



VFM (Virtual File Manager) v6.0

Transfer Data

May 29, 2007

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 migrate and replicate CIFS and NFS data seamlessly.

The guide is divided into the following sections:

- **Chapter 1 Introduction** –Provides information that you should consider before transferring data and explains the four policies that VFM provides to transfer data.
- **Chapter 2 Migration Policies** – Explains how to create and manage policies to transfer CIFS and NFS data.
- **Chapter 3 Replication Policies** – Explains how to create and manage policies to distribute CIFS and NFS data to heterogeneous devices using VFM replication.
- **Appendix A Specifying migration tasks from a file** – Explains how you can customize a migration policy.
- **Appendix B Example Task Creation Batch file** – Provides an example Task Creation batch file that you can specify in a migration policy.
- **Appendix C Example Task Description file** – Provides an example Task Description file.
- **Appendix D VFM Documentation** – 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
	Indicates that the information applies only to DFS roots on Windows 2000 or Windows Server 2003, Enterprise Edition.
	Indicates that the information applies only to domain-based roots.
	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 .

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)

Note: To get your product build number and license information, click **About VFM** 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 VFM**.
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**.

Note: After you change the activation key, you must restart the VFM client.

Table of Contents

PREFACE	I
CHAPTER 1 INTRODUCTION.....	1
DATA MIGRATION CONSIDERATIONS.....	1
<i>Planning considerations</i>	<i>1</i>
<i>Transferring NTFS file security attributes.....</i>	<i>3</i>
<i>Transferring between dissimilar file systems.....</i>	<i>3</i>
<i>Specifying migration tasks from a file</i>	<i>4</i>
REPLICATION PLANNING CONSIDERATIONS	4
<i>How it works.....</i>	<i>4</i>
<i>VFM replication planning considerations.....</i>	<i>5</i>
<i>Transferring NTFS file security attributes.....</i>	<i>5</i>
<i>Permissions.....</i>	<i>5</i>
<i>Status monitoring.....</i>	<i>6</i>
CHAPTER 2 MIGRATION	7
MIGRATION LIMITATIONS	7
INTEGRATED NAMESPACE MANAGEMENT	7
MIGRATION USAGE OPTIONS	8
SCHEDULING CONSIDERATIONS	10
<i>Policy migration schedule</i>	<i>10</i>
<i>Using start dates.....</i>	<i>10</i>
<i>Running task creation scripts to add migration tasks to a policy.....</i>	<i>15</i>
<i>Scheduling scenarios</i>	<i>16</i>
CREATE A MIGRATION POLICY CHECKLIST	18
CREATE A MIGRATION POLICY	18
CONFIGURE THE INITIAL PHASE.....	19
SCHEDULE THE INITIAL PHASE	20
CONFIGURE THE INCREMENTAL PHASE	20
SCHEDULE THE INCREMENTAL PHASE	21
CONFIGURE THE FINAL PHASE.....	21
SCHEDULE THE FINAL PHASE	23
<i>To schedule the final phase.....</i>	<i>23</i>
SPECIFY NAMESPACE OPTIONS	23
<i>How automount files are updated.....</i>	<i>23</i>
<i>To specify namespace options.....</i>	<i>24</i>
SELECT REPLICATION OPTIONS.....	24
SPECIFY REPLICATION FILTERS	27
SELECT A REPLICATION AGENT	28
SELECT DIFFERENTIAL REPLICATION OPTIONS	29
SELECT FILE SECURITY REPLICATION OPTIONS.....	30
SELECT FILE ATTRIBUTE REPLICATION OPTIONS.....	32
ADD A MIGRATION TASK	33
<i>To add a migration task.....</i>	<i>35</i>
ADVANCE TO FINAL PHASE	38
<i>To specify the conditions for advancing to the final phase at the policy level.....</i>	<i>38</i>
<i>To specify the condition for advancing a migration task to the final phase</i>	<i>39</i>
SPECIFY MIGRATION TASKS FROM A FILE	40
RUN CUSTOMIZED COMMANDS DURING REPLICATION.....	40
UPDATE THE PROPERTIES OF A MIGRATION TASK.....	41
START A MIGRATION TASK IMMEDIATELY	41
START ALL TASKS FOR A POLICY IMMEDIATELY.....	41
VIEW PROGRESS OF A TASK.....	41

Transfer Data

USER ACTION REQUIRED	42
SUSPEND A MIGRATION TASK.....	43
REMOVE A TASK	43
DELETE A MIGRATION POLICY.....	43
CHAPTER 3 REPLICATION	45
CREATE A REPLICATION POLICY CHECKLIST	45
CREATE A REPLICATION POLICY.....	45
SELECT REPLICATION OPTIONS.....	46
SCHEDULE DATA DISTRIBUTION.....	47
<i>To schedule data distribution</i>	<i>47</i>
CHANGE DATA DISTRIBUTION PATH.....	47
CHANGE REPLICATION FOR A TARGET.....	48
ADD OR REPLACE THE SOURCE.....	48
ADD OR REPLACE THE TARGET.....	49
CHANGE THE REPLICATION SCHEDULE.....	50
START REPLICATION IMMEDIATELY	50
CANCEL A REPLICATION IN PROGRESS.....	50
VIEW INFORMATION ABOUT REPLICATION OPERATIONS.....	50
COMPARE CONTENTS OF TARGETS IN A REPLICATION POLICY	51
VIEW PROPERTIES OF A REPLICATION POLICY	51
DELETE A DESTINATION TARGET.....	51
DELETE A REPLICATION POLICY	51
APPENDIX A SPECIFYING MIGRATION TASKS FROM A FILE	53
APPENDIX B EXAMPLE TASK CREATION BATCH FILE	55
APPENDIX C EXAMPLE TASK DESCRIPTION FILE	57
APPENDIX D DOCUMENTATION.....	59
GLOSSARY	61
INDEX.....	65

Chapter 1 Introduction

This section provides instructions about how to create policies to transfer data. VFM provides two policies to move data. Select the policy that best meets your needs.

- **Migration policies** enable you to balance free space and data throughput capacity across volumes or to migrate data prior to maintaining or consolidating servers. Use migration policies to transfer data in heterogeneous environments. Before you create your first policy, see "Data migration considerations."
- **Replication policies** enable you to centrally manage tasks associated with replicating data across a network. You can select the optimal replication mechanism for your configuration, whether this is VFM for heterogeneous environments, SnapMirror for NetApp filers, or another mechanism suited to your environment. VFM provides status information about replication operations performed using the VFM replication agent. Use replication policies to replicate data from a central location to remote sites and from remote sites to a central location. Before you create your first policy, see "Replication planning considerations."

IMPORTANT: If you have users and applications that are tied to specific shares and file locations through hard-coded UNC pathnames, you need to take extra precautions to see that such links are maintained as part of the migration process. UNC references can occur in a variety of settings, including Microsoft Office documents, applications, and other files. The **UNCUpdate utility** is a tool you can use to report on files containing UNC references to storage resources under consideration for migration. **UNCUpdate** can report on UNC entries encountered in a wide variety of files. **UNCUpdate** enables you to update UNC references to target the storage resources at their new location. Contact Brocade for more information about this product.

Data migration considerations

Before you migrate data, it is important that you understand transfer results when different file systems, such as FAT or NFS, are involved.

Planning considerations

- Users and applications should not be accessing the source or destination folders while the migration is taking place.
- Do not use differential replication on a LAN. When operating on a high-speed network, the computational cost of calculating the file differences is more expensive in terms of time than simply copying the entire file from source to destination.
- When you perform full backups during a migration operation, backups can change archive attribute settings and trigger massive differences between source and destination. With VFM you can avoid this problem, because migration policies enable you to specify how you want to set attributes of destination files.

- When configuring replication options, consider carefully how much event detail you need. If you gather information about all files, it affects performance. As a compromise between performance and the need for feedback, it is recommended that you gather information only when an error is encountered.
- It is not recommended that you enable real-time virus scanning on filers because migration throughput can be severely reduced.
- When planning a migration, consider the available time you have to complete the operation and how much data you have to migrate. Use include and exclude lists sparingly because they add processing time to the migration operation.
- (CIFS only) By default, the security descriptor is copied every time a file is copied. Various options are provided for controlling how a file's security is copied during data migration. See "Select file security replication options" for details.
- Exclusive locks on files on source or destination prevent VFM from accessing those files for transfer. Find out who or what program is locking the files and make the files available to VFM.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended duplications of files and can interfere with VFM replication.
- API support limitations in the NetWare® server affect the operation of VFM. VFM can transfer data from Novell® to CIFS shares. Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained.
- (CIFS only) The account used by the VFM server, a replication agent, or a monitoring agent must have the following permissions:
 - **SeBackupPrivilege** ("Back up files and directories") - required to open source directories/files with backup semantics.
 - **SeRestorePrivilege** ("Restore files and directories") - required to open destination directories/files with backup semantics.
 - **SeSecurityPrivilege** ("Manage auditing and security log") - required to access SACLS (System Access Control Lists)
 - **SeTakeOwnershipPrivilege** ("Take ownership of files or other objects") - required to delete directories/files.
- (NFS only) The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports. The exports must be writable on the destination.

Note: When the VFM server shuts down, the migration policy and the migration tasks running on the server are canceled. VFM also tries to cancel replication jobs on the replication agent. The state of the migration policy and the migration task are persisted. If the server is shut down gracefully, when it restarts the migration policy starts again where it left off.

Transferring NTFS file security attributes

- When you transfer file security attributes, VFM looks up the trustees configured as the owner and those referenced in the source file's ACL. If the owner of the file is set to an unknown SID and you attempt a transfer, VFM reports error 1332, "No mapping between account names and security IDs was done." You can correct this by assigning ownership.
- 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.

Transferring between dissimilar file systems

Some file attribute information (for example, security attribute) might be lost when copying or moving files from one file system to another because of the inherent differences between the file systems. The following table displays transfer results between different file systems.

Source			Destination			
	NTFS W2K	NTFS NT4	FAT16	FAT32	NetApp /CIFS	NFS
NTFS W2K	Attributes maintained see Note 4	Some attributes lost	Some attributes lost See Notes 2 and 6	Some attributes lost See Notes 2 and 6	See Note 5	Some attributes lost See Notes 1 and 3
NTFS NT4	Attributes maintained	Attributes maintained	Some attributes lost See Notes 2 and 6	Some attributes lost See Notes 2 and 6	See Note 5	Some attributes lost See Notes 1 and 3
FAT16	Attributes maintained	Attributes maintained	Attributes maintained	Attributes maintained	See Note 5	Some attributes lost See Notes 1 and 3
FAT32	Attributes maintained	Attributes maintained	Some attributes lost	Attributes maintained	See Note 5	Some attributes lost See Notes 1 and 3
NetApp CIFS	Attributes maintained	Attributes maintained	Some attributes lost	Some attributes lost	Attributes maintained	Some attributes lost See Notes 1 and 3
NFS	See Note 3	See Note 3	See Note 3	See Note 3	See Note 3	See Note 3
NetWare	Some attributes lost See Note7	Some attributes lost See Note7	Some attributes lost See Note7	Some attributes lost See Note7	Some attributes lost See Note7	Not applicable

Note 1: NFS maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to an NFS file system from NTFS, their security and extended attributes are subsumed by NFS attributes. VFM UNIX replication agents perform NFS to NFS data transfer. Other file systems require redirector client software on the machine hosting the Windows replication agent.

Note 2: NTFS stores access times in UTC, while FAT stores access time in LocalTime. VFM does not adjust between UTC and LocalTime.

Note 3: Copying files from NFS sources is supported. This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.

Note 4: Copying hard links, reparsing points, or NTFS encrypted files are not supported.

Note 5: You can configure security on NetApp filers in a variety of ways. These variables affect how security attributes are maintained when copying to NetApp filers. When files are transferred from Windows systems to a NetApp filer and sparse or compressed attributes are set, the attributes have no effect on how the files are stored on the filer.

Note 6: FAT maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to a FAT file system from NTFS, their access times, extended attributes, indexing attributes, and alternate streams (including security) are lost when the files are written to the FAT destination.

Note 7: Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained. If the version of NetWare is 5.1 or earlier, the Novell client must be installed on the same machine on which the replication agent is installed. Windows and Novell accounts must be configured with identical user name and password. The VFM console and server components must also run under the same user name.

Specifying migration tasks from a file

Migration policies provide the option of running a batch file to add migration tasks to the policy. You can specify the batch file in the **Run task creation batch** field on the **General** tab of a data migration policy properties dialog box. For information about the batch file works, see Appendix A. For example task creation batch file and task description file, see Appendixes B and C respectively.

Replication planning considerations

You can choose the optimal replication mechanism for your configuration, whether it is the VFM agent for heterogeneous environments or another mechanism for homogeneous environments. VFM enables you to configure replication at the folder level.

How it works

- A replication policy is created and the schedule is defined.
- A single-target link is added to the policy as the source link. You can add the link during the creation of the replication policy, or drag it from the Logical View tree or from the right pane to the policy.
- A multi-target link is added to the policy as the destination link. You can add the link during the creation of the replication policy, or drag it from the Logical View tree or from the right pane to the policy. You can quickly add additional targets to the destination link.

- When you select a replication policy in the Admin View, several tabs appear in the right pane. The Replication Topology tab provides a graphical display of the distribution path of the policy. From this tab:
 - You can change the distribution path by a simple drag-and-drop operation.
 - You can view monitoring status of the links and each of their targets.
 - You can add additional targets to the destination link by right-clicking the link.

VFM replication planning considerations

VFM uses file-to-file copy to synchronize data across a network.

- Exclusive locks on files on source or destination prevent VFM from accessing the file for transfer. Find out who or what program is locking the files, and make the files available to VFM.
- It is not recommended that you perform full backups during a replication operation because backups can change archive attribute settings and trigger massive differences between source and destination.
- When configuring replication options, carefully consider how much event detail you need. If you gather information about all files, it affects performance. As a compromise between performance and the need for feedback, it is recommended that you gather information only when an error is encountered.
- When planning replication, consider the available time you have to complete the operation and how much data you have to replicate. Use include and exclude lists sparingly because they add processing time to the replication operation.
- It is not recommended that you enable real-time virus scanning on filers because replication throughput can be severely reduced.
- The benefits of differential replication are realized on slower links rather than on a LAN. Computational costs outweigh the benefits of differential replication on a LAN.
- A replication policy synchronizes from its source to a number of read-only copies.
- There is a replication time delay between primary and secondary contents. If a primary target becomes unavailable, there is no guarantee that all the contents of the primary target are replicated to the standby targets.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files and can interfere with VFM replication.
- The VFM client does not have to be running for replication policies to function.

Transferring NTFS file security attributes

For information about transferring file security attributes, see "Data migration considerations."

Permissions

For information about setting permissions for the account used by the VFM server, a replication agent, or a monitoring agent, see "Data migration considerations."

Status monitoring

When you use the VFM replication agent, status information about replication operations performed is available for the policy and its targets.

- If status monitoring is turned on for the links and folders contained in a policy, status indicators appear for the items on the **Replication** tab that appears when you select a policy.
- Monitoring DFS resources generates network traffic between the VFM console and the resources being monitored. Slower polling rates generate less traffic but provide less-timely notification of problems in the monitored namespace.

Chapter 2 Migration

Migration policies enable you to manage CIFS and NFS data migration easily. In addition to the data, attributes and permissions are transferred.

- Data transfer can be configured in phases. This enables you to copy data while users have access to the data and configure a final replication at a time you specify when users are temporarily denied access to the source data and are directed to the destination.
- The migration operation is divided into three phases.
 - In the **initial phase**, you can scan for potential problems and perform a baseline copy of the data.
 - In the **incremental phase**, you can copy the data incrementally and specify when to advance to the final phase.
 - In the **final phase**, you can update the namespace, prevent users from connecting to the data, delete the source data, perform a final copy of the data, and stop sharing source or export.
- Before you start creating a migration policy, see "Data migration considerations." The topic covers such information as the permissions the account used by the VFM server and monitoring agent must have.
- You can configure the migration operation to execute all three phases (a typical migration) or you can configure the migration operation to execute one or two phases. For more information, see "Migration usage options."
- Before you configure migration schedules, see "Scheduling considerations."

Migration limitations

The following files are not supported:

- Encrypted
- Reparse points
- Hard links

Integrated namespace management

When the source participates in the namespace, you can update the namespace.

DFS namespace management

- Retarget link - The destination folder becomes the new link target.
- Add link replicas - The destination folder is added as a new replica of the link.

UNIX namespace management

The references to the automount files you specify are updated.

Migration usage options

No.	Usage	Steps
1	Typical data migration (source data might be changing) - Optionally scan for potential problems before copying data. - Perform a baseline copy of the data. - Perform incremental copies of the data. - Advance to the final phase.	Run the initial phase, incremental, and final phases. - Select initial phase actions you want to run. - Specify when you want to run the initial phase. - Make sure that the incremental phase action "Copy the data" is selected. - Specify how often you want to run the incremental phase. - Specify when you want to advance to the final phase. - Select the final phase actions you want to run. - Specify when you want to run the final phase.
2	Data migration when the source data is static - Optionally scan for potential problems before copying data. - Copy the data. - Advance to the final phase.	Run the initial and final phases. - Select initial phase actions you want to run. - Specify when you want to run the initial phase. - Clear "Copy the data" action in the incremental phase. - Select the final phase actions you want to run. - Specify when you want to run the final phase.
3	Data migration when the baseline copy was performed using another method, such as tape backup.	Run the incremental and final phases. - Make sure that all the actions are cleared in the initial phase. - Make sure that the incremental phase action "Copy the data" is selected. - Specify how often you want to run the incremental phase. - Specify when you want to advance to the final phase.

No.	Usage	Steps
		5. Select the final phase actions you want to run. 6. Specify when you want to run the final phase.
4	Data migration when the source data is static and baseline copy was performed using another method, such as tape backup.	Run the final phase only. 1. Make sure that all the actions are cleared in the initial phase. 2. Clear "Copy the data" action in the incremental phase. 3. Select the final phase actions you want to run. 4. Specify when you want to run the final phase.
5	Do not copy data. Only scan for potential problems that might interfere with copying the data.	Run the initial phase only. 1. Select the initial phase action "Scan for potential problems copying the data." 2. Specify when you want to run the initial phase. 3. Clear all the incremental phase actions. 4. Clear all the final phase actions.
6	Only perform baseline and repeated incremental copies of the data. A cutover date has not been decided.	Run the initial and incremental phases. 1. Select initial phase actions you want to run. 2. Specify when you want to run the initial phase. 3. Make sure that the incremental phase action "Copy the data" is selected. 4. Select the condition "Never, Remain in the Incremental phase" for advancing to the final phase. 5. Specify how often you want to run the incremental phase.
7	Data migration when baseline copy was performed using another method, such as tape backup. A cutover date has not been decided.	Run the incremental phase only. 1. Make sure that all the actions are cleared in the initial

No.	Usage	Steps
		phase. 2. Make sure that the incremental phase action "Copy the data" is selected. 3. Select the condition "Never, Remain in the Incremental phase" for advancing to the final phase. 4. Specify how often you want to run the incremental phase.

Scheduling considerations

Before you configure schedules for each migration phase, it is important that you understand how the different schedule options work.

Policy migration schedule


- The policy migration is divided into three phases: initial, incremental, and final. Each phase has a schedule. These schedules control when each migration task associated with the policy executes the actions you specify in the migration phases.
- You can schedule when each migration phase starts. You can also schedule how often a migration task runs the action specified in the incremental phase and how often the final phase should attempt to run a migration task (the final phase can run only one time for each migration task).
- A policy can contain multiple incremental phase and final phase schedules. For example, you can set up a schedule to start the incremental phase every Monday through Friday at midnight, and you can set up another schedule to start the incremental phase every Saturday and Sunday at 6:00 A.M.
- You have the option of manually specifying when a migration task or multiple tasks advance to the final phase. For information about how advancing to the final phase works, see "Scheduling scenarios."

Using start dates

- Make sure that the incremental phase start date is set to run after the initial phase start date. The incremental phase will not run until the initial phase is completed. The incremental phase runs on the next interval run after the initial phase is completed.

Initial phase scheduled before incremental phase


The following example illustrates how a migration task runs through the migration phases when the initial phase is scheduled to start before the incremental phase.



- After the initial phase completes on 4/2, the incremental phase runs.
- The six scheduled incremental copies of the data are executed.

Incremental phase scheduled before initial phase

The following example illustrates how a migration task runs through the migration phases when the incremental phase is scheduled to start before the initial phase.



- The incremental phase, which was scheduled to start on 4/1, cannot start until 4/6 because the migration task must complete the initial phase before it can run the incremental phase.
- The incremental phase was scheduled to run six times but runs only two times. The first four incremental intervals never run.
- The incremental phase runs on 4/6, which is the next scheduled run of the phase

after the completion of the initial phase.


- Configure the final phase start date to run at intervals when you do not know when the incremental phase will be completed. If you set the final phase schedule to run only once and the incremental phase is running at the time that the final phase attempts to run, the final phase never runs. You have to manually advance the migration operation to the final phase.

Final phase scheduled to run at intervals

The following example illustrates how a migration task runs through the migration phases when the final phase is scheduled to attempt to run multiple times.

Scenario

- The incremental phase schedule is configured to run at 9:00 P.M. every day.
- The policy is configured to advance to the final phase on 4/12 at 3:00 P.M.
- The final phase schedule is configured to run on 4/12 at 9:00 P.M. and to attempt to run at 9:00 P.M. on 4/13, 4/14, and 4/15.




- It takes longer to complete the fourth incremental copy of the data.
- The fourth incremental copy of the data runs past 3:00 P.M. on 4/12. The migration task cannot advance to the final phase because the incremental phase is running at the time that the migration task was scheduled to advance to the final phase.
- The migration task completes the fourth incremental copy of the data.
- The final phase does not run at 9:00 P.M. on 4/12 because the fourth incremental copy of the data runs beyond the final phase scheduled start time.
- The migration task runs the final phase starting at 9:00 P.M. on 4/13. The final phase runs only one time.

Final phase not scheduled to run at intervals

The following example illustrates how a migration task runs through the migration phases when the final phase is scheduled to attempt to run only one time.

Scenario


- The incremental phase schedule is configured to run at 9:00 P.M. every day.
- The policy is configured to advance to the final phase on 4/12 at 3:00 P.M.
- The final phase schedule is configured to run on 4/12 at 9:00 P.M.



- It takes longer to complete the fourth incremental copy of the data.
- The fourth incremental copy of the data runs past 3:00 P.M. on 4/12. The migration task cannot advance to the final phase because the incremental phase is running at the time that the migration task was scheduled to advance to the final phase.
- The migration task cannot run the final phase at 9:00 P.M. on 4/12 because the fourth incremental copy is still running.
- The migration task completes the fourth incremental copy of the data.
- Because the final phase schedule was not configured to attempt to run at intervals, the final phase will never run. You must manually advance the migration task to the final phase.


- Do not schedule a large gap between the baseline copy and the incremental copy of the data.

The following example illustrates how a migration task runs through the migration phases when there is a large time gap between the baseline copy and incremental copy of the data.



- There is a nineteen-day gap between the baseline copy and the beginning of the incremental copy of the data. Within this time, the data might have changed significantly.
- As a consequence, the incremental copy of the data will take longer to complete.
- Do not schedule a large gap between the last incremental copy and the final copy of the data.

The following example illustrates how a migration task runs through the migration phases when there is a large time gap between the last incremental copy and the final copy of the data.



- There is a twelve-day gap between the last incremental copy and the beginning of the final copy of the data. Within this time, the data might have changed significantly.
- As a consequence, the final copy of the data will take longer to complete.

Running task creation scripts to add migration tasks to a policy

Task creation script runs every time the initial, incremental, and final phase schedules run. How you configure the phase migration schedules determines how often the script runs.

Scenario

- A task creation script automatically adds migration tasks to PolicyA.
- The initial phase is scheduled to start on 4/1.
- The incremental phase is scheduled to start on 4/3 and run daily.
- Migration tasks advance to the final phase on 4/6.
- The final phase is scheduled to start on 4/6 and run daily.

Execution of the task creation script

The script runs once on 4/1, 4/3, 4/4, and 4/5. It runs twice on 4/6 (one time when the incremental phase schedule runs and one time when the final phase schedule runs). After 4/6, the script runs every time the incremental and final phase schedules run.

On 4/1

1. PolicyA runs.
2. The task creation script runs.
3. Migration taskA is added to the policy.
4. Migration taskA runs the initial phase.

On 4/3

1. PolicyA runs.
2. The task creation script runs.
3. Migration taskB is added to the policy.
4. Migration taskA runs the incremental phase.

On 4/4

1. PolicyA runs.
2. The task creation script runs.
3. Migration taskC is added to the policy.
4. Migration taskA runs the incremental phase.

On 4/5

1. PolicyA runs.
2. The task creation script runs.
3. Migration taskD is added to the policy.
4. Migration taskA runs the incremental phase.

On 4/6

1. PolicyA runs.
2. The task creation script runs.
3. Migration taskE is added to the policy.
4. Migration taskA advances to the final phase and runs the final phase once.

Important: Migration tasks B, C, D, and E did not run because they were added to PolicyA *after* the initial phase schedule start date of 4/1. You must start the migration tasks manually.

Scheduling scenarios

The scheduling scenarios explain how the execution of the migration operation is affected by the options you select for advancing to the final phase.

Example:

You set the policy schedule to run as follows:

- You set up the initial phase schedule to perform a baseline copy of the data on Saturday and Sunday.
- You set up the incremental phase schedule to perform incremental copies of the data Monday through Thursday at 8:00 P.M.

Scenario 1: Manual advance to final phase while incremental phase is not running

You do not want the migration operation to advance to the final phase. In the incremental phase, you specify **Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually**. The policy runs. The baseline copy and incremental copies of the data are completed. On Thursday at 7:00 P.M., you want to transition the migration tasks to the final phase. At the time that you transition to the final phase, the incremental phase is not running. You right-click a migration task or multiple tasks and click **Advance to Final Phase**.

- If you select **Now**, the migration tasks transition immediately to the final phase.
- If you select **After the current Incremental Phase run**, the migration tasks transition to the final phase the next time the final phase is scheduled to run.
- If you select **On the first Final Phase interval after**, the migration tasks transition to the final phase the next time the final phase is scheduled to run after the transition date you specify.

Scenario 2: Manual advance to final phase while incremental phase is running

You do not want the migration operation to advance to the final phase. In the incremental phase, you specify **Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually**. The policy runs. The baseline copy is completed and incremental copies of the data are in progress. On Thursday at 9:00 P.M., you want to transition the migration tasks to the final phase. At the time that you transition to the final phase, the incremental phase is running. You right-click a migration task or multiple tasks and click **Advance to Final Phase**.

- Do not select **Now**. If you select **Now**, the tasks do not transition to the final phase. The policy continues to execute the incremental phase, and you get the error message "The policy is already in progress" after you click **OK** in the Advance to Final Phase dialog box.
- If you select **After the current Incremental Phase run**, the migration tasks transition to the final phase the next time the final phase is scheduled to run.
- If you select **On the first Final Phase interval after**, the migration tasks transition to the final phase the next time the final phase is scheduled to run after the transition date you specify.

Scenario 3: Run the incremental phase once and manually advance to the final phase

You want the migration operation to advance to the final phase after running the incremental phase one time. In the incremental phase, you specify **After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule**. On Tuesday at 9:00 P.M., you want to advance a migration task to the final phase. At the time that you transition the migration task to the final phase, the incremental phase is running. You right-click the migration task and click **Advance to Final Phase**.

- Do not select **Now**. If you select **Now**, the migration task does not transition to the final phase. The policy continues to execute the incremental phase, and you get the error message, "The policy is already in progress" after you click **OK** in the Advance to Final Phase dialog box.
- If you select **After the current Incremental Phase run**, the migration task transitions at the same time that it would transition if you had used the policy default.
- If you select **On the first Final Phase interval after**, the migration tasks transition to the final phase the next time the final phase is scheduled to run after the transition date you specify.

Create a migration policy checklist

To create your first policy, follow the basic steps listed below:

1. Create a migration policy.
2. Define how you want to execute all the migration tasks associated with the policy.
 - Configure the initial phase and schedule the phase.
 - Configure the incremental phase and schedule the phase
 - Configure the final phase and schedule the phase
3. Specify namespace options.
4. Add a migration task to the policy. A migration task identifies the source and destination. A policy can have multiple migration tasks.

Create a migration policy

Create a migration policy to migrate CIFS and NFS data between machines in your network. When you migrate CIFS data referenced from a namespace, your links are updated and the underlying data is migrated. You configure the migration phases you want to run and the actions within each phase that you want to execute for each policy migration task.

1. Perform one of the following actions:
 - On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and click **Migration Policy**.
 - Drag a machine from the **Physical View** to the **Migration Policies** folder, and click **New Policy**.
 - Right-click the **Data Movement Policies** folder, point to **New**, and click **Migration Policy**.
 - Right-click the **Migration Policies** folder and click **New Migration Policy**.
 - Right-click  **Admin View**, click **New**, point to **Data Movement Policy**, and click **Migration Policy**.
2. Type a word or phrase that best describes this policy. The name appears under the **Migration Policies** folder and on the **Properties** tab of items that are members of the policy.
3. In the **Use the Monitor on** list, select the machine that contains the agent that you want to use to process the migration policy. Normally, the current VFM server is used to process the migration 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 to be processed.
4. Select **Run task creation batch** to run a batch file that specifies the source and destination of data transfers. Type the path and name of a batch file or click **Browse** to locate the file in your network. For information about creating a migration tasks file, see "Data migration considerations."
5. Click **OK**.

Notes:

- If you uninstall a monitoring agent, ensure that migration policies that formerly used the uninstalled agent are reassigned to a currently available agent. Do this by selecting the desired agent through the properties dialog box for the policy.
- The task creation batch file runs every time a migration phase schedule runs. For example, if the initial phase of the policy is scheduled to run at 9:00 P.M. on Saturday and the incremental phase is scheduled to run at 10:00 P.M. every Monday through Friday, the task creation batch file will run at 9:00 P.M. on Saturday and 10:00 P.M. every Monday through Friday.
- For information about configuring batch files to run in VFM storage policies, see "Running batch files."
- For information about the example scripts provided by VFM, see "Batch files."

Configure the initial phase

Define how you want to run each migration task associated with the policy. In the first phase of the migration operation, you can run a script, scan for potential problems, and perform a baseline copy of the data. For each action, you can select to pause only if errors occur during the execution of that action, suspend the migration operation after completing the action, or continue the migration operation after completing an action.

1. Right-click the policy and click **Properties**.
2. Click **Initial Phase**.
3. (Optional) Select the **Run batch file** check box and type the name of a script file, or browse for a script file to run before each migration task is executed. VFM provides sample batch files that you can modify. The sample batch files are in the **Examples\Scripts\Migration** directory in the VFM server installation directory.
4. Select the **Scan for potential problems copying the data** check box to check for problems, such as loss of file attributes and permission problems, before data is migrated.
5. Select the **Perform a baseline copy** check box to perform the initial data transfer from the source to the destination.
6. Click **OK**.

Notes:

- When an action is suspended or paused after an error occurs, you must resume the migration operation manually. For instructions, see "User action required."
- For information about configuring batch files to run in VFM storage policies, see "Running batch files."
- For information about the example scripts provided by VFM, see "Batch files."

Schedule the initial phase

Specify the date and time at which you want to start the initial phase of the migration operation.

1. Right-click the policy and click **Properties**.
2. Click **Initial Phase Schedule**.
3. Select the day and time at which you want the initial phase to start.
4. Click **OK**.

Notes:

- If you want to skip the initial phase, clear all the actions.
- The default start date for all phases is the date and time the policy was created.
- For information about scheduling migration phases, see "Scheduling considerations."

Configure the incremental phase

In the second phase of the migration operation, you can copy the data incrementally and specify when you want to advance migration tasks to the final phase.

1. Right-click the policy and click **Properties**.
2. Click **Incremental Phase**.
3. Select the **Copy the data** check box to copy the data in increments after you have performed a baseline copy of the data.

Note: When you do not want the incremental phase to run, clear this option or disable the incremental phase schedule.

4. Specify the condition for advancing policy migration tasks to the final phase.
 - **Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually** - Migration tasks perform incremental copies indefinitely at the intervals specified by the incremental phase schedule. You must advance migration tasks to the final phase manually. For instructions about how to advance migration tasks manually, see "To specify the condition for advancing a migration task to the final phase."
 - **After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule** - Migration tasks advance to the final phase after each migration task runs the incremental phase one time.
 - **End the Incremental Phase for all tasks ...on...and advance the tasks according to the Final Phase schedule** - Migration tasks advance to the final phase on the date and time you specify and run the final phase the next time the final phase runs after the date and time you specify.
5. Click **OK**.

Schedule the incremental phase

Define when you want to start the incremental phase and how often you want to run the phase to copy data.

1. Right-click the policy and click **Properties**.
2. Click **Incremental Phase Schedule**.
3. Select the **Use migration schedule** check box to run the incremental phase. Clear the check box if you do not want to run the phase.
4. Click **New** to configure a schedule. The current day and time appear in the schedule list. To change the start time, select the schedule and click **Advanced**.
5. Change the date and time and click **OK**.
6. Specify how often you want to run the phase.
7. Click **OK**.

Notes:

- If "Use migration schedule" is selected and "Copy the data" in the Incremental Phase dialog box is not selected, the phase will run and the event history will show only a start and end date.
- The default start date for all phases is the day and time the policy was created.
- For information about scheduling migration phases, see "Scheduling considerations."

Configure the final phase

In the last phase of the migration operation, define the actions that you want to execute for each policy migration task. For each action, you can select to pause only if errors occur during the execution of the action, to suspend the migration operation after completing the action, or to continue the migration operation after completing an action.

1. Right-click the policy and click **Properties**.
2. Click **Final Phase**.
3. Select the **Update Namespace** check box when you want to update links targeting the source data. Select how you want to update links. Clear this option if you are migrating data that is not accessed by users through the namespace.
 - **by retargeting links** - The destination replaces the source as a target of a link.
 - **by adding online link targets** - The destination is added as an online target of a link.

Notes:

- The contents of the targets of a link are synchronized when the link is a member of a disaster recovery or replication policy.
- For links that are not members of a disaster recovery or replication policy, you must rely on other mechanisms when possible or you must manually ensure that the contents of the links targets stay synchronized.

- **by adding offline link targets** - The destination is added as an offline target of a link. Use this option when you are preparing links for a disaster recovery policy because only one target of a link can be online.
 - **by removing the link targets** - The link target pointing to the source is deleted.
4. **(CIFS only)** Select **Prevent user connections during the final copy** to keep the source folder shared but restrict user access to it during the migration process. The share is renamed (a \$ sign is appended to the name of the share). This option is available only when you are migrating CIFS data.
 5. Select **Do a final replication prior to deleting the source data** if the data transfer might be interrupted or if it is possible that source files changed since they were last copied to the destination. VFM performs data synchronization before the source data is deleted.
 6. Select **Delete the source data** to delete the source data after the data is migrated to the destination.
 7. Select **Stop sharing the source share or export** to stop sharing the source file or to remove the export. Select this option if you do not want the source data accessible through its former share or export, but you want to retain the data. Only the specified source share or export is unshared. If there are multiple shares or exports of the source directory, the other shares or exports remain in effect.
 8. Select **Run batch file** and type the name of a script file, or browse for a script file to run after each migration task is completed. VFM provides sample batch files that you can modify. The sample batch files are in the `Examples\Scripts\Migration` directory in the VFM server installation directory.
 9. Click **OK**.

Notes:

- When an action is suspended or paused after an error occurs, you must resume the migration operation manually. For instructions, see "User action required."
- For information about configuring batch files to run in VFM storage policies, see "Running batch files."
- For information about the example scripts provided by VFM, see "Batch files."

Schedule the final phase

Define when you want to start the final phase and how often you want to attempt to run the phase. Every policy migration task runs the final phase only one time. The phase will attempt to run based on the schedule you specify.

To schedule the final phase

1. Right-click the policy and click **Properties**.
2. Click **Final Phase Schedule**.
3. Select the **Use migration schedule** check box to run the final phase. Clear the check box if you do not want to run the final phase.
4. Click **New** to configure a schedule. The current day and time appear in the schedule list. To change the start time, select the schedule and click **Advanced**.
5. Change the data and time and click **OK**.
6. Specify how often you want to attempt to run the phase.
7. Click **OK**.

Notes:

- If "Use migration schedule" is selected and no actions are selected for the final phase, the phase will run and the event history will show only a start and end date. No executed actions will appear between the start and end date.
- The default start date for all phases is the date and time the policy was created.
- For information about scheduling migration phases, see "Scheduling considerations."

Specify namespace options

Specify how you want to update the namespace.

How automount files are updated

- The shell credentials for the machines hosting automount files must be set in VFM so that they can be managed. See "Add the shell credentials of a UNIX machine" or "Add the shell credentials of a filer" for instructions.
- During the data transfer process, the entries that referenced the source location are updated to reference the data at the new location on the destination. For example, you want to migrate data from SourceServer to NewServer.
 - **Source** - SourceServer.domain.com:/export/home
 - **Destination** - newServer.domain.com:/export/legacy/home

The automount file contains a line that references the source:

```
/home SourceServer:/export/home/myHome
```

When the migration task runs, the line in the automount file is updated to reference the destination:

```
/home NewServer.domain.com:/export/legacy/home/myHome
```

- VFM automatically runs a script named `namespaceupdate` if it is found in the

directory containing an automount file you are updating. VFM runs the script after the automount file is updated. The script can be used to perform custom processing of your automount files. For example, an updated automount file is named:

```
nis:/var/local/yp/automount.new
```

VFM searches for the namespaceupdate script file in the directory:

```
nis:/var/local/yp/namespaceupdate
```

If the script file is found, VFM executes it and returns the status.

- If the `namespaceupdate` script is not found in the directory containing an updated automount file, VFM runs the command `make` in the same directory as the automount file. `make` searches for a file named `Makefile` and executes the commands in the file. For example, if the automount file is in the directory `nis:/var/local/yp/automount.new`, the command `make` searches for `Makefile` in the directory `nis:/var/local/yp/Makefile`.
- If `Makefile` is not found in the same directory as the automount file, VFM shows an error status.

To specify namespace options

1. Right-click the policy and click **Properties**.
2. Click **Update Namespace Options**.
3. Perform the following actions:
 - If you are transferring CIFS data, add the roots where you will update links.
 - If you are transferring NFS data, you can update references in automount files. Add the automount files you want to update. The automount files must be specified with a UNIX fully qualified path; for example, `<hostname>:/<root>/<subfolder>/<file>`.
4. Click **OK**.

Select replication options

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

Note: To set application-wide default settings for use in creating new policies, 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. When you select **Delete orphaned files in destination folders**, files in the destination targets are deleted if they do not have a corresponding file in the source target. For information about what happens when this option is selected and a policy contains an exclusion list, see the "Note" for Step 15.
7. (CIFS only) In cases where 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, the current destination file is renamed, the temporary file is renamed to the actual file name, and the original destination file is deleted. This minimizes the chance of network errors adversely affecting file contents.
9. (CIFS only) VFM takes 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 saves time 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.
- 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 earlier, 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 you can enable the NTFS Change Journal 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 which files need to be transferred.

10. Select **Use snapshot data** when you want to take a snapshot of or shadow copy the volume in which the source data is located and use the snapshot 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 create, 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. Snapshot data resides on read-only volumes and NTFS Change Journal and file differential replication require that the data be on writable volumes.

11. Select **Override comparison timestamp resolution** when you want to use a one-second resolution to compare 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. Example
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 whose time stamps do not match the time stamps of files in the source target, and copy 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 files in the source target that are not found in the destination targets.
 - **If destination file is missing** - Only copy 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 only 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. Click **OK**.

Specify replication filters

Specify files and directories that you want to explicitly include in or exclude from the migration operation.

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 Filters**.
3. Type directory exclusions delimited by double quotes, or click **Expand** to specify multiple directories to exclude.
4. Type file exclusions delimited by double quotes, or click **Expand** to specify multiple files to exclude.

Note: When **Delete orphaned files in destination folders** is selected in a policy that contains an exclusion list, files that match the exclusion list are not copied and files on the destination that match the exclusion list are not deleted. The following example illustrates what happens to files when a policy contains an exclusion list and the **Delete orphaned files in destination folders** option is selected:

- PolicyA excludes .doc files. Therefore:
 - When PolicyA runs and finds .doc files in the source, the files are not transferred to the destination.
 - When PolicyA finds .doc files in the destination, the files are not deleted.

5. Type directory inclusions delimited by double quotes, or click **Expand** to specify multiple directories to include.
6. Type file inclusions delimited by double quotes, or click **Expand** to specify multiple files to include.
7. Click **OK**.

Notes:

- After you add an item to an exclusion list, specifying it in a subsequent include list does not include the item.
- Exclusions and inclusions are case-sensitive when they list directories and files on UNIX machines.

Select a replication agent

You have the option of selecting which replication agent will perform data migration. You can specify a replication agent on the source, the destination, or a proxy machine, or you can specify an agent group or enable VFM to select which agent to use.

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 **Agent Options**.
3. Select which replication agent will perform data migration.
 - **on the source** - Use the replication agent installed on the source or deploy an agent on the source if an agent is not already installed.
 - **on the destination** - Use the replication agent installed on the destination or deploy an agent on the destination if an agent is not already installed.
 - **on this machine or corresponding proxy** - Use the replication agent installed on the machine you specify or deploy an agent on the machine if an agent is not already installed. When you do not know the name of the proxy that will perform data migration, you can specify the machine to which you assigned a proxy and VFM finds the proxy machine. Use this option when data migration is between two machines on which replication agents cannot be installed, such as NetApp filers. You can type the name of a proxy machine or you can type the name of the filer, and VFM finds the proxy machine you assigned to the filer.
 - **in this agent group** - VFM uses a replication agent from the agent group you select. To create a replication agent group and add replication agents to the group, see "Create an agent group."
 - **Automatically selected** - VFM deploys a replication agent. For an explanation about how VFM automatically distributes replication agents, see "Distributed replication."
4. Click **OK**.

Select differential replication options

Select file differential replication options for the policy. You can send only the parts of a file that change instead of sending the entire file. Replication time and network traffic are decreased considerably.

Notes:

- When data is replicated using differential replication, both the source and the destination must have replication agents installed. VFM attempts to deploy replication agent to both the source and the destination. If VFM cannot deploy a replication agent to either the source or destination, you must designate a proxy.
- When NFS data is replicated using differential replication, the machines hosting the replication agents must be running the same operating system. The source and destination machines or a proxy machine and the source or destination machine must be running Solaris or must be running Linux.
- File differential replication is not recommended for files in compressed format (for example, JPEG, MP3, and MPEG) because a small change to the content of these files causes changes throughout the entire file.
- To set application-wide default replication differencing options for all new policies, click **Options** on the **Tools** menu.
- To specify a proxy machine, see "Specify a proxy machine."

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 Differencing**.
3. Select **Enable file differential replication**.
4. Change the minimum and maximum files sizes or accept the default options. Differencing is performed only on files that meet the size criteria.
5. Type file exclusions delimited by double quotes, or click **Expand** to specify multiple files to exclude.
6. Type file inclusions delimited by double quotes, or click **Expand** to specify multiple files to include.
7. Change how large the data to be transferred can be, or accept the default option. When the data exceeds the limits set, VFM replicates the entire file instead of differencing.
8. Change the differencing block size or accept the default option.
9. Select **Fail replication if differential replication is not available** if you want to cancel the replication of a file that cannot be transferred using differential replication.
10. Select **Use cached information if available** if you want to save hash sets. Some systems cannot take advantage of this option.
11. Click **OK**.

Select file security replication options

There are various options for controlling how a file's security is copied with a file during replication. By default, the security descriptor is copied every time a file is copied.

Notes:

- 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.
- 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 Security**.
3. (CIFS only) You can choose how you want to handle security descriptors when replicating files for the policy. Choose to copy security descriptors:
 - **Each time the file or folder is copied** - Copies security descriptors from source to destination each time the file or folder is copied. This is the default.
 - **Only if target file or folder does not exist** - Preserves security descriptors of files that exist in the destination and copies security descriptors of source files only if the files do not exist in the destination. Use this option if you want to preserve the security setting on files that already exist on the destination.
 - **Never** - Does not copy the security descriptor of source files. Use this option when you set security on the destination and you want the files to inherit the security of the destination rather than having the security match that of the source.
 - **Always copy security descriptor** - Copies the security descriptor from the source file to the destination file whether or not the files are different. When the files are different, the security descriptor and the contents of the files are copied. This option can be useful in scenarios in which file contents are already synchronized between the source and destination but permission changes have been made to source files after their contents have been replicated to the destination.

Note: The number of bytes transferred for files that had only their security descriptor copied (contents were not copied) might not be accurately reflected in the Event Details dialog box.

4. (CIFS only) Select **Process local trustees in security descriptor** to translate or remove security identifiers (SIDs) from security descriptors of files copied from the source. The default option is to translate the SID. If a local group is not found on the destination, a local group is created. If the SID cannot be translated, VFM translates it to the local administrators group in the destination. When transferring data to a domain controller, you must clear the **Process local trustee in security descriptor** check box. Domain controllers only have domain SIDs, and the processing of local SIDs is not applicable.
 - **Remove the SID from the security descriptor** - Deletes access control entries containing local SIDs from the destination copy of the file. The source files are not modified. Use this option when you do not want to migrate local security settings from the source to the destination. If the owner or primary group of a file or directory being transferred corresponds to a local SID, it is translated to the local administrators group in the destination because the owner and group IDs cannot be removed.
 - **Translate the SID to the same trustee name and type on the target** - Translates local SID in the source files to a local SID in the destination that corresponds to a trustee of the same name and type. Use this option to preserve local security settings in the destination.
5. (CIFS only) Select **Create local groups on the target if not found** if you have not already created local security groups on the destination. This is the default if **Translate the SID to the same trustee name and type on the target** is selected.
6. If the SID cannot be translated, select one of the following options:
 - **Leave it as it is** - Copies the source files without modifying the SID. Use this option to preserve the access control entry (ACE) even though the SID cannot be translated.
 - **Remove the SID from the security descriptor** - Deletes ACEs containing local SIDs. Use this option when the SIDs cannot be translated. If the owner or primary group of a file or directory corresponds to a local SID, it is translated to the local administrators group in the destination because the owner ID cannot be removed.
 - **Translate to local administrators group** - Replaces the SID read from the source files with the SID of the local administrators group in the destination. Use this option when you want the SID to be meaningful in the destination. This is the default if the SID cannot be translated. Security settings on the source files are not modified.

7. Choose whether you want to allow loss of security information.
 - (CIFS only) In cases where the destination targets do not support security information supported in the source target (for example, when you are copying from NTFS to FAT), you can select **Allow loss of security information**. This copies files to the destination targets that lose security information.
 - (NFS only) Whether this option is or is not selected, the contents of the source target are copied to the destination targets. When this option is selected and the security information of the source target cannot be copied to the destination targets, error messages are suppressed. When this option is not selected and the security information of the source target cannot be copied to the destination targets, you see error messages.
8. Click **OK**.

Select file attribute replication options

Select file attributes replication options for a policy.

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 **Attribute Replication**.
3. Determine whether the source and destination files have changed. You can select to **compare** files when you specify **Set**, **Clear**, or **Unchanged** as the action to perform on the destination files.
 - When **Compare Attribute** is selected and the source and destination files are the same, but the attributes of the source and destination files are different, the source is copied and the task you specified for the destination file is performed.
 - When **Compare Attribute** is not selected and the source and destination files are the same, but the attributes of the source and destination files are different, the source file is not copied.
4. Specify how you want to configure the attribute of destination files. For each attribute listed, select one of the following actions:
 - Preserve the attribute of the source files (**Match**).
 - Turn on the attribute of the destination files (**Set**).
 - Turn off the attribute of the destination files (**Clear**).
 - Do not change the destination files (**Unchanged**).

Note: When you are transferring NFS data and you clear the Read-Only attribute, write permission is granted to users and groups.
5. The replication is performed and the last access time of the source file is set back to what it was prior to the replication operation. When **Preserve last access time** is not selected, the replication operation causes the last accessed time stamp on the source target to be updated. Selecting this option incurs extra processing and slows replication performance.

Note: When you migrate data from Novell sources, you must clear the **Preserve last access time** option.

6. Specify the retention information for files copied to SnapLock volumes on NetApp filers. You can specify a time and date or the length of time files are to be locked after they are copied.

Notes:

- When you copy files to be locked on a SnapLock volume, you must **Set** the Read-Only attribute.
- For more information about SnapLock volumes, see the *Data ONTAP® 6.5 Data Protection Online Backup and Recovery Guide* (Part number 210-00476).

7. Click **OK**.

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

Add a migration task

Specify the source and destination of data transfers in a migration task. A migration task can contain only CIFS source and destinations or only NFS source and destinations, but not both protocols. You can associate multiple tasks with a given policy.

Notes:

CIFS

- 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.
- CIFS physical paths must be UNC paths; for example, `\\<machine>\<share>`.
- When you want to update the namespace after a migration, make sure that the UNC paths of the source and the link targets you want to update match. If the UNC paths do not match, the namespace cannot be updated. Example

This example illustrates what happens when the UNC path of a migration source and link targets that match the search criteria do not match.

- You create a migration task and specify a source and destination.

Source	Destination
\\Saturn.npb.com\c\$\Sales\Reports	\\Pluto.npb.com\Sales

- You scan for migration candidates and find one candidate.

Migration candidate
\\Saturn.npb.com\c\$\Sales\Reports\HoustonRegion

- The namespace contains a link named Sales that targets the migration candidate. Notice the difference in the UNC paths between the migration candidate and the link target.

Namespace link	Target of link
 Sales	\\Saturn.npb.com\Sales\Reports\HoustonRegion

- You run the policy to start the migration. VFM successfully migrates the data to the destination but does not update the link target because the UNC paths of the migration candidate and the link target do not match.

Destination after migration
\\Pluto.npb.com\Sales\HoustonRegion

Namespace link	Target of the link after migration is not updated to target the destination
 Sales	\\Saturn.npb.com\Sales\Reports\HoustonRegion

- After migration, the target for the link Sales does not contain any data because the data was migrated to the destination.
- To avoid this problem, the source path should match the link target path.

Correct source path
\\Saturn.npb.com\Sales\Reports

Migration candidate
\\Saturn.npb.com\Sales\Reports\HoustonRegion

Namespace link	After migration, the target of the link is updated to target the destination
 Sales	\\Pluto.npb.com\Sales\HoustonRegion

NFS

- The source and destination must be UNIX fully qualified paths (for example, *<hostname>:/<root>/<subfolder>/<file>*).
- NFS 3 or higher must be running on the source and destination.
- A replication agent must be installed on the source or destination or you can specify a default proxy.
 - For instructions about how to specify default Windows and UNIX proxy machines, see "Specify default replication agents."
 - For instructions about how to specify a Windows or UNIX proxy machine, see "Specify a proxy machine."
- During the data transfer process, the entries that referenced the source location are updated to reference the data at the new location on the destination. For example, you want to migrate data from *SourceServer* to *NewServer*

- **Source** - SourceServer.domain.com:/export/home
- **Destination** - newServer.domain.com:/export/legacy/home

The automount file contains a line that references the source.

```
/home SourceServer:/export/home/myHome
```

When the migration task runs, the line in the automount file is updated to reference the destination.

```
/home NewServer.domain.com:/export/legacy/home/myHome
```

- VFM automatically runs a script named **namespaceupdate** if it is found in the directory containing an automount file you are updating. VFM runs the script after the automount file is updated. The script can be used to perform custom processing of your automount files. For example, an updated automount file is named `nis:/var/local/yp/automount.new`, VFM searches for the **namespaceupdate** file in the directory `nis:/var/local/yp/namespaceupdate`. If the file is found, VFM executes it and returns the status.
 - If **namespaceupdate** is not found in the directory containing an updated automount file, VFM runs the command **make** in the same directory as the automount file. **make** searches for a file named **Makefile** and executes the commands in the file. For example, if the automount file is located in the directory `nis:/var/local/yp/automount.new`, the command **make** searches for a **Makefile** file in the directory `nis:/var/local/yp/Makefile`.
 - If **Makefile** is not found in the same directory as the automount file, VFM returns an error status.

To add a migration task

1. Perform one of the following actions:
 - Drag a share (UNC path) from Windows Explorer to a migration policy.
 - Right-click a migration policy and click **Add task to policy**.
 - Drag a machine, share, or export from the **Physical View** to a migration policy.
 - Drag a machine, share, or export from the **Physical View** to the **Migration Policies** folder. Select an existing policy or create a new policy.
 - Right-click a machine, share, or export in the **Physical View**, then click **New Migration Task**. Select an existing policy or create a new policy.
 - Select a machine in the **Physical View** to display a **Properties** tab in the right pane and drag a share or export to a migration policy.
 - Select a link in the **Logical View** to display a **Properties** tab in the right pane and drag a share under **Target** to a migration policy.
2. Type the source path in **Original Path** or click **Browse** to locate a shared folder, export, or machine.
3. Type the destination path or click **Browse** to locate a shared folder, export, or machine.
4. Click **Next**. If the source path is a machine, you can select whether to migrate all shares, migrate only shares referenced by the namespace, or migrate all exports. To continue, click **Next**.
5. Review the list of source and destinations.

- Clear the items that you do not want to include in the migration.
 - To change a destination path, click the destination path and type a new path or click **Browse**  to locate a new share or export.
 - (CIFS only) Clear **Warn if destination already contains files or folders** if you do not want to check for destination shares that are not empty.
6. Click **Next**. VFM checks for duplicate paths in other tasks. If a match is found, a dialog box lists the name of the policy containing the duplicate path and highlights the path. Click **Back** to make changes.
 7. Click **Next**. VFM checks for empty destinations (CIFS only) and destinations that do not exist (CIFS and NFS).
 8. If destinations that are not empty are found, the items appear in a list. You can continue or select each destination and type a different destination, browse for another destination, or remove the destination from the migration process.
 9. If destinations are found that do not exist, the destinations appear in a list with a suggested local path for each destination.
 - (CIFS only) To change the local path for multiple items, select the items and click **Rebase**. Type your changes and click **OK**. To change the local path for one item, select the item and type the new local path.
 - (NFS only) To change an export point, select it and type a new local path. The export point must be the same as or a parent of the destination export. The following example shows acceptable export points.

Destination export	Acceptable Export Points
moria:/home/directory1/directory2	/home /home/directory1 /home/directory1/directory2

- To accept all or some of the suggested local paths, select the items on the list and click **Create Destination Share (CIFS only)** or click **Create Destination Export (NFS only)**.

Notes:

- When you create an export on a NetApp filer, the path must begin with `/vol/` followed by a valid name of a volume (for example, `/vol/vol1/<directory>` or `/vol/vol0/<directory>/<subfolder>`). The volume must not already be exported.
- If part of the destination export path is already exported, the new export path is created under the current export and no new export is created.
- When you create an export on a UNIX machine, all the path elements specified are created.
- The permissions on the source are copied to the destination.
 - If data transfer is between different platforms (for example, Linux, Solaris, or a filer), VFM translates the applicable security permissions for the new export.
 - If the directory already exists, the source permissions are copied to the destination. The ability to set UID and GID permissions must be enabled on the NFS volume on which you want to create the export. The VFM server runs with NFS credentials of `root`. The export must be configured to enable the machine hosting VFM to log in as `root`. If you are overriding the NFS credentials set for VFM, you must be the current owner of the destination directory.
- To create exports on the destination machine, you must set shell credentials for UNIX machines and filers participating in data transfer. See "Change the shell credentials of a filer" or "Change the shell credentials of a UNIX machine" for instructions about how to set shell credential in VFM.

10. Click **Next**.
11. A list displays of source and destination paths that are validated. You can clear items that you do not want to migrate, or you can change the destination path of an item.
12. Click **Next**.
13. (CIFS only) If the policy has the action to update the namespace, a list of links to be updated appears. Clear any links that you do not want to update.
14. Select when you want the migration task to advance to the final phase.
 - **Never; Remain in Incremental Phase** - The migration task repeats the incremental phase indefinitely. You must manually advance the migration task to the final phase. For instructions about how to manually advance the migration task, see "User action required."
 - **After the current Incremental Phase run** - The migration task runs the incremental phase one time and advances to the final phase. The next time the final phase runs, the migration task executes the actions specified.
 - **On the first Final Phase interval after** - The migration task advances to the final phase the next time the final phase runs after the date and time you specify.
 - **According to the policy default** - The migration task advances to the final phase based on the settings you specified for the policy. For instructions about how to

configure the cutover date for the policy, see "Advance to final phase."

15. Click **Finish**.

Note: To monitor the progress of a migration task, see "View progress of a task."

Advance to final phase

Specify when you want each policy migration task  to advance to the final phase.

- Incremental copy of the data is performed for each migration task until the condition for advancing a migration task to the final phase is met. When the condition is met, a migration task advances to the final phase and executes the actions specified in the final phase the next time the final phase runs.
- The conditions for advancing to the final phase can be configured at the policy or migration task level.
- Before you specify the conditions for advancing migration tasks to the final phase, see "Scheduling considerations."

To specify the conditions for advancing to the final phase at the policy level

The conditions you specify for advancing to the final phase affect all migration tasks associated with the policy.

1. Double-click **Admin View**.
2. Double-click the **Data Movement Policies** folder.
3. Double-click the **Migration Policies** folder.
4. Right-click the policy you want to specify when you want to advance to the final phase and click **Properties**.
5. Click **Incremental Phase**.
6. Select one of the following options:
 - **Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually** - Migration tasks perform incremental copies indefinitely at the intervals specified by the incremental phase schedule. You must advance migration tasks to the final phase manually. For instructions about how to advance migration tasks manually, see "To specify the condition for advancing a migration task to the final phase."
 - **After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule** - Migration tasks advance to the final phase after each migration task runs the incremental phase one time.
 - **End the Incremental Phase for all tasks ... on...and advance the tasks according to the Final Phase schedule** - Migration tasks advance to the final phase on the day and time you specify. The next time the final phase runs after the date you specify, the migration tasks execute the actions specified in the final phase.
7. Click **OK**.

To specify the condition for advancing a migration task to the final phase

The condition for advancing to the final phase that you specify for a migration task overrides the condition that you specify at the policy level.

1. Perform one of the following actions:
 - Right-click a migration task appearing under a migration policy and click **Advance to Final Phase**.
 - Select a migration policy in the **Admin View** to show a **Properties** tab in the right pane. Under **Tasks**, right-click a migration task and click **Advance to Final Phase**.
2. Select one of the following options:
 - **Now** - Manually advance the migration task to the final phase immediately if the incremental phase is not running. If the incremental phase is running when you select this option, the policy continues to execute the incremental phase and you get an error message. For more information, see "Scheduling considerations."
 - **Never; Remain in Incremental Phase** - The migration task runs the Incremental phase indefinitely. You must manually advance the migration task to the final phase.
 - **After the current Incremental Phase run** - After the migration task performs the actions specified in the incremental phase one time, the migration task advances to the final phase. The next time the final phase runs, the migration task executes the actions specified in the final phase.
 - **On the first Final Phase interval after** - Select a date and time after which you want the migration task to advance to the final phase. The next time the final phase runs, the migration task executes the actions specified in the final phase.
 - **According to the policy default** - The migration task uses the option specified for the policy.
3. Click **OK**.

Notes:

To specify the condition for advancing to the final phase for multiple migration tasks, perform the following actions:

1. Select a migration policy in the **Admin View** to show a **Properties** tab in the right pane. Under **Tasks**, select multiple migration tasks.
2. Right-click the selected tasks and click **Advance to Final Phase**.
3. Continue with Step 2 in "To specify the conditions for advancing to the final phase for a migration task."

To view the condition for advancing to the final phase for a migration task, perform one of the following actions:

- Click a migration task appearing under a policy. The information appears in the **Properties** tab in the right pane.
- Right-click a migration task and click **Properties**.

Specify migration tasks from a file

Specify the source and destination of data transfers from a batch file.

1. Perform one of the following actions:
 - Right-click a migration policy and click **Properties**.
 - Select a migration policy to show a Properties tab in the right pane. Click  **Edit**.
2. Click **General** and select **Run task creation batch**.
3. Type the path and name of a batch file or click **Browse** to find a file in your network.
4. Click **OK**.

Note: See "Data migration considerations" for information about creating a migration tasks file.

Run customized commands during replication

You have the option of running batch files before and after a replication action during the migration operation.

1. Perform one of the following actions:
 - Right-click a migration policy and click **Properties**.
 - Select a migration 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 specify the batch files you want to run. 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 VFM monitoring is taking place
4. Click **OK**.

Notes:

- (Migration, Archival Migration, and Storage Load Balancing policies only) When you specify scripts to run before and after replication, the files run before and after the following migration actions:
 - Scan for potential problems copying the data (Initial Phase).
 - Perform a baseline copy (Initial Phase).
 - Copy the data (Incremental Phase).
 - Do a final replication prior to deleting the source data (Final Phase).
 - Delete the source data (Final Phase).
- For information about configuring batch files to run in VFM storage policies, see "Running batch files."
- For information about the example scripts provided by VFM, see "Batch files."

Update the properties of a migration task

Update the properties of a migration task before it starts. You can change the source and destination and determine which links targeting the source you want to update.

1. Right-click a  scheduled task and click **Properties**.
2. Update the task properties.
3. Click **OK**.

Start a migration task immediately

You can bypass the replication schedule and start a migration task manually.

Perform one of the following actions:

- Open a folder under a migration policy, double-click a machine, right-click a task, and click **Start task now**.
- Select a migration policy to show a **Properties** tab in the right pane. Under **Tasks**, right-click a task and click **Start task now**.

Start all tasks for a policy immediately

You can bypass the replication schedule and manually start all the tasks associated with a policy.

Perform one of the following actions:

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

View progress of a task

After you start a migration task, VFM provides detailed information about the progress of the task.

1. Select the task to display tabs in the right pane.
2. Click the **History** tab.
3. Click **Show Details** to view detailed information about the progress of a task.

Note: To create a report of a task history, click  **Save Report**. The report appears under **Reports** in the tree.

User action required

When a task  requires user action, a list of actions from which you can choose is provided. To execute an action, perform one of the following steps:

- Right-click a migration task, point to **User actions** and select the action you want to perform.
- Select a migration task to display tabs in the right pane. Click the **History** tab. Under **User action is required...**, click an action. The results of the action appear on the page.
- Right-click a migration task and click **Suspend task now**.

The following table lists the actions and provides an explanation for each action.

Action	Explanation
Go back to the first action <u>now</u> .	Restart the migration task from the beginning immediately, starting with the first checked action for the migration task. If you make changes to the selected actions for the migration task then use this option, the changed settings are used on the reprocessing of the migration task.
Go back to the first action <u>the next time the policy runs</u> .	Restart the migration task from the beginning when the task is scheduled to run, starting with the first checked action for the migration task. If you make changes to the selected actions for the migration task then use this option, the changed settings are used on the reprocessing of the migration task.
Retry the last action <u>now</u> .	Start the last action immediately.
Retry the last action <u>the next time the policy runs</u> .	Restart the migration task at the last action completed when the task is scheduled to run.
Continue with the next action <u>now</u> .	Start the next action immediately.
Continue with the next action <u>the next time the policy runs</u> .	Restart the migration task with the next action when the task is scheduled to run.
Suspend task <u>now</u> .	Cancel the current action and stop the migration process for the task. A suspended task appears with no status  . To take action on a suspended task, select the task to display tabs in the right pane. Click the History tab and select an action.

Suspend a migration task

Interrupt the current action and stop the migration process. A suspended task appears with no status .

Perform one of the following actions:

- Right-click a task and click **Suspend task now**.
- Select a policy folder to display a **Properties** tab in the right pane. Under **Tasks**, right-click a task and click **Suspend task now**.

Note: To resume or take another action on a suspended task, select the task to display tabs in the right pane. Click the **History** tab and select an action.

Remove a task

When a task is removed from a policy, the task history is also deleted. If you want to keep this information, select the task to display a **History** tab in the right pane. Print or save the information as a report.

1. Perform one of the following actions:
 - Right-click a migration task and click **Remove from Policy**.
 - Select the policy to show a **Properties** tab in the right pane. Under **Task**, right-click a task and click **Remove from Policy**.
2. Click **Yes**.

Delete a migration policy

When you delete a migration policy, the machines, shares, and exports that are members of the policy are not deleted from the namespace.

1. Perform one of the following actions:
 - Right-click the policy and click **Delete**.
 - Select the **Migration Policies** folder to view a list of policies in the right pane. Right-click a policy and click **Delete**.
2. Click **Yes**.

Note: If a migration task is in progress when you delete the policy, the migration task runs until it is complete.

Chapter 3 Replication

Replication policies enable you to transfer CIFS and NFS data. In addition to the data, you can transfer attributes and permissions.

Before you begin creating a replication policy, review the information in "Replication planning considerations." The topic covers detailed information, such as the permissions the account used by the VFM server and monitoring agent must have.

Create a replication policy checklist

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

1. Create a replication policy.
2. Add or replace the source.
3. Add or replace the target.
4. Optionally change the data distribution path.
5. Schedule data distribution.

Create a replication policy

Add a data distribution policy and centrally manage the policy from the Replication Policies folder in the Admin View. The VFM replication agent supports replication across heterogeneous Windows and UNIX devices.

1. Perform one of the following actions:
 - On the **Tools** menu, click **New Policy**, point to **Data Movement Policies**, and click **Replication Policy**.
 - Drag a link to the Replication Policies folder.
 - Right-click the **Data Movement Policies** folder, point to **New**, and click **Replication Policy**.
 - Right-click the **Replication Policies** folder and click **New Replication Policy**.
 - Right-click  **Admin View**, click **New**, point to **Data Movement Policy**, and click **Replication Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field. The name appears under the Replication Policies folder and on the **Properties** tab of items that are members of the policy.
3. In the **Source** field, type a CIFS or NFS physical path or the name of a  link that references a single CIFS target, or click **Browse** to locate the source. If the source is a single target link , it must be writable. The source can be a member of more than one replication policy.
4. In the **Target** field, type CIFS physical paths and/or links that reference CIFS targets or NFS physical paths. A target  link can reference multiple targets. However, a target can be a member of only one replication policy.

5. You can select another location where operation of the replication policy is performed. The default is to run from the VFM server. From the **Use the Monitor on** list, select the hostname of another monitoring agent.
6. Click **OK** to save the policy.

Notes:

- Data transfer must occur within the same protocol (for example, CIFS to CIFS or NFS to NFS).
- A policy can contain only CIFS source and destinations or only NFS source and destinations, but not both protocols.
- NFS physical paths must be UNIX fully qualified paths (for example, *<machine>:/<UNIX local path>*).
- CIFS physical paths must be UNC paths (for example, *\\<machine>\<share>*).
- **(CIFS only)** It is recommended that you add only links that have shared folders as their targets to replication policies. Due to replication agent location considerations, it is not recommended that you add links to a replication policy that has a DFS root or another link as its target.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files.
- If you uninstall a monitoring agent, ensure that replication policies that formerly used the uninstalled agent are reassigned to a currently available agent.

Select replication options

Configure replication, file security, and replication differencing options. 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

Note: 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.

Schedule data distribution

Select how and when you want to distribute content automatically from remote machines across a network.

- 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 or throttling 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 distribution

1. Right-click a policy and click **Properties**.
2. Click **Replication Schedule**.
3. Click **Use replication schedule** and configure the schedule.
5. Optionally, select **Use alternate schedule when the policy fails**, and type a frequency (in minutes) to run the schedule any time replication fails.
6. Click **OK**.

Notes:

- You can bypass the replication schedule and start the replication process manually whenever you choose. Right-click a replication policy and click **Start Replication Now**.
- The tool tips for Replication policies in the tree view display the time that a replication last occurred and the location of the agent handling the replication.

Change data distribution path

Reconfigure a data distribution path to move data more effectively to targets that are geographically close together. When you create a replication policy, data distribution is configured to be from the source target to each destination target. You can change the configuration to create a cascaded replication topology.

1. Select a replication policy.
2. Click the **Replication Topology** tab that appears in the right pane.
3. Drag a source target to the intended destination or child target to transfer data from the dragged target to the destination or child target.
4. Repeat Step 3 until you configure the data distribution the way you want it.
5. Click **Apply Changes**.

Change replication for a target

Normally all targets use the VFM replication agent. You can choose to use another replication.

1. Select a policy.
2. Click the **Replication Topology** tab that appears in the right pane.
3. Right-click a target path and click **Use other replication**.
4. Click **Apply Changes**.

Notes:

- Replication status monitoring is available only for targets using the VFM replication agent.
- If a source target is using another replication, all targets of the source must use the same replication. To change back to the VFM replication agent, you must first change the replication of the source target.

Add or replace the source

Easily add a source to a replication policy or replace the existing source of a policy.

1. Perform one of the following actions:
 - Drag a  single target link from the Logical View, a shared folder from the Physical View, a shared folder (UNC path) from Windows Explorer, or an export from the Physical View. Select **As the source path**.
 - If the policy did not have a source, the newly-added item becomes the source.
 - If the policy already had a source, the newly-added item replaces the current source.
 - Drag a shared folder from the Physical View or a shared folder (UNC path) from Windows Explorer. Select **Replace the existing source link's target**. The shared folder replaces the share of the policy's source link.
2. Click **OK**.

Notes:

- 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.
- Links that are already members of disaster recovery policies cannot be members of a replication policy because links contained in disaster recovery policies can have only one active target.
- An item that is already a destination in a replication policy cannot be a source of another replication policy.
- It is recommended that you add a link that has a shared folder as its target to a replication policy. Due to replication agent location considerations, it is not recommended that you add a link to a replication policy that has a DFS root or another link as its target.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files.
- For more information, see the "Notes" in "Create a replication policy."

Add or replace the target

Easily add a target to a replication policy or replace the existing target of a policy.

1. Perform one of the following actions:

- Drag a  single target link or  replicated link from the Logical View to a replication policy. Select **Add it to the list of destination paths** to add a new target or select **Replace existing destination path** and select the destination in the policy you want to replace from the list.
- Drag a shared folder from the Physical View or a shared folder (UNC path) from Windows Explorer to the replication policy. Select **Add it to the list of destination paths** to add a new target or select **As an additional target for an existing destination link** and select the destination link in the policy to which you want to add the share.
- Drag an export from the Physical View to the replication policy. Select **Add it to the list of destination paths** to add a new target.

2. Click **OK**.

Notes:

- Links that are already members of disaster recovery policies cannot be members of a replication policy because links contained in disaster recovery policies can have only one active target.
- A link that is already a destination link in a replication policy cannot be a destination link of another replication policy.
- It is recommended that you add only links to a replication policy that have shared folders as their targets. Due to replication agent location considerations, it is not recommended that you add links to a replication policy that has a DFS root or another link as its target.
- When configuring replication, do not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files.

Change the replication schedule

Change the time and date of a schedule or add another schedule.

1. Right-click a policy and click **Properties**.
2. Click **Replication Schedule**.
3. Update an existing schedule or add a new schedule.
4. Click **OK**.

Start replication immediately

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

Perform one of the following actions:

- Right-click a replication policy and click **Start Replication Now**.
- Select a replication policy to show tabs in the right pane. Right-click anywhere on the **Properties** or **Status** tab 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 on the right pane and click .

Cancel a replication in progress

Perform one of the following actions to stop a data distribution operation:

- Right-click a replication policy and click **Cancel Replication In Progress**.
- Select a replication 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 .

View information about replication operations

View information about replication operations configured in VFM. 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 **Replication Manager Policies** folder in the Admin View 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.

Compare contents of targets in a replication policy

Compare the contents of the source with the contents of each of the targets in the destination contained in a replication policy.

1. Right-click a replication policy, point to **Reports**, point to **Physical**, and click **File Differences**. The source target appears in the **Compare** field and the destination targets appear in the **To** field.
2. Click **Next**.
3. Type a name for your report and select how you want to save it.
4. Click **Finish**.

View properties of a replication policy

View the properties of a replication policy.

Perform one of the following actions:

- Right-click a policy and click **Properties**.
- Select a policy. A **Properties** tab appears on the right pane.

Delete a destination target

When you delete a target, it is removed from the replication policy. If the target is referenced by a link, it is also removed from the logical namespace. The target is not deleted from the machine on which it resides.

1. Select a policy and click the **Replication Topology** tab in the right pane.
2. Right-click a folder and click **Delete**.
3. Click **Yes**. A status report appears.

Delete a replication policy

Deleting a replication policy does not delete the targets that are members of the policy.

1. Perform one of the following actions:
 - Right-click the policy and click **Delete**.
 - Select the Replication Policies folder to view a list of policies in the right pane. Right-click a policy and click **Delete**.
2. Click **Yes**.
3. Click **Done**.

Appendix A Specifying migration tasks from a file

Migration policies provide the option of running a batch file to add migration tasks to a policy. You can specify the batch file in the **Run task creation batch** field on the **General** tab of a data migration policy properties dialog box.

How it works

- Each time the migration policy runs, it invokes the task creation batch file.
- The task creation batch file creates a task description file.
- After the batch file exits, the VFM server performs the following actions:
 - Reads the task description file created by the `.bat` file.
 - Adds the tasks described in the task description file to the policy.
 - Deletes the task description file.

This enables you to schedule repeated transfers of data for the locations described in the tasks added by the batch file.

The task creation batch file

VFM invokes the task creation batch file with the following arguments:

`<policy name> <taskfile name> [optional user-specified arguments]`

The `<policy name>` parameter is the name of the data migration policy that is invoking the task creation batch file. The `<taskfile name>` parameter specifies the name and location of a file to which the task creation batch file can write a list of migration tasks for VFM to process. VFM does not create the file. The batch job creates the task file in the task file format described below. VFM processes the contents of the `<taskfile name>` and processes data migration tasks described in it. After the task is complete, VFM deletes the `<taskfile name>`.

For example, if you specify `TaskCreation.bat` in `MyMigrationPolicy`, VFM invokes your batch file as shown below:

```
TaskCreation.bat MyMigrationPolicy c:\temp\taskfile0001.xml
```

You can also specify additional arguments to the batch file. If you specify the following in `MyMigrationPolicy`:

```
ComplexTaskCreation.bat arg1 arg2 arg3
```

VFM invokes your batch file as:

```
ComplexTaskCreation.bat MyMigrationPolicy  
c:\temp\taskfile1234.xml [arg1] [arg2] [arg3]
```

Task File Description Syntax

- `DFSMIGRATION` (element): required
 - Version: required, must be set to "1."
- `Task` (element): optional (though nothing will happen if you don't have at least one of these)

- **CurrentTarget** (attribute): required
 - Must be a valid UNC path with at least the \\server\share portions of the path specified.
 - Cannot be a DFS path.
 - Cannot be the same as **NewTarget**.
- **NewTarget** (attribute): required
 - Must be a valid UNC path with at least the \\server\share portions of the path specified.
 - Cannot be a DFS path.
 - Cannot be the same as **CurrentTarget**.
- **AutoDelete** (attribute): optional
 - If set to "1", the task is automatically deleted from the policy if it is completed successfully.
 - If not present or set to anything other than "1," the task is not automatically deleted.
- **CutoverDate** (attribute): optional
 - If set to "", the task setting for Advance to Final Phase is "Never".
 - If set to "NextRun," the task setting for Advance to Final Phase is "Next Policy Run."
 - If set to a date in the format "mm/dd/yyyy hh:mm:ss", the task setting for Advance to Final Phase is "Next Policy Run After" the date specified.
 - If the attribute is not present, the task setting for Advance to Final Phase is "Use the Policy Default."
- **Link** (element): optional
 - **Update DFS** must be selected. Click the **Actions** tab of the policy properties to verify that this action is selected.
 - Only links listed in the task file are updated when the policy runs.
 - **LinkName** (attribute): required
 - The full path of a link to be retargeted.
 - Must be a valid UNC path with at least the \\server\share\path portions specified (Note that specifying just \\server\share is invalid).
 - **CurrentTarget** (attribute): optional
 - Indicates the current replica of the link that is to be retargeted.
 - If not present, default value is the value of Task>CurrentTarget attribute.
 - If present, must indicate the same folder or a subfolder of the path specified in the Task>CurrentTarget attribute.

Appendix B Example Task Creation batch file

The following example shows a task creation batch file that creates the task description file illustrated in Appendix C.

```
@echo off
rem
rem *****
rem *****
rem FILE: CreateTasks.bat
rem
rem PURPOSE:
rem     Illustrate creating a task-description file
rem
rem
rem INPUT:
rem     %1 Name of the migration policy calling this BAT file.
rem     This is supplied from VFM.
rem     %2 Name of the output xml file. This is supplied from
rem     VFM.
rem
rem COPYRIGHT:
rem     Copyright (c) 2007 Brocade Communications Systems, Inc.
rem     All rights reserved.
rem
rem
rem *****/
rem
rem Next two lines set up the start of the XML file to be handed
rem to VFM
rem
echo ^<?xml version="1.0" encoding="UTF-8" ?^> > ExampleTasks.xml
echo ^<DFSMIGRATION Version="1"^> >> ExampleTasks.xml
rem Add first task
echo ^<Task CurrentTarget=^"\SERVERA\Test\Src1^"
NewTarget=^"\SERVERB\Test\Dst1^" ^/^> >> ExampleTasks.xml
rem Add second task
echo ^<Task CurrentTarget=^"\SERVERA\Test\Src2^"
NewTarget=^"\SERVERB\Test\Dst2^" AutoDelete=^"1^" ^/^> >>
ExampleTasks.xml
rem Add third task
echo ^<Task CurrentTarget=^"\SERVERA\Test\Src3^"
NewTarget=^"\SERVERB\Test\Dst3^" AutoDelete=^"0^" ^/^> >>
```

```

ExampleTasks.xml
rem Add fourth task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src4^"
NewTarget=^"\\SERVERB\Test\Dst4^" AutoDelete=^"1"^> >>
ExampleTasks.xml
echo ^<Link LinkName=^"\\ROOTSERVER\Root\Link^" ^/^> >>
ExampleTasks.xml
echo ^<Link LinkName=^"\\ROOTSERVER\Root\OtherLink^" >>
ExampleTasks.xml
echo CurrentTarget=^"\\SERVERA\Test\Src4\Subdir^" ^/^> >>
ExampleTasks.xml
echo ^</Task^> >> ExampleTasks.xml
rem
rem Terminate the XML file and copy it to where VFM told us to
put it
rem
echo ^</DFSMIGRATION^> >> ExampleTasks.xml
copy ExampleTasks.xml %2
rem
rem Delete our temporary file
rem
del ExampleTasks.xml

exit 0

```

Appendix C Example Task Description file

The following example shows the task description file created by the batch file illustrated in Appendix B that adds four tasks to a migration policy.

- The first task migrates data from \\SERVERA\Test\Src1 to \\SERVERB\Test\Dst1 and leaves the task in the policy after it is completed successfully.
- The second task migrates data from \\SERVERA\Test\Src2 to \\SERVERB\Test\Dst2; the task is automatically removed from the migration policy after the task is completed successfully.
- The third task migrates data from \\SERVERA\Test\Src3 to \\SERVERB\Test\Dst3 and explicitly specifies that the task is not to be removed from the policy after it is completed successfully.
- The fourth task migrates data from \\SERVERA\Test\Src4 to \\SERVERB\Test\Dst4, retargets \\ROOTSERVER\Root\Link to reference the data at its new location and retargets the link replica of \\ROOTSERVER\Root\OtherLink currently targeted at \\SERVERA\Test\Src4\Subdir to reference beneath the new location.

```
<?xml version="1.0" encoding="UTF-8" ?>
  <DFSMIGRATION Version="1">
    <Task CurrentTarget="\\SERVERA\Test\Src1"
NewTarget="\\SERVERB\Test\Dst1" />
    <Task CurrentTarget="\\SERVERA\Test\Src2"
NewTarget="\\SERVERB\Test\Dst2" AutoDelete="1" />
    <Task CurrentTarget="\\SERVERA\Test\Src3"
NewTarget="\\SERVERB\Test\Dst3" AutoDelete="0" />
    - <Task CurrentTarget="\\SERVERA\Test\Src4"
NewTarget="\\SERVERB\Test\Dst4" AutoDelete="1">
      <Link LinkName="\\ROOTSERVER\Root\Link" />
      <Link LinkName="\\ROOTSERVER\Root\OtherLink"
CurrentTarget="\\SERVERA\Test\Src4\Subdir" />
    </Task>
  </DFSMIGRATION>
```


Appendix D Documentation

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

No.	Title	File name	Contents
1.	<i>Advanced Configuration and Techniques</i>	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.	<i>Create and Manage Namespaces</i>	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.	<i>Getting Started</i>	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.	<i>Make Data Highly Available</i>	VFMMakeDataAvailable.pdf	Explains how to make data highly available by creating client, filer, and server disaster recovery policies
5.	<i>Manage Physical Resources</i>	VFMPhysicalResources.pdf	Explains how to use VFM to Work with NetApp filers, set replication properties, and work with Windows and UNIX machines
6.	<i>Manage Snapshots on Filers</i>	VFMSnapshots.pdf	Explains how to use VFM to manage snapshots on NetApp filers

No.	Title	File name	Contents
			using snapshot access and snapshot scheduling policies
7.	<i>Optimize Storage</i>	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.	<i>Reports</i>	VFMReports.pdf	Explains how to create and publish reports
9.	<i>Transfer Data</i>	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.	<i>Troubleshooting</i>	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 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

Add a replication policy	45
Add a source link to a replication policy	48
Add a target link to a replication policy	49
Advance to final phase	38
Automount files	
update references.....	23
Automount files	23

B

Baseline copy	19
Batch files.....	19, 40
block size.....	29, 46

C

Cancel a migration task	43
Cancel a task.....	42
Cancel replication.....	50
Change data distribution path.....	47
Change data distribution schedule	50
Compare contents of targets	51
Configure data distribution path.....	47
Content distribution	45, 47, 48, 50, 51
Contents of a replica.....	51
Contents of targets.....	51
Copy data.....	1, 18, 24, 33
Copy data incrementally.....	20
Create a replication policy.....	45
Cutover date	18, 33, 38

D

Data distribution	45, 47, 48, 50, 51
Data distribution policy	51
Data migration	1, 18, 30
Data move	45, 47, 48, 50, 51
Data transfer	
set rate	47
Data transfer	30
Data transfer	47
Delete	
a migration policy	43
a migration task.....	43
a replica.....	51
a replication policy	51
a target.....	51
Delete source data	21
DFS namespace	
update links	23
Differencing.....	29
Display replication operation information.....	50
Distribution path.....	47

F

FAT16	1
FAT32	1
File attributes	32
File differences	51
File system.....	1

H

Hash-based differencing	29, 46
-------------------------------	--------

I

Information about replication operations.....	4
---	---

L

Link target	
compare contents with other targets.....	51
List replication operations for a policy	50

M

Manually start data distribution	50
Migration destination.....	33
Migration operation	
final phase	21
final phase schedule	23
incremental phase.....	20
incremental phase schedule	21
initial phase	19
initial phase schedule	20
view progress	41
Migration operation	8
Migration operation	41
Migration policy	
add a task.....	40
cutover date	18
delete a policy	43
remove a task.....	43
replication options.....	24
run batch files.....	40
schedule data migration.....	20, 21, 23
security replication options	30
specify namespace options	23
start a task manually.....	41
start all task manually.....	41
suspend a task.....	43
update a migration task	41
Migration policy	33
Migration source.....	33
Migration task	
advance to final phase	38
cancel a task	42
cutover date	33
progress information	41

- remove from policy 43
- specify migration tasks from a file 40
- suspend a task.....43
- update a task.....41
- user action required.....42
- Migration task.....33
- Monitoring.....4
- Move data..... 1, 18, 24, 45, 47, 50, 51
- N**
- NetApp CIFS 1
- NFS.....1
- NTFS 1
- P**
- Perform final replication.....21
- Permissions.....1
- Planning considerations
 - replication 4
- Planning considerations.....4
- Ports.....4
- Prevent user connections during final copy.....21
- Privilege
 - backup 1
 - Ownership 1
 - restore.....1
 - security 1
- Privilege 1
- Properties
 - of a replication policy.....51
- Properties.....51
- R**
- Replace a source link of a replication policy....48
- Replace a target link of a replication policy 49
- Replica
 - compare contents of targets.....51
 - replication agent.....48
- Replication
 - planning considerations.....4
- Replication.....4
- Replication differencing 29
- Replication differencing options 46
- Replication operation information.....50
- Replication options
 - file attributes 32
 - file exclusions 24
 - file inclusions 24
- Replication options 24
- Replication options 32
- Replication options 46
- Replication policy
 - add or replace the source link.....48
 - add or replace the target link 49
 - cancel replication 50
 - change data distribution path.....47
 - change data distribution schedule..... 50
 - compare contents of targets..... 51
 - configure distribution path 47
 - create a policy 45
 - delete a policy 51
 - delete a target 51
 - hash-based differencing..... 29
 - replication agent..... 48
 - replication differencing 29
 - schedule data distribution..... 47
 - start replication immediately 50
 - view properties 51
 - view replication operation information 50
- Re-target a link 1
- Run batch files 40
- S**
- SAMBA 1
- Schedule data distribution 47
- Schedule migration 10, 20, 21, 23
- Script settings 40
- Security attributes 1, 30
- Security ID.....30, 46
- Security replication options 30
- Show replication operation information 50
- Show replication operations in progress 50
- Show replication operations with problems..... 50
- SID.....1, 30, 46
- Source link..... 48
- Specify migration tasks from a file..... 40
- Start a migration task manually 41
- Start all migration tasks manually 41
- Start data distribution..... 50
- Start replication 41
- Status monitoring..... 4
- Stop replication..... 50
- Suspend a migration task 43
- T**
- Target
 - delete a target 51
 - replication agent 48
- Target..... 48
- Target..... 51
- Target link of a replication policy..... 49
- Task 43
- Terminate replication..... 50
- Throttling 47
- Transfer data 1, 24, 30, 33
- Trustees 30, 46
- Typical migration 8
- U**
- Update the namespace 21
- User action required..... 42

View migration progress	41
View properties of a replication policy.....	51