VERITAS™

VERITAS Dynamic Multipathing

Increasing the Availability and Performance of the Data Path

VERITAS™

TABLE OF CONTENTS

I/O Path Availability and Performance

Whether accessing data on a direct attached storage (DAS) device or through a storage area network (SAN), the I/O path is as important to application availability and performance as the server and physical disks. A connectivity failure between server and storage will impact application availability just as much as a failure in a database or a physical component of the storage subsystem. And, insufficient bandwidth for I/O traffic creates performance problems every bit as significant as an overwhelmed switch in the IP network.

Avoiding a single point of failure in the I/O path requires redundancy. As with other components of the infrastructure, a backup must be available to support I/O traffic if the primary route to storage fails. However, simply adding extra host bus adapters (HBA) to a server does not solve this problem. To avoid the integrity-compromising duplicate disk and LUN images, which occur when multiple HBAs exist on a server, and to ensure I/Os are shepherded along the appropriate path to a storage device, sophisticated multipathing I/O management software must be available on the server.

Solutions to I/O path availability can also solve I/O path performance problems. With added connectivity between server and storage, the bandwidth available to I/O traffic also increases. Managed correctly, this additional bandwidth can help to eliminate bottlenecks in the I/O path and ensure consistent access speed between the server and storage. As with the availability problem, however, successfully exploiting the additional bandwidth requires intelligence at the server that can monitor every route to storage and dynamically alter the path of each I/O.

Dynamic Multipathing

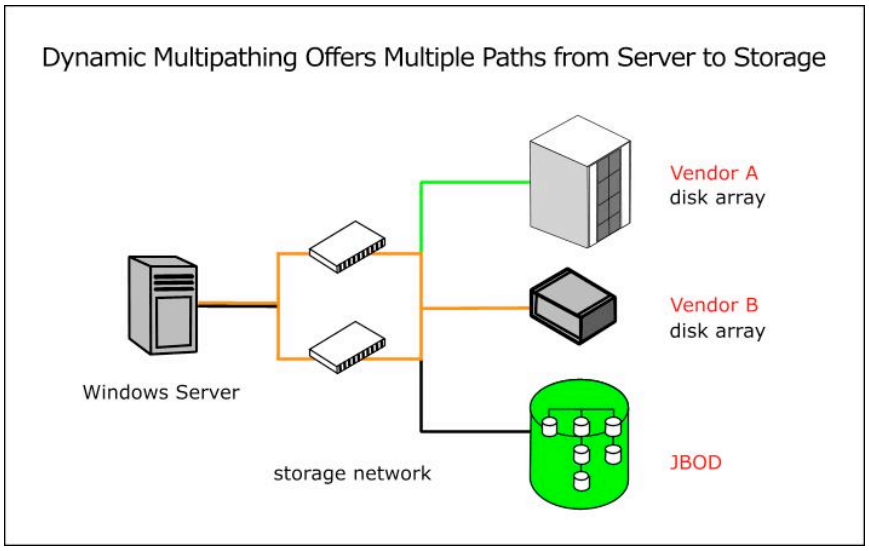

Figure 1: Dynamic Multipathing Offers Multiple Paths from Server to Storage

Dynamic multipathing (DMP) software provides the front-end intelligence to manage multiple I/O paths between a server and the LUNs and disks in a storage subsystem. Without DMP software the server operating system (OS) presents multiple images of a disk, or LUN, to applications — one image for each I/O path discovered by the OS. And, if traffic along the available paths is not carefully managed, I/O conflicts can cause data corruption problems.

The DMP software manages failover between the I/O paths in the event of an outage. Configuring each I/O path for full redundancy — no infrastructure component relied on by both paths for connectivity – eliminates the potential for a single point of failure. If connectivity along one path to a storage device is interrupted, the DMP software dynamically switches I/Os to a surviving path, allowing application access to continue unimpeded.

DMP software can also improve application performance. By leveraging the presence of additional paths into the storage infrastructure, the DMP software increases the available bandwidth for I/O traffic.

DMP solutions are available from most storage hardware vendors, including EMC and HP. Independent testing, however, has shown that the DMP software available with the VERITAS Storage Foundation Suite performs better than some of the competing array-based solutions.

VERITAS Storage Foundation for Windows Dynamic Multipathing

VERITAS Storage Foundation for Windows provides mission-critical, enterprise environments with sophisticated storage management features to reduce planned and unplanned outages and dramatically increase end-user access to data. Storage Foundation maximizes administrator productivity, ensures high application data availability, and provides the heterogeneous support needed to protect current and future investments in storage technology.

VERITAS Dynamic Multipathing is an optional component of VERITAS Storage Foundation for Windows offering fault tolerance and round-robin load balancing features to improve the robustness and performance of the I/O path. The DMP software supports an unlimited number of paths from a server into the storage subsystem, and manages the interface between the host server file system and two or more HBAs.

VERITAS Dynamic Multipathing Features

- Fault Tolerance
- Load Balancing
- Dynamic Recovery
- Automatic Path Recognition
- Granularity
- Maintenance Modes

Dynamic path recognition auto-detects the presence of I/O paths — either at reboot or on a rescan — and makes new paths available for failover and load balancing. At any time, only one path into the storage subsystem is accessible from the host server, avoiding data corruption problems.

VERITAS DMP eliminates I/O path single point of failure problems by managing all routes between the server and storage subsystem. If the primary I/O path fails, the VERITAS DMP software flags the problem path and activates a secondary route into the storage subsystem. All new I/Os are directed to the new path. The DMP software dynamically recovers I/Os that were in-transit on the failing path to ensure successful completion on the new path. Each path is continually monitored and as soon as a failed path becomes available again service is automatically restored. An indicator of path status is available instantaneously to the storage administrator through the VERITAS Enterprise Administrator (VEA) graphical user interface (GUI).

Administrators have the flexibility to enable or disable I/O paths at the device level — a disk or LUN — or for an entire storage array, allowing paths to be temporarily stopped during maintenance operations. This gives exceptionally granular control of the DMP configuration.

DMP Modes

VERITAS Storage Foundation for Windows supports three modes of DMP operation. An Active/Passive configuration, offers a conventional path failover setup, with one path being used for all I/O and an inactive second path held in reserve as a backup. In Active/Active mode all paths are available, and DMP software balances I/O traffic evenly between the resources. Active/Passive Concurrent is similar to Active/Active in that all paths to storage are used. However, unlike the round-robin approach of Active/Active, Active/Passive Concurrent mode binds each device to a specific path. The I/O load is then balanced across the available resources.

Active/Passive

An Active/Passive multipathing configuration provides a server with a designated preferred path for I/O traffic. The preferred path is always active and secondary paths are maintained in standby mode and only activated if the primary connection fails.

Configuring a disk array as Active/Passive implies that a LUN, or disk, can only be bound to one of the array's controllers at a time. If the array receives I/O for a LUN simultaneously from two different controllers it will shift ownership of the LUN between the controllers in a ping-pong effect that degrades performance. The VERITAS DMP software avoids this situation by implementing a policy for Active/Passive disk arrays that channels I/O traffic along the primary path, as long as it is available. I/O is sent along the secondary path only if the primary is fails.

Active/Active

Active/Active multipathing does not bind a controller to a specific path. In Active/Active mode I/O travels along any of the defined routes to the storage array.

The VERITAS DMP software performs load balancing in an Active/Active configuration by writing each consecutive I/O to a different path — working through the list of available paths sequentially. If a path fails it is simply omitted from the list of available paths until it becomes active again. This method is also known as a round-robin approach and improves performance by spreading I/O requests equally across all available paths.

Active/Passive Concurrent

An Active/Passive Concurrent multipathing configuration binds each LUN to a specific storage array controller. This allows I/O to be streamed to multiple controllers, along multiple paths, simultaneously. The performance benefits from Active/Passive

The DMP framework

Core DMP Driver (VxDMP.sys)

VxDMP.sys is an upper filter driver for Disk.sys and ScsiPort.sys. IOCTLs sent to these drivers are intercepted and redirected to the active I/O path. VxDMP.sys interfaces with DMP Array Support Libraries to identify LUNs belonging to an array, monitor path status, and perform failover.

Array Support Library (ASL)

An array specific set of common functions:

- Claim Device
  When a LUN is presented to the OS, DMP accesses ASLs to determine which array the device belongs to. If an ASL does not claim the LUN, it will not be managed by DMP. Once LUN ownership is established, the ASL identifies the array type (A/A or A/P), name, cabinet, and device serial number, and the path type (Active, Preferred…). If LUN ownership is involved, the Claim Device routine returns the path as A/P.

- Test Path
  Places DMP managed devices on a list so that path status can be monitored.

- Fail Over
  Performs a failover operation based on an array's specification

VxDMP Provider

Interfaces with VxDMP.sys through a private set of IOCTLs to get/set DMP information. Updates DMP relevant information displayed in the VEA GUI.

- Include/Exclude (array/device)
- Load Balancing (A/A, A/P)
- Monitor interval (default = 7 seconds)
- Set preferred path in A/P mode
- Enable/Disable path for maintenance
- Purge Disk

PnP5 Provider

Checks disk shadow whenever the VEA bus is initialized/re-initialized to see if a disk is a copy of another disk. It will create a disk object in the VEA bus if a disk is not a shadow (duplicate) of another disk.

VxDMP Path

A virtual path that is defined by Port Number, Path ID, Target ID and LUN ID.

Concurrent mode are similar to those of an Active/Active configuration, but combining an Active/Passive Concurrent DMP configuration with disk striping can improve I/O performance even further. By striping a volume across disks that have been bound to different controllers, I/Os will utilize the available paths regardless of sequence.

VERITAS DMP and Windows Clustering

Microsoft Windows clustering technologies — MSCS and VCS — use SCSI-2 reservations on disks to protect against data corruption and split brain — a situation where both servers in a cluster are active but think the other is down. The SCSI-2 reservations establish ownership of a disk along the entire I/O path. This presents a problem if the I/O path between a host and a disk is switched, as would occur during normal operations with an Active/Active DMP configuration. VERITAS Storage Foundation for Windows only supports Active/Passive and Active/Passive Concurrent DMP modes in a cluster environment, and the administrator must ensure that devices in a cluster are set to Active/Passive.

Although some storage array manufacturers claim support for Active/Active configuration in a clustered environment the solutions offered are not pure SCSI-2-based Active/Active multipathing. Vendors who make these claims usually have solutions that fall into one of two categories: pseudo Active/Active configurations that use multiple Active/Passive paths, each path going to selective disks; and, SCSI-3 host-based device reservations, which rely on the array supporting the SCSI-3 command set.

VERITAS Storage Foundation for Windows supports Active/Passive Concurrent multipathing — also known as Dual Active multipathing — in a clustered environment. This configuration supports DMP at the device level and provides the redundancy necessary for I/O path fault tolerance while allowing all paths to be fully utilized. For example, a clustered system with two paths into a six-disk storage array can be configure so that the first path is designated the preferred path for odd numbered disks and the second is designated the preferred path for even numbered disks, allowing both paths to be active in the cluster with each providing a failover target for the other.

Testing Multipathing Performance

Internal VERITAS Benchmark tests have shown that Active/Active multipathing configurations perform better than Active/Passive for applications that issue large I/O — greater than 64 KB. Applications that perform smaller writes — 8 KB or less — show negligible difference in performance between the two DMP configuration modes.

Comparing VERITAS DMP and EMC PowerPath

In January 2004 VERITAS commissioned the independent testing organization VeriTest to conduct a series of performance tests, comparing the VERITAS Storage Foundation for Windows multipathing solution against EMC's PowerPath product in a Microsoft Cluster Server environment. The tests reveal important differences between the two approaches to multipathing.

Using the Microsoft Benchcraft TPC-C software, provided with the Microsoft SQL Server TPC-C Benchmark Kit, VeriTest executed a series of tests using both out of the box and vendor recommended configuration settings for the two multipathing solutions.

VeriTest Testbed Configuration

VeriTest configured two identical testbeds to compare VERITAS Storage Foundation for Windows DMP and EMC's PowerPath. Each testbed consisted of the following:

- Dell PowerEdge 4400
  - Dual 1GHZ Pentium III Xeon processors
  - 2 GB of RAM
  - Windows Server 2003 Enterprise Edition, with all of the latest updates
  - Microsoft Cluster Server
- QLogic QLA-2342 HBA
- EMC Clariion CX600
- Brocade 3800 2Gb 16-port switch.

The Benchcraft TPC-C software generates two performance metrics. Tpm measures total number of transactions — new orders, payments, delivery, stock level, and order status — per minute. And, TpmC measures order entry processing only, that is, the number of new orders the system is able to process in one minute.

Key Findings

The VeriTest benchmark testing revealed that the VERITAS multipathing solution delivered better Tpm and TpmC scores, in both out-of-the-box and vendor-recommended configurations. VeriTest also noted that the VERITAS GUI interface was used for all VERITAS Storage Foundation for Windows configuration operations, whereas the EMC PowerPath product required configuration through a command line interface.

TpmC Scores	Out of the Box Configuration	Vendor-Recommended Configuration
VERITAS DMP	3661	3661
EMC PowerPath	3240	3328

Figure 2: VeriTest TpmC Results

Tpm Scores	Out of the Box Configuration	Vendor-Recommended Configuration
VERITAS DMP	8557	8555
EMC PowerPath	8508	7825

Figure 3: VeriTest Tpm Results

Conclusion

Dynamic multipathing solutions solve the problem of a single point of failure in the I/O path. These solutions also improve I/O performance by allowing applications to benefit from the increased bandwidth available from multiple paths between a server and the storage subsystem.

VERITAS Storage Foundation for Windows Dynamic Multipathing provides the fault-tolerance needed by highly available, mission-critical applications. And, as verified in independent benchmark testing by VeriTest, the VERITAS DMP solution offers highly competitive performance.

VERITAS Software Corporation

Corporate Headquarters
350 Ellis Street
Mountain View, CA 94043
650-527-8000 or 866-837-4827

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.