

# **Hot-Swap Procedures Guide**

## **D Class**

**HP 9000 Enterprise Servers**



**Part No. A3262-90051**

**Edition 3  
August 1998**

**E0898**

Printed in: USA

---

## Legal Notices

The information in this document is subject to change without notice.

*Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.* Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

**Restricted Rights Legend.** Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY  
3000 Hanover Street  
Palo Alto, California 94304  
U.S.A.

**Copyright Notices.** ©copyright 1983-98 Hewlett-Packard Company, all rights reserved.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

**Trademark Notices** UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

---

# Contents

Safety and Regulatory Information .....	vi
Preface .....	vi
<b>1. Hot-Swap Disk Module for HP 9000 D Class Enterprise Servers</b>	
Overview: .....	1
Hot-Swap Disk Features .....	2
Hot-Swap Bay .....	2
LVM Patches .....	4
Hot Swap Example .....	4
Hot-Swap Procedure .....	5
Step 1 .....	5
Step 2 .....	6
Step 3 .....	7
Step 4 .....	7
Step 5 .....	7
Step 6 .....	8
Step 7 .....	8
Hot Swap Procedure for Unattached Physical Volumes .....	9
Step 1 .....	9
Step 2 .....	9
Step 3 .....	9
Step 4 .....	10
Step 5 .....	10
Troubleshooting .....	11

---

## Contents

### List of Figures

Figure 1-1 . Hot-Swap Disk Module . . . . .	1
Figure 1-2 . Hot-Swap Disk Bay . . . . .	3
Figure 1-3 . Configuration Switch . . . . .	11

### List of Tables

Table 1-1. Example Configuration . . . . .	4
--	---

---

## Printing History

The manual printing date and part number indicate its current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The manual part number will change when extensive changes are made.

Manual updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

First Edition: October 1995

Second Edition: November 1996

Third Edition: August 1998

---

### NOTE

**Reader Comments.** We welcome your comments about our documentation. If you have editorial suggestions or recommended improvements for this document, please write to us. You can reach us through e-mail at: *hardwaredocs@cup.hp.com* or by sending your letter to: *Documentation Manager M/S 5657, Hewlett-Packard Company, 8000 Foothills Blvd., Roseville, CA 95747-6588 USA*. Please include the following information in your message:

- Title of the manual you are referencing.
- Manual part number (from the title page).
- Edition number or publication date (from the title page).
- Your name.
- Your company's name.

**SERIOUS ERRORS**, such as technical inaccuracies that may render a program or a hardware device inoperative, should be reported to your HP Response Center or directly to a Support Engineer.

## Safety and Regulatory Information

For your protection, this product has been tested to various national and international regulations and standards. The scope of this regulatory testing includes electrical/mechanical safety, radio frequency interference, ergonomics, acoustics, and hazardous materials. Where required, approvals obtained from third-party test agencies are shown on the product label. In addition, various regulatory bodies require some of the information under the following headings.

---

**WARNING**

---

**The WARNING sign denotes a hazard. It calls attention to a procedure, practice, or the like, which if not done correctly or adhered to, could result in injury. Do not proceed beyond a WARNING sign until the indicated conditions are fully understood and met.**

---

**CAUTION**

---

The CAUTION sign denotes a hazard. It calls attention to an operating procedure, practice, or the like, which, if not done correctly or adhered to, could damage or destroy part or all of the product. Do not proceed beyond a CAUTION sign until the indicated conditions are fully understood and met.

## Preface

This edition of the *Hot-Swap Procedure's Guide* is intended for trained and experienced personnel.

This guide contains technical information about HP 9000 D Class Enterprise Server. At the time of publication, all models may have Hot-Swap Disk Modules installed.

---

# 1

## Hot-Swap Disk Module for HP 9000 D Class Enterprise Servers

---


### Overview:

The HP 9000 D Class Enterprise Server implements the use of the newly developed HP Hot-Swap, hard disk drives on a common Fast/Wide SCSI bus. The hot-swap feature implies the ability to remove or add an inactive disk drive module to a system while DC power is still applied and the SCSI bus is still active.

Two versions of the Hot-Swap chassis have been implemented. The D2XX systems will offer an optional Hot-Swap upgrade which will allow the low end servers to replace their single-ended SCSI disks with the Fast/Wide SCSI hot-swap disk modules. This upgrade will supply a Hot-Swap module that will support up to two hot-swap disk modules. The D3XX systems will come standard with a Hot-Swap module that supports up to five hot-swap disk modules.

**Figure 1-1**

### Hot-Swap Disk Modules



pshtpl2

The hot-swap disk modules used in the HP 9000 D Class Enterprise Server are low-profile disks that can have 1, 2, 4, or 9 gigabyte capacity. Each module contains a low-profile disk, EMI shielding, and an interface PCA.

## Hot-Swap Disk Features

- The delay start feature prevents large surge currents from occurring when several disk drives spin up at the same time. Spin delays are also necessary for hot-swap event detection when LVM mirroring is used. The delay spin signal is hard strapped active on all but the top position which is normally the "root" drive.
- The delay spin feature causes the spin-up to delay for a time approximately equal to ten times the SCSI address selected for a given drive in seconds. Delay spin can be overridden if the software attempts to access a disk drive before it has timed out and begun its spin up. The disk in the top position will spin-up as soon as power is applied.
- The remote spin feature gives the controlling software complete control over the disk drive spindle motor. Software can accordingly issue appropriate SCSI commands to spin-up or spin-down on the drive. The current HP UNIX software does not support this feature; therefore, it is hard disabled at the Hot-Swap backplane.

The module addresses are hard wired on the PCA for each position in the Hot-Swap module. A switch position on the backplane selects one of two address ranges.


---

## Hot-Swap Bay

The hot-swap bay is common to both the D3XX system package and the D2XX optional add on product.



**Figure 1-2**      **Hot-Swap Disk Bay**



---

## LVM Patches

Customers using manual hot swap on HP-UX 10.01, will need to pick up two LVM patches to install on their system. There will be two documents included in every D Class Enterprise Server informing the customer of the patch and the manual procedure.

With HP-UX 10.10 or later, the LVM patches will be included in the OS.

---

## Hot Swap Example

The following example describes a particular system problem where the solution is to replace a hot-swap disk module.

### Example configuration:

Volume group /dev/vg00 contains the three disks, with the logical volume configuration as shown:

**Table 1-1**

### Example Configuration

Volume Description	Volume Description	Volume Description
lvol 1	lvol 3	lvol 4
lvol 2	lvol 4	lvol 5
lvol 3	lvol 5	

```
hardware address      52.6.0          52.5.0          52.4.0
device file (/dev/dsk/) c2t6d0         c2t5d0         c2t4d0
```

The system problem for this example is that the disk at hardware address 52.4.0 has a head crash, and as a result it is unusable. The steps described in the *Hot-Swap Procedure* section of this document outline a method that can be used to recover from this state.

1. All of the replaced disk's in-use extents, must belong to mirrored logical volumes which were created with the 'strict' option (-s).

2. You must have an up to date configuration backup file. This is done by default each time an LVM command changes the LVM configuration.

The default backup file's path is:

**`/etc/lvmconf/base_vg_name.conf`**.

3. The replacement disk must be the same product ID as the replaced one.

---

**NOTE**

---

HP often uses different manufacturers for disks having the same product number. The hotswap manual procedure will not update the disk driver's internal information to that of the replaced disk.

The replacement disk will have the same capacity and blocksize as the defective disk because they have the same product number. The only field that could be incorrect is the string specifying the vendor's name. This will not affect the behavior of the LVM. If it is desired to update the manufacturers' name, then the disk's volume group must be deactivated and reactivated.

4. You must have previously installed the following two software patches:
  - PHKL\_5840, kernel (HP-UX version 10.01 only)
  - PHCO\_5841, utility (HP-UX version 10.01 only)

---

## Hot-Swap Procedure

Follow these steps to replace a hot-swap disk module.

### Step 1

- Check if the LVM found the physical volume to be defective when the volume group was activated.
- The *vgchange* would have printed the following message on the console:

```
WARNING:
```

```
VGCHANGE:WARNING: COULDN'T ATTACH TO THE VOLUME GROUP PHYSICAL
```

## Hot-Swap Disk Module for HP 9000 D Class Enterprise Servers

### Hot-Swap Procedure

```
VOLUME "/DEV/DSK/cXtXdX"
```

```
THE PATH OF THE PHYSICAL VOLUME REFERS TO A DEVICE THAT DOES NOT EXIST, OR IS NOT CONFIGURED INTO THE KERNEL.
```

- If the status of the *vgchange* is unknown, you may check if this occurred by doing a *vgdisplay* command:

```
# vgdisplay <VG name>
```

For our example:

```
# vgdisplay /dev/vg00
```

- If the disk was defective at *vgchange* time, the following messages will be printed one or more times.

```
WARNING:
```

```
VGDISPLAY: WARNING: COULDN'T QUERY PHYSICAL VOLUME  
                "/DEV/DSK/cXtXdX"  
THE SPECIFIED PATH DOES NOT CORRESPOND TO PHYSICAL VOLUME  
ATTACHED
```

```
TO THE VOLUME GROUP.
```

```
VGDISPLAY: WARNING: COULDN'T QUERY ALL OF THE PHYSICAL VOLUMES  
.
```

- If you see these messages, the disk was defective at the time the volume group was activated.

Hot swapping a disk which was defective during activation requires a different sequence of commands. Skip to the alternative procedure, *"Hot Swap Procedure for Unattached Physical Volumes,"* at the end of this document.

Otherwise, your disk became defective after the *vgchange* and you must continue with step 2 of this procedure.

## Step 2

- Reduce any logical volumes that have mirror copies on the faulty disk so that they no longer mirror onto that disk (note the *-A n* option):

```
# lvreduce -m 0 -A n <LV name> /dev/dsk/cXtXdX (for 1 way mirroring)
```

OR

```
# lvreduce -m 1 -A n <LV name> /dev/dsk/cXtXdX (for 2 way mirroring)
```

For our example:

```
# lvreduce -m 0 -A n /dev/vg00/lvol4 /dev/dsk/c2t4d0
```

```
# lvreduce -m 0 -A n /dev/vg00/lvol5 /dev/dsk/c2t4d0
```

### Step 3

- Replace the faulty disk. Please refer to the appropriate administration guide for instructions on how to replace the disk.
- Do an *ioscan* on the replaced disk to insure that it is accessible and also as a double check that it is a proper replacement. (See note under Step 3, on page 5.)

For our example:

```
# ioscan /dev/dsk/c2t4d
```

### Step 4

- Restore the LVM configuration/headers onto the replaced disk from your backup of the LVM configuration:

```
# vgcfgrestore -n <volume group name> /dev/rdisk/cXtXdX
```

- where X is the Logical unit number of the disk that has been replaced.

For our example:

```
# vgcfgrestore -n /dev/vg00 /dev/rdisk/c2t4d0
```

### Step 5

- Attach the new disk to the active volume group with the *vgchange* command.

```
# vgchange -a y <volume group name>
```

For our example:

```
# vgchange -a y /dev/vg00
```

## Step 6

- If this disk is not a mirror of a root disk, then skip this step.
- The *mkboot* command must be run.

For our example:

```
# mkboot /dev/rdisk/c2t4d0
```

- After running the *mkboot* command, do an *lvlnboot -R* to relink the replaced disk into the Boot Data Reserved Area of all the Physical Volumes in the Volume Group.

```
# lvlnboot -R
```

## Step 7

- *Lvextend* the mirrors back onto the replaced disk. This may take several minutes as it will have to copy all the data from the original copy of the data to the mirrored extents. The logical volume(s) are still accessible to users' applications during this command.

```
# lvextend -m 1 <LV name> /dev/dsk/cXtXdX
```

OR

```
# lvextend -m 2 LV name /dev/dsk/cXtXdX for 3-way mirroring
```

For our example:

```
# lvextend -m 1 /dev/vg00/lvol4 /dev/dsk/c2t4d0
```

```
# lvextend -m 1 /dev/vg00/lvol5 /dev/dsk/c2t4d0
```

At this stage, your system should be fully functioning.

## Hot Swap Procedure for Unattached Physical Volumes

Follow these steps to replace a hot-swap disk module for unattached physical volumes.

### Step 1

- Replace the faulty disk.
- Do an `ioscan` on the replaced disk to insure that it is accessible and also as a double check that it is a proper replacement. (See note under Step 3, on page 5.)

For our example:

```
# ioscan /dev/dsk/c2t4d0
```

### Step 2

- Restore the LVM configuration/headers onto the replaced disk from your backup of the LVM configuration:

```
# vgcfgrestore -n <volume group name> /dev/rdisk/cXtXdX
```

- where `X` is the Logical unit number of the disk that has been replaced.

For our example:

```
# vgcfgrestore -n /dev/vg00 /dev/rdisk/c2t4d0
```

### Step 3

- Attach the new disk to the active volume group with the `vgchange` command.

```
# vgchange -a y <volume group name>
```

For our example:

```
# vgchange -a y /dev/vg00
```

## Step 4

- If this disk is not a mirror of a root disk, then skip this step.
- The mkboot command must be run.

For our example:

```
# mkboot /dev/rdisk/c2t4d0
```

- After running the mkboot command, do an *lvlnboot -R* to relink the replaced disk into the Boot Data Reserved Area of all the Physical Volumes in the Volume Group.

```
# lvlnboot -R
```

## Step 5

- Resynchronize the mirrors of the replaced disk. This may take several minutes as it will have to copy all the data from the original copy of the data to the mirrored extents. The logical volume(s) are still accessible to users' applications during this command.

```
# vgsync <VG name>
```

For our example:

```
# vgsync /dev/vg00
```




---

## Troubleshooting

The hot-swap PCA contains a four position dip switch. It is set at the factory for disk mirroring. If the system doesn't respond, check to make sure the switch is set to disk mirroring.

**Figure 1-3** Configuration Switch



hsswitch.tif

### Mirrored Root Switch Position:

The first position is labeled MIRRORED ROOT / NON-MIRRORED ROOT. When set to mirrored root, the top most disk module will exhibit a normal spin delay of either 40 or 50 seconds, dependent on the root address setting. Spin delays are necessary for modules included in LVM mirroring. The non-mirrored root setting will allow the top most disk to begin spin-up at power on. The default setting is mirrored root.

### Reset Switch Position:

The second position is labeled RESET DISABLED / RESET ENABLED. This switch determines if the hot-swap event reset circuitry is active or inactive. The default setting is reset enabled.

**Root Address Switch Position:**

The third position is labeled ROOT ADDR.5 / ROOT ADDR.4. This switch selects the SCSI address of the top most disk module (root disk) independent of the other disk module locations. The default setting is root addr.5.

**Address Switch Position:**

The fourth position is labeled ADDR.12-15 / ADDR.8-11. The lower four disk module positions in a Model D3xx system will either be 12, 13,14, and 15, or 8, 9, 01, and 11. In the Model D2xx system optional hot-swap bay, the second disk position will either be 12 or 8. The default setting is addr. 8-11.