Site Recovery Manager Administration Guide

vCenter Site Recovery Manager 4.0

EN-000182-00

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

Copyright © 2008, 2009 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About This Book

VMware® vCenter Site Recovery Manager (SRM) is an extension to VMware vCenter that enables integration with array-based replication, discovery and management of replicated datastores, and automated migration of inventory from one vCenter to another. SRM servers coordinate the operations of the replicated storage arrays and vCenter servers at the two sites so that, as virtual machines at one site (the protected site) are shut down, virtual machines at the other site (the recovery site) start up and, using the data replicated from the protected site, assume responsibility for providing the same services. Migration of protected inventory and services from one site to the other is controlled by a recovery plan that specifies the order in which virtual machines are shut down and started up, the compute resources they are allocated, and the networks they can access. SRM enables you to test a recovery plan, using a temporary copy of the replicated data, in a way that does not disrupt ongoing operations at either site.

Intended Audience

This book is intended for Site Recovery Manager administrators who are familiar with vSphere and its replicated datastores, and who want to configure protection for vSphere inventory. It may also be appropriate for other users who need to add virtual machines to protected inventory or verify that existing inventory is properly configured for use with SRM.

Document Feedback

VMware welcomes your suggestions for improving our documentation. If you have comments, send your feedback to docfeedback@vmware.com.

Technical Support and Education Resources

The following technical support resources are available to you. To access the current version of this book and other books, go to http://www.vmware.com/support/pubs.

Online and Telephone Support

To use online support to submit technical support requests, view your product and contract information, and register your products, go to http://www.vmware.com/support.

Customers with appropriate support contracts should use telephone support for the fastest response on priority 1 issues. Go to http://www.vmware.com/support/phone_support.html.

Support Offerings

To find out how VMware support offerings can help meet your business needs, go to http://www.vmware.com/support/services.

VMware Professional Services

VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting Services provides offerings to help you assess, plan, build, and manage your virtual environment. To access information about education classes, certification programs, and consulting services, go to http://www.vmware.com/services.

Administering VMware vCenter Site Recovery Manager 1

VMware vCenter Site Recovery Manager is a business continuity and disaster recovery solution that helps you plan, test, and execute a scheduled migration or emergency failover of vCenter inventory from one site to another.

This chapter includes the following topics:

Protected Sites and Recovery Sites

In a typical SRM installation, the protected site provides business-critical datacenter services. The recovery site is an alternate facility to which these services can be migrated.

The protected site can be any site where vCenter supports a critical business need. The recovery site can be located thousands of miles away or in the same room. The recovery site is usually located in a facility that is unlikely to be affected by environmental, infrastructure, or other disturbances that affect the protected site.

SRM has the following requirements for the VMware vSphere® configurations at each site:

- Each site must have at least one datacenter.
- The recovery site must support array-based replication with the protected site.
- The recovery site must have hardware, network, and storage resources that can support the same virtual machines and workloads as the protected site.
- At least one virtual machine must be located on a replicated datastore at the protected site. This datastore must be supported by a storage array that is compatible with SRM.
- The sites should be connected by a reliable IP network (storage arrays might have additional network requirements).
- The recovery site should have access to the same networks (public and private) as the protected site, although not necessarily the same range of network addresses.

Site Pairing

The protected and recovery sites must be paired before you can use SRM. Site pairing includes three main steps:

1 Exchange of authentication information between the two sites.

2 Discovery of the replicated storage arrays that support the protected site, and discovery of peer arrays at the recovery site.

3 Discovery of the replicated devices supported by the arrays, and mapping of these devices to datastores that support virtual machines.

SRM includes a wizard that guides you through the site-pairing process.

Site pairing requires vSphere administrative privileges at both sites. To initiate the site-pairing process, you must know the username and password of a vSphere administrator at each site.

Array-Based Replication

In array-based replication, one or more storage arrays at the protected site replicate data to peer arrays at the recovery site. Storage replication adapter (SRAs) enable integration of SRM with a wide variety of replicated arrays.

You can configure array-based replication between ESX hosts whether or not you use SRM. In fact, it is a good idea to set up this replication before you install and configure SRM, so that you can be certain it is working correctly before you install SRM and the necessary SRAs.

Storage Replication Adapters

Storage replication adapters are not part of an SRM release. They are developed and supported by your array vendor. You can download storage replication adapters and their documentation from http://www.vmware.com/download/srm/. Storage replication adapters downloaded from other sites are not supported by VMware. You must install an SRA specific to each array that you use with SRM on the SRM server host.

Replicated Storage Devices, Datastores, and Datastore Groups

When you use array-based replication, each storage array supports a set of replicated devices. On Storage Area Network (SAN) arrays that use connection protocols such as Fibre Channel and iSCSI, these devices are called LUNs (logical storage units comprising one or more physical devices). On NFS arrays, they are typically referred to as volumes. In every pair of replicated storage devices, one device is the replication source and the other is the replication target. Data written to the source device is replicated to the target device on a schedule controlled by the arrays' replication software. When you configure SRM to work with an SRA, the protected site is the replication source and the recovery site is the replication target.

A VMFS datastore provides storage for virtual machine files. By hiding the details of physical storage devices, datastores simplify the allocation of storage capacity and provide a uniform model for meeting the storage needs of virtual machines. Because any datastore can span multiple devices, SRM must ensure that all devices in a datastore are replicated before it can protect the virtual machines that use that datastore. When necessary, SRM aggregates datastores into datastore groups to accommodate virtual machines that span multiple datastores. The set of virtual machines that use a replicated datastore or datastore group is called a protection group. Any datastore that supports a protection group is considered a protected datastore.

About Bidirectional Operation

Any pair of sites can be configured to operate bidirectionally, so that each acts as a recovery site for the other. To support bidirectional operation, you must have arrays at each site that are configured as replication targets of the other site. This type of configuration is the only way to provide SRM services for both members of a site pair.

You cannot designate a third site as a recovery site for one that is already paired with another site. If you want to use SRM to provide business continuity and disaster recovery services at a recovery site, you must configure that site as a protected site that uses its own array managers to replicate data to the other member of the site pair. After site pairing is complete, configuring bidirectional operation requires you to follow the same site configuration procedures that are required for unidirectional operation, but you must do so for each site in each capacity. At recovery site that has not been configured for bidirectional operation, items that must be configured at a protected site remain unconfigured:

- Array Managers and Inventory Mappings are always listed as **Not Configured**.
- Protection Groups are listed as **No Groups Created**.

How Changes to Virtual Machine Storage Affect Protection

When you edit the properties of a virtual machine to add or change storage devices (such as hard disks or DVD drives) you can affect the protection of that machine if you connect it to a device that has storage on a datastore that is not replicated, or that is protected by a different protection group.

- If the new device is created on a replicated datastore that is not protected (not part of any protection group), the datastore is added to the virtual machine's protected datastore group and the virtual machine's protection is unaffected.
- If the new device is created on a replicated datastore that is protected by a different protection group, the virtual machine's protection is invalidated. If the new device is created on a nonreplicated datastore, the virtual machine's protection is invalidated.

If you use Storage VMotion to move a virtual machine to a nonreplicated datastore, or to a replicated datastore on an array that SRM has not been configured to manage (through an SRA), the virtual machine's protection is invalidated.

How Site Recovery Manager Computes Datastore Groups

The composition of a datastore group is determined by the set of virtual machines that have files on the datastores in the group, and by the devices on which those files are stored.

Virtual machine files are located on one or more vSphere datastores. Each datastore consists of one or more extents. Each extent corresponds to a single partition of a device on a storage array. Array replication is configured on a per-device basis; most arrays include some devices that are not replicated. SRM must ensure that all devices containing protected virtual machine files are replicated. During a recovery or test, SRM must failover all such devices together.

To solve this problem, SRM aggregates datastores into datastore groups. A datastore group consists of the minimal set of devices required to ensure that if any of a virtual machine's files is stored on a device in the group, all of the virtual machine's files are stored on devices that are part of the same group. For example, if a virtual machine has disks on two different datastores, then both datastores must be combined into a datastore group. Conditions that can cause datastores to be combined into a datastore group include:

- A virtual machine has files on two different datastores.
- Two virtual machines share an RDM device on a SAN array.
- Two datastores span extents corresponding to different partitions of the same device.
- A single datastore spans two extents corresponding to partitions of two different devices.
- Multiple devices belong to a consistency group defined on the storage array.

SRM computes datastore groups when you first configure your array managers. After that, the computation executes every time that a virtual machine is added to or removed from a datastore that is part of a group.

About Protection Groups and Recovery Plans

A protection group is a collection of virtual machines and templates that use the same replicated datastore or datastore group. A recovery plan specifies how the virtual machines in a protection group are recovered.

When the replicated devices that support a datastore group failover, that operation affects all of the virtual machines and templates that use the datastores in the group. Because of this, SRM considers all of those virtual machines and templates members of a single protection group. When you create a protection group, it initially contains only those members that store all of their files on the selected datastore group. You can add members to the protection group by creating them on that datastore, or by using Storage VMotion to move their storage onto it. You can remove a member from a protection group by moving it to another datastore.

Recovery Plans and Replicated Datastores

A recovery plan is like an automated runbook. It controls every step of the recovery process, including the order in which virtual machines are powered off or powered on, the network addresses that recovered virtual machines use, and so on. Recovery plans are flexible and easy to customize.

A recovery plan applies to one or more protection groups. A protection group can be specified in more than one recovery plan. For example, you could create one recovery plan to handle a planned migration of services from the protected site to the recovery site, and another to handle an unplanned event such as a power failure or natural disaster.

A protection group can be recovered by only one recovery plan at a time. If multiple recovery plans that specify the same protection group are tested or run simultaneously, only one recovery plan will actually be able to failover the protection group. Other running recovery plans that specify the same protection group report errors for that protection group and the virtual machines it contains. Other protection groups covered by those recovery plans are not affected by the errors.

Configuring and Maintaining the Protection of a Virtual Machine

Every virtual machine in a protection group must be configured in such a way that it can be added to vCenter inventory at the recovery site. At a minimum, each machine needs to be assigned to a resource pool, folder, and network that exist at the recovery site. An SRM administrator can specify defaults for these assignments. These defaults, called inventory mappings, are applied when the protection group is created, and can be reapplied as needed (for example, whenever you add a new virtual machine to the protected datastore). If you do not specify inventory mappings, you must configure them individually for each member of the protection group. Virtual machines that are on a protected datastore but improperly configured are not protected.

About Placeholder Virtual Machines and Inventory Mapping

For each virtual machine that you add to a protection group, SRM creates a placeholder at the recovery site. These placeholders are added to, and can be managed as part of, the protected site's inventory.

When you add a virtual machine or template to a protection group, SRM reserves a place for it in the recovery site's inventory by creating a subset of virtual machine files at the recovery site and then using that subset as a placeholder to register the virtual machine with the recovery site vCenter. The presence of these placeholders in recovery site inventory provides a visual indication to SRM administrators that the virtual machines are protected, and to vCenter administrators that the virtual machines can be powered on and start consuming local resources when SRM tests or runs a recovery plan.

No member of a protection group is protected until its placeholder has been created. Placeholders are not created until valid inventory mappings have been established by either applying the site’s inventory mappings to all members of a protection group or configuring mappings for individual members. If inventory mappings have been established for a site, you cannot override them by configuring the protection of individual virtual

machines. If you need to override inventory mappings for a few members of a protection group, use the vSphere Client to connect to the recovery site and edit the network settings of the placeholders or move them to a different folder or resource pool. If a member of a protection group loses its protection, its placeholder is removed from the recovery site until the protection has been restored.

Placeholders can be treated like any other members of the recovery site vCenter inventory, although they cannot be powered on. When a placeholder is created, its folder, network, and compute resource assignments are derived from inventory mappings established at the protected site. Its permissions are inherited from the protected virtual machine that it represents. A recovery site vCenter administrator can modify these assignments and permissions as necessary. Changes made to the placeholder override settings established by inventory mapping, and are preserved in the recovery site SRM database.

When a protected virtual machine is recovered by testing or running a recovery plan, its placeholder is unregistered, and the recovered virtual machine is registered in its place and powered on as directed by the recovery plan. After a recovery plan test completes, the placeholders are restored as part of the cleanup process.

Understanding Recovery and Test Recovery

A test recovery exercises nearly every aspect of a recovery plan, though several concessions are made to avoid disrupting ongoing operations. While a test recovery has no lasting effects on either site, a recovery has significant effects on both sites.

You can (and should) run test recoveries as often as needed. A test recovery does not affect replication or the ongoing operations of either site (though it might temporarily suspend selected local virtual machines at the recovery site). You can pause, resume, or cancel a recovery plan test at any time.

A recovery stops replication (after a final synchronization of the source to the target) and makes changes at both sites that require significant time and effort to reverse. Because of this, the privilege to test a recovery plan and the privilege to run a recovery plan must be separately assigned. A running recovery plan cannot be paused or canceled.

Table 1-1 lists the differences between testing and running a recovery plan.

Table 1-1. Differences Between Testing and Running a Recovery Plan

	Test a Recovery Plan	Run a Recovery Plan
Required privileges	**Site Recovery Manager > Recovery Plan > Test**	**Site Recovery Manager > Recovery Plan > Run**
Effect on virtual machines at protected site	none	Virtual machines are shut down in reverse priority order.
Effect on virtual machines at recovery site	Local virtual machines are suspended if required by the plan. Suspended virtual machines are restarted after the test is complete.	Local virtual machines are suspended if required by the plan.
Effect on replication	Temporary snapshots of replicated storage are created at the recovery site and the arrays are rescanned to discover them.	All replicated datastores are synchronized, then replication is stopped, and the target devices at the recovery site are made writable.
Network	Recovered virtual machines are connected to a test network.	Recovered virtual machines are connected to a datacenter network.
Interruption	Can be paused or canceled.	Must run to completion.

How SRM Interacts with DPM and DRS During Recovery

Distributed Power Management (DPM) is a VMware facility that manages power consumption by ESX hosts. Distributed Resource Scheduler (DRS) is a VMware facility that manages the assignment of virtual machines to ESX hosts. When DPM and DRS are enabled on a recovery site cluster, SRM temporarily disables DPM for the cluster and ensures that all hosts in it are powered on before recovery begins. After the recovery hosts have been powered on, SRM relies on DRS to manage the assignment of virtual machines to hosts in the cluster. After the recovery or test is complete, SRM re-enables DPM for the cluster, but the hosts in it are left in the running state so that DPM can power them down as needed.

Test Bubble Networks and Datacenter Networks

To facilitate testing, SRM can create a "test bubble" network to which recovered virtual machine are connected during a test. This network is managed by its own virtual switch, and in most cases recovered virtual machines can use it without having to change network properties such as IP address, gateway, and so on. A datacenter network, in contrast, is one that typically supports existing virtual machines at the recovery site. To use it, recovered virtual machines must conform to its network address availability rules (they must use a network address that can be served and routed by the network's switch, must be configured to use the correct gateway and DNS host, and so on). Recovered virtual machines that use DHCP can connect to this network without additional customization. Others require IP property customization and recovery plan steps that apply it.

Operational Limits of Site Recovery Manager

Each SRM server can support a maximum number of virtual machines, protection groups, and datastore groups.

Table 1-2 lists the limits for a single SRM server. When you create a protection group, SRM prevents you from exceeding the limits for protected virtual machines and protection groups. If a configuration created in an earlier version of SRM exceeds these limits, SRM displays a warning, but allows the configuration to operate. You should modify these configurations to bring them within the supported limits as soon as possible.

The limits for replicated datastore groups and running recovery plans are suggested and not enforced.

Table 1-2. SRM Protection Limits

Item	Maximum
Protected virtual machines in total	1000
Protected virtual machines in a single protection group	500
Protection groups	150
Datastore groups	150
Simultaneously running recovery plans	3

About Failback

A failback restores the original configuration of the protected and recovery sites after a failover. You can configure and execute a failback procedure when you are ready to restore services to the protected site.

Failback is a catch-all term for a collection of procedures that you can use to restore the original configuration of the protected and recovery sites after a failover. The specific procedures required depend on the nature of the preceding failover: a planned failover that leaves the protected site intact requires a different set of failback steps than an unplanned failover initiated before (or after) an event that compromises the protected site temporarily or permanently.

A typical failback has two phases. In the first phase, the protected and recovery sites switch roles, and the virtual machines are migrated from the recovery site to the protected site under the control of a recovery plan. In the second phase, the relationship of the protected and recovery sites is restored, so that future failovers migrate the protected virtual machines from the protected site to the recovery site. Alternately, the recovery site can be promoted to a protected site, and the protected site becomes the recovery site.

Configuring and executing a failback is a time-consuming task that requires downtime at the recovery site and changes to storage replication. After the failback is complete, restoring the protected site to its original role and enabling failover to the recovery site requires additional downtime and changes to storage replication.

SRM and vCenter

The SRM server operates as an extension to the vCenter server at a site. Because the SRM server depends on vCenter for a number of services, you must install and configure vCenter at a site before you install SRM.

SRM takes advantage of vCenter services, such as storage management, authentication, authorization, and guest customization. SRM also uses the standard set of vSphere administrative tools to manage these services.

How Changes to vCenter Inventory Affect SRM

Because SRM protection groups apply to a subset of vCenter inventory, changes to the protected inventory made by vCenter administrators and users can affect the integrity of SRM protection and recovery.

SRM depends on the availability of certain objects, such as virtual machines, folders, resource pools, and networks, in the vCenter inventory at the protected and recovery sites. Deletion of resources such as folders or networks that are referenced by recovery plans can invalidate the plan. Renaming or relocating objects in the vCenter inventory does not affect SRM, unless it causes resources to become inaccessible during test or recovery.

SRM can tolerate the following changes at the protected site without disruption:

- Modifying a protected virtual machine configuration, including adding, modifying, or removing devices. Changing a virtual machine's memory size on the protected site is not reflected on the recovery site if the virtual machine is already in a protection group.
- Relocating virtual machines.
- Deleting protected virtual machines.
- Deleting an object for which an inventory mapping exists.

SRM can tolerate the following changes at the recovery site without disruption:

- Deleting placeholder virtual machines.
- Moving placeholder virtual machines to a different folder, resource pool, or network.
- Deleting an object for which an inventory mapping exists.

SRM and the vCenter Database

If you update the vCenter installation that SRM extends, you must not overwrite the vCenter database during the update. Overwriting removes information that SRM has stored in that database and invalidates the current SRM installation.

About the Site Recovery Manager Database

The SRM server requires its own database, which it uses to store recovery plans, inventory information, and similar data.

The SRM database is a critical part of any SRM installation. The database must be initialized and a database connection created before you can install SRM. If you are updating SRM to a new release, you can use the existing database connection, but you must back up the database first, otherwise, you will not be able to revert to the previous release of SRM.

The SRM database at each site holds information about virtual machine configurations, protection groups, and recovery plans. SRM cannot use the vCenter database because it has different database schema requirements, though you can use the vCenter database server to create and support the SRM database. Each SRM site requires its own instance of the SRM database. The database must exist before SRM can be installed. If the SRM database at either site becomes corrupted, the SRM servers at both sites shut down.

When you install SRM, you specify the following information about how SRM connects to the database:

Connection Count — The initial connection pool size. If all connections are in use and a new one is needed, a connection is created as long as it does not exceed the maximum number of connections allowed. It is faster for SRM to use a connection from the pool than to create a new one.

Max Connections — The maximum number of connections to open to the database at one time. If the database administrator has restricted the number of connections that the database can have open, this value cannot exceed that number.

SRM Licensing

The SRM server requires a license key to operate. Each SRM server installs with an evaluation license that is valid for 60 days.

After the evaluation license expires, you cannot run a recovery plan or add a virtual machine to a protection group until you install a valid SRM license key. VMware recommends that you install an SRM license key as soon as possible after installing SRM. You can obtain a license key from your VMware sales representative.

SRM Authentication

All communications between SRM and vCenter servers take place over an SSL connection and are authenticated by public key certificates or stored credentials.

When you install an SRM server, you must choose either credential-based authentication or certificate-based authentication. You cannot mix authentication methods. The authentication method you choose when installing the SRM server is used to authenticate connections between the SRM servers at the protected and recovery sites, and between SRM and vCenter.

Certificate-Based Authentication

If you have or can acquire a PKCS#12 certificate signed by a trusted authority, use certificate-based authentication. Public key certificates signed by a trusted authority streamline many SRM operations and provide the highest level of security. Certificates used by SRM have special requirements. See "Requirements When Using Public Key Certificates," on page 15.

Credential-Based Authentication

If you are using credential-based authentication, SRM stores a user name and password that you specify during installation, and then uses those credentials when connecting to vCenter or another SRM server. SRM also creates a special-purpose certificate for its own use. This certificate includes additional information that you supply during installation. That information, an Organization name and Organization Unit name, must be identical for both members of an SRM server pair.

NOTE Even though SRM creates and uses this special-purpose certificate when you choose credential-based authentication, credential-based authentication is not equivalent to certificate-based authentication in either security or operational simplicity.

Certificate Warnings

If you are using credential-based authentication, attempts by the SRM server to connect to vCenter produce a certificate warning because the trust relationship asserted by the special-purpose certificates created by SRM and vCenter cannot be verified by SSL. The warning dialog allows you to specify a disposition for the current instance of the problem, for all instances of the problem when making connection to a specific host, or for all instances of the problem for all hosts. To avoid these warnings, use certificate-based authentication and obtain your certificate from a trusted certificate authority.

Requirements When Using Public Key Certificates

If you have installed SSL certificates issued by a trusted certificate authority (CA) on the vCenter server that supports SRM, the certificates you create for use by SRM must meet certain specific criteria.

While SRM uses standard PKCS#12 certificate for authentication, it places a few specific requirements on the contents of certain field of those certificates. These requirements apply to the certificates used by both members of an SRM server pair (the protected site and the recovery site).

- The certificates must have a Subject Name value constructed from:
  - A Common Name (CN) attribute, whose value must be the same for both members of the pair. A string such as "SRM" is appropriate here.
  - An Organization (O) attribute, whose value must be the same as the value of this attribute in the supporting vCenter server's certificate.
  - An Organizational Unit (OU) attribute, whose value must be the same as the value of this attribute in the supporting vCenter server's certificate.
- The certificate used by each member of an SRM server pair must include a Subject Alternative Name attribute whose value is the fully-qualified domain name of the SRM server host. (This value will be different for each member of the SRM server pair.) If you are using an openssl CA, modify the openssl configuration file to include a line like the following if the SRM server host's fully-qualified domain name is srm1.example.com:

  ```
  subjectAltName = DNS: srm1.example.com
  ```

  If you are using a Microsoft CA, refer to http://support.microsoft.com/kb/931351 for information on how to set the Subject Alternative Name.
- The certificate used by each member of an SRM server pair must include an "Extended Key Usage" attribute whose value is "serverAuth, clientAuth". If you are using an openssl CA, modify the openssl configuration file to include a line like the following:

  ```
  extendedKeyUsage = serverAuth, clientAuth
  ```

How SRM Uses Network Ports

SRM servers use several network ports to communicate with each other, with client plug-ins, and with vCenter. If any of these ports are in use by other applications or are blocked on your network, you must reconfigure SRM to use different ones.

Table 1-3 lists the default network ports the SRM uses for intrasite (between hosts at a single site) and intersite (between hosts at the protected and recovery sites) communications. You can change these defaults when you install SRM.

Table 1-3. How SRM Uses Network Ports

Default Port	Protocol	Endpoints
8095	SOAP	SRM server and vCenter server (intrasite only)
8096	HTTP	vCenter server (for plug-in download)
9007	SOAP	API clients

Site Recovery Manager Roles and Permissions

SRM uses vCenter roles and permissions but includes additional ones that allow fine-grained control over SRM-specific tasks and operations.

SRM and vCenter use the same authorization model. The set of permissions applied to or inherited by an object determine the operations that are allowed on the object and the list of roles that can perform those operations. To manage roles and permissions, you must log in to vCenter as an administrator.

NOTE To configure SRM, you must have vCenter permissions and SRM permissions. SRM-specific roles do not include vCenter privileges. Without vCenter privileges, you do not have adequate permissions to perform all SRM operations. In addition,vCenter roles do not provide any SRM privileges. Ensure that SRM users have vCenter and SRM specific roles as appropriate.

For more information about vCenter roles and permissions, see *Managing Users, Groups, Roles, and Permissions* in the vSphere Client Help.

Site Recovery Manager Roles

SRM adds the following roles to the ones already defined on the vCenter Server:

- Protection Groups Administrator—Set up and modify protection groups.
- Protection SRM Administrator—Pair the protected and recovery sites, and configure inventory mappings.
- Protection Virtual Machine Administrator—Set up and modify the protection characteristics of a protected virtual machine.
- Recovery Datacenter Administrator—View available datastores and customize recovered virtual machines.
- Recovery Host Administrator—Configure virtual machine components during recovery. If the recovery host is a cluster, this permission must be assigned for the cluster object itself and for every host in the cluster.
- Recovery Inventory Administrator—View customization specifications for the recovery site.
- Recovery Plans Administrator—Reconfigure protected and recovered virtual machines. Also grants the ability to set up and run recovery.

- Recovery SRM Administrator—Configure arrays and create protection profiles.
- Recovery Virtual Machine Administrator—Create virtual at the recovery site machines and add them to the resource pool. Also grants the ability to reconfigure and customize the recovery virtual machines when a recovery plan is run.

SRM Administrative Tasks and Required Privileges

Some SRM administrative tasks are performed at the protected site, some at the recovery site, and some at both sites. Specific privileges are required at each site. To perform all SRM administrative tasks at the protected and recovery sites, you must be a vSphere administrator or have specific permissions.

You do not need to propagate the permissions to child objects, unless specified.

You must have the following permissions at the protected site:

- Read-only at the vCenter root.
- Read-only at the datacenter inventory object.
- Protection Virtual Machine Administrator role at the virtual machine level (propagate).
- Protection SRM Administrator role at the SRM site recovery root level (propagate).
- Protection Groups Administrator role at the SRM protection groups level (propagate).

You must have the following permissions at the recovery site:

- Recovery Inventory Administrator role at the vCenter root.
- Recovery Datacenter Administrator role at the datacenter level (propagate).
- Recovery Host Administrator role at the host level (If the recovery host is a cluster, this permission must be assigned for the cluster object itself and for every host in the cluster.)
- Recovery Virtual Machine Administrator at the resource pool and folder levels (propagate).
- Recovery SRM Administrator at the SRM root level (propagate).
- Recovery Plans Administrator at the SRM recovery plans level (propagate).

You can grant selected roles or individuals the minimum set of privileges required to perform specific operations. You can grant these privileges in addition to the default permissions for a role rather than giving many users broad administrative powers.

Table 1-4 summarizes the permissions required for common SRM administrative tasks and the sites at which those permissions must be granted.

Table 1-4. Site Recovery Manager Administrative Tasks and Minimum Required Privileges

Task	Site	Minimum Required Privilege
Add new user and role	Both	**Permissions > Modify Role**
Assign access permission	Both	**Permissions > Modify Permission**
Change access permission	Both	**Permissions > Modify Permission**
Remove access permission	Both	**Administrator**
Connect (pair) sites	Both	**Site Recovery Manager > Protection SRM Administrator**
Modify advanced settings	Protected	**Site Recovery Manager > Protection SRM Administrator**
Modify advanced settings	Recovery	**Site Recovery Manager > Recovery SRM Administrator**
Configure or repair array managers	Both	**Site Recovery Manager > Array Manager > Configure**

Table 1-4. Site Recovery Manager Administrative Tasks and Minimum Required Privileges (Continued)

Task	Site	Minimum Required Privilege
Configure inventory preferences	Both	**Site Recovery Manager > Inventory Preferences > Create Mapping**
Create protection groups	Protected	**Site Recovery Manager > Protection Group > Create**
Create or modify a recovery plan	Recovery	**Site Recovery Manager > Recovery Plan > Create**
Edit a recovery plan	Recovery	**Site Recovery Manager > Recovery Plan > Modify**
Test a recovery plan	Recovery	**Site Recovery Manager > Recovery Plan > Test**
Run or remove a recovery plan	Recovery	**Site Recovery Manager > Recovery Plan > Run**

Installing and Updating Site Recovery Manager 2

You must install an SRM server at the protected site and also at the recovery site. After the SRM servers are installed, you can download the client plug-in from either server to any vSphere Client. You use the SRM client plug-in to configure and manage SRM at each site.

Prerequisites

SRM requires the support of a vCenter server at each site. The SRM installer must be able to connect with this server during installation. If you cannot install SRM on a dedicated server host, you can install it on the same host where vCenter Server is installed.

The SRM server host must meet the following hardware requirements:

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

For up-to-date information about supported platforms and databases, see the *Site Recovery Manager Compatibility Matrixes*, at http://www.vmware.com/support/pubs/srm_pubs.html.

This chapter includes the following topics:

Configuring the SRM Database

The SRM server requires its own database to store recovery plans, inventory information, and similar data. Before installing the SRM server, you must configure and initialize the SRM database.

If you are updating SRM to a new release, you can use the existing database. Back up the database first in case you need to revert after the upgrade.

The SRM database at each site holds information about virtual machine configurations, protection groups, and recovery plans. SRM cannot use the vCenter database because it has different database schema requirements, though you can use the vCenter database server to create and support the SRM database. Each SRM site requires its own instance of the SRM database. The database must exist before SRM can be installed. If the SRM database at either site becomes corrupted, the SRM servers at both sites shut down.

NOTE If you reinitialize the database after you install SRM, you must run the SRM installer in repair mode and specify a new database connection.

Microsoft SQL Server Configuration

A Microsoft SQL Server configuration must meet specific requirements to support SRM.

Microsoft SQL Server has the following configuration requirements when used as the SRM database:

- There are three requirements for the database schema:
  - It must be owned by the SRM database user (the database user name you supply when configuring the SRM database connection).
  - It must have the same name as the SRM database user.
  - It must be the default schema for the SRM database user.
- The SRM database user must have database administrator privileges.
- The SRM database user must be granted the following permissions:
  - bulk insert
  - connect
  - create table
  - create view
- If you are using Windows authentication, the SRM server and database server must run on the same host.
- If the SRM server and database server run on different hosts, you must use mixed mode authentication.
- If SQL Server is installed locally, you might need to disable the Shared Memory network setting on the database server.

Oracle Server Configuration

An Oracle Server configuration must meet specific requirements to support SRM.

Oracle Server has the following configuration requirements when used as the SRM database:

- The SRM database user (the database user name you supply when configuring the SRM database connection) must be granted the following permissions:
  - connect
  - resource
  - create session
  - create view

DB2 Server Configuration

A DB2 Server configuration must meet specific requirements to support SRM.

DB2 Server has the following configuration requirements when used as the SRM database:

- When creating the database instance, specify utf-8 encoding.
- Because DB2 uses Windows authentication, you must specify the database owner as a domain account.

Install the SRM Server

You must install an SRM server at the protected site and the recovery site as an extension to the site's vCenter Server.

Prerequisites

You must supply the following information during the installation:

- The hostname or IP address of the site's vCenter Server. The server must be running and accessible during SRM installation, and it must be in the same Windows domain as the SRM server host.
- The user name and password of the vCenter administrator.
- A user name and password for the SRM database. See "Configuring the SRM Database," on page 19.
- If you are using certificate-based authentication, the pathname to an appropriate certificate file. See "SRM Authentication," on page 14.

Procedure

1. Log in to the server host on which you are installing SRM.

   Log in as a local administrator.

2. Download the SRM installation file to a folder on the host, or open a folder on the network that contains this file.

3. Double-click the SRM installer icon to begin installation.

   If the installer detects an existing installation, verify that you want to update the existing installation, and then follow the procedure at "Update the SRM Server," on page 23.

4. Click **Next** on the Welcome to the installation wizard screen.

5. On the License Agreement page, select **I accept the terms in the license agreement** and then click **Next**.

6. On the Destination Folder page, select the folder in which you want to install SRM and click **Next.**

   The default installation folder for a new installation of SRM is `C:\Program Files\VMware\VMware vCenter Site Recovery Manager`. If you use a different folder, the pathname cannot be longer than 240 characters and cannot include non-ASCII characters.

7. On the VMware vCenter Server page, enter information about the vCenter server at the site where you are installing SRM and then click **Next**.

   - **vCenter Server Address**—Enter the hostname or IP address of the vCenter Server. If you use the hostname, enter it in lowercase. After installation is complete and you are configuring the connection between the protected and recovery sites, you must supply this hostname or IP address exactly as you enter it here.
   - **vCenter Server Port**—Accept the default or enter a different port.

- **vCenter Server Username**—Enter the user name of an administrator of the specified vCenter server.
- **vCenter Server Password** —Enter the password for the specified user name.

When you click **Next**, the installer contacts the specified vCenter server and validates the information you supplied.

8 On the Certificate Type Selection page, select an authentication method.

- To use credential-based authentication, select **Automatically generate certificate** and click **Next**. Enter text values for your organization and organization unit, typically your company name and the name of your group within the company.
- To use certificate-based authentication, select **Use a PKCS #12 certificate file** and click **Next**. Enter the path to the certificate file. The certificate file must contain exactly one certificate with exactly one private key matching the certificate. Enter the certificate password if necessary.

See "SRM Authentication," on page 14.

9 Enter the following additional information:

- **Local Site Name**—A name for this installation of SRM. A suggested name is generated for you, but you can specify any name you want, so long as it is not the same name that you use for another SRM installation with which this one will be paired.
- **Administrator E-mail**—The email address to which SRM administrative alerts and notifications are sent.
- **Additional E-mail**—An optional additional email address to which SRM administrative alerts and notifications are sent.
- **Local Host**—The name or IP address of the local host. This value is obtained by the SRM installer and need only be changed if it is incorrect (for example, if the local host has more than one network interface and the one detected by the SRM installer is not the one you want to use).
- **Listener Ports**—The SOAP and HTTP port numbers to use.
- **API Listener Port**—The SOAP port number for API clients to use.

The SRM installer supplies default values for these ports. Do not change them unless the defaults would cause port conflicts. See "How SRM Uses Network Ports," on page 16.

10 Enter the database configuration information and click **Next**.

- **Database Client**— Select a database client type from the pulldown control.
- **Data Source Name**— Select and existing DSN from the pulldown, or click **ODBC DSN Setup** to view existing DSNs or create a new one.
- **Username**—A user ID valid for the specified database.
- **Password**—The password for the specified user ID.
- **Connection Count**—The initial connection pool size.
- **Max Connections**—The maximum number of database connections that can be open simultaneously.

For more information about any of these values, see "About the Site Recovery Manager Database," on page 14.

NOTE If a database exists at the DSN that you provide, you are prompted to either use it or overwrite it.

---

11 Click **Install.**

12 When the wizard completes, click **Finish.**

What to do next

You can now install SRAs at each site. See "Install the Storage Replication Adapters," on page 23.

Install the Storage Replication Adapters

A Storage Replication Adapter (SRA) is a program provided by an array vendor that enables SRM to work with a specific kind of array. You must install an appropriate SRM on the SRM server hosts at the protected and recovery sites.

Prerequisites

- SRM server installation creates a directory in which you can install the SRAs. Install the SRM server before you install the SRAs.
- Your SRA might require the installation of other vendor-provided components. Some of these components might need to be installed on the SRM server host; others might require only network access by the SRM server.
- SRM might occasionally need to rescan storage arrays. You can improve array rescan times by changing default value of the `Scsi.RescanAllHbas` on ESX hosts. If rescan times on ESX hosts are longer than 10 minutes, you may want to set the value of this option to 1.
- Masking and zoning must be configured for replicated devices to remote ESX hosts for failover. VMware recommends that you configure storage to create clones or snapshots of the replicated devices. Snapshots or clones must be masked to the recovery site ESX hosts.

Procedure

1 Download the SRA.

  You can download storage replication adapters and their documentation from http://www.vmware.com/download/srm/. Storage replication adapters downloaded from other sites are not supported by VMware.

2 Install the SRA on each SRM server host.

  Storage replication adapters come with their own installation instructions. The adapter you are using must be installed on the SRM server host at the protected and recovery sites. You cannot use different adapters, or different versions of the same adapter, at these sites. Both members of an SRM site pair must use identical adapters.

3 Re-start the SRM service.

  The SRM service looks for SRAs when it starts up. If you add or change SRAs on a host, you must re-start the SRM server process on that host.

Update the SRM Server

When you update the Site Recovery Manager server, information about vCenter server connections, certificates, and database configuration is read from the existing installation and reused by the updated installation.

The update mode of the SRM installer provides a quick way to update the SRM server to a new release without changing any of the information that you provided for the current installation. If you need to change any of that information, including database connections, authentication method, certificate location, or administrator credentials, you must follow the update with a repair mode installation, or uninstall the existing release and then install the new release.

Prerequisites

Before you begin the update, back up your current SRM database. The update wizard requires you to verify that the database is backed up, and pauses until you confirm that it is.

Procedure

1 Log in to the server host on which you are installing SRM.

Log in as a local administrator.

2 Download the SRM installation file to a folder on the host, or open a folder on the network that contains this file.

3 Double-click the SRM installer icon to begin installation.

When prompted to verify that you want to update the existing installation, click **Yes**.

4 Click **Next** on the Welcome to the installation wizard page.

The wizard prompts you to verify that you have backed up the SRM database.

5 Click **Yes** to confirm that you have backed up the database and are ready to proceed with the installation.

The installer reads configuration data from the existing installation and uses it to complete the update. The update installs the same location as the previous installation. If any of the existing configuration information is invalid for the upgrade (for example, if the database is not accessible at the same DSN, or the vCenter server is not accessible at the same port), the update fails.

6 When the wizard completes, click **Finish.**

If the update replaced any open files, you are prompted to shut down and restart Windows.

What to do next

You can now install the updated client plug-in. See "Install the SRM Client Plug-In," on page 24.

Install the SRM Client Plug-In

To install the Site Recovery Manager client plug-in, use a vSphere Client to connect to the vCenter Server at the protected or recovery site, then download the plug-in from the server and enable it in the vSphere Client.

When you install the Site Recovery Manager server, The Site Recovery Manager client plug-in becomes available as a download from the vCenter server that the Site Recovery Manager server installation extends. You can download, install, and enable the SRM client plug-in on any host where a vSphere Client is installed.

Prerequisites

The Site Recovery Manager server must be installed at the protected and recovery sites.

Procedure

1 Start the vSphere Client and connect to vCenter Server at the protected or recovery site.

2 On the vSphere Client menu bar, click **Plugins > Manage Plugins**.

3 In the Available Plug-ins area of the Plugin Manager window, right-click the VMware vSphere Site Recovery plug-in and click **Download and Install**.

4 After the download completes, click **Next** on the plug-in installation wizard Welcome page to start the wizard.

5 Click **I accept the terms in the license agreement**, and click **Next.**

6 Click **Install**.

7 When the installation completes, click **Finish**.

If the installation replaced any open files, you are prompted to shut down and restart Windows.

Revert to a Previous Release

To revert to a previous release, uninstall the current SRM server release from the protected and recovery sites, uninstall the SRM plug-in, and restore the SRM database from the backup you made before you updated the SRM server. You can then install the previous release and use the restored database.

Prerequisites

Before you revert an SRM installation to a previous release, be sure that the current installation of vCenter supports that release. For information about vCenter releases that support SRM, see the *Site Recovery Manager Compatibility Matrixes*, accessible from http://www.vmware.com/support/pubs/srm_pubs.html. For information about reverting a vCenter installation, see the vSphere documentation.

Procedure

1 Uninstall SRM at the protected and recovery sites.

Where sites have been paired, SRM at both sites must be uninstalled. If you uninstall SRM from one member of a site pair, the database of the remaining member becomes inconsistent.

2 Uninstall the SRM plug-in from any vCenter Clients where it has been installed.

3 Restore the database used by the previous release, following the procedures documented by your database vendor.

4 Install the previous release of SRM.

Repair a Site Recovery Manager Server Installation

If you need to change any of the information you supplied when you installed the SRM Server, you can repair the installation and supply the changed information.

Installing the SRM server binds the installation to a number of values that you supply, including the vCenter server to extend, the SRM database DSN and credentials, the type of authentication you want to use, and so on. The SRM installer supports a repair mode that allows you to change any of the following values for the installation:

- The hostname or IP address of the site's vCenter Server.
- The username and password of the vCenter administrator.
- The username, password, and DSN for the SRM database.
- The type of authentication (certificate-based or credential-based), the authentication details, or both.

The installer's repair mode presents modified versions of most of the pages that are part of the SRM server installation. For more information about any of the repair options, see "Install the SRM Server," on page 21.

Procedure

1 Log in to the SRM server host.

Log in as a local administrator.

2 Open the Windows Add or Remove Software tool. Navigate to the entry for VMware vCenter Site Recovery Manager and click **Change** to start the installer in repair mode.

3 Click **Next** on the Welcome to the installation wizard screen.

4 Click **Repair** on the Program Maintenance Options page.

5 On the VMware vCenter Server page, enter the following information:

- **vCenter Server Username**—Enter the user name of an administrator of the specified vCenter server.
- **vCenter Server Password** —Enter the password for the specified user name.

You cannot use the installer's repair mode to change the vCenter server address or port. When you click **Next**, the installer contacts the specified vCenter server and validates the information you supplied.

6 On the Certificate Type Selection page, choose an authentication method and click **Next**.

- To leave the current authentication method unchanged, select **Use existing certificate**. If the installed certificate is not valid, this option is unavailable.
- To choose credential-based authentication, select **Automatically generate certificate**.
- To choose certificate-based authentication, select **Use a PKCS #12 certificate file**.

Unless you select **Use existing certificate**, you will be prompted to supply additional authentication details such as certificate location or strings to use for Organization and Organizational Unit. For more information, see "SRM Authentication," on page 14.

7 On the Database Configuration page, Enter the following database configuration information and click **Next**:

- **Data Source Name**— Select and existing DSN from the pulldown, or click **ODBC DSN Setup** to view existing DSNs or create a new one.
- **Username**—A user ID valid for the specified database.
- **Password**—The password for the specified user ID.
- **Connection Count**—The initial connection pool size.
- **Max Connections**—The maximum number of database connection open simultaneously.

For more information about any of these values, see "About the Site Recovery Manager Database," on page 14.You cannot use the installer's repair mode to change the database type. If the installer detects an existing database at the DSN you provide, it prompts you to either use it (preserving its contents) or overwrite it (destroying its contents).

8 On the Ready to Repair the Program page, click **Install** to repair the installation.

The installer makes the requested repairs and restarts the SRM server.

Configuring the Protected and Recovery Sites

3

After you have installed SRM at the protected and recovery sites, you must connect the two sites to create a site pair, configure the array managers at each site, and configure SRM at each site. You use the SRM client plug-in to administer SRM. Site pairing requires vSphere admnistrative privileges at both sites.

Prerequisites

Before you can connect the protected and recovery sites, you must:

1 Install an SRM server at each site.

2 Install the appropriate storage replication adapters on the SRM server hosts at both sites. The recovery site must be the replication target of arrays managed by the SRA at the protected site.

3 Download the SRM plug-in from an SRM server into the vSphere client that you want to use to administer SRM.

This chapter includes the following topics:

Create a Site Pair

Before you can use SRM, you must designate the protected site and the recovery site and connect them. The sites must authenticate with each other.

Prerequisites

Before you can connect the protected and recovery sites, you must install an SRM server at each site, then install and enable the SRM plug-in at a vSphere client from which you want to administer SRM.

NOTE If you are using credential-based authentication, several of the steps in this procedure produce certificate warnings. For more information about these warnings and options for dealing with them, see "SRM Authentication," on page 14.

Procedure

1 Open a vSphere client and connect to the vCenter server at the site that you want to designate as the protected site.

Log in as a vSphere administrator.

NOTE The recovery site must be the replication target of arrays managed by the SRA at the protected site.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Protection Setup area of the Summary window, navigate to the Connection line and click **Configure**.

4 On the Remote Site Information page, type the IP address or host name of the vCenter server at the recovery site and click **Next**.

NOTE If you are using credential-based authentication, you must enter exactly the same information here that you entered when installing the SRM server. If you entered an IP address in that step, enter it again here. If you entered a hostname in that step, enter it here in exactly the same way.

Port 80 is used for the initial connection to the remote site. After the initial HTTP connection is made, the two sites establish an SSL connection over port 443 for subsequent connections.

5 On the vCenter Server Authentication page, provide the vCenter administrator user name and password for the remote site and click **Next**.

If you are using credential-based authentication, you must enter exactly the same information here that you entered when installing the SRM server.

6 On the Compete Connections page, click **Finish** after all of the site paring steps have completed successfully.

The SRM and vCenter servers at the protected and recovery sites are connected. Connection information is saved in the SRM databases, and persists across logins and host restarts.

What to do next

After the sites are connected, you can configure the array managers.

Disconnect From a Protected or Recovery Site

You can use the Site Recovery Manager's Logout link to disconnect a vSphere Client from SRM so that you can reconnect as a different user.

In the **Protection Setup** area of the **Summary** window, navigate to the **Connection** line and click **Logout**.

The vSphere is disconnected from SRM, and a **Connect to VMware vCenter Site Recovery Manager** button is displayed. To reconnect, click the button and supply the requested credentials.

Install the SRM License Key

The SRM server requires a license key to operate. VMware recommends that you install an SRM license key as soon as possible after installing SRM.

Procedure

1 Open a vSphere client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 Right-click **Site Recovery** in the vSphere Client navigation pane and click **Advanced Settings**.

4 In the navigation pane of the Advanced Settings window, click **Licensing**.

5 Enter the SRM license key in the **Licensing.LicenseKey** text box

The first time you open the Licensing page, the evaluation key is displayed in the **Licensing.LicenseKey** text box.

6 Click **OK** to save your changes and close the Advanced Settings window.

7 Repeat the process to install a license key at the recovery site.

If you enter a valid license key, it is displayed in the **Licensing.LicenseKey** text box each time you open the Licensing page. If you enter an invalid license key, the previous valid license key (or the evaluation key, if no other valid key has ever been entered) is displayed the next time you open the Licensing page.

Configure Array Managers

After you have connected the protected site and recovery site, you must configure their respective array managers so that SRM can discover replicated devices, compute datastore groups, and initiate storage operations.

The array manager configuration wizard leads you through a number of steps:

- You provide SRM with connection information and credentials (if needed) for array management systems at the protected and recovery sites.
- SRM verifies that it can connect to arrays at both sites.
- SRM verifies that it can discover replicated storage devices on these arrays and identify the VMFS datastores that they support.
- SRM computes and verifies datastore groups based on virtual machine storage layout and any consistency groups defined by the storage array.

When the configuration process is complete, the wizard presents a list of replicated datastore groups. You typically configure array managers only once, after you have connected the protected and recovery sites. You do not need to reconfigure them unless array manager connection information or credentials have changed, or you want to use a different set of arrays.

Prerequisites

Before you configure the array managers at the protected and recovery sites, be sure that at least one virtual machine at the protected site is stored on a replicated device supported by an array for which you have installed an SRA. The array manager configuration wizard does not detect replicated devices unless they are part of a datastore that is home to at least one virtual machine.

You must also connect the protected and recovery sites (see "Create a Site Pair," on page 27).

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Protection Setup area of the Summary window, navigate to the Array Managers line and click **Configure**.

4 On the Protected Site Array Managers page of the Configure Array Managers wizard, click **Add**.

5 Make sure that the array manager type that you want SRM to use appears in the **Manager Type** field.

If more than one SRA has been installed on the SRM server host, click the drop-down arrow and select the manager type you want to use. If no manager type is displayed, no SRA has been installed on the SRM server host. For more information, see "Install the Storage Replication Adapters," on page 23.

6 Type a name for the array in the **Display Name** field of the Add Array Manager window.

Use any descriptive name that makes it easy for you to identify the storage associated with this array manager.

7 Fill in the remaining fields of the Add Array Manager window.

These fields are created by the SRA. For more information about how to fill them in, see the documentation provided by your SRA vendor.

8 Click **Connect** to validate the information you supplied and return the list of arrays that the selected array manager has discovered.

All discovered arrays are selected. Clear the selection of any array that you do not want SRM to use.

9 Click **OK**.

The array manager queries the selected arrays to discover which of their devices are replicated. Detailed information about the selected arrays and the number of replicated devices they support appears in the Replicated Array Pairs area of the Configure Array Managers window.

10 Click **Next** to configure array managers at the recovery site.

11 On the Recovery Site Array Managers page of the Configure Array Managers wizard, click **Add**.

The procedure for configuring these arrays is identical to the procedure for configuring the arrays at the protected site, described in steps Step 5 through Step 8.

12 Click **OK**.

The array manager at the recovery site queries the selected arrays to discover which of their devices are replicated, and displays detailed information about the selected arrays and the number of replicated devices they support in the **Replicated Array Pairs** area of the Configure Array Managers window. A green checkmark icon distinguishes arrays that have peers at the protected site.

13 Click **Next** to display the list of replicated datastore groups.

On the Review Replicated Datastores page, you can expand each datastore group to see the datastores it contains and the devices that they use. If the list of datastore groups is not what you expected, correct it before continuing.

NOTE Only those datastores used by at least one virtual machine are displayed. If no datastores are displayed, verify that the inventory of this vCenter includes at least one virtual machine that uses a datastore supported by the paired arrays.

---

14 Click **Finish** to complete the configuration of the array managers.

Configure Recovery Site Array Managers When the Protected Site Is Inaccessible

If you need to edit array manager details when the protected site is not accessible, use the Repair Array Managers function.

Normally, configuration of array managers requires access to both the protected and recovery sites. SRM provides a Repair Array Managers function that allows you to modify the recovery site array manager configuration even though the protected site is inaccessible. If the protected site is accessible, you can accomplish the same thing by following the procedures in "Configure Array Managers," on page 29.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Recovery Setup area of the Summary window, navigate to the Recovery Plans line and click **Repair Array Managers**.

4 On the Recovery Site Array Managers page, click the **Add**, **Remove**, or **Edit** button to change the array manager information for the recovery site.

Rescan Arrays to Detect Configuration Changes

SRM checks arrays for changes to device configurations every 24 hours. However, if needed, you can force an array rescan at any time.

Configuring array managers causes SRM to compute datastore groups based on the set of replicated storage devices it discovers. If you change the configuration of the array at either site to add or remove devices, SRM must rescan the arrays and recompute the datastore groups.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Protection Setup area of the Summary window, navigate to the Array Managers line and click **Configure**.

4 On the Protected Site Array Managers page of the Configure Array Managers wizard, click **Next**.

5 On the Recovery Site Array Managers page, click **Next**.

6 On the Review Replicated Datastores page, click **Rescan Arrays**.

7 Click **Finish** to complete the operation.

Configure Inventory Mappings

Inventory mappings establish recovery site defaults for the folders, networks, and resource pools to which recovered virtual machines are assigned. You create these mappings at the protected site, and they apply to all virtual machines in all protection groups at that site.

Inventory mappings are optional but recommended. They provide a convenient way to specify how resources at the protected site are mapped to resources at the recovery site. These mappings are applied to all members of a protection group when the group is created, and can be reapplied as needed (for example, when new members are added). If you do not create mappings, you must specify them individually for each virtual machine that you add to a protection group. A virtual machine cannot be protected unless it has valid inventory mappings for networks, folders, and compute resources. You do not need to specify inventory mappings for resources that are not used by protected virtual machines.

NOTE If inventory mappings have been established for a site, you cannot override them by configuring the protection of individual virtual machines. If you need to override inventory mappings for a few members of a protection group, use the vSphere Client to connect to the recovery site and edit the network settings of the placeholders or move them to a different folder or resource pool.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Protection Setup area of the Summary window, navigate to the Inventory Mappings line and click **Configure**.

   The Inventory Mappings page displays a tree of resources at the protected site and a corresponding tree of resources at the recovery site. For any protected site resource that does not have an inventory mapping, the corresponding item in the recovery site tree is listed as **None Selected**.

4 To configure mapping for a resource, right-click it in the Protected Site Resources column and click **Configure**.

5 Expand the top-level folder in the Configure Inventory Mapping window and navigate to the recovery site resource (network, folder, or resource pool) to which you want to map the protected site resource.

6 Select the resource and click **OK**.

   The selected resource is displayed in the Recovery Site Resources column, and its path relative to the root of the recovery site vCenter is displayed in the Recovery Site Path column.

7 To undo an inventory mapping, right-click it and click **Remove**.

What to do next

Create one or more protection groups. Inventory mappings are applied whenever a new protection group is created. New or changed mappings must be manually applied to existing protection groups.

Apply Inventory Mappings to All Members of a Protection Group

When you create a protection group, your inventory mappings are applied to all the virtual machines in it. If you change the mappings, add virtual machines to the protected datastore, or if the virtual machines lose their protection for any reason, you can reapply the mappings to all unconfigured virtual machines in one step.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

   Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 Select a protection group from the list, and click the **Virtual Machines** tab.

4 On the Virtual Machines page, click **Configure All**.

This procedure applies existing inventory mappings to all virtual machines that have a status of Not Configured.

What to do next

After this process completes, virtual machines that could not be configured have a status of Mapping Missing or Mapping Invalid. You must configure protection for these machines individually.

Configure Resource Mappings for a Virtual Machine

If you have not specified inventory mappings for your site, you must configure resource mappings for individual virtual machines. You can configure resource mappings only if site-wide inventory mappings have not been established.

If inventory mappings have been established for a site, you cannot override them by configuring the protection of individual virtual machines. If you need to override inventory mappings for a few members of a protection group, use the vSphere Client to connect to the recovery site and edit the network settings of the placeholders or move them to a different folder or resource pool.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, navigate to the protection group that includes the virtual machine that you want to configure.

4 On the Virtual Machines page, right-click a virtual machine and click **Configure Protection**.

If you have established inventory mappings, they are applied.

5 In the Edit Virtual Machine Properties window, configure mappings as needed.

For most virtual machines, the only required mappings are Folder, Compute Resource, and Network. You can also change resource mappings and other virtual machine properties. See "Configure Protection for a Virtual Machine or Template," on page 54.

Create Protection Groups

SRM organizes virtual machines into protection groups based on the datastore group that they use. All virtual machines in a protection group store their files on the same replicated datastore, and all failover together.

To create a protection group, you select a replicated datastore group to protect, and then specify a nonreplicated datastore at the recovery site where SRM can create placeholders for members of the protection group.

Prerequisites

Before you can create a protection group, you must connect the protected site and recovery site and then configure the array managers. To be protected, a virtual machine must have folder, network connection, and resource pool assignments that are valid at the recovery site. Unless you intend to configure these mappings individually for each member of the group, configure inventory mappings before you create protection groups.

NOTE You can include virtual machine templates in a protection group. Inventory mappings are applied to these templates, and they can be customized as needed. Protected templates, like protected virtual machines, appear as placeholders at the recovery site. If you convert a protected template to a virtual machine or convert a protected virtual machine to a template, the converted object loses its protection and must be reconfigured.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Protection Setup area of the Summary window, navigate to the Protection Groups line and click **Create**.

4 On the Name and Description page of the Create Protection Group wizard, type a name and optional description for the protection group, and click **Next**.

5 On the Select a Datastore Group page, select a datastore group from the list, and click **Next**.

The datastores listed were discovered when you configured the array managers. Each datastore in the list is replicated to the recovery site and supports at least one virtual machine at the protected site. When you select a datastore, the virtual machines that it supports are listed in the **VMs on the selected datastore group** field, and are automatically included in the protection group.

6 On the Datastore for Placeholder VMs page, select a datastore group from the list.

The datastores listed on this page exist only at the recovery site. None of them are replicated from the protected site. The datastore that you select is used to hold the files that constitute the placeholder virtual machines. These files are not large, so any datastore that is accessible to the recovery site host and cluster can be an appropriate choice.

7 Click **Finish** to create the protection group.

SRM creates a protection group that includes all of the virtual machines on the datastore you selected in Step 5. Placeholders are created and inventory mappings applied for each member of the group. If a group member cannot be mapped to a folder, network, and resource pool on the recovery site, it is listed with a status of Mapping Missing, and a placeholder cannot be created for it.

Edit a Protection Group

You can change the name of a protection group and its default recovery site datastore.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, navigate to the protection group that you want to edit.

4 Right-click the group and click **Edit.**

NOTE If you change the datastore for Placeholder VMs, that change applies only to new members added to the protection group. To change the datastore used by existing placeholders, connect to the recovery site with a vSphere Client and use the vSphere Datastores page to migrate the placeholder to a new datastore.

Adding and Removing Members of a Protection Group

When you create a protection group, it includes all the virtual machines on the selected datastore. You can add or remove protection group members by adding or moving virtual machines to the datastore, or by removing them from the datastore.

All virtual machines and templates that reside on a protected datastore are part of the protection group that applies to that datastore. There is no explicit add or remove operation to change group membership. The contents of the datastore implicitly specify the membership of the protection group.

- To add a new virtual machine or template to a protection group, create it on the protected datastore and then configure protection for it.
- To add an existing virtual machine to a protection group, use Storage VMotion to move it to the protected datastore and then configure protection for it.
- To remove a virtual machine or template from a protection group, remove it from the protected datastore.

NOTE When you add a virtual machine or template to a protected datastore, it has an initial status of **Not Configured** in the protection group. You must configure protection for the new group member by applying inventory mappings if they exist, or by configuring resource mappings for it individually.

Limitations on Recovery of Snapshots and Linked Clones

Array-based replication supports recovering VMware Virtual Consolidated Backup (VCB) snapshots, but it does not support recovering other types of snapshots or virtual machines configured as linked clones.

SRM cannot reliably recover virtual machine snapshots that are not created by VCB. A protection group can include virtual machines that have snapshots, but those virtual machines are not usable when recovered.

Virtual machines configured as linked clones are also not protected. You can include such virtual machines in protection groups, but only the parent is completely protected. The linked clones are not usable after recovery.

NOTE If you need to support the use of certain types of VCB snapshots at the recovery site (snapshots taken when the virtual machine is powered on or suspended), the ESX hosts at both sites must have compatible CPUs, as defined in the VMware knowledge base articles *VMotion CPU Compatibility Requirements for Intel Processors* (article 1991) and *VMotion CPU Compatibility Requirements for AMD Processors* (article 1992). The hosts must also have the same BIOS features enabled. If the servers' BIOS configurations do not match, they still show a compatibility error message even if they are otherwise identical. The two most common features to check are Non-Execute Memory Protection (NX / XD) and Virtualization Technology (VT / AMD-V).

Create a Recovery Plan

A recovery plan controls how virtual machines in a protection group are recovered. A basic recovery plan includes a number of prescribed steps that use default values to control how protection group members are migrated to the protected site. You can customize the plan to meet your needs. The plan is stored in the SRM database at the recovery site and executed by the SRM server at that site.

A simple recovery plan assigns all virtual machines in a protection group to two networks on the recovery site: a recovery network and a test network. The recovery network is used in an actual recovery. The test network is used only for testing the recovery plan and does not typically allow the recovered virtual machines to communicate on your corporate network or the Internet. SRM can create a test network for you, or you can create one yourself.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

   Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Recovery Setup area of the Summary window, navigate to the Recovery Plans line and click **Create**.

4 On the Recovery Plan Information page of the Create Recovery Plan wizard, type a name for the plan in the **Name** text box and add an optional description, and then click **Next**.

5 On the Protection Groups page, select one or more protection groups for the plan to recover, and click **Next**.

6 On the Response Times page, specify how long you want the recovery plan to wait for a response from a virtual machine after various recovery plan events, and then click **Next**.

Change Network Settings	If the virtual machine does not acquire the expected IP address within the specified interval after a recovery step that changes network settings, an error is reported and the recovery plan proceeds to the next virtual machine.
Wait for OS Heartbeat	If the virtual machine does not report an OS heartbeat within the specified interval after being powered on, an error is reported and the recovery plan proceeds to the next virtual machine.

NOTE Responses cannot be detected on virtual machines that do not have VMware Tools installed.

7 On the Configure Test Networks page, select a recovery site network to which recovered virtual machines connect during recovery plan tests, and then click **Next**.

By default, the test network is specified as Auto, which creates an isolated test network. If you would prefer to specify an existing recovery site network as the test network, click **Auto** and select the network from the drop-down menu.

8 On the Suspend Local Virtual Machines page, select the virtual machines at the recovery site that the recovery plan should suspend.

Suspending local virtual machines frees resources for use by recovered virtual machines. The virtual machines are suspended during a test recovery as well as during an actual recovery. After a test recovery, they are powered on again.

9 Click **Finish** to create the recovery plan.

Edit a Recovery Plan

You can change the properties of a recovery plan.

Edit a recovery plan if you want to change any of the properties that you specified when you created it. You can also customize a recovery plan by adding or changing recovery steps. For more information, see "Customizing a Recovery Plan," on page 46.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Recovery Setup area of the Summary window, navigate to the Recovery Plans line, right-click the plan that you want to edit, and select **Edit Recovery Plan**.

What to do next

After you have opened the plan for editing, you can change any of its properties. For more information, see "Create a Recovery Plan," on page 35.

Remove a Recovery Plan

You can remove a recovery plan if you no longer need it.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Recovery Setup area of the Summary window, navigate to the Recovery Plans line, right-click the plan that you want to remove, and select **Remove Recovery Plan**.

Test Recovery, Recovery, and Failback

4

After you have configured SRM at the protected and recovery sites, you can test your recovery plan without affecting services at either site. You can also run a recovery plan and, if necessary, configure the two sites for failback so that you can restore services at the protected site.

SRM makes it easy to test a recovery plan. The test does not disrupt replication or any ongoing activities at the protected site. Recovery plans that suspend local virtual machines do so for tests as well as for actual recoveries. With this exception, recovery plan tests do not disrupt activities at either site.

NOTE Permission to test a recovery plan does not include permission to run a recovery plan. Permission to run a recovery plan does not include permission to test a recovery plan. Each permission must be assigned separately.

This chapter includes the following topics:

Test a Recovery Plan

When you test a recovery plan, you use a test network and a temporary copy of replicated data at the recovery site. No operations are disrupted at the protected site.

Testing a recovery plan runs all the steps in the plan with the exception of powering down virtual machines at the protected site and forcing devices at the recovery site to assume mastership of replicated data. If the plan requests suspension of local virtual machines at the recovery site, they are suspended during the test recovery. A test recovery makes no other changes to the production environment at either site.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

Log in as a user who has permission to test a recovery plan.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, expand the **Recovery Plans** icon and click the recovery plan that you want to test.

4 In the Commands area of the Summary window, click **Test Recovery Plan**.

At the confirmation prompt, click **Yes** to proceed with the test.

5 Click the **Recovery Steps** tab to monitor the progress of the test and respond to messages.

The **Recovery Steps** tab displays the progress of individual steps. The Recent Tasks area reports the progress of the overall plan.

NOTE If the SRM server loses contact with the recovery site vCenter while a recovery plan is being tested or run, the recovery plan fails and displays the message `Error: The session is not authenticated`. If this happens during a test, cancel the test. If this happens during a recovery, manual cleanup will probably be required after the plan completes.

6 To clean up and finish the test, click **Continue** at the prompt.

SRM powers down and unregisters the protected virtual machines and then registers the placeholders again.

Pause, Resume, or Cancel a Test

You can pause, resume, or cancel a recovery plan test at any time.

When you pause or cancel a test, no new steps are started, and in-progress steps are stopped subject to the following rules

- Steps that cannot be stopped, such as powering on or waiting for a heartbeat, run to completion before the pause or cancellation completes.
- Steps that add or remove storage devices are undone by cleanup operations if you cancel or by subsequent steps if you pause and resume.

The time it takes to pause or cancel a test depends on the type and number of steps that are currently in progress. The time it takes to resume a test depends on the type and number of steps that were in progress when the pause was requested.

To pause, resume, or cancel a test, click the **Pause**, **Resume**, or **Stop** button on the recovery plan toolbar.

Run a Recovery Plan

When you run a recovery plan, all virtual machines in the plan are migrated to the recovery site and the protected site is shut down.

CAUTION Do not run a recovery plan unless you are prepared to support all recovered virtual machines and services at the recovery site for an indefinite length of time. A recovery plan makes significant alterations in the configurations of the protected and recovery sites. It also stops replication of all devices that support the protected datastores. Reversing these changes (using a failback procedure) takes significant time and effort and is likely to result in prolonged service downtime. Do not run any recovery plan that has not been thoroughly tested.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

Log in as a user who has permission to run a recovery plan.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, expand the **Recovery Plans** icon and click the recovery plan that you want to run.

4 In the Commands area of the Summary window, click **Run Recovery Plan**.

5 Review the information in the confirmation prompt, and when you are ready to proceed, select **I understand that this process cannot be undone** and click **Run Recovery Plan**.

6 To monitor the progress of the recovery and respond to messages, click the **Recovery Steps** tab.

The **Recovery Steps** tab displays the progress of individual steps. The Recent Tasks area reports the progress of the overall plan.

NOTE If the SRM server loses contact with the recovery site vCenter while a recovery plan is being tested or run, the recovery plan fails and displays the message `Error: The session is not authenticated`. If this happens during a test, cancel the test. If this happens during a recovery, manual cleanup will probably be required after the plan completes.

Configuring and Executing Failback

After a recovery plan has been run and the virtual machines in it are operating at the recovery site, you can configure and run a failback procedure, which migrates those virtual machines back to the protected site and prepares both sites for the next recovery or test.

NOTE Not all arrays support the operations required by failback. Before attempting a failback, see the documentation that accompanied your storage replication adapter.

After a failover has completed, there are significant changes at the protected and recovery sites:

- Array replication from the protected site to the recovery site has stopped. Devices at the recovery site are not configured as replication sources or targets.
- If the protected site is still operational, all protected virtual machines affected by the failover have been powered down.
- At the recovery site, all placeholder virtual machines have been replaced by powered-on virtual machines in the recovery site's vCenter inventory.

The virtual machines and the services that they provide are now accessible at the recovery site, but the recovery site itself is no longer protected. To protect the site, you must reconfigure SRM to designate a new recovery site and create the protection groups and recovery plans that are needed to facilitate recovery. If you intend to restore virtual machines and services to the original protected site, you must first configure it to be a recovery site. You then run a failback recovery plan that migrates the protected inventory from the original recovery site back to the original protected site. You can then reconfigure the two sites to resume their original roles.

If you cannot, or do not want to, restore the original protected site to its former status, you can establish a new recovery site. To do so, create the protection groups and recovery plans needed to protect the original recovery site, and then promote the old recovery site to a protected site.

Procedure

1 Review and Execute Post-Failover Cleanup Tasks on page 42

Before you can execute a failback, you must remove artifacts such as invalid protection groups and unneeded placeholders that are left over from the previous configuration.

2 Reconfigure Replication on page 42

Failover stops replication. Failback requires you to configure replication in reverse, from the recovery site to the protected site. Restoring the protected and recovery sites to their original roles requires you to configure replication from the protected site to the recovery site, as it was before the original failover was executed.

3 Reconfigure SRM to Enable Failback to the Protected Site on page 43

Before you can run a failback, you must create the protection groups and recovery plans required to migrate protected inventory from the recovery site back to the protected site.

4 Restore the Original Configuration on page 43

After a failback is complete, you can restore the original configuration so that the protected and recovery sites resume the roles they had before the failover.

Review and Execute Post-Failover Cleanup Tasks

Before you can execute a failback, you must remove artifacts such as invalid protection groups and unneeded placeholders that are left over from the previous configuration.

If the original protected site is intact after a failover, all the protected virtual machines are still part of its inventory, although they are powered off. You must remove those virtual machines before you can create the protection groups that are needed by the failback. You must also remove various other artifacts at both sites as part of preparing the sites to assume their new roles.

Procedure

1 If the protected site is intact after a failover and you want to restore it to its former status, clean up the protected site.

   a Verify that the SRM server host is operational. Verify that the SRM installation, including its database and vCenter server, has not been affected by the events that motivated the failover.

   b Open a vSphere Client and connect to the vCenter server at the protected site. Log in as a vCenter administrator.

   c Rescan the host bus adapters (HBAs) on the SRM server host.

   After the rescan completes, SRM lists the failed-over virtual machines and their protection groups as invalid because their storage is no longer replicated. (Replication was turned off by the failover.)

   d Delete the invalid virtual machines and protection groups.

2 Clean up the recovery site.

   a Open a vSphere Client and connect to the vCenter server at the recovery site. Log in as a vCenter administrator.

   b Remove the placeholder virtual machines from vCenter inventory.

Reconfigure Replication

Failover stops replication. Failback requires you to configure replication in reverse, from the recovery site to the protected site. Restoring the protected and recovery sites to their original roles requires you to configure replication from the protected site to the recovery site, as it was before the original failover was executed.

Reconfiguring replication is likely to require help from the team that manages vSphere storage for the two sites. The operations required are specific to the arrays that you are using. Generally, you must take the following steps:

- To prepare for a failback, configure the arrays so that the source devices are the ones located at the recovery site and the target devices are the ones located at the protected site.
- After the failback is complete and you are ready to have the protected site and recovery site resume their original roles, configure the arrays so that the source devices are the ones located at the protected site and the target devices are the ones located at the recovery site.

After you have configured replication as needed, force an immediate, one-time replication from the source to the target. This step is always required during a failback, but might not be needed when you are reconfiguring the protected and recovery sites to resume their original roles.

Reconfigure SRM to Enable Failback to the Protected Site

Before you can run a failback, you must create the protection groups and recovery plans required to migrate protected inventory from the recovery site back to the protected site.

After you have prepared both sites for failback, reconfigured array replication, and replicated the source devices at the recovery site to their targets at the protected site, you can create the environment needed for failback. You follow the same steps that you took when you configured SRM for failover, but with the roles of the two sites reversed. For the duration of the failback, the original recovery site becomes the protected site, and the original protected site becomes the recovery site. Because of this temporary role reversal, when the failback procedures refer to the protected or recovery site by name, instead access the sites that are currently playing those roles, not the sites that originally played them.

Procedure

1 Configure the array managers (see "Configure Array Managers," on page 29).

2 Configure inventory mappings at the recovery site (see "Configure Inventory Mappings," on page 31).

3 Create a protection group that includes all the virtual machines that you want to include in the failback (see "Create Protection Groups," on page 33).

4 Create a recovery plan that includes the protection group that you created (see "Create a Recovery Plan," on page 35).

5 Test the recovery plan (see "Test a Recovery Plan," on page 39).

6 After you have verified that the test recovery completed as planned, run the recovery plan (see "Run a Recovery Plan," on page 40).

   Running the recovery plan completes the failback. It powers off the virtual machines at the original recovery site. It then restores the failed-over virtual machines to the original protected site.

What to do next

After the failback completes, replication has again been turned off. The original protected site, while restored to its original role, is no longer protected by a recovery plan. To reinstate its protection, you must reconfigure array replication to use the protected site devices as the source and the recovery site devices as the targets. You then configure the array managers, inventory mappings, protection groups, and recovery plans, as you did when you first configured SRM.

Restore the Original Configuration

After a failback is complete, you can restore the original configuration so that the protected and recovery sites resume the roles they had before the failover.

Procedure

1 At the original recovery site (now restored to that role), clean up any artifacts that remain from the original failover and the subsequent failback.

   - Remove the recovered virtual machines from vCenter inventory and delete them from storage at the recovery site.
   - Remove the protection group and recovery plan that you created in "Reconfigure SRM to Enable Failback to the Protected Site," on page 43.
   - Remove the placeholder virtual machines created at the protected site by the failback.

2 Reconfigure array replication to use the protected site devices as the source and the recovery site devices as the targets.

   See "Reconfigure Replication," on page 42.

3 Configure the array managers (see “Configure Array Managers,” on page 29).

4 Configure the inventory mappings (see “Configure Inventory Mappings,” on page 31).

5 Create the protection groups (see in “Create Protection Groups,” on page 33).

6 Create the recovery plans (see “Create a Recovery Plan,” on page 35).

7 Test the recovery plan (see “Test a Recovery Plan,” on page 39).

Customizing Site Recovery Manager 5

In its default configuration, SRM enables a number of simple recovery scenarios. Advanced users can customize SRM to support a broader range of site recovery requirements.

The default protection and recovery capabilities of SRM can be appropriate for sites that have simple configurations or recovery objectives. Sites that have more complex requirements, such as many virtual machines, a variety of guest operating systems, and application-specific networking requirements, typically need to customize the recovery plans and modify the settings.

This chapter includes the following topics:

Assign Roles and Permissions

To provide additional control over your SRM environment, you can assign permission to perform certain operations on SRM objects to specific users or roles. Permission assignments apply on a per-site basis. Depending on the type of assignment, you might have to add the permission on both sites.

SRM augments vCenter roles and permissions with additional ones that allow fine-grained control over SRM-specific tasks and operations. You can use the SRM Assign Permissions window the same way that you use the Assign Permissions window in the vSphere Client. For more information, see *Managing Users, Groups, Roles, and Permissions* in the vSphere Client Help.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, right-click the **Site Recovery** icon or any of the icons under it and click **Add Permission**.

4 In the **Assigned Permissions** dialog box, select a role from the **Assigned Role** drop-down menu.

This menu displays all the roles that are available from SRM and vCenter. When the role appears, the privileges granted to the role are listed in the section below the role title.

5 To apply the selected role to all child objects of the selected inventory object, select **Propagate to Child Objects**.

6 To select the user or group for the role, click the **Add** button.

7 Identify the user or group.

   a From the **Domain** drop-down menu, select the domain where the user or group is located.

   b Either enter a name in the **Search** text box or select a name from the **Name** list.

   c Click **Add** and then click **OK** when finished.

8 Click **OK** to finish the task.

The list of permissions references all users and groups who have roles assigned to the object and where in the hierarchy those roles are assigned.

What to do next

Repeat the procedure to assign roles and permissions to users at the recovery site.

Customizing a Recovery Plan

You can customize a recovery plan to run commands, display messages that require a response, and change the recovery priority of protected virtual machines.

While a simple recovery plan–one that specifies only a test network to which the recovered virtual machines connect and response times that the test should expect–can provide an effective way to test an SRM configuration, most recovery plans intended for production use must be customized to suit specific needs. For example, a recovery plan for an emergency at the protected site is likely to be different from a planned migration of services from one site to another.

NOTE A recovery plan always reflects the current state of the protection groups that it recovers. If any members of a protection group display problems (for example, a status other than OK), you must correct the problems before you can make any changes to the recovery plan.

Recovery Plan Steps

A recovery plan runs a prescribed series of steps in a specific order. You cannot change the order or purpose of these steps, but you can insert your own steps that display messages and run commands.

Some recovery plan steps are executed during all recoveries, some are executed only during test recoveries, and some are always skipped during test recoveries. Understanding these steps, their order, and the context in which they run is important when customizing a recovery plan.

NOTE When you run a recovery plan, it starts by powering off the virtual machines at the protected site. Machines are powered off in reverse priority order (high-priority machines are powered off last). This step is omitted when you test a recovery plan.

Recovery Order

When a recovery plan runs, virtual machines in the high-priority group are recovered first, followed by the normal-priority group, the low-priority group, and the no-power-on group. Before a priority group is started, all machines in the next-higher priority group must have recovered or failed to recover .

Within a group, virtual machines are always recovered in the order specified by the recovery plan. High-priority virtual machines are recovered serially. Recovery of a machine in this group does not begin until its predecessor in the list has either been recovered (powered on and connected to the network) or has failed to recover within a specified period.

Virtual machines in all other priority groups are recovered serially per ESX host to enable a group of machines that spans several hosts to recover in parallel. During this type of recovery, machines on a specific ESX host are recovered in the order specified by the list, but the recovery order of the entire list is subject to the assignment of virtual machines to hosts. For example, if the first three normal-priority virtual machines are hosted on one ESX host and the fourth is hosted on a different ESX host, the fourth machine in the list might be recovered before the second or third.

Because vCenter limits the number of virtual machines that can be powered on in a single request, recovery plans cannot power on more than 20 virtual machines at a time even if more than 20 ESX hosts available.

Recovery Plan Time-Outs and Pauses

Several kinds of time-outs can occur during the execution of recovery plan steps. These time-outs cause the plan to pause for a specified interval to give the step time to complete.

For information about changing these defaults, see "Change Recovery Site Settings," on page 58.

NOTE Message steps force the plan to pause until they are acknowledged. Before you add a message step to a recovery plan, make sure that it is really necessary. Before you test or run a recovery plan that contains message steps, make sure that someone can monitor the plan's progress and respond to the messages as needed.

Steps That Are Part of All Recovery Plans

The following high-level steps occur during test recoveries and recoveries:

1. Initiate storage operations. The specific set of operations depends on the SRA and whether the recovery is being run as a test.
2. Unregister placeholder virtual machines and register recovered virtual machines with the recovery site vCenter. If any placeholder machines have had settings such as resource pool and memory allocation modified, those settings are applied to the recovered virtual machine.
3. Suspend local virtual machines at the recovery site if requested.
4. Recover high-priority virtual machines. For each virtual machine in the list, the plan runs each step listed in "Virtual Machine Recovery Steps," on page 47. High priority machines are always recovered in list order, regardless of how many physical ESX hosts are involved.
5. Recover normal-priority virtual machines. For each virtual machine in the list, the plan runs each step listed in "Virtual Machine Recovery Steps," on page 47. Machines in this priority group are recovered serially per ESX host.
6. Recover low-priority virtual machines. For each virtual machine in the list, the plan runs each step listed in "Virtual Machine Recovery Steps," on page 47. Machines in this priority group are recovered serially per ESX host.
7. Recover no-power-on virtual machines. For each virtual machine in the list, the plan runs steps 1 and 2 of "Virtual Machine Recovery Steps," on page 47.

At this point, the recovery is complete. If the recovery was run as a test, the plan pauses and prompts you to verify that the test was successful.

Virtual Machine Recovery Steps

Whenever a virtual machine is recovered, regardless of its priority, the recovery plan runs the following steps:

1. Apply IP customization and verify that it succeeded within the specified period.
2. Run any pre-power-on command or message steps.

3 Power on the virtual machine and verify that VMware Tools reports an OS heartbeat within the specified period.

4 Run any post-power-on command or message steps.

NOTE Post-power-on command steps provide an application-specific way to verify that a recovered virtual machine has all the capabilities that you expect. For example, after powering on a recovered database server, you could execute a simple database query from a script and declare the recovery complete (by having the script exit with status of 0) only if the script receives the expected response. In the absence of such additional steps, the virtual machine is considered to have been recovered if it powers on and connects to the network.

Steps That Are Not Executed During a Test Recovery

When you run a recovery plan, it starts by shutting down protected virtual machine at the protected site. Machines are shut down in reverse priority order (high-priority machines are shut down last). This step is omitted when you test a recovery plan.

Cleanup Steps That Are Executed Only During a Test Recovery

Clean up steps are performed after a recovery plan test. The steps begin executing after you have responded to the prompt displayed after the test completes.

1 Power off each recovered virtual machine.

2 Unregister the recovered virtual machines from the recovery site vCenter and re-register the placeholders.

3 Clean up replicated storage snapshots that were used by the recovered virtual machines during the test.

Guidelines for Writing Command Steps

When you create a command step to add to a recovery plan, make sure that it takes into account the environment in which it must run. Errors in a command step affect the integrity of a recovery plan, so test the command on the recovery site SRM server host before you add it to the plan.

All batch files or commands that you add to a recovery plan must meet the following requirements:

- You must start the Windows command shell using its full path on the local host. For example, to run a script located in `c:\alarmscript.bat`, use the following command line:

  ```
  c:\windows\system32\cmd.exe /c c:\alarmscript.bat
  ```

- Batch files and commands must be installed locally on the SRM server host at the recovery site.
- Batch files and commands must complete within 300 seconds. Otherwise, the recovery plan terminates with an error. To change this limit, see "Change Recovery Site Settings," on page 58.
- Batch files or commands that produce output that contains characters with ASCII values greater than 127 must use UTF-8 encoding. Only the final 4KB of script output is captured in log files and recovery history. Scripts that produce more output can redirect the output to a file rather than sending it to the standard output to be logged.

Execution Environment for Command Steps

Command steps run with the identity of the LocalSystem account on the SRM server host at the recovery site. When a command step runs, a number of environment variables are available for it to use. Table 5-1 lists the environment variables that are available to all command steps.

Table 5-1. Environment Variables Available to All Command Steps

Name	Value	Example
VMware_RecoveryName	Name of the recovery plan that is executing	"Plan A"
VMware_RecoveryMode	Recovery mode	"test" or "recovery"
VMware_VC_Host	Host name of the vCenter host at the recovery site	"vc_hostname.example.com"
VMware_VC_Port	Network port used to contact the vCenter host	"443"

The environment variables listed in Table 5-2 are also set if the command step is executing on a recovered virtual machine.

Table 5-2. Environment Variables Available to Command Steps Running on Recovered Virtual Machines

Name	Value	Example
VMware_VM_Uuid	UUID used by vCenter to uniquely identify this virtual machine	"4212145a-eeae-a02c-e525-ebba70b0d4f3"
VMware_VM_Name	Name of this virtual machine, as set at the protected site	"My New Virtual Machine"
VMware_VM_Ref	Managed object ID of the virtual machine	"vm-1199"
VMware_VM_GuestName	Name of the guest OS as defined by the VIM API	"otherGuest"
VMware_VM_GuestIp	IP address of the virtual machine, if known	"192.168.0.103"
Vmware_VM_Path	Path to this virtual machine in recovery site inventory	"[datastore-123] jquser-vm2/jquser-vm2.vmdk"

Customize Recovery Plan Steps

You can customize many recovery plan steps to extend the basic functions provided by a default recovery plan.

To customize recovery plan steps, open the Recovery Steps page of a recovery plan.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, navigate to Recovery Plans, and click the plan that you want to customize.

4 In the recovery plan window, click the **Recovery Steps** tab.

5 Right-click the step that you want to modify and select an option from the menu.

To export the entire plan as an HTML document for your reference, right-click any step and then click **Export**. To edit the properties of the plan, right-click any step and then click **Edit Recovery Plan**.

Specify Virtual Machine Recovery Priority

By default, all virtual machines in a new recovery plan are members of the normal priority group. Members of this group are recovered in the order that they were created on the protected datastore. You can move a virtual machine to a different priority group or to a different priority within a group.

Procedure

1 Open the Recovery Steps page for the plan, as described in "Customize Recovery Plan Steps," on page 49.

2 To display the virtual machines in the normal priority group, expand the Recover Normal Priority Virtual Machines step .

Unless you have modified recovery priorities, all virtual machines in the plan are included in the Recover Normal Priority Virtual Machines step.

NOTE In a recovery plan, a virtual machine is always listed under the name it had when the plan was created. If you change the name after the plan is created, the change is not reflected in the plan.

3 To raise the recovery priority of a virtual machine, right-click it and click **Move Up**.

You can move a virtual machine to a higher priority within its current group, or to a higher priority group.

4 To lower the recovery priority of a virtual machine, right-click it and click **Move Down**.

You can move a virtual machine to a lower priority within its current group, or to a lower priority group.

What to do next

Review the list of virtual machines in the Shutdown Virtual Machines at Protected Site step. Modifying the recovery priority of a virtual machine does not affect the priority with which it is powered off on the protected site. If you want to change the power off priority of a virtual machine, you must do so explicitly by moving it up or down in one of the Shutdown steps.

NOTE Shutdown steps are run in reverse priority order; high-priority virtual machines are powered off last.

Add Messages to a Recovery Plan

You can customize a recovery plan to include messages that are displayed in the vSphere client when the plan is tested or run. Each message adds a step to the recovery plan, and pauses the plan at that step until the message is acknowledged.

You can add message steps to any part of a recovery plan.

NOTE Message steps force the plan to pause until they are acknowledged. Before you add a message step to a recovery plan, make sure that it is really necessary. Before you test or run a recovery plan that contains message steps, make sure that someone can monitor the plan's progress and respond to the messages as needed.

Procedure

1 Open the Recovery Steps page for the plan, as described in "Customize Recovery Plan Steps," on page 49.

2 Right click the recovery plan step that you want the message to precede and click **Add Message** to open the Edit Message Step dialog box.

3 Type the message text and click **OK.**

The message is added to the recovery plan as a new step, and subsequent steps are renumbered. When you test or run the recovery plan, the plan pauses, displays the message, and waits for acknowledgment when it reaches this step.

Add Commands to a Recovery Plan

You can customize a recovery plan to include commands that are executed on the SRM server host at the recovery site when the plan is tested or run.

You can add command steps to any part of a recovery plan. When you create a command step to add to a recovery plan, make sure that it takes into account the environment in which it must run. For more information, see "Guidelines for Writing Command Steps," on page 48.

Procedure

1 Open the Recovery Steps page for the plan, as described in "Customize Recovery Plan Steps," on page 49.

2 Right-click the recovery plan step that you want the command to precede and click **Add Command**.

3 Type the command line and click **OK.**

The command is added to the recovery plan as a new step, and subsequent steps are renumbered. When you test or run the recovery plan, the plan executes the command line on the SRM server host at the recovery site when it reaches this step.

Customize the Recovery of an Individual Virtual Machine

You can configure a virtual machine in a recovery plan to use a prescribed customization specification, or to execute message or command steps when it is recovered.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 Expand the **Site Recovery** icon in the navigation tree, navigate to Recovery Plans, and click the plan that you want to customize.

4 In the recovery plan window, click the **Virtual Machines** tab.

5 Right-click a virtual machine in the list, and click **Configure**.

The Edit Virtual Machine Properties dialog box enables you to select a customization specification for the virtual machine, and also to add message and command steps that execute before or after the machine is powered on.

a Select a customization specification. Click **Browse** to see a list of customization specifications available from the vCenter at the recovery site.

You can also enter a description of the specification you apply. Only the IP properties from the selected specification are applied. All other properties in the specification are ignored. If you have used the `dr-ip-customizer.exe` command to customize virtual machines in the recovery plan, you do not need to specify that customization here.

b Click **Next** to add a message or command step that executes before the machine is powered on.

c Click **Next** to add a message or command step that executes after the machine is powered on.

Message and command steps added to the recovery steps for a virtual machine operate like message and command steps added to a recovery plan. For more information, see "Guidelines for Writing Command Steps," on page 48.

The customizations you specify are saved as properties of the placeholder virtual machine and then applied to the recovered virtual machine when a recovery plan is run or tested.

NOTE If you remove the protection of a virtual machine, all recovery customizations are lost.

Report IP Address Mappings for a Protection Group

The IP address map reporter generates an XML document describing the IP properties of protected virtual machines and their placeholders, grouped by site and recovery plan. This document can help you understand the network requirements of a recovery plan.

Because the IP address map reporter must connect to both sites, you can run the command at either site. You are prompted to supply the vCenter login credentials for each site when the command runs.

Procedure

1 Open a command shell on the SRM server host at either the protected or recovery site.

2 Change to the `C:\Program Files\VMware\VMware vCenter Site Recovery Manager\bin` directory.

3 Run the `dr-ip-reporter.exe` command, as shown in this example.

```
dr-ip-reporter.exe -cfg ..\config\vmware-dr.xml -out c:\tmp\report.xml
```

To restrict the list of networks to just the ones required by a specific recovery plan, include the `-plan` option on the command line, as shown in this example:.

```
dr-ip-reporter.exe -cfg ..\config\vmware-dr.xml -out c:\tmp\report.xml -plan Plan-B
```

NOTE The command normally asks you to verify the thumbprints presented by the certificates at each site. You can suppress the verification request by including the `-I` option.

Customize IP Properties for a Group of Virtual Machines

The IP property customization tool enables you to specify IP properties for any or all of the virtual machines in a recovery plan by editing a file that the tool generates.

SRM includes a tool that allows you to specify IP properties (network settings) for any or all of the virtual machines in a recovery plan by editing a comma-separated-value (CSV) file that the tool generates. Initially, this file includes a single row for each placeholder virtual machine in the plan. You can edit the file to add a row for each network adapter in each placeholder virtual machine and then customize the network settings for each adapter. When you are finished, you use the edited file as input to a command that creates customization specifications for the placeholder virtual machines.

The tool is named `dr-ip-customizer.exe`, and is installed in the `bin` subdirectory of the SRM installation directory.

Procedure

1 Open a command shell on the SRM server host at the recovery site.

2 Change directory to `C:\Program Files\VMware\VMware vCenter Site Recovery Manager\bin`.

3 Run the `dr-ip-customizer.exe` command, as shown in this example.

```
dr-ip-customizer.exe -cfg ..\config\vmware-dr.xml -csv c:\tmp\example.csv -cmd generate
```

In an SRM recovery plan that defines three placeholder virtual machines, the generated file might look like this:

```
VM ID,VM Name,Adapter ID,MAC Address,DNS Domain,Net BIOS,Primary WINS,Secondary WINS,IP
Address,Subnet Mask,Gateway(s),DNS Server(s),DNS Suffix(es)
shdw1,srm1,0,,,,,,,,,
shdw2,srm2,0,,,,,,,,,
shdw3,srm3,0,,,,,,,,,
```

The file consists of a header row that defines the meaning of each column, and a single row for each placeholder virtual machine found in the recovery plan. The only columns populated with values are:

- VM ID (the ID for the placeholder virtual machine)
- VM Name (the name of the placeholder virtual machine)
- Adapter ID (always 0, which designates global IP settings, not specific to any adapter)

All the other columns are empty.

4 Edit the generated file to customize IP properties for the virtual machines in the recovery plan.

This example shows the result of opening the output of `dr-ip-customizer` with a spreadsheet program and creating additional rows that define network settings for placeholder virtual machines in the recovery plan.

Table 5-3. IP Customization Spreadsheet

VM ID	VM Name	Adapter ID	MAC Address	DNS Domain	NetBIOS	Primary WINS	Secondary WINS	IP Address	Subnet Mask	Gateway(s)	DNS Server(s)	DNS Suffix(es)
shdw1	srm1	0									10.10.10.1	example.com
shdw1		1	00:1f:3a:38:29:9c	example.com				dhcp				
shdw2	srm2	0										
shdw2		1	00:1c:23:3d:b9:e3	example.com		10.10.10.10		10.13.99.4	255.255.0.0	10.10.10.100	10.10.10.1	
shdw2		1									10.10.10.2	
shdw3	srm3	0										
shdw3		1	00:1a:3f:b8:f3:79	example.com		10.10.10.10		10.13.99.5	255.255.0.0	10.10.10.100	10.10.10.1	
shdw3		1									10.10.10.2	

The following rules apply when you modify a CSV file created by the dr-ip-customizer utility.

- Commas are not allowed in any field.
- The VM Name field is intended as a reference for the user customizing the file. It is populated when the CSV file is created but ignored when the modifications are applied to the recovery plan. It cannot be used to rename a virtual machine.
- The only fields that you can modify for a row where Adapter ID is 0 are DNS Server(s) and DNS Suffix(es). These values, if specified, are inherited by all other adapters for that VM ID.

- To define properties for a specific adapter on a placeholder virtual machine, create a new row that contains that virtual machine's ID in the VM ID column and the adapter ID (the virtual PCI slot in which the adapter is installed on the placeholder virtual machine) in the Adapter ID column, then specify values for the other columns.
- To specify more than one value for a column, create an additional row for that adapter and include the value in the column in that row. In Table 5-3, additional rows define a secondary DNS server for the placeholder virtual machines shdw2 and shdw3.
- To customize a placeholder virtual machine as a DHCP client, enter dhcp in the IP Address field, as shown in the second row of Table 5-3. For any non-zero adapter ID that is not a DHCP client:
  - You must specify values for IP Address, Subnet Mask, Gateway(s), and DNS Server(s) unless global values for these properties exist (in the row for Adapter ID zero for that VM ID).
  - Global values, if specified, are overridden by values you specify for each non-zero adapter ID
  - The NetBIOS column, if not left empty, must contain one of the following strings: disableNetBIOS, enableNetBIOS, or enableNetBIOSViaDhcp.
  - If you are customizing multiple adapters for a virtual machine and want to be sure that the customizations in a specific row apply to a specific adapter, specify the adapter's MAC address as pairs of hexadecimal digits separated by the colon character. Character case is not considered.

5 Run `dr-ip-customizer.exe` to apply the customized IP properties.

Change directory to `C:\Program Files\VMware\VMware vCenter Site Recovery Manager\bin` and run the following command.

```
dr-ip-customizer.exe -cfg ..\config\vmware-dr.xml -csv c:\tmp\example.csv -cmd command
```

where command is one of the following commands:

- `create` – applies the customizations to the virtual machines listed in the `csv` file
- `drop` – removes the customizations from the virtual machines listed in the `csv` file
- `recreate` – applies revised customizations to virtual machines that have already been customized

You can include a `-verbose` option on any `dr-ip-customizer.exe` command line to log additional diagnostic messages.

The specified customizations are applied to all of the virtual machines named in the `csv` file during a recovery. (You do not need to select a customization specification for these machines when you edit their properties in a recovery plan.)

Configure Protection for a Virtual Machine or Template

You can edit the protection properties of any virtual machine or template in a protection group. You can change the resource mappings, attached storage devices and their datastores, and other properties that control the configuration with which the virtual machine is recovered.

You must configure protection for virtual machines that have a status of Not Configured, Mapping Missing, or Mapping Invalid.

NOTE If inventory mappings have been established for a site, you cannot override them by configuring the protection of individual virtual machines. If you need to override inventory mappings for a few members of a protection group, use the vSphere Client to connect to the recovery site and edit the network settings of the placeholders or move them to a different folder or resource pool.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, navigate to the protection group that includes the virtual machine that you want to configure.

4 On the Virtual Machines page, right-click a virtual machine and click **Configure Protection**.

5 In the Edit Virtual Machine Properties window, review and configure properties as needed.

a In the resource list, click **Folder** to review the recovery site folder to which this virtual machine is assigned.

If inventory mappings have not been established for this site, you can edit this property.

b Click **Next** to review the recovery site host to which the virtual machine is assigned.

If inventory mappings have not been established for this site, you can edit this property.

c Click **Next** to review the recovery site resource pool to which this virtual machine is assigned.

If inventory mappings have not been established for this site, you can edit this property.

d Click **Next** to review the recovery site networks to which this virtual machine is assigned.

If inventory mappings have not been established for this site, you can edit this property.

e Click **Next** to review the list of storage devices attached to the virtual machine and verify that they are all in the same datastore group or have appropriate storage on a nonreplicated datastore at the recovery site.

If any device has a Recovery Location that has a status of Not Configured, click **Browse** to find an appropriate datastore at the recovery site, or click **Detach** to detach the device during recovery.

f Click **Next** to review the datastore that you originally selected for the placeholder virtual machines in the protection group.

If inventory mappings have not been established for this site, you can edit this property.

g Select a customization specification. Click **Browse** to see a list of customization specifications available from the vCenter at the recovery site.

You can also enter a description of the specification you apply. Only the IP properties from the selected specification are applied. All other properties in the specification are ignored. If you have used the `dr-ip-customizer.exe` command to customize virtual machines in the recovery plan, you do not need to specify that customization here.

h Click **Next** to select a recovery priority group for the virtual machine. See "Specify Virtual Machine Recovery Priority," on page 50.

i Click **Next** to add a message or command step that executes before the machine is powered on.

j Click **Next** to add a message or command step that executes before the machine is powered on.

6 Click **Finish** to apply the new configuration to the selected virtual machine.

Repair Placeholder Virtual Machines After a Failed Test Recovery

If the vCenter Server at the recovery site becomes inaccessible during a test recovery, some virtual machines in a protection group might lose their protection configuration. Virtual machines in this state have a status of Needs Repair. You can repair these virtual machines to restore protection.

When the SRM server at the protected site is cleaning up after a test recovery, it must restore the placeholder virtual machines. Doing so requires the SRM server to contact the vCenter Server at the recovery site and retrieve configuration information for the placeholders. If the vCenter Server cannot be contacted, the placeholders cannot be restored correctly. Virtual machines in this condition are listed on the Protection Group page with a status of **Needs Repair**. After the vCenter server at the recovery site becomes accessible again, you can repair these virtual machines by clicking the **Repair All** button.

Prerequisites

The vCenter Server at the recovery site must be accessible to repair the virtual machines.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

   Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 In the Site Recovery tree view, expand the **Protection Groups** item.

   Protection groups that include virtual machines that need repair are highlighted with a warning icon.

4 Open the protection group and click the **Virtual Machines** tab.

   Each virtual machine that needs repair is listed with a status of Needs Repair.

5 Click **Repair All** to repair the virtual machines that have a status of **Needs Repair**.

The SRM server at the recovery site contacts the vCenter Server at the recovery site, retrieves protection configurations for the affected virtual machines, and applies those configurations, restoring the status of the machines to OK.

Configure SRM Alarms

SRM adds feature-specific alarms to the ones supported by vCenter. You can configure SRM alarms to send an email notification or trigger some other action.

vCenter provides a comprehensive and flexible alarm facility. As a vCenter extension, SRM can add its own alarms to the ones provided by vCenter. The SRM Alarms window lists all SRM alarm events and allows you to edit their settings to specify what action to take when an event triggers the alarm. None of the SRM alarms are configured by default to take any action. If you want to enable actions for any of them, you must configure them to do so.

NOTE If you want alarms to provide email notification, you must first configure vCenter mail sender settings. See the vCenter help.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the recovery site.

   Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 Click the **Alarms** tab to display the list of SRM alarms.

4 Right-click an alarm and click **Edit Settings**.

5 In the Edit Settings dialog box, click the Actions tab. In the Actions window, click **Add** to add an action.

The default action for every event is **Send a notification e-mail**. To change this action, click it and select a different action from the drop-down box. For more information about actions, see the vCenter help.

Working with Advanced Settings

The Advanced Settings dialogs enable you to view or change many custom settings for the SRM service.

The Advanced Settings dialog box provides an easy way for a user with adequate privileges to change a number of default values that affect the operation of various SRM features.

NOTE Changes you make in the Advanced Settings dialog boxes overwrite the contents of the SRM `vmware-dr.xml` configuration file on the SRM server host.

Procedure

1 Open a vSphere Client and connect to the vCenter server at the protected site.

Log in as a vSphere administrator.

2 On the vSphere Client Home page, click the **Site Recovery** icon.

3 Right-click **Site Recovery** in the vSphere Client navigation pane and click **Advanced Settings**.

4 In the navigation pane of the Advanced Settings window, click a setting category.

5 In the category window, make your changes.

6 Click **OK** to save your changes and close the Advanced Settings window.

7 Repeat the procedure as needed at the recovery site.

Guest Customization Settings

Change this setting only if instructed to do so by VMware Support.

Change License Key Settings

Use the Advanced Settings Licensing page to change the SRM license key or add a new one.

Procedure

1 Right-click **Site Recovery** in the vSphere Client navigation pane and click **Advanced Settings**.

2 In the navigation pane of the Advanced Settings window, click **Licensing**.

3 Enter the SRM license key in the **Licensing.LicenseKey** text box.

The first time you open the **Licensing** page, the evaluation key is displayed in the **Licensing.LicenseKey** field.

4 Click **OK** to save your changes and close the Advanced Settings window.

If you enter a valid license key, it is displayed in the **Licensing.LicenseKey** text box each time you open the Licensing page. If you enter an invalid license key, the previous valid license key (or the evaluation key, if no other valid key has ever been entered) is displayed the next time you open the Licensing page.

Change Recovery Site Settings

Use the Advanced Settings Recovery page to adjust default values for time-outs that occur when you test or run a recovery plan.

Several kinds of time-outs can occur during the execution of recovery plan steps. These time-outs cause the plan to pause for a specified interval to give the step time to complete.

- Command line timeout – By default, SRM allows 300 seconds for a command step to complete. If a command step takes longer than 300 seconds, the step terminates and the recovery plan fails with an error.
- Power state change timeout – By default, SRM allows 120 seconds for a virtual machine at the protected site to respond to a power-down request when testing or running a recovery plan. If the request does not complete in this interval, the plan skips to the next virtual machine in the list (or to the next step) and reports a recovery plan error.

Procedure

1 Right-click **Site Recovery** in the vSphere Client navigation pane and click **Advanced Settings**.

2 In the navigation pane of the Advanced Settings window, click **Recovery**.

3 Modify recovery site settings as needed.

- To change the command-line timeout, enter a new value in the **Recovery.calloutCommandLineTimeout** field. The new value applies to all command steps.
- To change the power state change timeout, enter a new value in the **Recovery.powerStateChangeTimeout** field. The new time-out value applies to all power state changes to virtual machines at the protected site.

4 Click **OK** to save your changes and close the Advanced Settings window.

Change SAN Provider Settings

The SAN provider is the interface between SRM and your storage replication adapter (SRA). Some SRAs require you to make changes to default SAN provider values. You can change the default timeout values and other behaviors of the SRM SAN provider.

For more information about these values, see the SRA documentation from your array vendor.

Procedure

1 Right-click **Site Recovery** in the vSphere Client navigation pane and click **Advanced Settings**.

2 In the navigation pane of the Advanced Settings window, click **SanProvider**.

3 Modify the SAN provider settings as needed.

- To change the length of time that SRM waits for a command issued by the SRA to complete, enter a new value in the **SanProvider.calloutCommandTimeout** text box.
- To force removal, upon successful completion of a test recovery, of the snap-xx prefix applied to recovered datastore names, select the **SanProvider.fixRecoveredDatastores** checkbox.
- To change the interval that SRM waits for a host to reconnect during a host bus adapter (HBA) rescan, enter a new value in the **SanProvider.hostReconnectTimeoutSec** text box.
- To change the number of HBA rescans that SRM executes when you test or run a recovery plan, enter a new value in the **SanProvider.hostRescanRepeatCount** text box.

- To change the interval that SRM waits for each HBA rescan to complete, enter a new value in the **SanProvider.hostRescanTimeoutSec** text box.
- To change the interval between datastore group computations, enter a new value in the **SanProvider.minLunGroupComputationInterval** text box.

4 Click **OK** to save your changes and close the Advanced Settings window.

Change Local Site Settings

SRM monitors consumption of resources on the SRM server host, and it raises an alarm when a resource threshold is reached. You can use the Advanced Settings **localSiteStatus** page to change the thresholds and the way the alarms are raised to suit the needs of your installation and administrative staff.

Procedure

1 Right-click **Site Recovery** in the vSphere Client navigation pane and click **Advanced Settings**.

2 In the navigation pane of the Advanced Settings window, click **localSiteStatus**.

3 Change the settings as needed.

- To change the interval at which SRM checks the CPU usage, disk space, and free memory at the local site, enter a new value in the **localSiteStatus.checkInterval** field.
- To change the interval that which SRM waits between raising alarms about CPU usage, disk space, and free memory at the local site, enter a new value in the **localSiteStatus.eventFrequency** field.
- To change the percentage of CPU usage that causes SRM to raise a high CPU usage event, enter a new value in the **localSiteStatus.maxCpuUsage** field.
- To change the percentage of free disk space that causes SRM to raise a low disk space event, enter a new value in the **localSiteStatus.minDiskSpace** field.
- To change the amount of free memory that causes SRM to raise a low memory event, enter a new value in the **localSiteStatus.minMemory** field.

4 Click **OK** to save your changes and close the Advanced Settings window.

Change Remote Site Settings

Use the Advanced Settings **remoteSiteStatus** page to modify default values that the SRM server at the site to which the vSphere client is currently connected uses to determine whether the SRM server at the remote site is available

SRM monitors the connection between the members of an SRM site pair (a protected site and its recovery site) and, by default, raises alarms when this connection is interrupted. You can change the criteria that cause a "remote site down" event and also change the way the related alarms are raised to suit the needs of your installation and administrative staff.

Procedure

1 Right-click **Site Recovery** in the vSphere Client navigation pane and click **Advanced Settings**.

2 In the navigation pane of the Advanced Settings window, click **remoteSiteStatus**.

3 Modify settings as needed.

- To change the interval at which SRM checks to see whether the SRM server at the remote site is available, enter a new value in the **remoteSiteStatus.checkInterval** field.
- To change the number of failed remote site status checks required to trigger a Remote Site Down alarm, enter a new value in the **remoteSiteStatus.panicDelay** field.

- To change the interval between Remote Site Down alarms, enter a new value in the **remoteSiteStatus.panicRepeatDelay** field.
- To change the number of remote site status checks to try before declaring the check a failure, enter a new value in the **remoteSiteStatus.warningDelay** field.

4 Click **OK** to save your changes and close the Advanced Settings window.

Avoiding Replication of Paging Files and Other Transient Data

While SRM allows you to replicate transient data such as Windows paging files or virtual machine swapfiles, such data need not be replicated. Preventing replication of such data avoids unnecessary consumption of network bandwidth.

In the default configuration, virtual machines that use replicated datastores have all of their storage replicated, including swap and paging files for which replication is unnecessary. These files, even if replicated, are overwritten when a recovered virtual machine is powered on. While it does no harm to place them on a replicated datastore, doing so adds unnecessary load to your replication infrastructure.

You can configure protected virtual machines to use local (nonreplicated) storage for swapfiles, Windows paging files, or both.

Specify a Nonreplicated Datastore for Swapfiles

Every virtual machine requires a swapfile, which is normally created in the same datastore as the other virtual machine files. When you use SRM, this datastore is replicated. To prevent swapfiles from being replicated, create them on a nonreplicated datastore.

If you are using a nonreplicated datastore for swapfiles, you must create a nonreplicated datastore for all protected clusters at both the protected and recovery sites. For more information, see the vSphere documentation.

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**.

Create a Nonreplicated Virtual Disk for Paging File Storage

You can avoid replication of a virtual machine's Windows paging file by creating a virtual disk on a nonreplicated datastore, configuring Windows to create its paging file on that disk, and configuring a nonreplicated copy of that disk at the recovery site.

In the default configuration, Windows creates its paging file on the system disk (typically `C:`). Paging files created on this disk are always replicated when the virtual machine uses a replicated datastore. You can configure any virtual machine to use a virtual disk on a nonreplicated datastore for its paging file. SRM detects that the virtual machine depends on a nonreplicated virtual disk (the paging file disk) and removes that virtual machine from its protection group until you make a copy of that virtual disk file at the recovery site for the recovered virtual machine to use.

NOTE To simplify creating a nonreplicated virtual disk for every virtual machine in a protection group, you can create a virtual disk file template and then clone it.

Procedure

1 At the protected site, select a nonessential virtual machine (one that you can power off during this procedure) or create a temporary virtual machine for this purpose.

   Because of differences in the NTFS file system among Windows releases, you must perform this procedure for each version of Windows that the protection group includes.

2 Power off the selected virtual machine and add a new disk to it.

   Create the disk file on a nonreplicated datastore at the protected site in a location where you typically store virtual machine templates. This disk becomes the template for all nonreplicated paging file disks. Create it with adequate capacity for a typical Windows paging file.

3 Power on the selected virtual machine and then create and format a partition on the new disk.

4 Configure the virtual machine to create its paging file on the new disk.

5 Power off the virtual machine and disconnect the new disk from it.

   You can clone this disk to provide nonreplicated paging file storage for the other virtual machines.

6 To make the template available for cloning at the recovery site, copy it to a folder in a nonreplicated datastore at the recovery site.

   You must copy the `.vmdk` file and its flat counterpart (for example, `pagedisk.vmdk` and `pagedisk-flat.vmdk`).

7 At the recovery site, use the `vmkfstools` command to create a clone of the copied disk.

   Create one clone for every placeholder virtual machine, but do not attach any clones to a virtual machine. The clones are assigned as part of the protection configuration process and are attached during recovery.

8 At the protected site, configure each protected virtual machine.

   a Use the `vmkfstools` command to clone the disk.

      Create the clone on a nonreplicated datastore at the protected site, and then copy it to a nonreplicated datastore at the recovery site with the original `.vmdk` file.

   b Connect the cloned disk to the virtual machine, and then power on the virtual machine and assign a drive to the cloned disk.

   c Configure the virtual machine to create its paging file on the cloned disk.

d Power off and then power or the virtual machine so that it writes its paging file to the new location on the cloned disk.

At this point, the protected virtual machine is writing its paging file to a disk on a nonreplicated datastore at the protected site. Until you specify a recovery site location for this disk, the virtual machine does not have a valid protection configuration.

e Assign recovery site storage for the paging file disk to one of the clones that you copied from the protected site.

See "Configure Protection for a Virtual Machine or Template," on page 54.

Initially, the paging file disk has a Recovery Location that is Not Configured. Click **Browse**, and then browse to the cloned `vmdk` file at the recovery site.

After you have configured the virtual machine to use the nonreplicated disk at the recovery site, SRM considers the virtual machine's storage properly configured and returns it to the protection group.

What to do next

After the changes at the protected site are replicated to the recovery site, you can test the recovery plan to verify that the recovered virtual machines are using the nonreplicated paging file.

After you reconfigure a virtual machine to use a paging file disk, you can delete the old, unused paging file from its system disk.

Troubleshooting SRM 6

If you have problems with storage replication, site pairing, or guest customization, you can try to troubleshoot the problem. To help identify the cause, you might need to collect SRM server or client log files to review or send to VMware Support.

Errors encountered during SRM operations are displayed in error dialogs or shown in the Recent Tasks window. Most errors also generate an entry in an SRM log files. It is important to check the recent tasks and log files for the recovery site and the protected site.

When searching for the cause of a problem, also check the VMware knowledge base at http://kb.vmware.com.

This chapter includes the following topics:

No Replicated Datastores Listed

If the Review Replicated Datastores page does not show any replicated datastores after you have configured the array managers, this usually indicates that there are no virtual machines on any datastore that uses the array you are configuring.

Problem

After you complete the Configure Array Managers task, the protected site and recovery site arrays are discovered successfully, but the Review Replicated Datastores page shows no replicated datastores.

Cause

The Review Replicated Datastores page shows only the replicated datastores that use the array and host at least one virtual machine. Replicated datastores that do not have virtual machines are not listed.

Solution

1 Ask your storage management team which datastores are hosted on the arrays that you have configured to work with SRM.

2 In the vSphere Client, create a virtual machine on each of the datastores managed by the arrays.

3 In the Protection Setup area of the SRM Summary window, navigate to the Array Managers line and click **Configure**.

4 In the Configure Array Mangers wizard, click **Next** on the Protected Site Array Managers page and then click **Next** on the Recovery Site Array Managers page.

The Review Replicated Datastores page should now display each replicated datastore that contains at least one virtual machine.

Inconsistent Mount Points Warning When Configuring NFS Arrays

When you are configuring an array manager for an NFS array, you might get an inconsistent mount points warning. Work with your storage team to correct the problem on the ESX hosts that mount the devices.

Problem

When you are configuring the array manager for an NFS array, the Review Replicated Datastores page shows one or more warnings similar to `Datastore datastore_name is mounted on multiple hosts with different NFS server addresses`.

Cause

The ESX Add Storage wizard allows an administrator to specify NFS mount points using either a host name or an IP address. If all ESX hosts do not specify the mount point for a specific NFS volume in the same way (using either the host name or IP address), SRM detects the inconsistency and does not allow the device to be used as a protected datastore.

Solution

Your storage management team can use the ESX Add Storage wizard to correct the inconsistency. Using the IP address rather than the host name when specifying NFS mount points can make it easier to avoid inconsistent mount point warnings.

Array Script Files Not Found

When SRM displays an error message indicating that array scripts cannot be found, it usually indicates that new SRA was installed but the SRM server was not restarted.

Problem

After you install a new SRA, SRM displays an error message containing the text "Unable to find any array script files."

Cause

You did not restart the SRM server after you installed the new SRA.

Solution

Restart the SRM server host, or use the Windows Service Control Manager to stop and restart the Site Recovery Manager service.

Expected Virtual Machine File Path Cannot Be Found

When a virtual machine file path cannot be found while a recovery plan is being tested or run, it often indicates that the virtual machine was created recently and its files have not yet been replicated to the recovery site.

Problem

While you are testing or running a recovery plan, an error of the form "Expected virtual machine file path path-name cannot be found" is logged.

Cause

This error usually occurs when a virtual machine has been recently created but its files have not yet been replicated to the recovery site. For instance, you have created a virtual machine at the protected site, added it to a protection group, and then tested or run a recovery plan that includes the new virtual machine. If the virtual machine files have not yet been replicated to the recovery site, the recovery plan cannot recover the virtual machine.

This problem can also occur if the virtual machine files have been replicated but then moved by a recovery site administrator using Storage vMotion or a similar tool.

Solution

Make sure that all virtual machines in a protection group have been replicated to the recovery site before you test or run a recovery plan for the protection group. The virtual machine files can be missing even if the corresponding placeholder virtual machine exists. Placeholders are created by SRM, not by array replication, and do not include the files necessary to recover a virtual machine.

Recovery Plan Time-Out During the Change Network Settings Step

When you are testing or running a recovery plan and it fails during the Change Network Settings step, verify that the time-out for this step is long enough, and that the requested network settings can be provided by the recovery network.

Problem

When attempting to recover a virtual machine, the Change Network Settings step of a recovery plan fails with a Timed Out error.

Cause

If the network settings for the virtual machine's guest operating system are being customized as part of recovery, this error indicates that the virtual machine failed to report the new network settings within the allotted amount of time.

Solution

1 Ensure that the latest version of VMware Tools is installed on all protected virtual machines.

   VMware Tools is required for IP customization.

2 Edit the recovery plan to increase the **Change Network Settings** timeout value.

   The default timeout is 300 seconds. A shorter timeout might not allow enough time for the guest operating system to boot and report that it has successfully changed its network settings. Some virtual machines might require a longer timeout than the default.

3 If the virtual machine is configured as a DHCP client, verify that a DHCP server is available on the recovery network.

4 If the virtual machine is configured with a static IP address, verify that the address is available and not in use on the recovery network.

What to do next

The customization process creates log files on each virtual machine. On Windows, these log files are written in the `C:\windows\temp\vmware-imc\` directory. On Linux, these log files are written in the `/var/log/vmware-imc/` directory. Review the log files for more information about errors that prevented the step from completing in time.

Collecting SRM Log Files

SRM creates several log files that contain information that can help VMware Support diagnose problems. You can use the SRM log collector to simplify log file collection.

The SRM server and client generate separate sets of log files. The SRM server log files contain information about the server configuration and messages related to server operations. The SRM client log files contain information about the client configuration and messages related to client plug-in operations. The SRM log collector retrieves the files and collects them in a compressed (zipped) folder on your desktop.

Collect SRM Server Log Files

You can collect SRM server log files into a log bundle.

Procedure

- To initiate the collection of SRM server log files from the Start menu:
  a. Log in to the SRM server host.
  b. Select **Start > Programs > VMware > VMware Site Recovery Manager > Generate vCenter Site Recovery Manager log bundle**.
- To initiate the collection of SRM server log files from the Windows command line:
  a. Start a Windows command shell on the SRM server host.
  b. Change directory to `C:\Program Files\VMware\VMware vCenter Site Recovery Manager\bin`.
  c. Run the following command.

     ```
     cscript srm-support.wsf
     ```

The individual log files are collected in a file named `srm-plugin-support-MM-DD-YYYY-HH-MM.zip`, where MM-DD-YYYY-HH-MM indicates the month, day, year, hour, and minute when the log files were created.

Collect an SRM Client Log Bundle

Use the Windows command line to initiate the collection of SRM client log files.

Procedure

1. Start a Windows command shell on the SRM client host.
2. Change directory to `C:\Program Files\VMware\VMware vCenter Site Recovery Manager\bin`
3. Run the following command:

   ```
   cscript srm-support.wsf
   ```

The individual log files are collected in a file named `srm-plugin-support-MM-DD-YYYY-HH-MM.zip`, where MM-DD-YYYY-HH-MM indicates the month, day, year, hour, and minute when the log files were created.

Index

R

S

V