

Lab Validation Report

EMC PowerPath vs. Windows Native MPIO

Automated Data Path Management, Failover and Recovery, and
Optimized Load Balancing

By Vinny Choinski

October 2011

Contents

Introduction	3
Background.....	3
EMC PowerPath Overview.....	4
ESG Lab Validation	5
Test Bed	5
Getting Started	5
EMC PowerPath for Windows vs. Windows Native MPIO	7
Performance	8
Advanced Monitoring and Management	11
ESG Lab Validation Highlights	13
Issues to Consider	13
The Bigger Truth	14
Appendix	15

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by EMC.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

Introduction

EMC PowerPath improves data center performance with automated data path management, failover and recovery, and optimized load balancing. While many operating systems include free, native multipath I/O (MPIO) capabilities, customers often don't know the difference between these offerings and PowerPath. This ESG Lab Validation report documents the results of hands-on testing conducted with the goal of comparing PowerPath's optimized load-balancing and other capabilities with the native MPIO included with Microsoft Windows.

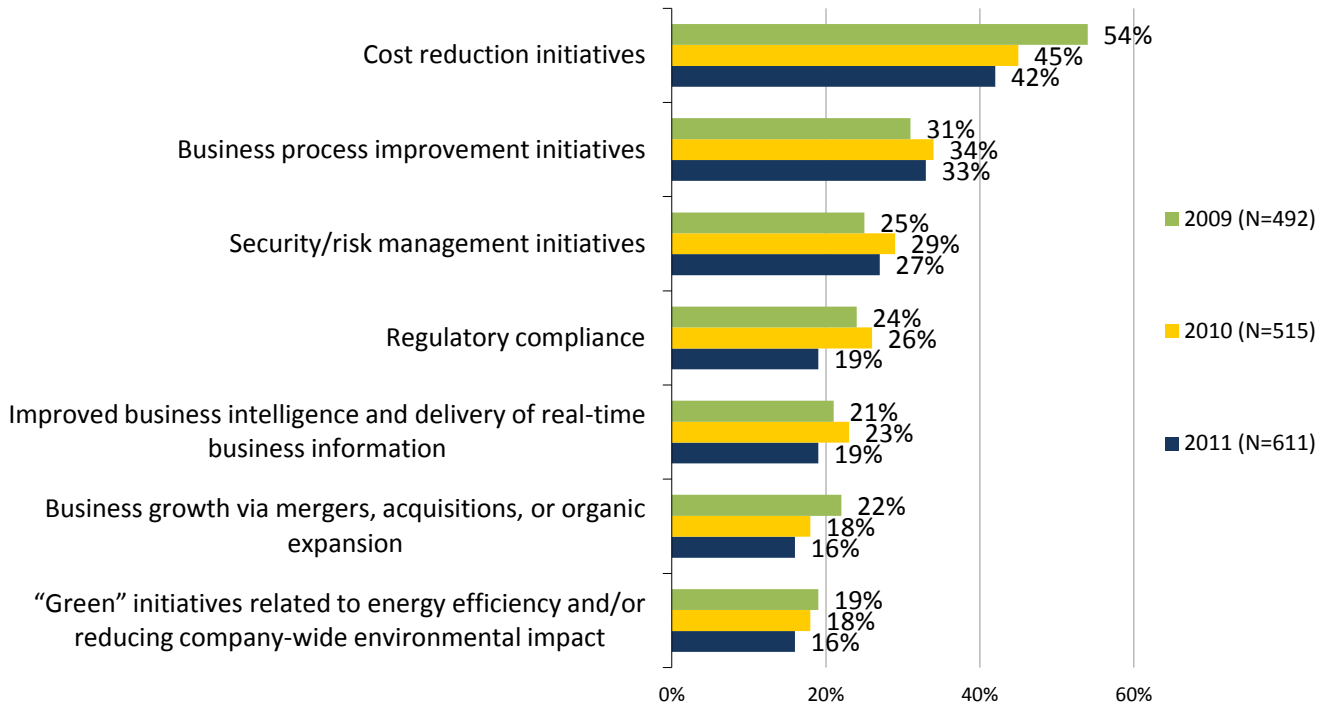
These tests were designed to simulate the real-world SAN conditions that organizations encounter every day, and were conducted using common tools to highlight the potential impact on application, SAN, and storage performance. While these test results represent what a user might experience, actual results will vary depending on the data center configuration.

Background

Organizations depend on consistent application performance and data availability to run business processes effectively, and ESG research indicates their willingness to invest in solutions that support these objectives. According to surveys of IT professionals in enterprise and midmarket companies, cost reduction and business process improvement initiatives have remained the top spending priorities for the past three years¹ (see Figure 1). Better performance and availability contribute to capital and operational cost reduction by increasing hardware optimization, minimizing IT tasks, and ensuring business uptime. Additionally, better performance and availability also deliver business process improvements by keeping mission-critical applications online and optimal to enable greater productivity.

Figure 1. Business Initiatives That Will Impact IT Spending Decisions, Three-year Trend

Which of the following business initiatives do you believe will have the greatest impact on your organization's IT spending decisions over the next 12-18 months? (Percent of respondents, three responses accepted)



Source: Enterprise Strategy Group, 2011.

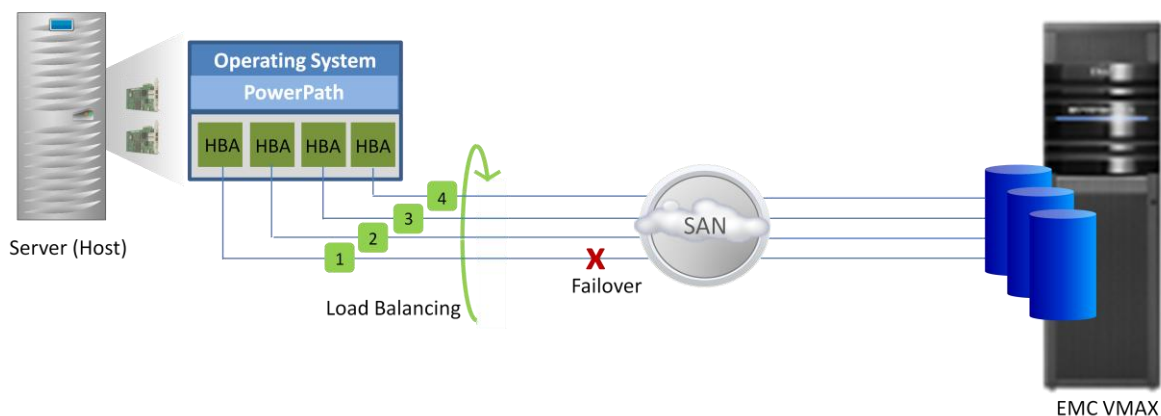
¹ Source: ESG Research Report, [2011 IT Spending Intentions Survey](#), January 2011.

EMC PowerPath Overview

EMC PowerPath optimizes path utilization in Fibre Channel, iSCSI, and Fibre Channel over Ethernet SANs to produce predictable, scalable, and consistent information access. It supports heterogeneous servers, operating systems, and storage (including non-EMC arrays) in virtual and physical environments. PowerPath provides efficient utilization of all data channels and eliminates the need for multiple, separate multipathing solutions for each operating system deployed. The multipathing license supports up to 32 paths per LUN from multiple HBAs to multiple storage ports.

Figure 2 illustrates PowerPath's functionality at a high level. The software resides on the host, sitting above the HBA, iSCSI TOE, or CNA for heterogeneous operating system and storage array support but below the application, database, and file system. As a result, it can work with raw devices, volume managers, or file systems and can perform full path management, including optimized load-balancing and automated data path failover and recovery. Each HBA provides a path through the SAN to the storage array; PowerPath balances I/O across the data paths. Should a path fail, PowerPath redirects I/O to the next best available path which re-balances the load.

Figure 2. EMC PowerPath Load-balancing and Failover



While PowerPath does support multiple load-balancing policies, its default policy is to select the *optimal* data path in normal and faulted conditions. It automates I/O load balancing on a host-by-host (server-by-server) basis and maintains statistics for all paths. For each I/O request, PowerPath selects the best of all available paths. It is not designed to evenly distribute I/O loads among paths because a path management solution based on even distribution is not necessarily effective or efficient. Instead, it optimizes the paths based on statistics and algorithms. If one data path is available but clogged, PowerPath moves I/O to paths that will maximize performance. Paths are selected based on pending I/O on that path, size of I/O, type of I/O, paths most recently used, and array-specific coding. PowerPath may also use the same data path to avoid continuous unnecessary switching, provided that path remains the best choice. PowerPath provides:

- **Proactive failure management.** Using sophisticated, mostly patented algorithms, PowerPath boosts application I/O rates through intelligent, self-sensing infrastructure management with ongoing testing of both active and inactive data paths. PowerPath can detect decaying paths that may threaten application performance or availability and make adjustments in the background. Inactive paths are placed back into the queue once they return to an active state, and decaying paths are utilized based on performance levels.
- **Automated load-balancing.** Intelligent algorithms enable PowerPath to select the best data path for every I/O request. Manual load-balancing is replaced by automated path management and optimization. Because it balances workloads across data paths, no single path can become overloaded and create a bottleneck while others have underutilized bandwidth. Additionally, ongoing PowerPath testing also includes load-based testing which ensures path viability for handling the data traffic presented to it.
- **Application availability and performance optimization.** PowerPath ensures continuous data access while optimizing server and path usage by protecting the I/O path and applications in the event of a path failure. Application-transparent failover insulates core business operations from disruption.


ESG Lab Validation

ESG performed hands-on testing at EMC's lab in Cambridge, Massachusetts and audited EMC tests conducted over a longer period of time. Testing was designed to compare the responses of EMC PowerPath for Windows versus native Windows 2008 R2 MPIO in a dynamic SAN environment that simulated real-world SAN conditions with changing workloads and sudden bottlenecks. Similar results are expected for all PowerPath-supported operating systems.

Test Bed

The test configuration is shown in Figure 3. Two Dell servers running Windows 2008 R2 were connected via Emulex HBAs to Symmetrix VMAX storage (10x10 GB disks with RAID 5). Each server had access to dual SANs created with Brocade 8 GB Fibre Channel switches. One server was configured with EMC PowerPath 5.5 for Windows while the other relied on native Windows MPIO with the default round robin policy. Five additional Windows 2008 R2 servers were used to generate simulated backup reading workloads to create resource contention.

Figure 3. EMC PowerPath for Windows vs. Windows Native MPIO Test Bed



Getting Started

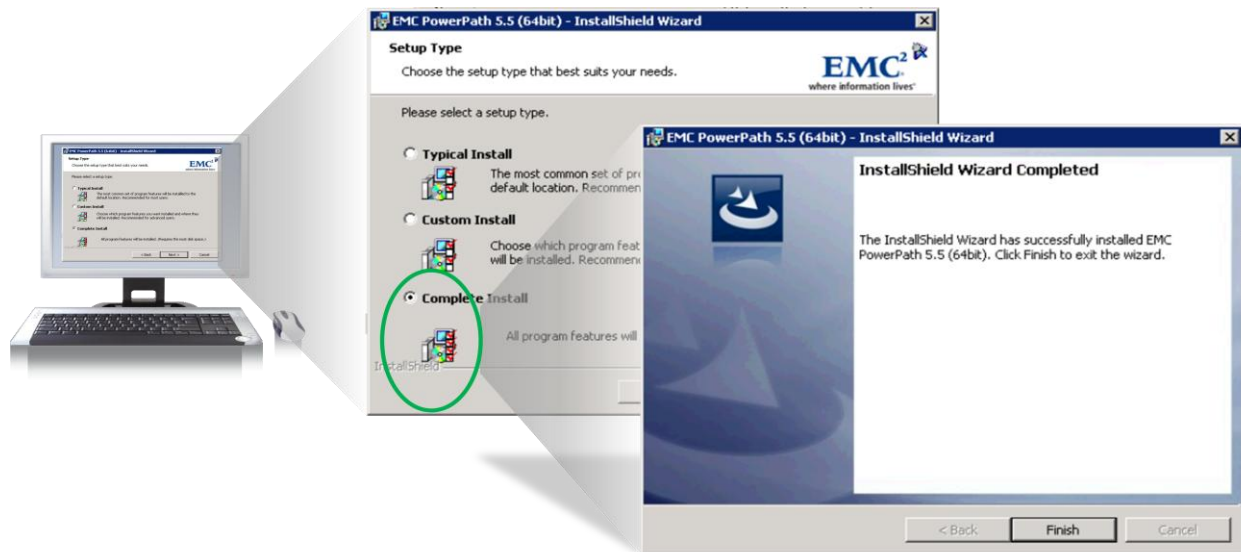
ESG Lab tested the installation procedure for EMC PowerPath 5.5 with a goal of assessing the ease of deploying optimized load balancing and failover on a new server.

ESG Lab Testing

Installation was fast and simple. ESG Lab performed the recommended *complete* installation (see Figure 4) which includes PowerPath Multipathing and PowerPath Migration Enabler. Only PowerPath Multipathing was tested. The installation had the ability to add other array types, but this capability was not included in the testing. The wizard-driven process was completed with seven mouse clicks and one reboot in approximately three minutes.

Please note that PowerPath Configuration Checker, included at no charge, uses EMC's internal website to run a utility on each host and generate a report verifying adherence to EMC Support Matrix standards for PowerPath, as well as recommended changes to ensure continued support for failover and load-balancing. The utility captures information for a variety of EMC tools in addition to Configuration Checker. ESG Lab verified that it took only three seconds to generate a report on the host.

Figure 4. Installation



Why This Matters

Ease of set up and installation minimize management effort and cost. As data centers grow in capacity and complexity, management tasks expand and IT organizations with constrained resources struggle to keep up. Consequently, administrators often retain product defaults for simplicity and time savings. This is one reason why native MPIO solutions default to round robin; its simplicity can preempt support calls. However, this simplicity is negated with native MPIO because different operating systems offer different features and installation procedures. As a result, administrators in heterogeneous environments must manage multiple installation processes that require multiple reboots. The Windows MPIO installation requires rebooting after configuring each feature and when configuring multiple data paths. It also requires reboots when adding different types of storage devices.

ESG Lab confirmed that PowerPath is easy to install and requires only a single reboot. By providing the ability to check configuration and versioning requirements before installation, PowerPath saves time and ensures a smooth installation the first time. No reboot is required to add different types of EMC or third-party storage arrays.

EMC PowerPath for Windows vs. Windows Native MPIO

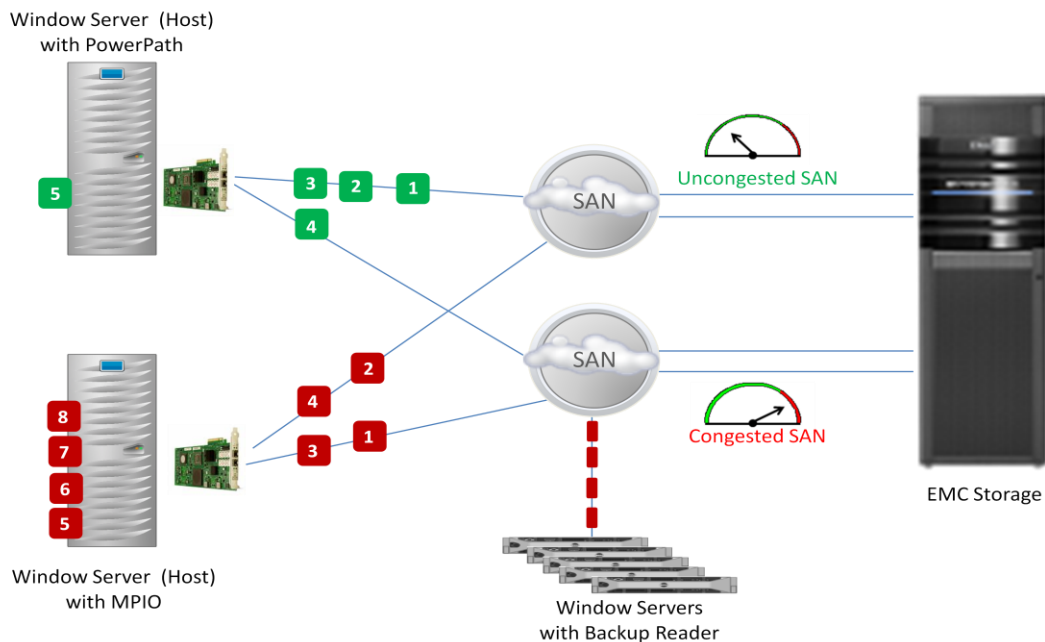
Native MPIO is included with most server operating systems, and while these native solutions offer some options, the default policy (which most customers retain) is round robin. A round robin policy distributes I/O among all data paths in sequence; failed paths may be taken out of the rotation, but clogged paths are not. Round robin considers all paths to be equal as long as they are active. EMC PowerPath's default policy pushes I/O through the optimal path based on algorithms and awareness of conditions across the SAN.

ESG Lab Testing

ESG Lab testing was designed to track how PowerPath and MPIO respond to changes in the environment. The primary Windows hosts ran various I/O profiles created using the Iometer workload generation utility.² Some I/O profiles focused heavily on I/O per second (IOPS) to simulate Exchange database, OLTP, and file-serving workloads; others were throughput-intensive to simulate backup, video-on-demand, decision support, and Exchange log traffic. Five additional Windows servers created a bottleneck in the data path using a load simulating a SAN-based backup.

As Figure 5 demonstrates, I/O requests (indicated by green and red numbers) were generated by both servers. The server configured with native MPIO round robin made no adjustments for resource contention, but continued to send I/O down both data paths in sequence. The result was SAN congestion and I/O backing up on the server. The server configured with EMC PowerPath handled the bottleneck differently, identifying the jammed data path and moving I/O to the more available path instead. I/O did not back up on the server and the SAN remained uncongested, resulting in an overall increase in throughput.

Figure 5. EMC PowerPath vs. Windows Native MPIO



² Iometer is a workload generating tool that creates traffic and measures performance characteristics of storage solutions. The Iometer application and documentation is freely available at <http://sourceforge.net/projects/iometer>.

Performance

Application performance often suffers in the presence of resource contention. Application workloads such as Exchange databases, 4K and 8K OLTP, and file systems are IOPS-intensive; performance depends on the number of I/Os per second that the data path can handle. Other application types such as backup, decision support, Exchange logs, and video-on-demand are throughput-intensive; performance depends on the data path’s ability to handle large file sizes. ESG Lab compared the performance of PowerPath with the performance of Windows native MPIO for these workloads.


ESG Lab Testing

ESG Lab tested simulated IOPS- and throughput-intensive workloads.³ Iometer was used to generate the primary server workloads while the jamming hosts executed backup reading tasks.

ESG Lab generated file serving, Exchange database, and 8k OLTP workloads on both servers, one configured with PowerPath and the other with MPIO. At the same time, five other Windows servers generated backup reading workloads to create resource contention. PowerPath determined that the jammed path, while active and available, was not optimal and automatically redirected I/O to better paths. The MPIO server, in contrast, retained its round robin distribution and continued to place I/O on the clogged data path.

As shown in Figure 6, PowerPath handled 22% more IOPS versus MPIO for OLTP workloads, 31% more for Exchange data, and 33% more for file serving.

Figure 6. *EMC PowerPath vs. Windows Native MPIO, IOPS-intensive Workloads*

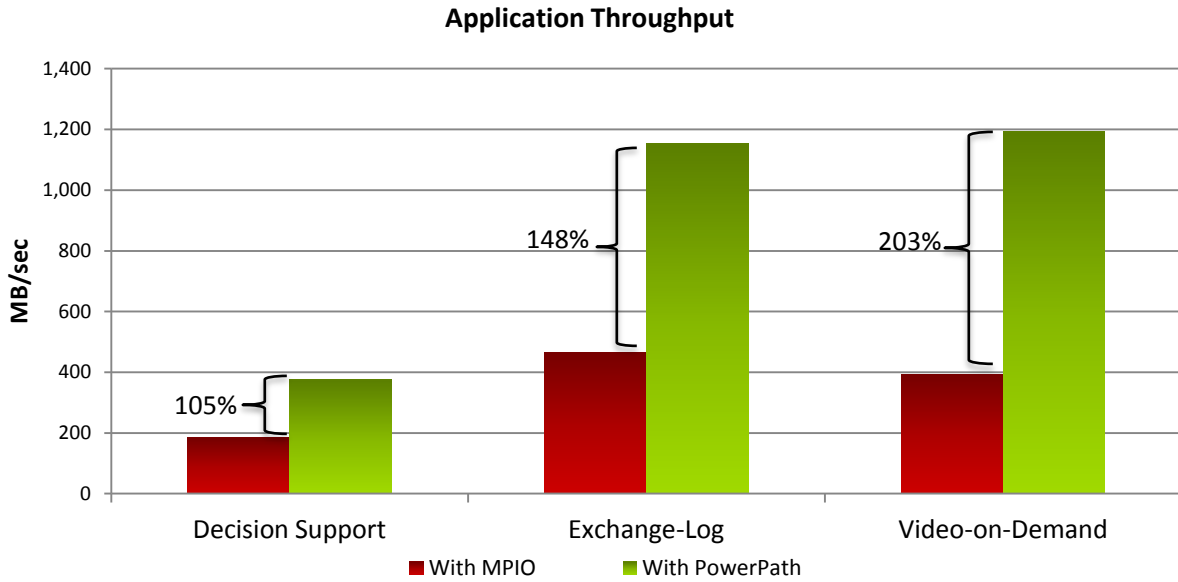


³ Note that application workload simulation often provides a clearer performance picture than actual customer workloads, which are likely to be affected by other data center characteristics.

The same tests were administered using decision support, Exchange log, and video-on-demand I/O loads on the primary servers. These simulate common occurrences such as users downloading large files that can degrade performance across the infrastructure. Figure 7 shows the results of these tests. The PowerPath advantage over MPIO was 105% for decision support, 148% for Exchange logs, and 203% for video-on-demand.


Figure 7. EMC PowerPath vs. Windows Native MPIO, Throughput-intensive Workloads

Result: Increased Throughput with PowerPath



Finally, ESG Lab tested the response time for Exchange database I/O on both the PowerPath- and MPIO-configured servers. The response time was 23% faster with PowerPath than with the native MPIO.

Figure 8. EMC PowerPath vs. Windows Native MPIO, Exchange Database Response Time



Why This Matters

Resource contention and SAN congestion are daily occurrences in today's dynamic data centers. They can cause performance and availability problems that threaten IT's ability to deliver on application service level agreements. Various application I/O profiles generate different loads and normal SAN operations such as application changes or adding/removing hosts create a dynamic environment. Performance degradation can arise from oversubscribed array ports, different I/O profiles zoned to the same storage ports, increasing queue length, misbehaving HBAs, degraded cables, etc. Recently, a popular topic of conversation has been contention caused by consolidated workloads in virtual server implementations; while it is certainly true that SAN problems can be exacerbated with consolidated workloads, the challenges of data path management and resource contention are quite conventional and related to everyday operations. A load-balancing solution that compensates and adjusts to the changing conditions in the data center can keep applications available and performing well.

Round robin MPIO policies aren't "smart." They simply alternate I/O among all paths without considering changes in server, network, or storage activities, array properties, HBA queue depths, I/O profiles, etc. In contrast, EMC PowerPath intelligently selects the optimal data path based on a complete view of paths from host through the SAN to storage port and disk, resulting in improved performance.

ESG Lab confirmed that PowerPath enabled better performance for both IOPS- and throughput-intensive workloads. Response time was faster and the SAN was less congested.

Advanced Monitoring and Management

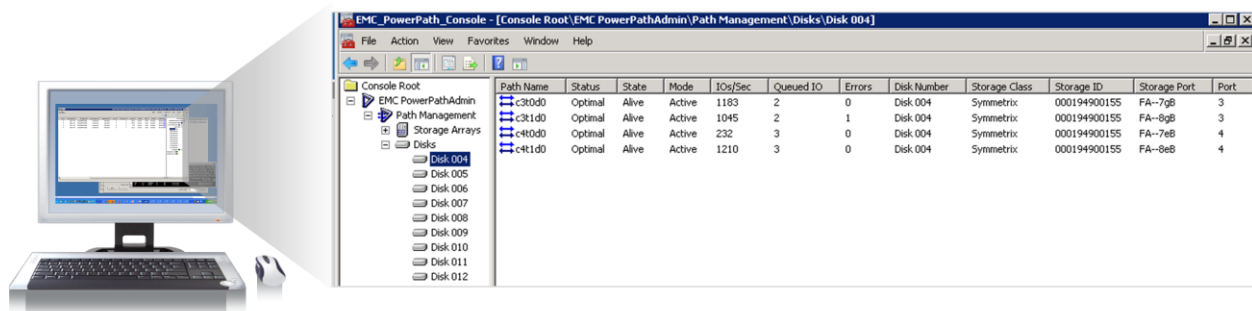
While this report is primarily about PowerPath performance compared to Microsoft Windows MPIO, it is worth noting some of PowerPath's advanced features.

The EMC PowerPath GUI is included with the installation and provides a central location for advanced monitoring and management. Individual disk paths, array ports, and adapters can be displayed, as can groups of disk paths. PowerPath 5.5 for Windows includes the PowerPath Administrator Console for local and remote monitoring of PowerPath for Windows hosts. Latency monitoring provides greater network response time visibility into the environment. Additionally, EMC has recently introduced PowerPath Viewer which consolidates details about large host environments into a single screen across hundreds of hosts running different operating systems. The Windows native MPIO solution includes a GUI, but most native solutions are only accessible by CLI, providing limited visibility and manageability.

ESG Lab Testing

Figure 9 shows the EMC PowerPath Administrator Console as seen through the Microsoft Management Console. This view is for an individual disk; four data paths and their details are viewable, including path status, IOPS, I/O queue, and number of path errors. The right side of the GUI shows what type of storage is in use, as well as array and storage port identification. Administrators can drill down for further detail by disk, array, or adapter.

Figure 9. EMC PowerPath Console

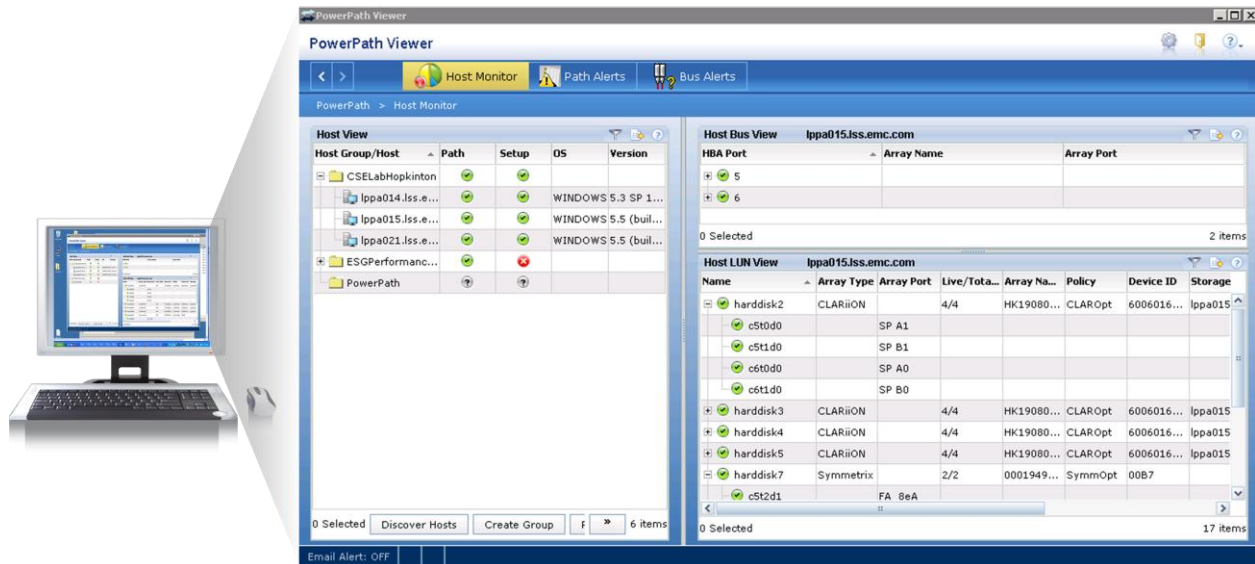


PowerPath also offers latency monitoring which can be seen through the PowerPath Monitor in the Microsoft Management Console or the command line interface. Server and storage administrators can identify, monitor, and proactively remedy path issues that may potentially impact application performance. With PowerPath latency monitoring, current or potential trouble spots can be identified by measuring I/O response times. Configurable threshold alerts enable a server or storage administrator to determine whether an application performance problem is caused by I/O issues.

For greater insight into data path topology, EMC PowerPath Viewer is a no-charge addition to the PowerPath installation that shows path status across numerous Windows, Linux, and VMware vSphere hosts. Viewer includes two components: the console and the host-based management component (for hosts other than ESX, which does not require it). The management component monitors host, LUN, data path, and bus events, and delivers information and alerts to the console over IP. Events are consolidated, so administrators can view and monitor PowerPath hosts from a single screen. Without Viewer, monitoring 20 data paths could necessitate logging into 20 separate machines. Also, multiple consoles can look at the same host; for example, virtualization and storage administrators can view the same information from different consoles. Servers can be grouped for viewing by operating system or application and Viewer can display host groups, LUNs, individual paths, and buses. Viewer shows the path heartbeat, remote ports, path and bus alerts, licensing state, and host versioning. E-mail alerts can be sent to administrators' consoles or smart phones in the event of a change in device status. The alert views consolidate path and bus alerts so an administrator can quickly identify problems.

Figure 10 shows PowerPath Viewer from the Host Monitor view. On the left, hosts are displayed within host groups; general path and setup health are displayed, as are operating systems and versions. ESG Lab created a path failure in the ESG performance host group; the red X indicates that there is a problem that the administrator can drill down into. On the right is detail for an individual host, with the screen split between the host bus and LUN views. With this detail, administrators can track errors, drill into alerts, and scan for hosts with management agents.

Figure 10. EMC PowerPath Viewer



Why This Matters

The ability to centrally monitor data paths for all physical and virtual hosts enhances visibility across the environment. Administrators no longer need to create elaborate scripts or decide not to monitor host paths. Data centers grow and change, adding complexity, especially as consolidation ratios increase. At the same time, IT is continually asked to do more with less. The ability to monitor data paths across the infrastructure provides a complete view of data paths impacting application performance while minimizing management time and effort. Also, in many data centers (particularly those implementing server virtualization), IT administrator roles are converging, and administrators are asked to manage across server, storage, virtualization, and network domains. Centralizing monitoring and consolidating the view from server through storage makes their jobs easier.

Native MPIO solutions only work for a single operating system; in heterogeneous environments, administrators must open individual GUIs which means they cannot gain an overall view of the environment. Limited GUIs and CLIs do not provide these overview and drilldown capabilities. While additional MPIO monitoring tools are sometimes available, they are difficult to install and most customers don't bother.

Administrators gain a better understanding with monitoring and management tools that provide visibility from the host through the SAN to the storage. This end-to-end visibility and manageability enables IT to meet aggressive service levels for application availability and performance, so problems can be located and addressed more quickly to eliminate application impact.

ESG Lab confirmed that PowerPath Console and PowerPath Viewer enable administrators to see all hosts centrally, view hosts by groups, and drill down into details including event alerts.

ESG Lab Validation Highlights

- ☑ ESG Lab confirmed that EMC PowerPath's intelligent load-balancing automatically selects the optimal data path for I/O in response to resource contention, providing an advantage over native Microsoft Windows MPIO for both IOPS- and throughput-intensive workloads.
- ☑ PowerPath's performance improvement over Windows MPIO was constant, ranging from more than 20% for OLTP and Exchange data base transactions to 203% for video-on-demand applications. PowerPath application response time was faster than MPIO when conditions degraded.
- ☑ EMC delivers advanced management functionality with the GUI-based console, including alerts and drilldown capability. In addition, PowerPath Viewer consolidates information about large numbers of hosts into a single screen, making it simple for administrators to gain a holistic view of data path health across the environment.

Issues to Consider

- ☑ The addition of PowerPath Viewer enables IT to view the current status of large numbers of hosts in one screen. A database would enable users to view performance statistics, show the impact of path failure on performance, and demonstrate historical trends that could aid in data center planning.
- ☑ Since PowerPath Viewer is new, ESG expects the utility to gain new features over time and hopes a central repository is in the product plans.

The Bigger Truth

No matter what dramatic and revolutionary technology advancements come along, one thing remains constant: end-users want the highest possible application performance and data availability—and it is IT's job to deliver. Data centers change, grow, and become more complex, but that doesn't change IT's mission. The challenge is to continue to provide the services required in the face of the growth and change that create resource contention. While many organizations consider implementing new performance improvement technologies such as solid state drives or faster storage and network protocols, these can be expensive and disruptive. In many cases, the IT department can greatly improve performance by simply optimizing the data paths already in place.

Shared infrastructure resources provide multiple servers with access to multiple storage arrays via a SAN. Data path failover and load balancing provide the opportunity to maximize availability and utilization of data center resources; however, all load balancing policies are not the same. Many organizations assume, incorrectly, that because the host OS includes native MPIO, they are getting the most out of their data paths. In addition, while some customers test their environments' ability to sustain a path failure, most cannot take the time and effort to perform tests that measure the performance impact of changing conditions.

ESG Lab compared EMC PowerPath 5.5 with Windows native MPIO in a round robin configuration. Testing confirmed that PowerPath, which automatically selects the optimal data path based on intelligent algorithms and a holistic view of the infrastructure, can deliver faster performance and higher availability than Windows native MPIO for both IOPS- and throughput-intensive workloads. PowerPath also provides far greater management and visibility features, including consolidating alerts and details in a single screen.

EMC PowerPath software has been delivering maximum application performance and availability for more than a decade; it runs on more than one million hosts at more than 30,000 customer sites. PowerPath may not be EMC's flashiest solution, but it can surely be called one of the most dependable. PowerPath does much of the everyday heavy lifting that keeps applications up and running and performing to users' satisfaction, and it remains the steady partner of administrators everywhere.

Appendix

Table 1. ESG Lab EMC PowerPath Test Bed

Servers	
(2) Dell Servers	Model R710 Windows 2008 R2 QLogic QLE2652 HBAs
(5) Dell Servers	Model R710 Windows 2008 R2 Emulex LP120002 HBAs
Storage	
(1) EMC VMAX Array	Microcode 5875.198.148 RAID5 LUNs
SAN	
(1) Brocade 1U Switch	Model DS-300B Firmware 6.1.0c
(1) Brocade 2U Switch	Model DS-5300B Firmware 6.1.0c
Software	
EMC PowerPath	Version 5.5
MPIO DSM	Windows 2008 R2
lometer	2008.06.18



Enterprise Strategy Group | **Getting to the bigger truth.**