



VFM (Virtual File Manager) v6.0

Getting Started

June 28, 2007

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

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

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

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

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

This product includes software developed by the Apache Software Foundation.

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

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

Trademark information

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

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

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

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

Preface

What the guide covers

This guide describes how you can use Network Appliance® VFM® technology to 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 Started gets you started using VFM to create your first namespace and solve your storage problems. This chapter explains how to 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, 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 Server® 2003, and DFS. It does not cover basic system or network administration topics.

Typographic conventions

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

code_text_in_italics – Variable; replace with actual value.

Using the online help

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

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

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

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

Support

Contact NetApp technical support, 24 x 7.

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

Support phone numbers:

888 4-NETAPP (US and Canada)

(1-888-463-8277)

(408) 822-4700

800-44-NETAPP (EMEA/Europe)

800-800-0800 (Asia/Pacific)

Notes:

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

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

Enter new license key

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

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

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

Table of Contents

PREFACE	I
CHAPTER 1 LEARNING ABOUT VFM.....	1
WHAT'S NEW IN THIS RELEASE	3
COMPONENTS	3
<i>Architecture overview.....</i>	<i>4</i>
<i>The desktop.....</i>	<i>5</i>
<i>Feature comparison checklist.....</i>	<i>6</i>
<i>Storage policies</i>	<i>9</i>
CHAPTER 2 WORKING WITH VFM.....	11
NAVIGATING YOUR DFS NAMESPACE	11
USING DRAG-AND-DROP	13
VIEW CONTENTS OF AN ITEM IN THE DFS TREE.....	15
FIND AN ITEM.....	15
LOCATE AN ITEM	16
<i>To locate a logical item in the network.....</i>	<i>16</i>
<i>To locate a physical item in the namespace.....</i>	<i>16</i>
<i>To locate a storage policy.....</i>	<i>16</i>
CHECK SYSTEM STATUS	16
SET DEFAULT TIMEOUTS	17
SET DEFAULT SNMP COMMUNITY NAME.....	17
WRITE AUDIT EVENTS TO THE NT EVENT LOG	17
SET DEFAULT POLICY OPTIONS.....	18
FILTER EVENT DISPLAY	18
SET UP MULTIPLE VIEWS	19
HIDE EMPTY ADMIN POLICIES FOLDERS	19
HIDE FOLDERS ON THE PROPERTIES TAB	19
DISABLE TOOL TIPS	20
PRINTER SETUP	20
CONNECT TO ANOTHER SERVER.....	20
LAUNCH EXTERNAL TOOLS	20
CHAPTER 3 PLANNING CONSIDERATION HIGHLIGHTS	21
NAMESPACE DESIGN CONSIDERATIONS	21
DISASTER RECOVERY PLANNING CONSIDERATIONS	22
MULTIPROTOCOL NAMESPACE PLANNING CONSIDERATIONS.....	23
SNAPMIRROR PLANNING CONSIDERATIONS	23
CHAPTER 4 GETTING STARTED	25
CREATE A LOGICAL NAMESPACE.....	26
<i>Create a namespace checklist.....</i>	<i>27</i>
QUICKLY POPULATE THE NAMESPACE.....	29
<i>Difference between namespace creation and permission based namespace policies.....</i>	<i>29</i>
<i>Quickly populate the namespace checklist</i>	<i>29</i>
MAKE THE NAMESPACE HIGHLY AVAILABLE	31
<i>Create a namespace availability policy checklist.....</i>	<i>31</i>
MAKE DATA HIGHLY AVAILABLE.....	33
<i>Create a client disaster recovery policy</i>	<i>34</i>
<i>Step 2 Add a link to the client disaster recovery policy</i>	<i>34</i>
<i>Create a filer disaster recovery policy.....</i>	<i>35</i>
<i>Create a server disaster recovery policy</i>	<i>36</i>
OPTIMIZE STORAGE.....	38

<i>Create an archival migration policy</i>	38
<i>Create a storage load balancing policy</i>	44
TRANSFER DATA	46
<i>Create a replication policy</i>	46
<i>Create a migration policy</i>	49
CREATE A MULTIPROTOCOL NAMESPACE.....	53
<i>Benefits of using a multiprotocol namespace policy</i>	53
<i>Create a multiprotocol namespace checklist</i>	53
BACK UP THE LOGICAL NAMESPACE.....	55
<i>Why you want to back up your namespace</i>	55
<i>Create a namespace backup policy checklist</i>	55
BACK UP DATA.....	58
<i>Create snapshot policies checklist</i>	58
CREATE AND PUBLISH A REPORT	61
<i>Create and publish a report checklist</i>	61
APPENDIX	65
VFM GUIDES	65
GLOSSARY	67
INDEX	71

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, and 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 files without having to know their physical location (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. Finally, 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. By first deploying a global namespace, you can use VFM data movement features to quickly and easily consolidate or expand file servers. The global namespace masks physical changes from users during consolidation, and VFM storage policies significantly reduce operator activities during the consolidation. VFM enables one-step data migration between storage devices to optimize capacity and achieve server consolidation 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 replicates only the changed data instead of replicating entire files and wasting WAN bandwidth. Key file attributes are automatically retained while consolidation of data across multiple file storage platforms is occurring, thus reducing the administrative workload. VFM enables administrators to set multiple migration schedules with bandwidth consumption settings and to 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.

What's new in this release

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

Phased migration enables you to stage the data migration operation. You can run automatically repeated data copy operations before running the final copy. You control when the migration operation moves from one phase to another. The feature is available for the archival migration, migration, and storage load balancing policies.

Quickly add multiple targets to a link to clone an existing link target to add multiple targets to a link at one time. A typical use of this feature is to quickly add multiple link targets pointing to WAFS edge devices that are based on an existing link target pointing to a WAFS core device.

Agent Management enhances the management of replication and monitoring agents and the VFM server including deployment credentials.

Replication Agent groups enable you to add multiple replication agents to an agent group and assign the group to a policy to distribute the data migration workload. This improves flexibility and throughput in many-to-one replication scenarios.

Namespace back up and restore supports new DFS attributes introduced in Windows Server 2003 release 2.

Components

This section provides a high-level overview of the VFM architecture, introduces the main components of the VFM desktop, compares the features of VFM-DM and VFM, 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."

To compare the features in VFM-ME (Migration Edition) with those available in VFM (Full or Enterprise Edition), see "Feature comparison checklist."

For information about storage policies, see "Storage policies."

Components

Architecture overview



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 policies and to maintain storage configuration and reporting information. The server is implemented as an NT service.

VFM monitoring agent – The agent allows distributed processing of policies. The monitoring agent is implemented as an NT service.

VFM replication agents – The agents carry out file transfers for policies. VFM provides Windows and UNIX replication agents. The agents are deployed by VFM. You can choose to deploy an agent manually

Windows replication agents are implemented as NT services. The agents transfer CIFS data.

UNIX replication agents are available for Red Hat® Enterprise 4.0 and Solaris® 10. The agents transfer NFS data.

The desktop


The left pane of the VFM desktop displays three views (Admin, Logical, and Physical), and Reports. 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, supporting machines, shares, and exports in **Reports**.

Display attributes of resources in the right pane.



Components

Feature comparison checklist

This table compares the features that are available in VFM-ME (Migration Edition) and VFM (Full or Enterprise Edition).

Feature	VFM-ME	VFM
Data Migration		
Transfer data in heterogeneous environments. Create, edit, and delete migration policies	•	•
Support bandwidth throttling	•	•
Support NTFS Change Journal	•	•
Support CIFS protocol	•	•
Options for controlling how a file's security is copied	•	•
Option to update the namespace	•	•
Option to run batch files before and after a migration	•	•
Schedule a migration	•	•
Start a migration manually	•	•
Manage Windows replication agents	•	•
Manage UNIX replication agents		•
Centrally manage tasks associated with replicating data across a network. Create, edit, and delete replication policies		•
Automatically balance the storage space among a group of storage locations containing directories that are targets of links in your namespace. Create, edit, and delete storage load balancing policies		•
Migrate directories containing older and less-frequently accessed data to secondary storage based on selection criteria you specify. Create, edit, and delete archival migration policies		•
Support differential replication		•
Support NFS protocol		•
Support SnapLock integration		•
Specify migration tasks from a file		•
Namespace		
Create, edit, and delete stand-alone DFS roots	•	•
Create, edit and delete domain-based DFS roots	•	•
Create, edit, and delete consolidation DFS roots	•	•
Create, edit, and delete DFS links	•	•
Create, edit, and delete shares	•	•

Feature	VFM-ME	VFM
Namespace		
Create and remove logical folders	•	•
Create domain-based root targets (replicas)	•	•
View the machines, links, and folders targeted by a root	•	•
View the machines and shares targeted by a link.	•	•
Toggle link targets online or offline	•	•
Check and clear the status of an item in the namespace	•	•
Expand a link	•	•
Monitor roots, links, and link targets		•
Fail over or fail back a link		•
Manually back up and restore a root		•
Create namespaces automatically. Create, edit, and delete namespace creation policies		•
Create namespaces based on users' security permissions to resources. Create, edit, and delete permission based namespace policies		•
Synchronize the logical structure of two or more roots or a root and a widelink namespace. Create, edit, and delete namespace availability policies		•
Create multiprotocol namespaces. Create, edit, and delete multiprotocol namespace policies.		•
Automatically add user home directories to a namespace. Create, edit, and delete home folder policies.		•
Automatically back a namespace. Create, edit, and delete namespace backup policies		•
Reports		
Generate file differences reports	•	•
Schedule reports	•	•
Start data collection manually	•	•
Run customized commands before and after a report schedule runs	•	•
Print a report	•	•
Create other reports		•
View a list of reports referencing a selected root, link, logical folder, machine, or share.		•
Publish reports		•
Archive reports		•

Components

Feature	VFM-ME	VFM
Data Disaster Recovery		
Detect a failure and automatically fail over clients to a replicated storage device. Create, edit, and delete client disaster recovery policies		•
Set up a filer or nearline storage device as a disaster recovery backup for one or more primary filers. Create, edit, and delete filer disaster recovery policies		•
Make data on servers participating in a namespace highly available. Create, edit, and delete server disaster recovery policies		•
Manage Physical Resources		
Manage machines running Windows 2000 and Windows Server 2003	•	•
Manage NetApp filers running Data ONTAP 6.5 and later.	•	•
Create a share	•	•
Create a volume	•	•
Refresh filer information	•	•
Provision NetApp filers		•
Centrally manage snapshots on NetApp filers. Create, edit, and delete snapshot access and snapshot scheduling policies.		•
Manage UNIX machines running Solaris 10 and RedHat Enterprise 4		•
Work with the application		
Connect to another server	•	•
Configure application access permissions	•	•
Check system status	•	•
View the properties of items in the namespace	•	•
View the status of selected items	•	•
View event history	•	•
Printer setup	•	•
Export/import application configuration files	•	•
Back up the application configuration files (using the client)		•
Launch external tools		•
Use drag-and-drop operations		•
View the contents of DFS links, shares, machines, or folders		•
View the machines, links, and folders targeted by a root		•
View the machines and shares targeted by a link.		•

Storage policies

You can 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 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 subfolders 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. These policies also enable you to 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 detect a failure and automatically fail over clients to a replicated storage device. One of the most significant advantages of the disaster recovery solution is the ability to designate a cascaded failover path. In a typical scenario, you designate critical data to fail over first to a local edge device then to a remote central device. This ensures high availability of the data whether there is a local server failure or an entire site outage. Failover takes place without any changes in login scripts or network configuration.

Filer disaster recovery policies enable you to make data on NetApp filers participating in your namespace highly available. You can 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. The policy 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. You can scan 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. You can 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 the migration is carried out automatically or you are notified of items matching your criteria so that you can migrate them later.

Components

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. Status information is provided about replication operations performed using the VFM replication agent. Data is replicated 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.

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 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 both 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. You can 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. You can 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 and restore files and folders from snapshot backups.

Snapshot scheduling policies enable you to schedule snapshots on NetApp filer volumes. You can 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 the reserve settings 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 to merge part of a root with an existing root.

Chapter 2 Working with VFM

This section provides general instructions about how to use and customize VFM. For example, you 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  Logical View to view its attributes in the right pane. Right-click a root on the tree or on the Properties tab of  Logical View to display a context menu.	
Storage	Select a storage	Click the name of a	

Navigating your DFS namespace

Item	Admin View	Logical View	Physical View
policy	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.	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 <u>not</u> 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 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, where you can drop them, and 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.
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.

Using drag-and-drop

General		
Drag	Drop on	To...
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 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.

General		
Drag	Drop on	To...
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 for which you want to search.
4. By default, a string is searched 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 information must be gathered 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 a 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 on which it is located.

To locate a 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 of time that the DFS roots or the targets of DFS links are cached on DFS clients. The changes affect only 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).

Set default SNMP community name

VFM uses "public" as the default SNMP community name. You can set an application-wide default value. Changes affect only 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: VFM shows events for items in the Admin, Logical, and Physical Views. Select an item in order 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

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 before VFM is shut down are still there when you restart VFM.

Events saved to a report are determined by the filter options that are 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.



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, the Admin policies folders always display. 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, folders appear 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

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

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.

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 section highlights issues 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 "Planning considerations" in Advanced Configuration and Techniques.

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 results 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 on. 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.

Disaster recovery planning considerations

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 on. 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 as a single point of failure the server that provides the namespace.

A single namespace server can handle about 2,000 to 4,000 users, depending on its 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 careful, thorough design of the virtual layer or namespace. The goal of the namespace is to present users and applications with a consistent long-term means to access data.

2. **Decide on the root configuration**
3. **Populate the namespace with links referencing network shares**
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 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 section serves as a tutorial to get you started using VFM to create your first namespace and storage policies.

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

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

- 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 and migration policies.

- 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.

Note: For system requirements, upgrade and installation instructions, and known issues, see the *VFM Release Notes*.

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.

Use VFM to virtualize and manage large-scale file storage in Windows environments.

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 or 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 `\\xyzcompany.com\salesInfo` is used heavily. You can use DFS to share the network load by associating the logical name *salesInfo* with the following servers:

```
\\npb.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 about 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 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 are creating 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.

Logical View

 \\ARGOS\dfsroot

Create a logical namespace

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

Quickly populate the namespace

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 those shares.

1. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Creation Policy**.
2. Click **Next**.
3. Type a word or phrase that best describes this policy in the **Name** field, and 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 click **Next**.
6. Select the domain or servers where you want to search for shares, and 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.

Quickly populate the namespace


8. Select **Search for shares matching the search string**, 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 click **Next**.
10. Select a template to arrange how links appear in the namespace. You can preview the namespace.
11. Click **Next**.
12. Verify the changes to the namespace, and click **Next**.
13. Click **New** and configure a schedule to run the policy, and 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 click **Permission Based Policy**.
2. Click **Next**.
3. Type a word or phrase that best describes this policy in the **Name** field and 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 click **Next**.
6. Select the domain or servers where you want to search for shares and 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 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 click **Next**.
11. Select a template to arrange how links appear in the namespace. You can preview the namespace.
12. Click **Next**.
13. Verify the changes to the namespace and click **Next**.
14. Click **New** and configure a schedule to run the policy and 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 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 refer DFS clients to local targets when the clients access 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 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.

Make the namespace highly available

5. Select the master root. This is the root with 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 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 with 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**.
5. Click **Done**.

Make data highly available

VFM provides three types of disaster recovery policies. Select which type you want to create.

Client disaster recovery policies enable you to set up disaster recovery based on links targeting data on heterogeneous devices. You can select to use VFM replication or another replication method to keep the data on the primary and standby servers synchronized.

Filer disaster recovery policies enable you to make data on NetApp filers participating in your namespace highly available. You can set up a filer or nearline storage device as a disaster recovery backup to one or more primary filers. A filer disaster recovery policy uses 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. The policy 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 Windows systems participating in your namespace highly available. You can monitor for new links targeting the primary server, replicate data to a secondary server, and automatically fail over all links to the secondary server when a failure is detected.

Make data highly available

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 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, to guide you through adding a second target to the link to make the data highly available. A  link with multiple targets appears under the policy.



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 the primary target can fail over to the alternate target should there be a problem with the primary 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 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 steps you through the process of 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**. The SnapMirror replication is created.
13. Click **Done**.
14. Click **Next**.
15. Select **Find new links** and schedule how often you want to monitor for new links.
16. Click **Next**.
17. Schedule how often you want to monitor the source link target. When the source link target becomes unavailable, VFM fails over to the destination link target. Click **Next**.

Make data highly available

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 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 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 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 a 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. One of the links that is a member of the policy is randomly selected to check the availability of the active target of the link. When the active target (a share on the primary server) is not available, all the links in the policy are failed over.

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 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

Archival migration and storage load balancing policies enable you to optimize your storage.

Archival migration policies enable you to set policies to control when data is migrated to secondary storage. You control whether the policy 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 criteria to evaluate the current size of disk space of storage locations you specify. You control whether the policy automatically load balances or provides migration suggestions so that you can decide what directories you want to migrate.

Create an archival migration policy

You can transfer CIFS or NFS directories. 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 checklist

To create your first archival migration policy to transfer CIFS directories, follow these steps:

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.

Transfer NFS directories checklist

To create your first archival migration policy to transfer NFS directories, follow these steps:

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.

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 subfolders. The shared folder named in the source path must contain subfolders because VFM scans for directories under the share as possible migration candidates. If the directories match the selection criteria, each subfolder 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.

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 to use to select directories to migrate.

1. On the **Tools** menu, point to **New Policy**, point to **Data Lifecycle Automation Policy**, and click **Archival Migration Policy**.
2. Type a word or phrase in the **Name** text box that best describes the policy.
3. Click **Add** to include a shared folder with subfolders as the source and a shared folder on another machine as the destination.

Optimize storage

4. Click **OK**. The policy and the source shared folder appear under the Data Lifecycle Automation Policies\Archival Migration Policies folder.

In the following example, the 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 **OK**.

Step 3 Specify the migration actions

After you specify the selection criteria, the next step is to define the migration operation which is divided into three phases: initial, incremental, and final. You will select actions in the initial and final phases that you want to run.

1. Right-click the policy you created and click **Properties**.
2. Click **Initial Phase**.
3. Make sure that **Perform a baseline copy** is selected.
4. Click **Final Phase**.
5. Make sure that the **Update namespace** check box is selected.
6. Make sure that **by retargeting links** is selected.
7. Click **OK**.

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. Click **Add** to list the DFS roots that contain links you want to update.
3. Click **OK**.

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. Migration tasks are created after scanning is complete.
3. Make sure that all the roots in your namespace are listed in the dialog box so that you can update links to reference the new location.
4. Click **OK**.

Step 6 Schedule scanning

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

1. Click **Policy Scan Schedule**.
2. Click **New**, select the schedule from the list provided, and configure the schedule.
3. Click **OK**.

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 when to run the initial and final phases to migrate directories that match the selection criteria.

1. Click **Initial Phase Schedule**.
2. Specify the day and time you want to run the initial phase.
3. Click **Final Phase Schedule**.
4. Click **New**, select the schedule from the list provided, and configure the final phase schedule.
5. Click **OK**.

Optimize storage

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.

The source and destination can be an export or a folder under an export. The policy scans for subfolders. The export named in the source path must contain subfolders because the policy scans for directories under the export as possible migration candidates. If the directories match the selection criteria, each subfolder 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
3. VFM replaces search path A with B	~~hostname:/root/subdirectoryA~~ \\hostname:/root/subdirectoryA	~~hostname:/root/subdirectoryB~~ \\hostname:/root/subdirectoryC

Notes:


For instructions about how to specify default Windows and UNIX proxy machines, see "Specify default replication agents."

For instructions about how to specify a Windows or UNIX proxy machine, see "Specify a proxy machine."

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 to use to select directories to migrate.

1. On the **Tools** menu, point to **New Policy**, point to **Data Lifecycle Automation Policy**, and click **Archival Migration Policy**.
2. Type a word or phrase in the **Name** text box that best describes the policy.
3. Click **Add** to include an export with subfolders 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 folder to secondary storage if the folder 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 **OK**.

Step 3 Specify the migration actions

After you specify the selection criteria, the next step is to define the migration operation which is divided into three phases: initial, incremental, and final. You will select actions in the initial and final phases that you want to run.

1. Right-click the policy you created and click **Properties**.
2. Click **Initial Phase**.
3. Make sure that **Perform a baseline copy** is selected.
4. Click **Final Phase**.
5. Make sure that the **Update namespace** check box is selected.
6. Make sure that **by retargeting links** is selected.
7. Click **OK**.

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 **OK**.

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. Migration tasks are created after scanning is complete.
3. Click **OK**.

Optimize storage

Step 6 Schedule scanning and migration

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

Step 6 Schedule scanning.

Step 7 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 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 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 **OK**.



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 **OK**.

Note: At least one of the shares selected 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 **New**, select the schedule from the list provided, and configure the schedule.
3. Click **OK**.

Step 4 Schedule migration

The last step is to schedule when to migrate directories referenced by links in your namespace to load balance the storage pool.

1. Click **Initial Phase Schedule**.
2. Specify the day and time you want to run the initial phase.
3. Click **Final Phase Schedule**.
4. Click **New**, select the schedule from the list provided, and configure the final phase schedule.
5. Click **OK**.

Transfer data

Select which type of data transfer policy you want to create.

Replication policies enable you to centrally manage tasks associated with distributing 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. Status information is provided about replication operations performed using the VFM replication agent.

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.

Create a replication policy

You can 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

Status information is provided about replication operations performed using a 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 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**. 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 have 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.



Transfer data

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 "Archive 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 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: You can 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

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; if the source is an NFS export, the destination must be an NFS export. A migration policy can contain both 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 Brocade 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 click **Migration Policy**.
2. Type a word or phrase that best describes this policy in the **Name** text box.
3. Click **OK**. A new policy appears under the Migration Policies folder.



Transfer data

Step 2 Select migration actions

After you create a policy, the next step is to define the migration operation which is divided into three phases: initial, incremental, and final.

1. Right-click the policy you created and click **Properties**.
2. Click **Initial Phase**.
3. Make sure that **Perform a baseline copy** is selected.
4. Click **Final Phase**.
5. Make sure that the **Update namespace** check box is selected.
6. Make sure that **by retargeting links** is selected.
7. Click **Apply**.

Step 3 Update namespace options

The next step is to specify the DFS root that contains the links you want to update.

1. Perform one of the following actions:
 - When migrating CIFS data, click **Add** to list the DFS roots that contain links you want to update.
 - When migrating NFS data, click **Add** to list the automount files you want to update.
2. Click **Apply**.

Step 4 Schedule migration

The next step is to schedule the initial and final migration phases.

1. Click **Initial Phase Schedule**.
2. Specify the day and time you want to run the initial phase.
3. Click **Final Phase Schedule**.
4. Click **New**, select the schedule from the list provided, and configure the final phase schedule.
5. Click **OK**.

Step 5 Add a task to the migration policy

The last step is to add tasks to the policy that identify source and destination paths. A wizard steps you through the process.


Notes:

Data transfer must occur within the same protocol. A migration task can only contain CIFS source and destination or NFS source and destination, but not both protocols.
A migration policy can contain CIFS migration tasks 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** text box 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, you see a list of the destination shares and a suggested local path for each share. Select all the items with no destination share and click **Create Destination Share**. The shares are created 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

The NFS source and destination must be UNIX fully qualified physical paths; for example, `<hostname>:/<root>/<subfolder>`.

1. Right-click the migration policy and click **Add task to policy**.
2. In the **Original Path** text box, type the path of an export, or click **Browse** to locate an export.
3. In the **Destination Path** text box, type the 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**. The exports are created 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.



Note: The machines hosting the source and destination must be running NFS 3 or higher.

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. You can 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 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 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.

Create a multiprotocol namespace


7. Select DFS as the master. Whenever the links in the DFS Logical Folder change, the automount file and corresponding export files are updated.
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 that 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 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. Whenever the contents of the automount file changes, the links in the DFS Logical Folder are updated.
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

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. You can 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:

- To duplicate a namespace on another server.
- To record the DFS configuration before making any changes so that you can restore the configuration if the changes don't function as desired.
- To easily maintain multiple DFS configurations for use in different conditions.
- To 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 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.

Back up the logical namespace



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 was 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 and restore data on NetApp filers, and optionally back up and restore the logical structure of the namespace. You can 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.

A common schedule is enforced 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 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. The policy 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: You can 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.


Back up data

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.

You can 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 to set up a flexible schedule. Filer-initiated snapshots enable you to manage the schedule and reserve space 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 at which 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

Create reports, publish them to a location you specify, process report data, and archive reports.

The Reports tree displays reports you generate.

A **Report** wizard 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. The data is passed 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 a report, and modify a 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.

Create and publish a report

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 add a machine to it. 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 for which you specify the location where you want to gather data.

1. Right-click the Reports tree , point to **Reports**, point to **Physical**, and 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 the data is gathered and the report is updated and published.

Appendix

VFM Guides

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

VFM Guides

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

Glossary

A

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

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

C

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

Client: The user interface component of the application.

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

D

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

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

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

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

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

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

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

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

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

E

Export: SEE NFS export

F

Fault-tolerant DFS: SEE Domain-based DFS root

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

FRS: Microsoft File Replication Service.

G

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

K

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

L

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

M

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

Migration task: Identifies the source and destination.

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

N

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

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

P

Pool: A group of storage locations.

Primary server: An active application server or monitoring agent.

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

Q

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

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

R

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

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

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

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

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

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

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

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

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

S

Secondary target: SEE Standby target.

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

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

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

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

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

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

T

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

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

U

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

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

W

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

Index

A

Add a DFS link.....	34
Add a link	13, 47
Add a shared folder replica.....	13
Add a source link to a replication policy	13, 47
Add a target link to a replication policy	13, 47
Add a task to a migration policy.....	51
Add links	47
Add multiple links	13
Admin View	
hide folders.....	19
Admin View	11
Admin View	19
Application server	
connect to another server	20
Applications.....	20
Application-wide options	17, 18
Architecture.....	4
Archival migration policy	
schedule to scan for directories	41
select migration task options.....	41
specify selection criteria.....	40
Attributes	
set up multiple views.....	19
Audit events.....	17

B

Back up a DFS root	55, 56
Browse contents of a file	59

C

Change status monitoring.....	18
Client console.....	4
Client disaster recovery policy	
add a link to a policy	34
schedule data replication	35
schedule status monitoring	34
Client disaster recovery policy	34
Community name	17
Content distribution.....	47, 48
Contents.....	15
Create a DFS root.....	27
Create a migration policy	49
Create a multiprotocol namespace policy.....	53
Create a namespace	21, 27

D

Data distribution.....	47, 48
Data move	47, 48
Data replication	23, 35
Data synchronization	35
Data transfer	23

Default storage policy options.....	18
Design a namespace.....	21
Develop a namespace	21
DFS link	
add a link	28, 34
add multiple links	28
find a link	15
logical location in DFS namespace	16
physical location on network.....	16
DFS namespace	
how to navigate your namespace.....	11
DFS namespace	28
DFS root	
back up a root.....	55
change cache time-out.....	17
create a root.....	27
DFS tree.....	15
Disable tool tips.....	20
Disaster recovery	
for NetApp filers	35
preparing for.....	22
Disk Manager	20
Display folders.....	19
DR planning considerations.....	22
DR policies	34
Drag a machine to a user-defined folder.....	13
Drag an item from Windows Explorer	13, 28
Drag UNC paths	13
Drag-and-drop support	13
Drives	
find a drive	15
logical location in DFS namespace	16
Drives	15
Drives	16
E	
Event history.....	17, 18
Event ID.....	17
Events	
filter records	18
Events	17
Events	18
External tools.....	20
F	
Filer disaster recovery policy	
create a policy	35
Filers	
disaster recovery policy.....	35
Filtering event records	18
Find a drive.....	15
Find a string.....	15
Folder	

find a folder	15	disaster recovery policy	35
hide folders in Admin View	19	NFS protocol	53
hide folders on the Properties tab	19	NT event log	17
logical location in DFS namespace	16	O	
physical location on network	16	Open external tools	20
show folders on the Properties tab	19	Options	18
view contents	15	P	
H		Physical location of a folder	16
Hide folders	19	Physical location of DFS link	16
L		Physical View	11
Launch external tools	20	Policies	9
Link target		Printer setup	20
set default cache timeout	17	Printing	20
view contents	15	Publish reports	62
Locate a string	15	R	
Logical structure of the namespace	55, 56	Replication	23, 35
Logical View	11	Replication agents	4
Look for a string	15	Replication policy	
M		add or replace the source link	47
Machine properties	17	add or replace the target link	47
Machines		schedule data distribution	48
set SNMP community name	17	Replication policy	47
view contents	15	Report publishing	62
Microsoft Management Console	20	Reports	
Migration operation	50	generate a report	62
Migration policy		publish a report	62
add a task	51	Restore data and namespace	59
create a policy	49	Run external tools	20
schedule data migration	50	S	
select migration actions	50	Schedule data distribution	48
MMC	20	Schedule data replication	35
Monitor link targets	34	Schedule migration	50
Monitor status	34	Schedule namespace synchronization	55
Move data	48	Schedule snapshots	60
Multiple views	19	Schedule status monitoring	34
Multiprotocol namespace policy		Search for	
create a policy	53	a string	15
schedule namespace synchronization	55	Search for	15
N		Set cache time-out	17
Namespace		Set SNMP community name	17
design considerations	21	Shared folder	
Namespace backup policy		view contents	15
back up the namespace	55	Show folders	19
schedule backups	56	SnapMirror policies	23
view contents of backup file	56	Snapshot access policy	
Namespace backup policy	55	access snapshot data	59
Namespace creation policy		browse contents of a snapshot file	59
create a policy	29	create a policy	58
Namespace synchronization	55	restore data and namespace	59
Navigate your namespace	11	Snapshot access policy	59
NetApp filers		SNMP community name	17
configure snapshots	58	Source link	47

Split right-hand pane	19
Status monitoring	34
Storage load balancing policy.....	44
Storage policies	9

T

Target link of a replication policy	47
Tool tips	
disable	20
Tool tips.....	20
Tools	
launch external tools	20

Tools	20
-------------	----

V

View contents of a file.....	15, 56
Views.....	19

W

Wide link namespace.....	32
Windows 2000.....	27, 34
Windows Explorer.....	28
Workspace	19