

HP StorageWorks P4500 SAS SAN vs. Dell EqualLogic PS6000XV: Feature and Functionality Comparison

Test report prepared under contract from Hewlett-Packard

Executive summary

Hewlett-Packard ("HP") commissioned VeriTest, the testing service of Lionbridge Technologies, Inc., to compare features and functionality between the HP StorageWorks P4500 SAS SAN Solutions and the Dell PS6000XV storage arrays. The 3-node HP P4500 SAN, a Virtualization SAN with an HP P4500 expansion node, has a current list price of \$96,600.00. The 2-node Dell PS6000XV currently has a list price of \$110,000. The goal of this study was to evaluate these devices in terms of high availability, capacity utilization under various common usage scenarios, and ease of upgrading the storage solution to meet future storage requirements.

Both HP and Dell provide a scalable iSCSI based solution, each claiming the following benefits: scalability, flexibility, cost effectiveness and high availability. Lionbridge developed a test methodology that compared these benefits between the two vendors to help end users better understand the real-world differences experienced with each product under test.

Both solutions are based on a scale-out technology, combining multiple nodes into a cluster and thus presenting a common storage pool. For these comparable prices, the configurations below are what were available for each product. Therefore, the hardware configurations were not identical with regard to number of nodes and disk spindles.

The HP P4500 SAN configuration consisted of a three (3) node cluster running SAN/iQ version 8.1.00.0047.0 with RAID5, and 2-way replication. Each node contained 12 450GB 15K RPM SAS drives with a total raw capacity of 16.2TB.

Key findings

- ❑ The 3-node HP P4500 SAN solution provides up to 32% greater usable storage capacity, ([see pg 13](#)).
- ❑ The 3-node HP P4500 SAN solution provides better High Availability features than the 2-node Dell PS6000XV SAN solution - at a 14% lower price. ([see pg 1](#))
- ❑ The 3-node HP P4500 SAN solution provides greater availability than the Dell EqualLogic 2-node PS6000XV solution. Under the three failure scenarios tested, NIC, Node and Double-disk, the HP P4500 SAN experienced no downtime, whereas the Dell PS6000XV experienced as much as 30 minutes of downtime ([see pgs 8-11](#)).
- ❑ When retaining data via snapshots the 3-node HP P4500 SAN solution has as much as 62% lower storage overhead, than the Dell EqualLogic 2-node PS6000XV ([see pg 15](#)).
- ❑ With its High Availability features, the 3-node HP P4500 SAN complements the VMware vSphere 4 Fault Tolerance feature, whereas the Dell PS6000XV does not provide High Availability ([see pg 13](#))
- ❑ Platform refresh of a Dell PS6000XV SAN take 10 times as long as for a HP P4500 SAN – 5 hours 39 minutes vs. 34 minutes ([see pg 18](#)).

The Dell PS6000XV configuration consisted of a two (2) node cluster running firmware version 4.1.3 with RAID5 striped across both nodes. Each node contained 16 450GB 15K RPM SAS drives, with dual controllers, and total raw capacity of 14.4TB.

To identify the differences between the two products under test, Lionbridge developed a test methodology that executed several common scenarios encountered by storage administrators. To best highlight these scenarios, Lionbridge used a fictitious company, called ACME Corp, and identified several requirements for comparing and contrasting the differences between the two products under test.

ACME Corp is a hypothetical mid-sized company providing storage for online records management. They are looking to upgrade their existing SAN solution to an iSCSI based centralized storage solution that will grow with the company needs and integrate with their existing VMware ESX environment. Aside from meeting current storage capacity requirements and future storage growth needs, High Availability ("HA") is of most importance to ACME Corp. ACME Corp has Service Level Agreements (SLA) with their customers in which any outage (defined as the customer being unable to access the data) results in penalties ACME Corp pays to their customer(s) affected by the outage. For ACME Corp every minute of downtime results in a \$1,000 penalty paid to each customer affected by the outage.

For each type of enterprise, there is a specific cost to any service interruption, of any length. But the general question "What does an outage cost?" has no meaningful answer. In this research, Gartner develops a method for estimating the cost of an outage for a particular organization. Read the full report "[Q&A: How Much Does an Hour of Downtime Cost?](#)" from 9/29/2009, to find out more.

In addition to HA, ACME Corp employs snapshots as a means to maintain readily available online backups to transactions. Therefore, the ability to maintain several snapshots, as well as the ease with which the process is managed are key factors in determining the storage solution selected. For Disaster Recovery ("DR") purposes, ACME Corp has a sister site located several miles away to act as a storage replication target. Remote replication efficiency is another item of importance to ACME Corp for determining the storage solution selected. Finally, as with many other companies, ACME Corp is facing the rapid growth in storage space requirements. Due to this fact, it is also important that ACME be able to easily and rapidly upgrade or scale the storage solution to meet additional capacity requirements.

Although the focus of this white paper's comparison is not on platform performance, performance is an important consideration when selecting the storage solution. To help end-users determine their options for backend storage when deploying an Exchange server, Microsoft maintains a benchmark for Exchange called ESRP that is a good proxy for random IO. Both vendors have published ESRP results, click [here](#) for further details or contact the vendors directly to discuss performance for other types of workloads. VMware also hosts a website with performance information, click [here](#) for further details.

To assist the end-user, Table 1 below highlights system specifications associated with platform performance capabilities.

Item	HP P4500 SAS	Dell PS6000XV
Nodes	3	2
Controllers	3 Active / 0 Passive	2 Active / 2 Passive
CPUs	3 Active / 0 Passive	2 Active / 2 Passive
RAID Configuration	6 RAID-5 (5+1) raid sets	4 RAID-5 + Hot Spare (6+1+1) raid sets
Disks	36 15K 450GB SAS disks 36 of 36 disks active for performance	32 15K 450GB SAS disks 28 of 32 disks active for performance
Formatted Usable Capacity	12.03 TB	10.98 TB
Cache / RAM	7.5 GB Active / 0 GB Passive	4 GB Active / 4 GB Passive
Bandwidth	6gbit (6 x 1gbit active)	8gbit (8 x 1gbit active)
Volume Connectivity	iSCSI Standard	iSCSI Standard
Volumes Supported / group	1500	1024
Hosts Supported / group	4000	512
Max Disks Supported / group	384	256

Table 1: Performance Related Platform Specifications

Testing revealed that the 3-node HP P4500 SAN Solution provided greater usable storage capacity, while also providing better high availability features than the 2-node Dell PS6000XV solution at a lower price point.

Testing also revealed that the 3-node HP P4500 SAN solution provided lower storage overhead when retaining data via snapshots and remote replication than the 2-node Dell PS6000XV solution.

The 3- node HP P4500 SAN solution required fewer steps and only 1/10 the time required to perform a platform refresh as compared to the 2-node Dell PS6000XV solution.

Finally, the high availability features of the HP P4500 SAN complemented the VMware vSphere 4 Fault Tolerance feature, thereby providing a better HA solution for ACME Corp.

Testing methodology

ACME Corp is looking to upgrade their existing SAN solution to an iSCSI based storage solution. They need a solution that meets the following criteria:

- Minimum usable capacity of 10 TB
- Provide the ability to survive the most common of hardware failures (i.e. network, system, controller and disk failures)
- Provide snapshot capabilities for online point-in-time data retention
- Provide remote replication capabilities for disaster recovery requirements
- Be compatible with and integrate with VMware vSphere 4, and specifically work with the Fault Tolerance configuration capabilities
- Provide an ease of upgrade path for future-proof needs

With the above requirements in mind, the HP P4500 SAN solution and the Dell PS6000XV arrays were viable options to consider.

Lionbridge configured a test environment that consisted of eight (8) HP BL460c G6 Blade Servers in a c3000 enclosure. The c3000 enclosure (backplane is shared) included the following:

- Two (2) Flex-10 interconnects shared by all blades
- One (1) 10GigE uplink for the internet connection
- One (1) 10GigE uplink for dedicated iSCSI traffic

For network connectivity a single 48 port HP ProCurve 3400cl was used. This switch contained two VLANs, one that was used for administrative, remote desktop protocol (RDP) access and internet connectivity. The second VLAN was configured to provide dedicated iSCSI data traffic.

One of the blades was configured with Microsoft Windows Server 2008 R2 x64. This blade was used to manage the storage arrays under test. It had the HP Centralized Management Console ("CMC") application installed to manage the HP P4500 SAN solution and the Dell EqualLogic Group Manager GUI to manage the Dell EqualLogic arrays.

Two blades were configured with VMware vSphere 4 as ESX host servers. Each ESX host was configured with a single Windows Server 2008 R2 x64 Virtual Machine ("VM"). Each VM was configured with two (2) vCPUs and 4 GB RAM. A single NIC was dedicated for virtual machine traffic on each host. One ESX host/VM was configured to connect to the P4500 SAN solution, the other configured to connect to the PS6000XV array. The ESX configuration was modified to connect to their respective storage solution using the built-in ESX software iSCSI initiator. IQN authentication was used to manage the iSCSI connections. The VM image was contained on a 250 GB volume, thick provisioned on each storage solution for their respective VM.

The HP P4500 SAN was configured using three (3) nodes. The cluster was configured into disk RAID5 with Network RAID configured in 2-way replication (as of publication known as Network RAID 10) for all volumes. Each P4500 SAN node was configured with a single 802.3ad bond using both embedded NICs.

The Dell PS6000XV was configured with a two (2) node, RAID5 striped configuration. No synchronous data replication was configured as this was not a feature available at the time of testing. Each EqualLogic PS6000XV node was configured with its default active/passive controller configuration, and each controller had four (4) 1GigE network ports configured.

For each of the test cases listed below, a single 100GB volume was created on each storage solution and presented to their respective VMs as a second hard disk drive via the ESX software iSCSI initiator. This feature was enabled in the ESX configuration and was setup to use IQN authentication. For the VMware Fault Tolerance test case only, the storage arrays were configured to permit multiple client connections to allow two ESX hosts to connect to the provisioned storage.

The industry standard benchmark utility IOMeter was used to generate workload to the 100GB LUN for each test case. The LUNs were kept in a raw, non-formatted state. All of the tests used an OLTP-like workload that online database applications typically encounter. Like many other industry standard OLTP benchmarks, including the majority of the top 10 TPC-C price/performance results, we used an 8KB request size with a mix of random read and write operations. IOMeter was configured with 8 outstanding IOs, 8KB request size, a mixture of 60% random read and 40% random write, with an 8KB alignment. The IOMeter workload was started five (5) minutes prior to each test case, and was allowed to run through completion of the test case.

Figure 1 below provides an overview of the test environment.

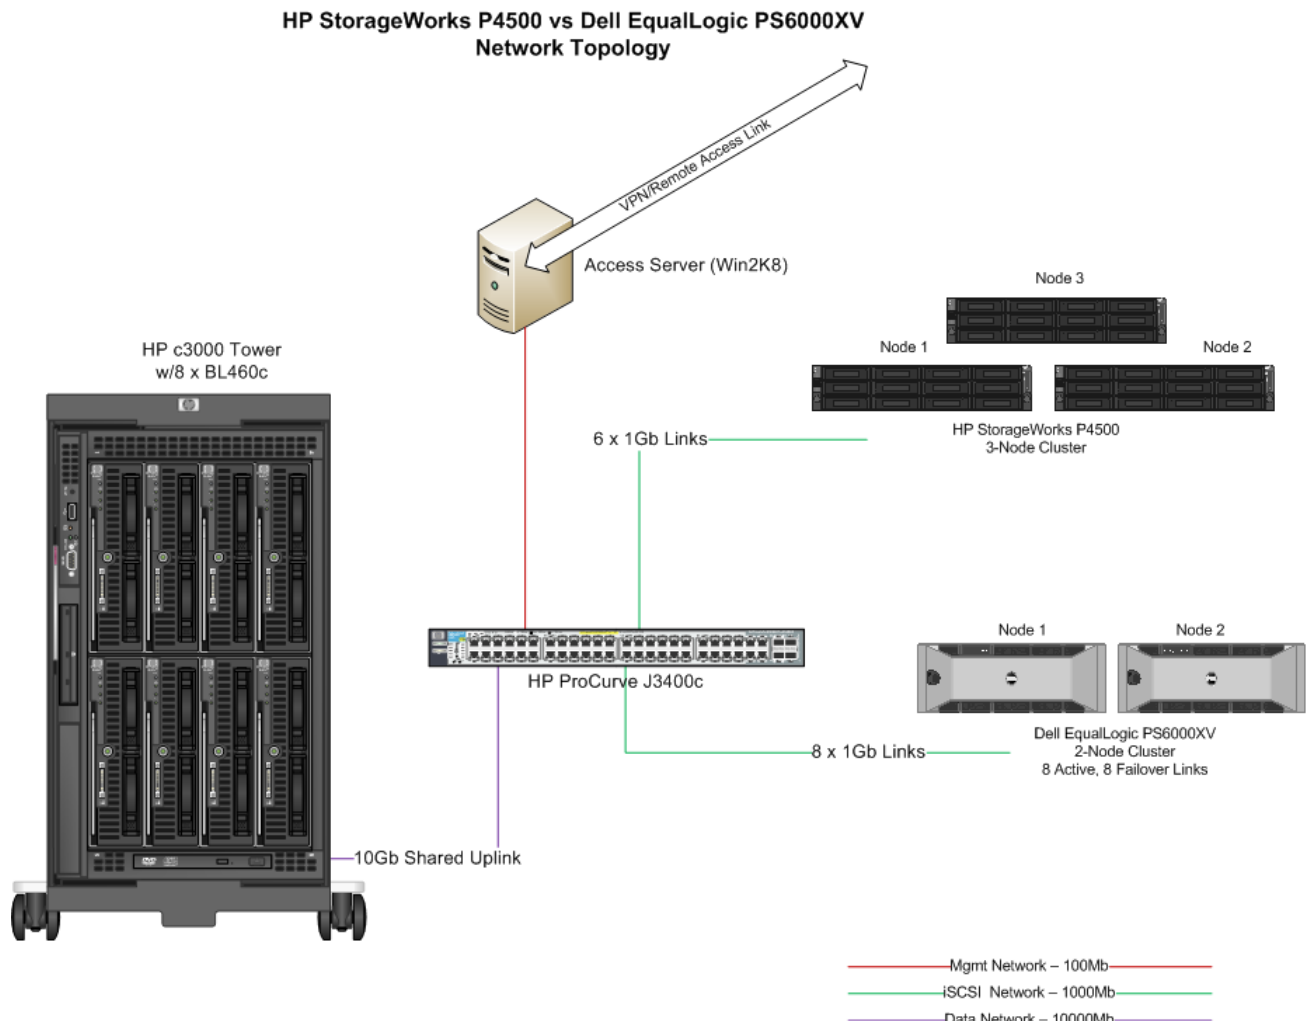


Figure 1: Test Environment Network Topology

Several test cases were developed to highlight key characteristics and to compare the HP P4500 SAN solution with the Dell PS6000XV storage arrays. These test cases were designed to address the requirements listed above for the fictitious ACME Corp. Each test case methodology is detailed in the sections below.

High Availability Comparison

It is important to understand what effect a hardware based failure can have on data availability. There are numerous points in the architecture of any storage solution where a hardware failure can occur. This section covers several common hardware failure scenarios. Lionbridge Test Engineers executed these hardware failure tests, and observed for the effect the failure had on the overall availability of the data. This section is broken into several subsections that focus on the specific hardware failure simulated.

NIC Failure

This test case was designed to observe behavior of each storage solution when a network failure occurred. In this case, a network failure was introduced by removing all network cables to one of the storage nodes. Hardware-based NIC only failure (for multiple NICs) is a very rare event, likely resulting from human error. Under a normal Best Practice configuration, the NICs would be connected to multiple switches which reduce the risk of a hardware failure causing the outage. For purposes of testing, this “human error” event was compared to see how each storage solution behaved under such an incident.

For this test case, a test engineer pulled all four cables from the primary controller on one of the EqualLogic nodes at the same time. Next, the network cables were pulled from one of the HP P4500 nodes to simulate the same type of failure on that platform. The test engineer observed the behavior of the VM and the IOMeter application. Also observed was the “Group Monitor” on the EqualLogic Group Admin GUI, as well as the “Performance Monitor” on the CMC. Each management application was also monitored for any errors reported. After five (5) minutes the network cables were reconnected and the arrays were allowed to return to optimal operation.

The ESX host for each VM was also monitored using vCenter. In addition, an active RDP connection to the VM was used to verify VM availability and IOMeter workload continued to function.

System Failure

This test case was designed to observe the behavior of each storage solution when a single node failure event occurred. Such a failure event can occur due to a number of reasons, such as power supply failures, system board failure, human error, etc.

For this test case, a test engineer simulated a system failure by pulling the power cable(s) from one of the nodes in each platform. The test engineer observed the behavior of the VM and the IOMeter application. Also observed was the “Group Monitor” on the EqualLogic Group Admin GUI, as well as the “Performance Monitor” on the CMC. Each management application was also monitored for any errors reported. After five (5) minutes the power cables were reconnected and the arrays were allowed to return to optimal operation.

The ESX host for each VM was also monitored using vCenter. In addition, an active RDP connection to the VM was used to verify VM availability and IOMeter workload continued to function.

Double-disk Failure

This test case was designed to observe the behavior of each storage solution in the event a double-disk failure (two disks on the same node) event occurred. Although relatively rare, double-disk failures do occur and it is important to understand the impact such an incident has on the availability of data.

For this test case, a test engineer simulated a double-disk failure by pulling two disks from the same node at the same time. Two odd numbered disks were pulled from one of the P4500 SAN nodes, as well as two odd numbered disks from one of the EqualLogic PS6000XV nodes. The test engineer observed the behavior of the VM and the IOMeter application. Also observed was the “Group Monitor” on the EqualLogic Group Admin GUI, as well as the “Performance Monitor” on the CMC. Each management application was also monitored for any errors reported. After five (5) minutes the disks were re-inserted and the arrays were allowed to return to optimal operation.

The ESX host for each VM was also monitored using vCenter. In addition, an active RDP connection to the VM was used to verify VM availability and IOMeter workload continued to function.

VMware Fault Tolerance Scenario

This test case was designed to demonstrate the HA capability of the storage solution when running within a VMware vSphere 4 ESX cluster with Fault Tolerance (“FT”). A two-node ESX cluster was configured with Fault Tolerance enabled. This cluster was configured to have access to both the P4500 SAN and PS6000XV storage clusters.

Future plans for ACME Corp are to move to a multi-site configuration in order to better provide data availability. By configuring into a multi-site configuration, the HA capabilities are more fully leveraged so that single points of failure (power distribution units, network infrastructure, racks, buildings, etc) are mitigated even further than the current configuration.

For this test case, a test engineer simulated a system failure and an ESX host failure at the same time. To accomplish this, the engineer removed the power cables from one of the ESX hosts as well as one of the nodes from each of the storage arrays. Testing was conducted serially. Power was pulled from the

ESX host first to allow the Fault Tolerance failover to occur. As this was running, the power was then pulled from one of the storage nodes under test. VMware FT was configured so that one ESX host was running the primary VM for the HP P4500 SAN and the other ESX host was running the primary VM for the Dell PS6000XV SAN. Prior to initiating the failure, the IOMeter OLTP script was executed on the VM to generate workload at the time of the failure.

To verify that the VMware FT configuration was sound, a baseline test was executed for each storage solution in which just one of the ESX hosts was failed by removing power to it. This initiated the FT failover process, and the Lionbridge Test Engineer verified the VM failed over properly and no outage was experienced by the running VM or the IOMeter OLTP workload that was generating active I/O at the time.

First, the P4500 SAN was tested by removing the power from the ESX host running the “live” (or primary) VM. Once failure was detected and the test engineer confirmed failover was initiated, the power was removed from one of the P4500 SAN nodes in the cluster. The test engineer then observed VMware vCenter and the VM console as well as the CMC GUI for any errors.

Second, the EqualLogic PS6000XV was tested by removing the power from the ESX host running the “live” (or primary) VM. Once failure was detected and the test engineer confirmed failover was initiated, the power was removed from one of the PS6000XV nodes in the cluster. The test engineer then observed VMware vCenter and the VM console as well as the Group Admin GUI for any errors.

Overall Capacity Utilization

ACME Corp has a current storage capacity requirement of 10TB. This is to cover existing data storage needs as well as expected growth over the next 12-14 months. Due to common budget approaches, the next opportunity to add additional storage will come during the next fiscal year.

This test case reviewed the available (usable) storage under different RAID configuration options available for each storage platform under test. This was used to determine the RAID options available to meet the minimum capacity requirements for ACME Corp of 10TB.

For this test case, the test engineer simply configured each storage array in the various RAID types offered in the platform to determine the available capacity for each type.

Capacity Utilization with Snapshots

This test case compared the amount of storage provisioned by each platform for snapshot retention. Also compared was the process of taking snapshots and the observed results.

For this test case, the test engineer created a 100GB volume, presented it to the respective VM as a second HDD as a raw device, and then filled it to 100% capacity using the IOMeter benchmark utility. For this data fill operation, a sequential write workload was configured for IOMeter with the following settings:

- Outstanding IOs: 8
- 100% Sequential Write
- 64K Request size
- 64K Alignment

An initial snapshot was then taken of the base data. IOMeter was run again using the OLTP workload to cause approximately 5% in data change. Then another snapshot was taken. This process was repeated a second and third time. For each of the iterations executed, the amount of storage provisioned to retain the snapshots was compared.

The EqualLogic solution was compared with both 100% snapshot reservation and also with a 200% reservation. This was required due to the manner in which the Dell EqualLogic handles snapshots by default. When snapshot reservation is exhausted, the default action with the PS6000XV is to remove the oldest snapshot. The Lionbridge Test Engineer observed that with 100% reservation, which is the default for the PS6000XV series, not all snapshot iterations could be retained following the above methodology.

Therefore, it was necessary to run a second test for the PS6000XV solution with a higher reservation amount so that all snapshots could be retained. This auto-deletion setting led to a situation in which a snapshot that was intended to be retained, was auto-deleted. It is important for an end user to understand the repercussions of this setting in order to make sure that required snapshots are not deleted inadvertently by the system.

Remote Replication Efficiency

This test case compared the amount of storage provisioned by each platform for remote replication retention. Also compared was the process of configuring and initiating remote replication.

For this test case, the test engineer created a 100GB volume, presented it to the respective VM as a second HDD as a raw device, and then filled it to 100% capacity using the IOMeter benchmark utility. For this data fill operation, a sequential write workload was configured for IOMeter with the following settings:

- Outstanding IOs: 8
- 100% Sequential Write
- 64K Request size
- 64K Alignment

The solution was configured with a single node, RAID5, no replication. A second cluster was created, containing a single node, also configured as RAID5, no replication. An initial remote copy operation was executed after the volume had been fully written. Then a 5% data change was introduced using the IOMeter OLTP workload. This was repeated a second and third time, each time a new remote copy was taken. The total amount of provisioned storage on the remote node was recorded.

The EqualLogic solution was configured with a single (source) node group in RAID5. A second (target) group was configured with a single node in RAID5, and the replication partnership was configured for one-way replication from the source to the target group.

A thin provisioned 100GB volume that was filled to 50% capacity was created. Then, an IOMeter OLTP workload was executed to cause 5% change in data. These changes were replicated to the replication target for each platform and the amount of provisioned space was compared.

Platform Refresh

This test case compared the process of refreshing the existing storage solutions to newer product offerings. As time goes by, storage capacity must increase to meet ever changing business requirements. It is imperative that any storage solution provide a means to update the underlying hardware without taking data offline. The process to accomplish the node upgrade and the amount of time required was compared.

For this test case, the P4500 cluster was configured with a two-node solution, RAID5, and no replication. This was then “upgraded” by adding two new nodes of equal capacity to the cluster and removing the “old” nodes.

The EqualLogic cluster was configured with a single node with RAID5. A second node was added to the group and then after the raid verification process completed, the “old” node was vacated.

Test results

This section provides all the results for the high availability, capacity utilization and feature comparison along with usability as experienced by the Lionbridge Test Engineer using the HP P4500 SAN and Dell PS6000XV storage arrays. Please refer to the testing methodology section of this report for complete details on how we conducted these tests.

High Availability Comparison

This section covers the results for test cases simulating several common hardware failure scenarios. Lionbridge Test Engineers executed these hardware failure tests and observed the effect the failure had

on the overall availability of the data. This section is broken into several subsections that focus on the specific hardware failure simulated.

It is important to note that the HP P4500 solution provides the same level of High Availability in a 2-node solution as it does with the 3-node solution using 2-way data replication. However, due to the less expensive acquisition cost, we were able to add a third node to the configuration that results in greater usable capacity, at a lower cost than the Dell EqualLogic 2-node solution. By adding a third node to the configuration, the HP solution meets all of the HA, capacity and feature requirements identified for ACME Corp.

High Availability Test Case Results

Table 2 below provides a summary of the results for each of the HA tests conducted by Lionbridge. Please refer to each test section below for detailed results.

HA Test	HP P4500 SAS		Dell PS6000XV	
	Result	Downtime Cost	Result	Downtime Cost
NIC Failure	Volumes Online	\$0	Volumes Offline	\$10K
System Failure	Volumes Online	\$0	Volumes Offline	\$13K
Double-disk Failure	Volumes Online	\$0	Volumes Offline	\$30K

Table 2: High Availability Test Case Results - Summary

NIC Failure

For this test case, a test engineer pulled all four cables at the same time from the primary controller on one of the EqualLogic arrays. Next, both network cables were pulled from one of the HP P4500 nodes to simulate the same type of failure on that platform. The test engineer observed the behavior of the VM and the IOMeter application. Also observed was the “Group Monitor” on the EqualLogic Group Admin GUI, as well as the “Performance Monitor” on the CMC. Each management application was monitored for any errors reported. After five (5) minutes the network cables were reconnected and the arrays were allowed to return to optimal operation.

The Lionbridge Test Engineer observed the following behavior for the HP P4500 SAN solution:

- The CMC GUI reported the effected node as offline
- 2-way replication went into a “degraded” state
- VMware ESX host did not report any issues
- The Virtual Machine continued to run without issue
- The IOMeter session running in the VM continued to execute without error

The end result of this test case was that under a NIC failure scenario on the P4500 SAN no outage was experienced. Data remained online and accessible during the entire time the node was in the failed state.

The Lionbridge Test Engineer observed the following behavior for the Dell PS6000XV solution:

- The Group Admin GUI reported the affected node as offline
- All striped volumes went offline
- The VMware ESX host reported connectivity problems to the VM images and IOMeter DataStors
- The Virtual Machine hung, and was unresponsive
- The Group Admin “Group Monitor” did not show any activity, and iSCSI sessions from the ESX host were disconnected.

The end result of this test case was that under a NIC failure scenario on the EqualLogic PS6000XV the data was offline and an outage had occurred.

After the cables were re-connected to the NICs, the Group Admin GUI reported the volumes back online within several seconds of the reconnect. After the cables were re-connected to the NICs, an HBA rescan on the ESX host had to be run in order to re-establish iSCSI connectivity and to bring the VM back online.

This total process led to a “downtime” of approximately 10 minutes, including the 5 minutes the cables were left unplugged.

Experiencing such a downtime can lead to significant SLA penalties for ACME Corp. The downtime experienced above would equate to approximately \$10,000 in SLA penalties with the Dell PS6000XV, but none with the HP P4500.

Figure 2 below shows the downtime cost associated with the above example when a NIC failure occurred.

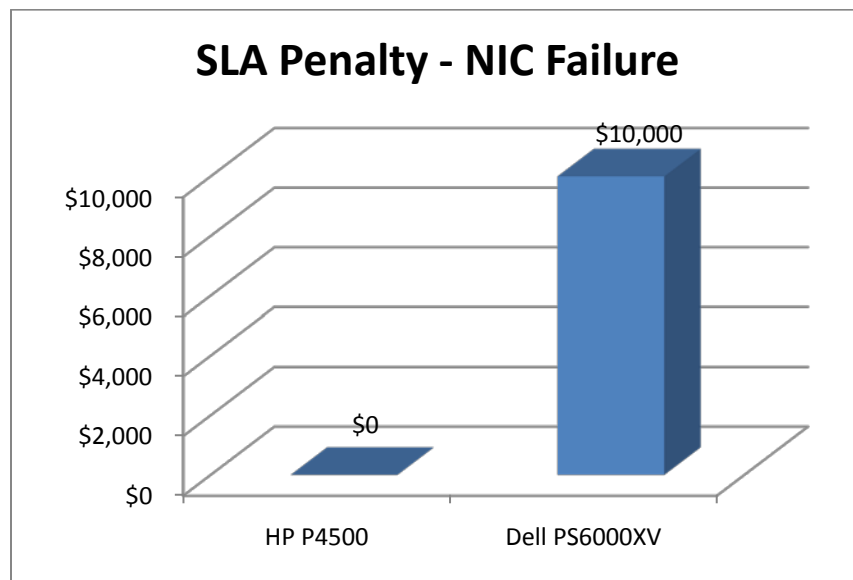


Figure 2: NIC Failure Outage – Cost in SLA Penalties

System Failure

For this test case, a test engineer simulated a system failure by pulling the power cable(s) from one of the nodes from each vendor. The test engineer observed the behavior of the VM and the IOMeter application. Also observed was the “Group Monitor” on the EqualLogic Group Admin GUI, as well as the “Performance Monitor” on the CMC. Each management application was monitored for any errors reported. After five (5) minutes the power cables were reconnected and the arrays were allowed to return to optimal operation.

The Lionbridge Test Engineer observed the following behavior for the HP P4500 SAN solution:

- The CMC GUI reported the affected node as offline
- 2-way replication went into a “degraded” state
- The VMware ESX host did not report any issues
- The Virtual Machine continued to run without issue
- The IOMeter session running in the VM continued to execute without error

The end result of this test case was that under a single node failure on the P4500 SAN, no outage occurred. The data remained online and accessible during the entire time the node was in the failed state.

The Lionbridge Test Engineer observed the following behavior for the Dell PS6000XV solution:

- The Group Admin GUI connection was lost and an automatic re-connection was initiated. It took approximately 70 seconds for the reconnection to complete.
- The node that was failed showed as offline in the Group Admin GUI
- All striped volumes, which included the VM image and IOMeter data LUN, were listed as offline
- I/O had stopped according to the Group Admin “Group Monitor” link

- The VM hung and new connections were unavailable.
- The IOMeter application appeared to be running but closer examination revealed it was running from memory and no actual iSCSI connectivity had continued.

The end result of this test case was that under a single node failure scenario on the EqualLogic PS6000XV the data was offline and an outage had occurred.

Once power was restored to the node, the volumes showed as “online” within one minute. However, the test engineer had to initiate an iSCSI HBA rescan in vCenter in order for the ESX host to re-establish iSCSI connectivity and for the VM to become online again. The total downtime experienced was approximately 13 minutes, including the 5 minutes the power was removed from the storage node.

Experiencing such an outage associated with a Node failure can lead to significant SLA penalties. The observed node failure outage would result in approximately \$13,000 in SLA penalties with the Dell PS6000XV solution, and no penalties with the HP P4500 solution.

Figure 3 below shows the downtime cost associated with the above example when a node failure occurred.

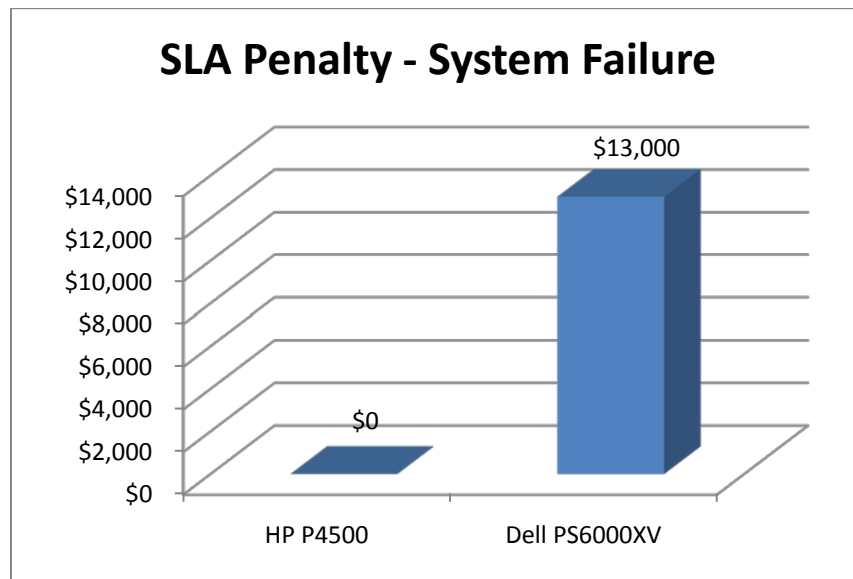


Figure 3: System Failure Outage – Cost in SLA Penalties

Double-disk Failure

For this test case, a test engineer simulated a double-disk failure by pulling two disks from the same node at the same time. Two odd numbered disks were pulled from one of the P4500 SAN nodes, and two odd numbered disks were pulled from one of the EqualLogic PS6000XV arrays. The test engineer observed the behavior of the VM and the IOMeter application. Also observed was the “Group Monitor” on the EqualLogic Group Admin GUI, as well as the “Performance Monitor” on the CMC. Each management application was monitored for any errors reported. After five (5) minutes the disks were re-inserted and the arrays were allowed to return to optimal operation.

The Lionbridge Test Engineer observed the following behavior for the HP P4500 SAN solution:

- The CMC GUI alerted that the node had disk failures
- The CMC GUI required a connection reset
- I/O paused for 1-2 seconds from IOMeter, then resumed
- No errors reported in the VM, nor IOMeter application
- After the disks were re-inserted, the node which experienced the failed disks remained listed as being in a “failed” state by the CMC GUI.

- I/O continued to be served and VM continued to function normally
- The test engineer power-cycled the “failed” node and when it returned it was reported as a normal state by the CMC GUI

Although more effort was required to bring the P4500 SAN node back to a “normal” state than in the other HA test cases, it is important to note that I/O continued to be served and there was no “downtime” detected by the VM or the IOMeter application.

The Lionbridge Test Engineer observed the following behavior for the Dell PS6000XV solution:

- EqualLogic Group Admin GUI alerted that the node suffered disk failures
- The Group Admin GUI became unresponsive for several second.
- All volumes went offline, which included the VM image and IOMeter data LUN
- I/O From IOMeter ceased
- The VM hung
- After the drives were re-inserted, the EQL Group Admin GUI reported the array had returned to a normal, fully functional state. However, the VM remained unresponsive.
- When attempting to rescan the LUNs from vCenter, the VM was still hung after 15 minutes post disk-return. The rescan hung indefinitely until the point of timeout.
- Ultimately the test engineer had to reboot the EqualLogic node that had experienced the disk failures, as well as the ESX host to clear out the hung iSCSI connections and re-establish a connection.
- This total downtime was measured at approximately 30 minutes

Although the VM did not appear to suffer any corruption due to the outage, it should be noted that in the event of a double-disk failure, it is expected that there would have been data loss on the IOMeter volume as data was being actively written at the time of failure.

In addition to an exposure to potential data loss, which can have tremendous cost implications, experiencing such an outage associated with a double-disk failure can lead to SLA penalties for ACME Corp. The observed node failure outage would result in approximately \$30,000 in SLA penalties with the Dell PS6000XV solution, and no penalties with the HP P4500 solution.

Figure 4 below shows the downtime cost associated with the above example when a double-disk failure occurred.

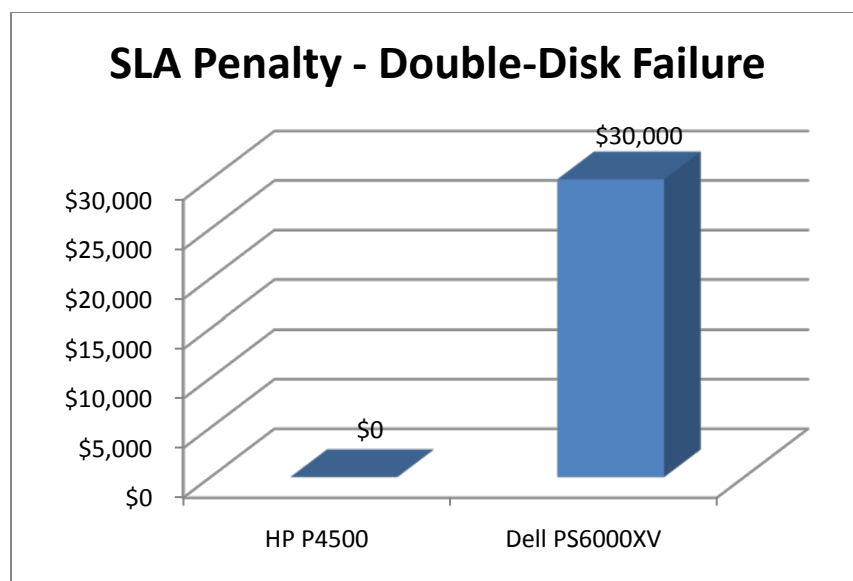


Figure 4: Double-Disk Failure Outage – Cost in SLA Penalties

VMware Fault Tolerance Scenario

This test case was developed to show a specific “real world” example of how the HP P4500 Virtualization SAN SAN’s capabilities address HA considerations, specifically with regard to its usage in a VMware installation. The P4500 Virtualization SAN has been designed to provide seamless integration with the VMware vSphere 4 Fault Tolerance feature.

For this test case, a test engineer simulated a node failure and an ESX host failure at the same time. To accomplish this, the test engineer removed the power cables from one of the ESX hosts as well as one of the nodes from each of the storage arrays. Testing was conducted serially. Power was pulled from the ESX host first to allow the Fault Tolerance failover to occur. As this was running, the power was then pulled from one of the storage nodes under test. VMware FT was configured so that one ESX host was running the primary VM for the HP P4500 SAN array and the other ESX host was running the primary VM for the Dell PS6000XV array. Prior to initiating the failure, the IOMeter OLTP script was executed on the VM to generate workload at the time of the failure.

For the HP P4500 SAN solution, the Lionbridge Test Engineer observed that when the ESX host was failed, the VMware FT process initiated as expected. Moments later, one of the P4500 SAN nodes was failed by removing power. However, the Virtual Machine remained online and fully functional. In addition, the IOMeter OLTP workload that was generating active I/O to the data volume also shared from the P4500 SAN cluster continued to function without issue. This result revealed that the P4500 SAN solution was able to provide the necessary HA capabilities to work seamlessly with the VMware vSphere 4 FT capabilities.

Conversely, with the Dell PS6000XV solution, the Lionbridge Test Engineer observed that when the ESX host was failed, the VMware FT process initiated as expected. However, when one of the EqualLogic PS6000XV nodes was failed by removing power, the VMware FT process failed because the VM became unavailable due to the VM volume located on the EqualLogic PS6000XV cluster going offline. In addition, the IOMeter OLTP workload ceased functioning as its volume had gone offline at the same time. After power was restored to the ESX host and the EqualLogic PS6000XV node, the iSCSI connection had to be rescanned from vCenter to reestablish the connectivity to the EqualLogic PS6000XV volumes. Then the VM was able to be powered back online.

Overall Capacity Utilization

For this test case, the test engineer simply configured each storage array in the various RAID types offered by the platform to determine the available capacity for each type.

For RAID5, the HP P4500 SAN solution provides 10% greater capacity than the Dell PS6000XV. When configured in RAID6, the HP solution provides for 2% greater capacity than the Dell PS6000XV. For RAID10 configurations, the HP P4500 SAN provides 32% greater usable capacity than the Dell PS6000XV.

Figure 5 below shows the pre-provisioned storage capacity available for each platform for several RAID configuration types. For RAID5, the HP P4500 provides 10% greater usable capacity than the Dell PS6000XV. When configured in RAID6, the HP P4500 provides 2% greater usable capacity. Finally, with a RAID10 configuration, the HP P4500 provides 32% greater usable capacity.

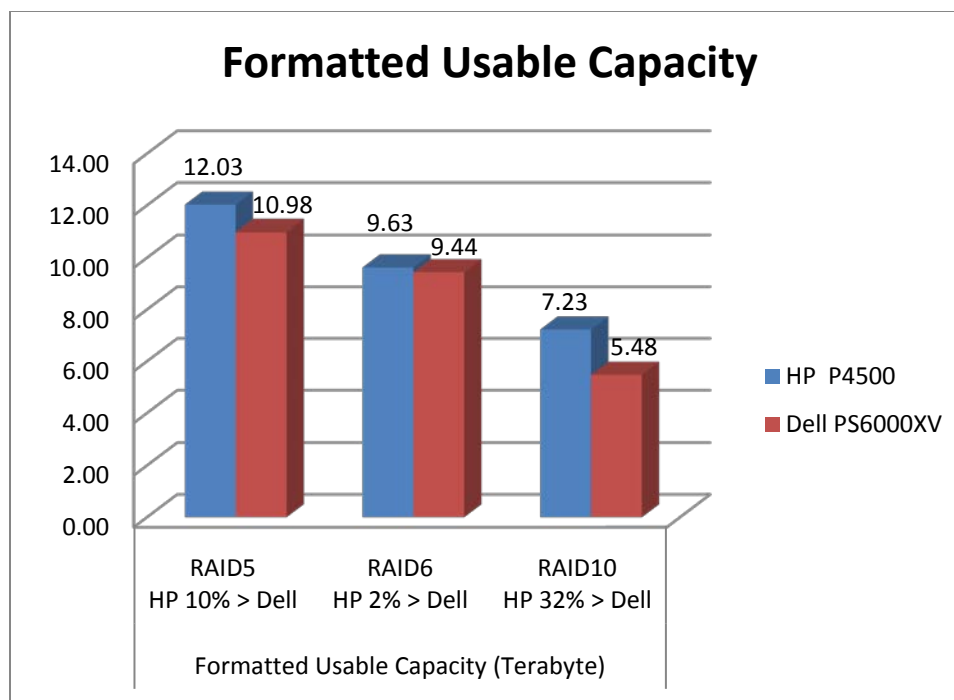


Figure 5: Formatted Usable Capacity Comparisons

Table 3 below represents the raw amount of usable storage available with various RAID configurations. The capacity listed is usable capacity after the RAID type is configured, but prior to any storage provisioning is allocated. As shown in Table 1, the HP P4500 SAN provides greater usable capacity at each of the RAID configurations available to both platforms.

Platform	RAID5	RAID6	RAID10
HP P4500 SAS	12.03 TB	9.63 TB	7.23 TB
Dell PS6000XV	10.98 TB	9.44 TB	5.48 TB

Table 3: Pre-provisioning Capacity at Available RAID Configurations

Table below represents the default space reservation policies for snapshot, data replication and thin provisioning.

Platform	Snapshots	Data Replication	Thin Provisioning
HP P4500 SAS	0%	100% (2-way)	1% (min)
Dell PS6000XV	100%	N/A	10% (min)

Table 4: Default Space Reservation for Snapshots, Replication and Thin Provisioning

With a minimum two node solution, the HP P4500 solution provides for data replication within its cluster. With the three node solution, the Lionbridge Test Engineer was able to configure 2-way replication on the HP solution, and still meet the minimum capacity requirements for ACME Corp. This means that data is written to two of the three nodes, thereby allowing the HP P4500 solution to handle single node failures, single network failures, and double-disk failures as shown above. This is not a feature of the Dell PS6000XV solution, thus, not tested. The two nodes would need to be split into two groups and a point-in-time replication partnership configured. However, this would not allow for enough capacity.

It is also interesting to note that the Lionbridge Test Engineer was only able to re-configure the RAID5 policy on the EqualLogic PS6000XV by removing the node from the Group and re-initializing it. The EqualLogic solution allows RAID configuration changes from one type to another type with certain restrictions. Figure 6:6 below is a screenshot which shows these policy restrictions.

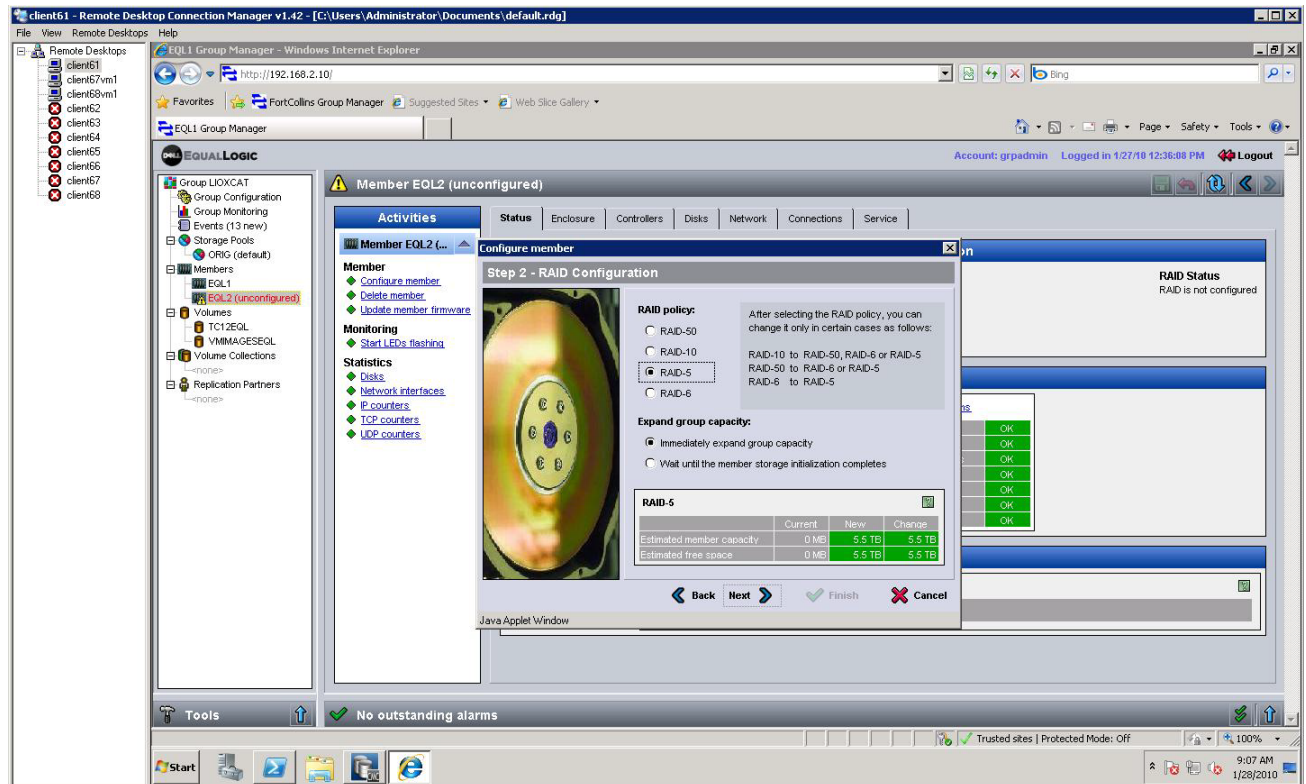


Figure 6: RAID Re-configuration Policies for the Dell Equallogic PS6000XV

Capacity Utilization Efficiency with Snapshots

This test case compared the amount of storage provisioned by each platform for snapshot retention. Also compared was the process of taking snapshots and the observed results.

One of the requirements for ACME Corp is that snapshots be taken to provide for readily available online backups of their data changes. Ideally, daily snapshots should be taken with the previous three (3) days of snapshots retained. On average, ACME Corp experiences approximately 5-6% data change per day. The Lionbridge Test Engineer executed a series of snapshots to measure how much capacity was required to maintain the required snapshots.

It is important to understand the basic difference between how each platform allocates snapshot space. The HP platform uses Thin Provisioning¹ to allocate snapshot space. The Dell EqualLogic solution uses a snapshot reservation policy in which a pre-defined amount of space, based on a percentage relative to the source volume. In addition, the Dell EqualLogic solution allocates capacity in 16MB blocks as opposed to the HP solution which writes data in 256KB blocks. This difference leads to a situation in which a large number of small writes (e.g. OLTP) can cause the provisioned space required on the Dell EqualLogic platform to grow at a much faster rate in order to retain the same amount of data changes.

¹ Thin Provisioning, in a shared storage environment, is a method for optimizing utilization of available storage. It relies on on-demand allocation of blocks of data versus the traditional method of allocating all the blocks up front (http://en.wikipedia.org/wiki/Thin_provisioning)

The Lionbridge Test Engineer observed that when the initial snapshot was taken on the HP platform, a new volume was created and labeled as a snapshot volume. The size of the snapshot volume was equal to the amount of data contained on the source volume. When the volume was initially snapped, this size was 100GB. The source volume then showed as 0% utilized. When data was subsequently written to the source volume, the utilization increased as it was filled with data. When subsequent snapshots are taken, the sizes of those snapshots were roughly equal to the amount of actual data that was written.

When taking snapshots on the Dell EqualLogic platform, the Lionbridge Test Engineer observed that when the initial snapshot was taken, no additional space utilization was reported beyond the source volume size and the reserved snapshot space. When subsequent data was written to the source volume, the amount of "Free Reserve" available in the snapshot reservation volume was reduced. The rate of this reduction occurred at a much higher rate than the amount of data that was actually written to the source volume.

The EqualLogic solution was compared with both a 100% snapshot reservation as well as a 200% reservation. The Lionbridge Test Engineer observed that with 100% reservation, which is the default for the PS6000XV series, not all snapshot iterations could be retained following the methodology below. Therefore, it was necessary to run a second test for the PS6000XV solution with a higher reservation amount so that all snapshots could be retained. This auto-deletion setting led to a situation in which a snapshot that was intended to be retained, was auto-deleted. It is important for an end user to understand the repercussions of this setting in order to make sure that required snapshots are not deleted inadvertently by the system.

This test case showed that the HP P4500 SAN platform provisioned less storage than the Dell PS6000XV when retaining the snapshot data. The following steps were executed to compare the process and provisioned space required for each platform:

HP P4500 SAN:

1. Created a 100GB LUN and presented it to the VM using the default snapshot reservation (thin)
2. Executed an IOMeter 64K Sequential Write to fill the volume to 100%
3. Took a base snapshot after the volume was filled
4. Ran an IOMeter OLTP workload until approximately 6% of the data had changed (specifically 5.85 GB) via write operations
5. Took a snapshot
6. Recorded the amount of snapshot space utilized
7. Ran an IOMeter OLTP workload until approximately 6% of the data had changed (specifically 5.85 GB) via write operations.
8. Took a snapshot
9. Recorded the amount of snapshot space utilized

Dell PS6000XV (Default snapshot reserve used):

1. Created a 100GB LUN and presented it to the VM using the default snapshot reservation of 100%
2. Ran an IOMeter 64K Sequential Write operation to fill the volume to 100%
3. Took a base snapshot after the volume was filled
4. Ran an IOMeter OLTP workload until approximately 5% data changed (approximately 5 GB) via write operations
5. Took a snapshot
6. Recorded the amount of snapshot space utilized
 - a. The first snapshot was auto-deleted per the default snapshot retention policy of the Dell EqualLogic platform.
7. Ran an IOMeter OLTP workload until approximately 5% data changed (approximately 5 GB) via write operations
8. Took a snapshot
9. Recorded the amount of snapshot space utilized.

Dell PS6000XV (Snapshot reserve set to 200% to ensure that all snapshots could be retained):

1. Deleted the previous snapshots
2. Set the snapshot reservation to 200% (200GB)
3. Ran an IOMeter 64K Sequential Write operation to fill the volume to 100%
4. Took a base snapshot after the volume was filled.
5. Ran an IOMeter OLTP workload until 5.5% of the data had changed (5.51 GB) via write operations
6. Took a snapshot
7. Recorded the amount of snapshot space utilized
8. Ran an IOMeter OLTP workload until 5.5% the data had changed (6.07 GB) via write operations
9. Took a snapshot
10. Recorded the amount of snapshot space utilized

As shown above, the first time the methodology was executed with default snapshot reservation settings the Dell PS6000XV was unable to retain all snapshot iterations. This caused a re-run of the test with a modified methodology that ensured that all snapshots could be retained. Due to thin provisioning of the snapshot reservation, the HP P4500 SAN platform did not have this issue.

When factoring in the initial data retention of the 100 GB fully filled LUN, the HP P4500 had an overall snapshot overhead retention rate of 1% (111.7 GB written vs. 112.6 GB provisioned), whereas the Dell PS6000XV had an overall snapshot provision overhead of 39% (111.58 GB written vs. 182.01 GB provisioned). In addition, by reserving 200% (200 GB), but only needing to provision 182 GB, the Dell solution resulted with approximately 18 GB (9%) of over-provisioned snapshot space.

Table 5 below presents a comparison of the snapshot provisioning observed for each platform.

Platform	Initial Snapshot Reservation (GB)	Snap1 Space Provisioned (GB)	Snap2 Space Provisioned (GB)	Snap3 Space Provisioned (GB)	Total Snap Space Provisioned (GB)	Total Snapshot Provisioned Overhead (%)
HP P4500	0	100	6.3	6.3	112.6	2%
Dell PS6000XV	200	.01	88.11	93.9	182.01	39%

Table 5: Snapshot Provisioning Comparison

Overall, testing revealed the Dell PS6000XV solution required 62% more provisioned space (182.01 GB vs 112.6 GB) to maintain the same number of snapshots as the HP P4500 required. This was observed at a data change rate of 5-6%.

It is also interesting to look at how much snapshot space is reserved when compared to the raw amount of data that was written between snapshot iterations. After the initial base snapshot was taken, both platforms were utilizing approximately the same amount of source volume and snapshot space. However, as subsequent data is written, and snapshots taken, the amount of space provisioned for retaining the snapshots, compared to how much data was written, is considerably different between the two platforms.

For example, for snapshot2, the raw data changed on the HP platform was 5.85 GB, with an allocated snapshot space of 6.3 GB; this equals a 7% overhead. This ratio is identical for the snapshot3 iteration. Another 5.85 GB of raw data change resulted in an additional 6.3 GB of allocated snapshot space

With the Dell EqualLogic platform, focusing on the second run which included enough reservation to retain all snapshots, for snapshot2, the raw data changed was 5.51 GB, with utilized snapshot space of 88.11 GB; this equals an overhead of 1500%. This ratio is repeated for snapshot3 iteration where 6.07

GB of raw data change occurred, and an additional 93.9 GB of utilized snapshot space. This equals an overhead of 1447%

Remote Replication Efficiency

This test case compared the amount of storage provisioned by each platform for remote replication retention. Also compared was the configuring and initiating of remote replication.

One of the requirements for ACME Corp is that data at their primary site be replicated to their Disaster Recovery (“DR”) site in real-time. Both the HP P4500 SAN and Dell PS6000XV solutions provide this capability. However, with the two-node Dell PS6000XV solution, both nodes have to be used in a striped RAID5 configuration to meet the minimum space requirements. This is also the situation with the HP P4500 SAN solution. For testing purposes, the capacity requirement was ignored in order to compare the differences between how remote replication functions for each platform.

Testing revealed a similar situation with remote replication utilization as was observed with snapshot utilization. The Dell EqualLogic platform required more provisioned space in order to retain the necessary remote replication copies than the HP platform. However, the overhead of provisioned space was not as high under remote replication utilization as it was for snapshot utilization.

Table 6 below presents the results observed for each platform.

	Dell PS6000XV		HP P4500 SAS	
	Provisioned Space	Data Copied	Provisioned Space	Data Copied
Full Snap 1	61.68 GB	49.83 GB	49.99 GB	49.91 GB
Snap 2 (5% △)	62.27 GB	3.42 GB	2.5 GB	2.28 GB
Snap 3 (5% △)	62.27 GB	2.92 GB	2.5 GB	2.35 GB
Snap 4 (5% △)	62.27 GB	2.83 GB	2.2 GB	2.06 GB
Total Space Allocated on Remote Site (Overhead)	62.27 GB	59.0 GB (5.25% Overhead)	57.19 GB	56.6 GB (1.03% Overhead)

Table 6: Remote Replication Provisioning & Utilization

As shown in the table above, the overhead associated with remote replication is lower for the HP platform than it is for the Dell EqualLogic platform. With a total of 59 GB of data written, the amount of provisioned space was 62.27 GB or 5.25% overhead for the Dell EqualLogic platform. With the HP platform a total of 56.6 GB of data was written, with the amount of provisioned space was 57.19 GB, or an overhead of 1.03%.

In addition, as was experienced with the Snapshot Capacity Utilization scenario compared above, the default action of the Dell PS6000XV was to remove the oldest volume replica when replication reservation space was diminished. This led to a situation in which the first volume replica was deleted on the target array. It is important for end users to understand the affect this setting has on day-to-day management of remote replication data retention and verify sufficient replication reserve has been allocated.

Platform Refresh

For this test case the HP P4500 cluster was configured with a two-node solution, RAID5, and no replication. This was then “refreshed” by adding two new nodes/arrays of equal size to the cluster and removing the “old” nodes/arrays.

The following steps were executed to complete the upgrade process on the HP platform:

1. Launch the CMC GUI
2. Use the CMC GUI discovery wizard to identify the two new nodes
3. Add the two “new” nodes to the management group containing the cluster to be updated

4. Delete the Virtual Manager.
5. Start a manager on one of the new nodes
6. Right-click on the cluster and select "edit cluster"
7. Click on "Add nodes" button
8. Select both nodes from the list. Click OK
9. Highlight the two original nodes from the list. Click on the "Remove Nodes" button. Nodes should be removed from list showing only the two new nodes. Click OK.
10. Restripe operation begins
11. When restripe completes, stop the managers on the old nodes and remove them from the management group.
12. Start a manager on the second "new" node
13. Create and start a virtual manager

Please note that the last two steps were required to satisfy CMC GUI recommended actions to meet Best Practices. However, data was fully available prior to executing these steps.

The Dell PS6000XV solution was configured with a single group, containing a single node. Then a new node with the same capacity was added to the group, and data was re-striped across the two nodes. Then the old node was vacated from the group. The following steps were executed to complete the upgrade process on the Dell PS6000XV platform:

1. Launch the EqualLogic Remote Setup wizard
2. Select "Initialize PS Series Array
3. The Discovery process runs and takes about 5 seconds. The "new" uninitialized node is discovered. Select the array and click next.
4. Enter the Member Name, IP Address, Subnet and Gateway. Select to "Join existing Group".
5. Enter the group information to join
6. The array is joined to the group. A confirmation window appears indicating success.
7. Click No when asked to start the GUI
8. Click Finish
9. Login to the EqualLogic Group Manager and click on the new member. This will launch the configuration GUI.
10. Select the Storage Pool (Default)
11. Configure RAID settings (RAID5)
12. Click on the Member and click on the Network tab.
13. Set the IP address and Netmask and enable the NIC for each of the three (3) remaining interfaces.
14. Check the Member status to confirm the RAID verify process is running.
15. When the RAID verify process completes, it is time to "vacate" the "old" node.
16. Right-click on the node to vacate under the "Member" category in the left-hand pane.
17. Select "delete member"
18. The node vacate process begins

The Lionbridge Test Engineer observed that in addition to more steps being required to complete a platform refresh, it also took considerably longer to complete the process. Whereas the HP P4500 platform completes the data re-stripe process to the new node(s) at once, the Dell PS6000XV platform was not able to do this. The new array must be added first, requiring data to be striped across the two arrays. Then, when the old array is vacated, the data must be re-written a second time to only the new array. Along with being much more time consuming, this process exposes the data to greater risk as the longer it takes to re-stripe the data the more at risk the data becomes if an outage of some sort were to occur during the process.

Table 7 below presents the number of steps and the amount of time it took to complete the platform refresh process for each platform.

Platform	Number of Steps	Time Required (Minutes)
HP P4500 SAN	13	34
Dell PS6000XV	18	339

Table 7: Number of Steps and Time Required Completing System Upgrade

Conclusion

Testing revealed that the HP P4500 SAN solution provides greater usable storage capacity while also providing better High Availability features than the Dell EqualLogic two node PS6000XV solution at a 14% lower price point.

Testing also revealed that the Dell PS6000XV solution required as much as 62% more provisioned space to retain the same amount of snapshot data as the HP P4500 SAN solution.

The HP P4500 SAS SAN solution requires five fewer steps and only 1/10 the time required to perform a platform refresh as compared to that of the Dell PS6000XV two array solution.

Finally, the High Availability features of the HP P4500 Virtualization SAN complement the VMware vSphere 4 Fault Tolerance feature thereby providing a better HA solution for ACME Corp.

The HP P4500 Virtualization SAN Virtual SAN solution better meets the requirements of a typical midsize company than the Dell PS6000XV solution.

Appendix

A. Hardware & Software Disclosures

- Storage
 - 3 x HP P4500 SAS SAN
 - SAN/iQ v8.1.00.0047.0
 - Six (6) total iSCSI connections (two per node)
 - 2 x Dell PS6000XV
 - Firmware 4.1.3
 - Eight (8) total iSCSI connections (4 active/4 passive per controller)
- Server
 - HP BL460c G6 Blade Server
 - Enclosure: C3000
 - Blades: 8
 - 2 x Flex-10 Interconnects
- Network
 - HP ProCurve 3400cl
 - Ports: 48
 - Speed: 1000Mb
 - VLANs: 2
 - 1 x Dedicated for iSCSI
 - 1 x Dedicated for Internet

VeriTest (www.veritest.com), the testing service of Lionbridge Technologies, Inc., provides outsourced testing solutions that maximize revenue and reduce costs for our clients. For companies who use high-tech products as well as those who produce them, smoothly functioning technology is essential to business success. VeriTest helps our clients identify and correct technology problems in their products and in their line of business applications by providing the widest range of testing services available.

VeriTest created the suite of industry-standard benchmark software that includes WebBench, NetBench, Winstone, and WinBench. We've distributed over 20 million copies of these tools, which are in use at every one of the 2001 Fortune 100 companies. Our Internet Benchmark service provides the definitive ratings for Internet Service Providers in the US, Canada, and the UK.

Under our former names of ZD Labs and eTesting Labs, and as part of VeriTest since July of 2002, we have delivered rigorous, objective, independent testing and analysis for over a decade. With the most knowledgeable staff in the business, testing facilities around the world, and almost 1,600 dedicated network PCs, VeriTest offers our clients the expertise and equipment necessary to meet all their testing needs.

For more information email us at info@veritest.com or call us at 919-380-2800.

Disclaimer of Warranties; Limitation of Liability:

VERITEST HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, VERITEST SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT VERITEST, ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL VERITEST BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL VERITEST'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH VERITEST'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.