215-03832_A0
1 March 2008

Network Appliance Virtual File Manager 6.1

Administrator's Guide

Copyright © 2001-2008 Brocade Communications Systems, Inc. All Rights Reserved.
Copyright © 2007-2008 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.

About This Document

This document is a procedural guide written to help storage area network (SAN) administrators install, configure, and use Network Appliance Virtual File Manager. This preface contains the following sections:

- Document conventions i
- Additional information ii

Document conventions

This section describes text formatting conventions and important notices formats.

Text formatting

The narrative-text formatting conventions that are used in this document are as follows:

bold text	Identifies command names Identifies the names of user-manipulated GUI elements Identifies keywords and operands Identifies text to enter at the GUI or CLI
italic text	Provides emphasis Identifies variables Identifies paths and Internet addresses Identifies document titles
`code` text	Identifies CLI output Identifies syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, **switchShow**. In actual examples, command lettercase is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive.

Notes and warnings

This documentation uses the following notes and statements.

NOTE
A note provides a tip, guidance or advice, emphasizes important information, or provides a reference to related information.

ATTENTION
An Attention statement indicates potential damage to hardware or data.

Additional information

For more information about our this product, go to http://www.netapp.com.

About This Document

Document conventions i
Additional information ii

Chapter 1 Introduction to VFM

In this chapter 1
What is VFM? 1
How VFM helps you 1
Simplifies storage management using a global namespace . . . 2
Enables non-disruptive expansion and consolidation 2
Provides intelligent replication and migration policies 2
Ensures business continuity with rapid disaster recovery 2
Centralizes routine filer configuration and management 3
Reduces ownership costs with a tiered storage architecture . . . 3
VFM components and architecture 3
Storage policies 4
Feature comparison checklist 6
Data migration features 7
Namespace features 7
Reporting features 8
Business continuity (Disaster recovery) features 9
Physical resource management features 9
Application and integration features 9

Chapter 2 Planning Considerations

In this chapter 11
VFM availability considerations 11
A convenient location for the VFM console 11
A reliable, highly available location for the VFM server 11
How large your VFM configuration is 12
Whether you have policies to run from remote location 12
Where the managed DFS roots are located 12
VFM network communication 12
Namespace design considerations 12
Plan the layout and design of the namespace 13
Select the root configuration you want to use 13
Populate the namespace with links to network shares 14
Instruct users to access data through the namespace 14
Data migration considerations 14
Data migration requirements 14
Transfer NTFS file security attributes 15
Transfer data between dissimilar file systems 16
Disaster recovery planning considerations 20
General disaster recovery considerations 20
Business continuity policy requirements 21
Resynchronizing SnapMirrors for disaster recovery failback . . 22

Replication planning considerations 22
Replication policy requirements 22
How replication policies work 23
Transfer NTFS file security attributes 23
Permissions 23
Status monitoring. 24

Namespace availability planning considerations 24
Namespace availability requirements 24
How namespace availability policies work 25
Options for setting up namespace availability policies 25

Namespace creation policy planning considerations. 26
Namespace creation policy requirements 26
How namespace creation policies work 26
Comparison with permission-based namespace policies. 27

Multiprotocol namespace planning considerations 27
Prerequisites 27
How multiprotocol namespace policies works. 27
UNIX namespace configuration 28

Home folder planning considerations 31

Chapter 3 Installing VFM

In this chapter 33

Installation options 33

System requirements 34
VFM server and monitoring agent requirements. 34
VFM client requirements 35
Windows replication agent requirements 36
UNIX replication agent requirements 37

Permission requirements 37

Filer requirements. 39

Installing VFM 39

Opening the VFM client for the first time. 41

Updating a VFM license key 41

RSH on Solaris and Linux computers 42

Upgrading VFM 42
Upgrading only the server or the server and client 42
Upgrading only the client. 42
Upgrading only the monitoring agent 43
Upgrading the monitoring agent and client 43
Upgrading replication agents 43

Uninstalling VFM components 45

Chapter 4 Permissions and the VFM client console

In this chapter 47

Permissions on accounts used with VFM components 47
Permission requirements for DFS roots 47
Permission requirements for the VFM client 47
Configuring VFM server access permissions 48
Configuring domain root and shared folder permissions 48

Configuring function-level access permissions 49
List of function-level permissions 49

Working with the VFM client console 50
View contents of an item in the DFS tree 50
Find an item 50
Locate an item 51
Check system status 51
Set default timeouts 52
Set default SNMP community name 52
Write audit events to the Windows event log 52
Set default policy options 53
Filter event display 53
Set up multiple views 53
Hide empty Admin policies folders 53
Hide folders on the Properties tab 54
Disable tool tips 54
Printer setup 54
Connect to another server 54
Integrate Tapestry FLM 54
Launch external tools 55

Getting started 55
Creating a logical namespace 55
Populating the namespace 57
Making the namespace highly available 59
Making the data highly available 61
Optimizing storage 64
Transferring data 69
Creating a multiprotocol namespace 74
Backing up the logical namespace 76
Backing up data 77
Creating and publishing a report 79

Chapter 5 Creating and Managing DFS Namespaces

In this chapter 81

Introduction 81
Checklist 81

Create namespaces 82
The Logical View tree 82
Basic DFS requirements 82
Recommended size limits 83
Additional resources 83

Populating the namespace 83
Create and manage roots 84
Create and manage links 98
Monitor status 113

Populate the namespace automatically 114
Configure a namespace creation policy 114
Update the schedule 116
Add or remove servers or domains 116
Add items to an exclusion list 117
Change the search type 117
Change the search string of a policy 118
Change how the links are arranged 118
Scan for items immediately 118
Review the list of new resources before adding links 119
Delete a namespace creation policy 119

Populate the namespace based on permissions 119
Configure a permission-based namespace policy 120
Update the search criteria 122
Other options 122
Delete a permission-based namespace policy 122

Make the namespace highly available 123
Create a namespace availability policy checklist 123
Namespace configuration options 123
Make DFS roots highly available 124
Create a widelink namespace 125
Namespace synchronization options 126
Schedule synchronization 126
Start synchronization immediately 127
Add additional roots or shares to a policy 127
Remove roots and widelink namespaces 128
Change the master root of a policy 128
Delete a namespace availability policy 128

Create multiprotocol namespaces 129
Create a multiprotocol namespace policy checklist 129
Create a multiprotocol namespace policy 129
Host a UNIX namespace on a filer 130
Configure UNIX clients to access the namespace 130
Select default UNIX settings 131
Add a sync point 132
Schedule synchronization 132
Run customized commands during synchronization 133
Update a multiprotocol namespace policy 133
Start namespace synchronization immediately 133
Delete a sync point 134
Delete a multiprotocol namespace policy 134

Virtualize user home directories 134
Create a home folder policy 134
Add a link to a user home folder 135
Add new users to a policy 136
View properties of a home folder policy 136
Refresh a home folder policy 136
Resolve discrepancies 137
Remove a user from a home folder policy 137
Move a user 137
Locate a user in the namespace 138
Locate a DFS root in the namespace 138
Delete a home folder policy 138

Back up the namespace 138
Create a namespace backup policy 139
Schedule backups 140
Browse contents of a backup 140
Remove a root from a namespace backup policy 140
Delete a namespace backup policy 141
Restore a DFS root 141

Chapter 6 Managing File Servers

In this chapter 143

Working with NetApp filers 143
Set SNMP community string 143
Change the shell credentials of a filer 144
Specify the ONTAPI credentials of a filer 145
Create a report of multiple filers 145
Enable or disable all SnapMirror backups 145
Filers not identified correctly 146
Create a SnapMirror on a filer 146
Create a volume on a filer 146
Create a qtree on a filer 147
Enable vFilers for filer recovery policy 147
Schedule snapshots on multiple filers 147
Enable or suspend snapshot schedules for multiple filers . . . 148
Take a snapshot of a volume immediately 149
Change the snapshot reserve for multiple volumes 149
Restore a snapshot 150

Setting replication properties 150
Disallow deployment of replication agents 150
Specify a proxy machine 151
Specify default replication agents 151
Set data transfer rate limit 152

Working with physical resources 152
The Physical View. 153
Add a machine to the physical view 154
Remove a machine from a folder 154
Specify NTFS Change Journal 155
Create shares. 155
Duplicate shares 156
Display drives 157
View attributes of a machine or drive. 157
View attributes of a shared folder. 158
View attributes of an export 158
Update security of a folder 158
Refresh a machine. 159
Change the shell credentials of a UNIX machine 160
Set NFS credentials 160

Chapter 7 Transferring Data

In this chapter 163

Introduction 163
Data migration considerations 164
Replication planning considerations 167

Migration policies 168
Migration limitations 169
Integrated namespace management 169
Migration usage options 170
Scheduling considerations 171
Create a migration policy checklist 178
Create a migration policy 178
Configure the initial phase 179
Schedule the initial phase 179
Configure the incremental phase 180
Schedule the incremental phase 180
Configure the final phase 181
Schedule the final phase 182
Specify namespace options 182
Select replication options 184
Specify replication filters 186
Select a replication agent 187
Select differential replication options 187
Select file security replication options 188
Select file attribute replication options 190
Add a migration task 191
Advance to final phase 195
Specify migration tasks from a file 197
Run customized commands during replication 197
Update the properties of a migration task 198
Start a migration task immediately 198
Start all tasks for a policy immediately 198
View progress of a task 199
User action required 199
Suspend a migration task 199
Remove a task 200
Delete a migration policy 200

Replication policies 200
Create a replication policy checklist 201
Create a replication policy 201
Select replication options for your policies 202
Schedule data distribution 202
Change data distribution path 203
Change replication for a target 203
Add or replace the source 203
Add or replace the target 204
Change the replication schedule 205
Start replication immediately 205
Cancel a replication in progress 205
View information about replication operations 205
Compare contents of targets in a replication policy 206
View properties of a replication policy 206
Delete a destination target 206
Delete a replication policy 206

Specifying migration tasks from a file 207
How it works 207
The task creation batch file. 207

Task Creation batch file 208

Example Task Description file 210

Archival Migration Policies 210
Planning considerations 211
Create a policy 212
Specify selection criteria 213
Configure migration phases 214
Specify namespace options 214
Select migration task options 214
Schedule scanning for directories to migrate 215
Schedule migration 215
Add source and destination paths 216
Exclude directories from migration. 218
To exclude a migration candidate: 218
Run customized commands 219
Scan for data immediately 219
Update the properties of a policy 220
Delete a policy 220
View a list of migration candidates. 220
View properties of a migration candidate 221
View a list of migration tasks 221
Add a migration task immediately 221
Start a migration task immediately 222
Additional migration task topics 222

Storage Load Balancing Policies 222
Create a storage load balancing policy 222
Add storage locations 225
Configure the initial phase 226
Configure the incremental phase 227
Configure the final phase 227
Schedule scanning. 228
Select replication options 228
Schedule migration 229
Start scanning immediately 229
Run customized commands during replication 229
Start migration immediately 230
View migration suggestions. 230
Manage migration tasks 231
Delete a storage load balancing policy. 231

Chapter 8 Creating and Managing Recovery Policies

In this chapter 233

Introduction 233

Disaster Recovery Planning 234
Planning considerations 234
Resynchronizing SnapMirrors for disaster recovery failback .235

Client Disaster Recovery 235
Create a client recovery policy 236
Add a link to a policy 237
Add multiple links to a policy. 238
Add a share to a client recovery policy 239
Select search options 239
Schedule target monitoring. 240
Schedule data replication 240
Select a SnapMirror 241
Failover and replication options 242
Select replication options 243
Select other replication options 246
Failover 246
Run customized commands during failover processing 249
Run customized commands during replication 250
Review the list of new resources. 251
Update properties of a client recovery policy. 251
Delete a client recovery policy 252
Add new shares to a link in a client recovery policy 252
Compare target contents in client recovery policy. 253
Add shared folder replicas to a link 253
Prioritize the targets of a link 254
Remove a link from a policy 254
View information about replication operations 255
Start replication immediately 255
Cancel a replication in progress 255

Filer Disaster Recovery 255
Create a filer recovery policy 256
Add volumes and qtrees to the policy 258
Fail over immediately. 258
Run customized commands 259
Review the list of new qtrees 259
Update the policy ignore list 259
View the status of a SnapMirror 260
View links targeting data on a filer 260
Update properties of a policy 260
Remove a filer, volume, or qtree from a policy. 261
Delete a filer recovery policy 261

Server Disaster Recovery 261
Create a server recovery policy. 261
Select resource monitoring options 262
Schedule resource monitoring 263
Schedule target monitoring and replication. 263
Select replication options 264
Add links to a server recovery policy. 264
Start monitoring for new links immediately 265
Assign a link for the policy to monitor. 265
Cancel resource monitoring in progress. 266
Manage replication operations. 266
Start failover or failback immediately. 266
Remove links from a server recovery policy. 267
Delete a server recovery policy 267

Chapter 9 Managing Snapshots on Filers

In this chapter 269

Snapshot Access Policies 269
Create a snapshot access policy 269
Take a snapshot immediately 270
Update captured namespaces 271
Specify which volumes participate in a policy 271
Browse snapshot data. 271
Restore data. 272
Restore namespace and data. 272
Re-create part of your namespace 273
Access snapshot data through the namespace 273
Access snapshot data directly 274
Change the snapshot schedule for all volumes. 274
Change the snapshot schedule for selected volumes 274
Capture the namespace of a snapshot 275
Update the properties of a snapshot access policy 275
Delete a snapshot access policy. 275

Snapshot Scheduling Policies 275
Create a snapshot scheduling policy checklist 276
Create a snapshot scheduling policy 276
Schedule snapshots 276
Update the properties of a snapshot scheduling policy 277
Suspend the snapshot schedule of a policy 277
Take a snapshot immediately 278
Delete a snapshot scheduling policy 278

Chapter 10 Working with Reports

In this chapter 279

Introduction 279

VFM reports. 280

Creating and Publishing Reports 282
Enable report publishing 282
Configuring default report settings. 283
Generate a report 283
Start data collection immediately. 284
Publish reports. 284
Run customized commands during index file generation. . . .285
Output report data to a file 286
Manually publish all reports 287

Grouping and Scheduling Reports. 287
Create a report schedule group 287
Create a report source group 288
Modify a report group 289
Move a source group 289
View reports assigned to a source group. 289
Update a report schedule 290
Delete a report source group 290
Delete a report schedule group 290

Managing Reports. 291
Organize reports 291
Save event history 291
Modify a report. 291
Configure report archiving. 292
Run customized commands 292
View the data collection log. 293
Update the format of a report. 293
Move a report. 295
View the contents of a report 296
Rename a report 296
Print a report 297
Delete a report. 297

Chapter 11 Managing Agents and the VFM Server

In this chapter 299

Replication agent overview. 299
What a replication agent is 299
Replication agent requirements 299
Distributed replication 300

Proxy computers and replication agents 302

Network ports in use. 302

Viewing status information 304

Manage replication agents 304
Deploying a replication agent 304
Setting replication agent default credentials 305
Stopping and restarting a replication agent 306
Pausing and resuming a replication agent 306
Upgrading a replication agent 306
Grouping replication agents 307
Changing the credentials of a replication agent 307
Changing the listening port of a replication agent 308
Purging a replication agent 309
Uninstalling replication agents 309

Installing and managing UNIX replication agents 310

Manage monitoring agents 314
Distributed monitoring 314
Checklist 315
Installing a monitoring agent 315
Changing monitoring agent credentials 315
Changing the listening port of a monitoring agent 316
Assigning a monitoring agent to a policy 316
Specifying a backup location for monitoring agent files 317
Purging a monitoring agent 317

Manage the application server 317
Checklist 318
Changing the application server credentials 318
Changing the application server listening port 319
Specifying a backup location for the application server 319

Chapter 12 Run Batch Files

In this chapter 321

Using batch files in VFM 321
How VFM processes batch files 321
Conditions under which scripts run 322
Testing scripts 322
Common script problems 322

Batch files you can run in VFM 323
Before updating index file 323
After updating index file 323
Criteria script 324
Migration task creation script 324
Pass report output to script 324
Pre-creation script 325
Pre-data gathering script 325
Post-data gathering script 326
Pre-Design phase script 326
Post-Finalization phase script 327
Pre-replication script 327
Post-replication script 328
Pre-failover script 328
Post-failover script 329
Pre-scan script 329
Post-scan script 330
Pre-snapshot script 330
Post-snapshot script 330
Pre-snapshot script 331
Post-snapshot script 331
Pre-sync script 331
Post-sync script 332
User decision script 332

Chapter 13 Server Availability

In this chapter 335

Plan for VFM availability 335

Selecting a backup method 335
Back up the namespace 336
Schedule snapshots of the state of the VFM server 336
Manually back up the VFM server configuration files 336

Back up and restore VFM configuration files 336
Back up VFM configuration files 336
Restore VFM configuration to a known state 337

Backing up VFM configuration files 337

Importing VFM configuration files 339

Restoring VFM on the primary machine 339

Performing subsequent failover on the same backup machine . . 340

Appendix A Using RSH and SSH to Execute Shell Commands

In this appendix 341

Using RSH 341
Enabling RSH as root on a Linux system 341
Enabling RSH as root on a Solaris system 343

Using SSH 343
SSH prerequisites 343
SSH best practices for VFM. 344
Configuring VFM to use SSH 345
Installing and configuring OpenSSH on Solaris 353
Installing and configuring OpenSSH on Red Hat 355
Configuring default remote shell credentials. 355

Appendix B Inclusion and Exclusion Tables

In this appendix. 357

Inclusion lists. 357
Inclusion rules 357
Examples of valid inclusion formats 358

Exclusion lists 359
Exclusion rules 359
Examples of valid exclusion formats 360

Appendix C Example Migration Tasks Script

In this appendix. 361

Specifying migration tasks from a file 361
Task creation batch file arguments 361
Task File Description Syntax 362

Example task creation batch file 363

Task description file 364

Appendix D Troubleshooting

In this appendix. 367

Finding solutions 367
General problems 368
DFS root problems 370
Report problems 371
Target problems 372
Windows XP and Windows 2000 incompatibility problem . . .372
Problems connecting to the VFM Server service 372
Problems selecting a monitoring agent 373
Replication and migration problems 373
NetApp filer problems 384
Browse namespace problems. 386
Windows NAS or storage server filer identification problems.386
Problems running batch files 386
Problems running a policy 387
Snapshot access policy problems. 387
Problems printing to PDF files. 388
Name of machine is incorrect. 388
Problems taking snapshots of the application 389
Permission-based namespace policy problems 389
Problem installing a UNIX replication agent. 390
Problems using NTFS Change Journal 390
Problem configuring the OpenSSH package on Solaris 391
Problem accessing a UNIX machine when using SSH. 391
Problem importing snapshots of the application. 391
Problems replicating data using VSS snapshots 392
Problems deleting items from the Logical View 392
Problem updating home directories 393

Migration, replication, and status indicators 393
Migration task indicators. 393
Replication status monitoring indicators 394
Status indicators 395

Exclusion and Inclusion Examples 399
Exclusion examples 399
Inclusion examples 401

Appendix E Event Descriptions

In this appendix. 403

1000 - 1099 Multiprotocol Namespace Event IDs 404

1100 - 1199 Home Folder Event IDs. 405

1200 - 1299 Disaster Recovery Event IDs 405

1300 - 1399 Replication Agent Event IDs. 407

1400 - 1499 Server Event IDs 408

1500 - 1699 SnapMirror Event IDs 408

1700 - 1799 Namespace Availability Event IDs 410

1800 - 1899 Multiprotocol Namespace Event IDs. 411

1900 - 1999 Service Event IDs 411

2000 - 2099 Migration Event IDs 412
2100 - 2199 Namespace Creation Event IDs 413
2200 - 2299 Snapshot Access Event IDs 414
2300 - 2399 Filer Disaster Recovery Event IDs 414
2400 - 2499 Snapshot Scheduling Event IDs. 416
2500 - 2599 Namespace Backup Event IDs 417
2600 - 2699 Logical Namespace Event IDs 417
2700 - 2799 Archival Migration Event IDs 419
2800 - 2899 Server Disaster Recovery Event IDs 420
2900 - 2999 Replication and Monitoring Agents Event IDs 422
3100 - 3199 Storage Load Balancing Event IDs. 423
3200 - 3299 Permission-Based Namespace Event IDs. 424

Chapter

Introduction to VFM

1

This section provides a high-level overview of VFM that includes information about how it works.

In this chapter

- What is VFM? 1
- How VFM helps you 1
- VFM components and architecture 3
- Storage policies 4

What is VFM?

VFM 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 the following data management functions:

- Heterogeneous migration
- Local and remote replication
- Client failover
- Capacity optimization
- Tiered data management
- Disk-to-disk archival
- 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 the NetApp SnapMirror® and Snapshot™ technology products to deliver best-in-class management capabilities for distributed file storage environments.

VFM simplifies enterprise file storage management in both multiprotocol client environments, such as in UNIX® and Window environments, and in heterogeneous server/filer environments, such as in environments with Windows, Windows network-attached storage (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.

How VFM helps you

VFM provides many important features to help you effectively manage your storage systems.

Simplifies 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 knowing the physical file locations. VFM simplifies enterprise storage management using a global namespace.

Users no longer need to access their data using drive letter mappings. The VFM global namespace permanently eliminates drive letter dependencies, giving administrators tremendous flexibility to expand, move, and reconfigure physical storage without affecting users. Finally, VFM achieves this by aggregating file storage across heterogeneous storage platforms, such as Windows and filer systems, and by providing a unified logical view of the storage. VFM enables role-based hierarchical administration and multi-user management.

Enables non-disruptive expansion and consolidation

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 storage resources. 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.

Provides intelligent 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 provides options to transfer 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.

Ensures business continuity with rapid disaster recovery

VFM provides 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. In addition, VFM provides a comprehensive automated solution for deploying and managing NearStore® solutions. You assign a NearStore device as the target for multiple primary filers and Windows servers, including Windows-attached storage-area networks (SANs). VFM automatically synchronizes qtrees and shares present on multiple source devices with the NearStore device.

Centralizes routine filer configuration and management

VFM enables the centralized configuration, viewing, scheduling, and management of filer snapshots and SnapMirror pairs across multiple filers. In this manner, VFM makes taking snapshots across multiple filers a one-step process and enables data availability based on business priorities. VFM also provides a consolidated view of multiple filers from a single console, and generates a range of descriptive reports that provide tremendous insight into the filer storage environment.

Reduces ownership costs with a tiered storage architecture

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.

VFM components and architecture

VFM includes the following components:

VFM client

Provides the VFM administrator interface or console and communicates with the VFM server. Multiple VFM clients can connect to the VFM server to manage the namespace and policies.

VFM server

Executes policies and maintains storage configuration and reporting information. The server is implemented as a Windows service.

VFM monitoring agents

Allows distributed processing of policies. Monitoring agents are implemented as a Windows service.

VFM replication agents

Executes file transfers for policies. VFM provides Windows and UNIX replication agents. You can quickly deploy agents with the VFM agent management tools.

- Windows replication agents are implemented as Windows services. Windows replication agents transfer CIFS data.
- UNIX replication agents are available for Red Hat® Enterprise 4 and Solaris® 10. Unix replication agents transfer NFS data.

VFM uses a distributed client/server architecture. This design provides a flexible and efficient means to manage Windows storage. The following figure shows the VFM architecture:

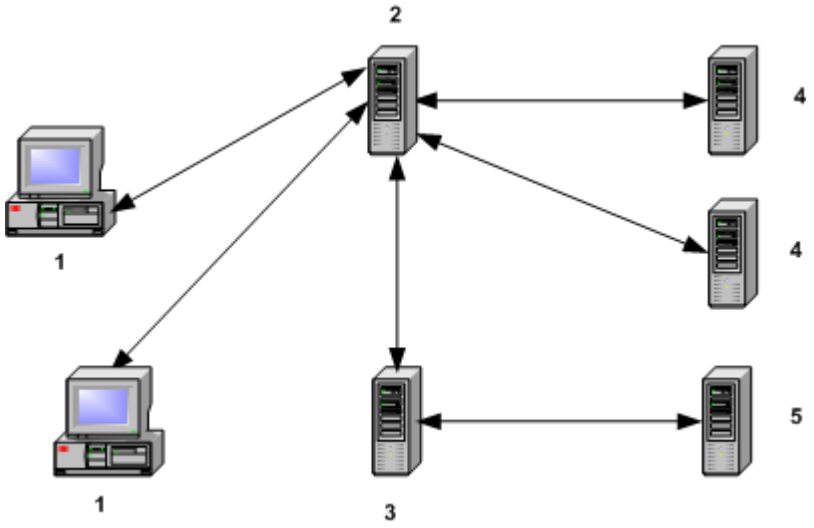

FIGURE 1 VFM architecture

1 VFM client
2 VFM server and DFS root
3 Monitoring agent
4 Windows replication agent
5 Unix replication agent

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 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 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 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.

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 or SnapMirror for NetApp filers. 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. 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 and in Active Directory. VFM adds user home directories to a logical namespace after querying Active Directory, 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.

Feature comparison checklist

The features provided by VFM depend on the version you have. The tables in the following sections identify the features available in VFM Enterprise Edition (EE) and VFM Migration Edition (ME).

Data migration features

The following table identifies the data migration features available in VFM Enterprise Edition (EE) and VFM Migration Edition (ME).

TABLE 1 Data migration features

Feature	VFM ME	VFM EE
Transfer data in heterogeneous environments. Create, edit, and delete migration policies.	X	X
Support bandwidth throttling.	X	X
Support NTFS Change Journal.	X	X
Support CIFS protocol.	X	X
Options for controlling how a file's security is copied.	X	X
Option to update the namespace.	X	X
Option to run batch files before and after a migration.	X	X
Schedule a migration.	X	X
Start a migration manually.	X	X
Manage Windows replication agents.	X	X
Manage UNIX replication agents.		X
Centrally manage tasks associated with replicating data across a network. Create, edit, and delete replication policies.		X
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.		X
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.		X
Support differential replication.		X
Support NFS protocol.		X
Support SnapLock integration.		X
Specify migration tasks from a file.		X

Namespace features

The following table identifies the namespace features available in VFM Enterprise Edition (EE) and VFM Migration Edition (ME).

TABLE 2 Namespace features

Feature	VFM ME	VFM EE
Create, edit, and delete stand-alone DFS roots.	X	X
Create, edit and delete domain-based DFS roots.	X	X
Create, edit, and delete consolidation DFS roots.	X	X
Create, edit, and delete DFS links.	X	X
Create, edit, and delete shares.	X	X

TABLE 2 Namespace features

Feature	VFM ME	VFM EE
Create and remove logical folders.	X	X
Create domain-based root targets (replicas).	X	X
View the machines, links, and folders targeted by a root.	X	X
View the machines and shares targeted by a link.	X	X
Toggle link targets online or offline.	X	X
Check and clear the status of an item in the namespace.	X	X
Expand a link.	X	X
Monitor roots, links, and link targets.		X
Fail over or fail back a link.		X
Manually back up and restore a root.		X
Create namespaces automatically. Create, edit, and delete namespace creation policies.		X
Create namespaces based on users' security permissions to resources. Create, edit, and delete permission based namespace policies.		X
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.		X

Reporting features

The following table identifies the reporting features available in VFM Enterprise Edition (EE) and VFM Migration Edition (ME).

TABLE 3 Reporting features

Feature	VFM ME	VFM EE
Generate file differences reports.	X	X
Schedule reports.	X	X
Start data collection manually.	X	X
Run customized commands before and after a report schedule runs.	X	X
Print a report.	X	X
Create other reports.		X
View a list of reports referencing a selected root, link, logical folder, machine, or share.		X
Publish reports.		X
Archive reports.		X

Business continuity (Disaster recovery) features

The following table identifies the business continuity features available in VFM Enterprise Edition (EE) and VFM Migration Edition (ME).

TABLE 4 Business continuity features

Feature	VFM ME	VFM EE
Detect a failure and automatically fail over clients to a replicated storage device. Create, edit, and delete client recovery policies.		X
Set up a filer or nearline storage device as a disaster recovery backup for one or more primary filers. Create, edit, and delete filer recovery policies.		X
Make data on servers participating in a namespace highly available. Create, edit, and delete server recovery policies.		X

Physical resource management features

The following table identifies the physical resource management features available in VFM Enterprise Edition (EE) and VFM Migration Edition (ME).

TABLE 5 Physical resource management features

Feature	VFM ME	VFM EE
Manage machines running Windows 2000 and Windows Server 2003.	X	X
Manage NetApp filers running Data ONTAP 6.5 and later.	X	X
Create a share.	X	X
Create a volume.	X	X
Refresh filer information.	X	X
Centrally manage snapshots on NetApp filers. Create, edit, and delete snapshot access and snapshot scheduling policies.		X
Manage UNIX machines running Solaris 10 and RedHat Enterprise 4.		X

Application and integration features

The following table identifies the application and integration features available in VFM Enterprise Edition (EE) and VFM Migration Edition (ME).

TABLE 6 Application and integration features

Feature	VFM ME	VFM EE
Connect to another server.	X	X
Configure application access permissions.	X	X
Check system status.	X	X
Integrate with FLM.	X	X
View the properties of items in the namespace.	X	X
View the status of selected items.	X	X
View event history.	X	X

TABLE 6 Application and integration features

Feature	VFM ME	VFM EE
Printer setup.	X	X
Export/import application configuration files.	X	X
Back up the application configuration files (using the client).		X
Launch external tools.		X
Use drag-and-drop operations.		X
View the contents of DFS links, shares, machines, or folders.		X
View the machines, links, and folders targeted by a root.		X
View the machines and shares targeted by a link.		X

Chapter

Planning Considerations

2

This section provides information you need to consider before you install VFM, create a namespace, move data, or make data highly available.

In this chapter

- VFM availability considerations 11
- Namespace design considerations 12
- Data migration considerations 14
- Disaster recovery planning considerations 20
- Replication planning considerations 22
- Namespace availability planning considerations 24
- Namespace creation policy planning considerations 26
- Multiprotocol namespace planning considerations 27
- Home folder planning considerations 31

VFM availability considerations

When you install VFM, consider the following aspects of your storage configuration to ensure VFM is highly available.

A convenient location for the VFM console

Typical locations for the console are the storage administrator's management desktop or on a system convenient to storage resources being managed by VFM.

A reliable, highly available location for the VFM server

The VFM server and any monitoring agents should typically be installed on servers rather than on desktop machines. These components must be running for policies to function. The risk with installing on desktop machines is that they might be shut down at the close of normal working hours rather than left operating at all times.

If the VFM server or monitoring agent for a policy becomes unavailable, neither disaster recovery failover nor VFM data synchronization can be performed because of the following reasons:

- VFM communicates with the DFS root to initiate a failover.
- VFM communicates directly with the storage targets to initiate data synchronization.

How large your VFM configuration is

If there is any chance that the number of connections used between VFM components in conjunction with the connections used by other processes on the system could exceed the limits on inbound connections imposed by Windows 2000 Professional or Windows XP Professional, install the VFM server on Windows 2000 Server or Windows Server 2003. For more information about inbound connection limits, refer to Microsoft® Knowledge Base article 122920 and the Microsoft Knowledge Base article 314882.

Whether you have policies to run from remote location

If you have policies that you want to run from a location that is separate from where the VFM server is installed, you have several important considerations. Based on the physical locations and network connectivity between components in your storage infrastructure, it might be advantageous to install monitoring agents at multiple locations. This deployment allows the monitoring workload to be distributed among multiple servers. Depending on the specifics of individual configurations, this deployment can also allow monitoring to take place over a LAN segment rather than requiring that it be performed over slower and more expensive WAN connections.

Where the managed DFS roots are located

Consider the location of VFM components and its connectivity in relation to the locations of the DFS root servers.

VFM network communication

VFM transfers file data using standard Microsoft networking protocols. The VFM components access remote files and directories through their UNC path names. This access requires that the normal ports needed for Microsoft networking must be opened in any firewall between the VFM client, VFM server, replication agent, and the servers hosting data files to be replicated or migrated.

In addition to the Microsoft networking ports, VFM uses two TCP ports for communication between its components. The VFM server and monitoring agent listen on TCP ports 6001 and 6005. The VFM replication agent listens on TCP port 6002. Therefore, these ports must be opened in any firewalls that are present between these VFM components.

Namespace design considerations

Before you create a namespace, you should perform the tasks summarized in the following topics:

- "Plan the layout and design of the namespace" on page 13
- "Select the root configuration you want to use" on page 13
- "Populate the namespace with links to network shares" on page 14
- "Instruct users to access data through the namespace" on page 14

Plan the layout and design of the namespace

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.

TABLE 7 Namespace Layout Alternatives

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

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

Just as a single namespace can provide different paths to the same data, you might consider creating multiple namespaces to serve the needs of different audiences. Examples of these could be namespaces for web content, software distribution, application configuration, and so forth. Keep the namespace simple.

In addition to the namespace organization decisions, reliability requirements must be considered during logical namespace design. You might want to replicate the namespace to allow it to be available from multiple servers. This eliminates the server that provides the namespace from being a single point of failure.

A single namespace server can handle about 2,000 to 4,000 users, depending on 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.

Select the root configuration you want to use

The root is the top of the logical hierarchy. You can create domain-based DFS roots, stand-alone DFS roots, consolidation DFS roots, and File Redirector roots. For a definition for each root and instructions about how to create the roots, refer to the online help.

Populate the namespace with links to network shares

The following list explains the ways that you can populate the namespace with links.

- Add a link dialog box enables you to add one link at a time.
- Drag UNC paths from Windows Explorer to a root in the Logical View.
- Add multiple links dialog box enables you to add multiple links and folders at a time. Use this method when you know where items are located on your network and most of the DFS shares that you want to link are located on a few machines. A wizard steps you through the process. You can drag items from Windows Explorer to a folder.
- Namespace Creation Policy wizard enables you to add multiple links and folders at a time. Use this method when you want to add links to multiple shares that are supported by different machines and when you want to keep the logical namespace aligned with subsequent changes in the underlying physical resources.
- Permission Based Namespace wizard enables you to add multiple links and folder at a time like the Namespace Creation Policy wizard. Use this method when you want to populate the namespace based on security permissions of shares.

Instruct users to access data through the 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.

Data migration considerations

This section discusses data migration requirements, the transfer of NTFS security attributes, and the transfer results when different file systems, such as FAT or NFS, are involved.

Data migration requirements

Before you create and run policies to move data, review the following list of requirements:

- Users and applications should not be accessing the source or destination folders while the migration is taking place.
- Do not use differential replication on a LAN. When operating on a high-speed network, the computational cost of calculating the file differences is more expensive in terms of time than simply copying the entire file from source to destination.
- When you perform full backups during a migration operation, backups can change archive attribute settings and trigger massive differences between source and destination. With VFM you can avoid this problem, because migration policies enable you to specify how you want to set attributes of destination files.
- When configuring replication options, consider carefully how much event detail you need. If you gather information about all files, it affects performance. As a compromise between performance and the need for feedback, it is recommended that you gather information only when an error is encountered.

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

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

Transfer NTFS file security attributes

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

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

Transfer data between dissimilar file systems

Some file attribute information, such as security attributes, might be lost when copying or moving files from one file system to another because of the inherent differences between the file systems. The following tables display transfer results between different file systems.

NTFS Windows 2000 is the source

The following table lists transfer results from NTFS Windows 2000 source files.

TABLE 8 NTFS Windows 2000 transfer results

Destination	Attribute Information
NTFS W2K	Attributes are maintained. Copying hard links, reparsing points, or NTFS encrypted files are not supported.
NTFS NT4	Some attributes are lost.
FAT16	• Some attributes are lost. • NTFS stores access times in UTC, while FAT stores access time in LocalTime. VFM does not adjust between UTC and LocalTime. • FAT maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to a FAT file system from NTFS, their access times, extended attributes, indexing attributes, and alternate streams (including security) are lost when the files are written to the FAT destination.
FAT32	• Some attributes are lost. • NTFS stores access times in UTC, while FAT stores access time in LocalTime. VFM does not adjust between UTC and LocalTime. • FAT maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to a FAT file system from NTFS, their access times, extended attributes, indexing attributes, and alternate streams (including security) are lost when the files are written to the FAT destination.
Filer CIFS	You can configure security on filers in a variety of ways. These variables affect how security attributes are maintained when copying to filers. When files are transferred from Windows systems to a filer and sparse or compressed attributes are set, the attributes have no effect on how the files are stored on the filer.
NFS	• Some attributes are lost. • NFS maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to an NFS file system from NTFS, their security and extended attributes are subsumed by NFS attributes. VFM UNIX replication agents perform NFS to NFS data transfer. Other file systems require redirector client software on the machine hosting the Windows replication agent. • Copying files from NFS sources is supported. This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.

NTFS Windows NT 4 is the source

The following table lists transfer results from NTFS Windows NT 4 source files.

TABLE 9 NTFS Windows NT 4 transfer results

Destination	Attribute Information
NTFS W2K	Attributes are maintained.
NTFS NT4	Attributes are maintained.

TABLE 9 NTFS Windows NT 4 transfer results

Destination	Attribute Information
FAT16	• Some attributes are lost. • NTFS stores access times in UTC, while FAT stores access time in LocalTime. VFM does not adjust between UTC and LocalTime. • FAT maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to a FAT file system from NTFS, their access times, extended attributes, indexing attributes, and alternate streams (including security) are lost when the files are written to the FAT destination.
FAT32	• Some attributes are lost. • NTFS stores access times in UTC, while FAT stores access time in LocalTime. VFM does not adjust between UTC and LocalTime. • FAT maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to a FAT file system from NTFS, their access times, extended attributes, indexing attributes, and alternate streams (including security) are lost when the files are written to the FAT destination.
Filer CIFS	You can configure security on filers in a variety of ways. These variables affect how security attributes are maintained when copying to filers. When files are transferred from Windows systems to a filer and sparse or compressed attributes are set, the attributes have no effect on how the files are stored on the filer.
NFS	• Some attributes are lost. • NFS maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to an NFS file system from NTFS, their security and extended attributes are subsumed by NFS attributes. VFM UNIX replication agents perform NFS to NFS data transfer. Other file systems require redirector client software on the machine hosting the Windows replication agent. • Copying files from NFS sources is supported. This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.

FAT16 is the source

The following table lists transfer results from FAT16 source files.

TABLE 10 FAT16 transfer results

Destination	Attribute Information
NTFS W2K	Attributes are maintained.
NTFS NT4	Attributes are maintained.
FAT16	Attributes are maintained.
FAT 32	Attributes are maintained.

TABLE 10 FAT16 transfer results

Destination	Attribute Information
Filer CIFS	You can configure security on filers in a variety of ways. These variables affect how security attributes are maintained when copying to filers. When files are transferred from Windows systems to a filer and sparse or compressed attributes are set, the attributes have no effect on how the files are stored on the filer.
NFS	• Some attributes are lost. • NFS maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to an NFS file system from NTFS, their security and extended attributes are subsumed by NFS attributes. VFM UNIX replication agents perform NFS to NFS data transfer. Other file systems require redirector client software on the machine hosting the Windows replication agent. • Copying files from NFS sources is supported. This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.

FAT32 is the source

The following table lists transfer results from FAT32 source files.

TABLE 11 FAT32 transfer results

Destination	Attribute Information
NTFS W2K	Attributes are maintained.
NTFS NT4	Attributes are maintained.
FAT16	Some attributes are lost.
FAT32	Attributes are maintained.
Filer CIFS	You can configure security on filers in a variety of ways. These variables affect how security attributes are maintained when copying to filers. When files are transferred from Windows systems to a filer and sparse or compressed attributes are set, the attributes have no effect on how the files are stored on the filer.
NFS	• Some attributes are lost. • NFS maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to an NFS file system from NTFS, their security and extended attributes are subsumed by NFS attributes. VFM UNIX replication agents perform NFS to NFS data transfer. Other file systems require redirector client software on the machine hosting the Windows replication agent. • Copying files from NFS sources is supported. This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.

Filer using the CIFS protocol is the source

The following table lists transfer results from Filer CIFS source files.

TABLE 12 File running CIFS protocol transfer results

Destination	Attribute Information
NTFS W2K	Attributes are maintained.
NTFS NT 4	Attributes are maintained.
FAT16	Some attributes are lost.

TABLE 12 File running CIFS protocol transfer results

Destination	Attribute Information
FAT32	Some attributes are lost.
Filer CIFS	Attributes are maintained.
NFS	• Some attributes are lost. • NFS maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to an NFS file system from NTFS, their security and extended attributes are subsumed by NFS attributes. VFM UNIX replication agents perform NFS to NFS data transfer. Other file systems require redirector client software on the machine hosting the Windows replication agent. • Copying files from NFS sources is supported. This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.

NFS is the source

The following table lists transfer results from NFS source files.

TABLE 13 NFS transfer results

Destination	Attribute Information
NTFS W2K	Copying files from NFS sources is supported This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.
NTFS NT4	Copying files from NFS sources is supported This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.
FAT16	Copying files from NFS sources is supported This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.
FAT32	Copying files from NFS sources is supported This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.
Filer CIFS	Copying files from NFS sources is supported This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.
NFS	Copying files from NFS sources is supported This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.

NetWare is the source

The following table lists transfer results from NetWare source files.

TABLE 14 NetWare source transfer results

Destination	Attribute Information
NTFS W2K	• Some attributes are lost. • When copying NetWare source files, some attributes are lost. Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained. If the version of NetWare is 5.1 or earlier, the Novell client must be installed on the same machine on which the replication agent is installed. Windows and Novell accounts must be configured with identical user name and passwords. The VFM console and server components must also run under the same user name.
NTFS NT4	• Some attributes are lost. • When copying NetWare source files, some attributes are lost. Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained. If the version of NetWare is 5.1 or earlier, the Novell client must be installed on the same machine on which the replication agent is installed. Windows and Novell accounts must be configured with identical user name and passwords. The VFM console and server components must also run under the same user name.
FAT16	• Some attributes are lost. • When copying NetWare source files, some attributes are lost. Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained. If the version of NetWare is 5.1 or earlier, the Novell client must be installed on the same machine on which the replication agent is installed. Windows and Novell accounts must be configured with identical user name and passwords. The VFM console and server components must also run under the same user name.
FAT32	• Some attributes are lost. • When copying NetWare source files, some attributes are lost. Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained. If the version of NetWare is 5.1 or earlier, the Novell client must be installed on the same machine on which the replication agent is installed. Windows and Novell accounts must be configured with identical user name and passwords. The VFM console and server components must also run under the same user name.
Filer CIFS	• Some attributes are lost. • When copying NetWare source files, some attributes are lost. Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained. If the version of NetWare is 5.1 or earlier, the Novell client must be installed on the same machine on which the replication agent is installed. Windows and Novell accounts must be configured with identical user name and passwords. The VFM console and server components must also run under the same user name.
NFS	You cannot copy NetWare source files to NFS destinations.

Disaster recovery planning considerations

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

General disaster recovery considerations

When planning for disaster recovery, consider the following areas:

- **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.

Business continuity policy requirements

Before you create and run a replication policy, review the following list of requirements:

- Links used in recovery policies must reside in Windows 2000 or Windows Server 2003 roots. This is because recovery policies rely on the ability to toggle link targets online and offline.
- The targets of links used in recovery policies must reference locations that have NTFS security.
- When you add a link to a recovery policy using SnapMirror replication, the link should have only two targets. Make sure that one target is the source and the other target is the destination of the SnapMirror.
- The account used by the VFM server or monitoring agent, as well as any accounts used to run the VFM client, must have permission to make modifications to the DFS root containing the links in the recovery policy. If the account lacks these permissions, VFM cannot control the online/offline state of the link targets.
- If a DFS root served by a single system is used and becomes unavailable, failover cannot be performed and users can access only the primary storage target until their referral times out. Adding additional servers to host copies of the namespace provides fault tolerance. To create copies of the namespace, you create DFS root replicas in the case of domain-based roots, or you can set up namespace availability policies to synchronize the logical structure of two or more stand-alone roots. An alternate approach for creating a highly available root is to host a stand-alone root on a Windows 2000 cluster server.
- If the VFM server or monitoring agent for a client or filer recovery policy becomes unavailable, neither failover nor data synchronization can be performed. User access to the primary storage target is not affected.
  - VFM communicates with the DFS root to initiate a failover.
  - VFM communicates directly with the storage targets to initiate data synchronization.
- The VFM client does not have to be running for recovery policies to function.
- Carefully plan where each VFM component (for example, client, server, monitoring agent, and replication agents) should be deployed.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files.
- Roots containing links participating in recovery policies can also be members of namespace availability policies. In this configuration, as with recovery policies containing links from replicated domain-based roots, the recommendation is that you configure the recovery policies for manual failover rather than for automatic failover if the links target read/write data rather than read-only data.
- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.

Resynchronizing SnapMirrors for disaster recovery failback

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

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

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

Replication planning considerations

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

Replication policy requirements

VFM uses file-to-file copy to synchronize data across a network. Before you create and run a replication policy, review the following list of requirements:

- Exclusive locks on files on source or destination prevent VFM from accessing the file for transfer. Find out who or what program is locking the files, and make the files available to VFM.
- It is not recommended that you perform full backups during a replication operation because backups can change archive attribute settings and trigger massive differences between source and destination.
- When configuring replication options, carefully consider how much event detail you need. If you gather information about all files, it affects performance. As a compromise between performance and the need for feedback, it is recommended that you gather information only when an error is encountered.

- When planning replication, consider the available time you have to complete the operation and how much data you have to replicate. Use include and exclude lists sparingly because they add processing time to the replication operation.
- It is not recommended that you enable real-time virus scanning on filers because replication throughput can be severely reduced.
- The benefits of differential replication are realized on slower links rather than on a LAN. Computational costs outweigh the benefits of differential replication on a LAN.
- A replication policy synchronizes from its source to a number of read-only copies.
- There is a replication time delay between primary and secondary contents. If a primary target becomes unavailable, there is no guarantee that all the contents of the primary target are replicated to the standby targets.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files and can interfere with VFM replication.
- The VFM client does not have to be running for replication policies to function.

How replication policies work

This section explains the steps necessary for creating and running a replication policy to move data.

1. Create a replication policy and define a schedule.
2. Add a single-target link to the policy as the source link. You can add the link during the creation of the replication policy or drag it from the Logical View tree or from the right pane to the policy.
3. Add a multi-target link to the policy as the destination link. You can add the link during the creation of the replication policy, or drag it from the Logical View tree or from the right pane to the policy. You can quickly add additional targets to the destination link.
4. Update the source or destination from the **Replication Topology** tab. When you select a replication policy in the Admin View, several tabs appear in the right pane. The **Replication Topology** tab provides a graphical display of the distribution path of the policy. From this tab, you can perform the following actions:
   - You can change the distribution path by a simple drag-and-drop operation.
   - You can view monitoring status of the links and each of their targets.
   - You can add additional targets to the destination link by right-clicking the link.

Transfer NTFS file security attributes

For information about transferring file security attributes, refer to "Data migration considerations" on page 14.

Permissions

For information about setting permissions for the account used by the VFM server, a replication agent, or a monitoring agent, refer to "Data migration considerations" on page 14.

Status monitoring

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

If status monitoring is turned on for the links and folders contained in a policy, status indicators appear for the items on the **Replication** tab that appears when you select a policy.

Monitoring DFS resources generates network traffic between the VFM console and the resources being monitored. Slower polling rates generate less traffic but provide less-timely notification of problems in the monitored namespace.

Namespace availability planning considerations

Namespace availability policies enable you to easily keep the logical structure of two or more DFS roots and/or widelink namespaces aligned.

By synchronizing roots, you can load balance the work on your DFS root server across multiple identical instances of the root. This also provides tolerance because the namespace is available from more than one server.

By synchronizing a master root with a widelink namespace, you can host your namespace on a filer and your CIFS clients can access the namespace directly on the filer.

Namespace availability requirements

Before you create and run a namespace availability policy, review the following list of requirements:

- Filers must run Data ONTAP 6.4 software or later to host a widelink namespace.
- For performance reasons and to avoid possible conflict with widelinks that might already exist under an existing share, it is recommended that you create and share a directory for hosting a widelink namespace instead of using an existing share. The contents of a share assigned to host a widelink namespace are scanned and all widelinks found on the share are managed.
- When synchronizing between roots in different domains, install the monitoring agent with an account that has privileges to read from the primary root and write to the secondary (slave) roots of the policy.
- When synchronizing between a master root and secondary roots in separate domains, you might need to set up multiple policies with a common master root and individual slave root in the domain in which the monitoring agent is privileged. In this configuration, the accounts for monitoring agents must have at least read privileges to access the master root. When running this configuration, you must use scheduled synchronizations instead of the manual start synchronization now operation. This is because the start synchronization now operation uses the monitoring agent on the VFM server due to the dynamic feedback provided by the manually initiated operation
- Links in a widelink namespace are restricted to having one target on line. Link timeout is set to 6 minutes by Data ONTAP, and comments are not supported.

How namespace availability policies work

This section explains the steps necessary for creating and running a namespace availability policy to keep the logical structure of two or more namespaces aligned.

1. Create a namespace availability policy and add the roots or root and CIFS share where you will host the widelink namespace. You must have at least two DFS roots or one DFS root and one CIFS share on a filer
2. Designate a root as the master.
3. Schedule synchronization.

   During synchronization, the logical structure of the secondary roots and widelink namespaces are compared with that of the master root. If they are not the same, the logical structures of the secondary namespaces are synchronized with that of the master root.

Options for setting up namespace availability policies

This section discusses options for setting up namespace availability policies to make your namespace highly available:

Users access the namespace through the widelink namespace, and the master root is only used for administrative purposes. Whenever changes need to be made to the logical structure of the namespace, the changes are made on the master root. When the policy runs, the logical structure of the master root is synchronized with the widelink namespace.

You can use the master root server of a namespace availability policy as the production server with the secondary root server held in reserve as backup in case the master root fails. If there is a failure, the secondary root server can assume the identity of master root server, enabling users to access the namespace at what they perceive as the normal location.

You can make users aware of both the master root server and secondary root server as hosts of the namespace so that if one server fails, they know to access the namespace through its peer.

You can run the namespace availability policies either alone or in conjunction with an IP load-balancing solution, such as NLB or an IP load-balancing switch.

You can use stand-alone DFS roots in a namespace availability policy in conjunction with DNS to provide the locality of reference benefits typically provided by domain-based DFS roots. To set up this configuration, follow these instructions:

1. Each region in which one of the roots in the namespace availability policy is located must have its own DNS server, and all clients and servers within a region access only their region's DNS server.
2. Disable zone transfers between DNS servers for the server record so that each DNS server maintains its individual notion of the IP address of the DFS root server for the region.
3. The servers hosting each root in the namespace availability policy must have the same host name but a different IP addresses. All users access the root through the same server name, but users are directed to the root closest to them.
4. Add the individual roots to the namespace availability policy using IP addresses rather than host names, so that VFM can access each of the roots being synchronized.
5. Use VFM to synchronize the logical structure of the roots.

Namespace creation policy planning considerations

Namespace creation policies and permission-based namespace policies provide automated synchronization between the logical namespace and supporting physical resources. The logical namespace is automatically synchronized with changes in physical resources.

The planning considerations for namespace creation policies also apply to permission-based namespace policies.

Namespace creation policy requirements

Before you create and run a namespace creation policy, review the following list of requirements:

- If a namespace creation policy scans for shares and finds a share that is already referenced by a link in the namespace, the policy does not create a link to the share. This enables you to move links created by the policy to different locations in your logical namespace and not have them recreated in their original location when the policy runs again.
- In addition to automatically adding links as part of scheduled execution of the policy, you can configure the policy to automatically delete links targeting shares that are removed from servers scanned by the policy. The shares for the servers in the policy are enumerated and links for shares that no longer exist on the server can be removed. You can configure the policy to remove links just from the subfolder in the logical namespace specified in the policy or to remove such links located anywhere in the root.
- During the following conditions, namespace creation policies might interfere with one another when items are added to the namespace:
  - Two or more policies are configured with selection criteria that allow search results to match more than one policy.
  - Two or more policies are configured to manage the same logical folder in the namespace.
  - Two or more policies are scheduled to run simultaneously.

  These conditions might result in sporadic failures when one of the policies attempts to add links to the namespace. These failures are benign and can be avoided by staggering the schedule of the policies involved.

How namespace creation policies work

The Namespace Creation wizard guides you through the process of configuring namespace creation policies. This section provides an overview of how a namespace creation policy works.

1. Name the policy and select whether you want the policy to automatically add links to the namespace that reference the shares that match the search criteria. You can also select whether you want the policy to automatically remove links whose shares no longer exist.
2. Select or create a DFS root to contain the links that the policy adds to the namespace.
3. Select the domain or servers that you want to search for shares.
4. Choose the location in your logical namespace where you want the links created by the policy to be added.
5. Enter the share search criteria. You can combine search criteria by separating search patterns with semicolons. The search string mkt*;sales* matches shares whose names begin with either `mkt` or `sales`.

6. Set the schedule to scan periodically for shares matching the search criteria.

   When a share matching the search criteria is found, a link is created and added to the namespace.

Comparison with 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 a search by group permissions, whereas namespace creation policies do not provide this option.

Multiprotocol namespace planning considerations

Multiprotocol namespace policies enable you to create a common namespace for UNIX® and Windows users to provide access to files stored on 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.

Prerequisites

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

- You have DFS roots containing links that target data on filers.
- You must be familiar with the NFS protocol and UNIX system.
- You must be familiar with configuring and managing filers.
- Your filers must be licensed for both CIFS and NFS protocols.

How multiprotocol namespace policies works

This section explains the steps necessary for creating and running a multiprotocol namespace policy.

1. Provide the name of a DFS root with links targeting data on filers.
2. Provide the CIFS UNC path name to the directory that contains the automount and exports files that configure the UNIX namespace.

   NOTE
   Direct and indirect automount maps are supported. The multiprotocol namespace enables users to access directories hosted on filers that are accessible through both NFS exports and CIFS shares.

3. Create sync points, UNIX and DFS namespace resources that are synchronized. Indicate the direction of the synchronization of the logical structure of the namespace. Synchronization can be configured from a DFS namespace to a UNIX namespace, or from a UNIX namespace to a DFS namespace. Select one of the following options:
   - If the UNIX namespace is the master, you provide the location of the automount files that configure the namespace.
   - If the DFS namespace is the master, you provide a location where VFM places the automount files and export files created by the multiprotocol namespace policy.
4. Schedule sync point synchronization to update the logical structure from each synch point master.

UNIX namespace configuration

The following example illustrates how a UNIX namespace is configured when the DFS namespace is the master and how NFS users access the same data on filers that CIFS users access.

Begin by creating a multiprotocol namespace policy to perform the following actions:

- Automatically create the necessary UNIX files, apply exports, and push automount files to NIS.
- Create a sync point that specifies the DFS and UNIX namespace resources you want to synchronize.
- Configure the policy to periodically synchronize from the DFS namespace to the UNIX namespace.

Provide the name of a DFS root with links targeting data on filers

In the following example, *National Sales* is a DFS logical folder in *DFS Root* that contains the *Marketing*, *Presentations*, *SalesReg1*, *SalesReg2*, *SalesReg3*, and *SalesReg4* links. Only the *SalesReg* links reference CIFS shares on filers. You configure a multiprotocol namespace policy to enable NFS users to also access the shares referenced by the *SalesReg* links.

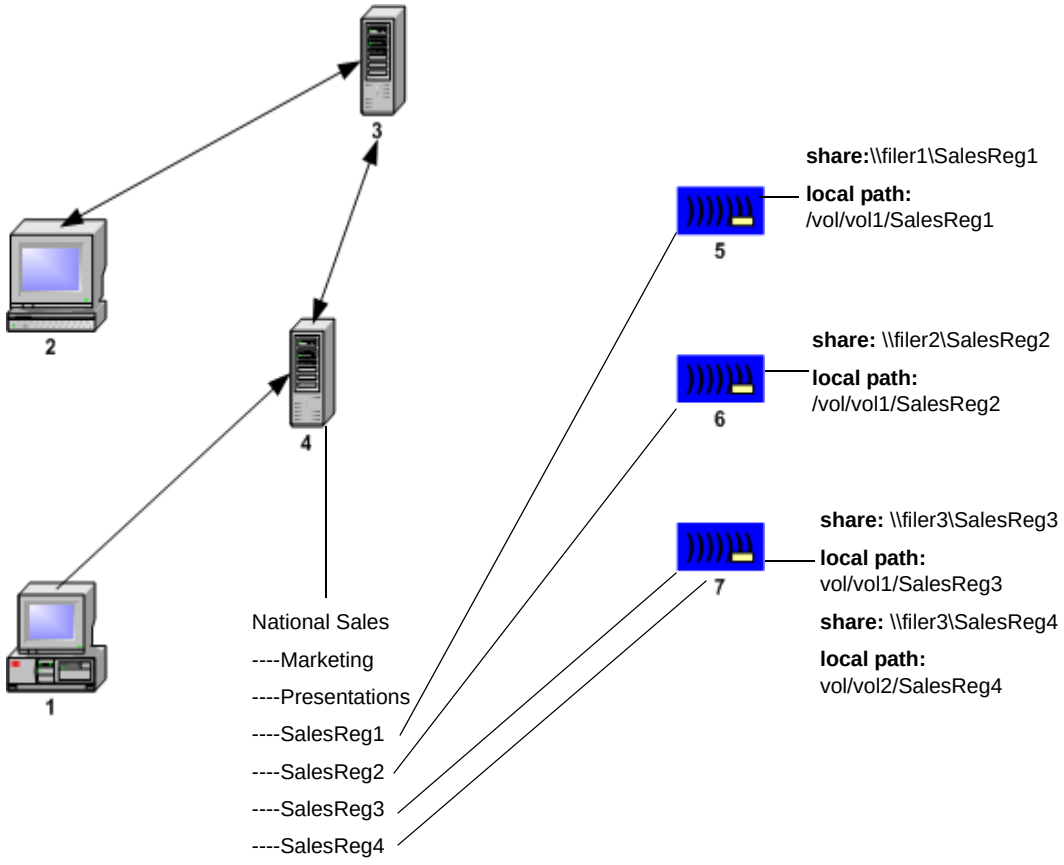

FIGURE 2 **DFS links targeting CIFS shares on filers**

L

1 CIFS user
2 VFM client
3 VFM server
4 DFS root
5 filer1
6 filer2
7 filer3

Configure the policy to create the UNIX namespace

The process is illustrated in the diagram listed below.

1. You designate a CIFS UNC path name on *filer3* as the location where the policy when it runs creates the automount and export files. You create a sync point named *MySyncPoint*, where *My logical folder* is the DFS namespace resource you designate as the master and *MyUNIXNamespace* is the name you give to the UNIX namespace resource that will be synchronized from *National Sales*.
2. The VFM server applies the exports files to *filer1*, *filer2*, and f*iler3*.
3. The VFM server notifies the NIS client you specified in the policy to copy from *filer3* the automount file named *MyUNIXNamespace*.
4. The NIS client copies *MyUNIXNamespace*.
5. The NIS client automatically pushes *MyUNIXNamespace* to the NIS database.

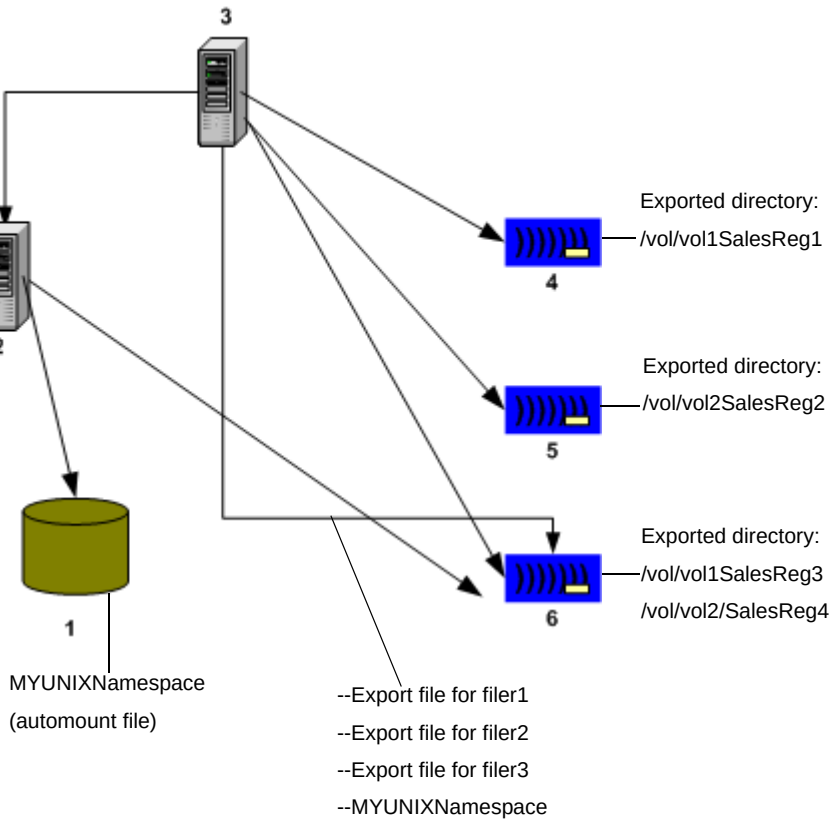

FIGURE 3 Create the UNIX namespace

1 NIS
2 NIS client
3 VFM server
4 filer1
5 filer2
6 filer3

NFS users access the same directories as CIFS users

Your NFS users can now access the same directories that your CIFS users access through the DFS namespace.

1. The NFS client obtains automount maps from NIS.
2. The NFS client automounts the NFS exports from the filers to access the same directories as CIFS users access through the DFS namespace.

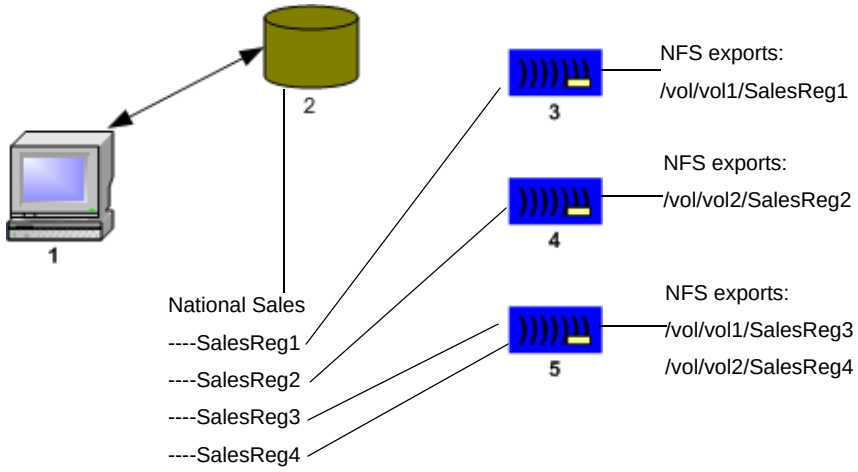

FIGURE 4 **UNIX namespace**

1 NFS user
2 NIS
3 filer1
4 filer2
5 filer3

When any of the *SalesReg* links in *My logical folder* change, the policy creates a new automount file named *MyUNIXNamespace* that VFM pushed to the NIS database.

NOTE
Automaster files are not supported. VFM does not support the Solaris-specific automount options, such as running executables and evaluating environment variables when determining the actual export to mount.

Home folder planning considerations

Home folder policies enable you to add home directories to your namespace. You can group users in home folders to accommodate users from different domains and to place home directories at different points in the namespace. When users access their home directories from the namespace, load balancing and data migration can be performed without affecting users.

Before you create a home folder policy, review DFS size limitations in light of the number of users to be added to the root, as well as the number of links that are already present in the root.

Decide how you want to arrange user home folders in the namespace. The following list of examples illustrate how you can arrange user home folders:

- Users grouped in a folder

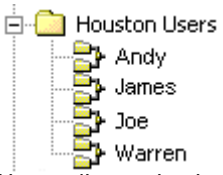

- Home directories in subfolders. This arrangement enables you to add links to the user home folder.

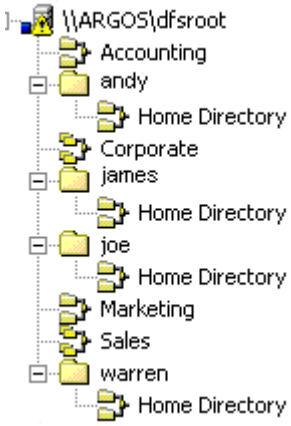

- Users grouped in a folder and home directories in subfolders. This arrangement enables you to add links to the user home folder.

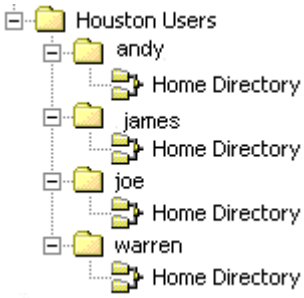

Before you move a user's share, review "Data migration considerations" on page 14. When you delete a home folder policy, all the user links are deleted from the namespace. The application changes the users' account information so that their home folders reference the shares or directories that were targeted by the links in the home folder policy before you deleted the policy.

Chapter

Installing VFM

3

This section explains VFM installation options, system requirements, and permission requirements. It provides information about how to install and upgrade VFM and configure remote shell (RSH) to communicate with computers running Linux and Solaris. The section also discusses how you can make VFM highly available.

In this chapter

- Installation options 33
- System requirements 34
- Permission requirements 37
- Filer requirements 39
- Installing VFM 39
- Opening the VFM client for the first time 41
- Updating a VFM license key 41
- RSH on Solaris and Linux computers 42
- Upgrading VFM 42
- Uninstalling VFM components 45

Installation options

You can select one of the following component combinations when you install VFM:

Client and Server	Install both the VFM server and the VFM client on a single computer.
Client	Install only the VFM client. You can install the VFM server on a remote computer or highly available computer and then access the VFM server from a different computer, such as your desktop computer.
Server	Install only the VFM server. You can install the VFM server on a remote computer or highly available computer and manage VFM from a different computer, such as your desktop computer.
Monitoring Agent	Install only the monitoring agent. You can install a monitoring agent to distribute VFM policy monitoring tasks across multiple computers.
Client and Monitoring Agent	Install the VFM client and monitoring agent components on a single system.

System requirements

The section provides information about system requirements for VFM components. Ensure your computer meets the hardware, operating system, and software requirements before installing VFM components.

VFM server and monitoring agent requirements

The following table lists VFM server and monitoring agent requirements.

TABLE 15 VFM server and monitoring agent requirements

Component	Minimum Requirements
Hardware	• 2 GHz or faster processor • 2 GB of memory • 500 MB available disk space • 75 MB additional available disk space for a typical or server-only installation with MSDE
Operating System	One of the following operating systems if you want to manage multiple Windows NT 4 or Windows 2000 roots on separate computers: • Microsoft Windows 2000 Server with SP4 • Microsoft Windows XP Professional with SP1 or SP2 One of the following operating systems if you want to manage multiple roots on a single Windows Server 2003 computer: • Microsoft Windows Server 2003 with SP1 • Microsoft Windows Server 2003 R2 Standard Edition • Microsoft Windows Server 2003 R2 Enterprise Edition • Microsoft Windows XP Professional with SP1 or SP2 If you install the VFM server on a computer running Windows XP Professional with SP1 or SP2, the number of connections used by VFM components and the connections used by other processes on the system could exceed the limits on inbound connections imposed by Windows XP Professional. If the inbound connection limits are exceeded, VFM components will experience communication problems. Use the Computer Management MMC snapin (*compmgmt.msc*) to review the number of connections in use. If you exceed the inbound connection limits, install the VFM server on a Windows 2000 Server or Windows Server 2003 computer. For more information about Windows XP Professional connection limits, refer to Microsoft Knowledge Base Article 314882, available on the Microsoft web site at www.microsoft.com.

TABLE 15 VFM server and monitoring agent requirements

Component	Minimum Requirements
Additional Software	• Microsoft .NET Framework 2.0 • Internet Explorer 5.5 or 6.0 • Windows Script Host 5.6 One of the following Microsoft SQL Server versions: • Microsoft SQL Server 2000 Desktop Engine (MSDE) with SP3 or SP4 • Microsoft SQL Server 2000 Standard Edition with SP3 or SP4 • Microsoft SQL Server 2000 Enterprise Edition with SP3 or SP4 • Microsoft SQL Server 2005 Enterprise Edition
Additional Configuration	• The VFM server must be installed on local fixed NTFS disks. • VFM uses the Microsoft SQL Server osql Utility to generate reports. If you do not install the VFM server on a computer where Microsoft SQL Server is installed, you must install the osql Utility on the VFM server computer to generate VFM reports. • Before you upgrade your VFM installation, if you are using a non-default instance of MSDE or SQL Server Express, you must complete some steps to prepare your installation for upgrade. Otherwise, the upgrade process may fail with SQL connection errors. For more information, refer to "Upgrading VFM" on page 42.

VFM client requirements

The following table lists VFM client requirements.

TABLE 16 VFM client requirements

Component	Minimum Requirements
Hardware	• 2 GHz or faster processor • 2 GB of memory • 500 MB available disk space
Operating System	One of the following operating systems if you want to manage multiple Windows NT 4 or Windows 2000 roots on separate computers: • Microsoft Windows 2000 Server with SP4 • Microsoft Windows XP Professional with SP1 or SP2 One of the following operating systems if you want to manage multiple roots on a single Windows Server 2003 computer: • Microsoft Windows Server 2003 with SP1 • Microsoft Windows Server 2003 R2 Standard Edition • Microsoft Windows Server 2003 R2 Enterprise Edition • Microsoft Windows XP Professional with SP1 or SP2

TABLE 16 VFM client requirements

Component	Minimum Requirements
Additional Software	• Microsoft .NET Framework 2.0 • Internet Explorer 5.5 or 6.0 • Windows Script Host 5.6
Additional Configuration	• The VFM client must be installed on local fixed NTFS disks. • Minimum video display of 1,024 by 768 and 256 colors. For the best display, use more than 256 colors. • Use the Remote Desktop Protocol (RDP) to connect to remote computers. Other products that provide remote desktop access, such as Virtual Network Computing (VNC), may exhibit screen rendering problems when connected with remote VFM components. • The user account you use to run the VFM client must have permissions to write to the following directories: *`\Documents and Settings\All Users\Application Data\Logs\`* *`\Documents and Settings\All Users\Application Data\NetApp\`* If you do not have permissions to write to these directories, VFM displays errors when you run VFM reports or set VFM system options.

Windows replication agent requirements

You can deploy a Windows replication agent to a computer running one of the following Windows operating systems:

- Microsoft Windows 2000 Server with SP4
- Microsoft Windows Server 2003 with SP1
- Microsoft Windows Server 2003 R2 Standard Edition
- Microsoft Windows Server 2003 R2 Enterprise Edition
- Microsoft Windows XP Professional with SP1 or SP2

NOTE
If a computer to which you want to deploy a Windows replication agent is running Microsoft VSS, you must install Microsoft .NET 2.0 on that computer before you deploy the replication agent.

The VFM server service account used to deploy a Windows replication agent must have Administrator permissions on the machine to which you deploy the replication agent and must be able to perform the following actions on the replication agent computer:

- Create, read, and write registry keys
- Write to the *Admin$* share
- Log on as a service to the replication agent computer

If the computer to which you want to deploy a replication agent is running Windows XP Professional with SP2, ensure the firewall on that computer is not blocking TCP port 6002. The replication agent uses TCP port 6002 to communicate with the VFM Server service.

UNIX replication agent requirements

You can deploy a Unix replication agent to a computer running one of the following UNIX operating systems:

- Red Hat Enterprise Linux 4
- Solaris 10

The VFM server service account must have root permissions on the computer where you deploy the replication agent.

If a UNIX replication agent is not installed on the source or destination computer, you must install Network File System (NFS) 3 or later on the source or destination computer.

If the UNIX replication agent is installed on the source or destination computer, you do not need to install NFS on the source or destination computer where the replication agent is installed. The UNIX replication agent can directly read the files on the source or destination computer.

The VFM server service account must have read and write access on the */opt* folder on the root of the UNIX computer. If the VFM service account tries to install the UNIX replication agent on a read-only file system on the UNIX computer, the following error messages are generated.

Sun Solaris 10	`pkgadd: ERROR: Unable to make directory </opt/URA>`
Red Hat Enterprise Linux 4	`error: unpacking of archive failed on file /opt/ura/bin:` `cpio: mkdir failed - No such file or directory`

You must install the packages shown in the following table on the replication agent computer. The version of the packages can be later than the packages listed in the table.

TABLE 17 Required operating system packages

Operating System	Packages
Sun Solaris 10	• SUNWcsl • SUNWlibms • SUNWcsu • SUNWloc • SUNWgzip
Red Hat Enterprise Linux 4	• glibc-2.3.4-2 • glibc-common-2.3.4-2 • bash-3.0-19.2 • coreutils-5.2.1-31 • gzip-1.3.3-9 • grep-2.5.1-16 • procps-3.2.3-7 • Red Hat package manager version 4 or greater.

Permission requirements

The following list identifies permission requirements for VFM accounts:

Installation account

The user account used to install VFM must be a member of the local Administrators group on the computer where you install VFM.

VFM server and management agent service accounts

The VFM server service account and the management agent service account must have the following permissions:

- **SeBackupPrivilege** (Back up files and directories) - required to open source directories and files with backup semantics.
- **SeRestorePrivilege** (Restore files and directories) - required to open destination directories and files with backup semantics.
- **SeSecurityPrivilege** (Manage auditing and security log) - required to access system access control lists.
- **SeTakeOwnershipPrivilege** (Take ownership of files or other objects) - required to delete directories and files.
- *If you are managing standalone DFS roots*, the service account must be a member of the local Administrators group on the computer hosting the root, and the service account must have permissions to write to the root share on the computer.
- *If you are managing domain-based DFS roots*, Microsoft Active Directory must be installed and functioning properly. The service account must have Administrator permissions on both the domain controller computer and on server computer that hosts the DFS root. For more information, refer to the Microsoft Distributed File System: Frequently Asked Questions document, available on the Microsoft web site at www.microsoft.com.
- Log on as a service permission
- Membership in the local Administrators group on the computer hosting the Microsoft SQL Server database used by the VFM server
- Permissions to write to the following directories:

  `\Documents and Settings\All Users\Application Data\Logs\`

  `\Documents and Settings\All Users\Application Data\NetApp\`

Replication agent service account

The replication agent service account requires the following permissions on the replication agent computer:

- Membership in the local Administrators group on the computer running the replication agent
- Log on as a service permission

The account used by the VFM server, a replication agent, or a VFM monitoring agent must also have the following permissions:

- **SeBackupPrivilege** (Back up files and directories) - required to open source directories and files with backup semantics.
- **SeRestorePrivilege** (Restore files and directories) - required to open destination directories and files with backup semantics.
- **SeSecurityPrivilege** (Manage auditing and security log) - required to access system access control lists.
- **SeTakeOwnershipPrivilege** (Take ownership of files or other objects) - required to delete directories and files.

NOTE
Some Network Attached Storage (NAS) systems may require the replication agent service account to be a member of the local Administrators and Backup Operators groups on the computer.

Filer requirements

If you plan to use VFM to manage NetApp filers, ensure your filer is configured appropriately.

- The VFM service account requires membership in the local Administrators group to access the *\\filer\C$* share, and to read and write files in the */etc* directory of the filer.
- You must provide the remote shell (RSH) user name and password for filers known to VFM. VFM uses RSH to communicate with filers when managing SnapMirrors and qtrees. You can configure this information through the computer properties dialog box for the filer.
- Source filers running the Data ONTAP 7.2 or later operating system must have the **snapmirror.access** option set to **legacy**.

Installing VFM

Before you install VFM, ensure the computer on which you plan to install VFM meets the requirements. For more information, refer to "System requirements" on page 34.

To install VFM

1. Log on to the computer where you want to install the VFM using a user account that is a member of the local Administrators group on the computer.
2. Install the Microsoft .NET Framework 2.0, available at www.microsoft.com.
3. Browse to the location where you saved the installation package.
4. Run the *Autorun.exe* file.
5. Click **Install VFM**.
6. On the Welcome dialog box, click **Next**.
7. Review the terms of the license agreement. If you agree to the terms, click **I accept the license agreement**, and then click **Next**.
8. If you want to review the Release Notes, click **View Release Notes**. When you are done, click **Next**.
9. Select which components you want to install, and then click **Next**. In most cases, install both the VFM server and client.

   NOTE
   If you want to change the combination of product components installed on the computer, you must first uninstall the existing components on that computer before reinstalling new components.

10. Specify where you want the setup program to install files, and then click **Next**.

11. Specify a location for application data storage, and then click **Next**.

    ATTENTION
    To avoid problems with future upgrades, specify an application data storage location that is in a different directory than where you are installing the VFM server. Specify an empty folder in the local NTFS file system where the VFM service account has write permissions. If you are installing the monitoring agent, specify a different database than the one used by the VFM server.

12. In the **Serial Number** field, type the serial number.
13. In the **License Type** field, select your license type.
14. In the **Activation Key** field, type your activation keys, and then click **Next**.
15. *If you are installing the monitoring agent, or you are installing the VFM client on a different computer than the VFM server computer,* type the name of the remote VFM server computer in the **Server** field, and then click **Next**.

    NOTE
    Make sure you accurately enter the name of the remote VFM server computer. If the specified name is incorrect, the monitoring agent or VFM client cannot communicate with the server.

16. *If you want to create a desktop shortcut*, select the **Create a desktop shortcut to this product** check box.
17. Click **Next**.
18. Specify the logon credentials for the VFM service, and then click **Next**.
19. *If you want to use an existing Microsoft SQL Server instance*, complete the following actions:
    a. Click **Use an existing Microsoft SQL Server instance**, and then click **Next**.
    b. Specify the name of the Microsoft SQL Server database you want to use.
    c. *If you want to connect to the database using Windows authentication*, click **Windows authentication**.
    d. *If you want to connect to the database using Microsoft SQL Server authentication*, click **Microsoft SQL Server authentication**, and then specify a login ID and password.
    e. Click **Next**.
20. *If you want to use MSDE for your database*, click **Install and use Microsoft SQL Server Desktop Engine (MSDE)**, and then click **Next**.

    NOTE
    You can use MSDE for product evaluations and normal product use. If you anticipate heavy usage of reporting features in VFM, consider using Microsoft SQL Server.

21. On the Review Installation Settings dialog box, click **Next**.
22. Click **Finish**, and then click **Exit** to close the autorun program.

Opening the VFM client for the first time

The first time you open the VFM client, VFM prompts you to select the DFS root, host server, or domain. Once you specify this information or start working with a DFS root, you do not need to specify it again. You can click **Cancel** to access the VFM client without selecting a DFS root to manage.

To open the VFM client

1. Log on to a computer where the VFM client is installed.
2. In the **Network Appliance** program folder, select **VFM > VFM Client**.
3. *If you have not previously opened the client*, select the DFS root, host server, or domain and then click **OK**.

Updating a VFM license key

You may need to update your license key to extend an evaluation period, when you purchase the product following a product evaluation, or when you purchase additional licenses.

NOTE
The evaluation license of VFM places a 60-minute limit on the amount of data you can transfer. This limit is for each run of a policy.

To update your license key

1. Contact Network Appliance to request a new license key. For more information about contacting Network Appliance, click **Support** on the **Help** menu.
2. On the **Help** menu, click **About VFM**.
3. Click **Edit Serial Numbers**.
4. In the **License Type** field, select whether you have an evaluation license or you have a purchased license. If you have a purchased license, select **Full**.
5. *If you selected Evaluation in the License Type field*, type your new activation key in the **Activation Key** field.
6. *If you selected Full in the License Type field*, select a product number from the **Product** list, and then type your new activation key in the **Activation Key** field for that product.
7. Click **OK**.

NOTE
After you change your license key, you need to restart the VFM client.

RSH on Solaris and Linux computers

By default VFM uses RSH to execute shell commands that configure and obtain data from NetApp filers and other non-Windows systems, including Solaris and Linux computers. On Solaris and Linux computers, ensure RSH is enabled as root. Where the environment does not support the use of RSH, you can configure VFM to use Secure Shell (SSH) to communicate with filers in place of RSH. For more information, refer to "Using RSH and SSH to Execute Shell Commands" on page 341.

Upgrading VFM

VFM 6.1 supports upgrades from VFM 5.8. and 6.0. The following are the various installation options available in VFM and the steps required to upgrade different components.

Upgrading only the server or the server and client

Before you upgrade the VFM server or the server and client, ensure the computer on which you are performing the upgrade meets the requirements for those components. Have the VFM serial number and license key information available and saved on a remote location before you begin the upgrade process.

To upgrade only the server or the server and client

1. Log on to the computer where the VFM server is installed using a user account that is a member of the local Administrators group on the computer.
2. Cancel all policies currently running or let them finish running. *If you are upgrading a server only,* cancel all policies from a client connected to the server.
3. In the policy properties for each policy, clear the **Use...schedule** check box to ensure no scheduled policies will attempt to run during the upgrade.
4. Back up the namespace. Right-click the **Logical View** and click **Backup**. Select all the roots that are used in your environment and click **Back Up**.
5. Take an application snapshot to back up the current configuration. On the **Tools** menu, click **Take Application Snapshot Now**. Save the existing configuration locally on the VFM server computer. You can use a replication policy or Windows Explorer to manually copy the application snapshot to a remote location to have a second copy of the configuration.
6. Close the VFM client.
7. Stop the VFM Server service.
8. Perform the upgrade. For more information, refer to "Installing VFM" on page 39.

Upgrading only the client

Before you upgrade the VFM client, ensure the computer on which you are performing the upgrade meets the requirements for the client component.

To upgrade only the client

1. Log on to the computer where the VFM client is installed using a user account that is a member of the local Administrators group on the computer.
2. Close the VFM client.
3. Perform the upgrade. For more information, refer to "Installing VFM" on page 39.

Upgrading only the monitoring agent

Before you upgrade the VFM monitoring agent, ensure the computer on which you are performing the upgrade meets the requirements for the monitoring agent component.

To upgrade only the monitoring agent

1. On the VFM server computer to which the monitoring agent connects, stop the VFM Server service.
2. Log on to the computer where the VFM monitoring agent is installed using a user account that is a member of the local Administrators group on the computer.
3. On the VFM monitoring agent computer, stop the VFM Server service.
4. Perform the upgrade. For more information, refer to "Installing VFM" on page 39.
5. On the VFM server computer, start the VFM Server service.

Upgrading the monitoring agent and client

Before you upgrade the VFM monitoring agent and client, ensure the computer on which you are performing the upgrade meets the requirements for those components.

To upgrade the monitoring agent and the client

1. On the VFM server computer to which the monitoring agent connects, stop the VFM Server service.
2. Log on to the computer where the VFM monitoring agent and client are installed using a user account that is a member of the local Administrators group on the computer.
3. On the VFM monitoring agent computer, stop the VFM Server service.
4. Close the VFM client.
5. Perform the upgrade. For more information, refer to "Installing VFM" on page 39.
6. On the VFM server computer, start the VFM Server service.

Upgrading replication agents

Replication agents used by VFM are upgraded automatically by the VFM server during normal status checks. You can perform the following steps to verify that replication agents have been upgraded, and to upgrade any agents that were not upgraded automatically.

To upgrade replication agents

1. Close the VFM client.
2. Restart the VFM Server service using the Windows service control manager.
3. Reopen the VFM client.
4. On the **Tools** menu, click **Agent Management**.
5. Select the replication agent in the Host list to view its properties.
6. Check the current version of the replication agent.
7. *If the replication agent was not upgraded,* right-click the agent and click **Upgrade**.
8. Click **Alerts** to view any errors that might have occurred.
9. *If the upgrade fails using the Agent Management tool,* try to uninstall the replication agent and redeploy it.
10. *If you are unable to uninstall a CIFS replication agent using the Agent Management tool,* complete the following steps:
    a. Log on to the computer running the replication agent.
    b. Using the command prompt, browse to the following location:

       Windows\Brocade\Brocade StorageX\ReplicationAgent

    c. Run the following command:

       ```
       net stop storagexreplicationagent
       ```

    d. Run the following command:

       ```
       replicationagent.exe –uninstall
       ```

    e. Delete the contents of this directory and exit out of the command prompt.
    f. Using the service control manager, verify that the StorageXReplicationAgent service no longer exists.
    g. Verify that the following registry key does not exist. *If it does exist,* delete it.

       ```
       HKEY_LOCAL_MACHINE\SOFTWARE\Brocade\Brocade StorageX\ReplicationAgent
       ```

    h. Restart the computer and then deploy the new replication agent using the Agent Management tool.
11. *If you are unable to uninstall an NFS replication agent using the Agent Management tool,* complete the following steps:
    a. Uninstall the older version of the replication agent. For more information, refer to "Installing and managing UNIX replication agents" on page 310. Make sure you delete the application directory.
    b. Open the VFM client.
    c. On the **Tools** menu, click **Agent Management**.
    d. Click **Deploy**.
    e. Type the name of the computer that will host the replication agent and specify the account that will run the replication agent.

f. Click **OK**.

g. When the replication agent appears in the list, click the agent to verify the version.

Uninstalling VFM components

Uninstall VFM components when you no longer want to use the product.

To uninstall VFM components

1. Log on to a computer where the VFM component is installed.
2. Open Control Panel.
3. Open Add or Remove Programs.
4. Select **VFM**.
5. Click **Change/Remove**.
6. Click **Yes**.
7. Click **OK**.

Chapter

Permissions and the VFM client console

4

This section discusses the permissions required to run the VFM server and explains how to configure access permissions to the server. This section also discusses permissions required to create, delete, or modify roots and links in a namespace and explains how to configure function-level access permission to the VFM client. In addition, this section provides an overview of the VFM client console and a tutorial to get you started creating your first namespace and storage policies.

In this chapter

- Permissions on accounts used with VFM components 47
- Configuring function-level access permissions 49
- Working with the VFM client console 50

Permissions on accounts used with VFM components

This section discusses the permissions that are necessary to manage the VFM server, client, and root and links in a namespace.

Permission requirements for DFS roots

Permissions settings, such as lack of administrator permission on the hosts supporting roots or the domains containing domain-based roots, can cause problems when you attempt to administer these roots.

Make sure that you have administrator privileges for any servers hosting stand-alone roots and domains and servers hosting domain-based roots before you attempt to administer them using either VFM client or native tools. If you do not have the appropriate permissions, you will not be able to create, modify, or delete roots or links in a namespace.

- **Stand-alone DFS roots** - The account must be a member of local administrators on the system hosting the root and must have permissions to write to the root share on that system.
- **Domain-based DFS roots** - For more information, refer to the Microsoft DFS FAQ available at http://www.microsoft.com/windowsserver2003/community/centers/fileservices/dfsfaq.mspx?pf=true.

For more information about account permissions for installing and running VFM, refer to "Permission requirements" on page 37.

Permission requirements for the VFM client

The account running the VFM client must have permissions to write to the

```
All Users\Application Data\NetApp\VFM\GuiConsole
```

directory tree. Users whose accounts lack write permissions to these directories will experience errors when attempting to run VFM reports or set VFM system options.

The VFM server treats the account used to run the server as having special privileges. When the VFM client runs with the same account as the VFM server, normal VFM access checks are bypassed. The account is always granted Login and Config Access Controls permissions by the VFM server. This is useful in cases where normal VFM administrator accounts have been denied login and you need to regain access for them.

When you view or configure VFM security from VFM clients running on different machines, there might be differences in how security settings appear. The following example illustrates what you might see.

In the following explanation, SERVERMACHINE is used to refer to the machine hosting the VFM server.

- The VFM client on `SERVERMACHINE` is used to restrict VFM security access so that only `SERVERMACHINE\Administrators` can perform management operations in VFM.
- When the VFM client hosted on machine `SECONDCLIENTMACHINE` connects to the VFM server on `SERVERMACHINE` and views application security settings, the display shows `SECONDCLIENTMACHINE\Administrators` and not `SERVERMACHINE\Administrators`.
- This is only a display issue. This happens because the SID of the local administrators group is well known as part of the Windows infrastructure. The VFM server running on `SERVERMACHINE` will prevent operations by accounts other than those in the `SERVERMACHINE\Administrators group`.

Lack of appropriate permissions can prevent policies from functioning correctly.

Configuring VFM server access permissions

By default, the server is configured to allow access to everyone, meaning that anyone who has access to the VFM client can connect to and manage the VFM server. You should restrict access to the server to only those administrators tasked with VFM administration.

To configure VFM server access permissions:

1. On the **Tools** menu, click **Options**.
2. Click **System Options**.
3. Click **Set Security**.
4. Add users and groups
5. Allow or deny login permission.
6. Click **OK**.
7. Click **OK**.

Configuring domain root and shared folder permissions

Configure system access permissions for domain-based roots and shared folders referenced by links in the namespace.

To configure domain-based roots and shared folders access permissions

1. Right-click a root in the Logical View or a shared folder in the Physical View and click **Properties**.
2. Click the **Security** tab.
3. Add users and groups.
4. Allow or deny permissions.
5. Click **OK**.
6. Click **OK**.

Configuring function-level access permissions

Access restrictions can be configured at the operation or function level. For example, you can allow or deny a user or group the ability to change storage policies, such as creating or deleting a policy. You can configure application-wide, function-level access permissions, and configure permissions for a DFS root in the Logical View, application snapshots, and policies in the Admin View.

To configure function-level access permissions

Select one of the following actions:

- To configure application-wide permissions, click **Options** on the Tools menu, click **System Options**, then click **Set Security**.
- To configure permissions for a DFS root, right-click the root and click **Properties**, then click the **Application Security** tab.
- To configure permissions for all policies, right-click **Admin View**, and click **Application Security**.
- To configure permissions for a policy category or type, right-click the item and click **Application Security**.

List of function-level permissions

The following table lists the function-level access permissions you can allow or deny users and groups.

TABLE 18 Function-level access permissions

Permission	Description
Full Control	Allow or deny superuser access permissions.
Login	Allow or deny login to the VFM server.
Config Access Control	Allow or deny the ability to configure access permissions.
Modify Logical View	Allow or deny the ability to manipulate items, such as links, roots, and virtual folders in the tree.
Modify Admin View	Allow or deny the ability to manipulate storage policies, such as create or delete a policy, change a policy, or make application-wide changes to all policies.
Execute Admin Policies	Allow or deny storage policy operations, such as running a policy, manually failing over links, or canceling a job in progress.

TABLE 18 Function-level access permissions

Permission	Description
Modify Physical View	Allow or deny changes to items in the Physical View, such as filers, servers, shares, and folders in the tree.
Manipulate Reports	Allow or deny the ability to manipulate reports, such as create, delete, refresh, and rename a report.

Working with the VFM client console

This section provides a general overview of the VFM client console and how to quickly customize and use the console. The left pane of the desktop provides several main items:

Admin View

Allows you to create and manage storage policies.

Logical View

Allows you to review and manage the logical layout of your DFS resources.

Physical View

Allows you to review and manage the physical layout of your DFS resources.

Reports

Allows you to view, record, and publish information about the namespace, policies, supporting machines, shares, and exports.

File Lifecycle Manager (FLM)

Allows you to integrate with FLM to manage files from their creation through to their eventual disposal.

When you select an item in the left pane, VFM displays attributes of that item in the right pane.

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

To find an item in the DFS tree

1. Click anywhere on the tree.
2. 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. 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. Double-click a shared folder.

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 then click **Locate**.
   - Click the **References** tab, right-click a reference, and then 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. For more information, refer to 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**.

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

Write audit events to the Windows event log

You can write VFM audit events to the Windows 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**.

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

Set default policy options

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

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

NOTE
To update options for a single policy, right-click the policy in the **Admin View**, and then 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 that information and select another item in the tree. The attributes of the second selected item appear in the other window.

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.

Integrate Tapestry FLM

Integrate FLM with VFM to automate the lifecycle of files from their creation to their eventual disposal.

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

Launch external tools

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

Perform one of the following actions:

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

Getting started

This section provides a tutorial to get you started using VFM to create your first namespace and storage policies. The following checklist outlines the tasks involved in this process and provides cross references to sections with additional information about each task.

TABLE 19 Getting Started Checklist

	Task		For more information, refer to...
❑	1	Create a logical namespace.	"Creating a logical namespace" on page 55
❑	2	Populate the namespace.	"Populating the namespace" on page 57
❑	3	Make the namespace highly available by making DFS roots highly available and by creating a widelink namespace on a filer.	"Making the namespace highly available" on page 59
❑	4	Make data highly available by creating client, filer, and server recovery policies.	"Making the data highly available" on page 61
❑	5	Optimize storage by creating archival migration and storage load balancing policies.	"Optimizing storage" on page 64
❑	6	Transfer data by creating replication and migration policies.	"Transferring data" on page 69
❑	7	Create a multiprotocol namespace to allow NFS and CIFS clients to access data on filers through a logical namespace.	"Creating a multiprotocol namespace" on page 74
❑	8	Automatically back up the logical structure of your namespace by creating a backup policy.	"Backing up the logical namespace" on page 76
❑	9	Back up data on filers by creating snapshot access and snapshot scheduling policies.	"Backing up data" on page 77
❑	10	Create and publish a report.	"Creating and publishing a report" on page 79

Creating 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.

Before creating your first namespace, refer to "Namespace design considerations" on page 12 for more information.

Create a namespace

The following steps outline the process for creating your first namespace. The first step is to create a DFS root to host the namespace. In this sample, you are creating a stand-alone DFS root on a Windows 2000 or Windows Server 2003. Next, populate the namespace by adding links to the DFS root. Then, review and verify the results to check how the namespace you are creating appears to your users.

To create your first namespace

1. Create a DFS root by completing the following steps:
   a. On the toolbar, click **Create a DFS Root**. The New DFS Root Wizard appears.
   b. Click **Next**.
   c. Select **Create a stand-alone DFS Root**, and then click **Next**.
   d. 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 backslashes (\\).
   e. Click **Next**.
   f. Select **Use an existing share**, and then select a share on the server.
   g. Click **Next**.
   h. Type descriptive text in the **Comment** field, and then click **Next**.
   i. Click **Finish**. The root appears under the **Logical View** tree.
2. Add links to the DFS root by dragging one or more network shares (UNC path) from Windows Explorer to the DFS root you created. Single target links appear under the root in the **Logical View**.
3. Review and verify the results by completing the following steps:
   a. Select the root in the **Logical View**.
   b. In the right pane, click the **Explorer** tab and view the results.

In Windows Explorer, the links you created appear as folders under the root.

Populating 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.

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.

You can quickly populate the namespace using a namespace creation policy or a permission-based namespace policy.

Using 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 then click **Next**.
4. Select a DFS root, and then click **Next**.
5. Click **New Folder**, type a meaningful name for the folder in the **Name** field, and then click **Next**.
6. Select the domain or servers where you want to search for shares, and then click **Next**.
7. Type the search string. You can use the `*` or `?` DOS wildcards. You can combine search strings by separating them with semi-colons. For example, entering `mkt*;sales*` finds all shares on the selected domains or servers whose names begin with either `mkt` or `sales`.
8. Select **Search for shares matching the search string**, and then click **Next**. A list of shares that match your search pattern appears.
9. Clear the items that you do not want to add to the namespace, and then click **Next**.
10. Select a template to arrange how links appear in the namespace, and then click **Next**. You can preview the namespace.
11. Verify the changes to the namespace, and then click **Next**.
12. Click **New**, configure a schedule to run the policy, and then click **Next**.
13. Click **Finish**. The new policy appears under the Namespace Creation Policies folder.

Using 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 then click **Next**.
4. Select a DFS root, and then click **Next**.

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

Making 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 is root synchronization?

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 is a namespace availability policy?

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.

What are the 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.

Make the namespace highly available

To make the namespace highly available, you need to determine whether you will synchronize the logical structure of two DFS roots, or synchronize the logical structure of a DFS root with a widelink namespace. To synchronize two DFS roots, you need to first make the DFS roots highly available. To synchronize a DFS root with a widelink namespace, you need to first create a namespace on a filer. Then, you need to schedule the synchronization to make sure that changes made to the logical structure of the master root are synchronized with the secondary root or widelink namespace.

To make the namespace highly available

1. If you are synchronizing two DFS roots, create a namespace availability policy and add the two DFS roots to the policy by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Availability Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field.
   c. Click **Add** and type the path of a root, or click **Browse** to locate a root in your network.
   d. Click **Add** and type the path of the other root, or click **Browse** to locate the other root in your network.
   e. 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.
   f. Click **OK**. A new policy appears under the namespace **Availability Policies** folder.
2. If you are synchronizing a DFS root with a widelink namespace, create a widelink namespace contained in a CIFS share that is hosted on a NetApp filer by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Availability Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field.
   c. Click **Add** and type the path of a root, or click **Browse** to locate a root in your network.
   d. 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.
   e. 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.
   f. Click **OK**. A new policy appears under the namespace **Availability Policies** folder. The master root and the widelink namespace appear under the policy.
3. Schedule how frequently to synchronize the logical structure by completing the following steps:
   a. Right-click the namespace availability policy, and then click **Properties**.
   b. Click **Schedule**.
   c. Select the schedule options you want.
   d. Click **OK**.
   e. Click **Done**.

Making the data highly available

VFM provides the following types of recovery policies. Select which type you want to create.

Client recovery policies

Enable you to set up 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 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 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 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.

Before you create your first recovery policy, refer to "Disaster recovery planning considerations" on page 20 for more information.

Creating a client recovery policy

After you create a client recovery policy, you need to add a link to the policy. Next, you need to configure status monitoring for the link so 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 recovery policy. Then, make sure that data changes made to the online target are synchronized with the offline targets. Schedule how frequently you want to replicate the data.

1. Create a client recovery policy by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Business Continuity Policy**, and then click **Client Recovery Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field.
   c. Select the VFM replication method.
   d. Click **OK**. The new policy appears under the **Client Recovery Policies** folder.
2. Add a link to the client recovery policy by completing the following steps:
   a. Double-click the **Business Continuity Policies** folder in the **Admin View**, and then double-click the **Client Recovery Policies** folder to display the policy you created.
   b. 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.
3. Configure monitoring of the client recovery policy by completing the following steps:

a. Right-click the client recovery policy, and then click **Properties**.

b. Click **Monitoring Schedule**.

c. Schedule status monitoring to begin at a specific time and date, and to repeat at a stated frequency.

d. Click **OK**.

4. Configure replication of the client recovery policy by completing the following steps:

   a. Right-click the client recovery policy, and then click **Properties**.

   b. Click **Replication Schedule**.

   c. Select **Use Schedule**.

   d. Select the schedule option you want, and then click **OK**.

Creating a filer recovery policy

The VFM Filer Recovery Policy Creation wizard guides you through the process of creating a filer recovery policy. Create a filer recovery policy to make data on NetApp filers highly available.

1. On the **Tools** menu, click **New Policy**, point to **Business Continuity 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 then click **Next**.
10. If a destination does not exist on the destination filer for the source qtree, click the row, and then 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**.

18. Click **Next**.
19. Click **Finish**. The filer recovery policy appears under the **Business Continuity Policy\Filer Recovery Policies** folder.

Creating a server recovery policy

First, create the server recovery policy to make the data on the server highly available. Next, select the namespaces you want to monitor for links targeting the primary server. After you select the namespaces, monitor the namespaces for new links targeting the primary server.

After you configure monitoring for new resources on the selected namespaces, you need 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.

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. Schedule how frequently you want to replicate data.

Before you create your first recovery policy, refer to "Disaster recovery planning considerations" on page 20 for more information.

To create your first server recovery policy

1. Create a server recovery policy to make data on a server highly available by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Business Continuity Policy**, and click **Server Recovery Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field.
   c. Type the name of the primary server.
   d. Type the name of the secondary server, and then click **OK**. The new policy appears under the **Server Recovery Policies** folder.
2. Select the namespaces to monitor by completing the following steps:
   a. Right-click the server recovery policy, and then click **Properties**.
   b. Click **Resource Monitoring Options**.
   c. Click **Add All Roots From the Logical View**.
   d. Click **OK**.
3. Configure monitoring for new resources by completing the following steps:
   a. Right-click the server recovery policy, and then click **Properties**.
   b. Click **Resource Monitoring Schedule**.
   c. Schedule resource monitoring to begin at a specific time and date, and to repeat at a stated frequency.
   d. 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.
4. Configure monitoring of a server recovery policy by completing the following steps:

a. Right-click the policy, and then click **Properties**.

b. Click **Monitoring Schedule**.

c. Schedule status monitoring to begin at a specific time and date, and to repeat at a stated frequency.

d. Click **OK**.

5. Configure replication of the server recovery policy by completing the following steps:

a. Right-click the server recovery policy, and then click **Properties**.

b. Click **Replication Schedule**.

c. Select **Use Schedule**.

d. Select the schedule options you want.

e. Click **OK**.

Optimizing 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, such as 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.

You need to first use an archival migration policy to transfer the directories. You can transfer CIFS or NFS directories. Then, use a storage load balancing policy to manage the storage usage and load levels. An archival migration policy can contain CIFS to CIFS and NFS to NFS source and destination pairs.

Using an archival migration policy to transfer CIFS directories

Create an archival migration policy to automatically migrate CIFS data from primary to secondary storage. Configure the policy with the criteria you want to use to select directories to migrate. After you specify the selection criteria, you need 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.

Review the following considerations:

- You cannot configure more than one search path with the same source.
- 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.

To create your first archival migration policy to transfer CIFS directories

1. Create an archival migration policy by completing the following steps:
   a. On the **Tools** menu, point to **New Policy**, point to **Data Lifecycle Automation Policy**, and click **Archival Migration Policy**.
   b. Type a word or phrase in the **Name** text box that best describes the policy.
   c. Click **Add** to include a shared folder with subfolders as the source, and a shared folder on another machine as the destination.
   d. Click **OK**. The policy and the source shared folder appear under the **Data Lifecycle Automation Policies\Archival Migration Policies** folder.
2. Specify the selection criteria 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 by completing the following steps:
   a. Right-click the policy, and then click **Properties**.
   b. Click **Migration Criteria**.
   c. Select **Migrate directories when**.
   d. Select **% of files have not been accessed within the last**.
   e. Change the percent to **85** for this example.
   f. Change the number of days to **30** or this example, and then click **OK**.
3. Specify the migration actions by completing the following steps:
   a. Right-click the policy you created, and then click **Properties**.
   b. Click **Initial Phase**.
   c. Make sure that **Perform a baseline copy** is selected.
   d. Click **Final Phase**.
   e. Make sure that **Update namespace** is selected.
   f. Make sure that **by retargeting links** is selected, and then click **OK**.
4. Specify how you want to update the namespace options after the data is migrated by completing the following steps
   a. Click the **Update Namespace Options** tab.
   b. Click **Add** to list the DFS roots that contain links you want to update.
   c. Click **OK**.
5. Specify the migration tasks you want the policy to execute by completing the following steps:

a. Click **Task Options**.

b. Make sure that **Automatically add migration tasks for migration candidates** is selected. Migration tasks are created after scanning is complete.

c. 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.

d. Click **OK**.

6. Schedule how frequently you want the policy to scan for migration candidates by completing the following steps:

   a. Click **Policy Scan Schedule**.

   b. Click **New**, select the schedule from the list provided, and then configure the schedule.

   c. Click **OK**.

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

7. Schedule when to run the initial and final phases to migrate directories that match the selection criteria by completing the following steps:

   a. Click **Initial Phase Schedule**.

   b. Specify the day and time you want to run the initial phase.

   c. Click **Final Phase Schedule**.

   d. Click **New**, select the schedule from the list provided, and then configure the final phase schedule.

   e. Click **OK**.

Using an archival migration policy to transfer NFS directories

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

Review the following considerations:

- 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.

To create your first archival migration policy to transfer NFS directories

1. Create an archival migration policy by completing the following steps:
   a. On the **Tools** menu, point to **New Policy**, point to **Data Lifecycle Automation Policy**, and click **Archival Migration Policy**.
   b. Type a word or phrase in the **Name** text box that best describes the policy.
   c. Click **Add** to include an export with subfolders as the source and an export on another machine as the destination.
   d. Click **OK**. The policy and the source shared folder appear under the **Data Lifecycle Automation Policies\Archival Migration Policies** folder.
2. Specify the selection criteria to automatically migrate a folder to secondary storage if the folder contains more than 200 files by completing the following steps:
   a. Right-click the policy, and then click **Properties**.
   b. Click **Migration Criteria**.
   c. Select **Migrate directories when**.
   d. Select **Directory contains more than...files**.
   e. Change the file number to **200** for this example, and then click **OK**.
3. Specify the migration actions by completing the following steps:
   a. Right-click the policy you created, and then click **Properties**.
   b. Click **Initial Phase**.
   c. Make sure that **Perform a baseline copy** is selected.
   d. Click **Final Phase**.
   e. Make sure that **Update namespace** is selected.
   f. Make sure that **by retargeting links** is selected, and then click **OK**.
4. Update the namespace options to list the automount files you want to update after the data is migrated by completing the following steps:
   a. Click the **Update Namespace Options** tab.
   b. Click **Add**.
   c. Type the UNIX fully qualified path of an automount file. You can type multiple paths to add multiple automount files.
   d. Click **OK**.
5. Specify the tasks you want the policy to execute by completing the following steps:
   a. Click **Task Options**.
   b. Make sure that **Automatically add migration tasks for migration candidates** is selected. Migration tasks are created after scanning is complete.
   c. Click **OK**.
6. Schedule how frequently you want the policy to scan for migration candidates by completing the following steps:

a. Click **Policy Scan Schedule**.

b. Click **New**, select the schedule from the list provided, and then configure the schedule.

c. Click **OK**.

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

7. Schedule when to run the initial and final phases to migrate directories that match the selection criteria by completing the following steps:

   a. Click **Initial Phase Schedule**.

   b. Specify the day and time you want to run the initial phase.

   c. Click **Final Phase Schedule**.

   d. Click **New**, select the schedule from the list provided, and then configure the final phase schedule.

   e. Click **OK**.

Creating a storage load balancing policy

Create a storage load balancing policy to automatically load balance storage locations you specify in the policy. Specify the criteria you want to use to evaluate the current size of disk space and migrate directories under the storage locations that are link targets. Next, add the storage locations you want to load balance. At least one of the shares selected must be a link target or must contain a directory that is a link target. Then, 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. The last step is to schedule when to migrate directories referenced by links in your namespace to load balance the storage pool.

To create your first storage load balancing policy:

1. Create a storage load balancing policy by completing the following steps.

   a. On the **Tools** menu, point to **New Policy**, point to **Data Movement Policy**, and click **Storage Load Balancing Policy**.

   b. Type a word or phrase in the **Name** field that best describes the policy.

   c. Click **Add**, select a DFS root in your namespace that contains links targeting directories you want to load balance, and then click **OK**.

   d. Click **OK**.

2. Add storage locations by completing the following steps.

   a. Click **Pool**. A new dialog box appears on the right pane.

   b. Click **Add** and select a share from the **Physical View** or your network. You can repeat this step to add multiple shares.

   c. Click **OK**.

3. Schedule scanning by completing the following steps.

   a. Click **Policy Scan Schedule**.

   b. Click **New**, select the schedule from the list provided, and configure the schedule.

c. Click **OK**.

4. Schedule migration by completing the following steps.
   a. Click **Initial Phase Schedule**.
   b. Specify the day and time you want to run the initial phase.
   c. Click **Final Phase Schedule**.
   d. Click **New**, select the schedule from the list provided, and configure the final phase schedule.
   e. Click **OK**.

Transferring 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.

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.

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.

Using replication policies to distribute CIFS data

When you create a replication policy to distribute CIFS data, review the following considerations:

- 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.

To create your first replication policy to distribute CIFS data

1. Create a replication policy to distribute CIFS data by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and click **Replication Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field. The name appears under the **Replication Policies** folder.
   c. Click **OK**. The new policy appears under the **Replication Policies** folder. The policy shows an error status because the source and targets are not yet added to the policy.
2. Add the source from which data is replicated by completing the following steps:
   a. Drag a share under a machine that appears in the **Physical View** to the replication policy. The option **As the source path** is selected.
   b. Click **OK**.
3. Add a destination that receives the data being replicated by completing the following steps:
   a. Drag a share under a machine in the **Physical View** to the replication policy.
   b. Select **Add it to the list of destination paths**.
   c. Click **OK**. The error status on the policy is cleared because you have added a 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. The destination may be a link that references a single or multiple CIFS targets or a CIFS physical path.
4. Schedule how frequently you want the replication to take place by completing the following steps:
   a. Right-click the replication policy, and then click **Properties**.
   b. Click **Replication Schedule**. By default, **Use replication schedule** is selected.
   c. Click **New**, and then configure the schedule.
   d. Click **OK**.

Using replication policies to distribute NFS data

Before you create a replication policy to distribute NFS data, refer to the NFS requirements in "Archival Migration Policies" on page 210.

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**

To create your first replication policy to distribute NFS data

1. Create a replication policy to distribute NFS data by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and click **Replication Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field. The name appears under the **Replication Policies** folder.
   c. In the **Source** field, type the physical path of an NFS export, or click **Browse** to locate an export in your network.
   d. 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.
   e. Click **OK**. The new policy appears under the **Replication Policies** folder.
2. Add the source from which data is replicated by completing the following steps:
   a. Drag a share under a machine that appears in the **Physical View** to the replication policy. The option **As the source path** is selected.
   b. Click **OK**.
3. Add a destination that receives the data being replicated by completing the following steps:
   a. Drag a share under a machine in the **Physical View** to the replication policy.
   b. Select **Add it to the list of destination paths**.
   c. Click **OK**. You can click the **Replication Topology** tab in the right pane to see a graphical view of how data is distributed by the policy.
4. Schedule how frequently you want the replication to take place by completing the following steps:
   a. Right-click the replication policy, and then click **Properties**.
   b. Click **Replication Schedule**. By default, **Use replication schedule** is selected.
   c. Click **New**, and then configure the schedule.
   d. Click **OK**.

Creating 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.

NOTE
If you have users and applications that are tied to specific shares and file locations through hard-coded UNC path names, 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.

To create your first migration policy

1. Create a migration policy to move data referenced by links in your namespace to a new server by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and click **Migration Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** text box.
   c. Click **OK**. A new policy appears under the **Migration Policies** folder.
2. Select migration actions to define the three phases of the migration operation by completing the following steps:
   a. Right-click the policy you created, and then click **Properties**.
   b. Click **Initial Phase**.
   c. Make sure that **Perform a baseline copy** is selected.
   d. Click **Final Phase**.
   e. Make sure that **Update namespace** is selected.
   f. Make sure that **by retargeting links** is selected, and then click **Apply**.
3. Update the namespace options to specify the DFS root that contains the links you want to update by completing the following steps:
   a. When migrating CIFS data, click **Add** to list the DFS roots that contain links you want to update.
   b. When migrating NFS data, click **Add** to list the automount files you want to update.
   c. Click **Apply**.
4. Schedule the initial and final migration phases by completing the following steps:
   a. Click **Initial Phase Schedule**.
   b. Specify the day and time you want to run the initial phase.
   c. Click **Final Phase Schedule**.
   d. Click **New**, select the schedule from the list provided, and then configure the final phase schedule.
   e. Click **OK**.
5. When migrating CIFS data, add tasks to the migration policy that identify source and destination paths by completing the following steps:

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

a. 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.

b. Type the name of a machine containing CIFS data in the **Destination Path** text box, and then click **Next**.

c. Select **Only migrate paths referenced by the logical namespace**, and then click **Next**. A list of source and destination folders appears.

d. Click **Next**. The policy checks for destination shares that do not exist.

e. 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 then click **Create Destination Share**. The shares are created on the destination machine.

f. 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.

g. Validate the list of links, and then click **Next**.

h. 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.

6. When migrating NFS data, add tasks to the migration policy that identify source and destination paths by completing the following steps:

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

a. Right-click the migration policy, and then click **Add task to policy**.

b. In the **Original Path** text box, type the path of an export, or click **Browse** to locate an export.

c. In the **Destination Path** text box, type the path of an export, or click **Browse** to locate an export.

d. Click **Next**. The original and destination paths appear in a list. You have the option of changing the destination path.

e. Click **Next**.

f. 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 then click **Create Destination Export**. The exports are created on the destination machine.

g. Click **Next**.

h. 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.

Creating 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.

Create a multiprotocol namespace 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. Before you create your first multiprotocol namespace policy, refer to "Multiprotocol namespace planning considerations" on page 27 for more information.

To create your first multiprotocol namespace policy where the DFS namespace is the master

1. Create a multiprotocol namespace policy to allow NFS and CIFS clients to access data on NetApp filer through a logical namespace by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Multiprotocol Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field.
   c. 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.
   d. Type a CIFS UNC path name for the directory that will contain files created by VFM to configure the UNIX namespace.
   e. Click **New** to select the data that you want UNIX clients to access.
   f. Type a word or phrase that best describes the data to be synchronized. This name appears on the **Sync Points** list.
   g. Select DFS as the master. Whenever the links in the DFS Logical Folder change, the automount file and corresponding export files are updated.

h. Type the path of a DFS folder that contains links referencing data on NetApp filers that you want UNIX clients to access.

i. Type a name for the automount file that VFM will create, and then click **OK**.

j. Click **OK**. The policy appears under the **Namespace Policies\Multiprotocol Policies** folder in the **Admin View**.

2. Schedule synchronization 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 by completing the following steps:

   a. Right-click the multiprotocol namespace policy, and then click **Properties**.

   b. Click **Schedule**.

   c. Click **Add**, and then configure the schedule.

   d. Click **OK**.

Create a multiprotocol namespace 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. Before you create your first multiprotocol namespace policy, refer to "Multiprotocol namespace planning considerations" on page 27 for more information.

To create your first multiprotocol namespace policy where the UNIX namespace is the master

1. Create a multiprotocol namespace policy to allow NFS and CIFS clients to access data on NetApp filer through a logical namespace by completing the following steps:

   a. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Multiprotocol Policy**.

   b. Type a word or phrase that best describes this policy in the **Name** field.

   c. 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.

   d. Type a CIFS UNC path name for the directory that contains the automount files describing the UNIX namespace.

   e. Click **New** to select the data that you want DFS clients to access.

   f. Type a word or phrase that best describes the data to be synchronized. This name appears on the **Sync Points** list.

   g. Select UNIX as the master. Whenever the contents of the automount file changes, the links in the DFS Logical Folder are updated.

   h. Type the name of a DFS folder that contains links referencing data on NetApp filers that UNIX clients are accessing.

   i. 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.

   j. Click **OK**.

   k. Click **OK**. The policy appears under the **Namespace Policies\Multiprotocol Policies** folder in the **Admin View**.

2. Schedule synchronization 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 by completing the following steps:
   a. Right-click the multiprotocol namespace policy, and then click **Properties**.
   b. Click **Schedule**.
   c. Click **Add**, and then configure the schedule.
   d. Click **OK**.

Backing 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 do not 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

To create your first namespace backup policy

1. Create a namespace backup policy to back up the logical structure of the DFS roots participating in the policy by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Backup Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field.
   c. Click **Add** and type the path of a root, or click **Browse** to select one of the roots that appears in your namespace. You can repeat this step to add more roots to the policy.
   d. Click **OK**. The policy appears under the **Namespace Policies\Backup Policies** folder in the **Admin View**.
2. Schedule how frequently you want to back up the roots by completing the following steps:
   a. Right-click the namespace backup policy, and then click **Properties**.
   b. Click **Schedule**.
   c. Select **Use Schedule**.
   d. Click **Fill in current time and date**.
   e. Select **Perform weekly backups**.

f. Click **OK**.

3. Browse the content of the backup file to check the results of the backup operation by completing the following steps:
   a. On the **Tools** menu, click **Browse Backups**.
   b. Double-click the policy you created. **Daily**, **Hourly**, and **Weekly** folders appear.
   c. Double-click the folder that contains the backup file.
   d. Double-click the backup file to view its contents. A tree that looks like the **Logical View** appears on the left pane.
   e. Click an item in the tree to view its properties in the right-pane.
   f. When you finish browsing, click **OK**.

Backing 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 a snapshot policy

NetApp filers must have the snapshot feature turned on.

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.

To create your first snapshot access and snapshot scheduling policies

1. 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 by completing the following steps:
   a. On the **Tools** menu, click **New Policy**, point to **Snapshot Policy**, and click **Access Policy**.
   b. Type a word or phrase that best describes this policy in the **Name** field.
   c. Click **Add All Roots From the Logical View**.

d. 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.

2. View the snapshot data to validate the data referenced by the captured namespace by completing the following steps:

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

   b. Select one of the files to show tabs on the right pane.

   c. Click the **Explorer** tab. The left pane on the **Explorer** tab shows the captured namespace.

   d. Click a link to show the data referenced by the link in 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.

3. Restore a file to its current location by completing the following steps:

   a. Right-click a captured namespace file under the **Captured Namespace** folder, and then click **Restore**. A list of the snapshot data to be restored appears.

   b. Select **Restore the Snapshot Files**.

   c. Click **OK**. The data in the active file system is replaced with the data in the snapshot.

4. Create a snapshot scheduling policy to centrally manage the snapshots on the volumes referenced in the snapshot accessing policy you created by completing the following steps:

   a. Right-click the snapshot access policy, and then click **Set Snapshot Schedule for Referenced Volumes**.

   b. Click **Next**.

   c. Select **Add this volume to a Snapshot Scheduling policy**, and then click **Next**.

   d. Click **New Policy**.

   e. Type a word or phrase that best describes the policy in the **Name** field.

   f. Click **Filer Settings**.

   g. Select **Manage filer-initiated snapshot schedule**.

   h. Type how many snapshots you want to retain.

   i. Select the hours at which you want to schedule snapshots.

   j. Optionally type the amount of space you want to reserve for snapshots on the referenced volumes.

   k. Click **OK**.

   l. Select the new policy, and then click **Next**.

   m. Click **Finish**. A new policy appears under the **Snapshot Policies\Scheduling Policies** folder.

Creating and publishing 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 guides 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.

In the following 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 if you already have a machine in the **Physical View**.

You will 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.

To create and publish your first report

1. If you do not already have a machine in the tree, add a machine to the **Physical View** tree by completing the following steps:
   a. Right-click **Physical View** in the tree, and then click **Add folder**.
   b. Type a name for the folder, and then click **OK**.
   c. Drag a machine from Windows Explorer to the folder you created.
2. Enable report publishing by completing the following steps:
   a. On the **Tools** menu, click **Options**.
   b. Click **Reporting**.
   c. Select **Enable Report Publishing**.
   d. 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.
   e. Click **OK**.
3. Create a Share Properties report by completing the following steps:
   a. Right-click the **Reports** tree, point to **Reports**, point to **Physical**, and click **Share and Export Properties**.
   b. The columns that will appear in the report are selected. Click **Next**.
   c. Select **Publish report**, and then click **Next**.
   d. Select the **Report Sources** report source group that appears under the **Monthly** report schedule group, and then click **Edit**.
   e. Type the UNC path of the machine you added to the **Physical View**, or click **Browse** to find it.
   f. Click **OK**.
   g. Click **Next**.
   h. Optionally type a different name for the report.

i. 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.

4. Gather data and publish the report by completing the following steps:

   a. Expand the **Report** tree to locate the report source group to which you assigned the report you created.

   b. Right-click the report source group, and then click **Start data collection now**. Wait a few minutes while the data is gathered and the report is updated and published.

Chapter

Creating and Managing DFS Namespaces

5

In this chapter

- Introduction 81
- Create namespaces 82
- Populating the namespace 83
- Populate the namespace automatically. 114
- Populate the namespace based on permissions 119
- Make the namespace highly available. 123
- Create multiprotocol namespaces 129
- Virtualize user home directories. 134
- Back up the namespace 138

Introduction

This guide explains how to create and manage a logical namespace. It begins with instructions about how to create DFS roots and links and how to quickly populate the namespace with namespace creation and permission-based policies. It provides instructions about how you can make the namespace highly available and how to automatically back up your namespace by creating namespace availability and namespace backup policies respectively. The section also explains how to create multiprotocol namespaces for CIFS and NFS clients and how to manage the life cycle of user home directories.

Checklist

- For more information and basic requirements about DFS and instructions about how to create namespaces, refer to "Create namespaces" on page 82.
- For more information about how to automatically populate your namespace, refer to "Populate the namespace automatically" on page 114.
- For more information about how to automatically populate your namespace based on group permissions you specify in a policy selection criteria, refer to "Populate the namespace based on permissions" on page 119.
- For more information about how to keep aligned the logical structure of two or more DFS roots or a DFS root and widelink namespaces, refer to "Make the namespace highly available" on page 123.
- For more information about how to create and manage namespaces for CIFS and NFS clients, refer to "Create multiprotocol namespaces" on page 129.
- For more information about how to manage the life cycle of user home directories, refer to "Virtualize user home directories" on page 134.

- For more information about how to back up the logical structure of a namespace, refer to "Back up the namespace" on page 138.

Create namespaces

This section explains how to create and manage a DFS namespace. The section is arranged according to the elements that make up a namespace, such as roots, links, and link targets. Status monitoring is also covered in this section.

The Logical View tree

The Logical View tree shows the roots, links, and logical folders in your namespace. You manage your namespace from this tree.

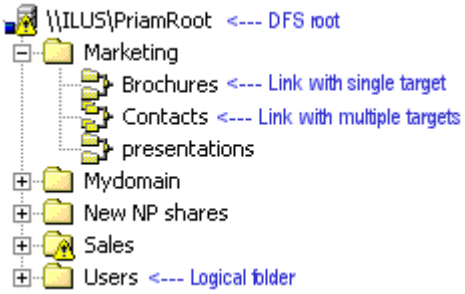

Basic DFS requirements

- You can create DFS roots on Windows 2000 and Windows Server 2003 servers. Domain-based roots replicated through Microsoft File Replication Service (FRS) must be created on partitions formatted with Windows 2000 or Windows Server 2003 NTFS on servers that are part of a Windows 2000 or Windows Server 2003 Active Directory™ Domain.
- You can add a DFS link under the DFS root to reference any root, shared folder, or subfolder of a shared folder.
- Linking from one DFS link to another DFS link is possible only if the second link is in a separate namespace. For example, you cannot have `\\Server\DFS_root\DFS_link1\DFS_link2` unless `DFS_link2` lies under another DFS root.
- You can replicate each DFS root (domain-based only) and shared folder (domain-based and stand-alone) to eliminate single points of failure.
- Administering a stand-alone DFS root requires Local Administrator privileges on the server hosting the root.
- Administering a domain-based DFS root requires Domain Administrator privileges for the domain hosting the root and Local Administrator privileges on each of the servers hosting root targets.
- A consolidation DFS root is a special stand-alone DFS root.

Recommended size limits

TABLE 20 Recommended size limits for DFS roots

Item	Maximum number
Characters per file path	DFS has no inherent limitation on the number of characters in any item of the namespace. Microsoft Win32 APIs have a MAX PATH limitation of 260 characters and applications exceeding this limit fail attempting to access UNCs or directory paths.
Replica members in a domain-based set	Refer to article Q232613 in the Microsoft Windows knowledge base.
DFS roots per server	One root for Windows 2000 servers. Multiple roots for a server running Windows Server 2003 Enterprise Edition. Other DFS-compatible systems might support additional roots per server.
DFS roots per domain	Refer to article Q232613 in the Microsoft Windows knowledge base.
DFS links or shared folders in a DFS root	Refer to article Q232613 in the Microsoft Windows knowledge base.

Additional resources

To read more about DFS, refer to the following list of recommended web sites.

- DFS documentation about Windows Server 2003
  - http://www.microsoft.com/windowsserver2003/community/centers/fileservices/dfsfaq.mspx?pf=true
- DFS documentation about Windows 2000
  - http://www.microsoft.com/WINDOWS2000/techinfo/administration/fileandprint/dfsbp.asp

Populating the namespace

There are several ways to populate the namespace. Select a method that best meets your needs.

- Add a Link dialog box enables you to add one link at a time.
- Drag UNC paths from Windows Explorer to a root in the Logical View.
- Add multiple links dialog box enables you to add multiple links and folders at a time. Use this method when you know where items are located on your network and most of the DFS shares that you want to link are located on a few machines. A wizard steps you through the process. You can drag items from Windows Explorer to a folder.
- Namespace Creation Policy wizard enables you to add multiple links and folders at a time. Use this method when you want to add links to multiple shares that are supported by different machines and when you want to keep the logical namespace aligned with subsequent changes in the underlying physical resources.
- Permission Based Namespace wizard enables you to add multiple links and folder at a time like the Namespace Creation Policy wizard. Use this method when you want to populate the namespace based on security permissions of shares.

Create and manage roots

The first step in creating a namespace is to create roots. The root is the top of the logical hierarchy. It can be domain-based or stand-alone. You can create domain-based DFS roots, stand-alone DFS roots, and a special stand-alone root named consolidation DFS root. This section explains how to create and manage these roots.

NOTE
Windows Server 2003, Enterprise Edition, can host multiple DFS roots. To manage this configuration, VFM must be installed on Windows XP Professional or on Windows Server 2003.

The following table compares the features of each root.

TABLE 21

Domain-based DFS root	Stand-alone DFS root	Consolidation DFS root
Microsoft Active Directory must be configured before you can create the root.	Does not use Microsoft Active Directory.	Does not use Microsoft Active Directory.
Must be hosted on a domain member server or domain controller.	Can be hosted on three types of Windows 2000 servers: a stand-alone server, a member server, or a domain controller.	• A special type of stand-alone root that must be hosted on a member server and not a domain controller. • Must be hosted on Windows Server 2003. It is recommended that you host the root on Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.
Can have multiple root targets.	Can only have one root target.	Can only have one root target.
Can be the target of a link.	Can be the target of a link.	• Cannot be the target of a link. • Does not support link names with backslash (\) characters, such as `abc\def`.

TABLE 21

Domain-based DFS root	Stand-alone DFS root	Consolidation DFS root
Shared folders can be hosted on any type of Microsoft server, or any other server software for which a Windows NT-based client is available. This includes any version of Windows NT Server or Windows NT Workstation, NetWare®, Windows 95, Windows 98, Windows 2000, Windows Server 2003, Windows for Workgroups, LAN Manager, or NFS. Only Windows 2000 and Windows Server 2003 clients support access to domain-based DFS roots in Microsoft Active Directory. You must update Windows NT 4.0, Windows 95, and Windows 98 clients with appropriate service packs before they can connect to domain-based roots. Windows 95 and Windows 98 clients can only negotiate referrals for Server Message Block (SMB) shares. All other targets appear as empty directories.	Shared folders can be hosted on any type of Microsoft server, or any other server software for which a Windows NT-based client is available. This includes any version of Windows NT Server or Windows NT Workstation, NetWare, Windows 95, Windows 98, Windows 2000, Windows Server 2003, Windows for Workgroups, LAN Manager, or NFS. Windows 95 and Windows 98 clients can only negotiate referrals for SMB shares. All other targets appear as empty directories.	
Supports multiple target links.	Supports multiple target links.	
Can make use of FRS replication.	FRS replication not available.	FRS replication not available.

Create a stand-alone DFS root

A Stand-alone DFS root 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.

NOTE
For more information about how to make a stand-alone DFS root highly available, refer to "Make the namespace highly available" on page 123.

Prerequisites

- Microsoft DFS server software (Distributed File System Service dfssvc.exe) must be running on the server that hosts the DFS root.
- You must have administrator permissions for the server where you create the DFS root.
- Windows 2000 servers can host only one DFS root.
- VFM creates roots on Windows 2000 and Windows Server 2003 servers.
- Windows Server 2003, Enterprise Edition, can host multiple DFS roots. To manage this configuration, VFM must be installed on Windows XP Professional or on Windows Server 2003.

To create a stand-alone DFS root

1. Perform one of the following actions:
   - Right-click **Logical View** and click **Add a new DFS root**.
   - Click **Create DFS root**.
   - On the **Tools** menu, click **Add a new DFS root**.
2. Click **Next**.
3. In the **Server name** field, type the name of the server where you want to create the DFS root, or click **Browse** to select an available server.
4. Click **Next**.
5. Select **Create a Stand-alone DFS root** and click **Next**.
6. Select an existing shared folder or create a new shared folder.
   - If you clicked **Use an existing shared folder**, type the path and name of the shared folder or select an available shared folder.
   - If you clicked **Create a new shared folder**, type the path and name of the shared folder in the appropriate fields.

   The name of the DFS root is the same as the shared folder you selected.
7. Click **Next**.
8. In the **Comment** field, optionally type a word or phrase.
9. In the **Clients cache this referral for** field, optionally change the duration (in seconds) that the root is cached on a DFS client, then click **Next**.
10. Verify the information then click **Finish**. The new DFS root appears in the Logical View tree.

Create a domain-based DFS root

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

Prerequisites

- Microsoft DFS server software (Distributed File System Service `dfssvc.exe`) must be running on the server that hosts the DFS root.
- Microsoft Active Directory must already be operational.
- The account used for VFM must have administrator permissions for both the domain and the server where you create the DFS root.

- When creating domain-based roots, it is strongly recommended that you run VFM on a system that is a member of the domain that hosts the root. Name resolution problems and permission problems often manifest themselves in scenarios where the administrative machine resides in a Windows NT 4 domain and is attempting to create a domain-based root in a separate Active Directory domain.
- Windows 2000 servers can host one only DFS root.
- Windows Server 2003, Enterprise Edition, can host multiple DFS roots. To manage this configuration, VFM must be installed on Windows XP Professional or on Windows Server 2003.

To create a domain-based DFS root

1. Perform one of the following actions:
   - Right-click **Logical View** and click **Add a new DFS root**.
   - Click **Create DFS root**.
   - On the **Tools** menu, click **Add a new DFS root**.
2. Click **Next**.
3. In the **Server name** field, type the name of the server where you want to create the DFS root or click **Browse** to select an available server.
4. Click **Next**.
5. Select **Create a domain-based DFS root** and click **Next**.
6. Select an existing shared folder or create a new shared folder.
   - Best practices dictate that the share name used should match the name of the root that is being created.
   - If you clicked **Use an existing shared folder**, type the path and name of the shared folder or select an available shared folder.
   - If you clicked **Create a new shared folder**, type the path and name of the shared folder in the appropriate fields.
7. Click **Next**.
8. In the **DFS root name** field, optionally type a name for the root (follow the file naming conventions for the operating system of the server).
9. Optionally type a word or phrase in the **Comment** field.
10. In the **Clients cache this referral for** field, optionally change the duration (in seconds) that the root is cached on a DFS client then click **Next**.
11. Verify the information.
12. Click **Finish**. The new DFS root appears in the DFS Tree window.

Create a consolidation DFS root

Consolidation DFS roots are special stand-alone roots that enable you to move files referenced in UNC paths embedded in links, in line-of-business applications, and in other places without affecting users. The consolidation root redirects the user from the old location referenced by the UNC path to the new location. For more information and an illustration about how consolidation roots work, refer to "How consolidation DFS roots work" on page 91

Prerequisites

- Microsoft DFS server software (Distributed File System Service `dfssvc.exe`) must be running on the server that hosts the DFS root.
- You must install a software update that enables you to host a consolidation root on the server you select. Contact Microsoft for the software update. After you apply the software update, you must restart the Distributed File System service. For more information about consolidation roots, refer to Microsoft Knowledge Base Article 829885.
- Before you configure the consolidation root, you must rename the server to be consolidated. For more information, refer to "How consolidation DFS roots work" on page 91 for clarification.
- The consolidation roots must be installed on Windows Server 2003. It is recommended that you host the root on Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.
- The consolidation root must be hosted on a member server and not a domain controller because the server consolidation logic is disabled if the root is hosted on a domain controller.
- The NETBIOS name of the server that hosts the consolidation root must be the same as the host name part of its fully qualified domain name (FQDN).
- The consolidation of UNC paths from two servers that have the same host name is not supported. For example, you cannot consolidate two servers that are named `server1.npd.com` and `server1.<subdomain>.npd.com`.
- The account used for VFM must have administrator permissions for the server where you create the DFS root.
- Windows Server 2003, Enterprise Edition, can host multiple DFS roots. To manage this configuration, VFM must be installed on Windows XP Professional or on Windows Server 2003.

To clarify the instructions about how to configure a consolidation DFS root, the following terms are used:

- The server that hosts the consolidation DFS root is named *ConsolidationServer*.
- The server to be consolidated is named *Server1*. *Server1* is renamed to *OldServer1* before the consolidation DFS root is created.
- Data is migrated from *OldServer1* to *ConsolidationServer*.

To create a consolidation DFS root

1. Perform one of the following actions:
   - Right-click **Logical View** and click **Add a new DFS root**.
   - Click **Create DFS root**.
   - On the **Tools** menu, click **Add a new DFS root**.
2. Click **Next**.
3. In the **Server name** field, type the name of the Windows Server 2003 server that will host the DFS root or click **Browse** to select an available server. For example, you would type *ConsolidationServer*.
4. Click **Next**.
5. Select **Create a DFS Consolidation root** and click **Next**.
6. In the **Server name** field, type the original name of the CIFS server that you want to consolidate. For example, you would type *Server1*.

7. In the **Local Path of New Share** field, type the path of the share on the server hosting the root that will contain the root configuration information.
8. Click **Next**. VFM uses the name of the server you want to consolidate as the name of the share of the consolidation root. The share name is preceded by a hash (#) character. For example, *\\ConsolidationServer\#Server1* would be the name of the consolidation DFS root.
9. Optionally, type a word or phrase that best describes the root in the **Comment** field.
10. Optionally, change the duration (in seconds) that the root is cached on a DFS client and then click **Next**.
11. Click **Finish**. The root appears in the Logical View tree.
12. After you create the root, you must perform the following actions:

TABLE 22 Consolidating DFS roots

Action		Description
1	Make sure that **Automatically update DHCP client information in DNS** is selected on the DHCP server.	• If you use a Microsoft Windows 2000 DHCP server, make sure that the **Automatically update DHCP client information in DNS** check box is selected only at the server level and not selected at the DHCP scope level. • If you use a Microsoft Windows 2003 DHCP server, make sure that the **Enable DNS dynamic updates according to the settings below** check box is selected only at the server level and not selected at the DHCP scope level. • To clear the DHCP scope level, perform the following actions: 1 On the DHCP server, click **Start**, point to **Programs**, point to **Administrative Tools**, and click **DHCP**. 2 Expand the DHCP server name, right-click the first Scope on the list, and click **Properties**. 3 Click the **DNS** tab. 4 For a Microsoft Windows 2000 DHCP server, clear the **Automatically update DHCP client information in DNS** check box. For a Microsoft Windows 2003 DHCP server, clear the **Enable DNS dynamic updates according to the settings below** check box. 5 Repeat Steps 2 through 4 for each DHCP scope.
2	On the Windows Internet Name Service (WINS) server, verify that static WINS entries do not exist for the server to be consolidated. For example, you would make sure that WINS entries do not exist for *Server1*. If WINS entries exist, delete them. For more information about how to delete static WINS entries, contact Microsoft.	

TABLE 22 Consolidating DFS roots

Action	Description
3	From a command-line prompt on the server hosting the consolidation DFS root, type the following command: `ipconfig /registerdns`
4	Restart the Server service on the server hosting the consolidation DFS root. If you stop and restart the Server service instead of restarting it, you must start the DFS service and Computer Browser service. Whenever you stop the Server service, the DFS and Computer Browser services are also stopped.

Optionally, you can log a system event every time that the DFS root consolidation server successfully resolves a referral. The log enables you to see which clients are accessing the consolidated data through the old server name.

To log a system event each time the DFS root consolidation server resolves a referral

1. Locate then click the following registry key in `HKEY_LOCAL_MACHINE`:

   `SYSTEM\CurrentControlSet\Services\Dfs\Parameters\Replicated`

2. On the **Edit** menu, point to **New**, and click **DWORD Value**.
3. Type **LogServerConsolidation** and press **Enter**.
4. In the **Value** data box, type **1** and click **OK**. This will log the logical path and the client's IP address to the system event log as event ID: 14318. Allocate sufficient space to the system log.

Microsoft recommends that you set the system event log to at least 64 megabytes. Make sure to set the system event log to overwrite events as needed.

- To configure a consolidation DFS root in a cluster server environment, refer to the Microsoft Knowledge Base Article 829885.
- After you create the consolidation DFS root, migrate the data to the new location by creating a migration policy and create links under the consolidation root that target shares in the new location that will be redirected. For example, a link under the consolidation root *\\ConsolidationServer\#Server1* would target the location of a share on *ConsolidationServer*. For more information and instructions about how to configure a migration policy and add links to the consolidation root, refer to "Create a migration policy" on page 178 and "Add multiple links" on page 101.
- Links in a consolidation root cannot be referenced under a logical folder. Links must always be directly under the consolidation root. If you try to access a link under a logical folder, you get the following error:

  `No network provider accepted the given path.`

Correct configuration of a consolidation root. Links are listed directly under the root.

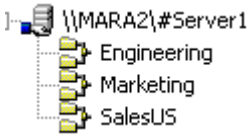

Incorrect configuration of a consolidation root. The targets of links Engineering and Marketing cannot be accessed because they are configured under the Consolidation Data logical folder.

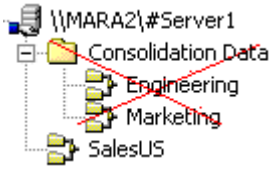

- The target of a link of a consolidation root must be a share. A consolidation root cannot be the target of a link of a DFS root. Refer to the following illustration.

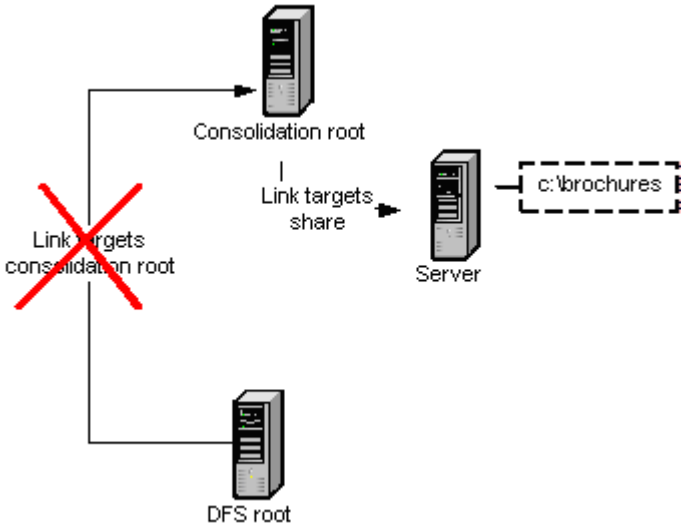

How consolidation DFS roots work

The following diagrams illustrate how consolidation DFS roots enable you to migrate data without affecting UNC path references embedded in user documents.

When to use a consolidation root

You might want to configure a consolidation root when you need to move data from a server that contains files that are referenced in users' documents. When the files are moved to the new location, all UNC path references to the files on the old server must be updated to point to the new location. A consolidation root enables you to move the referenced files without having to update UNC path references.

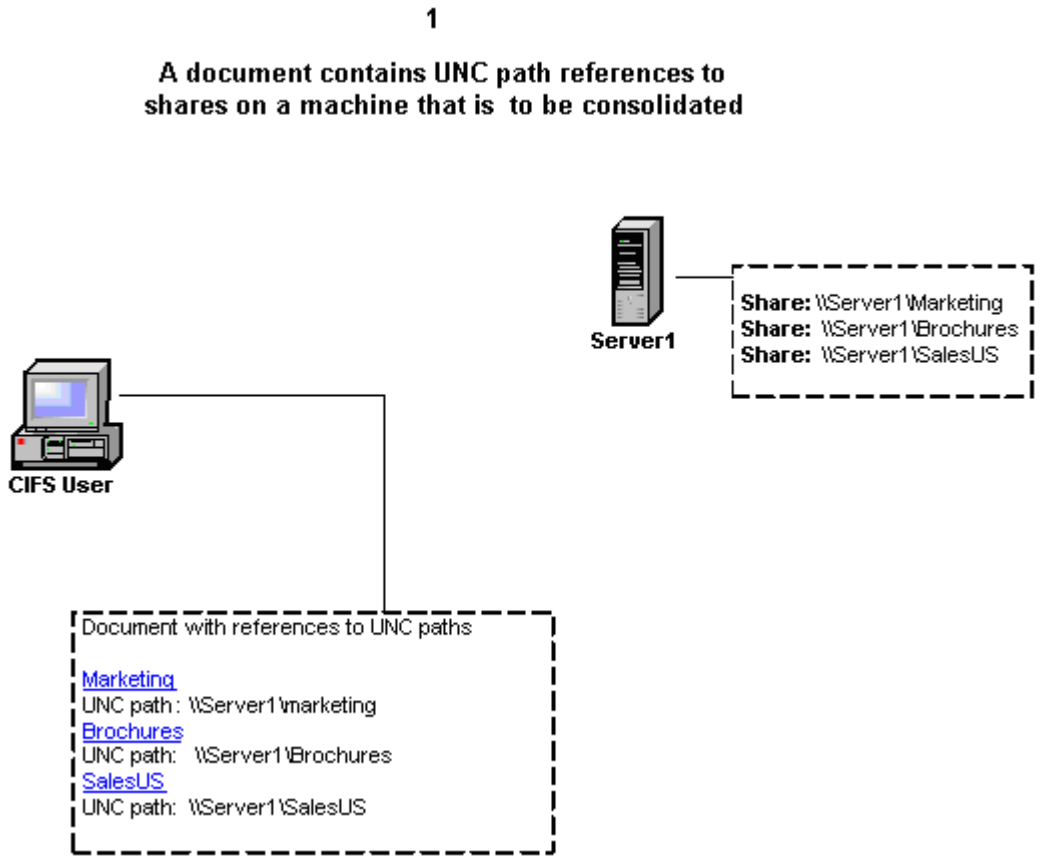

Rename the server you want to consolidate

Before you configure a consolidation root, rename the server you want to consolidate.

Create a consolidation DFS root

VFM provides a wizard that steps you through the process of creating a consolidation root. There are other steps you must perform when configuring a consolidation root. For example, you must install a Microsoft software update on the machine that will host the root. For more information and detailed instructions, refer to "Create a consolidation DFS root" on page 87.

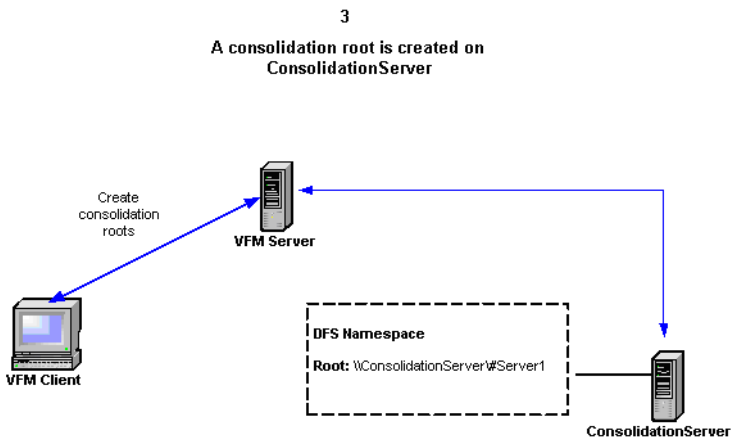

Move the data from old server location to the new location

Create a migration policy to move the data to the new location and create links to shares in the new location.

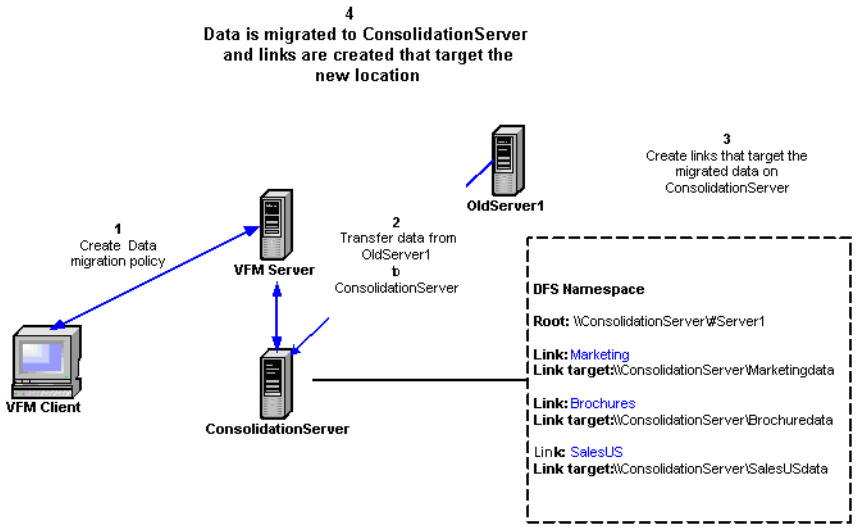

Old UNC path references are redirected

The old UNC path references to shares in the old server are redirected to the new location of the shares.

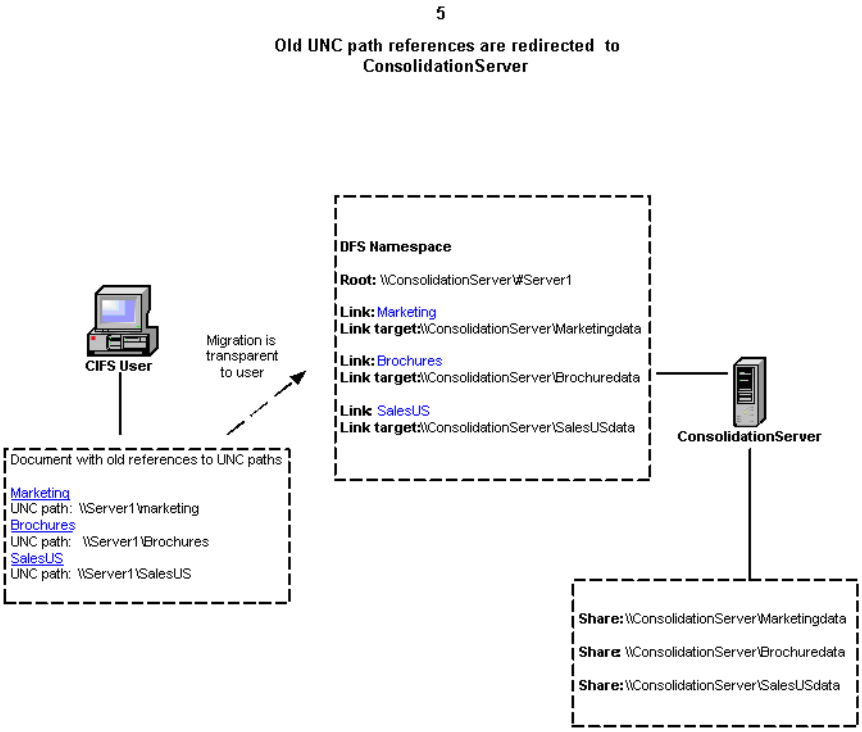

Display an existing DFS root

Display DFS roots that already exist on your network.

1. Perform one of the following actions:
   - Right-click **Logical View** and click **Display an existing DFS root**.
   - Click **Display an existing DFS root** on the toolbar.
   - On the **Tools** menu, click **Display an existing DFS root**.
2. Perform one of the following actions:
   - Type the name of the DFS root, its host server, or the domain.
   - Select the domain that contains the DFS root, and select the root. If a server is configured so that it does not appear in a browser, it does not appear in this list. You can still display the server in the Logical View tree by typing the name of the root or server in the **DFS Root, Host Server**, or **Domain** field.
3. Click **OK**.

NOTE
To view all the roots in the namespace, click **Logical View**.

Monitor DFS root status

By default, DFS roots and links are monitored. You can turn off status monitoring for a root.

1. Perform one of the following actions:
   - Click the root. On the **Properties** tab that appears in the right pane, click **Edit**.
   - Right-click the root and click **Properties**.
2. Click the **Status Monitoring** tab, and clear the **Monitor status every** check box.

Set permissions on a domain-based DFS root

Set or change permissions for a domain-based DFS root.

1. Right-click the root and click **Properties**.
2. Click the **Security** tab.

Modify properties of a DFS root

Edit the properties of a root displayed in the Logical View tree.

1. Perform one of the following actions:
   - Click the DFS root to display the **Properties** tab in the right pane, and click **Edit**.
   - Right-click the DFS root and click **Properties**.
2. Update the information and click **OK**.

View attributes of a DFS root

Click a root to display the following tabs in the right pane.

- **Properties** - Shows the characteristics or parameters of the selected root and enables you to edit certain fields. Right-click an item to display a context menu.
- **Status** - Shows information for any monitored item (for example, root, link, and target). By default, roots and links are monitored. You can choose to enable monitoring of targets. Right-click a root and click **Properties**, and click the **Status Monitoring** tab.
- **Replication Tasks** - Shows information about replication operations for items forming part of the root.
- **Explorer** - Shows the namespace, as a user browsing with Windows Explorer would see it. Right-click a root share, a single target link, or replicated link to display a context menu.
- **Dependencies** - Shows the underlying resources, such as systems, shares, and disks, forming part of the root. Right-click an item to display a context menu.
- **Events** - Shows the history of events for the item selected.
- **Reports** - Shows a list of reports generated that reference the root.

Back up a DFS root

You can manually back up the logical structure of a root.

1. Perform one of the following actions:
   - Right-click a root and click **Backup**.
   - Right-click **Logical View** and click **Backup**.
   - On the **Tools** menu, click **Manual backup**.

2. Select the DFS roots you want to back up.
3. Accept the default name for the backup file or type a new name.
4. Click **Backup**.
5. Click **OK**. The backup file is saved on the VFM server.

NOTES:

- To check whether the backup operation is complete, click the **Events** tab for each root that was backed up.
- You can set up a policy to automatically back up selected DFS roots. On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, click **Backup Policy**.

Refresh a DFS root

Update the contents of a root in the Logical View tree.

Perform one of the following actions:

- Select the root and press **F5**.
- Select the root and click **Refresh** on the main toolbar.
- Right-click the root and click **Refresh**.
- Select the root and click **Refresh** on the **View** menu.

Fresh information is gathered from the network for all the items that make up the DFS root.

NOTES:

- Gathering fresh information for the entire DFS root takes extra time if the DFS root is large.
- To refresh all the items in the namespace, click **Refresh All** on the **View** menu.

Check status of a DFS root

If needed, you can manually check the status of a DFS root to ensure that the root is still valid.

Right-click the DFS root and click **Check status**.

NOTE
To remove the status display from a root, right-click the root and click **Clear status**.

Remove a DFS root from display

When the **Logical View** displays too many roots to manage effectively, or you complete the management of a root, you can remove the root from appearing in the VFM console. This does not delete the root from the server.

Right-click the root and click **Remove** from display.

Delete a DFS root

When you delete a root, it is removed from the server. The configuration files are removed, but files that were referenced by the root are not deleted. When you delete a root, the clients can no longer browse it using the namespace.

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

Create a root replica

Specify additional servers to host copies of the logical structure of a namespace to provide fault tolerance and load balancing. Only domain-based DFS roots can contain more than one root target.

1. Perform one of the following actions:
   - Right-click the root and click **Add replica**.
   - Click **Logical View** to display a **Properties** tab. Right-click a root and click **Add replica**.
2. Type the name of the server that will host the root target, or click **Browse** to find the server.
3. Click **Next**.
4. Perform one of the following actions:
   - Select **Use an existing share** and select a shared folder from the list.
   - Select **Create a new share** and type a local path and the name for the shared folder. The application creates the shared folder on the server specified.
5. Click **Finish**. The root is now represented in the Logical View as root with multiple targets.

NOTES:

- You can create only one DFS root on Windows 2000 servers. Windows Server 2003, Enterprise Edition, can host multiple DFS roots.
- When creating or selecting a share, best practices dictate that the share name match the name of the root.
- To provide fault tolerance and load balancing of stand-alone roots, create a namespace availability policy.

View the replicas of a root

Display a list of root targets for a domain-based DFS root.

1. Click the DFS root in the Logical View tree.
2. Click the **Properties** tab that appears in the right pane. The targets of the root appear under **Replicas**.

NOTES:

- Right-click a root target to display a context menu.
- Double-click a root target to display where it is physically located on your network. The Physical View tree expands to display the machine where the target resides.

View contents of a root replica

View the contents of domain-based root targets. To browse the contents of a root target, you must make it the active target.

1. In the Logical View tree, select a domain-based root target.
2. Click the **Explorer** tab in the right pane.

3. Right-click a domain-based root target, point to **Target**, and click the target you want to activate. The contents of the replica appear in the right window.

NOTES:

- You can also view the contents of a link, share, or folder. In the Logical View tree or Physical View tree, right-click an item, point to **External Tools**, and click **Explorer**.
- Changing the active target only changes the target the VFM console machine is viewing. It does not change what other users of the namespace.

Compare contents of a root target

Compare the contents of a domain-based root target.

1. Click a root under the Logical View to display its properties in the right pane.
2. Under **Replicas**, select the targets that you want to compare.
3. Right-click, point to **Reports**, point to **Physical**, and click **File Differences**. The first target on the list that you selected appears in the **Compare** field.
4. Click **Next**.
5. Type a name for the report and the location at which you want to save it.
6. Click **Finish**.

NOTE
You can also compare contents of shares in the Physical View. Right-click a share, point to **Reports**, and click **File Differences**.

Delete a root replica

You can delete a domain-based DFS root replica and not delete the root as long as the domain-based root contains more than one root replica. If all root replicas are deleted, the root is deleted.

1. Select a root in the Logical View. A **Properties** tab appears in the right pane.
2. Under **Replicas**, right-click a shared folder and click **Delete**.
3. Click **Yes**.

Create and manage links

The next step in creating a namespace is to create links. Links form a connection to one or more shared folders, other links, or to another root. This section provides instructions about how to create and manage links and link targets.

NOTE
The Namespace Creation Policy and Permission Based Policy wizards enable you to add multiple links and folders at a time. Use these methods when you want to add links to multiple shares that are supported by different machines and when you want to keep the logical namespace aligned with subsequent changes in the underlying physical resources.

Add a DFS link

Add links to the namespace using any one of the following operations.

- Drag a share from Windows Explorer to a DFS root in the Logical View tree
- Drag an item from the Physical View tree to a DFS root in the Logical View tree
- Use the Add Link dialog box to perform the following tasks:
  - To make a shared folder or subfolder the target of a link
  - To add a link to a folder in the namespace
  - To add a link to multiple folders in the namespace
  - To add a link to multiple roots

A link can appear below a root or links can be grouped into folders. The target of a DFS link must be one of the following types:

- 1 to N shared folders
- 1 to N stand-alone DFS paths
- Only 1 domain-based DFS path

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

To drag a share from Windows Explorer to a DFS root in the Logical View tree

Drag a network share (UNC path) from Windows Explorer to a DFS root in the Logical View tree.

NOTES:

- You cannot drag local folders (for example, c:\<path>) to a DFS root from Windows Explorer.
- If you want to drag shares from your local system into the DFS namespace, browse Windows Explorer to your local host through the Network Neighborhood so that Explorer displays a UNC path name to your local share before performing the drag-and-drop operation.

To drag an item from the Physical View tree to the Logical View tree

Drag a shared folder from the Physical View tree to a DFS root in the Logical View tree.

To make a shared folder or subfolder the target of a link

1. Perform one of the following actions:
   - Select a root in the Logical View tree and click **Add link**.
   - Right-click a DFS root in the Logical View tree and click **Add link**.
   - Right-click a DFS root on the **Properties** tab of **Logical View** and click **Add link**.
   - Select a DFS root in the Logical View tree and click **Add link** on the **Tools** menu.
   - Drag a network share (UNC path) to the Logical View.
2. Under **Users connect to**, type the name of the link.
3. Click **Add**.
4. Under **Send the user to one of these shared folders or directories**, perform one of the following actions:

- Type the name of a folder.
- Click **Browse** to find a share or subfolder in your network or the Physical View tree.
- Drag a network share and drop it into the **Send the user to one of these shared folders or directories** field.

5. Click **OK**.
6. In the **Comment** field, optionally type a word or phrase that describes the link.
7. In **Clients cache this link for**, type how long (in seconds) a reference to the link is cached on the DFS client.
8. Click **OK**. The DFS link appears under the root selected.

To add a link to a folder

Group links in logical folders. Create a folder on a DFS root server and add links to it.

1. Perform one of the following actions:
   - Select a root in the Logical View tree and click **Add link.**
   - Right-click a DFS root in the Logical View tree and click **Add link**.
   - Right-click a DFS root on the **Properties** tab of **Logical View** and click **Add link**.
   - Select a root in the Logical View tree and click **Add link** on the **Tools** menu.
2. Under **Users connect to**, type a name for the folder and the name of the link; for example: *foldername\linkname or folder1\folder2\linkname*.
3. Click **Add**.
4. Under Send **the user to one of these shared folders or directories**, type the name of a folder, or click **Browse** to find a share or subfolder in your network or the Physical View tree.
5. Click **OK**.
6. In **Comment**, type a word or phrase that describes the link.
7. In **Clients cache this link for**, type how long (in seconds) a reference to the link is cached on the DFS client.
8. Click **OK**. The folder appears under the DFS root selected.

To add a link to multiple folders

Add the same link to multiple folders in the Logical View.

1. Select a folder with subfolders to display a **Properties** tab in the right pane.
2. On the **Properties** tab, right-click the selected subfolders and click **Add links**.
3. Under **Users connect to**, type the name of the link.
4. Click **Add**.
5. Under **Send the user to one of these shared folders or directories**, perform one of the following actions:
   - Type the name of a share.
   - Click **Browse** to find a share or subfolder in your network or the Physical View tree.
   - Drag a network share to **Send the user to one of these shared folders or directories**.
6. Click **OK**.

7. In **Comment**, type a word or phrase that describes the link.
8. In **Clients cache this link for**, type how long (in seconds) a reference to the DFS link is cached on the DFS client.
9. Click **OK**. The link appears under the folders selected.

To add a link to multiple roots

Add the same link to multiple roots in your namespace.

1. Select **Logical View to display a Properties tab** in the right pane.
2. On the **Properties** tab, right-click the selected roots and click **Add links**.
3. Under **Users connect to**, type the name of the link.
4. Click **Add**.
5. Under **Send the user to one of these shared folders or directories**, perform one of the following actions:
   - Type the name of a share.
   - Click **Browse** to find a share or subfolder in your network or the Physical View tree.
6. Drag a network share to **Send the user to one of these shared folders or directories**.
7. Click **OK**.
8. In **Comment**, type a word or phrase that describes the link.
9. In **Clients cache this link for**, type how long (in seconds) a reference to the link is cached on the DFS client.
10. Click **OK**. The link appears under the roots selected.

NOTES:

- To display hidden shared folders, select the **Show hidden shares** check box.
- By default, only server machines are displayed. To also display workstations, select **Show workstation computers**.
- The following reserved device names cannot be used as the name of a link: CON, PRN, AUX, CLOCK$, NUL, COM1, COM2, COM3, COM4, COM5, COM6, COM7, COM8, COM9, LPT1, LPT2, LPT3, LPT4, LPT5, LPT6, LPT7, LPT8, and LPT9.

Add multiple links

You can visualize the namespace you are designing and see the namespace as you build it. The Logical View tree is updated as you add a link or logical folder.

NOTE

Device names cannot be used as the name of a link. For more information, refer to "Add a DFS link" on page 99.

To add multiple links

Add multiple links to the DFS root selected or to an existing folder in the namespace.

1. Perform one of the following actions:
   - Right-click a DFS root in the Logical View and click **Add multiple links**.

- Right-click a DFS root on the **Properties** tab of **Logical View** and click **Add multiple links**.
- Select the root and click **Add multiple links** on the **Tools** menu.

2. To add links, perform one of the following actions:
   - Drag network shares (UNC paths) from Windows Explorer to the tree view in the Add Multiple Links dialog box.
   - Select the DFS root or a folder and click **Add Links**.
   - Right-click the DFS root or folder and click **Add links**.

   A dialog box appears, enabling you to browse your network or Physical View tree.
3. Expand the tree you want to search and select an item or items.

NOTES:

- To display hidden shared folders, select the **Show hidden shares** check box.
- By default, only displays server machines are displayed. To also display workstations, select Show workstation computers.
- To view a context menu, right-click a shared folder.

4. Click **OK**. The links you created appear in bold print in the namespace tree.
5. Repeat Steps 2 through 4 to create additional links.
6. Click a link to display its properties.
7. Optionally update the link properties
   - **Name** - Type a new name for the link.
   - **Send the user to this shared folder or directory** - Change the target of the link. Type the address of a new target or click **Browse** to locate another target.
   - **Comment** - Add an optional description of the link.
   - **Clients cache this link for -** Set the duration (in seconds) the target of the link is cached on a DFS client.
8. Click **OK** to apply the changes to the namespace. A status window displays the changes being made.
9. Click **Done**.

To add multiple links to a folder

Create folders to group links in a common directory in the Logical View. You can also create a folder when you add a link in the Add a Link dialog box.

1. Select the DFS root or a folder in the namespace of the Add Multiple Links dialog box.
2. Click **New Folder**.
3. Type a name for the folder and press **Enter**.
4. Perform one of the following actions to add links to the folder:
   - Select the folder and click **Add Links**.
   - Right-click the folder and click **Add Links**.
   - Drag network shares from Windows Explorer.
   - Cut and paste mapped network drives or items with valid UNC paths to the folder from Windows Explorer.

NOTES:

- You cannot drag or paste local folders from Windows Explorer.
- If you want to drag shares from your local system into the DFS namespace, browse Windows Explorer to your local host through the Network Neighborhood so that Explorer displays a UNC path name to your local share before performing the drag-and-drop operation.

To move a link or folder to a new location

After a link or folder is added to the namespace of the Add Multiple Links dialog box, you can move the item to a new location in the tree.

Perform one of the following actions:

- Drag the link or folder to a new location on the tree.
- Right-click the link or folder you want to move and click **Cut**. Right-click the link or folder where you want to move the item and click **Paste**.

Create new shared folders to link

Select the machine or machines where you want to create shares and add the shares to a new or existing link.

1. Perform one of the following actions:
   - On the **Tools** menu, click **Create shares**.
   - Right-click a link in the Logical View and click **Create shares**.
   - Right-click a machine in the Physical View tree and click **Create shares**.
2. Click **Next**.
3. Select **Regular Share** and click **Next**.
4. Type the information requested and click **Next**.
5. Select the machine or machines where you want to create the folder and click **Next**.
6. Select an option and click **Next**.
7. Perform one of the following actions:
   - If the path of the link appears, click **Next**.
   - If the path of the link does not appear, type the path or click **Browse** to locate the link in the namespace, and click **Next**.
8. Click **Finish**. The share is created on the machines and added to the link.

Expand a link

Replace a link targeting a share with links targeting the subdirectories contained in the original share. When the share contains data in its root, you can migrate the data and create a link to target the data. This enables you to access the data through the namespace.

Expand a link option enables you to perform the following:

- Move data in subdirectories to a different machine to load balance.
- Make the data in subdirectories available from different locations in the namespace.

NOTE
When expanding namespaces, the DFS limitations on numbers of links contained in a root need to be considered. For more information, refer to "Basic DFS requirements" on page 82

1. Right-click a link whose target contains subdirectories and click **Expand**.
2. If the same link appears in other locations of the namespace, a list of identical links appears. You can expand all, none, or some of the identical links. After you make your choice, click **OK**.
3. The subdirectories are identified and the actions that will be performed is explained. Click **Yes**.
4. Click **Done**. The original link is deleted, a logical folder with the same name as the original link is created, and new links are added to the logical folder that target the subdirectories.
5. If the original target contains data along with subfolders, you have the option to migrate the data to another location to make the data accessible through the namespace. To migrate the data, click **Yes** when the dialog box appears.
6. Perform one of the following actions:
   - Select an existing migration policy and continue with Step 10
   - Click **New Policy** to create a new migration policy and continue with Step 7.
7. Type the name for the policy.
8. Click **OK**.
9. Select the new policy.
10. Click **Next**. The original and destination paths appear.
11. Click **Next**.

    Click **Finish**. A link targeting the data appears under the same folder as the links targeting the subdirectories. The link name ends with `_data`.

Add a logical folder to the namespace

You can create folders in your namespace. The folders appear alphabetically in the Logical View tree. You can add links to the folders and move the folders and their links to another folder or DFS root.

1. Perform one of the following actions:
   - Right-click a DFS root or folder in the Logical View and click Add folder.
   - Select a DFS root to display a **Properties** tab in the right pane. Right-click anywhere on the page and click Add folder.
2. Type a name for the folder and click **OK**.

NOTES:

- The permissions on a DFS root share controls whether you can create logical folders in the namespace. Windows Server 2003 defaults share permissions to read-only and Windows 2000 Server defaults share permissions to allow everyone full access.
- A folder and its links can be moved to another folder or DFS root. Drag the folder to its new location.

- Dragging a folder moves only its logical structure but does not move data physically resident in the folder on the DFS root server. Best practices dictate that data should not reside physically on the DFS root server because this can increase the load on the server by forcing it to serve files as well as the logical namespace. If you have data in folders in your DFS namespace, use the VFM migration features to move it to the new location where you transferred the folder.
- To create a subfolder, right-click the newly created folder and click **Add folder**.
- To add links to the folder, right-click the folder and click **Add link** or **Add multiple links**. You can also drag links to the folder.

Rename a DFS link or folder

To change the name of a link or folder in the Logical View tree perform one of the following actions:

- Click a link or logical folder slowly to display a text box and type a new name.
- Right-click the link or folder and click **Properties**, type a new name, and click **OK**.
- Click the link to display a **Properties** tab in the right pane and click **Edit**. Type a new name and click **OK**.

NOTES:

- If the namespace contains links that are identical to the link you renamed, you are given the option of renaming all the identical links in one easy step.
- Device names cannot be used as the name of a link. For more information, refer to "Add a DFS link" on page 99.

Monitor DFS links

To get status monitoring information about links, you must go through the DFS root that contains the links. If status monitoring is turned on for a DFS root, you can also turn on status monitoring for its links.

1. Perform one of the following actions:
   - Right-click the root that contains the links you want to monitor and click **Properties**.
   - Click the root to display the **Properties** tab in the right pane and click **Edit**.
2. Click the **Status Monitoring** tab and select **Monitor links every**.
3. Enter how often (in minutes) you want to check the status of each link when it is online.
4. Enter how often (in minutes) you want to check the status of each link when it is offline.

View attributes of a DFS link

Select a DFS link to display the following tabs in the right pane.

- **Properties** - Shows the characteristics or parameters of the selected link and enables you to edit certain fields. Right-click an item to display a context menu.
- **Status** - Shows information about the link. This tab also shows information about targets when target monitoring is turned on. By default, roots and links are monitored. To enable monitoring of targets, edit the properties of the DFS root to which the link belongs.
- **Replication Tasks** - Shows information about link target replication operations.

- **Explorer** - Shows the namespace as a user browsing with Windows Explorer would see it. Right-click a root share, a single target link, or a replicated link to view a context menu.
- **Dependencies** - Shows the underlying resources, such as systems, shares, and disks that form part of the link. Right-click an item on the tab to display a context menu.
- **Events** - Shows the history of events for the item selected.
- **Reports** - Shows a list of reports generated that reference the link and the targets of the link.

Modify properties of a DFS link

Edit the properties of a DFS link in the Logical View tree.

1. Perform one of the following actions:
   - On the **Properties** tab of a link, click **Edit**.
   - Right-click the DFS link and click **Properties**.
2. Update the information and click **OK**. If the namespace contains links that are identical to the link you modified, you have the option of updating all the identical links in one easy step.

NOTES:

- When you browse for a folder, you have the option of creating, renaming, deleting, and sharing new folders, and removing a share.
- Right-click a shared folder to view a context menu.
- Double-click a shared folder and right-click a folder or anywhere in the window to view a context menu.
- 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 updated 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. For more information, refer to Microsoft Knowledge base article Q222604.

Refresh a DFS link

Update the contents of a DFS link in the Logical View tree. Fresh information is gathered from the network for all the targets of the DFS link.

Perform one of the following actions:

- Select the link and press **F5**.
- Select the link and click **Refresh** on the main toolbar.
- Right-click the link and click **Refresh**.
- Select the link and click **Refresh** on the **View** menu.

NOTE
To refresh all the items in the namespace, click **Refresh All** on the **View** menu.

Check the status of a DFS link

If needed, you can manually check the status of a DFS link to ensure that it is still valid.

Right-click the DFS link and click **Check status**.

NOTES:

- To remove the status display from a link, right-click the link, and click **Clear status**.
- If VFM is actively monitoring the status of a link or its targets, clearing status only takes effect after VFM finishes checking status for the link.

Delete a DFS link

When you delete a DFS link, the DFS logical structure is removed but the targets of the DFS link are not deleted.

1. Perform one of the following actions:
   - Select the DFS link and click **Delete** on the main toolbar.
   - Right-click the link and click **Delete**.
   - To delete all the links under a logical folder, right-click a logical folder and click **Delete**.
2. Click **Yes**.

NOTES:

- Links contained in any of the Admin policies must be removed from the policy before they can be deleted from the root.
- When you want to delete all the links under a logical folder, it is more efficient to delete the logical folder than it is to select all the links under the logical folder and delete the links.
- When a root runs on Microsoft Windows Server 2003 and you delete all the links under a logical folder, the logical folder is also removed. In Windows 2000, logical folders are not removed. This is a Microsoft Windows behavior. The following example illustrates this behavior.

The logical folder *Marketing* contains two links.

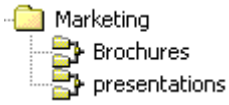

DFS root on Windows Server 2003. When links *Brochures* and *presentations* are deleted, the logical folder *Marketing* is also deleted.

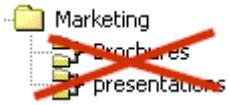

DFS root on Window 2000. When links *Brochures* and *presentations* are deleted, the logical folder *Marketing* is not deleted.

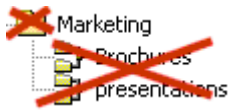

Manage link targets

- A link can have as targets:
  - 1 to N shared folders
  - 1 to N stand-alone DFS paths
  - Only 1 domain-based DFS path.
- When a link has two or more targets, the roots or shared folders that are the targets of the link provide alternate paths for the link.

View the target of a link

To display the target or targets (replicas) of a link:

1. Click the link in the Logical View tree.
2. On the **Properties** tab, scroll to **Targets**.

NOTES:

- Right-click a target to display a context menu.
- Double-click a target to display where it is physically located on your network. The Physical View tree expands to display the machine where the target resides.
- Targets that are members of a client recovery policy have radio buttons instead of check boxes next to them because only one target can be online. Replication status information appears to the right of each target in addition to the storage status information available for all targets.

Add targets to a link

Add multiple copies of a shared folder to a link. This provides fault tolerance and load balancing to control network traffic at the share level.

- There are two ways that you can add link targets.
  - Quickly add multiple targets at one time. Use this method when the path pattern of the target you want to rebase is the same as the path pattern of the target you want to add to the link. For example, you can rebase the link target `\\secure.npb.com\engineering` with `\\saturn.ppn.com\engineering` but you cannot rebase the link target `\\secure.npb.com\engineering` with `\\saturn.ppn.com\dept\engineering` because they do not have the same path pattern.
  - Add target paths and update how long a reference to the link is cached on the client. Use this method when the path pattern of the existing link target is different from the path pattern of the target you want to add to the link. For example, you have an existing link target `\\secure.npb.com\engineering` and you want to add another link target `\\saturn.ppn.com\dept\engineering`.
- The target of a link must be one of the following types:
  - 1 to N shared folders
  - 1 to N stand-alone DFS paths
  - Only 1 domain-based DFS path

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

To quickly add multiple link targets

1. Right-click a link and click **Clone**.
2. Perform one of the following actions:
   - When the link contains multiple targets, select the target that you want to rebase and click **Next**.
   - When the link contains one target, continue with Step 3.
3. Select which part of the UNC path of the link target you want to rebase.
4. Perform one of the following actions:
   - Click **Browse** to find a machine in your network.
   - Type the name of a machine or multiple machines that you want the link to target. Separate machine names with a semicolon (;). Machine names can be written in any of the following formats:

Format	Example
machine name	`secure;saturn`
UNC path	`\\secure.npb.com;\\saturn.ppn.com`
path	`secure.npb.com;\\saturn`

5. Click **Next**. VFM shows the target you want to rebase and the new link targets that will be created.
6. Click **Finish**.

To add a link target

1. Right-click a link and click **Add replica**.
2. Click in **Send the user to one of these share folders or directories** and perform one of the following actions:
   - Type the path of a folder and continue with Step 6.
   - Click **Browse** to find a folder on the network and continue with Step 3.
3. Expand the tree for the domain or click the domain.
4. Double-click a machine to display a list of shared folders.
5. Click **OK**.
6. To add another target, click **Add** and repeat Step 2.
7. Click **OK**.

NOTES:

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

Change storage status of a link target

- A link target on DFS roots hosted on Windows 2000 or Windows Server 2003 can be taken offline. The action does not remove the target from the link. The target is temporarily unavailable to DFS clients (not included in referrals to the DFS client).
- Links with multiple targets that are members of recovery policies can only have one target online. Targets of those links appear in VFM with radio buttons.
- Referrals for the link do not contain information about the offline target, although clients already accessing the target continue to do so until they request a fresh referral.

To take a link target offline

1. Click the link that references the shared folder to display the **Properties** tab in the right pane.
2. Under **Targets** on the **Properties** tab, perform one of the following actions:
   - Clear the check box next to a target.
   - Select a radio button next to a target to take the online target offline. Targets that are members of a client recovery policy have radio buttons instead of check boxes because only one target can be online. When you select a target, you perform a manual failover. The selected target becomes the active target.
   - Right-click the folder and click **Toggle Storage Offline/Online**.

NOTE
The offline status indicator appears on the folder.

To put a link target online

1. Click the link that references the shared folder. A **Properties** tab appears in the right pane.
2. Under **Targets** on the **Properties** tab, perform one of the following actions:
   - Select the check box or radio button next to a target. Targets that are members of a client recovery policy have radio buttons instead of check boxes because only one target can be online. When you select a target, you perform a manual failover. The selected target becomes the active target.
   - Right-click the offline folder and click **Toggle Storage Offline/Online**.

To change the storage status of all or selected targets

Quickly change the storage status of all or selected targets in a folder or root.

1. Click a root or folder in the Logical View tree to display the **Properties** tab in the right pane.
2. Under **Links**, select the links of the targets you want to change.
3. Right-click and click Toggle all targets online or **Toggle all targets offline**.

Monitor link targets

If status monitoring is turned on for links, you can turn on status monitoring for the targets of the links. Monitoring link targets provides additional information about the health of your namespace configuration by checking whether the targets of links are available.

1. Perform one of the following actions:
   - Right-click the root and click **Properties**
   - Click the root to display the **Properties** tab and click **Edit**.

2. Click the **Status Monitoring** tab and select **Monitor targets every**.
3. Enter how often (in minutes) you want to check the status of each target when it is online.
4. Enter how often (in minutes) you want to check the status of each target when it is offline.

Check status of a link target

If needed, you can manually check the status of a link target to ensure that it is still valid.

1. In the Logical View tree, click the link that references the shared folder to display the **Properties** tab in the right pane.
2. Under **Targets** on the **Properties** tab, right-click the target and click **Check status**.

NOTE
To remove the status display from a link target, right-click the target and click **Clear status**.

View size of data in the namespace

View the size and amount of space used by data in the namespace.

Perform one of the following actions:

- Right-click a root, link, or logical folder in the Logical View and click **Properties**.
- Right-click a share or directory in the Physical View and click **Properties**.

The size of the information contained in that portion of the logical namespace displays, along with the number of files, folders, and links when it is applicable.

Refresh a shared folder

You can update the information about a shared folder.

NOTES:

- To refresh the information about all the items in both the Logical View and Physical View, click **Refresh all** on the **View** menu.
- To refresh the information about all the items in the logical or physical views, right-click **Logical View** or **Physical View** in the tree and click **Refresh**.

Perform one of the following actions:

- Right-click a shared folder and click **Refresh**.
- Select a shared folder and press **F5**.
- Select a shared folder and click **Refresh** on the main toolbar.
- Select a shared folder and click **Refresh** on the **View** menu.

View targets replicated with FRS

Check the File Replication Service (FRS) status of the targets of a domain-based DFS root.

1. Click the root or link to display the **Properties** tab in the right pane.
2. VFM returns one of the following values for FRS Replication:
   - **Unknown** - Status monitoring is not turned on for the root or VFM has not checked status.
   - **On** - All the targets for the root have FRS replication turned on.
   - **Off** - FRS replication is not turned on for the root.

- **Partial** - Only some of the targets of the root have FRS replication turned on. (DFS roots or DFS links only)

NOTE
VFM replication does not rely on FRS.

View contents of a link target

View the contents of DFS link replicas. To browse the contents of a replica, you must make it the active target.

1. In the Logical View tree, click a DFS link with replicas.
2. Click the **Explorer** tab in the right pane.
3. Right-click a DFS link replica, point to **Target**, and click the replica you want to activate. The contents of the replica appear in the right window.

NOTES:

- You can also view the contents of a link, share, or folder. In the Logical View tree or Physical View tree, right-click an item, point to **External Tools**, and click **Explorer**.
- Changing the active target only changes the target the VFM console machine is viewing. It does not change what other users of the namespace.

Compare contents of link targets

Compare the contents of the targets of a link.

1. Click a link under the Logical View to display its properties in the right pane.
2. Under **Targets**, select the targets you want to compare.
3. Right-click, point to **Reports**, point to **Physical**, and click **File Differences**. The first target on the list that you selected appears in the **Compare** field.
4. Click **Next**.
5. Type a name for the report and the location where you want to save it.
6. Click **Finish**.

NOTE
You can also compare contents of shares in the Physical View. Right-click a share, point to **Reports**, and click **File Differences**.

Delete a link target

When a link contains multiple targets, you can delete a target without deleting the link from the namespace. The target is only removed from the logical structure or namespace. The target is not deleted from the machine where it resides.

1. In the Logical View tree, click the link that references the shared folder to display the **Properties** tab in the right pane. Under **Targets** on the **Properties** tab, right-click the target and click **Delete**.
2. Click **Yes**.

Monitor status

Monitoring the status of roots, links and targets

- The status of roots and links are monitored.
- You can control whether and how frequently the status of items in the namespace are checked.
- Monitoring the status of link targets is optional.
- When links and shares are monitored, the status report for a root includes information about all the items that make up the root. For example, a root with unavailable targets has a status report that includes a list of the targets and their status.
- Links in a client recovery policy show **unknown** status until monitoring has run at least once for the policy. Monitoring is necessary to perform failover of links that participate in recovery policies.
- On the **Status** tab of a root, right-click an item to view menu options.
- Monitoring resources generates network traffic between the VFM components and the resources being monitored. Slower polling rates generate less traffic but provide less timely notification of problems in the monitored namespace.

Monitoring status of data replication

- When you use VFM replication in a replication policy, replication status monitoring is available to the policy and its targets.
- Place the cursor over an item to quickly display information about the item.
- If status monitoring is turned on for the links and folders contained in a replication policy, status indicators are displayed for the items on the **Replication** tab that appears in the right pane when you select a policy.

Modify status monitoring properties

When you create a namespace, the status of the root and links contained in the root are monitored. You can control whether and how frequently the namespace is monitored. You can also select **Link targets for monitoring**.

1. Right-click a root and click **Properties**.
2. Click the **Status Monitoring** tab.
3. Perform one of the following actions:
   - To turn off monitoring the root, clear the **Monitor status every** check box.
   - To change how often the root is monitored, click the arrows to change **minutes when online** and **minutes when offline**.
4. Select the **Monitor links every** check box to monitor the status of every link in the namespace. Click the arrows to change **minutes when online** and **minutes when offline**.
5. Select the **Monitor targets every** check box to monitor the status of every target in the namespace. Click the arrows to change **minutes when online** and **minutes when offline**.

NOTES:

- You can also update status monitoring by clicking a root to display the **Properties** tab in the right pane.
- Monitoring resources generates network traffic between the VFM components and the resources being monitored. Slower polling rates generate less traffic but provide less timely notification of problems in the monitored namespace.

Populate the namespace automatically

Namespace creation policies enable you to easily populate and maintain a 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 namespace. You can configure policies to report changes or to update the namespace automatically to reflect the changes detected by the policy.

This section provides instructions about how to configure and manage namespace creation policies. Before you create your first policy, refer to "Configure a namespace creation policy" on page 114 for more information.

Configure a namespace creation policy

Configure a namespace creation policy to automatically populate the logical namespace and to keep the logical namespace aligned with subsequent changes in the underlying storage resources.

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Creation Policy**.
   - Right-click the **Namespace Policies** folder in Admin View, point to **New**, and click **Creation Policy**.
   - Right-click the **Creation Policies** folder and click **New Creation Policy**.
   - Right-click **Admin View**, click **New**, 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 (optional).
4. Clear **Automatically modify namespace after scanning** if you do not want to automatically create links for items that match the search criteria. After the policy runs, the items are listed under **New Resources** on the **Properties** tab that appears in the right pane when the policy is selected. You can create links for the items you want to include in the namespace.
5. Select one of the following:
   - Search for items not already targeted from this policy's folder when you do not want to target shares, subfolders, NetApp filer volumes, or qtrees that are already targeted by links in the folder. This is the default.
   - Search for items not already targeted from this root when you do not want to target shares, subfolders, NetApp filer volumes, or qtrees that are already targeted by links in the DFS root.
6. Select **Remove links whose shares or folders no longer exist** when you want to remove links targeting shares or subdirectories when a scan detects that the items no longer exist.

7. Select **Remove links from the entire root** when the links created by the policy are listed under a folder and you want to scan the entire root and remove links targeting the shares or subdirectories that no longer exist.
8. From the **Use the Monitor** on list, optionally select the machine where another monitoring agent is located.
9. Select whether you want to limit the number of links that are added every time the policy runs. Accept the limit specified or change the limit. When this option is cleared, adds links are added for all the candidates found. Microsoft recommends limits to the number of links that are added to a DFS root. For more information about recommended size limits, refer to "Basic DFS requirements" on page 82.
10. Click **Next**.
11. Select an existing DFS root or click **New** to create a new DFS root to contain the links, and click **Next**.
12. Group the links in a folder or add them under the root. Perform any of the following actions:
    - In the **Name** field, type a different name for the folder.
    - Click **New Folder** to create a new folder or subfolder.
    - Click **Delete** to remove the selected folder.
    - Click the root to add the links under the root.
13. Click **Next**.
14. Select the domain or domains that you want to search or select devices, and click **Next**.
15. Type a word or pattern in the **Search String** field. You can use DOS wild-card characters (* or ?) in the search string. You can use the asterisk (*) as a substitute for zero or more characters in a name and the question mark (?) as a substitute for a single character in a name. 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.
16. Select a search type.
17. If you selected **Search shares matching the search string for folders matching the following path**, type a folder search string.

NOTE
To search for folders, the VFM account must have Read and List Folder Contents permissions for the shares and subdirectories that match the search criteria.

---

18. (Optional) Select **Allow Partial Matches** if you want to search for partial matches.
19. (Optional) Select **Linked through target share** if the target folder is also shared when you want to create links that directly target shared subdirectories that match the folder search string.
20. (Optional) Search for hidden and administrative shares.
21. Click **Next**. A list of items that match your search string appears.
22. Clear the check box of items that you do not want to add, and click **Next**.
23. Select a template to arrange how links appear in the namespace. You can preview the namespace.
24. Click **Next**.

25. Rearrange the links, if desired. You can move a link or folder, add a folder, add more links, remove a link or folder from the namespace, and update the properties of a link.
26. Click **Next**.
27. Select **Use Schedule**.
28. Select how frequently you want to scan for items and click **Next**.
29. Click **Finish**. The new policy appears under the Creation Policies folder.

NOTES:

- If you uninstall a monitoring agent, ensure that namespace creation policies that formerly used the uninstalled agent are reassigned to a currently available agent. Right-click a namespace creation policy and click **Properties**. On the **General** tab, select a monitoring agent from the **Use this Monitor on** list.
- The Namespace Creation wizard creates links that reference volumes or qtrees regardless of their share status. Links that reference volumes or qtrees that cannot be accessed because they are not shared appear as broken links in the namespace.
- Permission-based namespace policies enable you to populate the namespace based on share security permissions. For more information about how to create a permission-based namespace policy, refer to "Configure a permission-based namespace policy" on page 120

Update the schedule

Make changes to the schedule of a namespace creation policy.

1. Perform one of the following actions:
   - Right-click a namespace creation policy and click **Properties**.
   - Select a namespace creation policy to show a **Properties** tab in the right pane, and click **Edit**.
2. Click **Schedule**.
3. Perform any of the following actions:
   - Click **Add** and configure a new schedule.
   - Select an existing schedule from the list and update it.
   - Select **Use alternate schedule when the policy fails** to attempt to run the policy at a stated frequency until the policy runs successfully.
4. Click **OK**.

Add or remove servers or domains

Add or remove servers and domains that a namespace creation policy scans.

1. Perform one of the following actions:
   - Right-click a namespace creation policy and click **Properties**.
   - Select a namespace creation policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Domains**.

3. Perform any of the following:
   - Click **Add** to browse for domains and servers to include in the policy.
   - To remove a server or domain from the policy, select it and click **Remove**. If a server being removed from the policy has shares being referenced by links, you have the option of removing the links from the namespace.
4. Click **OK**.

Add items to an exclusion list

Maintain an exclusion list of resources you do not want to reference from the namespace. The items on the exclusion list are ignored when a search is performed for resources to add to the namespace.

1. Perform one of the following actions:
   - Right-click a namespace creation policy and click **Properties**.
   - Select a namespace creation policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Target Search Ignore List**.
3. Click **Add** and select items to include in the list.
4. Click **OK**.

NOTE
When you create a new policy using the Namespace Creation wizard and you clear items from the search results list, the cleared items are added to the Ignore List of the policy.

Change the search type

Select a different item to search to update the namespace. You can search for shares, subfolders, NetApp filer volumes, or qtrees. The next time the policy runs, it scans for the new items.

1. Perform one of the following actions:
   - Right-click the namespace creation policy and click **Properties**.
   - Select a namespace creation policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Target Search**.
3. In the Search Type list, select a different item to search.
4. Optionally update the search string.
5. If you selected **Search Folders** as the search type, type a folder search string. The policy searches for subdirectories matching the folder search string of shares matching the string in the **Search string** field.
   - If you want to search for partial matches, select **Allow Partial Matches**.
   - If you want to create links that directly target shared subdirectories found, select **Linked through target share** if the target folder is also shared.
6. Click **OK**.

Change the search string of a policy

Change the search criteria for a namespace creation policy.

1. Perform one of the following actions:
   - Right-click a namespace creation policy and click **Properties**.
   - Select a namespace creation policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Target Search**.
3. In the **Search string** field, type a new word or pattern to match shares, volumes, or qtrees.
4. Optionally update the folder search string if **Search Folders** is the search type.
5. Click **OK**.

Change how the links are arranged

Change how new links are arranged in the namespace. Select a different template for a namespace creation policy. The next time the policy runs, the new links are arranged according to the new template. Existing links are not affected.

1. Perform one of the following actions:
   - Right-click the namespace creation policy and click **Properties**.
   - Select a namespace creation policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Target Search**.
3. Select a new template from the Hierarchy Template list. A graphic view of how new links will be arranged in the namespace appears in the lower part of the dialog box.
4. Click **OK**.

Scan for items immediately

You can bypass the namespace creation policy schedule and manually start the scanning process any time you want.

1. Perform one of the following actions:
   - Right-click a namespace creation policy and click **Start Scanning Now**.
   - Select a namespace creation policy to show a **Properties** tab in the right pane, right-click anywhere on the tab, and click **Start Scanning Now**.
2. Click **Scan**. If items that match the search criteria are found, they appear under the DFS root that contains the links.
3. Perform one of the following actions:
   - Click **Create** to add the links that reference the items to the namespace.
   - Click **Cancel** if you do not want to add the links or if you want to have the option of selecting which links in the list you want to add to the namespace.

4. To review the list of new items, select the policy to show the **Properties** tab on the right pane. Under New Resources, right-click an item and perform one of the following actions:
   - Click **Add as Link** to create a link and add it to the namespace.
   - Click **Ignore** if you do not want to reference the item from the namespace. The item is added to the policy Ignore list. The next time the policy is scanned, these items do not show up as new resources.

Review the list of new resources before adding links

You can set up a namespace creation policy so that you can always review scanning results and determine which items you want to reference from the namespace.

1. Perform one of the following actions:
   - If you are creating a new namespace creation policy, clear **Create links after scanning**.
   - If you want to change an existing policy, right-click the policy and click **Properties**. On the **General** tab, clear **Create links after scanning**.
2. When items are scanned, search results are listed under New Resources on the **Properties** tab of the policy. To display the **Properties** tab, select the policy and the tab appears on the right pane.
   - To reference an item from the namespace, right-click the item and click **Add as Link**.
   - If you do not want to reference an item in the namespace and you do not want to show the item again after the next scan, click Ignore. The item is added to the Ignore list for the policy. To view the list, right-click the policy, click **Properties**, then click **Ignore List**.

Delete a namespace creation policy

To delete a namespace creation policy, perform one of the following actions:

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

Populate the namespace based on permissions

Permission-based namespace policies enable you to populate a logical namespace based on the group permissions of shared folders and their subdirectories. Like namespace creation policies, permission-based namespace policies enable you to easily populate and maintain a logical namespace. You can use these policies, like namespace creation policies, 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.

This section provides instructions about how to configure and manage permission-based namespace policies.

NOTE
The VFM account must have permission to read the security descriptor of the shares and underlying folders you want the policy to scan.

Configure a permission-based namespace policy

Configure a permission-based namespace policy to automatically populate the logical namespace and to keep the logical namespace aligned with subsequent changes in the security permissions of underlying storage resources. Permission-based namespace policies enable you to group shares in the namespace based on security permissions. The VFM account must have permission to read the security descriptor of the shares and underlying folders you want the permission-based namespace policy to scan.

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Permission Based Policy**.
   - Right-click the Namespace Policies folder in Admin View, point to **New**, and click **Permission Based Policy**.
   - Right-click the Permission Based Policies folder and click **New Permission Based Policy**.
   - Right-click **Admin View**, click **New**, 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.
4. Clear **Automatically modify namespace after scanning** if you do not want to automatically create links for items that match the search criteria. After the policy runs, the items are listed under **New Resources** on the **Properties** tab that appears in the right pane when the policy is selected. You can create links for the items you want to include in the namespace.
5. Select one of the following:
   - Search for items not already targeted from this policy's folder when you do not want to target shares and subfolders that are already targeted by links in the folder. This is the default.
   - Search for items not already targeted from this root when you do not want to target shares and subfolders that are already targeted by links in the DFS root.
6. Select **Remove links whose shares no longer exist** when you want to remove links targeting shares when a scan detects that the shares no longer exist.
7. Select **Remove links from the entire root** when you want to scan the entire root and remove links targeting shares that no longer exist.
8. From the Use the Monitor on list, optionally select the machine where another monitoring agent is located.
9. Select whether you want to limit the number of links that are added every time the policy runs. Accept the limit specified or change the limit. When this option is cleared, links are added for all the candidates found. Microsoft recommends limits to the number of links that are added to a DFS root. For more information about recommended size limits, refer to the overview for DFS.
10. Click **Next**.

11. Select an existing DFS root or click **New** to create a new DFS root to contain the links, and click **Next**.
12. Group the links in a folder or add them under the root. Perform any of the following actions:
    - In the **Name** field, type a different name for the folder.
    - Click **New Folder** to create a new folder or subfolder.
    - Click **Delete** to remove the selected folder.
    - Click the root to add the links under the root.
13. Click **Next**.
14. Select the domain or domains that you want to search or select devices, and click **Next**.
15. Click **Add** to browse for security groups that have access to resources on the domain or servers you selected.
16. Type a word or pattern in the **Search string** field. You can use DOS wild-card characters (* or ?) in the search string. You can use the asterisk (*) as a substitute for zero or more characters in a name and the question mark (?) as a substitute for a single character in a name. 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.
17. Select a search type.
18. If you selected **Search shares matching the search string for folders matching the following path**, type a folder search string.

NOTE
To search for folders, the VFM account must have Read and List Folder Contents permissions for the shares and subdirectories that match the search criteria.

---

19. (Optional) select **Allow Partial Matches** if you want to search for partial matches. For more information about how partial matches works, refer to "Change the search type" on page 117.
20. (Optional) select to search hidden and administrative shares.
21. Click **Next**. A list of items that match your search criteria appears.
22. Clear the check box of items that you do not want to add, and click **Next**.
23. Select a template to arrange how links appear in the namespace. You to preview the namespace.

NOTE
If you selected domains to search for resources and the domains contain groups with the same name (for example, `domain1\sales` and `domain2\sales`), make sure that you select a template that sorts by domain, such as Domain-Group Name-Machine-Shares, Machines-Domain-Group Name-Shares, or Domain-Group Name-Shares. You cannot create links with the same name under the root or a logical folder of the namespace.

---

24. Click **Next**.
25. Optionally rearrange the links. You can move a link or folder, add a folder, add more links, remove a link or folder from the namespace and update the properties of a link.
26. Click **Next**.

27. Select **Use Schedule**.
28. Select how frequently you want to scan for items and click **Next**.
29. Click **Finish**. The new policy appears under the Permission Based Policies folder.

NOTES:

- If you uninstall a monitoring agent, ensure that the permission-based policies that formerly used the uninstalled agent are reassigned to a currently available agent. Right-click a permission-based policy and click **Properties**. On the **General** tab, select a monitoring agent from the **Use this Monitor on** list.
- Namespace creation policies enable you to populate the namespace with links that reference shares, volumes, and qtrees. For more information about how to create namespace creation policies, refer to "Configure a namespace creation policy" on page 114

Update the search criteria

To change security groups

1. Perform one of the following actions:
   - Right-click the permission-based namespace policy and click **Properties**.
   - Select a permission-based namespace policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Groups**.
3. Perform one of the following actions:
   - Click **Add** to include additional groups
   - Select a group and click **Remove** to delete the item from the list.
4. Click **OK**.

Other options

For more information about how to update the permission-based namespace policy, refer to the following namespace creation policy topics:

- "Update the schedule" on page 116
- "Add or remove servers or domains" on page 116
- "Add items to an exclusion list" on page 117
- "Change how the links are arranged" on page 118
- "Scan for items immediately" on page 118
- "Review the list of new resources before adding links" on page 119

Delete a permission-based namespace policy

- To delete a permission-based namespace policy, perform one of the following actions:
- Right-click the policy and click **Delete**.

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

Make the namespace highly available

Namespace availability policies enable you to keep the logical structure of two or more DFS roots and/or a DFS root and widelink namespaces aligned.

- By synchronizing 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.
- By synchronizing a DFS root and a widelink namespace, your CIFS clients can access the namespace hosted on a filer.
- A namespace availability policy provides flexibility in configuring your namespace to make it highly available. For more information, refer to "Namespace configuration options" on page 123.

This section provides instructions about how to create and manage policies to synchronize the logical structure of DFS roots and widelink namespaces to make them highly available. Before you create your first policy, refer to "Create a namespace availability policy checklist" on page 123 for more information.

NOTE
You can create recovery policies to make data highly available.

Create a namespace availability policy checklist

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

1. Make DFS roots highly available or create a widelink namespace.
2. Schedule synchronization.

Namespace configuration options

You have the flexibility of configuring your namespace in several ways to make it highly available.

- You can host the namespace on a filer.
  - You configure a namespace availability policy with a master DFS root and at least one secondary CIFS share on a NetApp filer where you will host a widelink namespace.
  - The logical structure of the DFS root is synchronized with the widelink namespace.
  - Your CIFS clients access the namespace on the filers.
- You can host the namespace on a filer and Windows systems, or only Windows systems. For more information, refer to the following illustration.
  - You configure a namespace availability policy with a master DFS root and at least one secondary CIFS share on a filer and/or secondary DFS root on a Windows system.
  - The logical structure of the master DFS root is synchronized with the secondary root and widelink namespace hosted on the secondary share.

- Your CIFS clients access the namespace on a filer and/or Windows system.

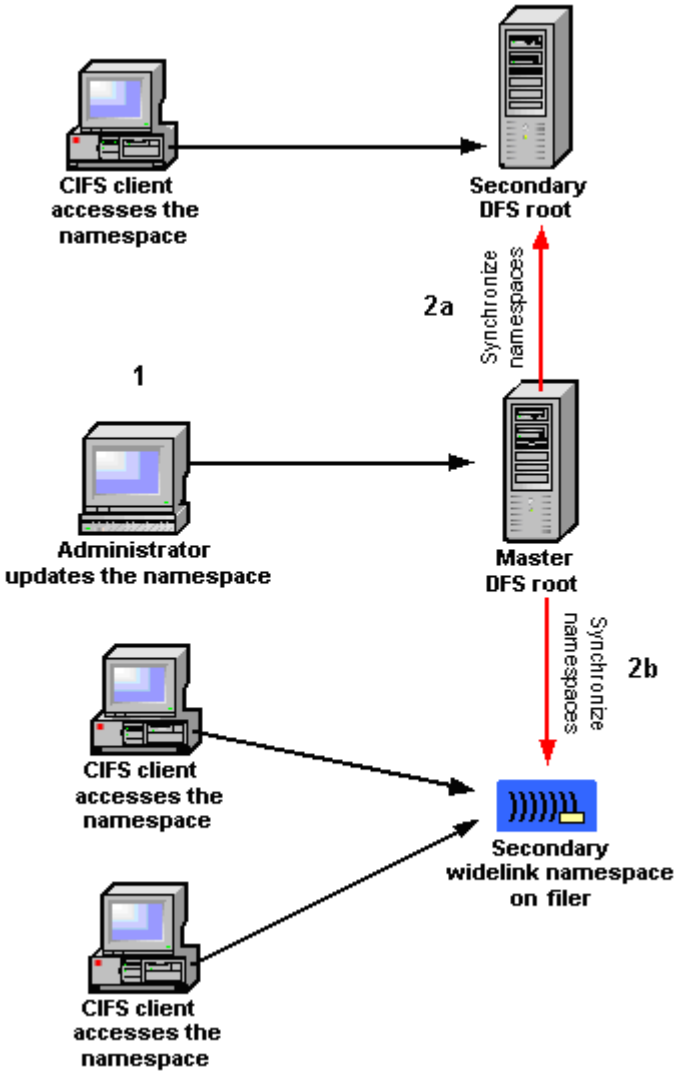

NOTE
Namespace availability policies do not support creating widelink namespaces on shares hosted on vFilers.

Make DFS roots highly available

Namespace availability policies enable you to synchronize the logical structure of two or more DFS roots

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, then click **Availability Policy**.
   - Right-click the **Namespace Policies** folder in the Admin View, point to **New**, and then click **Availability Policy**.
   - Right-click the **Availability Policies** folder and click **New Availability Policy**.

- Right-click **Admin View**, click **New**, 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 more roots. The policy must have at least two roots.
5. Select the master root. Secondary roots mirror the logical structure of the master root.
6. Select **Synchronize** after scanning.
7. Click **OK**.

NOTES:

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

Create a widelink namespace

Namespace availability policies enable you to create widelink namespaces on NetApp filers and synchronize the logical structure of a master DFS root with the widelink namespace. This enables CIFS clients to access the namespace directly through a filer instead of a Windows server.

NOTES:

- The filer where you want to create a widelink namespace must be running Data ONTAP 6.4 or later.
- Namespace availability policies do not support creating widelink namespaces on shares hosted on vFilers.

To create a widelink namespace

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, then click **Availability Policy**.
   - Right-click the **Namespace Policies** folder in the Admin View, point to **New**, and then click **Availability Policy**.
   - Right-click the **Availability Policies** folder and click **New Availability Policy**.
   - Right-click **Admin View**, click **New**, point to **Namespace Policy**, and click **Availability Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field (optional).
3. Click **Add** in the **Synchronize the following roots** field and type the path of the root, or click **Browse** to locate a root in your network.
4. Click **Add** in the **Synchronize the following NetApp filer shares** field and type the path of the share, or click **Browse** to locate the share in your network.
5. Select **Synchronize after scanning**.

6. Click **OK**. When the policy runs, a widelink namespace hosted on the CIFS share specified in the policy is created and the logical structure of the master DFS root is synchronized with the widelink namespace.

NOTE
For performance reasons and to avoid possible conflict with widelinks that might already exist under an existing share, it is recommended that you create and share a directory for hosting a widelink namespace instead of using an existing share. the contents of a share assigned to host a widelink namespace is scanned and all widelinks found on the share are managed.

Namespace synchronization options

You can choose whether the logical structure of the roots and widelink namespaces should be automatically synchronized each time a namespace availability policy runs, or whether to be informed when the logical structure of the roots and widelink namespaces are not the same.

To automatically synchronize at specified times

Make sure that the **Synchronize after scanning** properties option is selected for a policy then schedule synchronization.

To scan the roots for differences at specified times

Make sure that the **Synchronize after scanning** properties option is cleared for a policy and that you scheduled synchronization. The roots and widelink namespaces are scanned according to the schedule and you are provided the scanning results, but the logical structure of the roots and widelink namespaces is not modified.

To manually synchronize root structures

Use the **Start synchronization now** option to manually initiate the synchronization of the logical structure of the roots and widelink namespace in the policy. If you did not select the **Synchronize after scanning** properties option, you can synchronize after scanning is complete. Manually initiated scans or synchronizations are performed directly by the VFM server rather than by a monitoring agent to which the policy is assigned. This is because of the additional status information and user interaction provided by manually initiated scans.

Schedule synchronization

To make sure that changes made to the logical structure of the master root are synchronized with the secondary root and/or secondary widelink namespace, schedule how frequently you want to synchronize the logical structure.

1. Perform one of the following actions:
   - Right-click the namespace availability policy and click **Properties**.
   - Select a namespace availability policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Schedule Monitoring**.
3. Click **New** to configure a new schedule.
4. Optionally, select **Use alternate schedule when the policy fails** and type a frequency, in minutes, to run the schedule any time replication fails.

5. Click **OK**.

Start synchronization immediately

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

Manually initiated scans or synchronizations are performed directly by the VFM server rather than by a monitoring agent to which the policy is assigned. This is because of the additional status information and user interaction provided by manually initiated scans.

1. Perform one of the following actions:
   - Right-click a namespace availability policy and click **Start Synchronization Now**.
   - Click a namespace availability policy to show a **Properties** tab in the right pane. Right-click anywhere on the **Properties** tab and click **Start Synchronization Now**.
   - Click the **Availability Policies** folder to show a list of policies in the right pane. Right-click a policy and click **Start Synchronization Now**.

   The synchronization operation scans the logical structure of the roots and widelink namespaces in the policy and, if differences are detected between the master and secondary roots and widelink namespaces, changes are applied to the secondary roots and widelink namespaces to align their structure with that of the master root.
2. Expand the tree to view scanning results. If an error message appears, correct the problem before synchronizing. If no error messages appear, click **Synchronize**.

NOTES:

- This menu item is disabled if a synchronization operation is in progress.
- To view detailed information about the synchronization operation, click the policy to show an **Events** tab in the right pane. Double-click the **Synchronization completed successfully** event and expand the tree under **Details**.

Add additional roots or shares to a policy

A namespace availability policy must contain at least two DFS roots or a root and a widelink namespace. A policy can contain an unlimited number of roots and widelink namespaces.

Drag a root or CIFS share to a policy

You can add DFS roots and CIFS shares that will host widelink namespaces.

- Drag a root from the Logical View or right pane to a namespace availability policy.
- Drag a CIFS share on a filer from the Physical View tree or right pane to a namespace availability policy. The filer hosting the share must be running Data ONTAP 6.4 or later.

Add a root or share to a policy

1. Perform one of the following actions:
   - Right-click the namespace availability policy and click **Properties**.
   - Select a namespace availability policy and click **Edit** on the **Properties** tab that appears in the right pane.

2. Click **General**.
3. Click **Add** and type the path of a root, or click **Browse** to locate a root in your network.
4. 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 on a filer in your network.
5. Repeat Steps 3 and 4 to add more roots and shares.
6. Click **OK**.

NOTE
When you add a CIFS share to a namespace availability policy, the filer is added to the All Resources and NetApp Filer Resources folders in the Physical View.

Remove roots and widelink namespaces

If you no longer want a DFS root or widelink namespace to participate in a namespace availability policy, perform one of the following actions:

- Right-click a root in the Logical View tree and click **Remove from Namespace Availability Policy**.
- Select a namespace availability policy to display a **Properties** tab in the right pane. Right-click a root or widelink namespace on the **Properties** tab and click **Remove from Namespace Availability Policy**.
- Double-click a namespace availability policy, right-click a root or widelink namespace under the policy, and click **Remove from Namespace Availability Policy**.

Change the master root of a policy

To designate the master root for a synchronization policy, perform one of the following actions:

- Select a namespace availability policy to display a **Properties** tab in the right pane. On the **Properties** tab, select a root.
- Right-click a namespace availability policy and click **Properties**. On the General dialog box, select a root.

Delete a namespace availability policy

When you delete a namespace availability policy, namespace synchronization is no longer available. DFS roots and widelink namespaces that were members of the policy can still be accessed.

1. Perform one of the following actions:
   - Right-click a namespace availability policy and click **Delete**.
   - Select the **Availabilities Policies** folder to display a **Properties** tab in the right pane. Right-click a policy and click **Delete**.
2. Click **Yes**.

Create multiprotocol namespaces

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.

Create a multiprotocol namespace policy checklist

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

1. Create a multiprotocol namespace policy.
2. Host a UNIX namespace on a filer.
3. Configure UNIX clients to access the namespace.
4. Select default UNIX settings.
5. Add a sync point.
6. Schedule synchronization

Create a multiprotocol namespace policy

Create a multiprotocol namespace for CIFS and NFS clients. You can automatically synchronize the namespaces.

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Multiprotocol Policy**.
   - Right-click the **Namespace Policies** folder in Admin View, point to **New**, and click **Multiprotocol Policy.**
   - Right-click the **Multiprotocol Policies** folder and click **New Multiprotocol Policy**.
   - Right-click **Admin View**, click **New**, point to Namespace Policy, and click **Multiprotocol Policy**.
2. In the **Name** field, type a word or phrase that best describes this policy. The name appears under the Multiprotocol Policies folder and on the **Properties** tab of items that are members of the policy.
3. Type the name of a DFS root that contains links that target data on NetApp filers, or click **Browse** to locate a root in your namespace.
4. Perform one of the following actions:
   - If synchronization is from a UNIX namespace to a DFS namespace, type a CIFS UNC path name for the directory that contains the automount files describing the UNIX namespace.
   - If synchronization is from a DFS namespace to a UNIX namespace, type a CIFS UNC path name for the directory that will contain files created by VFM to configure the UNIX namespace.
   - Click **Browse** to locate the directory.

NOTES:

- The CIFS UNC path name you type in the **UNIX Files Location** text box must be accessible to the VFM server.
- The multiprotocol namespace enables users to access directories hosted on NetApp filers that are accessible through both NFS exports and CIFS shares.

5. From the **Use the Monitor on** list, optionally select the machine where another monitoring agent is located.
6. (Optional) Click **Select default share security** to select default share access permissions for all shares referenced by links contained the DFS namespace.
   - Click **Add** to select the users and groups to whom you want to assign permissions.
   - Select the permissions you want to allow each user or group.
   - Click **OK**.
7. Click **OK**.

NOTES:

- Direct and indirect automount maps are supported. Auto_master files are not supported.
- The DFS and UNIX namespaces in a multiprotocol namespace policy must reference data on NetApp filers.
- VFM does not support the Solaris-specific automount options, such as running executables and evaluating environment variables when determining the actual export to mount.

Host a UNIX namespace on a filer

Automatically copy the contents of the export files it creates to the */etc/exports* file and to run the command `exportfs -a` to make the new exports visible to NFS clients.

1. Perform one of the following actions:
   - Right-click a multiprotocol policy and click **Properties**.
   - Select a multiprotocol policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Scripts**.
3. Select **Automatically apply exports**.
4. Click **OK**.

Configure UNIX clients to access the namespace

You can automatically copy automount files to NIS.

1. Perform one of the following actions:
   - Right-click a multiprotocol policy and click **Properties**.
   - Select a multiprotocol policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Scripts**.
3. Select **Push automount files to NIS**.

4. Type the name of an NIS client machine.
5. Type the name of a directory on the NIS machine where the automount files are temporarily copied.
6. Type the path on the NIS machine where the make program is located.
7. Click **OK**.

NOTES:

- Before you run the multiprotocol namespace policy, you must add the NIS server to the Physical View. For more information, refer to "Add a machine to the physical view" on page 154.
- NFS clients that are making use of the automount files must have the directory that is specified in the automount files.

Select default UNIX settings

Set options to control NFS client access to the data.

1. Perform one of the following actions:
   - Right-click a multiprotocol namespace policy and click **Properties**.
   - Select the Multiprotocol Namespace Policies folder to display a **Properties** tab in the right pane, then click **Edit**.
2. Click **Default UNIX Settings**. Select how you want to set up access restrictions. You can select from several export settings options.

   TABLE 23 UNIX export setting options

To...	Then do this...
Allow only specific client machines access to the data	Select **Restricted Access**. In the **Allow Access** field, type the names of the machines, subnets, or netgroups that are allowed access.
Allow read-only access to the data to all machines	Select **Read-Only**.
Allow read-write access to specific client machines and read-only to all other machines	Select **Read-Only**. In the **Allow Write Access** field, type the name of the client machines, subnets, or netgroups that are allowed read-write access.
Prevent root access to the data	Select **No Root Access**.
Allow root access to the data to all clients	Select **All Root Access**.
Allow root access to the data to only specific machines	Select **No Root Access**. In the **Always allow Root Access** field, type the name of the client machines, subnets, or netgroups that always are allowed root access.

3. (Optional) Change how entries specified by the automount map are mounted. The following list specifies the VFM default mount options:
   - Mount resources read-write.
   - Allow set user ID execution.

- Soft mount, return an error if the server does not respond.
- Do not allow keyboard interrupts of a process waiting on a hard mount.

4. Click **OK**.

Add a sync point

Add synchronization points between the master and secondary multiprotocol namespaces. The secondary namespace is updated when the logical structure of the master namespace changes.

1. Perform one of the following actions:
   - Right-click a multiprotocol namespace policy and click **Properties**.
   - Select a multiprotocol namespace policy to show a **Properties** tab in the right pane and click **Edit**.
2. In the General dialog box, click **New**.
3. Type a word or phrase that describes the data synchronized between the DFS and UNIX namespaces.
4. Indicates the direction of the synchronization of the logical structure of the namespace.
   - If DFS is selected, whenever the links in the DFS Logical Folder change, the automount file and corresponding export files are updated.
   - If UNIX is selected, whenever the contents of the automount file changes, the links in the DFS Logical Folder are updated.
5. Type the name of the DFS namespace folder that contains links targeting data on NetApp filers or click **Browse** to select the folder contained in the DFS root you specified in the **DFS Root** field.
6. In the **Automount File** field, perform one of the following actions:
   - If DFS is the master, type a name for the automount file that will be create.
   - If UNIX is the master, 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.
7. Click **OK**. The new sync point appears in the **Sync Points** list.
8. Repeat Steps 2 through 7 to create additional synchronization points.
9. Click **OK**.

NOTE
The DFS and UNIX namespaces in a multiprotocol namespace policy must reference data on NetApp filers.

Schedule synchronization

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. Perform one of the following actions:
   - Right-click the multiprotocol policy and click **Properties**.

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

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

Run customized commands during synchronization

Specify batch files to run before and after namespace synchronization.

1. Perform one of the following actions:
   - Right-click a multiprotocol namespace policy and click **Properties**.
   - Select a multiprotocol namespace policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Scripts**.
3. Select **Run Batch File Before Sync** and type the path where the script file is located, or click **Browse** to locate the file in your network.
4. Select **Run Batch File After Sync** and type the path where the script file is located, or click **Browse** to locate the file in your network.
5. Click **OK**.

NOTE
For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Update a multiprotocol namespace policy

Make changes to a multiprotocol namespace policy.

1. Perform one of the following actions:
   - Right-click a multiprotocol namespace policy and click **Properties**.
   - Select the Multiprotocol Namespace Policies folder to display a **Properties** tab in the right pane, then click **Edit**.
2. Change the appropriate information.
3. Click **OK**.

Start namespace synchronization immediately

You can bypass the synchronization schedule and manually start the namespace synchronization process any time you want.

1. Perform one of the following actions:
   - Right-click a multiprotocol namespace policy and click **Start policy now**.
   - Select the Multiprotocol Policies folder to show a **Properties** tab in the right pane. Right-click a policy and click **Start policy now**.

2. Click **Done**.

Delete a sync point

When you delete a sync point from a multiprotocol namespace policy, synchronization between the master and secondary resources is canceled. UNIX and DFS clients continue to access their respective namespaces as they are currently configured.

1. Perform one of the following actions:
   - Right-click the policy and click **Properties**.
   - Select the policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **General**.
3. Select the sync point you want to delete and click **Remove**.
4. Click **OK**.

Delete a multiprotocol namespace policy

When you delete a multiprotocol namespace policy, both DFS and UNIX clients can still access their respective namespaces as they are currently configured. Namespace synchronization is no longer available.

1. Perform one of the following actions:
   - Right-click a multiprotocol namespace policy and click **Delete**.
   - Select the Multiprotocol Policies folder to display a **Properties** tab in the right pane. Right-click the policy and click **Delete**.
2. Click **Yes**.

Virtualize user home directories

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

Create a home folder policy

When users access their home directories from the namespace, you can load balance and migrate user data without affecting users. VFM enables you to group users in home folders to accommodate users from different domains and to place home directories at different points in the namespace. You can also organize users' home directories as subfolders to facilitate the creation of personal namespaces for users.

Before you begin, review DFS size limitations on the number of links contained in a DFS root.

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Home Folder Policy**.

- Right-click the **Namespace Policies** folder in the Admin View, point to **New**, and then click **Home Folder Policy**.
- Right-click the **Home Folder Policies** folder and click **New Home Folder Policy**.
- Right-click **Admin View**, click **New**, point to **Namespace Policy**, and click **Home Folder Policy**.

2. Click **Next**.
3. Type a word or phrase that best describes this policy in the **Name** field. The name appears under **Home Folder Policies** and on the **Properties** tab of items that are members of the policy.
4. In the **Domain** field, type the name of the domain that contains the home directories, or click **Browse** to select a domain.
5. In the **DFS Root** field, type the name of the root that you want to add users' home directories, or click **Browse** to select a root.
6. In the **Link Template** fields, optionally type the name of folders where you want to group the home directories, and/or type the common name for the link that targets the users' home directory. If you select not to group users in folders, the users appear alphabetically by name in the Logical View under the DFS root selected.
7. Click **Next**.
8. Select the users you want to add to the namespace.
9. Select **Show users without a home folder** if you want to add users who do not have a home directory.
10. Click **Next**.
11. Click **Finish**.

Add a link to a user home folder

Customize user access to data.

Assumption

The home folder policy that contains the users was created to show user home directories in subfolders.

To add a link to a user home folder

1. Right-click a user home folder and click **Add link**.
2. Type the name you want to give the link in the **User connect to** field.
3. Click **Add** and type the name of a shared folder or directory, or click **Browse** to display a list of domains from which you can select a share or subfolder.
4. Click OK. The new data is accessed as a subfolder on the user's home folder.

To add a common link to multiple home folders

This organization can be useful, for example, when you need to give a set of users access to a directory containing information about a project on which they are collaborating.

1. Select the root that contains the user home folders to display a **Properties** tab in the right pane.
2. Select the user home folders listed under **Links**.
3. Right-click and click **Add link**.
4. Type the name you want to give the link in the **User connect to** field.
5. Click **Add** and type the name of a shared folder or directory, or click **Browse** to display a list of domains from which you can select a share or subfolder.
6. Click **OK**. The new data is accessed as a subfolder on the selected users' home folders.

NOTE
You can update identical links in one easy step. When you rename, add, remove targets, or change the timeout of a link that appears in several places in the namespace, a list of the identical links displays and you can update all the links simultaneously.

Add new users to a policy

Add new users to an existing home folder policy. If the users have home directories, they are added automatically to the namespace.

1. Perform one of the following actions:
   - Right-click the home folder policy to which you want to add users and click **Add Users**.
   - Right-click anywhere on the **Properties** tab of a home folder policy and click **Add Users**.
2. Select the users you want to add.
3. Select **Show users without a home** folder if you want to add users who do not have a home directory.
4. Click **Next**.
5. Click **Finish**.

View properties of a home folder policy

View the properties of a home folder policy.

1. Open **Home Folder Management** in the Admin View.
2. Perform one of the following actions:
   - Click a home folder policy. A **Properties** tab appears in the right pane. The list of users also appears in this tab.
   - Right-click a home folder policy and click **Properties**.

Refresh a home folder policy

Update the information for users in a home folder policy to reflect changes made in the domain.

Perform one of the following actions:

- Under the **Home Folder Policies** folder, right-click the home folder policy you want to update and click Refresh.

- Right-click anywhere on the **Properties** tab of the home folder policy and click **Refresh**.

NOTE
Refreshing the policy reevaluates how users named in the policy access their home folder but does not make any changes to the user accounts or to the links in the namespace.

Resolve discrepancies

Update user information and apply the rules you set for a home folder policy.

Discrepancies that VFM resolves

- Users on the list with home directories but without a DFS link - VFM adds links for the users to the namespace.
- Users who were deleted from the domain - VFM removes the users from the namespace but keeps their name on the list.
- Users originally added to the list that did not have home directories - VFM checks the domain, updates information for each user, and applies rules.

To resolve a discrepancy
Perform one of the following actions:

- Right-click anywhere on the **Properties** tab of a home folder policy and click **Apply Rules**.
- Right-click a home folder policy and click **Apply Rules**.

Remove a user from a home folder policy

When you remove a user from a home folder policy, the user is only removed from the policy and the namespace. The user's data is not removed from the domain. The home directory setting in the user's profile is updated so that the user can access his home directory share.

1. Open **Home Folder Management** in the Admin View.
2. Perform one of the following actions:
   - Click a home folder policy. Under **Users** on the **Properties** tab that appears in the right pane, right-click a user and click **Remove**.
   - Right-click a home folder policy and click **Remove Users**. Select the users you want to remove. Click **Next** and click **Finish**.

Move a user

Move a user's share to another location and make this move transparent to the user. Before you move a user's share, refer to "Data migration considerations" on page 14 for more information.

1. Select a home folder policy to display the **Properties** tab in the right pane.
2. Under **Users**, right-click the user you want to move and click **Data migration**.
3. Use the Data Migration wizard to move the user's data and retarget the link.

Locate a user in the namespace

Find where a user is located in the namespace.

1. Click the home folder policy that contains the user. A **Properties** tab appears in the right pane.
2. Find the user in the list and click the link under DFS Link.

Locate a DFS root in the namespace

Find where the DFS root of a home folder policy is located in the namespace.

1. Click the home folder policy that contains the user. A **Properties** tab appears in the right pane.
2. Click the link next to Base DFS Path.

NOTE
You can also locate a DFS root in the namespace by clicking **Logical View**. A **Properties** tab appears in the right pane that shows a list of all the roots in the namespace. Double-click the root you want to locate.

Delete a home folder policy

When you delete a home folder policy, all the user links are deleted from the namespace. VFM changes the users' account information so that their home folders point directly to the targets that used to be pointed by links.

1. Right-click the policy you want to delete and click **Delete**.
2. Click **Yes**.

Back up the namespace

Create a namespace backup policy to back up the logical structure of DFS roots so that you can restore a root to a known state in the event that it is deleted or changes to the root need to be rolled back.

Backup

- Only the links of the DFS roots selected are backed up. Data referenced by links in the namespace and empty logical folders are not backed up.
- Permissions on logical folders are not backed up.
- The backup operation records the configuration data that was gathered about the namespace.
- If you have cascaded roots, you must back up each root separately.

Restore

- Recreates the links contained in a backed-up DFS root onto a DFS root.
- Restores selected items in a backup file.

- You might encounter errors if you attempt to restore a backup configuration to a different type of root than it was originally created on. This can occur if you attempt to restore the backup set for a Windows 2000 DFS root to a Windows NT 4 DFS root. You can skip the problem steps in the confirmation dialog boxes and restore those aspects of the backup that apply to Windows NT 4 DFS roots.

NOTE
When you export your VFM configuration, all your DFS root backups are included with the export data. This enables you to restore a root that might no longer be available after you import the VFM configuration.

Namespace backup policy checklist

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

1. Create a namespace backup policy.
2. Schedule backups.

Create a namespace backup policy

Back up automatically the logical structure of your namespace so that you can quickly restore it if it becomes necessary.

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Namespace Policy**, and click **Backup Policy**.
   - Right-click the **Namespace Policy** folder in the Admin View, point to **New**, then click **Backup Policy**.
   - Right-click **Admin View**, click **New**, point to **Namespace Policy**, and click **Backup Policy**.
   - Select the **Namespace Policies** folder to show a **Properties** tab in the right pane. Right-click the **Backup Policies** folder and click **New 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 locate a root in your namespace.
4. Repeat Step 3 to add more roots.
5. Click **OK**.

NOTES:

- When the DFS root you are backing up and the VFM server are hosted on machines running Windows Server 2003 release 2 or later, the attributes of the root are backed up.

  DFS root attributes restored:

  - Link target failover priority
  - Link target costing
  - DFS root permission settings
  - Ability to enable or disable root scalability mode
  - Ability to configure site costing
- You can also manually back up a root. On the Tools menu, click Manual backup.

Schedule backups

Schedule how frequently you want to back up the roots participating in a policy.

1. Perform one of the following actions:
   - Right-click a namespace backup policy and click **Properties**.
   - Select a namespace backup policy to show a **Properties** tab in the right pane and click **Edit**.
2. Click **Schedule**.
3. Select **Use Schedule**.
4. Type the time and date you want to begin making backups
5. Select how frequently you want to make backups, and type how many backups to keep.
6. Click **OK**.

NOTE
If you enable multiple schedules (for example, hourly, nightly, and weekly backups) and backups initiate at the same time, only the longer period backup is initiated. For example, if you schedule hourly and nightly backups that initiate at the same time, only the nightly backup is initiated. Multiple simultaneous backups are not initiated to avoid duplicate network traffic and the maintenance on duplicate backups.

Browse contents of a backup

Quickly view the contents of a backup.

1. On the Tools menu, click **Browse Backups**. A list of backup policies appears.
2. Double-click the policy that contains the backup file. Daily, Hourly, and Weekly folders appear.
3. Double-click the folder that contains the backup file. A list of backup files appears.
4. Double-click the backup file you want to view. 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**.

Remove a root from a namespace backup policy

Remove a DFS root from selected namespace backup policies.

1. Right-click the root and click **Remove from Namespace Backup Policy**.
2. Select the policies from which you want to remove the root.
3. Click **OK**.

NOTE
If the root is the only root participating in a policy, the root is not removed from the policy. If you no longer want the policy to back up the root, delete the policy.

Delete a namespace backup policy

When you delete a namespace backup policy, all the backup files associated with the policy are deleted.

1. Perform one of the following actions:
   - Right-click a namespace backup policy and click **Delete**.
   - Select the **Backup Polices** folder to show a **Properties** tab in the right pane. Right-click the policy you want to delete, and click **Delete**.
   - Select a namespace backup policy to show a **Properties** tab in the right pane. Right-click anywhere on the page, and click **Delete**.
2. Click **Yes**.

To delete a root and retain the namespace backups for the root

To help clarify the following instructions, *TempRoot* is the name of a root you add temporarily to the policy. *DeleteRoot* is the root you want to delete, and *NewRoot* is the root that you use to restore the contents of *DeleteRoot*.

1. Add *TempRoot* to the namespace backup policy that contains the root you want to delete.
2. Remove *DeleteRoot* from the policy, and delete the root.
3. Add *NewRoot* to the policy.
4. Restore the contents of *DeleteRoot* to *NewRoot*.
5. Remove *TempRoot* from the namespace backup policy.

Restore a DFS root

Restore the logical structure of a DFS root that you previously backed up. The DFS links of the root are restored.

1. Perform one of the following actions:
   - On the **Tools** menu, click **Restore from backup**.
   - Right-click **Logical View** and click **Restore**.
   - Right-click a DFS root and click **Restore**.
2. Click **Next**. A list of all the roots in your namespace appears.
3. Select the root to which you will restore the links from a backup file.
4. Click **Next**.
5. Type the path of a backup file, or click **Browse** to find the location of the file.
6. Click **Next**. If the backup file contains the logical structure of more than one root, a list of the roots appears. Select the root you want to restore.
7. Click **Next**. The target DFS root is compared with the backup root and shows the links that are different and the action that will be taken for each link.
8. Clear the links that you do not want to restore.
9. (Optional) Click **View Backup** for a detailed view of each item in the backup root.

10. (Optional) Select **Remove existing links** to remove links in the target root that are not in the backup root.
11. (Optional) Select **Restore DFS Root Properties** to restore the attributes of roots. The roots and the VFM server must be hosted on machines running Windows Server 2003 release 2 or later.

    DFS root attributes restored:

    - Link target failover priority
    - Link target costing
    - DFS root permission settings
    - Ability to enable or disable root scalability mode
    - Ability to configure site costing
12. Click **Next**. A window displays the name of the backup file, the name of the DFS root you selected from your backup file, and the name of the DFS root you selected as the target of the restore process.
13. Click **Finish**.

Chapter

Managing File Servers

6

This section provides instructions about how to use VFM to manage physical resources that participate in the namespace or in VFM policies.

In this chapter

- Working with NetApp filers 143
- Setting replication properties 150
- Working with physical resources 152

Working with NetApp filers

VFM enables you to integrate your NetApp filers into your namespace.

- NetApp filers in the Physical View tree
  - Integrated management of NetApp filers - VFM provides integrated in-context access to the FilerView® management tool. When a NetApp filer is selected in the Physical View, a **Manage** tab hosting the FilerView UI is added to the right pane.
  - Display of NetApp-specific information - VFM displays information from the NetApp filer, including Data ONTAP version, filer volume names, filer system status information, and NETBIOS aliases in use on the filer.
- NetApp vFiler in the Physical View tree
  - Unlike a NetApp filer, a vFiler does not display a **Manage** tab in the right pane when it is selected in the Physical View. vFilers are managed through their hosting filer.
  - The platform information for a vFiler is Data ONTAP without any revision number. This is because vFilers do not support SNMP.
- If VFM misidentifies a filer or vFiler, right-click the machine, click **Properties**, and select the correct machine identity setting.
- If filers are misidentified as either generic machines or as vFilers, or if FilerView does not appear in the **Manage** tab, refer to "Finding solutions" on page 367 for more information.

Set SNMP community string

VFM uses the SNMP community string in the process of identifying NetApp filers and UNIX machines and uses `public` as the default value. You can set the value for a specific machine.

1. In the Physical View tree, right-click a machine and click **Properties**.
2. On the **General** tab, type a name in the **SNMP community string** field.
3. Click **OK**.

NOTE
To set application-wide defaults, click **Options** on the **Tools** menu and click **System Options**.

Change the shell credentials of a filer

To manage filers that are members of a policy, you must configure VFM with the system or administrative user account and password of NetApp filers. VFM normally uses RSH to communicate with filers. You can choose to use SSH.

NOTES:

- **RSH**

  When the user name and password are left blank, VFM uses `root` as the user name with no password. You can specify different credentials for a filer or UNIX machine. You can also change the default RSH shell credentials that all new filers and UNIX machines added to VFM will use.

- **SSH**

  Before you can use SSH, you must configure VFM to use the remote shell. For more information, refer to "Using SSH" on page 343, which provides guidance about how to perform the following tasks:

  - Install and configure OpenSSH.
  - Generate a key-pair.
  - Copy the public key to the filers using SSH.
  - Cache the keys.
  - Provide the location of the `plink` and `pscp` executables.

- To set default RSH or SSH shell credentials that all new filers and UNIX machines will use, refer to "Configuring default remote shell credentials" on page 355. Machines already known to VFM are not affected.

To change the shell credentials of a filer

1. Right-click the filer either in the Physical View or beneath the policy in which it is contained, and then click **Properties**.
2. Click the **Shell** tab.
3. Select a type in the **Shell Type** list.
4. Perform one of the following actions:
   - *If you select RSH,* type the username and password that VFM will use to communicate with the filer.
   - *If you select SSH,* type the location of the private key and the user name that VFM will use to communicate with the filer. When VFM uses SSH to communicate with UNIX machines and filers, SSH is configured without a password or passphrase.
5. Click **OK**.
6. Right-click the filer in the Physical View and click **Refresh Filer Information**.

Specify the ONTAPI credentials of a filer

VFM uses ONTAPI to communicate with filers. Normally, VFM uses the DCE/RPC protocol. If you want to use HTTP to communicate with a filer or the filer is not running CIFS, you must update the ONTAPI credential information for the filer. If the filer is not known to VFM, add the filer to the Physical View tree before you complete the following instructions.

1. Right-click the filer in the Physical View and click **Properties**.
2. Click the **NetApp Filer** tab.
3. In the **Connection Type** list, select **HTTP**.
4. Type the credentials that VFM will use to communicate with the filer.
5. Click **OK**.

NOTE
When the DCE/RPC protocol is used to communicate with the filer, the credentials of the account used to run the VFM server and client are used.

Create a report of multiple filers

Create NetApp-specific reports containing information about multiple filers.

1. Perform one of the following actions:
   - On the **Reports** menu, point to **NetApp** then click a report.
   - Right-click **Reports** or **Physical View**, point to **NetApp**, and click a report.
   - Right-click the Namespace Resources folder under the Physical View tree or on the **Properties** tab of Physical View, point to **NetApp**, and click a report.
2. Select any or all filers.
3. Click **Next**.
4. Optionally, type a name for the report.
5. Click **Finish**.

Enable or disable all SnapMirror backups

You can enable or disable all the SnapMirror® backups on a filer.

1. Right-click a filer and click **Properties**.
2. Click the **NetApp Filer** tab.
3. Perform one of the following actions:
   - To enable all the SnapMirrors, select **Enable SnapMirrors**.
   - To disable all the SnapMirrors, clear **Disable SnapMirrors**.
4. Click **OK**.
5. Click **Done**.

Filers not identified correctly

You can change a machine setting if VFM does not correctly identify a NetApp filer or vFiler in the Physical View.

1. Right-click the machine and click **Properties**.
2. Select the appropriate machine identity.
3. Click **OK**. When a filer is correctly identified, additional properties tabs appear.

Create a SnapMirror on a filer

Configure a SnapMirror on a NetApp filer. To manage filers, you must configure VFM with the system or administrative user account and password of the filer. You must have write access to `C$\etc` on the filer. For more information, refer to "Disaster recovery planning considerations" on page 20.

1. Right-click the filer in the Physical View and click **Properties**.
2. Click the **SnapMirrors** tab.
3. Click **Create**.
4. Select the source and destination.
5. Click the **Options and Schedule** tab and set up the replication schedule.
6. Click **OK**.
7. Click **Done**.

Create a volume on a filer

Configure a traditional volume on a NetApp filer.

1. Perform one of the following actions:
   - Right-click a filer and click **Create Volume**.
   - Right-click the Volumes folder under a filer in the Physical View tree and click **Create Volume**.
   - Select a filer in the Physical View tree to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Create Volume**.
   - Select the Volume folder under a filer in the Physical View tree to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Create Volume**.
2. In the **New Volume Name** field, type the name of the volume. The name can contain alphanumeric characters and underscores (_). The first character must be a letter or an underscore.
3. In the **RAID Group Size field**, type the number of disks per RAID group in the volume. The maximum value is 28, and the default value is 14.
4. Select a language from the list provided.
5. Select one of the following options:

- **Automatic Disk Selection** enables the system to decide which disks to add, based on your response to the **Number of Disks** to add to the volume and the size of the disks to be added.
- **Manual Disk Selection** enables you to select the disks to add to the new volume.

6. Click **Create**.

NOTE
To manage filers, you must configure VFM with the system or administrative user RSH account and password of the filer.

Create a qtree on a filer

Configure a qtree on a NetApp filer. To manage filers, you must configure VFM with the system or administrative user RSH account and password of the filer. For more information, refer to "Change the shell credentials of a filer" on page 144

1. Right-click a filer in the Physical View and click **Properties**.
2. Click the **QTrees** tab and click **Create**.
3. Select the filer.
4. Select the volume on the filer in which you want to create the qtree.
5. Type a name for the qtree.
6. Select the appropriate security style.
7. Select to enable or disable `oplocks`.
8. Click **OK**.

Enable vFilers for filer recovery policy

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

1. If the vFiler does not appear in a folder in the Physical View, add the vFiler to a folder.
2. Right-click the vFiler and click **Properties**.
3. Click the **NetApp Filer** tab.
4. Type the name of the filer hosting the vFiler in **Hosting Filer**, or select the filer from the list provided.
5. Click **OK**.
6. Click **Done**.

Schedule snapshots on multiple filers

This is a convenient way to change the snapshot schedule and reserve setting on multiple filers without requiring you to repeat the changes on each individual filer.

To manage filers, you must configure VFM with the system or administrative user account and password of the filer. You must also have write access to `C$\etc` on the filer.

NOTE
You are warned if some of the volumes on the selected filers are members of other snapshot scheduling policies. When a snapshot scheduling policy manages filer-initiated snapshots and the snapshot schedule of a volume is changed outside the policy, VFM restores the schedule to the policy settings the next time the policy schedule runs.

To schedule snapshots on multiple filers

1. Perform one of the following actions:
   - Right-click a folder in the Physical View and click **Set snapshot schedule**.
   - Select a folder in the Physical View to show a **Properties** tab in the right pane. Select multiple filers on the **Properties** tab, right-click and click **Set snapshot schedule**.
2. Perform one of the following actions:

TABLE 24 Setting a snapshot schedule

If...	Then do this...
You want to manage the snapshots, select **Add this volume to a Snapshot Scheduling** policy.	When you select an existing policy Click **Next**. The name of the policy and a list of the volumes to be added to the policy appears. Click **Finish**. The volumes on the selected filers are added to the policy. When you click **New Policy** to create a new policy Type a word or phrase that best describes the policy in the **Name** field. Click the **Virtual File Manager Settings** or **Filers Settings** tab and set the schedule for the policy. Click **OK**. A status report appears. Click **Done**. The new policy appears on the list of snapshot scheduling policies. Click **Next**. The name of the policy and a list of the volumes to be added to the policy appears. Click **Finish**. The volumes on the selected filers are added to the new policy.
You want the volumes managed by their respective filers, select **Set the schedule for this volume**, but don't add it to a policy.	Click **Next**. Normally, **Set the snapshot schedule** and **Enable scheduled snapshots** options are selected. Type the number of weekly, daily, and hourly snapshots to retain. Select the time you want to initiate snapshots on the volumes and click **Next**. Perform one of the following actions: Select **Don't modify the snapshot reserve** if you want to retain the current reserve on the selected volumes. Select **Set the snapshot reserve** if you want all the selected volumes to have the same reserve. Type the amount of space on the volumes to reserve for snapshots. Click **Next**. A list of the changes to be made appears. Click **Finish**.

Enable or suspend snapshot schedules for multiple filers

You can stop the snapshot schedule from running on volumes on multiple filers. The schedules are not deleted.

1. Perform one of the following actions:
   - Right-click a folder in the Physical View and click **Set snapshot schedule**.

- Select a folder in the Physical View to show a **Properties** tab in the right pane. Select multiple filers on the **Properties** tab, and right-click and click **Set snapshot schedule**.

2. Click **Next**. The option Set the schedule for this volume, but don't add it to a policy is selected.
3. Click **Next**.
4. Select or clear the **Enable scheduled snapshots** check box, depending on whether you want to enable or suspend the snapshot schedule for the selected items and click **Next**.
5. Click **Next**. A list of the changes to be made appears.
6. Click **Finish**.

NOTES:

- Suspending a snapshot schedule does not prevent you from manually taking a snapshot. Select a folder in the Physical View that contains filers to show a **Properties** tab in the right pane. Select the filers with volumes you want to take a snapshot of, and right-click and click Take snapshot now.
- Existing snapshots on the selected volumes are not deleted when you suspend the schedule.

Take a snapshot of a volume immediately

Manually take a snapshot of NetApp filers.

1. Perform one of the following actions:
   - To take a snapshot of all the volumes on a NetApp filer, right-click the filer and click **Take snapshot now**.
   - To take a snapshot of selected volumes on a NetApp filer, select the filer to show a **Properties** tab in the right pane. Under **Volumes**, select the volumes, right-click, and click **Take snapshot now**.
   - To take a snapshot of the volumes on multiple NetApp filers, right-click the folder in the Physical View that contains filers and click **Take snapshot now**.
2. Type a name for the snapshot.
3. Optionally, select **Run batch before snapshot** and type the UNC path of a script file that you want to run.
4. Optionally, select **Run batch after snapshot** and type the UNC path of a script file that you want to run.
5. Click **OK**.

Change the snapshot reserve for multiple volumes

You can change the snapshot reserve for volumes on multiple filers.

1. Perform one of the following actions:
   - To change the snapshot reserve for all the volumes on a NetApp filer, right-click the filer and click **Set snapshot schedule**.

- To change the snapshot reserve for selected volumes on a NetApp filer, select the filer to show a **Properties** tab in the right pane. Under **Volumes**, select the volumes, right-click, and click **Set snapshot schedule**.
- To change the snapshot reserve for the volumes on multiple NetApp filers, right-click the folder in the Physical View that contains filers and click **Set snapshot schedule**.

2. Click **Next**.
3. Make sure that the option **Set the schedule for this volume**, but do not add it to this policy is selected. Click **Next**.
4. Select **Don't modify the snapshot schedule** and click **Next**.
5. Make sure that the option **Set the snapshot reserve** is selected. Type the amount of space the volumes must reserve for snapshots.
6. Click **Next**. Verify the changes that are to be made to the volumes.
7. Click **Finish**.

Restore a snapshot

Restore snapshot data to its original location on a NetApp filer.

1. Select a filer or volume in the Physical View to show tabs in the right pane.
2. Click the **Snapshots** tab.
3. Right-click a snapshot and click **Restore**.
4. Click **Yes**.

Setting replication properties

This section provides instructions about how to set replication properties on machines participating in the namespace or a storage policy.

- To specify a Windows or UNIX machine as a proxy to host a replication agent instead of having the agent deployed to a Windows or UNIX machine participating in a policy, refer to "Specify a proxy machine" on page 151
- To prevent deployment of a replication agent to a machine, refer to "Disallow deployment of replication agents" on page 150
- To specify application-wide default Windows and UNIX replication agents, refer to "Specify default replication agents" on page 151
- To set bandwidth throttling for a machine hosting a replication agent, refer to "Set data transfer rate limit" on page 152

Disallow deployment of replication agents

Replication agents are installed on Windows or UNIX machines participating in policies that transfer CIFS or NFS data. You can prevent deployment of agents to individual machines.

1. On the Physical View tree, right-click a machine and click **Properties**.
2. Click the **Agent** tab.

3. Clear **Allow replication agent to be deployed to this host**.
4. Click **OK**.

NOTE
For more information about replication agents, such as system requirements and instructions about how to manually deploy agents, refer to "Managing Agents and the VFM Server" on page 299.

Specify a proxy machine

VFM deploys Windows and UNIX replication agents to machines participating in policies that transfer data. You can specify a proxy machine to host the replication agent instead of having the agent deployed to a machine participating in the policy.

NOTE
If the machine to which you want to deploy a replication agent is running Windows XP SP 2, you must disable the firewall on the machine to which you want to deploy the replication agent, and then restart the machine. For more information about Windows and UNIX replication agent requirements, refer to "System requirements" on page 34.

1. In the Physical View tree, right-click a machine and click **Properties**.
2. Click the **Agent** tab.
3. Perform one of the following actions:
   - Select **Use default proxy** to use the machine you specified as the application-wide default system option.

     NOTE
     To set application-wide default Windows and UNIX replication agents, refer to "Specify default replication agents" on page 151

   - Select **Use this specified proxy** to select a machine to which you want to deploy a replication agent. Type the path of the machine in the appropriate text box, or click **Browse** to locate the machine in your network.
4. Click **OK**.

Specify default replication agents

You can specify application-wide default proxy replication agents for policies that transfer data. You can specify a replication agent for Windows and UNIX. When VFM cannot deploy a Windows or UNIX agent onto the source or destination, VFM uses the default agents you specify.

1. On the **Tools** menu, click **Options**.
2. Select **Systems Options**.
3. Under **Default Replication Agents**, specify a machine hosting a Windows or UNIX replication agent or a machine to which you want to deploy a replication agent during data transfer.
4. Click **OK**.

Set data transfer rate limit

Set the replication data transfer rate limit or bandwidth throttling for a machine that is hosting or will host a replication agent.

1. In the Physical View tree, right-click a machine and click **Properties**.
2. Click the **Agent** tab.
3. Select **Allow replication agent to be deployed to this host**.
4. Under **Throttling for replication agent on this host**, select the business hours for weekdays and weekends.
5. Type or select a rate, in bytes per second, for business and non-business hours.
6. Click **OK**.

NOTE
If you change throttling on a machine hosting a replication agent while the agent is replicating data, the changes are not implemented for the task currently running. Changes in throttling values appear in the replication tasks view and replication ToolTips the next time the task runs.

Working with physical resources

This section provides instructions about managing machines from the Physical View.

The Physical View

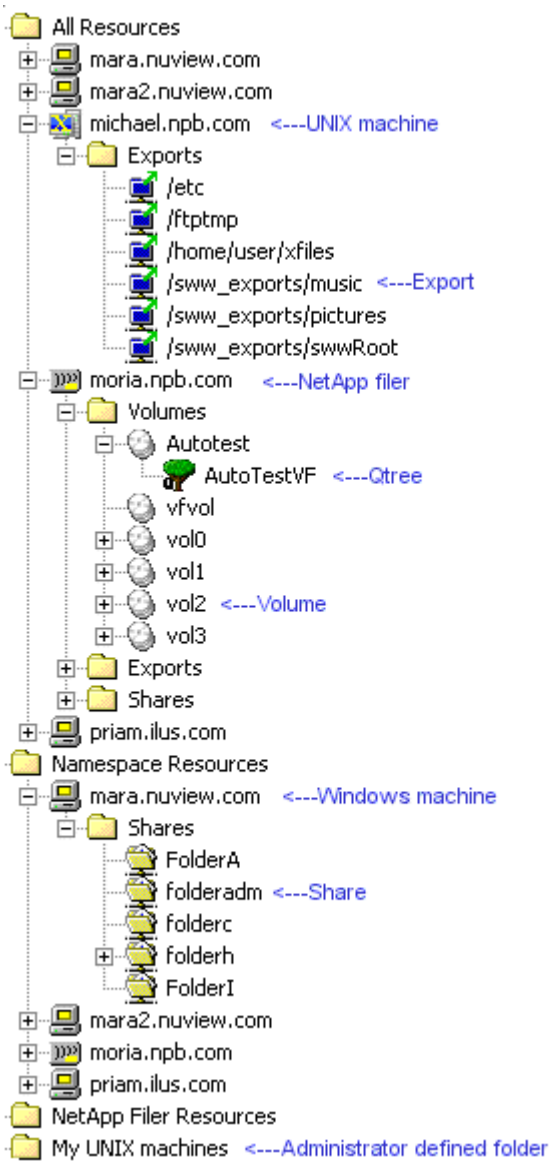

The Physical View shows machines participating in the namespace or a policy and machines you add to folders that you create in the Physical View tree.

- The **Namespace Resources** folder in the Physical View tree is automatically populated with a list of the machines participating in the logical namespace. You can create additional folders in the Physical View to organize your view of the machines. You can add machines to the folders you create whether or not the machines participate in the logical namespace.
- Machines in the Physical View tree not participating in the namespace are identified by a unique icon.
- A machine can appear in multiple folders.
- You can only add and remove machines from Physical View folders that you create. VFM automatically maintains the contents of the following folders:
  - All Resources

  - Namespace Resources
  - NetApp Filer Resources
- After you add a machine to a user-defined folder in the Physical View tree, you can drag a shared folder on the machine to a root in the Logical View to create a link in the namespace. When you do this, the machine is added to the **Namespace Resources** folder.

NOTE
You cannot drag an export in the Physical View to a root in the Logical View because links cannot reference exports.

- NetApp filers appear in the Physical View tree with a NetApp-specific icon indicating whether they participate in the namespace.
- NetApp vFiler images appear in the Physical View tree with a NetApp-specific icon indicating whether they participate in the namespace.
- UNIX machines appear in the Physical View with a unique icon.

Add a machine to the physical view

Create a folder in the Physical View to contain the machines that you want to add.

1. Right-click **Physical View** and click **Add folder**.
2. Type a name for the folder and click **OK**.
3. Perform one of the following actions:
   - Drag a machine from the **Namespace Resources** or **All Resources** folder to the folder you created.
   - Drag a machine from Windows Explorer to the folder you created in the Physical View tree or on the **Properties** tab of **Physical View**.
   - Right-click the folder you created and click **Add Machine**. Use this option to add UNIX NFS machines to the Physical View.

NOTES:

- It is recommended that you reference machines using their NETBIOS name rather than IP address or fully qualified DNS name.
- If a machine appears in the Physical View with a name that is not what you expect, check whether DNS reverse lookups are configured correctly for the machine.

Remove a machine from a folder

Remove a machine from a user-defined folder and delete user-defined folders in the Physical View tree.

To remove a machine

1. In the Physical View tree, open a user-defined folder.
2. Right-click a machine and click **Remove from folder**.

NOTES:

- If the machine participates in the namespace, it continues to appear in the **Namespace Resources** folder.
- If the machine is not participating in the namespace but appears in more than one user-defined folder, it is removed only from the folder you indicated.

To delete a user-defined folder

Right-click a folder and click **Delete**.

Specify NTFS Change Journal

You can enable NTFS Change Journal for NTFS volumes on machines running Windows 2000 and later. When this option is enabled, VFM deploys a replication agent on the machine where the source data resides. This is a deviation from the standard VFM replication agent deployment algorithm.

1. Right-click a machine in the Physical View tree and click **Properties**.
2. Click the **Agent** tab.
3. Click **Allow replication agent to be deployed to this host**.
4. Click **Use NTFS Change Journal on this host**.
5. Click **OK**.

NOTE

For information about problems you might encounter using NTFS Change Journal, refer to "Problems using NTFS Change Journal" on page 390.

Create shares

Create a share on one or more machines and set share and folder permissions. In addition, you can add the shares to a new or existing link, or create new links for each share. This is a quick way of setting up shares for use in a recovery or replication policy.

1. Right-click a machine in the Physical View tree and click **Create shares**.
2. Click **Next**.
3. Type the appropriate information and click **Next**. When creating shares on a NetApp filer, you must use the NetApp naming convention (for example, `c:\vol\voln\<directory name>`, where `n` is the appropriate volume number and `<directory name>` is the name of the directory to be shared.
4. Select the machines where you want to create the share and click **Next**.
5. Perform one of the following actions:
   - Select **Do not modify any DFS links** if you do not want to add the shares to a link.
   - Select one of the other options if you want to add the shares to the namespace. Click **Next**. Select a DFS root or DFS link, depending on the option you selected.
6. Click **Next**.
7. Click **Finish**.

You can also create shares on a machine by performing one of the following actions:

- Drag a share from a machine in the Physical View tree to another machine in the Physical View tree.
- Drag a share on the **Properties** tab of any item in the Logical View or Physical View to a machine in the Physical View tree.

Duplicate shares

Copy multiple shares from one server to another server. You can specify whether to copy all the shares on the source server or only the shares referenced in your namespace. You can also run batch files before and after each share is created.

NOTE
Data is not migrated. To migrate data, create a migration policy.

---

1. On the **Tools** menu, click **Create shares**.
2. Click **Next**.
3. Select **Duplicate Shared From** and type the name of the source and destination servers or browse for the servers in your network.
4. Click **Next**.
5. Select one of the following options:
   - **Only duplicate shares referenced by the logical namespace** - Select this if you want to duplicate only the shares on the source server that are referenced in your namespace.
   - **Duplicate all shares** - Select this if you want to duplicate all the shares on the source server. This option enables you to duplicate hidden shares.
6. Select **Run pre-creation batch file** and type the path of the batch file or click **Browse**. The batch file runs before each share is created. Refer to the notes at the end of this topic for information about running batch files in VFM and more information about the batch file.
7. Select **Run post-creation batch file** and type the path of the batch file or click **Browse**. The batch file runs after each share is created. Refer to the notes at the end of this topic for information about running batch files in VFM and more information about the batch file.
8. Click **Next**.
9. Review the list of source and destination shares.
   - Clear the shares that you do not want to duplicate.
   - To change a destination path, click the destination path and type a new path or click **Browse** to locate a new share.
   - To change the destination path for multiple selected shares, click **Rebase**.
   - Clear **Warn if destination already contains files or folders if you do not want to check** for destination shares that are not empty.
10. Click **Next**. The policy checks for empty destination shares and destination shares that do not exist.
11. If destination shares that are not empty are found, the shares appear in a list. You can continue or select each share and type a different destination share, browse for another share, or remove the share from the duplication process. If destination shares that do not exist are found, the shares appear in a list with a suggested local path for each share.

- To change the local path for multiple items, select the items and click **Rebase**. Type your changes and click **OK**.
- To change one local path, select the item and type a new path.
- To accept all or some of the suggested local paths, select the items on the list and click **Create Destination Share**.

12. Click **Next**.
13. A list of source and destination paths that are validated display. You can clear items that you do not want to duplicate or change the destination path of an item.
14. Click **Next**.
15. Perform one of the following actions:
    - Select **Do not modify any DFS links** if you do not want to add the shares to a link.
    - Select one of the other options if you want to add the shares to the namespace. Click **Next**. Select a DFS root or DFS link, depending on the option you selected.
16. Click **Next**.
17. Click **Finish**.

NOTE
For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Display drives

By default, servers and shared folders are displayed in the Physical View tree. You can select to display the drives.

1. On the **Tools** menu, click **Options**.
2. Select the **Show drives in physical views** check box.
3. Click **OK**.

NOTE
Displaying drives in the Physical View requires gathering additional data from servers; this process can slow the display of information in VFM when managing a large configuration or configurations in which servers are inaccessible.

View attributes of a machine or drive

View the properties and references of a machine or drive that appear in the Physical View tree. In the Physical View tree, click the machine or drive to display the following tabs in the right pane:

- **Properties** - Displays the characteristics or parameters of the selected machine or drive. Right-click an item on the tab to view a context menu.
- **Status** - Shows information concerning attempts by VFM to communicate with the machine or filer.
- **Replication Tasks** - Shows information about replication operations.

- **Snapshots** - Shows information about snapshots. This tab only appears if a folder contains a NetApp or Windows Powered NAS/Windows Storage Server filer or if the item selected is a filer or a volume on a filer.
- **References** - Shows the DFS link referencing the machine or shared folder. Right-click a root or link to view a context menu.
- **Events** - Shows the history of events for the item selected.
- **Manage** - Only appears if the device selected is a NetApp or Windows Powered NAS/Windows Storage Server filer. You can access FilerView or the Windows Powered NAS/Windows Storage Server Web Management UI from this tab.
- **Reports** - Shows a list of reports generated that reference the machine or filer.

View attributes of a shared folder

View the attributes of a CIFS shared folder. The Physical View tree can contain shared folders that participate in the namespace and shared folders that do not participate in the namespace.

In the Physical View tree, click a machine, click the **Shares** folder, and click the share. The following tabs appear in the right pane:

- **Properties** - Displays the characteristics or parameters of the selected share. Right-click anywhere on the tab to view a context menu.
- **Replication Tasks** - Shows information about replication operations.
- **References** - Shows the DFS links that reference the shared folder. Right-click a root or link to view a context menu.
- **Events** - Shows the history of events for the item selected.
- **Reports** - Shows a list of reports generated that reference the shared folder.

NOTE
The application displays share information that is returned by the system hosting the share. Information such as local path and connection limit might not be displayed for a share if the system hosting the share does not provide this information.

View attributes of an export

In the Physical View tree, click a machine, click the **Exports** folder, and click the export. The following tabs appear in the right pane:

- **Properties** - Displays the name and security permissions of the selected export. Right-click anywhere on the tab to view a context menu.
- **Replication Tasks** - Shows information about replication operations.
- **Events** - Shows information about the export.
- **Reports** - Shows a list of reports generated that reference the export.

Update security of a folder

Update permissions for a shared folder.

1. Perform one of the following actions:
   - Click a link in the Logical View. On the **Properties** tab, right-click a target.
   - Click a domain-based root in the Logical View. On the **Properties** tab, right-click a replica.
   - Right-click a folder in the Physical View.
   - Click a replication policy under the **Replication Policies** folder in Admin View. Click the **Replication** tab that appears in the right pane. Right-click a shared folder.
2. Click **Properties**.
3. Click the **Security** tab.
4. Update permissions and click **OK**.

NOTE
If the `Unable to display security information` message appears when you click the **Security** tab, you do not have permission to view or edit the security of the folder.

Refresh a machine

Update the information about a machine and all the shared folders on the machine.

NOTES:

- To refresh the information about all the items in both the Logical View and Physical View, click **Refresh all on the View** menu.
- To refresh the information about all the items in the Logical View or Physical View, right-click Logical View or Physical View in the tree and click **Refresh**.

To refresh a machine

Perform one of the following actions:

- Right-click the machine and click **Refresh**.
- Select a machine and press **F5**.
- Select a machine and click **Refresh** on the main toolbar.
- Select a machine and click **Refresh** on the **View** menu.

To refresh all the machines in a folder

Update the information about all the machines in the **Namespace Resources** folder or a user-defined folder in the Physical View.

Perform one of the following actions:

- Right-click the **Namespace Resources** folder or a user-defined folder and click **Refresh**.
- Select the **Namespace Resources** folder or a user-defined folder and press **F5**.
- Select the **Namespace Resources** folder or a user-defined folder and click **Refresh** on the main toolbar.
- Select the **Namespace Resources** folder or a user-defined folder and click **Refresh** on the **View** menu.

Change the shell credentials of a UNIX machine

To manage UNIX machines that are members of a policy, you must configure VFM with the system or administrative user account and password of the machines. Normally, VFM uses RSH to communicate with UNIX machines. You can choose to use SSH.

NOTES:

- RSH

  When user name and password are left blank, VFM uses root as the user name with no password. You can specify different credentials for a filer or UNIX machine. You can also change the default RSH shell credentials that all new filers and UNIX machines added to VFM will use.

- SSH

  Before you can use SSH, you must configure VFM to use the remote shell. For more information, refer to "Using SSH" on page 343, which provides guidance about how to perform the following tasks:

  - Install and configure OpenSSH.
  - Generate a key-pair.
  - Copy the public key to the filers using SSH.
  - Cache the keys.
  - Provide the location of the `plink` and `pscp` executables.

- To set default RSH or SSH shell credential that all new filers and UNIX machines will use, refer to "Configuring default remote shell credentials" on page 355. Machines already known to VFM are not affected.

To change the shell credentials of a UNIX machine

1. Right-click the UNIX machine either in the Physical View or beneath the policy in which it is contained, and then click **Properties**.
2. Click the **Shell** tab.
3. Select a type in the **Shell Type** list.
4. Perform one of the following actions:
   - *If you select RSH,* type the username and password that VFM will use to communicate with the machine.
   - *If you select SSH,* type the location of the private key and the user name that VFM will use to communicate with the filer. When VFM uses SSH to communicate with UNIX machines and filers, SSH is configured without a password or passphrase.
5. Click **OK**.

Set NFS credentials

Normally, UNIX replication agents use `root(0)` when communicating with the source and destination during NFS data transfers. You can change the NFS credentials for individual filers and UNIX machines appearing in the Physical View. You can also set default NFS credentials for all new UNIX machines and filers added to the Physical View. The NFS credentials of filers and UNIX machines already appearing in the Physical View are not changed.

To set NFS credentials for a filer or UNIX machine

1. Right-click the UNIX machine either in the Physical View or beneath the policy in which it is contained and click **Properties**.
2. On the **General** tab, type a different Default User ID and/or Default Group ID.
3. Click **OK**.

To set default NFS credentials

1. On the **Tools** menu, click **Options**.
2. Click **System Options** to show a dialog box in the right pane.
3. Under **NFS Settings**, type a different Default User ID and/or Default Group ID.
4. Click **OK**.

Chapter

Transferring Data

7

This section provides instructions about how to create policies to transfer data.

In this chapter

- Introduction 163
- Migration policies 168
- Replication policies 200
- Specifying migration tasks from a file 207
- Task Creation batch file. 208
- Example Task Description file. 210
- Archival Migration Policies 210
- Storage Load Balancing Policies 222

Introduction

VFM provides two policies to move data. Select the policy that best meets your needs.

- **Migration policies** enable you to balance free space and data throughput capacity across volumes or to migrate data prior to maintaining or consolidating servers. Use migration policies to transfer data in heterogeneous environments. Before you create your first policy, refer to "Data migration considerations" on page 164 for more information.
- **Replication policies** enable you to centrally manage tasks associated with replicating data across a network. You can select the optimal replication mechanism for your configuration, whether this is VFM for heterogeneous environments, SnapMirror for NetApp filers, or another mechanism suited to your environment. VFM provides status information about replication operations performed using the VFM replication agent. Use replication policies to replicate data from a central location to remote sites and from remote sites to a central location. Before you create your first policy, refer to "Replication planning considerations" on page 167 for more information.

ATTENTION
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.

Data migration considerations

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

Planning considerations

- Users and applications should not be accessing the source or destination folders while the migration is taking place.
- Do not use differential replication on a LAN. When operating on a high-speed network, the computational cost of calculating the file differences is more expensive in terms of time than simply copying the entire file from source to destination.
- When you perform full backups during a migration operation, backups can change archive attribute settings and trigger massive differences between source and destination. With VFM you can avoid this problem, because migration policies enable you to specify how you want to set attributes of destination files.
- When configuring replication options, consider carefully how much event detail you need. If you gather information about all files, it affects performance. As a compromise between performance and the need for feedback, it is recommended that you gather information only when an error is encountered.
- It is not recommended that you enable real-time virus scanning on filers because migration throughput can be severely reduced.
- When planning a migration, consider the available time you have to complete the operation and how much data you have to migrate. Use include and exclude lists sparingly because they add processing time to the migration operation.
- (CIFS only) By default, the security descriptor is copied every time a file is copied. Various options are provided for controlling how a file's security is copied during data migration. For more information, refer to "Select file security replication options" on page 188.
- Exclusive locks on files on source or destination prevent VFM from accessing those files for transfer. Find out who or what program is locking the files and make the files available to VFM.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended duplications of files and can interfere with VFM replication.
- API support limitations in the NetWare® server affect the operation of VFM. VFM can transfer data from Novell® to CIFS shares. Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained.
- (CIFS only) The account used by the VFM server, a replication agent, or a monitoring agent must have the following permissions:
  - **SeBackupPrivilege** (Back up files and directories) - required to open source directories/files with backup semantics.
  - **SeRestorePrivilege** (Restore files and directories) - required to open destination directories/files with backup semantics.
  - **SeSecurityPrivilege** (Manage auditing and security log) - required to access SACLs (System Access Control Lists)
  - **SeTakeOwnershipPrivilege** (Take ownership of files or other objects) - required to delete directories/files.

- (NFS only) The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports. The exports must be writable on the destination.

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

Transferring NTFS file security attributes

- When you transfer file security attributes, VFM looks up the trustees configured as the owner and those referenced in the source file's ACL. If the owner of the file is set to an unknown SID and you attempt a transfer, VFM reports error 1332, No mapping between account names and security IDs was done. You can correct this by assigning ownership.
- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.

Transferring between dissimilar file systems

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

TABLE 25 Transfer results between differing file systems

Source	Destination					
	NTFS W2K	NTFS NT4	FAT16	FAT32	NetApp/CIFS	NFS
NTFS W2K	Attributes maintained See Note 4	Some attributes lost	Some attributes lost See Notes 2 and 6	Some attributes lost See Notes 2 and 6	See Note 5	Not applicable. Copying data from CIFS to NFS is not supported.
NTFS NT4	Attributes maintained	Attributes maintained	Some attributes lost See Notes 2 and 6	Some attributes lost See Notes 2 and 6	See Note 5	Not applicable. Copying data from CIFS to NFS is not supported.
FAT16	Attributes maintained	Attributes maintained	Attributes maintained	Attributes maintained	See Note 5	Not applicable. Copying data from CIFS to NFS is not supported.
FAT32	Attributes maintained	Attributes maintained	Some attributes lost	Attributes maintained	See Note 5	Not applicable. Copying data from CIFS to NFS is not supported.

TABLE 25 Transfer results between differing file systems

Source	Destination					
NetApp/ CIFS	Attributes maintained	Attributes maintained	Some attributes lost	Some attributes lost	Attributes maintained	Not applicable. Copying data from CIFS to NFS is not supported.
NFS	Not applicable. Copying data from CIFS to NFS is not supported.	Not applicable. Copying data from CIFS to NFS is not supported.	Not applicable. Copying data from CIFS to NFS is not supported.	Not applicable. Copying data from CIFS to NFS is not supported.	Not applicable. Copying data from CIFS to NFS is not supported.	See Note 3
NetWare	Some attributes lost See Note7	Some attributes lost See Note7	Some attributes lost See Note7	Some attributes lost See Note7	Some attributes lost See Note7	Not applicable

NOTES:

- Note 1: NFS maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to an NFS file system from NTFS, their security and extended attributes are subsumed by NFS attributes. VFM UNIX replication agents perform NFS to NFS data transfer. Other file systems require redirector client software on the machine hosting the Windows replication agent.
- Note 2: NTFS stores access times in UTC, while FAT stores access time in LocalTime. VFM does not adjust between UTC and LocalTime.
- Note 3: Copying files from NFS sources is supported. This is typically done by using the VFM UNIX replication agent. The account under which the UNIX replication agent is running must have root-equivalent access to the source and destination exports.
- Note 4: Copying hard links, reparse points, or NTFS encrypted files are not supported.
- Note 5: You can configure security on NetApp filers in a variety of ways. These variables affect how security attributes are maintained when copying to NetApp filers. When files are transferred from Windows systems to a NetApp filer and sparse or compressed attributes are set, the attributes have no effect on how the files are stored on the filer.
- Note 6: FAT maintains a very limited set of attributes relative to the set maintained by NTFS. When files are transferred to a FAT file system from NTFS, their access times, extended attributes, indexing attributes, and alternate streams (including security) are lost when the files are written to the FAT destination.
- Note 7: Security settings are updated to reflect the destination security. Legacy settings from the Novell source are not maintained. If the version of NetWare is 5.1 or earlier, the Novell client must be installed on the same machine on which the replication agent is installed. Windows and Novell accounts must be configured with identical user name and password. The VFM console and server components must also run under the same user name.

Specifying migration tasks from a file

Migration policies provide the option of running a batch file to add migration tasks to the policy. You can specify the batch file in the **Run task creation batch** field on the **General** tab of a data migration policy properties dialog box. For more information about the batch file works, refer to "Specifying migration tasks from a file" on page 207. For example task creation batch file and task description file, refer to "Task Creation batch file" on page 208 and "Example Task Description file" on page 210.

Replication planning considerations

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

How it works

- A replication policy is created and the schedule is defined.
- A single-target link is added to the policy as the source link. You can add the link during the creation of the replication policy, or drag it from the Logical View tree or from the right pane to the policy.
- A multi-target link is added to the policy as the destination link. You can add the link during the creation of the replication policy, or drag it from the Logical View tree or from the right pane to the policy. You can quickly add additional targets to the destination link.
- When you select a replication policy in the Admin View, several tabs appear in the right pane. The **Replication Topology** tab provides a graphical display of the distribution path of the policy. From this tab:
  - You can change the distribution path by a simple drag-and-drop operation.
  - You can view monitoring status of the links and each of their targets.
  - You can add additional targets to the destination link by right-clicking the link.

VFM replication planning considerations

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

- Exclusive locks on files on source or destination prevent VFM from accessing the file for transfer. Find out who or what program is locking the files, and make the files available to VFM.
- It is not recommended that you perform full backups during a replication operation because backups can change archive attribute settings and trigger massive differences between source and destination.
- When configuring replication options, carefully consider how much event detail you need. If you gather information about all files, it affects performance. As a compromise between performance and the need for feedback, it is recommended that you gather information only when an error is encountered.
- When planning replication, consider the available time you have to complete the operation and how much data you have to replicate. Use include and exclude lists sparingly because they add processing time to the replication operation.

- It is not recommended that you enable real-time virus scanning on filers because replication throughput can be severely reduced.
- The benefits of differential replication are realized on slower links rather than on a LAN. Computational costs outweigh the benefits of differential replication on a LAN.
- A replication policy synchronizes from its source to a number of read-only copies.
- There is a replication time delay between primary and secondary contents. If a primary target becomes unavailable, there is no guarantee that all the contents of the primary target are replicated to the standby targets.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files and can interfere with VFM replication.
- The VFM client does not have to be running for replication policies to function.

Transferring NTFS file security attributes

For more information about transferring file security attributes, refer to "Data migration considerations" on page 164.

Permissions

For more information about setting permissions for the account used by the VFM server, a replication agent, or a monitoring agent, refer to "Data migration considerations" on page 164.

Status monitoring

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

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

Migration policies

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

- Data transfer can be configured in phases. This enables you to copy data while users have access to the data and configure a final replication at a time you specify when users are temporarily denied access to the source data and are directed to the destination.
- The migration operation is divided into three phases.
  - In the initial phase, you can scan for potential problems and perform a baseline copy of the data.
  - In the incremental phase, you can copy the data incrementally and specify when to advance to the final phase.

  - In the final phase, you can update the namespace, prevent users from connecting to the data, delete the source data, perform a final copy of the data, and stop sharing source or export.
- Before you start creating a migration policy, refer to "Data migration considerations" on page 164 for more information. The topic covers such information as the permissions the account used by the VFM server and monitoring agent must have.
- You can configure the migration operation to execute all three phases (a typical migration) or you can configure the migration operation to execute one or two phases. For more information, refer to "Migration usage options" on page 170.
- Before you configure migration schedules, refer to "Scheduling considerations" on page 171 for more information.

Migration limitations

The following files are not supported:

- Encrypted
- Reparse points
- Hard links

Integrated namespace management

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

DFS namespace management

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

UNIX namespace management

The references to the automount files you specify are updated.

Migration usage options

TABLE 26 Options for Migration Usage

No.	Usage	Steps
1	**Typical data migration (source data might be changing)** • Optionally scan for potential problems before copying data. • Perform a baseline copy of the data. • Perform incremental copies of the data. • Advance to the final phase.	Run the initial phase, incremental, and final phases. 1 Select initial phase actions you want to run. 2 Specify when you want to run the initial phase. 3 Make sure that the incremental phase action **Copy the data** is selected. 4 Specify how often you want to run the incremental phase. 5 Specify when you want to advance to the final phase. 6 Select the final phase actions you want to run. 7 Specify when you want to run the final phase.
2	**Data migration when the source data is static** • Optionally scan for potential problems before copying data. • Copy the data. • Advance to the final phase.Run the initial and final phases.	Select initial phase actions you want to run. 1 Specify when you want to run the initial phase. 2 Clear **Copy the data** action in the incremental phase. 3 Select the final phase actions you want to run. 4 Specify when you want to run the final phase.
3	Data migration when the baseline copy was performed using another method, such as tape backup.	Run the incremental and final phases. 1 Make sure that all the actions are cleared in the initial phase. 2 Make sure that the incremental phase action **Copy the data** is selected. 3 Specify how often you want to run the incremental phase. 4 Specify when you want to advance to the final phase. 5 Select the final phase actions you want to run. 6 Specify when you want to run the final phase.
4	**Data migration when the source data is static and baseline copy was performed using another method, such as tape backup.**	Run the final phase only. 1 Make sure that all the actions are cleared in the initial phase. 2 Clear **Copy the data** action in the incremental phase. 3 Select the final phase actions you want to run. 4 Specify when you want to run the final phase.
5	**Do not copy data. Only scan for potential problems that might interfere with copying the data.**	Run the initial phase only. 1 Select the **Scan for potential problems copying the data** initial phase action. 2 Specify when you want to run the initial phase. 3 Clear all the incremental phase actions. 4 Clear all the final phase actions.

TABLE 26 Options for Migration Usage

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

Scheduling considerations

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

Policy migration schedule

- The policy migration is divided into three phases: initial, incremental, and final. Each phase has a schedule. These schedules control when each migration task associated with the policy executes the actions you specify in the migration phases.
- You can schedule when each migration phase starts. You can also schedule how often a migration task runs the action specified in the incremental phase and how often the final phase should attempt to run a migration task (the final phase can run only one time for each migration task).
- A policy can contain multiple incremental phase and final phase schedules. For example, you can set up a schedule to start the incremental phase every Monday through Friday at midnight, and you can set up another schedule to start the incremental phase every Saturday and Sunday at 6:00 A.M.
- You have the option of manually specifying when a migration task or multiple tasks advance to the final phase. For more information about how advancing to the final phase works, refer to "Scheduling scenarios" on page 176.

Using start dates

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

Initial phase scheduled before incremental phase

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

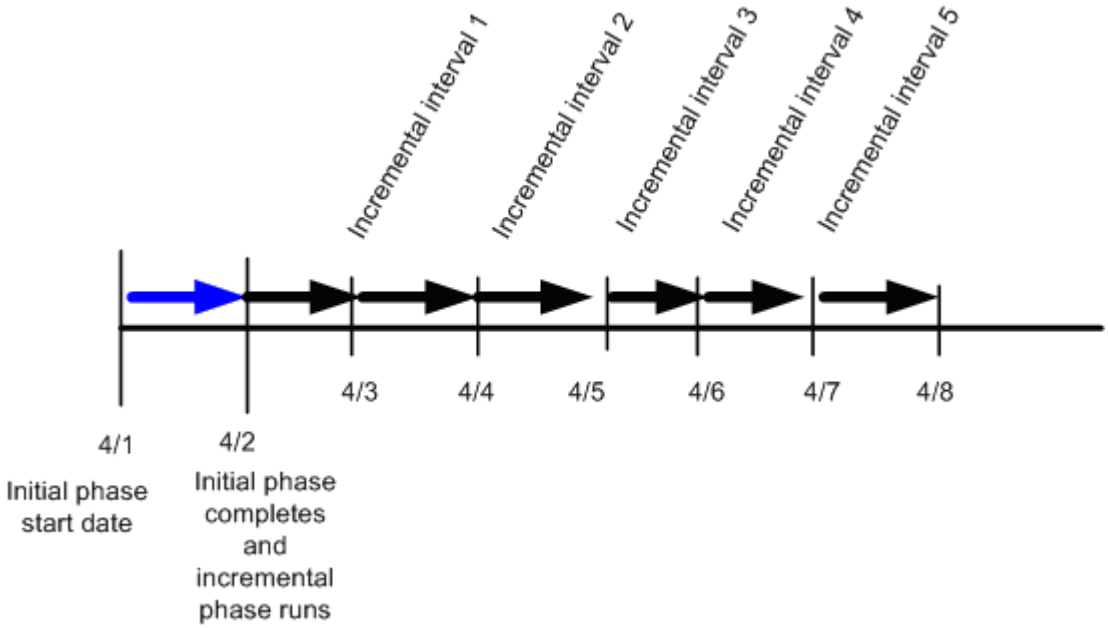

- After the initial phase completes on 4/2, the incremental phase runs.
- The six scheduled incremental copies of the data are executed.

Incremental phase scheduled before initial phase

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

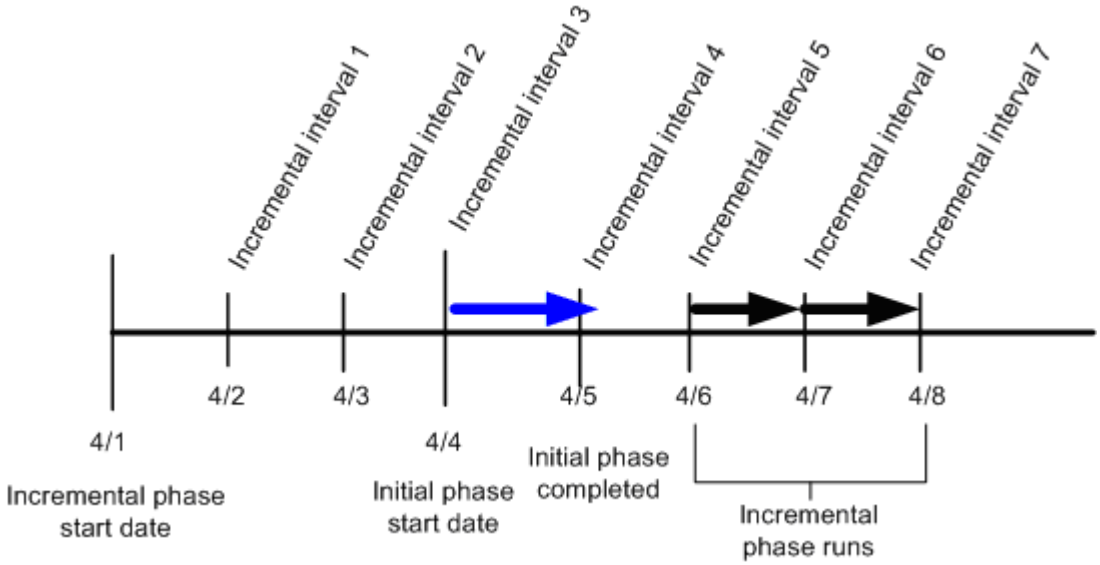

- The incremental phase, which was scheduled to start on 4/1, cannot start until 4/6 because the migration task must complete the initial phase before it can run the incremental phase.

- The incremental phase was scheduled to run six times but runs only two times. The first four incremental intervals never run.
- The incremental phase runs on 4/6, which is the next scheduled run of the phase after the completion of the initial phase.

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

  Final phase scheduled to run at intervals

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

  Scenario

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

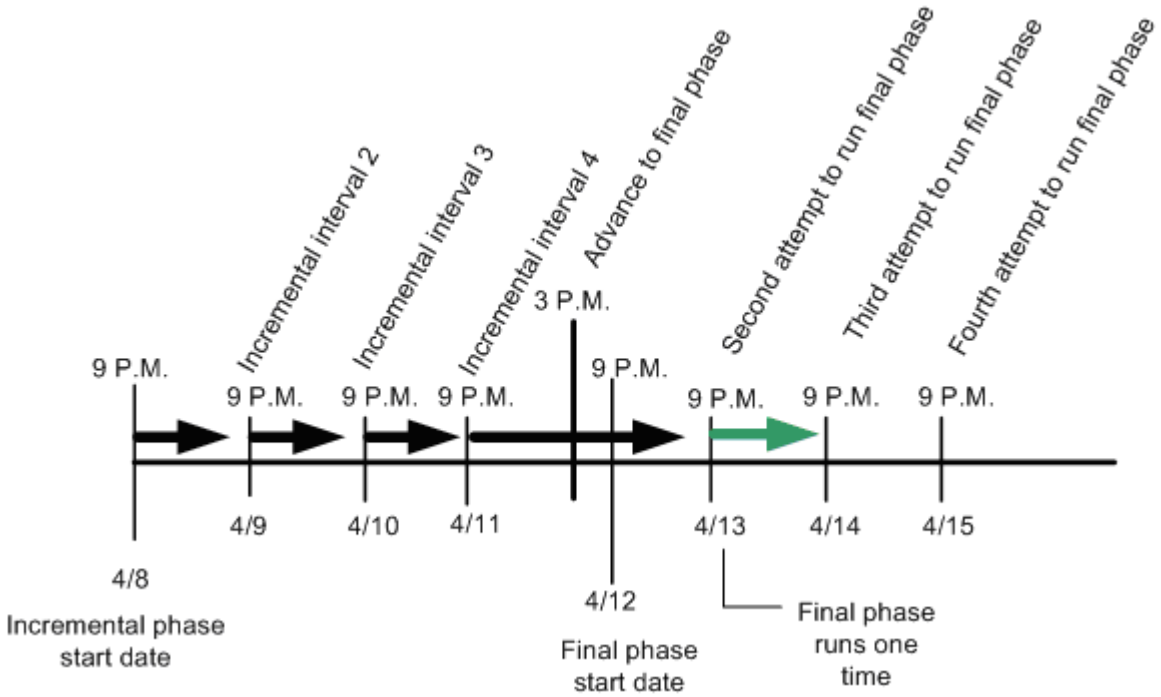

  - It takes longer to complete the fourth incremental copy of the data.
  - The fourth incremental copy of the data runs past 3:00 P.M. on 4/12. The migration task cannot advance to the final phase because the incremental phase is running at the time that the migration task was scheduled to advance to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

Final phase not scheduled to run at intervals

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

Scenario

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

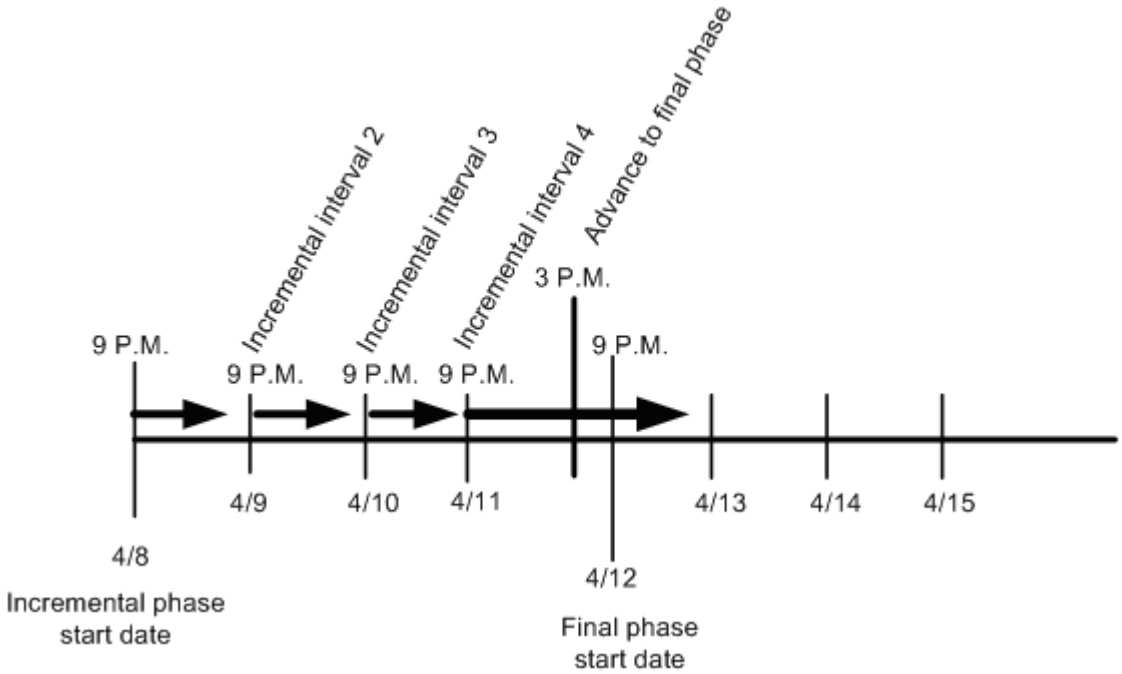

- It takes longer to complete the fourth incremental copy of the data.
- The fourth incremental copy of the data runs past 3:00 P.M. on 4/12. The migration task cannot advance to the final phase because the incremental phase is running at the time that the migration task was scheduled to advance to the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

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

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

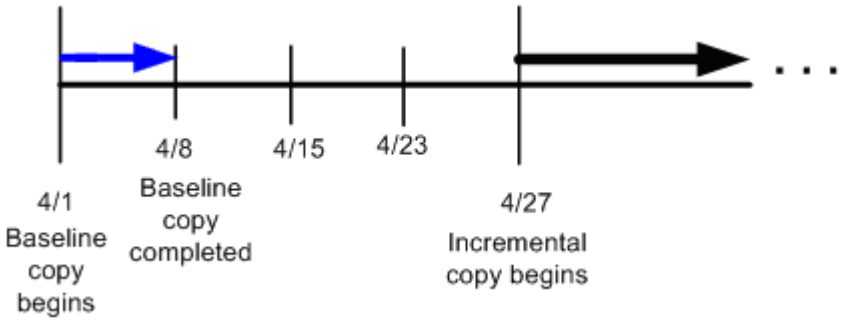

- There is a nineteen-day gap between the baseline copy and the beginning of the incremental copy of the data. Within this time, the data might have changed significantly.
- As a consequence, the incremental copy of the data will take longer to complete.

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

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

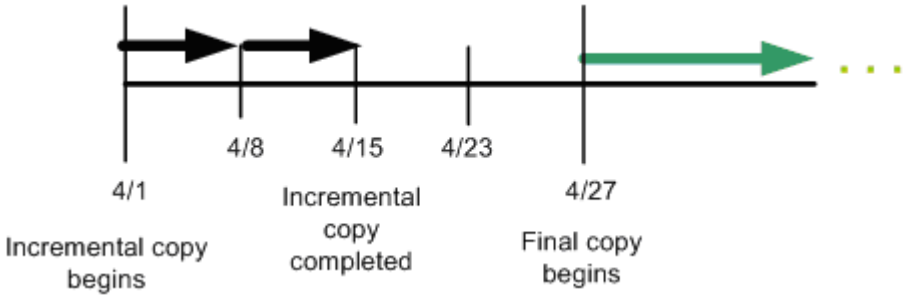

- There is a twelve-day gap between the last incremental copy and the beginning of the final copy of the data. Within this time, the data might have changed significantly.
- As a consequence, the final copy of the data will take longer to complete.

Running task creation scripts to add migration tasks to a policy

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

Scenario

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

Execution of the task creation script

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

On 4/1

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

On 4/3

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

On 4/4

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

On 4/5

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

On 4/6

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

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

Scheduling scenarios

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

Example:
You set the policy schedule to run as follows:

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

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

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

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

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

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

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

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

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

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

Create a migration policy checklist

To create your first migration policy, perform the tasks in the following sections. For more information, refer to the following topics:

1. "Create a migration policy" on page 178
2. Define how you want to execute all the migration tasks associated with the policy.
   - "Configure the initial phase" on page 179
   - "Schedule the initial phase" on page 179
   - "Configure the incremental phase" on page 180
   - "Schedule the incremental phase" on page 180
   - "Configure the final phase" on page 181
   - "Schedule the final phase" on page 182
3. "Specify namespace options" on page 182
4. "Select replication options" on page 184
5. "Add a migration task" on page 191. A migration task identifies the source and destination. A policy can have multiple migration tasks.

Create a migration policy

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

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Data Movement Policy**, and click **Migration Policy**.
   - Drag a machine from the Physical View to the **Migration Policies** folder, and click **New Policy**.
   - Right-click the **Data Movement Policies** folder, point to **New**, and click **Migration Policy**.
   - Right-click the **Migration Policies** folder and click **New Migration Policy**.
   - Right-click **Admin View**, click **New**, point to **Data Movement Policy**, and click **Migration Policy**.
2. Type a word or phrase that best describes this policy. The name appears under the **Migration Policies** folder and on the **Properties** tab of items that are members of the policy.
3. In the **Use the Monitor on** list, select the machine that contains the agent that you want to use to process the migration policy. Normally, the current VFM server is used to process the migration policy. If you installed monitoring agents on other machines, you can select one of the locations where monitoring agents are installed as the location from which this policy is to be processed.
4. Select **Run task creation batch** to run a batch file that specifies the source and destination of data transfers. Type the path and name of a batch file or click **Browse** to locate the file in your network. For more information about creating a migration tasks file, refer to "Data migration considerations" on page 164.

5. Click **OK**.

NOTES:

- If you uninstall a monitoring agent, ensure that migration policies that formerly used the uninstalled agent are reassigned to a currently available agent. Do this by selecting the desired agent through the properties dialog box for the policy.
- The task creation batch file runs every time a migration phase schedule runs. For example, if the initial phase of the policy is scheduled to run at 9:00 P.M. on Saturday and the incremental phase is scheduled to run at 10:00 P.M. every Monday through Friday, the task creation batch file will run at 9:00 P.M. on Saturday and 10:00 P.M. every Monday through Friday.
- For information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Configure the initial phase

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

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

NOTES:

- When an action is suspended or paused after an error occurs, you must resume the migration operation manually. For more information, refer to "User action required" on page 199.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Schedule the initial phase

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

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

4. Click **OK**.

NOTES:

- If you want to skip the initial phase, clear all the actions.
- The default start date for all phases is the date and time the policy was created.
- For more information about scheduling migration phases, refer to "Scheduling considerations" on page 171.

Configure the incremental phase

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

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

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

4. Specify the condition for advancing policy migration tasks to the final phase.
   - **Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually** - Migration tasks perform incremental copies indefinitely at the intervals specified by the incremental phase schedule. You must advance migration tasks to the final phase manually. For more information about how to advance migration tasks manually, refer to "To specify the condition for advancing a migration task to the final phase" on page 196.
   - **After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule** - Migration tasks advance to the final phase after each migration task runs the incremental phase one time.
   - **End the Incremental Phase for all tasks ...on...and advance the tasks according to the Final Phase schedule** - Migration tasks advance to the final phase on the date and time you specify and run the final phase the next time the final phase runs after the date and time you specify.
5. Click **OK**.

Schedule the incremental phase

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

1. Right-click the policy and click **Properties**.
2. Click **Incremental Phase Schedule**.
3. Select the **Use migration schedule** check box to run the incremental phase. Clear the check box if you do not want to run the phase.

4. Click **New** to configure a schedule. The current day and time appear in the schedule list. To change the start time, select the schedule and click **Advanced**.
5. Change the date and time and click **OK**.
6. Specify how often you want to run the phase.
7. Click **OK**.

NOTES:

- If **Use migration schedule** is selected and **Copy the data** in the Incremental Phase dialog box is not selected, the phase will run and the event history will show only a start and end date.
- The default start date for all phases is the day and time the policy was created.
- For more information about scheduling migration phases, refer to "Scheduling considerations" on page 171.

Configure the final phase

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

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

     The contents of the targets of a link are synchronized when the link is a member of a recovery or replication policy. For links that are not members of a recovery or replication policy, you must rely on other mechanisms when possible or you must manually ensure that the contents of the links targets stay synchronized.
   - **by adding offline link targets** - The destination is added as an offline target of a link. Use this option when you are preparing links for a recovery policy because only one target of a link can be online.
   - **by removing the link targets -** The link target pointing to the source is deleted.
4. (CIFS only) Select **Prevent user connections during the final copy** to keep the source folder shared but restrict user access to it during the migration process. The share is renamed (a $ sign is appended to the name of the share). This option is available only when you are migrating CIFS data.
5. Select **Do a final replication** prior to deleting the source data if the data transfer might be interrupted or if it is possible that source files changed since they were last copied to the destination. VFM performs data synchronization before the source data is deleted.
6. Select **Delete the source data** to delete the source data after the data is migrated to the destination.

7. Select **Stop sharing the source share** or export to stop sharing the source file or to remove the export. Select this option if you do not want the source data accessible through its former share or export, but you want to retain the data. Only the specified source share or export is unshared. If there are multiple shares or exports of the source directory, the other shares or exports remain in effect.
8. Select **Run batch file** and type the name of a script file, or browse for a script file to run after each migration task is completed. VFM provides sample batch files that you can modify. The sample batch files are in the `Examples\Scripts\Migration` directory in the VFM server installation directory.
9. Click **OK**.

NOTES:

- When an action is suspended or paused after an error occurs, you must resume the migration operation manually. For more information, refer to "User action required" on page 199.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Schedule the final phase

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

To schedule the final phase

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

NOTES:

- If **Use migration schedule** is selected and no actions are selected for the final phase, the phase will run and the event history will show only a start and end date. No executed actions will appear between the start and end date.
- The default start date for all phases is the date and time the policy was created.
- For more information about scheduling migration phases, refer to "Scheduling considerations" on page 171.

Specify namespace options

Specify how you want to update the namespace.

How automount files are updated

- The shell credentials for the machines hosting automount files must be set in VFM so that they can be managed. For more information, refer to "Change the shell credentials of a UNIX machine" on page 160 and "Change the shell credentials of a filer" on page 144.
- During the data transfer process, the entries that referenced the source location are updated to reference the data at the new location on the destination. For example, you want to migrate data from `SourceServer` to `NewServer`.
  - **Source** - `SourceServer.domain.com:/export/home`
  - **Destination** - `newServer.domain.com:/export/legacy/home`

  The automount file contains a line that references the source:

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

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

  ```
  /home NewServer.domain.com:/export/legacy/home/myHome
  ```
- VFM automatically runs a script named `namespaceupdate` if it is found in the directory containing an automount file you are updating. VFM runs the script after the automount file is updated. The script can be used to perform custom processing of your automount files. For example, an updated automount file is named:

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

  VFM searches for the `namespaceupdate` script file in the directory:

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

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

To specify namespace options

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

Select replication options

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

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

1. Perform one of the following actions:
   - Right-click a policy and click **Properties**.
   - Select a policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Replication Options**.
3. Click **Optimize For Speed** if you want to automatically set replication options when the time it takes to transfer the data is very important.
4. Click **Optimize For Security** if you want to automatically set replication options when the completeness and accuracy of the data transfer is very important.
5. Clear **Include subfolders** if you do not want to copy the source target subfolders and their contents to the destination targets.
6. When you select **Delete orphaned files in destination folders**, files in the destination targets are deleted if they do not have a corresponding file in the source target. For more information about what happens when this option is selected and a policy contains an exclusion list, refer to the Note for Step 15.
7. (CIFS only) In cases where the destination targets do not support auxiliary file streams supported by the source target, you can select **Allow loss of additional file streams** if you want to copy files to the destination targets that lose auxiliary streams.
8. When you select **Enable Copy-in-place**, the source file directly replaces the destination file. This decreases the time required to perform a replication operation. When this option is not selected, the source file is copied to a temporary file in the destination directory, the current destination file is renamed, the temporary file is renamed to the actual file name, and the original destination file is deleted. This minimizes the chance of network errors adversely affecting file contents.
9. (CIFS only) VFM takes advantage of the NTFS Change Journal on the replication source to determine which files are to be replicated rather than scanning the entire source to detect files that have changed. This option saves time when you are replicating large directory trees with limited changes scattered throughout the directory hierarchy.

NOTES:

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

- Use of this option requires that VFM deploy a replication agent on the machine where the source data resides and enable the Change Journal on the NTFS volume hosting the source data. The machine properties dialog box enables you to control whether the replication agent can be deployed to a machine and whether you can enable the NTFS Change Journal for volumes on the machine.
- Use of this option means that changes made on the source are replicated to the destination and changes made on the destination are not overwritten by data from the source unless the corresponding file on the source changes as well. This differs from standard VFM replication, wherein changes made on the destination are detected and overwritten with data from the source.
- In cases where the NTFS Change Journal wraps, a full scan of the source and destination must be performed to determine which files need to be transferred.

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

    To reduce the number of snapshots that must be create, try to group sources in the same volume under the same policy. When **Use snapshot data** is selected, you cannot enable NTFS Change Journal or file differential replication using cached hash information. Snapshot data resides on read-only volumes and NTFS Change Journal and file differential replication require that the data be on writable volumes.

11. Select **Override comparison timestamp resolution** when you want to use a one-second resolution to compare time stamps between the source and destination. By default, the lowest resolution between the source and destination file systems is used when comparing time stamps.

12. In the **Copy files** list, select one of the following options:

    - **If destination file is different or missing** - Replace files in the destination targets whose time stamps do not match the time stamps of files in the source target, and copy files in the source target that are not found in the destination targets.
    - **If destination file is older or missing** - Replace files in the destination targets that have time stamps older than files in the source target, and copy files in the source target that are not found in the destination targets.
    - **If destination file is missing** - Only copy files in the source target that are not found in the destination targets.

13. Type the number of times and the frequency, in seconds, that you want the policy to try to open a file.

14. Type how long, in minutes, a policy should run unsuccessfully before it is aborted.

15. Select one of the following options in the **Events Details** list:

    - **Only list files that encounter errors** - Provide a detailed report of only the files that encounter errors during the replication process.
    - **Don't list files** - Do not provide a report of the files included in the replication process.
    - **List all files** - Provide a detailed report of all the files included in the replication process.

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

16. Click **OK**.

You can perform additional tasks to customize replication activities. For more information, refer to the following topics:

- "Specify replication filters" on page 186
- "Select a replication agent" on page 187
- "Select differential replication options" on page 187
- "Select file security replication options" on page 188
- "Select file attribute replication options" on page 190

Specify replication filters

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

1. Perform one of the following actions:
   - Right-click a policy and click **Properties**.
   - Select a policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Replication Filters**.
3. Type directory exclusions delimited by double quotes, or click **Expand** to specify multiple directories to exclude.
4. Type file exclusions delimited by double quotes, or click **Expand** to specify multiple files to exclude.

NOTES:

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

  PolicyA excludes .doc files. Therefore:

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

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

NOTES:

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

Select a replication agent

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

1. Perform one of the following actions:
   - Right-click a policy and click **Properties**.
   - Select a policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Agent Options**.
3. Select which replication agent will perform data migration.
   - **on the source -** Use the replication agent installed on the source or deploy an agent on the source if an agent is not already installed.
   - **on the destination** - Use the replication agent installed on the destination or deploy an agent on the destination if an agent is not already installed.
   - **on this machine or corresponding proxy** - Use the replication agent installed on the machine you specify or deploy an agent on the machine if an agent is not already installed. When you do not know the name of the proxy that will perform data migration, you can specify the machine to which you assigned a proxy and VFM finds the proxy machine. Use this option when data migration is between two machines on which replication agents cannot be installed, such as NetApp filers. You can type the name of a proxy machine or you can type the name of the filer, and VFM finds the proxy machine you assigned to the filer.
   - **in this agent group** - VFM uses a replication agent from the agent group you select.
   - **Automatically selected** - VFM deploys a replication agent.
4. Click **OK**.

Select differential replication options

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

NOTES:

- When data is replicated using differential replication, both the source and the destination must have replication agents installed. VFM attempts to deploy replication agent to both the source and the destination. If VFM cannot deploy a replication agent to either the source or destination, you must designate a proxy.

- When NFS data is replicated using differential replication, the machines hosting the replication agents must be running the same operating system. The source and destination machines or a proxy machine and the source or destination machine must be running Solaris or must be running Linux.
- File differential replication is not recommended for files in compressed format (for example, JPEG, MP3, and MPEG) because a small change to the content of these files causes changes throughout the entire file.
- To set application-wide default replication differencing options for all new policies, click **Options** on the **Tools** menu.
- To specify a proxy machine, refer to "Specify a proxy machine" on page 151.

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

Select file security replication options

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

NOTES:

- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.
- To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.

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

     NOTE
     The number of bytes transferred for files that had only their security descriptor copied (contents were not copied) might not be accurately reflected in the Event Details dialog box.
4. (CIFS only) Select **Process local trustees in security descriptor** to translate or remove security identifiers (SIDs) from security descriptors of files copied from the source. The default option is to translate the SID. If a local group is not found on the destination, a local group is created. If the SID cannot be translated, VFM translates it to the local administrators group in the destination. When transferring data to a domain controller, you must clear the **Process local trustee in security descriptor** check box. Domain controllers only have domain SIDs, and the processing of local SIDs is not applicable.
   - **Remove the SID from the security descriptor** - Deletes access control entries containing local SIDs from the destination copy of the file. The source files are not modified. Use this option when you do not want to migrate local security settings from the source to the destination. If the owner or primary group of a file or directory being transferred corresponds to a local SID, it is translated to the local administrators group in the destination because the owner and group IDs cannot be removed.
   - **Translate the SID to the same trustee name and type on the target** - Translates local SID in the source files to a local SID in the destination that corresponds to a trustee of the same name and type. Use this option to preserve local security settings in the destination.
5. (CIFS only) Select **Create local groups on the target if not found** if you have not already created local security groups on the destination. This is the default if Translate the SID to the same trustee name and type on the target is selected.
6. If the SID cannot be translated, select one of the following options:

- **Leave it as it is -** Copies the source files without modifying the SID. Use this option to preserve the access control entry (ACE) even though the SID cannot be translated.
- **Remove the SID from the security descriptor** - Deletes ACEs containing local SIDs. Use this option when the SIDs cannot be translated. If the owner or primary group of a file or directory corresponds to a local SID, it is translated to the local administrators group in the destination because the owner ID cannot be removed.
- **Translate to local administrators group -** Replaces the SID read from the source files with the SID of the local administrators group in the destination. Use this option when you want the SID to be meaningful in the destination. This is the default if the SID cannot be translated. Security settings on the source files are not modified.

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

Select file attribute replication options

Select file attributes replication options for a policy.

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

NOTE
When you are transferring NFS data and you clear the Read-Only attribute, write permission is granted to users and groups.

5. The replication is performed and the last access time of the source file is set back to what it was prior to the replication operation. When **Preserve last access time** is not selected, the replication operation causes the last accessed time stamp on the source target to be updated. Selecting this option incurs extra processing and slows replication performance.

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

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

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

7. Click **OK**.

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

Add a migration task

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

CIFS Notes

- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.
- CIFS physical paths must be UNC paths; for example, `\\<machine>\<share>`.
- When you want to update the namespace after a migration, make sure that the UNC paths of the source and the link targets you want to update match. If the UNC paths do not match, the namespace cannot be updated. Example:
- The following example illustrates what happens when the UNC path of a migration source and link targets that match the search criteria do not match.

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

  TABLE 27

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

You scan for migration candidates and find one candidate.

TABLE 28

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

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

TABLE 29

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

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

TABLE 30

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

TABLE 31

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

After migration, the target for the link Sales does not contain any data because the data was migrated to the destination.

To avoid this problem, the source path should match the link target path.

TABLE 32

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

TABLE 33

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

TABLE 34

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

NFS Notes

- The source and destination must be UNIX fully qualified paths (for example, `<hostname>:/<root>/<subfolder>/<file>`).

- NFS 3 or higher must be running on the source and destination.
- A replication agent must be installed on the source or destination or you can specify a default proxy.
  - For more information about how to specify default Windows and UNIX proxy machines, refer to "Specify default replication agents" on page 151
  - For more information about how to specify a Windows or UNIX proxy machine, refer to "Specify a proxy machine" on page 151.
- During the data transfer process, the entries that referenced the source location are updated to reference the data at the new location on the destination. For example, you want to migrate data from `SourceServer` to `NewServer`
  - **Source** - `SourceServer.domain.com:/export/home`
  - **Destination** - `newServer.domain.com:/export/legacy/home`

  The automount file contains a line that references the source.

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

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

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

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

To add a migration task

1. Perform one of the following actions:
   - Drag a share (UNC path) from Windows Explorer to a migration policy.
   - Right-click a migration policy and click **Add task to policy**.
   - Drag a machine, share, or export from the Physical View to a migration policy.
   - Drag a machine, share, or export from the Physical View to the **Migration Policies** folder. Select an existing policy or create a new policy.
   - Right-click a machine, share, or export in the Physical View, then click **New Migration Task**. Select an existing policy or create a new policy.
   - Select a machine in the Physical View to display a **Properties** tab in the right pane and drag a share or export to a migration policy.

- Select a link in the Logical View to display a **Properties** tab in the right pane and drag a share under **Target to a migration policy**.

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

   TABLE 35

Destination export	Acceptable export points

`moria:/home/directory1/directory2`	`/home`
	`/home/directory1`
	`/home/directory1/directory2`

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

NOTES:

- When you create an export on a NetApp filer, the path must begin with `/vol/` followed by a valid name of a volume (for example, `/vol/vol1/<directory>` or `/vol/vol0/<directory>/<subfolder>`). The volume must not already be exported.

- If part of the destination export path is already exported, the new export path is created under the current export and no new export is created.
- When you create an export on a UNIX machine, all the path elements specified are created.
- The permissions on the source are copied to the destination.
  - If data transfer is between different platforms (for example, Linux, Solaris, or a filer), VFM translates the applicable security permissions for the new export.
  - If the directory already exists, the source permissions are copied to the destination. The ability to set UID and GID permissions must be enabled on the NFS volume on which you want to create the export. The VFM server runs with NFS credentials of root. The export must be configured to enable the machine hosting VFM to log in as root. If you are overriding the NFS credentials set for VFM, you must be the current owner of the destination directory.
- To create exports on the destination machine, you must set shell credentials for UNIX machines and filers participating in data transfer. For more information about how to set shell credentials, refer to "Change the shell credentials of a filer" on page 144 or "Change the shell credentials of a UNIX machine" on page 160.

10. Click **Next**.
11. A list displays of source and destination paths that are validated. You can clear items that you do not want to migrate, or you can change the destination path of an item.
12. Click **Next**.
13. (CIFS only) If the policy has the action to update the namespace, a list of links to be updated appears. Clear any links that you do not want to update.
14. Select when you want the migration task to advance to the final phase.
    - **Never; Remain in Incremental Phase** - The migration task repeats the incremental phase indefinitely. You must manually advance the migration task to the final phase. For more information about how to manually advance the migration task, refer to "User action required" on page 199.
    - **After the current Incremental Phase run** - The migration task runs the incremental phase one time and advances to the final phase. The next time the final phase runs, the migration task executes the actions specified.
    - **On the first Final Phase interval after** - The migration task advances to the final phase the next time the final phase runs after the date and time you specify.
    - **According to the policy default** - The migration task advances to the final phase based on the settings you specified for the policy. For more information about how to configure the cutover date for the policy, refer to "Advance to final phase" on page 195.
15. Click **Finish**.

NOTE
For more information about how to monitor the progress of a migration task, refer to "View progress of a task" on page 199.

Advance to final phase

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

- Incremental copy of the data is performed for each migration task until the condition for advancing a migration task to the final phase is met. When the condition is met, a migration task advances to the final phase and executes the actions specified in the final phase the next time the final phase runs.
- The conditions for advancing to the final phase can be configured at the policy or migration task level.
- Before you specify the conditions for advancing migration tasks to the final phase, refer to "Scheduling considerations" on page 171 for more information.

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

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

1. Double-click **Admin View**.
2. Double-click the **Data Movement Policies** folder.
3. Double-click the **Migration Policies** folder.
4. Right-click the policy you want to specify when you want to advance to the final phase and click **Properties**.
5. Click **Incremental Phase**.
6. Select **one of the following options:**
   - **Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually** - Migration tasks perform incremental copies indefinitely at the intervals specified by the incremental phase schedule. You must advance migration tasks to the final phase manually. For more information about how to advance migration tasks manually, refer to "To specify the condition for advancing a migration task to the final phase" on page 196.
   - **After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule** - Migration tasks advance to the final phase after each migration task runs the incremental phase one time.
   - **End the Incremental Phase for all tasks ... on...and advance the tasks according to the Final Phase schedule** - Migration tasks advance to the final phase on the day and time you specify. The next time the final phase runs after the date you specify, the migration tasks execute the actions specified in the final phase.
7. Click **OK**.

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

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

1. Perform one of the following actions:
   - Right-click a migration task appearing under a migration policy and click **Advance to Final Phase**.
   - Select a migration policy in the Admin View to show a **Properties** tab in the right pane. Under **Tasks**, right-click a migration task and click **Advance to Final Phase**.
2. Select one of the following options:

- **Now** - Manually advance the migration task to the final phase immediately if the incremental phase is not running. If the incremental phase is running when you select this option, the policy continues to execute the incremental phase and you get an error message. For more information, refer to "Scheduling considerations" on page 171.
- **Never; Remain in Incremental Phase** - The migration task runs the Incremental phase indefinitely. You must manually advance the migration task to the final phase.
- **After the current Incremental Phase run** - After the migration task performs the actions specified in the incremental phase one time, the migration task advances to the final phase. The next time the final phase runs, the migration task executes the actions specified in the final phase.
- **On the first Final Phase interval after** - Select a date and time after which you want the migration task to advance to the final phase. The next time the final phase runs, the migration task executes the actions specified in the final phase.
- **According to the policy default** - The migration task uses the option specified for the policy.

3. Click **OK**.

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

Perform one of the following actions:

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

Specify migration tasks from a file

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

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

NOTE

For more information about creating a migration tasks file, refer to "Data migration considerations" on page 164.

Run customized commands during replication

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

1. Perform one of the following actions:
   - Right-click a migration policy and click **Properties**.
   - Select a migration policy and click **Edit** on the **Properties** tab that appears in the right pane.

2. Click **Replication Scripting**.
3. Select the options you want and specify the batch files you want to run. The batch file can be specified either as the UNC path to a file accessible on your network or as a local path to a file residing on the server where VFM monitoring is taking place
4. Click **OK**.

NOTES:

- (Migration, Archival Migration, and Storage Load Balancing policies only) When you specify scripts to run before and after replication, the files run before and after the following migration actions:
  - Scan for potential problems copying the data (Initial Phase).
  - Perform a baseline copy (Initial Phase).
  - Copy the data (Incremental Phase).
  - Do a final replication prior to deleting the source data (Final Phase).
  - Delete the source data (Final Phase).
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Update the properties of a migration task

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

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

Start a migration task immediately

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

Perform one of the following actions:

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

Start all tasks for a policy immediately

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

Perform one of the following actions:

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

View progress of a task

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

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

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

User action required

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

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

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

TABLE 36 Tasks with User Action Required

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

Suspend a migration task

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

Perform one of the following actions:

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

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

Remove a task

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

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

Delete a migration policy

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

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

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

Replication policies

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

Before you begin creating a replication policy, review the information in "Replication planning considerations" on page 167. The topic covers detailed information, such as the permissions the account used by the VFM server and monitoring agent must have.

Create a replication policy checklist

To create a policy, perform the tasks in the following sections. For more information, refer to the following topics:

1. "Create a replication policy" on page 201
2. "Add or replace the source" on page 203
3. "Add or replace the target" on page 204
4. "Select replication options for your policies" on page 202
5. "Schedule data distribution" on page 202
6. "Change data distribution path" on page 203

Create a replication policy

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

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Data Movement Policies**, and click **Replication Policy**.
   - Drag a link to the **Replication Policies** folder.
   - Right-click the **Data Movement Policies** folder, point to **New**, and click **Replication Policy**.
   - Right-click the **Replication Policies** folder and click **New Replication Policy**.
   - Right-click **Admin View**, click **New**, point to **Data Movement Policy**, and click **Replication Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field. The name appears under the **Replication Policies** folder and on the **Properties** tab of items that are members of the policy.
3. In the **Source** field, type a CIFS or NFS physical path or the name of a link that references a single CIFS target, or click **Browse** to locate the source. If the source is a single target link, it must be writable. The source can be a member of more than one replication policy.
4. In the **Target** field, type CIFS physical paths and/or links that reference CIFS targets or NFS physical paths. A target link can reference multiple targets. However, a target can be a member of only one replication policy.
5. You can select another location where operation of the replication policy is performed. The default is to run from the VFM server. From the **Use the Monitor on** list, select the host name of another monitoring agent.
6. Click **OK** to save the policy.

NOTES:

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

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

Select replication options for your policies

Configure replication, file security, and replication differencing options. For more information, refer to the following topics:

- "Select replication options" on page 184
- "Specify replication filters" on page 186
- "Select a replication agent" on page 187
- "Select differential replication options" on page 187
- "Select file security replication options" on page 188
- "Select file attribute replication options" on page 190

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

Schedule data distribution

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

- Before you schedule data replication, review your network bandwidth and infrastructure. For example, review the rate of data changes and link speeds.
- You to set the data transfer rate or throttling for a machine. In the Physical View tree, right-click a machine, click **Properties**, and click the **Agent** tab.
- To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.

To schedule data distribution

1. Right-click a policy and click **Properties**.
2. Click **Replication Schedule**.
3. Click **Use replication schedule** and configure the schedule.

4. Optionally, select **Use alternate schedule when the policy fails**, and type a frequency (in minutes) to run the schedule any time replication fails.
5. Click **OK**.

NOTES:

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

Change data distribution path

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

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

Change replication for a target

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

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

NOTES:

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

Add or replace the source

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

1. Perform one of the following actions:
   - Drag a single target link from the Logical View, a shared folder from the Physical View, a shared folder (UNC path) from Windows Explorer, or an export from the Physical View. Select **As the source path**.

- If the policy did not have a source, the newly-added item becomes the source.
- If the policy already had a source, the newly-added item replaces the current source.
- Drag a shared folder from the Physical View or a shared folder (UNC path) from Windows Explorer. Select **Replace the existing source link's target**. The shared folder replaces the share of the policy's source link.

2. Click **OK**.

NOTES:

- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.
- Links that are already members of recovery policies cannot be members of a replication policy because links contained in recovery policies can have only one active target.
- An item that is already a destination in a replication policy cannot be a source of another replication policy.
- It is recommended that you add a link that has a shared folder as its target to a replication policy. Due to replication agent location considerations, it is not recommended that you add a link to a replication policy that has a DFS root or another link as its target.
- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files.
- For more information, refer to "Create a replication policy" on page 201.

Add or replace the target

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

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

NOTES:

- Links that are already members of recovery policies cannot be members of a replication policy because links contained in recovery policies can have only one active target.
- A link that is already a destination link in a replication policy cannot be a destination link of another replication policy.
- It is recommended that you add only links to a replication policy that have shared folders as their targets. Due to replication agent location considerations, it is not recommended that you add links to a replication policy that has a DFS root or another link as its target.

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

Change the replication schedule

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

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

Start replication immediately

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

Perform one of the following actions:

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

NOTES:

- This menu item is disabled if a data replication operation is in progress.
- To view information about any replication in progress, click the **Replication Tasks** tab on the right pane, and then click the blue circle.

Cancel a replication in progress

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

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

NOTE
To view information about any replication in progress, click the **Replication Tasks** tab, and then click the blue circle.

View information about replication operations

View information about replication operations configured in VFM. By default, the information for all replication operations appears. You can filter the information to see only replication operations with problems or only replication operations in progress.

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

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

Compare contents of targets in a replication policy

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

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

View properties of a replication policy

View the properties of a replication policy.

Perform one of the following actions:

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

Delete a destination target

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

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

Delete a replication policy

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

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

Specifying migration tasks from a file

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

How it works

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

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

The task creation batch file

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

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

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

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

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

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

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

VFM invokes your batch file as:

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

Task File Description Syntax

`DFSMIGRATION` (element): required

- Version: required, must be set to **1**.

`Task` (element): optional (though nothing will happen if you do not have at least one of these)

- `CurrentTarget` (attribute): required

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

Task Creation batch file

The following example shows a task creation batch file that creates the task description file illustrated in "Example Task Description file" on page 210.

```
@echo off
rem ******************************************************************************
rem FILE: CreateTasks.bat
rem
rem PURPOSE:
rem    Illustrate creating a task-description file
```

```
rem
rem
rem INPUT:
rem    %1 Name of the migration policy calling this BAT file.  This is supplied
from VFM.
rem    %2 Name of the output xml file.  This is supplied from VFM.
rem
rem COPYRIGHT:
rem    Copyright (c) 2007 Brocade Communications Systems, Inc.  All rights
reserved.
rem
rem ****************************************************************/
rem
rem Next two lines set up the start of the XML file to be handed to VFM
rem
echo ^<?xml version="1.0" encoding="UTF-8" ?^> > ExampleTasks.xml
echo ^<DFSMIGRATION Version="1"^> >> ExampleTasks.xml
rem Add first task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src1^"
NewTarget=^"\\SERVERB\Test\Dst1^" ^/^> >> ExampleTasks.xml
rem Add second task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src2^"
NewTarget=^"\\SERVERB\Test\Dst2^" AutoDelete=^"1^" ^/^> >> ExampleTasks.xml
rem Add third task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src3^"
NewTarget=^"\\SERVERB\Test\Dst3^" AutoDelete=^"0^" ^/^> >> ExampleTasks.xml
rem Add fourth task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src4^"
NewTarget=^"\\SERVERB\Test\Dst4^" AutoDelete=^"1^"^>   >> ExampleTasks.xml
echo ^<Link LinkName=^"\\ROOTSERVER\Root\Link^" ^/^> >> ExampleTasks.xml
echo ^<Link LinkName=^"\\ROOTSERVER\Root\OtherLink^" >> ExampleTasks.xml
echo CurrentTarget=^"\\SERVERA\Test\Src4\Subdir^" ^/^> >> ExampleTasks.xml
echo ^<^/Task^> >> ExampleTasks.xml
rem
rem Terminate the XML file and copy it to where VFM told us to put it
rem
echo ^</DFSMIGRATION^> >> ExampleTasks.xml
copy ExampleTasks.xml %2
rem
rem Delete our temporary file
rem
del ExampleTasks.xml
```

```
exit 0
```

Example Task Description file

The following example shows the task description file created by the batch file illustrated in the "Task Creation batch file" on page 208 that adds four tasks to a migration policy.

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

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

Archival Migration Policies

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

DFS namespace management

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

UNIX namespace management

- The references to the automount files you specify are updated.

Planning considerations

Before you create archival migration policies to move directories, review the following information:

- Data transfer must occur within the same protocol (for example, CIFS to CIFS or NFS to NFS).
- A policy can contain CIFS source and destination pairs and NFS source and destination pairs.
- If you uninstall a monitoring agent, ensure that archival migration policies that formerly used the uninstalled agent are reassigned to a currently available agent. Do this by selecting the desired agent through the properties dialog box for the policy.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.
- To configure default options for all archival migration policies, click **Options** on the **Tools** menu and click **Archival Migration** or any menu item under Archival Migration. The changes affect only new policies. Existing policies are not affected.

CIFS requirements

This section provides information you need to know when you add CIFS source and destination paths.

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

  TABLE 37

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 path (for example, `\\server\share\subdir`) or a logical path (for example, a path in the namespace).
- Links selected as the source must have one target.
- The destination must be a physical path.
- When the source is a physical path, the policy scans for subfolders. The shared folder named in the source path must contain subfolders because the policy 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, the policy scans the target of a link. If the link target matches the selection criteria, it is listed as a migration candidate.

NFS requirements

This section provides information you need to know when you add NFS source and destination paths.

- The source and destination must be UNIX fully qualified paths; for example:

  ```
  <hostname>:/<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.

TABLE 38

Steps		Source	Destination
1	Search path A is con-figured.	hostname:/root/subdirectoryA	hostname:/root/subdirectoryB
2	Search path B is con-figured.	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

Archival migration policy checklist

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

1. Give the policy a name and add source and destinations.
2. Specify the migration criteria that determines whether a directory should be migrated.
3. Configure migration phases.
4. Schedule migration.
5. (Optional) Specify how you want to update the namespace.
6. Define how you what to manage migration candidates.
7. Schedule when the policy scans for directories to migrate.

You can specify replication, file security, and replication differencing options for CIFS data. You can also maintain exclusion and inclusion lists for CIFS and NFS data.

Create a policy

Create an archival migration policy to automatically migrate directories from primary to secondary storage. You configure the policy with the criteria you want to use to select directories to migrate. You can transfer CIFS and NFS directories. When you transfer CIFS directories, you can automatically update the namespace to reference the data at the new location. When you transfer NFS directories, you can automatically update references to the directories in automount files.

Before you create your first archival migration policy, review the requirements.

1. Perform one of the following actions:
   - On the **Tools** menu, point to **New Policy**, point to **Data Lifecycle Automation Policy** and click **Archival Migration Policy**.
   - Right-click the **Data Lifecycle Automation Policies** folder, point to **New**, and click **Archival Migration Policy**.
   - Right-click the **Archival Migration Policies** folder and click **New Archival Migration Policy**.
2. Type a word or phrase in the **Name** field that best describes this policy. The name appears under the **Data Lifecycle Automation Policies\Archival Migration Policies** folder and on the **Properties** tab of items that are members of the policy.

3. Click **Add** to include the source and destination you want to migrate.
4. Repeat Step 3 to include additional source and destination paths.
5. (Optional) Type directory exclusions delimited by double quotes, or click **Expand** to specify multiple directories to exclude. Refer to the table of examples of valid formats in the Notes section of "Exclude directories from migration" on page 218.
6. (Optional) Type directory inclusions delimited by double quotes, or click **Expand** to specify multiple directories to include. Refer to the table of examples of valid formats in the Notes section of "Exclude directories from migration" on page 218.
7. In the **Use the Monitor on** list, select the machine hosting the agent that you want to use to process the policy. By default, the current VFM server is used to process the policy. If you installed monitoring agents on other machines, you can select one of the locations where monitoring agents are installed as the location from which this policy is processed.
8. (Optional) Select **Run batch before scan** and type the path of or browse for a script you want to run before scanning for migration candidates. VFM provides sample batch files that you can modify. The sample batch files are in the `Examples\Scripts \Archival` directory in the VFM server installation directory.
9. (Optional) Select **Run batch after scan** and type the path of or browse for a script you want to run after scanning is complete. VFM provides sample batch files that you can modify. The sample batch files are in the `Examples\Scripts\Archival directory` in the VFM server installation directory.
10. Click **OK**.

Specify selection criteria

Specify the migration criteria that determine whether a directory should be migrated. You can configure the criteria provided by VFM, or you can specify your own criteria. You can migrate all directories found or search for directories that match the search criteria you specify.

1. Perform one of the following actions:
   - Right-click an archival migration policy and click **Properties**.
   - Select an archival migration policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Migration Criteria** to display a dialog box on the right pane.
3. Normally, Always migrate is selected when you want to migrate all the directories found in the source listed in the policy. This option enables you to migrate directories without using migration criteria. The size and size on disk is calculated and displayed for each migration candidate. Calculating this data might take time, depending on the size of the migration candidates. Network traffic is generated as the information is gathered.
4. Select **Migrate directories when** to search for directories that match the criteria you specify. When you specify multiple items, the items are connected with OR Boolean logic.
   - **% of files have not been accessed within the last...** - Migrate directories when a specified percentage of the files in the directories are not accessed within a specified time.
   - **% of files have not been modified within the last...** - Migrate directories when a specified percentage of the files in the directories have not been modified within a specified time.

- **(CIFS only) % of files have not been created within the last...** - Migrate directories when a specified percentage of the files in the directories have not been created within a specified time.
- **(CIFS only) Directory is more than...old** - Migrate directories that have reached a specified age.
- **Directory size or size on disk is greater than MB** - Migrate directories with size or size on disk greater than a specified size.
- **Directory contains more than... files** - Migrate a directory that contains more than a specified number of files.
- **Confirmed by user selection criteria script** - Migrate a directory based on the criteria in the script file listed. VFM provides a sample batch file that you can modify.

5. Click **OK**.

NOTES:

- A sample migration criteria batch file is located in the `Examples\Scripts\Archival` directory in the VFM server installation directory.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Configure migration phases

Define how you want to run all the migration tasks associated with the policy. The migration operation is divided into three phases; initial, incremental, and final. For more information about how to configure the phases, refer to the following migration policy topics:

- "Configure the initial phase" on page 179
- "Configure the incremental phase" on page 180
- "Configure the final phase" on page 181

Specify namespace options

If you selected to update the namespace as part of the migration process, you must specify namespace options. For more information, refer to the Migration policies topic, "Specify namespace options" on page 182.

Select migration task options

Select how you want to execute the policy. Decide how you what to manage migration candidates.

1. Perform one of the following actions:
   - Right-click an archival migration policy and click **Properties**.
   - Select an archival migration policy and click on the **Properties** tab that appears in the right pane.
2. Click **Task Options** to display a dialog box in the right pane.

3. Select **Automatically add migration tasks for migration candidates** when you want to automatically create migration tasks after scanning is complete. When this option is cleared, you see a list of migration candidates. You can manually create migration tasks for all or selected candidates.
4. Select **Calculate total size for migration candidates** when you want to calculate and display the size and size on disk for each migration candidate. Calculating this data might take time, depending on the size of the migration candidates. Network traffic is generated as the information is gathered.
5. Select **Automatically delete completed tasks** when you want to delete tasks after they are completed. Clear this option when you want to keep a record of completed migration tasks. VFM provides a visual cue to indicate that a task was completed successfully.
6. Select **Only migrate data referenced through the namespace (CIFS only)** when you want to search only for directories referenced by links in the roots in the **Search these DFS roots for links targeting migration candidates** list. Clear this option when you want to search for all directories that match the selection criteria. You see a list of migration candidates and which candidates are referenced by links.
7. Click **Add** to include DFS roots in the **Search these DFS roots for links targeting migration candidates (CIFS only)** list. As part of the migration, the links are updated to reference the new location. Users are shielded from changes in the location of the data. For more information, refer to the Notes in "Add source and destination paths" on page 216.
8. Click **OK**.

Schedule scanning for directories to migrate

Normally the policy automatically scans for directories that match the policy migration criteria and shows a list of migration tasks when scanning is complete.

1. Perform one of the following actions:
   - Right-click an archival migration policy and click **Properties**.
   - Select an archival migration policy and click on the **Properties** tab that appears in the right pane.
2. Click **Policy Scan Schedule**.
3. Click **New** and configure a new schedule or select a schedule from the list to update.
4. Click **OK**.

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

Schedule migration

Schedule when you want to run the initial, incremental, and final phases for all the migration tasks of a policy. Before you configure migration schedules, refer to the information in "Scheduling considerations" on page 171. For more information about how to schedule the migration, refer to the following migration policy sections:

- "Schedule the initial phase" on page 179

- "Schedule the incremental phase" on page 180
- "Schedule the final phase" on page 182

To configure replication options, such as, file security, replication differencing, and file attributes options, refer to the following migration policy sections:

- "Select replication options" on page 184
- "Specify replication filters" on page 186
- "Select a replication agent" on page 187
- "Select differential replication options" on page 187
- "Select file security replication options" on page 188
- "Select file attribute replication options" on page 190

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

Add source and destination paths

Quickly include additional source and destination paths to an archival migration policy by dragging items to the policy. A wizard steps you through the process. You can drag items from the Logical View, Physical View, or Windows Explorer to a policy.

NOTES:

- Review the requirements for transferring CIFS and NFS data.
- When you want to update the namespace after a migration, make sure that the UNC paths of the source and the link targets you want to update match. If the UNC paths do not match, the namespace cannot be updated. Example

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

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

  TABLE 39

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

- You scan for migration candidates and find one candidate.

  TABLE 40

Migration candidates
`\\Saturn.npb.com\c$\Sales\Reports\HoustonRegion`

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

TABLE 41

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

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

TABLE 42

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

TABLE 43

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

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

TABLE 44

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

TABLE 45

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

TABLE 46

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

1. Drag one of the following items to an archival migration policy:
   - A logical folder or link in the Logical View
   - A shared folder in the Physical View
   - An export in the Physical View
   - A shared folder or subfolder of a share in Windows Explorer

   A dialog box appears with the path of the item you dragged in the source field.
2. Type a destination path or click **Browse** to find a destination.
3. Click **Next**.

4. Optionally clear the items you do not want to add to the policy.
5. Optionally select an item and click **Rebase** to change the destination path.
6. Click **Next**.
7. Click **Finish**.

Exclude directories from migration

Maintain an exclusion list of directories that you do not want to migrate.

1. Perform one of the following actions:
   - Right-click an archival migration policy and click **Properties**.
   - Select an archival migration policy to show a **Properties** tab in the right pane. Click **Edit**.
2. Click **General**.
3. Type directory exclusions delimited by double quotes, or click **Expand** to specify multiple directories to exclude. The UNC paths and UNIX fully qualified paths can include DOS wild cards (* or ?).
4. Click **OK**.

NOTE
When you exclude a migration candidate, the path of the candidate is added automatically to the policy exclusion list.

To exclude a migration candidate:

1. Select the archival migration policy to show a **Properties** tab in the right pane.
2. Right-click a migration candidate and click **Exclude**.
   - You can also maintain an inclusion list for a policy. Directory inclusions follow the same rules as directory exclusions.
   - After a migration candidate appears in the exclusion list, you cannot include it as a candidate by adding it to an inclusion list. You must first remove the candidate from the exclusion list.
   - The following table provides a list of valid exclusion and inclusion formats and includes an explanation for each format.

TABLE 47 Valid migration exclusions and inclusion formats

Item	Format	Description
1	*	Match everything.
2	*exp*	Match any path containing the string `exp`.
3	Listing???	Match folders whose name begins with `listing` followed by three characters.
4	G*	Match folders whose name begins with g.

TABLE 47 Valid migration exclusions and inclusion formats

Item	Format	Description
5	"G*" "L*	Match folders whose name begins with G and folders whose name begins with L.
6	Windows: \\server\directory\folder UNIX: server:/root/directory	Exclude/include absolute path.

Run customized commands

Specify batch files to run before and after migration criteria evaluation, before and after a replication operation, and before and after a migration operation.

To run batch files before and after migration criteria evaluation (scanning)

1. Perform one of the following actions:
   - Right-click an archival migration policy and click **Properties**.
   - Select an archival migration policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **General** to display a dialog box in the right pane.
3. Select **Run batch file before scan** or **Run batch file after scan**.
4. Type the name of or browse for the batch file that you want to run.
5. Click **OK**.

To run batch files before and after a replication action

- For more information about how to specify scripts to run before and after a replication action of a migration operation, refer to "Run customized commands" on page 219.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Scan for data immediately

Manually scan for directories that match the migration criteria of an archival migration policy. By default, the policy displays a list of migration tasks.

1. Perform one of the following actions:
   - Right-click an archival migration policy and click **Start Scan Now**.
   - Select an archival migration policy to show a **Properties** tab in the right pane. Right-click anywhere on the tab and click **Start Scan Now**.
   - Select the **Archival Migration Policies** folder to show a list of policies in the right pane. Right-click a policy and click **Start Scan Now**.
2. To view a list of migration tasks, select the policy you scanned to show a **Properties** tab in the right pane.

NOTE
If you cleared the policy task option to automatically add migration tasks for migration candidates, the policy adds migration candidates after scanning. This enables you select which candidates you want to migrate. Candidates you do not select for migration are automatically added to the policy exclusion list.

Update the properties of a policy

Update the properties of an archival migration policy.

1. Perform one of the following actions:
   - Right-click an archival migration policy and click **Properties**.
   - Select an archival migration policy and click **Edit** on the **Properties** tab that appears on the right pane.
2. Click a menu item to show a dialog box.
3. Click **OK** when you finish updating the policy.

NOTE
To configure default options for all archival migration policies, click **Options** on the **Tools** menu, and click **Archival Migration** or any item under Archival Migration. Changes you make affect only new policies. Existing policies are not affected.

Delete a policy

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

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

View a list of migration candidates

View the results of a criteria evaluation and select which candidates you want to migrate.

1. Right-click an archival migration policy and click **Properties**.
2. Click **Task Options**.
3. Clear **Automatically add migration tasks for migration candidates**.
4. Click **OK**.
5. Right-click the policy and click **Start Scan Now**.
6. When scanning is complete, select the policy to show a **Properties** tab in the right pane.

7. Use the scroll bar to browse the list of migration candidates.

View properties of a migration candidate

View the source and destination of a migration candidate. The appearance of some properties depends on task options you select (for example, size and size on disk.)

1. Select an archival migration policy to show a **Properties** tab in the right pane.
2. On the **Properties** tab, right-click a migration candidate and click **Properties**.
3. Close the dialog box when you finish viewing the information about the candidate.

NOTE
Calculating size and size on disk might take time, depending on the size of the migration candidates, and it generates network traffic as the information is gathered.

View a list of migration tasks

View all the migration tasks for a policy and a source.

To view all the migration tasks for a policy
Select the policy to show a **Properties** tab in the right pane.

To view the migration tasks for a source
Perform one of the following actions:

- Select the source folder under an archival migration policy to show a **Properties** tab in the right pane.
- Open the source folder to show a list of tasks under the folder.

NOTES:

- If the policy task option to automatically delete completed tasks is selected, the migration task is removed after the task is completed successfully.
- VFM uses visual cues to indicate the state of a migration task. For more information, refer to "View progress of a task" on page 199.

Add a migration task immediately

You can manually add a migration task to an archival migration policy.

1. Select an archival migration policy to show a **Properties** tab in the right pane.
2. Right-click a migration candidate and click **Add Task**. The candidate is removed from the **Migration Candidates** list and a new migration task for the candidate appears under **Tasks**.

To view a list of migration candidates, you must clear the **Automatically add migration tasks for migration candidates task** option by completing the following steps:

1. Right-click the policy and click **Properties**.
2. Click **Task Options**.
3. Make sure that **Automatically add migration tasks for migration candidates** is not selected.

Start a migration task immediately

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

Perform one of the following actions:

- Double-click an archival migration policy, double-click a source to show a list of migration tasks, right-click a task, and click **Start task now**.
- Select an archival migration policy to show a **Properties** tab in the right pane. Under **Tasks**, right-click a task and click **Start task now**.

Additional migration task topics

For more information, refer to the following migration policy sections:

- "Update the properties of a migration task" on page 198
- "Start all tasks for a policy immediately" on page 198
- "View progress of a task" on page 199
- "User action required" on page 199
- "Suspend a migration task" on page 199
- "Remove a task" on page 200

Storage Load Balancing Policies

To create your first load balancing policy, perform the tasks in the following sections. For more information, refer to the following topics:

1. "Create a storage load balancing policy" on page 222
2. "Add storage locations" on page 225
3. (Optional) "Configure the initial phase" on page 226
4. (Optional) "Configure the incremental phase" on page 227
5. (Optional) "Configure the final phase" on page 227
6. "Schedule scanning" on page 228
7. "Select replication options" on page 228
8. "Schedule migration" on page 229

Create a storage load balancing policy

A storage load balancing policy enables you to automatically balance the storage load among a group of storage locations you specify.

1. Perform one of the following actions:
   - On the **Tools** menu, point to **New Policy**, point to **Data Movement Policy**, and click **Storage Load Balancing Policy**.
   - Right-click **Admin View**, point to **New**, point to **Data Movement Policy**, and click **Storage Load Balancing Policy**.

- Under Admin View, right-click the **Data Movement Policies** folder, point to **New**, and click **Storage Load Balancing Policy**.
- Under the **Data Movement Policies** folder, right-click the **Storage Load Balancing Policies** folder and click **New Storage Load Balancing Policy**.

2. Type a word or phrase that best describes this policy. The name appears under the **Storage Load Balancing Policies** folder and on the **Properties** tab of items that are members of the policy.
3. (Optional) Change the default **High-mark**. The High-mark is the maximum percentage of used disk space allowed. When disk use exceeds this limit, the disk is excluded as a destination.
4. (Optional) Change the default **Low-mark**. The Low-mark is the minimum percentage of used disk space allowed. When disk use falls below this limit, the disk is excluded as a source.
5. (Optional) Change the default **Threshold**. The threshold is the maximum allowed variance between the current disk usage (percent full) and the ideal balance. A pool of storage devices in a Load Balancing policy is not balanced when the variance between the current disk space usage level of any drive in the pool containing targets of links and the ideal balance exceeds the threshold. This imbalance triggers a load-balancing operation. The policy reports the current variance in the event manifest logged when the scan phase is complete. The note at the end of this topic provides an example that illustrates the threshold load balancing setting.
6. Select the **Load balance data even when no machine is above high mark** option when you want any device in the pool above the low-mark to be considered as a possible input or source device. When this option is not selected, the policy considers only devices above the high-mark as possible input or source devices.
7. Select the **Automatically add migration tasks for data move suggestions** option when you want the policy to automatically balance storage. When this option is not selected, the policy shows a list of link targets and suggests where to migrate data. You can accept or delete a suggestion. When you accept a suggestion, the policy creates a migration task.
8. Click **Add** to list a root that contains link targets that you want to load balance.
9. Repeat Step 8 to add more roots.
10. In the **Use the Monitor on** list, select the machine that contains the agent that you want to use to process the policy. By default, the current VFM server is used to process the policy. If you installed monitoring agents on other machines, you can select one of the locations where monitoring agents are installed as the location from which this policy is processed.
11. Click **OK**.

NOTE
If you uninstall a monitoring agent, ensure that storage load balancing policies that formerly used the uninstalled agent are reassigned to a currently available agent. Do this by selecting the desired agent through the properties dialog box for the policy.

Threshold load balancing example

The following example illustrates the threshold load balancing setting. In this example a storage load balancing policy monitors the disk space usage of drive C on machineA and machineB. Share1 and Share2 are link targets on machineA and Share3 is a link target on machineB. The ideal balance of the pool was determined to be 41 percent (41%).The policy threshold is set to 30 percent (30%):

1. The policy monitors the disk space of drive C on machineA and machineB and finds that the used disk space on machineA is 80 percent (80%) and the used disk space on machineB is 20 percent (20%). The variance between the percentage used disk space and the ideal balance for machineA is 39 percent (39%) and for machineB it is twenty-one percent (21%). machineA triggers a load-balancing attempt because the variance from the ideal balance (39%) exceeds the threshold.

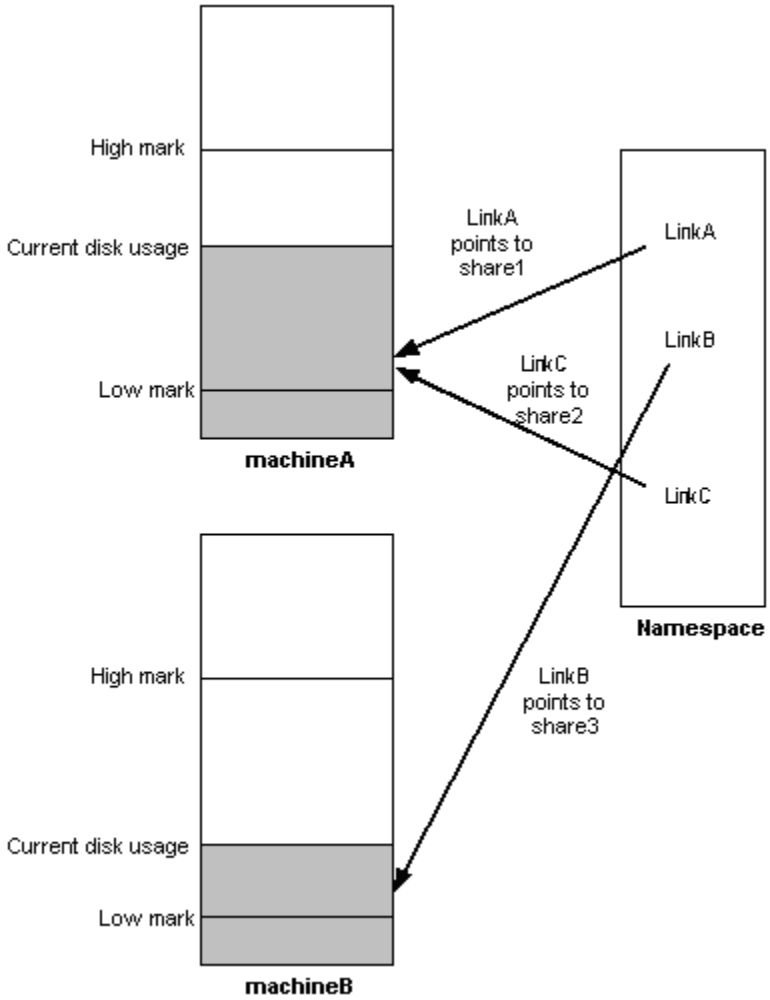

2. The policy evaluates the space usage and determines that to load balance the devices, the data contained in Share1 must migrate to machineB. A migration task is created. When the migration schedule runs, the data contained in Share1 is migrated to the C drive on machineB and linkA pointing to Share1 is retargeted to point to Share1 on machineB.

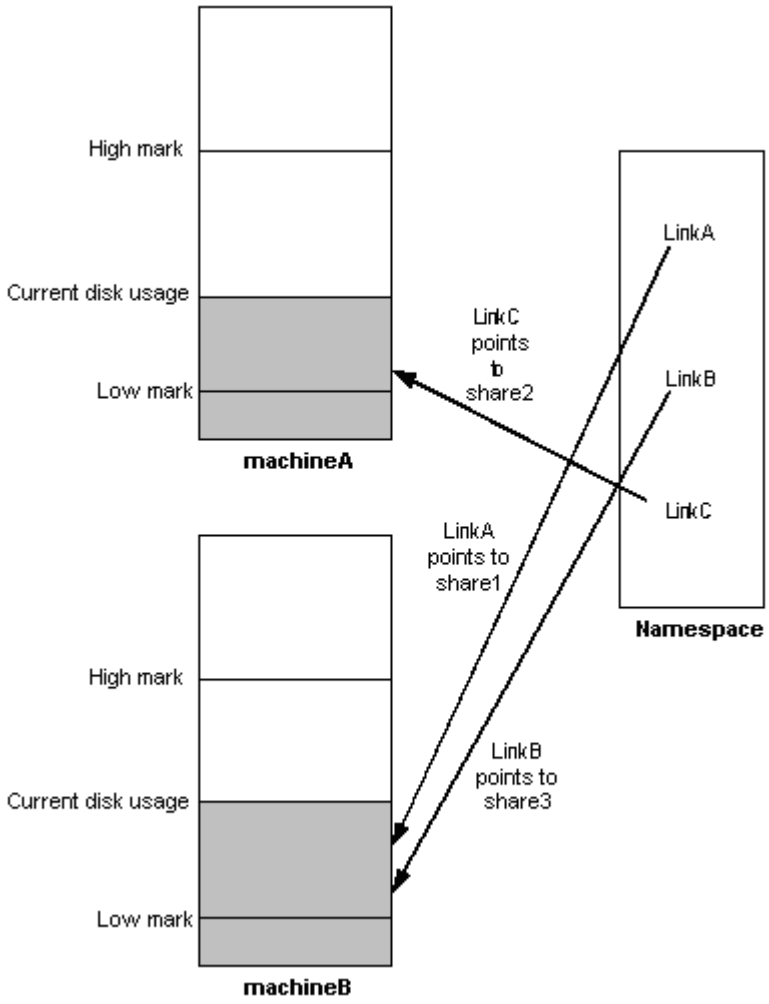

Add storage locations

Add the storage locations that you want the policy to load balance.

1. Right-click the policy and click **Properties**.
2. Click **Pool**.
3. Click **Add** and select a shared folder in the Physical View or a shared folder, volume, or qtree in your network.
4. Repeat Step 3 to add at least one more shared folder.
5. Optionally select whether data is migrated to or from the storage location during load balancing:
   - **Policy decides at run-time** - When the policy evaluates disk space usage, it decides whether the storage location should be an input (source) or output (destination) device.
   - **Input device only** - When the policy evaluates disk space usage, data is always migrated from this storage location. The storage location must be a link target or must contain at least one link target.

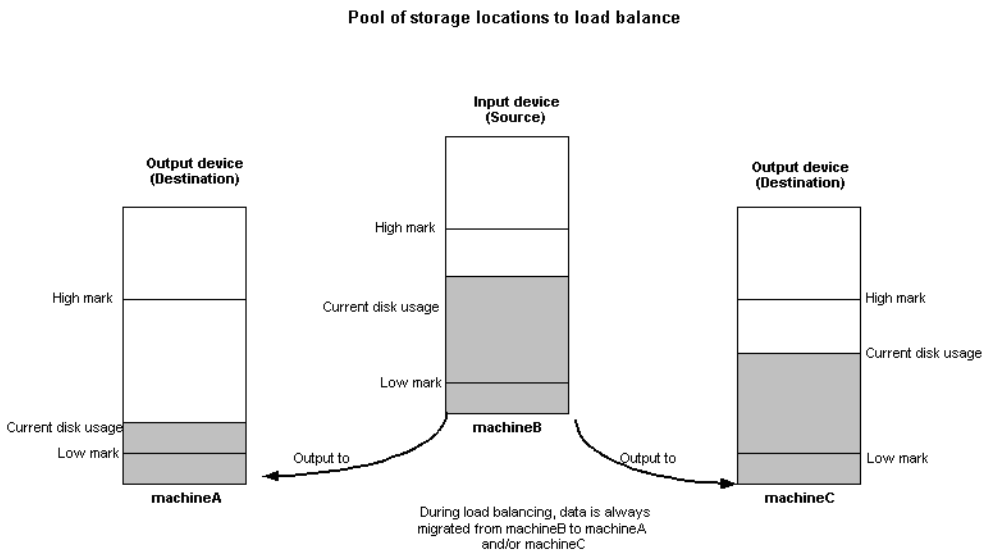

- **Output device only** - When the policy evaluates disk space usage, data is always migrated to the storage location.

6. Click **OK**.

NOTES:

- Storage load balancing policies do not support qtree quotas on pool members.
- When you do not designate any storage location in the pool to be an input (source) device, the storage location with the highest load becomes the input device when the policy load balances the pool.
- When you designate more than one storage location as an input (source) device, only one storage location is selected as the input device when the policy load balances the pool.
- If an output (destination) device exceeds the high-mark specified by the policy, another storage location is selected as the output device. A storage location that exceeds the high-mark cannot be an output device.
- If an input (source) device falls below the low-mark specified by the policy, another storage location is selected as the input device. A storage location that falls below the low-mark cannot be an input device.

Configure the initial phase

You have the option of specifying a script to run before each migration task is executed. The policy automatically scans for potential problems copying the data and performs a baseline copy of the data.

1. Right-click the policy and click **Properties**.
2. Click **Initial Phase**.
3. (Optional) Select the **Run batch file** check box and type the name of a script file, or browse for a script file. VFM provides sample batch files that you can modify. The sample batch files are in the `Examples\Scripts\Migration` directory in the VFM server installation directory.

4. (Optional) Select to pause only if errors occur during the execution of the script, suspend the migration operation after completing the script, or continue the migration operation after completing the script.

NOTES:

- When an action is suspended or paused after an error occurs, you must resume the migration operation manually. For more information, refer to "User action required" on page 199.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Configure the incremental phase

You have the option of copying data incrementally and specifying when you want the migration operation to advance to the final phase. This action affects all the migration tasks associated with the policy. Data is copied every time the incremental schedule runs until the conditions you specify for advancing to the final phase are met.

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

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

4. Specify the condition for advancing migration tasks to the final phase.
   - **Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually** - Migration tasks perform incremental copies indefinitely at the intervals specified by the incremental phase schedule. You must advance migration tasks to the final phase manually. For more information about how to advance migration tasks manually, refer to "To specify the condition for advancing a migration task to the final phase" on page 196.
   - **After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule** - Migration tasks advance to the final phase after each migration task runs the incremental phase one time.
   - **After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule** - Migration tasks advance to the final phase on the date and time you specify and run the final phase the next time the final phase runs after the date and time you specify.
5. Click **OK**.

Configure the final phase

You have the option of specifying a script to run after each migration task is completed.

1. Right-click the policy and click **Properties**.
2. Click **Final Phase**.

3. (Optional) Select **Run batch file** and type the name of a script file, or browse for a script file to run after each migration task is completed. VFM provides sample batch files that you can modify. The sample batch files are in the `Examples\Scripts\Migration` directory in the VFM server installation directory.

NOTES:

- When an action is suspended or paused after an error occurs, you must resume the migration operation manually. For more information, refer to "User action required" on page 199.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Schedule scanning

Schedule how frequently you want the storage load balancing policy to scan the current size of disk space of the storage locations to determine whether load balancing is required. If load balancing is required, the policy suggests directories to migrate.

1. Perform one of the following actions:
   - Right-click a storage load balancing policy and click **Properties**.
   - Select a storage load balancing policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Policy Scan Schedule**.
3. Click **New** and configure a new schedule or select a schedule from the list to update.
4. Click **OK**.

NOTE
During scanning, storage locations that contain parent and child directories that are targets of links are ignored.

Storage location example that contains parent and child directories

A DFS root `\\ilus.com\PriamRoot` is a member of a storage load balancing policy. The root contains the following links:

- `LinkA` targeting directory `\\server\share\folder1`
- `LinkB` targeting subfolder `\\server\share\folder1\subfolder`

When scanning for storage locations to migrate, `LinkA` and `LinkB` are ignored.

Select replication options

Configure replication, file security, and replication differencing options for a storage load balancing policy as well as for other storage policies. For more information, refer to the following migration policy sections:

- "Select replication options" on page 184
- "Specify replication filters" on page 186
- "Select a replication agent" on page 187

- "Select differential replication options" on page 187
- "Select file security replication options" on page 188
- "Select file attribute replication options" on page 190

Schedule migration

Schedule how frequently you want to migrate directories referenced by links in your namespace to load balance the storage pool. Before you configure migration schedules, refer to the information in "Scheduling considerations" on page 171. For more information about how to schedule the migration, refer to the following migration policy sections:

- "Schedule the initial phase" on page 179
- "Schedule the incremental phase" on page 180
- "Schedule the final phase" on page 182

Start scanning immediately

Manually scan the current size of disk space of the storage locations to determine whether load balancing is required. If load balancing is required, the policy shows a list of suggested directory migrations from which you can select the directories you want to migrate. You can manually exclude suggestions, or you can leave rejected suggestions in the list. The current list of suggestions is replaced with a new list when the scheduled policy scan runs or you manually scan again.

1. Perform one of the following actions:
   - Right-click a storage load balancing policy and click **Start Scan Now**.
   - Select a storage load balancing policy to show a **Properties** tab in the right pane. Right-click anywhere on the tab and click **Start Scan Now**.
   - Select the **Storage Load Balancing Policies** folder to show a list of policies in the right pane. Right-click a policy and click **Start Scan Now**.
2. To view a list of migration suggestions, select the policy you scanned to show a **Properties** tab in the right pane.

NOTE
If you selected the policy option to **Automatically add migration tasks for data move suggestions**, migration tasks are added automatically after scanning. To view the migration tasks, double-click the policy to show the **All Tasks folder**.

Run customized commands during replication

Specify batch files to run before and after migration criteria evaluation and before and after a replication operation.

1. Perform one of the following actions:
   - Right-click a migration policy and click **Properties**.
   - Select a migration policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Replication Scripting**.

3. Select the options you want and specify the batch files you want to run. The batch file can be specified either as the UNC path to a file accessible on your network or as a local path to a file residing on the server where VFM monitoring is taking place
4. Click **OK**.

NOTES:

- (Migration, Archival Migration, and Storage Load Balancing policies only) When you specify scripts to run before and after replication, the files run before and after the following migration actions:
  - Scan for potential problems copying the data (Initial Phase).
  - Perform a baseline copy (Initial Phase).
  - Copy the data (Incremental Phase).
  - Do a final replication prior to deleting the source data (Final Phase).
  - Delete the source data (Final Phase).
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Start migration immediately

You can bypass the migration schedule and start migration tasks manually.

To start one migration task

Double-click the **All Tasks folder** under a storage load balancing policy, right-click a migration task and click **Start Migration Now**.

To start all migration tasks for a policy immediately

Perform one of the following actions:

- Right-click a storage load balancing policy and click **Start Migration Now**.
- Select a storage load balancing policy to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Start Migration Now**.
- Select the **Storage Load Balancing Policies** folder to show a **Properties** tab in the right pane. Right-click a policy and click **Start Migration Now**.

View migration suggestions

Normally, a storage load balancing policy scans the current size of disk space of the storage locations to determine whether load balancing is required. If load balancing is required, the policy shows a list of suggested directory migrations.

1. Select a storage load balancing policy to show a **Properties** tab in the right pane. The suggested directory migrations are listed under **Data Move Suggestion**.
2. To view detailed information about a suggestion, right-click the suggestion and click **Properties**.
3. Perform one of the following actions:

- To accept a suggestion and create a migration task to migrate a directory to the suggested location and retarget the link referencing the directory, right-click the suggestion and click **Add Task**. The migration task appear in the **All Tasks** folder under the policy.
- To refuse a suggestion, right-click it and click **Exclude**.

NOTES:

- If you selected the Automatically add migration tasks for data move suggestions policy option, migration tasks are added after scanning.
- A suggestion that is excluded might appear as a suggestion on subsequent policy scans.

Manage migration tasks

To manage migration tasks, refer to the following migration policy sections:

- "Update the properties of a migration task" on page 198
- "View progress of a task" on page 199
- "User action required" on page 199
- "Suspend a migration task" on page 199
- "Remove a task" on page 200

Delete a storage load balancing policy

If a migration task is in progress when you delete a storage load balancing policy, the migration task runs until it is complete.

1. Perform one of the following actions:
   - Under the **Storage Load Balancing Policies** folder, right-click a policy and click **Delete**.
   - Click a policy to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Delete**.
   - Click the *Storage Load Balancing Policies* folder to show a **Properties** tab in the right pane. Right-click a policy and click **Delete**.
2. Click **Yes**.

Chapter

Creating and Managing Recovery Policies

8

This chapter provides detailed instructions about how to create and manage business continuity policies (recovery policies) to make data highly available.

In this chapter

- Introduction 233
- Disaster Recovery Planning 234
- Client Disaster Recovery 235
- Filer Disaster Recovery 255
- Server Disaster Recovery 261

Introduction

VFM provides three types of recovery policies. Before you create a policy, refer to "Disaster recovery planning considerations" on page 20 for more information.

Select the policy that best meets your needs.

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

- Server recovery policies enable you to make data on servers participating in your namespace highly available. Use this policy when data on a primary server must be made highly available. When a failure of the primary server is detected, VFM automatically fails over all links in the policy to a secondary server.
  - This policy enables you to configure a server as a disaster recovery backup for one or more primary servers.
  - You can monitor for new links targeting the primary server to add to the policy and to replicate data to a secondary server.

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

Disaster Recovery Planning

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

Planning considerations

- Links used in recovery policies must reside in Windows 2000 or Windows Server 2003 roots. This is because recovery policies rely on the ability to toggle link targets online and offline.
- When you add a link to a recovery policy using SnapMirror replication, the link should have only two targets. Make sure that one target is the source and the other target is the destination of the SnapMirror. For more information about configuring mirrored qtrees as targets of a client recovery link, refer to "Add a link to a policy" on page 237.
- The account used by the VFM server or monitoring agent, as well as any accounts used to run the VFM client, must have permission to make modifications to the DFS root containing the links in the recovery policy. If the account lacks these permissions, VFM cannot control the online/offline state of the link targets.
- If a DFS root served by a single system is used and becomes unavailable, failover cannot be performed and users can access only the primary storage target until their referral times out. Adding additional servers to host copies of the namespace provides fault tolerance. To create copies of the namespace, you create DFS root replicas in the case of domain-based roots, or you can set up namespace availability policies to synchronize the logical structure of two or more stand-alone roots. An alternate approach for creating a highly available root is to host a stand-alone root on a Windows 2000 cluster server.
- If the VFM server or monitoring agent for a client or filer recovery policy becomes unavailable, neither failover nor data synchronization can be performed. User access to the primary storage target is not affected.
  - VFM communicates with the DFS root to initiate a failover.
  - VFM communicates directly with the storage targets to initiate data synchronization.
- The VFM client does not have to be running for recovery policies to function.
- Carefully plan where each VFM component (for example, client, server, monitoring agent, and replication agents) should be deployed.

- When configuring replication, you should not configure overlapping source and destination directory hierarchies. This can result in unintended repeated duplication of files.
- Roots containing links participating in recovery policies can also be members of namespace availability policies. In this configuration, as with recovery policies containing links from replicated domain-based roots, the recommendation is that you configure the recovery policies for manual failover rather than for automatic failover if the links target read/write data rather than read-only data.
- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.

Resynchronizing SnapMirrors for disaster recovery failback

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

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

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

Client Disaster Recovery

To create a policy, perform the following tasks. For more information, refer to the information in the following sections:

1. "Create a client recovery policy" on page 236
2. "Add a link to a policy" on page 237

3. "Schedule target monitoring" on page 240
4. "Schedule data replication" on page 240

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

Create a client recovery policy

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

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

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Business Continuity Policy**, and click **Client Recovery Policy**.
   - Drag a link to the **Client Recovery Policies** folder.
   - Right-click the **Business Continuity Policies** folder in the Admin View, point to **New**, and click **Client Recovery Policy**.
   - Right-click the **Client Recovery Policies** folder and click **New Client Recovery Policy**.
   - Right-click **Admin View**, click **New**, point to **Business Continuity Policy**, and click **Client Recovery Policy**.
2. Type a word or phrase that best describes this policy in the **Name** field. The name appears under the **Client Recovery Policies** folder and on the **Properties** tab of items that are members of the policy.
3. Select a replication method. Use VFM replication when you are using heterogeneous devices.
4. Select **Automatically failover** when you want to automatically make a standby target the active target when the active target is not available.
5. Select **Failover all links** when you want to automatically switch to the standby target of all the links in the policy when any active target that is a member of the policy is not available. This option is automatically selected when the replication method is SnapMirror.
6. Select **Update associated policies on failover** when links in the client recovery are part of a multiprotocol namespace or namespace availability policy, and you want to create new automount and export files referencing the new active targets after a failover.
7. Select a link in the **Only monitor link** list if you want to monitor only the link selected. This option is not available if you selected SnapMirror replication.
8. You can select another monitoring location where monitoring of the targets of the client recovery policy is performed. The default is to monitor from the VFM server. From the **Use the Monitor on** list, select the host name of another monitoring agent.

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

9. Click **OK**.

Add a link to a policy

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

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

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

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

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

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

The targets of a link in a recovery policy are sorted in alphabetical order and set offline. You must specify which target is the active or online target and sort the targets in the order in which you want them to fail over. To change the order in which targets fail over, refer to "Prioritize the targets of a link" on page 254 for more information.

To add a link with multiple targets

Perform one of the following actions:

- Drag a link with multiple targets from the Logical View tree to the policy.
- Drag a link with multiple targets from the **Properties** tab of an item in the Logical View to the policy.
- Right-click a policy under the **Business Continuity Policies** folder and click **Add Link to policy**. Select a link with multiple targets. Click **OK**.
- Click the **Business Continuity Policies** folder to display a **Properties** tab in the right pane. Right-click a policy and click **Add Link to policy**. Select a link with multiple targets. Click **OK**.

To add a link with a single target

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

Perform one of the following actions:

- Drag a single target link from the Logical View to the policy.
- Drag a single target link from the **Properties** tab of an item in the Logical View to the policy.
- Right-click the policy and click **Add Link to policy**. Select a single target link. Click **OK**.
- • Click the **Business Continuity Policies** folder to display a **Properties** tab in the right pane. Right-click a policy and click **Add Link to policy**. Select a single target link. Click **OK**.

Add multiple links to a policy

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

- You can add multiple links to a policy by dragging the selected DFS links to the policy.
- The targets of a link in a recovery policy are sorted alphabetically. The first target becomes the active or online target and all other targets of the link are placed offline. To change which target is active, refer to "Prioritize the targets of a link" on page 254 for more information.

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

1. Drag a logical folder or DFS root to the policy.
2. Perform one of the following actions:
   - Right-click a client recovery policy and click **Properties**.
   - Select a client recovery policy to show a **Properties** tab in the right pane. Click **Edit**.
3. Click **Search Options**.
4. Select any of the following options:
   - **Automatically add newly found links** - Search for new links and automatically adds them to the policy.
   - **Automatically remove links from the policy when links are removed from a logical folder or DFS root** - Remove links from the policy that are removed from a logical folder or root that is a member of the policy.
   - **Automatically set the first link target online** - Set the first target of a link to always be online. To see the targets of a link, click the link. The list of targets appears in the **Properties** tab under **Target**.
   - **Ignore single-target links** - Do not search for links with only one target.
5. Optionally, add items to the **Ignore List** that you do not want to include in the policy. When the policy searches for new links in logical folders and roots, it ignores the items on the list.

6. Click **OK**. When the policy runs, it searches for links that match the search criteria. If you selected the option to Automatically add newly found links, the links found are added to the policy. If you did not select this option, the links are listed under **New Resources** on the **Properties** tab of the policy. You have the option of adding the links to the policy or the Ignore list.

NOTES:

- Select a policy to view the **Properties** tab in the right pane.
- To set default search options for all new client recovery policies, click **Options** on the **Tools** menu and click **Search Options** under **Disaster Recovery**.

To add selected links

1. In the Logical View, click a root or folder to display its properties in the right pane.
2. Perform one of the following actions:
   - Under **Links** on the **Properties** tab of a root, select multiple links and drag them to a client recovery policy.
   - Under **Children** on the **Properties** tab of a folder, select multiple links and drag them to a client recovery policy.
3. If some or all the links you selected have single targets, the Disaster Recovery Link Target Wizard provides you with additional options. When you finish going through the wizard, a status report appears.
4. If all the links you selected have multiple targets, a status report appears.

Add a share to a client recovery policy

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

- If the shares participate in the namespace, the links that point to the shares are added to the policy.
- If the shares do not participate in the namespace, the Disaster Recovery Link Target wizard assists you in creating links.
- If a share is the only target of a link, the Disaster Recovery Link Target wizard assists you in adding another target to the link.

Select search options

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

1. Select **Automatically add newly found links** when you want to search for new links and automatically add them to the policy.
2. Select **Automatically remove links from the policy** when links are removed from a logical folder or DFS root that is a member of the policy when you want to remove links from the policy that are removed from a logical folder or root that is a member of the policy.

3. Select **Automatically set the first link target online** when you want the first target of a link to always be online. To see the targets of a link, click the link. The list of targets appears in the **Properties** tab under **Target**.
4. Select **Ignore single-target links** when you do not want to search for links with only one target.

Schedule target monitoring

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

- To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.
- Status monitoring of recovery policies is performed in the VFM server or monitoring agent. Recovery policies are monitored even if the VFM console is closed.
- When link targets are monitored, VFM checks whether files in the target directory can be enumerated.
- Links in a recovery policy show **unknown** status until monitoring runs at least once for the policy.

To schedule target monitoring

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

Schedule data replication

Schedule data replication for a recovery policy using VFM replication.

- Before you schedule data replication, review your network bandwidth and infrastructure. For example, review the rate of data changes and link speeds.
- You to set the data transfer rate for a machine. In the Physical View tree, right-click a machine, click **Properties**, and click the **Agent** tab.
- To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.

To schedule data replication

1. Perform one of the following actions:
   - Right-click a recovery policy and click **Properties**.

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

2. Click **Replication Schedule**.
3. Click **New** and configure the schedule.
4. Select **Use alternate schedule** when the policy fails and type a frequency (in minutes) to run the schedule any time replication fails. Any time a policy runs and fails, attempts are made to run the policy at a stated frequency until the policy runs successfully.
5. Click **OK**.

Select a SnapMirror

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

1. Perform one of the following actions:
   - Right-click a client recovery policy and click **Properties**.
   - Select a client recovery policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **SnapMirror Replication**.
3. Click **Select a SnapMirror**.
4. Perform one of the following actions:
   - Select a filer from the list provided. The list includes all the NetApp filers in the Physical View. If the filer contains SnapMirrors and they do not appear, click **Refresh Filer SnapMirror Information**. Click **Done**. Select an existing SnapMirror.
   - To add a new SnapMirror, click **Create**. Type the source and destination information and click **OK**. Click **Done**.
5. Click **OK**.
6. To break the source/mirror relationship when a failover occurs, select **Break SnapMirror on Failover**. When the SnapMirror is broken, the destination data becomes writable. If you do not select this option, the destination data remains read-only after a failover. A quiesce is performed before the SnapMirror is broken.
7. Click **OK**.

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

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

Failover and replication options

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

Failover options

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

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

TABLE 48 Policy failover options

Client recovery policy

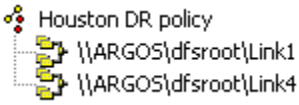

Link1

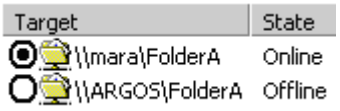

Target	State
\\mara\FolderA	Online
\\ARGOS\FolderA	Offline

Link4

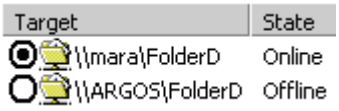

Target	State
\\mara\FolderD	Online
\\ARGOS\FolderD	Offline

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

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

Replication options

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

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

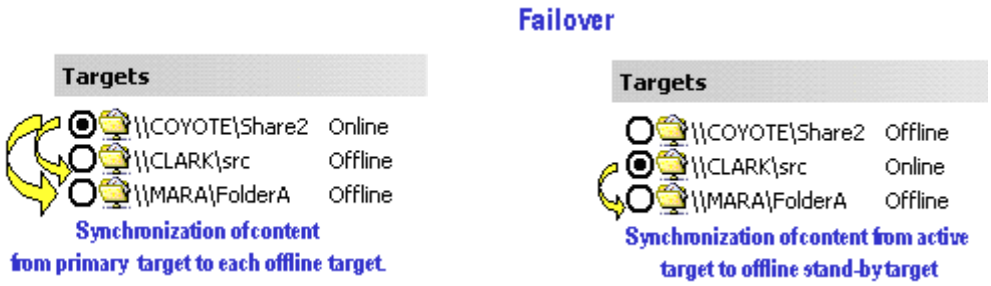

- **Replicate all targets.** The content of the online target is synchronized to each standby target, regardless of the order in which the targets are arranged on the **Properties** tab of the link that contains the targets. Replication is from the active target to all the standby targets.

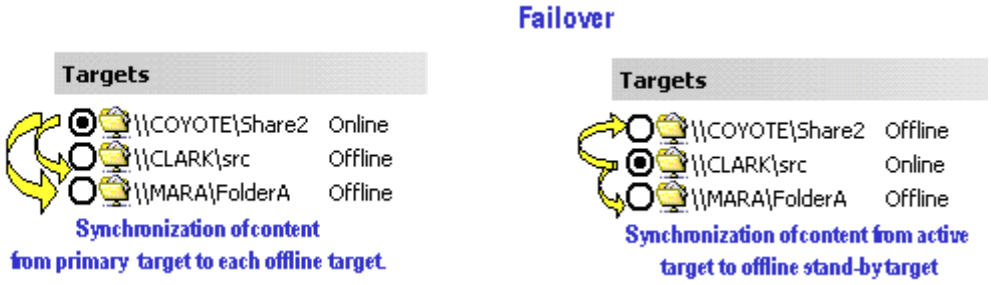

- You have the option to synchronize the contents of the active target with the standby target when a failover or failback is initiated.

NOTE
If a client recovery policy is using SnapMirror replication, refer to "Disaster Recovery Planning" on page 234 for more information about how to resynchronize SnapMirrors for disaster recovery failback.

Select replication options

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

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

1. Perform one of the following actions:
   - Right-click a policy and click **Properties**.
   - Select a policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Replication Options**.

3. Click **Optimize For Speed** if you want to automatically set replication options when the time it takes to transfer the data is very important.
4. Click **Optimize For Security** if you want to automatically set replication options when the completeness and accuracy of the data transfer is very important.
5. Clear Include subfolders if you do not want to copy the source target subfolders and their contents to the destination targets.
6. Select **Delete orphaned files in destination folders** if you want to delete files in the destination targets that do not have a corresponding file in the source target.
7. If the destination targets do not support auxiliary file streams supported by the source target, you can select **Allow loss of additional file streams** if you want to copy files to the destination targets that lose auxiliary streams.
8. When you select **Enable Copy-in-place**, the source file directly replaces the destination file. This decreases the time required to perform a replication operation. When this option is not selected, the source file is copied to a temporary file in the destination directory, renames the current destination file, renames the temporary file to the actual file name, and deletes the destination file. This minimizes the chance of network errors adversely affecting file contents.
9. VFM can take advantage of the NTFS Change Journal on the replication source to determine which files are to be replicated rather than scanning the entire source to detect files that have changed. This option offers time savings when you are replicating large directory trees with limited changes scattered throughout the directory hierarchy. Review the following considerations:
   - The NTFS Change Journal is supported on Windows 2000 NTFS, which is available on Windows 2000 and later versions of the operating system.
   - When you enable the NTFS Change Journal on a volume, the size of the volume is used to determine the amount of space dedicated to the change journal. For each 50 GB of space on the volume, 128 MB is assigned for use by the change journal up to a maximum of 10 GB. On Windows 2000 SP2 or below, the amount of space dedicated to the change journal is always 128 MB, regardless of the size of the volume.
   - Use of this option requires that VFM deploy a replication agent on the machine where the source data resides and enable the Change Journal on the NTFS volume hosting the source data. The machine properties dialog box enables you to control whether the replication agent can be deployed to a machine and whether the NTFS Change Journal can be enabled for volumes on the machine.
   - Use of this option means that changes made on the source are replicated to the destination and changes made on the destination are not overwritten by data from the source unless the corresponding file on the source changes as well. This differs from standard VFM replication, wherein changes made on the destination are detected and overwritten with data from the source.
   - In cases where the NTFS Change Journal wraps, a full scan of the source and destination must be performed to determine what files need to be transferred.
10. When **Use snapshot data** is selected, a snapshot or a shadow copy is taken of the volume in which the source data is located and the snapshot is used as the source. This reduces the possibility of not being able to transfer source files that are held open for exclusive access by an external application at the time of replication. The source data must be on a NetApp filer (CIFS or NFS protocol) or a machine running the Windows Storage Server 2003 or Windows Server 2003 file system.

To reduce the number of snapshots that must be created, try to group sources in the same volume under the same policy. When **Use snapshot data** is selected, you cannot enable NTFS Change Journal or file differential replication using cached hash information because snapshot data resides on read-only volumes and NTFS Change Journal and file differential replication require that the data be on writable volumes.

11. When **Override comparison timestamp resolution** is selected, a one second resolution is used when comparing time stamps between the source and destination. By default, the lowest resolution between the source and destination file systems is used when comparing time stamps.
12. In the **Copy files** list, select one of the following options:
    - **If destination file is different or missing** - Replace files in the destination targets if the time stamp does not match the time stamp of files in the source target, and copy the files in the source target that are not found in the destination targets.
    - **If destination file is older or missing** - Replace files in the destination targets that have time stamps older than files in the source target, and copy the files in the source target that are not found in the destination targets.
    - **If destination file is missing** - Copy only files in the source target that are not found in the destination targets.
13. Type the number of times and the frequency, in seconds, that you want the policy to try to open a file.
14. Type how long, in minutes, a policy should run unsuccessfully before it is aborted.
15. Select one of the following options in the **Events Details** list:
    - **Only list files that encounter errors** - Provide a detailed report of the files that encounter errors during the replication process.
    - **Don't list files** - Do not provide a report of the files included in the replication process.
    - **List all files** - Provide a detailed report of all the files included in the replication process.

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

16. Select one of the following options in the **Post-Failover** list:
    - **Replicate all targets after failover** - Synchronize the content of the online target to each standby target regardless of the order in which the targets are arranged on the **Properties** tab of the link that contains the targets. Replication is from the active target to all the standby targets.
    - **Replicate only to lower priority targets after failover** - Synchronize the content of the online target to each standby target in the order in which the targets are arranged on the **Properties** tab of the link that contains the targets. Replication is from the active target down to the standby targets.
17. Click **OK**.

Select other replication options

Configure file security, replication differencing, and file attributes options for a client recovery policy as well as other storage policies. For more information, refer to the following migration policy sections:

- "Specify replication filters" on page 186
- "Select a replication agent" on page 187
- "Select file security replication options" on page 188
- "Select differential replication options" on page 187
- "Select file attribute replication options" on page 190

NOTES:

- NFS data transfer is not available for client recovery policies.
- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result is different file permissions between the source and destination.

Failover

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

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

Set up automatic failover

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

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

1. Perform one of the following actions:
   - Right-click a client recovery policy and click **Properties**.
   - Click a client recovery policy and click on the **Properties** tab that appears in the right pane.
2. Click **General** and select **Automatically fail over**.
3. Click **OK**. All the links contained in the client recovery policy fail over automatically.

Manually fail over one link

The following tasks guide you through the process.

To manually fail over a link displayed in the client

1. Click a link under a client recovery policy or click the link in the Logical View tree. The **Properties** tab for the link appears in the right pane.
2. Under **Target**, select a target below the active target.
3. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active target to the target you selected. Failover is not performed. This enables you to synchronize data before you make an offline target the active target.
   - Click **Change to fail over** to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.

To manually fail over a link using the right-click menu

1. Right-click the link under a recovery policy or right-click the link in the Logical View tree, and click **Properties**.
2. Select a target below the active target.
3. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active target to the target you selected. Failover is not performed. This enables you to synchronize data before you make an offline target the active target.
   - Click **Change** to fail over to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.
4. Click **OK**.

To manually fail over a link under a client recovery policy

1. Click a client recovery policy. The **Properties** tab for the policy appears in the right pane.
2. Under **Links**, right-click a link and click **Failover**.
3. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active target to the target below the active target. Failover is not performed. This enables you to synchronize data before you make an offline target the active target.
   - Click **Change** to fail over to the target below the active target without replicating contents. Select this option if you are sure that the data contained in the offline target is synchronized with the active target.

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

Manually fail over all links

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

1. Perform one of the following actions:
   - Right-click a client recovery policy and click **Failover all Links**.
   - Click a client recovery policy to display the **Properties** tab in the right pane. Under **Links**, select all the links, right-click, and click **Failover**.
2. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active targets to the standby target. Failover is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.
   - Click **Change** to fail over to the standby targets without replicating contents. Select this option if you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

Manually fail back a link

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

To manually fail back a link in a client recovery policy

1. Click a client recovery policy to display the **Properties** tab in the right pane.
2. Under **Links**, right-click a link and click **Failback**. Failback is performed to the standby target above the active target in the order in which the targets are arranged.
3. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active target to the target above the active target. Failback is not performed. This enables you to synchronize data before you make an offline target the active target.
   - Click **Change** to fail back to the target above the active target without replicating contents. Select this option if you are sure that the data contained in the offline target is synchronized with the active target.

To manually fail back a link using the right-click menu

1. Right-click the link under a recovery policy or right-click the link in the Logical View tree, and click **Properties**.
2. Select a target above the active target.
3. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active target to the target you selected. Failback is not performed. This enables you to synchronize data before you make an offline target the active target.
   - Click **Change** to fail back to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.
4. Click **OK**.

To manually fail back a link displayed in the client

1. Click a link under a client recovery policy or click the link in the Logical View tree. The **Properties** tab for the link appears in the right pane.
2. Under **Target**, select a target above the active target.
3. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active target to the target you selected. Failback is not performed. This enables you to synchronize data before you make an offline target the active target.
   - Click **Change** to fail back to the target selected without replicating contents. Select this option if you are sure that the data contained in the target you want to activate is synchronized with the active target.

Manually fail back all links

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

1. Perform one of the following actions:
   - Right-click a client recovery policy and click **Failback all Links**.
   - Click a client recovery policy to display the **Properties** tab in the right pane. Under **Links**, select all the links, right-click, and click **Failback**.
2. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active targets to the standby target. Failback is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.
   - Click **Change** to fail back to the standby targets without replicating contents. Select this option if you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

NOTE
If a client recovery policy is using SnapMirror replication, refer to "Disaster Recovery Planning" on page 234 for more information about how to resynchronize SnapMirrors for disaster recovery failback.

Run customized commands during failover processing

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

1. Perform one of the following actions:
   - Right-click a recovery policy and click **Properties**.
   - Select a recovery policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Failover Scripting**.

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

NOTES:

- When you select **Run On Manual Failover**, the scripts you specify to run before and after a failover will also run after a manual failback.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Run customized commands during replication

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

1. Perform one of the following actions:
   - Right-click a client recovery policy and click **Properties**.
   - Select a recovery policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Click **Replication Scripting**.
3. Select the options you want and type or locate the batch files you want to run.
4. Click **OK**.

For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Add items to an ignore list

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

1. Perform one of the following actions:
   - Right-click a client recovery policy and click **Properties**.

- Select a client recovery policy to show a **Properties** tab in the right pane. Click **Edit**.

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

NOTES:

- The policy only searches for resources when a client recovery policy contains logical folders and/or DFS roots.
- You can also add links to the Ignore list from the **New Resources** list of a policy. For more information, refer to "Review the list of new resources" on page 251.

Review the list of new resources

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

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

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

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

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

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

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

Update properties of a client recovery policy

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

1. Perform one of the following actions:
   - Right-click a client recovery policy and click **Properties**.
   - Click a client recovery policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Make the necessary changes and click **OK**.

Delete a client recovery policy

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

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

Add new shares to a link in a client recovery policy

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

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

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

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

To create new shares to add to a recovery policy

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

6. *If you want to add shares to a new DFS link*, complete the following steps:
   a. Add the new shares as targets to a new DFS Link
   b. Click **Next**.
   c. Select the DFS root that hosts the link and click **Next**.
   d. Optionally, select a folder in which to group the link.
   e. Click **Next**.
   f. If you do not want to create more links, click **Next**.
7. *If you want to add shares to an existing DFS link*, complete the following steps:
   a. Add the new shares as targets to an existing DFS Link
   b. a. Click **Next**.
   c. b. Select the link that targets the shares and click **Next**.
8. Select the client recovery policy of which you want the link to be a member and click **Next**.
9. Click **Finish**. The shares are created and added to the link and the link appears under the recovery policy folder.

Compare target contents in client recovery policy

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

1. Open the **Business Continuity Policies** folder under Admin View.
2. Double-click a client recovery policy.
3. Perform one of the following actions:
   - Right-click a link under the recovery policy, point to **Reports**, point to **Physical**, and click **File Differences**.
   - Click a link under the client recovery policy to display its **Properties** tab. Select all the targets listed under **Targets**. Right-click, point to **Reports**, then click **File Differences**.

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

Add shared folder replicas to a link

Add multiple copies of a shared folder to a DFS link to provide fault tolerance. The target of a link must be one of the following types:

- 1 to N shared folders
- 1 to N stand-alone DFS paths
- Only 1 domain-based DFS path

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

1. Open the **Business Continuity Policies** folder in the Admin View, right-click a link under a policy, and click **Add replica**.
2. Under **Send the user to one of these share folders or directories**, perform one of the following actions:
   - Type the address of a folder.
   - Click **Browse** to find a folder on the network. Expand the tree for the domain or click the domain. Double-click a machine to display a list of shared folders. Click **OK**.
3. To add another replica, click **Add** and repeat Step 2.
4. Click **OK**. A window displays the status for each shared folder to be added.

NOTES:

- The contents of a share folder replica set are synchronized only if the replica set is a member of a recovery or replication policy.
- For shared folder replica sets that are not a member of a recovery or replication policy, you must either rely on other mechanisms when possible or manually ensure that the contents of the shared folders stay synchronized.
- Use the Ctrl or Shift key to select multiple link targets at one time.

Prioritize the targets of a link

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

NOTE
Make sure that the active or online target appears as the top or is the first target. For more information about how links are failed over in recovery policies, refer to "Failover options" on page 242 and "Replication options" on page 242.

---

1. Select the target which should be active or online.
2. Specify the order in which you want the targets to fail over. Perform one of the following actions:
   - Right-click the link under the recovery policy or right-click the link in the Logical View tree and click **Properties**. Arrange the targets by dragging them to their new location, or right-click a target and click **Move up** or **Move down**. Click **OK**.
   - Click the link under the recovery policy or click the link in the Logical View tree. The **Properties** tab appears in the right pane. Under **Targets**, arrange the targets by selecting the target you want to move and clicking the up or down arrow. Click **Done**.

Remove a link from a policy

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

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

View information about replication operations

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

- To view information about replication operations for all policies, select the **Business Continuity Policies** folder and click the **Replication Tasks** tab in the right pane.
- To view information about replication operations for one policy, select the policy and click the **Replication Tasks** tab in the right pane.

Start replication immediately

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

Perform one of the following actions:

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

NOTES:

- This menu item is disabled if a data replication operation is in progress.
- To view information about any replication in progress, click the **Replication Tasks** tab in the right pane, and then click the blue circle.
- When you configure a remote monitoring agent for a client recovery policy, there is a short time lag between when you make the selection and when you can manually initiate replication on the policy through the **Start Replication Now** option. Attempts to initiate replication during this window return errors.

Cancel a replication in progress

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

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

NOTE
To view information about any replication in progress, click the **Replication Tasks** tab, and then click the blue circle.

Filer Disaster Recovery

A wizard steps you through the process of creating a filer recovery policy. To create a policy, follow the steps in "Create a client recovery policy" on page 236.

Create a filer recovery policy

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

NOTES:

- Before you begin, make sure that the destination filer contains volumes.
- If you are creating a new qtree SnapMirror, the destination qtree cannot exist.
- VFM supports vfiler as the source but not as the destination.

To create a filer recovery policy

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Business Continuity Policy**, and then click **Filer Recovery Policy**.
   - Right-click the **Business Continuity Policies** folder in the Admin View, click **New**, and then click **Filer Recovery Policy**.
   - Right-click the **Filer Recovery Policies** folder and click **New Filer Recovery Policy**.
   - Right-click **Admin View**, click **New**, point to **Business Continuity Policy**, and then click **Filer Recovery Policy**.
2. Click **Next**.
3. Type a word or phrase that best describes this policy. The name appears under the **Business Continuity Policies\Filer Recovery Policies** folder and on the **Properties** tab of items that are members of the policy.
4. Type the name of the destination filer or click **Browse** to select a filer known to VFM.
5. Optionally select the name of a volume on the destination filer.
6. By default, Break SnapMirror on failover is not selected. When a failover occurs, the destination data remains read-only. Select to break the source/mirror relationship when a failover occurs, if you want the destination data to be writable. A quiesce is performed before the SnapMirror is broken.
7. Select **Update associated policies on failover** if links in the filer recovery are part of a multiprotocol namespace or namespace availability policy, and you want to create new automount and export files referencing the new active targets after a failover.
8. By default, **Automatically failover** is selected. All the links targeting data on source volumes and qtrees are monitored. When the active target for a link is not available, VFM automatically switches to a standby target of the link. Status monitoring must be scheduled. If you clear this option, you are notified when an active target is not available.
9. By default, the current VFM server is used to process the policy. If you installed monitoring agents on other machines, you can select one of the locations where monitoring agents are installed as the location from which this policy is processed.

10. Click **Next**. A list of the DFS roots in your namespace appears.
11. Optionally remove a DFS root from the list that does not contain links targeting data on source NetApp filers. During the failover process, users accessing data targeted by links in the selected roots are redirected to the standby secondary target on the destination filer.
12. Click **Next**.
13. Click **Add** and select the source volumes and/or qtrees.
14. Click **Next**.
15. If a destination for a SnapMirror source does not exist, click the down arrow under the destination column and select **Add New SnapMirror**.
16. Configure the destination for the SnapMirror and click **OK**. The source and destination are mirrored.
17. Repeat Steps 15 and 16 for each SnapMirror source that does not have a destination.
18. Click **Next**.
19. Select **Find new Links** if you want to monitor for new links targeting data on the source volumes.
20. Select **Find new QTrees** if you want to monitor for qtrees added to a source filer.
21. Specify how often you want to monitor for new resources.
22. If the policy fails to run at the scheduled time, you are given the option of running the schedule again at a stated frequency. Select **Use alternate schedule when the policy fails** and type how frequently you want the schedule to run.
23. If you want to allow a link to fail a specified number of times before failing over to the standby target, select **Do not failover until link fails at least...times** and type how often it can fail. This option enables you to avoid a failover in response to transient errors.
24. Click **Next**.
25. Specify how often you want to monitor the availability of the volumes and qtrees that are member of the policy. If you selected Automatically failover, you must schedule status monitoring of links.
26. Optionally select **Use alternate schedule when the policy fails** if you want to attempt to run the policy at a stated frequency until the policy runs successfully.
27. Click **Next**.
28. Optionally specify batch files to run during disaster recovery failover processing. Separate entries are provided for scripts to run before and after the link targets are failed over. Type the path of the batch file or click **Browse** to locate the file in your network.
29. Optionally select **Run On Manual Failover** if you want to run a batch file when you manually failover link targets.
30. Click **Next**.
31. Click **Finish**. The policy checks for links targeting the source and automatically creates shares on the destination.

NOTES:

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

- For more information about how to resynchronize SnapMirrors for disaster recovery failback, refer to "Disaster Recovery Planning" on page 234.

Add volumes and qtrees to the policy

Add volumes and qtrees to an existing filer recovery policy.

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

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

Fail over immediately

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

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

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

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

Perform one of the following actions:

- Right-click the filer or volume under a filer recovery policy folder and click **Failover All**.
- Select the filer or volume to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Failover All**.
- To manually fail over the links targeting data on a qtree participating in a filer recovery policy, right-click the qtree and click **Failover**.

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

NOTE
For more information about how to resynchronize SnapMirrors for disaster recovery failback, refer to "Disaster Recovery Planning" on page 234.

Run customized commands

Specify batch files to run during filer recovery failover processing. Separate entries are provided for scripts to be run before and after the link targets are failed over. For more information, refer to "Run customized commands during failover processing" on page 249.

Review the list of new qtrees

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

To add a new qtree to the policy

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

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

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

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

NOTE
To review a filer recovery ignore list, refer to "Update the policy ignore list" on page 259 for more information.

Update the policy ignore list

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

1. Perform one of the following actions:
   - Right-click a filer recovery policy and click **Properties**.

- Select a filer recovery policy to show a **Properties** tab in the right pane. Click **Edit**.

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

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

- When you add a qtree on a volume with multiple qtrees, the qtrees that were not selected are added to the ignore list.
- When the policy scans for new qtrees and you select a group of qtrees to add to the policy, the new qtrees not added to the policy are automatically added to the ignore list.

View the status of a SnapMirror

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

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

View links targeting data on a filer

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

1. Under a filer recovery policy, select a filer, volume, or qtree to show a list of links on the **Properties** tab that appears in the right pane.
2. Double-click a link to view its properties.
   - View information about replication operations
   - You can view information about replication operations for all policies or just one policy. By default, the information for all replication operations appears. You can filter the information to see only replication operations with problems or only replication operations in progress.

- To view information about replication operations for all policies, select the **Business Continuity Policies** folder and click the **Replication Tasks** tab in the right pane.
- To view information about replication operations for one policy, select the policy and click the **Replication Tasks** tab in the right pane.

Update properties of a policy

Update the properties of a filer recovery policy.

1. Perform one of the following actions:
   - Right-click a filer recovery policy and click **Properties**.
   - Click a filer recovery policy and click **Edit** on the **Properties** tab that appears in the right pane.
2. Make the necessary changes and click **OK**.

Remove a filer, volume, or qtree from a policy

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

Perform one of the following actions:

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

Delete a filer recovery policy

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

1. Perform one of the following actions:
   - Right-click the policy and click **Delete**.
   - Select the policy to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Delete**.
2. Click **Yes**.

Server Disaster Recovery

To create a server recovery policy, perform the following tasks. For more information, refer to the following topics:

1. "Create a server recovery policy" on page 261
2. "Select resource monitoring options" on page 262
3. "Schedule resource monitoring" on page 263
4. "Schedule target monitoring and replication" on page 263

Create a server recovery policy

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

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

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

NOTES:

- The account under which the VFM server is running must have the appropriate permissions to create shares and subdirectories on the secondary server.
- Server recovery policies do not use SnapMirror replication. If your primary and secondary servers are filers, it is recommended that you create a filer recovery policy.
- If you uninstall a monitoring agent, ensure that recovery policies that formerly used the uninstalled agent are reassigned to a currently available agent.

Select resource monitoring options

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

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

- **Add secondary target, use the primary target share as a template**. A share is created on the secondary server with the same name, security settings, and local path as the target on the primary server.
- **Add secondary target, create a new directory and share**. Type a local path and a name for the share is created. Select this option when you want to create a secondary target with a local path that is different from the local path of the primary target. VFM uses the share name of the primary target. If the share already exists on the secondary target, the existing share is used. If the existing share contains files, the files are replaced during the replication process.

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

Schedule resource monitoring

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

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

Schedule target monitoring and replication

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

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

To schedule data replication on the server recovery policy, refer to the following client recovery policy topics:

- "Schedule target monitoring" on page 240
- "Schedule data replication" on page 240

Select replication options

Configure replication, file security, and replication differencing options for a server recovery policy as well as for other storage policies. For more information, refer to the following migration policy topics:

- "Select replication options" on page 184
- "Specify replication filters" on page 186
- "Select a replication agent" on page 187
- "Select file security replication options" on page 188
- "Select differential replication options" on page 187
- "Select file attribute replication options" on page 190

NOTES:

- NFS data transfer is not available for server recovery policies.
- When transferring data between NTFS source and destination, you need to think about inherited permissions and how they will affect permissions on the source and destination. If the source and destination are inheriting different permission settings from their parents, the result will be different file permissions between the source and destination.

Add links to a server recovery policy

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

To add links to the Active Links list

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

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

NOTE
If you want to automatically add these links to the **Active Links** list, select **Automatically add conformant links to the Active Links list** on the Resource Monitoring Options dialog box. For more information, refer to "Select resource monitoring options" on page 262.

To repair links in the Nonactive Links list

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

- To view information about the error status of a link, click the **Status** tab for the policy.
- To update the links, perform one of the following actions:
  - Double-click the path of a link in the **Status** tab.
  - Right-click a link in the **Properties** tab and click **Properties**.
  - Double-click a link in the **Properties** tab then click **Edit**.

NOTE

You can configure the policy to remove extra targets on links found that contain targets on the primary and secondary servers. You can also configure the policy to add a second target to links found that do not contain a target on the secondary server. For more information, refer to "Select resource monitoring options" on page 262.

To drag a link to a server recovery policy

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

NOTES:

- When a link in the **Nonactive Links** list no longer has an error status, you can drag the link to the **Active Links** list.
- If you want to remove a link from the **Active Links** list, perform one of the following actions:
  - Drag the link to the **Nonactive Links** list.
  - Right-click the link and click **Move to Nonactive Links** list.
- You can configure the policy to automatically remove links you drag to the policy that do not have a target on the primary server. For more information, refer to "Select resource monitoring options" on page 262.

Start monitoring for new links immediately

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

Perform one of the following actions:

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

Assign a link for the policy to monitor

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

1. Perform one of the following actions:
   - Right-click a server recovery policy and click **Properties**.
   - Select a server recovery policy to show a **Properties** tab in the right pane and click **Edit**.
2. On the General dialog box, select a link from the **Monitor link** list.
3. Click **OK**.

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

Cancel resource monitoring in progress

You can cancel resource monitoring that is in progress.

To check whether monitoring is in progress

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

To cancel resource monitoring in progress
Perform one of the following actions:

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

Manage replication operations

Replication operations in server recovery policies are managed in the same manner as in client recovery policies. For more information, refer to the following topics:

- "View information about replication operations" on page 255
- "Start replication immediately" on page 255
- "Cancel a replication in progress" on page 255

Start failover or failback immediately

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

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

To fail over all links manually

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

1. Right-click a server recovery policy and click **Failover all Links**.
2. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active targets to the standby target. Failover is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.
   - Click **Change** to fail over to the offline targets without replicating contents. Select this option when you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

To manually fail back all links

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

1. Right-click a server recovery policy and click **Failback all Links**.
2. Perform one of the following actions:
   - Click **Replicate** to replicate contents from the active targets to the offline targets. Failback is not performed. This enables you to synchronize data before you make the offline targets the active targets in each link.
   - Click **Change** to fail back to the offline targets without replicating contents. Select this option when you are sure that the data contained in the targets you want to activate is already synchronized with the active targets in each link.

Remove links from a server recovery policy

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

- Select links that appear under the policy, right-click, and click **Remove from Disaster Recovery policy**.
- Select a policy to show a **Properties** tab in the right pane. Select the links under the Active Links list, right-click, and click **Remove from Disaster Recovery policy**.
- Select a policy to show a **Properties** tab in the right pane. Select a link or links under the **Nonactive Links** list, right-click, and click **Remove from Disaster Recovery policy**.
- Right-click a link in the Logical View tree that is a member of a server recovery policy and click **Remove from Disaster Recovery policy**.

Delete a server recovery policy

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

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

NOTE
You should not delete a policy that contains active links. To remove active links from a server recovery policy, refer to "Remove links from a server recovery policy" on page 267 for more information.

Chapter

Managing Snapshots on Filers

9

This section provides instructions about how to create snapshot access and snapshot scheduling policies to manage snapshots on NetApp filers.

In this chapter

- Snapshot Access Policies 269
- Snapshot Scheduling Policies. 275

Snapshot Access Policies

Snapshot access policies enable you to centrally manage snapshots on multiple NetApp filers. You can take snapshots of data and the logical structure of the namespace that references the data. You can take snapshots of volumes participating in your namespace and restore files and folders from snapshot backups.

To create a snapshot access policy, refer to the following topics and perform these tasks

1. "Create a snapshot access policy" on page 269
2. "Browse snapshot data" on page 271
3. "Restore data" on page 272

Create a snapshot access policy

Snapshot access policies enable you to back up and restore data on NetApp filers and the logical structure of the namespace that references the data. You select DFS roots, links, and folders in your namespace. A captured namespace is created that targets all scheduled snapshots on the volumes in the selected namespace.

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Snapshot Policies**, and click **Access Policy**.
   - Right-click the **Snapshot Policies** folder in the Admin View, point to **New**, and click **Access Policy**.
   - Right-click the **Access Policies** folder and click **New Access Policy**.
   - Drag a link, root, or folder in the Logical View that references data on NetApp filers to the **Access Policies** folder.
   - Drag a NetApp filer or volume in the Physical View to the **Access Policies** folder.
   - Right-click **Admin View**, click **New**, point to **Snapshot Policy**, and click **Access Policy**.

2. In the **Name** field, type a word or phrase that best describes this policy. The name appears under the **Access Policies** folder and on the **Properties** tab of items that are members of the policy.
3. Perform one of the following actions:
   - Click **Add to type** or select a DFS root, link, or folder that references data on filers.
   - Click **Add All Roots From the Logical View** to include all the DFS roots displayed in your namespace.
4. Optionally select **Limit selection to the following volumes** and click **Add to type the host and volume name** (for example, `secure:/vol/vol0`), or click **Browse** to locate the volume on your network or the Physical View tree. This enables you to select a subset of the volumes referenced from the namespace.
5. In the **Use the Monitor on** list, select the machine that contains the agent that you want to use to process the policy. Normally, the current VFM server is used to process the policy. If you installed monitoring agents on other machines, you can select one of the locations where monitoring agents are installed as the location from which this policy is processed.
6. Click **OK**. The policy searches for snapshots on the volumes referenced in the selected namespace, captures the namespace targeting each snapshot, and displays the captured namespace files under a **Captured Namespaces** folder under the policy.

NOTES:

- Filers participating in a snapshot access policy must have **Show Snapshot** enabled in their CIFS configuration.
- For more information about snapshots, refer to the *Data ONTAP® System Administration Data Protection Guide*.
- Monitoring agents are installed using the VFM media and uninstalled through normal Windows mechanisms for removing programs.

Take a snapshot immediately

You can override a snapshot schedule and manually create a snapshot any time you want.

To take a snapshot and capture an associated namespace

1. Perform one of the following actions:
   - Right-click a snapshot access policy and click **Snapshot referenced volumes now**.
   - Select a snapshot access policy to show tabs in the right pane. Right-click anywhere on the **Properties** tab and click **Snapshot referenced volumes now**.
2. Type the name of the snapshot.
3. Optionally run batch files before and after the snapshot process.
4. Click **OK**. The captured namespace appears under the **Captured Namespaces\Manual** folder.

To take a snapshot of one or selected volumes participating in the policy

1. Perform one of the following actions:
   - Right-click a volume under a filer folder and click **Take snapshot now**.

- Select the **Volumes** folder under a snapshot access policy to show a **Properties** tab in the right pane. Select the volumes of which you want to take snapshots and click **Take snapshot now**.

2. Type the name of the snapshot.
3. Optionally run batch files before and after the snapshot process.
4. Click **OK**. To view the snapshot, double-click the volume and click the **Snapshots** tab. Select a snapshot on the list and click **Refresh**.

NOTE
When you take a snapshot of a volume referenced in a snapshot access policy, the namespace is not capture. To capture the namespace of the snapshot, refer to "Capture the namespace of a snapshot" on page 275 for more information.

Update captured namespaces

VFM periodically checks for new filer-initiated and VFM-initiated snapshots on volumes participating in a snapshot access policy and updates the captured namespaces. You can update captured namespaces manually if you do not want to wait for the scheduled updates.

Perform one of the following actions:

- Right-click a snapshot access policy and click **Update Volumes and Captured Namespaces**.
- Select the **Snapshot Policies\Access Policies** folder to show a **Properties** tab in the right pane. Right-click a snapshot access policy and click **Update Volumes and Captured Namespaces**.

Specify which volumes participate in a policy

Specify which volumes referenced in your namespace participate in a snapshot access policy. You can select a subset of the volumes referenced in your namespace.

1. Perform one of the following actions:
   - Right-click a snapshot access policy and click **Properties**.
   - Select a snapshot access policy to show a **Properties** tab in the right pane. Click **Edit**.
2. Select **Limit selection** to the following volumes.
3. Click **Add** and type the path of a volume (for example, `safe:/vol/vol3)` on a NetApp filer containing data targeted by a link in your namespace, or click **Browse** to find a volume.
4. Repeat Step 3 to add more volumes.
5. Click **OK**.

Browse snapshot data

Browse the contents of a captured namespace and the snapshot data it references.

1. Open a folder under the **Captured Namespaces** folder of a snapshot access policy to show a list of captured namespace files.
2. Select one of the files to show the **Explorer** tab in the right pane.

3. 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.

NOTES:

- If you have the appropriate program, you can double-click a file to view its contents.
- You can copy snapshot data to another location so users can access the data directly. For more information, refer to "Access snapshot data directly" on page 274.
- If you have problems viewing files in the **Explorer** tab, refer to Appendix D, "Troubleshooting" for more information.

Restore data

Restore snapshot data to its original location. The snapshot overrides corresponding files in the original location. The content of a file is restored to the state that existed at the time the snapshot was taken. If a corresponding file does not exist in the original location, the file in the snapshot is copied to the original location.

1. Perform one of the following actions:
   - Right-click a captured namespace under the **Captured Namespace** folder and click **Restore**.
   - Select a snapshot access policy to show tabs on the right pane. Click the **Explorer** tab. Right-click an item on the left pane or right pane and click **Restore**.
2. Select **Restore the Snapshot Files**.
3. Click **OK**.

NOTES:

- You can restore the data on one or more files by selecting items in the **Explorer** tab of a snapshot access policy.
- You can browse a file before you restore it. For more information, refer to "Browse snapshot data" on page 271.
- If you have problems restoring files, refer to "Troubleshooting" on page 367 for more information.
- You can copy snapshot data to another location so users can access the data directly. For more information, refer to "Access snapshot data directly" on page 274.

Restore namespace and data

Restore a captured namespace to a root, an existing folder, or new folder in your namespace and restore snapshot data to its original location. The restored namespace and content of files referenced by the namespace are restored to the state that existed at the time the snapshot was taken. If a corresponding file does not exist in the original location, the file in the snapshot is copied to the original location.

1. Perform one of the following actions:
   - Right-click a captured namespace under the **Captured Namespace** folder and click **Restore**.
   - Select a snapshot access policy to show tabs on the right pane. Click the **Explorer** tab. Right-click an item on the left pane or right pane and click **Restore**.

2. Select **Restore the Captured Namespace and the Snapshot Files**.
3. Type the name of a root or folder in your namespace where you want to create the namespace to reference the current active file system, or click **Browse** to find the root or folder.
4. Click **OK**. The namespace is restored to the location you specified and the data is restored to the active file system.

Re-create part of your namespace

Recreate a portion of your namespace. You can restore the captured namespace to a root, an existing folder, or new folder in your namespace. The restored namespace references the current data.

1. Perform one of the following actions:
   - Right-click a captured namespace under the **Captured Namespace** folder and click **Restore**.
   - Select a snapshot access policy to show tabs on the right pane. Click the **Explorer** tab. Right-click an item on the left pane or right-click a folder or link in the right pane and click **Restore**.
2. Select **Restore the Captured Namespace**.
3. Type the name of a root or folder in your namespace where you want to create the namespace to reference the current data, or click **Browse** to find the root or folder.
4. Select **New namespace will target active file system**.
5. Click **OK**.

Access snapshot data through the namespace

Make snapshot data accessible to users through the namespace. You can copy the captured namespace to a root, an existing folder, or a new folder in your namespace. The links in the restored namespace reference the snapshot data.

1. Perform one of the following actions:
   - Right-click a captured namespace under the **Captured Namespace** folder and click **Restore**.
   - Select a snapshot access policy to show tabs on the right pane. Click the **Explorer** tab. Right-click an item on the left pane or right-click a folder or link in the right pane and click **Restore**.
2. Select **Restore the Captured Namespace**.
3. Type the name of a root or folder in your namespace where you want to create the namespace to reference the data in the snapshot, or click **Browse** to find the root or folder.
4. Make sure **New namespace will target snapshot data** is selected.
5. Click **OK**.

NOTES:

- For more information about how to browse the contents of a snapshot file, refer to "Browse snapshot data" on page 271.

- For more information about how to copy snapshot data to another location so users can access the data directly, refer to "Access snapshot data directly" on page 274.

Access snapshot data directly

You can copy snapshot data to another location so that users can access the data directly.

1. Select a snapshot access policy to show tabs in the right pane.
2. Click the **Explorer** tab.
3. Click a link in the right or left pane to show the files referenced by the link.
4. Drag a file to Windows Explorer.

NOTE
A snapshot access policy enables you to make snapshot data available to users through the namespace. For more information, refer to "Access snapshot data through the namespace" on page 273.

Change the snapshot schedule for all volumes

You should ensure that snapshot schedules for the volumes participating in your snapshot access policy are aligned so that snapshots of all the volumes occur as close to the same times as possible. You can quickly update the snapshot schedule for the volumes referenced in a snapshot access policy. A wizard steps you through the process of creating a new snapshot scheduling policy or adding the volumes to an existing policy.

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**. Adding volumes to a snapshot scheduling policy enables you to manage the schedule of the snapshots on all volumes participating in the policy.
4. Click **Next**.
5. Perform one of the following actions:
   - Click **New Policy**. Type a word or phrase that best describes the policy in the **Name** field. Then, select whether you want VFM or the filer to perform the snapshots.
   - Click the **Virtual File Manager Settings** tab and select **Use Virtual File Manager-initiated snapshot**. Type how often you want VFM to take snapshots.
   - Click the **Filer Settings** tab and select **Manage filer-initiated snapshot schedule**. Type how many snapshots you want to retain and how often you want the filer to take snapshots, and then click **OK**. Select the new policy and click **Next**.
   - Select an existing policy and click **Next**.
6. Click **Finish**. A new policy appears under the **Snapshot Policies\Scheduling Policies** folder.

Change the snapshot schedule for selected volumes

Change the snapshot schedule for some of the volumes participating in a snapshot access policy.

1. Select the policy to show a **Properties** tab in the right pane.
2. Under **Volumes**, select the volumes for which you want to change the snapshot schedule.
3. Right-click the selected volumes and click **Set snapshot schedule**.
4. Continue with Step 2 of "Change the snapshot schedule for all volumes" on page 274.

Capture the namespace of a snapshot

Capture the namespace of a manual snapshot of a volume. The volume must be referenced by the namespace and participate in a snapshot access policy.

1. Right-click the policy and click **Properties**.
2. Click **Capture Namespaces**.
3. Click **Add** and type the name of a manual snapshot.
4. Click **OK**. The captured namespace appears under the **Captured Namespaces\Manual** folder.

Update the properties of a snapshot access policy

Add or remove volumes or namespace items from the policy or change the monitoring agent for the policy.

1. Perform one of the following actions:
   - Right-click the policy and click **Properties**.
   - Select the policy to show a **Properties** tab in the right pane. Click **Edit**.
2. Update the appropriate information.
3. Click **OK**.

Delete a snapshot access policy

When a snapshot access policy is deleted, all the captured namespace files are also deleted. Snapshots of the data are not deleted.

1. Perform one of the following actions:
   - Right-click the policy and click **Delete**.
   - Select the **Snapshot Policies** folder to display a **Properties** tab in the right pane. Under the **Snapshot policies** folder, right-click the policy and click **Delete**.
2. Click **Yes**.

Snapshot Scheduling Policies

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 reserve setting on multiple filers without having to repeat the changes on each individual filer.

Create a snapshot scheduling policy checklist

To create your first policy, refer to the following topics and perform these tasks:

1. "Create a snapshot scheduling policy" on page 276
2. "Schedule snapshots" on page 276

Create a snapshot scheduling policy

Create a snapshot scheduling policy to schedule snapshots on multiple filers.

1. Perform one of the following actions:
   - On the **Tools** menu, click **New Policy**, point to **Snapshot Policy**, and click **Scheduling Policy**.
   - Right-click the **Snapshot Policies** folder in the Admin View, point to **New**, and click **Scheduling Policy**.
   - Right-click **Admin View**, click **New**, point to **Snapshot Policy**, and click **Scheduling Policy**.
2. In the **Name** field, type a word or phrase that best describes this policy. The name appears under the **Snapshot Policies\Scheduling Policies** folder and on the **Properties** tab of items that are members of the policy.
3. Click **Add** and type the host and volume name (for example, `secure:/vol/vol0)`, or click **Browse** to locate the volume on your network or the Physical View tree.
4. Repeat Step 3 to add more volumes to the policy.
5. You can select another location where the policy is monitored. The default is to monitor from the VFM server. From the **Use the Monitor on** list, optionally select the machine hosting another monitoring agent.
6. Click **OK**.

NOTES:

- Monitoring agents are installed using the VFM media and uninstalled through normal Windows mechanisms for removing programs.
- If you uninstall a monitoring agent, ensure that policies that formerly used the uninstalled agent are reassigned to a currently available agent. Right-click a snapshot access policy and click **Properties**. On the **General** tab, select a new monitoring agent from the **Use the Monitor on** list.

Schedule snapshots

Set the schedule for filer-initiated as well as VFM-initiated snapshots. A snapshot scheduling policy can have both types of snapshots.

Perform one of the following actions:

- Right-click a snapshot scheduling policy and click **Properties**.
- Select a snapshot scheduling policy to show a **Properties** tab in the right pane. Click **Edit**.

To set the schedule for VFM-initiated snapshots

VFM-initiated snapshots enable you to run batch files before and after a snapshot.

1. Click **Virtual File Manager Settings**.
2. Select **Use Virtual File Manager-initiated snapshot**.
3. Type the number of snapshots you want to retain. By default, the latest four snapshots are retained.
4. Optionally, select **Run batch before snapshot** and type the UNC path of a script file that you want to run.
5. Optionally, select **Run batch after snapshot** and type the UNC path of a script file that you want to run.
6. Type the time and date when and how frequently to take snapshots.

For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

To set the schedule for filer-initiated snapshots

Filer-initiated snapshots enable you to configure snapshots and snapshot reserves on multiple volumes and maintain the filer schedule settings. If a filer-initiated snapshot schedule is changed outside the policy, VFM restores the schedule back to the settings in the policy the next time the policy runs.

1. Click the **Filer Settings** tab.
2. Select **Manage filer-initiated snapshot schedule**.
3. Type the number of weekly, nightly, and hourly snapshots you want to retain. By default, the two most recent weekly snapshots and the four most recent nightly and hourly snapshots are retained.
4. Select the hours at which you want to start taking snapshots.
5. Optionally, select **Apply the same Snapshot Reserve** to all volumes and type the amount of space to reserve for snapshots on all volumes in the policy.

Update the properties of a snapshot scheduling policy

Add, edit, or remove volumes from the policy and change the snapshot schedule.

1. Perform one of the following actions:
   - Right-click the policy and click **Properties**.
   - Select the policy to show a **Properties** tab in the right pane. Click **Edit**.
2. Update the appropriate information.
3. Click **OK**. A status report appears.
4. Click **Done**.

Suspend the snapshot schedule of a policy

You can stop a snapshot schedule from running on the volumes that are members of a policy. The schedule is not deleted.

1. Right-click the snapshot scheduling policy and click **Properties**.
2. Perform one of the following actions:

- If the schedule is VFM-initiated, click **Virtual File Manager Settings** and clear **Use Virtual File Manager-initiated snapshot**.
- If the schedule is filer-initiated, click **Filer Settings** and clear **Enable scheduled snapshots**.

3. Click **OK**. A status report appears.
4. Click **Done**.

NOTES:

- Existing snapshots on the selected volumes are not deleted when you suspend the schedule.
- Suspending a snapshot schedule does not prevent you from manually taking a snapshot. Right-click the policy and click **Take snapshot now**.

Take a snapshot immediately

Take a manual snapshot of one or more volumes in a snapshot scheduling policy.

1. Perform one of the following actions:
   - Double-click a policy under the **Scheduling Policies** folder to show the volumes that are members of the policy. Right-click a volume and click **Take snapshot now**.
   - Select a policy to show a **Properties** tab in the right pane. Select one or multiple volumes, right-click, and click **Take snapshot now**.
2. Type a name for the snapshot.
3. Optionally, select **Run batch before snapshot** and type the UNC path of a script file that you want to run.
4. Optionally, select **Run batch after snapshot** and type the UNC path of a script file that you want to run.
5. Click **OK**.
6. Click **Done**.

Delete a snapshot scheduling policy

When you delete a snapshot scheduling policy, the schedule for VFM-initiated snapshots is also deleted. The management of filer-initiated snapshots returns to the respective filers. Changes made to a filer-initiated snapshot after the deletion of the policy remain in effect. The parameters set by the policy are no longer enforced.

1. Perform one of the following actions:
   - Right-click the policy and click **Delete**.
   - Select the **Scheduling Policies** folder to show a **Properties** tab in the right pane. Right-click the policy and click **Delete**.
2. Click **Yes**.
3. Click **Done**.

Chapter

Working with Reports

10

This section provides information about how you can create reports, publish them to a location you specify, process report data, and archive reports.

In this chapter

- Introduction 279
- VFM reports 280
- Creating and Publishing Reports 282
- Grouping and Scheduling Reports 287
- Managing Reports 291

Introduction

The Reports tree shows reports of recorded events or instructions processed that you generate or save. 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.

- When you create a report, you assign it to a report source group that specifies the locations where you want to gather information for the report. You can assign multiple reports to a report source group.
- Report schedule groups specify when and how often data is gathered. VFM provides default report schedule groups. You can use the default settings or change the settings. A report schedule group can contain multiple report source groups.
- When you create a report, it appears under the report source group to which it is assigned and under the **Generated Reports** folder. You can organize reports in folders under the **Generated Reports** folder.
- You can view the contents of the report and modify them.

VFM reports

VFM provides the following types of reports:

- Reports generated from data periodically gathered by report source groups and maintained by VFM. These reports appear under the report source groups to which they are assigned and under the **Generated Reports** folder.
- Reports of recorded information, such as events. These reports are not scheduled and appear only under the **Generated Reports** folder.

The following table lists reports you can generate. These reports are assigned to report source groups.

TABLE 49 Administrative reports

Administrative reports	Descriptions
Business Continuity Policy Status	Enables you to review links contained in the selected recovery policies and to check storage and replication status for each target of a policy.
Policy References (Admin View Selection)	Enables you to view the resources referenced in selected policies.
Replication Policy Status	Enables you to review source and destination links contained in the selected replication policy and to check storage and replication status.
Server Configuration	Enables you to review your licensing information and the storage on each machine. For Windows computers, VFM uses WMI to collect storage information. If WMI is not available, the **Storage Used** column represents the total used volume size for all volumes that have at least one share. For UNIX computers, the **Storage Used** column represents the total used space for NFS exports.

TABLE 50 Logical namespace reports

Logical Namespace reports	Descriptions
Namespace Properties	• Provides an overview of physical items, such as servers and link targets (shared folders) that form part of the selected DFS root. • Enables you to quickly review machines participating in the selected root and to check available disk space. • Provides the option of gathering information about the size of the logical structure of the namespace (also referred to as BLOB size) NOTE: Gathering BLOB size information can take a long time, especially for stand-alone roots that contain a large number of links.
Physical References (Logical View Selection)	Enables you to view the DFS links or roots that reference the targets in the namespace. The Report wizard enables you to filter the report by items you select in the Logical View.
Replica Properties	Provides an easy means for scanning a namespace and viewing links that are highly available because they have multiple targets and viewing links that are not highly available because they have only one target.

TABLE 51 Physical resources reports

Physical resources reports	Description
File Differences	Provides an easy way to compare the differences that might exist in the contents of network shares.
Server Properties	Enables you to view the machines participating in your namespace.
Share Connections	Provides an easy way to view how many users are currently connected to the items selected. This information is useful when you are considering taking a share offline or rebooting a machine for maintenance.
Share Permissions	Enables you to view the permissions of the shares selected.
Share and Export Properties	Includes information about shares and exports, such as local path, capacity, and space usage.

TABLE 52 File system reports

File system reports	Description
Archived Files	Lists files that do not have their archive attribute set.
File Ages	Lists files arranged by age and file size for a specified location.
File Summary By Type	Summarizes the numbers of files and space consumed based on the types of the files.
Files Not Backed Up	Lists files modified within a specified time and whose archive attribute is set.
Large Files	Lists files larger than a specified size in the selected locations. You can exclude files from the report that have been accessed recently.
New Files	Lists files that were created within a specified time.
Orphaned Files	Lists files whose owner is unknown (displayed as a raw SID) in your environment.
Over-utilized File Systems	Lists file systems with small amounts of free space remaining.
Recently Used Files	Lists the files accessed within a specified time.
Stale Files	Lists files that have not been accessed within a specified time.

TABLE 52 File system reports

File system reports	Description
Under-utilized File Systems	Lists file systems containing a large amount of free space.
Usage By Owner	Lists the amount of space consumed by files owned by each user.

TABLE 53 NetApp filer reports

NetApp filer reports	Description
Filer CIFS Configuration	Enables you to review information about the CIFS set up on a filer.
Filer CIFS Statistics	Enables you to review CIFS protocol usage information by filer. The report contains information about the physical filer and all vFilers hosted on the filer.
Filer File System Status	Enables you to review statistics about each volume on a filer.
Filer Properties	Enables you to view the filers participating in your namespace. Contains details about the filers, including model, product ID, status, up time, and operating system, and whether the filer is licensed for CIFS and NFS protocols.
Filer Protocol Usage	Enables you to review CIFS and NFS protocol operations for a filer.
Filer Qtrees by Volume	Enables you to review information about available qtrees by volume for a filer.
Filer RAID	Enables you to review RAID group configuration information by volume for a filer.
Filer SnapMirror Status	Enables you to review statistics about SnapMirror status on selected filers.
Filer Snapshot	Enables you to review information about available snapshots by volume for a filer.
Filer vFiler Status	Enables you to review statistics about each vFiler on selected filers.

Creating and Publishing Reports

This section explains how you can generate and publish a report. You must enable report publishing before you can publish a report. For more information, see "Enable report publishing" on page 282.

To create and publish a report, complete the following tasks.

1. "Generate a report" on page 283
2. "Start data collection immediately" on page 284
3. "Publish reports" on page 284

Enable report publishing

You can publish reports to a location you specify. Every time you publish a report, VFM generates an index file that is automatically copied to the same location as the published report. You must configure global report publishing settings before you can publish a report you generate.

1. On the **Tools** menu, click **Options**.
2. Under Systems **Options**, click **Reporting**.
3. Select the **Enable Report Publishing** check box.

4. Type the UNC path of the share where you want to publish reports and copy the updated index file, or click **Browse** to find a share.
5. Select **Publish reports by default** to set report publishing as the default report output for all new reports you create.
6. (Optional) Type the name of an XSL style sheet you want to use for the index file. A default style sheet is available in the VFM installation directory in the `Examples\Stylesheets\Report` publishing folder.
7. (Optional) Type the name of a batch file, or click **Browse** to find a batch file to run before the index file is updated.
8. (Optional) Type the name of a batch file, or click **Browse** to find a batch file to run after the index file is updated.
9. Click **OK**.

ATTENTION
Make sure that the location you specify to publish reports is empty. VFM does not check for duplicate file names.

For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Configuring default report settings

You can define report settings for all new reports.

To configure application-wide default report settings

1. On the **Tools** menu, click **Options**.
2. Click **Reporting**.
3. Update the settings and click **OK**. Only new reports contain the updated settings. Existing reports are not affected.

Generate a report

Gather information about the namespace and network resources, such as servers, shares, and exports.

- You specify the information that will appear on the report.
- You can process the data gathered by the report to run a batch file you specify and, if you enabled report publishing, you can publish the report.
- You assign the report to a source group and specify whether you want to archive the report.
- VFM provides numerous report types.
- Reports you create appear in the Report tree under the source group to which they are assigned and under the **Generated Reports** folder.

To generate a report

1. Perform one of the following actions:
   - On the **Reports** menu, point to a category, and click a report type.
   - Right-click the Reports tree, point to **New Report**, point to a category, and click a report.
   - Right-click a report source group, point to **New Report**, point to a category, and click a report.
2. The Report wizard steps you through the process of creating the report. For help on this dialog box, click **Help**.

NOTES:

- When creating file system reports, do not select overlapping shares when you specify the source for a report. For example, do not select C$ and a shared subdirectory of the C drive together in the scope of a report. This example results in double-reporting of the items beneath the shared subdirectory.
- You can also create a report of a specific item on the Logical View tree.
  - Right-click the item in the Logical View tree, point to **Reports**, and click the report type.
  - To create a report of a specific target, click a link or root to display the **Properties** tab. Under **Target**, right-click a shared folder, point to Reports, then click the report type.
- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.
- For more information about how to enable report publishing, refer to "Enable report publishing" on page 282.
- VFM provides default report schedule and source groups. For more information about how to update a report group, refer to "Modify a report group" on page 289.

Start data collection immediately

When you manually start data collection, the data gathered and the reports are updated based on the new data. When you start data collection for a report schedule group, all the report source groups assigned to the schedule group gather new data and update their respective reports. You can start data collection for a report source group.

Perform one of the following actions:

- Right-click the report schedule group or report source group that contains the reports you want to refresh and click **Start data collection now**.
- Select a report schedule group or report source group to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Start data collection now**.

Publish reports

Before you can publish reports, you must enable report publishing. Reports are published in HTML.

1. Perform one of the following actions:
   - Right-click a report and click **Properties**.
   - Select a report to show its contents in the right pane. Click **Edit**.
2. Click **Next** until you see the Report Output dialog box.

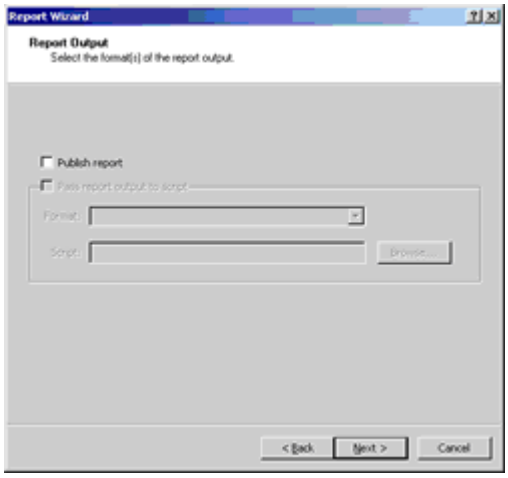

3. Select the **Publish report** check box.
4. Click **Next** until you see the Report Name dialog box.
5. Click **Finish**. The next time new data is gathered, the contents of the report are updated and the report is published.

In the Reports tree, published reports are distinguished from reports that are not published. Published reports in the Reports tree are synchronized with their published location. When you delete a published report from the Reports tree, the report is also deleted from the location where it is published.

To publish all reports designated for publication, refer to "Manually publish all reports" on page 287 for more information. You can configure report publishing as the default option for all the reports you create.

To configure report publishing as the default option

1. On the **Tools** menu, click **Options**.
2. Click **Reporting**.
3. Select **Publish reports by default**.
4. Click **OK**.

Run customized commands during index file generation

When you publish reports, you can run batch files before and after the index file is updated.

1. On the **Tools** menu, click **Options**.
2. Click **Report Publishing**.
3. Select **Run batch before updating index** file and type the name of a batch file, or click **Browse** to find a batch file in your network you want to run before VFM updates the index file.
4. Select **Run batch after updating index file** and type the name of a batch file, or click **Browse** to find a batch file in your network to run after VFM updates the index file.
5. Click **OK**.

For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Output report data to a file

Specify a batch file to process report data and select the format you want to output the data. You can specify a batch file for a single report or for all reports you generate. A copy of the report is stored in VFM and is available in the **Reports** tree.

To specify a batch file for a single existing report

1. Right-click a report and click **Properties**.
2. Click **Next** until you see the Report Output dialog box.

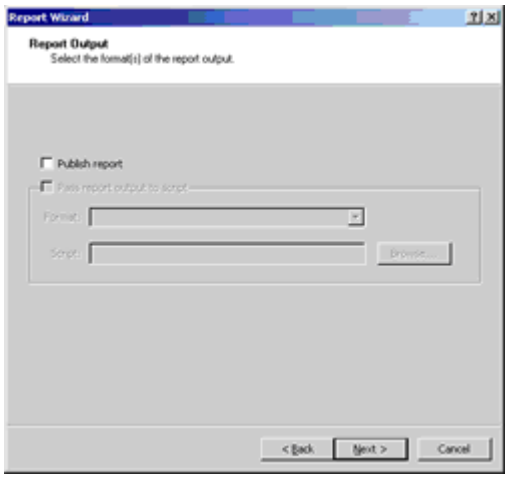

3. Select **Pass report output to script**.
4. In the **Format** list, select the format in which you want to output the data.
5. Type the name of a batch file that will process the data, or click **Browse** to find the file.
6. Click **Next** until you see the Report Name dialog box.
7. Click **Finish**. The batch file runs the next time the schedule for the report runs.

To specify a batch file for a new report

Point to a report category on the Reports menu and click a report type. The Report wizard steps you through the process of creating the report and specifying a batch file.

To specify a batch file for all reports

You can specify a default batch file to process the data for all reports.

1. On the **Tools** menu, click **Options**.
2. Click **Reporting**.
3. Select **Pass report output to script**.
4. In the Format list, select the format in which you want to output the data.
5. Type the name of a batch file that will process the data, or click **Browse** to find the file in your network.

6. Click **OK**.

For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Manually publish all reports

You can publish all the reports appearing in the Reports tree that are enabled for publishing. This feature is useful when you change the location where reports are published and want to populate the new location with reports.

Perform one of the following actions:

- Right-click **Reports** and click **Publish Reports Now**.
- Select **Reports** to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Publish Reports Now**.

NOTE
The reports are published with the data gathered the last time the report schedule groups ran.

Grouping and Scheduling Reports

You can group and schedule reports to simplify data gathering and to schedule these activities during off hours.

Create a report schedule group

Create a new report schedule group to run the source groups that are assigned to it. A report schedule group can be configured to start gathering information 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 schedule groups enable you to run batch files before and/or after a schedule runs.

1. Right-click **Reports** and click **New Report Schedule Group**.
2. Type a name for the schedule group.
3. Select the **Use schedule** check box.
4. Click **New** and configure how often you want the assigned report source groups to gather data and run the reports.
5. Optionally, click **Scripts** to type the name of batch files you want to run before and/or after the report schedule group runs.
6. Click **OK**. The report schedule group appears under the Reports tree.

NOTES:

- VFM provides several report schedule groups. You can use their default settings or you can modify them.
- When you create a report schedule group, a default report source group is created. You must edit the default source group and specify locations at which you want to gather information.

- For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.
- For more information about how to run a schedule or source group batch file, refer to "Run customized commands during index file generation" on page 285.

Create a report source group

Create a new report source group to periodically gather and store information about your storage environment. You produce reports based on the data that is gathered. A report source group specifies the location from which you want to gather the data. It also enables you to specify batch files to run before and after the information is gathered. Report source groups are assigned to report schedule groups. You can assign multiple report source groups to a report schedule group.

1. Right-click a report schedule and click **New Report Source Group**.
2. Type the name of the source group.
3. Click **Add** and type the path of the item where you want to gather data. Refer to the Note below for format examples.
4. Repeat Step 3 to add more locations.
5. Optionally, click **Scripts** to type the name of batch files you want to run before and/or after the report source group runs.
6. Click **OK**.

Format Examples

The following examples illustrate the formats you should use when adding items to the list.

- An Admin View policy

  Format: `$/Admin View/<item>`

  Example: `$/Admin View/Data Movement Policies/Migration Policies/My migration policy`
- All Migration policies

  Format: `$/Admin View/<item>`

  Example: `$/Admin View/Data Movement Policies/Migration Policies`
- (CIFS only) A link or logical folder in the Logical View

  Format: `\\<domain>\<DFS root>\<link or logical folder>`

  Example: `\\ilus.com\priamroot\marketing`
- (CIFS only) A directory under a share on a machine

  Format: `\\<machine name>\<share>\<directory>`

  Example: `\\mars.my_company.com\sales\FirstQtr2003`
- A qtree on a filer

  Format: `<filer name>:/vol/<volume name>/qtree`

  Example: `moria:/vol/vol0/EngineeringRpts`
- A volume on a filer

Format: `<filer name>:/vol/<volume name>`

Example: `moria:/vol/vol0`

- (NFS only) An export

  Format: `<hostname>:/<root>/<subdirectory>`

  Example: `moria:/vol/marketing`

Modify a report group

Changes you make to a schedule or resource group affect all the reports that are members of the groups.

1. Perform one of the following actions:
   - Right-click a report schedule or source group and click **Properties**.
   - Select a report schedule or source group to show a **Properties** tab in the right pane. Click **Edit** or right-click anywhere on the page and click **Properties**.
2. Make the necessary updates and click **OK**. For more information, refer to "Create a report schedule group" on page 287 and "Create a report source group" on page 288.

Move a source group

When you move a source group to another schedule group, all the reports assigned to the source group move to the new schedule group.

1. Perform one of the following actions:
   - Right-click a source group and click **Move to schedule group**.
   - Select the schedule group that contains the source group you want to move to show a **Properties** tab in the right pane. Right-click the source group you want to move and click **Move to schedule group**.
2. Select the schedule group you want to move the source group, or click **Create schedule group** to move the source group to a new schedule group.
3. Click **Finish**.

View reports assigned to a source group

View the list of reports that are generated from data gathered by a report source group.

Perform one of the following actions:

- Expand a source group in the Report tree to show all the reports that are assigned to it.
- Select a report source group to show a **Properties** tab in the right pane. The reports assigned to it appear under **Reports**.

NOTE
All the reports you create also appear alphabetically under the **Generated Reports** folder.

Update a report schedule

To update when and how often a report is generated, you must update the schedule group to which the report belongs. Updating the schedule of a schedule group affects all reports that are members of the group.

1. Perform one of the following actions:
   - Right-click the report schedule group and click **Properties**.
   - Select a report schedule group to show a **Properties** tab in the right pane. Click **Edit**.
2. Click **Schedule**.
3. Click **New** and configure your schedule.
4. Click **OK**.

NOTE
The data-gathering process might take a significant amount of time, depending on the report type and volume of data to be gathered. It is not recommended that you run a report more frequently than the interval required to gather data for a report schedule group.

Delete a report source group

When you delete a report source group, the reports assigned to the group and their archives are deleted.

1. Perform one of the following actions:
   - Under a report schedule group, right-click a source group, and then click **Delete**.
   - Select a report schedule group to show a **Properties** tab in the right pane. Right-click a source group and click **Delete**.
2. Click **Yes**.

Delete a report schedule group

When you delete a report schedule group, all the source groups, reports, and archives assigned to it are deleted.

1. Perform one of the following actions:
   - Right-click a report schedule group and click **Delete**.
   - Click the Report tree to show a **Properties** tab in the right pane. Right-click a report schedule group and click **Delete**.
   - Select a schedule group to show a **Properties** tab in the right pane. Right-click anywhere on the page and click **Delete**.
2. Click **Yes**.

Managing Reports

There are many additional tasks you can perform to meet your reporting needs. These tasks help you customize reporting.

Organize reports

Add folders to organize your reports in the **Generated Reports** folder, which lists all your reports in alphabetical order.

1. Right-click the **Generate Reports** folder and click **New Folder**.
2. Type a meaningful name for the folder.

NOTE
When you create a report, it appears under the report source group to which it is assigned and under the **Generated Reports** folder.

Save event history

Save a current list of events for an item in the namespace or for supporting servers, shares, and exports. The report appears under the **Generated Reports** folder in the Reports tree.

1. Click an item under Admin View, Logical View, or Physical View.
2. Click the **Events** tab that appears in the right pane.
3. Click **Save** to save a copy of the event list.
4. Type a name for the report.
5. Optionally, select an existing folder or click **New Folder** to create a new folder in which to save the report.
6. Click **Finish**.

NOTE
To view all the events for a view (Admin, Logical, or Physical), click the view on the tree.

Modify a report

Update the properties of a report.

1. Expand the Reports tree.
2. Perform one of the following actions:
   - In the **Generated Reports** folder, right-click the report you want to change and click **Properties**.
   - Select the report source group to which the report is assigned to show a list of reports on the **Properties** tab that appears in the right pane. Right-click the report you want to change and click **Properties**.
   - Open the report and click **Properties** on the report toolbar.

   The Report wizard steps you through the process of updating the report.

Configure report archiving

You can specify how many copies of a report you want to keep. You can specify this option for all reports or a single report.

To configure application-wide report archiving

Configure a default value for all reports. Existing reports are not affected. Only reports you create after configuring this value are affected.

1. On the **Tools** menu, click **Options**.
2. Click **Reporting**.
3. Select **Enable report archiving** and type the number of copies you want to keep for all reports.
4. Click **OK**.

To configure report archiving for a report

You can change the default report archiving value for new or existing reports. When you create a new report, the Report wizard provides the option of changing the value for report archiving. To change report archiving for an existing report, complete the following steps:

1. Right-click a report and click **Properties**.
2. Click **Next** until the Report Name dialog box of the Report wizard appears.
3. Select **Enable report archiving** and type the number of copies you want to keep for the report. When this number is exceeded, the oldest report is deleted.
4. Click **Finish**.

Run customized commands

Run batch files before and after report schedule or report source groups run. When a schedule group and its source groups run batch files, VFM executes the batch files in a specific order. The order in which the scripts are executed is illustrated in the Notes section of this topic.

1. Right-click a report schedule or report source group and click **Properties**.
2. Click **Scripts**.
3. Select Run batch before gathering data and type the name of the batch file you want to run, or click **Browse** to find the file in you network.
4. Select Run batch after gathering data and type the name of the batch file you want to run, or click **Browse** to find the file in your network.
5. Click **OK**.

For more information about configuring batch files to run in VFM and about the example scripts provided by VFM, refer to "Run Batch Files" on page 321.

Batch File Order Example

The following example illustrates the order in which batch files run when they are specified in a schedule group and the source groups assigned to it. In the example, *SourceA*, and *SourceB* report source groups are assigned to the report schedule group named *Schedule1*. All three groups are configured to run batch files.

1. *Schedule1* starts and runs the batch file specified to run before gathering information.
2. *SourceA* starts. The source group performs the following actions:
   a. Runs the batch file specified to run before information is gathered.
   b. Gathers information.
   c. Runs the batch file specified to run after the information is gathered.
3. *SourceB* starts. The source group performs the following actions:
   a. Runs the batch file specified to run before information is gathered.
   b. Gathers information.
   c. Runs the batch file specified to run after the data is gathered.
4. *Schedule1* runs the batch file specified to run after all the source groups gather information.

View the data collection log

You can view information about the data gathered for reports. The Data Collection Log is refreshed every time you refresh all the reports assigned to a report source group and every time the report schedule group runs.

1. Select a report source group to show tabs on the right pane.
2. Click the **Data Collection Log** tab.
3. Select an item in the log and click **Properties**, or double-click an item to view detailed information.

To sort the information appearing in the log by a column value, in descending or ascending order, click the heading of that column.

By default, VFM shows all the data in the log. You can filter the information that appears.

- To view information about only the items that were not gathered because of errors, click **Error**.
- To view information about only the items that encountered some failures, click **Warning**.
- To view information about only the items that were completed successfully, click **Completed**.

Update the format of a report

Some reports enable you to specify how the data is arranged on the report. The type of information that is gathered for a report determines whether this feature is available.

1. Right-click a report and click **Properties**.
2. Click **Next** until the Report Columns dialog box of the Report wizard appears.

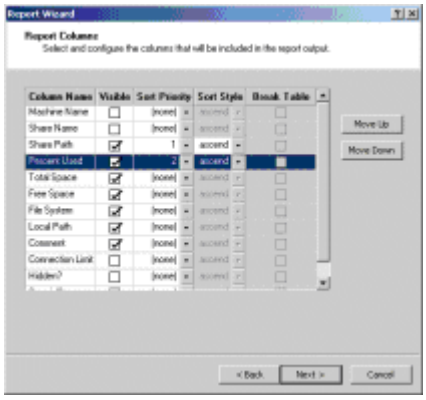

3. Select the data that should appear on the report. Select or clear the check boxes under **Visible**.
4. Arrange the column headings using the move keys.
5. Define how you want to sort the data. Begin by selecting the highest sort priority and optionally continue with the next level of priority.
6. Select how you want the data sorted (for example, ascending [A to Z] or descending [Z to A]). You must configure the sort priority of a column heading before you can configure the sort style.
7. Optionally arrange the data into multiple tables. You can break the table by up to three sort priorities. When the value in a selected column changes, the wizard creates a new table.
8. Click **Next**.
9. Click **Finish**.

Report formatting examples

The following examples illustrate how you can format a report using the Report wizard.

Example 1

The vertical list of selected column headings is arranged from left to right on the report table. In the following example, the **Break Table** option is not selected and the report is created with only one table. The table is sorted by **Share Name** and sub-sorted by **Total Space**.

TABLE 54 Report wizard

Column Name	Visible	Sort priority	Sort style	Break table
Machine Name				
Share Name	Yes	1	ascending	
Share Path	Yes			
Percent Used				
Total Space	Yes	2	ascending	
Free Space	Yes			

TABLE 55 Report output

Share Name	Share Path	Total Space	Free Space
Marketing	`\\machine5\Marketing`	10.0 MB	2.0 MB
Marketing	`\\machine2\Marketing`	20.0 MB	6.0 GB
Reports	`\\machine2\Reports`	10.0 MB	6.0 GB
Reports	`\\machine1\Reports`	20.0 MB	6.0 GB
Vouchers	`\\machine2\Vouchers`	5.0 MB	1.0 GB
Vouchers	`\\machine1Vouchers`	25.0 MB	2.0 GB

Example 2

In the following example, the **Break Table** option is selected for **Share Name**. A new table is created every time the **Share Name** changes. Each table is sorted by **Total Space**.

TABLE 56 Report wizard

Column Name	Visible	Sort priority	Sort style	Break table
Machine Name				
Share Name	Yes	1	ascending	Yes
Share Path	Yes			
Percent Used				
Total Space	Yes	2	ascending	
Free Space	Yes			

TABLE 57 Report output

	Share Path	Total Space	Free Space
Marketing	`\\machine5\Marketing`	10.0 MB	2.0 MB
	`\\machine2\Marketing`	20.0 MB	6.0 GB
Reports	`\\machine2\Reports`	10.0 MB	6.0 GB
	`\\machine1\Reports`	20.0 MB	6.0 GB
Vouchers	`\\machine2\Vouchers`	5.0 MB	1.0 GB
	`\\machine1\Vouchers`	25.0 MB	2.0 GB

Move a report

You can move a report or multiple reports to another report source group under the same or a different schedule group.

To move one report to another source group

Perform one of the following actions:

- Drag the report from anywhere in the Report tree to another source group.

- Right-click the report and click **Move to report source group**. Select the source group you want to move the report or click **Create source group** to move the report to a new source group. Click **Finish**.
- Select the source group to which the report you want to move is assigned. A list of reports appears on the **Properties** tab in the right pane. Drag a report to another source group in the Report tree.

To move multiple reports to another source group

1. Perform one of the following actions:
   - Click the source group to which the reports you want to move are assigned. A list of reports appears on the **Properties** tab in the right pane. Select the reports you want to move.
   - Under the **Generated Reports** folder or a subfolder, select the reports you want to move.
2. Right-click and click **Move to report source group**.
3. Select the source group to which you want to move the reports or click **Create source group to move the reports to a new source group**.
4. Click **Finish**.

NOTE
To move all the reports in a source group to another schedule group, refer to "Move a source group" on page 289 for more information.

View the contents of a report

There are several ways to view the contents of a report. Perform one of the following actions:

- Expand the **Generated Reports** folder. Click the report you want to view.
- Expand a report resource group and click the report you want to view.
- Click a report resource group to show a **Properties** tab in the right pane. Double-click a report.

The content of the report appears in the right pane.

NOTE
The time and date the report was last updated are displayed as a tool tip.

Rename a report

Change the name of a report.

1. Perform one of the following actions:
   - Slowly click the name of the report.
   - Right-click a report and click **Rename**.
2. Type a new name.
3. Press **Enter**.

Print a report

Print a report that is opened on the right pane.

Perform one of the following actions:

- Open a report and click **Print** on the report toolbar.
- Right-click anywhere on the opened report and click **Print**.

NOTE
To designate and configure the printer that VFM uses for printing reports, click **Print Setup** on the **File** menu.

Delete a report

Published reports in the Reports tree are synchronized with their published location. When you delete a published report, the report is also deleted from the location where it is published. To delete a report, perform one of the following actions:

- Expand the Reports tree, double-click the **Generated Reports** folder, select the report, and click **Delete** on the main toolbar.
- Expand the Reports tree, double-click the **Generated Reports** folder, right-click the report, and click **Delete**.
- Expand a source group, right-click a report, and click **Delete**.
- If the report is opened in the right pane, right-click the report and click **Delete**.
- If the report is opened in the right pane, click **Delete** on the report toolbar.

NOTES:

- When the VFM server is uninstalled, all the reports appearing in the Reports tree are deleted.
- When you delete a report, all copies of the report are removed from the **Reports** tree. For example, if you delete a report in the **Generated Reports** folder, the report is also removed from all source groups to which it is assigned.
- When you delete a report, all the archives for the report are deleted.

Chapter

Managing Agents and the VFM Server

11

Use Agent Management to manage the VFM server, monitoring agents, and Windows and UNIX replication agents. Agent Management enables you to deploy replication agents and perform management tasks, such as change credentials and listening port for replication and monitoring agents and the VFM server.

In this chapter

- Replication agent overview 299
- Proxy computers and replication agents 302
- Network ports in use 302
- Viewing status information 304
- Manage replication agents 304
- Installing and managing UNIX replication agents 310
- Manage monitoring agents 314
- Manage the application server 317

Replication agent overview

This section defines replication agents, discusses the requirements for using replication agents, and explains how the agents are distributed by VFM.

What a replication agent is

Replication agents perform file transfer for policies. The agents are deployed automatically by VFM as part of policies and file differences reports. The Agent Management tool enables you to deploy agents manually. VFM deploys the following two types of replication agents:

- Windows replication agents are implemented as 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.

Replication agent requirements

The following list provides information about replication agents that you should know before you manually deploy or manage agents.

- You can set the data transfer rate limits for machines that participate in data transfers by using the properties dialog box for the machine in the VFM Physical View.

- The account used by VFM to deploy a Windows replication agent needs to have administrator privileges on the machine that receives the replication agent. Windows replication agents are deployed through the machine's ADMIN$ share and are installed as a service. Both of these operations require administrator privileges on the system where they take place. The account used by VFM to deploy a UNIX replication agent needs to have root privileges on the machine that receives the replication agent.
- If you change the password on the account where the VFM service is installed, the Service account password and the password of all deployed replication agents and monitoring agents must be updated. *If you do not update passwords, already deployed agents cannot perform replication and monitoring*. Refer to "Changing the credentials of a replication agent" on page 307 for instructions about how to update credentials.
- When a replication operation uses NTFS Change Journal, the replication agent is deployed to the machine hosting the source target. NTFS Change Journal is only applicable to Windows files systems.
- (CIFS only) When data is replicated using differential replication, VFM deploys a replication agent to both the source and destination. If VFM cannot deploy a replication agent to either the source or destination, you must designate a proxy for that location.
- (CIFS only) If the machine on which you want to deploy a replication agent is running Windows XP SP 2, you must disable the firewall on that machine and then restart the machine.

Distributed replication

This section explains how VFM automatically distributes Windows and UNIX replication agents to transfer data. The section also provides instructions about how you can view the status of an agent.

The following list illustrates the algorithm used by VFM to deploy replication agents.

1. When you move data, the replication agent is deployed to the machine hosting the destination target.

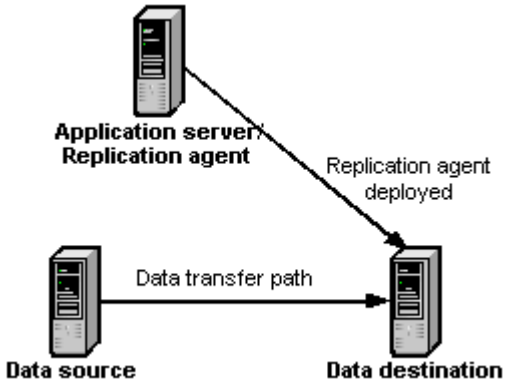

2. If VFM cannot deploy the agent to this machine, it attempts to deploy the agent to the machine hosting the source target.

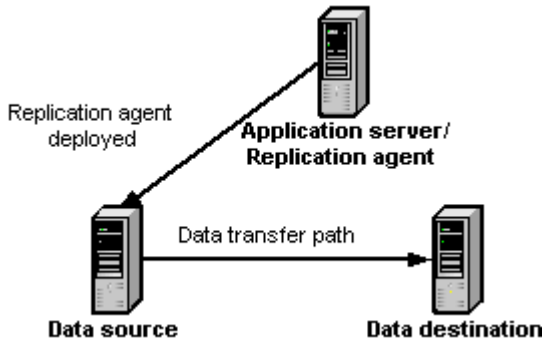

3. If VFM cannot deploy the agent to either the destination or source target, it attempts to use the agent on the proxy machine of the destination target, if one is configured.

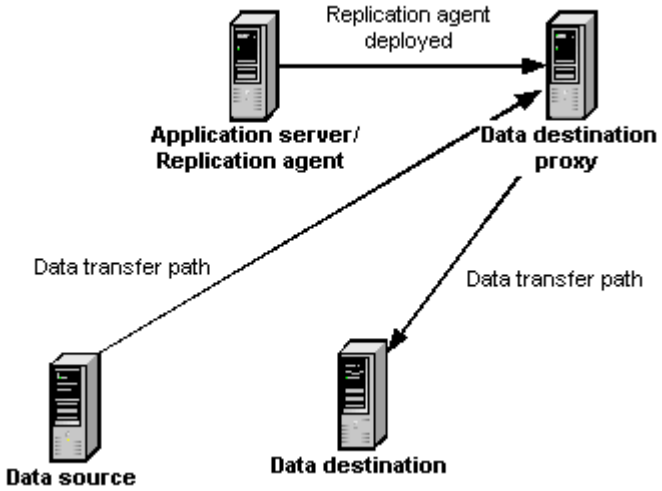

4. If VFM cannot deploy the agent to the proxy machine of the destination target, it attempts to use the agent on the proxy machine of the source target, if one is configured.

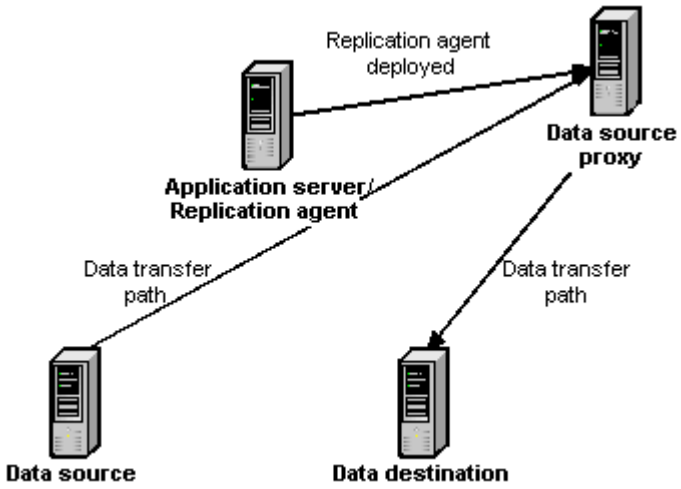

5. If VFM cannot use any of the agents listed above, VFM uses the replication agent on the VFM server or the monitoring agent to transfer data.

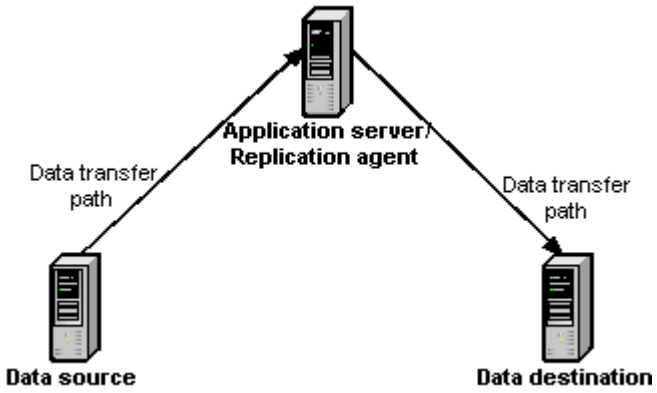

Proxy computers and replication agents

VFM enables you to specify a proxy computer for a replication agent that is used during data transfer when a replication agent cannot be deployed on the source or destination. The following list specifies the reasons for using a proxy computer:

- Use a proxy machine to avoid transferring data across a WAN.
- When the VFM or is not on the same LAN as the source and destination targets and you cannot deploy the replication agent to the destination or the source target, assign a proxy machine. If the replication agent on the VFM server was used in this scenario, the data transfer path would be from the source target to the VFM server and from the VFM server to the destination target. By deploying a proxy replication agent, you can avoid transferring the data across the WAN during the transfer operation.

Network ports in use

This section provides information about the network ports used by VFM, DFS, filers with NFS enabled, and other systems. The Agent Management tool enables you to change the listening port of replication and monitoring agents and the VFM server. Before you change the listening port for an agent or the VFM server, it is recommended that you review the list of network ports already in use.

Network ports used by DFS

The following table lists the network ports used by DFS and VFM.

TABLE 58 Network ports used by DFS

Service Name	UDP	TCP
NetBIOS Datagram Service Remote Server Message Block (SMB)	138	
NetBIOS Session Service		139
LDAP Serve	389	389
Procedure Call (RPC)		135
Server Message Block (SMB)	445	445

Network ports used by VFM

The following table lists the network ports used by VFM.

TABLE 59 Network ports used by VFM

Service Name	UDP	TCP
VFM Server or monitoring agent		6001 and 6005
Replication Agent		6002

Network ports used on filers with NFS enabled

The following table lists the network ports used on filers with NFS enabled.

TABLE 60 Network ports used on filers with NFS enabled

Service Name	UDP	TCP
NFS mount daemon (mountd)	602	603
NFS status daemon (statd,statmon)	604	605
NFS lock manager (lockd,nlockmgr)	606	607
NFS quota daemon (quotad,rquotad)	608	

Network ports used by other systems

The following table lists network ports used by other systems.

TABLE 61 Network ports used by other systems

Service Name	UDP	TCP
NFS mount daemon (mountd)	611	612
NFS status daemon (statd,statmon)	613	614
NFS lock manager (lockd,nlockmgr)	615	616
NFS quota daemon (quotad,rquotad)	617	

NOTE
The VFM server and UNIX Replication Agent communicate with storage devices using NFS during NFS migrations. The VFM server uses SNMP when identifying filers and uses RSH or SSH when issuing some commands to UNIX machines and filers.

Viewing status information

View status information about a replication or monitoring agent or the VFM server.

To view the status of an agent or the application server

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click a replication agent, monitoring agent, or the VFM server, and then click **Check Health**.
5. Click the **Alerts** tab to view errors.

Manage replication agents

This section provides instructions about how to deploy and manage Windows and UNIX replication agents.

Deploying a replication agent

VFM automatically deploys Windows and UNIX replication agents to manage data transfers. You can manually deploy agents.

To deploy a replication agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, click **Deploy**.
5. In the **Host** box, type the name or UNC path of the machine where you want to deploy a replication agent or click **Browse** to find a machine in your network. VFM determines whether to deploy a Windows or UNIX replication agent.
6. (Optional) Type a different port. VFM normally uses port 6002 to communicate with replication agents.

   NOTE
   If you change the listening port, make sure that the port you select enables the VFM server to communicate through any firewalls with the VFM client and replication agents the server deployed.

7. Specify the credential that VFM will use to communicate with the replication agent. Perform one of the following actions:
   - Select **The default replication agent user name** when you want to use the default credentials specified by VFM or use the default credentials you specified.

     NOTE
     For application default credentials, deployed Windows replication agents use the VFM server credentials. Deployed UNIX replication agents use root with no password. For user specified default credentials, to specify default credentials for Windows and UNIX replication agents that are deployed, refer to"Setting replication agent default credentials" on page 305.

   - Select **The user name and password specified below** and type new credentials.
8. Click **OK**.
9. Click the **Alerts** tab to verify that the replication agent was deployed successfully.
10. Click **Refresh** in the **Host** view to display the new replication agent.

Setting replication agent default credentials

When you specify default Windows and UNIX replication agent credentials, all replication agents deployed after you change the credentials will use the default credentials. Replication agents that are already deployed are not affected.

To set default credentials

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the Host view, click **Manage Credentials**.
5. Click **Change** for the appropriate system.
6. Type a new user name.
7. Type the password.
8. Type the password again to confirm.
9. Click **OK**.
10. Repeat step 5 through step 9 to change the credentials for replication agents deployed to another computer.
11. Click **OK**.

NOTE
To change the credentials of deployed replication agents, refer to "Changing the credentials of a replication agent" on page 307.

Stopping and restarting a replication agent

When you stop a replication agent while it is transferring data, the data transfer terminates. Restarting the agent does not restart the data transfer. The data transfer restarts when you run the policy manually or during the next scheduled run of the policy.

To stop an agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click a replication agent and click **Stop**.

To restart an agent

In the Host view, right-click the replication agent and click **Start**.

Pausing and resuming a replication agent

When you pause a replication agent while it is transferring data, the data transfer is suspended and resumes when the agent starts again.

To pause an agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click a replication agent and click **Pause**.

To resume an agent

To resume the agent, right-click the agent in the **Host** view, and then click **Resume**.

Upgrading a replication agent

When you upgrade to a new version of VFM, you *must* upgrade all deployed replication agents. If the replication agents are not upgraded, the policies that use the agents to perform data transfer will be disabled.

To upgrade an agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click a replication agent and click **Upgrade**.
5. To view the progress of the update, click **Tasks** in the Status view.
6. To view information about errors that might occur, click **Alerts** in the Status view.

NOTE
You must use the VFM media to upgrade installed monitoring agents. If you get the error message `Network Resource Not Found` when you try to upgrade the replication agent installed on the same machine as the VFM server, refer to "General problems" on page 368.

Grouping replication agents

When you want to distribute a migration workload or increase the availability of agents during migration, assign replication agents to groups and then assign the agent group to a policy.

When the policy runs, VFM selects a replication agent from the group that is not running a migration operation to run the migration. When a machine hosting a replication agent is not available, VFM uses another agent within the group to run the migration. When a migration operation contains multiple migration tasks, VFM distributes the workload among the replication agents in the group. A replication agent can belong to only one agent group.

There are no limits to the number of agents that an agent group can contain.

To create a replication agent group

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, click a replication agent.
5. In the **Properties** view, perform one of the following actions:
   - Type a name for the group in the **Group** text box.
   - Click the down arrow, click **New**, and type a name for the group.
6. (Optional) Type a word or phrase that best describes the group.
7. Click **OK**. The replication agent is assigned to the group and the group name is added to the list.

To assign a replication agent to an existing group

1. In the **Host** view, click a replication agent.
2. In the **Properties** view, click the **Group** arrow and select a group from the list.

To view the replication agents assigned to a group

In the **Host** view, click the **Group** column to sort the list alphabetically by agent group.

NOTE
Migration, replication, archival migration, storage load balancing, and data user management policies enable you to select which replication agent or agent group performs data migration.

Changing the credentials of a replication agent

Agent Management enables you to change the credentials of deployed Windows and UNIX replication agents. Before you change the credentials of an agent, review the following list of requirements:

- When deployed Windows replication agents are using the credentials of the VFM server and you change the credentials of the server, you must update the credentials of all the Windows replication agents deployed by the server. If you do not update passwords, you cannot perform replication and monitoring on already deployed agents.
- When you change the default agent credentials for Windows or UNIX replication agents, you must update the credentials of deployed agents that are using the default credentials. If you do not update the credentials, you cannot perform replication and monitoring on the agents.
- If you installed a monitoring agent and you change the credentials of the agent, you must update the credentials of all the replication agents deployed by the monitoring agent.
- You can change the credentials for one agent or you can change the credentials for multiple agents and the VFM server.

To change the credentials of an agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click a replication agent and click **Change Credentials**.
5. Select one of the following options:
   - The default replication agent user name. Run the replication agent using the default replication agent credentials. To specify default replication agent credentials for all Windows and UNIX replication agents that you deploy, refer to "Setting replication agent default credentials" on page 305.
   - One of the credentials in the credentials vault. VFM stores all the credentials that are used or were used to run replication and monitoring agents or the VFM server. Select one of these credentials to run the replication agent
   - The credentials specified below. Before you specify new credentials for the replication agent, make sure that the account has the correct permissions and privileges.

Changing the listening port of a replication agent

VFM uses the TCP port 6002 for replication agents and TCP ports 6001 and 6005 for the VFM server. You have the option of changing the ports. Make sure that the port you select enables the replication agent to communicate with the VFM server through firewalls.

To change the listening port of an agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. *If you want to change the listening port of multiple replication agents*, perform the following actions:
   a. In the **Host** view, select the replication agents.
   b. In the **Properties** view, select **Multiple Values** in the **Port** list and type the new port number.
   c. Click **Save**.

5. *If you want to change the listening port of a replication agent*, perform the following actions:
   a. In the **Host** view, select a replication agent.
   b. In the **Properties** view, type the new port number.
   c. Click **Save**.

Purging a replication agent

Remove a replication or monitoring agent from the VFM database. Use this option when you cannot uninstall an agent because the machine hosting the agent does not exist or was renamed.

To purge an agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click the agent you want to remove from the list and click **Purge**. The agent is removed from the database and from the list of agents.

Uninstalling replication agents

When VFM or a monitoring agent is uninstalled, replication agents deployed by the VFM server or monitoring agent are not deleted. You have the option of uninstalling replication agents.

To uninstall a replication agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click the replication agent and click **Uninstall**.
5. Select one of the following options:
   - **Uninstall using the same credentials that were used to deploy each of the selected replication agents**. Select this option when the credentials used to run the replication agent have the security permissions to uninstall the agent.
   - **Uninstall the selected agents using the following credentials**. Select this option when the credentials used to run the replication agent do not have the security permissions to uninstall the agent.
6. If you selected to uninstall the replication agent using different credentials, type the credentials.
7. Click **OK**.
8. Click the **Alerts** tab to verify that the replication agent was uninstalled successfully.
9. In the **Host** view, click **Refresh** to remove the replication agent from the list.

NOTE
If you have problems uninstalling a replication agent, check whether the option **Allow replication agent to be deployed to this host** is cleared in the properties of the machine hosting the agent.

Installing and managing UNIX replication agents

This section describes how you can install UNIX replication agents and manage the agents using command line commands. VFM automatically deploys UNIX replication agents on the source or destination when transferring NFS data. You can use the Agent Management tool to manually deploy and manage UNIX and Windows replication agents.

NOTE
For instruction about how to deploy Windows and UNIX replication agents, refer to "Deploying a replication agent" on page 304.

Prerequisite

Before you install a UNIX replication agent, review account permissions for running VFM. Refer to "Replication agent requirements" on page 299.

Working with UNIX agents on Red Hat Enterprise systems

Use these instructions to install and manage replication agents on Red Hat systems.

To install a replication agent on a Red Hat system

1. To get a copy of the installation package, perform one of the following actions:
   - In the VFM installation directory, click the `Agents\ReplicationAgent\Linux\` folder.
   - Extract the installation package from the VFMInstall.zip or VFMInstall.exe file.
2. Copy the installation package *vfmra-product version number-product build number.i386.rpm*, where *product version number* is the current version of VFM and *product build number* is the current build of VFM, to the UNIX machine that will host the replication agent.
3. *View details of the installation package.* In these instructions, product version **99.9** is used and build number **1** is used. You must substitute these numbers with the current version and build of VFM. To view details of the installation package without installing, type the following command:

   ```
   rpm -qilp vfmra-99.9-1.i386.rpm
   ```

4. *Install the replication agent.* Type the command listed below to install the UNIX replication agent This determines where the binaries and log files of the replication agent are kept. In the following example, the binaries are installed in /opt/URA/bin and the logs and registry are installed in /opt/URA/logs.

   ```
   rpm -ihv vfmra-99.9-1.i386.rpm
   ```

5. *Configure the replication agent.* After you install the replication agent, you must configure it. The following example shows how to configure the agent.

   ```
   /etc/init.d/ura configure
   ```

You are prompted to enter the following information:

- The TCP port where the replication agent and the VFM server listen.
- The name of your VFM servers.

You can type the name or IP address of the VFM server. You must update the replication agent whenever the information you entered changes. Perform the following actions:

a. Stop the replication agent

b. Configure the agent

c. Start the agent again.

6. *Start the replication agent.* Type the following command:

```
/etc/init.d/ura start
```

To verify that the replication agent is running, type the following command:

```
pgrep -fl Repl
```

To stop a replication agent on a Red Hat system

Type the following command to stop the replication agent:

```
/etc/init.d/ura stop
```

To remove a replication agent on a Red Hat system

1. Stop the replication agent.
2. Type the following command to remove the replication agent:

```
rpm -evv vfmra-99.9-1
```

3. Type the following command to remove the directory /opt/URA. This directory contains the configuration files created by the agent.

```
rm -Rf  /opt/URA
```

To upgrade a replication agent on a Red Hat system

To install a new version of the package, you must uninstall the old package and then install the new package:

1. Stop the replication agent.
2. Type the following command to remove the replication agent:

```
rpm -evv vfmra-99.9-1
```

3. Install the updated replication agent. For instructions, refer to the topic "To install a replication agent on a Red Hat system" on page 310

To show all the packages installed, type the following command:

```
rpm -qa | grep vfm
```

Working with UNIX replication agents on Solaris systems

VFM deploys UNIX replication agents to the directory */opt/URA*. The following instructions show how to install a replication agent on this directory. To install a replication agent to another directory, refer to "To install a UNIX replication agent to another directory on Solaris" on page 313.

The package name is vfmra-*product version number-product build number*.pkg, where *product version number* is the current version of VFM and *product build number* is the current build of VFM. In these instructions, product version **99.9** is used and product build number **1** is used. You must substitute these numbers with the current version and build of VFM.

To install a replication agent on a Solaris system

1. To get a copy of the installation package, perform one of the following actions:
   - In the VFM installation directory, click the `Agents\ReplicationAgent\Solaris` folder.
   - Extract the installation package from the VFMInstall.zip or VFMInstall.exe file.
2. Copy the installation package to the UNIX machine that will host the replication agent.
3. Unzip the package file. At a shell prompt, type the following command. Make sure you substitute **99.9** with the current version of VFM and **1** with the current build of VFM.

   ```
   gunzip vfmra-99.9-1.pkg.gz
   ```

4. Type the following command:

   ```
   pkgadd -d ./vfmra-99.9-1.pkg
   ```

   The replication agent is installed in the */opt/URA* directory.
5. After you install the replication agent, you must configure it. The following example shows how to configure the agent.

   ```
   /etc/init.d/ura configure
   ```

   You are prompted to enter the following information:
   - The TCP port where the replication agent and the VFM server listen
   - The name of your VFM servers. You can type the name or IP address of the application server.

   You must update the replication agent whenever the information you entered changes. Perform the following actions:

   a. Stop the replication agent.

   b. Configure the agent.

   c. Start the agent again.
6. Type the following command to start the replication agent:

   ```
   etc/init.d/ura start
   ```

   To verify that the replication agent is running, type the following command:

   ```
   ps -ef|grep| Repl
   ```

To view packages currently installed

- To view all packages, type the following command: `pkginfo`

- To view detailed information about a package, type the following command:

  ```
  pkginfo -l NTAPura
  ```

To search for the replication agent package

Type the following command:

```
pkginfo | grep NTAP
```

To stop the replication agent

Type the following command:

```
/etc/init.d/ura stop
```

To remove a replication agent

1. Stop the replication agent.
2. Type the following command to remove the replication agent:

   ```
   pkgrm NTAPura
   ```

3. Type the following command to remove the directory `/opt/URA`. This directory contains the configuration files created by the agent.

   ```
   rm -Rf  /opt/URA
   ```

To update a replication agent

To install a new version of the package, you must uninstall the old package and then install the new package.

1. Stop the replication agent.
2. Type the following command to remove the replication agent:

   ```
   pkgrm NTAPura
   ```

3. Install the updated replication agent. For instructions, refer to the topic "To install a replication agent on a Solaris system" on page 312.

To install a UNIX replication agent to another directory on Solaris

VFM automatically deploys UNIX replication agents to the directory `/opt/URA`. You can install a replication agent to another directory.

NOTE

If you install a UNIX replication agent to another directory and VFM upgrades the replication agent, the new replication agent is installed in the default directory `/opt/URA`.

---

1. To get a copy of the installation package, perform one of the following actions:
   - In the VFM installation directory, click the `Agents\ReplicationAgent\Solaris` folder.
   - Extract the installation package from theVFMInstall.zip or VFMInstall.exe file.
2. Copy the installation package to the UNIX machine that will host the replication agent.

3. Unzip the package file. At a shell prompt, type the following command. Make sure you substitute **99.9** with the current version of VFM and **1** with the current build of VFM.

   ```
   gunzip vfmra-99.9-1.pkg.gz
   ```

4. Type the following command:

   ```
   # pkgadd -a none -d ./vfmra-99.9-1.pkg
   ```

   You will see the message `Enter path to package base directory [?,q]`.

5. Type the location of the directory, for example:

   ```
   /opt/Software/URA
   ```

6. After you install the replication agent, you must configure it. The following example shows how to configure the agent.

   ```
   /etc/init.d/ura configure
   ```

   You are prompted to enter the following information:

   - The TCP port where the replication agent and the VFM server listen
   - The name of your VFM servers. You can type the name or IP address of the application server.

   You must update the replication agent whenever the information you entered changes. Perform the following actions:

   a. Stop the replication agent.

   b. Configure the agent.

   c. Start the agent again.

7. Type the following command to start the replication agent:

   ```
   etc/init.d/ura start
   ```

   To verify that the replication agent is running, type the following command:

   ```
   ps -ef|grep| Repl
   ```

Manage monitoring agents

This section explains why you should consider distributing monitoring agents and provides instruction about how to install and manage the agents on Windows machines.

Distributed monitoring

Consider distributing monitoring agents based on your network topology. You can select another monitoring agent that you installed at another location and configure it to work with your VFM server. Monitoring agents should have the same network connectivity as clients for the link targets they monitor.

Checklist

VFM enables you to install and manage distributed monitoring agent.

- To install an agent on a machine, refer to "Installing a monitoring agent" on page 315.
- To update credentials of a monitoring agent, refer to "Changing monitoring agent credentials" on page 315.
- To change the listening port of a monitoring agent, refer to "Changing the listening port of a monitoring agent" on page 316.
- To assign a monitoring agent to run a policy, refer to "Assigning a monitoring agent to a policy" on page 316.
- To specify where you want to store the configuration files of a monitoring agent and the policies assigned to the agent, refer to "Specifying a backup location for monitoring agent files" on page 317.
- To remove a monitoring agent from the VFM database, refer to "Purging a monitoring agent" on page 317.
- To view information about a deployed agent, refer to "Viewing status information" on page 304.

Installing a monitoring agent

The system requirements for installing and running a monitoring agent are the same as those for the VFM server. Before you install a monitoring agent, review the VFM system requirements. For more information, refer to "System requirements" on page 34.

To install a monitoring agent

Use the VFM media to install a monitoring agent.

NOTE
Before you uninstall a monitoring agent, make sure that policies assigned to the agent are assigned to another monitoring agent or the VFM server because the policies are disabled after the monitoring agent is uninstalled.

Changing monitoring agent credentials

When you change the password on the account where a monitoring agent is installed, you must update the credentials of the monitoring agent. If you do not update credentials, the monitoring agent ceases to function correctly because Windows security prevents the service from authentication with the system. Before you change the credentials of a monitoring agent, review the following list of requirements:

- When you upgrade to a newer version of VFM, you must upgrade all installed monitoring agents.
- You specify the credentials for a monitoring agent during the installation process of the agent. This account information is used when the monitoring runs.
- You can change the credentials for one agent or you can change the credentials for multiple agents (replication and monitoring) and the VFM server.

To change the credentials of a monitoring agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click the agent and click **Change Credentials**.
5. Select one of the following options:

   TABLE 62 Change credentials

Option	Description
One of the credentials in the credentials vault	VFM stores all the credentials that are used or have been used to run replication and monitoring agents or the VFM server. Select one of these credentials to run the monitoring agent.
The credentials specified below	Before you specify new credentials for the monitoring agent, make sure that the account has the correct permissions and privileges.

6. Click **OK**.

Changing the listening port of a monitoring agent

Make sure that the port you select enables the monitoring agent to communicate through firewalls with the VFM client and all the replication agents deployed by the monitoring agent.

To change the listening port

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, click the VFM server or a monitoring agent.
5. In the **Properties** view, type the new TCP port.

Assigning a monitoring agent to a policy

The VFM server runs policies. You can select a monitoring agent installed on another machine to run a policy.

Uninstalling a VFM monitoring agent disables any policies assigned to that agent.

Reinstalling the agent on a given host does not automatically reenable policies using the agent on that host. To reenable these policies, you must assign the newly installed agent to run the policies by reselecting the desired agent to be used by the policy through the properties dialog box for the policies.

To assign a monitoring agent to a policy

1. In the **Admin View**, right-click a policy and click **Properties**.
2. Click **General**.

3. On the **Use the Monitor on** list, select the host name of the monitoring agent where monitoring is performed. The list displays the host names of the VFM server and all the VFM monitoring agents that are installed and configured to communicate with the VFM server.
4. Click **OK**.

Specifying a backup location for monitoring agent files

When you schedule backups of the application, the configuration files of the application and all installed monitoring agents are saved. The configuration files of monitoring agents are saved in folders under the location you specify to save the application configuration files. You can specify a different location to store snapshots of the state of a monitoring agent and the policies assigned to the agent.

To specify a backup location

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, click a monitoring agent to show its properties.
5. In **Snapshot Save location**, type the UNC path of the location where you want to save the configuration files of the monitoring agent or click **Browse** to find a location. This should be a location that is different than the location you specified to save the application.

For snapshot location requirements, refer to "Specifying a backup location for the application server" on page 319.

Purging a monitoring agent

Remove a monitoring agent from the VFM database. Use this option when you cannot uninstall an agent because the machine hosting the agent does not exist or was renamed.

To purge an agent

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click the agent you want to remove from the list and click **Purge**. The agent is removed from the database and from the list of agents.

Manage the application server

This section provides information about how you can change the VFM server credentials, the listening port of the server, specify a location to back up the server configuration files, and stop and restart the server.

Checklist

- To update the account information for the VFM server, refer to "Changing the application server credentials" on page 318.
- To change the listening port for the VFM server, refer to "Changing the listening port of a monitoring agent" on page 316.
- To specify the location where you want to back up the VFM configuration files, refer to "Specifying a backup location for the application server" on page 319.

Changing the application server credentials

During the initial installation process of VFM, you specify the account and password for the VFM server or monitoring agent. This account information is used when these servers run, as well as when they deploy replication agents to transfer data.

Before you change the credential of the application server, review the following list of requirements:

- When you change the password on the account where the VFM server is installed, you must update the credentials of the server. If you do not update credentials, the server ceases to function correctly because Windows security prevents the service from authenticating with the system.
- After updating the credentials of the server, you must update the password of all the replication deployed by the server. Refer to "Changing the credentials of a replication agent" on page 307.
- You can change the credentials for the VFM server only or you can change the credentials for the server and multiple agents (replication and monitoring).

To change the application server credentials

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, right-click the VFM server and click **Change Credentials**.
5. Select one of the following options

   TABLE 63

Option	Description
One of the credentials in the credentials vault	VFM stores all the credentials that are used or were used to run replication and monitoring agents or the VFM server. Select one of these credentials to run the server.
The credentials specified below	Before you specify new credentials for the VFM server, make sure that the account has the correct permissions and privileges.

6. Click **OK**.

NOTE
During the process of changing the server credentials, the VFM server service is stopped and restarted.

Changing the application server listening port

Make sure that the port you select enables the server to communicate through firewalls with the VFM client and all the replication agents deployed by the server.

To change the listening port

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, click the VFM server.
5. In the **Properties** view, type the new TCP port.

Specifying a backup location for the application server

Specify the location where you want to save the VFM configuration files. Before you back up the application server, review the following list of requirements:

- The location selected for the application snapshot must be writable by the account under which MSDE is running and the account under which the VFM server is installed. By default, MSDE runs as LocalSystem. Share permissions and directory permissions configured on the snapshot save location must allow both LocalSystem and the account under which the VFM server is installed to write to the directory.
- If the destination directory already exists, it must be empty.
- Application snapshots cannot be saved on a mapped network drive.
- Store the exported files in an appropriate location where you can access them if the machine hosting the VFM server fails.
- For instructions about how to schedule backups or manually back up the application, refer to "Backing up VFM configuration files" on page 337.

To specify a backup location

1. On the **Tools** menu, click **Agent Management**.
2. Click **Control Panel**.
3. In the **Control Panel** view, click **Agent Management**.
4. In the **Host** view, click the VFM server to show its properties.
5. In **Snapshot Save location**, type the UNC path of the location where you want to save the configuration files of the VFM server or click **Browse** to find a location.

11 Manage the application server

Chapter

Run Batch Files

12

Configure batch files to execute before and after you run a storage policy, run a report, or take a snapshot of the application.

In this chapter

- Using batch files in VFM 321
- Batch files you can run in VFM 323

Using batch files in VFM

Batch files enable you to customize policies, such as sending an email, posting an event in the Windows event log, or sending an SNMP trap in response to a failure. This section discusses how VFM processes batch files, the conditions under which batch files run, testing scripts, and common script problems that you might encounter.

How VFM processes batch files

This section describes how batch files are processed and explains how you must enter the UNC path of batch files you want to run in VFM.

The batch files are run by VFM on the machine from which the targets of a policy are being monitored by VFM, whether by the VFM server or by a monitoring agent to which the policy is assigned. Batch files related to reports are run by the VFM server.

VFM uses the directory where a batch file is located as the default/current directory where the command is executed. If a batch file is specified through a UNC path, the TMP directory of the server executing the policy is used as the default/current directory.

VFM checks the return value from the batch script that is run before failover. If the batch script returns failure, VFM does not continue the failover process. This allows customized checking to be performed before the actual failover takes place.

If the UNC path to the batch file contains spaces, you must enclose the entire path in quotation marks. You must follow the same rules that apply when running commands from a Windows command prompt.

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

Conditions under which scripts run

Scripts run under the following conditions:

- The script specified in the **Run Batch Before Replication** field, located in the properties of policies that move data, runs *before* the following migration actions:
  - Scan for potential problems copying the data in the initial phase.
  - Perform a baseline copy in the initial phase.
  - Copy the data in the incremental phase.
  - Do a final replication prior to deleting the source data in the final phase.
  - Delete the source data in the final phase.
- The script specified in the **Run Batch After Replication** field, located in the properties of policies that move data, runs after the migration actions listed previously.
- The script specified in the initial phase runs before the actions in the initial phase run, and the script specified in the final phase runs after the actions in the final phase are completed.
- When the script specified in the **Run Batch Before Replication** field or the script specified in the *initial phase* runs without errors, the script specified in **Run Batch After Replication** field or the script specified in the *final phase* is executed.
- When the script specified in the **Run Batch Before Replication** field or the script specified in the *initial phase* returns an error, the replication is not executed and the script specified in the **Run Batch After Replication** field or the script specified in the *final phase* is not executed.

Testing scripts

Test the scripts to verify that they are functioning as expected. You must do this by specifying small test directories for the source and destination of the synchronization task and configuring the scripts to write diagnostic output to a text file for review. After the scripts function correctly in the test configuration, you can configure them for use in actual policies, and you can reduce the amount of diagnostic information logged by the scripts to match the requirements of your environment.

Common script problems

This section discusses common problems you might encounter when working with scripts.

The path to the batch file must be wrapped in quotation marks if that path contains spaces. If you browse to the script location, VFM automatically enters the quotation marks. If you enter the path to the batch file manually, you must enter the quotation marks.

Batch files run with their working directory set to the directory containing the batch file when they are specified by a local path (for example, `c:\bats\myfile.bat`). When the batch file is referenced by a UNC path, its working directory is the system32 of the server or monitoring agent that runs the batch file.

Any file invoked by VFM must be a batch file. It is possible to have a top-level batch file invoke other tools, such as `jscript` files or other executables.

The arguments passed to the script are slightly different than expected. It can be very helpful to have your script log into a text file the time it was invoked along with the arguments it received.

Multiple copies of a script might invoke simultaneously, such as pre- and post-migration task batch files, when tasks are running at the same time. This can also occur in disaster recovery and replication scripts.

VFM uses the exit value of the script to assess whether the script succeeded or failed. Scripts must set the return value correctly to signal either success or failure.

If your script is not producing the desired results, you can insert different exit codes throughout the script to determine where exactly the script is exiting. You can also insert logging statements showing progress through the script to also trace execution.

Scripts must not prompt for user input. There is no console available to display the script prompt at run time. Scripts that prompt for input and the policy executing the script will hang indefinitely and might cause the application to hang. This is true for any prompt, whether in an error path or in a normal execution path in the script.

VFM *does not validate the path* to the script you specify. It is your responsibility to ensure that the path is accurate and accessible to the VFM server or monitoring agent that invoke the script.

Some antivirus programs display a message every time a script is invoked, inquiring whether you want to enable the script to run. VFM server components execute scripts for snapshots taken of the application and batch files configured to run with storage policies. When scripts are invoked by VFM server components, the checks of the antivirus program interfere with the operation of the server. You should configure your antivirus program to enable scripts used by VFM to run unimpeded.

Batch files you can run in VFM

VFM provides example scripts in the installation folder under `Examples\Scripts`. These examples rely on the presence of additional third-party software, such as an RSH client and other utilities, to function.This section provides detailed information about scripts you can create and run in VFM.

Before updating index file

The script is a report publishing option. The UNC path of the script is entered in the **Run batch before updating index file** field on the **Tools** > **Options** > **Reporting** dialog box.

Argument

The path to the index.xml file that VFM creates when publishing report.

Exit codes

0	Success
all others	Error

After updating index file

The script is a report publishing option. The UNC path of the script is entered in the **Run batch after updating index file** field on the **Tools** > **Options** > **Reporting** dialog box.

Argument

The path to the index.xml file that *VFM* creates when publishing report.

Exit codes

0	Success
all others	Error

Criteria script

The script is an archival migration option. The UNC path of the script is entered in the **Confirmed by user selection criteria script** field on the **Migration Criteria** tab of the policy properties.

Argument

1. Policy name
2. UNC of the path to evaluate
3. first user argument

Exit codes

0	Success
all others	Error

Migration task creation script

The script is a migration policy option. The UNC path of the script is entered in the **Run task creation batch** field on the **General** tab of the policy properties.

Argument

1. Policy name
2. Local path to a file to which script can output a list of tasks to be added. The local path is relative to the VFM server or monitoring agent that invokes the script.
3. first user argument

Exit codes

0	Success
all others	Error

Pass report output to script

The script is an option in the Report wizard. The UNC path of the script is entered in the **Script** field on the **Report Output** dialog box of the Reports wizard.

Argument

1. The path to a TMP file containing the report in the selected format. It is expected that the script file will copy the contents of the file to a new location for subsequent processing. VFM deletes the temporary file after the report output script file runs.Local path to a file to which script can output a list of tasks to be added. The local path is relative to the VFM server or monitoring agent that invokes the script.
2. The data collection process shows the following states:
   - OK: All the data required for the report was successfully collected
   - Warning: Not all the data required for the report was collected. The information passed to the script file might not be complete. You need to refresh the associated report group to force collection of the data.
   - Error: Problems occurred while collecting the data for the report. The information passed to the script file might not be complete
3. first user argument

Exit codes

0	Success
all others	Error

Pre-creation script

The script is an option in the **Share Creation** wizard. When you select to create duplicate shares, the UNC path of the script is entered in the **Run pre-creation batch file** field.

Argument

1. UNC path of source share
2. UNC path of destination share
3. Local path of the destination share

Exit codes

0	Success
all others	Error

Pre-data gathering script

The script is an option in the report schedule group and report source group. The UNC path of the script is entered in the **Run batch before gathering data** field on the **Scripts** dialog box of the **Report Schedule Group** or **Report Source Group** properties.

Argument

1. Report source group name
2. first user argument

Exit codes

0	Success
all others	Error

Post-data gathering script

The script is an option in the report schedule group and report source group. The UNC path of the script is entered in the **Run batch after gathering data** field on the **Scripts** dialog box of the **Report Schedule Group** or **Report Source Group** properties.

Argument

1. Report source group name
2. first user argument

Exit codes

0	Success
all others	Error

Pre-Design phase script

The script is an option in the archival migration, migration, and storage load balancing policies. The UNC path of the script is entered in the **Run batch before migration task** field on the **Task Scripting** tab of the policy properties.

Argument

1. Path of source
2. Path of destination
3. first user argument

Exit codes

0	Success
all others	Error

Post-Finalization phase script

The script is an option in the archival migration, migration, and storage load balancing policies. The UNC path of the script is entered in the **Run batch after migration task** field on the **Task Scripting** tab of the policy properties.

Argument

1. Path of source
2. Path of destination
3. copy=[pass|fail|skip]
4. updatedfs=[pass|fail|skip]
5. delete=[pass|fail|skip]
6. first user argument

Exit codes

0	Success
all others	Error

Pre-replication script

The script is an option in the archival migration, client disaster recovery, migration, replication, server disaster recovery, and storage load balancing policies. The UNC path of the script is entered in the **Run Batch Before Replication** field on the **Replication Script** tab of the policy properties.

Argument

1. UNC path of link (for replication policies this is the destination link)
2. UNC path of source
3. UNC path of destination
4. first user argument

Exit codes

0	Success
all others	Error

Post-replication script

The script is an option in the archival migration, client disaster recovery, migration, replication, server disaster recovery, and storage load balancing policies. The UNC path of the script is entered in the **Run Batch After Replication** field on the **Replication Script** tab of the policy properties.

Argument

1. UNC path of link (for replication policies this is the destination link)
2. UNC path of source
3. UNC path of destination
4. The time that it took to perform the operation
5. Result of the replication:
   - Pass: Operation succeeded
   - Fail: Operation encountered errors
   - Cancel: Operation canceled by the user
   - Expired: Operation was not completed in the scheduled time window
6. Number of files copied
7. Number of bytes copied
8. Number of files already in sync
9. Number of bytes already in sync
10. first user argument

Exit codes

0	Success
all others	Error

Pre-failover script

The script is an option in the client disaster recovery, filer disaster recovery, and server disaster recovery policies. The UNC path of the script is entered in the **Run Batch before failover** field on the **DR Script Settings** tab of the policy properties.

Argument

1. Policy name
2. (Not used. Will be empty)

3. UNC path of the link being failed over
4. Invocation type [/failover|/failback]
5. first user argument

Exit codes

0	Success
all others	Error

Post-failover script

The script is an option in the client disaster recovery, filer disaster recovery, and server disaster recovery policies. The UNC path of the script is entered in the **Run Batch after failover** field on the **DR Script Settings** tab of the policy properties.

Argument

1. Policy name
2. UNC path of active target
3. UNC path of the link being failed over
4. Invocation type [/failover|/failback]
5. first user argument

Exit codes

0	Success
all others	Error

Pre-scan script

The script is an option in the archival migration policy. The UNC path of the script is entered in the **Run batch before scan** field on the **General** tab of the policy properties.

Argument

1. Policy name
2. Local path to a file to which script can output a list of tasks to be added. The local path is relative to the VFM server or monitoring agent that invokes the scriptUNC path of the link being failed over
3. first user argument

Exit codes

0	Success
all others	Error

Post-scan script

The script is an option in the archival migration policy. The UNC path of the script is entered in the **Run batch after scan** field on the **General** tab of the policy properties.

Argument

1. Policy name
2. first user argument

Exit codes

0	Success
all others	Error

Pre-snapshot script

The script is an option when you take a snapshot of the application. The UNC path of the script is entered in the **Run batch before snapshot** field on the **Tools** > **Options** > **Application Snapshot Options** dialog box.

Argument

first user argument

Exit codes

0	Success
all others	Error

Post-snapshot script

The script is an option when you take a snapshot of the application. The UNC path of the script is entered in the **Run batch after snapshot** field on the **Tools** > **Options** > **Application Snapshot Options** dialog box.

Argument

first user argument

Exit codes

0	Success
all others	Error

Pre-snapshot script

The script is an option in the snapshot scheduling policy. The UNC path of the script is entered in the **Run batch before snapshot** field on the **VFM Settings** tab of the policy properties.

Argument

1. Policy name
2. File name
3. Volume name
4. first user argument

Exit codes

0	Success
all others	Error

Post-snapshot script

The script is an option in the snapshot scheduling policy. The UNC path of the script is entered in the **Run batch before snapshot** field on the **VFM Settings** tab of the policy properties.

Argument

1. Policy name
2. File name
3. Volume name
4. first user argument

Exit codes

0	Success
all others	Error

Pre-sync script

The script is an option in the multiprotocol namespace policy. The UNC path of the script is entered in the **Run Batch Before Sync** field on the **Scripts** tab of the policy properties.

Argument

1. Policy name
2. Automount files location
3. first user argument

Exit codes

0	Success
all others	Error

Post-sync script

The script is an option in the multiprotocol namespace policy. The UNC path of the script is entered in the **Run Batch After Sync** field on the **Scripts** tab of the policy properties.

Argument

1. Policy name
2. Automount files location
3. first user argument

Exit codes

0	Success
all others	Error

User decision script

The script is an option in the client disaster recovery, filer disaster recovery, and server disaster recovery policies. The UNC path of the script is entered in the **Run User Decision Script** field on the **DR Script Settings** tab of the policy properties.

Argument

1. Policy name
2. Link name
3. Online target
4. Offline target
5. first user argument

Exit codes

0	Do not fail over
1	Do fail over
all others	Error (logged in event log)

Chapter

Server Availability

13

This section explains how to make your VFM server highly available. It discusses how to plan for VFM availability and explains the methods that you can use to make the server highly available.

In this chapter

- Plan for VFM availability 335
- Selecting a backup method 335
- Back up and restore VFM configuration files 336
- Backing up VFM configuration files 337
- Importing VFM configuration files 339
- Restoring VFM on the primary machine. 339
- Performing subsequent failover on the same backup machine 340

Plan for VFM availability

When you install VFM, consider the following aspects of your storage configuration:

- A reliable, highly available location for the VFM server software.
- How large is your VFM configuration.
- Whether you have storage policies that you want to run from locations that are different from where the VFM server is installed.
- Where the DFS roots that are managed by VFM are located.

For more information about planning for VFM availability, refer to "VFM availability considerations" on page 11.

Selecting a backup method

Select a method for backing up and restoring VFM server configuration files and the logical structure of all the roots in your namespace.

Use VFM to back up the namespace and schedule snapshots of the state of the VFM server or manually back up and restore the server files.

NOTE
An additional component of any availability plan is maintaining backups of the actual data contained on the servers participating in the DFS namespace.

Back up the namespace

Create a namespace backup policy to save backups of the logical structure of all the DFS roots in your namespace. This enables you to recreate the DFS namespace if a DFS root server is lost along with the VFM server machine. For more information, refer to "Create a namespace backup policy" on page 139.

Schedule snapshots of the state of the VFM server

You can schedule snapshots of the state of the VFM server. Snapshots can be used to restore the VFM server to a known state after unintended changes or disk failure on the system where the server is installed. When the VFM server is not available, you can restore the latest snapshot on another machine where VFM is installed and have it take over as the VFM server. For more information, refer to "To back up VFM configuration files automatically" on page 338.

Manually back up the VFM server configuration files

You can manually back up VFM configuration files. For more information, refer to the following tasks:

- "To back up VFM configuration files manually when you have access to the VFM client and VFM server is running" on page 338
- "To back up VFM configuration file manually when you do not have access to the VFM client or the VFM server is not running" on page 338

Back up and restore VFM configuration files

This section explains how you can back up and restore VFM configuration files. The section provides information that you should consider before you back up and restore the files.

Back up VFM configuration files

This section describes the information that VFM backs up and provides information about the backup process.

The backup contains the following information:

- The contents of your VFM tree view, including administrative policies, roots selected for display in the Logical View, Physical View contents, and reports.
- Any backup files you created of the logical structure of DFS roots in your namespace.
- The state of the VFM server.

You can back up the VFM configuration files to a VFM installation on another machine that you can bring up if the primary VFM machine becomes unavailable.

When the primary machine becomes available, you can back up the configuration files from the backup machine and restore the files back to the primary VFM machine.

When backing up data, you may want to create a namespace backup policy to save backups of the logical structure of all the DFS roots in your namespace. This enables you to re-create the DFS namespace if a DFS root server is lost along with the VFM server machine.

When backing up VFM information, you must save copies of any batch files you configured to run as part of your policies. You must make these files available to the backup server before it is activated. If you install the batch files in a different location on the backup server from where they were on the primary server, you must edit your policies to reference the batch files in their new locations

Restore VFM configuration to a known state

You can use the server configuration files you backed up to restore the VFM configuration on your primary machine to a known state, if necessary.

This section describes the process for restoring the VFM configuration files.

1. Choose to automate or manually back up VFM configuration files. For more information refer to "Selecting a backup method" on page 335.
2. If the machine where VFM is installed becomes unavailable, you install VFM on another machine and import the VFM configuration files to the backup machine. The backup machine takes ownership of all monitoring agents in the imported data.
3. When the primary VFM machine becomes available, you can perform one of the following actions:
   - Manually take a snapshot of the VFM configuration information on the backup VFM machine and specify the primary VFM machine as the location where the snapshot is saved. Use the VFM command-line utility to import the files to the primary VFM machine. For more information, refer to "Importing VFM configuration files" on page 339.
   - Use the VFM command-line utilities to export the VFM configuration files from the backup machine and import them back to the primary VFM machine. For more information, refer to "Restoring VFM on the primary machine" on page 339.

ATTENTION
You must stop all VFM processes, including any remote clients accessing the VFM server, before you can perform an import operation.

Backing up VFM configuration files

Back up your VFM server configuration files so that you can continue to manage your namespace if the machine on which the VFM server is installed becomes unavailable.

Before you export the VFM configuration files, review the following list of requirements:

- When you specify a snapshot location for the VFM server in the Agent Management tool, VFM uses the location specified in Agent Management instead of the location you specify when scheduling application snapshots.
- The location you type at which to manually back up the VFM configuration files overrides the snapshot location you specify for the VFM server in the Agent Management tool.
- The location selected for the application snapshot must be writable by the account under which MSDE is running and the account under which the VFM server is installed. By default, MSDE runs as LocalSystem. Share permissions and directory permissions configured on the snapshot save location must allow both LocalSystem and the account under which the VFM server is installed to write to the directory.
- The destination directory must exist and must be empty.

- Application snapshots cannot be saved on a mapped network drive.
- Store the exported files in an appropriate location where you can access them if the machine hosting the VFM server fails.

To back up VFM configuration files automatically

1. On the **Tools** menu, click **Options**.
2. Under System Options, click **Application Snapshot Options**.
3. Select Automatically snapshot the application.
4. Type the UNC path to the location at which you want to save the configuration files, or click **Browse** to locate the machine in your network.
5. Type the number of snapshots to retain.
6. Optionally, type the UNC path of the batch files that you want to run before and after the snapshot process.
7. Type how frequently you want to back up the files.
8. Click **OK**.

To back up VFM configuration files manually when you have access to the VFM client and VFM server is running

1. On the **Tools** menu, click **Take Application snapshot now**.
2. Type the UNC path to the location at which you want to save the configuration files, or click **Browse** to locate the machine in your network. You cannot save the snapshot on a mapped network drive.
3. Type the name of the folder where the files are saved.
4. Click **OK**.

To back up VFM configuration file manually when you do not have access to the VFM client or the VFM server is not running

1. Close the VFM client.
2. On the machine where the VFM server is running, start a command prompt.
3. Change directory to the VFM installation *Tools\Export* subfolder.
4. Type `export` *directory name*, where *directory name* is the path to the directory where you want to save the VFM configuration files.
5. Press Enter.

NOTE
Some antivirus programs might interfere with the ability of VFM to execute scripts. For more information, refer to "Using batch files in VFM" on page 321.

Importing VFM configuration files

If the machine where the VFM server is installed becomes unavailable, you can perform a manual failover to a backup machine. The following instructions assume that you previously exported your VFM configuration files using one of the methods described in this document. For more information, refer to "Backing up VFM configuration files" on page 337

Before you import the VFM configuration files, review the following list of requirements:

- You must stop all VFM processes, including any remote clients accessing the VFM server, before you can perform an import operation.
- When the SQL service is running on the Local Network Account on the backup machine, perform one of the following actions:
  - Save the exported snapshot on the local drive of the backup machine before you import the snapshot to VFM.
  - Make sure that the SQL service is running under an account that has full privileges across all domains.

To import VFM configuration files

1. Install VFM on a backup machine.
2. Close the VFM client.
3. On the machine where the VFM server is installed, start a command prompt and change to the VFM installation *Tools\Export* subfolder.
4. Type **`import`** *`directory name`* where *`directory name`* is the path to the directory in which the backup files are located.
5. Press Enter.
6. Start the VFM client and validate that the VFM configuration is correct by reviewing current policies and namespaces.
7. Review the properties of policies to make sure that policies are assigned to the proper server or monitoring agent.

Restoring VFM on the primary machine

When you are ready to resume running VFM on the primary machine, perform the following instructions. These instructions assume that you are running the VFM server on a backup machine and that you used the instructions described in this document to import the VFM configuration files to the backup machine.

To restore VFM on the primary machine

1. On the backup VFM server machine, perform the following actions:
   a. Export the VFM configuration files. For more information, refer to "Backing up VFM configuration files" on page 337.
   b. On the machine where the VFM server is installed, start a command prompt and change to the VFM installation `Tools\Export` subfolder.

c. Type `server disable` and press Enter to stop and disable the VFM server on the backup machine. The primary and backup VFM servers *cannot be active simultaneously*.

2. Ensure that VFM is operational and continue with step 2 in "Importing VFM configuration files" on page 339.

NOTE
The path you specify for the import operation should be the location of the files that you exported from the backup machine.

Performing subsequent failover on the same backup machine

This section explains how you can perform subsequent fail overs on the same backup machine whenever the primary VFM server machine is not available.

The following tasks must be completed before you can perform subsequent failovers:

- VFM server is stopped and disabled on the backup machine.
- You previously exported your VFM configuration files. For more information, refer to "Backing up VFM configuration files" on page 337.
- You followed the instructions in "Restoring VFM on the primary machine" on page 339 to resume operation on the VFM primary machine.

To perform subsequent failovers on the same VFM backup machine

1. Start a command prompt on the backup machine and change to the VFM installation *Tools\Export* subfolder.
2. Type `server enable` and press Enter to enable and start the VFM server on the backup machine.
3. Type **`import`** *`directory name`* where *`directory name`* is the path to the directory in which the backup files are located.
4. Press Enter.
5. Start the VFM client and validate that the VFM configuration is correct by reviewing current policies and namespaces.
6. Review the properties of policies to make sure that policies are assigned to the proper server or monitoring agent.

Appendix

Using RSH and SSH to Execute Shell Commands

A

VFM executes shell commands to configure and obtain data from NetApp filers and other non-Windows systems, such as Solaris and Linux systems. VFM usually uses RSH to execute these commands. Where the environment does not support the use of RSH, you can configure VFM to use SSH communications in place of RSH.

In this appendix

- Using RSH 341
- Using SSH. 343

Using RSH

VFM communicates with non-Windows computers using RSH. RSH must be enabled as root on the non-Windows machines. This section provides instructions about how to enable RSH as root on Linux and Solaris systems. You also have the option of using SSH to communicate with non-Windows computers. For more information, refer to "Using SSH" on page 343.

Enabling RSH as root on a Linux system

For VFM to communicate with Linux systems using RSH, the Linux systems must enable RSH as root.

To enable RSH as root on a Linux system

1. Make sure that the rshd server is installed. Type the following command to verify that the rshd server is installed:

   ```
   rpm -qa | grep rsh
   ```

   After executing the command, look for a package name, such as *rsh-server-0.17-14*. If you do not have this package, you can get it from the Red Hat media.

2. Edit */etc/xinetd.d/rsh.*

3. Find the line:

   ```
   disable = yes
   ```

   and change it to:

   ```
   disable = no
   ```

4. Restart Xinetd. Type the following commands:

   ```
   /etc/init.d/xinetd stop
   /etc/init.d/xinetd start
   ```

5. Edit the file *.rhosts* in the roots home directory or create it if it does not exist. Add the following line:

   ```
   + +
   ```

   Make sure that there are no leading spaces and that there is only one space between the two plus symbols. This format is important. Instead of the above line you can also add:

   ```
   machine1.npb.com +user1
   ```

   where *machine1.npb.com* is the name of the machine where the VFM server is installed and *user1* is the account used by the VFM server. This enables the VFM server to access the Linux system using RSH.

   NOTE
   The file *.rhosts* must have permission 600 and must be owned by root.

6. Edit the */etc/hosts.allow* file or create it if it does not exist. Add the following line at the end of the file and make sure that there are no leading spaces:

   ```
   ALL: ALL
   ```

7. Make sure that the file */etc/hosts. deny* is empty, or has only lines that begin with pound (#). These are comments and are ignored. If there are existing entries in *hosts.deny/hosts.allow*, check with your system administrator about allowing connection to rsh from specific hosts.

8. Edit the file */etc/pam.d/rsh* and change the following line:

   ```
   auth  required   pam_rhosts_auth.so
   ```

   to

   ```
   auth  sufficient pam_rhosts_auth.so
   ```

9. Add entries for `rsh`, `rlogin` and `rexec` in */etc/securetty* as shown in the following excerpt:

   ```
   tty10
   tty11
   rsh
   rlogin
   rexec
   ```

10. Check whether the kernel firewall is blocking incoming or outgoing packets. Type the following command:

    ```
    iptables --list
    ```

    If rules are listed, the firewall might be blocking connection attempts. Ask the firewall administrator to allow RSH ports through the firewall.

11. Verify RSH from another machine. Type the following command:

    ```
    rsh your host -l root "ls -la"
    ```

    where *your host* is the name of the machine to which you are trying to connect. If you see the directory listing, RSH as root is configured correctly.

Enabling RSH as root on a Solaris system

For VFM to communicate with Solaris systems using RSH, the Solaris systems must enable RSH as root.

To enable RSH as root on a Solaris system

1. Ensure that the *SUNWcsu* package is installed. The following command should display information about the package if it is installed:

   ```
   pkginfo -l SUNWcsu
   ```

2. Ensure that the r-services (for example, RSH, REXEC, and RLOGIN) are enabled in the file */etc/inetd.conf*.
3. Ensure that the file */etc/default/login* has *CONSOLE* variable commented out.
4. Edit the file *.rhosts* in the roots home directory or create it if it does not exist. Add the following line:

   ```
   + +
   ```

   Make sure that there are no leading spaces and that there is only one space between the two plus symbols. This format is important. Instead of the above line you can also add:

   ```
   machine1.npb.com +user1
   ```

   where *machine1.npb.com* is the name of the machine where the VFM server is installed and *user1* is the account used by the VFM server. This enables the VFM server to access the Solaris system using RSH.

   NOTE
   Make sure that the file has permission 600 and is owned by root.

5. Verify RSH from another machine. Type the following command:

   ```
   rsh your host -l root "ls -la"
   ```

   where *your host* is the name of the machine to which you are trying to connect. If you see the directory listing, RSH as root is configured correctly.

Using SSH

You can use SSH instead of RSH to communicate with non-Windows computers. The following sections define how to configure VFM to use SSH.

SSH prerequisites

Before you configure VFM to use SSH, make sure that the VFM client and server and computers known to VFM and involved in data transfer meet the following requirements:

- The following programs from PuTTY suite 0.55 must be installed on the VFM machine:
  - *Plink.exe*
  - *Puttygen.exe*
  - *pscp.exe*

- *pageant.exe*
- *putty.exe*

  The *putty.exe* program is optional, but useful in cases where Telnet is disabled on a filer or other computer with which you want to have a CLI session.

NOTES:

- If the VFM client, server, and monitoring agent are installed on different machines, you must install the PuTTY Suite programs on all three machines.
- The *Putty* folder is usually installed under *Program Files*, but it is recommended that you move the *Putty* folder to the root of the *C:* drive. Having a space in paths to *plink.exe*, *pscp.exe*, and the private key can cause problems in VFM, but *C:\putty* works well.
- Putty is case sensitive, so make sure you use the right case when you enter the user name.

- A UNIX machine must be running OpenSSH. For more information about how to install OpenSSH on Solaris 10 and Red Hat Enterprise 4, refer to "Installing and configuring OpenSSH on Solaris" on page 353 or "Installing and configuring OpenSSH on Red Hat" on page 355.
- A filer must have the SecureAdmin package installed. To verify that a filer has this package, look for the directory named */etc/sshd*. Data ONTAP 7.0 or later includes built-in SecureAdmin functionality with SSL protocol.
- When VFM uses SSH to communicate with UNIX machines and filers, SSH must be configured without a password or passphrase. However, note that Data ONTAP 6.3 does not support passwordless access through SSH.

For more information about SSH, OpenSSH, and PuTTY, refer to the following list of recommended references:

- Information about SSH implementation in UNIX and Windows:

  Barrett, David J., and Richard Silverman. *The Secure Shell: The Definitive Guide*. Sebastopol: O'Reilly Media, Inc., 2001.

- Information about SSH and PuTTY for Windows:

  http://www.jfitz.com/tips/ssh_for_windows.html

- Information about OpenSSH:

  http://www.openssh.org

You can obtain a copy of PuTTY from the following web site:
http://ftp.gnome.org/pub/putty/putty-0.55/x86/

SSH best practices for VFM

Review the following best practices for configuring SSH with VFM:

- Filer or remote host keys are cached in the registry of the currently logged in user. Since the VFM service account is the account used to do work on behalf of the user, you should log in to the server using the VFM service account.
- You can enable SSH1, SSH2, or both for NetApp filers. However, if you select both, only SSH2 requests will receive responses, and SSH1 requests will error out. If you just want to test and cache keys for both SSH1 and SSH2, ensure you enable only SSH1 on the filer before you cache the keys.

- You can set up SSH in two locations in VFM. The most important location is the master one under **Tools > Options > System Options > Remote Shell**. The master System Options settings override the settings in **Physical View > Filer > Properties > Shell**, at least where the private key is concerned. So, always make sure that the private key you are using is the correct one in the master System Options. However, if you set the Shell Type to be RSH in the Physical View properties, that setting overrides the SSH default option in the master settings.

Configuring VFM to use SSH

This section outlines the steps required to configure VFM to use SSH. Complete the steps that are applicable to your environment.

VFM communicates with UNIX computers

The instructions you follow will depend on whether the VFM client and server are installed on the same computer or separate computers.

To configure SSH when the client and server are installed on the same computer

The VFM server runs all your policies.

1. Generate a key-pair on the computer where the VFM client and server are installed. For more information, refer to "Generating a key-pair" on page 348.
2. Copy the public key to each UNIX computer known to VFM. For more information, refer to "Copying the public key to a UNIX computer" on page 349.
3. Cache the public keys onto each UNIX computer known to VFM. For more information, refer to "Caching the public keys" on page 352.
4. Add the private keys into memory on the VFM computer. For more information, refer to "Adding private keys to pageant" on page 353.

To configure SSH when the client and server are installed on separate computers

It is assumed that you installed monitoring agents to run some of your policies.

1. Generate a key-pair on the VFM client machine. For more information, refer to "Generating a key-pair" on page 348.
2. Generate a key-pair on the VFM server machine. For more information, refer to "Generating a key-pair" on page 348.
3. Generate a key-pair on each machine you installed a VFM monitoring agent. For more information, refer to "Generating a key-pair" on page 348.
4. Copy the public key generated on the VFM client to the UNIX machines known to VFM. For more information, refer to "Copying the public key to a UNIX computer" on page 349.
5. Copy the public key generated on the VFM server to the UNIX machines known to VFM. For more information, refer to "Copying the public key to a UNIX computer" on page 349.
6. Copy the public key generated on each VFM monitoring agent to the UNIX machines known to VFM. For more information, refer to "Copying the public key to a UNIX computer" on page 349.
7. Cache the public keys on the VFM client onto the UNIX machines known to VFM. For more information, refer to "Caching the public keys" on page 352.

8. Cache the public keys on the VFM server onto the UNIX machines known to VFM. For more information, refer to "Caching the public keys" on page 352.
9. Cache the public keys on each VFM monitoring agent onto the UNIX machines known to VFM. For more information, refer to "Caching the public keys" on page 352.
10. Add the private keys into memory on the VFM client machine. For more information, refer to "Adding private keys to pageant" on page 353.
11. Add the private keys into memory on the VFM server machine. For more information, refer to "Adding private keys to pageant" on page 353.
12. Add the private keys into memory on each VFM monitoring agent. For more information, refer to "Adding private keys to pageant" on page 353.

VFM communicates with filers

The instructions you follow will depend on whether the VFM client and server are installed on the same or separate computers and whether the filer is running CIFS or NFS.

NOTE
You can configure VFM to communicate with only *filers* using SSH. VFM does not currently support SSH on *vfilers*.

To configure SSH when the client and server are installed on the same computer and filers are running CIFS

The VFM server runs all your policies.

1. Enable SSH1 or SSH2 on the CIFS filer. For more information, refer to "Enabling SSH1 or SSH2 on CIFS filers" on page 347.
2. Generate a key-pair on the machine where the VFM client and server are installed. For more information, refer to "Generating a key-pair" on page 348.
3. Copy the public key to each filer that is running CIFS. For more information, refer to "Copying the public key to a filer using CIFS" on page 351.
4. Cache the public keys onto the CIFS filers known to VFM. For more information, refer to "Caching the public keys" on page 352.
5. Add the private keys into memory on the VFM machine. For more information, refer to "Adding private keys to pageant" on page 353.

To configure SSH when the client and server are installed on separate computers and filers are running NFS

It is assumed that you installed monitoring agents to run some of your policies.

1. Generate a key-pair on the VFM client machine. For more information, refer to "Generating a key-pair" on page 348.
2. Generate a key-pair on the VFM server machine. For more information, refer to "Generating a key-pair" on page 348.
3. Generate a key-pair on each machine you installed a VFM monitoring agent. For more information, refer to "Generating a key-pair" on page 348.
4. Copy the public key generated on the VFM client to the NFS filers known to VFM. For more information, refer to "Copying the public key to a filer using NFS" on page 350.

5. Copy the public key generated on the VFM server to the NFS filers known to VFM. For more information, refer to "Copying the public key to a filer using NFS" on page 350.
6. Copy the public key generated on each VFM monitoring agent to the NFS filers known to VFM. For more information, refer to "Copying the public key to a filer using NFS" on page 350.
7. Cache the public keys on the VFM client onto the NFS filers known to VFM. For more information, refer to "Caching the public keys" on page 352.
8. Cache the public keys on the VFM server onto the NFS filers know to VFM. For more information, refer to "Caching the public keys" on page 352.
9. Cache the public keys on each VFM monitoring agent onto the NFS filers know to VFM. For more information, refer to "Caching the public keys" on page 352.
10. Add the private keys into memory on the VFM client machine. For more information, refer to "Adding private keys to pageant" on page 353.
11. Add the private keys into memory on the VFM server machine. For more information, refer to "Adding private keys to pageant" on page 353.
12. Add the private keys into memory on each VFM monitoring agents. For more information, refer to "Adding private keys to pageant" on page 353.

Enabling SSH1 or SSH2 on CIFS filers

Before you generate a key-pair on the VFM server and client computer, enable SSH1 or SSH2 on your CIFS filers.

To enable SSH1 or SSH2 on a CIFS filer

1. Use FilerView and click **Secure Admin**.
2. Under **SSH > Enable/Disable**, select **Enable SSH v1** or **Enable SSH v2**. You can enable them both together at the same time if appropriate for your environment.
3. You might have to configure SSH by Generating keys if used for the first time. You can use the default key sizes. Once you have generated keys, the following public and private key files are located in the *ETC$\sshd* folder:

   ```
   ssh_host_dsa_key (ssh v2 dsa private key)
   ssh_host_dsa_key.pub (ssh v2 dsa public key)
   ssh_host_key (ssh v1 private key)
   ssh_host_key.pub (ssh v1 public key)
   ssh_host_rsa_key (ssh v2 rsa private key)
   ssh_host_rsa_key.pub (ssh v2 rsa public key)
   ```

4. You can also run command line checks as follows:

   a. If you run `filer> options ssh`, the filer should return the following for SSH v2:

   ```
   ssh.access *
   ssh.enable on
   ssh.passwd_auth.enable on
   ssh.pubkey_auth.enable on
   ssh1.enable off
   ssh2.enable on
   ```

b. If you do not run `filer> options ssh`, run `secureadmin setup -f ssh`. Reply `y` to proceed and take the default key sizes, and then reply `y` to proceed again. Wait for the syslog message to the console stating that "... SSH setup is done. ..."

Generating a key-pair

When the VFM client and server are on different machines, you must generate key-pairs on both machines. If monitoring agents are running policies that transfer data, you must also generate key-pairs on the machines where the agents are installed.

To generate a key-pair

1. On the machine where the VFM client or server is running, start a command prompt and type the path of the *puttygen.exe* file.
2. Press Enter.

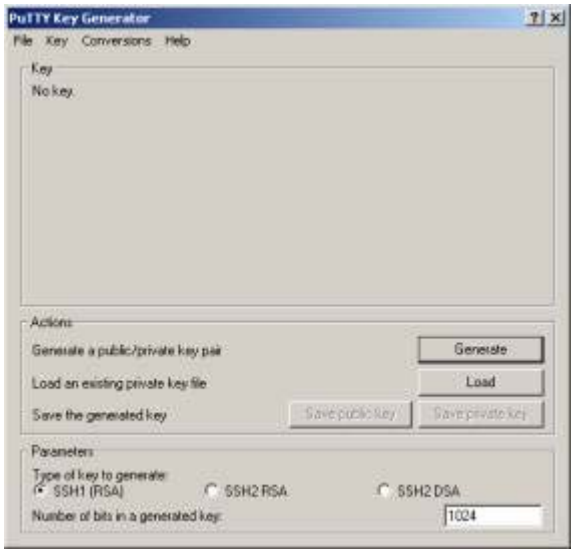

3. Under **Type of key to generate**, select the appropriate key type. For example, if you enabled SSH2 RSA on your CIFS filer, select **SSH2 RSA**.
4. Click **Generate**. Make sure you continually move the cursor around the PuTTY Key Generator dialog box, because puttgen does not run unless the cursor is moving. An SSH1 RSA or SSH2 RSA key with 1024 bits is generated.

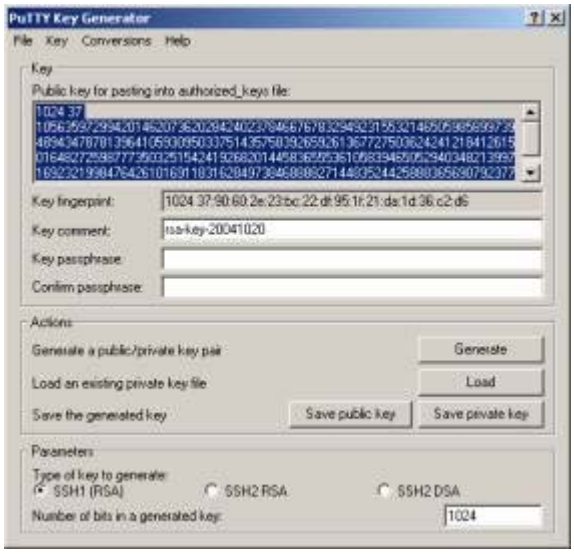

5. Click **Save private key** and specify a directory and file name. It is recommended that you keep all keys in one directory, such as the *C:\putty* installation directory. Provide a meaningful name for the key, such as *private_sshV2_1024_key.ppk*. Remember where this file is saved and the name of the file, because you will have to provide this information to VFM. For more information, refer to "Configuring default remote shell credentials" on page 355.
6. When the message **Are you sure you want to save this key without a passphrase to protect it?** appears, click **Yes**.
7. Select the public key that appears in the **Public key for pasting into authorized_keys file** field, right-click, and click **Copy**.
8. Click **Save Public Key** and specify a directory and file name. It is recommended that you keep all keys in one directory, such as the *C:\putty* installation directory. Provide a meaningful name for the key, such as *public_sshV2_1024_key*. There is no extension on the public key. Open this file with Notepad without word wrap on, and paste the public key from the clipboard into the top line of this file. The pasted public key must not have any line breaks in it. Remember where this file is saved and the name of the file.

Copying the public key to a UNIX computer

After you generate the key-pairs, you copy the public key to a UNIX computer.

To copy the public key to a UNIX computer

1. Login to the UNIX computer from your Windows computer.
2. Create a `.ssh` directory in the home directory of root. Execute the following commands:
   a. To create the `.ssh` directory, type the following command: `mkdir .ssh`
   b. To make the directory your current directory, type the following command:

      ```
      cd .ssh
      ```
3. Create a file named *authorized_keys*.

NOTE
If you generated a SSH2 (RSA) or SSH2 (DSA) key, the key can be stored in the *authorized_keys* or *authorized_keys2* file. Check with your administrator to see which file is being used in your environment. Wherever the filename *authorized_keys* appears in the instructions, substitute the filename *authorized_keys2*.

4. Open the *authorized_keys* file and paste the public key at the end of the file. Make sure that the public key appears in one line. The key should not be separated by line breaks.
5. Save and close the file.
6. Make sure that the following directories have the appropriate permissions:

TABLE 64 Directory permissions

Directory	Permissions
home	700 or 755
.ssh	700
authorized_keys	700 or 755

7. Check the sshd configuration of the UNIX machine to verify that public-key authentication is enabled. Look for the line:

   ```
   PubkeyAuthentication yes
   ```

   The line can begin with a pound symbol (#).

   NOTE
   On Red Hat Enterprise 3, the configuration file is in */etc/ssh/sshd_config*. On Solaris 8, the file is in */usr/local/etc/sshd_config*.

8. *If you changed the configuration,* restart the SSH daemon.

Copying the public key to a filer using NFS

After you generate the key-pairs, you copy the public key to a filer using NFS.

To copy the public key to a filer using NFS

1. Log in to the UNIX machine from your Windows machine.
2. Open a command prompt and type the following command:

   ```
   mount –t nfs filer name:/ /mnt
   ```

   where *filer name* is the name of the filer you want to copy the public key.

3. Make sshd your current directory. Type the following command:

   ```
   cd /mnt/etc/sshd
   ```

4. Perform one of the following actions:
   - Using a root account

     If the root directory does not exist, create it and then make it your current directory. Type the following command:

     ```
     mkdir root
     ```

If the root directory exists, make it your current directory. Type the following command:

```
cd root
```

- Using a non-root account

  If the directory does not exist, create it and then make it your current directory. Type the following command:

  ```
  mkdir non-root account name
  ```

  where *non-root account name* is the name of the directory you want to create. The name of the directory must be the same as the name of the account you are using to create the directory.

  If the directory exists, make it your current directory. Type the following command:

  ```
  cd non-root account name
  ```

  where *non-root account name* is the name of the directory. The name of the directory is the same as the name of the account you are using to configure SSH.

5. Create a directory named *.ssh*. Type the following command:

   ```
   mkdir .ssh
   ```

6. Create a file named *authorized_keys*.
7. Open the *authorized_keys* file and paste the public key at the end of the file. Make sure that the public key appears in one line. The key should not be separated by line breaks.
8. Save and close the file.

Copying the public key to a filer using CIFS

After you generate the key-pairs, you copy the public key to each filer using CIFS.

To copy the public key to a filer using CIFS

1. In Windows Explorer, type the fully qualified path of the filer.
2. Open the *filer name**etc$**sshd* directory, where *filer name* is the name of the filer to which you want to copy the public key.
3. Perform one of the following actions:
   - Using a root account

     If the root directory does not exist, create it and then make it your current directory. Type the command:

     ```
     mkdir root
     ```

     If the root directory exists, make it your current directory. Type the command:

     ```
     cd root
     ```

   - Using a non-root account

     If the directory does not exist, create it and then make it your current directory. Type the command:

     ```
     mkdir non-root account name
     ```

where *non-root account name* is the name of the directory you want to create. The name of the directory must be the same as the name of the account you are using to create the directory.

- If the directory exists, make it your current directory. Type the command:

  `cd non-root account name`

  where *non-root account name* is the name of the directory. The name of the directory is the same as the name of the account you are using to configure SSH.

4. Create a directory named *.ssh*. Type the command:

   `mkdir .ssh`

5. Create a file named *authorized_keys*.
6. Open the *authorized_keys* file and paste the public key at the end of the file. Make sure that the public key appears in one line. The key should not be separated by line breaks.

   NOTE
   This file can contain multiple keys of mixed SSH1 and SSH2 types from the same server or different servers.

7. Save and close the file.

Caching the public keys

After you copy the public key to the UNIX machine or filer, you need to run the *plink.exe* program on the VFM client and server machines. If monitoring agents are assigned to run policies that transfer data, you must also run the *plink.exe* program on the machines where the monitoring agents are installed. The *plink.exe* program caches the public keys in the registry. You must run this program before you can use VFM to transfer data.

NOTE
The `plink.exe` program must be executed using the account that runs VFM.

To cache the keys

1. On the machine where you generated the key-pairs, start a command prompt and type the path of the `plink.exe` file. Type the following command:

   `plink host -l user`

   where *host* is the fully qualified name of the UNIX machine or filer, such as *secure.npb.com*, and *user* is the account name used to log in to the UNIX machine or filer.

   NOTE
   You are prompted to add the key to the cache of the account that runs VFM.

2. Type the command *plink* again. For *host* type the fully-qualified name of the UNIX machine or filer, such as *secure.npb.com*.
3. Verify that you are not prompted again to add the key to the cache. If you are prompted, complete step 1 and step 2 again.

Once you have cached the public keys, you should export this branch of the registry to an easy-to-remember folder for importing each time you log on as a different user who needs access to VFM and SSH. For example, export to a file named *SshHostKeys.reg*.

Adding private keys to pageant

After you cache the public keys, you must run the *pageant.exe* program on the VFM client and server machines. This program copies the private keys into memory. The private keys are decoded so that VFM can communicate with a filer or UNIX machine without a password.

NOTE
If monitoring agents are assigned to run policies that transfer data, you must also run the `pageant.exe` program on the machines where the monitoring agents are installed.

To add private keys to pageant

1. On the machine where you generated the key-pairs, start a command prompt and type the path of the *pageant.exe* file. Type the following command:

   `pageant path`

   where *path* is the path to the private key and the name of the file containing the private key.

   NOTE
   The *pageant.exe* program must be executed using the account that runs VFM.

2. Verify that you are not prompted for a password. At the command prompt, type the following command:

   `plink host -l user ls`

   where *host* is the fully qualified name of a UNIX machine or filer (for example, *secure.npb.com*) and *user* is the account name used to login to the UNIX machine or filer.

Installing and configuring OpenSSH on Solaris

Before you install OpenSSH on a UNIX machine, you must install a patch for */dev/random/*. Contact Sun or go to the web site http://sunsolve.sun.com to get the latest version of patch 112438-03. After you install the patch, continue with the instructions listed below to get, install, and configure the OpenSSH package.

To install patch on Solaris

1. Download the patch and install it as indicated in the instructions.
2. Verify that */kernel/drv/random*, */dev/urandom*, and */dev/random* exist on the UNIX machine. If these files not exist, execute the following commands:

   ```
   ln –s /kernel/drv/random /dev/random
   ln –s /kernel/drv/random /dev/urandom
   ```

3. Reboot the machine.

To get the OpenSSH packages for SPARC

1. Download the following packages from sunfreeware.com:

   ```
   openssh-3.9pl –sol8-sparc-local.gz
   openssh-0.9.7d-sol8-sparc-local.gz
   zlib-1.2.1-sol8-sparc-local.gz
   libgcc-3.3-sol8-sparc-local.gz
   ```

2. Copy the packages to the UNIX machine.

To install the packages

1. Change to the directory where you copied the packages.
2. Type the following commands:

```
gunzip openssh-3.9pl-sol8-sparc-local.gz
gunzip openssl-0.9.7d-sol8-sparc-local.gz
gunzip zlib-1.2.1-sol8-sparc-local.gz
gunzip libgcc-3.3-sol8-sparc-local.gz
pkgadd –d openssh-3.8.1p1-sol8-sparc-local
pkgadd –d openssl-0.0.7d-sol8-sparc-local
pkgadd –d zlib-1.2.1-sol8-sparc-local
pkgadd –d libgcc-3.3-sol8-sparc-local
```

NOTE
`ssh` is installed in */usr/local/bin* and `sshd` is installed in */usr/local/sbin*. Make sure that */usr/local/bin* and */usr/local/sbin* are in your *PATH* environment variables.

To configure the sshd user and the /var/empty directory

Execute the following commands as root:

```
mkdir /var/empty
chown root:sys /var/empty
chmod 755 /var/empty
groupadd sshd
useradd –g sshd –c 'sshd privsep' –d /var/empty –s /bin/false sshd
```

NOTE
The directory */var/empty* should not contain any files.

To Install ssh and sshd

Each machine that you want to communicate with SSH must have an sshd daemon running.

1. Run the following three lines to create the key information for the server machine. Make sure that you have `/usr/local/bin` and `/usr/local/sbin` in your `PATH`. It may take a few minutes to complete the following commands. As root, type:

```
ssh-keygen –t rsal –f  /usr/local/etc/ssh_host_key –N ""
ssh-keygen –t dsa –f /usr/local/etc/ssh_host_dsa_key –N ""
ssh-keygen –t rsa –f /usr/local/etc/ssh_host_rsa_key –N ""
```

2. Create a script named *sshd* to start the sshd daemon. Place the script in the */etc/init.d* directory.

```
#!/sbin/sh
#
#
[ ! –d /usr/local/sbin ] && exit
case "$1" in
'start')
        /usr/local/sbin/sshd
         ;;
'stop')
       /usr/bin/pkill –x –u 0 'sshd'
```

```
        ;;
*)
       echo "usage: $0 { start | stop }"
       exit 1
       ;;
esac
exit 0
```

3. Run the following commands as root:

```
chown root /etc/init.d/sshd
chgrp sys /etc/init.d/sshd
chmod 555 /etc/init.d/sshd
ln –s /etc/init.d/sshd /etc/rc2.d/S98sshd
```

To start the process

The final step is to start the sshd daemon. Execute the following command as root:

```
/etc/rc2.d/S98sshd start
```

If you want to stop the sshd daemon, type the following command:

```
/etc/rc2.d/S98sshd stop
```

Installing and configuring OpenSSH on Red Hat

Use the following instructions to install and configure OpenSSH on Red Hat Enterprise 3.

To install and configure OpenSSH on Red Hat

1. Get the latest OpenSSH server packages.
2. Copy the appropriate packages for Red Hat Enterprise 3 to the UNIX machine.
3. Change to the directory where you copied the packages.
4. Install the OpenSSH server packages by typing the command `Rpm –ivh` *`package filename`*, where *package filename* is the name of the package you downloaded.
5. Configure Red Hat to start the SSH service automatically when the system reboots by typing the following command:

   ```
   Chkconfig –level 35 sshd on
   ```

6. Start the SSH service by typing the following command:

   ```
   /etc/init.d/ssh start
   ```

Configuring default remote shell credentials

Before you can configure default shell credentials, you must configure VFM to use SSH. For information about how to configure SSH, refer to "Configuring VFM to use SSH" on page 345.

After you complete the steps to configure VFM to use SSH, you must configure default remote shell credentials for all new filers and UNIX machines added to the Physical View. These credentials enable VFM to use *plink.exe* in `ssh` mode for remote shell operations and use *pscp.exe* to deploy replication agents.

To configure default remote shell credentials

1. On the **Tools** menu, click **Options**.
2. Click **Remote Shell**.
3. Select a default shell type from the list.
4. *If you select RSH as the default shell,* type the location of the rsh and rcp executables and the user name and password that VFM will use to communicate with filers and UNIX machines. When you leave the user name and password blank, VFM uses root as the user name with no password.
5. *If you select SSH as the default shell,* enter the following information:
   a. In the **SSH Executable Location** field, type the location of *plink.exe*.
   b. In the **SCP Executable Location** field, type the location of *pscp.exe*.
   c. In the **Private Key Location** field, type the location of the private key. Make sure this path does not have any spaces.
   d. In the **User Name** field, type the user name or account name that VFM will use to communicate with filers and UNIX machines. This is usually the root user. When VFM uses SSH to communicate with UNIX machines and filers, SSH is configured without a password or passphrase.

Once you have configured default remote shell credentials, you can:

- Change the remote SSH shell credentials for a filer in the Physical View. For more information, refer to "Change the shell credentials of a filer" on page 144.
- Change remote SSH shell credentials for a UNIX machine in the Physical View. For more information, refer to "Change the shell credentials of a UNIX machine" on page 160.

Appendix

Inclusion and Exclusion Tables

B

This section explains the rules for using file and directory inclusion and exclusion lists in VFM policies that move data. It provides examples of inclusion and exclusion list formats.

In this appendix

- Inclusion lists 357
- Exclusion lists 359

Inclusion lists

File and directory inclusion lists can be configured in policies that move data, such as migration and archival migration policies. Inclusion lists are configured in the **Replication Filters** dialog box of the properties of a policy.

Inclusion rules

When creating inclusion lists, remember the following points:

- Exclusion terms are processed before inclusion terms. Files or directories matching an exclusion term are not processed further and cannot be included by a subsequent inclusion term.
- Directory inclusions are search expressions comprised of folder names or folder paths that can be used to include any folders matching the specified inclusion criteria.
- File inclusions are search expressions comprised of file names or extensions that can be used to include files matching the specified inclusion criteria.
- DOS wildcards can be used alone or in concert with a character string to include file or folders of similar name or type.
  - To match a string of characters of any length including an empty string, type an asterisk (`*`). For example, `s*d` finds `sad`, `suddenly`, and `Saturday`.
  - A question mark (?) can be used to match any single character in the search text. For example, `s?t` finds `set` and `sit`, but not `Saturday`.
- To include a subfolder and its contents, specify a backward slash (\) for Windows or a forward slash (/) for UNIX after the directory name.
- To include just a subfolder and not its contents, omit the trailing backward slash (\) or forward slash (/).
- When you enable **Delete orphaned files in destination folders** in the Replication Options dialog box of a policy that moves data, the inclusion list affects how VFM handles orphan files on the destination. Orphan files on the destination that match the inclusion term are deleted.

- You must delimit multiple directory or file inclusion parameters with double quotes or you can click **Expand** to specify multiple files, file types, or directories.
- When the **Include subfolders** option in the Replication Options dialog box of a policy that moves data is enabled and you specify file inclusions, VFM traverses subfolders for files to include in the destination. Subfolders are always transferred to the destination even when they are empty.

Examples of valid inclusion formats

The following table provides examples of valid inclusion formats and descriptions for each format.

TABLE 65 Example of valid inclusion formats

Format	Applied to	Description
*	Files or Directories	When this example is used specifically as a directory inclusion parameter, only folders (no files) of the specified parent will be included. When it is used as a file inclusion parameter, all files and all folders will be included.
.	Files	All files contained within the specified parent and their associated child folders will be included.
\ (Windows) / (UNIX)	Directories	When used as a directory inclusion parameter, all folders contained within the specified parent are included.
* (Windows) /* (UNIX)	Files or Directories	When used as a directory inclusion parameter, this example will include all folders. When used as a file inclusion parameter, this example will include all files and related folders.
. (Windows) /*.* (UNIX)	Files	When used as a file inclusion parameter, this string includes all files contained in the parent.
*.tmp	Files	This example can be used to include files that contain the literal character string `.tmp`.
listing	Files or Directories	Include Files or Directories named `listing`.
listing.txt	Files	Match files or directories named `listing.txt`.
listing???.txt	Files	This example can be used to include all text files with names that begin with the string `listing` followed by any three characters.
listing*.txt	Files	Used to include all text files with names that begin with the literal character string `listing`.
\listing (Windows) /listing (UNIX)	Directories	When prefaced with the \ character, this example will traverse the directory tree and include folders named `listing`.
listing\ or \listing\ (Windows) listing/ or /listing/ (UNIX)	Directories	This example will traverse the directory tree and include folders that are sub directories of the `listing` directory.

TABLE 65 Example of valid inclusion formats

Format	Applied to	Description
*\listing (Windows) */listing (UNIX)	Directories	Include subfolders named listing.
*\listing\file.txt (Windows) */listing/file.txt (UNIX)	Files	This example can be used to include files named "file.txt" located in a subfolder named "listing."

Exclusion lists

File and directory exclusion lists can be configured in policies that move data, such as migration and archival migration policies. Exclusion lists are configured in the **Replication Filters** dialog box of the properties of a policy.

Exclusion rules

When creating exclusion lists, remember the following points:

- Exclusion terms are processed before inclusion terms. Files or directories matching an exclusion term are not processed further and cannot be included by a subsequent inclusion term.
- Directory exclusions are search expressions comprised of folder names and/or folder paths that can be used to exclude any folders matching the specified exclusion criteria.
- File exclusions are search expressions comprised of file names or extensions that can be used to exclude files matching the specified exclusion criteria.
- DOS wildcards can be used alone or in concert with a character string to exclude file or folders of similar name or type.
  - To match a string of characters of any length including an empty string, type an asterisk (`*`). For example, `s*d` finds `sad`, `suddenly`, and `Saturday`.
  - A question mark (?) can be used to match any single character in the search text. For example, `s?t` finds `set` and `sit`, but not `Saturday`.
- To exclude a subfolder and its contents, the trailing backward slash (\) or forward slash (/) is optional. The subfolder and its contents are excluded whether or not the trailing backward slash or forward slash is present.
- When you enable **Delete orphaned files in destination folders** in the Replication Options dialog box of a policy that moves data, the exclusion list affects how VFM handles orphan files on the destination. Orphan files on the destination that match the exclusion term are not deleted, while orphan files not matching the exclusion term are deleted.
- You must delimit multiple directory or file exclusion parameters with double quotes or you can click **Expand** to specify multiple files, file types, or directories.
- When the **Include subfolders** option in the Replication Options dialog box of a policy that moves data is enabled and you specify file exclusions, VFM traverses subfolders for files to exclude from the destination. Subfolders are always transferred to the destination even when they are empty.

Examples of valid exclusion formats

The following table provides examples of valid file and folder exclusion formats and descriptions for each format.

TABLE 66 Examples of valid file and folder exclusion formats

Format	Applied to	Description
*	Files or Directories	When this example is used specifically as a directory exclusion parameter, all files and folders will be excluded. When used specifically as a file exclusion parameter, folders are still included, but all files contained in those folders will be excluded.
.	Files	All files contained within the specified parent and their associated child folders are excluded.
\ (Windows) / (UNIX)	Directories	When the slash is used as a directory exclusion parameter, all files and folders contained within the specified parent and child folders are excluded.
* (Windows) /* (UNIX)	Files or Directories	When used as a directory exclusion parameter, this example will exclude all files and all folders. When used as a file exclusion parameter, this example will result in the exclusion of all files.
. (Windows) /*.* (UNIX)	Files	When used as a file exclusion parameter, this string excludes all files contained in the parent.
*.tmp	Files	This example can be used to exclude files that contain the literal character string `.tmp`.
listing	Files or Directories	Exclude files or directories named `listing`.
listing.txt	Files	Exclude files named `listing.txt`.
listing???.txt	Files	This example can be used to exclude files that begin with the string `listing` followed by any three characters and then the literal character string `.txt`.
listing*.txt	Files	Used to exclude all text files with names that begin with the literal character string `listing`.
\listing (Windows) /listing (UNIX)	Directories	When prefaced with the \ character, this example will traverse the directory tree and exclude folders named `listing`.
listing\ or \listing\ (Windows) listing/ or /listing/ (UNIX)	Directories	This example will traverse the directory tree and exclude folders that are subfolders of the `listing` directory.
*\listing (Windows) */listing (UNIX)	Directories	Exclude subfolders named `listing`.
*\listing\file.txt (Windows) */listing/file.txt (UNIX)	Files	This example can be used to exclude files named `file.txt` located in folders named `listing`.

Appendix

Example Migration Tasks Script

C

This section explains how to create task creation batch files to add migration tasks to a migration policy. It provides example task creation and task description files that you can use as a model to create these files in your environment.

In this appendix

- Specifying migration tasks from a file 361
- Example task creation batch file 363
- Task description file. 364

Specifying migration tasks from a file

You have the option of specifying the source and destination in a file instead of having a migration policy scan for candidates to migrate. Migration policies provide the option of running a batch file to add migration tasks to a policy. You can specify the batch file in the **Run task creation batch** field on the **General** tab of a data migration policy properties dialog box.

The following list describes how task scripts are processed in VFM.

- Each time the migration policy runs, it invokes the task creation batch file.
- The task creation batch file creates a task description file.
- After the batch file exits, the VFM server performs the following actions:
  - Reads the task description file created by the batch file.
  - Adds the tasks described in the task description file to the policy.
  - Deletes the task description file. Deleting the task description file enables you to schedule repeated transfers of data for the locations described in the tasks added by the batch file.

Task creation batch file arguments

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

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

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

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

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

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

```
ComplexTaskCreation.bat [arg1] [arg2] [arg3]
```

VFM invokes your batch file as:

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

Task File Description Syntax

This section describes the elements and attributes of a task creation batch file. For an example task creation batch file, refer to "Example task creation batch file" on page 363.

DFSMIGRATION element

The `DFSMIGRATION` element is required. The word `Version:` is required and must be set to `1`. The element does not have attributes.

Task element

The `Task` element is optional. If you do not include this element with at least one of the following attributes in the batch file, nothing is processed.

CurrentTarget

- The attribute is required.
- Must be a valid UNC path with at least the computer and share portions of the path specified, for example, *server**share*.
- Cannot be a DFS path.
- Cannot be the same as NewTarget

NewTarget

- The attribute is required.
- Must be a valid UNC path with at least the computer and share portions of the path specified, for example, *server**share*.
- Cannot be a DFS path.
- Cannot be the same as `CurrentTarget`

AutoDelete

- The attribute is optional.
- When set to `1`, the task is automatically deleted from the policy if it is completed successfully.
- If not present or set to anything other than `1`, the task is not automatically deleted

CutoverDate

- The attribute is optional.
- If set to “”, the task setting for **Advance to Final Phase** is **Never**.
- If set to `NextRun`, the task setting for **Advance to Final Phase** is **Next Policy Run**.
- If set to a date in the format “mm/dd/yyy hh:mm:ss”, the task setting for **Advance to Final Phase** is **Next Policy Run After** the date specified.
- If the attribute is not present, the task setting for **Advance to Final Phase** is **Use the Policy Default**.

Link element

The `Link` element is optional. Only links listed in the batch file are updated when the policy runs. The **Update DFS** option in the **Actions** tab of the policy properties must be selected. If you include this element, you must include the `LinkName` attribute.

LinkName

- The attribute is required.
- The full path of a link to be retargeted.
- Must be a valid UNC path with at least the computer and share portions of the path specified, for example, *server**share*. Specifying only *server**share* is invalid.

CurrentTarget

- The attribute is optional.
- Indicates the current replica of the link that is to be retargeted.
- If not present, the default value is the value of the `CurrentTarget` attribute of the `Task` element.
- If present, the attribute must indicate the same folder or a subfolder of the path specified in the `CurrentTarget` attribute of the `Task` element.

Example task creation batch file

The following example shows a task creation batch file that creates a task description file. For more information about task description files, refer to “Task description file” on page 364.

```
@echo off
rem ****************************************************************************
rem FILE: CreateTasks.bat
rem
rem PURPOSE:
rem    Illustrate creating a task-description file
rem
rem
rem INPUT:
rem    %1 Name of the migration policy calling this BAT file.  This is supplied
from VFM.
rem    %2 Name of the output xml file.  This is supplied from VFM.
rem
rem COPYRIGHT:
rem    Copyright (c) 2007 Brocade Communications Systems, Inc.  All rights
reserved.
```

```
rem
rem
*****************************************************************************/
rem
rem Next two lines set up the start of the XML file to be handed to VFM
rem
echo ^<?xml version="1.0" encoding="UTF-8" ?^> > ExampleTasks.xml
echo ^<DFSMIGRATION Version="1"^> >> ExampleTasks.xml
rem Add first task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src1^"
NewTarget=^"\\SERVERB\Test\Dst1^" ^/^> >> ExampleTasks.xml
rem Add second task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src2^"
NewTarget=^"\\SERVERB\Test\Dst2^" AutoDelete=^"1^" ^/^> >> ExampleTasks.xml
rem Add third task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src3^"
NewTarget=^"\\SERVERB\Test\Dst3^" AutoDelete=^"0^" ^/^> >> ExampleTasks.xml
rem Add fourth task
echo ^<Task CurrentTarget=^"\\SERVERA\Test\Src4^"
NewTarget=^"\\SERVERB\Test\Dst4^" AutoDelete=^"1^"^>   >> ExampleTasks.xml
echo ^<Link LinkName=^"\\ROOTSERVER\Root\Link^" ^/^> >> ExampleTasks.xml
echo ^<Link LinkName=^"\\ROOTSERVER\Root\OtherLink^" >> ExampleTasks.xml
echo CurrentTarget=^"\\SERVERA\Test\Src4\Subdir^" ^/^> >> ExampleTasks.xml
echo ^<^/Task^> >> ExampleTasks.xml
rem
rem Terminate the XML file and copy it to where VFM told us to put it
rem
echo ^</DFSMIGRATION^> >> ExampleTasks.xml
copy ExampleTasks.xml %2
rem
rem Delete our temporary file
rem
del ExampleTasks.xml

exit 0
```

Task description file

This section provides an example of a task description file with four tasks. The task description file is created by the task creation batch file. For more information about task creation batch files, refer to "Example task creation batch file" on page 363.

The following list describes the four migration tasks included in the example task description file listed below.

- The first task migrates data from *\\SERVERA\Test\Src1 to \\SERVERB\Test\Dst1* and leaves the task in the policy after it is completed successfully. This task uses the Advance to Final Phase settings from the policy.
- The second task migrates data from*\\SERVERA\Test\Src2 to \\SERVERB\Test\Dst2*; the task is automatically removed from the migration policy after the task is completed successfully. This task advances to the final phase after 10:00 PM on 11/03/2007.
- The third task migrates data from *\\SERVERA\Test\Src3 to \\SERVERB\Test\Dst3* and explicitly specifies that the task is not to be removed from the policy after it is completed successfully. This task never advances to the final phase.

- The fourth task migrates data from *\\SERVERA\Test\Src4 to \\SERVERB\Test\Dst4*, retargets *\\ROOTSERVER\Root\Link* to reference the data at its new location and retargets the link replica of *\\ROOTSERVER\Root\OtherLink* currently targeted at *\\SERVERA\Test\Src4\Subdir* to reference beneath the new location. This task advances to the final phase the first time it runs.

The following example illustrates a task description file containing the four migration tasks describes above.

```
<?xml version="1.0" encoding="UTF-8" ?>
  <DFSMIGRATION Version="1">
  <Task CurrentTarget="\\SERVERA\Test\Src1" NewTarget="\\SERVERB\Test\Dst1" />
  <Task CurrentTarget="\\SERVERA\Test\Src2" NewTarget="\\SERVERB\Test\Dst2"
AutoDelete="1" CutoverDate="11/03/2006 22:00:00"/>
  <Task CurrentTarget="\\SERVERA\Test\Src3" NewTarget="\\SERVERB\Test\Dst3"
AutoDelete="0" Cutover Date=""/>
  <Task CurrentTarget="\\SERVERA\Test\Src4" NewTarget="\\SERVERB\Test\Dst4"
AutoDelete="1" CutoverDate="NextRun">
  <Link LinkName="\\ROOTSERVER\Root\Link" />
  <Link LinkName="\\ROOTSERVER\Root\OtherLink"
CurrentTarget="\\SERVERA\Test\Src4\Subdir" />
  </Task>
  </DFSMIGRATION>
```

C Task description file

Appendix

Troubleshooting

D

This section provides a list of problems that you might encounter and explains how to fix the problems. It also includes a list of events logged by VFM and lists of status indicators used throughout the application.

In this appendix

- Finding solutions 367
- Migration, replication, and status indicators. 393
- Exclusion and Inclusion Examples 399

Finding solutions

Start by checking the health of your VFM configuration then browse the list of problems in the table below.

- Are the VFM Server and VFM Replication Agent services running on the machines where the VFM server or monitoring agent components are installed? Click Agent Management on the Tools menu to view the list of services installed on the system and ensure that the VFM Server and VFM Replication Agent services are enabled and started.
- Does sufficient free space exist on the drive containing the **Documents and Settings** directories on the system where VFM is installed?
- Does the account used by the VFM Server service have appropriate privileges to manage the selected DFS roots and access the source and destination of any data transfers?
- Does the account used to run the VFM console have administrator privileges on the system where it is installed?

General problems

TABLE 67 General problems and solutions

If...	Then, do this...
VFM seems slow to respond, or takes up a large amount of CPU.	Check how many links are present in the DFS configuration. Larger configurations require VFM to gather additional information. Verify that targets of links are reachable from the VFM console machine. Verify that network browsing works correctly for regular Windows tools, such as Windows Explorer. • Check what level of status monitoring is enabled (root, link, storage) and what frequency monitoring is taking place. • Enabling the **show drives in physical views** option causes VFM to gather additional information and can slow operation, especially when target machines are not reachable from the VFM console machine.
You experience access denied errors while setting VFM system options.	The account running the VFM client must have permissions to write to the directory tree `All Users\Application Data\NetApp\VFM\GuiConsole`. Write permission for these directories is typically granted to members of the local administrators group and the local power users group.
VFM reports an error about expired evaluation copy.	Your evaluation copy expired. Contact your supplier to request an extension or to purchase a copy of VFM. This error message can result when the currently logged in user attempting to run the VFM console is not a member of the administrators group on the local system.
VFM refuses to install.	VFM currently supports installation only on Windows 2000 SP 2 or later. Check whether the problem is installing MSDE. Microsoft KB article 315083 details various steps to take when troubleshooting a failed MSDE installation attempt.
You cannot drag-and-drop items out of Windows Explorer into VFM.	Items dragged to VFM must be network paths. You can map shares from remote machines as drives on your machine or address them as UNC paths. For local drives on your machine, you must address them as UNC paths `(\\server\dir\)`.
You can only rearrange items displayed in **Bold** in the Add Multiple Links dialog box.	This is normal behavior. Use the VFM Logical View to make such changes when needed.
You encounter errors creating logical folders in your namespace	• Check that the permissions on the DFS root share allow the account running the VFM console the ability to make changes to the share. • Windows Server 2003 creates shares with default permissions of Everyone>Read. Windows 2000 creates shares with default permissions of Everyone>Full. The VFM server account must have permissions to write to the DFS root share to create logical directories in the namespace.

TABLE 67 General problems and solutions

If...	Then, do this...
You changed your DFS configuration but client machines still access the old locations.	Clients cache referrals for a period of time. You can flush the client cache on Windows 2000 and later by performing one of the following actions: • Use the DFS tab of Windows Explorer Property dialog box to "clear history". • Run `dfsutil` with appropriate options on the client machine. • You must reboot clients running older operating systems, such as Windows NT 4, Windows 98, or Windows 95, to clear the cache.
VFM shows the wrong date and time for a NetApp filer message appearing in the Event View.	Check whether you correctly set up remote shell credentials on the filer. Normally, VFM uses RSH to communicate with NetApp filers and other non-Windows systems. When the remote shell credentials on the filer are not properly configured, VFM shows the filer message with the same date and time as the VFM server.
You see the error message **Network resource not found** when you try to upgrade the replication agent installed on the same machine as the VFM server.	• Check whether the VFM 5.5 server was uninstalled and the VFM 5.5 R2 server was installed. • Before you upgrade the 5.5 replication agent, open a command prompt and change to the directory where the agent is installed and type the following command: `ReplicationAgent.exe -uninstall`
The VFM desktop does not display correctly. The left pane is blank.	• Check whether you are using the Virtual Network Computing (VNC) desktop protocol to connect to remote machines. • It is highly recommended that you use Remote Desktop to communicate with remote machines. Remote Desktop Protocol (RDP) is a multi-channel protocol that allows a user to connect to a computer running Microsoft Terminal Services. The server listens by default on TCP port 3389.
You installed the VFM client on the same machine that you previously installed a monitoring agent. When you select an item from a view on the left pane, the right pane is blank.	If you installed the VFM client after you installed the monitoring agent, perform the following actions to correct the problem: 1 Close the VFM client. 2 Start a command-line prompt window. 3 Change directory to the VFM installation directory. 4 Run the command: `REGSVR32 GuiControls.dll` You will see a message confirming that the registration was successful

DFS root problems

TABLE 68 DFS root problems and solutions

If...	Then, do this...
You do not see your roots in the VFM browse dialog box. You might see one of the following error messages: • Unable to display the DFS Root <hostname>. The network resource was not found. • **Unavailable** status in the UI and the root shows an error status.	• Check whether network browsing is working correctly in Windows Explorer and correct any browsing problems. • Type the root name directly rather than browsing. • Check whether the DFS service is started on the machines hosting the roots. Microsoft Windows Server 2003 SP 1 ships with the DFS service stopped as the default.
Changes made to domain roots (for example, creation or deletion) take a long time to appear in VFM.	Information about domain roots is hosted in Active Directory. There is a time lag while Active Directory replicates the changes. This time lag is affected by the Active Directory configuration.
You cannot create a DFS root using VFM or you cannot create links on a root. You might see one of the following error messages: • Error: the network resource was not found. • An error occurred while trying to take the following action. Action: Create DFS root \\hostname\rootname on server 'hostname' using share <name of share>. • The DFS Root, link or replica could not be added.	Check network connectivity by pinging the server the root is created on from the VFM machine. • Check that the account running VFM has appropriate privileges for the domain and servers where the domain-based root is being created or server where the stand-alone root is being created. • Check whether forward and reverse DNS lookups are accurate for the machine on which you want to create the root. • Check whether you are trying to use a share to create the root that contains a root. • Check whether you are trying to use a folder to create the root that is a sub-directory of a share that contains a root. • Check whether the DFS service is started on the machine where you want to create the root. • Check whether the DFS service is started on the machine hosting the root where you want to create a link.
You cannot add links to a DFS root.	• Check that the account running VFM has appropriate permissions for the DFS root being manipulated. • Check whether forward and reverse DNS lookups are accurate for the machine hosting the DFS root. • Verify that all root replica share names of a domain-based root match the name of the DFS root. This is particularly important for roots hosted on Windows Server 2003 systems.
Attribute sets are different between different types of roots.	The attributes that are available and capable of being set vary between Windows NT 4 roots, Windows 2000 stand-alone roots and domain-based roots. This is an artifact of DFS itself on these platforms and not a problem with VFM.

TABLE 68 DFS root problems and solutions

If...	Then, do this...
You have problems creating a domain-based root with replicas. After the domain-based root is created • It appears as a single replica root • The status of the root is displayed as inaccessible	The creation of and operation on domain-based roots involves interactions with Active Directory. The Active Directory replication frequency affects how quickly changes appear in the VFM console. • Name resolution issues can affect the creation of domain-based roots. You might need to refer to the domain-based root by its fully-qualified domain name when accessing it when attempts to access it by its NETBIOS domain name are unsuccessful or result in inaccurate data. • Verify that all root replica share names of a domain-based root match the name of the DFS root. This is particularly important for roots hosted on Windows Server 2003 systems.
Morphed filenames appear in the namespace of domain root replicas	Do not enable FRS at the root level. The Microsoft DFS FAQ http://www.microsoft.com/windowsserver2003/techinfo/overview/dfsfaq.mspx describes this symptom in more detail.
You get the following error when accessing a link contained in a consolidation root: `No network provider accepted the given path`	• Check whether the link is under a logical folder. You must configure links in consolidation roots directly under the root. They cannot be referenced under logical folders. This restriction applies only to consolidation roots. • For more information about consolidation roots, refer to Microsoft Knowledge Base Article 829885.
You see the error message `The request is not supported` when attempting to create a DFS root.	This typically occurs when attempting to create a second DFS root on Windows 2000 or Windows Server 2003 Standard Edition server. Windows 2000 and Windows Server 2003 Standard Edition supports only one DFS root per computer. Windows Server 2003 Enterprise Server and Windows Server 2003 Datacenter support multiple DFS roots per server.

Report problems

TABLE 69 Report problems and solutions

If...	Then, do this...
You experience errors attempting to run VFM reports.	The account running the VFM client must have permissions to write to the directory tree `All Users\Application Data\NetApp\VFM\GuiConsole`. Write permission for these directories is typically granted to members of the local administrators group and the local power users group.
Reports are slow to generate.	• Check the General problems for information about when VFM seems slow to respond. • Check whether you created a new report group when you created the report. Before VFM can create the report, the report group must gather the data. To cut down on report generation, create report groups to gather data before you assign and run reports based on the data gathered. • Create multiple reports instead of one large report. • Schedule report groups so that reports assigned to the report groups are available when needed for review.

Target problems

TABLE 70 Target problems and solutions

If...	Then, do this...
VFM does not enumerate (browse) shares hosted on NetWare®, Win9x, and WinME systems.	These operating systems do not support the system calls required by VFM to gather information about their shares. You can add shares from these systems to your namespace by dragging their UNC paths from Windows Explorer and dropping them into any of the following: VFM Logical View, Add Link dialog box, Add Multiple Links dialog box
When selecting targets for links, you do not see any entries in the Browse dialog box	You might have browsed into a leaf directory and there are no subfolders available at this level to which you can target links.
You see the error message, `There were no targets for disaster recovery link on a disaster recovery policy.`	Links that are members of recovery policies must have at least two replicated targets. Only one target can be online. When the server containing the active target becomes unavailable, VFM fails over the link to the replicated target and makes it the online target.
You see the error message, `A new link was detected for the disaster recovery policy but the link cannot be added for the stated reasons.`	• Check whether the root containing the link is hosted on Windows NT 4, which does not support setting link target state. • Check whether the link is already in another recovery policy. • For more information, refer to "Disaster recovery planning considerations" on page 20.

Windows XP and Windows 2000 incompatibility problem

TABLE 71 Windows XP and Windows 2000 incompatibility problems and solutions

If...	Then, do this...
Links for a stand-alone root installed on a machine running Windows 2000 do not appear when VFM is running on Windows XP.	Upgrade Windows 2000 to Service Pack 2 on the machine containing the stand-alone root.

Problems connecting to the VFM Server service

TABLE 72 VFM server service connection problems and solutions

If...	Then, do this...
The connection was refused by the VFM Server service.	• The VFM Server service might have been stopped through the Windows Service Control Manager. Restart the VFM Server service. • The machine hosting the VFM Server service might not be reachable by the VFM console. Check your network connectivity and the machine hosting the VFM Server service. • Check whether the account configured for the VFM Server service is in the local administrators group on the system where it is installed. • The VFM client reports access denied error attempting to connect to the VFM server. Contact your VFM administrator to obtain access to the VFM server for your account.
The VFM client does not show updates.	Check whether there is a loss of network connection between the VFM client and server. Clear this problem by exiting and reentering the VFM client.

Problems selecting a monitoring agent

TABLE 73 Monitoring agent selection problems and solutions

If...	Then, do this...
A monitoring agent is installed but does not appear as a choice in the **Use the Monitor on** list when configuring a client recovery policy.	• Check that the machines on which the monitoring agent and the VFM Server service are running have network connectivity with one another. • Check that the monitoring agent is configured with the correct host name for the machine where the VFM Server service is running. This value is stored in the registry at *HKLM\Software\Brocade\VFM\Services\Server\Primary.* • Check that no firewalls are blocking TCP port 6001 between the monitoring agent and the VFM Server service. **Note:** The table that follows this one lists the network ports used by DFS and VFM.

TABLE 74 Network ports used by DFS

Service Name	UDP	TCP
NetBIOS Datagram Service	13	
NetBIOS Session Service		139
LDAP Server	389	389
Remote Procedure Call (RPC)		135
Server Message Block (SMB)	445	445

TABLE 75 Network ports used by VFM

Service Name	UDP	TCP
VFM Server or monitoring agent		6001 and 6005
VFM Replication Agent		6002

NOTES:

- The VFM server and UNIX Replication Agent communicate with storage devices using NFS during NFS migrations.
- The VFM server uses SNMP when identifying Network Appliance® filers.
- The VFM server uses RSH or SSH when issuing some commands to UNIX machines and Network Appliance filers.

Replication and migration problems

When performing replication operations, you may encounter any of the following common errors:

- Network resource not found
- Unexpected network error
- Error copying source

These errors can be caused by denial of service (DoS) network attacks or other network connection failures. For more information about denial of service attacks and other network issues, refer to the following Microsoft Knowledge Base articles:

http://support.microsoft.com/kb/898468/en-us

http://support.microsoft.com/kb/843515/

http://support.microsoft.com/kb/324270

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
VFM reports error messages indicating that the replication agent did not respond either in reports or in relation to replication tasks and your VFM server is installed on Windows 2000 Professional or Windows XP.	• Check the number of inbound connections on the machine hosting the VFM server. The VFM server accepts inbound connections from replication agents in response to commands sent from the server to the replication agents. If these connections in conjunction with other connections existing on the machine exceed the limits imposed by Windows 2000 Professional and Windows XP, the VFM server cannot receive responses from its replication agents and reports errors. Microsoft Knowledge Base article 122920 documents these limits for Windows 2000, and Microsoft Knowledge Base article 314882 describes the limits in Windows XP. If the connection count limits are an issue in your environment, run the VFM server on a Windows 2000 Server to avoid this problem. • Check whether the replication agent can communicate with the machine hosting the VFM server using the NETBIOS host name. Verify that the replication agent can ping the machine hosting the VFM server using the host name. If it is necessary, add an entry to the hosts file on the machine hosting the replication agent, or edit resolv.conf (UNIX replication agents only) so that the machine hosting the replication agent can communicate with the machine hosting the VFM server by using its NETBIOS host name.
VFM documentation mentions NFS file systems but you do not see them when you browse the network from your VFM console machine.	Windows clients need additional software to access NFS file systems. This additional software can be Microsoft tools, such as Services for UNIX® (either client or server), or various third-party NFS redirectors.
VFM reports an error when attempting to copy files when you can successfully copy them using Windows Explorer.	VFM uses the backup read/write APIs to get the most accurate transfer of the source file and its associated attributes. Using these APIs requires specific privileges that your account might lack on the source or destination machines. Although Windows Explorer might copy the file data, attributes such as those governing security might be changed during the course of the transfer. • When copying files with summary information from Windows 2000 NTFS to Windows NT 4 NTFS, you must select “allow loss of additional file streams” to perform the transfer. Windows 2000 NTFS creates a sparse file stream associated with the file when summary information is added to the file. Windows NT 4 NTFS does not support sparse files. • Check network connectivity between machines hosting the replication agent and those hosting the VFM server or monitoring agent.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
VFM Server service or VFM Replication Agent service fail to start.	• Check the Windows event log for errors related to startup failure. • Check the Service Control Manager to see whether the VFM Server or VFM Replication Agent was stopped or the startup of either was disabled. • Check whether the account used by the service was removed from the domain. • Check whether the password was changed on the account used by the service. • Validate that the account and password information used by the service is correct. • Check whether the account used by the service has “logon as service” privilege enabled. • Check whether the VFM Server or VFM Replication Agent is installed on the system. **Note:** When you install VFM, you are prompted for account and password for the VFM Server. The server caches password information using Windows 2000 LSA (Local Security Authority) facilities. The server uses the account and the cached password when deploying replication agents. To manually update information for the VFM server and deployed replication agents, click **Agent Management** on the **Tools** menu.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
Deployment of replication agent fails.	• Check whether the system being deployed to is up and reachable. This can be done by pinging it as well as using the net view command to view the shares that it publishes on the network. • Check whether there are domain trust issues preventing deployment. • Check whether there are network routing and infrastructure issues preventing network connectivity in both directions between the system hosting the server or monitoring agent and the system that hosts the replication agent. • Check whether the intended host of the replication agent is being communicated over a VPN connection. ICMP packets typically do not travel over VPN connections. This restriction prevents the VFM server from correctly identifying the machine type of the intended host machine. It is also possible that the network administrator has restricted ICMP traffic on non-VPN connected machines. • Verify that the target system running one of the following platforms: Windows 2000, Windows Server 2003, or Windows XP. Agents cannot be pushed to other platforms. • If the machine on which you want to deploy a replication agent is running Windows XP SP 2, disable the firewall on that machine and then restart the machine. • If the target system is running Windows NT 4 (versions of VFM 5.5 or earlier), check whether MSXML3 is installed on the system. You can do this by checking for the existence of `msxml3.dll` in the system32 directory. MSXML3 is available for download from the Microsoft web site. • Check whether the normal `admin$` shares are published on the destination system. Try viewing the `admin$` share for the target system using Windows Explorer. • Check whether the account under which the VFM Server service is running has privileges to write to the destination `admin$` share to install the replication agent. You can do this by using the Windows Run As feature to run Windows Explorer and browse the admin$ share on the server. • Verify that there is space on the destination system to receive the replication agent. The replication agent requires approximately 700 KB to be installed. • Verify that the account under which the VFM Server service is running has privileges to install and run as a service on the destination machine. • Check whether there are any firewalls between the system hosting the VFM Server or VFM monitoring agent services and the intended destination system for the replication agents that are blocking the port where the VFM Server service listens (TCP ports 6001 or 6005). • Check whether any process is already listening on the port the VFM Server and VFM Replication Agent services use and thus preventing VFM from operating. • Check whether the Remote Registry Service (`regsvc.exe`) is started on the machine where the replication agent is to be deployed. Stopping this service interferes with the deployment and operation of the Replication Agent.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
	• Check whether the system running on the machine where you want to deploy a replication agent was upgraded from Windows 2000 to Windows 2003. After an upgrade from Windows 2000 to Windows 2003, the Local Service account might not have the required permissions to access certain registry keys. Refer to the following Microsoft Knowledge Base article: http://support.microsoft.com/kb/555335/en-us • If the intended destination is a Microsoft Cluster Server Cluster, additional steps are required to configure the replication agent to run on the cluster. Contact support for this configuration information. • Check whether you are being denied remote access to the registry on the machine where you want to deploy a replication agent. Check the network access settings on the machine where you want to deploy a replication agent. `Administrative Tools>Local Security Policy>Security ettings>Local Policies>Security Options>Network access: Remotely accessible registry paths and Network access:Remotely accessible registry paths and sub-paths.` • Check whether File and Print Sharing is enabled on the system where you want to deploy the replication agent. If this option is not enabled, VFM cannot deploy the replication agent. • If you are running a version of VFM that is later than 5.5, you cannot deploy replication agents to Windows NT 4 systems. If you migrated data to Windows NT 4 systems prior to version 5.7 of VFM and VFM deployed a replication agent to the Windows NT 4 system, you do not have to do anything. In the future, when transferring data to Windows NT 4 systems, specify a proxy machine. • Check whether the replication agent can communicate with the machine hosting the VFM server using the NETBIOS host name. Refer to the note in "Problems selecting a monitoring agent" on page 373 for a list of network ports used by DFS and VFM.
VFM cannot communicate with replication agents.	• Check whether there are any firewalls between the system hosting the VFM Server or VFM monitoring agent services and the intended destination system for the replication agents that are blocking the port where the VFM Replication Agent service listens (TCP port 6002). • Check whether the console system or agent system is running a personal firewall, such as ZoneAlarm or BlackIce and whether the appropriate ports are open on it to allow communication. • Check whether any process is already listening on the port the VFM Server and VFM Replication Agent services use, thus preventing VFM from operating. • Check that the system hosting the replication agent can ping the system hosting the VFM server using the server name. • Check whether the replication agent can communicate with the machine hosting the VFM server using the NETBIOS host name. Refer to the note in "Problems selecting a monitoring agent" on page 373 for a list of network ports used by DFS and VFM.
VFM does not display full information about shares hosted on my filer.	Depending on the filer implementation, VFM might not support the detailed level of information provided by Windows systems. VFM displays as much information as can be obtained about storage targets.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
Invalid security descriptor messages reported during migration operations.	• Check whether the source, the destination, and the machine hosting the DFS root are all in the same domain. If not, check whether trusts are set up appropriately among the systems. • Check whether the domain admin account under which the VFM Server service was installed has administrator privileges on the source and destination of the replication. • Check whether the files for which the problem is reported are owned by a local account on the source machine. This can cause this error to be reported. • It is possible that the security descriptor on the file is invalid. When logged into the account used by the VFM Server service, check whether you can view the security settings on the file from Windows Explorer. • Check whether the owner and primary group of the source files are valid on the destination machine. If the owner or primary group of a file is not present on the destination machine, warnings about invalid security descriptors can result. This can be the case if files and directories on the source machine have a local account as their owner or a local group as their primary group. You can use VFM replication options to control how local SIDs are processed during the replication process. • If the destination of the transfer is a NetApp filer, check the filer's `/etc/messages` file for any error messages indicating problems communicating with its domain controller.
Replication or migration operations produce error 1332 `No mapping between account names and security IDs was done`.	• The files are owned by an unknown SID. Use Windows Explorer or another tool to assign ownership of the problem file to a known trustee. • This condition can arise if the account owning the file is removed from the system.
A replication or migration job is canceled after 60 minutes.	Check whether you have an evaluation license for VFM. Evaluation licenses enable you to copy only 60 minutes of data.
Replication or migration operations produce the following errors: • Disk is full • The system cannot find the file specified	Use the following checks to diagnose environmental problems: • Check whether space is available on the destination. • Check whether the quotas are not limiting the ability of the replication agent to write to the destination. This applies to NTFS quotas in Windows systems or qtree quotas in NetApp filers. • If the destination is hosted on a NetApp filer, check whether the /etc/messages file of the filer for any errors related to MAXDIRSIZE. It might be necessary to increase the value of this parameter on the filer if the destination will contain a very large number of files. • If the destination is hosted on a NetApp filer, check whether any files have UNC paths longer than the limit of 256 characters set by Data ONTAP. This condition can occur if CIFS clients access data via shares of deeply nested subfolders and replication is attempted from the root of a volume, causing a longer UNC path to be used in referencing the file on the filer. • If the destination is hosted on a NetApp filer, check whether the destination volume has sufficient nodes available to accommodate the files being transferred.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
Replication and migration operations produce the following errors: • Resource not found • Path not found • Source not found • Destination not found • Destination could not be created	Use the following checks to diagnose network problems: • Check whether the source server is accessible from the machine hosting the replication agent. For example, you can ping, net use, or browse for the machine in Microsoft Network neighborhood. • Check whether the destination server is accessible from the machine hosting the replication agent. For example, you can ping, net use, or browse for the machine in Microsoft Network neighborhood. • Check that the source and destination shares exist and are accessible from the machine hosting the replication agent. For example, you can try accessing source or destination shares via Windows Explorer. • Check whether the VFM Server service account is not explicitly denied share permissions on destination shares. • Check whether the VFM Server service account is in Local Administrators group of the source or destination server. • Refer to the following Microsoft Knowledge Base articles: http://support.microsoft.com/kb/898468/en-us http://support.microsoft.com/kb/843515/ http://support.microsoft.com/kb/324270
Replication and migration operations produce the following errors: • Cannot delete source/destination. It is being used by another process • Cannot open source/destination. It is being used by another process	Use the following checks to diagnose problems with locked files: • Check whether files or directories in question are being accessed by another application. Close all applications accessing the files or directories and retry data transfer. Files that are commonly locked include Outlook data files such as PSTs, database files used by Access, SQL Server, FoxPro, or Exchange. • Check whether nested data transfers are not in progress at the same time. The following example illustrates this problem: LinkA references \\source\shareA and LinkB references \\source\shareA\folderB. LinkA and LinkB are source links in two different replication policies. Starting replication policies at the same time may lead to above errors. • Check whether the destination of the transfer is a filer, verify that any virus scanning software was disabled, and retry the transfer. When enabled, virus scanning software can cause problems with locked files during large replications.
Replication and migration operations produce the following errors: • "Unable to delete a file. Access is denied" • "Unable to delete source. Access is denied" • "Unable to delete destination. Access is denied" • "Access to a directory or file is denied" • "Access is denied" • "Access to source is denied" • "Access to destination is denied"	Permission problems can be diagnosed by checking whether the following conditions are true for your environment: • Check whether the VFM Server service account is in the Local Administrators group of the source server. • Check whether the VFM Server service account is in the Local Administrators group of the destination server. • Check whether the VFM Server service account is not explicitly denied share permissions on the source or destination share. • If the source or destination is a NetApp filer, use the command testdc to verify that the system can communicate with domain controllers. • Check the filer's `/etc/messages` file for errors related to communication with domain controllers.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
There are warnings when you replicate or migrate data from Windows 2000 NTFS to Windows NT 4 NTFS.	The Windows 2000 version of NTFS supports a superset of the file information supported on Windows NT 4 NTFS. This means that not all the information associated with a file on a Windows 2000 NTFS file system can be transferred to a Windows NT 4 NTFS file system. VFM warns when this situation is encountered. This enables you to specify whether to carry out the operation.
VFM reports access-denied errors when attempting to transfer files to a filer and vscan is in use on the filer.	• Check whether the virus scanning software is operating correctly. It is possible to configure virus scanning software using the vscan interface to prevent any files from being written to the filer when the virus scanning software is inoperable. • Some conditions that can cause problems for virus scanning software include changes to the account under which the virus scanning software runs or network connectivity problems between the filer and the virus scanning host.
Files and directories on the source and destination have different permission settings.	• Check inherited permissions on the source and destination directory trees. • Check whether the data transfer was interrupted before it was completed. VFM manipulates directory permissions on the destination during the migration process to ensure it has sufficient permissions to write to tightly ACL-ed destination directories. After VFM has completed its work in a directory tree, the permissions are set to the actual permission intended for the destination. If the transfer is interrupted, it is possible that extra permissions may be present on destination directories because VFM was prevented from completing the permission-setting phase of the transfer. • Check whether local trustee processing is enabled for the policy. Right-click the policy and click **Properties**. Then, click **Replication Security**. Local trustee processing is not applicable when the source of a transfer is a Domain Controller. Domain Controllers do not host local SIDs.
You see the error 1130 `Not enough server storage is available to process this command`.	Check whether the IRPStackSize parameter is set too low on the system reporting the error. Microsoft KB article 177078 describes the default values for this parameter and explains how to increase it through the registry for versions of Windows, including Windows Server 2003. http://support.microsoft.com/default.aspx?scid=kb;en-us;177078

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...

You have problems transferring data using the NFS protocol.	Check whether the machine where the UNIX replication agent is running can resolve the name of the NFS server you are trying to contact.
	• Log in to the machine where the UNIX replication is installed and ping the NFS server.
	• In a mixed Windows and UNIX environment, always use the fully qualified DNS name.
	Log in to the NFS server and run `exportfs` with no arguments. Check whether the path you are trying to contact or one of its parent directories appears with the rights you expect.
	Check the `exports` file on the NFS server you are trying to contact. This file is found in the following locations:
	• **Linux** - `/etc/exports`
	• **Solaris** - `/etc/dfs/dfstab`
	• **NetApp** - `<the name of the root volume>/etc/exports` OR `/etc/exports`
	Check whether the path you are trying to contact or one of its parent directories is in the `exports` file. If the path exists, check whether the rights granted are correct. To clarify the instructions that follow, SXmachine is used to represent the machine communicating with the UNIX server. The machine can be the VFM server or UNIX replication agent. The entry for the machine is its IP address or fully qualified DNS name.
	• **Linux** - Search for an entry that contains *(rw) or SXmachine(rw). Within the parentheses, there can be other rights, however, it is very important that `rw` (read/write) access is granted.
	• **Solaris and NetApp** - Entries are separated by commas. Search for the entry `rw`, `rw=SXmachine`, or `root=SXmachine`. Make sure that `ro` or `ro=SXmachine` are not specified. The VFM machines might appear in a colon-separated list with other machines.
	Normally, VFM uses root (uid 0). You can specify a different uid for a UNIX machine or filer. When the uid is not root, the replication agent might not be able to copy some security information and therefore reports data transfer failures. To suppress these error reports, perform the following actions:
	1 Right-click the policy that contains the UNIX machine or filer and click **Properties**.
	2 Click **Security Replication**.
	3 Select **Allow loss of security information**.
	4 Click **OK**.
	When an export is used as root (uid 0), special entries are required in the exports file.
	• **Linux** - The VFM machine entry must include `no_root_squash` or `anonuid=0` (for example, `rw,no_root_squash`). Make sure that `root_squash` is not specified.
	• **Solaris and NetApp** - When the VFM machine is not given explicit root access, the entry `anon=0` must be included in the exports file.
	After changing the exports file, run the following commands:
	`exportfs -ua`
	`exportfs -a`
	To check what the NFS machine is exporting, execute the command `exportfs`.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
	Check the UNIX security mode-bit rights on the NFS machine and compare them with the rights of the files you are trying to access. Check whether the replication agent can communicate with the machine hosting the VFM server using the NETBIOS host name. For NetApp filers, check whether the volume/qtree is set to use UNIX security. If the security mode is NTFS, VFM cannot perform attribute settings on the filer volume or qtree. Perform the following actions: 1 The UNIX user must be mapped to an NTFS user in the file `/etc/usermap.cfg`. 2 Select **Allow loss of security information** in the properties of the policy. This suppresses error messages reported by the UNIX replication agent.

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...

Links are not updated after a migration	Make sure that the UNC paths of the source and the link targets you want to update match. If the UNC paths do not match, VFM cannot update the namespace.
	This example illustrates what happens when the UNC path of a migration source and link targets that match the search criteria do not match:
	• You create a migration task and specify a source and destination.
	Source
	`\\Saturn.npb.com\c$\Sales\Reports`
	Destination
	`\\Pluto.npb.com\Sales`
	• You scan for migration candidates and find one candidate.
	Migration candidate
	`\\Saturn.npb.com\c$\Sales\Reports\HoustonRegion`
	• The namespace contains a link named Sales that targets the migration candidate. Notice the difference in the UNC paths between the migration candidate and the link target.
	Namespace link
	`Sales`
	Target of link
	`\\Saturn.npb.com\Sales\Reports\HoustonRegion`
	• You run the policy to start the migration. VFM successfully migrates the data to the destination but does not update the link target because the UNC paths of the migration candidate and the link target do not match.
	Destination after migration
	`\\\Pluto.npb.com\Sales\HoustonRegion`
	Namespace link
	`Sales`
	Target of the link after migration is not updated to target the destination
	`\\Saturn.npb.com\Sales\Reports\HoustonRegion`
	• After migration, the target for the link Sales does not contain any data because the data was migrated to the destination.
	To avoid this problem, the source path should match the link target path.
	Correct source path
	`\\Saturn.npb.com\Sales\Reports`
	Migration candidate
	`\\Saturn.npb.com\Sales\Reports\HoustonRegion`
	Namespace link
	`Sales`
	After migration, the target of the link is updated to target the destination
	`\\Pluto.npb.com\Sales\HoustonRegion`

TABLE 76 Replication and migration problems and solutions

If...	Then, do this...
The Event Details dialog box shows that more data was copied than what you expected	The Event Details dialog box does not accurately show bytes transferred when only the security descriptor is copied to the destination. This occurs when you select the security replication option to **Always copy security settings.**
You cannot uninstall a replication agent using the VFM Agent Management tool.	Check whether the option **Allow replication agent to be deployed to this host** is selected in the properties of the machine hosting the replication agent. 1 In the Physical View, right-click the machine and click **Properties**. 2 Click the **Agent** tab. 3 Make sure that **Allow replication agent to be deployed to this host** is selected. 4 Click **OK**. 5 Go back to Agent Management and uninstall the replication agent.

NetApp filer problems

TABLE 77 NetApp filer problems and solutions

If...	Then, do this...
NetApp filers are identified as generic machines.	Check network connectivity between the VFM console machine and filer. VFM provides a batch file to which you can pass arguments to check the filer. The **FilerCheck.bat** file is in the `Tools\FilerCheck` directory in the VFM server installation directory. At a command prompt type **Filercheck** `<filer> <username:password> <community name>`, where *filer* is the name or IP address of the filer, *username:password* is the username and password used to access the filer through RSH, and *community name* is the SNMP community name used to access the filer through SNMP. Filercheck performs the following checks: • Checks whether the current account has administrator privileges on the filer or vFiler. The lack of these privileges prevents accurate identification. • Checks whether the current account can access `\\filer\IPC$` through a `net use` command. If `IPC$` cannot be accessed, VFM cannot complete the filer identification process. • If you have access to an account with administrator privileges on the filer, you can explicitly map a drive on your VFM console machine to the `\\filer\c$` share using that account to administer the filer through VFM.
NetApp filers are identified as vFilers rather than as filers.	Check that SNMP is enabled on the filer and that the community string used by VFM to communicate with the filer is correct. VFM uses “public” as the default community string. You can override this through the **Tools**>**Options** dialog box, as well as on a per-machine basis through the machine properties dialog box for each specific machine. VFM provides a batch file to which you can pass arguments to check the filer. The **FilerCheck.bat** file is in the `Tools\FilerCheck` directory in the VFM server installation directory. At a command prompt, type **Filercheck** *<filer> <username:password> <community name>*, where *filer* is the name or IP address of the filer, *username:password* is the username and password used to access the filer through RSH, and *community name* is the SNMP community name used to access the filer through SNMP.
FilerView does not appear as expected in the Manage tab.	Check whether you can browse to `http://filer/na_admin` using Internet Explorer on the VFM machine to get to the FilerView Web user interface.

TABLE 77 NetApp filer problems and solutions

If...	Then, do this...
Attempting a filer management operation, such as creating or initializing a SnapMirror® function or showing a filer in the Admin View, causes the following errors: • Device command didn't report an output • SnapMirror is too busy, please try again later	• The NetApp filers are under an extremely heavy load. Retry the operation when the load on the filers has decreased. • Check whether the `/etc/snapmirror.conf` file is 0 bytes. If `snapmirror.conf` does not contain the normal comment showing when the file was generated by Data ONTAP software, add a placeholder comment or delete the file. • Check the `file /etc/messages` on the filer for SnapMirror-related error messages or RSH-related error messages. **Note:** VFM does not support automatic creation of SnapMirrors on separate network interfaces other than the primary interface used by VFM to communicate with the filers.
Refreshing filer information or filer event information results in `access denied` errors.	VFM provides a batch file to which you can pass arguments to check the filer. The **FilerCheck.bat** file is in the `Tools\FilerCheck` directory in the VFM server installation directory. At a command prompt type: **Filercheck** *<filer> <username:password> <community name>* where *filer* is the name or IP address of the filer, *username:password* is the user name and password used to access the filer through RSH, and *community name* is the SNMP community name used to access the filer through SNMP. Filercheck performs the following checks: • Checks whether the current account has administrator access to the c$ share on the filer. To manually check access, try to browse the c$ share using Windows Explorer (for example, `\\<filename>\c$`). • Checks whether the current account has administrator privileges on the filer or vFiler. Lack of these privileges results in `access denied` errors. • If RSH reports: `<filename>: Permission denied. RSH: cannot establish connection,` the filer is configured with the administrator name and possibly the administrator password. • If RSH reports something like: `NetApp Release 6.3: Wed Aug 7 23 : 04 : 57 PDT 2002,` the RSH command succeeded against the filer. If the filer is configured with a user name and password, you must enter the user name and password on the Filer Account tab of the machine's properties in VFM. If the RSH check is indeterminate, you can use telnet to check whether the filer is prompting for a user name and password. From a command prompt on the VFM client machine, telnet to the filer and check whether the filer prompts for a user name and password.
SnapMirrors do not appear in the properties of a NetApp filer.	Check whether the SnapMirrors were created using a command-line prompt. • When a SnapMirror is created using FilerView, the SnapMirror configuration is added into a file named **snapmirror.conf**. • When FilerView receives the command to display SnapMirrors on a filer, the product reads the information in the **snapmirror.conf** file. • When a SnapMirror is created using a command-line prompt, the configuration of the SnapMirror is not added to the **snapmirror.conf** file.

Browse namespace problems

TABLE 78 Browse namespace problems and solutions

If...	Then, do this...
Attempts to browse or manage the DFS namespace yield “RPC server not available” error messages.	Check whether the DFS service is running on the machine hosting the DFS root.
Attempts to browse the namespace yield “access denied” errors.	Check whether the current account has privileges to access the share hosting the DFS root, as well as to the root itself in the case of domain roots. If access denied errors are reported when attempting to browse folders in the namespace, check the permissions of the shares hosting these folders as well as the permissions on the folders themselves. When accessing resources through the namespace, the credentials used for the access are those used when the user initially accesses the namespace. This can cause problems in environments where users typically use explicit credentials to map network drives.

Windows NAS or storage server filer identification problems

TABLE 79 Windows powered NAS or storage server filer identification problems and solutions

If...	Then, do this...
Windows Powered NAS or Windows Storage Server filers are being identified as generic machines.	Check whether the current account has administrator privileges on the Windows Powered NAS or Windows Storage Server machine by using RegEdit.exe to connect to the filer's registry. Inability to open the filer's registry prevents VFM from accurately identifying the filer as a Windows Powered NAS or Windows Storage Server filer.

Problems running batch files

TABLE 80 Batch file execution problems and solutions

If...	Then, do this...
Problems are encountered trying to run batch files configured to run with VFM storage policies.	• For more information about running batch files in VFM and about the example scripts provided by VFM, refer to “Run Batch Files” on page 321. • Save your BAT file log to a text file. You can examine the contents of the text file to diagnose potential problems. Adding a line such as `echo %date% %time%` The arguments passed from server are: `%* >> Decision.txt` will record all the arguments passed to your BAT file along with the date and time that the script ran. This information is appended to the file Decision.txt located in the directory where the BAT file is running
You get the following error when you know that you correctly typed the path of the batch file: `Script was not found`	Check whether the script path contains spaces. If the path contains spaces, you must enclose the path in double quotes. If you browse for the script, VFM encloses the path in double quotes for you.
You specify a script, run the policy, and the script does not run.	Check whether the script path is valid. VFM does not validate the path to the script you specify. It is your responsibility to ensure that the path is accurate and accessible to the VFM server or monitoring agent that invoke the script.

Problems running a policy

TABLE 81 Policy execution problems and solutions

If...	Then, do this...
You see the following error messages: • The namespace backup policy failed. • Manually running the namespace backup policy failed.	• Check whether the account running VFM has appropriate permissions for the DFS root being manipulated. • Check whether VFM can communicate with the DFS root. Check network connectivity by pinging the server where the root is located from the VFM machine. • Check whether there is a problem communicating with the VFM database.
You see the following error messages: • The server disaster recovery policy was unable to retrieve the active link information • The server disaster recovery policy failed to activate the primary server target for the link • The server disaster recovery policy failed to activate the secondary server target for the link	• Check whether VFM can communicate with the servers that contain the link targets. • Check whether the account running VFM has appropriate permissions for the DFS root being manipulated. • Check whether VFM can communicate with the DFS root. Check network connectivity by pinging the server where the root is located from the VFM machine.
You see the following error messages: • The policy scan for the storage load balancing policy encountered problems • Failed updating DFS links for a task in the storage load balancing policy	• Check whether the account running VFM has appropriate permissions for the DFS root being manipulated. • Check whether VFM can communicate with the DFS root. Check network connectivity by pinging the server where the root is located from the VFM machine. • Check whether VFM can communicate with the storage resources. • Check whether the account running VFM has appropriate permissions to access the storage resources.

Snapshot access policy problems

TABLE 82 Snapshot access policy problems and solutions

If...	Then, do this...
The restore operation is not working.	Check whether the volume hosting the files has the .snapshot directory visible. VFM cannot restore data hosted on a volume where the.snapshot directory is not visible.

TABLE 82 Snapshot access policy problems and solutions

If...	Then, do this...
You cannot view files on the Explorer tab of a snapshot access policy.	• The **.snapshot** directory on the volume hosting the data might not be visible. Make sure that the operating system is configured to show hidden files. • To access snapshots from CIFS clients, type the following command on a command prompt: `\\<filename>\<share>\.snapshot` Where *filename* is the name of the filer you are using and *share* is the name of the share you want to access.
You encounter `network resource not found` errors when attempting to use the SnapRestore function to retrieve items from the Snapshot Access Namespace Explorer view.	Check to see whether there is a mismatch between the exact local path on disk of the shared directory and the local path entered when creating the path. These two items must match exactly, including upper/lower case of the path components, for SnapRestore to work. You can correct this by removing and re-creating the share using the exact name of the local path and re-creating your Snapshot access policy.

Problems printing to PDF files

TABLE 83 PDF file printing problems and solutions

If...	Then, do this...
You cannot print from VFM to PDF files.	Check whether VFM was upgraded using terminal services. You need to upgrade from the machine console and not through terminal services.

Name of machine is incorrect

TABLE 84 Incorrect machine name problems and solutions

If...	Then, do this...
The name of a machine in the Physical View is incorrect.	Check whether DNS forward and reverse lookups are configured correctly for the machine.

Problems taking snapshots of the application

TABLE 85 Application snapshot problems and solutions

If...	Then, do this...
You see access denied errors while trying to take a snapshot of the application.	Check the permissions on the snapshot save location (both share and local directory) to ensure that both the account under which MSDE and the account under which the VFM server are running have permission to write to the snapshot save location.
The application snapshot fails to complete	• Check whether the MSDE database runs with the LOCALSYSTEM account. Application snapshots cannot be written to UNC path hosted on a different machine. • Use the Service Control Manager to change the logon account for the appropriate MSDE services to an account with privileges to write to a network share. • It is not recommended that MSDE run as the same account as the VFM server because the account used by the VFM server has administrator privileges. If MSDE is attacked by programs such as Slammer, the damage might be lessened if MSDE is running as a minimally privileged account.

Permission-based namespace policy problems

TABLE 86 Permission-based namespace policy problems and solutions

If...	Then, do this...
Links were not created when you ran the policy.	Check whether you tried to add multiple links with the same name. If you selected domains to search for resources and the domains contain groups with the same name (for example, `domain1\sales` and `domain2\sales`), make sure that you select a template that sorts by domain, such as Domain-Group Name-Machine-Shares, Machines-Domain-Group Name-Shares, or Domain-Group Name-Shares. You cannot create links with the same name under the root or a logical folder of the namespace. • Check whether the NTFS permissions for a folder hosting a share to which you want the policy to create link contain `<hostname>\users`. Remove the `<hostname>\users` group from the NTFS permissions on the folder and run the policy. • Check whether the resources you wanted to add to the namespace have the local users group in their permissions. Remove the local user group from the permissions for each folder and run the scan again.

Problem installing a UNIX replication agent

TABLE 87 UNIX replication agent installation problems and solutions

If...	Then, do this...
You get the following error message while trying to install a UNIX replication agent: • On Solaris: pkgadd: ERROR: Unable to make directory `</usr/local/URA>` • On Linux: error: unpacking of archive failed on file `/usr/local/ura/bin`: cpio: mkdir failed - No such file or directory	Check whether the file system where you are installing the replication agent is mounted read-write. The file system where a replication agent is installed must be mounted read-write.

Problems using NTFS Change Journal

TABLE 88 NTFS Change Journal usage problems and solutions

If...	Then, do this...
You see one of the following error messages: • Change Journal could not be used because this is the first time the job ran. • Change Journal could not be used because it is not active and auto-enabling is not permitted. • Change Journal could not be used due to an unexpected failure.	The NTFS Change Journal is supported on Windows 2000 NTFS which is available on Windows version 2000 and later of the Windows operating system. • When you enable the NTFS Change Journal on a volume, the size of the volume is used to determine the amount of space dedicated to the change journal. For each 50 GB of space on the volume, 128 MB is assigned for use by the change journal up to a maximum of 10 GB. On Windows 2000 SP2 or below, the amount of space dedicated to the change journal is always 128 MB, regardless of the size of the volume. • Use of this option requires that VFM deploy a replication agent on the machine where the source data resides and enable the change journal on the NTFS volume hosting the source data. The VFM machine properties dialog enables you to control whether the replication agent can be deployed to a machine and whether VFM can enable the NTFS change journal for volumes on the machine. • Use of this option means that changes made on the source are replicated to the destination and changes made on the destination are not overwritten by data from the source unless the corresponding file on the source changes as well. This differs from standard VFM replication, wherein changes made on the destination are detected and overwritten with data from the source. • In cases where the change journal wraps, VFM must perform a full scan of the source and destination to determine what files need to be transferred.

Problem configuring the OpenSSH package on Solaris

TABLE 89 OpenSSH configuration problems and solutions on Solaris

If...	Then, do this...
You see the error message `Privilege Separation user sshd does not exist`	Edit the **/etc/passwd** file and add the following line at the bottom of the file: `sshd:x:74:74:Privilege-separated SSH:/var/empty/sshd:/sbin/nologin` **Note**: The command should not have line breaks. It should be written in one line.

Problem accessing a UNIX machine when using SSH

TABLE 90 UNIX machine access problems and solutions when using SSH

If...	Then, do this...
You are prompted for a password when you access a UNIX machine when VFM is configured to use SSH	If you generated a SSH2 (RSA) or SSH2 (DSA) key, check whether you created the file `authorized_keys2`. • Check whether you generated the keys without a password or passphrase. • Check whether OpenSSH is installed on the UNIX machine. • Make sure that you have the appropriate permissions. • Check whether the private keys were added into the memory of the VFM client and server machines. To check whether the private keys were added to the pageant key list, right-click the pageant icon on the system tray and click **View Keys**. If the private key is not listed, click **Add key**.

Problem importing snapshots of the application

TABLE 91 Application snapshot importing problems and solutions

If...	Then, do this...

You see the following error message when you try to import a snapshot of the application:
`ERROR: Cannot execute command: "C:\Program Files\Brocade\VFM\Server.exe""-restore:\\jsmith\myshare\mydatabase"`
`ERROR: desc: Run status: -1`
	Check whether the SQL service is running on the Local Network Account. When the SQL service is running on the Local Network Account on the backup machine, perform one of the following tasks:
	• Save the exported snapshot on the local drive of the backup machine before you import it to VFM.
	• Make sure that the SQL service is running under an account that has full privileges across all domains.

Problems replicating data using VSS snapshots

TABLE 92 Data replication problems using VSS snapshots

If...	Then, do this...
Time-out errors occur in volume shadow copy service writers, and shadow copies are lost during backup of high input/output activity	Apply a hot fix update to the Volume Shadow Copy service (VSS) for Microsoft Windows Server 2003 which is available at the following Web site: http://support.microsoft.com/kb/833167#8 The update addresses this and other known VSS problems.

Problems deleting items from the Logical View

TABLE 93 Logical view deletion problems and solutions

If...	Then, do this...
• You cannot delete a root, link, or link target from the Logical View. • You get the following error message when you try to delete a link from the Logical View: `Unable to delete link because the link is being reference by another object` • You get the following error message when you try to delete the target of a link: `Deleting the last replica for a DFS link will cause the link to be deleted. This link is being referenced by another object and may not be deleted.`	The item that you want to delete is a member of a policy. Before you can delete the item, you must remove it from the policy or policies that reference the item.

Problem updating home directories

TABLE 94 Home directory updating problems and solutions

If..	Then, do this...
You run a user data management policy after you manually delete a link managed by the policy. In the Alerts tab, you see the following error message: `"Namespace link was not found. This error will prevent you from applying the namespace options."`	The links that a user data management policy creates in the namespace should always be managed by the policy. If you manually delete a link that was created by a user data management policy, the policy cannot manage the home directory the link targeted because the location of the directory is no longer available. **Links targeting home directories contained in Windows systems** • If you deleted a link targeting the home directory of an active user, manually update the user's Active Directory Home Folder property. Type the correct physical UNC path and run the policy again. • If you deleted a link targeting the home directory of a disabled user and the user data management policy was configured to migrate or retarget links of disabled users, update the user's Active Directory Home Folder property. Type the correct physical UNC path and run the policy again. • If you deleted a link targeting the home directory of a deleted user and the user data management policy was configured to migrate or retarget the links of deleted users, you must create a VFM migration policy to migrate the data. • If you deleted a link targeting the home directory of an unmanaged user and the user data management policy was configured to update the user's Home Folder property in Active Directory, you must manually update the Home Folder property with the physical UNC path to the user's data. **Links targeting home directories contained in NetApp filers in which the /etc/cifs_homedir.cfg file is configured** If you deleted a link targeting the home directory of any user, the user data management policy creates the link again the next time the policy runs.

Migration, replication, and status indicators

The following sections describe the various icons and graphic indications used throughout the product.

Migration task indicators

The following list specifies status indicators that are used as visual cues.

TABLE 95 Indicators for migration tasks

Source Share	Description
	The task is scheduled to be processed.
	User action is required.
	A task was completed successfully or was canceled.

Replication status monitoring indicators

The following list specifies status indicators that are used as visual cues.

TABLE 96 Indicators for migration tasks

Replication policy	Description	Example
	Replication is in progress or completed.	
	Warning: • A link in the policy is in an abnormal state. • A target in the policy is in an abnormal state.	• Replicas could be offline. • Access permissions prevent VFM from completing some aspects of a link. • A target might be unreachable due to network or machine problems. • Access permissions on a target might prevent VFM from accurately assessing the state of the target.
	Error - VFM detected a serious problem.	• The policy has no source link or no destination link. • The last synchronization attempt failed. • The source link has more than one target.
	Unknown - VFM cannot determine the state information for one or more items in the policy.	

TABLE 97

Target	Description	Examples
	In progress - Data is currently being moved to the target.	
✓	Completed - Data move to the target was completed successfully.	All files were copied. No files needed to be copied.
	Warning - Data move to the target encountered some failures.	• Some files failed to be copied. • No files were copied, but some failed to copy due to sharing violations. • All files were copied, but some files had warnings.
	Error - The target is in an abnormal state.	The target might be unreachable due to network or machine problems. Access permissions on the target might prevent VFM from accurately assessing the state of the target. There were problems with a DFS root. No files were copied for reasons other than sharing violation.
?	Unknown - VFM cannot determine the state information for the target.	

Status indicators

The following tables identify status indicators that are used as visual cues.

DFS root

TABLE 98 DFS root status indicators and examples

DFS root	Description	Example
 Note: Single and replicated DFS roots have the same status symbols.	**Online**	
	Warning - VFM detected that one or more elements being monitored are not operating correctly.	• A DFS link has some replica shares that are unavailable. • Access permissions prevent VFM from monitoring the status of some portion of the root. • FRS is enabled for some, but not all of the replicas for this domain-based root.
	Error - VFM detected a serious problem.	• The internal data structure of the DFS root is corrupted. • All the replicas for a link are offline. • The DFS root is offline or does not respond to queries.
	Unknown - VFM cannot determine the state information for one or more DFS elements.	

DFS link

TABLE 99 DFS link indicators and examples

DFS link	Description	Example
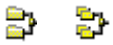 **Note**: Single and replicated DFS links have the same status symbols.	**Online**	
	Warning - VFM detected that one or more shared folder replicas being monitored are not operating correctly.	• A DFS link has some replica shares that are unavailable. • Access permissions prevent VFM from monitoring the status of some portion of the link. • FRS is enabled for some, but not all of the replicas for this link in a domain-based root. • A single-target link is a member of a recovery policy.
	Error - VFM detected a serious problem.	All the replicas for a link are unavailable.
	Unknown - VFM cannot determine the state information for one or more replicas.	

Shared folder

TABLE 100 Shared folder status indicators and examples

Shared folder	Description	Example
	Online	
	Warning - VFM detected that the shared folder being monitored is not operating correctly.	• The shared folder is unavailable. • Access permissions prevent VFM from monitoring the status of the shared folder.
	Error - VFM detected a serious problem.	The shared folder is offline.
	Unknown - VFM cannot determine the state information of the shared folder.	

Logical folder

TABLE 101 Logical folder status indicators and examples

Logical folder	Description	Example
	Online	
	Warning - VFM detected that the folder being monitored is not operating correctly.	• The shared folder is unavailable. • Access permissions prevent VFM from monitoring the status of the shared folder.
	Error - VFM detected a serious problem.	The shared folder is offline.
	Unknown - VFM cannot determine the state information of the folder.	

Machine

TABLE 102 Machine status indicators and examples

Machine	Description	Example
Note: All the machines have the same status symbols.	Information about the machine was updated without errors or **Unknown** - The machine was not refreshed.	
	Refreshing - Information about the machine is being updated.	
	Error - VFM detected a problem.	• There is a problem refreshing the server information (for example, comment, IP address or platform) • There is a problem refreshing the NetApp filer information (for example, volume or SnapMirror) • There is a problem refreshing the NetApp filer events

Volume

TABLE 103 Volume status indicators and examples

Volume	Description	Example
	Information was updated without errors or **Unknown** - The volume was not refreshed.	
	Refreshing - Information about the volume is being updated.	

Qtree

TABLE 104 Qtree status indicators and examples

Qtree	Description	Example
	Information about the qtree was updated without errors or **Unknown** - The volume was not refreshed.	
	Refreshing - Information about the qtree is being updated.	

Export

TABLE 105 Export status indicators and examples

Export	Description	Example
	Online	
	Warning - VFM detected that the export being monitored is not operating correctly.	• The export is unavailable. • Access permissions prevent VFM from monitoring the status of the export.
	Error - VFM detected a serious problem.	The export is offline.
	Unknown - VFM cannot determine the state information of the export.	

VFM server and monitoring agent

TABLE 106 VFM server and monitoring agent status indicators and examples

VFM server and monitoring agent	Description	Example
	No problems were detected	
	Warning - VFM server is in an unknown state.	
	Error - VFM detected a problem.	• The database data file is full. • The database log file is full.

Exclusion and Inclusion Examples

The following sections provide examples of how exclusion and inclusion lists work.

Exclusion examples

When creating exclusion lists, remember the following points:

- Exclusion terms are processed before inclusion terms. Files or directories matching an exclusion term are not processed further and cannot be included by a subsequent inclusion term.
- Directory exclusions are search expressions comprised of folder names and/or folder paths that can be used to exclude any folders matching the specified exclusion criteria.
- File exclusions are search expressions comprised of file names or extensions that can be used to exclude files matching the specified exclusion criteria.
- DOS wildcards can be used alone or in concert with a character string to exclude file or folders of similar name or type.
  - To match a string of characters of any length including an empty string, type an asterisk (`*`). For example, `s*d` finds `sad`, `suddenly`, and `Saturday`,
  - A question mark (?) can be used to match any single character in the search text. For example, `s?t` finds set and `sit`, but not `Saturday`.
- To exclude a subfolder and its contents, the trailing backward slash or forward slash is optional (The subfolder and its contents are excluded whether or not the trailing backward slash or forward slash is present).
- When you enable “delete orphans” as a replication option, the exclusion list affects how VFM handles orphan files on the destination. Orphan files on the destination that match the exclusion term are not deleted, while orphan files not matching the exclusion term are deleted.
- You must delimit multiple directory or file exclusion parameters with double quotes or you can click **Expand** to specify multiple files, file types, or directories.
- When the **Include subfolders** replication option is enabled and you specify file exclusions, VFM traverses subfolders for files to exclude from the destination. Subfolders are always transferred to the destination even when they are empty.

The following table provides examples of valid file and folder exclusion formats and descriptions for each format.

TABLE 107 Valid file and folder exclusion formats and descriptions

Item	Format	Applies to	Description
1	*	Files or Directories	When this example is used specifically as a directory exclusion parameter, all files and folders will be excluded. When used specifically as a file exclusion parameter, folders are still included, but all files contained in those folders will be excluded.
2	*.*	Files	All files contained within the specified parent and their associated child folders are excluded.
3	\ (Windows) / (UNIX)	Directories	When the slash is used as a directory exclusion parameter, all files and folders contained within the specified parent and child folders are excluded.
4	* (Windows) /* (UNIX)	Files or Directories	When used as a directory exclusion parameter, this example will exclude all files and all folders. When used as a file exclusion parameter, this example will result in the exclusion of all files.
5	*.* (Windows) /*.* (UNIX)	Files	When used as a file exclusion parameter, this string excludes all files contained in the parent.
6	*.tmp	Files	This example can be used to exclude files that contain the literal character string `.tmp`.
7	listing	Files or Directories	Exclude files or directories named `listing`.
8	listing.txt	Files	Exclude files named `listing.txt`.
9	listing???.txt	Files	This example can be used to exclude files that begin with the string `listing` followed by any three characters and then the literal character string `.txt`.
10	listing*.txt	Files	Used to exclude all text files with names that begin with the literal character string `listing`.
11	\listing (Windows) /listing (UNIX)	Directories	When prefaced with the \ character, this example will traverse the directory tree and exclude folders named `listing`.
12	listing\ or \listing\ (Windows) listing/ or /listing/ (UNIX)	Directories	This example will traverse the directory tree and exclude folders that are subfolders of the `listing` directory.
13	*\listing (Windows) */listing (UNIX)	Directories	Exclude subfolders named `listing`.
14	*\listing\file.txt (Windows) */listing/file.txt (UNIX)	Files	This example can be used to exclude files named `file.txt` located in folders named `listing`.

Inclusion examples

When creating inclusion lists, remember the following points:

- Exclusion terms are processed before inclusion terms. Files or directories matching an exclusion term are not processed further and cannot be included by a subsequent inclusion term.
- Directory inclusions are search expressions comprised of folder names or folder paths that can be used to include any folders matching the specified inclusion criteria.
- File inclusions are search expressions comprised of file names or extensions that can be used to include files matching the specified inclusion criteria.
- DOS wildcards can be used alone or in concert with a character string to include file or folders of similar name or type.
- To match a string of characters of any length including an empty string, type an asterisk (`*`). For example, `s*d` finds `sad`, `suddenly`, and `Saturday`,
- A question mark (?) can be used to match any single character in the search text. For example, `s?t` finds `set` and `sit`, but not `Saturday`.
- To include a subfolder and its contents, specify a backward slash \ (Windows) or a forward slash / (UNIX) after the directory name.
- To include just a subfolder and not its contents, omit the trailing backward slash or forward slash.
- When you enable “delete orphans” as a replication option, the inclusion list affects how VFM handles orphan files on the destination. Orphan files on the destination that match the inclusion term are deleted.
- You must delimit multiple directory or file inclusion parameters with double quotes or you can click **Expand** to specify multiple files, file types, or directories.
- When the **Include subfolders** replication option is enabled and you specify file inclusions, VFM traverses subfolders for files to include in the destination. Subfolders are always transferred to the destination even when they are empty.

The following table provides examples of valid inclusion formats and descriptions for each format.

TABLE 108 Valid file and folder inclusion formats and descriptions

Item	Format	Applies to	Description
1	*	Files or Directories	When this example is used specifically as a directory inclusion parameter, only folders (no files) of the specified parent will be included. When it is used as a file inclusion parameter, all files and all folders will be included.
2	*.*	Files	All files contained within the specified parent and their associated child folders will be included.
3	\ (Windows) / (UNIX)	Directories	When used as a directory inclusion parameter, all folders contained within the specified parent are included.
4	* (Windows) /* (UNIX)	Files or Directories	When used as a directory inclusion parameter, this example will include all folders. When used as a file inclusion parameter, this example will include all files and related folders.
5	*.* (Windows) /*.* (UNIX)	Files	When used as a file inclusion parameter, this string includes all files contained in the parent.

TABLE 108 Valid file and folder inclusion formats and descriptions

Item	Format	Applies to	Description
6	*.tmp	Files	This example can be used to include files that contain the literal character string `.tmp`.
7	listing	Files or Directories	Include Files or Directories named `listing`.
8	listing.txt	Files	Match files or directories named `listing.txt`.
9	listing???.txt	Files	This example can be used to include all text files with names that begin with the string `listing` followed by any three characters.
10	listing*.txt	Files	Used to include all text files with names that begin with the literal character string `listing`.
11	\listing (Windows) /listing (UNIX)	Directories	When prefaced with the \ character, this example will traverse the directory tree and include folders named `listing`.
12	listing\ or \listing\ (Windows) listing/ or /listing/ (UNIX)	Directories	This example will traverse the directory tree and include folders that are sub directories of the “listing” directory.
13	*\listing (Windows) */listing (UNIX)	Directories	Include subfolders named “listing.”
14	*\listing\file.txt (Windows) */listing/file.txt (UNIX)	Files	This example can be used to include files named “file.txt” located in a subfolder named “listing.”

Appendix

Event Descriptions

E

Events logged as a result of errors are always recorded in the Windows event log. You can record the event log changes made to the global namespace or to VFM settings. To select this option, click **Options** on the **Tools** menu, click **Systems Options**, and select **Write audit events to NT event log**.

VFM events listed in the following tables with `Optional` under the **NT Event Log Entry** column are logged to the Windows event log when **Write audit events to NT event log** is selected.

In this appendix

- 1000 - 1099 Multiprotocol Namespace Event IDs 404
- 1100 - 1199 Home Folder Event IDs 405
- 1200 - 1299 Disaster Recovery Event IDs 405
- 1300 - 1399 Replication Agent Event IDs 407
- 1400 - 1499 Server Event IDs 408
- 1500 - 1699 SnapMirror Event IDs 408
- 1700 - 1799 Namespace Availability Event IDs 410
- 1800 - 1899 Multiprotocol Namespace Event IDs 411
- 1900 - 1999 Service Event IDs 411
- 2000 - 2099 Migration Event IDs 412
- 2100 - 2199 Namespace Creation Event IDs 413
- 2200 - 2299 Snapshot Access Event IDs 414
- 2300 - 2399 Filer Disaster Recovery Event IDs 414
- 2400 - 2499 Snapshot Scheduling Event IDs 416
- 2500 - 2599 Namespace Backup Event IDs. 417
- 2600 - 2699 Logical Namespace Event IDs 417
- 2700 - 2799 Archival Migration Event IDs. 419
- 2800 - 2899 Server Disaster Recovery Event IDs 420
- 2900 - 2999 Replication and Monitoring Agents Event IDs. 422
- 3100 - 3199 Storage Load Balancing Event IDs 423
- 3200 - 3299 Permission-Based Namespace Event IDs 424

1000 - 1099 Multiprotocol Namespace Event IDs

TABLE 109 Event IDs 1000 to 1099

Event ID	NT Event Log Entry	Description
1000	Optional	Replication policy was created.
1001	Optional	Replication policy was deleted.
1002	Optional	Replication policy was renamed.
1003	Optional	The links associated with a replication policy were changed.
1004	Optional	The replication schedule for a replication policy were changed.
1005	Optional	The replication options for a replication policy were changed.
1006	Optional	The distribution plan for a replication policy was changed.
1007	Optional	Evaluation of a replication policy was started by the user.
1008	No	Replication started for a replication policy.
1009	No	Replication completed successfully for a replication policy.
1010	Yes	Replication failed for a replication policy.
1011	Yes	Replication could not be started for a replication policy.
1012	Yes	Replication could not be started for a replication policy because one or both of the links was unavailable, or in an invalid state.
1013	No	Replication could not be started for a replication policy because none of the replication targets are configured to use VFM replication.
1015	Optional	The schedule for retrying failed attempts at a replication policy was changed. The policy shows the old and new schedule.
1016	Optional	The replication options for a replication policy were changed. The policy shows the old and new values for the options that were changed.
1017	Optional	The monitoring agent for a replication policy was changed. The policy shows the old and new monitoring agent.
1018	No	Replication was canceled for the replication policy.
1019	Yes	Replication was completed with warnings for replication policy.
1020	Optional	The paths associated with the replication policy were changed.
1021	Yes	Replication was canceled because of snapshot problems. Failed replicating %3 with %2 using agent %4. The server was not able to take a snapshot.
1022	Yes	Replication failed because the agent on machine %1 did not respond after receiving the job. Failed replicating %3 with %2 using agent %1.
1023	Yes	Replication failed because the policy is configured with options not compatible with snapshot-based replication. Incompatible Option: %1
1024	Yes	There was an error trying to create a snapshot. The error message is : %1 Additional error information: Server Name: %2 Volume: %3 Policy Name: %4 Exposed Path: %5

1100 - 1199 Home Folder Event IDs

TABLE 110 Event IDs 1100 to 1199

Event ID	NT Event Log Entry	Description
1100	Optional	A home folder policy was created.
1101	Optional	A home folder policy was deleted.
1102	Optional	A home folder policy was renamed.
1103	Optional	A home folder policy was modified.
1104	Optional	Evaluation of a home folder policy was started by the user.

1200 - 1299 Disaster Recovery Event IDs

TABLE 111 Event IDs 1200 to 1299

Event ID	NT Event Log Entry	Description
1200	Optional	A recovery policy was created.
1201	Optional	A recovery policy was deleted.
1202	Optional	A recovery policy was renamed.
1203	Optional	A DFS link was added to a recovery policy.
1204	Optional	A DFS link was removed from a recovery policy.
1205	Optional	A link target was added to a DFS link in a recovery policy.
1206	Optional	A link target was removed from a DFS link in a recovery policy.
1207	Optional	Options were changed for a recovery policy.
1208	Optional	The monitoring schedule for a recovery policy was changed.
1209	Optional	The target priority was modified for a recovery policy.
1210	Optional	The status was changed for a DFS link in a recovery policy.
1211	Optional	The replication schedule for a recovery policy was changed.
1212	Optional	Evaluation of a recovery policy was started by the user.
1213	No	Replication was started for a recovery policy.
1214	No	Replication was completed for a recovery policy.
1215	Yes	Replication failed for a recovery policy.
1216	Optional	The name was changed for a DFS link in a recovery policy.
1217	Yes	A DFS link in a recovery policy was failed over manually.
1218	Yes	A DFS link in a recovery policy was failed back manually.
1219	Yes	A DFS link in a recovery policy failed over automatically.
1220	Optional	A DFS link in a recovery policy had its active target manually set to another target.
1221	Optional	The status of a DFS link was changed in a recovery policy.

TABLE 111 Event IDs 1200 to 1299

Event ID	NT Event Log Entry	Description
1223	Optional	The schedule for retrying failed attempts at monitoring policies for a recovery policy was changed. The policy shows the old and new schedule.
1224	Optional	The schedule for retrying failed replication jobs for a recovery policy was changed. The policy shows the old and new schedule.
1225	Optional	The options for a recovery policy were changed. The policy shows the old and new failover values.
1226	Optional	The failover script settings for a recovery policy were changed. The policy shows the old and new script settings.
1227	Optional	The monitoring agent for a recovery policy was changed. The policy shows the old and new monitoring agent.
1228	Optional	The replication options for a recovery policy were changed. The policy shows old and new values for the options that were changed.
1229	Yes	Error executing a command for a recovery policy.
1230	Yes	There were no targets for a disaster recovery link on a recovery policy.
1231	Yes	There was an error getting information for a disaster recovery link.
1232	Yes	There was an error. More than one target online for a link in a recovery policy.
1233	Yes	No target was online for a link in a recovery policy.
1234	Yes	Unable to [fail back, fail over, post-fail over, or pre-fail over] a disaster recovery link policy.
1235	Yes	There was an error setting a target online for a disaster recovery link.
1236	Yes	There was an error setting a target offline for a disaster recovery link.
1237	Optional	The SnapMirror replication options for the recovery policy were changed.
1238	Yes	The link needs to be manually failed over.
1239	Yes	There was an error executing a pre-failover command for the recovery policy.
1240	Yes	There was an error executing a post-failover command for the recovery policy.
1243	Yes	Replication was completed with warnings for the client recovery policy.
1244	Optional	The fail count for the recovery policy was changed.
1245	Optional	The link target is unreachable for the recovery policy. No failover is needed because the fail count has not been reached.
1246	Optional	The User Decision Script settings for the recovery policy were changed.
1247	Yes	Error executing the User Decision Script for the recovery policy.
1248	Optional	The User Decision Script is missing for the recovery policy.
1249	Optional	The link needs to be failed over for the recovery policy because the User Decision Script returned a warning.
1250	Optional	The search options settings for the recovery policy were changed.
1251	Optional	The search ignore list was changed for the recovery policy.
1252	Optional	The following search paths were added to the recovery policy.
1253	Optional	The following search paths were removed from the recovery policy.

TABLE 111 Event IDs 1200 to 1299

Event ID	NT Event Log Entry	Description
1254	Optional	Cannot locate and remove the following search paths from the recovery policy.
1255	Optional	The include/exclude list settings for the recovery policy were changed.
1256	Optional	Internal errors occurred executing the User Decision Script for the recovery policy.
1257	Optional	The User Decision Script requested a failover for a link in the recovery policy but no failover is needed because the fail count was not reached.
1258	Optional	A new link was detected for the recovery policy but the link cannot be added for the stated reasons.
1259	Optional	The link target for the link specified is not available.
1260	Optional	Evaluation of Recovery Policy was started.
1261	Optional	Evaluation of Recovery Policy was cancelled.

1300 - 1399 Replication Agent Event IDs

TABLE 112 Event IDs 1300 to 1399

Event ID	NT Event Log Entry	Description
1300	Optional	Copied data successfully from one target to another target.
1301	Yes	Failed to copy data from one target to another target.
1302	Optional	Moved data successfully from one target to another target.
1303	Yes	Failed to move data from one target to another target.
1304	Optional	Successfully synchronized data from one target to another target.
1305	Yes	Failed to synchronize data from one target to another target.
1306	Yes	The replication agent on a specified server is not responding.
1307	Yes	The server could not queue a replication job for the agent.
1308	Optional	Successfully deleted files.
1309	Yes	File delete operation failed.
1310	Yes	There was an error executing the pre-replication script.
1311	Yes	There was an error executing the post-replication script.
1312	Yes	Synchronization failed because the replication agent could not be installed.
1313	Yes	The delete operation failed because the replication agent could not be installed.
1314	Yes	The server could not queue a delete job for the agent.
1315	Yes	The replication agent on a specified server is not responding to a delete operation request.
1316	Yes	Synchronization timed out waiting for the replication agent to accept the job.
1317	Yes	The delete operation timed out waiting for the replication agent to accept the job.
1318	Optional	Synchronization was canceled by the user.

TABLE 112 Event IDs 1300 to 1399

Event ID	NT Event Log Entry	Description
1319	Optional	The delete operation was canceled by the user.
1320	Optional	Synchronization was completed with warnings.
1321	Optional	The delete operation was completed with warnings.
1322	Yes	Synchronization failed. Differential replication was required but two Replication Agents could not be found.
1323	Optional	Change Journal could not be used due to journal discontinuity.
1324	Optional	Change Journal could not be used because this is the first time the job ran.
1325	Optional	Change Journal could not be used because it is not active and auto-enabling is not permitted.
1326	Yes	Change Journal could not be used due to an unexpected failure.
1327	Optional	Change Journal could not be used because the feature is not available on this host.
1328	Optional	Change Journal could not be used because the replication is not being executed from the source host.

1400 - 1499 Server Event IDs

TABLE 113 Event IDs 1400 to 1499

Event ID	NT Event Log Entry	Description
1400	Optional	The application snapshot was created.
1401	Optional	The application snapshot failed to be created at the time specified.
1402	Optional	Policy schedule was changed.
1403	Optional	The policy alternate schedule was changed.
1404	Optional	The name resolution is ambiguous. This is usually caused by inconsistencies between the forward and reverse DNS lookups.

1500 - 1699 SnapMirror Event IDs

TABLE 114 Event IDs 1500 to 1699

Event ID	NT Event Log Entry	Description
1500	Optional	A SnapMirror on the filer was enabled.
1501	Optional	A SnapMirror on the filer was disabled.
1502	Optional	A SnapMirror on the filer was quiesced.
1503	Optional	A SnapMirror on the filer was resumed.
1504	Optional	A SnapMirror on the filer was created.

TABLE 114 Event IDs 1500 to 1699

Event ID	NT Event Log Entry	Description
1505	Optional	A SnapMirror on the filer was initialized.
1506	Optional	A SnapMirror on the filer was broken off.
1507	Optional	A SnapMirror on the filer was updated.
1508	Optional	A SnapMirror on the filer was aborted.
1509	Optional	A SnapMirror on the filer was resynchronized.
1510	Optional	A SnapMirror on the filer was deleted.
1511	Optional	The qtree on the volume for the filer was created.
1512	Optional	The security on the qtree on the volume for a filer was changed.
1513	Optional	The oplocks on the qtree on the volume for the file were changed.
1514	Optional	The qtree on the volume for the filer was removed.
1515	Optional	The volume for the filer was set to offline status.
1516	Optional	The volume for the filer was set to online status.
1517	Optional	The volume for the filer was set to restricted status.
1518	Optional	The qtree on the volume for the filer was modified.
1519	Optional	A SnapMirror on the filer was modified.
1520	Optional	The snapshot schedule was enabled for the volume on the filer.
1521	Optional	The snapshot schedule was changed for the volume on the filer. Old and new values: • Number of weekly snapshots to keep • Number of daily snapshots to keep • Number of hourly snapshots to keep • Hourly schedule
1522	Optional	The snapshot reserve was changed for the volume on the filer.
1523	Optional	The snapshot was created for the volume on the filer.
1524	Optional	The snapshot was renamed for the volume on the filer.
1525	Optional	The snapshot was deleted for the volume on the filer.
1526	Optional	The snapshot schedule was disabled for the volume on the filer.
1527	Optional	Snapshot was restored to its original location in the active file system.
1528	Optional	Snapshot was not restored to its original location in the active file system.
1529	Optional	Snapshot was restored to <location>.
1530	Optional	Snapshot was not restored to <location>.
1531	Optional	Snapshot was restored to volume.
1532	Optional	Snapshot was not restored to volume.
1533	Optional	Starting restore of snapshot files.
1534	Optional	Completed restore of snapshot files.
1535	Optional	Failed to restore snapshot files.

TABLE 114 Event IDs 1500 to 1699

Event ID	NT Event Log Entry	Description
1600 (obsolete)	Optional	A SnapMirror policy was created.
1601 (obsolete)	Optional	A SnapMirror policy was deleted.
1602 (obsolete)	Optional	A SnapMirror policy was renamed.
1603 (obsolete)	Optional	A SnapMirror policy was changed.

1700 - 1799 Namespace Availability Event IDs

TABLE 115 Event IDs 1700 to 1799

Event ID	NT Event Log Entry	Description
1700	Optional	A namespace availability policy was created.
1701	Optional	A namespace availability policy was deleted.
1702	Optional	A namespace availability policy was renamed.
1703	Optional	A namespace availability policy was changed.
1704	Optional	The monitoring schedule for a namespace availability policy was changed. The policy shows the old and new schedule.
1705	Optional	The schedule for retrying failed attempts at monitoring the namespace availability policy was changed. The policy show the old and new schedule.
1706	Optional	The roots for a namespace availability policy were changed.
1707	Optional	Evaluation of a namespace availability policy was started by the user.
1708	Optional	The options for a root synchronization policy were changed. The policy shows the old and new Auto Sync values.
1709	Optional	The status was changed in a namespace availability policy.
1710	Optional	The status was changed for a root in a namespace availability policy.
1711	Optional	The monitoring agent for a namespace availability policy was changed. The policy shows the old and new monitoring agent.
1712	Optional	Synchronization started for a namespace availability policy. This indicates whether synchronization was automatic or manual (forced synchronization).
1713	Optional	Synchronization was completed successfully for a namespace availability policy.
1714	Yes	Synchronization scan failed for a namespace availability policy.
1715	Optional	Synchronization was started for a namespace availability policy.
1716	Optional	Synchronization was completed successfully for a namespace availability policy.
1717	Yes	Synchronization failed for a namespace availability policy.
1718	Yes	Synchronization was completed with errors for a namespace availability policy.

1800 - 1899 Multiprotocol Namespace Event IDs

TABLE 116 Event IDs 1800 to 1899

Event ID	NT Event Log Entry	Description
1800	Optional	The multiprotocol namespace policy was created.
1801	Optional	The multiprotocol namespace policy was deleted.
1802	Optional	The multiprotocol namespace policy was renamed.
1803	Optional	The multiprotocol namespace policy was modified.
1804	Optional	The schedule for the multiprotocol namespace policy was changed.
1805	Optional	The monitoring agent for the multiprotocol namespace policy was changed.
1806	Optional	The multiprotocol namespace policy was started.
1807	Optional	The multiprotocol namespace policy was completed.
1808	Yes	The multiprotocol namespace policy failed to be completed.
1809	Optional	Sync point in the multiprotocol namespace policy was started.
1810	Optional	Sync point in the multiprotocol namespace policy was completed.
1811	Yes	Sync point in the multiprotocol namespace policy failed to be completed.
1812	Optional	The Sync Point was updated in the multiprotocol namespace policy.
1813	Optional	Sync Point was removed from the multiprotocol namespace policy.
1814	Optional	The share security defaults are updated for the multiprotocol namespace policy.
1815	Optional	The export and mount defaults are updated for the multiprotocol namespace policy.
1816	Optional	Evaluation of the multiprotocol namespace policy was started by the user.
1819	Yes	Error executing pre-sync script.
1820	Yes	Error executing post-sync script.

1900 - 1999 Service Event IDs

TABLE 117 Event IDs 1900 to 1999

Event ID	NT Event Log Entry	Description
1900	Yes	Successfully installed a service.
1901	Yes	Failed to install a service.
1902	Yes	Successfully upgraded a service.
1903	Yes	Failed to upgrade a service.
1904	Yes	Successfully uninstalled a service.
1905	Yes	Failed to uninstall a service.
1906	Yes	A service was started.
1907	Yes	A service was stopped.

TABLE 117 Event IDs 1900 to 1999

Event ID	NT Event Log Entry	Description
1908	Yes	A service failed.
1909	Yes	There was an exception in a service.
1910	Yes	The service failed to start up.

2000 - 2099 Migration Event IDs

TABLE 118 Event IDs 2000 to 2099

Event ID	NT Event Log Entry	Description
2000	Optional	A data migration policy was created.
2001	Optional	A data migration policy was deleted.
2002	Optional	A data migration policy was renamed.
2003	Optional	A DFS link was added to the data migration policy.
2004	Optional	The schedule was changed for the data migration policy.
2005	Optional	Evaluation of the data migration policy was started by the user.
2006	Optional	Evaluation of the data migration policy task was started by the user.
2007	Optional	Options for the data migration policy were changed.
2008	Optional	The replication options for the data migration policy were changed.
2009	Optional	A task was added to the data migration policy.
2010	Optional	The monitoring agent for the data migration policy was changed.
2011	Optional	A task was removed from the data migration policy.
2012	Yes	A batch file running as part of a migration task failed.
2013	Yes	A migration task error occurred while updating DFS.
2014	Optional	The evaluation of a data migration task was suspended by the user.
2015	Optional	The task creation batch file for the migration policy was changed.
2016	Yes	The batch file is missing for the task in the migration policy.
2017	Yes	The task encountered errors while being added to the migration policy.
2018	Optional	The monitoring agent for the migration policy was changed. Because the existing monitoring agent was unknown (usually an indication that the agent was uninstalled), some historical data about existing migration tasks might be lost.
2019	Optional	There was a failure updating automount files for the task in the migration policy.
2020	Optional	The automount file list for the migration policy was changed.
2021	Optional	The Namespace list for the migration policy was changed.
2022	Optional	The task scripting options for the migration policy were changed.
2023	Optional	The task finalization options for the migration policy were changed.

TABLE 118 Event IDs 2000 to 2099

Event ID	NT Event Log Entry	Description
2024	Optional	Evaluation of the migration policy was started.
2025	Optional	Evaluation of the migration policy was cancelled.

2100 - 2199 Namespace Creation Event IDs

TABLE 119 Event IDs 2100 to 2199

Event ID	NT Event Log Entry	Description
2100	Optional	A namespace creation policy was created.
2101	Optional	A namespace creation policy was deleted.
2102	Optional	A namespace creation policy was renamed.
2103	Optional	A namespace creation policy was changed.
2106	Optional	The roots for the namespace creation policy were changed.
2107	Optional	Evaluation of the namespace creation policy was started by the user.
2108	Optional	The options on the namespace creation policy were changed.
2109	Optional	The status of the namespace creation policy was changed.
2110	Optional	The status of a root in the namespace creation policy was changed.
2111	Optional	The monitoring agent for the namespace creation policy was changed.
2112	Optional	A scan was started for the namespace creation policy.
2113	Optional	A scan succeeded for the namespace creation policy.
2114	Yes	A scan failed for the namespace creation policy.
2115	Optional	A namespace modification started for the namespace creation policy.
2116	Optional	A namespace modification succeeded for the namespace creation policy.
2117	Yes	A namespace modification failed for the namespace creation policy.
2118	Optional	The search list for domains and servers was changed for a namespace creation policy.
2119	Optional	The Ignore list was changed for a namespace creation policy.
2120	Yes	Synchronization was completed with errors for a namespace creation policy.
2121	Optional	The search type for the namespace creation policy was changed.
2122	Optional	The namespace hierarchy template for the namespace creation policy was changed.
2123	Optional	The folder search string for the namespace creation policy was changed.
2124	Optional	The **Add Link Option** for the namespace creation policy was changed.
2125	Optional	The scanning process for the namespace creation policy found the items listed. The policy could not suggest a link to add because the **Add Link Options** is enabled and the policy scan has already suggested links to add.

TABLE 119 Event IDs 2100 to 2199

Event ID	NT Event Log Entry	Description
2126	Yes	Scanning process for the namespace creation policy encountered errors. The policy scan for the permission-based namespace policy encountered errors when it retrieved permission information.
2127	Optional	The options for the namespace creation policy changed.
2128	Optional	The scan for the namespace creation policy encountered warnings because we are unable to create links using custom hierarchy template for the items specified.

2200 - 2299 Snapshot Access Event IDs

TABLE 120 Event IDs 2200 to 2299

Event ID	NT Event Log Entry	Description
2200	Optional	The snapshot access policy was created.
2201	Optional	The snapshot access policy was deleted.
2202	Optional	The snapshot access policy was renamed.
2203	Optional	The snapshot access policy was modified.
2204	Optional	The monitoring agent for the snapshot access policy was changed.
2205	Optional	A manually created snapshot was added to the Snapshot Access policy.
2206	Yes	The snapshot access policy failed.
2207	Optional	The snapshot access policy finished.
2208	Yes	The snapshot access policy failed for the reasons stated.
2209	Optional	An evaluation of the snapshot access policy was started by the user.
2210	Optional	The snapshot access policy was canceled.
2211	Optional	The snapshot access policy was started.

2300 - 2399 Filer Disaster Recovery Event IDs

TABLE 121 Event IDs 2300 to 2399

Event ID	NT Event Log Entry	Description
2300	Optional	Recovery policy was created.
2301	Optional	Recovery policy was deleted.
2302	Optional	Recovery policy was renamed.
2303	Optional	Recovery policy was changed.
2304	Optional	The resource monitoring schedule for the recovery policy was changed.

TABLE 121 Event IDs 2300 to 2399

Event ID	NT Event Log Entry	Description
2305	Optional	The schedule for retrying the failed resource monitoring policy for a recovery policy was changed.
2306	Optional	The options for a recovery policy were changed. Old Values: • Destination • Auto Failover • Break On Failover New Values: • Destination • Auto Failover • Break On Failover
2307	Optional	Status was changed for a recovery policy.
2308	Optional	The resource monitor options for a recovery policy were changed. Old Values: • Find Links • Find QTrees • Automatically SnapMirror New Values: • Find Links • Find QTrees • Automatically SnapMirror
2309	Optional	The Ignore list was changed for a recovery policy.
2310	Optional	The SnapMirror was changed for a recovery policy.
2311	Optional	The Root list was changed for a recovery policy.
2312	Optional	The evaluation of the filer recovery policy was started by the user.
2313	Optional	The policy was removed from the filer recovery policy.
2314	Yes	There was an error breaking the SnapMirror while adding a target to a link for a filer recovery policy.
2315	Yes	There was an error adding a share while adding a target to a link for a filer recovery policy.
2316	Yes	There was an error resynchronizing the SnapMirror while adding a target to a link for a filer recovery policy.
2317	Yes	There was an error adding a target to a link for a filer recovery policy.
2318	Optional	The fail count for the filer recovery policy was changed.
2319	Optional	The User Decision Script settings for the filer recovery policy were changed.
2320	Optional	Failure was detected for the resource on the filer of the filer recovery link, but no failover is needed because the fail count has not been reached.
2321	Optional	The resource on the filer needs to be failed over for the filer recovery policy because the fail count has been exceeded.
2322	Optional	There was an error executing the User Decision Script for the filer recovery policy.
2323	Optional	The User Decision Script is missing for the filer recovery policy.
2324	Optional	The resource on the filer needs to be failed over for the filer recovery policy because the User Decision Script returned an error code.

TABLE 121 Event IDs 2300 to 2399

Event ID	NT Event Log Entry	Description
2325	Optional	An internal error occurred executing the User Decision Script for the filer recovery policy.
2326	Optional	Evaluation of Recovery policy was started.
2327	Optional	Resource monitoring of Recovery policy was started.
2328	Optional	Evaluation of Recovery policy was cancelled.
2329	Optional	Resource monitoring of Recovery policy was cancelled.

2400 - 2499 Snapshot Scheduling Event IDs

TABLE 122 Event IDs 2400 to 2499

Event ID	NT Event Log Entry	Description
2400	Optional	Snapshot scheduling policy was created.
2401	Optional	Snapshot scheduling policy was deleted.
2402	Optional	Snapshot scheduling policy was renamed.
2403	Optional	The schedule for the snapshot schedule policy was changed.
2404	Optional	The options for the snapshot scheduling policy were changed. Values: • Number of VFM-initiated snapshots to keep • Batch before snapshot • Batch after snapshot
2405	Optional	The options for the snapshot scheduling policy were changed. Values: • Manage filer snapshot reserve • Manage filer snapshot schedule • Enable scheduled snapshots • Number of weekly snapshots to keep • Number of daily snapshots to keep • Number of hourly snapshots to keep • Hourly schedule
2406	Optional	The monitoring agent for the snapshot scheduling policy was changed.
2407	Optional	The volume list for the snapshot scheduling policy was changed.
2408	Optional	Evaluation of the snapshot scheduling policy was started.
2409	Optional	Evaluation of the snapshot scheduling policy was cancelled.

2500 - 2599 Namespace Backup Event IDs

TABLE 123 Event IDs 2500 to 2599

Event ID	NT Event Log Entry	Description
2500	Optional	The namespace backup policy was created.
2501	Optional	The namespace backup policy was deleted.
2502	Optional	The namespace backup policy was renamed.
2503	Optional	The namespace backup policy was changed.
2504	Optional	The monitoring agent for the namespace backup policy was changed.
2505	Optional	The monitoring schedule for the namespace backup policy was changed.
2506	Optional	The roots for the namespace backup policy were changed.
2507	Optional	The hourly backup schedule for the namespace backup policy was changed.
2508	Optional	The daily backup schedule for the namespace backup policy was changed.
2509	Optional	The weekly backup schedule for the namespace backup policy was changed.
2510	No	The number of hourly backups to keep for the namespace backup policy was changed.
2511	Optional	The number of daily backups to keep for the namespace backup policy was changed.
2512	Optional	The number of weekly backups to keep for the namespace backup policy was changed.
2513	Optional	The namespace backup policy was started.
2514	Optional	The namespace backup policy was canceled.
2515	Optional	The namespace backup policy was completed successfully.
2516	Yes	The namespace backup policy failed.
2517	Optional	Manual namespace backup was started.
2518	Optional	Manual namespace backup was canceled.
2519	Optional	Manual namespace backup was completed successfully.
2521	Yes	The namespace backup policy failed.
2522	Yes	Manual namespace backup failed for the reasons specified.
2523	Optional	The namespace backup policy finished.
2524	Optional	Manual namespace backup finished.

2600 - 2699 Logical Namespace Event IDs

TABLE 124 Event IDs 2600 to 2699

Event ID	NT Event Log Entry	Description
2600	Optional	The DFS root was created.
2601	Optional	The DFS root was deleted.
2602	Optional	DFS root was restored from the backup.

TABLE 124 Event IDs 2600 to 2699

Event ID	NT Event Log Entry	Description
2603	Optional	Replica was added to the DFS root.
2604	Optional	Replica was removed from the DFS root.
2605	Optional	Comment for the DFS root was changed.
2606	Optional	Timeout was changed for the DFS root.
2607	Optional	Monitoring for the DFS root was changed.
2608	Optional	Logical folder was added to the DFS root.
2609	Optional	Logical folder was deleted from the DFS root.
2610	Optional	Logical folder was renamed on the DFS root.
2611	Optional	Link was added to the DFS root.
2612	Optional	Link was deleted from the DFS root.
2613	Optional	Link was renamed on the DFS root.
2614	Optional	Storage was added to a link in the root.
2615	Optional	Storage was removed from the link.
2616	Optional	Timeout was changed for the DFS link.
2617	Optional	Comment was changed for the DFS link.
2618	Optional	Storage state was changed on the link.
2619	Optional	The share was created on the server.
2620	Optional	The share was deleted from the server.
2621	Optional	The share was renamed on the server.
2622	Optional	The share was created on the server based on a share.
2623	Optional	Directory was created.
2624	Optional	Directory was cleaned.
2625	Optional	Share comment was changed.
2626	Optional	Maximum connections value was changed on the share.
2627	Optional	Directory was renamed in <pathname>.
2628	Optional	Restoration from back to namespace root was completed with an error.
2629	Optional	Referral properties have changed for namespace root.
2630	Optional	Referral properties have changed for namespace link.
2631	Optional	Root replica target priority has changed for namespace root.
2632	Optional	Link storage target priority has changed for namespace link.
2633	Optional	Namespace root management permission has changed for namespace root.
2634	Optional	Snapshot has changed for namespace link.

2700 - 2799 Archival Migration Event IDs

TABLE 125 Event IDs 2700 to 2799

Event ID	NT Event Log Entry	Description
2700	Optional	Archival migration policy was created.
2701	Optional	Archival migration policy was deleted.
2702	Optional	Archival migration policy was renamed.
2703	Optional	The DFS link was added to the archival migration policy.
2704	Optional	The replication schedule or the archival migration policy was changed.
2705	Optional	Evaluation of the archival migration policy was started by the user.
2706	Optional	Evaluation of the archival migration task was started by the user.
2707	Optional	The options for the archival migration policy were changed.
2708	Optional	The replication options for the archival migration policy were changed.
2709	Optional	The task was added to the archival migration policy.
2710	Optional	The monitoring agent for the archival migration policy was changed.
2711	Optional	The task was removed from the archival migration policy.
2712	Yes	The batch file failed with a specified exit code for a task in the archival migration policy.
2713	Yes	The task failed to update DFS links in the archival migration policy.
2714	Optional	The evaluation of the archival migration policy task was suspended by the user.
2715	Optional	The scan schedule for the archival migration policy was changed.
2716	Yes	The batch file is missing for the task in the archival migration policy.
2717	Optional	The script settings for the archival migration policy were changed.
2718	Optional	The scan policy settings for the archival migration policy were changed.
2719	Optional	The archive rules for the archival migration policy were changed.
2720	Optional	The list of resources in the archival migration policy to scan were changed.
2721	Optional	The candidates were added to the exclusion list for the archival migration policy.
2722	Optional	The scan for movable resources in the archival migration policy ran successfully.
2723	Yes	The scan for movable resources in the archival migration policy encountered problems.
2724	Yes	The task encountered errors while being added to the archival migration policy.
2725	Optional	The monitoring agent for the archival migration policy was changed. Because the existing monitoring agent was unknown (usually an indication that the agent was uninstalled), some historical data about existing migration tasks might be lost.
2726	Optional	There was a failure updating the automount file for the task in the archival migration policy.
2727	Optional	The policy scan evaluation of the archival migration policy was started by the user.
2728	Optional	The policy migration of the archival migration policy was started by the user.
2729	Optional	The automount file list for the archival migration policy was changed.
2730	Optional	The namespace list for the archival migration policy was changed.

TABLE 125 Event IDs 2700 to 2799

Event ID	NT Event Log Entry	Description
2731	Optional	The task-scripting options for the archival migration policy were changed.
2732	Optional	The task finalization options for the archival migration policy were changed.
2733	Optional	The evaluation of the archival migration policy was started.
2734	Optional	The evaluation of the archival migration policy was cancelled.

2800 - 2899 Server Disaster Recovery Event IDs

TABLE 126 Event IDs 2800 to 2899

Event ID	NT Event Log Entry	Description
2800	Optional	The server recovery policy was created.
2801	Optional	The server recovery policy was deleted.
2802	Optional	The server recovery policy was renamed.
2803	Optional	The DFS link was added to the server recovery policy.
2804	Optional	The DFS link was removed from the server recovery policy.
2805	Optional	The link target was added to the DFS link in the server recovery policy.
2806	Optional	The link target was removed from the DFS link in the server recovery policy.
2808	Optional	The status was changed for the DFS link in the server recovery policy.
2809	Optional	The replication schedule for the server recovery policy was changed.
2810	Optional	The evaluation of the server recovery policy was started by the user.
2811	Optional	Replication was started for the server recovery policy.
2812	Optional	Replication was completed successfully for the server recovery policy.
2813	Yes	Replication failed for the server recovery policy.
2814	Optional	The name was changed for the DFS link in the server recovery policy.
2815	Optional	The schedule for retrying failed monitoring policies for the server recovery policy was changed.
2816	Optional	The schedule for retrying failed replication jobs for the server recovery policy was changed.
2817	Optional	The options for the server recovery policy were changed.
2818	Optional	The failover script settings for the server recovery policy were changed.
2819	Optional	The monitoring agent for the server recovery policy was changed.
2820	Optional	The replication options for the server recovery policy were changed.
2821	Optional	There was an error executing a command for the server recovery policy.
2822	Optional	There are no targets for a link on the server recovery policy.
2823	Optional	There was an error getting information for the server disaster recovery link.
2824	Yes	More than one target is online for the server disaster recovery link.

TABLE 126 Event IDs 2800 to 2899

Event ID	NT Event Log Entry	Description
2825	Optional	No target is online for the server disaster recovery link.
2827	Yes	There was an error setting a target online for the server disaster recovery link.
2828	Yes	There was an error setting a target offline for the server disaster recovery link.
2829	Optional	The link needs to be manually failed over.
2830	Optional	There was an error executing a pre-failover command for the server recovery policy.
2831	Optional	There was an error executing a post-failover command for the server recovery policy.
2832	Optional	Replication was completed with warnings for the server recovery policy.
2833	Optional	Resource Monitoring Options were changed for the server recovery policy.
2836	Optional	The fail count for the server recovery policy was changed.
2837	Optional	The User Decision Script settings for the server recovery policy was changed.
2838	Optional	Resource monitoring for the server recovery policy removed an extra link target.
2839	Yes	Resource monitoring for the server recovery policy failed to remove an extra link target.
2840	Optional	Resource monitoring for the server recovery policy added a secondary link target.
2841	Yes	Resource monitoring for the server recovery policy failed to add a secondary link target.
2842	Optional	Removed the active link for the server recovery policy.
2843	Optional	Removed the nonactive link for the server recovery policy.
2844	Optional	Changed an existing nonactive link to an active link for the server recovery policy.
2845	Optional	Added a new active link for the server recovery policy.
2846	Optional	Added a new nonactive link for the server recovery policy.
2847	Optional	The server recovery policy could not determine the master link.
2848	Optional	The master link for the server recovery policy could not be reached.
2849	Optional	The master link for the server recovery policy is not in a proper state to monitor for failover.
2850	Optional	There was an error executing the User Decision Script for the server recovery policy.
2851	Optional	The link needs to be failed over for the server recovery policy because the User Decision Script returns the following message.
2852	Optional	The User Decision Script for the server recovery policy requested a failover for the link, but no failover is needed because fail count has not been reached.
2853	Optional	There was an internal error when the User Decision Script was executed for the server recovery policy.
2854	Optional	The User Decision Script is missing for the server recovery policy.
2855	Optional	The link target is unreachable for the server recovery policy, but no failover is needed because fail count has not been reached.
2856	Yes	The server recovery policy was unable to monitor the master link to determine whether failover was needed.
2857	Yes	The server recovery policy was unable to retrieve the active link information.
2858	Optional	The primary server target was activated for the link by the server recovery policy.

TABLE 126 Event IDs 2800 to 2899

Event ID	NT Event Log Entry	Description
2859	Optional	The secondary server target was activated for the link by the server recovery policy.
2860	Yes	The server recovery policy failed to activate the primary server target for the link.
2861	Yes	The server recovery policy failed to activate the secondary server target for the link.
2862	Optional	The link for the server recovery policy is not in a proper state for failover.
2863	Optional	An existing active link was changed to a nonactive link for the server recovery policy.
2864	Optional	The evaluation of the server recovery policy was started.
2865	Optional	The resource monitoring of the server recovery policy was started.
2866	Optional	The evaluation of the server recovery policy was cancelled.
2867	Optional	The resource monitoring of the server recovery policy was cancelled.

2900 - 2999 Replication and Monitoring Agents Event IDs

TABLE 127 Event IDs 2900 to 2999

Event ID	NT Event Log Entry	Description
2900	Optional	The agent installation started.
2901	Yes	The agent installation failed.
2902	Optional	The agent installation succeeded.
2910	Optional	The agent upgrade started.
2911	Yes	The agent upgrade failed.
2912	Optional	The agent upgrade succeeded.
2920	Optional	The uninstall of the agent started.
2921	Yes	The uninstall of the agent failed.
2922	Optional	The uninstall of the agent succeeded.
2923	Optional	The agent on the specified machine is installed but not running.
2924	Yes	A error occurred during agent deployment.
2925	Optional	A warning occurred during agent deployment.
2926	Optional	Agent deployment information.
2927	Optional	The agent selected for the job has an incompatible version number. Please check your agent reports and update those agents that do not match the server version.
2928	Yes	An error occurred opening the required registry key.
2929	Yes	An error occurred in querying required registry key value.
2930	Yes	Invalid expected value for registry value returned.

3100 - 3199 Storage Load Balancing Event IDs

TABLE 128 Event IDs 3100 to 3199

Event ID	NT Event Log Entry	Description
3100	Optional	Storage load balancing policy was created.
3101	Optional	Storage load balancing policy was deleted.
3102	Optional	Storage load balancing policy was renamed.
3103	Optional	Storage load balancing policy was changed.
3104	Optional	The monitoring agent for the storage load balancing policy was changed.
3105	Optional	The scan schedule for the storage load balancing policy was changed.
3106	Optional	The migration schedule for the storage load balancing policy was changed.
3107	Optional	The replication options for the storage load balancing policy were changed.
3108	Optional	The options for the storage load balancing policy were changed.
3109	Optional	The DFS Root List for the storage load balancing policy was changed.
3110	Optional	The Storage Pool List for the storage load balancing policy was changed.
3111	Optional	The policy scan for the storage load balancing policy ran successfully.
3112	Yes	The policy scan for the storage load balancing policy encountered problems.
3113	Yes	The task encountered errors while being added to the storage load balancing policy.
3114	Optional	The monitoring agent for the storage load balancing policy was changed because the existing monitoring agent was unknown (usually an indication that the agent has been uninstalled). Some historical data about existing migration tasks might have been lost.
3115	Optional	Evaluation of the storage load balancing policy was started by the user.
3116	Optional	Evaluation of the storage load balancing policy task was started by the user.
3117	Optional	The task was added to the storage load balancing policy.
3118	Optional	The task was removed from the storage load balancing policy.
3119	Optional	Evaluation of the storage load balancing policy task was suspended by the user.
3120	Optional	Failed updating DFS links for a task in the storage load balancing policy.
3121	Optional	The following data move suggestions are excluded for the storage load balancing policy.
3122	Optional	The policy scan for the storage load balancing policy encountered the following problems.
3123	Optional	Updating of automount files has failed for a task in the storage load balancing policy.
3124	Optional	The task-scripting options for the storage load balancing policy have changed.
3125	Optional	The task finalization options for the storage load balancing policy have changed.
3126	Optional	The task finalization options for the storage load balancing policy task have changed.
3127	Optional	Evaluation of the storage load balancing policy was started.
3128	Optional	Evaluation of the storage load balancing policy was cancelled.
3129	Optional	Migration was begun for storage load balancing policy.
3130	Optional	Migration was cancelled for storage load balancing policy.

3200 - 3299 Permission-Based Namespace Event IDs

TABLE 129 Event IDs 3200 to 3299

Event ID	NT Event Log Entry	Description
3200	Optional	Permission-based namespace policy was created.
3201	Optional	Permission-based namespace policy was deleted.
3202	Optional	Permission-based namespace policy was renamed.
3203	Optional	Permission-based namespace policy was changed.
3204	Optional	The monitoring schedule for the permission based namespace policy was changed.
3205	Optional	The alternate monitoring schedule for the permission based namespace policy was changed.
3206	Optional	The roots for permission based namespace policy was changed.
3207	Optional	Evaluation of the permission based namespace policy was started by the user.
3208	Optional	The options for the permission based namespace policy were changed.
3209	Optional	The status was changed for the permission based namespace policy.
3210	Optional	The status was changed for the root in the permission based namespace policy.
3211	Optional	The monitoring agent for the permission based namespace policy was changed.
3212	Optional	Scanning was started for the permission based namespace policy.
3213	Optional	Scanning completed successfully for the permission based namespace policy.
3214	Yes	Scanning failed for the permission based namespace policy.
3215	Optional	The modification of the namespace started for the permission based namespace policy.
3216	Optional	The modification of the namespace completed successfully for the permission based namespace policy.
3217	Yes	The modification of the namespace failed for the permission based namespace policy.
3218	Optional	The Domain/Server search list was changed for the permission based namespace policy.
3219	Optional	The Ignore list for the permission based namespace policy was changed.
3220	Yes	The modification of the namespace was completed with errors for the permission based namespace policy.
3221	Optional	The search type for the permission based namespace policy was changed.
3222	Optional	The namespace hierarchy template for the permission based namespace policy was changed.
3223	Optional	The folder search string for the permission based namespace policy was changed.
3224	Optional	The Users and Groups List for the permission based namespace policy was changed.
3225	Yes	The scanning process for the permission based namespace policy encountered an error when retrieving permission information.
3226	Optional	The **Add Link Option** for the permission based policy was changed.
3227	Optional	The scanning process for the permission based namespace policy found the listed items but the policy did not suggest links to add because the **Add Link Option** is enabled and the policy scan already suggested links to add.

TABLE 129 Event IDs 3200 to 3299

Event ID	NT Event Log Entry	Description
3228	Yes	The policy scan for the permission based namespace policy encountered errors when retrieving file system permission information for the items listed.
3229	Yes	The policy scan for the permission based namespace policy encountered errors when retrieving CIFS sharing permission information for the items listed.
3230	Yes	The policy scan for the permission based namespace policy encountered an error when searching.
3231	Optional	The options for the permission based namespace policy changed.
3232	Optional	The scan for the permission based namespace policy encountered warnings because we are unable to create links using custom hierarchy template for the items specified.