

NetApp® Disaster Recovery Adapter 1.4.3

For use with VMware® vCenter Site Recovery Manager

Installation and Administration Guide

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>
Part number: 215-05229_A0
March 2010

Contents

Copyright information	7
Trademark information	9
About this guide	11
Audience	11
Terminology	11
Where to enter commands	12
Keyboard and formatting conventions	13
Special messages	14
How to send your comments	14
What Disaster Recovery Adapter is	15
What is new in this release	15
What Site Recovery Manager does	15
What Site Recovery Manager supports using Disaster Recovery Adapter	16
SnapMirror modes supported by Disaster Recovery Adapter	16
Components of a disaster recovery environment	16
What TestFailover-Start does	18
What TestFailover-Stop does	18
What Failover does	19
Disaster Recovery Adapter installation	21
System requirements for using Disaster Recovery Adapter	21
Limitations of Disaster Recovery Adapter	22
Setting up the storage system before running Disaster Recovery Adapter in a NAS environment	23
Setting up the storage system before running Disaster Recovery Adapter in a SAN environment	24
Downloading Disaster Recovery Adapter	25
Installing Disaster Recovery Adapter	25
Recommendations when using Disaster Recovery Adapter	26
Configuring the primary ESX server to access LUNs	26
Configuring Disaster Recovery Adapter	27
Disaster Recovery Adapter configuration file	28
Failover testing	31

Failback management by using Site Recovery Manager	33
Scenario: Running failback manually	33
Guidelines for using the Data ONTAP API over SecureHTTP	37
Setting up the disaster recovery scripts to use the Data ONTAP API over SecureHTTP	37
Disaster Recovery Adapter requirements for vFiler support	39
Troubleshooting error messages by using the log files	41
Error messages when executing the command.pl script	41
Missing the OutputFile parameter in XML input	41
Missing the LogLevel parameter in XML input	42
Command is not specified in XML input	42
Error while executing script	42
Error messages when executing the common.pm script	42
Could not get the host name	43
Unable to set HTTPS transport	43
Error messages when executing the DiscoverArrays.pl script	43
Creating server element failed	43
Collecting storage IP address failed	44
Invalid IP address present in the NFS access address list	44
Error messages when executing the DiscoverLuns.pl script	44
Could not find any exported NFS	44
Error while discovering NFS exports	45
Skipping path nfs_path as it has multiple peers	45
Could not discover any replicated exports	45
Error messages when executing the failover.pl script	45
Could not find any IP address for NFS	46
Skipping export nfs_path as input does not specify any valid VMkernel IP to be exported to	46
Skipping export nfs_path as there are no VMkernel IP to export	46
Skipping export nfs_path as this is not a valid snapmirror relation	47
Export nfs_path does not exist	47
Could not find base Snapshot copy for relation_path; Skipping nfs_path	47
Test volume clone creation for testClonevolume failed	47
Exporting path nfs_path failed	48
Onlining export nfs_path failed	48
Skipping export nfs_path as its not failed over	48

Error while exporting VMkernel and/or additional IP to the export
 nfs_path. Skipping export nfs_path 48
Destroying test Volume clone_volume failed 49
Index 51

Copyright information

Copyright © 1994–2010 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp; the NetApp logo; the Network Appliance logo; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

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 <http://www.ibm.com/legal/copytrade.shtml>.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. 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.

About this guide

You can use your product more effectively when you understand this document's intended audience and the conventions that this document uses to present information.

This guide describes how NetApp Disaster Recovery Adapter enables interaction between VMware vCenter Site Recovery Manager and the storage controller. It also provides details about the adapter installation, and system configuration and requirements. The document includes troubleshooting information.

Next topics

[Audience](#) on page 11

[Terminology](#) on page 11

[Where to enter commands](#) on page 12

[Keyboard and formatting conventions](#) on page 13

[Special messages](#) on page 14

[How to send your comments](#) on page 14

Audience

This document is written with certain assumptions about your technical knowledge and experience.

This guide is for system administrators using the array scripts with VMware vCenter Site Recovery Manager. You must be familiar with ESX servers, VMware vCenter Site Recovery Manager for Windows, and storage systems running Data ONTAP software. The guide assumes that you are familiar with how to configure these systems and how the NFS, CIFS, HTTP, and SAN, and NAS protocols (such as iSCSI, NFS, and FC) are used for file sharing or transfers. The guide does not cover basic system or network administration topics; they describe the various scripts and how the scripts are used during failover along with storage systems running Data ONTAP.

Terminology

To understand the concepts in this document, you might need to know how certain terms are used.

Storage terms

LUN (Logical Unit Number)

Refers to a logical unit of storage identified by a number.

storage controller

Refers to the component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also

sometimes called *controllers*, *storage appliances*, *appliances*, *storage engines*, *heads*, *CPU modules*, or *controller modules*.

storage system	Refers to the hardware device running Data ONTAP that receives data from and sends data to native disk shelves, third-party storage, or both. Storage systems that run Data ONTAP are sometimes referred to as <i>filers</i> , <i>appliances</i> , <i>storage appliances</i> , <i>V-Series systems</i> , or <i>systems</i> .
third-party storage	Refers to back-end storage arrays, such as IBM, Hitachi Data Systems, and HP, that provide storage for storage systems running Data ONTAP.

Cluster and high-availability terms

active/active configuration	In the Data ONTAP 7.2 and 7.3 release families, refers to a pair of storage systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning. Also sometimes referred to as <i>active/active pairs</i> . In the Data ONTAP 7.1 release family and earlier releases, this functionality is referred to as a <i>cluster</i> .
cluster	In the Data ONTAP 7.1 release family and earlier releases, refers to a pair of storage systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i> .

Where to enter commands

You can use your product more effectively when you understand how this document uses command conventions to present information.

You can perform common administrator tasks in one or more of the following ways:

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the graphical user interface.
- You can enter Windows, ESX, HP-UX, AIX, Linux, and Solaris commands at the applicable client console.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the client graphical user interface.
Your product documentation provides details about how to use the graphical user interface.
- You can enter commands either at the switch console or from any client that can obtain access to the switch using a Telnet session.

In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

Keyboard and formatting conventions

You can use your product more effectively when you understand how this document uses keyboard and formatting conventions to present information.

Keyboard conventions

Convention	What it means
The NOW site	Refers to <i>NetApp On the Web</i> at http://now.netapp.com/ .
<i>Enter, enter</i>	- Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. - Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
type	Used to mean pressing one or more keys on the keyboard.

Formatting conventions

Convention	What it means
<i>Italic font</i>	- Words or characters that require special attention. - Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters "arp -d" followed by the actual name of the host. - Book titles in cross-references.
Monospaced font	- Command names, option names, keywords, and daemon names. - Information displayed on the system console or other computer monitors. - Contents of files. - File, path, and directory names.

Convention	What it means
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Special messages

This document might contain the following types of messages to alert you to conditions that you need to be aware of.

Note: A note contains important information that helps you install or operate the system efficiently.

Attention: An attention notice contains instructions that you must follow to avoid a system crash, loss of data, or damage to the equipment.

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 e-mail to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the name of your product and the applicable operating system. For example, *FAS6070—Data ONTAP 7.3*, or *Host Utilities—Solaris*, or *Operations Manager 3.8—Windows*.

What Disaster Recovery Adapter is

NetApp Disaster Recovery Adapter (disaster recovery scripts) 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 replicated datastores, and to fail over or test failover the replicated datastores on which the virtual machines reside.

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

Next topics

[What is new in this release](#) on page 15

[What Site Recovery Manager does](#) on page 15

[What Site Recovery Manager supports using Disaster Recovery Adapter](#) on page 16

[SnapMirror modes supported by Disaster Recovery Adapter](#) on page 16

[Components of a disaster recovery environment](#) on page 16

What is new in this release

Disaster Recovery Adapter 1.4.3 is a unified adapter that you can use in either a NAS or a SAN environment.

Starting with Disaster Recovery Adapter 1.4.3, the following features are supported:

- Support for vFiler units as a storage array
- Support for non-quieted Snapshot copy recovery by using the new `use_smvi_snapshot` option.

Related concepts

[Disaster Recovery Adapter requirements for vFiler support](#) on page 39

[Disaster Recovery Adapter configuration file](#) on page 28

What Site Recovery Manager does

Site Recovery Manager uses Disaster Recovery Adapter supported by NetApp 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 leverages VMware Infrastructure to build, automate, and test disaster recovery plans for the data center. Site Recovery Manager

integrates with VMware Infrastructure and VMware vCenter Server (formerly VMware VirtualCenter) and provides a seamless view of your environment that includes both virtual infrastructure and disaster recovery management.

Site Recovery Manager uses the following disaster recovery scripts supported by NetApp:

- DiscoverArrays.pl—This script discovers the array name, array ID, storage-port information, and so on.
- DiscoverLuns.pl—This script discovers all the datastores that are replicated on the primary storage array.
- failover.pl—This script performs test-failover-start, test-failover-stop, or failover.

Related concepts

[*Components of a disaster recovery environment*](#) on page 16

What Site Recovery Manager supports using Disaster Recovery Adapter

Using Disaster Recovery Adapter, Site Recovery Manager supports SAN storage environments such as VMFS (iSCSI and FC) and RDM (iSCSI and FC), and NAS storage environments such as NFS and active/active configuration.

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

SnapMirror modes supported by Disaster Recovery Adapter


Disaster Recovery Adapter supports asynchronous qtree SnapMirror, asynchronous volume SnapMirror, and synchronous SnapMirror.

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

Components of a disaster recovery environment

A disaster recovery environment consists of the storage array, ESX servers, virtual machines, Site Recovery Manager, vCenter Server, and Disaster Recovery Adapter.

The following illustration shows the components of a disaster recovery environment.



The adapter provides array-specific support for Site Recovery Manager:

- Array discovery (DiscoverArrays.pl)

DiscoverArrays.pl performs the storage array discovery operation. This operation discovers the storage arrays that the administrator has selected as part of the Site Recovery Manager recovery policy during a disaster. The operation discovers the array ID and the port information about the storage system.

For NetApp storage in a NAS environment, the array ID is the host name of the storage system and the port is the NFS IP address.

For NetApp storage in a SAN environment, the array ID is the host name of the storage system and the port is the FC WWNN or the iSCSI IQN.

This script also discovers the model and product type of the storage array.

- Replicated device discovery (DiscoverLuns.pl)

In a NAS environment, DiscoverLuns.pl discovers replicated and exported datastores on the primary storage array. To discover an exported volume or qtree, you should export the volume or qtree to an IP address and replicate the volume or qtree by using the SnapMirror technology. The operation discovers the datastores that are exported and then checks the datastores that are replicated. Finally, the operation reports the replicated datastores along with their replication destination storage array and destination export path.

In a SAN environment, DiscoverLuns discovers the LUNs that are replicated and mapped to an initiator group (igroup) of ostype vmware. To discover a LUN, you should map the LUN to an igroup of ostype vmware and replicate the LUN by using the SnapMirror technology. The operation discovers the datastores that are mapped and then checks the datastores that are

replicated. Finally, the operation reports the replicated datastores along with their replication destination storage array and destination LUN path.

- Test failover and actual failover (failover.pl)

The failover operation performs one of the following three activities, depending on the Site Recovery Manager input:

- TestFailover-Start
- TestFailover-Stop
- Failover

Note: These activities are highly critical because they modify the customer environment.

Test failover creates a FlexClone volume from the replicated base Snapshot copy at the destination volume for testing.

Next topics

[*What TestFailover-Start does*](#) on page 18

[*What TestFailover-Stop does*](#) on page 18

[*What Failover does*](#) on page 19

What TestFailover-Start does

When a replicated export path in a NAS environment or a replicated LUN path in a SAN environment is provided in the secondary storage array, the TestFailover-Start operation tests the actual failover to verify that the failover works.

The operation first confirms whether the export path or the LUN path from the Site Recovery Manager input is mirrored. If the path is a SnapMirror destination, then the operation finds the base Snapshot copy of the volume and then creates a clone of the volume by using FlexClone technology.

In a NAS environment, the operation exports the path by adding the vmkernel IP addresses to the exports file. Now, a clone of the mirrored datastore is ready for discovery. The operation displays the cloned exported path to Site Recovery Manager.

In a SAN environment, the operation then maps the LUNs of the FlexClone volume to the initiator groups of ostype vmware. The operation displays the cloned LUN path to Site Recovery Manager.

Note: This operation does not affect the ongoing SnapMirror replication; a clone of the destination is used to verify the disaster recovery process. Therefore, you can perform a nondisruptive test to check disaster recovery. You can create a cloned volume with no space guarantee because the cloned volume is not used during actual disaster recovery and is destroyed.

What TestFailover-Stop does

The TestFailover-Stop operation cleans the environment after a TestFailover-Start operation. This operation takes the cloned volume offline and then destroys the cloned volume. This step marks the end of the nondisruptive test to check disaster recovery.

What Failover does

The Failover operation aborts any ongoing transfers to the SnapMirror destination. The operation then quiesces the destination and breaks the SnapMirror relationship. When the SnapMirror relationship is broken, the destination volume is ready for read/write operations.

The Failover operation is the actual failover at the time of disaster recovery. In a NAS environment, the operation exports the input paths by adding the vmkernel IP address to the exports file. In a SAN environment, the operation maps the input LUN paths to the initiator groups of ostype vmware.

The operation then brings the datastore online and displays the failed-over datastores to Site Recovery Manager for discovery.

Note: The Site Recovery Manager input tells Disaster Recovery Adapter the paths that need to be exported using specific vmkernel IP addresses in a NAS environment. Alternatively, in a SAN environment, it provides the LUN paths that need to be mapped using specific initiators. Disaster Recovery Adapter has to follow the input specifications.

Disaster Recovery Adapter installation

You should ensure that you meet the system requirements to install Disaster Recovery Adapter.

Next topics

[*System requirements for using Disaster Recovery Adapter*](#) on page 21

[*Limitations of Disaster Recovery Adapter*](#) on page 22

[*Setting up the storage system before running Disaster Recovery Adapter in a NAS environment*](#) on page 23

[*Setting up the storage system before running Disaster Recovery Adapter in a SAN environment*](#) on page 24

[*Downloading Disaster Recovery Adapter*](#) on page 25

[*Installing Disaster Recovery Adapter*](#) on page 25

[*Recommendations when using Disaster Recovery Adapter*](#) on page 26

[*Configuring the primary ESX server to access LUNs*](#) on page 26

[*Configuring Disaster Recovery Adapter*](#) on page 27

[*Disaster Recovery Adapter configuration file*](#) on page 28

System requirements for using Disaster Recovery Adapter

Before you run Disaster Recovery Adapter, ensure that you meet the system requirements.

The system requirements include the following:

- Data ONTAP 7.2.4 or later is installed on the following storage systems:
 - FAS200 series
 - FAS20xx
 - FAS30xx
 - FAS31xx
 - FAS60xx
 - V30xx
 - V31xx
 - V60xx

For more information about supported platforms, see the supportability matrix on the NOW site.

Note: In a SAN environment, if you are upgrading to Disaster Recovery Adapter 1.4.3, you must ensure that your storage system is running Data ONTAP 7.2.4 or later.

- Support for the NearStore option.
- FlexClone license to test failover.

- SnapMirror license is enabled on the storage system running Data ONTAP.
- In a NAS environment, NFS services are running.
- In a SAN environment, iSCSI, FC, or both services are running.
- Site Recovery Manager is installed on Windows.
In a SAN environment, Site Recovery Manager supports VMFS (iSCSI and FC) and RDM (iSCSI and FC) storage environments.
In a NAS environment, Site Recovery Manager supports NFS storage environments.
- Perl 5.8 is installed in the system in which Site Recovery Manager is installed.

Note: If the Site Recovery Manager installation does not include Perl 5.8, you can install Perl separately with the default configuration. For more information about downloading Perl, see the Perl download page.

For more information about installing SSL on Perl, see the knowledgebase article on the NOW site.

- If you want to use secureHTTP, you must set up SecureAdmin and ensure that SSL or SSH is running.

By default, the Perl script calls the Data ONTAP API over HTTP.

Note: SecureAdmin does not require a license.

- To keep the corresponding port open, the firewall (if any) between Data ONTAP and Site Recovery Manager must be configured because SOAP, HTTP, or SSL is used.

Related information

[Perl script download page](#)

<http://now.netapp.com>

Limitations of Disaster Recovery Adapter

When you use Disaster Recovery Adapter (disaster recovery scripts) for Site Recovery Manager, exports with multiple replication destinations cannot be failed over. Also, SyncMirror replication is not supported.

- Exports or LUNs that have multiple replication destinations cannot be failed over.
- The TestFailover operation might fail if there is a volume with the same name that it tries to create. This step might destroy the existing volume during TestFailover cleanup.
- The TestFailover operation is not supported on traditional volumes because FlexClone technology is not available for traditional volumes.
- In NAS environment, exports from the storage array must have values in the `rw=` field of export permissions.
The disaster recovery scripts do not discover exports that are exported as read-only by using `ro` or exports that are exported as read/write to all hosts from the `rw` field.
- SyncMirror replication between Site Recovery Manager sites (as in MetroCluster) is not supported.

Note: A MetroCluster system might be either the source or destination of a SnapMirror relationship; however, Site Recovery Manager does not support management of SyncMirror failover between MetroCluster nodes.

For specific information about recommended configurations and best practices when using different replication modes, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Setting up the storage system before running Disaster Recovery Adapter in a NAS environment

You need to set up the system before running Disaster Recovery 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 Disaster Recovery Adapter for Site Recovery Manager.
3. Ensure that the NFS exports on the primary storage controller are exported to the ESX hosts by using the `rw=` field 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 will not be discovered with the default export setting of `rw` to all hosts.

4. Ensure that the ESX hosts at the protected site have VMkernel ports that are capable of accessing the IP addresses used to serve NFS exports from the primary storage controller and that the datastores are mounted.
5. Ensure that the datastores at the protected site contain virtual machines that are registered with vCenter Server.
6. Configure the primary ESX server to access the exports in the primary storage system.
7. Replicate the volume or qtree containing the exports to the secondary storage system.
8. On the primary storage system, enter the `snapmirror status` command to confirm that the concerned export is involved in only one replication relationship and that the relationship status displays `Source`.
9. On the secondary storage system, enter the `snapmirror status` command to confirm that the relationship state is `SnapMirrored`.

For information about managing a SnapMirror relationship, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

10. Ensure that the ESX hosts at the secondary storage system have a VMkernel port that is capable of accessing the IP addresses used to serve NFS exports from the secondary storage controller.

Related tasks

[Installing Disaster Recovery Adapter](#) on page 25

Setting up the storage system before running Disaster Recovery Adapter in a SAN environment

You need to set up the system before running Disaster Recovery 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 Disaster Recovery Adapter for Site Recovery Manager.
3. Configure the primary ESX server to access the LUNs in the primary storage system.
4. Replicate the volume or qtrees containing the LUNs to the secondary storage system.
5. On the primary storage system, enter the `snapmirror status` command to confirm that the concerned LUN is involved in only one replication relationship and that the relationship status displays Source.
6. On the secondary storage system, enter the `snapmirror status` command to confirm that the relationship state is SnapMirrored.

For information about managing a SnapMirror relationship, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

7. Create igroups, either iSCSI, FC, or both, of ostype *vmware* on the secondary storage system.
8. Add secondary ESX server initiators to this igroup.
9. Enter the `igroups show` command to confirm the addition of the secondary ESX server initiators to the igroup.

You can allow more than one initiator to see the LUNs that have failed over. Site Recovery Manager informs the disaster recovery script about all initiators that should see these LUNs.

Note: Ensure that you add the initiators to igroups of ostype *vmware* on the secondary storage system.

Downloading Disaster Recovery Adapter

You can download Disaster Recovery Adapter (disaster recovery scripts) from the NOW site.

Steps

1. Log in as administrator and download the applicable file directly to your target system. Alternatively, download it to a PC host and copy it to the target system.
2. After downloading or copying the .zip file to your target system, unzip the file.

Related tasks

[Installing Disaster Recovery Adapter](#) on page 25

Related information

<http://now.netapp.com>

Installing Disaster Recovery Adapter

You can install Disaster Recovery Adapter from the location where you downloaded the scripts.

About this task

Note: In a NAS environment, if you are upgrading from an earlier version of Disaster Recovery Adapter, you must reconfigure the protection groups and recovery plans on Site Recovery Manager after installing the adapter.

For more information about reconfiguring protection groups and recovery plans on Site Recovery Manager, see the Site Recovery Manager documentation.

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.

Note: The version of Disaster Recovery Adapter is available at the following registry location: HKEY_LOCAL_MACHINE\SOFTWARE\NetApp Inc\NetApp Disaster Recovery Adapter\Version.

3. Click **Finish**.

Recommendations when using Disaster Recovery Adapter

When you use Disaster Recovery Adapter with Site Recovery Manager, you must ensure that the secondary storage system does not contain volumes with the same name as that of the volume that TestFailover creates.

- In a SAN environment, ensure that the primary and secondary ESX servers can discover LUNs on storage systems by using iSCSI, FC, or both.
- In a NAS environment, ensure that the primary and secondary ESX servers are connected to exports on the primary and secondary storage systems.
- On the secondary storage system, ensure that there are no volumes with names that match the name of the volume that TestFailover creates.
- When removing SnapMirror relationships to avoid multiple replication destinations, if you delete the SnapMirror relationship on the secondary storage system, you should also delete the relationship on the primary storage system.
- When adding active/active configuration systems to the Site Recovery Manager configuration, you should add each NetApp controller as an array in Array Managers.
- In a NAS environment, while adding storage systems to Array Managers, you should use the NFS IP Addresses field.

In this field, enter the IP address of the storage controller that is used to serve NFS from that controller. If the controller serves NFS to the ESX servers on multiple IP addresses, you should add multiple array NFS IP addresses in the field, separated by commas.

Configuring the primary ESX server to access LUNs

You can configure the primary ESX server to access the LUNs in the primary storage system.

Steps

1. Create igroups, either iSCSI, FC, or both, of ostype *vmware* on the primary storage system.
2. Add the primary ESX server initiator to the igroup.

Configuring Disaster Recovery Adapter

You can configure Disaster Recovery Adapter by using the **Configure Array Managers** wizard of Site Recover Manager.

Steps

1. In Site Recovery Manager, click the **Configure** link next to Array Managers in the Site Recovery Summary tab.

2. To add a storage array, click **Add**.

The first array that is configured is Protection Site Array.

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

- a. **Display Name**—Enter a name to identify the storage controller.
- b. **Manager Type**—Select Data ONTAP storage system.
- c. **IP Address**—Enter the management IP address of the storage controller.

Example

10.10.20.54.

- d. **NFS IP Addresses**—Enter the IP address of the controller that is used to serve NFS to the ESX hosts at the site where this controller is located.

Example

192.168.2.100,192.168.3.100

Note: This field is applicable only in NAS environments.

The adapter requires that you enter the IP addresses that are used by the storage controller to serve NFS to the ESX hosts. You can enter multiple IP addresses, separated by commas. This configuration allows for 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 NFS exports by using the IP addresses entered for managing the storage controller.

- e. **Username**—Enter the user name of an account with privileges to this storage controller.
- f. **Password**—Enter the password of the user account.

Note: If the storage system that is being added is in an active/active configuration, each controller must be added separately as an array.

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

6. Click **OK**.
7. Follow the prompts of the **Configure Array Managers** wizard to add Recovery Site Arrays.

Disaster Recovery Adapter configuration file

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

The following are some of the options that can be configured in the adapter configuration file:

- `volume_guarantee` (none/file/volume; default value is none)
The option controls the value of the volume guarantee on all volumes created for disaster recovery testing.
- `use_smvi_snapshot` (on/off; default value is off)
By default, for Site Recovery Manager disaster recovery and disaster recovery tests, the adapter recovers the volumes to the last replication point transferred by the SnapMirror technology. When the `use_smvi_snapshot` option is set to on, the adapter recovers the volumes to the last Snapshot copy that was created by SMVI. To support this feature, the option to create a VMware consistency Snapshot copy must be unselected in the SMVI. This option is useful if custom scripts are used to quiesce applications inside the virtual machine and the point in time of the recovered image must be the one that was created by SMVI.

Option specific to a SAN environment

- `lun_space_reservation` (yes/no; default value is no)
This is a provisioning management option that checks whether the LUNs recovered for disaster recovery testing are set with space reservations enabled or disabled.

Options specific to a NAS environment

- `root` (yes/no; default value is yes)
By default, the adapter exports all volumes that have root access to the ESX hosts. If your environment does not use the `root` field on exports, you can set this option to no, preventing this access. This 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` (yes/no; default value is no)
Specifies whether the adapter exports all volumes with the `nosuid` property.
- `anon` (*anon uid*)
By default, the adapter does not export volumes that are using the `anon` flag. To force volumes to be exported with an anonymous security identifier, you should enter the desired `anon` user ID.

For example: anon=0

- `sec` (default value is `sys`)
Allows you to customize the NFS export security type.

By default, during failover or TestFailover, the adapter exports volumes only to those IP addresses directed by Site Recovery Manager. The following three options of `ontap_config.txt` allow some additional customization of export hosts:

- `read-write-hosts`
A comma-separated list of additional hosts to be added to the `rw` field of the NFS exports.
- `read-only-hosts`
A comma-separated list of additional hosts to be added to the `ro` field of the NFS exports.
- `root-hosts`
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. This is useful if you need to ensure that an additional host is granted root access for the purpose of virtual disk file management or backup.

Failover testing

You need to test the failover operation to ensure that it works properly when you encounter a problem.

When testing failover, the TestFailover operation creates a cloned volume on the secondary storage system with the prefix `testfailoverClone_nss_v10745371_`, followed by the parent volume name. The test cloned volume that is created for the *secondary_volume* is named *testfailoverClone_nss_v10745371_secondary_volume*.

After testing failover, the TestFailover cleanup script tries to destroy the *testfailoverClone_nss_v10745371_secondary_volume* volume. Ensure that you do not have user-created volumes with names prefixed with *testfailoverClone_nss_v10745371_*. You can verify the names of existing volumes in the secondary storage system by running the `vol status` command.

Failback management by using Site Recovery Manager

You need to manage failback manually because Site Recovery Manager does not have an automated push-button mechanism. Failback management is also manual because the protected site might have completely different hardware and network configurations after a disaster occurs.

To achieve failback, you should perform a series of manual steps that you should run in a reverse direction in the failover workflow.

Scenario: Running failback manually

In a failback scenario, you can use Site Recovery Manager as a failback tool to run failback manually. You should run failback to restore the original primary protected site to its pre-failover state. You can manage failback like any planned server migration.

About this task

In the following failback scenario, Site Recovery Manager is the failback tool in the protected site after running the recovery plan at the recovery site in recovery mode for recovered virtual machines.

Name	Description
Site A	original protected site
Site B	original recovery site
Storage systemA	original protected storage system
Storage systemB	original recovery storage system

Assume that there is an original protected site (Site A) and an original recovery site (Site B). After failover occurs, the recovery site becomes the acting protected site (Site B). By using the failback procedure, the original protected site becomes the acting recovery or failback site (Site A).

Steps

1. Ensure that the protected virtual machines at Site A are no longer accessed by any users or applications.
2. Shut down all the protected virtual machines that were part of the recovery plan at Site B.
3. Perform one of the following steps:

If the environment is...	Then...
SAN	Perform a cleanup of the failover LUN at Site B that contained virtual machine configuration files created during the protection group setup at Site A.
NAS	Perform a cleanup of the failover export at Site B that contained virtual machine configuration files created during the protection group setup at Site A.

4. Delete protection groups from Site A.
5. Perform a remove from inventory operation for all the protected virtual machines from Site A that were part of the recovery plan at Site B.
6. Perform one of the following steps:

If...	Then...
You want to recover storage systemA's volume	Enter the following command from storage systemA to resynchronize the SnapMirror relationship: snapmirror resync -S storage systemB:volB storage systemA:volA
You do not want to recover storage systemA's volume	Enter the following command from storage systemA to initialize the SnapMirror relationship: snapmirror initialize -S storage systemB:volB storage systemA:volA

7. In a SAN environment, ensure that there is an igroup of ostype *vmware* on the original protected storage system. If not, create an igroup of ostype *vmware*.
8. Verify that the two Site Recovery Manager servers are paired with each other; if they are not, pair them with each other.
9. At Site B, configure the array manager and inventory mapping.
10. At Site B, create protection groups for virtual machines to be failed back. In a SAN environment, select a spare LUN from Site A for failback inventory mapping. In a NAS environment, select a spare export from Site A for failback inventory mapping.
11. At Site A, create a recovery plan for the protection groups that are created on Site B.
12. From Site A, test the recovery plan.
13. If the test is successful, then run the recovery plan.
14. Shut down all the protected virtual machines at Site A that were failed back during the failback recovery plan.
15. Perform one of the following steps:

If the environment is...	Then...
SAN	Perform a cleanup of the failback LUN at Site A that contained the virtual machine configuration files created during the protection group setup at Site B.
NAS	Perform a cleanup of the failback export at Site A that contained the virtual machine configuration files created during the protection group setup at Site B.

Guidelines for using the Data ONTAP API over SecureHTTP

The Data ONTAP API can use either HTTP or SecureHTTP. By default, the disaster recovery scripts use the Data ONTAP API over HTTP. However, the editable configuration SSL parameter enables you to use the Data ONTAP API over SecureHTTP.

Setting up the disaster recovery scripts to use the Data ONTAP API over SecureHTTP

You can set up Disaster Recovery Adapter to use the Data ONTAP API over SecureHTTP.

Steps

1. In the primary and secondary storage systems, set up SecureAdmin and enable the SSL and SSH services.
2. Enter the `secureadmin status` command to confirm that the SSL and SSH services are running.
3. Add the following line to the `ontap_config.txt` file: `ssl=on`
4. In the Perl environment, install the `Net::SSLeay` module.

HTTPS support uses `Net::SSLeay.pm`. `Net::SSLeay` is not part of Perl 5.8 that is shipped with Site Recovery Manager. You need to compile `Net::SSLeay.pm` from the source `Net_SSLeay.pm-1.xx.tar.gz` from the Comprehensive Perl Archive Network. For instructions about how to compile and install `Net::SSLeay.pm` from the source, see the Readme file available in `Net_SSLeay.pm-1.xx.tar.gz`.

Related information

<http://www.cpan.org/>
KB 54012

Disaster Recovery Adapter requirements for vFiler support

For Disaster Recovery Adapter to support vFiler units, you must add each vFiler unit to the adapter as a separate array and your storage system must be running Data ONTAP 7.2.3 or later.

The following are the requirements for vFiler support by the adapter:

- The storage system must be running Data ONTAP 7.3.2 or later.
- You must add each vFiler unit to the adapter as a separate array.
- Both the source and destination vFiler units must be online.
- SnapMirror relationships must be defined within the destination vFiler unit and not on vfiler0.
- Only iSCSI or NFS connected datastores are supported in vFiler units.
- In a SAN environment, the recovery site ESX hosts must have established iSCSI sessions to the recovery vFiler units before you execute the Disaster Recovery Adapter test or failover. The adapter does not start the iSCSI service in the vFiler unit.
- In a SAN environment, the recovery site vFiler units must have an igroup of type *vmware* before you execute the Disaster Recovery Adapter test or failover. The adapter does not create the *vmware* igroup in the vFiler unit.
- A single virtual machine must not have data on multiple vFiler units.
- Before performing a test failover operation on a vFiler unit, you must set the `vfiler.vol_clone_zapi_allow` option of the base vFiler unit to on.

Troubleshooting error messages by using the log files

You can use the log files to troubleshoot problems that you might encounter when using Disaster Recovery Adapter. You can find the log files at C:\Documents and Setting\All Users\Application Data\VMware\VMware Site Recovery Manager\Logs\. There might be multiple log files with names such as vmware-dr-log number.log.

The Site Recovery Manager user interface does not display the error log messages.

In case of any adapter failure or unexpected behavior, you should set the log level to `trivia` in Site Recovery Manager and run the failed operation again. For details about running the failover operation, see the Site Recovery Manager documentation. Then, you should examine the log messages.

Note: The solutions provided assume that you are running the disaster recovery scripts in a failover environment.

Next topics

[Error messages when executing the `command.pl` script](#) on page 41

[Error messages when executing the `common.pm` script](#) on page 42

[Error messages when executing the `DiscoverArrays.pl` script](#) on page 43

[Error messages when executing the `DiscoverLuns.pl` script](#) on page 44

[Error messages when executing the `failover.pl` script](#) on page 45

Error messages when executing the `command.pl` script

The `command.pl` script invokes `discoverArrays.pl`, `discoverLuns.pl`, or `failover.pl` based on the input.

Next topics

[Missing the `OutputFile` parameter in XML input](#) on page 41

[Missing the `LogLevel` parameter in XML input](#) on page 42

[Command is not specified in XML input](#) on page 42

[Error while executing script](#) on page 42

Missing the `OutputFile` parameter in XML input

Message	Missing the <code>OutputFile</code> parameter in XML input
Log level	Quiet

Cause	The output file parameter is missing in the XML input file from Site Recovery Manager.
Corrective action	Check the configuration settings in Site Recovery Manager.

Missing the LogLevel parameter in XML input

Message	Missing the LogLevel parameter in XML input
Log level	Quiet
Cause	The LogLevel parameter is missing in the XML input file from Site Recovery Manager.
Corrective action	Check the configuration settings in Site Recovery Manager.

Command is not specified in XML input

Message	Command is not specified in XML input
Log level	Quiet
Cause	The input command from Site Recovery Manager is not part of discoverArrays, discoverLuns, failover, or testFailover operations.
Corrective action	Check the configuration settings in Site Recovery Manager.

Error while executing script

Message	Error while executing script
Log level	Quiet
Cause	The error occurs when the script is run in a Perl environment.
Corrective action	Add the path to perl.exe to the system path and install the scripts in the correct directory.

Error messages when executing the common.pm script

The common.pm script contains common routines that are used by other scripts.

Next topics

Could not get the host name on page 43

Unable to set HTTPS transport on page 43

Could not get the host name

Message	Could not get the host name
Log level	Quiet
Cause	The storage system does not have a valid host name.
Corrective action	Configure a valid host name for the storage system.

Unable to set HTTPS transport

Message	Unable to set HTTPS transport
Log level	Quiet
Cause	The SSL and SSH services might not be running in the storage system.
Corrective action	Run the <code>secureadmin status</code> command. If the services are not running, configure the <code>secureadmin</code> service and start the SSH and SSL services.

Error messages when executing the DiscoverArrays.pl script

The DiscoverArrays.pl script discovers the array name, array ID, storage-port information, and so on.

Next topics

[Creating server element failed](#) on page 43

[Collecting storage IP address failed](#) on page 44

[Invalid IP address present in the NFS access address list](#) on page 44

Creating server element failed

Message	Creating server element failed
Log level	Quiet
Cause	This might occur because of the following reasons: • Low memory in the vCenter Server host. • Cannot connect to the NetApp storage system.

Collecting storage IP address failed

Message	Collecting storage IP address failed
Log level	Quiet
Cause	Either the storage controller does not have configured IP addresses or has incorrect IP address entered in the NFS IP Addresses list.
Corrective action	Verify that the storage controller has at least one configured IP address and the list of IP addresses entered in the NFS IP Addresses list is correct.

Invalid IP address present in the NFS access address list

Message	Invalid IP address present in the NFS access address list
Log level	Verbose
Cause	Either the list of IP addresses entered in the NFS IP Addresses list is not in the correct format or one or more IP addresses entered in the list do not match the configured IP addresses of the storage controller.
Corrective action	Verify that the NFS IP Addresses list is in the correct format. All the IP addresses entered in the NFS IP Addresses list should be configured and present on the storage controller.

Error messages when executing the DiscoverLuns.pl script

The DiscoverLuns discovers replicated and exported datastores on the primary storage array.

Next topics

[*Could not find any exported NFS*](#) on page 44

[*Error while discovering NFS exports*](#) on page 45

[*Skiping path nfs_path as it has multiple peers*](#) on page 45

[*Could not discover any replicated exports*](#) on page 45

Could not find any exported NFS

Message	Could not find any exported NFS
Log level	Quiet
Cause	Either there are no NFS exports on the storage controller or the NFS exports are not exported using the rw field.

Corrective action Verify that there are NFS exports on the storage controller and the NFS exports are exported using the rw field.

Error while discovering NFS exports

Message Error while discovering NFS exports

Log level Quiet

Cause The NFS service might not be running.

Corrective action Verify and ensure that the NFS service is running on the storage controller.

Skipping path nfs_path as it has multiple peers

Message Skipping path nfs_path as it has multiple peers

Log level Quiet

Cause NFS exports that have been replicated multiple times is not currently supported.

Corrective action Ensure that the concerned NFS export has only one replication peer.

Could not discover any replicated exports

Message Could not discover any replicated exports

Log level Quiet

Cause The NFS exports might not be replicated.

Corrective action Replicate the concerned NFS exports.

Error messages when executing the failover.pl script

The failover.pl script performs test-failover-start, test-failover-stop, or the actual failover depending on the Site Recovery Manager input.

Next topics

Could not find any IP address for NFS on page 46

Skipping export nfs_path as input does not specify any valid VMkernel IP to be exported to on page 46

Skipping export nfs_path as there are no VMkernel IP to export on page 46

Skipping export nfs_path as this is not a valid snapmirror relation on page 47

Export nfs_path does not exist on page 47

Could not find base Snapshot copy for relation_path; Skipping nfs_path on page 47

Test volume clone creation for testClonevolume failed on page 47

Exporting path nfs_path failed on page 48

Onlining export nfs_path failed on page 48

Skipping export nfs_path as its not failed over on page 48

Error while exporting VMkernel and/or additional IP to the export nfs_path. Skipping export nfs_path on page 48

Destroying test Volume clone_volume failed on page 49

Could not find any IP address for NFS

Message	Could not find any IP address for NFS
Log level	Quiet
Cause	The storage controller does not have configured IP addresses.
Corrective action	Verify that the storage controller has at least one configured IP address.

Skipping export nfs_path as input does not specify any valid VMkernel IP to be exported to

Message	Skipping export nfs_path as input does not specify any valid VMkernel IP to be exported to
Log level	Quiet
Cause	The XML input file from Site Recovery Manager does not specify any valid storage device on the storage controller.
Corrective action	Configure the Site Recovery Manager. Verify the Site Recovery Manager input from the Site Recovery Manager logs.

Skipping export nfs_path as there are no VMkernel IP to export

Message	Skipping export nfs_path as there are no VMkernel IP to export
Log level	Quiet
Cause	The XML input file from Site Recovery Manager does not specify a valid vmkernel IP address to export the NFS storage devices.
Corrective action	Configure the Site Recovery Manager. Verify the Site Recovery Manager input from the Site Recovery Manager logs.

Skipping export nfs_path as this is not a valid snapmirror relation

Message	Skipping export nfs_path as this is not a valid snapmirror relation
Log level	Quiet
Cause	The nfs_path is not replicated.
Corrective action	Check the snapmirror status command output to confirm whether the nfs_path is from a replication relation. Replicate the nfs_path.

Export nfs_path does not exist

Message	Export nfs_path does not exist
Log level	Quiet
Cause	The storage device nfs_path might not be present on the storage controller.
Corrective action	Verify that the nfs_path storage device is present on the storage controller.

Could not find base Snapshot copy for relation_path; Skipping nfs_path

Message	Could not find base Snapshot copy for relation_path; Skipping nfs_path
Log level	Quiet
Cause	The relation_path might not be replicated.
Corrective action	Check the snapmirror status command output to confirm whether the relation_path is from a replication relation.

Test volume clone creation for testClonevolume failed

Message	Test volume clone creation for testClonevolume failed
Log level	Quiet
Cause	This might occur because a volume with the name <i>testfailoverClone_nss_v10745371_Volumename</i> already exists. The script tries to create a cloned volume and to prefix the original volume name with <i>testfailoverClone_nss_v10745371_Volumename</i> .
Corrective action	Rename the existing volume.

Exporting path nfs_path failed

Message	Exporting path nfs_path failed
Log level	Quiet
Cause	The exporting of the nfs_path failed due to some reason.
Corrective action	Check if NFS is running and verify if the NFS exporting options are correct.

Onlining export nfs_path failed

Message	Onlining export nfs_path failed
Log level	Quiet
Cause	The onlining of the volume failed due to some reason.
Corrective action	Check if the nfs_path exists.

Skipping export nfs_path as its not failed over

Message	Skipping export nfs_path as its not failed over
Log level	Quiet
Cause	For a successful failover, the nfs_path must be online, writable, and NFS exported.
Corrective action	Run the operation with loglevel=trivia and identify the problem from the logs.

Error while exporting VMkernel and/or additional IP to the export nfs_path. Skipping export nfs_path

Message	Error while exporting VMkernel and/or additional IP to the export nfs_path. Skipping export nfs_path
Log level	Quiet
Cause	The nfs_path was not exported to the Site Recovery Manager input vmkernel IP addresses and the user-configured NFS exporting options.
Corrective action	Run the operation with loglevel=trivia and identify the problem from the logs.

Destroying test Volume clone_volume failed

Message	Destroying test Volume clone_volume failed
Log level	Quiet
Cause	The cloned volume was busy or being used by some other subsystem.
Corrective action	Retry destroying the clone_volume later.

Index

C

command not specified 42
 command.pl 41
 common.pm 42

D

Data ONTAP API 37
 Disaster Recovery Adapter
 configuration 27
 downloading 25
 Installation 25
 limitations in a NAS environment 22
 overview 15
 disaster recovery components
 DiscoverArrays.pl 16
 DiscoverLuns.pl 16
 failover.pl 16
 disaster recovery scripts
 recommendation 26
 discoverArrays.pl 43
 discoverLuns.pl 44

E

error message
 log files 41
 executing script 42

F

failover 18, 31
 failover.pl 45

H

host name 43

L

license requirements 21
 LogLevel parameter 42

LUN access
 configuring ESX server 26
 primary storage system 26

M

managing failback
 running a failback 33
 scenario:manual failback 33

O

output file parameter 41

S

server element 43
 set HTTPS 43
 setup scripts
 Secure HTTP 37
 SSL and SSH services 37
 using the Data ONTAP API 37
 Site Recovery Manager
 disaster recovery scripts
 DiscoverArrays.pl 15
 Discoverexports.pl 15
 failover.pl 15
 SnapMirror
 asynchronous qtree 16
 asynchronous volume 16
 synchronous 16
 system requirements 21
 system setup
 NAS environment 23
 SAN environment 24

T

testClonevolume fail 47
 testFailover 26
 TestFailover-Start 18
 TestFailover-Stop 18
 testing 31

