

Dell EMC ME4 Series Storage Replication Adapter for vSphere

User's Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: Installing and configuring the SRA.....	4
About VMware Site Recovery Manager.....	4
Planned migration.....	4
Disaster recovery.....	4
Protected sites and recovery sites.....	4
SRM requirements.....	5
Configuring the ME4 Series storage systems.....	5
Configure replication.....	5
Install SRM software.....	6
Install the SRA software.....	6
Configure SRM.....	6
Chapter 2: Using SRM for disaster recovery.....	7
Volume discovery.....	7
Create a recovery plan.....	7
Test a recovery plan.....	7
Failover and failback.....	8
Automatic failover.....	8
Reprotection.....	9
Automated failback.....	9
Chapter 3: Troubleshooting.....	10
Chapter 4: Best practices.....	13

Installing and configuring the SRA

The Dell EMC ME4 Series Storage Replication Adapter (SRA) for vSphere enables full-featured use of the VMware vCenter Site Recovery Manager (SRM) version 6.5 or later. Combining the ME4 Series storage system's replication functionality with the vCenter SRM, the SRA provides an automated solution for implementing and testing disaster recovery between geographically separated sites. It also enables you to use SRM for planned migrations between two sites.

Topics:

- [About VMware Site Recovery Manager](#)
- [Protected sites and recovery sites](#)
- [SRM requirements](#)
- [Configuring the ME4 Series storage systems](#)
- [Install SRM software](#)
- [Install the SRA software](#)
- [Configure SRM](#)

About VMware Site Recovery Manager

The 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).

Two types of recovery are available—planned migration and disaster recovery.

Planned migration

Planned migration is the orderly decommissioning of virtual machines at the protected site and commissioning of equivalent machines at the 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. 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 enables you to 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.

Protected sites and recovery sites

In a typical SRM installation, a protected site provides business-critical data center services. The protected site can be any site where vCenter supports a critical business need.

The recovery site is an alternative facility to which these services can be migrated. The recovery site can be located thousands of miles away. 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.

NOTE: The ME4 Series SRA connects VMware SRM with the storage system's replication functionality, so you might encounter different terminology that has similar meanings. For example, the VMware user interface and documentation typically refer to protected and recovery sites. The ME Storage Manager (MESM) user interface and replication documentation refer to primary and secondary volumes and sites.

SRM requirements

A typical SRM configuration involves two geographically separated sites with TCP/IP connectivity, the protected site and the recovery site. The protected site is the site that is being replicated to the recovery site for disaster recovery. Each site contains a Dell EMC ME4 Series storage system, VMware ESX servers, a Virtual Center (vCenter) server, and an SRM server running the SRM.

After you have set up the protected site and the recovery site, and installed the necessary infrastructure for networking between the two sites, you can install and configure the software. For more information, see [Configuring the ME4 Series storage systems](#) on page 5.

Configuring the ME4 Series storage systems

If your ME4 Series storage systems are not already configured:

1. Follow the installation instructions in your *Dell EMC ME4 Series Storage System Deployment Guide*.
2. Ensure that both storage systems have the same host interface configuration (iSCSI or FC or hybrid FC/iSCSI).

Configure replication

NOTE: Ensure the storage system name, user credentials, and IP addresses for both storage systems are set before configuring the SRA. The SRA uses the same user credentials for both the local and remote storage system, so if the manage user does have the same password on both sites, create a new user ID for the SRA with manage on both systems. See [Best practices](#) on page 13 for additional setup information.

1. Use the ME Storage Manager (MESM) to configure replication software, following the instructions in the replication section of the Administrator's Guide, including the following settings for SRA:
 - snapshot-count: 3 (or higher)
 - snapshot-history: both
 - snapshot-retention: high
 - queue-policy: queue-latest
 - (optional) snapshot-basename: same-as-volume-name

NOTE: Setting the basename as indicated makes troubleshooting easier because replication snapshots will have the same name as the base volume with _nnnn appended (indicating the replication generation number).
2. Use the MESM on each system to define the other system in the replication set as a remote system.
3. Use the MESM to perform at least one replication.
4. Use the MESM to schedule replications from the protected site to the recovery site. Doing so ensures that, in the event of a disaster that disables the protected site, damages hardware, or damages files, SRM can use the most recently replicated copy at the recovery site for disaster recovery. It is important, when using scheduled replication replications, to verify that the source of the most recent replication was in a valid state.

Install SRM software

You must install an SRM server at the protected site and also at the recovery site. After the SRM servers are installed, download the SRM client plug-in from either SRM server using the **Manage Plugins** menu from your vSphere Client. Use the SRM client plug-in to configure and manage SRM at each site.

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.

To install the SRM software:

1. Set up vCenter Server at each site.
2. Create a single data center in each instance of vCenter Server.
3. Add the local hosts to this data center.
4. Download VMware Site Recovery Manager software using the Download Product link at the following link on the VMware website:

<https://my.vmware.com/web/vmware/downloads>

5. Install VMware Site Recovery Manager 6.5 or later at each site, following the instructions in the VMware *Site Recovery Manager Administration* guide.

You can locate the guide and release notes for your SRM version on the VMware documentation website:

<https://docs.vmware.com/en/Site-Recovery-Manager/index.html>

6. Do not configure SRM at this time. First, [Install the SRA software](#) on page 6.

Install the SRA software

Download and install the Dell EMC ME4 Series SRA software for your version of VMware SRM:

1. Go to <https://www.dell.com/support>.
2. Locate the SRA software on the ME4 Series series storage system support page.
3. Open and run the SRA setup file.
After the SRA is installed at each site, you can configure SRM, which enables the plug-in to discover the volumes replicated between the sites.

Configure SRM

After you have both SRM and SRA installed, the **Getting Started** tab of the main SRM window guides you through the steps necessary to configure it. For detailed SRM configuration instructions, see the VMware publication *Site Recovery Manager Administration* guide.

Configuring ME4 Series storage systems in SRM requires the following:

- The IP addresses of the ME4 Series storage systems.
- A user name and a password for each storage system. This is the user name and password as configured in the MESM.

Make the following changes to your SRM settings:

- Set `storageProvider.autoResignatureMode` to 1 (required).
- Set `storageProvider.hostRescanRepeatCnt` to 2 (required).
- Set `Storage.commandTimeout` to 1200 seconds (recommended).

Using SRM for disaster recovery

After storage system replication and VMware SRM software are configured at local and remote sites and you have configured at least one replication set, use MESM to schedule replications. Then use SRM to create and test one or more recovery plans. At this point, SRM is able to provide disaster recovery, failover and failback, and reprotect operations.

The VMware *Site Recovery Manager Administration* guide provides detailed instructions and information regarding these operations. The guide is available at the [VMware support website](#).

Topics:

- [Volume discovery](#)
- [Create a recovery plan](#)
- [Test a recovery plan](#)
- [Failover and failback](#)
- [Automatic failover](#)
- [Reprotection](#)
- [Automated failback](#)

Volume discovery

SRM obtains information from the SRA about what volumes are being replicated. SRM then compares that list to the volumes it recognizes in a VMware environment.

For SRM planned migrations in non-disaster situations, SRM can initiate a replication to ensure that the replicated data is current.

For disaster recovery situations, SRM attempts to create a current replication. If this is not possible because, for instance, the protected site is offline, SRM uses the most recent replication available at the remote site.

Use the replication scheduler to regularly perform replications to minimize data loss in the event of a disaster, or regularly create SRM planned migrations. In either case, ensure that the volumes to be replicated from the protected site are in a valid state so that the most recent replication at the remote site can be used in production.

For instructions on how to configure replication schedules, see the chapter on using replication in the *Dell EMC ME4 Series Storage System Administrator's Guide*.

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. Recovery plans indicate how virtual machines in one or more protection groups are restored at the recovery site.

The **Recovery** tab of the main SRM window guides you through the steps necessary to create, test, and run a recovery plan. For detailed instructions, see the VMware publication *Site Recovery Manager Administration* guide.

Test a recovery plan

You can automatically create a non-disruptive, isolated testing environment on the recovery site by using replication and connecting virtual machines to your isolated testing network. You can also save test results for viewing and export at any time.

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

You should run test recoveries as often as needed. Testing a recovery plan does not affect replication or the ongoing operations of either site (though it might temporarily suspend the 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 and the target. Note that 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.

When SRM test failovers to the recovery site are requested, SRM performs the following steps:

1. Determines the latest recovery point for each replicated volume.
2. Creates a writeable test snapshot for each recovery point, with a name in the form **sra***nnnnnn* where *nnnnnn* is a monotonically increasing number.
3. Maps the test snapshots to the appropriate ESXi hosts on the recovery site.

When testing stops, the test snapshots are unmapped and deleted.

Failover and failback

Failback is the process of setting the replication environment back to its original state at the protected site prior to failover. Failback with SRM is an automated process that occurs after recovery. This makes the failback process of the protected virtual machines relatively simple in the case of a planned migration. If the entire SRM environment remains intact after recovery, failback is done by running the reprotect recovery steps with SRM, followed by running the recovery plan again, which moves the virtual machines configured within their protection groups back to the original protected SRM site.

In disaster scenarios, failback steps vary with respect to the degree of failure at the protected site. For example, the failover could have been due to a storage system failure or the loss of the entire data center. The manual configuration of failback is important because the protected site may have a different hardware or SAN configuration after a disaster. Using SRM, after failback is configured, it can be managed and automated like any planned SRM failover. The recovery steps can differ based on the conditions of the last failover that occurred. If failback follows an unplanned failover, a full data re-mirroring between the two sites may be required. This step usually takes most of the time in a failback scenario.

All recovery plans in SRM include an initial attempt to synchronize data between the protection and recovery sites, even during a disaster recovery scenario.

During the disaster recovery, an initial attempt will be made to shut down the protection group's virtual machines and establish a final synchronization between the sites. This is designed to ensure that virtual machines are static and quiescent before running the recovery plan, in order to minimize data loss wherever possible. If the protected site is no longer available, the recovery plan will continue to execute and will run to completion even if errors are encountered.

This new attribute minimizes the possibility of data loss during a disaster recovery, balancing the requirement for virtual machine consistency with the ability to achieve aggressive recovery-point objectives.

Automatic failover

SRM automates the execution of recovery plans to ensure accurate and consistent execution. Through the vCenter Server you can gain full visibility and control of the process, including the status of each step, progress indicators, and detailed descriptions of any error that occurs.

In the event of a disaster when an SRM actual failover is requested, the SRA will perform the following steps:

1. Select the replicated volumes.
2. Identify and remove any incomplete remote copies that are in progress and present the most recently completed Remote Copy as a primary volume.
3. Convert remote volumes into primary volumes and configure authentication for ESXi hosts to mount them.

If an actual failover does not run completely for any reason, the failover can be called many times to try to complete the run. If, for example, only one volume failed to restore and that was due to a normal snapshot being present, the snapshot could be manually deleted and the failover be requested again.

Reprotection

After a recovery plan or planned migration is executed, there are often cases where the environment must continue to be protected against failure in order to ensure its resilience or to meet all disaster recovery objectives.

Reprotection is an SRM extension to recovery plans for use only with storage system replication. It enables the environment at the recovery site to establish synchronized replication and protection of the original environment.

After failover of the recovery site, choosing to reprotect the environment will establish synchronization and attempt to replicate the data between the protection groups running at the recovery site and at the previously protected primary site.

This capability to reprotect an environment ensures that environments are protected against failure even after a site recovery scenario. It also enables automated failback to a primary site following a migration or failover.

Automated failback

You can set up an automated failback workflow to return the entire environment to the primary site from the recovery site.

The failback happens after the reprotection has ensured that data replication and synchronization are established to the original primary site.

Automated failback runs the same workflow that was used to migrate the environment to the protected site. It ensures that the critical systems encapsulated by the recovery plan are returned to their original environment. The workflow executes only if reprotection is successfully completed. Failback is only available with storage system replication.

Failback ensures the following:

- All virtual machines that were initially migrated to the recovery site are moved back to the primary site.
- Environments that require that disaster recovery testing be done with live environments with genuine migrations can be returned to their initial site.
- Simplified recovery processes enable a return to standard operations after a failure.
- Failover can be done in case of disaster or in case of planned migration.

Troubleshooting

VMware vCenter Server uses the SRA to display a detailed error message each time a recovery step fails.

The SRA also creates a log file called `sra.log` that shows each SRM event and each CLI command that occurs on the ME4 Series storage systems. Examining the error messages and this log file will often provide enough information to rectify errors. For additional assistance, contact VMware for support.

Table 1. SRA error messages and suggested actions

Message number	Message	Suggested action
1002	VMware Site Recovery Manager version 6.5 was not found on this system.	Install VMware SRM 6.5 or later and then rerun the SRA installation procedure.
1003	XML output to "{file}" failed: {error}	Ensure that the specified file location exists, has adequate free space, and is writable.
1004	Install option is not supported on this system	See the SRA installation instructions.
1005	A native version of Perl must be used when invoking this option.	Ensure that you are using the <code>Perl.exe</code> version installed with the VMware SRM software.
1006	Timed out waiting for volume {volume} to appear on the storage system {arrayname} at {file}:{line}.	Verify that the specified volume has been created on the storage system and retry the operation.
1007	Storage system '{systemName}' is not licensed for use with this SRA.	Contact your storage system vendor to verify that this system is supported and to request replication and SRA license keys.
1008	No WWN found for volume "{primary}".	Verify that the specified volume is configured for replication.
1009	discoverDevices: Could not determine WWN for temporary snapshot "{serialNumber}" ({name}).	Check to see whether the specified snapshot was left over from a previous test and can be deleted.
1010	Cannot find recovery point for temporary snapshot "{serialNumber}" ({name}).	Check to see whether the specified snapshot was left over from a previous test and can be deleted.
1011	discoverDevices: could not find WWN for promoted volume "{secondaryName}" ({secondary}).	Check the status of the specified volume and the health of the storage system and then retry the operation.
1013	No valid sync point exists for {volume}.	In MESM, use the Snapshots table to verify that the specified volume has been completely replicated from the protected site. For more information, see the Administrator's Guide.
1014	Could not export a snapshot for volume {vol}.	A snapshot previously created by the SRA already exists for the specified volume. Only one exported snapshot is allowed per replication destination volume. Delete the existing snapshot and retry this operation.
1018	unknown or missing PeerId parameter '{PeerId}' in {command} request.	Ensure that each storage system reports the name of its replication peer(s) correctly, and that the storage system names have not changed since SRM was configured. If the storage system name has been changed, delete and recreate remote system entries on each storage system as necessary. If the problem continues after restarting SRM, recreate the storage system pair configuration in SRM.

Table 1. SRA error messages and suggested actions (continued)

Message number	Message	Suggested action
1020	Could not find peer volume for local volume {localsn}.	Ensure that the specified volume has been set up as part of a replication set.
1021	Invalid or missing parameters in SRM '{cmd}' request received by the SRA.	Verify that the replication sets, remote systems, and SRM configuration are correct.
1022	Invalid or unknown ArrayId '{ArrayId}' in {cmd} request.	Ensure that the storage controller system names and IP addresses have not been reconfigured since SRM was configured.
1023	Failed to open lock file {filename}.	Check file and directory permissions for the specified filename.
1024	Unknown or missing DeviceId parameter '{DeviceId}' in {command} request.	Verify that SRM and the SRA are configured correctly. Also check the health of storage system and network paths between the SRM host and both storage systems.
1025	No valid sync point found for volume {vol} during the {command} operation.	The operation failed on this volume because no valid sync point exists for the volume. In MESM, use the Snapshots table to verify that the specified volume has been completely replicated from the protected site. For more information, see the Administrator's Guide.
1026	Timed out waiting for replication set for volume {volume} to transition to conflict status on storage system {arrayname} at {file}:{line}.	Verify that the specified volume has been created on the storage system and retry the operation.
1027	The SRA syncOnce command timed out waiting for replication images for volume(s) [{volumes}] to start on the storage system.	Check to make sure that the storage system is healthy, and repeat the operation if necessary to ensure that the volumes are replicated.
1028	No SRA snapshot found for volume '{DeviceId}' in {command} request.	The SRA failed to export the snapshot in a previous testFailoverStart operation, or the snapshot has already been removed, or the snapshot was not found due to a problem communicating with the management port on the storage system.
1029	An existing SRA snapshot {snapshot} must be removed before the testFailoverStart function can be performed on {volume}.	Remove snapshot volume {snapshot} before trying the test failover operation again.
1030	reverseReplication cannot be performed on target volume {volume} because original protected volume {target} is still mapped on the remote storage system {remoteArray}	Ensure that both storage systems ({localArray} and {remoteArray}) and their corresponding SRM servers are running and manageable over the network.
1101	Failed to log in to storage system at {url} ({response})	Ensure that storage system IP addresses are configured correctly and that the storage system is reachable from the SRM host. Also, if any storage system IP addresses have changed, it may be necessary to delete and recreate the remote system definitions on one or both storage systems.
1102	Execution of command "{cmd}" failed on storage system at {ipAddr}: {err}	If the error message did not specify the reason for the failure, open the specified address with a web browser to check the health of the storage system.
1103	No IP addresses specified for MC for command "{cmd}"	Verify that the IP addresses for the storage system are configured correctly on the storage system and on the host.
1104	Response from storage system at {ipAddr} did not include status indication.	Check the health of the storage system and restart the management controller if necessary.

Table 1. SRA error messages and suggested actions (continued)

Message number	Message	Suggested action
1105	Failed to run command "{cmd}" on storage system at {system}: {err}	Verify the IP address configuration on the storage system and on the host, and check network connectivity.
2001	Volume {volume}({name}) is already unmapped.	SRM requested that a volume be prepared for failover, but the volume is already prepared.
2002	No data found for {volume}replication image {imageSn} ({err}).	Verify that replication has started for volume {volume}.
2003	querySyncStatus: No data found for replication image {imageSn} for volume {vol} ({err}).	Verify that replication has started for the specified volume.

NOTE: You can expect to see certain errors in the log file when commands are executed to ensure that volumes are in a particular state if the volumes are already in that state. These errors are -3395 (Replication is not active on this secondary volume) and -10306 (Unable to set the specified volume as the primary volume because the specified volume is already a primary volume). You can safely disregard these error messages if they occur under these circumstances.

Best practices

Specific guidelines and recommendations for using the SRA and replication software in conjunction with the VMware SRM disaster recovery solution include the following:

- Prepare a plan in advance for how you will re-establish replication schedules in the event of a site failover. After performing a reverse-replication operation, you must set up replication schedules in order to ensure periodic replication of data from the new source volumes back to the original source site. Alternatively, you can initiate replication manually if appropriate.
- Try to group virtual machines with related backup requirements or schedules on the same datastore volume, since replication occurs on a per-volume basis. For example, if some virtual machines do not need to be replicated to a remote site, or need to be replicated less frequently, do not store them on the same datastore volume as virtual machines which must be replicated frequently, to avoid replicating data unnecessarily.
- The SRA only supports replication between identical hardware models. For example, replication between an iSCSI-only system and a FC/iSCSI hybrid system is not supported.
- Avoid mapping replication volumes to LUN 0 to avoid issues with dynamically mapping and unmapping LUNs, due to special management functionality assigned to LUN 0. You can map volumes to LUN 0 if those volumes are not expected to be mapped and unmapped automatically the way replication volumes are, such as local datastores that are not replicated.
- Replication volumes should be mapped with the same LUN number on all hosts.
- Do not use the same LUN number for different volumes that are mapped to different hosts.
- Failover operations will cause read-write host mappings for replication volumes to be converted to read-only, and restoring replication will convert all read-only mappings for the same volume to read-write. Be careful not to create read-only mappings of replication volumes such as for data mining purposes. If a read-only mapping of a replication volume is required, consider creating a non-replicated hardware or software snapshot of the volume.
- The SRA might create host entries on the storage system to keep track of remote IP or FC addresses. Do not delete host entries whose name starts with "SRA." However, you may rename them to be more descriptive.
- Replication set basenames for replicated volumes (assigned when creating the replication set) should be no more than 23 bytes long to allow for suffixes to be appended when creating replication snapshots. 23 bytes allows for up to 23 ASCII characters, but non-ASCII UTF-8 characters require more than one byte each.
- Do not change the name of replication snapshots or change replication-set basenames except to conform to these best practices. The SRA depends on consistency between the replication-set basenames and the snapshot names.