



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

Red Hat® Linux® Configuration Guide

© 2004 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.

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

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

Linux is a registered trademark of Linus Torvalds.

Qlogic is a registered trademark of Qlogic Corporation.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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
MK-92RD127-1	July 2002	Revision 1, supersedes and replaces MK-92RD127-0
MK-92RD127-2	October 2002	Revision 2, supersedes and replaces MK-92RD127-1
MK-92RD127-3	January 2003	Revision 3, supersedes and replaces MK-92RD127-2
MK-92RD127-4	April 2003	Revision 4, supersedes and replaces MK-92RD127-3
MK-92RD127-5	August 2003	Revision 5, supersedes and replaces MK-92RD127-4
MK-92RD127-6	October 2003	Revision 6, supersedes and replaces MK-92RD127-5
MK-92RD127-7	March 2004	Revision 7, supersedes and replaces MK-92RD127-6

Source Documents for this Revision

- RAID400_RAID450_Public_Mode_Rev[1].2c
OPEN_for_RAID400_RAID450_Public_Mode_Rev[1].4a_ENG

Changes in this Revision

- Updated Comments section (Preface)
- Updated trademark date (Front Matter)
- Added note that if you are using Linux Kernel 2.4.18 or later, public system option mode 236 must be **ON** when you are recognizing 9 or more LUNs, and **OFF** when you are recognizing 8 or fewer LUNs (section 2.1)
- Added system option information (see section 2.3 and Table 2.1)
- Removed System Option 249 (section 2.3)

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 Manager/LUN Security User's Guide*, MK-92RD105
- *Hitachi Lightning 9900™ V Series LUN Expansion (LUSE) and Virtual LVI/LUN (VLL) 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 Series (9900V) 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 Series array subsystem, and
- the user is familiar with the Red Hat® Linux® operating system, including commands, utilities, and file systems.

Note: The term “9900V” refers to the entire Hitachi 9900V subsystem family, unless otherwise noted. Please refer to the *Hitachi Freedom Storage™ 9900V User and Reference Guide* (MK-92RD100) for further information on the 9900V disk array subsystems.

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 21-08-2x and higher.

Comments

Please send us your comments on this document. Make sure to include the document title, number, and revision. Please refer to specific section(s) and paragraph(s) whenever possible.

- E-mail: doc.comments@hds.com
- Fax: 858-695-1186
- Mail:
Technical Writing, M/S 35-50
Hitachi Data Systems
10277 Scripps Ranch Blvd.
San Diego, CA 92131

Thank you! (All comments become the property of Hitachi Data Systems Corporation.)

Contents

Chapter 1	Overview of Red Hat® Linux® Configuration	1
1.1	Red Hat® Linux® Configuration	1
1.2	Hitachi Lightning 9900™ V Series Subsystem	1
1.3	Device Types and Configuration Procedures	2
Chapter 2	Preparing for New Device Configuration.....	5
2.1	Configuration Requirements	5
2.2	Installing the 9900V Subsystem.....	6
2.3	System Option Modes	7
2.4	Preparing to Connect the 9900V Subsystem.....	8
2.4.1	Setting the Host Mode for the 9900V Ports.....	8
2.4.2	Configuring the 9900V Fibre-Channel Ports	9
2.5	Connecting the 9900V Subsystem	11
2.6	Configuring the Host Fibre-Channel Adapters.....	12
2.6.1	Sample Instructions (Qlogic®)	13
2.6.2	Sample Instructions (Emulex® LP9000).....	14
2.7	Verifying Device Recognition	15
Chapter 3	Configuring the New Devices	17
3.1	Setting the Number of Logical Units.....	17
3.2	Partitioning the Devices	19
3.3	Creating, Mounting, and Verifying the File Systems	20
3.3.1	Creating the File Systems	20
3.3.2	Creating the Mount Directories	20
3.3.3	Mounting the New File Systems	20
3.3.4	Verifying the File Systems	20
3.3.5	Setting the Auto-Mount Parameters.....	21
Chapter 4	SNMP Remote Subsystem Management	23
Chapter 5	Troubleshooting	25
5.1	Troubleshooting.....	25
5.2	Calling the Hitachi Data Systems Support Center	26
	Acronyms and Abbreviations.....	27
	Index.....	29

List of Figures

Figure 2.1	Setting the Host Mode.....	8
Figure 2.2	Setting the Fibre-Channel Port Parameters	10
Figure 2.3	Verifying New Device Recognition	15
Figure 2.4	Verifying the Device Files	15
Figure 3.1	Setting the Number of LUs (LILO)	18
Figure 3.2	Setting the Number of LUs (GRUB).....	18
Figure 3.3	Setting the Emulex Driver	18
Figure 3.4	Mounting the New Devices	20
Figure 3.5	Verifying the File System	20
Figure 3.6	Setting the Auto-Mount Parameters	21
Figure 4.1	9900V SNMP Environment.....	23

List of Tables

Table 1.1	9900V Device Specifications for Red Hat® Linux®	3
Table 2.1	System Option Modes for Open-System Connectivity	7
Table 2.2	Fibre Parameter Settings on the 9900V	9
Table 2.3	Available AL-PA Values.....	10
Table 2.4	Requirements for the 9900V Devices	12
Table 2.5	Sample SCSI Path Worksheet.....	16
Table 5.1	Troubleshooting	25

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® operations includes:

- Setting the number of LUs (see section 3.1)
- Partitioning the devices (see section 3.2)
- Creating, mounting, and verifying the file systems (see section 3.3)

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 Hitachi Lightning 9900™ V Series 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.

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 OPEN-x logical units (LUs) (e.g., OPEN-3) are disk devices of predefined sizes. The 9900V subsystem currently supports OPEN-3, OPEN-9, OPEN-E, OPEN-L, OPEN-8, and OPEN-V 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.

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

- Setting the number of LUs (see section 3.1)
- Partitioning the devices (see section 3.2)
- Creating, mounting, and verifying the file systems (see section 3.3)

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-8	Disk	HITACHI	OPEN-8	14351040	512	9966	15	96	7007
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-V	Disk	HITACHI	OPEN-V	Max 125827200 Note 3	512	Note 4	15	128	Note 5
OPEN-3*n	Disk	HITACHI	OPEN-3*n	4806720*n	512	3338*n	15	96	2347*n
OPEN-8*n	Disk	HITACHI	OPEN-8*n	14351040*n	512	9966*n	15	96	7007*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-V*n	Disk	HITACHI	OPEN-V*n	Max 125827200 Note 3	512	Note 4	15	128	Note 5
OPEN-3 VLL	Disk	HITACHI	OPEN-3-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-8 VLL	Disk	HITACHI	OPEN-8-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-9 VLL	Disk	HITACHI	OPEN-9-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-E VLL	Disk	HITACHI	OPEN-E-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-V VLL	Disk	HITACHI	OPEN-V	Note 3	512	Note 4	15	96	Note 5
OPEN-3*n VLL	Disk	HITACHI	OPEN-3*n-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-8*n VLL	Disk	HITACHI	OPEN-8*n-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-9*n VLL	Disk	HITACHI	OPEN-9*n-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-E*n VLL	Disk	HITACHI	OPEN-E*n-CVS	Note 3	512	Note 4	15	96	Note 5
OPEN-V*n VLL	Disk	HITACHI	OPEN-V*n	Note 3	512	Note 4	15	128	Note 5

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 MB = 1024² bytes rather than 1000² bytes.

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

of blocks = (# of data cylinders) × (# of heads) × (# of sectors per track)

(The number of sectors per track is 128 for OPEN-V and 96 for the other emulation types.)

Example: For an OPEN-3 VLL 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 Virtual LVI/LUN volume (including VLL 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-x VLL volume = (except for OPEN-V):

of cylinders = ↑ (capacity (MB) × 1024/720) ↑

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

of cylinders = ↑37 × 1024/720↑ = ↑52.62↑ = 53 cylinders

- The number of data cylinders for an OPEN-V VLL volume =

of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑

Example: For an OPEN-V VLL volume with capacity = 50 MB:

of cylinders = ↑50 × 16/15↑ = ↑53.33↑ = 54 cylinders

- The number of data cylinders for a VLL LUSE volume = (except for OPEN-V):

of cylinders = ↑ (capacity (MB) × 1024/720) ↑ × n

Example: For an OPEN-3 VLL 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 an OPEN-V VLL LUSE volume =

of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑ × n

Example: For an OPEN-V VLL LUSE volume with capacity = 50 MB and n = 4:

of cylinders = ↑50 × 16/15↑ × 4 = ↑53.33↑ × 4 = 54 × 4 = 216

- The number of data cylinders for a 3390-3A/C =

of cylinders = (number of cylinders) + 9

- The number of data cylinders for a 3390-3B VLL volume =

of cylinders = (number of cylinders) + 7

Note 5: The size of an OPEN-3/8/9/E VLL volume is specified by capacity in MB, not number of cylinders. The size of an OPEN-V VLL volume can be specified by capacity in MB or number of cylinders. The user specifies the volume size using the 9900V Virtual LVI/LUN software. For more information, please see *Hitachi Lightning 9900™ V Series LUN Expansion (LUSE) and Virtual LVI/LUN (VLL) User's Guide*, MK-92RD104.

Chapter 2 Preparing for New Device Configuration

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.
 - **Note:** If you are using Linux Kernel 2.4.18 or later, public system option mode 238 must be **ON** when you are recognizing 9 or more LUNs, and **OFF** when you are recognizing 8 or fewer LUNs.
- **Red Hat® Linux® server**. Please refer to the Red Hat® Linux® user documentation for server hardware and configuration requirements.
- **Red Hat® Linux® operating system**, version 7.1 or 7.2.
 - Root (superuser) login access to the host system.
- **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 fibre-channel 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 2.6 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
 - 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 section 2.4.2).

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

2. **LUN Manager installation.** The user 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 Lightning 9900™ V Series Remote Console - Storage Navigator User's Guide* (MK-92RD101). For further information on LUN Manager, please refer to the *Hitachi Lightning 9900™ V Series LUN Manager User's Guide* (MK-92RD105).

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 System Option Modes

To provide greater flexibility and enable the 9900V to be tailored to unique customer operating requirements, additional operational parameters, or system option modes, are available for the 9900V subsystem. At installation, the 9900V modes are set to their default values (**OFF**), so make sure to discuss these settings with your Hitachi Data Systems team. The 9900V modes can only be changed by the Hitachi Data Systems representative.

Table 2.1 show the public 9900V system option modes. **Note:** This 9900V mode information was current at the time of publication of this document but may change. Please contact your Hitachi Data Systems representative for the latest information on the 9900V System Option Modes.

Table 2.1 System Option Modes for Open-System Connectivity

Mode	Level	Host Mode	Platform	HA Software	Description and Usage
236	Optional	0x00	Linux	Linux	For Linux (kernel 2.4.18 or later) The number of recognized LUNs is changed in the kernel version 2.4.18 or later. ON: Nine or more LUNs are recognized. OFF: Eight or less LUNs are recognized. Micro-code level : MAIN Ver. 21-04-04/00 and higher.
254	Optional	0x00 & 0x09	All	Veritas Database Editions	Change the response of reserve conflict status to Read Capacity, Verify, and Start Stop Unit. ON: Normal end is reported. OFF: Reserve conflict status is reported. Micro-code level : MAIN Ver. 21-06-22/00 and higher. HDS engineers should refer to the appropriate level of ECN DKC4601-M044.

2.4 Preparing to Connect the 9900V Subsystem

Before the 9900V is connected to the Linux[®] system, you must perform the following tasks:

- Set the host mode for the 9900V fibre-channel ports (see section 2.4.1), and
- Configure the 9900V fibre-channel ports (see section 2.4.2).

After the 9900V subsystem has been connected to the Linux[®] system (see section 2.5), you need to perform the following tasks before configuring the new devices:

- Configure the host fibre-channel adapters connected to the 9900V (see section 2.6), and
- Verify new device recognition (see section 2.7).

2.4.1 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 Setting the Host Mode

2.4.2 Configuring the 9900V Fibre-Channel Ports

You need to configure the 9900V fibre-channel ports to define the fibre parameters (see Table 2.2, Figure 2.2, and Table 2.3). You will use the LUN Manager remote console software to configure the 9900V fibre-channel ports. For instructions on using the LUN Manager software, please refer to the *Hitachi Lightning 9900™ V Series LUN Manager User's Guide* (MK-92RD105).

Note: The 9900V subsystem supports up to 512 LUs per fibre-channel port (256 per host group), but connectivity with Linux® systems is limited to 16 LUs per FC port.

Fibre topology. Figure 2.2 shows the LUN Manager panel used to define the port parameters, 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 on loop ID conflict: The AL-PAs should be unique for each device on the loop to avoid conflicts.

Table 2.2 Fibre Parameter Settings on the 9900V

Fabric Parameter	Connection Parameter	Provides:
ON	FC-AL	<i>Not Supported</i>
ON	Point-to-Point	F-port (fabric loop)
OFF	Point-to Point	<i>Not Supported</i>
OFF	FC-AL	NL-port (private arbitrated loop) *

* The QLA2200F supports the NL-port.

Note: Please contact Hitachi Data Systems for information on port topology configurations supported by each host bus adapter/switch combination. Some adapters use the FL-port setting, and some adapters use the NL-port setting.




Figure 2.2 Setting the Fibre-Channel Port Parameters

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	

2.5 Connecting the 9900V Subsystem

After you have configured the 9900V fibre-channel ports (see section 2.4), 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.
 - b) When shutdown is complete, power off the Red Hat® Linux® display.
 - c) Power off all peripheral devices except for the 9900V subsystem.
 - d) 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.

Note: Get ready to access the host bus adapter utility while the system is booting up.

2.6 Configuring the Host Fibre-Channel Adapters

After connecting the 9900V subsystem, you need to configure the fibre-channel adapter(s) connected to the 9900V. The host bus adapter (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 host fibre-channel adapters for operation with the 9900V subsystem.

- The queue depth 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.

2.6.1 Sample Instructions (Qlogic®)

Qlogic® 2200F, 2310F, 2340F adapter. The following sample instructions apply to the Qlogic® 2200F, 2310F, 2340F adapter. Qlogic driver **06.01.00** is recommended for use with the 9900V. For other adapters, please refer to the user documentation for the adapter.

To configure a Qlogic® 2200F, 2310F, 2340F 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.**
 - LUNs Per Target: **32** (GRUB used as Bootloader. Refer to section 3.1 and to the note in section 2.4.1)
5. Go to **<Configuration Settings>**, select **<Extended Firmware Settings>**. Verify the following settings:
 - Connection Options: **0 (FC_AL)**
3 (Point to Point For Switch Connection)
6. 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.
7. 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: Make sure to add the Qlogic® driver to the kernel (driver **06.01.00** is recommended).

2.6.2 Sample Instructions (Emulex® LP9000)

Emulex® LP9000 adapter. The following sample instructions apply to the Emulex® LP8000, LP9000, and LP9002L adapter. For other adapters, please refer to the user documentation for the adapter.

To configure a Emulex® LP9000 fibre-channel adapter connected to the 9900V:

1. Open the file **lpfc.conf.c*** and verify the following settings:

lpfc_fcfabric_tmo	255
lpfc_linkdown_tmo	255
lpfc_nodey_tmo	255
lpfc_lun_queue_depth	256/LUN ≤ Queue depth ≤ 32
lpfc_topology	2 (Point To Point)
	4 (FC_AL)

*: lpfc.conf.c exists in the same directory as the installed Emulex® driver.

2. Create the Emulex driver. (Refer to the Emulex driver user manual for instructions.)

2.7 Verifying Device Recognition

The last step in preparing for 9900V device configuration is to verify that the host system recognizes the new devices. The 9900V ports must already be configured (see section 2.4), and the HBA(s) connected to the 9900V must already be configured (see section 2.6).

To verify new device recognition:

1. Display the devices using the **dmesg** command (see Figure 2.3).
2. Record the device file name for each new 9900V device. You will need this information when you partition the devices (see section 3.2). See Table 2.5 for a sample SCSI path worksheet.
3. The device files are created under the **/dev** directory. Verify that a device file was created for each new 9900V disk device (see Figure 2.4).

```
# 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

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 ↩ Logical unit number

        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 device (TID 0, LUN 0) and the HITACHI OPEN-9 device (TID 0, LUN 1) are recognized by the Red Hat® Linux® server.

Figure 2.3 Verifying New Device Recognition

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

Figure 2.4 Verifying the Device Files

Table 2.5 Sample SCSI Path Worksheet

LDEV (CU:LDEV)	Device Type	LUSE (*n)	VLL (MB)	Device File Name	Path	Alternate Path
0:00					TID:____ LUN:____	TID:____ LUN:____
0:01					TID:____ LUN:____	TID:____ LUN:____
0:02					TID:____ LUN:____	TID:____ LUN:____
0:03					TID:____ LUN:____	TID:____ LUN:____
0:04					TID:____ LUN:____	TID:____ LUN:____
0:05					TID:____ LUN:____	TID:____ LUN:____
0:06					TID:____ LUN:____	TID:____ LUN:____
0:07					TID:____ LUN:____	TID:____ LUN:____
0:08					TID:____ LUN:____	TID:____ LUN:____
0:09					TID:____ LUN:____	TID:____ LUN:____
0:0A					TID:____ LUN:____	TID:____ LUN:____
0:0B					TID:____ LUN:____	TID:____ LUN:____
0:0C					TID:____ LUN:____	TID:____ LUN:____
0:0D					TID:____ LUN:____	TID:____ LUN:____
0:0E					TID:____ LUN:____	TID:____ LUN:____
0:0F					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. **Note:** Configuration of the 9900V devices should only be performed by a Linux® system administrator and/or a Hitachi Data Systems representative. Configuration requires superuser/root access to the host system. If you have questions or concerns, please contact the Hitachi Data Systems Support Center.

The activities involved in 9900V device configuration are:

- Setting the number of logical units (see section 3.1),
- Partitioning the new devices (see section 3.2), and
- Creating, mounting, and verifying the file systems for the new devices (see section 3.3).

3.1 Setting the Number of Logical Units

Important: The 9900V supports up to 512 logical units per fibre-channel port (256 per host group), but the Red Hat® Linux® system only supports a maximum of 16 LUNs in one system. If other devices already exist on different host adapters, the number of available LUs will be reduced accordingly.

To set the number of LUs:

1. Edit the `/etc/modules.conf` file to add the following line, for example:

```
options scsi_mod max_scsi_luns=16
```
2. To set the Emulex Driver, as shown in Figure 3.3, add the following line to the `/etc/modules.conf` file:

```
Alias scsi_hostadapter lpfcdd
```
3. To activate the above modification, make an image file for booting, for example:

```
# mkinitrd /boot/initrd-2.4.x.scsiluns.img `uname -r`
```
4. To change the setting of Bootloader, use one of the following methods:
 - a) **LILO used as Bootloader.** You need to edit the `lilo.conf` file as shown in Figure 3.1, and then execute the `lilo` command to activate the `lilo.conf` setting with selecting the label. For example: `# lilo`
 - b) **GRUB (GRand Unified Bootloader) is used as Bootloader.** You need to edit the `/boot/grub/grub.conf` file as shown in Figure 3.2.
5. Reboot the system.

<pre>image=/boot/vmlinuz-qla2x00 label=Linux-qla2x00 append="max_scsi_luns=16" # initrd=/boot/initrd-2.4.x.img initrd=/boot/initrd-2.4.x.scsiluns.img root=/dev/sda7 read-only #sbin/lilo</pre>	← <i>Comment out this line.</i> ← <i>Add this line.</i>
---	---

Figure 3.1 **Setting the Number of LUs (LILO)**

<pre>kernel /boot/vmlinuz-2.4.x ro root=/dev/hda1 # initrd /boot/initrd-2.4.x.img initrd /boot/initrd-2.4.x.scsiluns.img</pre>	← <i>This line is commented out.</i> ← <i>Add this line.</i>
--	--

Figure 3.2 **Setting the Number of LUs (GRUB)**

<pre>Alias scsi_hostadapter lpfcdd</pre>	← <i>Add this to</i> <i>/etc/modules.conf</i>
--	--

Figure 3.3 **Setting the Emulex Driver**

3.2 Partitioning the Devices

After the setting the number of logical units, you can set the partitions.

To partition the 9900V disk devices:

1. Enter `fdisk/dev/<device_name>` (for example, `fdisk/dev/sda`, where `/dev/sda` is the device file 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 (1-4), 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.
Note: 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 new 9900V disk 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.

To create the file system, execute the **mkfs** 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, execute the **mkdir** command:

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

3.3.3 Mounting the New File Systems

Use the **mount** command to mount each new file system (see example in Figure 3.4). The first parameter of the **mount** command is the device file name (**/dev/sda1**), and the second parameter is the mount directory, as shown in Figure 3.4.

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

Figure 3.4 Mounting the New Devices

3.3.4 Verifying the File Systems

After mounting the file systems, you should verify the file systems (see example in Figure 3.5).

```
# 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-LU00
#
```

Figure 3.5 Verifying the File System

3.3.5 Setting the Auto-Mount Parameters

To set the auto-mount parameters, edit the `/etc/fstab` file (see example in Figure 3.6).

# cp -ip /etc/fstab /etc/fstab.standard	← <i>Make a backup of /etc/fstab.</i>
# vi /etc/fstab	← <i>Edit /etc/fstab.</i>
:	
/dev/sda1 /9900V-LU00 ext2 defaults 0 2	← <i>Add new device.</i>

Figure 3.6 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.

When a SIM occurs, the 9900V 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 failover and SNMP management software on the UNIX®/PC server host. 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 further information on 9900V SNMP, please refer to 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.




Figure 4.1 9900V SNMP Environment

Chapter 5 Troubleshooting

5.1 Troubleshooting

The Hitachi Lightning 9900™ V Series subsystem provides continuous data availability. 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 contact your Hitachi Data Systems representative, 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 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 mkfs . Make sure that the LU is properly connected and partitioned.
The file system is not mounted after rebooting.	Make sure that the system was restarted properly. Make sure that the auto-mount information in the /etc/fstab file is correct.
LUNs cannot be recognized beyond 16.	Red Hat® Linux® versions 7.1 and 7.2 support a maximum of 16 LUNs in one system. If there are other devices on other host adapters, the maximum number is reduced accordingly.

5.2 Calling the Hitachi Data Systems Support Center

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

The worldwide Hitachi Data Systems Support Centers are:

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

Acronyms and Abbreviations

AL	arbitrated loop
AL-PA	arbitrated loop physical address
blk	block
CU	control unit
CVS	custom volume size
ESCON®	Enterprise System Connection (IBM trademark for optical channels)
ExSA™	Extended Serial Adapter™
FC	fibre-channel
FCA	fibre-channel adapter
FCP	fibre-channel protocol
GB	gigabytes
Gbps	gigabits per second
HBA	host bus adapter
I/O	input/output
LDEV	logical device
LU	logical unit
LUN	logical unit, logical unit number
LUSE	LUN Expansion
LVI	logical volume image
MB	megabytes
OFC	open fibre control
PA	physical address
PC	personal computer system
PCI	power control interface
RAID	redundant array of independent disks
SCSI	small computer system interface
SIM	service information message
SNMP	simple network management protocol
SVP	service processor
TID	target ID
VLL	Virtual LVI/LUN

Index

9

9900V

- connecting to host system, 11
- overview of subsystem, 1
- prerequisites for connecting, 8
- subsystem installation, 6

A

adapter

- configuring Emulex® LP9000, 14
- configuring Qlogic® 2200F, 2310F, 2340F, 13
- adapter hard loop ID (Qlogic® setting), 13
- AL-PA values, 9, 10
- auto-mount parameters, 23

B

BIOS, 12

C

- configuration requirements, 5
- configuring
 - 9900V fibre-channel ports, 8-10
 - adapter (Emulex® LP9000), 14
 - adapter (Qlogic® 2200F, 2310F, 2340F), 12-14
- connection parameter (9900V port), 9
- cylinders, calculating number of, 4

D

- device file names, recording, 15
- device files, verifying, 15
- device recognition, verifying, 15
- device types, 2
 - LUSE, 2
 - specifications, 3
 - VLL (CVS), 2
 - VLL LUSE, 2
- driver (HBA), 13

E

- etc/fstab file, 23

F

- fabric parameter (9900V port), 9
- FC-AL setting (9900V port), 9
- fibre-channel adapters, configuring, 12-14
- file systems, 22-23
- FL-port setting (HBA), 9
- frame size (Qlogic® setting), 13

H

- hard loop ID (Qlogic® setting), 13
- host adapter BIOS (Qlogic® setting), 13
- host bus adapters, configuring, 12-14
- host mode, 8

I

- installing the 9900V subsystem, 6
- IOCB allocation (queue depth), 12

L

- lilo.conf file, 19
- Linux®
 - power off/on, 11
 - supported versions, 5
- LIP reset (Qlogic® setting), 13
- login retry count (Qlogic® setting), 13
- loop ID, 9
- loop reset delay (Qlogic® setting), 13
- LUN Manager software, 6
- LUSE devices, 2

M

- MAXTAGS, 12
- mount directories, 22
- mounting the file systems, 22

N

- NL-port setting (HBA), 9
- number of LUs, setting, 19

P

- partitioning the devices, 21
- point-to-point setting (9900V port), 9
- port down retry count (Qlogic® setting), 13
- ports
 - addresses (AL-PA), 9
 - configuring, 8-10
 - host mode, 8
 - topology parameters, 9

Q

- Qlogic® 2200F, 2310F, 2340F adapters, 12-14
- queue depth, 12

S

- SCSI path worksheet
 - table, 16
- SNMP, 25
 - illustration, 25

system option modes, 6

T

topology parameters, 9

troubleshooting, 27

 calling the Support Center, 28

V

VLL devices, 2

VLL LUSE devices, 2