

where information lives

EMC²
where information lives

PowerPath Configuration and Administration



EMC Global Education

© 2004 EMC Corporation. All rights reserved. These materials may not be copied without EMC's written consent.

Mercer Road
Global Education & Productivity

1

Welcome to PowerPath Configuration and Administration.

PowerPath Configuration and Administration

After completing this course, you will be able to:

- **Technical Overview**
 - Describe PowerPath architecture
 - Identify PowerPath load balancing policies
 - Describe PowerPath behavior difference in Active-Active and Active-Passive storage systems
 - Identify key features of PowerPath Volume Manager
 - Identify key features of PowerPath Volume Mobility
- **Configuration**
 - Describe PowerPath connectivity and configuration in both SCSI and Fibre Channel environments
 - Identify most configuration changes that affect PowerPath
- **PowerPath Management**
 - Manage PowerPath license
 - Configure PowerPath
 - Monitor the device under PowerPath control
 - Restore dead path
 - Customize load balancing and device priority
 - Secure and reuse custom PowerPath settings
 - Remove a path from PowerPath control

The objectives for this course are shown here. Please take a moment to read them.

Technical Overview

Section One



EMC Global Education

© 2004 EMC Corporation. All rights reserved. These materials may not be copied without EMC's written consent.

Mercer Road
Global Education & Productivity

3

This section provides a technical overview of PowerPath.

where information lives

EMC²
 where information lives

PowerPath

Automatic, Non-disruptive, Optimized

PowerPath Delivers:



**Intelligent Path
Management**



**Integrated Volume
Management**

PowerPath Features Are:

- **Automatic**
 - Intelligent, self-sensing infrastructure management to reduce cost and complexity
- **Non-disruptive**
 - Application-transparent functionality insulates core business operations
- **Optimized**
 - Superior performance with current IT investment

EMC Global Education
 © 2004 EMC Corporation. All rights reserved.

Mercer Road
 Enhancing Education & Productivity

4


PowerPath combines multiple path I/O capabilities, automatic load balancing, path failover, and logical volume management functions into one integrated solution to work across your heterogeneous server and storage environment. PowerPath maximizes application availability, optimizes performance, and automates online storage management while reducing complexity and cost, all from one powerful data path management solution.

Automatic: PowerPath uses sophisticated algorithms to increase application I/O rates through Symmetrix, CLARiiON, and non-EMC arrays with automatic data path load balancing, allowing for greatest efficiency and throughput. PowerPath's Volume Manager capability simplifies disk administration tasks, thereby further reducing the total cost of ownership through high-level commands that hide storage complexity. It automatically manages workloads and volume expansion.

Non-disruptive: PowerPath provides users with application-to-storage automation, seamless data migration, and automatic import of Volume Groups, and it simplifies the growth of logical volumes while applications remain online. PowerPath optimizes server and data path utilization by avoiding downtime. PowerPath Volume Manager has enhanced its current offering and now provides application-transparent volume mobility between heterogeneous storage arrays.

Optimized: By leveraging your server, SAN, and storage assets, PowerPath maximizes your investment by increasing storage utilization.

PowerPath Implementation



EMC Global Education
 © 2004 EMC Corporation. All rights reserved.

 Mercer Road
 Unified Education & Productivity

5

PowerPath provides continuous application availability for the customer's environment, making sure that applications remain online all the time.

PowerPath supports Symmetrix, CLARiiON, Hitachi Data Systems Lightning, HP StorageWorks XP, and IBM Enterprise Storage Server (Shark) storage arrays.

Please consult the "EMC Support Matrix" for current operating system and array support.

PowerPath Features – page 1 of 2

- Automatic path failover
- Automatic detect and restore of failed components
- Dynamic load balancing
- Volume Manager operations*
- Automatic recognition of LUN hardware growth*
- Seamless import of same-host BCV devices*

Automatic path failover: By taking advantage of all paths, PowerPath increases the ability of server access to data on the Symmetrix, CLARiiON, and non-EMC arrays. PowerPath moves I/O workloads across multiple paths to ensure fastest possible I/O speed.

Automatic detect and restore of failed components: PowerPath can automatically detect a path failure using Auto Detect. For instance, if a cable is mistakenly dislodged, PowerPath will take all existing I/O that was going down that particular path and re-route it to another active path, so you maintain application availability and access to your data continuously. Once the cable is re-attached, another feature of PowerPath, called Auto Restore, automatically restores path access permitting data flow down the path again.

Dynamic load balancing: When the host needs access to data on the Symmetrix, CLARiiON, and non-EMC arrays, PowerPath uses the quickest and easiest route or path to access the data. PowerPath supports several load balancing policies compared to other vendors. In the rare instance of a path failure, PowerPath will re-issue I/O to an alternate path maintaining data availability and ensuring information access. PowerPath offers policy-based dynamic path management that accelerates information access and provides high availability.

Volume manager operations: PowerPath's Volume Manager performs other functions like:

- Physical-to-logical device abstraction
- Slicing — Take a disk and “slice” into smaller pieces (i.e., creating volumes out of it)
- Striping — Data can be striped across multiple underlying physical disks. Striping enhances performance, since multiple I/Os are processed in parallel.
- Mirroring — One or more copies of data can be created and kept in sync with the source data. Data can be mirrored across multiple logical devices, which can be distributed throughout the storage network. Mirroring improves data availability; if one mirror fails, PowerPath automatically directs I/O to the other mirror image. PowerPath supports software and hardware mirroring.

PowerPath Features – page 2 of 2

- Application-transparent volume mobility*
 - User-selectable storage allocation policies*
 - Online configuration and management
 - Common HBA driver support
 - Write throttling
 - Consistency Group support (Symmetrix only)
- * New features in PowerPath V4.0 and beyond

Automatic recognition of LUN hardware growth: When a LUN is expanded, PowerPath recognizes the LUN growth and automatically expands the Storage Element corresponding to that LUN, while applications remain online.

Seamless import of same-host BCV devices: One of PowerPath's unique Volume Manager advantages is its ability to import a copy of a Volume Group to the same host. This can be used to backup the volumes in a Volume Group on the same host, using a hardware mirrored copy of the Volume Group. Same-host import is more automated (easier and less error-prone) on PowerPath than on VERITAS or other volume managers. For instance, PowerPath automatically re-labels volumes.

Application-transparent volume mobility: PowerPath Volume Manager has enhanced its offering to enable you to move or change the structure of a PowerPath Volume Manager volume without application disruption.

User-selectable storage allocation policies: PowerPath Volume Manager provides users with these policies (more to come in future releases).

- Storage Element policy specifies how storage is allocated when a volume is created or expanded.
- Most space policy uses Storage Elements with the most free space first. This optimizes for parallel I/O performance.
- Best Fit policy uses Storage Elements with the least free space first. This optimizes for space utilization by minimizing fragmentation.


Software mirrors on redundant hardware: Controls what happens if the user tries to create a software mirrored volume on a Storage Element that is hardware redundant (refers to hardware mirroring or hardware-implemented RAID levels that are redundant). This policy helps you avoid using redundant hardware for software mirrors, which is not an optimum use of your system's storage capacity. For instance, if you already have a mirrored disk on the Symmetrix system and you try to create an additional mirror using PowerPath Volume Manager, it can be configured to issue a warning or prevent the action.

Online configuration and management: You can have some paths active and some passive (or "standby"). This means that the passive paths will only be used if the active paths fail. A user can segment applications. For example, in an eight-path configuration running two applications, you can have each path active for one application and passive for the other. PowerPath allows the user to make configuration changes while applications remain online, meaning you can reconfigure the load balancing policy, take paths offline and online, hot repair, etc. In addition, PowerPath can automatically resize or move a volume. Storage Elements can be added to or deleted from a Volume Group in real time without requiring application downtime.

PowerPath also supports common HBA driver connectivity on the Symmetrix, CLARiiON, and non-EMC arrays, which means that one host can share the same HBA to one or all arrays. PowerPath also supports write throttling, which is a way to minimize or eliminate database read timeout errors, by throttling the write errors so that a read error can be processed.

For a group of Symmetrix devices, PowerPath will maintain the integrity of a database distributed across multiple SRDF units. PowerPath also supports the ability of booting from a mirrored R2 device in an SRDF configuration.

Path Fault without PowerPath




EMC Global Education
 © 2004 EMC Corporation. All rights reserved.

Mercer Road
 Enter Education & Productivity

8

Without PowerPath or another path failover software package, the loss of a channel (red dotted line) means one or more applications may stop functioning. This can be caused by the loss of a Host Bus Adapter, Channel Director, Storage Processor, FC-AL hub, or a failed (kicked) cable. In a standard UNIX environment, this is a single point of failure. In this case, all I/O that was heading down the path highlighted in red is now lost.

Path Fault with PowerPath



EMC Global Education
© 2004 EMC Corporation. All rights reserved.

 Mercer Road
Global Education & Recruitment Ltd.

9

PowerPath is critical in this situation. PowerPath will ensure that the I/O heading down the failed path is re-directed (keeping in mind load characteristics and priorities) to the most suitable alternative path. Unlike other path management products that will failover all I/O to the next available path (overburdening the path), PowerPath uses intelligence to decide where to failover the I/O.

To PowerPath, the loss of a channel means the load must be balanced across n-1 channels. This doesn't mean that performance will not be affected, but the applications will continue to operate at their best without down time.

Because PowerPath spreads the additional load across the remaining channels, the system will continue to operate, however, in a “degraded” mode. This takes the pressure off the staff to make corrections immediately. Since some O/S servers do not support hot plugging of boards, the replacement of an HBA may require bringing the server down. PowerPath allows the IT staff to schedule the repair at a convenient time.


If you are concerned with degraded mode behavior, the PowerPath GUI provides the ability to disable and re-enable channels on the fly so you can test the performance sensitivity of your configuration.

Simple channel failover products work in pairs, i.e., one channel is the backup for the other. In the diagram, two pairs of channels would be set up. If one channel fails, all the load for that channel goes to its “twin,” essentially doubling its load. With PowerPath, the load is spread across the three remaining channels.

The objective is to make channel failover a non-issue. There are low cost products that provide two channel failover. By making queue management the issue, PowerPath is not competing with channel failover products and establishes the value of the PowerPath solution. The ability to compensate for a lost channel by spreading the I/Os across multiple channels instead of dumping all of it on one requires multi-channel support.

I/O without PowerPath

Queues Out of Balance



EMC Global Education
 © 2004 EMC Corporation. All rights reserved.

 Mercer Road
 Enhancing Education & Productivity


10

This diagram depicts a snapshot of the system at a moment in time. The depth of the I/O queues is very unbalanced. Host applications sitting on top of deep queues are not getting the data they need. If this was the average loading, the System Administrator would reconfigure the system to balance the load better. In any system, there will be times when the load is unbalanced due to one application receiving heavy I/O requirements.

In this instance, two of the applications are currently causing high I/O traffic. At this point, two channels are overloaded (depicted by the red disks on the outside, and the pending request stack) while two other channels are lightly loaded. In a while, the requests will have been handled and the system will return to a more balanced load. In the meantime, the applications are being “data starved” and the users or applications are experiencing less than optimal performance.

I/O with PowerPath

Queues in Balance



EMC Global Education
© 2004 EMC Corporation. All rights reserved.

 Mercer Road
Unite Education & Productivity

11


With PowerPath in the system, applications transparently access PowerPath instead of the SD (SCSI driver) devices. PowerPath allocates the requests across the available channels, reducing bottlenecks and improving performance. This diagram shows a similar snapshot, with PowerPath using multiple channels to minimize the queue depth on all channels.

Since the Channel Directors or Storage Processors are writing to cache and not to disks, any Channel Director/Storage Processor can handle any request. This allows PowerPath to constantly tune the server to adjust to changing loads from the applications running on the server.

PowerPath improves the performance of the server, enabling it to make better use of the storage. This results in better application performance and less operational resources spent on the care and feeding of the system and more (financial) value from your server investment.

PowerPath does not manage the I/O queues; it manages the placement of I/O requests in the queue.

Performance Tuning and Active/Standby Channel Policies



Available bandwidth can be segmented per application

EMC Global Education
© 2004 EMC Corporation. All rights reserved.

Mercer Road
Enter Education & Productivity

12

You can have some paths active and some passive (what we call "standby"). This means that the passive paths will only be used if the active paths fail; thus, you can also segment applications if you wish. For example, in an eight-path configuration running two applications, you can have each path active for one application and passive for the other. All of this can be done through PowerPath dynamically, meaning you can reconfigure the load balancing policy, take paths offline and online, hot repair, etc.

In the case where you segment an application, you can dynamically add and remove paths as well. For example, if you had eight paths partitioned into four for one application and four for another, and one application needed more bandwidth, say during batch, you can, through scripts, add more paths during the batch run and take them away when you're done.

The actual command is `powermt set mode=standby/active adapter=adapter# dev=integer/all`

Here are examples of two practical applications of this:

- If you have two applications on a server that are both very busy, but you don't want the activity of one application to impact the performance of the other application, you could set up the two applications to have their own sets of channels. I/O for those two channels would be balanced among the channel groups only, and if one application was very busy, the other application would not be impacted from an I/O standpoint. In the event of all the defined channels for one Channel Group failing, I/O would be re-routed to the other Channel Group's channels.
- If a host has two types of channels, for example, Fibre Channel and SCSI, you would have to segregate channels by channel type. It is not possible for PowerPath to load balance an application across both Fibre Channel and SCSI together, but by using channel segregation, a user can set up the environment so that PowerPath would load balance across the SCSI channels for one set of devices and Fibre Channel for the other set of devices. Then, in the event of a failure, all the devices defined to the failed paths would use the other available paths for I/O, because they were standby paths for those devices.

PowerPath Load Balancing Policies

- **Symm_opt/CLAR_opt/Adaptive (default)**
 - I/O requests are balanced across multiple paths based on composition of reads, writes, and user-assigned device/application priorities
- **Round_Robin**
 - I/O requests are distributed to each available path in turn
- **Least_I/Os**
 - I/O requests are assigned to the path with the fewest number of requests in the queue
- **Least_Blocks**
 - I/O requests are assigned to the path with the fewest total blocks in the queue
- **Request**
 - Path failover only
- **No Redirect**
 - Disable path failover and load balancing (same as Transparent=on)
 - Default for Symmetrix when there is no license key; not supported on CLARiiON
- **Basic Failover**
 - CDE-type functionality (SPA–SPB failover)
 - Default for CLARiiON when there is no license key; not supported on Symmetrix

Symm_opt/CLAR_opt/Adaptive —Choose paths based on path load and logical device priority. Device priority is host based priority and not Symmetrix/CLARiiON level priority. Adaptive is the default priority on non-EMC arrays and performs similarly to Symm_opt and CLAR_opt.

Round Robin —I/O requests are assigned to each available path in rotation.

Least I/Os —Load balance is based on the number of pending I/Os. I/O requests are routed to the path with the fewest queued requests, regardless of total block volume.

Least Blocks —Load balance is based on the number of blocks in pending I/Os. I/O requests are routed to the path with the fewest queued blocks, regardless of the number of requests involved.

Request —For native devices, this policy uses the path that would have been used if PowerPath were not installed. For pseudo devices, it uses one arbitrary path for all I/O. For all devices, path failover is in effect, but load balancing is not.

No Redirect —Neither load balancing nor failover is in effect. If this algorithm is set on a failed path and a native device is used, I/O errors will occur when I/O is directed to that path. This policy is valid only on Symmetrix arrays.


Basic failover —Load balancing is not in effect. I/O routing on failure is limited to one host bus and one port on each Storage Processor. This policy is required for a non-disruptive upgrade. It protects against Storage Processor and back-end failures, but not against HBA or host loop failures. PowerPath is required to prevent HBA and/or host loop failures.

You can change the policies using the powermt set policy command, or within the GUI on Windows.

Active-Active vs. Active-Passive

CLARiiON's "LUN Ownership Model"

- Dual-port disks, redundant FC-AL loops
 - All disks have paths to both Storage Processors
- CLARiiON LUNs are "owned" by one Storage Processor
 - One Storage Processor services I/O for LUN
 - Second path to LUN is passive



* Symmetrical Volume Access:
 volume accessible (active) on
 all paths

* Asymmetrical Volume (LU)
 Access: LU accessible (active)
 on one Storage Processor
 at a time

The picture shows the number of paths connected from a host directly to a CLARiiON. There are four available paths to the CLARiiON on this host. The green lines show active paths to a single LUN. The red lines show the passive paths to the same LUN.

PowerPath enhances CLARiiON high availability with I/O path failover software. During normal operations, the host utilizes two high-speed data paths.

In the event of a failure of a HBA, Storage Processor, or cable, PowerPath automatically reroutes traffic in seconds.

By taking advantage of both paths under normal conditions, PowerPath increases the availability of the server to access data on the CLARiiON. All this is accomplished while the application remains online.

Access to a LUN takes place on one or more active paths through the default Storage Processor. An equivalent set of passive paths is also available via the peer Storage Processor.

Native Device

- A native device is the device that the host OS created if PowerPath is not installed
 - For example:
 - /dev/rdisk/c1t2d3s2 on Solaris
 - /dev/rhdisk2 on AIX
- PowerPath may alter the behavior of a native device, e.g., load balancing

The operating system creates native devices to represent and provide access to logical devices. A native device is path specific (as opposed to path independent) and represents a single path to a logical device. The device is native in that it is provided by the operating system for use with applications.

On Windows, native devices are not exposed to users. PowerPath does not support native devices On AIX. Thus applications that do not use AIX LVM volumes must be reconfigured to use a pseudo device or LVM volumes for PowerPath to take effect.


PowerPath operates at a lower level in the I/O stack than the Logical Volume Manager. Thus, the load balancing policy of the LVM may be affected by that of PowerPath .

Pseudo Device

What's a Pseudo Device?

- One or more native paths to the same storage volume are presented as a single pseudo device
- Provides load balancing, failover, and consistency group support
- Required on some platforms to avoid name collisions

PowerPath Driver



A PowerPath pseudo device represents a single logical device which can contain any number of physical paths to the same storage device. The pseudo device in this example has 3 native paths to the physical drive. The use of pseudo devices is not required for PowerPath to take effect, except for AIX in which case the applications must be reconfigured.

PowerPath Load Balancing with CLARiiON

LUN Ownership Model with CLARiiON

- Load balancing
 - Load balancing and failover across active channels
 - Trespass to passive channels in event of all paths/Storage Processor failover
 - Only on systems with two or more FC-SW (CLARiiON FC4700 and above) ports per Storage Processor
 - For HBAs on active paths only
 - Not across Storage Processors
- Customer selects best algorithm for application requirements
 - CLAR_opt is usually recommended

The LUN ownership model means that only one Storage Processor can own a particular LUN at any one time. For the other Storage Processor to gain access to the LUN, it must take ownership of the LUN.

Some people describe this as “active-passive” (on active-passive storage systems, “available” paths are those paths leading to the active Storage Processor for each logical device). However, there can be many LUNs on a CLARiiON and each of these LUNs can be assigned to either Storage Processor. PowerPath can distribute the I/O for all logical devices over all paths shared by those logical devices, so all paths are equally burdened.

PowerPath load balances I/O on a host-by-host basis. It maintains statistics on all I/O for all paths. For each I/O request, PowerPath intelligently chooses the least-burdened available path, depending on the load balancing and failover policy in effect. If an appropriate policy is specified, all paths in a PowerPath system have approximately the same work load.

If no preference is given, PowerPath will choose what it thinks is the best method of load balancing.

Failover Redefined with PowerPath

- Use all active paths before doing a trespass
- When all active paths are exhausted, trespass on a LUN by LUN basis
- Scope of LUNs failing over depends on component
 - If there is a failure to an HBA or cable connecting the HBA to a Storage Processor, hub, or switch, just the affected hosts' LUNs fail over
 - If there is a failure to a Storage Processor, all LUNs assigned to that Storage Processor will fail over if they are accessed

Failover occurs on a LUN-by-LUN basis. This helps avoid any unnecessary trespassing of devices that can cause possible performance issues and degradation.


A multi-path failover will re-route the I/O through an alternate path to the LUN through the Storage Processor that is the primary owner. If the LUN is still inaccessible after going through all possible paths, a LUN trespass will occur. Trespass means that LUN ownership has been transferred to the alternate Storage Processor.

The manner in which logical units fail will also depend on the mode of failure within the data path. If there is a failure of the HBA, or anywhere between the server and the array, just the affected host's devices would trespass over to the other Storage Processor, so they are accessible to the server through the alternate data path.

If there is a Storage Processor failure, all devices assigned to that Storage Processor that received I/O from the requesting host will fail over to the surviving Storage Processor.

PowerPath Failover with CLARiiON

- Path fault
- HBA fault
- Storage Processor fault



PowerPath protects the customer's environment in case of a path, HBA, or Storage Processor failure.

Path fault —If a path or a cable is accidentally bumped or pulled, PowerPath will automatically fail over the path to an alternate path. Once the path(s) have been restored, PowerPath automatically recognizes this and resumes sending I/O down those paths.


HBA fault —If an HBA is lost at the front end, PowerPath will fail over the entire path to an alternate HBA. Once the HBA is brought back online, PowerPath recognizes this and resumes sending I/O through the original HBA.

Storage Processor fault —If SP-B fails (see diagram), PowerPath will stop all I/O that was currently headed from the server to SP-B and trespasses it over to SP-A; once SP-B is brought back online, PowerPath recognizes that SP-B is now available and resumes sending I/O down to SP-B.

In the event that any of the above scenarios takes place, PowerPath ensures that applications remain continuously available.

CLARiiON—No PowerPath License

- Provides back-end (SP) failover, non-disruptive upgrades, and single HBA configuration support for single-attach hosts
- Is included in CLARiiON Utility Kit
- Non-HA environment
- No multi-pathing capabilities
- The HBA must see both SPs



EMC Global Education
 © 2004 EMC Corporation. All rights reserved.

Mercer Road
 Senior Education & Productivity

20

A special case for non-high availability environments. PowerPath replaces CDE functionality, is provided on all CLARiiON arrays, and provides single HBA configuration support, back-end failover support, and non-disruptive upgrade support.


This does not have any multi-pathing or load balancing capabilities, and should not be used in a high availability environment.

where information lives

EMC²
where information lives

Intelligent Path Management for SANs

- Automatic detection of SAN switch-to-storage path failures
- Maintain maximum host HBA performance



EMC Global Education
© 2004 EMC Corporation. All rights reserved.

Mercer Road
Sustaining Education & Productivity

21

PowerPath supports end-to-end path failover. If there is a path failure between switch and storage (as shown in the slide), PowerPath recognizes the loss of connectivity and automatically fails over the path to an alternate path between the switch and the storage.


In the past, a switch-to-storage path failure would have taken down the entire associated end-to-end path. This would have propagated back up to the server and caused an HBA failover. Today, PowerPath allows all the host HBAs to remain fully active, ensuring the utmost in application performance. Once the path has been repaired, PowerPath automatically restores the path connection.

PowerPath ensures that SAN failures are managed locally and there's maximum application performance.

where information lives

EMC²
where information lives

PowerPath Volume Management

EMC Global Education
© 2004 EMC Corporation. All rights reserved.Mercer Road
Unite Education & Productivity

22

PowerPath now includes host-based logical volume management capabilities, which presents a simple, flexible view of data that is independent of the underlying physical storage. Using a volume manager, Storage Administrators place physical storage devices that share a common configuration into logical pools called Volume Groups. The administrator then uses these Volume Groups to manage the physical storage independently of how applications and users access the data.

PowerPath's Volume Manager feature lets users build their filesystems utilizing multiple storage systems without application interruption. PowerPath ensures that users utilize all the power and investment made in existing server and storage investments.

A Storage Element is a physical piece of storage corresponding to a LUN. A Storage Element can be part of or all of a LUN, but it cannot cross LUNs. Each data image is stored on one or more Storage Elements.

Creating a Volume Group: Volumes exist in the context of Volume Groups. A Volume Group can be created from LUNs identified by the user. Users can select and provision LUNs, make them visible to a host, and add them to a Volume Group on that host. LUNs are mapped one-to-one to Storage Elements in a Volume Group. Storage Elements can be added to, or removed from, a Volume Group at any time when not in use. However, Storage Elements can't be shared across different volumes.


Creating a volume (shown in slide): To create a volume using PowerPath's Volume Manager capability, the user first identifies a Volume Group in which to create the new volume. The user has created new volumes from free Storage Elements in the Volume Group. As you can see, Volume 1 comprises all of Storage Element 2 and part of Storage Element 3 (in yellow). Volume 2 comprises all of Storage Element 1 and the remaining part of Storage Element 3 (blue).

When creating volumes, the user needs to specify attributes of name, size, and volume structure (for instance, software mirroring). The administrator is not required to select which free space in the Volume Group is used, but does have the option to specify which Storage Elements compose a volume. By default, PowerPath Volume Manager uses Storage Elements with the most free space first. Alternately, users can specify that PowerPath use the Storage Elements with the least possible free space (allowing for the best fit).

where information lives

EMC²
where information lives

Dynamic LUN Expansion

EMC Global Education
© 2004 EMC Corporation. All rights reserved.Mercer Road
Enter Education & Productivity

23

Volumes can be expanded while applications remain online using free space from one or more Storage Elements. In the slide, Volume 2 is expanded to use free space from Storage Element 3. An advantage of PowerPath's Volume Manager feature over other volume managers is its support for dynamic LUN expansion. In this figure, LUN 3 is expanded. PowerPath recognizes the LUN growth and automatically expands Storage Element 3 that corresponds to that LUN. The administrator can then choose to incorporate the added free space in the Storage Element by expanding an existing volume or adding a volume to the Volume Group. In the figure, the administrator expanded Volume 2. With other volume managers, when LUNs are expanded, the volume manager cannot use the extra free space automatically; manual Volume Group configuration is required.

A volume can be expanded using the PowerPath CLI to specify the new size. PowerPath uses free Storage Elements in the Volume Group to extend the size of the volume. EMC ControlCenter has a Graphical User Interface to monitor filesystems and to send an alert before they run out of free space. However, it is not required because PowerPath has a CLI that performs all similar functions.


Users can set up rules to cause ControlCenter to do the following:

- Provision new space when needed (from a designated, SAN-wide pool of available space)
- Expand or add LUNs (causing PowerPath to automatically expand corresponding Storage Elements)
- Direct PowerPath to expand volumes accordingly

where information lives

EMC²
where information lives

Simplify Import of BCV Devices to Other Servers



EMC Global Education
© 2004 EMC Corporation. All rights reserved.

Mercer Road
Unite Education & Productivity

24

PowerPath's Volume Manager feature supports volume import and deport to allow a Volume Group to be moved from one host to another. These functions also support versioning on PowerPath Volume Manager volumes.


Import—Storage Elements can be imported to a host as one Volume Group. On import, the host recognizes both the Volume Group and its individual volumes. PowerPath Volume Manager re-labels the Volume Group and the volumes in it when needed, to avoid naming conflicts with other volumes and Volume Groups on the host. On import, a new name can be specified for an existing Volume Group. PowerPath's Volume Manager automatically sets up paths for the volumes in the Volume Group and registers the volumes for use on the importing host. A Volume Group can be imported to a host in a remote location to provide remote failover.

Scenario—In the slide, after creating the mirror, the administrator splits the mirror from the original data, makes the mirrored LUNs visible on another host (Host B, in the figure), and then imports the mirrored data as a Volume Group on the other host.

where information lives

EMC²
where information lives

Seamless Import of Same-Host BCV Devices



SE = Storage Element

EMC Global Education
© 2004 EMC Corporation. All rights reserved.Mercer Road
Enter Education & Productivity


25

PowerPath's Volume Manager supports volume import and deport to allow a volume to be split and copied back to the same host as the source data.

Same-host Import —Of special note in PowerPath's Volume Manager feature is its ability to import a copy of a Volume Group to the same host. This can be used, for example, to backup the volumes in the Volume Group on the same host, using a hardware-mirrored copy of the Volume Group. PowerPath's Volume Manager automatically relabels volumes as a part of the import process.

Scenario —In the slide, after creating the mirrors, the administrator splits the mirror from the original data and makes the mirrored LUNs visible to the same volumes ("after," in the figure). PowerPath recognizes that there are two similar copies of data and, to prevent any data corruption (because the operating system would be confused over seeing two similar copies of data), it re-labels the Volume Group to "Volume Group 2."


Hardware-based Mirroring



In this figure, LUNs 1, 2, and 3 are mirrored remotely onto LUNs 4, 5, and 6 using LUN replication software (such as SRDF on Symmetrix systems). Hardware mirrors can be made accessible to the same host system as the source data, or to a different host system. To do hardware mirroring, all Storage Elements in a Volume Group must be from the same type of storage system (for example, Symmetrix systems).

You benefit from hardware-based mirroring, which is a simple and very effective mirroring option, which insulates the application from mirroring as it lowers your TCO.


Software-based Mirroring



When a software mirror, a new image, is added, users cannot access it independently. All reads are directed to a volume's preferred read image; this is either the source data (by default) or a mirror image designated by the user. Writes are propagated to the source data and all mirror images of it. All software mirrors of a volume are equivalent. When doing software mirroring, all volumes (source and mirror) must be in the same Volume Group, and all Storage Elements must be visible to the same host. The source and mirrored volumes, however, can be on different storage systems and have different layouts if desired.

- Pros:
 - Customer now have a way of putting heterogeneous Storage Elements (LUNs) into a single volume.
 - Customers don't require hardware services from storage.
- Cons:
 - Uses CPU resources on the host.
 - Each implementation is separately managed.

Manage Volume Striping



Data can be interleaved, or striped, across multiple underlying physical disks. If data from a volume is striped across four disks, for example, four I/Os are processed in parallel. Striping enhances performance, since multiple I/Os are processed in parallel. Striping is also known as RAID 0. (RAID stands for Redundant Array of Independent Disks or Redundant Array of Inexpensive Disks.)


If striping on a volume is desired, it is set up when the volume is created. In this figure, a volume is striped across three physical disks (LUNs). Data can be written to, or read from, the three LUNs in a round-robin fashion. For example, after starting a write to LUN 1, the system starts a write to LUN 2 without waiting for the prior write to complete.

The number, or width, of stripe columns can be changed by moving a volume to another volume that is set up with the desired striping; this can be done while applications are online. (Re-striping cannot be done in place.)

PowerPath Volume Mobility

Mobility Process:

1. An application accessing data through PPVM
2. A new array introduced into the SAN
3. Application data being migrated to the new array
- Ability to test application I/O
4. All I/O re-directed to the new array
5. No disruption to the application



EMC Global Education
 © 2004 EMC Corporation. All rights reserved.

Mercer Road
 Online Education & Productivity

29


PowerPath Volume Mobility allows an online application migration to from one storage to another. This is an integrated feature of PowerPath 4.X.

Scenario: The customer's lease on their existing Hitachi array is complete. They are planning to swap out the Hitachi and replace that with a Symmetrix array. PowerPath Volume Mobility can assist in this process.

- Step 1—System Administrator accesses data under the control of PowerPath Volume Manager.
- Step 2—The new Symmetrix array is introduced to the SAN.
- Step 3—System Administrator initializes the data movement and the host writes data to both the Symmetrix and Hitachi array simultaneously.
- Step 4—System Administrator tests application I/O throughput to the new array and determines if the redirection of data is optimal. At this point, the user will commit the data movement to the new array (Symmetrix) and sever the connection to the old array (Hitachi). The old array is retired and all application I/O throughput is being sent to the new array.
- Step 5—The entire mobility process takes place while the application remains online. The application doesn't care that the physical data has moved from one storage location to another. The application only needs to know that the logical path name has changed and now points to the new location.

PowerPath Volume Mobility

- An extension of PowerPath Volume Manager functionality that allows a user to:
 - Move a PowerPath Volume Manager volume from one set of SEs within a Volume Group to another set of SEs within the same Volume Group
 - Change the structure of a PowerPath Volume Manager volume
 - Applications using the volume remain *online*. *No down time*



Here is the set of tasks that a System Administrator would need to undergo to move or restructure a PowerPath Volume Manager volume.

- Step 1—Complete tasks before moving or restructuring a volume (provision target storage devices and add target storage devices to a PowerPath Volume Manager Volume Group)
- Step 2—Set up the mobility operation. This indicates what the volume will look like and where it will be located once the operation is complete.
- Step 3—Synchronize the target Storage Elements with the source Storage Elements.
- Step 4—When the volume is in “source selected state”, the user redirects read operations to the target Storage Elements to try out the new volume.
- Step 5—Once the user has determined to make the target Storage Elements the permanent storage for the volume, the user commits that operation.
- Step 6—When the mobility operation is complete, the source Storage Elements are now part of the Volume Group’s “free space”.

PowerPath Configuration Considerations

Section Two




EMC Global Education
© 2004 EMC Corporation. All rights reserved. These materials may not be copied without EMC's written consent.

Mercer Road
Global Education & Productivity

31

This next section covers PowerPath configuration considerations.

PowerPath Connectivity




PowerPath works with Fibre Channel physical connections. With Windows 2000 only, PowerPath also works with SCSI physical connections. In a Fibre Channel environment, Fibre Channel cables connect each Fibre Channel HBA to a port on a Fibre Channel interface on the storage system or to a Fibre Channel hub or switch.

(Hubs are not supported on all storage systems. See the EMC Support Matrix.)

In a SCSI bus environment, cables connect each SCSI HBA to a port on a SCSI interface on the storage system. Fibre Channel hubs insulate viable nodes from most hardware failures that can occur on another node in the loop. Switches provide similar functionality plus additional bandwidth for the duration of a transmission between nodes. To add switched fabric topology support to storage-system environments, EMC supports EMC Connectrix and other switches, providing greater levels of throughput, connectivity, and manageability. Working within the arbitrated loop and switched fabric topologies, PowerPath enhances high availability and performance.


SCSI Configuration



PowerPath supports SCSI physical connections for Windows 2000 hosts only. When configuring the SCSI physical connections for a PowerPath environment, follow these rules:


- For PowerPath to provide failover and load balancing, the host must have at least two HBAs.
- Only paths from identical HBAs can be mapped to the same logical device. The HBAs must be comparable in every way; they cannot even be different revisions of the same HBA.
- Logical devices must be properly configured on the SCSI bus.
- A SCSI HBA must connect to one and only one interface port. That is, no path to a single logical device can share either an HBA or an interface port.
- EMC requires single-initiator connections from the host to the storage system.

Single Fabric Configuration



In a single-fabric configuration, a host with one HBA can have multiple paths to a logical device. A single-HBA host has four paths to each logical device. While this does not provide maximum availability, it offers back-end path redundancy and load balancing, reduces zoning complexity, and enables multiple hosts to share a storage-system port.

Multiple Fabric Configuration



PowerPath supports multiple paths between an HBA and a logical device. This can offer higher availability and/or better performance, and it usually simplifies zoning.

Multiple connections between hosts and multiple fabrics can insulate I/Os from fabric-wide failure. PowerPath delivers maximum system-wide availability when dual HBAs in each server connect to separate fabrics. An application that cannot use a path on one fabric can fail over to a different fabric, protecting the application from a fabric-wide outage.

PowerPath supports load balancing and redundancy across the storage ports in a fabric. Be aware, however, that the number and complexity of connection points in a fabric multiply rapidly in a multipath configuration. The number of connections you create depends on the bandwidth required.


This Multiple Fabric configuration has several advantages:

- If either fabric fails, both hosts can still access both logical devices, through the other fabric.
- If either interface fails, both hosts can still access both logical devices, through the other interface.
- If either HBA in a host fails, that host can still access both logical devices, through the other HBA on that host.

If multiple hosts access the same storage-system ports, PowerPath does not add any access control. Normal sharing considerations apply, using products such as SAN Manager, Volume Logix, Access Logix, Oracle RAC, and clustering software.

If multiple HBAs are mapped to the same LUN. They must be identical both hardware and software revisions.

Path States and Modes



Paths can be in either of two modes, Active or Standby. PowerPath is designed to use all active paths in a normal condition. There are two states for a particular path, Alive or Dead. PowerPath manages the state of each path to each logical device independently. It also issues a sequence of I/Os (path test), specifically to ascertain the viability of a path. If a path test fails, PowerPath marks the path “Dead” and stops sending I/O to it.

After a path fails, PowerPath continues testing it periodically, to determine if it is fixed. If the path passes a test, PowerPath restores it to service and resumes sending I/O to it. The storage system, host, and application remain available while the path is restored. The time it takes to do a path test varies. Testing a working path takes milliseconds. Testing a failed path can take several seconds depending on the type of failure.

PowerPath restores dead paths automatically. However, this task is a low priority task which can be preempted in a very active environment. PowerPath provides a CLI interface to immediately restore the paths.

Configuration Changes Affecting PowerPath

- Zoning or LUN masking changes
- Adding or removing Fibre Channel switches
- Adding or removing HBAs or storage-system ports
- Adding or removing logical devices

These changes require a PowerPath reconfiguration. Many platform-specific and some PowerPath reconfiguration operations will fail if the affected logical devices are in use (either just marked alive or with I/O in progress). After any reconfiguration, you must monitor the outcome of individual reconfiguration steps and confirm that the resulting changes are as expected, before relying on the new configuration. Otherwise, some paths may not be as expected. For example, on platforms that support native devices, it is possible to inadvertently write to the wrong logical device (thinking that the native device is associated with a different logical device).

PowerPath Management

Section Three



EMC Global Education

© 2004 EMC Corporation. All rights reserved. These materials may not be copied without EMC's written consent.

Mercer Road
Global Education & Productivity

38

In this last section we will take a look at PowerPath management.

Managing PowerPath License

```
emcpreg [-list|-install]  
        [-add|-remove|-check] key
```

- Must be done during or after the installation of PowerPath
- One simple command to manage the license
- License key modification does not require a host reboot
- Alternatively, use the following command to quickly validate the license
 - `powermt check_registration`

Only a user with root or administrator privileges can run `emcpreg` command. It manages PowerPath license registration. On some platforms, `emcpreg` is invoked automatically by the PowerPath installation processes. On other platforms, you must invoke the command manually. Once installation is complete, `emcpreg` can be invoked manually on any supported platform.

Configuring PowerPath

```
powermt config
```

- Discover physical devices available to the host
 - The OS must first see the devices before this step works properly
- Does not remove devices or paths no longer available to the host
 - These non-existent paths will be presented as “Dead”
- Non-intrusive command that can be repeatedly executed

`Powermt config` configures all detected logical devices on Symmetrix, CLARiiON, and third party storage as PowerPath devices and adds these devices to the PowerPath configuration. This command is automatically executed as a part of the boot process on the supported platforms. You need to run `powermt config` only if more paths are added after boot time.

The execution of `powermt config` varies slightly depending on the different platform.

On AIX, if `powermt config` encounters an error, a general AIX error message is displayed, indicating there is a problem configuring a device. The text of the message depends on the underlying error. For more detailed PowerPath error information, we recommend AIX users run `cfgmgr -vl powerpath0` instead of `powermt config`.

The `cfgmgr` command performs the same function as `powermt config` but generates more detailed error messages. Also on AIX, `powermt config` hangs if either the `/` or the `/tmp` file system is full.

On Linux and HP-UX, `powermt config` runs at boot time from `/etc/inittab` to configure PowerPath. On Solaris, `powermt config` runs at boot time from startup scripts under `/etc/rcS.d`. On AIX, PowerPath is configured at boot time using ODM (Object Data Manager) configuration rules. Also on Linux, before running `powermt config`, make sure the required special native SCSI disk block device files exist in the `/dev` directory.

Monitoring PowerPath

```
powermt display [paths|ports]
                [class={symm|CLARiiON|hitachi|hpxp|ess|all}]
                [every=seconds] [dev=device|all]
```

- The most often used command
- The *every* option will display path in a loop
- The *paths* option displays all paths between the HBAs on a host and the storage-system ports that the HBAs can access
- The *ports* displays the state of the storage-system ports that are visible to the HBAs under PowerPath control
- All supported storage classes are displayed by default

PowerPath gives users a great flexibility to monitor the environment. As a result, the `powermt display` command can display devices in many different ways, allowing users to customize the command to suit their particular need. The default display mode is to list all devices visible to the host.

Zoning Effect on “powermt display”

- Logical device count does not match the number of total paths

```
powermt display
Symmetrix logical device count=32
=====
- Host Bus Adapters - ---- I/O Paths --- ----- Stats -----
### HW Path      Summary Total Dead  IO/sec Q-IOs Errors
=====
2303 port3\path0  optimal 128  0    -    0    0
2304 port4\path0  optimal 128  0    -    0    0
```

In a switched topology, HBAs may be zoned to see multiple storage-system ports. `powermt display` cannot display virtual HBAs (multiple instances of an HBA, one per port). As a result, the sum of the values in the I/O Paths Total column may exceed the logical device count value. In addition, in the Host Bus Adapters column, only one instance of an HBA is shown for all paths in the I/O Paths Total column. These discrepancies are due to Fibre Channel switch zoning of HBAs to multiple storage-system ports.

The I/O Paths Total column shows the number of paths associated with each HBA. Although there are multiple storage-system ports visible to each HBA, the output displays one instance of each HBA and its hardware path (in the Host Bus Adapters column). In the figure, HBAs 2303 and 2304 are fibre connections that can access identical logical devices; both display a value of 128 in the I/O Paths Total column. In reality, there are 32 logical devices. Due to zoning, however, these 32 logical devices are shared among several ports (in this case, 4 ports); therefore, the number of logical devices in the output is 128. This is an example of 1:4 zoning; each HBA has access to 4 Fibre Channel interfaces.

Restoring Dead Paths

- PowerPath automatically restores dead paths by default
- Done periodically
- There may be a delay in accessing a recovered path. To avoid this delay, run `powermt restore` after a path is physically restored
- A restored path could temporarily stay “dead” until PowerPath sends the next I/O down that path
- Non-intrusive command; No reboot required

`Powermt restore` tests and restores specified paths. It issues test I/Os and responds to the test results as follows:

- If a live path passes the test, `powermt restore` does nothing.
- If a dead path passes the test, `powermt restore` marks it alive; PowerPath can now use it. In addition, other dead paths on the same HBA and/or storage system port may be tested.
- If a live path fails the test, `powermt restore` marks it dead and prints a warning (every time the path fails the test). In addition, other paths that share the same HBA and port may be marked dead, and other paths that share only the HBA or only the port may be tested.
- If a dead path fails the test, `powermt restore` prints a warning (every time the path fails the test).

Load balancing and Failover Policy (1 of 2)

Policy	Storage Systems	Load Balancing	Failover
ad (Adaptive)	Hitachi, HP, IBM	✓	✓
bf (Basic failover)	CLARiiON(SP failover)		✓
co (CLARiiON Optimization)	CLARiiON	✓	✓
lb (Least Block)	All	✓	✓
li(least I/O)	All	✓	✓

Adaptive - I/O requests are assigned to paths based on an algorithm that takes into account path load and logical device priority. This policy is valid only for Hitachi Lightning, HP xp, and IBM ESS storage systems and is the default policy for them on platforms with a valid PowerPath license.

Basic Failover - Load balancing is not in effect. I/O routing on failure is limited to one HBA and one port on each storage processor. This policy protects against SP and back-end failures and allows non-disruptive upgrades to work when running PowerPath without a license key. It does not protect against HBA or host loop failures. This policy is valid only for CLARiiON storage systems and is the default policy for them, on platforms without a valid PowerPath license.

CLARiiON Optimization - I/O requests are assigned to paths based on an algorithm that takes into account path load and logical device priority. This policy is valid only for CLARiiON storage systems and is the default policy for them, on platforms with a valid PowerPath license.

Least Block - Load balance is based on the number of blocks in pending I/Os. I/O requests are routed to the path with the fewest queued blocks, regardless of the number of requests involved.

Least I/O - Load balance is based on the number of pending I/Os. I/O requests are routed to the path with the fewest queued requests, regardless of total block volume.

Load balancing and Failover Policy (2 of 2)

Policy	Storage Systems	Load Balancing	Failover
nr(NoRedirect)	Symmetrix, Hitachi, HP, IBM		
re(Request)	All		✓
rr (RoundRobin)	All	✓	✓
so (Symmetrix Optimization)	Symmetrix	✓	✓

No Redirect - Neither load balancing nor path failover is in effect. If nr is set on a pseudo device and the I/O path fails, data errors can occur. If nr is set on a native device and the I/O path fails, data errors will occur. This policy is valid only for Symmetrix, Hitachi Lightning, HP XP, and IBM ESS storage systems and is the default policy for them on platforms with a valid PowerPath Base license and on platforms without a valid PowerPath license.

Request - For native devices, this policy uses the path that would have been used if PowerPath were not installed. For pseudo devices, it uses one arbitrary path for all I/O. For all devices, path failover is in effect, but load balancing is not. This is the default policy for CLARiiON storage systems on platforms with a valid PowerPath Base license.

Round Robin - I/O requests are assigned to each available path in rotation.

Symmetrix Optimization - I/O requests are routed to paths based on an algorithm that takes into account path load and logical device priority. Load is a function of the number, size, priority, and type of I/O queued on each path. This policy is valid only for Symmetrix storage systems and is the default policy for them, on platforms with a valid PowerPath license.

Setting Load Balancing Policy

```
powermt set policy=li|lb|rr|re
           [class={symm|CLARiiON|hitachi|hpxp|ess|all}]
           [dev=device|all]
powermt set policy=so|nr [dev=device|all]
powermt set policy=bf|co [dev=device|all]
powermt set policy=ad|nr [class=hitachi|hpxp|ess] [dev=device|all]
```

- Policy is set per storage device
- By default, set to “so” for Symmetrix and “co” for CLARiiON
- Non-intrusive command
- Changes take place dynamically

`Powermt set policy` sets the load balancing and failover policy for devices configured in PowerPath. One storage devices on a host could have a different load balancing policy than another device. For instance, a host could be connected to both the Symmetrix and CLARiiONs. Then, users can set the policy for Symmetrix devices to “so” and to “co” for CLARiiON devices. This command is non-intrusive and can be executed at anytime. Changes will occur dynamically without requiring a reboot.

Setting Device Priority

```
powermt set priority=priority#  
[class={symm|CLARiiON|hitachi|hp|xpress|all}]  
[dev=device|all]
```

- Priority is set per storage device
- Boost I/O for device with higher priority number
- Ranges from 0 – 10; 10 has highest priority
- Load balancing must be “so” or “co”

Priority allows you to tune the bandwidth for specific logical devices, to improve application performance. To prevent one application from taking bandwidth away from another, set a higher priority for one application's logical devices than the other application's logical devices. All the system's bandwidth is available when the high-load application peaks.

`Powermt set priority` sets the I/O priority for devices. I/O to lower-priority logical devices tends to be delivered to paths not used by higher-priority logical devices. This command is relevant only for those devices whose load balancing and failover policy is set to Symmetrix optimization, CLARiiON optimization, or Adaptive. Although it sets device priorities with other policies, it has no effect on load balancing on devices. If the policy for a given device is anything other than Symmetrix optimization, CLARiiON optimization, or Adaptive, and a priority for that device is set, the priority setting has no effect; however, if the policy for that device is later changed to Symmetrix optimization, CLARiiON optimization, or Adaptive, the previously set priority becomes relevant when the policy change occurs.

Saving PowerPath Custom Settings

```
powermt save [file=pathname]
```

- The settings are saved in powermt.custom file
 - On Windows: %SystemRoot%\system32\drivers\etc\
 - On Unix: /etc
- Should be done before and after customizing load balancing policy and device priority
- Use meaningful name for the setting

By default, `powermt save` saves the PowerPath settings in the file `powermt.custom`. You can specify an alternate file. For each configured logical device, `powermt save` records its serial numbers, policy, priority, write throttle setting, write throttle queue depth, and pseudo device name (for platforms that support pseudo devices). For each configured path (alive or dead), `powermt save` records its mode. For each storage-system class, `powermt save` records its periodic autorestore setting.

For performance tuning purposes, you can use `powermt save` to save different configurations under different filenames. If a configuration is saved with `powermt save` under the default filename, the configuration loads automatically at boot time. (The exception is AIX, which always loads default settings at boot time.) If the default filename is not used, the configuration does not load automatically at boot time; in this case, you must load the configuration manually with `powermt load`.

Loading Custom Settings

```
powermt load  
[class={symm|CLARiiON|hitachi|hpxp|ess|all}]  
[file=pathname]
```

- May need to run `powermt config` after this command on some platforms
- Custom files with name other than `powermt.custom` must be loaded manually after the OS boot process

`Powermt load` the following settings from a saved configuration file applies to the current configuration : PowerPath device policies and priorities; write throttle setting and write throttle queue depth; periodic autorestore setting; and path modes. The file is saved when `powermt save` is run manually and, on some platforms, at boot time. (On Windows, PowerPath Administrator prompts to save the file after configuration changes are made in the GUI.)

Paths added since the last save remain as is. If the configuration file has information about dead paths, they are added in a failed state, allowing them to be restored later. If the file has information about paths that were removed since the last save, they are added.

Removing a Path from PowerPath Control

```
powermt remove  
[class={symm|CLARiiON|hitachi|hpxp|ess|all}] [force]  
<hba=[hba#|all] | dev=[path|device|all] >
```

- Do not remove the last path to a logical device unless you plan to remove the logical device entirely, as data access will be interrupted
- HBA can be removed only after all paths on it have been removed
- Does not physically remove the storage device

`Powermt remove` deletes the specified path (or paths) from PowerPath's list of configured paths. It does not delete the logical device that the paths refer to. As long as the logical device remains visible on the system, the logical device can be reconfigured via `powermt config`.

If a PowerPath device is open, `powermt remove` does not remove that device or the last path to that device. If a device is in a PPVM volume group, PowerPath does not remove the last path to that device.

Course Summary

Key points covered in this course:

- PowerPath Technical Overview
 - Architecture
 - Load balancing policies
 - Behavior differences in Active-Active and Active-Passive storage systems
 - Key features of PowerPath Volume Manager
 - Key features of PowerPath Volume Mobility
- PowerPath Configuration
 - Connectivity and configuration in both SCSI and Fibre Channel environments
 - The most configuration changes that affect PowerPath
- PowerPath Management
 - Managing PowerPath licenses
 - Configuring PowerPath
 - Monitoring the device under PowerPath control
 - Restoring dead path
 - Customizing load balancing and device priority
 - Securing and reusing custom PowerPath settings
 - Removing a path from PowerPath control

Key Points covered in this course are shown here. Please take a moment to review them.

where information lives

EMC²
where information lives

EMC Global Education
© 2004 EMC Corporation. All rights reserved.

 **Mercer Road**
Unite Education & Productivity

52

Thank you for your attendance. This ends our training session on PowerPath Configuration and Administration.