$

3.3.3 Creating the Logical Volumes

After the desired new volume groups have been created, you can begin creating the logical volumes for each new disk device on the 9900 subsystem.

Note: Do not create logical volumes for the HMDE devices (e.g., 3390-3A/B/C, OPEN-x-HMDEoto).

To create the logical volumes:

1. Create the logical volume using the **lvcreate -L** command, and specify the volume size and volume group for the new logical volume (see Figure 3.16). The HP® system assigns the logical volume numbers automatically (lv01, lv02, lv03, ...). Use the capacity values specified in Table 1.1 for the size parameter (OPEN-K = 1740, OPEN-3 = 2344, OPEN-8 = 7004, OPEN-9 = 7004, OPEN-E = 13888). To calculate **S1** for CVS, LUSE, and CVS LUSE volumes:

Use the **vgdisplay** command to display the physical extent size (**PE Size**) and usable number of physical extents (**Free PE**) for the volume (see Figure 3.17). Calculate the maximum size value (in MB) as follows: $S1 = (PE\ Size) \times (Free\ PE)$

2. Verify that the logical volume was created correctly using the **lvdisplay** command (see Figure 3.18). If desired, you can wait until all logical volumes have been created and then verify all volumes at one time by using the * wildcard character with the **lvdisplay** command (e.g., **lvdisplay /dev/vg06/lvol***).
3. Repeat steps (1) and (2) for each logical volume to be created. You can only create one logical volume at a time, but you can verify more than one logical volume at a time.

Note: To delete a logical volume, use the **lvremove** command (e.g., **lvremove /dev/vgmn/lvolx**). To increase the size of an existing logical volume, use the **lvextend** command (e.g., **lvextend -L size /dev/vgmn/lvolx**). To decrease the size of an existing logical volume, use the **lvreduce** command (**lvreduce -L size /dev/vgmn/lvolx**).

```
# lvcreate -L 2344 /dev/vg06                                ← Create new logical volume.
                               ↖ Size of volume = 2344 MB (OPEN-3)
Logical volume "/dev/vg06/lvol1" has been successfully created with character device
"/dev/vg06/r1vol1".
Logical volume "/dev/vg06/lvol1" has been successfully extended.
Volume Group configuration for /dev/vg06 has been saved in /etc/lvmconf/vg06.cof.
```

Figure 3.16 Creating a Logical Volume

```

# vdisplay /dev/vg01

--- Volume groups ---
VG Name          /dev/vg01
VG Write Access  read/write
VG Status        available
Max LV           255
Cur LV          0
Open LV          0
Max PV           16
Cur PV          1
Act PV           1
Max PE per PV    1016
VGDA             2
PE Size (Mbytes) 4
Total PE         586
Alloc PE         0
Free PE          586
Total PVG        0

```

← *Physical extent size.*

← *Number of physical extents.*

This example shows the following information for /dev/vg01: physical extent size = 4, usable number of physical extents = 586. Therefore, maximum size value = $4 \times 586 = 2344$

Figure 3.17 Calculating the Volume Size for CVS, LUSE, and CVS-LUSE Devices

```

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

```

← *Verify new logical volume.*

← $2344 = 586 \times 4 = \text{OPEN-3}$

← *LE = logical extent*

← *PE = physical extent*

Figure 3.18 Verifying a Logical Volume

3.4 Creating the File Systems

Once the logical volumes have been created, you can create the file system for each new logical volume on the 9900 subsystem. The default file system types are:

- HP-UX[®] OS version 10.20 = hfs or vxfs, depending on entry in **/etc/defaults/fs** file.
- HP-UX[®] OS version 11.0 = vxfs

Note: Do not create file systems for the HMDE devices (e.g., 3390-3A/B/C, OPEN-x-HMDEoto).

To create the file system on a new logical volume:

1. Create the file system using the **newfs** command with the logical volume as the argument. Figure 3.19 shows an example of creating the file system for an OPEN-3 volume. Figure 3.20 shows an example of creating the file system for an OPEN-9 volume. Figure 3.21 shows examples of specifying the file system type (vxfs or hfs) with the **newfs** command.
2. Repeat step (1) for each new logical volume on the 9900 subsystem.

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

Figure 3.19 Creating a File System (default file system, OPEN-3 shown)

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

Figure 3.20 Creating a File System (default file system, OPEN-9 shown)

```
# newfs -F vxfs /dev/vg06/rlvol1                        ← Specify file system type.
:
# newfs -F hfs /dev/vg06/rlvol12
```

Figure 3.21 Specifying the File System Type

3.5 Setting the Device Parameters

When the device files are created, the HP-UX[®] system sets the IO time-out parameter to its default value (20 seconds) and the queue depth parameter to its default value (2 or 8). You need to change these values for all new disk devices on the 9900 subsystem.

Note: Do not change the device parameters for the HMDE devices (e.g., 3390-3A/B/C).

3.5.1 Setting the IO Time-Out Parameter

The IO time-out parameter for the disk devices on the 9900 subsystem must be set to **60 seconds**. To change the IO time-out parameter:

1. Verify the current IO time-out value using the **pvdisplay** command (see Figure 3.22).
2. Change the IO time-out value to 60 using the **pvchange -t** command (see Figure 3.23).
3. Verify that the new IO time-out value is 60 seconds (use **pvdisplay**) (see Figure 3.24).
4. Repeat steps (1) through (3) for each new disk device on the 9900 subsystem.

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

Figure 3.22 Checking the Current IO Time-Out Value

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

Figure 3.23 Changing the IO Time-Out Value

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

Figure 3.24 Verifying the New IO Time-Out Value

3.5.2 Setting the Queue Depth Parameter

The HP-UX[®] system automatically sets the queue depth to the default value, which may be either 2 or 8 depending on the installed HP[®] options and drivers. The queue depth for the 9900 disk devices must be set as specified in Table 3.4. The **scsictl** command allows you to view and change the queue depth parameter for each device. However, this command only allows you to change the queue depth for one volume at a time, and the queue depth will be reset to the default value the next time the system is restarted. Therefore, you also need to create and register a start-up script to set the queue depth for the 9900 disk devices each time the system is restarted.

Note: Do not set the queue depth for the HMDE devices (e.g., 3390-3A/B/C).

Table 3.4 Queue Depth Requirements for the 9900 Devices

| Parameter | Required Value |
|----------------------|----------------|
| Queue depth per LU | ≤ 8 |
| Queue depth per port | ≤ 256 |

Note: You can adjust the queue depth for the 9900 devices later as needed (within the specified range) to optimize the I/O performance of the 9900 devices.

To set the queue depth parameter for the 9900 devices:

1. If you cannot shut down and restart the system at this time, use the **scsictl** command to set the queue depth for each 9900 device (see Figure 3.25). The **scsictl** commands to set queue depth should be registered as HP-UX[®] start-up script for future reboot.
2. Check the **/sbin/init.d** and **/sbin/rc1.d** directories to see if the script name **queue** is already used (link name **Sxxxqueue** or **Kxxxqueue**) (see Figure 3.26). Choose a unique name for the start-up script as follows:
 - a) If there is no script named **queue** and no link file named **Sxxxqueue** or **Kxxxqueue**, use the name **queue** for the new script, and go to step (3).
 - b) If the script **queue** and the link file **Sxxxqueue** or **Kxxxqueue** already exist and the script is already used for setting the queue depth for other previously installed 9900 subsystems, check the script file to make sure that the queue depth is set to the desired number (per Table 3.4), and add a line for each new 9900 disk device. If necessary, restart the HP[®] system to set the queue depth for the new volumes.
 - c) If the script **queue** and the link file **Sxxxqueue** or **Kxxxqueue** already exist and the script is not used for setting the queue depth for the 9900, use some other name for the new queue-depth script for the 9900 (e.g., **hitachi_q**), and go to step (3).

Note: If the link **Sxxxqueue** and/or **Kxxxqueue** already exists but there is no script file named **queue**, delete the link file(s), use the name **queue** for the new script, and go to step (3).
3. Choose a unique three-digit number for the link name. This number cannot be used in any other links. The link name is derived as follows: **S** stands for “start up script,” **K** stands for “kill script,” the three-digit number is unique to each link, and the script file name follows the three-digit number (e.g., **S890queue** or **S890hitachi_q**).

4. Create and register the new start-up script for the 9900. See section 3.5.3 for an example of the start-up script.
5. Shut down and restart the HP® system, so that the new start-up script sets the queue depth for the 9900 disk devices to the specified value (per Table 3.4).
6. After restarting the system or setting the queue depths manually, verify the queue depth for each 9900 disk device using the **scsictl** command (see Figure 3.27).

```
# /usr/sbin/scsictl -m queue_depth=2 -a /dev/rdisk/c0t6d0 ← Set queue depth per Table 3.4.
                                     ↩ Character-type device file
# /usr/sbin/scsictl -m queue_depth=2 -a /dev/rdisk/c0t6d1
# /usr/sbin/scsictl -m queue_depth=2 -a /dev/rdisk/c0t6d2
# /usr/sbin/scsictl -m queue_depth=2 -a /dev/rdisk/c0t6d3
:
:
# /usr/sbin/scsictl -m queue_depth=8 -a /dev/rdisk/c0t8d0
```

Figure 3.25 Changing the Queue Depth for the 9900 Disk Devices

```
# ls /sbin/init.d                                     ← Check for QUEUE.
OspfMib      clean_ex  dfs      hpether   names     nis.server savecore  swconfig
SnmpHpunix   clean_tmps diagnostic iforls    ncs       pd        sendmail  syncer
:
clean_adm    ddfa      hparray   mrouterd  nis.client rwhod     swcluster xntpd
# ls /sbin/rc1.d                                     ← Check for SxxxQUEUE and KxxxQUEUE.
K230audio    K340xntpd  K420dfs   K475rarpd K630named  S420set_date
K240auditing K356vjed   K430dce   K480rdpd  K660net    S440savecore
K250envd     K358egcd   K435OspfMib K490gated K700nettl  S500swap_start
K258diagnostic K360kks    K435SnmpHpunix K500inetd K770ptydaemon S520syncer
K270cron     K370vt     K435SnmpMib2 K510mrouterd K780syslogd
K278pd       K380xfs    K440SnmpMaster K570nfs.client K900swagentd
K280lp       K390rbootd K450ddfa   K580nis.client S100localmount
K290hparray   K400iforls K460sendmail K590nis.server S320hostname
K300acct     K410ncs    K470rwhod  K600nfs.core S400set_prvgrp
```

Figure 3.26 Checking the Existing Script Names and Link Names

```
# /usr/sbin/scsictl -a /dev/rdisk/c0t6d0 ← Verify new queue depth.
                                     ↩ Character-type device file
immediate_report = 0; queue_depth = 2 ← Queue depth = 2.
:
:
# /usr/sbin/scsictl -a /dev/rdisk/c0t8d0 ← Verify new queue depth.
immediate_report = 0; queue_depth = 8 ← Queue depth = 8.
```

Figure 3.27 Verifying the Queue Depth for the 9900 Disk Devices

3.5.3 Creating and Registering the Queue Depth Start-Up Script

The **queue** (or **hitachi_q**) start-up script sets the queue depth to 2 for all new volumes (SCSI disk devices) on the 9900 subsystem each time the HP® system is restarted. If the **queue** script already exists for an older (previously installed) 9900 subsystem, check the script file to verify that the queue depth value is set to the desired value (see Table 3.4), and add a line for each new 9900 volume. If the script does not already exist, you must create and register the script as shown below. You can use the UNIX® **vi** editor or any other similar text editor.

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

```
# cp /sbin/init.d/template /sbin/init.d/queue          ← Copy start-up script template file.
# vi /sbin/init.d/queue                                ← Edit script file as shown below.
-----file(/sbin/init.d/queue)-----
# !/sbin/sh
#
# @(#) $Revision: 78.1 $
#
# NOTE:   This script is not configurable!  Any changes made to this
#         script will be overwritten when you upgrade to the next
#         release of HP-UX.
#
# WARNING: Changing this script in any way may lead to a system that
#         is unbootable.  Do not modify this script.
#
# <Insert comment about your script here>
#
# Allowed exit values:
#   0 = success; causes "OK" to show up in checklist.
#   1 = failure; causes "FAIL" to show up in checklist.
#   2 = skip; causes "N/A" to show up in the checklist.
#       Use this value if execution of this script is overridden
#       by the use of a control variable, or if this script is not
#       appropriate to execute for some other reason.
#   3 = reboot; causes the system to be rebooted after execution.
# Input and output:
#   stdin is redirected from /dev/null
#   stdout and stderr are redirected to the /etc/rc.log file
#   during checklist mode, or to the console in raw mode.
#
PATH=/usr/sbin:/usr/bin:/sbin
export PATH
#
# NOTE: If your script executes in run state 0 or state 1, then /usr
#       might not be available.  Do not attempt to access commands or
#       files in /usr unless your script executes in run state 2 or
#       greater.  Other file systems typically not mounted until run
#       state 2 include /var and /opt.
#
rval=0
#
# Check the exit value of a command run by this script.  If non-zero,
# the exit code is echoed to the log file and the return value of this
# script is set to indicate failure.
```

```

set_return() {
    x=$?
    if [ $x -ne 0 ]; then
        echo "EXIT CODE: $x"
        rval=1 # script FAILED
    fi
}
# Kill the named process(es).
# $1=<search pattern for your process>

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

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

```

← Edit text here.

```

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

    # Check to see if this script is allowed to run...
    if [ "$CONTROL_VARIABLE" != 1 ]; then
        rval=2
    else

        # Execute the commands to stop your subsystem
        :
    fi

```

← Delete these lines.

```

        /usr/sbin/scsictl -m queue_depth=2 /dev/rdisk/c0t6d0
        /usr/sbin/scsictl -m queue_depth=2 /dev/rdisk/c0t6d1
        /usr/sbin/scsictl -m queue_depth=2 /dev/rdisk/c0t8d0
        :
    ;;
'stop')

```

← Add one line for each
← new disk device.
←

```

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

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

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

# ls /sbin/rc1.d
K230audio      K340xntpd      K420dfs        K475rarpd      K630named      S420set_date
K240auditing   K356vjed       K430dce        K480rdpd       K660net        S440savecore
K250envd       K358egcd       K4350spfMib    K490gated      K700nettl      S500swap_start
K258diagnostic K360kks        K435SnmpHpunix K500inetd      K770ptydaemon  S520syncer
K270cron       K370vt         K435SnmpMib2   K510mrouted    K780syslogd
K278pd        K380xfs        K440SnmpMaster K570nfs.client K900swagentd
K280lp        K390rbootd     K450ddfa       K580nis.client S100localmount
K290hparrray   K400iforls     K460sendmail   K590nis.server S320hostname
K300acct       K410ncs        K470rwhod      K600nfs.core   S400set_prvgrp

# ln -s /sbin/init.d/queue /sbin/rc1.d/S890queue

```

← **Check link names.**

← **Create link file.**

↩ **Make sure this file name does not already exist.**

3.6 Creating the Mount Directories

After the file systems have been created and the device parameters have been set, the next step in configuring the new 9900 volumes is to create the mount directory for each volume. Make sure to choose a unique name for each mount directory which identifies the logical volume.

To create the mount directories:

1. Create the mount directory using the **mkdir** command with the new mount directory name as the argument (see Figure 3.28).

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

2. Verify the new mount directory using the **ls -x** command (see Figure 3.28).
3. Repeat steps (1) and (2) for each new device on the 9900 subsystem.

| | |
|--------------------|-------------------------------|
| # mkdir /9900-LU00 | ← Create new mount directory. |
| # ls -x | ← Verify new mount directory. |
| 9900-LU00 | bin dev device etc export |
| floppy | home hstsboof kadb kernel lib |
| # | |

Figure 3.28 Creating and Verifying a Mount Directory

3.7 Mounting and Verifying the File Systems

After the mount directories have been created, you can mount the file system for each new logical volume and then verify the file systems.

To mount and verify the file systems:

1. Use the **mount** command to mount the file system for the volume (see Figure 3.29).
Note: If you need to unmount a file system, use **umount** (e.g., `umount /9900-LU00`).
2. Repeat step (1) for each new logical volume on the 9900 subsystem.
3. Verify that the file systems are correct using the **bdf** command (see Figure 3.30). Make sure that the capacity (listed under **Kbytes**) is correct for each device.
4. As a final verification, perform some basic UNIX[®] operations (e.g., file creation, copying, and deletion) on each logical device to make sure that the new devices on the 9900 subsystem are fully operational (see Figure 3.31).

```
# mount /dev/vg06/lvol1 /9900-LU00                                ← Mount file system.
      ↵ Block-type lvol name ↵ Mount directory name
#
```

Figure 3.29 Mounting a File System

```
# bdf                                                                ← Verify file systems.
Filesystem      Kbytes    used    avail  %used  Mounted on
/dev/vg00/lvol1  59797    59364      0   100%  /
:
/dev/vg06/lvol1  2348177      9 2113350    0%  /9900-LU00    ← OPEN-3
/dev/vg07/lvol1  2348177      9 2113350    0%  /9900-LU01    ← OPEN-3
/dev/vg08/lvol1  7052764      9 6347478    0%  /9900-LU02    ← OPEN-9
```

Figure 3.30 Verifying the File Systems

```
# mount /dev/vg06/lvol1 /9900-LU00                                ← Mount LUN.
# cd /9900-LU00                                                    ← Go to LUN mount directory.
# cp /bin/vi /9900-LU00/vi.back1                                   ← Copy any file to LUN.
# ll                                                                ← Verify file copy.
drwxr-xr-t  2 root    root      8192 Mar 15 11:35 lost+found
-rwxr-xr-x  1 root    sys      217088 Mar 15 11:41 vi.back1
# cp vi.back1 vi.back2                                           ← Copy file again.
# ll                                                                ← Verify second file copy.
drwxr-xr-t  2 root    root      8192 Mar 15 11:35 lost+found
-rwxr-xr-x  1 root    sys      217088 Mar 15 11:41 vi.back1
-rwxr-xr-t  1 root    sys      217088 Mar 15 11:52 vi.back2
# rm vi.back1                                                     ← Delete first test file.
# rm vi.back2                                                     ← Delete second test file.
```

Figure 3.31 Final Verification of the File System for One Volume

3.8 Setting and Verifying the Auto-Mount Parameters

The final step in configuring the 9900 volumes for LVM operations is to set up and verify the auto-mount parameters for each new volume. The **/etc/fstab** file contains the auto-mount parameters for the logical volumes. If you do not plan to auto-mount the 9900 devices, you can skip this section.

To set and verify the auto-mount parameters:

1. Edit the **/etc/fstab** file to add a line for each new volume (SCSI disk device) on the 9900 subsystem (see Figure 3.32). Table 3.5 shows the auto-mount parameters.
2. Reboot the HP® system after you are finished editing the **/etc/fstab** file. If you cannot reboot at this time, issue the **mount -a** command.
3. Verify the device file systems again using the **bdf** command (refer to Figure 3.30).

| | | | | | | | |
|---|-----------------|------------|--------|----------|---|---|-------------------------------|
| # cp -ip /etc/fstab /etc/fstab.standard | | | | | | | ← Make backup before editing. |
| # vi /etc/fstab | | | | | | | ← Edit the file (vi shown). |
| | /dev/vg00/lvol1 | / | hfs | rw | 0 | 1 | # root |
| | /dev/vg00/lvol2 | swap | ignore | sw | 0 | 0 | # primary swap |
| | : | | | | | | |
| | /dev/vg06/lvol1 | /9900-LU00 | hfs | defaults | 0 | 2 | # 9900-LU00 |
| | /dev/vg06/lvol2 | /9900-LU01 | vxfs | defaults | 0 | 2 | # 9900-LU01 |
| ① | ② | ③ | ④ | ⑤ | ⑥ | ⑦ | ← See Table 3.5. |

Figure 3.32 Setting the Auto-Mount Parameters

Table 3.5 Auto-Mount Parameters

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

Chapter 4 Failover and SNMP Configuration

The 9900 subsystem supports industry-standard products and functions which provide host and/or application failover, I/O path failover, and logical volume management (LVM). For the HP-UX[®] environment, the 9900 subsystem currently supports the following products and functions (please contact your Hitachi Data Systems representative for the latest information):

- HP[®] Multi-Computer/ServiceGuard (MC/ServiceGuard) software for application failover.
- Alternate link for I/O path failover (included in the HP-UX[®] OS).
- Logical volume management (included in the HP-UX[®] OS).

The 9900 subsystem also supports the industry-standard simple network management protocol (SNMP) for remote subsystem management from the UNIX[®]/PC server host. SNMP is used to transport management information between the 9900 subsystem and the SNMP manager on the host. The SNMP agent for the 9900 subsystem sends status information to the host(s) when requested by the host or when a significant event occurs.

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, please refer to the user documentation, or contact the vendor's technical support.

4.1 Application Failover

The 9900 subsystem supports the MC/ServiceGuard application failover feature of the HP-UX[®] operating system. The MC/ServiceGuard software is sold separately. Please contact HP[®] for the latest information on the MC/ServiceGuard product.

After 9900 device configuration is complete, make sure to configure the MC/ServiceGuard software on the host server(s) as needed to recognize the new devices on the 9900 subsystem. For assistance with MC/ServiceGuard operations, please refer to the user documentation, or contact HP[®] technical support.

4.2 I/O Path Failover

The 9900 subsystem supports the alternate link path failover function of the HP-UX[®] operating system. After 9900 device configuration is complete, make sure to configure alternate links to the 9900 devices as needed using the **vgextend** command. For assistance with alternate link operations, please refer to the HP-UX[®] user documentation (UNIX[®] man pages), or contact HP[®] technical support.

4.3 SNMP Remote Subsystem Management

SNMP is a part of the TCP/IP protocol suite that supports maintenance functions for storage and communication devices. The 9900 subsystem utilizes SNMP to transfer status and management commands to the UNIX®/PC server host via the 9900 SVP (see Figure 4.1). When the SNMP manager requests status information or when a service information message (SIM) occurs, the SNMP agent on the 9900 SVP notifies the SNMP manager on the UNIX®/PC server host. Notification of 9900 error conditions is made in real time, providing UNIX® and PC server users with the same level of monitoring and support available to S/390® mainframe users. The SIM reporting via SNMP enables the user to monitor the 9900 subsystem from the UNIX®/PC server host without having to check the Remote Console PC for remote SIMs (R-SIMs).

When a SIM occurs, the 9900 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 UNIX®/PC server host. For assistance with SNMP manager configuration on the UNIX®/PC server host, please refer to the user documentation, or contact the vendor's technical support.

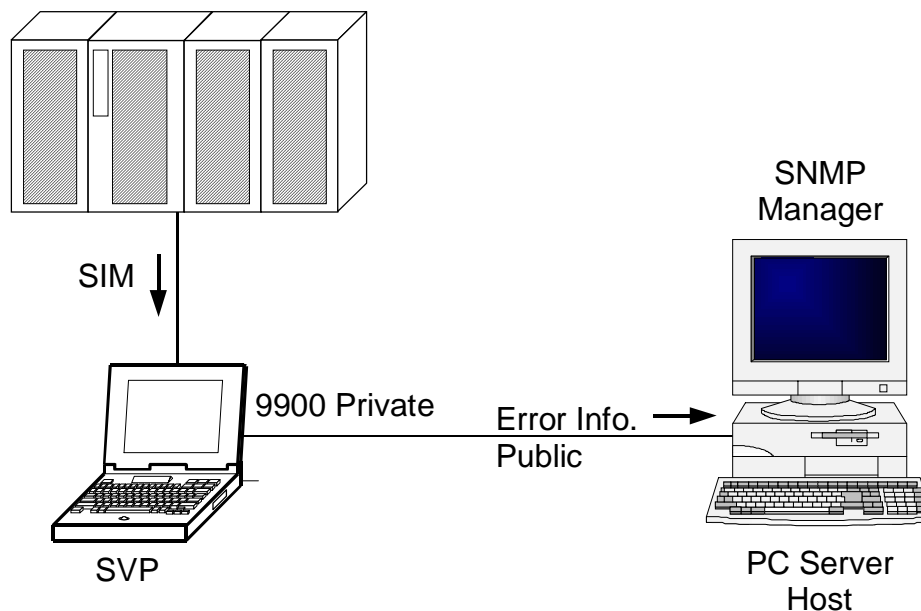


Figure 4.1 9900 SNMP Environment

Chapter 5 Troubleshooting

5.1 Troubleshooting

The Hitachi Lightning 9900™ disk array subsystem provides continuous data availability. For troubleshooting information for the 9900 subsystem, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide* (MK-90RD008).

Table 5.1 lists potential error conditions during HP-UX® installation and configuration and provides instructions for resolving each condition. If you are unable to resolve an error condition, please consult your Hitachi Data Systems representative, 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 9900 subsystem are ON. Make sure that the fibre-channel cables are correctly installed and firmly connected. Make sure that the LUNs are properly configured. The LUNs for each SCSI TID must start at 0 and continue sequentially without skipping any numbers. |
| A physical volume cannot be created (PVCREATE command). | Make sure that the 9900 devices are properly formatted. Make sure that the character-type device file exists. Make sure that the correct character-type device file name is used with pvcreate . |
| A volume group cannot be created (VGCREATE command). | Make sure that the directory for the new volume group exists. Make sure that the control file exists. Make sure that the correct major # (64) and minor # are used with mknod . Make sure that the block-type file exists and is entered correctly with vgcreate . Make sure that the physical volume is not already allocated to another volume group. |
| A logical volume cannot be created (LVCREATE command). | Make sure that the specified capacity is not greater than 4096 MB. Make sure that the capacity of the volume group is not less than the capacity of the partitioned logical volume. |
| File system cannot be created (newfs). | Make sure that the character-type device file is entered correctly with newfs . |
| The file system is not mounted after rebooting. | Make sure that the system was restarted properly. Make sure that the auto-mount information in the /etc/fstab file is correct. |
| The HP® system does not reboot properly after hard shutdown. | If the HP® system is powered off without executing the shutdown process, wait three minutes before restarting the HP® system. This allows the 9900's internal time-out process to purge all queued commands so that the 9900 is available (not busy) during system startup. If the HP® system is restarted too soon, the 9900 will continue trying to process the queued commands, and the HP® system will not reboot successfully. |

5.2 Calling the Hitachi Data Systems Support Center

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including the circumstances surrounding the error or failure and the exact content of any error messages displayed on the Remote Console PC and/or logged at the host. Please check the R-SIMs (remote service information messages) logged on the Remote Console PC, and note the reference codes and severity levels of the recent R-SIMs.

The worldwide Hitachi Data Systems Support Centers are:

- Hitachi Data Systems North America/Latin America
San Diego, California, USA
1-800-348-4357
- Hitachi Data Systems Europe
Contact Hitachi Data Systems Local Support
- Hitachi Data Systems Asia Pacific
North Ryde, Australia
011-61-2-9325-3300

Appendix A Acronyms and Abbreviations

| | |
|---------|---|
| AL | arbitrated loop |
| AL-PA | arbitrated loop physical address |
| CVS | custom volume size (also called Virtual LUN) |
| ESCON | Enterprise System Connection (IBM trademark for optical channels) |
| FC | fibre-channel |
| FCA | fibre-channel adapter |
| FC-AL | fibre-channel arbitrated loop |
| FCP | fibre-channel protocol |
| H/W | hardware |
| HMDE | Hitachi Multiplatform Data Exchange |
| HP | Hewlett-Packard Company |
| I/O, IO | input/output |
| LDEV | logical device |
| LU | logical unit |
| LUN | logical unit number, logical unit |
| LUSE | LU Size Expansion |
| LVM | Logical Volume Manager, logical volume management |
| MPE | maximum number of physical extents |
| OFC | open fibre control |
| P-P | point-to-point |
| PA | physical address |
| PC | personal computer system |
| PE | physical extent |
| PV | physical volume |
| R-SIM | remote service information message |
| RAID | redundant array of independent disks |
| SAM | System Administration Manager |
| SCSI | small computer system interface |
| SIM | service information message |
| SNMP | simple network management protocol |
| SVP | service processor |
| TID | target ID |

Appendix B SCSI TID Map 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 HP-UX[®] 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/rdisk/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 B.1 identifies the fixed mappings between the TID values assigned by the HP-UX[®] system and the FC native addresses (AL_PA/SEL_ID) for FC adapters. For each device file (/dev/dsk/c*t*d*), the **c** value is the adapter number, and the **t** value is the target ID. The **c** value depends on the server configuration; a different value is assigned per each column of Table B.1.

Note: The mapping defined in Table B.1 cannot be guaranteed under the following conditions:

- When 9900 devices and other types of devices are connected in the same loop,
- When information for unused devices remains in server system, or
- When multiple ports participate in the same arbitrated loop.

Table B.1 AL-PA to SCSI TID Mapping (t value) for HP-UX® Systems

| AL-PA | tvalue | AL-PA | tvalue | AL-PA | tvalue | AL-PA | tvalue | AL-PA | tvalue | AL-PA | tvalue | AL-PA | tvalue |
|-------|--------|-------|--------|-------|--------|-------|--------|-------|--------|-------|--------|-------|--------|
| EF | 0 | CD | 0 | B2 | 0 | 98 | 0 | 72 | 0 | 55 | 0 | 3A | 0 |
| E8 | 1 | CC | 1 | B1 | 1 | 97 | 1 | 71 | 1 | 54 | 1 | 39 | 1 |
| E4 | 2 | CB | 2 | AE | 2 | 90 | 2 | 6E | 2 | 53 | 2 | 36 | 2 |
| E2 | 3 | CA | 3 | AD | 3 | 8F | 3 | 6D | 3 | 52 | 3 | 35 | 3 |
| E1 | 4 | C9 | 4 | AC | 4 | 88 | 4 | 6C | 4 | 51 | 4 | 34 | 4 |
| E0 | 5 | C7 | 5 | AB | 5 | 84 | 5 | 6B | 5 | 4E | 5 | 33 | 5 |
| DC | 6 | C6 | 6 | AA | 6 | 82 | 6 | 6A | 6 | 4D | 6 | 32 | 6 |
| DA | 7 | C5 | 7 | A9 | 7 | 81 | 7 | 69 | 7 | 4C | 7 | 31 | 7 |
| D9 | 8 | C3 | 8 | A7 | 8 | 80 | 8 | 67 | 8 | 4B | 8 | 2E | 8 |
| D6 | 9 | BC | 9 | A6 | 9 | 7C | 9 | 66 | 9 | 4A | 9 | 2D | 9 |
| D5 | 10 | BA | 10 | A5 | 10 | 7A | 10 | 65 | 10 | 49 | 10 | 2C | 10 |
| D4 | 11 | B9 | 11 | A3 | 11 | 79 | 11 | 63 | 11 | 47 | 11 | 2B | 11 |
| D3 | 12 | B6 | 12 | 9F | 12 | 76 | 12 | 5C | 12 | 46 | 12 | 2A | 12 |
| D2 | 13 | B5 | 13 | 9E | 13 | 75 | 13 | 5A | 13 | 45 | 13 | 29 | 13 |
| D1 | 14 | B4 | 14 | 9D | 14 | 74 | 14 | 59 | 14 | 43 | 14 | 27 | 14 |
| CE | 15 | B3 | 15 | 9B | 15 | 73 | 15 | 56 | 15 | 3C | 15 | 26 | 15 |
| | | | | | | | | | | | | | |

Appendix C Configuring the Boot Disk on the 9900

You can configure the HP-UX[®] 11.0 boot disk on the 9900 subsystem under the following conditions:

- **OS version:** The HP-UX[®] OS version must be 11.0.
Firmware version: The HP[®] 9000 firmware (BIOS) version must be 38.40 or later.
- The HP-UX[®] system must recognize the 9900 LU which will be the HP-UX[®] boot disk.
- The HP-UX[®] system must also recognize LU 0 which is in the same hardware path group as the 9900 boot disk.

Note: In the following instructions, “LU X” refers to the 9900 LU on which the HP-UX[®] 11.0 operating system will be installed (the boot disk), and X represents the LU number (LUN) of the LU (e.g., if LU 0 will be the boot disk, X = 0; if LU 3 will be the boot disk, X = 3).

To set up the HP-UX[®] boot disk on the 9900 subsystem:

1. **Checking the status and hardware path of LU X** (see Figure C.1). Check the status and hardware path of the 9900 LU X (LU X will be the new boot disk) using the **ioscan -fnC disk** command. The status of LU X must be **claimed**, and the hardware path of LU X must exist. Also check the status and hardware path of LU 0 in the same hardware path group as LU X. The status of LU 0 must be **claimed**, and the hardware path of LU 0 must exist. If you verified both LU X and LU 0, write down the hardware path of LU X, and shut down the HP[®] 9000 server (**shutdown -h now**).
- Note:** If you cannot verify LU 0, you cannot set up the 9900 as the boot disk, because HP-UX[®] requires the existence of LU 0 in the same hardware path group as the boot disk.
2. **Installing HP-UX[®] 11.0 on the 9900.** After shutting down the HP[®] system, install the HP-UX[®] OS version 11.0 on LU X of the 9900. For information on HP-UX[®] 11.0 installation, please refer to the HP-UX[®] 11.0 user documentation. When HP-UX[®] OS installation is complete, shut down the HP[®] 9000 server (**shutdown -h now**).
3. **Confirming the HP[®] 9000 BIOS configuration** (see Figure C.2). After shutdown is complete, restart the HP[®] 9000 server. If the **To discontinue, press any key within 10 seconds** message is displayed, press the **Esc** button. Confirm that the HP[®] 9000 firmware (BIOS) version is 38.40 or later. If not, you cannot set up the 9900 as the HP-UX[®] boot disk.

Note: If you did not press the **Esc** button within 10 seconds, the system boots from the old boot disk, and you will not be able to set up the 9900 boot disk. If this happens, you must reboot the system (**reboot**) and make sure to stop the boot process using the **Esc** button.

4. **Setting the New Boot Disk** (see Figure C.3). After confirming the firmware (BIOS) version, enter the **path** command at the **Main Menu** prompt, and then enter the **path primary** command specifying the hardware path of LU X (refer to step 1). After the **path primary** command is entered, the system displays back the description. Confirm the new configuration by entering the **path** command again. If the configuration is not correct, set the new configuration again using the **path primary** command. If the configuration is correct, turn off the HP[®] 9000 server's power switch.

(Instructions continue on the next page.)

5. **Booting from the 9900.** Wait a few minutes after power-off, and then restart the HP® 9000 server. Confirm that the HP® system is booting from LU X on the 9900 (see Figure C.4).

```
# ioscan -fnC disk
```

| Class | I | H/W Path | Driver | S/W State | H/W Type | Description |
|-------|-----|------------------|-----------------|-------------------|----------|--------------------------|
| ===== | | | | | | |
| disk | 0 | 8/4.5.0 | sdisk | CLAIMED | DEVICE | SEAGATE ST34573WC |
| | | | /dev/dsk/c0t5d0 | /dev/rdisk/c0t5d0 | | |
| disk | 1 | 8/4.8.0 | sdisk | CLAIMED | DEVICE | SEAGATE ST34573WC |
| | | | /dev/dsk/c0t8d0 | /dev/rdisk/c0t8d0 | | |
| disk | 2 | 8/4.9.0 | sdisk | CLAIMED | DEVICE | SEAGATE ST34573WC |
| | | | /dev/dsk/c0t9d0 | /dev/rdisk/c0t9d0 | | |
| disk | 4 | 8/12.8.0.3.0.0.0 | sdisk | CLAIMED | DEVICE | HITACHI OPEN-3 |
| | | | /dev/dsk/c4t0d0 | /dev/rdisk/c4t0d0 | | |
| disk | 317 | 8/12.8.0.3.0.0.1 | sdisk | CLAIMED | DEVICE | HITACHI OPEN-3 |
| | | | /dev/dsk/c4t0d1 | /dev/rdisk/c4t0d1 | | |
| disk | 318 | 8/12.8.0.3.0.0.2 | sdisk | CLAIMED | DEVICE | HITACHI OPEN-3 |
| | | | /dev/dsk/c4t0d2 | /dev/rdisk/c4t0d2 | | |
| | . | . | . | . | . | . |
| | . | . | . | . | . | . |
| disk | 323 | 8/12.8.0.3.0.0.7 | sdisk | CLAIMED | DEVICE | HITACHI OPEN-3 |
| | | | /dev/dsk/c4t0d7 | /dev/rdisk/c4t0d7 | | |
| disk | 3 | 8/16/5.2.0 | sdisk | CLAIMED | DEVICE | TOSHIBA CD-ROM XM-5701TA |
| | | | /dev/dsk/c1t2d0 | /dev/rdisk/c1t2d0 | | |
| # | | | | | | |

Figure C.1 Checking the Status and Path of LU X and LU O

| | | |
|--|-------------------|--|
| Firmware Version | 38.40 | ← Verify firmware (BIOS) version. |
| Duplex Console IO Dependent Code (IODC) revision 1 | | |
| | | |
| (c) Copyright 1996-1999, Hewlett-Packard Company All rights reserved | | |
| Processor Speed | Coprocessor State | Cache Size |
| | | |
| 160 MHz | Functional | 64 KB |
| | | 1 MB additional |
| Model: D330/1 | | |
| Available memory (bytes) : 536870912 | | |
| Primary boot path: 8/4.5 (dec) | | |
| Alternate boot path: 8/16/5.0 (dec) | | |
| Console path: 8/0/0.0 (dec) | | |
| Keyboard path: 8/16/7.0 (dec) | | |
| Processor is booting from first available device. | | |
| To discontinue, press any key within 10 seconds. | | |
| | | ← Press Esc within 10 seconds. |

Figure C.2 Confirming the Firmware (BIOS) Version

```

Main Menu: Enter command > path                                ← Enter path command.

Primary boot path : 8/4.5 (dec)                                  ← Default configuration.
                   08/04.5 (hex)
Alternate boot path : 8/16/5.0 (dec)
                   08/10/05.0 (hex)
Console path : 8/0/0.0 (dec)
               08/0/0.0 (hex)
Keyboard path : 8/16/7.0 (dec)
               08/10/07.0 (hex)

Main Menu: Enter command > path primary 8/12.8.3.0.0.X        ← Set new boot disk.

If boot disk = LU 0:
  Primary boot path : 8/12.8.3 (dec)                            ← Special short expression for LU 0.
                   08/0c.8.3 (hex)
If boot disk = LU X:
  Primary boot path : 8/12.8.3.0.0.X (dec)
                   08/0c.8.3.0.0.X (hex)

Main Menu: Enter command > path                                ← Verify new configuration.

If boot disk = LU 0:
  Primary boot path : 8/12.8.3 (dec)                            ← Special short expression for LU 0
                   08/0c.8.3 (hex)
If boot disk = LU X:
  Primary boot path : 8/12.8.3.0.0.X (dec)                      ← New configuration.
                   08/0c.8.3.0.0.X (hex)
Alternate boot path : 8/16/5.0 (dec)
                   08/10/05.0 (hex)
Console path : 8/0/0.0 (dec)
               08/0/0.0 (hex)
Keyboard path : 8/16/7.0 (dec)
               08/10/07.0 (hex)

```

Figure C.3 Setting the New Boot Disk

```

Trying Primary Boot Path
.....
Booting ...
Boot IO Dependent Code (IODC) revision 144

HARD Booted.

ISL booting hpux

Boot
:disc(8/12.8.0.3.0.0.X)/stand/vmunix.                          ← Verify new boot disk.

4325948 + 561152 + 368416 start 0x175b68

```

Figure C.4 Confirming Boot from 9900 LU X

Appendix D Reference Information for SAM

The HP® System Administrator Manager (SAM) (see Figure D.1) allows you to perform the HP® system administration functions, including setting up users and groups, configuring the disks and file systems, performing auditing and security activities, and editing the system kernel configuration. This appendix provides instructions for:

- Using SAM to configure the 9900 disk devices (see section D.1), and
- Using SAM to set the maximum number of volume groups (see section D.2).

Note: The information in this appendix is provided for reference use only and is not intended to be complete. For further information on using SAM, please refer to the online help and/or user documentation for this product, or contact HP® technical support.

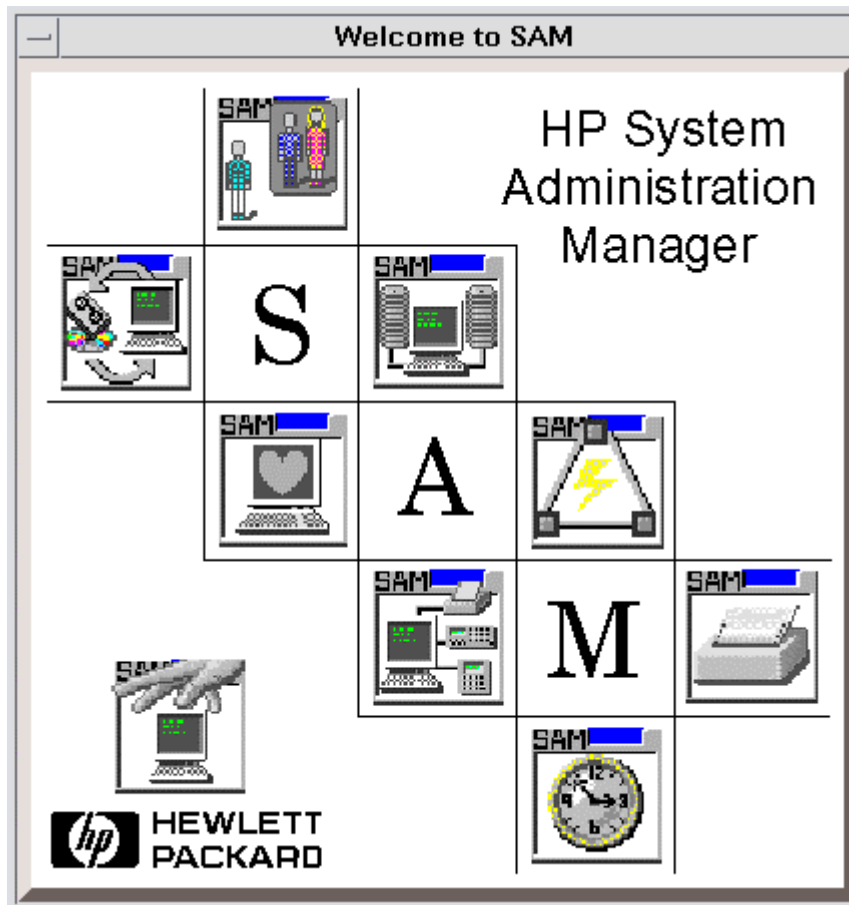


Figure D.1 Welcome to SAM

D.1 Configuring the 9900 Devices Using SAM

The SAM Areas window (see Figure D.2) displays the system administration functions and allows you to select the desired function. The **Disks and File Systems** function (selected in Figure D.2) allows you to configure the new disk devices for LVM operations.

Note: Do not configure the HMDE devices (e.g., OPEN-x-HMDEoto) for LVM operations.

To configure the new 9900 SCSI disk devices using SAM:

1. Select the **Disks and File Systems** function, and select **Disk Devices** (see Figure D.3).
2. Verify that the new 9900 devices are displayed on the Disk Devices panel.
3. Select the device to configure, select the **Actions** menu, select **Add**, and then select the **Using the Logical Volume Manager** option.
4. On the Add a Disk Using LVM panel, select **Create or Extend a Volume Group...**
5. On the Create a Volume Group panel, enter the name of the new or existing volume group to assign the new device to, and then select **OK** twice. The Add a Disk Using LVM panel now displays the volume group name.
6. Select **Add New Logical Volume...** to open the Create New Logical Volumes panel.
7. On the Create New Logical Volumes panel, enter the name, size (e.g., 2344 MB for OPEN-3), and mount directory for the new logical volume. Select **Add**, and then select **OK** twice.
8. Repeat steps (3) through (7) for each new disk device on the 9900 subsystem.

D.2 Setting the Maximum Number of Volume Groups Using SAM

The HP® system kernel specifies the maximum number of volume groups that can be created (default = 10). The **Kernel Configuration** function (next to **Disks and File Systems** in Figure D.2) allows you to change the maximum number of volume groups as needed, for example, to accommodate the new devices on the 9900 subsystem.

To change the maximum number of volume groups using SAM:

1. Select the **Kernel Configuration** function, and then select **Configurable Parameters**.
2. On the Configurable Parameters panel, select the **maxvgs** parameter, select the **Actions** menu, and then select **Modify Configurable Parameter...**
3. On the Modify Configurable Parameter panel, enter the desired maximum number of volume groups in the **Formula/Value** field, and then select **OK**.
4. On the Configurable Parameters panel, make sure that none of the parameters is selected, and then select the **Actions** menu, and select **Create New Kernel**.
5. When the confirmation panel opens, select **Yes** to create the new kernel (or select **No** to return to the Configurable Parameters panel).
6. When the Reboot the System panel opens, select **OK** to move the new kernel into place and reboot the system.

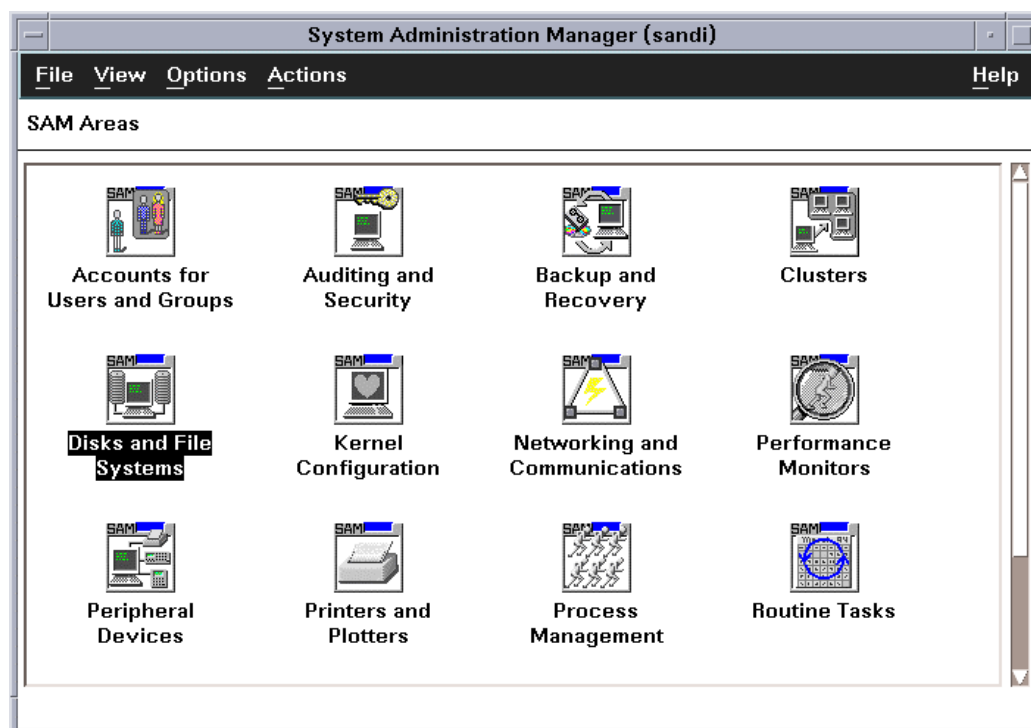


Figure D.2 SAM Areas

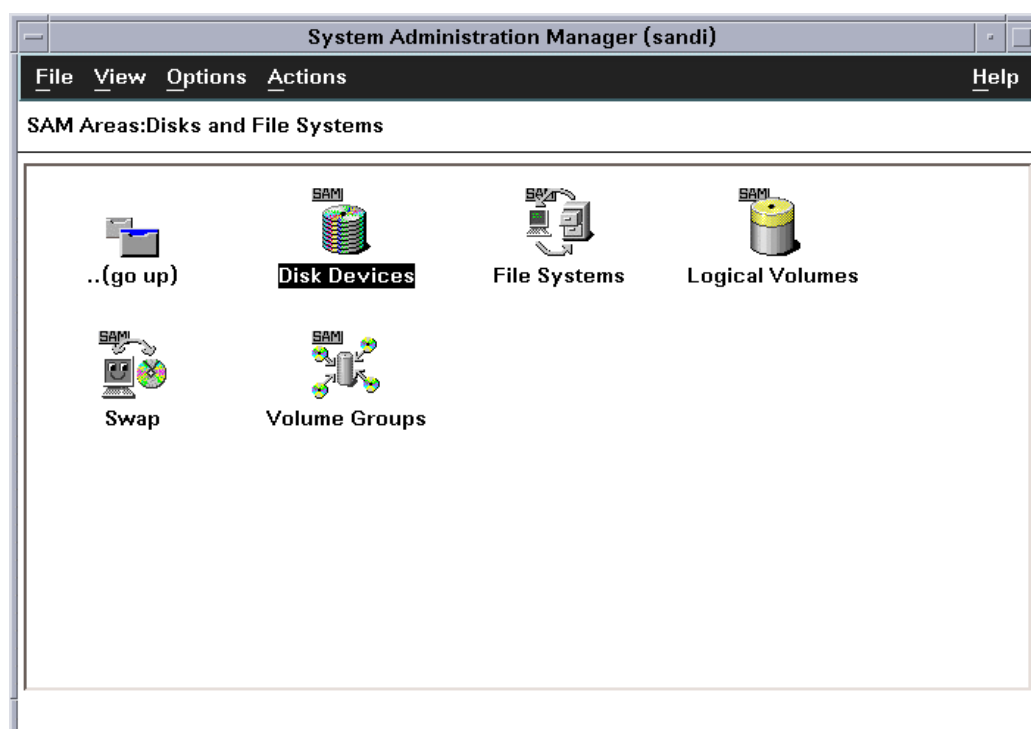


Figure D.3 SAM Areas: Disks and File Systems

Appendix E Online Device Installation

This appendix provides instructions for configuring new 9900 SCSI disk devices without rebooting the HP-UX[®] system. Please use the normal disruptive device configuration procedure in the following cases:

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

To perform online device installation and configuration:

1. Verify that the new devices on the 9900 subsystem are ready to be configured. The Hitachi Data Systems representative should have completed hardware installation and verified the normal status of the new devices (refer to section 2.2).
2. Make sure that you are logged in as **root**.
3. Enter the **insf -e** command to perform online device recognition. The **insf -e** command creates device files for the new devices on the existing fibre busses (refer to section 3.2.1).
4. Configure the new 9900 disk devices for HP-UX[®] operations as described in sections 3.2 through 3.8. For new HMDE devices, you only need to verify the device files and driver. Do not partition or create a file system on any HMDE device.
5. Remember to configure the application failover, path failover (i.e., **vgextend**), and/or SNMP software on the HP[®] system as needed to recognize the new disk devices.

