



**Hitachi Freedom Storage™
Thunder 9200™**

**HP-UX® Host Installation Guide
(SCSI)**

© 2001 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.

Hitachi Freedom Storage and Thunder 9200 are trademarks of Hitachi Data Systems.

HP and HP-UX are registered trademarks of Hewlett-Packard Company.

Microsoft and Windows NT are registered trademarks of Microsoft Corporation in the United States and/or 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.

All other brand or product names are or may be trademarks or service marks of 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-91DF542-0	June 2001	Initial Release

Source Document Revision Level

The following source document was used to produce this Thunder 9200 host installation guide:
Hitachi Disk Array Subsystem Installation Manual (SCSI), revision 2.

Preface

The *Hitachi Freedom Storage™ Thunder 9200™ HP-UX® Host Installation Guide* describes and provides instructions for configuring the devices on the Hitachi Thunder 9200™ disk array subsystem for operation with the HP-UX® 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 Hitachi Thunder 9200™ array subsystem,
- the user is familiar with the HP-UX® operating system and adapters, and
- the user is familiar with the UNIX® file system, system commands, and utilities.

Note: The term “9200” refers to the entire Hitachi Thunder 9200™ subsystem family, unless otherwise noted. Please refer to the *Hitachi Freedom Storage™ Thunder 9200™ User and Reference Guide* (MK-90DF504) for further information on the 9200 disk array subsystems.

Note: For further information on the HP-UX® operating system, please consult the HP-UX® user documentation, or contact HP® 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 Thunder 9200™ HP-UX® Configuration

1.1	HP-UX® Configuration	1
1.2	Hitachi Thunder 9200™ Array Subsystem	1

Chapter 2 Preparing for New Device Configuration

2.1	Configuration Requirements	3
2.2	Installing the 9200 Subsystem	4
2.3	Preparing for New Device Configuration	4
2.3.1	Setting Logical Units	4
2.3.2	Setting File System Sizes	4

Chapter 3 Configuring the New Devices

3.1	Verifying Device Files	5
3.2	Creating Volume Groups	6
3.3	Creating a File System	7
3.3.1	Recognizing LUs from all Ports	8
3.3.2	Installing the Host System through the Sub-host	9
3.4	Connecting the Dual Controller Multiple Ports Configuration	10
3.4.1	Verifying Recognition of the Subsystem from all Connected Ports	10
3.4.2	Setting an Alternate Link	11
3.5	Mounting and Verifying the File System	12
3.6	Setting and Verifying the Auto-Mount Parameters	12
3.7	Verifying the File Systems	13

Chapter 4 Troubleshooting

4.1	Troubleshooting	15
4.2	Calling the Support Center	15

Appendix A	Acronyms and Abbreviations	17
-------------------	---	-----------

List of Figures

Figure 3.1	Verifying the Device Files.....	5
Figure 3.2	Creating Volume Groups Creating Logical Volumes	6
Figure 3.3	Creating Logical Volumes.....	7
Figure 3.4	Creating a File System.....	7
Figure 3.5	Recognizing LUs in the Single Controller, Multiple Ports Configuration	8
Figure 3.6	Installing the Host System through the Sub-host	9
Figure 3.7	Dual Controller Multiple Ports: Verifying Subsystem Recognition.....	10
Figure 3.8	Setting Alternate Links.....	11
Figure 3.9	Creating a Directory for Mounting a File System	12
Figure 3.10	Checking the File System	12
Figure 3.11	Setting the Auto-Mount Parameters	12
Figure 3.12	Verifying the File System.....	13

List of Tables

Table 2.1	File Limitations	4
-----------	------------------------	---

Chapter 1 Overview of Thunder 9200™ HP-UX® Configuration

1.1 HP-UX® Configuration

This document describes the requirements and procedures for connecting the 9200 subsystem to an NT® server and configuring the new 9200 devices for operation with the HP-UX® server operating system. The Hitachi Data Systems representative performs the physical installation of the 9200 subsystem. The user prepares for 9200 subsystem installation, and then configures the new 9200 devices with assistance as needed from the Hitachi Data Systems representative.

Configuration of the 9200 disk devices for HP-UX® operations includes:

- Setting LUs and file sizes (see chapter 2),
- Installing the 9200 subsystem (see chapter 2),
- Creating and formatting volumes and groups (see chapter 3), and
- Verifying file system operations (see chapter 3).

1.2 Hitachi Thunder 9200™ Array Subsystem

The Hitachi Freedom Storage™ Thunder 9200™ 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 9200 subsystem provides continuous data availability, high-speed response, scaleable connectivity, and expandable capacity for PC server and open-system storage. The 9200 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.

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

The requirements for 9200 HP-UX[®] configuration are:

- Hitachi Thunder 9200[™] subsystem, all-open or multiplatform configuration:
 - The 9200 Resource Manager is used to configure the SCSI ports.

Note: The availability of 9200 features and devices depends on the level of microcode installed on the 9200 subsystem.

- HP-UX[®] operating system, version 10.20 or 11.0. **Important:** Please contact HP[®] technical support to make sure the most current OS patches are installed on the HP[®] system(s).
 - **root** log-in (superuser) access to the HP[®] system is required.

Note: Hitachi Data Systems plans to support future releases of the HP-UX[®] operating system. For the latest information on HP-UX[®] version support, please contact your Hitachi Data Systems account team.

2.2 Installing the 9200 Subsystem

The 9200 subsystem comes with all hardware and cabling required for installation.

Note: The Hitachi Data Systems representative must use the 9200 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 for New Device Configuration

Before setting up or installing the disk array subsystem, you should be familiar with the following activities:

- Setting logical units,
- Setting file system sizes, and
- Configuring the adapter(s).

2.3.1 Setting Logical Units

The maximum number of logical units (LUs) that can be set on a single disk array subsystem in HP-UX[®] is 8 for the Hitachi Freedom Storage[™] Thunder 9200[™].

2.3.2 Setting File System Sizes

See Table 2.1 for maximum file sizes.

Table 2.1 File Limitations

OS Version	Maximum Capacity
HP-UX 10.20	131072 MB (128 GB)
HP-UX 10.20	131072 MB (128 GB)
HP-UX 11.0	1 TB (JFS)

Chapter 3 Configuring the New Devices

Configuration of the new 9200 devices for HP-UX[®] operations involves the following activities:

- Verifying Device Files
- Creating Volume Groups,
- Creating logical volumes,
- Mounting and verifying the file systems, and
- Setting Auto Mount.

Note: Do not use the HP[®] System Administrator Manager (SAM) for system administrator functions. Alternate links cannot be created through SAM, and SAM will not recognize LUs that are recognized with the **ioscan** command.

3.1 Verifying Device Files

Use the **ioscan** command to verify the device files (see Figure 3.1).

# ioscan -nfc disk ↵							
Class	I	H/W Path	Driver	S/W State	H/W Type	Description	
=====							
disk	62	10/0.2.0	sdisk	CLAIMED	DEVICE	HITACHI	DFXXX
			/dev/dsk/c0t2d0	/dev/rdsk/c0t2d0			
disk	63	10/0.2.1	sdisk	CLAIMED	DEVICE	HITACHI	DFXXX
			/dev/dsk/c0t2d1	/dev/rdsk/c0t2d1			
disk	0	10/0.6.0	sdisk	CLAIMED	DEVICE	SEAGATE	ST34573WC
			/dev/dsk/c0t6d0	/dev/rdsk/c0t6d0			
disk	1	10/0.8.0	sdisk	CLAIMED	DEVICE	SEAGATE	ST34573WC
			/dev/dsk/c0t8d0	/dev/rdsk/c0t8d0			
disk	2	10/0.9.0	sdisk	CLAIMED	DEVICE	SEAGATE	ST34573WC
			/dev/dsk/c0t9d0	/dev/rdsk/c0t9d0			
disk	3	10/0.10.0	sdisk	CLAIMED	DEVICE	SEAGATE	ST34573WC
#							

Figure 3.1 Verifying the Device Files

3.2 Creating Volume Groups

Use the **pvcreate** command to create volume groups (see Figure 3.2).

```
# pvcreate /dev/rdisk/c7t0d0 ↵
Physical volume "/dev/rdisk/c7t0d0" has been successfully created.
# ls /dev/vg* ↵           ← Confirms the volume group.

/dev/vg00:
group   lvol2   lvol4   lvol6   lvol8   rlvol2   rlvol4   rlvol6   rlvol8
lvol1   lvol3   lvol5   lvol7   rlvol1   rlvol3   rlvol5   rlvol7

/dev/vg01:
group   lvol1   lvol2   lvol3   rlvol1   rlvol2   rlvol3

/dev/vg02:
group
# mkdir /dev/vg03 ↵      ← Creates a directory for volume group.

# ls -l /dev/vg* | grep group ↵
crw-r----- 1 root    sys      64 0x000000 Apr 18 22:51 group
crw-rw-rw- 1 root    sys      64 0x010000 Aug  3 12:51 group
crw-rw-rw- 1 root    sys      64 0x020000 Aug  3 12:02 group
# mknod /dev/vg03/group c 64 0x030000 ↵      ← Creates a control file.

# vgcreate /dev/vg03 /dev/dsk/c7t0d0 ↵      ← Creates a volume group.
Volume group "/dev/vg03" has been successfully created.
Volume Group configuration for /dev/vg03 has been saved in /etc/lvmconf/vg03.conf
#
# vgsdisplay /dev/vg03 ↵      ← Confirms the created volume group.
--- Volume groups ---
VG Name                /dev/vg03
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 0
Open LV                 0
Max PV                 16
Cur PV                 1
Act PV                  1
Max PE per PV          1016
VGDA                    2
PE Size (Mbytes)        4
Total PE                511
Alloc PE                0
Free PE                 511
Total PVG                0
Total Spare PVs         0
Total Spare PVs in use  0
#
```

Figure 3.2 Creating Volume Groups Creating Logical Volumes

You may divide volume groups into logical volumes. Use the **lvcreate** command (see Figure 3.3).

```
# lvcreate -L 2044 /dev/vg03 ↵
Logical volume "/dev/vg03/lvol1" has been successfully created with
character device "/dev/vg03/rlvol1".
Logical volume "/dev/vg03/lvol1" has been successfully extended.
Volume Group configuration for /dev/vg03 has been saved in /etc/lvmconf/vg03.conf
#
# lvdisplay /dev/vg03/lvol1 ↵
--- Logical volumes ---
LV Name                /dev/vg03/lvol1
VG Name                /dev/vg03
LV Permission          read/write
LV Status              available/syncd
Mirror copies          0
Consistency Recovery   MWC
Schedule               parallel
LV Size (Mbytes)       2044
Current LE             511
Allocated PE           511
Stripes                0
Stripe Size (Kbytes)   0
Bad block              on
Allocation             strict
IO Timeout (Seconds)   default
#
```

Figure 3.3 Creating Logical Volumes

3.3 Creating a File System

Create a file system using the **newfs** command (see Figure 3.4).

```
# newfs /dev/vg03/rlvol1 ↵
newfs: /etc/default/fs is used for determining the file system type
version 3 layout
2093056 sectors, 2093056 blocks of size 1024, log size 1024 blocks
unlimited inodes, 2093056 data blocks, 2091440 free data blocks
64 allocation units of 32768 blocks, 32768 data blocks
last allocation unit has 28672 data blocks
first allocation unit starts at block 0
overhead per allocation unit is 0 blocks
#
```

Figure 3.4 Creating a File System

3.3.1 Recognizing LUs from all Ports

Use the **ioscan** command to verify LUs from all connected ports (see Figure 3.5).

```
# ioscan -nfc disk ↵
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
=====						
disk	62	10/0.2.0	sdisk	CLAIMED	DEVICE	HITACHI DFXXX
			/dev/dsk/c0t2d0		/dev/rdisk/c0t2d0	
disk	63	10/0.2.1	sdisk	CLAIMED	DEVICE	HITACHI DFXXX
			/dev/dsk/c0t2d1		/dev/rdisk/c0t2d1	
disk	0	10/0.6.0	sdisk	CLAIMED	DEVICE	SEAGATE ST34573WC
			/dev/dsk/c0t6d0		/dev/rdisk/c0t6d0	
disk	1	10/0.8.0	sdisk	CLAIMED	DEVICE	SEAGATE ST34573WC
			/dev/dsk/c0t8d0		/dev/rdisk/c0t8d0	
disk	2	10/0.9.0	sdisk	CLAIMED	DEVICE	SEAGATE ST34573WC
			/dev/dsk/c0t9d0		/dev/rdisk/c0t9d0	
disk	3	10/0.10.0	sdisk	CLAIMED	DEVICE	SEAGATE ST34573WC
#						

Figure 3.5 Recognizing LUs in the Single Controller, Multiple Ports Configuration

3.3.2 Installing the Host System through the Sub-host

After making a normal installation from the master host, it is possible to use disks from the sub-host if you import Volume Groups so that disks are recognized from the sub-host. It should be noted, however, that data is not guaranteed if simultaneous access is made from hosts of the both systems. Basically, access from only one of the hosts is a must even if you execute an installation (import) from both hosts.

```
# pvcreate /dev/rdisk/c7t0d0 ↵
Physical volume "/dev/rdisk/c7t0d0" has been successfully created.
# ls /dev/vg* ↵

/dev/vg00:
group   lvol2   lvol4   lvol6   lvol8   rlvol2   rlvol4   rlvol6   rlvol8
lvol1   lvol3   lvol5   lvol7   rlvol1   rlvol3   rlvol5   rlvol7

/dev/vg01:
group   lvol1   lvol2   lvol3   rlvol1   rlvol2   rlvol3

/dev/vg02:
group

# mkdir /dev/vg03 ↵

# ls -l /dev/vg* | grep group ↵
crw-r-----  1 root      sys          64 0x000000 Apr 18 22:51 group
crw-rw-rw-   1 root      sys          64 0x010000 Aug  3 12:51 group
crw-rw-rw-   1 root      sys          64 0x020000 Aug  3 12:02 group
# mknod /dev/vg03/group c 64 0x030000 ↵

# vgcreate /dev/vg03 /dev/dsk/c7t0d0 ↵
Volume group "/dev/vg03" has been successfully created.
Volume Group configuration for /dev/vg03 has been saved in /etc/lvmconf/vg03.conf
#
# vgsdisplay /dev/vg03 ↵
--- Volume groups ---
VG Name                /dev/vg03
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 0
Open LV                 0
Max PV                  16
Cur PV                 1
Act PV                  1
Max PE per PV           1016
VGDA                    2
PE Size (Mbytes)        4
Total PE                511
Alloc PE                0
Free PE                 511
Total PVG                0
Total Spare PVs         0
Total Spare PVs in use   0
#
```

Figure 3.6 Installing the Host System through the Sub-host

3.4 Connecting the Dual Controller Multiple Ports Configuration

Avoid alternate paths between ports of the same controller and setting alternate paths between the controller.

3.4.1 Verifying Recognition of the Subsystem from all Connected Ports

Use the **ioscan** command to verify recognition of the subsystem (see Figure 3.7).

```
# ioscan -nfc disk ↵
Class      I  H/W Path      Driver      S/W State  H/W Type  Description
=====
disk       4  10/12/5.2.0  sdisk       CLAIMED    DEVICE    TOSHIBA CD-ROM XM-5701TA
                        /dev/dsk/c3t2d0  /dev/rdisk/c3t2d0
#
```

Figure 3.7 Dual Controller Multiple Ports: Verifying Subsystem Recognition

3.4.2 Setting an Alternate Link

Use the **vgextend** command to set alternate links (see Figure 3.8).

```
# vgextend vg03 /dev/dsk/c6t1d0 ↵
Volume group "vg03" has been successfully extended.
Volume Group configuration for /dev/vg03 has been saved in /etc/lvmconf/vg03.conf
# vgsdisplay -v vg03 ↵

--- Volume groups ---
VG Name                /dev/vg03
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 1
Open LV                 1
Max PV                 16
Cur PV                 1
Act PV                 1
Max PE per PV          1016
VGDA                    2
PE Size (Mbytes)        4
Total PE                511
Alloc PE                511
Free PE                 0
Total PVG                0
Total Spare PVs         0
Total Spare PVs in use  0

--- Logical volumes ---
LV Name                /dev/vg03/lvol1
LV Status               available/syncd
LV Size (Mbytes)        2044
Current LE              511
Allocated PE            511
Used PV                 1

--- Physical volumes ---
PV Name                /dev/dsk/c7t0d0
PV Name                /dev/dsk/c6t1d0 Alternate Link
PV Status               available
Total PE                511
Free PE                 0
Autoswitch              On
#
```

Figure 3.8 Setting Alternate Links

3.5 Mounting and Verifying the File System

Create a directory for mounting a file system using the **mkdir** command with the directory name as the argument (see Figure 3.9).

```
# mkdir /array1 ↵  
#
```

Figure 3.9 Creating a Directory for Mounting a File System

Check the file system using the **mount** command with the device file as the argument (see Figure 3.10).

```
# mount /dev/vg03/lvol1 /array1 ↵  
#
```

Figure 3.10 Checking the File System

3.6 Setting and Verifying the Auto-Mount Parameters

Use the **vi /etc/fstab** command to register the logical volumes that will perform Auto-Mount.

```
# vi /etc/fstab ↵  
# System /etc/fstab file. Static information about the file systems  
# See fstab(4) and sam(1M) for further details on configuring devices.  
/dev/vg00/lvol3      /          vxfs delaylog 0 1  
/dev/vg00/lvol1     /stand     hfs defaults  0 1  
/dev/vg00/lvol4     /home      vxfs delaylog 0 2  
/dev/vg00/lvol5     /opt       vxfs delaylog 0 2  
/dev/vg00/lvol6     /tmp       vxfs delaylog 0 2  
/dev/vg00/lvol7     /usr       vxfs delaylog 0 2  
/dev/vg00/lvol8     /var       vxfs delaylog 0 2  
/dev/vg01/lvol1     /home1     vxfs delaylog 0 2  
/dev/vg03/lvol1     /array1    vxfs delaylog 0 2
```

Figure 3.11 Setting the Auto-Mount Parameters

3.7 Verifying the File Systems

Verify the file system using the **bdf** command (see Figure 3.12).

```
# bdf ↵
Filesystem      kbytes  used   avail  %used  Mounted on
/dev/vg00/lvol3 307200 143283 153686 48%    /
/dev/vg00/lvol1 203413 50711  132360 28%    /stand
/dev/vg00/lvol8 819200 467392 329910 59%    /var
/dev/vg00/lvol7 921600 555094 343623 62%    /usr
/dev/vg00/lvol6 65536  1229   60348  2%     /tmp
/dev/vg00/lvol5 614400 135130 449361 23%    /opt
/dev/vg01/lvol1 4190208 1519545 2503787 38%    /home1
/dev/vg00/lvol4 204800 1214   190868 1%     /home
/dev/vg03/lvol1 2093056 1614   1960734 0%     /array1
```

Figure 3.12 Verifying the File System

Chapter 4 Troubleshooting

4.1 Troubleshooting

The Hitachi Freedom Storage™ Thunder 9200™ array subsystem provides continuous data availability. For troubleshooting information for the 9200 subsystem, please refer to the *Hitachi Freedom Storage™ Thunder 9200™ User and Reference Manual* (MK-90DF504).

4.2 Calling the 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).

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

Appendix A Acronyms and Abbreviations

FCA	FC adapter
HP	Hewlett-Packard
LU	logical unit
MB	megabyte(s)
PC	personal computer system
RAID	redundant array of independent disks
SAM [®]	System Administrator Manager [®]
SIM	service information message
TB	terabyte(s)

