VFM (Virtual File Manager) v6.0

Make Data Highly Available

June 28, 2007

Publication Number: 215-03205_A0

Copyright © 2007 Brocade Communications Systems, Inc. All Rights Reserved.

Copyright © 2007 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this book 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 owners.

Software derived from copyrighted material of Brocade Communications Systems, Inc. is subject to the following license and disclaimer:

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

This product includes software developed by the Apache Software Foundation.

The product described in this manual may be protected by one or more U.S. 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

Brocade®, Tapestry™ StorageX®, and associated logos are registered trademarks or trademarks of Brocade Communications Systems, Inc. Other product and company names are trademarks of their respective owners.

Brocade Communications Systems, Inc, 1745 Technology Drive, San Jose, CA 95110 (www.brocade.com).

NetApp® and the Network Appliance design are registered trademarks of Network Appliance, Inc. in the United States, Canada, the European Union, and Japan. Network Appliance is a registered trademark of Network Appliance, Inc. in Monaco and a trademark of Network Appliance, Inc. in the United States and Canada. SnapMirror is a registered trademark of Network Appliance, Inc. in the United States.

VFM® and the VFM logo are registered trademarks of Network Appliance, Inc.

Preface

What the guide covers

This guide describes how you can use Network Appliance® VFM® technology to simplify file management and increase data availability while reducing storage costs. It describes how to create and manage disaster recovery policies to make your data highly available.

The guide is divided into the following sections:

Chapter 1 Introduction explains how you can use VFM to make your data highly available. It provides information you should consider when planning for disaster recovery.

Chapter 2 Client Disaster Recovery Policies explains how to create and manage client disaster recovery policies.

Chapter 3 Filer Disaster Recovery Policies explains how to create and manage filer disaster recovery policies.

Chapter 4 Server Disaster Recovery Policies explains how to create and manage server disaster recovery policies.

Appendix provides a list and description of the VFM documents.

Glossary provides definitions of terms used in this guide.

Audience

This guide is for network administrators who are familiar with Windows NT®, Windows® 2000, Windows Server® 2003, and DFS. It does not cover basic system or network administration topics.

Typographic conventions

Bold text - Items on a dialog box or page on the VFM console that you are instructed to click.

`code_text_in_italics` – Variable; replace with actual value.

Using the online help

Certain icons appear to the left of a topic title to indicate that the information provided applies only to a specific operating system, type of DFS root, or data transfer protocol.

Icon	Description
Win2K	Indicates that the information applies only to DFS roots on Windows 2000 or Windows Server 2003, Enterprise Edition.
domain root	Indicates that the information applies only to domain-based roots.
W2K3	Indicates that the information applies only to DFS roots on Windows Server 2003.

Definitions of new terms are available wherever they appear in help topics. The first time a new term appears in a help topic, it appears as a green hyperlink.

Context help contains detailed information about a dialog box or view, and might include links to related topics. Press **F1** or click **Help**.

Support

Contact NetApp technical support, 24 x 7.

Support is either through the web at http://now.netapp.com or by phone. For details, see http://www.netapp.com/support.

Support phone numbers:

888 4-NETAPP (US and Canada)

(1-888-463-8277)

(408) 822-4700

800-44-NETAPP (EMEA/Europe)

800-800-0800 (Asia/Pacific)

Notes:

You can also contact technical support by clicking Support on the Help menu.

To get your product build number and license information, click About Virtual File Manager on the Help menu.

Enter new license key

Enter a new license key for the VFM client and server components when you want to extend the VFM evaluation period or when you purchase additional licenses from VFM.

1. Contact NetApp to request a new license key.
2. On the **Help** menu, click **About Virtual File Manager**.
3. Click **Edit Serial Numbers**.

4. In the License Type list, click the appropriate item.
5. If you selected "Evaluation" as the license type, type the new activation key.
6. If you selected "Full" as the license type, select a product number from the Product list and type the new activation key.
7. Click **OK**.

Table of Contents

PREFACE ..... **I**

CHAPTER 1 INTRODUCTION **1**

DISASTER RECOVERY PLANNING CONSIDERATIONS 2
Planning considerations *2*
Resynchronizing SnapMirrors for disaster recovery failback *3*

CHAPTER 2 CLIENT DISASTER RECOVERY **5**

CREATE A CLIENT DISASTER RECOVERY POLICY 5
ADD A LINK TO A POLICY 6
To add a link with multiple targets *7*
To add a link with a single target *7*
ADD MULTIPLE LINKS TO A POLICY 7
To add DFS links contained in a logical folder or DFS root *8*
To add selected links *8*
ADD A SHARE TO A CLIENT DISASTER RECOVERY POLICY 9
SELECT SEARCH OPTIONS 9
SCHEDULE TARGET MONITORING 9
To schedule target monitoring *10*
SCHEDULE DATA REPLICATION 10
To schedule data replication *10*
SELECT A SNAPMIRROR 11
FAILOVER AND REPLICATION OPTIONS 12
Failover options *12*
Replication options *12*
SELECT REPLICATION OPTIONS 13
SELECT OTHER REPLICATION OPTIONS 16
FAILOVER 17
Set up automatic failover *17*
Manually fail over one link *17*
Manually fail over all links *18*
Manually fail back a link *19*
Manually fail back all links *20*
RUN CUSTOMIZED COMMANDS DURING FAILOVER PROCESSING 20
RUN CUSTOMIZED COMMANDS DURING REPLICATION 21
ADD ITEMS TO AN IGNORE LIST 21
REVIEW THE LIST OF NEW RESOURCES 22
UPDATE PROPERTIES OF A CLIENT DISASTER RECOVERY POLICY 23
DELETE A CLIENT DISASTER RECOVERY POLICY 23
ADD NEW SHARES TO A LINK IN A CLIENT DISASTER POLICY 23
To add existing shares to a link that is a member of a disaster recovery policy *24*
To create new shares to add to a disaster recovery policy *24*
COMPARE CONTENTS OF TARGETS IN A CLIENT DISASTER RECOVERY POLICY 25
ADD SHARED FOLDER REPLICAS TO A LINK 25
PRIORITIZE THE TARGETS OF A LINK 26
REMOVE A LINK FROM A POLICY 26
VIEW INFORMATION ABOUT REPLICATION OPERATIONS 26
START REPLICATION IMMEDIATELY 27
CANCEL A REPLICATION IN PROGRESS 27

CHAPTER 3 FILER DISASTER RECOVERY **29**

CREATE A FILER DISASTER RECOVERY POLICY 29
To create a filer disaster recovery policy *29*

ADD VOLUMES AND QTREES TO THE POLICY 31
FAIL OVER IMMEDIATELY 31
RUN CUSTOMIZED COMMANDS 32
REVIEW THE LIST OF NEW QTREES 32
To add a new qtree to the policy *32*
To add a new qtree to the ignore list of the policy *33*
UPDATE THE POLICY IGNORE LIST 33
VIEW THE STATUS OF A SNAPMIRROR 33
VIEW LINKS TARGETING DATA ON A FILER 34
VIEW INFORMATION ABOUT REPLICATION OPERATIONS 34
UPDATE PROPERTIES OF A POLICY 34
REMOVE A FILER, VOLUME, OR QTREE FROM A POLICY 34
DELETE A FILER DISASTER RECOVERY POLICY 35

CHAPTER 4 SERVER DISASTER RECOVERY 37

CREATE A SERVER DISASTER RECOVERY POLICY 37
SELECT RESOURCE MONITORING OPTIONS 38
SCHEDULE RESOURCE MONITORING 39
SCHEDULE TARGET MONITORING AND REPLICATION 39
SELECT REPLICATION OPTIONS 39
ADD LINKS TO A SERVER DISASTER RECOVERY POLICY 40
To add links to the Active Links list *40*
To repair links in the Nonactive Links list *40*
To drag a link to a server disaster recovery policy *41*
START MONITORING FOR NEW LINKS IMMEDIATELY 41
ASSIGN A LINK FOR THE POLICY TO MONITOR 41
CANCEL RESOURCE MONITORING IN PROGRESS 42
To check whether monitoring is in progress *42*
To cancel resource monitoring in progress *42*
MANAGE REPLICATION OPERATIONS 42
START FAILOVER OR FAILBACK IMMEDIATELY 42
To fail over all links manually *42*
To manually fail back all links *43*
REMOVE LINKS FROM A SERVER DISASTER RECOVERY POLICY 43
DELETE A SERVER DISASTER RECOVERY POLICY 43

APPENDIX 45

VFM GUIDES 45

GLOSSARY 47

INDEX 51

Chapter 1 Introduction

This guide provides detailed instructions about how to create and manage disaster recovery policies to make data highly available. VFM provides three types of disaster recovery policies. Before you create a policy, review "Disaster recovery planning considerations."

Select the policy that best meets your needs.

Client disaster recovery policies enable you to detect a failure and automatically fail over clients to a replicated storage device. Use this policy when data is scattered among multiple servers.

This policy provides the greatest flexibility. It enables you to control how monitoring and failover occur.

Links contained in the policy must have a least two replicated targets and might contain multiple offline targets.

One of the most significant advantages of this policy is the ability to designate a cascaded failover path. A typical scenario is one in which you designate critical data to fail over first to a local edge device then to a remote central device. This ensures high availability of the data whether there is a local server failure or an entire site outage. Failover takes place without changing login scripts or network configuration.

Filer disaster recovery policies enable you to make data on NetApp filers participating in your namespace highly available.

Set up a filer or nearline storage device as a disaster recovery backup for one or more primary filers.

Leverage SnapMirror® replication to replicate data to a standby volume or qtree. When a failure on the primary storage is detected, an automatic failover to the standby storage is initiated. VFM periodically checks for new links and qtrees on the primary filer and creates the new data configuration on the backup filer.

Server disaster recovery policies enable you to make data on servers participating in your namespace highly available. Use this policy when data on a primary server must be made highly available. When a failure of the primary server is detected, VFM automatically fails over all links in the policy to a secondary server.

This policy enables you to configure a server as a disaster recovery backup for one or more primary servers.

You can monitor for new links targeting the primary server to add to the policy and to replicate data to a secondary server.

Note: You can create namespace availability policies to make the logical structure of the namespace highly available.

Disaster recovery planning considerations

This section discusses planning considerations for client and filer disaster recovery policies and failback procedures when you are using SnapMirror replication in a disaster recovery policy.

Planning considerations

Links used in disaster recovery policies must reside in Windows 2000 or Windows Server 2003 roots. This is because disaster recovery policies rely on the ability to toggle link targets online and offline.

When you add a link to a disaster recovery policy using SnapMirror replication, the link should have only two targets. Make sure that one target is the source and the other target is the destination of the SnapMirror. See "Add a link to a disaster recovery policy" for information about configuring mirrored qtrees as targets of a client disaster recovery link.

The account used by the VFM server or monitoring agent, as well as any accounts used to run the VFM client, must have permission to make modifications to the DFS root containing the links in the disaster recovery policy. If the account lacks these permissions, VFM cannot control the online/offline state of the link targets.

If a DFS root served by a single system is used and becomes unavailable, failover cannot be performed and users can access only the primary storage target until their referral times out. Adding additional servers to host copies of the namespace provides fault tolerance. To create copies of the namespace, you create DFS root replicas in the case of domain-based roots, or you can set up namespace availability policies to synchronize the logical structure of two or more stand-alone roots. An alternate approach for creating a highly available root is to host a stand-alone root on a Windows 2000 cluster server.

If the VFM server or monitoring agent for a client or filer disaster recovery policy becomes unavailable, neither failover nor data synchronization can be performed. User access to the primary storage target is not affected.

VFM communicates with the DFS root to initiate a failover.

VFM communicates directly with the storage targets to initiate data synchronization.

The VFM client does not have to be running for disaster recovery policies to function.

Carefully plan where each VFM component (for example, client, server, monitoring agent, and replication agents) should be deployed.

When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files.

Roots containing links participating in disaster recovery policies can also be members of namespace availability policies. In this configuration, as with disaster recovery policies containing links from replicated domain-based roots, the recommendation is that you configure the disaster recovery policies for manual failover rather than for automatic failover if the links target read/write data rather than read-only data.

When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.

Resynchronizing SnapMirrors for disaster recovery failback

When a failover breaks the SnapMirror and redirects users to the alternate target, you can use the following procedures to update the repaired source and reestablish the original configuration of the SnapMirror. For detailed information, see the *Data ONTAP® System Administration Data Protection Guide*.

In the example below, LinkAB is a replicated link participating in a disaster recovery policy. LinkAB targets the source and destination of a SnapMirror. FilerA:volA, the source, is the primary target online and FilerB:volB, the destination, is the secondary offline target. On failover, the policy breaks the SnapMirror and redirects users to FilerB:volB. In VFM, LinkAB shows target FilerA:volA offline and target FilerB:volB online.

1. Make the original source into a read-only destination by performing one of the following actions:

   If the original source is recoverable, execute the `snapmirror resync` command on FilerA.

   If the original source is not recoverable, make a new volA on FilerA and execute the command `snapmirror initialize` on FilerA to initialize FilerA:volA from FilerB:volB.
2. Using VFM, manually fail back LinkAB to FilerA:volA which is now the read-only destination of the SnapMirror.
3. Transfer the latest data from FilerB:volB to FilerA:volA by executing the command `snapmirror update` on FilerA.
4. Execute the command `snapmirror break` on FilerA to make FilerA:volA, the read-only destination, writable.
5. Execute the command `snapmirror resync` on FilerB to make FilerB:volB (the original destination) the destination again. In VFM, LinkAB shows target FilerA:volA online and FilerB:volB offline.

Chapter 2 Client Disaster Recovery

To create a policy, follow the basic steps listed below:

1. Create a client disaster recovery policy.
2. Add a link to the policy.
3. Schedule status monitoring.
4. Schedule replication.

Note: The Brocade product Tapestry™ MyView® provides the option to optimize namespaces for disaster recovery integration by managing link targets. Tapestry MyView reduces the number of links in a namespace which reduces the time it takes to fail over links. You can create a client disaster recovery policy containing links in the namespace that are managed by Tapestry MyView.

Create a client disaster recovery policy

Provide a name for the policy, select automatic failover, select a replication method, and set up status monitoring and data replication. A policy can contain multiple links with replicated targets, logical folders, and DFS roots.

Note: Roots containing links participating in client disaster recovery policies can also be members of namespace availability policies.

1. Perform one of the following actions:

   On the **Tools** menu, click **New Policy**, point to **Disaster Recovery Policy**, and click **Client Recovery Policy**.

   Drag a link to the Client Recovery Policies folder.

   Right-click the **Disaster Recovery Policies** folder in the Admin View, point to **New**, and click **Client Disaster Recovery Policy**.

   Right-click the **Client Recovery Policies** folder and click **New Client Recovery Policy**.

   Right-click

   , click **New**, point to **Disaster Recovery Policy**, and click **Client Recovery Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field. The name appears under the Client Recovery Policies folder and on the Properties tab of items that are members of the policy.
3. Select a replication method. Use VFM replication when you are using heterogeneous devices.
4. Select **Automatically failover** when you want to automatically make a standby target the active target when the active target is not available.
5. Select **Failover all links** when you want to automatically switch to the standby target of all the links in the policy when any active target that is a member of the policy is not available. This option is automatically selected when the replication method is SnapMirror.
6. Select **Update associated policies on failover** when links in the client disaster recovery are part of a multiprotocol namespace or namespace availability policy, and you want to

create new automount and export files referencing the new active targets after a failover.

7. Select a link in the **Only monitor link** list if you want to monitor only the link selected. This option is not available if you selected SnapMirror replication.
8. You can select another monitoring location where monitoring of the targets of the client disaster recovery policy is performed. The default is to monitor from the VFM server. From the **Use the Monitor on** list, select the host name of another monitoring agent.

   Note: If you uninstall a monitoring agent, ensure that policies that formerly used the uninstalled agent are reassigned to a currently available agent. Right-click a client disaster recovery policy and click **Properties**. On the General tab, select a monitoring agent from the **Use the Monitor on** list.

9. Click **OK**.

Add a link to a policy

Add a link to a client disaster recovery policy.

When you add a link to a client disaster recovery policy using SnapMirror replication, the link should have only two targets. Make sure that one target is the source and the other target is the destination of the SnapMirror.

When configuring mirrored qtrees as primary and secondary targets of a client disaster recovery link, you must complete the following steps:

1. Create the SnapMirror for the qtree.
2. Quiesce and break the SnapMirror.
3. Share the mirror qtree as a CIFS share.
4. Create a link that targets both the source and destination shares of the mirror.
5. Add the link to the client disaster recovery policy.
6. Resync the SnapMirror.

It is recommended that you add only links that have shared folders as their targets to disaster recovery policies. Due to replication agent location and DFS referral considerations, it is not recommended that you add links to client disaster recovery policies that have a DFS root or another link as their target.

The account used by the VFM server or monitoring agent must have permission to make modifications to the DFS root containing the link or links in the disaster recovery policy. If the account lacks these permissions, VFM cannot control the online/offline state of the link targets.

You must "specify a secondary server for failover" when the primary and secondary targets of a link are hosted on a mix of Windows systems and filers. The reason for this is that local path naming conventions are different between Windows and filers.

The targets of a link in a disaster recovery policy are sorted in alphabetical order and set offline. You must specify which target is the active or online target and sort the targets in the order in which you want them to fail over. To change the order in which targets fail over, see "Prioritize the targets of a link."

To add a link with multiple targets

Perform one of the following actions:

Drag a link with multiple targets from the Logical View tree to the policy.

Drag a link with multiple targets from the **Properties** tab of an item in the Logical View to the policy.

Right-click a policy under the **Disaster Recovery Policies** folder and click **Add Link to policy**. Select a link with multiple targets . Click **OK**.

Click the **Disaster Recovery Policies** folder to display a **Properties** tab in the right pane. Right-click a policy and click **Add Link to policy**. Select a link with multiple targets . Click **OK**.

To add a link with a single target

When you add a link with a single target to a client disaster recovery policy, you must add another target to the link or create a copy of the existing target on another machine.

Perform one of the following actions:

Drag a single target link from the Logical View to the policy.

Drag a single target link from the **Properties** tab of an item in the Logical View to the policy.

Right-click the policy and click **Add Link to policy**. Select a single target link . Click **OK**.

Click the **Disaster Recovery Policies** folder to display a **Properties** tab in the right pane. Right-click a policy and click **Add Link to policy**. Select a single target link . Click **OK**.

Add multiple links to a policy

You can configure a client disaster recovery policy to search for DFS links in a logical folder or a DFS root that is a member of the policy. When the policy runs, it searches for links that match the selection criteria of the policy. You have the option of automatically adding the new links to the policy or viewing a list of new links and selecting which links to add to the policy.

You can add multiple links to a policy by dragging the selected DFS links to the policy.

The targets of a link in a disaster recovery policy are sorted alphabetically. The first target becomes the active or online target and all other targets of the link are placed offline. To change which target is active, see "Prioritize the targets of a link."

To add DFS links contained in a logical folder or DFS root

1. Drag a logical folder or DFS root to the policy.
2. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Properties**.

   Select a client disaster recovery policy to show a **Properties** tab in the right pane. Click Edit.
3. Click **Search Options**.
4. Select any of the following options:

   Automatically add newly found links - Search for new links and automatically adds them to the policy.

   Automatically remove links from the policy when links are removed from a logical folder or DFS root - Remove links from the policy that are removed from a logical folder or root that is a member of the policy.

   Automatically set the first link target online - Set the first target of a link to always be online. To see the targets of a link, click the link. The list of targets appears in the **Properties** tab under **Target**.

   Ignore single-target links - Do not search for links with only one target.
5. Optionally, add items to the **Ignore List** that you do not want to include in the policy. When the policy searches for new links in logical folders and roots, it "ignores" the items on the list.
6. Click **OK**. When the policy runs, it searches for links that match the search criteria. If you selected the option to **Automatically add newly found links**, the links found are added to the policy. If you did not select this option, the links are listed under **New Resources** on the **Properties** tab of the policy. You have the option of adding the links to the policy or the Ignore list.

Notes:

Select a policy to view the **Properties** tab in the right pane.

To set default search options for all new client disaster recovery policies, click **Options** on the **Tools** menu and click **Search Options** under Disaster Recovery.

To add selected links

1. In the Logical View, click a root or folder to display its properties in the right pane.
2. Perform one of the following actions:

   Under **Links** on the **Properties** tab of a root, select multiple links and drag them to a client disaster recovery policy.

   Under **Children** on the **Properties** tab of a folder, select multiple links and drag them to a client disaster recovery policy.
3. If some or all the links you selected have single targets, the **Disaster Recovery Link Target Wizard** provides you with additional options. When you finish going through the wizard, a status report appears.
4. If all the links you selected have multiple targets, a status report appears.

Add a share to a client disaster recovery policy

Drag a share from the **Physical View** tree, drag multiple shares from the **Properties** tab of a machine, or drag a machine from the **Physical View** tree to add all the shares on the machine to a client disaster recovery policy.

If the shares participate in the namespace, the links that point to the shares are added to the policy.

If the shares do not participate in the namespace, the **Disaster Recovery Link Target** wizard assists you in creating links.

If a share is the only target of a link, the **Disaster Recovery Link Target** wizard assists you in adding another target to the link.

Select search options

Customize the search process. When a client disaster recovery policy contains logical folders and/or DFS roots, the policy automatically monitors the folders and roots for new links and shows a list of links found. You have the option of selecting which links you want to add to the policy.

1. Select **Automatically add newly found links** when you want to search for new links and automatically add them to the policy.
2. Select **Automatically remove links from the policy when links are removed from a logical folder or DFS root that is a member of the policy** when you want to remove links from the policy that are removed from a logical folder or root that is a member of the policy.
3. Select **Automatically set the first link target online** when you want the first target of a link to always be online. To see the targets of a link, click the link. The list of targets appears in the **Properties** tab under **Target**.
4. Select **Ignore single-target links** when you do not want to search for links with only one target.

Schedule target monitoring

Schedule status monitoring of all link targets in a client disaster recovery policy. Status monitoring must be scheduled when automatic failover is selected for a policy.

To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.

Status monitoring of disaster recovery policies is performed in the VFM server or monitoring agent. Disaster recovery policies are monitored even if the VFM console is closed.

When link targets are monitored, VFM checks whether files in the target directory can be enumerated.

Links in a disaster recovery policy show "unknown" status until monitoring runs at least once for the policy.

To schedule target monitoring

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Properties**.

   Select a client disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.
2. Click **Target Monitoring Schedule**.
3. Click **New** and configure the schedule or select a schedule from the list to update.
4. If the policy fails to run at the scheduled time, you have the option of running the schedule again at a stated frequency. Select **Use alternate schedule when the policy fails** and type how frequently you want the schedule to run.
5. Click **OK**.

Schedule data replication

Schedule data replication for a disaster recovery policy using VFM replication.

Before you schedule data replication, review your network bandwidth and infrastructure. For example, review the rate of data changes and link speeds.

You to set the data transfer rate for a machine. In the Physical View tree, right-click a machine, click **Properties**, and click the **Agent** tab.

To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.

To schedule data replication

1. Perform one of the following actions:

   Right-click a disaster recovery policy and click **Properties**.

   Select a disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.
2. Click **Replication Schedule**.
3. Click **New** and configure the schedule.
4. Select **Use alternate schedule when the policy fails** and type a frequency (in minutes) to run the schedule any time replication fails. Any time a policy runs and fails, attempts are made to run the policy at a stated frequency until the policy runs successfully.
5. Click **OK**.

Select a SnapMirror

Select a SnapMirror for a client disaster recovery policy using SnapMirror replication. If one does not exist, you can create a new one. SnapMirrors are configured on destination filers.

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Properties**.

   Select a client disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.
2. Click **SnapMirror Replication**.
3. Click **Select a SnapMirror**.
4. Perform one of the following actions:

   Select a filer from the list provided. The list includes all the NetApp filers in the Physical View. If the filer contains SnapMirrors and they do not appear, click **Refresh Filer SnapMirror Information**. Click **Done**. Select an existing SnapMirror.

   To add a new SnapMirror, click **Create**. Type the source and destination information and click **OK**. Click **Done**.
5. Click **OK**.
6. To break the source/mirror relationship when a failover occurs, select **Break SnapMirror on Failover**. When the SnapMirror is broken, the destination data becomes writable. If you do not select this option, the destination data remains read-only after a failover. A quiesce is performed before the SnapMirror is broken.
7. Click **OK**.

Note: When configuring mirrored qtrees as primary and secondary targets of a disaster recovery link, you must complete the following steps:

1. Create the SnapMirror for the qtree.
2. Quiesce and break the SnapMirror.
3. Share the mirror qtree as a CIFS share.
4. Create a link that targets both the source and destination shares of the mirror.
5. Add the link to the disaster recovery policy.
6. Resync the SnapMirror.

Failover and replication options

There are several options that you can use to execute failover and replication for link targets that are members of a client disaster recovery policy.

Failover options

The following illustration explains the failover options that you can select for a policy.

For this example, assume that you have a client disaster recovery policy named Houston DR policy with two links, and each link has two targets. The primary target of both links is online.

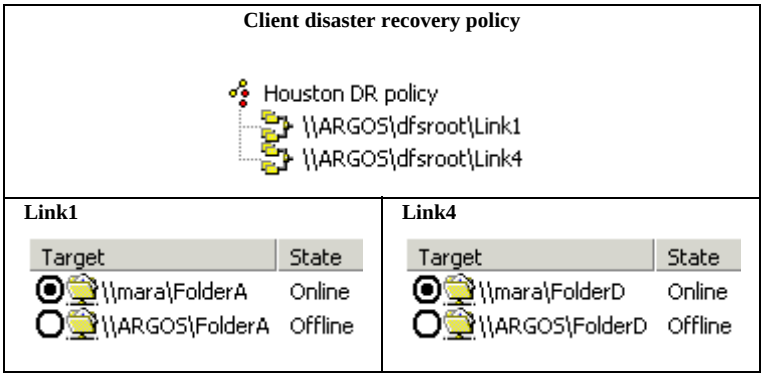

Now assume that the `\\mara\FolderA` of Link1 target becomes unavailable.

If you only selected **Automatically failover,** Link1 would fail over to \\Argos\FolderA, and Link4 would continue to have \\mara\FolderD as the online target.

If you selected **Automatically failover** and **Failover all links**, Link1 would fail over to \\ARGOS\FolderA and Link4 would fail over to \\ARGOS\FolderD.

If you selected **Automatically failover**, **Failover all links**, and selected to **Only monitor Link1**, Link1 would fail over to \\ARGOS\FolderA, and Link4 would fail over to \\ARGOS\FolderD. However, if you selected to only monitor Link4, failover would not initiate. This option is useful when all the link targets are contained in a single SnapMirror because the SnapMirror is broken as the result of a failover. If only one link is monitored, the amount of resources used in monitoring is reduced.

Replication options

The following list explains the two ways that you can synchronize the content of the link targets of a disaster recovery policy after a failover.

Replicate only to lower priority targets. The content of the online target is synchronized to each standby target in the order in which the targets are arranged on the **Properties** tab of the link that contains the targets. Replication is from the active target down to the standby targets.

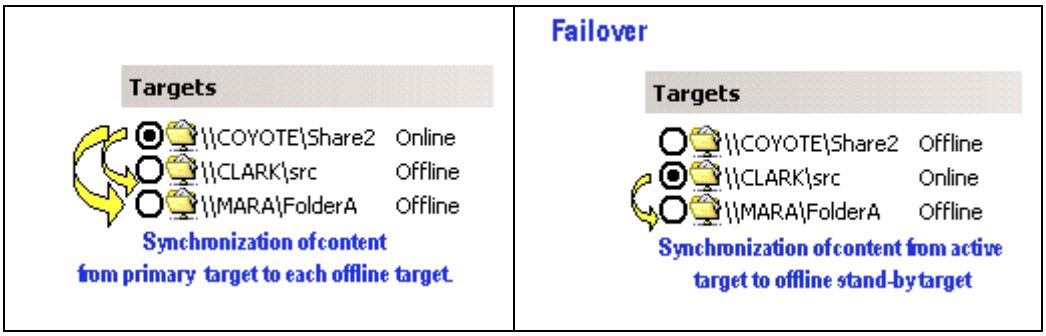

Replicate all targets. The content of the online target is synchronized to each standby target, regardless of the order in which the targets are arranged on the **Properties** tab of the link that contains the targets. Replication is from the active target to all the standby targets.

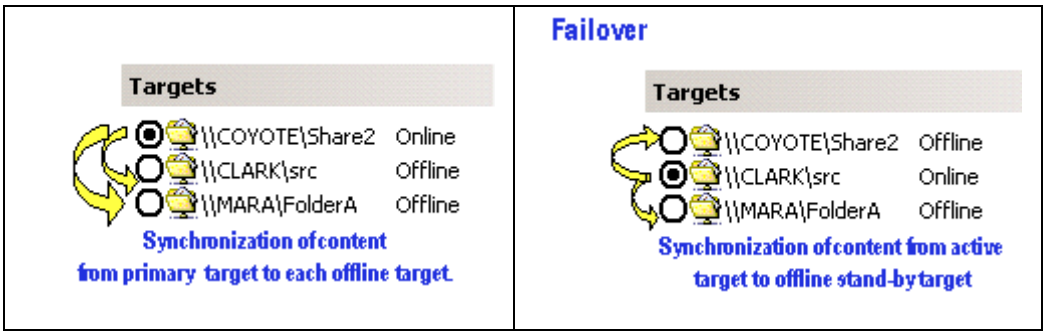

You have the option to synchronize the contents of the active target with the standby target when a failover or failback is initiated.

Note: If a client disaster recovery policy is using SnapMirror replication, see "Disaster recovery planning considerations" for information about how to resynchronize SnapMirrors for disaster recovery failback.

Select replication options

Select the replication options for a policy or enable the policy to select the options based on the optimization you want.

Note: To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.

1. Perform one of the following actions:

   Right-click a policy and click **Properties**.

   Select a policy and click Edit on the **Properties** tab that appears in the right pane.
2. Click **Replication Options**.
3. Click **Optimize For Speed** if you want to automatically set replication options when the time it takes to transfer the data is very important.

4. Click **Optimize For Security** if you want to automatically set replication options when the completeness and accuracy of the data transfer is very important.
5. Clear **Include subfolders** if you do not want to copy the source target subfolders and their contents to the destination targets.
6. Select **Delete orphaned files in destination folders** if you want to delete files in the destination targets that do not have a corresponding file in the source target.
7. If the destination targets do not support auxiliary file streams supported by the source target, you can select **Allow loss of additional file streams** if you want to copy files to the destination targets that lose auxiliary streams.
8. When you select **Enable Copy-in-place**, the source file directly replaces the destination file. This decreases the time required to perform a replication operation. When this option is not selected, the source file is copied to a temporary file in the destination directory, renames the current destination file, renames the temporary file to the actual file name, and deletes the destination file. This minimizes the chance of network errors adversely affecting file contents.
9. VFM can take advantage of the NTFS Change Journal on the replication source to determine which files are to be replicated rather than scanning the entire source to detect files that have changed. This option offers time savings when you are replicating large directory trees with limited changes scattered throughout the directory hierarchy.

   Notes:

   The NTFS Change Journal is supported on Windows 2000 NTFS, which is available on Windows 2000 and later versions of the operating system.

   When you enable the NTFS Change Journal on a volume, the size of the volume is used to determine the amount of space dedicated to the change journal. For each 50 GB of space on the volume, 128 MB is assigned for use by the change journal up to a maximum of 10 GB. On Windows 2000 SP2 or below, the amount of space dedicated to the change journal is always 128 MB, regardless of the size of the volume.

   Use of this option requires that VFM deploy a replication agent on the machine where the source data resides and enable the Change Journal on the NTFS volume hosting the source data. The machine properties dialog box enables you to control whether the replication agent can be deployed to a machine and whether the NTFS Change Journal can be enabled for volumes on the machine.

   Use of this option means that changes made on the source are replicated to the destination and changes made on the destination are not overwritten by data from the source unless the corresponding file on the source changes as well. This differs from standard VFM replication, wherein changes made on the destination are detected and overwritten with data from the source.

   In cases where the NTFS Change Journal wraps, a full scan of the source and destination must be performed to determine what files need to be transferred.

10. When **Use snapshot data** is selected, a snapshot or a shadow copy is taken of the volume in which the source data is located and the snapshot is used as the source. This reduces the possibility of not being able to transfer source files that are held open for exclusive access by an external application at the time of replication. The source data must be on a NetApp filer (CIFS or NFS protocol) or a machine running the Windows Storage Server 2003 or Windows Server 2003 file system.

Notes:

To reduce the number of snapshots that must be created, try to group sources in the same volume under the same policy.

When "Use snapshot data" is selected, you cannot enable NTFS Change Journal or file differential replication using cached hash information because snapshot data resides on read-only volumes and NTFS Change Journal and file differential replication require that the data be on writable volumes.

11. When **Override comparison timestamp resolution** is selected, a one second resolution is used when comparing time stamps between the source and destination. By default, the lowest resolution between the source and destination file systems is used when comparing time stamps.
12. In the **Copy files** list, select one of the following options:

    If destination file is different or missing - Replace files in the destination targets if the time stamp does not match the time stamp of files in the source target, and copy the files in the source target that are not found in the destination targets.

    If destination file is older or missing - Replace files in the destination targets that have time stamps older than files in the source target, and copy the files in the source target that are not found in the destination targets.

    If destination file is missing - Copy only files in the source target that are not found in the destination targets.
13. Type the number of times and the frequency, in seconds, that you want the policy to try to open a file.
14. Type how long, in minutes, a policy should run unsuccessfully before it is aborted.
15. Select one of the following options in the **Events Details** list:

    Only list files that encounter errors - Provide a detailed report of the files that encounter errors during the replication process.

    Don't list files - Do not provide a report of the files included in the replication process.

    List all files - Provide a detailed report of all the files included in the replication process.

Note: It is not recommended that you select "List all files" when you are migrating large numbers of files because it will slow down replication. When migrating large numbers of files, it is recommended that you run a file differences report after the migration to validate potential differences between the source and destination.

16. Select one of the following options in the **Post-Failover** list:

    Replicate all targets after failover - Synchronize the content of the online target to each standby target regardless of the order in which the targets are arranged on the **Properties** tab of the link that contains the targets. Replication is from the active target to all the standby targets.

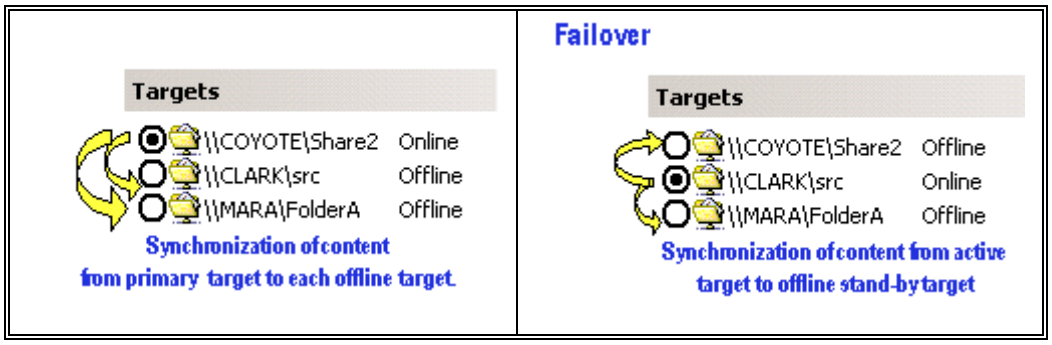

Replicate only to lower priority targets after failover - Synchronize the content of the online target to each standby target in the order in which the targets are arranged on the **Properties** tab of the link that contains the targets. Replication is from the active target down to the standby targets.

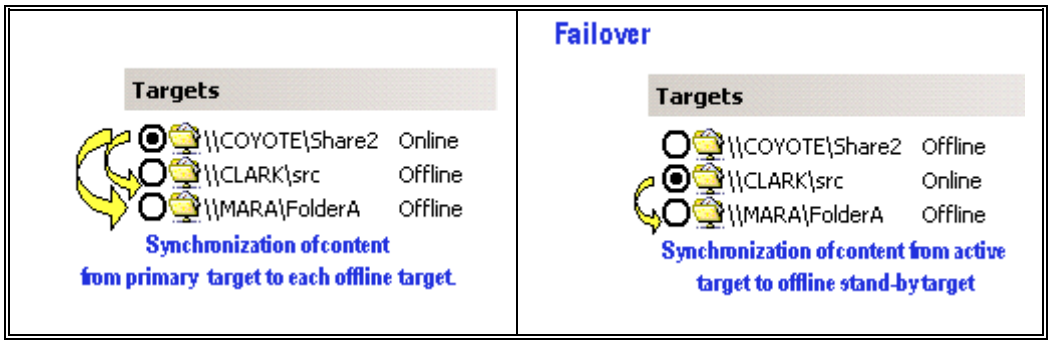

17. Click OK.

Select other replication options

Configure file security, replication differencing, and file attributes options for a client disaster recovery policy as well as other storage policies. See the following migration policy topics:

Specify replication filters

Select a replication agent

Select file security replication options

Select differential replication options

Select file attribute replication options

Notes:

NFS data transfer is not available for client disaster recovery policies.

When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result is different file permissions between the source and destination.

Failover

Set up automatic failover, manually fail over, or fail back one or multiple links.

The account used by the VFM server or monitoring agent, as well as any accounts used to run the VFM client, must have permission to make modifications to the DFS root containing the link or links in the disaster recovery policy. If the account lacks these permissions, VFM cannot control the online/offline state of the link targets.

Set up automatic failover

Automatic failover is performed in the order in which the targets are arranged. For example, if the first target on the list becomes unavailable, the link fails over to the second target on the list. If a given standby target cannot be activated, VFM tries to activate the next standby target in priority order.

When you set up automatic failover, you must also set up monitoring of the targets.

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Properties**.

   Click a client disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.
2. Click **General** and select **Automatically fail over**.
3. Click **OK**. All the links contained in the client disaster recovery policy fail over automatically.

Manually fail over one link

Perform one of the following actions:

Click a link under a client disaster recovery policy or click the link in the Logical View tree. The **Properties** tab for the link appears in the right pane.

1. Under **Target**, select a target below the active target.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active target to the target you selected. Failover is not performed. This enables you to synchronize data before you make an offline target the active target.

   Click **Change** to fail over to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.

Right-click the link under a disaster recovery policy or right-click the link in the Logical View tree, and click **Properties**.

1. Select a target below the active target.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active target to the target you selected. Failover is not performed. This enables you to synchronize data before you make an offline target the active target.

   Click **Change** to fail over to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.
3. Click **OK**.

Click a client disaster recovery policy. The **Properties** tab for the policy appears in the right pane.

1. Under **Links**, right-click a link and click **Failover**.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active target to the target below the active target. Failover is not performed. This enables you to synchronize data before you make an offline target the active target.

   Click **Change** to fail over to the target below the active target without replicating contents. Select this option if you are sure that the data contained in the offline target is synchronized with the active target.

Note: When you select to perform a failover from a menu option, VFM activates the next standby target in the order in which the targets are arranged.

Manually fail over all links

Manually fail over all or a selected number of links in a client disaster recovery policy. Failover is performed to the next standby target in the order in which the targets are arranged.

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Failover all Links**.

   Click a client disaster recovery policy to display the **Properties** tab in the right pane. Under **Links**, select all the links, right-click, and click **Failover**.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active targets to the standby target. Failover is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.

   Click **Change** to fail over to the standby targets without replicating contents. Select this option if you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

Manually fail back a link

Manually fail back a link in a disaster recovery policy. Failback is performed to the standby target above the active target in the order in which the targets are arranged in the link.

Perform one of the following actions:

Click a client disaster recovery policy to display the **Properties** tab in the right pane.

1. Under **Links**, right-click a link and click **Failback**. Failback is performed to the standby target above the active target in the order in which the targets are arranged.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active target to the target above the active target. Failback is not performed. This enables you to synchronize data before you make an offline target the active target.

   Click **Change** to fail back to the target above the active target without replicating contents. Select this option if you are sure that the data contained in the offline target is synchronized with the active target.

Right-click the link under a disaster recovery policy or right-click the link in the Logical View tree, and click **Properties**.

1. Select a target above the active target.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active target to the target you selected. Failback is not performed. This enables you to synchronize data before you make an offline target the active target.

   Click **Change** to fail back to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.
3. Click **OK**.

Click a link under a client disaster recovery policy or click the link in the Logical View tree. The **Properties** tab for the link appears in the right pane.

1. Under **Target**, select a target above the active target.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active target to the target you selected. Failback is not performed. This enables you to synchronize data before you make an offline target the active target.

   Click **Change** to fail back to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.

Manually fail back all links

Manually fail back all or a selected number of links in a client disaster recovery policy. Failback is performed to the standby target above the active target in the order in which the targets are arranged.

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Failback all Links**.

   Click a client disaster recovery policy to display the **Properties** tab in the right pane. Under **Links**, select all the links, right-click, and click **Failback**.

2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active targets to the standby target. Failback is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.

   Click **Change** to fail back to the standby targets without replicating contents. Select this option if you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

Note: If a client disaster recovery policy is using SnapMirror replication, see "Disaster recovery planning considerations" for information about how to resynchronize SnapMirrors for disaster recovery failback.

Run customized commands during failover processing

Specify batch files to run during disaster recovery failover processing. Separate entries are provided for scripts to be run before and after the link targets are failed over.

1. Perform one of the following actions:

   Right-click a disaster recovery policy and click **Properties**.

   Select a disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.

2. Click **Failover Scripting**.
3. Select **Run Batch Before Failover** when you want to run a batch file before a link is failed over. For example, if ten links must be failed over, the batch file runs ten times, once before each link fails over. The batch file can be specified either as the UNC path to a file accessible on your network or as a local path to a file residing on the server where monitoring is taking place. Type the path of a batch file or click **Browse** to locate a batch file on your network. Optionally select **Run On Manual Failover** if you want to run the batch file every time you perform a manual failover.
4. Select **Run Batch After Failover** when you want to run a batch file after a link is failed over. For example, if ten links are failed over, the batch file runs ten times, once after each link fails over. The batch file can be specified either as the UNC path to a file accessible on your network or as a local path to a file residing on the server where monitoring is taking place. Type the path of a batch file or click **Browse** to locate a batch file on your network. Optionally select **Run On Manual Failover** if you want to run the batch file every time you perform a manual failover.

5. Select **Run User Decision Script** when you want to run a batch file that determines when a link is considered failed. When VFM monitors a link, it runs the script. The batch file can be specified either as the UNC path to a file accessible on your network or as a local path to a file residing on the server where monitoring is taking place. Type the path of a batch file or click **Browse** to locate a batch file on your network
6. Click **OK**.

Notes:

When you select **Run On Manual Failover**, the scripts you specify to run before and after a failover will also run after a manual failback.

For detailed information about configuring batch files to run in VFM storage policies, see "Running batch files."

For more information about the example scripts provided by VFM, see "Batch files."

Run customized commands during replication

Specify batch files to run before and after a replication operation for disaster recovery policies using VFM replication.

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Properties**.

   Select a disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.
2. Click **Replication Scripting**.
3. Select the options you want and type or locate the batch files you want to run.
4. Click **OK**.

Notes:

For detailed information about configuring batch files to run in storage policies, see "Running batch files."

For more information about the example scripts provided by VFM, see "Batch files."

Add items to an ignore list

Maintain an ignore list of resources you do not want to add to the client disaster recovery policy. The items on the list are ignored when a search is performed for links to add to the policy.

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Properties**.

   Select a client disaster recovery policy to show a Properties tab in the right pane.

   Click Edit.
2. Click **Search Options**.
3. Type the path of a root, logical folder, or link, or click **Add** to browse for an item in your namespace to include in the list.
4. Click **OK**.

Notes:

The policy only searches for resources when a client disaster recovery policy contains logical folders and/or DFS roots.

You can also add links to the Ignore list from the New Resources list of a policy. See "Review the list of new resources."

Review the list of new resources

By default, the policy searches for new links and shows a list of links found. This enables you to select which links you want to add to a client disaster recovery policy.

Note: The policy only searches for links when a client disaster recovery policy contains logical folders or DFS roots.

1. Select a client disaster recovery policy to show a Properties tab in the right pane.
2. If the policy ran and new links were found, the links are listed under **New Resources**.
3. Right-click a link and perform one of the following actions:

   Click **Add to policy** to add the link to the policy.

   Click **Ignore** to add the link to the policy ignore list.

Notes:

If you want to automatically add new links to a client disaster recovery policy, perform the following actions:

1. Right-click a policy and click **Properties**.
2. Click **Search Options**.
3. Select **Automatically add newly found links**.
4. Click **OK**.
5. Click **Done**.

If you add a single-target link to the policy, the link appears under the root or logical folder with a warning status . To add another target to the link, perform the following actions:

1. Right-click the link and click **Add replica**.
2. Type the path of a shared folder or click **Add** to browse your network.
3. Click **OK**.
4. Click **Done**.

Update properties of a client disaster recovery policy

Update the data replication or monitoring schedule, change the name of a policy or change automatic failover.

1. Perform one of the following actions:

   Right-click a client disaster recovery policy and click **Properties**.

   Click a client disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.
2. Make the necessary changes and click **OK**.

Delete a client disaster recovery policy

Deleting a client disaster recovery policy does not delete the links or targets that are members of the policy from the namespace.

1. Right-click the policy and click **Delete**.
2. Click **Yes**. A status report appears.

Add new shares to a link in a client disaster policy

Add additional shares to a client disaster recovery policy. You can add existing shares or create new shares to add to a client disaster recovery policy.

To add existing shares to a link that is a member of a disaster recovery policy

Drag the share from the Physical View to the link in the policy or use the Share Creation wizard.

1. Right-click a link under a client disaster recovery policy and click **Create shares**.
2. Click **Next**.
3. Select **Disaster Recovery Share** and click **Next**.
4. Type the information requested and click **Next**.
5. Select the machines where you want to create the folder and click **Next**. You must create a disaster recovery share on at least two machines.
6. Make sure that the **Add the new shares as targets to an existing DFS Link** option is selected and click **Next**.
7. Make sure that the path of the disaster recovery link is displayed and click **Next**.
8. Select the client disaster recovery policy that contains the link from the list provided and click **Next**.
9. Click **Finish**. The share is created on the machines you selected and added to the disaster recovery link.

To create new shares to add to a disaster recovery policy

1. Perform one of the following actions:

   From the **Tools** menu, click **Create shares**.

   Right-click a link under a client disaster recovery policy and click **Create shares**.
2. Click **Next**.
3. Select **Disaster Recovery Share** and click **Next**.
4. Type the information requested and click **Next**.
5. Select the machines where you want to create the folder and click **Next**. You must create a disaster recovery share on at least two machines.
6. Select one of the following options:

   Add the new shares as targets to a new DFS Link

   a. Click **Next**.
   b. Select the DFS root that hosts the link and click **Next**.
   c. Optionally, select a folder in which to group the link.
   d. Click **Next**.
   e. If you do not want to create more links, click **Next**.

   Add the new shares as targets to an existing DFS Link

   a. Click **Next.**
   b. Select the link that targets the shares and click **Next**.
7. Select the client disaster recovery policy of which you want the link to be a member and click **Next**.

8. Click **Finish**. The shares are created and added to the link and the link appears under the disaster recovery policy folder.

Compare contents of targets in a client disaster recovery policy

Compare the contents of the online target with the contents of each of the standby targets contained in a disaster recovery policy.

1. Open the **Disaster Recovery Policies** folder under Admin View.
2. Double-click a client disaster recovery policy.
3. Perform one of the following actions:

   Right-click a link under the disaster recovery policy, point to **Reports**, point to **Physical**, and click **File Differences**.

   Click a link under the client disaster recovery policy to display its **Properties** tab. Select all the targets listed under **Targets**. Right-click, point to **Reports**, then click **File Differences**.

   The online target appears in the **Compare** field and the standby targets appear in the **To** field.
4. Click **Next**.
5. Type a name for your report and select how you want to save it.
6. Click **Finish**.

Add shared folder replicas to a link

Add multiple copies of a shared folder to a DFS link to provide fault tolerance.

Note: The target of a link must be one of the following types:

1 to N shared folders

1 to N stand-alone DFS paths

Only 1 domain-based DFS path

Configuring other combinations of link targets might prevent link target failover from functioning correctly.

1. Open the **Disaster Recovery** folder in the Admin View, right-click a link under a policy, and click **Add replica**.
2. Under **Send the user to one of these share folders or directories**, perform one of the following actions:

   Type the address of a folder.

   Click **Browse** to find a folder on the network. Expand the tree for the domain or click the domain. Double-click a machine to display a list of shared folders. Click **OK**.
3. To add another replica, click **Add** and repeat Step 2.
4. Click **OK**. A window displays the status for each shared folder to be added.

Notes:

The contents of a share folder replica set are synchronized only if the replica set is a member of a disaster recovery or replication policy.

For shared folder replica sets that are not a member of a disaster recovery or replication policy, you must either rely on other mechanisms when possible or manually ensure that the contents of the shared folders stay synchronized.

Use the **Ctrl** or **Shift** key to select multiple link targets at one time.

Prioritize the targets of a link

When you add a link with multiple targets to a disaster recovery policy, the targets are sorted in alphabetical order and all the targets of the link are set offline. You must specify which target is active and the order in which you want the targets to fail over.

Note: Make sure that the active or online target appears as the top or is the first target. For an explanation about how links are failed over in disaster recovery policies, see "Failover and replication options."

1. Select the target which should be active or online.
2. Specify the order in which you want the targets to fail over. Perform one of the following actions:

   Right-click the link under the disaster recovery policy or right-click the link in the Logical View tree and click **Properties**. Arrange the targets by dragging them to their new location, or right-click a target and click **Move up** or **Move down**. Click **OK**.

   Click the link under the disaster recovery policy or click the link in the Logical View tree. The **Properties** tab appears in the right pane. Under **Targets**, arrange the targets by selecting the target you want to move and clicking the up or down arrow. Click **Done**.

Remove a link from a policy

Removing a link from a disaster recovery policy does not delete the link from the namespace.

Right-click a link under a client disaster recovery policy and click **Remove from Disaster Recovery policy**. A status report appears.

View information about replication operations

You can view information about replication operations for all policies or just one policy. By default, the information for all replication operations appears. You can filter the information to see only replication operations with problems or only replication operations in progress.

To view information about replication operations for all policies, select the **Disaster Recovery Policies** folder and click the **Replication Tasks** tab in the right pane.

To view information about replication operations for one policy, select the policy and click the **Replication Tasks** tab in the right pane.

Start replication immediately

You can bypass the replication schedule and manually start the data replication process any time you want.

Perform one of the following actions:

Right-click a disaster recovery policy and click **Start Replication Now**.

Click a disaster recovery policy to display a **Properties** tab in the right pane. Right-click the **Links** heading and click **Start Replication Now**.

Notes:

This menu item is disabled if a data replication operation is in progress.

To view information about any replication in progress, click the **Replication Tasks** tab in the right pane and click .

When you configure a remote monitoring agent for a client disaster recovery policy, there is a short time lag between when you make the selection and when you can manually initiate replication on the policy through the "Start Replication Now" option. Attempts to initiate replication during this window return errors.

Cancel a replication in progress

Perform one of the following actions to stop a replication operation:

Right-click a client disaster recovery policy and click **Cancel Replication In Progress**.

Select a disaster recovery policy to show tabs in the right pane. Right-click anywhere on the **Properties** or **Status** tab and click **Cancel Replication In Progress**.

Note: To view information about any replication in progress, click the **Replication Tasks** tab and click .

Chapter 3 Filer Disaster Recovery

A wizard steps you through the process of creating a filer disaster recovery policy. To create a policy, follow the steps in "Create a filer disaster recovery policy."

Create a filer disaster recovery policy

You can easily set up SnapMirror replication from one or more filers to another filer or nearline storage device. A filer disaster recovery policy can monitor for new resources on the source filer. You select the new resources that will be added to the policy and you mirror the new resources to the destination filer. When the policy runs, it checks for new resources targeted by links and automatically creates corresponding shares on the destination which are then added as standby targets on the links. The policy can also monitor for the availability of a primary link target on the source filer and automatically fail over the link to the standby target on the destination filer when the primary target is not available.

Notes:

Before you begin, make sure that the destination filer contains volumes.

If you are creating a new qtree SnapMirror, the destination qtree cannot exist.

To create a filer disaster recovery policy

1. Perform one of the following actions:

   On the **Tools** menu, click **New Policy**, point to **Disaster Recovery Policy**, and then click **Filer Recovery Policy**.

   Right-click the Disaster Recovery Policies folder in the Admin View, click **New**, and then click **Filer Recovery Policy**.

   Right-click the Filer Recovery Policies folder and click **New Filer Recovery Policy**.

   Right-click Admin View, click **New**, point to **Disaster Recovery Policy**, and then click **Filer Recovery Policy**.
2. Click **Next**.
3. Type a word or phrase that best describes this policy. The name appears under the **Disaster Recovery Policies\Filer Recovery Policies** folder and on the **Properties** tab of items that are members of the policy.
4. Type the name of the destination filer or click **Browse** to select a filer known to VFM.
5. Optionally select the name of a volume on the destination filer.
6. By default, **Break SnapMirror on failover** is not selected. When a failover occurs, the destination data remains read-only. Select to break the source/mirror relationship when a failover occurs, if you want the destination data to be writable. A quiesce is performed before the SnapMirror is broken.

7. Select **Update associated policies on failover** if links in the filer disaster recovery are part of a multiprotocol namespace or namespace availability policy, and you want to create new automount and export files referencing the new active targets after a failover.
8. By default, **Automatically failover** is selected. All the links targeting data on source volumes and qtrees are monitored. When the active target for a link is not available, VFM automatically switches to a standby target of the link. Status monitoring must be scheduled. If you clear this option, you are notified when an active target is not available.
9. By default, the current VFM server is used to process the policy. If you installed monitoring agents on other machines, you can select one of the locations where monitoring agents are installed as the location from which this policy is processed.
10. Click **Next**. A list of the DFS roots in your namespace appears.
11. Optionally remove a DFS root from the list that does not contain links targeting data on source NetApp filers. During the failover process, users accessing data targeted by links in the selected roots are redirected to the standby secondary target on the destination filer.
12. Click **Next**.
13. Click **Add** and select the source volumes and/or qtrees.
14. Click **Next**.
15. If a destination for a SnapMirror source does not exist, click the down arrow under the destination column and select **Add New SnapMirror**.
16. Configure the destination for the SnapMirror and click **OK**. The source and destination are mirrored.
17. Repeat Steps 15 and 16 for each SnapMirror source that does not have a destination.
18. Click **Next**.
19. Select **Find new Links** if you want to monitor for new links targeting data on the source volumes.
20. Select **Find new QTrees** if you want to monitor for qtrees added to a source filer.
21. Specify how often you want to monitor for new resources.
22. If the policy fails to run at the scheduled time, you are given the option of running the schedule again at a stated frequency. Select **Use alternate schedule when the policy fails** and type how frequently you want the schedule to run.
23. If you want to allow a link to fail a specified number of times before failing over to the standby target, select **Do not failover until link fails at least ...times** and type how often it can fail. This option enables you to avoid a failover in response to transient errors.
24. Click **Next**.
25. Specify how often you want to monitor the availability of the volumes and qtrees that are member of the policy. If you selected **Automatically failover**, you must schedule status monitoring of links.
26. Optionally select **Use alternate schedule when the policy fails** if you want to attempt to run the policy at a stated frequency until the policy runs successfully.
27. Click **Next**.
28. Optionally specify batch files to run during disaster recovery failover processing. Separate entries are provided for scripts to run before and after the link targets are failed over. Type the path of the batch file or click **Browse** to locate the file in your network.
29. Optionally select **Run On Manual Failover** if you want to run a batch file when you

manually failover link targets.

30. Click **Next**.
31. Click **Finish**. The policy checks for links targeting the source and automatically creates shares on the destination.

Notes:

To enable a vFiler to participate in a filer disaster recovery policy, the filer that hosts the vFiler must be known to VFM.

For information about how to resynchronize SnapMirrors for disaster recovery failback, see "Disaster recovery planning considerations."

Add volumes and qtrees to the policy

Add volumes and qtrees to an existing filer disaster recovery policy.

1. Perform one of the following actions:

   Right-click a filer disaster recovery policy and click **Add Volumes/QTrees**.

   Select a filer disaster recovery policy to show a properties tab in the right pane. Right-click anywhere on the properties tab and click **Add Volumes/QTrees**.
2. Click **Add** and select a volume from the list of NetApp filers.
3. Click **OK**.
4. Repeat Steps 2 through 3 to add more volumes to the list.
5. Optionally select qtrees on the selected volumes.
6. Click **Next**. A list of source volumes and qtrees appears. If a source is mirrored in the destination filer, the destination appears.
7. If a source is not mirrored, you can create the SnapMirror. Select a row and click **Add New SnapMirror** from the list. Type the destination for the SnapMirror. Click **OK**.
8. Click **Finish**.

Note: To enable a vFiler to participate in a filer disaster recovery policy, the filer that hosts the vFiler must be known to VFM.

Fail over immediately

Manually fail over to the standby offline target of all links participating in a filer disaster recovery policy. Users are redirected to the mirrored data on the destination filer. If the option to break the SnapMirror on failover is not selected on the policy, the data on the destination filer is read-only.

To manually fail over all the links participating in a filer disaster recovery policy, perform one of the following actions

Right-click the policy and click **Failover All**.

Select the policy to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Failover All**.

To manually fail over all the links targeting data on one filer or volume participating in a filer disaster recovery policy

Perform one of the following actions:

Right-click the filer or volume under a filer disaster recovery policy folder and click Failover All.

Select the filer or volume to show a Properties tab in the right pane. Right-click anywhere on the page and click Failover All.

To manually fail over the links targeting data on a qtree participating in a filer disaster recovery policy, right-click the qtree and click **Failover**.

To manually fail back a link or multiple links, execute the **Failback** and **Failback All** commands in the same manner as is described for Failover.

Note: See "Disaster recovery planning considerations" for information about how to resynchronize SnapMirrors for disaster recovery failback.

Run customized commands

Specify batch files to run during filer disaster recovery failover processing. Separate entries are provided for scripts to be run before and after the link targets are failed over. See the following client disaster recovery policy topics:

"Run customized commands during failover processing"

"Run customized commands to determine a failed link"

Review the list of new qtrees

When a filer disaster recovery policy is set to monitor for new qtrees on the source volumes, VFM displays a list of new qtrees and enables you to determine whether you want to add them to the policy.

To add a new qtree to the policy

1. Select a filer disaster recovery policy to show a **Properties** tab in the right pane.
2. Under **New QTrees**, select the qtrees you want to add to the policy, right-click, and click **Add To Policy**.
3. Click **Next**.
4. If a qtree is not mirrored to the destination filer, you must create the SnapMirror. Click the row under the **Destination** column where a new qtree appears and click **Add New SnapMirror**.
5. Select or type the destination qtree, click **OK**, and click **Done**.
6. Click **Finish**.

To add a new qtree to the ignore list of the policy

When a scan is performed for new qtrees, the policy ignores the qtrees listed in the policy ignore list.

1. Select a filer disaster recovery policy to show a **Properties** tab in the right pane.
2. Under **New QTrees**, select the qtrees you want to add to the policy ignore list, right-click, and click **Ignore**.

Note: To review a filer disaster recovery ignore list, see "Update the policy ignore list."

Update the policy ignore list

Add or remove qtrees from the ignore list of a filer disaster recovery policy.

1. Perform one of the following actions:

   Right-click a filer disaster recovery policy and click **Properties**.

   Select a filer disaster recovery policy to show a **Properties** tab in the right pane. Click Edit.
2. Click **Ignored QTrees**.
3. To add a qtree on a source volume that you do not want to participate in the policy, click **Add**. Select the qtree and click **OK**.
4. To remove a qtree from the list, select a qtree and click **Remove**.
5. Click **OK**.

Note: Qtrees are automatically added to the ignore list of the policy in two ways:

When you add a qtree on a volume with multiple qtrees, the qtrees that were not selected are added to the ignore list.

When the policy scans for new qtrees and you select a group of qtrees to add to the policy, the new qtrees not added to the policy are automatically added to the ignore list.

View the status of a SnapMirror

View the status of SnapMirrors associated with a filer disaster recovery policy:

To view all the SnapMirrors associated with a filer disaster recovery policy, select the policy to display a list of SnapMirrors on the **Properties** tab that appears in the right pane.

To view all the SnapMirrors associated with a source filer, select the filer under the filer disaster recovery policy to display a list of SnapMirrors on the **Properties** tab that appears in the right pane.

To view the SnapMirror associated with a source qtree, double-click a filer under the filer disaster recovery and select a qtree to display the SnapMirror on the **Properties** tab that appears in the right pane.

View links targeting data on a filer

View the links in your namespace that target data on filers participating in a filer disaster recovery policy.

1. Under a filer disaster recovery policy, select a filer, volume, or qtree to show a list of links on the **Properties** tab that appears in the right pane.
2. Double-click a link to view its properties.

View information about replication operations

You can view information about replication operations for all policies or just one policy. By default, the information for all replication operations appears. You can filter the information to see only replication operations with problems or only replication operations in progress.

To view information about replication operations for all policies, select the **Disaster Recovery Policies** folder and click the **Replication Tasks** tab in the right pane.

To view information about replication operations for one policy, select the policy and click the **Replication Tasks** tab in the right pane.

Update properties of a policy

Update the properties of a filer disaster recovery policy.

1. Perform one of the following actions:

   Right-click a filer disaster recovery policy and click **Properties**.

   Click a filer disaster recovery policy and click Edit on the **Properties** tab that appears in the right pane.
2. Make the necessary changes and click **OK**.

Remove a filer, volume, or qtree from a policy

Removing a filer, volume, or qtree from a filer disaster recovery policy does not delete the SnapMirrors.

Perform one of the following actions:

Right-click a filer, volume, or qtree under a filer disaster recovery policy and click **Remove From Filer DR Policy**.

Select a filer, volume, or qtree under a filer disaster recovery policy to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Remove From Filer DR Policy**.

Delete a filer disaster recovery policy

Deleting a filer disaster recovery policy does not delete the SnapMirrors on the destination filer.

1. Perform one of the following actions:

   Right-click the policy and click **Delete**.

   Select the policy to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Delete**.

2. Click **Yes**.

Chapter 4 Server Disaster Recovery

To create a server disaster recovery policy, follow the basic steps listed below:

1. Create a server disaster recovery policy.
2. Select resource monitoring options.
3. Schedule resource monitoring.
4. Schedule status monitoring and replication.

Create a server disaster recovery policy

Create a disaster recovery policy to make data on servers highly available. You can configure the policy to monitor for new resources on the primary server, schedule replication, and automatically fail over when a failure is detected.

Notes:

Roots containing links participating in server disaster recovery policies can also be members of namespace availability policies.

1. Perform one of the following actions:

   On the **Tools** menu, click **New Policy**, point to **Disaster Recovery Policy**, and click **Server Recovery Policy**.

   Right-click the **Disaster Recovery Policies** folder in the Admin View, point to **New**, and click **Server Disaster Recovery Policy**.

   Right-click the **Server Recovery Policies** folder and click **New Server Recovery Policy**.

   Right-click Admin View, click **New**, point to **Disaster Recovery Policy**, and click **Server Recovery Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field. The name appears under the **Server Recovery Policies** folder and on the **Properties** tab of items that are members of the policy.
3. Type the name of a server or NetApp filer in the **Primary** and **Secondary Server** fields.
4. Select Automatically failover when you want to automatically fail over all the links in the policy when a failure is detected.
5. Select Update associated policies on failover if links in the server disaster recovery are part of a multiprotocol namespace or namespace availability policy and you want to create new automount and export files referencing the new active targets after a failover.
6. You can select another monitoring location where monitoring of the targets of the server disaster recovery policy is performed. The default is to monitor from the VFM server. From the **Use the Monitor on** list, select the hostname of another monitoring agent.

Notes:

The account under which the VFM server is running must have the appropriate permissions to create shares and subdirectories on the secondary server.

Server disaster recovery policies do not use SnapMirror replication. If your primary and secondary servers are filers, it is recommended that you create a filer disaster recovery

policy.

If you uninstall a monitoring agent, ensure that disaster recovery policies that formerly used the uninstalled agent are reassigned to a currently available agent.

Select resource monitoring options

Select the namespaces you want to monitor for links targeting the primary server. When new links are found, you can configure what you want to do with them.

1. Perform one of the following actions:

   Right-click a server disaster recovery policy and click **Properties**.

   Select a server disaster recovery policy to show a **Properties** tab in the right pane. Click Edit.
2. Click **Resource Monitoring Options**.
3. Click **Add** and type the name of a DFS root or logical folder in your namespace that you want VFM to monitor for links targeting the primary server.
4. Repeat Step 3 to add more items to the list.
5. Select what you want to do when it finds links that do not have a target on the secondary server. Normally VFM shows the link with an error status and does not modify the link. You can select one of the other options:

   Add secondary target, use the primary target share as a template. A share is created on the secondary server with the same name, security settings, and local path as the target on the primary server.

   Add secondary target, create a new directory and share. Type a local path and a name for the share is created. Select this option when there is a possibility that the secondary server might have a share with the same name.
6. Select what you want to do when it finds links with more than two targets. Normally VFM shows the links with an error status and does not remove the targets that are not on the primary or secondary servers. You can select to remove all but the targets on the primary and secondary servers.
7. Normally all new links found are placed in a Nonactive Links list. Links in the Nonactive Links list are not monitored or replicated. Select **Automatically add conformant links to the Active Links list** if you want to automatically add new links with two targets (one target each in the primary and secondary servers) to the Active Links list. Links in this list are monitored and replicated.
8. You can add links to a server disaster recovery policy by dragging them to the policy. This method enables you to add links to the Nonactive Links list that do not have targets on the primary server. After the links are added to the Nonactive Links list, you can add a target on the primary server and move the links to the **Active Links list**. Select **Automatically remove links with no Primary Server target** if you want to automatically remove all the links you drag to the policy that do not have a target on the primary server.
9. Click **OK**.

Schedule resource monitoring

Set up the schedule to periodically monitor for new links on the primary server.

1. Perform one of the following actions:

   Right-click a server disaster recovery policy and click **Properties**.

   Select a server disaster recovery policy to show a **Properties** tab on the right pane. Click Edit.
2. Click **Resource Monitoring Schedule**.
3. Click **Add** and configure the schedule.
4. If the policy fails to run at the scheduled time, you have the option of running the schedule again at a stated frequency. Select **Use alternate schedule when the policy fails** and type how frequently you want the schedule to run.
5. Click **OK**.

Schedule target monitoring and replication

You can schedule status monitoring for the link appearing in the **Monitor field** in the **General** dialog box. When a failure is detected on the link, all the links in the policy are failed over.

Note: When you do not select a link to monitor, a link is randomly selected from among the links that are members of the policy.

To schedule data replication on the server disaster recovery policy, see the following client disaster recovery policy topics:

Schedule target monitoring

Schedule data replication

Select replication options

Configure replication, file security, and replication differencing options for a server disaster recovery policy as well as for other storage policies. See the following migration policy topics:

Select replication options

Specify replication filters

Select a replication agent

Select file security replication options

Select differential replication options

Select file attribute replication options

Notes:

NFS data transfer is not available for server disaster recovery policies.

When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.

Add links to a server disaster recovery policy

You can choose how you want to add links to a server disaster recovery policy. The policy scans for links with a target on the primary server and adds those links to a Nonactive Links list.

To add links to the Active Links list

When links are in the **Active Links** list, the link targets are monitored and the data from the online target is replicated to the offline target. Only links with two targets, one each on the primary and secondary servers, can be moved to the **Active Links** list. To move a link from the **Nonactive Links** list to the **Active Links** list, perform one of the following actions:

Drag the link to the **Active Links** list.

Right-click the link and click **Move to Active Links List**.

Note: If you want to automatically add these links to the **Active Links** list, select **Automatically add conformant links to the Active Links list** on the **Resource Monitoring Options** dialog box. See "Select resource monitoring options" for instructions.

To repair links in the Nonactive Links list

Links with an error status are added to the Nonactive Links list when they do not have a target on the secondary server and when they have more than two targets. You must update the properties of the links before you can add them to the Active Links list.

To view information about the error status of a link, click the **Status** tab for the policy.

To update the links, perform one of the following actions:

Double-click the path of a link in the **Status** tab.

Right-click a link in the **Properties** tab and click **Properties**.

Double-click a link in the **Properties** tab then click Edit

Note: You can configure the policy to remove extra targets on links found that contain targets on the primary and secondary servers. You can also configure the policy to add a second target to links found that do not contain a target on the secondary server. See "Select resource monitoring options" for instructions.

To drag a link to a server disaster recovery policy

You can add links manually to a server disaster recovery policy. You can drag a link to a policy and the link is added to the **Nonactive Links** list.

Notes:

When a link in the **Nonactive Links** list no longer has an error status, you can drag the link to the **Active Links** list.

If you want to remove a link from the **Active Links** list, perform one of the following actions:

Drag the link to the **Nonactive Links** list.

Right-click the link and click **Move to Nonactive Links list**.

You can configure the policy to automatically remove links you drag to the policy that do not have a target on the primary server. See "Select resource monitoring options" for instructions.

Start monitoring for new links immediately

You can manually start monitoring for new links on the primary server.

Perform one of the following actions:

Right-click a server disaster recovery policy and click **Start Resource Monitoring Now**.

Select a server disaster recovery policy to show a **Properties** tab in the right pane. Right-click anywhere on the tab and click **Start Resource Monitoring Now**.

Select the **Server Recovery Policies** folder to show a **Properties** tab in the right pane. Right-click a policy and click **Start Resource Monitoring Now**.

Assign a link for the policy to monitor

Server disaster recovery policies monitor the status of one link. When a failure is detected for the link, all the links are failed over to the offline targets in the secondary server. If you do not select the link you want to monitor, a link is randomly selected when the policy runs.

1. Perform one of the following actions:

   Right-click a server disaster recovery policy and click **Properties**.

   Select a server disaster recovery policy to show a **Properties** tab in the right pane and click Edit.
2. On the General dialog box, select a link from the **Monitor link** list.
3. Click **OK**.

Note: The **Monitor link** list contains all the links that are members of the policy. If no links appear on the list, you can add links to the policy by manually searching for new resources or scheduling resource monitoring.

Cancel resource monitoring in progress

You can cancel resource monitoring that is in progress.

To check whether monitoring is in progress

1. Select a policy to show a **Properties** tab in the right pane.
2. Scroll down to **Resource Monitoring in Progress**.

To cancel resource monitoring in progress

Perform one of the following actions:

Right-click the policy and click **Cancel Resource Monitoring In Progress**.

Select the policy to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Cancel Resource Monitoring In Progress**.

Select the **Server Recovery Policies** folder to show a **Properties** tab in the right pane. Right-click a policy and click **Cancel Resource Monitoring In Progress**.

Manage replication operations

Replication operations in server disaster recovery policies are managed in the same manner as in client disaster recovery policies. See the following topics:

View information about replication operations

Start replication immediately

Cancel a replication in progress

Start failover or failback immediately

You can manually fail over or fail back all the links in a server disaster recovery policy.

Note: The account used by the VFM server or monitoring agent, as well as any accounts used to run the VFM client, must have permission to make modifications to the DFS root containing the links in the server disaster recovery policy. If the account lacks these permissions, VFM cannot control the online/offline state of the link targets.

To fail over all links manually

Manually fail over all the links in a server disaster recovery policy. Failover is performed to the offline target in the secondary server.

1. Right-click a server disaster recovery policy and click **Failover all Links**.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active targets to the standby target. Failover is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.

   Click **Change** to fail over to the offline targets without replicating contents. Select this option when you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

To manually fail back all links

Manually fail back all links in a server disaster recovery policy. Failback is performed to the offline target in the primary server.

1. Right-click a server disaster recovery policy and click **Failback all Links**.
2. Perform one of the following actions:

   Click **Replicate** to replicate contents from the active targets to the offline targets. Failback is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.

   Click **Change** to fail back to the offline targets without replicating contents. Select this option when you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

Remove links from a server disaster recovery policy

To remove links from a server disaster recovery policy, perform any of the following actions:

Select links that appear under the policy, right-click, and click **Remove from Disaster Recovery policy**.

Select a policy to show a **Properties** tab in the right pane.

1. Select the links under the **Active Links** list, right-click, and click **Remove from Disaster Recovery policy**.
2. Select a link or links under the **Nonactive Links** list, right-click, and click **Remove from Disaster Recovery policy**.

Right-click a link in the Logical View tree that is a member of a server disaster recovery policy and click **Remove from Disaster Recovery policy**.

Delete a server disaster recovery policy

Deleting a server disaster recovery policy does not delete the links or targets that are members of the policy from the namespace.

1. Right-click the policy and click **Delete**.
2. Click **Yes**.

Note: A policy cannot be deleted if it contains active links. To remove active links from a server disaster recovery policy, see "Remove links from a server disaster recovery policy."

Appendix

VFM Guides

The following table lists the VFM guides. The list is arranged in ascending order by document title.

No.	Title	File name	Contents
1.	*Advanced Configuration and Techniques*	VFMAdvConfig.pdf	Explains how to use VFM to configure security permissions Explains how to manage replication and monitoring agents Explains how to run batch files in policies and reports Explains how to make the VFM server highly available
2.	*Create and Manage Namespaces*	VFMNamespace.pdf	Explains how to create a DFS namespace (roots, links, and link targets), how to quickly populate the namespace, make the namespace highly available, create multiprotocol namespaces, virtualize user home directories, and back up the namespace.
3.	*Getting Started*	VFMGettingStarted.pdf	A tutorial to get you started using VFM. It explains how to install VFM and steps you through the process of creating a logical namespace, making the namespace highly available, making data highly available, transferring CIFS and NFS data, creating multiprotocol namespaces, virtualizing user home directories, and backing up data and the namespace.
4.	*Make Data Highly Available*	VFMMakeDataAvailable.pdf	Explains how to make data highly available by creating client, filer, and server disaster recovery policies
5.	*Manage Physical Resources*	VFMPhysicalResources.pdf	Explains how to use VFM to Work with NetApp filers, set replication properties, and work with Windows and UNIX machines

No.	Title	File name	Contents
6.	*Manage Snapshots on Filers*	VFMSnapshots.pdf	Explains how to use VFM to manage snapshots on NetApp filers using snapshot access and snapshot scheduling policies
7.	*Optimize Storage*	VFMOptimizeStorage.pdf	Explains how to create and manage archival migration policies to optimize storage. Explains how to load balance Windows servers by creating storage load balancing policies
8.	*Reports*	VFMReports.pdf	Explains how to create and publish reports
9.	*Transfer Data*	VFMTransferData.pdf	Explains how to transfer CIFS and NFS data by creating migration and replication policies Explains how to transfer data between NetApp filers by creating SnapMirror policies
10.	*Troubleshooting*	VFMTroubleshooting.pdf	Includes diagnostic information, such as troubleshooting, event IDs, migration, replication, and status indicators Provides examples of file exclusions and inclusions

Glossary

A

ACE: Access Control Entry, an entry in an Access Control List.

ACL: Access Control List, a list associated with a file that contains information about which users or groups have permission to access or modify the file.

C

CIFS: Common Internet File System. A file-sharing protocol based on the Microsoft Server Message Block (SMB). It enables you to share file systems over the Internet and Intranet. The CIFS protocol is used primarily for file sharing by Windows platforms.

Client: The user interface component of the application.

Consolidation DFS root: A special stand-alone DFS root that has its configuration information stored locally on a member server. The root must be hosted on Windows Server 2003.

D

Dependencies: The servers, drives, and shared folders that together contain the data published through the DFS namespace

DFS: Microsoft Distributed File System consists of software residing on network servers and clients that transparently links shared folders on different file servers into a single namespace for improved load balancing and data availability. NOTE: This should not be confused with Transarc's DFS product, which is a separate distributed file system product that is not interoperable with the Microsoft product.

DFS link: Located under the namespace root. Forms a connection to one or more shared folders, links or to another root.

DFS namespace: The logical view of shared file resources seen by users.

DFS path: A UNC path starting at a DFS root.

DFS root: The share at the top of the namespace topology. The starting point for the links and shared folders that make up the namespace.

DFS topology: Overall logical hierarchy of a namespace, including elements such as roots, links, shared folders, and replica sets.

DFS tree: The consolidated DFS view provided by the application, which includes the Admin View, Logical View, Physical View, and Reports.

Domain-based DFS root: A DFS root that has its configuration information stored in Active Directory. The root can have multiple root targets, which offers fault tolerance and load sharing at the root level.

E

Export: SEE NFS export

F

Fault-tolerant DFS: SEE Domain-based DFS root

FQDN: Fully-qualified Domain Name. A host, domain name, and a top-level domain. For example, www.brocade.com is a fully-qualified domain name where www is the host, .brocade is the second-level domain, and .com is the top-level domain.

FRS: Microsoft File Replication Service.

G

GID: Group Identification Number. A unique number assigned to identify a group of users in UNIX environments.

K

Key-pair: Consists of a public key and a corresponding private key. It is safe to distribute a public key, but a private key must be kept secret.

L

Link target: A DFS link can have as targets: 1 to N shared folders, OR 1 to N stand-alone DFS paths, OR only 1 domain-based DFS path.

M

Migration candidates: Items that match migration criteria in an archival migration policy. You can select which items you want to migrate. Items not selected for migration are automatically added to the exclusion list of the archival migration policy.

Migration task: Identifies the source and destination.

Monitoring agent: The component that can be installed to evaluate storage policies from a different machine than the one where the application server is installed.

N

NFS: Network File System developed by Sun Microsystems. A file-sharing protocol that enables you to share file systems over the Internet and Intranet. NFS is used with both UNIX and non-UNIX computers.

NFS export: A directory that is published for access by NFS clients. The export specifies a directory to be exported and which machines have access to the directory. NFS exports are similar to CIFS shares in that they provide access to directories.

P

Pool: A group of storage locations.

Primary server: An active application server or monitoring agent.

Primary target: In a disaster recovery policy, this is the active target to which users are referred by the DFS root.

Q

qtree: A special directory of the root of a volume that behaves like a virtual subvolume with special attributes.

Quiesce: Wait for all existing transfers to both volumes and qtrees to complete and then block any further updates.

R

RAID: Redundant Array of Independent Disks enables file access even when a disk in a RAID group is damaged.

Rebase: To change the destination path proposed by the application to the actual path to where the data is to be migrated.

References: The DFS link targeting the selected physical item. For example, when you click a server under the Physical View tree, properties and associated DFS links appear.

Replica: A DFS root or DFS shared folder that provides an alternate path for a DFS link. NOTE: In Windows NT 4 terminology, this is referred to as an alternate volume.

Replica set: Two or more DFS roots or DFS shared folders that provide alternate paths for a DFS link.

Replication agent: The component deployed by the application server or monitoring agent during the configuration of storage policies that transfer CIFS and NFS data, or when you run a file differencing report.

Report schedule group: Enables you to set a schedule to automatically gather information. A report schedule group can be configured to start during off hours to offload network and computing resources during business hours. You can produce multiple reports from the data gathered without generating additional network traffic.

Report source group: Specifies the location where you want to gather information. It is assigned to a report schedule group. A report source group also enables you to specify batch files to run before and/or after the information is gathered.

Report type: Determines what information is gathered. Report types are assigned to report source groups.

S

Secondary target: SEE Standby target.

Server: The component that responds to requests from the client component, monitors the DFS namespace, and processes storage policies.

Shared folder: File or folder in the DFS namespace that can be accessed by users who have proper permissions. NOTE: In Windows NT 4, this is referred to as a DFS volume, leaf, or branch.

SID: Security identifier that uniquely identifies a Windows user or group.

Stand-alone DFS root: A DFS root that has its configuration information stored locally on the host server. The root has a single root target. When the root target is unavailable, the data referenced by links under the root is inaccessible.

Standby target: A target in a disaster recovery policy that is offline. During a failover, a standby target can become the active target.

Storage location: A combination of a server UNC path and the local path pointing to a directory in the file system. For example, \\rama:c:\ is a storage location representing the C drive on machine "rama" and \\secure.npb.com:/vol/vol0/myqtree is a storage location representing the qtree "myqtree" on volume vol0 on secure.npb.com.

T

Target: LINK TARGET: A shared folder, a DFS link, or a DFS root that is referenced by a DFS link. ROOT TARGET: A shared folder that is referenced by a DFS root. SEE ALSO Link target.

Trustee: The user account, group, or logon session to which an Access Control Entry (ACE) applies.

U

UID: User Identification Number. A unique number assigned to identify a user in UNIX environments.

UNIX file permissions: Control access to files and directories. Permissions can be applied to the owner of a file, a group, or everyone on the UNIX network.

W

Widelink namespace: A logical namespace hosted on a NetApp filer CIFS share that provides the benefits of a namespace without requiring additional Windows servers.

Index

A

Active target 17
Add a client disaster recovery policy.....5
Add a DFS link.....9
Add a link 9
Add links 6, 7
Arrange targets 26
Automatic failover.....17, 26

B

Batch files.....20, 21

C

Cancel replication.....42
Change journal 13
Change properties of a client disaster recovery policy.....23
Client disaster recovery policy
- add a link to a policy 6, 9, 25
- add items to ignore list 21
- add multiple links.....7
- add new shares 23
- automatically add new links.....9
- automatically remove links 9
- cancel replication 27
- compare contents of targets.....25
- create a policy 5
- delete a policy 23
- planning.....12
- prioritize targets 26
- remove a link.....26
- replication operation information.....26, 34
- review search results 22
- run batch files.....20, 21
- schedule data replication 10
- schedule status monitoring.....9
- search options.....9
- select replication options.....13
- select SnapMirror replication 11
- start replication immediately.....27
- update properties 23

Compare contents of targets 25
Content replication 27
Contents of targets.....25
Create a client disaster recovery policy 5
Create a server disaster recovery policy 37

D

Data move 27
Data replication 10, 27
Data synchronization 10, 12
Delete
- a client disaster recovery policy 23
- a filer disaster recovery policy 35
- a link from a client disaster recovery policy 26

Delete..... 23
DFS link
- add new shares 23
- assign folders..... 25
- remove a link from a policy 26

DFS root
- add to client disaster recovery policy 7

DFS root 7
Disaster recovery
- for NetApp filers 31, 33
- preparing for..... 2
- resynchronizing SnapMirrors for failback..... 2

Disaster recovery 31
DR policies 23

E

Exclusion list 21
Exclusions..... 13

F

Failback 12, 31, 42
Failover.....12, 17, 20, 26, 31, 32, 42
File differences 25
File security options..... 39
File transfer
- exclusions..... 13
- inclusions 13

File transfer..... 13
Filer disaster recovery policy
- add new qtrees to policy..... 32
- add volumes and qtrees 31
- delete a policy 35
- manual failover..... 31
- remove a filer from a policy 34
- remove a qtree from a policy..... 34
- remove a volume from a policy..... 34
- replication operation information.....26, 34
- run batch files..... 32
- update ignore list 33
- update properties 34
- view links targeting data on filers 34
- view SnapMirror status 33

Filer disaster recovery policy..... 31
Filers
- remove from a filer disaster recovery policy34

Folder
- add a folder..... 25
- add multiple folders..... 25
- create a share 25
- permissions..... 25

I

Ignore list.....21, 32, 33
Inclusions 13

L

Link target
  compare contents with other targets.....25
Link targeting data on filers.....34
Logical folder
  add to client disaster recovery policy.....7
Logical folder 7

M

Manual failover 17, 31
Manually start data replication 27
Monitor link targets 9
Monitor status.....9
Monitoring.....38
Move data.....27

N

NetApp filers
  links targeting data on filers.....34
NTFS Change Journal 13

P

Planning considerations
  disaster recovery 2
Planning considerations.....2
Primary target.....17, 26
Prioritize targets 26
Properties
  of a client disaster recovery policy.....23
  of a filer disaster recovery policy.....34
Properties.....23
Properties.....34

Q

qtree
  add to ignore list.....32, 33
  remove from a filer disaster recovery policy34
qtree.....32
qtree.....33
qtree.....34

R

Replica
  compare contents of targets.....25
Replication
  SnapMirror replication.....11
Replication.....10
Replication.....11
Replication differencing options 39
Replication scripts 21
Resource monitoring
  cancel monitoring in progress 42
Resource monitoring..... 42
Resource monitoring options..... 38
Resynchronizing SnapMirrors for disaster
  recovery failback..... 2
Review search results 22
Run batch files.....20, 21

S

Schedule data replication..... 10
Schedule replication 23, 39
Schedule resource monitoring 39
Schedule status monitoring.....9, 23, 39
Script settings 21
Scripts..... 32
Search options 9
Security information..... 13
Server disaster recovery policy
  add links to policy 40
  cancel a replication in progress 42
  cancel resource monitoring in progress..... 42
  check link status 40
  create a policy 37
  delete a policy 43
  Drag links to policy..... 40
  fail back all links 42
  fail over all links..... 42
  file security replication options 39
  manually monitor for new links 41
  monitor status of a link..... 41
  remove links from policy 43
  replication differencing options..... 39
  replication options..... 39
  resource monitoring options..... 38
  schedule replication..... 39
  schedule resource monitoring..... 39
  schedule status monitoring..... 39
  start replication immediately..... 42
  status monitoring..... 41
  view information about replication operations
    42
Server disaster recovery policy..... 40
Servers
  disaster recovery policy..... 37
Servers..... 37
Set up a client disaster recovery policy 5
Shared folder
  add a shared folder replica..... 25
  assign to a link..... 25
  create a share..... 25
  create shares 23
  permissions..... 25
SnapMirror
  select SnapMirror replication 11
  view status..... 33
SnapMirror replication 29

Start replication 27, 42
Status monitoring 9, 23, 41
Synchronization of data.....12

U

Update properties of a client disaster recovery policy.....23

V

vFilers
add to filer disaster recovery policy 29
vFilers..... 29
View target of a link 34
Volumes
remove from a filer disaster recovery policy34

W

Windows 2000..... 5