



**Hitachi Freedom Storage™  
Lightning 9900™ V Series**

**Red Hat® Linux® Configuration Guide**



© 2002 Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

**Notice:** No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi Data Systems Corporation.

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

This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.

## Trademarks

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

Red Hat is a registered trademark of Red Hat, Inc.

Linux is a registered trademark of Linus Torvalds.

Qlogic is a trademark of Qlogic Corporation.

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

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

Adaptec is a registered trademark or trademark of Adaptec, Inc.

JNI is a trademark or registered trademark of JNI Corporation.

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

## Notice of Export Controls

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

## Document Revision Level

Revision	Date	Description
MK-92RD127-0	April 2002	Initial Release

## Source Document Revision Level

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

- Hitachi DKC310 Disk Subsystem SCSI/Fibre Installation Manual, Red Hat® Linux, revision 1.1.
- *Hitachi Lightning 9900™ V LUN Manager User's Guide* (MK-92RD105)
- *Hitachi Lightning 9900™ V Red Hat® Linux® Configuration Guide* (MK-92RD105).

## Referenced Documents

- *Hitachi Lightning 9900™ V Series User and Reference Guide*, MK-92RD100.
- *Hitachi Lightning 9900™ V Series Remote Console - Storage Navigator User's Guide*, MK-92RD101.
- *Hitachi Lightning 9900™ V Series LUN Management User's Guide*, MK-92RD105.
- *Hitachi Lightning 9900™ V Series LUN Expansion (LUSE) and Virtual LVI/LUN User's Guide*, MK-92RD104.
- *Hitachi Lightning 9900™ V Series Hitachi RapidXchange User's Guide*, MK-91RD052.

## Preface

The *Hitachi Lightning 9900™ V Series Red Hat® Linux® Configuration Guide* describes and provides instructions for installing and configuring the devices on the Lightning 9900™ V 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™ V 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 “9900V” to refer to any 9900V subsystem configured for UNIX®/PC server operations, including the all-open and multiplatform models.

For further information on the Lightning 9900™ V array subsystem, please refer to the *Hitachi Lightning 9900™ V Series User and Reference Guide (MK-92RD100)*, 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™ V subsystem and its features and options.

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

### Microcode Level

This document revision applies to 9900V microcode version 20-01-XX and higher.

### COMMENTS

Please send us your comments on this document: [doc.comments@hds.com](mailto:doc.comments@hds.com).

Make sure to include the document title, number, and revision.  
Please refer to specific page(s) and paragraph(s) whenever possible.  
(All comments become the property of Hitachi Data Systems Corporation.)

**Thank you!**



# Contents

<b>Chapter 1</b>	<b>Overview of Red Hat® Linux® Configuration.....</b>	<b>1</b>
1.1	Red Hat® Linux® Configuration.....	1
1.2	The Lightning 9900™ V Array Subsystem .....	2
1.3	Device Types and Configuration Procedures.....	3
<b>Chapter 2</b>	<b>Preparing for New Device Configuration .....</b>	<b>7</b>
2.1	Configuration Requirements .....	7
2.2	Installing the 9900V Subsystem.....	9
2.3	Preparing to Connect the 9900V Subsystem .....	10
2.3.1	Setting the Host Mode for the 9900V Ports .....	10
2.3.2	Configuring the 9900V Fibre-Channel Ports .....	11
2.4	Connecting the 9900V Subsystem .....	14
2.5	Configuring the Host Fibre-Channel Adapters.....	15
2.6	Verifying Device Recognition .....	18
<b>Chapter 3</b>	<b>Configuring the New Devices .....</b>	<b>21</b>
3.1	Set the Number of Logical Units.....	21
3.2	Partitioning Devices .....	22
3.3	Creating, Mounting and Verifying the File Systems .....	23
3.3.1	Creating the File Systems .....	23
3.3.2	Creating the Mount Directories.....	23
3.3.3	Creating the Mount Table .....	23
3.3.4	Verifying the File System.....	24
3.3.5	Setting the Auto-Mount Parameters .....	24
<b>Chapter 4</b>	<b>SNMP Remote Subsystem Management.....</b>	<b>25</b>
<b>Chapter 5</b>	<b>Troubleshooting.....</b>	<b>27</b>
5.1	Troubleshooting .....	27
5.2	Calling the Hitachi Data Systems Technical Support Center .....	28
<b>Appendix A</b>	<b>Acronyms and Abbreviations .....</b>	<b>29</b>

## List of Figures

Figure 2.1	The Add New Host Group Panel .....	10
Figure 2.2	LUN Manager/Port Panel, Port Tab.....	12
Figure 2.3	Verifying New Device Recognition .....	18
Figure 2.4	Verifying Device Files .....	18
Figure 3.1	Setting the number of LUs .....	21
Figure 3.2	Creating the Mount Table .....	23
Figure 3.3	Verifying the File System.....	24
Figure 3.4	Setting the Auto-Mount Parameters .....	24
Figure 4.1	9900V SNMP Environment .....	25

## List of Tables

Table 1.1	9900V Device Specifications for Red Hat® Linux® .....	4
Table 2.1	Fibre Parameter Settings for the 9900V .....	12
Table 2.2	Fibre Parameter Settings on the 9900V .....	13
Table 2.3	Available AL-PA Values .....	13
Table 2.4	Requirements for the 9900V Devices .....	15
Table 2.5	Sample SCSI Path Worksheet.....	19
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™ V subsystem to a Red Hat® Linux® system and configuring the new Lightning 9900™ V devices for operation with the Red Hat® Linux® operating system. The Hitachi Data Systems Customer Service representative performs the physical installation of the 9900V subsystem. The user prepares for 9900V subsystem installation and configures the new 9900V devices with assistance as needed from the Hitachi Data Systems representative.

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

- Setting the host mode and configuring the 9900V fibre-channel ports (see section 2.3).
- Connecting to the 9900V 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 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 Red Hat® Linux® user documentation, or contact technical support.

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

## 1.2 The Lightning 9900™ V Array Subsystem

The Hitachi Lightning 9900™ V 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 9900V subsystem provides continuous data availability, high-speed response, scaleable connectivity, and expandable capacity for PC server and open-system storage. The 9900V 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™ V subsystem can be configured with fibre-channel, FICON™, and/or ExSA™ (Extended Serial Adapter™) ports to provide connectivity with S/390® mainframe hosts as well as open-system hosts. For further information on the 9900V subsystem, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ V User and Reference Guide* (MK-92RD100), or contact your Hitachi Data Systems account team.

For further information on the 9900V subsystem, please refer to the *Hitachi Lightning 9900™ V Series User and Reference Guide* (MK-92RD100), or contact your Hitachi Data Systems account team, or visit Hitachi Data Systems online (<http://www.hds.com>).

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

## 1.3 Device Types and Configuration Procedures

The 9900V subsystem allows the following types of logical devices (LDEVs) to be installed and configured for operation with the Red Hat® Linux® operating system. Table 1.1 lists the device specifications for the 9900V devices.

**OPEN-x Devices** The standard OPEN-x logical units (LUs) (e.g., OPEN-3, OPEN-9) are disk devices of predefined sizes. The 9900V subsystem currently supports OPEN-3, OPEN-9, OPEN-E, and OPEN-L 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 9900V 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 9900V subsystem using fewer LU numbers (LUNs). For further information on the LUSE feature, please refer to the *Hitachi Lightning 9900™ V Series LUN Expansion and Virtual LVI/LUN User's Guide* (MK-92RD104).

**VLL Devices (OPEN-x VLL)** The Virtual LVI/LUN (VLL) feature of the 9900V subsystem enables you to “slice up” a single LU into several smaller LUs to best fit your application needs and improve host access to frequently used files. For further information on the Virtual LVI/LUN feature, please refer to the *Hitachi Lightning 9900™ V Series LUN Expansion (LUSE) and Virtual LVI/LUN User's Guide* (MK-92RD104). **Note:** The OPEN-L LU does not support Virtual LVI/LUN.

**VLL LUSE Devices (OPEN-x\*n VLL)** The VLL LUSE devices combine Virtual LVI/LUN devices (instead of standard OPEN-x LUs) into LUSE devices. The Virtual LVI/LUN feature is used first to create custom-size devices, and then the LUSE feature is used to combine (concatenate) these Virtual LVI/LUN devices. The user can combine from 2 to 36 Virtual LVI/LUN devices into one VLL LUSE device. For example, an OPEN-3 LUSE volume that was created from ten OPEN-3 Virtual LVI/LUN volumes would be designated as an OPEN-3\*10 VLL device. **Note:** The OPEN-L LU does not support Virtual LVI/LUN.

**Table 1.1 9900V Device Specifications for Red Hat® Linux®**

Device Type (Note 1)	Category	Vendor Name	Product Name	# of Blocks (512-byte blk)	Sector Size (bytes)	# of Data Cylinders	# of Heads	# of Sectors per Track	Capacity MB (Note 2)
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-E	Disk	HITACHI	OPEN-E	28452960	512	19759	15	96	13893
OPEN-L	Disk	HITACHI	OPEN-L	71192160	512	49439	15	96	34761
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-E*n	Disk	HITACHI	OPEN-E*n	28452960*n	512	19759*n	15	96	13893*n
OPEN-L*n (n=2 to 14)	Disk	HITACHI	OPEN-L*n	71192160*n	512	49439*n	15	96	34761*n
OPEN-3 VLL	Disk	HITACHI	OPEN-3-CVS	<b>Note 3</b>	512	<b>Note 4</b>	15	96	<b>Note 5</b>
OPEN-9 VLL	Disk	HITACHI	OPEN-9-CVS	<b>Note 3</b>	512	<b>Note 4</b>	15	96	<b>Note 5</b>
OPEN-E VLL	Disk	HITACHI	OPEN-E-CVS	<b>Note 3</b>	512	<b>Note 4</b>	15	96	<b>Note 5</b>
OPEN-3*n VLL	Disk	HITACHI	OPEN-3*n-CVS	<b>Note 3</b>	512	<b>Note 4</b>	15	96	<b>Note 5</b>
OPEN-9*n VLL	Disk	HITACHI	OPEN-9*n-CVS	<b>Note 3</b>	512	<b>Note 4</b>	15	96	<b>Note 5</b>
OPEN-E*n VLL	Disk	HITACHI	OPEN-E*n-CVS	<b>Note 3</b>	512	<b>Note 4</b>	15	96	<b>Note 5</b>

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

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

**Note 3:** The number of blocks for a Virtual LVI/LUN volume is calculated as follows:

$$\# \text{ of blocks} = (\# \text{ of data cylinders}) \times (\# \text{ of heads}) \times (\# \text{ of sectors per track})$$

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

$$\# \text{ of blocks} = (53 \text{ cylinders}) \times (15 \text{ heads}) \times (96 \text{ sectors per track}) = 76320$$

**Note 4:** The number of data cylinders for a Virtual LVI/LUN volume (including VLL LUSE volumes) is calculated as follows:

( $\uparrow$ ... $\uparrow$  means that the value should be rounded up to the next integer)

- The number of data cylinders for an OPEN-\* VLL volume =  
 $\# \text{ of cylinders} = \uparrow (\text{capacity (MB)}) \times 1024/720 \uparrow$   
Example: For an OPEN-3 VLL volume with capacity = 37 MB:  
 $\# \text{ of cylinders} = \uparrow 37 \times 1024/720 \uparrow = \uparrow 52.62 \uparrow$  (rounded up to next integer) = 53 cylinders
- The number of data cylinders for a 3390-3A/C VLL volume =  
 $\# \text{ of cylinders} = (\text{number of cylinders}) + 9$
- The number of data cylinders for a 3390-3B VLL volume =  
 $\# \text{ of cylinders} = (\text{number of cylinders}) + 7$

**Note 5:** The size of an OPEN-x VLL volume is specified by capacity in MB, not by number of cylinders. The user specifies the volume size.



## Chapter 2 Preparing for New Device Configuration

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

- Installing the 9900V subsystem (see section 2.2).
- Setting the host mode and configuring the fibre-channel ports (see section 2.3).
- Connecting the 9900V 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 9900V Red Hat® Linux® configuration are:

- **Hitachi Lightning 9900™ V Series subsystem**, all-open or multiplatform configuration.
  - The 9900V LUN Manager software is used to configure the fibre-channel (FC) ports. If the remote LUN Manager feature is not installed, please contact your Hitachi Data Systems account team for information on LUN and fibre-channel configuration services.  
**Note:** The availability of 9900V features and devices depends on the level of microcode installed on the 9900V subsystem.
  - The 9900V LUN Manager software is used to configure the fibre-channel (FC) ports. For additional instructions, please refer to the *Hitachi Lightning 9900™ V Series LUN Manager User's Guide (MK-92RD105)*. 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.
- **Red Hat® Linux® PC server**. Please refer to the Red Hat® Linux® user documentation for PC server hardware and configuration requirements.
- Red Hat® Linux® operating system, versions 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.
- Root (superuser) login access to the host system.

Continued on the next page

- **Fibre-channel adapters.** Make sure to install all utilities/tools for the adapter. Do not install/load the driver(s) yet (see below). For information on driver requirements for the adapters, please refer to the user documentation for the adapter or contact the vendor.
  - The 9900V subsystem supports: 2 Gbps fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface and multimode optical cables with LC connectors; and 1 Gbps fibre-channel interface, including shortwave non-OFC optical interface and multimode optical cables with SC connectors. Do not connect any OFC-type fibre-channel interface to the 9900V 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). For other information on supported fibre-channel adapters and driver requirements, please refer to the user documentation for the adapter or contact the vendor.
- **Drivers.** Do not install/load the driver(s) yet. Section 2.5 provides instructions on loading the drivers with the correct parameters and options for operation with the 9900V subsystem.

## 2.2 Installing the 9900V Subsystem

The 9900V subsystem comes with all hardware and cabling required for installation. Installation of the 9900V 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, Virtual LVI/LUN and VLL LUSE 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.
  - Configuring the 9900V fibre port topology. The fibre topology parameters for each 9900V fibre-channel port depend on the type of device to which the 9900V port is connected, and the type of port. Determine the topology parameters supported by the device, and set your topology accordingly (see sections 2.3 and 2.5).

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

2. **LUN Manager installation.** The user or Hitachi Data Systems representative can perform this activity. You will use the LUN Manager software to configure the 9900V fibre-channel ports. For instructions on installing the LUN Manager software, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ V Remote Console - Storage Navigator User's Guide* (MK-92RD101). For further information on LUN Manager, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ V LUN Manager User's Guide* (MK-92RD105).

**Note:** If the remote LUN Manager feature is not installed, the Hitachi Data Systems representative can configure fibre-channel 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 9900V Subsystem

### 2.3.1 Setting the Host Mode for the 9900V Ports

The 9900V ports have special modes that must be set for the connected operating system. **Note:** The required host mode setting for 9900V Red Hat® Linux® operations is **00**. Use the LUN Manager software to ensure that the host mode for each fibre port connected to the Red Hat® Linux® system is **00**.

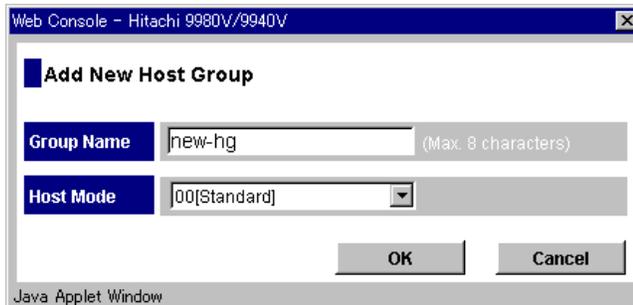


Figure 2.1 The Add New Host Group Panel

## 2.3.2 Configuring the 9900V Fibre-Channel Ports

You need to configure the 9900V fibre-channel ports to define the fibre topology parameters (e.g., arbitrated loop) and port addresses. The LUN Manager software enables you to configure the 9900V fibre-channel ports. For additional instructions, please refer to the *Hitachi Lightning 9900™ V Series LUN Manager User's Guide (MK-92RD105)*. 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.** Figure 2.1 shows the LUN Manager/Port panel, Port Tab which shows the port name, port type, host speed, port address (Loop ID), fabric switch (on or off) and topology (type of connection). Figure 2.2 explains the fibre-parameter settings, and Table 2.2 explains the settings on this panel. You will select the appropriate settings for each 9900V 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 9900V port is also important.

**Port address.** In fabric environments, the port addresses are assigned automatically by fabric switch port number and are not controlled by the 9900V port settings. In arbitrated loop environments, the port addresses are set by entering an AL-PA (arbitrated-loop physical address, or loop ID). Table 2.3 shows the available 9900V AL-PA values, which range 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 9900V port into a SCSI TID.

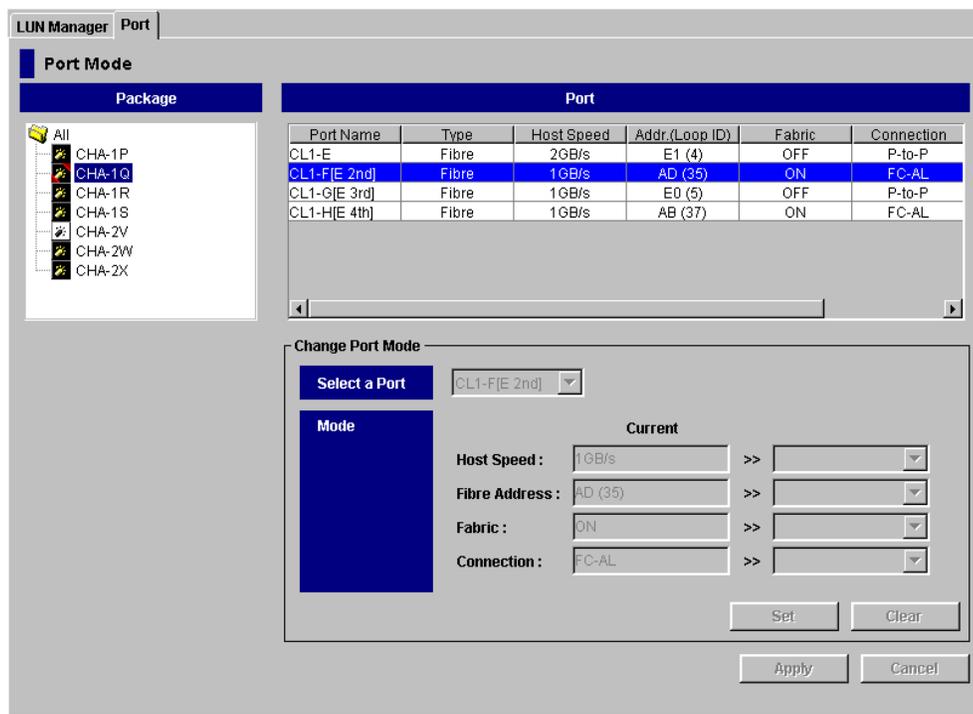
**Note:** The AL-PAs should be unique for each device on the loop to avoid conflicts.

**Table 2.1 Fibre Parameter Settings for the 9900V**

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

**Topology parameters.** Use the LUN Manager software to define the topology parameters and port ID for each fibre-channel port. For additional instructions, please refer to the *Hitachi Lightning 9900™ V Series LUN Manager User's Guide (MK-92RD105)*. Your Hitachi Data Systems representative can also 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:** The 9900V subsystem supports up to 256 LUs per fibre-channel port, but connectivity with Linux® systems is limited to 16 LUs per FC port.



**Figure 2.2 LUN Manager/Port Panel, Port Tab**

**Table 2.2 Fibre Parameter Settings on the 9900V**

Fabric Parameter	Connection Parameter	Provides:
ON	FC-AL	FL Port (Public Arbitrated Loop) (1)
ON	Point-to-Point	Not Supported
OFF	Point-to Point	Not Supported
OFF	FC-AL	NL-port (private arbitrated loop) (2)

(1) The FCE-6410 supports the FL-port.

(2) The QLA2200F supports the NL-port.

**Note:** Some adapters use the FL-port setting, and some adapters use the NL-port setting. Please ask your Hitachi Data Systems representative for assistance with these settings.

**Table 2.3 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 9900V Subsystem

After you have configured the 9900V fibre-channel ports, you are ready to connect the 9900V to the host system(s).

To connect the 9900V 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:
  - a) Shut down the host system.
  - a) When shutdown is complete, power off the Red Hat® Linux® display.
  - b) Power off all peripheral devices except for the 9900V subsystem.
  - c) Power off the host system. You are now ready to connect the 9900V subsystem.
3. **Connect the 9900V to the host system.** The Hitachi Data Systems representative installs the fibre cables between the 9900V and the host system. **Note:** The Hitachi Data Systems representative must use the 9900V 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:
  - a) Power on the host system display.
  - b) Power on all peripheral devices. The 9900V should already be on, and 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 9900V.
  - d) Power on and boot up the host system connected to the 9900V.

## 2.5 Configuring the Host Fibre-Channel Adapters

After connecting the 9900V subsystem and recording the device numbers for the new devices, you are ready to configure the fibre-channel adapter(s) connected to the 9900V. The host bus adapter 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 fibre-channel adapters for operation with the 9900V subsystem.

- The requirements for 9900V devices are specified in Table 2.4.
- The BIOS might need to be disabled to prevent the system from trying to boot from the 9900V.
- 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 9900V devices.

**Table 2.4** Requirements for the 9900V 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 9900V devices later as needed (within the specified range) to optimize the I/O performance of the 9900V 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 9900V:**

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 **<Host Adapter Settings>**, and then verify the following settings:
  - Host Adapter BIOS: **Disabled**
  - Frame size: **1024**
  - Loop Reset Delay: **2**
  - Adapter Hard Loop ID: **Enabled**
  - Hard Loop ID: **Larger than any port addresses**
4. Go back to **<Configuration Settings>**, select **<Advanced Adapter Settings>**, and then verify the following settings:
  - Enable LIP Reset: **No**
  - Login Retry Count: **100**
  - Port Down Retry Count: **30**
  - IOCB Allocation: **as specified in Table 2.4.**
5. 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.
6. Repeat the preceding steps for each Qlogic® fibre-channel adapter connected to the 9900V subsystem. When you are finished configuring Qlogic® adapters, exit the HBA setup utility.

**Note:** Be sure to add the Qlogic® driver to the kernel.

The following sample instructions apply to the JNI™ FCE6410 adapter. For instructions on configuring other adapters, refer to the user documentation for the adapter.

**To configure a JNI™ FCE6410 fibre-channel adapter connected to the 9900V:**

1. Log in as an administrator.
2. To load a `jnifc-2.2.14-5.0smp-vs` module, edit the `/etc/conf.modules` by adding the following line: `alias scsi_hostadapter jnifc-2.2.14-5.0smp-vs`

**Note:** `jnifc-2.2.14-5.0smp-vs` is an example of a module name for smp kernel. If you use a different module name, change the name in the example accordingly.

**Note:** When you are loading the Qlogic® driver, comment out the description to not load the JNI driver.

3. To set the topology with a directory including `jnifc-2.2.14-5.0smp-vs.o`, execute the following command: `insmod jnifc-2.2.14-5.0smp-vs.o jnil116x='lport:0'`

**Note:** There is no way to verify the current setting on a module that has already been loaded. To change the setting, unload the driver and execute the following command:  
`/sbin/rmod jnifc-2.2.14-5.0smp-vs.o.`

4. Execute the following command to create boot image:  
`/sbin/mkinitrd /boot/jni_nport.img "2.2.14-5.0smp"`

**Note:** `jni_nport.img` is an example. The file name is arbitrary.

5. Add the following lines to `/etc/lilo.conf`:

```
image=/boot/jni_nport.img
    label= jni_nport
    root=/dev/sda7
    read-only
```

6. Execute `/sbin/lilo`.
7. Reboot the server.

## 2.6 Verifying Device Recognition

The last step in preparing for 9900V device configuration is to verify that the host system recognizes the new devices. Hitachi Data Systems recommends that the devices should be installed and formatted with the fibre ports configured before the host system is powered on. 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
        ↖ Device file name of this disk = /dev/sda ↗ Number of logical unit

  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
:
:
```

In this example, the HITACHI OPEN-3 (SCSI target ID is 0 and LUN is 0) and the HITACHI OPEN-9 (SCSI target ID is 0 and LUN is 1) are recognized by Red Hat® Linux® server.

**Figure 2.3 Verifying New Device Recognition**

Device files are created under a directory called `/dev`. Verify these files as follows:

```
# ls -l /dev | more
:
brw-rw---- 1 root   disk    8,  0 May  6 1998 sda           Major number =8
                                           Minor number = 0
                                           Device file = sda
```

**Figure 2.4 Verifying Device Files**

Record the device file names for the new devices. You will need this information when you change the device parameters (see Chapter 3). See Table 2.5 for a sample SCSI path worksheet.

**Table 2.5 Sample SCSI Path Worksheet**

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



## Chapter 3 Configuring the New Devices

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

**Important:** The 9900V supports up to 16 logical units per fibre-channel port, but 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.

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
```

**Figure 3.1** Setting the number of LUs

## 3.2 Partitioning Devices

After the device parameters have been changed, you can set the partitions. In a Red Hat® 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.

To partition the devices:

1. Type `fdisk/dev/<device_name>` (for example, `fdisk/dev/sda`, where `sda` is the device name).
1. Select `p` to display the present partitions.
2. 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.
3. Select `w` to write the partition information to disk and complete the `fdisk` command.
4. Other commands that you might want to use include:
  - To remove partitions, select `d`.
  - To stop a change, select `q`.
5. 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. Execute the following command:

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

#### 3.3.2 Creating the Mount Directories

To create the mount directories, type the following command:

```
# mkdir /9900V-LU00
```

#### 3.3.3 Creating the Mount Table

The file system can be mounted by using the “`mount`” command. In the following example, the first parameter of the `mount` command should be the device file name (`/dev/sda1`) and the second parameter should be the mount directory, as shown in the following example:

```
# mount /dev/sda1 /9900V-LU00 ← Mount directory name
    ↖ Device file name
#
```

Figure 3.2 Creating the Mount Table

### 3.3.4 Verifying the File System

After you create the mount table, you should verify the file system, as follows:

```
# 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%    /9900V-LJ00
#
```

Figure 3.3 Verifying the File System

### 3.3.5 Setting the Auto-Mount Parameters

Edit the “`/etc/fstab`” file and set the auto-mount as follows:

```
# cp -ip /etc/fstab /etc/fstab.standard.1 ... take a backup of /etc/fstab.
# vi /etc/fstab.1
:
/dev/sda1          /9900V-LJ00 ext2 defaults 0 2
```

Figure 3.4 Setting the Auto-Mount Parameters

## 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 9900V subsystem utilizes SNMP to transfer status and management commands to the UNIX®/PC server host via the 9900V 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 9900V SVP notifies the SNMP manager on the UNIX®/PC server. Notification of 9900V 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 9900V subsystem from the UNIX®/PC server host.

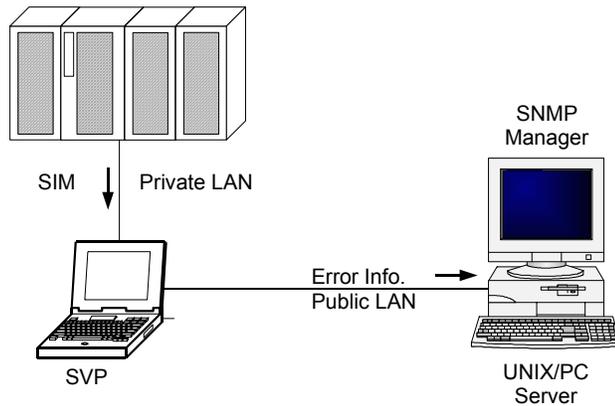


Figure 4.1 9900V SNMP Environment

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 more information on SNMP, see the *Remote Console - Storage Navigator User's Guide* (MK-92RD101). For assistance with specific SNMP configuration issues related to the 9900V subsystem, please contact your Hitachi Data Systems representative.



# Chapter 5 Troubleshooting

## 5.1 Troubleshooting

For troubleshooting information on the 9900V subsystem, please refer to the *Hitachi Lightning 9900™ V Series User and Reference Guide (MK-92RD100)*. Table 5.1 lists potential error conditions during 9900V 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.

**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 9900V subsystem are ON. Make sure that LUSE devices are not intermixed with normal LUs on the same fibre-channel port. Make sure that the LUNs are properly configured. The LUNs for each target ID must start at 0 and continue sequentially without skipping any numbers.
The file system cannot be created.	Make sure that the device name is entered correctly with <b>mkfs</b> . 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 <b>/etc/fstab</b> file is correct.
LUNs cannot be recognized beyond 16	A Red Hat® 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, 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

CruiseControl	CruiseControl performs automatic relocation of volumes to optimize performance.
CU	Control Unit. The 9900V subsystem supports a maximum of 16 logical control unit (CU) images, numbered sequentially from 0 to F. Each CU image controls up to 256 LDEVs.
Custom Access	A feature that allows a non-administrator to be assigned update access to one or more of the restricted <del>Remote Console</del> functions.
DASD	Direct-Access Storage Device
DKC	Disk Controller. The 9900V disk controller provides up to sixteen logical control unit (CU) images and supports 3990-6, 3990-6E, and 2105-F20 disk controller emulation.
DKU	Disk Array Unit. The 9900V subsystem has up to six disk array frames containing the storage components (disk drive arrays) of the subsystem.
ESCON®	Enterprise System Connection
Export File	The Export File function allows you to export data (e.g., monitoring data used by Performance Monitor or Hitachi TrueCopy) to data files.
FD	floppy disk
FD Dump Tool	This function downloads the 9900V configuration information onto a floppy diskette or a hard disk drive, and is generally used for troubleshooting purposes.
FICON™	Fibre Connection
FlashAccess	FlashAccess (Dynamic Cache Residence) enables you to store specific high-usage data directly in cache memory to provide virtually immediate data availability.
GB	gigabyte(s)
Gb	gigabit
Gb/s	gigabits per second
(HiStar)	Hierarchical Star Network architecture improves the total performance of internal data transfer by using high-speed crossbar switches.
HMBR	Hitachi Multiplatform Backup/Restore
Java™ applet program	The Web Client Java™ applet program runs on a web browser. When a user accesses and logs on to the desired SVP, the Web Client Java™ applet is downloaded from the SVP. The Web Client Java™ applet program runs on a web browser.
JVM™	Java Virtual Machine™ is the Web Client Java™ applet program that is installed in each SVP and runs using web browsers to provide a user-friendly interface for the 9900V functions.
kB	kilobyte(s)

LAN	local-area network
LBA	logical block address
LDEV	logical device. An LDEV used by mainframe hosts can be called a device, logical volume image (LVI) or a volume. An LDEV used by open-system hosts is called a logical unit (LU).
LU	An LDEV used by open-system hosts is called a logical unit (LU). Open-system fibre interfaces access LUs that are mapped to one or more LDEVs.
LUN	logical unit number is an identifying number for an LU.
LUN Manager	LUN Manager enables you to configure the 9900V fibre-channel ports for operational environments, and restrict host access to LUs.
LUSE	LUN Expansion. This function allows you to concatenate two or more volumes into a larger volume.
LVI	Logical Volume Image (also called device emulation)
MB	megabyte(s)
MIB	message information block
Open Volume Management	A suite of options that includes Virtual LVI/LUN and LUSE (LUN Expansion). Virtual LVI/LUN divides a logical volume for open-system into two or more volumes. LUN Expansion allows you to concatenate two or more volumes into a larger volume.
Parity group	A set of hard disk drives that have the same capacity and are treated as one group. A parity group contains both user data and parity information, which allows the user data to be accessed in the event that one or more of the drives within the group are not available.
Performance Manager	An option that includes Performance Monitor, CruiseControl, and Priority Access. The Performance Monitor option monitors subsystem and volume activity, and the Hitachi CruiseControl option performs automatic relocation of volumes to optimize performance, and Priority Access***
RMI™	Remote Method Invocation. RMI™ is a remote procedure call, which allows Java™ objects stored in the network to be run remotely.
R-SIM	remote service information message (generated by the 9900V when it detects an error or service requirement).
SIM	service information message (generated by a subsystem when it detects an error or service requirement).
ShadowImage	A program product that allows you to maintain subsystem-internal copies of all user data for purposes such as data backup and duplication.
SNMP	simple network management protocol (part of the TCP/IP protocol suite)
SSID	storage subsystem ID. The 9900V is configured with one SSID for each 64 devices, and up to four SSIDs for each CU image.
SVP	Service Processor (this is the notebook computer that is inside the RAID450).
TCP/IP	transmission control protocol/internet protocol
TID	target ID
Trap	An SNMP agent initiates trap operations when R-SIMs occur, in order to send the R-SIMs to the SNMP manager (see Figure 4.1). An SNMP agent can be configured to deliver traps to more than one SNMP manager.

TrueCopy	TrueCopy is a program product that allows you to perform host-free remote copy operations between 9900V subsystems in different locations for data backup and disaster recovery purposes.
UCB	unit control block
User account list	The user account list includes user information such as user ID, password, and write permission for each 9900V option.
VLL	Virtual LVI/ LUN is a program product that enables you to configure custom-size logical device images and logical units, which are smaller than standard-size devices.
Volser	volume serial number (mainframe volume identifier, not related to the LDEV ID)



# Index

## 9

### 9900V

- fibre topology, 11
- fibre topology settings (table), 13
- instructions for connecting to host system, 14
- overview, 2
- port address, 11
- prerequisites for connecting, 10
- setting the host mode, 10
- subsystem installation, 9
- topology parameters, 12

## A

### adapter

- configuring
  - sample instructions (JNI™ FCE6410), 17
  - sample instructions (Qlogic® 2200F), 16

### Add New Host Group panel

- illustration, 10

### AL-PA values

- table, 13

## C

### configuration

- requirements, in general, 7

### configuring

#### adapter

- sample instructions (JNI™ FCE6410 2200F), 17
- sample instructions (Qlogic® 2200F), 16
- fibre-channel adapters, 15
- fibre-channel ports, 11

### configuring new devices

- creating the file system, 23
- creating the mount directories, 23
- creating the mount table, 23
- partitioning devices, 22
- setting number of LUs, 21
  - illustration, 21
- setting the auto-mount procedures, 24

### creating

- mount table
  - illustration, 23

## D

### device recognition

- verifying, 18

### device types

- configuration procedures, in general, 3
- definition

### OPEN-x, 3

### OPEN-x VLL, 3

### OPEN-x\*n (LUSE), 3

### OPEN-x\*n VLL (VLL LUSE), 3

- specifications (table), 4
- driver, 8

## E

- execution throttle (queue depth) requirements, 15

## F

### fibre-channel adapters

- configuring, 15

### fibre-channel ports

- configuring, 11

### file systems

- creating, in general, 23-24

## I

### installing the 9900V subsystem, 9

### IOCB allocation (queue depth)

- requirements, 15

## L

### LUN Manager/Port Panel, Port Tab

- illustration, 12

## N

### new device configuration

- creating the file system, 23
- creating the mount directories, 23
- creating the mount table, 23
- setting the auto-mount parameters, 24

## O

### overview

- Red Hat® Linux® configuration, 1-5
- overview of 9900V, 2

## Q

### queue depth (execution throttle)

- requirements, 15

## S

### SCSI path worksheet

- table, 19

### setting

- auto-mount parameters
  - illustration, 24

### SNMP

- illustration, 25

in general, 25

## T

troubleshooting

contacting the technical support center, 28  
in general, 27

## V

verifying

device files

illustration, 18

device recognition, 18

illustration, 18

file system

illustration, 24