



**Hitachi Freedom Storage™
Thunder 9500™ V Series
IBM® AIX® Host Installation**

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.

Freedom Storage and Thunder 9500 are trademarks of Hitachi Data Systems Corporation.

AIX, IBM, POWERserver, POWERstation, RISC System/6000, RS/6000, S/390, SP, and System Management Interface Tool (SMIT) are either registered trademarks or trademarks of International Business Machines Corporation (IBM) in the United States and other countries.

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.

LP8000 is a registered trademark of Emulex Corporation. Lightpulse is a trademark of Emulex 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-92DF617-P	September 2002	Preliminary Release
MK-92DF617-P2	December 2002	Supersedes and replaces MK-92DF617-P1
MK-92DF617-0	February 2003	Supersedes and replaces MK-92DF617-P2

Source Document Revision Level

The following source documents were used to produce this 9500V host installation guide:

- *Hitachi Freedom Storage™ Thunder 9200™ IBM® AIX® Host Installation Guide*, MK-91DF544
- Hitachi Data Systems' review of this document

Changes in this Revision

- Added recommended parameters (section 2.3.1)

Preface

The *Hitachi Freedom Storage™ Thunder 9500™ V Series IBM® AIX® Host Installation Guide* describes and provides instructions for installing and configuring the devices on the 9500V disk array subsystem for operation with an IBM® AIX® operating system. This configuration guide assumes that:

- the user has a background in data processing and understands direct-access storage device subsystems and their basic functions,
- the user is familiar with the 9500V array subsystem,
- the user is familiar with the IBM® AIX® operating system and the IBM® RISC System/6000 (RS/6000®), POWERstation®, POWERserver®, and/or SP system,
- the user is familiar with the AIX® Journaled File System, system commands, and utilities.

Note: The term “9500V” refers to the entire Hitachi Freedom Storage™ Thunder 9500™ V Series subsystem family, unless otherwise noted. Please refer to the *Hitachi Freedom Storage™ Thunder 9500™ V Series User and Reference Guide* (MK-92DF601) for further information on the 9500V disk array subsystems.

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

COMMENTS

Please send us your comments on this document: 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

Chapter 1	Overview of IBM® AIX® Configuration	1
1.1	IBM® AIX® Configuration	1
1.2	The 9500V Array Subsystem	1
Chapter 2	Preparing for New Device Configuration	3
2.1	Configuration Requirements	3
2.2	Installing the 9500V Subsystem	4
2.3	Preparing to Connect the 9500V	5
2.3.1	Setting the Host-Specific Parameters for the 9500V Ports	5
2.4	Connecting the 9500V Subsystem	7
2.5	Verifying New Device Recognition	7
Chapter 3	Configuring the New Devices	9
3.1	Changing the Device Parameters	10
3.1.1	Changing Device Parameters for the 9500V	10
3.1.2	Special Instructions for Using an Emulex Lightpulse™ LP8000® Adapter	13
3.2	Assigning the New Devices to Volume Groups and Setting the Partition Size	16
3.3	Creating, Mounting, and Verifying the File Systems	18
3.3.1	Creating the File System	18
3.3.2	Mounting and Verifying the File Systems	21
Chapter 4	Troubleshooting	23
4.1	Troubleshooting	23
4.2	Calling the Hitachi Data Systems Technical Support Center	24
Appendix A	Acronyms and Abbreviations	25
Appendix B	Online Installation and Deinstallation of Devices	27
Appendix C	Commands for the AIX Command Line	29

List of Figures

Figure 2.1	Verifying New Device Recognition.....	7
Figure 3.1	Verifying the Device Parameters Using the lsattr -E -1 hdiskx Command.....	12
Figure 3.2	Verifying the Device Parameters Using the lscfg -vl hdisk1 Command	12
Figure 3.3	Changing Device Parameters When Using an Emulex Lightpulse™ LP8000® Adapter (Command Line)	14
Figure 3.4	Changing Device Parameters When Using an Emulex Lightpulse™ LP8000® Adapter (SMIT®)	15
Figure 3.5	Assigning Devices to Volume Groups and Setting the Partition Size	17
Figure 3.6	Adding a Journaled File System Using SMIT®	19
Figure 3.7	Verifying Creation of Journaled File System.....	19
Figure 3.8	Determining the Maximum File System Size	20
Figure 3.9	Verifying the Auto-Mounted File Systems	21
Figure 3.10	Final File System Verification	21

List of Tables

Table 2.1	Fibre-Channel Adapters	3
Table 2.2	Device Data Table (Sample)	8
Table 3.1	R/W Time-Out and Queue Type Requirements.....	10
Table 3.2	Partition Sizes for Disk Devices	17
Table 4.1	Troubleshooting	23
Table C.1	Commands and Usage	29

Chapter 1 Overview of IBM® AIX® Configuration

1.1 IBM® AIX® Configuration

This document describes the requirements and procedures for connecting the 9500V subsystem to an IBM® AIX® system and configuring the new 9500V devices for operation with the IBM® AIX® operating system. The Hitachi Data Systems representative performs the physical installation of the 9500V subsystem. The user prepares for 9500V subsystem installation and configures the new 9500V devices with assistance as needed from the Hitachi Data Systems representative.

- Configuration of the 9500V disk devices for IBM® AIX® operations includes:
- Changing the device parameters (see section 3.1),
- Assigning the new devices to volume groups and setting the partition size (see section 3.2),
- Creating the Journaled File Systems (see section 3.3.1), and
- Mounting and verifying the file systems (see section 3.3.2).

1.2 The 9500V Array Subsystem

The Hitachi 9500V 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 9500V subsystem provides continuous data availability, high-speed response, scaleable connectivity, and expandable capacity for PC server and open-system storage. The 9500V 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 9500V subsystem can be configured with fibre-channel ports to provide connectivity with open-system hosts. For further information on the 9500V subsystem, please refer to the *Hitachi Freedom Storage™ Thunder 9500™ V Series User and Reference Guide*, or contact your Hitachi Data Systems account team.

For further information on the 9500V subsystem, please refer to the *Hitachi Freedom Storage™ Thunder 9500™ V Series User and Reference Guide*, or contact your Hitachi Data Systems account team.

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

The requirements for 9500V IBM® AIX® configuration are:

- 9500V subsystem, all-open or multiplatform configuration:

Note: You will use the Resource Manager 9500V software to configure the 9500V fibre ports. For instructions on using the Resource Manager, please refer to the *Hitachi Thunder 9200™ Resource Manager 9200 User and Reference Guide*, (MK-90DF504).

Note: The availability of 9500V features and devices depends on the level of microcode installed on the 9500V subsystem.

- IBM® RS/6000®, POWERstation®, POWERserver®, or SP series system.
- IBM® AIX® operating system, version 4.3.3. **Important:** Please contact IBM® to make sure the most current OS patches are installed on the IBM® system(s).

Note: Hitachi Data Systems plans to support future releases of IBM® AIX®. For the latest information on AIX® version support, contact your Hitachi Data Systems account team.

- **Root** (superuser) login access to the IBM® system.
- Fibre-channel adapters. Make sure to install all utilities, tools, and drivers that come with the adapter(s).
 - The 9500V subsystem supports full-speed (100 MB/s) fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface, and multimode optical cables with SC connectors. Do not connect any OFC-type fibre-channel interface to the 9500V 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 4.2).
 - **Note:** If you plan to connect different types of servers to the 9500V subsystem via the same fabric switch, you must use the **zoning** function of the fabric switch.
 - The fibre-channel adapters and drivers listed in Table 2.1 have been verified for use with the 9500V subsystem. For information on the driver requirements for other fibre-channel adapters, please refer to the user documentation for the adapter or contact the vendor.

Table 2.1 Fibre-Channel Adapters

Fibre-Channel Adapter	Driver(s)	AIX® OS Version
Emulex Lightpulse™ LP8000®	V4.0.2.2	AIX 4.3.3
IBM 6227	AIX bundle	4.3.3/5.1
IBM6228	AIX bundle	4.3.3/5.1

2.2 Installing the 9500V Subsystem

The 9500V subsystem comes with all hardware and cabling required for installation. Installation of the 9500V subsystem involves the following activities:

1. **Hardware installation.** The Hitachi Data Systems representative performs hardware installation as specified in the 9500V Maintenance Manual. Follow all precautions and procedures in the 9500V maintenance manual. Check all specifications to ensure proper installation and configuration. Hardware installation includes:

- Assembling all hardware and cabling.
- Loading the latest microcode updates for full fibre-channel support.
- Using Resource Manager 9500V to create needed RAID groups, and to create and format the LUNs.
- Installing the fibre-channel adapters and cabling. **Note:** The total length of the fibre cables that are attached to each fibre-channel adapter must not exceed 500 meters (1,640 feet). Do not connect any OFC-type connector to the 9500V subsystem. Do not connect/disconnect fibre-channel cabling that is being actively used for I/O. This can cause the IBM® AIX® system to hang. Always confirm that the devices on the fibre cable are offline before connecting/disconnecting the fibre cable.

9500V fibre-channel port: The fibre topology parameters for each 9500V fibre-channel port depend on the type of device to which the 9500V port is connected. Determine the topology parameters supported by the device, and set your topology accordingly (see section 2.3.2). The type of 9500V port is also important.

- Connect the subsystem

Note: The Hitachi Data System representative must use the 9500V Maintenance Manual during all installation activities. Follow all precautions and procedures in the maintenance manual, and always check all specifications to ensure proper installation and configuration.

2.3 Preparing to Connect the 9500V

Before the 9500V is connected to the AIX® system, you must perform the following tasks:

- Set the host-specific parameters for the 9500V fibre-channel port(s) (see section 2.3.1), and
- Configure the 9500V fibre-channel ports (see section 2.3.1).

You will use the Resource Manager 9500V software to configure the 9500V fibre ports. For instructions on using the Resource Manager, please refer to the *Hitachi 9200™ Resource Manager 9200 User and Reference Guide* (MK-90DF504).

2.3.1 Setting the Host-Specific Parameters for the 9500V Ports

The 9500V ports must be configured for the connected operating system.

Fibre Topology. You need to configure the 9500V FC ports to define the fibre topology parameters and port addresses. The 9500V subsystem supports up to 512 LUs and 128 host groups per port (with 256 LUNs per port host). You will select the appropriate settings for each 9500V 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.

For an IBM AIX configuration with HDLM and HACMP using a 6227 or 6228 HBA connected to the 9500 through a 3200 Brocade switch, the recommended parameters are:

Data Share mode = use
UA (06/2A00) suppress mode enable
Logical unit reset mode enable
NACA enable
Wolfpack mode enable (for online LUN formatting)

Note: The online format will only function in AIX with Wolfpack mode ON and NO HDLM.

Note: If you plan to connect different types of servers to the 9500V via the same fabric switch, you must use either the **zoning** function of the switch or Hitachi SANTinel (LUN security).

Port address. In fabric environments, the port addresses are assigned automatically by fabric switch port number and are not controlled by the 9500V port settings. In FC arbitrated-loop (FCAL) environments, the port addresses are set by entering an AL-PA (arbitrated-loop physical address, or loop ID, or port address). The host communicates with the devices comprising the loop with 8-bit AL-PA.

Available AL-PA values and the corresponding SCSI TID addresses are listed in Appendix B. The number of available port addresses is 126. (There are 127 port addresses, but address 00H is reserved for fibre connection.) Fibre-channel protocol uses the AL-PAs to communicate on the fibre-channel link, but the software driver of the platform host adapter translates the AL-PA value assigned to the 9500V port to a SCSI TID.

Devices communicate with hosts using individual port addresses. However, hosts map SCSI protocol to fibre channel devices. The hosts access the device's LUs using the device files `/dev/hdisk/*` and `/dev/rhdisk/*`. SCSI and fibre-channel devices are accessed the same way; however, the device files for SCSI and fibre-channel devices are configured differently.

2.4 Connecting the 9500V Subsystem

After you have configured the 9500V fibre-channel ports, you are ready to connect the 9500V subsystem to the IBM® AIX® system. The 9500V comes with all hardware and cabling required for connection to the host system(s).

To connect the 9500V subsystem to the IBM® system:

1. **Verify subsystem installation.** The Hitachi Data Systems representative verifies the fibre-port address configuration and the status of the fibre-channel adapters and LUNs (normal).
2. **Connect the 9500V to the IBM® system.** Install the fibre cables between the 9500V and the IBM® system.

To recognize the LUNs, run the **cfgmgr** command.

2.5 Verifying New Device Recognition

The final step before configuring the new 9500V disk devices is to verify that the host system recognizes the new devices. The host system automatically creates a device file for each new device recognized. Hitachi Data Systems recommends that the devices should be installed and formatted with the fibre ports configured before the host system is powered on. If the system is not restarted, the user must issue the **cfgmgr** command to force the system to check the buses for new devices.

To verify new device recognition:

1. Log in to the host system as **root**.
2. Display the system device data using the **lsdev -C -c disk** command (see Figure 2.1).
3. Make a blank table (see Table 2.2 for a sample table) for recording the 9500V device data. The table must include the device file name, bus number, TID, LUN, and device type for each new device. Use the **lsattr -El hdiskx** command to view the TID and LUN for fibre disks if the location code does not include them.
4. Record the device information for all new devices in your device data table (see Table 2.2). You will need this information in order to change the device parameters.

```
# lsdev -C -c disk                                     ← Display device data.
hdisk0 Available 10-68-00-0,0 16 Bit SCSI Disk Drive   ← Hardware address.
hdisk1 Available 00-01-00 N/A                          ← New device.
hdisk2 Available 00-01-00-2,1 N/A                      ← New device.
  ↩ Device file name = hdiskx.
:
#
```

Figure 2.1 Verifying New Device Recognition

Table 2.2 Device Data Table (Sample)

Device File Name	Hardware Address
hdisk1	
hdisk2	
hdisk3	
hdisk4	
hdisk5	
hdisk6	
hdisk7	
hdisk8	
hdisk9	
and so on...	

Chapter 3 Configuring the New Devices

Configuration of the 9500V disk devices is performed by the user and requires root (superuser) access to the AIX® system. The host modes for the 9500V fibre ports must already be set, and the 9500V fibre ports must already be configured. If any of these steps is performed after the AIX® system is powered on, you must stop and restart the system or run `cfgmgr` before configuring the new devices.

Configuration of the 9500V SCSI disk devices for IBM® AIX® operations includes:

- Changing the disk device parameters (see section 3.1),
- Assigning the new disk devices to volume groups and setting the partition size (see section 3.2),
- Creating, mounting and verifying the Journaled File Systems (see section 3.3), and

Online device installation: For information on configuring newly installed 9500V devices without rebooting the AIX® system, please refer to Appendix C.

Note on the term “SCSI disk”: The 9500V logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

3.1 Changing the Device Parameters

3.1.1 Changing Device Parameters for the 9500V

Note: If you are using the Emulex Lightpulse™ LP8000® adapter, see section 3.1.2 for adapter-specific instructions

When the device files are created, the IBM® system sets the device parameters to the system default values. You must change the read/write (r/w) time-out, queue type, and queue depth parameters for each new 9500V device. Table 3.1 specifies the r/w time-out and queue type requirements for the 9500V devices as well as the queue depth requirements.

AIX® uses the Logical Volume Manager (LVM) (accessed from within SMIT®) to manage data storage. You can use either SMIT® or the AIX® command line to perform this procedure. Make sure to set the device parameters for all 9500V devices.

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

Parameter Name	Default Value	Required Value for 9500V
Read/write time-out	30	60
Queue type	Simple	Simple
Queue depth per LU ¹	1	≤ 32

Note 1: You can adjust the queue depth for the 9500V devices later as needed (within the specified range) to optimize the I/O performance of the devices.

Instructions for SMIT®

To change the device parameters using SMIT®:

1. Enter **smit** on the command line to start SMIT®.
2. On the SMIT® System Management screen, select **Devices** to bring up the Devices screen.
3. Select **Fixed Disk** to bring up the Fixed Disk screen.
4. Select **Change/Show Characteristics of a Disk** to go to the Disk screen.
5. Select the desired device from the **Disk** menu to bring up the Change/Show Characteristics of a Disk screen.
6. Enter the desired queue depth (refer to Table 3.1), queue type (simple), and read/write time-out value (60), and then press **Enter** to complete the parameter changes.
7. Repeat steps (5) and (6) for each new device on the 9500V subsystem.
8. Verify that the parameters for all devices were successfully changed using the **lsattr -E -l hdiskx** command.

Instructions for the AIX® Command Line

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

1. At the AIX® command line prompt, enter `lsattr -E -l hdiskx` to display the parameters for the specified device.
 - `'hdiskx'` is the device file name, (for example, `hdisk2`), or
 - use `smit/Devices/Fixed Disk/Change Show Characteristics of a Disk`.
2. Change the device parameters by using the `chdev -l hdiskx -a rw_timeout='60' -a q_type='simple' -a queue_depth='x'` command. (`x` is used to indicate the desired queue depth, within the limits specified in Table 3.1.)
3. Repeat steps (1) and (2) for each new device on the 9500V.
4. Verify that the parameters for all devices were successfully changed using either the `lsattr -E -l hdiskx` command.

Note: A list of commands is provided in Appendix D.

```
#lsattr -E -l hdiskx
scsi_id      0xef                SCSI ID
lun_id       0x0                 LUN ID
location     Location Label
ww_name      0x500490e802757500    FC World Wide Name for this LUN
pvid         000432871c6bbceb00000000000000000    Physical volume identifier
queue_depth  2                   Queue DEPTH
q_type       simple           Queuing TYPE
q_err        yes            Use QERR bit
clr_q        no             Device CLEARS its Queue on error
rw_timeout   60                   READ/WRITE time out value
start_timeout 60                   START unit time out value
reassign_to  120                  REASSIGN time out value
```

Figure 3.1 Verifying the Device Parameters Using the `lsattr -E -l hdiskx` Command

```
#lscfg -vl hdisk1
DEVICE      LOCATION      DESCRIPTION
hdisk1      20-58-01      Other FC SCSI Disk Drive
Manufacturer.....HITACHI
Machine Type and Model.....DF500           Type of device emulation
ROS Level and ID.....20202020
```

Figure 3.2 Verifying the Device Parameters Using the `lscfg -vl hdisk1` Command

3.1.2 Special Instructions for Using an Emulex Lightpulse™ LP8000® Adapter

When you are using an Emulex Lightpulse™ LP8000® host adapter, you must delete any existing 9500V disk devices, change the setting of the device parameters, then re-create the 9500V disk devices.

Instructions for the AIX® Command Line

1. Delete existing device parameters using the `rmdev -l hdiskx -d` command.
2. Change the LUN queue depth parameter by using the `chdev -l lpfcx -a lun_queue_depth='1'` command.
 - `lpfcx` is the host adapter's device file name, (e.g. `lpfc0`)
3. Change the linkdown parameter using the `chdev -l lpfcx -a linkdown_tmo='z'` command.
 - If you are using alternate path software, `z = 60`.
 - If you are not using alternate path software, `z = 180`.
4. Change the `nodev_holdio` parameter by using the `chdev -l lpfcx -a nodev_holdio='1'` command.
5. Repeat the above steps for each installed LP8000® adapter.
6. Verify that the parameters were successfully changed using the `lsattr -E -l lpfcx` command.
7. Since the devices (hdiskx) were deleted, run `cfmgr` to rediscover the hdiskx.

Figure 3.3 illustrates how to use the command line to display device parameters when you are using an Emulex Lightpulse™ LP8000® Adapter.

```
#lsattr -E -l
bus_intr_lvl          83          N/A False
intr_priority        4           N/A False
bus_mem_addr         0xa8000000  N/A False
bus_io_addr          0xffffc00   N/A False
dma_window           64M         N/A True
xmt_que_size         256         N/A True
topology              4           N/A True
ip_class              3           N/A True
fcp_class             3           N/A True
sli                   2           N/A True
network_on            0           N/A True
fcp_on                1           N/A True
num_iocbs             512         N/A True
num_bufs              512         N/A True
fcpfabric_tmo        0           N/A True
use_adisc             0           N/A True
first_check           1           N/A True
linkdown_tmo         180         N/A True
scan_down             1           N/A True
fabric_reg            1           N/A True
log_only              1           N/A True
log_verbose           0           N/A True
post_ip_buf           128         N/A True
ack0                  0           N/A True
automap               1           N/A True
no_device_delay       0           N/A True
lun_queue_depth       1           N/A True
tgt_queue_depth       0           N/A True
zone_rscn             0           N/A True
nodev_holdio         1           N/A True
check_cond_err        0           N/A True
delay_rsp_err         0           N/A True
rscn_adisc            0           N/A True
rscn_ns_delay         0           N/A True
dqfull_throttle       1           N/A True
```

Figure 3.3 Changing Device Parameters When Using an Emulex Lightpulse™ LP8000® Adapter (Command Line)

Figure 3.4 illustrates how to use SMIT® to change the device parameters for an Emulex Lightpulse™ LP8000® adapter.

```

                                Characteristics of a Fibre Channel Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[MORE...15]                                [Entry Fields]
  Extra FCP timeout for fabrics                [0]
  Use ADSIC instead of PDISC for discovery     [0]
  Linkdown Timeout for recovery             [180]
  Scandown ALPAs on Discovery                  [1]
  Issue Fabric Registration                    [1]
  Logging to log file only                     [1]
  Logging verbosity flag                       [0]
  Buffers to post to IP ring                   [128]
  ACK0 support                                 [0]
  Auto assignment of scsi ids                   [1]
  No FCP device failed I/O Usec delay          [0]
  Number of FCP commands per LUN           [1]
  Number of FCP commands per target           [0]
  Force NameServer check for RSCN DIDs         [0]
  Hold I/O errors if device disappears     [1]
  Delay for NameServer login on RSCN           [0]
  Treat check condition as FCP RSP err         [0]
  Treat FCP RSP err as no_device_delay        [0]
  Send ADISC instead of PLOGI for RSCN        [0]
  Apply change to DATABASE only                no
[BOTTOM]

F1=Help          F2=Refresh          F3=Cancel          F4=List
Esc+5=Reset      Esc+6=Command        Esc+7=Edit         Esc+8=Image
Esc+9=Shell      Esc+0=Exit           Enter=Do

```

Figure 3.4 Changing Device Parameters When Using an Emulex Lightpulse™ LP8000® Adapter (SMIT®)

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

After the device parameters have been changed, you can assign the new SCSI disk devices to new or existing volume groups. The partition size is also set during this procedure. Table 3.2 specifies the partition sizes for standard LUs.

Note: If you are configuring storage devices for databases that use a “raw” partition, do not assign those devices to volume groups.

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

1. At the AIX® command line prompt, enter **smit** to start SMIT®. This brings up the System Management screen. **Note:** If SMIT® is not installed, please refer to the IBM® AIX® user guide for instructions on assigning new devices to volume groups using AIX® commands.
2. Select **System Storage Management (Physical & Logical Storage)** to bring up the System Storage Management screen.
3. Select **Logical Volume Manager** to bring up the Logical Volume Manager screen.
4. Select **Volume Groups** to bring up the Volume Group screen.
5. Select **Add a Volume Group** to bring up the Add a Volume Group screen.
6. The Add a Volume Group panel (see Figure 3.5) allows you to assign one or more devices (physical volumes) to a new volume group and set the physical partition size. To assign one or more devices to a volume group and set the partition size:
 - a) Place the cursor in the **VOLUME GROUP name** entry field. Enter the name of the new volume group (e.g., 9200vg0). A volume group can contain multiple hdisk devices, depending on the application.
 - b) Place the cursor in the **Physical partition SIZE in megabytes** field, and press the **F4** key. When the size menu appears, select the correct partition size for the device(s).
 - c) Place the cursor in the **PHYSICAL VOLUME names** entry field. Enter the device file name(s) for the desired device(s) (e.g., hdisk1), or press **F4** and select the device file name(s) from the list.
 - d) Place the cursor in the **Activate volume group AUTOMATICALLY** entry field, and enter **yes** to activate the volume group automatically at system restart.
Note: If you are using HACMP, enter **no**.
7. After selecting the volume group, partition size, and physical volume(s) on the Add a Volume Group screen, press the **Enter** key.
8. When the confirmation screen opens, select **Yes** to assign the specified device(s) to the specified volume group with the specified partition size.
9. The Command Status screen now opens. To ensure that the devices have been assigned to a volume group, wait for **OK** to appear on the Command Status line. To continue creating volume groups, press **F3** until the Add a Volume Group screen appears.
10. Repeat steps (2) through (9) until all new disk devices are assigned to a volume group.

```

                                Add a Volume Group
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
VOLUME GROUP name                [9200vg0]          ← Enter volume group.
Physical partition SIZE in megabytes 4              ← Enter partition size.
PHYSICAL VOLUME names            [hdisk1]          ← Enter device file name(s).
Activate volume group AUTOMATICALLY yes           ← Enter no for HACMP.
  at system restart
Volume Group MAJOR NUMBER        []
*Create VG Concurrent Capable?
*Auto-varyon in Concurrent Mode?
F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit      F8=Image
F9=Shell    F10=Exit        Enter=Do

```

* These lines are added in AIX® version 4.3.

Figure 3.5 Assigning Devices to Volume Groups and Setting the Partition Size

Table 3.2 Partition Sizes for Disk Devices

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

3.3 Creating, Mounting, and Verifying the File Systems

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

Note: If you are configuring storage devices for databases that use a “raw” partition, you will create a logical volume only.

3.3.1 Creating the File System

1. At the AIX® command line prompt, enter **smit** to start SMIT®. This brings up the System Management screen. **Note:** If SMIT® is not installed, please refer to the IBM® AIX® user guide for instructions on creating file systems using AIX® commands.
2. Select **System Storage Management (Physical & Logical Storage)** to bring up the System Storage screen.
3. Select **File Systems** to bring up the File System screen.
4. Select **Add/Change/Show/Delete File Systems** to bring up the Add/Change screen.
5. Select **Journalled File Systems** to bring up the Journalled File System screen.
6. Select **Add a Standard Journalled File System** to bring up the Volume Group Name screen.
7. Move the cursor to the selected volume group, and press the **Enter** key.
8. Select the desired value, and then press the **Enter** key to bring up the Add a Journalled File System screen (see Figure 3.6).
9. Place the cursor in the **SIZE of file system** field, and enter the desired file system size (see Table 3.2).
10. Place the cursor in the **Mount Point** field, and enter the desired mount point name (e.g., /9200_VG00). Please record the mount point name and file system size. You will be asked to input this information again.
11. Place the cursor in the **Mount AUTOMATICALLY** field. Enter **yes** to auto-mount the file systems. **Note:** If you are using HACMP, do not set the file systems to auto-mount.
12. Place the cursor in the **Number of bytes per inode** field, and enter the correct value for the selected device.
13. Make sure that the file system size, mount point name, auto-mount options, and number of bytes per inode are correct, and press the **Enter** key to create the Journalled File System.
14. The Command Status screen now appears. To make sure that the Journalled File System has been created, wait for **OK** to appear on the Command Status line (see Figure 3.7).
15. Repeat steps (2) through (14) for each Journalled File System that you want to create. To continue creating Journalled File Systems press the **F3** key until you return to the Add a Journalled File System screen.
16. To exit SMIT®, press the **F10** key.

```

                                Add a Journaled File System
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
Volume group name                9200vg0
SIZE of file system (in 512-byte blocks) [4792320]
MOUNT POINT                      [/9200VG00]      ← Enter mount point name.
Mount AUTOMATICALLY at system restart? yes          ← Enter no for HACMP.
PERMISSIONS                      read/write
Mount OPTIONS                    []
Start Disk Accounting?           no
Fragment Size (bytes)           4096
Number of bytes per inode       4096      ← See Table 3.7..
Compression algorithm           no
*Allocation Group Size (Mbytes)
F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit      F8=Image
F9=Shell     F10=Exit       Enter=Do

```

* This line is added in AIX® version 4.3.

Figure 3.6 Adding a Journaled File System Using SMIT®

```

                                COMMAND STATUS
Command : OK          stdout : yes          stderr : no

      Before command completion, additional instructions may appear below.

Based on the parameters chosen, the new /9200VG00 JFS file system
is limited to a maximum size of 134217728 (512 byte blocks)
New Filesystems size is 4792320

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

```

Figure 3.7 Verifying Creation of Journaled File System

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

Calculate the file system size for these devices as follows:

1. Display the number of free physical partitions (FREE PPs) and physical partition size (PP SIZE) using the `lsvg` command (see Figure 3.8).
2. Calculate the maximum size of the file system as follows:
 $(\text{FREE PPs} - 1) \times (\text{PP SIZE}) \times 2048$

Example for a device is shown in Figure 3.8:

The maximum file system size is: $(733 - 1) \times (64) \times 2048 = 95944704$

```
# lsvg 9200vg0
VOLUMEGROUP: 9200vg0          VG IDENTIFIER: 0083665612e98521
VG STATE:      active          PP SIZE:      64 megabyte(s)
VG PERMISSION: read/write      TOTAL PPs:    733 (46912 megabytes)
MAX LVs:      256             FREE PPs:     733 (46912 megabytes)
LVs:          0               USED PPs:     0 (0 megabytes)
OPEN LVs:     0               QUORUM:       2
TOTAL PVs:    1               VG DESCRIPTORS: 2
STALE PVs:    0               STALE PPs     0
ACTIVE PVs:   1               AUTO ON:      yes
Concurrent:   Non-Capable     Auto-Concurrent: Disabled
VG Mode:      Non-Concurrent
```

Figure 3.8 Determining the Maximum File System Size

3.3.2 Mounting and Verifying the File Systems

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

To mount and verify the file systems:

1. At the AIX[®] command line prompt, enter: `mount <mount_point_name>` (e.g. `mount/9200_VG00`).
2. Repeat step (1) for each new file system.
3. Verify the size of the file systems you have created using the `df` command.
Note: The file system capacity is listed in 512-byte blocks by default. To list capacity in 1024-byte blocks, use the `df -k` command.
4. Verify that the new devices and file systems are fully operational by performing some basic operations (e.g., file creation, copying, deletion) on each device (see Figure 3.9).
5. At the next system restart, verify that the file systems have successfully auto-mounted by using the `mount` or `df` command to display all mounted file systems (see Figure 3.10). Any file systems that were not auto-mounted can be set to auto-mount using the SMIT[®] Change a Journaled File System screen. **Note:** If you are using HACMP, do not set the file systems to auto-mount.

```
# cd /9200VG00
# cp /smit.log /9200VG00/smit.log.back1
# ls -l 9200VG00
-rw-rw-rw- 1 root system 375982 Nov 30 17:25 smit.log.back1
# cp smit.log.back1 smit.log.back2
# ls -l
-rw-rw-rw- 1 root system 375982 Nov 30 17:25 smit.log.back1
-rw-rw-rw- 1 root system 375982 Nov 30 17:28 smit.log.back2
# rm smit.log.back1
# rm smit.log.back2
```

← Go to mount point.
 ← Copy file.
 ← Verify file copy.
 ← Copy file again.
 ← Verify copy again.
 ← Remove test file.
 ← Remove test file.

Figure 3.9 Verifying the Auto-Mounted File Systems

```
# df
File system 512-blocks free %Used Iused %Iused Mounted on
/dev/hd4 8192 3176 61% 652 31% /
/dev/hd2 1024000 551448 46% 6997 5% /usr
/dev/hd9var 8192 5512 32% 66 6% /var
/dev/hd3 24576 11608 52% 38 0% /tmp
/dev/hd1 8192 7840 4% 17 1% /home
/dev/lv00 4792320 4602128 4% 16 1% /9900VG00
/dev/lv01 4792320 4602128 4% 16 1% /9200VG01
/dev/lv02 14401536 13949392 4% 16 1% /9200VG02
```

← List mounted file systems.
 ← New filesystem.
 ← New filesystem.
 ← New filesystem.

Figure 3.10 Final File System Verification

Chapter 4 Troubleshooting

4.1 Troubleshooting

For troubleshooting information on the 9500V subsystem, please refer to the *Hitachi Freedom Storage™ Lightning 9900™ User and Reference Guide (MK-90RD008)*. Table 4.1 lists potential error conditions during 9500V IBM® AIX® configuration and provides instructions for resolving each condition. If you are unable to resolve an error condition, please ask your Hitachi Data Systems Customer Service representative for help, or call the Hitachi Data Systems Support Center for assistance.

Table 4.1 Troubleshooting

Error Condition	Recommended Action
The logical devices are not recognized by the system.	Make sure that the READY indicator lights on the 9500V subsystem are ON. The LUNs for each port must start at 0 and continue sequentially without skipping any numbers. RSD: LUN # need not start at 0 or contiguous. Run <code>cfgmgr</code> to recheck the fibre channel for new devices.

4.2 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, and the exact content of any messages displayed. The worldwide Hitachi Data Systems Technical Support Centers are:

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

Appendix A Acronyms and Abbreviations

AL	arbitrated loop
AL-PA	arbitrated loop physical address
FC	fibre channel
FCA	fibre channel adapter
FC-AL	fibre channel-arbitrated loop
FCP	fibre channel protocol
HACMP	High Availability Cluster Multi-Processing
HBA	host bus adapter
HRX	Hitachi RapidXchange
I/O, IO	input/output
IBM	International Business Machines Corporation
LU	logical unit
LUN	logical unit number
LUSE	LUN Expansion
LVM	Logical Volume Manager
PC	personal computer system
PCI	peripheral component interconnect
r/w	read/write
RISC	reduced-instruction-set computer
SCSI	small computer system interface
SMIT®	System Management Interface Tool
TID	target ID

Appendix B Online Installation and Deinstallation of Devices

After initial installation and configuration of the 9500V subsystem, additional devices can be installed or de-installed online without having to restart the system. After online installation, the disk device parameters for the new volumes/LUNs must be changed as specified in chapters 2 and 3.

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

Appendix C Commands for the AIX Command Line

Table C.1 lists the commands and their usage.

Table C.1 Commands and Usage

Commands	Use
<code>lsdev -Cc disk</code>	List all disks on the system.
<code>lscfg -vL hdiskx</code>	List VPD information for the disk.
<code>lsattr -El hdiskx</code>	List attributes of a disk device.
<code>lsattr -El lpfcx</code>	List attributes of the Emulex adapter.
<code>lsattr -El fcsx</code>	List attributes of the IBM adapter (can also use <code>lsattr -Elfscix</code>).
<code>Chdev -l hdiskx -a rw_timeout=60 -a q_type='simple' -a queue_depth='x'</code>	<i>Change the default attributes of the disk device</i>
<code>cfgmgr</code>	Discover devices added after IPL.
<code>errpt</code>	Show the AIX error log.
<code>Rmdev -dl hdiskx</code>	Remove a disk device logically from the system.
For Emulex LP8000	
<code>Chdev -l lpfcx -a lun_queue_depth='1' -a linkdown_tmo='z' -a nodev_holdio='1'</code> Where, Z = 60 if using alternate path software Z = 180 if not using alternate path software	Change the default attribute of the Emulex adapter.

