

TOTAL COST COMPARISON IN SMALL TO MEDIUM BUSINESS SCENARIOS: VMWARE VSPHERE VS. MICROSOFT HYPER-V

Simpler management. Better value.

VMware provided greater operational efficiency for a lower overall cost compared to Microsoft.®

vmware®



Total cost of ownership (TCO) is the ultimate measure to compare IT infrastructure platforms, as it incorporates the purchase and support costs of the platform along with ongoing operational and management expenses. The operational efficiency built into your software stack can greatly affect your bottom line—once you have procured and implemented your platform, operational costs for administration and maintenance can easily balloon. A solution that streamlines and automates routine maintenance tasks can increase uptime and save an organization time and money. In this study, we use results from the VMware Cost-Per-Application calculator and examine the operational expenses of the two platforms using simulated workflow scenarios to provide a hypothetical TCO comparison.

In our labs at Principled Technologies, we compared the automated administration capabilities of two common virtualization platforms, VMware vSphere® 5 and Microsoft Windows Server® 2008 R2 SP1 Hyper-V™, in several scenarios. It took significantly less time to complete common administrative tasks in our scenarios with the VMware solution, potentially reducing post-acquisition operational expense costs in the five operational tasks we tested using vSphere Enterprise Plus Edition by as much as 91 percent over a two-year period compared to the Microsoft solution. Additionally, when we narrowed the scenarios to those using features available on vSphere Standard Edition, we found the operational savings in those scenarios could save up to 91 percent over the Microsoft solution.



When we combine the operational cost savings VMware vSphere provides with the capital expenses the VMware Cost-Per-Application calculator predicts, we find that VMware virtualization platforms can provide substantially lower two-year total cost of ownership compared to Microsoft platforms. We approximated operating expenses for two small and medium business scenarios with 100 VMs: a business using VMware vSphere Enterprise Plus Edition and a business using VMware vSphere Standard Edition. For vSphere Standard Edition cost calculations, we included only the scenarios with Standard edition-supported features.

Figure 1: In the 100-VM Enterprise Plus Edition scenarios we tested, VMware had 91 percent lower operational costs over a two-year period.





Figure 2: In the 100-VM scenario Standard Edition scenarios we tested, VMware had 91 percent lower operational costs over a two-year period.



SELECTING A COMPLETE SOLUTION

When choosing a virtualization platform, considering all costs—both acquisition-related and operational—is essential. An organization must account for not only the cost of acquiring licenses and software, but also the cost associated with the time a system administrator will devote to maintenance and management tasks within each environment. Since system administrator time can be more valuable when used on strategic IT initiatives that deliver a competitive edge for their organization instead of routine maintenance, it is always beneficial to reduce operational administrative costs. As these operational costs add up over time, they can become a significant portion of overall costs for a small or medium business. We discuss both acquisition and operational costs below. For acquisition estimates, we used the VMware Virtualization Cost-Per-Application Calculator on VMware’s Web site at <http://www.vmware.com/technology/whyvmware/calculator/>.

Acquisition costs

As verified by Principled Technologies’ 2011 testing,¹ VMware vSphere offers significant advantages that can lead to higher VM density than Microsoft Hyper-V. Higher VM density translates directly to reduced capital costs for a virtualization platform because the customer needs fewer hypervisor hosts and management servers to support a population of virtualized applications. The VMware Cost-Per-Application calculator factors the vSphere VM density advantage into comparisons with solutions based on Microsoft Hyper-V and System Center to show that at higher VM densities, VMware can provide acquisition costs lower than that of Microsoft for hardware, software, management components, data center space, and power and cooling.

Management and maintenance scenario summary

To test the management and maintenance functionality for each platform, we chose a number of representative operational tasks that a large organization running vSphere Enterprise Plus Edition would carry out throughout the course of a typical two-year period. These scenarios include the following:

- Shifting virtual machine workloads for host maintenance
- Adding new volumes and redistributing VM storage
- Isolating a storage-intensive “noisy neighbor” VM
- Provisioning new hosts
- Performing non-disruptive disaster recovery testing

Additionally, we calculated the costs and savings for businesses choosing to run vSphere Standard Edition. In this situation, we used only two operational tasks—Shifting


¹ <http://www.principledtechnologies.com/clients/reports/VMware/vsphere5density0811.pdf>

virtual machines workloads for host maintenance and Performing non-disruptive disaster recovery testing—as these features are available with Standard Edition.

After timing each scenario, we estimated how many times IT staff would complete each of these routine maintenance tasks during a typical two-year period. We calculated costs for a small to medium business with 100 VMs. The default output from the VMware Cost-Per-Application calculator assumes a density advantage of 50 percent more VMs for VMware over Microsoft, but we chose a more conservative estimate of 25 percent and used those VM densities as guidelines in our pricing estimates. Therefore, for acquisition cost purposes, we estimated 15 VMs per VMware vSphere server and 12 VMs per Microsoft Hyper-V server.

Using our density approximations, time estimates, and the number of iterations for each task, we then calculated person-hours and the cost of those person-hours using standard IT salary and benefits rates to determine the administrative savings an organization could realize using VMware vSphere. We performed this calculation two ways: in our five test scenarios for vSphere Enterprise Plus Edition and with two test scenarios for vSphere Standard Edition. In both cases, we used 100 as the target number of VMs. Using the representative tasks and scenarios we chose for 100 VMs, the VMware solution could save \$19,886 in management costs over a two-year period compared to a comparable solution from Microsoft when running vSphere Enterprise Plus Edition (See Figure 3). When running 100 VMs with vSphere Standard Edition and using the two chosen operational tasks, the savings would be \$16,867.

Figure 3: Using VMware products can lower your operational cost by as much as \$19,886 over the course of a two-year period compared to comparable Microsoft offerings (100-VM scenarios, Enterprise Plus Edition).



Shifting virtual machine workloads for host maintenance

Firmware upgrades, BIOS updates, and hardware replacements often require short periods of server downtime. To perform this routine maintenance, an administrator must first offload the virtual machines running on those servers to other servers to keep the infrastructure running. This time to migrate VMs from source to destination servers requires valuable hands-on time from the administrator; the faster these migrations happen, the better. Figure 4 depicts the live migration process.

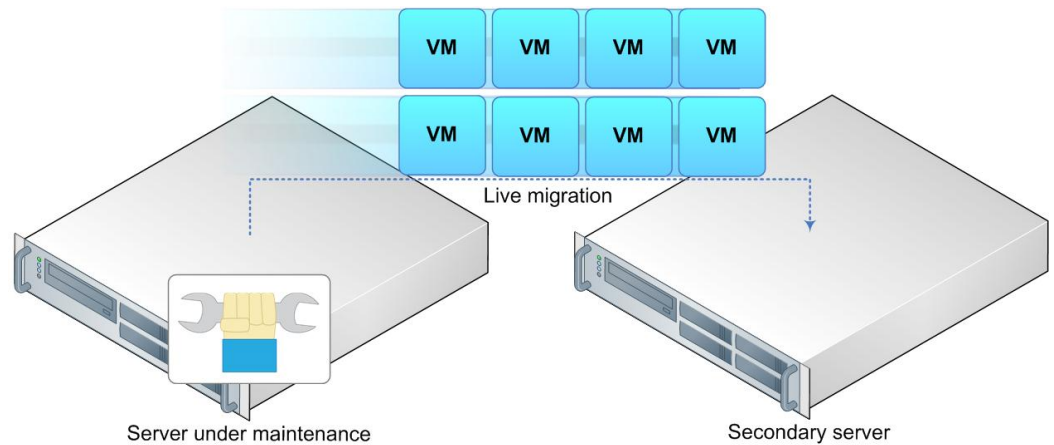


Figure 4: VM live migration time is critical during a server maintenance event.

To test this scenario for both VMware and Microsoft, we placed six VMs, each with 10 GB of RAM, on each server in our three-server cluster and ran a medium database workload on each of the 18 VMs. We then measured the time it took one server in the cluster to enter maintenance mode, evacuate all its VMs to the two remaining servers, and then migrate the VMs back to the original server. We performed these tests using both the VMware solution and the Microsoft solution. We found that the solution running VMware vSphere 5 reduced the time to complete the shifting of the VM workloads by 79 percent over the Microsoft solution. Figures 5 and 6 show the time it took to complete each task needed to perform physical maintenance on a server. We provide further details in [Appendix C](#). This feature is included in both vSphere Enterprise Plus and vSphere Standard editions, and therefore is included in both calculations in this report.

Time to shift VM workloads




Figure 5: It took 79 percent less time to shift the VM workloads using the VMware solution than it did with the Microsoft solution. Lower numbers are better.

Task	VMware solution	Microsoft solution
Time to fully migrate all VMs off one node and enter maintenance mode	01:06	07:56
Time to exit maintenance mode	00:01	00:14
Time to migrate VMs back	01:09	02:55
Total without boot	02:16	11:05

Figure 6: Times, in minutes:seconds, to complete the live migration relating to performing physical maintenance on one server.

Adding new volumes and redistributing VM storage

If your business is growing, the increasing numbers of VMs and data in your environment mean that you will need new storage. System admins must frequently add

new storage capacity, which requires them to redistribute existing VM storage to new storage or to re-provision existing storage. Using the available features in each platform, we timed how long it would take to redistribute VM storage after new storage capacity had been added into a cluster. The goal of storage expansion was to expand the overall cluster capacity and relieve preexisting datastores that were nearing capacity.

The features available to each platform differ slightly in this scenario. On VMware vSphere, we used VMware Storage Distributed Resource Scheduler (Storage DRS), a fully automated solution. Because an equivalent feature does not exist in the Microsoft platform, on Microsoft Hyper-V, we used a combination of manual decision-making by an administrator and System Center Virtual Machine Manager (SCVMM) to perform the Quick Storage Migration.

With VMware Storage DRS, the end user experiences no downtime (see Figure 7); therefore, we did not factor in any additional time to the scenario besides administrator UI data entry and confirmation times. With Microsoft SCVMM Quick Storage Migration, a brief “save state” occurs on the VM, causing downtime to the applications inside that VM. Therefore, we determined that for each of those VMs, additional administrator time was needed not only for the physical move of the VM files, but also for the inevitable coordination effort with application stakeholders and business users. This would be necessary to ensure that users were prepared for the downtime during the migration window.





Figure 7: VMware Storage DRS efficiently and automatically handles the addition of new storage tiers.

We discovered that performing this management operation took 95 percent less time with VMware as compared to Microsoft, due to VMware Storage DRS automation and the lack of downtime with the VMware solution. Figure 8 shows the time it took for each solution to migrate VM storage. Figure 9 and [Appendix D](#) provide a breakdown of each task we performed and the time required for completion. We did not measure the time necessary to implement the new tray of storage, as it was the same for both platforms. Nor did we measure the actual storage migration time as we assume administrators would let this operation run automatically. This feature is not included in vSphere Standard edition, and therefore is included only in the vSphere Enterprise Plus edition calculations in this report.

Figure 8: It took 95 percent less time to add a new volume and redistribute VM storage using the VMware solution than it did with the Microsoft solution. Lower numbers are better.



VMware solution		Microsoft solution	
Task	Time	Task	Time
		1. Plan for the brief but inevitable downtime with Quick Storage Migration. We assume 15 minutes of coordination time per VM, and a density of six VMs on the affected volume to be migrated.	1:30:00*
1. On a host, rescan the iSCSI Software adapter for the new LUN on the new storage tier.	0:02:10	2. On each host, connect to the new LUN using iSCSI initiator. We assume three hosts.	0:01:07
2. Add the new LUN as a datastore to the cluster.	0:01:40	3. Using disk management on one of the hosts, create a new simple volume using the new LUN.	0:00:36
		4. Bring the LUN online on each host.	0:02:12
3. Add the new datastore to the preexisting datastore cluster.	0:00:23	5. Using failover clustering services on the management server, add the disk to the	0:01:10

VMware solution		Microsoft solution	
		cluster and add it to cluster shared volumes.	
		6. Assess administrator time necessary to manually calculate how many migrations are necessary to balance LUN capacity using the new storage tier. We assume 1 minute per VM, and six VMs on the affected volume.	0:06:00*
4. Click “Run Storage DRS” to start the redistribution of the VMs using the new storage tier.	0:00:10	7. Using SCVMM and quick storage migration, queue each quick storage migration using the built-in wizard.	0:02:01
Total	0:04:23	Total	1:43:06

Figure 9: Times, in hours:minutes:seconds, to complete the tasks relating to adding a new datastore and redistributing VM storage. (*=estimated)

Isolating a storage-intensive VM

Both VMware and Microsoft virtualization solutions implement some degree of resource management when it comes to CPU and RAM. However, when a particular user’s VMs overwhelm storage I/O resources, IT staff must isolate this “noisy neighbor” in order to distribute resources properly for other users. For VMware, this isolation process involves enabling storage I/O control and capping the VM IOPS within the vCenter Server console. As was the case with the previous storage scenario, Hyper-V has no equivalent feature. For Hyper-V to fully isolate the VM, the VM’s virtual disks must be offloaded to different physical storage. Figure 10 shows how VMware Storage I/O control works.




Figure 10: VMware Storage I/O control easily isolates and caps VMs’ storage bandwidth.

We isolated and redistributed resources from the noisy neighbor using both solutions, and found that it took 97 percent less time to do so using the VMware solution compared to the Microsoft solution (see Figure 11). VMware vSphere Storage I/O Control was able to quickly isolate the user, where Microsoft’s manual isolation approach took significantly longer. We provide the detailed steps we used in Figure 12 and in [Appendix E](#). This feature is not included in vSphere Standard edition, and therefore is included only in the vSphere Enterprise Plus edition calculations in this report.

For our comparison, on the Microsoft side, we assume no additional costs for purchasing new storage hardware for isolation. We assume the company has existing storage that they can re-provision for this isolation event. In our lab, we re-provisioned additional iSCSI storage, but similar steps would exist for provisioning additional Fibre Channel trays and fabric.

Time to isolate a storage-intensive VM




Figure 11: It took 97 percent less time to isolate a storage-intensive VM using the VMware solution than it did with the Microsoft solution. Lower numbers are better.

VMware solution		Microsoft solution	
Task	Time	Task	Time
1. Enable Storage I/O Control on each datastore to balance I/O usage across VMs.	0:00:24	1. Install new NICs on each of three hosts, migrating the VMs off each host before shutting down.	0:50:27
2. Adjust the advanced Storage I/O Control setting for the congestion threshold.	0:00:50	2. Rack and cable the new storage tray.	0:10:00
3. Adjust single VM disk shares.	0:00:24	3. Configure the storage array for initial use using a serial connection, creating a new storage group and new storage pool, and using a separate IP subnet from the current	0:02:03

VMware solution		Microsoft solution	
		storage for complete fabric isolation.	
4. Adjust single VM virtual disk IOPS limit.	0:00:27	4. Use the EQL Web management console to configure LUN(s) on the new tray.	0:05:10
		5. Update the necessary drivers for the new NICs on each host.	0:14:24
		6. Configure each new NIC for iSCSI (MTU, IP addresses) on each host.	0:06:33
		7. Using the iSCSI initiator, connect to the new LUN(s) on each host in the cluster.	0:01:48
		8. Using disk management, bring the LUN(s) online and format on each host in the cluster.	0:02:42
		9. In Failover Clustering Services, add the new disks as a cluster disk(s).	0:00:21
		10. Add the disk(s) to cluster shared volumes.	0:00:26
		11. Using SCVMM, move the noisy VM(s) to the new disk with the quick storage migration feature.	0:00:38
Total	0:02:05	Total	1:34:32

Figure 12: Times, in hours:minutes:seconds, to complete the tasks relating to redistributing resource from a noisy neighbor VM.

Provisioning new hosts

Provisioning new hosts in a data center environment is a constant requirement if your business is growing, or even if your business is simply refreshing your hardware. Each solution has automated tools to accomplish the provisioning task. In our testing, we set up both platforms' automated solutions: for VMware we used VMware vSphere Auto Deploy (see Figure 13), and for Microsoft Hyper-V we used System Center Configuration Manager 2007 R3 bare metal deployment task sequence.




Figure 13: VMware Auto Deploy quickly deploys new diskless hosts.

Using VMware Auto Deploy provisioned new hosts more quickly than using Microsoft SCCM 2007 R3—by up to 78 percent—and without the use of onboard storage (see Figure 14). We provide the detailed steps we followed in Figure 15 and

[Appendix F](#). This feature is not included in vSphere Standard edition, and therefore is included only in the vSphere Enterprise Plus edition calculations in this report.

Time to provision new hosts




Figure 14: It took 78 percent less time to provision new hosts using the VMware solution than it did with the Microsoft solution. Lower numbers are better.

VMware solution		Microsoft solution	
Task	Time	Task	Time
1. Click Apply Host Profile.	0:00:05	1. Enter license and log into the domain.	0:00:45
2. Answer profile questions.	0:01:03	2. Connect LUNs via iSCSI Initiator.	0:01:44
3. Wait until host is configured and ready.	0:01:45	3. Bring disks online via Disk management.	0:00:36
		4. Create four new virtual networks for Hyper-V.	0:02:12
		5. Join host to the cluster.	0:02:04
Total	0:02:53	Total	0:07:21

Figure 15: Times, in hours:minutes:seconds, to complete the tasks relating to provisioning new hosts.

Performing non-disruptive disaster recovery testing

We set out to test a non-disruptive disaster recovery plan, where each step of the process causes no downtime, retargeting of production workloads, or production networking changes. For VMware, we used VMware Site Recovery Manager, and for Microsoft we used two distinct site clusters and a manual runbook procedure. We opted not to use a geographically stretched Hyper-V failover cluster, because their distance limitations can make them unsuitable for some use cases and there is no way to perform disaster recovery testing scenarios without disrupting or altering the production workload.

In our testing, we measured the time it took to perform a complete non-disruptive disaster recovery test using VMware Site Recovery Manager, then measured

or approximated the equivalent actions using the Microsoft solution. For our time calculation scenarios, we assume the organization has two SAN systems and 100 VMs, but only 20 of the VMs in each scenario are tier 1 and must be tested for DR purposes.

In our configuration, the non-disruptive test of a disaster recovery scenario using VMware is 94 percent less time-consuming to perform than that of Microsoft (see Figure 16). We provide the detailed steps we followed in Figure 17 and [Appendix G](#). This feature is compatible with both vSphere Enterprise Plus and vSphere Standard editions, and therefore is included in both calculations in this report.

Time to test non-disruptive DR recovery




Figure 16: It took 92 percent less time to perform a non-disruptive test for disaster recovery using the VMware solution than it did with the Microsoft solution.

VMware solution		Microsoft solution	
Task	Time	Task	Time
1. Time cost - Monthly maintenance of <i>wizard</i> -based recovery plan. ²	1:00:00*	1. Time cost - Monthly maintenance of <i>script</i> -based metadata for VM syncing, boot order preferences, and IP address changes that must occur on recovery.	10:00:00*
2. In vCenter Server, within the SRM plug-in, right-click your recovery plan and choose Test.	0:00:10	2. Pause SAN replication. ³	0:00:50
		3. Modify DNS or WAN to ensure no traffic flows to DR site. ⁴	0:10:00*
		4. Configure storage snapshots and volumes for DR test. ⁵	0:22:00

² We assume script-based recovery plans require 10x more time to maintain than graphical wizard-based recovery plans.

³ For a small to medium business, we assume only two SANs are relevant; therefore, we multiplied our original “pause” hand timing step (0:00:25) by two.

⁴ Estimated time to approximate networking staff adjusting configuration on networking hardware. We assume a flat 10-minute cost for this process.

VMware solution		Microsoft solution	
		5. On each host, online the disks. ⁶	0:03:02
		6. For each volume, attach to the cluster hosts. ⁷	0:20:00
		7. Run prepared scripts for VM power on and IP addressing. Perform DR testing. ⁸	0:30:00*
3. In vCenter Server, within the SRM plug-in, right-click your recovery plan and choose "Cleanup"	0:00:10	8. Run prepared scripts to power down VMs. ⁹	0:01:00*
		9. Cleanly remove volumes from DR cluster. ¹⁰	0:11:20
		10. For each volume on each host, offline the disks, and disconnect the disks. ¹¹	0:00:20
		11. Clean up and revert storage configuration from DR test. ¹²	0:11:20
		12. Revert DNS or WAN for normal operation. ¹³	0:10:00*
		13. Unpause SAN replication. ¹⁴	0:00:50
Total	1:00:20	Total	12:00:42

Figure 17: Times, in hours:minutes:seconds, to complete the tasks relating to provisioning new hosts. (*=estimated)

⁵ This time will differ by SAN vendor. Our manual process on the Dell EqualLogic storage in our lab was to mimic the automated process that VMware performed. We manually promoted the DR replica set to a volume, which automatically created writeable snapshots for DR testing. We assume 10 volumes per SAN, and two DR SANs; therefore, we multiplied our original time (0:01:06) times 20.

⁶ We assume 20 of our 100 VMs are tier 1 protected VMs. We also assume a host density for Microsoft of 12 VMs per host, which amounts to seven hosts (20/12=1.66, which requires two hosts). Therefore, we multiplied our original time (0:01:31) by two.

⁷ We assume 10 volumes per SAN, and two DR SANs; therefore, we multiplied our original time (0:01:00) by 20.

⁸ We assume a flat 30-minute cost for this process.

⁹ We assume a flat 1-minute cost for this process.

¹⁰ We assume 10 volumes per SAN, and two DR SANs; therefore, we multiplied our original time (0:00:34) by 20.

¹¹ We assume two hosts (see footnote 6), each only connecting to one volume. Therefore, we multiplied our original time (0:00:10) by two.

¹² This time will differ by SAN vendor. Our manual process on the Dell EqualLogic storage in our lab was to mimic the automated process that VMware performed. We manually removed the writeable snapshots on the storage, and then demoted volume to a replica set for DR replication. We assume 10 volumes per SAN, and two DR SANs; therefore, we multiplied our original time (0:00:34) by 20.

¹³ Estimated time to approximate networking staff adjusting configuration on networking hardware. We assume a flat 10-minute cost for this process.

¹⁴ We assume two SANs in our sample organization. Therefore, we multiplied our original "unpause" hand timing step (0:00:25) by two.

CALCULATING TWO-YEAR OPERATIONAL COSTS FOR THE SCENARIOS

To calculate the management operational costs of the two solutions, we timed how long it took to perform the tasks in each of our management scenarios with VMware and Hyper-V. We were able to perform each set of tasks faster with VMware. Figure 18 shows our tested times and the time savings for VMware vs. Hyper-V on the test scenarios. For our vSphere Standard edition calculations, we used only Scenario 1 and Scenario 5.

Scenario	VMware solution	Microsoft solution	Savings with VMware solution
Scenario 1: Shifting virtual machines workloads for host maintenance*	0:02:16	0:11:05	0:08:49
Scenario 2: Adding new volumes and redistributing VM storage	0:04:23	1:43:06	1:38:43
Scenario 3: Isolating a storage-intensive VM	0:02:05	1:34:32	1:32:27
Scenario 4: Provisioning new hosts	0:02:53	0:07:21	0:04:28
Scenario 5: Performing non-disruptive disaster recovery testing*	1:00:20	12:00:42	11:00:22

Figure 18: Time savings in hours:minutes:seconds for VMware compared to Hyper-V on test scenarios. Times and savings are for one iteration of each scenario on our tested server. *Scenarios we used for vSphere Standard Edition calculations.

To illustrate how these time savings can affect an organization’s bottom line, we assumed two example 100-VM cost scenarios. In the first, we compare the Microsoft solution to a solution running vSphere Enterprise Plus Edition, and in the second we compare the Microsoft solution to a solution running vSphere Standard Edition. Both environments assume a VM density of 15 VMs per server for VMware vSphere servers and 12 VMs per server for Microsoft Hyper-V servers. We then calculated the cost savings for a small to medium business that chooses VMware vSphere over Microsoft Hyper-V and must repeat many of these scenarios throughout a typical two-year period. For the vSphere Enterprise Plus edition calculations, we used all five scenarios to estimate operational savings. For the vSphere Standard edition calculations, we used two scenarios (Scenario 1 and Scenario 5) to estimate operational savings. We assumed the tasks would be carried out by a senior system administrator and calculated costs based on that individual’s salary plus benefits.¹⁵ Each minute of that Senior System Administrator’s time is valued at \$1.02. Figure 19 shows the cost per iteration using the times in Figure 19 multiplied by \$1.02.

¹⁵ The average national base salary for a senior system administration was \$88,599 and total compensation was \$126,662 according to salary.com on March 5, 2012. Total compensation includes base salary, employer contributions for bonuses, Social Security, 401k and 401b, disability, healthcare, and pension, and paid time off. We calculated the average cost per minute for a Senior Systems Administrator at that salary at \$1.02 based on 52 forty-hour weeks.

Scenario	VMware cost per iteration	Microsoft cost per iteration
Scenario 1: Shifting virtual machines workloads for host maintenance*	\$2.32	\$11.30
Scenario 2: Adding new volumes and redistributing VM storage	\$4.47	\$105.16
Scenario 3: Isolating a storage-intensive VM	\$2.12	\$96.42
Scenario 4: Provisioning new hosts	\$2.94	\$7.50
Scenario 5: Performing non-disruptive disaster recovery testing*	\$61.54	\$735.11

Figure 19: Cost for one iteration of each scenario using each solution. *Scenarios we used for vSphere Standard Edition calculations.

We then estimated the number of times the system administrator would need to carry out these tasks per two-year period for each scenario. To estimate the number of tasks per two-year period, we factored in the number of VMs (100), the aforementioned VM densities by platform, and industry experience to come up with reasonable estimates of maintenance events, storage additions, deployments, and so on. Below, we present the assumptions we used to calculate the number of events for cost comparisons.

Shifting virtual machine workloads for host maintenance¹⁶

We assume quarterly firmware and BIOS checks per server, and a hardware failure rate of 5 percent of the total servers per year. This equates to 58 events for this scenario for VMware and 74 events for Microsoft, over a two-year period. We used the following calculations:

VMware solution

$(100 \text{ VMs} / 15 \text{ VMs per server}) = 7 \text{ servers}$
 $(7 \text{ servers} * 4 \text{ quarters}) + \text{round}(7 \text{ servers} * 0.05 \text{ failure rate}) = 29 \text{ events per year}$
 $29 \text{ events} * 2 \text{ years} = 58 \text{ events}$

Microsoft solution

$(100 \text{ VMs} / 12 \text{ VMs per server}) = 9 \text{ servers}$
 $(9 \text{ servers} * 4 \text{ quarters}) + \text{round}(9 \text{ servers} * 0.05 \text{ failure rate}) = 37 \text{ events per year}$
 $37 \text{ events} * 2 \text{ years} = 74 \text{ events}$

Savings per two-year period = $(74 * \text{Microsoft event cost}) - (58 * \text{VMware event cost})$

¹⁶ This scenario was used in both vSphere Enterprise Plus and vSphere Standard calculations.

Adding new volumes and redistributing VM storage¹⁷

In this scenario, we assume only two storage systems, each of which requires that a new LUN be provisioned once monthly, resulting in 24 events for this scenario over a two-year period.

Isolating a storage-intensive VM¹⁸

We did not factor in cost requirements for new hardware, only the time it took to provision the hardware for the isolation event. We assume that a small to medium business with 100 VMs would require at least one isolation event every four months, for six events per two-year period.

Provisioning new hosts¹⁹

We assume that a business would refresh one-third of its hosts annually and assume an additional 10 percent growth rate. We calculated the number of deployment events as follows:

VMware solution

$(100 \text{ VMs} / 15 \text{ VMs per server}) = 7 \text{ servers}$
 $\text{round}(7 \text{ servers} * .33) + \text{round}(7 \text{ servers} * 0.1) = 4 \text{ events}$
 $4 \text{ events} * 2 \text{ years} = 8 \text{ events}$

Microsoft solution

$(100 \text{ VMs} / 12 \text{ VMs per server}) = 9 \text{ servers}$
 $\text{round}(9 \text{ servers} * .33) + \text{round}(9 \text{ servers} * 0.1) = 4 \text{ events}$
 $4 \text{ events} * 2 \text{ years} = 8 \text{ events}$

$\text{Savings per two-year period} = (8 * \text{Microsoft cost}) - (8 * \text{VMware cost})$

Performing non-disruptive disaster recovery testing²⁰

We assume a monthly test of disaster recovery, for 24 events per two-year period.

Figure 20 shows the estimated number of events and the subsequent savings per two-year period that a small to medium business running 100 VMs could realize when choosing VMware and managing these scenarios; that value is the product of the estimated events value multiplied by the savings per iteration value in the previous figure.

¹⁷ This scenario was used only in vSphere Enterprise Plus calculations

¹⁸ This scenario was used only in vSphere Enterprise Plus calculations

¹⁹ This scenario was used only in vSphere Enterprise Plus calculations

²⁰ This scenario was used in both vSphere Enterprise Plus and vSphere Standard calculations.

100-VM small to medium business scenario		
Scenario	Total events per two-year period	Savings per two-year period
Scenario 1: Shifting virtual machines workloads for host maintenance*	VMware: 58 Microsoft: 74	\$701.64
Scenario 2: Adding new volumes and redistributing VM storage	24	\$2,416.56
Scenario 3: Isolating a storage-intensive VM	6	\$565.80
Scenario 4: Provisioning new hosts	VMware: 8 Microsoft: 8	\$36.48
Scenario 5: Performing non-disruptive disaster recovery testing*	24	\$16,165.68
Total savings		\$19,886.16

Figure 20: Estimated operational cost savings based on these scenarios when using VMware vs. Microsoft with 100 VMs over a two-year period. *Scenarios we used for vSphere Standard Edition calculations.

CALCULATING ACQUISITION AND CAPITAL COSTS

We used the VMware Cost-Per-Application Calculator to calculate the acquisition costs of virtualization platforms needed to support a business with 100 VMs. We used the following as inputs to the calculator: 100 VMs, “Typical” workload profile, “Server B” configuration, iSCSI storage, use of virtual management servers, and average electricity and real estate costs. We ran the calculator twice, the first time choosing vSphere Enterprise Plus edition, the second time choosing vSphere Standard edition. We used results from each of those outputs in this report.

The VMware Cost-Per-Application Calculator normally factors in a 50 percent VM density advantage for vSphere over Hyper-V, but we use a more conservative 25 percent advantage for VMware (12 VMs per host for Microsoft, 15 VMs per host for VMware). With those assumptions at 100 VMs, the VMware Cost-Per-Application Calculator finds that the VMware platform requires seven vSphere hosts, while the Microsoft platform requires nine Hyper-V hosts. Management servers were virtualized and required no additional physical hosts. See the VMware Cost-Per-Application Calculator methodology paper²¹ for details. Additionally, we factor in VMware Site Recovery Manager Standard Edition acquisition cost and two years of support for our 20 protected VMs in each VM count scenario at a cost of \$5,538. The calculated costs of hardware (servers, networking, and storage), software (virtualization, management, OS licenses, VMware vCenter Site Recovery Manager) and data center infrastructure with two years of support are as follows:

- VMware vSphere Enterprise Plus edition: \$251,938
- VMware vSphere Standard edition: \$204,238
- Microsoft: \$235,216

²¹ <http://www.vmware.com/go/costperapp-calc-methods>

CALCULATING TOTAL COST OF OWNERSHIP

We calculated the two-year total cost of ownership as the sum of the platform acquisition costs generated by the VMware Cost-Per-Application Calculator and the operational costs of the scenarios we evaluated for our hypothetical 100-VM business (see Figures 21 and 22).

100-VM scenario – vSphere Enterprise Plus edition, five scenarios tested		
	VMware solution	Microsoft solution
Two-year CAPEX (hardware, software, and support costs)	\$251,938	\$235,216
Two-year operating expenses (from five evaluated scenarios)	\$1,755	\$21,641
Two-year TCO	\$253,693	\$256,857

Figure 21: Two-year total cost of ownership for the two solutions at 100 VM count.

100-VM scenario – vSphere Standard edition, two scenarios tested		
	VMware solution	Microsoft solution
Two-year CAPEX (hardware, software, and support costs)	\$204,238	\$235,216
Two-year operating expenses (from five evaluated scenarios)	\$1,612	\$18,479
Two-year TCO	\$205,850	\$253,695

Figure 22: Two-year total cost of ownership for the two solutions at 100 VM count.

The results show that VMware’s lower operational costs can lead to a lower TCO for the VMware platform compared to Microsoft, when considering the five scenarios we tested for vSphere Enterprise Plus edition and the two scenarios we tested for vSphere Standard edition. However, these scenarios are only a small subset of the typical operational requirements of an organization, and other studies of cross-industry IT spending show that annual operational expenses are over two times capital expenses.²² This means the impact of operational cost savings for platform technologies such as virtualization may be multiplied well beyond the totals for the common tasks we include in this analysis. Therefore, organizations may find that additional features of VMware vSphere 5—such as a single unified management interface in vCenter, hot-add CPU for guest VMs, VM-to-host and VM-to-VM affinity capabilities, and VM storage tier placement automation—could lead to further operational time savings.

²² <http://storage.networksasia.net/content/migrating-cloud-beware-prickly-financial-situations>

WHAT WE TESTED

About VMware vSphere 5

vSphere 5 is the latest virtualization platform from VMware. vSphere 5 allows companies to virtualize their server, storage, and networking resources, achieving significant consolidation ratios, all while gaining significant management time savings as we demonstrate in this paper. To learn more about VMware vSphere 5, visit <http://www.vmware.com/products/vsphere/overview.html>.

About Microsoft Windows Server 2008 R2 Hyper-V

Microsoft Windows Server 2008 R2, Microsoft's server operating system platform, includes the Hyper-V hypervisor for virtual infrastructures. The management products included in the Microsoft solution are System Center Virtual Machine Manager 2008 R2, which enables centralized management of physical and virtual IT infrastructure, and System Center Configuration Manager 2007, which enables deployment and other features.

IN CONCLUSION

We designed our assumptions based on a small-to medium-business using 100 VMs. However, for businesses that grow and increase their virtual workloads over time, the potential for greater operational savings exists, as administrators are able to take advantage of cost-saving features of VMware vSphere at a greater scale. For example, in a separate report, we highlight the potential two-year operational savings of a 1,000 VM data center.²³

Managing a virtualized infrastructure that runs continuously inevitably requires some degree of maintenance from IT staff. Any time that can be saved when performing routine maintenance tasks through system automation and capable management features frees IT staff to concentrate on ways to help your business grow. In the scenarios we tested, using the VMware solution had the potential to reduce administrative labor costs by as much as 91 percent compared to using similar offerings from Microsoft.

When we added the expected operational efficiency cost savings to the hardware acquisition estimates provided by the VMware Cost-Per-Application Calculator, we found that the VMware solution could provide a lower total cost of ownership over two years compared to the Microsoft solution.

²³ http://www.principledtechnologies.com/clients/reports/vmware/vSphere_Hyper-V_TCO_0412.pdf

APPENDIX A – SERVER AND STORAGE CONFIGURATION INFORMATION

Figures 23 and 24 present configuration information for our test servers and storage.

System	3 x Dell PowerEdge R710 servers
Power supplies	
Total number	2
Vendor and model number	Dell Inc. N870P-S0
Wattage of each (W)	870
Cooling fans	
Total number	5
Vendor and model number	Nidec UltraFlo™ RK385-A00
Dimensions (h x w) of each	2.5 x2.5
Volts	12
Amps	1.68
General	
Number of processor packages	2
Number of cores per processor	6
Number of hardware threads per core	2
CPU	
Vendor	Intel®
Name	Xeon®
Model number	X5670
Stepping	B1
Socket type	FCLGA 1366
Core frequency (GHz)	2.93
Bus frequency	6.4 GT/s
L1 cache	32 KB + 32 KB (per core)
L2 cache	6 x 256 KB (per core)
L3 cache	12 MB (shared)
Platform	
Vendor and model number	Dell PowerEdge R710
Motherboard model number	OYDJK3
BIOS name and version	Dell Inc. 6.0.7
BIOS settings	Default
Memory module(s)	
Total RAM in system (GB)	96
Vendor and model number	M393B1K70BH1-CH9
Type	PC3-10600
Speed (MHz)	1,333
Speed running in the system (MHz)	1,333
Timing/Latency (tCL-tRCD-tRP-tRASmin)	9-9-9-24
Size (GB)	8
Number of RAM module(s)	12
Chip organization	Double-sided
Rank	Dual

System	3 x Dell PowerEdge R710 servers
Microsoft OS	
Name	Windows Server 2008 R2 SP1
Build number	7601
File system	NTFS
Kernel	ACPI x64-based PC
Language	English
VMware OS	
Name	VMware vSphere 5.0.0
Build number	469512
File system	VMFS
Kernel	5.0.0
Language	English
Graphics	
Vendor and model number	Matrox® MGA-G200ew
Graphics memory (MB)	8
RAID controller	
Vendor and model number	PERC 6/i
Firmware version	6.3.0-0001
Cache size (MB)	256
Hard drives	
Vendor and model number	Dell ST9146852SS
Number of drives	4
Size (GB)	146
RPM	15,000
Type	SAS
Onboard Ethernet adapter	
Vendor and model number	Broadcom® NetXtreme® II BCM5709 Gigabit Ethernet
Type	Integrated
10Gb Fibre adapter for vMotion scenario	
Vendor and model number	Intel Ethernet Server Adapter X520-SR1
Type	Discrete
Quad-port Ethernet adapter for Storage I/O Control scenario	
Vendor and model number	Intel PRO/1000 Quad Port LP SVR Adapter
Type	Discrete
Optical drive(s)	
Vendor and model number	TEAC DV28SV
Type	DVD-ROM
USB ports	
Number	6
Type	2.0

Figure 23: Detailed configuration information for our test servers.

Storage array	Dell EqualLogic™ PS5000XV storage array
Arrays	3
Number of active storage controllers	1
Number of active storage ports	3
Firmware revision	5.1.2
Switch number/type/model	Dell PowerConnect™ 5448
Disk vendor and model number	Dell ST3600057SS/ST3450856SS/ST3600002SS
Disk size (GB)	600/450/600
Disk buffer size (MB)	16
Disk RPM	15,000/15,000/10,000
Disk type	6.0 Gbps SAS / 3.0 Gbps SAS/ 6.0 Gbps SAS
EqualLogic Host Software for Windows	Dell EqualLogic Host Integration Tools 3.5.5
EqualLogic Host Software for VMware	Dell EqualLogic Multipathing Extension Module (MEM) 1.1

Figure 24: Detailed configuration information for our test storage.

APPENDIX B – TEST HARDWARE SETUP

For our testing, we configured three servers, each with two RAID 1 volumes, each consisting of two 146GB SAS drives. On the first volume, we installed a default installation of Windows Server 2008 R2, enabled the Hyper-V role, and performed additional Microsoft Failover Clustering steps, as specified by Microsoft documentation. On the second volume of each server, we installed a default installation of VMware vSphere 5 (ESXi). We updated each operating system with current drivers and updates.

To ensure hardware had no effect on time measurements, we used the same three servers for both VMware and Microsoft. To switch between platforms between each test scenario, we toggled the boot volume by using the PERC 6/i RAID controller to control which operating system environment to boot.

For external storage, we used three Dell EqualLogic PS5000XV arrays, each containing 16 drives and configured for RAID 10 mode. One array contained 10K SAS drives, and the remaining two arrays contained 15K SAS drives. We created one storage pool and assigned the slower drive tray to it. We created an additional storage pool and assigned the two faster drive trays to it. We spread VM files equally amongst both storage pools for each platform.

We cabled each Dell EqualLogic PS5000XV array to a Dell PowerConnect 5448 switch via their three available ports, and cabled each server to the same switch using two onboard server NICs for iSCSI traffic. We configured each operating system and the switch for iSCSI optimizations, such as jumbo frames, as specified by each vendor's documentation. For specific configurations on the Dell PowerConnect 5448, we used the recommended settings from the Dell EqualLogic Configuration Guide.

For each platform, we created 30 VMs, each installed with Windows Server 2008 R2 as the guest OS. We configured each VM with 2 vCPUs and 10 GB of RAM. We configured four attached virtual disks: a 13GB disk for the OS, and three additional virtual disks (25GB, 15GB, and 4GB). Altogether, each VM had 57GB of storage attached. For networking, each VM had one virtual network connection using the hosts' network connection to a Dell PowerConnect 6248, the switch we used for VM and management traffic.

For scenario-specific changes after the base installations, see the individual methodology appendices below.

APPENDIX C – SCENARIO 1: SHIFTING VIRTUAL MACHINE WORKLOADS FOR HOST MAINTENANCE

VMware vSphere 5: Additional configuration after installation for Scenario 1

To perform this test, we configured our vMotion network and VMware DRS to handle the offloading and evacuation of VMs from one host to the others in the cluster. The following steps outline how to configure the additional 10Gb Intel X520-SR1 server adapter on each host for vMotion and enable VMware vSphere Distributed Resource Scheduler (DRS) on the test cluster. On each host, the Intel NIC was cabled to a 10Gb-capable switch.

1. Log into the vCenter server via the vSphere client.
2. On the vCenter console, click Hosts and Clusters.
3. On the left, click the first host.
4. Click the Configuration tab, and click Networking.
5. Click Add Networking...
6. On the Connection Type screen, select VMkernel, and click Next.
7. On the VMkernel-Network Access screen, select the Intel X520-SR1 server adapter, and ensure that no other adapters are selected.
8. Click Next.
9. On the VMkernel-Connection Settings screen, enter a name for the new connection in the Network Label box.
10. Check the box Use this port group for vMotion, and click Next.
11. On the VMkernel-IP Connection Settings screen, enter a new IP address and subnet mask for the new network object, preferably on a separate subnet than the VM and storage networks.
12. Click Next.
13. On the Ready to Complete screen, click Finish.
14. Repeat steps 4-13 on the remaining two hosts, using the same IP subnet for each as the first host's vMotion network.
15. To ensure the network settings are correct on each host, migrate a VM to and from each host.
16. To enable DRS, right-click the test cluster, and click Edit Settings...
17. On the Cluster Features screen, check the box beside Turn On vSphere DRS.
18. For our purposes, we left all DRS settings at default. Click OK.

VMware vSphere 5: Running the test for Scenario 1

In our test, we assigned six VMs to each host for a total of 18 VMs. We used a medium workload database benchmark to fill each VM's RAM and to run during the test measurement period. We placed one host into maintenance mode and allowed vSphere DRS to evenly distribute the VMs from the original host to the two remaining hosts using vSphere vMotion. Below are the steps we performed to complete the test. We timed and recorded each step and report those results in the body of the report.

1. Begin the benchmark to fill the VM's RAM allotment.
2. With all VMs distributed evenly at six per host and with the benchmark running, right-click one host, and click Enter Maintenance Mode.
3. When the Confirm Maintenance Mode prompt appears, uncheck the box to ensure that no other migration activity will occur during the test, and click OK.
4. Once all six VMs are migrated to the remaining two hosts, right-click the host that is in maintenance mode, and click Shutdown. At this point we stopped timing momentarily, simulating a system upgrade or hardware replacement and reboot.

5. Once the host connects back to the cluster, begin timing again, right-click the host, and click Exit Maintenance Mode.
6. Once the host is out of maintenance mode, select the cluster, and click the Virtual Machines tab.
7. Multi-select the six VMs that migrated to the other two hosts, and click Migrate...
8. On the Select Migrate Type screen, select Change host, and click Next.
9. On the Select Destination screen, expand the cluster, select the recently booted host, and click Next.
10. On the Ready to Complete screen, click Finish.

Microsoft Hyper-V: Additional setup after installation for Scenario 1

To perform this test, we configured our live migration network to handle the offloading and evacuation of VMs from one host to the others in the cluster. The following steps outline how to configure the additional 10Gb Intel X520-SR1 server adapter on each host for Microsoft Live Migration. After installing the new NICs, we installed the latest drivers on each host. From each host, the Intel NIC was cabled to a 10Gb-capable switch. Because a maintenance mode feature is not native to Hyper-V and Failover Clustering Services, we set up a separate server on our domain running Microsoft System Center Virtual Machine Manager and used this to invoke the maintenance mode event. We assume SCVMM 2008 R2 installation is performed prior to testing and the hosts have been added to SCVMM 2008 R2.

1. Log into the first host, and open Network and Sharing Center.
2. Click Change adapter settings.
3. Right-click the new Intel NIC, and click Properties.
4. Select Internet Protocol Version 4, and click Properties.
5. Enter an IP address and subnet mask for the new network connection. Make sure to assign an IP on a separate subnet than the domain and storage networks.
6. Repeat steps 1-5 on each of the remaining two hosts.
7. Open Server Manager, and expand Features→Failover Cluster Manager→*cluster name*.
8. On the left side, click Networks, and ensure the new network subnet has been added to the cluster.
9. On the left side, expand Services and Applications, and click a VM.
10. In the center pane, right-click the VM object, and click Properties.
11. Click the Network for Live Migration tab.
12. Check the box next to the new network, and click OK.

Microsoft Hyper-V: Running the test for Scenario 1

In our test, we assigned six VMs to each host for a total of 18 VMs. We used a medium workload database benchmark to fill each VM's RAM and to run during the test measurement period. We placed one host into maintenance mode and allowed SCVMM to evenly distribute the VMs from the original host to the two remaining hosts using Microsoft Live Migration. Below are the steps we performed to complete the test. We timed and recorded each step and report those results in the body of the report.

1. Begin the benchmark to fill the VM's RAM allotment.
2. With the VMs evenly distributed and the benchmark running, log into the SCVMM server, and open the SCVMM Administrator Console.
3. On the left, click Virtual Machines.
4. Right-click the target host, and click Start maintenance mode.
5. On the Start Maintenance Mode Screen, select the Live Migration option, and click Start Maintenance mode.
6. Once all the VMs have been migrated, log into the target host, and shut it down. At this point we stopped timing momentarily, simulating a system upgrade or hardware replacement and reboot.
7. Once the host becomes available in Failover Cluster Manager, begin timing again. Click Services and Applications, and right-click one of the previously migrated VMs.

8. Click Live migrate virtual machine to another node, and select the target server.
9. Repeat steps 7-8 for the remaining nine VMs, and include the time for the last one to finish migrating.

APPENDIX D – SCENARIO 2: ADDING NEW VOLUMES AND REDISTRIBUTING VM STORAGE

VMware vSphere 5: Additional setup after installation for Scenario 2

In this scenario, we simulated a volume nearing capacity, the addition of new storage to a cluster, and the redistribution of VM storage after that event. To simulate a low storage capacity incident, we used Storage vMotion to artificially load two of our storage LUNs near top capacity with VMs. Via the Dell EqualLogic console, we then created one additional LUN on a separate storage tier to be used as our newly added storage to relieve the two full LUNs. We sized the new LUN at 1 TB to ensure capacity requirements were met.

We then created a datastore cluster using the two full LUNs and enabled VMware vSphere Storage DRS. Below are the steps required to create a SDRS enable datastore cluster.

1. Log into the vCenter server via a vSphere client.
2. On the vCenter home page, click Datastores and Datastore Clusters.
3. On the left, right-click the data center object, and click New Datastore Cluster.
4. Enter a name for the new datastore cluster, and leave the Turn on Storage DRS checkbox selected.
5. Click Next.
6. On the SDRS Automation screen, select Fully Automated, and click Next.
7. On the SDRS Runtime Rules screen, leave the default I/O Metric Inclusion and Threshold settings. Note the Utilized space Threshold is set to 80 percent.
8. Click Next.
9. On the Select Hosts and Clusters screen, check the box next to the test cluster.
10. On the Select Datastores screen, select the two identical datastores that were loaded near capacity as described previously, and click Next.
11. On the Summary screen, click Finish.

VMware vSphere 5: Running the test for Scenario 2

We timed and recorded the following steps and report those times in the body of the report.

1. Select a host in the test cluster, and click the Configuration tab.
2. Under the Hardware heading, click Storage Adapters.
3. Right-click the configured iSCSI Software Adapter, and click Rescan.
4. Once the Rescan VMFS task has completed, click Storage.
5. In the upper right corner, click Add Storage...
6. On the Select Storage Type screen, select Disk/LUN, and click Next.
7. On the Select Disk/LUN screen, select the new LUN, and click Next.
8. On the File System Version screen, leave the default set to VMFS-5, and click Next.
9. On the Current Disk Layout screen, click Next.
10. On the Properties screen, enter a new for the new datastore, and click Next.
11. On the Disk/LUN-Formatting screen, select Maximum available space, and click Next.
12. On the Ready to Complete screen, review the datastore settings, and click Finish.
13. Once the Create VMFS Datastore task completes, return to the Datastores and Datastore Clusters screen.
14. Right-click the previously created datastore cluster, and click Add Storage...
15. Select the new datastore, and click OK.
16. Once the datastore has been added to the datastore cluster, select the datastore cluster, and click the Storage DRS tab.

17. By default, Storage DRS runs automatically every 8 hours. To manually initiate a Storage DRS action, click Run Storage DRS in the upper right corner.
18. Because the original two datastores in the cluster are near capacity (over the 80 percent threshold) and there is now new capacity added to the datastore cluster, Storage DRS will now make recommendations to bring the two datastores under the 80 percent threshold by moving VMs to the new datastore. To begin moving the VMs, click Apply recommendations.

Note: In our testing, exactly three VMs from each of the original two datastores moved to the new datastore for a total of six VMs. We did not include the time it took for the migration, because Storage DRS automates the rest of the process and requires no more system administrator interaction.

Microsoft Hyper-V: Additional setup after installation for Scenario 2

In this scenario, we simulated a volume nearing capacity, the addition of new storage to a cluster, and the redistribution of VM storage after that event. As in the VMware portion of this scenario, we loaded two identical LUNs near capacity with VMs. Prior to running this test, using the Dell EqualLogic Web console, we created a new LUN on an existing storage tier to add to our cluster. We sized the new LUN at 1 TB to ensure capacity requirements were met. As with the prior scenario, we used SCVMM 2008 R2. No additional steps were necessary to the SCVMM setup, as the Quick Storage Migration is a default feature.

Microsoft Hyper-V: Running the test for Scenario 2

The following steps outline how we conducted this test. We include two time estimates into the overall timing for this test that are not covered in the steps below. The first is the time it takes a system administrator to coordinate with VM application stakeholders and business users due to the brief VM downtime that is incurred, since Quick Storage Migration enters a save state during the migration. The second is the time it takes the system administrator to manually determine VM placement to relieve the capacity issues.

1. Log into the first host.
2. Click Start→Administrative Tools→iSCSI Initiator.
3. Click the Targets tab, and click Refresh.
4. Connect to the new LUN, checking the Enable multi-path box.
5. Click OK to close the iSCSI Initiator Properties window.
6. Open the Server Management console.
7. On the left side, click Storage→Disk Management.
8. Right-click the newly connected disk, and click Initialize.
9. Once the disk is initialized, right-click the disk, and click New Simple Volume.
10. Run through the New Simple Volume Wizard, and format the new disk.
11. On the remaining two servers, run steps 1-8, right-click the new disk, and click Online.
12. Log into the first host, and open Server Manager.
13. On the left side, expand Features→Failover Cluster Manager→*Cluster Name*
14. On the left side, click Storage, and click Add a disk.
15. Select the new disk, and click OK.
16. Once the disk is done being added, click Cluster Shared Volumes on the left side.
17. Click Add Storage.
18. Select the newly added cluster disk, and click OK.
19. Log into the SCVMM server, and open the SCVMM Administration console.
20. On the left-hand, click Virtual Machines.
21. Select the cluster.

22. Right-click one of the target VMs, and click Migrate Storage.
23. Change the VM and four VHDs file path to the new storage LUN, and click Next.
24. On the Summary screen, click Move to begin the migration.
25. Repeat steps 22-24 for the remaining five VMs.

APPENDIX E – SCENARIO 3: ISOLATING A STORAGE-INTENSIVE VM

VMware vSphere 5: Additional setup after installation for Scenario 3

In this scenario, we simulated the storage isolation of a VM. Because VMware vSphere Storage I/O Control is a basic feature of vSphere 5, no additional setup was required to complete this test scenario.

VMware vSphere 5: Running the test for Scenario 3

We timed and recorded the following steps and report those times in the body of the report.

1. While logged into the vCenter server via a vSphere Client, select the host where the target VM is located.
2. Click the Virtual Machines tab, and right-click the target VM.
3. Click Edit Settings...
4. Click the Resources tab, and select Disk.
5. Adjust the IOPS for each disk to limit the resources allowed for each disk, and click OK.

Microsoft Hyper-V: Additional setup after installation for Scenario 3

In this scenario, we simulated the storage isolation of a VM. Microsoft does not offer specific tools to fully isolate VM IOPs. Therefore, there is no additional OS setup to prepare for this scenario. In our measurements, the process for isolating the target VM required the manual addition of a storage tray and storage network to the cluster, which we detail below.

Microsoft Hyper-V: Running the test for Scenario 3

To begin this scenario we had to put each host in maintenance mode to offload their VMs prior to installing the new NICs. Since this includes steps from scenario 1 (the VM offload and live migration timings), we used those timings as part of this scenario.

We timed and recorded the tasks we list below. Because we first put the server in maintenance mode to offload VMs (as we did in Scenario 1), we include that offloading time in this scenario as well.

Because this action resembled the first scenario, we included those times in the overall timing of this scenario, with the addition of installing the new hardware. After installing the new NICs, we racked and cabled a new storage tray in our test cluster. We then configured the new tray and created a new LUN via the EqualLogic web management console. With this all in place, we then powered the servers back on, updated the drivers for each new NIC, and configured each NIC with an IP address on the same subnet as the new storage and enabled Jumbo Frames on each NIC. Below are the step-by-step instructions for the rest of the test.

1. Log into the first host.
2. Click Start→Administrative Tools→iSCSI Initiator.
3. Click the Discovery tab, and add the IP address for the storage group to the list of Discover Portals.
4. Click the Targets tab, and click Refresh.
5. Connect to the new LUN, checking the Enable multi-path box.
6. Click OK to close the iSCSI Initiator Properties window.
7. Open the Server Management console.
8. On the left side, click Storage→Disk Management.
9. Right-click the newly connected disk, and click Initialize.
10. Once the disk is initialized, right-click the disk, and click New Simple Volume.
11. Run through the New Simple Volume Wizard, and format the new disk.
12. On the remaining two servers, run steps 1-8, right-click the new disk, and click Online.
13. Log into the first host, and open Server Manager.

14. On the left side, expand Features→Failover Cluster Manager→*Cluster Name*.
15. On the left side, click Storage, and click Add a disk.
16. Select the new disk, and click OK.
17. Once the disk is done being added, click Cluster Shared Volumes on the left side.
18. Click Add Storage.
19. Select the newly added cluster disk, and click OK.
20. Once the new disk has been added to Cluster Shared Volumes, log into the SCVMM server previously setup for the Live Migration and Quick Storage Migration scenarios.
21. Use Quick Storage Migration to move the target VM to the new storage disk, and include the time it takes for the migration to finish.

APPENDIX F – SCENARIO 4: PROVISIONING NEW HOSTS

VMware vSphere 5: Additional setup after installation for Scenario 4

Installing VMware Auto Deploy

VMware Auto Deploy requires a TFTP server. For our Auto Deploy procedure, we used <http://www.solarwinds.com>. We installed the TFTP server on the vCenter server. Auto Deploy also requires the installation of vSphere PowerCLI. For more information on PowerCLI, visit <http://www.vmware.com/support/developer/PowerCLI/>.

We configured TFTP to apply an offline ESXi bundle to any PXE boot server that was in the range of 192.168.20.190-199. After that, we powered on our target system.

1. Log onto the vCenter via vSphere Client.
2. From the VMware vCenter 5 install media, click Autorun.
3. Click VMware Auto Deploy.
4. Select English, and click Next.
5. At the install wizard welcome screen, click Next.
6. Agree to the license agreement, and click Next.
7. Select 2 GB for the repository size, and click Next.
8. Enter the vCenter IP address, for user name, type `administrator`, and enter the password for the administrator account.
9. Use the default server port 6501, click Next.
10. Select the option to Use the IP address of the server to identify auto deploy on the network, and click Next.
11. Click Install.
12. Click Finish.
13. In the vSphere client, click Plug-ins, and click Manage plug-ins...
14. Right-click Auto Deploy, and click Enable.
15. Ignore the security warning, and click the box next to the text that reads Install this certificate and do not display any security warnings about this host.
16. Close the Plug-in manager.
17. In the vSphere client browse to home → Administration → Auto Deploy → vCenter.
18. Click the Download the TFTP boot zip link.
19. Extract the TFTP boot files to the TFTP server (vCenter).

Configuring Auto Deploy ESXi software depot and deployment rule.

1. Download the ESXi 5.0 Offline Bundle from www.vmware.com.
2. Open PowerCLI.
3. Type `Set-ExecutionPolicy Unrestricted`
4. Type `Connect-VIServer vCenter IP Address`
5. Type `Add-EsxSoftwareDepot <location of the file in step 1>`
6. Type `Get-EsxImageProfile` and make note of the name for standard (example: ESXi-5.0.0-20111104001-standard).
7. Type `new-deployrule -name "IP-deployrule" -item "<name from step 6>","cluster name" -Pattern "ipv4=192.168.20.190-192.168.20.199"`
8. Type `Add-DeployRule -DeployRule "IP-deployrule"`

Creating a host profile

1. Open vSphere client and navigate to Hosts and Clusters.

2. Right-click one the preconfigured hosts, and select Host Profiles→Create Profile from Host.
3. Create a name for the new profile, and click Next.
4. On the Ready to Complete screen, click Finish.
5. Navigate to the vCenter home page, and click Host Profiles.
6. Right-click the newly created profile, and click Attach host/cluster.
7. On the Attach Host/Cluster, select the test cluster, click attach, and click OK.

VMware vSphere 5: Running the test for scenario 4

We powered on the target host and allowed it to PXE boot to the ESXi image from Auto Deploy before performing the following steps. Each step below was timed and recorded.

1. Once the target host has been added to the cluster by Auto Deploy, right-click the host, and click Apply Host Profile→Apply Profile...
2. When the Apply Profile wizard appears, click Next on the Software iSCSI Initiator Selection screen.
3. On the Initiator IQN screen, click Next to keep the default IQN for the iSCSI adapter.
4. On the Initiator Alias screen, enter an alias for the iSCSI adapter, and click Next.
5. On the Determine how MAC address for vmknic should be decided screen, enter a new MAC address, and click Next.
6. On the IPv4 address screen, enter a valid IP address and subnet for the storage networking, and click Next.
7. Repeat steps 5-6 for the second iSCSI vmknic.
8. On the Configuration Tasks summary screen, click Finish.
9. Allow the Apply host configuration task to complete, and add the time for the task to complete to the overall time for the test.

Microsoft Hyper-V: Additional setup after installation for scenario 4

In our testing, we used Microsoft System Center Configuration Manager 2007 R3 (SCCM) for comparison with VMware Auto Deploy. We installed SCCM on a VM running Windows Server 2008 R2 SP1. We then configured a PXE service point in SCCM to deploy a preconfigured Windows Sever 2008 R2 SP1 image to our target server. Lastly, we created a task sequence using specific driver and deploy packages from Dell. For specific details on installation steps and best practices for SCCM see <http://technet.microsoft.com/en-us/library/bb735860.aspx>. Below we provide the steps for configuring our specific task sequence and for advertising the new task sequence via the SCCM PXE service point.

Creating a task sequence

1. In the Configuration Manager console, navigate to System Center Configuration Manager→Site Database→Computer Management→Operating System Deployment→Task Sequences.
2. Right-click Task Sequences.
3. Select Bare Metal Server Deployment→Create a Dell PowerEdge Server Deployment Template.
4. Under Server Hardware Configuration, choose Set RAID config wizard.
5. Under Network (Admin) Account, enter the domain\user with domain admin credentials.
6. Under Operating System Installation, choose Use an OS WIM image, then for the Operating System Package select <do not select now>, and for the Package with Sysprep.inf info choose <do not select now>.
7. Click Create then click Close.
8. Right-click the task sequence and choose edit.
9. Click OK on the Missing Objects prompt.
10. Select the Set RAID Config (wizard).
11. Click the View button to open the Array Builder.
12. Right-click Each Controller, and choose Delete.
13. Click Controllers, and choose New Controller.

14. Choose Select any controller with exactly 2 disks. Click OK.
15. Right-click Embedded Controller, and choose Delete. We only want one controller rule defined.
16. Expand the newly created controller.
17. Click No variable conditions defined.
18. Click Arrays/New Array and select RAID 1: Mirrored pair, and click OK
19. At the error handling rule drop down, select Fail the task if any controller does not match a configuration rule.
20. Click OK to close the Array Builder.
21. Select Apply Operating Image.
22. Under Apply operating system from a captured image, click Browse.
23. Select the imported Operating System image file.
24. Uncheck the Use an unattended or sysprep answer file for a custom installation checkbox.
25. Click Apply Windows Settings and change Server licensing to Per server.
26. Click Enable the account and specify the local administration password, and enter credentials.
27. Set the time zone.
28. Click Apply Driver Packages.
29. Click Browse, select the appropriate x64 Dell PowerEdge R710 driver package, and click OK.
30. Click Apply Network Settings.
31. Click the star button to add a new network IP setting.
32. Enter a name for the setting based on the purpose of the first physical NIC (domain in our case).
33. Select Use the following addresses, and add an IP address for the domain.
34. Click OK.
35. Repeat steps 31-34 for the remaining three network connections, configuring storage settings for the last two NICs.
36. Click Ok to close the Task Sequent Editor.

Advertising the task sequence

1. In the Configuration Manager console, navigate to System Center Configuration Manager → Site Database → Computer Management → Operating System Deployment → Task Sequences.
2. Right-click the task sequence, and select Advertise.
3. At the General screen, click Browse and select the appropriate collection.
4. Select the Make this task sequence available to boot media and PXE box, and click Next.
5. At the Schedule screen, click the new toolbar icon to add a Mandatory assignment.
6. Select Assign immediately after this event, select As soon as possible, and click OK.
7. Click Next.
8. Select Access content directory from a distribution point when needed by the running task sequence.
9. Select When no local distribution point is available, use a remote distribution point.
10. Select When no protected distribution point is available, use an unprotected distribution point.
11. Click Next.
12. At the Interaction screen, click Next.
13. At the security screen, click Next.
14. At the Summary screen, click Next.
15. Click Close.

Microsoft Hyper-V: Running the test for scenario 4

We powered on the target host and allowed it to PXE boot with DHCP to the Windows Server 2008 R2 image from SCCM before performing the following steps. The image we created and imported into SCCM contained all roles, features, drivers, and updates preinstalled, so after the image was applied, we configured the storage and added the server to the cluster. Each step below was timed and recorded.

1. Once the system finishes booting to the new Windows Server 2008 R2 image, enter the license information, and click Next.
2. At the login screen, log into the domain, and allow the desktop to load.
3. Click Start→Administrative Tools→iSCSI Initiator.
4. Click the Discovery tab, and add the IP address for the storage group to the list of Discover Portals.
5. Click the Targets tab, and click Refresh.
6. Connect each of the 4 Hyper-V LUNs, checking the Enable multi-path box.
7. Click OK to close the iSCSI Initiator Properties window.
8. Open the Server Management console.
9. On the left side, click Storage→Disk Management.
10. Right-click each of the newly connected disks, and click Online.
11. On the left side, click Roles→Hyper-V→Hyper-V Manager.
12. On the right-hand side, click Virtual Network Manager...
13. In the Virtual Network Manager window, click New virtual network.
14. Select External, and click Add.
15. In the new virtual network properties screen, enter a name for the network that matches the other virtual network names for the other hosts in the cluster.
16. Under connection type, ensure that the external network select is the same network used for access to the domain.
17. Click Ok to finish creating the new virtual network.
18. On the left side, click Features→Failover Cluster Manager.
19. In the center pane, click Manage a cluster.
20. Enter the domain name of the Hyper-V cluster, and click OK.
21. On the left side, expand the target cluster.
22. Right-click Node, and click Add Node.
23. In the Add Node Wizard screen, enter the name of the newly deployed host, and click Add.
24. Once the server name appears in the Selected servers list, click Next.
25. On the Confirmation screen, click Next.

APPENDIX G – SCENARIO 5: PERFORMING NON-DISRUPTIVE DISASTER RECOVERY TESTING

VMware vSphere 5: Additional setup after installation for Scenario 5

For this scenario, we setup a second VMware vSphere 5 cluster with two additional Dell PowerEdge R710s and two trays of Dell EqualLogic PS Series arrays. We also setup a second vCenter server to manage the new cluster. In our lab environment, we attached this second cluster to the same internal network and assigned IPs on the same subnet as our original test bed for both storage and the cluster network.

We configured replication between our two EqualLogic array groups and allowed replication to synchronize our VMware volumes to the secondary site. With replication configured, we setup up new SQL Server database on each vCenter server to be used for VMware vSphere Site Recovery Manager (SRM) data. We then installed and configured SRM, the vCenter SRM plugin, and the Dell EQL replication adapter for SRM on each vCenter server and ensured that the two sites were communicating, both at the vCenter level and the storage level. For detailed steps on how to set EqualLogic replication and SRM, reference the document titled Disaster recovery with Dell EqualLogic PS series SANs and VMware vSphere Site Recovery Manager 5. To download the document from Dell, see the following site:

<http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=2248>. We setup a recovery plan prior to testing.

VMware vSphere 5: Running the test for scenario 5

Before running the test, we logged into the vCenter server via vSphere client and opened the Site Recovery plugin, providing the secondary site credentials. In our timed results, we include an estimate of the recovery plan maintenance cost, plus the timed results from the steps below. Since the test and cleanup require no administrator interaction or supervision, we did not time the completion of the tasks.

1. On the left-hand side, click Recovery Plans.
2. Select the test recovery plan on the left-hand side.
3. Click the Recovery Steps tab.
4. To start the test, click Test.
5. Once the test completes, click Cleanup to return the secondary site to a pre-recovery state.

Microsoft Hyper-V: Additional setup after installation for scenario 5

We also configured the Hyper-V volumes for replication using the same storage techniques on our primary and secondary sites. Refer to the document cited in the above VMware scenario for EqualLogic replication configuration. For the Hyper-V secondary site cluster, we installed and configured a Hyper-V cluster on the same Dell PowerEdge R710 hardware.

Microsoft Hyper-V: Running the test for scenario 5

In our time estimates of the Microsoft disaster-recovery test scenario, we based many of the timings on estimates, due to the fact that these situations would be driven by customized scripts and would vary from organization to organization. We did manually time the elements of the topology that were available to us, such as configuring and de-configuring SAN replication, which we provide the steps for below. We note in the body of the report our assumptions and estimates for the remaining times.

The steps below regarding SAN replication were timed and recorded and the estimated times were added.

1. Using the Dell EqualLogic web-based manager to manage the primary site storage group, click Replication.
2. Select the replication partner name, and click both Pause outbound and Pause inbound.
3. Repeat steps 1-2 while managing the secondary site storage group.
4. While still logged into the secondary site storage group management console and viewing the Replication tab, expand Inbound replicas, and right-click the first replica object.
5. Click Promote to volume.
6. On the Volume options screen, check Keep ability to demote to replica set, and click Next.
7. On the iSCSI Access screen, set the desired access settings for the volume, and click Next.
8. On the Summary screen, click Finish.
9. Once the replica has been promoted, the page will jump to Volumes.
10. Select the most recent snapshot associated with the new volume, and click the Access.
11. Click Add, and enter in the proper access settings.
12. Click OK.
13. Under Activities, click Set snapshot online.
14. On the first host in the secondary site cluster, open iSCSI Initiator, and connect to the new volume, checking the box to enable multi-path.
15. Once connect, use Disk Management to bring the new disk online, and assign a drive letter.
16. Repeat steps 14-15 on the remaining cluster host.
17. Using Failover Cluster Manager, expand the cluster name, and click Storage.
18. Click Add a disk, select the new disk, and click OK.
19. On the left-hand side, click Cluster Shared Volumes.
20. Click Add storage, select the new cluster disk, and click OK.
21. Complete steps 1-20 for each replicated volume.
22. At this point scripts would be used to boot VMs and assign IP addresses. Since we did not use these scripts, move on to the storage cleanup steps.
23. Using Failover Cluster Manager in the Cluster Shared Volumes folder, select each cluster shared volume, and click Remove from Cluster Shared Volumes.
24. Click Yes.
25. On the left-hand side, click Storage.
26. Select each cluster disk, and click Delete.
27. Click Yes.
28. Open Disk Management, and set each external storage disk to Offline.
29. Open iSCSI initiator, and disconnect from each volume.
30. Using the EqualLogic Web-based manager, log into the secondary site storage group.
31. Click Volumes.
32. Right-click each snapshot, and click Set snapshot offline.
33. Right-click each volume, and click Set offline.
34. Right-click each volume, and click Demote to replica set.
35. Click Yes.
36. Once that tasks finishes, the page will jump to replication.
37. Select the replication partner name, and click both Resume outbound and Resume inbound.
38. Repeat step 38 for the primary site storage group.

ABOUT PRINCIPLED TECHNOLOGIES



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

We provide industry-leading 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 our clients assess how it will fare against its competition, its performance, its market readiness, 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 is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. 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 PRINCIPLED TECHNOLOGIES, INC., 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 PRINCIPLED TECHNOLOGIES, INC. 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 PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.
