VMware vCenter™ Site Recovery Manager 4.0 Performance and Best Practices for Performance

Architecting Your Recovery Plan to Minimize Recovery Time

WHITE PAPER

vmware®

Table of Contents

Introduction

About VMware vCenter Site Recovery Manager 4.0

VMware vCenter™ Site Recovery Manager (SRM) 4.0 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 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 datastores setup for replication between the protected and recovery sites. Site Recovery Manager provides protection for the operating systems and applications encapsulated by virtual machines running on a VMware ESX™ host. A Site Recovery Manager server must be installed both at the protected and recovery site. The protected and recovery sites must each be managed by their own vCenter Server.

Furthermore, VMware vCenter Site Recovery Manager 4.0 supports VMware vSphere™, shared recovery site, and NFS.

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 and executing a disaster recovery plan.

The storage provider whose storage replication adapters work in conjunction with Site Recovery Manager to enable a simple and a fully automated test or real recovery fulfills the RPO. 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 and disaster recovery solution.

About This Paper

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

This white paper addresses various dimensions on which the recovery time depends:

- Recoveries with iSCSI, FC, and NFS storage
- Number of virtual machines and protection groups associated with a recovery plan
- Virtual machine to protection group relation
- Recovery site performance in a cluster with DPM and DRS
- 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 are suggested in applicable areas 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 Server hosts at protected and recovery sites.

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

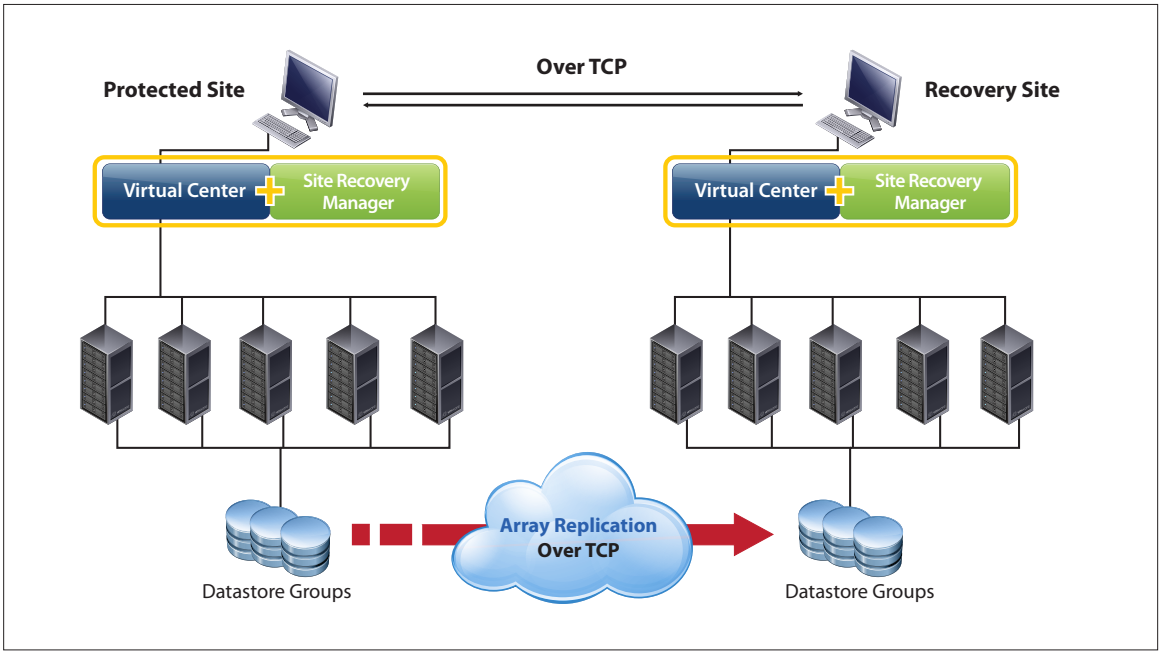

Fig 1. Illustration of the Site Recovery Manager Testbed Environment

Hardware/Software Configuration

Experimental Setup

Site Recovery Manager 4.0 , VMware vCenter™ 4.0 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 and 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 4.0

VMware vCenter Server software: VMware vCenter 4.0

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 8.0 VSA nodes - SRA_8.0.00.1682

Site Recovery Manager Server Configuration Recommendation and Database Sizing
Below are the minimum hardware resource requirements for Site Recovery Manager 4.0 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 4.0 uses 400KB RAM per protected virtual machine for the recovery site during recoveries.

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 test or real recovery run, and more. Some of the disk space usage is permanent in nature while some of it is transient—e.g. the space required for temporary transactional data during a running recovery.

It is recommended that the Site Recovery Manager database be installed as close to the Site Recovery Manager server as possible—as it reduces the Round Trip Times (RTT) between both of them. This way, recovery time performance will not suffer greatly because of shorter round trips to the database server.

You can use the same database server to support the 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 mentioned parameters.

Summary of performance improvements made in Site Recovery Manager 4.0 over Site Recovery Manager 1.0

The following improvements are representative of the performance data collected from the lab environment described in the Introduction section. The scale of improvement may vary in your lab environment as it depends on the scale of your inventory[4] and setup configuration[5].

UI improvements

- UI responsiveness improvement for Inventory Mapping screen with a large scale inventory
- UI responsiveness improvements for creating protection groups with a large number of virtual machines
- UI performance speedup for displaying large recovery plans

Recovery time improvements

- Increased parallelism for recovering virtual machines - 2 "recovery virtual machine" operations started per host instead of 1.
- Noticeable performance improvement for deleting swap files during recoveries—available with ESX 4.0 U1
- Noticeable performance improvement for preparing storage during recovery plan execution via efficient database access for a large number of virtual machines
- Noticeable performance improvement for all concurrent sub-steps during recovery plan execution via efficient database access for a large number of virtual machines
- Noticeable performance improvement in searching datastores during recovery plan execution—available with VC 4.0 U1
- Distribution of shadow virtual machines across hosts on the recovery site during protection improves recovery time performance.

Other

- Noticeable performance improvement in recovery plan creation

Test and real recovery time performance has improved in Site Recovery Manager 4.0 when compared to Site Recovery Manager 1.0 because of the recovery time improvements mentioned above.

---

4 Inventory in reference includes, but not limited to, virtual machines, protection groups, datastores, clusters, network and folders.

5 Setup configuration in reference includes, but not limited to, hardware configuration and RTT between Site Recovery Manager Server and Site Recovery Manager database.

1. 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 following operations:

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

Recovery involves following operations:

- Creating a recovery plan
- Test recovery
- Real recovery

The following graph depicts average baseline latencies for major operations.

Note: The actual numbers can vary in a real deployment and the numbers presented here are from a specific setup. The virtual machines used for all the data presented in this white paper did not have any IP customization specs or any waiting for heartbeats during the recovery. The motivation was to exercise Site Recovery Manager behavior during a recovery with different settings.

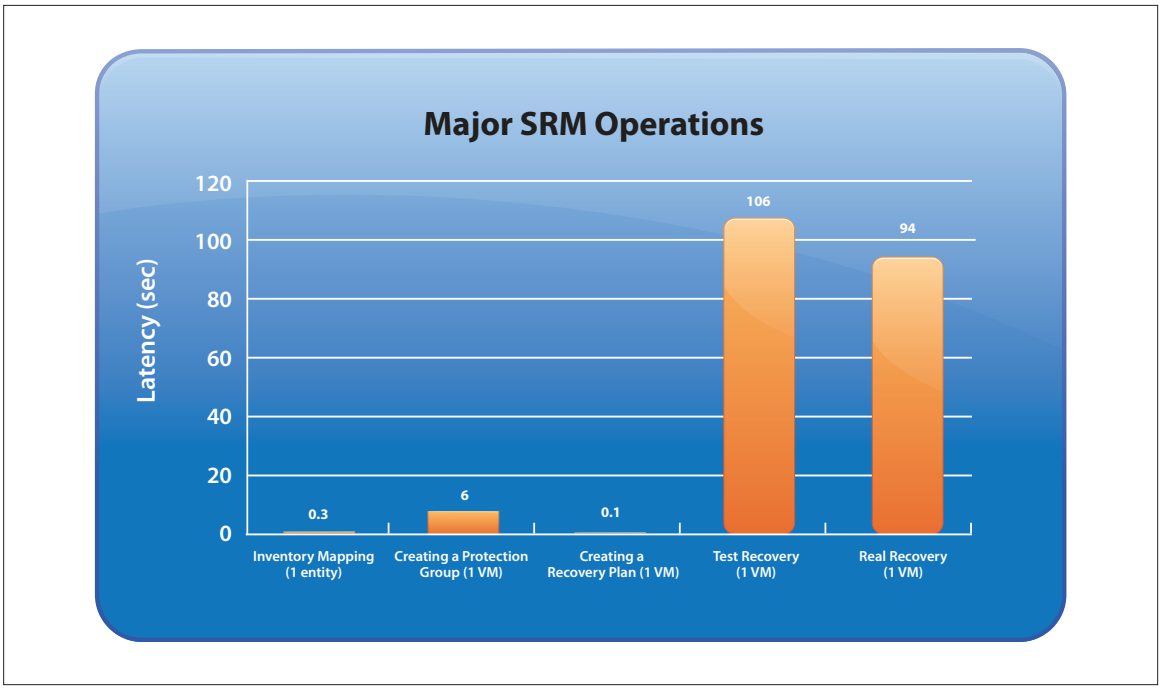

Fig 2. Basic Site Recovery Manager operations—latency[6] overview

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.

6 In this white paper, latency is defined as the time taken for a certain operation to finish.

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.

2. 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, such as powering off production virtual machines at protected site (if protected site is still available, for example in case of a planned migration).

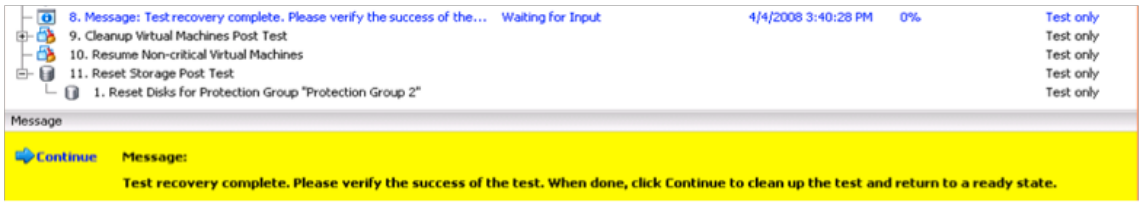

Fig 3. Steps executed only in test recovery

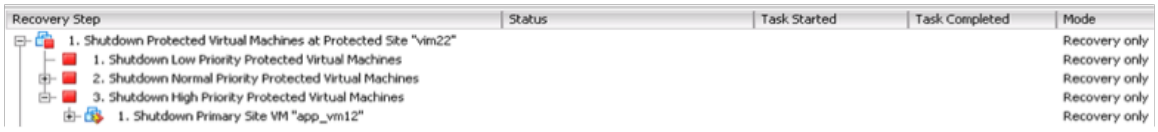

Fig 4. Step executed only in real recovery (in case of datacenter migration)

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, such as network latencies, resource availability, and storage I/O. For instance, recovery operation latency during the "Prepare Storage" step in a test recovery differs from the latency during a real recovery since the two recovery modes entail different storage level operations.

3. Site Recovery Manager Scalability Performance

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

3.1. Protection (Scaling with a number of virtual machines)

Site Recovery Manager 4.0 allows protecting a maximum of 1000 virtual machines. Site Recovery Manager scales well with an increasing number of virtual machines in a protection group. Majority time for this operation is spent in creating placeholder virtual machines on the recovery site.

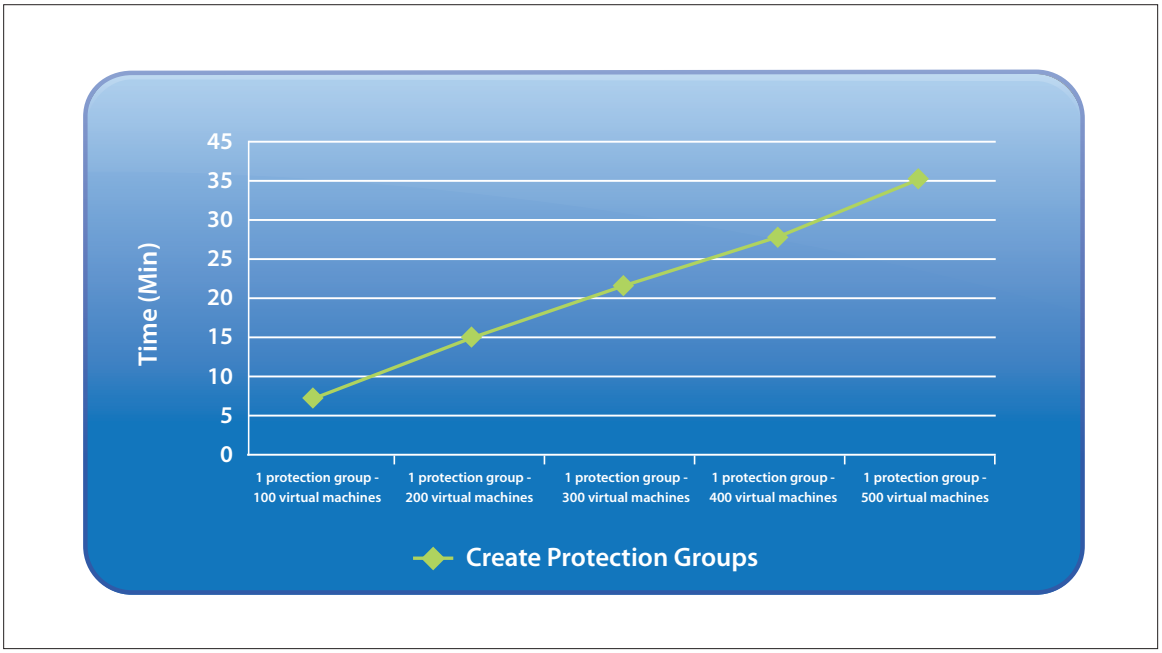

Fig 5. Protection group creation time: scaling with virtual machines under a single protection group

3.2. Recovery (Scaling with a number of virtual machines and scaling with a number of protection groups)

Site Recovery Manager 4.0 allows recovering a maximum of 1000 virtual machines and a maximum of 150 protection groups. As mentioned previously, test recovery time is higher than real recovery time. The same can be observed by scaling the number of virtual machines.

Site Recovery Manager 4.0 provides support for replicated NFS storage as well as iSCSI and FC. In figure 6 and 7, you will find the statistics we gathered using software iSCSI replicated storage.

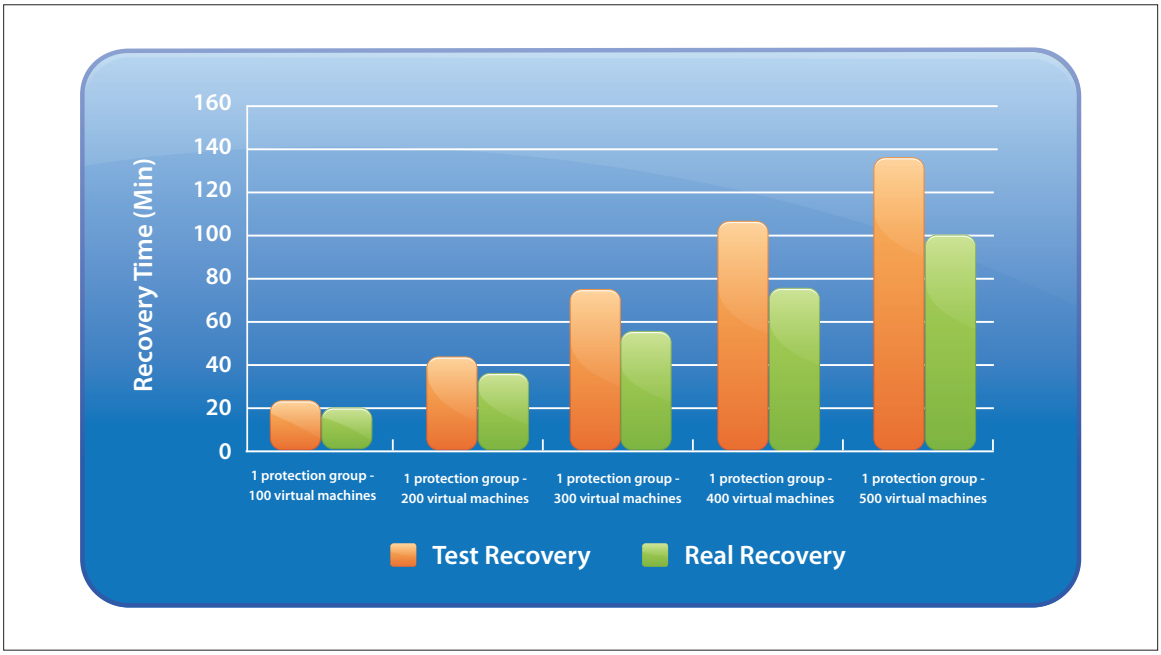

Fig 6. Test and real recovery time: scaling with virtual machines under a single protection group—Software iSCSI

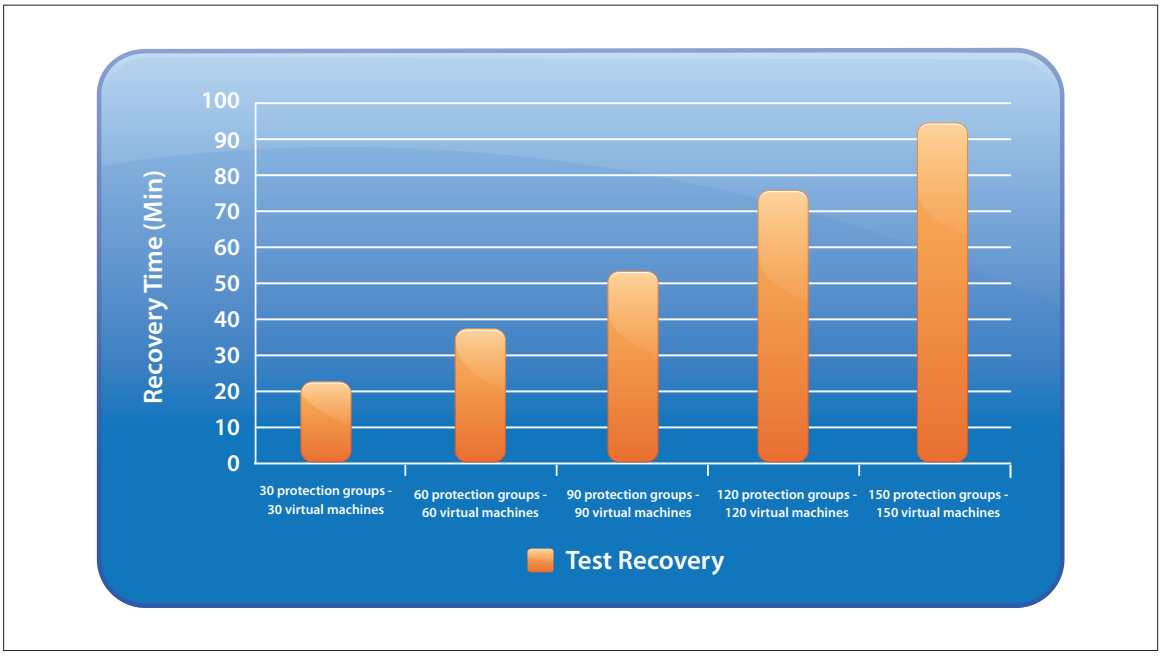

Fig 7. Test recovery time: scaling with protection groups (1 virtual machine per protection group)

3.3. 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 RTT.

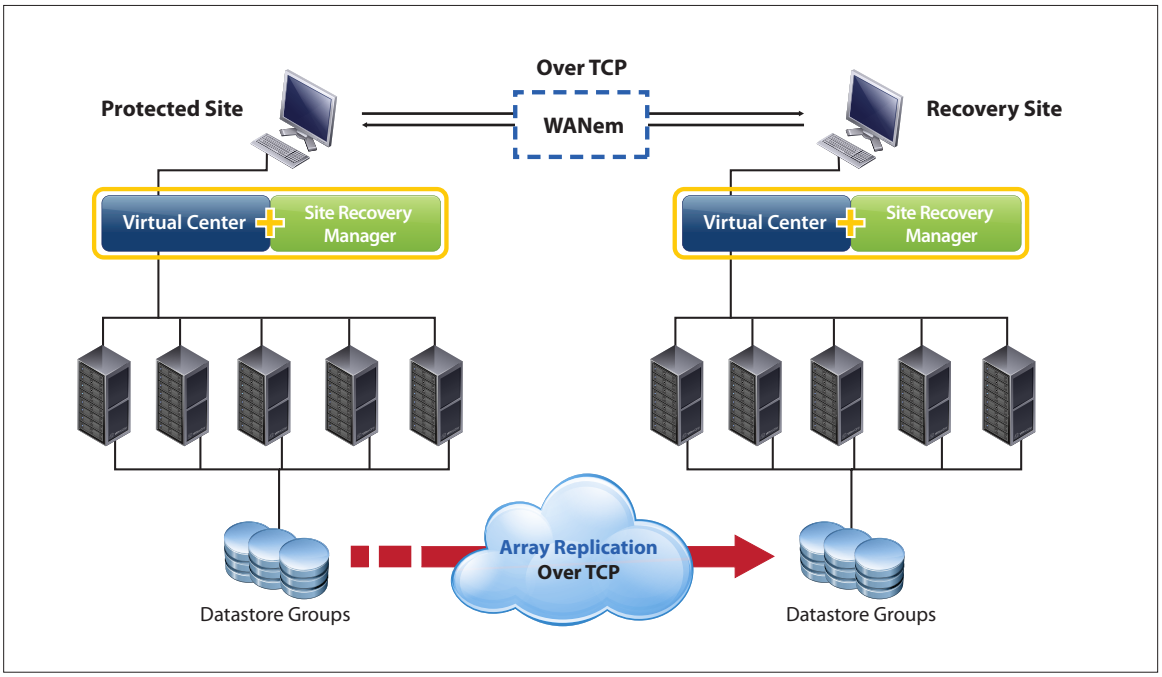

Fig 8. Illustration of the Site Recovery Manager Testbed Environment with a high latency network

Network RTT 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 RTT of 100, 250, 400 milliseconds.

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

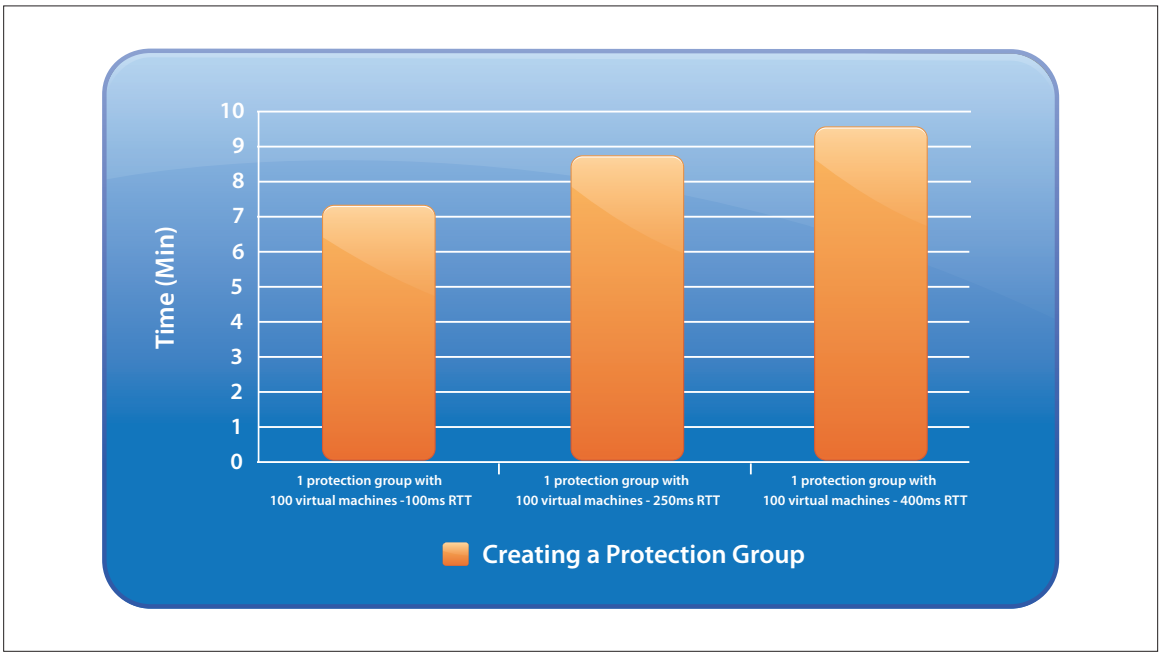

Fig 9. Creating protection groups on a high latency network with different RTT between protected and recovery site

3.3.2. Recoveries on a high latency network

Real recoveries are affected by high latency networks because of the operation to shutdown remote virtual machine 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). A high latency network between a protected and a recovery site should not affect test recoveries. Thus, network latency between the protected and the recovery sites needs to be considered when estimating real recovery time.

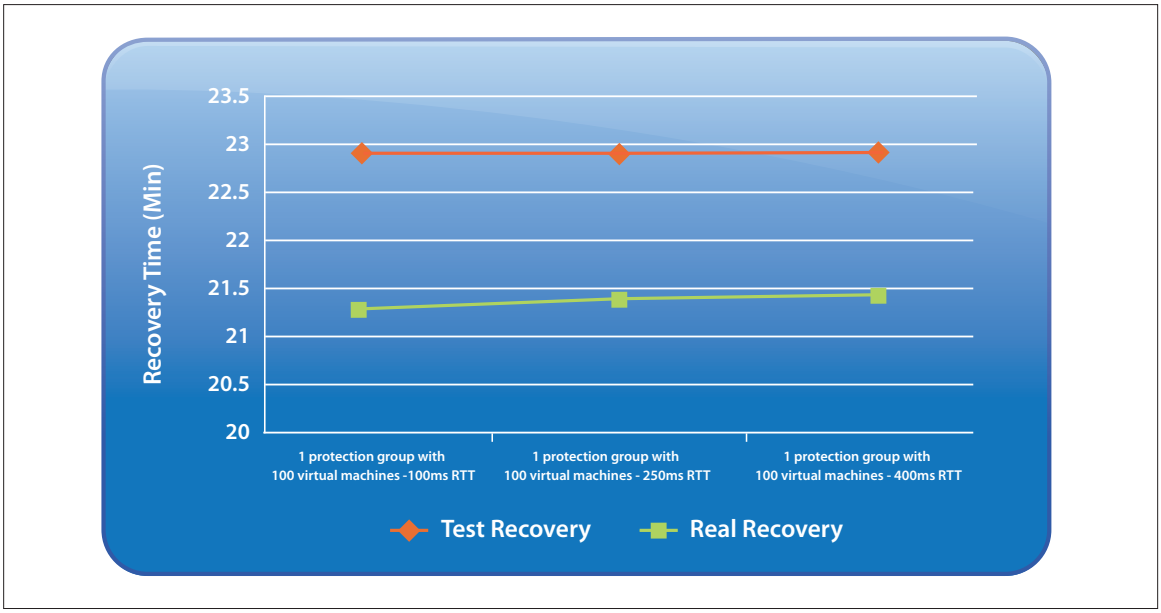

Fig 10. Test and real recovery time on a high latency network with different RTT

4. Architecting Recovery Plans (From a performance/recovery time perspective)

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

4.1. 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.

The following 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).

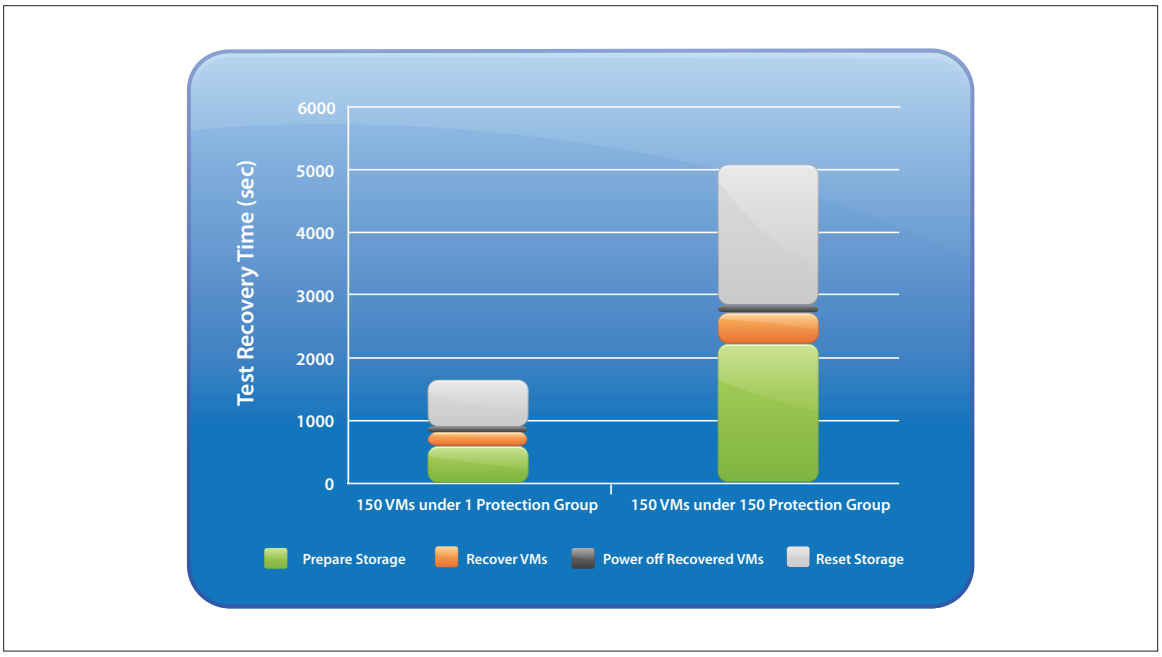

Fig 11. Virtual machine to protection group relation

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

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.

4.2. Recovery time – iSCSI/FC vs. NFS

When working with iSCSI/FC storage, Site Recovery Manager initiates a rescan on the HBAs on the hosts used for the recovery to make the storage available to all the hosts during a recovery (test and real). These rescan calls are issued in parallel across all recovery site hosts.

When working with NFS storage, Site Recovery Manager mounts the replicated NFS volumes/snapshots on the ESX hosts during the recovery. Site Recovery Manager 4.0 issues all the mount/un-mount calls in a serial manner per host.

This means that if it takes X seconds to mount a single replicated NFS volume on a single recovery site host and you have 20 volumes to mount across 10 hosts in the recovery site cluster, then it will take 10*20*X seconds to mount all the volumes across all the hosts during a recovery(both test and real). Similar behavior is expected for unmounting volumes.

For large scale recoveries with a large number of hosts and NFS volumes hosting the protected virtual machines, it might take more time to mount or un-mount NFS volumes across all the recovery site hosts as compared to rescanning all host HBAs for iSCSI or FC.

Key Takeaways:

- It is a good practice to have fewer, but larger, NFS volumes so that the time taken to mount a large number of volumes decreases during the recovery.
- This might also translate to fewer protection groups on your setup – which will help even more in reducing the recovery time as mentioned in Section 4.1.

4.3. Placeholder virtual machines placement and VMware® Distributed Resource Scheduler (DRS) behavior on the recovery site

Site Recovery Manager 4.0 creates placeholder virtual machines at the recovery site for each protected virtual machine on the protected site when creating protection groups. The placement of these placeholder virtual machines is a major factor in Site Recovery Manager determination of where the virtual machines will be powered on during recovery. A rule of thumb is to distribute them evenly on all the available ESX hosts.

Site Recovery Manager 4.0 distributes all the shadow virtual machines in a random fashion across all hosts within a cluster. This is the default behavior and is not dependent on any kind of cluster settings. During a recovery (test and real) each of these placeholder virtual machines are replaced by their respective recovered virtual machines registered from the recovered datastore.

Site Recovery Manager 4.0 initiates a default of 2 Recovery virtual machines operations per host. It can concurrently start up to a max of 18 such operations during a recovery. That is why it becomes important to ensure that placeholder virtual machines are distributed across all hosts planned to be used for recovery to obtain the max level of parallelism Site Recovery Manager offers for recovery operations. This will help in decreasing the recovery time.

For more details on job throttling during a recovery refer to Section 4.3.1.

If you are adding new hosts to the recovery site after the virtual machines have already been protected, you should migrate shadow virtual machines (drag and drop) across the newly added hosts to fully leverage the new hardware for recovery in order to decrease recovery time.

Site Recovery Manager 4.0 also offers host failure resiliency (i.e. if any ESX host hosting a placeholder virtual machine does not respond during a recovery because of any reason such as that the ESX host does not have access to the recovered datastore, then Site Recovery Manager selects other available hosts to recover that virtual machine on).

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. In this case, you do not need to manually distribute shadow virtual machines across all the hosts.

If VMware DRS is not enabled on the recovery cluster, concurrency during recovery will still be maintained because of the default random shadow virtual machine placement during virtual machine protection across the hosts within the recovery cluster.

4.3.1. 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 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 18 or 2 times the number of ESX hosts used for recovery, whichever is less. Adding more than 9 hosts to your recovery site will not speed up the Recovery virtual machine step. These additional hosts will be used for recovering virtual machines, but the maximum number of hosts that can be used simultaneously by Site Recovery Manager during a recovery is 9.

Key Takeaways:

- It is a good practice to have VMware DRS enabled on a recovery site. Migrations might be initiated as DRS tries to load balance the cluster during the recovery.
- If DRS is not enabled, ensure that all recovery site hosts are being used for hosting placeholder virtual machines. This will ensure maximum parallelism during recovery. Evenly distribute all placeholder virtual machines manually across all hosts in this case.

4.4. Standby hosts on the recovery site and enabling DPM

Site Recovery Manager 4.0 works with vSphere to recover virtual machines on a DPM enabled cluster.

If the recovery site cluster has DPM enabled, Site Recovery Manager works with DPM to power on any standby hosts in the cluster. Site Recovery Manager then works with DRS to use these hosts for recovering virtual machines during a test and real recovery.

For the case presented below, a real recovery with 100 virtual machines was carried out on a recovery site cluster (with DPM disabled and enabled) with 5 hosts – out of which 2 were placed in standby mode.

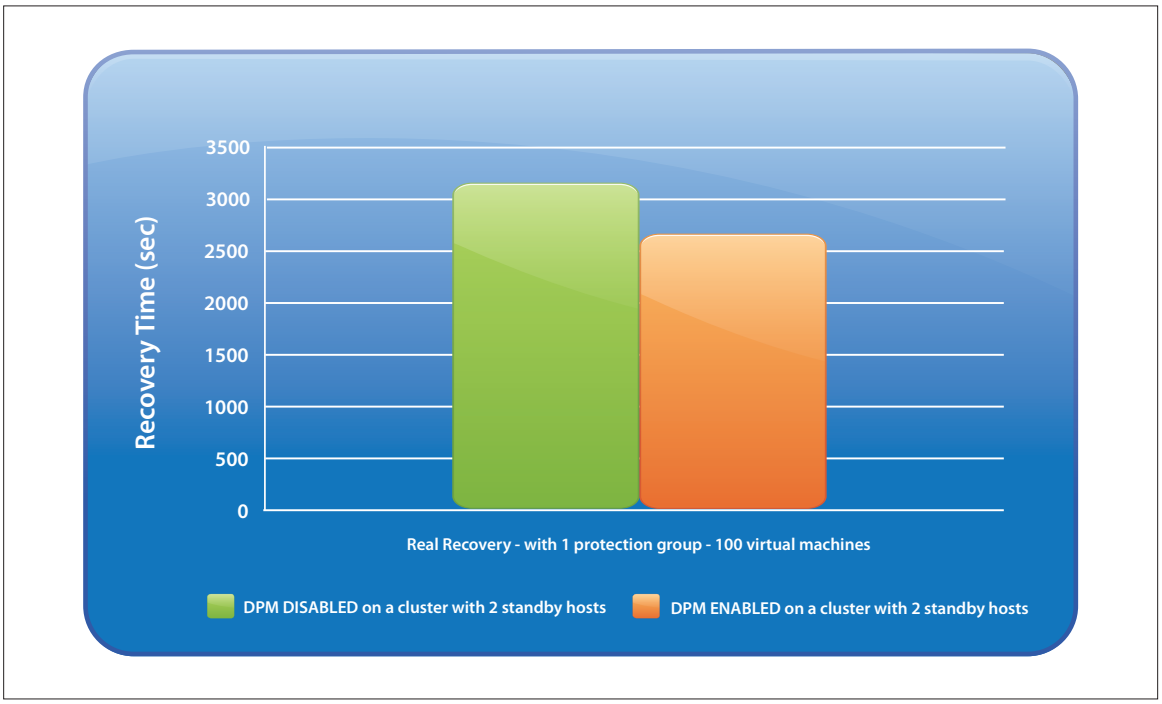

Fig 12. Recovery time benefits with VMware DPM: recovering 100 virtual machines under a single protection group

Recovery with DPM enabled completed faster as Site Recovery Manager had 2 extra powered-on hosts to use for recovering virtual machines concurrently.

Virtual machine recovery starts only after Site Recovery Manager has finished bringing all ESX hosts out of the standby mode. These standby ESX hosts are powered on concurrently. This power-on process creates an overhead, which is the maximum of the time taken to bring any one host out of standby mode. In general the aforementioned overhead is relatively small compared to the gain in recovery time performance due to the availability of more ESX hosts. Since this overhead does not vary as the number of virtual machines increases, you will reap more performance benefits if you have a larger number of protected virtual machines.

Key Takeaways:

- It is a good practice to enable DPM on recovery site clusters if your recovery site hosts are in standby mode. More hosts lead to more concurrency for recovering virtual machines and thus results in shorter recovery time.
- If the hosts are in a standby mode and DPM is not enabled, bring the hosts out of standby mode manually and drag and drop shadow virtual machines on them.
- When protecting virtual machines (i.e. creating protection groups) ensure that the recovery site hosts mapped under the respective inventory are in a proper powered on state otherwise Site Recovery Manager will not use those hosts to create placeholder virtual machines.

4.5. 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.

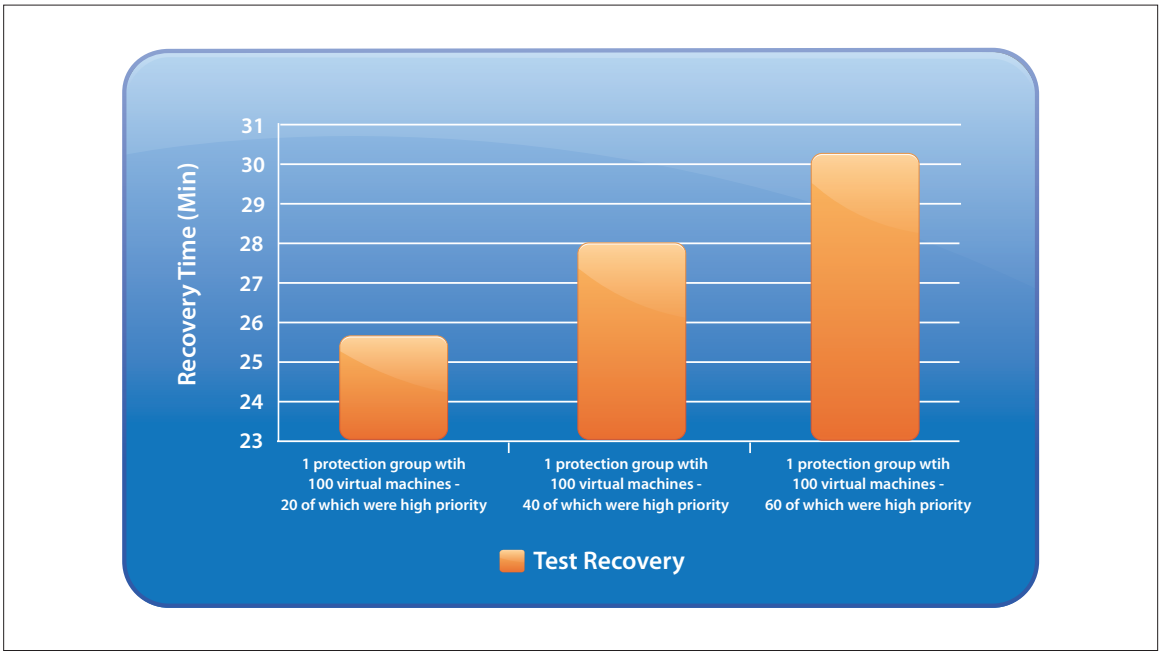

Fig 13. Test recovery: varying the number of high priority virtual machines

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.

As an alternative to placing some virtual machines in the high priority group and some virtual machines in the normal priority group in a single recovery plan, one can separate out all the virtual machines to be recovered into 2 logical groups – Group1 with the first level of virtual machines to be recovered, Group2 with the second level of virtual machines to be recovered (i.e. dependent upon the first level virtual machines in Group1).

Virtual machines in Group1 can be placed in normal priority and virtual machines in Group2 can be placed in low priority. Both of these priority groups allow parallel recovery of virtual machines. Since the low priority group will start after the normal priority group, you will still be maintaining the dependency between the 2 logical groups of virtual machines without sacrificing the recovery time. Note that dependency is maintained between priority groups and there is no concept of dependency across virtual machines within a single priority group.

Key Takeaways:

- It is important to chart out the dependencies 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.

- As an alternative to placing virtual machines in high priority group, one can separate all the virtual machines to be recovered in 2 logical groups – Group1 with level 1 virtual machines, Group2 with level 2 virtual machines, which are dependent upon virtual machines in Group1. Group1 virtual machines can then be placed in the normal priority group and Group2 virtual machines in the low priority group within the same recovery plan. This should maintain dependency across both logical groups and will still reduce the recovery time by introducing more concurrency for both priority groups. Note that dependency is maintained between priority groups and there is no concept of dependency across virtual machines within a single priority group.
- Suspending virtual machines on the recovery site will also impact recovery time.

4.6. Advanced Settings/VMware Tools

Site Recovery Manager 4.0 provides the capability to change the advanced settings[7]. One advanced setting that is related to Site Recovery Manager performance is SanProvider.minLunGroupComputationInterval.

Whenever the protected site inventory changes, Site Recovery Manager will perform a new LUN Group computation. If you are adding multiple virtual machines to the datacenter or changing any inventory in general, you can get Site Recovery Manager to wait before doing another computation by setting SanProvider.minLunGroupComputationInterval to the approximate time taken to make the change. For example, setting SanProvider.minLunGroupComputationInterval to N tells Site Recovery Manager that there should be at least N seconds in between any two consecutive LUN Group Computation tasks. This setting is intended to make LUN Group Computation tasks less frequent when there are a lot of inventory changes going on.

VMware strongly recommends that VMware Tools be installed in all protected virtual machines. 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 Recovery.powerStateChangeTimeout in the advanced settings to zero.

  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.

---

7 Refer to the subsection "Working with Advanced Settings" in section 5 of Site Recovery Manager Administration Guide for more details on Advanced Settings.

4.7. Specify a Non-replicated Datastore for Swap Files

Every virtual machine requires a swap file, which is normally created in the same datastore as the other virtual machine files. When you use Site Recovery Manager, this datastore is replicated. To prevent swap files from being replicated, create them on a non-replicated datastore.

If you are using a non-replicated datastore for swap files, you must create a non-replicated datastore for all protected clusters at both the protected and recovery sites.

Procedure

1. In the vSphere Client, right-click an ESX cluster and click **Edit Settings**.
2. In the Settings window for the cluster, click **Swapfile Location** and select Store the swapfile in the datastore specified by the host, then click **OK**.
3. For each host in the cluster, select a nonreplicated datastore.
   a. Click the **Configuration** tab.
   b. On the Swapfile Location line, click **Edit**.
   c. In the Virtual Machine Swapfile Location window, select a nonreplicated datastore and click **OK**.

Recovery Time advantages

Site Recovery Manager makes remote calls to vCenter Server to delete any swap files found on a replicated datastore during a recovery on the recovery site. If swap files reside on non-replicated datastores, then this step will be skipped and it will help speed up the recovery. This will also avoid wasting of network bandwidth during replication between the 2 sites.

5. 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:

- It is recommended that Site Recovery Manager database be installed as close to the Site Recovery Manager server as possible – such that it reduces the RTT between both of them. This way recovery time performance will not suffer greatly because of round trips to the database server.
- 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.
- It is a good practice to have fewer but larger NFS volumes so that the time taken to mount a large number of such volumes decreases during the recovery. This might also translate to fewer protection groups on your setup leading to reduced recovery time.
- It is a good practice to have VMware DRS enabled on a recovery site. Migrations might be seen as DRS tries to load balance the cluster during the recovery.
- It is a good practice to enable DPM on recovery site clusters if your recovery site hosts are in standby state. More hosts lead to more concurrency for recovering virtual machines and thus results in shorter recovery time. If DPM is not enabled and the hosts are in a standby state, bring the hosts out of standby mode manually and drag and drop shadow virtual machines on them.

- 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. As an alternative to placing virtual machines in high priority group, one can separate all the virtual machines to be recovered in 2 logical groups – Group1 with level 1 virtual machines and Group2 with level 2 virtual machines, which are dependent upon virtual machines in Group1. Group1 virtual machines can then be placed in normal priority group and Group2 virtual machines in the low priority group within the same plan. This should maintain dependency across both logical groups and will still reduce the recovery time by introducing more concurrency for both priority groups. Note that dependency is maintained between priority groups and there is no concept of dependency across virtual machines within a single priority group.
- 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.
- Specify a non-replicated datastore for swap files. This avoids wasting network bandwidth during replication between 2 sites and remote calls to vCenter Server during a recovery to delete swap files for all virtual machines, which in turn helps in speeding up the recovery.

6. Appendix

6.1. Acknowledgements

Thanks to Arturo Fagundo, David Levy, Maria Basmanova, Win Crofton, Fabio Rojas, 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.

6.2. 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.

6.3. References

- VMware vCenter Site Recovery Manager Documentation

  http://www.vmware.com/support/pubs/srm_pubs.html

- VMware vCenter Site Recovery Manager 4.0 Evaluator's Guide

  http://www.vmware.com/files/pdf/vcenter-srm-evaluators-guide.pdf

- Site Recovery Manager Administration Guide

  http://www.vmware.com/pdf/srm-admin.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 online.

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