# AssuredSAN 3000 Series Storage Replication Adapter Software User Guide

Adobe PostScript™

# Contents

# Figures

# Tables

# About this guide

The Dot Hill AssuredSAN 3000 Series Storage Replication Adapter (SRA) Software for VMware® vCenter Site Recovery Manager (SRM) version 5.1 or 5.0 enables full-featured use of the VMware SRM with AssuredSAN 3000 Series storage systems. Combining AssuredSAN 3000 Series AssuredRemote replication software with VMware SRM provides an automated solution for implementing and testing disaster recovery between geographically separated sites. This white paper provides information about configuring and using the SRA with VMware vCenter Site Recovery Manager (SRM).

## Intended audience

This guide is intended for system administrators who possess extensive knowledge of host hardware, AssuredSAN 3000 Series storage systems, and VMware Site Recovery Manager (SRM). SRM administrators should also be familiar with vSphere and its replication technologies such as host-based replication and replicated datastores.

## Prerequisites

Prerequisites for using this product include knowledge of:

- Network administration
- Storage system configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- VMware products and services, especially SRM.

## Related documentation

**Table 1**  Related documentation

| For information about | See |
|---|---|
| Enhancements, known issues, and late-breaking information not included in product documentation | Release Notes |
| Overview of product shipkit contents and setup tasks | Getting Started[*] |
| Regulatory compliance and safety and disposal information | AssuredSAN 3000 Series Product Regulatory Compliance and Safety[*] |
| Installing and using optional host-based software components (CAPI Proxy, MPIO DSM, VDS Provider, VSS Provider, SES Driver) | AssuredSAN 3000 Series Installing Optional Software for Microsoft Windows® Server |
| Recommendations for using optional data-protection features (AssuredSnap, AssuredCopy, AssuredRemote) | AssuredSAN 3000 Series Using Data Protection Software |
| Using a rackmount bracket kit to install an enclosure into a rack | AssuredSAN 3000 Series Rackmount Bracket Kit Installation[*] *or* AssuredSAN 3000 Series 2-Post Rackmount Bracket Kit Installation[*] |
| Product hardware setup and related troubleshooting | AssuredSAN 3000 Series Setup Guide |
| Obtaining and installing a license to use licensed features | AssuredSAN 3000 Series Obtaining and Installing a License Certificate File |
| Using the web interface to configure and manage the product | AssuredSAN 3000 Series RAIDar User Guide |
| Using the command-line interface (CLI) to configure and manage the product | AssuredSAN 3000 Series CLI Reference Guide |

**Table 1** Related documentation (continued)

| For information about | See |
|---|---|
| Event codes and recommended actions | AssuredSAN 3000 Series Event Descriptions Reference Guide |
| Identifying and installing or replacing field-replaceable units (FRUs) | AssuredSAN 3000 Series FRU Installation and Replacement Guide |

\* Printed document included in product shipkit.

For additional information, see Dot Hill's Customer Resource Center web site: crc.dothill.com.

# Document conventions and symbols

**Table 2** Document conventions

| Convention | Element |
|---|---|
| Blue text | Cross-reference links and e-mail addresses |
| Blue, underlined text | Web site addresses |
| **Bold font** | • Key names<br>• Text typed into a GUI element, such as into a box<br>• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes |
| *Italics font* | Text emphasis |
| `Monospace font` | • File and directory names<br>• System output<br>• Code<br>• Text typed at the command-line |
| `Monospace, italic font` | • Code variables<br>• Command-line variables |
| `Monospace, bold font` | Emphasis of file and directory names, system output, code, and text typed at the command line |

⚠ **WARNING!**   Indicates that failure to follow directions could result in bodily harm or death.

△ **CAUTION:**   Indicates that failure to follow directions could result in damage to equipment or data.

📝 **IMPORTANT:**   Provides clarifying information or specific instructions.

**NOTE:**   Provides additional information.

☼ **TIP:**   Provides helpful hints and shortcuts.

# 1 Installing and configuring the SRA

The Dot Hill AssuredSAN 3000 Series SRA for VMware SRM enables full-featured use of VMware Site Recovery Manager 5.1 or 5.0. Combining AssuredRemote replication with VMware SRM 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.

## About VMware Site Recovery Manager

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).

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 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 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 datacenter 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:** Because the Dot Hill AssuredSAN 3000 Series SRA connects VMware SRM with AssuredSAN 3000 Series AssuredRemote replication software, you might encounter different terminology that has similar meanings. The VMware user interface and documentation typically refer to protected and recovery sites. The AssuredSAN 3000 Series RAIDar user interface and AssuredRemote 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 Dot Hill AssuredSAN 3000 Series storage system, VMware ESX servers, a Virtual Center (vCenter) Server, and a SRM server running VMware Site Recovery Manager SRM 5.1 or 5.0 software.
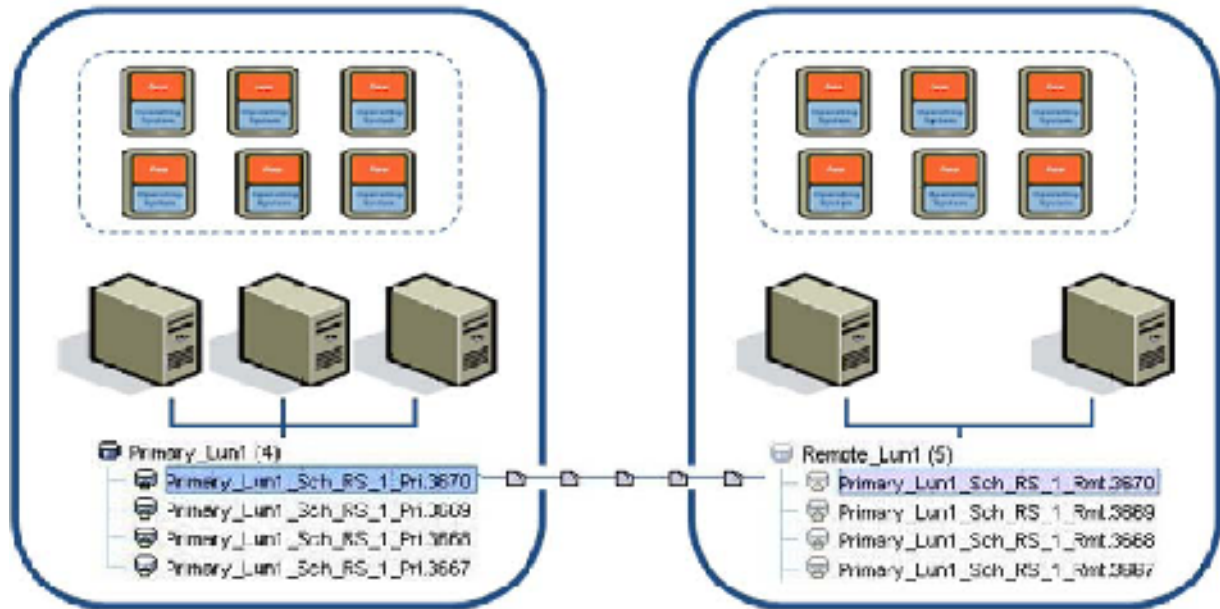


**Figure 1** Typical SRM configuration showing a protected site and recovery site

Once 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. An overview of the necessary steps is shown below, with the titles of the appropriate documents where you can find detailed instructions. See the Reference appendix on page 25 for links to the document locations.

# Installing and configuring AssuredSAN 3000 Series storage systems

Ensure that both storage systems have the same host interface (iSCSI or FC or hybrid) configuration.

If your AssuredSAN 3000 Series storage systems are not already configured:

1. Follow the installation instructions in the *AssuredSAN 3000 Series Setup Guide*.
2. Ensure that both storage systems have the same host interface configuration (iSCSI or FC or Hybrid FC/iSCSI).
3. Ensure that replication, snapshot, and SRA licenses are installed and enabled on both storage systems as described in the *AssuredSAN 3000 Series RAIDar User Guide*.

## Using RAIDar's Replication Setup Wizard

1. Use the Replication Setup Wizard in RAIDar to configure AssuredRemote software, following the instructions in Chapter 6, "Using AssuredRemote to replicate volumes" of the *AssuredSAN 3000 Series RAIDar User Guide*, to do the following:
   a. Select the primary volume, which is an existing volume or snapshot to replicate.
   b. Specify whether the replication mode will be local or remote. If replication will be to a remote system that has not been added already to the local system, you can add it. To do so, you must know the user name and password of a user with Manage role on that system, and the system's IP address.
   c. Select the secondary volume. You can select an existing volume prepared for replication or create a volume in an existing vdisk that has sufficient available space for the replicated data.

  **d.** Confirm your changes and apply them.

 **2.** Use RAIDar on each system to define the other system in the replication set as a remote system.

 **3.** Use RAIDar to perform at least one replication.

 **4.** Optionally, use RAIDar 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 AssuredRemote replications, to verify that the source of the most recent replication was in a valid state.

 An alternative approach is to use SRM's planned migration capabilities to create regular replications.

# Installing 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.

 **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 5.1 or 5.0 software from:
  https://my.vmware.com/web/vmware/downloads.

 **5.** Install VMware Site Recovery Manager 5.1 or 5.0 at each site, following the instructions in the *VMware Site Recovery Manager Administration Guide.*

See the *VMware vCenter Site Recovery Manager Release Notes f*or additional SRM requirements.

Do not configure SRM at this time.

# Installing the SRA

 **1.** Download the Dot Hill AssuredSAN 3000 Series software for the most recent version of VMware SRM 5.1 or 5.0.

 For VMware SRM 5.1.1, that website is:

 https://my.vmware.com/web/vmware/details?downloadGroup=SRM511&productId=291.

 For VMware SRM 5.0.2, that website is:

 https://my.vmware.com/web/vmware/details?downloadGroup=SRM502&productId=238.

 The SRA is also available from Dot Hill systems at:

 http://dothill.com/vmware-sra.

 **2.** Install the 3000 Series SRA on the SRM server at each site.

  **a.** Open `assuredsan-sra-2.1.`*`xy`*`.zip` with Microsoft Windows Explorer.

  **b.** Run `assuredsan-sra/setup.exe` to install the SRA.

 The installation process is simple and straightforward. Once the SRA is installed at each site you can configure SRM.

# Configuring SRM

Once you have both SRM and the 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 *Site Recovery Manager Administration Guide.*

Configuring AssuredSAN 3000 Series arrays in SRM requires the following:

- The IP addresses of the AssuredSAN 3000 Series arrays.
- A user name and a password for each array.

---

**NOTE:** This is the AssuredSAN 3000 Series user name and password as configured in RAIDar.

---

When configuring SRM, rescan so that it detects the AssuredSAN SRA and is able to discover volumes replicated between sites.

# 2 Using SRM for disaster recovery

Once AssuredSAN 3000 Series AssuredRemote replication software and VMware SRM software are configured and licensed at local and remote sites and you have configured at least one replication set, use RAIDar 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, which are summarized below.

## Array and volume discovery

SRM obtains information from the 3000 Series SRA about what volumes are being replicated by the AssuredRemote software. SRM then compares that list to the volumes it recognizes in a VMware environment.

For SRM planned migrations in non-disaster situations, SRM ensures that the replication 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 AssuredRemote 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 Chapter 6, "Using AssuredRemote to replicate volumes" in the *AssuredSAN 3000 Series RAIDar User Guide*. For more information about using AssuredRemote, see *AssuredSAN 3000 Series Using Data Protection Software*. You can download these manuals from the location shown in Appendix A, "Reference," on page 25.

AssuredSAN 3000 Series SRA documentation, including the most recent version of this manual, is available at http://dothill.com/vmware-sra.

## Creating 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, refer to the *Site Recovery Manager Administration Guide*.

## Testing a recovery plan

You can automatically create a non-disruptive, isolated testing environment on the recovery site by using AssuredRemote 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, the 3000 Series SRA will perform the steps listed.

1. Select the replicated volumes.
2. Identify the latest complete Remote Copy snapshot.
3. Delete any temporary writable space on that snapshot to ensure an unedited snapshot is presented to ESX hosts.
4. Configure authentication for ESX hosts to directly mount snapshots.
5. When testing stops, to conserve space on the SAN, delete the temporary writable space that was used during the test.

# 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 will move 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 an array 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.

SRM reprotection is an extension to recovery plans for use only with array-based 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

An automated failback workflow can be run to return the entire environment to the primary site from the recovery site.

This will happen after the reprotection has ensured that data replication and synchronization are established to the original primary site.

Failback will run the same workflow that was used to migrate the environment to the protected site. It will ensure that the critical systems encapsulated by the recovery plan are returned to their original environment. The workflow will execute only if reprotection is successfully completed. Failback is only available with array-based replication.

Failback ensures the following:

- All virtual machines that were initially migrated to the recovery site will be 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 will enable a return to standard operations after a failure.
- Failover can be done in case of disaster or in case of planned migration.

# 3 Troubleshooting

VMware vCenter Server uses the 3000 Series SRA to present a detailed error message each time a recovery step fails.

The 3000 Series SRA also creates a log file called `sra.log` that shows each SRM event and each CLI command that occurs on the AssuredSAN 3000 Series storage systems. Examining the error messages and this log file will often provide enough information to rectify errors.

**Table 3** SRA error messages and suggested actions

| Message number | Message | Suggested action |
|---|---|---|
| 1002 | VMware Site Recovery Manager version *x.x* was not found on this system. | Install VMware SRM 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 | Refer to the SRA installation instructions. |
| 1005 | A native version of Perl must be used when invoking this option. | Ensure that you are using the `Perl.exe` version installed with the VMware SRM software. |
| 1006 | Timed out waiting for volume {volume} to appear on array {arrayname} at {file}:{line}. | Verify that the specified volume has been created on the array and retry the operation. |
| 1007 | Array '{systemName}' is not licensed for use with this SRA. | Contact your array vendor to verify that this array is supported and to request AssuredRemote 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 array and then retry the operation. |
| 1013 | No valid sync point exists for {volume}. | In RAIDar, use the Snapshot Properties table to verify that the specified volume has been completely replicated from the protected site. See Chapter 6, "Using AssuredRemote to replicate volumes" in the *AssuredSAN 3000 Series RAIDar User 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 array reports the name of its replication peer(s) correctly, and that the array names have not changed since SRM was configured. If the array name has been changed, delete and recreate remote system entries on each array as necessary. If the problem continues after restarting SRM, recreate the array pair configuration in SRM. |
| 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. |

**Table 3**   SRA error messages and suggested actions

| 1021 | Invalid or missing parameters in SRM '{cmd}' request received by 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 array management 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 array and network paths between the SRM host and both arrays. |
| 1024 | unknown or missing DeviceId parameter '{DeviceId}' in {command} request. | Verify that SRM and the SRA are configured correctly. Also check the health of array and network paths between the SRM host and both arrays. |
| 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 RAIDar, use the Snapshot Properties table to verify that the specified volume has been completely replicated from the protected site. See Chapter 6, "Using AssuredRemote to replicate volumes" in the *AssuredSAN 3000 Series RAIDar User Guide.* |
| 1026 | Timed out waiting for replication set for volume {volume} to transition to conflict status on array {arrayname} at {file}:{line}. | Verify that the specified volume has been created on the array and retry the operation. |
| 1026 | Timed out waiting for replication set for volume {volume} to transition to online status on array {arrayname} at {file}:{line}. | Verify that the specified volume has been created on the array and retry the operation. |
| 1027 | The SRA syncOnce command timed out waiting for replication images for volume(s) [{volumes}] to start on the array. | Check to make sure that the array 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 array. |
| 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 remote array {remoteArray} | Ensure that both arrays ({localArray} and {remoteArray}) and their corresponding SRM servers are running and manageable over the network. |
| 1101 | Failed to log in to array at {url} ({response}) | Ensure that array IP addresses are configured correctly and that the array is reachable from the SRM host. Also, if any array IP addresses have changed, it may be necessary to delete and recreate the remote system definitions on one or both arrays. |
| 1102 | Execution of command "{cmd}" failed on array 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 array. |

**Table 3** SRA error messages and suggested actions

| 1103 | No IP addresses specified for MC for command "{cmd}" | Verify that the IP addresses for the array are configured correctly on the array and on the host. |
|---|---|---|
| 1104 | Response from array at {ipAddr} did not include status indication. | Check the health of the array and restart the management controller if necessary. |
| 1105 | Failed to run command "{cmd}" on array at {system}: {err} | Verify the IP address configuration on the array 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.

# 4 Best practice recommendations

See the documents in the table below for information about configuring AssuredRemote software to create and schedule replications that SRM can leverage, using the 3000 Series SRA, and for recommendations on their use. These documents are available from Dot Hill's Customer Resource Center at http://crc.dothill.com.

**Table 4** AssuredSAN 3000 Series information

| For information about | See |
|---|---|
| Product hardware setup and related troubleshooting | *AssuredSAN 3000 Series 3000 Series Setup Guide* |
| Obtaining and installing a license to use licensed features | *AssuredSAN 3000 Series 3000 Series Obtaining and Installing a License Certificate File* |
| Using the web interface to configure and manage the product | *AssuredSAN 3000 Series 3000 Series RAIDar User Guide* |
| Recommendations for using optional data-protection features (AssuredSnap, AssuredCopy, AssuredRemote) | *AssuredSAN 3000 Series 3000 Series Using Data Protection Software* |

Specific recommendations for using the 3000 Series SRA and AssuredRemote software in conjunction with VMware SRM disaster recovery solution include the following:

- Prepare a plan ahead of time for how you will re-establish replication schedules in the event of a site failover. After performing a reverseReplication operation, the user must set up replication schedules in order to ensure periodic replication of data from the new source volumes back to the original source site. Alternately, 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.

- Each array must have a unique alphanumeric "System Name" assigned which does not use any non-alphanumeric characters other than "." or "-".

- Each array must have remote systems defined for each of the remote arrays to or from which data is being replicated. The SRA depends on these remote system names being defined for basic functionality.

- The SRA only supports replication between identical hardware models. For example, replication between an iSCSI-only system and a FC/iSCSI Hybrid 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 array to keep track of remote IP or FC addresses. Do not delete host entries whose name starts with "SRA."

# 5 Reference

This guide provides general information about using the 3000 Series SRA. Specific information, including installation and configuration details, is provided and updated in these manuals.

## VMware documentation

The *VMware Site Recovery Manager Administration Guide* is the definitive manual for SRM, including installation and configuration information. It is available at:

http://pubs.vmware.com/srm-51/topic/com.vmware.ICbase/PDF/srm-admin-5-1.pdf.

*VMware vCenter Site Recovery Manager Release Notes* cover the following topics:

• What's New
• Localization
• Compatibility
• Installation and Upgrade Notes
• SRM SDKs
• Open Source Components
• Caveats and Limitations
• Known Issues

The release notes are available at:

https://www.vmware.com/support/srm/srm-releasenotes-5-1-0.html.

*Compatibility Matrixes for vCenter Site Recovery Manager* show compatibility (or incompatibility) between VMware vCenter Site Recovery Manager with platforms, database software, guest operating systems, storage partners, and other VMware solutions. These matrixes are available at:

http://www.vmware.com/support/srm/srm-compat-matrix-5-1.html.

http://www.vmware.com/pdf/srm_compat_matrix_5_0.pdf.

An overview of the steps involved in creating, testing, and executing recovery plans, including examples of VMware management screens, is available at:

http://www.vmware.com/products/site-recovery-manager/screens.html.

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.

Other VMware vCenter Site Recovery Manager documentation is available at:

http://vmware.com/support/pubs/srm_pubs.html.

## Dot Hill AssuredSAN 3000 Series documentation

Dot Hill AssuredSAN 3000 Series product information is available from the AssuredSan & R/Evolution Products folder on the Dot Hill Customer Resource Center at:

http://crc.dothill.com.

The AssuredSAN 3000 Series SRA software is available at http://dothill.com/vmware-sra, along with links to support information and product documentation, including the most recent version of this manual.

# Glossary

| | |
|---|---|
| **CHAP** | Challenge-Handshake Authentication Protocol. |
| **chunk size** | The amount of contiguous data that is written to a vdisk member before moving to the next member of the vdisk. |
| **compatible disk** | A disk that has enough capacity to replace a failed disk and is the same type (SAS or SATA). |
| **dedicated spare** | A disk that is reserved for use by a specific vdisk to replace a failed disk. See compatible disk. |
| **default mapping** | Host-access settings that are configured when a volume is created, and that apply to all hosts that are not explicitly mapped to that volume using different settings. See also explicit mapping and masking. |
| **drive spin down (DSD)** | A power-saving feature that monitors disk activity in the storage system and spins down inactive SAS and SATA disks, based on user-selectable policies. |
| **dual-port disk** | A disk that is connected to both controllers so its data path is fault-tolerant. |
| **dynamic spare** | An available compatible disk that is automatically assigned, if the dynamic spares option is enabled, to replace a failed disk in a redundant vdisk. See compatible disk. |
| **EC** | Expander Controller. The processor (located in the SAS expander in each controller module and expansion module) that controls the SAS expander and provides SES functionality. See also EMP. |
| **EMP** | Enclosure management processor. An EC subsystem that provides SES data such as temperature, power supply and fan status, and the presence or absence of disks. |
| **explicit mapping** | Access settings for a host to a volume that override the volume's default mapping. See also default mapping and masking. |
| **FC** | Fibre Channel interface protocol. |
| **global spare** | A disk that is reserved for use by any redundant vdisk to replace a failed disk. See compatible disk. |
| **host** | An external port that the storage system is attached to. The external port may be a port in an I/O adapter in a server, or a port in a network switch. |
| **image ID** | A globally unique serial number that identifies the point-in-time image source for a volume. All volumes that have identical image IDs have identical data content, whether they be snapshots or stand-alone volumes. |
| **IQN** | iSCSI Qualified Name. |
| **iSCSI** | Internet SCSI interface protocol. |
| **iSNS** | Internet Storage Name Service. |
| **jumbo frame** | In an iSCSI network, a frame that can contain 9000 bytes for large data transfers. A normal frame can contain 1500 bytes. |
| **leftover** | The state of a disk that has been automatically excluded from a vdisk, and is no longer needed by the vdisk after the vdisk is reconstructed. |
| **loop** | Fibre Channel Arbitrated Loop (FC-AL) topology. |
| **masking** | Volume-mapping settings that specify no access to that volume by hosts. See also default mapping and explicit mapping. |
| **master volume** | A volume that is enabled for snapshots and has an associated snap pool. |
| **MC** | Management Controller. The processor (located in a controller module) that is responsible for human-computer interface and computer-computer interface functions, and interacts with the SC. |

| | |
|---|---|
| **metadata** | Data in the first sectors of a disk drive that stores all disk, vdisk, and volume specific information including vdisk membership or spare identification, vdisk ownership, volumes and snapshots in the vdisk, host mapping of volumes, and results of the last media scrub. |
| **network port** | The Ethernet port on a controller module through which its Management Controller is connected to the network. |
| **point-to-point** | Fibre Channel Point-to-Point topology. |
| **primary site** | A primary site is a production environment that provides business-critical datacenter services. In a replication environment, the primary site can be any site where Dot Hill AssuredSAN 3000 Series AssuredRemote replication software supports a critical business need. Another term for protected site. See also primary volume. |
| **primary volume** | The volume that is the source of data in a replication set and that can be mapped to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The primary volume exists in a primary vdisk in the primary (or local) storage system. |
| **protected site** | A protected site provides business-critical datacenter services. The protected site can be any site where vCenter supports a critical business need. See also protected site. |
| **proxy volume** | A virtual volume in the local system that represents a volume in a remote system. Proxy volumes are used internally by the controllers to perform actions such as transferring replication data. |
| **recovery site** | A recovery site is an alternative facility in a replication environment to which services and data can be migrated from a protected site. 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. |
| **remote replication** | Asynchronous (batch) replication of block-level data from a volume in a primary system to a volume in one or more secondary systems by creating a replication snapshot of the primary volume and copying the snapshot data to the secondary systems via Fibre Channel or iSCSI links. The capability to perform remote replication is a licensed feature (AssuredRemote). |
| **remote system** | A management object on the local system that enables the MCs in the local system and in the remote system to communicate and exchange data. |
| **replication image** | A conceptual term for replication snapshots that have the same image ID in primary and secondary systems. These synchronized snapshots contain identical data and can be used for disaster recovery. |
| **replication set** | Associated primary and secondary volumes that are enabled for replication and that typically reside in two physically or geographically separate storage systems. See primary volume and secondary volume. |
| **replication snapshot** | A special type of snapshot, created by the remote replication feature, that preserves the state of data of a replication set's primary volume as it existed when the snapshot was created. For a primary volume, the replication process creates a replication snapshot on both the primary system and, when the replication of primary-volume data to the secondary volume is complete, on the secondary system. Replication snapshots are unmappable and are not counted toward a license limit, although they are counted toward the system's maximum number of volumes. A replication snapshot can be exported to a regular, licensed snapshot. See also replication sync point. |
| **replication sync point** | The state of a replication snapshot whose corresponding primary or secondary snapshot exists and contains identical data. For a replication set, four types of sync point are identified: the only replication snapshot that is copy-complete on any secondary system is the "only sync point"; the latest replication snapshot that is copy-complete on any secondary system is the "current sync point"; the latest replication snapshot that is copy-complete on all secondary systems is a "common sync point"; a common sync point that has been superseded by a new common sync point is an "old common sync point." |
| **SAS** | Serial Attached SCSI interface protocol or disk-drive architecture. |
| **SATA** | Serial ATA disk-drive architecture. |

| | |
|---|---|
| **SC** | Storage Controller. The processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. |
| **secondary site** | A backup or recovery site in a replication environment. The secondary site is usually located in a facility that is unlikely to be affected by environmental, infrastructure, or other disturbances that affect the protected site. Another term for recovery site. See also secondary volume. |
| **secondary volume** | The volume that is the destination for data in a replication set and that is not accessible to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The secondary volume exists in a secondary vdisk in a secondary (or remote) storage system.

The contents of a secondary volume are in a constant state of flux and are not in a consistent state while a replication is in process. Only snapshots that are associated with a secondary volume are data consistent. |
| **secret** | For use with CHAP, a password that is shared between an initiator and a target to enable authentication. |
| **SES** | SCSI Enclosure Services. |
| **single-port disk** | A disk that is connected to both controllers so its data path is not fault-tolerant. A single-port disk's type is shown as SAS-S or SATA-S. |
| **snap pool** | A volume that stores data that is specific to snapshots of an associated master volume, including copy-on-write data and data written explicitly to the snapshots. A snap pool cannot be mapped. |
| **snapshot** | A "virtual" volume that preserves the state of a master volume's data as it existed when the snapshot was created. Data associated with a snapshot is recorded in both the master volume and in its associated snap pool. A snapshot can be mapped and written to. The capability to create snapshots is a licensed feature (AssuredSnap). Snapshots that can be mapped to hosts are counted against the snapshot-license limit, whereas transient and unmappable snapshots are not. |
| **SSD** | Solid-state drive. |
| **ULP** | Unified LUN Presentation. A RAID controller feature that enables a host to access mapped volumes through any controller host port. ULP incorporates Asymmetric Logical Unit Access (ALUA) extensions. |
| **unwritable cache data** | Cache data that has not been written to disk and is associated with a volume that no longer exists or whose disks are not online. If the data is needed, the volume's disks must be brought online. If the data is not needed it can be cleared, in which case it will be lost and data will differ between the host and disk. Unwritable cache is also called orphan data. |
| **vdisk** | A "virtual" disk comprising the capacity of one or more disks. The number of disks that a vdisk can contain is determined by its RAID level. |
| **volume** | A portion of the capacity of a vdisk that can be presented as a storage device to a host. |
| **volume copy** | An independent copy of the data in a volume. The capability to create volume copies is a licensed feature (AssuredCopy) that makes use of snapshot functionality. |
| **WWN** | World Wide Name. A globally unique 64-bit number that identifies a node process or node port. |
| **WWNN** | World Wide Node Name. A globally unique 64-bit number that identifies a node process. |
| **WWPN** | World Wide Port Name. A globally unique 64-bit number that identifies a node port. |

# Index