HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Lightning 9900™

Red Hat® Linux® 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-90RD028-P	June, 2000	Preliminary Release
MK-90RD028-0	November, 2000	Initial Release, supersedes and replaces MK-90RD028-P

Source Document Revision Level

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

- *Hitachi DKC310 Disk Subsystem SCSI/Fibre Installation Manual, Linux,* revision 1.1.
- *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide,* MK-90RD006-1.
- ECN TW-90RD042 (Technical Writing ECN).
- ECN TW-90RD043 (Technical Writing ECN).

Preface

The *Hitachi Freedom Storage™ Lightning 9900™ Red Hat® Linux® Configuration Guide* describes and provides instructions for installing and configuring the devices on the Lightning 9900™ array subsystem for operation with the Red Hat® Linux® 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 Red Hat® Linux® operating system, including commands, utilities and file systems.

Note: This document will use the term "9900" to refer to any 9900 subsystem configured for UNIX/PC server operations, including the all-open and multiplatform models.

For further information on the Lightning 9900™ array subsystem, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide* (MK-90RD008), or contact your Hitachi Data Systems account team. The Hitachi Data Systems worldwide web site (http://www.hds.com) also provides information on the Lightning 9900™ subsystem and its features and options.

For further information on Red Hat® Linux®, please consult the Red Hat® user documentation, or contact Red Hat® technical support.

Contents

Chapter 1 Overview of Red Hat® Linux® Configuration 1

1.1 Red Hat® Linux® Configuration 1
1.2 The Lightning 9900™ Array Subsystem 1
1.3 Device Types and Configuration Procedures 2

Chapter 2 Preparing for New Device Configuration 7

2.1 Configuration Requirements 7
2.2 Installing the 9900 Subsystem 8
2.3 Preparing to Connect the 9900 Subsystem 9
2.3.1 Setting the Host Mode for the 9900 Ports 9
2.3.2 Configuring the 9900 Fibre-Channel Ports 10
2.4 Connecting the 9900 Subsystem 13
2.5 Configuring the Host Fibre-Channel Adapters 14
2.6 Verifying Device Recognition 16

Chapter 3 Configuring the New Devices 19

3.1 Set the Number of Logical Units 19
3.2 Partitioning Devices 20
3.3 Creating, Mounting and Verifying the File Systems 23
3.3.1 Creating the File Systems 23
3.3.2 Create the Mount Directories To create the mount directories, enter the following: 23
3.3.3 Create the Mount Table 23
3.3.4 Verify the File System 24
3.3.5 Set the Auto-Mount Parameters 24

Chapter 4 SNMP Remote Subsystem Management 25

Chapter 5 Troubleshooting 27

5.1 Troubleshooting 27
5.2 Calling the Hitachi Data Systems Technical Support Center 28

Appendix A Acronyms and Abbreviations 29

List of Figures

Figure 2.1 Set Host Mode Panel 9
Figure 2.2 Fibre Parameter Panel 11
Figure 2.3 Verifying Device Recognition 16

Figure 4.1 9900 SNMP Environment 25

List of Tables

Table 1.1 9900 Device Specifications for Red Hat® Linux® 4
Table 1.2 Volume Usage for Device Types 5

Table 2.1 Fibre Topology Settings on the 9900 Remote Console PC 12
Table 2.2 Available AL-PA Values 12
Table 2.3 Requirements for the 9900 Devices 14
Table 2.4 Sample SCSI Path Worksheet 17

Table 3.1 Partition Sizes for Standard LUs 20
Table 3.2 Partition Sizes for LUSE Devices 21
Table 3.3 Partition Sizes for CVS LUSE Devices 21

Table 5.1 Troubleshooting 27

Chapter 1 Overview of Red Hat® Linux® Configuration

1.1 Red Hat® Linux® Configuration

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

Configuration of the 9900 SCSI disk devices for Red Hat® Linux® includes:

- Setting the number of LUs (see section 3.1).
- Partitioning the devices (see section 3.2).
- Creating the file systems (see section 3.3.1).
- Creating the mount directories (see section 3.3.2).
- Creating the mount table (see section 3.3.3).
- Certifying the file system (see section 3.3.4).
- Setting the auto mount procedures (see section 3.3.5).

For further information on Red Hat® Linux® please refer to the Linux® user documentation, or contact Red Hat® technical support.

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 (LDEVs) to be installed and configured for operation with the Red Hat® Linux® operating system. **Error! Reference source not found.** lists the device specifications for the 9900 devices. Table 1.2 shows the volume usage (i.e., file system or raw device) for the 9900 devices.

OPEN-x Devices. The OPEN-x logical units (LUs) (e.g., OPEN-3, OPEN-9) are disk devices of predefined sizes. The 9900 subsystem currently supports OPEN-3, OPEN-8, OPEN-9, and OPEN-K 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 that can be from 2 to 36 times larger than standard OPEN-x LUs. The Logical Unit Size Expansion (LUSE) feature of the 9900 subsystem enables you to configure these custom-size devices. LUSE devices are designated as OPEN-x*n, where x is the LU type (e.g., OPEN-9*n) and $2 \leq n \leq 36$. For example, a LUSE device created from ten OPEN-3 LUs would be designated as an OPEN-3*10 disk device. This capability enables the server host to combine logical devices and access the data stored on the 9900 subsystem using fewer LU numbers (LUNs). For further information on the LUSE feature, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006).

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

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

Configuration of the 9900 disk devices for Red Hat® Linux® operations includes:

- Setting the host mode and configuring the 9900 fibre-channel ports (see section 2.3),
- Connecting to the 9900 subsystem (see section 2.4),
- Configuring the host fibre-channel adapters (see section 2.5),
- Verifying device recognition (see section 2.6),
- Setting the number of logical units (see section 3.1),
- Partitioning the devices (see section 3.2), and
- Creating, mounting and verifying the file systems (see section 3.3).

Table 1.1 9900 Device Specifications for Red Hat® Linux®

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	Disk	HITACHI	OPEN-3	4806720	512	3338	15	96	2347
OPEN-9	Disk	HITACHI	OPEN-9	14423040	512	10016	15	96	7042
OPEN-K	Disk	HITACHI	OPEN-K	3661920	512	2543	15	96	1788
OPEN-8	Disk	HITACHI	OPEN-8	14351040	512	9966	15	96	7007
OPEN-3*n	Disk	HITACHI	OPEN-3*n	4806720*n	512	3338*n	15	96	2347*n
OPEN-9*n	Disk	HITACHI	OPEN-9*n	14423040*n	512	10016*n	15	96	7042*n
OPEN-K*n	Disk	HITACHI	OPEN-K*n	3661920*n	512	2543*n	15	96	1788*n
OPEN-8*n	Disk	HITACHI	OPEN-8*n	14351040*n	512	9966*n	15	96	7007*n
OPEN-3 CVS	Disk	HITACHI	OPEN-3-CVS	*Note 3*	512	*Note 4*	15	96	*Note 5*
OPEN-9 CVS	Disk	HITACHI	OPEN-9-CVS	*Note 3*	512	*Note 4*	15	96	*Note 5*
OPEN-K CVS	Disk	HITACHI	OPEN-K-CVS	*Note 3*	512	*Note 4*	15	96	*Note 5*
OPEN-8 CVS	Disk	HITACHI	OPEN-8-CVS	*Note 3*	512	*Note 4*	15	96	*Note 5*
OPEN-3*n CVS	Disk	HITACHI	OPEN-3*n-CVS	*Note 3*	512	*Note 4*	15	96	*Note 5*
OPEN-9*n CVS	Disk	HITACHI	OPEN-9*n-CVS	*Note 3*	512	*Note 4*	15	96	*Note 5*
OPEN-K*n CVS	Disk	HITACHI	OPEN-K*n-CVS	*Note 3*	512	*Note 4*	15	96	*Note 5*
OPEN-8*n CVS	Disk	HITACHI	OPEN-8*n-CVS	*Note 3*	512	*Note 4*	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.

Table 1.2 Volume Usage for Device Types

Category	Device Type	Volume Usage
Disk	OPEN-x, OPEN-x CVS, OPEN-x*n LUSE, OPEN-x*n CVS LUSE	File System

Note 2: 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 3: 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 4: The number of data cylinders for a CVS volume (including CVS LUSE volumes) is calculated as follows:

(↑…↑ means that the value should be rounded up to the next integer)

- The number of data cylinders for an OPEN-* CVS volume =
  # of cylinders = ↑ (capacity (MB) specified on the 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 3390-3A/C or 3380-KA/C CVS volume =
  # of cylinders = (number of cylinders specified on the 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 the Remote Console PC) + 7

Note 5: 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

Before you can configure the 9900, you must take the following steps:

- Installing the 9900 subsystem (see section 2.2),
- Setting the host mode and configuring the fibre-channel ports (see section 2.3),
- Connecting the 9900 subsystem (see section 2.4),
- Configuring the host fibre-channel adapters (see section 2.5),
- Verifying device recognition (see section 2.6).

2.1 Configuration Requirements

The requirements for 9900 Red Hat® Linux® configuration are:

- 9900 subsystem.
  - The 9900 Remote Console PC and LUN Manager software are used to configure the fibre-channel (FC) ports. 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). If the Remote LUN Manager software is not installed, the Hitachi Data Systems representative can configure the fibre-channel ports for you using the SVP of the subsystem. Please contact your Hitachi Data Systems account team for further information on LUN fibre-channel configuration services.
  - ***Note:*** The availability of 9900 features and devices (e.g., CVS, 3390-3C, OPEN-8) depends on the level of microcode installed on the 9900 subsystem.
- 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 supports full-speed (100 MB/s) fibre-channel interface, including short-wave non-OFC (open fibre control) optical interface and multimode optical cables with SC connectors.
  - For information on supported FC adapters, optical cables, hubs, and fabric switches, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center (see section 0).

  Note: Do not connect any QLogic® 2200F fibre-channel driver (version 2-11) connector to the 9900 subsystem.
- Red Hat® Linux® operating system, version 6.0 and 6.2.

  Note: Hitachi Data Systems plans to support future releases of Red Hat® Linux®. This document will be updated as needed to cover version-specific information.
- Superuser (root) login access to the host system.

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 this activity, which includes:
   - Assembling all hardware and cabling.
   - Loading the latest microcode and SVP updates for full fibre support.
   - Installing and formatting the logical devices (LDEVs) using the SVP. (This requires complete LDEV configuration information from the user, including the desired number of OPEN-x, LUSE, and CVS 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 install/de-install fibre-channel cabling while Red Hat® Linux® is active. This can cause the system to hang. Always confirm that Red Hat® Linux® is shut down before installing/de-installing fibre cabling.
   - ***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 0). 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 fibre-channel ports for the 9900 devices. For instructions on installing the Remote Console PC for the 9900, please refer to *Hitachi Freedom Storage™ Lightning 9900™ Remote Console User's Guide* (MK-90RD003). 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). If the Remote LUN Manager software is not installed, the Hitachi Data Systems representative can configure the fibre-channel ports for you using the SVP of the subsystem. Please contact your Hitachi Data Systems account team for further information on LUN fibre-channel configuration services.

2.3 Preparing to Connect the 9900 Subsystem

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 Linux® operations is **00**. Use the LUN Manager remote console software to ensure that the host mode for each fibre port connected to the Linux® system is 00.

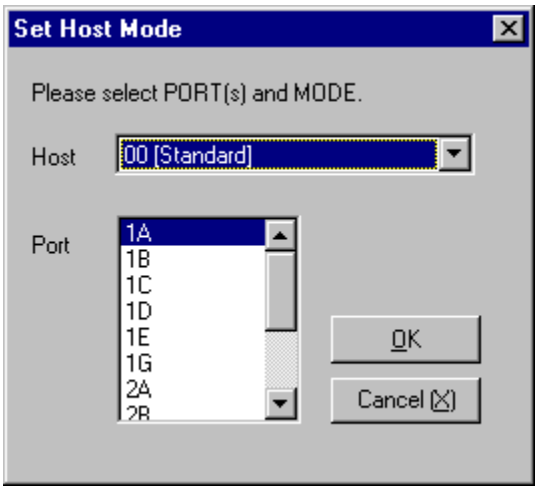

Figure 2.1 Set Host Mode Panel

2.3.2 Configuring the 9900 Fibre-Channel Ports

You need to configure the 9900 FC ports to define the fibre topology parameters (e.g., arbitrated loop) and port addresses. The LUN Manager remote console software enables you to configure the 9900 FC ports. For additional instructions, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006). If the Remote LUN Manager software is not installed, the Hitachi Data Systems representative can configure the fibre-channel ports for you using the SVP of the subsystem. Please contact your Hitachi Data Systems account team for further information on LUN fibre-channel configuration services.

Fibre topology. . **Error! Reference source not found.** shows an example of the Fibre Parameter panel (part of the LUN Manager software), and explains the settings on this panel. You will select the appropriate topology settings for each 9900 FC port based on the device to which the port is connected. **Error! Reference source not found.** shows the fibre topology settings for the 9900. Determine which topology parameters are supported by the device, and set your topology accordingly.

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). **Error! Reference source not found.** 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 SCSI software driver of the platform host adapter translates the AL-PA value assigned to the 9900 port into a SCSI TID.

Note: The AL-PAs should be unique for each device on the loop to avoid conflicts.

Topology parameters. Use the LUN Manager software to define the topology parameters (refer to **Error! Reference source not found.**and **Error! Reference source not found.**) and port ID for each fibre-channel port. For additional instructions, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006). If the Remote LUN Manager software is not installed, the Hitachi Data Systems representative can configure the fibre-channel ports for you using the SVP of the subsystem. Please contact your Hitachi Data Systems account team for further information on LUN fibre-channel configuration services. When you are finished configuring the fibre channel ports, you can exit the LUN Manager software.

Note: In a Linux environment, the 9900 subsystem supports up to 16 LUs per fibre-channel port.

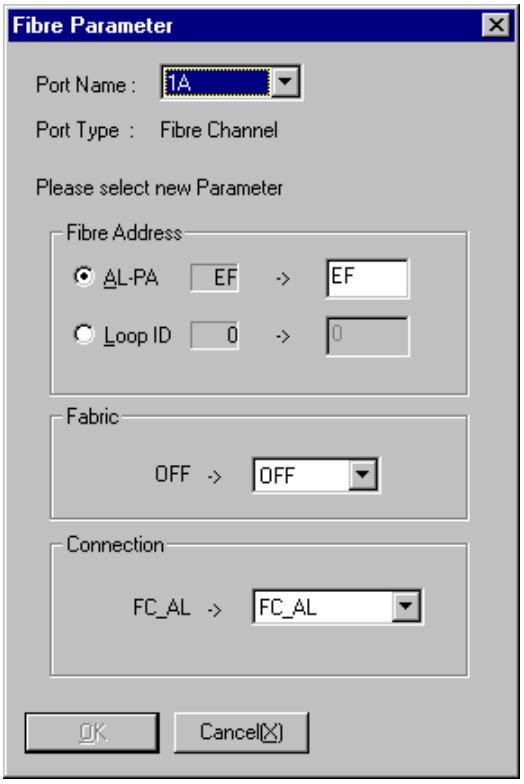

Figure 2.2 Fibre Parameter Panel (from the LUN Manager Software; the settings can vary)

Table 2.1 Fibre Topology Settings on the 9900 Remote Console PC

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

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 to the host system(s).

To connect the 9900 subsystem to the host system:

1. **Verify subsystem installation**. The Hitachi Data Systems representative verifies that the status of the fibre ports and LDEVs is normal.
2. **Shut down and power off the host system**. The user should perform this activity. You must shut down and power off the host system before connecting the 9900:
   a) Shut down the host system.
   b) When shutdown is complete, power off the Red Hat® Linux® display.
   c) Power off all peripheral devices except for the 9900 subsystem.
   d) Power off the host system. You are now ready to connect the 9900 subsystem.
3. **Connect the 9900 to the host system**. The Hitachi Data Systems representative installs the fibre cables between the 9900 and the host 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 host system**. The user should perform this activity. To power on the host system after connecting the 9900:
   a) Power on the host system display.
   b) Power on all peripheral devices. The 9900 should already be on. The fibre-channel ports should already be defined. If not, the host 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 host system connected to the 9900.

2.5 Configuring the Host Fibre-Channel Adapters

After connecting the 9900 subsystem and recording the device numbers for the new devices, you are ready to configure the fibre-channel adapter(s) connected to the 9900. The HBA setup utilities allow you to configure the adapter settings while the system is booting up. The host bus adapters have many configuration options. This section provides the following minimum requirements for configuring FC adapters for operation with the 9900 subsystem.

- The requirements for 9900 devices are specified in **Error! Reference source not found.**.
- The BIOS may need be disabled to prevent the system from trying to boot from the 9900.
- In addition to the queue depth and BIOS, several other parameters (e.g., FC fabric) may also need to be set. Please refer to the user documentation that came with your HBA to determine whether other options are required to meet your operational requirements.

Note: Make sure to use the same settings and device parameters for all 9900 devices.

Table 2.3 Requirements for the 9900 Devices

Parameter	Required Value
IOCB Allocation (Queue depth) per LU	≤ 32 per LU
IOCB Allocation (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 9900 devices.

The following sample instructions apply to the QLogic® 2200F adapter. For instructions on configuring other adapters, refer to the user documentation for the adapter.

To configure a QLogic® 2200F fibre-channel adapter connected to the 9900:

1. While the system is booting up, launch the HBA setup utility as follows: when the message **Press <Alt-Q> to Run QLogic Fast! Utility** appears, press **Alt-Q**.
2. Select the QLogic® adapter to **configure**.
3. Go to **<Configuration Settings>**, select **Advanced Adapter Settings>**, and then verify the following settings:
   - Host Adapter BIOS: Disabled
   - Loop Reset Delay: 2
   - Adapter Hard Loop ID: Enabled
   - Hard Loop ID: Larger than any port addresses
   - Enable LIP Reset: No
   - Login Retry Count: 100
   - Port Down Retry Count: 30
   - IOCB Allocation same as specified in **Error! Reference source not found.**
4. Verify all other required settings for your operational environment. For example, the QLogic® adapter defaults to eight **LUNs per target**. Refer to the user documentation for the adapter as needed.
5. Repeat the preceding steps for each QLogic® FC adapter connected to the 9900 subsystem. When you are finished configuring QLogic® adapters, exit the HBA setup utility.
6. ***Note:*** Be sure to add the Qlogic® driver to the Kernel.

2.6 Verifying Device Recognition

The last step in preparing for 9900 device configuration is to verify that the host system recognizes the new devices. Hitachi Data Systems recommends that the devices already be installed and formatted before the host system is powered on.

1. To verify new device recognition:

```
# dmesg | more↵
:
:
scsi0 : Qlogic QLA2200 PCI to Fibre Channel Host Adapter: 0 device 14 irq 11
        Firmware version:  1.17.26, Driver version  2.11 Beta          Host Adapter name

scsi : 1 host.
  Vendor: HITACHI        Model: OPEN-3            Rev: 0111
  Type:   Direct-Access                      ANSI SCSI revision: 02
Detected scsi disk sda at scsi0, channel 0, id 0, lun 0
                                                    Number of logical unit
                Device file name of this disk is /dev/sda

  Vendor: HITACHI        Model: OPEN-9            Rev: 0111
  Type:   Direct-Access                      ANSI SCSI revision: 02
Detected scsi disk sdb at scsi0, channel 0, id 0, lun 1
:
:
```

2. Device files are created under a directory of **"/dev"**. Verify these files as follows:

```
# ls -l /dev | more↵
:
brw-rw----   1 root     disk       8,   0 May  6  1998 sda
```

Figure 2.3 Verifying Device Recognition

Record the device file names for the new devices. You will need this information when you change the device parameters. (see **Error! Reference source not found.** for a sample SCSI path worksheet).

Table 2.4 Sample SCSI Path Worksheet

LDEV (CU:LDEV)	Device Type	LUSE (✓)	CVS (✓)	Device File Name	Path 1	Alternate Path(s)		
0:00					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:01					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:02					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:03					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:04					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:05					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:06					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:07					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:08					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:09					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0A					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0B					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0C					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0D					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0E					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
0:0F					TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____	TID:____ LUN:____
and so on...								

Chapter 3 Configuring the New Devices

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

- Setting the number of Logical Units (see section 3.1)
- Partitioning the disk devices (see section 3.2),
- Creating, mounting and verifying the file systems (if necessary) for the disk devices (see section 3.3),

3.1 Set the Number of Logical Units

The 9900 supports up to 16 logical units per fibre-channel port. To set the number of LUs, you need to edit the **lilo.conf** file as follows:

```
image=/boot/vmlinuz-qla2x00
       label=linux-qla2x00
       append="max_scsi_luns=16"
 root=/dev/sda7
 read-only
#sbin/lilo
```

Important: A Linux environment only supports a maximum of 16 LUNs in one system. If there are other devices with different host adapters, the number of available LUNs will be reduced accordingly.

3.2 Partitioning Devices

After the device parameters have been changed, you can set the partitions. In a Linux environment, one LU can be divided into a maximum of four primary partitions. Instead of primary partitions, you can make a maximum of one extended partition.

Error! Reference source not found. specifies the partition sizes for the OPEN-x devices. **Error! Reference source not found.** specifies the partition sizes for the CVS LUSE devices (OPEN-x*n CVS).

Table 3.1 Partition Sizes for Standard LUs

Device Type		Partition Size
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
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

Table 3.2 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
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

Table 3.3 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-on	256

To partition the devices:

1. Enter **fdisk/dev/device_name** (for example, **fdisk/dev/sda,** where **sda** is the device name).
2. Select **p** to display the present partitions.
3. Select **n** to make a new partition. You can make up to four primary partitions, or as an alternative, you can make one extended partition. The extended partition can be divided into a maximum of 11 logical partitions, which can be assigned partition numbers from 5 to 15.
4. Select **w** to write the partition information to disk and complete the **fdisk** command.
5. Other commands that you might want to use include:
   - To remove partitions, select **d**.
   - To stop a change, select **q**.
6. Repeat the above steps for each device.

3.3 Creating, Mounting and Verifying the File Systems

3.3.1 Creating the File Systems

After you have partitioned the devices, you can create the file systems, making sure that they are appropriate for the primary and/or extended partition for each logical unit.

```
# mkfs /dev/sda1          (where/dev/sda1 is device file of primary partition number 1.)
```

3.3.2 Create the Mount Directories To create the mount directories, enter the following:

```
# mkdir /9900-LU00↵
```

3.3.3 Create the Mount Table

The file system can be mounted by using the **"mount"** command. In the following example, the first parameter of **"mount"** command should be device file name (**/dev/sda1**) and second parameter should be the mount directory. An example is as follows:

```
# mount /dev/sda1 /9900-LU00↵
    │              └──────────   Mount directory name
    └──       Device file name
#
```

3.3.4 Verify the File System

```
# df -h
Filesystem            Size      Used     Avail    Used%    Mounted on
/dev/sda1             1.8G      890M      866M     51%      /
/dev/sdb1             1.9G      1.0G      803M     57%      /usr
/dev/sdc1             2.2G       13k      2.1G      0%      /9900-LU00
#
```

3.3.5 Set the Auto-Mount Parameters

Edit the **"/etc/fstab"** file and set the auto-mount as follows:

```
# cp -ip /etc/fstab /etc/fstab.standard↵       .... take a backup of /etc/fstab.
# vi /etc/fstab↵
:
/dev/sda1          /9900-LU00 ext2    defaults   0   2
```

Chapter 4 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. Notification of 9900 error conditions is made in real time, providing the UNIX/PC server user with the same level of monitoring and support available to the mainframe user. The SIM reporting via SNMP enables the user to monitor the 9900 subsystem from the UNIX/PC server host without having to check the Remote Console PC for SIMs.

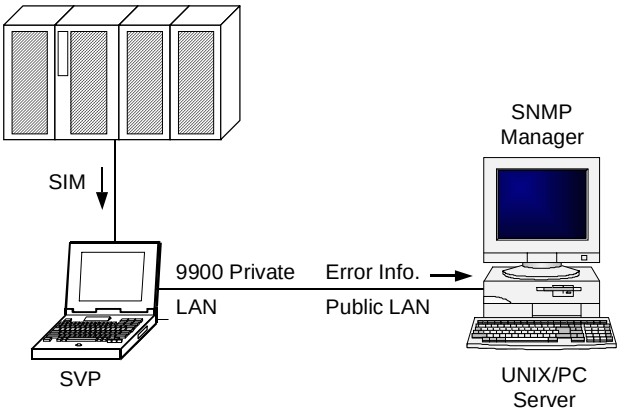

Figure 4.1 9900 SNMP Environment

The SNMP agent only initiates trap operations, which alert the SNMP manager when a SIM occurs. The SNMP manager requests information from the SNMP agent and receives the traps in response to the query.

For assistance with SNMP manager configuration on the UNIX/PC server, please refer to the user documentation for the SNMP management software, or contact the vendor's technical support. For assistance with specific SNMP configuration issues related to the 9900 subsystem, please contact your Hitachi Data Systems representative.

Chapter 5 Troubleshooting

5.1 Troubleshooting

For troubleshooting information on the 9900 subsystem, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide* (BO-98DD845). Table 5.1 lists potential error conditions during 9900 Red Hat® Linux® 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.

Important: When the CHT is replaced, you must pull out all fibre cables that are connected to it.

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 LUSE devices are not intermixed with normal LUs on the same fibre-channel port. Make sure that all new devices have been added to the **/kernel/drv/sd/conf** file.
The file system cannot be created.	Make sure that the device name is entered correctly with **mkfs**. Make sure that the LU is properly connected and labeled.
The file system is not mounted after rebooting.	Make sure the system was restarted properly. Make sure that the auto-mount information in the **/etx/fstab** file is 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. ***Note***: SCSI is not supported on the 9900 except via gateways (bridges). The fibre HBAs require SCSI target IDs, because they exist on a SCSI bus in the host. The fibre IDs can be set automatically through the open-system host.
LUNs cannot be recognized beyond 16	A Linux environment supports a maximum of up to 16 LUNs in one system. If there are other devices with other host adapters, the maximum number is reduced accordingly.

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

AL	arbitrated loop
AL-PA	arbitrated loop physical address
CDE	common desktop environment
ESCON®	Enterprise System Connection (IBM® trademark for optical channels)
FC	Fibre channel
FWD	fast-wide differential
GLM	gigabit link module
HBA	host bus adapter
I/O	input/output
LDEV	logical device
LU	logical unit
LUN	logical unit number
LUSE	LU Size Expansion
LVM	Logical Volume Manager
OFC	open Fibre control
PCI	peripheral component interconnect
r/w	read/write
RISC	reduced-instruction-set computer
R-SIM	remote service information message
SCSI	small computer system interface
SIM	service information message
SNMP	simple network management protocol
SVP	service processor (this is the laptop PC that is part of the 9900 subsystem)
TCP/IP	transmission control protocol/internet protocol
TID	target ID
TOV	timeout value
UWD	ultra-wide differential
VOLSER	volume serial number

Index

9

9900
- fibre topology, 10
- fibre topology settings (table), 14
- instructions for connecting, 15
- overview, 1
- port address, 10, 11
- prerequisites for connecting to, 9
- setting the host mode, 9
- subsystem installation, 8

A

available AL-PA values
- table, 14

C

configuration
- requirements, in general, 7

configuring
- fibre-channel adapters, 16
- fibre-channel ports, 10
- OLogic® adapter (instructions), 17

configuring new devices
- creating the file system, 26
- creating the mount directories, 26
- creating the mount table, 26
- in general, 21–27
- partition sizes, CVS LUSE configuration(table), 23
- partition sizes, in general (table), 22
- partition sizes, LUSE configuration(table), 23
- partitioning devices, 22
- setting number of LUs, 21
- setting the auto-mount procedures, 27
- verifying the file system, 27

D

device recognition
- verifying, 18

device types
- configuration procedures, in general, 2
- definition
  - OPEN-x, 2
  - OPEN-x*n (LUSE), 2
  - OPEN-x*n CVS (CVS LUSE), 2
  - OPEN-x-CVS, 2
- specifications (table), 4
- volume usage (table), 5

E

execution throttle (queue depth)
- requirements, 16

F

fibre-channel adapters
- configuring, 16

fibre-channel ports
- configuring, 10

file systems
- creating, in general, 26–27

I

installing the 9900 subsystem, 8

N

new device configuration
- creating the file system, 26
- creating the mount directories, 26
- creating the mount table, 26
- partition sizes, CVS LUSE configuration(table), 23
- partition sizes, in general (table), 22
- partition sizes, LUSE configuration(table), 23
- setting the auto-mount parameters, 27
- verifying the file system, 27

O

overview
- Linux configuration, 1–5

overview of 9900, 1

Q

queue depth (execution throttle)
- requirements, 16

R

requirements
- execution throttle (queue depth), 16

S

SCSI path worksheet
- table, 19

Set Fibre Topology panel
- illustration, 13

Set Host Mode panel
- illustration, 9

SNMP
- remote subsystem management, 29

SNMP environment
- illustration, 29

T

troubleshooting
contacting the technical support center, 32
in general, 31

V

verifying
device recognition, 18