Principled Technologies®

TEST REPORT
APRIL 2010

Virtualized Exchange 2010 performance: Comparing the Dell EqualLogic PS4000XV iSCSI SAN, HP LeftHand P4500 SAS Virtualization SAN Solution, and a Fibre Channel SAN

Executive summary

Dell Inc. (Dell) commissioned Principled Technologies (PT) to compare the performance of the following mid-range storage area network (SAN) solutions using a heavily virtualized Exchange 2010 workload:

- Dell™ EqualLogic™ PS4000XV 14.4TB SAS iSCSI SAN array (Dell EqualLogic SAN)
- HP LeftHand P4500 SAS Virtualization SAN Solution (HP LeftHand SAN)
- FC SAN-2009 Fibre Channel Array (Fibre Channel SAN)

We used the primary storage performance tests in the Microsoft® Exchange Server 2010 Jetstress tool for our comparison. Jetstress reports the input/output operations per second (IOPS) that a given Microsoft Exchange 2010 storage solution provides while maintaining satisfactory responsiveness. Our goal was to determine which SAN solution provided better performance value for transactional applications supporting business operations.

KEY FINDINGS

- The Dell EqualLogic PS4000XV iSCSI SAN array delivered the best performance value on our virtualized Exchange workload tests: 118% greater total IOPS per active disk than the HP LeftHand SAN and 29% greater total IOPS per active disk than the Fibre Channel SAN. (See Figure 1.)
- The Dell EqualLogic PS4000XV iSCSI SAN delivered 181% more useable MB per dollar than the HP LeftHand SAN and 69% more useable MB per dollar than the Fibre Channel SAN. (See Figure 2.)
- The Dell EqualLogic PS4000XV iSCSI SAN delivered 110% greater total IOPS per dollar than the HP LeftHand SAN and 127% greater total IOPS per dollar than the Fibre Channel SAN. (See Figure 3.)

We followed each vendor's documented best practices for keeping important data highly available. We used hot spares with the Dell EqualLogic SAN and the Fibre Channel SAN. The hot spares enable automatic recovery and allow for pre-failure copy-out that can avoid performance effects from a failing drive. For the HP LeftHand SAN, we did not use hot spares, so all disks were active for maximum performance but there would not be proactive copy-out of a failing disk. The Fibre Channel solution had 24 disks active for data, with additional disks reserved for other internal storage processes and one hot spare. The Dell EqualLogic solution had 28 active disks with four hot spares. The HP solution had 36 active disks with no hot spares.

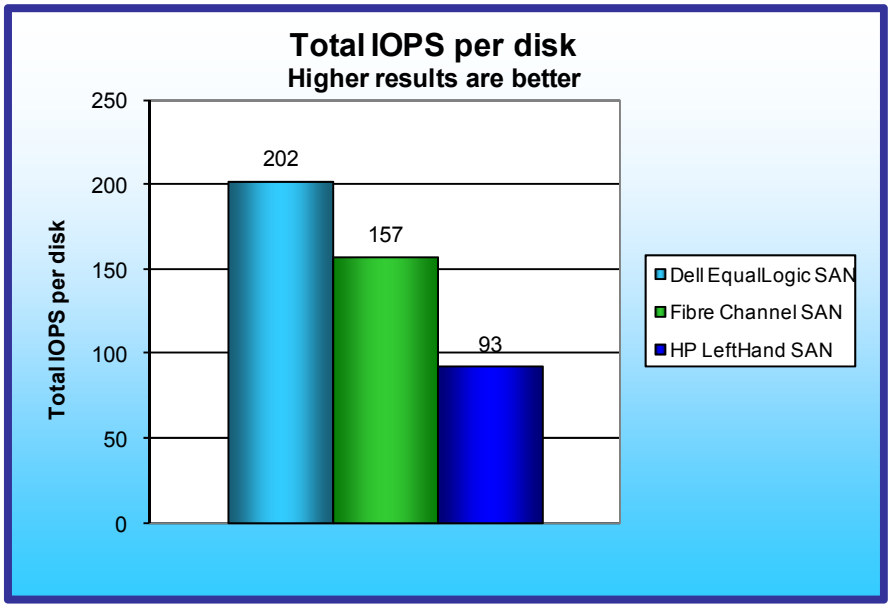

Figure 1: Total IOPS-per-disk results for the storage arrays during the Jetstress test. We calculate results by dividing total IOPS for all virtual machines by number of active disks.

We configured each SAN to use a comparable proportion of its usable capacity so that seek

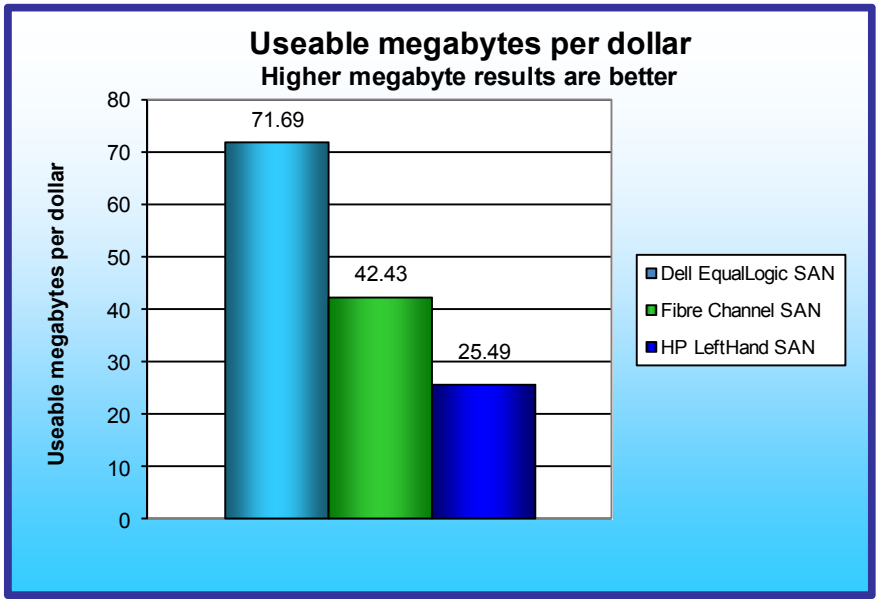

Figure 2: Useable megabytes-per-dollar results for the storage arrays during the Jetstress test. Results are useable megabytes divided by the list price of the array.

distances on the disks would be comparable. We ran the tests three times to ensure repeatability and we report the results from the run that produced the median total IOPS.

The Dell EqualLogic SAN achieved 5,662 total IOPS using 28 active disks, for 202 IOPS per disk, whereas the HP LeftHand SAN achieved 3,334 total IOPS using 36 active disks, for 93 IOPS per disk, and the Fibre Channel SAN achieved 3,771 total IOPS using 24 active disks, for 157 IOPS per disk. As Figure 1 shows, the Dell EqualLogic SAN achieved 118 percent more IOPS per disk than the HP LeftHand SAN, and 29 percent more IOPS per disk than the Fibre Channel SAN. In a larger environment, the IOPS per disk for each of the total physical disks can be close to the IOPS per disk for each of the active disks. For instance, a continuously staffed data center would not require hot spares and, in the case of the Fibre Channel SAN, such as data center would not need to reserve additional disks for other purposes as it added more capacity to the controllers.

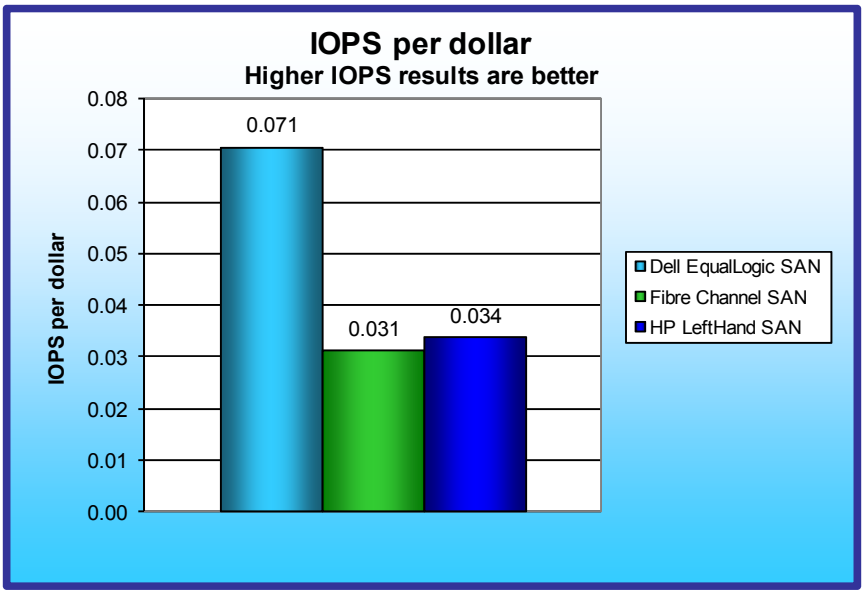

Figure 3: Total IOPS-per-dollar results for the storage arrays during the Jetstress test. Results are total IOPS divided by the list price of the array.

Figure 2 shows the useable megabytes per dollar for the three storage configurations. The Dell EqualLogic SAN provided 5,608 GB of useable space at a list price of $80,098, and delivered 71.69 useable MB per dollar. This is 181 percent more than the HP LeftHand SAN, which supplied 2,464 GB of useable space at a list price of $99,000, and delivered 25.49 useable MB per dollar; and is 69 percent greater than the Fibre Channel SAN, which supplied 5,020 GB of useable space at a list price of $121,141, and delivered 42.43 useable MB per dollar.[1]

Figure 3 shows the total IOPS per dollar for the three storage configurations. The Dell EqualLogic SAN delivered 0.071 total IOPS

[1] We obtained list prices for the HP LeftHand SAN, the Fibre Channel SAN, and the Dell EqualLogic SAN configurations from resellers. Prices include installation.

per dollar; 110 percent more than the HP LeftHand SAN, which delivered 0.034 total IOPS per dollar, and 127 percent more than the Fibre Channel SAN, which delivered 0.031 total IOPS per dollar.

Configuration for the Dell EqualLogic PS4000XV

We used the recommended configuration and configured the two Dell EqualLogic PS4000XV arrays into one storage pool using the default settings. We created all the LUNs for the database and logging inside this storage pool.

Configuration for the HP LeftHand P4500 SAN

We acquired the HP LeftHand P4500 SAN prior to the newer G2 systems becoming publicly available. While we did not test with a G2 system, we can make the following observations:

- During our Exchange Jetstress 2010 tests, the G1 CPU utilization remained under 40 percent. Because of this, it may be reasonable to expect that the G2 systems would perform similarly to the G1 systems in the same tests.
- During our Exchange Jetstress 2010 tests, the G1 systems only used one third of the bandwidth provided by the six 1Gb connections. Based on this, it may be reasonable to assume that using 10Gb NICs would not significantly increase performance in the same tests.

We also chose to configure our HP LeftHand storage with three-way replication because that was the only mode the user manual recommended for keeping online data that is also important to preserve. The HP user manual for version 8.0 firmware indicated that if our data was either critical or mission critical, we should configure with a priority setting of "Redundancy." Using the Redundancy priority setting, had we made only two copies, the data would have been taken offline whenever any one motherboard was offline. Had we selected the priority setting of “Availability” instead of the recommended Redundancy setting, we could have run two-way replication and still kept the volumes online when any one motherboard went offline. However, when configuring a SAN for important applications, one should follow the user manual’s best practices. We did not want to jeopardize our data by having it kept online in situations where the user manual states that critical data volumes should be offline.

We also considered that companies might choose to use two-way replication with a priority of Availability for economic reasons. Alternatively, they might select the Availability setting due to a lack of understanding of best practices. In the end, we chose three-way replication with a priority of Redundancy because if companies value their data's integrity and availability, they should configure their storage in accordance with the user manual's best practices.

Configuration for the Fibre Channel SAN

In the case of the Fibre Channel SAN, we found conflicting best practices documents regarding how many Exchange log disks we should configure, with one document recommending 2.5 times as many disks as the other document. We then tested the arrays using all three different possible configurations.

Both best practice documents for Exchange avoided putting Exchange data or logs on disks that have special functions in the storage besides storing user data. Nonetheless, we also ran a third test and put the logs on those disks to ensure fairness.

Configuration 1: Typical best practice configuration

The configurations in the Exchange best practice document on the support site used a ratio of approximately one log disk to four database disks. With this configuration, there was one hot spare, 5 disks for software functions not holding data, 20 disks for data, and 4 disks for logs. This configuration yielded 3,008 IOPS on 24 active disks, 125 IOPS per active disk, and 100 IOPS per physical disk.

Configuration 2 (conflicting best practice configuration): Maximized IOPS/disk dedicated to Exchange in an array that might also host other applications

Publicly available ESRP reports use a configuration with only about 1 log disk to 10 database disks. This proved to be a better but much more complicated configuration.

A single RAID 10 set of 22 disks was not an option because of the limitations of the array software. Breaking the RAID up into two sets of 11 disks was not an option because RAID 10 requires even numbers of disks. We needed one set of 10 disks and one set of 12 disks. The best practice guide recommended starting with even smaller building blocks because for expansion, one would need to expand using the same size building block one started with. Therefore, we built what is in effect a pool of 12 disks out of three sets of 4 disks each by striping the data across multiple smaller LUNs.

Having LUNs spread across either 10 disks or 12 disks, we then divided the users so that we tested 10/22nds of the users in one RAID group and 12/22nds in the other. This configuration yielded the 3,771 IOPS on 24 active disks, 157 IOPS/disk.

Configuration 3 (not a documented best practice): Maximized IOPS/dollar for an Exchange-only configuration

Other best practice documents indicated that we could put some data with non-trival IO load on the disks that the array would otherwise reserve for software purposes, but that we should not put as much IO load on those disks as we did on other disks. Therefore, we performed a third test for absolute maximum performance with no hot spares, 24 database disks, and with the logs on the remaining 6 disks that the array also used for internal software purposes. Depending on a company's needs, this configuration may be appropriate.

This configuration yielded the highest IOPS, 4,310, on 30 active disks, 144 IOPS per disk.

Figure 4 shows the performance comparison of the three configurations.

	Fibre Channel SAN configuration 1	Fibre Channel SAN configuration 2	Fibre Channel SAN configuration 3
Test configurations			
Number of active disks tested	24	24	30
Total IOPS	3,008	3,771	4,310
Total IOPS/active disk	125	157	144

Figure 4: The performance comparison of the Fibre Channel SAN configurations.

The second Fibre Channel configuration had the highest total IOPS per active disk in our tests. We therefore chose to use this configuration for our testing.

Workload

The purpose of the Jetstress primary storage performance tests is to exercise the storage arrays using the maximum sustainable Exchange 2010 type of I/O. The tests produce results that show how long it takes the storage under load to respond to an I/O request. The tests do not have a single workload; instead, testers create a simulated Exchange mailbox profile that defines the Jetstress workload. We calculated values for three Jetstress test parameters that define the simulated Exchange mailbox profile: target database IOPS per user, average mailbox size, and number of mailboxes.

Note: These results are not comparable to Exchange 2007 or Jetstress 2007. Microsoft has significantly modified the Exchange 2010 I/O workload from Exchange 2007. They have increased the size of each database I/O operation and changed to a more random and even less sequential I/O pattern.

- **Target database IOPS per user.** We based our tested user profile for this report on the Microsoft profile for heavy Exchange 2010 cached mode users. These users average 20 sent and 80 received messages per day. This load averages 0.32 database IOPS for each user. We used this

average with 20 percent headroom as our targeted database IOPS per user for a total of 0.4 IOPS target.

- **Mailbox size.** We used an average mailbox size that would create databases equal to 60 percent of the total formatted capacity of each array.
- **Number of mailboxes.** We used 6,000 users total (1,500 per virtual machine [VM]) on every array to create a consistent configuration across all the arrays. We then manually increased the thread count to increase the IOPS achieved by each array to indicate performance.

The overall Jetstress test result is a pass/fail rating based on whether the system's performance fell within acceptable latency thresholds. All the results we report are from runs that passed.

Appendix D provides further details about the Jetstress tests.

Test procedures

We first ran the Jetstress disk subsystem test with automatic tuning to identify a thread count, and then ran the initial performance test using that thread count. If the system passed the initial performance test run, we continued to retest using higher thread counts to push IOPS to the maximum, stopping when Jetstress reported a failing user experience rating. If the system failed the initial performance test run, we retested using lower thread counts until Jetstress reported a passing user experience rating. If the test failed using just a single thread, we reduced the number of mailboxes until Jetstress gave a passing user experience rating. This process identified the highest IOPS score that the system could achieve while receiving a passing user experience rating. We then performed two additional runs using the settings from the run that produced those results. We ran each test for 2 hours, the default run time for a Jetstress 2010 test. Jetstress 2010 reports results in IOPS.

We gathered the following results from the Jetstress 2010 report:

- Database disk reads per second
- Database disk writes per second
- Achieved IOPS (sum of the database disk reads and writes per second)
- Log writes per second (log IOPS)
- Total IOPS (sum of the achieved IOPS and the log IOPS)

The results we report in Figures 1 through 3 are from the run that produced the median total IOPS. See Appendix A for complete Jetstress 2010 test results.

Appendix A: Test results

Figure 5 provides test results for the storage arrays, as well as information about the simulated Exchange configuration that defines the test workload. We ran the Jetstress primary storage performance test three times for each SAN.

	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
Test workload			
Number of disks tested	28	36	24
Number of mailboxes	6,000	6,000	6,000
Total useable space (GB) (1GB=2^30 bytes)	5,608	2,464	5,020
Initial database size per VM (GB)	842	368	754
Initial database storage capacity utilization	60%	60%	60%
Total IOPS (higher numbers are better)			
VM 1 Achieved IOPS - Run 1	1,157.564	661.131	717.382
VM 1 Log IOPS - Run 1	105.173	99.844	124.500
VM 1 Total IOPS - Run 1	1,262.737	760.975	841.882
VM 2 Achieved IOPS - Run 1	1,238.330	757.809	800.244
VM 2 Log IOPS - Run 1	109.930	106.543	132.772
VM 2 Total IOPS - Run 1	1,348.260	864.352	933.016
VM 3 Achieved IOPS - Run 1	1,403.259	755.708	909.498
VM 3 Log IOPS - Run 1	115.194	107.417	145.139
VM 3 Total IOPS - Run 1	1,518.453	863.125	1,054.637
VM 4 Achieved IOPS - Run 1	1,278.119	748.305	892.710
VM 4 Log IOPS - Run 1	102.168	106.594	140.962
VM 4 Total IOPS - Run 1	1,380.287	854.899	1,033.672
Total IOPS – Run 1	**5,509.737**	**3,343.351**	**3,884.172**
VM 1 Achieved IOPS - Run 2	1,216.025	656.304	640.440
VM 1 Log IOPS - Run 2	102.937	99.449	116.743
VM 1 Total IOPS - Run 2	1,318.962	755.753	757.183
VM 2 Achieved IOPS - Run 2	1,349.440	746.856	770.972
VM 2 Log IOPS - Run 2	108.703	106.451	131.278
VM 2 Total IOPS - Run 2	1,458.143	853.307	902.250
VM 3 Achieved IOPS - Run 2	1,368.253	763.539	967.615
VM 3 Log IOPS - Run 2	107.221	108.861	149.363
VM 3 Total IOPS - Run 2	1,475.474	872.400	1,116.978
VM 4 Achieved IOPS - Run 2	1,306.046	744.890	855.950
VM 4 Log IOPS - Run 2	103.836	107.171	139.097
VM 4 Total IOPS - Run 2	1,409.882	852.061	995.047
Total IOPS – Run 2	**5,662.461**	**3,333.521**	**3,771.458**

	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
VM 1 Achieved IOPS - Run 3	1,245.751	640.440	605.136
VM 1 Log IOPS - Run 3	105.545	96.490	110.305
VM 1 Total IOPS - Run 3	1,351.296	736.930	715.441
VM 2 Achieved IOPS - Run 3	1,325.986	731.936	715.273
VM 2 Log IOPS - Run 3	106.231	102.861	124.485
VM 2 Total IOPS - Run 3	1,432.217	834.797	839.758
VM 3 Achieved IOPS - Run 3	1,383.371	750.516	945.826
VM 3 Log IOPS - Run 3	110.665	108.137	143.709
VM 3 Total IOPS - Run 3	1,494.036	858.653	1,089.535
VM 4 Achieved IOPS - Run 3	1,308.200	749.940	830.580
VM 4 Log IOPS - Run 3	107.932	106.917	135.042
VM 4 Total IOPS - Run 3	1,416.132	856.857	965.622
Total IOPS – Run 3	**5,693.681**	**3,287.237**	**3,610.356**
Additional results from median run			
Achieved IOPS	5,263.31	2,872.83	3,096.82
Database disk reads/second	3,590.18	1,864.84	2,217.47
Database disk writes/second	1,538.47	1,023.86	1,017.51
Log writes/second	430.37	414.41	513.54
Achieved IOPS per disk	188.00	79.80	106.80
Database disk reads/second per disk	128.22	51.80	76.46
Database disk writes/second per disk	54.95	28.44	35.09
Log writes/second per disk	15.40	11.50	17.70
Total IOPS per disk	202.20	92.60	130.10

Figure 5: Jetstress primary storage performance test results for the storage arrays.

Appendix B: Test environment

We created a test bed in a climate-controlled room for each storage system. Each test bed included the following components, which we connected with Cat6 Ethernet cables:

- Microsoft Exchange Server
  - One Dell PowerEdge™ R710 server installed with the following software:
    - vSphere 4.0 Update 1
  - Each VM had the following software installed on it:
    - Microsoft Windows Server® 2008 R2
    - Jetstress 14.00.0639.012
    - Exchange 2010 14.00.0639.019 ESE binaries
- Switch
  - One Dell PowerConnect™ 5448 switch
- Storage systems under test
  - Two Dell EqualLogic PS4000XV 14.4TB SAS iSCSI SAN arrays
  - HP LeftHand P4500 10.8TB SAS Virtualization SAN Solution with three full nodes (one additional node to use three-way replication for mission critical data)
  - One Fibre Channel Array with two trays of storage

Figure 6 provides highlights of the Dell PowerEdge R710 server configuration. See Figure 8 for complete configuration details.

Dell PowerEdge R710 rack server	
Processors	Two quad-core Intel® Xeon® Processor X5570s at 2.93 GHz
Memory	48 GB, 6 x 8 GB, 1,333 MHz
Internal disk	Two 73GB, RAID 1, 10K RPM Seagate ST973402SS SATA 3.0Gbps drives
Network	Intel PRO/1000 PT Quad Port Server Adapter
Operating system	vSphere 4.0 Update 1 w/ Windows Server 2008 R2
Test software	Jetstress 2010 Beta 14.00.0639.012 with Exchange 2010 14.00.0639.019 ESE binaries

Figure 6: Dell PowerEdge R710 server configuration.

Storage arrays under test

The storage arrays we tested each have fully redundant components. To obtain basic controller redundancy, HP requires two nodes and requires that the array copy each block of data to each node ("2-way replication"). The HP user manual states that, if data is critical, we should configure the volumes with a priority setting of "redundancy." We configured the volumes with three-way replication. The user manual also indicated that, with this setting, our data would go offline if two nodes went offline. Had we selected the priority setting of "availability," the data could remain online in case of two-node failure, but selecting this setting would contradict the user manual recommendation for mission critical data. In this three-way replicated configuration, critical data remains online if any one controller goes offline.

We used two Dell EqualLogic PS4000XV arrays to provide a similar price-point comparison to the three-node HP LeftHand SAN configuration and Fibre Channel SAN configuration.

We also used two trays of storage from a Fibre Channel SAN vendor to provide a similar price-point comparison to the Dell EqualLogic and HP LeftHand storage solutions. We configured all three SANs with 450GB 15K RPM drives.

Figure 7 summarizes the systems we tested, their disk counts, the amount of useable storage, and the RAID technology we used.

System components	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
Arrays	Dell EqualLogic PS4000XV SAS iSCSI SAN array	HP LeftHand P4500 SAS Virtualization SAN Solution	FC SAN-2009 Fibre Channel Array
Disks	32 450GB 15K RPM SAS disks total (28 active disks)	36 450GB 15K RPM SAS disks total for the three shelves (36 active disks)	30 450GB FC 15K RPM disks total for the two shelves (29 active disks)
Formatted storage capacity (1GB=2^30 bytes)	5,608 GB	2,464 GB	5,020 GB
Connection	iSCSI	iSCSI	FC
Multi-pathing	EqualLogic HIT (Host Integration Toolkit) 3.2 with DSM	SAN/iQ® 6.6 Virtual IP Load Balancing v8.1 with DSM	Yes
RAID technology	RAID 10	RAID 10+2 (RAID 10 in each array with three-way replication required for basic controller redundancy).	RAID 10

Figure 7: Storage system configuration information.

Appendix C provides additional details about the storage arrays we tested.

Microsoft Exchange Server

Figure 8 provides detailed configuration information for the server we used for testing.

Server	Dell PowerEdge R710
General dimension information	
Height (inches)	3.50
Width (inches)	17.50
Depth (inches)	27.00
U size in server rack (U)	2
Power supplies	
Total number	1
Wattage of each (W)	570
Cooling fans	
Total number	5
Dimensions (h x w) of each	2.50” x 2.50”
Voltage (V)	12
Amps (A)	1.60
General processor setup	
Number of processor packages	2
Number of cores per processor package	4
Number of hardware threads per core	2
CPU	
Vendor	Intel

Server	Dell PowerEdge R710
Name	Xeon X5570
Stepping	D0
Socket type	LGA1366
Core frequency (GHz)	2.93
L1 cache	4 x 32 KB + 32 KB
L2 cache	4 x 256 KB
L3 cache (MB)	8
Platform	
Vendor and model number	Dell PowerEdge R710
Motherboard model number	0M233H
Motherboard revision number	13
BIOS name and version	Dell 1.3.6 (12/14/2009)
BIOS settings	Default
Memory modules	
Total RAM in system (GB)	48
Vendor and model number	Samsung M393B1K70BH1-CH9
Type	PC3-10600R
Speed (MHz)	1,333
Speed in the system currently running @ (MHz)	1,333
Timing/latency (tCL-tRCD-iRP-tRASmin)	9-9-9-9
Size (GB)	8
Number of RAM modules	6
Chip organization	Double-sided
Hard disk	
Vendor and model number	Seagate ST973402SS
Number of disks in system	2
Size (GB)	73
Buffer size (MB)	16
RPM	10,000
Type	SAS
Network card/subsystem	
Vendor and model number	Intel PRO/1000 PT Quad Port
Type	Discrete

Figure 8: Detailed system configuration information for the server we used for testing.

Dell PowerConnect 5448 switch

The test bed included one Dell PowerConnect 5448 switch. Figure 9 presents details on the switch.

Component	Details
Switch	One Dell PowerConnect 5448 switch
Firmware	Dell PowerConnect 5448 v2.0.0.41
Frame size	Jumbo frames enabled on all connected devices. Size set to 9,014 bytes.
iSCSI optimization	iSCSI optimization disabled
Storm control	Broadcast Only on all ports

Figure 9: Dell PowerConnect 5448 switch configuration.

Individual test beds

The remainder of this appendix shows diagrams of the components in each of the test beds.

Dell EqualLogic SAN

The Dell EqualLogic PS4000XV array has two controllers, with two 1Gbps iSCSI ports each. One controller is active; the secondary controller on each array is in standby failover mode in case of controller failure. We used two iSCSI ports per controller to connect the arrays to the iSCSI switch. We configured one NIC on the R710 with the ESX iSCSI Initiator for the operating system volume. We configured the remaining three NIC's as pass through NIC's to each virtual machine and installed the EqualLogic HIT KIT with the Microsoft iSCSI Initiator on each virtual machine. Figure 10 shows the wiring diagram for the server, switch, and storage for the Dell EqualLogic SAN.

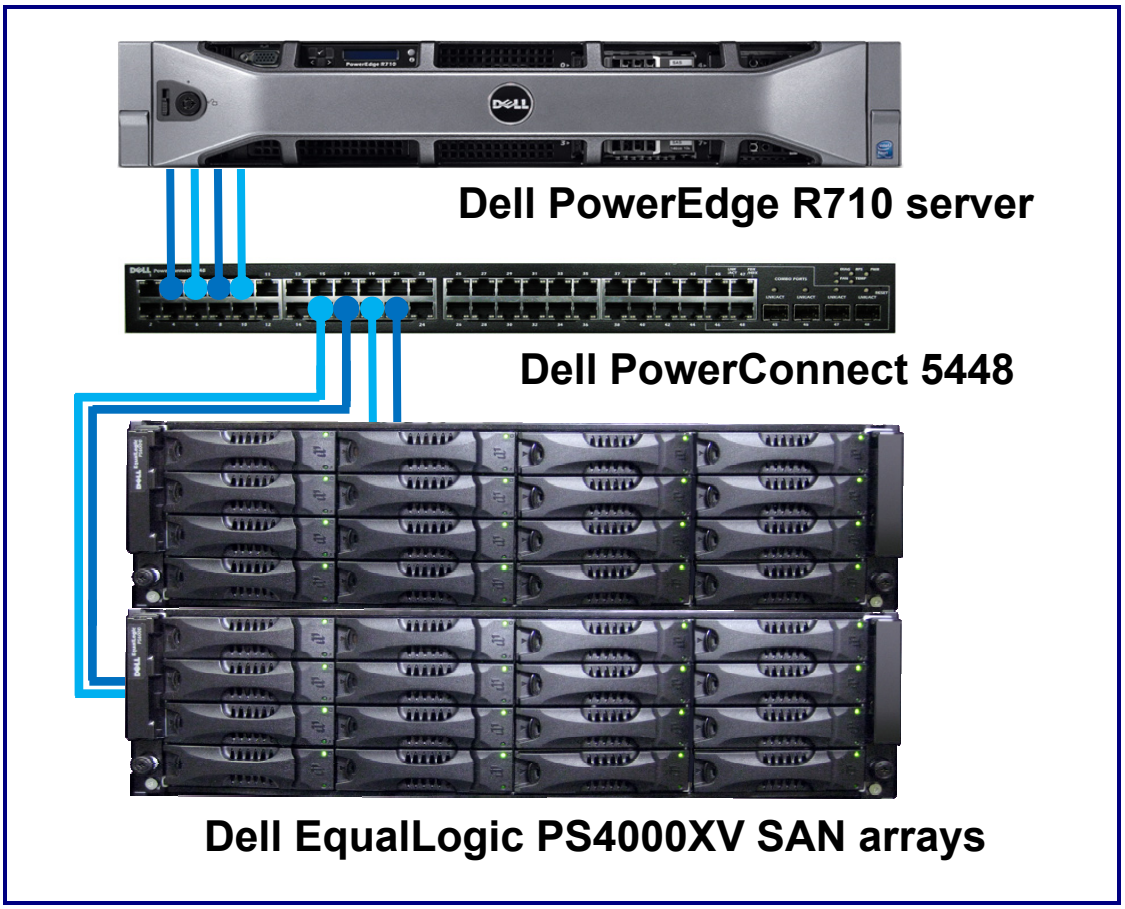

Figure 10: Dell EqualLogic SAN wiring diagram for the server, switch, and storage.

HP LeftHand SAN

The HP LeftHand P4500 Virtualization SAN Solution has one controller per node, with two 1Gbps iSCSI ports per controller. We used two iSCSI ports per controller to connect the arrays to the iSCSI switch. We configured one NIC on the R710 with the ESX iSCSI Initiator for the operating system volume. We configured the remaining three NIC's as pass through NIC's to each virtual machine and installed the HP SAN/iQ 6.6 Virtual IP Load Balancing v8.1 with DSM with the Microsoft iSCSI Initiator on each virtual machine. Figure 11 shows the wiring diagram for the HP LeftHand SAN.

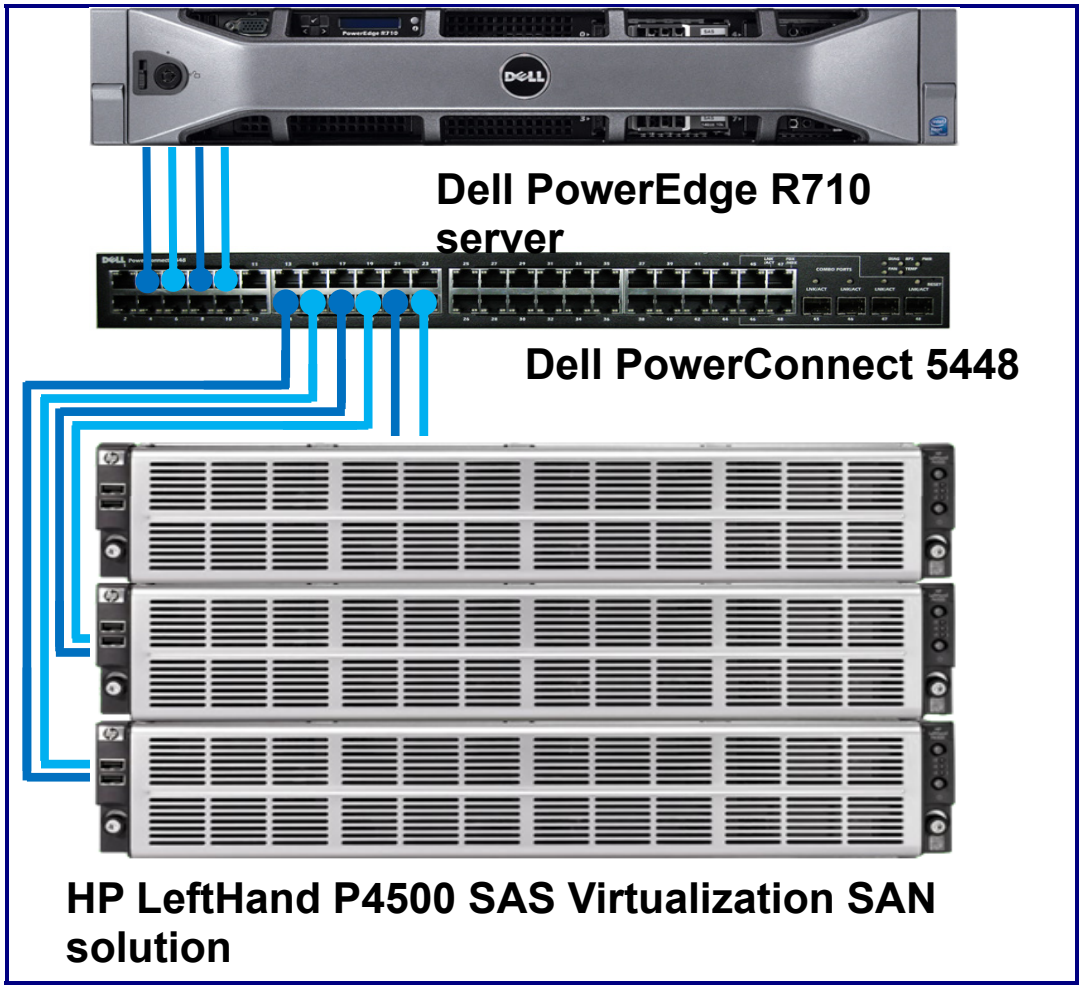

Figure 11: HP LeftHand SAN wiring diagram for the server, switch, and storage.

Fibre Channel SAN

The Fibre Channel SAN has two controllers per array, with two 8Gbps FC ports per controller. We used one 8Gb Fibre Channel port per controller to connect the array directly to the R710 Dual Port 8Gb Fibre Channel card.

Appendix C: Storage configuration and setup

This section includes the parameters we used for simulated Exchange configuration calculations, the results of those calculations for the tested arrays, and the steps we took to configure the arrays.

Simulated Exchange configuration calculations

We defined the simulated Exchange configuration using the following guidelines:

- Initial database size, equal to 60 percent of the total formatted capacity
- To allow for growth in the database files, we allocated 80 percent of useable disk capacity for the database LUNs and the rest to log LUNS. We used 100 percent of the available space in order to ensure that neither array was gaining an unfair advantage by not formatting the entire capacity. Note: In a production environment, installers should follow each vendor's best practices when configuring storage.
- One mailbox per user

We calculated the simulated Exchange configuration using the parameters in Figure 12. Figure 13 shows the simulated Exchange configuration, Figure 14 shows the primary storage hardware, Figure 15 shows the primary storage software on the host server, and Figure 16 shows the primary storage disk configuration.

Item	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
IOPS per mailbox	0.4 (0.32 is typical of a heavy user profile plus 20% headroom)	0.4 (0.32 is typical of a heavy user profile plus 20% headroom)	0.4 (0.32 is typical of a heavy user profile plus 20% headroom)
Database files capacity utilization percentage	60%	60%	60%
Database LUN capacity utilization percentage	80%	80%	80%
Mailbox size quota (MB)	575	251	515
Mode	Cached	Cached	Cached
Number of hosts	4	4	4

Figure 12: Simulated Exchange configuration parameters.

Simulated Exchange configuration

System	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
Number of Exchange mailboxes we simulated	6,000	6,000	6,000
Number of hosts	4	4	4
Mailbox size (MB)	575	251	515
Number of mailboxes/VM	1,500	1,500	1,500
Simulated profile: I/Os per second per mailbox	0.4 (0.32 plus 20% headroom)	0.4 (0.32 plus 20% headroom)	0.4 (0.32 plus 20% headroom)
Database LUN size (GB)	1,120	490	1,150
Log LUN size (GB)	230	70	104
Initial total database size for performance testing (GB)	3,368	1,472	3,016
Percentage formatted capacity used by Exchange database	60%	60%	60%

Figure 13: Simulated Exchange configuration.

Primary storage hardware

System	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
Total number of disks we tested in solution	32 (28 active)	36	30 (22 in use for database, 2 for logging, 5 for array OS, and 1 for logging)
Storage connectivity (Fibre Channel, SAS, SATA, iSCSI)	iSCSI	iSCSI	Fibre Channel
Storage model and OS/firmware revision	Dell EqualLogic PS4000XV Firmware v4.3.2	HP LeftHand P4500 Firmware 8.1.00.0047.0	-
Storage memory	4 GB (2 GB per controller)	4 GB (2 GB per controller)	-
Number of storage controllers	4 (2 active)	3	2
Number of storage ports	4 active	6 active (2 per controller)	2 active (1 per controller)

Figure 14: Primary storage hardware.

Primary storage software on host server

System	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
Server HBA/NIC driver	1.0.0.11 (vmxnet3)	1.0.0.11 (vmxnet3)	N/A
Multi-pathing	EqualLogic HIT (Host Integration Toolkit) 3.2	SAN/iQ 6.6 Virtual IP Load Balancing v8.1	N/A
Host OS	Windows Server 2008 x64 R2 SP2	Windows Server 2008 x64 R2 SP2	Windows Server 2008 x64 R2 SP2
Server ESE.dll file version	14.00.0639.019	14.00.0639.019	14.00.0639.019

Figure 15: Primary storage software on host server.

Primary storage disk configuration

System	Dell EqualLogic SAN	HP LeftHand SAN	Fibre Channel SAN
Disk type, speed, and firmware revision	32x Seagate SAS 15k, ST3450856SS, XQH6	36x Seagate SAS 15k, ST3450856SS, HPD5	30x FC 15K
Raw capacity per disk (GB)	450	450	450
Number of physical disks in test	32	36	24
Total raw storage capacity (GB)	14,400	16,200	10,800
Raid level	RAID 10	RAID 10+2 (RAID 10 in each node with three-way replication required for basic controller redundancy)	RAID 10
Total formatted capacity (GB) (1GB=2^30 bytes)	5,608	2,464	5,020
Database capacity utilization	60%	60%	60%

Figure 16: Primary storage disk configuration (Mailbox Store and transactional logs).

Setting up the Dell PowerEdge R710 Server

We installed VMware vSphere 4 update 1 on the host server.

Installing VMware vSphere 4 update 1

1. Insert the vSphere 4 update 1 DVD, and restart the computer.
2. While booting, press F11 to enter Boot Menu.
3. Press the down arrow to navigate to the appropriate boot device, and press Enter.
4. To start ESX in graphical mode, press Enter.
5. On the Welcome screen, click Next.
6. Click the checkbox to accept the terms of the license agreement, and click Next.
7. Accept the default keyboard layout, and click Next.
8. On the Custom Drivers screen, choose No to install custom drivers, and click Next.
9. You will receive a warning if you chose No. Click Yes to proceed. Drivers will load at this point.
10. Click Next.
11. Enter your serial number now, or choose to enter one later. Click Next.

12. On the Network Configuration screen, choose the NIC that you will use for system tasks. Click Next.
13. On the next Network Configuration screen, set your IP addressing information, subnet mask, DNS, and hostname. Click Next.
14. On the Setup Type screen, choose standard setup.
15. Select the virtual disk you wish to install ESX on, and click Next.
16. On the Data Loss warning screen, click OK.
17. On the Time Zone Settings screen, choose your appropriate time zone, and click Next.
18. On the Date and Time screen, modify the date and time as you wish, and click Next.
19. Assign a root password for ESX. Optionally, add additional users if you wish, and click Next.
20. On the Summary screen, click Next to install.
21. When the installation is complete, click Next.
22. Click Finish to reboot the system.

Enabling secure shell (ssh) access for the root user

1. Log in as root to the ESX console.
2. Type the following command to change to the appropriate directory:
   `# cd /etc/ssh`
3. Edit the sshd_config file using vi. Use the following command:
   `# vi sshd_config`
4. Press the down arrow key to move the cursor to the PermitRootLogin line, and move the cursor to the word no. Press the i key to insert text.
5. Type the word `yes` and delete the word no.
6. Press the Escape key to return to command mode.
7. Type the following command to save and exit the file:
   `:wq`
8. Type the following command to restart the ssh service:
   `# service sshd restart`

Installing vSphere on a management workstation

1. Configure your management workstation to be on the same subnet as your ESX server. In Internet Explorer, navigate to the IP address of your ESX Server.
2. Accept certificate warnings, and continue.
3. Click vSphere to download the vSphere client.
4. Choose Run to install.
5. Choose your language, and click OK.
6. On the Welcome screen, click Next.
7. Accept the license terms, and click Next.
8. Enter user and organization details, and click Next.
9. Optionally, click the checkbox to install the update utility. Click Next.
10. On the Destination Folder screen, click Next.
11. Click Install.
12. Click OK.

Connecting the storage to the host

We attached the database volumes on the Dell EqualLogic PS4000XV storage to the Dell PowerEdge R710 server by connecting both NICs available on each Dell EqualLogic P4000XV array to a Dell PowerConnect 5448 switch. We used three physical NICs from the Intel PRO/1000 PT dedicated to iSCSI traffic to connect the Microsoft iSCSI Initiator within each VM to their respective database volumes. We used the remaining NIC on the Intel PRO/1000 for the VMkernel iSCSI traffic.

Adding the iSCSI storage adapter in ESX

1. On the management workstation, click Start, All Programs, VMware, VMware vSphere client.
2. Click the Configuration tab.
3. Click Networking.

4. Click Add Networking…
5. Select the VMkernel radio button, and click Next.
6. Deselect any selected NICs that vSphere shows as down. Select the first of the available four NICs that vSphere does not show as down in the Speed column. Click Next.
7. Assign a label to the network. Click Next.
8. Enter a valid IP address and subnet. Click Next.
9. Click Finish.
10. Click No when the application prompts you to enter a default gateway.
11. Click the Configuration tab.
12. Click Storage Adapters under Hardware pane.
13. Click iSCSI Software Adapter, and click Properties.
14. Click Configure.
15. Click Enabled.
16. Click OK.
17. After the application enables the device, click the Static Discovery tab.
18. Click Add.
19. Enter the IP address of the storage group. Additionally, enter the initiator target name for the large LUN you will use for operating system virtual disks, which is accessible from the Dell EqualLogic Group Manager web application.
20. Click Close.
21. Click Yes when the application prompts you to rescan for changes.

Adding the OS storage LUN as a datastore to ESX storage

1. In vSphere, click the Configuration tab.
2. Click Storage in the Hardware pane.
3. Click Add Storage…
4. Choose the 200GB Disk/LUN. Click Next.
5. Select the appropriate LUN. Click Next.
6. Click Next.
7. Enter a name for the datastore. Click Next.
8. Click Next to accept the defaults for block size.
9. Click Finish to complete the addition of the LUN.

Setting up the iSCSI NICs with vSphere

1. Click the Configuration tab.
2. Click Networking.
3. Click Add Networking.
4. Click Virtual Machine.
5. Choose to Create a virtual switch. Assign the relevant NIC in the system. Click Next.
6. Assign a network name.
7. Click Finish.
8. Repeat steps 1 through 7 twice.

Creating the virtual machine with vSphere

1. Click Start, All Programs, VMware, VMware vSphere client.
2. Enter the IP address or hostname, user name, and password. Click Login.
3. Click the Virtual Machines tab.
4. Right-click, and choose New Virtual Machine.
5. Choose Custom, and click Next.
6. Assign a name to the virtual machine. Click Next.
7. Select a datastore for the virtual machine files. We chose the OS Datastore. Click Next.
8. Choose Virtual Machine Version. Click Next.
9. Choose Microsoft Windows, then Microsoft Windows Server 2008 R2 (64-bit). Click Next.
10. Choose four virtual processors. Click Next.
11. Choose 8192 RAM. Click Next.

12. Click None for the number of NICs. Click Next.
13. Choose LSI Logic SAS as the SCSI controller. Click Next.
14. Choose to create a new virtual disk. Click Next.
15. Make the OS virtual disk size 40 GB, and click Next.
16. Keep the default virtual device node, and click Next.
17. Click Finish.
18. Right-click the VM, and choose Edit Settings.
19. On the Hardware tab, click Add…
20. Choose Ethernet Adapter, and click Next.
21. Choose VMXNet 3, and click Next.
22. Repeat steps 19 through 21 twice more to add the iSCSI dedicated NICs.
23. Click Finish, and click OK.

Installing the guest operating system

1. Insert the installation DVD for Windows Server 2008 R2 into the DVD drive.
2. In vSphere, right-click the virtual machine, and choose CD/DVD drive.
3. Click the Host Device radio button, and choose the appropriate drive.
4. Click OK.
5. Right-click the machine, and choose Power, Power On.
6. Right-click the machine, and choose Open console.
7. Choose the language, time and currency, and keyboard input. Click Next.
8. Click Install Now.
9. Choose Windows Server Standard (Full Installation). Click Next.
10. Accept the license terms, and click Next.
11. Click Custom.
12. Click the Disk, and click Drive options (advanced).
13. Click New, Apply, Format, and click Next.
14. Let the installation process continue. The VM will reboot several times.
15. After the installation completes, click OK to set the Administrator password.
16. Enter the administrator password twice, and click OK.
17. When the operating systems finishes booting, choose VM, Guest, Install/Upgrade VMware Tools.
18. On the information Install VMware Tools screen, click OK.
19. On the autoplay menu, click Run setup.exe.
20. On the Welcome screen, click Next.
21. On the Setup Type screen, choose Typical, and click Next.
22. Click Install.
23. On the various Windows Security screens, click Install each time.
24. Click Finish to complete the installation.
25. Click Yes to reboot the VM.
26. Log into the VM after reboot.
27. Click Start, Control Panel, and double-click System.
28. Click Change Settings.
29. Click Change.
30. Enter the new computer name, and click OK.
31. Click OK to restart, click Close, and click Restart Now.

Installing Jetstress

We followed this process to install Jetstress on the test server:

1. Download the Microsoft Exchange Server Jetstress Tool (64-bit)
2. Run Jetstress.msi.
3. Click Next.
4. Accept the terms of the License Agreement, and click Next.
5. Click Next.
6. Click Next.
7. Click Close.

8. Copy the following files from an Exchange 2010 installation disk:
    - ese.dll
    - eseperf.dll
    - eseperf.hxx
    - eseperf.ini
9. Paste the files in C:\Program Files\Exchange Jetstress.
10. Run JetstressWin.exe.
11. Click Start new test.
12. Click Exit.

Configuring networking in the VM

1. Power on the VM, and open the console.
2. Click Start, Control Panel, Network Connections, and click Manage Network Connections.
3. Right-click the NIC, and choose properties.
4. Select TCP/IP (v4), and choose properties.
5. Set the IP address, subnet, gateway, and DNS server for this NIC, which will handle outgoing server traffic (e.g., the public IP). Click OK, and click Close.
6. Repeat steps 2 through 5 twice more, but these NICs should be on the same network segment as the Dell EqualLogic storage.

Dell EqualLogic SAN

1. Using the command line or web interface, setup a storage group with the following RAID policies:
    a. PS4000XV-1: RAID-10
    b. PS4000XV-2: RAID-10
2. Create a storage pool that you name database with the following members:
    a. PS4000XV-1
    b. PS4000XV-2
3. Enable all network connections on the PS4000's using the following IP address scheme:
    a. IP Address: 192.168.1.## (## being any number between 10 and 40)
    b. Subnet Mask: 255.255.255.0
4. Create one 200GB volume in the storage pool with no snapshot reserve, and name it `OS`.
5. Create four 1120GB volumes in the storage pool with no snapshot reserve, and name each of them `VM#-DB` (where `#` is the number of the VM).
6. Create four 230GB volumes in the storage pool with no snapshot reserve, and name each of them `VM#-Log` (where `#` is the number of the VM).
7. Enable shared access to the iSCSI target from multiple initiators on the volume.
8. Create an access control record for the volume without specifying any limitations.
9. Create a read-only SNMP community name to use for group monitoring.

HP LeftHand SAN

1. Assign a name and IP to each node using the following IP address scheme:
    a. IP Address: 192.168.1.## (## being any number between 10 and 40)
    b. Subnet Mask: 255.255.255.0
2. Open the HP LeftHand Centralized Management Console.
3. Create a Management Group named `grpmanager`.
4. Create a cluster named `grpcluster` with the following nodes and RAID policies:
    a. P4500-1: RAID 10
    b. P4500-2: RAID 10
    c. P4500-3: RAID 10
5. Assign a virtual IP Address of `192.168.1.101`.
6. Create one 200GB volume in the cluster with three-way replication, and name it `OS`.
7. Create four 490GB volumes in the cluster with three-way replication, and name each of them `VM#-DB` (where `#` is the number of the VM).

8. Create four 70GB volumes in the cluster with three-way replication, and name each of them `VM#-Log` (where # is the number of the VM).
9. Create a new server, name it `HP-Server`, set the Initiator Node Name to the INN of the ESX Server, and assign the OS volume to it.
10. Create four new servers, name each of them `HP#` (where # is the number of the VM), set the Initiator Node Name to the INN of the individual VM, and assign the OS volume to it.

Appendix D: Jetstress 2010 test procedures

Jetstress 2010 performance tests

Jetstress 2010 is a tool that Microsoft developed to help administrators verify the performance and stability of the disk subsystem, the subsystem Exchange stresses most, before putting their Exchange server into a production environment. Jetstress simulates the Exchange database and log file loads produced by a specific number of users. Jetstress tests both responsiveness and throughput, giving a pass/fail rating for responsiveness and reporting throughput in I/O operations per second. Jetstress uses read and write latency as its primary measure of responsiveness. The results also report these latency values.

The Jetstress reliability test measures the reliability of the storage array, primary storage performance, and streaming backup/recovery performance. We ran the primary storage performance tests only. The primary storage performance testing identifies the maximum sustainable Exchange I/O that the storage system can handle while providing acceptable responsiveness over a 2-hour period.

Setting up the Jetstress 2010 test

We set up our test environment and installed Microsoft Windows 2008 R2 Standard Edition on the test server. To get starting values for the first Jetstress performance test run, we ran the Jetstress disk subsystem throughput test and used its estimated thread count as the thread count for the first Jetstress performance test.

We ran the Jetstress performance test multiple times, and with different thread counts, on each storage solution. Our goal was to find the peak IOPS result for each solution. We define the peak result as the result of the run that had the highest IOPS but that also provided a passing Jetstress score for user experience.

We followed a different process based on whether the first run passed or failed. If the first run passed, we repeated the test with higher thread counts until the test failed because of unsatisfactory responsiveness. If the first run failed, we decreased the thread count until the run passed. We did so to find the highest transactional I/O load that would get a passing Jetstress score.

Preparing for the test

To prepare for the test, we first ran a Jetstress disk subsystem throughput test to find a starting point for tuning the mailbox profile.

1. Run JetstressWin.exe.
2. Click Start new Test, and click Next.
3. Select Create a new test configuration file.
4. Name the file `ArrayName_Initial_Run.xml`, and click Next.
5. Select Test disk subsystem throughput, and click Next.
6. Change the value to 60 in the Size the database storage capacity percentage box.
7. Leave the default value of 100 in the Target IOPS using throughput capacity percentage box, and click Next.
8. Select Performance, and click Next.
9. Set the test duration to 2 hours, and click Next.
10. Leave the number of databases set to 1, and click Next.
11. Select Create new databases.
12. Click Prepare Test.
13. Once you have prepared all the databases, click Execute Test on every VM.
14. Once the test has finished, record the thread count from the disk subsystem throughput. We use this value as the thread count for the first Jetstress test run below.

Running the test

Before testing, we ensured that all Ethernet ports were active and connected at 1 Gbps to the iSCSI switches for maximum performance and availability. We then followed this process for each test:

1. Run JetstressWin.exe.

2. Click Start new Test, and click Next.
3. Select Create a new test configuration file.
4. Name the file `ArrayName_Run#.xml` and click Next.
5. Select Test an Exchange mailbox profile.
6. Enter a description in the text box, and click Next.
7. Set the number of mailboxes to 6,000.
8. Type `.4` for IOPS/mailbox.
9. Set the Mailbox size to 575 or 251 MB depending on the array.
10. Check the Suppress tuning and use thread count (per-storage group) checkbox.
11. For the first run, enter the number of threads that the automatic tuning from the Disk subsystem throughput test chose. For subsequent runs, use the thread count we calculated in Step 19. Click Next.
12. Select the Performance test type, and click Next.
13. Set the test duration to 2 hours.
14. Set the number of databases to 1.
15. Click the white box under … for Group1 Database1, and enter E: for the database volume and F: for the log volume.
16. Select Create new databases, and click Next.
17. Repeat steps 1 through 16 on the three other VM's.
18. Click Execute Test on every VM.
19. Wait for the test to finish; then check the output files to see if the run reported errors or failed by exceeding the Jetstress latency thresholds for the log or database files.
    a. If the run had errors or the latency exceeded the thresholds, proceed based on whether the run prior to this run passed or also failed:
        i. If the previous run was successful, its results show the maximum IOPS score. Save those results as Run 1 results.
        ii. If the previous run also had errors or exceeded the latency thresholds, decrease the number of threads by 1, and perform another run.
    b. If the run succeeds, perform the following steps:
        i. Save from the following files:
            - The Jetstresswin configuration file used (e.g., ServerName_Run#.xml)
            - Performance_(TimeStamp).html
            - Performance_(TimeStamp).blg
            - DBChecksum_(TimeStamp).html
            - DBChecksum_(TimeStamp).blg
        ii. Record the items that we show in Figure 5: Jetstress primary storage performance test results for the storage arrays.
        iii. Check the results for the run prior to this run, and proceed based on whether that run passed or failed.
            1. If the previous run was successful, increase the number of threads by 1, and perform another run.
            2. If the previous run had errors or exceeded the latency thresholds, the current runs results show the maximum IOPS score. Save those results as the Run 1 results.
20. Repeat the test two more times using the thread count you used in the run that produced the results you saved as Run 1 results in Step 19.

About Principled Technologies

We provide industry-technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help you assess how it will fare against its competition, its performance, whether it's ready to go to market, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC 27703
www.principledtechnologies.com
info@principledtechnologies.com