



VMware vCenter Site Recovery Manager[™] 5.0 Performance and Best Practices

Performance Study

TECHNICAL WHITE PAPER

Table of Contents

Introduction.....	3
About Site Recovery Manager.....	3
Performance Considerations in a Site Recovery Manager Environment.....	3
Test Environment.....	4
Hardware/Software Configuration	5
Site Recovery Manager Server Configuration Recommendation	6
Summary of Performance Improvements Made in SRM 5.0 over SRM 4.0.....	6
Basic Operation Latency Overview	7
Site Recovery Manager Scalability Performance.....	9
Protection (Scaling with a number of virtual machines)	9
Recovery (Scaling with a number of virtual machines and scaling with a number of protection groups).....	9
Architecting Recovery Plans (From a performance/recovery time perspective).....	12
Virtual Machine to Protection Group Relation	12
Recovery Time – iSCSI/FC vs. NFS.....	13
Placeholder VM Placement and DRS Behavior on the Recovery Site	13
Configuring Job Throttling Parameters for VM Power Operations.....	14
Standby Hosts on the Recovery Site and Enabling DPM	15
High Priority Virtual Machines and Suspending Virtual Machines.....	16
Advanced Settings/VMware Tools	16
Specify a Nonreplicated Datastore for Swap Files	19
Recommendations.....	20
References	20

Introduction

VMware vCenter™ Site Recovery Manager (SRM) 5.0 provides business continuity and disaster recovery protection for VMware virtual environments. Protection can range from virtual machines (VMs) residing on a single, replicated datastore to all the VMs in a datacenter and includes protection for the operating systems and applications running in the VM.

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

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

- Number of virtual machines and protection groups associated with a recovery plan
- Improvements to the performance of recovering multiple protection groups within a single recovery plan
- Leveraging DPM and DRS for a better recovery
- Configuration of various recovery plan parameters
- Priority assignment of virtual machines in the recovery plan

Furthermore, we suggest best practices in applicable areas so that you can optimize the recovery time using SRM.

About Site Recovery Manager

SRM requires a protected site and a recovery site and an SRM server must be installed at each site. Additionally, each site must be managed by its own vCenter Server.

On the protected site, you configure a protection group, which is a group of virtual machines that fail over together. On the recovery site, you add the protection group, which is a collection of VMs that can be recovered simultaneously.

SRM 5.0 supports NFS, iSCSI, and FC storage and supports two forms of replication: array-based replication (ABR) in which the storage subsystem manages VM replication, and host-based replication (vSphere replication) in which ESXi manages VM replication.

SRM automatically discovers datastores set up for array-based replication between the protected and recovery sites. SRM supports 1,000 VMs for ABR.

vSphere replication (VR) replicates only the most recent data in changed disk areas to increase network efficiency and eliminates the ABR requirement for having identical storage arrays across sites. SRM supports 500 VMs for VR.

Performance Considerations in a Site Recovery Manager Environment

Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are the two most important performance metrics you need to keep in mind while designing and executing a disaster recovery plan.

- RPO defines the point in time at which data must be restored to meet service level agreements.
- 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.

For array-based replication, RPO is fulfilled by the replication schedules configured on the storage array. For vSphere replication, you set the RPO using the SRM plugin in the vSphere Client. The minimum RPO you can set

with VR is 15 minutes. The VR algorithm adjusts the replication schedule dynamically in order to fulfill the RPO. Site Recovery Manager helps you meet RTO by minimizing datacenter recovery time, which is crucial for any business continuity or disaster recovery solution.

Test Environment

Figure 1 shows 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 a different number of ESXi hosts at protected and recovery sites.

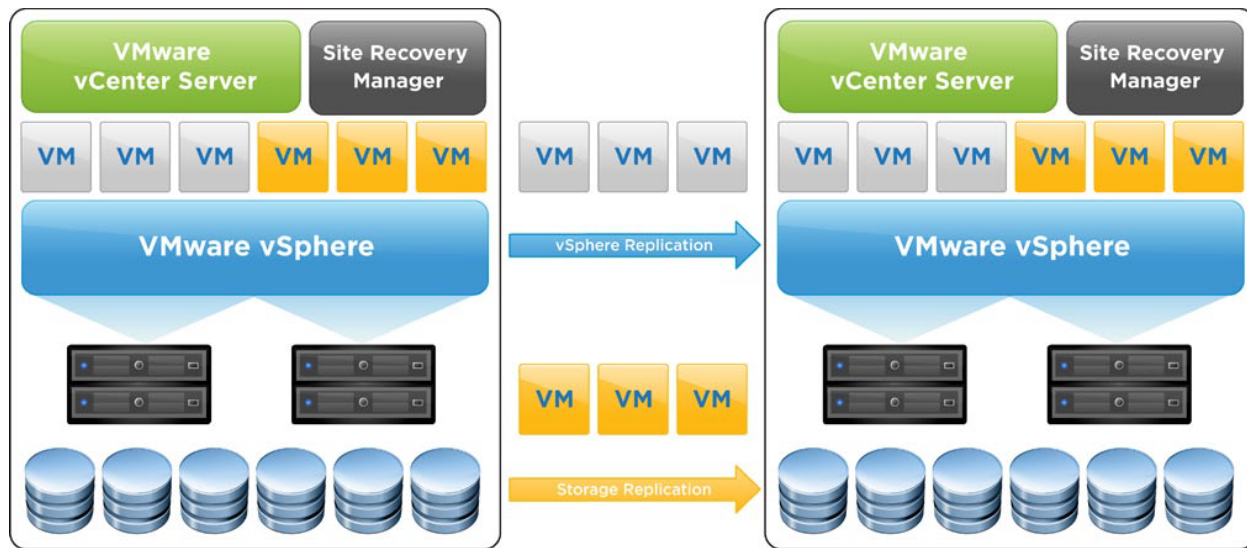


Figure 1. Illustration of the Site Recovery Manager test bed environment

Hardware/Software Configuration

Site Recovery Manager 5.0, vCenter 5.0, and ESXi 5.0 were used for performance measurements.

- **Site Recovery Manager and vCenter Server – Protected Site and Recovery Site Configuration**
Site Recovery Manager, VMware vCenter™ servers, and SRM database servers were installed in individual VMs with the following configuration:
 - CPUs: 4 vCPUs
 - RAM: 10GB
 - OS: Windows Server 2003 x64
 - Site Recovery Manager Server software: Site Recovery Manager 5.0
 - VMware vCenter Server software: VMware vCenter 5.0
 - vCenter database was installed in the same virtual machine as vCenter Server
 - Microsoft SQL Server 2005 was used to host both SRM and vCenter databases
- **vRMS, vR, vRMS Database Server – Protected Site and Recovery Site Configuration**
vR servers were installed in VMs with the following configuration:
 - CPUs: 1 vCPU
 - RAM: 512MBvRMS servers were installed in VMs with the following configuration:
 - CPUs: 2 vCPUs
 - RAM: 4GB
 - OS: Novell SUSE Linux Enterprise 11 (64 bit)vRMS database servers were installed in VMs with the following configuration:
 - CPUs: 2 vCPUs
 - RAM: 5GB
 - OS: Windows Server 2003 x64
 - Microsoft SQL Server 2005 was used to host these databases.
- **ESXi System – 5 ESXi 5.0 hosts on Protected Site, 5 ESXi 5.0 hosts on Recovery Site**
 - Host Computer: Dell PE2650
 - CPUs: Intel Xeon CPU 3.06 GHz - 8 cores
 - RAM: 48GB
 - Network: NetXtreme BCM5703 Gigabit Ethernet 1Gbps
 - ESXi software: VMware ESXi 5.0
- **Storage**
 - 2 NetApp FAS3140 arrays
 - System version : 7.3.3.3

Site Recovery Manager Server Configuration Recommendation

In case you decide to install SRM and vCenter servers in a VM, here is a vCPU sizing guideline for the VM that helps achieve optimal performance for the SRM workload. Note that after executing a planned migration and a reprotect operation, your primary site will become your secondary site and the recommendations for the secondary site that are mentioned here will be applied during failback to your primary site and vice-versa. Also note that the vCPU for the secondary site is more than that for the primary site for 650-1,000 VM operations.

ARRAY BASED REPLICATION SETUP CONSISTS OF	PRIMARY SITE SRM VM VCPUS	SECONDARY SITE SRM VM VCPUS	PRIMARY SITE VCENTER VM VCPUS	SECONDARY SITE VCENTER VM VCPUS
< 10-50 VMs	2	2	2	2
~ 50-300 VMs	4	4	4	4
~ 300-650 VMs	4	4	4	4
~ 650-1,000 VMs	4	4	4	6

Table 1. vCPU sizing guidelines

Site Recovery Manager server uses one database each: one on the protected site and one on the recovery site to store information. 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.

We recommend that the SRM database be installed as close to the SRM server as possible, such that it reduces the round-trip time between both of them. This way, recovery time performance will benefit 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, although it is a good practice to maintain separate database servers for each database instance.

Database size depends on:

- 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

Summary of Performance Improvements Made in SRM 5.0 over SRM 4.0

The following improvements are representative of the performance data collected from our lab environment described in the Introduction section. The scale of improvement may vary in your lab environment because it depends on the scale of your inventory¹ and setup configuration².

¹ Inventory includes, but is not limited to, virtual machines, protection groups, datastores, clusters, network and folders.

² Setup configuration includes, but is not limited to, hardware configuration and RTT between the SRM server and SRM database.

- **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 VMs
 - UI performance speedup for displaying large recovery plans
- **Protection improvements**
 - Significant performance improvement for protecting VMs and creating protection groups
- **Recovery time improvements**
 - Increased parallelism for recovering VMs
 - Significant performance improvement for preparing storage during recovery plan execution through the use of parallelized operations
 - Significant performance improvement for powering on VMs during recovery plan execution through the use of parallelized operations

Test and real recovery performance has improved significantly in SRM 5.0 when compared to SRM 4.0 because of the recovery time improvements previously mentioned.

Basic Operation Latency Overview

VMware vCenter Site Recovery Manager provides two major operations: protection and recovery.

Protection involves following operations common for both ABR and VR except for the replication setup:

- Replication setup
 - Array Manager configuration for ABR
 - Configuring VM replication schedule for VR
- Inventory mapping
 - Inventory mappings are associations between resource pools, virtual machine folders, networks at the protected site and their destination counterparts at the recovery site.
- Creating a protection group and protecting VMs within that group
 - A protection group is a group of VMs that will be failed over together to the recovery site during testing or recovery.

Recovery involves the following operations which are common for both ABR and VR, except for reprotect which is supported only for ABR:

- Creating a recovery plan
 - A recovery plan is the complete set of steps needed to recover (or test the recovery of) the protected VMs in one or more protection groups
- Planned migration
 - During a planned migration, SRM attempts to shut down protection site VMs and replicate outstanding changes to the recovery site before proceeding with the failover sequence.
- Unplanned migration
 - During an unplanned migration, SRM proceeds directly with the failover sequence without attempting to shut down protection site VMs and replicate changes to the recovery site.
- Reprotect for ABR only
 - Reprotect involves a reversal of direction of replication, and automatic reprotection of protection groups.

NOTE: The actual numbers can vary in a real deployment and the numbers presented here are from a specific setup. The VMs used for the performance data presented in the "Major SRM 5.0 operations" figure were configured for IP customization. The VMs used for presenting the rest of the performance data in this white paper

did not have any IP customization specs or any waiting for heartbeats during the recovery. The motivation was to exercise SRM server behavior during a recovery with different settings.

The following graph depicts average baseline latencies for major operations.

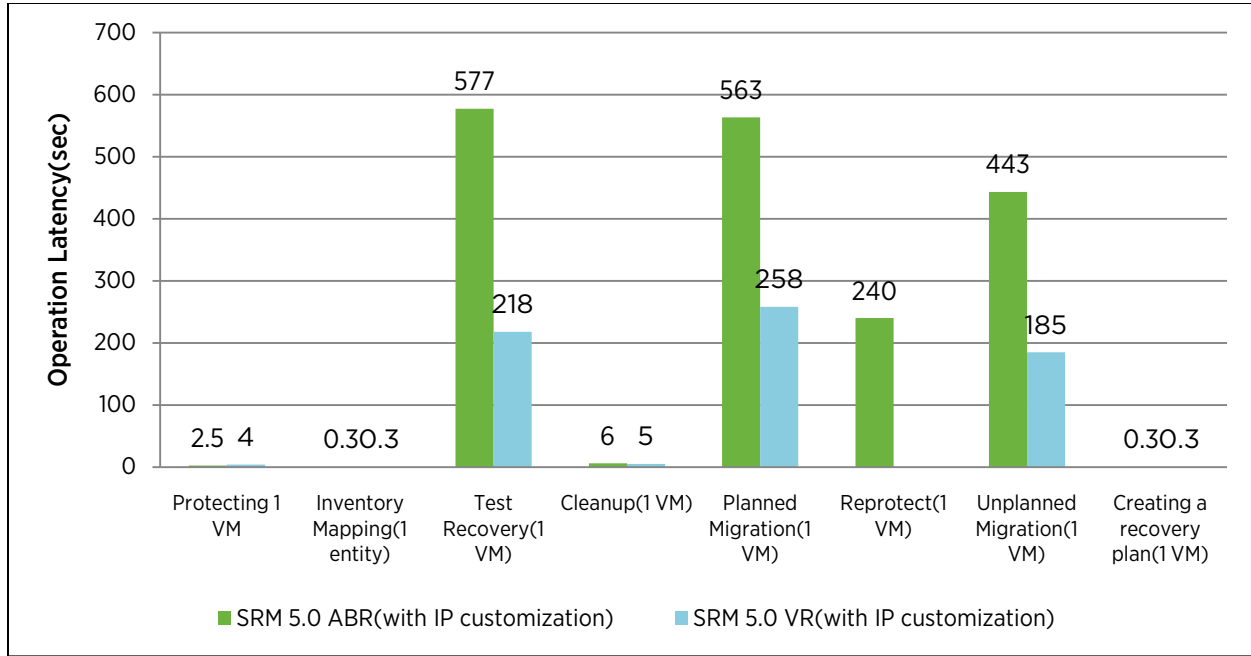


Figure 2. Major SRM 5.0 operations - latency overview

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

Site Recovery Manager Scalability Performance

Protection (Scaling with a number of virtual machines)

Site Recovery Manager 5.0 allows protecting a maximum of 1,000 virtual machines for ABR and 500 VMs for VR.

Figure 3 presents some trending data for the protection of VMs with ABR.




Figure 3. Protection group creation time: scaling with virtual machines under a single protection group

Here, the linear trend line identifies ideal scaling, and the trend line for creating protection groups and protecting VMs resembles the ideal trend. This shows that SRM scales well with an increasing number of virtual machines in a protection group.

Our testing also showed the majority of time for creating protection groups and protecting VMs is spent in creating placeholder virtual machines on the recovery site.

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

Site Recovery Manager 5.0 allows recovering a maximum of 1,000 virtual machines and a maximum of 150 protection groups.

Figures 4 and 5 show the statistics we gathered using NFS storage.

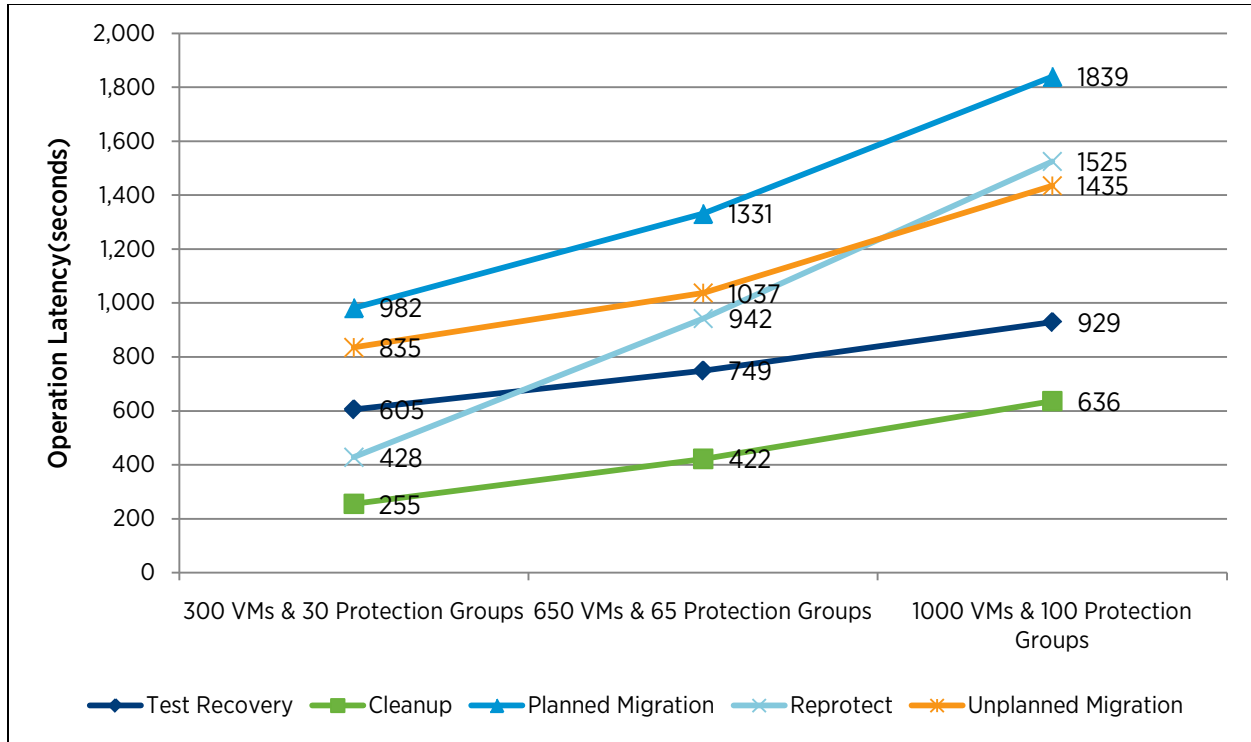


Figure 4. ABR trending experiments: scaling with protection groups and virtual machines

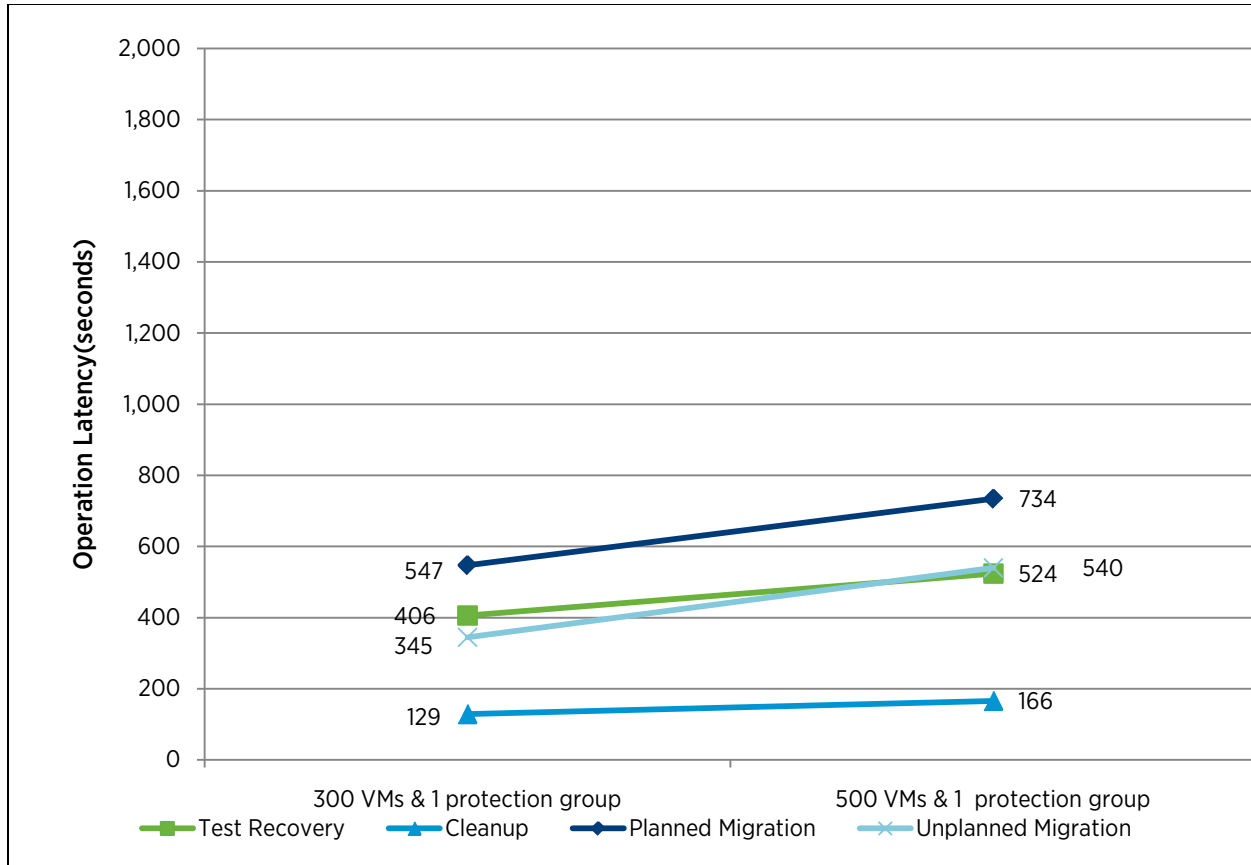


Figure 5. vSphere replication trending experiments: scaling with virtual machines

While the performance of SRM 4.0 was linear or super-linear in nature, the trend line for all scalability-related trending experiments with SRM 5.0 is sub-linear or linear in nature which indicates improved performance in SRM 5.0 compared to SRM 4.0.

For both ABR and VR, planned migration takes more time than unplanned migration. This is because planned migration shuts down VMs on the protected site and synchronizes storage before recovering the virtual machines on the recovery site.

We also compared SRM 5.0 performance with SRM 4.0 performance with the same setup and inventory and noticed an improvement of 8x to 15x for the operations shown in Figure 5.

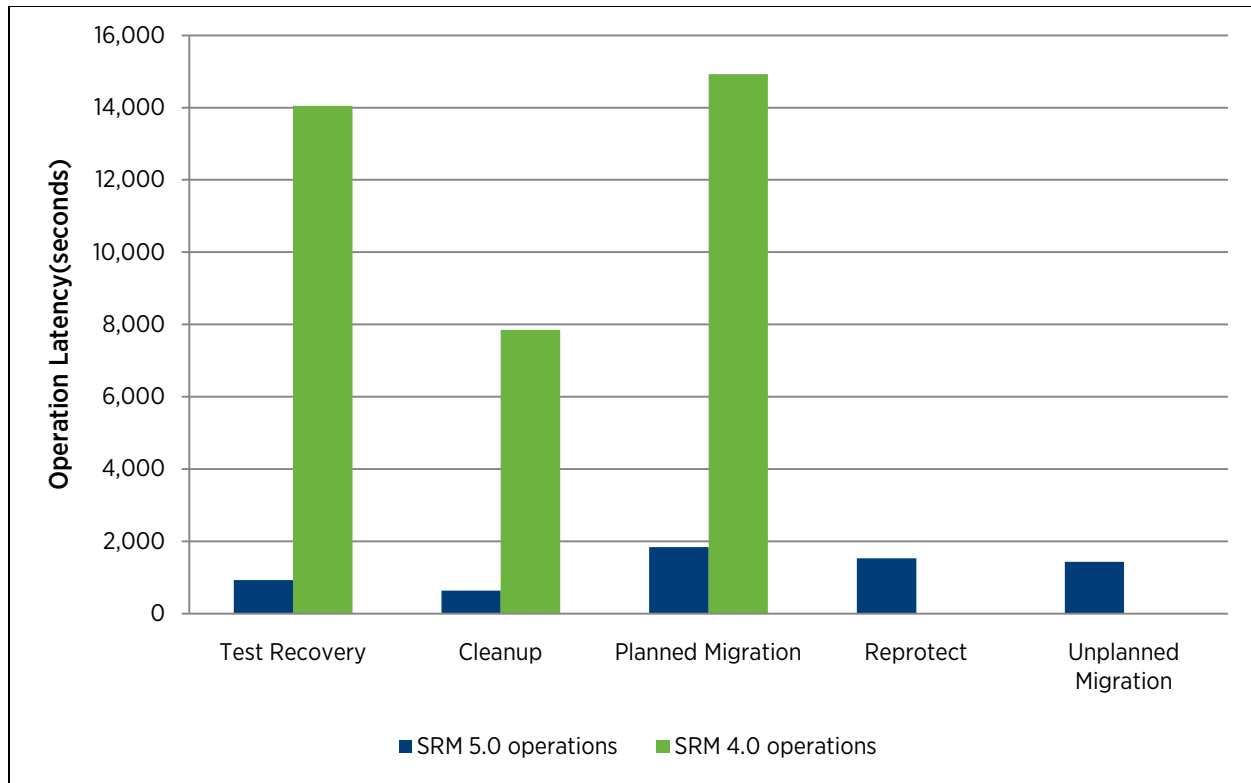


Figure 6. SRM 5.0 performance vs SRM 4.0 performance (both with 1,000 VMs)

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

This section presents Site Recovery Manager best practices in certain areas. These best practices will help you architect efficient recovery plans that minimize recovery time.

Virtual Machine to Protection Group Relation

With ABR, 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. SRM 4.0 executed this operation sequentially for each protection group. As a result, adding more protection groups to a recovery plan as opposed to adding more virtual machines was more costly from a recovery time perspective. This also depended on the underlying storage used for replication between the two sites. SRM 5.0 parallelizes the communication with the underlying storage, resulting in a significant performance improvement.

We compared recovery time measurements with ABR for a test recovery for 300 virtual machines in a single protection group vs. 300 virtual machines in 30 protection groups (10 virtual machines per protection group).

The test recovery performance was almost similar despite the fact that one recovery plan had more protection groups when compared to the other.

- **Key takeaway:**

Adding protection groups to a recovery plan does not increase the recovery time by a large factor.

Recovery Time – iSCSI/FC vs. NFS

When working with NFS storage, SRM mounts the replicated NFS volumes/snapshots on the ESXi hosts during the recovery. SRM 4.0 issued all the mount/unmount calls in a serial manner per host.

SRM 5.0 issues all the mount/unmount calls in a parallelized fashion across all hosts being used in the recovery plan. However it still limits the calls to two parallel calls 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 multiple volumes to mount across multiple hosts in the recovery site cluster, then it will roughly take $((\text{unique_volumes_to_be_mounted_across_all_hosts})/2) * X$ seconds to mount all the volumes across all the hosts during a recovery (both for test and real recoveries). Similar behavior is expected for unmounting volumes.

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

When working with iSCSI/FC storage, SRM 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.

- **Key takeaway:**

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.

Placeholder VM Placement and DRS Behavior on the Recovery Site

When you create a protection group for a set of virtual machines on the protected site, SRM creates placeholder virtual machines at the recovery site for each protected virtual machine. During a recovery, SRM replaces each of these placeholder virtual machines with full versions of the virtual machines, which are recovered from the datastore. SRM 5.0 automatically leverages DRS (independent of any cluster/DRS settings) to optimally place the virtual machines across various ESXi hosts in the cluster, even taking advantage of newly added hosts. This ensures that your cluster remains balanced after the recovery completes. This also helps cluster and VM performance during a boot storm—when many VMs are powered on at the same time after they are recovered.

SRM 4.0 used to initiate a default of two “Recovery VM” operations per host. It could concurrently start up to a maximum of 18 such operations during a recovery. SRM 5.0 initiates a default of unlimited “Recovery VM” operations across all hosts. This helps tremendously in decreasing the recovery time.

For more details on configuring this “unlimited” amount of Recovery VM operations during a recovery to a specific value which limits these throttling jobs, see the next section, [Configuring Job Throttling Parameters for VM Power Operations](#).

SRM 5.0 also offers host failure resiliency—if any ESXi server hosting a placeholder VM does not respond during a recovery (for example, if the ESXi host does not have access to the recovered datastore), then SRM selects other available hosts.

- **Key takeaway:**

It is a good practice to have vSphere DRS enabled on the recovery site. SRM 5.0 leverages DRS to reserve sufficient resources during the recovery in order to successfully power on all VMs.

Configuring Job Throttling Parameters for VM Power Operations

SRM offers certain parameters to control the boot and shutdown operations per cluster and per host.

You might want to set these parameters to an appropriate value that would be in line with the capabilities of the underlying systems being used for the disaster recovery. This gives you some control over the boot storm executed by an SRM-initiated recovery.

Here is some data which throws light on how much guest boot up latencies add up to the overall recovery time:

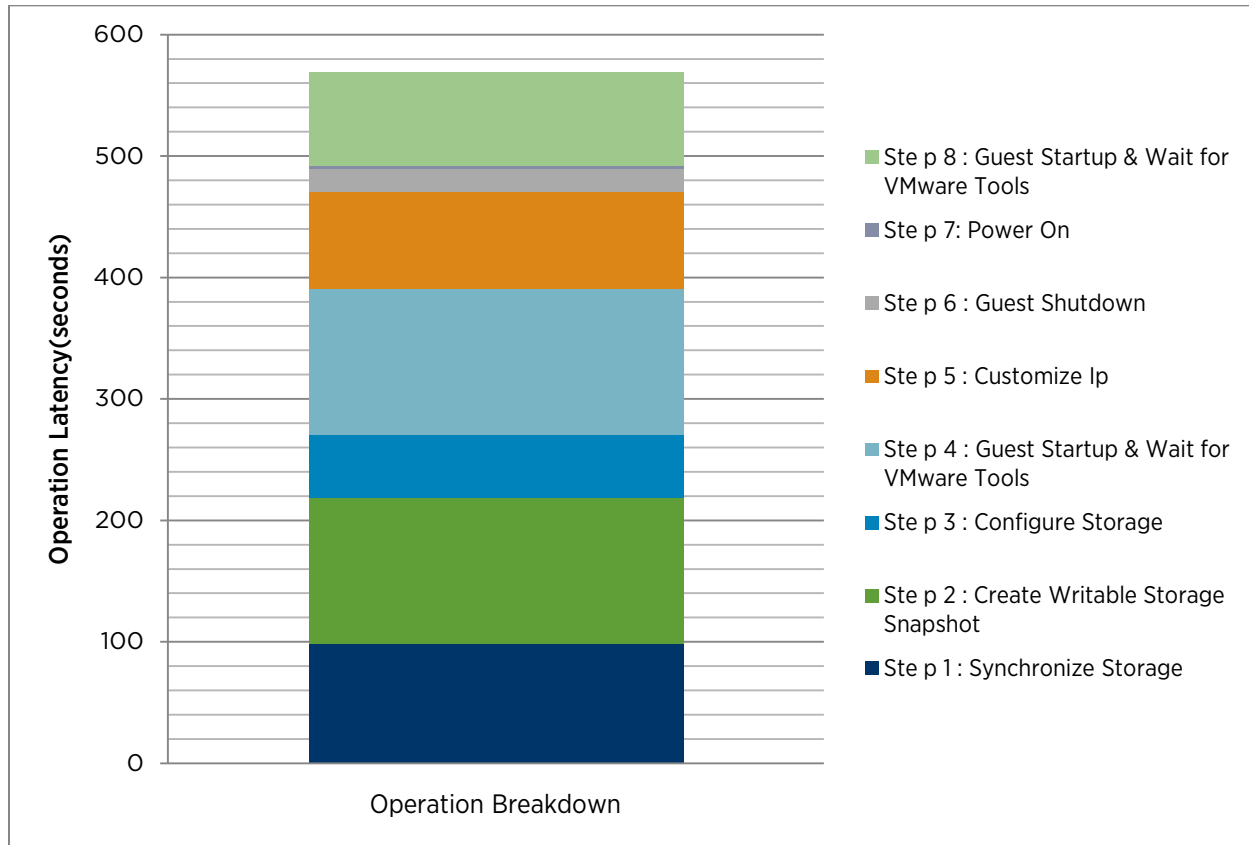


Figure 7. Single VM test recovery with ABR

As seen in the Figure 7, Steps 4 and 8—both of which involve starting up the guest and waiting for VMware Tools—take up a considerable chunk of time from the overall recovery time. With a boot storm, these latencies could increase as a result of I/O bottlenecks or any other resource bottlenecks your platform might encounter.

To configure parameters that will help control the boot storm effectively:

1. Locate the SRM folder, and in it find the `config` folder.
2. You should now be able to find the `vmware-dr.xml` file.
3. Use a text editor like Notepad to edit this file.
4. Look for the section in this file that is denoted by `<Config>`.
5. Add the following lines between the `<Config>` and `</Config>` tags.

```
<defaultMaxBootAndShutdownOpsPerCluster>20</defaultMaxBootAndShutdownOpsPerCluster>
<defaultMaxBootAndShutdownOpsPerHost>20</defaultMaxBootAndShutdownOpsPerHost>
```

Note that 20 is just a sample value that we've used here. Consider the performance of your underlying platform when choosing an appropriate value. The following steps describe one way to do this:

1. Set these values to a specific number.
2. Run a test recovery.
3. If you notice the following error messages:

Error - Cannot complete customization, possibly due to a scripting runtime error or invalid script parameters.

Error - An error occurred when uploading files to the guest VM.

Error - Timed out waiting for VMware Tools after 600 seconds.

Then you can:

- a. Run Cleanup, decrease the config values, and go to Step 2.

If you do not notice any errors and feel that your platform is undercommitted, then:

- b. Run Cleanup, increase the config values, and go to Step 2.

Based on the above steps, try to find a sweet spot for these config values such that you can gain optimum RTO while not completely overcommitting the platform.

If you want to increase SRM log retention, refer to:

<http://blogs.vmware.com/uptime/2011/04/increasing-srm-log-retention.html>

Standby Hosts on the Recovery Site and Enabling DPM

Site Recovery Manager 5.0 works with vSphere to recover VMs even on standby hosts.

If the recovery site cluster has a couple of standby hosts, SRM works with DPM to power on any standby hosts in the cluster. SRM then works with DRS to use these hosts for recovering VMs during a test and real recovery. SRM 5.0 brings DPM out of automatic mode while VMs are being powered on during a recovery in order to prevent hosts from being placed in a standby state. It resets DPM cluster settings to their original state after the recovery has successfully completed. This process with DPM occurs whether or not you have DPM enabled for your cluster.

VM recovery starts only after SRM has finished with bringing all ESXi hosts out of standby mode. These standby ESXi hosts are powered on concurrently. This power-on process creates an overhead, which is the maximum amount of 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 ESXi hosts. Since this overhead does not vary as the number of VMs increases, you will reap more performance benefits if you have a larger number of protected VMs.

- **Key takeaways:**

- More hosts lead to more concurrency for recovering VMs and so results in a shorter recovery time.
- When protecting VMs (creating protection groups) ensure that the recovery site hosts mapped under the respective inventory are in a proper powered-on state; otherwise, SRM will not use those hosts to create placeholder VMs.
- SRM 5.0 places DPM in a manual mode while VMs are being powered on during a recovery in order to prevent hosts from being placed in a standby state. It resets DPM to its original state after the recovery has successfully completed.

High Priority Virtual Machines and Suspending Virtual Machines

In a recovery plan, the virtual machines being recovered can be assigned to five different priority groups. SRM 5.0 also provides the functionality of setting dependencies across individual VMs.

The recovery time does increase with an increase in the dependency across individual VMs or groups of VMs. This applies to both real and test recoveries.

Chaining groups of VMs together is a better idea than chaining individual VMs. This means that "priority groups" should be used first to determine the dependency and startup order of the VMs instead of using the "VM dependencies" functionality directly, because individual VMs starting up sequentially affect the RTO. Grouping VM dependencies in priority groups is usually the best and the safest idea because VMs within each priority group will be started in parallel.

You can also configure SRM to suspend VMs during a recovery. Note that suspending VMs is a resource intensive operation which might take time to complete depending upon the configuration of the VMs being suspended. Your overall RTO might increase if a lot of VMs are being suspended. However, the upside to this is that suspending VMs frees up platform resources, which can then be used to recover other VMs that are part of your recovery plan.

- **Key takeaways:**

- It is important to chart out the dependencies and priorities between virtual machines to be recovered so that only a certain number of required virtual machines can be assigned individual dependencies. Such dependencies impact recovery time.
- Configuring VM dependencies across priority groups instead of setting per VM dependencies is highly recommended because VMs within each priority group will be started in parallel.
- Suspending virtual machines on the recovery site will also impact recovery time.

Advanced Settings/VMware Tools

SRM offers certain advanced settings that you can configure on each site. These settings can affect the performance of general SRM operations and critical operations like test and real recoveries.

To configure advanced settings:

1. From the vSphere Client, right click the site on which you want to configure a particular setting.

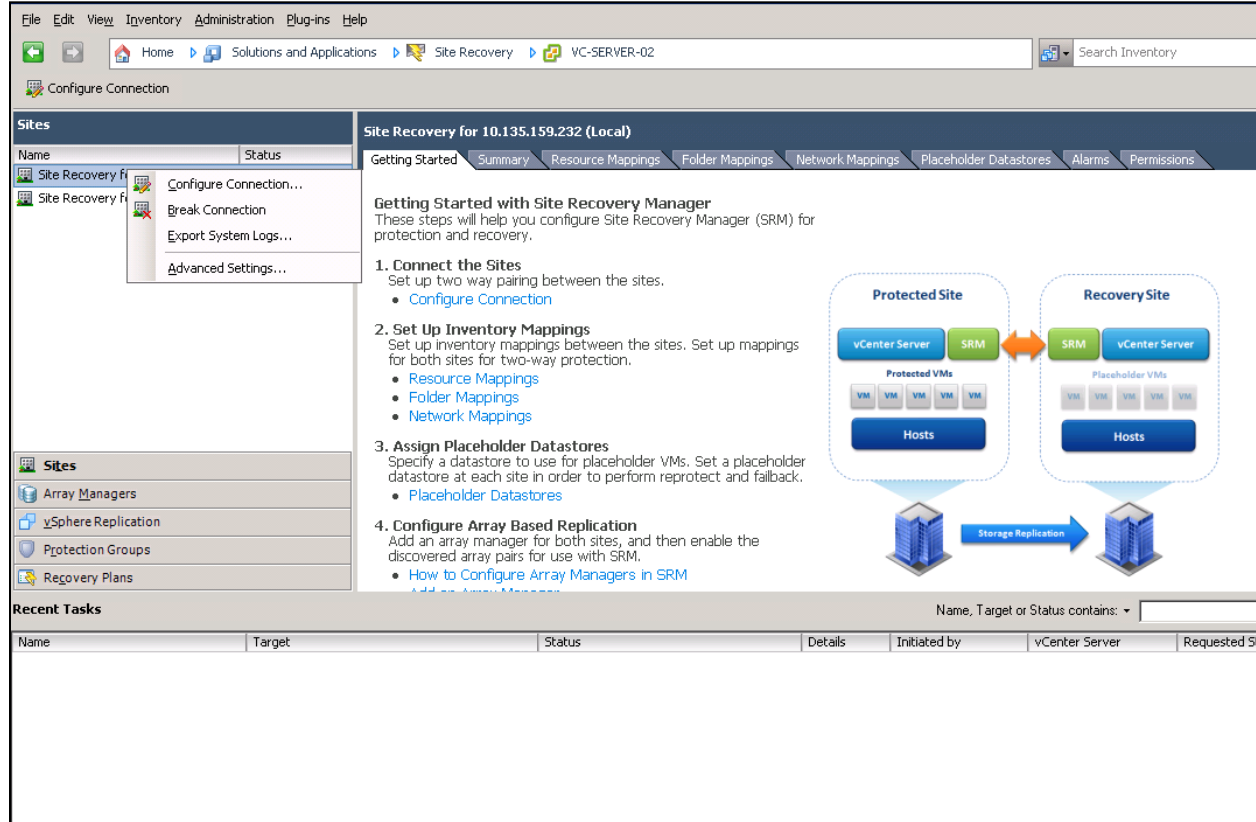


Figure 8. Choose site in left panel, then choose Advanced Settings

2. Click **Advanced Settings** and browse through the various settings to determine which ones you want to change. Once that is done, click **OK**.

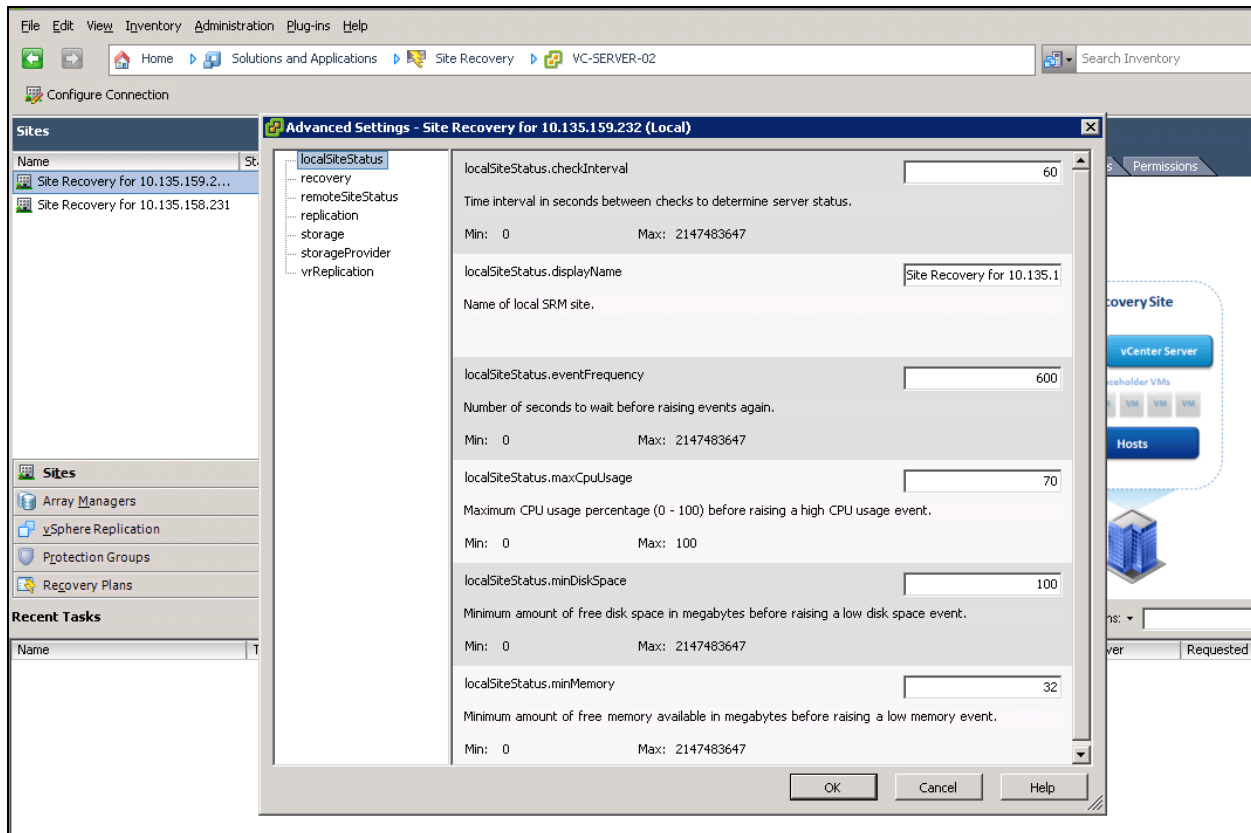


Figure 9. Advanced Settings page

Here are some settings that you might want to be aware of:

- Whenever the protected site inventory changes, SRM will perform a new LUN Group computation. If you are adding multiple VMs to the datacenter or changing any inventory in general, you can get SRM to wait before doing another computation by setting **storage.minDsGroupComputationInterval** to the approximate time taken to make the change. For example, setting **storage.minDsGroupComputationInterval** to a number value tells SRM that there should be at least that many 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 SRM recovery operations depend on the proper installation of VMware Tools in the protected virtual machines to carry out the following tasks:
 - Wait for the OS heartbeat while powering on the virtual machine and wait for a network change while reconfiguring the recovered virtual machine. SRM depends on VMware Tools to report the OS heartbeat and completion of the network change. However, if you do not have VMware Tools installed on any of the protected virtual machines, you can choose to set the timeout values for **recovery.powerOnTimeout** and **recovery.customizationTimeout** to zero (0).
 - Wait for virtual machines to shut down on the protected site.
 - During a planned migration, Site Recovery Manager tries to gracefully shut down the virtual machines on the protected site. Before Site Recovery Manager forcibly powers a virtual machine off, it tries to shut down the guest OS. If your intention is to power off the virtual machines without gracefully shutting down

the guest OS, you can set **recovery.skipGuestShutdown** to true in the Advanced Settings³ menu.

*NOTE: If the VMs do not have VMware Tools installed and the guest shutdown timeout is set to a non-zero value, then your recovery will not proceed beyond the “Shutdown VMs at Protected Site” step. When your VMs do not have VMware Tools installed, you will have to set **recovery.skipGuestShutdown** to true if you want your recovery plan to make any progress.*

Specify a Nonreplicated 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 SRM, 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. To do so:

1. In the vSphere Client, right-click an ESXi 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**, and 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

SRM makes remote calls to vCenter Server for deleting 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 is skipped, which speeds up the recovery. This will also avoid wasting network bandwidth during replication between the two sites. This is especially true if you're using NFS storage.

³ Refer to the subsection “Working with Advanced Settings” in section 8 of Site Recovery Manager Administration Guide (http://www.vmware.com/pdf/srm_admin_5_0.pdf) for more details on Advanced Settings.

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 the SRM database be installed as close to the SRM server as possible, such that it reduces the round-trip time between both of them. This way recovery time performance will not suffer greatly because of round trips to the database server.
- 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 DRS enabled on a recovery site.
- More hosts lead to more concurrency for recovering VMs which results in a shorter recovery time.
- Also, before protecting VMs, bring recovery site hosts out of standby mode so that they get leveraged for creating placeholder VMs.
- Configuring VM dependencies across priority groups instead of setting per VM dependencies is usually the best idea because VMs within each priority group will be started in parallel.
- 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. Refer to [Advanced Settings/VMware Tools](#) for more information.
- 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 two sites and reduces remote calls to vCenter Server during a recovery to delete swap files for all VMs, which in turn helps in speeding up the recovery.

References

- **VMware vCenter Site Recovery Manager Documentation**
https://www.vmware.com/support/pubs/srm_pubs.html
- **VMware vCenter Site Recovery Manager Evaluator's Guide**
<https://www.vmware.com/files/pdf/products/SRM/VMware-vCenter-Site-Recovery-Manager-Evaluation-Guide.pdf>
- **Site Recovery Manager Administration Guide**
https://www.vmware.com/pdf/srm_admin_5_0.pdf
- **VMware vCenter Site Recovery Manager Resources for Business Continuity**
<http://www.vmware.com/products/site-recovery-manager/resource.html>

About the Author

Aalap Desai is a Member of Technical Staff at VMware. He has been working on the performance of VMware vCenter Site Recovery Manager. Aalap received his master's degree in Computer Science from Syracuse University.

Acknowledgements

Thanks to Deepak Bobbarjung, John Liang, Glenn McElhoe, Michael White, Julie Brodeur and the SRM team for their input and reviews on various drafts of this paper.

