

PRELIMINARY RELEASE



**Hitachi Freedom Storage™
Thunder 9500™ V Series
Linux® Host Installation Guide**

PRELIMINARY RELEASE

© 2002 Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

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

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

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

Trademarks

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

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

Linux is a registered trademark of Linus Torvalds.

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, 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-92DF649-P	September 2002	Preliminary Release
MK-92DF649-P1	November 2002	Supercedes and replaces MK-92DF649-P

Source Documents for this Revision

The following source document was used to produce this Thunder 9500V host installation guide:

- *Hitachi Freedom Storage™ Thunder™ 9200 Red Hat® Linux® Host Installation Guide, (MK-91DF557)*
- HDS review of this document

Preface

The *Hitachi Freedom Storage™ Thunder 9500™ V Series Linux® Host Installation Guide* describes and provides instructions for installing and configuring the devices on the Thunder 9500V array subsystem for operation with the 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 Thunder 9500V array subsystem,
- the user is familiar with the Linux® operating system, including commands, utilities and file systems.

Note: The term “9500 V Series” 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 Linux®, please contact your Linux vendor’s 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!

PRELIMINARY RELEASE

Contents

Chapter 1	Overview of 9500V Linux® Configuration.....	1
1.1	Linux® Configuration	1
1.2	9500V Array Subsystem	1
Chapter 2	Preparing for New Device Configuration	3
2.1	Configuration Requirements	3
2.2	Installing the 9500V Subsystem.....	3
2.3	Preparing for New Device Configuration	3
2.3.1	Setting Logical Units	4
2.3.2	Verifying the Subsystem	5
Chapter 3	Configuring the 9500V Devices	7
3.1	Creating Partitions.....	7
3.2	Creating File Systems	9
3.3	Mounting and Verifying the File Systems.....	10
3.4	Setting and Verifying the Auto-Mount Parameters	12
3.5	Checking the File Systems.....	13
Chapter 4	Troubleshooting.....	15
4.1	Troubleshooting	15
4.2	Calling the Support Center	15
Appendix A	Acronyms and Abbreviations	17

List of Figures

Figure 2.1	Registering the Logical Units	4
Figure 2.2	Verifying New Device Recognition.....	5
Figure 3.1	Setting the Partitions.....	8
Figure 3.2	Creating the File Systems.....	9
Figure 3.3	Mounting and Verifying the File System	11
Figure 3.4	Setting and Verifying the Auto-Mounting Parameters	12
Figure 3.5	Checking the File Systems	13

List of Tables

Table 3.1	Auto-Mount Parameters	12
Table 4.1	Troubleshooting	15

Chapter 1 Overview of 9500V Linux® Configuration

1.1 Linux® Configuration

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

Configuration of the 9500V disk devices for Linux® operations includes:

- Installing the 9500V subsystem (see Chapter 3),
- Setting LUs and file sizes (see Chapter 2),
- Creating and formatting partitions (see Chapter 3),
- Setting Auto Mount (see Chapter 3), and
- Verifying file system operations (see Chapter 3).

1.2 9500V Array Subsystem

The 9500V 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, scalable 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.

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

The requirements for 9500V Linux® configuration are:

- Hitachi Freedom Storage™ Thunder 9500™ V Series subsystem.
 - The 9500V Resource Manager is used to configure the fibre ports.
- Linux® server. For information on server hardware requirements, please refer to the Linux® user documentation, or contact Linux® technical support.
- Any distribution of Linux® operating system, with kernel version 2.2.x or above.

Note: Hitachi Data Systems plans to support future releases of Linux®. This document will be updated as needed to cover version-specific information.
- Superuser (root) login access to the host system

2.2 Installing the 9500V Subsystem

The 9500V subsystem comes with all hardware and cabling required for installation.

Note: The Hitachi Data Systems 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 for New Device Configuration

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

- Setting logical units, and
- Verifying the subsystem.

2.3.1 Setting Logical Units

The maximum number of logical units (LUs) that can be set on a single disk array subsystem is 64 for the Hitachi Freedom Storage™ Thunder 9500™ V Series. Contact your Hitachi Data Systems representative for information

When using more than one logical unit, you must let the operating system know by registering them in the `/etc/lilo.conf` file to pass parameters that will adjust the behavior of the SCSI portion of the kernel (see Figure 2.1). The `lilo.conf` file is located under `/etc`. After editing `lilo.conf`, you will need to run it under `/sbin`. Type `/sbin lilo`.

```
# uname -r ↵          ← Check a version of the kernel.
2.2.16-22
#
# /sbin/mkinitrd /boot/newinitrd-img 2.2.16-22 ↵      ← Create the Ramdisk image.
#
# vi /etc/lilo.conf ↵
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
message=/boot/message
linear
default=linux

image=/boot/vmlinuz-2.2.16-22
label=old
initrd=/boot/initrd-2.2.16-22.img
read-only
root=/dev/sda5

image=/boot/vmlinuz-2.2.16-22          ← Add information to have IUs recognized.
label=linux                            "
initrd=/boot/newinitrd-img            "
read-only                             "
root=/dev/sda5                        "
append="max_scsi_luns=8"              "
#
# lilo -t ↵          ← Verify the lilo file.
Added old
Added linux *
The boot sector and the map file have *NOT* been altered.
#
# lilo ↵          ← Install the lilo file.
Added old
Added linux *
#
```

Figure 2.1 Registering the Logical Units

2.3.2 Verifying the Subsystem

The last step in preparing for 9500V device configuration is to verify that the host system recognizes the new devices. Hitachi Data Systems recommends that the devices already be installed and formatted before the host system is powered on. To verify new device recognition, see Figure 2.2.

```
# dmesg ↵
:
scsi : 0 hosts.
scsi : detected total.
md.c: sizeof(md_p_super_t) = 4096
Partition check:
 hda: hda1 hda2 < hda5 hda6 hda7 hda8 >
autodetecting RAID arrays
autorun ...
... autorun DONE.
VFS: Mounted root (ext2 filesystem) readonly.
Freeing unused kernel memory: 64k freed
Adding Swap: 298328k swap-space (priority -1)
Emulex LightPulse FC SCSI/IP 4.10g
!lpfc0:031:Link Up Event received Data: 1 1 1 2
scsi0 : Emulex LPFC (DRAGONFLY) SCSI on PCI bus 00 device 50 irq 9
scsi : 1 host.
 Vendor: HITACHI   Model: DF500F           Rev:      ← Verifies the Hitachi 9200 system.
 Type:   Direct-Access           ANSI SCSI revision: 03
Detected scsi disk sda at scsi0, channel 0, id 1, lun 0
 Vendor: HITACHI   Model: DF500F           Rev:
 Type:   Direct-Access           ANSI SCSI revision: 03
Detected scsi disk sdb at scsi0, channel 0, id 1, lun 1
 Vendor: HITACHI   Model: DF500F           Rev:
 Type:   Direct-Access           ANSI SCSI revision: 03
Detected scsi disk sdc at scsi0, channel 0, id 1, lun 2
 Vendor: HITACHI   Model: DF500F           Rev:
 Type:   Direct-Access           ANSI SCSI revision: 03
Detected scsi disk sdd at scsi0, channel 0, id 1, lun 3
 Vendor: HITACHI   Model: DF500F           Rev:
 Type:   Direct-Access           ANSI SCSI revision: 03
Detected scsi disk sde at scsi0, channel 0, id 1, lun 4
SCSI device sda: hdwr sector= 512 bytes. Sectors= 40960 [20 MB] [0.0 GB]
 sda: unknown partition table
SCSI device sdb: hdwr sector= 512 bytes. Sectors= 61440 [30 MB] [0.0 GB]
 sdb: unknown partition table
SCSI device sdc: hdwr sector= 512 bytes. Sectors= 81920 [40 MB] [0.0 GB]
 sdc: unknown partition table
SCSI device sdd: hdwr sector= 512 bytes. Sectors= 102400 [50 MB] [0.1 GB]
 sdd: unknown partition table
SCSI device sde: hdwr sector= 512 bytes. Sectors= 204920832 [100059 MB] [100.1 GB]
 sde: unknown partition table
SCSI device sde: hdwr sector= 512 bytes. Sectors= 204920832 [100059 MB] [100.1 GB]
 sde: sdel
SCSI device sde: hdwr sector= 512 bytes. Sectors= 204920832 [100059 MB] [100.1 GB]
 sde: sdel
```

Figure 2.2 Verifying New Device Recognition

Chapter 3 Configuring the 9500V Devices

Configuration of the new 9500V devices for Linux® operations involves the following activities:

- Creating and formatting partitions,
- Creating file systems
- Setting Auto-Mount

3.1 Creating Partitions

After the device parameters have been changed, you can set the partitions. In a Linux® environment, one LU can be divided into a maximum of four primary partitions. Instead of primary partitions, you can make a maximum of one extended partition. The extended partition can then be divided into multiple logical partitions of arbitrary size.

Use the `fdisk/dev/sdX` commands to set disk geometry information as shown in Figure 3.1.

PRELIMINARY RELEASE

```
# cfdisk /dev/sdb ↵
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.

Command (m for help): p ↵          ← Check the current status.

Disk /dev/sdb: 255 heads, 63 sectors, 653 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot   Start     End   Blocks  Id System
Command (m for help): n ↵          ← Start creating the partitions.
Command action
  e   extended
  p   primary partition (1-4)
p ↵                                  ← Create the primary partitions.
Partition number (1-4): 1 ↵        ← Specify partition number.
First cylinder (1-653, default 1): 1 ↵ ← Specify first cylinder.
Last cylinder or +size or +sizeM or +sizeK (1-653,default 653):200 ↵ ← Specify partition size.

Command (m for help): n ↵
Command action
  e   extended
  p   primary partition (1-4)      ← Create the extended partition.
e ↵
Partition number (1-4): 2 ↵
First cylinder (201-653, default 201): 201 ↵
Last cylinder or +size or +sizeM or +sizeK (201-653, default 653): 653 ↵

Command (m for help): n ↵
Command action
  l   logical (5 or over)
  p   primary partition (1-4)      ← Create logical partitions in the extended partition.
l ↵
First cylinder (201-653, default 201): 201 ↵
Last cylinder or +size or +sizeM or +sizeK (201-653, default 653): 653 ↵

Command (m for help): p ↵          ← Verify created partitions.

Disk /dev/sdb: 255 heads, 63 sectors, 653 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot   Start     End   Blocks  Id System
/dev/sdb1         1       200  1606468+  83  Linux
/dev/sdb2        201       653  3638722+   5  Extended
/dev/sdb5        201       653  3638691   83  Linux

Command (m for help): w ↵          ← Write the partition information.
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: If you have created or modified any DOS 6.x
partitions, please see the cfdisk manual page for additional
information.
Syncing disks.
#
```

Figure 3.1 Setting the Partitions

3.2 Creating File Systems

Use the `mke2fs /dev/sdX` command to create the file systems as shown in Figure 3.2.

```
# mke2fs /dev/sdb1 ↵
mke2fs 1.18, 11-Nov-1999 for EXT2 FS 0.5b, 95/08/09
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
200928 inodes, 401617 blocks
20080 blocks (5.00%) reserved for the super user
First data block=0
13 block groups
32768 blocks per group, 32768 fragments per group
15456 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912

Writing inode tables: 0/13_____ 1/13_____ 2/13_____ 3/13_____ 4/13_____ 5/13_____
6/13_____ 7/13_____ 8/13_____ 9/13_____10/13_____11/13_____12/13_____ done
Writing superblocks and filesystem accounting information: done
#
```

Figure 3.2 Creating the File Systems

3.3 Mounting and Verifying the File Systems

To mount the file system, use the `mkdir` *dir*.

After you have created the mount directories, you can mount and verify the file systems for the new SCSI disk devices. The file system for each logical partition should be mounted and verified to ensure that all new logical units are fully operational.

To mount and verify the file systems for the newly installed 9500V devices (see Figure 3.3).

1. Mount the file system using the **mount** command. Make sure to use the correct block-type device file name and mount directory for the device/partition. For example, to mount the file `/dev/sda1` with the mount directory `/9200_LU00` enter:

```
mount /dev/sda1 /9200_LU00
```

If you need to unmount a file system, use **umount** command (e.g., **umount /9200_LU00**).

Note: If you have already set the auto-mount parameters, you do not need to specify the block-type device file, only the mount directory.

2. Repeat step (1) for each partition of each newly installed SCSI disk device.
3. Display the mounted devices using the **df -h** command, and verify that all new SCSI disk devices are displayed correctly.
4. As a final verification, perform some basic UNIX® operations (e.g., file creation, copying, and deletion) on each LU to make sure that the new file systems are fully operational.

PRELIMINARY RELEASE

```

# mount /dev/sda1 /9200_LU00                                ← Mount file system.
                    ↖ Block-type device file name
# mount /dev/sda1 /9200_LU01                                ← Mount next file system.
                    ↖ Mount directory name
# mount /dev/sda1 /9200_LU02                                ← Mount next file system.
:
:
#
# df -h                                                    ← Display file systems.
File system      Kbytes    used    avail  capacity  Mounted on
/dev/dsk/c0t3d0s0  28775    27706     0    100%      /
/dev/dsk/c0t3d0s6 269191   234897   7384    97%      /usr
/proc              0         0         0     0%      /proc
fd                 0         0         0     0%      /dev/fd
/dev/dsk/c0t3d0s4s 57567    29515   22302    57%      /var
swap              142204    20    142184    0%      /tmp
/dev/dsk/c0t3d0s7 462119   206000  209909    50%      /export/home
/dev/dsk/c0t3d0s5  47975    42059    1126    97%      /opt
/dev/dsk/clt2d0s0 2256436     9  2030787    0%      /9200_LU00 ← Verify file systems.
/dev/dsk/clt2d1s0 2256436     9  2030787    0%      /9200_LU01 ← 9200 device
/dev/dsk/clt2d2s0 6774358     9  6548709    0%      /9200_LU02 ← 9200 device
:
# mount /dev/dsk/clt2d0s0 /9200_LU00                        ← Mount file system.
# cd /9200_LU00                                             ← Go to mount directory.
# cp /bin/vi /9200_LU00/vi.back1                            ← Copy a file.
# ls -l                                                     ← Verify the file copy.
drwxr-xr-t  2 root    root      8192 Mar 15 11:35  lost+found
-rwxr-xr-x  1 root    sys      2617344 Mar 15 11:41  vi.back1
# cp vi.back1 vi.back2                                    ← Copy file again.
# ls -l                                                     ← Verify file copy again.
drwxr-xr-t  2 root    root      8192 Mar 15 11:35  lost+found
-rwxr-xr-x  1 root    sys      2617344 Mar 15 11:41  vi.back1
-rwxr-xr-t  1 root    sys      2617344 Mar 15 11:52  vi.back2
# rm vi.back1                                             ← Remove test files.
# rm vi.back2                                             ←

```

Figure 3.3 Mounting and Verifying the File System

3.4 Setting and Verifying the Auto-Mount Parameters

The final step in configuring the new 9500V devices is to set up and verify the auto-mount parameters for each new 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 Linux® system, verify that the new devices were auto-mounted using the `df` (or `df -h`) command. If the system does not reboot properly, check auto-mount setting later, or check the `/etc/fstab` file by mounting using the desired mount point (for example, mount `/mnt/scsi/LU00`).

```
# vi /etc/fstab ↵
LABEL=/                /                ext2    defaults    1 1
LABEL=/boot            /boot            ext2    defaults    1 2

①                      ②                ③          ④                ⑤ ⑥ ← Refer to Table 3.1.
/dev/cdrom              /mnt/cdrom       iso9660  noauto,owner,ro 0 0
/dev/fd0                /mnt/floppy      auto    noauto,owner    0 0
none                   /proc            proc    defaults      0 0
none                   /dev/pts         devpts  gid=5,mode=620 0 0
/dev/sda6               swap             swap    defaults      0 0
/dev/sdb1               /array1          ext2    defaults      1 3 ← Refer to Table 3.1.
```

Figure 3.4 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 (for example, ext2, xfs, or reiserfs)
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

3.5 Checking the File Systems

Check the newly created file systems as shown in Figure 3.5.

```
# df -k
Filesystem          1k-blocks      Used    Available  Use%  Mounted on
/dev/sda5           6585928      937336    5314040   15%   /
/dev/sda2            23333         2949      19180    14%   /boot
/dev/sdb1           1581200         20    1500860    1%   /array1
#
```

Figure 3.5 Checking the File Systems

Chapter 4 Troubleshooting

4.1 Troubleshooting

The Hitachi Freedom Storage™ Thunder 9500™ V Series array subsystem provides continuous data availability. For troubleshooting information for the 9500V subsystem, please refer to the *Hitachi Freedom Storage™ Thunder 9500™ V Series User and Reference Guide* (MK-92DF601).

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). Please check the SVP service information messages (SIMs) using Web access, and note the reference codes and severity levels of the recent messages.

Table 4.1 lists potential error conditions during 9500V 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 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. Make sure that the fibre cables are correctly installed and firmly connected. Recheck the fibre buses for new devices (using diag command) For a listing of SCSI devices, type cat/proc/scsi .
File system cannot be created (newfs command)	Create the standard journalled filesystem for Linux® which consists of mkfs.ext2 /dev/sda1; tune2fs -j /dev/sda1 . Alternatively, use another filesystem that suits system needs such as mkfs.xfs or mkfs.reiserfs . Verify that logical unit is correctly labeled by with the dmesg command.
The file system is not mounted after rebooting.	Make sure that the system was restarted properly. Make sure that the file system attributes are correct. Verify that /etc/vfstab is correctly edited.
The 9500V red alarm light is on.	Please contact the Hitachi Data Systems Support Center.
System hangs, or devices are declared and then system hangs.	Make sure that the target IDs are set 0 through 6 and 8 through 15, and target ID 7 has been reserved for the SCSI controller card.

PRELIMINARY RELEASE

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

LUN	logical unit number
MB	megabytes
OFC	open fibre control
PC	personal computer system
RAID	redundant array of independent disks

