VERITAS™

WHITE PAPER

VERITAS Dynamic Multipathing (DMP)

VERITAS

Table of Contents

Introduction/Background

What is Dynamic Multipathing (DMP)?

DMP is a method of providing two or more hardware paths to a single drive. For example, the physical hardware can have at least two paths, such as **c1t1d0** and **c2t1do**, directing I/O to the same drive. VERITAS Volume Manager™ arbitrarily selects one of the two names and creates a single device entry, then transfers data across both paths to spread the I/O.

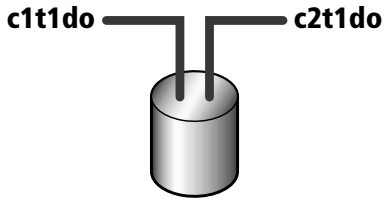

Figure 1: Two I/O paths to a single drive or array

DMP is enabled by default. Volume Manager detects multipath systems with the Universal Worldwide Device Identifiers (WWD IDs) and manages multipath targets, such as disk arrays, which define policies for using more than one path.

Benefits of DMP

High availability: DMP provides greater reliability through a path failover mechanism. In the event of a loss of one connection to a disk, the system continues to access the critical data over the other sound connections until the failed path is replaced.

Improved performance: DMP provides greater I/O throughput by balancing the I/O load uniformly across multiple I/O paths to the disk device.

Volume Manager 2.5.x vs. Volume Manager 3.x

On Volume Manager 2.5.x, DMP supports 42 paths. On Volume Manager 3.x and beyond, DMP supports unlimited paths.

Active/Active Disk Arrays

"Active/Active" is a DMP-specific term for disk arrays that allow I/Os to disks through multiple paths simultaneously, without causing any performance overhead. For these disk arrays, Volume Manager follows the round-robin policy for issuing I/Os across available paths in the Volume Manager 2.5.x release(s).

VERITAS Volume Manager uses the "balanced path" policy to distribute I/Os across available paths in the Volume Manager 3.x release(s). Load balancing ensures that I/O throughput can be increased by utilizing the full bandwidth of all paths. Sequential I/Os starting within a certain range are sent down the same path to optimize I/O throughput using disk track caches. However, large sequential I/Os that do not fall within this range are distributed across multiple paths to take advantage of load balancing.

These disk arrays are used in load-balancing mode by distributing I/Os across all available paths. Some examples of these types of disk array are: A5X00 (SENA) disk array(s) from Sun, the SPARC Storage Array (SSA), EMC Symmetrix, Hitachi 7700E.

Active/Passive Disk Arrays

"Active/Passive" is a DMP-specific term for disk arrays in which controllers "own" Logical Units (LUNs)/disks. Only the controller which "owns" the LUN issues I/Os to these LUNs. The controller that "owns" a LUN is called the "primary" path and the alternate controller is termed as the "secondary" path to the LUN.

Accessing a disk/LUN through multiple available paths on such disk arrays is not allowed. If a LUN on an Active/Passive disk array is accessed simultaneously via multiple paths, the "ownership" of the LUN shifts back and forth across the controllers, causing the "ping-pong" effect. Because changing ownership is a time-consuming operation, this can cause immense performance degradation.

For such disk arrays, DMP policy is to use the available primary path as long as it is accessible. Shift I/Os to the secondary path only when the primary path fails. This is called "failover" or "standby" mode of operation for I/Os. Examples of such disk arrays are: Hitachi 5700E, Hitachi 5800E, Nike (Model 10, 20), Galaxy and Purple (T300).

Scheduling I/Os

Volume Manager 2.5.x

DMP follows the round-robin policy of scheduling I/Os across available paths. I/Os within 64K on either side of the last I/O, are sent down the same path. This sequential I/O detection is done in Volume Manager 2.5.x for A5000 disks only.

If DMP cannot identify multiple paths on Volume Manager 2.5.5 on Solaris 2.6, it may lack the proper license to turn on multipathing. Or the disk(s) are not returning unique serial number information in Standard Inquiry data in bytes 36-47.

Volume Manager 3.x

DMP follows a "balanced path policy" for scheduling I/Os across available paths for Active/Active disk arrays such as EMC. Balance random I/Os across all the available paths using the following method:

```
iopath = ((bp->b_blkno >> 7) % dmpp->tpaths);`
```

The above means that all I/Os starting within a certain 64K range go through the same path. I/Os that start within other 64K ranges on the disk will be routed through a different path.

"dmpp->tpaths" indicates the total number of paths to this disk and "iopath" is the path used for the I/O. I/Os are sent down only those paths that are enabled.

In other words, the I/O load is balanced across all the available paths and sequential I/Os within a 64K range, are sent down the same path.

Path Failure

If a path fails, alternate paths are used. If all access paths have failed, the disk is considered to have failed. A disk driver failure will not cause a DMP failure.

The flow of event notification during I/O failure is as follows:

1. Device driver/HBA controller issues SCSI III commands to disk and command fails
2. OS kernel trying to issue write fails, generates an error
3. Based on the error code, the DMP driver takes the correct action and fails over the I/O to the alternate HBA
4. I/O is reissued down alternate path

The failover process can be very fast for an Active/Active disk array like the XP256 or EMC Symmetrix. The DMP driver itself does not take any time to switch over. However, the total time for failover is dependent on how long the underlying disk driver retries the command before giving up. On HP, the user can set the PowerFail timeout value for that disk.

DMP allows the administrator to indicate to the DMP subsystem in Volume Manager whether the connection is repaired or restored. This is called DMP reconfiguration. The reconfiguration procedure also allows the detection of newly added devices, as well as devices that are removed after the system is fully booted (if the operating system detects them properly).

Disabling DMP

NOTE : The following procedures will require a system reboot for the changes to take effect.

Typically, DMP is disabled when other multipathing drivers cannot co-exist with DMP, or when DMP does not handle certain devices properly.

Disabling DMP on an HP system

1. Stop vxconfigd using the command "vxdctl stop"
2. Save the "/stand/system" file as "/stand/system.vxdmp" by executing the command "cp /stand/system /stand/system.vxdmp"
3. Save the "/stand/vmunix" file as "/stand/vmunix.vxdmp" by executing the command "cp /stand/vmunix /stand/vmunix.vxdmp"
4. Edit the "/stand/system" file and remove the "vxdmp" entry
5. Run the "/etc/vx/bin/vxdmpdis" script

If all of the steps are completed successfully, reboot the system. When the system comes up, DMP will have been removed. Verify this by running the vxdmpadm command. The following message should appear:

Volume Manager:vxdmpadm: ERROR: vxdmp module is not loaded on the system. Command invalid.

Also the command "vxdisk list <da_name>" should not display any multipathing information.

Disabling DMP on a Solaris system

In Volume Manager 2.5.x and Volume Manager 3.x, the user can disable DMP by following these steps:

1. Remove the directories /dev/vx/dmp and /dev/vx/rdmp and any files in these directories
2. Use the following commands to link the /dev/vx/dmp and /dev/vx/rdmp directories to the /dev/dsk and /dev/rdsk directories respectively:

   # ln -s /dev/dsk /dev/vx/dmp

   # ln -s /dev/rdsk /dev/vx/rdmp
3. Remove the line "forceload: drv/vxdmp" from the /etc/system file by commenting out that entry to "* forceload: drv/vxdmp" OR by deleting that line
4. Go to the /kernel/drv directory and execute the following:

   # mv ./vxdmp ./vxdmp.orig

Considerations for Disabling DMP in Volume Manager 3.1.1

Adopting the current method of removing the DMP layer altogether is not flexible. This takes away the multipathing functionality that DMP provides for other devices in the system.

Two new features, co-existence with third-party multipathing solutions and Platform Independent Device Naming, require that the DMP driver always should be present in the system. In Volume Manager 3.1.1, customers will be able to prevent Volume Manager (DMP) from multipathing some or all devices on the system, without removing the DMP layer.

The user will be presented with interfaces through "vxinstall" and "vxdiskadm" to do the following:

- Suppress devices from Volume Manager's view
- Unsuppress devices that were previously suppressed from Volume Manager's view
- Prevent devices from being multipathed by VERITAS DMP
- Allow VERITAS DMP multipathing for devices that were prevented earlier

The user can specify the devices for the above operations using the following mechanisms:

- Using a VendorID:ProductID combination, e.g. EMC:SYMMETRIX
- Using controller name, e.g. c1
- Specifying path names, e.g. c1t0d9
- All devices on the system

Files are used to keep the exclude information the user specified (through vxinstall or vxdiskadm).

- Devices that the user has suppressed from Volume Manager's view are kept in the file /etc/vx/vxvmexclude
- Devices that the user has specified not be multipathed are kept in the file /etc/vx/vxdmpexclude

Enabling DMP on an HP System (where DMP is fully disabled)

1. Stop vxconfigd using the command "vxdctl stop"
2. Save the "/stand/system" file as "/stand/system.vxdmp" by executing the command "cp /stand/system /stand/system.vxdmp"
3. Save the "/stand/vmunix" file as "/stand/vmunix.vxdmp" by executing the command "cp /stand/vmunix /stand/vmunix.vxdmp"
4. Edit the "/stand/system" file and add the "vxdmp" entry to it (after the "VxVm" entry)
5. Run the "/etc/vx/bin/vxdmpen" script

If all of the steps are completed successfully, reboot the system. When the system comes up, DMP should be enabled. Verify this by running the vxdmpadm command. This command should show multipathing information. Also, the command "vxdisk list <da_name>" should show multipathing information.

Retrieving Information on DMP Nodes

The vxdmpadm utility is an administrative interface to the DMP subsystem. Using the vxdmpadm utility, you can:

- List all controllers connected to disks attached to the host
- List all the paths connected to a particular controller
- List all paths under a DMP device
- Retrieve the name of the DMP device corresponding to a path
- Enable or disable a host controller
- Rename an enclosure

Supported Arrays

Solaris 2.6, 7, 8

DMP-supported arrays are a subset of arrays that Volume Manager supports. For a current and accurate list, please refer to **www.veritas.com**.

- EMC Symmetrix when configured in the *Common Serial Number Mode* before installing VERITAS Volume Manager
- HP SureStore E Disk Array XP256/XP512
- IBM Enterprise Storage Servers (ESS)
- Hitachi Data Systems 5700E Disk Array Subsystem
- Hitachi Data Systems 5800E/7700E Disk Array Subsystem
- Sun StorEdge A5x00 Array
- Sun StorEdge T3 Array
- JBOD (Just a Bunch of Disks)
- Seagate disks that return unique serial numbers in standard SCSI inquiry data

- Storage Computer OmniRAID disk array. To multipath Storage Computer disk arrays connected to the system while using Volume Manager 3.1.1, you must assign a unique system name for each disk array connected to a machine. The RAID-5 Users Manual at the Web site **www.storage.com** describes how to set up a system name for Storage Computer disk arrays.
- ECCS Synchronix Array

HP-UX 11.0 with the XSWGR1100 Extension Pack (May 1999), and HP-UX 11i

DMP-supported arrays are a subset of arrays which Volume Manager supports. For a current and accurate list, please refer to **www.veritas.com**.

- HP AutoRAID disk array with SCSI
- EMC Symmetrix with fibre channel and SCSI when configured in the *Common Serial Number Mode* before installing VERITAS Volume Manager
- HP SureStore E Disk Array XP256
- HP FC1010D with fibre channel
- HP NIKE disk array models 10, 20, 30 with fibre channel and SCSI
- Hitachi Data Systems 5700E disk array subsystem with SCSI
- Hitachi Data Systems 7700E disk array subsystem with fibre channel and SCSI
- JBOD

VERITAS™

V E R I T A S W H I T E P A P E R

VERITAS Software Corporation
Corporate Headquarters
350 Ellis Street
Mountain View, CA 94043
650-527-8000 or 800-327-2232

For additional information about VERITAS Software, its products, or the location of an office near you, please call our corporate headquarters or visit our Web site at
www.veritas.com

VER06-DMPWPR-0000 • 90-01455-399