WHITE PAPER

VMware vCenter™ Site Recovery Manager Performance and Best Practices for Performance

Architecting Your Recovery Plan to Minimize Recovery Time

vmware®

Table of Contents

Introduction

About VMware vCenter™ Site Recovery Manager

VMware vCenter™ Site Recovery Manager provides business continuity and disaster recovery protection for virtual environments. Protection can extend from individual replicated datastores to an entire virtual site.

In a Site Recovery Manager environment, there are two sites involved – a protected site and a recovery site. Protection groups[1] that contain protected virtual machines are configured on the protected site and these virtual machines can be recovered by executing the recovery plans on the recovery site.

Site Recovery Manager leverages array based replication between a protected site and a recovery site. The workflow that is built into Site Recovery Manager automatically discovers which datastores are setup for replication between the protected and recovery sites. Site Recovery Manager provides protection for the operating systems and applications encapsulated by the virtual machines running on an ESX host. A Site Recovery Manager server must be installed at the protected site and at the recovery site. The protected and recovery sites must each be managed by their own VMware vCenter™ Server.

Performance Considerations in a Site Recovery Manager Environment

Recovery Point Objective (RPO)[2] and Recovery Time Objective (RTO)[3] are the two most important performance metrics IT administrators need to keep in mind while designing/executing a disaster recovery plan.

RPO is fulfilled by the storage provider whose storage replication adapters work in conjunction with Site Recovery Manager to enable a simple and a fully automated test/real recovery. On the RTO front, Site Recovery Manager provides capability for IT administrators to minimize recovery time for a datacenter recovery – which is crucial for any business continuity/disaster recovery solutions.

About This Paper

The goal of this whitepaper is to provide you with Site Recovery Manager performance data and recommendations so that you can architect an efficient recovery plan that minimizes the recovery time for your environment.

This whitepaper addresses various dimensions on which the recovery time depends:

- Number of virtual machines and protection groups associated with a recovery plan
- Virtual machine to protection group relation
- Recovery site performance in a cluster
- Configuration of various recovery plan parameters
- Priority assignment of virtual machines in the recovery plan
- High latency network between protected and recovery sites

Furthermore, best practices in applicable areas are suggested so that you can optimize the recovery time using Site Recovery Manager.

[1] A protection group is a group of virtual machines that failover together. After the protection groups are created at the protected site, they (and their virtual machines) must also be added to recovery plans on the recovery site to complete the Site Recovery Manager setup.

[2] Recovery Point Objective (RPO) describes the acceptable amount of data loss measured in time.

[3] The Recovery Time Objective (RTO) is the duration of time and a service level within which a business process must be restored after a disaster (or disruption) in order to avoid unacceptable consequences associated with a break in business continuity.

Reference Setup Environment

Below is the setup used for all Site Recovery Manager experiments presented in this white paper.

Note: Site Recovery Manager does not impose similar hardware requirements across both sites. You can have different number of ESX hosts at protected and recovery sites.

In this particular setup, both Site Recovery Manager and VMware vCenter™ Server were installed on the same physical machine (this applies to both the protected and the recovery site).

Figure 1 Illustration of the Site Recovery Manager testbed environment

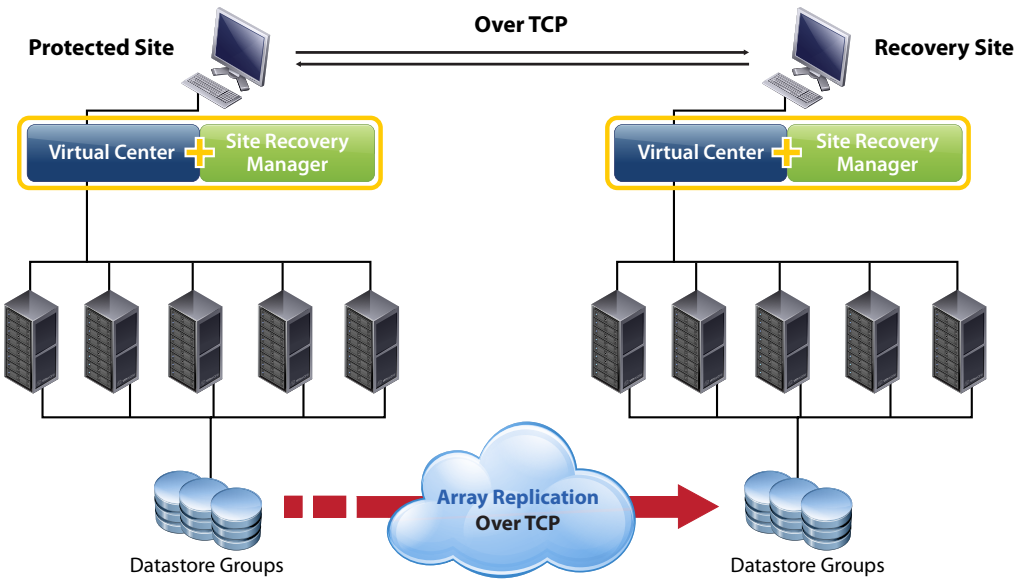

Hardware/Software Configuration

Experimental Setup

Site Recovery Manager 1.0 U1, VMware vCenter™ 2.5 U4, and VMware® ESX™ 3.5 U4 were used for performance measurements. WANem 2.1 was used for simulating a high-latency network with packet drops.

Site Recovery Manager and VMware vCenter™ Server – Protected Site & Recovery Site Configuration

Host computer: Dell Power Edge R900
CPUs: 2x Quad Core E7310 Xeon, 1.6GHz
RAM: 20GB
Network: Broadcom BCM5708C NetXtreme II GigE
Site Recovery Manager Server software: Site Recovery Manager 1.0 U1
VMware vCenter Server software: VMware vCenter 2.5 U4

ESX System – 5 ESX 3.5 hosts on Protected Site, 5 ESX 3.5 hosts on Recovery Site

Host computer: Dell PE2650
CPUs: Intel Xeon CPU 3.06 GHz
RAM: 8GB
Network: NetXtreme BCM5703 Gigabit Ethernet 100 mbps
ESX Software: VMware ESX 3.5 U4

Network Simulation Software
WANem v2.1

Storage
LeftHand SANIQ 7.0 VSA nodes - SRA_7.0.01.6066

Site Recovery Manager Server Configuration Recommendation and Database Sizing

Below are the minimum hardware resource requirements for Site Recovery Manager 1.0 on both sites:

- Processor – 2.0GHz or higher Intel or AMD x86 processor
- Memory – 2GB minimum
- Disk Storage – 2GB minimum
- Networking – Gigabit recommended

Site Recovery Manager uses a database on both protected and recovery sites to store information. The protected site Site Recovery Manager database stores data regarding the protection group settings and protected virtual machines, while the recovery site Site Recovery Manager database stores information on recovery plan settings, results for testing recovery plans, results for running a real recovery, transactions made during a running test/real recovery and much more. Some of the disk space usage is permanent in nature while some of it is transient, like the space required for temporary transactional data during a running recovery.

You can use the same database server to support the VMware vCenter™ database instance and the Site Recovery Manager database instance.

Database size is dependent upon:

- Number of protected virtual machines
- Number of protection groups
- Number of recovery plans
- Transient data written during test and real recoveries
- Extra steps added to the recovery plan
- Etc.

Refer to the Tools section in the Site Recovery Manager resources page for more information on how to size your Site Recovery Manager database for Oracle and SQL Server based on the above mentioned parameters.

Basic Operation Latency Overview

VMware vCenter™ Site Recovery Manager revolutionizes the way disaster recovery plans are designed and executed. This involves two simple steps: protection and recovery.

Protection involves the following operations:

- Array manager configuration
- Inventory mapping
- Creating a protection group

Recovery involves the following operations:

- Creating a recovery plan
- Test recovery
- Real recovery

Figure 2 depicts average baseline latencies for each operation.

Note: The actual numbers can vary in a real deployment and the numbers presented here are from a specific setup.

Figure 2 Basic Site Recovery Manager operations – latency[4] overview

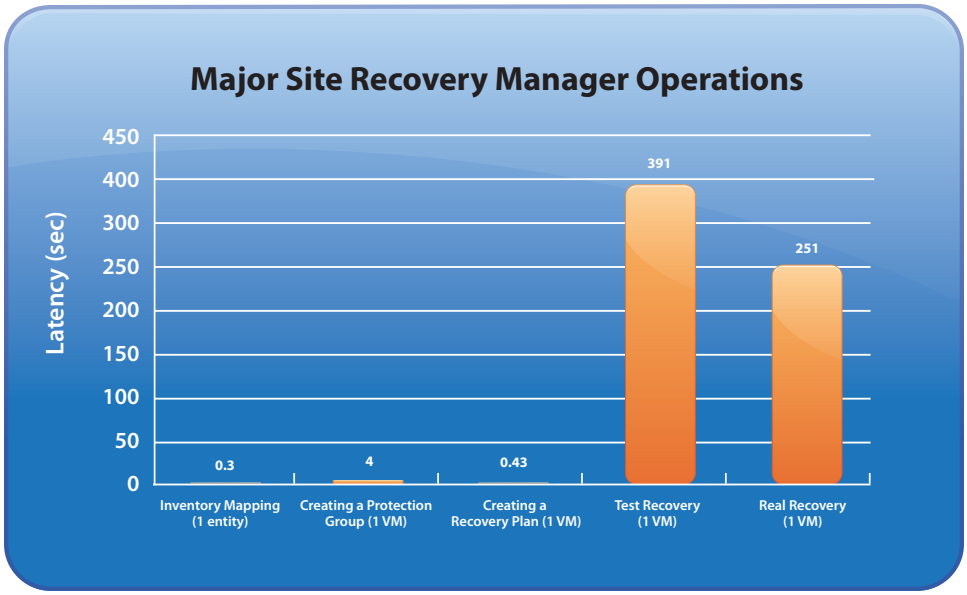

Creating a protection group: Protecting protected site virtual machine(s) and creating placeholder virtual machine(s) on the recovery site.

Inventory mapping: Mapping inventory entities like networks, compute resources, and virtual machine folders between the protected and the recovery site.

Creating a recovery plan: Creating a recovery plan with protection group(s) consisting of protected virtual machine(s).

Test recovery: Executing a Test recovery.

Real recovery: Executing a Real recovery.

Test Recovery Time vs. Real Recovery Time

With any disaster recovery solution, you need to have an estimation of the recovery time. In addition to providing real recovery, Site Recovery Manager supports a non-disruptive test recovery mode. The empirical data gathered from the setup show that real recovery time is capped by test recovery time. Test recovery time is higher than the real recovery time as the former involves reverting the datacenter back to its original state. This involves recovery site operations like powering off test virtual machines, resetting the storage, and replacing test virtual machines with placeholder virtual machines (see Figure 3). A real recovery on the other hand includes steps not executed in a test recovery, which are powering off production virtual machines at protected site (if protected site is still available, for example, in case of a planned migration).

[4] In this whitepaper, latency is defined as the time taken for a certain operation to finish.

Figure 3 Steps executed only in test recovery

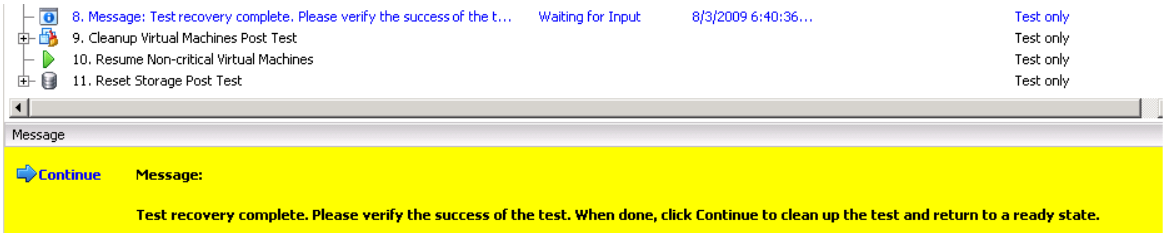

Figure 4 Steps executed only in real recovery (in case of datacenter migration)

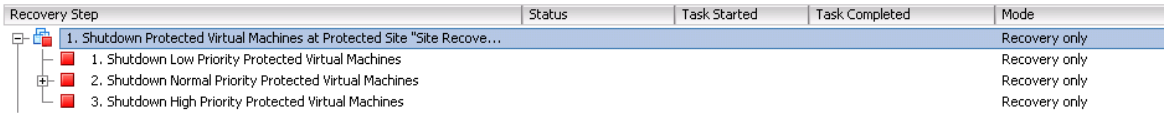

As a general rule, you can gauge the real recovery time by measuring the time taken for a test recovery until the message prompt step (step 8 in Figure 3). In case of planned datacenter migration, you need to account for the time taken to shutdown the virtual machines (step 1 in Figure 4) on the protected site as well.

Estimating an accurate RTO is relatively difficult as there are a lot of variable parameters to consider – network latencies, resource availability, and storage I/O to name a few. **Note:** Recovery operation on a storage level may take more or less time for a test recovery as opposed to a real recovery. Latencies for this operation are storage dependent.

Site Recovery Manager Scalability

All the following operations were executed over a low latency network.

Note: These numbers are just for reference and will differ for each environment.

Protection (Scaling With a Number of Virtual Machines)

Site Recovery Manager allows protecting a maximum of 500 virtual machines. Site Recovery Manager scales well with an increasing number of virtual machines in a protection group. The majority of time for this operation is spent in creating placeholder virtual machines on the recovery site.

Figure 5 Protection group creation time: scaling with virtual machines under a single protection group

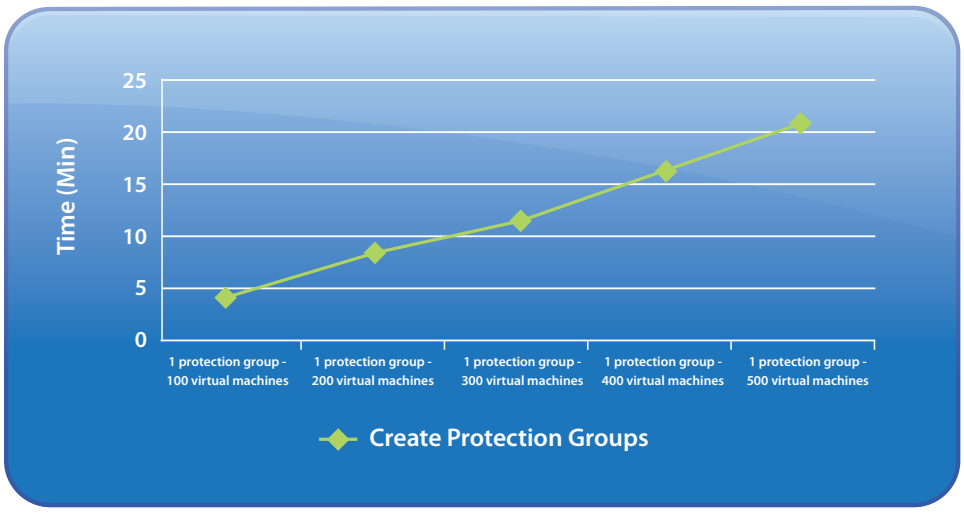

Recovery (Scaling With a Number of Virtual Machines and Scaling With a Number of Protection Groups)

Site Recovery Manager allows recovering a maximum of 500 virtual machines and a maximum of 150 protection groups. As mentioned previously, test recovery time is higher than real recovery time. And the same can be observed by scaling the number of virtual machines.

Figure 6 Test and real recovery time: scaling with virtual machines under a single protection group

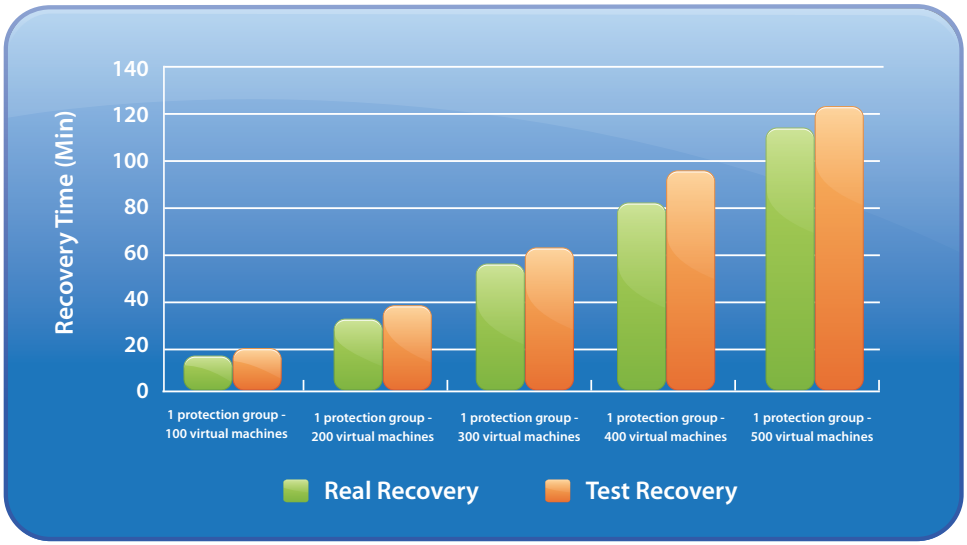

Figure 7 Test recovery time: scaling with protection groups (1 virtual machine per protection group)

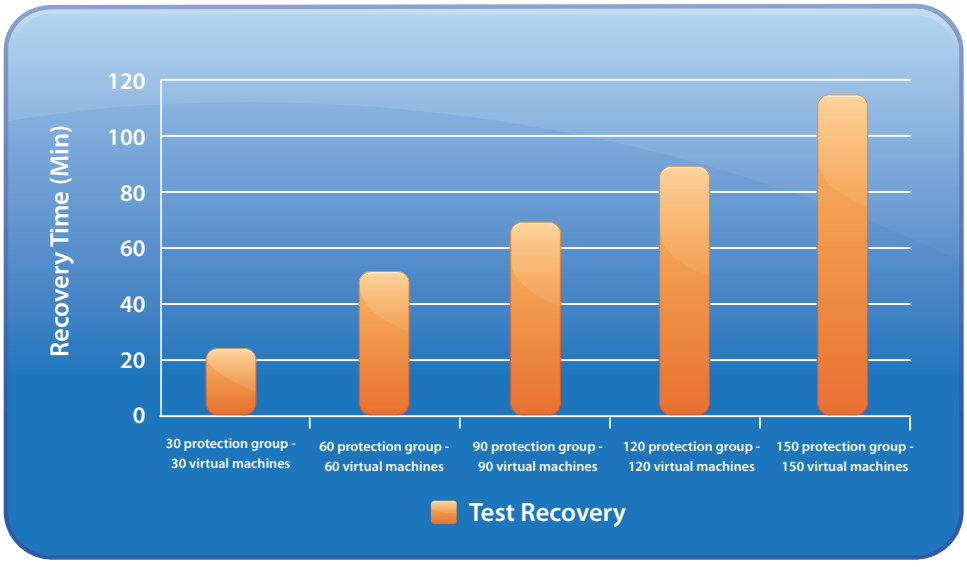

Note: Site Recovery Manager supports running three simultaneous recovery plans.

High Latency Network

Here is the setup used for simulating a high latency network across both sites.
WANem 2.1 was used to simulate a high latency network with different Round Trip Times (RTT).

Figure 8 Illustration of the Site Recovery Manager Testbed Environment with a high latency network

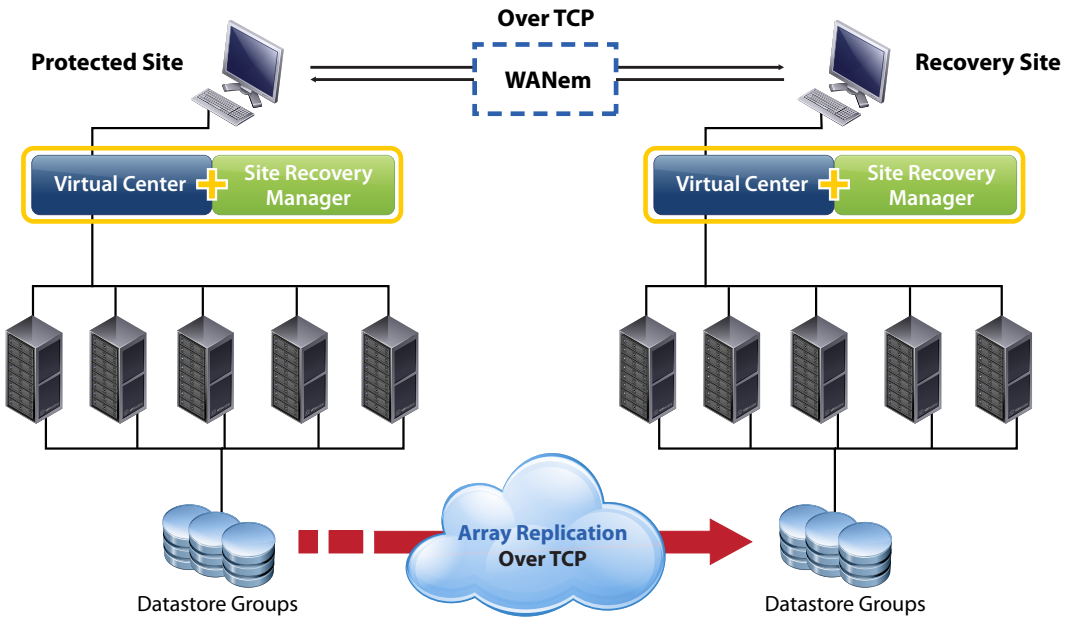

Network RTTs are typically 75 to 100 milliseconds for intra-U.S. networks, about 250 milliseconds for transatlantic networks, and 320 to 430 milliseconds for satellite networks.

Latencies for creating protection groups and real recoveries show some impact with a high latency network between the protected and the recovery site.

The following experiments were carried out with RTTs of 100, 250, 400 milliseconds.

Creating Protection Groups on a High Latency Network

Latencies for creating protection groups are affected by a high latency network between the protected and the recovery site.

Creating a protection group involves creation and monitoring of placeholder virtual machines on the recovery site – this involves a number of Remote Procedure Calls (RPC).

Figure 9 Creating protection groups on a high latency network with different RTTs between protected and recovery site

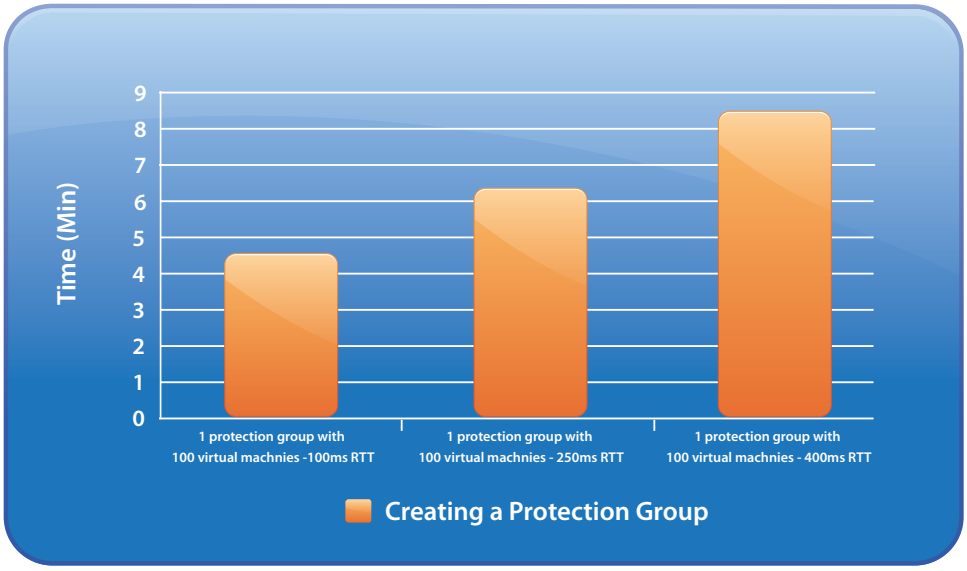

Recoveries on a High Latency Network

Real recoveries are affected by high latency networks because of the operation to shutdown remote virtual machines on the protected site. This applies to planned migration only and not to real disasters (as there would be no access to the protected site virtual machines in this case). Test recoveries should not be affected by a high latency network between a protected and a recovery site. Thus, network latency between the protected and the recovery sites needs to be considered when estimating real recovery time.

Figure 10 Test and real recovery time on a high latency network with different RTTs

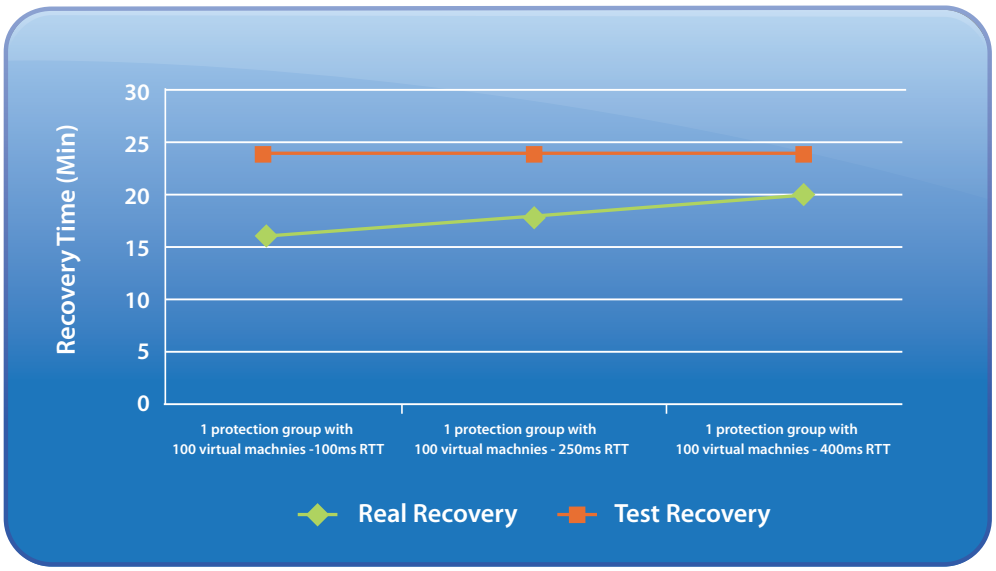

Architecting Recovery Plans (From a Performance/Recovery Time Perspective)

In this section, Site Recovery Manager best practices in certain areas are provided to help you architect efficient recovery plans that minimize recovery time.

Virtual Machine to Protection Group Relation

For each protection group included in a recovery plan, Site Recovery Manager needs to communicate with the underlying storage to create snapshots of replicated LUNs (or promote replicated LUNs in case of real recovery) in that protection group and present them to recovery site hosts. Currently, this operation is executed sequentially for each protection group. As a result, adding more protection groups to a recovery plan as opposed to adding more virtual machines is more costly from a recovery time perspective. This is also dependent upon the underlying storage used for replication between the two sites.

Compared below are recovery time measurements for a test recovery for 150 virtual machines in a single protection group vs. 150 virtual machines in 150 protection groups (1 virtual machine per protection group).

Figure 11 Virtual machine to protection group relation

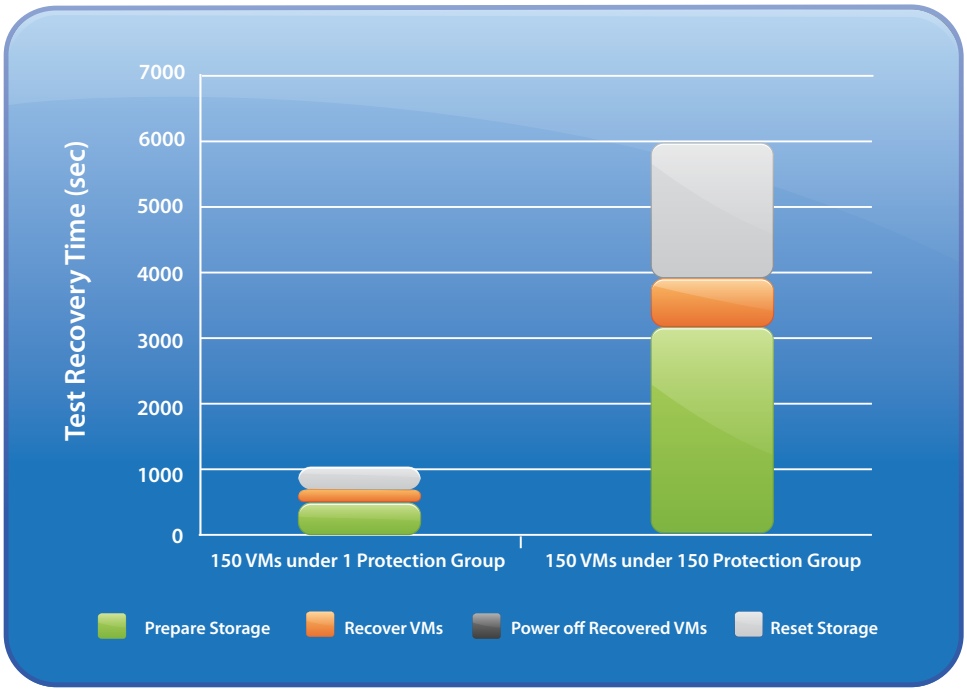

As shown in Figure 11, for a specific environment, the former case outperforms the latter. These numbers can change significantly depending on the environment.

Note: The operation to power off all recovered virtual machines in this experiment took an average of 58 seconds for both cases mentioned in Figure 11. It is not visible in Figure 11 as it constitutes a small percentage of the overall recovery time.

Key Takeaways:

- Grouping virtual machines in fewer protection groups enables faster test and real recoveries, provided those virtual machines have no constraints preventing them from being grouped under similar protection groups.
- The above recommendation applies to both test and real recoveries.

  Note: The step to reset storage is not a part of the real recovery.

- Adding a couple of protection groups does not necessarily increase the recovery time by a large factor. The recovery time overhead of adding a single protection group with a single virtual machine is more than adding a single virtual machine under an existing protection group.

Enabling VMware® Distributed Resource Scheduler (DRS) on the Recovery Site

Site Recovery Manager creates placeholder virtual machines at the recovery site for each protected virtual machine on the protected site. During a recovery (test and real) each of these placeholder virtual machines are replaced by a virtual machine registered from the recovered datastore.

If VMware DRS is enabled on the recovery cluster, then Site Recovery Manager will utilize VMware DRS to load balance the recovered virtual machines across the hosts to ensure optimal performance and RTO.

The graph below compares test recovery times with VMware DRS enabled and VMware DRS disabled on the recovery site. The recovery plan consists of one protection group and 100 virtual machines. A worst-case scenario is simulated for the VMware DRS disabled case, where all virtual machines are recovered on a single host. For the VMware DRS enabled case, VMware DRS is set to automatic mode with "moderate" aggressiveness. VMware DRS then intelligently recommends placement for virtual machines being recovered, which improves recovery time significantly because of less contention on the ESX Servers on the recovery site.

Figure 12 Recovery time benefits with VMware DRS: recovering 100 virtual machines under a single protection group

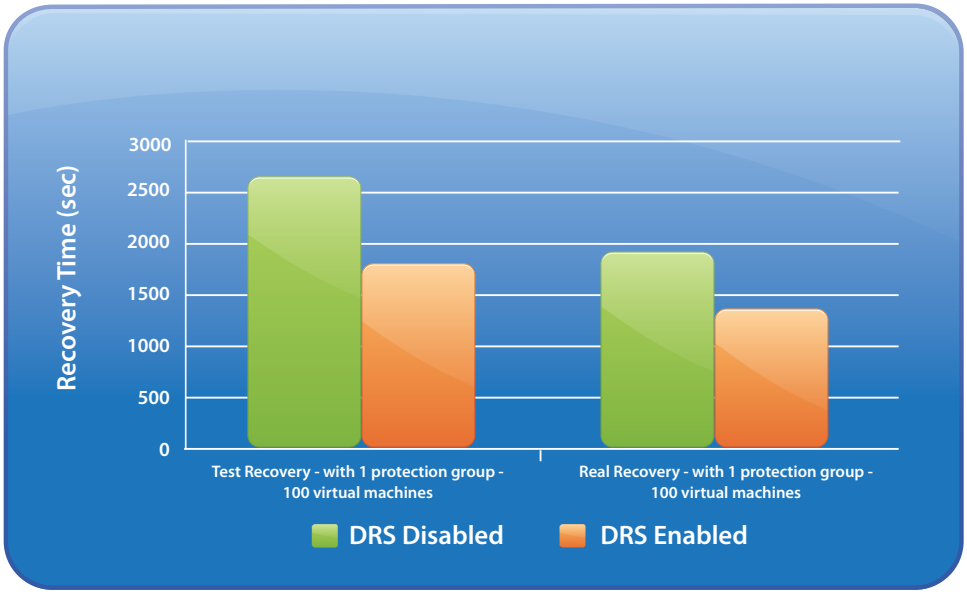

Key Takeaways:

- It is a good practice to have VMware DRS enabled on a recovery site.
- If VMware DRS is not enabled on a cluster, then it is a good practice to manually distribute placeholder virtual machines evenly across hosts. This will help in distributing the load across hosts when recovered virtual machines are powered on and will in turn improve performance and recovery time. To do this, drag and drop the placeholder virtual machines across desired hosts.

High Priority Virtual Machines and Suspending Virtual Machines

In a recovery plan, the virtual machines being recovered can be assigned as high, normal, or low priority virtual machines.

The assignment of the order of these virtual machines and the total number of virtual machines placed in each category will affect the recovery time.

For example, shown in Figure 13 are the recovery times for executing a test recovery with increasing number of high priority virtual machines.

Figure 13 Test recovery: varying the number of high priority virtual machines

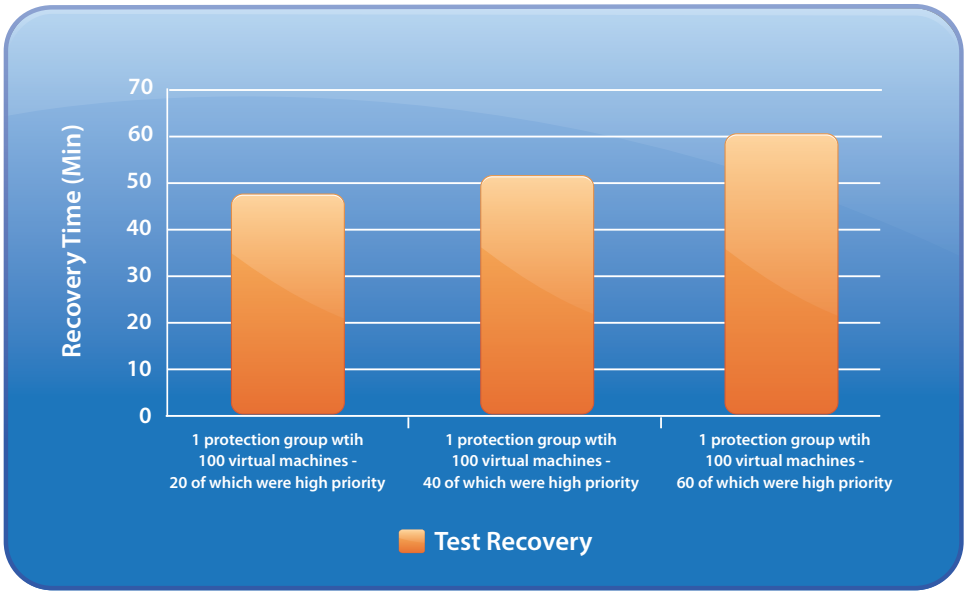

As you can see, the recovery time does increase with the number of virtual machines placed in high priority, as all high priority operations to recover virtual machines will be executed sequentially. This applies to both real and test recoveries.

Key Takeaways:

- It is important to chart out the dependencies/priorities between virtual machines to be recovered here so that only a certain number of required virtual machines can be assigned as high priority. It does impact recovery time – test as well as real recovery.
- Similarly, suspending virtual machines on the recovery site will also impact recovery time.

VMware Tools

The installation of VMware tools in all protected virtual machines is strongly recommended. Many Site Recovery Manager recovery operations depend on the proper installation of VMware Tools in the protected virtual machines.

- Wait for OS heartbeat while powering on virtual machine and wait for network change while reconfiguring recovered virtual machine.

  Site Recovery Manager depends on VMware Tools to report OS heartbeat and finish of network change. If you do not have VMware Tools installed on any of the protected virtual machines, you can choose to set the timeout values for **Wait for OS Heartbeat** and **Wait for Network Change** to zero (0).

- Wait for virtual machines to shutdown on the protected site.

  During a real recovery, Site Recovery Manager tries to gracefully shutdown the virtual machines on the protected site. Before Site Recovery Manager forcibly powers a virtual machine off, it tries to shutdown the guest OS. If your intention is to power off the virtual machines without gracefully shutting down the guest OS, you can set the timeout value called **powerStateChangeTimeout** in the vmware-dr.xml configuration file to zero. See below for an example:

```
<Recovery>
<powerStateChangeTimeout>0</ powerStateChangeTimeout>
</Recovery>
```

  VMware vCenter Site Recovery Manager Service needs to be restarted after making this change in the configuration file.

  Note: This applies only for virtual machines with VMware Tools installed and the timeout is automatically set to zero (0) if no VMware Tools are installed.

Recommendations

VMware vCenter™ Site Recovery Manager provides advanced capabilities for disaster recovery management, non-disruptive testing, and automated failover. The following performance recommendations have been made in this paper:

- Grouping virtual machines under fewer protection groups enables faster test and real recoveries, provided those virtual machines have no constraints preventing them from being grouped under similar protection groups.
- Enable VMware DRS at recovery site. This will ensure optimal performance and recovery time as VMware DRS will load balance the recovered virtual machines across the hosts.
- Manually distribute placeholder virtual machines evenly across recovery hosts if VMware DRS is not enabled. This will help in distributing the load across hosts for all recovery virtual machine operations, which will in turn improve performance and recovery time.
- It is important to chart out the dependencies between, and priorities for, virtual machines to be recovered so that only a certain number of required virtual machines can be assigned as high priority. It does impact test and real recovery time. Similarly, suspending virtual machines on the recovery site will also impact recovery time.
- It is strongly recommended that VMware Tools be installed in all protected virtual machines in order to accurately acquire their heartbeats and network change notification.
- Make sure any internal script or call out prompt does not block recovery indefinitely.

Appendix

Job Throttling During a Recovery

To alleviate resource pressure on Site Recovery Manager Server and vCenter Server, Site Recovery Manager performs job throttling by limiting the maximum number of concurrent operations initiated during a recovery.

Most of the sub-steps in a recovery plan are comprised of units of execution, each of which is dispatched to a single ESX host via VMware vCenter™ Server. The following sub-steps support concurrent executions:

- Shutdown protected virtual machines on the protected site (real recovery only)
- Suspend non-critical virtual machines
- Recover normal priority virtual machines
- Recover low priority virtual machines
- Recover no power on virtual machines
- Resume non-critical virtual machines (test recovery only)
- Cleanup virtual machines post test (test recovery only)

At any time, the number of sub-steps executed concurrently is the minimum of 16 and the number of ESX hosts used for recovery.

Acknowledgements

Thanks to Arturo Fagundo, David Levy, Maria Basmanova, Site Recovery Manager team, John Liang, Rajit Kambo, Lee Dilworth, Desmond Chan, Michael White, Site Recovery Manager field staff, and VMware Technical Marketing for their input and reviews on various drafts of this paper.

About the Author

Aalap Desai is a MTS Performance Engineer at VMware. He has been working on the performance project for VMware vCenter™ Site Recovery Manager. Aalap received his Masters in Computer Science from Syracuse University.

References

- VMware vCenter™ Site Recovery Manager Documentation:
  http://www.vmware.com/support/pubs/srm_pubs.html
- VMware vCenter™ Site Recovery Manager Evaluator's Guide:
  http://www.vmware.com/pdf/srm_10_eval_guide.pdf
- VMware vCenter™ Site Recovery Manager Resources for Business Continuity:
  http://www.vmware.com/products/srm/resource.html
  Tools for sizing the Site Recovery Manager Database based on the inventory size can be found under the "Tools" section.

VMware vCenter Site Recovery Manager 1.0
Performance and Best Practices for Performance
Source: Performance Engineering, AD
Revision: 20090731

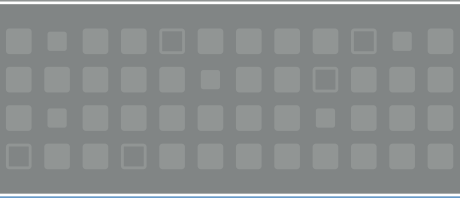

VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

 VMW_09Q3_vCenter_Performance_WP_P16_R1

vmware®