

LAB VALIDATION REPORT

Manageable, Affordable, and Automated Disaster Recovery Solutions

SunFire X4450 Enterprise Servers

Sun Storage 6000 Array Family

VMware vCenter Site Recovery Manager

By Brian Garrett
With Tony Palmer

May, 2009

Table of Contents

Table of Contents	i
Introduction	1
<i>The Problem.....</i>	<i>1</i>
<i>The Solution.....</i>	<i>2</i>
ESG Lab Validation.....	4
<i>Getting Started.....</i>	<i>4</i>
<i>Testing a Disaster Recovery Plan.....</i>	<i>8</i>
<i>Running a Disaster Recovery Plan</i>	<i>9</i>
<i>Failback</i>	<i>12</i>
<i>Performance Analysis.....</i>	<i>13</i>
ESG Lab Validation Highlights.....	17
Issues to Consider	17
ESG Lab's View	18
Appendix.....	19

ESG Lab Reports

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

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

Introduction


A growing number of organizations are deploying networked storage in conjunction with server virtualization to consolidate, reduce costs, and improve the flexibility and availability of mission-critical applications. IT managers choosing networked storage for virtualized environments are looking for easy and cost effective disaster recovery without sacrificing performance. This ESG Lab report presents the results of hands on testing of VMware vCenter Site Recovery Manager, SunFire X4450 servers, and SAN-attached Sun Storage 65/6780 series deployed in a highly virtualized data center.

The Problem

The use of server virtualization technology is on the rise among organizations of all sizes and in all industries around the world. In a recent ESG survey of current and planned users of server virtualization, 66% of organizations had already deployed the technology, while 16% plan to do so.¹ Given the impressive economic benefits of server virtualization, the glut of affordable and under-utilized processing power, and growing power and cooling issues in the data center, ESG predicts that brisk adoption of server virtualization will continue for the foreseeable future.

For the most critical applications, a majority of organizations (63%) cannot tolerate more than four hours of downtime without significant adverse impact to their business.² Thanks to virtualization, a growing number of mission-critical applications can now be moved to an x86 environment, but as more critical applications rely on a virtualized infrastructure, affordable and testable disaster recovery is essential.

FIGURE 1. DOWNTIME TOLERANCE



While the benefits of server virtualization and networked storage are clearly compelling, IT managers are faced with the long-standing challenge of maintaining availability as they try to manage a consolidated mix of real-world

¹ Source: ESG Research Report, *Data Center Spending Intentions*, March 2009.

² Source: ESG Research Report, *Data Protection Market Trends*, February 2008.

applications running on a virtualized infrastructure. This holds true across organizations of all sizes, regardless of the number of virtual servers deployed. Forty-six percent of virtualization users report that they currently run “Tier-1” applications on virtual machines and 33% plan to in the future.

The Solution

This report explores how servers and storage from Sun, combined with Site Recovery Manager software from VMware, can be used to automate and simplify disaster recovery testing and failover in a virtual server environment. As shown in Figure 2, ESG Lab testing was performed with SunFire X4450 Servers, Sun Storage 6000 family arrays, and VMware vCenter Site Recovery Manager.

FIGURE 2. SUN SERVERS, SUN STORAGE, AND VMWARE VCENTER SITE RECOVERY MANAGER



SunFire X4450 servers³ are designed to optimize performance, density, and power efficiency to meet the demanding requirements of highly virtualized data centers. The 2RU SunFire X4450 delivers compact, energy efficient performance and expandability.

- Up to four six-core 64bit Xeon CPUs provide the processing power needed to serve large numbers of virtualized applications in parallel.
- Six PCI-Express x8 high-performance I/O expansion slots provide the flexibility needed to meet the I/O and bandwidth requirements of a wide variety of applications.
- Standard Integrated Lights Out Manager simplifies system management and monitoring at no extra cost.

The Sun Storage 6780 array⁴ is a high-performance SAN-attached storage system architected to meet the demanding storage requirements of highly virtualized data centers. Based on a seventh generation architecture, the Sun Storage 6780 has the performance, scalability, and data reliability that customers have come to expect from this class of storage system.

- Predictably scalable and balanced application performance for virtualized environments.
- Modular, multi-dimensional scalability enables organizations to independently add or replace host interfaces, grow capacity, and add cache on the fly.
- Dual active, hot swappable controllers and non-disruptive firmware upgrades ensure high storage availability.
- Sun StorageTek Common Array Manager (CAM) software provides unified management across multiple Sun Storage product lines.
- Sun StorageTek Data Snapshot, Sun StorageTek Data Volume Copy and Sun StorageTek Data Replicator enable integrated disaster recovery.

³ Product details: <http://www.sun.com/servers/x64/x4450/>

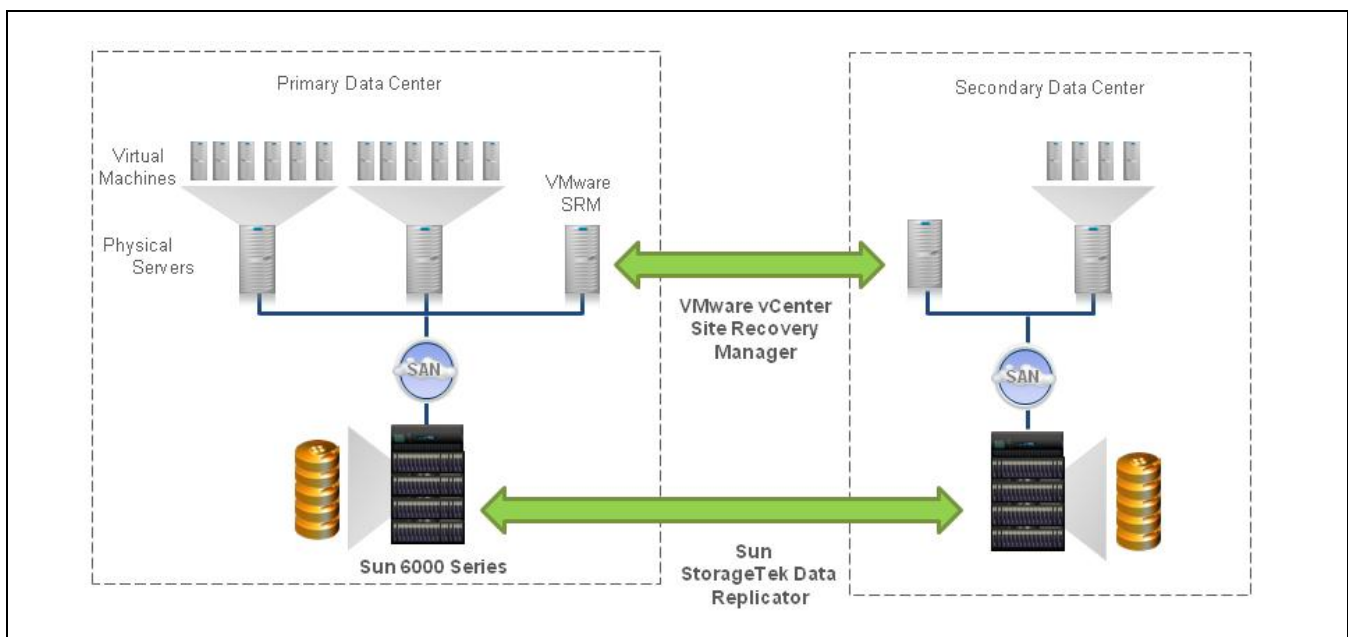
⁴ Product details: http://www.sun.com/storage/disk_systems/midrange/6780/

VMware vCenter Site Recovery Manager⁵ is an integrated element of VMware Infrastructure and VMware vSphere that works in concert with the replication capabilities of qualified storage arrays to simplify and automate disaster recovery management, non-disruptive testing, and automated failover.

- Recovery plans managed from VMware vCenter Server decrease planned and unplanned downtime.
- Automated testing of recovery plans improves business continuity.
- Single button recovery plan execution protects important applications and simplifies disaster recovery.
- Affordable implementation of disaster recovery allows companies of all sizes to protect their critical applications.

Before we get started with the presentation of how this solution was tested, let's begin with a quick explanation of how the pieces fit together. As shown in Figure 3, critical applications running on four of twelve virtual machines in a primary data center are being mirrored to a secondary data center over a wide area network (WAN). VMware software running in each data center works in concert with the field-proven remote replication capabilities of the Sun 65/6780 series disk arrays to automate the test and execution of a disaster recovery plan. More specifically, Sun StorageTek Data Replicator and Data Snapshot software running on Sun Storage 65/6780 series arrays integrate with VMware vCenter Site Recovery Manager to provide testing, automated failover, and recoverability services.

FIGURE 3. SUN 65/6780 SERIES STORAGE, INTEGRATED WITH VMWARE VCENTER SRM



Sun provides and supports a key piece of software for installation on the VMware vCenter SRM server at no additional cost: the Sun SRM Replication Adapter (SRA). It is used for communication and to pass replication commands between VMware Site Recovery Manager and the storage system. The combined solution eliminates the need for custom scripts as it simplifies defining, testing, and executing a disaster recovery plan. A wizard-driven VMware user interface automates the definition and maintenance of the storage, network, and processing resources needed to quickly resume operations at the secondary data center after a disaster.

The balance of this report represents the results of ESG Lab testing of Sun Storage 65/6780 arrays working in concert with SunFire X4450 Servers and VMware vCenter Site Recovery Manager.

⁵ VMware SRM Administration guide: http://www.vmware.com/pdf/srm_10_admin.pdf

ESG Lab Validation

Hands-on testing was performed at LSI's facility located in Beaverton, Oregon. A real-world mix of applications running in a highly consolidated virtual server environment was used to demonstrate the automated recoverability of applications at a secondary data center up to 100 kilometers (62 miles) away.

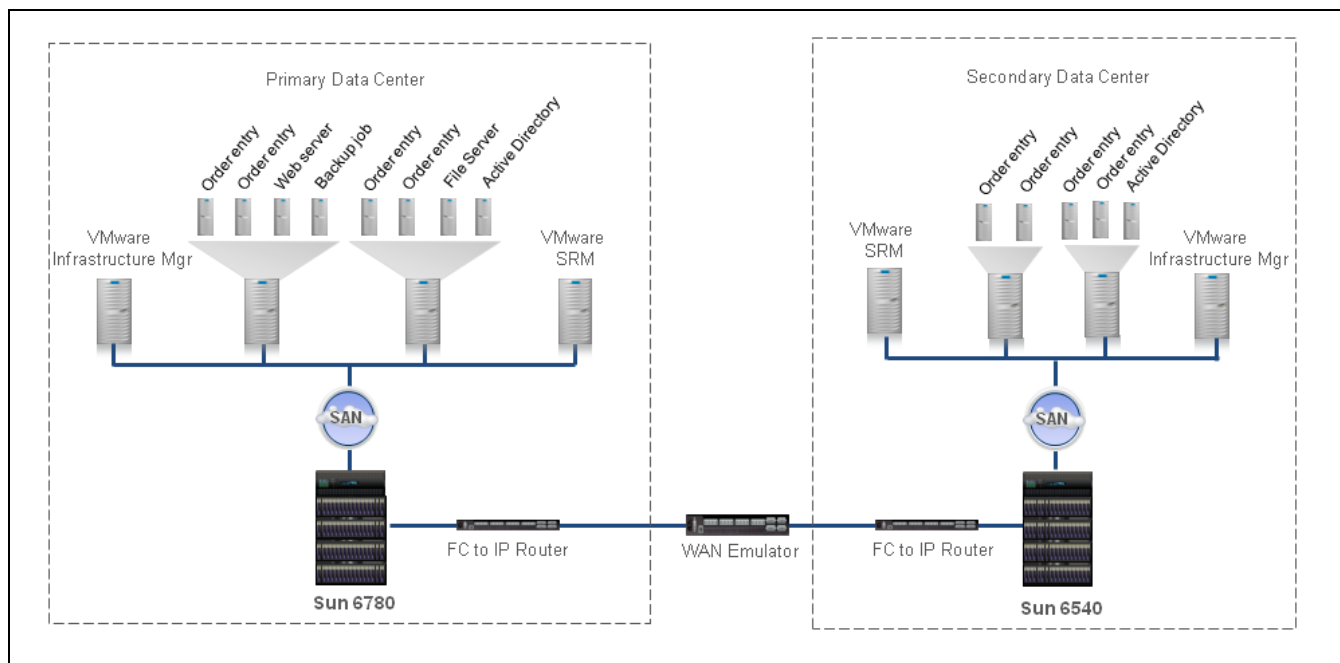
Getting Started

The test bed used during ESG Lab testing is summarized in Figure 4. Disaster recovery plans were created to manage the failover, and failback, of order entry database applications powered by Oracle 11g. Order entry applications were run in parallel with a number of other simulated real-world workloads running in virtual machines in the primary data center, including a web server, a file server, and a backup job.

A pair of SunFire X4450 servers was deployed in the primary data center. Each was equipped with two quad-core 2.13 GHz processors, 32 GB of RAM, and four dual port 4 Gbps FC host bus adapters. Two dual-core SunFire X4200 servers were deployed in the secondary data center. Single core servers running VMware Infrastructure Manager and VMware vCenter Site Recovery Manager were deployed in the primary and secondary data centers. Sun Site Replication Adapter was installed on each of the VMware vCenter SRM servers.

A Sun Storage 6780 array was deployed in the primary data center and a previous generation Sun StorageTek 6540 array was deployed in the secondary data center. Fibre Channel to IP routers were used to connect the two data centers to a wide area network simulated with an Empirix WAN emulator.⁶

FIGURE 4. THE ESG LAB TEST BED



VMware ESX version 3.5, update 3, was used to deploy virtual machines in each of the data centers. All of the virtual machines ran the Windows 2003 Enterprise x64, SP2 operating system. The Sun 6780 and 6540 arrays ran firmware version 7.30.22.

⁶ See the Appendix for more configuration details.

The Sun 6780 was deployed with eighty 300 GB 15K RPM FC drives configured as ten eight-drive RAID 1+0 groups. Thirty 30 GB LUNs were allocated from the RAID groups, which were then presented to the servers. VMware raw device mapping was used to present the thirty LUNs to the guest operating system. LUNs were striped as dynamic drives using the Windows disk administrator utility and presented to Oracle.

A variety of workloads were tested in the primary data center:

- **Order Entry Database application:** Four virtual machines used the Quest Benchmark Factory tool and Oracle 11g to simulate an online order entry database application. Each virtual machine was used to simulate the activity of 100 online interactive database users. Oracle database data was striped over seven LUNs and Oracle flash recovery area (FRA) data was striped over three LUNs.
- **Employee database application:** A custom Oracle 11g database application was used to create a simple employee database application. This application was used to verify the integrity of the Oracle database after a failover.
- **Web Server:** The industry standard Iometer utility was used to simulate web server traffic. The IO definition was composed of random reads of various block sizes as documented in the Appendix. The web server Iometer profile used for this test was originally distributed by Intel, the author of Iometer. Iometer has since become an open source project.⁷ Iometer application workloads were run against a single 30 GB LUN presented as a physical drive.
- **File Server:** The Iometer utility was used to generate web server traffic.
- **Backup job:** The Iometer utility was used to simulate a single stream of backup read traffic.

Creating a Recovery Plan

ESG Lab testing began with a review of the steps required to configure a VMware vCenter SRM disaster recovery plan:

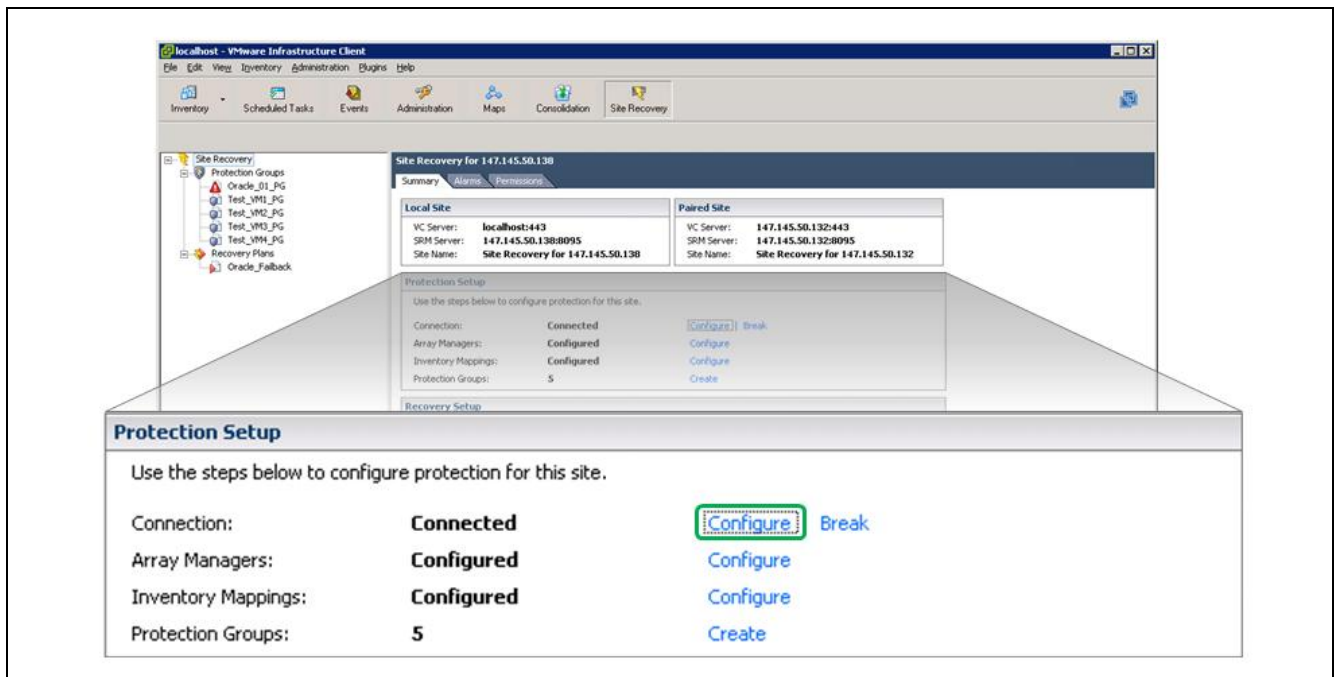
1. Install VMware vCenter SRM software on a server in each data center
2. Install Sun SRA software on the VMware vCenter SRM server in each data center
3. Launch the VMware vCenter SRM GUI to create a disaster recovery plan
4. Provide the IP address of the key components in the primary data center, including the Sun disk array
5. Provide a list of virtual machines to failover

VMware vCenter SRM does the rest; including discovering the virtual machine resources (memory, CPU, network, and drive resources) needed to automatically recreate a working application environment after a disaster.

ESG found that defining a recovery plan is an intuitive, wizard-driven process using the VMware vCenter SRM *Protection Setup* interface shown in Figure 5.

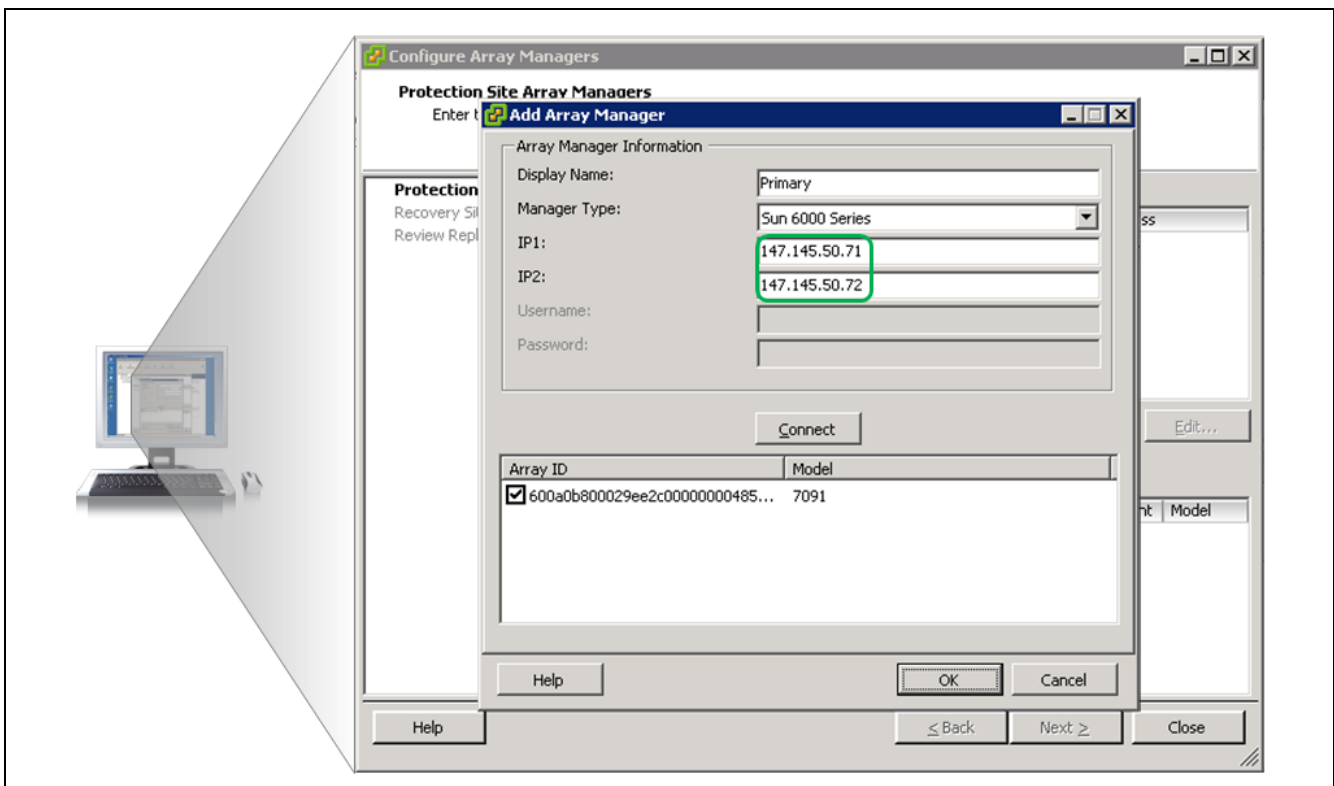
⁷ www.sourceforge.net/projects/iometer

FIGURE 5. CONFIGURING VMWARE VCENTER SITE RECOVERY MANAGER



For example, the interface shown in Figure 6 was used to define a line of communication between VMware vCenter SRM and the Sun 6780 during the *Array Manager* phase of the wizard-driven configuration process.

FIGURE 6. ADDING AN ARRAY MANAGER



Fifteen minutes after getting started with the wizard-driven VMware vCenter SRM *Protection Setup* interface, a disaster recovery plan was defined for five virtual machines running critical database applications and Microsoft Active Directory. ESG Lab found that the process was intuitive and straightforward. Similar to a wizard-driven software installation process on a Windows server, most of the time was spent reviewing automatically generated settings and hitting *Next* or *OK*. Other than deciding which virtual machines to protect (the order entry applications), the virtual machines they rely on (Microsoft Active Directory), and the IP addresses of the management interfaces used during the failover (the Sun 6780 in the primary data center), the configuration process was automated and extremely intuitive.

ESG Lab noted that the wizard-driven *Protection Setup* interface can not only be used to configure a new recovery plan, but can also be used to modify an existing plan. For example, an existing protection group can be modified to include additional virtual machines, and the dependent applications running on those machines, during a failover.

It was also noted that changes to the definition of virtual machines that are part of an existing disaster recovery plan are handled automatically by VMware vCenter SRM. For example, if more storage resources have been added to an existing virtual machine (e.g., another LUN was added to a database application), the disaster recovery plan is updated automatically to include those new resources. This simplifies the task of keeping a disaster recovery plan up to date, compared to the traditional manual process of updating scripts.

Why This Matters

IT organizations running mission-critical applications need to guard against service interruptions. An interruption could be something common, such as a server failure, a disk drive failure, a software error, data corruption, a computer virus, or “pilot” error. It could also be something more disastrous, such as a fire, a flood, a natural disaster, a pandemic, terrorism, or a blackout. As a growing number of organizations standardize on the use of virtualized environments for mission-critical applications, rapid and reliable disaster recovery solutions are needed now more than ever.

Due to the high cost of disaster recovery, multi-site disaster recovery has traditionally been reserved for mission-critical applications running on high end servers with high-end storage. These environments are typically managed with manual scripting and regular disaster recovery testing. Keeping disaster recovery scripts up to date can be a challenge as new applications and storage resources are deployed. As a growing number of organizations deploy Tier 1 applications on virtual servers, affordable and easily testable DR becomes critical.

ESG Lab found that setting up a VMware vCenter SRM disaster recovery plan with SunFire X4450 servers and Sun Storage 6000 family arrays is quick, intuitive, and automated. Fifteen minutes after getting started with a wizard driven process, the disaster recovery plan was ready for testing.

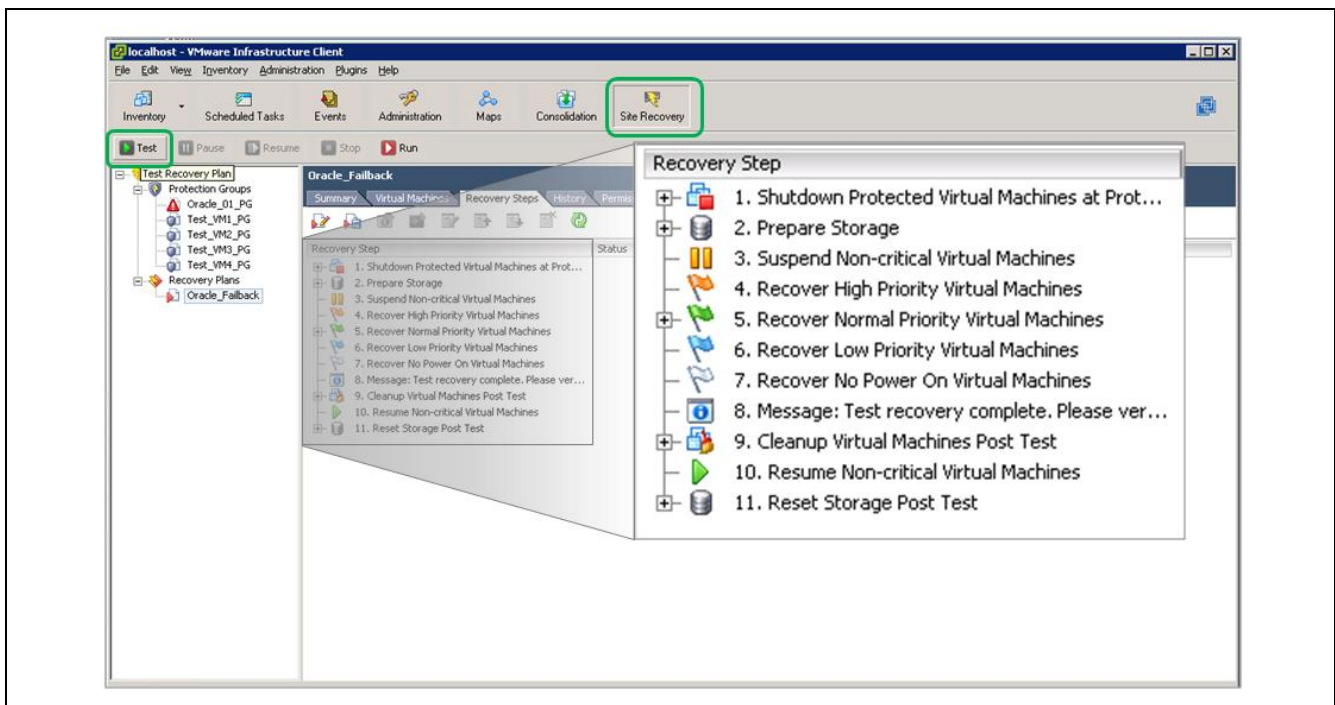
Testing a Disaster Recovery Plan

Testing a disaster recovery plan ensures that critical IT services can be resumed at an alternate site in the event of a disaster. VMware vCenter SRM, working in concert with SunFire servers and the Sun Storage 6000 array family, supports simple one-click testing of a disaster recovery plan.

ESG Lab Testing

The graphical user interface (GUI) used to monitor and control VMware vCenter SRM during ESG Lab testing is shown in Figure 7. The GUI was accessed by clicking the *Site Recovery* icon from the centrally managed VMware infrastructure client in the secondary data center. A single mouse click on the *Test* icon was all that was needed to start testing the plan.

FIGURE 7. ONE CLICK TEST OF A RECOVERY PLAN



Eighteen minutes later, the automated test of the disaster recovery plan had completed without error. Later on during the ESG Lab Validation, a test of a disaster recovery plan failed due to ESG operator error. In this case, the test failure helped rectify the fact that disk mirroring was accidentally left suspended after a performance test.

Why This Matters

Traditional disaster recovery testing is resource intensive, complex, and often fails. At a high level, the failures are due to the complexity of manually maintaining run books and scripts as server, network, and storage resources are replicated over a wide area. At a low level, disaster recovery tests usually fail due to missed resource dependencies that were added since the last disaster recovery test (e.g., a new storage device). The dynamic and fast changing nature of virtual server environments exacerbates the challenges of disaster recovery testing.

ESG Lab has validated that a VMware vCenter SRM, working in concert with SunFire servers and Sun Storage 6000 arrays, provides one-click non-disruptive automated testing of a disaster recovery plan. Non-disruptive testing allows IT to test DR during business hours—reducing both the cost and inconvenience of overnight or weekend tests.

Running a Disaster Recovery Plan

Testing a VMware disaster recovery plan verifies that the resources required to perform a failover are correctly configured and available. Running the plan actually executes the failover.

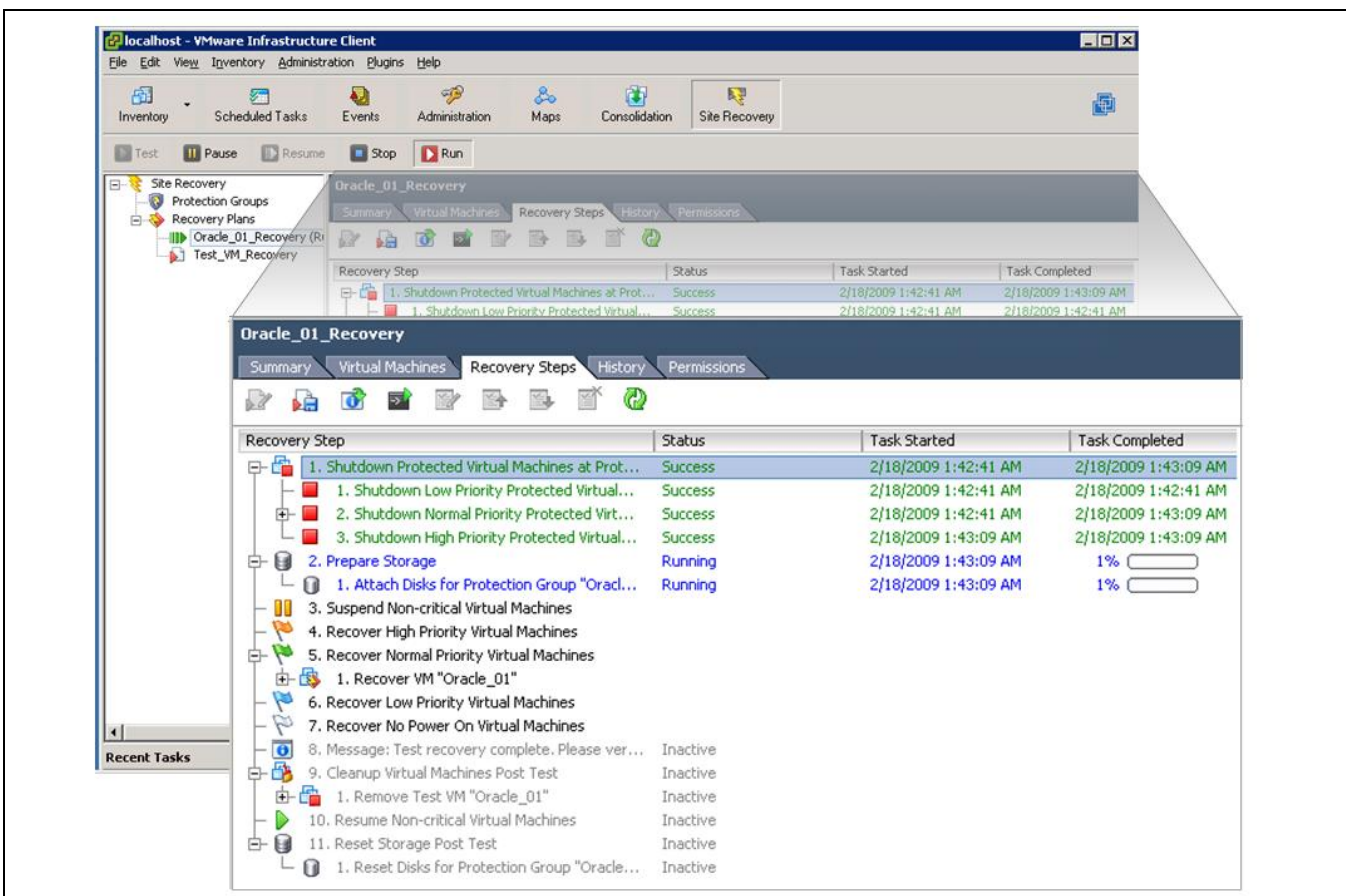
ESG Lab Testing

VMware recovery plans were run during ESG Lab testing to fail over critical Oracle database applications from a primary datacenter with a Sun 6780 disk array to a secondary data center with a Sun 6540 disk array. Two tests were performed:

1. **A failover with both data centers up and operational.** This test was performed to confirm that the virtual machines protected by a VMware recovery plan (and the database applications running in those machines) can be successfully restarted at the secondary data center.
2. **A failover after a catastrophic failure in the primary center.** Server power in the primary data center was dropped as a database application was making updates. This test was used to confirm that the applications can be restarted with no data loss after a disaster.

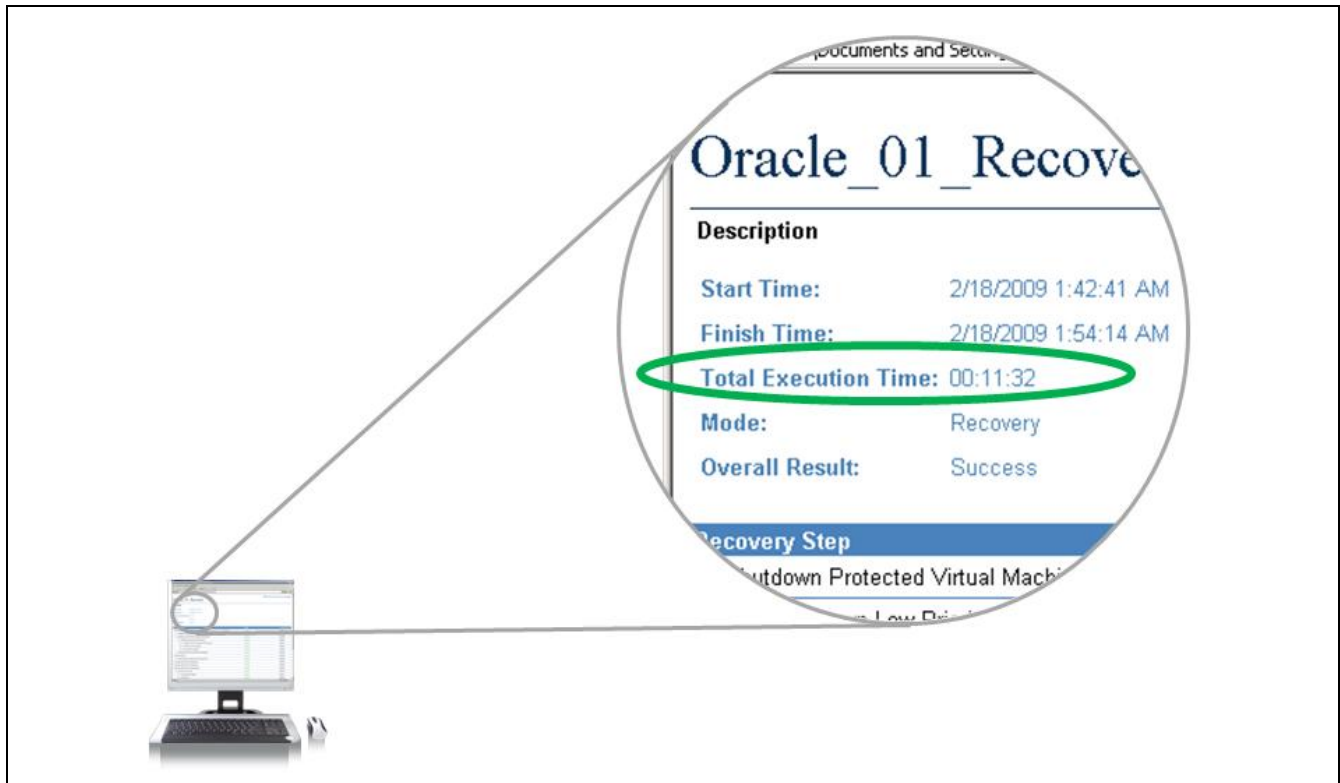
Figure 8 shows the progress of the first recovery moments after the *Run* button was clicked. ESG Lab noted that progress was very easy to follow. Successfully completed steps are depicted in green, the currently executing step is blue, and any failed steps are shown in red.

FIGURE 8. RECOVERY PLAN IN ACTION



The fully automated recovery completed in less than 12 minutes as shown in Figure 9. Three minutes after the failover had completed, an Oracle virtual machine was booted and running a query.

FIGURE 9. SUCCESSFUL FAILOVER



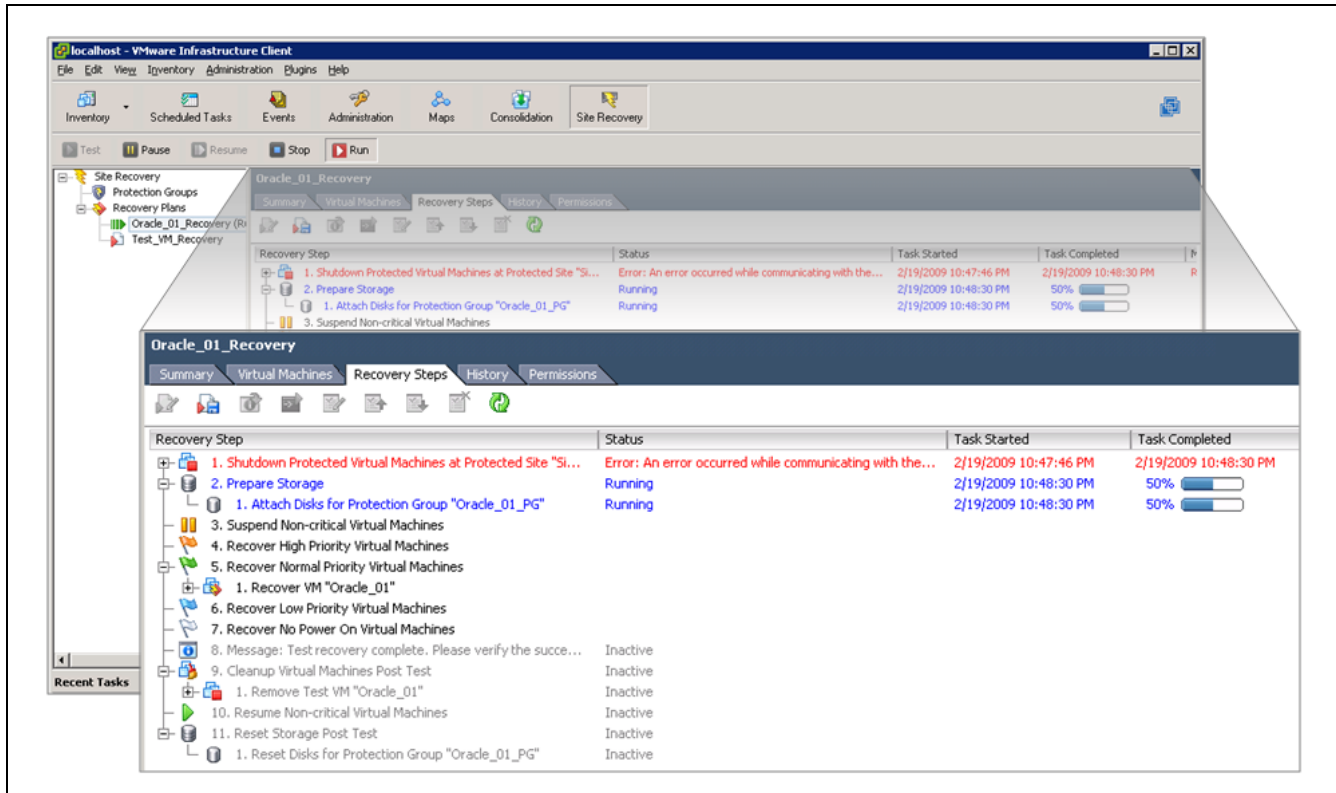
A custom Oracle application was used to verify that no data is lost after a disaster. This test relied on the fact that ESG Lab had configured Sun StorageTek Data Replicator in synchronous remote mirroring mode. Synchronous mirroring waits until writes are completed at the primary and remote sites before returning a successful status update to the application. This mode of mirroring is used to ensure that the data at a recovery site is an exact replica of the data at a production data center. More specifically, it ensures that no data will be lost if a disaster occurs.

An Oracle database with a table containing employee contact information was created. A script was used to read a pre-populated contact from an existing table and add it to a new table once every second. A separate Oracle console was used to run a SQL query, which showed rows as they were being added to the table. A sequentially incrementing employee number was used to keep track of the most recently added row.

A catastrophic power failure in the primary data center was simulated by pulling the power cables to the server as it was processing database transactions. As expected, the Oracle application stopped abruptly. The last employee added to the database before the primary site lost power was noted.

A VMware recovery plan was run. Once again, a single mouse click was all that was needed to start the recovery. A VMware screenshot taken shortly after the recovery began is shown in Figure 10. Note that there is a key difference between this recovery and the previous recovery with both data centers up and operational (Figure 8). During this test, the shutdown of the virtual machines in the primary data center failed because the primary data center had no power. This first phase of the recovery plan, which failed as expected, is shown in red.

FIGURE 10. FAILING OVER AFTER A SIMULATED DISASTER AT THE PRIMARY SITE



The fully automated recovery completed in less than 12 minutes, just as in the first test, shown in Figure 9. Three minutes after the failover had completed, an Oracle virtual machine was booted and the last employee added to the database in the primary data center before the power failure was confirmed as present in the secondary.

Why This Matters

Recovering from a disaster using traditional backup software methods can take days. Disk-based remote mirroring can be used to cut the recovery time from days to minutes, but traditional solutions which rely on scripts and manual operations can be complex. When sprinklers are running and cell phones are ringing, complex manual operations are the last thing an IT staffer needs to deal with.

ESG Lab has verified that VMware vCenter SRM, in conjunction with Sun 6780 remote replication software, can be used to automate the recovery of critical applications after a disaster. Less than fifteen minutes after a power failure in a primary data center, database transactions were up and running at a remote recovery site with zero data loss.

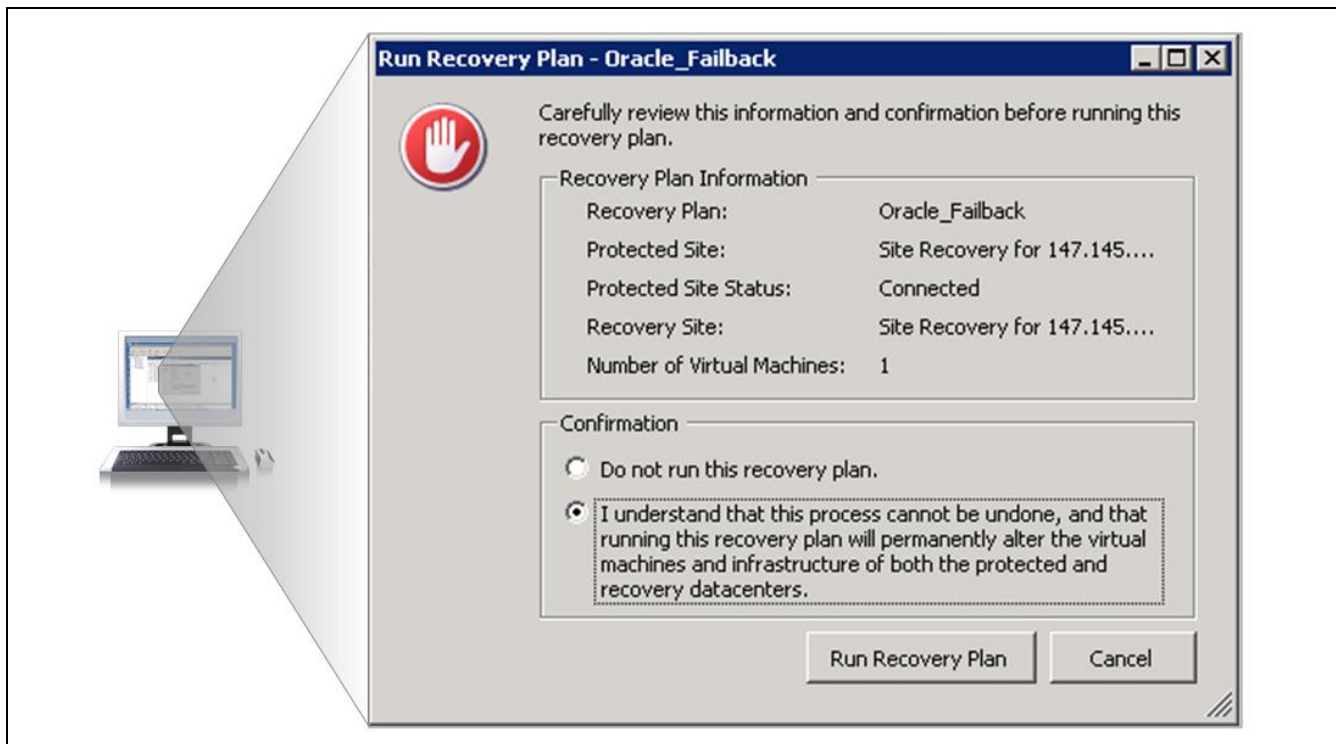
Failback

“Failback” is a term used to describe the process of restoring normal operation at a primary data center after operations have failed over to a recovery site. While VMware vCenter SRM *does not* automatically create a failback plan based on the definition of an existing recovery plan, VMware vCenter SRM *does* support the definition of a recovery plan for automated failback.

ESG Lab Testing

The servers in the primary data center were powered on after the simulated disaster and failover to the secondary site. The VMware infrastructure manager console at the primary site was used to define a new recovery plan using the wizard-driven process described earlier in this report (see Figure 5). Except for the swapping of IP addresses for the VMware and storage resources at the secondary data center, the process was exactly the same. The confirmation screen VMware presented before the failback recovery plan was initiated is shown in Figure 11.

FIGURE 11. RUNNING A FAILBACK PLAN



Less than thirty minutes after getting started with the definition of an *Oracle Failback* plan, database transactions were being serviced from the primary data center.

Why This Matters

Restoring operations in a rebuilt or new data center is typically the very last step in a disaster recovery process. It is a common misconception that VMware vCenter SRM does not support the failback capabilities needed to automate this necessary step in the recovery process. ESG Lab has confirmed that the wizard-driven process used for failover can also be used for failback.

Performance Analysis

Cost and performance are key concerns when designing and implementing a disaster recovery solution. The recurring cost of WAN bandwidth and the potentially low bandwidth of the WAN compared to LANs and SANs in the data center are at the root of a key performance concern: How will performing remote mirroring over a WAN impact application performance during normal operation? These concerns are magnified when moving to a consolidated infrastructure with a growing number of applications deployed on virtual servers attached to a shared pool of Sun Storage 6000 series arrays.

ESG Lab Testing

ESG Lab tested an order entry database application running in virtual machines with a goal of answering the following questions:

- Does remote mirroring have a noticeable impact on virtual machine CPU utilization?
- Does remote mirroring have an impact on the throughput of an order entry database application?
- How does remote mirroring affect response times for order entry users?
- What is the performance impact when the data centers are located 100 Km apart?

An Sun 6780 array was deployed in the primary data center and a previous generation Sun 6540 array was deployed at the recovery site. The drive and host interface layout for both arrays was designed with a goal of dedicating a 4 Gbps FC interface and a separate pool of drives to each virtual machine. VMware vCenter SRM, working in concert with Sun StorageTek Data Replicator, was used to create a protection plan for four of the eight virtual machines running an order entry database application. The industry standard Iometer utility was used to simulate less critical application activity on the balance of the virtual machines.

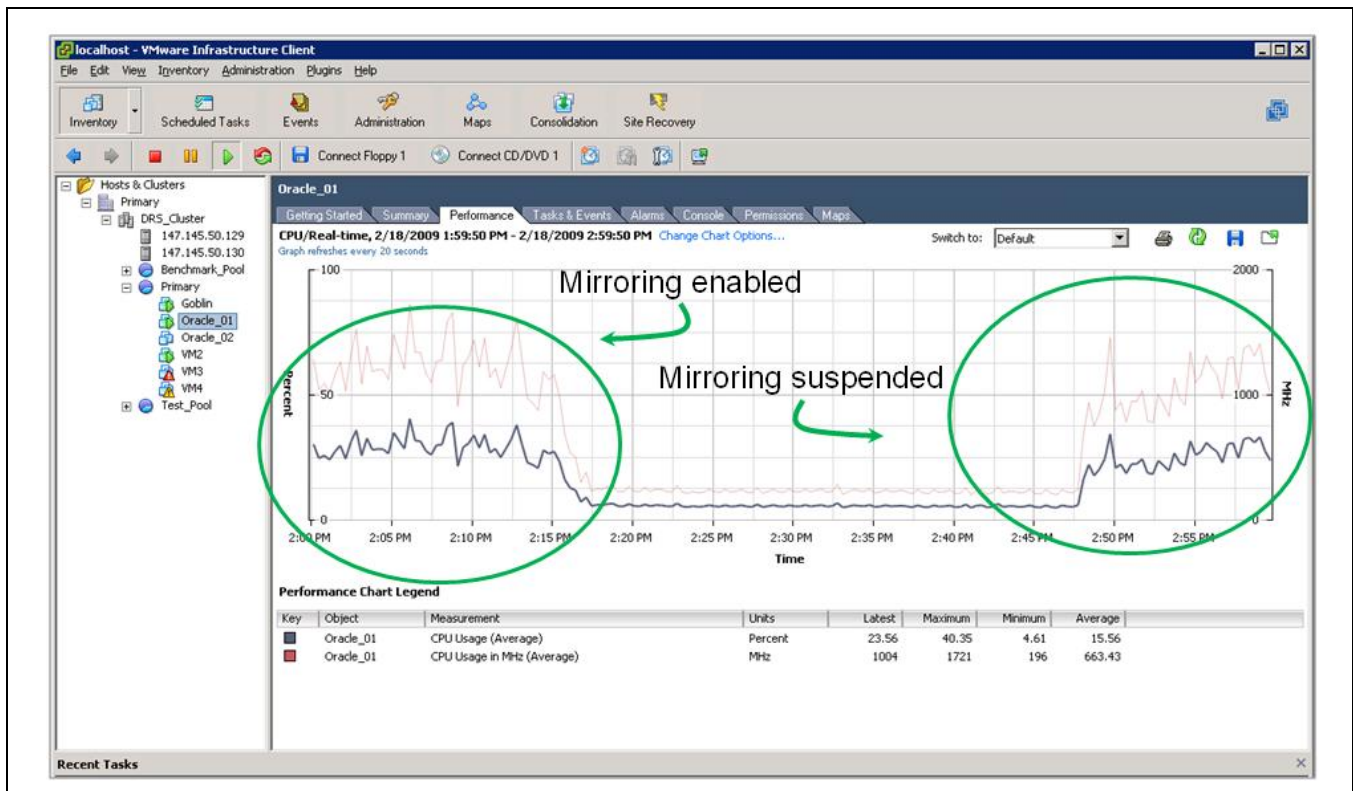
Quest Benchmark Factory software was used to simulate a typical online transaction processing (OLTP) database application. The test was designed to simulate hundreds of terminal operators executing database transactions. The application was designed to simulate the principal activities (transactions) of an order-entry system, including entering and delivering orders, recording payments, checking the status of orders, and monitoring stock levels. Like the transactions themselves, the frequency of the individual transactions was modeled after realistic scenarios.

The most frequent transaction consisted of entering a new order which, on average, was comprised of ten different items. Recording a payment received from a customer was another frequently used transaction. Less frequently executed were requests for the status of a previously placed order, a batch of ten orders for delivery, or a query of the system for potential supply shortages. The number of fully processed orders per minute and the average transaction response time was recorded.

Each of the four virtual machines running Quest Benchmark Factory and Oracle 11g was populated with a scaling factor of 1,000 and simulated 100 order entry users. Tests were run for 20 minutes with data collection beginning after a five minute ramp period. The arrays were configured for Metro Mirroring through a FC to IP router over a 1 Gbps WAN connection. An Empirix WAN emulator was used to simulate replication between two arrays located in different facilities within a metropolitan region. A baseline set of performance tests was run with mirroring suspended. Remote mirroring was enabled for the next pair of tests: the first with no WAN delay and the second with a one millisecond delay to determine the effect of replicating between two data centers approximately 100 Km (62 miles) apart.

Each of the virtual machines was configured with 2 GHz of CPU and 4 GB of RAM. The CPU utilization of the virtual machines was monitored with remote mirroring enabled and suspended while the order entry application was running. As shown in Figure 12, there was a noticeable, but manageable, difference in CPU utilization (e.g., a peak CPU utilization of 38% with mirroring suspended increased slightly to 41% with mirroring enabled).


FIGURE 12. MINIMAL CPU UTILIZATION OVERHEAD



Four hundred simulated order entry users were tested within four virtual machines as lometer workloads ran in parallel. lometer was used to simulate less critical web server, file server, and backup jobs running in virtual machines that were not being replicated. While the amount of IO activity generated by these less critical applications was significant, it had no discernable impact on the number of transactions per minute that the order entry applications could sustain.

Transactions per minute as reported by the Benchmark Factory utility were compared as tests were run with mirroring suspended (no mirroring), mirroring enabled (mirroring within a data center), and mirroring enabled with a one millisecond WAN delay (mirroring over 100 Km). The results are summarized in Figure 13.

FIGURE 13. MINIMAL DATABASE TRANSACTION OVERHEAD




What the Numbers Mean

- The number of order entry database applications that can be processed per minute is a measure of the amount of work that the application infrastructure can sustain.
- In this example, the amount of work that can be done by 400 users serviced by Oracle 11g order entry applications running four virtual machines (VM1, VM2, VM3, VM4) is compared.
- The amount of work that could be performed by each virtual machine drops slightly when mirroring is enabled (0.41% when mirroring within the data center, 0.61% when mirroring over 100 Km).

The average transaction response time per order entry transaction was recorded with mirroring enabled and disabled. The results are summarized in Figure 14.

FIGURE 14. SUB-SECOND RESPONSE TIMES



What the Numbers Mean

- The average transaction time increases as expected as database writes are replicated synchronously to a remote recovery site.
- While transaction response times increased noticeably (111% when mirroring within the data center, 174% when mirroring over 100 Km), all of the transaction response times recorded by ESG Lab were less than half a second, which can't be perceived by an end-user.
- Given the value of being able to recover a critical application at a remote site minutes after a disaster, higher sub-second response times and an aggregate performance impact of less than 1% as recorded during ESG Lab testing are minimal and manageable.

Why This Matters

ESG research indicates that performance is a key concern when deploying applications in a highly consolidated environment. With multiple applications relying on a shared infrastructure, there is a concern that performance requirements can't be met predictably. As a matter of fact, 51% of ESG survey respondents that have already deployed virtual servers connected to networked storage report that performance is their top concern.⁸ Replicating virtualized applications over a wide area network increases the risk and concern. If you care enough to replicate an application, you probably care about its performance as well. Protection at the expense of customer satisfaction is not a viable option.

For the configuration and workloads tested by ESG Lab, remote mirroring introduced a measurable, but imperceptible, performance impact for an online order entry application with hundreds of interactive users.

⁸ Source: ESG Research Report, *The Impact of Server Virtualization on Storage*, December 2007.

ESG Lab Validation Highlights

- ☑ Fifteen minutes after getting started with a wizard-driven VMware process, a disaster recovery plan was defined for virtualized applications running on SunFire X4450 servers and a Sun 6780 disk array.
- ☑ Sun StorageTek Data Replicator, configured in Metro Mirror mode, provided real-time replication of critical data to a recovery site located 100 Km away.
- ☑ Testing a disaster recovery plan using Sun StorageTek Data Replicator and Sun StorageTek Data SnapShot was started with a single mouse click from a VMware GUI.
- ☑ Fifteen minutes after starting an automated recovery, Oracle 11g was up and running in a virtual machine at the remote site.
- ☑ No database transactions were lost after a simulated disaster.
- ☑ A VMware recovery plan was used to fail back to the primary data center.
- ☑ Remote mirroring to a recovery site 100 Km away had no perceptible performance impact during benchmark testing with 400 simulated order entry database users.
- ☑ The Sun StorageTek Data Replicator and Sun StorageTek Data SnapShot software were very easy to setup and manage for both sites using Sun StorageTek Common Array Manager (CAM).

Issues to Consider

- ☑ Oracle 11g was used during ESG Lab testing as an example of a critical application that could be protected with VMware vCenter SRM and Sun StorageTek Data Replicator in a virtualized environment. While the configuration and methodology presented in this report was not meant to document best practices, similar testing can be performed with applications that you feel are critical in your environment (e-mail, for example).
- ☑ While customers around the world have deployed Oracle on VMware ESX Server, there is a lot of misinformation in the market about Oracle's VMware support policies. If you are considering Oracle in a VMware environment, ESG Lab recommends that you refer to Oracle Metalink (<https://metalink.oracle.com>) or other authorized Oracle sources for the official support policy specific to the Oracle products you want to virtualize.
- ☑ The performance results presented in this document are based on industry-standard benchmarks deployed in a controlled environment. Due to the many variables in each production data center environment, capacity planning and testing are needed to validate the configuration of storage systems and wide area networks when implementing a disaster recovery solution in your environment.

ESG Lab's View

Server virtualization is being deployed by a growing number of organizations to lower costs, improve resource utilization, provide non-disruptive upgrades, and increase availability. Each of these benefits is fundamentally enabled by de-coupling servers, applications, and data from specific physical assets. Storage virtualization takes those very same benefits and extends them from servers to the underlying storage domain—bringing IT organizations one step closer to the ideal of a completely virtualized IT infrastructure.

As a growing number of organizations harness the power of server and storage virtualization to move closer to that vision, IT managers are concerned about their ability to recover virtualized infrastructure after a disaster. The changing nature of virtual server environments poses challenges for organizations that have historically relied on manual scripting for remote mirroring and disaster recovery. The complexity of keeping track of application dependencies increases dramatically in a fast growing virtual server environment. And once you think you have it under control, it can change in an instant as applications move within the virtual infrastructure. VMware vCenter Site Recovery Manager, working in concert with Sun StorageTek Data Replicator, is a simply elegant disaster recovery solution expressly designed with these concerns in mind. A simple GUI interface is all that is needed to control the definition, testing, and execution of a disaster recovery plan.

ESG Lab has decades of experience working with disk-based disaster recovery solutions. Compared to the manual scripting and complexity we've come to expect from traditional solutions, we were amazed by the automated simplicity of VMware vCenter SRM. Configuring a complex recovery plan involving multiple virtual machines was wizard-driven and easy. Fifteen minutes after starting a failover with a single mouse click, the entire application environment was up and running at a remote recovery site. Thirty minutes later, the same wizard-driven process was used to fail back to the primary data center. From an end-user perspective, the recovered environment felt exactly the same—regardless of the data center delivering services. The network addresses, logins, and operating system preferences were the same. No application data was lost. And last, but not least, for the Oracle 11g order entry application tested by ESG Lab, there was no noticeable difference in application performance—even as data was replicated to a recovery site 100 Km away.

If your company has been on the sidelines, crossing fingers in hopes that a disaster never strikes, ESG Lab recommends that you consider VMware vCenter Site Recovery Manager, SunFire x64 Servers, and the Sun Storage 6000 family of arrays with Sun StorageTek Data Replicator—an automated disaster recovery solution for virtual server environments.

Appendix

TABLE 1. TEST BED OVERVIEW

Database	Oracle 11g
Database workload generator	Quest Benchmark Factory, version 5.5
General purpose I/O workload generator	Iometer version 2008-10-23
Physical servers (primary site)	Two SunFire X4450 servers (2.13 GHz dual proc/quad core, 32 GB RAM)
Physical servers (remote site)	Two AMD Opteron 252 (3 GHz dual core, 64 GB RAM)
Fibre Channel Storage Area Network (SAN)	Four FC switches: Brocade Silkstorm 4100 4 Gbps Eight FC HBA's: QLE2462 dual port 4 Gbps
Wide Area Network (WAN)	Two Brocade 7500E FC to IP SAN routers One Empirix WAN emulator
Server Virtualization software	VMware ESX Server 3.5, update 3
Sun Site Replication Adapter	Version 01.00.32.13
Guest operating systems	Windows 2003 Enterprise x64 Edition, SP 2 4GB of RAM per VM
Fibre Channel Storage Array (primary site)	Sun 6780, firmware 7.30.22
Fibre Channel Storage Array (remote site)	Sun 6540, firmware 7.30.22
Logical drive configuration	Eighty 300 GB 15K RPM FC drives configured as ten 30 GB RAID-1 LUNs which were presented to Oracle. Each 30 GB LUN was comprised of eight 300 GB Drives in a 4+4 RAID-10 configuration. VMware raw device mapping was used to present thirty LUNs to the guest operating system. LUNs were striped as dynamic drives using the Windows disk administrator utility. Oracle database data was striped over seven LUNs, Oracle FRA data was striped over 3 LUNs. Iometer LUNs were configured in the same manner (30 GB RAID-1 LUNs). Each Iometer workload ran against a single LUN presented as a physical drive within a Windows 2003 virtual machine.

TABLE 2. IOMETER WORKLOAD DEFINITIONS

Web Server	One worker, sixteen outstanding IOs, 100% random reads of assorted block sizes: 22% 0.5 KB, 15% 1 KB, 8% 2 KB, 23% 2 KB, 15% 8 KB, 2% 16 KB, 7%, 74 KB, 1% 128 KB, 1% 6512 KB
Backup job	One worker, one outstanding IO, 80% sequential reads, 100% 8 KB
File Server	One worker, sixteen outstanding IOs, 80% random reads of assorted block sizes: 10% 0.5 KB, 5% 1 KB, 5% 2 KB, 60% 4 KB, 2% 8 KB, 4% 16 KB, 4% 32 KB, 10% 64 KB



20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com