

*PRELIMINARY RELEASE*



**Hitachi Freedom Storage™  
Thunder 9500™ V Series**

**Compaq® Tru64™ UNIX® 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.

AlphaServer Series, Compaq, TruCluster, and Tru64 are registered trademarks or trademarks of Compaq Computer Corporation.

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-92DF645-P	September 2002	Preliminary Release

## Source Document Revision Level

The following source document was used to produce this 9500V configuration guide:

- *Hitachi Freedom Storage™ Thunder 9200™ Compaq® Tru64™ UNIX® Host Installation Guide, MK-90DF533*
- *Hitachi Freedom Storage™ Lightning 9900™ V Series Compaq® Tru64™ UNIX Configuration Guide, MK-90RD021*

## Preface

The *Hitachi Freedom Storage™ Thunder 9500™ V Series Compaq® Tru64™ UNIX® Host Installation Guide* describes and provides instructions for configuring the devices on the Thunder 9500™ V array subsystem for operation with the Tru64™ operating system. This document 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 Tru64™ UNIX® operating system (OS).

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

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

### COMMENTS

Please send us your comments on this document: [doc.comments@hds.com](mailto: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

<b>Chapter 1</b>	<b>Overview of 9500V Tru64™ UNIX® Configuration .....</b>	<b>1</b>
1.1	9500V Tru64™ UNIX® Configuration .....	1
1.2	9500V Array Subsystem .....	1
<b>Chapter 2</b>	<b>Preparing for New Device Configuration .....</b>	<b>3</b>
2.1	Configuration Requirements .....	3
2.2	Installing the 9500V Subsystem.....	4
2.3	Preparing to Connect the 9500V Subsystem .....	4
2.3.1	Setting the Host-Specific Parameters for the 9500V Ports.....	4
2.4	Connecting the 9500V Subsystem to the Tru64™ UNIX® System .....	6
2.5	Verifying New Device Recognition .....	7
2.6	Rebooting the Tru64™ System .....	8
<b>Chapter 3</b>	<b>Configuring the 9500V Devices .....</b>	<b>9</b>
3.1	Verifying the Device Files .....	10
3.2	Enabling Command Tag Queuing .....	12
3.3	Setting the I/O Timeout Value.....	13
3.4	Registering the Disk Types .....	15
3.5	Writing the Partition Labels.....	15
3.6	Creating the File Systems .....	16
3.6.1	UNIX® File System (UFS) .....	16
3.6.2	Advanced File System (AdvFS).....	16
3.7	Creating and Verifying the Mount Directories .....	17
3.8	Mounting and Verifying the File Systems.....	18
3.9	Setting and Verifying the Auto-Mount Parameters .....	20
<b>Chapter 4</b>	<b>Troubleshooting.....</b>	<b>21</b>
4.1	Calling the Hitachi Data Systems Support Center .....	22
<b>Appendix A</b>	<b>Acronyms and Abbreviations .....</b>	<b>23</b>
<b>Appendix B</b>	<b>Online Device Installation .....</b>	<b>25</b>

## List of Figures

Figure 2.1	Verifying 9500V Device Recognition (UNIX® prompt, Tru64™ UNIX® 5.0A and 5.1)	8
Figure 3.1	Changing the I/O Timeout Value .....	13
Figure 3.2	Reconfiguring the Kernel.....	14
Figure 3.3	Writing the Partition Labels for Tru64™ UNIX® 5.0A and 5.1 .....	15
Figure 3.4	Allocating Multiple Disk Partitions.....	16
Figure 3.5	Creating and Verifying the Mount Directories .....	17
Figure 3.6	Verifying the File Systems .....	19
Figure 3.7	Verifying File System Operation .....	19
Figure 3.8	Setting and Verifying the Auto-Mounting Parameters .....	20

## List of Tables

Table 2.1	Fibre Port Addressing.....	5
Table 3.1	Queue Depth Requirements for the 9500V Devices.....	12
Table 3.2	Auto-Mount Parameters .....	20
Table 4.1	Troubleshooting .....	21

# Chapter 1 Overview of 9500V Tru64™ UNIX® Configuration

## 1.1 9500V Tru64™ UNIX® Configuration

This document describes the requirements and procedures for connecting the 9500V subsystem to a Tru64™ system and configuring the new 9500V devices for operation with the Tru64™ operating system.

Configuration of the 9500V disk devices for Tru64™ operations includes:

- Verifying the device files (see section 3.1),
- Enabling command tag queuing (see section 3.1),
- Setting the I/O timeout value (see section 3.3),
- Registering the disk types (see section 3.4),
- Writing the partition labels (see section 3.5),
- Creating the file systems (see section 3.6),
- Creating and verifying the mount directories (see section 3.7),
- Mounting and verifying the file systems (see section 3.8), and
- Setting and verifying the auto-mount parameters (see section 3.9).

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

## 1.2 9500V Array Subsystem

The 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.

For further information on the 9500V subsystem, please refer to the *Hitachi Freedom Storage™ Thunder 9500™ V Series User and Reference Guide* (MK-92DF601), or contact your Hitachi Data Systems account team.



## Chapter 2 Preparing for New Device Configuration

### 2.1 Configuration Requirements

The requirements for 9500V Tru64™ configuration are:

- **9500V subsystem.**

- The Resource Manager 9500V software is required to configure the fibre-channel (FC) ports.

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

- **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, 5.0A, and 5.1.** **Important:** Please contact Compaq® technical support to make sure the most current OS patches are installed on the Tru64™ UNIX® system(s).

**Note:** For further information on supported Tru64™ versions, please contact Hitachi Data Systems.

- **High-availability (HA) software.** The 9500V currently supports the following software products. Please contact your Hitachi Data Systems account team for the latest information on supported software products.

- VERITAS® Cluster Server™ for host/application failover.
- Hitachi Dynamic Link Manager and VERITAS® Volume Manager™ for path failover and logical volume management.

- **Fibre-channel adapters.** Make sure to install all utilities, tools, and drivers that come with the adapter(s).

- The 9500V subsystem supports full-speed (1 and 2 Gb/s), shortwave, non-OFC (open fibre control) optical fibre-channel interface and multimode optical cables with SC and/or LC connectors. Do not connect any OFC-type fibre-channel interface to the 9500V subsystem.

## 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:

**Hardware installation.** Perform hardware installation as specified in Hitachi Data Systems' documentation and Hitachi, Ltd. source documentation. Follow all precautions and procedures in this documentation. Check all specifications to ensure proper installation and configuration. Hardware installation includes:

- Assembling all hardware and cabling.
- Upgrading to the latest microcode level.
- Creating RAID groups and LUNs and formatting LUNs using the Resource Manager 9500V software. For information and instructions on using Resource Manager, please refer to the *Hitachi Freedom Storage™ Thunder 9200™ Resource Manager User's Guide* (MK-91DF552).
- Installing the fibre-channel adapters and cabling.

**9500V FC 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.

## 2.3 Preparing to Connect the 9500V Subsystem

Before connecting the 9500V subsystem, you must perform the following tasks:

- Set the host-specific parameters for the 9500V fibre-channel port(s) (see section 2.3.1), and
- Verify host bus adapter installation

You will use the Resource Manager 9500V software to configure the 9500V fibre ports. For instructions on using Resource Manager, please refer to the *Hitachi Freedom Storage™ Thunder 9200™ Resource Manager User's Guide* (MK-91DF552).

### 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 (see Table 2.1). The 9500V subsystem supports up to 64 LUs. 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.

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

**PRELIMINARY RELEASE**

**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. See Table 2.1.

Table 2.1 shows the available AL-PA values and the corresponding SCSI TID address. 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/dsk/c*t*d*` and `/dev/rdisk/c*t*d*`. SCSI and fibre-channel devices are accessed the same way; however, the device files for SCSI and fibre-channel devices are configured differently.

**Table 2.1 Fibre Port Addressing**

AL-PA	TID														
EF	0	CD	16	B2	32	98	48	72	64	55	80	3A	96	23	112
E8	1	CC	17	B1	33	97	49	71	65	54	81	39	97	23	113
E4	2	CB	18	AE	34	90	50	6E	66	53	82	36	98	1F	114
E2	3	CA	19	AD	35	8F	51	6D	67	52	83	35	99	1E	115
E1	4	C9	20	AC	36	88	52	6C	68	51	84	34	100	1D	116
E0	5	C7	21	AB	37	84	53	6B	69	4E	85	33	101	1B	117
DC	6	C6	22	AA	38	82	54	6A	70	4D	86	32	102	18	118
DA	7	C5	23	A9	39	81	55	69	71	4C	87	31	103	17	119
D9	8	C3	24	A7	40	80	56	67	72	4B	88	2E	104	10	120
D6	9	BC	25	A6	41	7C	57	66	73	4A	89	2D	105	0F	121
D5	10	BA	26	A5	42	7A	58	65	74	49	90	2C	106	08	122
D4	11	B9	27	A3	43	79	59	63	75	47	91	2B	107	04	123
D3	12	B6	28	9F	44	76	60	5C	76	46	92	2A	108	02	124
D2	13	B5	29	9E	45	75	61	5A	77	45	93	29	109	01	125
D1	14	B4	30	9D	46	74	62	59	78	43	94	27	110		
CE	15	B3	31	9B	47	73	63	56	79	3C	95	26	111		

## 2.4 Connecting the 9500V Subsystem to the Tru64™ UNIX® System

After you have configured the 9500V fibre ports and host fibre-channel adapter(s), you are ready to connect the 9500V subsystem to the Tru64™ UNIX® system. The 9500V subsystem comes with all the hardware and cabling required for connection to the host system(s).

To connect the 9500V subsystem to the host system:

1. **Verify subsystem installation.** The Hitachi Data Systems representative verifies that the status of the fibre ports and LDEVs is normal. The Hitachi Data Systems representative should also check the fibre device parameters to make sure that all 9500V LDEVs are unique for each host system.
2. **Shut down and power off the host system.** The user should perform this activity. You must shut down and power off the host system before connecting the 9500V:
  - a) Shut down the Tru64™ UNIX® system.
  - a) When shutdown is complete, power off the Tru64™ UNIX® display.
  - b) Power off all peripheral devices except for the 9500V subsystem.
  - c) Power off the host system. You are now ready to connect the 9500V subsystem.
3. **Connect the 9500V to the host system.** The Hitachi Data Systems representative installs the fibre cables between the 9500V and the host system. **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.
4. **Power on and boot up the host system.** The user should perform this activity. To power on the host system after connecting the 9500V:
  - a) Power on the Tru64™ UNIX® display.
  - b) Power on all peripheral devices. The 9500V should already be on. The fibre-channel ports should already be configured (see section 2.3). If not, the host system may need to be restarted in order to recognize the new devices.
  - c) Confirm the ready status of all peripheral devices, including the 9500V subsystem.
  - d) Power on the Tru64™ UNIX® system. **Do not boot the system yet.** You may want to verify new device recognition from the boot prompt (see next section).

## 2.5 Verifying New Device Recognition

After connecting the 9500V 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. If not, the user must shut down and restart the system to allow the system to recognize the new devices.

To verify that the Tru64™ UNIX® system recognizes the new devices on the 9500V subsystem:

1. If desired, you can verify new 9500V fibre channel adapter recognition before booting the Tru64™ UNIX® system using the **show device | more** command. This listing of the device information is not as easy to read as the listing produced by the **scu show edt** command (see step 4). Make sure that all new fibre channel adapters are recognized by the Tru64™ UNIX® system.
2. Boot the Tru64™ UNIX® system as usual (e.g., **boot**).
3. Log in to the Tru64™ UNIX® system as **root**.
4. Display the device information using the **scu show edt** command (see Figure 2.1). Make sure that all new 9500V devices, are recognized by the Tru64™ UNIX® system.
5. Record the device file names. You will need the device file name information when you configure the new devices.

**Note:** The sample screens shown here may not be exactly the same for all Tru64™ UNIX® systems.

```
# scu show edt ↵

CAM Equipment Device Table (EDT) Information:

Device: RZ28D   Bus: 3, Target: 0, Lun: 0, Type: Direct Access
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
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
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: DFXXXXF Bus: 4, Target: 1, Lun: 0, Type: Direct Access
Device: DFXXXXF Bus: 4, Target: 1, Lun: 1, Type: (not present)
Device: DFXXXXF Bus: 4, Target: 1, Lun: 2, Type: (not present)
Device: DFXXXXF Bus: 4, Target: 1, Lun: 3, Type: (not present)
Device: DFXXXXF Bus: 4, Target: 1, Lun: 4, Type: (not present)
Device: DFXXXXF Bus: 4, Target: 1, Lun: 5, Type: (not present)
Device: DFXXXXF Bus: 4, Target: 1, Lun: 6, Type: (not present)
Device: DFXXXXF Bus: 4, Target: 1, Lun: 7, Type: (not present)

#
```

Figure 2.1 Verifying 9500V Device Recognition (UNIX® prompt, Tru64™ UNIX® 5.0A and 5.1)

## 2.6 Rebooting the Tru64™ System

After setting the disk and device parameters, setting and recognizing the LUs, and configuring the HBAs, you are ready to reboot the Tru64™ System.

To reboot the Tru64™ System, go to the command prompt, and enter: **reboot -- -r**  
 The -r option tells the system to rebuild the devices. Using **reboot** by itself will not build the devices on the newly installed 9500V subsystem.

After rebooting, you are now ready to configure the new LUs.

## Chapter 3      Configuring the 9500V Devices

After 9500V installation and connection are complete, the devices on the 9500V subsystem are ready to be configured for use. Configuration of the 9500V devices is performed by the user and requires root access to the Tru64® system. The activities involved in device configuration are:

- Verifying the device files (see section 3.1),
- Enabling command tag queuing (see section 3.1),
- Setting the I/O timeout value (see section 3.3),
- Registering the disk types (see section 3.4),
- Writing the partition labels (see section 3.5),
- Creating the file systems (see section 3.6),
- Creating and verifying the mount directories (see section 3.7),
- Mounting and verifying the file systems (see section 3.8), and
- Setting and verifying the auto-mount parameters (see section 3.9).

**Troubleshooting:** Chapter 4 provides troubleshooting information.

**Online device installation:** Appendix B provides instructions for online 9500V device installation.

### 3.1 Verifying the Device Files

The Tru64™ UNIX® system should create the device files for all new devices automatically during startup. Please make sure these files were created. When the device files are created, the system automatically reads the disk partitioning information directly from the disks.

The Tru64™ UNIX® device file naming convention is different than the device file naming conventions used by other UNIX®-based platforms. For Tru64™ UNIX® 4.0F/G, the character-type device file name has the form **rrzXYZ** (**rzXYZ** for block-type device file name), 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

For example: **rrzc18a** = bus 2, TID 2, LUN 2, partition a

**rrz12c** = bus 1, TID 4, LUN 0, partition c

For Tru64™ UNIX® 5.0A and 5.1, the character-type device file name has the form **/dev/rdisk/dskXY** (**/dev/disk/dskXY** for block-type device file name), where:

X = device instance number (assigned by the operating system)

Y = partition = a through h

For example: **/dev/rdisk/dsk8a** = disk 8 partition is a character device

**/dev/disk/dsk8a** = disk 8 partition is a character device

To verify that the device files for the new 9500V devices were successfully created:

1. Display the character-type device files. The device files are listed by bus number.
2. Make sure that character-type device files were generated for each new 9500V device. For example, for LUN 2 at TID 2 on bus 2, verify that **rrzc18(a-h)** are present. For Tru64™ UNIX® 5.0A, verify that a device file was made for LUN2 at TID 2 on bus 2.
3. Display the block-type device files.
4. Make sure that block-type device files were generated for each new 9500V device. For example, for LUN 2 at TID 2 on bus 2, make sure that **rzc18(a-h)** are present. For Tru64™ UNIX® 5.0A, verify that a device file was made for LUN2 at TID 2 on bus 2.

**PRELIMINARY RELEASE**

## 3.2 Enabling Command Tag Queuing

You need to add the 9500V device information to the device definition file (`/etc/DDR.dbase`) to enable SCSI command tag queuing for all new 9500V devices, including SCSI disk devices and HRX (raw) devices. Table 3.1 lists the queue depth requirements for 9500V devices.

**Table 3.1 Queue Depth Requirements for the 9500V Devices**

Parameter	Required Value
Queue depth per LU (TagQueueDepth)	<= 32 per LU  = 32 per LU (for Tru64 5.0A and 5.1)
Queue depth per port	≤ 256 per port

**Note:** 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 9500V devices for 4.0F/G.

To enable SCSI command tag queuing for the 9500V 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 9500V device information.
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 9500V device information. The device information you just entered should be displayed.

### 3.3 Setting the I/O Timeout Value

For Tru64™ UNIX® systems, the non-I/O timeout value is 10 seconds. You need to edit the `/sys/data/cam_data.c` file to change this value to 60 seconds. To set the I/O timeout value:

1. Edit the `/sys/data/cam_data.c` file (e.g., vi editor) to change `u_long cdisk-to-def` to **60** (see Figure 3.1).
2. After editing the `/sys/data/cam_data.c` file, reconfigure the kernel (see Figure 3.2).
3. For final verification, reboot the system (now or later), and make sure that the system operates normally with the newly reconfigured kernel. If the system does not reboot correctly, reboot using the backup copy of the kernel (**boot -fi vmunix.backup**), and check the edited files and entered values.

**Note:** Use caution when doing online volume installation. If you cannot reboot the system, check the reconfigured kernel later or use normal disruptive installation procedure.

```
# cp /sys/data/cam_data.c /sys/data/cam_data.c.backup          ← Back up the file.
# vi /sys/data/cam_data.c
:
:
/*
 * Changeable disk driver timeouts. Cdisk_to_def for non read/write
 * commands to the disks (test unit ready, mode select etc.)
 * cdisk_io_def Hard fixed disk timeout value for i/o, and
 * cdisk_io_rmb for removable media disks since they are slow.
 */
u_long cdisk_to_def = 60; /* 60 seconds */                    ← Change 10 [sec] to 60 [sec].
u_long cdisk_io_def = 60; /* 60 seconds (Tagged commands) */
u_long cdisk_io_rmb = 120; /* 120 seconds (slow removables) */
:
:
: wq                                                         ← Enter : wq! to save and exit.
#
```

Figure 3.1 Changing the I/O Timeout Value

## PRELIMINARY RELEASE

```
# cp /vmunix /vmunix.backup ← Make a backup copy first.
# doconfig -c DEC8400 ← Specify the config file (usually system name).
*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***

Saving /sys/conf/DEC8400 as /sys/conf/DEC8400.bak

Do you want to edit the configuration file? (y/n) [n]: y ← If 'n', go to ①.

Using /usr/dt/bin/dtpad to edit the configuration file. Press return when ready,
or type 'quit' to skip the editing session: ← If 'quit', go to ①.
Warning: locale not supported by C library, locale unchanged

*** PERFORMING KERNEL BUILD *** ① ← Start here if you
Working...Thu Feb 19 14:20:06 JST 1998 entered 'n' or 'quit' above.

The new kernel is /sys/DEC8400/vmunix
# mv /sys/DEC8400/vmunix /
#
```

**Figure 3.2** Reconfiguring the Kernel

### 3.4 Registering the Disk Types

Register the disk types of the 9500V subsystem. If the partition size or offset is different for each LU, you should register a different disk type name for each. Partition c must be allocated to all area of the LU.

### 3.5 Writing the Partition Labels

Use the **disklabel** command to label the partition for each new device (see Figure 3.3). Check that no errors are found in the partition settings after the labeling.

**Note:** You can also edit the disk partition size using the **disklabel** command with option **-e**. When **disklabel -e** command is executed, the system's editor opens. When done editing, save the file, and execute the **disklabel** command again to save your changes.

```
# disklabel -r /dev/rrz33a ↵
# /dev/rrz33a:
type: SCSI
disk: DFXXXF
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 3.3 Writing the Partition Labels for Tru64™ UNIX® 5.0A and 5.1

## 3.6 Creating the File Systems

the Advanced File System (AdvFS) to overcome the size and speed limitations of the UNIX<sup>®</sup> file system. The type of file system is application-dependent. If you are not sure which file system is right for your setup, contact Compaq<sup>®</sup> customer support.

**WARNING:** Do not create file systems for the multiplatform devices or for SCSI disk devices which will be accessed as raw devices.

### 3.6.1 UNIX<sup>®</sup> File System (UFS)

Versions 4.0F/G: `newfs <device_file_name>`

**Note:** Use the character-type device file name (rrzXYZ). For example, to create the file system for device `rrz33a`, enter: `newfs /dev/rrz33a`

### 3.6.2 Advanced File System (AdvFS)

To create an advanced file system:

1. You can either create a new AdvFS domain, or you can add a new fileset to an existing domain. To create a new domain: `mkfdmn <device_file_name> <domain_name>`
  - Versions 4.0F/G: Use the block-type device file name (rzXYZ). For example, to create an AdvFS domain called `dmn1` for device `rz12c`, enter: `mkfdmn /dev/rz12c dmn1`
  - Version 5.0A: `mkfdmn /dev/disk/dsk10c dmn1`

To add a new fileset to an existing domain: `mkfset <domain_name> <fileset_name>`  
For example, to create `fileset1` in `domain1`, enter: `mkfset domain1 fileset1`

2. Repeat steps (1) and (2) for each advanced file system you need to create.

**Note:** To allocate multiple disk partitions to a domain, the AdvFS must be installed. Installing the utilities will be able to allocate multiple disk partitions to the domain using the `addvol` command as shown in Figure 3.4. You can also do this work after completing mount operations as described later in this chapter.

```
# addvol /dev/rzb8c dmn1          ← Tru64™ UNIX® 4.0F/G
                               ↖ Block-type device file and domain name.
# addvol /dev/disk/dsk10c dmn1   ← Tru64™ UNIX® 5.0A and 5.1
#
```

Figure 3.4 Allocating Multiple Disk Partitions

### 3.7 Creating and Verifying the Mount Directories

After you have created the file systems, you can create mount directories for the new 9500V SCSI disk devices. Make sure to choose a unique name for each mount directory, which identifies the device being mounted.

**Note:** Do not create mount directories for the multiplatform devices (e.g., 3390-3A/B/C) or for devices which will be accessed as raw devices.

To create and verify the mount directories for the new devices (see Figure 3.5):

1. Use the **mkdir** command to create a new mount directory.
2. Use the **ls** command to verify the new mount directory.
3. Repeat steps (1) and (2) for each desired new mount directory.

```
# mkdir /9500V_LU2           ← Create new mount directory.
#ls -x                      ← Verify new mount directory.
.mrg ..DXsession          9500V_LU2c ← New mount directory appears.
CDROMDXclock DXprint
DXsessionMaila           bin
cdf3dev      df300      df300-2
df300-3df300-4          df300-5 etc
homellib      lnn
mdecmnt       opt      osf_boot
picassoproc   real.profile sbin
subsysys      tcb      tmp
usrvar        vmunix   vmunix.stdrd
```

**Figure 3.5** Creating and Verifying the Mount Directories

### 3.8 Mounting and Verifying the File Systems

After the file systems and mount directories have been created, you can mount each new 9500V SCSI disk device and verify the file system.

To mount a device with a UNIX<sup>®</sup> file system (UFS):

1. Mount the device by entering: `mount <device_file_name> <mount_directory>`

For example, to mount device `rz12c` with mount directory name `9200_LU2c`, enter:  
`mount /dev/rz12c /9200_LU2c (for 4.0F/G)`

For example, to mount device `dsk10c` with mount directory name `9200_LU2c`, enter:  
`mount /dev/disk/dsk10c /9200_LU2c (for 5.0A and 5.1)`

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

2. Assign the appropriate ownership and permissions by entering:  
`chown <owner>:<group> *<device_file_name>*`  
For example, to assign ownership to `rz12c` with owner `Oracle`, group `dba`, enter:  
`chown oracle:dba *rz12c*`
3. Repeat steps (1) and (2) for each new 9500V device with a UNIX<sup>®</sup> file system.

To mount a device with an advanced file system (AdvFS):

1. Mount the device by entering:  
`mount -t advfs <domain_name>#<fileset_name> <mount_directory>`  
For example, to mount the device with mount directory `9200_LU2c` enter:  
`mount -t advfs domain1#fileset1 /9200_LU2c`
2. Assign the appropriate ownership and permissions by entering:  
`chown <owner>:<group> *<device_file_name>*`  
For example, to assign ownership to `rz12c` with owner `Oracle`, group `dba`, enter:  
`chown oracle:dba *rz12c*`
3. Repeat steps (1) and (2) for each new 9500V device with an advanced file system.

After mounting all new 9500V devices, you need to verify that the new file systems were created correctly and are functioning properly. To verify the new file systems:

Display all mounted file systems (`df` command), and verify that the new file systems were created correctly (see Figure 3.6).

**PRELIMINARY RELEASE**

**Note:** To view capacity in kilobytes instead of 512-byte blocks, use the **df -k** command.

- a) Go to the new device directory: **cd /<mount\_directory>**
- b) Copy a file from the root directory to the new device:  
**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>**

1. 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 3.6** Verifying the File Systems

```
# cd /mnt      ← Go to mount directory.
# cp /vmunix vmunix.back1      ← Copy a file.
# ls -l      ← List files.
total 7240
-rwxr-xr-x  1  root  system  7397832  Nov 14 15:08  vmunix.back1
# cp /vmunix vmunix.back2      ← Copy a file again.
# ls -l      ← List files.
total 14480
-rwxr-xr-x  1  root  system  7397832  Nov 14 15:08  vmunix.back1
-rwxr-xr-x  1  root  system  7397832  Nov 14 15:12  vmunix.back2
# rm vmunix.back1 vmunix.back2      ← Remove test files.
#
```

**Figure 3.7** Verifying File System Operation

### 3.9 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 SCSI disk device. The `/etc/fstab` file contains the auto-mount parameters for the disk devices.

To set and verify the auto-mount parameters (see Figure 3.8):

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. Table 3.2 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.

3. 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
/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 3.8 Setting and Verifying the Auto-Mounting Parameters

Table 3.2 Auto-Mount Parameters

Parameter #	Name	Enter
①	Device to mount	Block type device filename
②	Mount point	Mount directory name
③	File system	Type of file system (e.g., UFS, AdvFS)
④	Mount options	Options (i.e., rw for read-write)
⑤	Frequency dump in days	# of days (e.g., 1, 2, 3)
⑥	File system check (fsck)	Order of performing file system checks

## Chapter 4 Troubleshooting

The 9500V disk array subsystems provide continuous data availability.

Table 4.1 lists potential error conditions during 9500V Tru64™ configuration and provides instructions for resolving each condition. If you are unable to resolve an error condition, please contact your Hitachi Data Systems representative or VAR 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 <b>diag</b> command). Verify the contents of <b>/kernel/drv/sd.conf</b> file.
File system cannot be created ( <b>newfs</b> command)	Check that character-type device file is specified for <b>newfs</b> command. Verify that logical unit is correctly labeled by UNIX® <b>format</b> 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 <b>/etc/vfstab</b> 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.
File system/superblock error	Run <b>fsck</b> .
Wrong magic number.	Label the disk using the <b>format</b> utility.

## **4.1 Calling the Hitachi Data Systems Support Center**

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

Run the getconfig script and send its resulting output file to Hitachi Data Systems when placing a support request call. Instructions are provided in Appendix E.

The worldwide Hitachi Data Systems Support Centers are:

- Hitachi Data Systems North America/Latin America  
San Diego, California, USA  
1-800-348-4357  
1-619-537-3000
- 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

AdvFS	advanced file system
AL-PA	arbitrated-loop physical address
CVS	custom volume size
ESCON®	Enterprise System Connection (IBM® trademark for optical channels)
FC	fibre channel
HRX	Hitachi RapidXchange
I/O, IO	input/output
LDEV	logical device
LU	logical unit
LUN	logical unit number
OFC	open fibre control
OS	operating system
PC	personal computer system
RAID	redundant array of independent disks
RCU	remote control unit (used for Hitachi TrueCopy operations)
r/w	read/write
SCSI	small computer system interface
SNMP	simple network management protocol
SVP	service processor
TID	target ID
UFS	UNIX® file system
WWN	worldwide name



## Appendix B Online Device Installation

This appendix provides instructions for online installation of new devices. After initial installation and configuration of the 9500V subsystem, additional devices can be installed or de-installed online without having to restart the Tru64™ system. After online installation, the device parameters for new volumes must be changed to match the LUs defined under the same fibre-channel port. This procedure should be performed by the system administrator (i.e., super-user).

When the TOV for 9500V volumes needs to be changed, you must edit the `/etc/system` file and reboot to make the file change effective. Please use the normal disruptive installation procedure for this case. Tru64™ must be booted once by the `boot -r` command with the TIDs and LUNs for new volumes defined in `/kernel/drv/sd.conf` file for non-disruptive volume installation. It is one way to prepare future online volume installation that whole TIDs and LUNs, i.e. TID=0~15 and LUN=0~7 are defined in the Tru64™ `kernel/drv/sd.conf` file and the system is once booted with `boot -r` command. In case of fibre channel, you should define LUN=0~119 in `/kernel/drv/sd.conf` file for the path (TID) which you will install additional LUNs.

1. The Tru64™ system must be booted once by `boot -r` command with the TIDs and LUNs for new volumes defined in `/kernel/drv/sd.conf` file for non-disruptive volume installation. If not, the Tru64™ system must be rebooted to recognize new volumes.
2. Check existing SCSI TIDs.
3. 9500V setup:
  - a) Ensure the latest microcode is loaded. Non-disruptive version up requires alternate path.
  - b) Install CHF and LUN, and connect fibre cable if necessary.
  - c) Execute online LU installation from the Resource Manager software, or Remote Console PC.
  - d) Verify 9500V SCSI path configuration.
4. Execute the following commands (as super-user) to recognize the new devices.
 

```
# /usr/sbin/drvconfig
# /usr/sbin/devlinks
# /usr/sbin/disks
# /usr/ucb/ucblinks
```
5. After online installation, the device parameters (for example, TOV, queue depth) for the new volumes must be changed to match the LUs defined under the same fibre-channel port.
6. Configure the new devices (partition, label, file system, auto-mount).

