NetApp® FAS/V-Series Storage Replication Adapter 2.1

For VMware® vCenter Site Recovery Manager

Installation and Administration Guide

Contents

What Storage Replication Adapter is

Storage Replication Adapter is a storage vendor specific plug-in to VMware vCenter Site Recovery Manager that enables interaction between Site Recovery Manager and the storage controller.

The adapter interacts with the storage controller on behalf of Site Recovery Manager to discover storage arrays and their associated datastores and RDM devices, which are connected to vSphere. The adapter manages failover and test-failover of the virtual machines associated with these storage objects.

For more information about the new and enhanced features of Site Recovery Manager, see the Site Recovery Manager documentation.

What Site Recovery Manager does

Site Recovery Manager uses Storage Replication Adapter to discover arrays and replicated and exported datastores, and to fail over or test failover of datastores.

Site Recovery Manager uses virtualization to provide end-to-end disaster recovery management and automation across the entire data center. Site Recovery Manager uses VMware infrastructure to build, automate, and test disaster recovery plans for the data center. Site Recovery Manager integrates with vSphere and provides a seamless view of your environment that includes both virtual infrastructure and disaster recovery management.

Related concepts

What Site Recovery Manager supports using Storage Replication Adapter

Using Storage Replication Adapter, Site Recovery Manager supports SAN storage environments, VMFS (iSCSI and FC) and RDM (iSCSI and FC); and NAS storage environments, NFS and HA pair.

For more information about Site Recovery Manager, see the Site Recovery Manager documentation on the VMware site.

SnapMirror modes supported by Storage Replication Adapter

Storage Replication Adapter supports asynchronous qtree SnapMirror, asynchronous volume SnapMirror, and synchronous SnapMirror.

For asynchronous volume SnapMirror, the destination system must be running a Data ONTAP version that is the same as or later than that of the source system if the source and destination systems belong to different Data ONTAP release families. Data ONTAP 7.3 and Data ONTAP 8.0 are examples of different release families.

There is no version requirement if the source and destination systems belong to the same Data ONTAP release family. For example, SnapMirror volume replication is possible from a source system running Data ONTAP 7.3.2 to a destination system running Data ONTAP 7.3.1.

For synchronous or semi-synchronous volume SnapMirror, the source and destination systems must be running the same Data ONTAP version.

There is no version requirement for qtree SnapMirror.

For more information about SnapMirror modes, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Support for IPv6 in Storage Replication Adapter

Starting with Storage Replication Adapter 2.1, IPv6 is supported for management access and data transport over NFS and iSCSI. IPv6 addressing provides expanded routing and addressing capabilities.

You can enable IPv4 and IPv6 on both the primary and secondary disaster recovery sites. You can enable IPv4 or IPv6 on the storage controllers, provided IPv4 and IPv6 are enabled on the associated ESX host. You must enable IPv6 on the storage array.

Components of a disaster recovery environment

A disaster recovery environment consists of the storage array, ESX or ESXi servers, virtual machines, Site Recovery Manager, vSphere Client, and Storage Replication Adapter.

The following illustration shows the components of a disaster recovery environment:
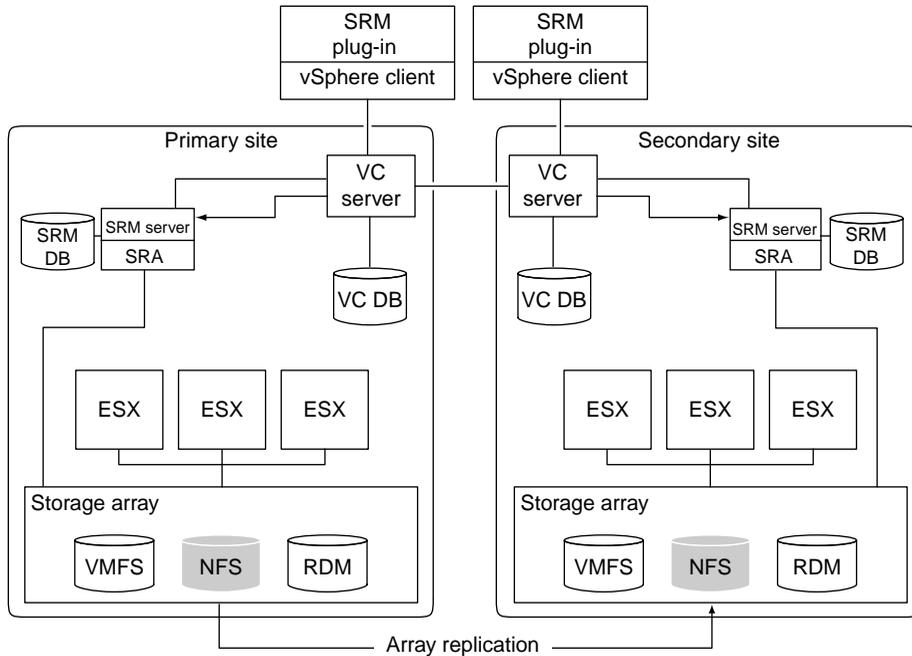
Array replication

Role of Storage Replication Adapter in a disaster recovery environment

Storage Replication Adapter provides array-specific support for Site Recovery Manager by following the input specifications from Site Recovery Manager.

The adapter enables Site Recovery Manager to execute the following workflows:

- Discovery of arrays and replicated devices
- Test recovery
- Recovery
- Reprotect

Discovery of arrays and replicated devices

Storage Replication Adapter discovers pairs of storage arrays that replicate data between the protected and recovery sites. It also discovers the storage devices that are replicated. You can select these devices (*datastores*) as part of the Site Recovery Manager recovery policy.

The discovery operation identifies the source and destination array ID, name and version of the replication software, and model and vendor of the storage array. The operation also discovers all the replicated devices in a pair of storage arrays along with their destination export path or LUN path and device ID. If you want Storage Replication Adapter to discover specific volumes, you can specify the

volumes in the **Volume include list**/**Volume exclude list** of Array Manager in Site Recovery Manager.

- For storage in a NAS environment, the array ID is the host name of the storage system and the device ID is the export path.
- For storage in a SAN environment, the array ID is the host name of the storage system and the device ID is the FC WWN or the iSCSI IQN.

To discover an exported volume or qtree in a NAS environment, you should export the volume or qtree to a vmkernel IP address and replicate the volume or qtree by using SnapMirror technology. This operation discovers the devices that are mapped and then verifies whether these devices are replicated.

To discover a LUN in a SAN environment, you should map the LUN to an igroup of OS type *vmware* and replicate the LUN by using SnapMirror technology. This operation discovers the devices that are mapped and then verifies whether these devices are replicated.

How a test recovery operation works

You should understand how a test recovery operation works so that you can run a test recovery operation to ensure that it works during a disaster or a planned migration. A test recovery operation does not affect an ongoing SnapMirror replication.

During a test recovery operation, Storage Replication Adapter creates a FlexClone volume from the latest Snapshot copy of the SnapMirror destination volume and displays the cloned export path or LUN path to Site Recovery Manager.

The test recovery operation creates a cloned volume on the secondary storage system with the prefix *testfailoverClone_nss_v10745371_*, followed by the parent volume name. For example, the cloned volume that is created for the *secondary_volume* is named *testfailoverClone_nss_v10745371_secondary_volume*.

> **Note:** As volume creation does not support qtrees with special characters such as, space, @ ! $ % ^ & * ( ) + = ? < > - in their names, Storage Replication Adapter creates a clone volume replacing the special characters with underscore (_). However, colon (:) in qtree name is not supported.

At the end of a test recovery operation, the FlexClone volume is destroyed. You must ensure that you do not have user-created volumes with names prefixed with *testfailoverClone_nss_v10745371_*. You can verify the names of the volumes in the secondary storage system by running the `vol status` command.

By default, Storage Replication Adapter creates a cloned volume with no space guarantee because the cloned volume is not used during the actual recovery and is destroyed.

In a NAS environment, the test recovery operation exports the path of the FlexClone volume by adding the vmkernel IP addresses to the `/etc/exports` file. Thus, a clone of the mirrored datastore is ready for discovery. The operation displays the cloned export path to Site Recovery Manager.

In a SAN environment, the test recovery operation maps the LUNs of the FlexClone volume to an igroup of OS type *vmware*. The operation verifies whether an igroup is present for the specific

initiator type and ID. If an igroup is present, the cloned LUN is mapped to that igroup. If an igroup does not exist for a particular initiator, Storage Replication Adapter automatically creates an igroup for that initiator and maps the cloned LUN to that igroup. The operation displays the cloned LUN path to Site Recovery Manager.

Automatically created igroups have names prefixed with *srmtest_v10745371*, followed by the access group name and initiator type received from the Site Recovery Manager input. For example, an igroup created for an FC initiator of access group *host-group-A* has the name *srmtest_v10745371_host-group-A_FC_random_string*.

When to run a recovery operation

You can perform a recovery operation for planned migration of services or for unplanned disaster recovery.

Disaster recovery

During disaster recovery, Storage Replication Adapter cancels any ongoing transfers to the SnapMirror destination, quiesces the destination, and breaks the SnapMirror relationship. When the SnapMirror relationship is broken, the destination volume is ready for read/write operations.

Planned migration

During a planned migration, the recovery operation shuts down the virtual machines at the protected site. Storage Replication Adapter performs a SnapMirror update and synchronizes the storage systems, quiesces the destination, and breaks the SnapMirror relationship. When the SnapMirror relationship is broken, the destination volume is ready for read/write operations.

In a NAS environment, the recovery operation exports the input paths by adding the vmkernel IP address to the `/etc/exports` file. In a SAN environment, the operation maps the input LUN paths to the igroups of OS type *vmware*. If an igroup does not exist for a specific initiator type and ID, Storage Replication Adapter automatically creates an igroup for that initiator and maps that LUN to that igroup. The operation then brings the datastores online and displays the failed-over datastores to Site Recovery Manager.

Automatically created igroups have names in the format access-group_initiator-type_random-string. For example, an igroup created for an FC initiator of access group host-group-A has the name host-group-A_FC_random_string.

Note: You must manually reverse the replication if problems are encountered in the database and if Site Recovery Manager fails to provide replication details for the reprotect operation.

What a reprotect operation does

After a recovery operation, you can run the reprotect operation to replicate the changes made at the recovery site to the original protected site.

During a reprotect operation, Storage Replication Adapter reverses the direction of the existing SnapMirror relationships and starts replicating the changes made at the recovery site to the original protected site.

Before running the reprotect operation, you must have reviewed the SnapMirror Requirements for Storage Replication Adapter section.

Related concepts

[SnapMirror requirements for Storage Replication Adapter](#) on page 11

SnapMirror requirements for Storage Replication Adapter

You must ensure that you meet all the SnapMirror requirements before running a reprotect operation.

For a reprotect operation to succeed, you must ensure that the names of the storage systems are unique. You must also ensure that the names of the volumes that are part of separate SnapMirror relationships are different for the source and destination. For example, the volume names for two different SnapMirror relationships should be similar to the following:

- One SnapMirror relationship with source volume named *srcd16* and destination volume named *destd14*
- The second SnapMirror relationship with source volume named *srcd18* and destination volume named *destd16*

By default, a reprotect operation imports the SnapMirror settings from the original recovery site to the current recovery site. However, the compression and wsize (window size) settings are not imported. You must update these settings manually in the /etc/snapmirror.conf file of the storage controller at the recovery site.

Note: The default value for snap autodelete is set to try to prevent Data ONTAP from deleting the SnapMirror Snapshot copies. You must retain the default setting of snap autodelete because if the setting is modified and the SnapMirror Snapshot copies are deleted, then Storage Replication Adapter cannot reverse and resynchronize the replication from the affected volume.

Failback (using reprotect and recovery workflows)

Failback procedure enables you to restore the original protected and recovery sites after a recovery operation.

To perform failback after recovery, you must execute the following workflows:

1. Reprotect

Reverses the replication direction of the existing SnapMirror relationships from the recovery site to the protected site.

If the devices and the SnapMirror relationships do not exist after a disaster, you should create new devices at the protected site and initialize the SnapMirror relationship from the recovery site to the original protected site.

2. Planned migration

   Migrates virtual machines from the recovery site back to the protected site.

3. Reprotect

   Reverses the replication direction from the protected site to the recovery site, thus restoring the original direction of replication between the protected and recovery sites.

Installing or upgrading Storage Replication Adapter

You should ensure that you meet the system requirements to install Storage Replication Adapter. You can upgrade to 64-bit Storage Replication Adapter 2.1.

System requirements for using Storage Replication Adapter

Before you run Storage Replication Adapter, you must meet certain requirements related to Data ONTAP and Site Recovery Manager.

* Site Recovery Manager must be installed on Windows.
  Site Recovery Manager supports VMFS (iSCSI and FC), RDM (iSCSI and FC), and NFS storage environments.
* Your storage system must be running a supported version of Data ONTAP operating in 7-Mode (Data ONTAP 7.2.4 or later).
  For configuring vFiler units as storage arrays, your storage system must be running Data ONTAP 7.3.2 or later.
* FlexClone license can be enabled.
  The license is optional because a test failover operation can be performed when the option `replication_disruptive_test_failover` is set to `on` in the `ontap_config` file.
* SnapMirror license must be enabled on the storage system.
* In a NAS environment, if the `nfs.ipv6.enable` option is set to `on`, then NFS services must be restarted.
* In a SAN environment, iSCSI, FC, or both the services must be running.
* Storage Replication Adapter must be installed on the Site Recovery Manager server at the protected and recovery sites.
* Perl 5.14.2 must be installed on the system on which Site Recovery Manager is installed.

  Note: If the Site Recovery Manager installation does not include the required Perl version, you can install Perl separately with the default configuration. For more information about downloading Perl, see the Perl download page.

* To use secureHTTP, SecureAdmin must be set up and SSL or SSH must be running.
  For more information about installing SSL on Perl, see KB article 1012531.
  By default, Storage Replication Adapter calls the Data ONTAP API over HTTP.

  Note: SecureAdmin does not require a license.

* The 80 or 443 port corresponding to the HTTP or SSL protocol must be open on the firewall.
  The port is between the Site Recovery Manager server and the storage system running Data ONTAP.

- For Data ONTAP 8.2.x, the cluster should be enabled with IPv6 before enabling IPv6 for data LIF and cluster-management LIF.

For a detailed list of the supported software version and storage systems, see the Interoperability Matrix.

For information about the NetApp platforms that are supported in the VMware environment, see the NetApp-VMware platform compatibility matrix on the VMware site.

Related information

VMware Interoperability Matrix: www.vmware.com/resources/compatibility/search.php? deviceCategory=srax
Perl script download page: www.perl.com/download.csp
NetApp Support Site: support.netapp.com
How to enable SSL communication between VMware vCenter Site Recovery Manager and NetApp FAS Storage Arrays: kb.netapp.com/support/index?page=home
Interoperability Matrix: support.netapp.com/matrix

Downloading Storage Replication Adapter

Before you install Storage Replication Adapter, you have to download the installation file from the NetApp Support Site.

Steps

1. Log in to the NetApp Support Site and click the **Downloads** tab.

2. In the **Downloads** page, select **Software**.

3. From the list of products, select **FAS/V-Series Storage Replication Adapter for VMware SRM**.

4. Follow the instructions on the page until you reach the download page.

5. Download the file in one of the following ways:

   - Download it directly to your target system.
   - Download it to a PC host and copy it to the target system.

After you finish

You can unzip the file and install Storage Replication Adapter.

Related tasks

Installing Storage Replication Adapter on page 15

Related information

[NetApp Support Site: support.netapp.com](support.netapp.com)

Installing Storage Replication Adapter

You can install Storage Replication Adapter from the location where you downloaded the scripts.

Steps

1. Use Windows Explorer to navigate to the directory to which you unzipped the downloaded file, and then double-click the setup.msi file.

    The installer finds the installed location of Site Recovery Manager from the registry and copies these scripts to that location.

2. Select the **Complete install** option.

    The Storage Replication Adapter version is available at the following registry location: `HKEY_LOCAL_MACHINE\SOFTWARE\NetApp, Inc\NetApp FAS/V-Series Storage Replication Adapter`.

3. Click **Finish**.

Upgrade considerations for Storage Replication Adapter

If you are using Site Recovery Manager 5.1, you can upgrade to the 64-bit Storage Replication Adapter.

Starting with Storage Replication Adapter 2.1, VMware datastores and RDM devices that are replicated in the same SnapMirror relationship are reported to the Site Recovery Manager as a Site Recovery Manager consistency group. This ensures that the devices are contained in the same protection group to fail back and to fail over together.

> **Note:** The Site Recovery Manager consistency group does not guarantee application consistency between different virtual machines sharing the same SnapMirror replication relationship.

If multiple VMware datastores or RDM devices are sharing the same SnapMirror replication relationship but are contained in separate protection groups, then you must re-create these protection groups so that the datastores or RDM devices sharing the same SnapMirror replication relationship are in the same protection group. You can re-create the protection groups before or after the upgrade procedure. Site Recovery Manager considers the affected protection groups as invalid after the Storage Recovery Adapter upgrade process is complete.

Supported configurations for a consistency group

Following are the supported configurations and scenarios for a consistency group:

*   NFS volumes with multiple qtrees in a qtree SnapMirror relationship

- NFS volumes with multiple qtrees in a volume SnapMirror relationship
- Volumes with multiple qtrees and each qtree with a single LUN in a volume SnapMirror relationship
- Volumes with multiple qtrees and each qtree with multiple LUNs in a volume SnapMirror relationship
- Volumes with a single qtree with a single LUN in a qtree SnapMirror relationship
- SAN volumes with multiple LUNs in a volume SnapMirror relationship

Upgrading to 64-bit Storage Replication Adapter

You can upgrade to 64-bit version of the Storage Replication Adapter 2.1.

Before you begin

- If you have created any configuration settings in the `ontap_config.txt` file of the earlier version of the adapter, create a copy of that file to use those configuration settings in Storage Replication Adapter 2.1.
- You must have downloaded the Site Recovery Manager installer from the VMware site.

Steps

1. Upgrade Site Recovery Manager using the Site Recovery Manager installer.

   For more information about Site Recovery Manager, see the Site Recovery Manager documentation on the VMware site.

2. Install Storage Replication Adapter 2.1.

3. Configure the required settings in the `ontap_config.txt` file.

4. Pair the Site Recovery Manager sites.

5. Reload Storage Replication Adapter 2.1 from the Site Recovery Manager user interface.

6. Configure Storage Replication Adapter 2.1 in the **Array Manager** wizard of Site Recovery Manager.

7. Verify that all the required replicated datastores are discovered by Site Recovery Manager.

Setting up Storage Replication Adapter

You must set up the storage system before running Storage Replication Adapter for Site Recovery Manager. You can customize Storage Replication Adapter by changing the configuration files.

Setting up the storage system in a NAS environment

You must set up the system before running Storage Replication Adapter for Site Recovery Manager.

Steps

1. Install Site Recovery Manager.

   For information about installing Site Recovery Manager, see the Site Recovery Manager documentation on the VMware site.

2. Install Storage Replication Adapter on the Site Recovery Manager server at the protected and recovery sites.

3. Ensure that the datastores at the protected site contain virtual machines that are registered with vCenter Server.

4. Ensure that the ESX or ESXi hosts at the protected site have mounted the NFS exported volumes from the storage controller.

5. Ensure that valid addresses, such as the IP address, host name, FQDN, are entered in the **NFS Addresses** field when adding arrays to Site Recovery Manager by using the **Array Manager** wizard.

6. Verify that the ESX or ESXi hosts at the secondary storage system have a VMkernel port that can access the IP addresses that are used to serve NFS exports from the secondary storage controller or a vFiler unit.

   You can do this by using the `vmkping nfs_ip_address` command on the ESXi shell or ESX service console of each ESXi or ESX host.

   `nfs_ip_address` is one of the NFS IP addresses on the storage controller or a vFiler unit.

   You can verify whether access to the ESXi shell is disabled by using the `ping vmkernel_ip_address` command on the console of the storage controller or a vFiler unit.

   `vmkernel_ip_address` is one of the IP addresses on the VMkernel port that the ESX or ESXi host uses to access the storage controller or a vFiler unit.

7. Ensure that the NFS exports on the primary storage controller are exported to the ESX or ESXi hosts by using the values in the `rw` and `root` fields and that the export is contained in the `/etc/exports` file.

> **Note:** If the export is created manually and not listed in the `/etc/exports` file, the exports are not discovered with the default export setting of `rw` to all the hosts.

8. Ensure that the volume or qtree that contains the exports is replicated to the secondary storage system.

   If VMware datastores are provisioned in qtrees, then you should use a qtree SnapMirror operation to replicate the datastores. If VMware datastores are provisioned in volumes, then you should use a volume SnapMirror operation to replicate the datastores.

   Storage Replication Adapter does not support SnapMirror operations of non-qtree data and therefore the adapter does not discover Snap Mirror relationships. For example, the following format is not discovered: `source_system:/vol/source_volume/- dest_system:/vol/dest_volume/qtree_name`. The hyphen (-) indicates non-qtree data in the specified volume.

   If you have specified the IP address of the source storage array for the SnapMirror relationship, you must set `use_ip_for_snapmirror_relation` to `on` and map the IP addresses of both the source and destination storage arrays to their host names in the `ip_hostname_mapping.txt` file at the protected and recovery sites.

   If you have specified a SnapMirror connection name as the source storage array for the SnapMirror relationship, you must ensure that you use the connection name that matches the host name of the source storage system.

   If you use a SnapMirror connection name that does not match the host name of the source storage system, then mapping information must be added in the `snapmirror_connection.txt` file at the protected and recovery sites. The primary site includes the mapping information of the destination storage array with the connection name, and the recovery site includes the mapping information of the source storage array with the connection name.

   > **Note:** If you do not add mapping information in any of the protected and recovery sites, Storage Replication Adapter ignores the connection name and proceeds without performing the replication operation. If you have added the mapping information in the recovery site but have omitted to add it in the source site, then the reprotect operation fails. For more information see IP to host name mapping information not found topic.

9. On the primary storage system, enter the `snapmirror status` command to confirm that the export is part of only one replication relationship and that the relationship status is `Source`.

10. On the secondary storage system, enter the `snapmirror status` command to confirm that the relationship state is `snapmirrored`.

    For more information about SnapMirror modes, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

11. If the `snapmirror status` command on the destination storage system shows the IP address of the source storage array, configure the `use_ip_for_snapmirror_relation` option in the configuration file or reconfigure the SnapMirror relationship to use the host name instead of the IP address.

Related concepts

What the Storage Replication Adapter configuration file does on page 21

Related references

IP to host name mapping information not found on page 30

Related information

NetApp Support Site: support.netapp.com

Setting up the storage system in a SAN environment

You must set up the system before running Storage Replication Adapter for Site Recovery Manager.

Steps

1. Install Site Recovery Manager.

   For information about installing Site Recovery Manager, see the Site Recovery Manager documentation on the VMware site.

2. Install Storage Replication Adapter on the Site Recovery Manager server at the protected and recovery sites.

3. Ensure that the primary ESX or ESXi hosts are connected to LUNs in the primary storage system.

4. Ensure that the LUNs, iSCSI, FC, or both, are in igroups of OS type *vmware* on the primary storage system.

5. Ensure that the ESX or ESXi hosts at the recovery site have appropriate FC or iSCSI connectivity to the secondary storage controller or a vFiler unit.

   You can do this either by verifying that the ESX or ESXi hosts have local LUNs connected on the secondary storage controller or a vFiler unit, or by using the `fcp show initiators` or `iscsi show initiators` commands on the secondary storage controllers.

6. Ensure that the volume or qtree that contains the LUNs is replicated to the secondary storage system.

   If VMware datastores are provisioned in qtrees, then you should use qtree SnapMirror operation to replicate the datastores. If VMware datastores are provisioned in volumes, then you should use volume SnapMirror operation to replicate the datastores.

   If you have specified the IP address of the source storage array for the SnapMirror relationship, you must set `use_ip_for_snapmirror_relation` to `on` and map the IP addresses names of both the source and destination storage arrays to their host names in the `ip_hostname_mapping.txt` file at the protected and recovery sites.

If you have specified a SnapMirror connection name as the source storage array for the SnapMirror relationship, you must ensure that you use connection name that matches the host name of the source storage system.

If you use a SnapMirror connection name that does not match the host name of the source storage system then mapping information must be added in the `snapmirror_connection.txt` file at the protected and recovery sites. The primary site includes the mapping information of the destination storage array with the connection name, and the recovery site includes the mapping information of the source storage array with the connection name.

> **Note:** If you do not add mapping information in any of the protected and recovery sites, Storage Replication Adapter ignores the connection name and proceeds without performing replication operation. If you have added the mapping information in the recovery site but have omitted to add in the source site then reprotect operation fails. For more information see IP to host name mapping information not found topic.

7. On the primary storage system, enter the `snapmirror status` command to confirm that the LUN is part of only one replication relationship and that the relationship status is `Source`.

8. On the secondary storage system, enter the `snapmirror status` command to confirm that the relationship state is `snapmirrored`.

   For more information about SnapMirror modes, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

9. If the `snapmirror status` command on the destination storage system shows the IP address of the source storage array, configure the `use_ip_for_snapmirror_relation` option in the configuration file or reconfigure the SnapMirror relationship to use the host name instead of the IP address.

10. When using vFiler units as storage arrays in Site Recovery Manager, ensure that the iSCSI service is started on the vFiler units on the secondary storage system.

Related concepts

What the Storage Replication Adapter configuration file does on page 21

Related references

IP to host name mapping information not found on page 30

Related information

NetApp Support Site: support.netapp.com

What the Storage Replication Adapter configuration file does

The Storage Replication Adapter configuration file enables you to customize some of the disaster recovery behaviors, including the security settings during Site Recovery Manager recovery or test recovery. By default, the configuration file is located at `install_dir\Program Files\VMware \VMware vCenter Site Recovery Manager\storage\sra\ONTAP\ontap_config.txt`.

The following are some of the options that you can configure in the adapter configuration file:

- `use_ip_for_snapmirror_relation` (value can be `on` or `off` (default))

  If you have specified the IP address or the connection name of the storage array for the SnapMirror relationship, you must set this option to `on` and map the IP addresses or connection names of both the source and destination storage arrays to their host names (*host_name = IP_address/connection_name*) in the `ip_hostname_mapping.txt` file at the protected and recovery sites.

  For example, protected site A and recovery site B have storage arrays with host names hostA (IP address: 10.10.0.1) and hostB (IP address: 10.10.0.2), respectively. If you have established a SnapMirror relationship from site A to site B using the IP addresses of the storage arrays, you must have the following entry in the `ip_hostname_mapping.txt` file at both site A and site B:

  hostA = 10.10.0.1
  hostB = 10.10.0.2

  The `ip_hostname_mapping.txt` file must have the same entries in both the protected and recovery sites. Also, the entries are case-sensitive.

  The mapping of the IP address or the connection address to its host name is required because Storage Replication Adapter uses the host name to identify the source and destination storage arrays.

  When this option is set to `on`, Storage Replication Adapter reports only those devices that use the IP address or connection address for the SnapMirror relationship. When it is set to `off`, Storage Replication Adapter reports only those devices that use the host name for the SnapMirror relationship.

- `import_replication_settings` (value can be `on` (default) or `off`)

  This option imports the SnapMirror settings at the original recovery site to the required recovery site during the reprotect operation. However, compression and wsize (window size) settings are not imported and those settings must be updated manually in the `/etc/snapmirror.conf` file.

- `volume_guarantee` (value can be `none`, `file`, or `volume` (default))

  This option controls whether the volume is guaranteed some amount of disk space.

- `replication_disruptive_test_failover` (value can be `on` or `off` (default))

  The FlexClone license is optional if the value is set to `on` because FlexClone volumes are not created. The replication is disrupted during disaster recovery test failover operations for all executed recovery plans when `replication_disruptive_test_failover` is set to `on`.

- `isipv4` (values can be `yes` or `no` (default))

  This option is used to support the mounting of datastores by host name and FQDN instead of IPv6. You should modify the `isipv4` option in `ontap_config.txt`, depending on whether the IP address is IPv6 or IPv4.

  You must add the IPv4 address in the NFS Address field if `isipv4` is set to `yes`. If the host name or FQDN is added, then Storage Replication Adapter resolves the host name with the IPv4 address.

  You must add the IPv6 address in NFS Address field if `isipv4` is set to `no`. If the host name or FQDN is added, then Storage Replication Adapter resolves the host name with the IPv4 address.

Option specific to a SAN environment

- `lun_space_reservation` (value can be `yes` or `no` (default))

  This is a provisioning management option that verifies whether the LUNs recovered for disaster recovery testing are set with space reservations enabled or disabled.

Options specific to a NAS environment

- `root` (value can be `yes` (default) or `no`)

  By default, the adapter exports all the volumes that have root access to the ESX or ESXi hosts. If your environment does not use the `root` field on exports, you can set this option to `no`, preventing this access. The `root` option is useful if you have configured your environment to use a nonroot user to access datastores. You might want to prevent root access to datastores recovered by Site Recovery Manager.

- `nosuid` (value can be `yes` or `no` (default))

  This option is used to specify whether the adapter exports all the volumes with the `nosuid` property.

- `anon` (*anon uid*)

  This option is used to force volumes to be exported with an anonymous security identifier. You should enter the required `anon` user ID—for example, `anon=0`. By default, the adapter does not export volumes that are using the `anon` option.

- `sec` (default value is `sys`)

  This option enables you to customize the NFS export security type.

By default, during a recovery or test recovery operation, the adapter exports volumes only to those IP addresses directed by Site Recovery Manager. The following three options of `ontap_config.txt` enable you to perform some additional customization of export hosts:

- `read-write-hosts`

  This option is to add a comma-separated list of additional hosts to be added to the `rw` field of the NFS exports.

- `read-only-hosts`

  This option is used to add a comma-separated list of additional hosts to be added to the `ro` field of the NFS exports.

- `root-hosts`

This option is used to add a comma-separated list of additional hosts to be added to the root field of the NFS exports. This option is not dependent on the value of the `root=yes/no` option. The `root-hosts` option is useful if you have to ensure that an additional host is granted root access for the purpose of virtual disk file management or backup.

Configuring Storage Replication Adapter

You can configure Storage Replication Adapter by using the Array Manager wizard of Site Recover Manager.

Steps

1. In Site Recovery Manager, click **Array Managers**, and then click **Add Array Manager**.

2. Enter the following information to describe the array in Site Recovery Manager:

   a) Enter a name to identify the storage controller in **Display Name**.
   b) In the **SRA Type** field, select **NetApp FAS/V-Series Storage Replication Adapter**.
   c) Enter the IP address of the storage controller in **IP Address of Storage System**.

   ### Example

   ```
   10.10.20.54
   ```

   d) In **NFS Addresses**, enter the IP address, host name, or DNS name of the controller that is used to serve NFS to the ESX or ESXi hosts at the site where this controller is located.

   ### Example

   ```
   192.168.2.100, f3170-192-75
   ```

   You must enter the addresses that are used by the storage controller to serve NFS to the ESX or ESXi hosts. You can enter multiple addresses, separated by commas. These NFS addresses allow the use of private storage networks that are separate from the network that is used to administer the controller. If you do not enter any addresses in this field, Site Recovery Manager attempts to mount the NFS exports by using the IP addresses entered for managing the storage controller.

   e) Enter the volumes to be discovered in **Volume include list**.

   You can either enter the full volume name or strings at the protected site and the replicated destination volume at the recovery site.

   ### Example

   If you want to discover the volume *src_vol1* that is in a SnapMirror relationship with volume *dst_vol1*, you must specify either *src_vol1* or *src* in the **Volume include list** at the protected site and *dst_vol1* or *dst* in the volume include list at the recovery site.

   f) Enter the volumes to be excluded from discovery in **Volume exclude list**.

You can either enter the full volume name or strings at the protected site and the replicated destination volume at the recovery site.

Example

If you want to exclude the volume *src_vol1* that is in a SnapMirror relationship with volume *dst_vol1*, you must specify either *src_vol1* or *src* in the protected site and *dst_vol1* or *dst* in the recovery site.

Following are some of the scenarios for Volume include and Volume exclude list:

- If both the Volume include list and Volume exclude list fields have data, Storage Replication Adapter discovers all volumes with names that match one of the sub-strings that are specified in the Volume include list and also discovers all volumes with names that do not have any of the sub-strings that are specified in the Volume exclude list
- If both the Volume include list and Volume exclude list do not contain data, all the volumes are discovered
- If Volume include list is empty and Volume exclude list has data then all volumes with names that do not match the sub-strings specified in exclude list are discovered
- If Volume exclude list is empty and Volume include field has data then all volumes with names that match the sub-strings specified in include list are discovered

Example

Volumes on storage arrays are as follows: srm_sql_oracle_db2, srm_exchange, srm_1, sql_2, access

- Data in Volume include list is as follows: srm, oracle, sql
- Data in Volume exclude list is as follows: db2, exchange

Storage Replication Adapter discovers all volumes with names that match the sub-strings specified in the Volume include list. Next, Storage Replication Adapter filters or removes all volumes with names that match the sub-string specified in the Volume exclude list. Volumes discovered are as follows:

- srm_1
- sql_2

    g) Enter the user name of an account with privileges to this storage controller in **Username**.

    h) Enter the password of the user account in **Password**.

If the storage system that is being added is in an HA pair, each controller must be added separately as an array.

3. Click **Next**.

4. Verify that the array is discovered and displayed at the bottom of the **Add Array Manager** window.

5. Click **Finish**.

Storage Replication Adapter requirements for vFiler unit support

You must meet certain requirements so that Storage Replication Adapter can support vFiler units, such as adding each vFiler unit to the adapter as a separate array and ensuring that your storage system is running Data ONTAP 7.3.2 or later.

You must meet the following requirements for vFiler unit support:

- The storage system must be running Data ONTAP 7.3.2 or later.
- Each vFiler unit must be added to the adapter as a separate array.
- Both the source and destination vFiler units must be online.
- The `httpd.admin.enable` option on the vFiler unit must be enabled.
- SnapMirror relationships must be defined in the destination vFiler context and not in the destination physical vfiler0 context.
- iSCSI or NFS connected datastores must be supported for vFiler units.
- SCSI service must be started on the vFiler unit.
  The adapter does not automatically start the iSCSI service.
- In a SAN environment, the recovery site ESX or ESXi hosts must have established iSCSI sessions to the recovery vFiler units before you perform the test recovery or recovery operation.
- A single virtual machine must not have data on multiple vFiler units.
- Before performing a test recovery operation on a vFiler unit, the `vfiler.vol_clone_zapi_allow` option must be turned on for vfiler0.
- IPv6 must be supported for vFiler units.
- You must have set the `nfs.ipv6.enable` ption to `ON` and restarted the NFS service.

  Note: SSL and synchronous SnapMirror mode are not supported for vFiler units.

Error messages in Storage Replication Adapter

You might encounter an error message when using Storage Replication Adapter with Site Recovery Manager. The user interface of Site Recovery Manager displays the error log messages and possible solutions.

You can also use the log files of Site Recovery Manager and the log files of the controller to troubleshoot issues that you might encounter when using Storage Replication Adapter.

Address of the storage array is not reachable

Message	Address of the storage array is not reachable
Cause	Storage array might be down or IP address entered might be incorrect.
Corrective action	Ensure that the storage array is up and running and the IP address of the storage array is reachable through the command line interface.

An error occurred during host configuration

Message	An error occurred during host configuration
Cause	The device might not be exported or mapped correctly.
Corrective action	For a NAS device, ensure that it is exported with proper rules and vmkernel IP addresses. Also ensure that the NFS service is started. For a SAN device, ensure that it is mapped to an igroup of OS type vmware.

An error occurred during host configuration and failed to recover datastore

Message	Failed to recover datastore. An error occurred during host configuration. Recovery of device already completed. You must configure the SnapMirror relationship again to recover the device.
Cause	The SnapMirror relationship for the device might not exist.
Corrective action	Ensure that the SnapMirror relationship exists for the device and that the relationship status is either snapmirrored or quiesced.

Array ID and peer array ID are the same

Message   `Array ID and peer array ID are the same`

Cause   SRA does not support this type of configuration.

Corrective action Ensure that the protected site array and the recovery site array are different.

Authentication failed

Message   `Authentication failed`

Cause   Invalid user name or password.

Corrective action Ensure that you provide a valid user name and password.

Creation of a temporary writable copy of the device failed

Message   `Creation of a temporary writable copy of the device failed`

Cause   The FlexClone license might not be enabled for the storage array.

Corrective action Ensure that the FlexClone license is enabled for the storage array.

Device found is neither of the SAN type nor of the NAS type

Message   `Device found is neither of the SAN type nor of the NAS type.`

Cause   Device might not exist on the storage array.

Corrective action Ensure that the device exists on the storage array and is of type NAS or SAN.

Device is already tested for recovery

Message   `Device is already tested for recovery`

Cause   The test recovery operation is either completed or is already running on the device.

Corrective action	If a test recovery operation is running, wait for the operation to complete and then clean up the test environment. If the test recovery is already completed, clean up the test environment.

Device is not in the snapmirrored state

Message	`Device is not in the snapmirrored state`
Cause	The SnapMirror relationship might be broken or uninitialized.
Corrective action	Ensure that the replication state of the device is snapmirrored for synchronization to work.

Device is not replicated

Message	`Device is not replicated`
Cause	This error message can occur if any one of the following conditions exists:

- SnapMirror relationship is not configured on the device.
- The device provided in the input is not online.
- The device provided in the input does not exist.

Corrective action	Ensure that a SnapMirror relationship is configured on the device. Also ensure that the device exists on the storage system and is online.

Device synchronization did not complete properly

Message	`Device synchronization did not complete properly`
Cause	The device synchronization process might have been disrupted because of network failure.
Corrective action	Ensure that the storage array hosting the device is connected to the network and is accessible to its peer storage array.

Failed to promote replica devices

Message	`Failed to promote replica devices. Base snapshot not found. Ensure that the SnapMirror relationship exists for the device. Verify whether base snapshot exists for SnapMirror relationship.`

Cause	The volume might not be online or the base Snapshot copy might be deleted.
Corrective action	Ensure that the specified volume is online and that the base Snapshot copy exists for the SnapMirror relationship.

Failed to recover datastore while the pseudo and actual path names are the same

Message	`Failed to recover datastore.`
Cause	The pseudo path name and actual path name are the same.
Corrective action	Ensure that the pseudo and actual path names are different in `/etc/exports`.

Failed to recover datastore

Message	`Failed to recover datastore. Recovered device not found after HBA rescan.`
Cause	FC and iSCSI adapters and vmkernel IP might not be configured on ESX servers.
Corrective action	Ensure that at least one initiator is configured on the ESX server.

FlexClone is not licensed

Message	`FlexClone is not licensed`
Cause	The FlexClone license might not be enabled on the storage system.
Corrective action	Ensure that the FlexClone license is enabled on the storage system.

Importing SnapMirror Connection information failed

Message	`Failed to reverse replication for device. Importing Snapmirror Connection information failed. Please make sure the snapmirror_mapping file is updated.`
Cause	Host name of the storage system is not mapped to the connection name in the `snapmirror_mapping.txt` file.

Corrective action	Map the host names of the storage systems to connection names in the snapmirror_mapping.txt file to ensure that the protected site has an entry for the destination storage array and the recovery site has an entry for the source storage array.

Incorrect array ID

Message	Incorrect array ID
Cause	Incorrect array ID in the SRM input.
Corrective action	Verify the SRM logs.

Incorrect synchronization ID received as input from SRM

Message	Incorrect synchronization ID received as input from SRM
Cause	Synchronization ID might not be present on the storage system.
Corrective action	Check the SRM logs to verify the synchronization ID.

Invalid SnapMirror configurations

Message	Invalid SnapMirror configurations
Cause	SnapMirror configurations specified in the /etc/snapmirror.conf file might be incorrect.
Corrective action	Ensure that the SnapMirror configurations are valid in the configuration file.

Related concepts

What the Storage Replication Adapter configuration file does on page 21

IP to host name mapping information not found

Message	IP to host name mapping information not found
Cause	Host name of the storage system is not mapped to its IP address in the ip_hostname_mapping.txt file.

Corrective action	Map the host names of the storage systems at the protected and recovery sites to its IP addresses in the `ip_hostname_mapping.txt` file located at the protected and recovery sites.

Related concepts

Location for reverse replication is not source

Message	`Location for reverse replication is not source`
Cause	Reverse replication might be performed at the recovery site.
Corrective action	Check the logs to ensure that reverse replication is running on the source. Reverse replication is performed at the protected site (original recovery site) to reverse the replication direction after a failover.

LUN could not be mapped as requested

Message	`LUN could not be mapped as requested. LUN cannot be mapped to both ALUA enabled FCP igroup and ALUA disabled iSCSI igroup at the same time. SRA does not support mixed FC/iSCSI configuration on same ESX host.`
Cause	When both iSCSI and FCP are enabled at the same time on the ESX host at a recovery site, SRM attempts to map the LUN to both igroups that have different asymmetric logical unit addressing (ALUA) settings (iSCSI with default ALUA disabled and FCP with ALUA enabled).
Corrective action	Ensure that only one protocol, either FC or iSCSI, is enabled at the same time on the recovery ESX hosts because Storage Replication Adapter does not support a mixed ALUA configuration.

No SnapMirror relationship found

Message	`No SnapMirror relationship found`
Cause	The SAN or NAS devices might not be mirrored on separate storage arrays.
Corrective action	Ensure that at least one SAN or NAS device is mirrored and the source and destination are on different storage arrays. SRA does not support SnapMirror relationships on a single storage array.

Peer array ID provided in the SRM input is incorrect

Message	`Peer array ID provided in the SRM input is incorrect`
Cause	SnapMirror relationship might not exist for the peer array ID.
Corrective action	Check the SRM logs to verify the peer array ID.

Recovery of device already completed

Message	`Recovery of device already completed`
Cause	SnapMirror relationship of the device might be in broken.
Corrective action	Configure the SnapMirror relationship to recover the device.

SAN or NAS device not found

Message	`SAN or NAS device not found`
Cause	The SAN or NAS device might not be configured properly.
Corrective action	Ensure that the SAN device is configured and mapped to an igroup of OS type *vmware* and the NAS device is configured and exported with `rw` rules.

SnapMirror replication is in progress

Message	`SnapMirror replication is in progress`
Cause	Synchronization as part of test recovery, recovery, or reprotect operations might already be running for the device.
Corrective action	Wait for the replication to be completed. Check the status of the SnapMirror by using the `snapmirror status` command.

SnapMirror resynchronization failed

Message	`SnapMirror resynchronization failed`
Cause	The devices might not be configured correctly.

Corrective action	Ensure that the devices are configured correctly. Check the log files for more details.

SnapMirror resynchronization failed because of various reasons

Message	`SnapMirror resynchronization failed. Ensure that the device is in online state and is mapped or exported properly / SnapMirror resynchronization failed. Ensure that the devices are configured correctly. Check the log files for more details.`
Cause	The source volume is offline.
Corrective action	Ensure that the volume is online.

SnapMirror version could not be traced

Message	`SnapMirror version could not be traced`
Cause	Unable to get the Data ONTAP version of the storage array.
Corrective action	Ensure that the `version` command on the storage array displays the correct Data ONTAP version.

Storage IP is not defined or valid

Message	`Storage IP is not defined or valid`
Cause	The IP address provided for the storage system is either null or invalid.
Corrective action	Provide a valid IP address for the storage system.

Storage system configuration error

Message	`Storage system configuration error`
Cause	FCP, iSCSI, and storage IP might not be configured properly.
Corrective action	Ensure that the storage system has proper configuration of FCP, iSCSI, and storage IP.

System information: host name, model, and vendor not found

Message	`System information: host name, model, and vendor not found`
Cause	SRA was not able to collect the storage array details.
Corrective action	Ensure that the `hostname` command displays the host name for the storage array. Also ensure that the model and vendor is defined for the storage array.

The access group initiator information is missing in the SRM input

Message	`The access group initiator information is missing in the SRM input`
Cause	SRA supports dynamic access restriction. Therefore, it expects SRM to provide the access group information in the input.
Corrective action	Ensure that FC and iSCSI initiators and vmkernel IP addresses are configured on ESX servers. Check SRM logs for more details.

The SnapMirror license is not present

Message	`The SnapMirror license is not present`
Cause	The SnapMirror license might not be enabled.
Corrective action	Ensure that the SnapMirror license is enabled.

The SRA command issued was not found

Message	`The SRA command issued was not found`
Cause	Command issued by SRM is not supported by SRA.
Corrective action	Check the command line options and logs.

The temporary device is in improper state

Message	`The temporary device is in improper state`
Cause	The temporary device might be in restricted state.
Corrective action	Manually make the device offline.

The temporary writable copy created is not consistent

Message	`The temporary writable copy created is not consistent`
Cause	The temporary writable copy might not be online.
Corrective action	Ensure that the temporary writable copy is in the online state.

The temporary writable copy might not have enough space

Message	`The temporary writable copy might not have enough space`
Cause	Enough space might not available for the creation of temporary writable copy.
Corrective action	Ensure that you reserve enough space for creating the temporary writable copy.

The XML input to SRA from SRM is incorrect

Message	`The XML input to SRA from SRM is incorrect`
Cause	XML input might not be in correct format.
Corrective action	Check the SRM logs for more details.

Unable to add the initiators passed in the SRM input to an initiator group

Message	`Unable to add the initiators passed in the SRM input to an initiator group`
Cause	The initiators passed in the SRM input might not adhere to the required naming convention.

Corrective action	Ensure that the initiators passed in the SRM input adhere to the required naming convention. For example, ensure that FC has a proper WWPN or WWNN. iSCSI has a proper IQN, and NFS has a proper IPv4 address.

Unable to create initiator group on the storage array based on the access group information

Message	`Unable to create initiator group on the storage array based on the access group information`
Cause	The access group information to the SRA might not contain a proper alphanumeric name.
Corrective action	Check the SRM logs to ensure that a proper alphanumeric name is specified as part of the access group information to the SRA.

Unable to destroy the temporary writable copy

Message	`Unable to destroy the temporary writable copy`
Cause	The temporary device might be in use.
Corrective action	Ensure that temporary writable copy is not used by other operations.

Unable to fail over the specified device

Message	`Unable to fail over the specified device`
Cause	Either the device is not in the online state or is not mapped properly.
Corrective action	Ensure that the device is in the online state and is mapped or exported properly.

Unable to make the device online

Message	`Unable to make the device online`
Cause	The device might be in an inconsistent state.
Corrective action	Use the Data ONTAP commands to bring the device online.

Glossary

You should be familiar with basic storage concepts, terms, and technologies when using Storage Replication Adapter with VMware vCenter Site Recovery Manager.

Aggregate	In a NetApp FAS system, an aggregate is a collection of many physical disks that are collected into RAID groups to provide performance and disk redundancy. In a NetApp V-Series system, an aggregate is a collection of LUNs that have been provisioned on a supported third-party storage platform to provide storage for the NetApp system.
FlexClone volume	FlexClone volumes are writable, point-in-time copies of a parent FlexVol volume. FlexClone volumes and their parent volumes share the same disk space for any common data. Therefore, creating a FlexClone volume is instantaneous and requires no additional disk space until changes are made to the FlexClone volume or its parent. When performing a test recovery in VMware vCenter Site Recovery Manager, FlexClone volumes are created in the testing environment, and these FlexClone volumes are removed when testing is complete.
FlexVol volume	A FlexVol volume is provisioned from the available storage in an aggregate. FlexVol volumes are flexible and can be increased or decreased in size dynamically without impacting or disrupting the environment. A single aggregate can contain many FlexVol volumes. A FlexVol volume is not tied to any specific set of disks in the aggregate and is striped across all the disks in the aggregate.
Initiator	An initiator is a port that is used to connect to a LUN. It can be on an FC or an iSCSI software or hardware adapter in a host.
Initiator group (igroup)	An igroup provides LUN-masking capability to the storage system. An igroup contains one or more LUNs with their defined LUN IDs and the host initiators that are allowed to access those LUNs.
LUN	A LUN is a block-based storage object provisioned within a FlexVol volume. LUNs are presented to an ESX or ESXi host through the iSCSI or FC protocol. A LUN can be formatted with VMFS so that it can be used as a VMware datastore to contain virtual machines, or it can be configured as an RDM device and formatted with the appropriate guest OS file system.
NFS export	An NFS export is a FlexVol volume or individual qtree that has been shared for use as a NAS datastore by the ESX or ESXi hosts.
Qtree	A qtree is a subunit of a FlexVol volume. From a client perspective (CIFS or NFS), a qtree appears as a folder or directory. Qtrees can be used to organize data or to apply different security styles.
SnapMirror	Replication software that performs automated file system replication of a volume onto a separate disk or storage system for data protection and disaster recovery.

SnapMirror runs as part of the Data ONTAP operating system on the storage controller. SnapMirror replication can be configured at the FlexVol level or at the qtree level. Volume SnapMirror can run in synchronous or asynchronous mode; qtree SnapMirror can run in asynchronous mode only.

SnapMirror operations

- SnapMirror initialize: The initial baseline transfer of data from a source volume or qtree to a destination.
- SnapMirror update: A manual update of the SnapMirror destination.
- SnapMirror quiesce or break: SnapMirror quiesce or break is performed on the SnapMirror destination to convert the destination to a writable volume or qtree.
- SnapMirror resync: After a SnapMirror destination is made writable, the SnapMirror resynchronization operation reestablishes the SnapMirror relationship. The operation also resynchronizes the contents of the source and destination volumes or qtrees, without repeating the initial transfer. SnapMirror resync is also used to reverse the direction of a SnapMirror relationship.

SnapMirror relationship

Pairing of a source FlexVol volume to a destination FlexVol volume or the pairing of a source qtree to a destination qtree.

Snapshot copy A point-in-time, read-only image of the storage system volume.

Copyright information

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at *www.ibm.com/legal/copytrade.shtml*.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index