HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Lightning 9900™

SGI™ IRIX® Configuration Guide

Document Revision Level

Revision	Date	Description
MK-90RD024-0	November 2000	Initial Release

Source Document Revision Level

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

- *Hitachi DKC310 Disk Subsystem Fibre Installation Manual, SGI IRIX 6.5.7*, revision 1.
- *Hitachi Lightning 9900™ LUN Manager User's Guide*, MK-90RD006-1.

Preface

The *Hitachi Freedom Storage™ Lightning 9900™ SGI™ IRIX® Configuration Guide* describes and provides instructions for configuring the devices on the Hitachi Lightning 9900™ disk array subsystem for operation with the SGI™ IRIX® 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 SGI™ servers and the fibre-channel adapters, and
- the user is familiar with the SGI™ IRIX® operating system and the UNIX® file system, system commands, and utilities.

Note: The term "9900" refers to the entire Hitachi Lightning 9900™ subsystem family, unless otherwise noted.

For further information on the 9900 disk array subsystems, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide* (MK-90RD008), or contact your Hitachi Data Systems account team. You can also visit Hitachi Data Systems online at http://www.hds.com.

Note: For further information on the SGI™ IRIX® operating system, please consult the IRIX® user documentation, or contact SGI™ customer support services.

Contents

Chapter 1 Overview of 9900 SGI™ IRIX® Configuration

1.1 9900 SGI™ IRIX® Configuration 1
1.2 Hitachi Lightning 9900™ Subsystem 1
1.3 Device Types and Configuration Procedures 2

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements 5
2.2 Installing the 9900 Subsystem 6
2.3 Preparing for 9900 Device Configuration 7
2.3.1 Setting the Host Mode for the 9900 Ports 7
2.3.2 Configuring the 9900 Fibre-Channel Ports 8
2.4 Connecting the 9900 Subsystem to the SGI™ System 10

Chapter 3 Configuring the New Devices

3.1 Verifying New Device Recognition 12
3.1.1 Device Files and WWNs 12
3.2 Partitioning the Disk Devices 14
3.3 Enabling Command Tag Queuing 16
3.4 Creating the File Systems 18
3.4.1 EFS File System 18
3.4.2 XFS™ File System 18
3.5 Creating the Mount Directories and Mounting the Devices 19
3.6 Verifying the File Systems 19
3.7 Setting the Auto-Mount Parameters 20

Chapter 4 Failover and SNMP Configuration

4.1 Host/Application Failover 21
4.2 SNMP Remote Subsystem Management 22

Chapter 5 Troubleshooting

5.1 Troubleshooting 23
5.2 Calling the Support Center 24

Appendix A Acronyms and Abbreviations 25

Appendix B SCSI TID Map for Fibre-Channel Adapters 27

List of Figures

Figure 2.1 Setting the Host Mode 7
Figure 2.2 Fibre Parameter Panel (from the LUN Manager Software) 9

Figure 3.1 Verifying New Device Recognition 13
Figure 3.2 Partitioning and Labeling the Disk Devices (continues on the next page) 14
Figure 3.2 Partitioning the Disk Devices (continued from previous page) 15
Figure 3.3 Verifying a Partition. 15
Figure 3.4 **Sar** Command 16
Figure 3.5 Enabling Command Tag Queuing and Setting the Queue Depth 17
Figure 3.6 Verifying the New File Systems 19
Figure 3.7 Setting the Auto-Mount Parameters 20

Figure 4.1 9900 SNMP Environment. 22

List of Tables

Table 1.1 9900 Device Specifications 3

Table 2.1 Fibre Parameter Settings on the 9900 Remote Console PC 8
Table 2.2 Available AL-PA Values 9

Table 3.1 Queue Depth Requirements for the 9900 Devices 16
Table 3.2 Auto-Mount Parameters 20

Table 5.1 Troubleshooting 23

Table B.1 AL-PA to SCSI TID Mapping (t value) for SGI™ IRIX® Systems 28

Chapter 1 Overview of 9900 SGI™ IRIX® Configuration

1.1 9900 SGI™ IRIX® Configuration

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

Configuration of the 9900 disk devices for SGI™ IRIX® operations includes:

- Verifying new device recognition (see section 3.1),
- Partitioning the devices (see section 3.2),
- Enabling command tag queuing (CTQ) (see section 3.3),
- Creating the file systems (see section 3.4),
- Creating the mount directories and mounting the devices (see section 3.5),
- Verifying the file systems (see section 3.6), and
- Setting the auto-mount parameters (see section 3.7).

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

1.2 Hitachi Lightning 9900™ Subsystem

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

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

1.3 Device Types and Configuration Procedures

The 9900 subsystem allows the following types of logical devices (also called volumes) to be installed and configured for operation with the SGI™ IRIX® operating system. Table 1.1 lists the device specifications 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 custom-size LUs which are smaller than standard OPEN-x LUs. The Custom Volume Size (CVS) feature of the 9900 subsystem (also called Virtual LUN and Virtual LVI) enables you to configure CVS devices. The CVS capability enables you to "slice up" a single LU into several smaller LUs to best fit the application needs and improve host access to frequently used files. For further information on the CVS feature, please refer to the *Hitachi Lightning 9900™ Virtual LVI/LUN User's Guide* (MK-90RD005).

Configuration of the 9900 disk devices for SGI™ IRIX® operations includes:

- Verifying new device recognition (see section 3.1),
- Partitioning the devices (see section 3.2),
- Enabling command tag queuing (CTQ) (see section 3.3),
- Creating the file systems (see section 3.4),
- Creating the mount directories and mounting the devices (see section 3.5),
- Verifying the file systems (see section 3.6), and
- Setting the auto-mount parameters (see section 3.7).

Table 1.1 9900 Device Specifications

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

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

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

Note 2: The 9900 command device (used for Hitachi Command Control Interface operations) is distinguished by **-CM** on the product name (e.g., OPEN-3-CM, OPEN-3-CVS-CM).

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

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

The requirements for 9900 SGI™ IRIX® configuration are:

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

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

- **Server:** SGI™ O2 system, OCTANE®, Onyx2®, Challenge®, ORIGIN™ 200, or ORIGIN™ 2000 system. Please contact your Hitachi Data Systems account team for further information on server hardware requirements.

- **Operating system:** SGI™ IRIX® operating system version 6.5.8. ***Important:*** Please contact Silicon Graphics® to make sure that the most current OS patches are installed on the SGI™ systems(s).

  Note: For further information on SGI™ IRIX® version support, please contact your Hitachi Data Systems account team.

  Note: **Root** (superuser) login access to the SGI™ IRIX® system is required.

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

2.2 Installing the 9900 Subsystem

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

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

     Note: The SGI™ IRIX® system can only recognize up to 32 LUs (per port). Make sure to determine the number of 9900 LUs based on this restriction.
   - Installing the fibre 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 SGI™ IRIX® system to hang. Always confirm that the devices on the fibre cable are offline before connecting/disconnecting the fibre cable.

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

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

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

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

2.3 Preparing for 9900 Device Configuration

Before the 9900 is connected to the SGI™ system, you must perform the following tasks:

- Set the host mode for the 9900 fibre port(s) (see section 2.3.1), and
- Configure the 9900 fibre ports (fibre only) (see section 2.3.2).

You will use the LUN Manager remote console software to set the host modes for the 9900 ports and configure 9900 fibre ports. For instructions on using the LUN Manager software, please refer to the *Hitachi Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

After completing these steps, you will shut down the SGI™ system, connect the 9900 subsystem, and then restart the SGI™ system (see section 2.4).

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. The required host mode setting for 9900 SGI™ IRIX® operations is **00** (standard mode, default setting). Use the LUN Manager remote console software to ensure that the host mode for each port connected to the IRIX® system is 00 (see Figure 2.1).

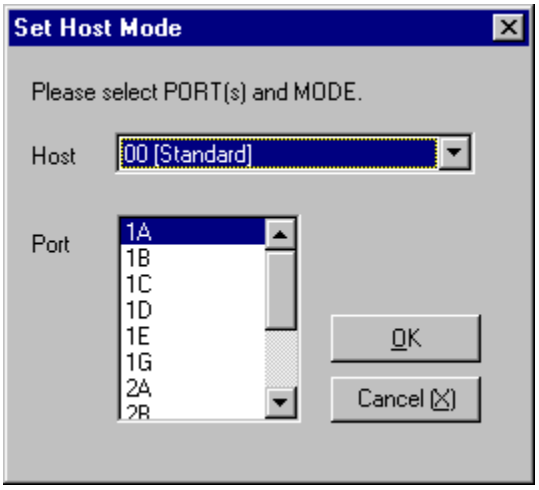

Figure 2.1 Setting the Host Mode

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 256 devices per port, but the SGI™ IRIX® system only recognizes 32 devices per port.

Fibre topology. Figure 2.2 shows the Fibre Parameter panel (part of the LUN Manager software), and Table 2.1 explains the settings on this panel. You will select the appropriate settings for each 9900 FC port based on the device to which the port is connected. Determine the topology parameters supported by the device, and set your topology accordingly. The type of 9900 port is also important.

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 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 B for a description of the AL-PA-to-TID translation.

Note on loop ID conflict: The SGI™ system assigns port addresses from lowest (01) to highest (EF). To avoid loop ID conflict, assign the port addresses from highest to lowest (i.e., starting at EF). The AL-PAs should be unique for each device on the loop to avoid conflicts. Do not use more than one port address with the same TID in same loop (e.g., addresses EF and CD both have TID 0, refer to Appendix B for the TID-to-AL-PA mapping).

Table 2.1 Fibre Parameter Settings on the 9900 Remote Console PC

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

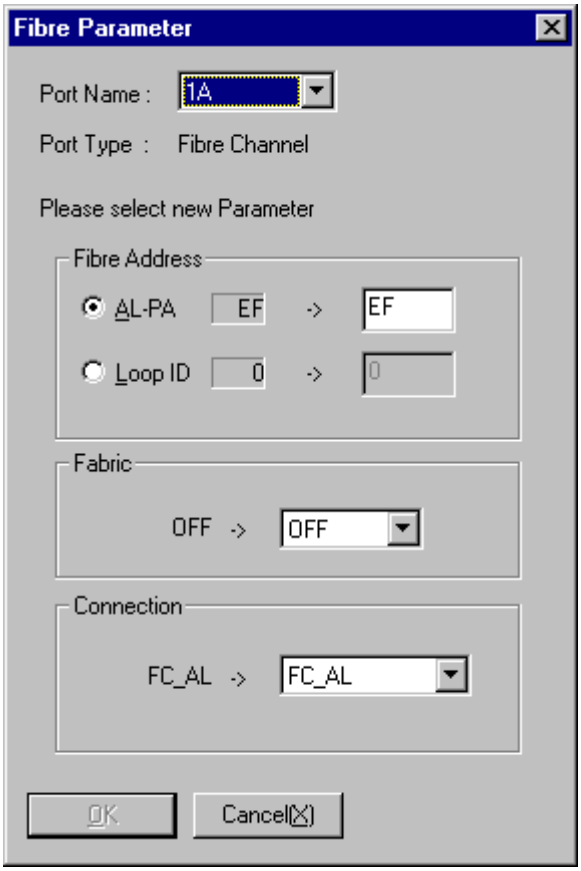

Figure 2.2 Fibre Parameter Panel (from the LUN Manager Software)

Table 2.2 Available AL-PA Values

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

2.4 Connecting the 9900 Subsystem to the SGI™ System

After you have set the host modes for the 9900 ports and configured the 9900 fibre ports, you are ready to connect the 9900 subsystem to the SGI™ system. The 9900 subsystem comes with all hardware and cabling required for connection to the host system(s).

To connect the 9900 subsystem:

1. **Verify subsystem installation**. The Hitachi Data Systems representative verifies the fibre-port address configuration and the status of the FC adapters and LDEVs (normal).
2. **Shut down the SGI™ system**. The user should perform this activity. You must shut down and power off the SGI™ system before connecting the 9900:
   a) Shut down the SGI™ system as usual (e.g., **shutdown -y -g0 -i0**).
   b) When shutdown is complete, power off the SGI™ display.
   c) Power off all peripheral devices.
   d) Power off the SGI™ system.
3. **Connect the 9900 to the SGI™ system**. The Hitachi Data Systems representative installs the fibre-channel cables between the 9900 and the SGI™ system.

   Note: The Hitachi Data Systems representative must use the 9900 maintenance manual during all installation activities. Follow all precautions and procedures in the maintenance manual, and always check all specifications to ensure proper installation and configuration.
4. **Power on the SGI™ system**. The user or Hitachi Data Systems representative can perform this activity. To power on the SGI™ system after connecting the 9900 subsystem:
   a) Power on the SGI™ system display.
   b) Power on all peripheral devices. The 9900 subsystem should already be on, the host modes should already be set, and the fibre-channel ports should already be configured. If the host modes or FC ports are configured after the SGI™ system is powered on, the system must be restarted.
   c) Confirm the ready status of all devices.
   d) Power on the SGI™ system. ***Note:*** Some SGI™ systems may require you to use the System Maintenance Menu to start the system.

Chapter 3 Configuring the New Devices

Configuration of the new 9900 disk devices is performed by the user and requires **root** access to the SGI™ IRIX® system. The host modes for the 9900 ports must already be set, and the 9900 fibre ports must already be configured (see section 2.3). If any of these steps is performed after the SGI™ system is powered on, you must halt and restart the system before configuring the new 9900 devices.

The activities involved in 9900 device configuration are:

- Verifying new device recognition (see section 3.1),
- Partitioning the devices (see section 3.2),
- Enabling command tag queuing (CTQ) (see section 3.3),
- Creating the file systems (see section 3.4),
- Creating the mount directories and mounting the devices (see section 3.5),
- Verifying the file systems (see section 3.6), and
- Setting the auto-mount parameters (see section 3.7).

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

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

3.1 Verifying New Device Recognition

The first step in configuring the 9900 devices for SGI™ IRIX® operations is to verify that the SGI™ system recognizes the new devices on the 9900 subsystem. The SGI™ IRIX® system creates device files for new devices automatically during server startup.

WARNING: Make sure that LUN 0 is defined for each target ID. If LUN 0 is not defined, SGI™ IRIX® will not recognize the rest of the LUs on that target ID.

To verify that the SGI™ IRIX® system recognizes the new 9900 devices:

1. Log in to the SGI™ system as **root**.
2. Display the peripheral device information using the **hinv** command (see Figure 3.1).
3. Verify that the system recognizes all new 9900 devices.

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

3.1.1 Device Files and WWNs

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

The format for device file names in SGI™ IRIX® is:

- FC-AL (and parallel SCSI):
  /dev/rdsk/dkscontroller#ddrive#{spartition#|vh|vol}
  /dev/rdsk/dkscontroller#ddrive#llun#{spartition#|vh|vol}
  /dev/dsk/dkscontroller#ddrive#spartition#
  /dev/dsk/dkscontroller#ddrive#llun#spartition#
  Example: /dev/rdsk/dks8d0l8s0

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

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

Figure 3.1 Verifying New Device Recognition

3.2 Partitioning the Disk Devices

After new device recognition has been verified, you need to partition the new SCSI disk devices using the **fx** utility (see Figure 3.2). After setting the partitions for a device, verify the partitions using the **prtvtoc** command (see Figure 3.3).

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

Note: Do not change partitions 8, 9, or 10. Set the partitions (except partition 10) so as not to a partition is on top of other partition. Because partition 8 usually uses parts between 0 and 6, set that base of first setting partition is 6.

WARNING: Do not partition or label devices which will be accessed as raw devices.

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

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

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

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

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

Figure 3.2 Partitioning and Labeling the Disk Devices (continues on the next page)

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

capacity is 4806656 blocks

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

Figure 3.2 Partitioning the Disk Devices (continued from previous page)

```
# prtvtoc /dev/dsk/dks8d0l15s0          ← Verifying partition 0 of LUN15 on TID0, controller 8.
```

Figure 3.3 Verifying a Partition

3.3 Enabling Command Tag Queuing

Command tag queuing (CTQ) must be enabled to optimize the performance of the 9900 devices. Since CTQ is disabled by default in IRIX, you need to enable it and set the queue depth for each 9900 logical device (including HMDE devices) using the **fx** utility. Table 3.1 lists the queue depth requirements for the 9900 devices.

Table 3.1 Queue Depth Requirements for the 9900 Devices

Parameter	Required Value
Queue depth per LU	≤ 32
Queue depth per port	≤ 256

Note: You can adjust the queue depth for the 9900 devices later as needed (within the specified range) to optimize the I/O performance of the 9900 devices. If I/O response time will be long, you must adjust queue depth parameter. You can check the response time using the **sar** command (see Figure 3.4) and set queue depth within 10 seconds.

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

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

Figure 3.4 Sar Command

To enable CTQ and set the queue depth for the 9900 devices (see Figure 3.5):

1. Start the **fx** disk utility and select the desired device to configure.
2. Once the device is selected and the **fx>** prompt reappears, enter **/label/set/para** to set the command tag queuing and queue depth options.
3. When prompted, enter **enable** to enable CTQ, and enter the desired CTQ depth (e.g., 8).
4. When prompted, enter **yes** to modify the drive parameters as specified.
5. Exit the **fx** utility, and enter **yes** to write out (save) the changes to the drive parameters.
6. Repeat steps (1) through (5) for each 9900 disk device, including OPEN-x, CVS, LUSE, and HMDE devices.

```
# fx -x "dksc(8,0,2)"                                    ← Start fx and enter 9900 device.

fx version 6.5, Jan 11, 2000
...opening dksc(8,0,2)
...drive selftest...OK
fx: Warning:  invalid label from disk driver, ignored
Scsi drive type == HITACHI      OPEN-3-CVS      5244
...creating default bootinfo
...created default partitions, use /repartition menu to change
...creating default volume directory

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

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

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

label info has changed for disk dksc(8,0,2). write out changes? (yes) yes   ← Enter "yes".
```

Figure 3.5 Enabling Command Tag Queuing and Setting the Queue Depth

3.4 Creating the File Systems

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

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

3.4.1 EFS File System

To create an EFS file system:

1. Use the **mkfs** command to create an EFS file system. For example, to create an EFS file system for controller 8, drive (TID) 0, logical unit 15, partition 0, enter:
   mkfs /dev/rdsk/dks8d0l15s0
2. Repeat step (1) for each device partition on which you want to create an EFS file system.

3.4.2 XFS™ File System

To create an XFS™ file system:

1. Use the **mkfs** command to create an XFS™ file system. For example, to create an XFS™ file system for controller 8, drive (TID) 0, logical unit 15, partition 0, enter:
   mkfs –xfs -d name=/dev/rdsk/dks8d0l15s0 -b size=1k -l internal,size=10m
2. Repeat step (1) for each device partition on which you want to create an XFS™ file system.

3.5 Creating the Mount Directories and Mounting the Devices

After you have created the file systems for the new 9900 SCSI disk devices, you can create the mount directories and mount the new devices. Make sure to choose a unique directory name which identifies the logical volume.

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

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

1. Create the desired new mount directories using the **mkdir** command. For example, to create a mount directory for logical unit 0 on the 9900, enter:
   mkdir /9900_LU00
2. Mount all new 9900 devices using the **mount** command. For example, to mount partition 0 of LUN 15, drive (TID) 0, controller 8, enter:
   mount /dev/rdsk/dks8d0l15s0 /9900_LU00

3.6 Verifying the File Systems

Verify the file systems for the new 9900 disk devices using the **df** command (see Figure 3.6). Make sure that the capacity value (**Kbytes**) for each device is correct.

```
# df -k                                                          ← List file systems.
File system             Type    Kbytes     use    avail  %use  Mounted on
/dev/root               xfs     969857  414702   555155   43%   /
/dev/dsk/dks8d0l15s0    xfs    1048576      13  1048563    0%   /9900_LU0          ← New device.
:
:
#
```

Figure 3.6 Verifying the New File Systems

3.7 Setting the Auto-Mount Parameters

The final step in configuring the 9900 devices for SGI™ IRIX® operations is to set the auto-mount parameters for the new devices. You will add each device to be auto-mounted to the system auto-mount table (**/etc/fstab** file). If you do not plan to auto-mount any of the 9900 devices, you can skip this section.

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

1. First make a backup copy of the mount table:
   cp /etc/fstab /etc/fstab.backup
2. Add each desired new device to the mount table as shown in Figure 3.7. Table 3.2 describes the auto-mount parameters.

```
# cp /etc/fstab  /etc/fstab.backup                                  ← Make backup.
# vi /etc/fstab                                                     ← Edit mount table.
/dev/root            /             xfs  rw,raw=/dev/rroot  0  0
/dev/dsk/dks8d0l15s0 /DKC310_LU00  xfs  rw,noquota         0  1     ← Enter new device.
          ①             ②          ③       ④             ⑤  ⑥       ← See Table 3.2.
:
```

Figure 3.7 Setting the Auto-Mount Parameters

Table 3.2 Auto-Mount Parameters

No.	Description
①	Device to mount (device file name).
②	Mount point (mount directory).
③	File system (FS) type.
④	Mount options (usually [rw,noquota]).
⑤	Enhance – enter 0 for 9900 devices.
⑥	fsck pass – order in which FS checks are to be performed.

Chapter 4 Failover and SNMP Configuration

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

- IRIS FailSafe™ and SGI™ Advanced Cluster Environment (ACE) software for application/host failover.
- The XLV extended logical volume manager for IRIX.

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

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

4.1 Host/Application Failover

The 9900 subsystem supports the IRIS FailSafe™ and SGI™ Advanced Cluster Environment (ACE) software products for the SGI™ IRIX® operating system. Please contact SGI™ customer support services for the latest information on failover products and solutions.

After 9900 device configuration is complete, make sure to configure the failover software on the host server(s) as needed to recognize the new devices on the 9900 subsystem. For assistance with IRIS FailSafe™ and/or SGI™ ACE operations, please refer to the user documentation, or contact SGI™ customer support services.

4.2 SNMP Remote Subsystem Management

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

When a SIM occurs, the 9900 SNMP agent initiates trap operations, which alert the SNMP manager of the SIM condition. The SNMP manager receives the SIM traps from the SNMP agent, and can request information from the SNMP agent at any time.

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

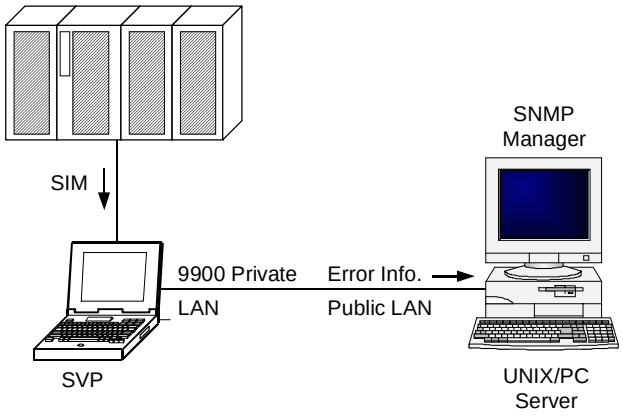

Figure 4.1 9900 SNMP Environment

Chapter 5 Troubleshooting

5.1 Troubleshooting

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

Table 5.1 lists potential error conditions during SGI™ IRIX® installation and configuration and provides instructions for resolving each condition. If you are unable to resolve an error condition, please consult your Hitachi Data Systems Customer Service representative, or contact the Hitachi Data Systems Support Center.

Table 5.1 Troubleshooting

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Make sure that the READY indicator lights on the 9900 subsystem are ON. Make sure that the fibre cables are correctly installed and firmly connected. Make sure that the LUNs are properly configured. The LUNs for each target ID must start at 0 and continue sequentially without skipping any numbers. Make sure that the SCSI IDs (unit IDs) on each bus are unique. Do not connect two devices with the same SCSI ID on the same bus.
The SGI™ system does not reboot properly after hard shutdown	If the SGI™ system is powered off without executing the shutdown process, wait three minutes before restarting the SGI™ system. This allows the 9900's internal time-out process to purge all queued commands so that the 9900 is available (not busy) during system startup. If the SGI™ system is restarted too soon, the 9900 will continue trying to process the queued commands, and the SGI™ system will not reboot successfully.
The SGI™ system hangs.	Make sure that the target IDs are set 0 through 6 and 8 through 15 and that target ID 7 has been reserved for the fast-wide SCSI controller card.

5.2 Calling the Support Center

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

The worldwide Hitachi Data Systems Support Centers are:

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

Appendix A Acronyms and Abbreviations

AL	arbitrated loop
AL-PA	arbitrated loop physical address
CTQ	command tag queuing
CXFS™	clustered version of XFS™ file system
EFS	IRIX® standard file system
ESCON®	Enterprise System Connection (IBM trademark for optical channels)
FC	fibre-channel
FCA	fibre-channel adapter
FC-AL	fibre-channel arbitrated loop
FCP	fibre-channel protocol
fx	IRIX® disk utility
HMDE	Hitachi Multiplatform Data Exchange
HP	Hewlett-Packard Company
I/O, IO	input/output
LDEV	logical device
LU	logical unit
LUN	logical unit number, logical unit
LUSE	LU Size Expansion
LVM	Logical Volume Manager, logical volume management
OFC	open fibre control
P-P	point-to-point
PA	physical address
PC	personal computer system
R-SIM	remote service information message
RAID	redundant array of independent disks
SCSI	small computer system interface
SGI	Silicon Graphics, Incorporated
SIM	service information message
SNMP	simple network management protocol
SVP	service processor
TID	target ID
XFS™	IRIX® extended file system
XLV	extended logical volume manager

Appendix B SCSI TID Map for Fibre-Channel Adapters

When an arbitrated loop (AL) is established or re-established, the port addresses are assigned automatically to prevent duplicate TIDs. With the SCSI over fibre-channel protocol (FCP), there is no longer a need for target IDs in the traditional sense. SCSI is a bus-oriented protocol requiring each device to have a unique address since all commands go to all devices. For fibre channel, the AL-PA is used instead of the TID to direct packets to the desired destination. Unlike traditional SCSI, once control of the loop is acquired, a point-to-point connection is established from initiator to target. To enable transparent use of FCP, the SGI™ IRIX® system "maps" a TID to each AL-PA.

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

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

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

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

Table B.1 AL-PA to SCSI TID Mapping (t value) for SGI™ IRIX® Systems

AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value	AL-PA	t value
EF	0	CD	0	B2	0	98	0	72	0	55	0	3A	0	25	0
E8	1	CC	1	B1	1	97	1	71	1	54	1	39	1	23	1
E4	2	CB	2	AE	2	90	2	6E	2	53	2	36	2	1F	2
E2	3	CA	3	AD	3	8F	3	6D	3	52	3	35	3	1E	3
E1	4	C9	4	AC	4	88	4	6C	4	51	4	34	4	1D	4
E0	5	C7	5	AB	5	84	5	6B	5	4E	5	33	5	1B	5
DC	6	C6	6	AA	6	82	6	6A	6	4D	6	32	6	18	6
DA	7	C5	7	A9	7	81	7	69	7	4C	7	31	7	17	7
D9	8	C3	8	A7	8	80	8	67	8	4B	8	2E	8	10	8
D6	9	BC	9	A6	9	7C	9	66	9	4A	9	2D	9	0F	9
D5	10	BA	10	A5	10	7A	10	65	10	49	10	2C	10	08	10
D4	11	B9	11	A3	11	79	11	63	11	47	11	2B	11	04	11
D3	12	B6	12	9F	12	76	12	5C	12	46	12	2A	12	02	12
D2	13	B5	13	9E	13	75	13	5A	13	45	13	29	13	01	13
D1	14	B4	14	9D	14	74	14	59	14	43	14	27	14	00	-
CE	15	B3	15	9B	15	73	15	56	15	3C	15	26	15		