NetworkAppliance®

The evolution of storage.™

VFM™ (Virtual File Manager) 5.6 Getting Started Guide

Part Number 215-01121

Copyright © 2001-2005 NuView, Inc. All rights reserved.

Copyright © 2003-2005 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 NuView, Inc. is subject to the following license and disclaimer:

Network Appliance and NuView, Inc. reserve the right to change any products described herein at any time, and without notice. Network Appliance and NuView assume no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance and NuView, Inc. 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 NuView, Inc.

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

NuView®, StorageX®, and associated logos are registered trademarks or trademarks of NuView, Incorporated. Other product and company names are trademarks of their respective owners.

NuView Incorporated, 738 Highway 6 South, Suite 850, Houston, Texas 77079 (www.nuview.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 trademarks of Network Appliance, Inc.

Table of Contents

PREFACE ..... **III**

CHAPTER 1 LEARNING ABOUT VFM **1**

COMPONENTS OF VFM 3
Architecture overview *3*
The desktop *4*
Storage policies *5*
WHAT'S NEW IN THIS RELEASE 7

CHAPTER 2 WORKING WITH VFM **9**

HOW TO WORK WITH VFM 9
Navigating your DFS namespace *9*
Using drag-and-drop *11*
View contents of an item in the DFS tree *13*
Find an item *14*
Locate an item *14*
Check system status *15*
Set default timeouts *15*
Set default SNMP community name *16*
Write audit events to the NT event log *16*
Set default policy options *16*
Filter event display *16*
Set up multiple views *17*
Hide empty Admin policies folders *17*
Hide folders on the Properties tab *18*
Disable tool tips *18*
Printer setup *18*
Connect to another server *18*
Integrate FLM *18*
Launch external tools *19*

CHAPTER 3 PLANNING CONSIDERATION HIGHLIGHTS **21**

NAMESPACE DESIGN CONSIDERATIONS 21
DISASTER RECOVERY PLANNING CONSIDERATIONS 23
MULTIPROTOCOL NAMESPACE PLANNING CONSIDERATIONS 23
SNAPMIRROR PLANNING CONSIDERATIONS 23

CHAPTER 4 GETTING STARTED **25**

INSTALL VFM 25
System requirements *25*
Install VFM *28*
CREATE A LOGICAL NAMESPACE 31
Benefits of using DFS *31*
Step 1 Create a DFS root *32*
Step 2 Add links *33*
Step 3 Check the results *33*
QUICKLY POPULATE THE NAMESPACE 34
Difference between namespace creation and permission based namespace policies *34*
Configure a namespace creation policy *34*
Configure a permission based policy *35*
MAKE THE NAMESPACE HIGHLY AVAILABLE 36
Step 1 Make DFS roots highly available *37*
Step 1 Create a namespace on a filer *37*

Step 2 Schedule synchronization *38*
MAKE DATA HIGHLY AVAILABLE 38
Create a client disaster recovery policy *39*
Create a filer disaster recovery policy *40*
Create a server disaster recovery policy *41*
OPTIMIZE STORAGE 43
Create an archival migration policy *43*
Create a storage load balancing policy *49*
TRANSFER DATA 50
Replication policies *50*
Distribute CIFS data *51*
Distribute NFS data *53*
Create a migration policy *53*
Configure SnapMirror *57*
VIRTUALIZE USER HOME DIRECTORIES 59
Create a home folder policy *60*
CREATE A MULTIPROTOCOL NAMESPACE 61
Step 1 Create a multiprotocol namespace policy *61*
BACK UP THE LOGICAL NAMESPACE 63
Step 1 Create a backup policy *63*
Step 2 Schedule backups *64*
Step 3 Browse your backup file *64*
BACK UP DATA 65
Step 1 Create a snapshot access policy *66*
Step 2 View the snapshot data *66*
Step 3 Restore data *67*
Step 4 Schedule snapshots *67*
CREATE AND PUBLISH A REPORT 68
Step 1 Add a machine to the Physical View tree *69*
Step 2 Enable report publishing *69*
Step 3 Create a Share Properties report *70*
Step 4 Gather data and publish the report *70*

APPENDIX 71

GLOSSARY 73

INDEX 76

Preface

What the guide covers

This guide describes how you can use VFM™ (Virtual File Manager) technology to simplify file management and increase data availability while reducing storage costs. It describes how to create and manage a Microsoft® Distributed File System (DFS) global namespace, prepare for rapid disaster recovery, migrate and replicate data seamlessly, virtualize user home directories, configure SnapMirror® software, and manage NetApp® filer activities.

The guide is divided into the following sections:

- **Chapter 1 Learning about VFM** provides an overview of the benefits of using VFM and introduces the VFM components and new features in this release.
- **Chapter 2 Working with VFM** - Explains how to customize VFM.
- **Chapter 3 Planning consideration highlights** provides information that you should consider when designing a namespace, planning for disaster recovery, creating multiprotocol namespaces, and transferring data using SnapMirror.
- **Chapter 4 Getting Start** gets you started using VFM to create your first namespace and solve your storage problems. This chapter explains how to install VFM, create a logical namespace, quickly populate the namespace, make the namespace highly available, make data highly available, transfer CIFS and NFS data, create a multiprotocol namespace, virtualize user home directories, back up data and the namespace, and create and publish a report.
- **Appendix** - Provides a list and description of the VFM guides.
- **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 XP Professional, 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.
- *<text in italics>* - Substitute the text within the greater than and less than symbols with the value requested.

Using the online help

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

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

- In the table of contents, a new topic is represented as a page with a red asterisk , and an entire section of new topics is represented as a book with a red asterisk .
- 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.
- When performing text searches, you can use the DOS wildcard " * " (asterisk).
- Context help contains detailed information about a dialog box or view, and might include links to related topics. Press **F1** or click **Help**.

Support

Contact NetApp technical support, 24 x 7.

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

Support phone numbers:

- 888 4-NETAPP (US and Canada)
- (1-888-463-8277)
- (408) 822-4700
- 800-44-NETAPP (EMEA/Europe)
- 800-800-0800 (Asia/Pacific)

Notes:

- You can also contact technical support by clicking **Support** on the **Help** menu.
- To get your product build number and license information, click **About Virtual File Manager** on the **Help** menu.

Enter new license key

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

licenses from VFM.

1. Contact NetApp to request a new license key.
2. On the **Help** menu, click **About Virtual File Manager**.
3. Click **Edit Serial Numbers**.
4. In the License Type list, click the appropriate item.
5. If you selected "Evaluation" as the license type, type the new activation key.
6. If you selected "Full" as the license type, select a product number from the Product list and type the new activation key.
7. Click **OK**.

Chapter 1 Learning about VFM

Network Appliance™ VFM™ (Virtual File Manager) is software for unifying and managing distributed file storage. VFM is an integrated data management solution that logically aggregates distributed file storage and provides administrators with policies that automate data management functions, such as heterogeneous migration, local and remote replication, client failover, capacity optimization, tiered data management, disk-to-disk archival, and data classification and reporting. VFM enables administrators to perform these tasks with minimal manual involvement and without causing client downtime. VFM is tightly integrated with NetApp SnapMirror® and Snapshot™ products to deliver best-in-class management capabilities for distributed NetApp file storage environments.

VFM simplifies enterprise file storage management in both multiprotocol client environments (for example, UNIX® and Windows) and heterogeneous server/filer environments (for example, NetApp, Windows, Windows NAS, EMC® Celerra). VFM directly addresses your current needs and those of your users by increasing data availability, minimizing client downtime, optimizing storage capacity, and dramatically simplifying storage management.

- **VFM simplifies enterprise storage management using a 'Global Namespace'**

  Deployment of a global namespace is the key to effective, efficient management of distributed file storage. A namespace enables clients to access the files without having to know the physical location of the files (just as they access a web site without knowing the underlying IP address of the site). Users no longer need to access their data using drive letter mappings. The VFM global namespace permanently eliminates drive letter dependencies, thereby providing administrators with tremendous flexibility in expanding, moving, and reconfiguring physical storage without affecting users. VFM achieves this by aggregating file storage across heterogeneous storage platforms (Windows and NetApp), and by providing a unified logical view of the storage. VFM enables role-based hierarchical administration and multi-user management.

- **VFM makes possible non-disruptive expansion or consolidation of servers and storage**

  VFM enables you to seamlessly consolidate data from multiple heterogeneous file servers across heterogeneous file storage platforms. The global namespace masks physical changes from users during consolidation, and VFM storage policies significantly reduce operator activities during the consolidation. By first deploying a global namespace, you can use VFM data movement features to quickly and easily consolidate or expand file servers. VFM enables one-step data migration between storage devices to optimize capacity and achieve server consolidation. This can be done even when devices are from different vendors and are geographically dispersed. By using VFM, you can perform multiple simultaneous consolidation jobs and reconfigure file storage without affecting how users access the files, because VFM automatically and transparently redirects users to the files in their new locations.

- **VFM enables intelligent data movement through replication and migration policies**

  VFM enables seamless replication and data migration between storage architectures. With VFM, you can automate and manage complex replication and data migration jobs quickly and reliably. VFM only replicates the changed data, instead of replicating entire files and wasting WAN bandwidth. Key file attributes are automatically retained while

consolidating data across multiple file storage platforms, thus reducing the administrative workload. VFM enables administrators to set multiple migration schedules with bandwidth consumption settings and modify file attributes during migration.

- **VFM enables rapid disaster recovery (DR) by supplying near-instantaneous failover for users and filers across locations**

  The true test of any disaster recovery solution is in how quickly clients regain access to data after a failure. In the event of a disaster, VFM shortens the restore window from hours or days to minutes without requiring client reconfiguration or network rerouting. Additionally, VFM provides you with a comprehensive automated solution for deploying and managing NearStore® solutions. You assign a NearStore device as the target for multiple primary filers and/or Windows servers (including Windows-attached SAN) and VFM automatically synchronizes qtrees and shares present on multiple source devices with the NearStore device.

- **VFM centralizes the configuration and management of routine filer operations including snapshots, SnapMirrors, and qtrees**

  VFM enables the centralized configuration, viewing, scheduling, and management of NetApp filer snapshots and SnapMirror pairs across multiple filers. VFM can initiate ‘logical’ snapshots, which, in turn, trigger physical snapshots on multiple filers. In this manner, VFM makes taking snapshots across multiple filers a ‘one-step’ process and enables data availability based on business priorities (business priorities are reflected by way of the logical namespace). VFM also provides a consolidated view of multiple NetApp filers from a single console, and generates a range of descriptive reports that provide tremendous insight into the filer storage environment.

- **VFM facilitates the creation of a ‘tiered’ storage architecture that results in reduced costs of storage ownership**

  VFM enables the creation of ‘storage tiers’, while maintaining unified access to all tiers of storage through the global namespace. User access to data is not disrupted or affected by data transfers across the storage tiers. VFM accomplishes this by policy-driven movement of data automatically across heterogeneous devices, based on criteria such as age, usage, and size. This feature enables you to reduce the costs of storage in the enterprise by placing important data on higher-cost storage, and less important data on cheaper storage.

  VFM provides policy-based load balancing of storage across devices in a defined ‘pool’ of physical devices. VFM integrates with SnapLock™ functionality, to provide data management features for compliance data. VFM also enables rich data classification and reporting that can help IT make business-driven decisions for their data.

Components of VFM

This section provides a high-level overview of the VFM architecture, introduces the main components of the VFM desktop, and describes the storage policies that you can create.

- For a high-level overview of the architecture, see "Architecture overview."
- For an introduction to the VFM desktop, see "The desktop."
- For information about storage policies, see "Storage policies."

Architecture overview

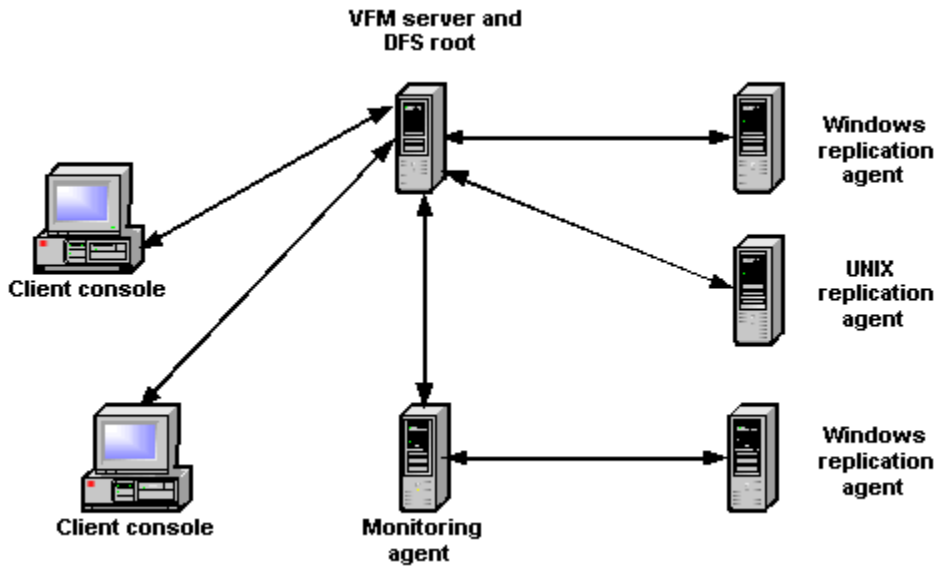

Figure 1 Distributed client/server architecture

VFM is implemented using distributed client/server architecture. This design provides a flexible and efficient means to manage Windows storage. The application includes the following components:

- **VFM client** – The client provides the VFM administrator interface or console. It communicates with the VFM server as well as with the DFS roots you are managing. Multiple VFM clients can connect to the VFM server to manage the namespace and policies.
- **VFM server** – The server acts on behalf of the VFM client to carry out VFM policies and to maintain storage configuration and reporting information. The server is implemented as an NT service. Another machine can be configured as a standby server to take over the role of the primary server if it becomes unavailable.
- **VFM monitoring agent** – The agent allows distributed processing of VFM policies. The monitoring agent is implemented as an NT service.
- **VFM replication agents** – The agents carry out file transfers for VFM policies. VFM

provides Windows and UNIX replication agents. The agents are deployed by VFM. You can choose to manually deploy an agent.

- Windows replication agents are implemented as NT services. The agents transfer CIFS data.
- UNIX replication agents are available for RedHat® 9, RedHat Enterprise 3.0, and Solaris® 8. The agents transfer NFS data.

The desktop

The left pane of the VFM desktop displays three views (Admin, Logical, and Physical), Reports, and File Lifecycle Manager. When an item from the left pane or tree is selected, attributes of the item are displayed in the right pane. The VFM desktop provides a graphical view of your resources in several ways.

- Create and manage storage policies in the **Admin View**.
- Simultaneously view the logical and physical layout of your DFS resources while you build or update your namespace in the **Logical** and **Physical Views**, respectively.
- View, record, and publish information about the namespace, policies, and supporting machines, shares, and exports in **Reports**.
- Integrate **File Lifecycle Manager** (FLM) with VFM to manage files from their creation to their eventual disposal. For instructions about how to integrate FLM with VFM, see "Integrate FLM."
- Display attributes of resources in the right pane.

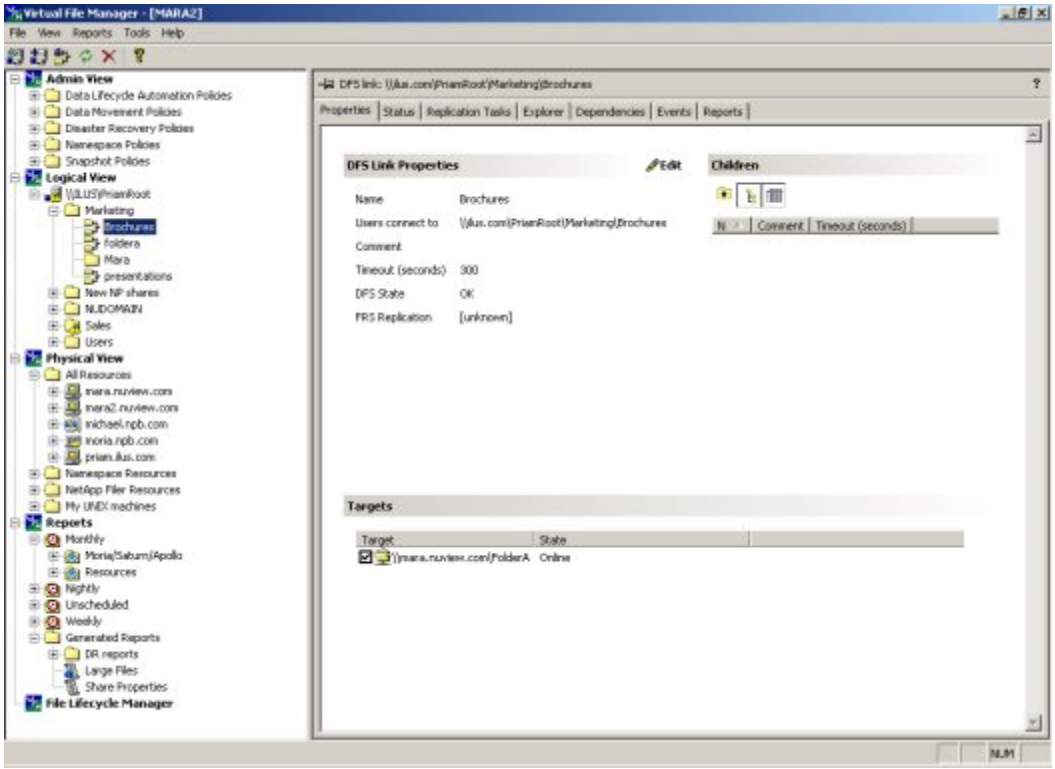

Figure 2 The VFM desktop

Storage policies

VFM makes it easy for you to manage storage using automated, intelligent storage policies. The following list describes the policies that you can create.

- **Namespace creation policies** enable you to easily populate and maintain a logical namespace. You can use VFM namespace creation features interactively to perform the initial population of your namespace. After the namespace is in place, the policies can automatically run on a scheduled basis to scan for changes in the physical resources supporting the logical namespace. You can configure policies to report changes, or to update the logical namespace automatically to reflect the changes detected by the policy.
- **Permission based namespace policies** enable you to populate your namespace based on the security permissions of shared folders and subdirectories of shares that you want links in the namespace to reference. These policies work the same as namespace creation policies.
- **Namespace availability policies** enable you to easily keep the logical structure of two or more DFS roots aligned and enable you to easily create widelink namespaces on NetApp filers. When you synchronize roots, you can balance the load on your DFS root servers across multiple identical instances of the root. This provides fault tolerance because the namespace is available from more than one server. A logical namespace hosted on a NetApp filer provides the benefits of a namespace without requiring additional Windows servers.
- **Client disaster recovery policies** enable you to set up VFM to detect a failure and automatically fail over clients to a replicated storage device. One of the most significant advantages of the VFM disaster recovery solution is the ability to designate a cascaded failover path. A typical scenario is one in which you designate critical data to fail over first to a local edge device and then to a remote central device. This ensures high availability of the data whether there is a local server failure or an entire site outage. Failover takes place without changing login scripts or network configuration.
- **Filer disaster recovery policies** enable you to make data on NetApp filers participating in your namespace highly available. VFM makes it easy to set up a filer or nearline storage device as a disaster recovery backup to one or more primary filers. VFM leverages SnapMirror replication to replicate data to a standby volume or qtree. When a failure on the primary storage is detected, an automatic failover to the standby storage is initiated. VFM periodically checks for new links and qtrees on the primary filer and creates the new data configuration on the backup filer.
- **Server disaster recovery policies** enable you to make data on servers participating in your namespace highly available. VFM makes it easy to monitor for new links targeting the primary server to add to the policy, replicate data to a secondary server, and automatically fail over all links when a failure of the primary server is detected.
- **Archival migration policies** enable you to optimize your storage. VFM enables you to set policies to control when old or unused data is migrated to secondary storage. Infrequently used data can be migrated to lower cost storage to free up space on primary servers. You control whether VFM automatically carries out the migration or simply notifies you of items matching your criteria so that you can migrate them later.
- **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 for your environment. VFM provides status

information about replication operations performed using the VFM replication agent. VFM replicates data from a central location to remote sites and from remote sites to a central location.

- **Migration policies** enable you to balance free space and data throughput capacity across volumes, increase reliability and availability by creating alternate copies of data, or migrate data prior to maintaining or consolidating servers. Migration policies work in heterogeneous environments.
- **SnapMirror policies** enable you to create and centrally manage replication between NetApp filers in your network. You can integrate SnapMirror data replication with your namespace. VFM provides a central view of events from all managed NetApp filers. You receive detailed reporting of filer configuration and status information.
- **Storage load balancing policies** enable you to automatically balance the storage space among a group of storage locations containing directories that are targets of links in your namespace. The policy periodically evaluates the space distribution of the group of storage locations, calculates a solution, and automatically creates and runs migration tasks to move directories and retarget links to balance the storages. You can view migration suggestions and decide which suggestions you want to use to create migration tasks.
- **Multiprotocol namespace policies** enable you to create a common namespace for UNIX and Windows users to provide access to files stored on NetApp filers. This presents a unified view of files to users. A multiprotocol namespace policy enables you to build namespaces for NFS clients to mirror the namespace presented to CIFS clients, and to build namespaces for CIFS clients to mirror the namespace presented to NFS clients. VFM enables you to schedule namespace synchronization.
- **Home folder policies** enable you to manage user storage in the logical namespace as well as in Active Directory and NT domain user information. VFM adds user home directories to a logical namespace after querying Active Directory or an NT domain controller, thus virtualizing the user's storage. By virtualizing home directories, you gain flexibility to make changes to physical storage without affecting users’ access to data.
- **Snapshot access policies** enable you to centrally manage snapshots on multiple NetApp filers. VFM makes it easy for you to take a snapshot of data and the logical structure of the namespace that references the data. You take snapshots of volumes participating in your namespace. VFM makes it easy to restore files and folders from snapshot backups.
- **Snapshot scheduling policies** enable you to schedule snapshots on NetApp filer volumes. VFM enables you to centrally manage filer-initiated and VFM-initiated snapshots on multiple filers with a single operation. This is a convenient way to change the snapshot schedule and reserve setting on multiple filers without requiring you to repeat the changes on each individual filer.
- **Namespace backup policies** enable you to automatically back up the logical structure of multiple DFS roots. If the server hosting a backed up root becomes unavailable, you can re-create the links contained in a backed-up root onto a DFS root in the Logical View. The policies enable you to recreate an entire root or merge part of a root with an existing root.

What's new in this release

This release of VFM provides the following new and enhanced features:

- **Volume Shadow Copy Services (VSS) and snapshot replication integration** enables replication of locked files. When replicating from a shadow copy or snapshot, the file contents can be transferred even when the source file is locked. This can aid in the replication of directories where files remain locked by applications, such as Outlook PST files that are locked while the Outlook client is running.
- **Agent Management** enables you to centrally manage replication agents, such as update credentials and manually deploy an agent. This feature also enables you to update credentials for the VFM server and all installed monitoring agents.
- **Option to use SSH remote shell** to enable VFM to communicate more securely with NetApp filers and UNIX machines.
- **New replication best practice options** enable you to optimize data transfer for speed or integrity.
- **Better access to troubleshooting information** enables you to solve problems quickly. Events provide links to related troubleshooting information.

Chapter 2 Working with VFM

This chapter provides general instructions about how to use and customize VFM.

How to work with VFM

Learn how to navigate in VFM, set default policy options, configure the desktop for multiple views, and write audit events to the NT event log.

Navigating your DFS namespace

Quickly display the attributes or a context menu for any item appearing in any of the views on the console.

Item	Admin View	Logical View	Physical View
Link	**Client disaster recovery policy** Select a link in the tree or double-click a link on the Properties tab of a policy to view its attributes in the right pane **Replication policy** Double-click a link on the Replication Topology tab to view its attributes in the right pane. Right-click a link anywhere in this view to display a context menu.	Select a link in the tree to view its attributes in the right pane. Double-click a link on the Properties tab of a DFS root or folder to view its attributes in the right pane. Right-click a link anywhere in this view to display a context menu.	Double-click a link on the Properties tab of a shared folder to view its attributes in the right pane. Right-click the link on the Properties tab of a shared folder to display a context menu.
Root		• Double-click a root on the tree or on the Properties tab of to view its attributes in the right pane. • Right-click a root on the tree or on the Properties tab of	

Item	Admin View	Logical View	Physical View
		to display a context menu.	
Storage policy	• Select a storage policy in the tree to view its attributes in the right pane. • Right-click a storage policy in the tree to view a context menu.	Click the name of a storage policy on the Properties tab of a link participating in the policy to view its attributes.	
Shared folder	**Replication policy** Double-click a shared folder on the Replication Topology tab of a policy to view its attributes.	Double-click a shared folder target on the Properties tab of a link to view its attributes. Double-click a shared folder replica on the Properties tab of a root to view its attributes. Right-click a shared folder anywhere in this view to display a context menu.	Select a shared folder in the tree to view its attributes in the right pane. Double-click a shared folder on the Properties tab of a machine to view its attributes in the right pane. Right-click a shared folder anywhere in this view to display a context menu. **Note:** Shared folders that participate in the namespace are yellow. Shared folders that do not participate in the namespace are blue.
Machine			Select a machine in the tree to view its attributes. **Note:** Machines with shared folders that participate in the namespace are

Item	Admin View	Logical View	Physical View
			yellow. Machines with no shared folders participating in the namespace are blue.

Using drag-and-drop

VFM provides extensive drag-and-drop support to make it easy to carry out common namespace management tasks. The following tables illustrate what items you can drag and where you can drop them, and explains the results of these operations.

General		
Drag	**Drop on**	**To...**
Share	DFS root	Add a link to the root that targets the share.
Share	DFS link	Add a target to the link.
Share	Client disaster recovery policy	Add a link to a policy. The Disaster Recovery Link Target wizard steps you through the process of creating a link to the share and adding another share to the link.
Share	Replication policy	Perform one of the following tasks: • Create a new source link if one does not exist. • Replace the existing source link's target. • Create a new target link if one does not exist. • Add to the existing target link.
Share, directory, or machine	Machine or share	Create a new migration task.
Share, directory, or machine	Migration policy	Create a new migration task.

General		
Drag	**Drop on**	**To...**
UNC directory or share path from Windows Explorer	DFS root	Add a link to the root that targets the share or folder.
DFS root	Namespace availability policy	Add the root to the policy.
DFS root	Namespace Availability Policies folder	Create a new namespace availability policy.
DFS root	Multiprotocol Namespace Policies folder	Create a new multiprotocol namespace policy.
DFS link	Disaster Recovery Policies folder	Create a new client disaster recovery policy for the link.
DFS link	Replication Policies folder	Create a new replication policy that contains the link.
DFS link	Client disaster recovery policy	Add the link to a policy. If the link does not have replicated targets, the Disaster Recovery Link Target wizard steps you through the process of adding another target to the link.
DFS link	Replication policy	Perform one of the following tasks: • Add the source link of a policy. • Add the target link of a policy. • Add the replicas of the link to the target link.
DFS link	Folder in the Logical View	Perform one of the following tasks: • Copy a link to another root. • Move a link within a root.
Machine	Client disaster recovery policy	Add multiple links to a policy. The Disaster Recovery Link Target wizard steps you through the process of adding links to target the shares on the machine.
Machine	Physical View folder	• Copy a machine from the Namespace

General		
Drag	**Drop on**	**To...**
		Resources folder in the Physical View to another folder in the Physical View. • Copy or move the machine to other folders in the Physical View.
NetApp filer	Client disaster recovery policy	Add multiple links to a policy. The Disaster Recovery Link Target wizard steps you through the process of adding links to target the shares on the filer.
NetApp filer	SnapMirror policy	Add the filer to the SnapMirror policy.
Report	Report folder	Move the report to another folder under Reports.
Report folder	Report folder	Move the folder within a folder under Reports.

Replication Topology (Example: Source target distributes data to targetA and targetB.)		
Drag	**Drop on**	**To...**
Destination targetB	Destination targetA	Change the data distribution path so that data is copied from source target to targetB and from targetB to targetA.
Source target	Destination targetA	Change the data distribution path so that data is copied from the source target to targetA and targetB.

View contents of an item in the DFS tree

You can view the contents of a DFS link, share, machine or folder from either the Logical View or Physical View.

- Items in the Logical View display an **Explorer** tab when you click them. Open the **Explorer** tab and click the item to display its contents in the right pane of the tab.
- Right-click an item, point to **External Tools**, and click **Explorer**.

Find an item

Find an item in the DFS tree.

1. Click anywhere on the tree.
2. Perform one of the following actions:
   - Press **Ctrl+F**.
   - On the **Tools** menu, click **Find**.
3. In the **Find what** field, type the word or phrase that you want to search.
4. By default, VFM searches for a string in all sections of the tree. Clear a check box under **Look at** to remove a section from the search.
5. Select **Search Comments** if you want to search in the **Comments** field of items in the tree. The Find operation takes a significant amount of time if comments are included in the search because VFM must gather information from the network.
6. Click **Find**. The first item that meets your search criteria is highlighted in the tree.
7. Perform one of the following actions to find the next instance of the search string.
   - Press **F3**.
   - On the **Tools** menu, click **Find Next**.

Locate an item

Find where a logical item in your namespace is physically located on your network or find where a physical item in your network is logically located in the namespace.

To locate logical item in the network:

1. In the **Logical View** tree, click a DFS link or root to display the **Properties** tab in the right pane.
2. Perform one of the following actions:
   - Double-click a shared folder.
   - Right-click a shared folder and click **Locate**.

   In the **Physical View** tree, the shared folder is highlighted under the machine where it is located.

To locate physical item in the namespace:

1. In the **Physical View** tree, click a machine to display all the shared folders on the machine that are part of the namespace.
2. Click the shared folder that you want to find in the **Logical View** tree. Tabs appear in the right pane.
3. Perform one of the following actions:
   - On the **Properties** tab, double-click a reference.
   - On the **Properties** tab, right-click a reference and click **Locate**.
   - Click the **References** tab, right-click a reference and click **Locate**.

   In the **Logical View** tree, the link referencing the shared folder is highlighted and tabs for

the link appear in the right pane.

To locate a storage policy:

Quickly view the properties of the storage policy to which an item in the **Logical View** belongs.

On the **Properties** tab of an item in the **Logical View**, click the name of the policy under **Properties**.

Check system status

View the status of the VFM server and the monitoring agent if the agent is installed on a different server. If problems are detected, a warning or error status symbol appears on the status bar. To view status information, perform one of the following actions:

- Double-click the status icon on the status bar.
- On the **Tools** menu, click **System Status**.

Note: If the status bar does not appear, click **Status Bar** on the **View** menu.

Set default timeouts

Set the duration that the DFS roots or the targets of DFS links are cached on DFS clients. The changes only affect new roots and links that you create after making the change in this dialog box. Existing roots and links are not affected.

1. On the **Tools** menu, click **Options**.
2. Click **System Options**.
3. Under Default Timeouts, type the cache time-out in seconds for DFS roots and link targets.
4. Click **OK**.

Notes:

- When the cache period elapses, the DFS client must access the root share to update the shared folder information. Entering a lower value increases associated network traffic. A higher value reduces associated network traffic at the possible expense of accurate share information.
- If the cache time-out value is set too high, client computers might not pick up changes to DFS referrals unless they are restarted (See Microsoft Knowledge base article Q222604).
- Windows NT 4 DFS roots do not support changing the cache time-out.

Set default SNMP community name

VFM uses "public" as the default SNMP community name. You can set an application-wide default value. Changes only affect machines that become part of the namespace after the update. Machines participating in the namespace before the update are not affected.

1. On the **Tools** menu, click **Options**.
2. Click **System Options**.
3. Type a name in the **Community Name** field.
4. Click **OK**.

Note: The SNMP community name is used in the process of identifying NetApp filers.

Write audit events to the NT event log

You can write VFM audit events to the NT event log.

1. On the **Tools** menu, click **Options**.
2. Click **System Options**.
3. Select **Write audit events to NT event log**.
4. Click **OK**.

Note: By default, VFM shows events for items in the Admin, Logical, and Physical Views. Select an item to display an Events tab in the right pane.

Set default policy options

Set default parameters to use when creating new storage policies. Changes you make do not affect policies created before the changes were made.

1. On the **Tools** menu, click **Options**.
2. Click a policy or any item under a policy to show a dialog box in the right pane.

Note: To update options for a single policy, right-click the policy in the Admin View and click **Properties**.

Filter event display

By default, VFM displays all the event history for items in the logical and physical layout of your DFS resources. You can filter this information.

1. On the **View** menu, click **Filter event records**.
2. Select the options you want.
3. Click **OK**.

Notes:

- Event filter option settings are global. The options apply to all items in the logical and physical layout.
- Event filter option settings are persistent. Any options you set up when VFM is shut down are still there when you restart VFM.
- Events saved to a report are determined by the filter options in effect when you save the report. Changes to event filter option settings do not affect any existing event reports.
- You can use the filtering facility to prevent older events from appearing so that it is easier to see just the newest events.

Set up multiple views

Split the right pane to view and edit the attributes of more than one item in the tree.

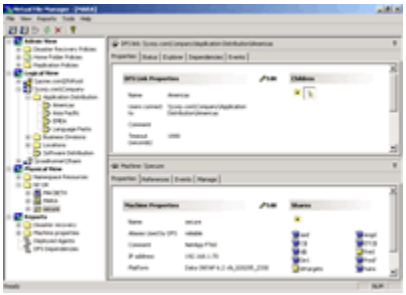

Figure 3 Split view

1. Select an item in the tree.
2. On the **View** menu, click **Split right pane**. The right pane splits into two windows, both of which display the attributes of the item selected.
3. Click the **Pin** on one of the windows to "keep" the information and select another item in the tree. The attributes of the second selected item appear in the other window.

Note: When the attributes of an item are pinned, the **Pin** looks like this .

Hide empty Admin policies folders

By default, VFM always displays the Admin policies folders. You can hide folders that do not have policies.

1. On the **Tools** menu, click **Options**.
2. Clear **Show Empty Admin Folders**.
3. Click **OK**.

Hide folders on the Properties tab

By default, VFM displays folders on the **Properties** tab of a DFS root or link with a root as the target. You can hide the folders and display links as a flat list.

1. In the Logical View, click a DFS root or link.
2. Click the **Properties** tab in the right pane.
3. Under Children, click the **Folder** button to prevent folders from appearing.

Disable tool tips

By default, tool tips appear for items in the Logical View, machines in the Physical View, and policies in the Admin View. You can disable this feature.

1. On the **Tools** menu, click **Options**.
2. Click **User Options**.
3. Clear **Show Tool Tips**.
4. Click **OK**.

Printer setup

Select a default printer and set up appropriate parameters.

On the **File** menu, click **Print Setup**.

Connect to another server

Connect to another VFM server. Multiple administrators can work simultaneously with VFM to monitor replication, migration, and disaster recovery jobs.

1. On the **File** menu, click **Connect to**.
2. Type the path of the machine hosting the VFM server to which you want to connect.
3. Click **OK**. Your VFM client now points to the server you specified.

Integrate FLM

Integrate File Lifecycle Manager with VFM to automate the lifecycle of files from their creation to their eventual disposal.

1. On the **Tools** menu, click **Options**.
2. On the User Options dialog box, select **Show Lifecycle Manager** and type the URL for the FLM server.
3. Click **OK**.
4. Click **File Lifecycle Manager** on the left pane of the VFM desktop to display the application on the right pane.

Launch external tools

Launch other tools, such as Microsoft Management Console (MMC), Windows Explorer, and Disk Manager, while working with VFM.

Perform one of the following actions:

- On the **Tools** menu, point to **External Tools** and click the application that you want to open.
- Right-click any item on the DFS namespace tree, point to **External Tools**, and click the application that you want to open.

Chapter 3 Planning consideration highlights

This chapter highlights things you should consider when you create a logical namespace, plan for disaster recovery, create a multiprotocol namespace, and configure SnapMirrors using VFM.

Review the following topics before you create your first namespace and storage policies:

- Namespace design considerations
- Disaster recovery planning considerations
- Multiprotocol namespace planning considerations
- SnapMirror planning considerations

Note: Detailed information about logical namespaces and storage policies can be found in the VFM Advanced Configuration and Techniques Guide or the online help.

Namespace design considerations

1. Namespace layout and design considerations

- To achieve the full benefit of storage virtualization, you must consider the namespace layout and design. VFM provides unlimited flexibility about how you design the namespaces, but for best effect, you should work closely with your user communities to design namespaces that meet their needs now and in the future. Namespaces are designed around some organizing principle, which can be geographical, functional, departmental, or something else altogether.

Namespace Layout Alternatives

Organizing Principle	Examples
• Corporate function	• HR, Sales, Marketing
• Geographical location	• US, Europe, Asia
• Department or division	• Corporate, Regional Divisions

- The flexibility inherent in the logical namespace concept makes it possible for you to create multiple namespace branches to access the same data. For example, consider a company with sites throughout the Americas. One namespace branch might be organized by geographic region. Another namespace branch might be organized by corporate function, that is, Sales, HR, and so forth. Both branches of the namespace resolve to the same data; you simply give users more than one way to find the information. This flexibility is very powerful in that it enables you to quickly respond to changes brought on by corporate reorganizations, employee moves, company mergers, and more. You can

organize your namespace using the Active Directory domain and OU (Organizational Unit) model as a prototype.

- Just as a single namespace can provide different paths to the same data, you might consider creating multiple namespaces to serve the needs of different audiences. Examples of these could be namespaces for web content, software distribution, application configuration, and so forth. Keep the namespace simple.
- In addition to the namespace organization decisions, reliability requirements must be considered during logical namespace design. You might want to replicate the namespace to allow it to be available from multiple servers. This eliminates the server that provides the namespace from being a single point of failure.
- A single namespace server can handle about 2,000 to 4,000 users, depending on the hardware configuration. The load on a namespace server is similar to the load on a DNS server. Size your namespace server the same way you size your DNS server. It is also important to place your namespace close to the users.
- Because it is desirable to minimize changes to the namespace after it is deployed, it is important that you design the namespace thoughtfully from the onset. The key to successful creation of a virtual file system is in the thoughtful, thorough design of the virtual layer or namespace. The goal of the namespace is to present users and applications with a long-term consistent means to access data.

2. **Decide on the root configuration**

   Create a stand-alone or domain-based root. The similarities and differences between the two types of DFS roots are illustrated in "Manage a DFS root" in the VFM Namespace Guide and the online help.

3. **Populate the namespace with links referencing network shares**

   VFM provides several ways to populate the namespace. Select a method that best meets your needs.

   - **Add a link dialog box** enables you to add one link at a time.
   - **Drag UNC paths** from Windows Explorer to a root in the DFS tree.
   - **Add multiple links dialog box** enables you to add multiple links and folders at a time. Use this method when you know where items are located on your network and most of the DFS shares that you want to link are located on a few machines. A wizard steps you through the process. You can drag items from Windows Explorer to a folder.
   - **Namespace Creation Policy wizard** enables you to add multiple links and folders at a time. Use this method when you want to add links to multiple shares that are supported by different machines and when you want to keep the logical namespace aligned with subsequent changes in the underlying physical resources.
   - **Permission based namespace wizard** enables you to add multiple links and folder at a time like the namespace creation policy wizard. Use this method when you want to populate the namespace based on security permissions of shares.
   - **Home Folder Policy Creation wizard** enables you to create links to users' home directories from information stored in Active Directory or an NT4 domain.

4. **Instruct users to access data through the logical namespace.**

   After users are comfortable with the namespace, you might want to hide shares to ensure that they access data through the logical namespace rather than by using the underlying share names. Rename the shares or place a dollar sign ($) at the end of the share name to hide

them. Until the majority of the users access the data through the namespace, you will not get the total benefit of a namespace.

Disaster recovery planning considerations

When planning for disaster recovery, consider the following:

- **Namespace availability -** If the namespace must be highly available, options such as domain root replicas or hosting the root on a clustered server should be evaluated.
- **Data availability -** How resilient the infrastructure storing the data in the namespace is and whether shares are hosted on RAID volumes or clustered servers.
- **Client and server connectivity -** Check whether network paths between clients and data are configured for fault-tolerance.

Multiprotocol namespace planning considerations

The following list details prerequisites and restrictions to implementing a multiprotocol namespace policy:

- You have DFS roots containing links that target data on NetApp filers.
- You must be familiar with the NFS protocol and UNIX system.
- You must be familiar with configuring and managing NetApp filers.
- Your filers must be licensed for both CIFS and NFS protocols.
- Auto_master files are not supported.
- VFM does not support the Solaris-specific automount options, such as running executables and evaluating environment variables when determining the actual export to mount.

SnapMirror planning considerations

For detailed information concerning such matters as prerequisites to replicating a volume and restrictions that apply to the use of SnapMirror, see the following documents:

- Data ONTAP™ System Administrator's Guide
- Data ONTAP Data Protection Guide

Chapter 4 Getting Started

This chapter illustrates how VFM provides efficient storage management. It serves as a tutorial to get you started using VFM to create your first namespace and manage your storages and data.

Before you begin, review the information under "Planning considerations highlights." The information covers things you should consider when creating namespaces and storage policies.

This section guides you through the process of using VFM to perform the following tasks:

- Install VFM
- Create a logical namespace.
- Quickly populate the namespace.
- Make the namespace highly available by making DFS roots highly available and by creating a widelink namespace on a NetApp filer.
- Make data highly available by creating client, filer, and server disaster recovery policies.
- Optimize storage by creating archival migration and storage load balancing policies.
- Transfer data by creating replication, migration, and SnapMirror policies.
- Virtualize user home directories by creating a home folder policy.
- Create a multiprotocol namespace to allow NFS and CIFS clients to access data on NetApp filers through a logical namespace.
- Automatically back up the logical structure of your namespace by creating a backup policy.
- Back up data on NetApp filers by creating snapshot access and snapshot scheduling policies.
- Create and publish a report.

Install VFM

This section explains system requirements, permissions the account used by the VFM components must have, Windows and UNIX replication agent requirements, and installation instructions.

System requirements

Before installing VFM, check that the system where the VFM components are installed meets the following requirements:

- Microsoft Windows® 2000 SP 2 or later or Windows XP Professional to manage roots on NT® 4 or Windows 2000 servers.
- Microsoft Windows XP Professional or Windows Server™ 2003 to manage multiple roots on a single Windows Server 2003.
- Internet Explorer™ 5.5 or later.
- Windows Script Host 5.6 or later.

- Disk space:
  - VFM self-extracting file needs 155 MB.
  - Application components: client only - 35 MB; server only - 200 MB; monitor only - 125 MB.

  Notes:
  - During installation, add 80 MB to the VFM component disk space requirements.
  - As you use VFM, additional disk space is needed to store items such as published reports and namespace backups.

- VFM components must be installed on local fixed NTFS disks.
- Minimum video display of 1,024 by 768 and 256 colors. For the best display, it is strongly recommended that you use more than 256 colors.

Notes:
- If the number of connections used between VFM components and the connections used by other processes on the system could exceed the limits on inbound connections imposed by Windows 2000 Professional or Windows XP Professional, install the VFM server on a Windows 2000 server or Windows Server 2003. These limits are documented in the Microsoft Knowledge Base Articles 122920 and 314882.
- VFM installs and uses the Microsoft SQL Server™ Desktop Engine as part of the VFM server or VFM monitoring agent.
- VFM installs the Microsoft .NET Framework during VFM installation if it is not already on the system.
- The evaluation license of VFM limits the amount of data you can transfer.

Permissions

The account used by the VFM server, a replication agent, or a VFM 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

Windows replication agent deployment requirements

- VFM deploys a replication agent to machines running Windows 2000, Windows Server 2003, Windows XP file systems, or Windows NT 4 systems where MSXML3 is installed.

- If the machine to which you want to deploy a replication agent is running Windows XP SP 2, you must perform the following actions:
    1. Disable the firewall on the machine you want to deploy the replication agent.
    2. Restart the machine.

UNIX replication agent installation requirements

You must install a replication agent on the source or destination when transferring NFS data. If you cannot install an agent on either the source or destination, install an agent on a proxy machine.

General requirements

- The UNIX machine where you deploy a replication agent must be running RedHat 9, RedHat Enterprise 3, or Solaris 8 operating system.
- The account used by VFM to deploy a UNIX agent must have root privileges on the machine where you deploy the replication agent.
- NFS 3 or later must be installed on the source and destination.
- The file system where the replication agent will be installed must be mounted read-write. If you attempt to install the replication agent on a read-only file system, the following error message appears:
    - On Solaris:
        - pkgadd: ERROR: Unable to make directory </usr/local/URA>
    - On Linux:
        - error: unpacking of archive failed on file /usr/local/ura/bin:
        - cpio: mkdir failed - No such file or directory

Solaris 8

The following packages must be installed on the machine where the replication agent will be deployed:

- SUNWcsl
- SUNWlibms
- SUNWcsu
- SUNWloc
- SUNWgzip

RedHat 9

The following packages must be installed on the machine where the replication agent will be deployed:

- glibc-2.3.2-27.9.7

- bash-2.05b-20.1
- coreutils-4.5.3-19.0.2
- glibc-common-2.3.2-27.9.7
- glibc-2.3.2-27.9.7
- gzip-1.3.3-9
- grep-2.5.1-7.8
- procps-2.0.11-6
- Redhat package manager version 4 or greater.

Note: The version of the packages can be later than the ones listed.

RedHat Enterprise 3

The following packages must be installed on the machine where the replication agent will be deployed:

- glibc-2.3.2-95.20
- glibc-common-2.3.2-95.20
- bash-2.05b-29
- coreutils-4.5.3-26
- gzip-1.3.3-9
- grep-2.5.1-16
- procps-2.0.13-9.2E
- Redhat package manager version 4 or greater.

Note: The version of the packages can be later than the ones listed.

Install VFM

VFM provides flexibility in how components are deployed in your environment.

Account permissions and privileges for installing and running VFM

- The account installing VFM from media must have administrator privileges on the local system.
- The account used for the VFM server, which is specified during installation, must have the following permissions and privileges:
    - Appropriate permissions to administer the DFS roots that are managed
    - Administrator privileges on any machines hosting either source or destination of data replications
    - "Run as service" privilege
- The account used when running the VFM console should be an administrator on the local system and have the appropriate permissions to administer the DFS roots that are managed.
- When administering SnapMirror information, the account running VFM must have

administrator privileges on any managed filers to access the *\\filer\C$* share and read and write files in the */etc* directory of the filer. VFM also uses RSH to communicate with filers when managing snapmirrors and qtrees. The RSH username and password for VFM to use when communicating with filers is configured through the machine properties dialog box for the filer.

- VFM also uses RSH to communicate with UNIX systems, such as Solaris and Linux. RSH must be enabled on these systems as root. See the VFM Advanced Configuration and Techniques Guide for instructions about how to enable RSH as root on Solaris and Linux systems.

Connection limits

- The VFM server accepts inbound connections from replication agents in response to commands sent from the server to the replication agents. If these connections, in conjunction with other connections existing on the machine, exceed the limits imposed by Windows 2000 Professional and Windows XP, the VFM server cannot receive responses from its replication agents and reports errors. If the connection count limits are an issue in your environment, run the VFM server on a Windows 2000 Server or Window Server 2003 to avoid this problem.
- Microsoft Knowledge Base article 122920 documents these limits for Windows 2000, and Microsoft Knowledge Base article 314882 describes the limits in Windows XP.

Installation options

The following list explains the different VFM configurations that you can set up.

Note: VFM components must be installed on local fixed NTFS disks.

- **Client and server**. This is the typical configuration. The VFM components are installed on a single system.

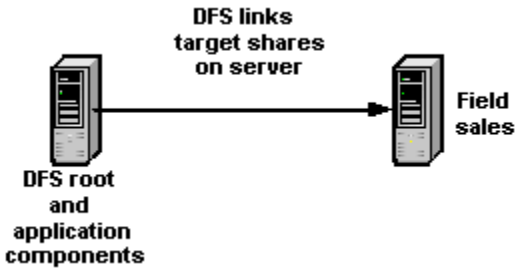

Figure 4 Typical Configuration

- **Client**. You are prompted to enter the location of the remote VFM server. This choice is useful if you installed the VFM server on an infrastructure machine and want access to the VFM console from your desktop machine.
- **Server**. This choice is useful if you want to install the VFM server on a highly available machine in your server room and manage VFM from your administrative desktop machine.
- **Monitoring agent**. You are prompted to enter the location of the remote VFM server. This choice is useful if you want to distribute VFM policy monitoring chores across

multiple machines. If the information (hostname) of the remote VFM server is not entered accurately, the monitoring agent cannot communicate with the server.

CAUTION:

- Never rename a machine while the VFM server,replication agent, or monitoring agent is installed on it. This can cause problems managing agents.
- To rename the machine where the VFM server or monitoring agent is installed:
  1. Export your VFM configuration files.
  2. Uninstall VFM.
  3. Rename the machine.
  4. Reinstall VFM.
  5. Import the configuration files from Step 1.

See "Export configuration files" in the VFM console online help for detailed instructions.

Upgrading from VFM 4.0

- If you generated reports, the reports cannot upgrade to a later version of VFM because the reporting feature changed in version 5.0. You will have to generate new reports after you upgrade.
- The following disaster recovery and replication sample batch files distributed with VFM were updated in version 5.0. If you are using any of these scripts, review the information about batch files in the online help before you run them after you upgrade.

Disaster Recovery

- PostEmail.bat
- PostEventlog.bat
- PostSnmpTrap.bat
- PreEmail.bat
- PreEventlog.bat
- PreSnmpTrap.bat
- PreStopServices.bat
- PreSystemsUp.bat

Replication

- PostEmail.bat
- PostEventLog.bat
- PostSnmpTrap.bat
- PostStartService.bat

To install VFM

VFM is available from the web or a CD. Select one of the following methods:

VFM web executable

1. Close all programs.
2. Double-click the VFM executable or zip file you downloaded from the web.

VFM CD

1. Close all programs.
2. Insert the VFM CD into your CD-ROM drive. If Autorun is enabled on your system, the installation starts automatically when you shut the drive. You can skip Steps 3 and 4.
3. From the **Start** menu, click **Run**.
4. Type **D:\Autorun** (substitute the appropriate letter of your CD-ROM drive for D).

Note: Before you uninstall VFM, you must first uninstall the monitoring agents you deployed to other servers because uninstalling VFM does not remove deployed monitoring agents.

Create a logical namespace

The Microsoft Distributed File System (DFS) is a storage service that can help solve many complex problems in the Windows environment. Enterprises that deploy DFS can achieve immediate and significant file management benefits. In addition to making it easier to manage files, DFS provides a core set of capabilities on which you can build storage solutions.

NetApp VFM is software for virtualizing and managing large-scale file storage in Windows environments. Built on Microsoft DFS, VFM is a solutions platform that enables the integrated management of logical and physical storage elements.

Benefits of using DFS

In creating DFS, Microsoft provided a robust enterprise storage solutions platform. The following list describes the benefits of using DFS:

- **Unified namespace**

  A DFS namespace eliminates machine name dependencies. You can develop a namespace that links to shares on multiple servers to create a hierarchical structure that behaves like a single hard disk. Users do not need to know the names of servers and shares to navigate the namespace.

- **Location transparency**

  DFS simplifies data migration. Because users are shielded from knowing the physical location of servers and shares, you can move data to another server without having to re-configure applications and shortcuts, and without having to inform users of the new location where they can find their data.

- **Storage scalability**

  You can add servers and present the storage on the new servers as new folders within your namespace.

- **Load sharing**

  DFS enables you to associate a logical name in your namespace with multiple shared folders on different servers, even when the servers are in different sites. For example, the share *nuview.com**salesInfo* is used heavily. You can use DFS to share the network load by associating the logical name *SalesInfo* with the following servers:

  \\nuview.com\salesInfo

  \\priam.com\salesInfo

  \\stargate\salesInfo

- **Intelligent client caching**

  When a user requests access to a target that is a part of a DFS namespace, a referral containing the information of the target is cached on the client. The next time the client requires access to the target, the client uses the cached referral.

- **Simplified maintenance**

  When multiple shares on different servers are associated with a single logical name on your namespace, you can perform preventive maintenance, repairs, or upgrades on servers by disabling referrals to specific shares. While the referral to a share is disabled, DFS automatically routes new requests to the remaining shares that are online.

- **Security integration**

  Your DFS namespace does not require additional security. Share access through the namespace is determined by existing NTFS and share permissions. For example, users can only access shares for which they have NTFS or share permissions.

Note: Before creating your first namespace, review the information in the topic "Namespace design considerations."

Create a logical namespace checklist

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

1. Create a DFS root.
2. Add links to the DFS root.
3. Check the results.

Step 1 Create a DFS root

The first step in creating a logical namespace is to create a DFS root. Here you create a stand-alone DFS root on a Windows 2000 or Windows Server 2003.

1. On the toolbar, click **Create a DFS Root**. The New DFS Root Wizard appears.
2. Click **Next**.
3. Select **Create a stand-alone DFS Root** and click **Next**.

4. Type the name of the server that hosts the root or click **Browse** to locate the server in your network. The name must be preceded by two back slashes "\\".
5. Click **Next**.
6. Select **Use an existing share** and select a share on the server.
7. Click **Next**.
8. If desired, type descriptive text in the **Comment** field.
9. Click **Next**.
10. Click **Finish**. The root appears under the Logical View tree.

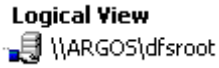

Step 2 Add links

After creating a root to host the namespace, the next step is to populate the namespace by adding links to it.

1. Drag a network share (UNC path) from Windows Explorer to the DFS root you just created.
2. Repeat Step 1 to add additional links. Single target links appear under the root.

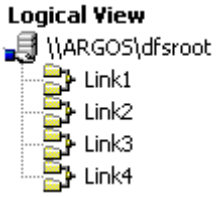

Step 3 Check the results

To check how the namespace you are creating appears to your users, you can use either the VFM Explorer tab or Windows Explorer.

To check the namespace in VFM:

1. Select the root in the Logical View.
2. In the right pane, click the **Explorer** tab.

To check the namespace in Windows Explorer:

Use Windows Explorer to browse the root you created. The links you created appear as folders in Windows Explorer.

Quickly populate the namespace

VFM namespace creation policies and permission based namespace policies enable you to easily populate and maintain a logical namespace. You can use policy features interactively to perform the initial population of your namespace. After the namespace is in place, the policies can automatically run on a scheduled basis to scan for changes in the physical resources supporting the logical namespace. You can configure policies to report changes, or to update the logical namespace automatically to reflect the changes detected by the policy.

Difference between namespace creation and permission based namespace policies

Namespace creation and permission based policies work the same except for the following features:

- Namespace creation policies enable you to search for shares, volumes, and qtrees that match the search string, whereas permission based namespace policies enable you to search for shares that match the search string.
- The permission based namespace policies enable you to filter your search by group permissions and search string, whereas namespace creation policies enable you to filter your search by search string.

Quickly populate the namespace checklist

Select the policy you want to create to populate your namespace. A wizard guides you through the steps to configure the policies:

- Configure a namespace creation policy.
- Configure a permission based namespace policy.

Configure a namespace creation policy

After creating a namespace, configure a namespace creation policy to quickly populate it. You configure the policy to find shares that match your search criteria and to automatically populate the namespace with links that reference the shares.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Creation Policy**.
2. Click **Next**.
3. Type a word or phrase that best describes this policy in the **Name** field, and then click **Next**.
4. Select a DFS root and click **Next**.
5. Click **New Folder** and type a meaningful name for the folder in the **Name** field, and then click **Next**.
6. Select the domain or servers where you want to search for shares, and then click **Next**.
7. Type the search string. You can use the DOS wildcards "*" or "?". You can combine search strings by separating them with semi-colons. For example, entering mkt*;sales* finds all shares on the selected domains or servers whose names begin with either mkt or

sales.

8. Select **Search for shares matching the search string**, and then click **Next**. A list of shares that match your search pattern appears.
9. Clear the items that you do not want to add to the namespace, and then click **Next**.
10. Select a template to arrange how links appear in the namespace. VFM enables you to preview the namespace.
11. Click **Next**.
12. Verify the changes to the namespace, and then click **Next**.
13. Click **New** and configure a schedule to run the policy, and then click **Next**.
14. Click **Finish**. The new policy appears under the Namespace Creation Policies folder.

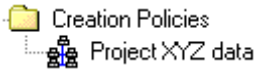

Configure a permission based policy

After creating a namespace, configure a permission based policy to populate the namespace based on security permissions. You configure the policy to find shares that match your search criteria and to automatically populate the namespace with links that reference the shares.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Permission Based Policy**.
2. Click **Next**.
3. Type a word or phrase that best describes this policy in the **Name** field, and then click **Next**.
4. Select a DFS root and click **Next**.
5. Click **New Folder** and type a meaningful name for the folder in the **Name** field, and then click **Next**.
6. Select the domain or servers where you want to search for shares, and then click **Next**.
7. Click **Add** to browse for security groups which have access to resources on the domain or servers you selected.
8. Type the search string. You can use the DOS wildcards "*" or "?". You can combine search strings by separating them with semi-colons. For example, entering mkt*;sales* finds all shares on the selected domains or servers whose names begin with either mkt or sales.
9. Select **Search for shares matching the search string** and then click **Next**. A list of shares that match your search pattern appears.
10. Clear the items that you do not want to add to the namespace, and then click **Next**.
11. Select a template to arrange how links appear in the namespace. VFM enables you to preview the namespace.
12. Click **Next**.
13. Verify the changes to the namespace, and then click **Next**.
14. Click **New** and configure a schedule to run the policy, and then click **Next**.
15. Click **Finish**. The new policy appears under the Permission Based Policies folder.

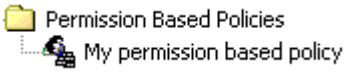

Make the namespace highly available

Namespace availability policies enable you to easily keep the logical structure of two or more DFS roots or a DFS root and a widelink namespace aligned. When you synchronize roots, you can load balance the work on your DFS root server across multiple identical instances of the root. This also provides fault tolerance because the namespace is available from more than one server. When you synchronize a root and widelink namespace, your CIFS clients can access the namespace on a NetApp filer.

- **What root synchronization is**

  Root synchronization involves keeping the logical structure of two or more DFS roots or a DFS root and widelink namespaces aligned. The layouts of the secondary roots and widelink namespaces are made identical to those of the primary or master root.
- **What a namespace availability policy is**

  By configuring a namespace availability policy, you can have VFM automatically scan a set of roots and apply changes to the logical structure of the secondary roots and widelink namespaces to keep them aligned with the logical structure of the primary root.
- **Benefits of synchronizing roots and widelink namespaces**
  - Improves namespace availability and facilitates load balancing of the DFS root.
  - Allows the creation of a DFS root whose replicas reside in different domains.
  - Can be used in conjunction with DNS configuration to provide locality of reference benefits to users accessing the DFS roots in the policy.
  - Enables CIFS clients to access the namespace on a filer.

Create a namespace availability policy checklist

To synchronize the logical structure of two DFS roots, follow the basic steps listed below:

1. Make DFS roots highly available.
2. Schedule synchronization.

To synchronize the logical structure of a DFS root and a widelink namespace, follow the basic steps listed below:

1. Create a namespace on a filer.
2. Schedule synchronization.

Step 1 Make DFS roots highly available

Create a namespace availability policy and add at least two DFS roots to the policy.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Availability Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Click **Add** and type the path of a root, or click **Browse** to locate a root in your network.
4. Repeat Step 3 to add another root.
5. Select the master root. This is the root against whose logical structure the other roots in the policy are compared. If automatic synchronization is enabled for the policy, this is the root whose logical structure is synchronized with the other roots in the policy.
6. Click **OK**. A new policy appears under the Namespace Availability Policies folder.

Step 1 Create a namespace on a filer

Create a widelink namespace contained in a CIFS share that is hosted on a NetApp filer.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Availability Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Click **Add** and type the path of a root, or click **Browse** to locate a root in your network.
4. Select the root. This is the root against whose logical structure the widelink namespace in the policy is compared. If automatic synchronization is enabled for the policy, this is the root whose logical structure is synchronized with the widelink namespace in the policy.
5. Click **Add** in the **Synchronize the following NetApp filer shares** field and type the path of a CIFS share, or click **Browse** to locate a CIFS share in your network.
6. Click **OK**. A new policy appears under the Namespace Availability Policies folder. The master root and the widelink namespace appear under the policy.

Step 2 Schedule synchronization

The next step is to make sure that changes made to the logical structure of the master root are synchronized with the secondary root or widelink namespace. You want to schedule how frequently to synchronize the logical structure.

1. Right-click the namespace availability policy and click **Properties**.
2. Click **Schedule**.
3. Select the schedule option you want.
4. Click **OK**.

Make data highly available

VFM provides three types of disaster recovery policies. Select the policy that meets your needs.

Select the policy that best meets your needs.

- **Client disaster recovery policies** enable you to set up VFM to detect a failure and automatically fail over clients to a replicated storage device. Use this policy when data is scattered among multiple servers.
    - This policy provides the greatest flexibility. It enables you to control how monitoring and failover takes place.
    - Links contained in the policy must have a least two replicated targets and might contain multiple offlline targets.
    - One of the most significant advantages of this policy is the ability to designate a cascaded failover path. A typical scenario is one in which you designate critical data to fail over first to a local edge device and then to a remote central device. This ensures high availability of the data whether there is a local server failure or an entire site outage. Failover takes place without changing login scripts or network configuration.
- **Filer disaster recovery policies** enable you to make data on NetApp filers participating in your namespace highly available.
    - VFM makes it easy to set up a filer or nearline storage device as a disaster recovery backup for one or more primary filers.
    - VFM leverages SnapMirror replication to replicate data to a standby volume or qtree. When a failure on the primary storage is detected, an automatic failover to the standby storage is initiated. VFM periodically checks for new links and qtrees on the primary filer and creates the new data configuration on the backup filer.
- **Server disaster recovery policies** enable you to make data on servers participating in your namespace highly available. Use this policy when data on a primary server must be made highly available. When a failure of the primary server is detected, VFM automatically fails over all links in the policy to a secondary server.
    - This policy enables you to configure a server as a disaster recovery backup for one or more primary servers.
- VFM makes it easy to monitor for new links targeting the primary server to add to the policy and to replicate data to a secondary server.

Create a client disaster recovery policy

To create your first client disaster recovery policy, follow the steps listed below:

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

Note: Review the topic "Disaster recovery planning considerations" before you create your first disaster recovery policy.

Step 1 Create a client disaster recovery policy

Create a client disaster recovery policy to make data highly available.

1. On the **Tools** menu, click **New Policy**, point to **Disaster Recovery Policy**, and then click **Client Recovery Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Select the VFM Replication method.
4. Click **OK**. The new policy appears under the **Client Recovery Policies** folder.

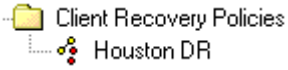

Step 2 Add a link to the client disaster recovery policy

After you create a client disaster recovery policy, the next step is to add links with targets.

1. Double-click the Disaster Recovery Policies folder in the Admin View, and then double-click the Client Recovery Policies folder to display the policy you created.
2. Drag a link from the Logical View to the policy you created. The **Disaster Recovery Link Target Wizard** appears, guiding you through adding a second target to the link to make the data highly available. After you complete the wizard, a status report appears, detailing the operations carried out.
3. Click **Done**. A link with multiple targets appears under the policy.

Note: Links on DFS roots hosted on Windows NT 4 systems cannot be members of a client disaster recovery policy because VFM must be able to control storage target state for the links.

Step 3 Configure monitoring of the client disaster recovery policy

After you add a link to the client recovery policy, the next step is to configure status monitoring for the link so that VFM can fail over to the alternate target should there be a problem with the online target. Schedule status monitoring of the link targets in the client disaster recovery policy.

1. Right-click the client disaster recovery policy and click **Properties**.
2. Click **Monitoring Schedule**.
3. Schedule status monitoring to begin at a specific time and date and to repeat continually at a stated frequency.
4. Click **OK**.

Step 4 Configure replication of the client disaster recovery policy

The last step is to make sure that data changes made to the online target are synchronized with the offline targets. You want to schedule how frequently you want to replicate data.

1. Right-click the client disaster recovery policy and click **Properties**.
2. Click **Replication Schedule**.
3. Select **Use Schedule**.
4. Select the schedule option you want.
5. Click **OK**.

Create a filer disaster recovery policy

The VFM Filer Recovery Policy Creation wizard guides you through creating a filer disaster recovery policy.

Create a filer disaster recovery policy

Create a filer disaster recovery policy to make data on NetApp filers highly available.

1. On the **Tools** menu, click **New Policy**, point to **Disaster Recovery Policy**, and then click **Filer Recovery Policy**.
2. Click **Next**.
3. Type a word or phrase that describes the policy in the **Name** field.
4. Type the name of the filer that will contain the SnapMirror destination or click **Browse** to select the filer.
5. Click **Next**. A list of DFS roots that contains links referencing volumes and qtrees on filers appears.
6. Click **Next**.
7. Click **Add** to select a volume containing the source data that is referenced by a link in your namespace.
8. Click **OK**.
9. Select a source qtree on a filer that contains data referenced by a link in your namespace and click **Next**.

10. If a destination does not exist on the destination filer for the source qtree, click the row and select **Add New SnapMirror**.
11. Create a qtree on the destination filer.
12. Click **OK**. VFM creates the SnapMirror replication.
13. Click **Done**.
14. Click **Next**.
15. Select **Find new links** and schedule how often you want VFM to monitor for new links.
16. Click **Next**.
17. Schedule how often you want VFM to monitor the source link target. When the target becomes unavailable, VFM fails over to the destination link target. Click **Next**.
18. Click **Next**.
19. Click **Finish**. The filer disaster recovery policy appears under the Disaster Recovery Policy\Filer Recovery Policies folder.

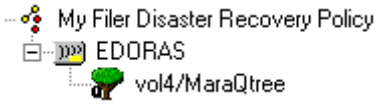

Create a server disaster recovery policy

To create your first server disaster recovery policy, follow the steps listed below:

1. Create a server disaster recovery policy.
2. Select the namespaces to monitor.
3. Configure monitoring for new resources.
4. Configure monitoring of the server disaster recovery policy.
5. Configure replication of the server disaster recovery policy.

Note: Review the topic "Disaster recovery planning considerations" before you create your first disaster recovery policy.

Step 1 Create a server disaster recovery policy

Create a server disaster recovery policy to make data on a server highly available.

1. On the **Tools** menu, click **New Policy**, point to **Disaster Recovery Policy**, and then click **Server Recovery Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Type the name of the primary server.
4. Type the name of the secondary server.
5. Click **OK**. The new policy appears under the **Server Recovery Policies** folder.

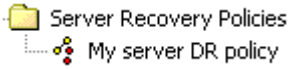

Step 2 Select the namespaces to monitor

After you create the server disaster recovery policy, the next step is to select the namespaces you want VFM to monitor for links targeting the primary server.

1. Right-click the server disaster recovery policy and click **Properties**.
2. Click **Resource Monitoring Options**.
3. Click **Add All Roots From the Logical View**.
4. Click **OK**.

Step 3 Configure monitoring for new resources

After you select the namespaces, the next step is to monitor the namespaces for new links targeting the primary server.

1. Right-click the server disaster recovery policy and click **Properties**.
2. Click **Resource Monitoring Schedule**.
3. Schedule resource monitoring to begin at a specific time and date and to repeat continually at a stated frequency.
4. Click **OK**. New links found on the primary server are listed under New Resources on the Properties tab. To show the Properties tab, select the policy.

Step 4 Configure monitoring of the server disaster recovery policy

After you configure monitoring for new resources on the selected namespaces, the next step is to configure status monitoring of the policy. VFM randomly selects one of the links that is a member of the policy to check the availability of the active target of the link. When the active target (a share on the primary server) is not available, VFM fails over all the links in the policy.

1. Right-click the policy and click **Properties**.
2. Click **Monitoring Schedule**.
3. Schedule status monitoring to begin at a specific time and date and to repeat continually at a stated frequency.
4. Click **OK**.

Step 5 Configure replication of the server disaster recovery policy

The last step is to make sure that data changes made to the active targets on the primary server are synchronized with the offline targets on the secondary server. You want to schedule how frequently you want to replicate data.

1. Right-click the server disaster recovery policy and click **Properties**.
2. Click **Replication Schedule**.
3. Select **Use Schedule**.
4. Select the schedule option you want.
5. Click **OK**.

Optimize storage

Use archival migration and storage load balancing policies to optimize your storage.

- **Archival migration policies** enable you to set policies to control when data is migrated to secondary storage. You control whether VFM automatically carries out the migration or simply notifies you of items matching your criteria so that you can migrate them later. Archival migration policies enable you to transfer CIFS and NFS data. Data transfer must be between the same protocol (for example, CIFS to CIFS and NFS to NFS)
- **Storage load balancing policies** enable you to set policies to evaluate the current size of disk space of storage locations you specify. You control whether VFM automatically load balances or provides migration suggestions so that you can decide what directories you want to migrate.

Create an archival migration policy

VFM enables you to transfer CIFS or NFS directories that match the selection criteria you specify in an archival migration policy.

- Data transfer must occur within the same protocol (for example, CIFS to CIFS or NFS to NFS).
- A policy can contain CIFS and NFS source and destination pairs.

Transfer CIFS directories

- You cannot configure more than one search path with the same source. This rule is illustrated in the following table:

Steps	Source	Destination
1. Search path A is configured.	\\serverA\shareA	\\serverB\shareB
2. Search path B is configured.	\\serverA\shareA	\\serverC\shareC
3. VFM replaces search path A with B	~~\\serverA\shareA~~ \\serverA\shareA	~~\\serverB\shareB~~ \\serverC\shareC

- The source can be a physical (for example, \\server\share\subdir) or logical path (for example, a path in the namespace).

- Links selected as the source must only have one target.
- The destination must be a physical path.
- When the source is a physical path, VFM scans for subdirectories. The shared folder named in the source path must contain subdirectories because VFM scans for directories under the share as possible migration candidates. If the directories match the selection criteria, each subdirectory is listed as a migration candidate.
- When the source is a logical path, VFM scans the target of a link. If the link target matches the selection criteria, it is listed as a migration candidate.

Transfer CIFS directories checklist

To create your first archival migration policy to transfer CIFS directories, follow the steps listed below:

1. Create an archival migration policy.
2. Specify the selection criteria.
3. Specify the migration actions.
4. Update namespace options.
5. Specify migration task options.
6. Schedule scanning.
7. Schedule migration.

Step 1 Create an archival migration policy

Create an archival migration policy to automatically migrate CIFS data from primary to secondary storage. Configure the policy with the criteria that you want VFM to use to select directories to migrate.

1. On the **Tools** menu, point to **New Policy**, point to **Data Lifecycle Automation Policy**, and then click **Archival Migration Policy**.
2. Type a word or phrase in the **Name** field that best describes the policy.
3. Click **Add** to include a shared folder with subdirectories as the source and a shared folder on another machine as the destination.
4. Click **OK**. The policy and the source shared folder appear under the Data Lifecycle Automation Policies\Archival Migration Policies folder.

   The following example archival migration policy shows three source candidates. The first source is a logical folder in the namespace, the second source is a folder that does not participate in the namespace, and the last source is a shared folder that participates in the namespace. All three are valid sources for an archival migration policy.

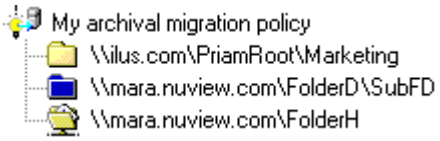

Step 2 Specify the selection criteria

The next step is to specify the selection criteria. In this example, you want to automatically migrate a directory to secondary storage when 85% of the files in the directory have not been accessed within the last 30 days.

1. Right-click the policy and click **Properties**.
2. Click **Migration Criteria**.
3. Select **Migrate directories when**.
4. Select **% of files have not been accessed within the last....**
5. Change the percent to "85."
6. Change the number of days to "30."
7. Click **Apply**.

Step 3 Specify the migration actions

After you specify the selection criteria, the next step is to define the migration operation. VFM provides a list of migration "actions" from which you can choose.

1. Click the **Actions** tab.
2. Make sure that "**Copy the data**" is selected.
3. Make sure that "**Update namespace**" is selected. When this option is selected, VFM only migrates directories referenced by links in the namespace.
4. Click **Apply**.

Step 4 Update namespace options

The next step is to specify how you want to update the namespace after the data is migrated.

1. Click the **Update Namespace Options** tab.
2. Make sure that "Retarget links" is selected. VFM retargets the links after migration to reference the new location. Users do not have to be told to access a new location.
3. Click **Apply**.

Step 5 Specify migration task options

After selecting the migration actions, the next step is to specify the tasks that you want the policy to execute.

1. Click **Task Options**.
2. Make sure that "Automatically add migration tasks for migration candidates" is selected. VFM creates migration tasks after scanning is complete.
3. Make sure that all the roots in your namespace are listed in the dialog box so that VFM can update links to reference the new location.
4. Click **Apply**.

Step 6 Schedule scanning

The next step is to schedule how frequently you want the policy to scan for migration candidates.

1. Click **Use policy scan schedule**.
2. Click **New** and configure the schedule.
3. Click **Apply**.

Note: If no migration candidates are found when the policy scans for directories, running the migration will have no effect.

Step 7 Schedule migration

The last step is to schedule how frequently you want VFM to migrate directories that match the selection criteria.

1. Click **Migration Schedule**. By default, **Use migration schedule** is selected.
2. Click **New** and configure the schedule.
3. Click **OK**.

Transfer NFS directories

- The source and destination must be UNIX fully qualified paths (for example, <machine>:/<UNIX local path>).
- NFS 3 or later 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.
    - See "Specify default replication agents" for instructions about how to specify default Windows and UNIX proxy machines.
    - See "Specify a proxy machine" for instructions about how to specify a Windows or UNIX proxy machine.
- The source and destination can be an export or a folder under an export. VFM scans for subdirectories. The export named in the source path must contain subdirectories because VFM scans for directories under the export as possible migration candidates. If the directories match the selection criteria, each subdirectory is listed as a migration candidate.
- You cannot configure more than one search path with the same source.

Steps	Source	Destination
1. Search path A is configured.	hostname:/root/subdirectoryA	hostname:/root/subdirectoryB
2. Search path B is configured	hostname:/root/subdirectoryA	\hostname:/root/subdirectoryC

Steps	Source	Destination
.		
3. VFM replaces search path A with B	~~hostname:/root/subdirectoryA~~ \hostname:/root/subdirectoryA	~~hostname:/root/subdirectoryB~~ \\hostname:/root/subdirectoryC

Transfer NFS exports checklist

To create your first archival migration policy to transfer NFS directories, follow the steps listed below:

1. Create an archival migration policy.
2. Specify the selection criteria.
3. Specify the migration actions.
4. Update namespace options.
5. Specify migration tasks options.
6. Schedule scanning and migration.

Step 1 Create an archival migration policy

Create an archival migration policy to automatically migrate NFS data from primary to secondary storage. Configure the policy with the criteria that you want VFM to use to select directories to migrate.

1. On the **Tools** menu, point to **New Policy**, point to **Data Lifecycle Automation Policy**, and then click **Archival Migration Policy**.
2. Type a word or phrase in the **Name** field that best describes the policy.
3. Click **Add** to include an export with subdirectories as the source and an export on another machine as the destination.
4. Click **OK**. The policy and the source shared folder appear under the Data Lifecycle Automation Policies\Archival Migration Policies folder.

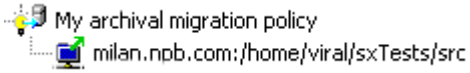

Step 2 Specify the selection criteria

The next step is to specify the selection criteria. In this example, you want to automatically migrate a directory to secondary storage the directory contains more than 200 files.

1. Right-click the policy and click **Properties**.
2. Click **Migration Criteria**.
3. Select **Migrate directories when**.

4. Select **Directory contains more than...files**.
5. Change the file number to "200."
6. Click **Apply**.

Step 3 Specify the migration actions

After you specify the selection criteria, the next step is to define the migration operation. VFM provides a list of migration "actions" from which you can choose.

1. Click the **Actions** tab.
2. Make sure that "**Copy the data**" is selected.
3. Make sure that "**Update namespace**" is selected. When this option is selected, VFM updates references in automounts you specify.
4. Click **Apply**.

Step 4 Update namespace options

The next step is to list the automount files you want to update after the data is migrated.

1. Click the **Update Namespace Options** tab.
2. Click **Add**.
3. Type the UNIX fully qualified path of an automount file.
4. Repeat Step 3 to add more automount files.
5. Click **Apply**.

Step 5 Specify migration task options

After selecting the migration actions, the next step is to specify the tasks that you want the policy to execute.

1. Click **Task Options**.
2. Make sure that "Automatically add migration tasks for migration candidates" is selected. VFM creates migration tasks after scanning is complete.
3. Click **Apply**.

Step 6 Schedule scanning and migration

The last two steps are to configure schedules to scan for directories and migrate the directories found. See the following topics under "Transfer CIFS directories." for instructions:

- Schedule scanning.
- Schedule migration.

Create a storage load balancing policy

To create your first storage load balancing policy, follow the steps listed below:

1. Create a storage load balancing policy.
2. Add storage locations.
3. Schedule scanning.
4. Schedule migration.

Step 1 Create a storage load balancing policy

Create a storage load balancing policy to automatically load balance storage locations you specify in the policy. Configure the policy with the criteria that you want VFM to use to evaluate the current size of disk space and migrate directories under the storage locations that are link targets.

1. On the **Tools** menu, point to **New Policy**, point to **Data Movement Policy**, and then click **Storage Load Balancing Policy**.
2. Type a word or phrase in the **Name** field that best describes the policy.
3. Click **Add,** select a DFS root in your namespace that contains links targeting directories you want to load balance, and click **OK**.
4. Click **Apply**.

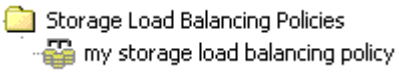

Step 2 Add storage locations

The next step is to add storage locations you want to load balance.

1. Click **Pool**. A new dialog box appears on the right pane.
2. Click **Add** and select a share from the Physical View or your network.
3. Repeat Step 2 to add another share.
4. Click **Apply**.

Note: At least one of the shares must be a link target or must contain a directory that is a link target.

Step 3 Schedule scanning

The next step is to schedule how frequently you want the policy to scan the current size of disk space of the storage locations, to determine whether load balancing is required.

1. Click **Policy Scan Schedule**.
2. Click **Add** and configure the schedule.
3. Click **Apply**.

Step 4 Schedule migration

The last step is to schedule how frequently you want VFM to migrate directories referenced by links in your namespace to load balance the storage pool.

1. Click **Migration Schedule**.
2. Click **Add** and configure the schedule.
3. Click **OK**.

Transfer data

VFM provides multiple data transfer policies. Select the policy that meets your needs.

- **Migration policies** enable you to balance free space and data throughput capacity across volumes, or 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 for your environment. VFM provides status information about replication operations performed using the VFM replication agent. Use replication policies to replicates 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."
- **SnapMirror policies** enable you to create and centrally manage replication between NetApp filers in your network. You can integrate SnapMirror data replication with your namespace. VFM provides a central view of events from all managed NetApp filers. You receive detailed reporting of filer configuration and status information.

Replication policies

VFM can help you easily and effectively manage replication for storage policies containing heterogeneous devices.

- **Data replication across heterogeneous devices**

  You can select the optimal replication mechanism for your configuration, whether this is VFM for heterogeneous environments, SnapMirror for NetApp filers, or another mechanism for homogeneous devices.

- **Status information about replication operations**

  VFM provides status information about replication operations performed using the VFM replication agent.

- **Allows horizontal scaling and load balancing**

  VFM provides you with tools to migrate data and namespace information in a single operation. These tools give you coordinated management of the logical and physical aspects of storage.

- **Causes no disruption to user access to data**

  The logical namespace shields users from changes in the underlying infrastructure. This simplifies data access for users and storage management for you.

- **Transfer CIFS or NFS data**

  Data transfer must occur within the same protocol (for example CIFS to CIFS or NFS to NFS). When transferring NFS data, the source and destination must be running NFS 3 or higher.

Select the data to distribute

Select which type of data you want to distribute:

- See "Distribute CIFS data."
- See "Distribute NFS data."

Distribute CIFS data

- CIFS physical paths must be UNC paths (for example, \\<machine>\<share>).
- 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 replication policies that have a DFS root or another link as their target.

Create a replication policy checklist

To create your first replication policy to distribute CIFS data, follow the steps listed below:

1. Create a replication policy.
2. Add the source.
3. Add a destination.
4. Configure replication.

Step 1 Create a replication policy

Add a replication policy to distribute CIFS data.

1. On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and then 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.
3. Click **OK**.

4. Click **Done**. The new policy appears under the **Replication Policies** folder.

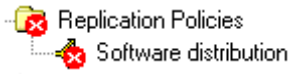

Note: The policy shows an error status because source or targets are not added to the policy.

Step 2 Add the source

The next step is to add a source from which data is replicated.

1. Drag a share under a machine that appears in the Physical View to the replication policy. The option **As the source path** is selected.
2. Click **OK**.

Step 3 Add a destination

After adding the source, the next step is to add a destination that receives the data being replicated.

1. Drag a share under a machine in the Physical View to the replication policy.
2. Select **Add it to the list of destination paths**.
3. Click **OK**. The error status on the policy is cleared because you added source and target to it. You can click the Replication Topology tab in the right pane to see a graphical view of how data is distributed by the policy.

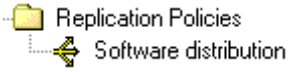

Note: The destination might be a link that references a single or multiple CIFS targets or a CIFS physical path.

Step 4 Configure replication

The last step is to schedule how frequently you want the replication to take place.

1. Right-click the replication policy and click **Properties**.
2. Click **Replication Schedule**. By default, **Use replication schedule** is selected.
3. Click **New** and configure the schedule.
4. Click **OK**.

Distribute NFS data

Before you create a replication policy to distribute NFS data, refer to the requirements in "Transfer NFS directories" in the section that discusses creating an archival migration policy.

Create a replication policy to distribute NFS data checklist

1. Create a replication policy.
2. Add a source. See the topic of the same name under "Distribute CIFS data."
3. Add a destination. See the topic of the same name under "Distribute CIFS data."
4. Configure replication. See the topic of the same name under "Distribute CIFS data."

Create a replication policy

Add a replication policy to distribute NFS data.

1. On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and then 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.
3. In the **Source** field, type the physical path of an NFS export or click **Browse** to locate an export in your network.
4. In the **Targets** field, Click **Add** and type the physical path of an NFS export or click **Browse** to locate an export in your network.
5. Click **OK**. The new policy appears under the **Replication Policies** folder

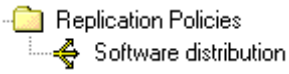

Note: VFM enables you to browse the following machines for exports:

- Machines that have UNIX replication agents installed
- Machines running the CIFS and NFS protocols
- Machines that you add to the Physical View

Create a migration policy

VFM migration policies provide powerful tools for use in migrating CIFS and NFS data. With this feature, you can perform the following tasks:

- Balance free space and data throughput capacity across volumes.
- Increase reliability and availability by creating alternate copies of data.
- Migrate data prior to maintaining or consolidating servers.

A migration policy can contain multiple tasks that identify source and destination pairs. A source and destination pair must be CIFS to CIFS or NFS to NFS. If the source is a CIFS shared folder,

the destination must be a CIFS shared folder, and if the source is an NFS export, the destination must be an NFS export. A migration policy can contain CIFS and NFS source and destination pairs.

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 NuView for more information about this product.

Migration policy checklist

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

1. Create a migration policy.
2. Select migration actions to define the migration operation.
3. Update namespace options.
4. Schedule when and how frequently you want to run the policy migration tasks.
5. Add a task to the policy that identifies source and destination paths. A policy can have multiple tasks.

Step 1 Create a migration policy

Create a migration policy to move data referenced by links in your namespace to a new server.

1. On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and then click **Migration Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Click **OK**. A new policy appears under the Migration Policies folder.

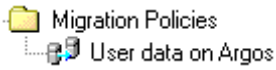

Step 2 Select migration actions

After you create a policy, the next step is to define the migration operation. VFM provides a list of migration "actions" from which you can choose.

1. Right-click the policy you created and click **Properties**.
2. Click **Actions**.
3. Make sure that "Copy the data" and "Update namespace" are selected.

Step 3 Update namespace options

The next step is to specify how you want to update the namespace after the data is migrated.

- When migrating CIFS data, make sure that "Retarget links" is selected.
- When migrating NFS data, click **Add** to list the automount files you want to update.

Step 4 Schedule migration

The next step is to schedule how frequently you want to run the policy to migrate data.

1. Click **Migration Schedule**. By default, **Use migration schedule** is selected.
2. Click **New** and configure the schedule.
3. Click **OK**.

Step 5 Add a task to the migration policy

After you create a migration policy, define the migration operation the policy executes, and specify how frequently you want the policy to run, the last step is to add tasks to the policy that identify source and destination paths. VFM provides a wizard that steps you through the process.

Notes:

- Data transfer must occur within the same protocol (for example, CIFS to CIFS or NFS to NFS).
- A migration task can only contain CIFS source and destinations or NFS source and destinations, but not both protocols.
- A migration policy can contain both CIFS and NFS migration tasks.

Migration task to transfer CIFS data

The CIFS source and destination must be UNC paths (for example, \\<machine>\<share>).

1. Drag the machine containing the source shares from the Namespace Resources folder in the Physical View to the migration policy you created. The Migration Task wizard appears.
2. Type the name of a machine containing CIFS data in the **Destination Path** field and click **Next**.
3. Select **Only migrate paths referenced by the logical namespace** and click **Next**. A list of source and destination folders appears.
4. Click **Next**. The policy checks for destination shares that do not exist.
5. If the shares do not exist on the destination machine, VFM shows you a list of the destination shares and suggests a local path for each share. Select all the items with no destination share and click **Create Destination Share**. VFM creates the shares on the destination machine.
6. Click **Next**. The policy warns you if there are destination shares that are not empty. You can type a new destination or click **Next**. An updated list of links appears.

7. Validate the list of links and click **Next**.
8. Click **Finish**. **All Tasks** and **Scheduled Tasks** folders appear under the policy you created. The task appears under the folders with a "clock" status indicating that the migration operation is scheduled.

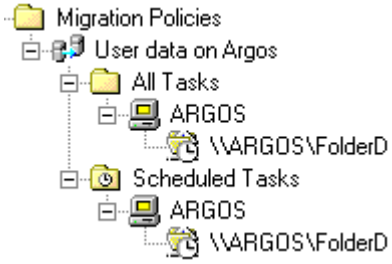

Migration task to transfer NFS data

1. Right-click the migration policy and click **Add task to policy**.
2. In the **Original Path** field, type the UNIX fully qualified physical path of an export, or click **Browse** to locate an export.
3. In the **Destination Path** field, type the UNIX fully qualified physical path of an export, or click Browse to locate an export.
4. Click **Next**. The original and destination paths appear in a list. You have the option of changing the destination path.
5. Click **Next**.
6. If the export does not exist on the destination machine, the policy shows you the destination export and suggests a local path for the export point. Select the export point and click **Create Destination Export**. VFM creates the exports on the destination machine.
7. Click **Next**.
8. Click **Finish.** The task appears under the folders with a "clock" status indicating that the migration operation is scheduled.

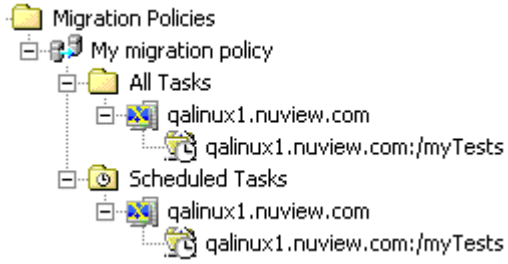

Notes:

- The NFS source and destination must be UNIX fully qualified paths (for example, <hostname>:/<root>/<subdirectory>).
- The machines hosting the source and destination must be running NFS 3 or higher.
- A replication agent must be installed on the source or destination or you can specify a default proxy.
    - See "Install a replication agent" for instructions about how to manually install a UNIX replication agent on the source, destination, or proxy machine.
    - See "Specify default replication agents" for instructions about how to specify default Windows and UNIX proxy machines.
    - See "Specify a proxy machine" for instructions about how to specify a Windows or UNIX proxy machine.

Configure SnapMirror

SnapMirror policies enable you to create and centrally manage replication between NetApp filers in your network.

Benefits of SnapMirror

SnapMirrors enable you to distribute, back up, and restore data on NetApp filers. The tasks for which SnapMirror is best suited are as follows:

- Recovering deleted items quickly
- Mirroring a volume on another NetApp filer
- Moving a volume among filers
- Rearranging data within a filer

Benefits of VFM SnapMirror Policies

- You can centrally manage all SnapMirror data replication on NetApp filers in your network.
- You can integrate SnapMirror data replication with your namespace.
- You have a central view of events from all managed NetApp filers.
- You receive detailed reporting of filer configuration and status information.

Note: Review the topic "SnapMirror planning considerations" before you create your first SnapMirror policy.

Create a SnapMirror policy checklist

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

1. Create a SnapMirror policy.
2. Configure replication.

Step 1 Create a SnapMirror policy

Create a SnapMirror policy to transfer data between volumes or qtrees on NetApp filers.

- Make sure that VFM can access the filers that are members of the policy. If you did not provide VFM with the system or administrative user account and password for each filer, you cannot manage the filers in your policy. To add a filer's account and password, first add the filer to the policy as stated below, right-click the filer, click **Properties**, and then click the **NetApp Filer** tab.
- You need to be able to access the *C$\etc* directory on the filer to read and write SnapMirror-related files. If the account running VFM does not have appropriate privileges to the filer shares, you must explicitly map to the share using a user name and password for a privileged account, and then add the filer to the SnapMirror policy.

1. On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and then click **SnapMirror Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Type the name of a filer that has SnapMirrors, or click ... **Browse** to locate a filer with SnapMirrors in the Physical View or network.
4. Click **OK**. The new policy appears under the **SnapMirror Policies** folder.

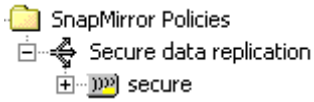

Step 2 Configure replication

After you create a SnapMirror policy and add a filer to the policy, the next step is to schedule replication on one of the snapmirrors on the filer. You want to make sure that data changes made to the source are synchronized with the read-only destination.

1. Right-click the policy and click **Properties**.
2. Click **SnapMirrors**.
3. Select a SnapMirror configuration and click **Modify**.
4. Click the **Options and Schedule** tab.
5. Select **Every Hour** on the **Repeat Mirror** list, assuming that you want to repeat data replication every hour.
6. Select **Quarter Past** on the **Start Mirror** list, assuming that you want to start replication 15 minutes after every hour.
7. Click **OK**.

Virtualize user home directories

VFM enables you to greatly simplify the management tasks associated with user home directory storage.

- **Centralized home folder management**

  Home directory management is centralized in the Admin View.
- **Simple home folder management**

  You can create different home folder policies to accommodate users from different domains and to place user home directories at different points in the namespace. Context menus provide easy access to all the tools necessary to manage policies of users.
- **Simple user data migration**

  VFM provides the Migration wizard to migrate links and underlying data in a single operation.
- **Customize user access to data**

  VFM enables you to group user home directories in folders. This enables you to individualize user access to data and provide access to data to multiple users in one easy step.
- **User database is kept synchronized with the storage configuration**

  VFM enables you to synchronize the logical namespace with home directory information contained in the Active Directory or NT4 domain controller. You can tell at a glance whether changes were made to the user accounts, such as adding, deleting, or changing user information. If problems are detected, VFM enables you to easily correct them.

Benefits of using VFM to create and manage home folders

- **Integrated management of user storage configuration**

  VFM manages user storage in logical namespace as well as in Active Directory and NT domain user information.
- **Virtualization of home directories**

  VFM adds user home directories to a logical namespace after querying Active Directory or an NT domain controller, thus virtualizing the user's storage.
  - By virtualizing home directories, you gain flexibility to make changes to physical storage without affecting users’ access to data.
  - By virtualizing home directories, you break the hard dependency between users and the location of data. By breaking this dependency, you can take advantage of the VFM flexibility and simple data access for users.
  - By virtualizing home directories, you no longer must update the Active Directory or Windows NT domain controller whenever the physical location of data is changed. You simply update the logical namespace to target the new location of the data.

How it works

The VFM Home Folder Policy Creation wizard guides you through the steps to creating a home folder policy.

Create a home folder policy

You can virtualize users' home folders by creating a home folder policy in VFM. This adds users' home directories to your namespace and updates your domain information (either in Active Directory or in an NT4 domain controller), so that users access their home directories through the logical namespace. When users access their home directories from the namespace, you can perform load balancing without affecting users and migration of user data is transparent to users.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Home Folder Policy**.
2. Click **Next**.
3. Type a word or phrase that best describes this policy in the **Name** field. The name appears under **Home Folder Policies**.
4. In the **Domain** field, type the name of the domain that contains the home directories, or click **Browse** to select a domain.
5. In the **DFS Root** field, type the name of the root that you want to add users' home directories to, or click **Browse** to select a root.
6. In the **Additional Path** field, optionally type the name of a folder where you want to group the home directories.
8. Click **Next**.
9. Select the users that you want to add to the namespace.
10. Click **Next**.
12. Click **Finish**. The policy appears under the **Home Folder Policies** folder and the links to the users' home folders appear in the Logical View.

    Home folder policy as it appears in the Admin View

    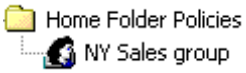

    Links to users' home folders as they appear in the Logical View

    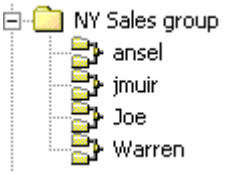

Create a multiprotocol namespace

Multiprotocol namespaces enable you to create a common namespace for UNIX and Windows users to provide access to files stored on NetApp filers. This presents a unified view of files to users. A multiprotocol namespace policy enables you to build namespaces for NFS clients to mirror the namespace presented to CIFS clients and to build namespaces for CIFS clients to mirror the namespace presented to NFS clients. VFM enables you to automatically synchronize the namespaces.

Benefits of using a multiprotocol namespace policy

- **UNIX clients can access data through the namespace** - Enables UNIX clients to access data without knowing its physical location.
- **Enables you to maintain a multiprotocol namespace** - Your CIFS and NFS clients can access data through a uniform namespace.
- **Enables you to automatically synchronize the namespace** - You can schedule how often you want VFM to synchronize the namespaces.

Note: Review the topic "Multiprotocol namespace planning considerations" before you create your first multiprotocol namespace policy.

Create a multiprotocol namespace checklist

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

1. Create a multiprotocol namespace policy.
2. Schedule synchronization.

Step 1 Create a multiprotocol namespace policy

Create a multiprotocol namespace to allow NFS and CIFS clients to access data on NetApp filer through a logical namespace.

Create a multiprotocol namespace policy where the DFS namespace is the master

When the DFS namespace is the master, the logical structure of the UNIX namespace is changed whenever the logical structure of the DFS namespace changes.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Multiprotocol Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Type the name of a DFS root that contains links to targets on NetApp filers, or click **Browse** to locate a root in your namespace.
4. Type a CIFS UNC path name for the directory that will contain files created by VFM to configure the UNIX namespace.
5. Click **New** to select the data that you want UNIX clients to access.
6. Type a word or phrase that best describes the data to be synchronized. This name appears on the **Sync Points** list.

7. Select DFS as the master. Any time the links in the DFS Logical Folder change, VFM updates the automount file and corresponding export files.
8. Type the path of a DFS folder that contains links referencing data on NetApp filers that you want UNIX clients to access.
9. Type a name for the automount file VFM will create.
10. Click **OK**.
11. Click **OK**. The policy appears under the Namespace Policies\Multiprotocol Policies folder in the Admin View.

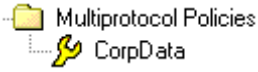

Create a multiprotocol namespace policy where the UNIX namespace is the master

When the UNIX namespace is the master, the logical structure of the DFS namespace is changed whenever the logical structure of the UNIX namespace changes.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Multiprotocol Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Type the name of a DFS root that contains links to targets on NetApp filers or click **Browse** to locate a root in the Physical View or your network.
4. Type a CIFS UNC path name for the directory that contains the automount files describing the UNIX namespace.
5. Click **New** to select the data that you want DFS clients to access.
6. Type a word or phrase that best describes the data to be synchronized. This name appears on the **Sync Points** list.
7. Select UNIX as the master. Any time the contents of the automount file changes, VFM updates the links in the DFS Logical Folder.
8. Type the name of a DFS folder that contains links referencing data on NetApp filers that UNIX clients are accessing.
9. Type the name of an automount file, or click **Browse** to select the file in the directory you specified in the **UNIX Files Location** field.
10. Click **OK**.
11. Click **OK**. The policy appears under the Namespace Policies\Multiprotocol Policies folder in the Admin View.

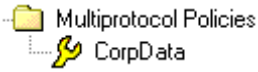

Step 2 Schedule synchronization

The next step is to make sure that changes made to the logical structure of the sync points are synchronized from the master namespace. Schedule how frequently you want to synchronize the logical structure.

1. Right-click the multiprotocol namespace policy and click **Properties**.
2. Click **Schedule**.
3. Click **Add** and configure the schedule.
4. Click **OK**.

Back up the logical namespace

VFM enables you to back up the logical structure of a DFS root or multiple roots. To manage namespace backups, you create a namespace backup policy and schedule how frequently you want to back up. VFM enables you to browse your backup files and restore a root.

Why you want to back up your namespace

You might want to back up the configuration of a DFS root or multiple roots for the following reasons:

- Duplicate a namespace on another server.
- Record the DFS configuration before making any changes so that you can restore the configuration if the changes don't function as desired.
- Easily maintain multiple DFS configurations for use in different conditions.
- Quickly restore a DFS root that becomes unavailable.

Create a namespace backup policy checklist

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

1. Create a backup policy.
2. Schedule backups.
3. Browse your backup file.

Step 1 Create a backup policy

Create a namespace backup policy to back up the logical structure of the DFS roots participating in the policy.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and then click **Backup Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Click **Add** and type the path of a root or click **Browse** to select one of the roots that appears in your namespace.

4. Repeat Step 3 to add more roots to the policy.
5. Click **OK**. The policy appears under the Namespace Policies\Backup Policies folder in the Admin View.

Step 2 Schedule backups

After you create a policy, the next step is to schedule how frequently you want to back up the roots.

1. Right-click the namespace backup policy and click **Properties**.
2. Click **Schedule**.
3. Select **Use Schedule**.
4. Click **Fill in current time and date**.
5. Select **Perform weekly backups**.
6. Click **OK**.

Step 3 Browse your backup file

Check the results of the backup operation. Browse the contents of the backup file that VFM created.

1. On the **Tools** menu, click **Browse Backups**.
2. Double-click the policy you created. Daily, Hourly, and Weekly folders appear.
3. Double-click the folder that contains the backup file.
4. Double-click the backup file to view its contents. A tree that looks like the Logical View appears on the left pane.
5. Click an item in the tree to view its properties on the right-pane.
6. When you finish browsing, click **OK**.

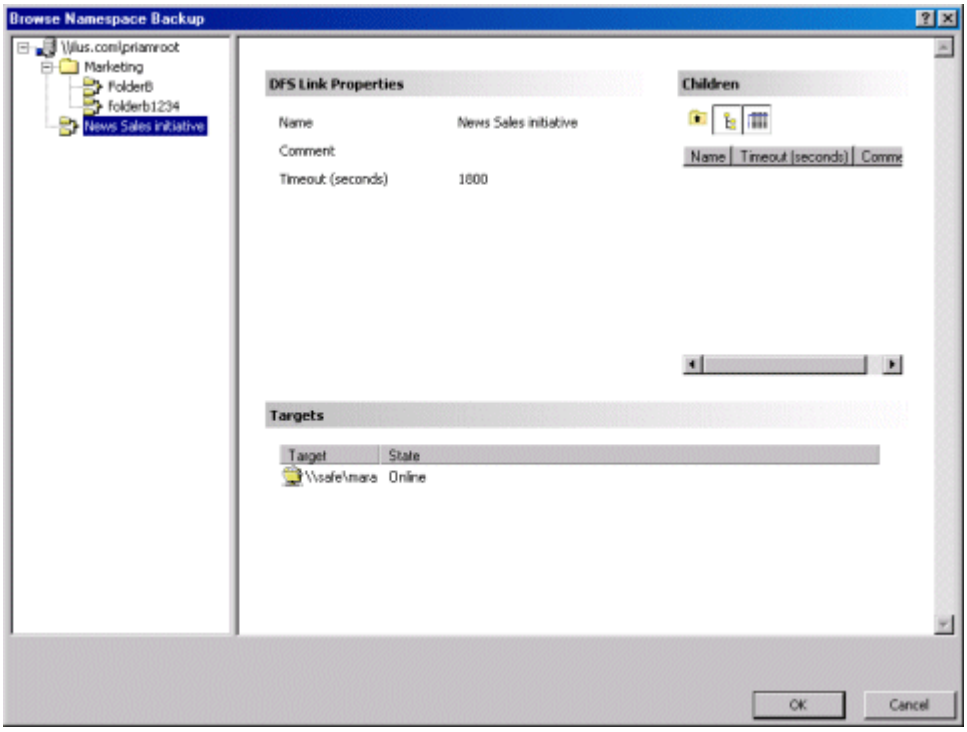

Back up data

Back up, restore data on filers, and optionally back up and restore the logical structure of the namespace. VFM enables you to create and centrally manage snapshots on multiple filers. VFM provides two snapshot policies:

- **Snapshot access policies** enable you to create new snapshots and manage filer-initiated snapshots.
- **Snapshot scheduling policies** enable you to schedule filer-initiated and VFM-initiated snapshots on all the filers in a policy.
    - VFM enforces a common schedule across multiple filers so that you are free from having to change snapshot parameters on individual filers and volumes.
    - VFM-initiated snapshots can run batch files before and after a snapshot.
    - Snapshot reserve settings can be centrally maintained and enforced for multiple volumes.

Create snapshot policies checklist

To create your first snapshot access and snapshot scheduling policies, follow the steps listed below:

1. Create a snapshot access policy.
2. View the snapshot data.

3. Restore the data.
4. Create a snapshot scheduling policy to schedule snapshots.

Step 1 Create a snapshot access policy

Create a snapshot access policy to back up data on filers with volumes that participate in your namespace and to capture the logical structure of the namespace targeting the data.

Prerequisite

NetApp filers must have the snapshot feature turned on.

1. On the **Tools** menu, click **New Policy**, point to **Snapshot Policy**, and then click **Access Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field.
3. Click **Add All Roots From the Logical View**.
4. Click **OK**. A new policy appears under the Snapshot Policies\Access Policies folder. VFM searches for snapshots on the volumes referenced in your namespace, captures the namespace referenced by each snapshot, and displays the captured namespace files under a Captured Namespaces folder under the policy.

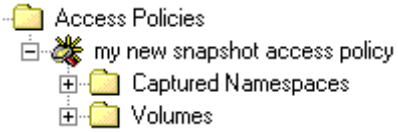

Step 2 View the snapshot data

After you create the snapshot access policy, the next step is to validate the data referenced by the captured namespace.

1. Open a folder under the Captured Namespaces folder to show a list of captured namespace files.

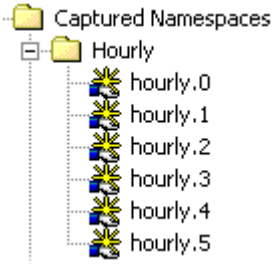

2. Select one of the files to show tabs on the right pane.
3. Click the **Explorer** tab.
4. The left pane on the Explorer tab shows the captured namespace. Click a link to show the data referenced by the link on the right pane.

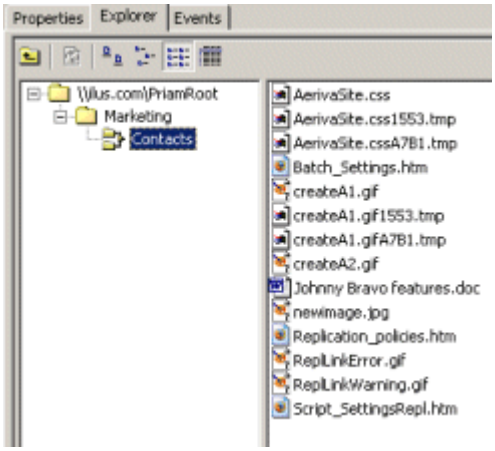

Note: VFM enables you to quickly browse the contents of a snapshot file if the appropriate application is available. Double-click a file to browse its content.

Step 3 Restore data

The next step is to restore a file to its current location.

1. Right-click a captured namespace file under the Captured Namespace folder and click **Restore**. A list of the snapshot data to be restored appears.
2. Select **Restore the Snapshot Files**.
3. Click **OK**. The data in the active file system is replaced with the data in the snapshot.

Step 4 Schedule snapshots

In the last step, you create a snapshot scheduling policy to centrally manage the snapshots on the volumes referenced in the snapshot accessing policy you created.

VFM enables you to centrally manage snapshots on all your NetApp filers. When you create a snapshot scheduling policy, you can create VFM-initiated and filer-initiated snapshots. VFM-initiated snapshots enable you to run batch files before and after a snapshot operation and enable you to set up a flexible schedule. Filer-initiated snapshots enable you to manage the schedule and reserve for snapshots on multiple filers.

1. Right-click the snapshot access policy and click **Set Snapshot Schedule for Referenced Volumes**.
2. Click **Next**.
3. Select **Add this volume to a Snapshot Scheduling policy** and click **Next**.
4. Click **New Policy**.
5. Type a word or phrase that best describes the policy in the **Name** field.
6. Click **Filer Settings**.
7. Select **Manage filer-initiated snapshot schedule**.
8. Type how many snapshots you want to retain.

9. Select the hours you want to schedule snapshots.
10. Optionally type the amount of space you want to reserve for snapshots on the referenced volumes.
11. Click **OK**.
12. Select the new policy and click **Next**.
13. Click **Finish**. A new policy appears under the Snapshot Policies\Scheduling Policies folder.

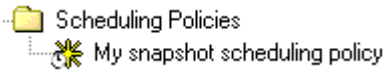

Create and publish a report

VFM enables you to create reports, publish them to a location you specify, process report data, and archive reports.

- The Reports tree displays reports you generate.
- VFM provides a Report wizard that steps you through the process of creating a report in Hypertext Markup Language (HTML) and publishing it with an updated index file. You can also select to process report data. VFM passes the data to a batch file you specify. For example, you can format and email the report to a distribution list and capture the contents of the report for use in trend analysis.
- When you create a report, a copy of the report appears alphabetically under the Generated Reports folder and under the report source group to which it is assigned. You can organize reports in folders, view the contents of the report, and modify the report.

Create and publish a report checklist

To create your first report, follow the steps listed below:

1. Add a machine to the Physical View tree.
2. Enable report publishing.

3. Create a Share Properties report.
4. Gather data and publish the report.

Note: Skip Step 1 if you already have a machine in the tree.

Step 1 Add a machine to the Physical View tree

In this example, you create a Share Properties report. For this report, you need to have a machine in the Physical View tree. If you do not have a machine in the tree, create a folder and then add a machine to the folder. Skip this step and go to "Enable report publishing" if you already have a machine in the Physical View.

1. Right-click **Physical View** and click **Add folder**.
2. Type a name for the folder and click **OK**.
3. Drag a machine from Windows Explorer to the folder you created.

Step 2 Enable report publishing

You must enable report publishing before you can publish reports.

1. On the **Tools** menu, click **Options**.
2. Click **Reporting**.
3. Select the **Enable Report Publishing** check box.
4. Type the UNC path of the share where you want to publish reports, or click **Browse** to find a share in the Physical View tree or your network.
5. Click **OK**.

Step 3 Create a Share Properties report

The next step is to use the Report wizard to create and publish a Share Properties report. You assign the report to a report source group that you specify the location where you want to gather data.

1. Right-click the Reports tree Reports , point to **Reports**, point to **Physical**, and then click **Share and Export Properties**.
2. The columns that will appear in the report are selected. Click **Next**.
3. Select the **Publish report** check box and click **Next**.
4. Select the "Report Sources" report source group that appears under the "Monthly" report schedule group and click **Edit**.
5. Type the UNC path of the machine you added to the Physical View, or click browse to find it.
6. Click **OK**.
7. Click **Next**.
8. Optionally type a different name for the report.
9. Click **Finish**. The report appears under the Generated Reports folder and under the report source group you modified. The report has no content because the schedule to gather data has not started.

Step 4 Gather data and publish the report

In the last step, you manually start data gathering and report publishing.

1. Expand the Report tree to locate the report source group to which you assigned the report you created.
2. Right-click the report source group and click **Start data collection now**. Wait a few minutes while VFM gathers the data and updates and publishes the report.

Appendix

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

No.	Title	File name	Contents
1.	Advanced Configuration and Techniques	AdvancedConfig.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.	Data Lifecycle Automation Policies Guide	DataLifecyclePolicies.pdf	Explains how to create and manage archival migration policies to optimize storage.
3.	Data Movement Policies Guide	DataMovePolicies.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 • Explains how to load balance Windows servers by creating storage load balancing policies
4.	Disaster Recovery Policies Guide	DRPolicies.pdf	Explains how to make data highly available by creating client, filer, and server disaster recovery policies
5.	Getting Started Guide	GettingStarted.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.
6.	Manage Physical	PhysicalResources.pdf	Explains how to use VFM to work

No.	Title	File name	Contents
	Resources Guide		with NetApp filers, set replication properties, and work with Windows and UNIX machines
7.	Namespace Guide	Namespace.pdf	Explains how to create a DFS namespace (roots, links, and link targets), namespace creation, permission based, namespace availability, multiprotocol namespace, home folder, and namespace backup policies
8.	Reports Guide	Reports.pdf	Explains how to create and publish reports
9.	Snapshot Policies Guide	SnapshotPolicies.pdf	Explains how to use VFM to manage snapshots on NetApp filers using snapshot access and snapshot scheduling policies
10.	Troubleshooting Guide	Troubleshooting.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.

Alternate volume: SEE Replica.

C

Child node: SEE DFS link.

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 branch: SEE Shared folder.

DFS leaf: SEE Shared folder.

DFS link: Located under the DFS root. Forms a connection to one or more shared folders, DFS links, or another DFS root. NOTE: In Windows NT 4 terminology, this is referred to as a child node, a junction point, or an inter-DFS link.

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

DFS path: A UNC path starting at a DFS root. Stand-alone DFS paths are rooted at a stand-alone root, while domain-based paths are rooted at a domain-based DFS root.

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

DFS topology: Overall logical hierarchy of a distributed file system, 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.

DFS volume: SEE Shared folder.

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

E

Export: SEE NFS export

F

Fault-tolerant DFS: SEE Domain-based DFS root

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

I

Inter-DFS link: SEE DFS link.

J

Junction point: SEE DFS link.

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 tasks: Items to be migrated during a scheduled migration or manually at any time.

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.

R

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

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 component during configuration of storage policies that transfer CIFS data, or when you run a file differencing report. NOTES: Deployment of replication agents can be controlled through a machine's properties dialog box accessible from the Physical View. For data transfer using the NFS protocol, the replication agent must be installed manually.

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 server: A machine that is standing by (not active) for the application server or a monitoring agent. During a failover, a standby server can become the active server.

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 DFS link....39
Add a link....11, 52
Add a shared folder replica....11
Add a source link to a replication policy....11, 52
Add a target link to a replication policy....11, 52
Add a task to a migration policy....55
Add links....52
Add multiple links....11
Admin View
  hide folders....17
Admin View....9
Admin View....17
Application
  system requirements....25
Application server
  connect to another server....18
Applications....19
Application-wide options....15, 16
Architecture....3
Archival migration policy
  schedule to scan for directories....46
  select migration task options....45
  specify selection criteria....45
Attributes
  set up multiple views....17
Audit events....16

B

Back up a DFS root....63, 64
Browse contents of a file....66

C

Change status monitoring....16
Client console....3
Client disaster recovery policy
  add a link to a policy....39
  schedule data replication....40
  schedule status monitoring....40
Client disaster recovery policy....39
Community name....16
Content distribution....51, 52
Contents....13
Create a DFS root....32
Create a migration policy....54
Create a multiprotocol namespace policy....61
Create a namespace....21, 32
Create a SnapMirror policy....58

D

Data distribution....51, 52
Data move....51, 52
Data replication....23, 40
Data synchronization....40
Data transfer....23, 58
Default storage policy options....16
Design a namespace....21
Develop a namespace....21
DFS link
  add a link....33, 39
  add multiple links....33
  find a link....14
  logical location in DFS namespace....14
  physical location on network....14
DFS namespace
  how to navigate your namespace....9
DFS namespace....33
DFS root
  back up a root....63
  change cache time-out....15
  create a root....32
DFS tree....14
Disable tool tips....18
Disaster recovery
  for NetApp filers....40
  preparing for....23
Disk Manager....19
Display folders....18
DR planning considerations....23
DR policies....39
Drag a machine to a user-defined folder....11
Drag an item from Windows Explorer....11, 33
Drag UNC paths....11
Drag-and-drop support....11
Drives
  find a drive....14
  logical location in DFS namespace....14
Drives....14
Drives....14

E

Event history....16
Event ID....16
Events
  filter records....16
Events....16
Events....16
External tools....19

F

File Lifecycle Manager....4, 18
Filer disaster recovery policy
  create a policy....40
Filers
  data transfer....58
  disaster recovery policy....40
Filtering event records....16
Find a drive....14
Find a string....14
FLM....4, 18
Folder

find a folder 14
hide folders in Admin View 17
hide folders on the Properties tab 18
logical location in DFS namespace 14
physical location on network 14
show folders on the Properties tab 18
view contents 13

H

Hide folders 17, 18

L

Launch external tools 19
Link target
set default cache timeout 15
view contents 13
Locate a string 14
Logical structure of the namespace 63, 64
Logical View 9
Look for a string 14

M

Machine properties 16
Machines
set SNMP community name 16
view contents 13
Microsoft Management Console 19
Migration operation 54
Migration policy
add a task 55
create a policy 54
schedule data migration 55
select migration actions 54
MMC 19
Monitor link targets 40
Monitor status 40
Move data 52
Multiple views 17
Multiprotocol namespace policy
create a policy 61
schedule namespace synchronization 63

N

Namespace
design considerations 21
Namespace backup policy
back up the namespace 63
schedule backups 64
view contents of backup file 64
Namespace creation policy
create a policy 34
Namespace synchronization 63
Navigate your namespace 9
NetApp filers
configure snapshots 66
data transfer 58
disaster recovery policy 40
NFS protocol 61
NT event log 16

O

Open external tools 19
Options 16

P

Physical location of a folder 14
Physical location of DFS link 14
Physical View 9
Policies 5
Printer setup 18
Printing 18
Publish reports 70

R

Replication 23, 40
Replication agents 3
Replication policy
add or replace the source link 52
add or replace the target link 52
schedule data distribution 52
Replication policy 52
Report publishing 69
Reports
generate a report 70
publish a report 70
Restore data and namespace 67
Run external tools 19

S

Schedule data distribution 52
Schedule data replication 40
Schedule migration 55
Schedule namespace synchronization 63
Schedule snapshots 67
Schedule status monitoring 40
Search for
a string 14
Search for 14
Set cache time-out 15
Set SNMP community name 16
Shared folder
view contents 13
Show folders 18
SnapMirror policies 23
SnapMirror policy
configure replication 58
create a policy 58
Snapshot access policy
access snapshot data 66
browse contents of a snapshot file 66
create a policy 66
restore data and namespace 67

Snapshot access policy 66
SNMP community name 16
Source link.....52
Split right-hand pane 17
Status monitoring 40
Storage load balancing policy.....49
Storage policies 5
System requirements 25

T

Target link of a replication policy 52
Tool tips
  disable 18
Tool tips.....18
Tools
  launch external tools..... 19
Tools..... 19

V

View contents of a file.....13, 64
Views..... 17

W

Wide link namespace..... 37
Windows 2000.....32, 39
Windows Explorer..... 33
Workspace 17