#SHAREorg

SHARE
Technology · Connections · Results

The New Path to EMC SCSI Devices

Gail Riley
EMC Corporation

March 10, 2014
Session14946

Test link: www.SHARE.org

SHARE in Anaheim

Agenda

- Introduction of PowerPath
- z Architectures Supported
- Types of Storage Devices
- Types of Configurations
- Inclusions and Limitations

What is EMC PowerPath?

Automate, optimize, standardize path management

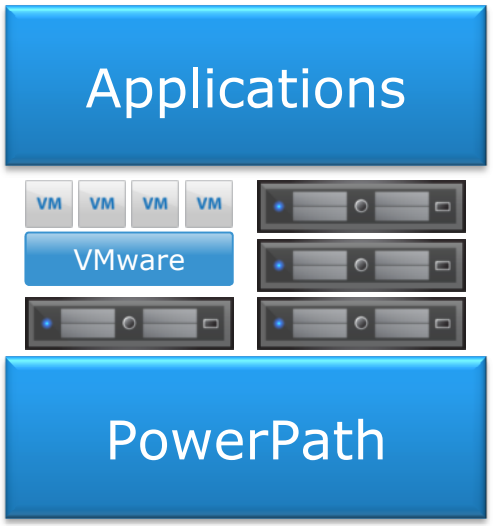

All data paths are active, optimized for load balancing

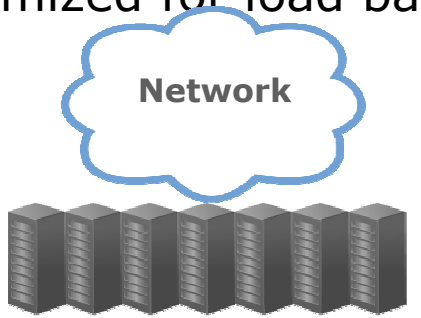

- Automate failover/ recovery for high availability
- Optimize load balancing for performance
- Standardize path management across physical and virtual
- **Now also available for Linux on System z!**

PowerPath History

Automate path failover and recovery, optimize load balancing

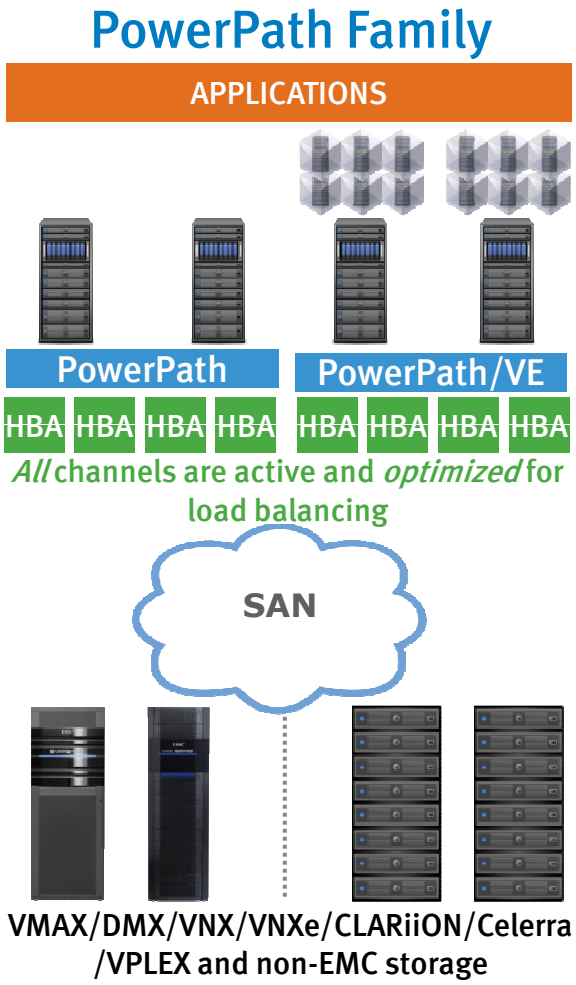

- Product maturity
  - Over 12 years' experience as leading path-management solution
    - Running on over 1,000,000 hosts
    - Over 30,000 installation sites
- Performance
  - Optimization of virtual and physical data centers
- Reliability
  - Time-tested, proven reliable solution across your enterprise
- Functionality
  - Uncompromised quality with advanced functionality
  - Failover and load balance IO
  - PPME – Migration Enabler (Non-Disruptive)
  - PPE – Data at rest Encryption
  - PP Viewer – Centralized Monitoring
  - Recent New Feature's
    - PMI – Path Management Insight
    - Flaky Path Detection

Automated Failover and Recovery

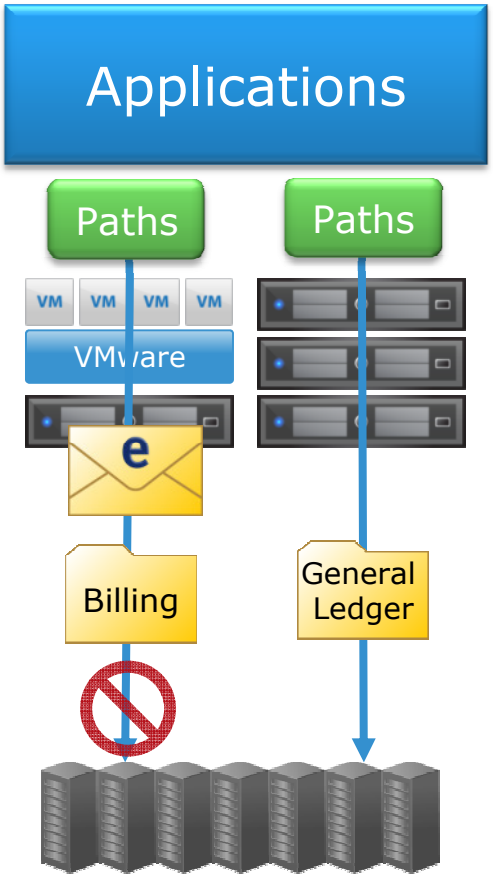

- PowerPath Patented algorithms detect inoperative paths automatically
- Failover I/O requests to *optimal* active paths
- Combine with sophisticated load balancing to ensure application performance

Path and Bus Testing

Automate path and bus diagnostics “behind the scenes”

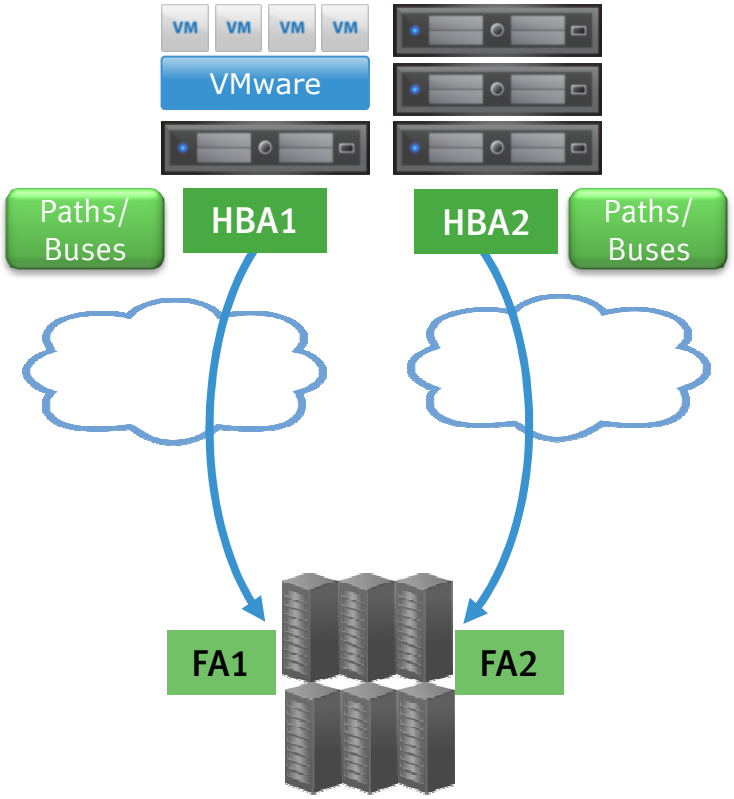

- Determine the operational status and viability of a data path
- Insulate applications from path testing
  - Adaptive to load on the host to not impact application performance
  - Path tests are scheduled with a lower priority than application I/O
  - Idle I/O paths are tested
  - Active I/O paths are not tested
- Activate and de-activate paths based on test success or failure

Optimized for EMC Platforms

Increase server, storage, and data path efficiency

- Assures tunable, predictable performance
- Optimize server, storage, and data path utilization
- Maximize data availability
- Unify management across physical and virtual environments

Intelligent Load Balancing

Physical environments

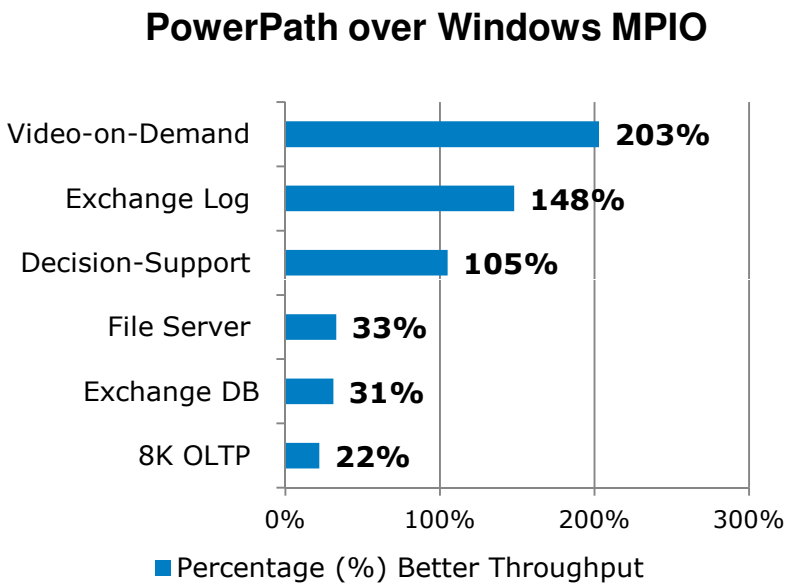

- Re-balance workloads automatically to optimize performance
- Monitor and diagnosis SAN traffic for proactive correction actions
  - Measure I/O response times against user-defined thresholds
  - Trigger SNMP events for alerts

Source: ESG Lab: EMC PowerPath vs. Windows Native MPIO, October 2011

PowerPath Exclusive Path Management Insight

- I/O metrics isolate issues to SAN or array
- LUN-based read/write metrics
  - Throughput (bytes/sec), average response time (milliseconds)
  - LUN I/Os per second in four buckets (unique to PowerPath)
- Path diagnostics metrics
  - Latency high and low water marks (milliseconds)
  - Latency distribution in four buckets (unique to PowerPath)
  - Numerous sample measuring to quantify path health
- Autostandby
  - Trigger to move a path to standby automatically

PowerPath 5.7 for Linux on System z - Architectures Supported

- Linux on System z as a guest virtual machine
- Linux on System z running native in a LPAR

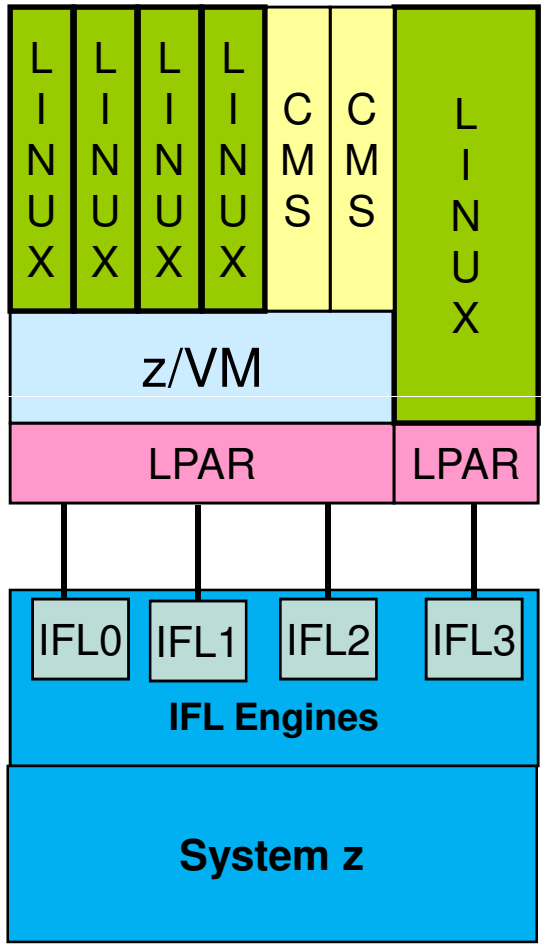

PowerPath 5.7 for Linux on System z – Storage Devices

- FBA (Fixed Block Architecture)
  - FCP (Fibre Channel Protocol)
  - SCSI devices
  - Multipathing is done by Linux on System z
  - FBA as z/VM edevice's are not supported
- CKD (Count Key Data)
  - CKD is not supported by PowerPath

PowerPath for Linux on System z – Supported Storage Configuration

- Boot device is a ECKD device and FBA SCSI LUNs
  - System z handles multipathing for ECKD devices
  - PowerPath will do multipathing for the SCSI LUNs
  - Considered as a local boot setup by PowerPath
- FBA LUNs only
  - Both boot device and data devices are FBA SCSI LUNs
  - Multipathing for all the FBA SCSI LUNs are done by PowerPath
  - Considered as a Boot From SAN setup by PowerPath

PowerPath for Linux on System z- Operating System Support

- Linux Operating Systems Supported
  - RHEL 5.8 (s390x)
  - RHEL 6.3 (s390x)
  - SLES 10 SP4 (s390x)
  - SLES 11 SP2 (s390x)

PowerPath 5.7 for Linux on System z - Array Support*

- VMAX Hardware
  - Only VMAX is supported
  - VMAX 20k, 40K
- VMAX Enginuity Software
  - Enginuity 5876

*Always check the latest support matrix before implementation

PowerPath 5.7 for Linux on System z - FileSystem Support

- Supported Filesystems
  - ext2 and ext3 are supported on all supported operating systems
  - ext4 on RHEL6.3 & SLES11 SP2
  - ReiserFS on SLES10 SP4 and SLES11 SP2
  - XFS on SLES11 SP2

PowerPath 5.7 for Linux on System z - Cluster and Volume Manager Support

- Volume Managers
  - Linux LVM (All supported OS'es)
- Cluster Support
  - None

PowerPath 5.7 for Linux on System z – Inclusions

- EMC PowerPath Viewer
  - Centralized monitoring utility providing a consolidated display of events
  - Allows you to view and monitor up to 1000 PowerPath hosts through a graphical user interface (GUI)
  - Ability to view hosts, host groups, LUNs, individual paths to each LUN, and buses

- Advanced Path Metrics (also known as PMI)
  - Supported (same functionality as PP 5.7 SP1 for Linux)
  - Includes latency and throughput information, autostandby (flaky path detection only)

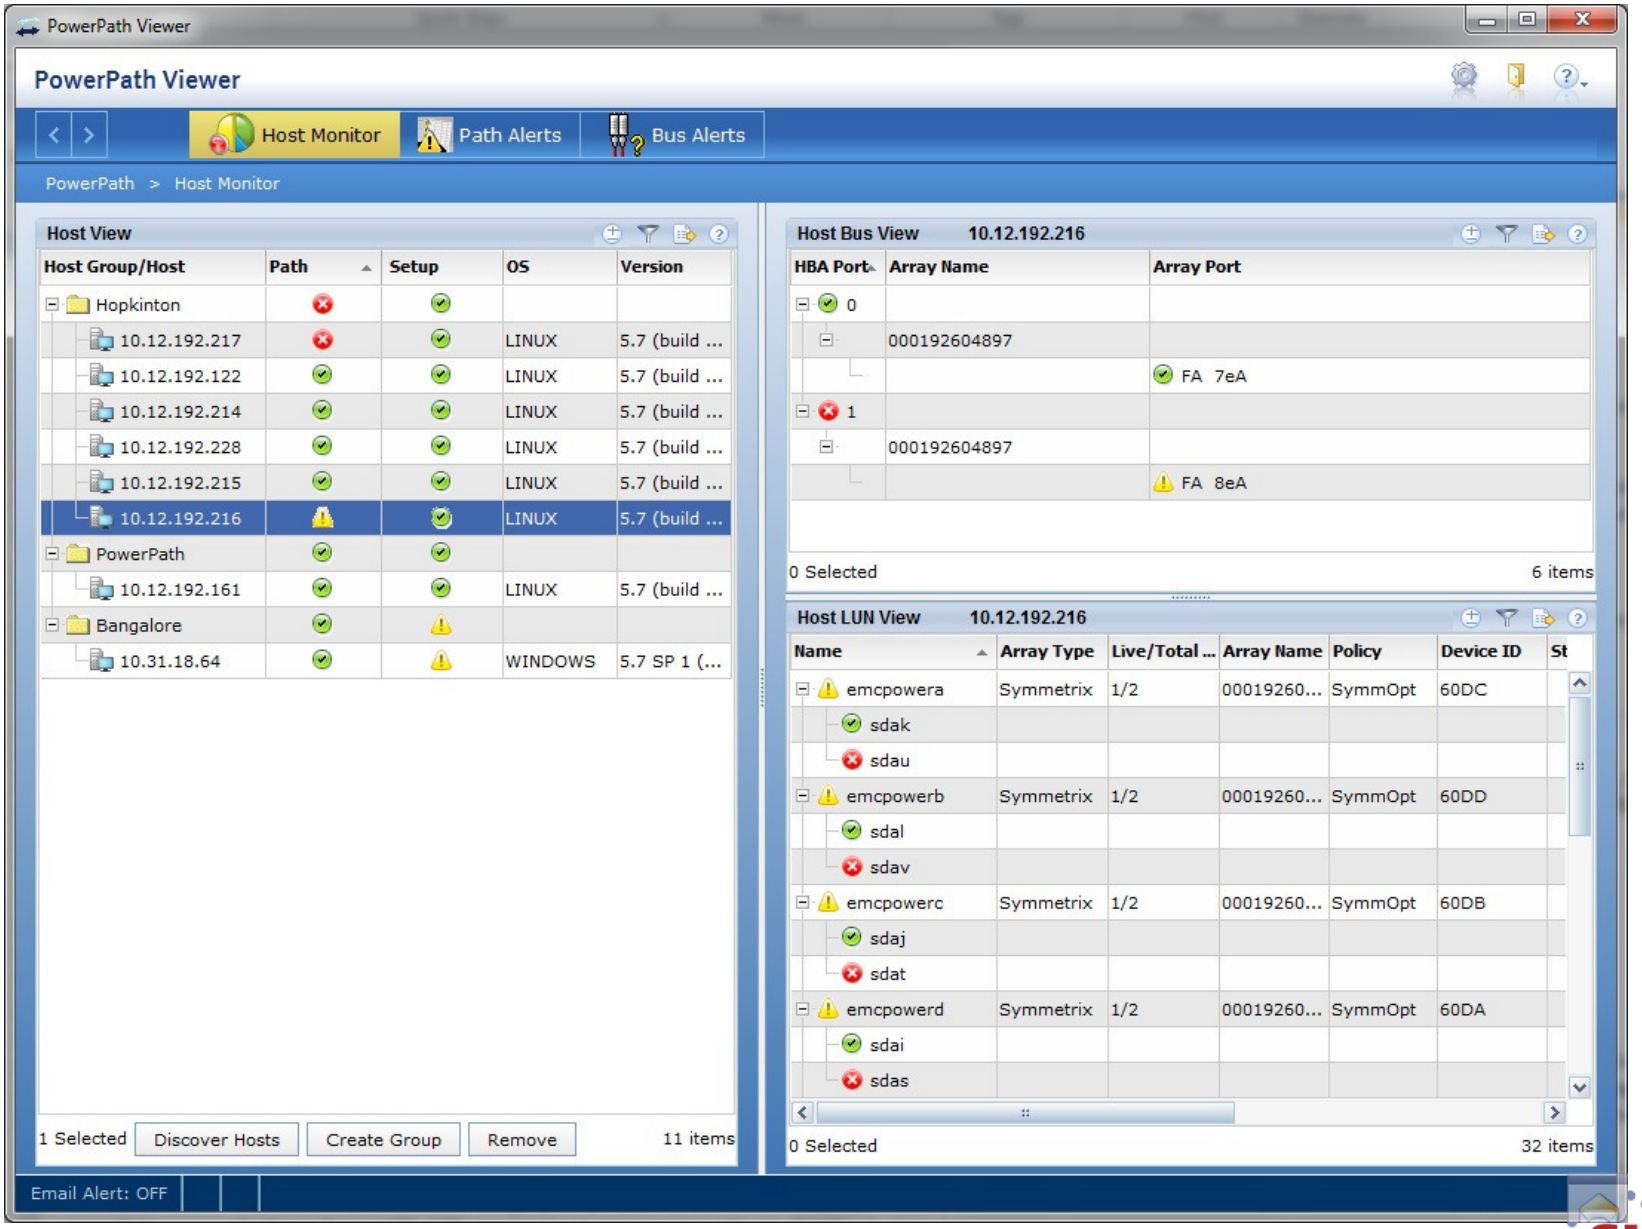
PowerPath Viewer
PowerPath Viewer
Host Monitor
Path Alerts
Bus Alerts
PowerPath > Host Monitor
Host View
Host Group/Host
Path
Setup
OS
Version
Hopkinton
10.12.192.217
10.12.192.122
10.12.192.214
10.12.192.228
10.12.192.215
10.12.192.216
PowerPath
10.12.192.161
Bangalore
10.31.18.64
LINUX
WINDOWS
5.7 (build ...
5.7 SP 1 (...
1 Selected
Discover Hosts
Create Group
Remove
11 items
Host Bus View
10.12.192.216
HBA Port
Array Name
Array Port
0
000192604897
FA 7eA
1
FA 8eA
0 Selected
6 items
Host LUN View
10.12.192.216
Name
Array Type
Live/Total ...
Array Name
Policy
Device ID
emcpowera
emcpowerb
emcpowerc
emcpowerd
Symmetrix
1/2
00019260...
SymmOpt
60DC
60DD
60DB
60DA
sdak
sdau
sdal
sdav
sdaj
sdat
sdai
sdas
0 Selected
32 items
Email Alert: OFF

PowerPath 5.7 for Linux on System z – Limitations

- PowerPath Migration Enabler (PPME)
  - Not supported (any flavor)
- Encryption
  - Not supported
- VNX, CLARiiON, and VPLEX arrays
  - Not supported

PowerPath Commands

- powermt - PowerPath management utility
- man pages available

powermt <command>
[class={all|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|generic}]

powermt check [hba=<hba#>|all] [dev=<path>|<device>|all] [class=<class>|all]
[force]
powermt check_registration
powermt config
powermt disable hba=<hba#>
powermt display [dev=<device>|all] [class=<class>|all] [every=<#seconds>]
[wide|width=<#col>]
powermt display hba_mode [class=<class>|all]
powermt display latency [dev=<device>|all] [class=<class>|all]
[every=<#seconds>] [wide|width=<#col>]
powermt display nonvirtual dev=<device>|all
……

powermt display dev

powermt display [dev=<device>|all] [every=<#seconds>]

```
Pseudo name=emcpowera
Symmetrix ID=000195701370
Logical device ID=0871
state=alive; policy=SymmOpt; queued-IOs=0
==============================================================================
   =
--------------- Host ---------------   - Stor -   -- I/O Path --    -- Stats --
   -
###  HW Path                I/O Paths    Interf.   Mode      State    Q-IOs
   Errors
==============================================================================
   =
   0 zfcp                    sda          FA 12eA   active    alive       0       0
   1 zfcp                    sdb          FA  5eA   active    alive       0       0
   3 zfcp                    sdw          FA 12eA   active    alive       0       0
   2 zfcp                    sdact        FA  5eA   active    alive       0       0
```

powermt display bus (FCP ports)

```
# powermt display bus
Symmetrix logical device count=205
============================================================================
----- Host Bus Adapters --------- ------ Storage System -----    - I/O Paths
###  HW Path                            ID           Interface     Total    Dead
============================================================================
  0 zfcp                          000195701370     FA 12eA          205       0
  1 zfcp                          000195701370     FA  5eA          204       0
  2 zfcp                          000195701370     FA  5eA          204       0
  3 zfcp                          000195701370     FA 12eA          204       0
```

powermt display hba_mode

```
 # powermt display hba_mode
Symmetrix logical device count=205
============================================================================
----- Host Bus Adapters  ------ I/O Paths -----  Stats
###  HW Path             Summary   Total   Dead  Q-IOs Mode
============================================================================
   0 zfcp                optimal     205      0      0  Enabled
   1 zfcp                optimal     204      0      0  Enabled
   2 zfcp                optimal     204      0      0  Enabled
   3 zfcp                failed      204    204      0 Disabled
```

- Once the path is healthy again it will be brought online automatically

powermt display dev

```
# powermt display dev=emcpowerdj
Pseudo name=emcpowerdj
Symmetrix ID=000195701370
Logical device ID=10DF
state=alive; policy=SymmOpt; queued-IOs=4
==============================================================================
----------Host ----------   - Stor -   -- I/O Path --   -- Stats ---
###  HW Path  I/O Paths    Interf.  Mode      State   Q-IOs Errors
==============================================================================
   2 zfcp        sdvt         FA  5eA  active    alive       2       0
   1 zfcp        sdwi         FA  5eA  active    alive       1       0
   0 zfcp        sdxu         FA 12eA  active    alive       1       0
   3 zfcp        sdaan        FA 12eA  active    dead        0       1
```

Display Symmetrix Ports

powermt display ports

```
Storage class = Symmetrix
================================================================================
---------- Storage System --------------     -- I/O Paths --    --- Stats ---
    ID               Interface          Wt_Q      Total     Dead    Q-IOs  Errors
================================================================================
    000195701370     FA  5eA            256        408        0        0        0
    000195701370     FA 12eA            256        409        0        0        0
```

Maximum number of write I/O requests that will be serviced before the I/O queue checks for any outstanding read I/O.

powermt config

- powermt config – configure paths to logical devices

1. Make physical path additions and zoning as required
2. Make new paths available to the host by using array and SAN management tools
3. Configure the newly added devices on Linux
4. Execute:

   powermt config

   powermt display

PowerPath Performance Monitoring

- powermt set path_latency_monitor={on|off}
- powermt set path_latency_threshold=<#seconds>|<#milliseconds>ms

- powermt set perfmon={on [interval=<#seconds>] | off}
- powermt display perf dev=<device>|all [continuous] [verbose] [xml]
- powermt display perf bus [continuous] [verbose] [xml]

powermt display perf bus verbose

```
ln192161:~ # powermt display perf bus verbose
Timestamp = 15:55:04 UTC, 07 Mar 2014
Sample Interval = 60
Symmetrix logical device count=205
=============================================================================
----- Host Bus Adapters ------- ------ Storage System -----     Retry   Error
###  HW Path                        ID           Interface       delta   delta
=============================================================================
 0 zfcp                          000195701370  FA 12eA              0       0
                                      Latency----Reads--Writes
                                       Low (ms)   0.177        -
                                      High (ms)     5.9        -
                                Lat <=   1ms (#)   34879        0
                          1ms < Lat <=  10ms (#)    4038        0
                         10ms < Lat <= 100ms (#)       0        0
                                Lat >  100ms (#)       0        0

.....
```

Summary

- History of PowerPath
- Linux and filesystem support
  - ext2, ext3, ext4, reiserfs, xfs
- FBA as SCSI device support
- Boot from CKD or boot from SAN both supported
- powermt commands

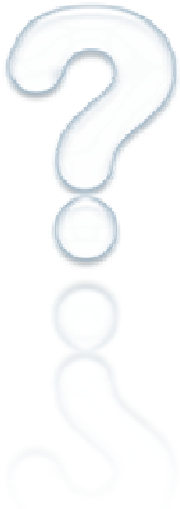

Q&A

THANK YOU