Advantages of Hitachi Adaptable Modular Storage 2000 Family Symmetric Active-active Controllers in vSphere 4 Environments

Lab Validation Report

By Bob Ingram

November 2009

Summary

IT administrators working with asymmetric active-active storage systems know that it's time consuming to make the workload balance and capacity changes necessary to meet service level agreements as demands for performance and capacity constantly grow. Workloads that vary according to time of day, day of the week, or season add another layer of complexity. Storage systems that use asymmetric active-active controllers also require close cooperation between storage and virtual infrastructure administrators and careful analysis of workloads to optimize the IT environment.

The Hitachi Adaptable Modular Storage 2000 family of storage systems has Hitachi Dynamic Load Balancing Controllers, which are symmetric active-active controllers that dramatically simplify deployment and management in VMware vSphere environments. The dynamic load balancing and symmetric active-active controllers simplify configuring the deployment because no additional load balancing software is needed. These features also ease management of the IT environment after deployment because the self-optimizing storage system hides complexity and adapts to changing workloads and simplifies adding capacity. Hitachi Data Systems laboratory testing demonstrates that the 2000 family's symmetric active-active controllers eliminate path thrashing and more equally distribute workload, improving system performance compared to asymmetric active-active and active-passive storage systems.

Symmetric active-active controllers on the 2000 family simplify configuration on vSphere 4 clustered environments by eliminating the need to select a preferred path. In addition, testing in Hitachi Data Systems labs shows that the native multipathing plugin (NMP) from VMware operates efficiently on the 2000 family using various algorithms and reduces complexities of configuring asymmetric active-active storage.

This paper is written for IT professionals like storage administrators, ESX host administrators and application administrators who are charged with managing large, dynamic environments. It assumes familiarity with SAN based storage systems, VMware ESX and general IT storage practices.

For best results use Acrobat Reader 8.0.

Contributors

The information included in this document represents the expertise, feedback and suggestions of a number of skilled practitioners. The author recognizes and sincerely thanks the following contributors and reviewers of this document:

- Patricia Brailey
- Alan Davey
- Patrick Flueckiger
- Michael Heffernan
- Ron-An Lee
- Lisa Pampuch
- Rob Simmons
- Eric Stephenson

Table of Contents

Advantages of Hitachi Adaptable Modular Storage 2000 Family Symmetric Active-active Controllers in vSphere 4 Environments

Validated Lab Report

In today's climate of quickly changing IT environments that must keep pace with growing demands for performance and capacity, it's time consuming for system administrators to make the workload balance and capacity changes necessary to meet service level agreements. Workloads that vary according to time of day, day of the week, or season add another layer of complexity. When storage systems use asymmetric active-active controllers, addressing these changes is a time-consuming and labor-intensive task that requires close cooperation between storage and virtual infrastructure administrators and careful analysis of workloads to optimize the IT environment. Sharing of SAN-based storage on virtualized environments like vSphere 4 brings new challenges. Because several servers might share storage, ports and logical units (LUs), configuring these resources for optimum performance and availability can be complex.

The Hitachi Adaptable Modular Storage 2000 family of storage systems has Hitachi Dynamic Load Balancing Controllers, which are symmetric active-active controllers that dramatically simplify deployment and management in VMware vSphere environments. In both standalone ESX and ESXi environments and clustered solutions using VMware High Availability (HA), Distributed Resource Scheduler (DRS) and Fault Tolerance (FT), the dynamic load balancing and symmetric active-active controllers simplify configuring the deployment because no additional load balancing software is needed. These features also ease management of the IT environment after deployment because the self-optimizing storage system hides complexity and adapts to changing workloads and simplifies adding capacity. Hitachi Data Systems laboratory testing demonstrates that the 2000 family's symmetric active-active controllers eliminate path thrashing and more equally distribute workload, improving system performance compared to asymmetric active-active and active-passive storage systems.

Symmetric active-active controllers on the 2000 family simplify configuration on vSphere 4 clustered environments by eliminating the need to select a preferred path. Contrast the 2000 family's simplicity to configuring asymmetric active-active storage systems on vSphere clustered environments, which can be complex. Path management on asymmetric active-active storage systems requires that all ESX hosts set the preferred path to the asymmetric active-active storage system so that communication is to the proper storage processor. In addition, performance considerations for Fibre Channel ports and storage processor can be time consuming to troubleshoot because of interactions between ESX hosts using the same logical unit (LU). None of this is necessary on the 2000 family.

In addition, testing in Hitachi Data Systems labs shows that the native multipathing plugin (NMP) from VMware operates efficiently on the 2000 family using various algorithms and reduces complexities of configuring asymmetric active-active storage.

This paper is written for IT professionals like storage administrators, ESX host administrators and application administrators who are charged with managing large, dynamic environments. It assumes familiarity with SAN based storage systems, VMware ESX and general IT storage practices.

Hitachi Adaptable Modular Storage 2000 Family

The Adaptable Modular Storage 2000 family systems are the only midrange storage systems with the Hitachi Dynamic Load Balancing Controller that provide integrated, automated hardware-based front to back end I/O load balancing thus eliminating many complex and time consuming tasks that storage administrators typically face. This ensures I/O traffic to back-end disk devices is dynamically managed, balanced and shared equally across both controllers. No other midrange storage products that scale beyond 100TB have a serial attached SCSI (SAS) drive interface. The new point-to-point back end design virtually eliminates I/O transfer delays and contention associated with Fibre Channel arbitration and provides significantly higher bandwidth and I/O concurrency.

The active-active Fibre Channel ports means the user does not need to be concerned with controller ownership. I/O is passed to the managing controller through cross-path communication. Any path can be used as a normal path. The Hitachi Dynamic Load Balancing controllers assist in balancing microprocessor load across the storage systems. If a microprocessor becomes excessively busy, the LU management automatically switches to help balance the microprocessor load.

For more information about the Adaptable Modular Storage 2000 family, see the Hitachi Data Systems Adaptable Modular Storage 2000 family Web site.

Symmetric Active-active vs. Asymmetric Active-active

Typical asymmetric active-active midrange storage systems can be very complex to configure in an ESX environment, especially when multiple ESX hosts are utilized in HA, DRS or FT clusters. If it is not done correctly, severe performance issues can occur and path failover might not perform as expected in the case of a hardware issue. If proper connectivity and ESX path management are not performed, path thrashing can occur. The administrative effort required for an asymmetric active-active storage system is far more intense than on a symmetric active-active storage system. The initial configuration of the ESX cluster and scaling and troubleshooting become complex and administrative overhead can be high.

The Hitachi Adaptable Modular Storage 2000 family symmetric active-active storage controllers do not experience path thrashing because controller ownership of the LU is not a factor. Because controller ownership does not exist, the ESX administrator only needs to select the path management algorithm desired. LUs shared in HA, DRS and FT clusters need no special considerations.

Table 1 highlights some of key differences between asymmetric active-active and symmetric active-active storage systems in VMware vSphere 4 environment.

Table 1. Asymmetric Active-activee and Symmetric Active-active Storage System Comparison

Comparison Criteria	*Asymmetric Active-active*	*Symmetric Active-active*
Preferred path setup	Yes	No
SAN path configuration	Each HBA requires two available connections to each storage controller. Total minimum connections zoned = 4.	Each HBA requires one connection to a storage controller. Total minimum connections zoned = 2.
Configuration data required	• LU number • LU controller ownership by WWN • VMware `vmhba` to owned controller	LU number
ESX datastore configuration steps	1. Assign LUs to storage controllers. 2. Balance LUs across both storage controllers. 3. Assign LUs to all ESX hosts in the cluster. 4. Assign path management algorithm. 5. Assign preferred path on all ESX hosts.	1. Assign LUs to storage controllers. 2. Assign LUs to all ESX hosts in the cluster. 3. Assign path to round robin algorithm.
Administrative overhead	• High • Assignment of LUs to preferred path might take several minutes per path • Scripting can be difficult as each LU must be addressed individually • Mistakes might cause performance issues on the entire cluster • Path thrashing might not be immediately apparent and might require attention at a later time	• Low • GUI can be used • Scripting is simpler because all LUs are identically configured • No risk of path thrashing
Documentation requirement	• High • All ESX hosts in the cluster must utilize the same preferred path • Each ESX host must be documented as to `vmhba` and storage controller	• Low • All LUs assigned round robin
Troubleshooting	• High • Path thrashing might cause severe performance issues in the ESX cluster • A failure in a single ESX host can affect the entire cluster • Incorrectly assigned preferred path might change LU storage controller ownership affecting multiple ESX hosts	• Low • Failures are localized to the ESX host and do not affect the cluster as a whole • No risk of path thrashing

Path Thrashing

With active-passive storage systems, if your server cannot access a LU, or if access is very slow, you might experience path thrashing (VMware also calls this LUN thrashing). Path thrashing might occur when two hosts access the LU through different controllers, making the LU unavailable to either host. The LU might continually switch ownership from controller 0 to controller 1 and back in a ping-pong manner depending on I/O patterns from the datastores. In severe cases, I/O is rejected by the storage processor and applications can fail.

In situations where the primary path fails and a single host is accessing the LU, this is an acceptable method for retaining data availability. However in a situation where several servers access the same LU, this can cause the LU ownership to change many times and possibly cause I/O failures. To prevent path thrashing, active-passive storage systems require at least four Fibre Channel connections as shown in Table 1. In the case of a path failure, the second connection on the same controller is available to process I/O and no controller ownership change is required.

Asymmetric active-active storage systems can experience performance degradation due to I/O being passed between controller.

Note that in an asymmetric active-active and active-passive environments, detailed design of path management is required. In addition, the storage administrator must monitor path failovers and the movement of virtual machines (VMs) across volumes using vMotion.

Asymmetric Active-active Configuration Challenges

Asymmetric active-active storage systems require careful consideration for LU mapping to all ESX hosts. If the VM guests containing vDisks on a LU are configured using VMware HA or DRS, all ESX hosts in the cluster must have the same LU configured at the same host LU number, or HA, DRS or FT might not perform as expected.

Configuring an Asymmetric Active-active Storage System on an ESX Cluster

To successfully connect an ESX cluster to an asymmetric active-active storage processor, a single LU path must be configured on all ESX hosts in the cluster. The same active I/O path must be configured on all ESX hosts to the active storage controller.

After this is complete and the ESX host discovers the LUs, path management must be configured so that the active I/O path is to the active storage controller for each LU on each ESX host. While this might be a simple task on a two-node ESX cluster with one shared LU, it can be a complex and time-consuming task on a typical three-to-four node ESX cluster.

For example, in a four-node ESX cluster with 50 LUs, you must configure paths 200 times in the vCenter GUI, or run more than 200 commands from the vSphere command-line interface. A mistake in the active I/O path for an LU can result in poor I/O performance due to path thrashing and the cause might be difficult to determine without examining all 200 paths.

Symmetric Active-active Configuration

The 2000 family's symmetric active-active controllers simplify deployments. They eliminate the risk of path thrashing and distribute the workload across all paths connected to a given host or cluster. Because I/O is processed by both controllers, I/O can continue through the remaining paths without interruption in the event of a path failure. With the 2000 family, only two paths are required for redundancy.

Configuring I/O paths in an ESX environment becomes a simple task when the symmetric active-active feature is available. Because you do not need to consider which path is the preferred path, you do not need complex charts to manage paths. Any of the paths can pass I/O with no penalty. Zoning configuration is less complex because no consideration is needed for the HBA to connect with both storage controllers. A single path from a host bus adapter (HBA) to a storage controller is the minimum requirement for redundancy. Failover is accomplished without the need to renegotiate a connection.

Hitachi Data Systems recommends using VMware's default storage array type plugin (SATP) round robin path management algorithm in ESX for ease of administration and balance of workload across the entire ESX cluster. Some storage administrators might choose to use fixed or most recently used (MRU) path policies. Both of these policies are acceptable, however round robin provides ease of administration and balance of workload across all Fibre Channel ports. For more information about these path selection policies, see VMware's *Fibre Channel SAN Configuration Guide*.

Test Results

To demonstrate the capabilities of the Adaptable Modular Storage 2000 family, Hitachi Data Systems configured a Microsoft Exchange environment and tested it using Microsoft Jetstress. The goal of these tests is to compare asymmetric active-active and symmetric active-active storage controllers using both MRU and round robin multipathing algorithms from VMware. The test results report IOPS on each Fibre Channel port, response time and microprocessor utilization on the storage controllers. The performance results are not intended to demonstrate actual performance capability of the 2000 family; instead, they are intended to compare the improved ease of management of the Adaptable Modular Storage 2000 family with symmetric active-active controllers to storage systems with asymmetric active-active controllers.

These tests demonstrate that symmetric active-active controllers perform well in all path policies under this conservative I/O load. At higher loads, the round robin policy is able to utilize all of the connected Fibre Channel HBAs for greater fabric bandwidth if needed.

VMware provides a generic SATP. The 2000 family's symmetric active-active storage controllers use the SATP module VMW_SATP_DEFAULT_AA plugin. Asymmetric active-active storage controllers use the VMW_SATP_DEFAULT_AP plugin. For more information about the plugins, see VMware's *Fibre Channel SAN Configuration Guide*.

Note: All laboratory testing reports idealized results. In production environments, results can be affected by many factors that cannot be predicted or duplicated in a lab. Conduct proof-of-concept testing using your target applications in a non-production, isolated test environment that is identical to your production environment. Following this best practice allows you to obtain results closest to what you can expect to experience in your deployment.

Test Methodology

Jetstress was used to demonstrate the advantages of dynamic load balancing and symmetric active-active controllers on the Adaptable Modular Storage 2000 family. Jetstress simulates Microsoft Exchange 2007 database and log file operations with a targeted number of users.

Testing compared a Hitachi Adaptable Modular Storage 2500 system against an asymmetric active-active storage system using similar configurations. The performance of each storage system was compared against the requirements for a successful Exchange implementation.

Testing collected the following metrics:

- IOPS, compared against the requirements of Exchange
- Response time (Microsoft defines a successful Exchange implementation response time as less than 20ms for each read operation)
- Microprocessor utilization on the storage system

Table 2 describes the Jetstress test configuration used for all tests.

Table 2. Jetstress Test Configuration

Configuration	Details
Disks used	18 300GB 15K RPM disks
RAID level	RAID-1+0 2D+2D for databases RAID-1 1D+1D for logs
Database LUs	3 535GB for databases 3 100GB for .ogs
ESX host	Dell R-905 with 128GB memory and 4 dual core processors
Virtual machine	Windows 2008 1 virtual processor 16GB memory

Table 3 lists Jetstress test parameters used for all testing.

Table 3. Jetstress Test Parameters

Parameter	Description
Test scenario	Exchange mailbox
Number of mailboxes	570
IOPS per mailbox	0.576
Mailbox size (MB)	2048
Test type	Performance
Test duration	2 hours

Figure 1 illustrates the RAID group layout used on the 2500.

Figure 1. Hitachi Adaptable Modular Storage 2500 RAID Group Layout

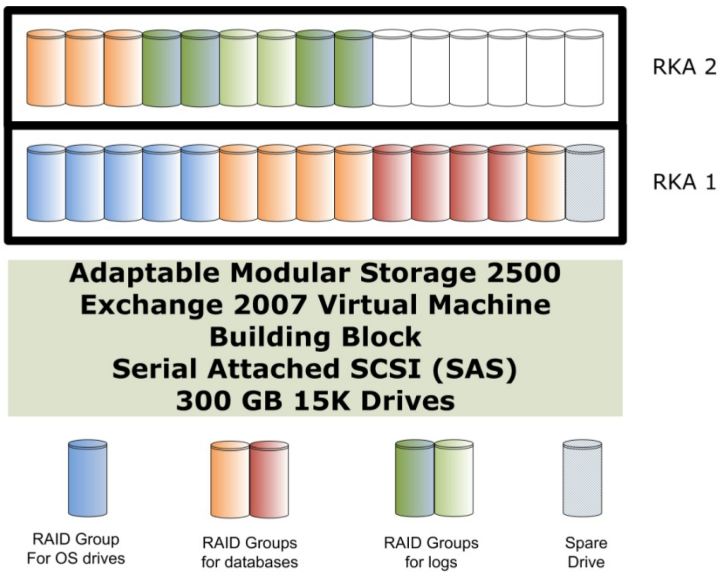

The asymmetric active-active storage system configuration closely matched the disk and RAID configuration used on the 2500 except that the disks were Fibre Channel attached.

Three tests were performed utilizing five LUs for Exchange databases:

- Single path, all I/O on a single Fibre Channel port using MRU path policy
- Balanced paths, I/O balanced between the two Fibre Channel ports using MRU path policy
- Dual paths, I/O balanced between the two Fibre Channel ports using round robin path policy

Test Analysis

Test analysis is separated into three sections:

- Port IOPS
- Response time
- Microprocessor utilization

The IOPS and response times are compared against Microsoft recommendations for Exchange Jetstress testing. Microprocessor utilization demonstrates how path thrashing affects the storage processors and how the 2000 family's Dynamic Load Balancing Controllers best utilize processor resources.

Test Analysis for Port IOPS on Asymmetric Active-active Controllers

The asymmetric active-active single path test run utilized LUs owned by controller 1 and port 1. The NMP was set to MRU. All I/O throughput exceeds the Jetstress requirements of 780 IOPS and less than 20ms response time.

Asymmetric active-active balanced paths distributed the LUs across both controllers and ports, with three LUs owned by controller 0 and 2 owned by controller 1. The NMP was set to MRU. This distributed the I/O across both ports but the load was not even due to the unequal number of LUs on the ports. All I/O throughput exceeds the Jetstress requirements.

The asymmetric active-active robin test used the same configuration as the symmetric active-active balanced path test, with the NMP set to round robin. Path thrashing occurs, causing degradation in performance. In this test the IOPS degrade by 31 percent. Figure 2 shows that although all I/O throughput exceeds the Jetstress requirements, response time does not meet the Jetstress requirement, therefore this test fails.

Figure 2. Port IOPS for Asymmetric Active-active Controllers

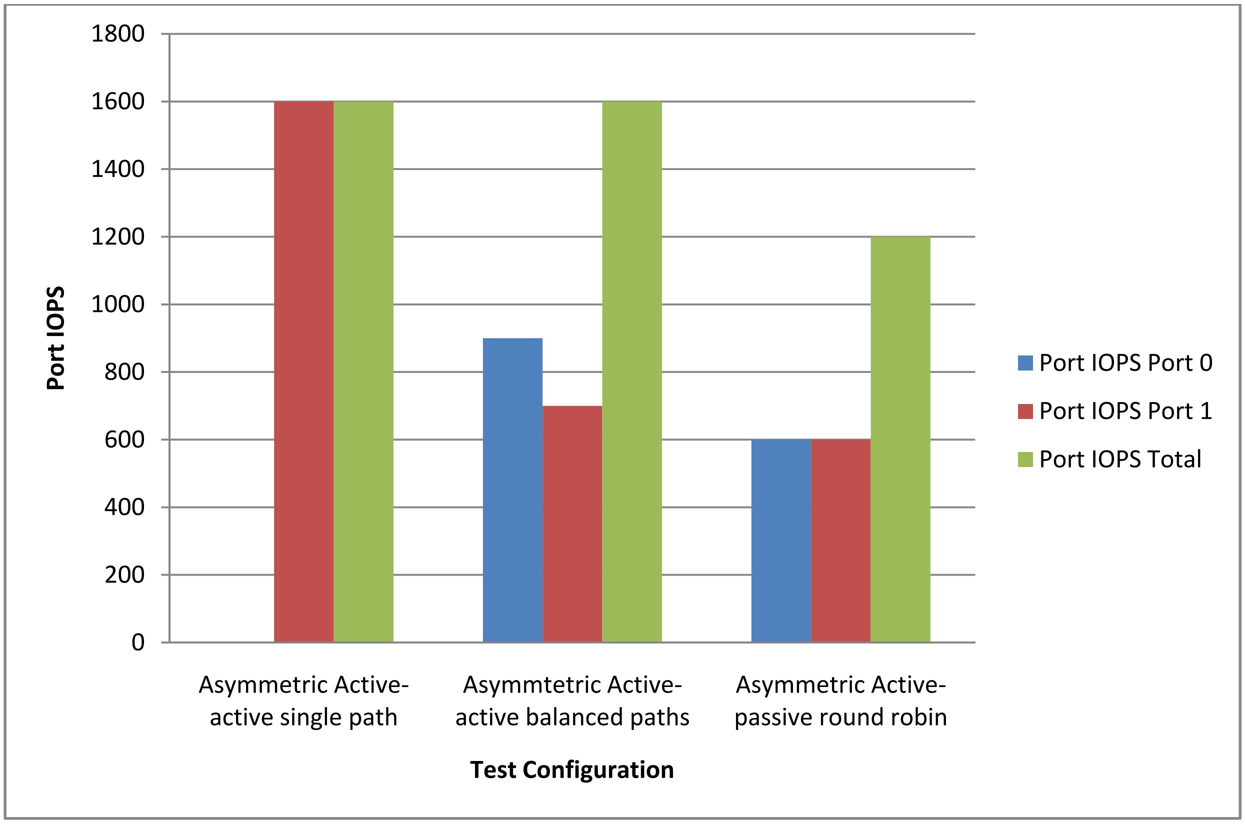

Test Analysis for Port IOPS on Symmetric Active-active Controllers

The symmetric active-active tests on the 2500 used the same test parameters as the asymmetric active-active tests described in the previous sections. Figure 3 shows that the round robin path policy performs as well as any of the other tests, but are simpler to configure because no analysis of the LUs and their preferred paths is required.

Figure 3. Port IOPS for Hitachi Adaptable Modular Storage 2500 with Symmetric Active-active Controllers

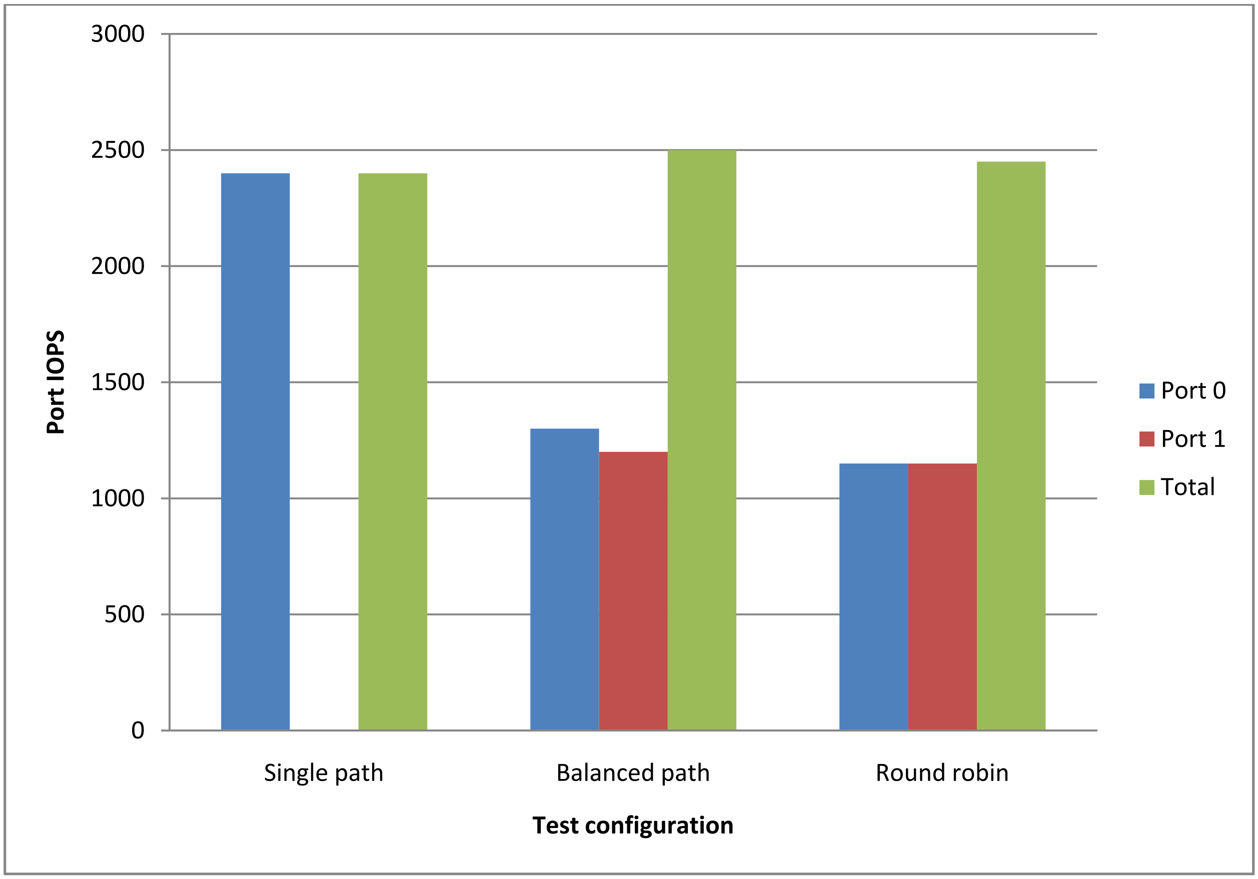

Test Analysis for Response Time on Asymmetric Active-active Controllers

Microsoft recommends that response time as measured by Jetstress be less than 20ms/read. Figure 4 shows that response times with both single path and balanced paths are in the acceptable range; however, the round robin average response time increases by 28 percent and fails to meet the response time requirement of less than 20ms.

Figure 4. Average Response Times for Asymmetric Active-active Controllers

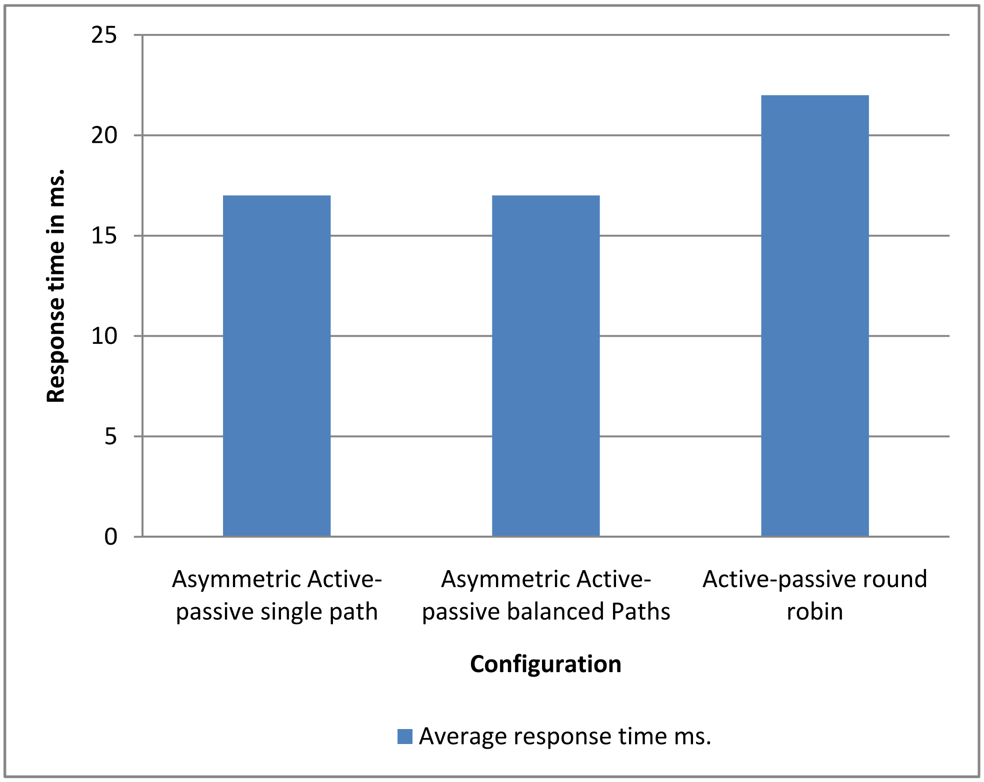

Test Analysis for Response Time on Active-active Controllers

Figure 5 shows that all three tests performed on the 2500 produced similar response time results. Because IOPS and response times are similar, the ease of setup, balance of Fibre Channel load and manageability of using round robin algorithm outweighs the other path policy options. Therefore, Hitachi Data Systems recommends using round robin for the simplicity of configuration within vSphere 4.

Figure 5. Average Response Times for Hitachi Adaptable Modular Storage 2500 with Symmetric Active-active Controllers

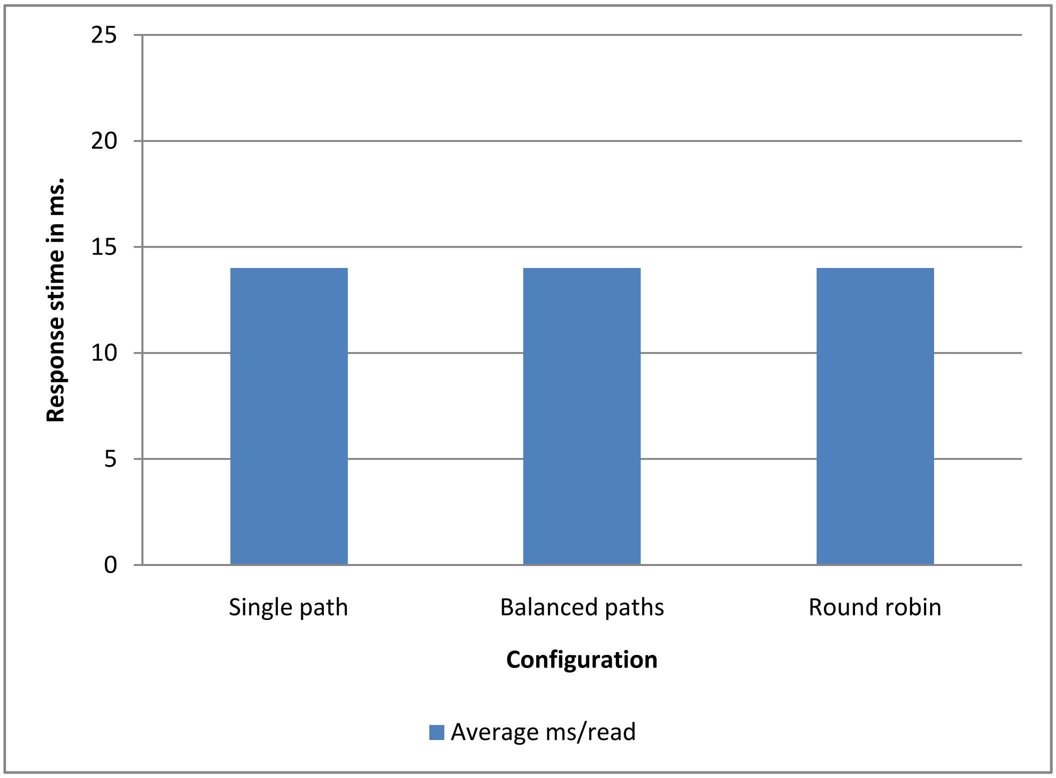

Test Analysis for Microprocessor Utilization on Asymmetric Active-active Controllers

Figure 6 shows microprocessor utilization rates for tests on asymmetric active-active storage.

In the asymmetric active-active single path test, only the microprocessor on controller 1 is utilized because all LUs are owned by controller 1. The asymmetric active-active balanced paths tests reduces microprocessor load by balancing across the controllers. This balance is not be equal because two LUs are owned by controller 0 and three are owned by controller 1. The asymmetric active-active round robin test produces very high microprocessor load, which is due to path thrashing.

Figure 6. Microprocessor Utilization Rates for Asymmetric Active-active Controllers

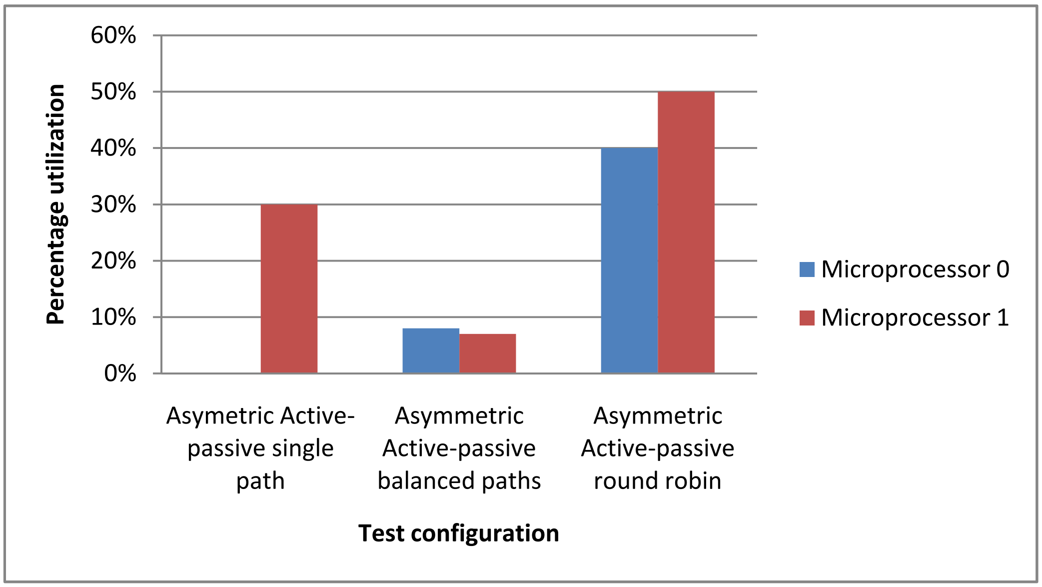

Test Analysis for Microprocessor Utilization on Symmetric Active-active Controllers

As shown in Figure 7, the symmetric active-active round robin tests demonstrate that the microprocessor load is far lower than the on the asymmetric active-active tests because the symmetric active-active controllers experience no path thrashing, which causes high utilization rates. When using the round robin path policy, the asymmetric active-active storage systems averaged 45 percent utilization, but the 2500 averaged a 20 percent utilization rate.

On the symmetric active-active single path test, the microprocessor load differs from the asymmetric active-active test. On the asymmetric active-active test, I/O is completely driven through controller 0; however, with Hitachi Dynamic Load Balancing controller technology, on the symmetric active-active test controller 1 assists in processing I/O and shares part of the workload. This shows the advantage of the Hitachi Dynamic Load Balancing controllers.

Load distribution in the symmetric active-active balanced path test is more equal than in the asymmetric active-active test, again due to dynamic load balancing. Note that no I/O was directed through controller 1; however the Hitachi Dynamic Load Balancing controller technology utilized the microprocessor to more equally distribute the load.

Figure 7. Microprocessor Utilization Rates for the Hitachi Adaptable Modular Storage 2500 with Symmetric Active-active Controllers

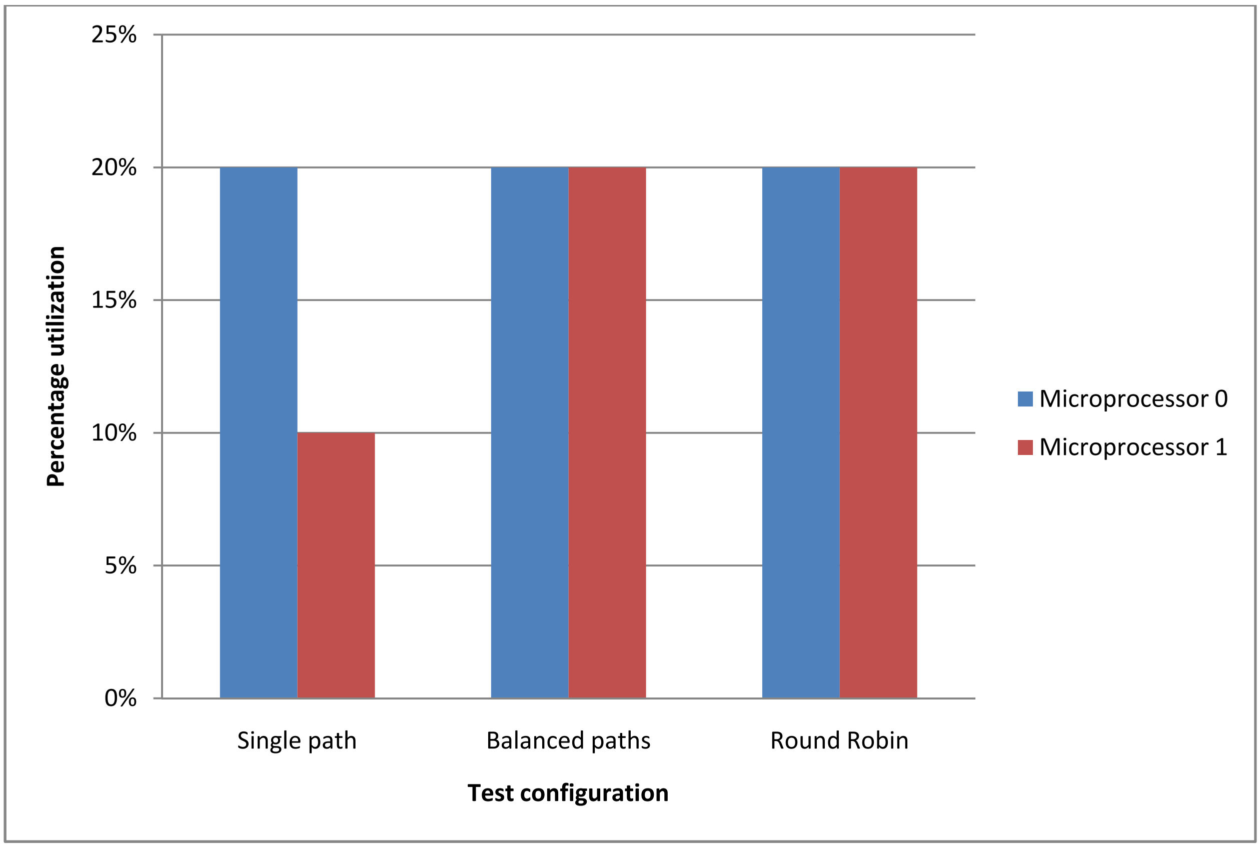

Conclusion

Hitachi Data Systems testing shows that Hitachi Adaptive Modular Storage 2000 family storage systems with symmetric active-active controllers provide throughput stability and sometimes increased performance as well as path redundancy with minimal administrative effort. Symmetric active-active storage systems like the 2000 family allow the use the round robin path policy of the VMware ESX I/O drivers and provide superior utilization of resources like Fibre Channel bandwidth and storage system microprocessors by ensuring that the I/O load is balanced between all available Fibre Channel connections.

For more information about the Hitachi Adaptable Modular Storage 2000 family, see the Hitachi Data Systems Web site, your sales representative or your channel partner.

Hitachi Data Systems Corporation

Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.uk@hds.com

AS-021-00 November 2009