HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Lightning 9900™

Sun® Solaris® Configuration Guide

Document Revision Level

Revision	Date	Description
MK-90RD017-P	May 2000	Preliminary Release
MK-90RD017-0	July 2000	Initial Release
MK-90RD017-1	November 2000	Revision 1, supercedes and replaces MK-90RD017-0

Source Document Revision Level

The following source document was used to produce this Lightning 9900(TM) configuration guide: *DKC310 Disk Subsystem SCSI/Fibre Installation Manual, Solaris,* rev. 6.1.

Preface

The *Hitachi Freedom Storage™ Lightning 9900™ Sun® Solaris® Configuration Guide* describes and provides instructions for configuring the devices on the Hitachi Lightning 9900™ array subsystem for operation with the Sun® Solaris® 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 Sun® Solaris® operating system and the Sun® SPARCstation, SPARCserver, SPARCcenter, and/or Ultra Series systems, 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 Lightning 9900™ User and Reference Guide* (MK-90RD008) for further information on the 9900 disk array subsystems.

For further information on the Sun® Solaris® operating system, please consult the Sun® Solaris® online help and/or user documentation, or contact Sun® technical support.

Contents

List of Figures

List of Tables

Chapter 1 Overview of Sun® Solaris® Configuration

1.1 Sun® Solaris® Configuration

This document describes the requirements and procedures for connecting the 9900 subsystem to a Sun® Solaris® system and configuring the new 9900 devices for operation with the Solaris® operating system. The Hitachi Data Systems Customer Service representative performs the physical installation of the 9900 subsystem. The user prepares for 9900 subsystem installation and configures the new 9900 devices with assistance as needed from the Hitachi Data Systems representative.

Chapter 2 provides instructions for preparing for 9900 device configuration. Chapter 3 provides instructions for configuring the new 9900 devices. Chapter 4 contains general information for middleware (e.g., host failover) and SNMP operations. Chapter 5 contains troubleshooting information for 9900 Sun® Solaris® configuration. For further information on the Sun® Solaris® operating system, please refer to the Solaris® user documentation, or contact Sun® technical support.

Configuration of the 9900 disk devices for Solaris® operations includes:

- Performing fibre channel installation (see section 2.3),
- Changing the timeout and queue_depth parameters (see section 2.5),
- Setting and recognizing LUNs (see section 3.1),
- Partitioning and labeling the disk (see section 3.2), and
- Creating and mounting the file system (see section 3.4).

1.2 Hitachi Lightning 9900™ Array Subsystem

The Hitachi Lightning 9900™ subsystem can be configured with fibre-channel ports and/or serial interface ports (compatible with ESCON® protocol) to provide connectivity with UNIX/PC-server hosts as well as S/390® mainframe hosts.

For further information on the 9900 subsystem, please refer to the *Hitachi Lightning 9900™ User and Reference Guide* (MK-90RD008), or contact your Hitachi Data Systems account team.

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

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 Sun® Solaris® 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 Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

CVS Devices (OPEN-x CVS). The CVS devices are disk devices which are used exclusively by the Sun® Solaris® host system. The Custom Volume Size (CVS) feature of the 9900 subsystem (also called Virtual LUN) enables you to configure custom-size LUs which are smaller than standard OPEN-x LUs. This 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 Virtual LUN/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 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 Solaris® operations includes:

- Performing fibre channel installation (see section 2.3),
- Changing the timeout and queue_depth parameters (see section 2.5),
- Setting and recognizing LUNs (see section 3.1),
- Partitioning and labeling the disk (see section 3.2), and
- Creating and mounting the file system (see section 3.4).

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 are not SCSI disk devices. 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 (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 multiplatform devices are ***not*** write-protected for UNIX/PC server access. Do not execute any write operation by the fibre-channel adapters on these devices. Do not create a partition or file system on these devices. This will overwrite any data on the HMDE device and also prevent the HMDE software from accessing the device.

Configuration of the 9900 HMDE devices for operation with Sun® Solaris® includes:

- Configuring the host fibre-channel adapters (see section 2.5),
- Verifying new device recognition (see section 2.7).

Note: When the Sun® system accesses the multiplatform devices, the message "Request sense couldn't get sense data" may be displayed on the remote console. Please disregard this message.

Table 1.1 9900 Device Specifications

Device Type *(Note 1)*	Category *(Note 2)*	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Cylinders	# of Heads	# of Sectors per Track	Capacity MB *(Note 3)*
OPEN-3	SCSI disk	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
OPEN-9	SCSI disk	HITACHI	OPEN-9	14423040	512	10016	15	96	7042
OPEN-K	SCSI disk	HITACHI	OPEN-K	3661920	512	2543	15	96	1788
OPEN-8	SCSI disk	HITACHI	OPEN-8	14351040	512	9966	15	96	7007
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	2347*n
OPEN-9*n	SCSI disk	HITACHI	OPEN-9*n	14423040*n	512	10016*n	15	96	7042*n
OPEN-K*n	SCSI disk	HITACHI	OPEN-K*n	3661920*n	512	2543*n	15	96	1788*n
OPEN-8*n	SCSI disk	HITACHI	OPEN-8*n	14351040*n	512	9966*n	15	96	7007*n
OPEN-E*n	SCSI disk	HITACHI	OPEN-E*n	28452960*n	512	19759*n	15	96	*Note 6*
OPEN-3 CVS	SCSI disk	HITACHI	OPEN-3-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-9 CVS	SCSI disk	HITACHI	OPEN-9-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-K CVS	SCSI disk	HITACHI	OPEN-K-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-8 CVS	SCSI disk	HITACHI	OPEN-8-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-E CVS	SCSI disk	HITACHI	OPEN-E-CVS	*Note 5*	512	*Note 6*	15	96	*Note 6*

Table 1.1 9900 Device Specifications (continued)

Device Type *(Note 1)*	Category *(Note 2)*	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Cylinders	# of Heads	# of Sectors per Track	Capacity MB *(Note 3)*
OPEN-3*n CVS	SCSI disk	HITACHI	OPEN-3*n-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-9*n CVS	SCSI disk	HITACHI	OPEN-9*n-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-K*n CVS	SCSI disk	HITACHI	OPEN-K*n-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-8*n CVS	SCSI disk	HITACHI	OPEN-8*n-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
OPEN-E*n CVS	SCSI disk	HITACHI	OPEN-E*n-CVS	*Note 5*	512	*Note 6*	15	96	*Note 6*
3390-3A	HMDE otm/mto	HITACHI	3390-3A	5820300	512	3345	15	116	2844
3380-KA	HMDE otm/mto	HITACHI	3380-KA	3833280	512	2662	15	96	1873
3390-3B	HMDEmto	HITACHI	3390-3B	5816820	512	3343	15	116	2842
3380-KB	HMDEmto	HITACHI	3380-KB	3830400	512	2660	15	96	1871
3390-3C	HMDEotm	HITACHI	OP-C-3390-3C	5820300	512	3345	15	116	2844
3380-KC	HMDEotm	HITACHI	OP-C-3380-KC	3833280	512	2662	15	96	1873
HMDE OPEN-3	HMDEoto	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
3390-3A CVS	HMDE otm/mto	HITACHI	3390-3A-CVS	*Note 4*	512	*Note 5*	15	116	*Note 6*
3380-KA CVS	HMDE otm/mto	HITACHI	3380-KA-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
3390-3B CVS	HMDEmto	HITACHI	3390-3B-CVS	*Note 4*	512	*Note 5*	15	116	*Note 6*
3380-KB CVS	HMDEmto	HITACHI	3380-KB-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
3390-3C CVS	HMDEotm	HITACHI	OP-C-3390-3C-CVS	*Note 4*	512	*Note 5*	15	116	*Note 6*
3380-KC CVS	HMDEotm	HITACHI	OP-C-3380-KC-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*
HMDE OPEN-3 CVS	HMDEoto	HITACHI	OPEN-3-CVS	*Note 4*	512	*Note 5*	15	96	*Note 6*

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 partitions and file systems for Solaris® operations. The multiplatform devices (e.g., 3390-3A/B/C) must be installed as raw devices and can only be accessed using HMDE. Do not create a file system on any device used for HMDE operations.

Note: The 9900 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-3*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 device capacity can sometimes be changed by the BIOS or host adapter board. Also, different capacities may be due to variations such as 1 MB = 1000^2 or 1024^2 bytes.

Note 4: 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 5: 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

Note 6: The size of an OPEN-x CVS volume is specified by capacity in MB, not by number of cylinders. The user specifies the volume size using the Remote Console PC.

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

The requirements for 9900 Sun® Solaris® configuration are:

- Hitachi Lightning 9900™ subsystem, all-open or multiplatform configuration.
  - The 9900 Remote Console PC and LUN Manager software are used to configure the fibre-channel (FC) ports. If the remote LUN Manager feature is not installed, please contact your Hitachi Data Systems account team for information on LUN and fibre-channel configuration services.
  - If a TACHYON chip is attached to a Jaycor FibreStar board, the TACHYON chip must be version 3.0 or higher.

  Note: The availability of 9900 features and devices depends on the level of microcode installed on the 9900 subsystem.
- Sun® system: Sun® SPARCstation series, Sun® SPARCserver series, Sun® SPARCcenter series, or Sun® Ultra series.
- Sun® Solaris® OS, version 2.5, 2.6, 7.0, 8.0.

  Important: Please contact Sun® to make sure that the most current OS patches are installed on the Sun® system(s). If you are using Solaris® version 2.5.1, please contact Jaycor to make sure that you have the necessary Sun® patches for TACHYON chip compatibility. If you are using the Sun® X6729A HBA with Solaris® 7.0, please contact Sun® to make sure that you have the necessary patches (107292-06, 107458-09, 107474-01, 107834-03).

  Note: Hitachi Data Systems plans to support future releases of the Solaris® operating system. This document will be updated as needed to cover version-specific information. For further information on Solaris® version support, please contact your Hitachi Data Systems account team.
  - **Root** login access to the Sun® system.
- 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, optical cables, hubs, and fibre switches, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center (see section 5.2).

- Drivers. The following FC adapters and drivers have been verified for use with the 9900 subsystem. For information on the driver requirements for other FC adapters, please refer to the user documentation for the adapter or contact the vendor.

FC Adapter	OS Version	Drivers / Patches
Jaycor FC-1063 32bit Sbus	2.6 7	2.2.0 HIT04 FCODE:V11.0.9 2.2.1.10 FCODE:V13.3.5
Jaycor FC64-1063 64bit Sbus	6 7 8	2.2.0 HIT04 FCODE:V11.0.9 2.2.1 HIT06.01 FCODE:V13.3.5 2.5 FCODE:V13.3.7
Jaycor FCI-1064 PCI bus	6 7	2.2.0 HIT04 2.2.1 HIT06.01
Emulex Light Pulse 7000	2.6	4.00
SUN PCI: X6729A	2.6	Bundled as part of the SUN OS. The following patches are mandatory: 107280-05 (For PCI bus) 105356-12
SUN PCI: X6729A	7	The following patches are mandatory: 107292-06 107458-09 107474-01 107834-03

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 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.
   - 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 Solaris® 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 2.3.4). 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 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 the 9900 ports for you using the SVP of the subsystem. Please contact your Hitachi Data Systems account team for further information on LUN configuration services.

2.3 Preparing to Connect the 9900 Subsystem

Before the 9900 subsystem is connected to your Solaris® system, you must perform the following tasks:

- Set the host mode for the 9900 fibre-channel port(s) (see section 2.3.1), and
- Configure the 9900 fibre-channel ports (see section 2.3.2),
- Verify and configure the host fibre-channel adapters (see sections 2.4 and 2.5),
- Set the disk and device parameters (see section 2.6), and
- Connect the 9900 subsystem to the Solaris® system (see section 2.7)

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).

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 Solaris® operations is **09**.

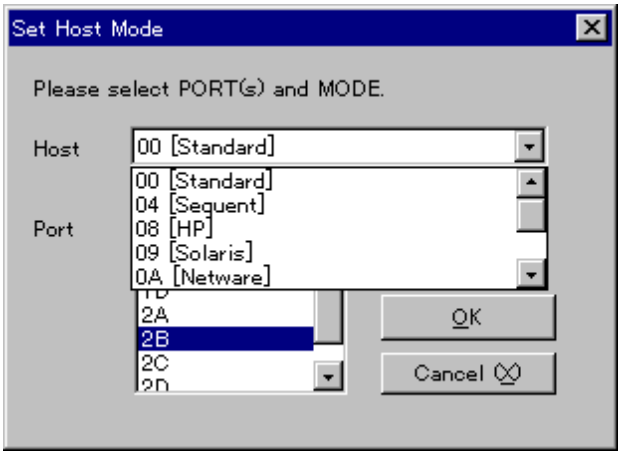

Figure 2.1 Setting the Host Mode

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.

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 TID. See Appendix C for a description of the AL-PA-to-TID translation.

Table 2.1 Fibre Topology Settings on the 9900 Remote Console PC

Fibre Parameter	Connection Parameter	Provides:
ON	FC-AL	FL-port (public arbitrated loop)
ON	Point-to-Point	F-port (fibre port)
OFF	FC-AL	AL-port (private arbitrated loop)
OFF	Point-to-Point	*Not supported*

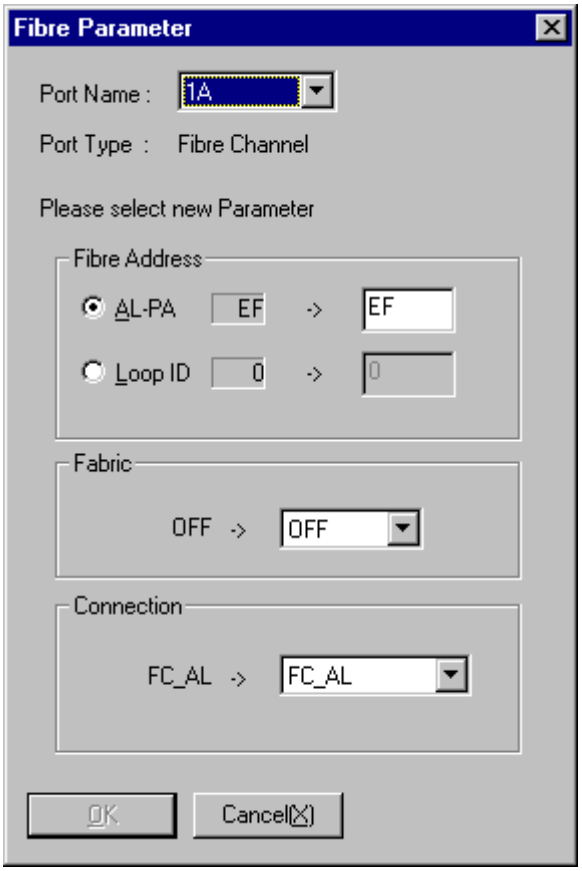

Figure 2.2 Set Fibre Topology Panel (from the LUN Manager Software)

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	

2.4 Verifying the Host Fibre-Channel Adapter Installation

Before the 9900 is connected to the Sun® system, you must verify the fibre host adapter installation. To ensure that the host configuration is correct, verify recognition of the Fibre Channel Adapter (FCA) and the FCA driver. Table 2.3 lists the driver requirements for the host adapters which have been verified with the 9900 subsystem. Please contact the Hitachi Data Systems Support Center for information on using other fibre-channel adapters with the 9900 subsystem.

Table 2.3 Driver Requirements for Fibre-Channel Adapters

Host Adapter Board	OS Ver.	Driver Version
Jaycor FC-1063 32-bit Sbus	2.6 7	2.2.0 HIT04 FCODE:V11.0.9 2.2.1.10 FCODE:V13.3.5
Jaycor FC64-1063 64-bit Sbus	2.6 7 8	2.2.0 HIT04 FCODE:V11.0.9 2.2.1.10 FCODE:V13.3.5 2.4.1
Jaycor FCI-1064 PCI bus	2.6 7 8	2.2.0 HIT04 FCODE:V11.0.9 2.2.1.10 FCODE:V13.3.5 2.4.1
Emulex Light Pulse 7000	2.6	4.00
Sun® PCI:X6729A	2.6 7.0	It is bundled as part of Solaris® operating system. The following paths are mandatory for 2.6: 107280-05 (for PCI bus) 105356-12 The following paths are mandatory for 7.0: 107292-06 107458-09 107474-01 107834-03

To verify the fibre-channel host configuration:

1. Log in to the Sun® system as **root**, and make sure that all existing devices are powered on and properly connected to the Sun® system.
2. Display the host configuration using the **dmesg** command (see Figures 2.2, 2.3, and 2.4). The fibre information (underlined in the following examples) includes the recognition of the **fibre channel adapter**, the **SCSI bus characteristics**, the **world wide name**, and the **FCA driver**. Make sure that the host recognizes these four classes. If this information is not displayed or if error messages are displayed, the host environment may not be configured properly.

```
# dmesg

Nov  9 23:14
ems, Inc.
mem = 65536K (0x4000000)
avail mem = 60129280
Ethernet address = 8:0:20:92:32:48
root nexus = Sun Ultra 1 SBus (UltraSPARC 167MHz)
sbus0 at root: UPA 0x1f 0x0 ...
espdma0 at sbus0: SBus0 slot 0xe offset 0x8400000
esp0:   esp-options=0x46
esp0 at espdma0: SBus0 slot 0xe offset 0x8800000 Onboard device sparc9 ipl 4
sd0 at esp0: target 0 lun 0
sd0 is /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@0,0
        <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
sd6 at esp0: target 6 lun 0
sd6 is /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@6,0
fca0: JNI Fibre Channel Adapter (1062 MB/sec), model FC                  <- Verify that
fca0:   SBus 1: IRQ 4: FCODE Version 11.0.9 [1a6384]: SCSI ID 125: AL PA 01  <- these items
fca0:   Fibre Channel WWN: 100000e0690000d5                              <- are listed.
fca0:   FCA Driver Version 2.2.HIT.03, Oct 09, 1999 Solaris 2.5, 2.6     <-
fca0:   All Rights Reserved.
fca0:   < Total IOPB space used: 1125824 bytes >
fca0:   < Total DMA space used: 565277 bytes >
root on /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@0,0:a fstype ufs
zs0 at sbus0: SBus0 slot 0xf offset 0x1100000 Onboard device sparc9 ipl 12
zs0 is /sbus@1f,0/zs@f,1100000
zs1 at sbus0: SBus0 slot 0xf offset 0x1000000 Onboard device sparc9 ipl 12
zs1 is /sbus@1f,0/zs@f,1000000
keyboard is </sbus@1f,0/zs@f,1000000> major <29> minor <2>
mouse is </sbus@1f,0/zs@f,1000000:b> major <29> minor <3>
stdin is </sbus@1f,0/zs@f,1000000> major <29> minor <2>
 . . . . . .
```

Figure 2.3 Displaying the Fibre Device Information (Jaycor FC-1063)

```
# dmesg

Nov  6 18:55
cpu0: SUNW,UltraSPARC-IIi (upaid 0 impl 0x12 ver 0x13 clock 333 MHz)
SunOS Release 5.6 Version Generic_105181-14 [UNIX(R) System V Release 4.0]
Copyright (c) 1983-1997, Sun Microsystems, Inc.
mem = 262144K (0x10000000)
avail mem = 256278528
Ethernet address = 8:0:20:a7:0:d9
root nexus = Sun Ultra 5/10 UPA/PCI (UltraSPARC-IIi 333MHz)
pci0 at root: UPA 0x1f 0x0
pci0 is /pci@1f,0
PCI-device: pci@1,1, simba #0
PCI-device: pci@1, simba #1
dad0 at pci1095,6460 target 0 lun 0
          :
          :
NOTICE:
Emulex LightPulse FC SCSI/IP 3.39m-COMBO                                 <- Verify that
                                                                         <- these items
NOTICE: Device Path for interface lpfc0:                                 <- are listed.
pci10df,f700 is /pci@1f,0/pci@1/fibre-channel@3                          <-
NOTICE: lpfc0: WWPN:10:00:00:00:c9:20:cb:b7 WWNN:10:00:00:00:c9:20:cb:b7 <-
dump on /dev/dsk/c0t0d0s1 size 131520K
```

Figure 2.4 Displaying the Fibre Device Information (Emulex LightPulse 7000)

```
# dmesg

Jan 24 09:48
ajor <37> minor <0>
mouse is </pci@1f,0/pci@1,1/ebus@1/su@14,3062f8> major <37> minor <1>
stdin is </pci@1f,0/pci@1,1/ebus@1/su@14,3083f8> major <37> minor <0>
SUNW,m64B0 is /pci@1f,0/pci@1,1/SUNW,m64B@2
m64#0: 1152x900, 2M mappable, rev 4750.7c
stdout is </pci@1f,0/pci@1,1/SUNW,m64B@2> major <8> minor <0>
se0 at ebus0: offset 14,400000
se0 is /pci@1f,0/pci@1,1/ebus@1/se@14,400000
SUNW,hme0: CheerIO 2.0 (Rev Id = c1) Found
SUNW,hme0 is /pci@1f,0/pci@1,1/network@1,1
SUNW,hme0: Using Internal Transceiver
SUNW,hme0: 100 Mbps half-duplex Link Up

. . . . . .
<HP-OPEN-3-5244 cyl 3336 alt 2 hd 15 sec 96>
fdthree0 at ebus0: offset 14,3023f0
fdthree0 is /pci@1f,0/pci@1,1/ebus@1/fdthree@14,3023f0
ecpp0 at ebus0: offset 14,3043bc
ecpp0 is /pci@1f,0/pci@1,1/ebus@1/ecpp@14,3043bc
power0 at ebus0: offset 14,724000
power0 is /pci@1f,0/pci@1,1/ebus@1/power@14,724000
pseudo-device: winlock0
winlock0 is /pseudo/winlock@0
atapicd2 at pci1095,6460 target 2 lun 0
atapicd2 is /pci@1f,0/pci@1,1/ide@3/atapicd@2,0
pseudo-device: lockstat0
lockstat0 is /pseudo/lockstat@0
pseudo-device: vol0
vol0 is /pseudo/vol@0
pseudo-device: llc10
llc10 is /pseudo/llc1@0
SUNW,CS42310 at ebus0: offset 14,200000
SUNW,CS42310 is /pci@1f,0/pci@1,1/ebus@1/SUNW,CS4231@14,200000
ifp0:   LIP reset occured; cause f8f7
ifp0:   LIP occured; cause f8f7
ifp0:   Loop up
ifp0:   Loop reconfigure in progress
ifp0:   Loop reconfigure done
pci1077,2100 is /pci@1f,0/pci@1/SUNW,ifp@3
ifp0:   Chip 2100 Rev 3; Firmware Version: 1.17.30
ssd12 at pci1077,2100: name w50000e10fff62200,0, bus address ef
ssd12 is /pci@1f,0/pci@1/SUNW,ifp@3/ssd@w50000e10fff62200,0
        <HP-OPEN-3-5245 cyl 3336 alt 2 hd 15 sec 96>
ssd13 at pci1077,2100: name w50000e10fff62200,1, bus address ef
ssd13 is /pci@1f,0/pci@1/SUNW,ifp@3/ssd@w50000e10fff62200,1
        <HP-OPEN-3-5245 cyl 3336 alt 2 hd 15 sec 96>
. . . . . .
```

Figure 2.5 Displaying the Fibre Device Information (Sun® Host Adapter)

2.5 Configuring the Host Fibre-Channel Adapter

After connecting the 9900 subsystem and rebooting the Sun® server, you are ready to configure the fibre-channel adapter(s) connected to the 9900. The host bus adapters have many configuration options. This section provides minimum requirements for configuring FC adapters for operation with the 9900 subsystem. The following sample instructions apply to the Jaycor, Emulex, and Sun adapters. For instructions on configuring other adapters, refer to the user documentation for the adapter.

2.5.1 Setting the FCA File

To configure the Jaycor FC-1063, FC64-1063 and FCI-1064, set the fca*.conf file (in /kernel/drv/ directory), and add the following descriptions in the configuration file:
timeout_reset_enable = 1
link_recovery_delay = 500
failover = 180;
def_port_binding = "$xxxxxx";
def_wwn_binding = " xxxxxxxxxxxxxxxx ";

When using VERITAS® Volume Manager (VxVM) Dynamic Multi Pathing (DMP), set the following values:
Recovery_attempts = 5
Failover = 60

When you use driver 2.5 with Solaris® version 8, change some parameters to following value.
ip_disable = 1;
def_wwpn_binding = "xxxxxxxxxxxxxxxx";
def_wwnn_binding = "xxxxxxxxxxxxxxxx";
def_port_binding = "$xxxxxx";
fca_verbose = 1;

Configure these parameters as follows:

- For 32-bit SBus adapters, configure the **/kernel/drv/fca.conf** file.
- For 64-bit SBus adapters, configure the **/kernel/drv/fca-pci.conf** file.

2.5.2 Editing the /kernel/drv/fca-pci.conf file

To edit the /kernel/drv/fca-pci.conf file:

1. Log in as **root**.
2. Make a backup of fca-pci.conf: **cp /kernel/drv/fca-pci.conf /kernel/drv/fca-pci.bk**
3. Edit the fca-pci.conf file: **vi /kernel/drv/fca-pci.conf**
4. Enter: **timeout_reset_enable = 1;** (see Figure 2.6).
5. Enter: **link_recovery_delay = 500;** (see Figure 2.7**)**
6. Enter: **recovery_attempts = 5;** (see Figure 2.8**)**

7. Enter: **failover = 180;** (see Figure 2.9)
8. Enter: **ip_disable = 1;** (This parameter is only for driver v2.5.) (see Figure 2.10)
9. Enter: **def_wwn_binding = "xxxxxxxxxxxxxxxx";** (before driver v2.5) (see Figure 2.11)
10. Or Enter: **def_wwpn_binding = "xxxxxxxxxxxxxxxx";** (for driver v2.5) (see Figure 2.12)
11. Enter: **def_wwnn_binding = "xxxxxxxxxxxxxxxx";** (for driver v2.5) (see Figure 2.13)
12. Enter: **`def_port_binding = "$xxxxxx";`**
13. Enter: **`fca_verbose = 1;`**
14. Save your changes and exit the text editor
15. Shut down and reboot to establish I/O TOV setting.

```
# Configuration flag timeout_reset_enable
# Type: boolean; default: 0
# When a command times out to a target, the driver can attempt to clear the problem in two
ways:
#   1. (timeout_reset_enable = 1) Reset the target.
#   2. (timeout_reset_enable = 0) Send an abort exchange (ABTS_LS) request to the target.
timeout_reset_enable = 1;
```

Figure 2.6 Example of set timeout_reset_enable = 1

```
# Configuration flag link_recovery_delay
# Type: unsigned int; default: 100
# Set delay between link up state and login recovery.
# This delay helps ensure link stability before recovery of communications to ports.
# The delay is in milliseconds with a 10 millisecond resolution.
link_recovery_delay = 500;
```

Figure 2.7 Example of set link_recovery_delay = 500

```
# Configuration variable recovery_attempts
# Type: integer; count; default: 5
# Defines the number of times login_recovery is attempted before failed
recovery_attempts = 5;
```

Figure 2.8 Example of set recovery_attempts = 5

```
# Configuration variable failover
# Type: integer; seconds; default: 180
# Defines the number of seconds after link failure before
# failing all pending commands on targets.
failover = 180;
```

Figure 2.9 Example of failover = 180

```
# Configuration flag ip_disable
# Type: boolean; defalt; 0 (false)
# if false (0), then the IP side of the driver is enabled
# if true (1), the the IP side of the driver is completely disabled
ip_disable = 1;
```

Figure 2.10 Example of ip_disable=1

```
# Configuration flag def_wwXn_binding where X is either n for node or p for port.
# Type: string; default: "$xxxxxxxxxxxxxxxx" (means WWXN is "static don't care")
# Sets the 16 digit hexidecimal default wwXn binding for every target/lun instance which does
not
# explicitly define one.
#       - A "$" preceding the string indicates static binding enabled
#       - A "x" in place of a digit indicates "don't care" for that digit
# *See technote for details on wwn bindings
def_wwpn_binding  = "xxxxxxxxxxxxxxxx";
def_wwnn_binding  = "xxxxxxxxxxxxxxxx";
```

Figure 2.11 Example of def_wwpn_binding and def_wwnn_binding

```
# Configuration flag def_port_binding
# Type: string; default: "xxxxxx" (means PORT is "non-static don't care")
# Sets the 6 digit hexidecimal default port binding for every target/lun instance which does
not
# explicitly define one.
#       - A "$" preceding the string indicates static binding enabled
#       - A "x" in place of a digit indicates "don't care" for that digit
# *See technote for details on port bindings
def_port_binding = "$xxxxxx";
```

Figure 2.12 Example of def_port_binding

```
# Configuration flag fca_verbose
# Type: boolean; default: 1
# Determines how many messages are displayed directly to the console.
#       - A "0" will quiet the driver messages to the console, but still print them to the
#         system message log.
#         (NOTE: a "boot -v" will override this setting and make the driver verbose again)
#       - A "1" will make the driver print all messages to the console and to the system
#         message log.
fca_verbose = 1;
```

Figure 2.13 Example of fca_verbose=1

2.5.3 Connecting to Fibre Switch

Perform the following steps to enable fibre switch connection:

1. Edit the /kernel/drv/fca.conf file as follows: **fca_nport = 1**.
2. The Jaycor port adapter should be connected to a Fibre switch F-Port. The DKC fibre channel port should connect to a fibre switch FL-Port.
3. Use the Zoning function when attempting to connect multiple types of servers to the 9900 via the same fibre switch.

2.5.4 Connecting to FC-AL

If there are multiple workstations/servers on same loop and two or more of the host's ports have the same AL_PA , change the AL_PA using the following procedure:

1. For FCI-1064 (see Figure 2.14):
   a) Backup the fca-pci.conf file: **cp /kernel/drv/fca-pci.conf /kernel/drv/fca-pci.bk**
   b) Edit the fca-pci.conf file
   c) Set **scsi-initiator-id = 0xZZ**; (0xZZ:0-7d hex)
   d) Save your changes and exit the text editor
   e) Shut down and reboot.

```
## fca-pci.conf - JNI FCA-PCI DRIVER (Solaris SCSI HBA) CONFIGURATION FILE

# Configuration variable scsi-initiator-id
# Type: integer, 0-125; default: none (must be explicitly set for PCI driver)
# Defines the adapters SCSI ID (and hence FC AL_PA) on the loop
scsi-initiator-id = 0x7d;
```

Figure 2.14 Example of /kernel/drv/fca-pci.conf

2. For FC-1063, FC64-1063:
   a) For Openboot mode, enter init 0.
   b) To disable autoboot, enter the following command at the **ok** prompt
   setenv auto-boot? False
   c) Set the **fcode-debug?** attribute to **true**. To view the OpenBoot environment variables, enter **printenv**. If the value of **fcode-debug**? is false, enter **setenv fcode-debug**? true
   d) Enter the following command to show a list of all SCSI devices: **probe-scsi-all**
   e) Select the device for which you want to change the SCSI-ID, and issue the following command (note there must be a space after the first "):
   " **/sbus@1f,0/fca@1,0**" **select-dev**
   f) If this fails, reset the system using the reset command, and issue the **select-dev** command again.
   g) To reconfigure the workstation to autoboot, enter: **setenv auto-boot? true**
   h) Reset the system by entering the **reset** command.

Note: There must be a storage device or loopback plug attached to the adapter. Enter the Openboot command to change the SCSI-ID, and enter **y** and the new SCSI ID (see Figure 2.16): **set-scsi-id**

```
Sure you want to change it? (y/n)  ...Enter `y'
Se ok set-scsi-id
SCSI Initiator ID set = 7D
t SCSI Initiator ID, 1 to 125(7d), enter 2 hex digits: xx  ...Enter new ID
SCSI ID set to xx
ok
```

Figure 2.15 Example of SCSI ID Change

2.5.5 Configuration for Emulex Lightpulse LP7000

1. Set the configuration in lpfc.conf file by setting the following variable in lpfc.conf: **avoid-resets=0**
2. Connecting to FC-AL: If there are multiple workstations/servers on the same Loop and two or more hosts' ports have the same AL_PA , change the AL_PA by changing the driver parameter in the lpfc.conf file.

Note: lpfcX-assign-alpa=AL_PA where X is the HBA number and AL_PA is a valid fibre channel named AL_PA.

3. How to edit /kernel/drv/lpfc.conf file:
   - make a backup of /kernel/drv/flpfc.conf: **cp /kernel/drv/lpfc.conf /kernel/drv/lpfc.bk**
   - edit /kernel/drv/lpfc.conf
   - set lpfcX-assign-alpa=AL_PA: **avoid-resets=0;** (see Figure 2.15.)

```
Solaris LightPulse lpfc (SCSI) / lpfn (IP) driver: global initialized data.
        :
        :
# Set a ALPA for interface, only valid if topology=4
# lpfc0-assign-alpa=2;  Request ALPA 2 for lpfc0
lpfc0-assign-alpa=0x3;
        :
        :
# Set avoid-resets to 1, to avoid having the target driver
# send BUS RESET commands down to the HBA driver.
# Default is 1.
avoid-resets=0;
```

Figure 2.16 Example of /kernel/drv/lpfc.conf

4. Save your changes and exit the text editor
5. Shutdown and reboot.

2.5.6 Configuration for Sun® X6729A Adapter

1. Modify the configuration in the ssd.conf file:
   - Edit /kernel/drv/ssd.conf file, and make a backup of /kernel/drv/flpfc.conf:
   - **cp -ip /kernel/drv/ssd.conf /kernel/drv/ssd.conf.standard**
2. Edit /kernel/drv/ssd.conf

Note: With X6729A, values higher than LUN=16 are not available.

3. Save changes and exit the text editor
4. Shutdown and reboot
5. Modify configuration in ses.conf file
   - Edit /kernel/drv/ses.conf file:
   - Make a backup of /kernel/drv/ses.conf: **cp -ip /kernel/drv/ses.conf /kernel/drv/ses.conf.standard**
   - Edit /kernel/drv/ses.conf: **vi /kernel/drv/ses.conf**
   - Please refer to Figure 2.14 and add the underlined part into your configuration file.

6. Save changes and exit the text editor
7. Shutdown and reboot.

```
# Copyright (c) 1995, by Sun Microsystems, Inc.
# All rights reserved.
#
#ident   "@(#)ssd.conf 1.5     97/02/20 SMI"

name="ssd" parent="SUNW,pln" port=0 target=0;
name="ssd" parent="SUNW,pln" port=0 target=1;
name="ssd" parent="SUNW,pln" port=0 target=2;
. . . . . .
name="ssd" parent="SUNW,pln" port=0 target=15;
name="ssd" parent="SUNW,pln" port=1 target=0;
name="ssd" parent="SUNW,pln" port=1 target=1;
name="ssd" parent="SUNW,pln" port=1 target=2;
. . . . . .
name="ssd" parent="SUNW,pln" port=5 target=14;
name="ssd" parent="SUNW,pln" port=5 target=15;
name="ssd" parent="sf" target=0;
name="ssd" parent="ifp" target=127;
name="ssd" parent="ifp" target=0 lun=0;
name="ssd" parent="ifp" target=0 lun=1;
name="ssd" parent="ifp" target=0 lun=2;
. . . . . .
name="ssd" parent="ifp" target=6 lun=10;
name="ssd" parent="ifp" target=6 lun=11;
```

Figure 2.17 Example of /kernel/drv/ssd.conf

```
#
# Copyright (c) 1996, by Sun Microsystems, Inc.
# All rights reserved.
#
#
#ident  "@(#)ses.conf 1.1     97/02/10 SMI"
#

name="ses" parent="sf"
      target=15;

name="ses" parent="SUNW,pln" port=0 target=15;
name="ses" parent="SUNW,pln" port=1 target=15;
name="ses" parent="SUNW,pln" port=2 target=15;
name="ses" parent="SUNW,pln" port=3 target=15;
name="ses" parent="SUNW,pln" port=4 target=15;
name="ses" parent="SUNW,pln" port=5 target=15;

name="ses" class="scsi"
      target=15 lun=0;
name="ses" parent="ifp" target=127;
name="ses" parent="ifp" target=0 lun=0;
name="ses" parent="ifp" target=0 lun=1;
name="ses" parent="ifp" target=0 lun=2;
name="ses" parent="ifp" target=0 lun=3;
name="ses" parent="ifp" target=0 lun=4;
name="ses" parent="ifp" target=0 lun=5;
. . . . . .
```

Figure 2.18 Example of /kernel/drv/ses.conf

2.6 Setting the Disk and Device Parameters

The queue depth parameter (**max_throttle**) for the 9900 devices must be set as specified in Table 2.4. The required I/O time-out value (TOV) for 9900 devices is 60 seconds (default TOV=60). If the I/O TOV has been changed from the default, you must change it back to 60 seconds by editing the **sd_io_time** parameter in the **/etc/system** file.

In addition to queue depth and I/O TOV, you should also verify all other required settings for your operational environment (e.g., FC fibre support). You must also make sure that the device parameters are the same for all 9900 devices. ***Note:*** For fibre-channel, the settings in the system file are effective for the whole system, not for just the adapter board.

Table 2.4 Max Throttle (Queue Depth) Requirements for the 9900 Devices

Parameter	Required Value
Queue depth per LU (sd_max_throttle)	≤ 32 per LU
Queue depth per port	≤ 256 per port

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 I/O TOV and queue depth for the 9900 devices:

1. Make a backup of /etc/system: **cp /etc/system /etc/system.old**
2. Edit /etc/system
3. Add the following to /etc/system: **set sd:sd_io_time=0x3c**
4. Add the following to /etc/system: **set sd:sd_max_throttle = x** (for **x** refer to Table 2.4)
5. Save your changes and exit the text editor
6. Shutdown and reboot to apply the above I/O TOV setting.

```
*ident  "@(#)system     1.15    92/11/14 SMI" /* SVR4 1.5 */
*
* SYSTEM SPECIFICATION FILE
*
:
*       To set a variable named 'debug' in the module named 'test_module'
*
*               set test_module:debug = 0x13
                set sd:sd_io_time = 0x3c                    ← Add this line to /etc/system
```

Figure 2.19 Setting the I/O TOV

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

set sd:sd_max_throttle = 8                                  ← Add this line to /etc/system
```

Figure 2.20 Setting the Max Throttle (Queue Depth)

2.7 Connecting the 9900 Subsystem to the Sun® System

The 9900 subsystem comes with all the hardware and cabling required for connection to the host system(s). Connection of the 9900 subsystem involves the following activities:

1. **Verify subsystem installation**. The Hitachi Data Systems representative verifies that the status of the SCSI/fibre adapters and LDEVs is NORMAL.
2. **Connect the 9900 to the Sun® system**. The Sun® system must be powered OFF before the 9900 is connected. The Hitachi Data Systems representative installs the fibre-channel cables between the 9900 and the Sun® system.
3. **Power on the Sun® system**. The user or Hitachi Data Systems representative can perform this activity. To power on the Sun® system after connecting the 9900:
   a) Power on the Sun® system display.
   b) Power on all peripheral devices. The 9900 should already be on, the fibre-channel ports should already be configured, and the driver configuration file and system configuration file should already be edited. If the SCSI paths are defined after the Sun® system is powered on, the system must be restarted in order to recognize the new devices.
   c) Confirm the ready status of all peripheral devices, including the 9900.
   d) Power on the Sun® system.
4. **Boot the Sun® system**. When the **OK** prompt appears, boot the system using the **boot -r** command. The **-r** option tells the system to rebuild the devices. Using **boot** by itself will not build the devices on the newly installed 9900.

Chapter 3 Configuring the 9900 Devices

After 9900 installation and connection are complete, the devices on the 9900 subsystem are ready to be configured for use. Configuration of the 9900 devices is performed by the user and requires root access to the Sun® system. The activities involved in device configuration are:

- Setting and recognizing LUNs (see section 3.1),
- Verifying new device recognition (see section 3.2)
- Partitioning and labeling the new devices (see section 3.3),
- Creating and mounting the file system (see section 3.4),
- Creating the file system (see section 3.4.1),
- Creating and verifying the mount directories (see section 3.4.2),
- Mounting and verifying the file systems (see section 3.4.3,
- Setting and verifying the auto-mount parameters (see section 3.4.4).

3.1 Setting and Recognizing LUNs

Once the 9900 is installed and connected, you must set and recognize the new LUs by adding the 9900 logical devices to the **sd.conf** file (/kernel/drv/sd.conf). The **sd.conf** file includes the SCSI TID and LUN for all LDEVs connected to the Sun® system. After editing the **sd.conf** file, you will halt the system and reboot.

To set and recognize LUNs:

1. Log in as root, and make a back up of the /kernel/drv/sd.conf file by entering the command: `cp -ip /kernel/drv/sd.conf /kernel/drv/sd.conf.standard`
2. Edit the file /kernel/drv/sd.conf
3. Edit the /kernel/drv/sd.conf file as shown in Figure 3.1. Make sure to make an entry (SCSI TID and LUN) for every new device being added to the Sun® system.
4. Exit the vi editor by entering the command: `ESC + :wq`
5. Shutdown the Sun® system by entering the command: `halt`
6. Reboot the Sun® system by entering the command: `boot -r`
7. Log in to the Sun® system as root, and verify that the system recognizes the 9900 by entering the command: `dmesg | more`. Figure 3.2 shows fibre device recognition.
8. Verify that the vendor name, product name, and number of blocks match the values listed in Figure 3.2.

```
# cp -ip /kernel/drv/sd.conf /kernel/drv/sd/conf/standard        ← Copy the /kernel/drv/sd.
:                                                                  conf file.
#
# vi /kernel/drv/sd.conf                                         ← Enter the file (vi shown).
#ident  "@(#)sd.conf    1.8     93/05/03 SMI"
name="sd" class="scsi"
        target=0 lun=0;

name="sd" class="scsi"
        target=1 lun=0;

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

name="sd" class="scsi"                                          ← Add this information for
        target=2 lun=1;                                         ← all assigned target IDs
                                                                   and LUNs. (*See note.)
name="sd" class="scsi"
        target=3 lun=0;

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

Program terminated
Type help for more information
OK

volume management starting.
The system is ready.

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

Note: The SCSI class type name is used because the SCSI driver is used for fibrechannel.

Figure 3.1 Setting and Recognizing LUNs

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

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

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

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

Figure 3.2 Fibre Device Recognition

3.2 Verifying Recognition of New Devices

After system start-up, log in as root and verify that the Sun® system recognizes the 9900 using the **dmesg | more** command. Verify that the displayed vendor names, product names, and number of blocks match the values displayed in Figure 3.3. If the results are different than the intended system configuration, SVP path definition or fibre cabling might be wrong.

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

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

Figure 3.3 Verifying New Devices

The example above indicates that there are two new disks on fca@1: target ID is 2, LUNs are 0 and 1, vendor name is “HITACHI”, product name is “OPEN-3” and the number of blocks is 4806720. LUN=0 and 1 are assigned as the device names sd192 and sd193, respectively. Other disks’ details:

vendor name “HITACHI”, product name “OPEN-9” and 14423040 512-byte blocks.
vendor name “HITACHI”, product name “3390-3B” and 5822040 512-byte blocks.
vendor name “HITACHI”, product name “3390-3A” and 5825520 512-byte blocks.
vendor name “HITACHI”, product name “3380-KB” and 3833280 512-byte blocks.
vendor name “HITACHI”, product name “3380-KA” and 3836160 512-byte blocks.

3.3 Partitioning and Labeling the New Devices

After the new devices have been recognized by the Sun® system, you can begin partitioning and labeling the devices. All new devices must be partitioned and labeled using the Sun® **format** utility (see ***WARNING*** below), including all SCSI disk devices and HMDE devices. Each SCSI disk device (e.g., OPEN-x) can have more than one partition. Each HMDE device (e.g., 3390-3A) must have one partition of fixed size. The disk partitioning and labeling procedure involves the following tasks: defining and setting the disk type, setting the partition(s), labeling the disk, and verifying the disk label.

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

WARNING: Be extremely careful when using the Sun® **format** utility. Do not use any **format** commands not described in this document. The Sun® **format** utility is designed for Sun® disks. Some **format** commands are not compatible with the 9900 and can overwrite the data on the disk. The 9900 will not respond to the **format** command (9900 disks are formatted using the SVP), and will not report any defect data in response to the **defect** command.

To partition and label the devices/disks on the newly installed 9900:

1. Enter **format** at the root prompt to start the **format** utility (see Figure 3.4).
   a) Verify that all new devices are displayed. If not, exit the **format** utility (**quit** or **Ctrl-d**), and then make sure that the SCSI/fibre-to-LDEV paths were defined for all devices and that all new devices were added to the driver configuration file). See Chapter 5 for troubleshooting information.
   b) Write down the character-type device file names (e.g., c1t2d0) for all of the new 9900 devices. You will need this information later to create the file systems.
2. When you are asked to specify the disk, enter the number (from the list) for the device to be partitioned and labeled. Remember the device type of this device (e.g., OPEN-3).
3. When you are asked if you want to label the disk, enter **n** for “no”.
4. After the format menu is displayed, enter **type** to display the disk types. The disk types are listed in Table 1.1 (vendor name + product name, e.g., HITACHI OPEN-3, HITACHI 3380-KB-CVS).
5. If the disk type for the selected device is already defined, enter the number for that disk type, and go to step (8). ***Note:*** Do not use HITACHI-OPEN-x-0315, HITACHI-3390-3A/B-0315, or HITACHI-3380-KA/B-0315. These disk types are created automatically by the Sun® system and cannot be used for the 9900 devices.
6. If the disk type for the selected device is not already defined, enter the number for **other** to define a new disk type.
7. Enter the disk type parameters for the selected device using the data provided above. Make sure to enter the parameters exactly as shown in Figure 3.5.
8. When you are asked if you want to label the disk, enter **n** for “no”.
9. After the format menu is displayed, enter **partition** to display the partition menu.

10. Enter the desired partition number, and then enter the partition parameters as shown in Figure 3.6. Refer to Tables 3.2-3.6 for partition sizes.
11. Enter **print** at the **partition>** prompt to display the current partition table.
12. Repeat steps (10) and (11) as needed to set the desired partitions for the selected device. ***Note:*** This step does not apply to the multiplatform devices (e.g., 3390-3A/B/C), because these devices can only have one partition of fixed size.
13. When you are finished setting the partitions for the selected device, enter **label** at the **partition>** prompt, and then enter **y** to label the device (see Figure 3.6).

    Note: The Sun® system displays the following warnings when an HMDE device (e.g., 3390-3A/B/C) is labeled. You can ignore these warnings.

    ```
    Warning: error warning VTOC.
    Warning: no backup labels.
    Label failed.
    ```
14. Enter **quit** to exit the **partition** utility and return to the **format** utility.
15. At the **format>** prompt, enter **disk** to display the available disks. Make sure that the disk you just labeled is displayed with the proper disk type name and parameters.
16. Repeat steps (2) through (15) for each new device to be partitioned and labeled. Once a 9900 disk type is defined (e.g., HITACHI OPEN-3), you can label all devices of that same type without having to enter the parameters (i.e., skipping steps (6) and (7)). For this reason, you may want to label the devices by type, for example, labeling all OPEN-3 devices, then all OPEN-9 devices, and so on until all new devices have been partitioned and labeled.
17. When you are finished partitioning and labeling the disks and verifying the disk labels, exit the **format** utility (**quit** or **Ctrl-d**).

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

c1t2d0: configured with capacity of 2.29GB   (OPEN-3)  ← These devices are not yet labeled.
c1t2d1: configured with capacity of 2.29GB   (OPEN-3)  ←
c2t4d0: configured with capacity of 6.88GB   (OPEN-9)  ←
c2t5d0: configured with capacity of 2.77GB   (3390-3B) ←
c2t6d0: configured with capacity of 2.78GB   (3390-3A) ←
c2t8d0: configured with capacity of 1.82GB   (3380-KB) ←
c2t9d0: configured with capacity of 1.83GB   (3380-KA) ←
  ⇑ These character-type device file names are used later to create the file systems.

AVAILABLE DISK SELECTIONS:

       0. c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>              ← Already labeled.
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0
       1. c0t3d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>              ←Already labeled.
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@3,0
       2. c1t2d0 <HITACHI-OPEN-3-52-34                              ← Not yet labeled:
                                  ↖ Product version
                 └─↖ LUN       ↖ Product ID
                         ↖Vendor ID
                   └─↖ Target Id
                 └─↖ Logical Controller ID
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@2,0    OPEN-3, TID=2, LUN=0
       3. c1t2d1 <HITACHI-OPEN-3-52-34                              ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@2,1    OPEN-3, TID=2, LUN=1
       4. c1t4d0 <HITACHI-OPEN-9-52-34                              ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@4,0    OPEN-9, TID=4, LUN=0
       5. c1t5d0 <HITACHI-3390-3B-52-34                             ←Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@5,0   3390-3B, TID=5, LUN=0
       6. c1t6d0 <HITACHI-3390-3A-52-34                             ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@6,0   3390-3A, TID=6, LUN=0
       7. c1t8d0 <HITACHI-3380-KB-52-34                             ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@8,0   3380-KB, TID=8, LUN=0
       8. c1t9d0 <HITACHI-3380-KA-52-34                             ← Not yet labeled:
          /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@9,0   3380-KA, TID=9, LUN=0
Specify disk (enter its number): 2                                  ← Select device.
selecting c1t2d0
[disk formatted]
Disk not labeled.  Label it now ? n                                 ← Enter "n" for no.
:
#
```

Figure 3.4 Verifying New Devices for Disk Partitioning

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

# format> type                                          ← Enter type.
:
AVAILABLE DRIVE TYPES
        0. Auto configure
             :
        14. SUN2.1G
        15. HITACHI-OPEN-3-0315                         ← Do not select this disk type.
        16. other
Specify disk type (enter its number):16                 ← Enter number for "other" to define.
Enter number of data cylinders:3336                     ← Enter value from Table 3.1 (*Note 2)
Enter number of alternate cylinders[2]:2                ← Enter value from Table 3.1
Enter number of physical cylinders[3338]:                 (press Enter for default)
Enter number of heads:15                                ← Enter value from Table 3.2-3.6
Enter number of physical heads[defaults]:                 (press Enter for default)
Enter number of data sectors/track:96                   ← Enter value from Table 3.2-3.6
Enter number of physical sectors/track[defaults]:         (press Enter for default)
Enter rpm of drive [3600]:10000                         ← Enter value from Table 3.1 (*Note 1)
Enter format time[defaults]:                              (press Enter for default)
Enter cylinder skew[defaults]:                            (press Enter for default)
Enter track skew[defaults]:                               (press Enter for default)
Enter track per zone[defaults]:                           (press Enter for default)
Enter alternate tracks[defaults]:                         (press Enter for default)
Enter alternate sectors[defaults]:                        (press Enter for default)
Enter cache control[defaults]:                            (press Enter for default)
Enter prefetch threshold[defaults]:                       (press Enter for default)
Enter minimum prefetch[defaults]:                         (press Enter for default)
Enter maximum prefetch[defaults]:                         (press Enter for default)
Enter disk type name(remember quotes):"HITACHI OPEN-3"  ← Enter name from Table 1.1
selecting c1t2d0
[disk formatted]
No defined partition tables.
Disk not labeled.  Label it now ? n                     ← Enter "n" for no.
format>
```

Figure 3.5 Defining and Setting the Disk Type

```
format> disk

AVAILABLE DISK SELECTIONS
        0. c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
           /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0
        1. c0t3d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
           /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@3,0
        2. c1t2d0 <HITACHI OPEN-3 cyl 3336 alt 2 hd 15 sec 96> ...already labeled
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,0
        3. c1t2d1 <HITACHI-OPEN-3-52-34 .....>                  ...not yet labeled
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,1
        4. c1t4d0 <HITACHI-OPEN-9-52-34 .....>                  ...not yet labeled
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@4,0
        5. c1t5d0 <HITACHI-3390-3B-52-34 ....>                  ...not yet labeled
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@5,0
        6. c1t6d0 <HITACHI-3390-3A-52-34 ....>                  ...not yet labeled
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@6,0
        7. c1t8d0 <HITACHI-3380-KB-52-34 ....>                  ...not yet labeled
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@8,0
        8. c1t9d0 <HITACHI-3380-KA-52-34 ....>                  ...not yet labeled
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@9,0

Specify disk (enter its number): 3


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

AVAILABLE DRIVE TYPES
        0.Auto configure
           :
        13. SUN1.3G
        14. SUN2.1G
        15. HITACHI-OPEN-3-52-34
        16. HITACHI OPEN-3
        17. other
Specify disk type (enter its number):16                       ← Enter the number for
:                                                                the desired drive type.
selecting c0t2d0
[disk formatted]
No defined partition tables.
Disk not labeled.  Label it now ? n                           ← Enter n for no.
format>
```

Figure 3.6 Setting the Partition(s)

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

PARTITION MENU
       0      - change '0' partition
       1      - change '1' partition
       2      - change '2' partition
       3      - change '3' partition
       4      - change '4' partition
       5      - change '5' partition
       6      - change '6' partition
       7      - change '7' partition
       select - select a predefined table
       modify - modify a predefined partition table
       name   - name the current table
       print  - display the current table
       label  - write partition map and label to the disk
       quit
partition> 0                                               ← Select partition number.
Part      Tag     Flag      Cylinders     Size       Blocks
  0  unassigned    wm        0 -           0          (0/0/0)

Enter partition id tag [root]:                             ← Press enter for default.
Enter partition permission flags [wm]:                     ← Press enter for default.
Enter new starting cyl [0]:                                ← Press enter for default.
Enter partition size [0b, 0c, 0.00mb]:3336c                ← Enter size (Table 3.1).
partition> print                                           ← Display partition table.
:                                                             (*See note.)
Current partition table (unnamed)
Part        Tag    Flag      Cylinders     Size              Blocks
  0        root    wm        0 -           0          (0/0/0)         0
  1        swap    wm        0 -           0          (0/0/0)         0
  2      backup    wu        0 - 3335      2.29 GB    (3336/0/0)   4803840
  3  unassigned    wu        0 -           0          (0/0/0)         0
  4  unassigned    wm        0 -           0          (0/0/0)         0
  5  unassigned    wm        0 -           0          (0/0/0)         0
  6         usr    wm      336 - 3335      204 GB     (2970/0/0)   4276800
  7  unassigned    wm        0 -           0          (0/0/0)         0
```

Note: In Solaris® version 2.6, make sure that partitions 1 and 6 are set to a partition size of 0. Solaris® version 2.5.1 will display partition 6 differently.

Figure 3.6 Setting the Partition(s) (continued)

```
PARTITION MENU
       0      - change '0' partition
       1      - change '1' partition
       2      - change '2' partition
       3      - change '3' partition
       4      - change '4' partition
       5      - change '5' partition
       6      - change '6' partition
       7      - change '7' partition
       select - select a predefined table
       modify - modify a predefined partition table
       name   - name the current table
       print  - display the current table
       label  - write partition map and label to the disk
       quit
partition> label                                                  ← Label the disk.
Ready to label disk, continue? Y                                  ← Enter Y for yes.
  * (see note at bottom)
partition> quit                                                   ← Return to format.
format> disk                                                      ← Display disks.

AVAILABLE DISK SELECTIONS
        0. c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
           /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0
        1. c0t3d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
           /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@3,0
        2. c1t2d0 <HITACHI OPEN-3 cyl 3336 alt 2 hd 15 sec 96>    ← Verify disk label.
                                                    ↖ Track size.
                                                ↖ Number of heads.
                                            ↖ Number of alternate cylinders.
                                        ↖ Number of data cylinders.
                                ↖ Disk type name.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,0
        3. c1t2d1 <HITACHI-OPEN-3-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@0,10000/sd@2,1
        4. c1t4d0 <HITACHI-OPEN-9-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@4,0
        5. c1t5d0 <HITACHI-3390-3B-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@5,0
        6. c1t6d0 <HITACHI-3390-3A-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@6,0
        7. c1t8d0 <HITACHI-3380-KB-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@8,0      (*Note 3)
        8. c1t9d0 <HITACHI-3380-KA-0315 ..................... >     ← Not yet labeled.
           /iommu@f,e0000000/sbus@f,e0001000/....,isp@1,10000/sd@9,0


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

Note: The Sun® Solaris® system displays the following warnings when an HMDE device (e.g., 3390-3A) is labeled. You can ignore these warnings: Warning: error warning VTOC. Warning: no backup labels. Label failed.

Figure 3.7 Labeling the Disk and Verifying the Disk Label

NOTES for Figures 3.4 through 3.7:

***Note1:** The 9900 reports the RPM of the physical disk drive in response to the type subcommand parameter.

***Note2:** The cylinder number of the 3390-3B is 3346, and the 9900 returns '3346 cylinder' to the Mode Sense command, and '5822040 blocks' (Maximum LBA 5822039) to Read capacity command. When 3390-3B is not labeled yet, Solaris® displays 3344 data cylinders and 2 alternate cylinders. When 3390-3B is labeled by the Solaris® format type subcommand, use 3340 for data cylinder and 2 for alternate cylinder. This is similar to the 3390-3B CVS.

***Note3:** The cylinder number of 3380-KB is 2662, and 9900 returns '2662 cylinder' to the Mode Sense command, and '3833280 blocks' (Maximum LBA 3833279) to the Read capacity command.

When 3380-KB is not labeled yet, Solaris® displays 2660 data cylinders and 2 alternate cylinders. When 3380-KB is labeled by the Solaris® format type subcommand, 2656 for data cylinder and 2 for alternate cylinder. This is similar to 3380-KB CVS.

Table 3.1 9900 Device Geometry Parameters

Device Type	# of Data Cylinders	# of Alternate Cylinders	RPM	Partition Size (sample)
OPEN-3	3336	2	10,000	3336c
OPEN-9	10014	2	10,000	10014c
OPEN-K	2541	2	10,000	2541c
OPEN-8	9964	2	10,000	9964c
OPEN-E	19757	2	10,000	19757c
OPEN-3*n	N1*	2	10,000	N4*
OPEN-9*n	N5*	2	10,000	N8*
OPEN-K*n	N9*	2	10,000	N12*
OPEN-8*n	N26*	2	10,000	N29*
OPEN-E*n	N30*	2	10,000	N33*
OPEN-x CVS	See Table 1.1	2	10,000	See Table 1.1
OPEN-3*n CVS	N22*	2	10,000	N25*
OPEN-9*n CVS	N22*	2	10,000	N25*
OPEN-K*n CVS	N22*	2	10,000	N25*
OPEN-8*n CVS	N22*	2	10,000	N25*
OPEN-E*n CVS	N22*	2	10,000	N25*
3390-3A	3346	2	10,000	3346c
3380-KA	2662	2	10,000	2662c
3390-3B	3340	2	10,000	3340c
3380-KB	2656	2	10,000	2656c
3390-3C	3346	2	10,000	3346c
3380-KC	2662	2	10,000	2662c
HMDE OPEN-3	3336	2	10,000	3336c
3390-3A CVS	See Table 1.1	2	10,000	See Table 1.1
3380-KA CVS	See Table 1.1	2	10,000	See Table 1.1
3390-3B CVS	See Table 1.1	2	10,000	See Table 1.1
3380-KB CVS	See Table 1.1	2	10,000	See Table 1.1
3390-3C CVS	See Table 1.1	2	10,000	See Table 1.1
3380-KC CVS	See Table 1.1	2	10,000	See Table 1.1
HMDE OPEN-3 CVS	See Table 1.1	2	10,000	See Table 1.1

* ***Note:*** For the values indicated by Nxx (e.g., N15, N22), please refer to Tables 3.2-3.6.

Table 3.2 Geometry Parameters for OPEN-3*n LUSE Devices

n	Data Cylinders-N1 Partition Size-N4	Heads-N2	Blocks/Track-N3	Usable Blocks (N1+2)*N2*N3	Provided Blocks =3338*15*96*n	Diff.
2	6674	15	96	9613440	9613440	0
3	10012	15	96	14420160	14420160	0
4	13350	15	96	19226880	19226880	0
5	16688	15	96	24033600	24033600	0
6	20026	15	96	28840320	28840320	0
7	23364	15	96	33647040	33647040	0
8	26702	15	96	38453760	38453760	0
9	30040	15	96	43260480	43260480	0
10	16688	30	96	48067200	48067200	0
11	20026	33	80	52873920	52873920	0
12	20026	30	96	57680640	57680640	0
13	20026	39	80	62487360	62487360	0
14	23364	30	96	67294080	67294080	0
15	16688	45	96	72100800	72100800	0
16	26702	30	96	76907520	76907520	0
17	30040	34	80	81714240	81714240	0
18	30040	30	96	86520960	86520960	0
19	30040	38	80	91327680	91327680	0
20	16688	60	96	96134400	96134400	0
21	23364	45	96	100941120	100941120	0
22	30040	55	64	105747840	105747840	0
23	30040	46	80	110554560	110554560	0
24	20026	60	96	115361280	115361280	0
25	16688	45	160	120168000	120168000	0
26	20026	39	160	124974720	124974720	0
27	30040	45	96	129781440	129781440	0
28	23364	60	96	134588160	134588160	0
29	30040	58	80	139394880	139394880	0
30	16688	45	192	144201600	144201600	0
31	30040	62	80	149008320	149008320	0
32	26702	60	96	153815040	153815040	0
33	30040	55	96	158621760	158621760	0
34	30040	64	85	163428480	163428480	0
35	30040	56	100	168235200	168235200	0
36	30040	60	96	173041920	173041920	0

N1,N2,N3: Use value on Table 1.1.

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

Table 3.3 Geometry Parameters for OPEN-9*n LUSE Devices

n	Data Cylinders-N5 Partition Size-N8	Heads-N6	Blocks/Track-N7	Usable Blocks (N5+2)*N6*N7	Provided Blocks =10016*15*96*n	Diff.
2	20030	15	96	28846080	28846080	0
3	30046	15	96	43269120	43269120	0
4	30046	20	96	57692160	57692160	0
5	30046	25	96	72115200	72115200	0
6	30046	30	96	86538240	86538240	0
7	30046	35	96	100961280	100961280	0
8	30046	40	96	115384320	115384320	0
9	30046	45	96	129807360	129807360	0
10	30046	50	96	144230400	144230400	0
11	30046	55	96	158653440	158653440	0
12	30046	60	96	173076480	173076480	0
13	30046	52	120	187499520	187499520	0
14	30046	56	120	201922560	201922560	0
15	30046	60	120	216345600	216345600	0
16	30046	64	120	230768640	230768640	0
17	30046	34	240	245191680	245191680	0
18	30046	36	240	259614720	259614720	0
19	30046	38	240	274037760	274037760	0
20	30046	40	240	288460800	288460800	0
21	30046	42	240	302883840	302883840	0
22	30046	44	240	317306880	317306880	0
23	30046	46	240	331729920	331729920	0
24	30046	48	240	346152960	346152960	0
25	30046	50	240	360576000	360576000	0
26	30046	52	240	374999040	374999040	0
27	30046	54	240	389422080	389422080	0
28	30046	56	240	403845120	403845120	0
29	30046	58	240	418268160	418268160	0
30	30046	60	240	432691200	432691200	0
31	30046	62	240	447114240	447114240	0
32	30046	64	240	461537280	461537280	0
33	30985	64	240	475960320	475960320	0
34	31924	64	240	490383360	490383360	0
35	31298	63	256	504806400	504806400	0
36	31689	64	256	519225344	519229440	4096

N5,N6,N7: Use value on Table 1.1, Table 1.2.

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

Table 3.4 Geometry Parameters for OPEN-K*n LUSE Devices

n	Data Cylinders-N9 Partition Size-N12	Head-N10	Blocks/Track-N11	Usable Blocks (N9+2)*N10*N11	Provided Blocks =2543*15*96*n	Diff.
2	5084	15	96	7323840	7323840	0
3	7627	15	96	10985760	10985760	0
4	10170	15	96	14647680	14647680	0
5	12713	15	96	18309600	18309600	0
6	15256	15	96	21971520	21971520	0
7	17799	15	96	25633440	25633440	0
8	20342	15	96	29295360	29295360	0
9	22885	15	96	32957280	32957280	0
10	25428	15	96	36619200	36619200	0
11	27971	15	96	40281120	40281120	0
12	30514	15	96	43943040	43943040	0
13	15256	39	80	47604960	47604960	0
14	17799	30	96	51266880	51266880	0
15	12713	45	96	54928800	54928800	0
16	20342	30	96	58590720	58590720	0
17	22885	34	80	62252640	62252640	0
18	22885	30	96	65914560	65914560	0
19	22885	38	80	69576480	69576480	0
20	22885	40	80	73238400	73238400	0
21	22885	42	80	76900320	76900320	0
22	22885	44	80	80562240	80562240	0
23	22885	46	80	84224160	84224160	0
24	22885	48	80	87886080	87886080	0
25	22885	50	80	91548000	91548000	0
26	22885	52	80	95209920	95209920	0
27	22885	54	80	98871840	98871840	0
28	22885	56	80	102533760	102533760	0
29	22885	58	80	106195680	106195680	0
30	22885	60	80	109857600	109857600	0
31	22885	62	80	113519520	113519520	0
32	22885	64	80	117181440	117181440	0
33	27971	45	96	120843360	120843360	0
34	22885	34	160	124505280	124505280	0
35	22885	35	160	128167200	128167200	0
36	22885	36	160	131829120	131829120	0

N9,N10,N11: Use value on Table 1.1, Table 1.2.

N12: Use same value as N9. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (e.g. enter 5084c for OPEN-K*2).

Table 3.5 Geometry Parameters for OPEN-8*n LUSE Devices

n	Data Cylinders-N26 Partition Size-N29	Heads-N27	Blocks/Track-N28	Usable Blocks (N26+2)*N27*N28	Provided Blocks =9966*15*96*n	Diff.
2	19930	15	96	28702080	28702080	0
3	29896	15	96	43053120	43053120	0
4	29896	20	96	57404160	57404160	0
5	29896	25	96	71755200	71755200	0
6	29896	30	96	86106240	86106240	0
7	29896	35	96	100457280	100457280	0
8	29896	40	96	114808320	114808320	0
9	29896	45	96	129159360	129159360	0
10	29896	50	96	143510400	143510400	0
11	29896	55	96	157861440	157861440	0
12	29896	60	96	172212480	172212480	0
13	29896	52	120	186563520	186563520	0
14	29896	56	120	200914560	200914560	0
15	29896	60	120	215265600	215265600	0
16	29896	64	120	229616640	229616640	0
17	29896	34	240	243967680	243967680	0
18	29896	36	240	258318720	258318720	0
19	29896	38	240	272669760	272669760	0
20	29896	40	240	287020800	287020800	0
21	29896	42	240	301371840	301371840	0
22	29896	44	240	315722880	315722880	0
23	29896	46	240	330073920	330073920	0
24	29896	48	240	344424960	344424960	0
25	29896	50	240	358776000	358776000	0
26	29896	52	240	373127040	373127040	0
27	29896	54	240	387478080	387478080	0
28	29896	56	240	401829120	401829120	0
29	29896	58	240	416180160	416180160	0
30	29896	60	240	430531200	430531200	0
31	29896	62	240	444882240	444882240	0
32	29896	64	240	459233280	459233280	0
33	32614	60	242	473584320	473584320	0
34	29896	64	255	487935360	487935360	0
35	30655	64	256	502284288	502286400	2112
36	31531	64	256	516636672	516637440	768

N26,N27,N28 : Use value on Table 1.1.

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

Note: Data cylinders must be less than or equal to **32767**, heads must be less than or equal to **4**, blocks per track must be less than or equal to **256** when these values are specified as parameters of Solaris® format type subcommand. The whole data blocks of OPEN-3*2 ~ OPEN-3*36 can be used by above parameters.

Table 3.6 Geometry Parameters for OPEN-E*n LUSE Devices

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

N30,N31,N32 : Use value on Table 1.1.

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

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

Table 3.7 Geometry Parameters for OPEN-x*n CVS-LUSE Devices (Example)

#	Data Cylinders-N22 Partition Size-N25	Heads-N23	Blocks/Track-N24	Usable Blocks (N22+2)*N23*N24	Provided Blocks-N21	Diff.
1	98	15	96	144000	35MB×2 volumes ↑35×1024/720↑×2=100 100×15×96=144000	0
2	2590	15	96	3732480	50MB×36 volumes ↑50×1024/720↑×36=2592 2592×15×96=3732480	0
3	284	15	96	411840	100MB×2 volumes ↑100×1024/720↑×2=286 286×15×96=411840	0
4	5694	15	96	8202240	500MB×8 volumes ↑500×1024/720↑×8=5696 5696×15×96=8202240	0
5	22758	30	96	65548800	2000MB×2 volumes ↑2000×1024/720↑×16=45520 45520×15×96=65548800	0
6	27455	40	188	206476640	2800MB×36 volumes ↑2800×1024/720↑×36=143388 143388×15×96=206478720	2080

N21 # of blocks of LUSE composed by CVS volumes are calculated by :
N21 = N20 × (# of heads) × (# of sectors per track).

N22: N20 – 2, Use total cylinder - 2.

N23,N24: Use value on Table 1.1, Table 1.2.

N25: Use same value as N22.

3.4 Creating and Mounting the File Systems

After you have partitioned and labeled all new disks, you can create and mount the file systems for the SCSI disk devices (e.g., OPEN-x). Creating and mounting the file systems for the new SCSI disk devices involves:

- Creating the file systems (see section 3.4.1),
- Creating and verifying the mount directories (see section 3.4.2),
- Mounting and verifying the file systems (see section 3.4.3),
- Setting the auto-mount parameters (optional) (see section 3.4.4).

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

3.4.1 Creating the File Systems

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

1. Create the file system using the **newfs -C <maxcontig>** command (see Figure 3.8).
   a) Use 6 or one of the following multiples of 6 as the **maxcontig** value for all 9900 SCSI disk devices: 12, 18, 24, or 30. If 6 is used, Solaris® will access 48 kB as a unit (6 × 8 kB), which matches the track size of the 9900 OPEN-3, OPEN-8, OPEN-K, and OPEN-9 devices. These **maxcontig** values (6, 12, 18, 24, 30) optimize the 9900's I/O performance by keeping the I/O data range on one track. The **maxcontig** value that you choose depends on your applications, and you can always change the **maxcontig** parameter to a different value at any time.
   b) Use the character-type device file as the argument (e.g., **/dev/rdsk/c1t2d0s0**).
2. When the confirmation appears, verify that the device file name is correct. If so, enter **y** for yes. If not, enter **n** for no, and then repeat step (1) using the correct device file name.
3. Repeat steps (1) and (2) for each new SCSI disk device on the 9900 subsystem. Make sure to use the same **maxcontig** value for all 9900 devices.

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

Figure 3.8 Creating the File Systems

3.4.2 Creating and Verifying the Mount Directories

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

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

1. Go to the root directory (Figure 3.9).
2. Create the mount directory using the **mkdir** command. Choose a name for the mount directory which identifies both the logical volume and the partition. For example, to create a mount directory named 9900_LU00, enter: `mkdir /9900_LU00`

   Note: If you need to delete a mount directory, use the **rmdir** command (e.g., `rmdir /9900_LU00`).
3. Verify the new mount directory using the **ls -x** command.
4. Repeat steps (2) and (3) for each logical partition on each new SCSI disk device.

```
# cd                                                    ← Go to the root directory.
# pwd                                                   ← Display current directory.
/
# mkdir /9900_LU00                                      ← Create new mount directory.
# ls -x                                                 ← Verify new mount directory.
9900_LU00  bin    dev      device    etc      export correctly
floppy     home   hstsboof kadb      kernel   lib
#
```

Figure 3.9 Creating and Verifying a Mount Directory

3.4.3 Mounting and Verifying the File Systems

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

To mount and verify the file systems for the newly installed 9900 devices (see Figure 3.10).

1. Mount the file system using the **mount** command. Make sure to use the correct block-type device file name and mount directory for the device/partition. For example, to mount the file /dev/dsk/c1t2d0s0 with the mount directory /9900_LU00 enter:
   `mount /dev/dsk/c1t2d0s0 /9900_LU00`

   a) If you need to unmount a file system, use the **umount** command (e.g., `umount /9900_LU00`).

   Note: If you have already set the auto-mount parameters (see section 3.4.4), you do not need to specify the block-type device file, only the mount directory.

2. Repeat step (1) for each partition of each newly installed SCSI disk device.
3. Display the mounted devices using the **df -k** command, and verify that all new SCSI disk devices are displayed correctly. OPEN-x devices will display as either OPEN-3, OPEN-8, OPEN-K, or OPEN-9 devices.
4. As a final verification, perform some basic UNIX® operations (e.g., file creation, copying, and deletion) on each logical unit to make sure that the new file systems are fully operational.

```
# mount /dev/dsk/c1t2d0s0 /9900_LU00                          ← Mount file system.
                ⇑ Block-type device file name
# mount /dev/dsk/c1t2d1s0 /9900_LU01                          ← Mount next file system.
                            ⇑ Mount directory name
# mount /dev/dsk/c1t2d2s0 /9900_LU02                          ← Mount next file system.
 :
 :
#
# df -k                                                         ← Display file
                                                                  systems.
File system         Kbytes     used    avail  capacity   Mounted on
/dev/dsk/c0t3d0s0    28775    27706        0    100%     /
/dev/dsk/c0t3d0s6   269191   234897     7384     97%     /usr
/proc                    0        0        0      0%     /proc
fd                       0        0        0      0%     /dev/fd
/dev/dsk/c0t3d0s4s   57567    29515    22302     57%     /var
swap                142204       20   142184      0%     /tmp
/dev/dsk/c0t3d0s7   462119   206000   209909     50%     /export/home
/dev/dsk/c0t3d0s5    47975    42059     1126     97%     /opt
/dev/dsk/c1t2d0s0  2256436        9  2030787      0%     /9900_LU00    ← Verify file systems.
/dev/dsk/c1t2d1s0  2256436        9  2030787      0%     /9900_LU01    ← OPEN-3 device
/dev/dsk/c1t2d2s0  6774358        9  6548709      0%     /9900_LU02    ← OPEN-9 device
:
# mount /dev/dsk/c1t2d0s0 /9900_LU00                          ← Mount file system.
# cd /9900_LU00                                               ← Go to mount directory.
# cp /bin/vi /9900_LU00/vi.back1                              ← Copy a file.
# ls -l                                                       ← Verify the file copy.
drwxr-xr-t   2 root     root       8192 Mar 15 11:35  lost+found
-rwxr-xr-x   1 root     sys     2617344 Mar 15 11:41  vi.back1
# cp vi.back1 vi.back2                                        ← Copy file again.
# ls -l                                                       ← Verify file copy again.
drwxr-xr-t   2 root     root       8192 Mar 15 11:35  lost+found
-rwxr-xr-x   1 root     sys     2617344 Mar 15 11:41  vi.back1
-rwxr-xr-t   1 root     sys     2617344 Mar 15 11:52  vi.back2
# rm vi.back1                                                 ←  Remove test files.
# rm vi.back2                                                 ←
```

Figure 3.10 Mounting and Verifying the File System

3.4.4 Setting and Verifying the Auto-Mount Parameters

You can add any or all of the new SCSI disk devices to the **/etc/vfstab** file to specify the auto-mount parameters for each device. Once a device is added to this file, you can mount the device without having to specify its block-type device file name (e.g., mount /9900_LU00), since the **/etc/vfstab** file associates the device with its mount directory.

To set the auto-mount parameters for the desired devices (see Figure 3.11).

1. Make a backup copy of the **/etc/vfstab** file by entering:
   `cp /etc/vfstab /etc/vfstab.standard`
2. Edit the **/etc/vfstab** file to add one line for each device to be auto-mounted. Table 3.8 shows the auto-mount parameters. If you make a mistake while editing, exit the **vi** editor without saving the file, and then begin editing again.
3. Reboot the Sun® system after you are finished editing the **/etc/vfstab** file.
4. Display the mounted devices using the **df -k** command, and verify that the desired devices were auto-mounted.

```
# cp -ip /etc/vfstab /etc/vfstab.standard                    ← Make backup before editing.
# vi /etc/vfstab                                             ← Start the UNIX vi editor.
#device            device              mount       FS      fsck  mount   mount
#to mount          to fsck             point       type    pass  at boot options
 ①                 ②                   ③           ④       ⑤     ⑥       ⑦   ← See Table 3.8.
/proc              -                   /proc       procfs  -     no      -
fd                 -                   /dev/fd     fd      -     no      -
swap               -                   /tmp        tmpfs   -     yes     -
/dev/dsk/c0t3d0s0  /dev/rdsk/c0t3d0s0  /           ufs     1     no      -
/dev/dsk/c0t3d0s6  /dev/rdsk/c0t3d0s6  /usr        ufs     2     no      -
/dev/dsk/c0t3d0s7  /dev/rdsk/c0t3d0s7  /export     ufs     3     yes     -
/dev/dsk/c0t3d0s5  /dev/rdsk/c0t3d0s5  /opt        ufs     4     yes     -
/dev/dsk/c0t3d0s1  -                   -           swapfs  -     no      -
/dev/dsk/c1t2d0s0  /dev/rdsk/c1t2d0s0  /9900_LU00  ufs     5     yes     -   ← Add one line
/dev/dsk/c1t2d1s0  /dev/rdsk/c1t2d1s0  /9900_LU01  ufs     5     yes     -      for each LUN.
```

Figure 3.11 Setting the Auto-Mount Parameters

Table 3.8 Auto-Mount Parameters

Parameter #	Name	Enter:
①	Device to mount	Block-type device file name
②	Device to fsck	Character-type device file name
③	Mount point	Mount directory name
④	FS type	File system type (e.g., ufs)
⑤	Fsck pass	Order for performing file system checks
⑥	Mount at boot	Yes = auto-mounted at boot/mountall No = not auto-mounted at boot/mountall
⑦	Mount options	Desired mount options: - no options (typical) -ro read-only access (e.g., for 3390-3B and 3380-KB devices)

Chapter 4 Middleware and SNMP Configuration

The 9900 subsystem supports many industry-standard middleware products which provide host failover, I/O path failover, and logical volume management functions. For SUN® Solaris®, the 9900 supports the following middleware products and functions:

- First Watch ® for failover (see section 4.1)

The 9900 disk 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 SVP and the SNMP manager on the host. The SNMP agent on the SVP sends status information to the host(s) when requested by the host or when a significant event occurs.

4.1 Host Failover

The 9900 subsystem supports the VERITAS® First Watch® host failover products for SUN® Solaris®.

The user must make sure to configure the VERITAS® First Watch® software and any other middleware products (e.g., Tuxedo) as needed to recognize and operate with the newly attached 9900 devices. For assistance with VERITAS® First Watch® operations, please refer to the VERITAS® First Watch® user documentation, or contact VERITAS® technical support. For assistance with specific configuration issues related to the 9900 subsystem, please contact the Hitachi Data Systems Technical Support Center (see section 5.4).

4.2 Path Failover

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 SNMP Manager on 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. Notification of 9900 error conditions is made in real time, providing the UNIX/PC server user with the same level of monitoring and support available to the mainframe user. The SIM reporting via SNMP enables the user to monitor the 9900 subsystem from the UNIX/PC server host without having to check the Remote Console PC for 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 SUN® Solaris® server host. For assistance with SNMP manager configuration on the SUN® Solaris® server host, please refer to the user documentation, or contact the vendor's technical support.

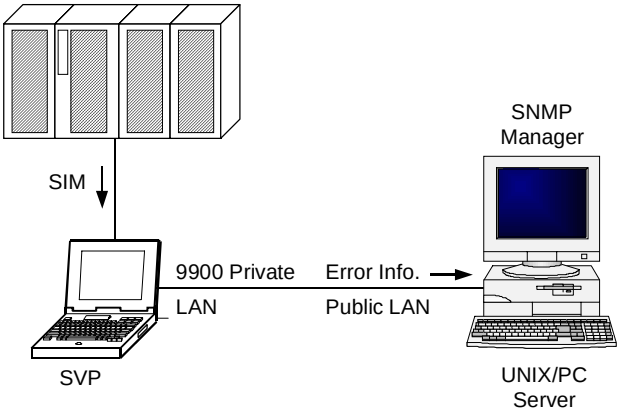

Figure 4.1 9900 SNMP Environment

Chapter 5 Troubleshooting

5.1 Troubleshooting

The Hitachi Lightning 9900™ array subsystems provide continuous data availability and is not expected to fail in any way. For troubleshooting information on the 9900 subsystem, please refer to the *Hitachi Freedom 9900 User and Reference Guide* (MK-90RD008).

Table 5.1 lists potential error conditions during 9900 Sun® Solaris® configuration and provides instructions for resolving each condition. If you are unable to resolve an error condition, please ask your Hitachi Data Systems Customer Service representative for help, or call the Hitachi Data Systems Support Center for assistance.

Table 5.1 Troubleshooting

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Make sure that the READY indicator lights on the 9900 subsystem are ON. Make sure that the SCSI/fibre cables are correctly installed and firmly connected. Make sure that the SCSI terminators are properly connected. Make sure that the SCSI target IDs are properly configured. The LUNs for each SCSI TID must start at 0 and continue sequentially without skipping any numbers. Make sure that the SCSI TIDs on each SCSI bus are unique. Do not connect two devices with the same SCSI TID on the same SCSI bus. Run **dmesg** to recheck the SCSI/fibre buses for new devices. Make sure that LUSE devices are not intermixed with normal LUNs or with multiplatform devices on the same SCSI port. Make sure that the maximum number of LUSE devices per SCSI port is 16.
The file system is not mounted after rebooting.	Make sure that the system was restarted properly. Make sure that the file system attributes are correct.
Physical volumes cannot be created.	Make sure that the 9900 logical devices are correctly formatted. Make sure that the character-type device file exists, and is correctly named.
Volume groups cannot be created.	Make sure that the directory for a volume group, the control file and block-type device files exist. Make sure that the block-type device file name is set correctly. Make sure that an already allocated physical volume has not been selected in the **vgcreate** command.
Logical volumes cannot be created.	Make sure that the volume capacity for OPEN-3 volumes is not greater than 2344 MB, and OPEN-9 volumes is not greater than 7040 MB. Make sure that the capacity of the volume group is not less than the total capacity of the partitioned logical volume.

Table 5.1 Troubleshooting (continued from previous page)

Error Condition	Recommended Action
A file system could not be created.	Make sure that the logical volume name is a character-type volume.
A file system is not mounted after rebooting.	Make sure that the system is powered on correctly. Make sure that /etc/fstab is correctly edited.
The Sun® system does not reboot properly after hard shutdown.	If the Sun® system is powered off without executing the shutdown process, wait three minutes before restarting the Sun® 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 Sun® system is restarted too soon, the 9900 will continue trying to process the queued commands, and the Sun® system will not reboot successfully.
9900 performs a self reboot because the subsystem was busy or it logged a panic message.	Reboot the Sun® system.
The 9900 subsystem responds Not Ready or the 9900 has displayed Not Ready and timed itself out.	Please contact the Hitachi Data Systems support center.
The Sun® system detects a parody error.	Check the SCSI adapter card and make sure that it has bee installed properly. Reboot the Sun® system.
System hangs, or devices are declared and then the system hangs.	Make sure that the target IDs are set 0 through 6 and 8 through 15 and target ID 7 has been reserved for the SCSI controller card.

5.2 Verbose Mode

One method of troubleshooting involves the "verbose" mode for the HBA configuration file. The following paragraphs are examples of error messages that may occur. A possible debugging method is to select the device and turn on verbose mode. Then attempt the boot process again. Verbose error messages will provide information which will help isolate the problem. Use the steps in Figure 5.1 to turn on the verbose flag. Figure 5.2 contains examples of Error Messages.

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

Figure 5.1 Turning on Verbose Flag

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

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

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

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

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

```
Error message:
Can't read disk label
Problem:
The selected target is not a Solaris filesystem.
```

```
Error message:
Nport init failed
Problem:
Card is connected to an arbitrated loop device, but wants to initialize as an NPORT. The
bootn0-wwn property has probably been set to a valid WWN.
```

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

```
Error message:
Can't open /packages
Problem:
The bootblk was not properly installed
```

Figure 5.2 Examples of Error Messages

5.3 Connection Problems

If problems occur when attempting connection of the 9900 subsystem and the Sun® Solaris® workstation, please check the items below during troubleshooting.

- 9900 Logical Devices cannot be recognized
  - Check that 9900 READY lamp is ON.
  - Confirm that target IDs are correctly set.
  - Confirm that SCSI ID is not duplicated with another devices.
  - Verify the contents of **/kernel/drv/sd.conf** file.
- File system cannot be created (**newfs** command)
  - Check that character-type is specified for device file of **newfs** command.
  - Verify that logical unit is correctly labeled by UNIX® **format** command.
- File system is not mounted after rebooting
  - Verify that system is powered on by correct procedure described.
  - Verify that **/etc/vfstab** is correctly edited.
- Reset ECC/LRC pinned track error status using the Solaris® **analyze** command as described in Appendix B.

5.4 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 host system (s). Please check the SIMs logged on the remote console PC, and note the reference codes and severity levels of the recent SIMs.

The worldwide Hitachi Data Systems Support Centers are:

- Hitachi Data Systems North America/Latin America
  San Diego, California, USA
  1-800-348-4357
  1-619-537-3000
- 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

DMP	Dynamic Multi Pathing
ESCON	Enterprise System Connection (IBM trademark for optical channels)
FCP	fibre channel protocol
HDS	Hitachi Data Systems Corporation
HMDE	Hitachi Multiplatform Data Exchange
I/O	input/output
IBM	International Business Machines Corporation
LDEV	logical device
LU	logical unit
LUN	logical unit number
LUSE	LU size expansion
LVM	Logical Volume Manager
PC	personal computer system
PCI	peripheral component interconnect
r/w	read/write
RISC	reduced-instruction-set computer
SCSI	small computer system interface
SIM	service information message
SNMP	simple network management protocol
SVP	service processor
TID	target ID
VxVM	VERITAS Volume Manager

Appendix B Pinned Track Recovery

To recover a pinned track on a 9900 device, use the **analyze** command in the format utility to write dummy data to the pinned track as shown in Figure B-1. This procedure is used for Sun® Solaris® 2.5.1 and 2.6 and may not be valid for other versions of the Solaris® OS. Make sure that you check the device number in the SCSI path configuration and not the 9900 LDEV ID.

WARNING: Data written on a pinned track or data on the partition containing the pinned track will be overwritten by the dummy data write process. The pinned track data or partition data should be restored using a backup file.

To recover a pinned track:

1. Log in to the Solaris® system as root.
2. Write dummy data to the pinned track using the **analyze** command (**analyze** is a subcommand of the **format** command).
3. Isolate LDEV number of the LDEV containing the pinned track by SVP.
4. Obtain the SCSI port number (1A~2H) and SCSI target ID and LUN that constructing the SCSI path from Solaris® to the LDEV.
5. Log in to the Solaris® system as root, and execute the **format** command.
6. Determine device number cxtxdx by port number and target ID and LUN and the board installed into the Sun® system to which the 9900 port is connected.
7. The whole track range must be specified by starting LBA and ending LBA.
8. One track has 96 blocks for OPEN-x, OPEN-x*n (LUSE), and OPEN-x CVS.
9. The pinned track will not be correctly reset if the whole track range is not specified.
10. Data written on the pinned track must be recovered by backup file because the **analyze** command writes dummy data on the pinned track.
11. The files or directories written with dummy data cannot be determined because of the structure of the UNIX® file system. All files and directories on the partition containing the pinned track must be recovered from the backup file.

```
host console login: root                                    ← Log in as root.
password: hitachi                                           ← Enter your password.
# format                                                    ← Enter format.
Searching for disks...done
AVAILABLE DISK SELECTIONS:
  :
Specify disk (enter its number):1                           ← Enter the pinned track
selecting  cxtxdx                                             disk number.
[disk formatted]

FORMAT MENU:
  :
format> analyze                                             ← Enter analyze.
  :
analyze> setup                                              ← Enter setup.
Analyze entire disk[yes]? N                                 ← Enter N for no.
Enter starting block number[0, 0/0/0]: starting LBA of the pinned track
Enter ending block number[5806479, 3336/14/115]: ending LBA of the track
Loop continuously[no] no                                    ← Enter no.
Enter number of passes[2]                                   ← Press enter for default.
Repair defective blocks[yes]                                ← Press enter for default.
Stop after first error[no]                                  ← Press enter for default.
Use random bit patterns[no]                                 ← Press enter for default.
Enter number of blocks per transfer[126, 0/1/10]: 1         ← Enter 1.
Verify media after formatting [yes]? no                     ← Enter no.
Enable extended messages[no]? yes                           ← Enter yes.
estore defect list[yes]?                                    ← Press enter for default.
Restore disk label[yes]?                                    ← Press enter for default.

analyze> write                                              ← Enter write.
Ready to analyze (will corrupt data). This takes a long time,
but is interruptable with CTRL-C. Continue? y               ← Enter y for yes.

      PASS 0 - pattern = c6dec6de
    cylinder number/head number/block counts

Total of 0 defective blocks repaired.
analyze> quit                                               ← Exit the analyze utility.
```

Figure B.1 Recovering a Pinned Track

Appendix C Fibre Port Addressing

In a FC-AL (fibre channel arbitrated loop) configuration, the host communicates with the devices comprising the loop with 8 bit AL-PA (arbitrated loop physical address, also called the port address). The number of available port addresses is 126. (There are 127 port addresses but address 00H is reserved for fibre connection.)

Devices communicate with hosts using individual port addresses. However, hosts map SCSI protocol to fibre channel devices. The hosts access the device's logical units using the device files /dev/dsk/c*t*d* and /dev/rdsk/c*t*d*. SCSI and fibre channel devices are accessed the same way however, the device files for SCSI and fibre channel devices are configured differently. Table C.1 lists the AL-PA and the corresponding SCSI address.

Table C.1 Fibre Port Addressing (continues on next page)

AL-PA	T value	AL-PA	T value	AL-PA	T value	AL-PA	T value
EF	0	CD	16	B2	32	98	48
E8	1	CC	17	B1	33	97	49
E4	2	CB	18	AE	34	90	50
E2	3	CA	19	AD	35	8F	51
E1	4	C9	20	AC	36	88	52
E0	5	C7	21	AB	37	84	53
DC	6	C6	22	AA	38	82	54
DA	7	C5	23	A9	39	81	55
D9	8	C3	24	A7	40	80	56
D6	9	BC	25	A6	41	7C	57
D5	10	BA	26	A5	42	7A	58
D4	11	B9	27	A3	43	79	59
D3	12	B6	28	9F	44	76	60
D2	13	B5	29	9E	45	75	61
D1	14	B4	30	9D	46	74	62
CE	15	B3	31	9B	47	73	63

Table C.1 Fibre Port Addressing (continued)

AL-PA	T value	AL-PA	T value	AL-PA	T value	AL-PA	T value
72	64	55	80	3A	96	23	112
71	65	54	81	39	97	23	113
6E	66	53	82	36	98	1F	114
6D	67	52	83	35	99	1E	115
6C	68	51	84	34	100	1D	116
6B	69	4E	85	33	101	1B	117
6A	70	4D	86	32	102	18	118
69	71	4C	87	31	103	17	119
67	72	4B	88	2E	104	10	120
66	73	4A	89	2D	105	0F	121
65	74	49	90	2C	106	08	122
63	75	47	91	2B	107	04	123
5C	76	46	92	2A	108	02	124
5A	77	45	93	29	109	01	125
59	78	43	94	27	110		
56	79	3C	95	26	111		

Appendix D Online Installation

This appendix provides instructions for online installation of new devices. After initial installation and configuration of the 9900 subsystem, additional devices can be installed or de-installed online without having to restart the Solaris® system. After online installation, the device parameters for new volumes must be changed to match the LUs defined under the same fibre-channel port (see section 3.2). This procedure should be performed by the system administrator (i.e., super-user).

Note: For additional instructions regarding online installation and deinstallation of LUs, please refer to the *DKC410 Maintenance Manual* (FE-90RD002) or the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

When the TOV for 9900 volumes needs to be changed, you must edit the **/etc/system** file and reboot to make the file change effective. Please use normal disruptive installation procedure for this case. Solaris® must be booted once by the **boot -r** command with the TIDs and LUNs for new volumes defined in **/kernel/drv/sd.conf** file for non-disruptive volume installation. It is one way to prepare future online volume installation that whole TIDs and LUNs, i.e. TID=0~15 and LUN=0~7 are defined in the Solaris® **kernel/drv/sd.conf** file and the system is once booted with **boot -r** command. In case of fibre channel, you should define LUN=0~119 in /kernel/drv/sd.conf file for the path (TID) which you will install additional LUNs.

Note: With the Sun® X6729A FC adapter, online installation is not available.

1. The Solaris® system must be booted once by **boot -r** command with the TIDs and LUNs for new volumes defined in **/kernel/drv/sd.conf** file for non-disruptive volume installation. If not, the Solaris® system must be rebooted to recognize new volumes.
2. Check existing SCSI TIDs.
3. Hitachi Lightning 9900™ setup:
   a) Ensure the latest microcode is loaded. Non-disruptive version up requires alternate path.
   b) Install CHF and LDEV, and connect fibre cable if necessary.
   c) Execute online LU installation from the SVP or Remote Console PC.
   d) Verify 9900 SCSI path configuration.
4. Execute the following commands (as super-user) to recognize the new devices.
   ```
   # /usr/sbin/drvconfig
   # /usr/sbin/devlinks
   # /usr/sbin/disks
   # /usr/ucb/ucblinks
   ```
5. After online installation, the device parameters (e.g., TOV, queue depth) for the new volumes must be changed to match the LUs defined under the same fibre-channel port.
6. Configure the new devices (partition, label, file system, auto-mount) as described in Chapter 3.

Appendix E Setting up the 9900 as Solaris ® 2.6 Boot Disk

E.1 Introduction

The 9900 can be set up as the boot disk of the Solaris® 2.6 OS instead of the built-in SCSI disk, if the following conditions are satisfied:

- HBA (host bus adapter) must be JNI 32-bit S-Bus FC Adapter.
- FCODE (firmware) version of HBA must be 11.0.11 or later.
- FCA driver version must be 2.2.0.HIT.04 or later.
- Fibre-channel topology must be FC-AL. The EPL mode of FC switching hub is not available.
- Open Boot version must be 2.15 or later.

Note: This functionality is currently supported for Solaris® 2.6 only. Please contact your Hitachi Data Systems representative for the latest information on platform/version support.

Install the JNI adapter card and the fibre channel disk drive on a Sparc workstation running the desired OS. Install the JNI software, and verify that the OS sees the fibre channel drives. Identify the controller and target for the installation of the Solaris® OS. The target ID must be in the range 0-9.

Note: The Solaris® system sometimes changes the adapter and controller instance number during a boot –r. If the target for the root file system is set to c2t0d0s0 in **/etc/vfstab**, but the reconfiguration boot changes the controller number to c3, the boot will fail. Before installing Solaris® on the fibre-channel driver, start with a clean system. Make sure the adapter is the first instance (fca0 or fcaw0). This may require a **boot -ar** to initialize the **/etc/path_to_inst** file.

Note the world wide name (wwn) of the target. The wwn is printed by the driver during the attach phase. It can be seen in the console window, or by viewing the /var/adm/messages file. You may use the target to world wide name binding feature for the driver to set the wwn in the /kernel/drv/sd.conf file. This is strongly recommended if the topology includes a switch or hub where configuration is subject to change. See the technotes accompanying the Solaris® device driver for details.

E.2 Partition Check

To configure the 9900 as the Solaris® boot disk, use the **format** command to verify that the target will accommodate the Sun® OS partitions. Verify that the sizes of the partitions on the target chosen for the Solaris® installation are large enough to copy the current OS partitions. Use the **format** command to examine the partitions of the system drives where the current OS resides (see Figure E.1).

```
# format
          (choose disk 0)
      format> partition
      partition> print
          (sample output:)
             Current partition table (original):
             Total disk cylinders available: 2733 + 2 (reserved cylinders)

             Part      Tag    Flag     Cylinders         Size            Blocks
             0        root     wm       0 -  133       99.45MB    (134/0/0)   203680
             1        swap     wu     134 -  241       80.16MB    (108/0/0)   164160
             2      backup     wm       0 - 2732        1.98GB    (2733/0/0) 4154160
             3  unassigned     wm       0                0         (0/0/0)          0
             4  unassigned     wm       0                0         (0/0/0)          0
             5  unassigned     wm       0                0         (0/0/0)          0
             6         usr     wm     242 -  761      385.94MB    (520/0/0)   790400
             7        home     wm     762 - 2732        1.43GB    (1971/0/0) 2995920

   quit partition, and choose the fibre channel disk on which you wish to install the OS.

      partition> quit
      format> disk
            AVAILABLE SELECTIONS ---
            specify disk:
            (Enter number of disk)
            1 (example)
     format> partition
     partition> print

            Repartition the drive (if necessary).
     partition> modify

            (exit format)
```

Figure E.1 OS System Drive Partitions

E.3 File System Creation

Creating file systems on the required partitions. In this example the OS is being copied to target 3 of the fibre channel drives on controller 1. To use a different target, change t3 to t(target id). To use a different controller, change c1 to c(controller number). Both the root partition (slice 0) and the /usr partition (slice 6) must be created. Other partitions (/export/home, /var, /opt) should also be created to mirror the current layout. Use the **newfs** command to create the filesystems (see Figure E.2).

```
# newfs     /dev/rdsk/c1t3d0s0
# newfs     /dev/rdsk/c1t3d0s6
# newfs     /dev/rdsk/c1t3d0s7
```

Figure E.2 Creating a New File System

E.4 Boot Block Installation

Install a bootblk on the root partition of the fibre channel drive. Use the **uname -a** command to determine the architecture of the workstation. This will determine the directory in /usr/platform where the bootblk will be found. The example below (see Figure E.3) gives sun4u as the architecture and subdirectory for /usr/platform. Then install a bootblk on the fibre channel target with the **installboot** command (see Figure E.4).

```
# uname -a
(RESPONSE:)
SunOS patriot 5.6 Generic sun4u Sparc SUNW,Ultra -1
```

Figure E.3 Installing a Boot Block

```
# /usr/sbin/installboot  /usr/platform/sun4u/lib/fs/ufs/bootblk  /dev/rdsk/c1t3d0s0
```

Figure E.4 Installing a Boot Block

E.5 Required Files Directories Copy

Use the **ufsdump** and **ufsrestore** command to copy the required files and directories from the current system to the fibre channel target. First mount the root directory of the fibre channel target. The example below (see Figure E.5) will create the required directory structure on the new target and copy the files. In this example the current OS is located on /dev/dsk/c0t0d0.
Note: There must be a space before and after '0f', '-' and 'rf'.

```
# mount   /dev/dsk/c1t3d0s0   /mnt
# ufsdump  0f  -  /dev/dsk/c0t0d0s0 | ( cd /mnt; ufsrestore  rf   -)
```

Figure E.5 Creating the Directory Structure for Root Directory

When the command completes, the target (c1t3d0s0) will have the complete image of the root partition. In the **/mnt/etc** directory, update the **vfstab** file to indicate the fibre channel target to be mounted during boot. Modify all partitions which will be located on the fibre channel target.

```
# cd /mnt/etc
# vi vfstab

Before Modification:
#device        device       mount        FS   fsck   mount     mount
#to mount      to fsck       point         type  pass   at boot     options
#
/dev/dsk/c0t0d0s1   -                   -      swap  -    no   -
/dev/dsk/c0t0d0s0   /dev/rdsk/c0t0d0s0  /       ufs  1    no   -
/dev/dsk/c0t0d0s6   /dev/rdsk/c0t0d0s6  /usr    ufs  1    no   -
swap                -                   /tmp  tmpfs  -    yes  -

After Modification:
#device        device       mount        FS   fsck   mount    mount
#to mount      to fsck      point         type  pass   at boot    options
#
/dev/dsk/c1t3d0s1   -                   -      swap  -    no   -
/dev/dsk/c1t3d0s0   /dev/rdsk/c1t3d0s0  /       ufs  1    no   -
/dev/dsk/c1t3d0s6   /dev/rdsk/c1t3d0s6  /usr    ufs  1    no   -
swap                -                   /tmp  tmpfs  -    yes  -
```

Figure E.6 Updating the vfstab File

Unmount the root partition, and repeat the procedure for the **/usr** partition in the current root directory (see Figure E.7). ***Note:*** There must be a space before and after '0f', '-' and 'rf'. Repeat the procedure for any other partitions to be mounted from the fibre channel target.

```
# umount /mnt

# mount /dev/dsk/c1t3d0s6 /mnt
# ufsdump  0f  -  /dev/dsk/c0t0d0s6 | ( cd /mnt; ufsrestore  rf  -)
# umount /mnt
```

Figure E.7 Creating the Directory Structure for /usr Partition

E.6 OpenBoot Modification

Halt the system to get to the OpenBoot environment. Create an OpenBoot alias to boot from the fibre channel drive. Issue the **printenv** command to see the default boot device in the environment variables.

```
# halt             # halt the system #
ok                  # now in the OpenBoot environment #
Ok  printenv
(RESPONSE:)
(You can see the following line.)
            boot-device    disk      disk
```

Figure E.8 Halt System

Issue the **devalias** command to see the device associated with the alias for the default boot device.

```
Ok  devalias                  # get a list of aliases #
(RESPONSE:)
            disk     /sbus/SUNW,fas@e,8800000/sd@0,0
```

Figure E.9 Devalias Command

Issue the **show-devs** command to see the device name for the fibre channel drive.

```
Ok  show-devs                 # get a list of devices #
(RESPONSE:)
                     /sbus@1f,0/SUNW,fas@e,8800000/sd
                     /sbus@1f,0/fca@1,0/sd
```

Figure E.10 Show-devs Command

Create a similar alias or the fibre channel drive with the **nvalias** command.

```
Ok  nvalias   fcadisk   /sbus@1f,0/fca@1/sd@3,0
```

Figure E.11 Nvalias Command

Save the alias with the **nvstore** command. Power cycle the workstation so that the alias is saved.

```
Ok   nvstore                       # save changes to nvram #
```

Figure E.12 Nvstore Command

Now boot from the fibre channel drive using the alias.

```
Ok   boot fcadisk -r               # boot from fibrechannel #
(RESPONSE:)
            The normal boot process begins.
```

Figure E.13 Boot Using Alias

If you want to set the fibre channel drive as the permanent boot drive, type the following command to check the default value.

```
Ok   printenv
(RESPONSE:)
(You can see the following line.)
            boot-device    disk      disk
```

Figure E.14 Checking Default Value

After checking the default value, type the following command to change the boot-device from the default scsi drive to the fibre channel drive.

```
Ok   setenv  boot-device  fcadisk
```

Figure E.15 Change Boot Device to Fibre-Channel Drive