

December 1998  
ECG001/1298

Prepared by High Availability  
Products Group

Compaq Computer Corporation

## Contents

### Managing Changes to Drive

<b>Ordering .....</b>	<b>3</b>
Symptoms .....	3
When does this situation occur? .....	3
Preventing the Issue .....	4
Simplifying Detection of the Issue .....	4
What Oracle Component is Affected? .....	4

### Removing a Shared Storage Array .....

### Adding a Shared Storage

### Array .....

### Adding or Removing Drives

<b>from a Shared Storage Array ...</b>	<b>11</b>
Replacing a Failed Drive .....	11
Adding a Drive to Increase Storage Capacity .....	12

# Modifying Physical Cluster Resources in a Compaq Parallel Database Cluster

**Abstract:** At some time during the life of a Compaq Parallel Database Cluster it is likely that some modifications to the physical resources of the cluster will be required. This paper describes how to perform some of the more likely modifications, such as:

- Managing Changes to Drive Ordering
- Removing a Shared Storage Array
- Adding a Shared Storage Array
- Adding or Removing Shared Storage Drives

Additional scenarios are described in Chapter 5 of the *Compaq Parallel Database Cluster Administrator Guide*.

## Notice

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, registered United States Patent and Trademark Office.

Netelligent, Armada, Cruiser, Concerto, QuickChoice, ProSignia, Systempro/XL, Net1, LTE Elite, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, Wonder Tools logo in black/white and color, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows, Windows NT, Windows NT Server and Workstation, Microsoft SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

Oracle is a registered trademark and Oracle8 is a trademark of Oracle Corporation

NetWare and Novell are registered trademarks and intraNetWare, NDS, and Novell Directory Services are trademarks of Novell, Inc.

Pentium is a registered trademark of Intel Corporation.

Copyright ©1998 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Modifying Physical Cluster Resources in a Compaq Parallel Database Cluster  
White Paper prepared by High Availability Products Group

First Edition (December 1998)  
Document Number ECG001/1298

## Managing Changes to Drive Ordering

The Compaq Parallel Database Cluster encounters difficulties when the order in which drives are brought online changes. For the drives associated with the fibre channel arrays, NT assigns disk numbers based on the order in which the arrays are powered on. Oracle uses these disk number assignments via the links created with the SETLINKS utility. The order in which drives are brought online directly affects the disk numbers assigned to the disks, which in turn affects whether Oracle8 Server can find its database files.

### Symptoms

The primary symptom of this issue is that the database will not start up correctly. Some other symptoms that might be seen are:

- The user may see Oracle errors when trying to mount the database.
- The database will mount correctly, however, the data will not be correct.
- The PGMS service will not start.
- When using Windows NT disk administrator to see the disk configuration:
  - one or more of the disks associated with the shared storage will not be seen or will be marked off-line; or
  - the disks will be shown in the wrong order

### When does this situation occur?

The situation occurs when the order in which the database's drives are brought online is different than the order in which the drives were brought online when the database cluster was originally set up. Following are some specific examples:

- When more than one shared storage array exists, if the arrays are powered on in a different order than they were when the cluster was initially configured, the order of the shared drives will change for all cluster nodes.
- When a drive is added, internally, to a cluster node, the order of the shared drives will change for that node.
- When an internal drive is removed from a cluster node, the order of the shared drives will change for that node.
- When a drive is added to a shared storage array, if the drive is added anywhere other than the last drive in the last array, the order of the drives will change for all cluster nodes.
- When an existing drive is removed from a shared storage array, if the drive is removed from anywhere other than the last drive in the last array, the order of the drives will change for all cluster nodes.
- When an array is added to the cluster's shared storage, if the array is added anywhere other than the last in the sequence of arrays, the order of the drives will change for all cluster nodes.

- When an existing array is removed from the cluster's shared storage, if the array was not the last one brought online, the order of the drives will change for all cluster nodes.

## Preventing the Issue

It is highly recommended that a power-on sequence for the shared storage arrays be determined during the initial configuration of the cluster. Thereafter, this power-on sequence must be followed exactly to ensure that each node sees the disk in a proper order in the Windows NT Disk Administrator utility. When adding shared storage (a physical drive or an array), make it the last drive (or array) to be brought online. In the case of adding an array, make sure it is the last array to be powered-on; and add it to the power-on sequence as such.

On each node use disk administrator to verify that the disks are brought on-line in the expected order before starting the Oracle instance. Also, verify that the PGMS service is set to start manually.

## Simplifying Detection of the Issue

A method that simplifies detection of the issue is to create logical partitions that are not all the same size. With different-sized logical partitions, disk administrator can be used to see the size of the partitions, which can be quickly and easily compared against the sizes written in the suggested Oracle worksheet. If the drive order in the worksheet does not compare with the drive order shown in disk administrator, the drives were brought online out of sequence.

For example, the shared disks associated with the database could have a small unused partition as the first partition on the disk. Assuming there are three disks associated with the database, a sample configuration might be:

- Harddisk1 Partition1 is sized to 11 MB
- Harddisk2 Partition1 is sized to 12 MB
- Harddisk3 Partition1 is sized to 13 MB

When the cluster is powered on, if disk administrator shows harddisk1, partition1 as 12MB or 13MB, the drives were brought on-line out of sequence. The cluster should be shut down and restarted correctly.

## What Oracle Component is Affected?

How does the drive ordering issue manifest itself in Oracle8 Server? The focal point of this issue resides in the Oracle symbolic link table files; frequently identified as ORALINKx.TBL. These files create symbolic links, which are used by Oracle8 Server to map specific Oracle database files to specific hard drive partitions. A link is required for each database file. A two node cluster requires nine database files and one cluster file, and therefore requires ten links. The first symbolic link table file is used by a two-node cluster. Each additional node requires two more database files, two more links, and one more symbolic link table file .

Oracle suggests using a worksheet to define the symbolic links. For more information about the worksheet, refer to the *Oracle Parallel Server Getting Started Release 8.0.5 for Windows NT manual*. An example of such a worksheet follows. The example assumes the cluster consists of two nodes – Node1 has two internal disk drives and Node2 has one internal disk drive.

Additionally, it assumes the shared storage consists of two fibre channel arrays - the first array has three RAID logical drives and the second array has two RAID logical drives. Therefore,

Node1 represents the drives in the first array as Harddisk2, Harddisk3, and Harddisk4, while Node2 represents them as Harddisk1, Harddisk2, and Harddisk3.

Thus, the configuration of the cluster as seen by disk administrator is:

Node1	Disk 0	An internal disk
	Disk 1	An internal disk
	Disk 2	Shared storage disk
	Disk 3	Shared storage disk
	Disk 4	Shared storage disk
	Disk 5	Shared storage disk
	Disk 6	Shared storage disk
Node2	Disk 0	An internal disk
	Disk 1	Shared storage disk
	Disk 2	Shared storage disk
	Disk 3	Shared storage disk
	Disk 4	Shared storage disk
	Disk 5	Shared storage disk

A symbolic link worksheet corresponding to the above configuration is in the following table. Note that partition1 on each disk is used for the small, identifiable partitions discussed in the *Simplifying Detection of the Issue* section above; and therefore, partition1 is not shown in the symbolic links worksheet.

**Table 1. Original Oracle Symbolic Links Worksheet**

Symbolic Link	Node1	Node2
OPS_log1t1	Harddisk2 Partition2	Harddisk1 Partition2
OPS_log2t1	Harddisk2 Partition3	Harddisk1 Partition3
OPS_sys01	Harddisk2 Partition4	Harddisk1 Partition4
OPS_usr01	Harddisk3 Partition2	Harddisk2 Partition2
OPS_rbs01	Harddisk3 Partition3	Harddisk2 Partition3
OPS_tmp01	Harddisk3 Partition4	Harddisk2 Partition4
OPS_log1t2	Harddisk4 Partition2	Harddisk3 Partition2
OPS_log2t2	Harddisk4 Partition3	Harddisk3 Partition3
OPS_cntr01	Harddisk5 Partition2	Harddisk4 Partition2
OPS_cmdisk	Harddisk6 Partition2	Harddisk5 Partition2

## What happens when the drive numbering order changes?

When the numbers associated with the shared drives change the symbolic link table files are no longer correct. Each of these files must be modified to reflect the changes in disk structure. Returning to the example worksheet given above, assume that two physical disks are added to the first fibre channel array, and that the physical disks are configured as a single RAID logical drive with at least 3 partitions. In disk administrator the ordering of the disks in the second array changes. The symbolic link table file must be changed.

---

**Note:** If the cluster consists of more than two nodes, the other symbolic link table files need to be modified as well.

---

Thus, the new configuration of the cluster as seen by disk administrator is:

Node1	Disk 0	An internal disk
	Disk 1	An internal disk
	Disk 2	Shared storage disk
	Disk 3	Shared storage disk
	Disk 4	Shared storage disk
	Disk 5	Shared storage disk
	Disk 6	Shared storage disk
	Disk 7	Shared storage disk
Node2	Disk 0	An internal disk
	Disk 1	Shared storage disk
	Disk 2	Shared storage disk
	Disk 3	Shared storage disk
	Disk 4	Shared storage disk
	Disk 5	Shared storage disk
	Disk 6	Shared storage disk

A symbolic link worksheet corresponding to the above configuration is in the following table. Note that partition1 on each disk is used for the small, identifiable partitions discussed in the *Simplifying Detection of the Issue* section above; and therefore, partition1 is not shown in the symbolic links worksheet.

**Table 2. Modified Oracle Symbolic Links Worksheet**

Symbolic Link	Node1	Node2
OPS_log1t1	Harddisk2 Partition2	Harddisk1 Partition2
OPS_log2t1	Harddisk2 Partition3	Harddisk1 Partition3
OPS_sys01	Harddisk2 Partition4	Harddisk1 Partition4
OPS_usr01	Harddisk3 Partition2	Harddisk2 Partition2
OPS_rbs01	Harddisk3 Partition3	Harddisk2 Partition3
OPS_tmp01	Harddisk3 Partition4	Harddisk2 Partition4
OPS_log1t2	Harddisk4 Partition2	Harddisk3 Partition2
OPS_log2t2	Harddisk4 Partition3	Harddisk3 Partition3
OPS_log1t3 (new volume)	Harddisk5 Partition2	Harddisk4 Partition2
OPS_log2t3 (new volume)	Harddisk5 Partition3	Harddisk4 Partition3
OPS_cntr01	Harddisk6 Partition2	Harddisk5 Partition2
OPS_cmdisk	Harddisk7 Partition2	Harddisk6 Partition2

Finally, after the symbolic link table file has been modified, it needs to be run through an Oracle utility called SETLINKS. Refer to the *Oracle Parallel Server Getting Started Release 8.0.5 for Windows NT manual* and to other appropriate Oracle8 documentation on the location and parameters of the SETLINKS utility.

## Removing a Shared Storage Array

Each cluster node, and the cluster as a whole, depends on the shared storage for the Oracle database, Oracle data logs, and possibly application program files. If the shared storage array is removed, all clustered applications that are dependent on that array will be unable to operate properly.

---

**IMPORTANT:** Removing a shared storage array may remove information that is vital to the database. The database administrator should determine if removal of the array requires that the database be recreated.

---

Before removing a shared storage array, perform the following steps:

1. On each node, shutdown the database. This can be accomplished with Oracle Enterprise Manager or some other management application. As well, it can be accomplished using an Oracle database administration tool, SVRMGR. Using SRVMGR from the command line, enter the following commands:
  1. C:\SVRMGR30
  2. SVRMGR> connect internal/<password>
  3. SVRMGR> shutdown
  4. SVRMGR>exit
2. Next, stop the oracle service with the following command:  
C:\Net stop oracleservice<ops\_instance\_sid>
3. Shutdown Windows NT Server and power off all the cluster nodes. Power off all of the fibre channel arrays.
4. For the fibre channel array that is being removed, physically unplug the fibre channel cables from the array and from the fibre channel hub.
5. When the cluster is ready to be restarted, ensure the fibre channel arrays are powered on first, in the same order they were powered on when the cluster was originally configured.

---

**IMPORTANT:** If the removed array was not last in the power-on sequence of the arrays, you will need to reconfigure the Oracle symbolic links. For information see the section in this paper entitled *Managing Changes to Drive Ordering*.

---

## Adding a Shared Storage Array

A fibre channel storage array cannot be dynamically added to the cluster. You must follow these steps to add another shared storage array to an existing cluster.

1. On each node, shutdown the database. This can be accomplished with Oracle Enterprise Manager or some other management application. As well, it can be accomplished using an Oracle database administration tool, SVRMGR. Using SRVMGR from the command line, enter the following commands:
  1. C:\SVRMGR30
  2. SVRMGR> connect internal/<password>
  3. SVRMGR> shutdown
  4. SVRMGR>exit
2. Next, stop the oracle service with the following command:  
C:\Net stop oracleservice<ops\_instance\_sid>
3. Shutdown Windows NT Server and power off all the cluster nodes. Power off all of the fibre channel arrays.
4. Insert the new SCSI drives into the fibre channel array.
5. Using the fibre channel cables, physically attach the fibre channel array to the fibre channel hub used for shared storage system.
  - See the *Compaq Fibre Channel Storage System User Guide* for details on how to install cables and shortwave Gigabit Interface Converters.
6. Restart the fibre channel arrays. To avoid the drive ordering issue explained in the *Managing Changes to Drive Ordering* section above, restart the original arrays first – in the same order they were started when the cluster was originally configured. Then power on the newly added array. If you change the power-on sequence of the fibre channel arrays be sure to modify the symbolic link table file(s) to reflect the changes in drive order.
7. Boot the primary cluster node and start Windows NT.
8. Log on and run the Compaq Array Configuration Utility to create the desired RAID logical drives and to configure them with the desired RAID level.
9. Run disk administrator to create an extended partition on each of the newly configured disks.
10. Using disk administrator create logical partitions within the newly created extended partitions.

---

**IMPORTANT:** Do not format the drives. Oracle8 Parallel Server uses RAW partitions, which requires that the drive not be formatted with any file system.

---

11. Use disk administrator to remove the drive letters that might be assigned to the newly created logical partitions.

12. The symbolic link table files may need to be modified. If so, run the symbolic link table files through the SETLINKS utility on each cluster node. Refer to the section in this chapter entitled *Managing Changes to Drive Ordering* for comprehensive information about this task. The symbolic link table files will need to be modified to properly assign the symbolic links if the disk number assignments (as seen by disk administrator) have changed.

---

**Note:** The number of symbolic link table files depends on the number of nodes in the cluster. All symbolic link table files must be modified.

---

13. Power on each cluster node.
14. Use the disk administrator remove the driver letters to the newly added partitions and to verify that the order of the disks is the order that is expected. If not, you will need to troubleshoot the problem as to why the expected order does not exist. See the troubleshooting section in the *Compaq Parallel Database Cluster Administrator Guide* for troubleshooting suggestions.
15. Perform the Oracle commands to associate the new data files to the database.

## Adding or Removing Drives from a Shared Storage Array

Two distinct situations exist in which you might add or remove drives from a shared storage array:

- When one of the drives fails. In this case, you are simply exchanging the failed drive for a new one.
- When you want to add capacity to your fibre channel storage system.

In both situations, it is assumed you are employing RAID levels 1,4, or 5 for all the devices in the storage unit.

---

**IMPORTANT:** The addition and removal of drives on the fibre channel array must follow certain rules that are interpreted by reading the LEDs on each drive within the storage array. Be sure to read the *Fibre Channel Storage System User Guide* to understand these rules. Failure to follow these rules may result in loss of data.

---

### Replacing a Failed Drive

The procedure for replacing a failed drive is completed totally within the fibre channel storage system. Neither Oracle8 Parallel Server nor Windows NT Server is aware of the change, and operation of each continues without interruption.

---

**IMPORTANT:** If the failed drive was not configured to use RAID levels 1,4, or 5 (i.e., your drives have no fault tolerance), you can lose some or all of the data on the failed drive.

---

See the *Fibre Channel Storage System User Guide* for instructions on replacing a failed drive.

## Adding a Drive to Increase Storage Capacity

It is feasible that during the life of your Parallel Database Cluster you will need to expand the capacity of your shared storage system. The following steps describe how to add a drive to the Compaq Fibre Channel Storage System and how to allocate it to Oracle8 Parallel Server.

1. Physically add the drive(s) to the fibre channel array.
2. Run the Compaq Array Configuration Utility to create the desired RAID logical drives and to configure them with the desired RAID level.

---

**Note:** You cannot increase the capacity of an existing Windows NT drive, but can create new partitions with the extra capacity furnished by the added drive(s). The new partitions can then be added to the Oracle database.

---

3. Start Windows NT and run disk administrator to create an extended partition on each of the newly configured drive volumes.
4. Using disk administrator, create logical partitions within the newly created extended partitions.

---

**IMPORTANT:** Do not format the drives. Oracle8 Parallel Server uses RAW partitions, which requires that the drive not be formatted with any file system.

---

5. Use disk administrator to remove the drive letters that may be assigned to the newly created logical partitions.
6. If the addition of the drives affects the order in which disk numbers are assigned to the shared drives, the symbolic link table file(s) will need to be modified to reflect the changes. Refer to the section in this chapter entitled *Managing Changes to Drive Ordering* for comprehensive information about this task.

---

**Note:** The number of symbolic link table files depends on the number of nodes in the cluster. All symbolic link table files must be modified. Update the symbolic link table file as needed for each node and then execute the SETLINKS utility to establish the proper symbolic links on each node.

---

7. Perform the necessary Oracle commands to associate the new data files to the database