HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Lightning 9900™

IBM® AIX® Configuration Guide

Notice of Export Controls

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

Document Revision Level

Revision	Date	Description
MK-90RD014-0	July 2000	Initial Release
MK-90RD014-1	November, 2000	Revision 1, supersedes and replaces MK-90RD014-0

Source Document Revision Level

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

- *Hitachi DKC310 Disk Subsystem SCSI/Fibre Installation Manual, IBM RS/6000 series, AIX Version 4.1, 4.2, 4.3*, revision 3.
- *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide*, MK-90RD006-1.
- ECN TW-90RD042 (Technical Writing ECN).
- ECN TW-90RD044 (Technical Writing ECN).
- ECN TW-90RD045 (Technical Writing ECN).

Preface

The *Hitachi Freedom Storage™ Lightning 9900™ IBM® AIX® Configuration Guide* describes and provides instructions for installing and configuring the devices on the Hitachi Lightning 9900™ disk array subsystem for operation with IBM® AIX® 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 Lightning 9900™ array subsystem,
- the user is familiar with the IBM® AIX® operating system and the IBM® RISC System/6000 (RS/6000)®, POWERstation®, POWERserver®, and/or SP system,
- the user is familiar with the AIX® Journaled File System, system commands, and utilities.

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

For further information on IBM® AIX®, please consult the IBM® AIX® online help and/or user documentation, or contact IBM® technical support.

Contents

List of Figures

List of Tables

Chapter 1 Overview of IBM® AIX® Configuration

1.1 IBM® AIX® Configuration

This document describes the requirements and procedures for connecting the 9900 subsystem to an IBM® AIX® system and configuring the new 9900 devices for operation with the IBM® AIX® operating system. The Hitachi Data Systems 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.

Configuration of the 9900 SCSI disk devices for IBM® AIX® operations includes:

- Changing the device parameters (see section 3.1),
- Assigning the new devices to volume groups and setting the partition size (see section 3.2),
- Creating the Journaled File Systems (see section 3.3.1), and
- Mounting and verifying the file systems (see section 3.3.2).

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 The Lightning 9900™ Array Subsystem

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

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

1.3 Device Types and Configuration Procedures

The 9900 subsystem allows the following types of logical devices (also called LDEVs) to be installed and configured for operation with the IBM® AIX® 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) are disk devices of predefined sizes. The 9900 subsystem currently supports OPEN-3, OPEN-8, OPEN-9, OPEN-K, and OPEN-E devices. Please contact your Hitachi Data Systems account team for the latest information on supported LU types.

LUSE Devices (OPEN-x*n) . The LUSE devices are combined LUs which can be from 2 to 36 times larger than standard OPEN-x LUs. The Logical Unit Size Expansion (LUSE) feature of the 9900 subsystem enables you to configure these custom-size devices. LUSE devices are designated as OPEN-x*n, where x is the LU type (e.g., OPEN-9*n) and $2 \leq n \leq 36$. For example, a LUSE device created from ten OPEN-3 LUs would be designated as an OPEN-3*10 disk device. This capability enables the server host to combine logical devices and access the data stored on the 9900 subsystem using fewer LU numbers (LUNs). For further information on the LUSE feature, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

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

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

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

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

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

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

Table 1.1 9900 Device Specifications for IBM® AIX® Operations (continues on the next page)

Device Type *(Note 1)*	Category *(Note 2)*	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Data Cylinders	# of Heads	# of Sectors per Track	Capacity MB *(Note 3)*
OPEN-3	SCSI Disk	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
OPEN-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	13893
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	13893*n
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 4*	512	*Note 5*	15	96	*Note 6*

Table 1.1 9900 Device Specifications for IBM® AIX® Operations (continued)

Device Type *(Note 1)*	Category *(Note 2)*	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Data Cylinders	# of Heads	# of Sectors per Track	Capacity MB *(Note 3)*
OPEN-3*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 4*	512	*Note 5*	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	2842F
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 file systems for IBM® AIX® operations. The HMDE devices (3390-3A/B/C, 3380-KA/B/C, OPEN-x-HMDEoto) must be installed as raw devices and can only accessed using HMDE. Do not create a partition or file system on any device used for HMDE operations.

Table 1.2 Volume Usage for Device Categories

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

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

Note 3: The device capacity can sometimes be changed by the BIOS or host adapter board. These device capacities are calculated based on 1 MB = 1024^2 bytes rather than 1000^2 bytes.

Note 4: The number of blocks for a 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 3) × (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 IBM® AIX® 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.
- IBM® RS/6000®, POWERstation®, POWERserver®, or SP series system.
- IBM® AIX® operating system, version 4.2 or 4.3. ***Important:*** Please contact IBM® to make sure the most current OS patches are installed on the IBM® system(s).

  Note: Hitachi Data Systems plans to support future releases of IBM® AIX®. For the latest information on AIX® version support, contact your Hitachi Data Systems account team.
- **Root** (superuser) login access to the IBM® system.
- Fibre-channel adapters. Make sure to install all utilities, tools, and drivers that come with the adapter(s).

  The 9900 subsystem supports full-speed (100 MB/s) fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface, and multimode optical cables with SC connectors. Do not connect any OFC-type fibre-channel interface to the 9900 subsystem. For information on supported FC adapters (FCAs), optical cables, hubs, and fabric switches, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center (see section 5.2).

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

  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	Driver(s)	AIX® OS Version
IBM® FC 6227	4.3.3.0 (Firmware: 2.22x1)	4.3.3 (32 bit / 64 bit)

2.2 Installing the 9900 Subsystem

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

1. **Hardware installation**. The Hitachi Data Systems representative performs hardware installation as specified in the 9900 Maintenance Manual. Follow all precautions and procedures in the 9900 maintenance manual. Check all specifications to ensure proper installation and configuration. Hardware installation includes:
   - Assembling all hardware and cabling.
   - Loading the latest microcode and service processor (SVP) updates for full fibre-channel support.
   - Installing and formatting the logical devices (LDEVs) using the SVP. Make sure to get the desired LDEV configuration information from the user, including the desired number of OPEN-x, LUSE, CVS, CVS LUSE, and multiplatform (HMDE) devices.
   - Installing the fibre-channel adapters and cabling. ***Note:*** 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 IBM® AIX® system to hang. Always confirm that the devices on the fibre cable are offline before connecting/disconnecting the fibre cable.

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

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

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

2.3 Preparing to Connect the 9900

Before the 9900 is connected to the AIX® 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).

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

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 that must be set for the connected operating system. ***Note***: The required host mode setting for 9900 AIX® operations is **0F**, which is **different** than the standard mode shown in the following panel. Use the LUN Manager remote console software to ensure that the host mode for each fibre port connected to the AIX® system is 0F.

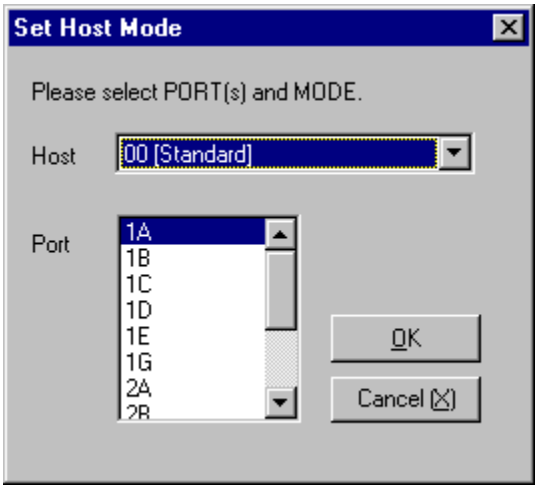

Figure 2.1 Set Host Mode Panel

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 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 target ID (TID). See Appendix B for a description of the AL-PA-to-TID translation.

Note on loop ID conflict: The AIX® 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 AL-PA-to-TID mapping).

Table 2.1 Fibre Parameter Settings on the 9900 Remote Console PC

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

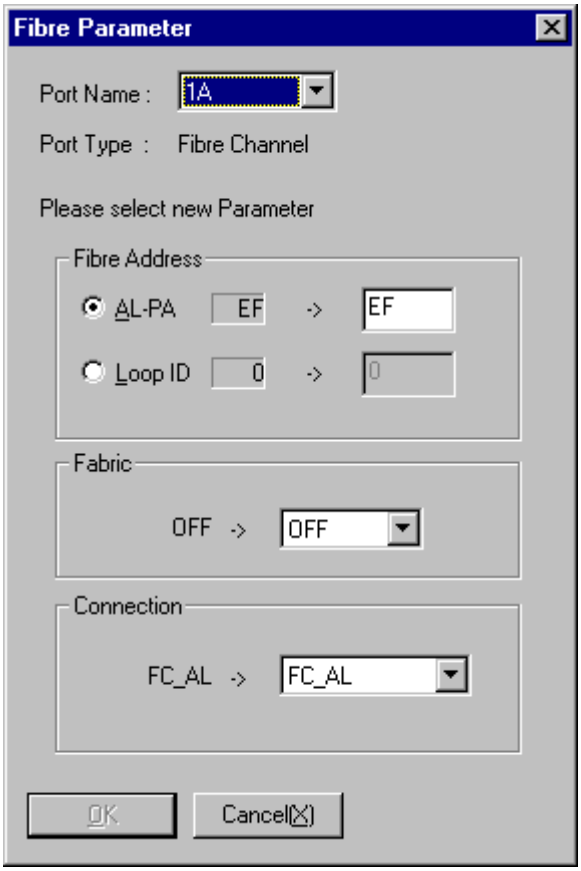

Figure 2.2 Fibre Parameter Panel (from the LUN Manager software)

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

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

To connect the 9900 subsystem to the IBM® system:

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 and power off the IBM® system**. The user should perform this activity. You must shut down and power off the AIX® system before connecting the 9900:
   a) Shut down the IBM® system.
   b) When shutdown is complete, power off the IBM® AIX® display.
   c) Power off all peripheral devices except for the 9900 subsystem.
   d) Power off the IBM host system. You are now ready to connect the 9900 subsystem.
3. **Connect the 9900 to the IBM® system**. The Hitachi Data Systems representative installs the fibre cables between the 9900 and the IBM® 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 and boot up the IBM® system**. The user should perform this activity. To power on the IBM® system after connecting the 9900:
   a) Power on the IBM® system display.
   b) Power on all peripheral devices. The 9900 subsystem should already be on, the host modes should already be set, and the fibre-channel ports should already be configured. If the host modes or fibre ports are configured after the IBM® system is powered on, the IBM® system may need to be restarted in order to recognize the new devices.
   c) Confirm the ready status of all peripheral devices, including the 9900.
   d) Power on and boot up the IBM® system connected to the 9900.

2.5 Verifying New Device Recognition

The final step before configuring the new 9900 disk devices is to verify that the host system recognizes the new devices. The host system automatically creates a device file for each new device recognized. The devices should already be installed and formatted and the fibre ports configured before the host system is powered on. If the system was not restarted, the user must issue the **cfgmgr** command to force the system to check the buses for new devices.

To verify new device recognition:

1. Log in to the host system as **root**.
2. Display the system device data using the **lsdev -C -s scsi** command (see Figure 2.3).
3. Verify that the system recognizes all new disk devices, including OPEN-x, LUSE, CVS, CVS LUSE, and HMDE devices. The devices are listed by device file name.
4. Make a blank table (see Table 2.3) for recording the 9900 device data. The table must include the device file name, bus number, TID, LUN, and device type for each new device.
5. Record the device information for all new devices in your device data table (see Table 2.3). You will need this information in order to change the device parameters.

```
# lsdev -C -s scsi                                          ∧ Display device data.
hdisk0  Available 00-00-00-0,0  2.0GB SCSI Disk Drive
cd0     Available 00-00-00-2,0  Multimedia CD-ROM Drive
rmt0    Defined   00-00-00-4,0  2.3 GB 8mm Tape Drive
hdisk1  Available 00-01-00-2,0  Other SCSI Disk Drive     ∧ New device.
hdisk2  Available 00-01-00-2,1  Other SCSI Disk Drive     ∧ New device.
   ⇑  Device file name = hdiskx.
 :
#
```

Note: This sample screen shows the following information:
The device hdisk1 is TID=2, LUN=0 on bus 1. The device hdisk2 is TID=2, LUN=1 on bus 1.

Figure 2.3 Verifying New Device Recognition

Note: You will need the device file names for the HMDE devices when you create the HMDE volume definition file (**datasetmount.dat**). For example, if **hdisk3** is a 3390-3B HMDE device, the entry for this volume in the HMDE volume definition file is:

```
\\.\PHYSICALDRIVE3 XXXXXX 3390-3B      (XXXXXX is the VOLSER)
```

Table 2.3 Device Data Table (Sample)

Device File Name	Bus No.	TID	LUN	Device Type	Alternate Path(s)	
hdisk1					TID:____ LUN:____	TID:____ LUN:____
hdisk2					TID:____ LUN:____	TID:____ LUN:____
hdisk3					TID:____ LUN:____	TID:____ LUN:____
hdisk4					TID:____ LUN:____	TID:____ LUN:____
hdisk5					TID:____ LUN:____	TID:____ LUN:____
hdisk6					TID:____ LUN:____	TID:____ LUN:____
hdisk7					TID:____ LUN:____	TID:____ LUN:____
hdisk8					TID:____ LUN:____	TID:____ LUN:____
hdisk9					TID:____ LUN:____	TID:____ LUN:____
and so on...						

Chapter 3 Configuring the New Devices

Configuration of the 9900 disk devices is performed by the user and requires root (superuser) access to the AIX® system. The host modes for the 9900 fibre ports must already be set, and the 9900 fibre ports must already be configured (refer to Chapter 2). If any of these steps is performed after the AIX® system is powered on, you must stop and restart the system before configuring the new devices.

Configuration of the 9900 SCSI disk devices for IBM® AIX® operations includes:

- Changing the device parameters (see section 3.1),
- Assigning the new devices to volume groups and setting the partition size (see section 3.2),
- Creating the Journaled File Systems (see section 3.3.1), and

Mounting and verifying the file systems (see section 3.3.2).

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

Online device installation. For information on configuring newly installed 9900 devices without rebooting the AIX® system, please refer to Appendix C.

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

3.1 Changing the Device Parameters

When the device files are created, the IBM® system sets the device parameters to the system default values. You must change the following device parameters for each new 9900 device: read/write (r/w) time-out, queue type, and queue depth. Table 3.1 specifies the r/w time-out and queue type requirements for the 9900 devices. Table 3.2 specifies the queue depth requirements for the 9900 devices.

AIX® uses the Logical Volume Manager (LVM) (accessed from within SMIT) to manage data storage. You can use either SMIT or the AIX® command line to perform this procedure. Make sure to set the parameters for the HMDE devices as well as the SCSI disk devices. Make sure to use the same settings and device parameters for all 9900 devices.

Table 3.1 R/W Time-Out and Queue Type Requirements

Parameter Name	Default Value	Required Value for 9900
Read/write time-out	30	60
Queue type	none	simple

Table 3.2 Queue Depth Requirements for the 9900 Devices

Parameter	Requirement
Queue depth per LU	≤ 32
Queue depth per port (MAXTAGS)	≤ 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 devices.

Instructions for SMIT

To change the device parameters using SMIT:

1. Enter **smit** on the command line to start SMIT.
2. On the SMIT System Management screen, select **Devices** to bring up the Devices screen.
3. Select **Fixed Disk** to bring up the Fixed Disk screen.
4. Select **Change/Show Characteristics of a Disk** to go to the Disk screen (see Figure 3.1).
5. Select the desired device from the **Disk** menu to bring up the Change/Show Characteristics of a Disk screen.
6. Enter the desired queue depth (refer to Table 3.2), queue type (simple), and read/write time-out value (60), and then press **Enter** to complete the parameter changes.
7. Repeat steps (5) and (6) for each new device on the 9900 subsystem.
8. Verify that the parameters for all devices were successfully changed using the **#lsattr -E -1 'hdiskx'** command.

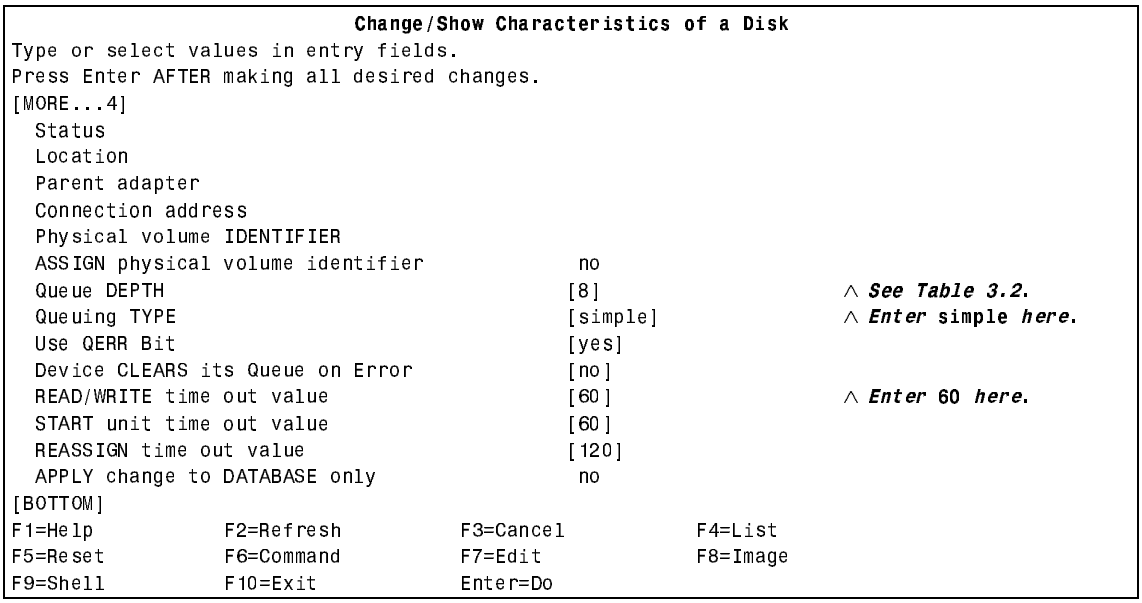

Figure 3.1 Changing the Device Parameters Using SMIT

Instructions for the AIX® Command Line

To change the device parameters from the AIX® command line:

1. At the AIX® command line prompt, enter **lsattr -E -l hdiskx** ('hdiskx' is the device file name, e.g., hdisk2) to display the parameters for the specified device.
2. Change the device parameters by entering:
   chdev -l hdiskx -a rw_timeout='60' -a q_type='simple' -a queue_depth='x'
   where x = desired queue depth within limits specified in Table 3.2.
3. Repeat steps (1) and (2) for each new device on the 9900.
4. Verify that the parameters for all devices were successfully changed using the **#lsattr -E -1 'hdiskx'** command.

3.2 Assigning the New Devices to Volume Groups and Setting the Partition Size

After the device parameters have been changed, you can assign the new SCSI disk devices to new or existing volume groups. The partition size is also set during this procedure. Table 3.3 specifies the partition sizes for standard LUs (OPEN-x devices). Table 3.4 specifies the partition sizes for LUSE devices (OPEN-x*n). Table 3.5 specifies the partition sizes for CVS LUSE devices (OPEN-x*n CVS).

Note: Do not assign the HMDE devices (e.g., 3390-3A/B/C) devices to volume groups. If you are configuring storage devices for databases that use a "raw" partition, do not assign those devices to volume groups.

To assign the SCSI disk devices to volume groups and set the partition size:

1. At the AIX® command line prompt, enter **smit** to start SMIT. This brings up the System Management screen. ***Note:*** If SMIT is not installed, please refer to the IBM® AIX® user guide for instructions on assigning new devices to volume groups using AIX® commands.
2. Select **System Storage Management (Physical & Logical Storage)** to bring up the System Storage Management screen.
3. Select **Logical Volume Manager** to bring up the Logical Volume Manager screen.
4. Select **Volume Groups** to bring up the Volume Group screen.
5. Select **Add a Volume Group** to bring up the Add a Volume Group screen.
6. The Add a Volume Group screen (see Figure 3.2) allows you to assign one or more devices (physical volumes) to a new or existing volume group and set the physical partition size. To assign one or more devices to a volume group and set the partition size:
   a) Place the cursor in the **VOLUME GROUP name** entry field. Enter the name of the new volume group (e.g., 9900vg0), or press **F4** and select an existing volume group. A volume group can contain multiple hdisk devices, depending on the application.
   b) Place the cursor in the **Physical partition SIZE in megabytes** field, and press the **F4** key. When the size menu appears, select the correct partition size for the device(s) as specified in Table 3.3 (standard LUs), Table 3.4 (LUSE), or Table 3.5 (CVS LUSE).
   c) Place the cursor in the **PHYSICAL VOLUME names** entry field. Enter the device file name(s) for the desired device(s) (e.g., hdisk1), or press **F4** and select the device file name(s) from the list.
   d) Place the cursor in the **Activate volume group AUTOMATICALLY** entry field, and enter **yes** to activate the volume group automatically at system restart.
   Note: If you are using HACMP, enter **no**.
7. After selecting the volume group, partition size, and physical volume(s) on the Add a Volume Group screen, press the **Enter** key.
8. When the confirmation screen opens, respond **Yes** to assign the specified device(s) to the specified volume group with the specified partition size.
9. The Command Status screen now opens. To ensure that the devices have been assigned to a volume group, wait for **OK** to appear on the Command Status line. To continue creating volume groups, press **F3** until the Add a Volume Group screen appears.
10. Repeat steps (2) through (9) until all new disk devices are assigned to a volume group.

```
                                  Add a Volume Group
Type or select values in entry fields.
Press Enter AFTER making all desired changes.
                                        [Entry Fields]
  VOLUME GROUP name                      [9900vg0]            ^ Enter volume group.
  Physical partition SIZE in megabytes      4                 ^ Enter partition size.
  PHYSICAL VOLUME names                  [hdisk1]             ^ Enter device file name(s).
  Activate volume group AUTOMATICALLY     yes                 ^ Enter no for HACMP.
    at system restart
  Volume Group MAJOR NUMBER              []
 *Create VG Concurrent Capable?
 *Auto-varyon in Concurrent Mode?
F1=Help          F2=Refresh         F3=Cancel          F4=List
F5=Reset         F6=Command         F7=Edit            F8=Image
F9=Shell         F10=Exit           Enter=Do
```

* These lines are added in AIX® version 4.3.

Figure 3.2 Assigning Devices to Volume Groups and Setting the Partition Size

Table 3.3 Partition Sizes for Standard LUs

Device Type	Partition Size
OPEN-K	2
OPEN-3	4
OPEN-8	8
OPEN-9	8
OPEN-E	16

Table 3.4 Partition Sizes for LUSE Devices

Device Type	LUSE Configuration	Partition Size (MB)
OPEN-K	OPEN-K	2
	OPEN-K*2	4
	OPEN-K*3-OPEN-K*4	8
	OPEN-K*5-OPEN-K*9	16
	OPEN-K*10-OPEN-K*18	32
	OPEN-K*19-OPEN-K*36	64
OPEN-3	OPEN-3	4
	OPEN-3*2-OPEN-3*3	8
	OPEN-3*4-OPEN-3*6	16
	OPEN-3*7-OPEN-3*13	32
	OPEN-3*14-OPEN-3*27	64
	OPEN-3*28-OPEN-3*36	128

Table 3.4 Partition Sizes for LUSE Devices

Device Type	LUSE Configuration	Partition Size (MB)
OPEN-8	OPEN-8	8
	OPEN-8*2	16
	OPEN-8*3-OPEN-8*4	32
	OPEN-8*5-OPEN-8*9	64
	OPEN-8*10-OPEN-8*18	128
	OPEN-8*19-OPEN-8*36	256
OPEN-9	OPEN-9	8
	OPEN-9*2	16
	OPEN-9*3-OPEN-9*4	32
	OPEN-9*5-OPEN-9*9	64
	OPEN-9*10-OPEN-9*18	128
	OPEN-9*19-OPEN-9*36	256
OPEN-E	OPEN-E	16
	OPEN-E*2	32
	OPEN-E*3,OPEN-E*4	64
	OPEN-E*5-OPEN-E*9	128
	OPEN-E*10-OPEN-E*18	256

Table 3.5 Partition Sizes for CVS LUSE Devices

Device Type	LU Size (MB)	Partition Size (MB)
OPEN-x*n CVS	35-1800	2
	1801-2300	4
	2301-7000	8
	7001-16200	16
	13201-32400	32
	32401-64800	64
	64801-126000	128
	126001 and higher	256

3.3 Creating, Mounting, and Verifying the File Systems

After you have assigned the SCSI disk devices to volume groups and set the partition sizes, you can create the file systems.

Note: Do not create file systems for the HMDE devices. If you are configuring storage devices for databases that use a "raw" partition, you will create a logical volume only.

3.3.1 Creating the File System

1. At the AIX® command line prompt, enter **smit** to start SMIT. This brings up the System Management screen. ***Note:*** If SMIT is not installed, please refer to the IBM® AIX® user guide for instructions on creating file systems using AIX® commands.
2. Select **System Storage Management (Physical & Logical Storage)** to bring up the System Storage screen.
3. Select **File Systems** to bring up the File System screen.
4. Select **Add/Change/Show/Delete File Systems** to bring up the Add/Change screen.
5. Select **Journaled File Systems** to bring up the Journaled File System screen.
6. Select **Add a Journaled File System** to bring up the Volume Group Name screen.
7. Move the cursor to the selected volume group, and press the **F7** key.
8. Select the desired value, and then press the **Enter** key to bring up the Add a Journaled File System screen (see Figure 3.3).
9. Place the cursor in the **SIZE of file system** field, and enter the desired file system size (see Table 3.6).
10. Place the cursor in the **Mount Point** field, and enter the desired mount point name (e.g., /9900_VG00). Please record the mount point name and file system size. You will be asked to input this information again.
11. Place the cursor in the **Mount AUTOMATICALLY** field. Enter **yes** to auto-mount the file systems. ***Note:*** If you are using HACMP, do not set the file systems to auto-mount.
12. Place the cursor in the **Number of bytes per inode** field, and enter the correct value for the selected device (see Tables 3.7-3.9).
13. Make sure that the file system size, mount point name, auto-mount options, and number of bytes per inode are correct, and press the Enter key to create the Journaled File System.
14. The Command Status screen now appears. To make sure that the Journaled File System has been created, wait for **OK** to appear on the Command Status line (see Figure 3.4).
15. Repeat steps (2) through (14) for each Journaled File System that you want to create. To continue creating Journaled File Systems press the **F3** key until you return to the Add a Journaled File System screen.
16. To exit SMIT, press the **F10** key.

```
                              Add a Journaled File System
Type or select values in entry fields.
Press Enter AFTER making all desired changes.
                                                [Entry Fields]
  Volume group name                               9900vg0
  SIZE of file system (in 512-byte blocks)      [4792320]          ^ See Table 3.6.
  MOUNT POINT                                   [/9900_VG00]       ^ Enter mount point name.
  Mount AUTOMATICALLY at system restart?         yes               ^ Enter no for HACMP.
  PERMISSIONS                                    read/write
  Mount OPTIONS                                  []
  Start Disk Accounting?                         no
  Fragment Size (bytes)                         4096
  Number of bytes per inode                     4096               ^ See Tables 3.7-3.9.
  Compression algorithm                          no
 *Allocation Group Size (Mbytes)
F1=Help           F2=Refresh          F3=Cancel           F4=List
F5=Reset          F6=Command          F7=Edit             F8=Image
F9=Shell          F10=Exit            Enter=Do
```

* This line is added in AIX® version 4.3.

Figure 3.3 Adding a Journaled File System Using SMIT

```
                                    COMMAND STATUS

Command : OK          stdout : yes          stderr : no

    Before command completion, additional instructions may appear below.

Based on the parameters chosen, the new /9900_VG00 JFS file system
is limited to a maximum size of 134217728 (512 byte blocks)
New Filesystems size is 4792320                          ^ 4792320 is displayed for OPEN-3.

F1=Help           F2=Refresh          F3=Cancel           F6=Command
F8=Image          F9=Shell            F10=Exit            /=Find
n=Find Next
```

Figure 3.4 Verifying Creation of Journaled File System

Table 3.6 Journaled File System Size

Device Type		Capacity (in 512-byte blocks)	Maximum File System Size *(See Note 1)* (in 512-byte blocks)
Standard LU	OPEN-K	3661920	*See Note 2.*
	OPEN-3	4806720	4792320
	OPEN-8	14351040	14319616
	OPEN-9	14423040	14401536
	OPEN-E	28452960	28409856
LUSE Device	OPEN-x*n	Refer to Table 1.1.	*See Note 2.*
CVS LUSE Device	OPEN-x*n CVS	Refer to Table 1.1.	*See Note 2.*

Note 1: When determining SIZE of File System at **Add a Journaled File System**, note that IBM® AIX® already uses an unspecified amount of disk space. You must determine the remaining size available for physical partitions.

Note 2: Calculate the file system size for these devices as follows:

1. Display the number of free physical partitions (FREE PPs) and physical partition size (PP SIZE) using the **lsvg** command (see Figure 3.5).
2. Calculate the maximum size of the file system as follows:
   (FREE PPs – 1) × (PP SIZE) × 2048

   Example for OPEN-3*20 LUSE device shown in Figure 3.5:
   The maximum file system size is: (733 – 1) × (64) × 2048 = 95944704

```
# lsvg 9900vg0
VOLUMEGROUP:    9900vg0              VG IDENTIFIER:   0083665612e98521
VG STATE:       active               PP SIZE:         64 megabyte(s)
VG PERMISSION:  read/write           TOTAL PPs:       733 (46912 megabytes)
MAX LVs:        256                  FREE PPs:        733 (46912 megabytes)
LVs:            0                    USED PPS:        0 (0 megabytes)
OPEN LVs:       0                    QUORUM:          2
TOTAL PVs:      1                    VG DESCRIPTORS:  2
STALE PVs;      0                    STALE PPs        0
ACTIVE PVs:     1                    AUTO ON:         yes
Concurrent:     Non-Capable          Auto-Concurrent: Disabled
VG Mode:        Non-Concurrent
```

Figure 3.5 Determining the Maximum File System Size

Table 3.7 Number of bytes per inode for LUSE devices

Device Type	LU Product Name	Number of Bytes per inode
OPEN-K	OPEN-K, OPEN-K*2-OPEN-K*36	4096
OPEN-3	OPEN-3, OPEN-3*2-OPEN-3*28	4096
	OPEN-3*29-OPEN-3*36	8192
OPEN-8	OPEN-8, OPEN-8*2-OPEN-8*9	4096
	OPEN-8*10-OPEN-8*18	8192
	OPEN-8*19-OPEN-8*36	16384
OPEN-9	OPEN-9, OPEN-9*2-OPEN-9*9	4096
	OPEN-9*10-OPEN-9*18	8192
	OPEN-9*19-OPEN-9*36	16384
OPEN-E	OPEN-E, OPEN-E*2-OPEN-E*4	4096
	OPEN-E*5-OPEN-E*9	8192
	OPEN-E*10-OPEN-E*18	16384

Table 3.8 Number of bytes per inode for CVS

Device Type	LU Product Name	Number of Bytes per inode
OPEN-x CVS	OPEN-3 CVS, OPEN-8 CVS, OPEN-9 CVS, OPEN-E CVS, OPEN-K CVS	4096

Table 3.9 Number of bytes per inode for CVS LUSE

Device Type	LU size in Megabytes	Number of Bytes per inode
OPEN-x*n CVS	35-64800	4096
	64801-126000	8192
	126001 and higher	16384

3.3.2 Mounting and Verifying the File Systems

After you have created the Journaled File Systems, you can mount the file systems and verify that the file systems were created correctly and are functioning properly.

To mount and verify the file systems:

1. At the AIX® command line prompt, enter: **mount <mount_point_name>** (e.g., mount /9900_VG00).
2. Repeat step (1) for each new file system. Refer to the list of mount point names you recorded in section 3.4.
3. Verify the size of the file systems you have created using the **df** command.

   Note: The file system capacity is listed in 512-byte blocks by default. To list capacity in 1024-byte blocks, use the **df -k** command.
4. Verify that the new devices and file systems are fully operational by performing some basic operations (e.g., file creation, copying, deletion) on each device (see Figure 3.6).
5. At the next system restart, verify that the file systems have successfully auto-mounted by using the **mount** or **df** command to display all mounted file systems (see Figure 3.7). Any file systems that were not auto-mounted can be set to auto-mount using the SMIT Change a Journaled File System screen. ***Note:*** If you are using HACMP, do not set the file systems to auto-mount.

```
# cd /9900_VG00                                                    ^ Go to mount point.
# cp /smit.log /9900_VG00/smit.log.back1                           ^ Copy file.
# ls -l 9900_VG00                                                  ^ Verify file copy.
-rw-rw-rw-   1   root  system      375982 Nov 30 17:25 smit.log.back1
# cp smit.log.back1 smit.log.back2                                 ^ Copy file again.
# ls -l                                                            ^ Verify copy again.
-rw-rw-rw-   1   root  system      375982 Nov 30 17:25 smit.log.back1
-rw-rw-rw-   1   root  system      375982 Nov 30 17:28 smit.log.back2
# rm smit.log.back1                                                ^ Remove test file.
# rm smit.log.back2                                                ^ Remove test file.
```

Figure 3.6 Verifying the Auto-Mounted File Systems

```
# df                                                                  ^ List mounted file systems.
File system  512-blocks      free  %Used  Iused  %Iused  Mounted on
/dev/hd4           8192      3176    61%    652     31%  /
/dev/hd2        1024000    551448    46%   6997      5%  /usr
/dev/hd9var        8192      5512    32%     66      6%  /var
/dev/hd3          24576     11608    52%     38      0%  /tmp
/dev/hd1           8192      7840     4%     17      1%  /home
/dev/lv00       4792320   4602128     4%     16      1%  /9900_VG00        ^ OPEN-3 device.
/dev/lv01       4792320   4602128     4%     16      1%  /9900_VG01        ^ OPEN-3 device.
/dev/lv02      14401536  13949392     4%     16      1%  /9900_VG02        ^ OPEN-9 device.
```

Figure 3.7 Final File System Verification

Chapter 4 Middleware and SNMP Configuration

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

- HACMP for host failover (see section 4.1).
- Hitachi Path Manager for path failover (see section 4.2).

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 Failover

The 9900 subsystem supports the HACMP host failover product for the IBM® AIX® operating system. The HACMP products are maintained by Availant™.

The user must make sure to configure the HACMP software and any other middleware products (e.g., Tuxedo) as needed to recognize and operate with the newly attached 9900 devices. For assistance with HACMP operations, please refer to the HACMP user documentation, or contact Availant™ technical support. For assistance with specific configuration issues related to the 9900 subsystem, please contact the Hitachi Data Systems Support Center (see section 5.2).

Note: HACMP does not provide a complete disaster recovery or backup solution, and is not a replacement for standard disaster recovery planning and backup/recovery.

4.2 Path Failover

The Hitachi Lightning 9900™ subsystem supports the Hitachi Path Manager (HPM) product for the IBM® AIX® operating system. For further information on HPM, please see *Hitachi Path Manager for AIX®* (BO-99DD871).

4.3 SNMP Remote Subsystem Management

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

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

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

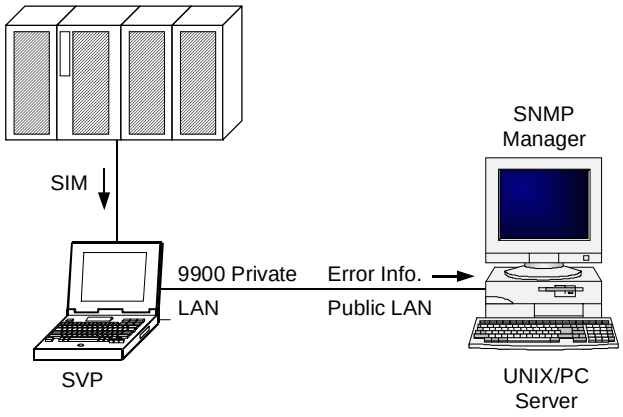

Figure 4.1 9900 SNMP Environment

Chapter 5 Troubleshooting

5.1 Troubleshooting

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 IBM® AIX® 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. The LUNs for each port must start at 0 and continue sequentially without skipping any numbers. Run **cfgmgr** to recheck the fibre channel for new devices. Make sure that LUSE devices are not intermixed with normal LUs or with HMDE devices on the same fibre channel port. Make sure that the LUNs are configured properly for each TID.
The file system is not mounted after rebooting.	Make sure the system was restarted properly. Make sure that the values listed under Journaled File System are correct.
System hangs, or devices are declared and then the system hangs.	Make sure 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 Calling the Hitachi Data Systems Technical Support Center

If you need to call the Hitachi Data Systems Technical Support Center, make sure to provide as much information about the problem as possible. Include the circumstances surrounding the error or failure, the exact content of any messages displayed on the remote console PC, and the severity levels and reference codes of the R-SIMs on the R-SIM panel. The worldwide Hitachi Data Systems Technical 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

CDE	common desktop environment
ESCON®	Enterprise System Connection (IBM® trademark for optical channels)
HACMP	High Availability Cluster Multi-Processing
HBA	host bus adapter
HMDE	Hitachi Multiplatform Data Exchange
HRC	Hitachi Remote Copy
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
SMIT®	System Management Interface Tool
SNMP	simple network management protocol
SVP	service processor
TCP/IP	transmission control protocol/internet protocol
TID	target ID
VOLSER	volume serial number

Appendix B SCSI TID Maps for Fibre-Channel Adapters

When an arbitrated loop (AL) is established or re-established, the port addresses are assigned automatically to prevent duplicate TIDs. When using 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 AIX® operating system "maps" a TID to each AL-PA.

Tables B.1 and B.2 identify the fixed mappings between the bus/TID/LUN addresses assigned by AIX® and the FC native addresses (AL_PA/SEL_ID) for FC adapters. There are two potential mappings depending on the value of the ScanDown registry parameter:

- For ScanDown = 0 (default) see Table B.1.
- For ScanDown = 1 see Table B.2.

Note: When 9900 devices and other types of devices are connected in the same arbitrated loop, the mappings defined in Tables B.1 and B.2 cannot be guaranteed.

Note: The Emulex® driver emulates six fibre buses per adapter to map all 126 possible AL-PAs to target IDs. The first bus (bus 0) is a dummy bus.

Table B.1 SCSI TID Map for Emulex® FC Adapter (ScanDown=0)

Bus #	TID	LUN	AL_PA	SEL_ID
0	0-31	0-7	NONE	NONE
1	0	0-7	0x01	0x7D
	1	0-7	0x02	0x7C
	2	0-7	0x04	0x7B
	3	0-7	0x08	0x7A
	4	0-7	0x0F	0x79
	5	0-7	0x10	0x78
	6	0-7	0x17	0x77
	7	0-7	0x18	0x76
	8	0-7	0x1B	0x75
	9	0-7	0x1D	0x74
	10	0-7	0x1E	0x73
	11	0-7	0x1F	0x72
	12	0-7	0x23	0x71
	13	0-7	0x25	0x70
	14	0-7	0x26	0x6F
	15	0-7	0x27	0x6E
	16	0-7	0x29	0x6D
	17	0-7	0x2A	0x6C
	18	0-7	0x2B	0x6B
	19	0-7	0x2C	0x6A
	20	0-7	0x2D	0x69
	21	0-7	0x2E	0x68
	22	0-7	0x31	0x67
	23	0-7	0x32	0x66
	24	0-7	0x33	0x65
	25	0-7	0x34	0x64
	26	0-7	0x35	0x63
	27	0-7	0x36	0x62
	28	0-7	0x39	0x61
	29	0-7	0x3A	0x60
	30	0-7	0x3C	0x5F
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
2	0	0-7	0x43	0x5E
	1	0-7	0x45	0x5D
	2	0-7	0x46	0x5C
	3	0-7	0x47	0x5B
	4	0-7	0x49	0x5A
	5	0-7	0x4A	0x59
	6	0-7	0x4B	0x58
	7	0-7	0x4C	0x57
	8	0-7	0x4D	0x56
	9	0-7	0x4E	0x55
	10	0-7	0x51	0x54
	11	0-7	0x52	0x53
	12	0-7	0x53	0x52
	13	0-7	0x54	0x51
	14	0-7	0x55	0x50
	15	0-7	0x56	0x4F
	16	0-7	0x59	0x4E
	17	0-7	0x5A	0x4D
	18	0-7	0x5C	0x4C
	19	0-7	0x63	0x4B
	20	0-7	0x65	0x4A
	21	0-7	0x66	0x49
	22	0-7	0x67	0x48
	23	0-7	0x69	0x47
	24	0-7	0x6A	0x46
	25	0-7	0x6B	0x45
	26	0-7	0x6C	0x44
	27	0-7	0x6D	0x43
	28	0-7	0x6E	0x42
	29	0-7	0x71	0x41
	30	0-7	0x72	0x40
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
3	0	0-7	0x73	0x3F
	1	0-7	0x74	0x3E
	2	0-7	0x75	0x3D
	3	0-7	0x76	0x3C
	4	0-7	0x79	0x3B
	5	0-7	0x7A	0x3A
	6	0-7	0x7C	0x39
	7	0-7	0x80	0x38
	8	0-7	0x81	0x37
	9	0-7	0x82	0x36
	10	0-7	0x84	0x35
	11	0-7	0x88	0x34
	12	0-7	0x8F	0x33
	13	0-7	0x90	0x32
	14	0-7	0x97	0x31
	15	0-7	0x98	0x30
	16	0-7	0x9B	0x2F
	17	0-7	0x9D	0x2E
	18	0-7	0x9E	0x2D
	19	0-7	0x9F	0x2C
	20	0-7	0xA3	0x2B
	21	0-7	0xA5	0x2A
	22	0-7	0xA6	0x29
	23	0-7	0xA7	0x28
	24	0-7	0xA9	0x27
	25	0-7	0xAA	0x26
	26	0-7	0xAB	0x25
	27	0-7	0xAC	0x24
	28	0-7	0xAD	0x23
	29	0-7	0xAE	0x22
	30	0-7	0xB1	0x21
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
4	0	0-7	0xB2	0x20
	1	0-7	0xB3	0x1F
	2	0-7	0xB4	0x1E
	3	0-7	0xB5	0x1D
	4	0-7	0xB6	0x1C
	5	0-7	0xB9	0x1B
	6	0-7	0xBA	0x1A
	7	0-7	0xBC	0x19
	8	0-7	0xC3	0x18
	9	0-7	0xC5	0x17
	10	0-7	0xC6	0x16
	11	0-7	0xC7	0x15
	12	0-7	0xC9	0x14
	13	0-7	0xCA	0x13
	14	0-7	0xCB	0x12
	15	0-7	0xCC	0x11
	16	0-7	0xCD	0x10
	17	0-7	0xCE	0x0F
	18	0-7	0xD1	0x0E
	19	0-7	0xD2	0x0D
	20	0-7	0xD3	0x0C
	21	0-7	0xD4	0x0B
	22	0-7	0xD5	0x0A
	23	0-7	0xD6	0x09
	24	0-7	0xD9	0x08
	25	0-7	0xDA	0x07
	26	0-7	0xDC	0x06
	27	0-7	0xE0	0x05
	28	0-7	0xE1	0x04
	29	0-7	0xE2	0x03
	30	0-7	0xE4	0x02
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
5	0	0-7	0xE8	0x01
	1	0-7	0xEF	0x00
	2	0-7	NONE	NONE
	3	0-7	NONE	NONE
	4	0-7	NONE	NONE
	5	0-7	NONE	NONE
	6	0-7	NONE	NONE
	7	0-7	NONE	NONE
	8	0-7	NONE	NONE
	9	0-7	NONE	NONE
	10	0-7	NONE	NONE
	11	0-7	NONE	NONE
	12	0-7	NONE	NONE
	13	0-7	NONE	NONE
	14	0-7	NONE	NONE
	15	0-7	NONE	NONE
	16	0-7	NONE	NONE
	17	0-7	NONE	NONE
	18	0-7	NONE	NONE
	19	0-7	NONE	NONE
	20	0-7	NONE	NONE
	21	0-7	NONE	NONE
	22	0-7	NONE	NONE
	23	0-7	NONE	NONE
	24	0-7	NONE	NONE
	25	0-7	NONE	NONE
	26	0-7	NONE	NONE
	27	0-7	NONE	NONE
	28	0-7	NONE	NONE
	29	0-7	NONE	NONE
	30	0-7	NONE	NONE
	31	0-7	NONE	NONE

Table B.2 SCSI TID Map for Emulex® FC Adapter (ScanDown=1)

Bus #	TID	LUN	AL_PA	SEL_ID
0	0-31	0-7	NONE	NONE
1	0	0-7	0xEF	0x00
	1	0-7	0xE8	0x01
	2	0-7	0xE4	0x02
	3	0-7	0xE2	0x03
	4	0-7	0xE1	0x04
	5	0-7	0xE0	0x05
	6	0-7	0xDC	0x06
	7	0-7	0xDA	0x07
	8	0-7	0xD9	0x08
	9	0-7	0xD6	0x09
	10	0-7	0xD5	0x0A
	11	0-7	0xD4	0x0B
	12	0-7	0xD3	0x0C
	13	0-7	0xD2	0x0D
	14	0-7	0xD1	0x0E
	15	0-7	0xCE	0x0F
	16	0-7	0xCD	0x10
	17	0-7	0xCC	0x11
	18	0-7	0xCB	0x12
	19	0-7	0xCA	0x13
	20	0-7	0xC9	0x14
	21	0-7	0xC7	0x15
	22	0-7	0xC6	0x16
	23	0-7	0xC5	0x17
	24	0-7	0xC3	0x18
	25	0-7	0xBC	0x19
	26	0-7	0xBA	0x1A
	27	0-7	0xB9	0x1B
	28	0-7	0xB6	0x1C
	29	0-7	0xB5	0x1D
	30	0-7	0xB4	0x1E
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
2	0	0-7	0xB3	0x1F
	1	0-7	0xB2	0x20
	2	0-7	0xB1	0x21
	3	0-7	0xAE	0x22
	4	0-7	0xAD	0x23
	5	0-7	0xAC	0x24
	6	0-7	0xAB	0x25
	7	0-7	0xAA	0x26
	8	0-7	0xA9	0x27
	9	0-7	0xA7	0x28
	10	0-7	0xA6	0x29
	11	0-7	0xA5	0x2A
	12	0-7	0xA3	0x2B
	13	0-7	0x9F	0x2C
	14	0-7	0x9E	0x2D
	15	0-7	0x9D	0x2E
	16	0-7	0x9B	0x2F
	17	0-7	0x98	0x30
	18	0-7	0x97	0x31
	19	0-7	0x90	0x32
	20	0-7	0x8F	0x33
	21	0-7	0x88	0x34
	22	0-7	0x84	0x35
	23	0-7	0x82	0x36
	24	0-7	0x81	0x37
	25	0-7	0x80	0x38
	26	0-7	0x7C	0x39
	27	0-7	0x7A	0x3A
	28	0-7	0x79	0x3B
	29	0-7	0x76	0x3C
	30	0-7	0x75	0x3D
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
3	0	0-7	0x74	0x3E
	1	0-7	0x73	0x3F
	2	0-7	0x72	0x40
	3	0-7	0x71	0x41
	4	0-7	0x6E	0x42
	5	0-7	0x6D	0x43
	6	0-7	0x6C	0x44
	7	0-7	0x6B	0x45
	8	0-7	0x6A	0x46
	9	0-7	0x69	0x47
	10	0-7	0x67	0x48
	11	0-7	0x66	0x49
	12	0-7	0x65	0x4A
	13	0-7	0x63	0x4B
	14	0-7	0x5C	0x4C
	15	0-7	0x5A	0x4D
	16	0-7	0x59	0x4E
	17	0-7	0x56	0x4F
	18	0-7	0x55	0x50
	19	0-7	0x54	0x51
	20	0-7	0x53	0x52
	21	0-7	0x52	0x53
	22	0-7	0x51	0x54
	23	0-7	0x4E	0x55
	24	0-7	0x4D	0x56
	25	0-7	0x4C	0x57
	26	0-7	0x4B	0x58
	27	0-7	0x4A	0x59
	28	0-7	0x49	0x5A
	29	0-7	0x47	0x5B
	30	0-7	0x46	0x5C
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
4	0	0-7	0x45	0x5D
	1	0-7	0x43	0x5E
	2	0-7	0x3C	0x5F
	3	0-7	0x3A	0x60
	4	0-7	0x39	0x61
	5	0-7	0x36	0x62
	6	0-7	0x35	0x63
	7	0-7	0x34	0x64
	8	0-7	0x33	0x65
	9	0-7	0x32	0x66
	10	0-7	0x31	0x67
	11	0-7	0x2E	0x68
	12	0-7	0x2D	0x69
	13	0-7	0x2C	0x6A
	14	0-7	0x2B	0x6B
	15	0-7	0x2A	0x6C
	16	0-7	0x29	0x6D
	17	0-7	0x27	0x6E
	18	0-7	0x26	0x6F
	19	0-7	0x25	0x70
	20	0-7	0x23	0x71
	21	0-7	0x1F	0x72
	22	0-7	0x1E	0x73
	23	0-7	0x1D	0x74
	24	0-7	0x1B	0x75
	25	0-7	0x18	0x76
	26	0-7	0x17	0x77
	27	0-7	0x10	0x78
	28	0-7	0x0F	0x79
	29	0-7	0x08	0x7A
	30	0-7	0x04	0x7B
	31	0-7	NONE	NONE

Bus #	TID	LUN	AL_PA	SEL_ID
5	0	0-7	0x02	0x7C
	1	0-7	0x01	0x7D
	2	0-7	NONE	NONE
	3	0-7	NONE	NONE
	4	0-7	NONE	NONE
	5	0-7	NONE	NONE
	6	0-7	NONE	NONE
	7	0-7	NONE	NONE
	8	0-7	NONE	NONE
	9	0-7	NONE	NONE
	10	0-7	NONE	NONE
	11	0-7	NONE	NONE
	12	0-7	NONE	NONE
	13	0-7	NONE	NONE
	14	0-7	NONE	NONE
	15	0-7	NONE	NONE
	16	0-7	NONE	NONE
	17	0-7	NONE	NONE
	18	0-7	NONE	NONE
	19	0-7	NONE	NONE
	20	0-7	NONE	NONE
	21	0-7	NONE	NONE
	22	0-7	NONE	NONE
	23	0-7	NONE	NONE
	24	0-7	NONE	NONE
	25	0-7	NONE	NONE
	26	0-7	NONE	NONE
	27	0-7	NONE	NONE
	28	0-7	NONE	NONE
	29	0-7	NONE	NONE
	30	0-7	NONE	NONE
	31	0-7	NONE	NONE

Appendix C Online Installation and Deinstallation of Devices

After initial installation and configuration of the 9900 subsystem, additional devices can be installed or de-installed online without having to restart the 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 3.1).

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

1. Log in to the IBM® system as **root**.
2. At the AIX® command line prompt, enter **smit** to start SMIT®. This brings up the System Management screen. ***Note:*** If SMIT® is not installed, please refer to the IBM® AIX® user guide for instructions on assigning new devices to volume groups using AIX® commands.
3. Select **Devices** to bring up the Devices screen.
4. Select **Install/Configure Devices Added After IPL** to bring up the Install/Configure Devices Added After IPL screen.
5. Select **INPUT device/directory** for software, and then press the **Enter** key. The AIX® system now scans the buses for new devices.
6. Verify new device recognition using the **lsdev -C -s scsi** command as described in section 2.5. Make sure to record the device file names for the new devices.
7. Configure the new devices for AIX® operations as described in Chapters 3 and 4.

Index

9

A

C

D

F

H

I

N

O

P

S

T

V