HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Thunder 9200™

Compaq Tru64™ UNIX® Host Installation Guide (SCSI)

Document Revision Level

Revision	Date	Description
MK-91DF545-0	July 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™ Compaq Tru64™ UNIX® Host Installation Guide* describes and provides instructions for installing and configuring the devices on the Thunder 9200™ array subsystem for operation with the Tru64™ UNIX® 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 Freedom Storage™ 9200 array subsystem,
- the user is familiar with the Tru64™ UNIX® operating system.

Note: The term "9200" refers to the entire Hitachi Freedom Storage™ Thunder 9200™ subsystem family, unless otherwise noted.

Note: In this document, "Tru64™ " refers to the Compaq Tru64™ UNIX® operating system.

For further information on the 9200 array subsystem, please refer to the *Hitachi Freedom Storage™ Thunder 9200™ User and Reference Guide* (MK-90DF504), or contact your Hitachi Data Systems account team. The Hitachi Data Systems worldwide web site (http://www.hds.com) also provides information on the Hitachi Thunder 9200™ subsystem and its features and options.

For further information on Tru64™, please consult the Compaq user documentation, or contact Compaq 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

List of Figures

List of Tables

Chapter 1 Overview of 9200 Tru64™ UNIX® Configuration

1.1 Tru64™ UNIX® Configuration

This document describes the requirements and procedures for connecting the 9200 subsystem to a Tru64™ system and configuring the new 9200 devices for operation with the Tru64™ UNIX® operating system. The Hitachi Data Systems Customer Service representative performs the physical installation of the 9200 subsystem. The user prepares for 9200 subsystem installation and configures the new 9200 devices with assistance as needed from the Hitachi Data Systems representative.

Configuration of the 9200 SCSI disk devices for Tru64™ operations includes:

- Verifying new device recognition (see section 3.1),
- Registering Disk Geometry Registration (see section 3.2),
- Creating the device files (see section 3.3),
- Writing the partition labels (see section 3.5),
- Creating and verifying the file systems (see section 3.6), and
- Setting and verifying the auto-mount parameters (see section 3.9).

Note: For further information on the Compaq Tru64™ UNIX® operating system, please consult the Tru64™ UNIX® user documentation, or contact Compaq technical support.

1.2 Hitachi Thunder 9200™ Array Subsystem

The Thunder 9200™ array subsystem supports concurrent attachment to multiple PC server and UNIX®-based 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.

1.3 Device Types and Configuration Procedures

The 9200 subsystem allows the following types of logical devices (also called volumes) to be installed and configured for operation with the Tru64™ UNIX® operating system. Table 1.1 lists the device specifications for the 9200 devices.

Table 1.1 9200 Device Specifications

RAID Configuration *(Note 1)*	# of Blocks (for each device type) *(Note 2)*			# of Sectors per Track
	4.1 GB	8.7 GB	17.8 GB	
RAID0	40,652,160	85,085,440	174,681,600	96
RAID1/RAIDA	16,260,884	34,034,176	69,872,640	96
RAID5 (4D + 1P)	32,521,728	68,068,352	140,161,536	96

Note 1: The availability of a specific 9200 device type depends on the level of microcode installed on the 9200 subsystem.

Note 2: The number of blocks (by drive and row) is calculated as follows:

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

Note 3: The number of data cylinders is calculated as follows:

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

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

Example: For a RAID 1 with 5 with 17.8 GB drives formatted as one LU:

174,681,600 ÷ (16 x 96) = 113,723 cylinders (not used)

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

The requirements for 9200 Tru64™ configuration are:

- Hitachi Thunder 9200™ subsystem.
  - The 9200 Resource Manager is used to configure the SCSI ports.
- Tru64™ UNIX® server. ***Note:*** For information on server hardware requirements, please contact the Compaq® technical support and/or the Hitachi Data Systems Support Center.
- Tru64™ UNIX® operating system, versions 4.0F, 4.0G, and 5.0A. ***Important:*** Please contact Compaq® technical support to make sure the most current OS patches are installed on the Tru64™ UNIX® system(s).

  Note: Hitachi Data Systems plans to support future releases of the Tru64™ UNIX® OS. This document will be updated as needed to cover version-specific information. For further information on version support, please contact your Hitachi Data Systems account team.
  - **root** log-in (superuser) access to the Tru64™ UNIX® system is required.

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:

- Recognizing logical units
- Setting the partitions
- Setting the target IDs

2.3.1 Recognizing Logical Units

The 9200 Tru64™ operating system recognizes up to 8 Logical Units (LUs).

2.3.2 Setting the Partitions

The 9200 Tru64™ operating system creates up to 8 partitions per Logical Unit (LU).

2.3.3 Setting the Target IDs

The system recognizes subsystems set to target IDs 8 through 15.

Chapter 3 Configuring the New Devices

After 9200 subsystem installation and connection are complete, the devices on the 9200 subsystem are ready to be configured for use. Configuration of the 9200 devices is performed by the user and requires superuser (root) access to the Tru64™ UNIX® system. The activities involved in 9200 Tru64™ UNIX® configuration are:

- Verifying new device recognition (see section 3.1),
- Create the device files (see section 3.3),
- Writing the partition labels (see section 3.5),
- Creating and verifying the file systems (see section 3.6), and
- Setting and verifying the auto-mount parameters (see section 3.9).

3.1 Verifying New Device Recognition

After connecting the 9200 subsystem to the Tru64™ UNIX® system, you need to verify that the system recognizes the new devices. The devices must be installed and formatted and the fibre channel ports configured before the host system is powered on.

To verify that the Tru64™ system recognizes new devices, use the **scu** `show edit` command as shown in Figure 3.1.

Verifying the 9200 and the LUNs should also be done at the hardware boot level. To verify this way, use the `show` **device|more** command.

```
# scu show edt ↵

CAM Equipment Device Table (EDT) Information:

    Device: RZ28D   Bus: 3, Target: 0, Lun: 0, Type: Direct Access  ← Bus 3, TID 0, LUN 0
    Device: RZ28D   Bus: 3, Target: 0, Lun: 1, Type: (not present)
    Device: RZ28D   Bus: 3, Target: 0, Lun: 2, Type: (not present)
    Device: RZ28D   Bus: 3, Target: 0, Lun: 3, Type: (not present)
    Device: RZ28D   Bus: 3, Target: 0, Lun: 4, Type: (not present)
    Device: RZ28D   Bus: 3, Target: 0, Lun: 5, Type: (not present)
    Device: RZ28D   Bus: 3, Target: 0, Lun: 6, Type: (not present)
    Device: RZ28D   Bus: 3, Target: 0, Lun: 7, Type: (not present)
    Device: RRD45   Bus: 3, Target: 4, Lun: 0, Type: Read-Only Direct Access ←TID 4, LUN 0
    Device: RRD45   Bus: 3, Target: 4, Lun: 1, Type: (not present)
    Device: RRD45   Bus: 3, Target: 4, Lun: 2, Type: (not present)
    Device: RRD45   Bus: 3, Target: 4, Lun: 3, Type: (not present)
    Device: RRD45   Bus: 3, Target: 4, Lun: 4, Type: (not present)
    Device: RRD45   Bus: 3, Target: 4, Lun: 5, Type: (not present)
    Device: RRD45   Bus: 3, Target: 4, Lun: 6, Type: (not present)
    Device: RRD45   Bus: 3, Target: 4, Lun: 7, Type: (not present)
    Device: TLZ09   Bus: 3, Target: 5, Lun: 0, Type: Sequential Access      ← TID 5, LUN 0
    Device: TLZ09   Bus: 3, Target: 5, Lun: 1, Type: (not present)
    Device: TLZ09   Bus: 3, Target: 5, Lun: 2, Type: (not present)
    Device: TLZ09   Bus: 3, Target: 5, Lun: 3, Type: (not present)
    Device: TLZ09   Bus: 3, Target: 5, Lun: 4, Type: (not present)
    Device: TLZ09   Bus: 3, Target: 5, Lun: 5, Type: (not present)
    Device: TLZ09   Bus: 3, Target: 5, Lun: 6, Type: (not present)
    Device: TLZ09   Bus: 3, Target: 5, Lun: 7, Type: (not present)
    Device: DFXXX   Bus: 4, Target: 1, Lun: 0, Type:  irect Access
    Device: DFXXX   Bus: 4, Target: 1, Lun: 1, Type: (not present)
    Device: DFXXX   Bus: 4, Target: 1, Lun: 2, Type: (not present)
    Device: DFXXX   Bus: 4, Target: 1, Lun: 3, Type: (not present)
    Device: DFXXX   Bus: 4, Target: 1, Lun: 4, Type: (not present)
    Device: DFXXX   Bus: 4, Target: 1, Lun: 5, Type: (not present)
    Device: DFXXX   Bus: 4, Target: 1, Lun: 6, Type: (not present)
    Device: DFXXX   Bus: 4, Target: 1, Lun: 7, Type: (not present)
```

Figure 3.1 Verifying 9200 Device Recognition

3.2 Registering Disk Geometry Information

Register disk geometry information in the **/etc/disktab** file as shown in Figure 3.2.

```
# cp /etc/disktab /etc/disktab.backup                    ← Make a backup.
# vi /etc/disktab                                        ← Edit the file.

DFXXXF|HITACHI DFXXX:
      :ty=winchester:dt=SCSI:ns#16:nc#10248:rm#7200:     Partition a:  oa = block size
      :oa#0:pa#8388608:ba#8192:fa#1024:                  pa = start sector
      :ob#8388608:pb#2097152:bb#8192:fb#1024:            ba = fixed at 8192
      :oc#0:pc#15740928:bc#8192:fc#1024                  ta = fixed at 1024
      :od#10485760:pd#1048576:bd#8192:fd#1024:
      :oe#11534336:pe#1048576:be#8192:fe#1024:
      :of#12582912:pf#204800:bf#8192:ff#1024:
      :og#12787712:pg#204800:bg#8192:fg#1024:
      :oh#12992512:ph#2748416:bh#8192:fh#1024:
```

Figure 3.2 Registering the Disk Geometry Information

Calculate ***nc*** as follows:

$$nc = \frac{\text{Logical Unit Blocks}}{[ns]\ (16) \times [nt]\ (96)}$$

3.3 Creating the Device Files

Two types of device files are created with the Tru64™ UNIX® system; block-type and character-type.

3.3.1 Block-type Device Files

Create block-type device files in the /dev directory as shown in Figure 3.3.

```
# mknod /dev/rzXYZ b [major number] [minor number]
```

Figure 3.3 Creating Block-type Device Files

Where,

X = b through h = LUN 1 through LUN 7 (no letter is used for LUN 0)
Y = bus number × 8 + SCSI TID
Z = partition = a through h

Major number = 8
Minor number = bus number x16384 + TID × 1024 + [LUN] × 64 + [Partition]
(Partition = a through h)

3.3.2 Character-type Device Files

Create character-type device files in the /dev directory as shown in Figure 3.4.

```
# mknod /dev/rzXYZ c [major number] [minor number]
```

Figure 3.4 Creating Character-type Device Files

Where,

X = b through h = LUN 1 through LUN 7 (no letter is used for LUN 0)
Y = bus number × 8 + SCSI TID
Z = partition = a through h

Major number = 8
Minor number = bus number x16384 + TID × 1024 + [LUN] × 64 + [Partition]
(Partition = a through h)

3.4 Adding New Device Information to the /etc/ddr.dbase File

You need to add the 9200 device information to the device definition file (**/etc/ddr.dbase**) to enable SCSI command tag queuing for all new 9200 devices, including SCSI disk devices and HMDE (raw) devices.

To enable SCSI command tag queuing for the 9200 devices:

1. Before editing, make a backup copy of the **/etc/ddr.dbase** file.
2. Edit the **ddr.dbase** file (e.g., using **vi** editor) to add the 9200 device information. For **ddr_dbase** parameters, see **Error! Reference source not found.**.
3. Update and validate the contents of the edited **ddr.dbase** file: **ddr_config -c**
   It is not necessary to reboot.
4. Verify the new 9200 device information.

3.5 Writing the Partition Label

Use the **disklabel** command to label the partition for each LU. See Figure 3.5.

```
# disklabel-rw /dev/rrz33a DFXXX ↵
```

Figure 3.5 Writing the Partition Label

After writing the partition labels, confirm them by entering **disklabel -r /dev/rrz33a DFXXX** as shown in Figure 3.6.

```
# disklabel -r /dev/rrz33a ↵
# /dev/rrz33a:
type: SCSI
disk: DFXXX
label:
flags:
bytes/sector: 512
sectors/track: 96
tracks/cylinder: 16
sectors/cylinder: 520
cylinders: 10248
sectors/unit: 15744512
rpm: 7200
interleave: 1
trackskew: 8
cylinderskew: 16
headswitch: 0# milliseconds
track-to-track seek: 0      # milliseconds
drivedata: 0

8 partitions:
#        size    offset    fstype   [fsize bsize   cpg]
  a:  83886080             unused    0     0        # (Cyl.    0 - 16131*)
  b:   2097152   8388608   unused    0     0        # (Cyl. 16131*- 20164*)
  c: 157409280             unused    0     0        # (Cyl.    0 - 30271*)
  d:   1048576  10485760   unused    0     0        # (Cyl. 20164*- 22181*)
  e:   1048576  11534336   unused    0     0        # (Cyl. 22181*- 24197*)
  f:    204800  12582912   unused    0     0        # (Cyl. 24197*- 24591*)
  g:    204800  12787712   unused    0     0        # (Cyl. 24591*- 24985*)
  h:   2748416  12992512   unused    0     0        # (Cyl. 24985*- 30271*)
#
```

Figure 3.6 Confirming the Partition Label

3.6 Creating and Verifying the File Systems

Now you can create a file system for the new 9200 SCSI disk devices. Assign a character-type device file to the device file as shown in Figure 3.7.

```
# newfs /dev/rrz33a
Warning: 1024 sector(s) in last cylinder unallocated
/dev/rrz33a: 8388608 sectors in 5462 cylinders of 16 tracks, 96 sectors
     4096.0MB in 342 cyl groups (16 c/g, 12.00MB/g, 2880 i/g)
super-block backups (for fsck -b #) at:
 32, 24704, 49376, 74048, 98720, 123392, 148064, 172736,
 197408, 222080, 246752, 271424, 296096, 320768, 345440, 370112,
     :
#
```

Figure 3.7 Creating a File System

3.7 Creating the Mount Directories

After you have created the file systems, the next step in configuring the 9200 devices is to create a mount directory for each SCSI disk device. Make sure to choose a unique name for each mount directory, which identifies the device being mounted.

To create the mount directories for the new devices:

1. Create a mount directory by entering: **mkdir /<mount_directory_name>**. See Figure 3.8 for an example.

```
# mkdir/array1
```

Figure 3.8 Creating a Mount Directory

3.8 Mounting the File Systems

1. Mount the file system by entering:
   mount <device_file_name> <mount_directory>

   For example, to mount device **rz12c** with mount directory name **9200_LU2c**, enter:
   mount /dev/rz33a /9200_LU2c

3.9 Setting and Verifying the Auto-Mount Parameters

The final step in configuring the new 9200 devices is to set up and verify the auto-mount parameters for each new SCSI disk device. The **/etc/fstab** file contains the auto-mount parameters for the disk devices.

To set and verify the auto-mount parameters:

1. Make a backup copy of the **/etc/fstab** file before editing.
2. Edit the **/etc/fstab** file to add a line for each new device to be auto-mounted.
3. Table 3.1 shows the auto-mount parameters required for each device. When you are finished editing the **/etc/fstab** file, save your changes and exit.

   Note: If you make a mistake while editing, exit the editor without saving your changes, and then begin editing again.
4. The next time you reboot the Tru64™ UNIX® system, verify that the new devices were auto-mounted using the **df** (or **df -k**) command. If the system does not reboot properly, check auto-mount setting later, or use the **mount -a** command to check the **/etc/fstab** file.

```
# cp -ip /etc/fstab /etc/fstab.standard          ← Back up the file first.
# vi /etc/fstab                                  ← Edit the file.

①           ②       ③       ④       ⑤ ⑥        ← Refer to Table 3.2.
/dev/rz24a  /       ufs     rw      1  1
/proc       /proc   procfs  rw      0  0
/dev/rz24g  /usr    ufs     rw      1  2
/dev/rz33a  /array1 ufs     rw      1  3        ← Add these lines when registering.
/dev/rz33b  /array2 ufs     rw      1  3              "
/dev/rz33d  /array3 ufs     rw      1  3              "
/dev/rz33e  /array4 ufs     rw      1  3              "
/dev/rz33f  /array5 ufs     rw      1  3              "
/dev/rz33g  /array6 ufs     rw      1  3              "
/dev/rz33h  /array7 ufs     rw      1  3              "

: wq!                                        ← When done, press ESC, save and exit.
```

Figure 3.9 Setting and Verifying the Auto-Mounting Parameters

Table 3.1 Auto-Mount Parameters

Parameter #	Name	Enter
①	Filesystems	Block type device filename
②	Directory	Mount directory name
③	File system	Type of file system (e.g., UFS)
④	Mount options	Options (for example, rw for read-write)
⑤	Frequency dump in days	# of days (for example, 1, 2, 3)
⑥	Fsck pass	Order of performing file system checks

3.9.1 Verifying the File System

To verify the new file systems:

1. Verify that the new file systems were created correctly by displaying all mounted file systems. To display all mounted file systems, enter df.

   Note: The default display for drive capacity is 512-byte blocks. To view the capacity in kilobytes rather than in 512-byte blocks, enter: **`df –k`** (see Figure 3.10).

2. Verify the operation of each new file system as follows:

   a) Go to the new device directory by entering: **`cd /<mount_directory>`**
   For example, **`cd /9200_LU2c`**

   b) Copy a file from the root directory to the new device by entering:
   `cp /<file_name> <file_name>.back1`

   For example, to copy file **vmunix** from the root directory to the 9200_LU2c device, enter: **`cp /vmunix vmunix.back1`**

   c) Copy a file to the new device again. For example, to copy the same file again, enter:
   `cp /vmunix vmunix.back2`

   d) List the files in the current directory by entering: **`ls -l`**

   e) Remove the files you copied by entering: **`rm <file_name>`**
   For example, enter **`rm vmunix.back1`** to remove the file **`vmunix.back1`**.

3. Repeat step (2) for each new file system.

```
# df -k ↵
Filesystem  1024-blocks      Used    Available       Capacity                Mounted on
/dev/rz24a       121063     78803    30153          73%                     /
/proc                 0         0        0          100%                    /proc
/dev/rz24g       901214    716524    94568          89%                     /usr
/dev/rz33a      4065690         1  3659120          1%                      /array1
/dev/rz33b      1016222         1   914598          1%                      /array2
/dev/rz33d       508103         1   457291          1%                      /array3
/dev/rz33e       508103         1   457291          1%                      /array4
/dev/rz33f        98999         1    89098          1%                      /array5
/dev/rz33g        98999         1    89098          1%                      /array6
/dev/rz33h      1332078         1  1198869          1%                      /array7
#
```

Figure 3.10 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 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

CVS	custom volume size
FC	fibre channel
FCA	fibre channel adapter
LU	logical unit
LUN	logical unit number
MB	megabyte(s)
OFC	open file control
PC	personal computer system
SCSI	small computer system interface
TID	target ID