



Site Recovery Manager Evaluation Guide

VMware BC/DR

Table of contents

Site Recovery Manager Evaluation Guide	3
Introduction	3
Overview	3
Terminology	3
About This Evaluation Guide	4
Requirements	5
Overview	5
Evaluation Workflow	7
Overview	7
Evaluation Checklist	7
Exercise	8
Exercise 1: Pairing Sites	8
Exercise 2: Configure Inventory Mappings	10
Exercise 3: Configure Placeholder Datastore	12
Exercise 4: Add Array Manager & Enable Array Pair	12
Exercise 5: Create a Protection Group	13
Exercise 6: Create a Recovery Plan	15
Exercise 7: Testing a Recovery Plan	17
Exercise 8: Running a Recovery Plan	18
Exercise 9: Reprotect a Recovery Plan & Fail Back	19
Exercise 10: Virtual Machine Recovery Properties	20
Conclusion	24
Overview	24

Site Recovery Manager Evaluation Guide

Introduction

Site Recovery Manager is a site migration and disaster recovery solution from VMware.

Overview

Site Recovery Manager is a site migration and disaster recovery solution from VMware. It is fully integrated with VMware vCenter Server™ and VMware vSphere® Web Client. Site Recovery Manager provides orchestration and non-disruptive testing of centralized recovery plans. Site Recovery Manager works in conjunction with various replication solutions including VMware vSphere Replication™ to automate the process of migrating and recovering virtual machine workloads.

Multiple recovery plans can be configured to migrate individual applications and entire sites providing finer control over what virtual machines are failed over and failed back. This also enables flexible testing schedules. For example, one application owner requires quarterly disaster recovery testing while another application owner must test once per month. This is easily accomplished with Site Recovery Manager.

Sites that share stretched storage can take advantage of zero-downtime virtual machine migrations. Site Recovery Manager can orchestrate the live migration of virtual machines using Cross-vCenter vMotion also known as “Long Distance vMotion.”

Storage policy protection groups enable automatic protection of virtual machines residing on array-replicated storage. Items such as networks, folders, and resource pools are mapped between sites in Site Recovery Manager to further automate the migration and recovery of virtual machines between sites. Utilizing VMware NSX™ universal logical switches with Site Recovery Manager enables automatic mapping of networks and virtual machine security policies across sites. NSX supports the spanning of layer 2 networks eliminating the need to customize virtual machine IP address settings during failover and migration. These features reduce complexity, improve reliability, and minimize recovery times.

Site Recovery Manager roles can be assigned to specific individuals and groups in vCenter Server. For example, an administrator might wish to allow several application owners to test recovery plans, but limit the actual migration and failover of virtual machines to just a few individuals in the organization. Site Recovery Manager also includes vCenter Server alarms for monitoring and alerting.

The key features provided by Site Recovery Manager are:

- Integration with third-party array replication technologies and vSphere Replication
- Protection for virtually any workload regardless of operating system and application
- Centralized recovery plans with predefined virtual machine startup sequences
- Workflows for planned migration, disaster avoidance, and disaster recovery
- Automated IP address customization
- Non-disruptive recovery plan testing
- Familiar management user interface in vSphere Web Client
- Roles with preconfigured permissions
- Detailed history reports for testing, migration, and failover auditing

Terminology

Recovery time objective (RTO): Targeted amount of time a business process should be restored after a disaster or disruption in order to avoid unacceptable consequences associated with a break in business continuity.

Recovery point objective (RPO): Maximum age of files recovered from backup storage for normal operations to resume if a system goes offline as a result of a hardware, program, or communications failure.

Array replication: Replication across one or more storage controllers, which eliminates the processing overhead from servers.

vSphere Replication: Host-based virtual machine replication technology created by VMware included with vSphere Essentials Plus Kit and higher editions.

Logical unit number (LUN): Number used to identify a logical unit, which is a device addressed by the SCSI protocol or Storage Area Network (SAN) protocols.

Consistency group: One or more LUNs or volumes that are replicated at the same time. When recovering items in a consistency group, all items are restored to the same point in time.

Failover: Method of recovering applications and services to a secondary system when the primary system experiences a failure or disaster.

Failback: Restoring applications and services from a secondary system back to the primary system after a failover has occurred.

Reprotect: Specific to Site Recovery Manager, the process of reversing the direction of replication and enabling recovery plans for a failback event.

Protected virtual machine: Virtual machine that is replicated from one site to another and is included in a Site Recovery Manager recovery plan for failover and failback.

Protected site: Site that contains protected virtual machines.

Recovery site: Site where protected virtual machines are recovered in the event of a failover.

NOTE: It is possible for the same site to serve as a protected site and recovery site when replication is occurring in both directions and Site Recovery Manager is protecting virtual machines at both sites.

Datastore group: One or more datastores that are treated as a unit in Site Recovery Manager. A common example is a consistency group in an array replication solution.

Protection group: Collection of protected virtual machines that are migrated or failed over as a unit.

Storage policy protection group: Protection group configured with a tag-based storage policy that enables automatic protection of a virtual machine in Site Recovery Manager simply by assigning the tag-based storage policy to the virtual machine.

Recovery plan: Documented process to recover a business IT infrastructure in the event of a disaster. A recovery plan in Site Recovery Manager includes one or more protection groups.

Storage replication adapter: Software components provided by array replication vendors that are installed on the Site Recovery Manager servers to enable communication between Site Recovery Manager and array replication solutions.

Placeholder virtual machine: Virtual machine created in the vCenter Server inventory at the recovery site when a virtual machine is protected by Site Recovery Manager. Placeholder virtual machines do not have virtual disks attached to it so the storage capacity consumed by placeholder virtual machines is very small.

Inventory mappings: In Site Recovery Manager, the default networks, folders, and resources for protected virtual machines to use at the recovery site.

NSX universal logical switch: Virtual switch that allows layer 2 networks to span multiple sites.

About This Evaluation Guide

The purpose of this document is to provide a structured guide for IT professionals to evaluate the primary features and benefits of using Site Recovery Manager to automate planned migration and disaster recovery workflows for applications and services running in virtual machines. The exercises in this guide should be completed in the order prescribed for best results. Some exercises have dependencies on previously completed items.

This guide does not contain detailed steps on performing activities such as installation and configuration since these steps are already included in the product documentation.

Requirements

Overview of the requirements in SRM

Overview

It is assumed the following items are already properly installed and configured in a non-production environment designated for this evaluation.

- Domain Name System (DNS) server with forward and reverse lookup enabled.
- Two or more vSphere hosts - a minimum of one designated for each site.
- Two vCenter Server virtual machines - one for each site.
- Two virtual machines with a supported Microsoft Windows operating system and Site Recovery Manager installed - one for each site.

Recommendation: Verify the Windows operating systems for the Site Recovery Manager host virtual machines are compatible with the Site Recovery Manager using the [VMware Compatibility Guides](#). Consult the [Site Recovery Manager Documentation](#) when installing and configuring Site Recovery Manager.

- Array replication or vSphere Replication deployed and configured for use in the evaluation environment.

Recommendation: While array replication supported by Site Recovery Manager can be used for this evaluation, vSphere Replication is recommended for simplicity and compatibility with a wide variety of storage types including VMware vSAN™. A minimum of one vSphere Replication virtual appliance must be deployed and configured for use with the vCenter Server instance at each site. For more information on deploying and configuring vSphere Replication, see the vSphere Replication documentation. vSphere Replication does not require installation of a storage replication adapter.

NOTE: Storage policy protection groups, cross-vCenter vMotion with stretched storage, and NSX integration require array-based replication. These features are not required to successfully complete the steps in this guide.

- One or more Windows and/or Linux based virtual machines with VMware Tools installed, which will be protected by Site Recovery Manager.
- Static IP addresses and DNS host (A) records assigned to all vSphere hosts and virtual machines in the evaluation environment.
- Reliable network connectivity between both sites and all components in the Site Recovery Manager evaluation environment. See [Network Ports for Site Recovery Manager \(2103394\)](#) and [TCP and UDP Ports required to access VMware vCenter Server, VMware ESXi and ESX hosts, and other network components \(1012382\)](#) for more information on network port requirements.
- Adequate storage capacity for all of the components above.

The figure below shows a logical diagram of how the evaluation environment can be configured. Network connectivity is required between the two sites, but they do not have to be geographically separated to satisfy the requirements of the evaluation exercises.




Figure 1: Example evaluation environment.

Recommendation: Use default settings for all components - installation paths, TCP port settings, and so on - wherever possible, to minimize complexity in the evaluation environment. Use consistent naming conventions, usernames, and passwords during evaluation environment deployment.

Recommendation : Use descriptive names for the components such as servers and port groups in a VMware virtualized environment. These names appear in the user interface and Site Recovery Manager history reports. Descriptive names improve the quality of these reports and ease troubleshooting. Use the same naming convention for items such as network port groups at the protected site and the recovery site, as this will simplify inventory mappings.

Evaluation Workflow

The following checklist can be used to track the progress of the evaluation at a high level.

Overview

The following exercises are covered in this document:

- 1. Pairing sites
- 2. Configure inventory mappings
- 3. Configure placeholder datastore
- 4. Add array manager and enable array pair (if using array replication)
- 5. Creating a protection group
- 6. Creating a recovery plan
- 7. Testing a recovery plan
- 8. Running a recovery plan
- 9. Performing fail-back
- 10. Customizing virtual machine recovery properties

The following checklist can be used to track the progress of the evaluation at a high level. The sections after the checklist provide more details on each exercise, including recommendations, documentation references, VMware Knowledge Base articles, and other resources. This document does not contain detailed, step-by-step instructions for completing the tasks in each exercise. These instructions are documented in items such as the Site Recovery Manager documentation. In most cases, one exercise is dependent on another one. For example, a recovery plan cannot be created until at least one protection group is created. Perform the exercises in the order documented in this guide.

Evaluation Checklist

Exercise

It is assumed that Site Recovery Manager has been installed in both sites, a replication solution has been deployed

Exercise 1: Pairing Sites

It is assumed that Site Recovery Manager has been installed in both sites, a replication solution has been deployed, and all virtual machines that will be protected by Site Recovery Manager are being replicated.

Site Recovery Manager is managed using vSphere Web Client. During the installation of Site Recovery Manager, a plugin is installed in vSphere Web Client and an icon labeled “Site Recovery” is displayed.




Figure 2. Site Recovery Manager in vSphere Web Client

The first step in configuring Site Recovery Manager is pairing sites. The most common configuration is pairing two sites - a protected site and a recovery site. That is the configuration prescribed for these evaluation exercises. Site Recovery Manager also supports “shared” sites that consist of a single vCenter Server and multiple Site Recovery Manager servers. An example of this is a shared recovery site for branch offices.




Figure 3. Shared Recovery Site for Branch Offices

See the Site Recovery Manager documentation for details on site pairing. After completing this step successfully, information on the paired sites is displayed in vSphere Web Client.

Site		Paired Site	
Name:	SiteA	Name:	SiteB
Client Connection:	✔ Connected	Client Connection:	✔ Connected
Server Connection:	✔ Connected	Server Connection:	✔ Connected
SRM Server:	10.156.189.5:9086	SRM Server:	10.156.189.7:9086
vCenter Server:	vcentersitea.vsanpe.vmware.com:443	vCenter Server:	vcentersiteb.vsanpe.vmware.com:443
SRM Server Build:	4613745	SRM Server Build:	4613745
Organization:	VMware, Inc.	Organization:	VMware, Inc.
Logged in as:	VSPHERE.LOCAL\Administrator	Logged in as:	VSPHERE.LOCAL\Administrator
VR Compatibility:	✔ 6.5.0.1706 - Compatible	VR Compatibility:	✔ 6.5.0.1706 - Compatible

Figure 4. Paired Sites in vSphere Web Client

Site Recovery Manager also provides a guide to track the progress and assist with configuration. Clicking on a step in the guide takes you to the user interface for that particular activity. A green checkmark is displayed when a step is completed.




Figure 5. Guide to configuring Site Recovery Manager in the User Interface

Exercise 2: Configure Inventory Mappings

Inventory mappings consist of three types: Resource mappings, folder mappings, and network mappings. These mappings provide default settings for recovered virtual machines. For example, a mapping can be configured between a network port group named "Production" at the protected site and a network port group named "Production" at the recovery site. As a result of this mapping, virtual machines connected to "Production" at the protected site will, by default, automatically be connected to "Production" at the recovery site.

There is no issue with having a port group at each site with the same name since each site is managed by a separate vCenter Server instance. Having port groups at each site with the same name eases Site Recovery Manager configuration. If port groups at the protected and recovery site have different names, the mappings must be created manually.

Select creation mode

Select the way in which you want to create mappings.

- ☒ Automatically prepare mappings for networks with matching names
- The system will automatically prepare mappings for networks with matching names under the selected network containers on "pod02vm01.pml.local" and "wdcpod06vm01.pml.local".
- ☐ Prepare mappings manually
- Manually select networks from "pod02vm01.pml.local" to be mapped to a specific network on "wdcpod06vm01.pml.local".

Figure 6. Network Mapping Configuration

Folder mappings can also be configured automatically when the names are the same. Resource mappings are configured manually. Reverse mappings can be created automatically. This provides default settings for both failover and failback operations.

Prepare reverse mappings

Select configured mappings for which to automatically create reverse mappings.

Automatically create reverse mappings on the paired site. This may override already existing mappings on the paired site. (Only for 1-1 mappings)

Select all applicable

Q Filter

	pod02vm01.pml.local	1 ▲	wdcpod06vm01.pml.local
☒	 pa > Production		wdc > Production

Figure 7. Reverse Mappings

It is possible to have multiple items at the protected site mapped to a single item at the recovery site. For example, two resource pools at the protected site can be mapped to a single resource pool at the recovery site. However, this prevents automatic creation of reverse mappings.

Recommendation: Provide the same name to folders and network port groups with similar functionality at the protected and recovery sites so that mappings can be prepared automatically. Use 1-1 mappings so that reverse mappings can be utilized. These practices will ease inventory-mapping configuration and minimize complexity in the environment.

Default network mappings can be overridden on a per-VM basis. Selecting a virtual machine in the protection group and clicking the Configure Protection button enables the user to change the virtual machine protection properties.




Figure 8. Change the Virtual Machine Protection Properties

Exercise 3: Configure Placeholder Datastore

Site Recovery Manager creates a placeholder virtual machine at the recovery site for every protected virtual machine. Placeholder virtual machines are contained in a datastore and registered with vCenter Server at the recovery site. This datastore is referred to as a “placeholder datastore”. Placeholder virtual machines do not have virtual disks (VMDK files) so they consume minimal storage capacity.

Create a small datastore that is accessible by all hosts at the recovery site for use as a placeholder datastore. Create a similar datastore at the protected site, as well. At least one placeholder datastore is required at each site to utilize the failover and failback functionality in Site Recovery Manager. If you are using array replication, do not configure replication for the placeholder datastores.

It is possible to configure multiple placeholder datastores at each site. Typically, one placeholder datastore at each site is sufficient. Multiple placeholder datastores may be beneficial in larger environments such as a site with multiple vSphere clusters or a shared recovery site.




Figure 9. Placeholder Datastores in the Site Recovery Manager User Interface

Exercise 4: Add Array Manager & Enable Array Pair

This exercise is necessary only if using array replication. If you are using vSphere Replication, this exercise can be skipped.

When using array replication, a storage replication adapter is required for the specific array replication solution to be used with Site Recovery Manager. Storage replication adapters are software components that are produced and supported by the array replication vendors. The Site Recovery Manager compatibility guide on VMware’s web site should be used to determine if a storage replication adapter is available for the array replication technology in the evaluation environment. Only storage replication adapters downloaded from vmware.com should be used to ensure compatibility and support.

Recommendation: Use vSphere Replication for this evaluation. While array replication has some advantages over vSphere Replication, it is also more complex to install and configure and it usually requires additional licensing from the array replication vendor. Using vSphere Replication has a number of advantages including simple management using vSphere Web Client, support for virtually any storage supported by vSphere including vSAN, the ability to configure replication on a per-VM basis, and vSphere

Replication is included with vSphere Essentials Plus Kit and higher editions. For a full comparison of array replication and vSphere Replication, see [SRM - Array Based Replication vs. vSphere Replication](#).

See “Configure Array Managers” in the Site Recovery Manager documentation for guidance when working with array managers. Thoroughly read the documentation - release notes, installation guide, etc. - that is typically included with a storage replication adapter. Most storage replication adapters have specific requirements that are outlined in this documentation, not in the Site Recovery Manager documentation.

After array managers have been successfully configured, paired, and enabled, information about the array replication can be seen in the Site Recovery Manager user interface. This information includes items such as the local and remote array names, local and remote devices, and the direction in which replication is currently occurring.

Local Array	Remote Array	Status	Local Array Manager	Remote Array Manager
 group01	group02	 Enabled	vsa01	vsa02

Array Pair: group01 - group02

Errors: None

Local Device	Status	Remote Device	Datastore
volume01	 Outgoing Replication	volume02	Local: [iscsi1]

Figure 10. Array Replication in the Site Recovery Manager User Interface

When using array replication, only virtual machines that will be protected by Site Recovery Manager should be placed on a replicated LUN or volume. Combining protected and unprotected virtual machines on the same replicated LUN or volume will produce warning messages in Site Recovery Manager.

Review [How Site Recovery Manager Computes Datastore Groups](#) in the Site Recovery Manager documentation to gain a thorough understanding of how datastore groups are composed when array replication is utilized.

Exercise 5: Create a Protection Group

Before a protection group can be created, replication must be configured. If you have not configured array replication or vSphere Replication for the virtual machines that Site Recovery Manager will protect, you will need to do so before proceeding.

Details on deploying and configuring vSphere Replication can be found in the [vSphere Replication documentation](#).

A protection group is a collection of one or more virtual machines that are failed over and failed back as a unit. In many cases, a protection group consists of multiple virtual machines that support a service such as an accounting system. For example, a service might consist of a database server, two application servers, and two web servers. In most cases, it is not be beneficial to fail over part of a service (only one or two of the servers in the example). All five servers would be included in a protection group to enable failover of the service.

Creating a protection group for each application or service also has the benefit of selective testing. With Site Recovery Manager, having a protection group for each application enables non-disruptive, low-risk testing of individual applications. Application owners can test disaster recovery plans, as needed.

Larger environments usually have higher numbers of applications. Creating a protection group for each application in these larger environments may not be practical and might exceed the maximum supported number of protection groups in Site Recovery Manager. Please see Operational Limits for [Site Recovery Manager \(2105500\)](#) for details.

There are other organizational methods to consider when creating protection groups. One is creating a protection group for each business unit - all virtual machines belonging to a specific business unit are placed in a protection group. Another method is grouping virtual machines together by application tier. For example, all database servers in one protection group, all middleware servers in a second protection group, and all client-facing servers in a third protection group. While these approaches have their limitations, they also reduce the number of protection groups to create and manage.

	Email	Business Intelligence	Accounting
Protection Group 1	Mailbox (database)	Warehouse	Records (database)
Protection Group 2	Hub Transport	Application	
Protection Group 3	Client Access	Web	Application

Figure 11. Creating Protection Groups by Application Tier

There is no recommendation for the number of protection groups to create as this varies with each organization depending on business and technical requirements. An organization must decide what method is best for its purposes. More protection groups increase the flexibility of testing and failover while fewer protection groups lowers complexity.




Figure 12. Protection Groups: The Balance of Flexibility and Complexity

When array replication is utilized with Site Recovery Manager, protection groups are created based on the datastore groups that are available. Consider this example: There are two LUNs that make up a consistency group in the array replication solution. Both LUNs contain virtual machines for a total of five virtual machines. The two LUNs are treated as a single datastore group in Site Recovery Manager. When a protection group is created, all 5 virtual machines are included. It is not possible to add a subset of virtual machines in the datastore group to a protection group.




Figure 13. Protection Groups and Datastore Groups with Array Replication

Recommendation: When using array replication, organize virtual machines that are to be tested and failed over together on the same array replicated LUN, volume, or consistency group. For example, an application consisting of five virtual machines should be grouped together to ensure consistent recovery. If there are other virtual machines that are not part of the application, they should be migrated to other LUNs or volumes. Having a mix of virtual machines protected by Site Recovery Manager and unprotected virtual machines in the same replicated LUN, volume, or consistency group will produce warning messages in Site Recovery Manager.

It is possible to utilize array replication and vSphere Replication in the same Site Recovery Manager environment. However, a protection group can only contain virtual machines replicated by one replication technology or the other. A recovery plan can contain a mixture of protection groups based on array replication and protections groups based on vSphere Replication.

Create Protection Group

✓ 1 Name and location

2 Protection group type

3 Virtual machines

4 Ready to complete

Protection group type
Select the protected site and protection group type.

Protected site: ☐ wdcpod06vm01.pml.local
☒ pod02vm01.pml.local

Replication type: ☐ Array Based Replication (ABR)
☒ vSphere Replication (VR)

Figure 14. Creating a Protection Group

After a protection has been created and the virtual machines have been successfully protected, placeholder virtual machines will be visible in the vCenter Server inventory at the recovery site. The presence of placeholders provides a visual indication to Site Recovery Manager administrators that virtual machines are protected. Placeholder virtual machines have a unique icon in vSphere Web Client.

Name	State	Status	Provisioned Space	Used Space	Host CPU
acct01	Powered Off	✓ Normal	2.78 GB	0 B	0 MHz
acct02	Powered Off	✓ Normal	2.78 GB	0 B	0 MHz
acct03	Powered Off	✓ Normal	1.28 GB	0 B	0 MHz
acct04	Powered Off	✓ Normal	2.78 GB	0 B	0 MHz
dc01	Powered On	✓ Normal	64.19 GB	29.75 GB	17 MHz
dc03	Powered On	✓ Normal	76.19 GB	32.33 GB	0 MHz
mysql01	Powered On	✓ Normal	20.42 GB	20.42 GB	0 MHz

Figure 15. Placeholder Virtual Machines.

For more details, see [Creating and Managing Protection Groups](#) in the Site Recovery Manager documentation.

Exercise 6: Create a Recovery Plan

After you configure Site Recovery Manager at the protected and recovery sites and you have created at least one protection group, you can create a recovery plan. A recovery plan controls every step of the recovery process, including the order in which virtual machines are recovered, IP address changes, and so on. Protection groups are created at the protected site. As one might expect, recovery plans are created at the recovery site.

A recovery plan must contain one or more protection groups. A protection group can be part of more than one recovery plan. For

example, there are two protection groups: Accounting and Email. Three recovery plans can be created: The Accounting recovery plan containing the Accounting protection group, the Email recovery plan containing the Email protection group, and the Entire Site recovery plan containing both protection groups.




Figure 16. Protection Groups and Recovery Plans

Creating a recovery plan for each application enables testing and failover of individual applications. Creating a recovery plan that includes all of the protection groups is useful when testing or failing over all applications at a site.

NOTE: Concurrently testing or running multiple recovery plans that contain the same protection group will produce error messages in the Site Recovery Manager.

One of the steps in creating a recovery plan is configuring a test network for each of the recovery networks at the recovery site. When testing a recovery plan, recovered virtual machines are connected to a corresponding test network, as defined in a recovery plan. By default, Site Recovery Manager uses an isolated network - a virtual switch with no external connectivity that is created dynamically when a recovery plan is tested. This is a simple and effective approach, but connectivity between virtual machines is limited to the virtual machines running on the same vSphere host.

Test networks

Select the networks to use while running tests of this plan.

Recovery Network	Test Network
 wdc > Production	 Isolated network (auto created)

Figure 17. Default Isolated Test Network in a Recovery Plan

The default setting, “Isolated network (auto created)”, can be changed to another port group available at the recovery site. For example, a Virtual LAN (VLAN) can be configured at the recovery site that is available to all vSphere hosts, but not routed to production networks. This enables connectivity between virtual machines on different hosts at the recovery site without interfering with other virtual machines that are not being tested.

Test networks

Select the networks to use while running tests of this plan.

Recovery Network	Test Network
 wdc > Production	 TestVLAN

Figure 18. Portgroup Configured as the Test Network in a Recovery Plan

Recommendation: When implementing Site Recovery Manager in a production environment, utilize a dedicated VLAN at the recovery site for recovery plan testing. This VLAN should be available to all vSphere hosts at the recovery site, but routing should not be enabled to any other networks to prevent interference with workloads outside of the Site Recovery Manager test environment. This is not a requirement for evaluation.


Additional virtual machine recovery properties can be configured such as creating virtual machine dependencies, adding visual prompts, and changing IP addresses. Virtual machine recovery properties will be covered in a later exercise.

See [Create a Recovery Plan](#) in the Site Recovery Manager documentation.

Exercise 7: Testing a Recovery Plan

After creating a recovery plan, it is beneficial to test the recovery plan to verify it works as expected. Site Recovery Manager features a non-disruptive testing mechanism to facilitate testing at any time. It is common for an organization to test a recovery plan multiple times after creation to resolve any issues encountered the first time the recovery plan was tested.

Verify the recovery plan is ready for testing or running by checking the “Plan status”. It should show “Ready”. Click the green arrow below “Description” to begin the test process.



Recovery Step	Status
1. Synchronize Storage	

Figure 19. Test Recovery Plan

When testing a recovery plan, there is an option to replicate recent changes, which is enabled by default. Replicating recent changes will provide the latest data for the testing process. However, it will also lengthen the amount of time required to recover virtual machines in the recovery plan, as replication has to finish before the virtual machines are recovered.

A question often asked is whether replication continues during the test of a recovery plan. The answer is yes. Site Recovery Manager utilizes snapshots - either array snapshots (or clones) with array replication or virtual machine snapshots with vSphere Replication - as part of the recovery plan test process. This approach allows powering on and modifying virtual machines recovered as part of the test while replication continues to avoid RPO violations.

Name
vrTestImage-GID-b2c22e43-c289-4573-a405-a5ea5c4959fe
vrTestImage-GID-68eac7cd-288e-4b7e-9507-8777ae5193d8
vrTestImage-GID-0e835187-2116-4dac-be37-3b4de09a5679
vrTestImage-GID-b6288da5-bad6-43a3-9b81-5eb163fde31a

Figure 20. vSphere Replication Test Image Folders on vSAN

Virtual machines that are in a recovery plan that is being tested will display unique icons in the vCenter Server inventory at the recovery site.



Name	State	Status	Provisioned Space	Used Space
acct01	Powered On	✓ Normal	34.19 GB	6.17 GB
acct02	Powered On	✓ Normal	34.19 GB	6.06 GB
acct03	Powered On	✓ Normal	32.67 GB	3.38 GB
acct04	Powered On	✓ Normal	34.19 GB	6.03 GB
dc01	Powered On	✓ Normal	64.19 GB	29.75 GB

Figure 21. Powered on Virtual Machines in a Recovery Plan Test

At this point, guest operating system administrators and application owners can log into their recovered virtual machines to verify functionality, perform additional testing, and so on. Site Recovery Manager easily supports recovery plan testing periods of varying lengths - from a few minutes to several days. However, longer tests tend to consume more storage capacity at the recovery site. This is due to the nature of snapshot growth as data is written to the snapshot.

Recommendation: Closely monitor storage capacity utilization at the recovery site during recovery plan tests, if capacity is limited. Configure vCenter Server alarms to alert administrators when free space is getting low on datastores at the recovery site.

See [Test a Recovery Plan](#) in the Site Recovery Manager documentation.

When testing is complete, a recovery plan must be “cleaned up”. This operation powers off virtual machines and removes snapshots associated with the test. Once the cleanup workflow is finished, the recovery plan is ready for testing or running.



Recovery Step	Status
1. Synchronize Storage	✓ Success

Figure 22. Recovery Plan Test Cleanup

See [Clean Up After Testing](#) a Recovery Plan in the Site Recovery Manager documentation.

History Reports

When workflows such as a recovery plan test and cleanup are performed in Site Recovery Manager, history reports are automatically generated. These reports document items such as the workflow name, execution times, successful operations, failures, and error messages. History reports are useful for a number of reasons including internal auditing, proof of disaster recovery protection for regulatory requirements, and troubleshooting. Reports can be exported to HTML, XML, CSV, or a Microsoft Excel or Word document.

Plan Name	Operation	Result	Date	Duration	User
Accounting	Cleanup	✓ Success	April 24, 2015 at 8:02:32 AM E...	13 s	VMWARE\jhunter
Accounting	Test	✓ Success	April 24, 2015 at 6:59:52 AM E...	1 m 22 s	VMWARE\jhunter

Figure 23. History Reports

See [View and Export a Recovery Plan History](#) in the Site Recovery Manager documentation.

Exercise 8: Running a Recovery Plan

Running a recovery plan differs from testing a recovery plan. Testing a recovery plan does not disrupt virtual machines at the protected site. When running a recovery plan, Site recovery Manager will attempt to shut down virtual machines at the protected site before the recovery process begins at the recovery site. Recovery plans are run when a disaster has occurred and failover is

required or when a planned migration is desired.



Figure 24. Run Recovery Plan

Clicking the Run Recovery Plan button opens a confirmation window requiring the selection of a recovery type - either a planned migration or a disaster recovery. In both cases, Site Recovery Manager will attempt to replicate recent changes from the protected site to the recovery site. It is assumed that for a planned migration, no loss of data is the priority. A planned migration will be cancelled if errors in the workflow are encountered. For disaster recovery, the priority is recovering workloads as quickly as possible after disaster strikes. A disaster recovery workflow will continue even if errors occur. The default selection is a planned migration.

After a recovery type is selected, the operator must also populate a confirmation checkbox as an additional safety measure. The idea behind this checkbox is to make sure the operator knows that he or she is running (not testing) a recovery plan.

The first step in running a recovery plan is the attempt to synchronize storage. Then, protected virtual machines at the protected site are shut down. This effectively quiesces the virtual machines and commits any final changes to disk as the virtual machines complete the shutdown process. Storage is synchronized again to replicate any changes made during the shutdown of the virtual machines. Replication is performed twice to minimize downtime and data loss. Once these steps have been completed, the recovery process at the recovery site is started.

If the protected site is offline due to a disaster, for example, the disaster recovery type should be selected. Site Recovery Manager will still attempt to synchronize storage as described in the previous paragraph. Since the protected site is offline, Site Recovery Manager will begin recovering virtual machines at the recovery site using the most recently replicated data.

See [Run a Recovery Plan](#) in the Site Recovery Manager documentation.

Since running a recovery plan is a disruptive operation, Site Recovery Manager administrators commonly limit the ability to run recovery plans to just a few people in the organization. This is accomplished through Site Recovery Manager roles and permissions that are added to vCenter Server when Site recovery Manager is installed. For example, an administrator can assign the “SRM Recovery Test Administrator” role to application owners allowing these individuals to test recovery plans for their applications, but not run recovery plans.

Recommendation: Considering the disruptive nature of running (not testing) a recovery plan, limit the permission to run a recovery plan to only a few individuals in the organization similar to the way root or domain administrator permissions are typically limited. All individuals with this permission should be fully trained on the operation of Site Recovery Manager. However, more than one person should have this permission to avoid a single point of failure.

There are several roles and permissions available. For more information on roles and permissions, see [Site Recovery Manager Privileges, Roles, and Permissions](#) in the Site Recovery Manager documentation.

Exercise 9: Reprotect a Recovery Plan & Fail Back

4. 9 Exercise 9: Reprotect a Recovery Plan & Fail Back

Site Recovery Manager features the ability to not only fail over virtual machine workloads, but also fail them back to their original site. However, this assumes that the original protected site is still intact and operational. An example of this is a disaster avoidance situation: The threat could be rising floodwaters from a major storm and Site Recovery Manager is used to migrate virtual machines from the protected site to the recovery site. Fortunately, the floodwater subsides before any damage was done leaving the protected site unharmed.

A recovery plan cannot be immediately failed back from the recovery site to the original protected site. The recovery plan must first undergo a reprotect workflow. This operation involves reversing replication and setting up the recovery plan to run in the opposite direction.




Figure 25. Reprotect Recovery Plan

Reprotecting a recovery plan can take a considerable amount of time depending on the number of protection groups and virtual machines in the recovery plan and the amount of data that must be replicated to resynchronize storage. Upon completion of the reprotect workflow, a history reports will be created and the recovery plan can be failed back. Essentially, the original recovery site becomes the protected site and the original protected site becomes the recovery site for the virtual machines in the recovery plan. Run the recovery plan to fail back the virtual machines to their original protected site.

NOTE: Be sure to reprotect a recovery plan after it has been run (virtual machines have been failed over or failed back). Failure to do this important step will prevent future testing and running of the recovery plan until the reprotect workflow has been run.

Recommendation: Test a recovery plan as soon as possible after a reprotect workflow has run to verify the recovery plan will work properly.

See [Reprotecting Virtual Machines After a Recovery](#) in the Site Recovery Manager documentation.

Exercise 10: Virtual Machine Recovery Properties

Site Recovery Manager includes several features to customize recovery for virtual machines. Examples include the options to change the IP address of virtual network interface cards, run scripts, and control the power state of virtual machines after they are recovered. It is also possible to organize virtual machines into one of five priority groups. Since these are recovery settings, these settings are accessed in the recovery plan user interface of Site Recovery Manager.




Figure 26. Change the Virtual Machine Recovery Settings

NOTE: Changes to virtual machine configuration properties apply to the virtual machine in all recovery plans. For example, if a virtual machine is configured as a member of priority group 3, it will appear in priority group 3 in all recovery plans.

In this exercise, there are no specific steps that must be performed. The recommendation is to read about and experiment with the various settings discussed below to understand the virtual machine recovery options available in Site Recovery Manager.

Priority Groups

There are five priority groups in Site Recovery Manager. As one might expect, the virtual machines in priority group 1 are recovered first, then the virtual machines in priority group 2 are recovered, and so on. This provides administrators one option for prioritizing the recovery of virtual machines. For example, the most important virtual machines with the lowest RTO are typically placed in the first priority group and less important virtual machines in subsequent priority groups. Another example is by application tier - database servers could be placed in priority group 2; application and middleware servers in priority group 3; client and web servers in priority group 4.

Dependencies

Virtual machines in the same priority group are started in parallel. There is no startup order guaranteed within a priority group unless one or more dependencies are defined. A dependency simply instructs one virtual machine to start before another. For example, a virtual machine named "acct02" can be configured to have a dependency on a virtual machine named "acct01" - Site Recovery Manager will wait until "acct01" is started before powering on "acct02". VMware Tools heartbeats are used to validate when a virtual machine has started successfully.

*VM Dependencies			
Start the following VMs before this VM.			
Virtual Machine	Status	Priority Group	Protection Group
 acct01	OK	3	Accounting

Figure 27. Virtual Machine Dependencies

NOTE: A virtual machine dependency is ignored if the virtual machines are not in the same priority group.

Recommendation: Limit the use of dependencies where possible to minimize the amount of time required to recover virtual machines.

Shutdown Actions

Shutdown actions apply to the protected virtual machines at the protected site during the run of a recovery plan. Shutdown actions are not used during the test of a recovery plan. By default, Site Recovery Manager will issue a guest OS shutdown, which requires VMware Tools and there is a time limit of five minutes. The time limit can be modified. If the guest OS shutdown fails and the time limit is reached, the virtual machine is powered off. Shutting down and powering off the protected virtual machines at the protected site when running a recovery plan is important for a few reasons:

- Quiesces the guest OS and applications before the final storage synchronization occurs
- Avoids the potential conflict of having virtual machines with duplicate network configurations (hostname, IP addresses) on the same network

Optionally, the shutdown action can be changed to simply power off virtual machines. Powering off virtual machines does not shut them down gracefully, but this option can reduce recovery times in situations where the protected site and recovery site maintain network connectivity during the run (not test) of a recovery plan. An example of this is a disaster avoidance scenario.

Recommendation: In most cases, minimizing risk and data loss are higher priorities than recovery time. Keep the default Shutdown Action setting of “Shutdown guest OS before power off” to properly quiesce the guest OS and applications, where possible, during a planned migration and disaster recovery.

Startup Actions

A startup action applies to a virtual machine that is recovered by Site Recovery Manager. Powering on a virtual machine after it is recovered is the default setting and this is typically not changed. In some cases, it might be desirable to recover a virtual machine, but leave it powered off. Startup actions are applied when a recovery plan is tested or run.

With the default setting of “Power on”, it is possible to configure the amount of time Site Recovery Manager waits for VMware Tools heartbeats before issuing an error message. VMware Tools heartbeats are used to validate a virtual machine started successfully. The default timeout value is five minutes. Changing the timeout value for this setting might be useful for virtual machines that take longer to start up. For example, if a virtual machine takes six minutes to fully boot, an error message would be produced even though the virtual machine is recovered without issue. Changing the timeout value to more than six minutes would eliminate this “false positive” error message.

Another configurable option in this section is the delay before running a post power on step, which will be covered next. A common example of a post power on step is running a script in the guest OS of a virtual machine. A delay might be needed to provide adequate time for a system service to start before running a script.

Pre and Post Power On Steps

Site Recovery Manager can run a command from the Site Recovery Manager server at the recovery site before and after powering on a virtual machine. A common use case is calling a script to perform actions such as making changes to DNS and modifying application settings on a physical server. Running a script inside of a virtual machine is also supported as a post power on step.

Site Recovery Manager can also display a visual prompt as a pre or post power on step. This prompt might be used to remind an operator to place a call to an application owner, modify the configuration of a router, or verify the status of a physical machine.

Plan status:	 Waiting for user input 68 %
Description:	The test has paused at a user prompt. You must dismiss the prompt to resume the test.
Prompts:	 Verify Database is Online Dismiss Verify the database services have started on UNIX machine uxdb01.vmware.local and the data warehouse is mounted.

Figure 28. Visual Prompt in a Recovery Plan

NOTE: A prompt pauses a recovery plan until an operator clicks “dismiss” in the Site Recovery Manager user interface.

IP Customization

The most commonly modified virtual machine recovery property is IP customization. The majority of organizations have different IP address ranges at the protected and recovery sites. When a virtual machine is failed over, Site Recovery Manager can automatically change the network configuration (IP address, default gateway, and so on) of the virtual network interface card(s) in the virtual machine. This functionality is available in both failover and fail-back operations.

There are multiple IP customization modes in Site Recovery Manager. For example, it is possible to create an IP customization rule that maps one range of IP addresses to another. In the figure below, an administrator has mapped 10.10.10.0/24 to 10.10.20.0/24.

Subnet:	10 . 10 . 10 . 0 / 24	10 . 10 . 20 . 0 / 24
Subnet mask:	255.255.255.0	255.255.255.0
Range:	10.10.10.0 - 10.10.10.255	10.10.20.0 - 10.10.20.255

Figure 29. IP Customization Rule

Continuing with the example above, a virtual machine containing one virtual network interface card with IP address 10.10.10.50 will be reconfigured to use IP address 10.10.20.50 when it is failed over. When the same virtual machine is failed back to the original protected site, Site Recovery Manager will change the IP address back to 10.10.10.50.

IP customization can also be configured manually for a virtual machine. Network settings such as the IP address, DNS server(s), and DNS suffixes can be modified on a per-VM basis for failover and fail-back.




Figure 30. Manual IP Customization

The two methods above are the most popular ways to customize network settings when migrating and failing over virtual machines. The third option is configuring customization using the DR IP Customizer Tool, which involves the use of a CSV file. For more information on this method, see [Customizing IP Properties for Multiple Virtual Machines By Using the DR IP Customizer Tool](#) in the Site recovery Manager documentation.

Recommendation: Install VMware Tools in all protected virtual machines to help ensure Site Recovery Manager can properly shut down supported guest OS's, report the power status of virtual machines using VMware Tools heartbeats, execute scripts inside of virtual machines, and perform IP customization.

Conclusion

The primary features of Site Recovery Manager are covered in this evaluation guide.

Overview

The primary features of Site Recovery Manager are covered in this evaluation guide. There are additional items such as implementing Site Recovery Manager across multiple sites, automation with [VMware vRealize™ Orchestrator™](#), and integration with [VMware vRealize Automation™](#) that are not discussed. More details on these solutions can be found in VMware Documentation and in other resources such as [VMware Blogs](#) and the [VMware Knowledge Base](#).

Site Recovery Manager provides the following benefits:

- Lower cost of disaster recovery by up to 50% by reducing the x86 hardware footprint, simplifying protection through VM-centric, policy-based replication, and decreasing operational expenses.
- Non-disruptive testing enables more frequent disaster recovery rehearsals, which minimizes risk.
- Recovery plans that can easily be updated and tested as new workloads are provisioned.
- Flexibility to protect nearly any workload that runs in a virtual machine.
- Automated customization of virtual machines to further improve recovery times and reduce risk.
- Detailed reporting to easily satisfy regulatory requirements for documentation of data protection and disaster recovery plans.

Hopefully, it is clear how these benefits are achieved after completing the exercises in this evaluation guide. More examples of how VMware customers have benefited from Site Recovery Manager can be found in [several customer success stories](#).

