

Site Recovery Manager Administration Guide

vCenter Site Recovery Manager 5.0

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see <http://www.vmware.com/support/pubs>.

EN-000706-03

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 © 2008–2012 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	7	
1	Administering VMware vCenter Site Recovery Manager	9
SRM Deployment	10	
Protected Sites and Recovery Sites	11	
Array-Based Replication	11	
vSphere Replication	12	
About Protection Groups and Recovery Plans	13	
Testing and Running a Recovery Plan	14	
About Reprotect	16	
About Failback	16	
About the Site Recovery Manager Database	16	
SRM and VMware vCenter Server	17	
SRM Licensing	18	
SRM Authentication	18	
Requirements When Using Public Key Certificates	19	
Understanding Roles and Permissions	20	
Assign Roles and Permissions	21	
SRM Roles Reference	21	
SRM Network Ports	23	
Connecting to SRM	24	
Operational Limits of Site Recovery Manager	24	
2	Installing and Updating Site Recovery Manager	27
Configuring the SRM Database	28	
Microsoft SQL Server Configuration	28	
Oracle Server Configuration	28	
DB2 Server Configuration	29	
About the vSphere Replication Management Database	29	
Configure the VRM Database	29	
Install the SRM Server	31	
Upgrading SRM	33	
Prepare for SRM Upgrade	34	
Update the SRM Server	35	
Upgrade the SRM Client Plug-In	36	
Configure the Upgraded SRM Installation	36	
SRM Migration Utility	37	
Install Storage Replication Adapters	38	
Install the SRM Client Plug-In	38	
Connect the Sites	39	
Revert to a Previous Release	40	
Repair or Modify the Installation of a Site Recovery Manager Server	40	

- Install the SRM License Key 42

- 3 Establishing Inventory Mappings and Placeholder Datastores 43**
 - Understanding Placeholder Datastores 43
 - Configure a Placeholder Datastore 44
 - Configure Datastore Mappings for vSphere Replication Management 44
 - Select Inventory Mappings 44

- 4 Configuring Array-Based Protection 47**
 - Configure Array Managers 47
 - Rescan Arrays to Detect Configuration Changes 48
 - Edit Array Managers 48

- 5 Installing vSphere Replication Servers 51**
 - Deploy a vSphere Replication Management Server 52
 - Configure vSphere Replication Management Server Settings 52
 - Configure VRMS Security Settings 53
 - Configure VRMS Network Settings 54
 - Configure VRMS System Settings 54
 - Configure vSphere Replication Management Connections 55
 - Deploy a vSphere Replication Server 55
 - Configure vSphere Replication Server Settings 56
 - Register a vSphere Replication Server 57

- 6 Creating Protection Groups and Replicating Virtual Machines 59**
 - Limitations to Protection and Recovery of Virtual Machines 59
 - Create Array-Based Protection Groups 60
 - Edit Array-Based Protection Groups 61
 - Create vSphere Replication Protection Groups 61
 - Edit vSphere Replication Protection Groups 62
 - Configure Replication for a Single Virtual Machine 62
 - Configure Replication for Multiple Virtual Machines 63
 - Replicate Virtual Machines Using Physical Couriering 64
 - Move a Virtual Machine to a New vSphere Replication Server 66
 - Apply Inventory Mappings to All Members of a Protection Group 66

- 7 Recovery Plans and Re-protection 67**
 - Create a Recovery Plan 67
 - Edit a Recovery Plan 68
 - Remove a Recovery Plan 68
 - Test a Recovery Plan 68
 - Cancel a Test or Recovery 69
 - Run a Recovery Plan 70
 - Understanding Re-protection 71
 - Re-protection Process 72
 - Re-protection State Reference 73

8	Customizing Site Recovery Manager	75
	Customizing a Recovery Plan	75
	Recovery Plan Steps	75
	Customize Recovery Plan Steps	77
	Customize the Recovery of an Individual Virtual Machine	80
	Customize IP Properties For an Individual Virtual Machine	80
	Report IP Address Mappings for a Protection Group	82
	Understanding Customizing IP Properties for Multiple Virtual Machines	82
	Configure Protection for a Virtual Machine or Template	86
	Configure Resource Mappings for a Virtual Machine	87
	Configure SRM Alarms	88
	Working with Advanced Settings	88
	Guest Customization Settings	88
	Change Recovery Site Settings	89
	Change Array-Based Storage Provider Settings	89
	Change Local Site Settings	90
	Change Remote Site Settings	91
	Change Storage Settings	91
	Change Replication Setting	92
	Change vSphere Replication Settings	92
9	Troubleshooting SRM	93
	Events and Alarms	93
	Site Status Events	94
	Protection Group Events	94
	Recovery Events	96
	SNMP Traps	97
	Storage and Storage Provider Events	98
	Licensing Events	101
	Permissions Events	101
	Collecting SRM Log Files	102
	Collect SRM Log Files Using the vSphere Client	102
	Collect SRM Server Log Files	102
	Features Are Unavailable When Deploying VRMS	103
	OVF Package is Invalid and Cannot be Deployed	103
	Connection Errors Between VRMS and SQL Cannot be Resolved	103
	Configuration of the VRMS Database Fails with DB2 Databases	104
	Index	105

About This Book

VMware® vCenter Site Recovery Manager (SRM) is an extension to VMware vCenter that delivers a business continuity and disaster recovery solution that helps you plan, test, and execute the recovery of vCenter virtual machines. SRM can discover and manage replicated datastores, and automate migration of inventory from one vCenter to another.

Intended Audience

This book is intended for Site Recovery Manager administrators who are familiar with vSphere and its replication technologies such as host based replication and replicated datastores. This solution serves the needs of administrators 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.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

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 (SRM) is a business continuity and disaster recovery solution that helps you plan, test, and execute the recovery of vCenter virtual machines between one site (the protected site) and another site (the recovery site).

You can configure SRM to work with several third-party disk replication mechanisms (array based replication) or with VMware vSphere Replication.

Two types of recovery are available.

Planned Migration Planned migration is the orderly decommissioning of virtual machines at the protected site and commissioning of equivalent machines at recovery site. For planned migration to succeed, both sites must be up and fully functioning.

Disaster Recovery Disaster recovery is similar to planned migration except it does not require that both sites be up. During a disaster recovery operation, failure of operations on the protected site are reported but otherwise ignored.

SRM coordinates the recovery process with the underlying replication mechanisms that the virtual machines at the protected site are shut down cleanly (in the event that the protected site virtual machines are still available) and the replicated virtual machines can be powered up. Recovery of protected virtual machines to the recovery site is guided by a recovery plan that specifies the order in which virtual machines are started up. The recovery plan also specifies network parameters, such as IP addresses, and can contain user-specified scripts that can be executed to perform custom recovery actions.

After a recovery has been performed, the running virtual machines are no longer protected. To address this reduced protection, SRM supports a reprotect operation for virtual machines protected on array-based storage. The reprotect operation reverses the roles of the two sites after the original protected site is back up. The site that was formerly the recovery site becomes the protected site and the site that was formerly the protected site becomes the recovery site.

SRM lets you test recovery plans. You can conduct tests using a temporary copy of the replicated data in a way that does not disrupt ongoing operations at either site. You can conduct tests after a reprotect has been done to confirm that the new protected/recovery site configuration is valid.

This chapter includes the following topics:

- [“SRM Deployment,”](#) on page 10
- [“Protected Sites and Recovery Sites,”](#) on page 11
- [“About the Site Recovery Manager Database,”](#) on page 16
- [“SRM and VMware vCenter Server,”](#) on page 17
- [“SRM Licensing,”](#) on page 18
- [“SRM Authentication,”](#) on page 18

- [“Understanding Roles and Permissions,”](#) on page 20
- [“SRM Network Ports,”](#) on page 23
- [“Connecting to SRM,”](#) on page 24
- [“Operational Limits of Site Recovery Manager,”](#) on page 24

SRM Deployment

You must complete several groups of tasks to configure SRM. You complete some tasks in all cases, and you complete some tasks only for vSphere Replication (VR) or for array based replication. If your environment will use both types of replication, consider all tasks, but if not, you might need to complete only a subset of the total possible set of tasks.

The set includes the following tasks:

- 1 Obtain the latest SRM software and any required patches.
- 2 Configure the SRM databases at each site.
- 3 Install SRM at the protected site.
- 4 Install SRM at the recovery site.
- 5 Pair sites.

In the case of using VR, complete the following tasks:

- 1 Deploy a vSphere Replication Management Server (VRMS) at the protected site.
- 2 Deploy VRMS at the recovery site.
- 3 Configure a database for VRMS at both sites.
- 4 Configure both VRMS servers using the Virtual Appliance Management Interface (VAMI).
- 5 Deploy a vSphere Replication Server (VRS) at the recovery site.
- 6 If bi-directional replication is required, deploy a VR server at the protected site.
- 7 Register VRS with VRMS.
- 8 Connect the two VRMS appliances between sites.

In the case of using array based replication, complete the following tasks at both sites:

- 1 Install Storage Replication Adapters (SRAs).
- 2 Configure array managers.

After you establish the required infrastructure for VR, arrays, or both, complete the following steps:

- 1 Configure inventory mappings.
- 2 Configure placeholder datastores.
- 3 If you are using VR, configure datastore mappings.
- 4 Create protection groups.
- 5 Protect virtual machines.
- 6 Create recovery plans.

Protected Sites and Recovery Sites

In a typical SRM installation, the protected site provides business-critical datacenter services. The recovery site is an alternative 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. Conversely, the recovery site can be in the same room as a way of establishing redundancy. 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.
- If you are using array-based replication, identical replication technologies must be available at both sites.
- The recovery site must have hardware, network, and storage resources that can support the same virtual machines and workloads as the protected site.
- The sites should be connected by a reliable IP network. If you are using array-based replication, ensure your network connectivity meets the arrays' network requirements.
- The recovery site should have access to comparable 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. SRM includes a wizard that guides you through the site-pairing process. You must establish a connection between the sites and you must provide authentication information for the two sites so they can exchange information. Site pairing requires vSphere administrative privileges at both sites. To initiate the site-pairing process, you must know the user name and password of a vSphere administrator at each site. If you are using vSphere Replication, pair vSphere Replication Management Servers similarly to how SRM sites are paired.

Array-Based Replication

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

If you plan to use array-based replication with SRM, establish replication before you install and configure SRM.

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/>. VMware does not support storage replication adapters downloaded from other sites. You must install an SRA specific to each array that you use with SRM on the SRM server host. SRM supports using multiple SRAs.

About Bidirectional Operation

You can use a single set of paired SRM sites to protect in both directions. Each site can simultaneously be both a protected site and recovery site but for a different set of virtual machines. This feature is not limited to array-based replication but when you are using array-based replication, any given one of the array's LUNs is only ever replicating in one direction. Two different LUNs in the same array can each be replicating in different directions from each other.

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 datastores are stored.

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 replication source is at the protected site and the replication target is at the recovery site.

A 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 backing the datastore are replicated before it can protect the virtual machines that use that datastore. SRM must ensure that all devices containing protected virtual machine files are replicated. During a recovery or test, SRM must handle all such devices together. To achieve this goal, SRM aggregates datastores into datastore groups to accommodate virtual machines that span multiple datastores. SRM regularly checks that datastore groups contain all necessary datastores to provide protection for appropriate virtual machines. When necessary, datastore groups are recalculated. For example, this may occur when new devices are added to a virtual machine, and those devices are stored on a datastore that was not previously a part of the datastore group.

A datastore group consists of the smallest 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, such as in the case of an MSCS cluster.
- 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. A consistency group is a collection of replicated devices where every state of the target set of devices existed at some point in time as the state of the source set of devices. Informally, the devices are replicated together such that when recovery happens using those devices, software accessing the targets do not see the data in a state it is not prepared to deal with.

vSphere Replication

In vSphere Replication (VR), SRM uses vSphere replication technologies to replicate data to servers at the recovery site.

vSphere Replication uses vSphere Replication Management Server (VRMS) to manage the VR infrastructure. VR requires installing the VR Server (VRS) virtual appliance and VRMS virtual appliance, both of which can be installed with SRM during the installation process. While VR does not require storage arrays, an VR storage replication source and target can be any regular storage device, including, but not limited to, storage arrays.

About Protection Groups and Recovery Plans

A protection group is a collection of virtual machines and templates. A recovery plan specifies how the virtual machines in a specified set of protection groups are recovered. In the case of virtual machines replicated using array-based replication, protection groups are composed of virtual machines that use the same replicated datastore group.

When you create a protection group for array-based replication, you specify array information and SRM computes the set of virtual machines. When you create a protection group for virtual machines that are replicated with vSphere Replication, you can add any virtual machines to the protection group.

With array-based replication, all of the virtual machines and templates on the datastores in the protection group's datastore group are recovered together. When you create a protection group, it initially contains only those virtual machines that store all of their files on one of the datastore groups associated with the protection group. You can add virtual machines to the protection group by creating them on one of the datastores that belong to the datastore groups associated with the protection group. You can also add virtual machines to the protection group by using Storage vMotion to move their storage onto one of the datastores that belong to the datastore groups associated with the protection group. You can remove a member from a protection group by moving the virtual machine's files to another datastore. A protection group can contain one or more datastore groups. However, a datastore group can belong to only one protection group.

Multiple Recovery Plans for the Same Protection Group

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 references one or more protection groups. A protection group can be specified in more than one recovery plan. For example, you can create one recovery plan to handle a planned migration of services from the protected site to the recovery site, and another plan to handle an unplanned event such as a power failure or natural disaster. Having these different recovery plans allow you to decide how recovery occurs.

You can use only one recovery plan at a time to recover a protection group. If multiple recovery plans that specify the same protection group are tested or run simultaneously, only one recovery plan can failover the protection group. Other running recovery plans that specify the same protection group report warnings for that protection group and the virtual machines it contains. These warnings explain that the virtual machines were failed over. Other protection groups that those recovery plans cover are not affected by the warnings.

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 vSphere inventory at the recovery site. Array-based replication requires that each virtual machine 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 protection group. 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 that are not configured or are 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 recovery 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 are 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 settings of the placeholders or move them to a different folder or resource pool.

You can treat placeholders like other members of the recovery site vCenter inventory, although they cannot be powered on. When a placeholder is created, its folder and compute resource assignments are derived from inventory mappings established at the protected site. A recovery site vCenter administrator can modify folder and compute resource assignments as necessary. Changes to a placeholder virtual machine network can be edited only in the inventory mappings. If no mapping for a network exists, the user can specify a new network when protecting the virtual machine. Changes made to the placeholder override settings established during the protection of the virtual machine and are preserved at the recovery site during the test and recovery.

When you recover a protected virtual machine by testing or running a recovery plan, its placeholder is replaced by the recovered virtual machine and powered on as directed by the recovery plan. After a recovery plan test finishes, the placeholders are restored as part of the cleanup process.

Testing and Running a Recovery Plan

Testing a recovery plan exercises nearly every aspect of a recovery plan, though several concessions are made to avoid disrupting ongoing operations. While testing a recovery plan has no lasting effects on either the protected or the recovery site, running a recovery plan has significant effects on both sites.

Run test recoveries as often as needed. Testing a recovery plan does not affect replication or the ongoing operations of either site. Testing a recovery plan might temporarily suspend selected local virtual machines at the recovery site if recoveries are configured to do so. You can cancel a recovery plan test at any time.

In the case of planned migrations, a recovery stops replication after a final synchronization of the source to the target. For disaster recoveries, virtual machines are restored to the most recent available state, as determined by the recovery point objective (RPO). After the final replication is completed, SRM 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.

You need different privileges when 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	Assign the Site Recovery Manager.Recovery Plans.Test permission from the Permissions tab.	Assign the Site Recovery Manager.Recovery Plans.Recovery permission from the Permissions tab.
Effect on virtual machines at protected site	None	Virtual machines are shut down in reverse priority order.

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

	Test a Recovery Plan	Run a Recovery Plan
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 cleaned up.	Local virtual machines are suspended if required by the plan.
Effect on replication	Temporary snapshots of replicated storage are created at the recovery site. For array based replication, the arrays are rescanned to discover them.	In the case of a planned migration, replicated datastores are synchronized, then replication is stopped, and the target devices at the recovery site are made writable. During a disaster recovery, the same steps are attempted, but if they do not succeed, the errors are ignored.
Network	If test networks are explicitly assigned, recovered virtual machines are connected to a test network. If virtual machine network assignment is auto , SRM assigns virtual machines to temporary networks that are not connected to any physical network.	Recovered virtual machines are connected to a datacenter network.
Interruption	Can be canceled.	May be canceled in some cases.

How SRM Interacts with DPM and DRS During Recovery

Distributed Power Management (DPM) is a VMware feature 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. DPM and DRS are not mandatory, but SRM supports both services and enabling them provides certain benefits when using SRM.

SRM temporarily disables DPM for the cluster and ensures that all hosts in it are powered on before recovery begins. 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. SRM registers virtual machines across the available ESX hosts in a round-robin order, to distribute the potential load as evenly as possible. SRM always uses DRS placement to balance the load intelligently across hosts before it powers on recovered virtual machines on the recovery site, even if DRS is disabled on the cluster. If DRS is enabled and in fully automatic mode, DRS might move other virtual machines to further balance the load across the cluster while SRM is powering on the recovered virtual machines, and DRS will continue to balance all virtual machines across the cluster after SRM has powered on the recovered virtual machines.

Test Bubble Networks and Datacenter Networks

SRM can create a test bubble network to which recovered virtual machines are connected during a test. SRM defaults to the Auto setting so that an accidental test recovery does not affect production. 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. These virtual machines 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 customization and recovery plan steps that apply the customization.

Virtual machines that must interact with each other should be failed over to the same test bubble network. For example, if a Web server accesses information on a database, those virtual machines should fail over together to the same network. This step enables testing of the function of the failed over virtual machines.

About Reprotect

With reprotect, you can protect recovered virtual machines after a recovery back to the original protected site, including reversing the direction of replication.

Reprotect uses the protection information that was established before a recovery to reverse the direction of protection. You can complete the Reprotect process after a recovery is finished. If the recovery finishes with errors, you must fix all errors and rerun the recovery, repeating this process until no errors occur.

IMPORTANT Reprotect is supported only for array-based replication. vSphere Replication (VR) reprotect is not supported. If a recovery plan contains VR groups, remove those groups before you run a reprotect operation.

About Failback

A failback is an optional procedure that restores the original configuration of the protected and recovery sites after a recovery. You can configure and run a failback procedure when you are ready to restore services to the protected site.

Failback is a term for a collection of procedures that you can use to restore the original configuration of the protected and recovery sites after a recovery. Anytime errors occur during a failback, you must resolve those errors and repeat the failback until the process completes successfully.

After a recovery has occurred, a failback can be completed. Failbacks have three phases. For example, at the start, A is the protected site and B is the recovery site. A recovery occurs, migrating the virtual machines to site B. You might choose to run a failback. At that point, the following phases occur:

- Perform a reprotect. The former recovery site, B, is made the protected site and information about protection is used to establish protection with A being the recovery site.
- Perform a planned migration. The virtual machines are recovered to site A. To avoid interruptions in virtual machine availability, you may want to run a test before actually completing the planned migration. If errors are identified by the test, these issues can be resolved before actually performing the planned migration.
- Perform a second reprotect, this time with site A being protected with site B as the recovery site.

About the Site Recovery Manager Database

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

The SRM database is a critical part of any SRM installation. The database must be created and a database connection established 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 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. Before you can install SRM, the database must exist .

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

Server Type	The type of database server being used.
DSN	The DSN (database source name) specifies a data structure that contains information about the SRM database that the ODBC driver needs to connect to that data source.

User Name and Password	This authentication information is required so that SRM can use 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. In most cases, it is not necessary to change this setting. Before changing this setting consult with your database administrator.
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. In most cases, it is not necessary to change this setting. Before changing this setting consult with your database administrator.

SRM and VMware vCenter Server

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

- 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:

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

SRM and the vCenter Database

If you update the vCenter installation that SRM extends, do not reinitialize the vCenter database during the update. SRM stores identification information about all vCenter objects in SRM's database. If you reinitialize the vCenter database, the identification data that SRM has stored no longer matches identification information in the new vCenter and objects are not found.

SRM and Other vCenter Server Solutions

You can run other VMware solutions such as vCenter Update Manager, vCenter Server Heartbeat, VMware Fault Tolerance, and vCenter CapacityIQ in deployments that you protect using SRM. However, use caution before connecting other VMware solutions to the vCenter Server instance to which the SRM server is connected. Connecting other VMware solutions to the same vCenter Server instance as SRM might cause problems when you upgrade SRM or vSphere. Check the compatibility and interoperability of these solutions with SRM before you deploy them.

SRM Licensing

The SRM server requires a license key to operate. You install each SRM server with an evaluation license that is valid for 60 days and supports protecting up to 75 virtual machines.

SRM uses the vSphere licensing infrastructure for license management. Additionally, vSphere needs to be licensed sufficiently for SRM to protect and recover virtual machines.

After the evaluation license expires, existing protection groups remain protected and can be recovered, but you cannot create new protection groups or modify existing ones until you obtain and assign a valid SRM license key. VMware recommends that you obtain and assign SRM license keys as soon as possible after installing SRM. You can obtain a license key from your VMware sales representative.

How License Keys Apply to Protected and Recovery Sites

SRM requires a license key that specifies the maximum number of protected virtual machines at a site. Larger licenses are often required when protecting large numbers of virtual machines.

- Install keys at one site to enable failover.
- Install keys at both sites to enable bidirectional operation including reprotection.

If your SRM Servers are connected with linked vCenter Servers, the SRM servers can share the same license key.

To obtain your license keys, go to the VMware Product Licensing Center (<http://www.vmware.com/support/licensing/index.html>).

SRM licensing checks for a valid license whenever you add a virtual machine to or remove a virtual machine from a protection group. If licenses are not in compliance, vSphere triggers a licensing alarm. VMware recommends that you configure alerts for triggered licensing events so that licensing administrators are notified by email.

SRM Authentication

All communications between SRM and vCenter Servers take place over a SSL connections 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 between SRM servers at different sites and between SRM and vCenter. By default, SRM uses credential-based authentication, but certificate-based authentication can alternatively be selected. 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 19.

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. 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. A warning allows you to verify the thumbprint of the certificate used by the other server and confirm its identity. 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 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 specific criteria.

While SRM uses standard PKCS#12 certificate for authentication, it places a few specific requirements on the contents of certain fields 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 the following components.
 - 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.) Because this name is subject to a case-sensitive comparison, use lowercase letters when specifying the name during SRM installation.
 - 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 `extendedKeyUsage` or `enhancedKeyUsage` 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
```
- The SRM certificate password must not exceed 31 characters.

Understanding Roles and Permissions

SRM provides disaster recovery by performing operations on behalf of users. These operations involve managing objects, such as recovery plans or protection groups, and performing operations, such as replicating or powering off virtual machines. SRM must be able to complete these tasks, when appropriate, and refuse to complete operations when they are not authorized. To achieve this goal, SRM uses permissions and roles.

the following are key terms related to permissions and roles.

Privilege	The right to perform an action. Examples of privileges include creating a recovery plan or modifying a protection group.
Role	A collection of privileges. Default roles are designed to provide the privileges associated with some user role such as users who will manage protection groups or complete recoveries.
Permissions	<p>A role granted to a particular user or group (also known as a principal) on some object. A permission is the intersection of role, object, and principal.</p> <p>A permission is the intersection of a privilege and an object. For example, the privilege to modify a protection group as it applies to a specific protection group in the inventory.</p>

SRM determines if the operation is permitted when protection is configured, rather than at the time the operation is to be completed. After SRM verifies that the appropriate permissions are assigned on vSphere resources, future actions are carried out on behalf of users by SRM using the vSphere administrator context.

For configuration operations, user permissions are validated when the operation is requested. Other operations require two phases of validation.

- 1 During configuration, SRM verifies that the user configuring the system has the required permissions to complete the configuration on the vCenter object. For example, a user must have permission to protect a virtual machine and use resources on a secondary vCenter Server that the recovered virtual machine would use.
- 2 The user executing the configuration must have permissions to complete the task. For example, a user must have permissions to execute a recovery plan. The task is then completed in the administrative context.

As a result, a user who completes a particular task, such as a failover, does not have to have permissions to act on vSphere resources. The action is authorized by the role, but is completed by SRM acting as an administrator. These operations are carried out using the administrator credentials provided during site pairing.

SRM maintains a database of permissions for internal SRM objects using a model similar to the one used by vCenter Servers. SRM verifies its own SRM privileges even on vCenter objects. For example, SRM checks for Recovery Use permission on the target datastore rather than multiple low-level permissions, such as Allocate space.

Assign Roles and Permissions

Permission assignments apply on a per-site basis. After installation only vCenter Administrators can log into SRM. To allow other users access, vCenter Administrators must grant them permissions in the SRM UI. You must add corresponding permission on both sites.

SRM requires permissions on vCenter objects as well as SRM objects. To configure permissions on the remote vCenter installation, start another instance of vSphere Client. You can change SRM permissions from the same UI on both sites after pairing. SRM augments vCenter roles and permissions with additional ones that allow detailed 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.

Procedure

- 1 Click **Sites**, and select the site for which you want to assign permissions.
- 2 Click the **Permissions** tab.
- 3 Right-click one of the items and click **Add Permission**.
- 4 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 Select **Propagate to Child Objects** to apply the selected role to all child objects of the selected inventory object.
- 6 Click the **Add** button.
- 7 Identify a user or group for the role.
 - 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 click **OK**.
- 8 Click **OK** to finish the task.

The list of permissions references all users and groups that 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.

SRM Roles Reference

SRM includes a set of roles. Each role is assigned a set of privileges, which enable the completion of actions.

Roles may have overlapping sets of privileges and actions. For example, both the SRM Administrator role and the SRM Protection Groups Administrator have the Create privilege for protection groups for Site Recovery Manager. This privilege enables them to complete one aspect of the set of tasks that make up managing protection groups.

The complete list of roles, the privileges granted to those roles, and the actions associated with those privileges are described in the following table.

Table 1-2. SRM Roles

Role	Privilege	Action
SRM Administrator	Site Recovery Manager > Advanced Settings > Modify	Configure advanced settings
	Site Recovery Manager > Array Manager > Configure	Configure connections
	Site Recovery Manager > Diagnostics > Export	Configure inventory preferences
	Site Recovery Manager > Inventory Preferences > Modify	Configure placeholder datastores
	Site Recovery Manager > Placeholder Datastores > Configure	Configure array managers
	Site Recovery Manager > Protection Group > Assign to Plan	Manage protection groups
	Site Recovery Manager > Protection Group > Create	Manage recovery plans
	Site Recovery Manager > Protection Group > Modify	Protect virtual machines
	Site Recovery Manager > Protection Group > Remove	Edit protection groups
	Site Recovery Manager > Protection Group > Remove from Plan	Remove protection groups
	Site Recovery Manager > Recovery History > View Deleted	
	Site Recovery Manager > Recovery History > Plans	
	Site Recovery Manager > Recovery Plan > Configure	
	Site Recovery Manager > Recovery Plan > commands	
	Site Recovery Manager > Recovery Plan > Create	
	Site Recovery Manager > Recovery Plan > Modify	
	Site Recovery Manager > Recovery Plan > Remove	
	Site Recovery Manager > Recovery Plan > Reprotect	
	Site Recovery Manager > Recovery Plan > Test	
	Site Recovery Manager > Remote Site > Modify	
	Virtual Machine > Replication > Protect	
	Virtual Machine > Replication > Stop	
	SRM Protection Groups Administrator	Site Recovery Manager > Protection Group > Create
Site Recovery Manager > Protection Group > Modify		Protect virtual machines
Site Recovery Manager > Protection Group > Remove		
Virtual Machine > Replication > Protect		
Virtual Machine > Replication > Stop		

Table 1-2. SRM Roles (Continued)

Role	Privilege	Action
SRM Recovery Plans Administrator	Site Recovery Manager > Protection Group > Assign to Plan Site Recovery Manager > Protection Group > Remove from Plan Site Recovery Manager > Recovery Plan > Configure Site Recovery Manager > Recovery Plan > Commands Site Recovery Manager > Recovery Plan > Create Site Recovery Manager > Recovery Plan > Modify Site Recovery Manager > Recovery Plan > Remove Site Recovery Manager > Recovery Plan > Test Resource > Recovery Use	Manage recovery plans Add protection groups to recovery plans Test recovery plans Cancel recovery plan test Edit virtual machine recovery properties
SRM Test Administrator	Site Recovery Manager > Recovery plan > Modify Site Recovery Manager > Recovery plan > Test	Test recovery plans Cancel recovery plans test Edit virtual machine recovery properties

SRM 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. Beyond these standard ports, you must also ensure the following requirements.

- vSphere Replication (VR) Servers and SRM servers have NFC traffic access to target ESX servers.
- Any network requirements of your particular array-based replication are met.

Table 1-3. SRM Network Ports

Default Port	Protocol or Description	Endpoints or Consumers
80		All traffic to SRM servers through the vCenter Server proxy.
8095	SOAP	SRM server and vCenter Server (intrasite only). This port must be accessible from the vCenter Server proxy system.
9085	HTTP	vCenter Server (for plug-in download). This port must be accessible from the vCenter Server proxy system.
9007	SOAP	Used by external API clients for task automation.

Table 1-4. How VRMS Uses Network Ports

Default Port	Protocol or Description	Endpoints or Consumers
80		All traffic to SRM servers through the vCenter Server proxy.
8043	SOAP	VRM and the vCenter Server proxy.
8080	VAMI Web UI	Administrator's web browser.

Table 1-5. How VR Servers Use Network Ports

Default Port	Protocol or Description	Endpoints or Consumers
8123	SOAP	Management traffic used by VRMS to manage the VR Servers.
5480	VAMI Web UI	Administrator's web browser.
31031	Initial replication traffic	From the ESX host at the protected site to the VR appliance at recovery site.
44046	Ongoing replication traffic	From the ESX host at the protected site to the VR appliance at recovery site.

Connecting to SRM

Use the vSphere Client to connect to and manage SRM. Ensure you connect using an account that has been paired with a role that has the necessary permissions.

SRM does not require that you connect to a specific SRM site in an SRM deployment. Changes can be made to the protected and recovery sites by connecting to a vCenter Server at either site. Completing administrative tasks on a SRM deployment begins with the following steps:

Procedure

- 1 Open a vSphere Client and connect to the vCenter Server for either the protected or recovery site.
Log in using an account that has been granted the permissions required to complete the desired task.
 - If sites are not paired, you must select the system to which to connect.
 - If sites are paired, you must provide the user name and password for both sites.
- 2 On the vSphere Client Home page, click the **Site Recovery** icon.

Once you have clicked the **Site Recovery** icon, complete the steps prescribed for the particular administrative task.

SRM adds several roles, each of which include permissions for completing SRM tasks. You can pair users with these particular roles, enabling them to complete tasks. For more information about the roles that SRM adds and the privileges required to complete tasks, see [“SRM Roles Reference,”](#) on page 21.

Operational Limits of Site Recovery Manager

Each SRM server can support up to a certain number of virtual machines, protection groups, datastore groups, vSphere Replication Management Server (VRMS) instances per host, and vSphere Replication Servers (VRS) servers per VRMS.

You can run one SRM Server per vCenter Server instance.

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

Table 1-6. SRM Protection Limits for Array Based Protection

Item	Maximum
Protected virtual machines per protection group	500
Protected virtual machines	1000
Protection groups per recovery plan	150
Datastore groups	150
Concurrent recoveries	10

Table 1-7. SRM Protection Limits for vSphere Replication Protection

Item	Maximum
Protected virtual machines per protection group	500
Protected virtual machines	500
Protection groups per recovery plan	250
Datastore groups	250
Concurrent recoveries	10

Table 1-8. SRM Deployment Limits for vSphere Replication

Item	Maximum
vSphere Replication Management Server (VRMS) appliances per vCenter Server instance	1
vSphere Replication Server (VRS) appliances registered to a VRMS appliance	5
Virtual machine replication schedules per VRS appliance	100

VMware recommends that you never use SRM to protect Active Directory domain controllers. Active Directory provides its own replication technology and restore mode, and these technologies can be used to handle disaster recovery situations.

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 SRM client plug-in from either SRM server using the Manage Plugins menu from your vSphere Client. You use the SRM client plug-in to configure and manage SRM at each site.

Prerequisites

SRM requires that a vCenter Server be installed at each site prior to installing SRM. The SRM installer must be able to connect with this server during installation. VMware recommends installing SRM on a system that is different from the system where vCenter Server is installed. If SRM and vCenter Server are installed on the same system, administrative tasks might become more difficult to perform. If you are upgrading SRM, only protection groups and recovery plans that are in a valid state are saved during the upgrade. Protection groups or recovery plans that are in an invalid state are discarded.

The system on which SRM is installed has the following hardware requirements:

- Processor – 2.0GHz or higher Intel or AMD x86 processor
- Memory – 2GB minimum
- Disk Storage – 5GB minimum
- Networking – Gigabit recommended

For current 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,”](#) on page 28
- [“About the vSphere Replication Management Database,”](#) on page 29
- [“Install the SRM Server,”](#) on page 31
- [“Upgrading SRM,”](#) on page 33
- [“Install Storage Replication Adapters,”](#) on page 38
- [“Install the SRM Client Plug-In,”](#) on page 38
- [“Connect the Sites,”](#) on page 39
- [“Revert to a Previous Release,”](#) on page 40
- [“Repair or Modify the Installation of a Site Recovery Manager Server,”](#) on page 40
- [“Install the SRM License Key,”](#) on page 42

Configuring the SRM Database

Each 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. Before attempting an SRM environment upgrade, ensure both SRM server databases are backed up. This helps ensure you can revert after the upgrade, if required.

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.

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

Microsoft SQL Server Configuration

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

SRM requires that the Microsoft SQL Server must have a 32-bit DSN. Microsoft SQL Server has the following configuration requirements when used as the SRM database.

- The database schema has three requirements:
 - It must be owned by the SRM database user (the database user name you supply when configuring the SRM database connection).
 - It must be the default schema for the SRM database user.
 - The DB schema name must be the same as the DB user name.
- You must grant the SRM database user the following permissions:
 - bulk insert
 - connect
 - create table
- If you are using Windows authentication, the database user account must be the same user account that you use to run the SRM service.
- If you are using SQL Authentication, you can leave the default local System user.
- 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.

- When creating the database instance, specify utf-8 encoding.
- You must grant the following permissions to the SRM database user (the database user name you supply when configuring the SRM database connection):
 - connect
 - resource

- create session

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, specify the database owner as a domain account.

About the vSphere Replication Management Database

To use vSphere Replication with SRM, you will need vSphere Replication Management (VRM) Servers. Each VRM Server requires its own database, separate from the SRM database.

The vSphere Replication Management (VRM) database is a critical part of any VRM installation. The database schema must be created before you can install VRM.

VRM cannot use the vCenter database because it has different database schema requirements, although you can use the vCenter database server to create and support the VRM database. Each VRM site requires its own instance of the VRM database. The database must exist before you can install VRM . If the VRM database at either site becomes corrupted, the VRM servers at both sites will shut down.

Configure the VRM Database

You must configure the vSphere Replication Management (VRM) database to enable VRM and vSphere Replication (VR).

To configure the common VRM Server (VRMS) database properties you have to go to the VRMS VAMI (Virtual Appliance Management Interface) by navigating your browser to the VRMS URL and port number, which is 8080. Alternately you can navigate to the VRMS VAMI by clicking on the **Configure VRM Server** link available in the SRM UI. You must configure each VRMS site separately. If you reinitialize the database after you deploy VRMS, you must go to the VRMS VAMI to re-setup the VRM to use the new database connection.

Common Database Configurations

Common database properties include the following.

- DB Type: Choose one of the supported database types.
 - Microsoft SQL Server
 - Oracle Server
 - DB2 Server
- DB Host: The database server URL.
- DB Port: When you select your database type a default port value will be suggested. You can keep it or change it to match your database server configuration.
- DB Username: The VRM database user.
- DB Password: The VRM database user password.
- DB Name: The VRM database schema name. Create the schema name in advance.
- DB URL: This URL is auto-generated and hidden by default. Advanced users might want to fine-tune other database properties.

Microsoft SQL Server Configuration

A Microsoft SQL Server configuration must meet specific requirements to support VRM. You must configure these requirements in Microsoft SQL Server.

- You can use either a named instance or the default instance of SQL Server.
- Enable TCP on the database instance.
- Use a static TCP port, for example set to the default of 1433. Alternatively, to use dynamic TCP ports, you must perform additional configuration.
 - Use a named instance of SQL Server rather than the default instance. You can only use dynamic ports with a named instance of SQL Server.
 - In the DB URL in the VRMS configuration interface, replace `port=port_number` with `instanceName=instance_name`.
 - Verify that the SQL Server Browser service is running.
 - The SQL Server Browser runs on port 1434. Use the PortQuery tool from a remote machine to check that the port on which the SQL Server Browser service runs is not blocked by a firewall.

```
PortQry.exe -n Machine_Name -p UDP -e 1434
```

- Because the VRMS and the database server run on different hosts, you must use mixed mode authentication and not Windows Authentication.
- The VRM database requires a security login with SQL Server Authentication.
- The VRM database login must be the database owner.
- Because it is the database owner, the login maps to the database user dbo and uses the dbo schema. Keep the dbo user and dbo schema settings.
- The VRM database user must have database administrator privileges.
- The VRM database user must have the following permissions:
 - bulk insert
 - connect
 - create table
 - create view

Oracle Server Configuration

An Oracle Server configuration must meet specific requirements to support VRM. Oracle Server has the following configuration requirements when used as the VRM database.

- When creating the database instance, specify utf-8 encoding.
- The VRM 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 VRM. DB2 Server has the following configuration requirements when used as the VRM database.

- When creating the database instance, specify utf-8 encoding.
- DB2 uses Windows authentication, so you must specify the database owner as a domain account.
- VRMS uses temporary tables, so you must verify that the user account that you use to log in to the VRMS database can create temporary tables. See [“Configuration of the VRMS Database Fails with DB2 Databases,”](#) on page 104.

Install the SRM Server

You must install an SRM server at the protected site and at the recovery site.

NOTE If you are upgrading an existing SRM installation, see [“Update the SRM Server,”](#) on page 35.

Prerequisites

Configure and start the SRM database service before you install the SRM server. See [“About the Site Recovery Manager Database,”](#) on page 16.

Provide a 32-bit System DSN. For information about creating a 32-bit DSN on a 64-bit system, see <http://kb.vmware.com/kb/1010401>.

Verify that you have the following information:

- The fully qualified domain name (FQDN) or IP address of the site’s vCenter server. The server must be running and accessible during SRM installation. Use the same type of addressing in all cases. Using FQDNs is preferred, but if that is not universally possible, use IP addresses for all cases. See [“Requirements When Using Public Key Certificates,”](#) on page 19.
- 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 28.
- If you are using certificate-based authentication, the pathname to an appropriate certificate file. See [“SRM Authentication,”](#) on page 18.

Procedure

- 1 Log in to the machine on which you are installing SRM.
Use an account with sufficient privileges. This account is often an Active Directory domain administrator, but can also be a local administrator.
- 2 Download the SRM installation file to a folder on the machine, or open a folder on the network that contains this file.
For faster starting and installing, copy the installation file to a local temporary folder.
- 3 Double-click the SRM installer icon.
If the installer detects an existing installation, verify that you want to update the existing installation. If you want to update an existing installation, stop installing the SRM server.
- 4 Click **Next**.
- 5 Select **I agree to the terms in the license agreement** and click **Next**.

- 6 Select the folder in which to install SRM and click **Next**.

The default installation folder for a new installation of SRM is C:\Program Files (x86)\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 (Optional) Select whether or not to install vSphere Replication functions.

Installing vSphere Replication enables additional functions. You can install vSphere Replication later using the Repair Installation option.

- 8 Enter information about the vCenter server at the site where you are installing SRM and click **Next**.

Option	Action
vCenter Server Address	Type the host name or IP address of the vCenter server. Enter the host name in lowercase letters. After installation is complete and you are configuring the connection between the protected and recovery sites, supply this host name or IP address exactly as you enter it here, because is subject to case-sensitive comparisons.
vCenter Server Port	Accept the default or enter a different port.
vCenter Server Username	Type the user name of an administrator of the specified vCenter server.
vCenter Server Password	Type the password for the specified user name. The password cannot be empty.

The installer contacts the specified vCenter server and validates the information you supplied.

- 9 Select an authentication method.

- To use credential-based authentication, select **Automatically generate certificate** and click **Next**. Type 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**. Type the path to the certificate file. The certificate file must contain exactly one certificate with exactly one private key matching the certificate. Type the certificate password.

- 10 Type the administrator and host configuration information.

Option	Description
Local Site Name	A name for this installation of SRM. A suggested name is generated, but you can type any name. It cannot be the same name that you use for another SRM installation with which this one will be paired.
Administrator E-mail	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	Name or IP address of the local host. This value is obtained by the SRM installer and needs to be changed only if it is incorrect. For example, the local host might have more than one network interface and the one detected by the SRM installer is not the one you want to use.
Listener Ports	SOAP and HTTP port numbers to use.
API Listener Port	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.

- 11 Type the SRM database configuration information and click **Next**.

Option	Action
Database Client	Select a database client type from the drop-down menu.
Data Source Name	Select an existing DSN from the drop-down menu. You can also click ODBC DSN Setup to view existing DSNs or create a system DSN.
Username	A user ID valid for the specified database.
Password	Password for the specified user ID.
Connection Count	Initial connection pool size.
Max Connections	Maximum number of database connections that can be open simultaneously.

- 12 Click **Install**.
- 13 When the installation is finished, click **Finish**.

What to do next

You can install SRAs at each site for array-based replication or install and configure vSphere Replication Management Servers for vSphere Replication. See [“Install Storage Replication Adapters,”](#) on page 38 or [“Deploy a vSphere Replication Management Server,”](#) on page 52.

Upgrading SRM

You can use the SRM 5.0 installer to update existing installations, while preserving work done to configure site protection.

SRM 5.0 supports upgrading existing SRM 4.1 and 4.1.1 deployments. Versions of SRM from before 4.1 must be upgraded to version 4.1 or 4.1.1 before upgrading to 5.0. SRM 4.1, 4.1.1, and 5.0 run only on 64-bit operating systems, but previous versions of SRM may be installed on 32-bit operating systems. Keep this in mind when you consider upgrading from versions of SRM from 4.0 and earlier.

For the supported upgrade paths for SRM 5.0 update releases, see the release notes for those update releases.

The SRM upgrade process preserves existing information about SRM configurations.

The SRM upgrade progresses through the following phases:

- Prepare for Upgrade
- Upgrade SRM Servers
- Upgrade the vSphere Client
- Configure the installation

To return to 4.1 or 4.1.1 after the SRM Server upgrade has completed, restore the database and re-install the SRM 4.1 or 4.1.1 server.

Information Supported for Transfer

Information from existing installations is preserved during upgrade to SRM 5.0. This includes the following:

- Datastore groups
- Protection groups
- Inventory mappings
- Recovery plans
- IP customizations for individual virtual machines
- Custom roles and their memberships

- SRM object permissions in vSphere
- Custom alarms and alarm actions
- Test plan histories
- Security certificates
- Mass IP customization files (CSVs)

New Behavior for 5.0

Because of the significant changes introduced in SRM 5.0, including modifications to the database schemas, an upgrade while the system is running is not possible.

As a result, the upgrade process involves taking an existing installation offline, upgrading the SRM server and vSphere Clients, and then migrating existing 4.1 inventory information to the SRM database.

Upgrade Workflow

The steps required to complete an upgrade are as follows:

- 1 Back up the SRM database at the protected site.
- 2 Upgrade vCenter Server to version 5.0 at the protected site.
- 3 Upgrade SRM server to version 5.0 at the protected site. Keep the Storage Report displayed by the installer.
- 4 Upgrade the database, vCenter Server, and SRM server on the recovery site.
- 5 At each site, install version 2.0 SRAs (compatible with SRM 5.0).
- 6 Upgrade the SRM plug-in used by the vSphere Client.
- 7 Pair the sites.
- 8 At each site, configure the array manager (using the Storage Report that was generated when the SRM 5.0 installer finished).
- 9 At only one site, run the `srm-migration.exe` utility using the `vmware-dr.xml` configuration file to import SRM 4.1 inventory.

As the `srm-migration.exe` utility completes its tasks, the exported XML data files are imported into the SRM installation.

Prepare for SRM Upgrade

You must complete this process for all SRM sites and assumes all SRM servers are version 4.1 or later.

As part of the upgrade process, vCenter Servers that support SRM servers must be upgraded to version 5.0, as well. For details about upgrading vCenter Servers, see the vSphere Upgrade Guide. Three upgrade approaches are available.

Table 2-1. vCenter Upgrade Paths for SRM

Upgrade Type	Description	Supported
In-place upgrade	The simplest upgrade path. This path involve upgrading those vCenter Servers associated with SRM before upgrading SRM.	Yes
Upgrade with migration	If you need to migrate an SRM server to a different host or virtual machine as part of the SRM upgrade, VMware recommends that you first uninstall the existing SRM 4.1 server (keeping the database contents) and then run the SRM 5.0 installer on the new host or virtual machine, using the existing database.	Yes
New installation with migration	New installations of vCenter Server are established and SRM is migrated to these new servers.	No. SRM cannot be migrated to a new installation of vCenter Servers. Unique object identifiers on the vCenter Server are not available if a new vCenter Server installation is used.

Prerequisites

To upgrade servers that are running older versions of SRM, upgrade to 4.1 .

Procedure

- 1 Log in to the SRM 4.1 machine.
- 2 Back up the site's database contents using tools provided with the database.
- 3 If vCenter Servers are not upgraded to 5.0, upgrade them now.

What to do next

Complete this process for all vCenter Servers. After upgrading all vCenter Servers, you are ready to upgrade the SRM Servers, as described in [“Update the SRM Server,”](#) on page 35.

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 (keeping the database) and then install the new release.

Prerequisites

Back up your current SRM database at both sites.

Procedure

- 1 Log in to the virtual machine on which you are installing SRM.
Log in using an account with sufficient privileges. This is often an Active Directory domain administrator, but may also be 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.

For faster starting and installing, copy the installation file to a local temporary folder.

- 3 Double-click the SRM installer icon to begin installation.
- 4 When prompted to verify that you want to update the existing installation, click **Yes**.
- 5 Click **Next**.
- 6 Click **Yes** to confirm that you backed up the database.

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.

- 7 When the wizard finishes, click **Finish**.

The installer completes the following tasks:

- The database schema is upgraded to support SRM 5.0.
- Setup data is exported to an XML file including information about all protection groups, recovery plans, and inventory mappings.
- A storage report is produced. Save or print the storage report. This information is required to complete the process of installing the SRAs in the [“Configure the Upgraded SRM Installation,”](#) on page 36.

What to do next

Complete this process for all sites and then upgrade vSphere Client plug-ins, as described in [“Upgrade the SRM Client Plug-In,”](#) on page 36.

Upgrade the SRM Client Plug-In

You must upgrade the client plug-in for all vSphere Clients that are used to manage SRM.

Prerequisites

Verify that vCenter Server 5.0 is available.

Procedure

- 1 Log in to the machine with vSphere Client installed.
- 2 Uninstall the SRM 4.1 plug-in, if installed.
- 3 Download the vSphere Client from vCenter Server and install or upgrade vSphere Client 5.0.
- 4 Connect to either site's vCenter Server.
- 5 Download, install, and start the SRM plug-in for vSphere Client.

What to do next

As necessary, complete this process for other vSphere Clients. After upgrading vSphere Clients, complete the tasks described in [“Configure the Upgraded SRM Installation,”](#) on page 36.

Configure the Upgraded SRM Installation

You must configure the upgraded components to establish a working SRM installation.

Prerequisites

Verify that you upgraded the vCenter Servers, SRM Servers, and vSphere Clients.

Procedure

- 1 Use the vSphere Client to pair the SRM servers.
- 2 Configure SRM storage.

Use the storage report generated by the installer for required information.

 - a Refresh SRAs on both the local and remote site.

SRM 5.0 requires different versions of SRAs from previous versions such as SRM 4.1.
 - b Create array managers for the SRAs.
 - c Enable replicated array pairs.
 - d Verify datastores are replicated as expected.
- 3 Use the `srm-migration.exe` utility to import the configuration that was exported to an XML file by the installer as described in [“Update the SRM Server,”](#) on page 35.

SRM Migration Utility

Use the SRM Migration utility (`srm-migration.exe`) to import SRM 4.1 inventory and, optionally, upgrade SRM 4.1 Mass IP Customization files (CSVs).

The SRM Migration tool performs the following tasks:

- Connect to local and remote vCenter and SRM servers.
- Verify SRM server pairing.
- Create inventory mappings.
- Create protection groups at the local site and protect virtual machines. This process re-uses placeholder virtual machines created with SRM 4.1.
- Create protection groups at the remote site. This task is applicable to bi-directional setup only.
- Create recovery plans at the remote site and protect virtual machines. This process re-uses placeholder virtual machines created with SRM 4.1.
- Create recovery plans at the local site. This task is applicable to bi-directional setup only.
- Link newly created recovery plans with the recovery plan history preserved in the database.

You can display the options for the `srm-migration` utility by running the command with no arguments. The help included with command line utility provides the following information.

```
srm-migration.exe takes the following arguments:
-cmd <command> : importConfig
-cfg <file path> : Full path to SRM server XML config file vmware-dr.xml
-lcl-csv-file <file path> : (Optional) File containing SRM 4.1 IP settings for Vms recovered at
the local site.
-lcl-csv-file <file path> : (Optional) File containing SRM 4.1 IP settings for Vms recovered at
the local site.
-lcl-usr <username> : local vCenter server username
-rem-usr <username> : remote vCenter server username
```

Install Storage Replication Adapters

If you are using array-based replication, you must install a Storage Replication Adapter (SRA). An 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 SRA on the SRM server hosts at the protected and recovery sites. If you are using vSphere Replication (VR), an SRA is not required.

Prerequisites

- You must read the documentation provided with your SRA. SRAs do not support all features that storage arrays support. The documentation included with your SRA normally details what is supported and required. For example, HP and EMC have very detailed physical requirements which must be met for the SRA to perform as expected.
- SRM server installation creates a directory in which you can install the SRAs. This directory is created in the SRM installation folder under `\storage\sra`. 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. For the latest information on such requirements, review the release notes and readme files for the SRAs you are installing.
- The storage array's capability to create snapshot copies of the replicated devices must be enabled. Refer to your array vendor's SRA documentation for details.

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. To properly install a vendor's SRA, install the same version of the SRA at both sites. Do not mix SRA versions.

- 3 Using the vSphere Client, connect to SRM, select **Array Managers** in the left pane, click the **SRAs** tab, and click **Reload SRAs**. This refreshes SRA information, allowing the discovery of the newly installed SRA.

What to do next

[“Configure Array Managers,”](#) on page 47 using the wizard to configure array managers so replication can occur as desired.

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. 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

Verify that the Site Recovery Manager server is 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 Select **Plugins > Manage Plugins**.
- 3 In the Available Plug-ins area, right-click the VMware vCenter Site Recovery plug-in and click **Download and Install**.
- 4 After the download finishes, click **Next** to start the wizard.
- 5 Click **I accept the terms in the license agreement**, and click **Next**.
- 6 Click **Install**.
- 7 When the installation finishes, click **Finish**.

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

Connect the Sites

Before you can use SRM, you must designate and connect sites. The sites must authenticate with each other. This is known as site pairing.

Prerequisites

Before performing any configuration activity such as site pairing, ensure that you have installed an SRM server at both sites, and that you have installed the SRM plug-in at a vSphere client from which you want to administer SRM.

NOTE If you are using credential-based authentication or if you are using an untrusted certificate, several of the steps in this procedure produce certificate warnings.

Procedure

- 1 Click **Sites** in the left pane and click **Configure Connection**.
- 2 On the Remote Site Information page, type the IP address or host name of the vCenter server at the recovery site and the port to which to connect 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 at the recovery site. 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 for subsequent connections.

- 3 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 at the recovery site. The wizard automatically completes several steps to establish the connection.

- 4 In the Remote vCenter Server window, enter credentials for the remote vCenter Server that is managing the SRM server at the other site.
- 5 On the Complete Connections page, click **Finish** after all of the site pairing steps have completed successfully.

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

Verify 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.
The database must be restored to both sites so they are synchronized.
- 4 Install the previous release of SRM.

Repair or Modify the Installation of a Site Recovery Manager Server

To repair an SRM server installation or to change the information that you supplied when you installed the SRM Server, you can make changes by running the SRM installer in repair or modify mode.

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, and so on. The SRM installer supports a modify mode that allows you to change the following values:

- The user name and password of the vCenter administrator
- The username, password, and connection numbers for the SRM database
- The type of authentication (certificate-based or credential-based), the authentication details, or both

The installer's modify mode presents modified versions of some of the pages that are part of the SRM server installation. For information about the repair options, see [“Install the SRM Server,”](#) on page 31.

Running the installer in repair mode fixes missing or corrupted files, shortcuts, and registry entries.



CAUTION If you activated automatic generation of SSL certificates on the SRM server, running the repair mode of the SRM installer updates the certificate thumbprint. Updating the thumbprint can affect the connection between the protected site and the recovery site. Do not run the SRM installer in repair mode on the protected site and on the recovery site simultaneously. Check the connection between the protected site and the recovery site after you run the installer in repair mode. For information about how to configure the connection between the protected site and the recovery site, see [“Connect the Sites,”](#) on page 39.

Prerequisites

Verify that you have administrator privileges on the SRM server or that you are a member of the Administrators group. If you are a member of the Administrators group but you are not an administrator, disable Windows User Account Control (UAC) before you attempt the change operation.

Procedure

- 1 Log in to the SRM server host.
- 2 Open the Windows Add or Remove Software tool.
- 3 Navigate to the entry for VMware vCenter Site Recovery Manager and click **Change** to start the installer in repair mode.
- 4 Click **Next**.
- 5 Select **Modify** or **Repair** and click **Next**.

Option	Description
Repair	Fixes issues in the SRM Server installation. Selecting Repair runs the installer in repair mode without any user input. Go to Step 10 .
Modify	Allows you to adjust the settings that you configured when you installed SRM Server. Selecting Modify opens a wizards that allows you to modify the installation settings. Go to Step 6 .

- 6 Type the vCenter Server username and password.

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.

- 7 Select an authentication method and click **Next**.

Option	Description
Leave the current authentication method unchanged	Select Use existing certificate . If the installed certificate is not valid, this option is unavailable.
Use credential-based authentication	Select Automatically generate certificate .
Use certificate-based authentication,	Select Use a PKCS #12 certificate file .

Unless you select **Use existing certificate**, you are prompted to supply additional authentication details such as certificate location or strings to use for Organization and Organizational Unit.

- 8 Provide the required database configuration information and click **Next**.

Option	Description
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.

- 9 Select **Use existing database** or **Recreate the database** and click **Next**.

Option	Description
Use existing database	Preserves the contents of the existing database.
Recreate the database	Overwrites the existing database and deletes its contents.

- 10 Click **Install** to repair or modify the installation.

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

NOTE If the SRM server fails to restart, check that the vCenter Server and SRM database server are running. Check also that the Server and Workstation Windows services are running on the SRM server host.

- 11 When the repair operation is finished and the SRM server restarts, log in to the SRM interface in the vSphere client to check the status of the connection between the protected site and the recovery site.
- 12 (Optional) If the connection between the protected site and the recovery site is broken, reconfigure the connection, starting from the SRM server that you updated.

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 you install SRM.

SRM uses the vSphere licensing infrastructure for license management. Additionally, vSphere itself needs to be licensed sufficiently for SRM to protect and recover virtual machines.

Procedure

- 1 Open a vSphere client and connect to a vCenter Server on which SRM is deployed.
- 2 On the vSphere client Home page, click **Licensing**.
- 3 For the report view, select **Asset**.
- 4 Right-click an SRM asset and select **Change license key**.
- 5 Select **Assign a new license key** and click **Enter Key**.
- 6 Enter the license key, type an optional label for the key, and click **OK**.
- 7 Click **OK**.

What to do next

Repeat the process so license keys are assigned to all appropriate sites.

Establishing Inventory Mappings and Placeholder Datastores

3

The protection provided by SRM is supported by placeholder datastores, datastore mappings, and inventory mappings.

This chapter includes the following topics:

- [“Understanding Placeholder Datastores,”](#) on page 43
- [“Configure Datastore Mappings for vSphere Replication Management,”](#) on page 44
- [“Select Inventory Mappings,”](#) on page 44

Understanding Placeholder Datastores

For each virtual machine in a protection group, SRM creates a placeholder at the recovery site. You can add to these placeholders, and manage them as part of the recovery site's inventory.

You must identify a datastore that SRM will use to store placeholder virtual machines. After you determine which datastore will hold placeholders, SRM reserves a place for protected virtual machines in the recovery site's inventory. This is done by creating a subset of virtual machine files on the specified datastore at the recovery site and then using that subset to register the placeholder virtual machine with the recovery site vCenter. The presence of these placeholder virtual machines in the 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. You must establish placeholder datastores at both sites. Having placeholder datastores at both sites enables re-protection by providing a location to store the identity and inventory location of the old production machine in an empty shell of a virtual machine. This placeholder virtual machine is created during the recovery workflow when the production virtual machine is deactivated. This placeholder virtual machine can then be used and eventually removed as the recovery process is completed.

Placeholder virtual machines behave like any other member of the recovery site vCenter inventory, though they cannot be powered on. When a placeholder is created, its folder and compute resource assignments are derived from inventory mappings established at the protected site. Changes made to the placeholder virtual machines in the recovery site inventory override settings established by inventory mapping. These settings are preserved during the recovery or test.

When evaluating datastores in which to establish placeholder datastores, consider the following:

- For clusters, the placeholder datastores must be visible to all hosts in the cluster.
- You cannot replicate placeholder datastores.
- Configure placeholders for both sites if failback is to be supported.

Configure a Placeholder Datastore

Specify a placeholder datastore for SRM to use for the storage of placeholder virtual machines.

Prerequisites

Verify that sites are connected and paired.

Procedure

- 1 Click **Sites** in the left pane, and select a site
- 2 Click the **Placeholder Datastores** tab.
- 3 Click **Configure Placeholder Datastore**.
- 4 Expand the folders to find a datastore to designate as the location for placeholder datastores, click the datastore, and click **OK**.

The selected placeholder datastore appears in the Datastore column. If the datastore is on a standalone host, the host name appears. If the datastore is on a host that is in a cluster, the cluster name appears.

NOTE If an array manager is replicating datastores, but the array manager is not configured with SRM, the option to select the replicated datastore might be available. Do not select replicated datastores. Previously configured and replicated datastores appear but you cannot select them.

Configure Datastore Mappings for vSphere Replication Management

Configure datastore mappings to determine which datastores are used to store replicated virtual machine disks and configuration files at the recovery site. Use datastore mappings when you configure vSphere Replication (VR) for virtual machines as a way to select the default destination datastores.

Datastore mappings are determined based on the source datastores of the virtual machines being configured for replication. A source datastore can be a single datastore containing a single virtual machine, or many datastores with many virtual machines with files spread across the datastores. When configuring replication for a single virtual machine, mappings may be overridden, but when configuring replication for multiple virtual machines, you use only datastore mappings, and you cannot override these mappings.

Datastores can be nested within objects, so you might need to expand folders or datacenters to find a datastore.

Procedure

- 1 Click **vSphere Replication** in the left pane, and select a site.
- 2 Click the **Datastore Mappings** tab, and select one or more source datastores.
- 3 Click **Configure Mapping**.
- 4 Browse to and select a datastore to map to at the recovery site.

Select Inventory Mappings

Mappings provide default locations and networks for use when placeholder virtual machines are initially created on 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. Inventory mappings 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, such as when new members are added. If you create a protection group and no mappings exist, you must configure each protected virtual machine individually. While SRM does not enforce an inventory mapping requirement, a virtual

machine cannot be protected unless it has some form of valid inventory mappings for networks, folders, and compute resources. Inventory mappings can be created at both protection and recovery site. After having established mappings at the protected site when configuring protection, you should also configure inventory mappings at the recovery site.

During reprotect, the virtual machines that were protected before failover are used to complete reprotection. The originally protected virtual machines and their backings are used to protect the virtual machines that were formerly the recovery virtual machines. For reprotection where devices are added after a virtual machine is failed over or when original production virtual machines were deleted, then mappings are used during reprotection. In most cases, the previously protected virtual machines and their device backings are used during reprotection.

Because placeholders do not support NICs, you cannot make changes to placeholder virtual machine network configurations.

Changes to Inventory Mappings do not affect virtual machines that are already protected by SRM. New mappings are only applied to newly added virtual machines or if users repair a lost placeholder for a particular virtual machine. Placeholder virtual machines are created in the location specified by inventory mappings, but you can move these virtual machines among folders and resource pools using the vSphere Client.

Procedure

- 1 Click **Sites** in the left pane and select the site for which to configure mappings.
- 2 Select a tab for a type of inventory object to configure.

You can choose among the **Resource Mappings** tab, **Folder Mappings** tab, and **Network Mappings** tab. The Mapping page displays a tree of resources at the protected site.

- 3 Select an inventory object and click **Configure Mapping**.
- 4 Expand the inventory items and navigate to the recovery site resource (network, folder, or resource pool) to which you want to map the protected site resource.
- 5 (Optional) Choose a number of options for how to establish the mapping:
 - Select an existing resource.
 - To create an object for a folder or resource mapping, click the button for creating a resource.

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

- 6 Repeat step [Step 2](#) through step [Step 4](#) for any resource types for which you want to establish mappings.

What to do next

Assign placeholder datastores. See [“Understanding Placeholder Datastores,”](#) on page 43.

Configuring Array-Based Protection

After you pair the protected and recovery sites, you must configure protection for virtual machines. If you are using array-based replication, configure storage replication adapters (SRAs) at each site, and then configure SRM at each site.

If you are using only vSphere Replication, array managers are not required, and you can go to [Chapter 5, “Installing vSphere Replication Servers,”](#) on page 51.

This chapter includes the following topics:

- [“Configure Array Managers,”](#) on page 47
- [“Edit Array Managers,”](#) on page 48

Configure Array Managers

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

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

- Select the array manager type and provide a name. When editing an array manager, you can change the name, but you cannot change the type.
- Provide SRM with connection information and credentials (if needed) for array management systems at the protected and recovery sites.

You typically configure array managers only once, after you connect the sites. You do not need to reconfigure them unless array manager connection information or credentials change, or you want to use a different set of arrays.

Prerequisites

- Connect the sites as described in [“Connect the Sites,”](#) on page 39.
- Install SRAs at both sites as described in [“Install Storage Replication Adapters,”](#) on page 38.

Procedure

- 1 Click **Array Managers**, and select the folder in which you want to configure array managers.
- 2 Click **Add Array Manager**.

- 3 Provide an SRA name and select an adapter type.
 - a Type a name for the array in the **Display Name** text box.
Use a descriptive name that makes it easy for you to identify the storage associated with this array manager.
 - b Ensure that the array manager type that you want SRM to use appears in the **SRA Type** field.
If more than one SRA has been installed on the SRM server host, click the drop-down arrow and select a manager type. If no manager type appears, either you need to rescan SRAs or no SRA has been installed on the SRM server host.
- 4 Provide the required information for the adapter selected.
The SRA creates these text boxes. For more information about how to fill them in, see the documentation provided by your SRA vendor. While text boxes vary among SRAs, common text boxes include IP address, protocol information, mapping between array names and IP addresses, and user name and password.
- 5 Finish the wizard.
- 6 Repeat steps [Step 1](#) through [Step 5](#) to configure an array adapter for the recovery site.
- 7 In the **Array Pairs** tab, select an array pair, and click **Enable**.
- 8 If array managers have been added, but no array pairs are visible, click **Refresh** to collect the latest information about array pairs.

Rescan Arrays to Detect Configuration Changes

SRM checks arrays for changes to device configurations every 24 hours. However, 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 Click **Array Managers** in the left pane and click **Configure**.
- 2 On the Protected Site Array Managers page, click **Next**.
- 3 Click **Next**.
- 4 Click **Rescan Arrays**.
- 5 Click **Finish** to complete the operation.

Edit Array Managers

Use the Edit Array Manager wizard to modify an array manager's name or other settings such as IP address or user name and password.

For more information about how to fill in the adapter fields, see the documentation provided by your SRA vendor. While fields vary among SRAs, common fields include IP address, protocol information, mapping between array names and IP addresses, and user names and passwords.

Procedure

- 1 Click **Array Managers** in the left pane, and select an array manager.
- 2 Click **Edit Array Manager**.

- 3 Modify the name for the array in the **Display Name** field.
Use a descriptive name that makes it easy for you to identify the storage associated with this array manager. The array manager type cannot be modified.
- 4 Modify the adapter information.
These fields are created by the SRA.
- 5 Click **Finish** to complete the modification of the array manager.

Installing vSphere Replication Servers

vSphere Replication (VR) uses replication technologies included in ESX Servers with the assistance of virtual appliances to replicate virtual machines between sites.

VR is provided by vSphere Replication Servers (VR Servers or VRS). VR Servers are managed by the vSphere Replication Management Server (VRMS). Both VRMS and VR Servers are virtual appliances. To use VR you must deploy exactly one VRMS server at each site and at least one VR Server at the recovery site. To enable replication in both directions, you must deploy at least one VRS at each site. To meet the load balancing needs of your VR you might want to deploy multiple VR Servers at each site. Each VRMS must be registered with a corresponding vCenter Server. For example, the primary site VRMS must be registered with the primary site vCenter Server.

Both the VRMS and VRS appliances provide a virtual appliance management interface (VAMI). You can use these interfaces to configure the VRMS database, as well as network settings, public-key certificates, and passwords for the appliances.

Before using VR, you need to configure the VR infrastructure including having managed IP addresses defined in the runtime settings at both sites and having a VRMS database installed. The getting started page provides guidance to ensure that you complete the installation and configuration process correctly.

- When installing SRM, make sure that you select the VR option. If you installed SRM and want to add VR, you can add that option by running the installer again.
- Pair SRM servers as described in [“Connect the Sites,”](#) on page 39.
- Deploy VRMS at both sites as described in [“Deploy a vSphere Replication Management Server,”](#) on page 52.
- Pair VRMS servers as described in [“Configure vSphere Replication Management Connections,”](#) on page 55.
- Deploy VR Servers as described in [“Deploy a vSphere Replication Server,”](#) on page 55.
- Configure VR Servers as described in [“Configure vSphere Replication Server Settings,”](#) on page 56.
- Register VR Servers as described in [“Register a vSphere Replication Server,”](#) on page 57.

This chapter includes the following topics:

- [“Deploy a vSphere Replication Management Server,”](#) on page 52
- [“Configure vSphere Replication Management Server Settings,”](#) on page 52
- [“Configure vSphere Replication Management Connections,”](#) on page 55
- [“Deploy a vSphere Replication Server,”](#) on page 55
- [“Configure vSphere Replication Server Settings,”](#) on page 56
- [“Register a vSphere Replication Server,”](#) on page 57

Deploy a vSphere Replication Management Server

Before using vSphere Replication (VR), you need to deploy a vSphere Replication Management Server (VRMS) from the SRM UI, thereby creating the framework to support VR. VR requires that VRMS be deployed at each site.

Prerequisites

- Verify that you have a static IP address.
- Verify that the SRM database has been set up.
- Verify that you enabled the VR option during SRM installation.

Procedure

- 1 Click **vSphere Replication** in the left pane.
- 2 Select one of the sites, which is indicated with a folder icon, and click the **Getting Started** tab.
- 3 Click **Deploy a vSphere Replication Management Server**.

A VRMS is deployed as a virtual appliance using the OVF wizard.

- 4 Click **OK** to start the OVF wizard.

Information about the OVF wizard appears in the *vSphere Virtual Machine Administration Guide*. The OVF wizard guides you through basic appliance configuration. When prompted by the OVF wizard, you must set the root password for the VRM Server appliance.

The VRMS is deployed. If problems occurred during the installation process, you can find the OVF for the VRMS server virtual appliance in the SRM installation directory. Use this file to manually deploy the appliance from the vSphere Client by selecting **Deploy OVF Template** from the **File** menu. The OVF file can be found in the `www` directory in the vCenter Server installation. For example, it might be found in `C:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\www\HMS_OVF10.ovf`.

Configure vSphere Replication Management Server Settings

You configure vSphere Replication Management Server (VRMS) settings through the virtual appliance management interface (VAMI). These settings are established during installation. You can modify the settings after the server is deployed.

The Configuration section of the **VRMS** tab is used to configure the following items:

- Database settings
- vCenter Server information
- SSL Certificate settings

Configure these settings before you use VR replication. Manage these settings using the VAMI after deploying VRMS. You can use the SRM interface to connect to the VRMS by clicking the **Getting Started** tab and clicking the **Configure VRMS** link. Alternatively, you can connect to VAMI on the Web by entering the server's IP address and the specify the port (8080) in a browser. A sample address might be `https://192.168.1.2:8080/`. The IP address might vary, but the port is always 8080.

Prerequisites

- Verify that the VRMS is powered on.
- You must have administrator privileges to configure VRMS.

Procedure

- 1 Connect to the VRMS Web interface.
- 2 Provide the user name and password for the server.
The user name is root and the password is configured during the OVF deployment of the VRMS.
- 3 Click the **VRMS** tab and click **Configuration**.
- 4 In Configuration mode, select manual configuration to specify a configuration or select **Configure from an Existing VRMS Database** to use a previously established configuration.
- 5 In the DB text boxes, provide information about the database for VRMS to use.
- 6 In the VRMS text boxes, provide the name and network address for the VRMS.
You must provide unique names for the VRMS on the protected site and on the recovery site.
- 7 In the VC text boxes, provide information about the vCenter Server being used with this installation.
Use the same address format, IP address or FQDN, that you used during installation.
- 8 Install an SSL certificate.

Option	Description
Enable the SSL Certificate Policy	Enabling the policy provides greater assurance, but might require additional public-key configuration. Enable the SSL Certificate Policy for greater assurance.
Generate a self-signed certificate or install an existing certificate	Using a self-signed certificate might provide reduced levels of trust and might not be suitable for environments where high-levels of compliance with security standards are required.

- 9 To apply changes, click **Save and Restart Service**.

Configure VRMS Security Settings

You may change the VRMS server password to meet security standards of your organization. Use this password to log into the VAMI or to log into the console.

Prerequisites

Ensure that the VRMS server is powered on.

Procedure

- 1 In the SRM interface, click the **Getting Started** tab and click **Configure VRMS Server** to connect to the VRMS server.

Alternatively, you can connect to the VRMS server web interface by entering the server's IP address and port 8080 in a browser. A sample address might be <https://192.168.1.2:8080/>.
- 2 Provide the user name and password for the server.
User name is always root.
- 3 Click the **VRMS** tab and click the **Security** button.
- 4 To change the root user password:
 - a Type the current password in the **Current Password** field.
 - b Type the new password in the **New Password** and the **Confirm New Password** text boxes.
VRMS does not support blank passwords.
 - c Click **Apply** to change the password.

- 5 Review the current SSL information.

Configure VRMS Network Settings

You can review current network settings and change address and proxy settings for the VRMS Server. You might make these changes to match network reconfigurations.

Procedure

- 1 In the SRM interface, click the **Getting Started** tab, and click **Configure VRMS Server** to connect to the VRMS server.

Alternatively, you can connect to the VRMS server web interface by entering the server's IP address and port 8080 in a browser. A sample address might be `https://192.168.1.2:8080/`.
- 2 Type the user name and password for the server.

User name is always root.
- 3 Click the **Network** tab.
- 4 Click the **Status** button to review current settings.
- 5 Click the **Address** button to review and modify address settings.
 - a Select DHCP or static IP addressing.

DHCP is not recommended if the IP address of the appliance might change if it reboots.
 - b Enter IP settings, DNS settings, and host name information.
 - c Click **Save Settings**. If you do not click **Save Settings**, changes are discarded.
- 6 Click the **Proxy** button to review and modify proxy settings.
 - a Enable **Use a proxy server** to use a proxy server.
 - b Type a proxy server name in the Proxy Server text box.
 - c Type a proxy port in the Proxy Port text box.
 - d (Optional) Type a proxy server user name and password.
 - e Click **Save Settings**. If you do not click **Save Settings**, changes are discarded.

Configure VRMS System Settings

Configure VRMS system settings if you need to administer or gather information about the vSphere Replication Manager Server (VRMS) appliance. Use the **System** tab to review information for a server virtual appliance. You can review information about: appliance vendor, name, and version; host name; and appliance operating system and version.

Procedure

- 1 In the SRM interface, click **Getting Started**, and click the **Configure VRMS Server** link.

Alternatively, you can connect to the VRMS server Web interface by entering the server's IP address and port 8080 in a browser. A sample address might be `https://192.168.1.2:8080/`.
- 2 Connect to the VRMS server.
- 3 Type the user name and password for the server.

User name is always root.
- 4 Click the **System** tab.

- 5 Click the **Information** button.

For this server virtual appliance, you can review information about:

- Appliance vendor
 - Appliance name
 - Appliance version
 - Hostname
 - Appliance operating system and version.
- a Click **View** for information about the OVF environment.
 - b Click **Reboot** to reboot the virtual appliance.
 - c Click **Shutdown** to shutdown the virtual appliance.

Shutting down the VRMS virtual appliance does not affect the replication of already configured virtual machines but it prevents the configuring replication of new virtual machines as well as modifying existing replication settings.

- 6 Click the **Time Zone** button and select a time zone from the **System Time Zone** drop-down list and click **Save Settings**.

Configure vSphere Replication Management Connections

Configure vSphere Replication Management Server (VRMS) connections between sites that support vSphere Replication (VR).

You can complete this process at either site that has VRMS servers installed. If you are using an untrusted certificate, certificate warnings might appear during the process.

Prerequisites

Verify that the VRMS servers are deployed and configured at two sites, both of which also have SRM servers installed and paired.

Procedure

- 1 Click **vSphere Replication** in the left pane, and select a site.
A site is indicated with a folder icon,
- 2 Click the **Summary** tab, and click **Configure VRMS Connection**.
- 3 Enter the administrator password for the remote vCenter Server.
- 4 Click **OK** to confirm connection between the servers.

The **Summary** tab for the **vSphere Replication** pane shows the servers are connected.

Deploy a vSphere Replication Server

Deploy a vSphere Replication Server (VRS) to enable vSphere Replication (VR). To use VR you must deploy exactly one VRMS server at each site and at least one VR Server at the recovery site. To enable replication in both directions, you must deploy at least one VRS at each site.

To meet the load balancing needs of your VR replication you may want to deploy multiple VR Servers at each site. To meet your load balancing needs you might want to deploy multiple VR Servers at each site. The protected site VRMS must be registered with the protected site vCenter Server. The recovery site VRMS must be registered with the recovery site vCenter Servers, and the VR servers must be registered with the VRMS at their site.

The VRS, which is a virtual appliance, is stored with the SRM server. VRS servers should be deployed at a site after a VRMS server is deployed and configured for that site. If problems occur when deploying the VRS using the SRM interface, the VRS OVF file can be found in the SRM installation directory. This file can be used to manually deploy the appliance from the vSphere Client by selecting **Deploy OVF Template** from the **File** menu. The OVF file is in the `www` directory in the vCenter Server installation. For example, it might be found in `C:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\www\vrserver_OVF10.ovf`.

This procedure involves connecting to SRM as described in [“Connecting to SRM,”](#) on page 24.

Procedure

- 1 Click **vSphere Replication**, and click the **Summary** tab,
- 2 Click **Deploy VR Server**.

The VRS is deployed as a virtual appliance using an OVF wizard.

- 3 Click **OK** to launch the OVF wizard.

Information about the OVF wizard is contained in deploying OVF templates in the *vSphere Virtual Machine Administration Guide*.

Configure vSphere Replication Server Settings

Configure vSphere Replication Server (VRS) settings using the virtual appliance management interface. VRS settings are established during installation. You can modify the settings after the server is deployed. Unlike VRMS, VRS does not require any additional configuration through the VAMI after deployment.

Using a self-signed certificate provides the benefit of public-key based encryption and authentication, though using such a certificate does not provide the level of assurance offered when using a certificate signed by a certificate authority.

Prerequisites

Verify that a VRS is installed, and that the server is powered on.

Procedure

- 1 In the SRM interface, select **vSphere Replication**.
- 2 Select a VR Server, and click the **Configure VR Server** link.

Alternatively, you can connect to the VRS Web interface by entering the server's IP address and port (5480) in a browser. A sample address might be `https://192.168.1.2:5480`. The IP address might vary, but the port is always 5480.

- 3 Log into the VR Server configuration interface as root.

The default root password is **vmware**.

- 4 Click the **VR Server** tab.
- 5 Change the Super User password.
- 6 (Optional) Click **Generate and Install** to generate and install a self-signed certificate.
- 7 Use an existing SSL certificate.
 - a Click the **Browse** button next to the **Certificate** text box to browse for an existing certificate.
 - b Click the **Browse** button next to the **Key** text box to browse for an existing private key.
 - c Click **Upload** to upload the specified certificate and key.

Register a vSphere Replication Server

You must register the vSphere Replication Server (VRS) with the vSphere Replication Management Server (VRMS) to enable vSphere Replication (VR).

Prerequisites

- Configure a VRMS as described in [“Configure vSphere Replication Server Settings,”](#) on page 56.
- Verify that the VRMS is paired as described in [“Configure vSphere Replication Management Connections,”](#) on page 55.

Procedure

- 1 Click **vSphere Replication** in the left pane, select a site, and click **Register VR Server**.

The Register VRS page displays a list of resources at the selected site.

- 2 Select a virtual machine in the inventory that is a working VRS and click **OK**.

If you are using an untrusted certificate, certificate warnings might appear.

What to do next

Create protection groups and replicate virtual machines as described in [Chapter 6, “Creating Protection Groups and Replicating Virtual Machines,”](#) on page 59.

Creating Protection Groups and Replicating Virtual Machines

6

After you configure a replication solution, you create protection groups. You can also customize the configuration of virtual machine replication and protection.

This chapter includes the following topics:

- [“Limitations to Protection and Recovery of Virtual Machines,”](#) on page 59
- [“Create Array-Based Protection Groups,”](#) on page 60
- [“Create vSphere Replication Protection Groups,”](#) on page 61
- [“Configure Replication for a Single Virtual Machine,”](#) on page 62
- [“Configure Replication for Multiple Virtual Machines,”](#) on page 63
- [“Replicate Virtual Machines Using Physical Couriering,”](#) on page 64
- [“Move a Virtual Machine to a New vSphere Replication Server,”](#) on page 66
- [“Apply Inventory Mappings to All Members of a Protection Group,”](#) on page 66

Limitations to Protection and Recovery of Virtual Machines

The protection and recovery by SRM of virtual machines in the suspended state, virtual machines with snapshots, and virtual machines that are linked clones is subject to limitations.

Protection and Recovery of Suspended Virtual Machines

When you suspend a virtual machine, vSphere creates and saves its memory state. When the virtual machine resumes, vSphere restores the saved memory state to allow the virtual machine to continue without any disruption to the applications and guest operating systems that it is running.

Protection and Recovery of Virtual Machines with Snapshots

Array-based replication supports the protection and recovery of virtual machines with snapshots, but with limitations.

You can specify a custom location for storing snapshot delta files by setting the `workingDir` parameter in VMX files. SRM does not support the use of the `workingDir` parameter.

Limitations also apply if you are running versions of ESX or ESXi Server older than version 4.1.

- If the virtual machine has multiple VMDK disk files, all the disk files must be contained in the same folder as the VMX file itself.

- If a virtual machine is attached to a Raw Disk Mapping (RDM) disk device, you must store the mapping file in the same folder as the VMX file. RDM snapshots are only available if you create the RDM mapping using Virtual Compatibility Mode.

If you are running an ESX or ESXi Server 4.1 or later, these limitations do not apply.

vSphere Replication supports the protection of virtual machines with snapshots, but you can only recover the latest snapshot.

Protection and Recovery of Linked Clone Virtual Machines

vSphere Replication does not support the protection and recovery of virtual machines that are linked clones.

Array-based replication supports the protection and recovery of virtual machines that are linked clones if all the nodes in the snapshot tree are replicated.

Create Array-Based Protection Groups

For array-based replication, SRM organizes datastore groups to collect all files associated with protected virtual machines. You then associate these datastore groups with protection groups. All virtual machines in a datastore group replicate their files together, and all failover together. You can have a virtual machine with files on different datastores. In such a case, SRM combines the datastores that contain files for a single virtual machine to create a datastore group.

A datastore group is the smallest unit of storage that can be failed over or tested independently. Several rules show how the groups are calculated.

- If a device is used by a datastore, all devices used by that datastore are combined.
- If a device is a part of a consistency group in the array, all devices in the consistency group are combined.
- If a virtual machine spans multiple datastores, all devices belonging to all such datastores are combined.

These rules are repeated until no more devices are added to the datastore group.

You can organize protection groups in folders. Different views in the Recovery interface display the names of the recovery groups, but they do not display the folder names. If you have two protection groups with the same name in different folders, it can be difficult to tell them apart in some views in the Recovery interface. Consequently, you should ensure that protection names are unique across all folders. In environments in which not all users have view privileges for all folders, to be sure of the uniqueness of protection group names, do not place protection groups in folders.

Procedure

- 1 Click **Protection Groups**, and click the **Create Protection Group**.
- 2 On the Select a Site and Protection Group Type page, select which site will be protected and select **Array Based Replication**.
- 3 Select a datastore group from the list, and click **Next**.

The datastore groups listed were discovered when you configured the array managers. Each datastore group in the list is replicated to the recovery site. When you select a datastore group, the virtual machines in that datastore group are listed in the lower **Virtual Machines on the Selected Datastore Group** pane, and are marked to include in the protection group after you establish the protection group.

- 4 Type a name and optional description for the protection group, and click **Next**.
- 5 Click **Finish** to create the protection group and begin the automatic protection of the specified virtual machines.

SRM creates a protection groups that you can use to protect virtual machines. After protection is established, placeholders are created and inventory mappings applied for each virtual machine in the group. If a virtual machine 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 is not created for it. Wait to ensure that the operations complete as expected. Make sure that the protection group was created and virtual machines were protected. The progress of those tasks can be monitored in the Recent Tasks panel of the vSphere Client.

What to do next

Create a protection plan with which to associate your protection groups. This process is described in [“Create a Recovery Plan,”](#) on page 67.

Edit Array-Based Protection Groups

You can change the name and description of an array-based protection group and add or remove datastore groups that are part of the protection group.

Procedure

- 1 Click **Protection Groups**, and select a protection group.
- 2 Click **Edit Protection Group**.
- 3 Edit the name or description of the protection group.
- 4 Add or remove virtual machines that are included in the protection group and click **OK**.

Create vSphere Replication Protection Groups

Create protection groups for vSphere Replication (VR) to enable protection.

You can organize protections groups in folders. Different views in the Recovery interface display the names of the recovery groups, but they do not display the folder names. If you have two protection groups with the same name in different folders, it can be difficult to tell them apart in some views in the Recovery interface. Consequently, you should ensure that protection names are unique across all folders. In environments in which not all users have view privileges for all folders, to be sure of the uniqueness of protection group names, do not place protection groups in folders.

Prerequisites

Establish vSphere Replication for virtual machines using the vSphere Client. See [“Configure Replication for a Single Virtual Machine,”](#) on page 62 or [“Configure Replication for Multiple Virtual Machines,”](#) on page 63.

Procedure

- 1 Click **Protection Groups** and click the **Create Protection Group**.
- 2 On the Select a Site and Protection Group Type page, select which site will be protected and select **vSphere Replication**.
- 3 Select virtual machines from the list and click **Next**.
Only virtual machines with VR enabled and not already in a protection group appear.
- 4 Type a name and optional description for the protection group, and click **Next**.
- 5 Click **Finish** to create the protection group and begin the automatic protection of the specified virtual machines.

What to do next

Next create a recovery plan with which to associate your protection groups. This process is described in [“Create a Recovery Plan,”](#) on page 67.

Edit vSphere Replication Protection Groups

You can change the name of a vSphere Replication (VR) protection group and the virtual machines that are protected in the group.

Procedure

- 1 Click **Protection Groups** in the left pane and right-click a protection group.
- 2 Select **Edit Protection Group**.
- 3 Edit the name or description of the protection group.
- 4 Add virtual machines to the protection group and click **OK**.

Configure Replication for a Single Virtual Machine

Individual virtual machines and their virtual disks can be protected by vSphere Replication (VR).

Prerequisites

To replicate a virtual machine using VR, you must establish the VR infrastructure at both sites. A minimal VR infrastructure requires one vSphere Replication Management Server (VRMS) at both sites and at least one VR Server (VRS) on the recovery site. SRM servers at each site must be paired, and VRMS must be paired as well.

Procedure

- 1 On the vSphere Client Home page, click **VMs and Templates**.
- 2 Browse the inventory to find the single virtual machine to be replicated using VR.
- 3 Right-click the virtual machine and click **vSphere Replication...**

- 4 Configure general replication settings.
 - a Use the Recovery Point Objective (RPO) slider or enter a value to configure the maximum amount of data that can be lost in the case of a site failure. The available range is from 15 minutes to 24 hours.

For example, a recovery point objective of one hour seeks to ensure that the virtual machine loses no more than one hour of data during the recovery. For smaller RPOs, less data is lost in a recovery, but more network bandwidth is consumed keeping the replica up to date.

NOTE Every time that a virtual machine reaches its RPO target, SRM records approximately 3800 bytes of data in the vCenter Server events database. If you set a low RPO period, this can quickly create a large volume of data in the database. To avoid creating large volumes of data in the vCenter Server events database, you should limit the number of days that vCenter Server retains event data. See [Configure Database Retention Policy](#) in the *vCenter Server and Host Management Guide*. Alternatively, set a higher RPO value.

- b Select a Guest OS Quiescing configuration.

The available quiescing types are determined by the virtual machine's operating system. Microsoft Volume Shadow Copy Service (VSS) quiescing is supported for Windows virtual machines running Windows Server 2003 or newer. VR does not support quiescing for Linux and older versions of Windows such as Windows XP or Windows 2000. For Windows 7 and Windows 2008 and newer virtual machines file-system level quiescing is supported, but application level quiescing is not.
 - c If no target file location is specified or to override the default determined by the datastore mappings, click **Browse** to select a target location for the virtual machine. If this option is selected, an empty, blank replica disk is created. To use non-blank disks, see [“Replicate Virtual Machines Using Physical Couriering,”](#) on page 64.

Either select a datastore or a folder within a datastore:

- Browse to select a datastore and click **OK**.
 - Enable **Specify datastore folder**, browse to a datastore, click **Browse**, browse to find the desired folder in the Browse Datastores window, then double-click the desired folder.
- 5 Select a replication destination for each media device for the virtual machine. Repeat this step for each device in the virtual machine.

The next pages are created dynamically. They might include multiple virtual drives, all of which you can configure uniquely. Configurable settings include whether the virtual drive is replicated, the virtual drive's replication destination, and information about how the replicated virtual drive is configured. If the disk is to be replicated, select a replication destination before proceeding.
 - 6 Accept the automatic assignment of a VRS, or select a particular server.
 - 7 Review the settings and click **Finish** to establish replication.

Configure Replication for Multiple Virtual Machines

Multiple virtual machines can be protected by vSphere Replication.

Prerequisites

To replicate virtual machines using VR, establish the VR infrastructure at both sites. A minimal VR infrastructure requires one vSphere Replication Management Server (VRMS) at both sites and at least one vSphere Replication Server (VRS) on the recovery site. SRM servers at each site must be paired, and VRMSs must be paired as well. If a virtual machine is powered off, replication begins when the virtual machine is powered on.

Before replicating multiple machines, configure datastore mappings in the SRM user interface. For more information, see [“Configure Datastore Mappings for vSphere Replication Management,”](#) on page 44. You configure the mappings so that information is available to SRM regarding the target datastore destinations for replication.

Procedure

- 1 On the vSphere Client Home page, click **VMs and Templates**.
- 2 Select a folder or datacenter in the left pane and click the **Virtual Machines** tab.
A list of virtual machines appears in the right pane.
- 3 On the right pane, right-click the virtual machines to be replicated and click **vSphere Replication**
- 4 On the Replication Settings page, configure general replication settings.

- a Use the Recovery Point Objective (RPO) slider or enter a value.

The RPO value determines the maximum amount of data that can be lost during the recovery. The available range is from 15 minutes to 24 hours.

NOTE Every time that a virtual machine reaches its RPO target, SRM records approximately 3800 bytes of data in the vCenter Server events database. If you set a low RPO period, this can quickly create a large volume of data in the database. To avoid creating large volumes of data in the vCenter Server events database, you should limit the number of days that vCenter Server retains event data. See [Configure Database Retention Policy](#) in the *vCenter Server and Host Management Guide*. Alternatively, set a higher RPO value.

- b Select the Guest OS Quiescing method.
- c Choose whether to enable **Initial copies of .vmdk files have been placed on the target datastore**.
This option searches the replicated datastore. If candidate files are found, confirm whether to use the found files as initial copies.
- 5 On the VRS page, accept the automatic assignment of a VRS or select a particular server.
- 6 Review the settings and click **Finish** to establish replication.

Replicate Virtual Machines Using Physical Couriers

Initial replication of .vmdk files can be made more efficient by physically couriers files on a storage device. This may also be known as sneakernet replication. Physical couriers may be required if it's not practical to copy the files over the network because the amount of data is large, the bandwidth available is small, or some combination of the two.

When replicating virtual machines, it is best to ensure that virtual machines are replicated to subdirectories within datastores. Copying disks are expected to work, as long as the method preserves the identity information stored inside the .vmdk file.

Prerequisites

To replicate a virtual machine using vSphere Replication (VR), you must establish the VR infrastructure at both sites. A minimal VR infrastructure requires one vSphere Replication Management Servers (VRMS) at both sites and at least one VR server on the recovery site. SRM servers at each site must be paired, and VRMS must be paired as well. If a virtual machine is powered off, replication begins when the virtual machine is powered on.

Procedure

- 1 Use the vSphere Client to connect to a vCenter Server that can manage the virtual machines to be physically couriers.

- 2 Click **Datastores**. In the left pane, browse to the datastore that contains the files for the virtual machine, select the datastore, and in the right pane, click **Browse this datastore**.
- 3 Select the folders for all virtual machines to be couriered, right-click the selection, and click **Download...**
- 4 Select a destination to which to copy the files and click **OK**.
- 5 Click **Yes**.
- 6 After the download finishes, transfer media containing the files to the paired site location to upload them.
- 7 On the vSphere Client Home page at the paired site, click **Datastores**. In the left pane, browse to the datastore that will contain the files for the virtual machine, select the datastore, and in the right pane, click **Browse this datastore**.
- 8 Select the folder that will contain the copies of the virtual machines, right-click the selection, and click **Upload Folder...**
- 9 Select the folder containing the virtual machines, and click **OK**.
- 10 On the production site, click **VMs and Templates** in the vSphere Client home page.
- 11 Find the virtual machine that will be replicated within the inventory. Right-click the virtual machine and click **Site Recovery Manager vSphere Replication...**
- 12 Configure general replication settings.
 - a Click **Browse...** to browse for a datastore to which to replicate the virtual machine. Click the **Browse** checkbox in the datastore selection dialog box, then double-click a destination folder to select it.

Select the folder you created when you uploaded the vmdk. Before selecting the destination directory, you must use the vSphere client to create the directory on the destination datastore.
 - b Use the Recovery Point Objective (RPO) slider or enter a value to configure the maximum amount of data that can be lost during the recovery.

For example, a recovery point objective of one hour seeks to ensure that no virtual machines lose more than one hour of data during the recovery. For smaller RPOs, less data is lost in a recovery, but more network bandwidth is consumed keeping replicas synchronized.
 - c Choose a Guest OS Quiescing configuration.

The available quiescing types are determined by the virtual machine's operating system. VSS quiescing is supported for Windows virtual machines running Windows XP or later. Linux does not support quiescing.
- 13 Select a replication destination for each media device for the virtual machine. Repeat this step for each device in the virtual machine.
 - a Select **Select an existing file to use as an initial copy** and click **Browse**.
 - b Browse to find the vmdk file that was physically transferred to the other site and double-click the file.
 - c Click **Yes** to confirm that the physically transferred copy of the vmdk will be overwritten in future replication.

The next pages are created dynamically. They may include multiple virtual drives, all of which can be configured uniquely. Configurable settings include whether the virtual drive is replicated, the virtual drive's replication destination, and information about how the replicated virtual drive is configured. You must specify a replication destination before proceeding.
- 14 In the VR Server page, either accept the automatic assignment of a VR Server, or select a particular server.
- 15 Review the settings and click **Finish** to establish replication.

Move a Virtual Machine to a New vSphere Replication Server

After establishing vSphere Replication (VR), you can move the virtual machines being replicated to other VR Servers. You might do this to complete maintenance tasks on existing VR servers or to balance the load on VR servers in the event that one server becomes overloaded with virtual machines.

Procedure

- 1 Click **vSphere Replication**, and select a site.
- 2 Click the **Virtual Machines** tab.
- 3 Right-click a virtual machine and select **Move to**.
- 4 Select an VR Server from the Move to Replication Server pane and click **Next**.
- 5 Review the information about the planned move, and click **Finish** to complete the change.

Apply Inventory Mappings to All Members of a Protection Group

When you create a protection group for either vSphere or array-based replication, 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, you can reapply the mappings to all unconfigured virtual machines in one step.

Procedure

- 1 Click **Protection Groups** in the left pane and click the **Virtual Machines** tab.
- 2 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 finishes, virtual machines that could not be configured have a status of Mapping Missing or Mapping Invalid. You must configure protection for these machines individually.

Recovery Plans and Reprotection

After you have configured SRM at the protected and recovery sites, you can create and test a recovery plan without affecting services at either site. You can also perform planned migrations or disaster recoveries by running a recovery plan, and, if necessary, reverse the role of your two sites by performing a reprotect.

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. You must assign each permission separately.

This chapter includes the following topics:

- [“Create a Recovery Plan,”](#) on page 67
- [“Test a Recovery Plan,”](#) on page 68
- [“Run a Recovery Plan,”](#) on page 70
- [“Understanding Reprotection,”](#) on page 71

Create a Recovery Plan

Create a recovery plan to establish how virtual machines are recovered. A basic recovery plan includes steps that use default values to control how virtual machines in a protection group are recovered at the recovery site. You can customize the plan to meet your needs. Recovery plans are different from protection groups in that recovery plans indicate how virtual machines in one or more protection groups are restored at the recovery site.

During tests, keep the virtual machine that is recovered during the test isolated from other machines in your environment. If duplicate machines are brought on-line and begin interacting with other machines in your production network, errors can occur. You can isolate virtual machines restored during test recoveries either by selecting **Auto**, which is an isolated network, or by selecting a network that was manually created, but which is not connected to other networks.

Procedure

- 1 Click **Recovery Plans**, and click **Create**.
- 2 On the Recovery Site page, choose the recovery site.
- 3 On the Select Protection Groups page, select one or more protection groups for the plan to recover, and click **Next**.

You might need to expand folders to find protection groups.

- 4 Select a recovery site network to which recovered virtual machines connect during recovery plan tests, and click **Next**.
- 5 Type a name for the plan in the **Recovery Plan Name** text box and add an optional description.
- 6 Click **Next**.
- 7 Review the summary information and click **Finish** to create the recovery plan.

Edit a Recovery Plan

You can change the properties of a recovery plan including plan name and description, test networks, and the contained protection groups. Edit a recovery plan to change any of the properties that you specified when you created it. You can edit recovery plans either from the protected site or the recovery site.

Procedure

- 1 Click **Recovery Plans**, and select the recovery plan you want to edit.
- 2 Click **Edit**.
After you open the plan for editing, you can change any of its properties.
- 3 Choose the recovery site.
- 4 Select one or more protection groups for the plan to recover, and click **Next**.
You might need to expand folders to find the protection groups.
- 5 Select a recovery site network to which recovered virtual machines connect during recovery plan tests, and click **Next**.
- 6 Type a name for the plan in the **Recovery Plan Name** text box and add an optional description.
- 7 Click **Next**.
- 8 Review the summary information and click **Finish** to make the specified changes to the recovery plan.

Remove a Recovery Plan

You can remove a recovery plan if you no longer need it. Removing a recovery plan also deletes the history of the plan.

Prerequisites

Export the history of the recovery plan. Removing a recovery plan also deletes the history of the plan.

Procedure

- 1 Click **Recovery Plans**.
- 2 Right-click the plan that you want to remove, and select **Remove Recovery Plan**.

Test a Recovery Plan

When you test a recovery plan, you use a test network and a temporary snapshot 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. A test makes no other changes to the production environment at either site.

Procedure

- 1 Click **Recovery Plans** in the left pane.
- 2 Click the recovery plan to test, and click **Test**.
- 3 Determine whether to enable **Replicate recent changes to recovery site**.
Enabling this option ensures that the recovery site has the latest copy of protected virtual machines, but the synchronization might take additional time.
- 4 Click **Next**.
- 5 Review the confirmation window and click **Finish**.
The wizard closes and the recovery plan test begins.
- 6 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 SRM initiates recovery steps in the prescribed order, with one exception. It does not wait for the Prepare Storage step to finish for all protection groups before continuing to the next steps.

- 7 When the recovery plan test completes, click **Cleanup**.
Running cleanup returns the protected virtual machines to their initial state and resets the recovery plan to the Ready state.
- 8 (Optional) If the cleanup encounters errors, run the cleanup again with the **Force Cleanup** option selected.
The **Force Cleanup** option cleans up and ignores any errors that might occur. If necessary, run cleanup several times with the **Force Cleanup** option selected, until the cleanup succeeds.

Cancel a Test or Recovery

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

When you cancel a test or recovery, 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 cancellation completes.
- Steps that add or remove storage devices are undone by cleanup operations if you cancel.

The time it takes to cancel a test or recovery depends on the type and number of steps that are currently in progress.

Procedure

- ◆ To cancel a test or recovery, click the **Cancel** button on the recovery plan toolbar.

Run a Recovery Plan

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

When you perform a planned migration, SRM attempts to replicate all virtual machines and gracefully shut down the protected machines. If errors occur during a planned migration, the plan pauses so that users can resolve errors. Replicating and shutting down virtual machines and providing opportunities to resolve errors makes it possible to reprotect the virtual machines. When you perform a disaster recovery, SRM attempts to shut down any virtual machines. If they cannot be shut down, the copies at the recovery site are still started, and automatic reprotection might not be possible.



CAUTION A recovery plan makes significant alterations in the configurations of the protected and recovery sites and it stops replication. Do not run any recovery plan that is not tested. In the case of array-based replication, recovered virtual machines and services might need to be supported at the recovery site for a period of time. Reversing these changes might cost significant time and effort and can result in prolonged service downtime.

Forced failover is intended for use in cases where storage arrays fail at the protected site and, as a result, protected virtual machines are unmanageable and cannot be shut down, powered off, or unregistered. In such a case, the system state cannot be changed for extended periods. To resolve this situation, you can force failover. Forcing failover does not complete the process of shutting down the virtual machines at the protected site. As a result, a split-brain scenario occurs, but the recovery might be completed quickly.

Running forced failover can affect the mirroring between the protected and the recovery storage arrays. After you run forced failover, you must check that mirroring is set up correctly between the protected array and the recovery array before you can perform further replication operations. If mirroring is not set up correctly, you must repair the mirroring by using the storage array software.

You can only use forced failover with array-based replication. Forced failover is not supported for vSphere Replication.

When forced failover is enabled, any outstanding changes on the protection site are not replicated to the recovery site before the failover sequence begins. Replication of the changes occurs according to the recovery point objective (RPO) period of the storage array. If a new virtual machine or template is added on the protection site and failover is initiated before the storage RPO period has elapsed, the new virtual machine or template does not appear on the replicated datastore and is lost. To avoid losing the new virtual machine or template, wait until the end of the RPO period before running the recovery plan with forced failover.

After the forced failover is finished and you have verified the mirroring of the storage arrays, you can resolve the issue that necessitated the forced failover. After the underlying issue is resolved, run planned migration on the recovery plan again, resolve any problems that occur, and rerun the plan until it finishes successfully. Running the recovery plan again does not affect the recovered virtual machines at the recovery site.

NOTE Forced failover is only available in SRM 5.0.1 and later.

Prerequisites

To use forced failover, you must first enable this function. You enable forced failover by enabling the `recovery.forcedFailover` setting as described in [“Change Recovery Site Settings,”](#) on page 89.

Procedure

- 1 Click **Recovery Plans** in the left pane and click the recovery plan to run.
- 2 In the Commands area, click **Recovery**.

- 3 Review the information in the confirmation prompt, and select **I understand that this process will permanently alter the virtual machines and infrastructure of both the protect and recovery datacenters.**
- 4 (Optional) If you enabled the forced failover function, you can select the **Forced Failover - recovery site operations only** check box.
- 5 Click **Next**.
- 6 Click **Start** to run the recovery plan.
- 7 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.

What to do next

You can reprotect virtual machines that use array-based replication. After a recovery plan is run and the virtual machines in it are operating at the recovery site, you can reprotect, which establishes protection of the virtual machines in the opposite direction. The virtual machines and the services that they provide are now accessible at the recovery site.

Understanding ReProtection

After a recovery, the recovery site becomes the new production site and is not protected. If a disaster occurs at the new production site, no other site is available to fail over to. A best practice is to protect the new production site to some other site immediately after a recovery. If the original production site is operational, you can use the original production site as a new recovery site to protect the new production site, effectively reversing the direction of protection. Reestablishing protection in the opposite direction by recreating all protection groups and recovery plans is time consuming and prone to errors. SRM provides an automated way to achieve reProtection.

ReProtection is available only in noncatastrophic failures. This means that the original vCenter servers, ESX Servers, SRM servers, and corresponding databases must be eventually recoverable.

For reprotect to be available, the following steps must first occur:

- 1 Run a planned migration and make sure that all steps of the recovery plan finish successfully. If errors occur during the recovery, resolve the problems that caused the errors and rerun the recovery. When you rerun a recovery, operations that succeeded previously are skipped. For example, successfully recovered virtual machines are not recovered again and continue running without interruption.
- 2 The original protected site must be available. The vCenter Server instances, ESXi Servers, SRM Server instances, and corresponding databases must all be recoverable.
- 3 If you performed a disaster recovery operation, you must perform a planned migration when both sites are running again. If errors occur during the attempted planned migration, you must resolve the errors and rerun the planned migration until it succeeds.

NOTE Reprotect is supported only for array-based replication. vSphere Replication (VR) reprotect is not supported. If a recovery plan contains VR groups, remove those groups before you run a reprotect operation.

ReProtection is not available if the following situations occur:

- Recovery plans cannot be completed without errors. For reprotect to be available, all steps of the recovery plan must finish successfully.
- VR groups are included in the recovery plan.

- The original site cannot be restored, for example if a physical catastrophe destroys the original site. To unpair and recreate the pairing of protected and recovery sites, both sites must be available. If you cannot restore the original protected site, you must reinstall SRM on both the protected and recovery sites.

Reprotection of MSCS and Fault Tolerant Virtual Machines

To use reprotection with Microsoft Cluster Server (MSCS) and fault tolerant virtual machines, the host machines on which the virtual machines run must meet certain criteria.

- You must run a fault tolerant virtual machine and its shadow on two separate ESXi Servers.
- You can run a cluster of MSCS virtual machines in the following possible configurations.
 - All the virtual machines of the cluster run on a single ESXi Server.
 - If the MSCS cluster is spread over more than one ESXi Server, you can only run one MSCS virtual machine per ESXi Server, with a maximum of two ESXi Servers in the cluster.

Because of these constraints, reprotection of MSCS and fault tolerant virtual machines is impossible if you do not enable VMware High Availability (HA) and Distributed Resource Scheduler (DRS). When moving MSCS and fault tolerant virtual machines across their primary and secondary sites during reprotection, you must enable HA and DRS, setting the affinity and anti-affinity rules as appropriate.

Reprotection Process

The process of completing reprotection includes two sets of tasks:

- Reversing the direction of protection groups using the new protected and recovery site configuration.
- Forcing a synchronization of storage from the new protected site (the original recovery site) to the new recovery site (to original protected site).

Reversing the Direction of Protection Groups

When a user initiates the reprotection process, SRM instructs the underlying arrays to reverse the direction of replication. After the replication is reversed, SRM creates placeholder virtual machines at the new recovery site (the original protected site).

When creating placeholder virtual machines, SRM uses the location of the original production virtual machine to determine where to create placeholder virtual machines. Placeholder virtual machines are created in a location that is defined as the placeholder location as part of inventory mapping. If the original production virtual machines are no longer available (for example because they were deleted by a user), SRM uses the reverse inventory mappings (from the original recovery site to the original production site) to determine the resource pools and folders for the placeholder virtual machines. Because of this, these reverse inventory mappings must be configured prior to running reprotect, otherwise reprotect may fail.

The files for the placeholder virtual machine are placed in the placeholder datastore defined for the original production site, not the datastore that held the original production virtual machine.

Forcing Data Synchronization

Forcing synchronization of data from the new protection site to the new recovery site ensures that the recovery site has a current copy of the production virtual machines running at the protection site. Forcing this synchronization ensures that recovery is immediately possible after the reprotect completes.

Manual Reprotection for vSphere Replication

Reprotection must be manually established for vSphere Replication. After a failover, existing replication is stopped. Existing replication must be unconfigured and then reconfigured in the opposite direction. It may be effective to use the original production virtual machine as a copy for physical couriering. In such a case, the disk names are different, so it is necessary to choose the target for each disk independently. Finally, configure protection groups and recovery plans in the opposite direction, thereby manually establishing reprotection.

Reprotection State Reference

As the process of reprotection proceeds, there are several states that can be observed in the recovery plan in the vSphere client SRM plug-in in the vSphere Client.

Table 7-1. Reprotection States

State	Description
Reprotect In Progress	SRM is attempting to complete the process of reprotection. If reprotection fails, the result is the state "Reprotect Interrupted".
Partial Reprotect	Occurs if multiple recovery plans share the same protection groups and some groups have already been reprotected in some plans, but some have not. If this occurs, run reprotect for partially reprotected plans and all groups that were not reprotected will be reprotected.
Incomplete Reprotect	Occurs because of failures during reprotect. For example, this state might occur due to a failure to synchronize storage or a failure to create placeholder virtual machines. <ul style="list-style-type: none"> ■ If a reprotect operation fails to synchronize storage, ensure sites are connected, review the reprotect step progress in the vSphere client SRM plug-in, and begin the reprotect task again. ■ If placeholder virtual machines are not created, recoveries are now possible, but you may wish to review the reprotect step progress in the vSphere client SRM plug-in, resolve any open issues, and begin the reprotect task again.
Reprotect Interrupted	If one of the SRM servers crashes during the reprotect, this state occurs. Make sure both SRM servers are running and begin the reprotect task again.

Customizing Site Recovery Manager

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 modify the settings.

This chapter includes the following topics:

- [“Customizing a Recovery Plan,”](#) on page 75
- [“Configure Protection for a Virtual Machine or Template,”](#) on page 86
- [“Configure Resource Mappings for a Virtual Machine,”](#) on page 87
- [“Configure SRM Alarms,”](#) on page 88
- [“Working with Advanced Settings,”](#) on page 88

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.

A simple recovery plan 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. However, you must customize most recovery plans intended for production use 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 series of steps that must be done in a specific order. You cannot change the order or purpose of the steps, but you can insert your own steps that display messages and run commands.

Recovery plan steps have several behaviors:

- Some steps are executed during all recoveries.
- Some steps are executed only during test recoveries.

- Some steps are always skipped during test recoveries.

Understanding the steps, their order, and the context in which they run is important when you customize 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, groups of virtual machines are started according to priority. Before a priority group is started, all machines in the next-higher priority group must have recovered or failed to recover. As long as dependencies have been met, the recovery engine attempts to power on as many virtual machines in parallel as vCenter supports.

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.

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 Not Executed During a Test

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

Clean up steps are performed after a recovery plan test. The steps begin executing after you respond to the prompt that appears after the test finishes.

- 1 Power off each recovered virtual machine.
- 2 Replace recovered virtual machines with placeholders, preserving their identity and configuration information.
- 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 89.

- 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 8-1](#) lists the environment variables that are available to all command steps.

Table 8-1. Environment Variables Available to All Command Steps

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

The environment variables listed in [Table 8-2](#) are also set if the command step is executing on a recovered virtual machine.

Table 8-2. Environment Variables Available to Command Steps Running on Recovered Virtual Machines

Name	Value	Example
<i>VMware_VM_Uuid</i>	UUID used by vCenter to uniquely identify this virtual machine.	4212145a-eeae-a02c-e525-ebba70b0d4f3
<i>VMware_VM_Name</i>	Name of this virtual machine, as set at the protected site.	My New Virtual Machine
<i>VMware_VM_Ref</i>	Managed object ID of the virtual machine.	vm-1199
<i>VMware_VM_GuestName</i>	Name of the guest OS as defined by the VIM API.	otherGuest
<i>VMware_VM_GuestIp</i>	IP address of the virtual machine, if known.	192.168.0.103
<i>VMware_VM_Path</i>	Path to this virtual machine's VMDK file.	[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 Click **Recovery Plans**, and click the plan that you want to customize.
- 2 Click the **Recovery Steps** tab.
- 3 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 click **Export**. To edit the properties of the plan, right-click any step and click **Edit Recovery Plan**.

Specify Virtual Machine Recovery Priority

By default, all virtual machines in a new recovery plan are members of the level 3 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 77.
- 2 To display the virtual machines in the normal priority group, expand the Recover Normal Priority Virtual Machines step.

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

Custom Recovery Steps

Custom recovery steps provide a way to run commands or to present messages. Custom recovery steps are executed during the recovery process and can be completed either on the SRM server machine or in a virtual machine that is being recovered.

You have several categories of custom recovery steps available:

- **Command recovery steps.** These custom recovery steps execute commands that run in their own processes. Command custom recovery steps have two types.
 - **Top-Level Commands.** Executed on the SRM server. For example, you might use these commands to power on physical devices or redirect network traffic.
 - **Per-VM Commands.** Associated with newly recovered virtual machines during the recovery process. Use these commands to complete configuration tasks after powering on a virtual machine. You can execute the commands pre-power on or post-power on. Commands that are configured to run after the virtual machine is recovered (meaning post-power on), can be run either on the SRM server machine or within the newly recovered virtual machine.
- **Message prompts.** Present a message in the SRM user interface. You can use this message to pause the recovery workflow and provide information to the user running the recovery plan. For example, users might be instructed to perform some manual recovery process or to verify steps. The only action users can take in direct response to a prompt is to click OK, which dismisses the message and allows the recovery to continue.

During the reprotect for a recovery plan, all custom recovery steps are preserved in the recovery plan. However, if a recovery or test is done for the recovery plan after a reprotect, these custom recovery steps are executed on the new recovery site (the original protected site). Typically, you can use prompt custom recovery steps directly without any modifications. However, you might need to modify command custom recovery steps after the reprotect, if these commands contain site-specific information, such as network configuration.

Recovery Step Outcomes

SRM attempts to complete all custom recovery steps, but some might fail to complete. The response to failures to complete a recovery step varies based on the recovery step type.

- **Command custom recovery steps.** By default, SRM waits for these commands to complete for five minutes. You can customize each command timeout. After a command is completed, the next recovery step in the recovery plan is executed. However, failures of custom commands have two different behaviors on the flow of the recovery plan.
 - **Top-Level commands.** If a recovery step fails, the failure is logged, a warning is shown in the recovery plan, and subsequent custom recovery steps continue to be executed.
 - **Per-VM commands.** Per-VM custom recovery steps are executed in sets either before or after a virtual machine is powered on. If a command fails, the remaining Per-VM custom recovery steps in the set are not executed. For example, if there are five pre-power on commands, and the third command in the pre-power on set fails, the remaining two pre-power on commands are skipped and SRM does not power on the virtual machine nor run any post-power on commands.
- **Prompt custom recovery steps.** Prompt custom recovery steps cannot fail. However, the execution of the recovery plan is paused until the user dismisses the prompt.

Configure Top-Level Commands and Prompts

You can add top-level commands and prompts anywhere in the recovery plan.

Procedure

- 1 Click **Recovery Plans** in the left pane, and select a recovery plan.
- 2 Click the **Recovery Steps** tab.
- 3 Select one of the steps and click the **Add Step** icon, the **Edit Step** icon, or the **Remove Step** icon. Some icons might be grayed-out, depending on where you clicked in the recovery steps.
- 4 If you clicked Add or Edit, configure the step.
 - a Select the type of Step to be completed. Choose from a step on the SRM server or a prompt.
 - b In the **Name** field, type a name for the step.
 - c In the **Content** field, enter the commands to be completed.
 - d For steps that are not prompts, you can modify the **Timeout** information.
- 5 For new steps, choose whether the step is placed before or after the selected step.

Configure Per-Virtual Machine Commands and Prompts

You can configure custom recovery steps for pre-power on and post-power on for a virtual machine. These steps are associated with the protected virtual machine in the same way that customization information is. If the same virtual machine is shared between plans, the commands and prompts are the same.

Procedure

- 1 Click **Recovery Plans** in the left pane, and select a recovery plan.
- 2 Click the **Virtual Machines** tab.

- 3 Select a virtual machine and click **Configure**.
- 4 Select Pre-Power On Steps or Post Power On Steps in the left pane, and click **Add**, **Edit**, or **Remove**.
- 5 If you click Add or Edit, configure the step.
 - a Select the type of Step to be completed: a step on the SRM server, a step on the virtual machine, or a prompt.
Complete pre-power on steps on the SRM server.
 - b In the **Name** field, type a name for the step.
 - c In the **Content** field, enter the commands to be executed.
 - d (Optional) For steps that are not prompts, modify the **Timeout** information.

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 Click **Recovery Plans** in the left pane, and click the plan that you want to customize.
- 2 Click the **Recovery Steps** tab.
- 3 Click the **Virtual Machines** tab.
- 4 Right-click a virtual machine in the list, and click **Configure**.
- 5 Click **IP Settings**.

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

- 6 Select the appropriate entry to add a message or command step that executes before the machine is powered on.
- 7 Select the appropriate entry 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. See [“Guidelines for Writing Command Steps,”](#) on page 76.

The customizations that you specify become associated with the protected virtual machine. As a result, the settings are shared between all recovery plans that apply to this virtual machine.

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

Customize IP Properties For an Individual Virtual Machine

You can customize IP settings for individual virtual machines for both the protected site and the recovery site. IP settings for the recovery site are used during the recovery or test from the protection site to the recovery site. IP settings for the protection site are used after reprotect during the recovery or test from the original recovery site to the original protection site. Customization settings are associated with protected virtual machines. As a result, if the same protected virtual machine is a part of multiple recovery plans, then all recovery plans will use a single copy of the customization settings.

IP customization in SRM 5.0 supports:

- IPv4 and IPv6 addressing.

- Different IP customizations for each site.
- DHCP, Static IPv4, or Static IPv6 addressing.
- Address customization for Windows and Linux virtual machines.
- Customizing multiple NICs for each virtual machine.

You can also apply IP customizations to multiple virtual machines.

SRM includes the ability for the vSwitches to be DVS based and, therefore, span hosts. You can create recovery plans that recover virtual machines across multiple recovery site hosts, contained within a quarantined test network. If you accept the default test network configured as **Auto**, then virtual machines that are recovered across hosts are placed in their own test bubble network. Each test bubble switch is isolated between hosts. As a result, virtual machines in the same plan are isolated when the recovery completes. If you want the virtual machines to be able to communicate, establish and select DVS switches or VLANs. Establishing an isolated VLAN that connects all hosts to each other but not to a production network makes it possible to more realistically test a recovery. To achieve connectivity among recovery hosts but isolation from the production network, use the following recommendations:

- Create DVS switches that are connected to an isolated VLAN that is private. Such a VLAN allows hosts and virtual machines to be connected, but to be isolated from production virtual machines. Use a naming convention that clearly designates the DVS is for testing use, and then select this DVS in the recovery plan test network column.
- Create test VLANs, providing no route back to protected site, on a physical network. Trunk test VLANs to recovery site vSphere clusters and create virtual switches for test VLAN IDs, again using a clear naming convention to identify these switches as being for testing use. Select these switches from the test recovery network column in recovery plan editor.

IP customization is configured as part of the process of configuring virtual machine recovery properties. If a NIC is not customized, it uses the same IP settings from the other site.

Procedure

- 1 Click **Recovery Plans**, and click the plan that you want to customize.
- 2 Click the **Virtual Machines** tab and click **Configure VM**.
- 3 Select the NIC for which you will modify IP Settings.
- 4 To customize settings, enable the **Customize IP settings during recovery** option.
- 5 Click **Configure Protection** or **Configure Recovery**, depending on which set of IP settings you want to configure.
- 6 Click the **General** tab to configure settings.
 - a Choose the type of addressing to be used.
Available options include DHCP, static IPv4, or static IPv6.
 - b For static addresses, enter an IP address, subnet information, and gateway server addresses.
Alternately, if the virtual machine is powered on and has VMware Tools installed, you can click **Update** to import current settings configured on the virtual machine.
- 7 Click the **DNS** tab to configure DNS settings.
 - a Choose how DNS servers are found.
You can use DHCP to find DNS servers or you can specify primary and alternate DNS servers.
 - b Enter a DNS suffix and click **Add** or select an existing DNS suffix and click **Remove, Move Up, or Move Down**.

- 8 Click the **WINS** tab to enter primary and secondary WINS addresses.
The **WINS** tab is available only when configuring DHCP or IPv4 addresses for Windows virtual machines.
- 9 Repeat [Step 5](#) through [Step 8](#) to configure recovery or protection settings, if required.
For example, if you configured IP settings for the protected site, you might want to configure settings for the recovery site.
- 10 Repeat the configuration process for other NICs, as required, beginning by choosing another NIC as described in step [Step 3](#).

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 information can help you understand the network requirements of a recovery plan.

Because the IP address mapping 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.

Understanding Customizing IP Properties for Multiple Virtual Machines

Manually configuring IP settings for many virtual machines at a recovery site can be time consuming and errors in configuration might occur. To facilitate the configuration process, SRM includes `dr-ip-customizer.exe`, which is installed in the `bin` subdirectory of the SRM installation directory. Use this tool to create or apply comma separated value (CSV) files containing information about networking configurations.

A challenge of representing virtual machine network configurations in a CSV file is that virtual machine configurations include hierarchical information. For example, a single virtual machine may contain multiple adapters, and each adapter may have multiple listings for elements such as gateways. The CSV format does not provide a system for hierarchical representations. As a result, each row in the CSV file that the DR IP Customizer generates may provide some or all pertinent information for a virtual machine.

For a virtual machine with a simple network configuration, all information can be included in a single row. In the case of a more complicated virtual machine, multiple rows might be required. In the example of virtual machines with multiple network cards or multiple gateways require multiple rows, each row in the CSV file includes identification information describing which virtual machine and adapter the information applies to. Information is aggregated together to be applied to the appropriate virtual machine.

To apply IPv6 customizations to virtual machines, the machine running `dr-ip-customizer` must have IPv6 enabled. IPv6 is not enabled by default on Windows XP or Windows 2003. Attempts to apply IPv6 address customizations from a machine without IPv6 enabled fail when the DR IP Customizer validates the CSV input.

Customize IP Properties for Multiple Virtual Machines

Use the `dr-ip-customizer.exe` tool on a computer with access to vCenter Servers in your environment to specify IP properties for any or all of the virtual machines in a recovery plan by editing a file that the tool generates. Configure IP settings for both protected and recovery sites so failback operations are easier to configure.

Procedure

- 1 Open a command shell on the SRM server host.
- 2 Change directory to `C:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\bin`.
- 3 Run the `dr-ip-customizer.exe` command.

The following is an example of how you might run the tool.

```
dr-ip-customizer.exe --cfg ..\config\vmware-dr.xml --cmd generate --out
"C:\MassIPCustCSVs\MassIPCust-generate-output.csv" --vc vc04.eng.example.com
```

- 4 Edit the generated file that was created at `C:\MassIPCustCSVs\MassIPCust-generate-output.csv` to customize IP properties for the virtual machines in the recovery plan.
- 5 Run `dr-ip-customizer.exe` to apply the customized IP properties.

For example, to apply settings from a CSV file, run the following command.

```
dr-ip-customizer.exe --cfg ..\config\vmware-dr.xml --cmd apply --csv
"C:\MassIPCustCSVs\MassIPCust-ipv6.csv" --vc vc04.eng.example.com
```

You can include a `--verbose` option on any `dr-ip-customizer.exe` command line to log additional diagnostic messages. Always use the same vCenter Server for `apply` and `drop` commands as the one used to generate the CSV. Virtual machine IDs for protected virtual machines are different at each site, so the CSV generated at one site should not be applied at a different site.

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

Make sure to consistently use the same `--vc` setting. You can work with either protected or recovery vCenter Server, but use the same one for both the generate and apply operations.

Dr Ip Customizer Reference

The `dr-ip-customizer.exe` tool includes several options. Whenever you use the DR IP Customizer, you must specify the location of the SRM server configuration XML file and the action to take.

DR IP Customizer Tool Syntax

The syntax of the DR IP Customizer is as follows:

```
dr-ip-customizer.exe --cfg <DR server configuration XML> --cmd <apply/drop/generate> [--csv
<existing CSV File>] [--out <New CSV file to be generated>] [--vc <VC Server Address>] [--ignore-
thumbprint] [--extra-dns-columns] [--verbose]
```

The options available with DR IP Customizer are:

- `-h [--help]` Display usage info.
- `--cfg path` Path to application XML configuration file.

- `--cmd arg` Command to execute:
 - `apply` Applies the network customization settings from the input CSV file to the recovery plans on the SRM servers.
 - `generate` Generates a basic CSV file for the all virtual machines in the recovery plans on the SRM servers.
 - `drop` Removes virtual machine recovery settings from virtual machines specified by the input CSV file.
- `--csv arg` Path to the CSV file. Read as input for the `apply` and `drop` commands.
- `-o [--out] arg` Output CSV file to use for the `generate` command. Will overwrite any existing contents.
- `--vc arg` VMware vCenter Server host name. The virtual machine Ids are different at each site.
- `-i [--ignore-thumbprint]` Ignore the server thumbprint confirmation prompt.
- `-e [--extra-dns-columns]` Must be specified if the input CSV file contains extra columns for DNS information.
- `-v [--verbose]` Enable verbose output.

CSV Structure Reference

You may omit values if no setting is needed. The columns used to apply values to virtual machine network configurations are as follows:

- **Adapter ID:** A unique identifier used to collect information from multiple rows together for application to a single virtual machine. Note that settings for entry for VM ID 0 are global settings. This means that values specified for VM ID 0 are applied to virtual machines that do not have values to override the global settings. This can be useful, for example, for configuring DNS server information or gateway information.
- VM Name
- vCenter Server
- Adapter ID
- DNS Domain
- Net BIOS
- Primary WINS
- Secondary WINS
- IP Address: This field is required. You can enter an address or the value *DHCP* .
- Subnet Mask
- Gateway(s)
- IPv6 Address
- IPv6 Subnet Prefix length
- IPv6 Gateway(s)
- DNS Server(s)
- DNS Suffix(es)

Guidelines

Follow these guidelines.

- You can export protected virtual machine information to use as a template, rather than author a new CSV file manually. The generated file can provide a structure that you can modify to suit the needs of your environment.
- Some restrictions exist on the ways in which DNS settings are applied.
 - In Windows
 - DNS Suffix : Global settings for all adapters
 - DNS Server : per-adapter setting (If user entered it in 'Adapter ID' 0 row, it will be treated as a global setting)
 - In Linux
 - DNS Suffix : Global settings for all adapters
 - DNS Server : Global settings for all adapters
- The DR IP Customizer validates that WINS settings are applied only to Windows virtual machines, but it does not validate NetBIOS settings.
- Use the minimum number of rows possible for each adapter.
- Use either 1 IPv4 or 1 IPv6 per NIC. Virtual machines may support multiple addresses.
- The IPv4 field cannot be left blank. In the IPv4 field you can specify DHCP if a static IPv6 address is to be used.
- Commas are not allowed in any field.
- 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. To ensure the additional row is associated with the intended virtual machine, copy the VM ID, VM Name, vCenter Server, Adapter ID column values.
- The NetBIOS column, if not left empty, must contain one of the following strings: disableNetBIOS, enableNetBIOS, or enableNetBIOSViaDhcp.

Example

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

```
VM ID,VM Name,VC Name,Adapter ID,Primary WINS,Secondary WINS,IP Address,Subnet
Mask,Gateway(s),DNS Server(s),DNS Suffix(es)
protected-vm-10301,example-vm-3-win,vc04.eng.example.com,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. All the other columns are empty.

After creating the csv file, the contents can be modified to configure settings for three network adapters for the virtual machine protected-vm-10301.

Table 8-3. Sample CSV File Contents

VM ID	VM Name	vCenter Server	Adapter ID	Primary Wins	Secondary Wins	IP Address	Subnet Mask	Gateway(s)	IPv6 Address	IPv6 Subnet Prefix Length	IPv6 Gateway(s)	DNS Server(s)	DNS Suffix(es)
protected-vm-10301	example-win	vc04.eng.example.com	0										example.com
protected-vm-10301	example-win	vc04.eng.example.com	0										example.com
protected-vm-10301	example-win	vc04.eng.example.com	1	192.168.1.5	192.168.1.6	192.168.1.9	255.255.255.0	192.168.1.1	dhcp				
protected-vm-10301	example-win	vc04.eng.example.com	1									192.168.1.1	6
protected-vm-10301	example-win	vc04.eng.example.com	2	192.168.1.5	192.168.1.6	192.168.1.10	255.255.255.0	192.168.1.1				192.168.1.1	7
protected-vm-10301	example-win	vc04.eng.example.com	3			DHCP			::ffff:192.0.0.12	32	::ffff:192.0.0.1		

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 or Mapping Missing.

If you are using array based replication, 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 add a device that is stored 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. You can use Storage VMotion to move a virtual machine to a datastore that is part of another protection group, though.

Procedure

- 1 Click **Protection Groups** in the left pane, navigate to the protection group that includes the virtual machine that you want to configure.
- 2 On the Virtual Machines tab, select a virtual machine and click **Configure Protection**.
- 3 In the Virtual Machine Properties window, review and configure properties as needed.
 - a Click **Recovery Folder** and specify an alternate destination folder.
 - b Click **Recovery Resource** and specify an alternate resource pool in which to place the recovered virtual machine.
 - c If configuring protection for a template, click **Recovery Host** and specify an alternate host to which to recover the virtual machine.
This step is only applicable for templates.
 - d Click **Recovery Network** and specify an alternate recovery network to which to restore the virtual machine.
- 4 Click **OK** to apply the new configuration to the selected virtual machine.

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 settings of the placeholders or move them to a different folder or resource pool.

Procedure

- 1 Click **Protection Groups**, and navigate to the protection group that includes the virtual machine that you want to configure.
- 2 On the Virtual Machines page, right-click a virtual machine and click **Configure Protection**.
If you established inventory mappings, they are applied.
- 3 Configure mappings as needed.

For most virtual machines, you can change the Folder and Compute Resource mappings. For more information, see [“Configure Protection for a Virtual Machine or Template,”](#) on page 86.

Configure SRM Alarms

SRM adds feature-specific alarms to the ones supported by vCenter. You can configure SRM alarms to send an email notification, send an SNMP trap, or run a script on the machine that contains the vCenter Server.

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 the action to take when an event triggers the alarm. None of the SRM alarms are configured by default to take any action. To enable actions for any of them, you must configure them to do so.

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

Procedure

- 1 In the left pane, click **Sites**, and select a site.
- 2 Click the **Alarms** tab to display the list of SRM alarms.
- 3 Right-click an alarm and click **Edit Settings**.
- 4 Click the **Actions** tab.
- 5 Click **Add** to add an action. The default action for every event is **Send a notification e-mail**.
- 6 (Optional) To change the default action, click it and select a different action from the drop-down list.
The default action for every event is **Send a notification e-mail**.
- 7 Click the **General** tab.
- 8 Check the **Enable this alarm** option to enable the actions for the alarm.

Working with Advanced Settings

Using the Advanced Settings, you can view or change many custom settings for the SRM service. The Advanced Settings dialog box provides a way for a user with adequate privileges to change a number of default values that affect the operation of various SRM features.

Procedure

- 1 Click **Sites** in the left pane, and right-click the site whose settings you want to change.
- 2 Click **Advanced Settings**.
- 3 , Click a setting category.
- 4 In the category window, make your changes.
- 5 Click **OK** to save your changes.
- 6 Repeat the procedure as needed at the recovery site.

Guest Customization Settings

Change these settings only if instructed to do so by VMware Support.

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. You might do this if tasks were failing to complete due to insufficient time.

Several types 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 finish.

Command line timeout By default, SRM allows 300 seconds for a command step to finish. 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 finish 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 Click **Sites** in the left pane, and right-click the site whose settings you want to change.
- 2 Click **Advanced Settings**.
- 3 Click **Recovery**.
- 4 Modify recovery site settings.
 - To change the command-line timeout, enter a new value in the **recovery.calloutCommandLineTimeout** text box. The new value applies to all command steps.
 - To change the customization timeout, enter a new value in the **recovery.customizationTimeout** text box.
 - To change the default priority, enter a new value in the **recovery.defaultPriority** text box.
 - To change the power off timeout, enter a new value in the **recovery.powerOffTimeout** text box. The new time-out value applies to power-off tasks for virtual machines at the protected site.
 - To enable or disable forced failover, select or deselect the **recovery.forcedFailover** check box.

NOTE Forced failover is only available in SRM 5.0.1 and later.

- To change the delay before powering on a virtual machine, enter a new value in the **recovery.powerOnDelay** text box. The new value applies to power-on tasks for virtual machines at the protected site.
 - To change the power state change timeout, enter a new value in the **recovery.powerOnTimeout** text box. The new power-on value applies to power-on tasks for virtual machines at the protected site.
 - Enable or disable **recovery.preserveCustPkg**.
 - Enable or disable **recovery.skipGuestShutdown** to complete or skip the guest shutdown.
- 5 Click **OK** to save your changes.

Change Array-Based Storage Provider Settings

For array-based replication, 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 Click **Sites** in the left pane, and right-click the site whose settings you want to change.
- 2 Click **Advanced Settings**.
- 3 Click **storageProvider**.
- 4 Modify the SAN provider settings.
 - To force removal, upon successful completion of a recovery, of the snap-xx prefix applied to recovered datastore names, select the **storageProvider.fixRecoveredDatastoreNames** check box.
 - To repeat host scans during testing and recovery, enter a new value in the **storageProvider.hostRescanRepeatCnt** text box. Some storage arrays require more than one rescan, for example to discover the snapshots of failed-over LUNs.

NOTE The `storageProvider.hostRescanRepeatCnt` parameter is available in SRM 5.0.1 and later. It is not available in SRM 5.0.

 - To change the interval that SRM waits for each HBA rescan to complete, enter a new value in the **storageProvider.hostRescanTimeoutSec** text box.
- 5 Click **OK** to save your changes.

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 your needs.

Procedure

- 1 Click **Sites** in the left pane, and right-click the local site whose settings you want to change.
- 2 Click **Advanced Settings**.
- 3 Click **localSiteStatus**.
- 4 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** text box.
 - To change the name for the local site, enter a new value in the **localSiteStatus.displayName** text box.
 - 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** text box.
 - 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** text box.
 - 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** text box.
 - To change the amount of free memory that causes SRM to raise a low memory event, enter a new value in the **localSiteStatus.minMemory** text box.
- 5 Click **OK** to save your changes.

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 your needs.

Procedure

- 1 Click **Sites** in the left pane, and right-click the remote site whose settings you want to change.
- 2 Click **Advanced Settings**.
- 3 Click **remoteSiteStatus**.
- 4 Modify the settings.
 - To change the number of failed pings before posting a site down event, enter a new value in the **remoteSiteStatus.panicDelay** text box.
 - To change the number of remote site status checks (pings) to try before declaring the check a failure, enter a new value in the **remoteSiteStatus.pingFailedDelay** text box.
 - 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.pingInterval** text box.
- 5 Click **OK** to save your changes.

Change Storage Settings

You can adjust SRM storage settings.

Procedure

- 1 Click **Sites** in the left pane, right-click a site, and click **Advanced Settings**.
- 2 In the navigation pane of the Advanced Settings window, click **Storage**.
- 3 Modify the storage settings as needed.
 - To change SRA update timeout, enter a new value in the **storage.commandTimeout** field.
 - To change the maximum number of concurrent SRA operations, enter a new value in the **storage.maxConcurrentCommandCnt** field.
 - To change the minimum amount of time in seconds between datastore group computations, enter a new value in the **storage.minDsGroupComputationInterval** field.
 - To change the interval between status updates for ongoing data synchronization operations, enter a new value in the **storage.querySyncStatusPollingInterval** field.
 - To change the interval between storage array discovery checks, enter a new value in the **storage.storagePingInterval** field.
 - To change the maximum amount of time permitted for data synchronization operations to complete, enter a new value in the **storage.syncTimeout** field.
- 4 Click **OK** to save your changes and close the Advanced Settings window.

Change Replication Setting

You can adjust replication settings to modify how long SRM waits for virtual machine placeholder creation to complete.

Procedure

- 1 Click **Sites**, and right-click the site whose settings you want to change.
- 2 Click **Advanced Settings**.
- 3 Click **replication**.
- 4 Change the **replication.placeholderVmCreationTimeout** setting to modify the number of seconds to wait when creating a placeholder virtual machine.
- 5 Click **OK** to save your changes.

Change vSphere Replication Settings

You can adjust the settings for vSphere Replication (VR) to specify different recovery point objectives.

Procedure

- 1 Click **Sites** in the left pane, and right-click the site whose settings you want to change.
- 2 Click **Advanced Settings**.
- 3 Click **vrReplication**.
- 4 Modify default recovery point objective (RPO) settings.
- 5 Click **OK** to save your changes.

Troubleshooting SRM

If you have problems with replication, site pairing, or guest customization, you can 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 appear in error dialog boxes or appear in the Recent Tasks window. Most errors also generate an entry in an SRM log files. 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:

- [“Events and Alarms,”](#) on page 93
- [“Collecting SRM Log Files,”](#) on page 102
- [“Features Are Unavailable When Deploying VRMS,”](#) on page 103
- [“OVF Package is Invalid and Cannot be Deployed,”](#) on page 103
- [“Connection Errors Between VRMS and SQL Cannot be Resolved,”](#) on page 103
- [“Configuration of the VRMS Database Fails with DB2 Databases,”](#) on page 104

Events and Alarms

SRM supports logging events and each event includes a corresponding alarm that can be triggered when the event occurs. This provides a way to track to health and functioning of your system and resolve potential issues before they impact the protection SRM provides.

Alarms are enabled using the Alarm Manager in the vSphere Client. SRM also includes SNMP traps.

Events can be categorized according to functional areas:

- Site Status
- Protection Group Events
- Recovery Events
- SNMP Traps
- Storage and Storage Provider Events
- Licensing Events
- Permissions Events

Connection Monitor Algorithm

Connection related events are produced according to this algorithm:

- 1 When administrative connection between two paired SRM servers is established SRM server that initiated the connection raises RemoteSiteUpEvent.
- 2 When SRM detects that a monitored connection is down it starts a periodic connection check. Time interval for these pings is configurable through the remoteSiteStatus.pingInterval advanced setting.
 - a Connection monitor skips a configurable number of failed pings. You can control this number through the remoteSiteStatus.pingFailedDelay advanced setting.
 - b When the number of skipped failed pings exceeds the allowed value SRM posts RemoteSitePingFailedEvent.
 - c When the number of skipped failed pings becomes more than another configurable limit SRM raises RemoteSiteDownEvent on each failed ping and stops raising RemoteSitePingFailedEvent events. The second limit of failed pings is configurable through remoteSiteStatus.panicDelay advanced setting.
 - d SRM continues to raise connection down events until connection is reestablished.

Site Status Events

Site status events provide information about the status and connection between protected and recovery sites.

Table 9-1. Site Status Events

Event Key	Event Description	Cause
UnknownStatusEvent	SRM server status is not available	
RemoteSiteDownEvent	Remote SRM site is down	This event is signaled when the SRM server loses its connection with the remote SRM server.
RemoteSitePingFailedEvent	Remote SRM site is not responding to pings	This may be due to failures at the site or network connectivity.
RemoteSiteCreatedEvent	Remote SRM site has been created	This occurs when a site is established.
RemoteSiteUpEvent	Remote SRM site is responsive	This occurs when the SRM server re-establishes its connection with the remote SRM server.
RemoteSiteDeletedEvent	Remote SRM site has been deleted.	

Protection Group Events

Protection Group events provide information about actions and status related to protection groups.

These events have three categories:

- Protection Group Replication Informational Events
- Protection Group Replication Warning Events
- Protection Group Replication Error Events

Table 9-2. Protection Group Replication Informational Events

Event Key	Event Description	Cause
ProtectionGroup > CreatedEvent	Created protection group.	Posted on both vCenter Servers in the completion of the Commit phase of creating a protection group.
ProtectionGroup > RemovedEvent	Removed protection group.	Posted on both vCenter Servers in the completion of the Commit phase of removing a protection group.
ProtectionGroup > ReconfiguredEvent	Reconfigured protection group.	Posted on both vCenter Servers in the completion of the Commit phase of reconfiguring a protection group.
ProtectedVmCreatedEvent	Virtual machine in group is configured for protection.	Posted on both vCenter Servers in the completion of the Commit phase of the protection of a virtual machine.
ProtectedVmRemovedEvent	Virtual machine in group is no longer configured for protection.	Posted on both vCenter Servers in the completion of the Commit phase of unprotecting a virtual machine.
ProtectedVmReconfiguredProtectionSettingsEvent	Reconfigured protection settings for virtual machine.	Posted on both vCenter Servers in the completion of the Commit phase of reconfiguring virtual machine protection settings.
ProtectedVmReconfiguredRecoveryLocationSettingsEvent	Reconfigured recovery location settings for virtual machine.	Posted on the Production site vCenter Server only on the successful completion of reconfiguring the recovery location settings for a protected virtual machine.
PlaceholderVmCreatedEvent	The placeholder virtual machine was created in the VMware vCenter Server inventory.	Posted on the Recovery site vCenter Server only when we create the placeholder virtual machine as a result of protection, repair.
PlaceholderVmCreatedFromOldProductionVmEvent	The placeholder virtual machine was created in the VMware vCenter Server inventory using the identity of the old production virtual machine.	Posted on the Recovery site vCenter Server only when we create the placeholder virtual machine as a result of swapping for the old production virtual machine during or after re-protection.

Table 9-3. Protection Group Replication Warning Events

Event Key	Event Description	Cause
VmNotFullyProtectedEvent	Virtual machine in group: One or more devices need to be configured for protection.	Posted on the Production site vCenter Server only upon device handling updating the recovery location settings with a non-empty unresolvedDevices set. This can be triggered by changes to the production virtual machine or during re-protection of a virtual machine.
PlaceholderVmUnexpectedlyDeletedEvent	Virtual machine in group: The placeholder virtual machine was removed from the VMware vCenter Server inventory.	Posted on the Recovery site vCenter Server only when we detect that the placeholder virtual machine was unexpectedly deleted or removed from the vCenter inventory.

Table 9-4. Protection Group Replication Error Events

Event Key	Event Description	Cause
ProductionVmDeletedEvent	Virtual machine in group: The production virtual machine has been removed from the virtual machineware vCenter Server inventory.	Posted when we detect that the protected virtual machine's production virtual machine has been deleted or removed from the vCenter inventory.
ProductionVmInvalidEvent	Virtual machine in group: Cannot resolve the file locations of the production virtual machine for replication.	Posted whenever we handle device or recovery location changes but notice that the provider cannot find the production virtual machine files in order to replicate them.

Recovery Events

Recovery events provide information about actions and status related to recovery processes.

Table 9-5. Recovery Events

Event Key	Event Description	Cause
RecoveryVmBegin	Recovery Plan has begun recovering the specified virtual machine.	Signaled when the recovery virtual machine was successfully created. If some error occurred before the virtual machine ID is known the event is not fired.
RecoveryVmEnd	Recovery Plan has completed recovering the virtual machine.	Signaled after the last post-power on script has completed, or after a recovery-stopping error has occurred for the virtual machine.
RecoveryPlanCreate	Recovery Plan has been created.	Signaled when a new plan is created or cloned. It will be sent to each vCenter server where the plan is hosted.
RecoveryPlanDestroy	Recovery Plan has been destroyed.	Signaled when a plan has been deleted from the site. Note that on the site where the plan has been requested to deleted there can be a significant delay, while it waits for the plan to be deleted at the other site. It will be sent to each vCenter server where the plan is hosted.
RecoveryPlanEdit	Recovery Plan was changed.	
RecoveryPlanExecuteTestBegin	Recovery Plan has begun a test.	Signaled on the recovery site when a recovery test is initiated.
RecoveryPlanExecuteTestEnd	Recovery Plan has completed a test.	Signaled on the recovery site when a recovery test has completed. If an error occurred it is available as described.
RecoveryPlanExecuteCleanupBegin	Recovery Plan has begun a test cleanup.	Signaled on the recovery site when a test cleanup is initiated.
RecoveryPlanExecuteCleanupEnd	Recovery Plan has completed a test cleanup.	Signaled on the recovery site when a test cleanup has completed. If an error occurred it is available as described.
RecoveryPlanExecuteBegin	Recovery Plan has begun a recovery.	Signaled on the recovery site when a recovery is initiated.
RecoveryPlanExecuteEnd	Recovery Plan has completed a recovery.	Signaled on the recovery site when a recovery has completed. If an error occurred it is available as described.

Table 9-5. Recovery Events (Continued)

Event Key	Event Description	Cause
RecoveryPlanExecuteReprotectBegin	Recovery Plan has begun a reprotect operation.	Signaled on the recovery site when a reprotect is initiated.
RecoveryPlanExecuteReprotectEnd	Recovery Plan has completed a reprotect operation.	Signaled on the recovery site when a reprotect has completed. If an error occurred it is available as described.
RecoveryPlanPromptDisplay	Recovery Plan is displaying a prompt and is waiting for user input.	Signaled on the recovery site when a prompt step is encountered. The key is a unique identifier for the prompt.
RecoveryPlanPromptResponse	Recovery Plan has received an answer to its prompt.	Signaled on the recovery site when a prompt step is closed.
RecoveryPlanServerCommandBegin	Recovery Plan has started to run a Command on the SRM server machine.	Signaled on the recovery site when SRM has started to run a Callout Command on the SRM server machine.
RecoveryPlanServerCommandEnd	Recovery Plan has completed the execution of a Command on the SRM server machine.	Signaled on the recovery site when SRM has finished running a Callout Command on the SRM server machine.
RecoveryPlanVmCommandBegin	Recovery Plan has started to run a Command on a recovered virtual machine.	Signaled on the recovery site when SRM has started to run a Callout Command on a recovered virtual machine.
RecoveryPlanVmCommandEnd	Recovery Plan has completed the execution of a Command on a recovered virtual machine.	Signaled on the recovery site when SRM has finished running a Callout Command on a recovered virtual machine.

SNMP Traps

SRM sends SNMP traps to community targets defined in vCenter. You can configure them using the vSphere Client. When you enter localhost or 127.0.0.1 as a target host for SNMP traps, SRM uses the IP address or host name of the vSphere server as configured by the SRM installer.

SNMP traps for SRM 5.0 are backward compatible with SRM 4.0 and later releases.

Table 9-6. SNMP Traps

Type	Description	Content
RecoveryPlanExecuteTestBeginTrap	This trap is sent when a Recovery Plan starts a test.	SRM site name, recovery plan name, recovery type, execution state.
RecoveryPlanExecuteTestEndTrap	This trap is sent when a Recovery Plan ends a test.	SRM site name, recovery plan name, recovery type, execution state, result status.
RecoveryPlanExecuteCleanupBeginTrap	This trap is sent when a Recovery Plan starts a test cleanup.	SRM site name, recovery plan name, recovery type, execution state.
RecoveryPlanExecuteCleanupEndTrap	This trap is sent a Recovery Plan ends a test cleanup.	SRM site name, recovery plan name, recovery type, execution state, result status.
RecoveryPlanExecuteBeginTrap	This trap is sent when a Recovery Plan starts a recovery.	SRM site name, recovery plan name, recovery type, execution state.
RecoveryPlanExecuteEndTrap	This trap is sent when a Recovery Plan ends a recovery.	SRM site name, recovery plan name, recovery type, execution state, result status.

Table 9-6. SNMP Traps (Continued)

Type	Description	Content
RecoveryPlanExecuteReprotectBeginTrap	This trap is sent when SRM starts the reprotect workflow for a Recovery Plan.	SRM site name, recovery plan name, recovery type, execution state.
RecoveryPlanExecuteReprotectEndTrap	This trap is sent when SRM has finished the reprotect workflow for a Recovery Plan.	SRM site name, recovery plan name, recovery type, execution state, result status.
RecoveryVmBeginTrap	This trap is sent when a Recovery Plan starts recovering a virtual machine.	SRM site name, recovery plan name, recovery type, execution state, virtual machine name, virtual machine UUID.
RecoveryVmEndTrap	This trap is sent when a Recovery Plan has finished recovering a virtual machine.	SRM site name, recovery plan name, recovery type, execution state, virtual machine name, virtual machine UUID, result status.
RecoveryPlanServerCommandBeginTrap	This trap is sent when a Recovery Plan starts the execution of a command callout on SRM server's machine.	SRM site name, recovery plan name, recovery type, execution state, command name.
RecoveryPlanServerCommandEndTrap	This trap is sent when a Recovery Plan has finished the execution of a command callout on SRM server's machine.	SRM site name, recovery plan name, recovery type, execution state, command name, result status.
RecoveryPlanVmCommandBeginTrap	This trap is sent when a Recovery Plan starts the execution of a command callout on a recovered virtual machine.	SRM site name, recovery plan name, recovery type, execution state, command name, virtual machine name, virtual machine UUID.
RecoveryPlanVmCommandEndTrap	This trap is sent when a Recovery Plan has finished the execution of a command callout on a recovered virtual machine.	SRM site name, recovery plan name, recovery type, execution state, command name, virtual machine name, virtual machine UUID, result status.
RecoveryPlanPromptDisplayTrap	This trap is sent when a Recovery Plan requires user input before continuing.	SRM site name, recovery plan name, recovery type, execution state, prompt string.
RecoveryPlanPromptResponseTrap	This trap is sent when a Recovery Plan no longer requires user input before continuing.	SRM site name, recovery plan name, recovery type, and execution state.

Storage and Storage Provider Events

Storage and storage provider events provide information about actions and status related storage or storage providers.

Table 9-7. SRA Events

Type	Description	Content
StorageAdapterLoaded	Loaded the specified SRA.	SRM detected new SRA either during startup or during user-initiated SRAs reload.
StorageAdapterReloadFailed	Failed to load SRA from the specified path.	SRM failed to reload previously known SRA either during startup or during user-initiated SRAs reload.
StorageAdapterChanged	Loaded new version of the specified SRA.	SRM detected that previously known SRA was upgraded.

Table 9-8. Array Manager Events

Type	Description	Content
StorageArrayManagerAdded	Created the specified array manager using the specified SRA.	User added an Array Manager.
StorageArrayManagerRemoved	Deleted the specified array manager.	User removed an Array Manager.
StorageArrayManagerReconfigured	Reconfigured the specified array manager.	User edited Array Manager properties.
StorageArrayManagerPingOk	Ping for the specified array manager succeeded.	SRM server successfully pinged an Array Manager.
StorageArrayManagerPingFailed	Failed to ping the specified array manager.	An error occurred during Array Manager ping.

Table 9-9. Array Pair Events

Type	Description	Content
StorageArrayPairDiscovered	Discovered replicated array pair with Array Manager.	User created Array Manager which discovered replicated array pairs.
StorageArrayPairEnabled	Enabled replicated array pair with Array Manager.	User enabled an Array Pair.
StorageArrayPairDisabled	Disabled replicated array pair with Array Manager.	User disabled an Array Pair.
StorageArrayPairPingOk	Ping for replicated array pair succeeded.	SRM server successfully pinged the array pair.
StorageArrayPairPingFailed	Failed to ping replicated array pair.	An error occurred during Array Pair ping.

Table 9-10. Datastore Events

Type	Description	Content
StorageDatastoreDiscovered	Discovered replicated datastore.	SRM server discovered replicated datastore.
StorageDatastoreLost	Specified datastore is no longer replicated.	User turned off replication of storage devices backing the datastore.
StorageRdmDiscovered	Discovered replicated RDM attached to specified virtual machine.	SRM server discovered replicated RDM.
StorageRdmLost	RDM attached to specified virtual machine is no longer replicated.	User turned off replication of the LUN backing the RDM.

Table 9-11. Protection Events

Type	Description	Content
StorageProviderDatastoreProtected	Protected datastore in specified protection group.	User included datastore in new or existing protection group.
StorageProviderDatastoreUnprotected	Unprotected specified datastore.	User removed datastore from protection group or deleted protection group which contained this datastore.
StorageProviderVmDiscovered	Discovered replicated virtual machine.	User created virtual machine on a replicated datastore.
StorageProviderVmLost	Specified virtual machine is no longer replicated	User migrated virtual machine off of the replicated datastore.

Table 9-11. Protection Events (Continued)

Type	Description	Content
StorageProviderDatastoreProtectionMissing	Replicated datastore needs to be included in specified protection group.	See description.
StorageProviderDatastoreProtectionConflict	Replicated datastore needs to be included in specified protection group but is included in an alternate protection group.	See description.
StorageProviderDatastoreReplicationLost	Datastore included in specified protection group is no longer replicated.	User turned off replication for devices backing the datastore.
StorageProviderGroupProtectionRestored	Protection has been restored for specified protection group.	The previous (non-empty) issues of a protection group are cleared.
StorageProviderVmDatastoreProtectionMissing	Datastore used by virtual machine needs to be included in specified protection group.	See description.
StorageProviderVmDatastoreProtectionConflict	Datastore used by specified virtual machine needs to be added to specified protection group, but is currently in use by an alternate protection group.	See description.
StorageProviderVmDatastoreReplicationLost	Datastore used by specified virtual machine and included in specified protection group is no longer replicated.	See description.
StorageProviderVmProtectionRestored	Protection for specified virtual machine in specified protection group has been restored.	The previous (non-empty) issues for a protected virtual machine are cleared. The event will not be posted when issues related to non-protected virtual machine are cleared
StorageProviderGgSpansProtectionGroups	Specified consistency group spans specified protection groups.	See description.
StorageProviderCgDatastoreMissingProtection	Datastore from specified consistency group needs to be included in specified protection group.	See description.
StorageProviderDatastoreSpansConsistencyGroups	Datastore spans devices from different consistency groups.	See description.
StorageProviderNfsDatastoreUrlConflict	NFS datastores mounted from specified volume have different URLs mounted from the remote host. The remote path has the specified URL, while the datastore mounted from the other host has the specified URL.	The same NFS volume is mounted using the different IP addresses of the same NFS server in two different datastores.

Licensing Events

Licensing events provide information about changes in SRM licensing status.

Table 9-12. Licensing Events

Type	Description	Content
LicenseExpiringEvent	The SRM License at the specified site expires in the specified number of days.	Every 24 hours, non-evaluation, expiring licenses are checked for the number of days left. This event is posted with the results.
EvaluationLicenseExpiringEvent	The SRM Evaluation License at the specified site expires in the specified number of days.	Every 24 hours, evaluation licenses are checked for the number of days left. This event is posted with the results.
LicenseExpiredEvent	The SRM license at the specified site license has expired.	Every 30 minutes, expired (non-evaluation) licenses will post this event.
EvaluationLicenseExpiredEvent	The SRM Evaluation License at the specified site license has expired.	Every 30 minutes, evaluation licenses will post this event.
UnlicensedFeatureEvent	The SRM license at the specified site is overallocated by the specified number of licenses.	Every 24 hours and upon the protection or unprotection of a virtual machine, this event will be posted if the total number of licenses exceeds the capacity in the license.
LicenseUsageChangedEvent	The SRM license at the specified site is using the specified number out of the total number licenses.	Every 24 hours and upon the protection or unprotection of a virtual machine, this event will be posted if the total number of licenses does not exceed the capacity in the license.

Permissions Events

Permission events provide information about changes to SRM permissions.

Table 9-13. Permissions Events

Type	Description	Content
PermissionsAddedEvent	Permission created for the entity on SRM.	A permission for the entity was created using the role specified. The IsPropagate flag indicates whether the permission is propagated down the entity hierarchy.
PermissionsDeletedEvent	Permission rule removed for the entity on SRM	A permission for the entity was deleted.
PermissionsUpdatedEvent	Permission changed for the entity on SRM.	A permission for the indicated entity was modified.

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 use different 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 plug-in logs are part of the general logs produced by the vSphere client. As a result, collecting the vSphere client log files (together with the vSphere log files) collects SRM plug-in log files. The SRM log collects or retrieves the files and collects them in compressed file (zipped) which is placed in a location that you choose.

SRM also provides for the collection of VRMS and VR logs as part of the SRM log bundle. Logs from vCenter servers and ESX servers that are part of your SRM system might also include information useful in diagnosing SRM issues.

Collect SRM Log Files Using the vSphere Client

SRM supports downloading logs for SRM, VRM, VR, VC, and ESX from a single site to a user-specified location. Use this information to understand and resolve issues. For best results, collect logs from each site.

Procedure

- 1 Click **Sites**, and select a site.
- 2 Click the **Summary** tab, and click **Export System Logs**.
- 3 Specify log collection settings. You may After selecting a location, click **Next**.
 - a In the Download Location field, enter a path, or click **Browse** to browse for a location.
You can also create a folder.
 - b (Optional) In Options, disable download of vSphere Replication (VR) log data.
VR system logs are downloaded by default. These logs include information about vSphere Replication Management (VRM), VR, and replication events.

The Downloading System Logs Bundles window provides information about the following:

- A list of each host system, the status of their log bundle download, and other details.
- Download Details provides information on the log bundle file name and destination for the log bundle file.

This process does not collect client logs. Client logs must be collected separately.

Collect SRM Server Log Files

You can collect SRM server log files into a log bundle to gather information that may be useful in diagnosing problems.

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 (x86)\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-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.

Features Are Unavailable When Deploying VRMS

When you deploy VRMS, you get an error about unavailable features.

Problem

When you deploy VRMS, an error appears about unavailable features.

Cause

This error is typically the result of the vCenter Management Web service being paused or stopped.

Solution

Attempt to start the vCenter Management Web service. If the service fails to start, confirm that Tomcat is running on the server. If the server is installed, but not running, try starting the server. If attempting to start the Tomcat server does not resolve the problem, the issue might be occurring because the vCenter Server has the wrong version of Java installed.

OVF Package is Invalid and Cannot be Deployed

When you attempt to deploy OVF for the vSphere Replication Server, the OVF package error might occur.

Problem

The error `OVF package is invalid and cannot be deployed` might appear while you attempt to deploy the vSphere Replication Management Server.

Cause

This problem is due to the vCenter Server port being changed from the default of 80.

Solution

If possible, change the vCenter Server port back to 80.

Connection Errors Between VRMS and SQL Cannot be Resolved

You get a connection error between the vSphere Replication Management Server (VRMS) and SQL that you cannot resolve.

Problem

VRMS might not be able to connect to SQL, and you have insufficient information to solve this problem.

Cause

This problem can be caused by several issues, and initially available information about the problem is insufficient to affect a resolution.

Solution

- 1 Use a file management tool to connect to the VRMS appliance.
For example, you might use SCP or WinSCP. Connect using the root account, which is the same account used to connect to VAMI.
- 2 Delete any files you find in `/opt/vmware/vrms/logs`.
- 3 Connect to VAMI and attempt to save the VRMS configuration.
This action recreates the SQL error.
- 4 Connect to the VRMS appliance again and find the `hms.log` file which is in `/opt/vmware/vrms/logs`.
This log file contains information about the error that just occurred. Use this information to troubleshoot the connection issue, or provide the information to VMware for further assistance.

Configuration of the VRMS Database Fails with DB2 Databases

vSphere Replication creates temporary tables in the vSphere Replication Management Server (VRMS) database. When you use a DB2 database, the database might require additional configuration.

Problem

If the VRMS user account that you use to log in to the VRMS database cannot create temporary tables, configuration of VRMS can fail, the VRMS database can become unstable, and replication can fail. If the temporary table space is not configured for this user account, the VRMS logs include error messages.

- `javax.persistence.PersistenceException: org.hibernate.exception.SQLGrammarException: could not insert/select ids for bulk delete`
- `DB2 SQL Error: SQLCODE=-204, SQLSTATE=42704`

This problem only occurs when you use a DB2 database.

Cause

The temporary table space was not configured for the user account that you use to connect to the VRMS database.

Solution

- 1 Run an SQL script to verify that the VRMS user account can create temporary tables in DB2 databases.

```
declare global temporary table testtable(foobar integer) on commit preserve rows not logged
```


If the script runs successfully, no further configuration is required. If it fails, you see a message like the following message:

A default table space could not be found with a page size of at least "4096" that authorization ID "HMS_PROT" is authorized to use. SQLCODE=-286, SQLSTATE=42727, DRIVER=4.11.69
- 2 If the script fails, run an SQL script to configure the temporary table space.

```
CREATE USER TEMPORARY TABLESPACE tbsp_temp_hms_prot MANAGED BY AUTOMATIC STORAGE GRANT USE OF TABLESPACE tbsp_temp_hms_prot TO USER HMS_PROT
```

Index

A

- administration, overview of **9**
- advanced setting dialog box, vSphere replication **92**
- advanced setting dialog boxes
 - replication **92**
 - storage **91**
- advanced settings dialog box, recovery site **89**
- advanced settings dialog boxes
 - guest customization **88**
 - local site **90**
 - remote site **91**
 - SAN provider **89**
- Advanced Settings dialog boxes **88**
- alarms, SRM-specific **88**
- array based protection group, edit **61**
- array based recovery plan, create **67**
- array managers
 - and storage replication adapters **47**
 - edit **48**
 - replicated device discovery **47**
 - to configure **47**
 - to rescan arrays **48**
- authentication
 - certificate warnings and **18**
 - methods used by Site Recovery Manager **18**

C

- callouts, *See also* recover steps
- certificate
 - public key **18**
 - requirements for **19**
 - to change type **40**
 - to update **40**
- certificate warning **18**
- client plug-in, upgrade **36**
- configure
 - custom recovery steps **79**
 - upgraded srm installation **36**
- connect to, srm **24**
- custom recovery steps, configure **79**
- customizing, IP properties **80**
- customizing SRM **75**

D

- database
 - backup requirements **35, 40**

- configuration details **28**
- Connection Count value **16**
- DB2 **29**
- Max Connections value **16**
- Microsoft SQL Server **28**
- Oracle **28**
- Site Recovery Manager **16**
- to change connection details **28, 40**
- vCenter **17**
- VRMS **29**
- database configuration, vSphere Replication **29**
- datastore
 - protected **11**
 - replicated **13**
- datastore group
 - how computed **12**
 - maximum number supported **24**
- datastore mappings, configure **44**
- dr-ip-customizer.exe, reference **83**

E

- environment variables **76**
- events
 - licensing **101**
 - permissions **101**
 - protection groups **94**
 - recovery **96**
 - site status **94**
 - storage **98**
 - storage provider **98**

F

- fallback, about **16**
- failover, effects of **70**
- feedback **7**
- forced failover **70**

H

- host based replication **12**

I

- installation
 - of storage replication adapter **38**
 - reverting to a previous release **40**
 - Site Recovery Manager Client plug-in **38**
 - Site Recovery Manager server **31**

- to repair **40**
- updating to a new release **35**
- inventory mappings
 - about **14**
 - and placeholders **14**
 - to apply **66**
 - to create **44**
 - to override **86, 87**
- IP address mappings
 - to customize **83**
 - to report **82**
- IP properties, customizing **80**

L

- license key, to install **42**
- licensing
 - about **18**
 - events **101**
 - license key **42**
- linked clones, limitations on recovery of **59**
- log files
 - collecting **102**
 - SRM server **102**
- logs, downloading **102**

M

- managing **24**

N

- network, test **14**
- network settings, vrms **54**

O

- OVF, cannot be deployed **103**

P

- pair, vrms **55**
- permissions
 - events **101**
 - to assign **21**
- permissions and roles, understanding **20**
- permissions required **24**
- physical couriering **64**
- placeholder datastore, add **44**
- placeholder datastores, understanding **43**
- placeholders, in vCenter inventory **14**
- plug-in
 - Site Recovery Manager Client **38**
 - to install **38**
- ports, used by SRM **23**
- protected site
 - configure array managers for **47**
 - configuring **47**
 - host compatibility requirements **11**
 - to designate **39**

- protection group
 - maximum number supported **24**
 - relationship to datastore group **13**
 - relationship to recovery plan **13**
 - to create **60**
- protection groups
 - events **94**
 - vr **61**

R

- recovery
 - customize for a virtual machine **80**
 - events **96**
 - steps **78**
 - test **67**
- recovery plan
 - cleanup **68**
 - command steps **76**
 - customizing **75**
 - force cleanup **68**
 - forced failover **70**
 - running **14, 70**
 - steps **75**
 - testing **14, 68**
 - time-outs **75**
 - to change properties of **68**
 - to customize steps **77**
 - to remove **68**
 - to report IP address mappings used by **82**
 - virtual machine recovery priority **75**
- recovery priority, virtual machine **75, 78**
- recovery site
 - configure array managers for **47**
 - configuring **47**
 - host compatibility requirements **11**
 - to designate **39**
- recovery test, to cancel **69**
- replicating, virtual machines **59**
- replication
 - and recovery **14**
 - array-based **11**
 - sneakernet **64**
- replication framework **43**
- reprotect **16**
- reprotection
 - process **72**
 - understanding **71**
- reprotection states **73**
- roles, to assign **21**

S

- security settings, vrms **53**
- settings, vrs **56**

- site
 - protected **11**
 - recovery **11**
- site pairing **39**
- Site Recovery Manager, and other vCenter Server Solutions **17**
- site status, events **94**
- snapshots, limitations on recovery of **59**
- sneakernet, replication **64**
- SNMP traps **97**
- SRA, See storage replication adapter
- SRM migration utility **37**
- srm roles **21**
- srm upgrade, preparation **34**
- steps, recovery **78**
- storage, events **98**
- storage provider, events **98**
- storage replication adapter
 - and array managers **47**
 - to download **38**
 - to install **38**
- support **7**
- suspended virtual machines, limitations on
 - recovery of limitations on recovery of **59**
- system settings, vrms **52**

T

- troubleshooting **93**

U

- understanding
 - permissions and roles **20**
 - reprotection **71**
- upgrade, client plug-in **36**
- upgrading, SRM **33**

V

- vCenter
 - and Site Recovery Manager **17**
 - to change connection information **40**
 - to change credentials used by Site Recovery Manager **40**
- virtual machine
 - customize IP properties for **83**
 - customize recovery of **80**
 - recovery priority **75, 78**
- vr, replicate virtual machines **62, 63**
- vr protection group, edit **62**
- vr server
 - deploy **55**
 - move **66**
- vrms, working with **51**
- vrms
 - network settings **54**
 - pair **55**
 - security settings **53**

- system settings **52**
- time zone **54**
- VRMS
 - database **29**
 - deploy **52**
 - features unavailable **103**
- VRMS and SQL, connection failure **103**
- VRMS database
 - configure temporary tables **104**
 - troubleshooting **104**
- vrs, settings **56**
- vSphere Replication
 - database configuration **29**
 - deployment limits **24**
- vSphere Replication Server (VRS), register **57**

W

- workflow **10**

