VERITAS Volume Manager 4.1 Hardware Notes

HP-UX 11i v2

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.

Warranty

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

U.S. Government License

Proprietary computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

VERITAS® is a registered trademark of VERITAS Software Corporation.

VERITAS File System™ is a trademark of VERITAS Software Corporation.

Publication History

The manual publication date and part number indicate its current edition. The publication date will change when a new edition is released. The manual part number will change when extensive changes are made.

To ensure that you receive the new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

- First Edition: December 2002, 5187-1878, HP-UX 11i Version 1(B.11.11)
  CD-ROM (Software Pack 11i December 2002) and Web (**http://docs.hp.com**)

- Second Edition: August 2003, 5971-4774, HP-UX 11i Version 2 (B.11.23)
  Web (**http://docs.hp.com**)

- Third Edition: December 2003, B3929-90015, HP-UX 11i Version 2 (B.11.23)
  Web (**http://docs.hp.com**)

- Fourth Edition: May 2005,5991-1143, HP-UX 11i Version 2 (B.11.23)

Getting Help

For license information (U.S. and Canadian Customers):

- Phone: 650-960-5111

- Email: hplicense@mayfield.hp.com

For license information (Europe)

- Phone :+33.(0)4.76.14.15.29

- Email: codeword_europe@hp-france-gen1.om.hp.com

For the latest patch information :

- http://itrc.hp.com

Preface

1. Hardware Notes

Preface

The *VERITAS Volume Manager 4.1 Hardware Notes* provides hardware information on VERITAS Volume Manager (VxVM) Release 4.1.

The purpose of this guide is to provide information on hardware requirements for VERITAS Volume Manager 4.1. It also describes how to use the VxVM command line interface for administering hardware with VxVM. Detailed descriptions of the VxVM commands and utilities, their options, and details on their use are located in the VxVM manual pages.

This guide assumes that you have a:

- Working knowledge of the HP-UX operating system

- Basic understanding of HP-UX system administration

- Basic understanding of storage management

NOTE	Most VERITAS Volume Manager commands require superuser or other appropriate privileges.

Conventions

The following table describes the typographic conventions used in this guide.

Table 1

Typeface	Usage	Examples
`monospace`	Computer output, file contents, files, directories, software elements such as command options, function names, and parameters	Read tunables from the `/etc/vx/tunefstab` file. See the *ls* (1) manual page for more information.
italic	New terms, book titles, emphasis, variables to be replaced by a name or value	See the *VERITAS Volume Manager 4.1 Hardware Notes* for details. The variable *ncsize* determines the value of...
`monospace (bold)`	User input; the "#" symbol indicates a command prompt	`# `**`mount -F vxfs /h/filesys`**
`monospace (bold and italic)`	Variables to be replaced by a name or value in user input	`# `**`mount -F `*`fstype`* *`mount_point`*

Symbol	Usage	Examples
`%`	C shell prompt	
`$`	Bourne/Korn/Bash shell prompt	
`#`	Superuser prompt (all shells)	

Table 1   (Continued)

Typeface	Usage	Examples
\	Continued input on the following line	`# `**`mount -F vxfs \`** **`/h/filesys`**
[]	In a command synopsis, brackets indicates an optional argument	`ls [ -a ]`
|	In a command synopsis, a vertical bar separates mutually exclusive arguments	`mount [suid | nosuid ]`

1 Hardware Notes

This document provides hardware support information for the VERITAS Volume Manager Release 4.1 on HP-UX 11i version 2 platform.

The following topics are covered in this document:

VERITAS Volume Manager Supported Platforms

The following platforms are supported for VERITAS Volume Manager 4.1 for HP-UX:

Table 1-1   Supported Platforms for VERITAS Volume Manager 4.1

HP Integrity Systems	cx2600, rx1600, rx2600, rx4640, rx5600, rx7600, and rx8600 series. Integrity Superdome, zx2000 and zx6000 workstations.
PA-RISC	rp2400, rp3340, rp4400, rp5400, rp7400, and rp8400 series Superdome 64-bit A, L, and N classes.

Contact HP Technical Support for more information.

Disk Array Configurations Supported for Use with DMP

For the most up to date list of supported Array Support Libraries (ASLs), visit the VERITAS support site at http://support.veritas.com, select Volume Manager > Volume Manager for UNIX, and click on Compatibility. Enter the product version and platform details, and then select HardWare Compatibility List. The Hardware Compatibility List includes a list of supported arrays, and a list of supported arrays together with a pointer to the appropriate Tech Note.

Arrays supported by the Dynamic Multipathing (DMP) feature are a subset of the arrays supported by VxVM. The following arrays can have multiple paths to the same disk. Arrays that are not listed in this document or on the VERITAS Support website may only be supported for single path access.

Arrays supported in Active/Active (A/A) mode:

- JBOD (Just a Bunch of Disks); see "JBOD Disk Array Support in DMP" on page 14

- Cisco MDS 9000

- EMC Symmetrix 5.0, 5.5 and 6.0; see "EMC Symmetrix Disk Arrays" on page 14

- Hitachi Data Systems 9970 and 9980 Disk Array Subsystem

- Hitachi TagmaStore USP

- HP Surestore Disk System FC10; see "HP Surestore Disk System FC10" on page 15

- HP StorageWorks Disk System DS2405

- HP Surestore Virtual Array 7100, 7110, 7400 and 7400

- HP Surestore Disk Array XP48, XP128, XP256 and XP1024; see "HP Surestore Disk Array XP48, XP256 and XP512" on page 16

Arrays supported in Active/Passive (A/P) mode:

- EMC Clariion CX400 and CX600; see"DMP Coexistence with EMC PowerPath" on page 17

- Hitachi Data Systems 9500 Disk Array Subsystem

- HP High Availability Fibre Channel Disk Array Model 30/FC

- HP Surestore Disk Array FC60; see"HP Surestore Disk Array FC60" on page 16

- IBM FAStT 500, 600, 700 and 900

- STK FlexLine FLX280

See "Administering Array Support Library Packages" on page 22 for details of how to add, remove or upgrade ASL packages.

JBOD Disk Array Support in DMP

NOTE	The Dynamic Multipathing feature works with JBOD drives only if the individual disks provide unique serial number information. If the individual disks do not provide unique serial numbers, even if the disks are configured for multipathing in hardware, Base-VXVM does not provide multipathing support.

DMP supports multipathing of JBOD disk arrays that are connected to a system running Base-VXVM. These JBOD disk arrays can contain disks from any vendor. However, DMP can correctly detect multiple paths to the disks in these disk arrays only if the following conditions are satisfied:

• The serial number field of the Standard SCSI INQUIRY DATA (bytes 36—47) should be world wide unique for any disk that is connected to a JBOD. This fact must be known to the user before the user attempts to configure DMP to add these disks to the JBOD category of DMP.

CAUTION	If any disk does note provide a unique serial number, DMP cannot detect multiple paths to the disk correctly and results are unpredictable. Contact the JBOD disk or disk array vendor to obtain correct information about unique serial numbers on disks.

• The JBOD disk arrays that require multipathing functionality from DMP should be of the Active/Active type. This means that all paths to all LUNs on the JBOD disk array should be simultaneously accessible through all paths, without causing any performance degradation.

EMC Symmetrix Disk Arrays

The Dynamic Multipathing feature (DMP) of the Volume Manager works with the EMC Symmetrix disk array only if the disk array is configured in the Common Serial Number Mode. This is because only the Common Serial Number Mode configuration provides unique device identification on different paths for DMP to configure its database. Without this mode, DMP cannot provide the multipathing functionality. DMP cannot access EMC Symmetrix disks as individual disks with a single path, it found by scanning the device information tree in the kernel.

Ultra SCSI and fast-wide differential (FWD) SCSI support has been tested on the EMC OSD (Open System Director) controller.

For the DMP functionality to work with the EMC Symmetrix disk array, the following conditions must be met:

- Common Serial Number Mode should be enabled for all channels that participate in DMP. This operation should be performed by the EMC Customer Engineer before installing the VERITAS software packages.

- Disable queue reset on unit attention mode must be set for all channels that participate in DMP. This operation should be performed by the EMC Customer Engineer before installing the VERITAS software packages.

- The minimum level of microcode on the EMC Symmetrix should be as shown in the following table.

Table 1-2   Minimum Level of Microcode on the EMC Symmetrix

Sym	Model	Microcde
4	3330, 3430 and 3700	5265 and 5266
4.8	3630, 3830 and 3930	5265 and 5266
5	8130, 8430 and 8730	5566

Contact your EMC Technical Support Representative for further instructions on configuring EMC hardware.

See your sales representative for further updates on hardware support.

HP Surestore Disk System FC10

Because of the way LUNs are distributed across controllers by HP-UX, the enable/disable controller feature of Base-VXVM does not function as expected for the following configurations of FC10 disk systems:

- multiple daisy chained FC10 disk systems connected to a single host

- multiple FC10 disk systems connected to a single host through a hub

If the `vxdmpadm disable ctlr=ctlr_name` command is used for the above configurations to disable I/O through a particular controller, I/O can be stopped through paths that were not intended by the user.

If each FC10 disk system is individually connected to the host (using separate HBAs), the above mentioned feature functions correctly and as expected.

HP Surestore Disk Array FC60

Although the HP Surestore Disk Array FC60 can be used as an active/active disk array in some configurations, with VERITAS Volume Manager for HP-UX it can only be used as an active/passive disk array.

HP Surestore Disk Array XP48, XP256 and XP512

HP Surestore Disk Arrays XP48, XP256 and XP512 are supported with Dynamic Multipathing (DMP) in active/active mode. If Base-VXVM encounters devices with identical VxVM Disk IDs, the Base-VXVM software will give a preference to selecting the P-Vol as opposed to the S-Vol to resolve the conflict. This disk selection policy works only for microcode revisions 5244 and above.

DMP Coexistence with EMC PowerPath

With this release of Base-VXVM, DMP can coexist with PowerPath. Depending on the scenario, you might need to install the EMC Cx600 ASL and its associated APM. To use DMP with PowerPath, you should be aware of the following scenarios.

- If you are installing Base-VXVM 4.1 and PowerPath is installed, you do not need to install the Cx600 ASL and its associated APM. The array must be configured in explicit fail-over mode.

- If you are installing Base-VXVM 4.1 and PowerPath is not installed, you must install the Cx600 ASL and its associated APM. The array can be in any mode.

The following table lists these scenarios and the required array modes:

Table 1-3   Scenarious and Required Array Modes

PowerPath	DMP	Array mode
Installed	`libvxemc.so` handles EMC arrays and DGC claiming internally. PP handles fail-over. No need to install C.x600 ASL or Cx600 APM	Explicit fail-over
Not installed; the array is not Cx600	`libvxemc.so` handles the EMC Symmetrix array. DMP handles multipathing.	-
Not installed, but the array is Cx600	DMP handle multipathing. Cx600 ASL and Cx600 APM must be installed	-

NOTE   If you are upgrading a system to Base-VXVM 4.1, which has PowerPath installed, and the Cx600 ASL and its associated Cx600 APM are also installed, you must uninstall both the Cx600 ASL and APM, otherwise the Cx600 will claim the disks and the PowerPath discs will not be identified.

If you are upgrading a system to Base-VXVM 4.1, which does not have PowerPath installed, but the Cx600 ASL and its APM are both installed, then the Cx600 ASL and its associated APM should not be uninstalled.

Dynamic LUN Expansion

The following form of the `vxdisk` command can be used to make VxVM aware of the

new size of a virtual disk device that has been resized:

```
# vxdisk [-f] [-g diskgroup] resize {accessname|medianame} \
[length=value]
```

The device must have a SCSI interface that is presented by a smart switch, smart array or RAID controller. Following a resize operation to increase the length that is defined for a device, additional disk space on the device is available for allocation. You can optionally specify the new size by using the length attribute.

If a disk media name rather than a disk access name is specified, the disk group must either be specified using the `-g` option or the default disk group will be used. If the default disk group has not been set up, an error message will be generated.

NOTE	A Storage Foundation license is required to use the `vxdisk` resize command.
	This facility is provided to support dynamic LUN expansion by updating disk headers and other VxVM structures to match a new LUN size. It does not resize the LUN itself.
	Any volumes on the device should only be grown after the device itself has first been grown. Otherwise, storage other than the device may be used to grow the volumes, or the volume resize may fail if no free storage is available.
	Resizing should only be performed on devices that preserve data. Consult the array documentation to verify that data preservation is supported and has been qualified. The operation also requires that only storage at the end of the LUN is affected. Data at the beginning of the LUN must not be altered. No attempt is made to verify the validity of pre-existing data on the LUN. The operation should be performed on the host where the disk group is imported (or on the master node for a cluster-shared disk group).
	Resizing of LUNs that are not part of a disk group is not supported.It is not possible to resize LUNs that are in the boot disk group (aliased as `bootdg`), in a deported disk group, or that are offline, uninitialized, being reinitialized, or in an error state.

CAUTION	Do not perform this operation when replacing a physical disk with a disk of a different size as data is not preserved.
	Before reducing the size of a device, any volumes on the device should first be reduced in size or moved off the device. By default, the resize fails if any subdisks would be disabled as a result of their being removed in whole or in part during a shrink operation.
	If the device that is being resized has the only valid configuration copy for a disk group, the -f option may be specified to forcibly resize the device.
	Resizing a device that contains the only valid configuration copy for a disk group can result in dataloss if a system crash occurs during the resize.
	Resizing a virtual disk device is a non-transactional operation outside the control of VxVM. This means that the resize command may have to be re-issued following a system crash. In addition, a system crash may leave the private region on the device in an unusable state. If this occurs, the disk must be reinitialized, reattached to the disk group, and its data resynchronized or recovered from a backup.

Extended Copy Service

The Extended Copy Service feature of VxVM automatically works in tandem with the extended copy engines from array vendors. When VxVM detects that the source and destination devices are enabled for extended copy, VxVM automatically off loads copy requests to the array's copy manager.

The benefits of VxVM Extended Copy Service are:

- Non-disruptive copy operations from disk to disk. The host server remains online during the copy and the data being copied remains accessible to the server.

- Server-free copy operation. The copy operation is done between the array subsystem and the target disk. The data copy operation does not use any CPU or I/O resources on the host server.

To see whether the Extended Copy Service feature is enabled on a disk, use the vxprint command as shown in the following example. The feature is enabled if an ecopy_enabled entry appears in the flags line.

```
# vxprint -l disk03

Disk group: rootdg

Disk: disk03

info: diskid=1234567890.1234.tungsten

assoc: device=sde

flags: autoconfig ecopy_enabled

device: pubpath=/dev/vx/dmp/sde

devinfo: publen=35354136 privlen=9167
```

NOTE	If required, you can use the -o noecopy option to turn off Extended Copy Service for each invocation of the vxplexatt, cp, mv and snapstart commands, and the vxsdmv command.

Enabling a Disk for ecopy Operation

To enable a disk for ecopy operation:

1. Install the Hardware Assisted copy license

2. Enable the Extended copy features in the array. (This procedure is vendor specific.)

3. Install the vendor ASL that supports the ecopy feature.

4. Shutdown and reboot the system.

Extended Copy Service Support for Hitachi 9900 and 9900V Arrays

To implement extended copy for the Hitachi 9900 and 9900V arrays, you must create the two files, `user_pwwn_file` and `user_luid_file`, that identify the Hitachi disk identification information. To create these files, run the `vxwwnluid` utility that is located in `/etc/vx/diag.d`. There are no parameters for the `vxwwnluid` utility, and the command must be executed as root.

- `/etc/vx/user_pwwn_file`

  This file contains the disk access name and the port world-wide-name (`pwwn`) for each disk in the array. For the hitachi arrays, both the source and the destination devices must have entries in the this file. The information for each disk in the array is defined on a single line. The disk access name and PWWN are separated by a single tab character.

  The following are sample entries from the `/etc/vx/user_pwwn_file` file:

  ```
  sde 50060e800404040b

  sdf 50060e800404040b

  sdg 50060e800404040b
  ```

- `/etc/vx/user_luid_file`

  This file contains the disk access names and their corresponding LUN numbers in the array. The information for each disk in the array is defined on a single line. The disk access name and the LUN are separated by a single tab character.

  The following are sample entries from the `/etc/vx/user_luid_file` file:

  ```
  sde 1 sdf 2 sdg 1
  ```

Administering Array Support Library Packages

Base-VXVM provides Dynamic MultiPathing (DMP) support for new disk arrays in the form of Array Support Library (ASL) software packages. You can obtain ASL packages for new arrays from:

- The VxVM release in the `veritas_enabled` directory.

- The Storage Array Vendor's support site.

- The VERITAS Technical Support site, http://support.veritas.com. Select "Volume Manager" from the Product Family menu and "Volume Manager on Unix" from the Product men. Then click on Knowledge Base Search and enter asl followed by the vendor and/or model name.

Base-VXVM 4.1 allows an ASL to be added to a running VxVM system (see"Adding an ASL package" on page 23 "). This means that you can add support for a particular disk array type (including multipathing support by Dynamic Multipathing, DMP) without having to: stop VxVM, reboot the system, or modify the VxVM package.

NOTE	For the installed ASL to be discovered correctly, the array LUNs must be visible to the operating system through multiple paths, and they must be appropriately configured to support DMP. For instructions on how to configure the array correctly, please contact the array vendor.

You can also remove ASL packages from a running VxVM system (see "Removing an ASL Package" on page 24"). If the associated disk array type contains active volumes, these will remain available during and after the addition or removal of the ASL package.

Only one version of an ASL for a specific type/model of disk array can be active in the system. To install a different version of an ASL, you need to first remove the previous version of the ASL package, then install the new package (see "Upgrading an ASL Package" on page 25").

Commands Used to Maintain ASL Packages

The following packaging commands are available for installing, removing, and listing ASL packages.

Table 1-4   **ASL Packages**

Install	Remove	Obtain Information
`swinstall`	`swremove`	`swlist`

See the appropriate manual page for more information on these commands.

Adding an ASL package

Only one version of an ASL can be active in the system. If you find a version of the ASL is already installed, you need to refer to "Upgrading an ASL Package" on page 25".

NOTE   Disk arrays that are supported by an ASL do not have to be connected to the system when the ASL package is installed.

To find out if package (ASL_name) is already installed; enter:

```
# swlist ASL_name
```

You will find the ASL packages in the
`/mount_point/storage_foundation/veritas_enabled` directory. They have the format, `packagename.gz.tar`.

To add an ASL package , proceed as follows:

1. Change directory to the ASL packages directory;

   ```
   # cd /mount_point/storage_foundation/veritas_enabled
   ```

2. Copy the required package to a temporary directory:

   ```
   # cp packagename.tar.gz /tmp
   ```

3. Change directory to the temporary directory:

   ```
   # cd /tmp
   ```

4. Unzip and extract the package:

   ```
   # gzcat packagename.tar.gz | tar -xvf -
   ```

5. Use `swinstall` to install the package:

```
# swinstall ASL_Name
```

If vxconfigd is active, invoke Device Discovery Layer (DDL) to allow the newly installed ASL libraries to claim the disk arrays:

```
# vxdctl enable
```

Otherwise DDL is started automatically the next time vxconfigd is activated.

6. Verify that disk array was claimed by the ASL_name library:

```
# vxdmpadm listenclosure all
```

```
# vxdisk list cxtxdxs2
```

```
# vxdmpadm getsubpaths ctlr=cx
```

NOTE   If disk array has multiple paths, use the above commands to verify that all paths to disk are claimed.

Removing an ASL Package

Before removing an ASL package, first find out if the package (ASL_name) is installed; enter:

```
# swlist ASL_name
```

1. To remove an ASL package, enter:

```
# swremove ASL_name
```

2. If vxconfigd is active, invoke Device Discovery Layer (DDL) through the vxdctl command to allow the arrays to be reclaimed:

```
# vxdctl enable
```

Otherwise DDL is started automatically the next time vxconfigd is activated.

3. Verify that disk array was claimed as "Disk".

```
# vxdmpadm listenclosure all
```

NOTE   If the disk array has multiple paths, they must all be claimed as "Disk" (use the vxddladmaddjbod command); otherwise, duplicate disk ID errors will be detected.

Upgrading an ASL Package

Before upgrading an ASL package (installing a new version of an already installed package), first find out if the old package (ASL_name) is already present; enter:

```
# swlist ASL_name
```

- Remove the package (ASL_name):

  ```
  # swremove ASL_name
  ```

- Proceed to add the new ASL package as described by "Adding an ASL package" on page 23".

The DMP Administartion Utility - `vxdmpadm`

The `vxdmpadm` utility administers the DMP functionality of VxVM. Refer to the *Administering Dynamic Multipathing (DMP)* chapter in the *VERITAS Volume Manager 4.1 Administrator's Guide*, and the `vxdmpadm`(1M) manual page for detailed descriptions of the options and attributes of this command.

You can use the `vxdmpadm` utility option `listctlr` to list all controllers connected to disks that are attached to the host. For example, to list all controllers connected to disks on the host, use the following command:

```
# vxdmpadm listctlr all
```

to display the output from the above command:

```
CTLR-NAME      ENCLR-TYPE     STATE     ENCLR-NAME

=======================================================

c0             OTHER          ENABLED   Disk

c1             Disk           ENABLED   Disk

c2             Disk           ENABLED   Disk
```

The previous displayed output shows that:

- The first controller `c0` is connected to disk(s) that are not under any recognized DMP category.
- The second and third controllers (`c1` and `c2`) are connected to a JBOD (indicated by an enclosure type of Disk).

All the above controllers are in the ENABLED state which indicates that they are available for I/O operations.

The state DISABLED is used to indicate that controllers are unavailable for I/O operations. Theunavailability can be due to a hardware failure or due to I/O operations being disabled on that controller by the System Administrator. The controller state can be changed by using the `vxdmpadm` utility.

To list all the paths that are connected to a particular controller, you can use the `getsubpaths` option with the `ctlr` attribute. For example, use the following command:

```
# vxdmpadm getsubpaths ctlr=c1
```

to display the output from the above command:

```
NAME STATE PATH-TYPE DMPNODENAME ENCLR-TYPE ENCLR-NAME

================================================================

c1t0d0      ENABLED    -c2t0d0     Disk        Disk0

c1t1d0      ENABLED    -c2t1d0     Disk        Disk

c1t2d0      ENABLED    -c2t2d0     Disk        Disk

c1t3d0      ENABLED    -c2t3d0     Disk        Disk

c1t4d0      ENABLED    -c2t4d0     Disk        Disk

c1t5d0      ENABLED    -c2t5d0     Disk        Disk

c1t6d0      ENABLED    -c2t6d0     Disk        Disk

c1t16d0     ENABLED    -c2t16d0    Disk        Disk

c1t17d0     ENABLED    -c2t17d0    Disk        Disk

c1t18d0     ENABLED    -c2t18d0    Disk        Disk

c1t19d0     ENABLED    -c2t19d0    Disk        Disk

c1t20d0     ENABLED    -c2t20d0    Disk        Disk

c1t21d0     ENABLED    -c2t21d0    Disk        Disk

c1t22d0     ENABLED    -c2t22d0    Disk        Disk
```

The output display shows the paths that are connected to the controller named c1.

You can use the getsubpaths option combined with the dmpnodename attribute to list all paths that are connected to a LUN (represented by a DMP device). For example, to list information about paths that lead to the LUN named c1t0d0, use the following command:

vxdmpadm getsubpaths dmpnodename=c1t0d0

to display the output from the above command:

```
NAME STATE      PATH-TYPE    CTLR-NAME ENCLR-TYPE ENCLR-NAME

================================================================

c2t0d0         DISABLED    -c2       Disk     Disk

c1t0d0         ENABLED     -c1       Disk     Disk
```

The listing above shows that the DMP device c1t0d0 has two paths to it that are named c1t0d0 and c2t0d0. Additional information indicates that only one of these paths is available for I/O operations. One of these paths is in the ENABLED state and the other is in the DISABLED state. Both paths are in a SEAGATE disk array.

To retrieve the name of the DMP device that controls a particular path, the getdmpnode option can be combined with the nodename attribute by using the following command:

vxdmpadm getdmpnode nodename=c1t0d0

to display the output from the above command:

```
ME        STATE       ENCLR-TYPE PATHS ENBL    DSBL ENCLR-NAME

==============================================================

c2t0d0    ENABLED     Disk      2      2    0    Disk
```

This example output shows that the physical path `c1t0d0` is owned by the DMP device `c2t0d0`, which has two paths to it.

Correcting the Effects of Swapping Paths

If you accidentally or intentionally swap the paths to an array by changing the N_Port_id of the affected disks, you must use the fcmsutl(1M) command to reconfigure the disks before attempting to rediscover the VxVM configuration.

To recover the situation, run the following command to fetch the N_Port_id of the affected devices:

```
# fcmsutil TD_device_file get remote all
```

The following is an example of running this command:

```
# fcmsutil /dev/td0 get remote all | grep N_Port_id
```

```
Target N_Port_id is = 0x00006b
```

```
Target N_Port_id is = 0x00006c
```

```
Target N_Port_id is = 0x00006d
```

```
Target N_Port_id is = 0x00006e
```

Then run fcmsutil with the replace_dsk option, for example:

```
for nport in 0x00006b 0x00006c 0x00006d 0x00006e
```

```
do
```

```
fcmsutil /dev/td0 replace_dsk $nport
```

```
done
```

You should run this script for both affected TD device files (for example, /dev/td0 and /dev/td1). If fcmsutil is not used, the commands ioscan, insf-e and vxdctl enable might take a long time to complete. When this situation occurs, there will be a message in the file /var/adm/syslog/syslog.log indicating that you must use fcmsutil to correct the problem.