IBM System Storage and Virtual File Manager

Life cycle management with VFM and N series

Intelligent data movement with VFM and N series

Data replication using VFM and N series

Alex Osuna
Murillo Bernardes
Silvio Martins
Shri Seshadri
Shahid Shaikh

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM System Storage and Virtual File Manager

October 2008

SG24-7597-00

Note: Before using this information and the product it supports, read the information in "Notices" on page ix.

First Edition (October 2008)

This edition applies to Data ONTAP Version 7.2.2 and above.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
DFS™
IBM®
Redbooks®
Redbooks (logo) ®
System Storage™
Tivoli®
TotalStorage®

The following terms are trademarks of other companies:

AMD, AMD Opteron, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

Snapshot, SecureAdmin, RAID-DP, FlexShare, WAFL, VFM, SyncMirror, SnapVault, SnapValidator, SnapRestore, SnapMover, SnapMirror, SnapManager, SnapLock, SnapDrive, NearStore, FlexVol, FlexClone, FilerView, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Solaris, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Access, Active Directory, Internet Explorer, Microsoft, MS, Outlook, PowerPoint, SharePoint, SQL Server, Windows NT, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

"Microsoft product box shot(s) reprinted with permission from Microsoft Corporation."

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redbooks® publication describes how IBM System Storage™ N series Virtual File Manager (VFM®) logically aggregates user file data distributed across heterogeneous environments and provides administrators with tools and policies to automate data management.

VFM is designed to provide data management functionality for server and storage consolidation, migration, remote office data management, and disaster recovery while avoiding disruption to users. VFM provides this functionality through automated policy-based data management leveraging a global namespace.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Tucson Center.

Alex Osuna is a Project Leader at the International Technical Support Organization, Tucson Center. He writes extensively and teaches IBM classes nationwide on all areas of IBM System Storage N series. Before joining the ITSO 3 years ago, Alex worked as a IBM Tivoli® Principal Systems Engineer specializing in IBM Tivoli storage. Alex has over 30 years in the IT industry focused mainly on storage. He holds over 10 certifications from Redhat, Microsoft®, and IBM.

Murillo Bernardes is a Software Engineer with IBM Linux® Technology Center in Brazil. He has 4 years of experience developing software mainly on system management, Linux virtualization, and server consolidation. He has a LPI1 certification and holds a BSc in Electrical Engineering from UNICAMP, Campinas, Brazil.

Silvio Martins is an IT Specialist in IBM Global Services in Brazil. He has 5 years of experience in the IT infrastructure field. He holds several product certifications from Microsoft. His areas of expertise include Microsoft infrastructure environments support.

Shri Seshadri is Director of Product Marketing at Brocade and is focused on file solutions in which VFM is a key product. Before this, Sri was VP Technology at NuView, a company purchased by Brocade in 2006. Sri's responsibilities included product management and product evangelization of all file products in the field. Sri also has about 15 years of IT consulting experience.

Shahid Shaikh is a System Software Engineer at the India Systems and Technology Labs, Pune Center. He has been working on the AIX® implementation of NFSv4 security, RAS, and deployment features. He also has worked on SAN and other storage technologies. He has over 3 years of experience in the IT industry. He holds a bachelor's degree in computer science and engineering.

Thanks to the following people for their contributions to this project:

Todd Cummings
Brocade

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

`ibm.com/redbooks/residencies.html`

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  `ibm.com/redbooks`

- Send your comments in an e-mail to:

  `redbooks@us.ibm.com`

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

1

Introduction

This chapter discusses enterprise data management challenges. It also introduces IBM Virtual File Manager and how it solves these data management challenges. We also describe the benefits of using IBM System Storage N series Virtual File Manager (VFM) in an enterprise environment.

The explosive growth of unstructured data and the associated proliferation of file servers and network attached storage (NAS) appliances have resulted in acute management challenges for IT administrators. as well as deteriorating data access for clients (users and application servers). Because every device represents an independent storage resource governed by its embedded file system, scaling these storage resources can become a significant challenge.

As presentations, spreadsheets, and other user- and application-generated files accumulate, network storage requirements continue to grow. Storage administrators must determine how much storage is available and how much is in use. They add new storage, balance loads, reconfigure existing storage, migrate and replicate data, prepare for disaster recovery, and consolidate storage, all without affecting users or incurring downtime. In addition, to find and access files stored on multiple machines, users must know the files' locations, must be able to connect to each machine, and must use file shares and file server names that can be cryptic and difficult to remember.

These challenges are described in the list that follows:

- Storage consolidation:

  Adding new storage to the network, migrating from Windows® NT to Windows 2003, consolidating storage and servers, and dealing with related challenges.

- Business continuity:

  While ensuring business continuity, administrators prepare for disaster recovery, try to avoid or minimize client downtime, and ensure the availability of business-critical data all the time.

- Storage optimization:

  With the increase in the numbers of users and the increase in users' data, performance of the storage might be degraded with many users accessing one device. This leads to managing load balancing among multiple storage devices for optimum performance. Increasing hardware acquisition cost is also a part of storage optimization-related challenges.

- Information life cycle management:

  This category of challenges deals with archivals and backups of user data and aligning them with regulatory compliance requirements.

- Remote site support:

  While supporting remote sites, enterprises try to centralize management of highly distributed data, try to reduce the cost of remote site backups, and try to conserve resources required to support the remote site.

- Data classification and reporting:

  Discovering storage capacity utilization and determining the business value of data are challenges related to data classification and reporting.

To add to the previously mentioned challenges, users find it confusing to find the data they want to access. The mapping of too many drive letters adds to the confusion. IT administrators find it challenging in such environments to ensure user access during ongoing data migration, ensure user access to files in the event of disaster or system failure, and efficiently manage multiple file types distributed across many devices and locations.

Figure 1-1 describes enterprise data management challenges.

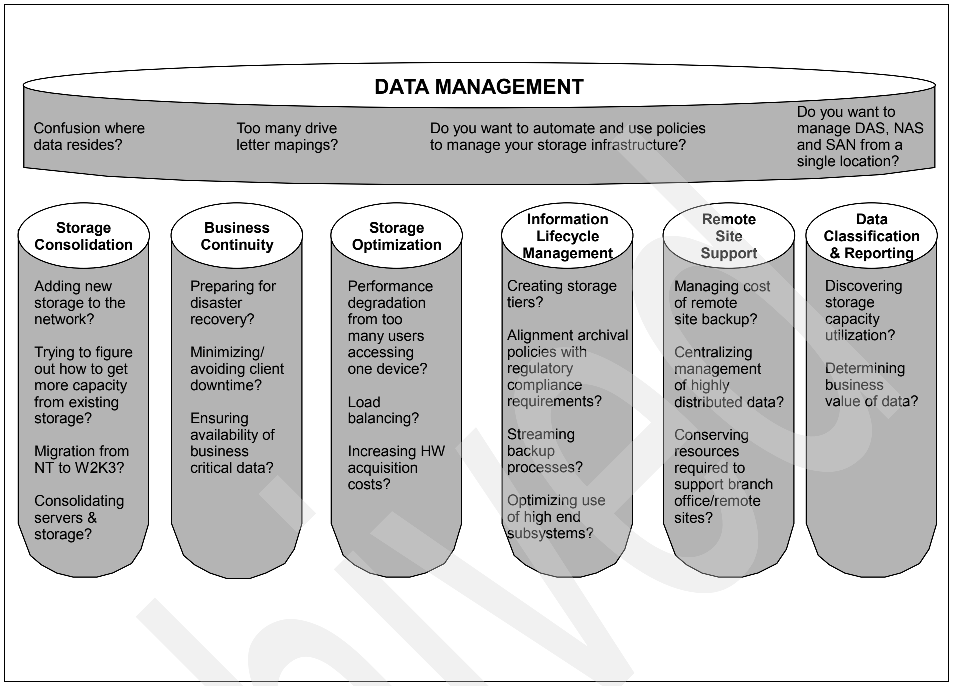
DATA MANAGEMENT
Confusion where data resides?
Too many drive letter mapings?
Do you want to automate and use policies to manage your storage infrastructure?
Do you want to manage DAS, NAS and SAN from a single location?
Storage Consolidation
Adding new storage to the network?
Trying to figure out how to get more capacity from existing storage?
Migration from NT to W2K3?
Consolidating servers & storage?
Business Continuity
Preparing for disaster recovery?
Minimizing/ avoiding client downtime?
Ensuring availability of business critical data?
Storage Optimization
Performance degradation from too many users accessing one device?
Load balancing?
Increasing HW acquisition costs?
Information Lifecycle Management
Creating storage tiers?
Alignment archival policies with regulatory compliance requirements?
Streaming backup processes?
Optimizing use of high end subsystems?
Remote Site Support
Managing cost of remote site backup?
Centralizing management of highly distributed data?
Conserving resources required to support branch office/remote sites?
Data Classification & Reporting
Discovering storage capacity utilization?
Determining business value of data?

Figure 1-1 Enterprise data management challenges

The solution to the problems presented in Figure 1-1 is file virtualization. File virtualization provides logical and location-independent views of distributed file storage across heterogeneous and geographically distributed storage devices.

Figure 1-2 depicts file virtualization.

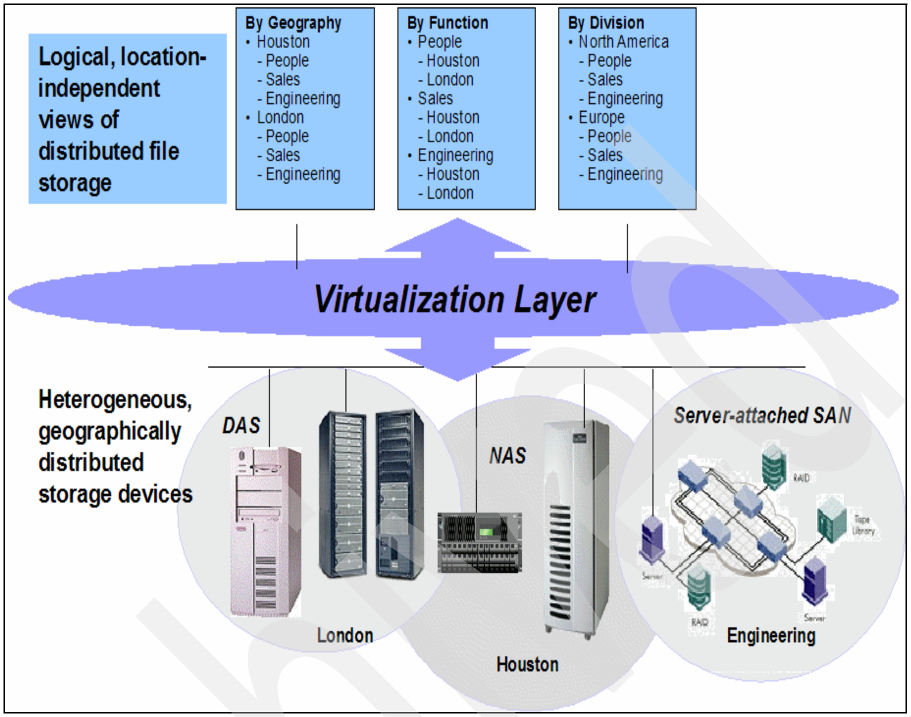

Figure 1-2 File virtualization

Users do not have to know where their data is coming from; for example, they do not have to know whether the data is on a storage area network (SAN), NAS, or direct-attached storage (DAS). They do not have to know whether the data is in Houston, London, or another location. When users travel from one place to another, their view of data is undisturbed. In addition, read-only data can be provided to different users from their respective closest server.

1.1 IBM Virtual File Manager

IBM System Storage N series Virtual File Manager (VFM) logically aggregates user file data distributed across heterogeneous environments and provides administrators with tools and policies to automate data management.

It is designed to provide data management functionality for server and storage consolidation, migration, remote office data management, and disaster recovery features, while avoiding disruption to users. It provides all this functionality through automated policy-based data management, leveraging a global namespace.

In Figure 1-3 we show different work folders for an enterprise; the work folders are distributed across heterogeneous hardware and geographies.

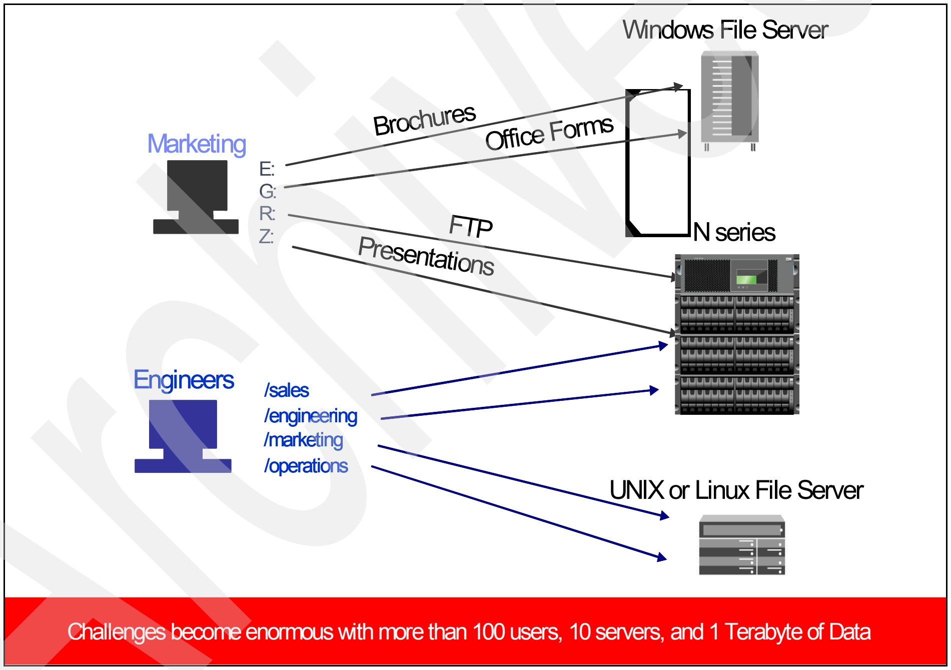

Figure 1-3 Without VFM

With VFM the configuration in Figure 1-3 on page 5 can be simplified. See Figure 1-4.

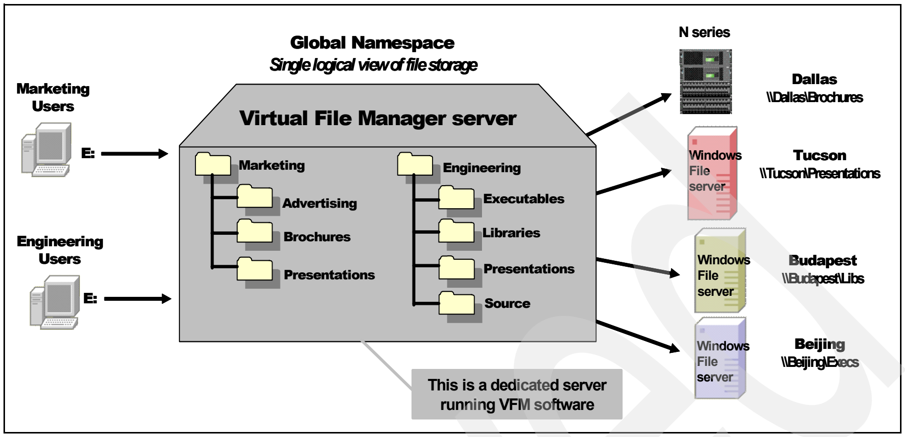

Figure 1-4 With VFM

With VFM users see a single logical view of file storage, making it easy to access files. VFM is an integrated solution that provides virtualized management of distributed files.

The VFM solution replaces many point products in an enterprise. See Figure 1-5.

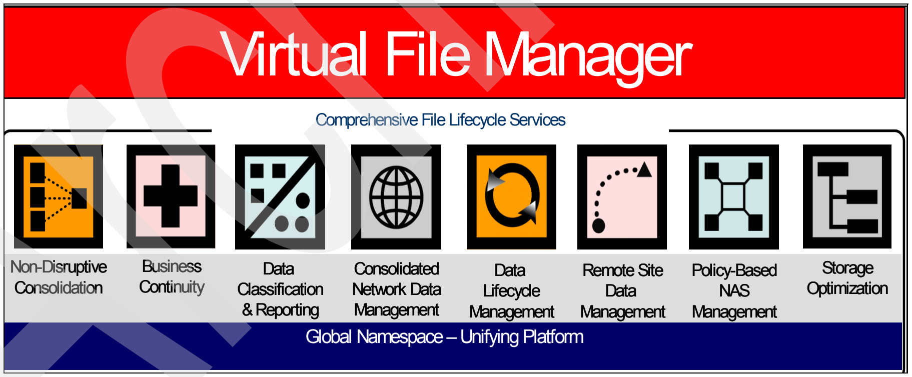
Virtual File Manager
Comprehensive File Lifecycle Services
Non-Disruptive Consolidation
Business Continuity
Data Classification & Reporting
Consolidated Network Data Management
Data Lifecycle Management
Remote Site Data Management
Policy-Based NAS Management
Storage Optimization
Global Namespace – Unifying Platform

Figure 1-5 VFM: an integrated solution

VFM is available in migration and enterprise editions. Figure 1-6 shows the difference between the VFM Migration Edition and the VFM Enterprise Edition.

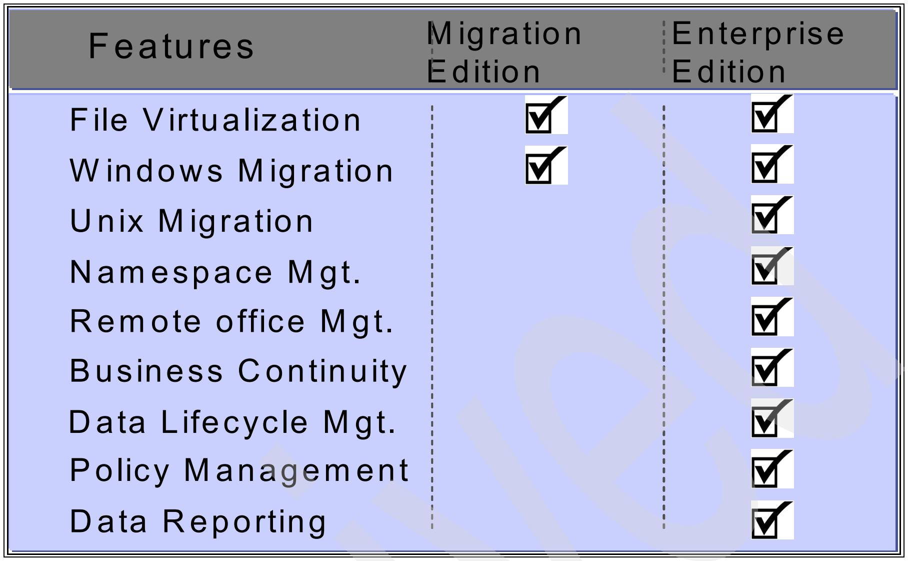

Features	Migration Edition	Enterprise Edition
File Virtualization	☑	☑
Windows Migration	☑	☑
Unix Migration		☑
Namespace Mgt.		☑
Remote office Mgt.		☑
Business Continuity		☑
Data Lifecycle Mgt.		☑
Policy Management		☑
Data Reporting		☑

Figure 1-6 VFM editions: capability matrix

1.2 Benefits of VFM

Using alternatively VFM benefits users in many ways. These benefits are described in the sections that follow.

Single view of entire file system

With VFM, the user has a simplified, comprehensive view of distributed data organized within a single logical directory (see Figure 1-7):

- Simple searches for files and data
- Intuitive drive and file names
- Changes to the physical data that do not affect the user

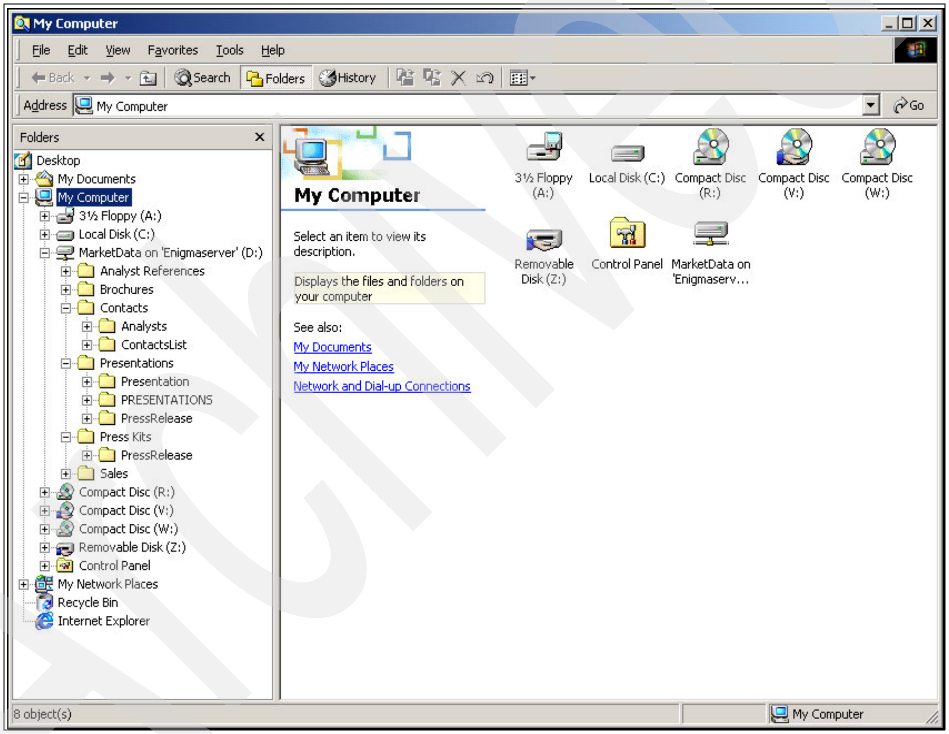

Select an item to view its description.

Displays the files and folders on your computer

Figure 1-7 Single view of entire file system

For example, suppose you are looking for a Microsoft Word document, and you do not know the drive letter to search.

File virtualization

VFM provides users with a location-independent method of accessing files. It also solves horizontal scaling of storage and drive letter mapping issues (see Figure 1-8). VFM can also keep the volume size in check within company-recommended maximum limits.

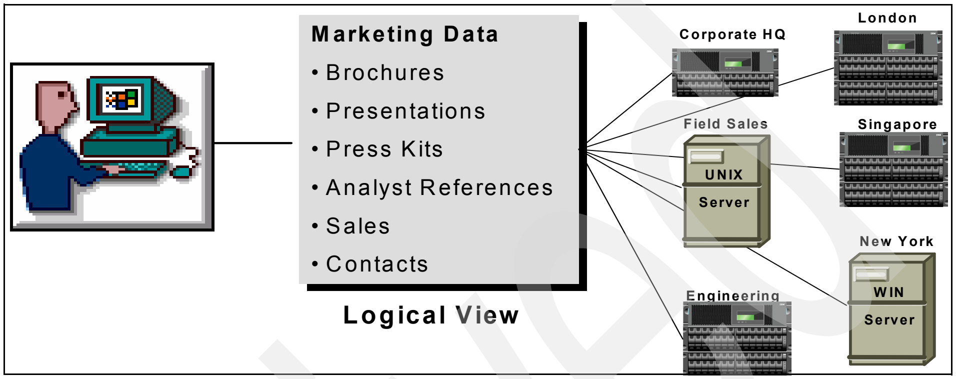

Figure 1-8 File virtualization

Transparent data migration

VFM makes it easy to move data and optimize existing storage without affecting users or applications. It also helps load balancing among multiple storage devices serving the same data (see Figure 1-9).

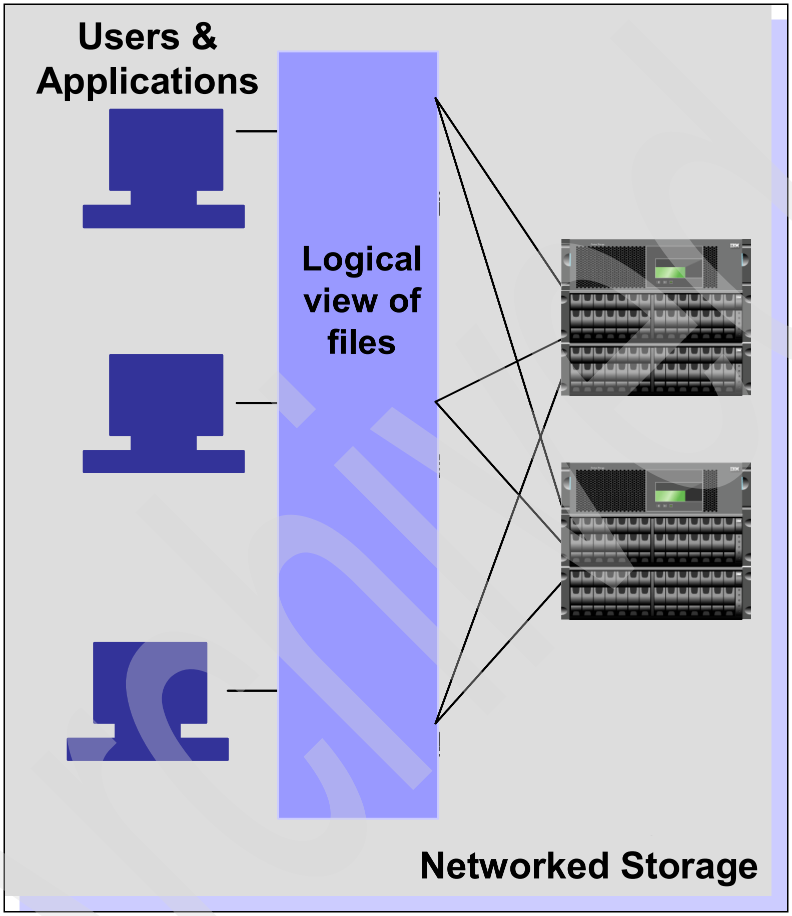

Figure 1-9 Transparent data migration

Higher availability

VFM provides rapid access to data after a disaster or storage reconfiguration, and it effectively reduces real and perceived downtime of clients (see Figure 1-10).

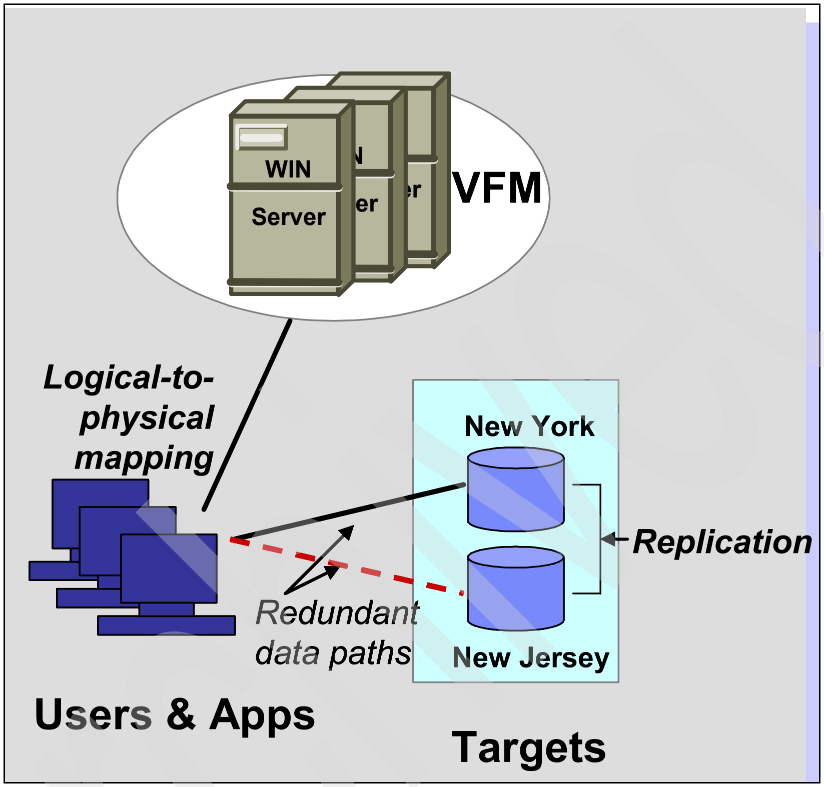

Figure 1-10 Higher availability

VFM is not in the data path between user and target. Thus all the data does not go through VFM, which prevents VFM from being a bottleneck.

Transparent data life cycle management

VFM provides sophisticated data life cycle management:

- Automated classification of distributed files
- Support of heterogeneous storage tiers
- Transparent data access
- Policy-based seamless creation of tiers of storage

Data classification and reporting

VFM assists IT administrators by providing an inventory of assets, reports, and so on (see Figure 1-11):

- Provides agent or agentless reporting
- Provides reporting mechanisms for reporting by individual server and across the enterprise
- Allows reuse of collected data for various reports
- Enables proactive decision making

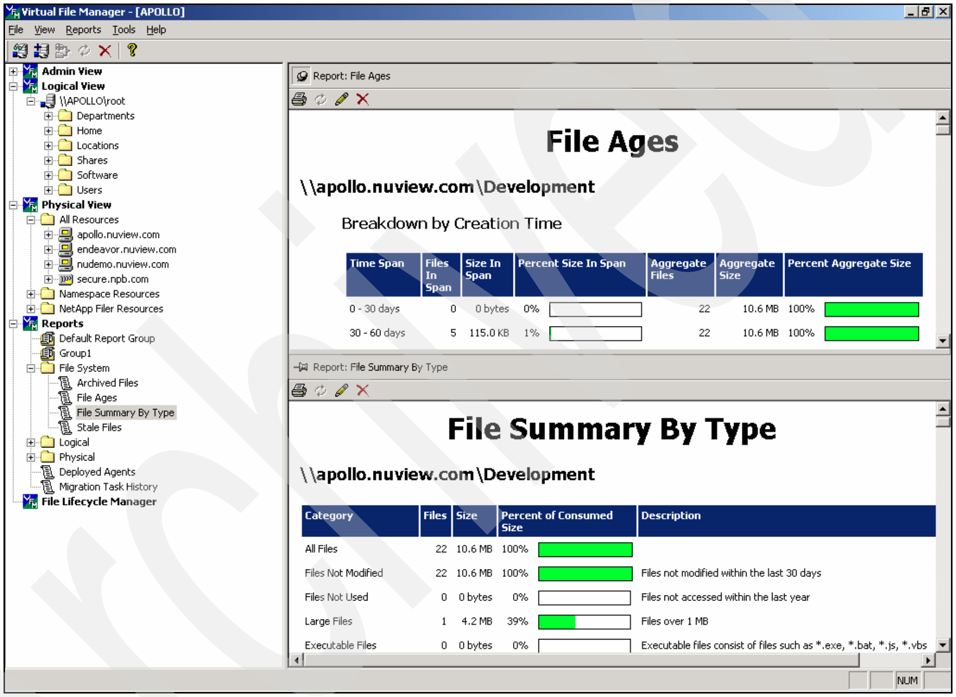

File Ages

\\apollo.nuview.com\Development

Breakdown by Creation Time

Time Span	Files In Span	Size In Span	Percent Size In Span	Aggregate Files	Aggregate Size	Percent Aggregate Size
0 - 30 days	0	0 bytes	0%	22	10.6 MB	100%
30 - 60 days	5	115.0 KB	1%	22	10.6 MB	100%

File Summary By Type

\\apollo.nuview.com\Development

Category	Files	Size	Percent of Consumed Size	Description
All Files	22	10.6 MB	100%	
Files Not Modified	22	10.6 MB	100%	Files not modified within the last 30 days
Files Not Used	0	0 bytes	0%	Files not accessed within the last year
Large Files	1	4.2 MB	39%	Files over 1 MB
Executable Files	0	0 bytes	0%	Executable files consist of files such as *.exe, *.bat, *.js, *.vbs

Figure 1-11 Data classification and reporting

Intelligent data movement

VFM provides policy-based migration and replication. Migration policies can be customized in a number of different ways for intelligent migration (see Figure 1-12).

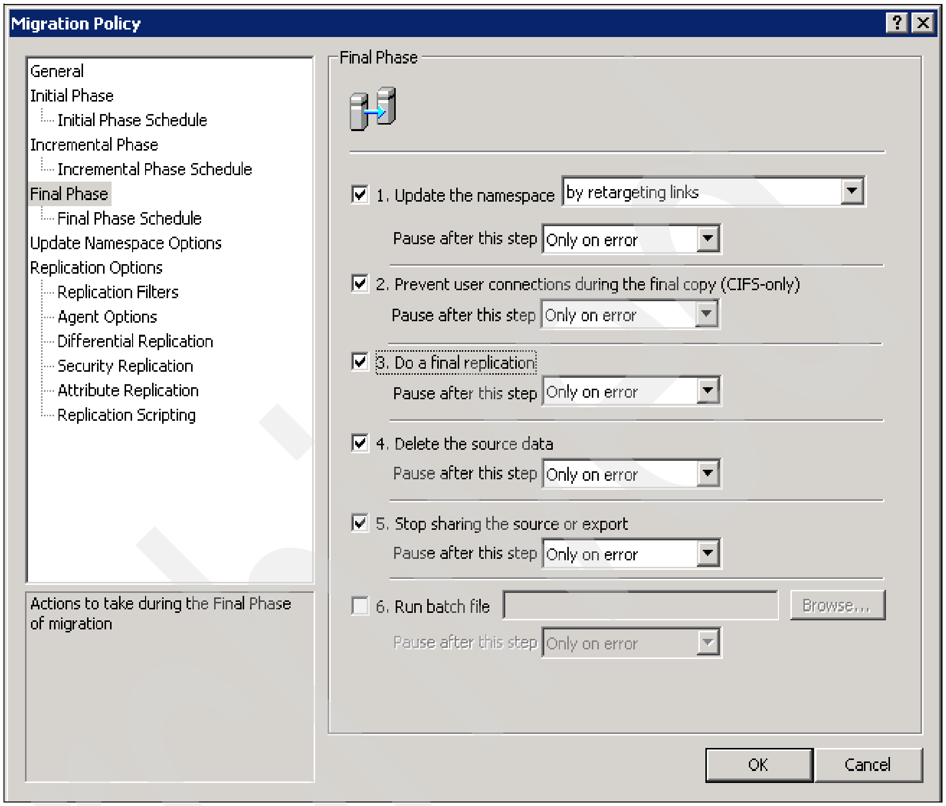

Migration Policy

- General
- Initial Phase
  - Initial Phase Schedule
- Incremental Phase
  - Incremental Phase Schedule
- Final Phase
  - Final Phase Schedule
- Update Namespace Options
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Actions to take during the Final Phase of migration

Final Phase

☑ 1. Update the namespace: by retargeting links

Pause after this step: Only on error

☑ 2. Prevent user connections during the final copy (CIFS-only)

Pause after this step: Only on error

☑ 3. Do a final replication

Pause after this step: Only on error

☑ 4. Delete the source data

Pause after this step: Only on error

☑ 5. Stop sharing the source or export

Pause after this step: Only on error

☐ 6. Run batch file

Pause after this step: Only on error

Figure 1-12 Policy-based migration

Policy-based migration provides greater control of data replication for business continuity processes and for synchronizing data (see Figure 1-13).

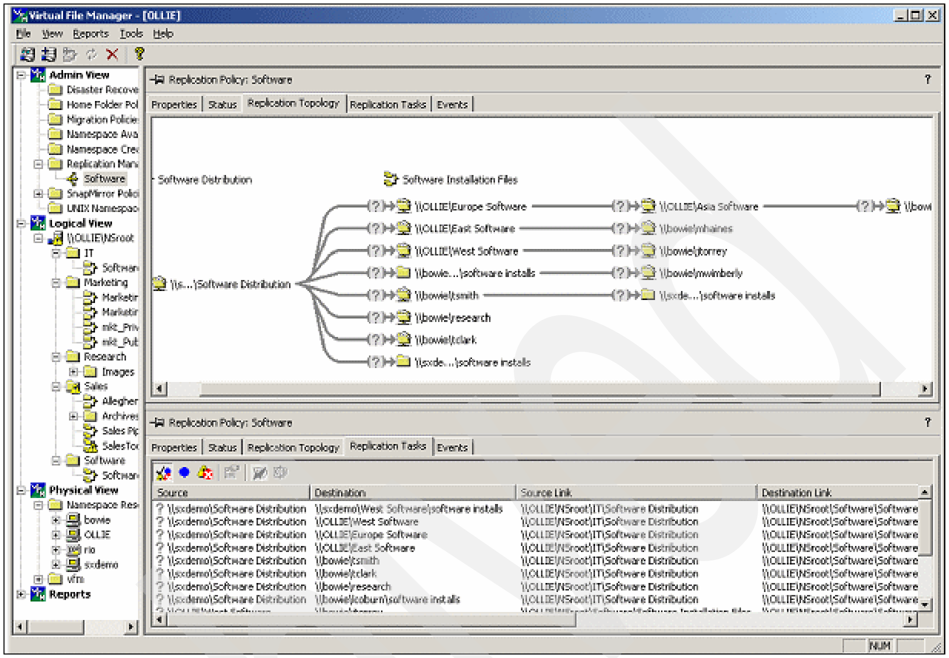

Replication Policy: Software

Source	Destination	Source Link	Destination Link
? \\sxdemo\Software Distribution	\\sxdemo\West Software\software installs	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software
? \\sxdemo\Software Distribution	\\OLLIE\West Software	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software
? \\sxdemo\Software Distribution	\\OLLIE\Europe Software	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software
? \\sxdemo\Software Distribution	\\OLLIE\East Software	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software
? \\sxdemo\Software Distribution	\\bowie\tsmith	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software
? \\sxdemo\Software Distribution	\\bowie\tclark	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software
? \\sxdemo\Software Distribution	\\bowie\research	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software
? \\sxdemo\Software Distribution	\\bowie\coburn\software installs	\\OLLIE\NSroot\IT\Software Distribution	\\OLLIE\NSroot\Software\Software

Figure 1-13 Policy-based replication

2

Introduction to IBM System Storage N series

In this chapter we introduce the IBM System Storage N series and describe the hardware and software.

The IBM System Storage N series is designed from the ground up as a standalone storage system. It provides a range of reliable, scalable storage solutions for a variety of storage requirements. These capabilities are achieved by using network access protocols such as NFS, Common Internet File System (CIFS), HTTP, and iSCSI as well as storage area technologies such as Fibre Channel. Utilizing built-in RAID technologies, IBM System Storage N series protects all data with options to add additional protection through mirroring, replication, snapshots, and backup. These storage systems are also characterized by simple management interfaces that make installation, administrating, and troubleshooting uncomplicated and straightforward.

Using this type of flexible storage solution includes the following advantages:

- Tune the storage environment to a specific application while maintaining the flexibility to increase, decrease, or change access methods with a minimum of disruption.
- React easily and quickly to changing storage requirements. If additional storage is required, you must be able to expand it quickly and with minimum disruption. When storage exists but is deployed incorrectly, the capability to

reallocate available storage from one application to another quickly and simply cannot be done.

- Maintain availability and productivity during upgrades. If outages are required, IBM System Storage N series keeps them to a minimum.
- Create effortless backup and recovery solutions that operate commonly across all data access methods.
- Provide file- and block-level services in a single system, helping to simply your infrastructure.
- Tune the storage environment to a specific application while maintaining its availability and flexibility.
- Change the deployment of storage resource nondisruptively, easily, and quickly. Online storage resource redeployment is possible.
- Provide easy and quick upgrade process. A nondisruptive upgrade is possible.
- Provide strong data-protection solutions and support online backup and recovery.

2.1 IBM N series hardware

In the following sections, we discuss the N series models available today (see Table 2-1 and Table 2-2 on page 18).

- N3000 series
- N5000 series
- N7000 series

Table 2-1 N series storage systems

A20 models	Max # of drives	Max capacity in TB
N3700	56	16.8
N3300	68	68
N3600	104	104
N5200	168	84
N5300	336	336
N5500	336	168
N5600	504	504
N7600	840	840
N7700	840	840
N7800	1008	1008
N7900	1176	1176

Table 2-2 IBM N series Gateway models

Use	Model	Capacity
Midrange	N5200 G10 &G20 	84TB
Midrange	N5300 G10 & G20 	336TB
Midrange	N5500 G10&G20 	80TB
Midrange	N5600 G10&G20 	504TB
Enterprise class	N7600 G10&G20 	840TB
Enterprise class	N7700 G10 & G20 	840TB
Enterprise class	N7800 G10 & G20 	1008TB
Enterprise class	N7900 G10 & G20 	1176TB

2.2 Comparing N series Gateway to N series storage systems

The following list compares the N series Gateway to the N series storage system. The two systems are identical as follows:

- Core NAS features and functionality
- iSCSI features and functionality
- FCP features and functionality
- Filer SAN host support matrix
- Behavior for the Write Anywhere File System (WAFL®)
- Data availability characteristics
- Data integrity characteristics
- Data management characteristics
- Serviceability characteristics

The N series Gateway can be further compared to the N series as follows:

- Each supports the same version of Data ONTAP®.
- The N5000 and N7000 series Gateway physical attributes are the same as the N5000 and N7000 models A10 and A20 storage systems.
- Differences exist in system initialization and storage expansion.
- The N series Gateway does not use the following features of Data ONTAP:
  - SnapLock® compliance
  - Nearstore option
- N series storage systems use disk storage provided by IBM only; the Gateway models support heterogeneous storage and IBM expansion units.
- Data ONTAP is enhanced to enable the Gateway series solution.
- A RAID array from a separate storage system can provide logical unit numbers (LUNs) to the Gateway:
  - Each LUN is equivalent to an IBM disk.
  - LUNs are assembled into aggregates and volumes, and then they are formatted with the WAFL file system, just like the IBM N series storage systems.

2.2.1 IBM N series A models hardware quick reference

Table 2-3 provides a hardware quick reference to the IBM N series A models.

Table 2-3 A models hardware quick reference

Function	N3700	N3300	N3600	N5200	N5500	N5300	N5600	N7600	N7700	N7800	N7900
Maximum raw capacity in TB A10 models	16	68	104	84	168	336	420	672	840	672	1176
Maximum raw capacity in TB A20 models	16	68	104	84	168	336	504	840	840	1008	1176
Fibre Channel disk drives	144 GB 10K RPM, 144 GB 15K RPM, 300 GB 10K RPM - EXN2000 144 GB 15K, 300 GB 10K - EXN4000										
SATA disk drives	250 GB 7.2K RPM., 320 GB 7.2K RPM., 500 GB 7.2K RPM, 750 GB 7.2K RPM, 1TB 7.2K RPM										

Function	N3700	N3300	N3600	N5200	N5500	N5300	N5600	N7600	N7700	N7800	N7900

Maximum number of disks	56	68	104	168	336	336	420 (A10)
							504 (A20)
								672(A10)
								840(A20)
									840	672 (A10)
										1008 (A20)
											1176

Expansion units supported	EXN1000 (SATA), EXN2000 (FC), EXN4000 4 Gbps (FC)										

Table 2-4 provides a quick reference to the N series G models.

Table 2-4 N series G models quick reference

Function	N5200	N5300	N5500	N5600	N7600	N7700	N7800	N7900
Maximum raw capacity in TB G10 models	84	336	84	504	672	840	672	1176
Maximum raw capacity in TB G20 models	84	336	84	504	840	840	1008	1176
Max number of LUNs on back-end disk storage array	168	252	336	420 for A10 and 504 for A20	840	840	1008	1176
Max LUN size in GB	1000	1000	500	1000	1000	1000	1000	1000
Maximum volume size in TB	16	16	16	16	16	16	16	16

2.2.2 IBM N series A and G models hardware quick reference

Table 2-5 provides a quick reference to the storage system.

Table 2-5 Storage system reference

Function	N3700	N3300	N3600	N5200	N5300	N5500	N5600	N7600	N7700	N7800	N7900
Network protocol support	NFS V2/V3/V4 over UDP or TCP,PCNFSD V1/V2 for (PC) NFS client authentication, Microsoft CIFS, iSCSI, FCP, VLD, HTTP 1.0, HTTP1.1 Virtual Host										
Other protocol support	SNMP, NDMP, LDAP, NIS, DNS										
Onboard I/O ports per node	2 x GbE 2 x Optical FC	4 x GbE 4 x Optical FC	4 x GbE 4 x Optical FC	4 x GbE 4 x FC 1 x LVD SCSI	4 x GbE 4 x FC 1 x LVD SCSI	4 x GbE 4 x FC 1 x LVD SCSI	4 x GbE 4 x FC (4 Gbps)	6 x GbE 8 x FC	6 x GbE 8 x FC	6 x GbE 8 x FC	6 x GbE 8 x FC
PCI expansion slots per node	N/A	N/A	1 x PCI-E	3 x PCI-X	3 x PCI-E	3 x PCI-X	3 x PCI-E	5 x PCI-E 3 x PCI-X	3 x PCI-E 6 X PCI-E	5 x PCI-E , 3 x PCI-X	3 x PCI-E 6 x PCI-E
NVRAM in MB per node	128	128	256	512	512	512	512	1024	1024	4096	4096
Memory in GB per node	1	1	2	2	4	4	8	16	32	32	64
Redundancy/ high availability	CompactFlash, dual-redundant hot-plug integrated cooling fans, hot-swappable autoranging power supplies, clustered storage controllers, hot-swappable disk bays										
Required rack space	3U	2U	4U	3U per node	3U per node	3U per node	3U per node	6U per node	6U per node	6U per node	6U per node
Processors (A10)	Two Broadcom MIPS-based	2.2 GHz 64-bit processors	2 2.2 GHz 64-bit processors	One 2.8 GHz Xeon	Two 1.8 GHz AMD ™	Two 2.8 GHz Xeon	Two AMD 1.8 GHz dual-core	Two 2.6 GHz AMD Opteron	Two 2.6 GHz AMD Opteron ™	Four 2.6 GHz AMD Opteron	Four 2.6 GHz AMD Opteron

Function	N3700	N3300	N3600	N5200	N5300	N5500	N5600	N7600	N7700	N7800	N7900
Processors (A20)	Four Broadcom MIPS-based	Two 2.2 GHz 64-bit processors	Four 2.2 GHz 64-bit processors	Two 2.8 GHz Xeon	Four 1.8 GHz AMD	Four 2.8 GHz Xeon	Four AMD 1.8 GHz dual-core	Four 2.6 GHz AMD Opteron	Four 2.6 GHz AMD Opteron	Eight 2.6 GHz AMD Opteron	Eight 2.6 GHz AMD Opteron

2.3 IBM N series standard software features

Table 2-6 lists licensed no-charge features available with the IBM N series.

Table 2-6 Standard software features

Data ONTAP	Operating system software optimizes data serving and allows multiple protocol data access.
FTP	File Transfer Protocol (FTP), a standard Internet protocol, is a simple way to exchange files between computers on the Internet.
Telnet	The Telnet protocol provides a general, bidirectional, eight-bit byte-oriented communications facility. It provides user-oriented command-line logon sessions between hosts.
Snapshot™	Enables online backups, providing near instantaneous access to previous versions of data without requiring complete, separate copies.
FlexVol®	FlexVol creates multiple flexible volumes on a large pool of disks. FlexVol provides dynamic, nondisruptive (thin) storage provisioning along with space- and time-efficiency. These flexible volumes can span multiple physical volumes without regard to size.
FlexShare™	FlexShare gives administrators the ability to leverage existing infrastructure and increase processing utilization without sacrificing the performance of critical business needs. With the use of FlexShare, administrators can confidently consolidate different applications and data sets on a single storage system. FlexShare gives administrators the control to prioritize applications based on how critical they are to the business.
Disk sanitization	Disk sanitization is the process of physically obliterating data by overwriting disks with specified byte patterns or random data in a manner that prevents recovery of current data by any known recovery methods. This feature enables you to carry out disk sanitization by using three successive byte overwrite patterns per cycle. By default six cycles are performed.
FilerView®	This Web-based administration tool enables IT administrators to fully manage N series storage systems from remote locations. It provides simple and intuitive Web-based single-appliance administration.

SnapMover®	SnapMover migrates data among N series clusters with no impact on data availability and no disruption to users.
AutoSupport	AutoSupport is a sophisticated, event-driven logging agent featured in the Data ONTAP operating software and inside each N series system, which continuously monitors the health of your system and issues alerts if a problem is detected. These alerts can also be in the form of e-mail.
SecureAdmin™	SecureAdmin is a Data ONTAP module that enables authenticated, command-based administrative sessions between an administrative user and Data ONTAP over an intranet or the Internet.
DNS	The N series supports using a host-naming file or a specified DNS server and domain.
Cluster	-Ensures high data availability for business-critical requirements by eliminating a single point of failure. -Must be ordered for A20 clustered configurations or upgrades from A10 to A20. -Active-active pairing delivers even more “nines to right of the decimal point.”
NIS	The N series does provide NIS client support and can participate in NIS domain authentication.
Integrated automatic RAID manager	The IBM N series and Data ONTAP provide integrated RAID management with RAID-Double Parity (default) and RAID 4.
iSCSI Host Attach Kit for AIX, Windows, Linux	A Host Support Kit includes support software and documentation for connecting a supported host to an iSCSI network. The support software includes programs that display information about storage, and programs to collect information needed by client support to diagnose problems.
NearStore® option	The NearStore option is a disk-based, secondary storage device for enterprise applications.
Advanced single instance storage	This functionality is designed to significantly improve physical storage efficiency and network efficiency by enabling the sharing of duplicate data blocks.

2.4 Optional software

Table 2-7 lists optional software available with the IBM N series.

Table 2-7 Optional software

CIFS	Provides file system access for Microsoft Windows environments.
NFS	Provides file system access for UNIX® and Linux environments.
HTTP	Hypertext Transfer Protocol enables a user to transfer displayable Web pages and related files.
FlexClone®	Designed to provide instant replication of data volumes/sets without requiring additional storage space at the time of creation.
Multistore	-Permits an enterprise to consolidate a large number of Windows, Linux, or UNIX file servers onto a single storage system. -Many “virtual filers” on one physical appliance ease migration and multidomain failover scenarios.
SnapLock	Provides non-erasable and non-rewritable data protection that helps enable compliance with government and industry records retention regulations.
SnapMirror®	-Remote mirroring software that provides automatic block-level incremental file system replication between sites. -Available in synchronous, asynchronous and semi synchronous modes of operation
SnapRestore®	Allows rapid restoration of the file system to an earlier point in time, typically in only a few seconds.
SnapVault®	Provide disk-based backup for N series systems by periodically backing up a snapshot copy to another system.
SnapDrive®	SnapDrive enables Windows and Unix applications to access storage resources on N series storage systems, which are presented to the Windows 2000 or later, operation system as locally attached disks. On UNIX it enables you to create storage on a storage system in the form of LUNs, file systems, logical volumes, or disk groups.
SnapManager®	Host software for managing Exchange, SQL Server®, and SAP® backup and restore. SnapManager software simplifies exchange data protection by automating processes to provide hands-off, worry-free data management. A new integrated GUI-based administration SnapManager Console has been introduced for Microsoft Office SharePoint® Server. It unifies the entire stack from applications to servers to storage for backup, recovery, and extraction of critical information.
SnapValidator®	For Oracle® deployments, SnapValidator can be used to provide an additional layer of integrity checking between the application and N series storage. SnapValidator allows Oracle to create checksums on data transmitted to N series storage for writes to disk and include the checksum as part of the transmission.

SyncMirror®	SyncMirror is synchronous mirror of a volume. It maintains a strict physical separation between the two copies of your mirrored data. In case of an error in one copy, the data is still accessible without any manual intervention.
Single Mailbox Recovery (SMBR) for Exchange	SMBR is a software option from SnapManager that is designed to take near-instantaneous online backups of Exchange databases, verify that the backups are consistent, and rapidly recover Exchange within levels: storage group, database, folder, single mailbox, or single message. The potential results are improved service to internal clients, reduced infrastructure expenses, and significant time savings for Exchange administrators.
Operations Manager	Operations Manager provides remote, centralized management of IBM N series data storage infrastructure, including global enterprise and storage network.
MetroCluster	MetroCluster software provides an enterprise solution for high availability over wide area networks.
Virtual File Manager (VFM)	IBM System Storage N series Virtual File Manager (VFM) software is a comprehensive solution for managing unstructured file data. It is designed to provide data management functionality for server and storage consolidation, migration, remote office data management, and disaster recovery features while avoiding disruption to users. It provides all this functionality through automated policy-based data management leveraging a global namespace.

2.5 Software quick reference

Table 2-8 provides a software quick reference.

Table 2-8 Software quick reference

Product/ feature/ function	**Included/ optional**	**N3000 A10 & A20**	**N5X00 A10 & A20**	**N5x00 G10 & G20**	**N7x00 A10 & A20**	**N7x00 G10 & G20**
Data ONTAP	Included	X	X	X	X	X
iSCSI protocol	Included	X	X	X	X	X
FTP protocol	Included	X	X	X	X	X
NDMP protocol	Included	X	X	X	X	X
FlexVol	Included	X	X	X	X	X
Snapshot	Included	X	X	X	X	X
SecureAdmin	Included	X	X	X	X	X

Product/ feature/ function	Included/ optional	N3000 A10 & A20	N5X00 A10 & A20	N5x00 G10 & G20	N7x00 A10 & A20	N7x00 G10 & G20
iSCSI Host Attach Kit for AIX, Windows, Linux	Included	X	X	X	X	X
FlexShare	Included				X	X
SnapMover	Included	X	X	X	X	X
CIFS protocol	Optional	X	X	X	X	X
NFS protocol	Optional	X	X	X	X	X
HTTP protocol	Optional	X	X	X	X	X
FCP protocol	Optional	X	X	X	X	X
FlexClone	Optional	X	X	X	X	X
Clustered failover	Optional	X(A20)	X(A20)	X(G20)	X(A20)	X(G20)
Multistore	Optional	X	X	X	X	X
SnapMirror	Optional	X			X requires special HBA card	X requires special HBA card
SnapRestore	Optional	X	X	X	X	X
Open Systems SnapVault (OSSV)	Optional	X	X	X	X	X
SnapVault	Optional	X	X	X	X	X
SnapDrive for Windows & UNIX: AIX, Solaris™, HP-UX, Linux	Optional	X	X	X	X	X
SnapValidator	Optional	X	X	X	X	X
SyncMirror	Optional		X	X	X	X
SnapManager for SQL Server	Optional	X	X	X	X	X

Product/ feature/ function	Included/ optional	N3000 A10 & A20	N5X00 A10 & A20	N5x00 G10 & G20	N7x00 A10 & A20	N7x00 G10 & G20
SnapManager for SAP	Optional					
SnapManager for Exchange	Optional	X	X	X	X	X
Single Mailbox Recovery (SMBR) for Exchange	Optional	X	X	X	X	X
Operations Manager Core, BC & SRM license	Optional	X	X	X	X	X
SnapLock Enterprise	Optional	X	X	X	X	X
MetroCluster A2X models only	Optional		X	X	X	X
Disk sanitization	Included	X	X		X	X
SnapLock compliance	Optional	X	X		X	
NearStore option bundle	Included	X	X	X	X	X
RAID 4, RAID-DP™	Included	X	X		X	
Advanced Single Instance Storage	Included		X	X	X	X
VFM	Optional	X	X	X	X	X

2.6 IBM N series storage systems A models

The A models of the IBM N series storage systems offer multiprotocol connectivity using internal storage or storage provided by expansion units (see Figure 2-1 on page 29). The IBM System Storage N series systems are designed to provide integrated block- and file-level data access, allowing concurrent operation in IP SAN (iSCSI), FC SAN, NFS, and CIFS environments. Other storage vendors might require the operation of multiple systems to provide this functionality. IBM N series systems are designed to avoid costly downtime, both planned and unplanned, and improve your access to mission-critical data, thereby helping you gain a competitive advantage.

The N series A models are a specialized, “thin server” storage system with a customized operating system, similar to a stripped down UNIX kernel and hereafter referred to as Data ONTAP. With a reduced operating system, many of the server operating system functions you are familiar with are not supported. The objective is to improve performance and reduce costs by eliminating unnecessary functions normally found in standard operating systems.

The N series comes with preconfigured software and hardware, and with no monitor or keyboard for user access, which is commonly termed a *headless* system. A storage administrator accesses the systems and manages the disk resources from a remote console using a Web browser or command line.

One of the typical characteristics of the N series product is its ability to be installed rapidly using minimal time and effort to configure the system. It is integrated seamlessly into the network. This approach makes the IBM N series product especially attractive when lack of time and skills are elements in the decision process.

The IBM N series A models are depicted in Figure 2-1.

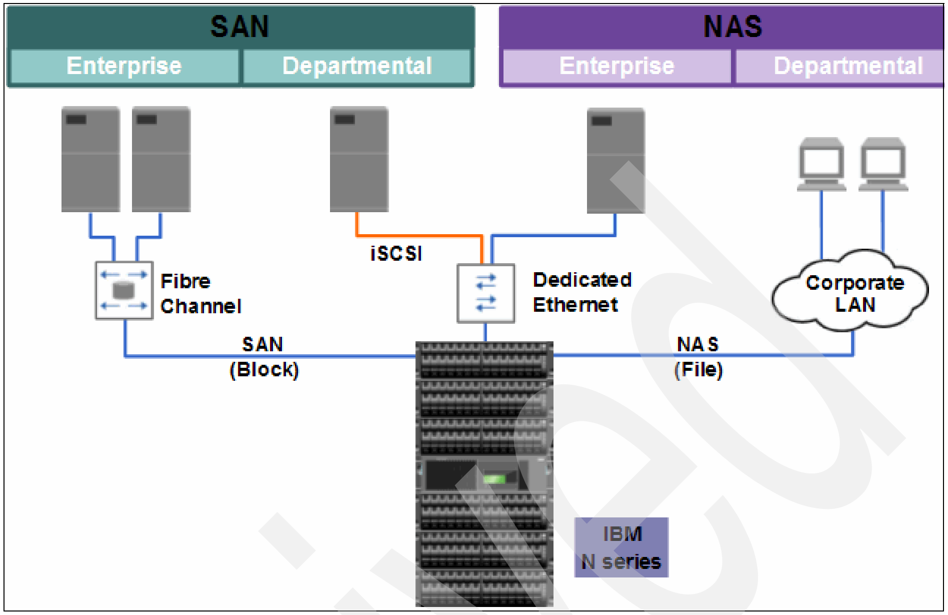

Figure 2-1 IBM N series A models

Drive flexibility

The IBM System Storage N series product is designed to provide network-attached storage in environments where clients must utilize their storage investment in a multifaceted environment. The IBM System Storage N series storage systems provide a tremendous amount of versatility by enabling this solution to be populated with both Fibre Channel disk drives and Serial Advanced Technology Attachment (SATA) disk drives. An N series populated with Fibre Channel disk drives can be suitable for mission-critical high-performance data transaction environments, whereas an N series populated with SATA disk drives can be attractive to clients who wish the use the platform for various scenarios, such as disk-to-disk backup, disaster recovery, data archive, or data in home directories that do not require high-performance transactional environments.

Table 2-9 provides guidance on environments where Fibre Channel, SAS, or SATA drives would be best suited.

Table 2-9 Drive positioning

Requirement	Fibre Channel drives	SAS drives	SATA drives
Online, high-performance, mission-critical production data repository	X	X	
Near-line storage used for tiered storage or infrequently accessed data	X	X	X
Data retention to help meet the needs of clients required to store data in non-erasable and non-rewritable (WORM) formats		X	X

Near-line storage

Two years ago the concept of near-line storage in the middle for disk staging was introduced. This enables organizations to do daily backups to disk and weekly or biweekly backups to tape, which reduces the amount of data that must be written to tape. Also, data is online for faster recovery. Another advantage near-line storage provides is you can leverage your existing investment in primary storage, your backup application, and tape libraries.

The IBM N series with SATA drives offers near-line storage. Figure 2-2 is an example of traditional disk-based backup and recovery. On the left is the primary storage, which is characterized by higher cost and very fast performance. On the far right are archive targets that have traditionally been tape or optical jukeboxes with reduced access times to read and write data.

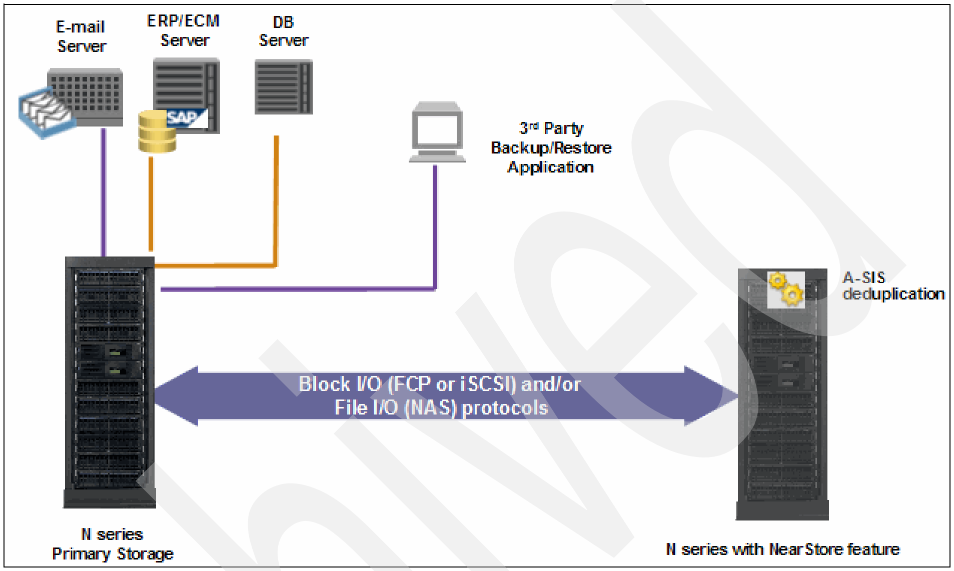
E-mail Server
ERP/ECM Server
SAP
DB Server
3rd Party Backup/Restore Application
A-SIS deduplication
Block I/O (FCP or iSCSI) and/or File I/O (NAS) protocols
N series Primary Storage
N series with NearStore feature

Figure 2-2   Near-line storage

2.6.1 IBM System Storage N3000 introduction

The N3000 systems are designed to provide primary and secondary storage for midsize enterprises. IT administrators can consolidate their fragmented application-based storage and unstructured data into one unified, easily managed, and expandable platform. N3000 systems offer integrated block- and file-level data access, intelligent management software, and data protection capabilities, such as higher-end N series systems, in a cost-effective package. N series innovations include Serial-Attached SCSI (SAS) drive support, expandable I/O connectivity, and onboard remote management.

The N3000 systems are designed as the entry point to the entire N series family. The systems provide the following key advantages:

- High availability leverages proven features including a high-performing and scalable operating system, data management software, and redundancy features.
- Backup and recovery features designed to support disk-based backup, with file- or application-level recovery with snapshot and SnapRestore software features.
- Simple replication and disaster recovery designed to provide an easy-to-deploy mirroring solution that is highly tolerant of WAN interruptions.
- Management simplicity self-diagnosing systems designed to enable on-the-fly provisioning.
- Versatile, single, integrated architecture designed to support concurrent block I/O and file serving over Ethernet and Fibre Channel SAN infrastructures.

The N3000 is compatible with the entire family of N series unified storage systems, which feature a comprehensive, top-to-bottom lineup of hardware and software designed to address a variety of possible deployment environments:

- N3700
  - 2863-A10 single filer
  - 2863-A20 clustered
- N3300
  - 2859-A10 single filer
  - 2859-A20 clustered
- N3600
  - 2862-A10 single filer
  - 2862-A20 clustered

The N3000 series supports Ethernet and Fibre Channel environments, enabling economical NAS, FC, and iSCSI deployments. The N3000 system functions as a “unification engine,” which is designed to enable you to simultaneously serve both file- and block-level data across a single or multiple networks. This functionality usually demands procedures that for some solutions require multiple, separately managed systems.

N3000 storage systems can offer significant advantages for distributed enterprises with remote and branch office sites. These organizations and others can leverage the SnapVault and SnapMirror software functions to implement a cost-effective data protection strategy by mirroring data back to a corporate data center.

No PCI adapter slots are on the N3300 and N3700 systems. No additional adapter options are supported for the N3300 and N3700 systems. One PCIe adapter slot is available per node on the N3600 storage system. For an A20 model, adapters must be added in pairs, one per node, so that both nodes are populated with one of the same type of PCIe adapter.

N3700

The N3700 storage system (see Figure 2-3) is a 3U solution designed to provide NAS and iSCSI functionality for entry to midrange environments. The basic N3700 offering is a single-node model A10, which is upgradeable to the dual-node model A20 and requires no additional rack space. The dual-node, clustered A20 is designed to support failover and failback functions to maximize reliability. The N3700 storage system can support 14 internal hot-plug disk drives with scalability provided through attachment to up to three expansion units, each with a maximum of 14 drives. The N3700 also has the capability to connect to a Fibre Channel tape for backup.

A list of supported Tape drives can be found at;

http://www.ibm.com/totalstorage/nas

Refer to the IBM System Storage and TotalStorage® N series interoperability matrix.

http://www-03.ibm.com/systems/storage/nas/interophome.html

Figure 2-3 N3700

The type of controller defines the model. Figure 2-4 shows a single control unit. The single node A10 uses a single control unit, and the dual-node clustered A20 uses two control units (see Figure 2-4).

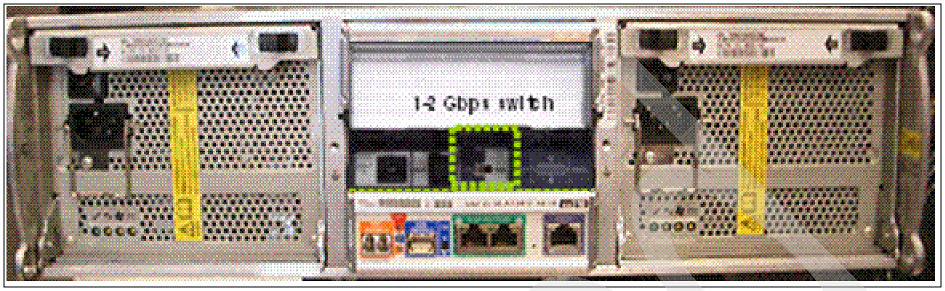

Figure 2-4 N3700 A10

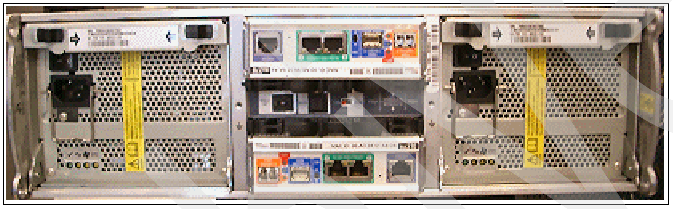

Figure 2-5 N3700 A20

The N3700 comes with redundant power supplies for higher reliability (see Figure 2-6).

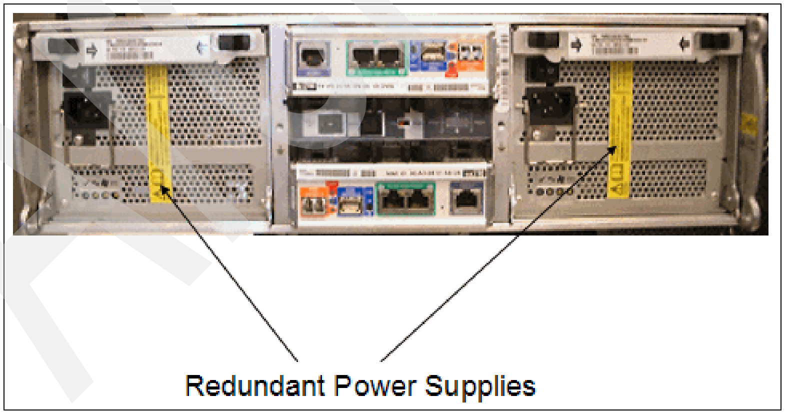

Figure 2-6 Redundant power supplies

From the rear of the N3700, you can see the diagnostic and operational LEDs on the power supply (Figure 2-7). Table 2-10 lists the LEDs and possible configurations.

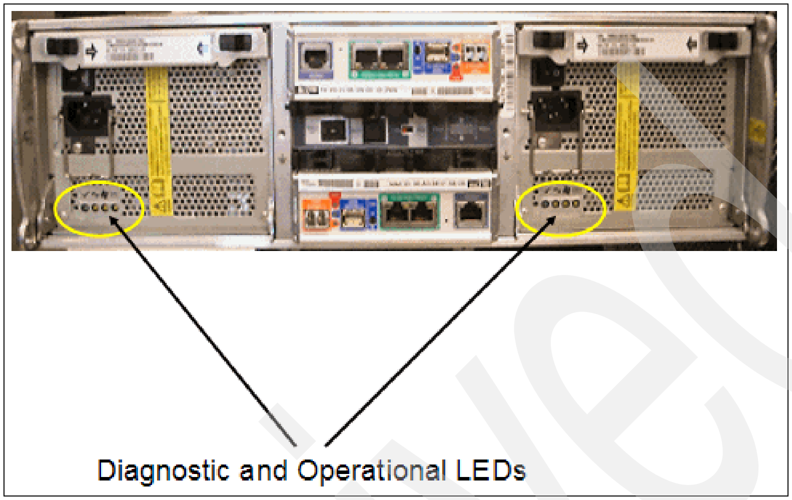

Figure 2-7 Diagnostic and operational LEDs

Table 2-10 LED status

LEDs visible from the rear of the system unit		
PSU status normal	On	Normal
AC missing for this PSU	Off	
Fan fault	Off	
Output voltage, current, temperature fault	Off	
PSU status normal	Off	Power supply failure
AC missing for this PSU	Off	
Fan fault	Off	
Output voltage, current, temperature fault	On	
PSU status normal	Off	Fan failure
AC missing for this PSU	Off	
Fan fault	On	
Output voltage, current, temperature fault	Off	

LEDs visible from the rear of the system unit		
PSU status normal	Off	No power to this PSU
AC missing for this PSU	On	
Fan fault	Off	
Output voltage, current, temperature fault	On	

The CPU module is shown on Figure 2-8; it controls connectivity to the storage and connectivity to the clients. If a power supply fails or is turned off while the other power supply is still providing DC power, both cooling fans continue to operate.

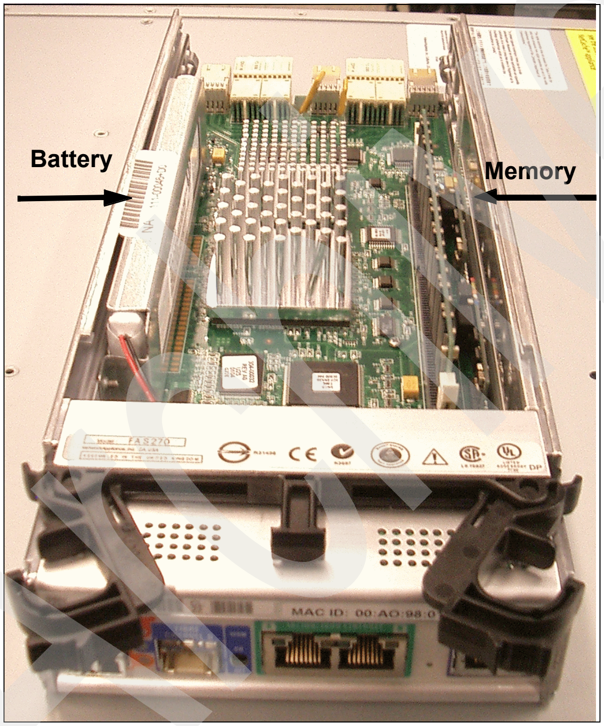

Figure 2-8 CPU tray module front view

The N3700 is based around a MIPS dual-core processor. It has 1 GB of system memory, of which 128 MB is defined as nonvolatile because it has a battery backup (see Figure 2-9 on page 37). The battery is a 3-cell Li-Ion.

Note: The NVRAM is a battery backed-up portion of the main system memory.

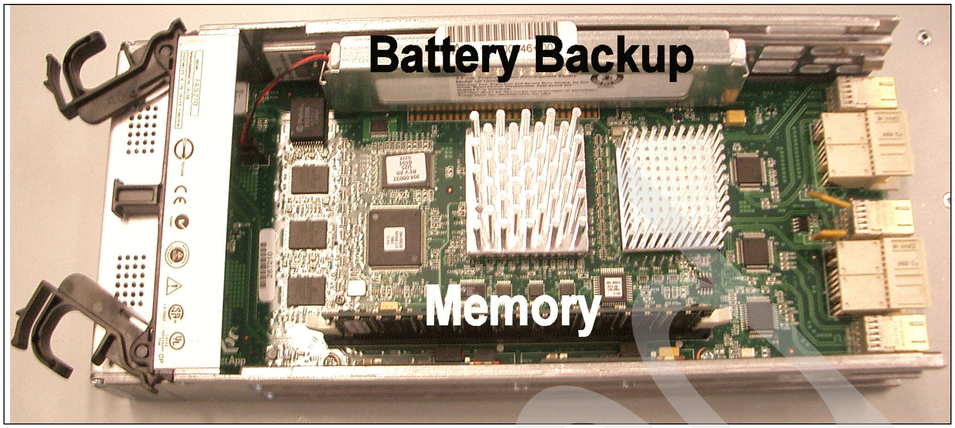

Figure 2-9 CPU tray module showing battery backup for memory

A 256 MB compact flash card is located on the bottom of the CPU tray module (see Figure 2-10). It contains a copy of the Data ONTAP operating system along with firmware. The operating system is also stored on each disk drive.

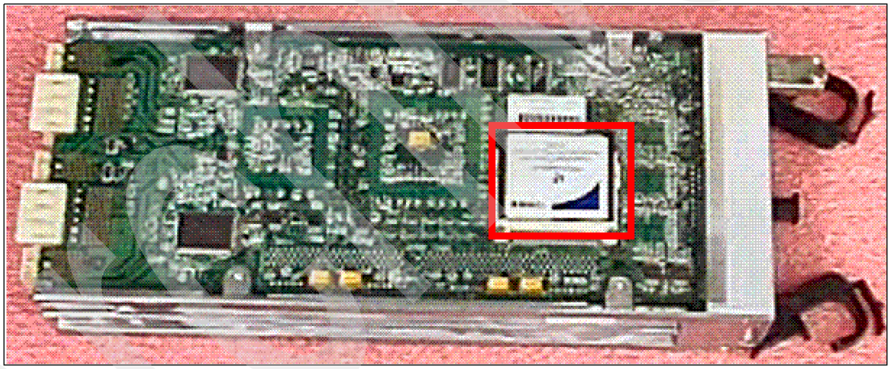

Figure 2-10 Bottom of CPU tray module showing compact flash card

Rear ports on the N3700 can be seen in Figure 2-11. Each CPU tray module has two integrated 2 Gbps Fibre Channel ports. Both of these ports are initially configured in “initiator mode” and do not use small form factor pluggable SFPs.

The first port, channel C, is optical and intended for direct or SAN attachment to a tape library. A standard LC-LC short wave optic cable must be used for this. The second port, channel B, is copper and is exclusively used for connecting an expansion unit. A special copper cable (option X6531-C) is used for this connection.

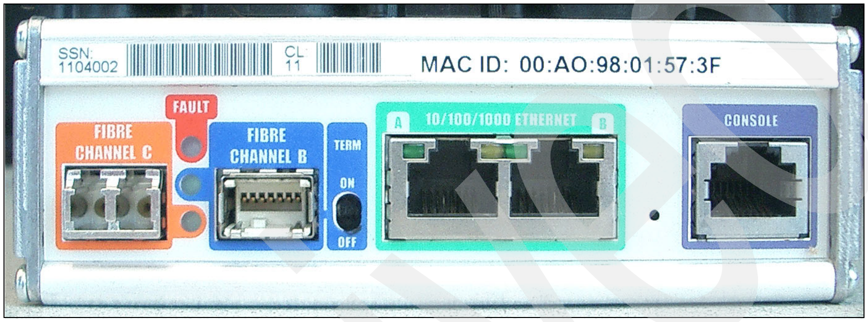

Figure 2-11 External ports on CPU tray module

The CPU tray module also contains two onboard 10/100/1000 MB copper Ethernet ports. Each port has two LED lights to indicate activity and speed. The final ports enable the connection of an ASCII terminal through an RJ45 to a DB-9 cable.

N3700 hardware features

The N3700 includes the following hardware features:

- 3U integrated storage system
- 3U optional storage expansion shelf, up to three
- Redundant hot plug power supplies
- Redundant cooling
- Integrated 10/100/1000 full-duplex Ethernet
- Two integrated Fibre Channel adapters
- Compact flash
- Diagnostic LEDs/OPS

Although ESH2 modules in the EXN2000 can support up to 84 drives per loop or 1 N3700 and 5 expansion shelves, only 56 drives, 3 expansion shelves, and 1 N3700 are supported. This firmware limitation for backward compatibility originates with the manufacturer for the previous shelf module the Loop Redundant Circuit (LRC). This module is not available on the N series.

Optional hardware

A second CPU tray supports cluster failover. Table 2-11 lists additional N3700 specifications.

Table 2-11 Additional N3700 specifications

Storage system specifications	N3700 A10	N3700 A20
Capable of clustered failover	No (requires upgrade to A20)	Yes
Max number of expansion units EXN1000/SATA disk drives or EXN2000/FC disk drives	3	3

N3300 and N3600

The N3300 and N3600 (see Figure 2-12 and Figure 2-13 on page 40) systems provide multiple I/O connectivity options, a small footprint to hold high-density SAS drives, and external expansion using low-cost SATA drives and Fibre Channel disks for production applications. The N3300 and N3600 utilize Data ONTAP Snapshot technology. SAS (Serial-Attached SCSI) is the next generation of SCSI, and it combines the advantages of parallel SCSI and serial FC. For further systems administration time and cost advantages, the systems come standard with remote onboard management capabilities to help simplify remote system monitoring, cycle power, execute firmware upgrades, enter console commands, and run diagnostics to help maintain the reliability of the system and your business-critical data.

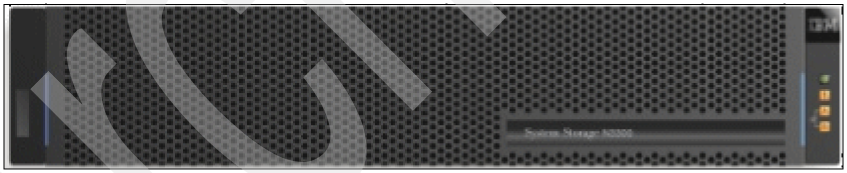

Figure 2-12 N3300

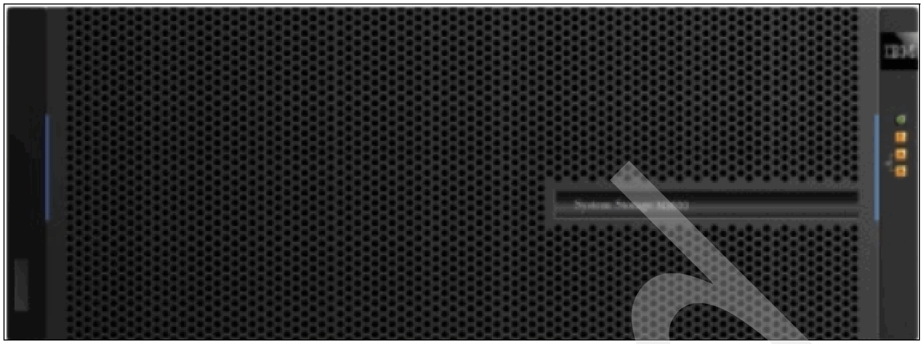

Figure 2-13 N3600

Figure 2-14 and Figure 2-15 show the rear panels of the N3300 and N3600. The single node A10 uses a single control unit, and the dual-node clustered A20 uses two control units (see Figure 2-5 on page 34).

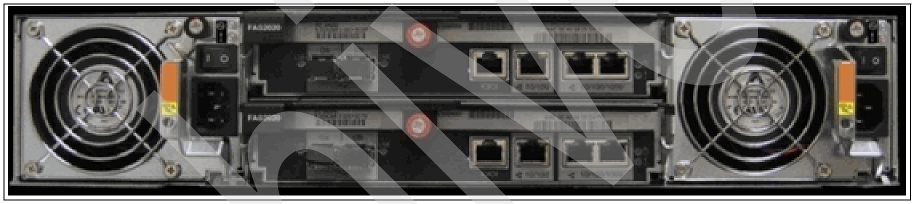

Figure 2-14 N3300 rear view

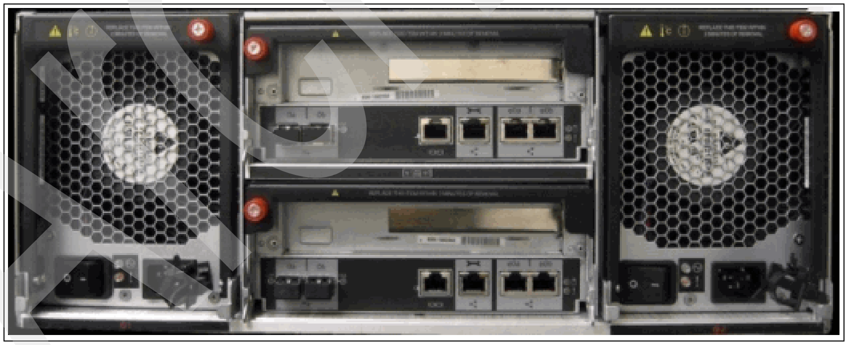

Figure 2-15 N3600 rear view

N3300 is a 2U high device. It has 12 internal SAS drive bays. It can support up to two external disk expansion units. Each controller has dual GB Ethernet ports and dual 4 Gbps Fibre Channel ports (see Figure 2-16). The N3300 also has one console port and one remote management port.

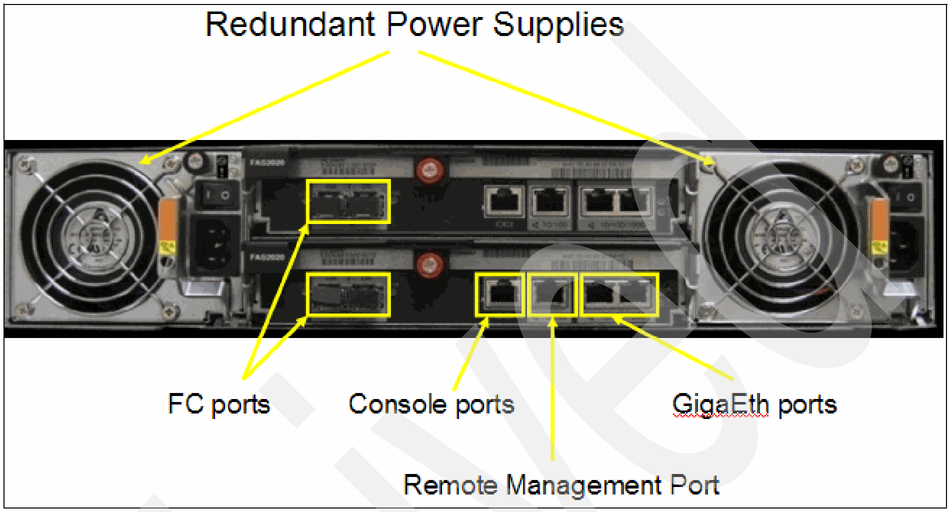

Figure 2-16 External ports on N3300

Note: The N3300 series supports SAS, FC, and SATA disk technologies. Twelve SAS disk drives are supported in the controller chassis. The N3300 can be configured with 0 disk drives in the controller and use the storage from disk expansion units such as the EXN1000 for SATA or EXN4000 for Fibre Channel disks.

The N3600 also has redundant power supplies (see Figure 2-17).

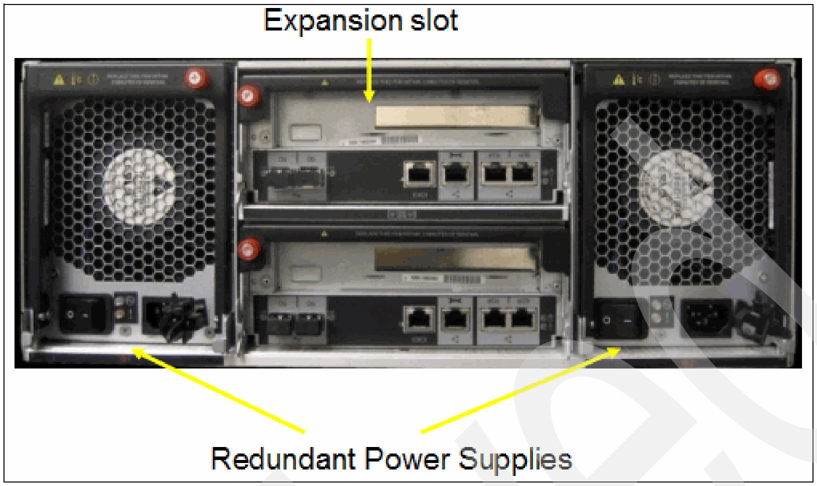

Figure 2-17 N3600 power supplies and expansion slot

The N3600 is a 4U high device. It has 20 Internal SAS drive bays. N3600 can support up to 6 external disk expansion units. Each controller has dual Gigabit Ethernet ports and dual 4 Gbps Fibre Channel ports (see Figure 2-18). The N3600 also has one console port and one remote management port. The N3600 has a PCIe slot on each controller.

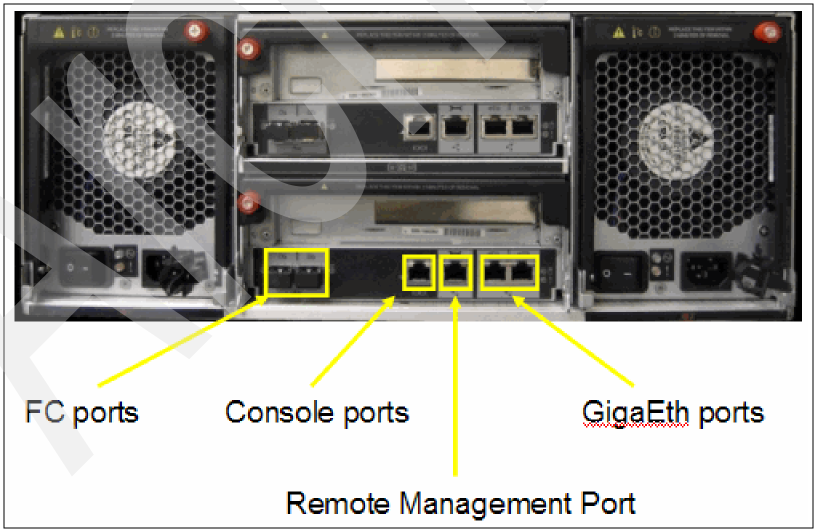

Figure 2-18 External ports on N3600

> **Note:** The N3600 series supports SAS, FC, and SATA disk technologies. Twenty SAS disk drives are supported in the controller chassis. The N3600 requires a minimum of six SAS drives in the controller chassis.

The key specifications for the N3300 and N3600 are as follows:

- Two U high (N3300) and 4U high (N3600)
- Up to two external disk expansion units for N3300 and up to six external disk expansion units for N3600
- High-performance SAS infrastructure
- Single controller or dual controller (for HA)
- Unified storage: iSCSI, NAS, Fibre Channel
- Each controller: dual Gigabit Ethernet ports and dual 4 Gbps Fibre Channel ports
- Onboard remote platform management
- Internal SAS drive bays

The N3000 series are small form-factor appliances that conserve scarce and valuable space in data centers or remote office locations. The N3000 is engineered for small to medium enterprises.

2.6.2 IBM System Storage N5000 introduction

The N5200, N5300, N5500, and N5600 are suitable for environments that demand data in high availability, high capacity, and highly secure data storage solutions. The IBM System Storage N5000 series offers an additional choice to organizations for enterprise data management. The IBM System Storage N5000 series is designed to deliver midrange to high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help support your efforts to increase reliability, simplify and unify storage infrastructure and maintenance, and deliver exceptional economy.

- The IBM N5000 A series comes in four models:
  - N5200
    - 2864-A10 single filer
    - 2864-A20 clustered
  - N5300
    - 2869-A10 single filer
    - 2869-A20 clustered

- – N5500
    - 2865-A10 single filer
    - 2865-A20 clustered
- – N5600
    - 2868 -A10 single filer
    - 2868 -A20 clustered

- FC or SATA (both can be used behind a single controller but not in the same drawer)

The N5200 and N5500 have no visible external differences. The differences are in maximum storage capacity and CPU processing power as listed in Table 2-4 on page 20. From the front, the N5600 and N5300 look very similar to the N5200 and N5500; some of the differences are seen from the rear (see Figure 2-19 and Figure 2-20 on page 45,) especially the absence of a LVD SCSI connector. The N5600 and N5300 (see Figure 2-21 on page 46) also use a BIOS prompt upon boot rather than a Common Firmware Environment (CFE) prompt.

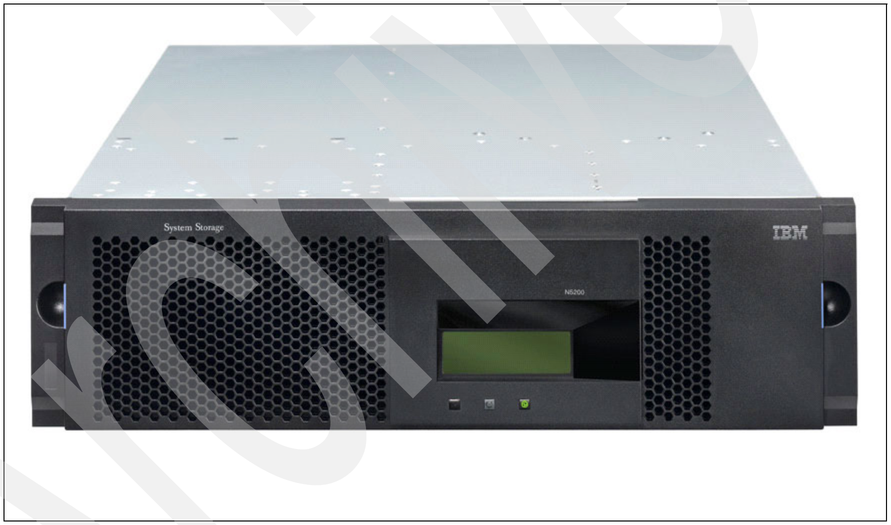

Figure 2-19   N5200

The N5000 A10 models come in a compact 3U rack-mountable unit that can coexist in the same rack as a EXN1000, EXN2000, or EXN4000 storage expansion unit (see Figure 2-20 and Figure 2-22 on page 46). The A20 models require 6U of space.

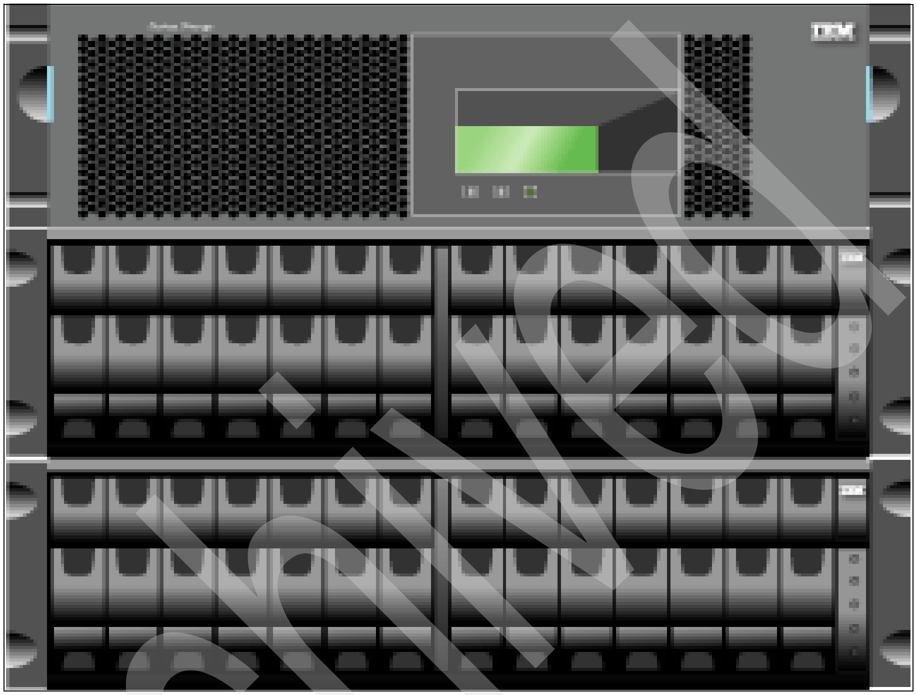

Figure 2-20   N5500 with EXN2000 shelves

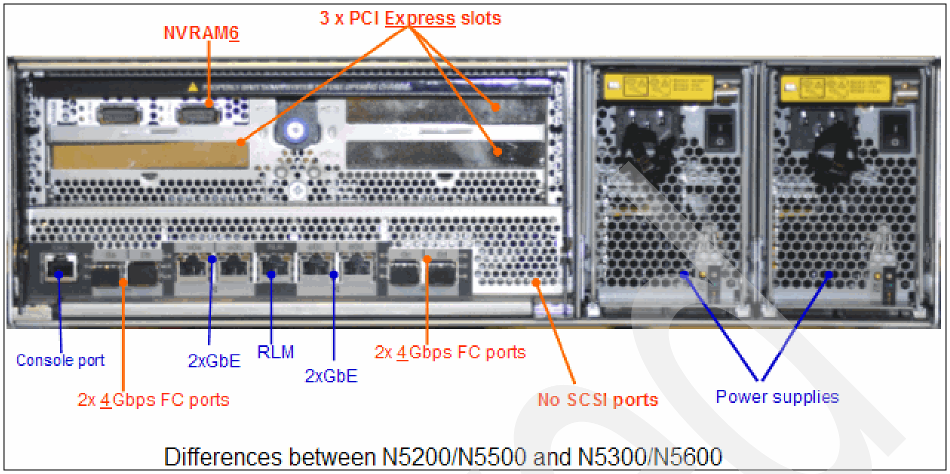

Figure 2-21 Rear view of N5600 and N5300

Depending on the model, one (N5200) or two (N5500) internal modules are available (see Figure 2-22).

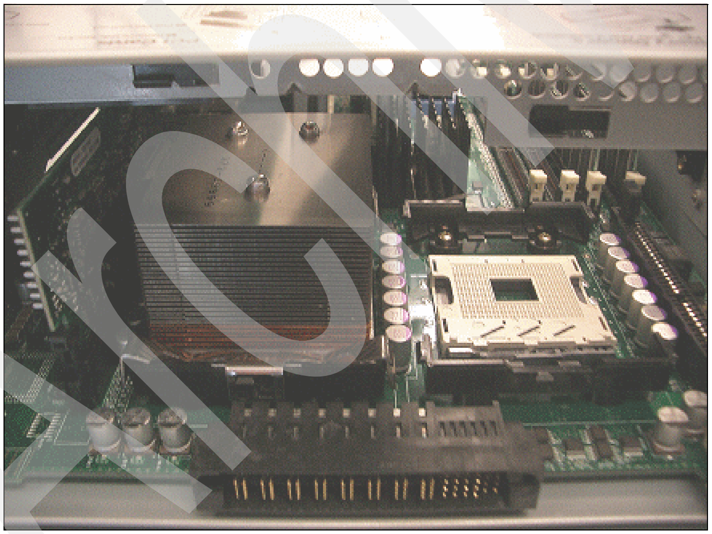

Figure 2-22 N5200: one CPU module

The easily accessible rear of the N5000 series provides I/O connectivity, power supply access, and status indications (see Figure 2-23).

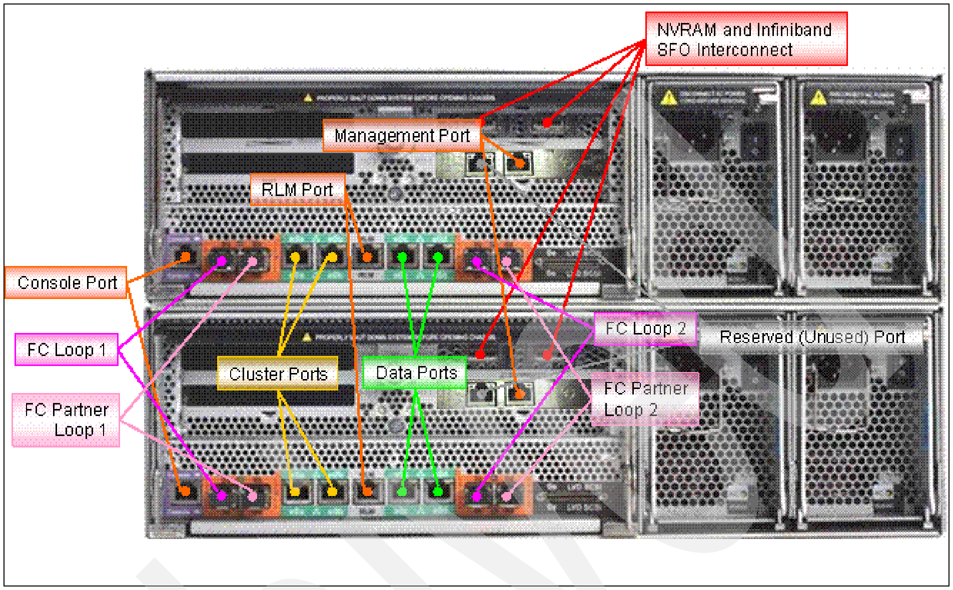

Figure 2-23 Rear view of N5500

The top view of the N5200 and N5500 is shown in Figure 2-24, displaying the modular design and field-replaceable unit capabilities.

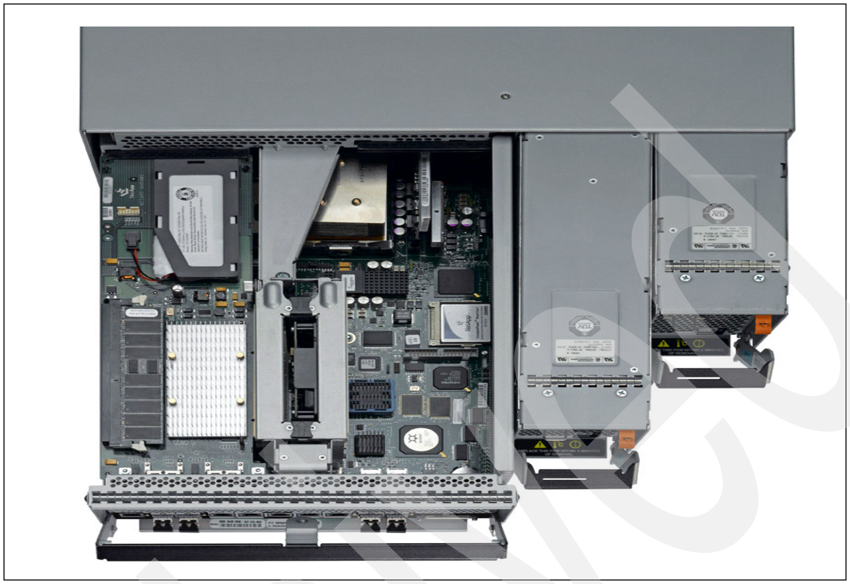

Figure 2-24 Top view of N5000

The motherboard of the N5200and N5500 is a self-contained unit holding components such as memory, CPU, and interfaces (see Figure 2-25).

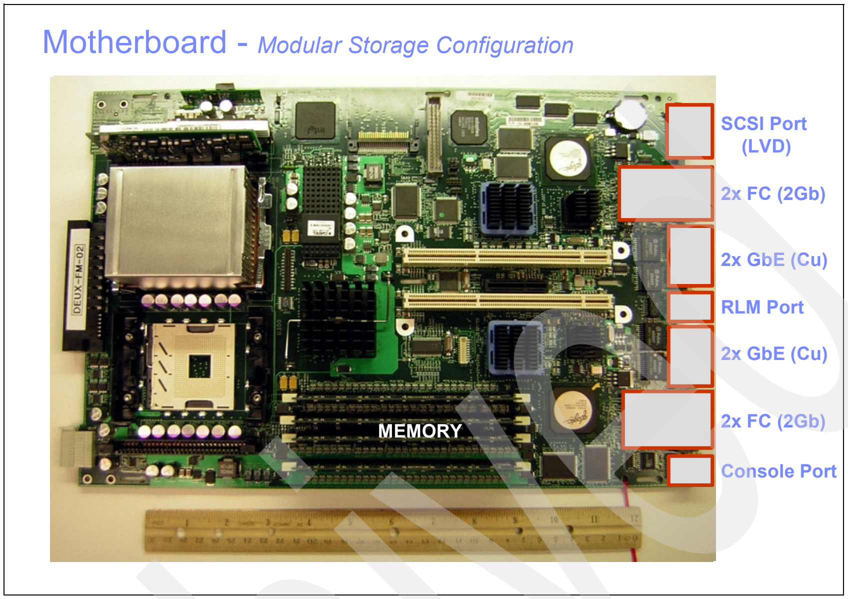

Figure 2-25 N5200 and N5500 motherboard

Table 2-12 lists the RAID group sizes.

Table 2-12 RAID group size in drive type

Model	FC-AL drives default	FC-AL drives maximum	ATA drives default	ATA drives maximum
RAID 4	8	14	7	7
RAID DP	16	28	14	16

Figure 2-26 depicts multiple storage options in the following models:

- N5200
- N5300
- N5500
- N5600

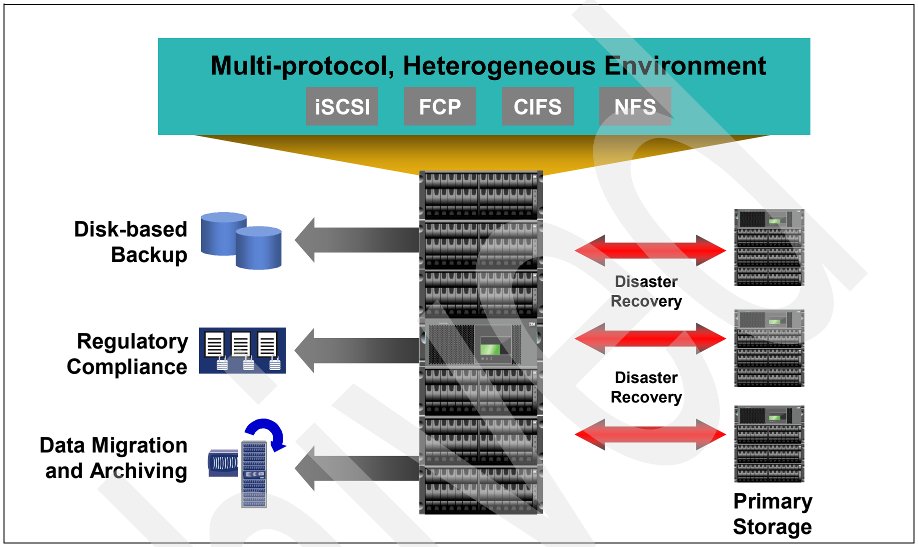

Figure 2-26 Multiple storage options

Multidisk drive options offer mission-critical, near-line, and compliance storage solutions.

2.6.3 IBM System Storage N7000 introduction

The IBM System Storage N7000 series offers additional choices to organizations facing the challenges of enterprise data management. The IBM System Storage N7000 series is designed to deliver high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help support your efforts to increase reliability, simplify and unify storage infrastructure and maintenance, and deliver exceptional economy.

- The IBM N7000 A series comes in four models:
  - N7600
    - 2866-A10 single node
    - 2866-A20 clustered

  - N7700
    - 2866-A11 single node
    - 2866-A21 clustered
  - N7800
    - 2867-A10 single node
    - 2867-A20 clustered
  - N7900
    - 2867-A11 single node
    - 2867-A21 clustered
- FC or SATA (both can be used behind a single controller but not in the same drawer)

Like its N5000 predecessor, the front of the N7000 series unit has an LCD display and the standard three LEDs, indicating system activity, status, and power (see Figure 2-27). Externally the N7600 and N7800 appear the same; the differences lie internally with increased CPU, memory, and NVRAM capability of the N7800 as compared to the N7600.

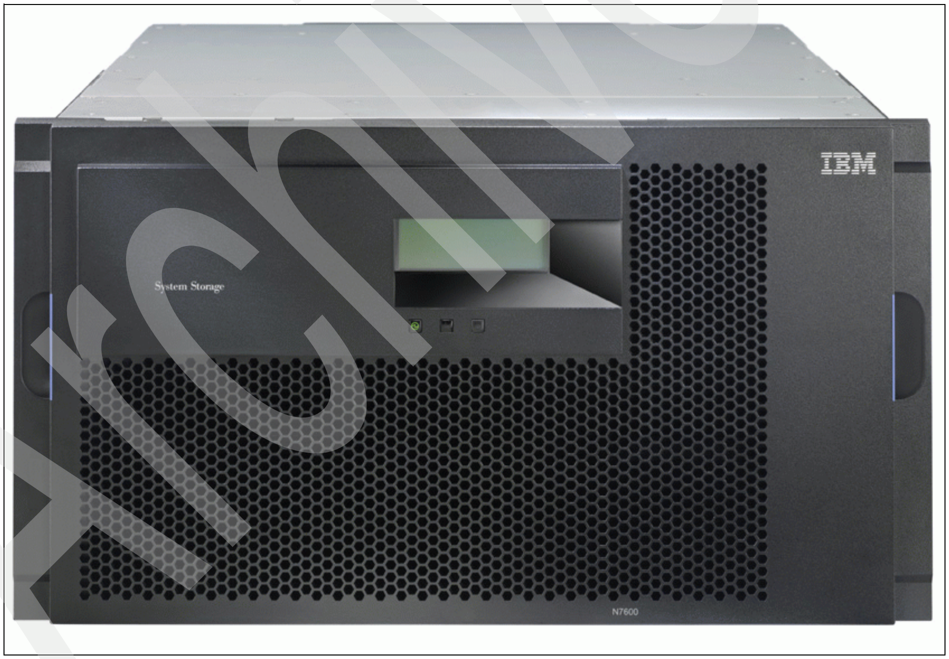

Figure 2-27 Front view of N7000

From the rear of the N7000, you can see the redundant power supplies, the NVRAM card, the Gigabit Ethernet interfaces as well as the Fibre Channel interfaces. The console port and RLM port are also located on the rear (see Figure 2-28).

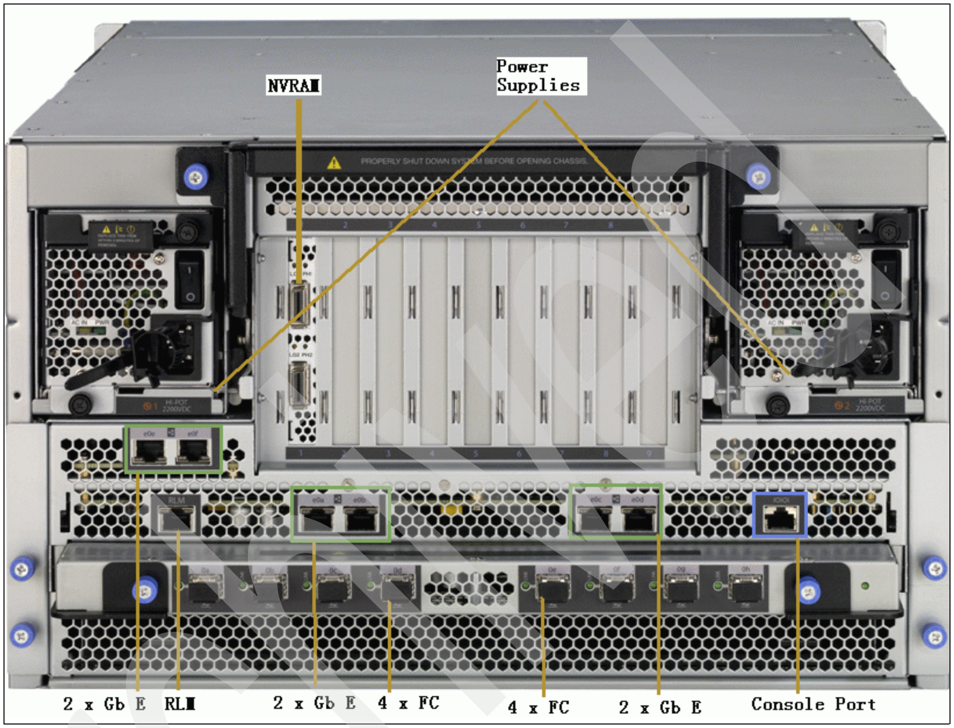

Figure 2-28 N7000 rear view

Each N7000 node requires 6U of rack space; each expansion unit requires 3U of rack space. Each N7000 node requires at least one expansion unit (see Figure 2-29).

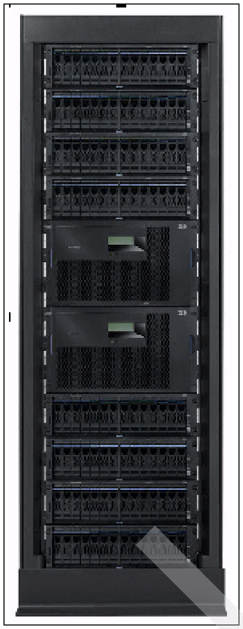

Figure 2-29 N7000 racked

A dual-node N7600 supports a maximum of 60 storage expansion units (EXN1000 or EXN2000, or both). A dual-node N7800 supports a maximum of 72 storage expansion units (EXN1000 or EXN2000, or both). (See Figure 2-30 on page 54.) A dual-node N7700 and N7900 supports a maximum of 60 and 84 storage expansion units, respectively. Each rack holds a maximum of 12 expansion units. The N7000 products are installed by IBM service or a qualified IBM Business partner and are not set up by customers.

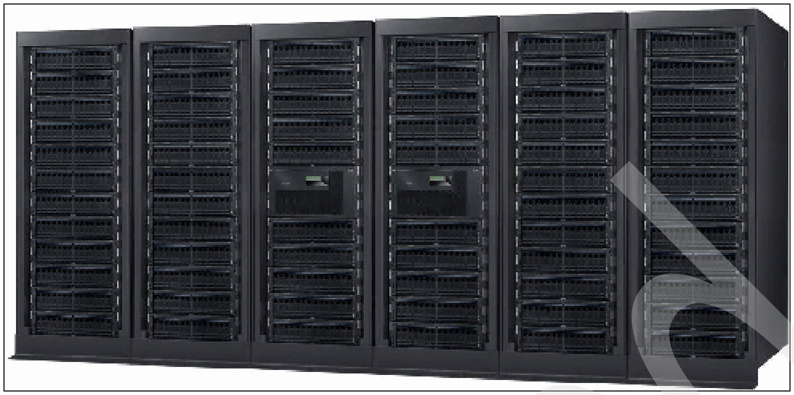

Figure 2-30 Clustered N7000 with multiple expansion units

When you remove the bezel, you see the CompactFlash card reader and, directly below it, the Remote LAN module, or RLM. The RLM is required in all N7000 series systems. The systems cannot boot unless the card is present. Also, you see five fan units. The fans are hot swappable and are numbered here for your reference (Figure 2-31).

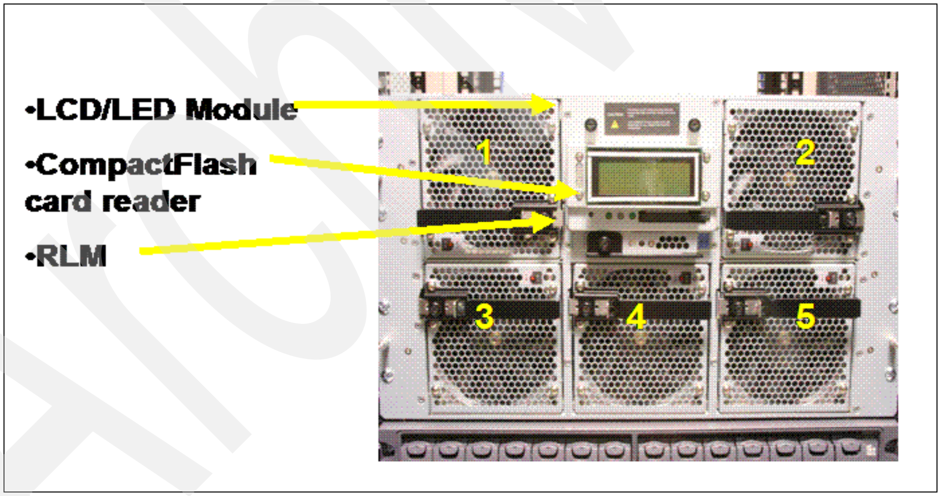

Figure 2-31 Front of N7000 with bezel removed

On the side of the system, notice the two handles on each side to help you lift the system (see Figure 2-32). The system is very heavy; fully loaded, it weighs 120 pounds. IBM recommends that before lifting the system, you remove the fan units and the two power supplies. This reduces the weight to slightly over 90 pounds. We recommend three people lift the system.

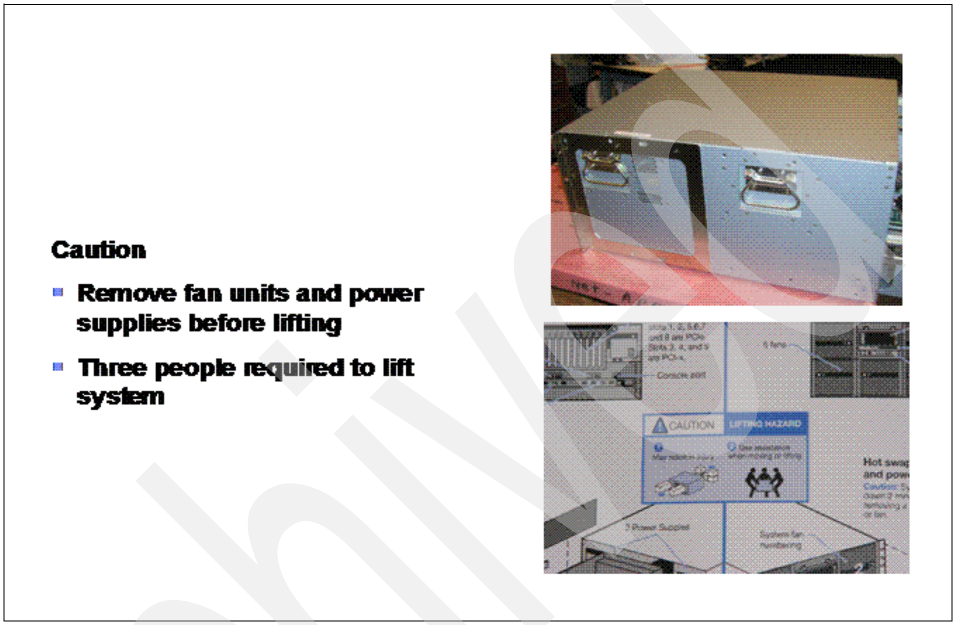

Caution

- **Remove fan units and power supplies before lifting**
- **Three people required to lift system**

Figure 2-32 Lifting N7000

The two hot swappable power supplies can be seen and removed from the rear of the N7000 (see Figure 2-33).

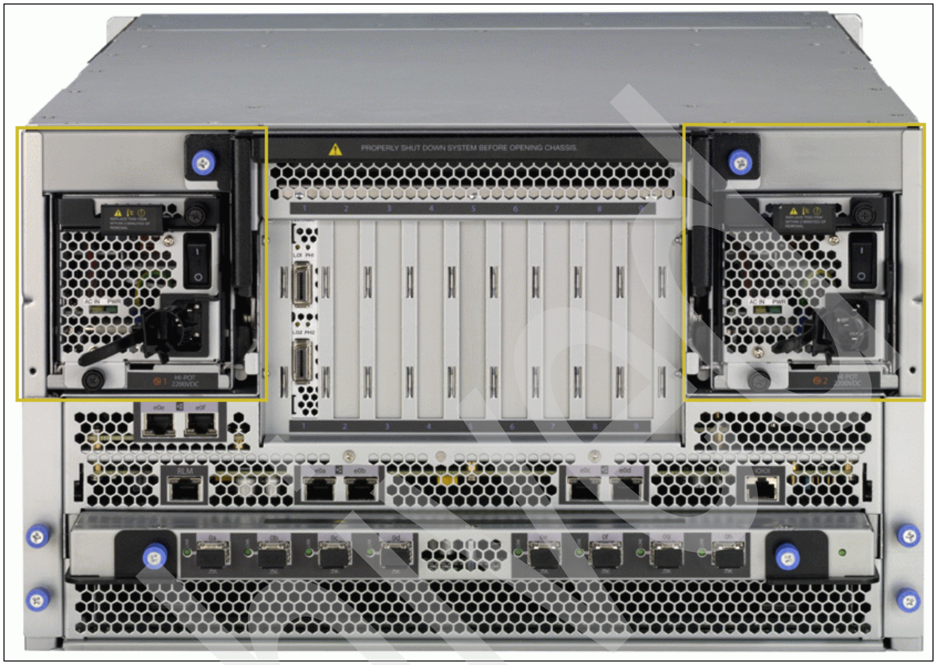

Figure 2-33 N7000 power supplies

In the middle are 9 PCI slots. They are numbered 1-9, from left to right, and all the slots are slot specific (see Figure 2-34 on page 57). When installing adapter cards, always refer to *IBM System Storage N7000 Series Hardware and Service Guide*, GC26-7953-05. You can find the guide here:

http://www.ibm.com/storage/support/nas/

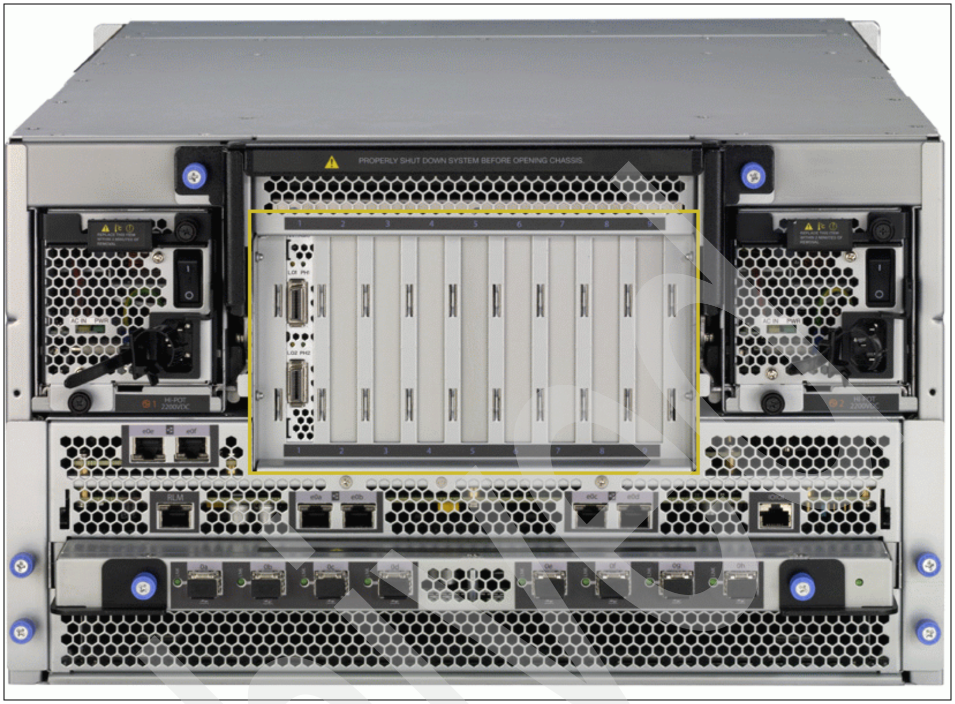

Figure 2-34 PCI slots

Below the PCI slots is a console port and RLM port (see Figure 2-35).

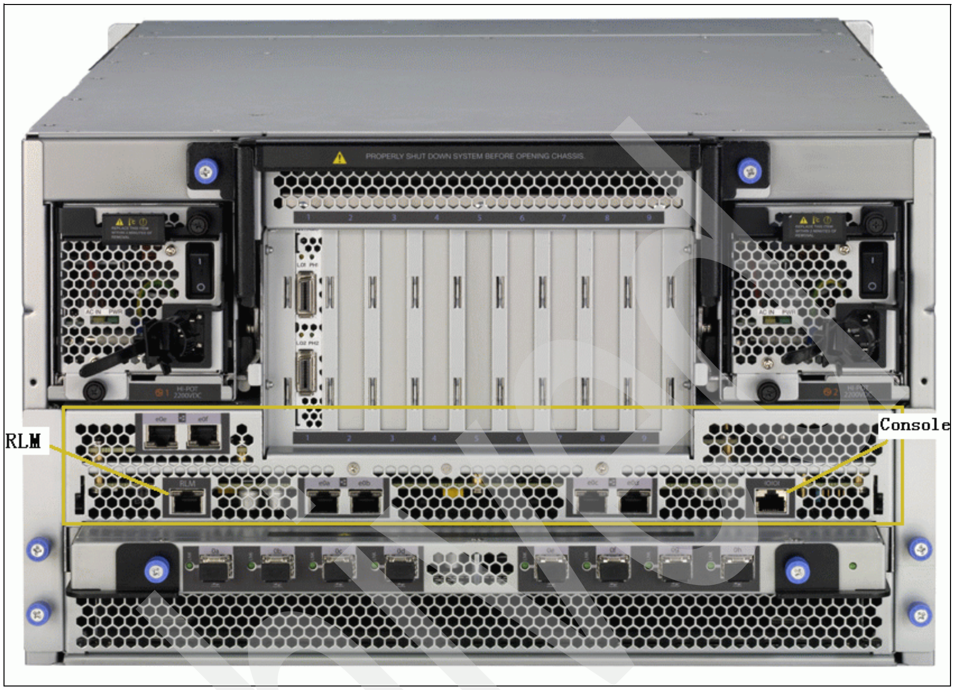

Figure 2-35 RLM and console ports

Below the Ethernet ports is the Fibre Channel tray, referred to as the *FC tray*, with eight onboard Fibre Channel ports (Figure 2-36). This tray is actually a field-replaceable unit.

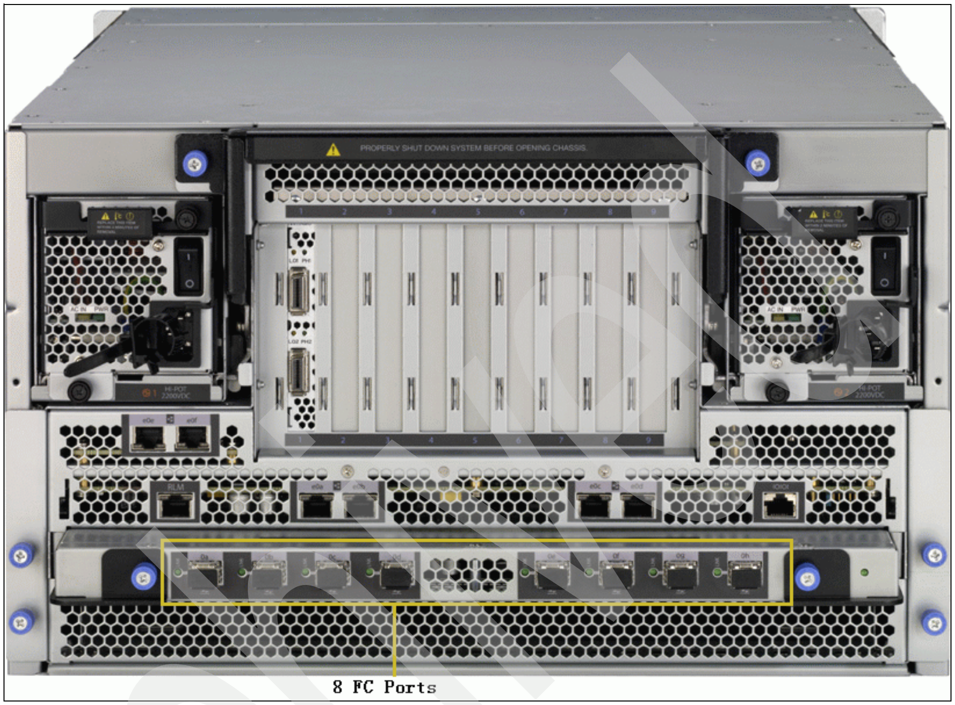

Figure 2-36 Fibre Channel ports

From inside the N7000 looking from the top (see Figure 2-37), you can see the PCI slots and system memory, but you cannot see the processors. They are on the other side of the motherboard tray. Recall that the N7900 and N7800 have eight CPUs, and the N7700 and N7600 have four processors.

From this perspective, you can see the nine PCI slots. Slots 3, 4, and 9 are black and represent PCI-X. Slots 1, 2, 5, 6, 7, and 8 are PCI-Express.

Notice the NVRAM6 adapter resides in slot 2 on this standalone system. On an active/active configuration, the NVRAM6 adapter resides in slot 1 and is used as the cluster interconnect card.

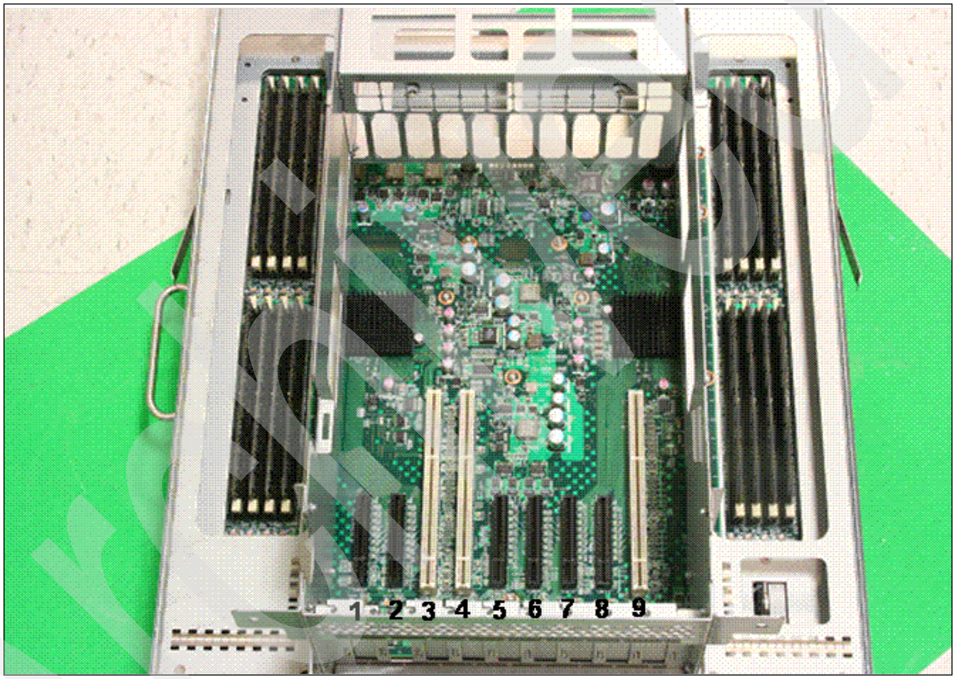

Figure 2-37 Top of N7000

Keep in mind that the N7800 and N7900 use an NVRAM6 adapter with 4 GB of memory, and the N7600 and N7700 use an NVRAM6 adapter with 1024 MB of memory (see Figure 2-38).

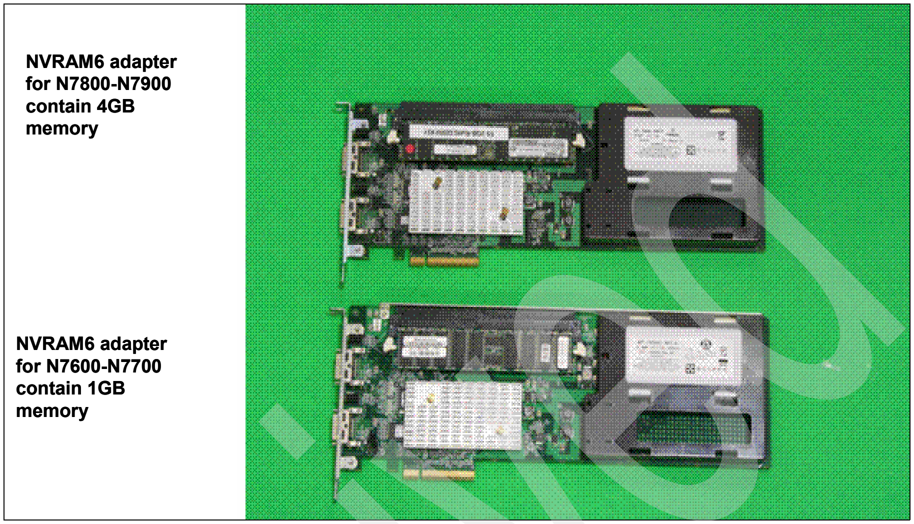
NVRAM6 adapter for N7800-N7900 contain 4GB memory
NVRAM6 adapter for N7600-N7700 contain 1GB memory

Figure 2-38 NVRAM

The N7000 includes new LEDs. The fan units, PCI slots, and memory DIMMs have LEDs to indicate a failed component (see Figure 2-39).

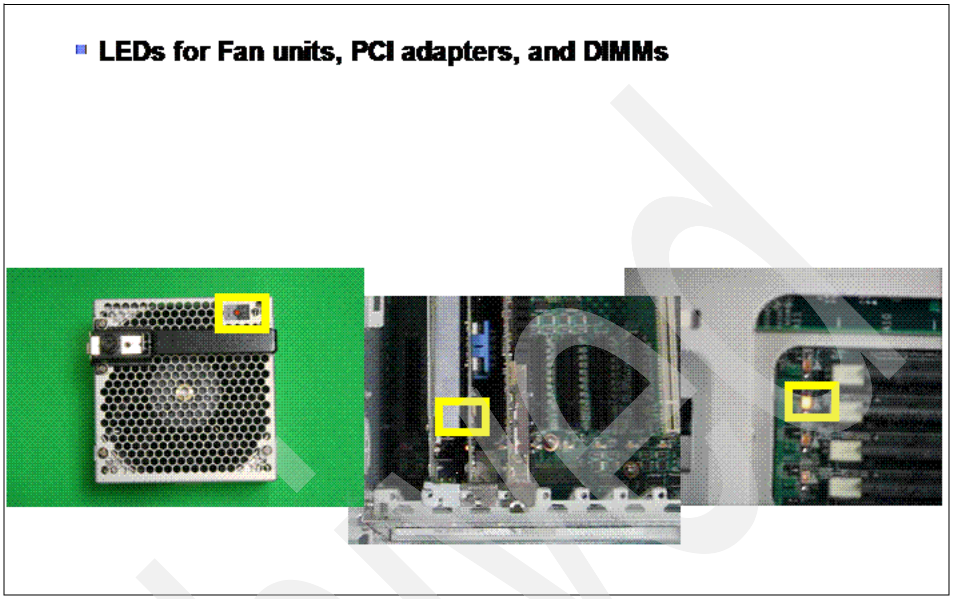

- LEDs for Fan units, PCI adapters, and DIMMs

Figure 2-39 New LEDs

Table 2-13 lists RAID group sizes.

Table 2-13 RAID group sizes in drive type

Model	FC-AL drives default	FC-AL drives maximum	ATA drives default	ATA drives maximum
N7600 RAID 4	8	14	7	7
N7600 RAID DP	16	28	14	16
N7700 RAID 4	8	14	7	7
N7700 RAID DP	16	28	14	16
N7800 RAID 4	8	14	7	7
N7800 RAID DP	16	28	14	16
N7900 RAID 4	8	14	7	7
N7900 RAID DP	16	28	14	16

2.7 IBM N series Gateway (G models)

The IBM System Storage N series Gateway, an evolution of the N5000 series product line, is a network-based virtualization solution that virtualizes tiered, heterogeneous storage arrays, enabling clients to leverage the dynamic virtualization capabilities available in Data ONTAP across multiple tiers of IBM and vendor acquired storage. Like all IBM N series storage systems, the IBM N series Gateway family is based on the industry-hardened Data ONTAP microkernel operating system, which unifies block and file storage networking paradigms under a common architecture and brings a complete suite of IBM N series advanced data management capabilities for consolidating, protecting, and recovering mission-critical data for enterprise applications and users.

The industry’s most comprehensive virtualization solution, the N series Gateway provides proven and innovative data management capabilities for sharing, consolidating, protecting, and recovering mission-critical data for enterprise applications and users and seamlessly integrates into mission-critical enterprise-class SAN infrastructures. These innovative data management capabilities when deployed with disparate storage systems simplify heterogeneous storage management.

The N series Gateway presents shares, exports, or LUNs built on flexible volumes that reside on aggregates. The N series Gateway is also a host on the storage array SAN. Disks are not shipped with the N series Gateway. The N series Gateway virtualizes storage array LUNs (which are treated as disks) through Data ONTAP, presenting a unified management interface.

The N series Gateway offers clients new levels of performance, scalability, and a robust portfolio of proven data management software for sharing, consolidating, protecting, and recovering mission-critical data. N series storage systems seamlessly integrate into mission-critical SAN environments and provide a simple, elegant data management solution that decreases management complexity, improves asset utilization, and streamlines operations to increase business agility and reduce total cost of ownership.

Leveraging storage

Organizations are looking for ways to leverage SAN-attached storage to create a consolidated storage environment for the various classes of applications and storage needs throughout their enterprise. These organizations are looking for ways to increase utilization, simplify management, improve consolidation, enhance data protection, enable rapid recovery, increase business agility, deploy heterogeneous storage services, and broaden centralized storage usage by provisioning SAN capacity for business solutions requiring NAS, SAN, or IP SAN data access (see Figure 2-40 on page 64).

These organizations have:

- Significant investments or a desire to invest in a SAN architecture
- Excess capacity or an attractive storage cost for SAN capacity expansion
- Increasing requirements for both block (FCP, iSCSI) and file (NFS, CIFS, and so on) access
- Increasing local or remote shared file services and file access workloads

They are seeking solutions to cost-effectively increase utilization; consolidate distributed storage, direct access storage, and file services to SAN storage; simplify storage management; and improve storage management business practices.

With Data ONTAP the N series Gateway now supports the attachment of heterogeneous storage systems as well as IBM expansion units of the type used with N series storage systems (see Figure 2-40 andFigure 2-41 on page 65).

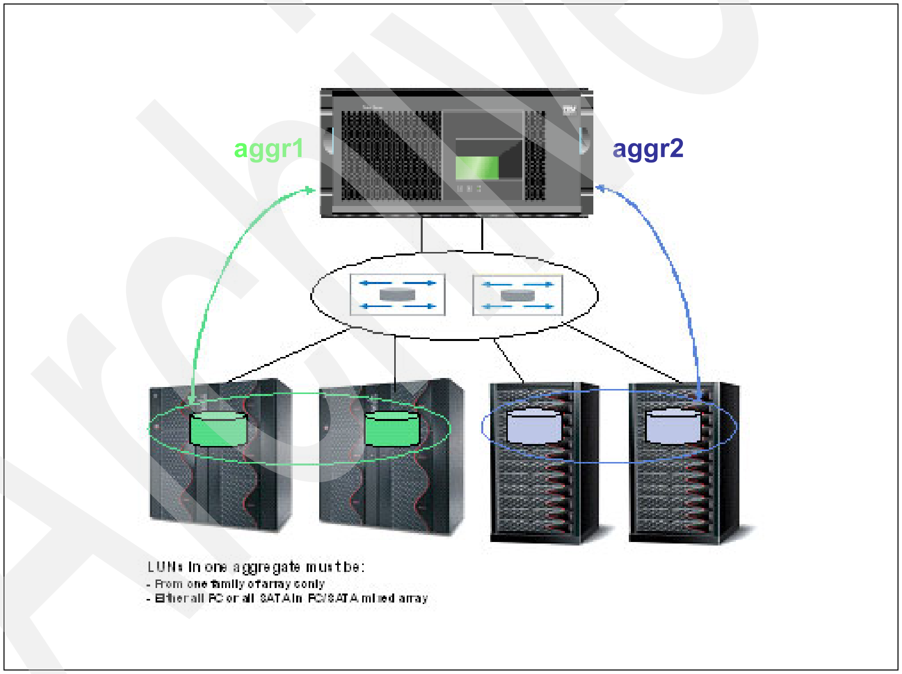
aggr1
aggr2
LUNs in one aggregate must be:
- From one family of array only
- Either all FC or all SATA in FC/SATA mixed array

Figure 2-40 Heterogeneous storage

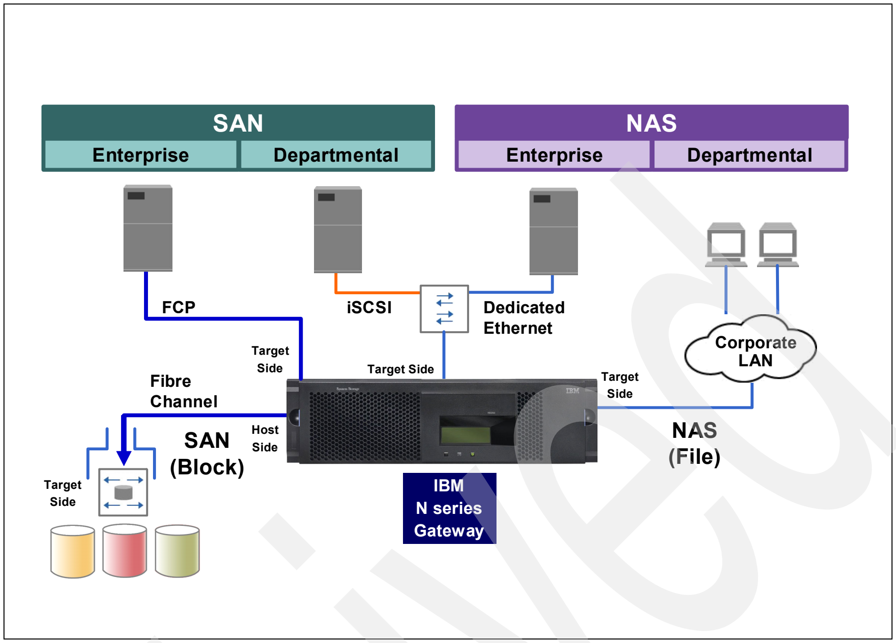

Figure 2-41 Gateway topology

Figure 2-42 shows the expansion units.

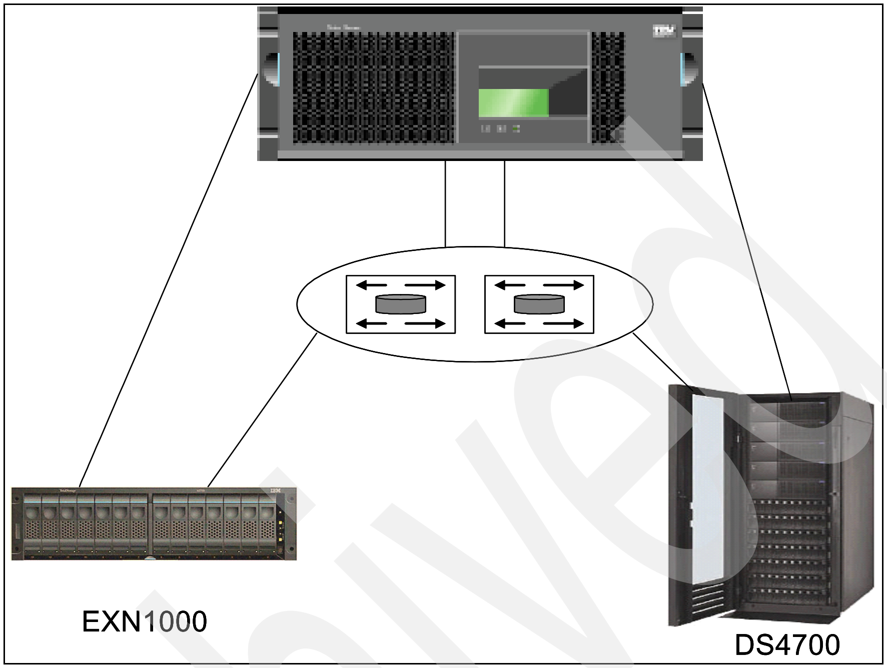

Figure 2-42 Storage systems with expansion units

Front-end and back-end implementations

A N series Gateway implementation can be thought of as a front-end implementation and a back-end implementation. The front-end setup includes configuring the N series Gateway for all protocols (NAS or FCP), implementing snap features (for example, snapshot, SnapMirror, SnapVault), and setting up backup including NDMP dumps to tapes. The back-end implementation includes all tasks required to set up the N series Gateway system up to the point where it is ready for Data ONTAP installation. These tasks include formatting array LUNs, assigning ports, setting up cabling and switch zoning, assigning LUNs to the V series system, creating aggregates, and loading Data ONTAP.

2.7.1 IBM N series Gateway highlights

IBM System Storage N series Gateway provides a number of key features that enhance the value and reduce the management costs of utilizing a SAN. An N series Gateway provides the following benefits:

- Simplifies storage provisioning and management
- Lowers storage management and operating costs

- Increases storage utilization
- Provides comprehensive simple-to-use data protection solutions
- Improves business practices and operational efficiency
- Transforms conventional storage systems into a better managed storage pool (see Figure 2-43)

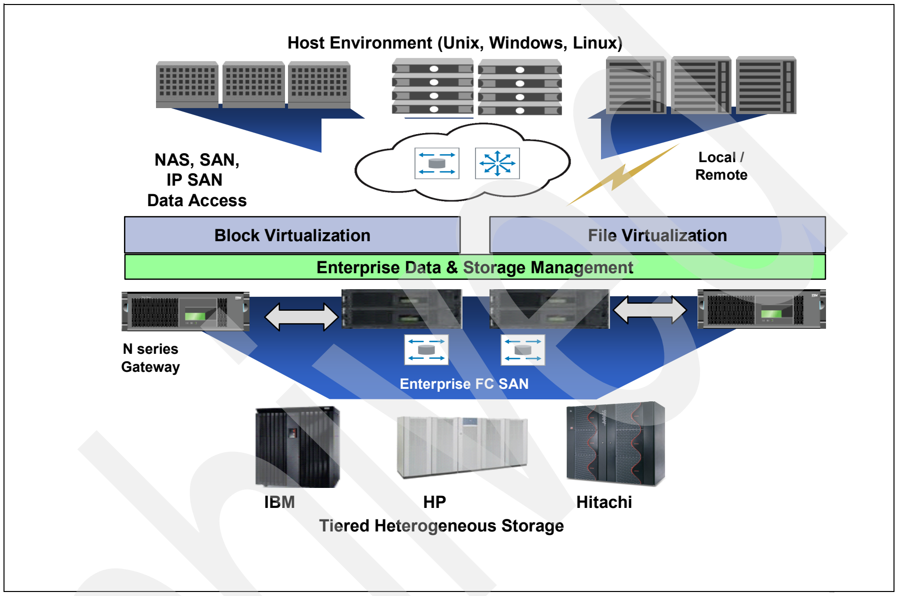

Figure 2-43 Tiered heterogeneous storage

2.7.2 Gateway RAID

The N series Gateway uses RAID0 on top of RAID1, RAID5, or RAID10 on RAID storage subsystems (see Figure 2-44). Physical disk operations such as scrubbing are disabled.

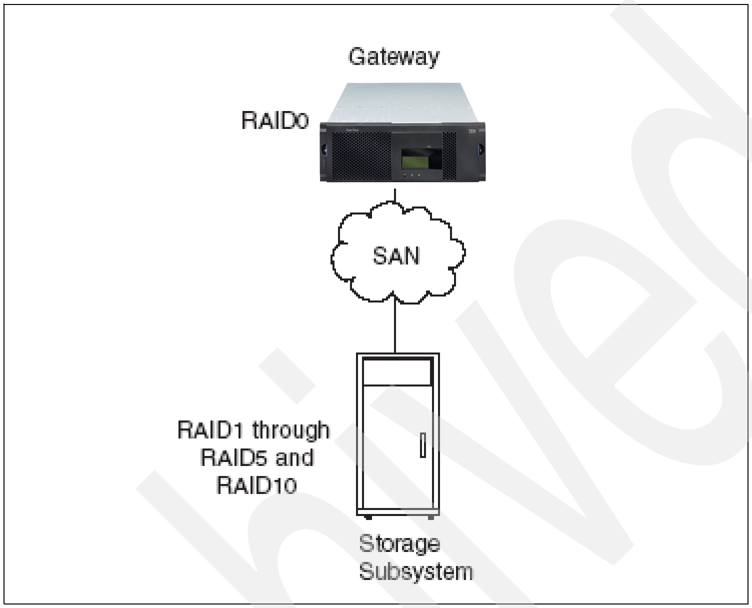

Figure 2-44 RAID configuration

RAID0 is used to write data. See Example 2-1 for a example of volume status. With the Gateway, it looks similar to what you see on a N series model A except for the RAID status.

Example 2-1 Volume status with Gateway volumes

```
itsotuc2*> vol status -v vol3
         Volume State      Status            Options
           vol3 online     raid0, flex     nosnap=on, nosnapdir=off,
                                             minra=off,
                                             no_atime_update=off,
                                             nvfail=off,
                                             snapmirrored=off,
                                             create_ucode=on,
                                             convert_ucode=on,
                                             maxdirsize=31457,
                                             fs_size_fixed=off,
                                             guarantee=volume,
                                             svo_enable=off,
                                             svo_checksum=off,
                                             svo_allow_rman=off,
                                             svo_reject_errors=off,
                                             fractional_reserve=100,
Containing aggregate: 'aggr0'

                Plex /aggr0/plex0: online, normal, active
                    RAID group /aggr0/plex0/rg0: normal
```

2.7.3 IBM N5200, N5300, N5500, and N5600 Gateway models

The N5000 Gateway models are a good value for those wishing to extend the reach of their SANs. The N5000 Gateway incorporates a variety of reliability and availability features designed to support high-demand operations. It houses hot swappable, redundant power supplies and fans, and supports multipath failover protection and host dual pathing between the unit and its SAN-attached storage device. In addition, the clustering feature between two storage systems is designed to help reduce system downtime.

From a hardware perspective, the G10 and G20 models are identical to the A10 and A20 models of the N5200, N5300, N5500, and N5600. The differences lie in the spectrum of Data ONTAP features supported and enabled.

- N5200
  - 2864-G10
  - 2864-G20 clustered model
- N5300
  - 2869-G10
  - 2869-G20 clustered model
- N5500
  - 2865-G10
  - 2865-G20 clustered model
- N5600
  - 2868-G10
  - 2868-G20 clustered model

Table 2-14 lists the Gateway models capacity.

Table 2-14 Gateway capacity

Model	Maximum capacity
2864-G10	50TB
2864-G20	50TB per node
2869-G10	126TB

Model	Maximum capacity
2869-G20	126TB
2865-G10	80TB
2865-G20	80 TB per node
2868	252TB
2868	252TB

Important: If you are going to enable the cf.takeover.on_panic option, ensure that a spare LUN is available for core dumps. If the cf.takeover.on_panic option is enabled and no spare LUN is available, no core dump file is produced on failure. (The cf.takeover.on_panic option controls whether a cluster partner immediately takes over for a panicked partner.)

Table 2-15 shows the number of LUNs supported by each model of the N series Gateway.

Table 2-15 LUNs supported by each model of N series Gateway

Model	Maximum number of LUNs
N5200 2864-G10 (noncluster model)	168
N5200 2864-G20 (cluster model)	For each node, single node N5200 2864-G10 values apply.
N5300 2869 - G10	252
N5300 2869 - G20	For each node, single node N5300 2869-G10 values apply.
N5500 2865-G10 (noncluster model)	336
N5500 2865-G20 (cluster model)	For each node, single node N5500 2865-G10 values apply.
N5600 2868-G10 (noncluster model)	504
N5600 2868-G20 (cluster model)	For each node, single node N5600 2868-G10 values apply.

2.7.4 IBM Gateway models N7600, N7700, N7800, and N7900

The IBM System Storage N7000 series Gateway models offer additional choices to organizations facing the challenges of enterprise data management. The IBM System Storage N7000 series is designed to deliver high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help support your efforts to increase reliability, simplify and unify storage infrastructure and maintenance, and deliver exceptional economy.

The IBM N series N7000 Gateway models deliver all the feature functionality that the N5000 series does but with increased processing, memory, NVRAM, and total storage capacity. The N7000 models are designed for the high end of enterprise environments. The N7000 series Gateway hardware is identical to the A1X and A2X models, with the difference being the enabled features and disk attachment by Data ONTAP.

The IBM N7000 G series comes in four models:

- N7600
  - 2866-G10 single node
  - 2866-G20 clustered
- N7700
  - 2866-G11 single node
  - 2866-G21 clustered
- N7800
  - 2867-G10 single node
  - 2867-G20 clustered
- N7900
  - 2867-G11 single node
  - 2867-G21 clustered

2.7.5 LUN sizing

N series Gateway support for LUN sizes is as follows:

- Maximum LUN size: 1024 GB
- Minimum LUN size: 1001 MB

Note: The Data ONTAP definition of a GB is as follows: One GB is equal to 1000 x 1024 x 1024 bytes. Therefore, the maximum LUN size that Data ONTAP supports means 1024 * 1000 * 1024 * 1024 = 1,048,576,000,000 bytes.

2.7.6 LUN mapping

Storage Subsystem LUNs are converted to disks for the IBM N series Gateway. Compared to the N series storage systems, the disk count is equivalent to the LUN count when only storage systems are attached to it.

Figure 2-45 is a example of an array LUN mapped to a Gateway disk.

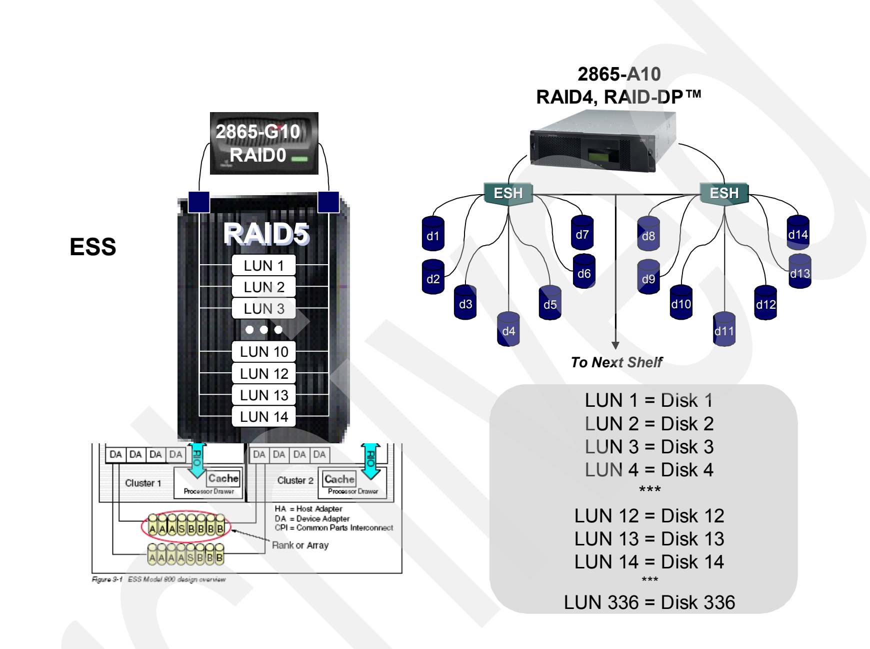

Figure 2-45 LUN to N series Gateway disk relationship

LUNs are added to the Gateway through the same volume wizard used on the N series A models (see Figure 2-46).

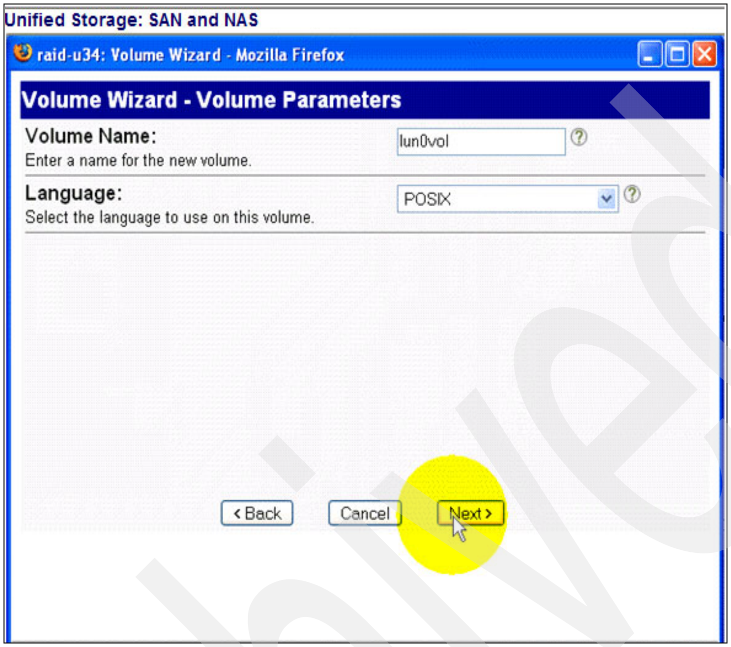

Unified Storage: SAN and NAS

raid-u34: Volume Wizard - Mozilla Firefox

Volume Wizard - Volume Parameters

Volume Name: lun0vol
Enter a name for the new volume.

Language: POSIX
Select the language to use on this volume.

Figure 2-46 Volume wizard

Note: Do not map LUN 0 to Gateway systems, even if LUN 0 is a storage LUN.

2.8 Interoperability between G and A models

Replication between SnapMirror on the G model and SnapMirror on the A model (Figure 2-47) includes async, semisync, and synchronous.

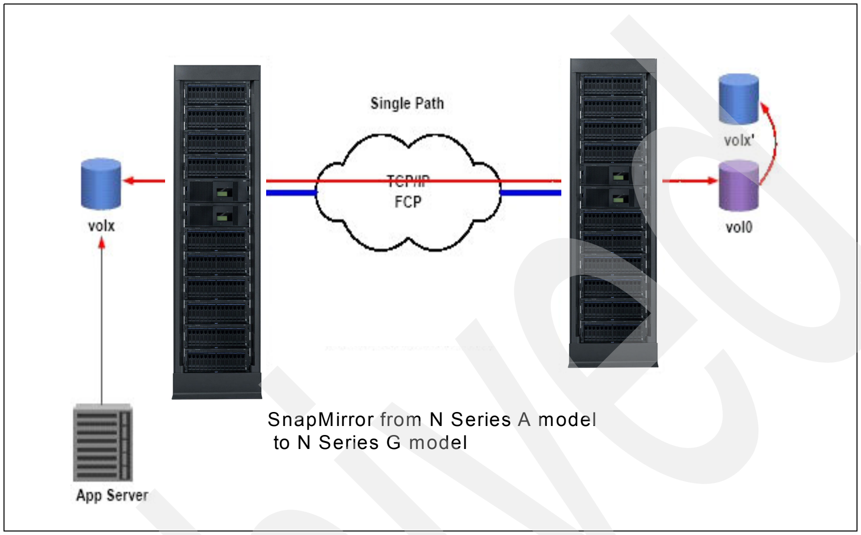

Figure 2-47 SnapMirror interoperability

Disk-to-disk backup from SnapVault is primary on the G model, while SnapVault is secondary on the A model (see Figure 2-48).

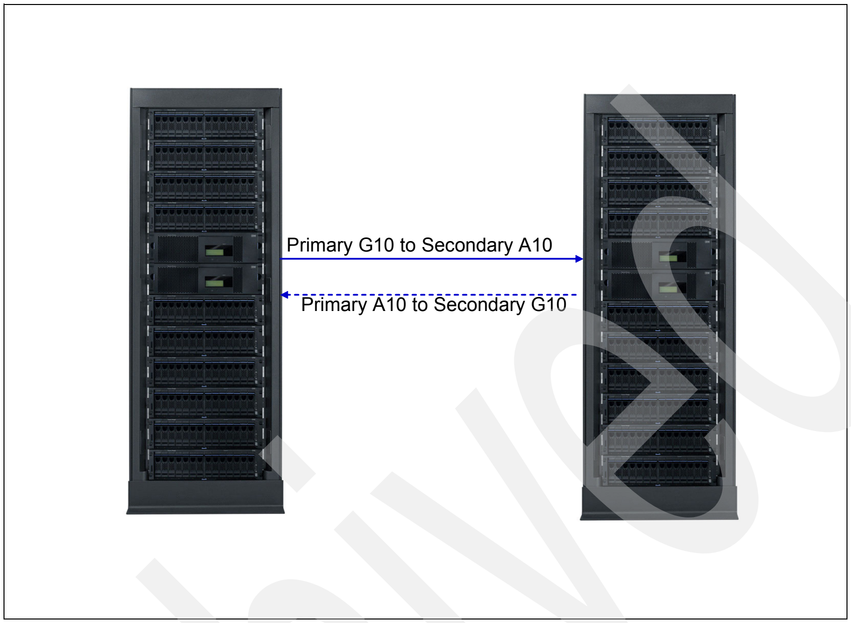

Figure 2-48 SnapVault interoperability

Disk-to-disk backup from SnapVault is primary on the A model, while SnapVault is secondary on the G model.

2.9 N series expansion units

Currently three disk storage expansion units are specifically designed for the IBM N series filers:

- IBM EXN4000 Fibre Channel disk storage expansion unit
- IBM EXN2000 Fibre Channel disk storage expansion unit
- IBM EXN1000 Serial Advanced Technology Attachment (SATA) storage expansion unit

Note: EXN expansion units are not intended for attachment to a Gateway.

Multiple EXN1000s, each having different SATA disk drive feature codes, can be attached to the same N series filer on the same Fibre Channel loop. Multiple EXN2000s and EXN4000s, each having different Fibre Channel disk drive feature codes, can be attached to the same N series filer on the same Fibre Channel loop. For the latest storage expansion unit support information, visit the following Web site:

http://www.ibm.com/storage/support/nas/

2.9.1 Intermixing EXN units with N series A models

EXN4000s and EXN2000s are both Fibre Channel disk storage expansion units. EXN4000 and EXN2000 expansion units can be mixed within the same loop, but the speed switches on all EXN4000s and EXN2000s must be set to the same speed (either 1 Gbps or 2 Gbps). (See Figure 2-49.) Intermixing Fibre Channel and SATA disk drives in a supported N series filer configuration is supported as follows:

- Intermixing Fibre Channel disk expansion units with SATA disk expansion units on the same loop is not supported.
- EXN4000s or EXN2000s (Fibre Channel disk drives) and EXN1000s (SATA disk drives) can be attached to the same N series filer only if the Fibre Channel disk expansion units (EXN4000s or EXN2000s) are on separate loops than the SATA disk expansion units (EXN1000s).

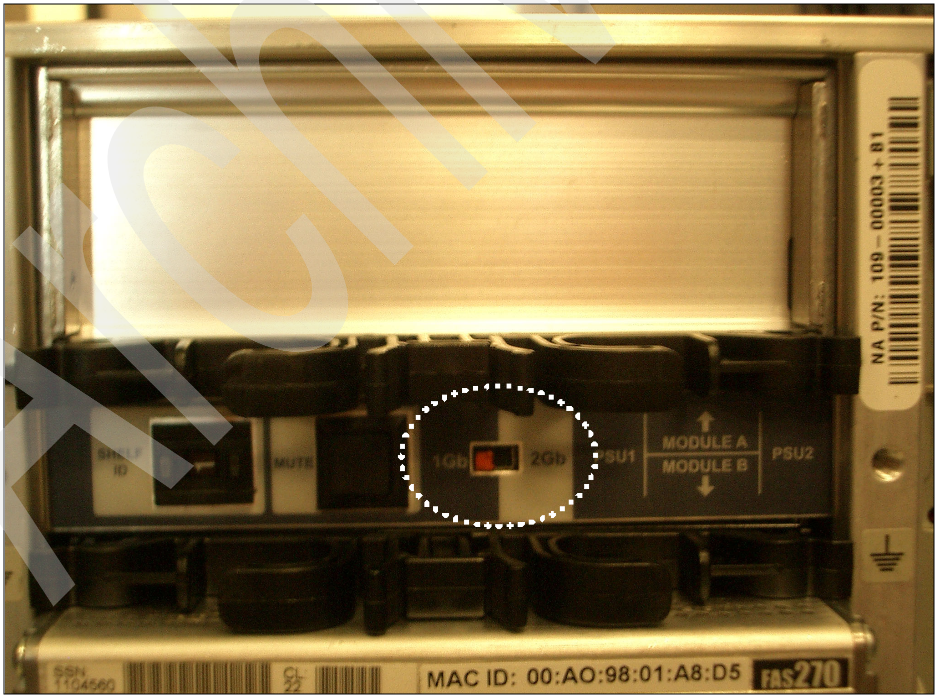

Figure 2-49 Speed switches

Note: Intermixing Fibre Channel and SATA disk drives in an N3700 configuration is not supported. Only N3300, N3600, N5000, and N7000 series models support intermixing of Fibre Channel and SATA disk drives in a configuration.

2.9.2 EXN2000

The EXN2000 is fibre expansion unit for the N series. The EXN2000 looks very similar to the N3700, but unlike the N3700, which has the CPU modules, the EXN2000 supports only the disk modules and the connectivity to them (see Figure 2-50).

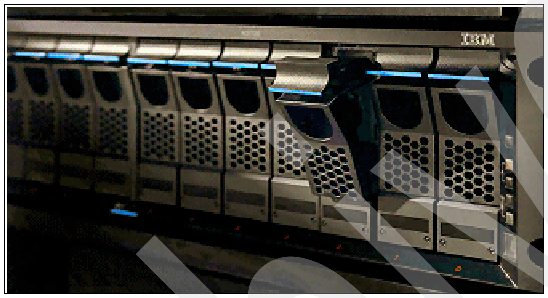

Figure 2-50 EXN2000

The EXN2000 is identical to the N3700 chassis except that the slot holding the CPU tray is replaced with an Electronically Switched Hub (ESH2). ESH2 provides a point-to-point connection to the drives (see Figure 2-51) rather than the traditional arbitrated loop. The maximum number of drives per shelf is unaffected by the capacity of the individual drive modules. Mixing of drives of different capacity in the same shelf is not recommended because of the effects it has on sparing, RAID groups, and flex volumes. The maximum number of drives on a loop are 84 or 6 shelves using the ESH2 module.

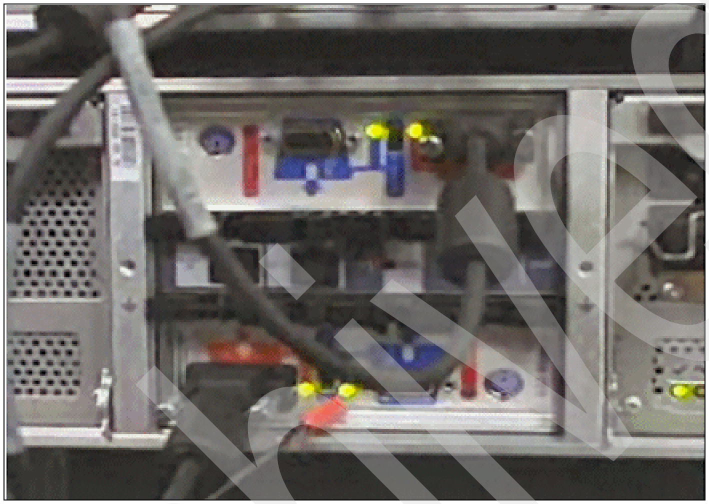

Figure 2-51 Rear of EXN2000: schematic of arbitrated loop versus switch hub

A switched hub architecture (see Figure 2-52 on page 79) has the benefit of additional availability, boosted performance in high I/O environments, and more powerful diagnostic abilities. Figure 2-53 on page 79 shows the ESH2 module. From a purely technical viewpoint, Fibre Channel loops support 126 devices. From a practical position, traditional FC-AL daisy-chain topologies (for example, loop resiliency circuits) require limits on the number of devices for performance reasons. The performance impact is directly attributable to loop overhead traffic. In the past, recommendations for LRC topologies is 56 devices per loop. Advanced FC-AL topologies allow the “cost” of loop overheads to be minimized, thereby increasing the number of supported disk drives. This is true for system configurations that include switched hub architecture.

Switched hub architecture is a hub and spoke arrangement with local neighborhoods surrounding each device. Loop overhead is minimized because traffic no longer flows through each disk drive. The hubs are capable of local communication to the disk drives and then more efficiently conveying this information to the storage system. Two Fibre Channel ports are on each module. The PS/2 port is for IBM service only and provides no functionality. The units have LED status lights that indicate speed and fault status and are hot swappable allowing maximum availability.

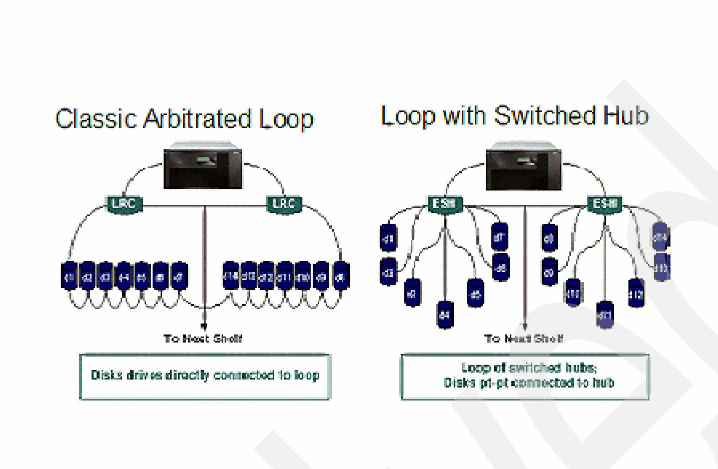

Figure 2-52 Switched hub architecture

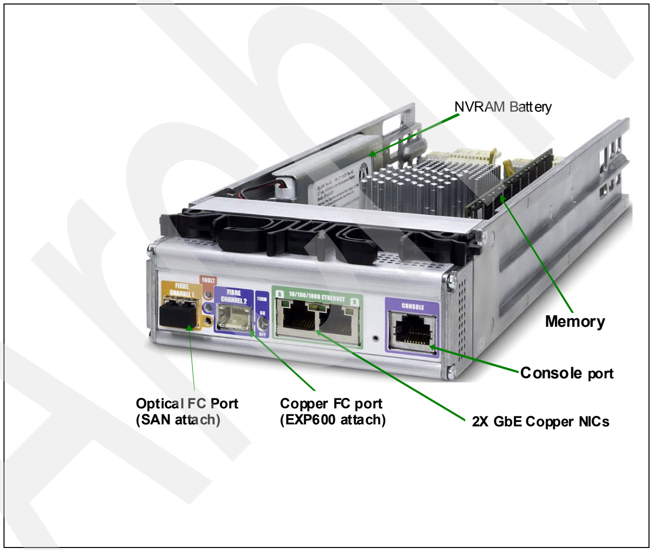

Figure 2-53 External ports on ESH2 module

The EXN2000 has been withdrawn from marketing as of May 22, 2007.

2.9.3 EXN1000

The EXN1000 uses the same shelf and hardware as the EXN2000 and EXN4000 so it has the same dimensions. It also supports the same number of disks per shelf (14); see Figure 2-54. The main differences are:

- Drive type: SATA versus Fibre Channel
- Interface module: the AT-FCX versus the ESH2

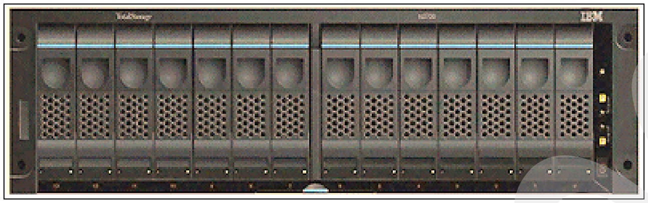

Figure 2-54 EXN1000 expansion unit

AT-FCX refers to the controller module (see Figure 2-55) of the SATA storage expansion unit.

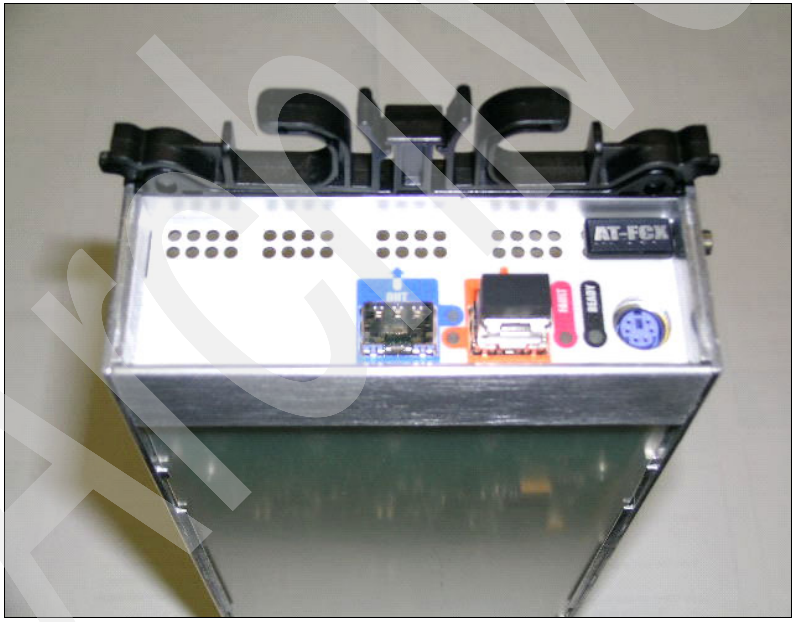

Figure 2-55 AT-FCX module

Data ONTAP supports up to 400 RAID groups per storage system or cluster. When configuring your aggregates, keep in mind that each aggregate requires at least one RAID group and that the total of all RAID groups in a storage system cannot exceed 400.

2.9.4 EXN4000

The EXN4000 uses the same shelf and hardware as the EXN2000 so the former has the same dimensions. EXN4000 also supports 14 disks per shelf, the same number as the EXN2000. EXN4000 uses ESH4 as its controller module. ESH4 refers to the third-generation, multiloop speed ESH module. ESH4 can function at 1 GB, 2 GB, 4 GB loop speed when it works with EXN4000. The ESH4 has LEDs that indicate whether the module is functioning normally (refer to Figure 2-58 on page 82) and whether any hardware problems exist. The LEDs also indicate the loop speed operation of the EXN4000. The main differences are as follows:

- A 4 Gbps capable Fibre Channel disk enclosure, that is, twice the maximum loop bandwidth of EXN2000
- Higher bandwidth for heavy sequential workload
- Fewer HBAs or slots used to achieve higher bandwidth needs

The EXN4000 FC storage expansion (see Figure 2-56 and Figure 2-57) unit runs at 2 Gbps FC when attached to systems that do not have 4 Gbps capability. It can be added to EXN2000 FC loops.

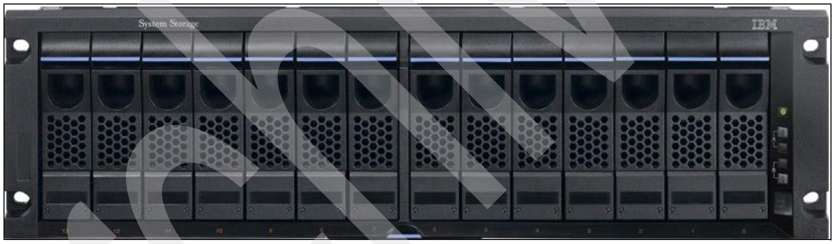

Figure 2-56 EXN4000 expansion unit

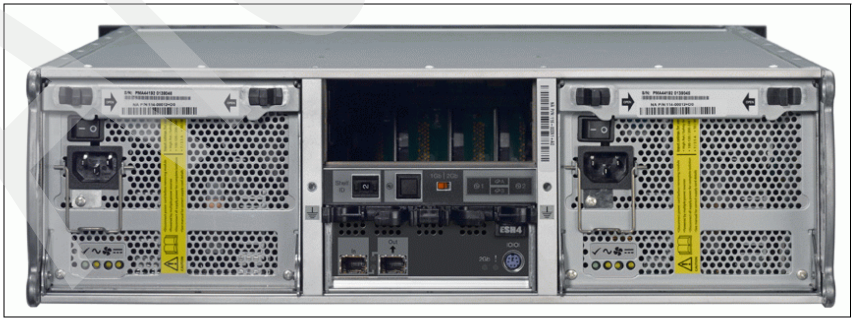

Figure 2-57 2xESH4, 2xPSU/fans

Figure 2-58 shows the location of the LEDs on the ESH4.

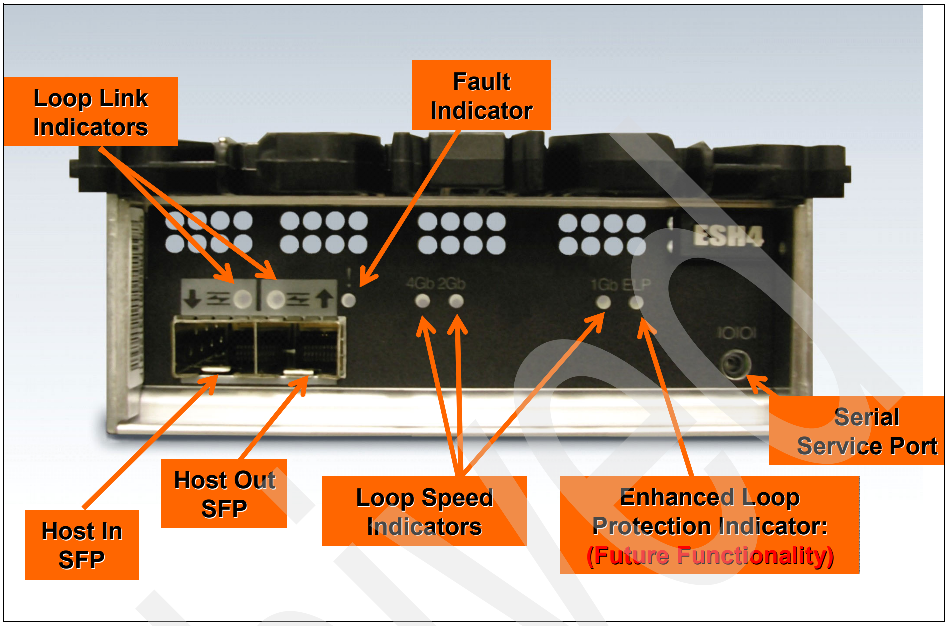

Figure 2-58 Location of the LEDs on an ESH4

EXN4000 is the replacement for the EXN2000 FC storage expansion unit.

3

Preparation

This chapter discusses several topics of interest, including the steps necessary to prepare your environment, while planning a VFM installation.

3.1 Understanding VFM components

VFM is split into several components, each with well-defined responsibilities, working together to provide this integrated solution. It is important to know each component's responsibilities in order to understand where each must run and to correctly plan your environment.

3.1.1 VFM server

The server centralizes the control of the VFM environment. It is responsible for storing and managing the configuration and distribution of all policies. It must be installed on a machine matching the requirements described in Chapter 4, "IBM VFM software and hardware requirements" on page 113.

3.1.2 VFM client

The client is the administrative GUI that administrators use to view and manage the namespace, policies, associated machines, shares, and reports (see Figure 3-1 on page 85).

You must always install the VFM server with a VFM client on the same machine, and you can also optionally install the VFM client on the administrator's workstations.

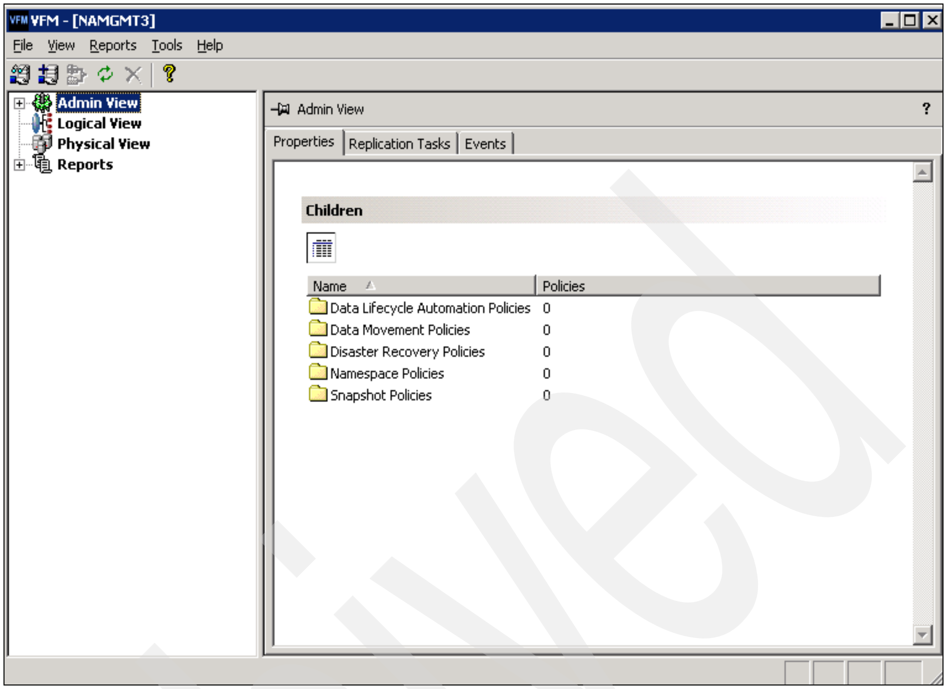

Name	Policies
Data Lifecycle Automation Policies	0
Data Movement Policies	0
Disaster Recovery Policies	0
Namespace Policies	0
Snapshot Policies	0

Figure 3-1 VFM client

3.1.3 VFM Replication Agent

The Replication Agents are the components responsible for moving data between two accessible CIFS shares or NFS exports. The Replication Agents perform the following functions:

- Validate the source and destination of a data movement policy
- Determine what must be moved as defined by the policy
- Transfer data
- Act in case of network failures to guarantee a successful operation
- Report the results and statistics to the VFM server

At the time of this publication, the Replication Agents available are for Microsoft Windows, Solaris 10, and RedHat Enterprise Linux V4. Windows Replication Agents can handle CIFS data transfers between shares, while UNIX and Linux replications agents can handle NFS data transfers between exports.

Replication Agents are also used to gather data for reports. Replication Agents are automatically deployed by the VFM server based on the existing policies. Replication Agents run under the context of the VFM service account.

3.1.4 VFM Monitoring Agent

The VFM Monitoring Agent performs the monitoring actions of the storage policies, as well as the actions associated with designated policies. The Monitoring Agent can be strategically located, based on data replication needs or credential considerations.

The best example of when the VFM Monitoring Agent must be used is when you have geographically distributed sites and need to maintain replicated data both locally and remotely. In this scenario, you can install the VFM server in one site and one Monitoring Agent on the other site, locally managing the designated policies.

3.1.5 VFM proxy

The proxy is used to communicate and control data transfer when a Replication Agent cannot or must not be deployed on the source or destination. N series storage and certain flavors of Linux and UNIX are examples of systems where a Replication Agent cannot be installed.

A good example for the necessity of a proxy is the migration of data between two N series machines, because Replication Agents cannot be installed on both sides of the replication.

3.2 Preparing your network environment

To ensure that each VFM component works properly, you must check and change some network-related items if they do not meet the requirements described in the sections that follow.

3.2.1 Ports

Each VFM component has to listen on TCP ports to properly communicate with the other components. The ports are not exactly the same for the different components. Ports and the components that use them are listed in Table 3-1.

Table 3-1 TCP ports used by each component

Component	Ports
Server	6001 and 6005
Replication agent	6002
Monitoring agent	6001 and 6005

Because two processes cannot use the same port at the same time, you must be sure that no other process is using these ports on the machine where you are about to install each component.

In the case of the Replication Agent, you must be sure that no other process is using the TCP port 6002 on all machines you plan to migrate to or from.

To determine which ports are already being used in a Microsoft Windows machine, complete the following steps:

1. Start a DOS command prompt (**Start** → **Programs** → **Accessories** → **Command Prompt**, or just **Start** → **Run** then type `cmd` and press **Enter**).
2. On the DOS prompt run **`netstat -nap tcp`**. A list of all used ports is displayed, as shown in Figure 3-2.

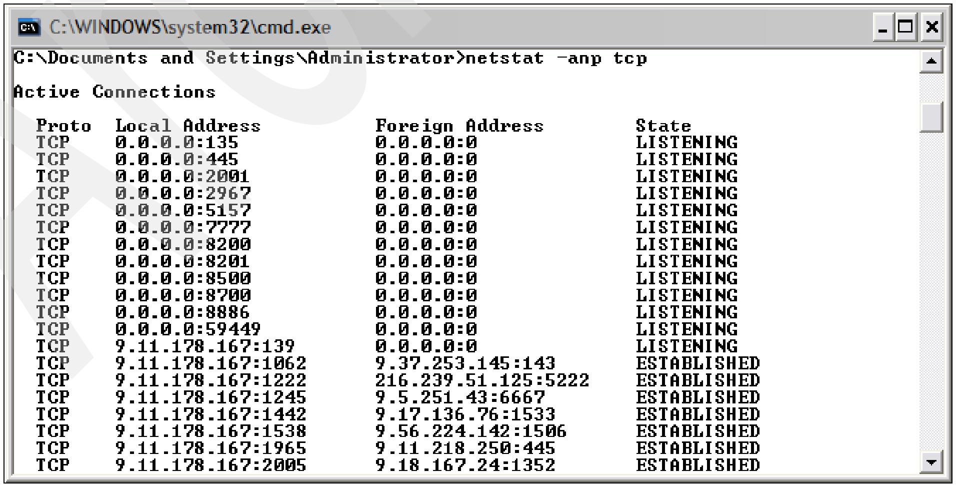

```
C:\Documents and Settings\Administrator>netstat -anp tcp

Active Connections

  Proto  Local Address          Foreign Address        State
  TCP    0.0.0.0:135            0.0.0.0:0              LISTENING
  TCP    0.0.0.0:445            0.0.0.0:0              LISTENING
  TCP    0.0.0.0:2001           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:2967           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:5157           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:7777           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:8200           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:8201           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:8500           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:8700           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:8886           0.0.0.0:0              LISTENING
  TCP    0.0.0.0:59449          0.0.0.0:0              LISTENING
  TCP    9.11.178.167:139       0.0.0.0:0              LISTENING
  TCP    9.11.178.167:1062      9.37.253.145:143       ESTABLISHED
  TCP    9.11.178.167:1222      216.239.51.125:5222    ESTABLISHED
  TCP    9.11.178.167:1245      9.5.251.43:6667        ESTABLISHED
  TCP    9.11.178.167:1442      9.17.136.76:1533       ESTABLISHED
  TCP    9.11.178.167:1538      9.56.224.142:1506      ESTABLISHED
  TCP    9.11.178.167:1965      9.11.218.250:445       ESTABLISHED
  TCP    9.11.178.167:2005      9.18.167.24:1352       ESTABLISHED
```

Figure 3-2 List of used TCP ports on Microsoft Windows

3. In the Local Address column in Figure 3-2 on page 87, the number after the colon is the port number. Check whether the component ports you are planning to install on that system are already in use.
4. If a component port is in use, you must stop the process that is using it.

> **Note:** The server and the Monitoring Agent cannot run on the same host because they use the same TCP ports.

3.2.2 DNS settings

VFM makes extensive use of DNS information about the other associated components. Because of this, you must verify that forward and reverse DNS zones are accurate for all systems where VFM components are to be deployed.

> **Note:** Incorrect DNS settings are one of the common causes of installation or operation failures.

You can use the command **`nslookup`** to do these checks. Its syntax is basically `nslookup <`*`hostname_to_check`*`>` to check forward zones (see Figure 3-3), or `nslookup <`*`ip_to_check`*`>` to check reverse zones (see Figure 3-4 on page 89).

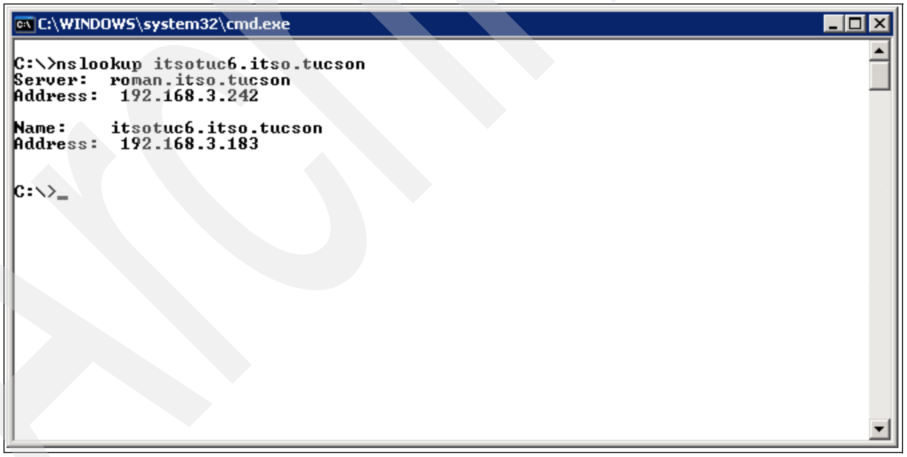

```
C:\>nslookup itsotuc6.itso.tucson
Server:  roman.itso.tucson
Address:  192.168.3.242

Name:    itsotuc6.itso.tucson
Address:  192.168.3.183


C:\>_
```

Figure 3-3 nslookup

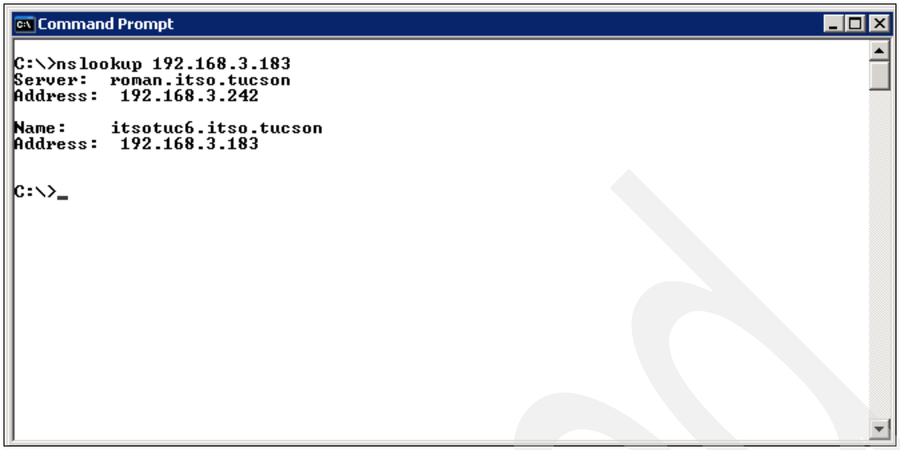

```
C:\>nslookup 192.168.3.183
Server:  roman.itso.tucson
Address:  192.168.3.242

Name:    itsotuc6.itso.tucson
Address:  192.168.3.183


C:\>_
```

Figure 3-4 reverse nslookup

Make sure you verify the accuracy of the forward and reverse DNS zones on all systems you are about to use with VFM as VFM server, console, Replication Agent, or Monitoring Agent. Several problems are expected in case DNS entries do not point correctly to each other's component addresses, and the logs do not always provide relevant information.

If you find any divergence between DNS information and current systems information, you must ensure that this divergence is resolved before deploying VFM.

3.3 Starting distributed file system service

Distributed file system (DFS™) service must run on the Microsoft Windows machines where VFM components run. You must verify that Windows is running and configured to automatically start on boot.

If DFS is not configured to run on the Windows machines, you can verify and fix it by going to **Start** → **Control Panel** → **Administrative Tools** → **Services** and complete the following steps:

1. Check whether the status of the DFS service is started and the startup type is Automatic (see Figure 3-5).

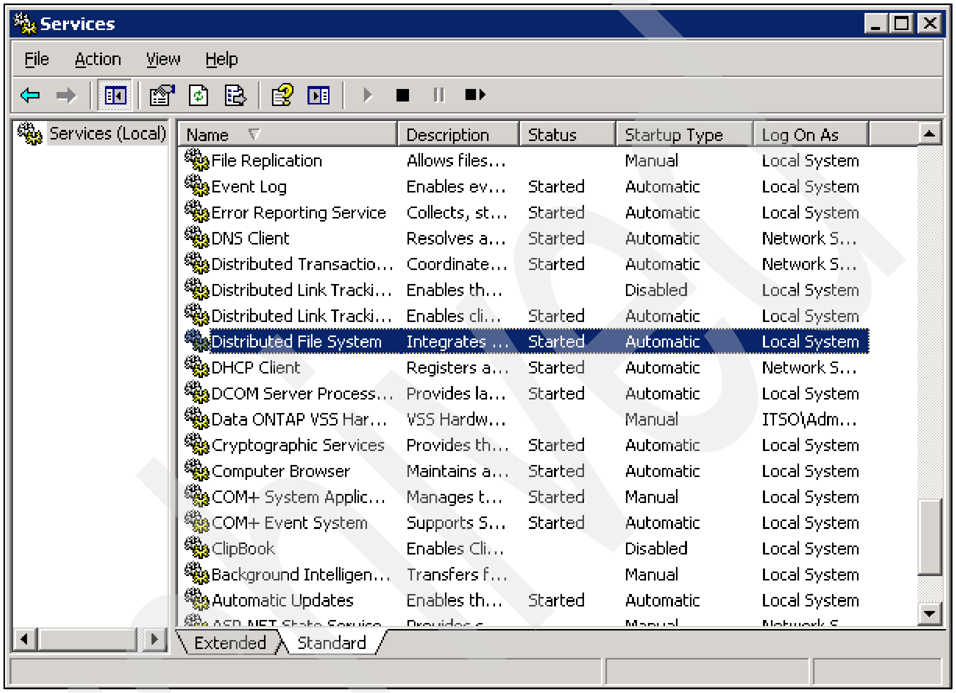

Name	Description	Status	Startup Type	Log On As
File Replication	Allows files...		Manual	Local System
Event Log	Enables ev...	Started	Automatic	Local System
Error Reporting Service	Collects, st...	Started	Automatic	Local System
DNS Client	Resolves a...	Started	Automatic	Network S...
Distributed Transactio...	Coordinate...	Started	Automatic	Network S...
Distributed Link Tracki...	Enables th...		Disabled	Local System
Distributed Link Tracki...	Enables cli...	Started	Automatic	Local System
Distributed File System	Integrates ...	Started	Automatic	Local System
DHCP Client	Registers a...	Started	Automatic	Network S...
DCOM Server Process...	Provides la...	Started	Automatic	Local System
Data ONTAP VSS Har...	VSS Hardw...		Manual	ITSO\Adm...
Cryptographic Services	Provides th...	Started	Automatic	Local System
Computer Browser	Maintains a...	Started	Automatic	Local System
COM+ System Applic...	Manages t...	Started	Manual	Local System
COM+ Event System	Supports S...	Started	Automatic	Local System
ClipBook	Enables Cli...		Disabled	Local System
Background Intelligen...	Transfers f...		Manual	Local System
Automatic Updates	Enables th...	Started	Automatic	Local System
ASP.NET State Service	Provides s...		Manual	Network S...

Figure 3-5 Windows Services list

2. If one of these properties differ, select the DFS service, right-click it, and then select **Properties** (see Figure 3-6).

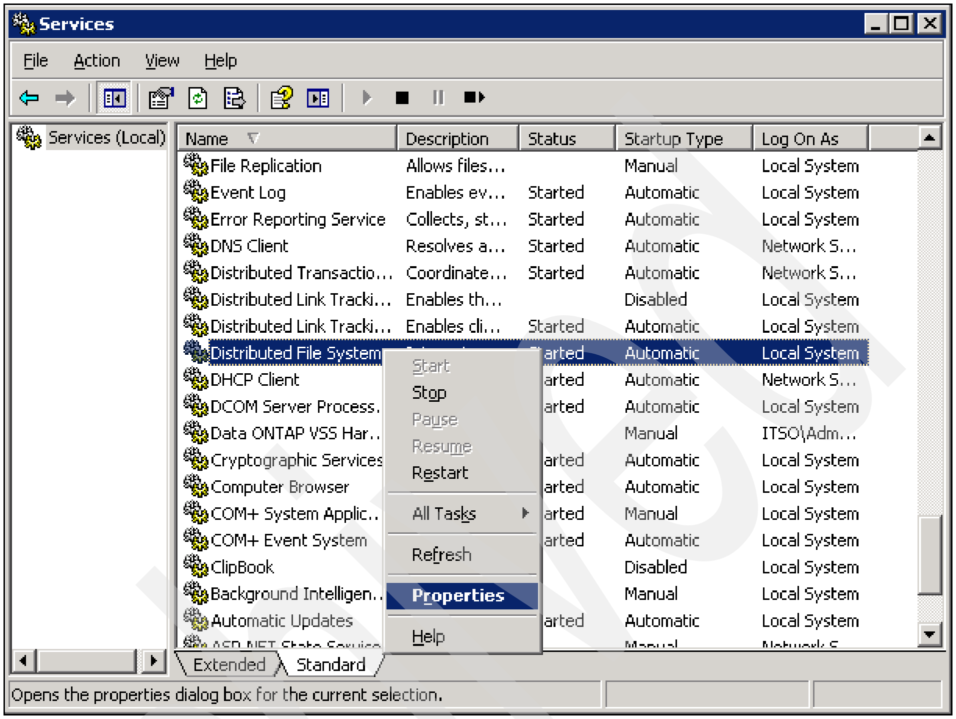

Name	Description	Status	Startup Type	Log On As
File Replication	Allows files...		Manual	Local System
Event Log	Enables ev...	Started	Automatic	Local System
Error Reporting Service	Collects, st...	Started	Automatic	Local System
DNS Client	Resolves a...	Started	Automatic	Network S...
Distributed Transactio...	Coordinate...	Started	Automatic	Network S...
Distributed Link Tracki...	Enables th...		Disabled	Local System
Distributed Link Tracki...	Enables cli...	Started	Automatic	Local System
Distributed File System		arted	Automatic	Local System
DHCP Client		arted	Automatic	Network S...
DCOM Server Process...		arted	Automatic	Local System
Data ONTAP VSS Har...			Manual	ITSO\Adm...
Cryptographic Services		arted	Automatic	Local System
Computer Browser		arted	Automatic	Local System
COM+ System Applic...		arted	Manual	Local System
COM+ Event System		arted	Automatic	Local System
ClipBook			Disabled	Local System
Background Intelligen...			Manual	Local System
Automatic Updates		arted	Automatic	Local System

Opens the properties dialog box for the current selection.

Figure 3-6 Windows Services list: right-click Distributed File System

3. In the Properties dialog box, you can start the service and configure it to start automatically (see Figure 3-7).

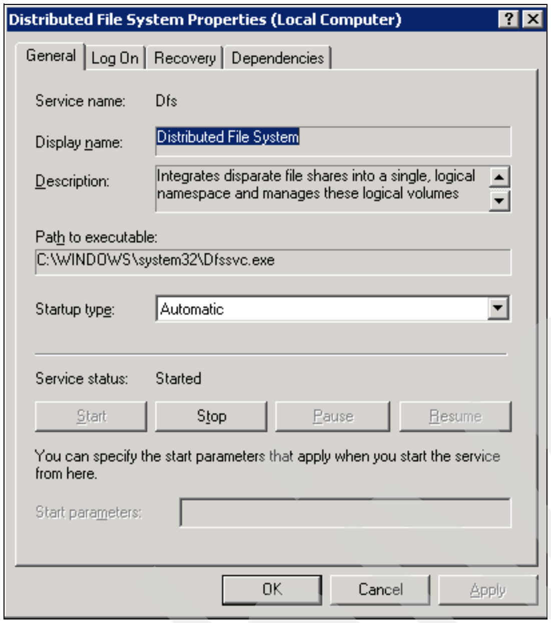

Distributed File System Properties (Local Computer)

Service name: Dfs

Display name: Distributed File System

Description: Integrates disparate file shares into a single, logical namespace and manages these logical volumes

Path to executable:
C:\WINDOWS\system32\Dfssvc.exe

Startup type: Automatic

Service status: Started

You can specify the start parameters that apply when you start the service from here.

Start parameters:

Figure 3-7 Distributed File System Properties

3.4 Creating a VFM service account

Because VFM has to handle system properties and data, its components require privileged access on the systems where they are located.

> **Note:** Regardless of the user using the VFM client, all VFM operations are performed using the VFM service account.

The access level varies based on the components, and you must identify accounts with the required privileges or create new ones.

The default domain administrators have all required privileges, but it must not be directly used as the VFM service account because its password might have to be changed. Changing the password causes unnecessary work to update the password on all VFM components already deployed.

These permissions and privileges are required:

- The account used by the VFM server, a Replication Agent, or a VFM Monitoring Agent must have the following permissions:
  - **SeBackupPrivilege** (“Back up files and directories”)

    Required to open source files and directories with backup semantics.
  - **SeRestorePrivilege** (“Restore files and directories”)

    Required to open destination files and directories 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 files and directories.
  - Appropriate permissions to administer the DFS roots that are managed. 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.
- Administrator privileges on any machines hosting either source or destination of data replications
- The service account must be run as service privilege.
- The account running the VFM client must have permissions to write to the All Users\Application Data\ directory tree. Users whose accounts lack write permissions to these directories experience errors when attempting to run VFM reports or set VFM system options.

We recommend the following steps to create the user on a domain for VFM service account and give it all necessary privileges:

1. Log on to your domain controller. Each of the following steps must be executed there.
2. Create the user:
   a. Go to **Start** → **Programs** → **Administrative Tools** → **Active Directory Users and Computers** (see Figure 3-8).

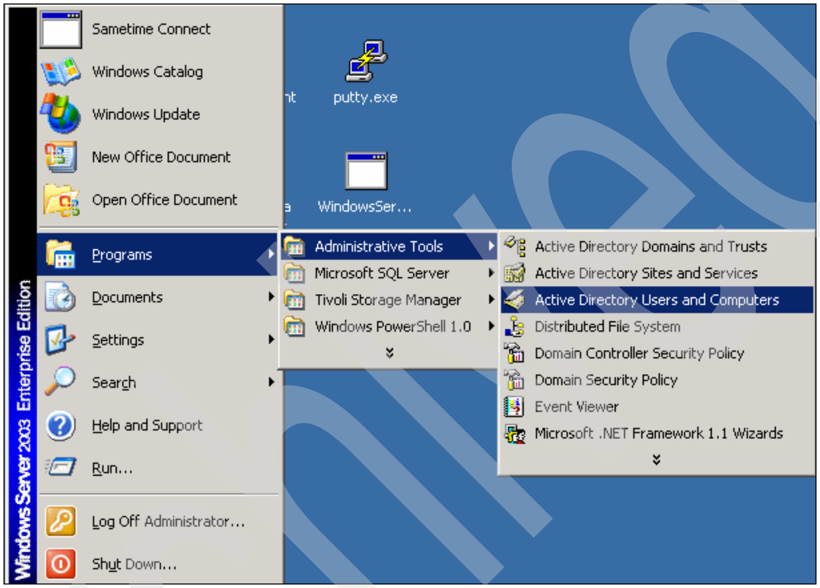

Figure 3-8 Windows menu to Active Directory Users and Computers

A window is displayed (see Figure 3-9) that enables you to manage the users and groups on your domain.

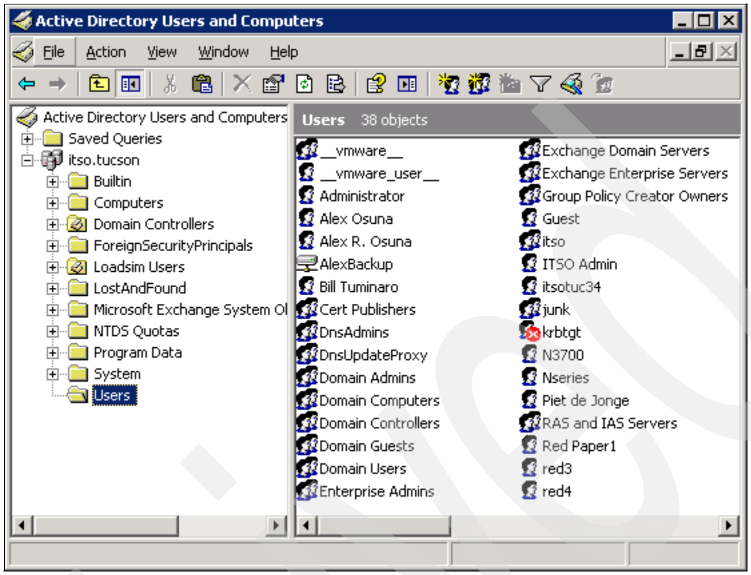

- Active Directory Users and Computers
  - Saved Queries
  - itso.tucson
    - Builtin
    - Computers
    - Domain Controllers
    - ForeignSecurityPrincipals
    - Loadsim Users
    - LostAndFound
    - Microsoft Exchange System Ol
    - NTDS Quotas
    - Program Data
    - System
    - Users

Users 38 objects

__vmware__
__vmware_user__
Administrator
Alex Osuna
Alex R. Osuna
AlexBackup
Bill Tuminaro
Cert Publishers
DnsAdmins
DnsUpdateProxy
Domain Admins
Domain Computers
Domain Controllers
Domain Guests
Domain Users
Enterprise Admins
Exchange Domain Servers
Exchange Enterprise Servers
Group Policy Creator Owners
Guest
itso
ITSO Admin
itsotuc34
junk
krbtgt
N3700
Nseries
Piet de Jonge
RAS and IAS Servers
Red Paper1
red3
red4

Figure 3-9 Active Directory Users and Computers: Users window

b. On the Microsoft Windows menu, go to **Action** → **New** → **User** (see Figure 3-10).

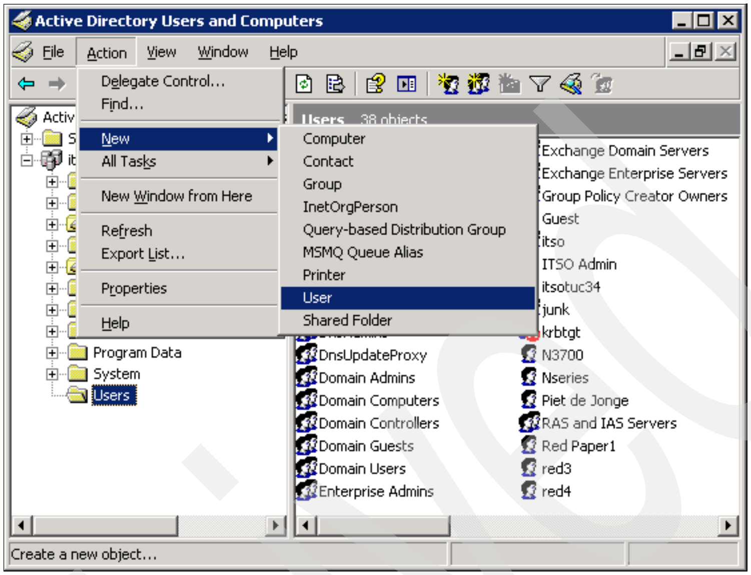

Create a new object...

Figure 3-10 Active Directory Users and Computer: new user menu

c. Fill in the dialogs, using your domain information and the user name `ibmvfm` (see Figure 3-11).

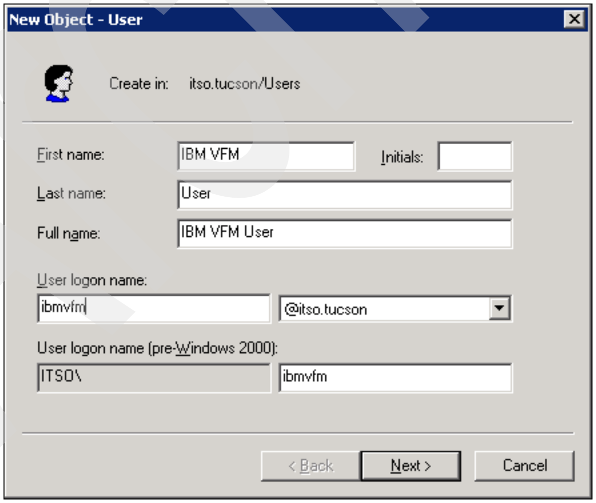

Create in: itso.tucson/Users

First name: IBM VFM Initials:

Last name: User

Full name: IBM VFM User

User logon name:
ibmvfm @itso.tucson

User logon name (pre-Windows 2000):
ITSO\ ibmvfm

Figure 3-11 New user dialog: basic information

d. Then check the **Password never expires** check box (see Figure 3-12).

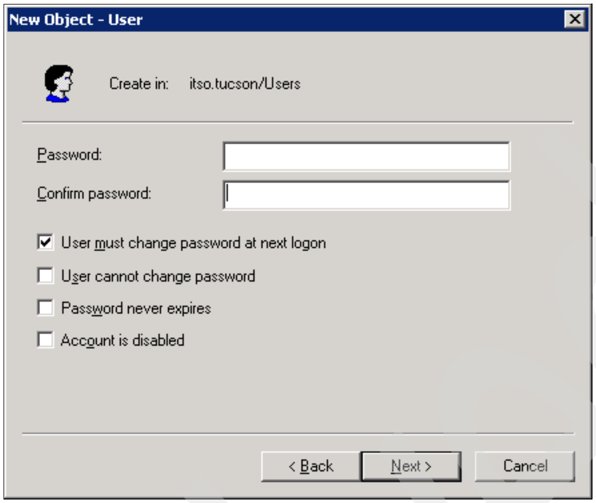

New Object - User

Create in: itso.tucson/Users

Password:

Confirm password:

☑ User must change password at next logon

□ User cannot change password

□ Password never expires

□ Account is disabled

Figure 3-12 New user dialog: password with default options

e. A message box is displayed (see Figure 3-13), indicating that the password will never expire. Click **OK**.

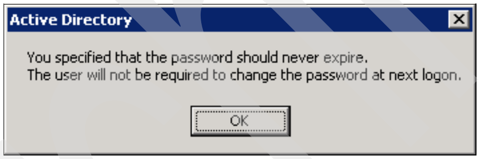

Active Directory

You specified that the password should never expire.
The user will not be required to change the password at next logon.

Figure 3-13 New user dialog: password never expires

f. Type the password in the specified fields and click **Next** (see Figure 3-14).

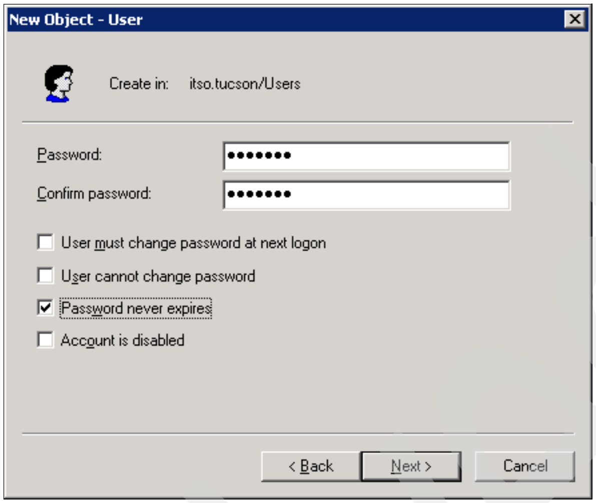

Create in: itso.tucson/Users

Password: •••••••

Confirm password: •••••••

□ User must change password at next logon

□ User cannot change password

☑ Password never expires

□ Account is disabled

Figure 3-14 New user dialog: password with never expire option

g. Click **Finish** (see Figure 3-15).

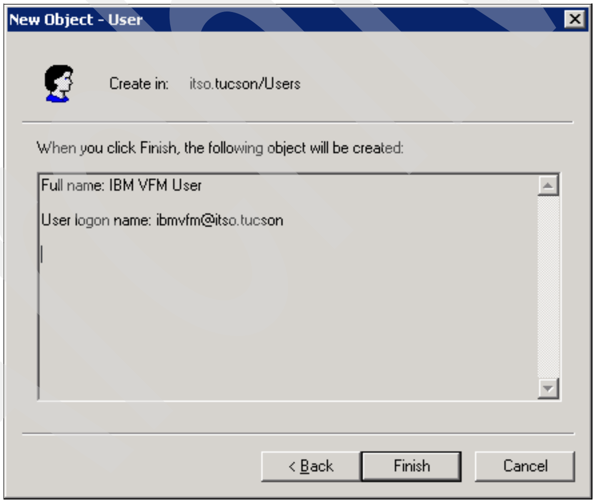

Create in: itso.tucson/Users

When you click Finish, the following object will be created:

Full name: IBM VFM User

User logon name: ibmvfm@itso.tucson

Figure 3-15 New user dialog: summary

The user is created, as shown on Figure 3-16.

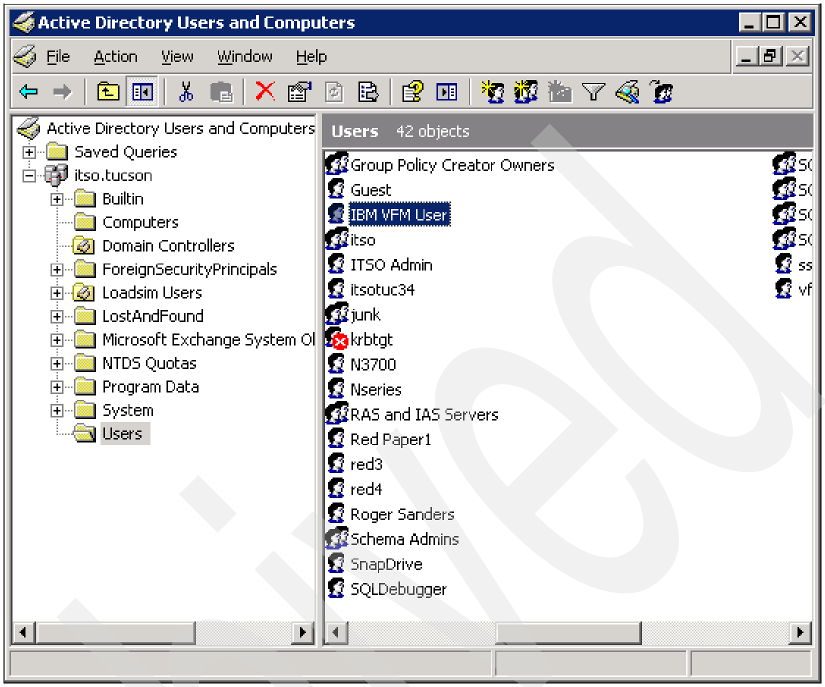

Active Directory Users and Computers
- Saved Queries
- itso.tucson
  - Builtin
  - Computers
  - Domain Controllers
  - ForeignSecurityPrincipals
  - Loadsim Users
  - LostAndFound
  - Microsoft Exchange System Ol
  - NTDS Quotas
  - Program Data
  - System
  - Users

Users 42 objects

Group Policy Creator Owners	SC
Guest	SC
IBM VFM User	SC
itso	SC
ITSO Admin	ss
itsotuc34	vf
junk	
krbtgt	
N3700	
Nseries	
RAS and IAS Servers	
Red Paper1	
red3	
red4	
Roger Sanders	
Schema Admins	
SnapDrive	
SQLDebugger	

Figure 3-16 Users list showing the alternatively created VFM user

3. Add user to the Domain Admins group:

   a. Right-click the user and go to **Properties** (see Figure 3-17).

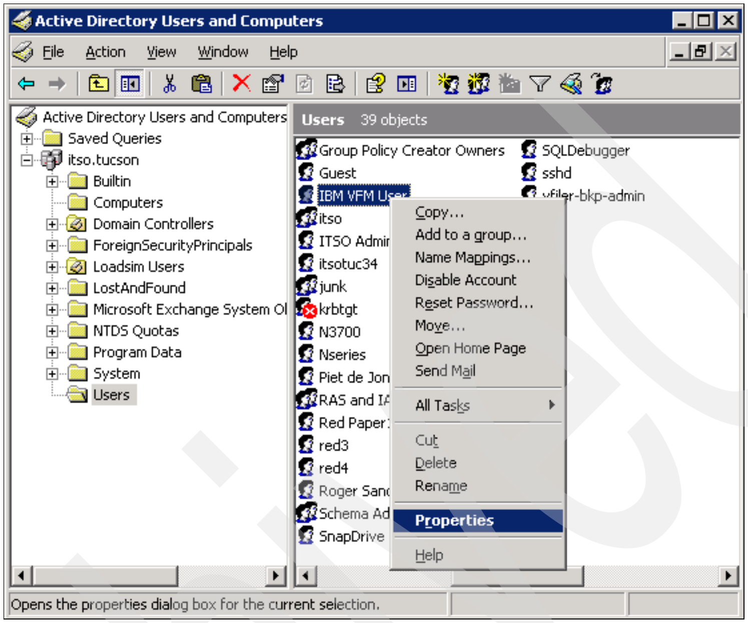

Opens the properties dialog box for the current selection.

Figure 3-17 VFM User: click Properties

b. Click the **Member Of** tab (see Figure 3-18).

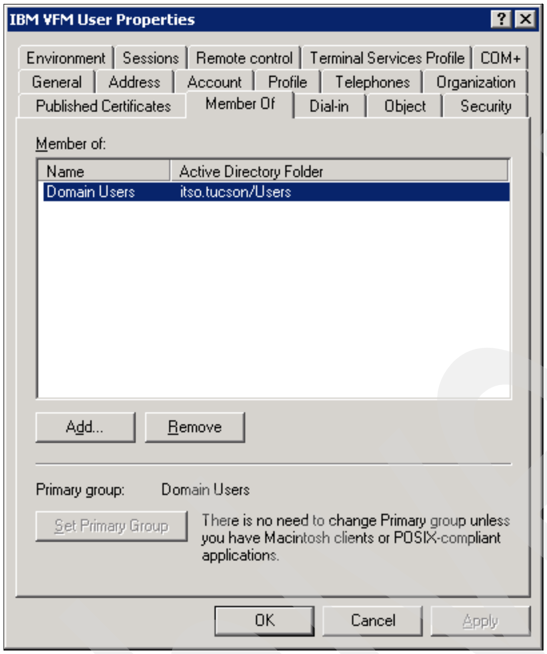

Member of:

Name	Active Directory Folder
Domain Users	itso.tucson/Users

Primary group: Domain Users

There is no need to change Primary group unless you have Macintosh clients or POSIX-compliant applications.

Figure 3-18 VFM User: original group list

c. Click the **Add** button and fill the Domain Admins file. Click the **Check Names** button, and then click **OK** (see Figure 3-19).

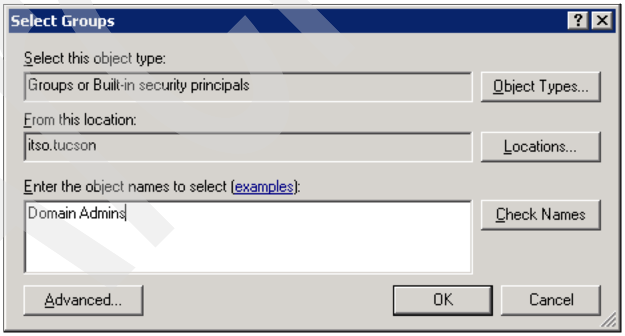

Select this object type: Groups or Built-in security principals

From this location: itso.tucson

Enter the object names to select (examples): Domain Admins

Figure 3-19 VFM User: Select Groups dialog

d. The new Groups list is displayed. Check that the new Domain Admins group is listed and click **OK** (see Figure 3-20).

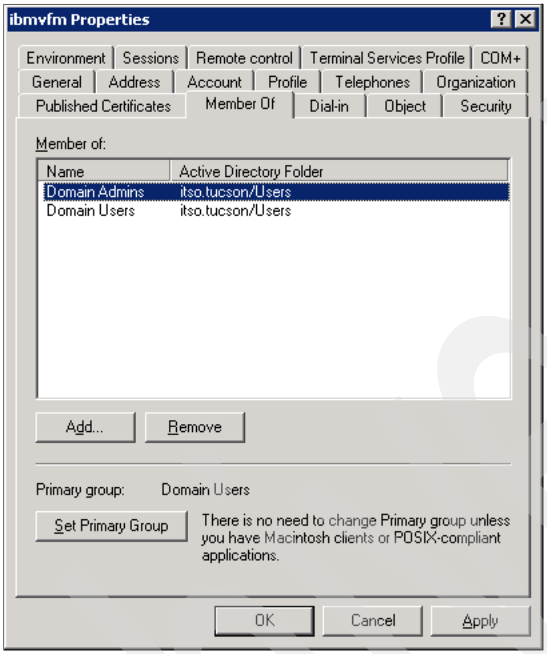

Member of:

Name	Active Directory Folder
Domain Admins	itso.tucson/Users
Domain Users	itso.tucson/Users

Primary group: Domain Users

There is no need to change Primary group unless you have Macintosh clients or POSIX-compliant applications.

Figure 3-20 VFM User: updated groups list

4. Provide the necessary privileges to the new user. Some privileges must be explicitly assigned to a user, and the VFM user requires five of them. They are:

   – Back up files and directories
   – Restore files and directories
   – Take ownership of files or other objects
   – Manage auditing and security log
   – Log on as a service

The process for setting one is exactly the same for setting all the others. We provide one example, and you can repeat it for the others.

Follow these steps to set privileges:

a. Go to **Start** → **Programs** → **Administrative Tools** → **Domain Security Policy**. On the left panel, choose **Local Policies** → **Users Rights assignment** (see Figure 3-21).

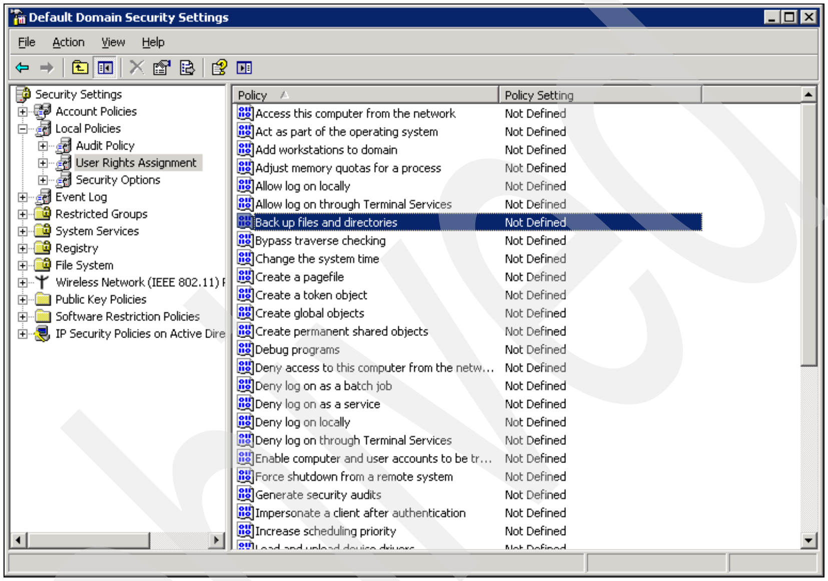

Policy	Policy Setting
Access this computer from the network	Not Defined
Act as part of the operating system	Not Defined
Add workstations to domain	Not Defined
Adjust memory quotas for a process	Not Defined
Allow log on locally	Not Defined
Allow log on through Terminal Services	Not Defined
Back up files and directories	Not Defined
Bypass traverse checking	Not Defined
Change the system time	Not Defined
Create a pagefile	Not Defined
Create a token object	Not Defined
Create global objects	Not Defined
Create permanent shared objects	Not Defined
Debug programs	Not Defined
Deny access to this computer from the netw…	Not Defined
Deny log on as a batch job	Not Defined
Deny log on as a service	Not Defined
Deny log on locally	Not Defined
Deny log on through Terminal Services	Not Defined
Enable computer and user accounts to be tr…	Not Defined
Force shutdown from a remote system	Not Defined
Generate security audits	Not Defined
Impersonate a client after authentication	Not Defined
Increase scheduling priority	Not Defined
Load and unload device drivers	Not Defined

Figure 3-21 Default Domain Security Settings window

b. Double-click the privilege you want to assign, for example, Back up files and directories (see Figure 3-22).

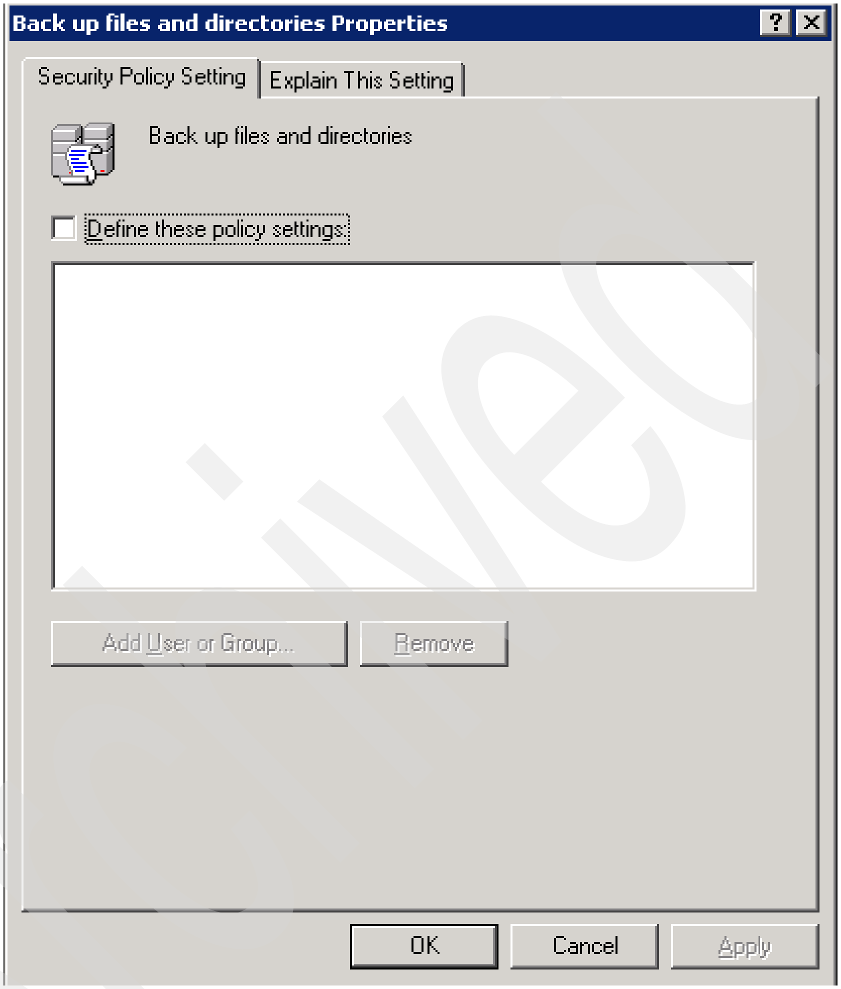

Back up files and directories Properties

Back up files and directories

□ Define these policy settings:

Figure 3-22 Privilege properties

c. Check the **Define these policy settings** check box and then click **Add User or Group**.

5. Add the user <*DOMAIN*>\ibmvfm and click **OK** (see Figure 3-23).

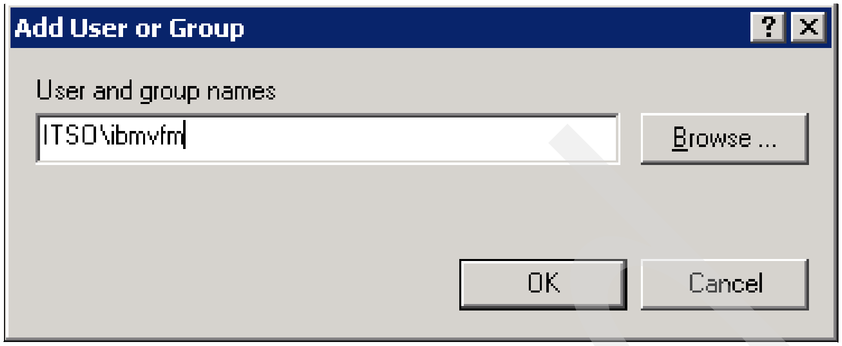

User and group names
ITSO\ibmvfm

Figure 3-23 Adding user to privilege

d. Check that the user was correctly added to the list and click **OK** (see Figure 3-24).

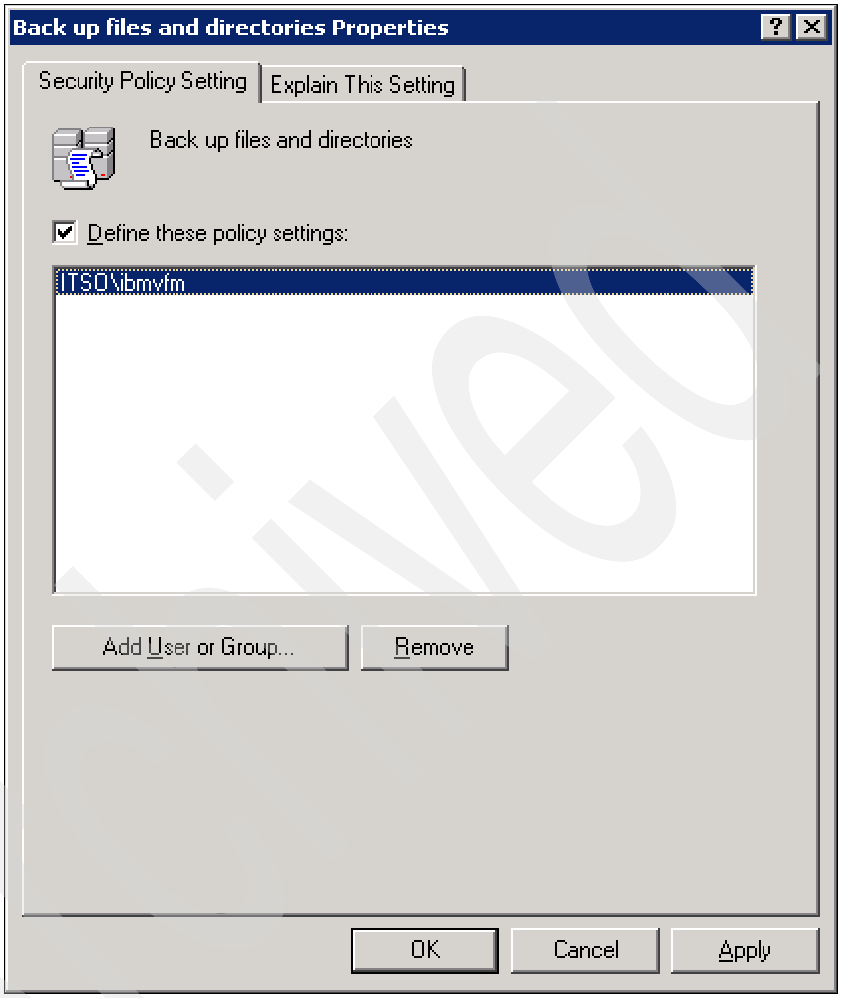

Back up files and directories

☑ Define these policy settings:

ITSO\ibmvfm

Figure 3-24 Privilege properties with ibmvfm user

e. Repeat the process in assigning the next privilege.

3.5 MSSQL

VFM needs MSSQL in order to operate. If you do not have a MSSQL server installed, VFM can install the desktop version (MSDE) for you.

MSDE has a 2 GB database limitation, and 2 GBs is sufficient for most of VFM use cases.

In this book, we show how to install the SQL server during VFM installation. If you plan to make extensive use of reports, you can consider installing a separate full MSSQL version, but that process is not covered here.

3.6 Preparing your N series storage system

VFM uses RSH to manage N series storage. Therefore you must make sure that RSH is enabled on the N series storage system you intend to manage with VFM.

You can perform this check or change the value by performing the following steps:

1. Using your Web browser, go to the Web administration of your filer (that is, http://<*filer_IP*>/na_admin/. You are asked for the root/administrator password (see Figure 3-25).

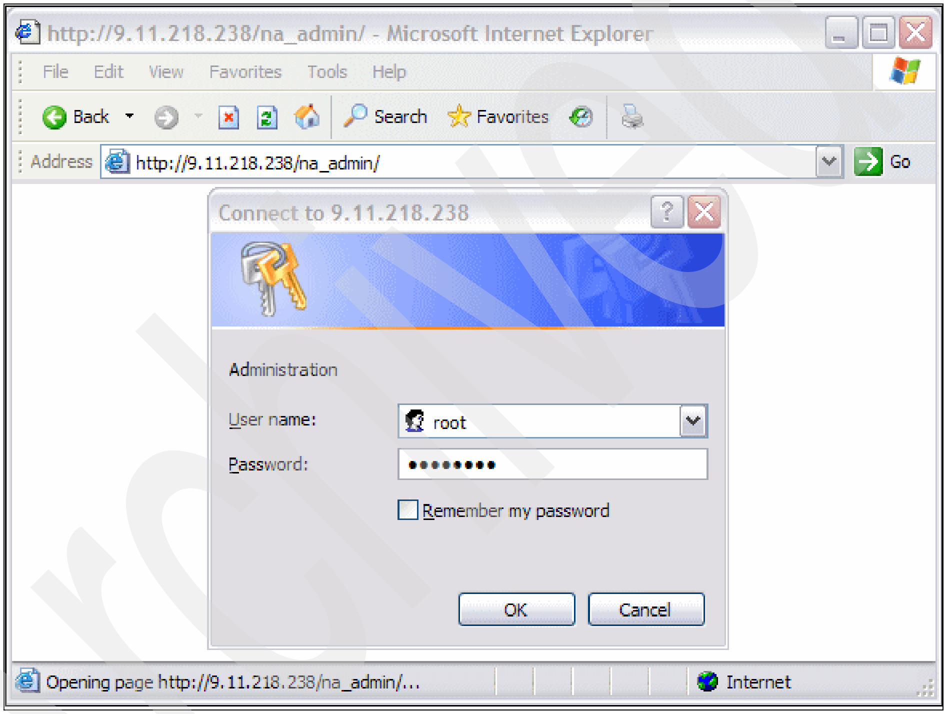
http://9.11.218.238/na_admin/ - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back
Search
Favorites
Address
http://9.11.218.238/na_admin/
Go
Connect to 9.11.218.238
OK
Cancel
Opening page http://9.11.218.238/na_admin/...
Internet

Administration

User name: root

Password: ••••••••

□ Remember my password

Figure 3-25 N series logon window

2. In the next window, click **FilerView** to go to the administration window (see Figure 3-26).

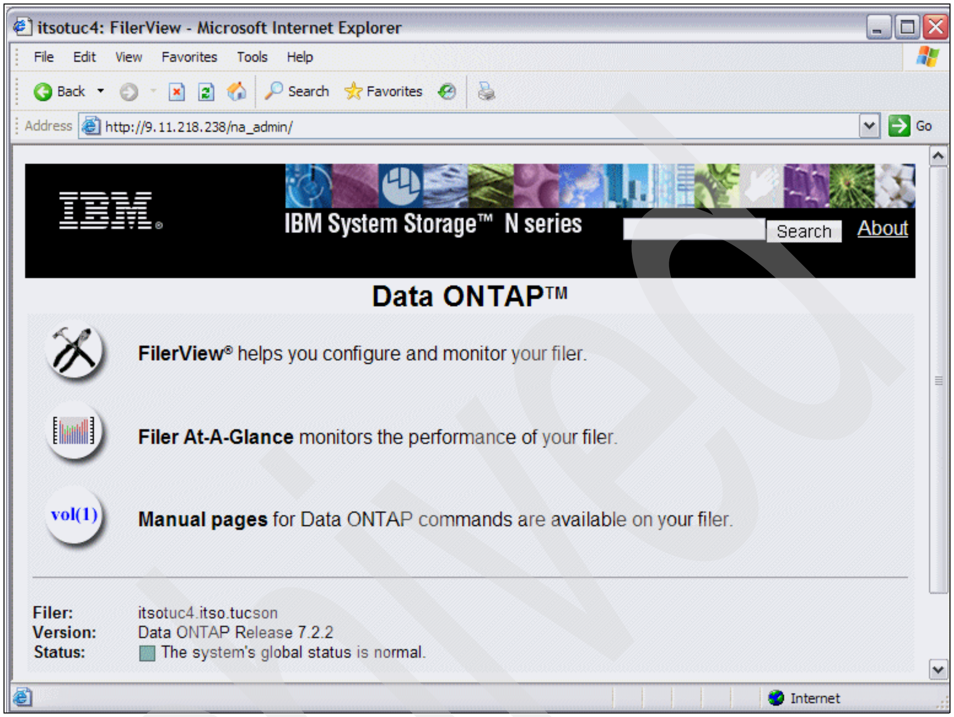
itsotuc4: FilerView - Microsoft Internet Explorer
File
Edit
View
Favorites
Tools
Help
Back
Search
Favorites
Address
http://9.11.218.238/na_admin/
Go
IBM
IBM System Storage™ N series
Search
About
Internet

Data ONTAP™

FilerView® helps you configure and monitor your filer.

Filer At-A-Glance monitors the performance of your filer.

Manual pages for Data ONTAP commands are available on your filer.

Filer: itsotuc4.itso.tucson
Version: Data ONTAP Release 7.2.2
Status: 🟩 The system's global status is normal.

Figure 3-26 N series first window

3. In the FilerView, select **Network** → **Configure** (see Figure 3-27).

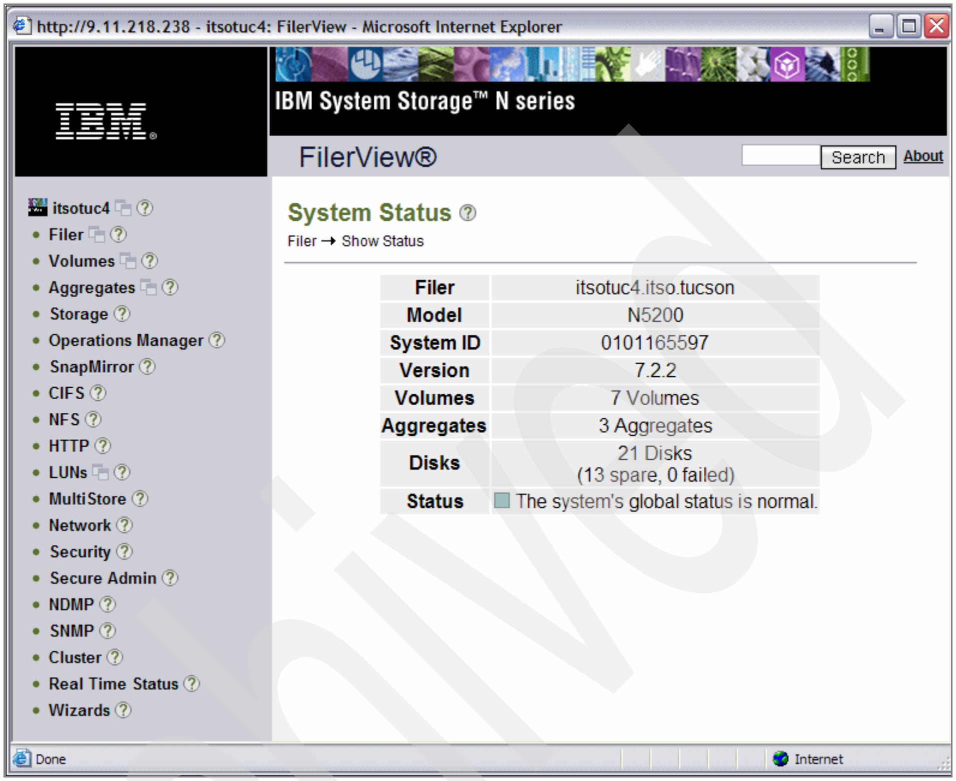

System Status

Filer → Show Status

Filer	itsotuc4.itso.tucson
Model	N5200
System ID	0101165597
Version	7.2.2
Volumes	7 Volumes
Aggregates	3 Aggregates
Disks	21 Disks (13 spare, 0 failed)
Status	🟩 The system's global status is normal.

Figure 3-27 FilerView

4. Check whether RSH is enabled. Apply the new settings (see Figure 3-28).

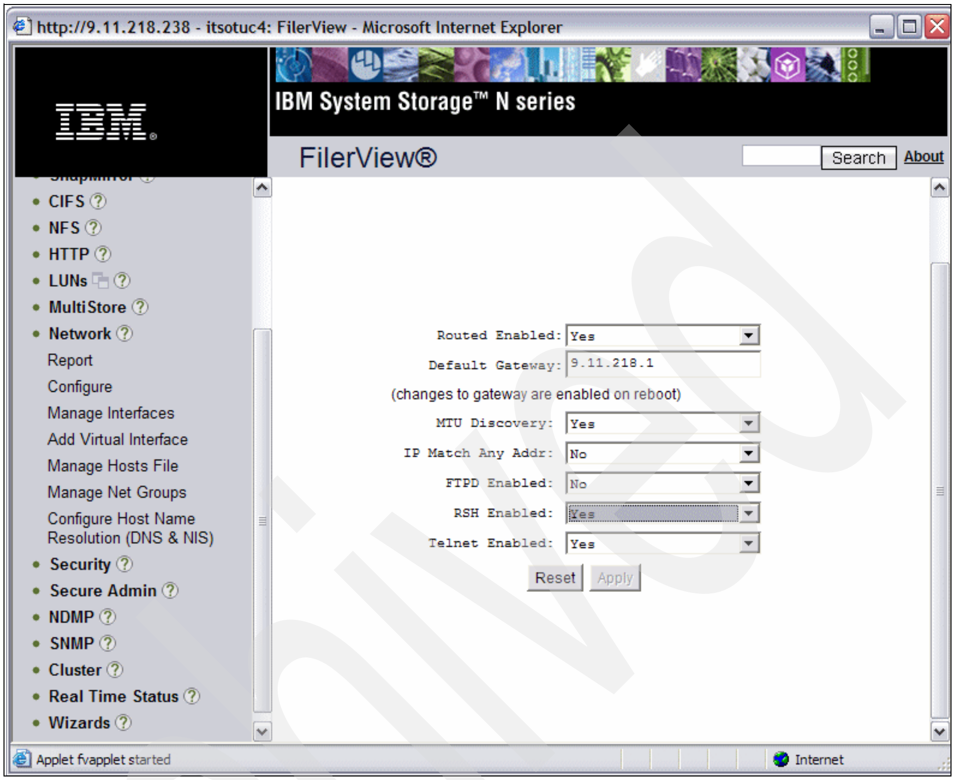

Routed Enabled: Yes

Default Gateway: 9.11.218.1

(changes to gateway are enabled on reboot)

MTU Discovery: Yes

IP Match Any Addr: No

FTPD Enabled: No

RSH Enabled: Yes

Telnet Enabled: Yes

Figure 3-28 N series network configuration

4

IBM VFM software and hardware requirements

This chapter describes the minimum system requirements for the Virtual File Manager (VFM) Version 6.0.0 server and client as well as the Replication Agent console. We discuss the following topics:

- Operating system requirements for:
  - VFM server or Monitoring Agent
  - VFM client
- Database engine (local or remote)
- Replication Agent
- Required software
- Recommended hardware

4.1 Operating system

Figure 4-1 depicts VFM installation supported systems.

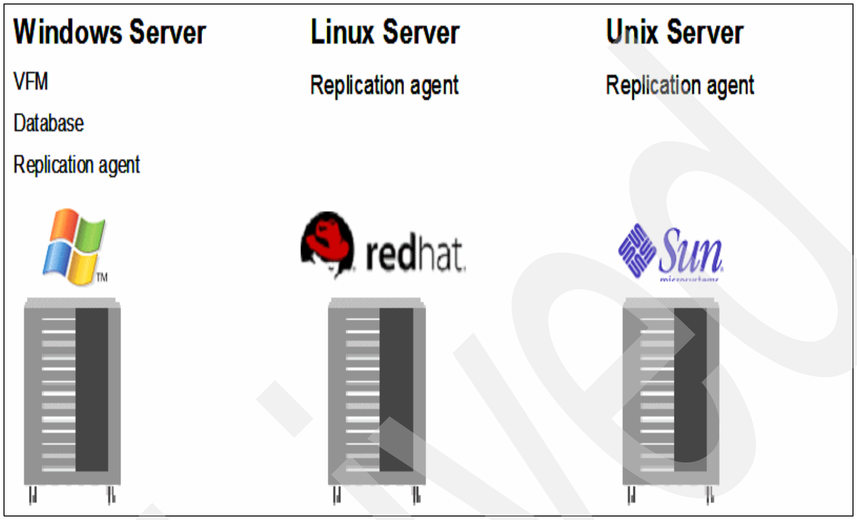

Figure 4-1 VFM installation supported systems

VFM installation supports the following systems:

- VFM server or Monitoring Agent
  - Microsoft Windows 2000 SP4.
  - Microsoft Windows 2003 Enterprise Server 2003 SP1, R2 (32 bit).
  - UNIX data movement agent can be installed on Solaris 10 and Red Hat Enterprise V4.
- VFM client
  - Microsoft Windows 2000 SP2, SP3, SP4 or Windows XP Professional, SP1, SP2
  - Microsoft XP Professional, SP1, SP2 or Windows Server® 2003, SP1, R2 to manage multiple roots on a single Windows Server 2003

4.2 Database engine (local or remote)

The database engine is supported on the following MSSQL versions.

- Microsoft SQL Server 2000 desktop engine (MSDE) SP3 (MSDE is provided with VFM.)
- Microsoft SQL Server 2000 Standard or Enterprise Edition SP3, SP4
- Microsoft SQL Server 2005

4.3 Replication Agent

The replication agent supports the following platforms.

- Microsoft Windows 2000 SP4 (Microsoft.NET Framework 2.0)
- Microsoft Windows 2003 Enterprise Server 2003 SP1, R2 (32 bit)
- Red Hat Enterprise Linux 4.0
- Solaris 10

4.4 Required software

The following software is required:

- Microsoft.NET Framework V2.0
- Internet Explorer® V5.5, V6, or V7
- Windows Script Host V5.6

4.5 Recommended hardware

The following hardware is recommended:

- 2 GB RAM or better
- 2 GHz or faster processor
- Disk space: 575 MB
- Minimum video display of 1024 x 768 and 256 colors

5

DFS namespace

This chapter describes the concepts of the DFS namespace and the steps to create a DFS namespace with VFM. The chapter is divided into the following topics:

- Namespace overview
  - Distributed file system (DFS) overview
  - Global namespace with VFM
- Creating DFS namespace with VFM
  - Domain-based DFS root
  - Standalone DFS root
  - Consolidation DFS root
  - Creating DFS root
  - Adding folders to create a hierarchy
  - Link overview
  - Adding a link
  - Adding a second link
  - Client view of the namespace

5.1 Namespace overview

This section provides an overview of the distributed file system (DFS) and the global namespace with VFM. In addition, we discuss their functionality as a storage management solution.

5.1.1 Distributed file system overview

One of the goals of most information technology (IT) groups is to manage file server resources efficiently while keeping them available and secure for users. As networks expand to include more users and servers, whether they are located in one site or in geographically distributed sites, administrators find it increasingly difficult to keep users connected to the files they need. On the one hand, distributing resources across a network makes them more available to more people and promotes cross-organizational efforts. Alternatively, storing files on different file servers located throughout an organization makes it difficult for users to know where to look for information. Administrators also find it difficult to keep track of all the servers and all of the people who use those servers. The task of swapping out an old server becomes a major communication chore when users across an organization must be notified to update links and file paths.

To help administrators address these problems, the distributed file system (DFS) was included in Microsoft Windows Server 2000 and Windows Server 2003. DFS enables administrators to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces.

A DFS namespace is a virtual view of shared folders in an organization. Using the DFS tools, you select which shared folders to present in the namespace, design the hierarchy in which those folders appear, and determine the names the shared folders show in the namespace. When a user views the namespace, the folders appear to reside on a single, high-capacity hard disk. Users can navigate the namespace without needing to know the server names or shared folders hosting the data.

DFS provides other benefits, including the following ones:

- Simplified data migration
- Increased availability of file server data
- Load sharing
- Security integration

VFM leverages the Microsoft Windows DFS namespace to enable transparent movement of files without affecting user access.

To learn more about DFS, see the following recommended Web site:

http://www.microsoft.com/dfs

5.1.2 Global namespace with VFM

A global namespace is a storage management solution that gives you a more flexible way to centrally manage your distributed resources. You can create simplified views of folders and files, called a *namespace*, regardless of where those files physically reside in a network. This concept is powerful because it means you can use a namespace to logically arrange and present data to users, irrespective of where the data is located. It also enables you to add, change, move, and reconfigure physical file storage without affecting how users view and access it.

A namespace is a means of pooling multiple file systems into a single, global file system. A global namespace can pool storage from multiple, heterogeneous storage types (DAS, SAN, or N series) and across different storage platforms (Microsoft Windows, Linux, and UNIX). (See Figure 5-1.)

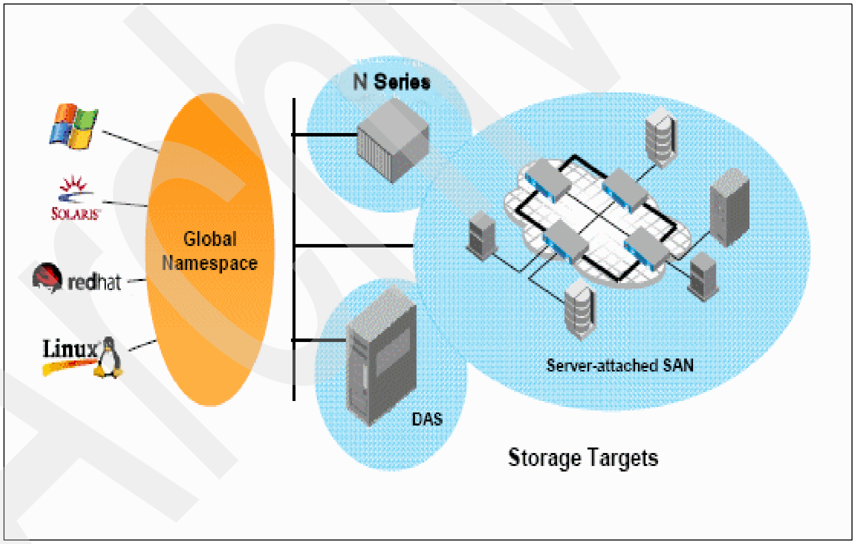

Figure 5-1 Global namespace view

With a global namespace in place, you can distribute files in a way that achieves the best performance and capacity utilization and enable clients to access them through the logical namespace. When storage is added or consolidated and files are moved or renamed, clients are automatically redirected to the files in their new location without ever having to know that they were moved. And most important, it permanently eliminates any need for desktop reconfiguration, drive letter remapping, or logon script modification when storage is reconfigured.

A common issue faced today is file access and complexity management. For instance, data that departmental groups require is spread across multiple shares and exports on multiple servers. Without a global namespace in place, users must physically map to each server and share, or mount, each export in order to access their data (see Figure 5-2).

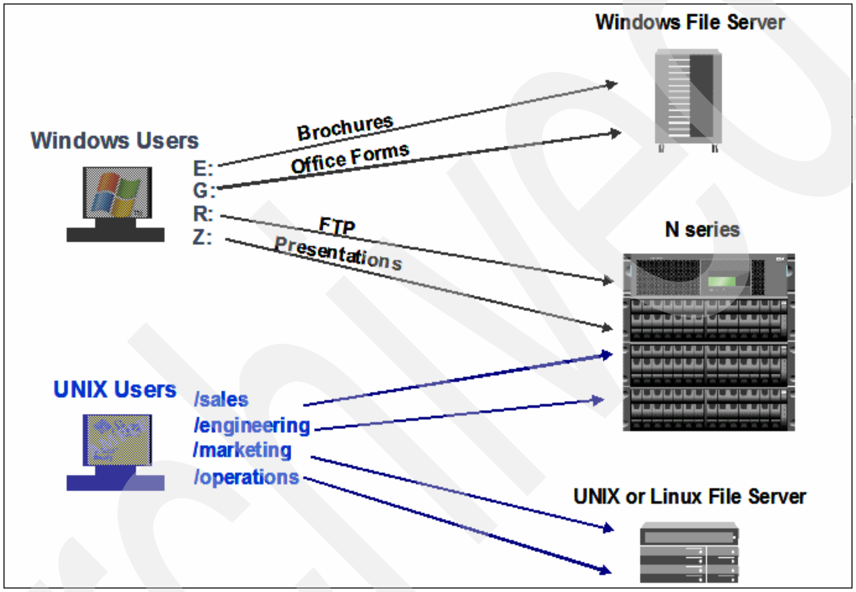

Figure 5-2 File access and management complexity

The VFM global namespace does for files what DNS does for networking: It provides a directory service. This directory service can deliver location-independent services to users and applications across multiple, heterogeneous, distributed file systems.

Once implemented, a global namespace enables users to access files in a logical, location-independent way, much like users access Web pages on the Internet. For example, when a user enters `http://www.IBM.com`, the user does not know the IP address for IBM, nor does the user care. Similarly, with a global namespace in place, the user accesses information using Internet Explorer and the location of files are transparent (see Figure 5-3 on page 121).

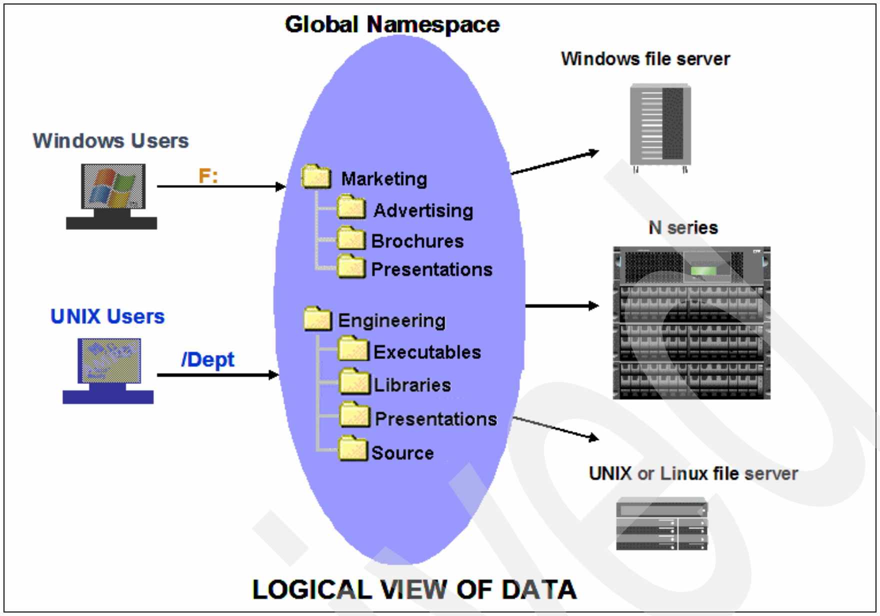

Figure 5-3 VFM global namespace as a solution

VFM utilizes the global namespace to solve a whole host of client issues. A global namespace is the foundation upon which the product and its applications are built. On top of the global namespace, a policy-based automation layer drives all of the functionality of the VFM solutions.

The global namespace offers the following main benefits:

- Aggregates all network file storage into a single, global file system
- Enables easy movement of data across storage architectures without affecting clients
- Makes it easy to add to and horizontally scale NAS
- Simplifies integration of N series in Microsoft Windows environments
- Provides centralized management and administration
- Enables creation of large distributed file systems
- Increases performance
- Shortens duration of backups
- Provides a common mount repository for data center and remote sites
- Enables nondisruptive migration and consolidation

- Facilitates expansion
- Enables tiers of storage for ILM
- Provides cost-effective business continuance
- Offers performance load balancing of N series and file servers

5.2 Creating DFS namespace with VFM

The first step in creating a namespace with VFM is to create a DFS root, the three types of which are: domain-based DFS root, standalone DFS root, and consolidation DFS root. This section describes each of the DFS root types as well as how to create a DFS namespace with VFM.

5.2.1 Domain-based DFS root

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

The prerequisites for domain-based DFS root configuration are as follows:

- 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 be operational.
- The account used for VFM must have administrator permissions for both the domain and the server where DFS root is created.
- We strongly recommend 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 an NT4 domain and is attempting to create a domain-based root in a separate Active Directory domain.
- Microsoft 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.

5.2.2 Standalone DFS root

A standalone DFS root stores its configuration information 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.

The configuration of standalone DFS root has the following prerequisites:

- Microsoft DFS server software (Distributed File System Service dfssvc.exe) must be running on the server that hosts the DFS root.
- Administrator permissions are required for the server where DFS root is created.
- Microsoft 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.

5.2.3 Consolidation DFS root

Consolidation DFS roots are special standalone roots that enable moving files referenced in UNC paths embedded in links, in line-of-business applications, and other places without affecting users. The consolidation root redirects the user from the old location referenced by the UNC path to the new location.

The configuration of consolidation DFS root has the following prerequisites:

- Microsoft DFS server software (Distributed File System Service dfssvc.exe) must be running on the server that hosts the DFS root.
- A software update must be installed to host a consolidation root on the server selected. After applying the software update, the Distributed File System Service must be restarted. For more information about consolidation roots, see the Microsoft Knowledge Base Article 829885. You can access the article here:

  `http://support.microsoft.com/kb/829885/en-us`
- The server to be consolidated must be renamed before configuring the consolidation root.
- The consolidation root must be installed on Microsoft Windows Server 2003. We recommend the root be hosted 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, two servers named server1.npd.com and server1.`<subdomain>`.npd.com cannot be consolidated.
- 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.

5.2.4 Creating a DFS root

DFS root consists in a share at the top of the namespace topology or the starting point for the links and shared folders that make up the namespace. Perform the following steps to install a namespace with VFM:

1. After launching VFM, right-click **Logical View** and choose **Add a new DFS root**, as shown in Figure 5-4.

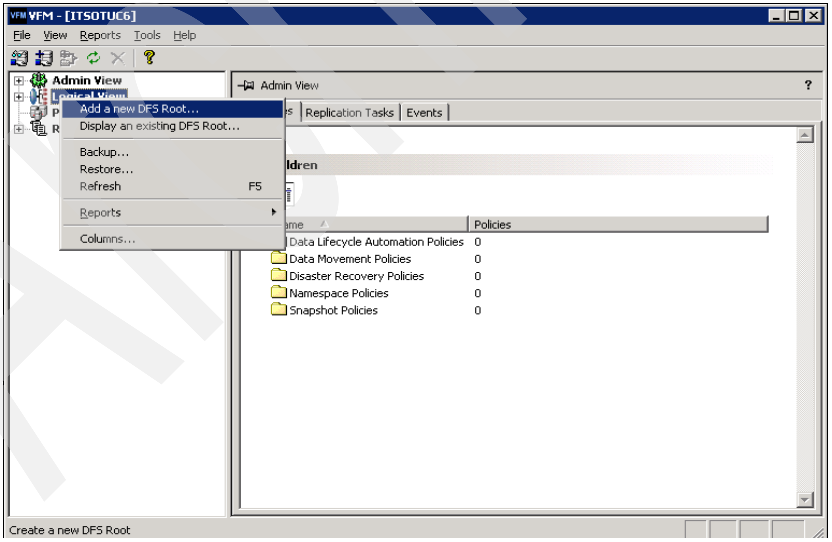

ame	Policies
Data Lifecycle Automation Policies	0
Data Movement Policies	0
Disaster Recovery Policies	0
Namespace Policies	0
Snapshot Policies	0

Figure 5-4 Add a new VFM root window

The New DFS Root Wizard window is displayed (see Figure 5-5).

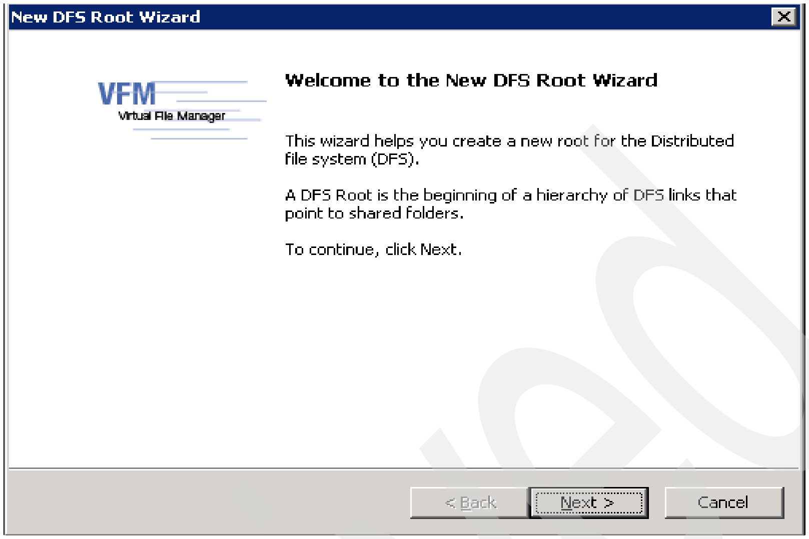

Welcome to the New DFS Root Wizard

This wizard helps you create a new root for the Distributed file system (DFS).

A DFS Root is the beginning of a hierarchy of DFS links that point to shared folders.

To continue, click Next.

Figure 5-5 New DFS Root Wizard

2. Click **Next**. A window requesting a name for the host server for the DFS root is displayed (see Figure 5-6).

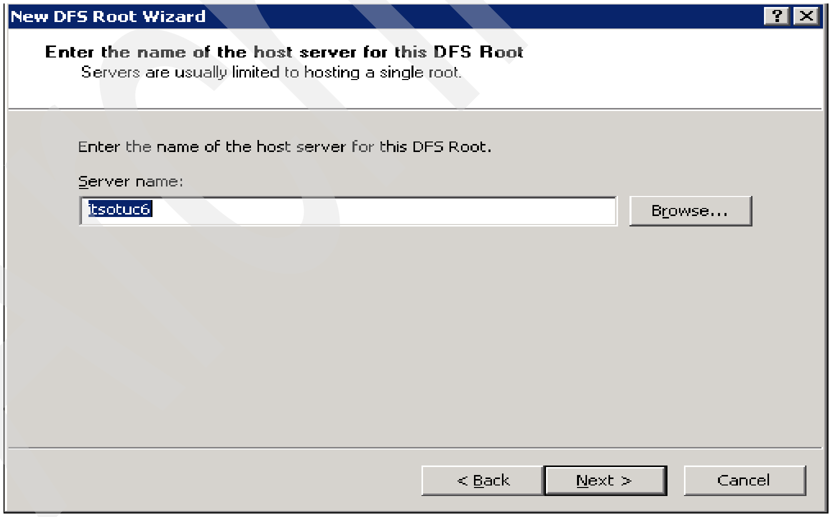

Enter the name of the host server for this DFS Root

Servers are usually limited to hosting a single root.

Enter the name of the host server for this DFS Root.

Server name: itsotuc6

Figure 5-6 Host server name window

3. After entering the name of the server or reaching it through the network using the Browse button, click **Next**. A DFS Root Type Wizard window is shown (see Figure 5-7).

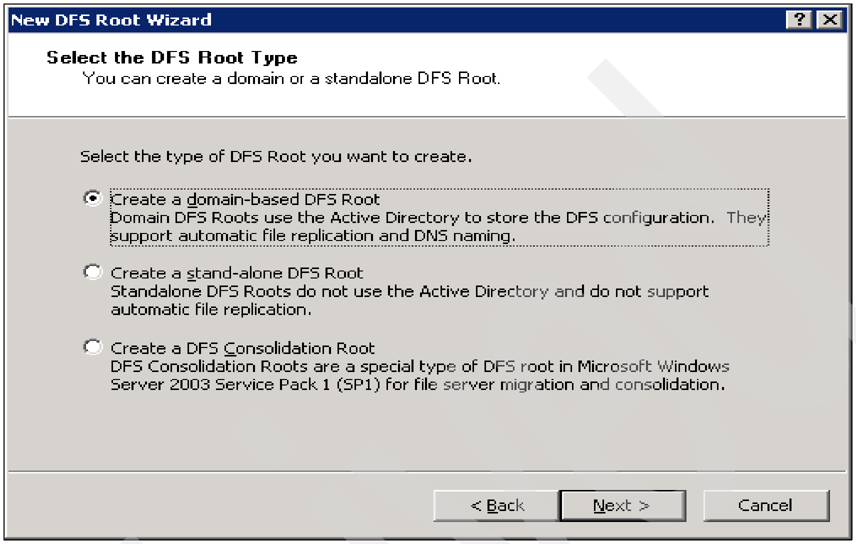

New DFS Root Wizard

Select the DFS Root Type

You can create a domain or a standalone DFS Root.

Select the type of DFS Root you want to create.

● Create a domain-based DFS Root
Domain DFS Roots use the Active Directory to store the DFS configuration. They support automatic file replication and DNS naming.

○ Create a stand-alone DFS Root
Standalone DFS Roots do not use the Active Directory and do not support automatic file replication.

○ Create a DFS Consolidation Root
DFS Consolidation Roots are a special type of DFS root in Microsoft Windows Server 2003 Service Pack 1 (SP1) for file server migration and consolidation.

Figure 5-7 DFS Root Type window

4. Select the type of DFS root that best applies (for this example, we use the first type, Create a default-based DFS root), then click **Next**. A window prompting you to select an existing share or create a new one is displayed (see Figure 5-8).

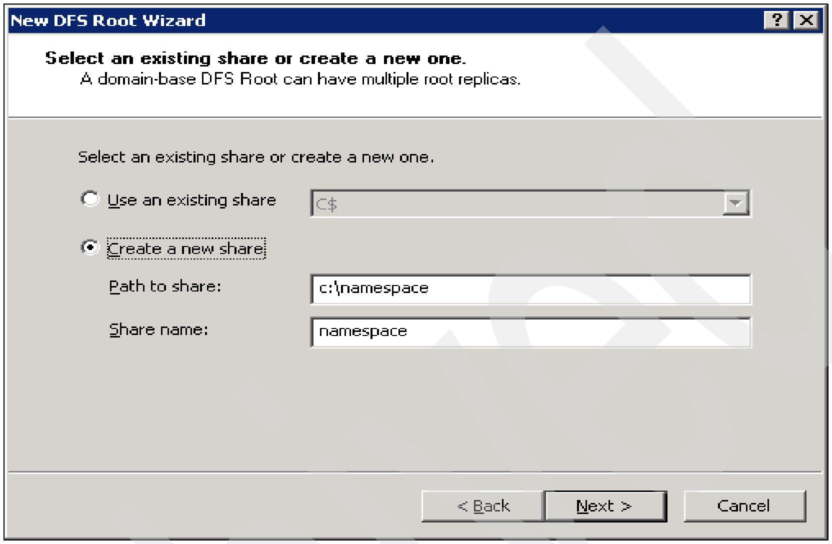

New DFS Root Wizard

Select an existing share or create a new one.

A domain-base DFS Root can have multiple root replicas.

Select an existing share or create a new one.

○ Use an existing share: C$

● Create a new share

Path to share: c:\namespace

Share name: namespace

Figure 5-8 Select or create a new share window

5. For a share already created, select the first option **Use an existing share** (in this example, we create another share named Namespace**).** After choosing the appropriate option, click **Next**. A window requesting a DFS root name is shown (see Figure 5-9).

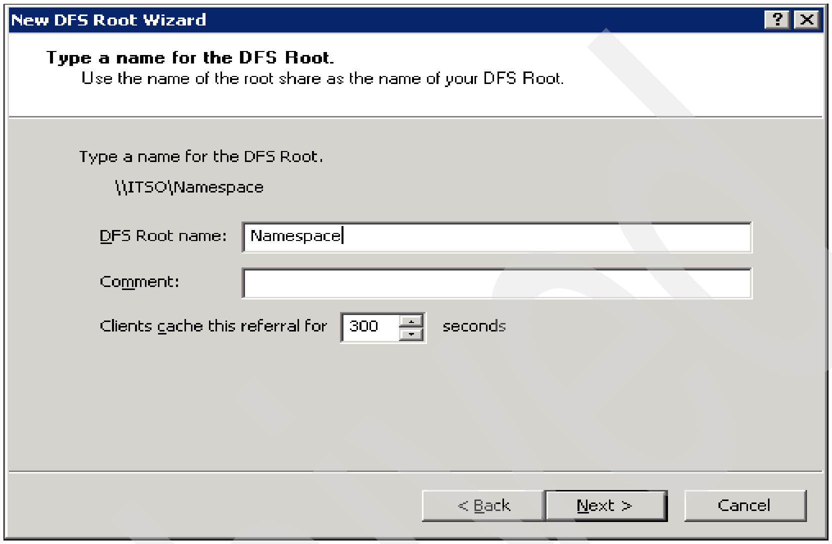

Type a name for the DFS Root.
Use the name of the root share as the name of your DFS Root.

Type a name for the DFS Root.
\\ITSO\Namespace

DFS Root name: Namespace

Comment:

Clients cache this referral for 300 seconds

Figure 5-9 New DFS root name window

6. Enter a DFS root name and comment that best apply (for our case, the DFS root name is Namespace), then click **Next**. The Completing the New DFS Root Wizard is shown (see Figure 5-10).

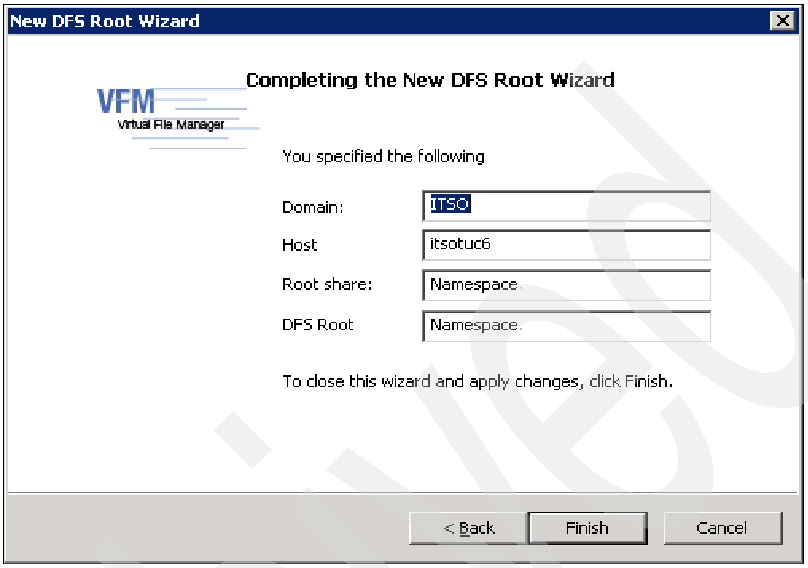

New DFS Root Wizard

Completing the New DFS Root Wizard

You specified the following

Domain: ITSO

Host itsotuc6

Root share: Namespace

DFS Root Namespace

To close this wizard and apply changes, click Finish.

Figure 5-10 Completing the New DFS Root Wizard

7. Click **Finish** to exit the wizard. On the VFM main console, expand the **Logical View** tree to check the DFS root namespace just created (see Figure 5-11).

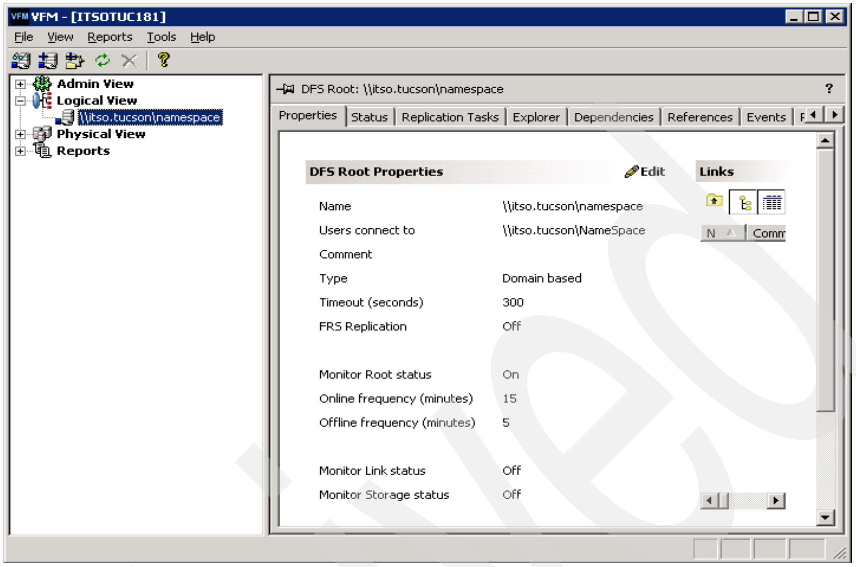

DFS Root Properties

Name	\\itso.tucson\namespace
Users connect to	\\itso.tucson\NameSpace
Comment	
Type	Domain based
Timeout (seconds)	300
FRS Replication	Off
Monitor Root status	On
Online frequency (minutes)	15
Offline frequency (minutes)	5
Monitor Link status	Off
Monitor Storage status	Off

Figure 5-11 View of DFS root namespace

5.2.5 Adding folders to create a hierarchy

You add folders to create a logical hierarchy of a namespace. To add folders on the DFS root just created, follow these steps:

1. On the VFM server window, expand **Logical View** and right-click the DFS root name that best applies (for this example, we use \\itso.tucson\namespace), then click **Add folder** as shown in Figure 5-12.

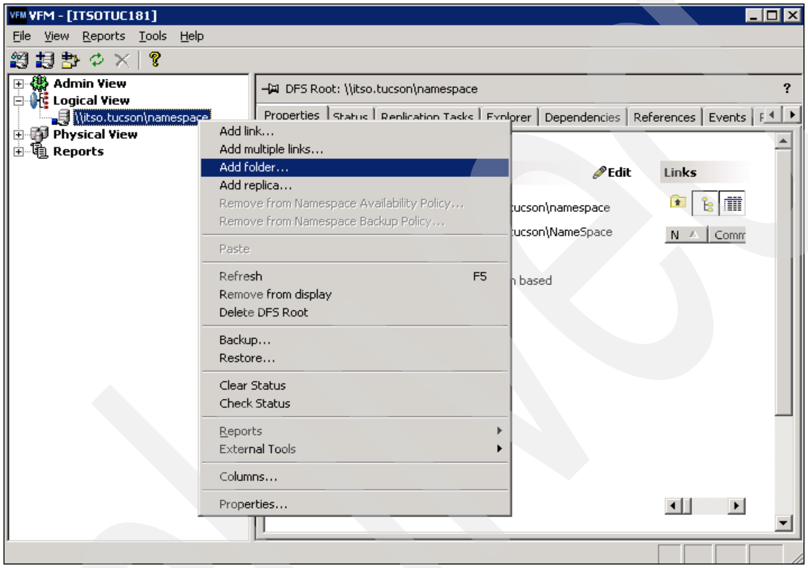

Figure 5-12 Adding a folder on the DFS root

A window prompting you for the name of the folder is shown (see Figure 5-13).

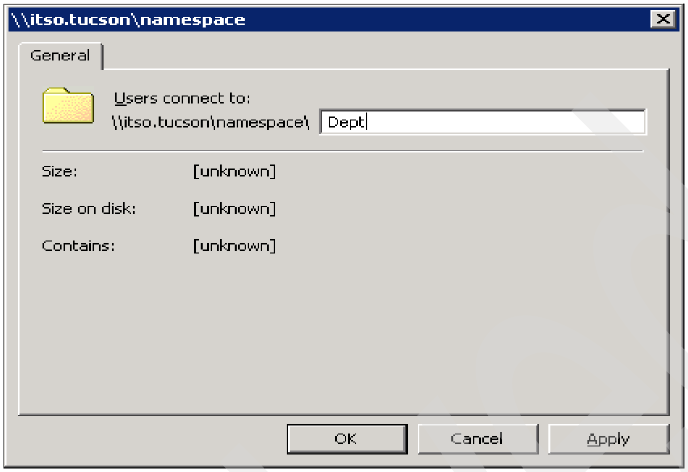

Users connect to:
\\itso.tucson\namespace\ Dept

Size: [unknown]

Size on disk: [unknown]

Contains: [unknown]

Figure 5-13 Folder name window

2. Type the name that best applies (for this example, the folder name is Dept) and then click **Apply**. On the VFM server window, the folder Dept is shown under the DFS root (\\itso.tucson\namespace) as shown in Figure 5-14.

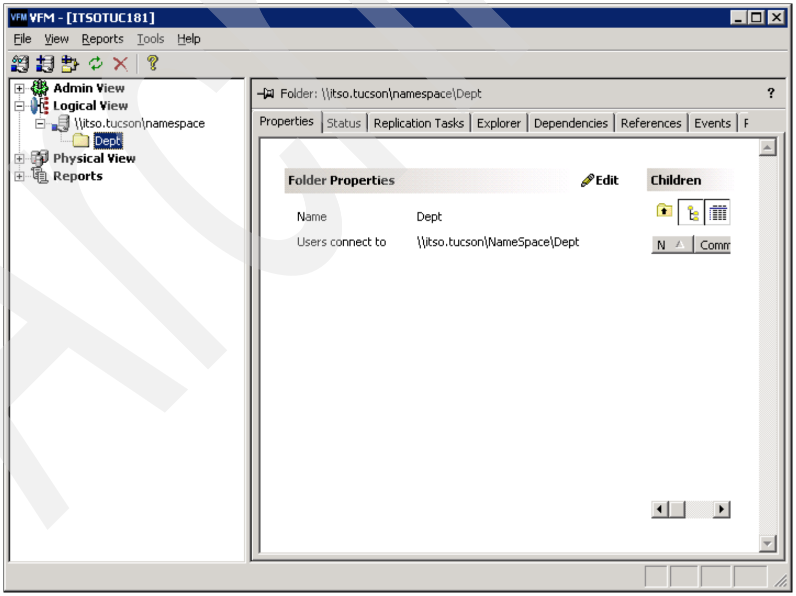

Name: Dept

Users connect to: \\itso.tucson\NameSpace\Dept

Figure 5-14 Dept folder view

5.2.6 Link overview

A DFS link is an element of the DFS namespace that can point to a share (also called the *link target*). This is called a *single-targeted DFS link* if the DFS link points to only one target share. A multitargeted DFS link is a DFS link that points to more than one target. We recommend a maximum of 256 DFS link targets.

Consider this scenario as an example of a multitargeted DFS link use. Suppose a multitargeted DFS link is in the DFS namespace called Finance with four targets:

- \\HoustonServer\finance
- \\LosAngeles\finance
- \\NewYork\finance
- \\London\finance

We expect the Houston users to be linked to the closest server, in this case \\houstonserver. We also expect the users in New York to be linked to \\newyork, and so on. This arrangement allows load balancing for all data across the set of servers. If all the data is identical across all servers and users are only reading data from the Finance share, this arrangement is beneficial.

DFS also allows a link target to be taken offline. If a link target is offline, it is not part of the referral given to the client. Clients are not directed to that link target because they do not know about it.

Multitargeted links used in this way are useful for read-only data distributed throughout an enterprise (for example, HR policies and software distribution shares). VFM uses the offline link target concept to enable multitargeted DFS link disaster recovery (DR). Instead of leaving all link targets online, VFM keeps one link target online and takes all other link targets offline. All users are directed to the online link target. If the online link target is unavailable, VFM forces one of the other link targets to become online and initiate the failover (automatically or manually, based on the policy). This way, you are unlikely to have a split-brain scenario, ensuring that all users update one copy of the data while the alternate copies of the data are available for failover.

5.2.7 Adding a link

A DFS link is located under the namespace root. It forms a connection to one or more shared folders, links, or to another root.

In the example that follows, perform these steps to link a DFS namespace to a share located on a different server:

1. On the VFM server Local view tree window, expand the DFS root that best applies (for our example, we use \\itso.tucson\namespace\Dept) and right-click **Add link** as shown in Figure 5-15.

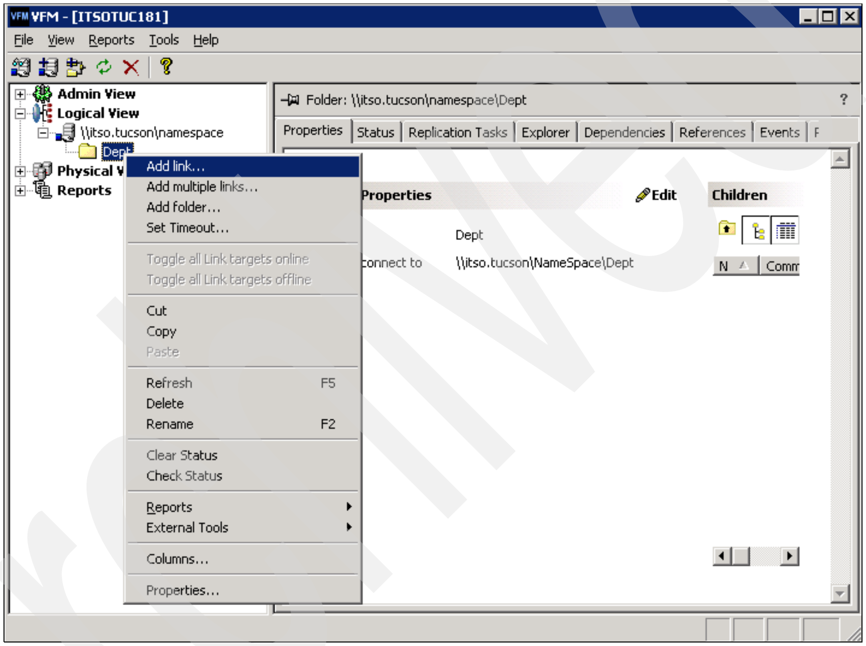

Dept

Connect to \\itso.tucson\NameSpace\Dept

Figure 5-15 Add link menu

2. A window prompting you for the name and the path to which the Dept folder must be linked is shown. In this example, the Dept folder is linked to a Marketing folder located in a N series machine (itsotuc4). Click **Add** and enter the folder name or browse the network path to reach the folder that best applies (in this example, \\itsotuc4.itso.tucson\Marketing) as shown in Figure 5-16.

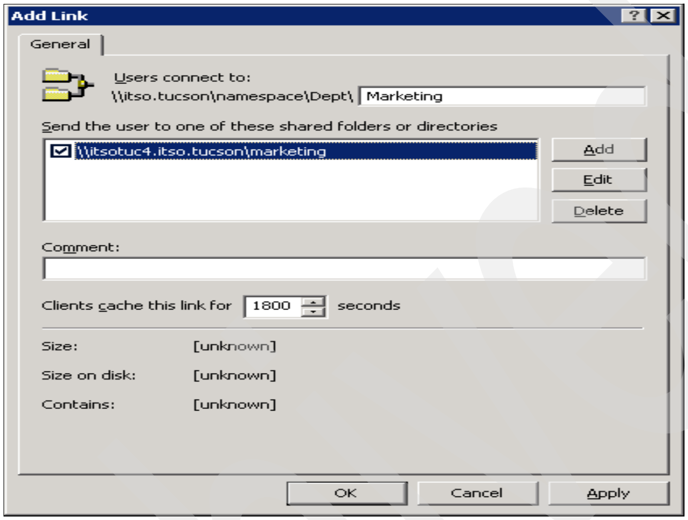

Users connect to:
\\itso.tucson\namespace\Dept\ Marketing

Send the user to one of these shared folders or directories
☑ \\itsotuc4.itso.tucson\marketing

Comment:

Clients cache this link for 1800 seconds

Size: [unknown]
Size on disk: [unknown]
Contains: [unknown]

Figure 5-16 Adding a link window

3. Click **OK** to finish. A Marketing link is shown under the Dept folder (see Figure 5-17).

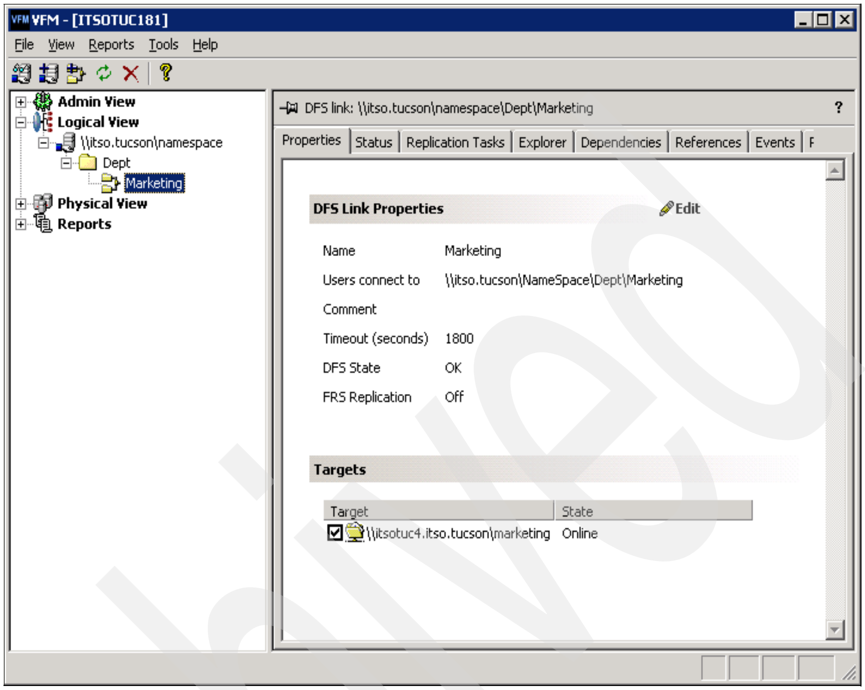

DFS Link Properties

Name	Marketing
Users connect to	\\itso.tucson\NameSpace\Dept\Marketing
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
☑ \\itsotuc4.itso.tucson\marketing	Online

Figure 5-17 Link view under Dept folder

5.2.8 Adding a second link

To create a second link, follow the steps described in the previous section (5.2.7, “Adding a link” on page 134). For the example in this section, the Dept folder is linked to a Sales folder located in a N series machine (itsotuc3) as shown (Figure 5-18 on page 137).

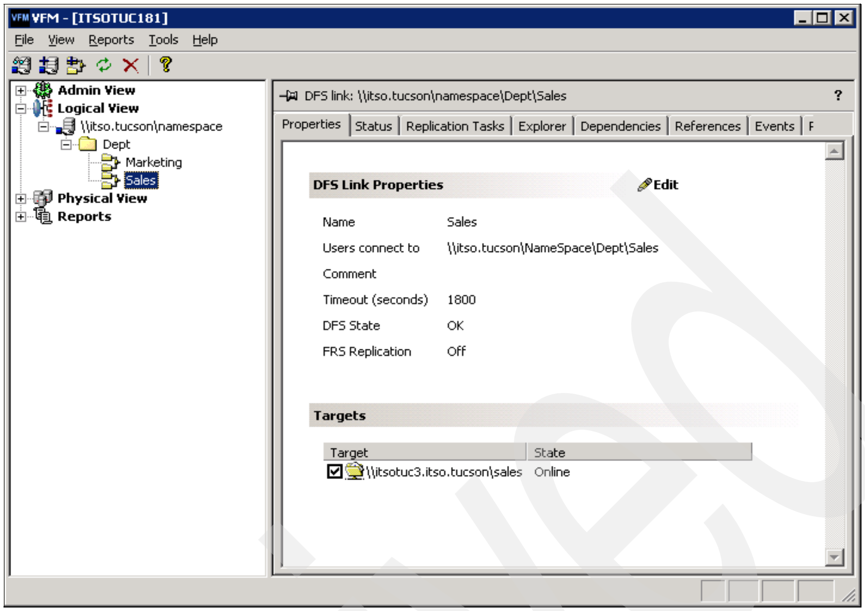

DFS link: \\itso.tucson\namespace\Dept\Sales

DFS Link Properties

Name	Sales
Users connect to	\\itso.tucson\NameSpace\Dept\Sales
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
☑ \\itsotuc3.itso.tucson\sales	Online

Figure 5-18 Second link view under Dept folder

5.2.9 Client view of the namespace

Once the links are created, users are able to browse them by mapping the DFS namespace path (in this example, \\itso.tucson\namespace\dept) as shown in Figure 5-19 on page 138.

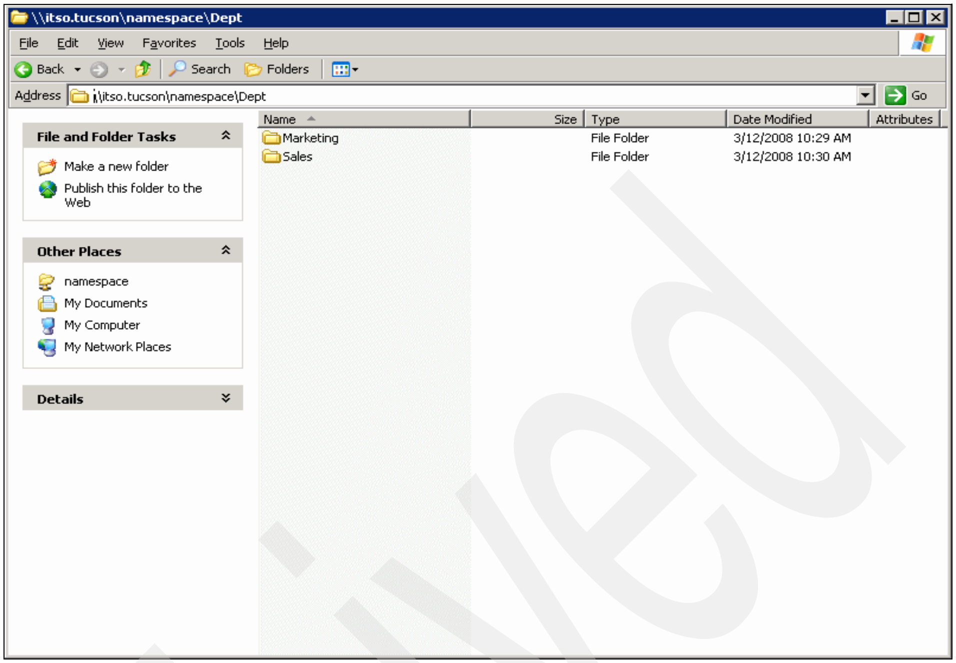

Address: \\itso.tucson\namespace\Dept

Name	Size	Type	Date Modified	Attributes
Marketing		File Folder	3/12/2008 10:29 AM	
Sales		File Folder	3/12/2008 10:30 AM	

Figure 5-19 Users view of the namespace

Note that the Marketing and Sales folders are located on different machines. In spite of that, users are able to map both folders within a single network path.

6

Installation

This chapter describes the installation of VFM components. It is important to understand each component's responsibilities as well as to follow all preparation steps described in Chapter 3, "Preparation" on page 83. Each section in this chapter represents a "setup type" on the installer and guides you through each one of the options.

You cannot install two VFM components by running the installer twice. If you run the installer in a system were VFM is already installed, the uninstaller is initiated instead. If you want to install more than one component in the same system, use the setup types containing multiple components.

6.1 Implementing a typical installation

The typical installation installs both the VFM server and the client.

> **Note:** This section describes the recommended way to install the server.

Perform the following steps to install the VFM server and client:

1. On the VFM Installation CD, look for the Setup.exe file and double-click it.
2. Click **Next** (see Figure 6-1).

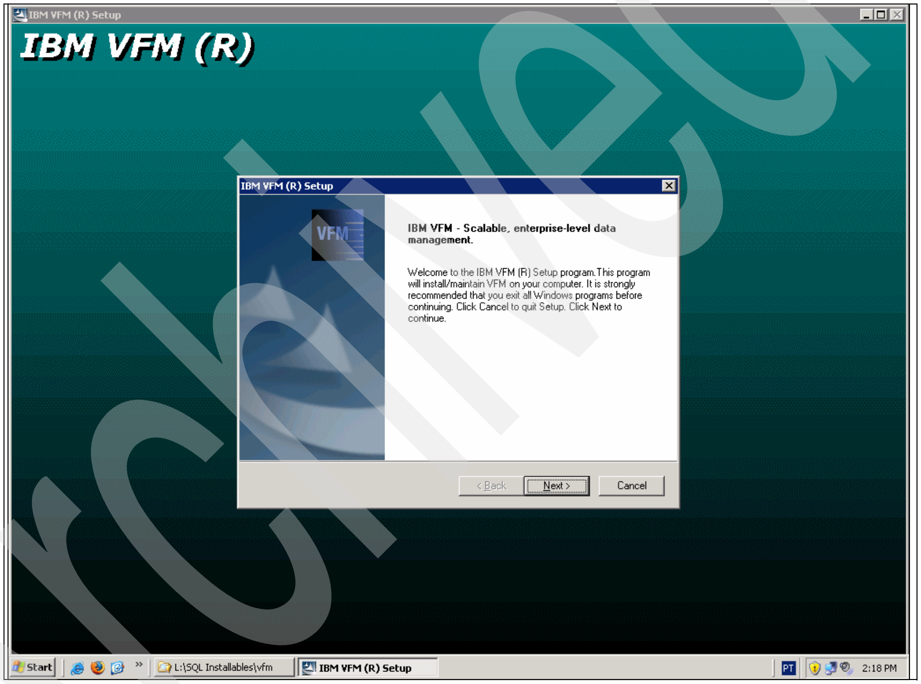

IBM VFM (R) Setup

IBM VFM - Scalable, enterprise-level data management.

Welcome to the IBM VFM (R) Setup program.This program will install/maintain VFM on your computer. It is strongly recommended that you exit all Windows programs before continuing. Click Cancel to quit Setup. Click Next to continue.

Figure 6-1 Typical installation: first window

3. Read and accept the license to proceed (see Figure 6-2).

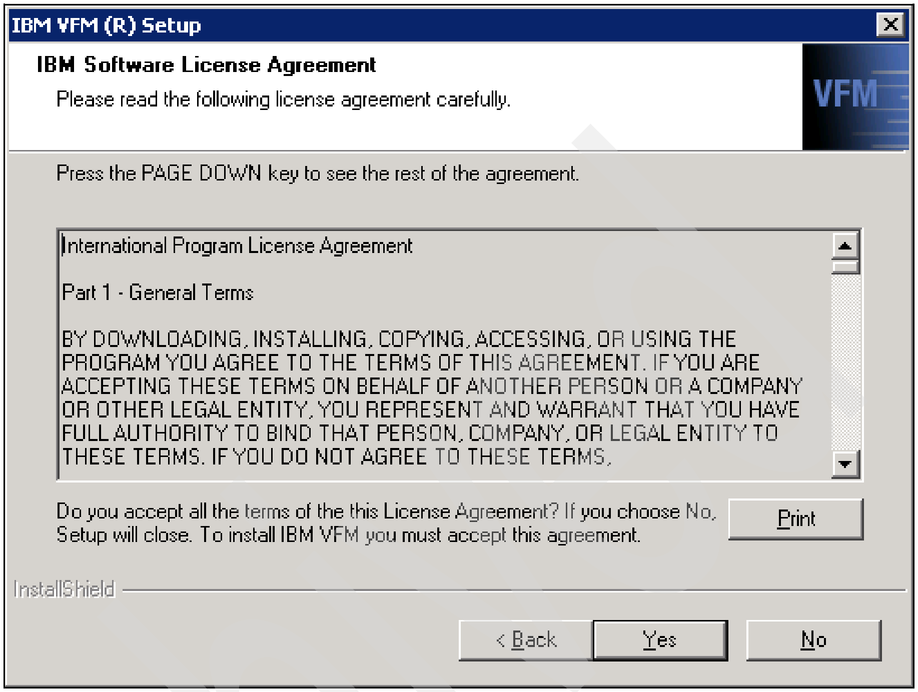

IBM VFM (R) Setup

IBM Software License Agreement

Please read the following license agreement carefully.

Press the PAGE DOWN key to see the rest of the agreement.

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, OR USING THE PROGRAM YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF ANOTHER PERSON OR A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON, COMPANY, OR LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

Do you accept all the terms of the this License Agreement? If you choose No, Setup will close. To install IBM VFM you must accept this agreement.

Figure 6-2 Typical installation: license agreement

4. Check the README notes and click **Yes** to proceed (see Figure 6-3).

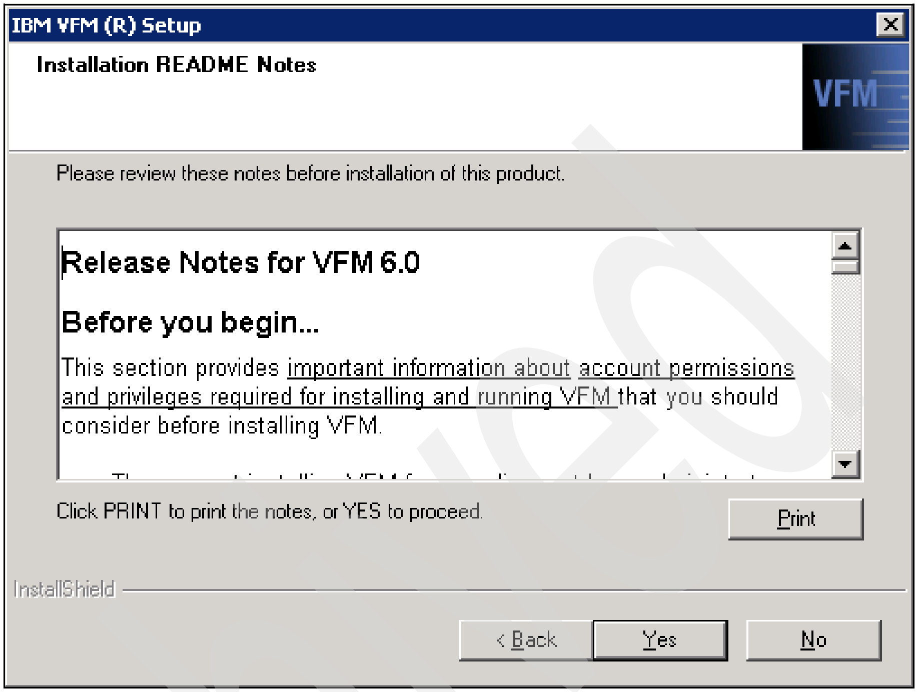

Please review these notes before installation of this product.

Release Notes for VFM 6.0

Before you begin...

This section provides important information about account permissions and privileges required for installing and running VFM that you should consider before installing VFM.

Click PRINT to print the notes, or YES to proceed.

Figure 6-3 Typical installation: README notes

5. Choose **Typical** and click **Next** (see Figure 6-4). The other setup types listed in the Choose a setup type window are discussed in the following sections.

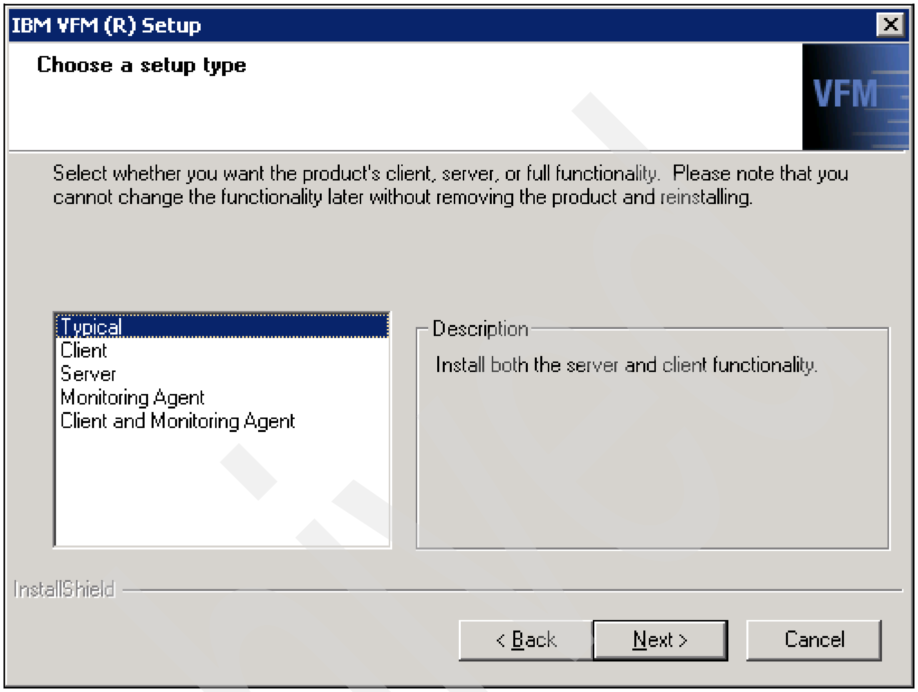

Select whether you want the product's client, server, or full functionality. Please note that you cannot change the functionality later without removing the product and reinstalling.

- Typical
- Client
- Server
- Monitoring Agent
- Client and Monitoring Agent

Description

Install both the server and client functionality.

Figure 6-4 Typical installation: choosing setup type

6. If you have an Evaluation license, fill in the form with the Serial Number and the Activation Key, then click **Next** (see Figure 6-5). Skip step 7 on page 145.

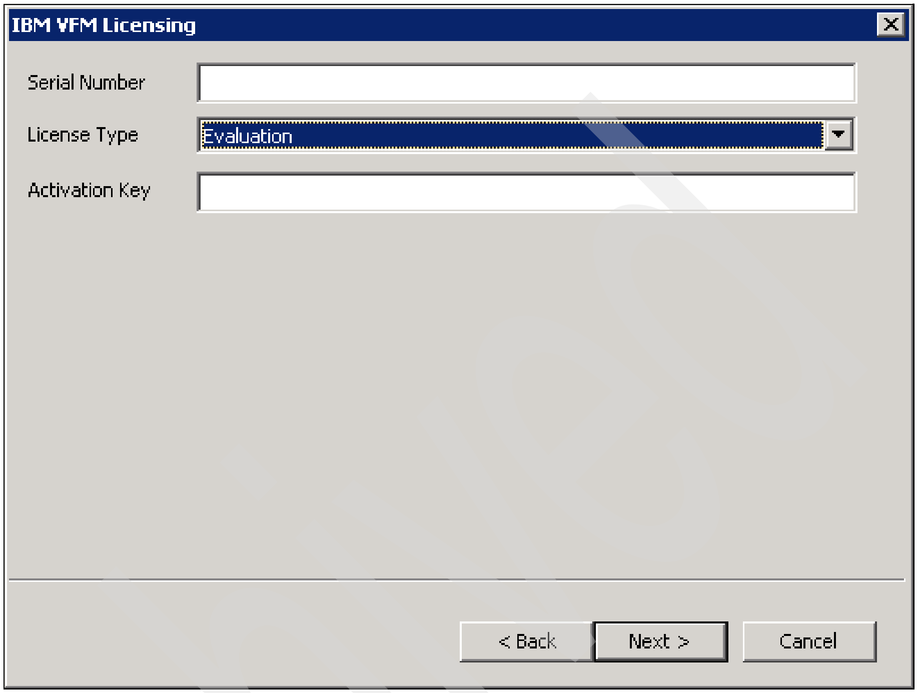

IBM VFM Licensing

Serial Number

License Type: Evaluation

Activation Key

Figure 6-5 Typical installation: licensing

7. If you have a Full license, choose **Full** in the License Type selection box, and the window shown in Figure 6-6 is displayed. Fill in the form with the Serial Number and all the Activation Keys you have received, then click **Next**.

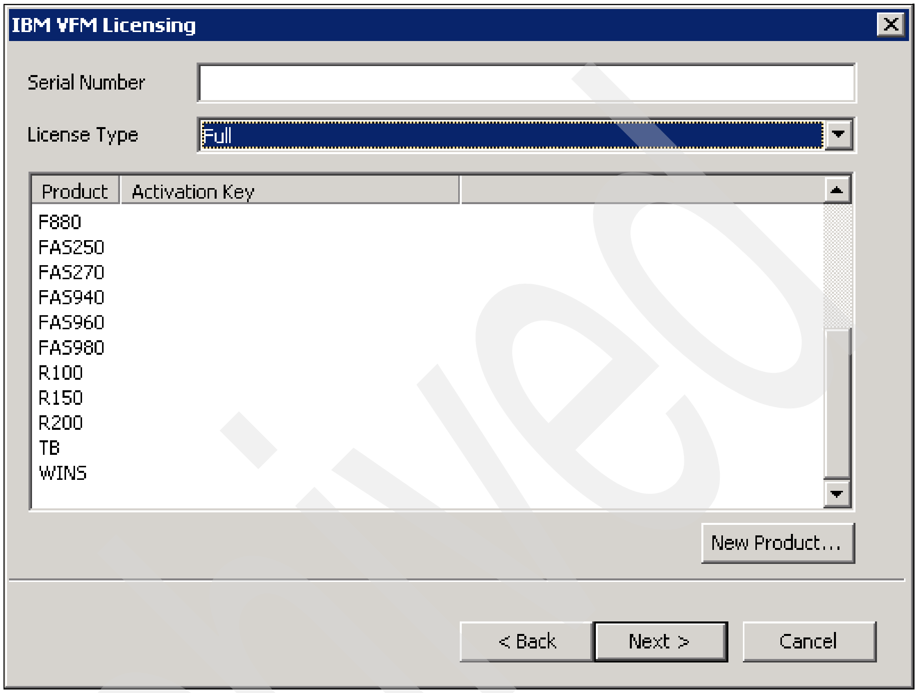

IBM VFM Licensing

Serial Number:

License Type: Full

Product	Activation Key
F880	
FAS250	
FAS270	
FAS940	
FAS960	
FAS980	
R100	
R150	
R200	
TB	
WINS	

Figure 6-6 Typical installation: full licensing

8. Now you can choose where VFM is installed (see Figure 6-7). It must be in a NTFS file system local to the host. The default value is acceptable in most cases.

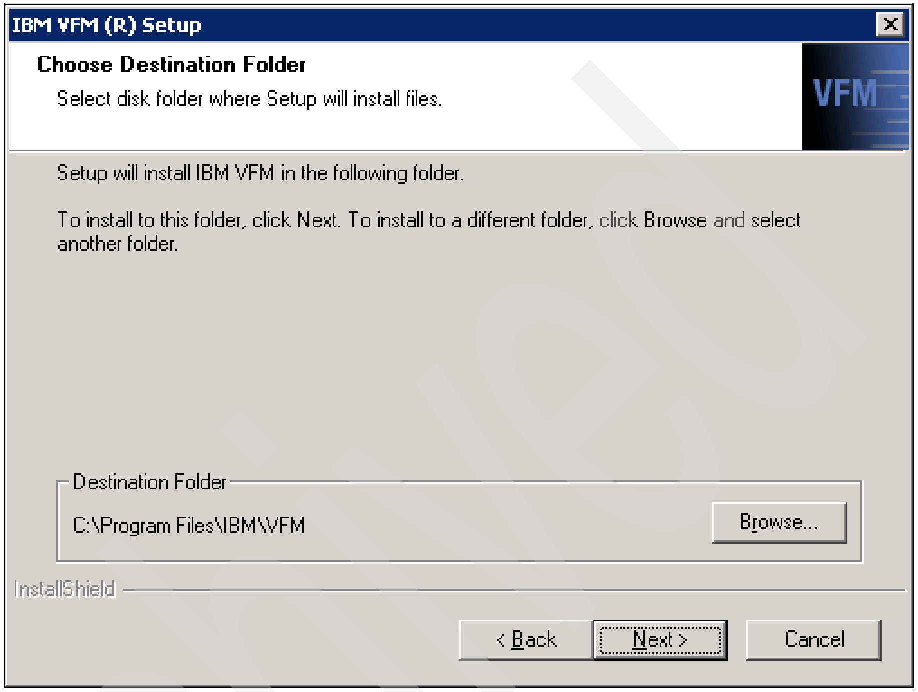

IBM VFM (R) Setup

Choose Destination Folder

Select disk folder where Setup will install files.

Setup will install IBM VFM in the following folder.

To install to this folder, click Next. To install to a different folder, click Browse and select another folder.

Destination Folder

C:\Program Files\IBM\VFM

Figure 6-7 Typical installation: choosing destination folder

9. VFM must store some files locally. At this point, you can modify where these files are stored (see Figure 6-8). The default value is acceptable in most cases.

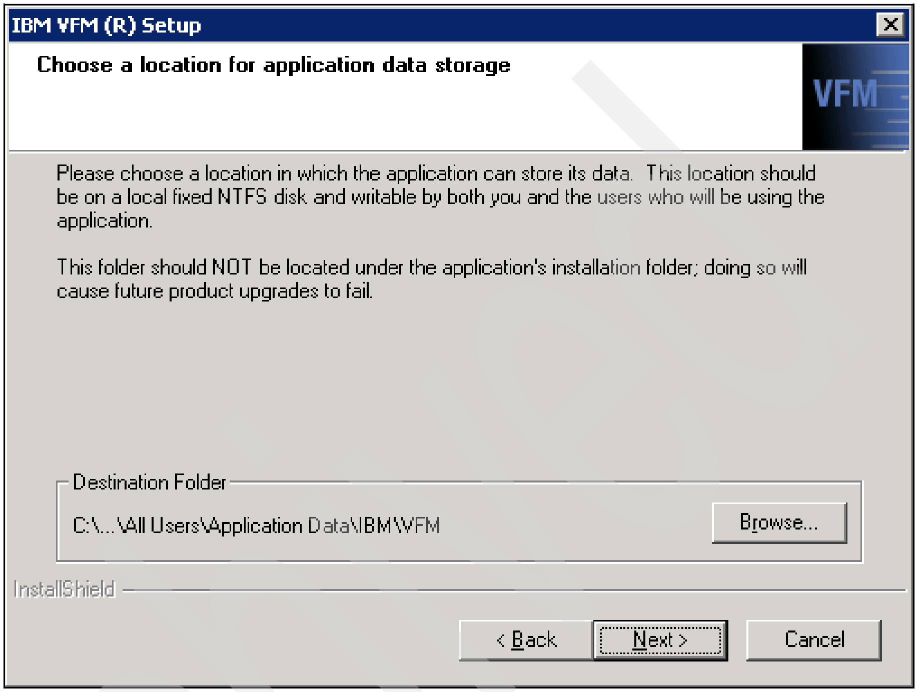

IBM VFM (R) Setup

Choose a location for application data storage

Please choose a location in which the application can store its data. This location should be on a local fixed NTFS disk and writable by both you and the users who will be using the application.

This folder should NOT be located under the application's installation folder; doing so will cause future product upgrades to fail.

Destination Folder

C:\...\All Users\Application Data\IBM\VFM

Figure 6-8 Typical installation: local data

10. Now choose whether or not you want to create a desktop shortcut for the VFM client (see Figure 6-9). Either way, the VFM client is available by choosing **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.

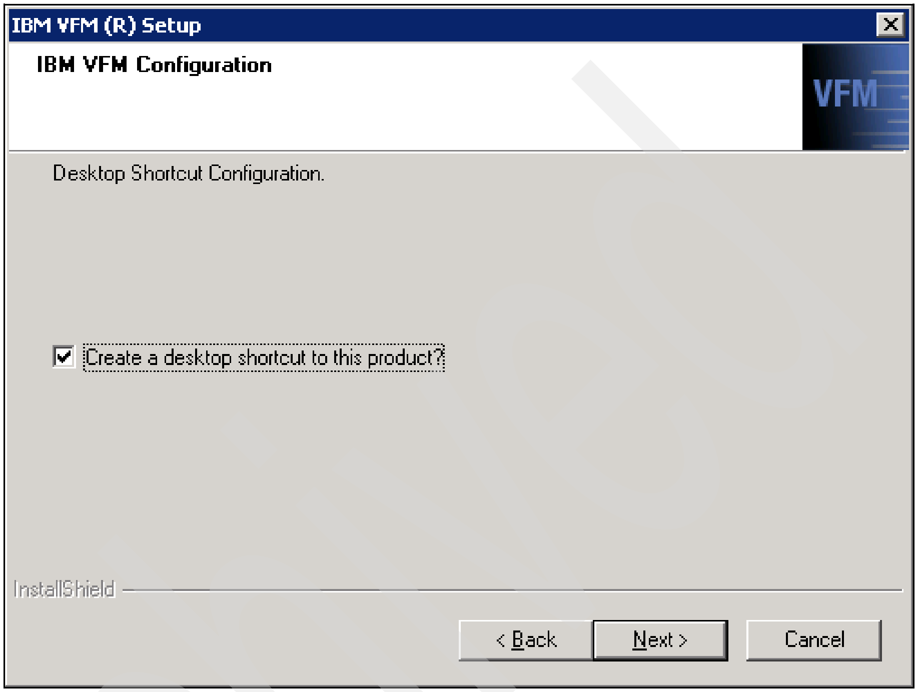

Desktop Shortcut Configuration.

☑ Create a desktop shortcut to this product?

Figure 6-9 Typical installation: desktop shortcut for VFM client

11. As stated in Chapter 3, “Preparation” on page 83, you must have a service account to run VFM. Enter the information about the account (see Figure 6-10). Remember to use the form DOMAIN\Username.

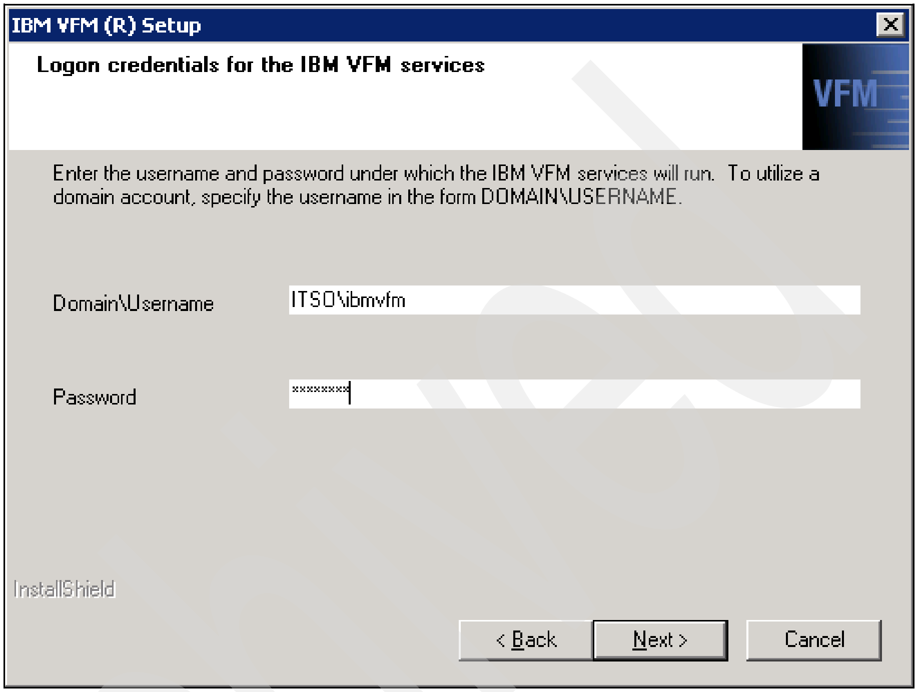

IBM VFM (R) Setup

Logon credentials for the IBM VFM services

Enter the username and password under which the IBM VFM services will run. To utilize a domain account, specify the username in the form DOMAIN\USERNAME.

Domain\Username: ITSO\ibmvfm

Password: ********

Figure 6-10 Typical installation: service account

VFM can automatically install the required database server on your machine. Refer to 3.5, “MSSQL” on page 107 for more information about the database server for VFM. If you are going to use a separate database server, refer to the steps in section 6.5, “Differences when using a separate MSSQL server” on page 188. In this section, we describe how to use the automatically installed server, MSDE.

To automatically install the database server, follow these steps:

1. Choose the option **Please install MSDE for me** and click **Next** (see Figure 6-11).

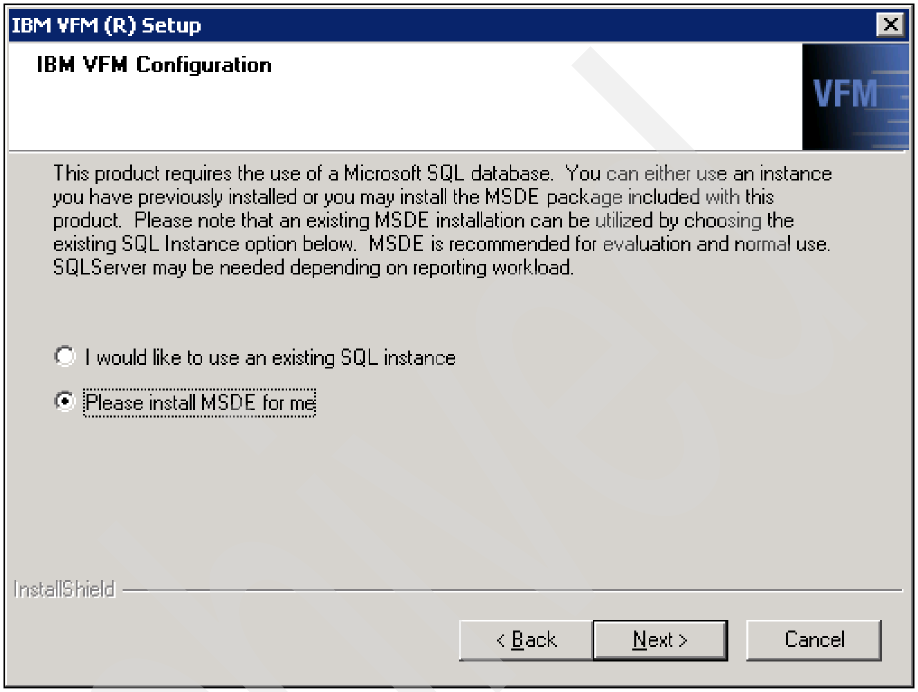

IBM VFM (R) Setup

IBM VFM Configuration

This product requires the use of a Microsoft SQL database. You can either use an instance you have previously installed or you may install the MSDE package included with this product. Please note that an existing MSDE installation can be utilized by choosing the existing SQL Instance option below. MSDE is recommended for evaluation and normal use. SQLServer may be needed depending on reporting workload.

○ I would like to use an existing SQL instance

● Please install MSDE for me

Figure 6-11 Typical installation: choosing SQL server

2. A summary of the chosen options is displayed. Click **Next** to proceed with the installation (see Figure 6-12).

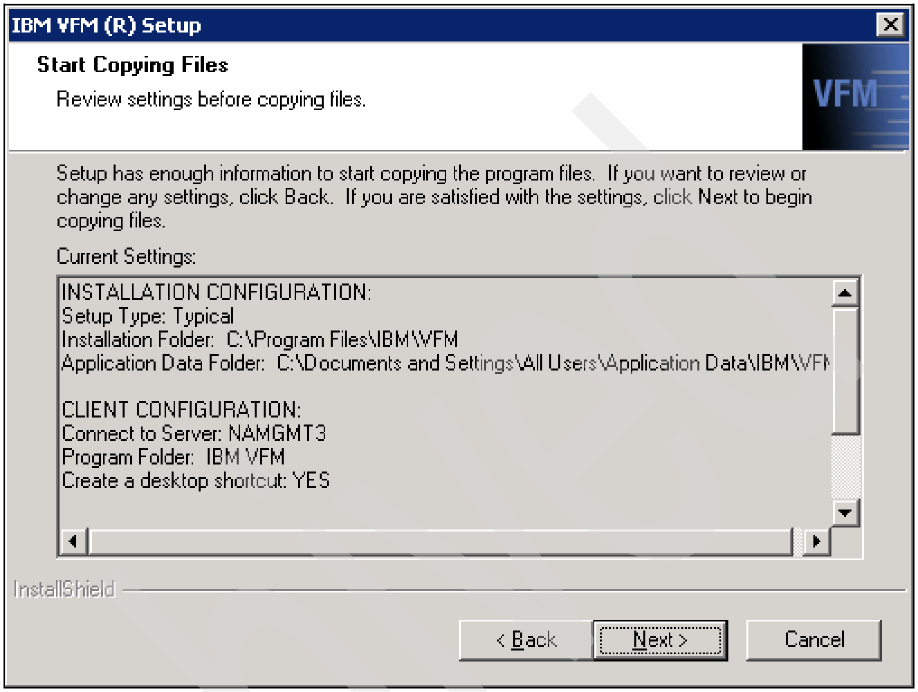

Setup has enough information to start copying the program files. If you want to review or change any settings, click Back. If you are satisfied with the settings, click Next to begin copying files.

Current Settings:

```
INSTALLATION CONFIGURATION:
Setup Type: Typical
Installation Folder:  C:\Program Files\IBM\VFM
Application Data Folder:  C:\Documents and Settings\All Users\Application Data\IBM\VFI

CLIENT CONFIGURATION:
Connect to Server: NAMGMT3
Program Folder:  IBM VFM
Create a desktop shortcut: YES
```

Figure 6-12 Typical installation: summary

3. At this point, the installation is complete; a window similar to that in Figure 6-13 is displayed. Click **Finish,** and VFM is ready to use.

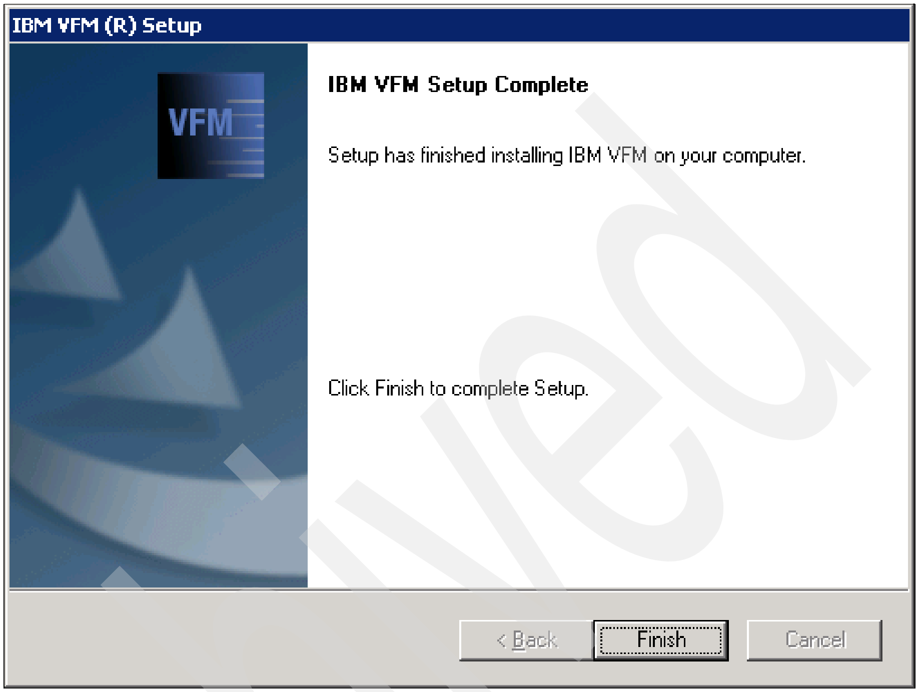
IBM VFM (R) Setup

IBM VFM Setup Complete

Setup has finished installing IBM VFM on your computer.

Click Finish to complete Setup.

Figure 6-13 Typical installation: finish

You can open the VFM client by going to **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.

6.2 Installing a client

With the Client option, it is possible to install the VFM client alone on other machines. This option can be used to install the VFM client on administrator machines.

> **Important:** The VFM client can be installed on several machines, but it must be always installed together with the VFM server using the Typical installation as described on section 6.1, “Implementing a typical installation” on page 140.

Follow these steps to install the client:

1. On the VFM Installation CD, find the Setup.exe file and double-click it.
2. Click **Next** (see Figure 6-14).

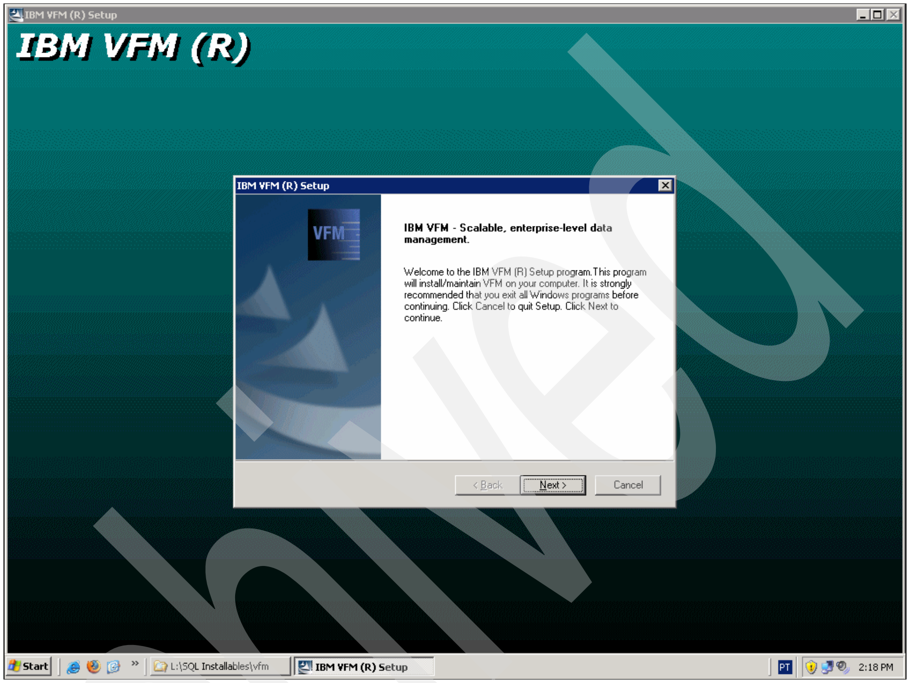

IBM VFM - Scalable, enterprise-level data management.

Welcome to the IBM VFM (R) Setup program.This program will install/maintain VFM on your computer. It is strongly recommended that you exit all Windows programs before continuing. Click Cancel to quit Setup. Click Next to continue.

Figure 6-14 Client installation: first window

3. Read and accept the license to proceed (see Figure 6-15).

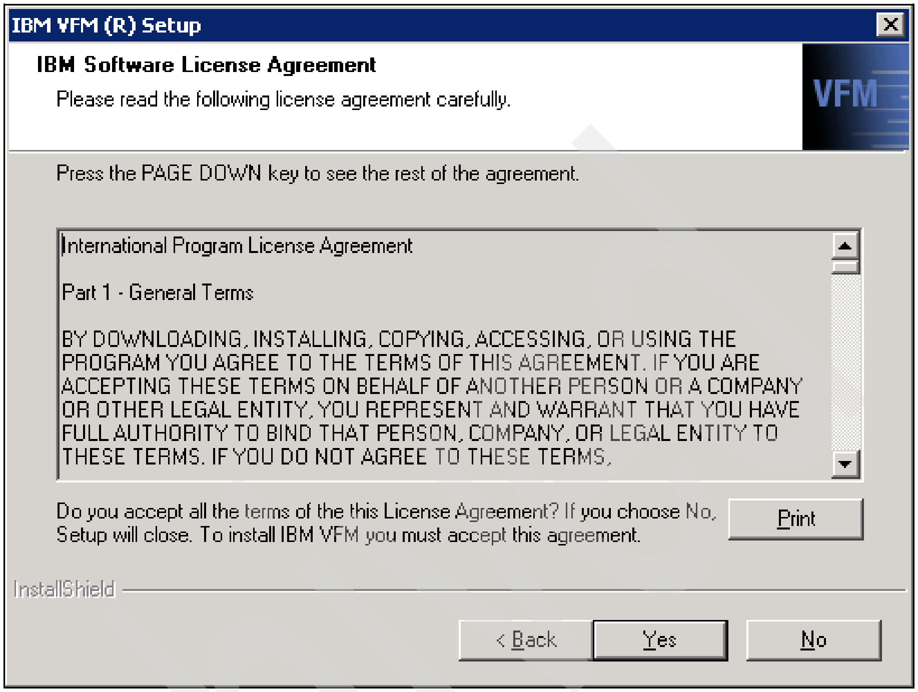

IBM Software License Agreement

Please read the following license agreement carefully.

Press the PAGE DOWN key to see the rest of the agreement.

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, OR USING THE PROGRAM YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF ANOTHER PERSON OR A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON, COMPANY, OR LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

Do you accept all the terms of the this License Agreement? If you choose No, Setup will close. To install IBM VFM you must accept this agreement.

Figure 6-15 Client installation: license agreement

4. Check the README notes and click **Yes** to proceed (see Figure 6-16).

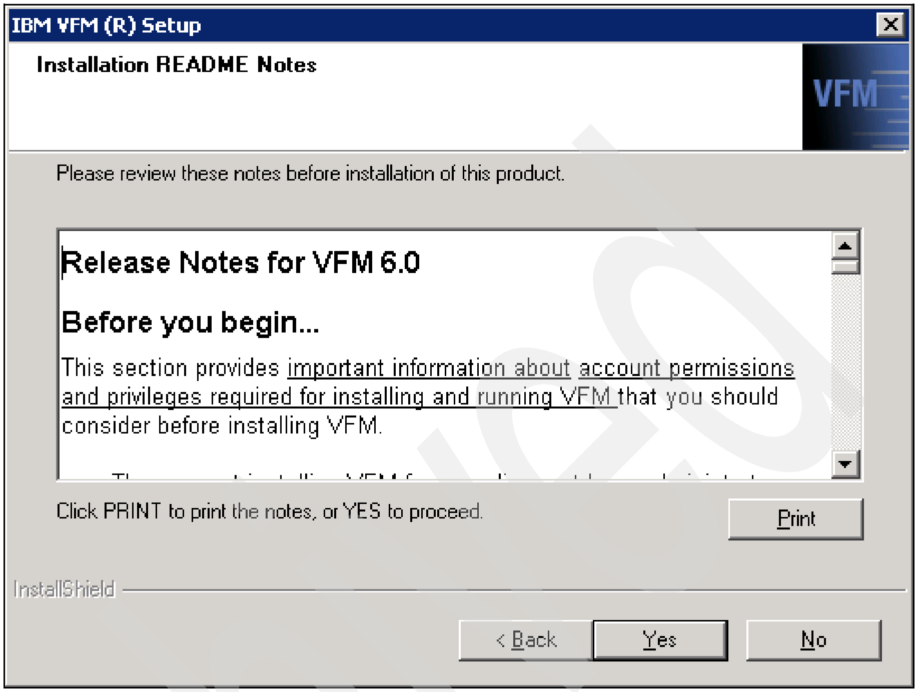

IBM VFM (R) Setup

Installation README Notes

Please review these notes before installation of this product.

Release Notes for VFM 6.0

Before you begin...

This section provides important information about account permissions and privileges required for installing and running VFM that you should consider before installing VFM.

Click PRINT to print the notes, or YES to proceed.

Figure 6-16 Client installation: README notes

5. Choose **Client** and click **Next** (see Figure 6-17). The other setup types are discussed in other sections in this chapter.

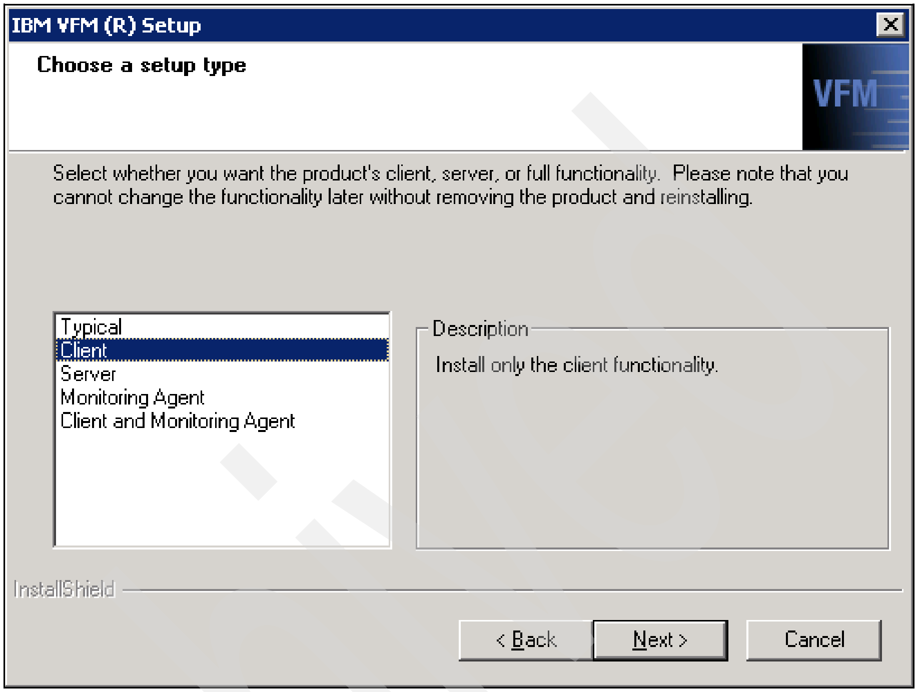

IBM VFM (R) Setup

Choose a setup type

Select whether you want the product's client, server, or full functionality. Please note that you cannot change the functionality later without removing the product and reinstalling.

- Typical
- Client
- Server
- Monitoring Agent
- Client and Monitoring Agent

Description

Install only the client functionality.

Figure 6-17 Client installation: choose setup type

6. Choose the VFM server to be used by default (see Figure 6-18).

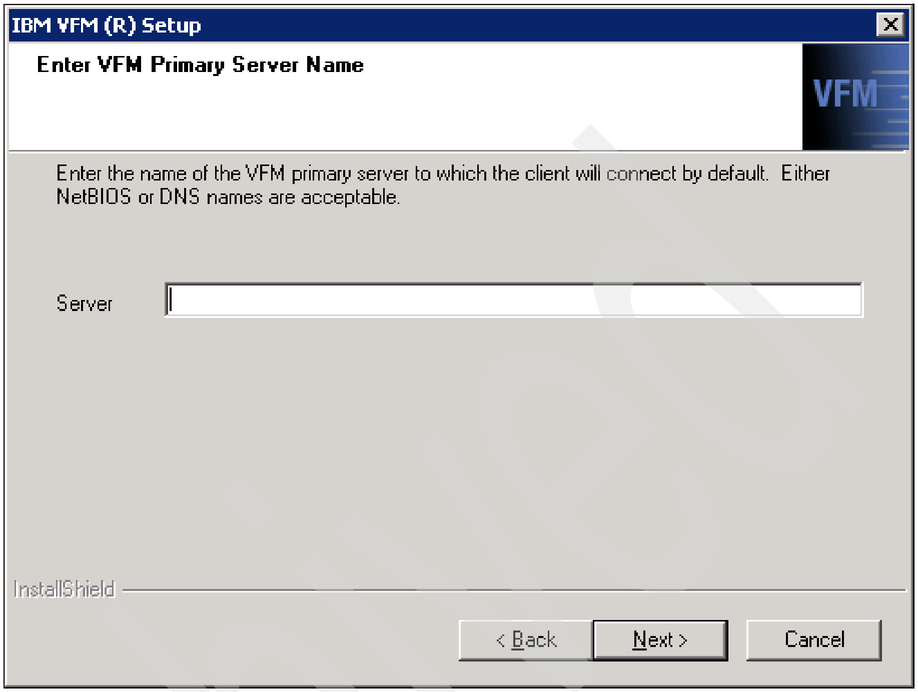

IBM VFM (R) Setup

Enter VFM Primary Server Name

Enter the name of the VFM primary server to which the client will connect by default. Either NetBIOS or DNS names are acceptable.

Server

Figure 6-18 Client installation: VFM server name

7. Choose whether or not you want to create a desktop shortcut for the VFM client (see Figure 6-19). Either way, the VFM client is available by choosing **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.

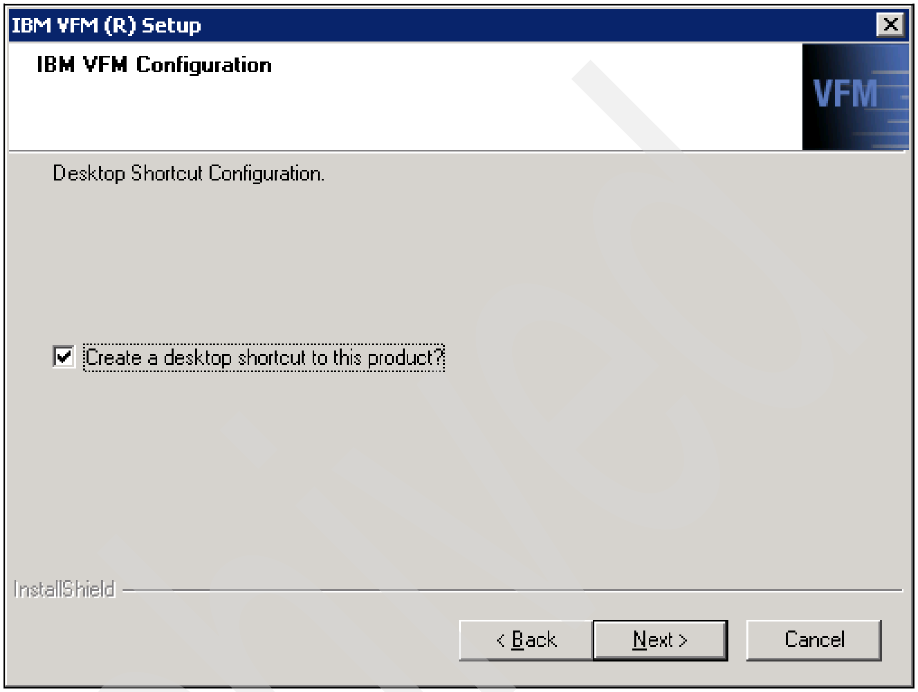

Desktop Shortcut Configuration.

☑ Create a desktop shortcut to this product?

Figure 6-19 Client installation: desktop shortcut

8. A summary of the chosen options is displayed. Click **Next** to proceed with the installation (see Figure 6-20).

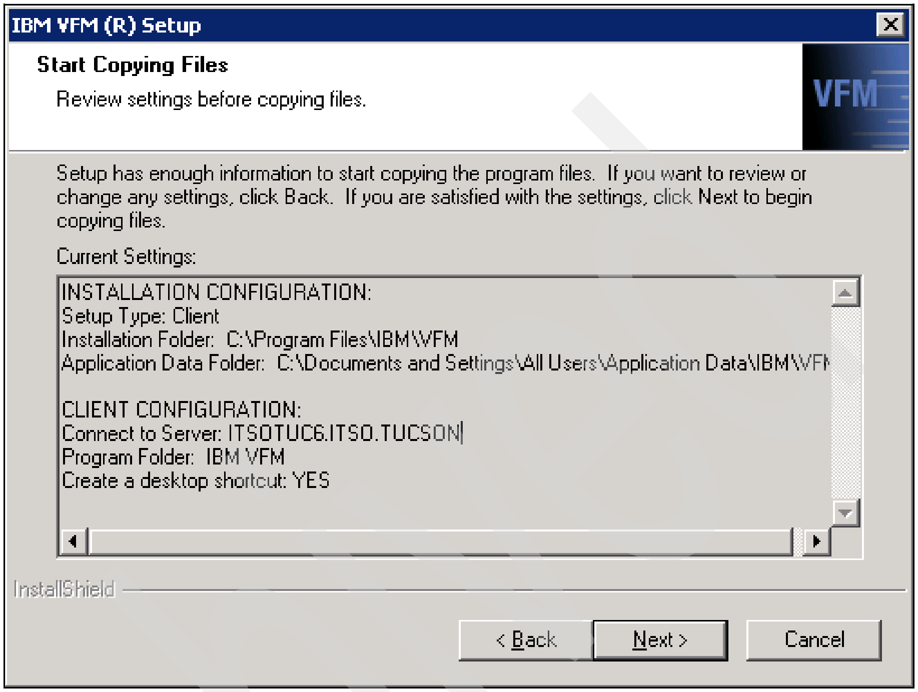

IBM VFM (R) Setup

Start Copying Files

Review settings before copying files.

Setup has enough information to start copying the program files. If you want to review or change any settings, click Back. If you are satisfied with the settings, click Next to begin copying files.

Current Settings:

```
INSTALLATION CONFIGURATION:
Setup Type: Client
Installation Folder:  C:\Program Files\IBM\VFM
Application Data Folder:  C:\Documents and Settings\All Users\Application Data\IBM\VFM

CLIENT CONFIGURATION:
Connect to Server: ITSOTUC6.ITSO.TUCSON
Program Folder:  IBM VFM
Create a desktop shortcut: YES
```

Figure 6-20 Client installation: summary

9. At this point, the installation is complete. A window similar to that shown in Figure 6-21 is displayed. Click **Finish**, and VFM is ready to use.

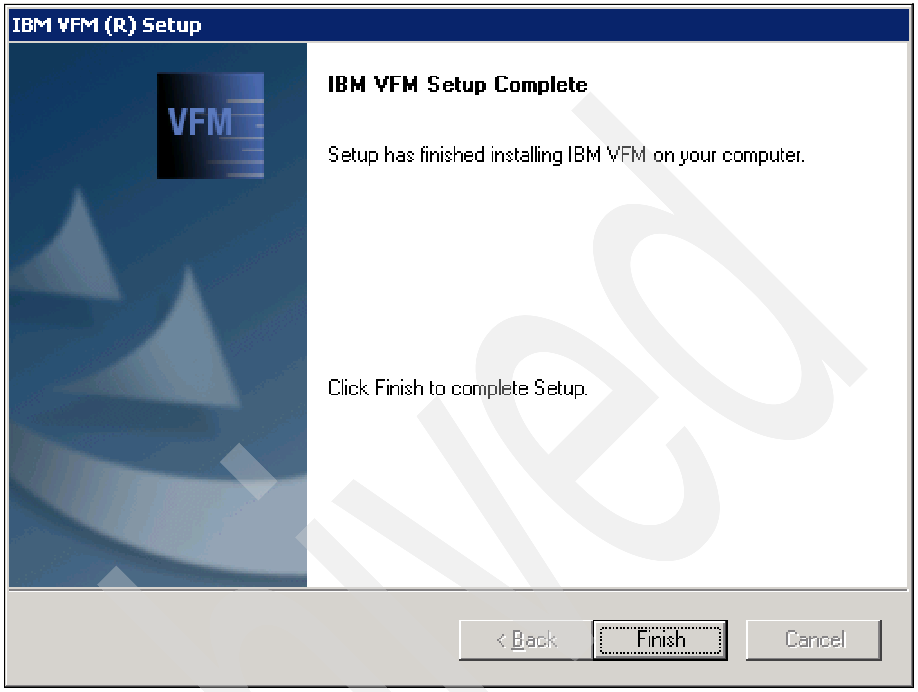

IBM VFM (R) Setup

IBM VFM Setup Complete

Setup has finished installing IBM VFM on your computer.

Click Finish to complete Setup.

Figure 6-21 Client installation: finish

You can open the VFM client by going to **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.

6.3 Installing a Monitoring Agent

With the Monitoring Agent option, you can install a VFM Monitoring Agent alone. We describe the installation in this section.

Follow these steps to install the Monitoring Agent:

1. On the VFM Installation CD, find the Setup.exe file and double-click it.
2. Click **Next** (see Figure 6-22).

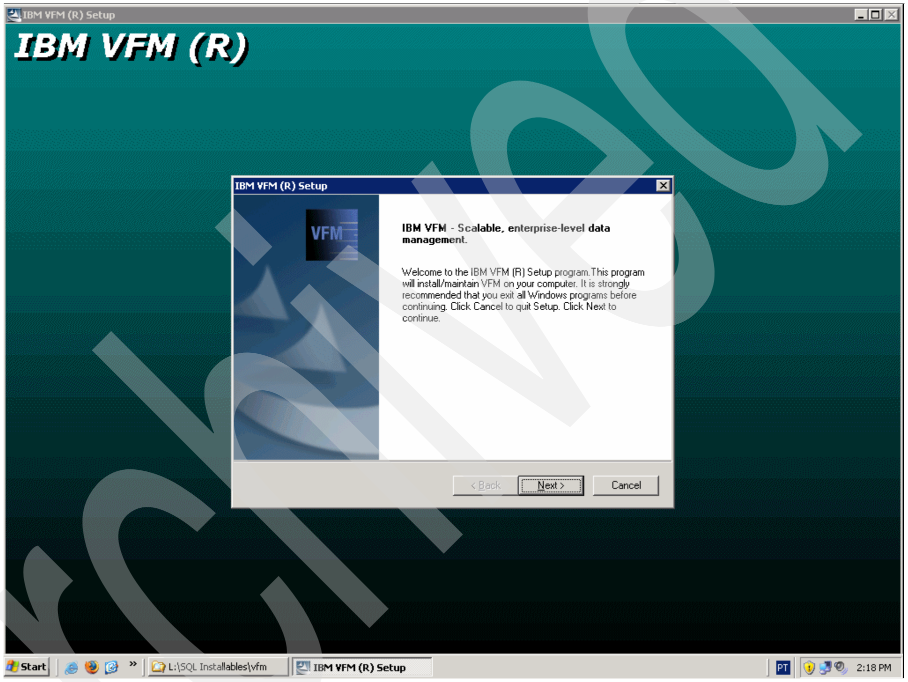

IBM VFM - Scalable, enterprise-level data management.

Welcome to the IBM VFM (R) Setup program.This program will install/maintain VFM on your computer. It is strongly recommended that you exit all Windows programs before continuing. Click Cancel to quit Setup. Click Next to continue.

Figure 6-22 Monitoring Agent installation: first window

3. Read and accept the license to proceed (see Figure 6-23).

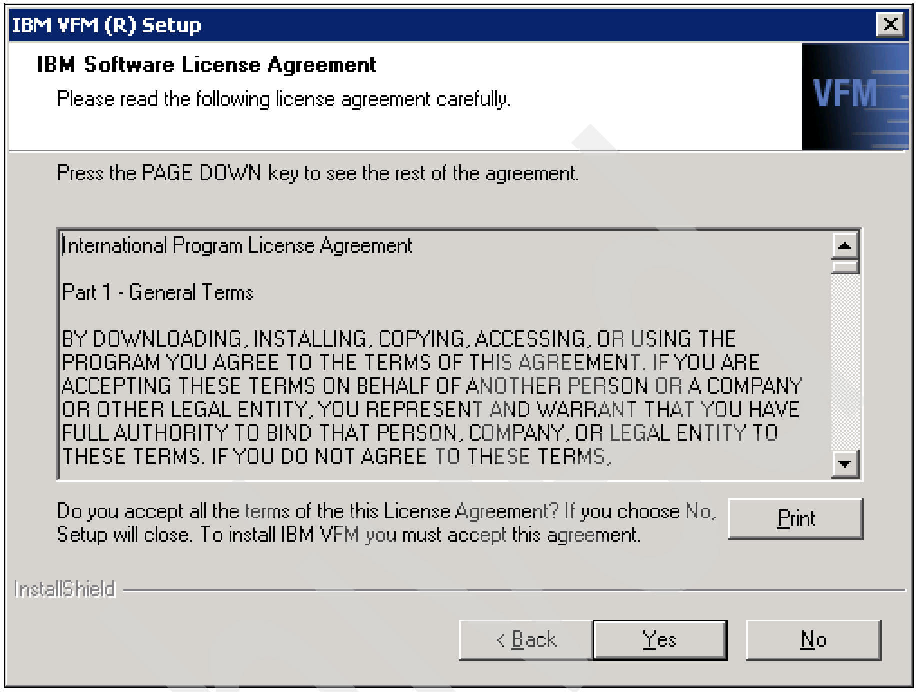
IBM VFM (R) Setup

IBM Software License Agreement

Please read the following license agreement carefully.

Press the PAGE DOWN key to see the rest of the agreement.

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, OR USING THE PROGRAM YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF ANOTHER PERSON OR A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON, COMPANY, OR LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

Do you accept all the terms of the this License Agreement? If you choose No, Setup will close. To install IBM VFM you must accept this agreement.

Figure 6-23 Monitoring Agent installation: license agreement

4. Check the README notes and click **Yes** to proceed (see Figure 6-24).

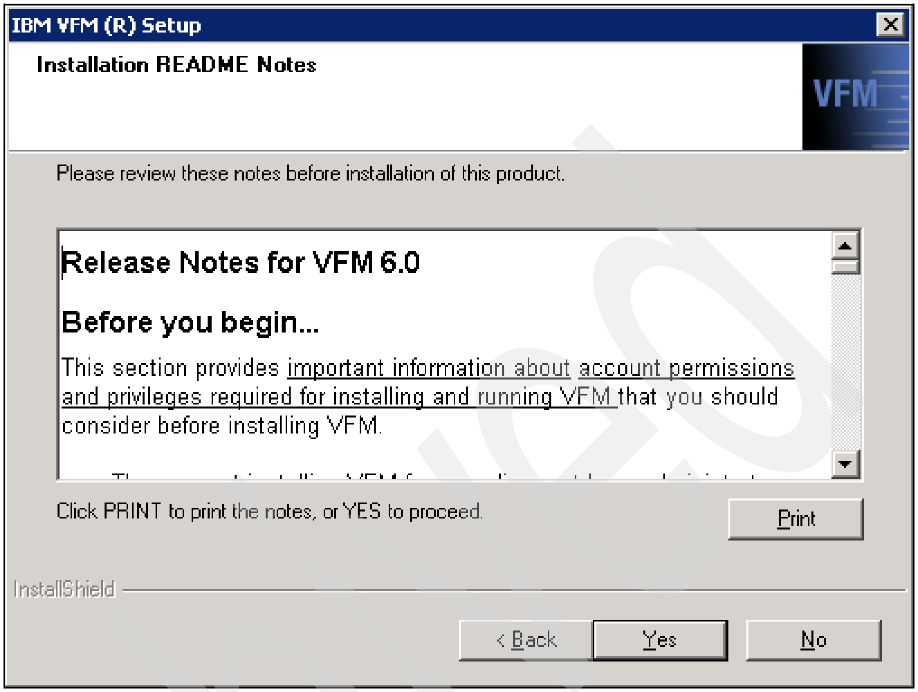

Please review these notes before installation of this product.

Release Notes for VFM 6.0

Before you begin...

This section provides important information about account permissions and privileges required for installing and running VFM that you should consider before installing VFM.

Click PRINT to print the notes, or YES to proceed.

Figure 6-24 Monitoring Agent installation: README notes

5. Choose **Monitoring Agent** and click **Next** (see Figure 6-25). The other setup types are discussed in other sections in this chapter.

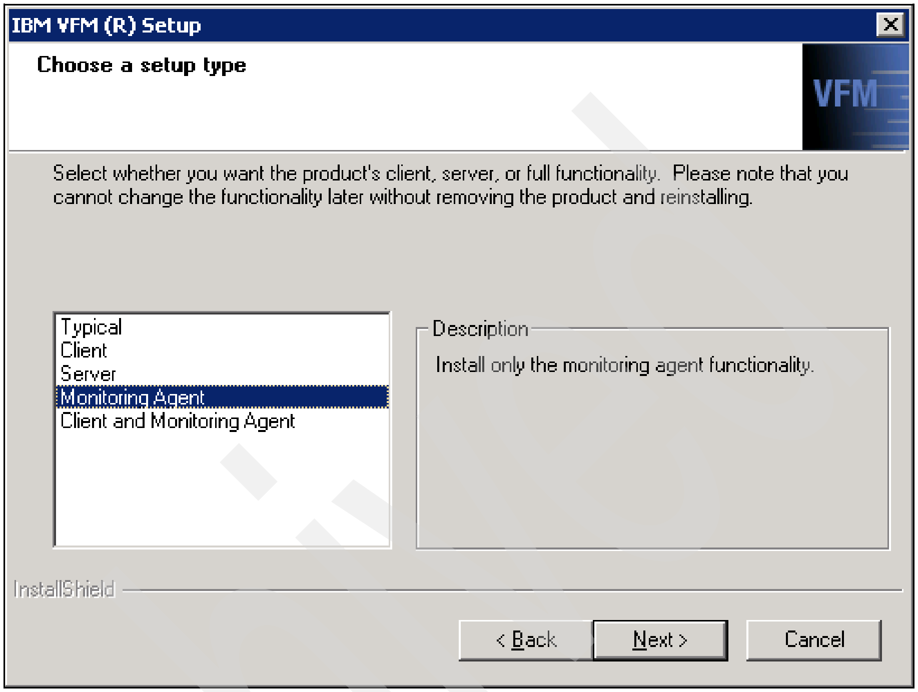

Choose a setup type

Select whether you want the product's client, server, or full functionality. Please note that you cannot change the functionality later without removing the product and reinstalling.

Typical
Client
Server
Monitoring Agent
Client and Monitoring Agent

Description

Install only the monitoring agent functionality.

Figure 6-25 Monitoring Agent installation: choosing setup type

6. If you have an Evaluation license, fill in the form with the Serial Number and the Activation Key, and then click **Next** (see Figure 6-26). Skip step 7 on page 166.

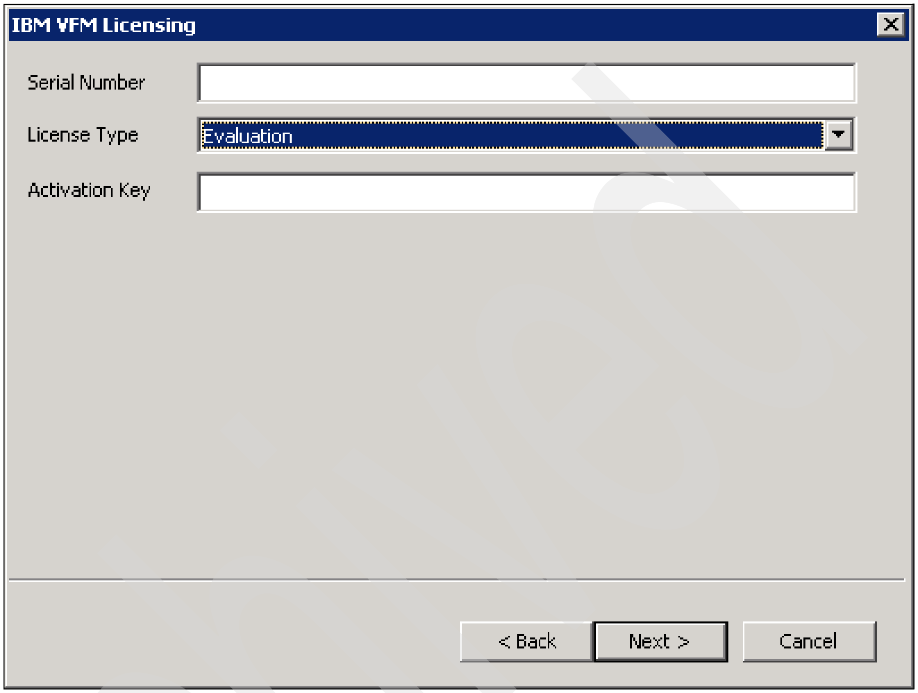

IBM VFM Licensing

Serial Number:

License Type: Evaluation

Activation Key:

Figure 6-26 Monitoring Agent installation: evaluation license

7. If you have a Full license, choose **Full** on the License Type selection box, and a window similar to the one in Figure 6-27 is displayed. Fill in the form with the Serial Number and all the Activation Keys you have received, then click **Next**.

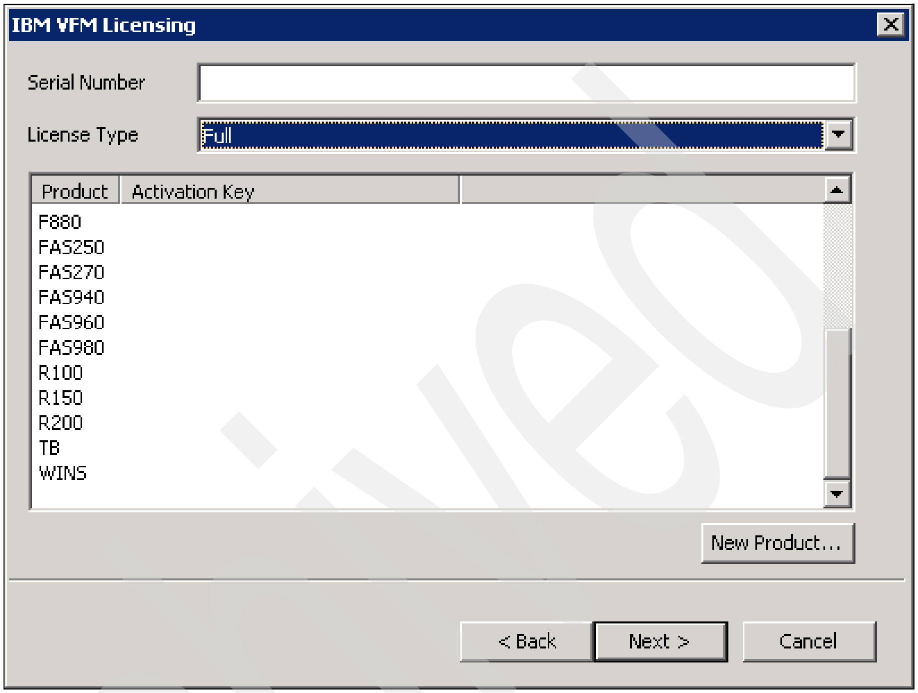

IBM VFM Licensing

Serial Number

License Type: Full

Product	Activation Key
F880	
FAS250	
FAS270	
FAS940	
FAS960	
FAS980	
R100	
R150	
R200	
TB	
WINS	

Figure 6-27 Monitoring Agent installation: full license

8. Choose where VFM will be installed (see Figure 6-28). It must be in an NTFS file system local to the host. The default value is acceptable in most cases.

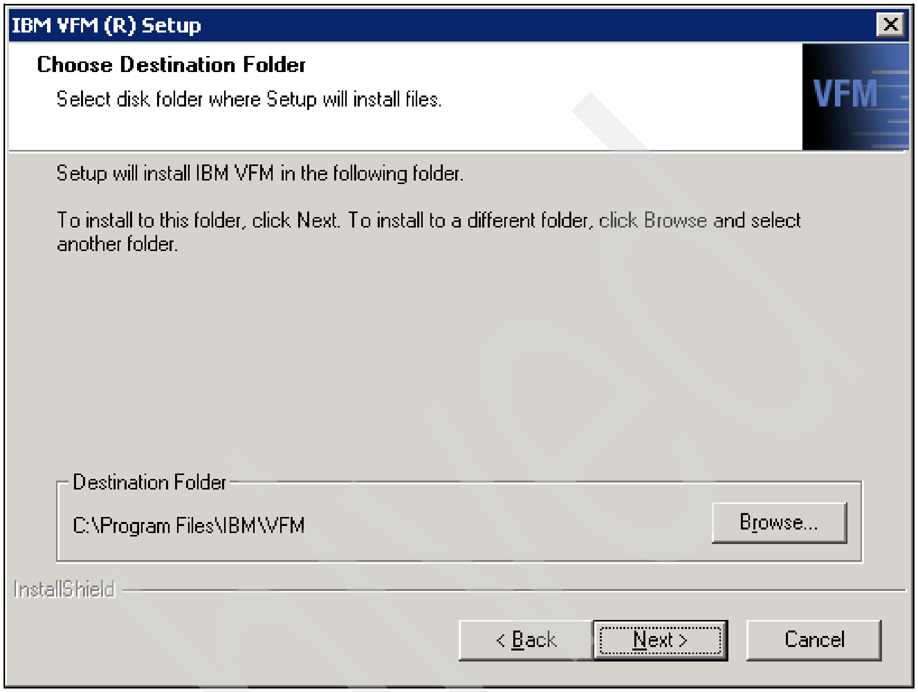

IBM VFM (R) Setup

Choose Destination Folder

Select disk folder where Setup will install files.

Setup will install IBM VFM in the following folder.

To install to this folder, click Next. To install to a different folder, click Browse and select another folder.

Destination Folder

C:\Program Files\IBM\VFM

Figure 6-28 Monitoring Agent installation: destination folder

9. VFM needs to store some files locally. At this point, you can modify where these files are stored (see Figure 6-29). The default value is acceptable in most cases.

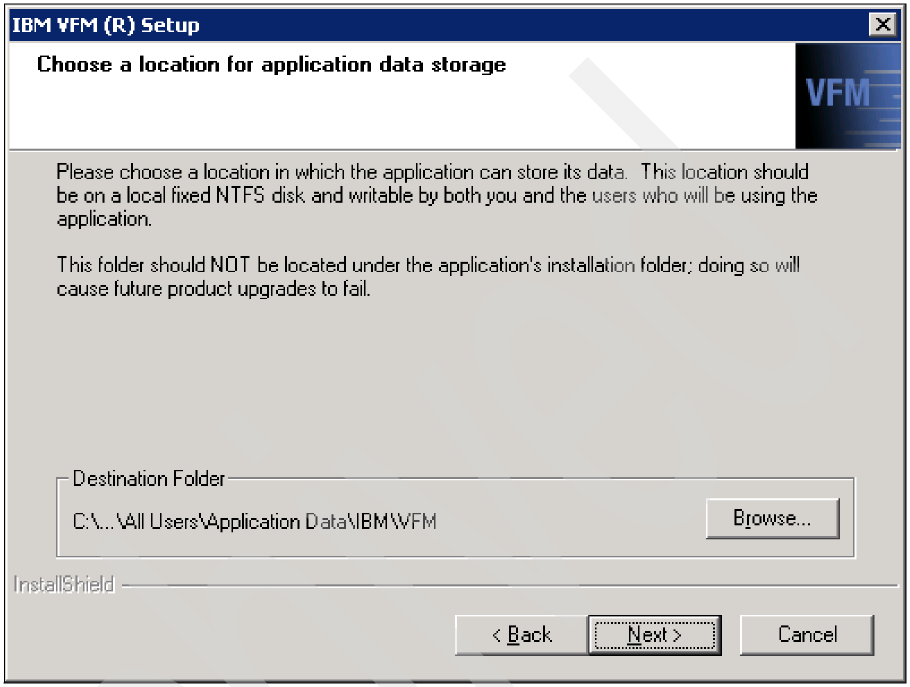

IBM VFM (R) Setup

Choose a location for application data storage

Please choose a location in which the application can store its data. This location should be on a local fixed NTFS disk and writable by both you and the users who will be using the application.

This folder should NOT be located under the application's installation folder; doing so will cause future product upgrades to fail.

Destination Folder

C:\...\All Users\Application Data\IBM\VFM

Figure 6-29 Monitoring Agent installation: local data

10. Choose the VFM server to connect to (Figure 6-30).

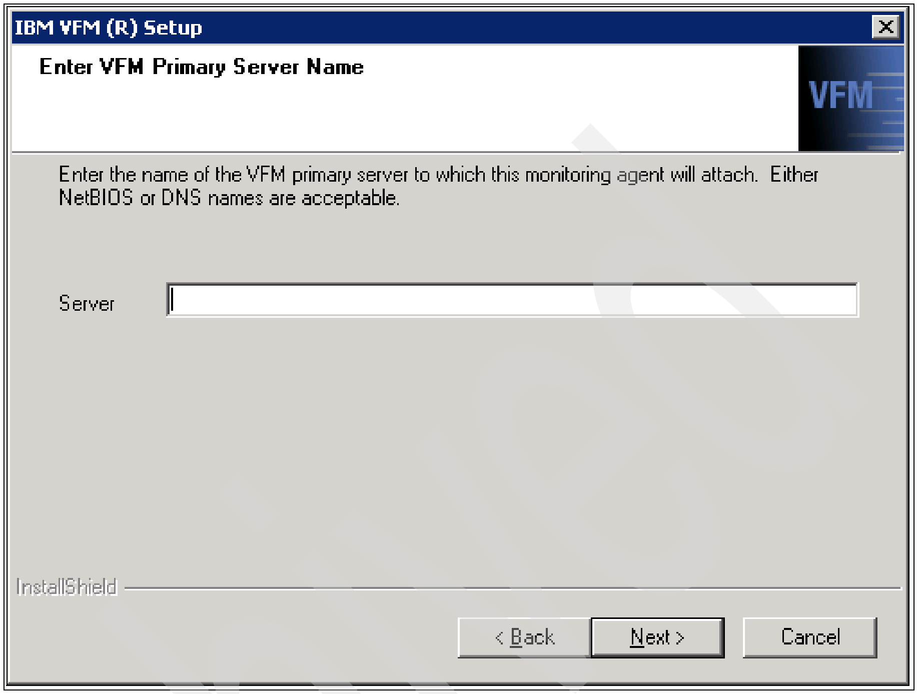

Enter VFM Primary Server Name

Enter the name of the VFM primary server to which this monitoring agent will attach. Either NetBIOS or DNS names are acceptable.

Server

Figure 6-30 Monitoring Agent installation: connecting to VFM server

11. As stated in Chapter 3, “Preparation” on page 83, you must have a service account to run VFM. Enter the information about the account created for it (see Figure 6-31). Remember to use the form DOMAIN\Username.

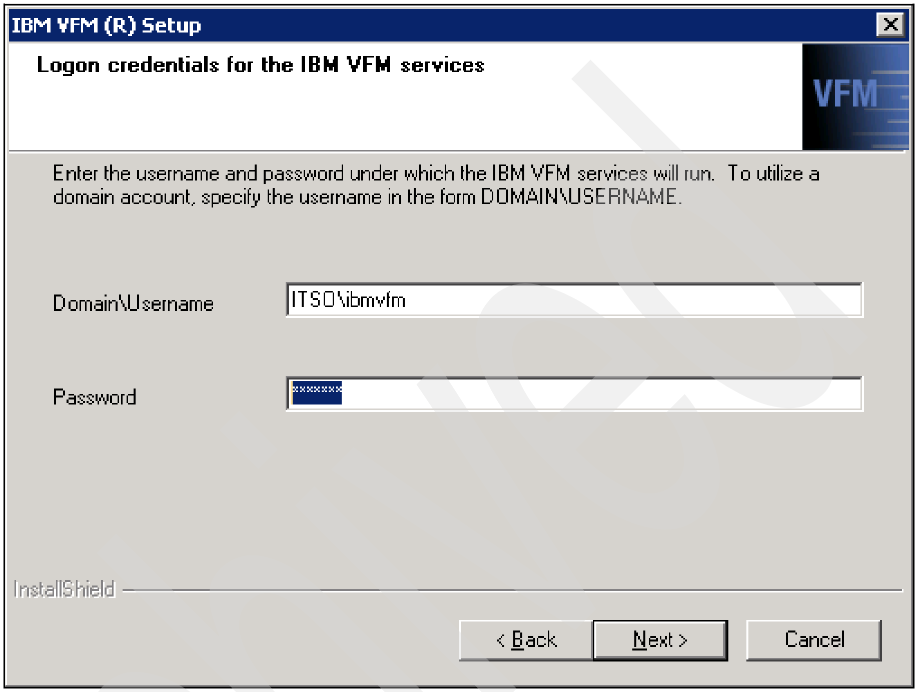

IBM VFM (R) Setup

Logon credentials for the IBM VFM services

Enter the username and password under which the IBM VFM services will run. To utilize a domain account, specify the username in the form DOMAIN\USERNAME.

Domain\Username: ITSO\ibmvfm

Password: *******

Figure 6-31 Monitoring Agent installation: establishing service account

VFM can automatically install the required database server on your machine. Refer to 3.5, “MSSQL” on page 107 for more information about the database server for VFM. If you intend to use a separate database server, refer to 6.5, “Differences when using a separate MSSQL server” on page 188 for steps describing how to do so. In this section, we describe how to use the automatically installed server, MSDE.

1. To automatically install the database server, choose the option **Please install MSDE for me** and click **Next** (see Figure 6-32).

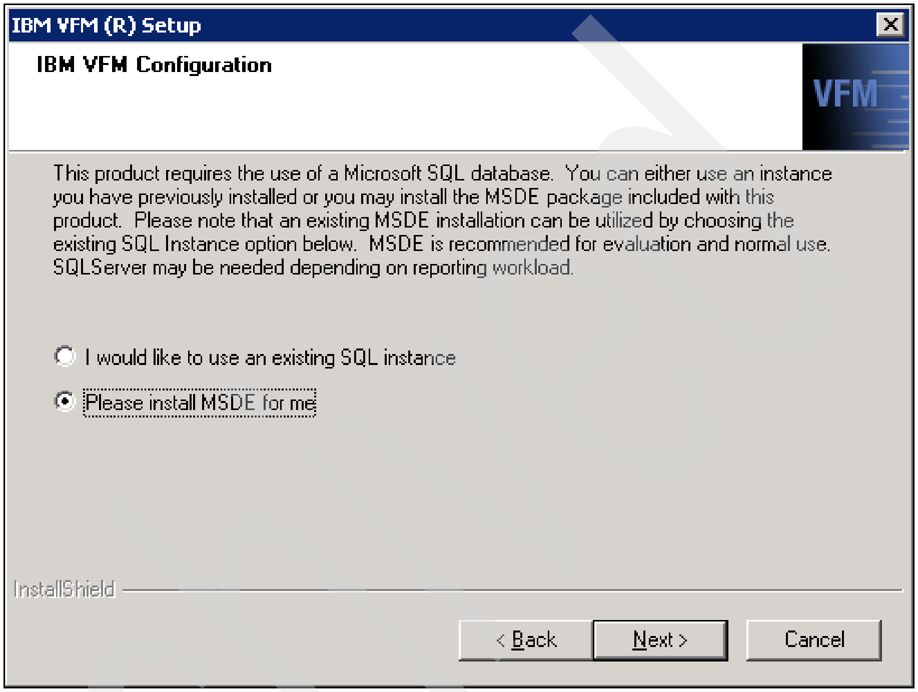

This product requires the use of a Microsoft SQL database. You can either use an instance you have previously installed or you may install the MSDE package included with this product. Please note that an existing MSDE installation can be utilized by choosing the existing SQL Instance option below. MSDE is recommended for evaluation and normal use. SQLServer may be needed depending on reporting workload.

○ I would like to use an existing SQL instance

● Please install MSDE for me

Figure 6-32 Monitoring Agent installation: installing database server

2. A summary of the chosen options is displayed. Click **Next** to proceed with the installation (see Figure 6-33).

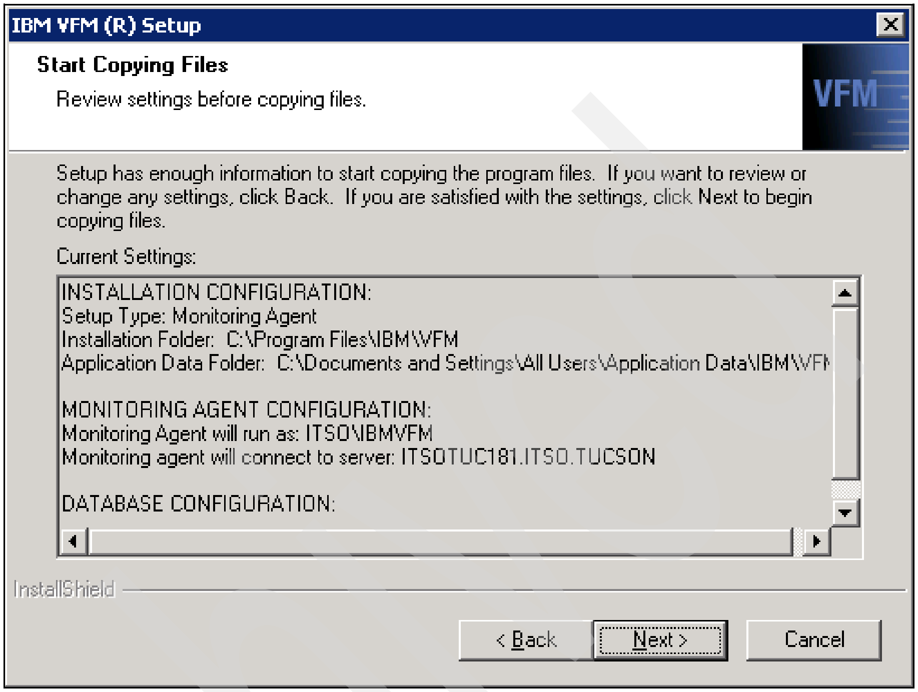

Setup has enough information to start copying the program files. If you want to review or change any settings, click Back. If you are satisfied with the settings, click Next to begin copying files.

Current Settings:

```
INSTALLATION CONFIGURATION:
Setup Type: Monitoring Agent
Installation Folder:  C:\Program Files\IBM\VFM
Application Data Folder:  C:\Documents and Settings\All Users\Application Data\IBM\VFM

MONITORING AGENT CONFIGURATION:
Monitoring Agent will run as: ITSO\IBMVFM
Monitoring agent will connect to server: ITSOTUC181.ITSO.TUCSON

DATABASE CONFIGURATION:
```

Figure 6-33 Monitoring Agent installation: summary

3. The installation is complete; a window similar to that shown in Figure 6-34 is displayed. Click **Finish**, and the VFM Monitoring Agent is ready to use.

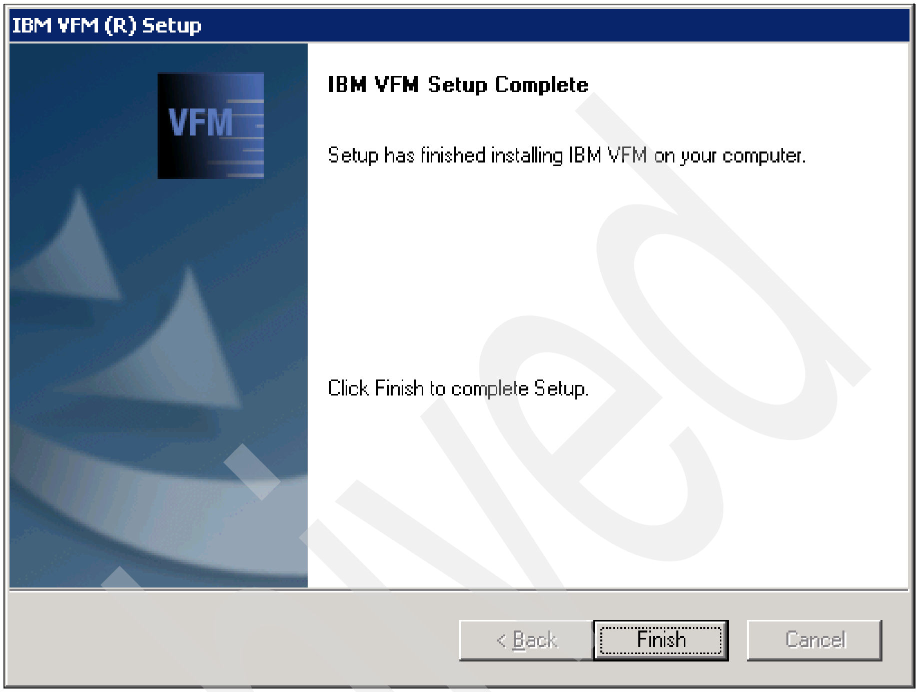

IBM VFM (R) Setup

IBM VFM Setup Complete

Setup has finished installing IBM VFM on your computer.

Click Finish to complete Setup.

Figure 6-34 Monitoring Agent installation: finish

6.4 Installing a client and Monitoring Agent together

You can install a client and a Monitoring Agent together on the same machine. Follow these steps:

1. On the VFM Installation CD, find the Setup.exe file and double-click it.
2. Click **Next** (see Figure 6-35).

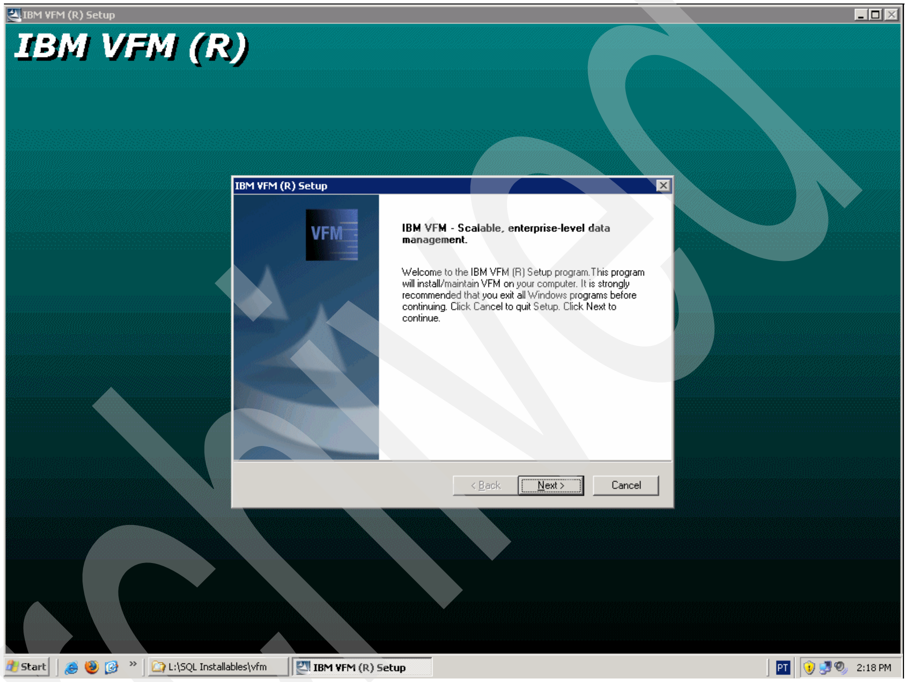

IBM VFM - Scalable, enterprise-level data management.

Welcome to the IBM VFM (R) Setup program.This program will install/maintain VFM on your computer. It is strongly recommended that you exit all Windows programs before continuing. Click Cancel to quit Setup. Click Next to continue.

Figure 6-35 Monitoring Agent and client: first window

3. Read and accept the license to proceed (see Figure 6-36).

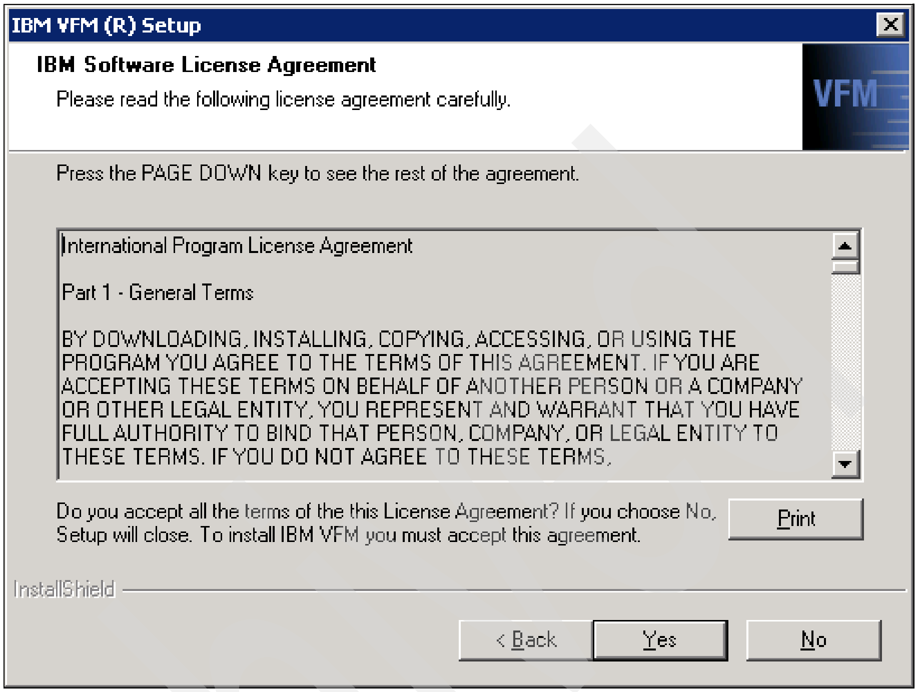

IBM Software License Agreement

Please read the following license agreement carefully.

Press the PAGE DOWN key to see the rest of the agreement.

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, OR USING THE PROGRAM YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF ANOTHER PERSON OR A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON, COMPANY, OR LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

Do you accept all the terms of the this License Agreement? If you choose No, Setup will close. To install IBM VFM you must accept this agreement.

Figure 6-36 Monitoring Agent and client: license agreement

4. Check the README notes, and click **Yes** to proceed (see Figure 6-37).

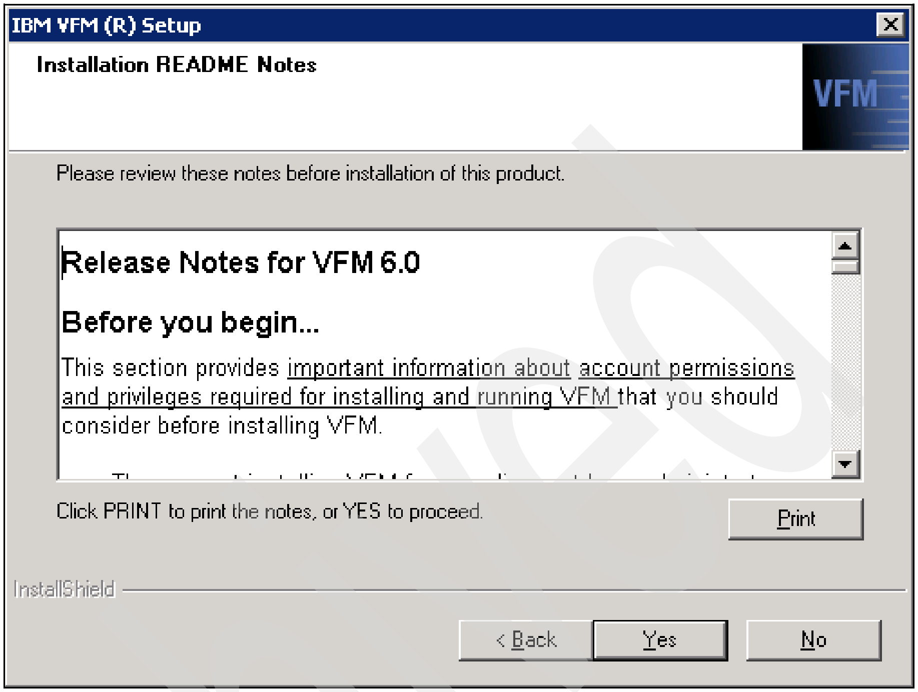

Installation README Notes

Please review these notes before installation of this product.

Release Notes for VFM 6.0

Before you begin...

This section provides important information about account permissions and privileges required for installing and running VFM that you should consider before installing VFM.

Click PRINT to print the notes, or YES to proceed.

Figure 6-37 Monitoring Agent and client: README notes

5. Choose **Client and Monitoring Agent** and click **Next** (see Figure 6-38). The other setup types are discussed in other sections in this chapter.

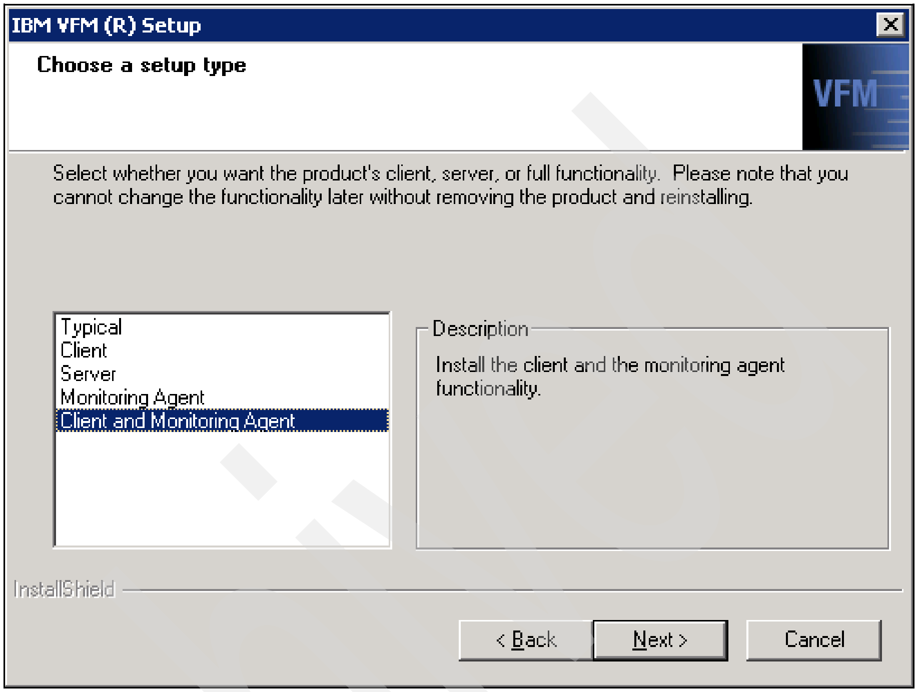

IBM VFM (R) Setup

Choose a setup type

Select whether you want the product's client, server, or full functionality. Please note that you cannot change the functionality later without removing the product and reinstalling.

Description

Install the client and the monitoring agent functionality.

Figure 6-38 Monitoring Agent and client: setup type

6. If you have an Evaluation license, fill in the form with the Serial Number and the Activation Key, then click **Next** (see Figure 6-39). Skip step 7 on page 179.

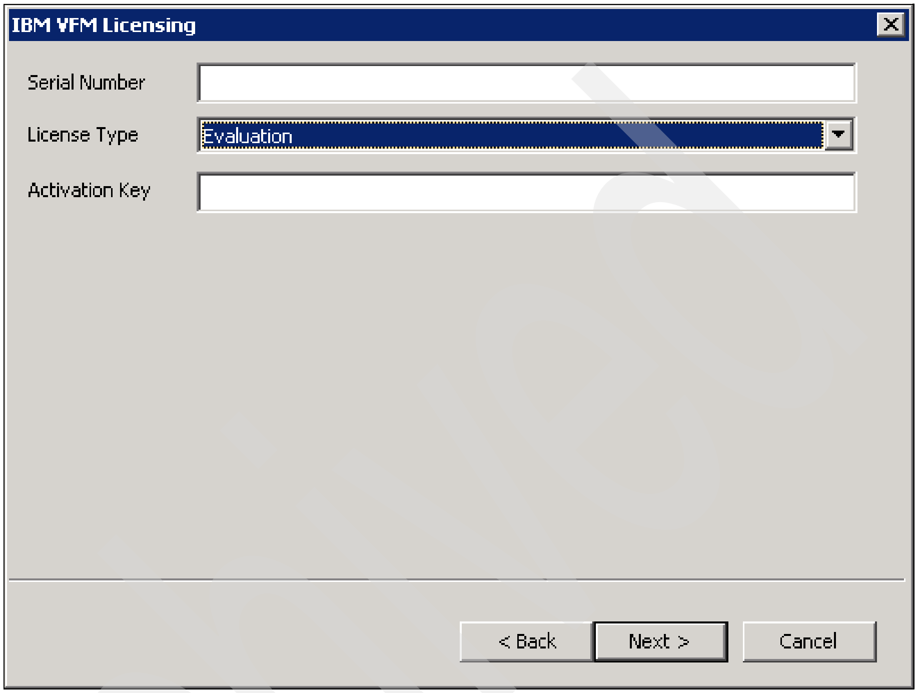

IBM VFM Licensing

Serial Number

License Type: Evaluation

Activation Key

Figure 6-39 Monitoring Agent and client: evaluation license

7. If you have a Full license, choose **Full** on the License Type selection box, and a window similar to the one shown in Figure 6-40 is displayed. Fill in the form with the Serial Number and all the Activation Keys you have received, then click **Next**.

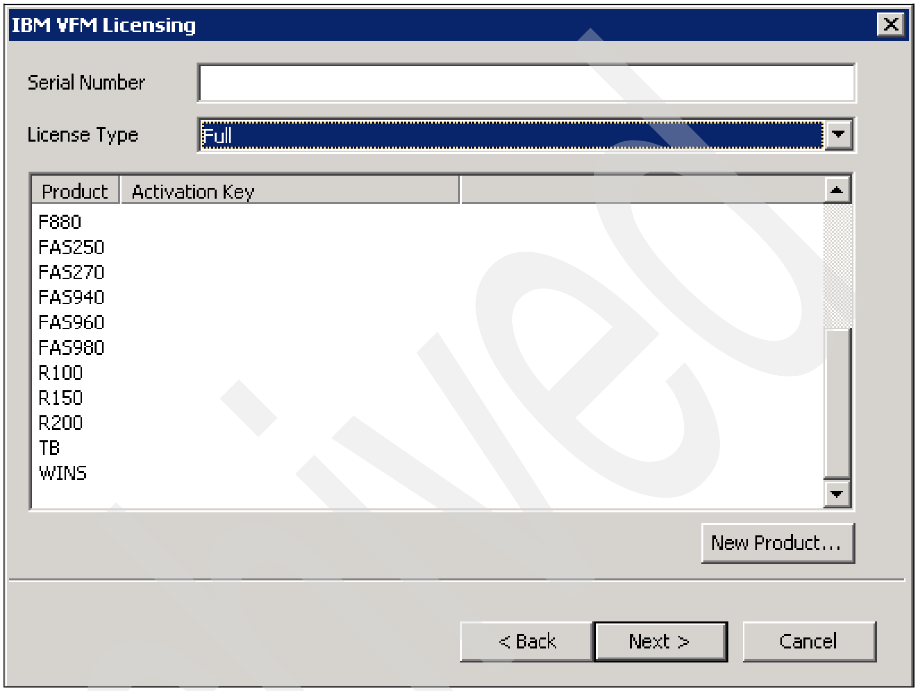

IBM VFM Licensing

Serial Number

License Type: Full

Product	Activation Key
F880	
FAS250	
FAS270	
FAS940	
FAS960	
FAS980	
R100	
R150	
R200	
TB	
WINS	

Figure 6-40 Monitoring Agent and client: full license

8. Choose where you intend to install VFM (see Figure 6-41). It must be in an NTFS file system local to the host. The default value is acceptable in most cases.

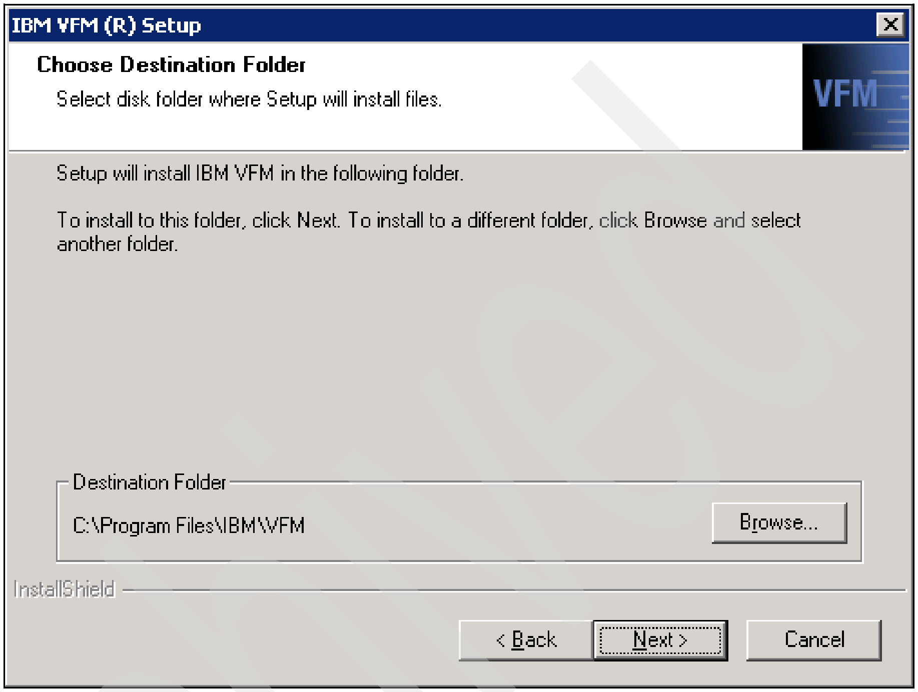
IBM VFM (R) Setup

Choose Destination Folder

Select disk folder where Setup will install files.

Setup will install IBM VFM in the following folder.

To install to this folder, click Next. To install to a different folder, click Browse and select another folder.

Destination Folder

C:\Program Files\IBM\VFM

Figure 6-41 Monitoring Agent and client: destination folder

9. VFM has to store some files locally. At this point, you can modify where these files are stored (see Figure 6-42). The default value is acceptable in most cases.

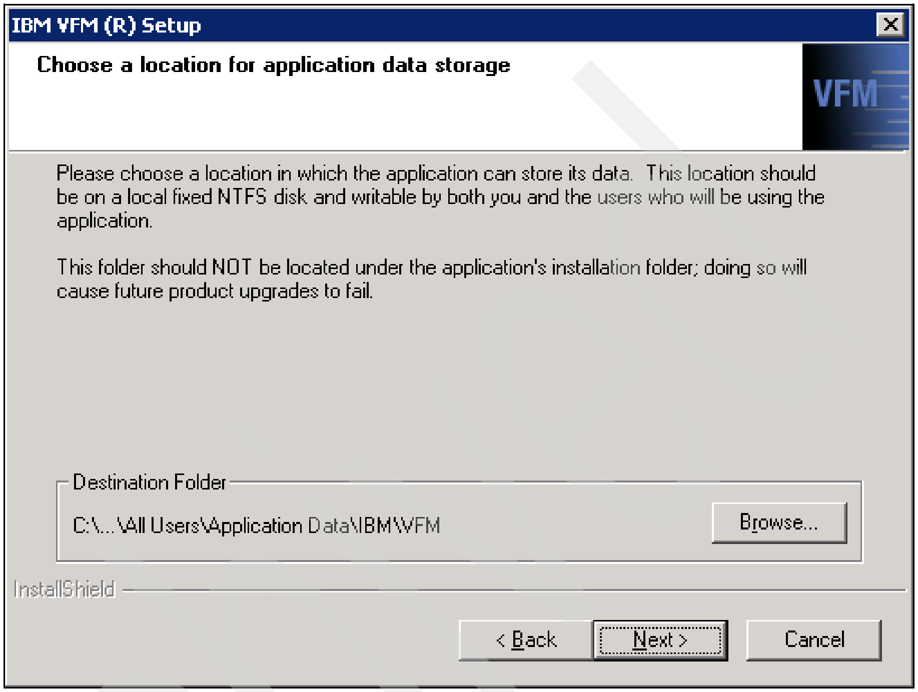

IBM VFM (R) Setup

Choose a location for application data storage

Please choose a location in which the application can store its data. This location should be on a local fixed NTFS disk and writable by both you and the users who will be using the application.

This folder should NOT be located under the application's installation folder; doing so will cause future product upgrades to fail.

Destination Folder

C:\...\All Users\Application Data\IBM\VFM

Figure 6-42 Monitoring Agent and client: local data

10.Now choose the VFM server to connect to (see Figure 6-43), then click **Next**.

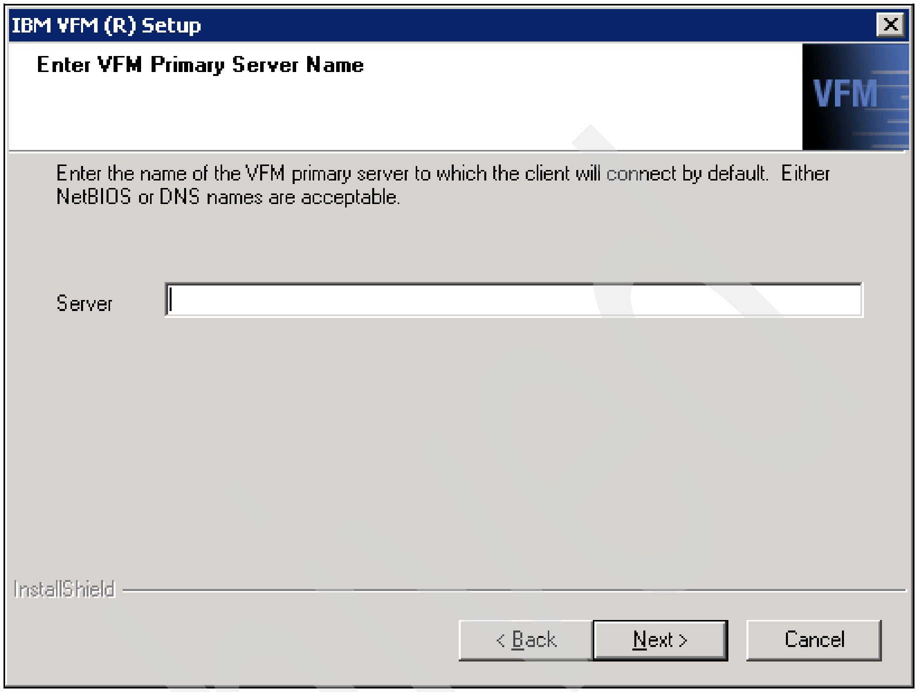

Enter VFM Primary Server Name

Enter the name of the VFM primary server to which the client will connect by default. Either NetBIOS or DNS names are acceptable.

Server

Figure 6-43 Monitoring Agent and client: server address

11. Choose whether or not you want to create a desktop shortcut for the VFM client (see Figure 6-44). Either way, the VFM client is available on **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.

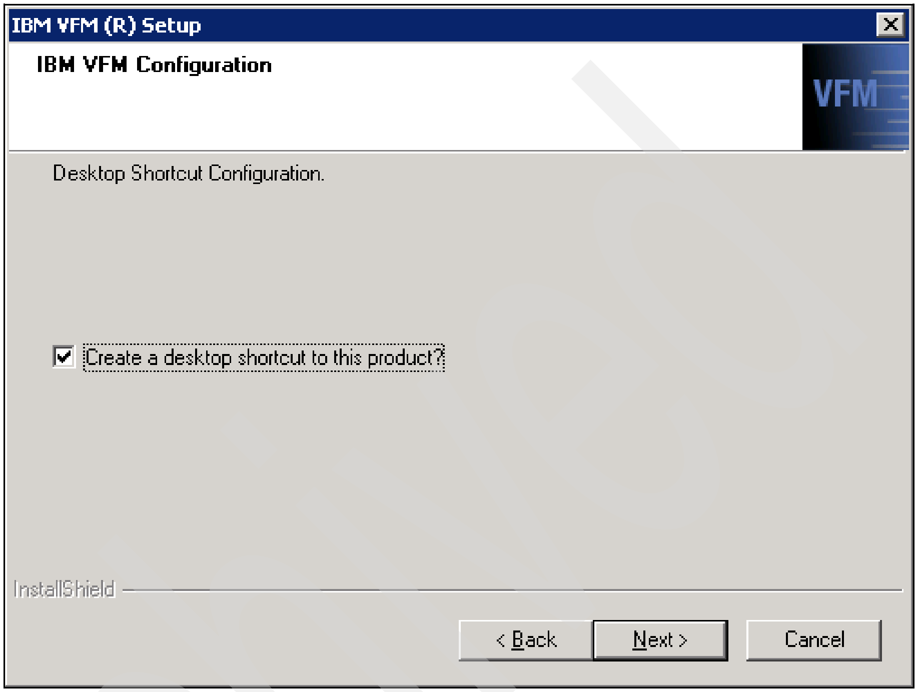

IBM VFM Configuration

Desktop Shortcut Configuration.

☑ Create a desktop shortcut to this product?

Figure 6-44 Monitoring Agent and client: desktop shortcut

12. As stated in Chapter 3, “Preparation” on page 83, you must have a service account to run VFM. Enter the information about the account created for it (see Figure 6-45). Remember to use the form DOMAIN\Username.

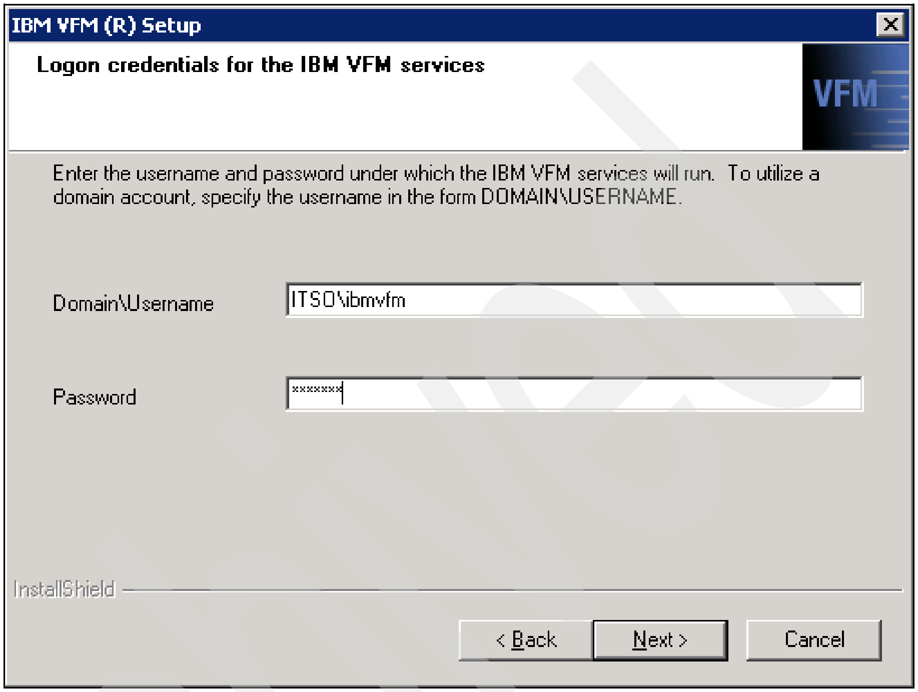
IBM VFM (R) Setup

Logon credentials for the IBM VFM services

Enter the username and password under which the IBM VFM services will run. To utilize a domain account, specify the username in the form DOMAIN\USERNAME.

Domain\Username: ITSO\ibmvfm

Password: *******

Figure 6-45 Monitoring Agent and client: VFM service account

13. VFM can automatically install the required database server on your machine. Refer to 3.5, “MSSQL” on page 107 for more information about the database server for VFM. If you intend to use a separate database server, refer to 6.5, “Differences when using a separate MSSQL server” on page 188 for steps

describing how to do so. In this section, we describe how to use the automatically installed server, MSDE.

To automatically install the database server, choose the option **Please install MSDE for me** and click **Next** (see Figure 6-46).

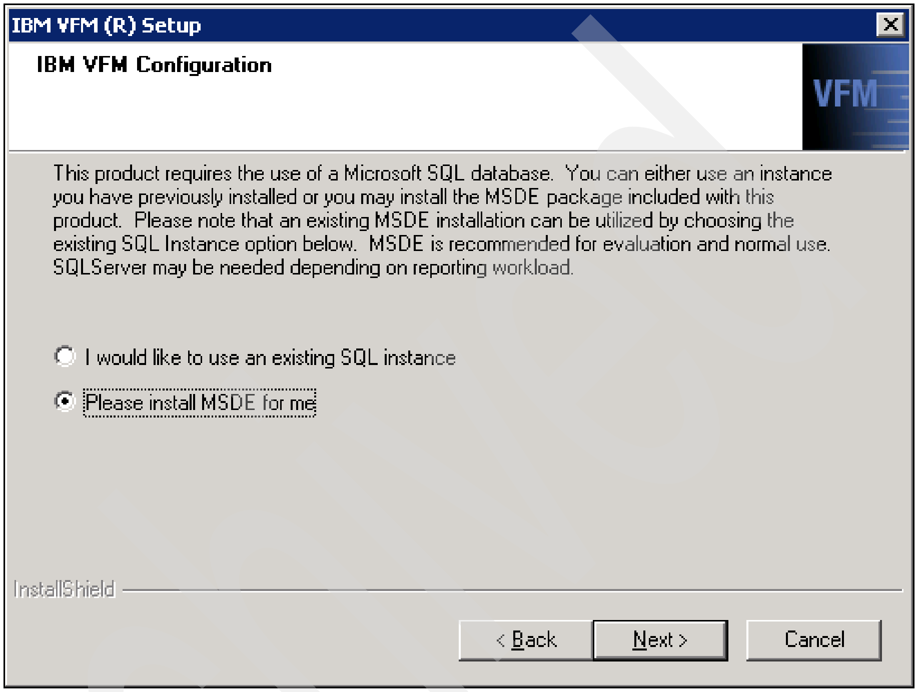

This product requires the use of a Microsoft SQL database. You can either use an instance you have previously installed or you may install the MSDE package included with this product. Please note that an existing MSDE installation can be utilized by choosing the existing SQL Instance option below. MSDE is recommended for evaluation and normal use. SQLServer may be needed depending on reporting workload.

○ I would like to use an existing SQL instance

● Please install MSDE for me

Figure 6-46 Monitoring Agent and client: SQL server

14. A summary of the chosen options. Click **Next** to proceed with installation (see Figure 6-47).

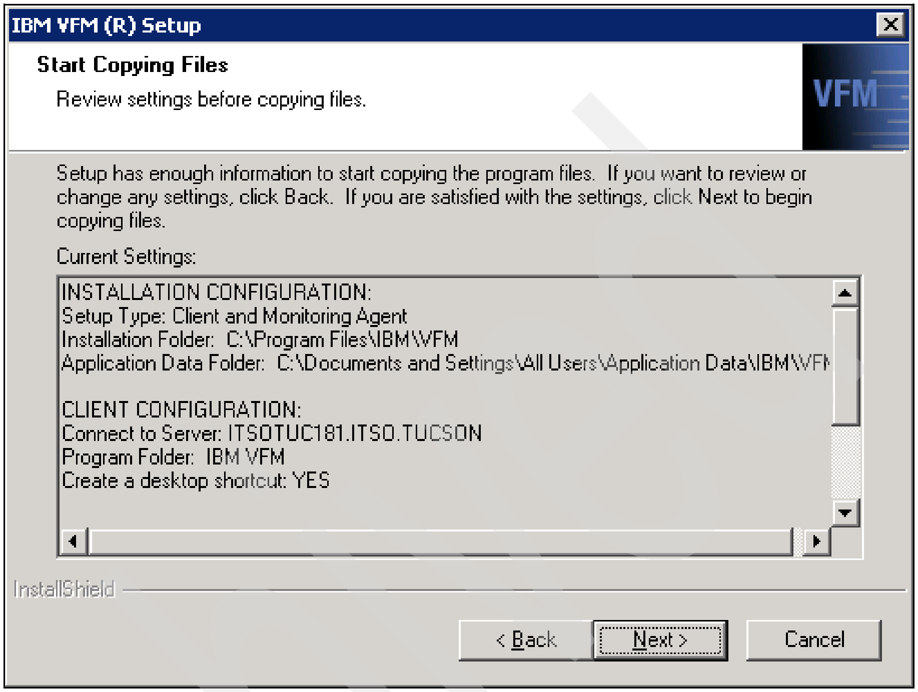

Start Copying Files

Review settings before copying files.

Setup has enough information to start copying the program files. If you want to review or change any settings, click Back. If you are satisfied with the settings, click Next to begin copying files.

Current Settings:

```
INSTALLATION CONFIGURATION:
Setup Type: Client and Monitoring Agent
Installation Folder:  C:\Program Files\IBM\VFM
Application Data Folder:  C:\Documents and Settings\All Users\Application Data\IBM\VFM

CLIENT CONFIGURATION:
Connect to Server: ITSOTUC181.ITSO.TUCSON
Program Folder:  IBM VFM
Create a desktop shortcut: YES
```

Figure 6-47 Monitoring Agent and client: summary

15. The installation is complete, and a window similar to that shown in Figure 6-48 is displayed. Click **Finish**, and the VFM Monitoring Agent and client are ready to use.

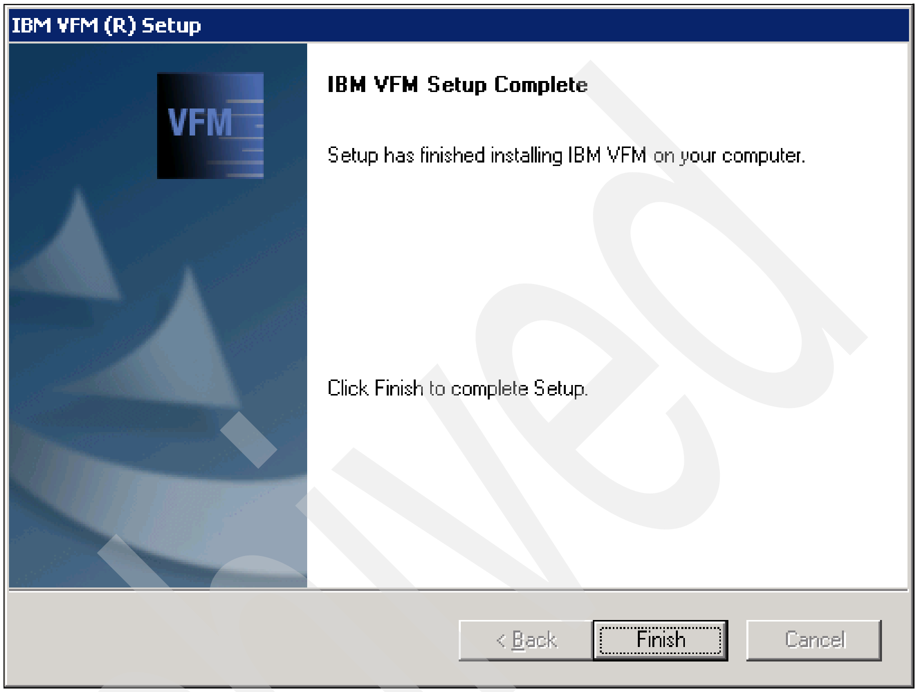

IBM VFM (R) Setup

IBM VFM Setup Complete

Setup has finished installing IBM VFM on your computer.

Click Finish to complete Setup.

Figure 6-48 Monitoring Agent and client: finish

You can open the VFM client by going to **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.

6.5 Differences when using a separate MSSQL server

When using a separate MSSQL server, you must follow some different steps than those described in previous sections. In this section, we briefly describe these steps.

Follow these steps when using a separate MSSQL server:

1. On the installation step regarding SQL server, you select **I would like to use an existing SQL instance** instead of **Please Install a MSDE for me** (see Figure 6-49).

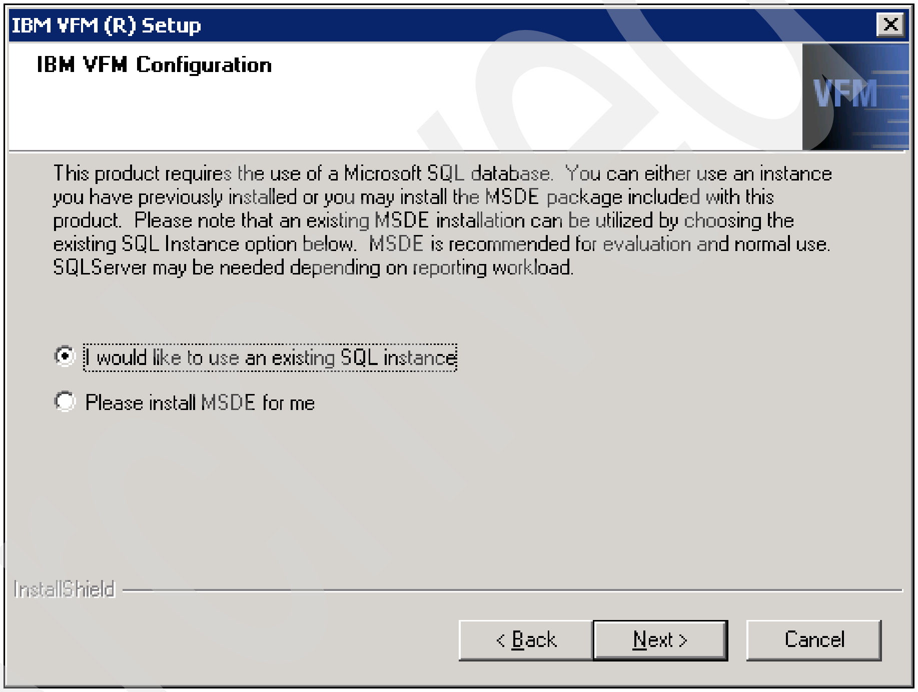

IBM VFM (R) Setup

IBM VFM Configuration

This product requires the use of a Microsoft SQL database. You can either use an instance you have previously installed or you may install the MSDE package included with this product. Please note that an existing MSDE installation can be utilized by choosing the existing SQL Instance option below. MSDE is recommended for evaluation and normal use. SQLServer may be needed depending on reporting workload.

● I would like to use an existing SQL instance

○ Please install MSDE for me

Figure 6-49 MSSQL: choosing an existing instance

2. You can specify the location of the database server and the authentication method to use (see Figure 6-50).

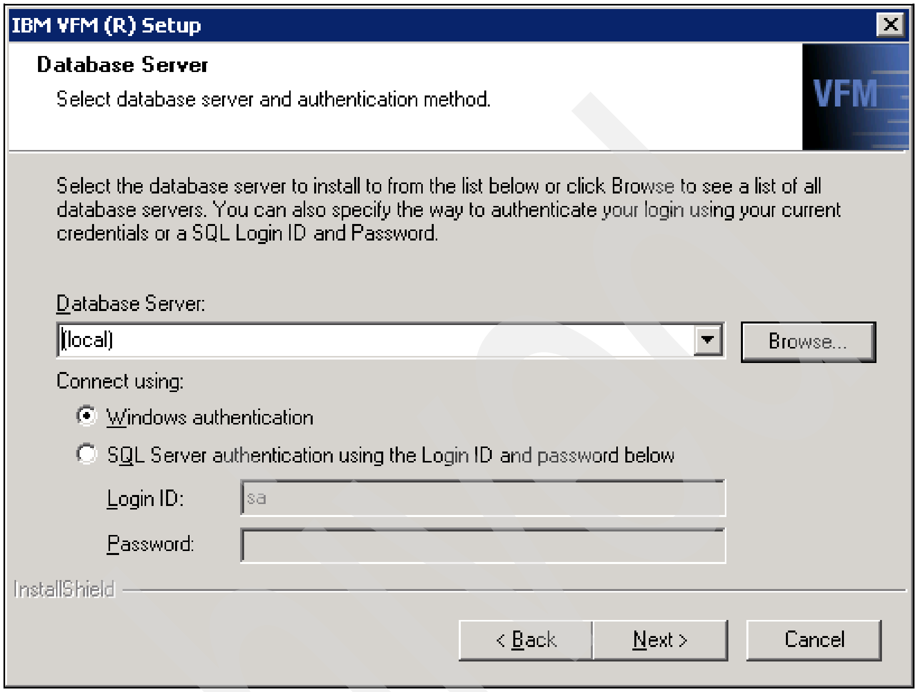

IBM VFM (R) Setup

Database Server

Select database server and authentication method.

Select the database server to install to from the list below or click Browse to see a list of all database servers. You can also specify the way to authenticate your login using your current credentials or a SQL Login ID and Password.

Database Server: (local)

Connect using:

● Windows authentication

○ SQL Server authentication using the Login ID and password below

Login ID: sa

Password:

Figure 6-50 MSSQL specifying database server

3. Choose a name for the VFM database (see Figure 6-51).

Note: Use different database names if you install VFM Monitoring Agents using the same database server, so that each VFM Monitoring Agent and the VFM server uses its own database.

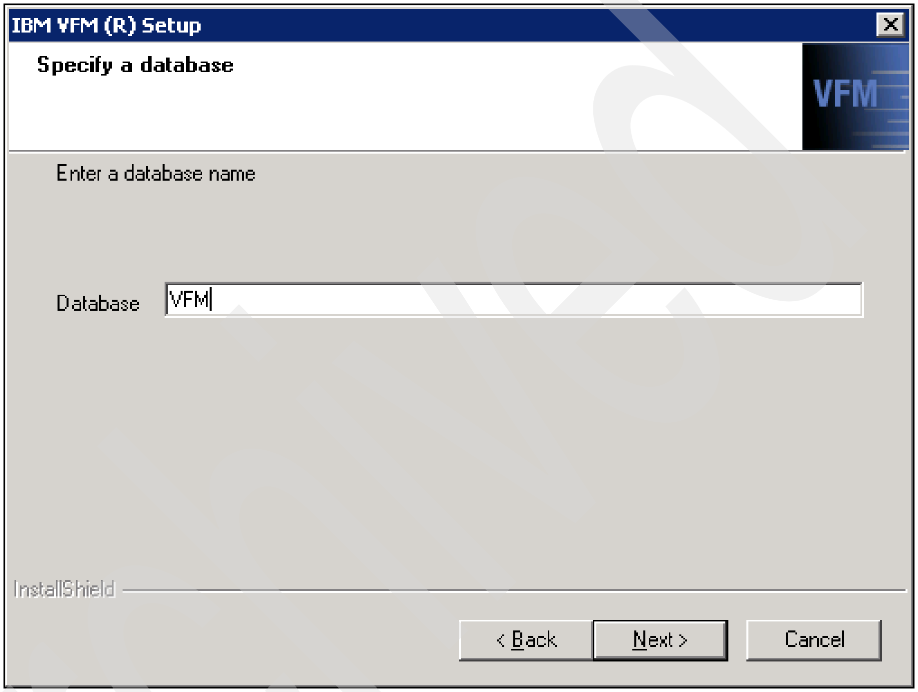

IBM VFM (R) Setup

Specify a database

Enter a database name

Database: VFM

Figure 6-51 MSSQL specifying database name

7

Intelligent data movement

IBM System Storage N series Virtual File Manager (VFM) software is a comprehensive solution for managing unstructured file data. It is designed to provide data management functionality for server and storage consolidation, migration, remote office data management, and disaster recovery features while avoiding disruption to users. It provides all of this functionality through automated policy-based data management leveraging a global namespace.

VFM is designed to provide nondisruptive data migration. VFM facilitates data movement and optimization of existing storage while avoiding impact on users. For users, VFM provides simplified and consistent access even when the underlying storage infrastructure changes.

The following client problems and scenarios are related to data movement:

- Moving from non-IBM storage to IBM storage
- Machine goes off lease (for example, old low-capacity drives are replaced by new high-capacity disks)
- Physical consolidation of servers to save on space, power, and licensing of many servers
- Logical consolidation: Moving relevant data scattered across different servers to one server
- Switching from NetWare to N series
- Storage capacity of servers is limited

VFM creates a global namespace that aggregates distributed files located on IBM N series storage systems to present a single logical “pool” of storage (see Figure 7-1). This helps you easily and quickly change, add, migrate, or consolidate storage while avoiding impact on users.

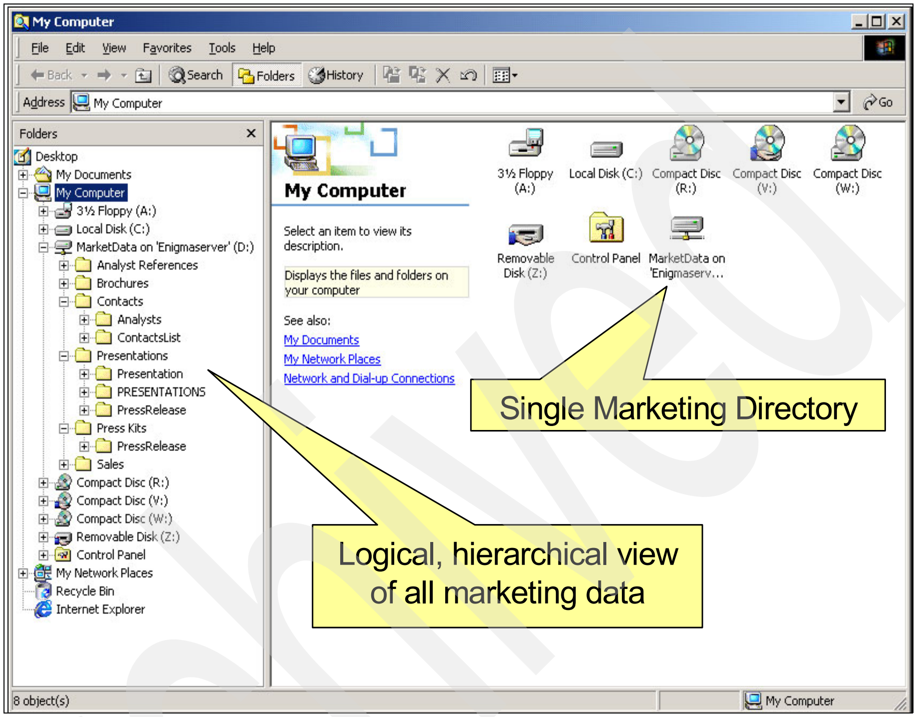

Select an item to view its description.

Displays the files and folders on your computer

Single Marketing Directory

Logical, hierarchical view of all marketing data

Figure 7-1 Single logical pool of storage

IBM System Storage N series Virtual File Manager (VFM) provides the following intelligent data movement features:

- Simplified and automated server migration
- Policy-based replication
- Nondisruptive storage expansion and consolidation

7.1 Simplified and automated server migration

VFM simplifies and automates migration across heterogeneous file systems. It also ensures the accuracy of copied files across heterogeneous file systems, including time stamps, security (local groups, secured files), SACL, data streams, and open files.

VFM also provides tools to automate migration. Administrators can automate the creation of shares, copying files, deleting files and shares, changing layout, and scheduling.

VFM maintains audit trails as a history of migration events (see Figure 7-2).

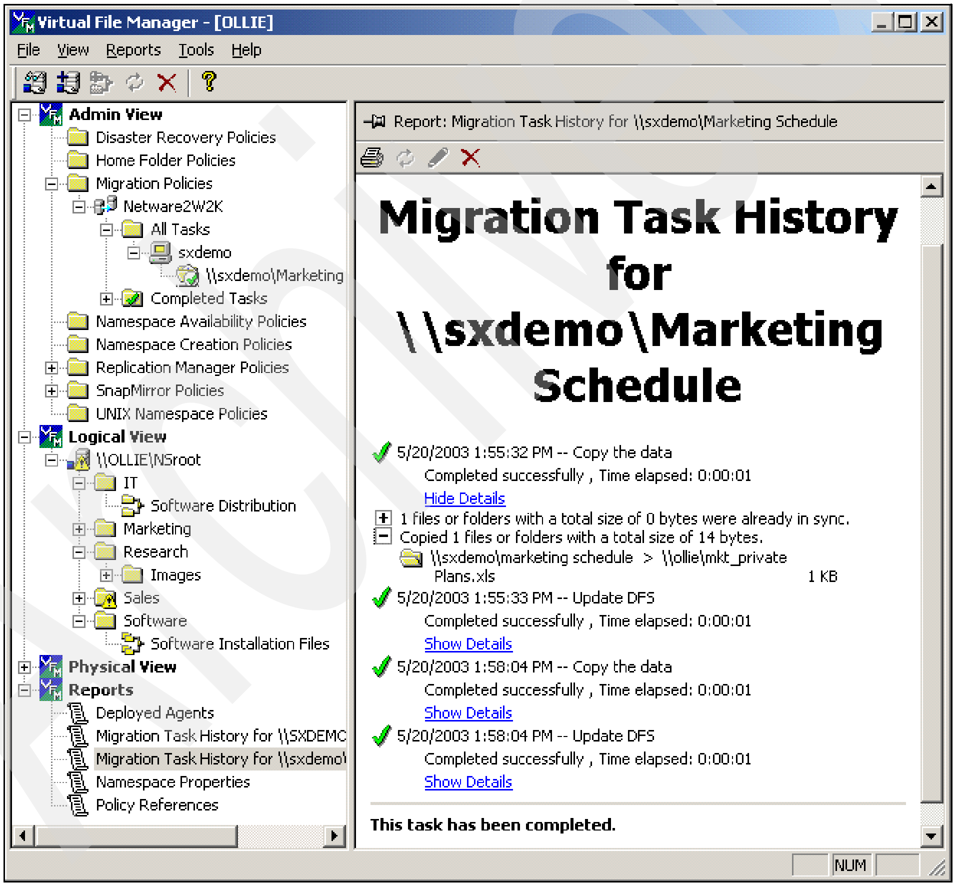

Migration Task History for \\sxdemo\Marketing Schedule

✓ 5/20/2003 1:55:32 PM -- Copy the data
Completed successfully , Time elapsed: 0:00:01
Hide Details
⊞ 1 files or folders with a total size of 0 bytes were already in sync.
⊟ Copied 1 files or folders with a total size of 14 bytes.
\\sxdemo\marketing schedule > \\ollie\mkt_private
Plans.xls 1 KB

✓ 5/20/2003 1:55:33 PM -- Update DFS
Completed successfully , Time elapsed: 0:00:01
Show Details

✓ 5/20/2003 1:58:04 PM -- Copy the data
Completed successfully , Time elapsed: 0:00:01
Show Details

✓ 5/20/2003 1:58:04 PM -- Update DFS
Completed successfully , Time elapsed: 0:00:01
Show Details

This task has been completed.

Figure 7-2 Maintain audit trails of migration task history

VFM provides substantial improvement in migration tasks performance over other traditional methods of migration. VFM facilitates migration by breaking the entire migration process into three phases:

Baseline copy
Creates a baseline copy before running multiple schedules of incremental copying. Also helps in scanning potential problems in the early stages of the migration process.

Incremental copy
Schedules incremental copies over a period of time for copying the changed data after creating the baseline copy.

Final phase
Creates a final copy of the data and updates the global name space.

VFM also provides a tool for defining and customizing migration policies in a number of different ways (see Figure 1-12 on page 13):

- Runs the migration as a reusable policy and schedules it; the administrator does not have to be present during the migration.
- Prescans for migration problems (ACL, file lock, and others).
- Automatically updates the global namespace; no desktop touch or logon scripts are required to remap drives for users.
- Hides or stops sharing source, ensuring that old data is no longer accessible.
- Runs any pre- or post-migration tasks by calling scripts. For example, VFM sends an e-mail when complete, stops and starts a service, renames files, and zips files before copying them over the network.
- Creates history reports to provide an audit trail of data sources and targets and of any migration errors.

7.2 Policy-based replication

VFM assists in creating and managing detailed replication policies resulting in greater control over the business continuance process. It helps in executing and scheduling hundreds of jobs simultaneously. (Refer to Figure 1-13 on page 14.)

VFM can also be used for remote site data management. It reduces the storage and management of resources at remote sites. VFM also supports one-to-many replication or data scattering (for example, software shares that are accessed by many and many-to-one replication or data gathering, eliminating the use of tape at the remote sites). (See Figure 7-3.)

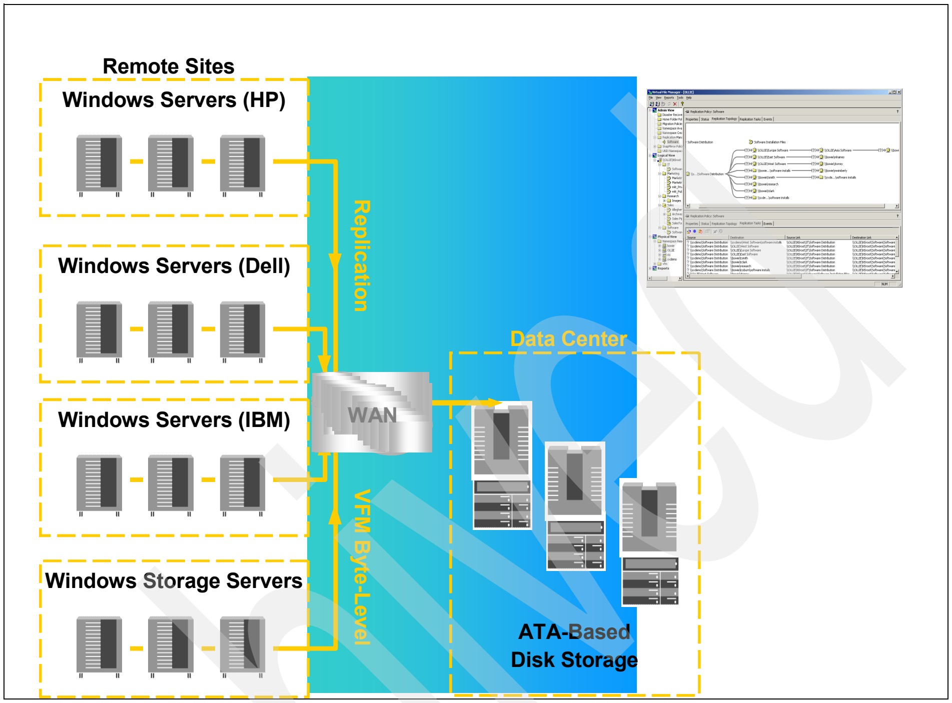

Figure 7-3 Heterogeneous replication at multiple sites

VFM uses byte-level differencing technology to reduce bandwidth consumption across slow links.

7.3 Nondisruptive storage consolidation and expansion

When the storage volume is full, users have the option of moving the entire data to large volumes or of growing the volume itself. Moving the entire data to large volumes is time consuming and requires downtime, and growing the volume in size can force backup and restore problems.

VFM makes it easy to expand and optimize storage without affecting users or applications with global namespace. Global namespace shields users from storage complexities. New storage can be added without affecting how users view or access it, and existing storage can be rebalanced without manual intervention. (See Figure 7-4.)

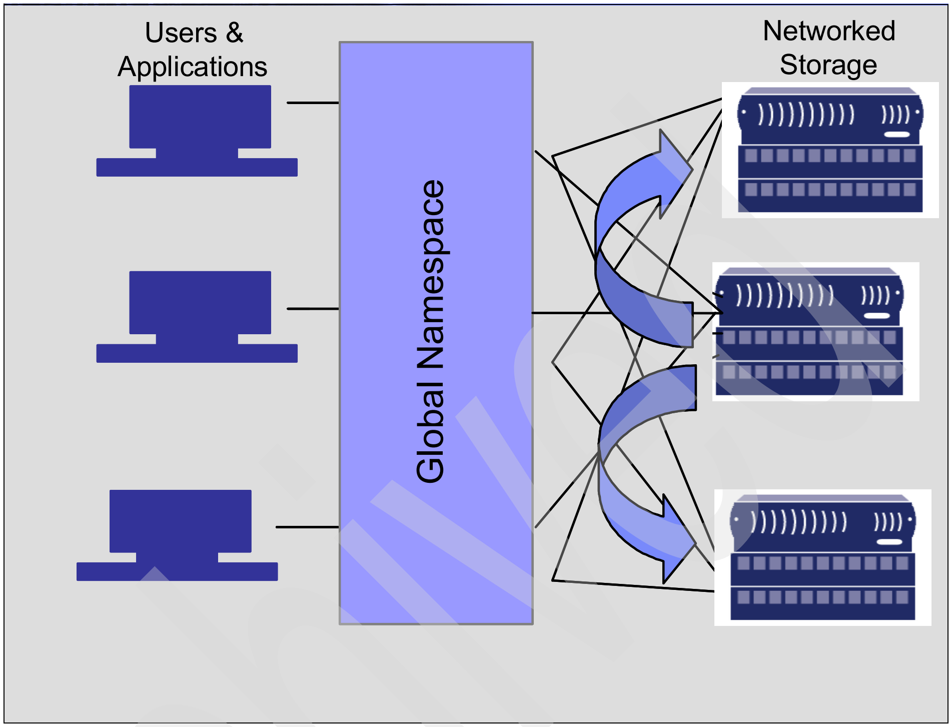

Figure 7-4 Nondisruptive storage expansion

Global namespace also helps in providing nondisruptive storage consolidation. (See Figure 7-5.)

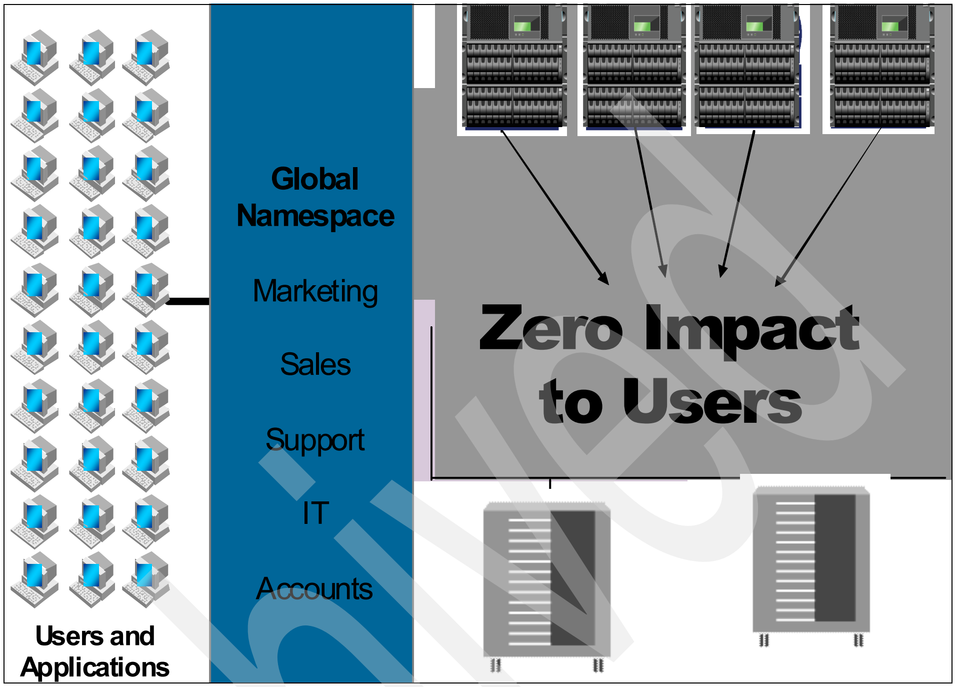

Figure 7-5 Nondisruptive storage consolidation

With VFM consolidation becomes effortless. (See Figure 7-6.)

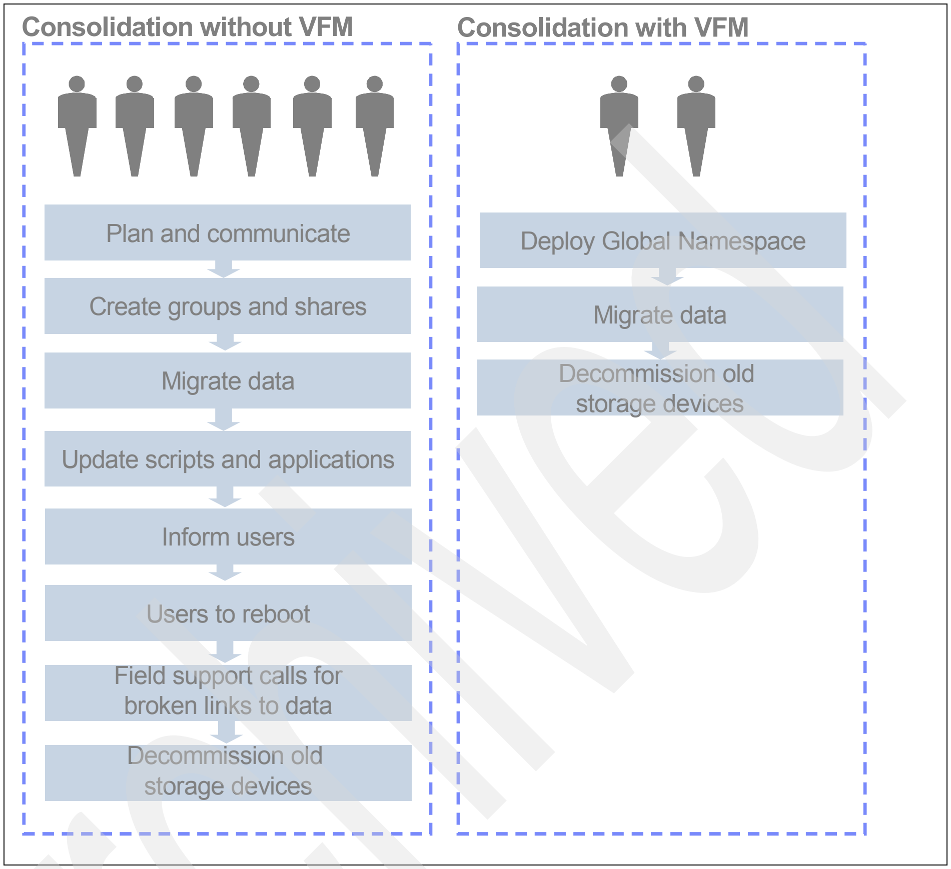
Consolidation without VFM
Plan and communicate
Create groups and shares
Migrate data
Update scripts and applications
Inform users
Users to reboot
Field support calls for broken links to data
Decommission old storage devices
Consolidation with VFM
Deploy Global Namespace
Migrate data
Decommission old storage devices

Figure 7-6   Effortless consolidation

Without VFM, it is difficult to find the broken paths until weeks and months later when you are performing consolidation.

8

Migration

IBM System Storage N series Virtual File Manager (VFM) provides the capability to migrate data between heterogeneous systems. VFM is a general-purpose data movement solution focused on unstructured file data.

A migration policy enables you to set up a migration strategy from one machine to another. You can also set up a migration policy to move data from one volume to another on the same machine. A migration policy can move CIFS data from one share or subfolder of a share to another. Creating a migration policy enables you to move NFS data from one NFS export to another. You can use a Replication Agent to move the data from one place to another. The Replication Agent is automatically deployed to perform the data migration. No administrator intervention is required to deploy a replication in a default configuration.

The goal of any migration is to move data from one location to another without disrupting users. Users must be unaware that data has moved. The same level of access must be maintained after the migration that existed before the migration.

The goal of a migration policy in VFM is to describe the conditions for a migration. You specify the source, the destination, and the controlling parameters of the policy. Once specified, VFM performs the actions of the policy based on the settings. Once complete, you can review the policy to obtain the status.

8.1 Client use cases

This chapter can help VFM users and clients in the following use cases:

- One-to-one migration
  - Move one share from Server1 to N series
  - Move all the shares from Server2 to N series
  - Retire old servers going off lease
  - Move from non-IBM storage to N series
- Many-to-one migration

  Logical consolidation of data from multiple servers onto one server
- One-to-many migration

  Spread data from one server onto different servers
- Many-to-many migration

8.2 Namespace

Migration is not just about moving the data from one location to another. While VFM can be used to move data that is not behind a namespace, the real benefit of VFM is to move data transparently without disrupting the users' view of the data. The result of providing this level of transparency is that the user does not have to know the physical location of the data. For example, a user does not have to know that the marketing data is on Server1. Separating the user from the physical location of the data is achieved by the namespace. Without a namespace, you might have to accompany the migration with changing logon scripts, sending e-mails to users, and so on. With a namespace in place, you can just update the namespace on completion of the migration and seamlessly point the user to the new destination (see Figure 8-1.)

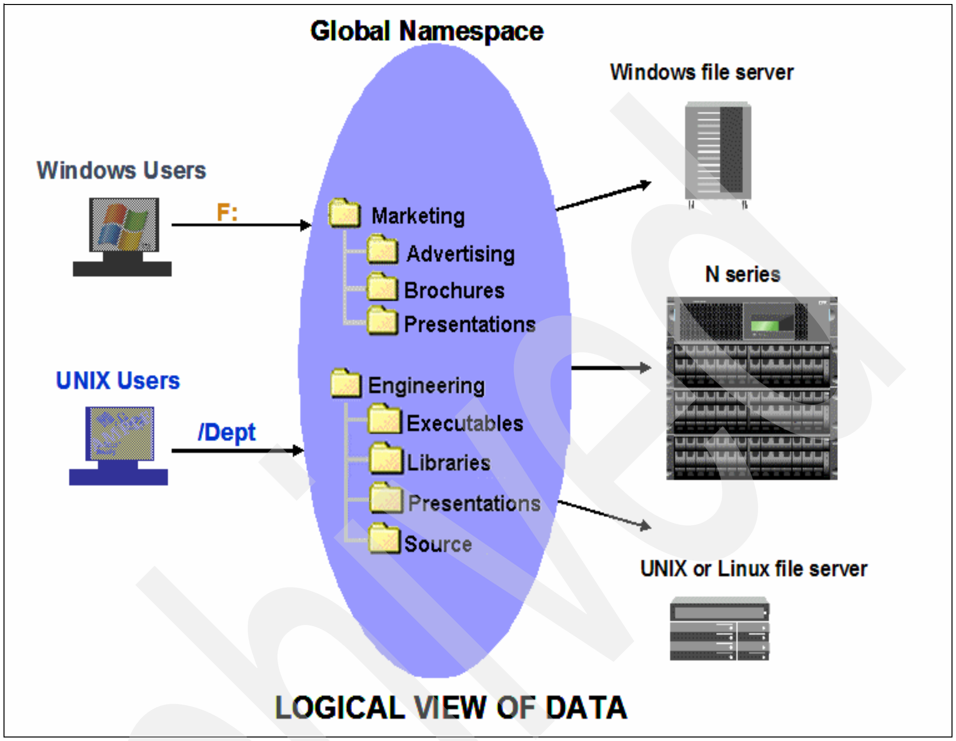

Figure 8-1 Namespace

8.3 Migration phases

A migration policy consists of three phases:

- Baseline or initial phase
- Incremental phase
- Final or cutover phase

Each phase runs on its own schedule, which enables you to schedule a policy at different times. It is possible to enable all three phases in a policy or skip any phase of the policy based on what is best for the client environment. Each phase includes a set of actions that can be performed in that phase. If an action is selected, it is performed during that phase. If an action is not selected, the action is not performed.

The initial phase is used to make a copy of the source data on the destination machine. This is achieved while the source is being used because the incremental phase copies the changes. It is best to schedule the baseline to run over a weekend. Based on the amount of data to be copied, the speed of the machine, the memory on the machine, and the network speed, the migration can take hours or days to complete. It is possible to start the baseline copy days or weeks before the cutover.

Once the baseline is complete, the incremental phase can be used to synchronize the source and destination. The incremental phase can be run nightly or as per the schedule set by you. On average, only a small percentage (usually 3% to 5%) of the source files change every day and are copied in the incremental phase. Any new files created on the source are also copied. Any files deleted on the source are removed from the destination.

You might want to run the final phase over the weekend where the last incremental copy is run before the cutover. During the cutover, the namespace can be updated (optional), and a final synchronization can be performed before the users are brought online. The time required for the final sync is determined by the amount of changed data to be copied and the time required to copy the data. This can take minutes on small data sets or hours on very large data sets.

In some circumstances, you can skip one or more of the phases. For example, if you are sure that no one is accessing the source data, you can configure a migration policy to execute the initial phase, skip the incremental phase, and then move directly to the final phase.

As another example, you might want to skip the initial phase if you use an external source to perform the baseline copy. You might choose to use a tape or DVD copy of the contents of the source to unload onto the destination. Once you make a copy on the destination, the source and destination are nearly identical to one another, and performing a baseline copy using VFM is pointless. In this case, you can unselect the initial phase and proceed with the nightly incremental copies. Sometimes you might or might not want to perform the incremental phase. You might just want to do the final cutover and copy the last set of changes after updating the namespace.

The policy settings offer a great deal of flexibility, providing a powerful way for you to configure the policy to perform automated data movement.

8.4 Policies and tasks

A migration policy specifies the overall criteria for moving data. A policy specifies the schedule, the data movement options, and the behavior of the migration (see Figure 8-2).

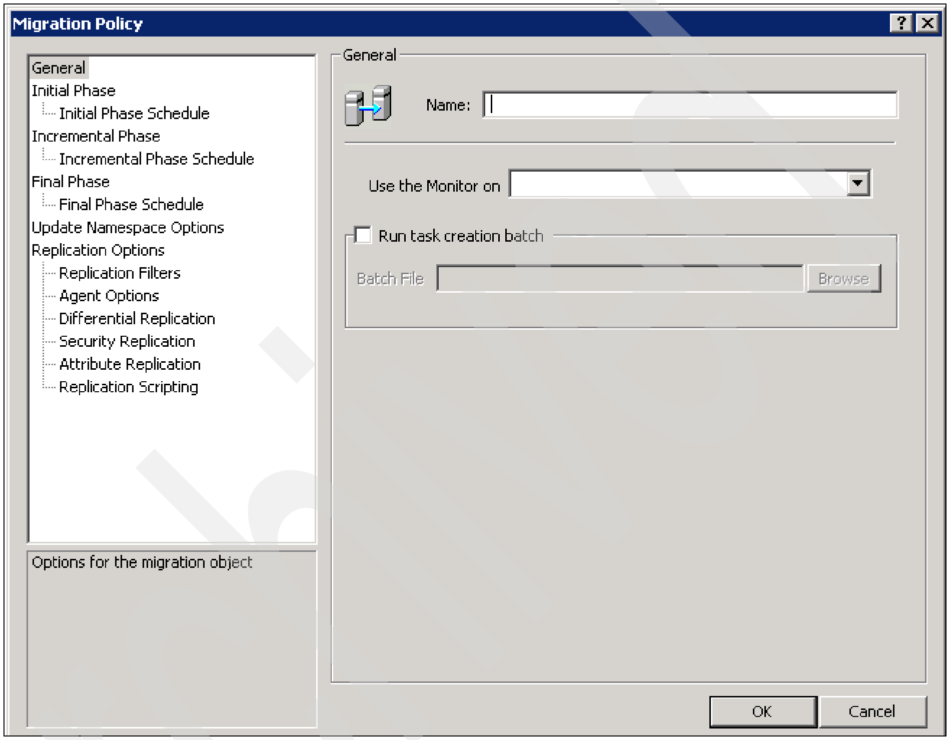

Figure 8-2 Migration Policy window

A policy is made up of tasks. Each task in a policy specifies a source and destination path. A task inherits the settings of the policy. The schedule for the policy specifies when a phase runs and the actions to be performed. The actions are performed on the tasks. Some tasks might be in one phase while other tasks in the same policy are in a different phase. A task can advance either automatically or manually from one phase to another. When a task completes the final phase, it is marked as complete.

8.5 Migration policy example 1

To move data from a Microsoft Windows 2003 source to a N series machine, you must create a migration policy. The migration policy specifies the phases, the settings, and the options. You then add tasks to the policy. You either schedule the job to run at a specific time or manually advance the policy. Figure 8-3 shows the Dept share with two links: Marketing and Sales.

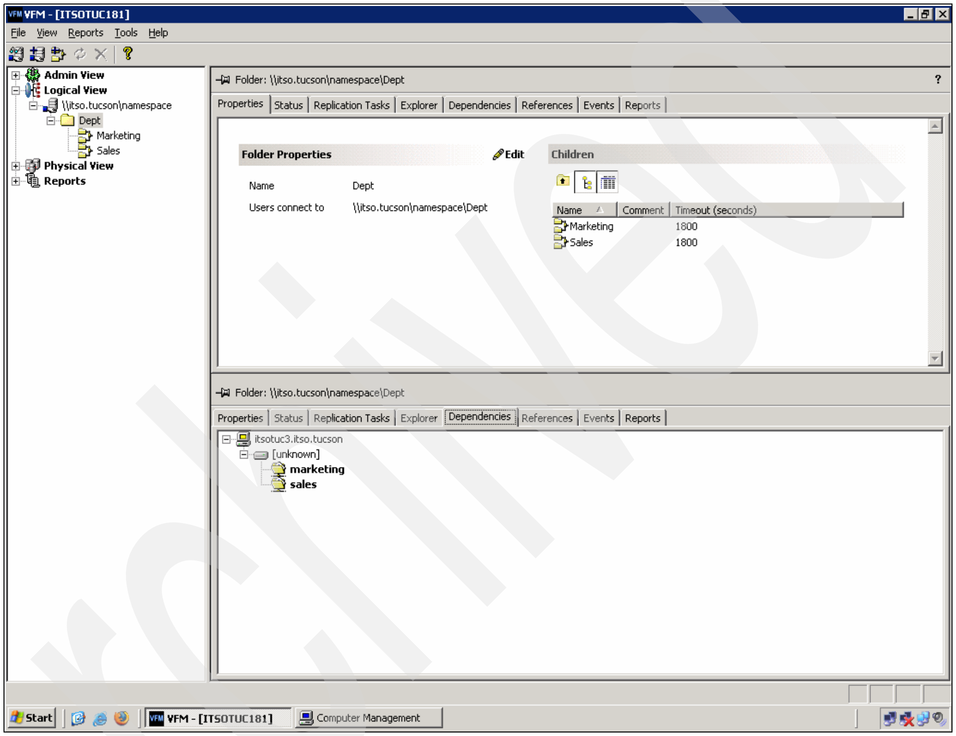

Folder: \\itso.tucson\namespace\Dept

Folder Properties

Name Dept

Users connect to \\itso.tucson\namespace\Dept

Children

Name	Comment	Timeout (seconds)
Marketing		1800
Sales		1800

Folder: \\itso.tucson\namespace\Dept

itsotuc3.itso.tucson
[unknown]
marketing
sales

Figure 8-3 Share with two links

8.5.1 Creating a migration policy

First you must create the migration policy (see Figure 8-4). Use the VFM client to create a migration policy. Specify the settings in the migration policy (see Figure 8-5 on page 206). You manually advance the policy from one phase to the next. You have a DFS link pointing to the source share. This DFS link is automatically updated to point to the destination on completion of the migration.

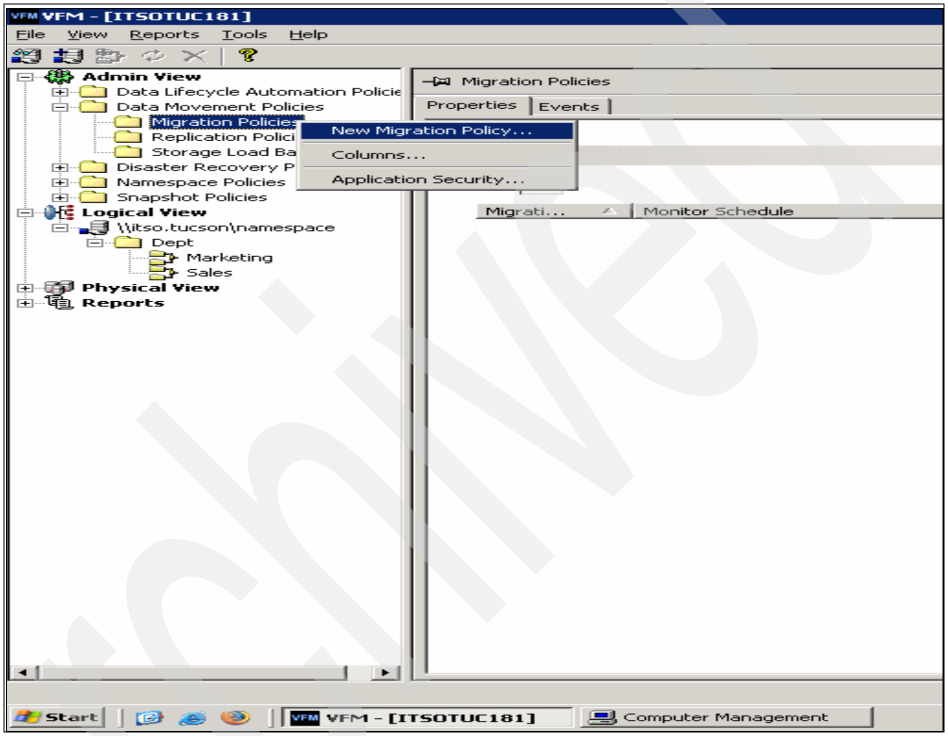

Figure 8-4 New migration policy

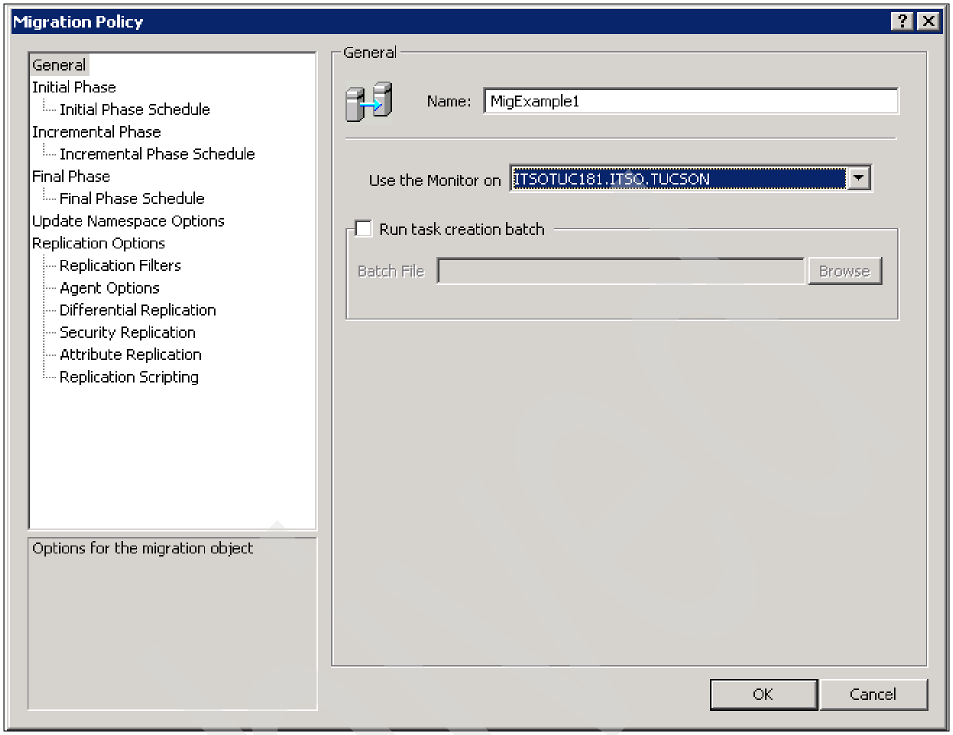

Name: MigExample1

Use the Monitor on ITSOTUC181.ITSO.TUCSON

□ Run task creation batch

Batch File

Figure 8-5 New migration policy: general

8.5.2 Initial phase

A migration policy can specify a set of operations to be carried out in the initial phase (see Figure 8-6 on page 207). You can run a batch file in the initial phase. For example, in the batch file, you might want to stop an application or send an e-mail. The batch file script is passed parameters as described in the Help text. The scripts are run by the VFM server under the context of the VFM service account. The scripts can be located on the local drive or from any network accessible share. A script runs once for each task in the policy. The policy can pause after running the script or continue. If you pause after this step, you have to manually advance the policy.

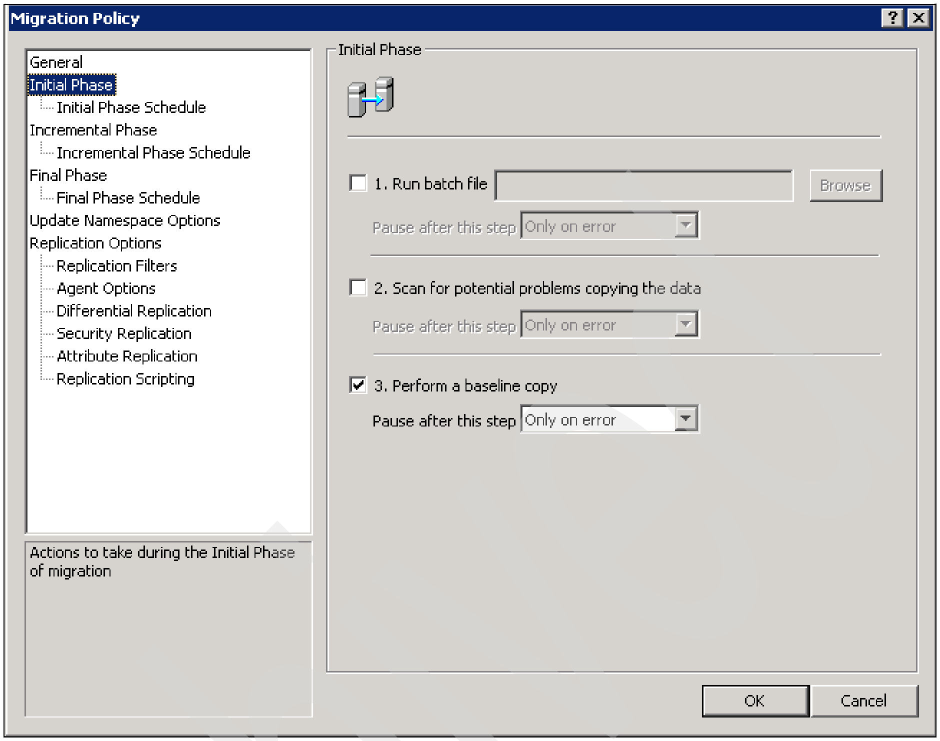

Actions to take during the Initial Phase of migration

Initial Phase

□ 1. Run batch file

Pause after this step: Only on error

□ 2. Scan for potential problems copying the data

Pause after this step: Only on error

☑ 3. Perform a baseline copy

Pause after this step: Only on error

Figure 8-6 Initial phase

VFM can also scan for potential problems. In this case, the server does a cursory examination of the files. It looks for locked files, checks whether it has permissions to the data, checks the destination of a volume type, and other tasks. This step is not an exhaustive one. For example, the disk on the destination might contain sufficient space at this time. However, when it is time to copy the data, sufficient space might be lacking. You might choose to pause at the end of this step. If you choose to pause, you have to manually advance to the next step.

You can also choose to perform a baseline copy of the data. All files on the source are copied to the destination. If any files are already on the destination, the action on the destination is determined by the settings of the policy. For example, if you check the Delete orphaned files on the destination check box, the existing files on the destination are cleared. If you select the Delete the orphaned files option, the Replication Agent makes the destination a mirror image of the source. Once the baseline is complete, the migration policy remains in this phase until it is time to move to the incremental phase.

Sometimes you might have to transfer a large amount of data from a source to a destination. This transfer can be across slow network links. It can be faster to transfer the data from an external means other than by using VFM to do the baseline copy. In this case, you can ship a tape or DVD of the source and unload it at the destination. VFM can be used to perform only the incremental replications. To skip the incremental phase, clear the check box to perform a baseline migration.

In this example, we select to perform a baseline copy (see Figure 8-6 on page 207).

The initial phase schedule specifies when the initial phase starts (see Figure 8-7). The initial phase specifies only a start time. The job runs after the start time passes or the job is manually advanced. If the scheduled time of the initial phase has passed, the job does not start and has to be manually run. The job stays in the initial phase until it is time to move to the incremental phase.

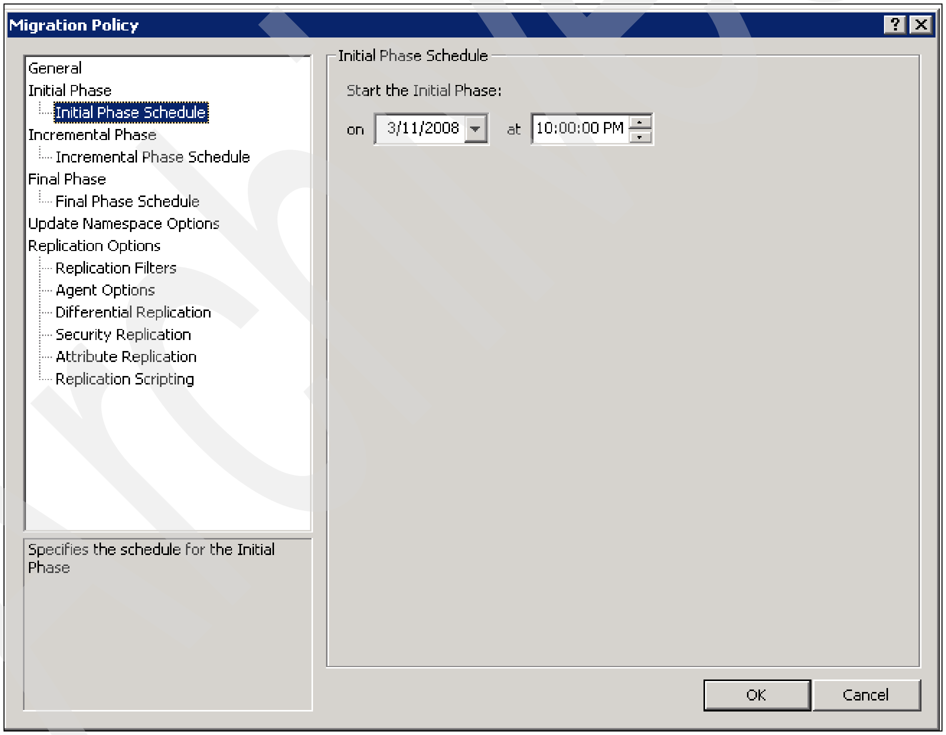

Specifies the schedule for the Initial Phase

Start the Initial Phase:

on 3/11/2008 at 10:00:00 PM

Figure 8-7 Initial Phase Schedule panel

8.5.3 Incremental phase

This phase is used to provide a mechanism to perform repeated copies of the data. As the data changes on the source, the data is copied to the destination at the next invocation of the policy's scheduled run. The incremental phase handles the creation of new files, changed files on the source, and deleted files on the source. The initial phase specifies when the incremental phase finishes and how the job advances to the final phase.

If the check box for Copy the Data is cleared, the incremental phase is disabled.

In this example, we select to perform an incremental copy. You can also select the radio button to never advance the tasks automatically (see Figure 8-8).

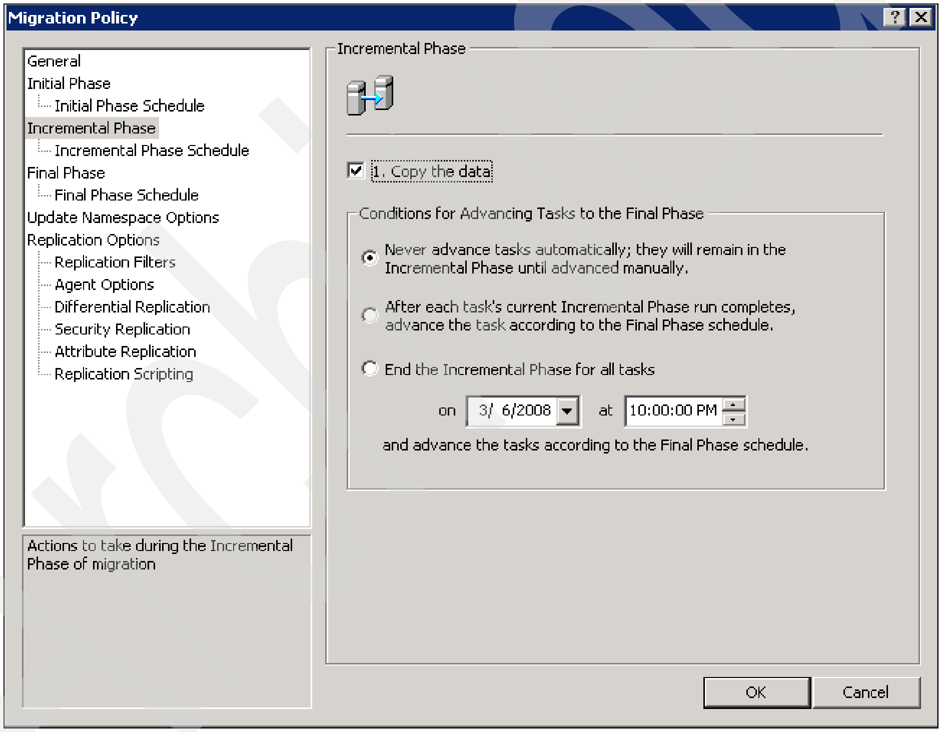

Incremental Phase

☑ 1. Copy the data

Conditions for Advancing Tasks to the Final Phase

● Never advance tasks automatically; they will remain in the Incremental Phase until advanced manually.

○ After each task's current Incremental Phase run completes, advance the task according to the Final Phase schedule.

○ End the Incremental Phase for all tasks

on 3/ 6/2008 at 10:00:00 PM

and advance the tasks according to the Final Phase schedule.

Actions to take during the Incremental Phase of migration

Figure 8-8 Incremental Phase panel

An incremental phase schedule in a migration policy is used to determine the frequency of migration (see Figure 8-9). In some instances, it might be satisfactory to copy data once a night. In other instances, you might want to resynchronize the data every four hours or every eight hours. The frequency of replication for any policy must always be based on the business need and business value of the data.

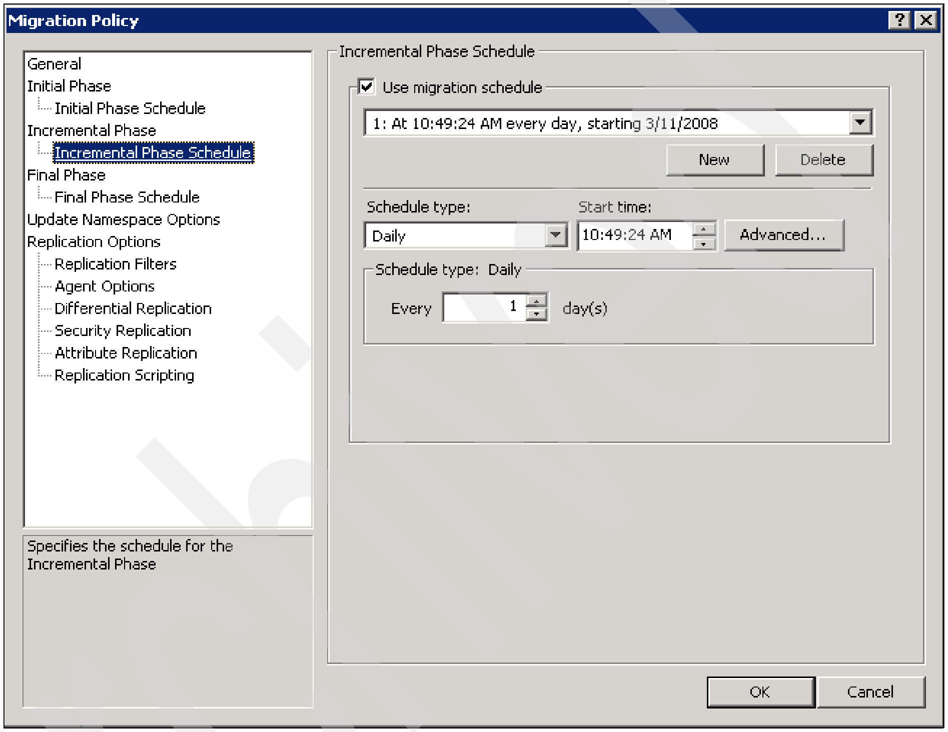

Specifies the schedule for the Incremental Phase

☑ Use migration schedule

1: At 10:49:24 AM every day, starting 3/11/2008

Schedule type: Daily

Start time: 10:49:24 AM

Schedule type: Daily

Every 1 day(s)

Figure 8-9 Incremental Phase Schedule panel

A migration schedule can be set so that data movement occurs only during the night. For example, you might want to move the data only between 8 p.m. and 6 a.m. You might not want to move the data when users are accessing the data during the day.

Do not schedule the incremental migration to occur more frequently than the time required to complete a run of the incremental phase. For example, if an incremental run takes one hour to complete, do not schedule the frequency to run every 30 minutes. If a policy is running and the next scheduled run of the policy elapses, the policy waits until the current run of the policy finishes. If a policy is running, the next run of the policy does not start until the current run finishes.

The VFM console does not have to be open when the policy runs. The VFM server initiates the policy. The policy is run based on the time settings and schedule of the VFM server, which might or might not be in the same time zone of the source or destination machines. Take care when setting the run times of a policy and take into account the source and destination time zones and where the VFM server is located.

Once the incremental phase finishes, the policy does not advance to the final phase until it is time to initiate the final phase.

8.5.4 Final phase

The final phase specifies the actions to be performed when you are almost ready to cut over (see Figure 8-10). This phase assumes that the source and destination are almost synchronized and that little data has to be copied in this phase. When the start time of the final phase passes, the actions selected in the final phase are performed. The selected actions are performed in order. You can choose to pause at the end of each step or continue until the end.

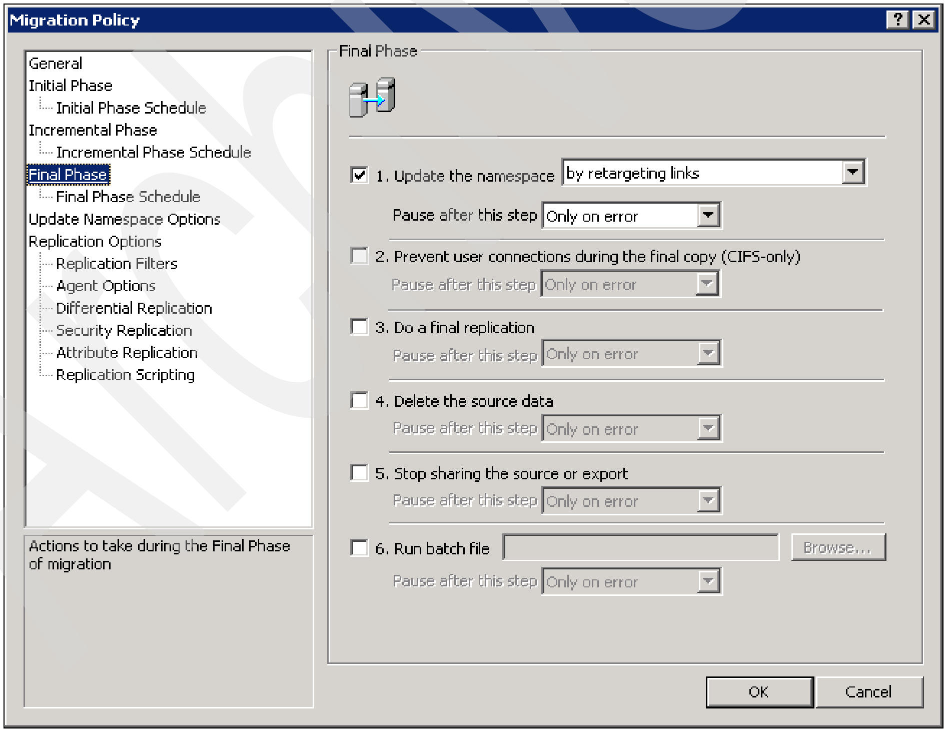
Migration Policy

- General
- Initial Phase
  - Initial Phase Schedule
- Incremental Phase
  - Incremental Phase Schedule
- Final Phase
  - Final Phase Schedule
- Update Namespace Options
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Actions to take during the Final Phase of migration

Final Phase

☑ 1. Update the namespace: by retargeting links
Pause after this step: Only on error

☐ 2. Prevent user connections during the final copy (CIFS-only)
Pause after this step: Only on error

☐ 3. Do a final replication
Pause after this step: Only on error

☐ 4. Delete the source data
Pause after this step: Only on error

☐ 5. Stop sharing the source or export
Pause after this step: Only on error

☐ 6. Run batch file
Pause after this step: Only on error

Figure 8-10 Final Phase panel

Note: Make sure that you run the final phase when no users have open files on the source. If open files are on the source, the files might not be copied correctly, or the user might be forcibly disconnected from the file.

Updating the namespace (that is, checking the Update the namespace option check box) is the most important step of the final phase. Any DFS links targeting the source share or folder are retargeted to point to the destination share or folder. This makes the migration transparent to the user. The user does not have to update drive mappings if they are being accessed through a DFS link. VFM automatically updates the link target, and the update causes the user to be redirected to the destination. Multiple links might be in the namespace that target the share being migrated. All links in the namespace that target the share are updated. Further, these links might be in one namespace or several namespaces. You can choose to pause after this step.

The “Prevent user connections during the final copy check box enables VFM to copy any open files. If a user has an open file, VFM changes the share name. This action is disruptive to the user and forcibly disconnects the user from the source share. Changing the share name enables VFM to copy the files without having any open files to worry about. VFM hides the source share by adding a few random characters to the share name and appending a dollar sign ($) to hide it. VFM then uses the newly created share as the source and moves the remaining data.

Note: Do not select the Prevent user connections during final copy check box if you are migrating a source that is a subfolder of a share. In that case, you might have other subfolders that are being accessed by the user through the share.

Checking the Do a final replication check box enables you to copy the last remaining files and synchronize the source and destination. The time required to perform the final replication is based on the amount of data to synchronize, the number of files already in sync, the number of files to copy, the network speeds available, and other considerations.

Checking the Delete the source data check box enables you to free up the space on the source server. Be careful about selecting this option because it is a destructive operation. Instead you might want to manually delete the data a few weeks after the cutover when you know the cutover was successful.

Checking the Stop sharing the source or export check box deletes the source share and no longer advertises it in the network. Users can no longer access the source share. This option is useful in a situation where users are bypassing the namespace to access the data using the physical paths. If you discontinue use of the physical share path, you are able to detect users that bypass the namespace. You can then encourage the users and applications to access the shares through the namespace.

Checking the Run batch file check box at the end of migration enables you to send e-mail to a set of users at the end of the migration. You can also use this option to restart an application from the new location (see Figure 8-11).

In our example, we select the following options:

- Update the namespace and retarget the links
- Prevent user connections during the final copy
- Do a final replication

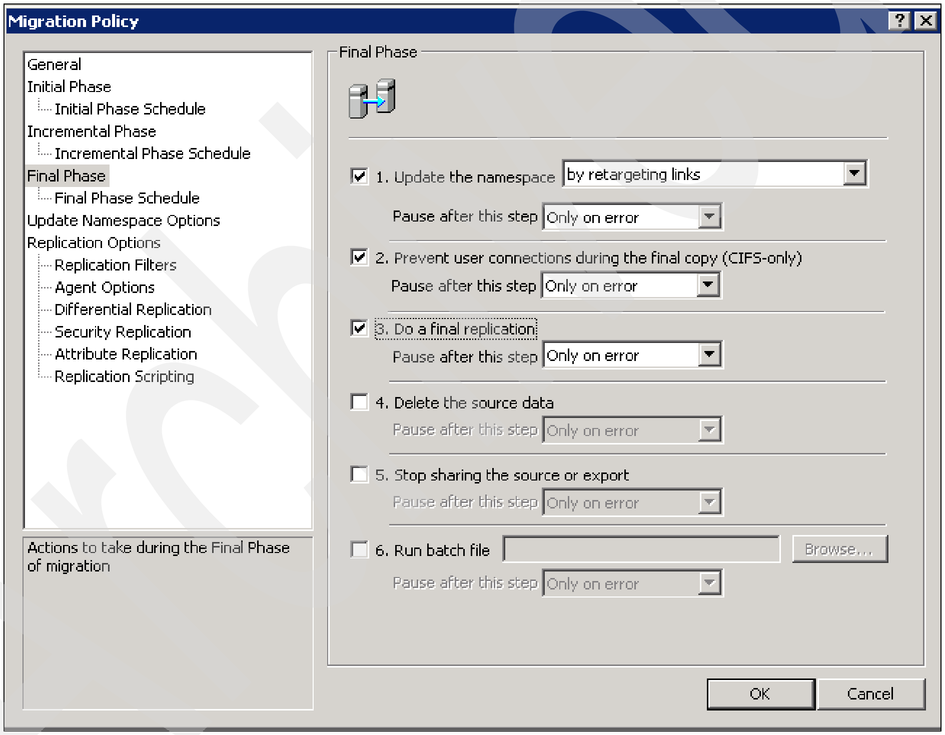

Actions to take during the Final Phase of migration

Final Phase

☑ 1. Update the namespace by retargeting links
Pause after this step Only on error

☑ 2. Prevent user connections during the final copy (CIFS-only)
Pause after this step Only on error

☑ 3. Do a final replication
Pause after this step Only on error

☐ 4. Delete the source data
Pause after this step Only on error

☐ 5. Stop sharing the source or export
Pause after this step Only on error

☐ 6. Run batch file
Pause after this step Only on error

Figure 8-11 Example 1: final phase parameters

The final phase schedule in a migration policy determines when to initiate the cutover (see Figure 8-12). It is best to schedule the final phase when no users are accessing the data. It is advisable to tell users to disconnect from the source share when the final cutover is scheduled. The VFM console does not have to be open when the final phase of the policy runs. The VFM server initiates the policy. The policy is run based on the time settings and schedule of the VFM server, which might or might not be in the same time zone of the source or destination machines. Take care when setting the run times of a policy and take into account the source and destination time zones where the VFM server is located.

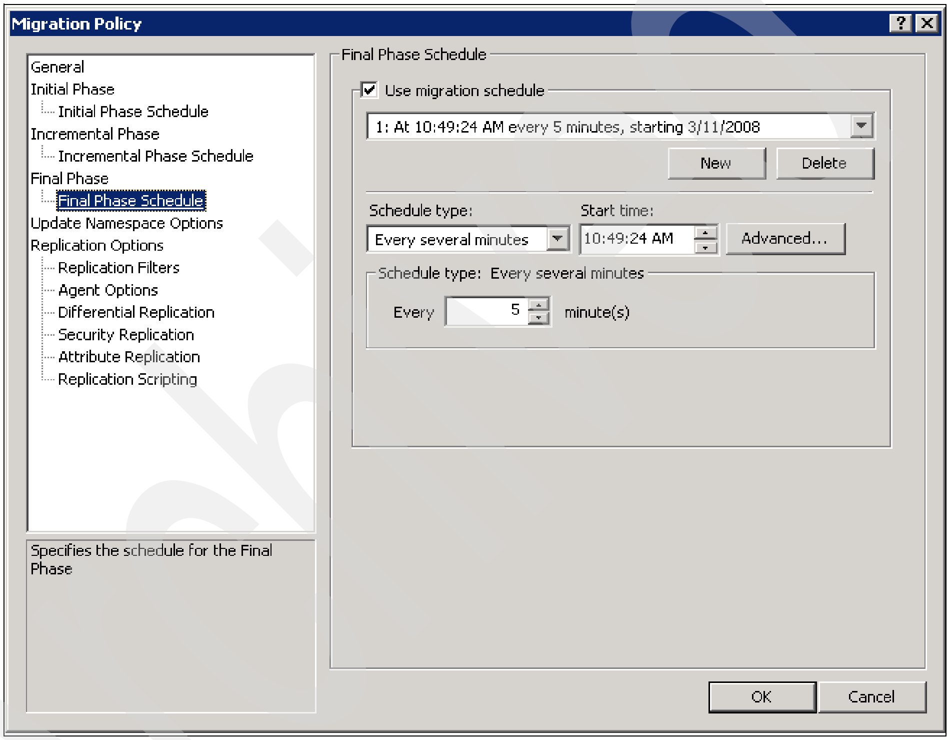

Final Phase Schedule

☑ Use migration schedule

1: At 10:49:24 AM every 5 minutes, starting 3/11/2008

Schedule type: Every several minutes

Start time: 10:49:24 AM

Schedule type: Every several minutes

Every 5 minute(s)

Specifies the schedule for the Final Phase

Figure 8-12 Final Phase Schedule panel

Once the final phase time elapses, the tasks move through each selected action of the final phase.

8.6 Namespace options

When VFM performs a data movement operation, the goal is to perform these operations so they are transparent to the user. The maximum amount of transparency is achieved when Microsoft Windows users access the data through the namespace. The Update Namespace Options specify the Windows namespaces (DFS root) to search to update DFS links. Each DFS root listed in the policy is used to determine if any DFS links reference the source share. These DFS links are automatically updated if the Update the namespace check box is selected on the Final Phase panel (see Figure 8-10 on page 211). If VFM does not know about a DFS root or namespace, the DFS links in those namespaces are not updated. Be careful that all the DFS roots listed are the ones you want to update and that you have the rights to update them (see Figure 8-13).

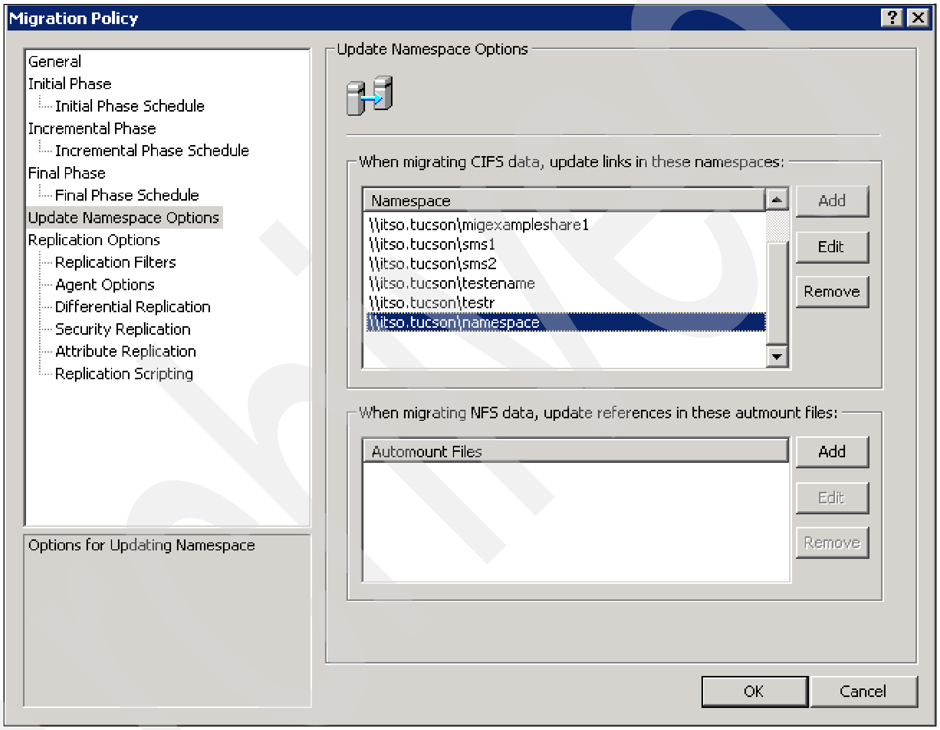

Migration Policy

- General
- Initial Phase
  - Initial Phase Schedule
- Incremental Phase
  - Incremental Phase Schedule
- Final Phase
  - Final Phase Schedule
- Update Namespace Options
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options for Updating Namespace

Update Namespace Options

When migrating CIFS data, update links in these namespaces:

Namespace
\\itso.tucson\migexampleshare1
\\itso.tucson\sms1
\\itso.tucson\sms2
\\itso.tucson\testename
\\itso.tucson\testr
\\itso.tucson\namespace

When migrating NFS data, update references in these automount files:

Automount Files

Figure 8-13 Update Namespace Options panel

If you are migrating NFS data, make sure to specify the automount files referencing the source paths. VFM updates these files to replace references to the old export with the new export. For example, suppose 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 auto.auto 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 refer the destination:

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

Make sure the automount file auto.auto is present under the Namespace list on the Update Namespace Options panel shown in Figure 8-13 on page 215.

8.7 Replication options

Replication options enable you to control data movement (see Figure 8-14). These options are similar to those provided by a replication policy (described in detail in Chapter 10, “Replication” on page 315).

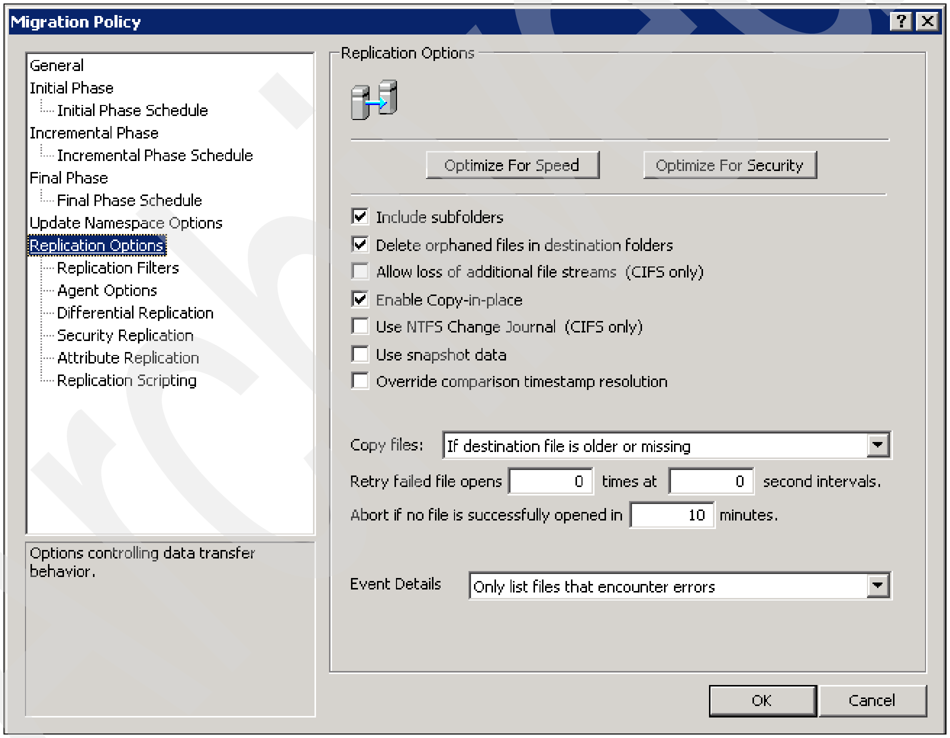

☑ Include subfolders
☑ Delete orphaned files in destination folders
□ Allow loss of additional file streams (CIFS only)
☑ Enable Copy-in-place
□ Use NTFS Change Journal (CIFS only)
□ Use snapshot data
□ Override comparison timestamp resolution

Copy files: If destination file is older or missing

Retry failed file opens 0 times at 0 second intervals.

Abort if no file is successfully opened in 10 minutes.

Event Details: Only list files that encounter errors

Figure 8-14 Replication Options panel

8.7.1 Include subfolders

A migration policy can be set to copy all files and folders or files at the top level. For example, you might want to get an idea about how long a migration might take by copying only the top-level files and folders. Or you might want to copy the top-level folders and then set security on the folders. Once you set the security, you might want to follow it with another migration that just copies the data and inherits the permissions set on the source. This is especially useful when you are doing a NetWare to N series migration.

8.7.2 Delete orphaned files in destination folders

The Delete orphaned files in destination folders option makes the source an identical copy of the destination. It deletes extra files on the destination. Do not select this option if you have multiple sources going to the same destination.

8.7.3 Enable Copy-in-place

Many optimizations in a data movement policy are a compromise between speed and safety. The Enable Copy-in-place option on the Replication Options panel (see Figure 8-14 on page 216) can help finish a migration policy faster with a slight security risk. Copy-in-place comes into effect during an incremental copy or a final phase copy. Assume the following scenario for a migration policy that moves files from one machine to another:

- A file called xyz.doc is on the source.
- The xyz.doc file is copied to the destination.
- The file is modified on the source.
- Because the file on the source is newer than the one on the destination, xyz.doc on the source must replace the file on the destination.

If the Enable Copy-in-place check box is not checked, the following actions occur:

- The file xyz.doc is copied to the destination as a temporary file.
- When the copy is complete, the file xyz.doc is deleted on the destination and the newly copied temporary file is renamed as xyz.doc (this is called a *safe rename technique*).

If Enable Copy-in-place is checked, the following actions occur:

- The file xyz.doc is truncated to 0 bytes on the destination.
- The file from the source is copied to the truncated file.

Checking Enable copy-in-place is faster because no rename operations are required. Copy-in-place is preferred when the source and destination are on the same LAN or connected by a high-speed network link. If the source and destination are connected by a low-speed link or the link is not reliable, you must clear the Enable Copy-in-place check box and allow the safe rename technique to work. This way, if the network link breaks when the copy is in progress, you are not left with a zero-length file on the destination.

8.7.4 Use NTFS change journal

Suppose you have a large source volume, and this source volume must be kept in sync with a destination volume during the incremental phase of a migration. If the source volume has millions of files, the time required to scan the source volume to determine the changed files can be significant. VFM can optimize the search of changed source files on Microsoft Windows machines by integrating with the NTFS change journal. The NTFS change journal (only available on Microsoft Windows) is a "file" that maintains a list of accessed, changed, created, and deleted files. If the Replication Agent scans the source NTFS change journal, it can detect the changed, deleted, and newly created files in a fraction of the time compared to doing a full file scan of TBs of source data. This way, an incremental migration resync proceeds much faster.

> **Note:** The change journal file can wrap, resulting in lost changes. The maximum change journal size on the volume is based on space you allocate. If the change journal wraps, VFM resorts to a full scan of the source to determine the changed files.

8.7.5 Use snapshot data

During any data movement policy run, one or more files can be open on the source. This is usually the case for PST files and Access databases where the user might have an exclusive lock on the files. If VFM can get a read lock on the files, it copies the files. Depending on how the policy is set up, VFM might either skip the opened files or copy the files as best as it can from the snapshot.

VFM can integrate with Volume Shadow Service (VSS) on Microsoft Windows 2003 or VFM snapshots. VFM can take a temporary snapshot of the source volume and copy all the files from the source snapshot. Once the files are copied, the source snapshot is automatically deleted. If the source is not Windows 2003 or N series, VFM can be configured to retry the failed copies. Files not copied are logged as errors and reported in the manifest based on the settings of the policy.

8.7.6 Allow loss of additional file streams

The Allow loss of additional file streams option (see Figure 8-14 on page 216) enables fine-grained control of the migration between heterogeneous machines. Some vendor machines support files streams while others do not. Some

Microsoft Windows applications store data in alternate data streams. In most cases, you might or might not know whether alternate data streams are used in the data files. You sometimes might or might not know whether the vendor supports file streams. If VFM cannot copy a file because the destination does not support streams, the only way to copy the file that has streams is to select the Allow loss of additional file streams check box. Otherwise, VFM cannot move the file with streams and reports it as an error.

8.7.7 Event details

The settings for the Event Details drop-down list box (shown in Figure 8-14 on page 216) has a significant effect on migration performance. The Event Details options enable a migration policy to provide a log of each file moved. The optimal setting is to list only files that encounter errors. If you choose to log each file moved, the amount of space required for the manifest is exceptionally large. Also, this manifest is stored with the Replication Agent. To view the manifest, it must be transferred to the server and then the client, which can take an inordinately long time. For test and demonstration migrations, it is acceptable to list all files so you can get comfortable using the product. For all production migration, we strongly recommend you list only files that have errors.

For the purposes our example (see Figure 8-14 on page 216), select the following options:

- Include subfolders
- Delete orphaned files in destination folder
- Enable copy in place
- Event details of only listing files with errors.

8.7.8 Replication filters

This option provides a more flexible way of controlling the files to be copied or excluded in a migration (see Figure 8-15) You might sometimes choose not to copy temporary files. Or you might want to move files, but the new policy prohibits MP3 files from being placed on the destination N series. Or perhaps you have a large share and want to exclude two top-level folders from being copied. Instead of specifying each folder to copy, you can choose to specify the folders to exclude.

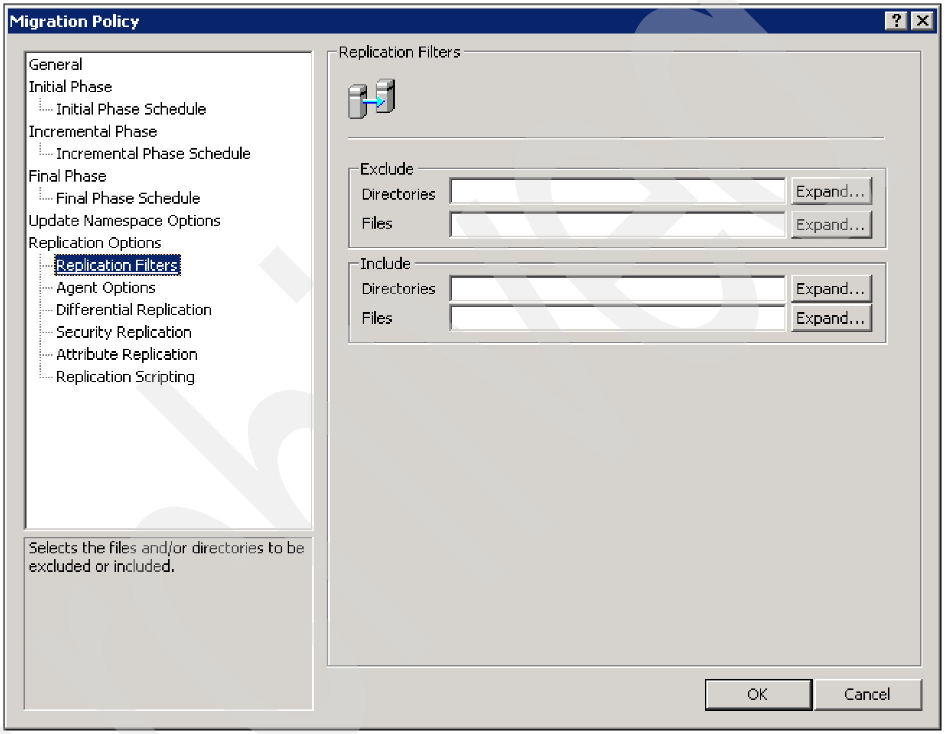

Selects the files and/or directories to be excluded or included.

Figure 8-15 Replication Filters panel

8.7.9 Agent options

By default, VFM uses the Replication Agent on the destination to pull the source of a migration policy from data (see Figure 8-16 on page 221). This method is preferred because the agent on the destination knows how much data it can handle based on the load on the destination machine. If the Replication Agent cannot run on the destination (for example, the destination is an N series machine), you can run the Replication Agent on the source and push data to the destination.

Other configurations are also possible when the agent is on a proxy machine and performing a three-way copy. Generally, three-way copies are slower than a direct copy when the Replication Agent is on the source or when the Replication Agent is on the destination. You need not manually install the Replication Agent. Based on the migration policy settings, the Replication Agent is automatically deployed to the specified machine. A Replication Agent can handle a number of data movement tasks at the same time although it is advisable to schedule the movements between a set of agents. An agent is multithreaded and can take advantage of multiprocessor systems.

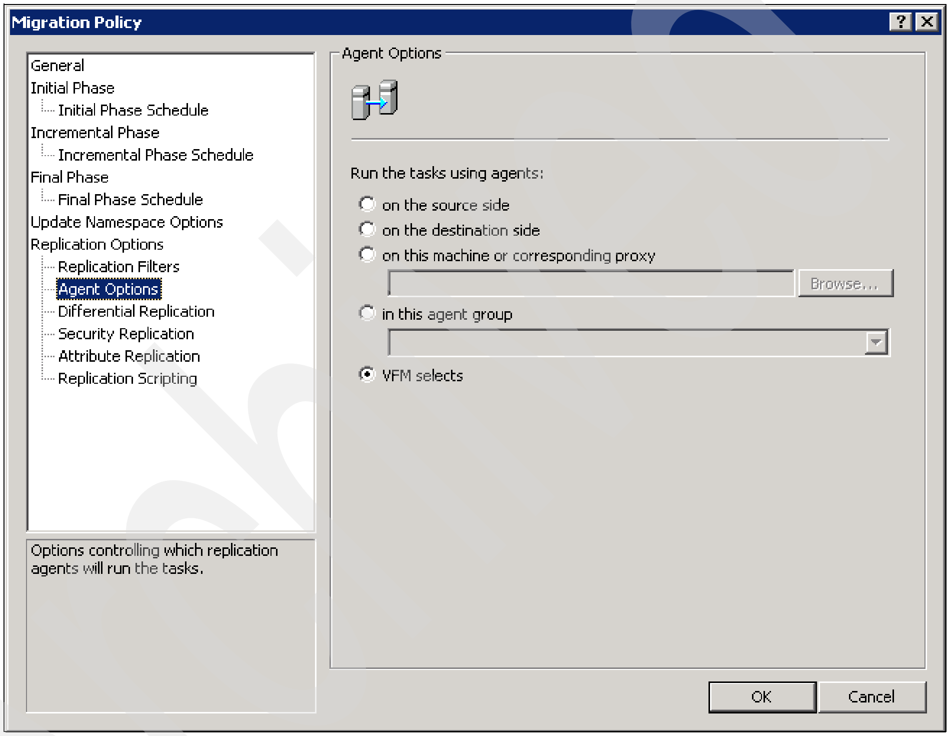

Options controlling which replication agents will run the tasks.

Agent Options

Run the tasks using agents:

○ on the source side
○ on the destination side
○ on this machine or corresponding proxy
○ in this agent group
● VFM selects

Figure 8-16   Agent Options panel

The Replication Agent runs under the credentials of the service account that VFM runs under, which is specified when the product is installed. The service account must have administrator rights on the source and destination machines. If the Replication Agent does not have security permissions on a folder or file, it is not able to copy the data. At times, it might be necessary to run the agent under a different account for a specific data movement policy. In that case, use the Agent Management tool to set the credentials.

A Replication Agent uninstalls itself after 30 days of inactivity. You can use the Agent Management tool to uninstall an agent.

The Replication Agent communicates with the VFM server and periodically passes status messages to it. This heartbeat happens every few minutes. Make sure that the network ports (6001, 6002, and 6005) are open for communication across all firewalls in the path between the VFM server and the Replication Agent; otherwise, the Replication Agent cannot communicate with the server. (Refer to Chapter 3, “Preparation” on page 83 for details.) If, for any reason, the Replication Agent cannot talk to the VFM server, the Replication Agent aborts the job, and the migration fails.

Agent grouping provides a way to share the load between multiple Replication Agents that are involved in a data movement policy. Assume that you are moving many TBs of data from one N series storage system to another. Instead of using one Replication Agent to perform all the migrations, it might be advisable to use a group of Replication Agents to move the data. You set up an agent group in the Agent Management interface. You specify which machines are members of the agent group.

When you configure a migration policy, you select the agent group to perform the migration. When the migration policy runs, VFM sends the task to the next available agent in the agent group. A round robin technique is used to determine which Replication Agent in the agent group performs the actions of task. Using agent groups prevents a Replication Agent from being overloaded while sharing the load across multiple agents to perform data migrations. If no agent of the group is available, the migration does not run.

For the purposes of our example, select the **VFM selects** radio button (see Figure 8-16 on page 221). Doing so enables VFM to select the Replication Agent to be used based on internal algorithms.

8.7.10 Differential replication or byte level replication (BLR)

VFM provides a feature using byte level replication to control data movement across slow links (see Figure 8-17 on page 223). BLR minimizes the data sent across the wire especially with slow network links. BLR does so by only sending the changed bytes. You might have a 10 MB Microsoft Word document file that is copied from the source to the destination share.

Suppose a user at the source changes one or two pages in the document. Instead of sending the10 MB changed file again from the source to the destination, VFM can be used to intelligently send the deltas across the wire. While the time required in computing the changed bytes in extremely large files can be significant, the network usage is minimized. You are sacrificing the extra CPU cycles to figure out the changed bytes so that the limited resource of network bandwidth is optimized. BLR requires two Replication Agents to work together to achieve a successful replication.

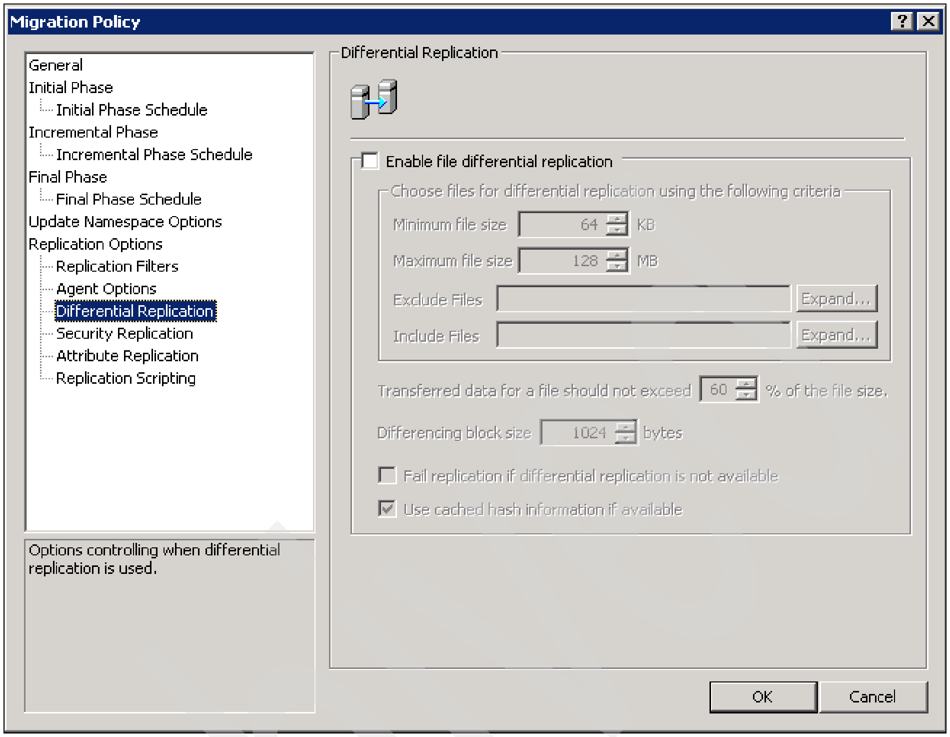

Differential Replication

☐ Enable file differential replication

Choose files for differential replication using the following criteria

Minimum file size 64 KB

Maximum file size 128 MB

Exclude Files

Include Files

Transferred data for a file should not exceed 60 % of the file size.

Differencing block size 1024 bytes

☐ Fail replication if differential replication is not available

☑ Use cached hash information if available

Options controlling when differential replication is used.

Figure 8-17 Differential Replication panel

BLR must only be used across slow links. In LAN speed networks, the overhead of BLR computation is more than the benefit it provides. BLR works when the destination Replication Agent computes rolling hash sets across all the blocks of the file at the destination. The destination Replication Agent sends the rolling hash set as a file signature to the source Replication Agent. This results in a small amount of data transferred from destination to source. The source Replication Agent computes the rolling hash set of the changed file at the source. When different hash blocks are encountered, the source Replication Agent computes an instruction set to send from the source to the destination. The instruction set is used to insert or remove data blocks from the file at the destination to make it look like the source. BLR over slow WAN links results in 80% savings in network traffic of slightly changed data files.

For the purposes of our example, we did not enable differential replication (see Figure 8-17).

8.7.11 Security replication

Every data migration requires that the files be copied accurately and completely. VFM provides several mechanisms to control the copying of files. Where possible, VFM signal errors when files were not copied completely or correctly (see Figure 8-18).

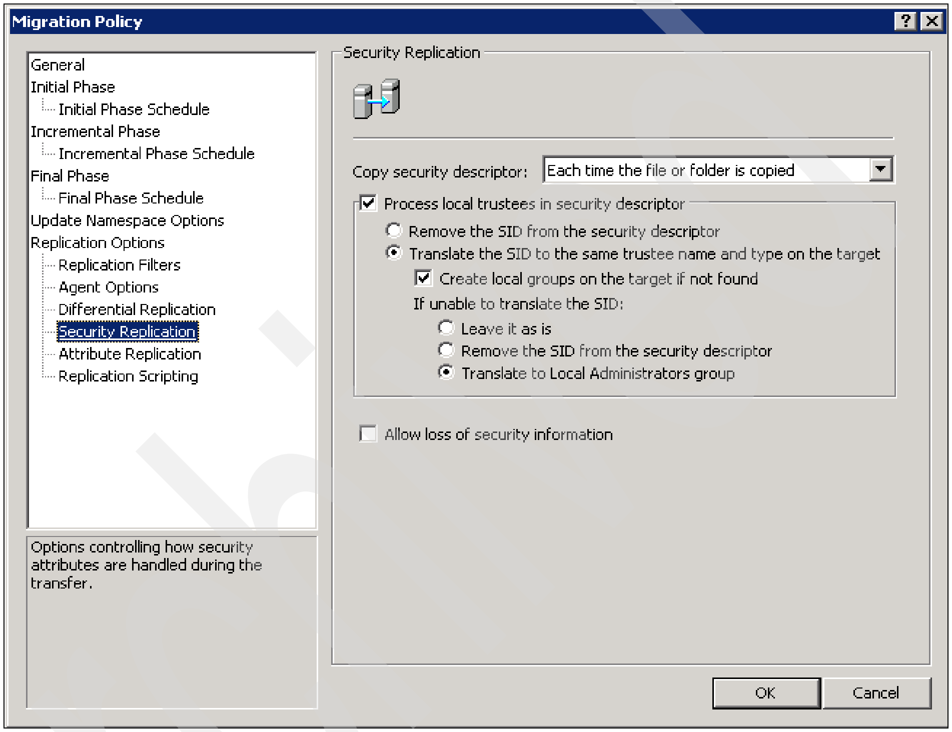

Options controlling how security attributes are handled during the transfer.

Security Replication

Copy security descriptor: Each time the file or folder is copied

☑ Process local trustees in security descriptor

○ Remove the SID from the security descriptor

● Translate the SID to the same trustee name and type on the target

☑ Create local groups on the target if not found

If unable to translate the SID:

○ Leave it as is

○ Remove the SID from the security descriptor

● Translate to Local Administrators group

□ Allow loss of security information

Figure 8-18 Security Replication panel

Many Active Directory deployments can use local groups to secure the access to file resources on the old server. In Microsoft Windows NT® 4 domains, Microsoft recommends the use of local groups. A local group is valid only on the source system on which it is defined. If the security on a file is controlled by the permissions of a local group and that file is copied to a destination machine, the security is lost. This means that the users who had access to the files on the old server might not have access to the migrated copy.

VFM provides a way to handle this situation as described on the Security Replication tab. VFM can create a local group on the destination with the same name as the local group on the source. All domain users and domain groups that are members of the old local group on the source server are added to the newly created local group of the destination server. This way, users that had access to the data to the old local group on the source server have access to the data on the new server through the newly created local group.

At times, you simply care about replicating the data from one location to another without regard to the security of the data on the new system. For example, if you replicate data from a Microsoft Windows server to a UNIX-style qtree on the new N series machine, it is advisable to check the Allow loss of security information check box (see Figure 8-18 on page 224).

For the purposes of our example, select the default check boxes and radio buttons shown in Figure 8-18 on page 224.

8.7.12 Attribute replication

You can control the attributes in a migration policy to compare two files to determine whether they are identical or different. By default, VFM compares the names, size, dates, and attributes. You can override this option by checking or clearing the check box for attribute replication. You can also specify how the attribute is set on the destination file after it has been copied (see Figure 8-19 on page 226). For example, you might want to clear the archive attribute of a copied file so that the file is backed up on the replication target.

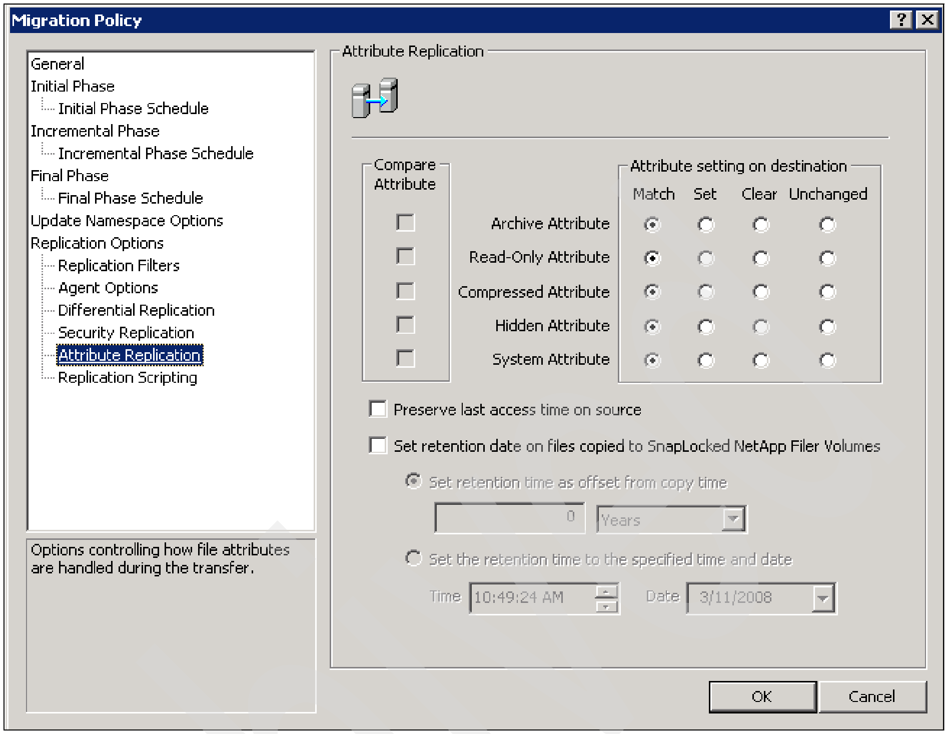

Migration Policy

- General
- Initial Phase
  - Initial Phase Schedule
- Incremental Phase
  - Incremental Phase Schedule
- Final Phase
  - Final Phase Schedule
- Update Namespace Options
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options controlling how file attributes are handled during the transfer.

Attribute Replication

Compare Attribute		Match	Set	Clear	Unchanged
☐	Archive Attribute	●	○	○	○
☐	Read-Only Attribute	●	○	○	○
☐	Compressed Attribute	●	○	○	○
☐	Hidden Attribute	●	○	○	○
☐	System Attribute	●	○	○	○

☐ Preserve last access time on source

☐ Set retention date on files copied to SnapLocked NetApp Filer Volumes

● Set retention time as offset from copy time

0 Years

○ Set the retention time to the specified time and date

Time 10:49:24 AM Date 3/11/2008

Figure 8-19 Attribute Replication panel

When moving files, you might want to preserve the last access time of the file on the source. Even though VFM copies the source file, it resets the last accessed time on the source file. This is useful when Hierarchical Storage Management (HSM) or other reporting software is based on the last access time. Overhead on the Replication Agent is involved in preserving the last access time on the source. So if this feature is not required, uncheck the Preserve last access time on source check box.

Once you have created a migration policy, you are ready to add tasks to the policy. (Right-click the newly created policy.)

8.7.13 Adding a task to migration policy

In this section, we describe adding a task to the migration policy. Creating a migration policy involves specifying the migration tasks that pertain to the policy. You specify the source server and share of the migration task (see Figure 8-20).

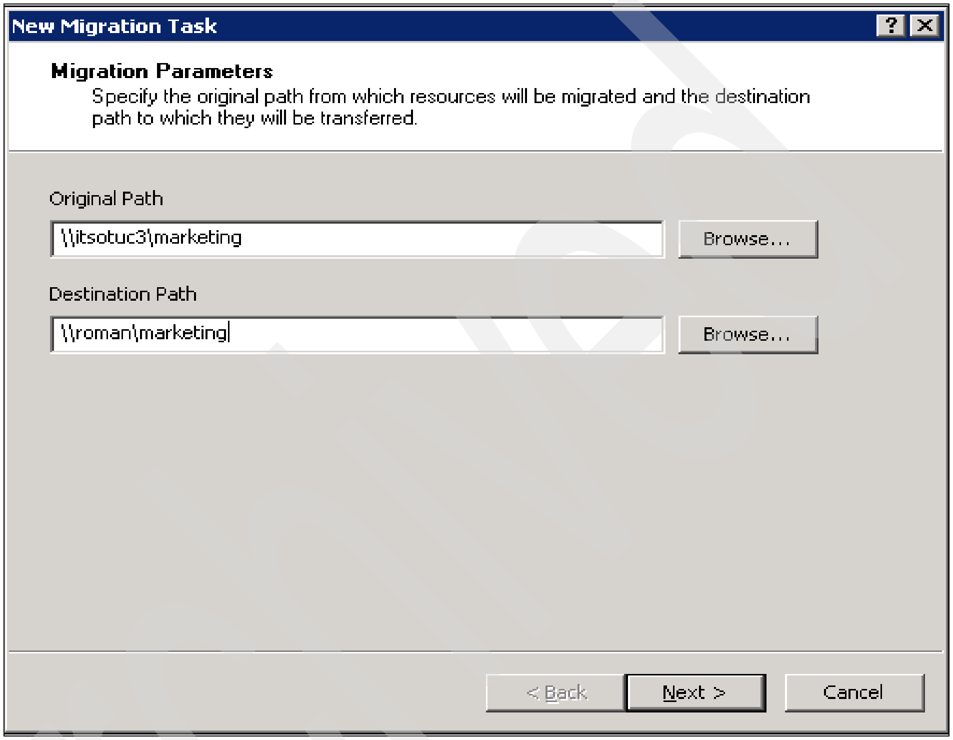

Migration Parameters

Specify the original path from which resources will be migrated and the destination path to which they will be transferred.

Original Path: \\itsotuc3\marketing

Destination Path: \\roman\marketing

Figure 8-20 Specifying source and destination

You specify the destination server and share of the N series machine that is the destination of the migration task. You might use an existing share as the destination. If the destination share does not exist, VFM creates it for you in the specified volume (see Figure 8-21 on page 228, Figure 8-22 on page 229, Figure 8-23 on page 230, and Figure 8-24 on page 231). Be sure to specify the correct share-level permissions.

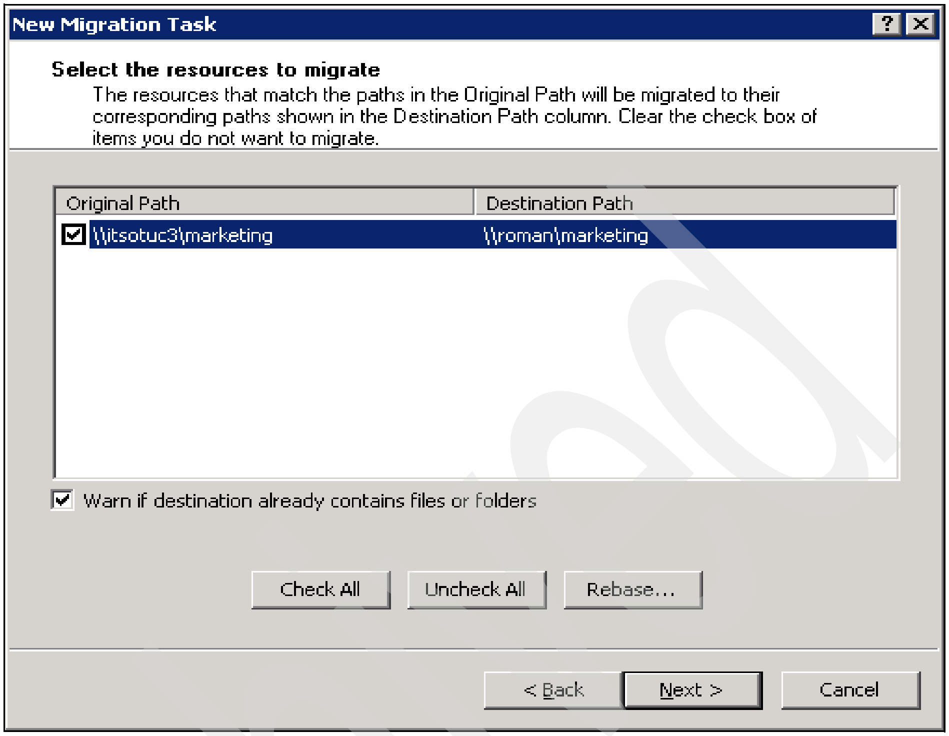

New Migration Task

Select the resources to migrate

The resources that match the paths in the Original Path will be migrated to their corresponding paths shown in the Destination Path column. Clear the check box of items you do not want to migrate.

Original Path	Destination Path
☑ \\itsotuc3\marketing	\\roman\marketing

☑ Warn if destination already contains files or folders

Figure 8-21 Selecting resource to migrate

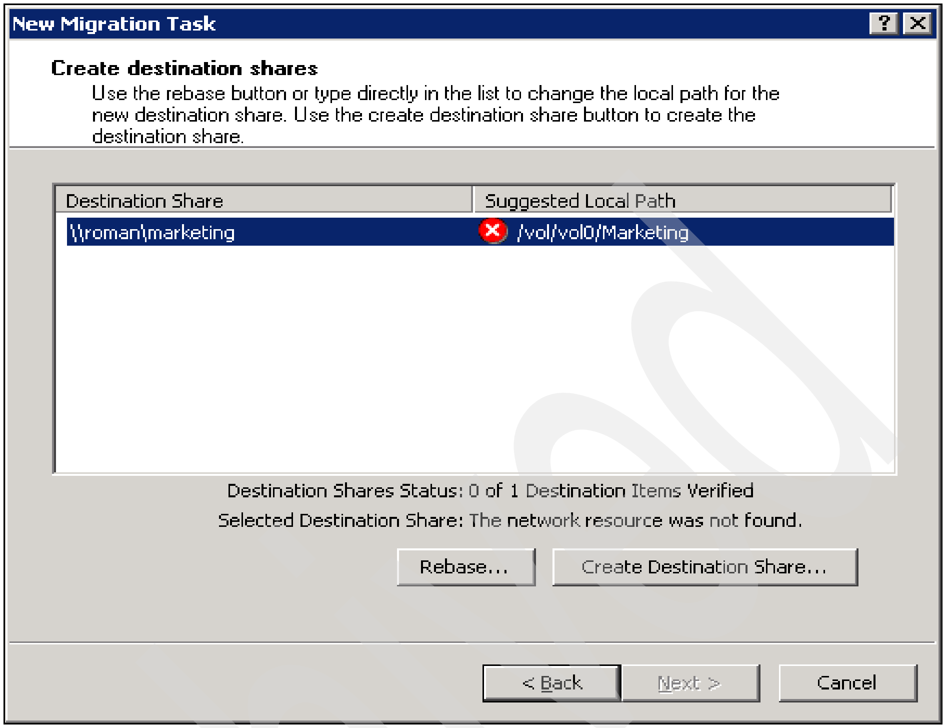

New Migration Task

Create destination shares

Use the rebase button or type directly in the list to change the local path for the new destination share. Use the create destination share button to create the destination share.

Destination Share	Suggested Local Path
\\roman\marketing	☒ /vol/vol0/Marketing

Destination Shares Status: 0 of 1 Destination Items Verified

Selected Destination Share: The network resource was not found.

Figure 8-22 Destination share

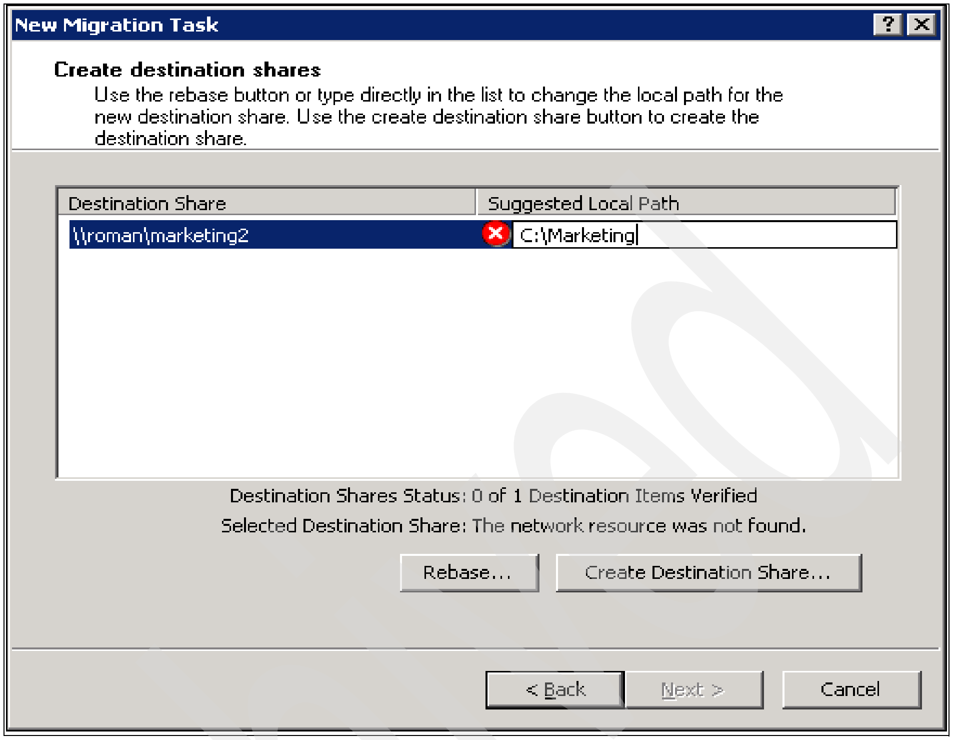
New Migration Task

Create destination shares

Use the rebase button or type directly in the list to change the local path for the new destination share. Use the create destination share button to create the destination share.

Destination Share	Suggested Local Path
\\roman\marketing2	☒ C:\Marketing

Destination Shares Status: 0 of 1 Destination Items Verified

Selected Destination Share: The network resource was not found.

Figure 8-23 Specifying desired name for destination share

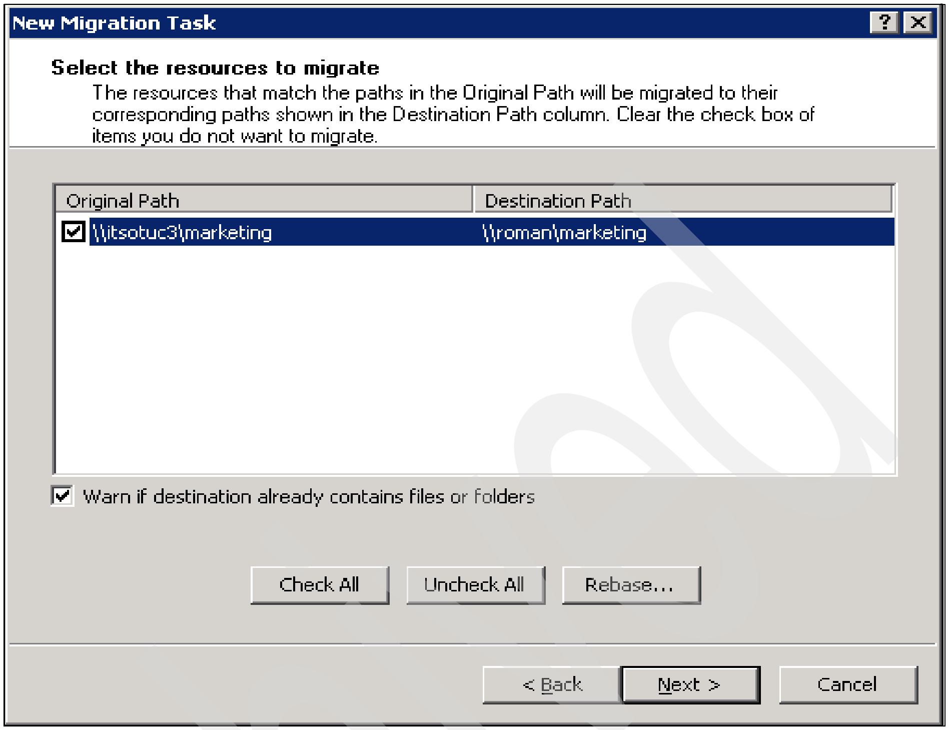

New Migration Task

Select the resources to migrate

The resources that match the paths in the Original Path will be migrated to their corresponding paths shown in the Destination Path column. Clear the check box of items you do not want to migrate.

Original Path	Destination Path
☑ \\itsotuc3\marketing	\\roman\marketing

☑ Warn if destination already contains files or folders

Figure 8-24 Destination share created

Set how you want the policy to transit from the incremental phase to the final phase (see Figure 8-25).

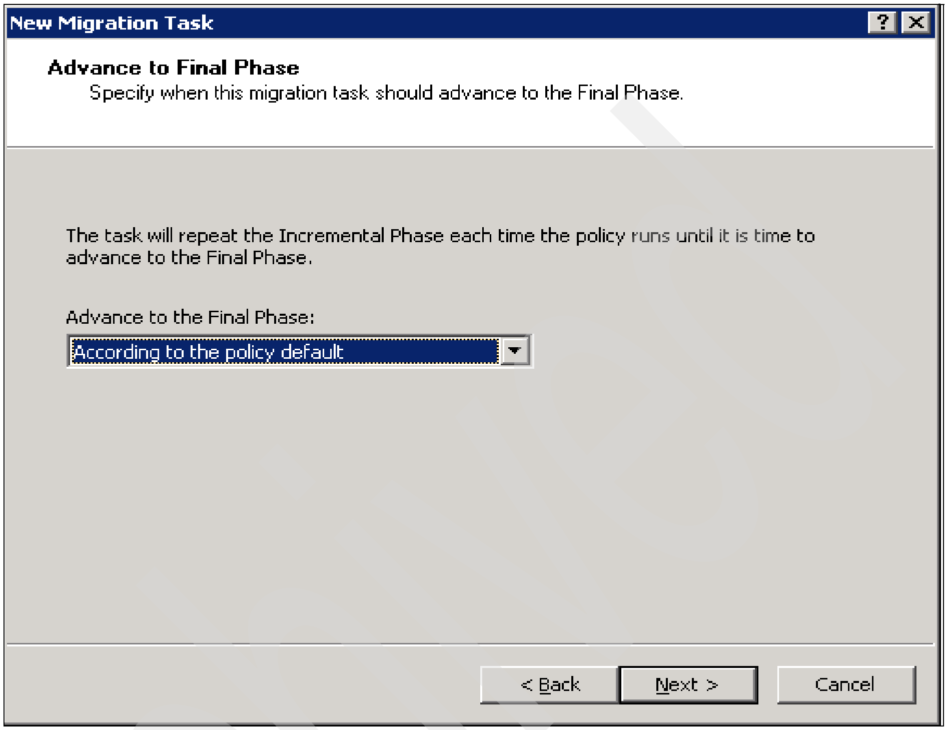
New Migration Task

Advance to Final Phase

Specify when this migration task should advance to the Final Phase.

The task will repeat the Incremental Phase each time the policy runs until it is time to advance to the Final Phase.

Advance to the Final Phase: According to the policy default

Figure 8-25 Transition to final phase

New task creation is complete at this stage (see Figure 8-26).

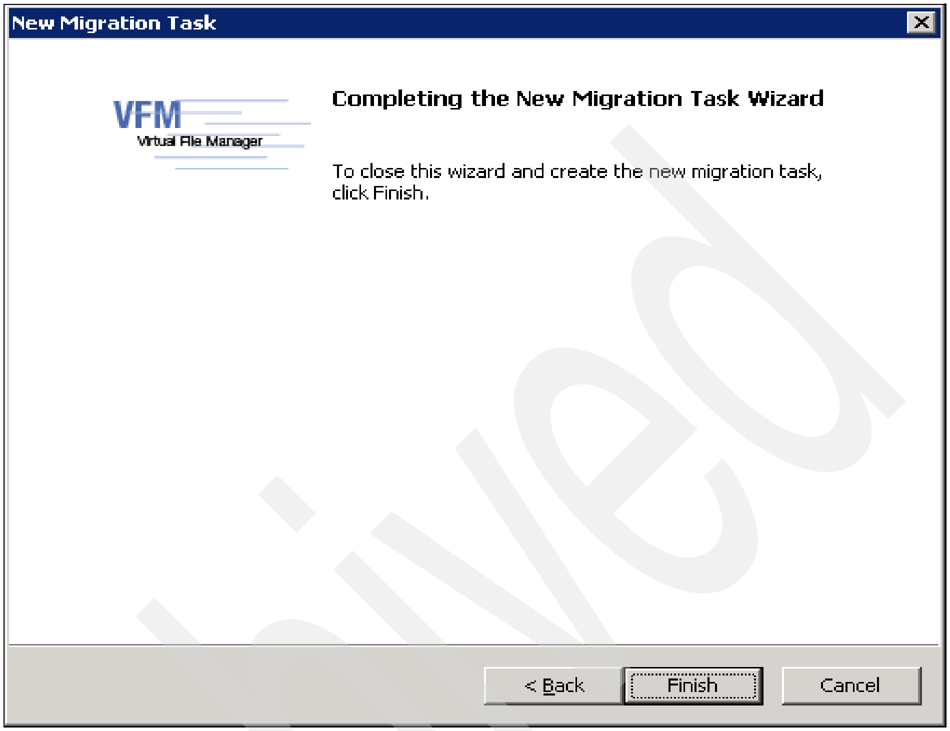

Completing the New Migration Task Wizard

To close this wizard and create the new migration task, click Finish.

Figure 8-26 Adding new migration task: complete

8.8 Migration policy example 2

In this section, we discuss running the migration policy. Once you have created the migration policy and added migration tasks to the policy (see Figure 8-27 on page 234), you are ready to run the policy.

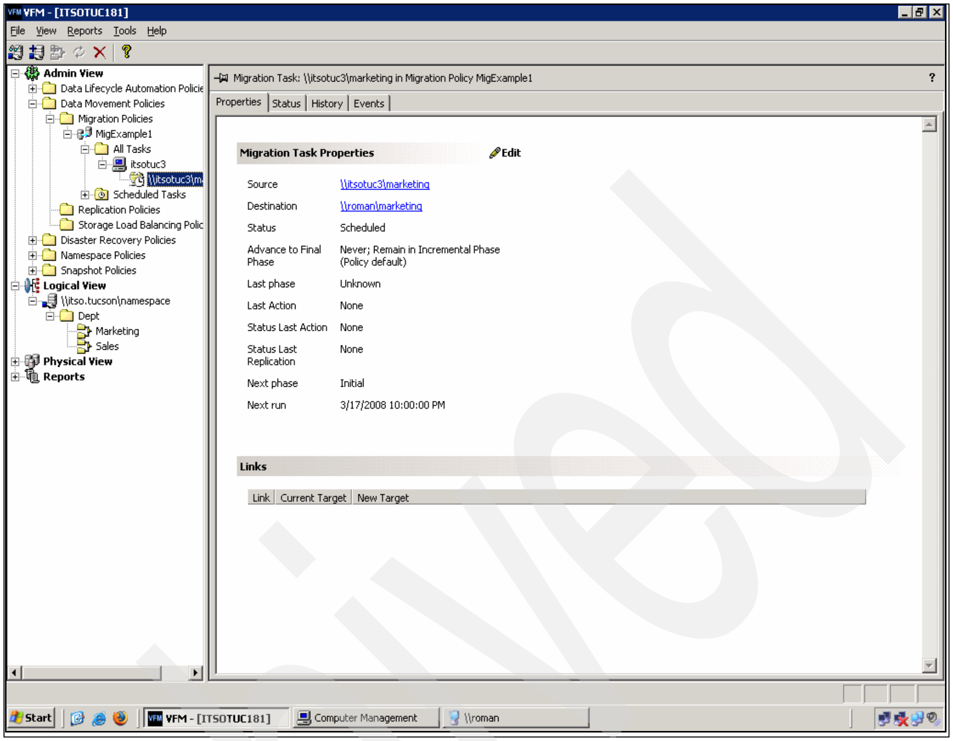

Migration Task: \\itsotuc3\marketing in Migration Policy MigExample1

Migration Task Properties

Source	\\itsotuc3\marketing
Destination	\\roman\marketing
Status	Scheduled
Advance to Final Phase	Never; Remain in Incremental Phase (Policy default)
Last phase	Unknown
Last Action	None
Status Last Action	None
Status Last Replication	None
Next phase	Initial
Next run	3/17/2008 10:00:00 PM

Links

Link	Current Target	New Target

Figure 8-27 New task created

The policy either can be invoked manually or can run automatically based on the schedule. For the purposes of this example, we run the policy manually through each phase of the migration policy (see Figure 8-28).

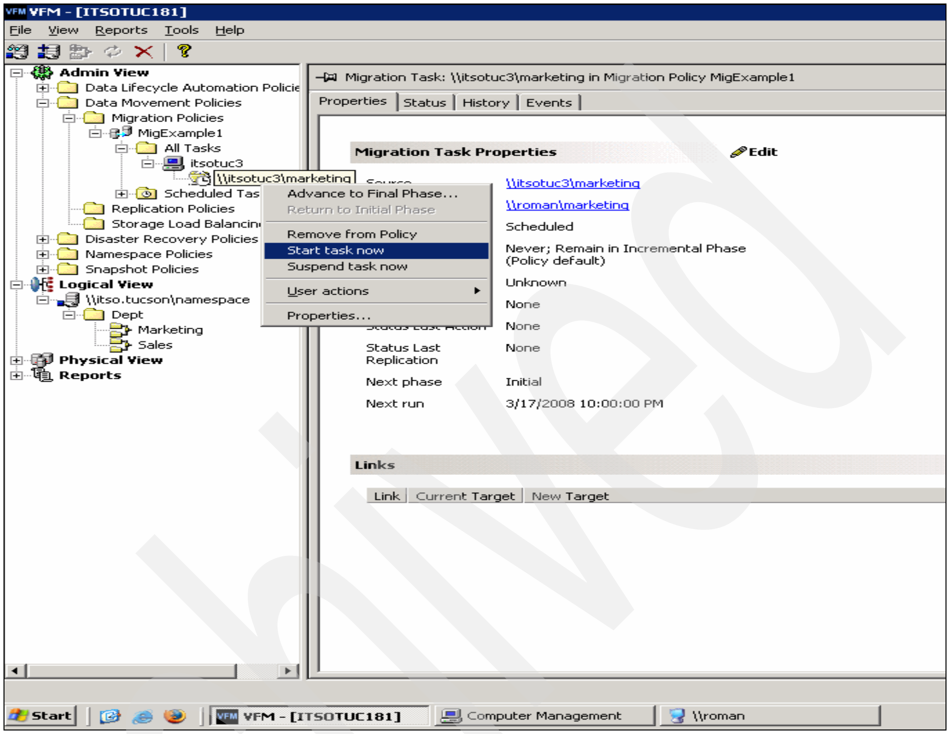

Migration Task: \\itsotuc3\marketing in Migration Policy MigExample1

Migration Task Properties

Source	\\itsotuc3\marketing
	\\roman\marketing
	Scheduled
	Never; Remain in Incremental Phase (Policy default)
	Unknown
	None
Status Last Action	None
Status Last Replication	None
Next phase	Initial
Next run	3/17/2008 10:00:00 PM

Figure 8-28   Starting the task manually

The initial phase is executed after starting it manually or at the scheduled time as set in the policy. The incremental phase is executed multiple times at the intervals specified in the policy. In this example, we are running it manually (see Figure 8-29).

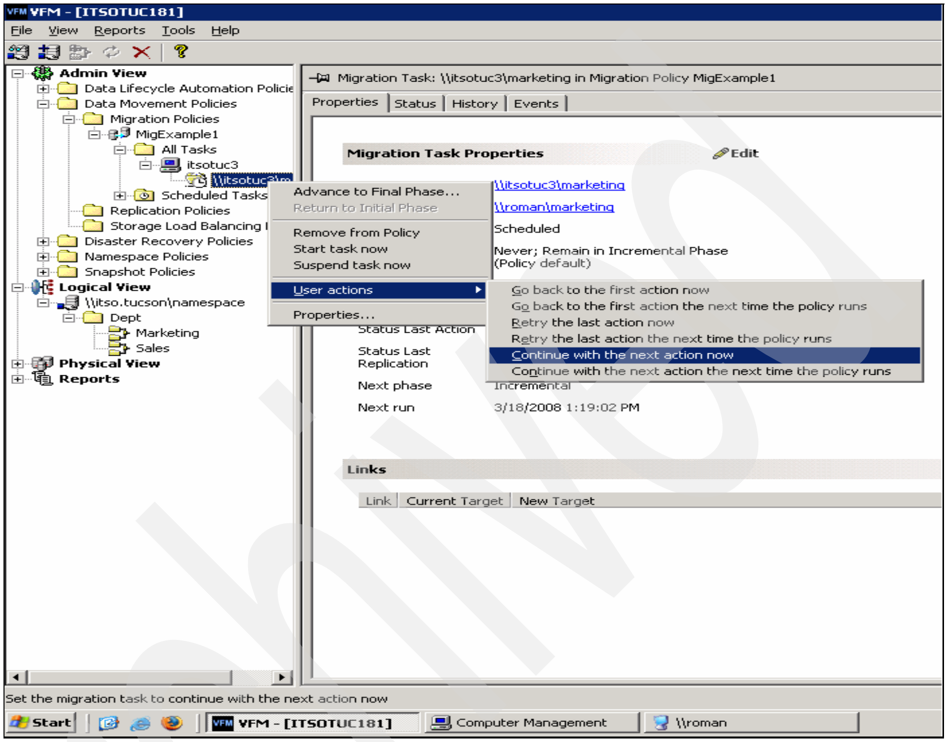

VFM - [ITSOTUC181]

File View Reports Tools Help

- Admin View
  - Data Lifecycle Automation Policie
  - Data Movement Policies
    - Migration Policies
      - MigExample1
        - All Tasks
          - itsotuc3
            - \\itsotuc3\m
        - Scheduled Tasks
    - Replication Policies
    - Storage Load Balancing I
  - Disaster Recovery Policies
  - Namespace Policies
  - Snapshot Policies
- Logical View
  - \\itso.tucson\namespace
    - Dept
      - Marketing
      - Sales
- Physical View
- Reports

Migration Task: \\itsotuc3\marketing in Migration Policy MigExample1

Properties | Status | History | Events

Migration Task Properties — Edit

	\\itsotuc3\marketing
	\\roman\marketing
	Scheduled
	Never; Remain in Incremental Phase (Policy default)
Status Last Action	
Status Last Replication	
Next phase	Incremental
Next run	3/18/2008 1:19:02 PM

Links

Link	Current Target	New Target

Advance to Final Phase...
Return to Initial Phase
Remove from Policy
Start task now
Suspend task now
User actions
Properties...
Go back to the first action now
Go back to the first action the next time the policy runs
Retry the last action now
Retry the last action the next time the policy runs
Continue with the next action now
Continue with the next action the next time the policy runs
Start
VFM - [ITSOTUC181]
Computer Management
\\roman

Set the migration task to continue with the next action now

Figure 8-29 Running migration policy: incremental phase

Check the history of the just concluded phase (see Figure 8-30).

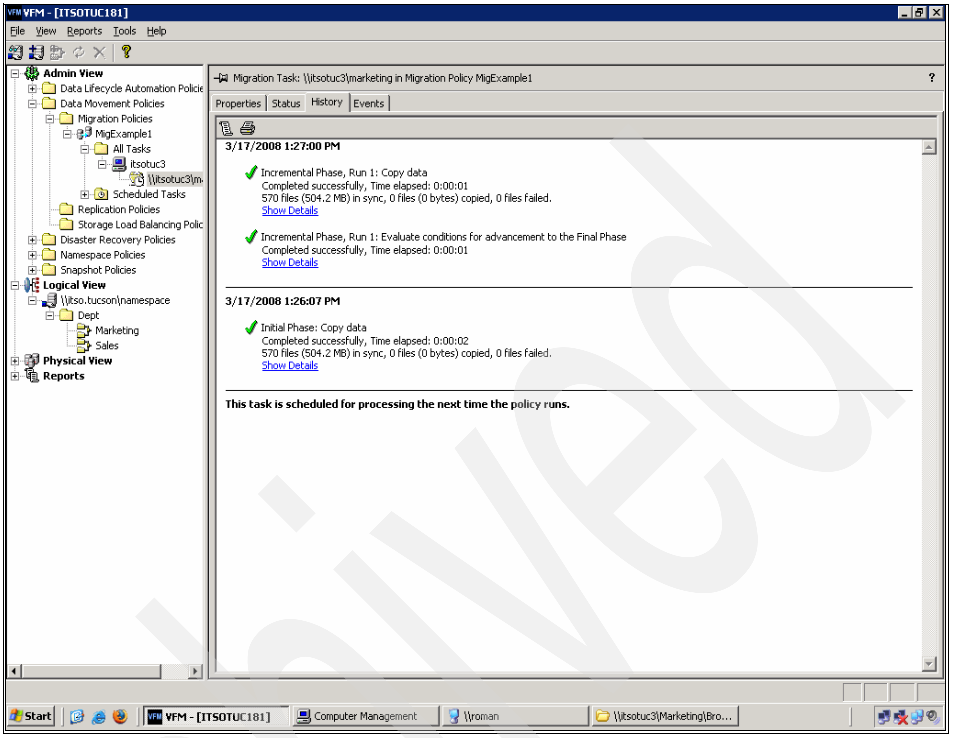

Migration Task: \\itsotuc3\marketing in Migration Policy MigExample1

3/17/2008 1:27:00 PM

✓ Incremental Phase, Run 1: Copy data
Completed successfully, Time elapsed: 0:00:01
570 files (504.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

✓ Incremental Phase, Run 1: Evaluate conditions for advancement to the Final Phase
Completed successfully, Time elapsed: 0:00:01
Show Details

3/17/2008 1:26:07 PM

✓ Initial Phase: Copy data
Completed successfully, Time elapsed: 0:00:02
570 files (504.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

This task is scheduled for processing the next time the policy runs.

Figure 8-30 History

The window shown in Figure 8-30 shows details about actions taken in each phase. These details include the action taken, number of files copied, and time elapsed during the action.

After the successful completion of the initial and incremental phases, the final phase is executed as scheduled in the policy. In this example, we are executing it manually (see Figure 8-31).

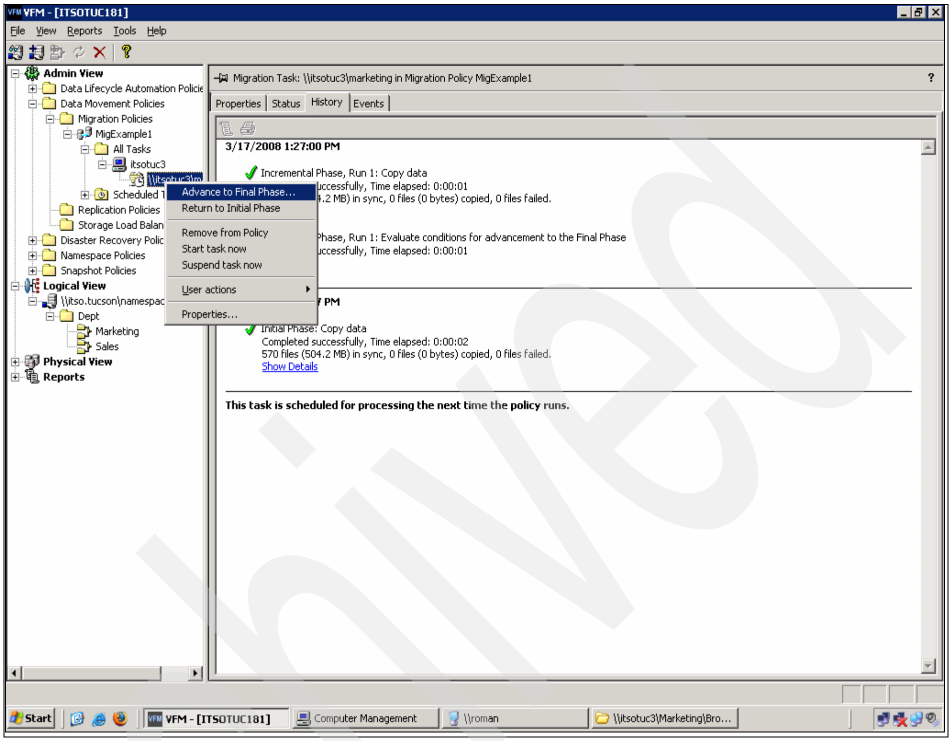

3/17/2008 1:27:00 PM

✓ Incremental Phase, Run 1: Copy data
uccessfully, Time elapsed: 0:00:01
4.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.

Phase, Run 1: Evaluate conditions for advancement to the Final Phase
uccessfully, Time elapsed: 0:00:01

7 PM

✓ Initial Phase: Copy data
Completed successfully, Time elapsed: 0:00:02
570 files (504.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

This task is scheduled for processing the next time the policy runs.

Figure 8-31 Moving to final phase

Instruct VFM to start the final phase (see Figure 8-32).

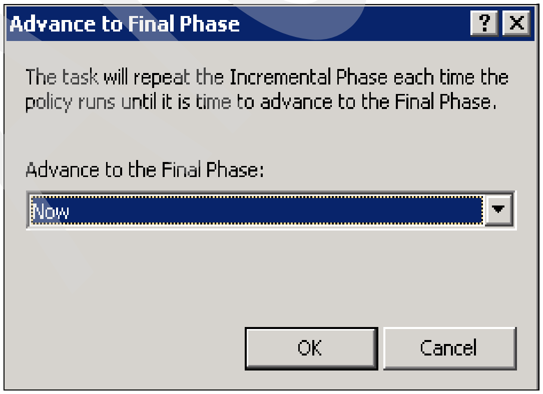

The task will repeat the Incremental Phase each time the policy runs until it is time to advance to the Final Phase.

Advance to the Final Phase: Now

Figure 8-32 Advance to Final Phase dialog

Check the history to determine whether the namespace update has completed successfully (see Figure 8-33).

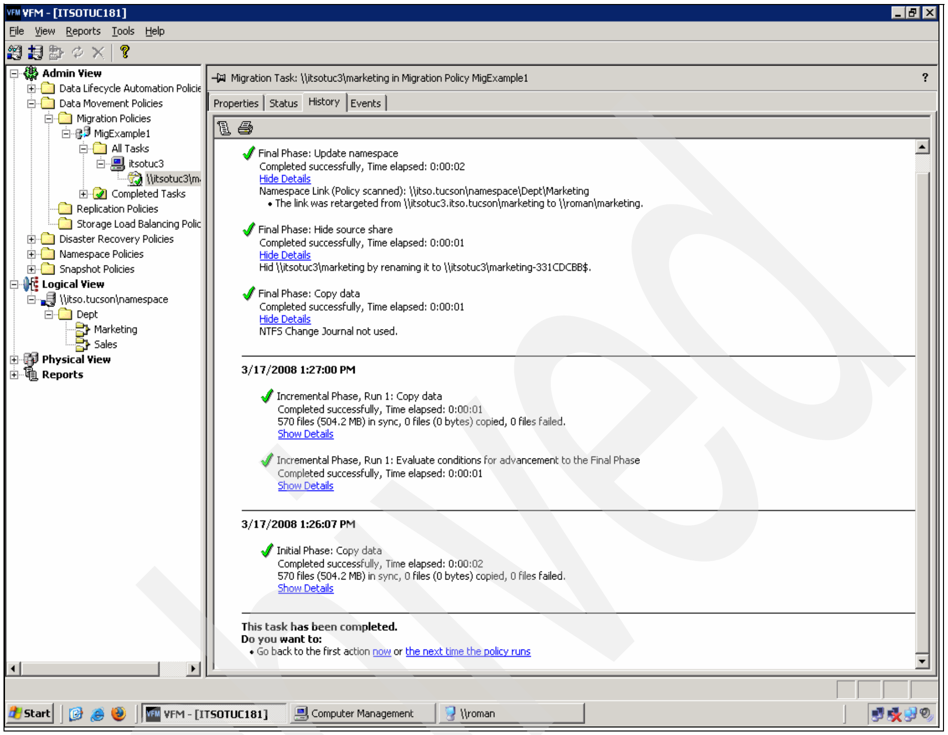

Migration Task: \\itsotuc3\marketing in Migration Policy MigExample1

✓ Final Phase: Update namespace
Completed successfully, Time elapsed: 0:00:02
Hide Details
Namespace Link (Policy scanned): \\itso.tucson\namespace\Dept\Marketing
- The link was retargeted from \\itsotuc3.itso.tucson\marketing to \\roman\marketing.

✓ Final Phase: Hide source share
Completed successfully, Time elapsed: 0:00:01
Hide Details
Hid \\itsotuc3\marketing by renaming it to \\itsotuc3\marketing-331CDCBB$.

✓ Final Phase: Copy data
Completed successfully, Time elapsed: 0:00:01
Hide Details
NTFS Change Journal not used.

3/17/2008 1:27:00 PM

✓ Incremental Phase, Run 1: Copy data
Completed successfully, Time elapsed: 0:00:01
570 files (504.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

✓ Incremental Phase, Run 1: Evaluate conditions for advancement to the Final Phase
Completed successfully, Time elapsed: 0:00:01
Show Details

3/17/2008 1:26:07 PM

✓ Initial Phase: Copy data
Completed successfully, Time elapsed: 0:00:02
570 files (504.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

This task has been completed.
Do you want to:
- Go back to the first action now or the next time the policy runs

Figure 8-33 Final phase history

Upon completion of all phases, you notice that:

- The namespace is updated to point the DFS link to the destination share of the migration policy (see Figure 8-34 and Figure 8-35 on page 241).

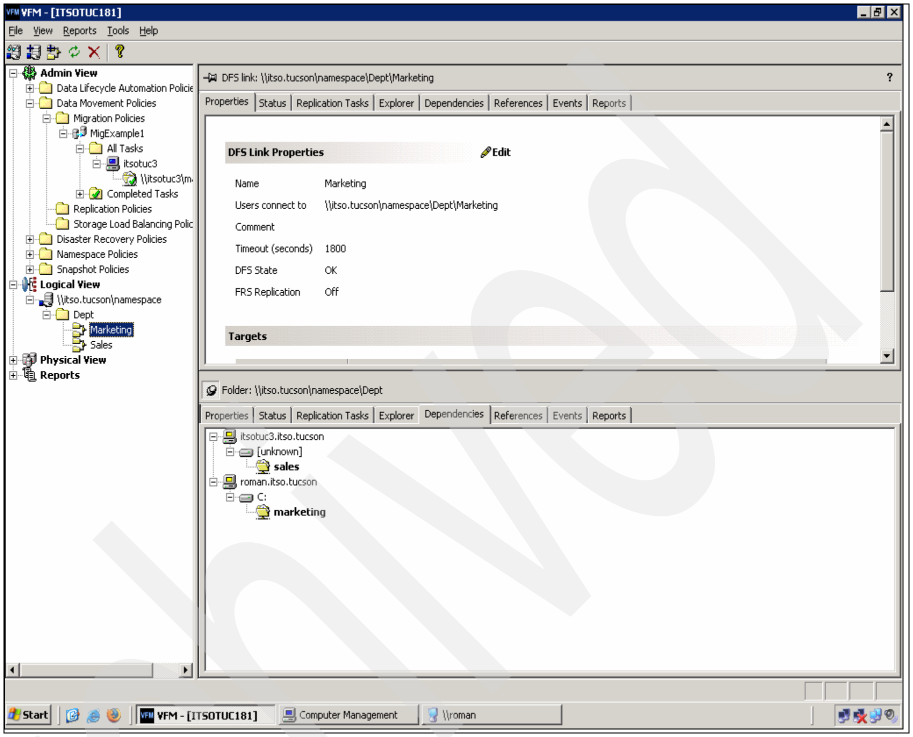

DFS Link Properties

Name Marketing
Users connect to \\itso.tucson\namespace\Dept\Marketing
Comment
Timeout (seconds) 1800
DFS State OK
FRS Replication Off

Figure 8-34 Dependency of DFS root changed

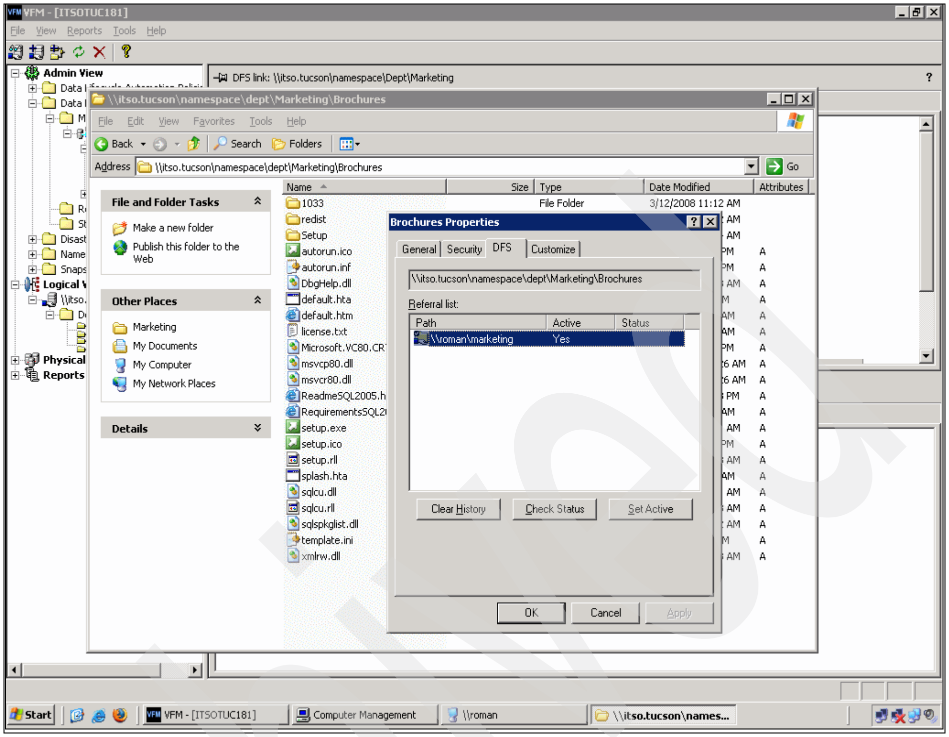

VFM - [ITSOTUC181]

DFS link: \\itso.tucson\namespace\Dept\Marketing

\\itso.tucson\namespace\dept\Marketing\Brochures

Address: \\itso.tucson\namespace\dept\Marketing\Brochures

Brochures Properties — DFS

\\itso.tucson\namespace\dept\Marketing\Brochures

Referral list:

Path	Active	Status
\\roman\marketing	Yes	

Figure 8-35 DFS link updated

- The manifest (shown in the History tab of the task) indicates how the task performed, the number of files in sync, the number of files copied, and so on. Any errors in the policy are also reported in the History tab.
- If set in the policy, the source share is disabled.

- The user can still access the files as before using the same namespace path even though the data has been relocated from one server to another (see Figure 8-36).

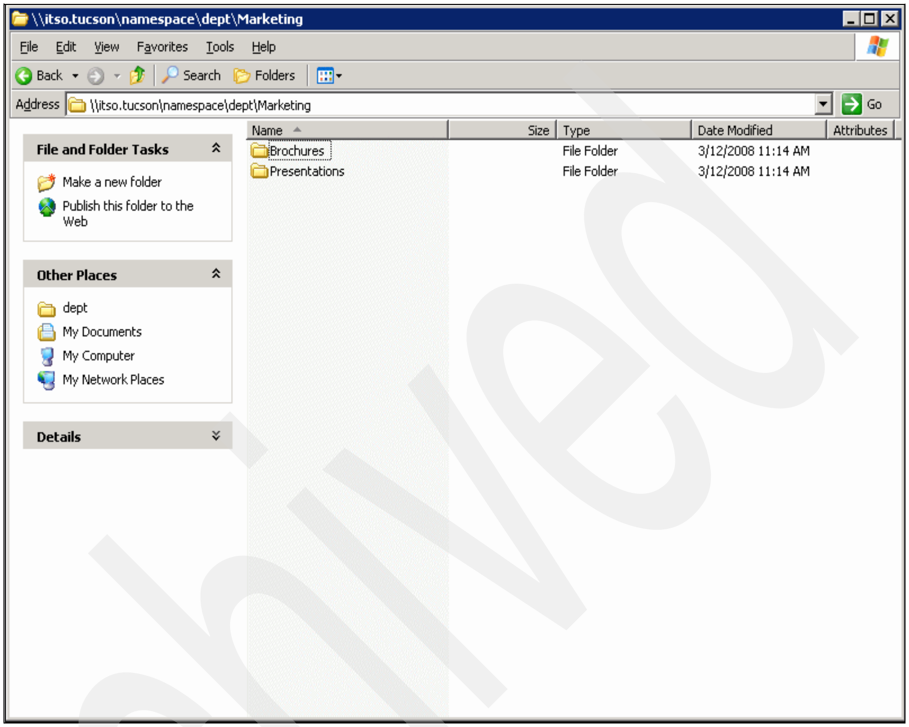

Name	Size	Type	Date Modified	Attributes
Brochures		File Folder	3/12/2008 11:14 AM	
Presentations		File Folder	3/12/2008 11:14 AM	

Figure 8-36 User data access

8.9 Migration policy example 3

In this example, we migrate all the shares from one server to another N series. A separate task is created for each share to be moved. At this point, you update the namespace, and all the shares are targeted to the new location. You can then decommission the old server.

Consider the namespace used in our first example with Marketing and Sales links under the Dept Folder of the namespace (see 8.5, “Migration policy example 1” on page 204).

We added one migration task in our first example to move the Marketing link to N series. At this point, we add one more migration task to move the Sales link similarly. Follow these steps:

1. Specify source and destination server names (see Figure 8-37).

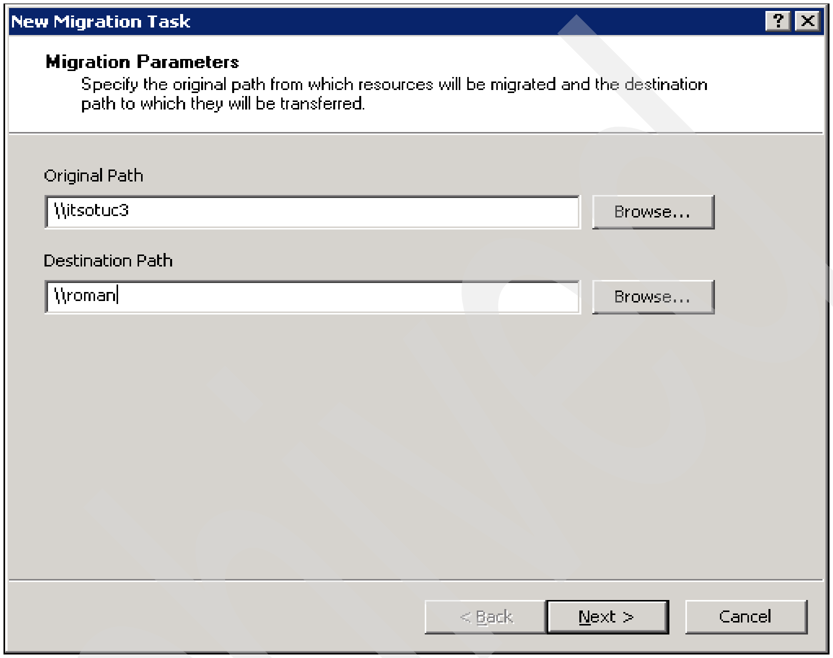
New Migration Task

Migration Parameters

Specify the original path from which resources will be migrated and the destination path to which they will be transferred.

Original Path: \\itsotuc3

Destination Path: \\roman

Figure 8-37 Adding a migration task: specifying source and destination

Specifying server names without specifying share names results in listing all the possible shares with specified servers (see Figure 8-38).

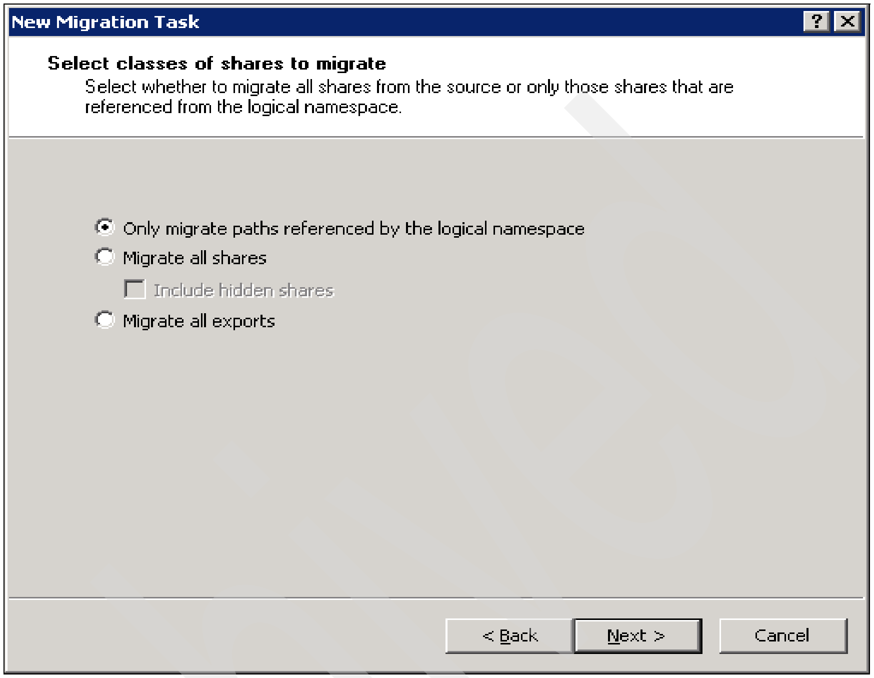
New Migration Task

Select classes of shares to migrate

Select whether to migrate all shares from the source or only those shares that are referenced from the logical namespace.

● Only migrate paths referenced by the logical namespace

○ Migrate all shares

□ Include hidden shares

○ Migrate all exports

Figure 8-38 Selecting classes of shares to migrate

2. Select the share name you want to migrate (see Figure 8-39).

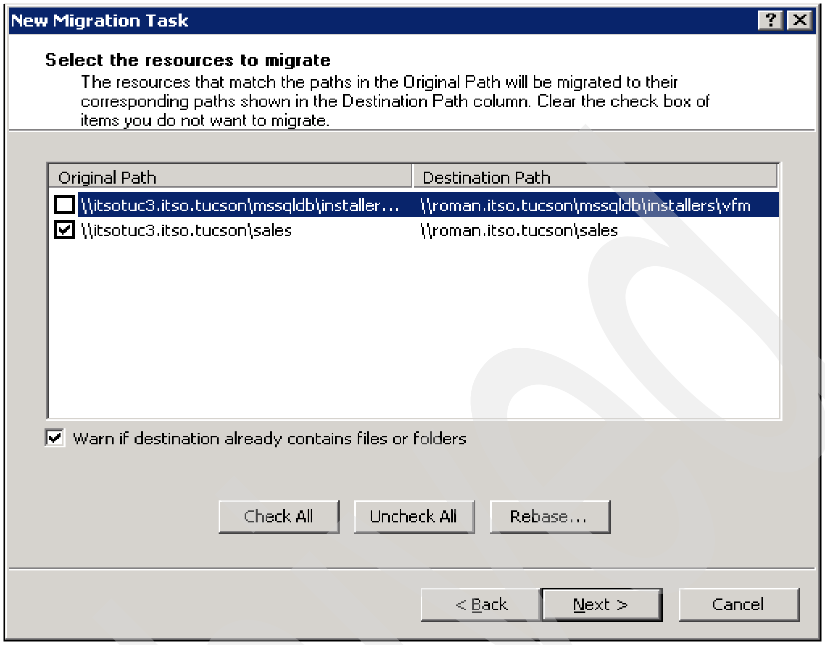

Select the resources to migrate

The resources that match the paths in the Original Path will be migrated to their corresponding paths shown in the Destination Path column. Clear the check box of items you do not want to migrate.

Original Path	Destination Path
□ \\itsotuc3.itso.tucson\mssqldb\installer...	\\roman.itso.tucson\mssqldb\installers\vfm
☑ \\itsotuc3.itso.tucson\sales	\\roman.itso.tucson\sales

☑ Warn if destination already contains files or folders

Figure 8-39 Selecting the share to migrate

VFM prompts you to create the destination shares (see Figure 8-40).

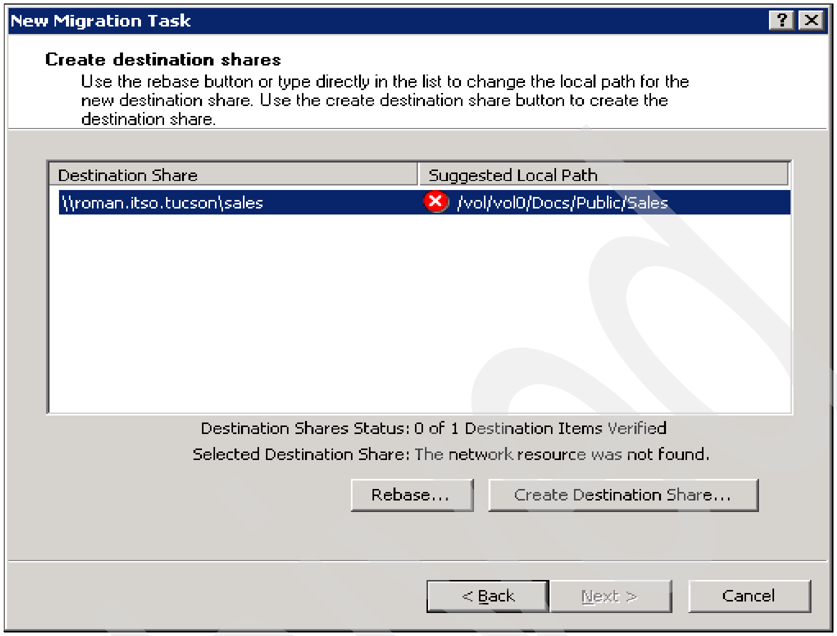

Create destination shares

Use the rebase button or type directly in the list to change the local path for the new destination share. Use the create destination share button to create the destination share.

Destination Share	Suggested Local Path
\\roman.itso.tucson\sales	✗ /vol/vol0/Docs/Public/Sales

Destination Shares Status: 0 of 1 Destination Items Verified

Selected Destination Share: The network resource was not found.

Figure 8-40 Destination share creation prompt

3. Specify the desired name for the destination share (see Figure 8-41).

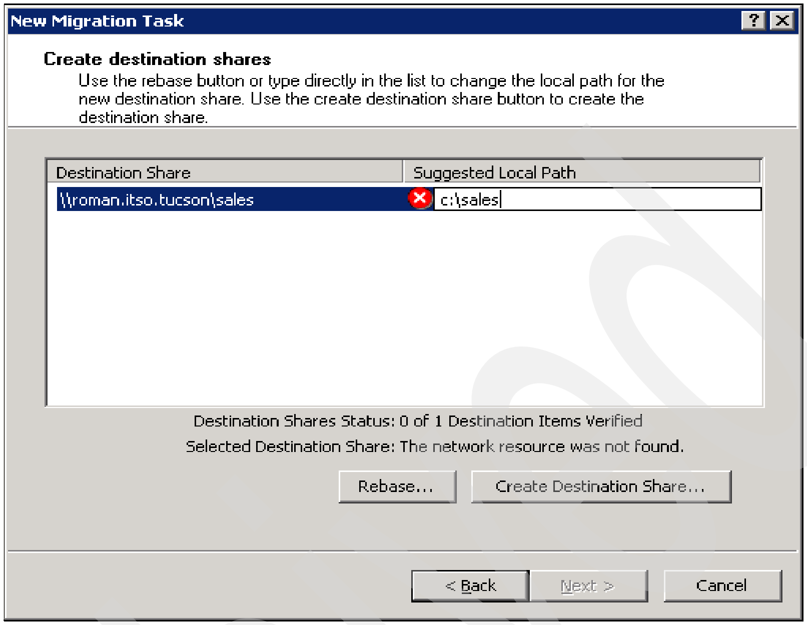

New Migration Task

Create destination shares

Use the rebase button or type directly in the list to change the local path for the new destination share. Use the create destination share button to create the destination share.

Destination Share	Suggested Local Path
\\roman.itso.tucson\sales	c:\sales

Destination Shares Status: 0 of 1 Destination Items Verified

Selected Destination Share: The network resource was not found.

Figure 8-41 Specifying name for destination share

4. Click **Create destination share** to create it (see Figure 8-42).

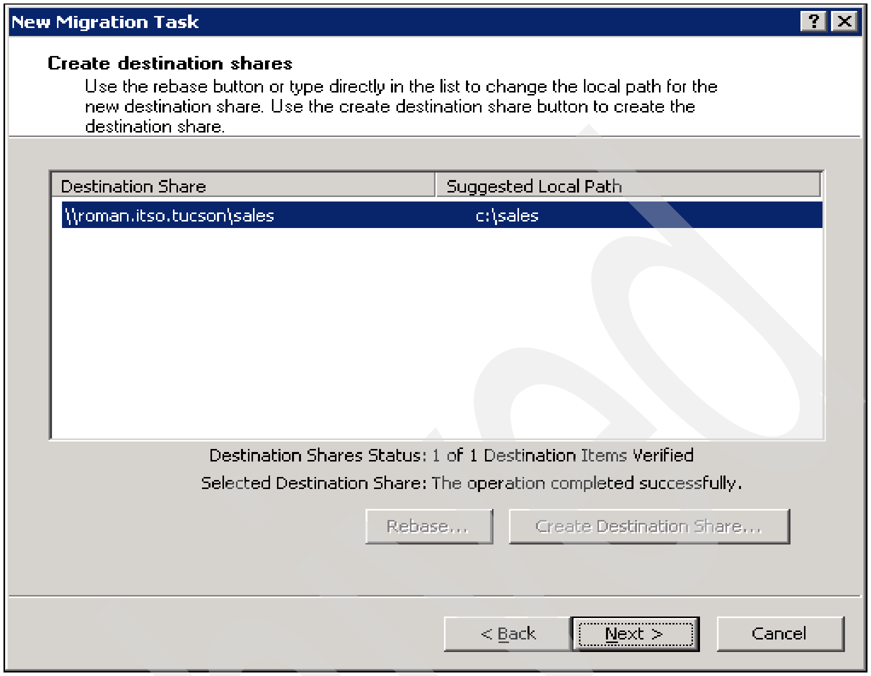

Create destination shares

Use the rebase button or type directly in the list to change the local path for the new destination share. Use the create destination share button to create the destination share.

Destination Share	Suggested Local Path
\\roman.itso.tucson\sales	c:\sales

Destination Shares Status: 1 of 1 Destination Items Verified

Selected Destination Share: The operation completed successfully.

Figure 8-42 Destination share created

Creation of the new migration task is complete (see Figure 8-43).

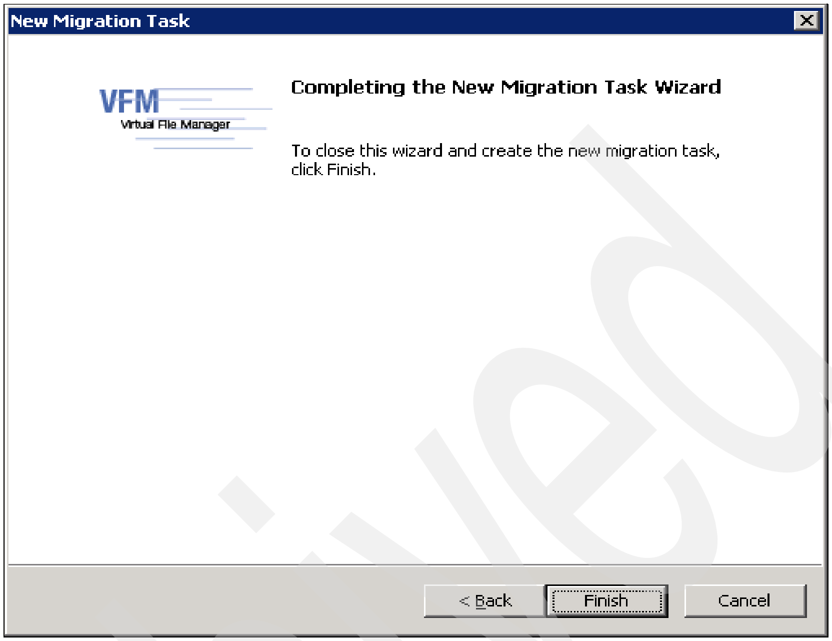

New Migration Task

Completing the New Migration Task Wizard

To close this wizard and create the new migration task, click Finish.

Figure 8-43 New migration task complete

5. Check the newly created migration task (see Figure 8-44).

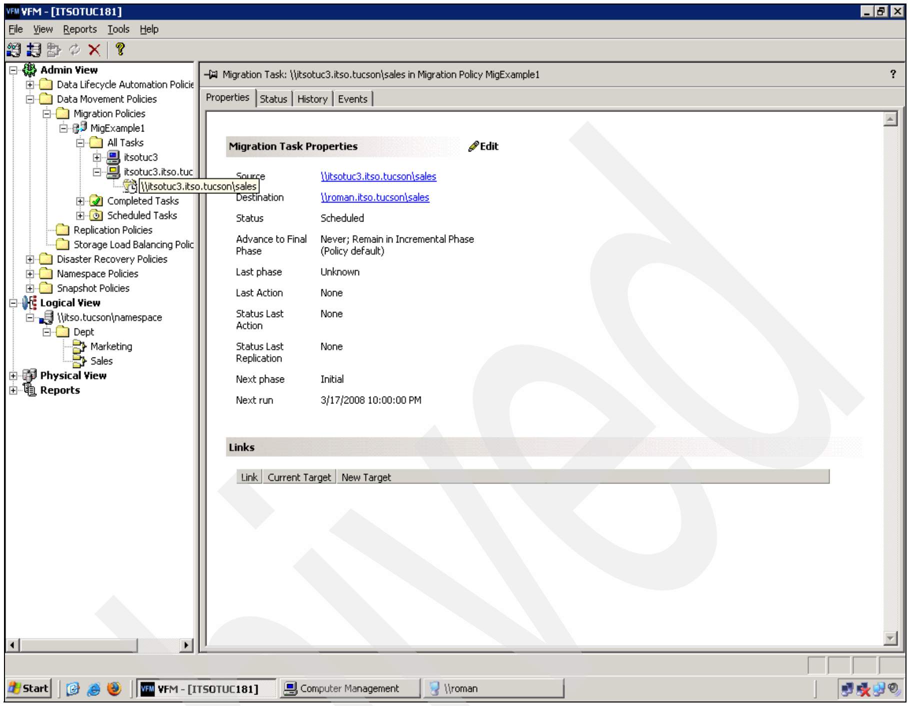

Migration Task: \\itsotuc3.itso.tucson\sales in Migration Policy MigExample1

Migration Task Properties

Source	\\itsotuc3.itso.tucson\sales
Destination	\\roman.itso.tucson\sales
Status	Scheduled
Advance to Final Phase	Never; Remain in Incremental Phase (Policy default)
Last phase	Unknown
Last Action	None
Status Last Action	None
Status Last Replication	None
Next phase	Initial
Next run	3/17/2008 10:00:00 PM

Links

Link	Current Target	New Target

Figure 8-44 Properties of migration task

Migration phases are initiated at the scheduled time of the migration policy. In this example, the migration phases are started manually, which is similar to our first example (see 8.5, “Migration policy example 1” on page 204).

6. After the completion of all three phases of the migration, check the history for the details on the migration phases (see Figure 8-45).

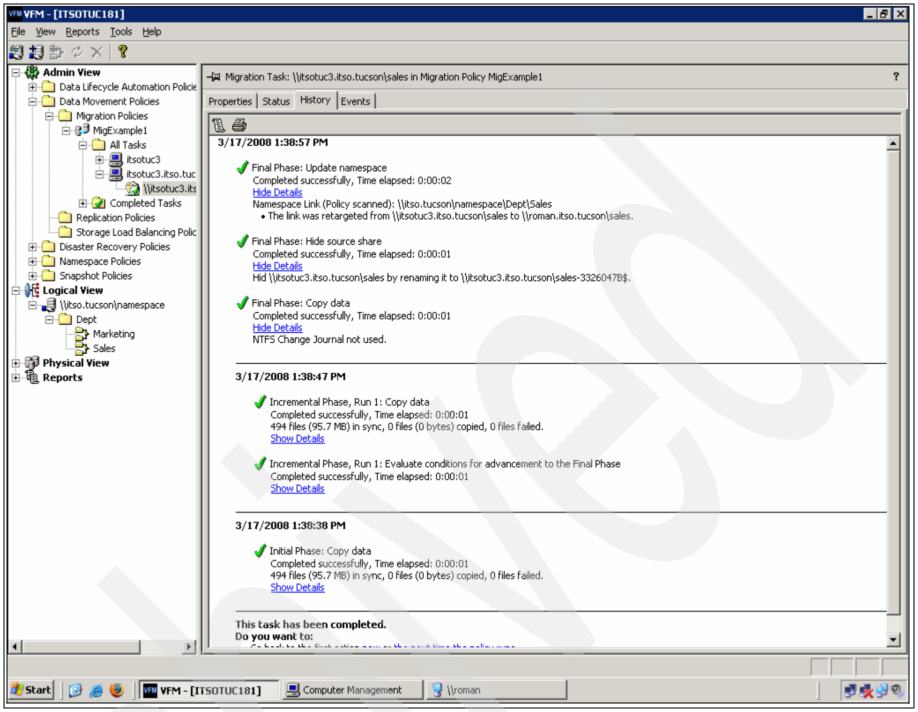

Migration Task: \\itsotuc3.itso.tucson\sales in Migration Policy MigExample1

3/17/2008 1:38:57 PM

✓ Final Phase: Update namespace
Completed successfully, Time elapsed: 0:00:02
Hide Details
Namespace Link (Policy scanned): \\itso.tucson\namespace\Dept\Sales
- The link was retargeted from \\itsotuc3.itso.tucson\sales to \\roman.itso.tucson\sales.

✓ Final Phase: Hide source share
Completed successfully, Time elapsed: 0:00:01
Hide Details
Hid \\itsotuc3.itso.tucson\sales by renaming it to \\itsotuc3.itso.tucson\sales-3326047B$.

✓ Final Phase: Copy data
Completed successfully, Time elapsed: 0:00:01
Hide Details
NTFS Change Journal not used.

3/17/2008 1:38:47 PM

✓ Incremental Phase, Run 1: Copy data
Completed successfully, Time elapsed: 0:00:01
494 files (95.7 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

✓ Incremental Phase, Run 1: Evaluate conditions for advancement to the Final Phase
Completed successfully, Time elapsed: 0:00:01
Show Details

3/17/2008 1:38:38 PM

✓ Initial Phase: Copy data
Completed successfully, Time elapsed: 0:00:01
494 files (95.7 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

This task has been completed.
Do you want to:

Figure 8-45 Migration of Sales share complete

7. Check the dependency tab to confirm the current location of the Sales share (see Figure 8-46).

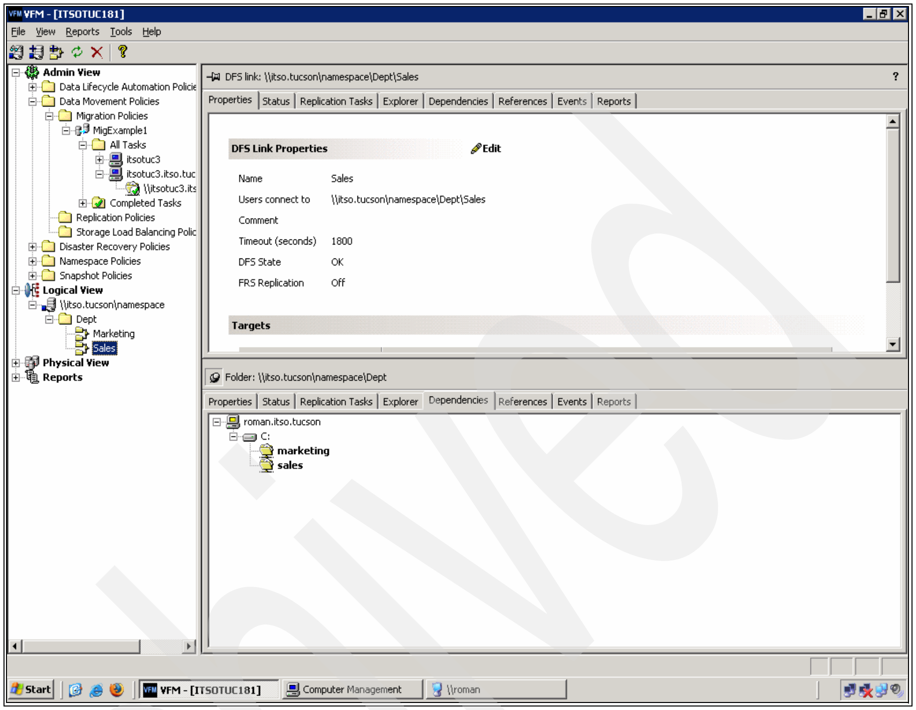

DFS link: \\itso.tucson\namespace\Dept\Sales

DFS Link Properties

Name Sales
Users connect to \\itso.tucson\namespace\Dept\Sales
Comment
Timeout (seconds) 1800
DFS State OK
FRS Replication Off

Targets

Folder: \\itso.tucson\namespace\Dept

roman.itso.tucson
C:
marketing
sales

Figure 8-46 Both shares moved

In the next section, we describe a client use case of moving all shares of the file server to N series.

8.10 Moving shares - example 4

In this section, we consider an example of moving some shares of a file server to one N series and the remaining shares to another N series. In other words, we spread data from one server to different servers.

The Dept folder contains the Marketing and Sales link. We want to move the Marketing shares on one N series and the Sales shares to another N series (see Figure 8-47).

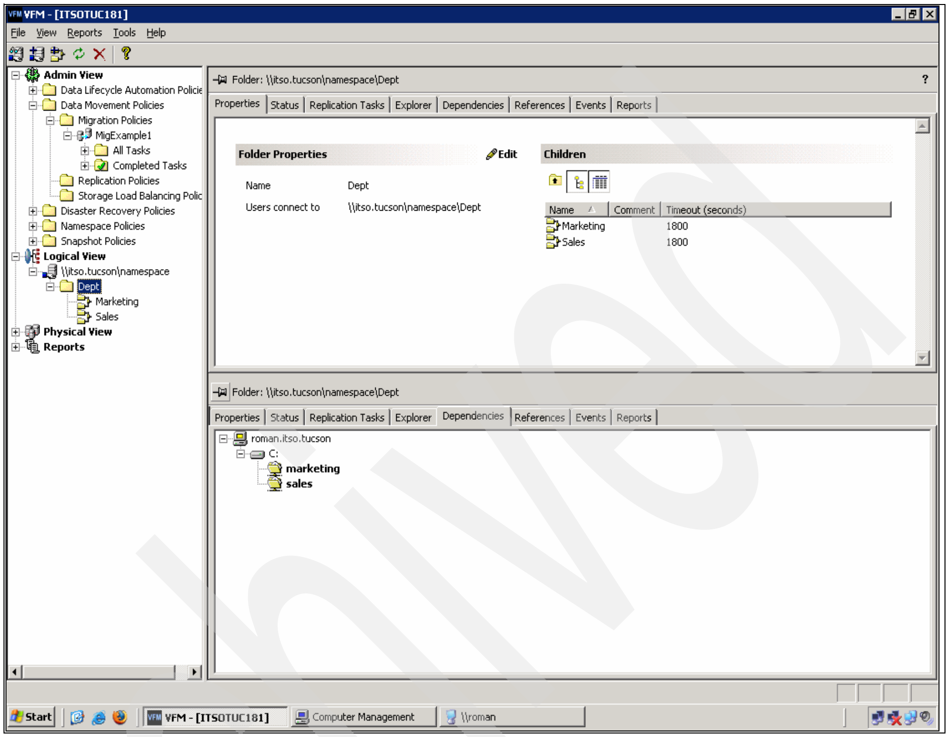
VFM - [ITSOTUC181]
File View Reports Tools Help
Admin View
Data Lifecycle Automation Policie
Data Movement Policies
Migration Policies
MigExample1
All Tasks
Completed Tasks
Replication Policies
Storage Load Balancing Polic
Disaster Recovery Policies
Namespace Policies
Snapshot Policies
Logical View
\\itso.tucson\namespace
Dept
Marketing
Sales
Physical View
Reports
Folder: \\itso.tucson\namespace\Dept
Properties Status Replication Tasks Explorer Dependencies References Events Reports
Folder Properties
Edit
Children
Folder: \\itso.tucson\namespace\Dept
Properties Status Replication Tasks Explorer Dependencies References Events Reports
roman.itso.tucson
C:
marketing
sales
Start
VFM - [ITSOTUC181]
Computer Management
\\roman

Name Dept

Users connect to \\itso.tucson\namespace\Dept

Name	Comment	Timeout (seconds)
Marketing		1800
Sales		1800

Figure 8-47 Moving shares

In our example, the Marketing and Sales shares are present on the file server Roman (see Figure 8-48).

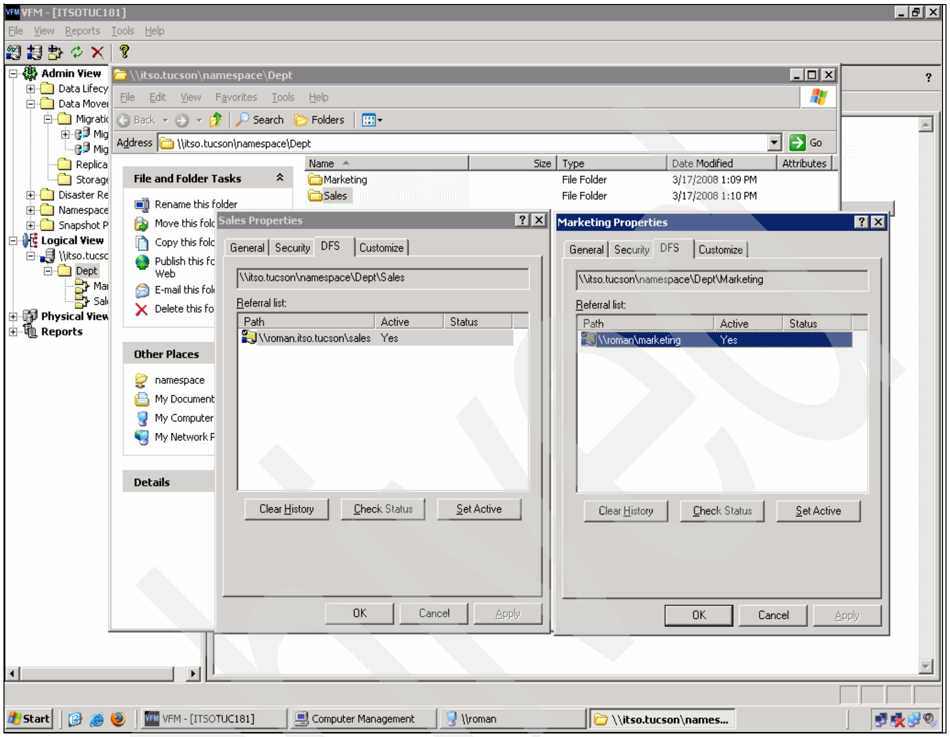

Address: \\itso.tucson\namespace\Dept

Name	Size	Type	Date Modified	Attributes
Marketing		File Folder	3/17/2008 1:09 PM	
Sales		File Folder	3/17/2008 1:10 PM	

Sales Properties: \\itso.tucson\namespace\Dept\Sales

Path	Active	Status
\\roman.itso.tucson\sales	Yes	

Marketing Properties: \\itso.tucson\namespace\Dept\Marketing

Path	Active	Status
\\roman\marketing	Yes	

Figure 8-48 Marketing and Sales share on Roman

To move the Marketing shares on one N series and the Sales shares to another N series, we complete the following steps:

1. Create a new migration policy MigExample3 as shown in Figure 8-49.

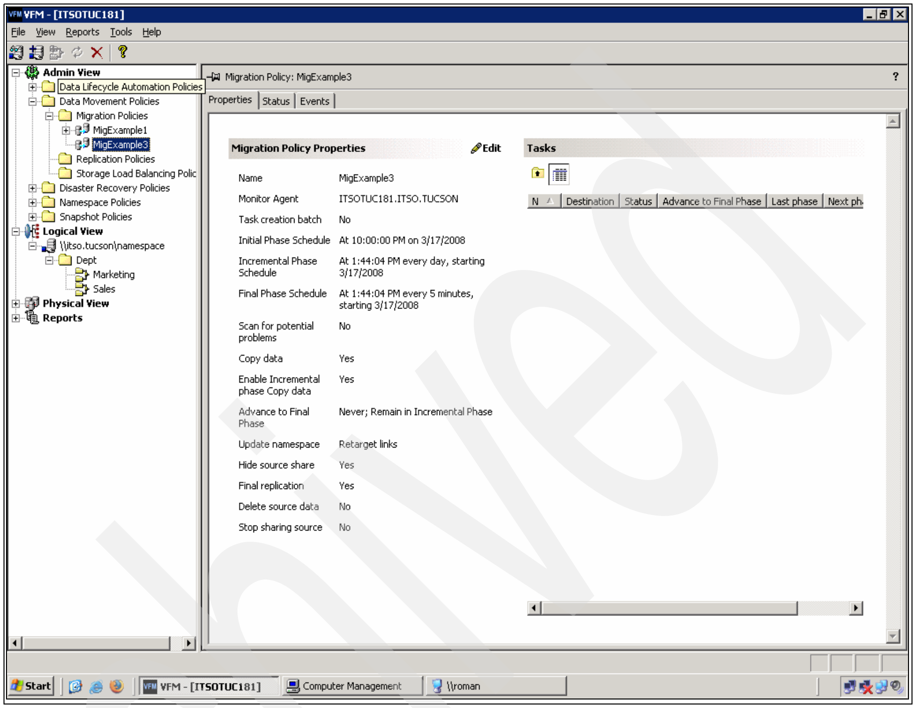

Migration Policy Properties	
Name	MigExample3
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Task creation batch	No
Initial Phase Schedule	At 10:00:00 PM on 3/17/2008
Incremental Phase Schedule	At 1:44:04 PM every day, starting 3/17/2008
Final Phase Schedule	At 1:44:04 PM every 5 minutes, starting 3/17/2008
Scan for potential problems	No
Copy data	Yes
Enable Incremental phase Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Update namespace	Retarget links
Hide source share	Yes
Final replication	Yes
Delete source data	No
Stop sharing source	No

Figure 8-49 Creating migration policy

2. Add new migration tasks for each share.

3. Move the Sales link on itsotuc3 N series (see Figure 8-50).

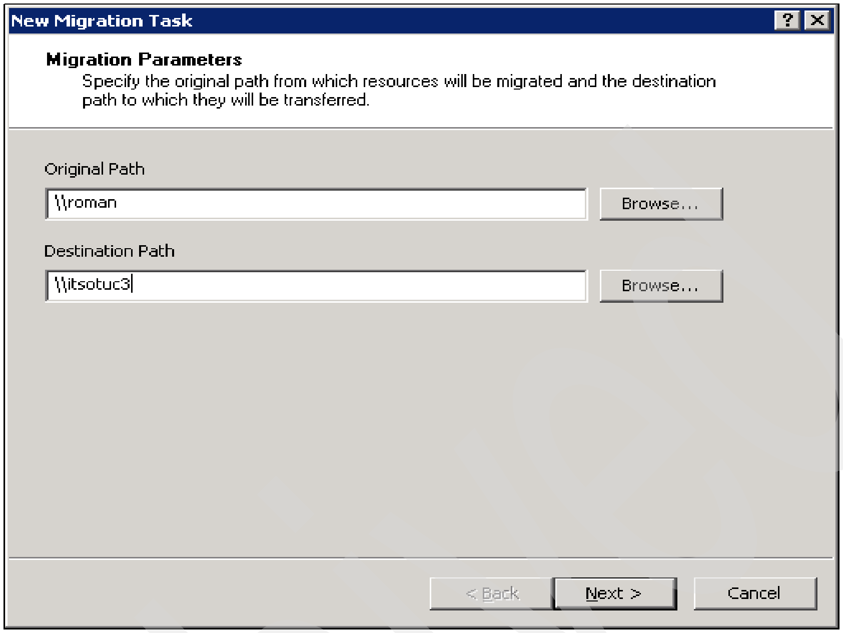

New Migration Task

Migration Parameters

Specify the original path from which resources will be migrated and the destination path to which they will be transferred.

Original Path: \\roman

Destination Path: \\itsotuc3

Figure 8-50 Task Sales share

VFM warns about duplication of the migration task path if found (see Figure 8-51).

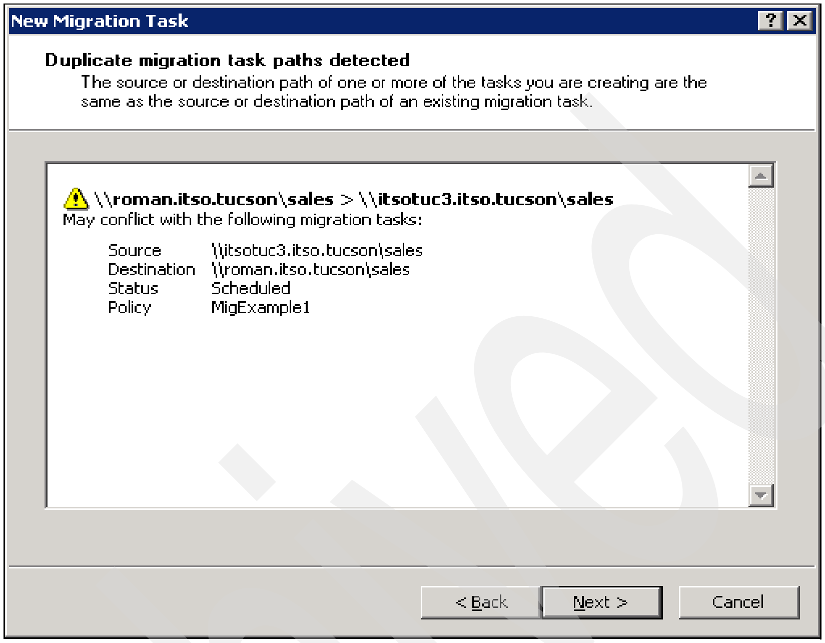

Duplicate migration task paths detected

The source or destination path of one or more of the tasks you are creating are the same as the source or destination path of an existing migration task.

\\roman.itso.tucson\sales > \\itsotuc3.itso.tucson\sales

May conflict with the following migration tasks:

Source	\\itsotuc3.itso.tucson\sales
Destination	\\roman.itso.tucson\sales
Status	Scheduled
Policy	MigExample1

Figure 8-51 Duplicate migration task paths

VFM also warns if the destination share is not empty (see Figure 8-52).

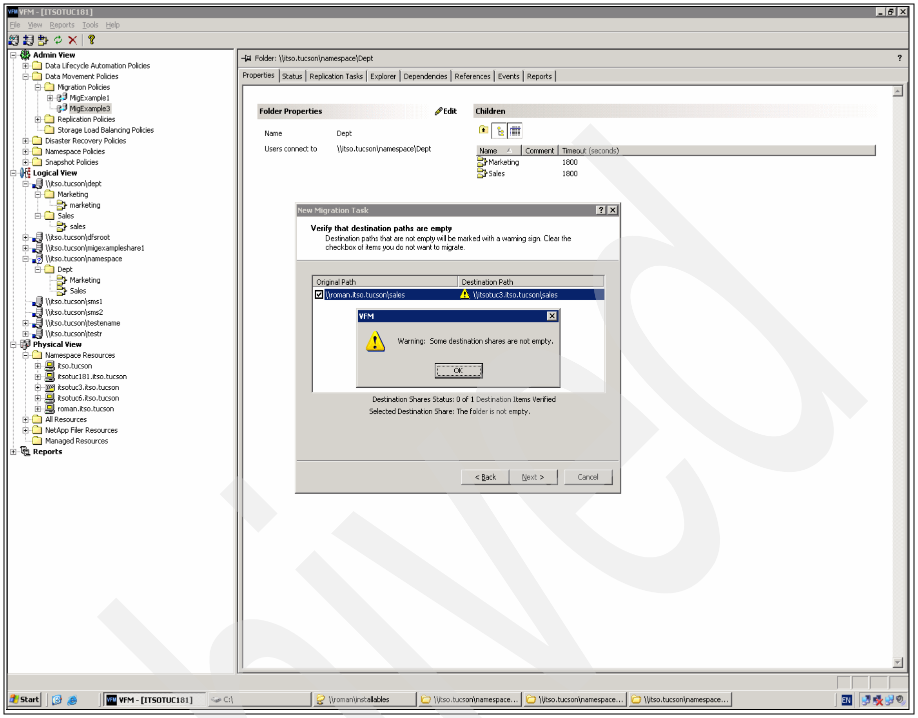

Folder: \\itso.tucson\namespace\Dept

Folder Properties

Name Dept

Users connect to \\itso.tucson\namespace\Dept

Children

Name	Comment	Timeout (seconds)
Marketing		1800
Sales		1800

New Migration Task

Verify that destination paths are empty

Destination paths that are not empty will be marked with a warning sign. Clear the checkbox of items you do not want to migrate.

Original Path	Destination Path
☑ \\roman.itso.tucson\sales	⚠ \\itsotuc3.itso.tucson\sales

VFM

Warning: Some destination shares are not empty.

Destination Shares Status: 0 of 1 Destination Items Verified

Selected Destination Share: The folder is not empty.

Figure 8-52 Destination share not empty

4. After completing the creation of the migration tasks as shown in example 1 (see 8.5, “Migration policy example 1” on page 204) and example 2 (see 8.8, “Migration policy example 2” on page 233), start the migration phases

manually. On completion of the final phase, check the history (see Figure 8-53).

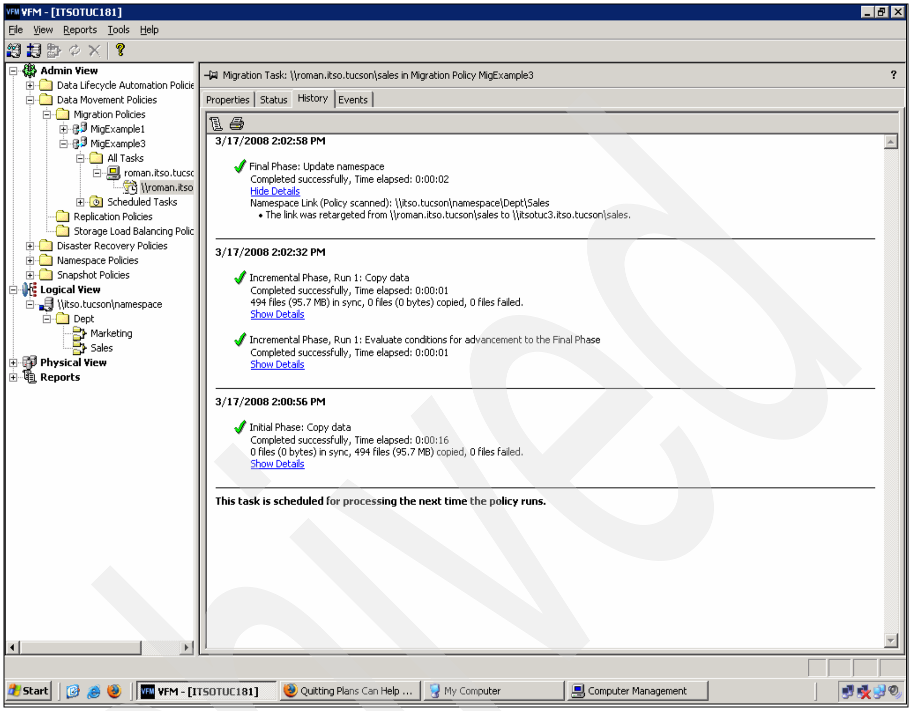

Migration Task: \\roman.itso.tucson\sales in Migration Policy MigExample3

3/17/2008 2:02:58 PM

✓ Final Phase: Update namespace
Completed successfully, Time elapsed: 0:00:02
Hide Details
Namespace Link (Policy scanned): \\itso.tucson\namespace\Dept\Sales
- The link was retargeted from \\roman.itso.tucson\sales to \\itsotuc3.itso.tucson\sales.

3/17/2008 2:02:32 PM

✓ Incremental Phase, Run 1: Copy data
Completed successfully, Time elapsed: 0:00:01
494 files (95.7 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

✓ Incremental Phase, Run 1: Evaluate conditions for advancement to the Final Phase
Completed successfully, Time elapsed: 0:00:01
Show Details

3/17/2008 2:00:56 PM

✓ Initial Phase: Copy data
Completed successfully, Time elapsed: 0:00:16
0 files (0 bytes) in sync, 494 files (95.7 MB) copied, 0 files failed.
Show Details

This task is scheduled for processing the next time the policy runs.

Figure 8-53 Sales share migration history

5. Create a new migration task for the Marketing folder. Move it to itsotuc4 N series (see Figure 8-54).

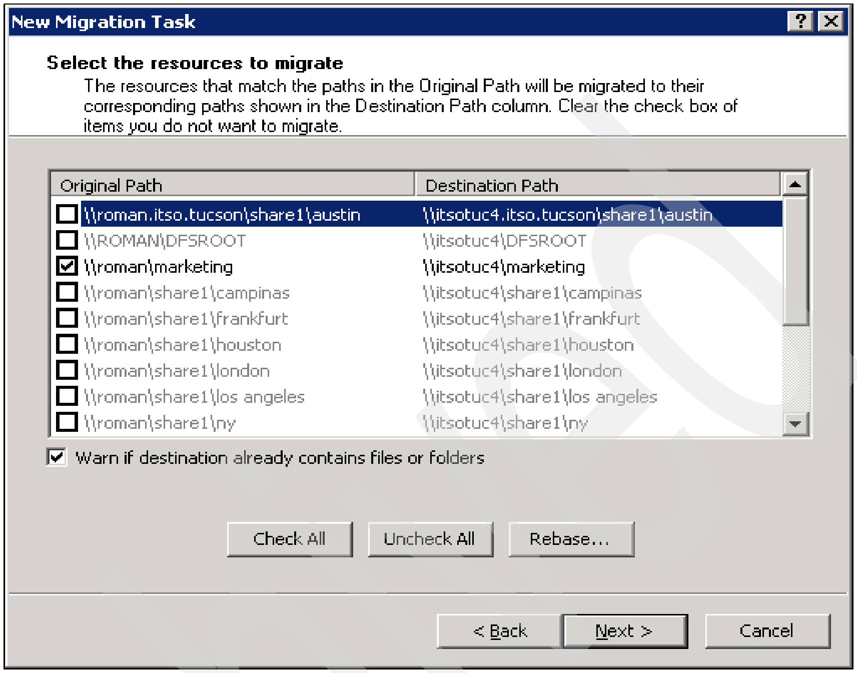

Select the resources to migrate

The resources that match the paths in the Original Path will be migrated to their corresponding paths shown in the Destination Path column. Clear the check box of items you do not want to migrate.

Original Path	Destination Path
□ \\roman.itso.tucson\share1\austin	\\itsotuc4.itso.tucson\share1\austin
□ \\ROMAN\DFSROOT	\\itsotuc4\DFSROOT
☑ \\roman\marketing	\\itsotuc4\marketing
□ \\roman\share1\campinas	\\itsotuc4\share1\campinas
□ \\roman\share1\frankfurt	\\itsotuc4\share1\frankfurt
□ \\roman\share1\houston	\\itsotuc4\share1\houston
□ \\roman\share1\london	\\itsotuc4\share1\london
□ \\roman\share1\los angeles	\\itsotuc4\share1\los angeles
□ \\roman\share1\ny	\\itsotuc4\share1\ny

☑ Warn if destination already contains files or folders

Figure 8-54 Migration task for Marketing share

6. After completing the task, invoke it manually.

7. Check the history of the migration phases for successful completion (see Figure 8-55).

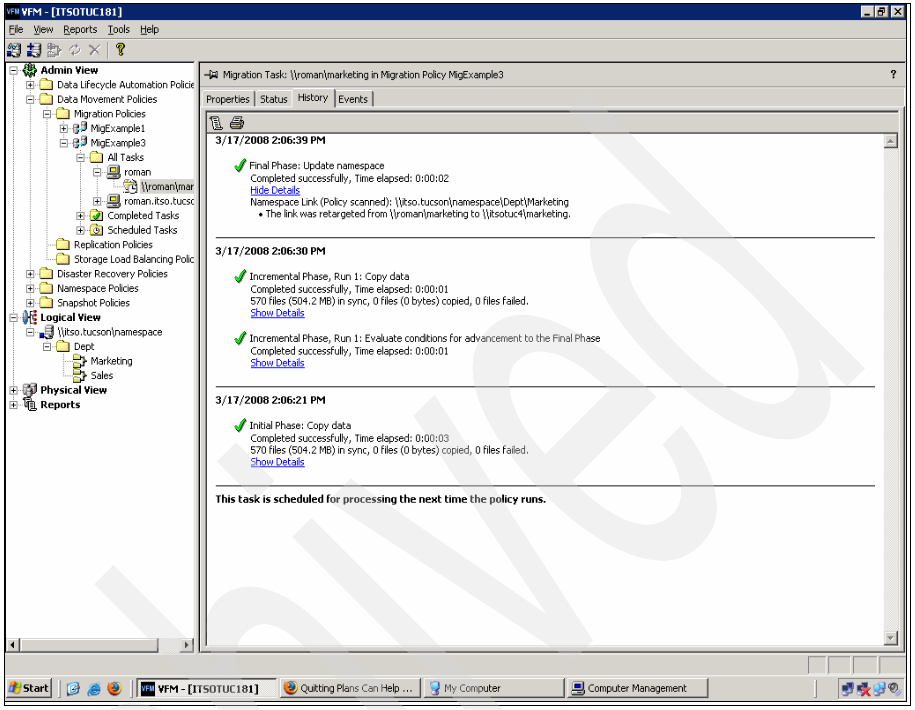

Migration Task: \\roman\marketing in Migration Policy MigExample3

3/17/2008 2:06:39 PM

✓ Final Phase: Update namespace
Completed successfully, Time elapsed: 0:00:02
Hide Details
Namespace Link (Policy scanned): \\itso.tucson\namespace\Dept\Marketing
- The link was retargeted from \\roman\marketing to \\itsotuc4\marketing.

3/17/2008 2:06:30 PM

✓ Incremental Phase, Run 1: Copy data
Completed successfully, Time elapsed: 0:00:01
570 files (504.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

✓ Incremental Phase, Run 1: Evaluate conditions for advancement to the Final Phase
Completed successfully, Time elapsed: 0:00:01
Show Details

3/17/2008 2:06:21 PM

✓ Initial Phase: Copy data
Completed successfully, Time elapsed: 0:00:03
570 files (504.2 MB) in sync, 0 files (0 bytes) copied, 0 files failed.
Show Details

This task is scheduled for processing the next time the policy runs.

Figure 8-55 History of Marketing share migration

8. Check the Dependency tab of the Dept folder under the DFS root namespace (see Figure 8-56).

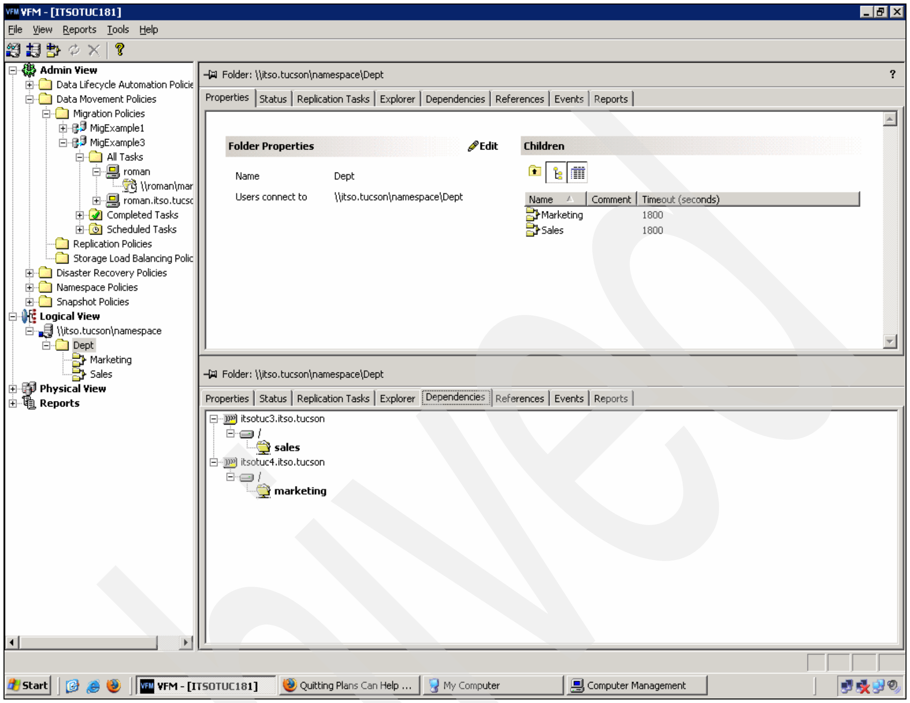

Name Dept

Users connect to \\itso.tucson\namespace\Dept

Name	Comment	Timeout (seconds)
Marketing		1800
Sales		1800

Figure 8-56 Shares moved to itsotuc3 and itsotuc4

9. Finally check whether the DFS links are retargeted on itsotuc3 and itsotuc4 (see Figure 8-57).

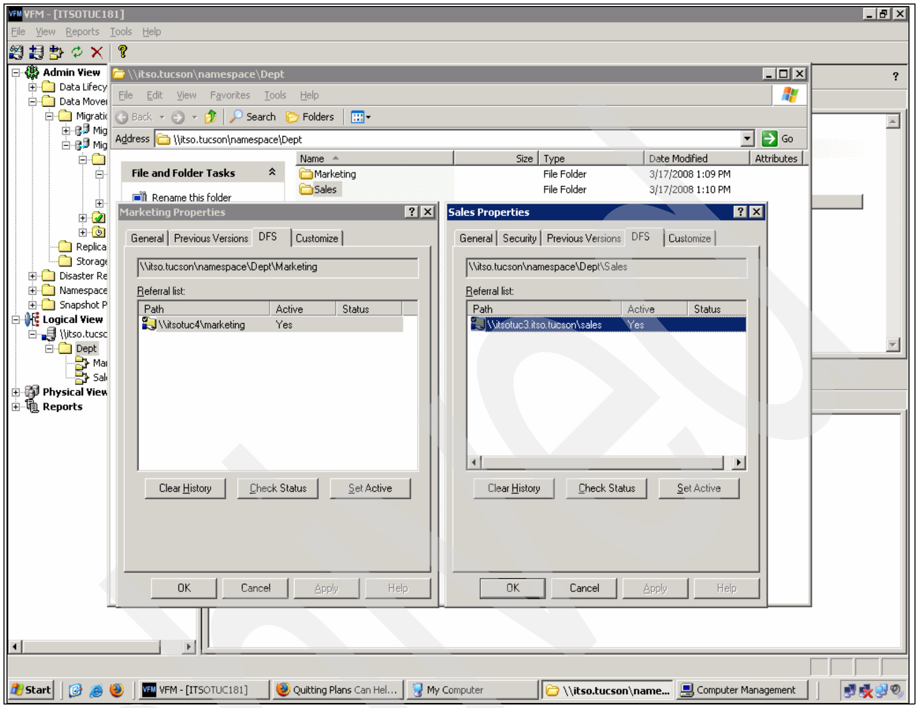

Address: \\itso.tucson\namespace\Dept

Name	Size	Type	Date Modified	Attributes
Marketing		File Folder	3/17/2008 1:09 PM	
Sales		File Folder	3/17/2008 1:10 PM	

Marketing Properties

\\itso.tucson\namespace\Dept\Marketing

Referral list:

Path	Active	Status
\\itsotuc4\marketing	Yes	

Sales Properties

\\itso.tucson\namespace\Dept\Sales

Referral list:

Path	Active	Status
\\itsotuc3.itso.tucson\sales	Yes	

Figure 8-57 DFS links

9

Chapter 9.

Life cycle management with VFM

VFM provides the capability to migrate old or unused data between heterogeneous systems. While VFM is a general-purpose data movement solution focused on moving unstructured file data, the policies can be used to move old data from one tier of storage to another. The archival migration policy enables you to move old data from one machine to another. You can also set up an archival migration policy to move data from one volume to another on the same machine.

For example, you can move infrequently used data from Fibre Channel storage to SATA storage. An archival migration policy can move CIFS data from one share to another. An archival migration policy can also move NFS data from one NFS export to another. A VFM Replication Agent is used to move the data from one place to another. The Replication Agent is automatically deployed by VFM to perform the data movement. No administrator intervention is required to deploy a Replication Agent in a default configuration.

The goal of any archival migration is to save space on the primary storage. In addition to saving space, you might want to save on backup costs where infrequently accessed or modified files are not backed up each week. The goal is to move data from one location to another without disrupting users. Users must be unaware that data has moved. They must maintain the same level of access before the migration and after the migration. Even though performance of the

SATA drives is not as efficient as the performance of the Fibre Channel drives, the data is still accessible to users, although with some latency.

The goal of an archival migration policy in VFM is to describe the conditions for the migration. You specify the source server and share, the destination server and share, and the controlling parameters of the policy. Once specified, you allow VFM to perform the actions of the policy based on the settings. When complete, you can review the policy to determine the status of the movement. You can also review the progress of the policy as it works through the various phases.

The following client use cases are described in this chapter:

- Move old project data from one server to another
- Move old project data from a Fibre Channel volume on an N series machine to a SATA volume on the same N series system storage
- Move disabled users' home drives to a less expensive tier of storage
- Move accounting data that has not been modified in three years to cheaper storage
- Provide a list of shares where the majority of data has not been changed within an extended period of time

9.1 Namespace

Archival migration is not just about moving data from one location to another or from one server to another. Data can be moved from one server to another in the same data center or across data centers. An archival migration policy requires data to be behind the namespace. The real benefit of VFM is the capability to move data transparently without disrupting the users' view of the data. The result of providing this level of transparency is that the user does not have to know the physical location of the data. Separating the user from the physical location of the data is achieved by the namespace. After performing an archival migration without a namespace, you might have to change logon scripts, send e-mails to the user, and complete other tasks. With a namespace in place, you can simply update the namespace on completion of the archival migration and seamlessly point the user to the new destination (see Figure 9-1).

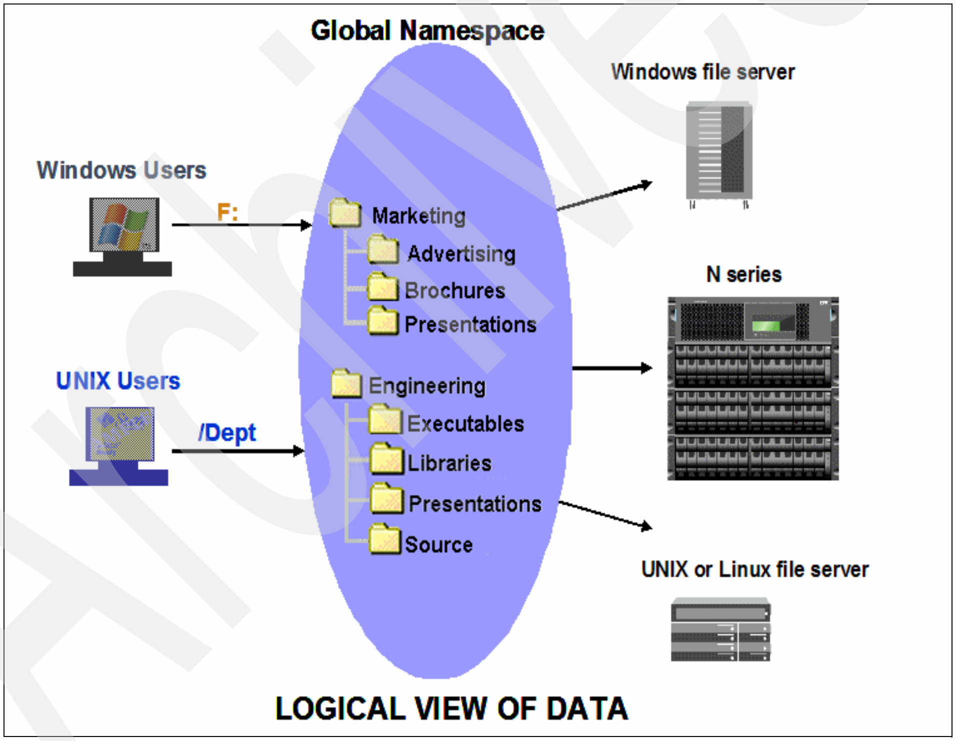

Figure 9-1 Namespace

In general, administrators have no detailed knowledge of the data in the shares of the file server. For example, you might never really know whether a project has finished or whether a particular share is no longer being actively used. Usually, only the data owner has such intimate knowledge of the data, but the user might not communicate that information to you. You continue to back up the data week after week. If you knew or could determine that certain storage areas were not frequently used, you could relocate the data and free up space. An archival migration policy is helpful in this situation.

9.2 Archival migration policy overview

An archival migration policy is broken up into the following steps:

1. Identify the candidate source shares to scan in order to determine which shares can be moved because they contain old or unused data.
2. Specify the intended sources and corresponding destinations.
3. Specify the migration criteria (the age) to determine which of the candidates shares satisfy the requirements to be moved.
4. Decide whether you want to automatically or manually add the migration candidates to a migration policy so that they can be moved.
5. Schedule and run the migration policy as you would any other migration policy.
6. Check the results of the migration and access to the moved data.

To obtain the maximum benefit of an archival migration policy, you must be familiar with the creation, setting, and running of a migration policy. The settings in a migration policy are a subset of the settings of an archival migration policy.

Note: It is important to understand the effects of archival migration. When users access data through the namespace, it has minimal effect on the archival migration policy. The files can be relocated transparently to the users. However, if some users are accessing data without going through the namespace (that is, users have direct access to the data, bypassing the namespace), then archiving the old data disrupts those users. Do not use an archival migration policy to move the data without knowing how users are accessing the data. (This situation is no different than a migration scenario. If you move a share from one location to another, user access to the share is disrupted if the user bypasses the namespace.)

9.3 Archival migration policy example 1

Suppose you have a DFS namespace called [\\namespace\root], and you have an Accounting folder in it (see Figure 9-2). This folder has several DFS links pointing to some data shares on a mixture of machines. Each DFS link points to accounting data for a specific year. For example, the DFS link 2005 points to the accounting share for year 2005. The DFS link 2006 points to the accounting share for year 2006, and so on. You want an archival migration policy to move the data if over 90% of the data has not been modified in the last year. If data is to be moved, it is sent to the cheaper storage without users knowing about it.

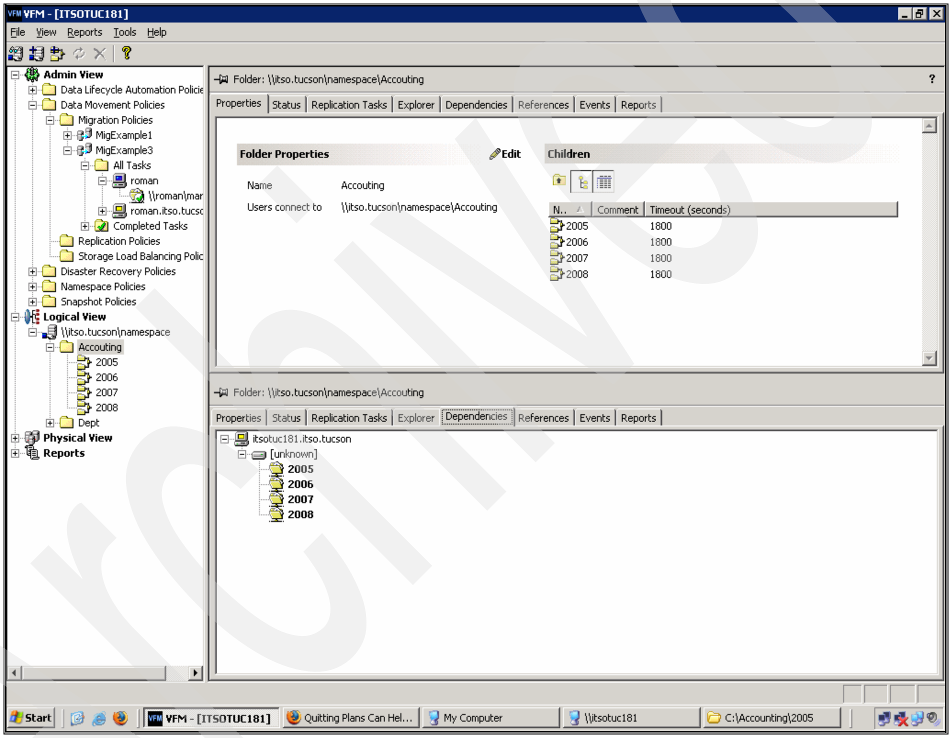

Folder: \\itso.tucson\namespace\Accouting

Folder Properties

Name Accouting

Users connect to \\itso.tucson\namespace\Accouting

Children

N..	Comment	Timeout (seconds)
2005		1800
2006		1800
2007		1800
2008		1800

Folder: \\itso.tucson\namespace\Accouting

itsotuc181.itso.tucson
[unknown]
2005
2006
2007
2008

Figure 9-2 Example 1 shares

> **Note:** The link 2005 is targeted to the itsotuc181 machine (see Figure 9-3 on page 270).

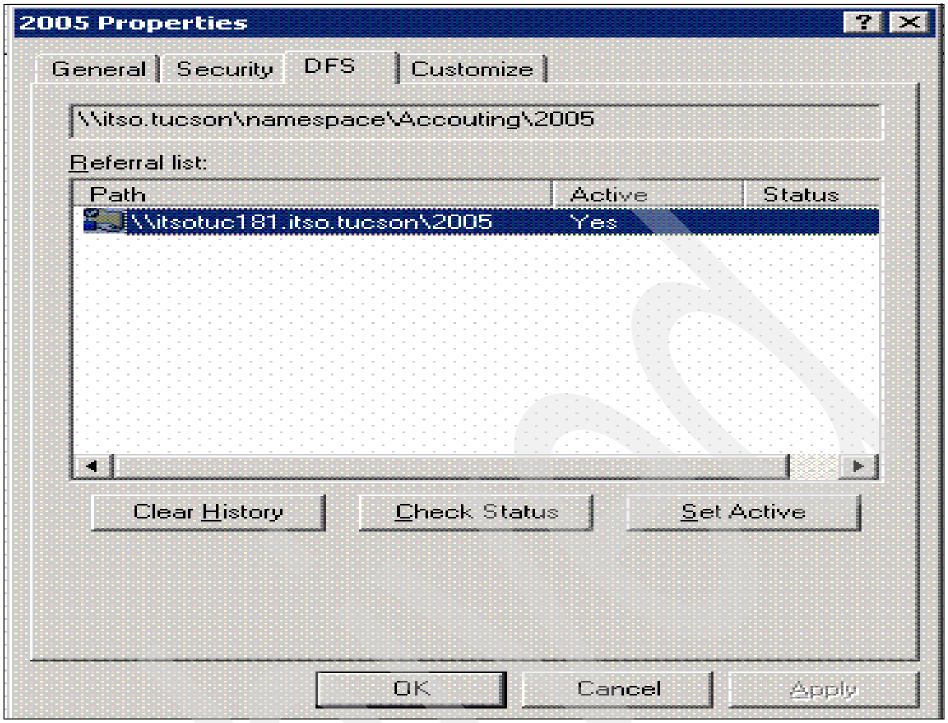

\\itso.tucson\namespace\Accouting\2005

Referral list:

Path	Active	Status
\\itsotuc181.itso.tucson\2005	Yes	

Figure 9-3 Example 1 DFS link

9.3.1 Archival migration policy general window

In the general window (see Figure 9-4 on page 271 and Figure 9-5 on page 272), you specify the name for the archival migration policy. The main purpose of this window is to provide the selection paths for the sources and destinations for the archival migration policy.

The source paths are the areas to be scanned. The source path can be specified as a logical path (a DFS path that has one or more links that make up the namespace). For example, you can specify an entire DFS root as the source path. In this case, all DFS links belonging to that root are processed to determine whether they are suitable targets for the archival migration. Or you can specify a logical folder under the DFS root as a source path. In this case, all DFS links under the logical folder are processed to determine whether they are an appropriate target for the archival migration. The broader the scope, the more source paths that are evaluated as possible migration candidates.

It is best to select a source path that meets your business objectives. Select a small subset of the data to be moved. Run the policy based on that subset. Evaluate the results and then expand the scope.

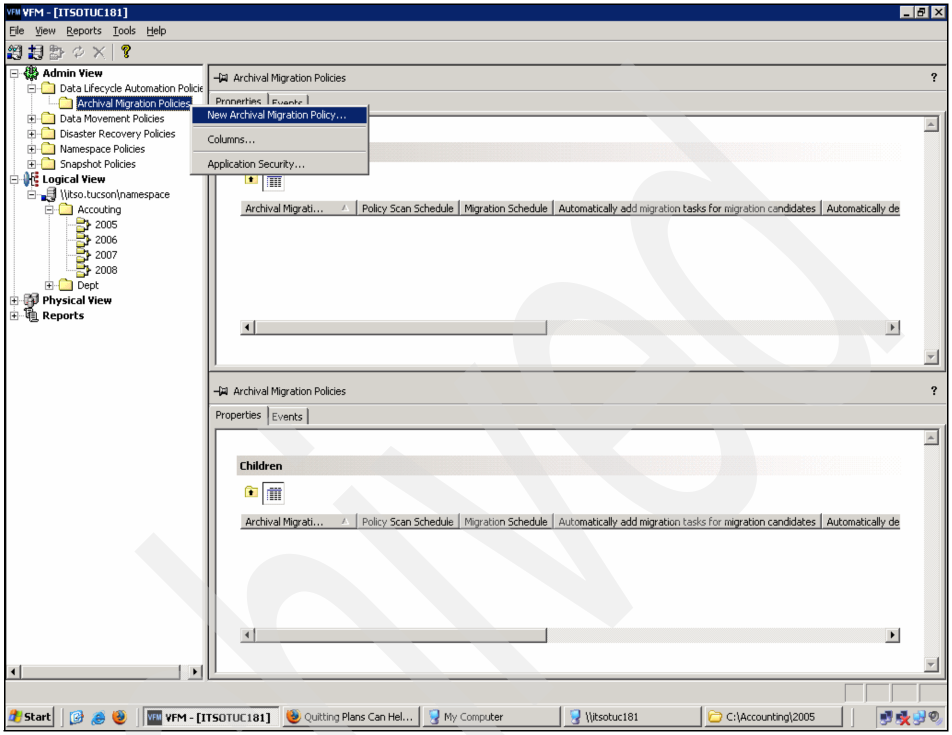

Figure 9-4 Creating a new archival migration policy

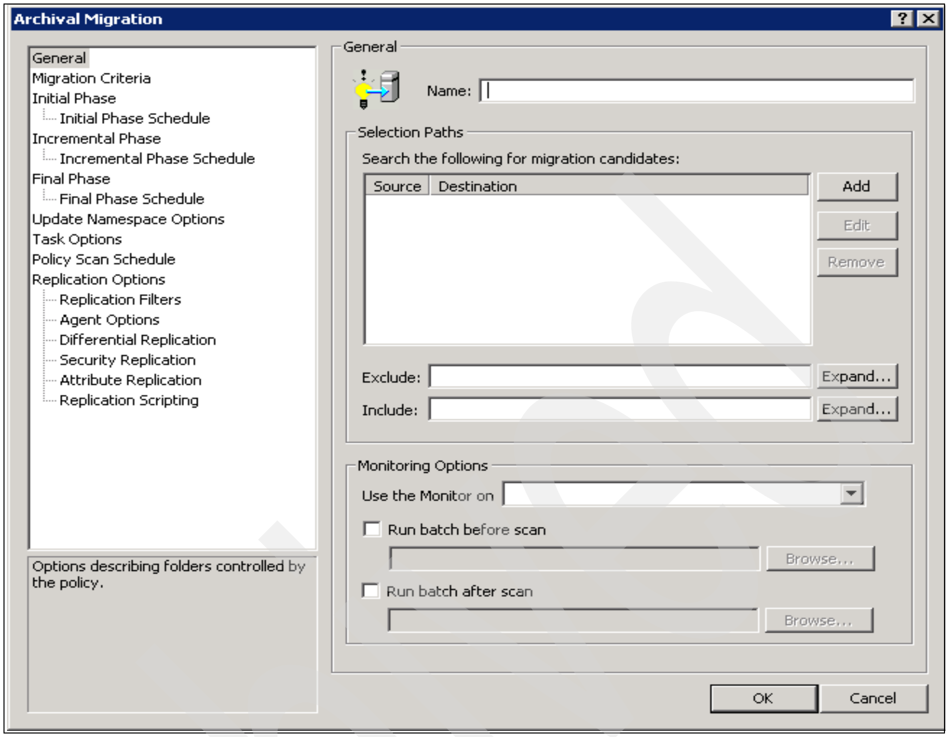

Figure 9-5 Archival Migration: General window

The destination path for the selection path must be a UNC path to a backend storage device. Make sure that the storage device has enough space to accommodate the data being migrated. When VFM moves the data, it is deleted from the source and moved to the destination.

A selection path can have multiple source and destination pairs. For example, you can move old accounting data from Server 1 to N series machine A. You can move old finance data from Server 2 to N series machine B. You can also move multiple sources to the same destination. You cannot specify the same source path to be moved to multiple destinations.

9.3.2 Creating an archival migration policy

Follow the steps in this section to create an archival migration policy:

1. If the destination folder does not exist, create the destination folder (see Figure 9-6).

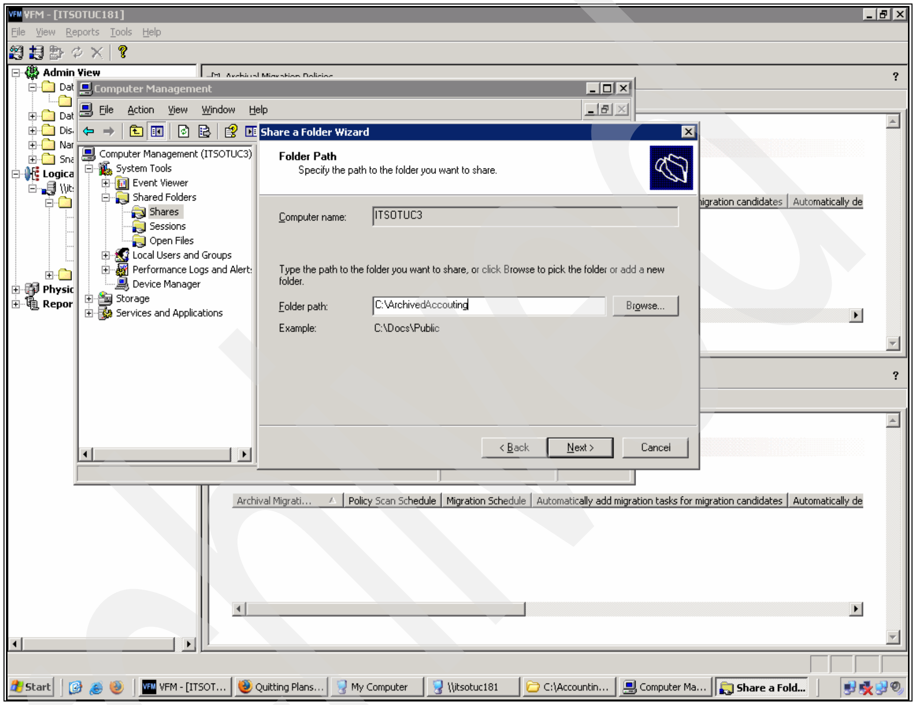

Folder Path

Specify the path to the folder you want to share.

Computer name: ITSOTUC3

Type the path to the folder you want to share, or click Browse to pick the folder or add a new folder.

Folder path: C:\ArchivedAccouting

Example: C:\Docs\Public

Figure 9-6 Creating destination folder

2. Grant full control access for the all users who require access to it (see Figure 9-7).

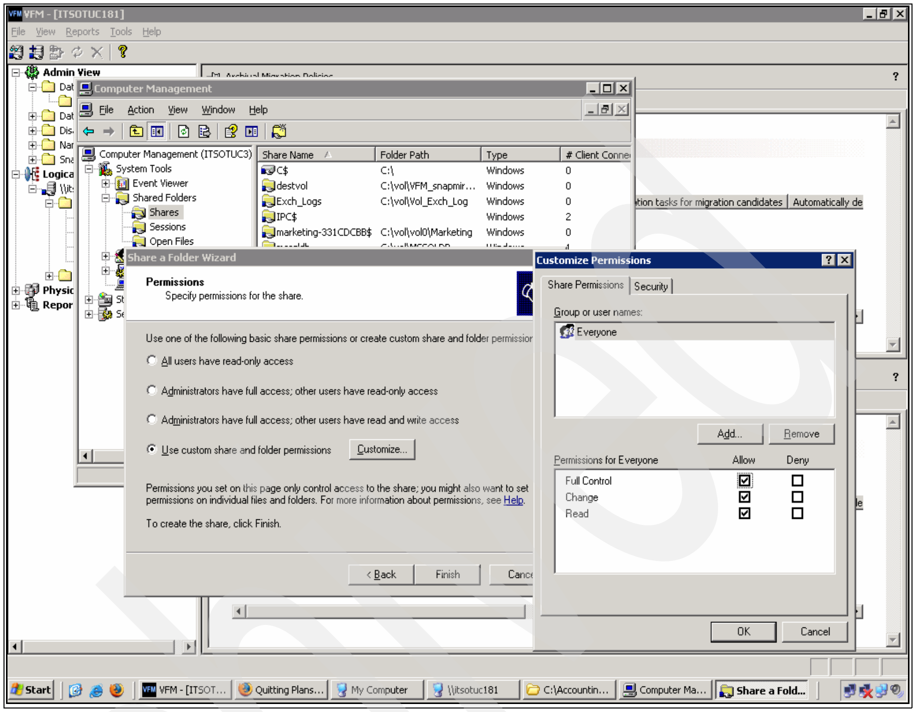

Share Name	Folder Path	Type	# Client Conne
C$	C:\	Windows	0
destvol	C:\vol\VFM_snapmir...	Windows	0
Exch_Logs	C:\vol\Vol_Exch_Log	Windows	0
IPC$		Windows	2
marketing-331CDCBB$	C:\vol\vol0\Marketing	Windows	0

Share a Folder Wizard

Permissions
Specify permissions for the share.

Use one of the following basic share permissions or create custom share and folder permissio

○ All users have read-only access

○ Administrators have full access; other users have read-only access

○ Administrators have full access; other users have read and write access

● Use custom share and folder permissions Customize...

Permissions you set on this page only control access to the share; you might also want to set permissions on individual files and folders. For more information about permissions, see Help.

To create the share, click Finish.

< Back Finish Cance

Customize Permissions

Share Permissions | Security

Group or user names:
Everyone

Add... Remove

Permissions for Everyone	Allow	Deny
Full Control	☑	□
Change	☑	□
Read	☑	□

OK Cancel

Figure 9-7 Granting access to the destination folder

3. Create a new archival migration policy. Specify a source and destination path.

 Specify a source path of \\namespace\root\accounting, and specify a destination path of \\nseries\archivedaccounting (see Figure 9-8 on page 275). If VFM encounters old data under the source path, that data becomes a candidate for migration that might be relocated to the archivedaccounting share.

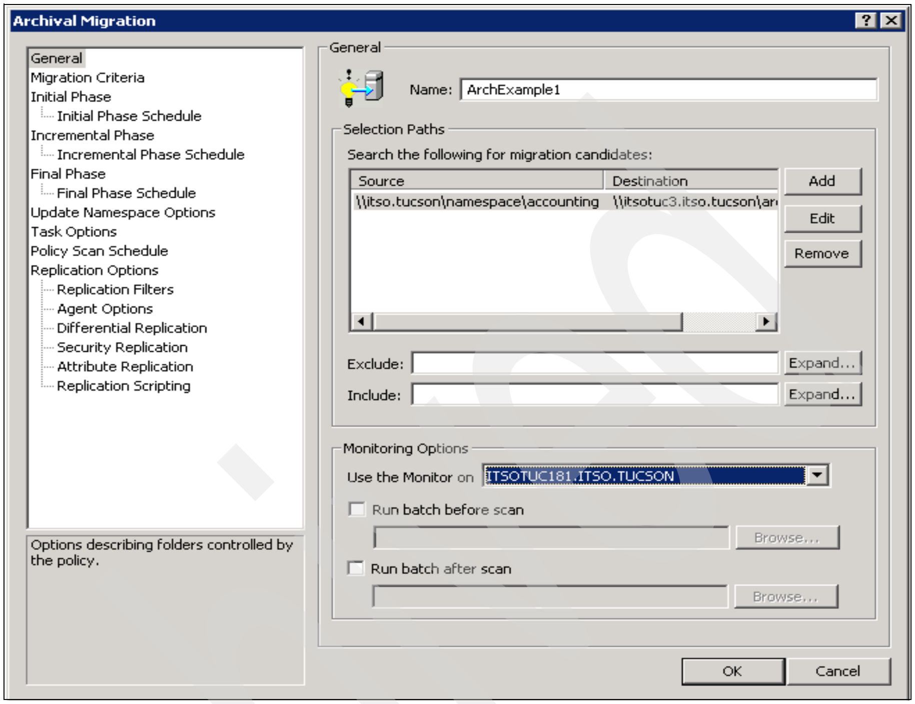

Options describing folders controlled by the policy.

Name: ArchExample1

Search the following for migration candidates:

Source	Destination
\\itso.tucson\namespace\accounting	\\itsotuc3.itso.tucson\ar

Exclude:

Include:

Use the Monitor on ITSOTUC181.ITSO.TUCSON

□ Run batch before scan

□ Run batch after scan

Figure 9-8 Example 1: specifying source and destination path

You can also specify batch scripts to run before or after the scan phase. These batch scripts are run by the VFM server under the credentials of the service account. The scripts can be specified as a local path or a UNC path. Make sure that the service account has the rights to run the script.

9.3.3 Migration criteria

The migration criteria distinguish an archival migration policy from a migration policy. For example, the simplest migration criterion is when you choose the option Always Migrate, which treats the archival migration policy as a migration policy (see Figure 9-9 on page 277).

You can have VFM move entire directory trees when a certain condition is satisfied. If the condition is satisfied and you have selected to move the data, VFM moves the data, and the namespace link is updated to point to the new location. A migration criterion is an ANDing of all the selected options. For example, you can have a criteria that says “if 90% of the files have not been modified in 6 months, then move the data.”

Remember, although some files can still be accessed, the entire directory is moved. Even the accessed files are relocated and accessed from the new location. This characteristic differentiates the archival migration policy from traditional Hierarchical Storage Management (HSM) software. Traditional HSM solutions move old, unused, or infrequently used files. The VFM archival migration moves entire folders even if some files are accessed.

It is best to keep your migration criteria simple. You do not want to be in a situation where you have migrated files based on the wrong criteria. For example, you might have migrated files based on the modified date on the file. However, the files in that share are constantly being accessed by users and applications even though the files have not changed. In this case, when you move the files from that share to the less expensive tier, the performance degrades to the detriment of users accessing the data.

It is best not to select a criterion where 100% of the files determine the criterion. For example, you can have a policy with this criterion: “100% of the files have not been modified in 6 months.” In this case, you might have 1 TB of data and 1 million files. If only three files have been modified in the last six months, this share is not selected for migration.

Use the archival migration policy to tell you which shares can be moved. It is a good idea to try “what-if” scenarios to better understand your data. Once you understand the data, you are in a better position to make informed decisions.

Migration criteria are flexible and powerful. Advanced administrators can determine their own criteria. For example, you might want a criterion to exclude temp files and only move data if no file is owned by a specific user. Any complicated criterion is supported by the “user selection criteria” script, which you can write. If the script returns 0, that source path is a candidate for migration; if the script returns 1, the source path is not a candidate for migration.

The user selection criteria script is run by the VFM server. The script can be specified using a local path or a remote UNC path. Make sure that the VFM service account has sufficient rights to look at the source data to decide the criteria.

9.3.4 Specifying migration criteria for archival migration policy

In our example, we specify the migration criteria by checking the box that specifies you move directories when “90% of the files have not been modified in the last year” (see Figure 9-9).

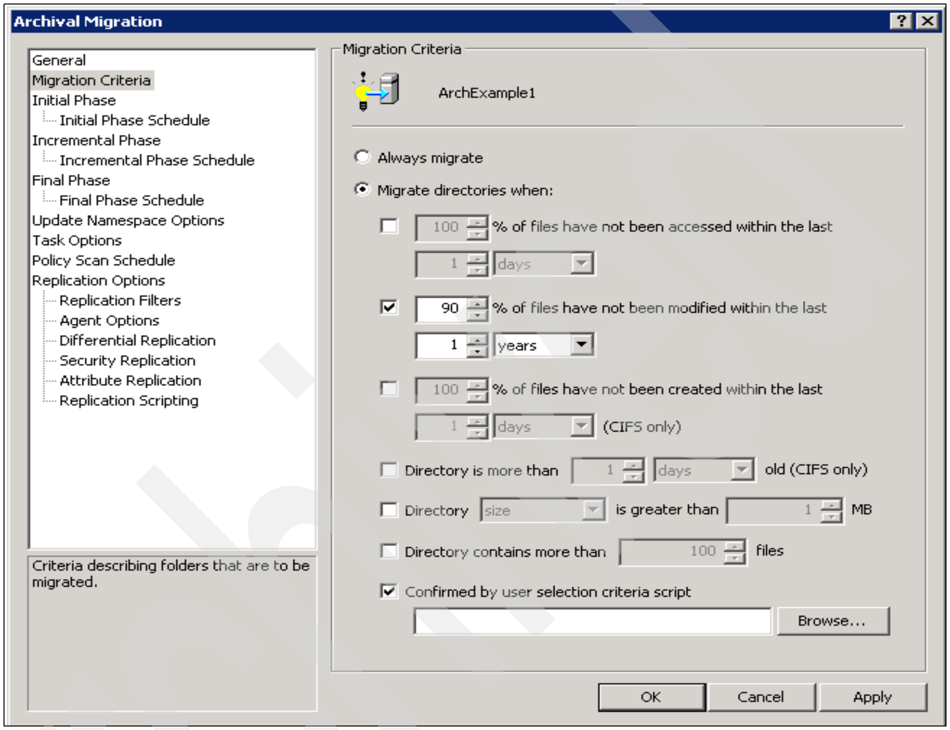
Archival Migration

General
Migration Criteria
Initial Phase
- Initial Phase Schedule

Incremental Phase
- Incremental Phase Schedule

Final Phase
- Final Phase Schedule

Update Namespace Options
Task Options
Policy Scan Schedule
Replication Options
- Replication Filters
- Agent Options
- Differential Replication
- Security Replication
- Attribute Replication
- Replication Scripting

Criteria describing folders that are to be migrated.

Migration Criteria

ArchExample1

○ Always migrate
● Migrate directories when:
□ 100 % of files have not been accessed within the last 1 days
☑ 90 % of files have not been modified within the last 1 years
□ 100 % of files have not been created within the last 1 days (CIFS only)
□ Directory is more than 1 days old (CIFS only)
□ Directory size is greater than 1 MB
□ Directory contains more than 100 files
☑ Confirmed by user selection criteria script

Figure 9-9 Migration Criteria panel

9.3.5 Migration phases of an archival migration policy

A migration policy has the following three phases:

- Initial phase
- Incremental phase
- Final phase

Each phase runs on its own schedule, which enables you to schedule the actions of a policy at different times. It is possible to complete all three phases (initial, incremental, and final) in a policy or skip any phase of the policy based on what is best for the environment.

Each phase has a set of preselected actions that can be performed in that phase. If an action is selected, it is performed in that phase. Multiple actions can be selected in a phase. Each selected action is performed in order.

The initial phase makes a copy of the source data on the destination machine. It is best to schedule the baseline to run over a weekend. Based on the amount of data to be copied, the baseline migration might take hours or days to complete. It is possible to start the baseline copy days or weeks before the cutover.

Once the baseline is complete, the incremental phase synchronizes the source and destination. The incremental phase can be run nightly or according to the schedule you set. Any new files created on the source are also copied.

You might want to run the final phase cutover on the weekend so that the last incremental copy is run before the actual cutover. During the cutover, the namespace can be updated and a final sync can be performed before users are brought online at the new destination. The time required for the final sync is determined by the amount of changed data to be copied and the time required to copy the data. This can take minutes on small data sets or hours on large data sets. Once the final sync is complete, the data is deleted from the source.

The policy settings have a great deal of flexibility, providing a powerful way to configure the policy to perform automated data movement.

9.3.6 Policies and tasks

A policy specifies the overall criteria for moving data. A policy specifies the schedule, the data movement options, and how the migration behaves. A policy is made up of tasks. Each task in a policy specifies a source and destination path. A task inherits the settings of the policy. The schedules for the policy specify when a phase runs and the actions to be performed. The actions are performed on the tasks. Some tasks can be in one phase while other tasks in the same policy could be in a different phase. A task can advance either automatically or manually from one phase to another. When a task completes the final phase, it is marked as complete.

9.3.7 Initial phase

A migration policy can specify a set of operations to be carried out in the Initial phase (see Figure 9-10 on page 280). You can run a batch file in the initial phase. For example, in the batch file, you can stop an application or send an e-mail. The batch file script is passed parameters as described in the Help text. The scripts are run by the VFM server under the context of the VFM service account. The scripts can be located on the local drive or from any network accessible share. A script runs once for each task in the policy.

VFM can also scan for potential problems. In this case, the server does a cursory examination of the files. It looks for locked files, checks to determine whether the server has permissions to the data, whether it knows the destination of a volume type, and other tasks. This step is not exhaustive. For example, a disk on the destination might contain sufficient space at this time. However, when it is time to copy the data, sufficient space might be lacking.

You can also choose to perform a baseline copy of the data. All files on the source are copied to the destination. If any files are already on the destination, the action on the destination is determined by the settings of the policy. For example, if you have checked Delete orphaned files on the destination check box, the existing files on the destination are cleared. If you select to delete the orphaned files, the Replication Agent makes the destination a mirror image of the source. Once the baseline is complete, the migration policy remains in this phase until it is time to move to the incremental phase.

For the purposes of this example, we select to perform a baseline copy (see Figure 9-10).

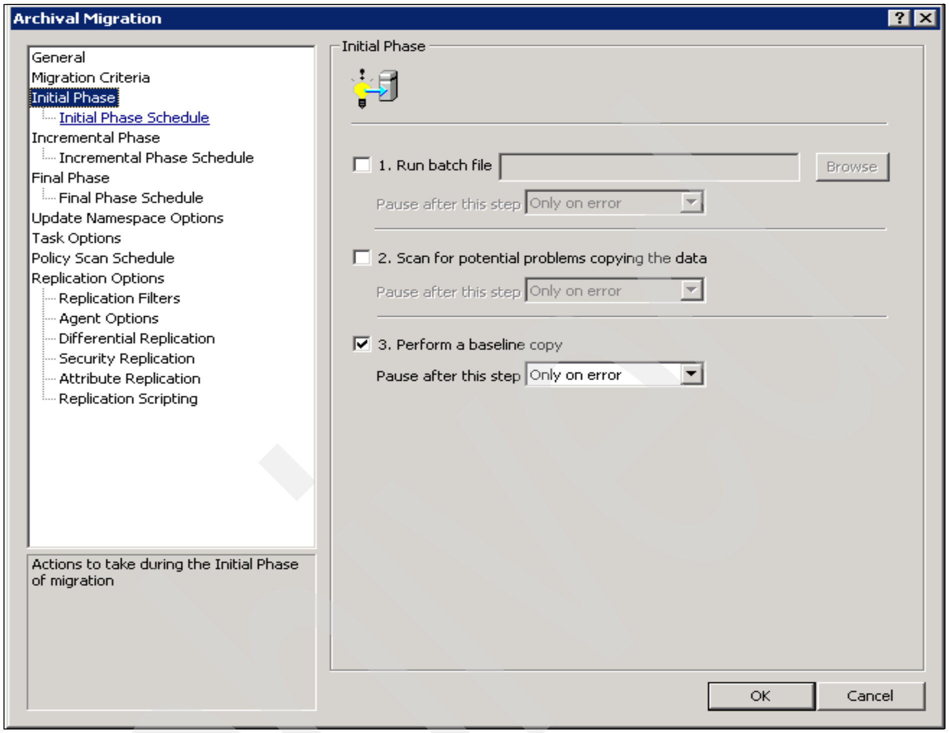

Actions to take during the Initial Phase of migration

□ 1. Run batch file

Pause after this step Only on error

□ 2. Scan for potential problems copying the data

Pause after this step Only on error

☑ 3. Perform a baseline copy

Pause after this step Only on error

Figure 9-10 Archival Migration: Initial Phase panel

9.3.8 Initial phase schedule

The initial phase schedule specifies when the initial phase starts. The schedule includes only a start time. The job runs only after the start time passes or the job is manually advanced. If the initial phase is past, the job does not start and has to be manually run. The job stays in the initial phase until it is time to move to the incremental phase.

9.3.9 Incremental phase

This phase is provides a mechanism to make repeated copies of the data. As the data changes on the source, the data is copied to the destination at the next invocation of the policy's scheduled run. The incremental phase schedule specifies when the incremental phase finishes and how the job advances to the final phase.

If you clear the Copy the Data check box, the incremental phase is disabled. For the purpose of this example, we leave the incremental phase as the default, that is, the Copy the Data check box is cleared.

9.3.10 Incremental phase schedule

An incremental phase schedule in a migration policy determines the frequency of migration. In some instances, it might be satisfactory to copy data once a night.

Do not schedule the incremental migration to occur more frequently than the time required to complete a run of the incremental phase. For example, if an incremental run takes one hour to complete, do not schedule the frequency to run every 30 minutes. If a policy is running and the next scheduled run of the policy elapses, the policy waits until the current run of the policy finishes. If a policy is running, the next run of the policy does not start until the current run finishes.

The VFM console does not have to be open when the policy runs. The VFM server initiates the policy. It is run based on the time settings and schedule of the VFM server, which might or might not be in the same time zone as the source or destination machines. Take care when setting the run times of a policy and take into account the source and destination time zones and the location of the VFM server.

Once the incremental phase finishes, the policy does not advance to the final phase until the time of the final phase elapses.

9.3.11 Final phase

The final phase specifies the actions to be performed when you are almost ready to cut over. This phase assumes that the source and destination are almost in sync and that little data is to be copied in this phase. When the start time of the final phase passes, the actions selected in the final phase are performed. The selected actions are performed in order. You can choose to pause at the end of each step or continue until the end (see Figure 9-11 on page 284).

Note: Make sure that you run the final phase when no users have open files on the source. If open files are on the source, the files cannot be copied correctly, or the user is forcibly disconnected from the files.

Setting the Update the namespace option (see Figure 9-11 on page 284) is the most important step of the final phase. Any DFS links targeting the source share or folder are retargeted to point to the destination share or folder. This makes the migration transparent to the user. The user does not have to update drive mappings if they are being accessed through a DFS link. VFM automatically updates the link target, which causes the user to be redirected to the destination for the archived data.

The Prevent user connections during the final copy option enables VFM to copy any open files. If a user has an open file, VFM changes the share name, which is disruptive to the user and forcibly disconnects the user from the source share. This enables VFM to copy the files without having any open files to worry about. VFM hides the source share by adding a few random characters to the share name and appending a dollar sign ($) to hide it. VFM then uses the newly created share as the source and moves the remaining data. See 9.6, “Archival migration policy example 4” on page 307.

The Do a final replication option enables you to copy the last remaining files and make the source and destination alike. The time required to do the final replication is based on the amount of data to synchronize, the number of files already in sync, the number of files to copy, the network speeds available, and other factors.

The Delete the source data option enables you to free up the space on the source server. This is the default operation on completion of a successful migration. Take care when selecting this option if you want to preserve the source data.

The Stop sharing the source or export option deletes the source share and no longer advertises it in the network. Users can no longer access the source share. If you discontinue the use of the physical share path, you can detect the users that bypass the namespace. You can then encourage the users and applications to access the shares through the namespace.

The Run batch file option at the end of migration enables you to send e-mail to a set of users at the end of the migration. You can also use this option to restart an application from the new location.

9.3.12 Final phase schedule

The final phase schedule in a migration policy determines when the cutover is initiated. It is best to schedule the final phase when no users are accessing the data. We recommend you tell users to disconnect from the source share when the final cutover is scheduled. The VFM console does not have to be open when the final phase of the policy runs.

The policy is initiated by the VFM server. The policy is run based on the time settings and schedule of the VFM server, which might or might not be in the same time zone of the source or destination machines. Take care when setting the run times of a policy and take into account the source and destination time zones and the location of the VFM server. When the final phase time elapses, the tasks move through each selected action of the final phase.

9.3.13 Archival migration final phase

For the purpose of this example, you select the options to enable the following actions (see Figure 9-11 on page 284):

- Update the namespace and retarget the links
- Prevent user connections during the final copy
- Do a final replication
- Delete the source data
- Stop sharing the source or export

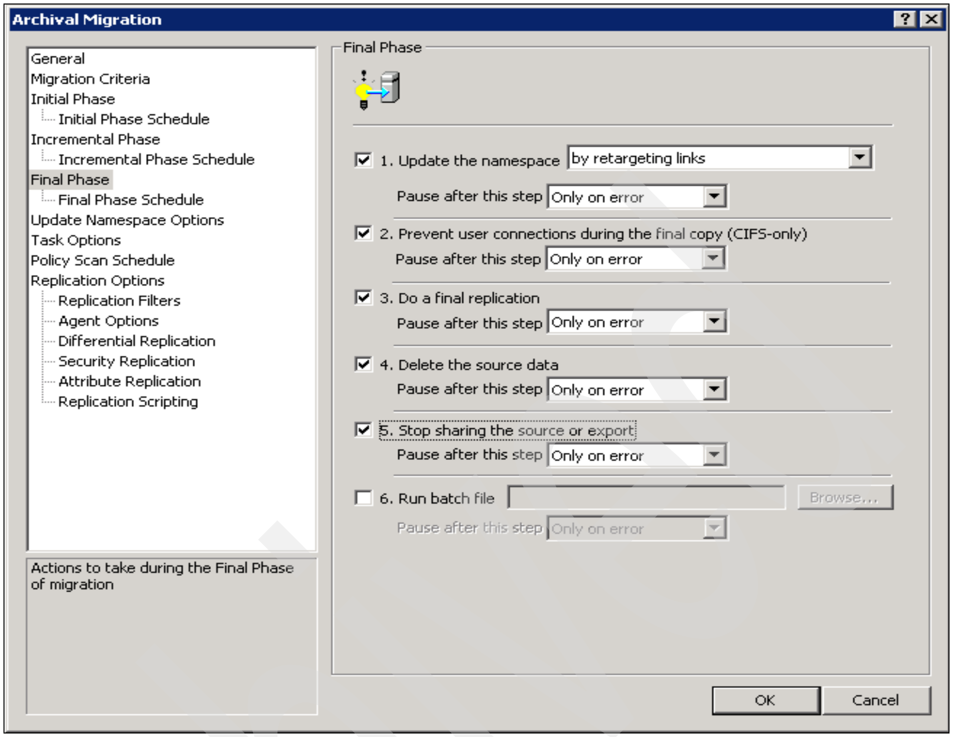

Actions to take during the Final Phase of migration

Final Phase

☑ 1. Update the namespace: by retargeting links
Pause after this step: Only on error

☑ 2. Prevent user connections during the final copy (CIFS-only)
Pause after this step: Only on error

☑ 3. Do a final replication
Pause after this step: Only on error

☑ 4. Delete the source data
Pause after this step: Only on error

☑ 5. Stop sharing the source or export
Pause after this step: Only on error

☐ 6. Run batch file
Pause after this step: Only on error

Figure 9-11 Archival Migration: Final Phase panel

9.3.14 Updating namespace options

When VFM performs an archival migration operation, the goal is to perform these operations so they are transparent to the user. The maximum amount of transparency is achieved when Microsoft Windows users access the archived data through the namespace. The Update namespace option specifies the Windows namespaces (DFS root) to search to update the DFS links.

Each DFS root listed in the policy determines whether any DFS links are referencing the source share. These DFS links are automatically updated if the final phase check box of Update the namespace is selected. If VFM does not know about some DFS roots, DFS links in those namespaces are not updated. Take care to ensure that all the DFS roots listed are the ones you want to update and that you have the rights to update them (see Figure 9-12 on page 285).

If you are migrating NFS data, make sure that the automount files referencing the source paths are specified. VFM updates these automount files to replace references to the old export with the new export.

For the purpose of this example, you select the default options on the window shown in Figure 9-12. Make sure the namespace under consideration for this example is listed in the Namespace list.

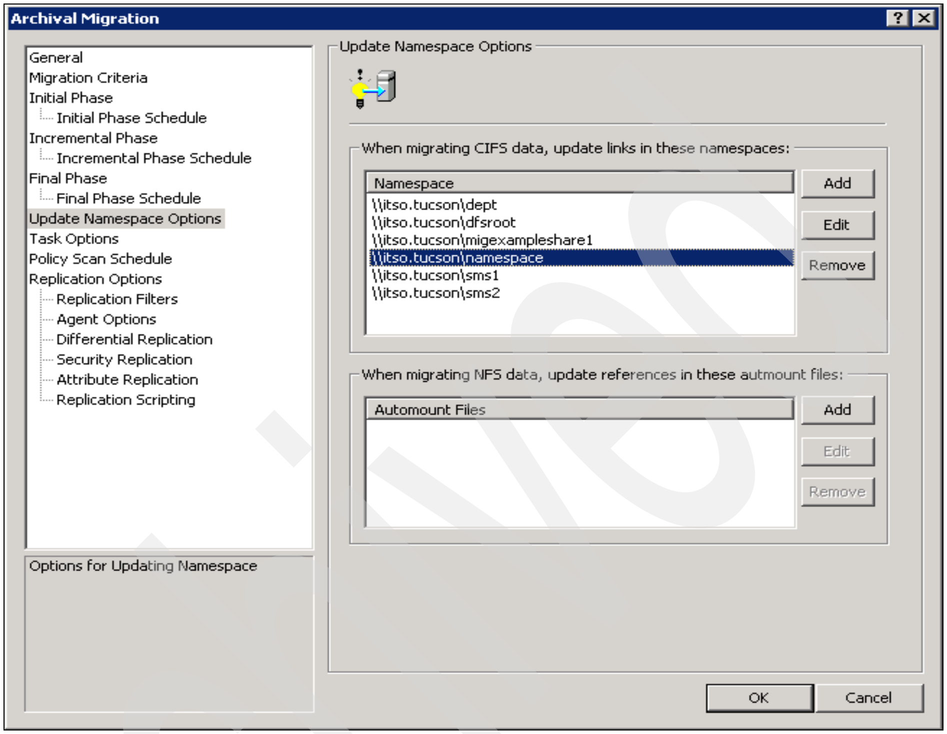

When migrating CIFS data, update links in these namespaces:

Namespace
\\itso.tucson\dept
\\itso.tucson\dfsroot
\\itso.tucson\migexampleshare1
\\itso.tucson\namespace
\\itso.tucson\sms1
\\itso.tucson\sms2

When migrating NFS data, update references in these autmount files:

Figure 9-12 Updating Namespace Options panel

9.3.15 Specifying task options

At this point, you have provided the search paths for the source and destination. You have specified the search criteria that help to determine the candidates for migration. You now have to specify what occurs when the search successfully finds candidates to be moved (see Figure 9-13 on page 287).

To add migration candidates as tasks, you can select to automatically or manually add migration tasks for the migration candidates (see the Automatically add migration tasks for migration candidate option in Figure 9-13 on page 287). It is best to manually add the migration tasks to the policy instead of automatically adding them. VFM might find candidates, and you cannot tell whether those candidates can be migrated without consulting the data owner. If you select to manually add the migration candidates, they show up on the migration candidates so that you can determine for which candidates you actually create tasks. For each candidate that you want to create a task, you right-click and create the task. You can multiselect migration candidates to be added as tasks.

If too many or too few migration candidates are shown, you might want to modify your scan criteria or your search paths. Either the criteria might be too stringent or too loose.

You also have the choice of deleting tasks that are completed (see the Automatically delete complete tasks option in Figure 9-13 on page 287). Either they can be automatically deleted, or you can manually delete them. You must always select to manually delete the tasks. If you automatically delete the tasks, you do not have a history of which tasks were moved and when. Tasks are only automatically deleted if they complete successfully.

The Only migrate data referenced by the namespace option (see Figure 9-13 on page 287) enables you to pare down the migration candidates to those behind the namespace. If no namespace link points to the source, that share is not selected as a candidate for migration.

We recommend you select the Only migrate data referenced by the namespace check box. Use caution when selecting a migration candidate to be added as a task. Even though only those shares referenced by a namespace are selected, you must make sure that no one is accessing the shares and bypassing the namespace. Otherwise, these users are disrupted.

For this example, clear the **Automatically add migration tasks for migration candidate** check box. Clear the **Automatically delete completed tasks** check box. Finally, check the **Only migrate data referenced through the namespace** check box (see Figure 9-13).

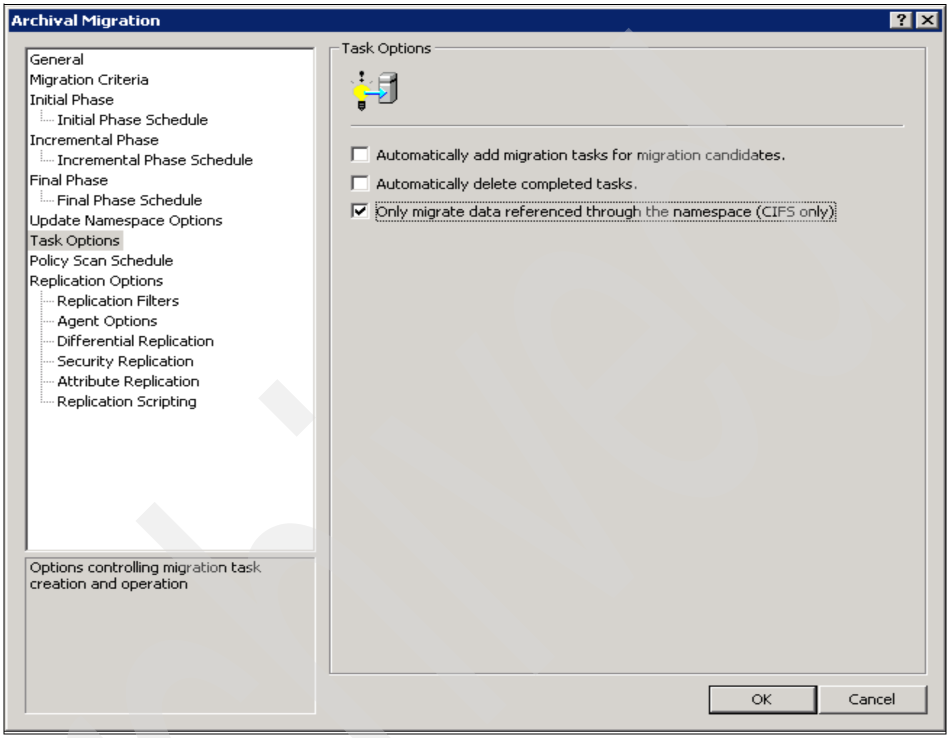

□ Automatically add migration tasks for migration candidates.

□ Automatically delete completed tasks.

☑ Only migrate data referenced through the namespace (CIFS only)

Figure 9-13 Archival Migration: Task Options panel

9.3.16 Specifying a policy scan schedule

An archival migration policy contains a number of schedules (see Figure 9-14 on page 288). All schedules (initial phase schedule, incremental phase schedule, and final phase schedule) deal with copying data. These schedules deal with data movement.

The policy scan schedule specifies when source paths are scanned to search for candidates that satisfy the migration criteria. The scan times can be different than the migration times. Generally, you want to scan the data during off hours. You want to look at the results to determine whether you need to refine the scan criteria and then redo the scan.

You can run the scan many times. Scans can be run manually or automatically. For example, you can run the scan once a week to identify the migration candidates. You might or might not choose to take action based on a scan.

In general, the same rules for the incremental phase schedule apply to the policy scan schedule.

For this example, uncheck the **Use policy scan schedule** check box. Instead of running the scan based on a schedule, you manually perform the scan (see Figure 9-14).

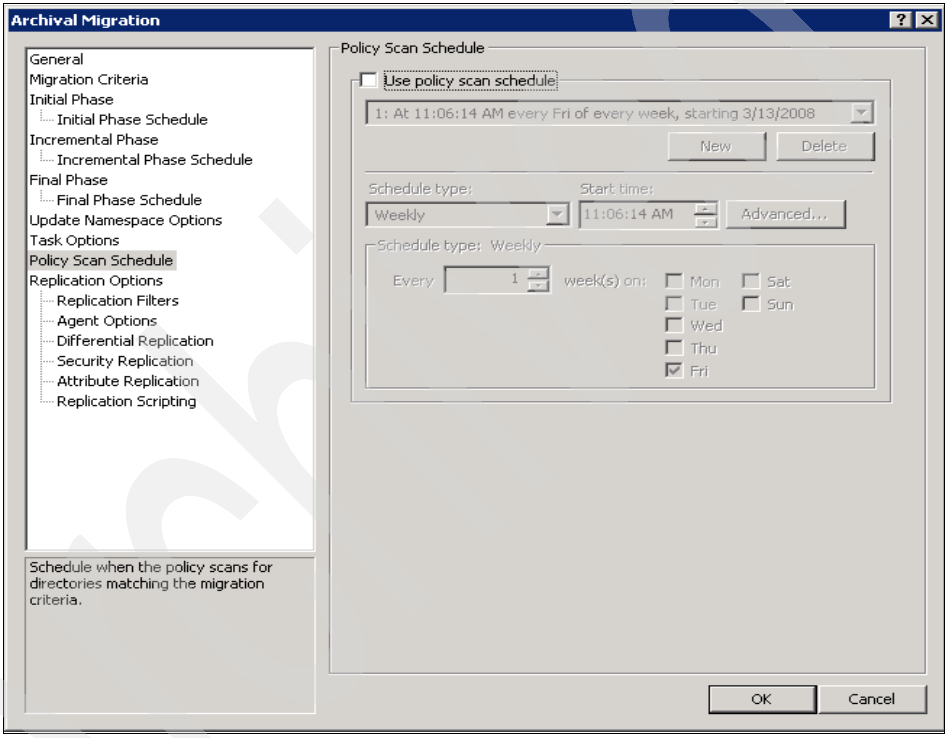

Schedule when the policy scans for directories matching the migration criteria.

□ Use policy scan schedule

1: At 11:06:14 AM every Fri of every week, starting 3/13/2008

Schedule type: Weekly

Start time: 11:06:14 AM

Schedule type: Weekly

Every 1 week(s) on: □ Mon □ Tue □ Wed □ Thu ☑ Fri □ Sat □ Sun

Figure 9-14 Policy Scan Schedule panel

9.3.17 Replication options

The replication options enable you to control data movement in an archival migration policy. The following options are similar to those provided by a migration policy (see 8.7, “Replication options” on page 216):

- **Include subfolders**

  An archival migration policy can be set to copy all files and folder or files at the top level. This is the default configuration to make sure all data is coped.

- **Delete orphaned files in destination folders**

  Makes the source an identical copy of the destination. This option deletes any extra files on the destination. Do not select this option if you have multiple sources going to the same destination.

- **Enable Copy-in-place**

  Many optimizations in a data movement policy are a compromise between speed and safety. The Copy-in-place option in the Replication Options panel can help a migration policy finish faster with a slight security risk. Copy-in-place comes into effect during an incremental copy or a final phase copy. Assume the following scenario for a migration policy that moves files from one machine to another:

  - An xyz.doc file is on the source.
  - The file is copied to the destination.
  - Assume the file is modified on the source.
  - Because the file on the source is newer than the destination, xyz.doc on the source must replace the file on the destination.

  If Copy-in-place is not checked, the following events occur:

  - The file xyz.doc is copied to the destination as a temporary file.
  - When the copy is complete, the file xyz.doc is deleted on the destination, and the newly copied temporary file is renamed xyz.doc (this process is called *safe rename technique*).

  If Copy-in-place is checked, the following events occur:

  - The file xyz.doc is truncated to 0 bytes on the destination.
  - The file from the source is copied to the truncated file.

  Checking copy-in-place is faster because no rename operations are required. Copy-in-place is preferred when the source and destination are on the same LAN or connected by a high-speed network link. If the source and destination are connected by a low-speed link or the link is not reliable, you must clear copy-in-place and allow the safe rename technique to work. This way, if the network link breaks when the copy is in progress, you are not left with a zero-length file on the destination.

- **Use of NTFS change journal**

  At times, you might have a large source volume. This source volume might have to be kept in sync with a destination volume during the incremental phase of an archival migration. If the source volume has millions of files, the time required to scan the source volume to determine the changed files can be significant.

  VFM can optimize the search of changed source files on Microsoft Windows machines by integrating with the NTFS change journal. The NTFS change journal (only available on Windows) is a “file” that maintains a list of accessed, changed, created, and deleted files. If the Replication Agent scans the source NTFS change journal, it can detect the changed, deleted, and newly created files in a fraction of the time compared to doing a full file scan of TBs of source data. This way, an incremental migration resynchronization proceeds much faster.

  > **Note:** The change journal file can wrap, thus resulting in lost changes. The maximum change journal size on the volume is based on space you have allocated. If the change journal wraps, VFM resorts to a full scan of the source to determine the changed files.

- **Use snapshot data**

  During any data movement policy run, one or more files can be open on the source. This is usually the case for PST files and Access databases where the user might have an exclusive lock on the file. If VFM can get a read lock on the file, it copies the file. Depending on how the policy is set up, VFM might either skip the opened file or copy the file as best as it can from the snapshot. VFM can integrate with Volume Shadow Service (VSS) on Microsoft Windows 2003 or VFM snapshots. VFM can take a temporary snapshot of the source volume and copy all the files from the source snapshot. Once the files are copied, the source snapshot is automatically deleted. If the source is not Windows 2003 or N series, You can force VFM to retry the failed copies. Files that are not copied are logged as errors and are reported in the manifest based on the settings of the policy.

- **Allow loss of additional file streams**

  This option enables fine-grained control of archival migration between heterogeneous machines. Some vendor machines support file streams, while others do not. Some Microsoft Windows applications store data in alternate data streams. In most cases, you might or might not know if alternate data streams are used in the data files. At times, you might or might not know whether the vendor supports file streams. If VFM cannot copy a file because the destination does not support streams, the only way to copy the file that has streams is to select the Allow loss of additional file streams check box. Otherwise, VFM cannot move the file with streams, and it is reported as an error.

- **Event Details**

  The settings for the Event Details drop box have significant effect on migration performance. Event details enable a migration policy to provide a log of each file moved. The optimal setting is to only list files that encounter errors. If you choose to log each file moved, the amount of space required for the manifest is exceptionally large. Also, this manifest is stored with the Replication Agent. To view the manifest, it must be transferred to the server and then the client, which can take an inordinately long time. For test and demo migrations, it is acceptable to list all files so you can get comfortable using the product. For all production migrations, we strongly recommend you only list files that have errors.

9.3.18 Replication filters

The Replication Filters option provides a more flexible way of controlling files to be copied or excluded in a migration. At times, you can choose not to copy temp files.

9.3.19 Agent options

By default, VFM uses the Replication Agent on the destination to pull from data the source of an archival migration policy. This method is preferred because the agent on the destination knows how much data it can handle based on the load on the destination machine. If the Replication Agent cannot run on the destination (for example, the destination is an N series machine), you can run the Replication Agent on the source and push data to the destination. Other configurations are also possible where the agent is on a proxy machine and performing a three-way copy. You need not manually install the Replication Agent. Based on the archival migration policy settings, the Replication Agent is automatically deployed to the specified machine. An agent is multithreaded and can take advantage of multiprocessor systems.

The Replication Agent runs under the credentials of the Service Account that VFM runs under. The Service Account is specified when the product is installed. This Service Account must have administrator rights on the source and destination machines. If the Replication Agent does not have security permissions to a folder or file, it is not able to copy the data. In that case, use the Agent Management tool to set the credentials.

The Replication Agent communicates with the VFM server and periodically passes status messages. This heartbeat happens every few minutes. Make sure that the network ports (6001, 6002, and 6005) are open for communication across all firewalls in the path between the VFM server and the Replication Agent; otherwise, the Replication Agent cannot communicate with the server (see Chapter 3, “Preparation” on page 83 for details). If, for any reason, the Replication Agent cannot talk to the VFM server, the Replication Agent aborts the job, and the migration fails.

Agent grouping provides a way to share the load between multiple Replication Agents involved in a data movement policy. Assume you are moving many TBs of data from one N series machine to another. Instead of using one Replication Agent to perform all the migrations, it might be advisable to use a group of Replication Agents to move the data. You do so by setting up an agent group in the Agent Management interface. You specify which machines are members of the agent group.

When you configure an archival migration policy, you select the agent group to perform the archival migration. When the archival migration policy runs, VFM sends the task to the next available agent in the agent group. A round robin technique is used to determine which Replication Agent in the agent group performs the actions of task. Using agent groups prevents a Replication Agent from being overloaded while sharing the load across multiple agents to perform data movements. If no agent of the group is available, the migration does not run.

For the purposes of this example, select the radio button **VFM selects**. This enables VFM to select the Replication Agent to be used based on internal algorithms.

9.3.20 Differential replication or byte level replication (BLR)

VFM provides byte level replication to control data movement across slow links. BLR (see Figure 9-15 on page 293) minimizes the data sent across the wire, especially with slow network links, by only sending the changed bytes. You can have a 10 MB Microsoft Word document file that is copied from the source to the destination share. Assume that the user at the source changes one or two pages in the document. Instead of sending the 10 MB changed file again from the source to the destination, VFM can be used to intelligently send the deltas across the wire. While the time required in computing the changed bytes in extremely large files can be significant, network usage is minimized. You are sacrificing the extra CPU cycles to figure out the changed bytes so that the limited resource of network bandwidth is optimized. BLR requires two Replication Agents to work together to make the replication happen successfully.

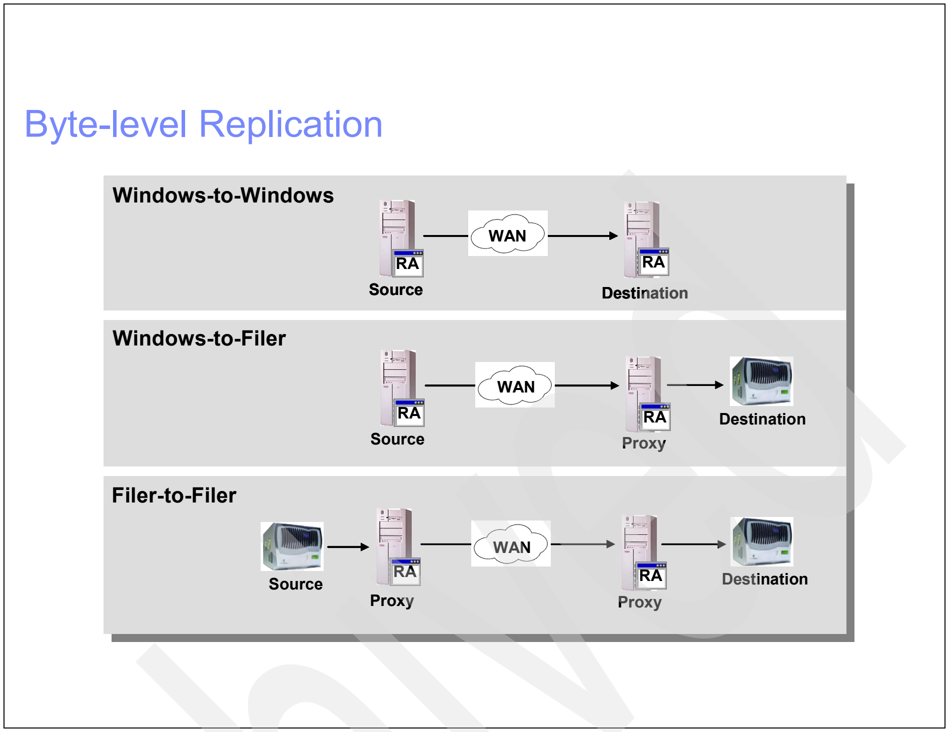

Byte-level Replication

Figure 9-15 Byte level replication

Use BLR only across slow links. In LAN speed networks, the overhead of BLR computation is more than the benefit it provides. BLR works when the destination Replication Agent computes rolling hash sets across all the blocks of the file at the destination. The destination Replication Agent sends the rolling hash set as a file signature to the source Replication Agent. This results in a small amount of data transferred from destination to source. The source Replication Agent computes the rolling hash set of the changed file at the source. When different hash blocks are encountered, the source Replication Agent computes an instruction set to send from the source to the destination. The instruction set is used to insert or remove data blocks from the file at the destination to make it look like the source. BLR over slow WAN links results in 80% savings in network traffic of slightly changed data files.

For the purposes of this example, do not enable differential replication.

9.3.21 Security replication

Every data migration requires that the files be copied accurately and completely. VFM provides several mechanisms to control the copying of files. When possible, VFM signals errors when files cannot be copied completely or correctly.

Many Active Directory deployments use local groups to secure the access to file resources on the old server. In NT4 domains, Microsoft recommends the use of local groups. A local group is valid only on the source system on which it is defined. If the security on a file is controlled by the permissions of a local group and that file is copied to a destination machine, the security is lost. This means that the users who had access to the files on the old server might not have access to the migrated copy.

VFM provides a way to handle this situation as described by the Security Replication tab. VFM can create a local group on the destination with the same name as the local group on the source. All domain users and domain groups that are members of the old local group on the source server are added to the newly created local group of the destination server. This way, users that had access to the data through the old local group on the source server have access to the data on the new server through the newly created local group.

At times, you care only about replicating the data from one location to another without regard to the security of the data on the new system. For the purposes of this example, select the default check boxes and radio buttons.

9.3.22 Attribute replication

In a migration policy, you can control whether attributes are used to compare whether two files are identical or different. By default, VFM compares the names, size, dates, and attributes (see Figure 9-30 on page 308). You can override this default by checking or clearing the check box for attribute replication. You can also specify how the attribute is set on the destination file after it has been copied

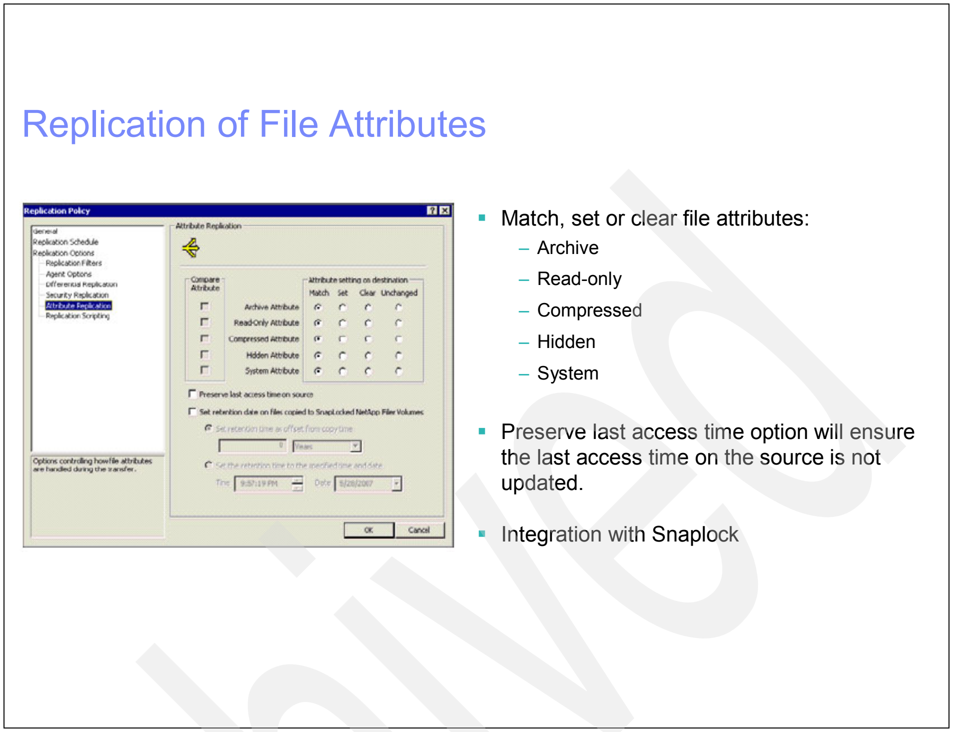

Replication of File Attributes

Options controlling how file attributes are handled during the transfer.

	Compare Attribute	Match	Set	Clear	Unchanged
Archive Attribute	□	●	○	○	○
Read-Only Attribute	□	●	○	○	○
Compressed Attribute	□	●	○	○	○
Hidden Attribute	□	●	○	○	○
System Attribute	□	●	○	○	○

□ Preserve last access time on source

□ Set retention date on files copied to SnapLocked NetApp Filer Volumes

● Set retention time as offset from copy time: 0 Years

○ Set the retention time to the specified time and date: Time 9:57:19 PM Date 5/28/2007

- Match, set or clear file attributes:
  - Archive
  - Read-only
  - Compressed
  - Hidden
  - System
- Preserve last access time option will ensure the last access time on the source is not updated.
- Integration with Snaplock

Figure 9-16 File attribute replication

9.4 Archival migration policy example 2

For the purposes of this example, select the following options:

- Replication options: Choose defaults.
- Replication filters: Choose defaults.
- Agent options: Choose defaults.
- Differential replication options: Choose defaults.
- Security Replication options: Choose defaults.
- Attribute replication options: You must clear the **Preserve last access time on source** check box because you are deleting the data on the source.
- Replication scripting options: Choose defaults.

Once the archival migration policy creation is complete, view the selected options for the same (see Figure 9-17).

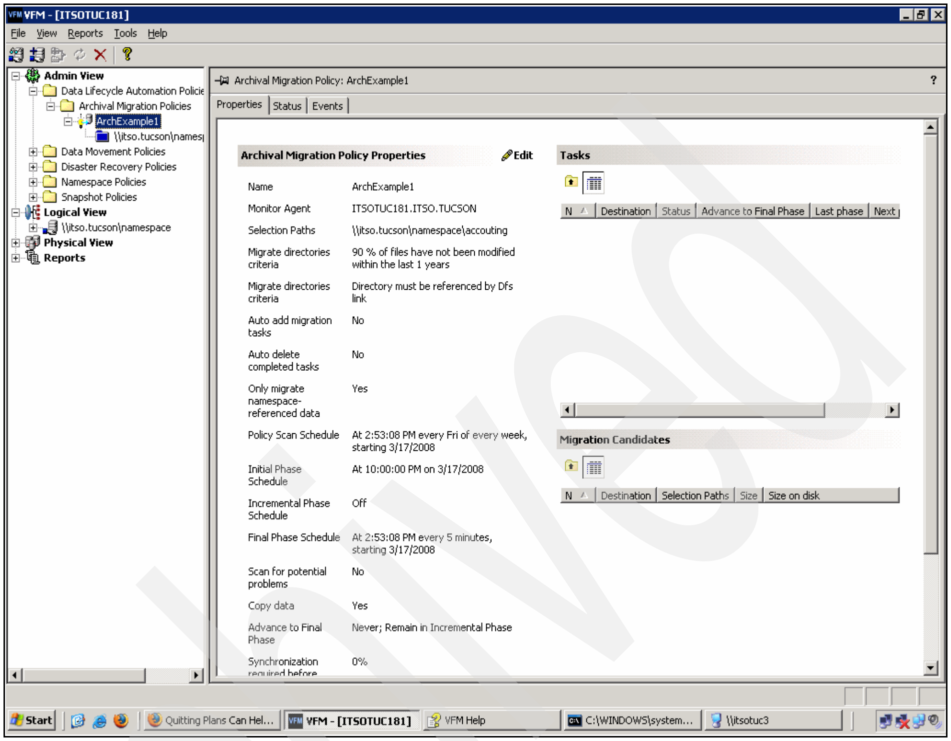

Archival Migration Policy Properties

Name	ArchExample1
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Selection Paths	\\itso.tucson\namespace\accouting
Migrate directories criteria	90 % of files have not been modified within the last 1 years
Migrate directories criteria	Directory must be referenced by Dfs link
Auto add migration tasks	No
Auto delete completed tasks	No
Only migrate namespace-referenced data	Yes
Policy Scan Schedule	At 2:53:08 PM every Fri of every week, starting 3/17/2008
Initial Phase Schedule	At 10:00:00 PM on 3/17/2008
Incremental Phase Schedule	Off
Final Phase Schedule	At 2:53:08 PM every 5 minutes, starting 3/17/2008
Scan for potential problems	No
Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Synchronization required before	0%

Figure 9-17 Archival Migration Policy Properties panel

Follow these steps to identify archival candidates:

1. Start the scan to find possible archival candidates (see Figure 9-18).

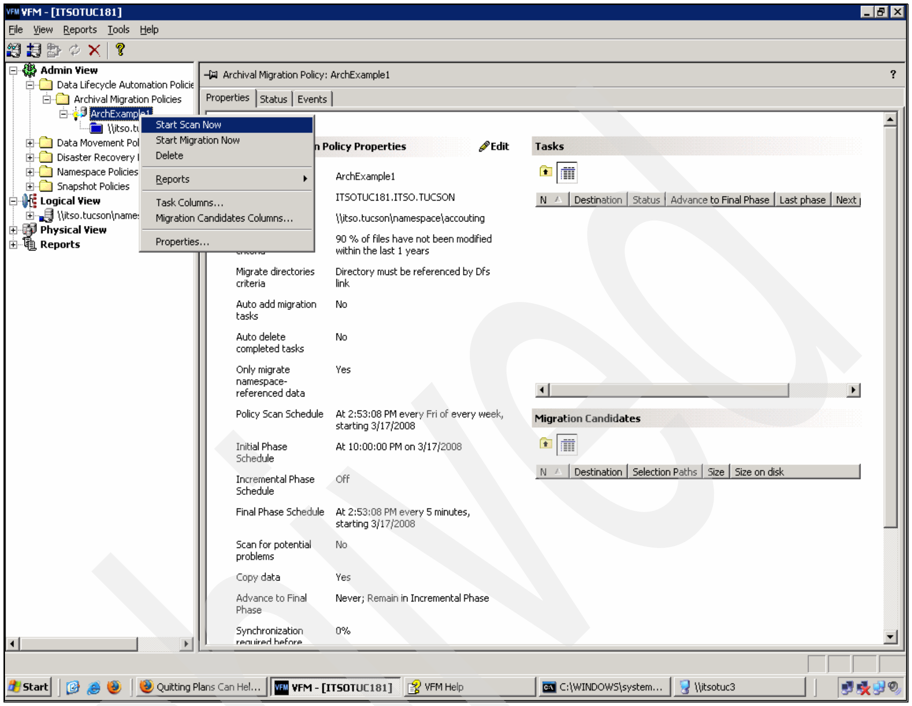

	ArchExample1
	ITSOTUC181.ITSO.TUCSON
	\\itso.tucson\namespace\accouting
criteria	90 % of files have not been modified within the last 1 years
Migrate directories criteria	Directory must be referenced by Dfs link
Auto add migration tasks	No
Auto delete completed tasks	No
Only migrate namespace-referenced data	Yes
Policy Scan Schedule	At 2:53:08 PM every Fri of every week, starting 3/17/2008
Initial Phase Schedule	At 10:00:00 PM on 3/17/2008
Incremental Phase Schedule	Off
Final Phase Schedule	At 2:53:08 PM every 5 minutes, starting 3/17/2008
Scan for potential problems	No
Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Synchronization required before	0%

Figure 9-18 Scanning for probable archival candidates

2. Check the results of the scan (see Figure 9-19).

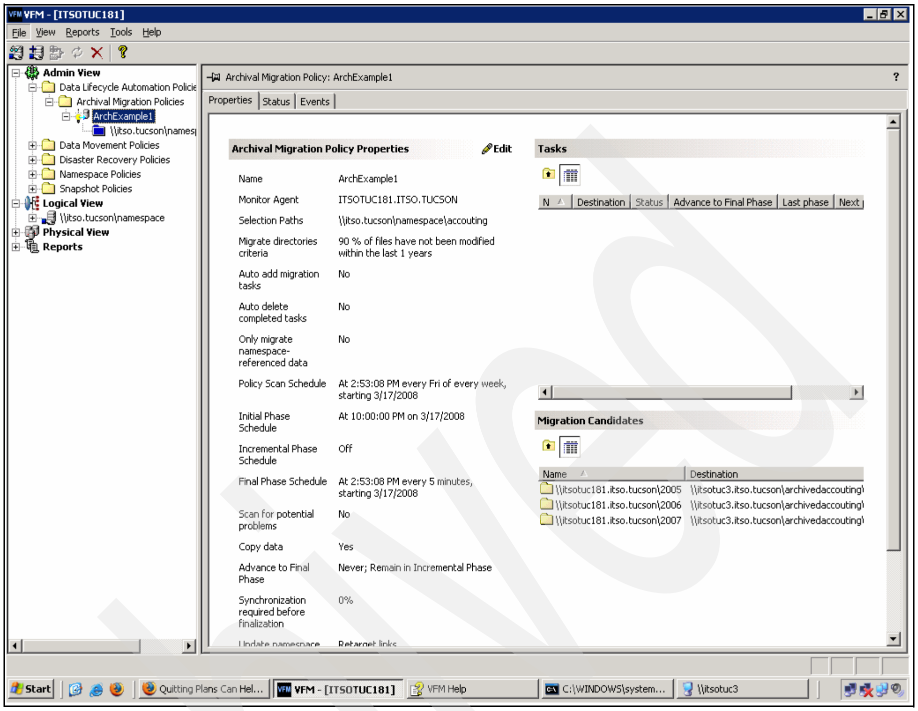

Archival Migration Policy Properties

Name	ArchExample1
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Selection Paths	\\itso.tucson\namespace\accouting
Migrate directories criteria	90 % of files have not been modified within the last 1 years
Auto add migration tasks	No
Auto delete completed tasks	No
Only migrate namespace-referenced data	No
Policy Scan Schedule	At 2:53:08 PM every Fri of every week, starting 3/17/2008
Initial Phase Schedule	At 10:00:00 PM on 3/17/2008
Incremental Phase Schedule	Off
Final Phase Schedule	At 2:53:08 PM every 5 minutes, starting 3/17/2008
Scan for potential problems	No
Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Synchronization required before finalization	0%
Update namespace	Retarget links

Tasks

N	Destination	Status	Advance to Final Phase	Last phase	Next

Migration Candidates

Name	Destination
\\itsotuc181.itso.tucson\2005	\\itsotuc3.itso.tucson\archivedaccouting\
\\itsotuc181.itso.tucson\2006	\\itsotuc3.itso.tucson\archivedaccouting\
\\itsotuc181.itso.tucson\2007	\\itsotuc3.itso.tucson\archivedaccouting\

Figure 9-19 Scan results

3. Select the share you want to move into the archive. Add a task for each of the selected archival candidates (see Figure 9-20).

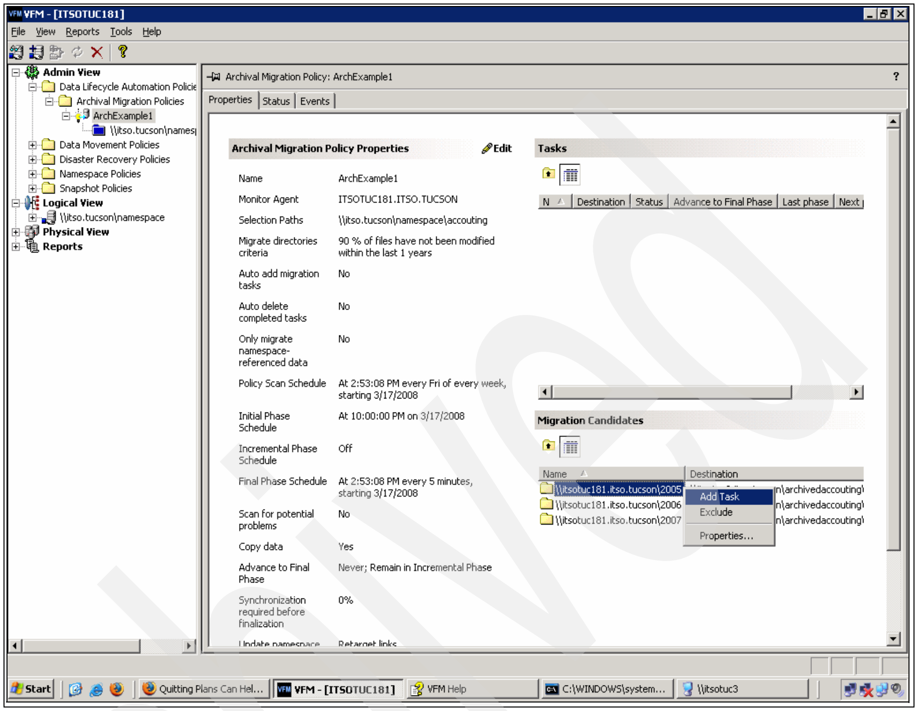

Archival Migration Policy Properties

Name	ArchExample1
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Selection Paths	\\itso.tucson\namespace\accouting
Migrate directories criteria	90 % of files have not been modified within the last 1 years
Auto add migration tasks	No
Auto delete completed tasks	No
Only migrate namespace-referenced data	No
Policy Scan Schedule	At 2:53:08 PM every Fri of every week, starting 3/17/2008
Initial Phase Schedule	At 10:00:00 PM on 3/17/2008
Incremental Phase Schedule	Off
Final Phase Schedule	At 2:53:08 PM every 5 minutes, starting 3/17/2008
Scan for potential problems	No
Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Synchronization required before finalization	0%
Update namespace	Retarget links

Migration Candidates

Name	Destination
\\itsotuc181.itso.tucson\2005	n\archivedaccouting\
\\itsotuc181.itso.tucson\2006	n\archivedaccouting\
\\itsotuc181.itso.tucson\2007	n\archivedaccouting\

Figure 9-20 Adding task for archival candidates

On successful completion of the task, the share moves from Migration Candidates to Tasks (see Figure 9-21).

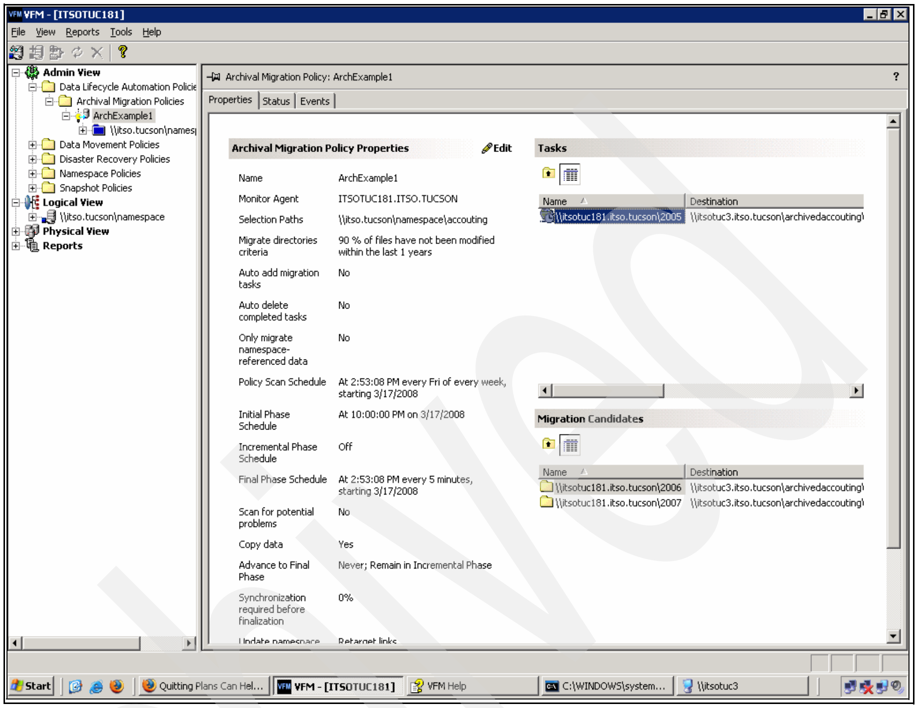

Archival Migration Policy: ArchExample1

Archival Migration Policy Properties

Name	ArchExample1
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Selection Paths	\\itso.tucson\namespace\accouting
Migrate directories criteria	90 % of files have not been modified within the last 1 years
Auto add migration tasks	No
Auto delete completed tasks	No
Only migrate namespace-referenced data	No
Policy Scan Schedule	At 2:53:08 PM every Fri of every week, starting 3/17/2008
Initial Phase Schedule	At 10:00:00 PM on 3/17/2008
Incremental Phase Schedule	Off
Final Phase Schedule	At 2:53:08 PM every 5 minutes, starting 3/17/2008
Scan for potential problems	No
Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Synchronization required before finalization	0%
Update namespace	Retarget links

Tasks

Name	Destination
\\itsotuc181.itso.tucson\2005	\\itsotuc3.itso.tucson\archivedaccouting\

Migration Candidates

Name	Destination
\\itsotuc181.itso.tucson\2006	\\itsotuc3.itso.tucson\archivedaccouting\
\\itsotuc181.itso.tucson\2007	\\itsotuc3.itso.tucson\archivedaccouting\

Figure 9-21 Selecting task for migration

4. Start the process of migration by starting the task manually (see Figure 9-22).

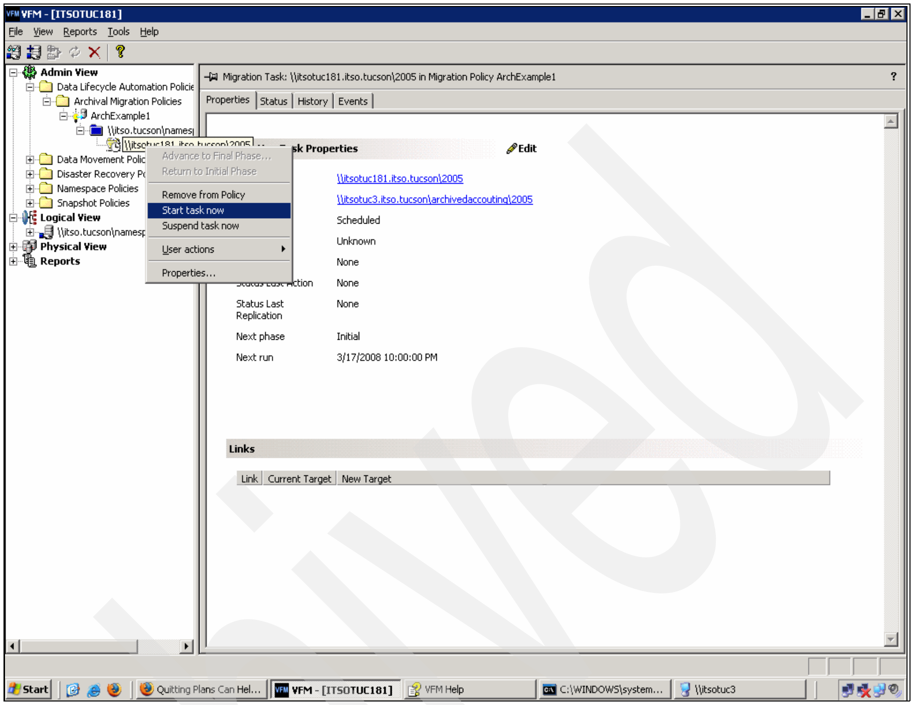

\\itsotuc181.itso.tucson\2005
\\itsotuc3.itso.tucson\archivedaccouting\2005
Scheduled
Unknown
None
Status Last Action None
Status Last Replication None
Next phase Initial
Next run 3/17/2008 10:00:00 PM

Figure 9-22 Starting task manually

Once you start the task, VFM executes the initial, incremental, and final phases as per the configurations of the policy (see Figure 9-23).

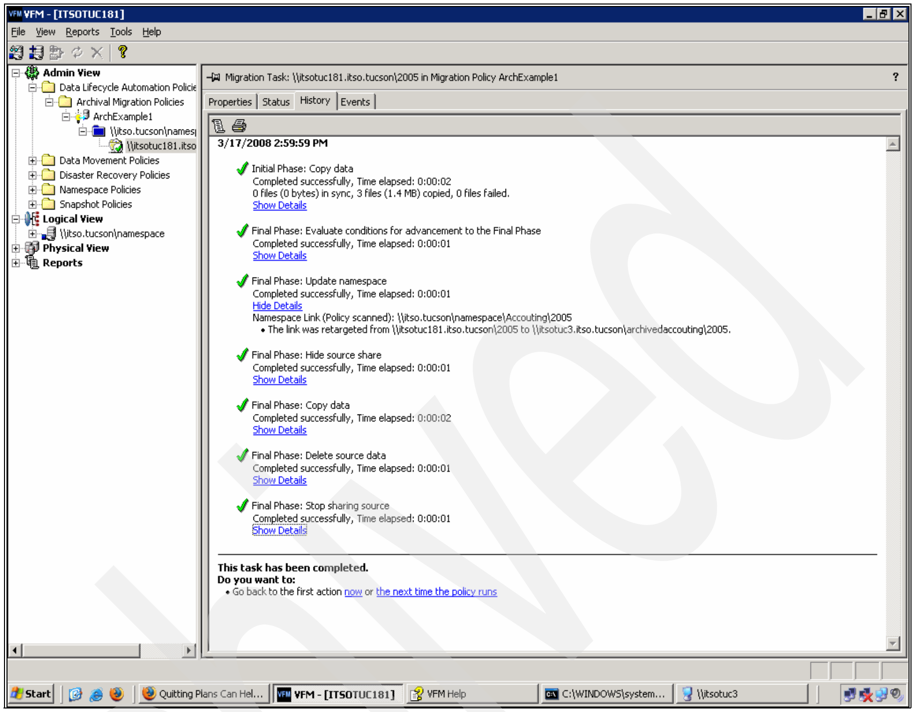

Migration Task: \\itsotuc181.itso.tucson\2005 in Migration Policy ArchExample1

3/17/2008 2:59:59 PM

✓ Initial Phase: Copy data
Completed successfully, Time elapsed: 0:00:02
0 files (0 bytes) in sync, 3 files (1.4 MB) copied, 0 files failed.
Show Details

✓ Final Phase: Evaluate conditions for advancement to the Final Phase
Completed successfully, Time elapsed: 0:00:01
Show Details

✓ Final Phase: Update namespace
Completed successfully, Time elapsed: 0:00:01
Hide Details
Namespace Link (Policy scanned): \\itso.tucson\namespace\Accouting\2005
- The link was retargeted from \\itsotuc181.itso.tucson\2005 to \\itsotuc3.itso.tucson\archivedaccouting\2005.

✓ Final Phase: Hide source share
Completed successfully, Time elapsed: 0:00:01
Show Details

✓ Final Phase: Copy data
Completed successfully, Time elapsed: 0:00:02
Show Details

✓ Final Phase: Delete source data
Completed successfully, Time elapsed: 0:00:01
Show Details

✓ Final Phase: Stop sharing source
Completed successfully, Time elapsed: 0:00:01
Show Details

This task has been completed.
Do you want to:
- Go back to the first action now or the next time the policy runs

Figure 9-23 History for migration phases

Share 2005 is moved from itsotuc181 to itsotuc3 (see Figure 9-24).

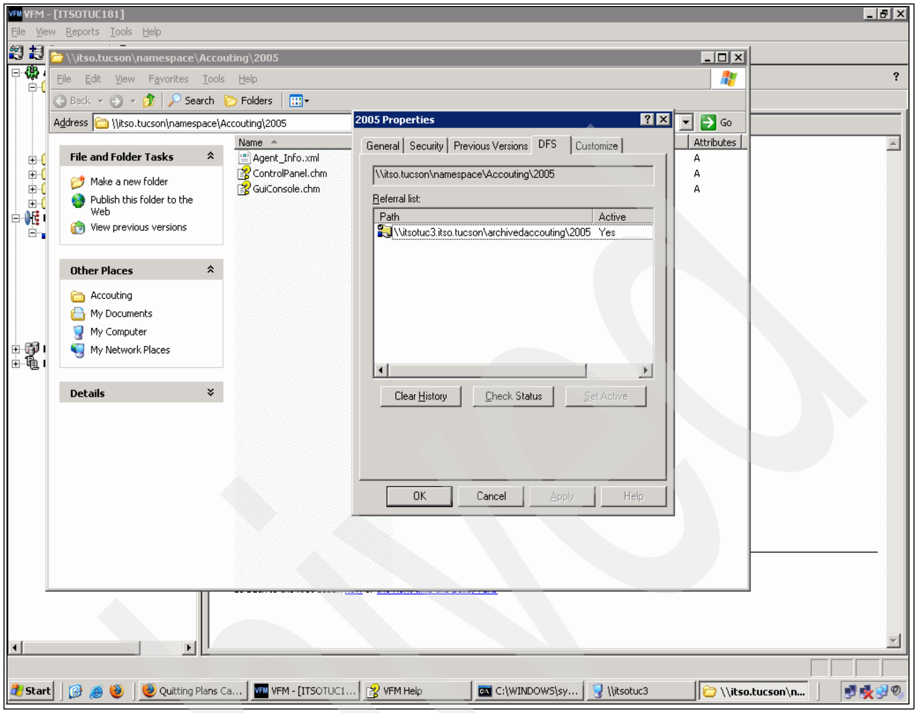

Address: \\itso.tucson\namespace\Accouting\2005

Name	Attributes
Agent_Info.xml	A
ControlPanel.chm	A
GuiConsole.chm	A

\\itso.tucson\namespace\Accouting\2005

Referral list:

Path	Active
\\itsotuc3.itso.tucson\archivedaccouting\2005	Yes

Figure 9-24 DFS link for 2005 share changed to itsotuc3

5. Also confirm the dependencies for the namespace and accounting folder (see Figure 9-25).

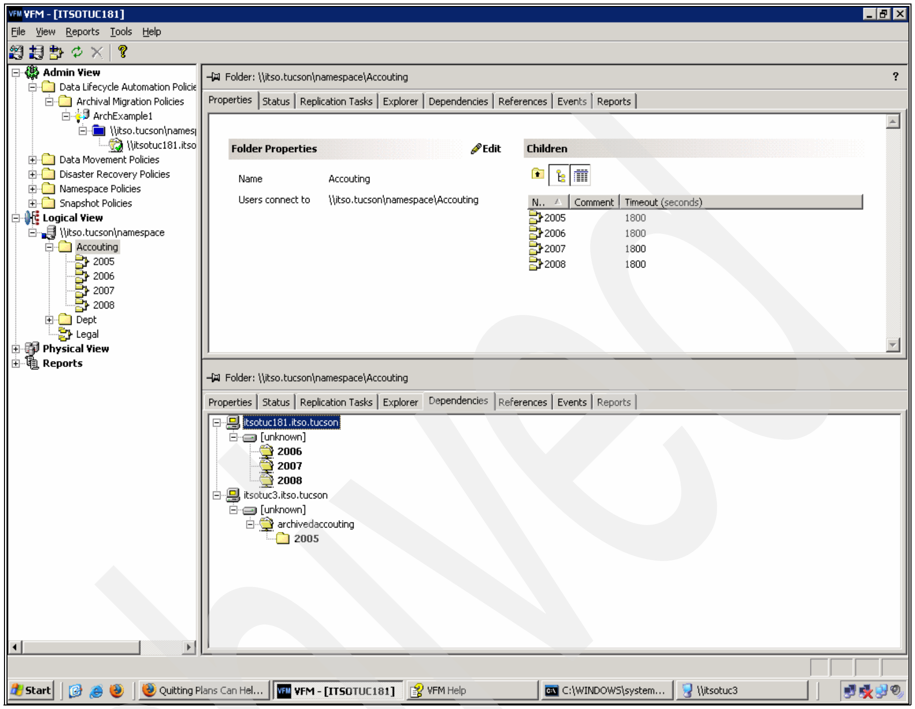

Name Accouting
Users connect to \\itso.tucson\namespace\Accouting

N..	Comment	Timeout (seconds)
2005		1800
2006		1800
2007		1800
2008		1800

Figure 9-25 Dependencies changed for share 2005

9.5 Archival migration policy example 3

Consider an environment where two file servers are used. The servers are used for storing different project data. These file servers are almost at their maximum storage capacity. We want to move certain projects to an archival server. You intend to archive the projects for which 90% of their files are unchanged since last year (see Figure 9-26 on page 305).

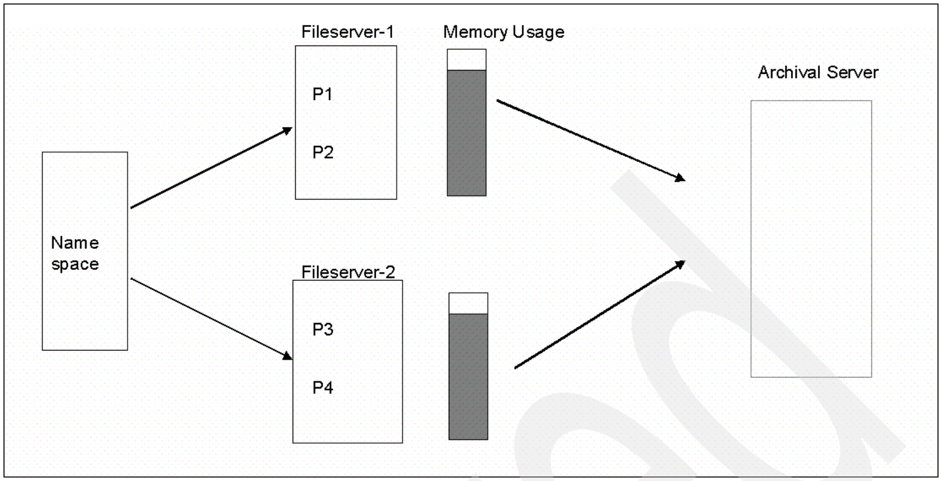

Figure 9-26 Storage usage before archive

Itsotuc181 and itsotuc3 are the two file servers (see Figure 9-27).

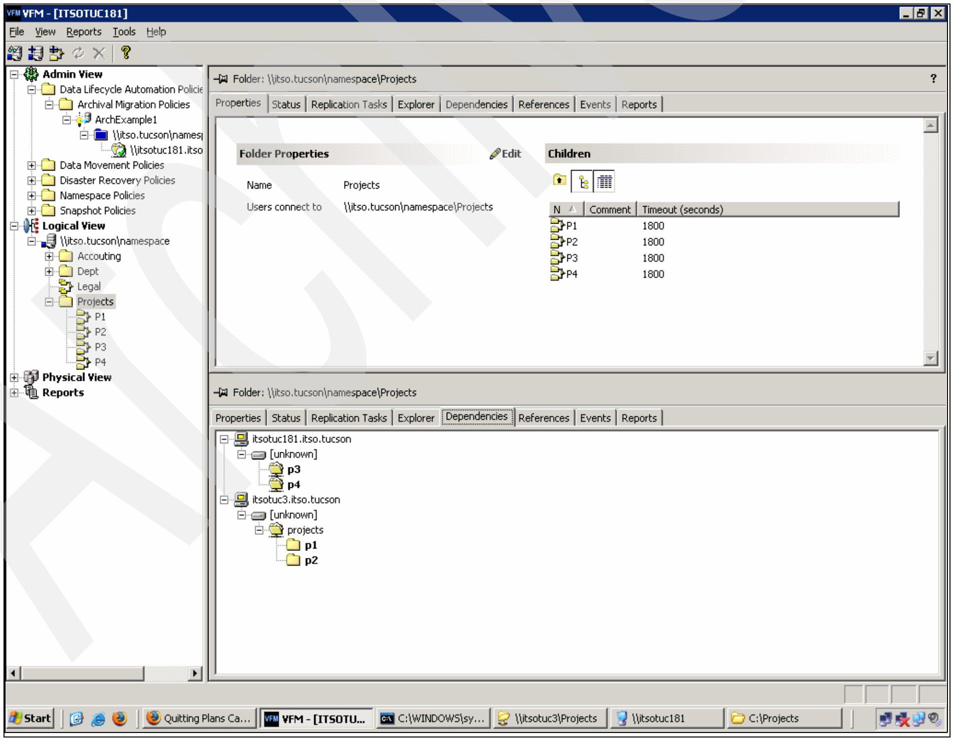

Folder: \\itso.tucson\namespace\Projects

Folder Properties

Name Projects

Users connect to \\itso.tucson\namespace\Projects

Children

N	Comment	Timeout (seconds)
P1		1800
P2		1800
P3		1800
P4		1800

Folder: \\itso.tucson\namespace\Projects

```
itsotuc181.itso.tucson
  [unknown]
    p3
    p4
itsotuc3.itso.tucson
  [unknown]
    projects
      p1
      p2
```

Figure 9-27 Before archival process

Using the procedures in 9.3, “Archival migration policy example 1” on page 269 as references, follow these steps:

1. Create a new archival migration policy to move projects that meet the criterion of 90% of files unchanged in the last year. For this example, P2 and P3 are those projects.
2. Add a task for both P2 and P3 to migrate to the archival server, which is itsotuc4 in this example.
3. Once the tasks are created, start them manually one by one. Observe the history for the successful completion of migration phases. On successful completion, projects P2 and P3 have moved to the itsotuc4 archival server (see Figure 9-28).

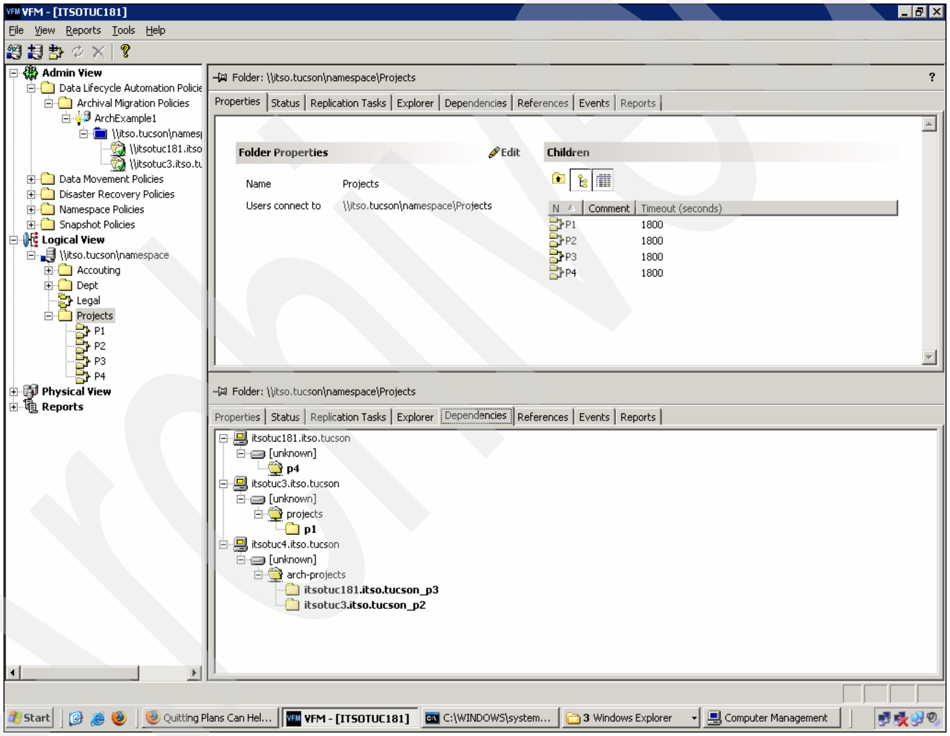
VFM - [ITSOTUC181]
File View Reports Tools Help
Admin View
Data Lifecycle Automation Policie
Archival Migration Policies
ArchExample1
\\itso.tucson\names
\\itsotuc181.itso
\\itsotuc3.itso.tu
Data Movement Policies
Disaster Recovery Policies
Namespace Policies
Snapshot Policies
Logical View
\\itso.tucson\namespace
Accouting
Dept
Legal
Projects
P1
P2
P3
P4
Physical View
Reports
Folder: \\itso.tucson\namespace\Projects
Properties Status Replication Tasks Explorer Dependencies References Events Reports
Folder Properties
Edit
Children
Folder: \\itso.tucson\namespace\Projects
Properties Status Replication Tasks Explorer Dependencies References Events Reports
Start
Quitting Plans Can Hel...
VFM - [ITSOTUC181]
C:\WINDOWS\system...
3 Windows Explorer
Computer Management

Name Projects

Users connect to \\itso.tucson\namespace\Projects

N	Comment	Timeout (seconds)
P1		1800
P2		1800
P3		1800
P4		1800

```
itsotuc181.itso.tucson
  [unknown]
    p4
itsotuc3.itso.tucson
  [unknown]
    projects
      p1
itsotuc4.itso.tucson
  [unknown]
    arch-projects
      itsotuc181.itso.tucson_p3
      itsotuc3.itso.tucson_p2
```

Figure 9-28 After archival process

The storage usage of the file servers is less now, and users can access the projects P2 and P3 through the same namespace (see Figure 9-29).

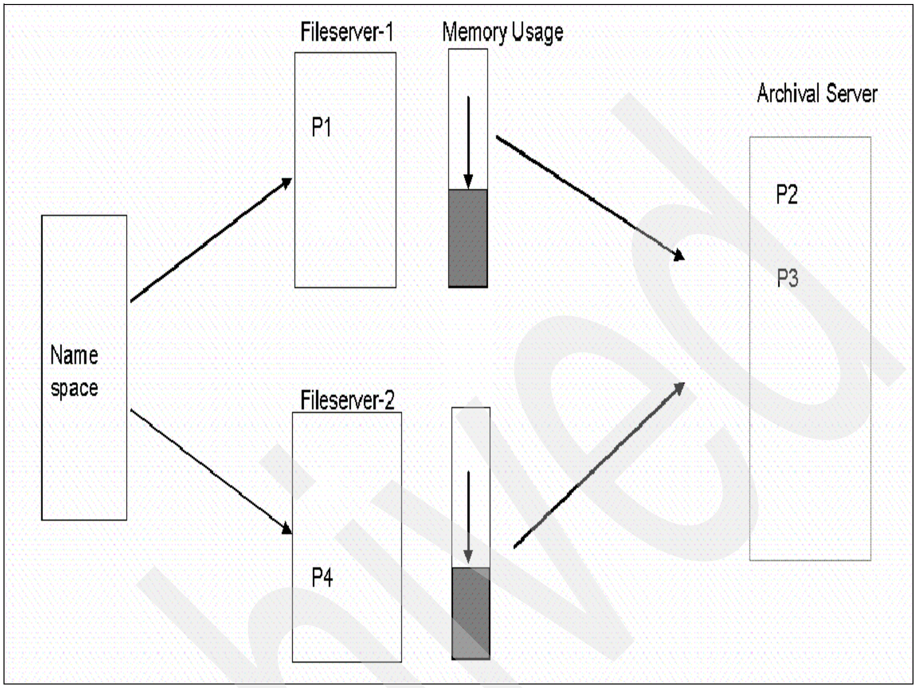

Figure 9-29 Storage usage after archive: example 3

9.6 Archival migration policy example 4

In this example, the user accesses the share while data is being archived. Share P6 is available under the Projects folder of the namespace (see Figure 9-30 on page 308).

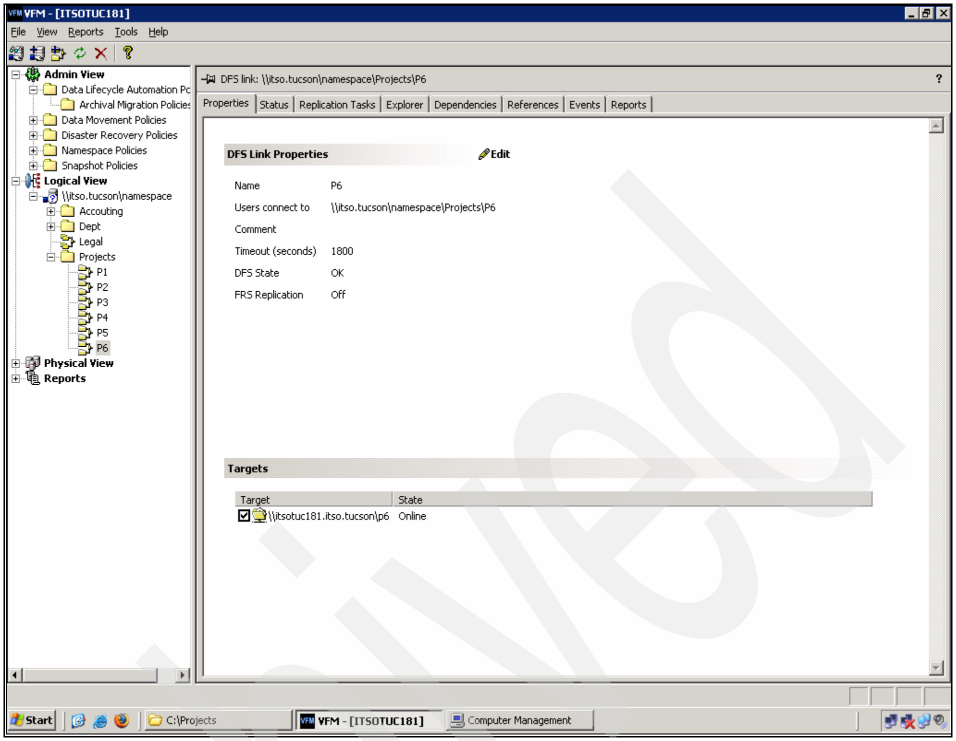

DFS Link Properties

Name	P6
Users connect to	\\itso.tucson\namespace\Projects\P6
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
☑ \\itsotuc181.itso.tucson\p6	Online

Figure 9-30 Archival migration policy: example 4

Share P6 is accessible through the namespace path (see Figure 9-31).

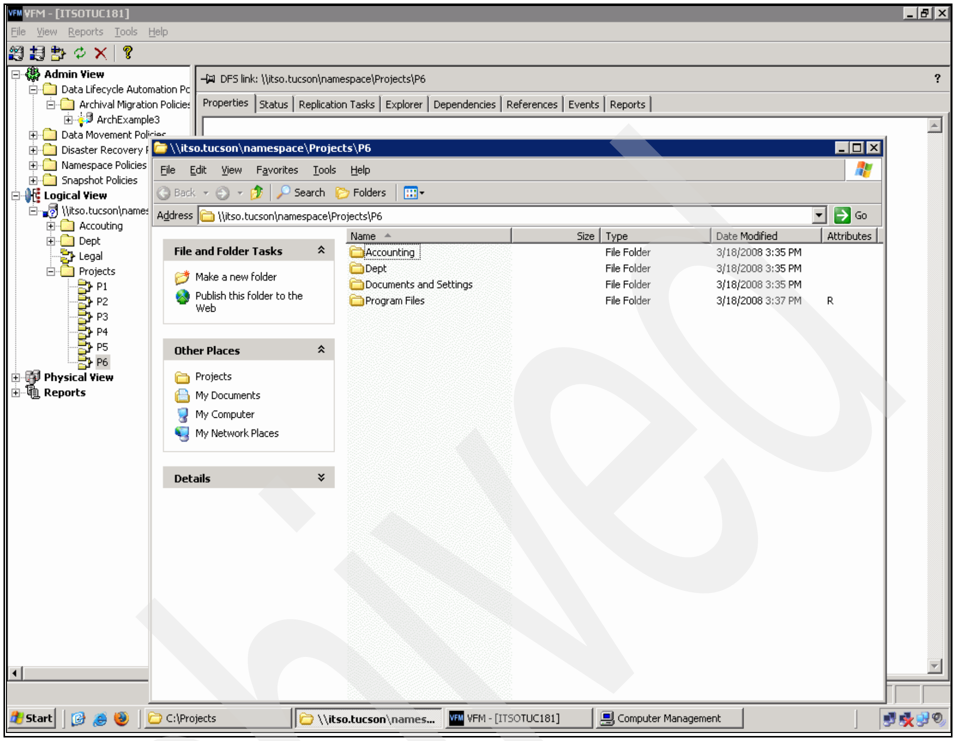

Name	Size	Type	Date Modified	Attributes
Accounting		File Folder	3/18/2008 3:35 PM	
Dept		File Folder	3/18/2008 3:35 PM	
Documents and Settings		File Folder	3/18/2008 3:35 PM	
Program Files		File Folder	3/18/2008 3:37 PM	R

Figure 9-31 Share P6 is accessible

Follow these steps to create the migration policy and start a scan:

1. Create a new migration policy ArchExample3, with a final phase setting as shown in Figure 9-32.

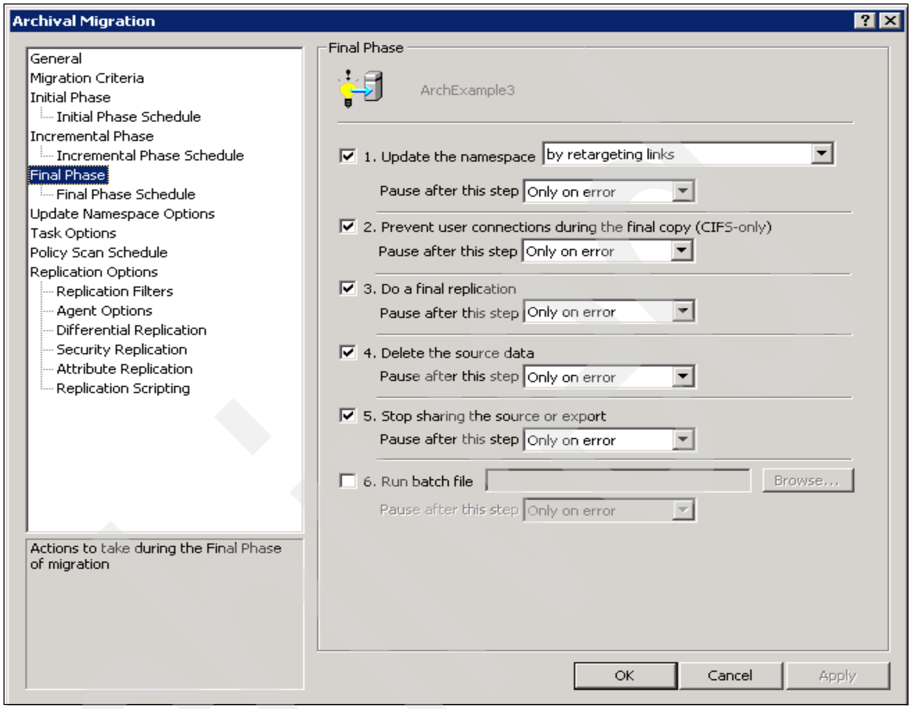

Actions to take during the Final Phase of migration

Final Phase

ArchExample3

☑ 1. Update the namespace: by retargeting links
Pause after this step: Only on error

☑ 2. Prevent user connections during the final copy (CIFS-only)
Pause after this step: Only on error

☑ 3. Do a final replication
Pause after this step: Only on error

☑ 4. Delete the source data
Pause after this step: Only on error

☑ 5. Stop sharing the source or export
Pause after this step: Only on error

☐ 6. Run batch file
Pause after this step: Only on error

Figure 9-32 Final phase settings

2. Select the default values for other parameters, and complete the Archival migration policy as explained in 9.3, “Archival migration policy example 1” on page 269.

3. Start the scan manually (see Figure 9-33).

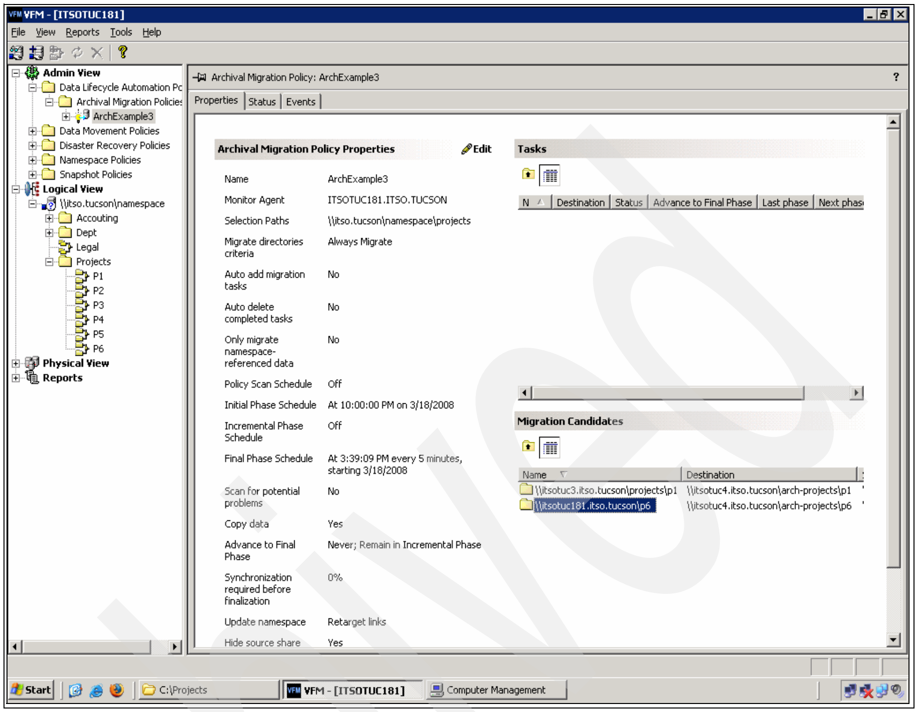

Archival Migration Policy Properties

Name	ArchExample3
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Selection Paths	\\itso.tucson\namespace\projects
Migrate directories criteria	Always Migrate
Auto add migration tasks	No
Auto delete completed tasks	No
Only migrate namespace-referenced data	No
Policy Scan Schedule	Off
Initial Phase Schedule	At 10:00:00 PM on 3/18/2008
Incremental Phase Schedule	Off
Final Phase Schedule	At 3:39:09 PM every 5 minutes, starting 3/18/2008
Scan for potential problems	No
Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Synchronization required before finalization	0%
Update namespace	Retarget links
Hide source share	Yes

Tasks

N	Destination	Status	Advance to Final Phase	Last phase	Next phas

Migration Candidates

Name	Destination
\\itsotuc3.itso.tucson\projects\p1	\\itsotuc4.itso.tucson\arch-projects\p1
\\itsotuc181.itso.tucson\p6	\\itsotuc4.itso.tucson\arch-projects\p6

Figure 9-33 Scan results for example 4

4. Add the archival migration task for share P6 (see Figure 9-34).

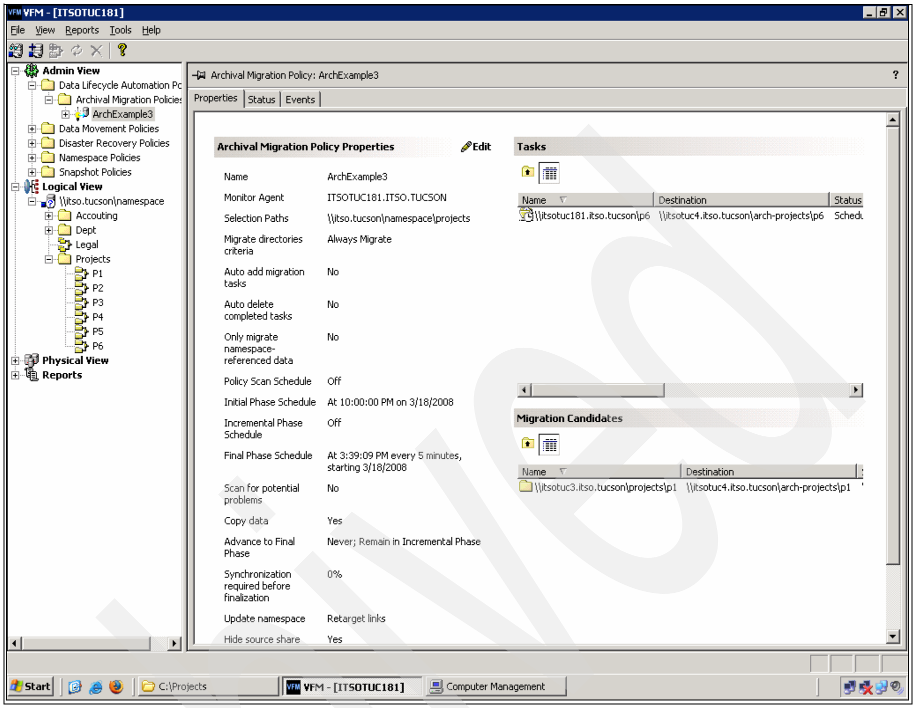

Archival Migration Policy Properties

Name	ArchExample3
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Selection Paths	\\itso.tucson\namespace\projects
Migrate directories criteria	Always Migrate
Auto add migration tasks	No
Auto delete completed tasks	No
Only migrate namespace-referenced data	No
Policy Scan Schedule	Off
Initial Phase Schedule	At 10:00:00 PM on 3/18/2008
Incremental Phase Schedule	Off
Final Phase Schedule	At 3:39:09 PM every 5 minutes, starting 3/18/2008
Scan for potential problems	No
Copy data	Yes
Advance to Final Phase	Never; Remain in Incremental Phase
Synchronization required before finalization	0%
Update namespace	Retarget links
Hide source share	Yes

Tasks

Name	Destination	Status
\\itsotuc181.itso.tucson\p6	\\itsotuc4.itso.tucson\arch-projects\p6	Schedu

Migration Candidates

Name	Destination
\\itsotuc3.itso.tucson\projects\p1	\\itsotuc4.itso.tucson\arch-projects\p1

Figure 9-34 Adding task for share P6

5. Start the task manually.

   As the archival migration is in progress, try to access the share P6. You receive an access error (see Figure 9-35).

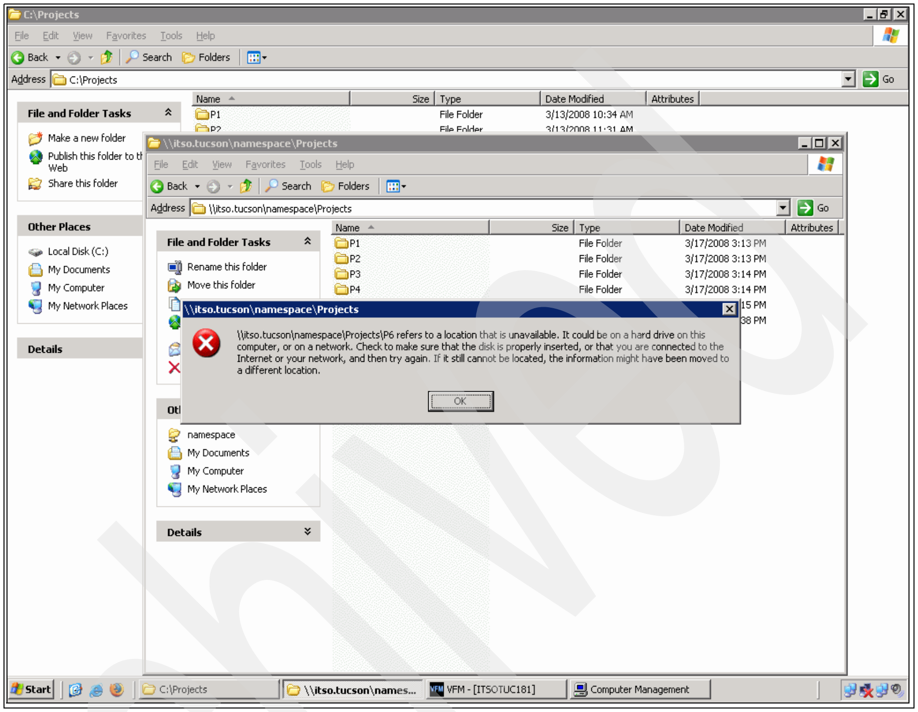

Name	Size	Type	Date Modified	Attributes
P1		File Folder	3/13/2008 10:34 AM	
P2		File Folder	3/13/2008 11:31 AM	

Name	Size	Type	Date Modified	Attributes
P1		File Folder	3/17/2008 3:13 PM	
P2		File Folder	3/17/2008 3:13 PM	
P3		File Folder	3/17/2008 3:14 PM	
P4		File Folder	3/17/2008 3:14 PM	

\\itso.tucson\namespace\Projects

\\itso.tucson\namespace\Projects\P6 refers to a location that is unavailable. It could be on a hard drive on this computer, or on a network. Check to make sure that the disk is properly inserted, or that you are connected to the Internet or your network, and then try again. If it still cannot be located, the information might have been moved to a different location.

OK

Figure 9-35 Access error

10

Replication

This chapter discusses how to use VFM for data replication. VFM provides the capability to replicate data between heterogeneous data sources. It is a general-purpose replication solution focused on unstructured file data.

Replication policies provide the means to synchronize data between two entities. VFM replication is asynchronous and schedule based. A replication policy is used to synchronize CIFS data between any two UNC shares or folders, or to synchronize data between any two NFS exports.

The machines involved can be located in the same data center (local replication), across data centers (remote replication), or between a remote office and a data center.

10.1 Using replication options

VFM provides numerous options for replicating data, offering flexibility and enabling you to choose only the proper options to efficiently use your resources while reliably replicating your data.

To use these options correctly, it is important to understand exactly what each means and does. In this section, we demonstrate the Replication Policy Wizard, showing you where you can find each option and explaining the most relevant ones in detail. Follow these steps:

1. To begin, open the VFM client by choosing **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.
2. Then go to **Admin View** → **Data Movement Policies** → **Replication Policies** (see Figure 10-1).

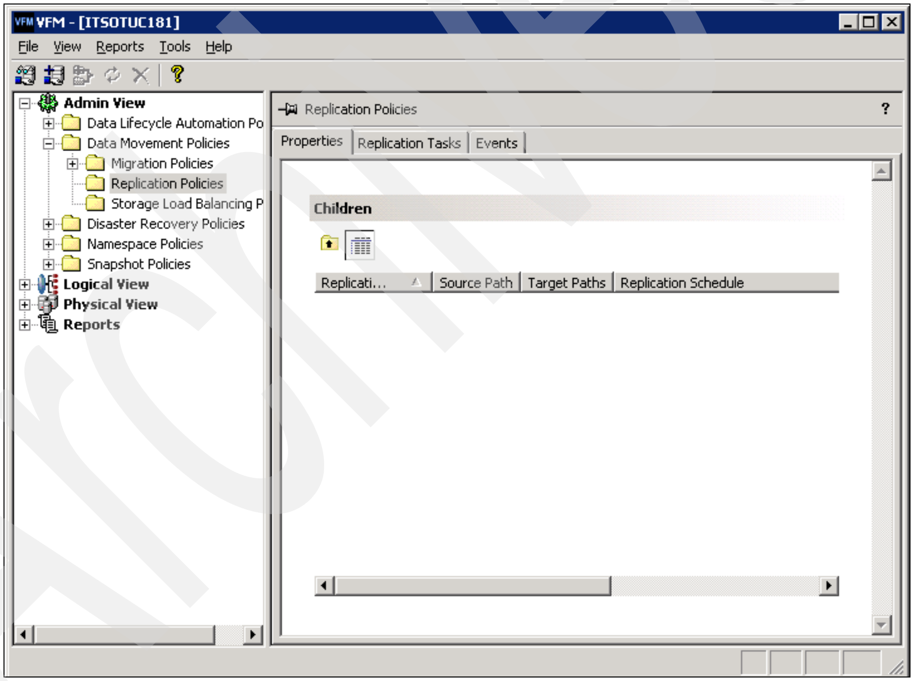

Figure 10-1 VFM client: Replication Policies panel

3. To create a new policy, right-click **Replication Policy**, then click **New Replication Policy** (see Figure 10-2).

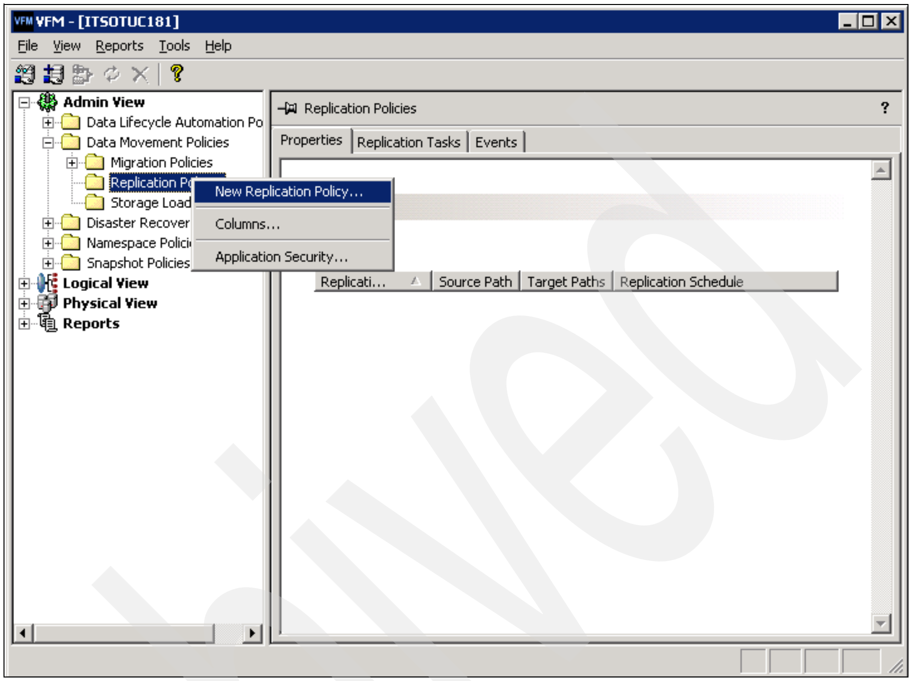

Figure 10-2 Replication Policies panel with right-click menu

10.1.1 General options

Figure 10-3 on page 318 shows the first window you access when creating a replication policy. In this section, we describe each option available:

- **Replication name**

  This first window enables you to name your replication policy, so that it can be identified and easily found later.

- **Replication sources and destinations**

  You also have the option to choose the source and destinations for this policy. A replication policy can have only one source. It can have one or more destinations.

  A replication source and the destination do not have to be part of a DFS namespace. A multitargeted DFS link cannot be a source in a replication policy. If you have a multitargeted DFS link, you can use that as the destination of the replication policy. In this case, VFM replicates data to each target of the multitargeted DFS link. Even if you take a DFS link target offline, VFM replicates data to that target. If you do not want to replicate data to a specific target, you need to remove the target from the replication policy.

As long as you specify a valid UNC path, a source or a destination path can be a share or a subfolder of a share.

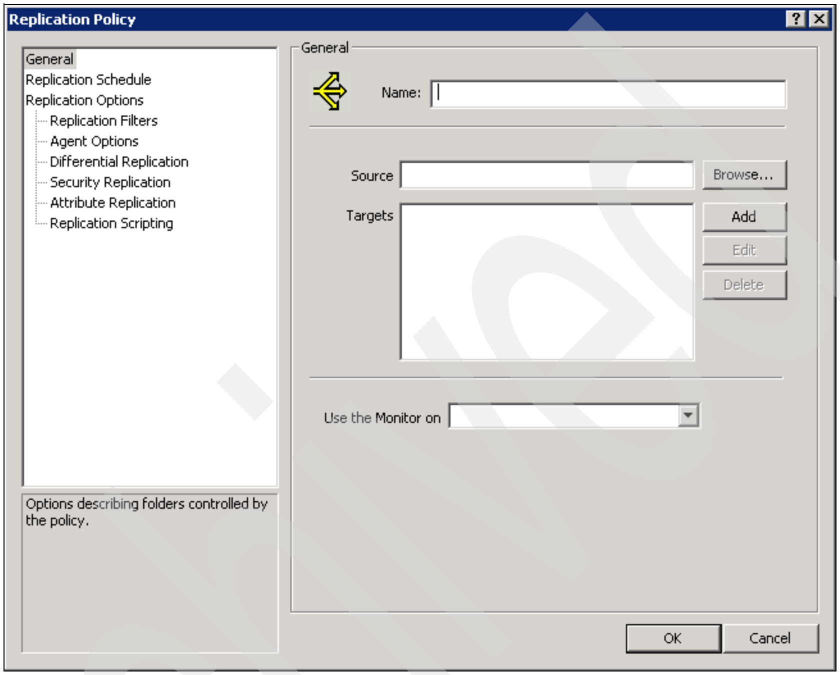

Figure 10-3 Replication Policy: General options

10.1.2 Replication schedule

A replication schedule in a replication policy determines the frequency of replication. In some instances, it can be appropriate to replicate data once a night. In other instances, you might need to replicate the data every four hours or every eight hours. For example, you can create a replication policy called Gold that is scheduled to run every four hours. You can have another replication policy called Silver that is scheduled to run every eight hours. The frequency of replication for any policy must always be based on your business needs and the business value of the data.

Do not schedule the replication to occur more frequently than the time required to complete a run of the replication. For example, if a replication policy takes one hour to complete, do not schedule the replication to run every 30 minutes. If a replication policy is running, and the next scheduled run of the policy elapses, the replication waits until the current run of the policy finishes. If a policy is running, the next run of the policy cannot start until the current run finishes.

The VFM server initiates the policy, so the VFM console does not have to be open. The policy is run based on the time settings and schedule of the VFM server, which might or might not be in the same time zone of the source or destination machines. Take care when setting the run times of a policy and take into account the source and destination time zones and the location of the VFM server.

As seen in Figure 10-4 on page 320, the following options are available on the Replication Schedule panel:

- **Schedule**

  Under the Use replication schedule check box, a drop-down list box enables you to choose the schedule you want to edit. By default one schedule is added when you start the replication policy creation.

- **New and Delete buttons**

  You can set up multiple schedules for a single replication policy. You can use the buttons New and Delete to create or delete a schedule on that policy.

  If you set up more than one schedule for a policy, it runs based on the frequency of each schedule. For example, you can set up multiple schedules to run every 8 hours on weekdays but every 24 hours on weekends.

- **Schedule type**

  With this drop-down list box, you can choose how often your policy runs. The option in the list box range from minutes to months, enabling you to even choose certain days in the week.

- **Use alternate schedules when the policy fails**

  Another advanced scheduling option is enabled through alternate schedules. You can set up an alternate schedule that runs if an error is issued when a certain job runs. This way, a job runs based on the primary schedule. When the job encounters an error (for example, if a network outage occurs or the source or destination reboots), the policy reschedules to run based on the settings for the alternate schedule. When the job runs successfully, scheduling reverts to the primary schedule. This option is useful in a situation where you want to run the job once a day. However, if the job fails, you want it to keep trying every hour until it succeeds.

Follow these steps to create the migration policy and start a scan:

1. On the left panel of the Replication Policy window, click the **Replication Schedule** link (see Figure 10-4).

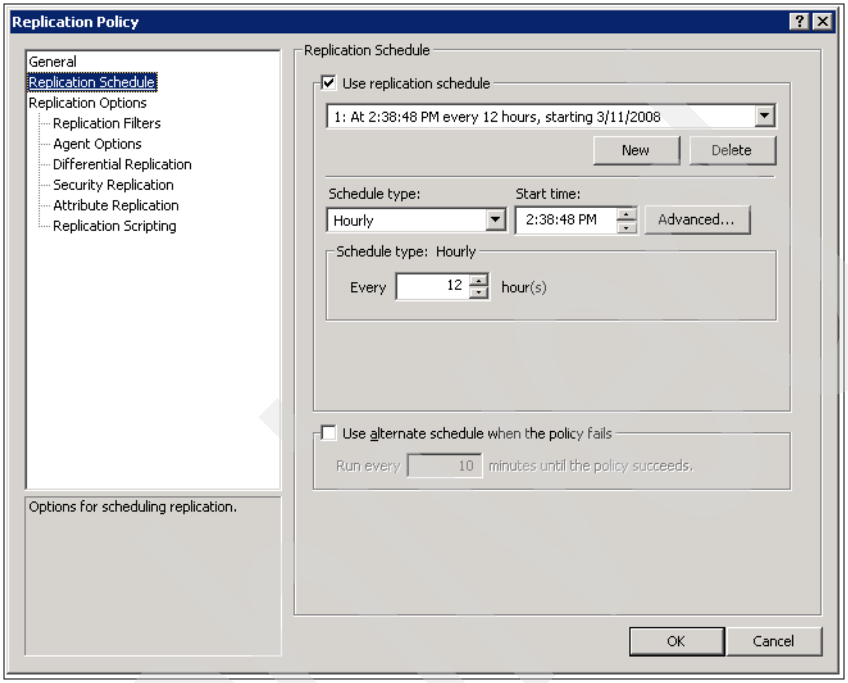

Replication Policy

- General
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options for scheduling replication.

Replication Schedule

☑ Use replication schedule

1: At 2:38:48 PM every 12 hours, starting 3/11/2008

Schedule type: Hourly

Start time: 2:38:48 PM

Schedule type: Hourly

Every 12 hour(s)

□ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

Figure 10-4 Replication Schedule panel

10.1.3 Replication options

On the Replication Options panel, the following options are available (see Figure 10-5 on page 324):

- **Include subfolders**
- **Delete orphaned files in destination folders**

  This option controls what occurs when extra files exist on the destination that do not correspond to any files on the source. If you select to delete the orphaned files, the Replication Agent makes the destination a mirror image of the source.

- **Allow loss of additional file streams (CIFS only)**

  In some instances, the destination file system might not support alternate data streams. Some Microsoft Windows applications store data in alternate data streams. Macintosh-created data files make extensive use of streams (called *resource forks*). In most cases, you might or might not know whether alternate data streams are used. If VFM cannot copy a file because the destination does not support streams, the only way to copy the file is to select the Allow loss of additional file streams check box.

- **Enable Copy-in-place**

  Many optimizations in a data movement policy are a compromise between speed and safety. The Copy-in-place option in the Replication Options panel can help a replication policy finish faster with a slight security risk. Copy-in-place comes into effect during a incremental copy. Assume the following scenario for a replication policy that moves files from one machine to another:

  - A file called xyz.doc is on the source.
  - The file is copied to the destination.
  - The file is modified on the source.
  - Because the file on the source is newer than the destination, xyz.doc on the source must replace the file on the destination.

  If Copy-in-place is not checked, the following events occur:

  - The file xyz.doc is copied to the destination as a temporary file.
  - When the copy is complete, the file xyz.doc is deleted on the destination, and the newly copied temporary file is renamed as xyz.doc.

  If Copy-in-place is checked, the following events occur:

  - The file xyz.doc is truncated to 0 bytes on the destination.
  - The file from the source is copied to the truncated file.

  Checking copy-in-place is faster because no rename operations are required. Copy-in-place is preferred when the source and destination are on the same LAN or connected by a high-speed network link.

  If the source and destination are connected by a low-speed link or the link is not reliable, you must clear the Copy-in-place check box and allow the safe rename technique to work. This way, if the network link breaks when the copy is in progress, you are not left with a zero-length file on the destination.

- **Use NTFS Change Journal (CIFS only)**

  The NTFS change journal is a “file” that maintains a list of accessed, changed, created, and deleted files.

  VFM can optimize the search of changed source files on Microsoft Windows machines by integrating with the NTFS change journal. If the Replication Agent scans the source NTFS change journal, it can detect the changed, deleted, and newly created files in a fraction of the time compared to fully scanning files containing TBs of data. This way, a replication resynchronization proceeds much faster.

> **Note:** The change journal file can wrap, thus resulting in lost changes. The maximum change journal size on the volume is based on space you allocate. If the change journal wraps, VFM resorts to a full scan of the source to determine the changed files.

> **Note:** The NTFS change journal is available only when the source is a Microsoft Windows machine. In all other cases, this option is silently ignored.

- **Use snapshot data**

  While a replication policy runs, one or more files can be open on the source. This is usually the case with PST files and Access databases where the user might have an exclusive lock on the file. If VFM can get a read lock on the file, it copies the file. Depending on how the policy is set up, VFM can either skip the opened file or copy the file as best it can from the snapshot.

  VFM can integrate with the Volume Shadow Service (VSS) on Microsoft Windows 2003 or IBM N series snapshots.

  If you expect that files are open, VFM can integrate with snapshots. VFM can take a temporary snapshot of the source volume and copy all the files from the source snapshot. Once the files are copied, the source snapshot is automatically deleted. If the source is not Windows 2003 or IBM N series, you must use the Retry failed file opens option.

  Files that are not copied are logged as errors and are reported in the manifest based on the settings of the policy. If one or more files are not copied successfully, the policy is reported as a unsuccessful run. This occurs because VFM chooses to warn you that the policy did not copy all files completely.

- **Copy files**

  When a policy runs (either based on the replication schedule or a manual run), VFM examine the source and destination to see which files have to be copied. VFM uses the file name, the date, the attributes, and the file size to determine whether two files are identical. Based on the policy settings, if two files are identical, no data is transferred. If the files on the source and destination are different, the setting in the Copy files drop-down list box determines whether a file is copied. This list box offers the following available options:

- If destination file is different or missing
- If destination is older or missing

At times, the file on the destination might be newer than the source. You must take care when setting the replication options for any policy.

- **Retry failed file opens**

  This option also determines how VFM proceeds when files are open on the source of the replication. This option forces the Replication Agents to retry the open files a number of times, with a specified interval between each try.

- **Abort if no file is successfully opened in**

  If the Replication Agent cannot open files for a certain time, the replication fails. With this option, you set how much time expires before the Replication Agent aborts.

- **Event Details**

  A replication policy can store a list of the files that failed during replication or a list of all files moved while a policy runs.

  We recommend you use the Only list files that encounter errors option, available from the Event Details list box, only when you really have to access a full list of replicated files.

  If you select to list all files during large replications, you can overload the Replication Agent and the VFM server by transferring that much data. Passing the file list of full migrations with millions of files generates a significant amount of network traffic and is an option that we do not recommend. The manifest file is stored on the machine running the Replication Agent and is transferred each time it is requested.

  If files cannot be copied between the source and the target, the policy is considered to complete with errors. Any open files that are not copied result in an error. If an alternate schedule exists, the policy runs based on the alternate schedule. If the VFM server cannot communicate with the Replication Agent performing the data movement, the VFM server forces the policy to an error state.

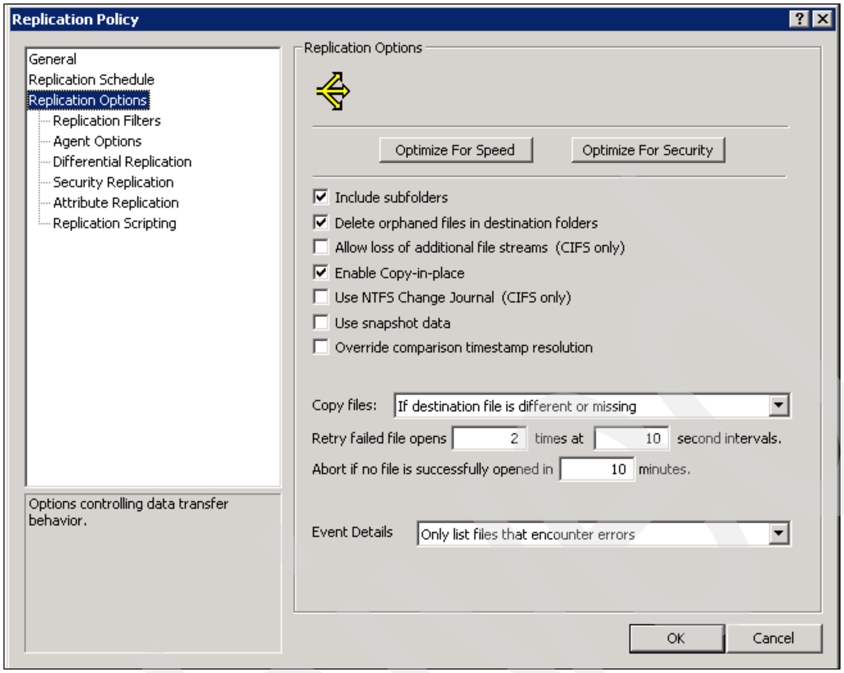

Options controlling data transfer behavior.

☑ Include subfolders
☑ Delete orphaned files in destination folders
☐ Allow loss of additional file streams (CIFS only)
☑ Enable Copy-in-place
☐ Use NTFS Change Journal (CIFS only)
☐ Use snapshot data
☐ Override comparison timestamp resolution

Copy files: If destination file is different or missing

Retry failed file opens 2 times at 10 second intervals.

Abort if no file is successfully opened in 10 minutes.

Event Details Only list files that encounter errors

Figure 10-5 Replication Options panel

10.1.4 Replication filters

On the Replication Filters panel as shown in Figure 10-6 on page 325, you can specify which directories and files you want to exclude from or include in the replication. You can use any DOS wild card characters.

> **Note:** When the replication option 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, but the corresponding files on the destination are not deleted.

For example, when a policy contains an exclusion list and the Delete orphaned files in destination folders option is selected, here is what happens to files. In this example, 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.

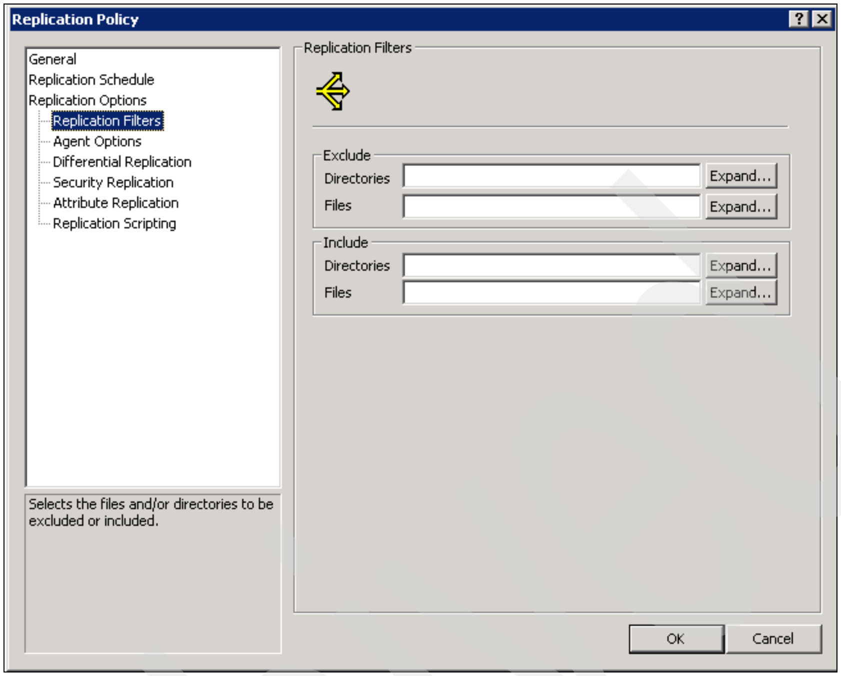

Selects the files and/or directories to be excluded or included.

Figure 10-6 Replication Filters panel

10.1.5 Agent options

On the Agent Options panel (see Figure 10-7 on page 327), you can choose where the Replication Agent runs. By default, VFM uses the Replication Agent on the destination to pull data from the source. This method is preferred because the agent on the destination knows how much data it can handle based on the load on the destination machine.

The following options are available from this panel:

- **On the source side**

  When the Replication Agent cannot run on the destination (for example, the destination is an IBM N series machine), you can use this option to force the agent to run on the source and push data to the destination.

- **On the destination side**

  This is the default behavior when you choose the VFM selects option (further down on the Agent Options panel), and this method is preferred because the agent on the destination knows how much data it can handle based on the load on the destination machine. With the On the destination side option, you can force this behavior, or make it clear that it is occurring.

- **On this machine or corresponding proxy**

  When Replication Agents cannot be deployed on the source or destination, you can use a Replication Agent on another machine to act like a proxy and make the replication possible.

  A typical scenario in which this option is helpful is when you want to migrate data from one N series to another. Because Data ONTAP does enable agents to be installed, you need a proxy. If you choose the VFM selects option (further down on the Agent Options panel) as the chosen option, VFM can automatically select a machine to act as the proxy, but it does not consider any geographical distance, link cost, or bandwidth because it does not have this information.

  As a best practice if Replication Agents cannot be deployed on the source or the destination, you must manually select a machine to act as a proxy.

  Remember that you do not need to manually install Replication Agents. Based on policy settings, the Replication Agent is automatically deployed to the specified machine, no matter whether it is a source, destination, or proxy.

- **In this agent group**

  Agent grouping provides a way to share the load between multiple Replication Agents. When the replication job runs, VFM sends the policy to the next available agent in the group. A round robin technique is used to determine which Replication Agent in the agent group performs the policy. Using the In this agent groups option prevents a Replication Agent from being overloaded, while sharing the load across multiple agents to perform replications.

- **VFM selects**

  This option is the default. Using this option, you can transfer the responsibility of choosing where to run the agent to VFM.

  VFM first tries the destination; if it fails, it tries the source. If it cannot use either source or destination, it uses the Replication Agent on the machine where VFM server is installed. Check the On this machine or corresponding proxy radio button as a best practice when you need a proxy.

Remember that Replication Agents run under the credentials of the service account that VFM runs under. This service account must have administrator rights on the source and destination machines.

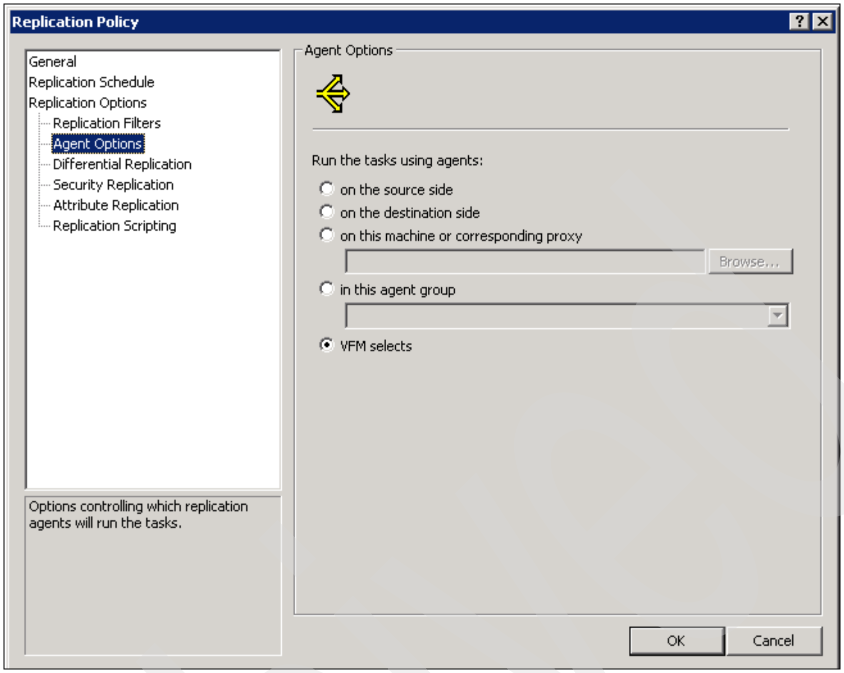

Replication Policy

- General
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options controlling which replication agents will run the tasks.

Agent Options

Run the tasks using agents:

○ on the source side
○ on the destination side
○ on this machine or corresponding proxy
○ in this agent group
● VFM selects

Figure 10-7 Agent options

10.1.6 Differential replication and byte level replication (BLR)

The options on the Differential Replication panel are shown in Figure 10-8 on page 328.

VFM provides the feature of byte level replication (BLR) to control data movement across slow links. BLR minimizes the data sent across the wire, especially with slow network links, by sending only the changed bytes.

You have a 10 MB Microsoft PowerPoint® file copied from the source to the destination share. Let us assume that the user at the source changes one or two slides in the presentation. Instead of sending the 10 MB changed file again from the source to the destination, VFM can be used to intelligently send the changes across the wire.

While the time required in computing the changed bytes in extremely large files can be significant, network usage is minimized. You sacrifice the extra CPU cycles to figure out the changed bytes so that the limited resource of network bandwidth is optimized. BLR requires two Replication Agents to work together to make the replication complete successfully.

Use BLR only across slow links. In LAN-speed networks, the overhead of BLR computation is more than the benefit it provides. BLR is also not effective for graphic files (such as JPG and BMP) where a single change in images is scattered in all parts of the file.

BLR works when the destination Replication Agent computes rolling hash sets across all the blocks of the file at the destination. The destination Replication Agent sends the rolling hash set as a file signature to the source Replication Agent. This results in a small amount of data being transferred from the destination to the source.

The source Replication Agent computes the rolling hash set of the changed file at the source. When different hash blocks are encountered, the source Replication Agent computes an instruction set to send from the source to the destination. The instruction set is used to insert or remove data blocks from the file at the destination to make it look like the source. BLR over slow WAN links results in an 80% savings in network traffic of slightly changed data files.

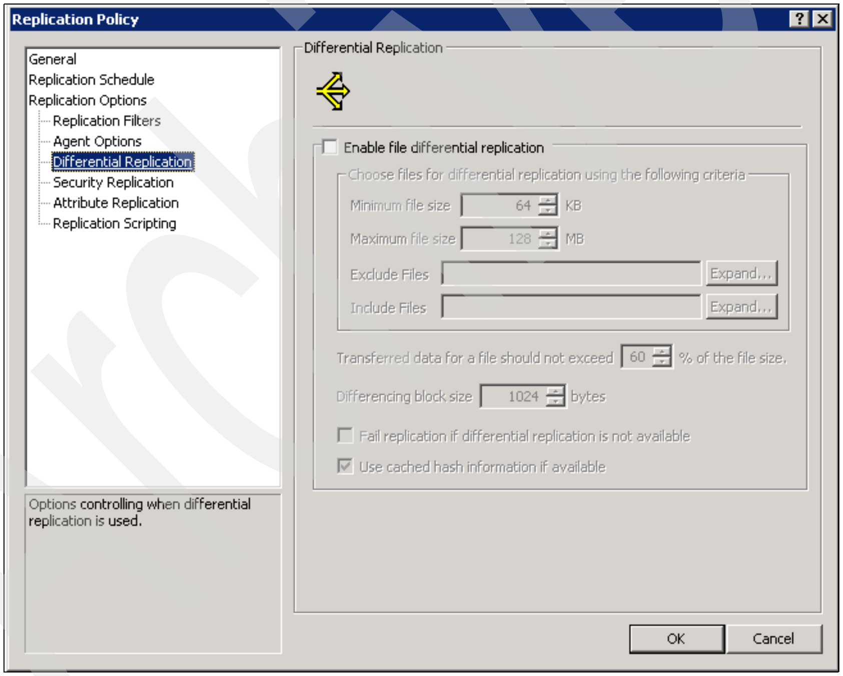

Options controlling when differential replication is used.

Differential Replication

□ Enable file differential replication

Choose files for differential replication using the following criteria

Minimum file size 64 KB

Maximum file size 128 MB

Exclude Files

Include Files

Transferred data for a file should not exceed 60 % of the file size.

Differencing block size 1024 bytes

□ Fail replication if differential replication is not available

☑ Use cached hash information if available

Figure 10-8 Differential Replication panel

10.1.7 Security replication

This section discusses the options available on the Security Replication panel as shown in Figure 10-9 on page 330.

Every replication requires that the files be copied accurately and completely. VFM provides several mechanisms to control file copying. Where possible, VFM signal errors when files were not copied completely or correctly.

Many Active Directory deployments can use local groups. A local group is valid only on the source system on which it is defined. If the security on a file is controlled by the permissions of a local groups and that file is copied to a destination machine, the security settings are not valid. This means that the users who had access to the files on the old server do not have access to the replicated copy.

The following options determine how VFM handles security replication (see Figure 10-9 on page 330):

- **Copy security descriptor**

  From the Copy security descriptor drop-down list box, you can choose when to copy the security information from files:

  - Each time the file or folder is copied.
  - Only if target file or folder does not exist.

    This means that the security information is copied only the first time the file is replicated. If you choose this option, be sure that it is appropriate for your environment.
  - Never.
  - Always copy security settings.
- **Remove the SID from the security descriptor**

  With this option, you can direct VFM to remove just the local groups from the file on the destination.
- **Translate the SID to the same trustee name and type on the target**

  VFM can also translate the security information for local groups. You can click the following radio boxes:

  - **Create local groups on the target if not found**

    With this option selected, VFM automatically creates a local group on the destination with the same name as the local group on the source.

    All domain users and domain groups that are members of the old local group are added to the newly created local group. This way, users that had access to the data through the old local group on the source server have access to the data on the new server through the newly created local group.

– **If unable to translate the SID**

  In case the SID cannot be translated, the chosen action is taken.

- **Allow loss of security information**

  At times, you care only about replicating the data from one location to another without regard to the security of the data on the new system. For example, if you replicate data from a Microsoft Windows server to a UNIX-style qtree on the new IBM N series, we recommend you check the Allow loss of security information check box.

  If the source machine or the destination machine is a Domain Controller, you must deselect this check box to process local groups (technically called *local trustees*). You do so because the Domain Controller does not support local groups.

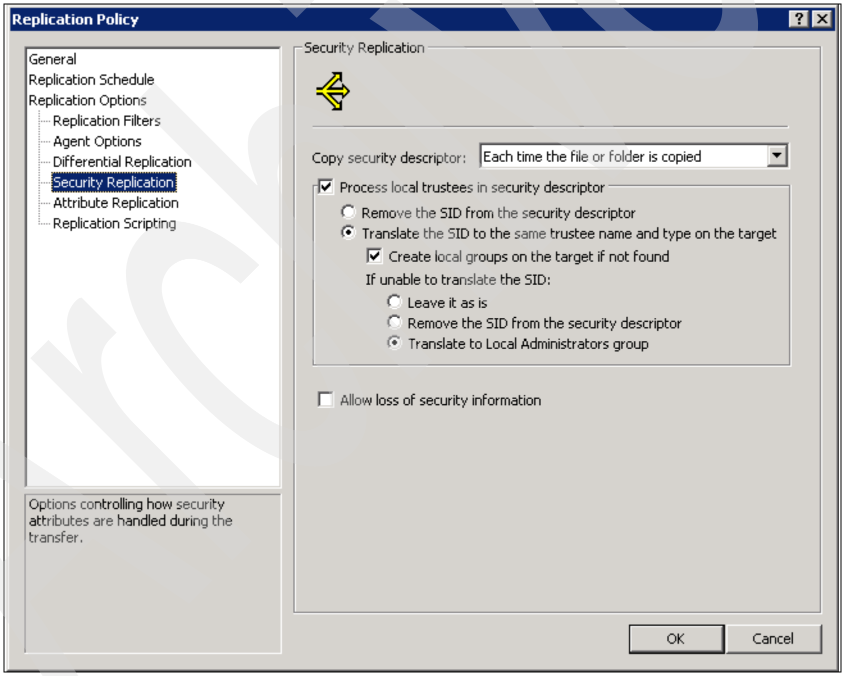

Options controlling how security attributes are handled during the transfer.

Security Replication

Copy security descriptor: Each time the file or folder is copied

☑ Process local trustees in security descriptor

○ Remove the SID from the security descriptor

● Translate the SID to the same trustee name and type on the target

☑ Create local groups on the target if not found

If unable to translate the SID:

○ Leave it as is

○ Remove the SID from the security descriptor

● Translate to Local Administrators group

□ Allow loss of security information

Figure 10-9 Security Replication panel

10.1.8 Attribute replication

In this section, we discuss the Attribute Replication options as shown in Figure 10-10. Using these options, you can change the file attributes to match your needs. We provide the following two examples:

- You might want to clear the archive attribute of a copied file so that the file can be backed up on the replication target.
- You might want to preserve the last access time of the file on the source.

Even though VFM copies the file, it resets the last accessed time on the file. This is useful when HSM or other reporting software is based on the last access time. Overhead is involved on the Replication Agent when preserving the last access time. So if you do not require this feature, uncheck the Preserve last access time on source check box.

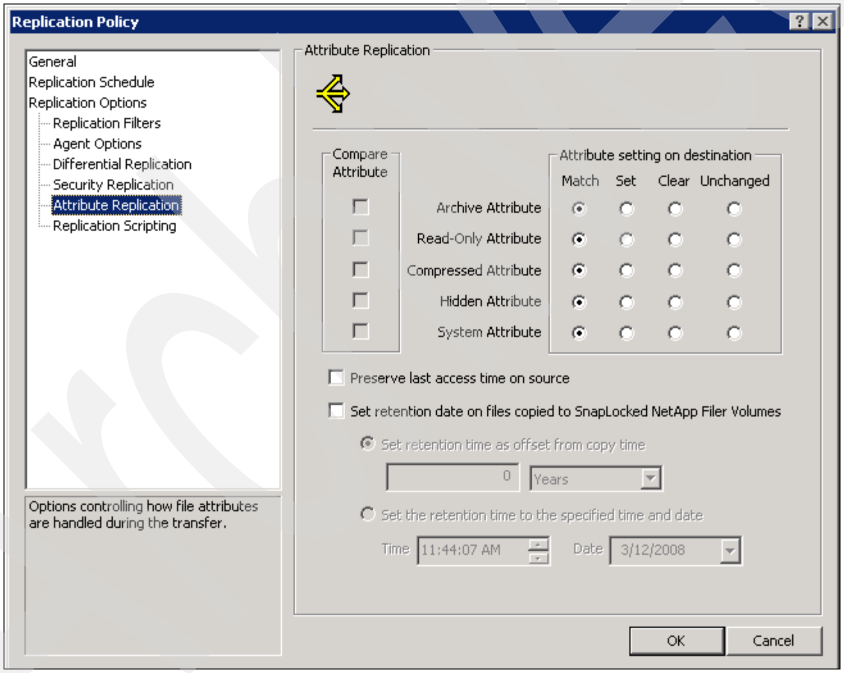

Options controlling how file attributes are handled during the transfer.

Attribute Replication

Compare Attribute		Attribute setting on destination			
		Match	Set	Clear	Unchanged
□	Archive Attribute	●	○	○	○
□	Read-Only Attribute	●	○	○	○
□	Compressed Attribute	●	○	○	○
□	Hidden Attribute	●	○	○	○
□	System Attribute	●	○	○	○

□ Preserve last access time on source

□ Set retention date on files copied to SnapLocked NetApp Filer Volumes

● Set retention time as offset from copy time

0 Years

○ Set the retention time to the specified time and date

Time 11:44:07 AM Date 3/12/2008

Figure 10-10 Attribute Replication panel

10.1.9 Replication scripting

You can configure VFM to enable you to run a script before or after a replication policy (see Figure 10-11). For example, you can choose to run the Run Batch Before Replication script to stop an application or send an e-mail. The Run Batch After Replication script enables you to restart the application or send an e-mail.

Replication scripts are passed parameters and can be any batch script. The scripts are run by the VFM server under the context of the VFM service account. The scripts can be located on the local drive or from any network-accessible share.

Note: The Run Batch After Replication script cannot run if the replication policy contains an error. A script runs once for each policy. Even if the policy has multiple hops, the script only runs once.

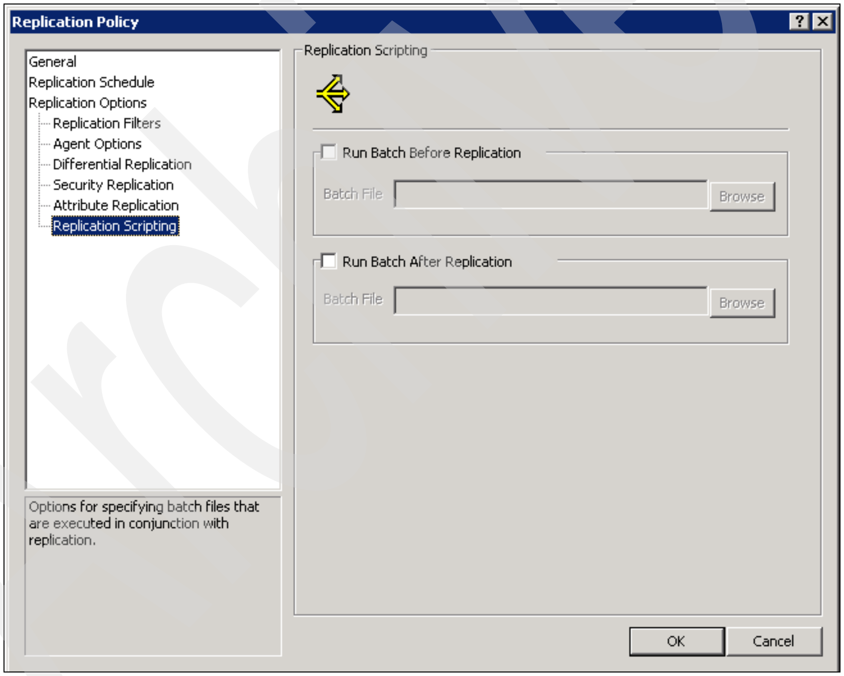

Options for specifying batch files that are executed in conjunction with replication.

☐ Run Batch Before Replication

Batch File

☐ Run Batch After Replication

Batch File

Figure 10-11 Replication Scripting panel

10.2 Typical examples of replication setup

Because of the numerous options available with replication, it is not possible to show all the possible configurations of replication. This section provide examples of typical replication setups.

10.2.1 Data gathering

Suppose company XYZ has several remote offices around the world, but one data center is responsible for all backups. With this configuration, the remote offices do not have to maintain tapes and specific staff for backups.

Every night the data on several shares is replicated from each office to the main data center. In this example, we show how to set up replication from two different locations to a single CIFS share on a N series storage system.

Company XYZ has the following three machines:

- roman.itso.tucson is the office server located in Rome.
- itsotuc6.itso.tucson is the São Paulo office server.
- itsotuc3.itso.tucson is the N series Storage System located in Tucson that centralizes the data of the Rome and São Paulo offices.

On N series is the share Remote Offices, which contains one folder for each remote office as shown in Figure 10-12.

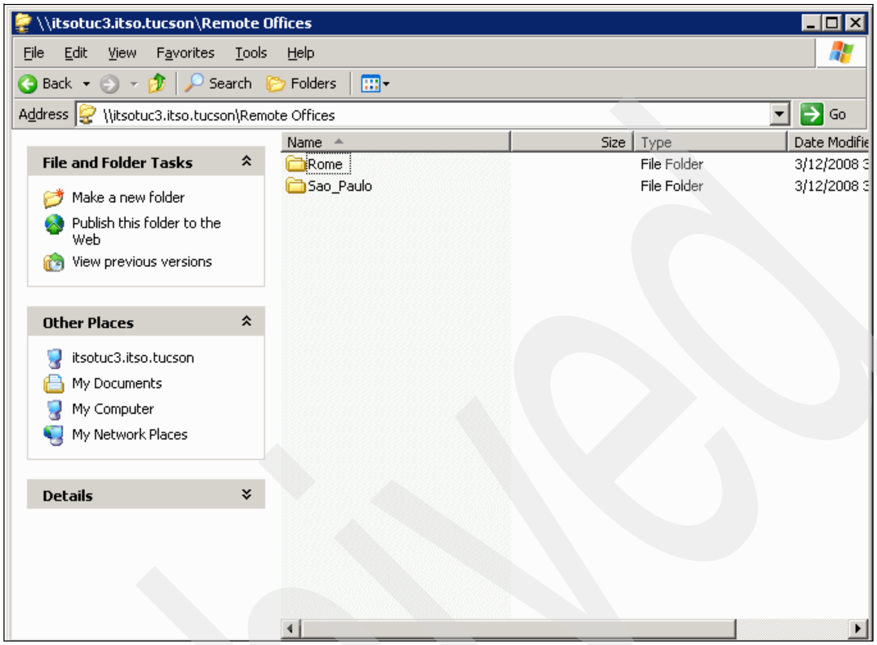

Name	Size	Type	Date Modifie
Rome		File Folder	3/12/2008 3
Sao_Paulo		File Folder	3/12/2008 3

Figure 10-12 Data gathering: directory structure on N series share

The server in Rome contains the Important Data share as shown in Figure 10-13. The office server in São Paulo has the share Main Data with the data displayed in Figure 10-14. All data in both shares must be replicated to the corresponding folder on the N series server. In this case, you must create one replication policy for each remote office whose data you want to replicate.

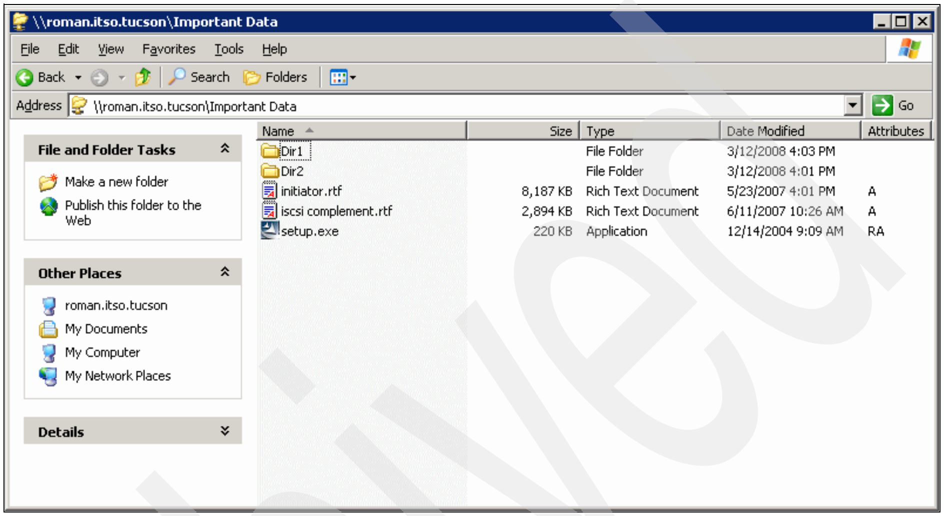

Name	Size	Type	Date Modified	Attributes
Dir1		File Folder	3/12/2008 4:03 PM	
Dir2		File Folder	3/12/2008 4:01 PM	
initiator.rtf	8,187 KB	Rich Text Document	5/23/2007 4:01 PM	A
iscsi complement.rtf	2,894 KB	Rich Text Document	6/11/2007 10:26 AM	A
setup.exe	220 KB	Application	12/14/2004 9:09 AM	RA

Figure 10-13 Rome server: Important Data contents

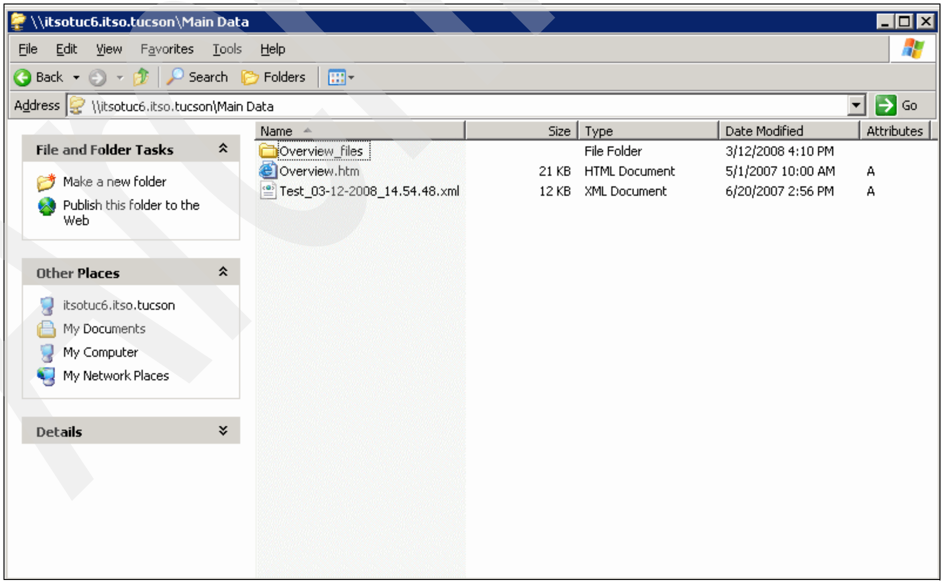

Name	Size	Type	Date Modified	Attributes
Overview_files		File Folder	3/12/2008 4:10 PM	
Overview.htm	21 KB	HTML Document	5/1/2007 10:00 AM	A
Test_03-12-2008_14.54.48.xml	12 KB	XML Document	6/20/2007 2:56 PM	A

Figure 10-14 São Paulo server: Main Data contents

Important: Before setting a replication policy, make sure that the service account used by VFM has Full Control set on all involved shares.

Creating the replication policy for the office server in Rome

To create a new replication policy, go to **Admin View** → **Data Movement Policies** → **Replication Policies** (see Figure 10-15).

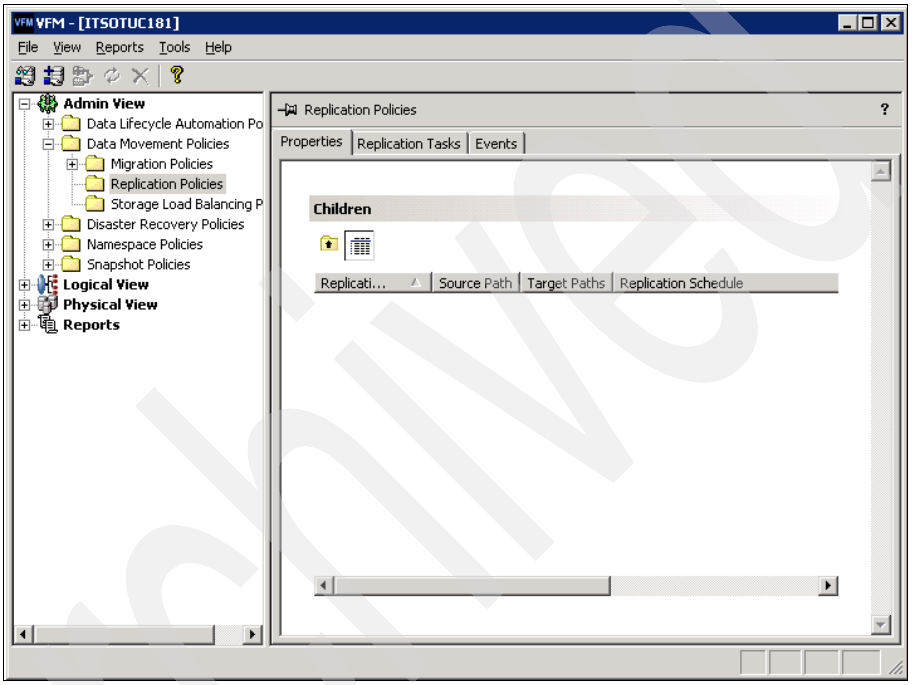

Figure 10-15 VFM client: Replication Policies panel

Follow these steps to create a new replication policy:

1. Right-click **New Replication Policy** (see Figure 10-16).

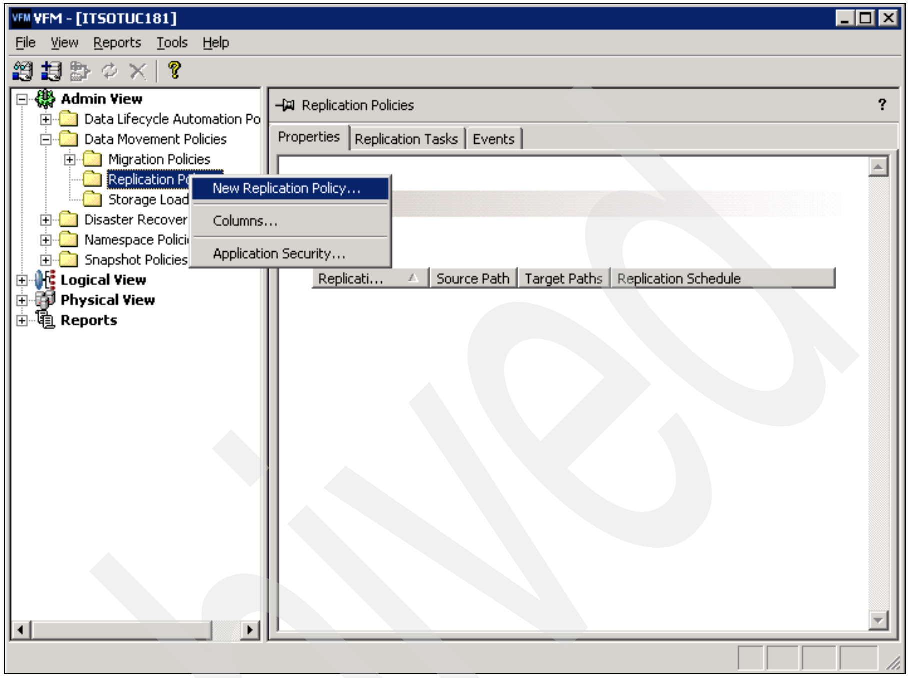

Figure 10-16 Replication Policies window with right-click menu

2. At this point, you can set the first parameters for this replication: Name, Source, and Targets. Fill in these three fields as shown in Figure 10-17.

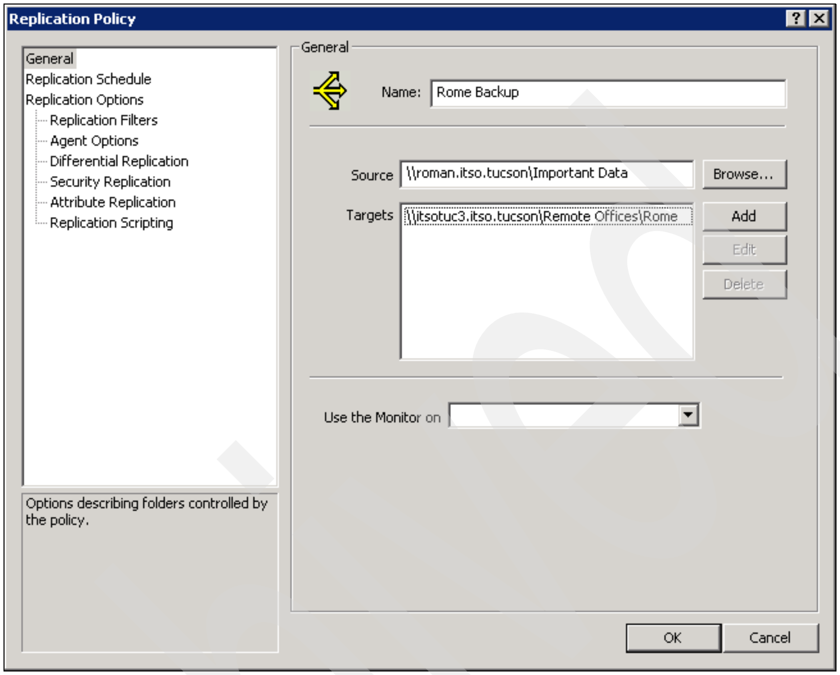
Replication Policy

Name: Rome Backup

Source \\roman.itso.tucson\Important Data

Targets \\itsotuc3.itso.tucson\Remote Offices\Rome

Use the Monitor on

Options describing folders controlled by the policy.

Figure 10-17 Rome server: new replication policy

3. In the panel on the left of the Replication Policy window, click the **Replication Schedule** link. In this example, we choose to perform a daily replication, starting at 10 p.m. (the time on the VFM server) as shown in Figure 10-18.

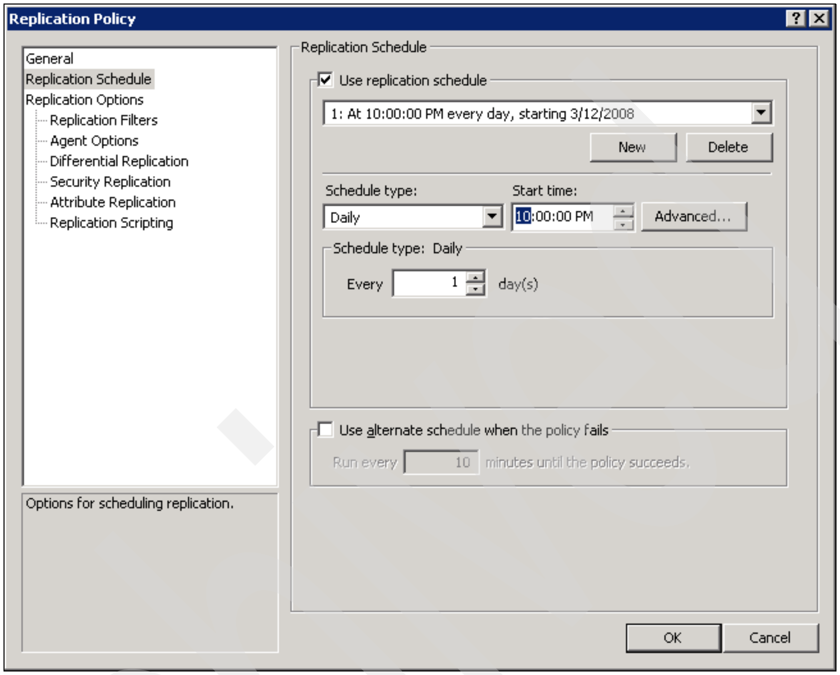

☑ Use replication schedule

1: At 10:00:00 PM every day, starting 3/12/2008

Schedule type: Daily

Start time: 10:00:00 PM

Schedule type: Daily

Every 1 day(s)

□ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

Figure 10-18 Rome server: replication schedule

4. In the Replication Options panel, check the values as shown in Figure 10-19.

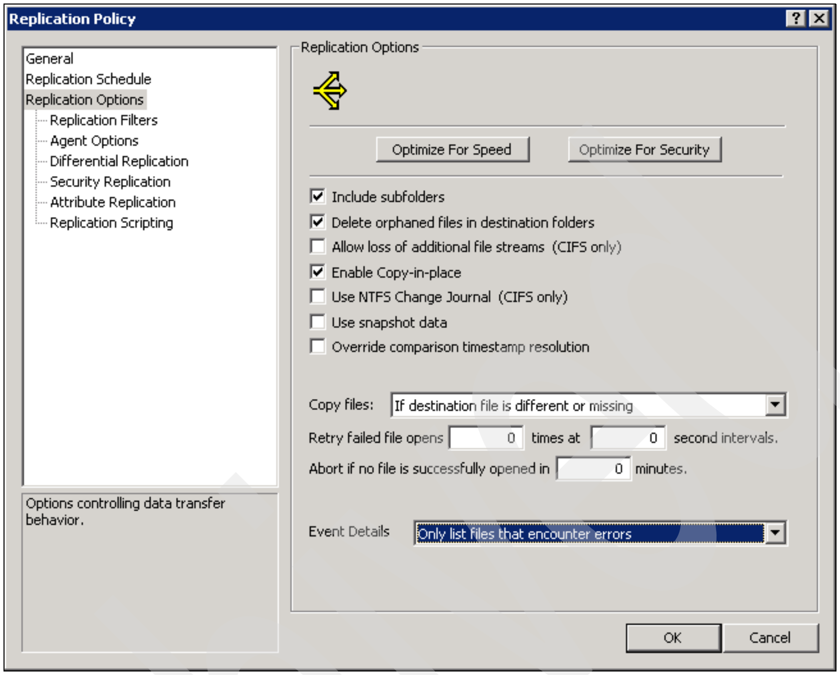

Replication Policy

General
Replication Schedule
Replication Options
- Replication Filters
- Agent Options
- Differential Replication
- Security Replication
- Attribute Replication
- Replication Scripting

Options controlling data transfer behavior.

Replication Options

Optimize For Speed | Optimize For Security

☑ Include subfolders
☑ Delete orphaned files in destination folders
☐ Allow loss of additional file streams (CIFS only)
☑ Enable Copy-in-place
☐ Use NTFS Change Journal (CIFS only)
☐ Use snapshot data
☐ Override comparison timestamp resolution

Copy files: If destination file is different or missing

Retry failed file opens 0 times at 0 second intervals.

Abort if no file is successfully opened in 0 minutes.

Event Details Only list files that encounter errors

OK | Cancel

Figure 10-19 Rome server: replication options

5. Click **OK** to create the policy and check that the policy is listed in Replication Policies as shown in Figure 10-20. In this case, we do not have to change any other value.

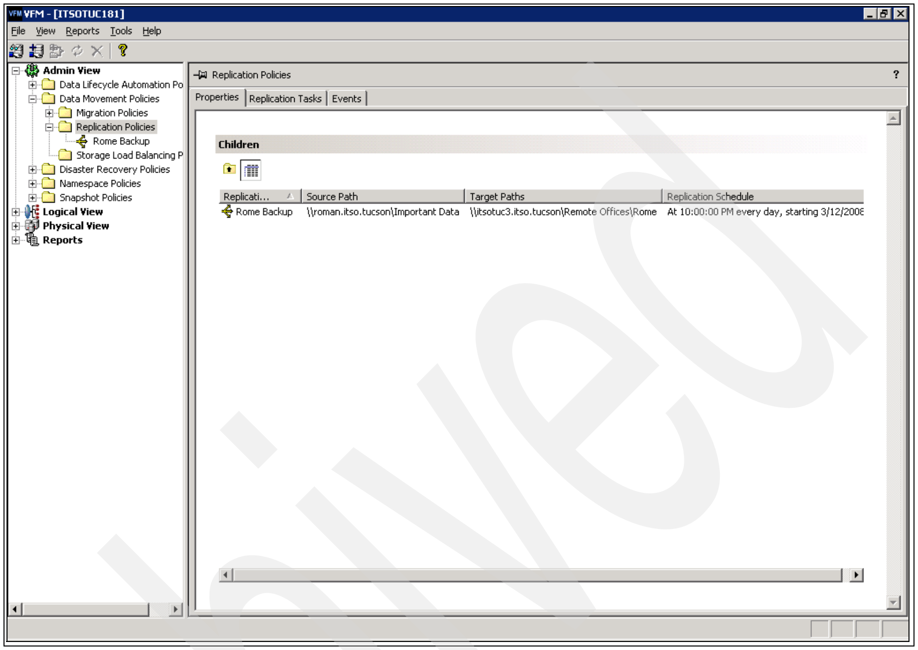

Replicati...	Source Path	Target Paths	Replication Schedule
Rome Backup	\\roman.itso.tucson\Important Data	\\itsotuc3.itso.tucson\Remote Offices\Rome	At 10:00:00 PM every day, starting 3/12/2008

Figure 10-20 Rome policy: replication policies list

6. Click the **Roman Backup** policy to view its properties (see Figure 10-21).

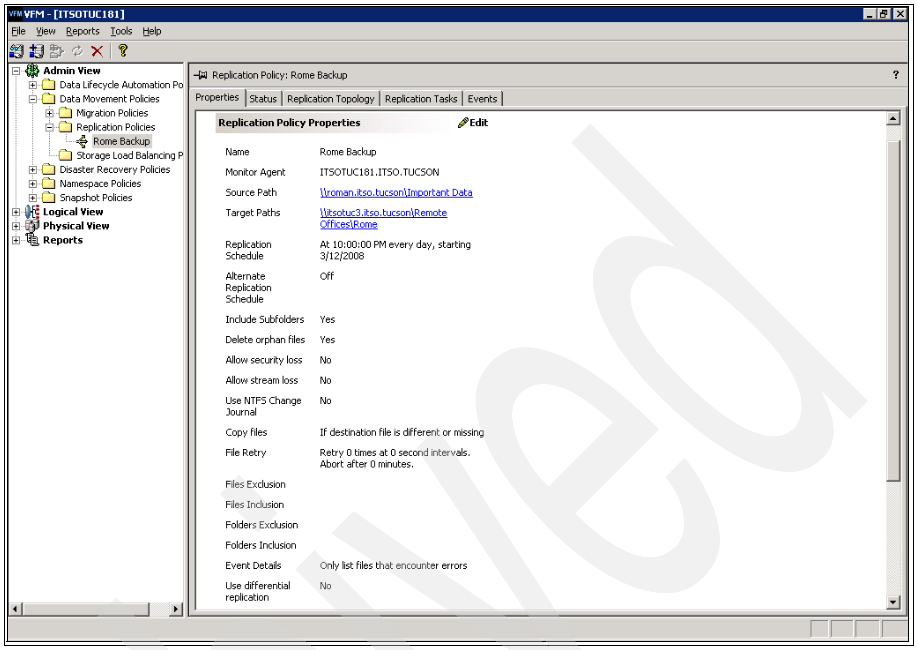

Replication Policy Properties

Name	Rome Backup
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Source Path	\\roman.itso.tucson\Important Data
Target Paths	\\itsotuc3.itso.tucson\Remote Offices\Rome
Replication Schedule	At 10:00:00 PM every day, starting 3/12/2008
Alternate Replication Schedule	Off
Include Subfolders	Yes
Delete orphan files	Yes
Allow security loss	No
Allow stream loss	No
Use NTFS Change Journal	No
Copy files	If destination file is different or missing
File Retry	Retry 0 times at 0 second intervals. Abort after 0 minutes.
Files Exclusion	
Files Inclusion	
Folders Exclusion	
Folders Inclusion	
Event Details	Only list files that encounter errors
Use differential replication	No

Figure 10-21 Rome backup properties

7. Click the **Replication Topology** tab to view a graphical representation of the topology of this replication policy (see Figure 10-22).

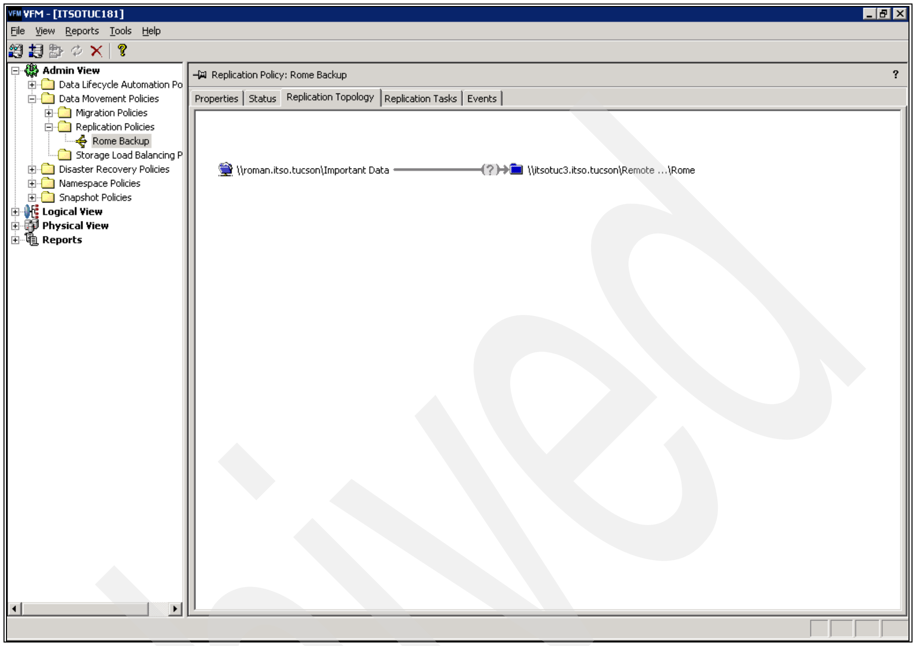

Figure 10-22 Rome backup topology

8. In the left panel, right-click the name of the policy. Then click **Start Replication now** (see Figure 10-23).

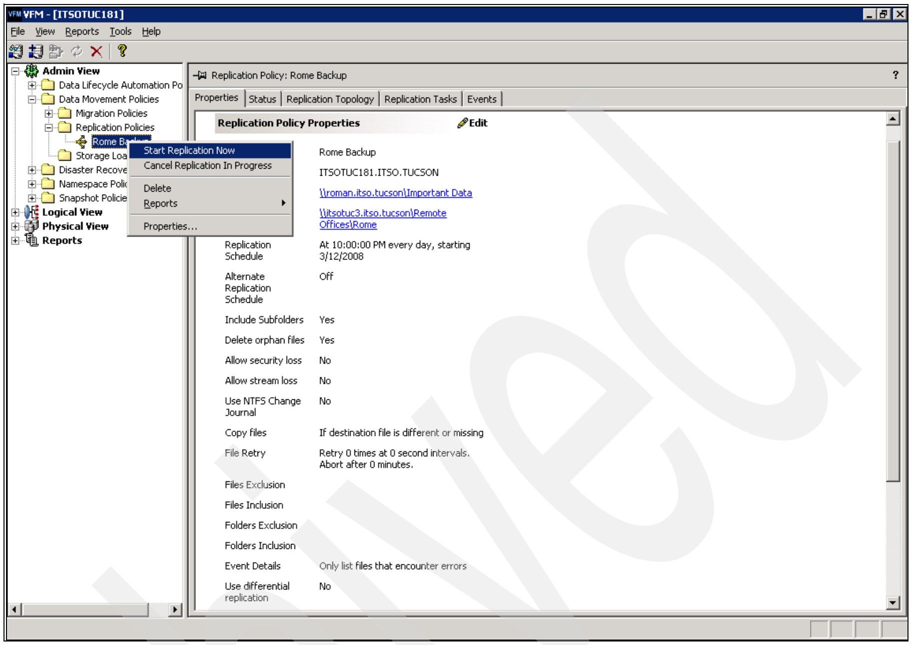

Replication Policy Properties

	Rome Backup
	ITSOTUC181.ITSO.TUCSON
	\\roman.itso.tucson\Important Data
	\\itsotuc3.itso.tucson\Remote Offices\Rome
Replication Schedule	At 10:00:00 PM every day, starting 3/12/2008
Alternate Replication Schedule	Off
Include Subfolders	Yes
Delete orphan files	Yes
Allow security loss	No
Allow stream loss	No
Use NTFS Change Journal	No
Copy files	If destination file is different or missing
File Retry	Retry 0 times at 0 second intervals. Abort after 0 minutes.
Files Exclusion	
Files Inclusion	
Folders Exclusion	
Folders Inclusion	
Event Details	Only list files that encounter errors
Use differential replication	No

Figure 10-23 Rome server: starting replication

9. Go back to the Replication Topology tab (see Figure 10-24). The blue circle indicates that the replication is being processed.

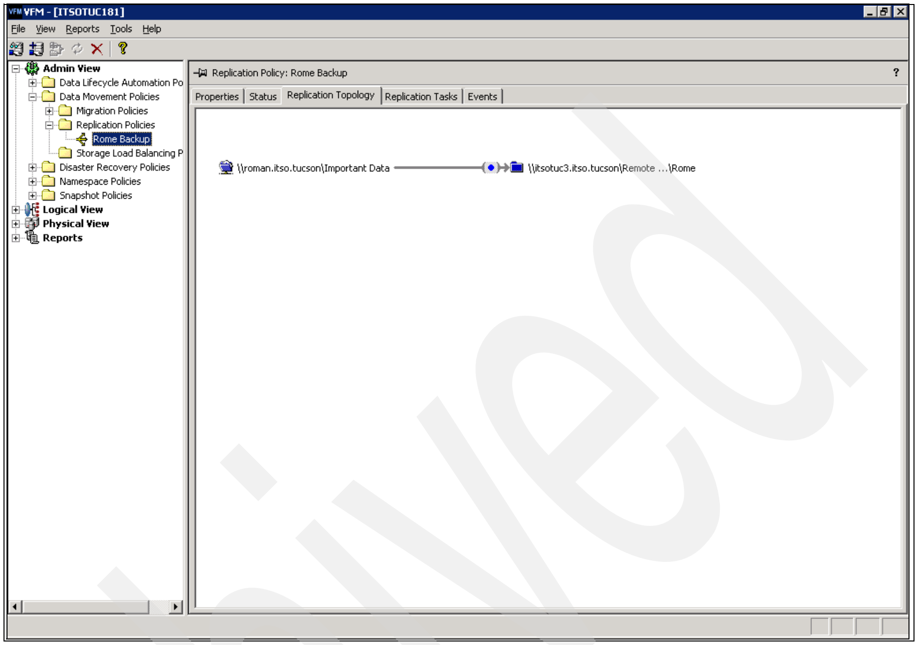

\\roman.itso.tucson\Important Data → \\itsotuc3.itso.tucson\Remote ...\Rome

Figure 10-24 Replication topology showing a ongoing replication

When the replication finishes, the final status is displayed as shown in Figure 10-25.

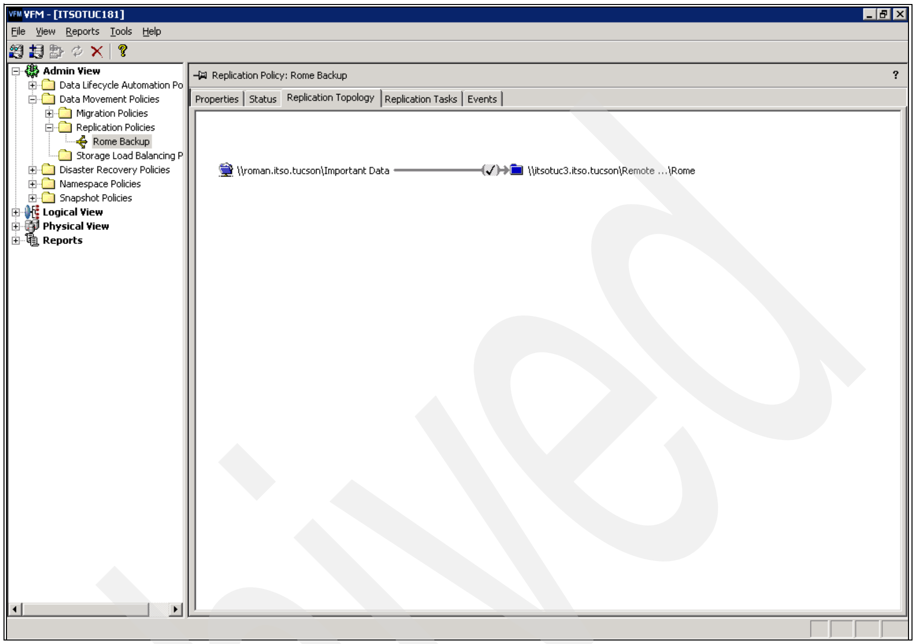

Figure 10-25 Rome replication topology after a successful replication

10. Access **Replication Tasks** (see Figure 10-26). Notice the contents of this window; one task for this replication policy is displayed.

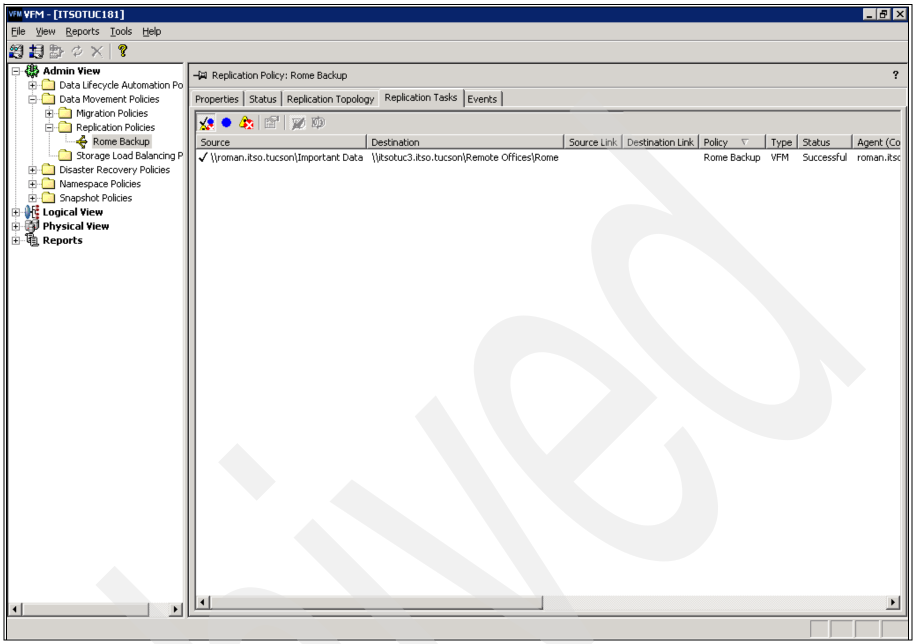

Source	Destination	Source Link	Destination Link	Policy	Type	Status	Agent (Co
✓ \\roman.itso.tucson\Important Data	\\itsotuc3.itso.tucson\Remote Offices\Rome			Rome Backup	VFM	Successful	roman.itsc

Figure 10-26 Rome backup replication policy

11. Check the contents of the destination share (see Figure 10-27). You can check whether it is the same as the original content (see Figure 10-13 on page 335).

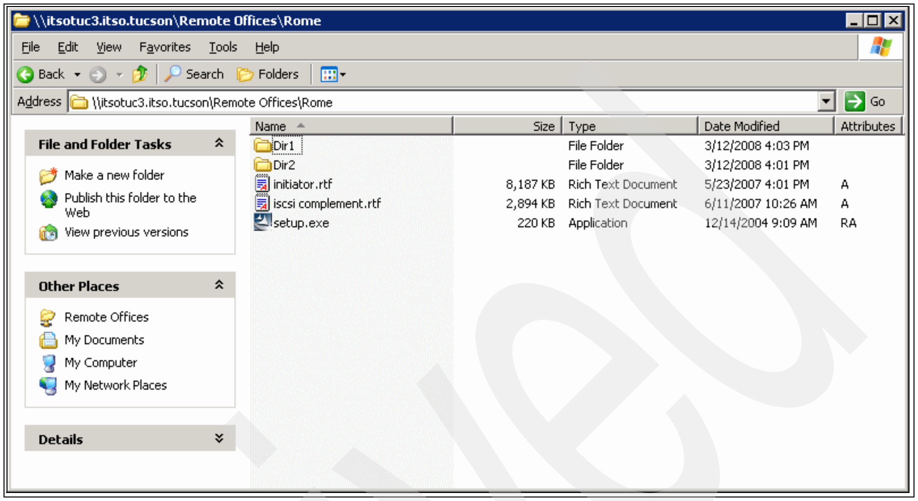

Name	Size	Type	Date Modified	Attributes
Dir1		File Folder	3/12/2008 4:03 PM	
Dir2		File Folder	3/12/2008 4:01 PM	
initiator.rtf	8,187 KB	Rich Text Document	5/23/2007 4:01 PM	A
iscsi complement.rtf	2,894 KB	Rich Text Document	6/11/2007 10:26 AM	A
setup.exe	220 KB	Application	12/14/2004 9:09 AM	RA

Figure 10-27 Rome backup: contents on destination

Creating the replication policy for the office server in São Paulo

In this section, we repeat the process for the remote office in São Paulo:

1. Right-click **Replication Policy** (see Figure 10-28).

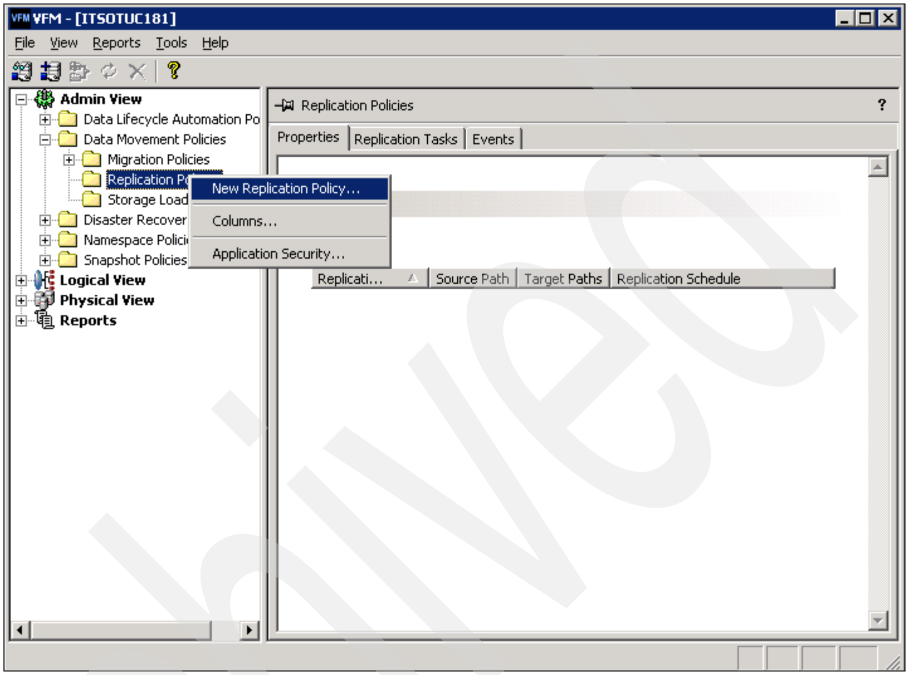

Figure 10-28 Replication Policies window with right-click menu

2. Set the parameters for this replication: **Name**, **Source**, and **Targets**. Fill in those three fields as shown in Figure 10-29.

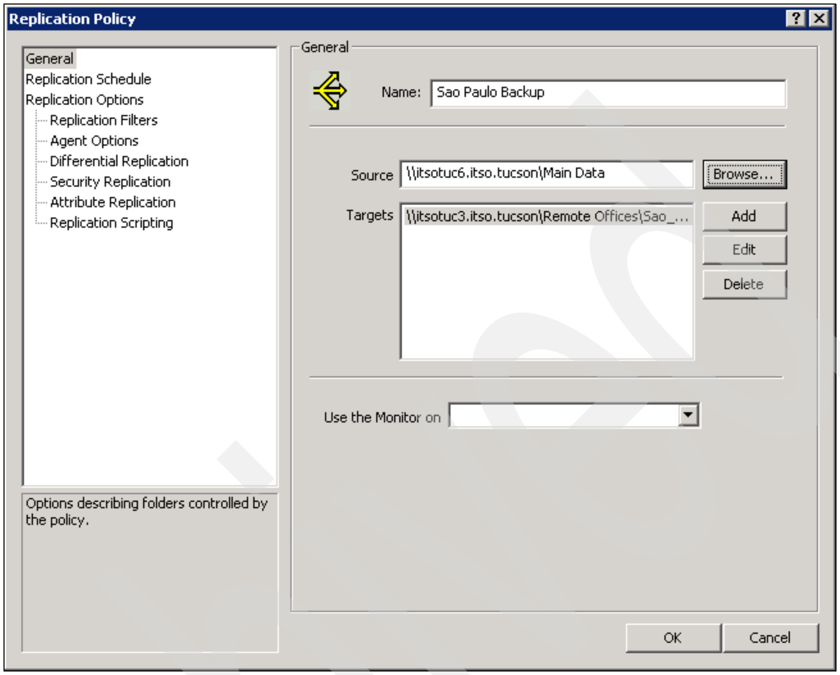

Options describing folders controlled by the policy.

General

Name: Sao Paulo Backup

Source \\itsotuc6.itso.tucson\Main Data

Targets \\itsotuc3.itso.tucson\Remote Offices\Sao_...

Use the Monitor on

Figure 10-29 São Paulo server: new replication policy

3. On the left of the Replication Policy window, choose **Replication Schedule**. In this example, we choose to perform a daily replication, starting at 1 p.m. (the time on the VFM server) as shown in Figure 10-30.

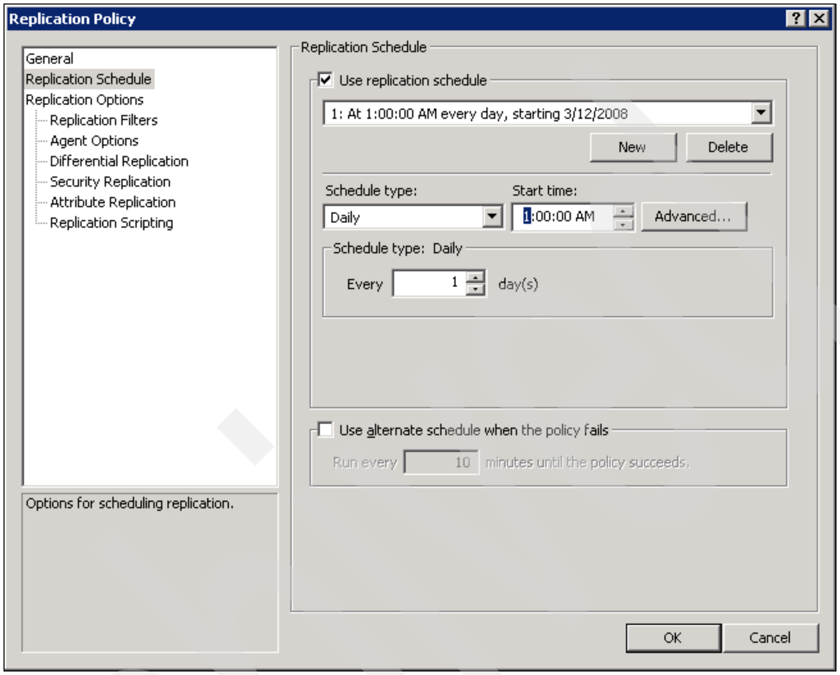

Replication Policy

- General
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options for scheduling replication.

Replication Schedule

☑ Use replication schedule

1: At 1:00:00 AM every day, starting 3/12/2008

Schedule type: Daily

Start time: 1:00:00 AM

Schedule type: Daily

Every 1 day(s)

□ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

Figure 10-30 São Paulo server: replication schedule

4. In the Replication Options panel, check the values as shown in Figure 10-31.

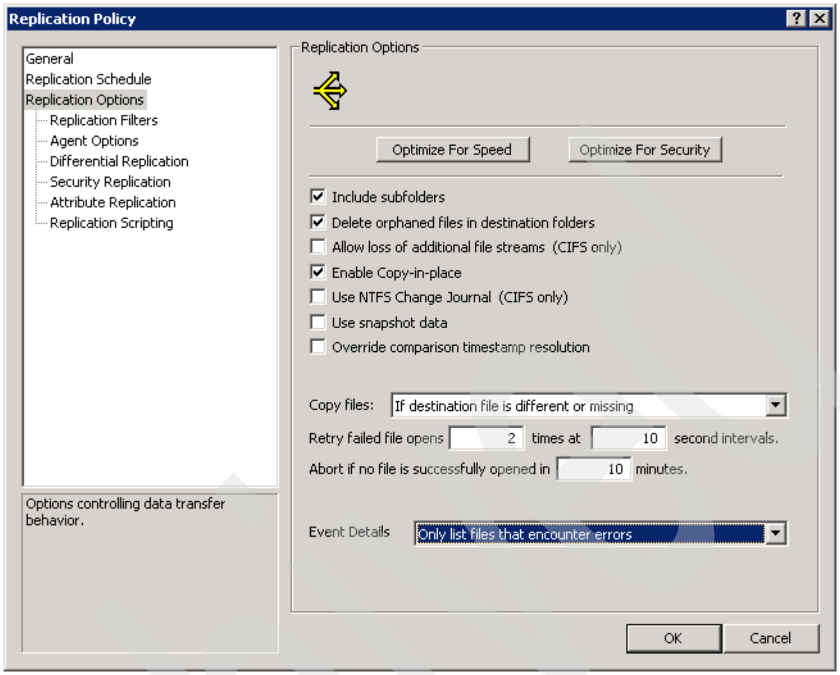

Options controlling data transfer behavior.

Replication Options

☑ Include subfolders
☑ Delete orphaned files in destination folders
☐ Allow loss of additional file streams (CIFS only)
☑ Enable Copy-in-place
☐ Use NTFS Change Journal (CIFS only)
☐ Use snapshot data
☐ Override comparison timestamp resolution

Copy files: If destination file is different or missing

Retry failed file opens 2 times at 10 second intervals.

Abort if no file is successfully opened in 10 minutes.

Event Details Only list files that encounter errors

Figure 10-31 São Paulo server: replication options

5. In this example, we consider other options. Click the **Security Replication** link and check whether the Process local trustees in security descriptor check box is checked. If it is checked, uncheck this check box. Refer to 10.1, “Using replication options” on page 316 for details about this option.

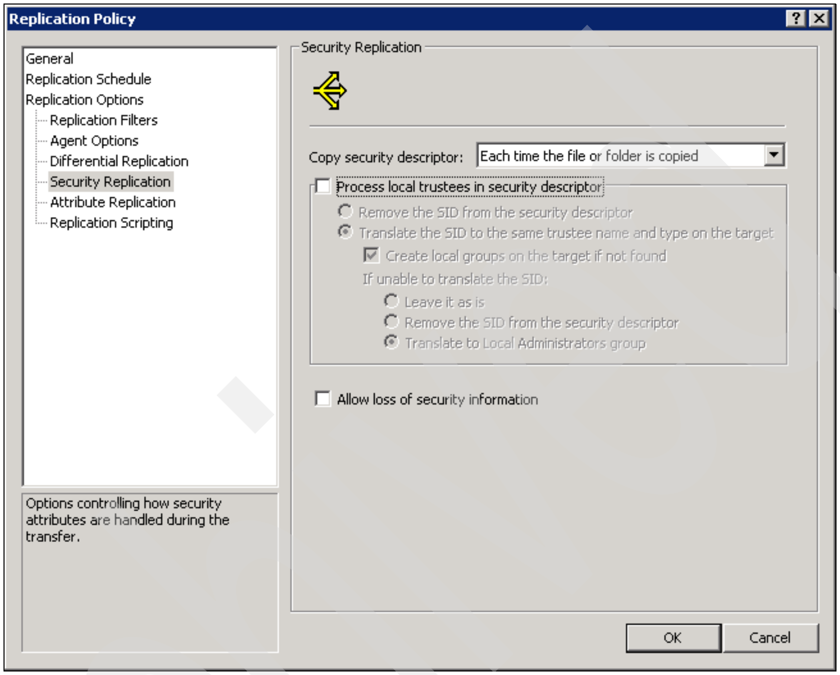

Options controlling how security attributes are handled during the transfer.

Security Replication

Copy security descriptor: Each time the file or folder is copied

□ Process local trustees in security descriptor

○ Remove the SID from the security descriptor

● Translate the SID to the same trustee name and type on the target

☑ Create local groups on the target if not found

If unable to translate the SID:

○ Leave it as is

○ Remove the SID from the security descriptor

● Translate to Local Administrators group

□ Allow loss of security information

Figure 10-32 São Paulo server: security replication

6. Click **OK** to create the policy and check that the policy is listed in the Replication Policies panel as shown in Figure 10-33. In this example, we do not change any other value.

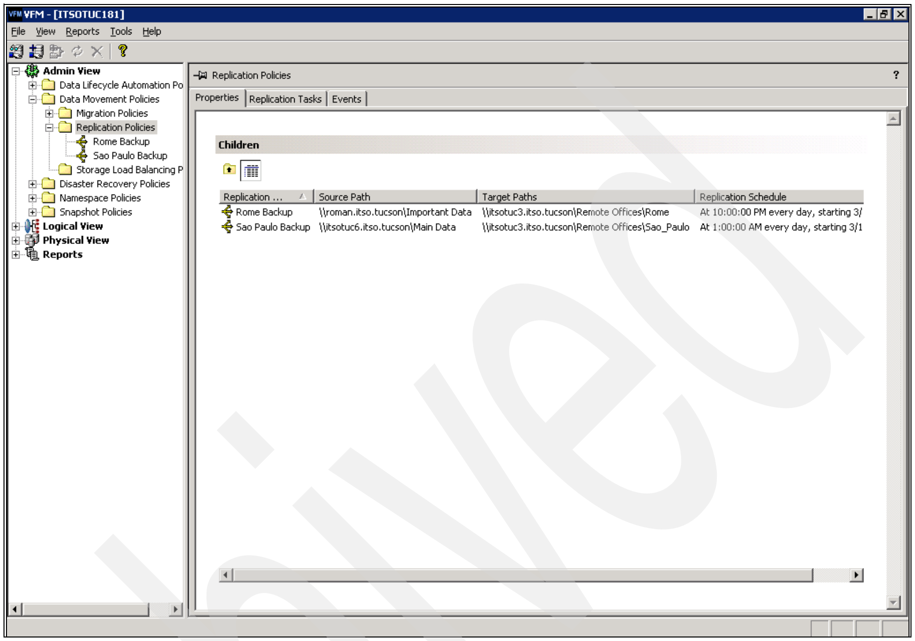

Replication ...	Source Path	Target Paths	Replication Schedule
Rome Backup	\\roman.itso.tucson\Important Data	\\itsotuc3.itso.tucson\Remote Offices\Rome	At 10:00:00 PM every day, starting 3/
Sao Paulo Backup	\\itsotuc6.itso.tucson\Main Data	\\itsotuc3.itso.tucson\Remote Offices\Sao_Paulo	At 1:00:00 AM every day, starting 3/1

Figure 10-33 Replication policies list with both Rome and São Paulo remote offices

7. Select the **Replication Topology** tab to view a graphical representation of the topology of this replication policy as shown in Figure 10-34.

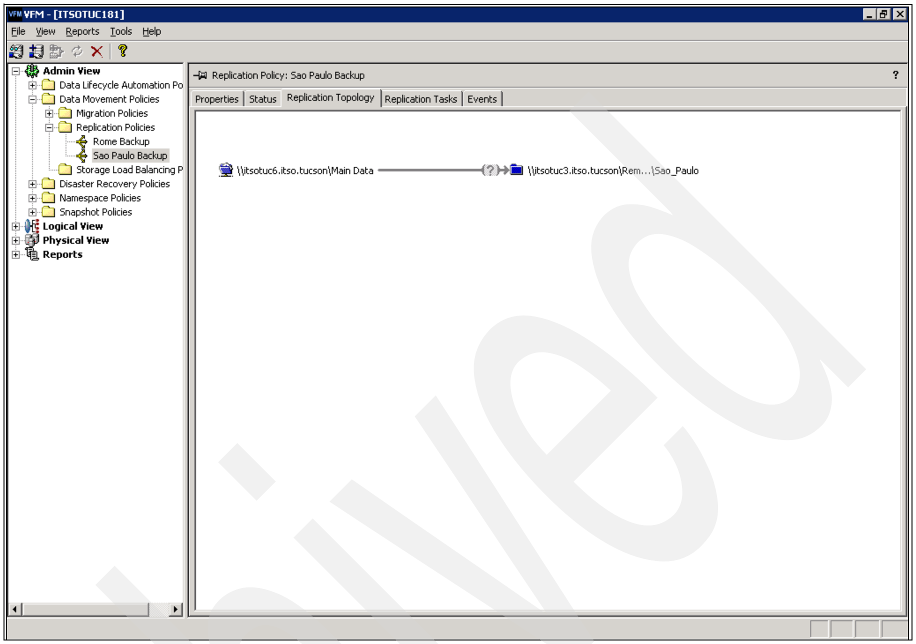

Figure 10-34   São Paulo server: replication topology

8. In the left panel, right-click the name of the policy. Then click **Start Replication now** (see Figure 10-35). This process is not required because you set a schedule for this policy, and it is always executed as you scheduled it. In this example, we use this option to demonstrate what happens when the replication is running and after it runs.

> **Note:** You can always start a replication manually if you want to.

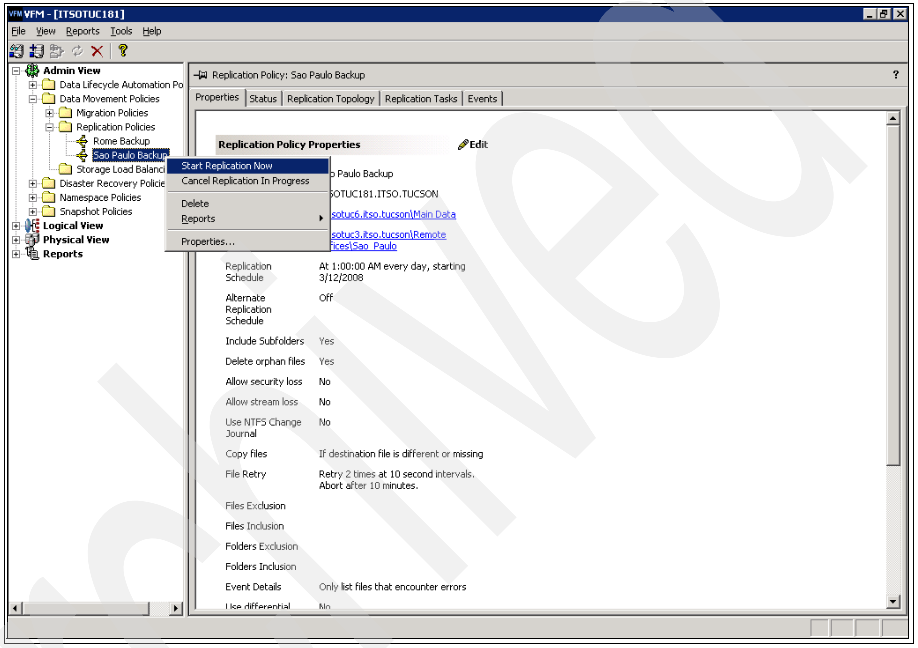

Replication Policy Properties

	o Paulo Backup
	SOTUC181.ITSO.TUCSON
	sotuc6.itso.tucson\Main Data
	sotuc3.itso.tucson\Remote fices\Sao_Paulo
Replication Schedule	At 1:00:00 AM every day, starting 3/12/2008
Alternate Replication Schedule	Off
Include Subfolders	Yes
Delete orphan files	Yes
Allow security loss	No
Allow stream loss	No
Use NTFS Change Journal	No
Copy files	If destination file is different or missing
File Retry	Retry 2 times at 10 second intervals. Abort after 10 minutes.
Files Exclusion	
Files Inclusion	
Folders Exclusion	
Folders Inclusion	
Event Details	Only list files that encounter errors
Use differential	No

Figure 10-35 São Paulo server: starting replication

9. Go back to the Replication Topology tab (see Figure 10-36). A blue circle indicates that the replication is being processed.

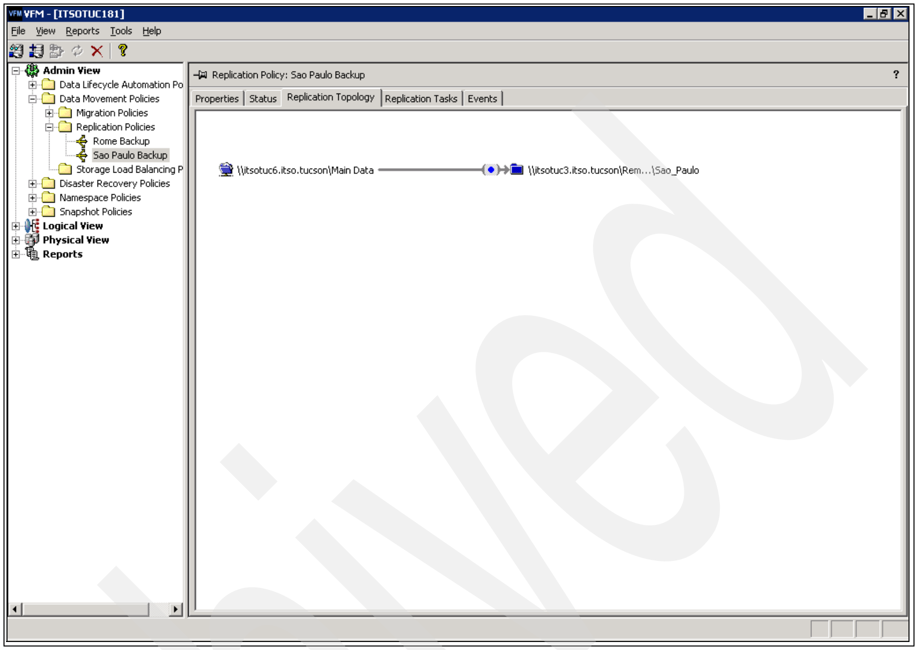

Figure 10-36 São Paulo topology during replication

10. On the Replication Tasks tab (see Figure 10-37), notice the contents of this window. One task is displayed for this replication policy.

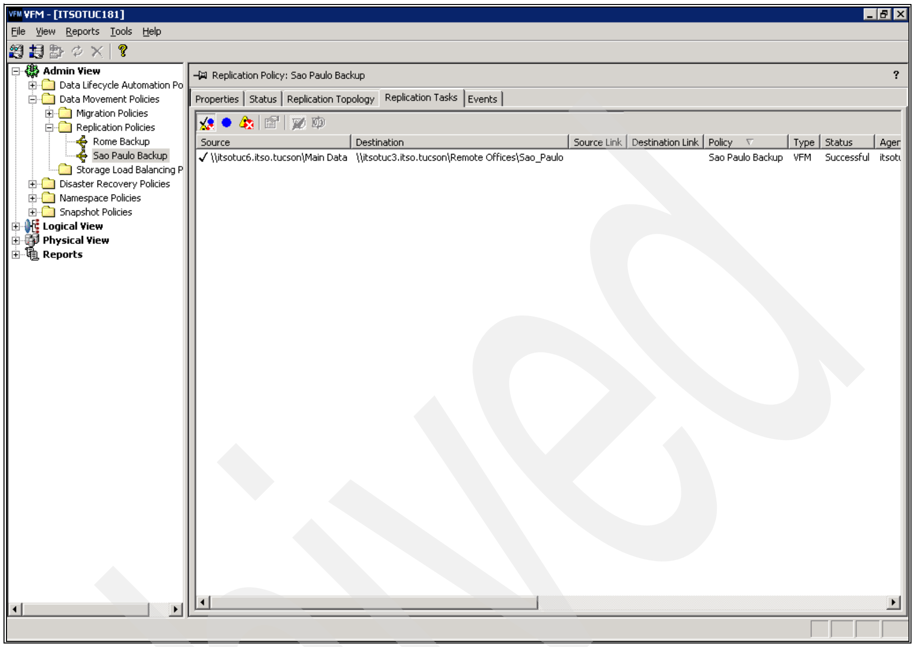

Source	Destination	Source Link	Destination Link	Policy	Type	Status	Ager
✓ \\itsotuc6.itso.tucson\Main Data	\\itsotuc3.itso.tucson\Remote Offices\Sao_Paulo			Sao Paulo Backup	VFM	Successful	itsotu

Figure 10-37 São Paulo server: replication tasks

11.When the replication finishes, you can view the final status as shown in Figure 10-38.

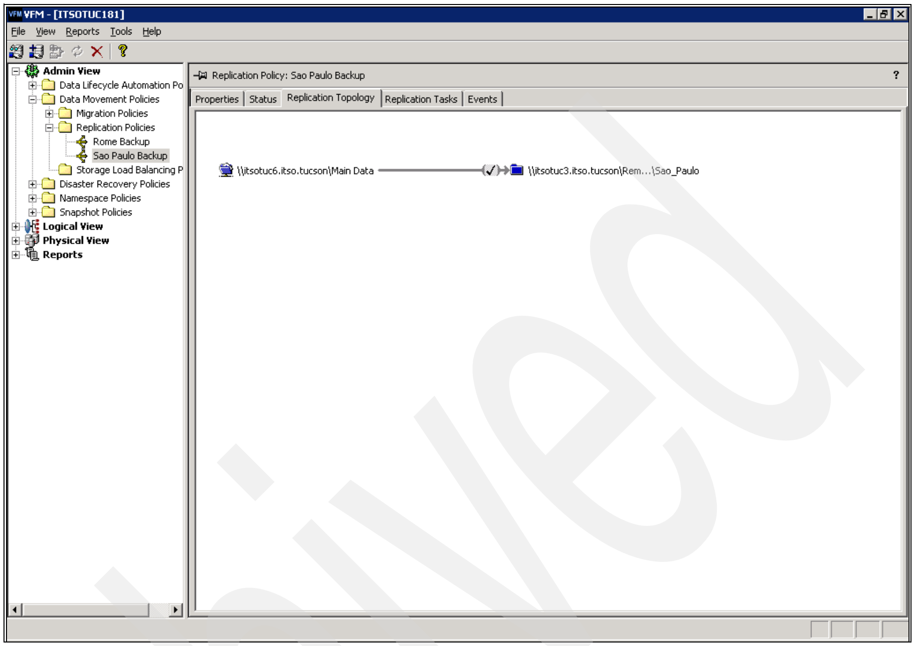

Figure 10-38 São Paulo replication topology after successful migration

12. Check the contents of the destination share (see Figure 10-39). You can verify that the replicated content matches the original content (see also Figure 10-14 on page 335).

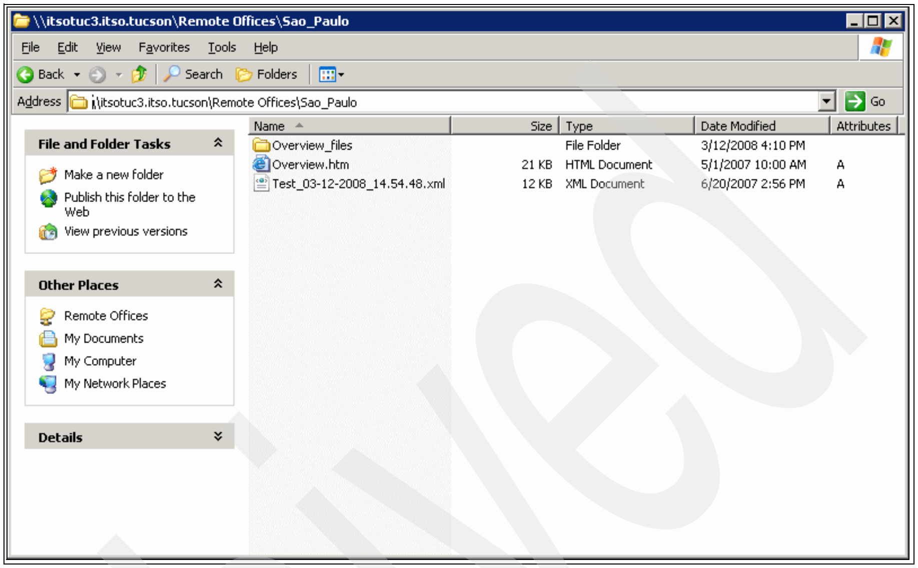
\\itsotuc3.itso.tucson\Remote Offices\Sao_Paulo

Address: \\itsotuc3.itso.tucson\Remote Offices\Sao_Paulo

Name	Size	Type	Date Modified	Attributes
Overview_files		File Folder	3/12/2008 4:10 PM	
Overview.htm	21 KB	HTML Document	5/1/2007 10:00 AM	A
Test_03-12-2008_14.54.48.xml	12 KB	XML Document	6/20/2007 2:56 PM	A

Figure 10-39 São Paulo backup: contents on destination

10.2.2 Data distribution

The IT department in Company XYZ is responsible for handling all aspects of software used by the company, that is, legal, new, upgrades, security updates, and so on. The department maintains a single repository listing the approved software and updates, and this repository is used by the entire company.

Because the company has several offices, each with several people that have to install the software (or update the already installed software), it make sense to maintain an updated copy of this repository, at least for the large offices.

In this example, we have the same three file servers (see 10.2.1, “Data gathering” on page 333):

- roman.itso.tucson is the Rome office server.
- itsotuc6.itso.tucson is the São Paulo office server.

- itsotuc3.itso.tucson is the N series located in Tucson. It contains the main software repository in the share Software Distribution (see Figure 10-40 and Figure 10-41 on page 362).

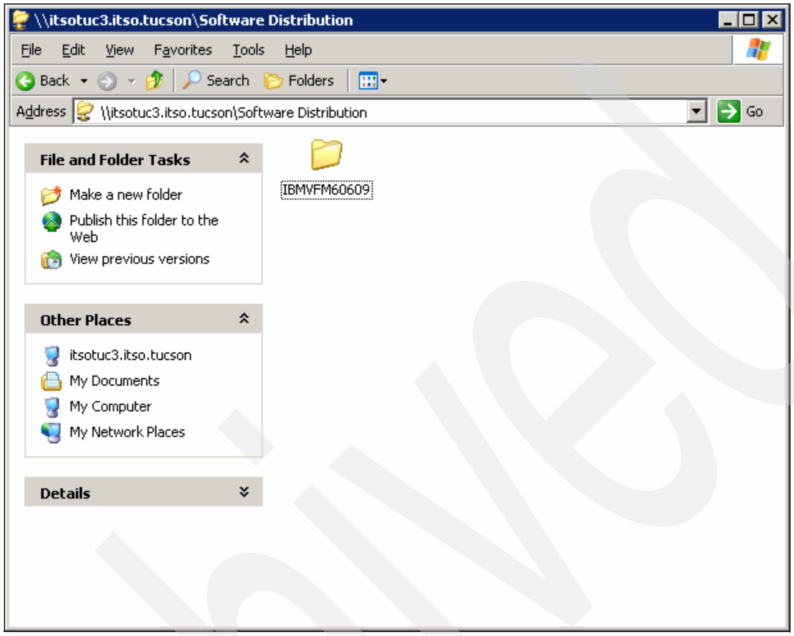

Figure 10-40 Original content on Software Distribution share

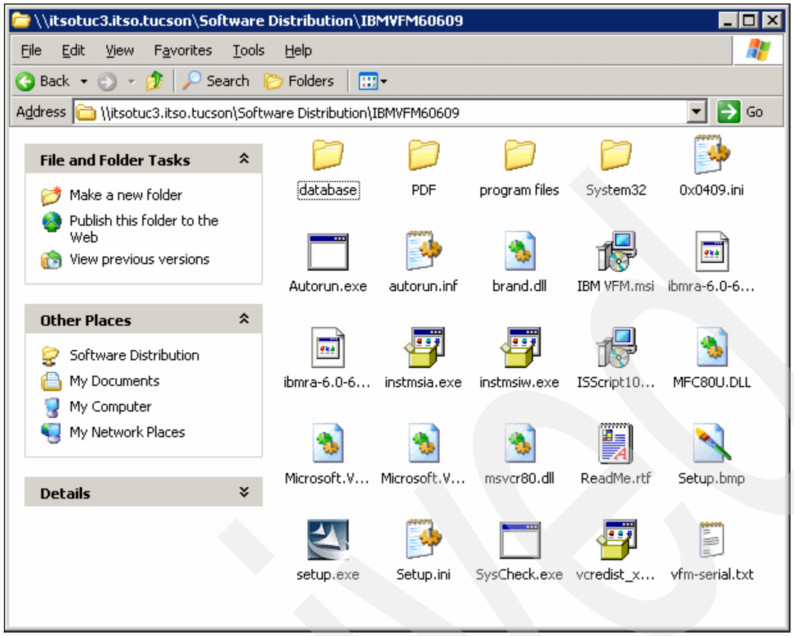

Figure 10-41 Original content on Software Distribution share: IBMVFM60609 folder

In this example, we demonstrate how to use VFM to create the shares on both the Rome and São Paulo servers and add them as a multitargeted DFS link, simplifying the replication and providing the benefits of a global namespace (refer to Chapter 5, “DFS namespace” on page 117).

In this scenario, the main repository is not supposed to be used directly by users; it is used only to replicate data to the office servers. Assume that the remote offices do not have the Software Distribution shares yet; thus we demonstrate how to use VFM to automatically create the share on all servers in a single operation. In addition, we recommend best practices for managing your resources on VFM.

Creating the Managed Resources folder

VFM enables you to create as many folders as you require in the Physical View panel. You create as many folders as required to organize all your servers, making it easier to manage them. In this example, we create a Managed Resources folder, and then add both the roman and itsotuc6 machines to it:

1. Go to the VFM client by choosing **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**.
2. Right-click **Physical View** and then click **Add Folder** as shown in Figure 10-42.

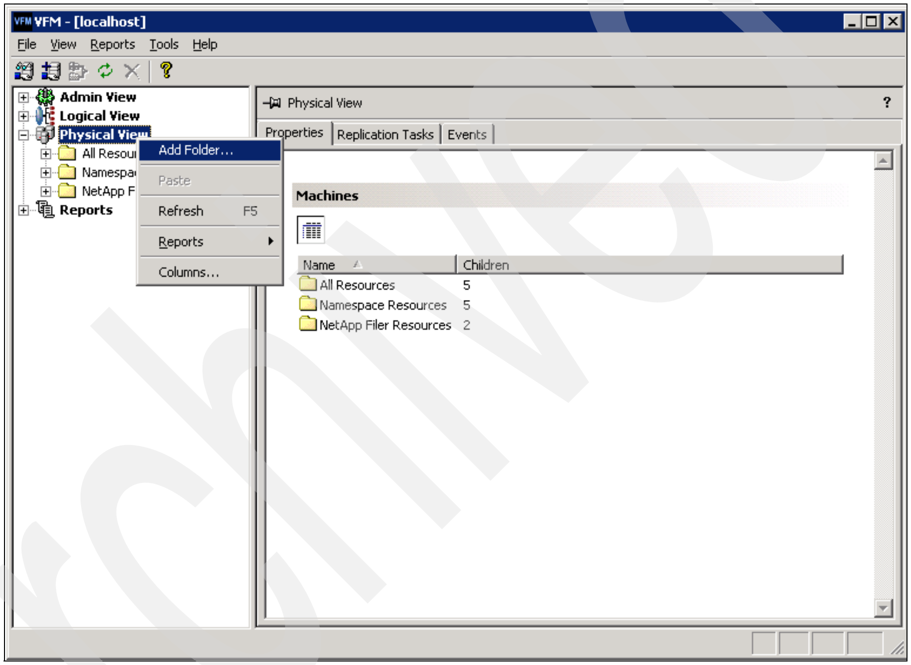

Name	Children
All Resources	5
Namespace Resources	5
NetApp Filer Resources	2

Figure 10-42 Physical view on VFM client

3. A new window opens, and you can set the folder name. Use Managed Resources as shown in Figure 10-43.

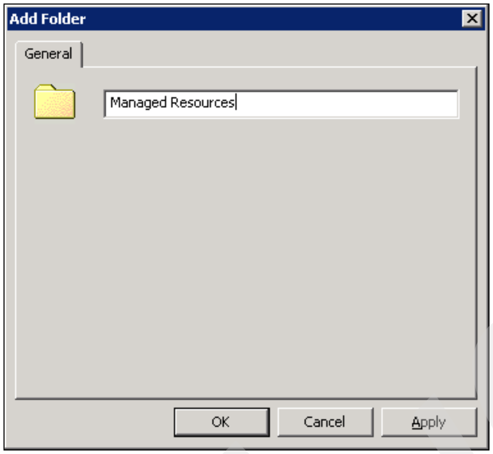

Figure 10-43 Adding Managed Resources folder

4. Right-click the newly created **Managed Resources** folder, then click **Add Machine** (see Figure 10-44).

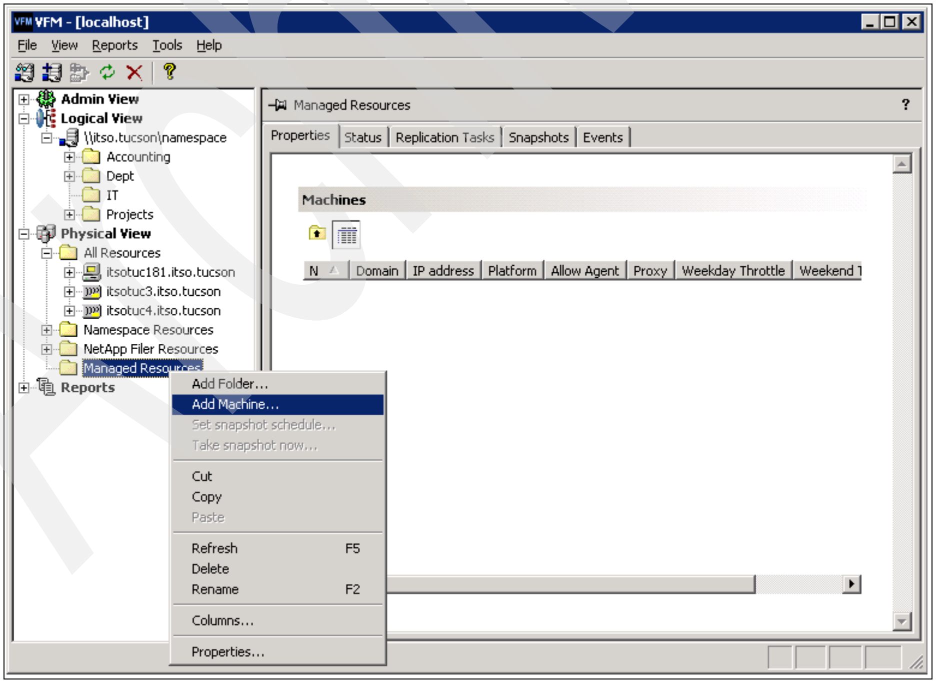

Figure 10-44 Add Machine on Managed Resources menu

5. In the new window, enter the first machine you want to add. In this example, it is roman.itso.tucson (see Figure 10-45). Then click **OK**.

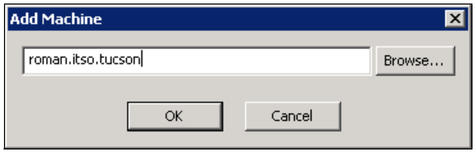

roman.itso.tucson

Figure 10-45 Adding first machine

6. Repeat step 5, but this time add the second machine itsotuc6.itso.tucson (see Figure 10-46).

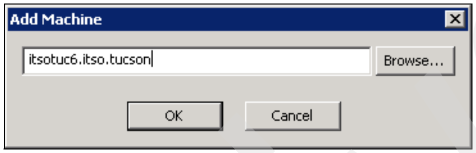

itsotuc6.itso.tucson

Figure 10-46 Adding second machine

You can see both machines added to the Managed Resources folder as shown in Figure 10-47.

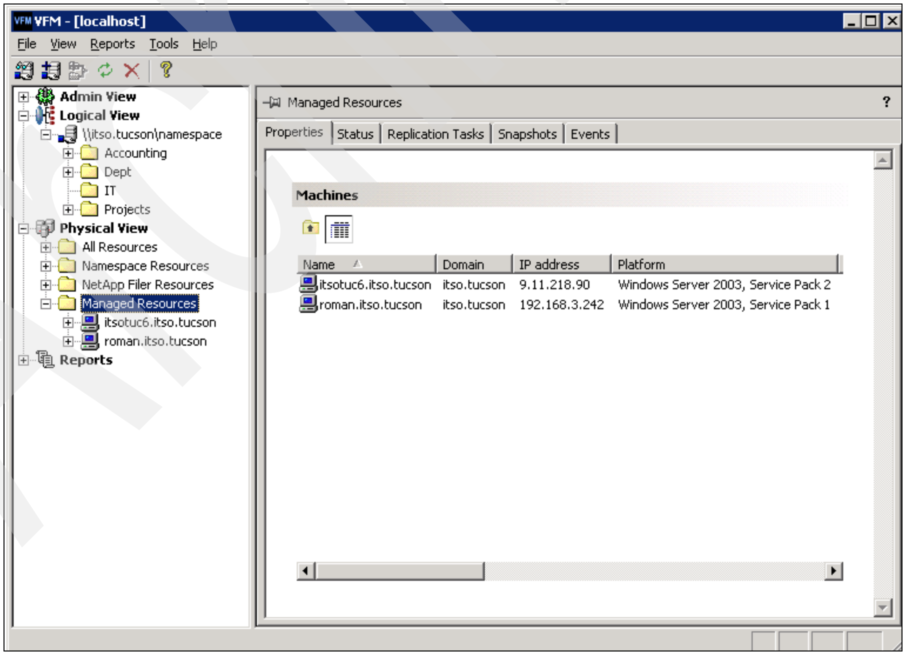

Name	Domain	IP address	Platform
itsotuc6.itso.tucson	itso.tucson	9.11.218.90	Windows Server 2003, Service Pack 2
roman.itso.tucson	itso.tucson	192.168.3.242	Windows Server 2003, Service Pack 1

Figure 10-47 Managed Resources folder with both roman and itsotuc6 machines

Creating and adding shares to the namespace

At this point, we create the shares on both machines, using VFM and also using the Managed Resources folder we just created:

1. On the VFM client, go to the Managed Resources folder, right-click one of the machines, and then click **Create Shares**. On Figure 10-48 we show this step on the itsotuc6 machine.

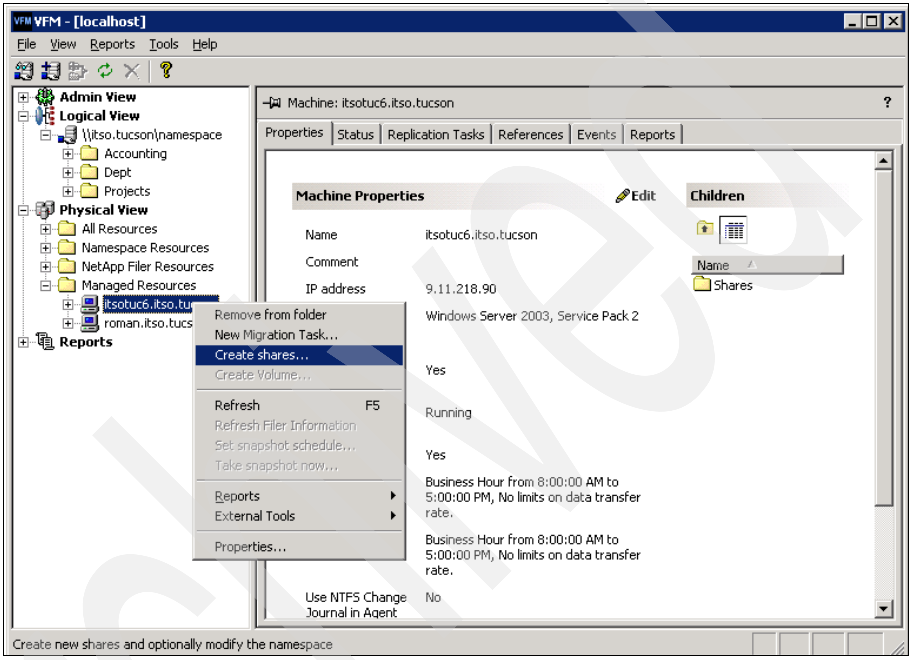

Name itsotuc6.itso.tucson

Comment

IP address 9.11.218.90

Windows Server 2003, Service Pack 2

Yes

Running

Yes

Business Hour from 8:00:00 AM to 5:00:00 PM, No limits on data transfer rate.

Business Hour from 8:00:00 AM to 5:00:00 PM, No limits on data transfer rate.

Use NTFS Change Journal in Agent No

Create new shares and optionally modify the namespace

Figure 10-48 Right-click itsotuc6 machine

2. The Share Creation Wizard is displayed (see Figure 10-49). Click **Next** to proceed.

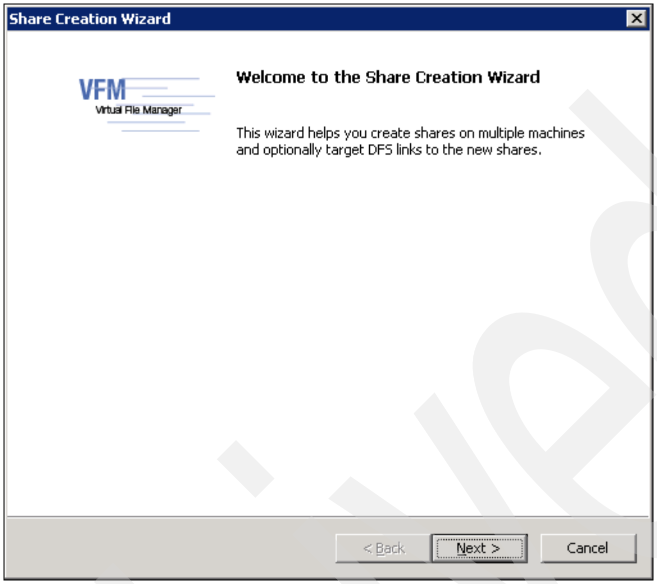

Welcome to the Share Creation Wizard

This wizard helps you create shares on multiple machines and optionally target DFS links to the new shares.

Figure 10-49 Share Creation Wizard

3. Choose **Regular Share** and click **Next** (see Figure 10-50).

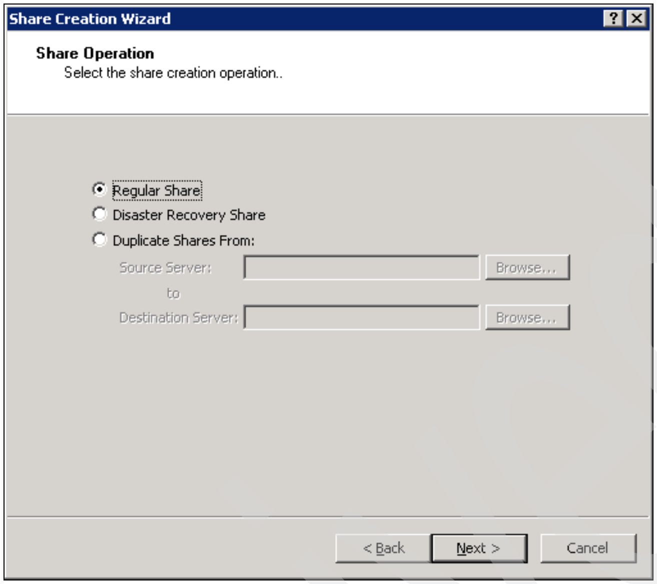

Share Creation Wizard

Share Operation

Select the share creation operation..

● Regular Share
○ Disaster Recovery Share
○ Duplicate Shares From:

Source Server: [ ] Browse...
to
Destination Server: [ ] Browse...

Figure 10-50 New share type

4. Enter the share information as shown in Figure 10-51.

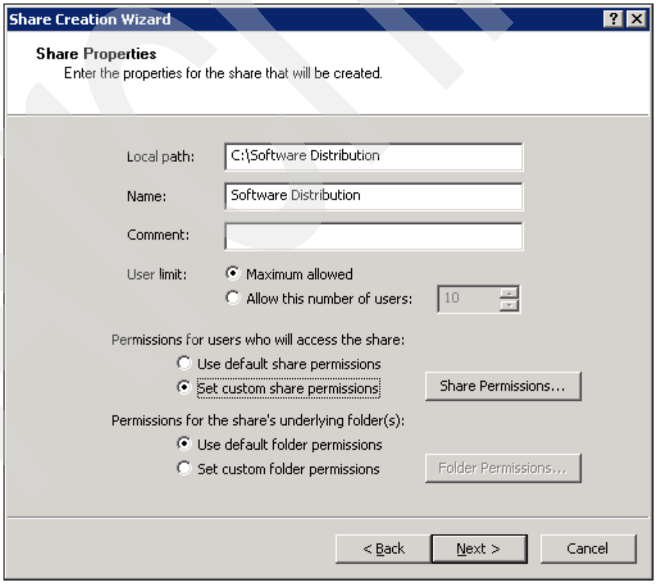

Share Creation Wizard

Share Properties

Enter the properties for the share that will be created.

Local path: C:\Software Distribution

Name: Software Distribution

Comment:

User limit:
● Maximum allowed
○ Allow this number of users: 10

Permissions for users who will access the share:
○ Use default share permissions
● Set custom share permissions

Permissions for the share's underlying folder(s):
● Use default folder permissions
○ Set custom folder permissions

Figure 10-51 New share properties

5. Select **Set custom share permissions**.
6. Then click **Share permissions**. You view the default share permissions as shown in Figure 10-52.

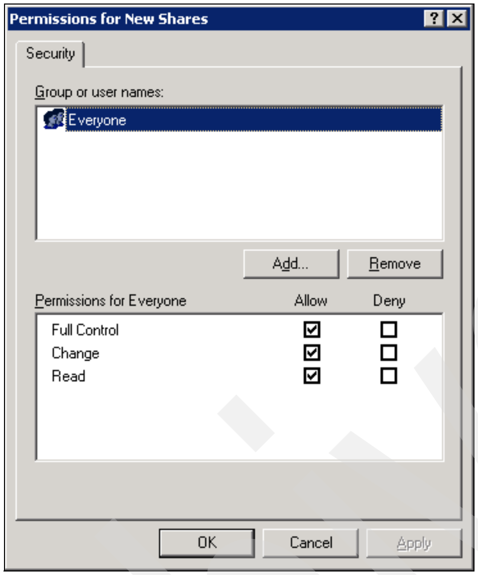

Group or user names:
Everyone

Permissions for Everyone	Allow	Deny
Full Control	☑	□
Change	☑	□
Read	☑	□

Figure 10-52 Default share permissions on Microsoft Windows 2003

7. Click **Add**, enter your VFM service account name, and then click **OK** (see Figure 10-53).

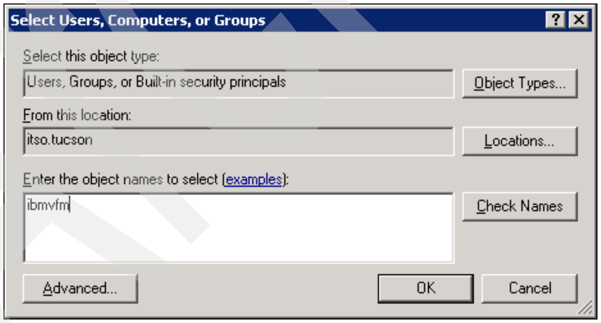

Select this object type:
Users, Groups, or Built-in security principals

From this location:
itso.tucson

Enter the object names to select (examples):
ibmvfm

Figure 10-53 Adding user to share permission

8. Check that the user was correctly added to the list.

9. Check the **Full Control** check box, then click **OK** (see Figure 10-54).

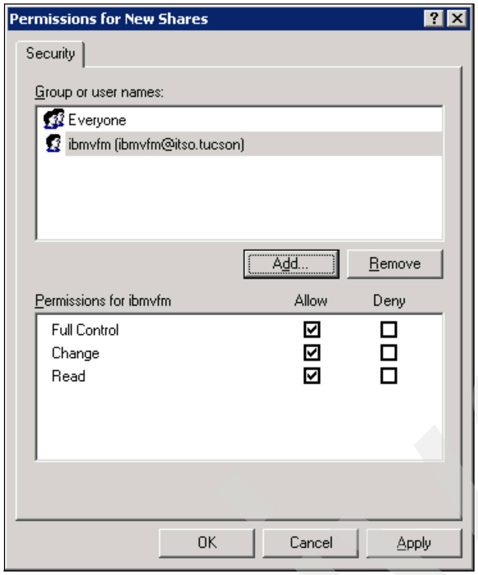

Group or user names:

Everyone

ibmvfm (ibmvfm@itso.tucson)

Permissions for ibmvfm	Allow	Deny
Full Control	☑	☐
Change	☑	☐
Read	☑	☐

Figure 10-54 Adding Full Control to the VFM service account

10. At this point, you can select exactly on which machines you plan to create the share. In our example, we use both the machines we added to the Managed Resources folder (see Figure 10-55), but you are not required to do so. You can also browse the machines on your domain and select them.

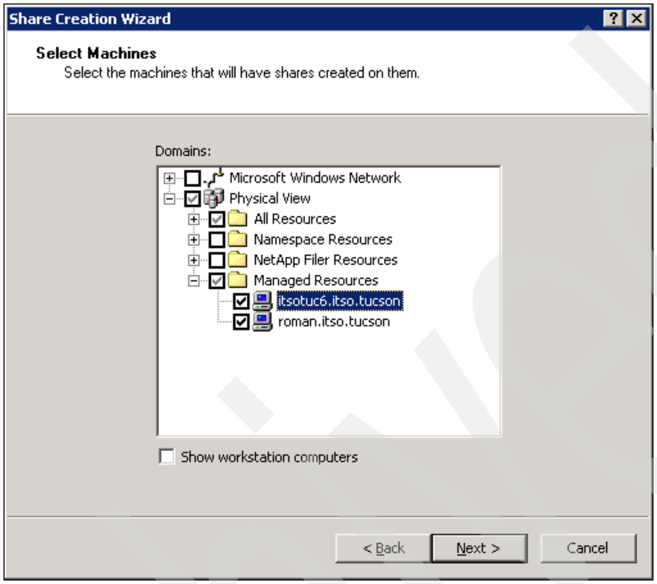

Share Creation Wizard

Select Machines

Select the machines that will have shares created on them.

Domains:

- ☐ Microsoft Windows Network
- ☑ Physical View
  - ☑ All Resources
  - ☐ Namespace Resources
  - ☐ NetApp Filer Resources
  - ☑ Managed Resources
    - ☑ itsotuc6.itso.tucson
    - ☑ roman.itso.tucson

☐ Show workstation computers

Figure 10-55 Selecting machines on which to create the shares

11. Because we want the benefits of the global namespace, we click the **Add the new shares, as targets, to a new DFS Link** radio button (see Figure 10-56).

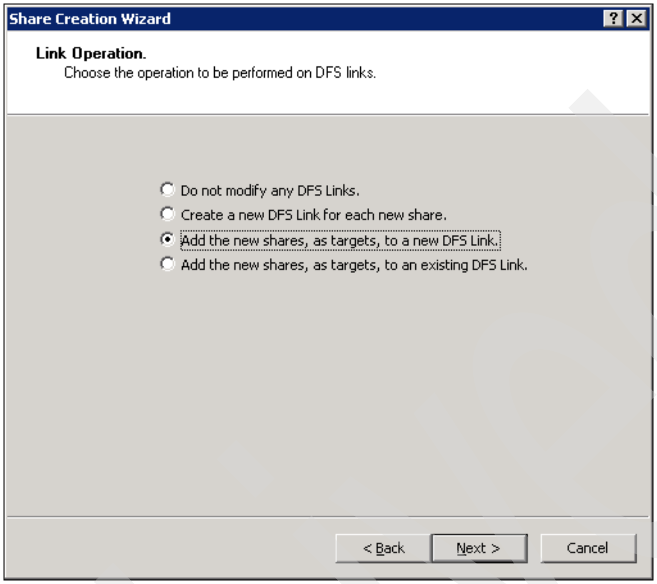

Share Creation Wizard

Link Operation.
Choose the operation to be performed on DFS links.

○ Do not modify any DFS Links.
○ Create a new DFS Link for each new share.
● Add the new shares, as targets, to a new DFS Link.
○ Add the new shares, as targets, to an existing DFS Link.

Figure 10-56 DFS link operation on new shares

12. Click **Next**.

13. Choose the DFS root you want to use and click **Next**. In our case, we choose the one shown in Figure 10-57.

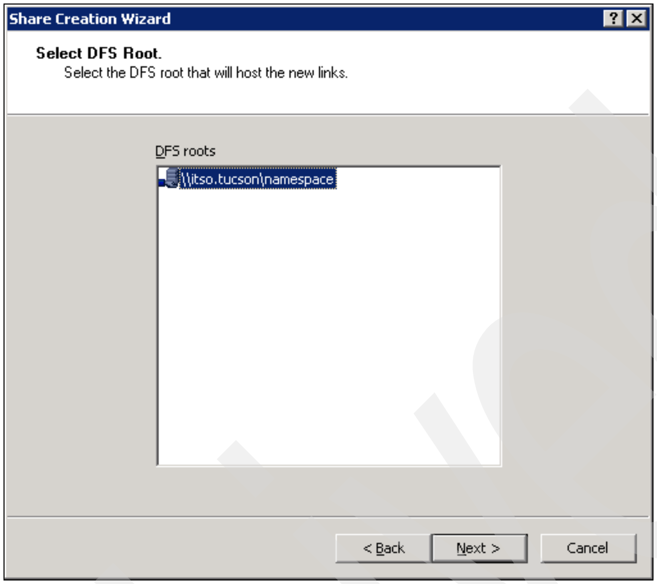

Share Creation Wizard

Select DFS Root.
Select the DFS root that will host the new links.

DFS roots

\\itso.tucson\namespace

Figure 10-57 Choosing DFS root for new link

14. The new link is created inside a folder on the namespace. VFM opens the dialog box displaying the New Folder folder, as shown in Figure 10-58. Change the folder name to IT as shown in Figure 10-59, then click **Next**.

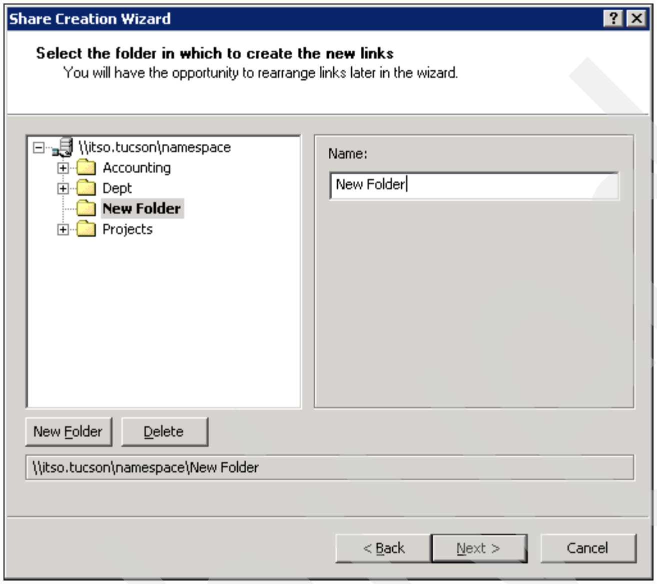

Share Creation Wizard

Select the folder in which to create the new links

You will have the opportunity to rearrange links later in the wizard.

\\itso.tucson\namespace
- Accounting
- Dept
- **New Folder**
- Projects

Name: New Folder

\\itso.tucson\namespace\New Folder

Figure 10-58   Default New Folder on namespace

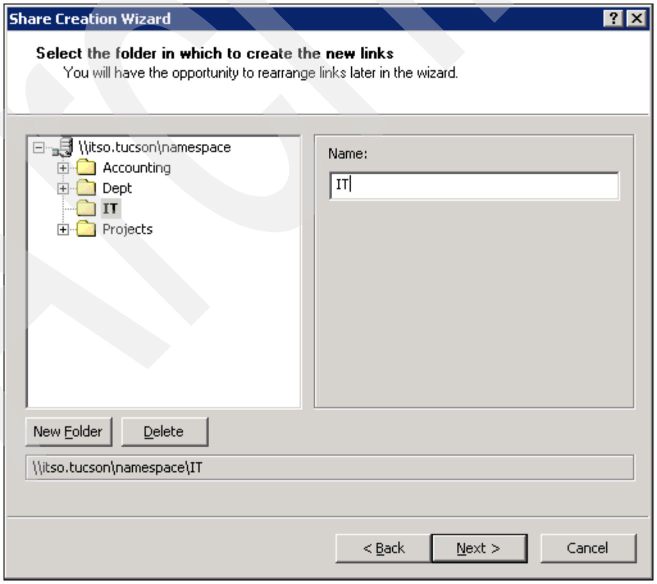

Share Creation Wizard

Select the folder in which to create the new links

You will have the opportunity to rearrange links later in the wizard.

\\itso.tucson\namespace
- Accounting
- Dept
- **IT**
- Projects

Name: IT

\\itso.tucson\namespace\IT

Figure 10-59   Creating IT folder on namespace

15. The folder is created, and the Share Creation Wizard displays the link that is about to be created (see Figure 10-60). Click **Next**.

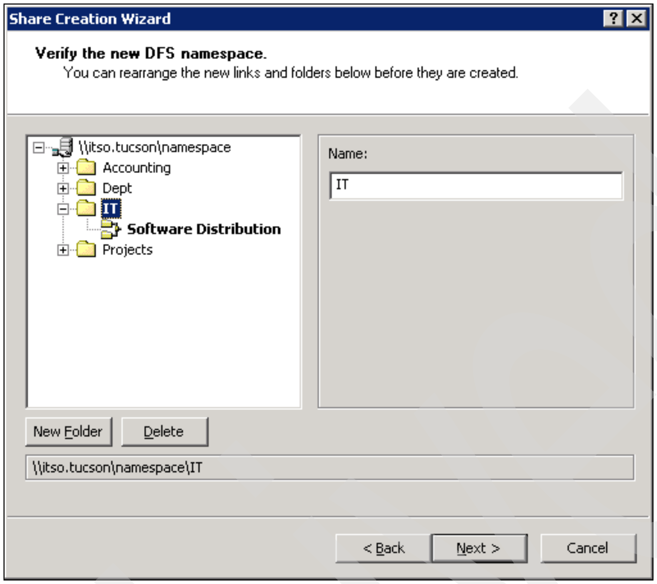

Share Creation Wizard

Verify the new DFS namespace.

You can rearrange the new links and folders below before they are created.

- \\itso.tucson\namespace
  - Accounting
  - Dept
  - IT
    - **Software Distribution**
  - Projects

Name: IT

\\itso.tucson\namespace\IT

Figure 10-60 Namespace with IT folder and new link

16. Click **Finish**. The shares are created, and the namespace is updated (see Figure 10-61).

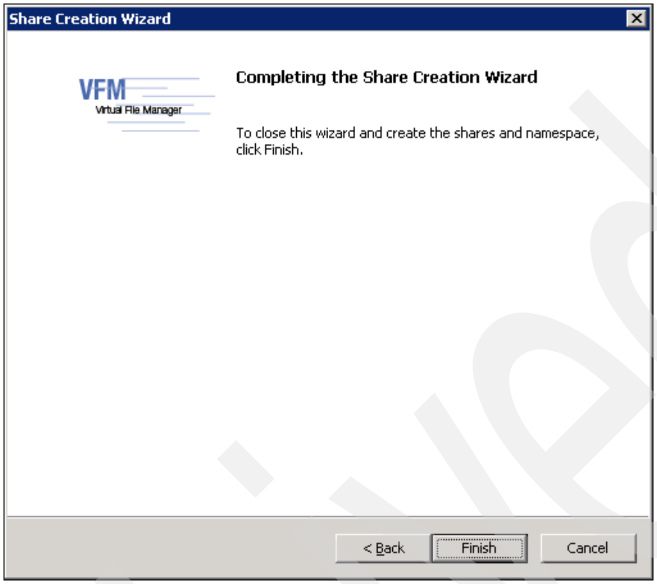

Completing the Share Creation Wizard

To close this wizard and create the shares and namespace, click Finish.

Figure 10-61 Completing Share Creation Wizard

17. Back in the main window, click **Logical View**, open the IT folder, and click the **Software Distribution** link. As you can see in Figure 10-62, it has two targets, the recently created shares on both roman and itsotuc6 machines.

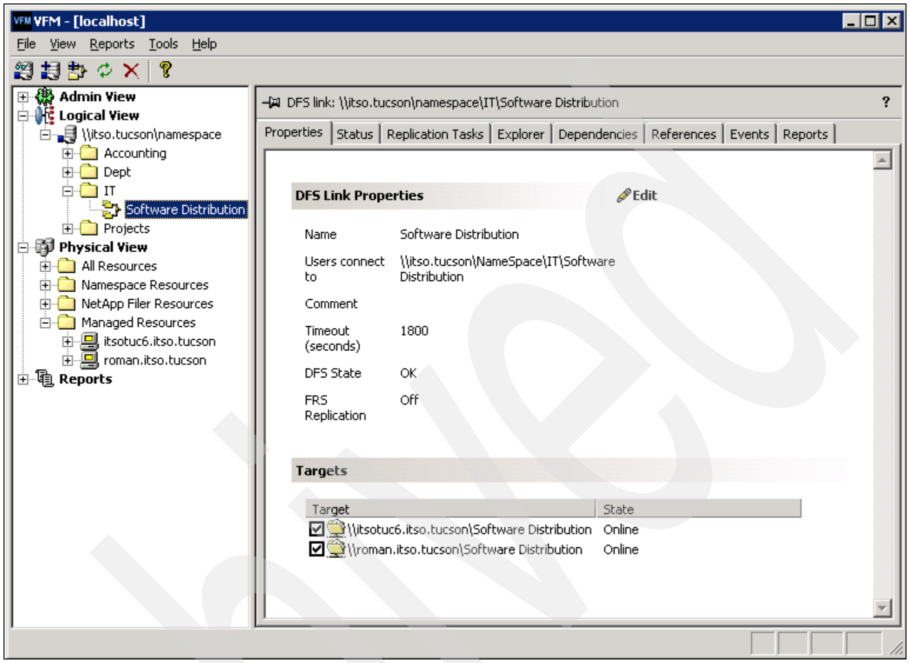

Name	Software Distribution
Users connect to	\\itso.tucson\NameSpace\IT\Software Distribution
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Target	State
☑ \\itsotuc6.itso.tucson\Software Distribution	Online
☑ \\roman.itso.tucson\Software Distribution	Online

Figure 10-62 Software Distribution link properties showing the two targets

Replicating data to the new shares

Now that the shares and the DFS link are created, we can effectively create the replication policy between them.

In this example, we use the DFS link as our destination. VFM always replicates data to all targets on that the DFS link. Follow these steps:

1. On the VFM client, go to **Admin View** → **Data Movement Policies** → **Replication Policies** (see Figure 10-63).

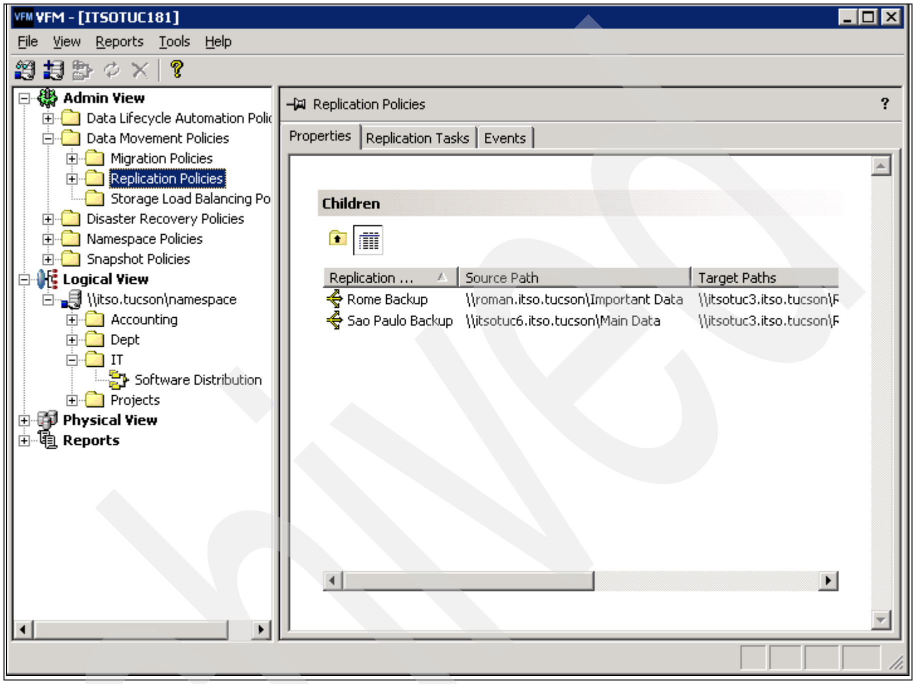

Replication ...	Source Path	Target Paths
Rome Backup	\\roman.itso.tucson\Important Data	\\itsotuc3.itso.tucson\F
Sao Paulo Backup	\\itsotuc6.itso.tucson\Main Data	\\itsotuc3.itso.tucson\F

Figure 10-63 Replication Policy list

2. Right-click **Replication Policy**, then click **New Replication Policy** (see Figure 10-64).

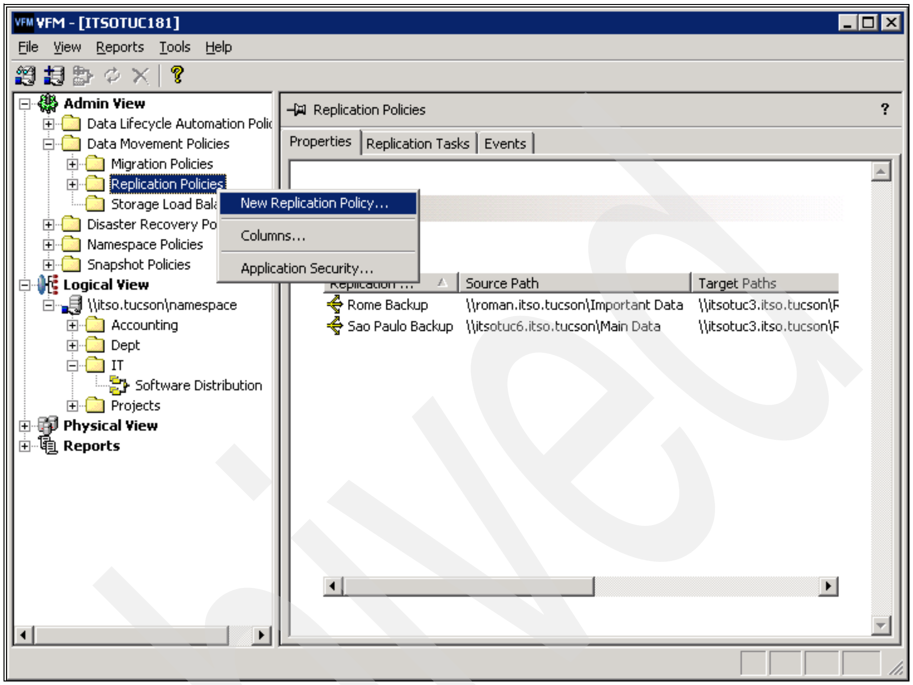

Replication ...	Source Path	Target Paths
Rome Backup	\\roman.itso.tucson\Important Data	\\itsotuc3.itso.tucson\F
Sao Paulo Backup	\\itsotuc6.itso.tucson\Main Data	\\itsotuc3.itso.tucson\F

Figure 10-64 New replication policy menu

3. Enter a name for this policy and the source and target information (see Figure 10-65). Note that our target is the DFS link. You can use both the NetBIOS or the fully qualified domain name on the DFS link.

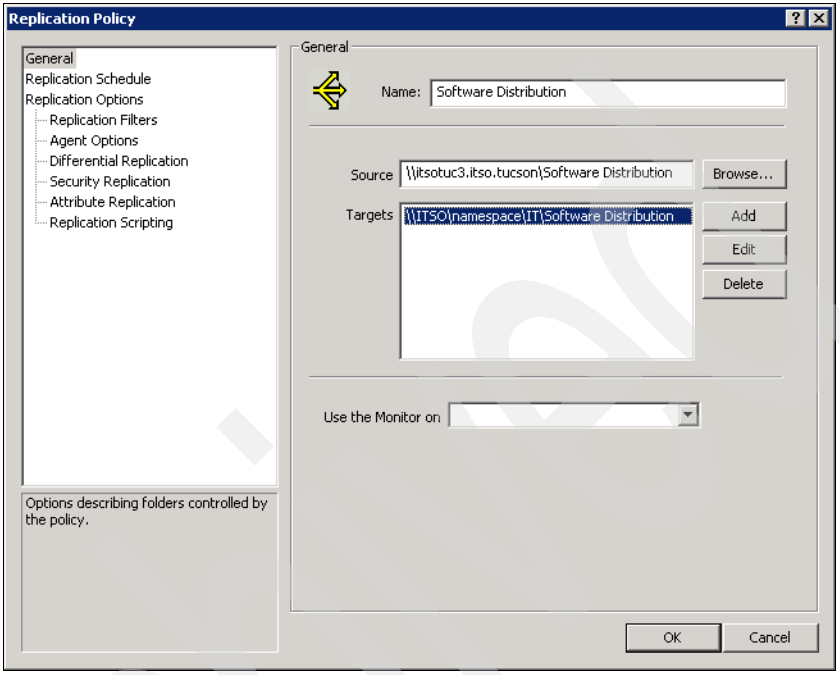

Name: Software Distribution

Source \\itsotuc3.itso.tucson\Software Distribution

Targets \\ITSO\namespace\IT\Software Distribution

Options describing folders controlled by the policy.

Figure 10-65 Software Distribution replication policy

4. Set the schedule for this replication. In this step, we replicate this data daily, starting at 10 p.m. as shown in Figure 10-66.

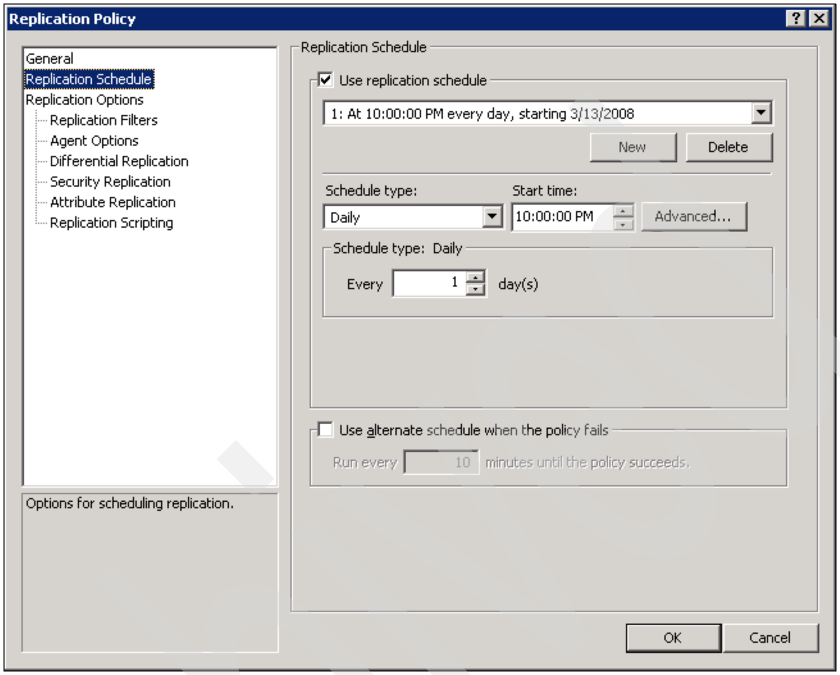

Replication Policy

Options for scheduling replication.

Replication Schedule

☑ Use replication schedule

1: At 10:00:00 PM every day, starting 3/13/2008

Schedule type: Daily

Start time: 10:00:00 PM

Schedule type: Daily

Every 1 day(s)

□ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

Figure 10-66 Software Distribution replication schedule

5. In the Replication Options panel, click the **Optimize for Security** button (see Figure 10-67). The difference between the options after you click this button is that no Copy-in-place operation is executed. You can read more about this option in 10.1, “Using replication options” on page 316. Because we want all replicas with exactly the same content, we click the **Delete orphaned files in destination folders** check box.

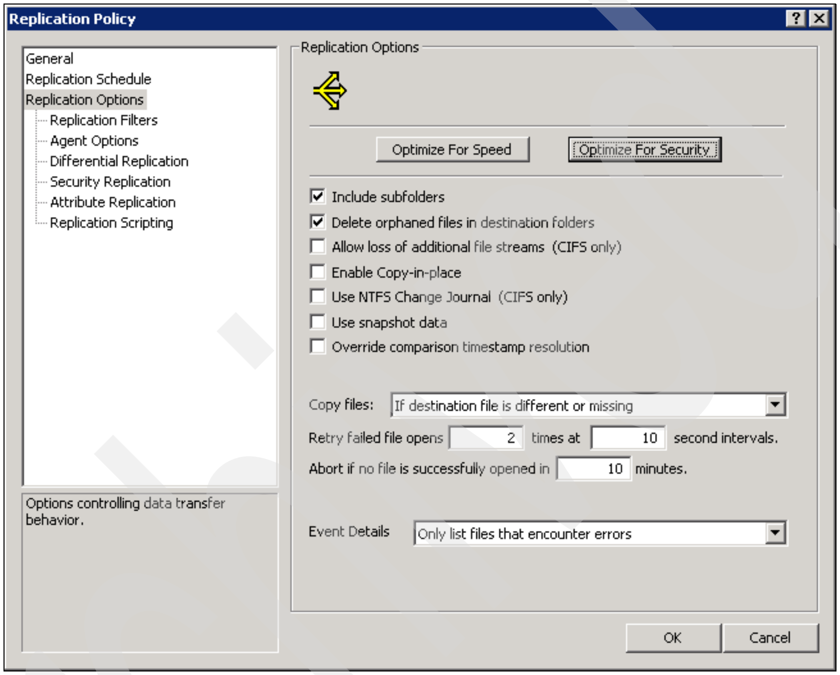

Options controlling data transfer behavior.

☑ Include subfolders
☑ Delete orphaned files in destination folders
☐ Allow loss of additional file streams (CIFS only)
☐ Enable Copy-in-place
☐ Use NTFS Change Journal (CIFS only)
☐ Use snapshot data
☐ Override comparison timestamp resolution

Copy files: If destination file is different or missing

Retry failed file opens 2 times at 10 second intervals.

Abort if no file is successfully opened in 10 minutes.

Event Details Only list files that encounter errors

Figure 10-67 Software Distribution replication options without copy-in-place

6. Click **Security Replication**, uncheck the **Process local trustees in security descriptor** check box (see Figure 10-68).

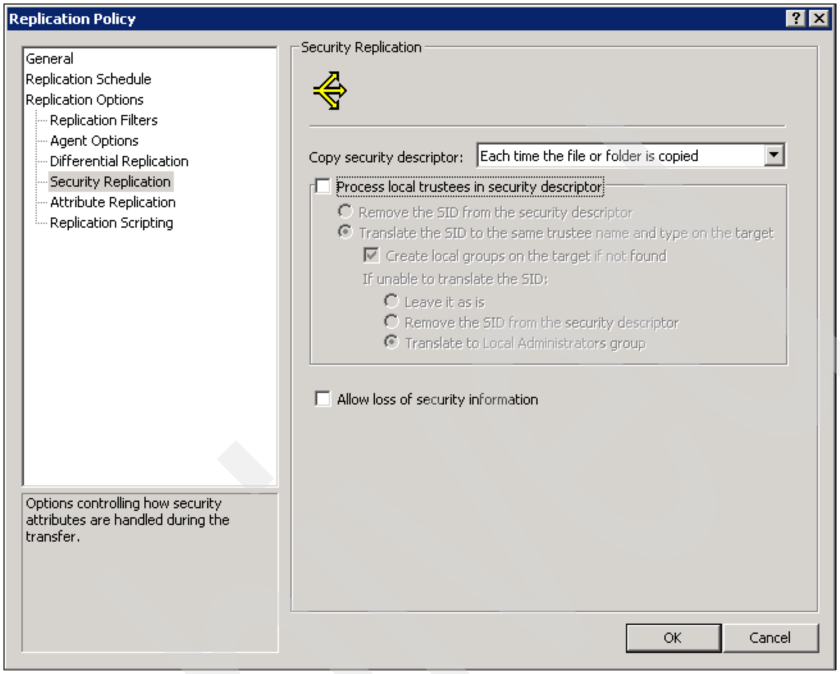

Options controlling how security attributes are handled during the transfer.

Security Replication

Copy security descriptor: Each time the file or folder is copied

□ Process local trustees in security descriptor

○ Remove the SID from the security descriptor

● Translate the SID to the same trustee name and type on the target

☑ Create local groups on the target if not found

If unable to translate the SID:

○ Leave it as is

○ Remove the SID from the security descriptor

● Translate to Local Administrators group

□ Allow loss of security information

Figure 10-68 Software Distribution security options

7. Click **OK** to create the policy.

   You can see the newly created policy on Figure 10-69.

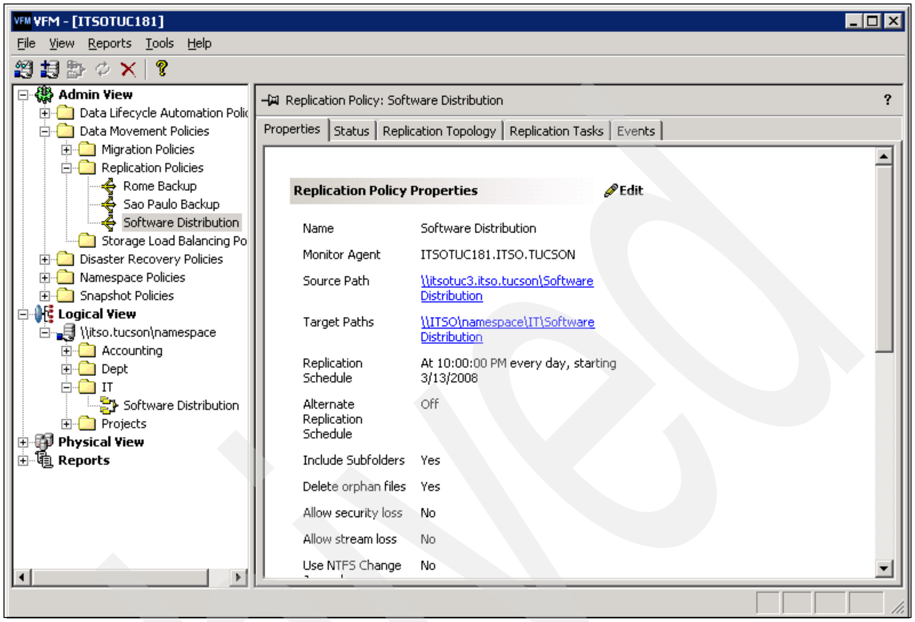

Replication Policy Properties

Name	Software Distribution
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Source Path	\\itsotuc3.itso.tucson\Software Distribution
Target Paths	\\ITSO\namespace\IT\Software Distribution
Replication Schedule	At 10:00:00 PM every day, starting 3/13/2008
Alternate Replication Schedule	Off
Include Subfolders	Yes
Delete orphan files	Yes
Allow security loss	No
Allow stream loss	No
Use NTFS Change	No

Figure 10-69 Software Distribution policy properties

On the Replication Topology tab (see Figure 10-70), you see a slightly more complex topology, with two targets. You can also check that VFM has all the targets on the link as destinations.

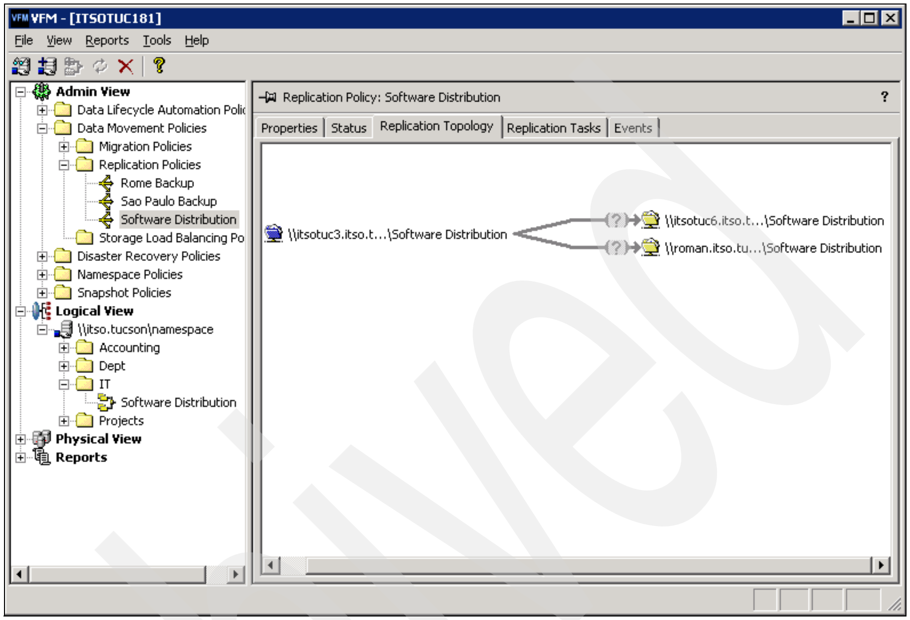

Figure 10-70 Software Distribution topology

11

Disaster recovery

This chapter describes using VFM to manage disaster recovery on CIFS shares.

11.1 Disaster recovery overview

High availability of data is critical for an enterprise. Downtime of a few minutes or a few hours can be significantly costly. Different levels of recovery can be provided for different data within an enterprise. Generally, structured data (applications, e-mail, databases, and so on) have a much higher requirement for uptime than unstructured file data.

VFM provides the capability to failover users for replicated file data between heterogeneous data sources. It is a general-purpose replication and disaster recovery (DR) solution focused on unstructured file data. The replication functionality provides a means to keep data in sync between two entities. In the event of a failover of the primary server, VFM can fail users to a DR site. VFM can use its own internal replication mechanism or rely on an external replication means (for example, SnapMirror). Most people look at VFM to extend DR capabilities to provide business continuance (BC) for file data. Thus, VFM is a DR and BC solution for file data.

Disaster recovery is focused on having a replicated copy of the data on another machine or in a different location. For some data sets, you might want synchronous replication of data. In other situations, semi-synchronous replication of data is called for. In such a case, replication can occur synchronously most of the time, but it is acceptable to have a brief replication lag. In most environments, asynchronous replication of file data is adequate. The frequency of asynchronous replication is governed by the business value of the data, the recovery time object (RTO), and the recovery point objective (RPO). Some clients have a mix where some data must be replicated synchronously, while other data can be replicated asynchronously. Even though you have a large quantity of file data, not all the data must be replicated and highly available.

Irrespective of how current the data is at the DR site, another important aspect is maintaining user access to the data after a failover. In case of a disaster, having the data available at the DR site is only one part of the equation. Connecting the user to the data is equally as important. Connecting users to the second copy of the data requires sending e-mails to the affected users, updating logon scripts, changing network parameters such as DNS and WINS entries, renaming servers at the DR site to match the server names on the source site, and more. The amount of work required to failover users during a disaster must be weighed against the likelihood of failure of the data located on the primary, the cost of the unavailability of data for a period of time, and the complexity of the DR solution.

VFM leverages the simplicity of the namespace to provide access to the data for users. Users never have to know the physical server name and share name of the data. If they access the data through the namespace, they are redirected to the appropriate data location. Because the users are shielded from the physical infrastructure, the failover scenario is simplified and can be transparent. In the event of a failure, the namespace is updated to point the users, automatically or manually, to the replicated copy of the data at the DR location. Users receive the new referral to the DR site and are automatically redirected. You might not have to send users e-mail, update logon scripts, or change network settings such as WINS or DNS entries. The time required for the failover is commensurate with the RPO and RTO of the data.

Failover scenarios are described in the following examples:

- Failure of a single volume on the source server
- Failure of a source server holding the data
- Failure of a entire site

The amount of data to be replicated to the DR site depends on the business value of the data. Not all the data has to be made highly available. Suppose a company has 10 TB of file data. Of the 10 TBs, only 1 TB is company-critical data that must be made highly available. The problem is that the 1 TB might be scattered across four servers. Instead of replicating an entire volume, you can replicate only a few shares.

In some scenarios, you replicate data within the data center, which is local replication of data. In other scenarios, you can perform remote replication of data. For example, your primary data center is New York, and your backup data center is in Los Angeles.

In other scenarios, you can perform a Local → Local → Remote replication to handle not just server failure but also site failure. In this example, if you lose a volume or a local server, you fail over to the other local server. If you lose a site, you fail over to the remote server. You can also combine the synchronous replication capabilities with asynchronous replication capabilities in the Local → Local → Remote failover scenario. You can replicate data synchronously or semi-synchronously among the servers at the local site but replicate data asynchronously among the servers at the remote site.

Another consideration is the extent of the disaster. For example, if a server hosting the data simply reboots, you might not want to fail over the users to a remote data center. You can tolerate the few minutes that it takes for the server to come back online and become available. However, if you have lost a volume that requires replacement parts that take hours or days to fix, you are more likely to choose to fail over the users to the replicated copy of the data.

This chapter discusses the following use cases:

- Making critical data highly available at a DR site in preparation for failover
- Integrating with SnapMirror to provide DR for N series volumes
- Replicating data from a remote office to a central data center for failover

11.2 Introduction to VFM disaster recovery policies

VFM provides three separate DR policies to provide a high availability solution for CIFS-based file data. The DR policies have different levels of flexibility and automation to simplify the configuration and operation of the policy.

- **Client recovery policy**

  This policy is the most generic, flexible, and powerful of the three DR policies. It allows granular failover of a single share from one machine to another. You can configure this policy to do a local → local failover, local → remote failover, or a local → local → remote failover. This policy does require some amount of setup to provide its flexibility.

- **Server recovery policy**

  This policy is used to replicate data from a server to a N series. It can replicate some or all shares of a server. If a new share is created on the remote server, it can be automatically added to the policy and made highly available, which is useful in a set-and-forget mode.

  The best use of this policy is when you have a remote office, and you want to have a DR copy of the data in a central location. You can replicate the data once a night. In the event of a failure of the remote office, users are redirected to the central site and have access to the data, although with some latency. When the server in the remote office is restored, the data can be replicated in the reverse direction, and user access to the remote server can be restored.

- **Filer recovery policy**

  This flexible policy is used to keep two N series storage systems in sync. It leverages a SnapMirror relation between the N series volumes to maintain replication of the data. In the event of a failure of the primary volume, the filer DR policy breaks the SnapMirror relationship, makes the target volume on the destination N series writable, and updates the namespace to point users to the new destination volume. The Filer DR policy requires some setup but is a powerful way to integrate with SnapMirror to make data highly available. It relies on SnapMirror to replicate the data and VFM to control the failover. Hence, it completes the use of SnapMirror to provide DR capabilities and VFM to provide the business continuity requirement.

11.3 Setting a client recovery policy

In this section, we discuss each VFM option when setting a client DR policy. We use a practical DR example to describe the VFM options.

On the machine Roman, you have a share called Legal. You want to replicate the Legal share onto the Itsotuc3 N series called Itsotuc3. You want the data replication from Roman to Itsotuc3 to run every eight hours. The users access the data through the namespace as \\domain\namespace\legal. If the primary link target located on Roman becomes unavailable, you want to fail over the users to the alternate link target share on Itsotuc3. You want to automatically perform the failover when the primary server goes down. (In production environments, you might want to do a manual failover.) You do not want to disrupt user access to legal data as a result of the failover.

A client DR policy can be used to set up the manual or automatic failover for data that resides in the namespace. It assumes that you already have the \\domain\namespace namespace. It also assumes that you have a multitargeted DFS link in the namespace called Legal (see 5.2.7, “Adding a link” on page 134). The two targets to the DFS link are called \\roman\legal and \\itsotuc3\legal, as shown in targets in Figure 11-1.

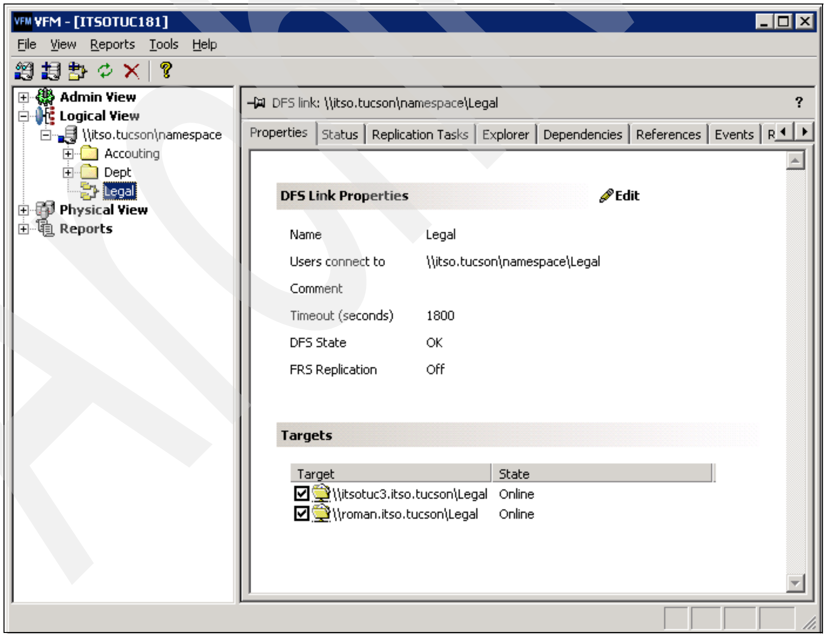

DFS Link Properties

Name	Legal
Users connect to	\\itso.tucson\namespace\Legal
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
☑ \\itsotuc3.itso.tucson\Legal	Online
☑ \\roman.itso.tucson\Legal	Online

Figure 11-1 VFM client displaying legal link properties

To access the client recovery policies on the VFM client, go to **Admin View** → **Disaster Recovery Policies** → **Client Recovery Policy** (see Figure 11-2).

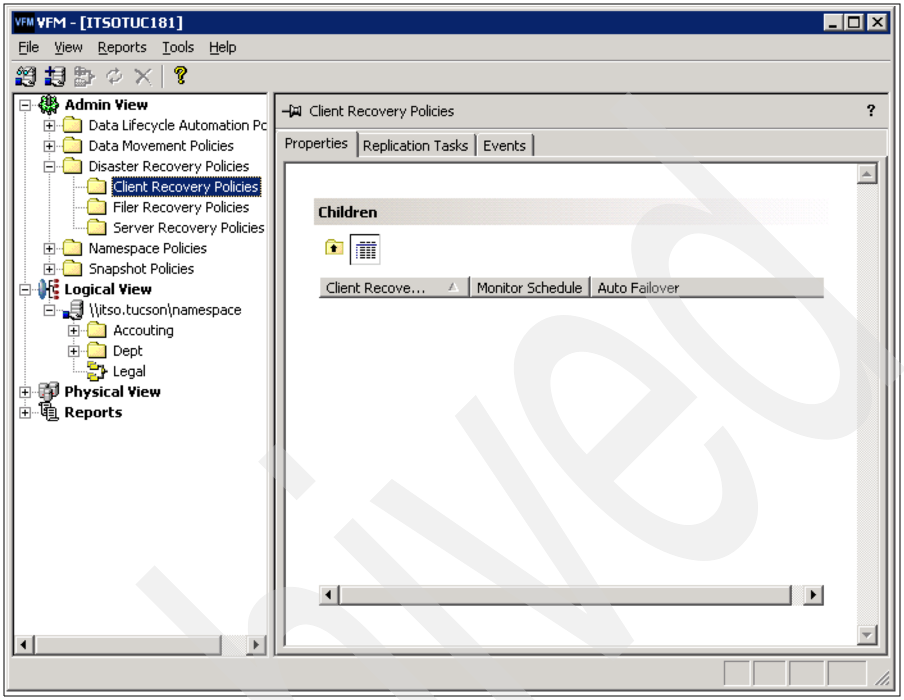

Figure 11-2 VFM Client Disaster Recovery Policies panel

Then right-click **Client Recovery Policy** (see Figure 11-3) and **New Client Recovery Policy**.

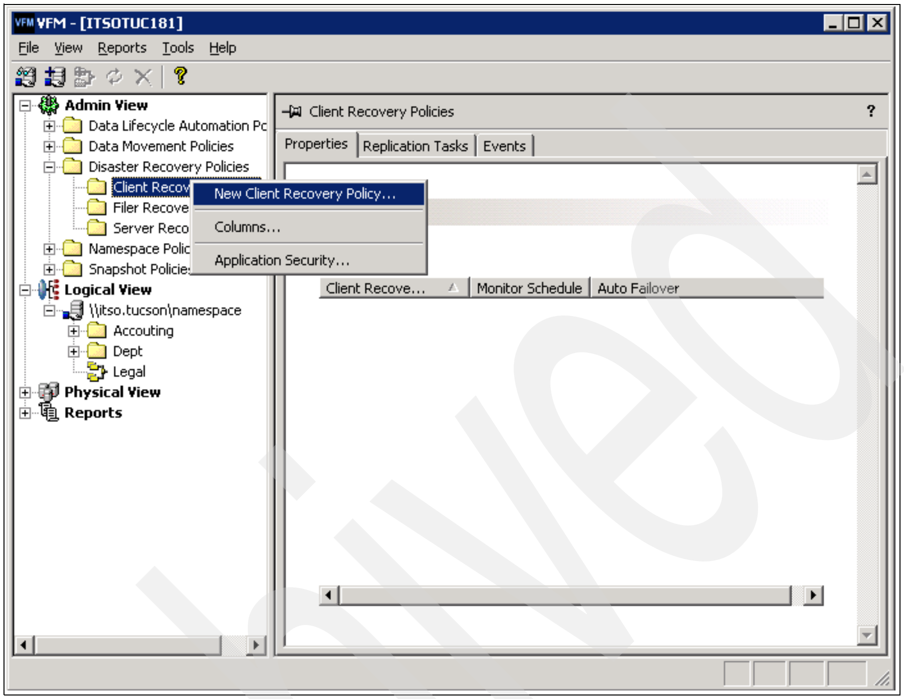

Figure 11-3 New Client Recovery Policy menu

In the following sections, we describe each option you can use in the process of setting up a client recovery policy creator.

11.3.1 General

The General panel is used to specify the name of the policy and some monitoring options (see Figure 11-4 on page 396):

- **Replication Method**

  This flexible method selects how the data is being replicated from the source to the destination. You can either choose to have VFM do the replication, or you can have SnapMirror replication between the source and destination. In this example, you are replicating data between a Microsoft Windows server and an N series device, so you select **VFM Replication** from the Replication Method list box.

- **Automatic failover**

  Monitoring options enables you to control how the failover is detected and how the failover behaves. You can select to automatically failover or to perform a manual failover. We recommend you perform a manual failover because you do not want a small network glitch to cause a false failover. Also, you want to be notified if something is wrong. After you investigate the problem, you then decide whether a failover is necessary.

  If you want to perform a manual failover, you select the DFS link in the policy, right-click it, and choose **Failover**. If you want to perform a manual failover of all the DFS links in a policy, you can also select the policy, right-click it, and select **Failover All Links**. Selecting this option initiates a failover of all the DFS links of the policy. A single click enables you to initiate a failover of hundreds of links in an automated fashion.

  > **Note:** It might be a while before users are notified of the new target. When VFM initiates a failover, it updates the DFS namespace. A domain-based DFS namespace is stored in Active Directory. VFM updates Active Directory with the information about the new DFS link target. The Domain Controllers replicate that information to all other Domain Controllers and other DFS root namespace servers. This process can take a while, depending on how the AD replication topology is configured, on network speeds, and on other factors.

- **Failover all links**

  A client DR policy can have hundreds of links in a single policy. Each link on the policy can point to a share on the same primary server. The secondary DR target for each link can be the same destination server. In other words, you are replicating data from one server to another for DR purposes. You can allow VFM to monitor each link in the policy individually. If the primary target of that specific link is unavailable, only that link is failed over. All other links are not failed over because only one link is monitored and unavailable.

  As an optimization, VFM enables you to monitor *only* one link in a policy. If that monitored link is unavailable, all links belonging to the policy are considered unavailable and failed over. This level of optimization reduces the amount of monitoring network traffic because you are essentially not monitoring each share of the machine. It is based on the premise that if one link target pointing to the source machine is unavailable, all link targets (on the same machine) are unavailable. It is assumed that if one link target is unavailable, the source machine is unavailable.

> **Note:** It is possible that using this option can produce a false positive. Perhaps someone inadvertently took the monitored share off line. In that case, all the links are failed over because the monitored link was not available.

If you select the **Failover all links** check box, the Only Monitor Link drop-down box becomes available, and you can select the link to be monitored.

- **Update associated policies on failover**

  This advanced feature enables you to update other policies (for example, a multiprotocol policy) when a change occurs in the DR policy initiated by the failover process.

- **Use the monitor on**

  This option enables you to distribute the monitoring and failover load associated with a DR policy. Normally, the VFM server monitors the health of the primary server to check for data availability. If the VFM server is extremely overloaded running other policies, you can choose to move the monitoring responsibilities to another server. This server has the role of a VFM Monitoring Agent. Only Monitoring Agents set to communicate to the current VFM server are listed in the Use the monitor on drop-down box of a client DR policy. Refer to Chapter 5, “DFS namespace” on page 117 for information about installing a VFM Monitoring Agent.

  Another time you can use the VFM Monitoring agent is when the VFM server is far from the source and destination locations of a DR policy. For example, your VFM server might be in Houston, your primary server hosting the data is in New York, and your DR site is in San Francisco. In this case, you might not want the VFM server in Houston monitoring the New York server to check for health and availability.

  The placement of the VFM server or Monitoring Agent is important for effective failover. The server checking the health of the source must be close to the user and essentially simulates user access to the data. If the VFM server cannot get to the primary data, it is assumed that users cannot get to the primary data. That might or might not be a correct assumption. This is why it is important not to do an automatic failover. You never want to run into a split brain scenario where half the users are accessing the data from the primary server, and the other half are accessing the data from the DR site because a failover was initiated automatically or manually. It is better to err on the side of caution and initiate a failover only when you are sure that the primary data is inaccessible.

For the purposes of this example, you clear the **Failover all links** and **Updating associated policies on failover** check boxes. Check the **Automatically failover** check box, and select the machine you are running the VFM server on from the Use the Monitor on list box, as shown in Figure 11-4.

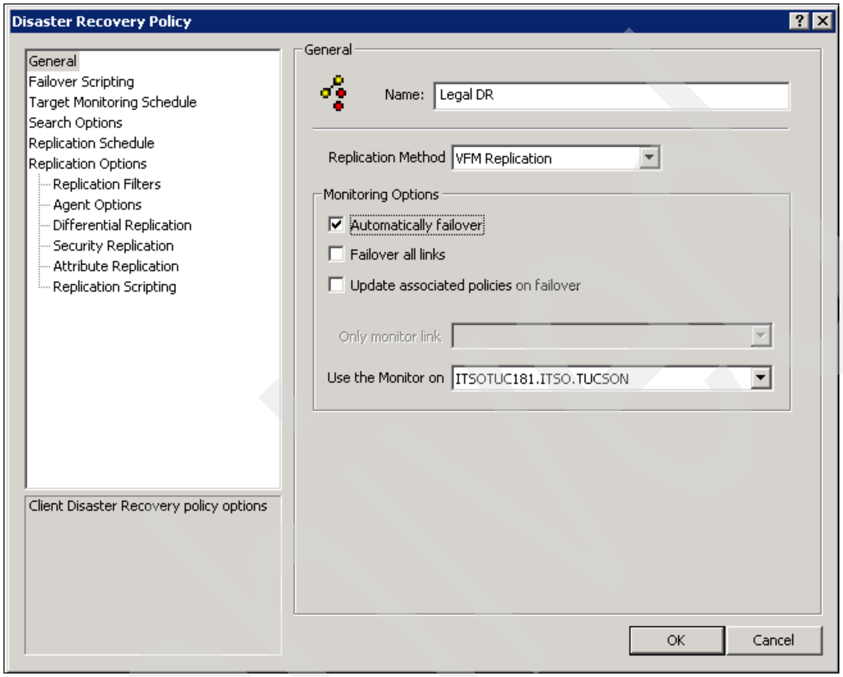

Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Client Disaster Recovery policy options

General

Name: Legal DR

Replication Method: VFM Replication

Monitoring Options

☑ Automatically failover

☐ Failover all links

☐ Update associated policies on failover

Only monitor link:

Use the Monitor on: ITSOTUC181.ITSO.TUCSON

Figure 11-4 Client recovery policy: General panel

Important: Make sure you select a monitor (from the Use the monitor on check box) to be used because if it is not set, the DR policy does not work properly, and a warning sign is displayed.

11.3.2 Failover scripting

Failover scripting is a powerful way of running scripts before or after a failover (see Figure 11-5 on page 397). For example, you can be notified if the primary data server fails. You can then decide to take action associated with the failover. You can write a batch script to send you an e-mail or write a message to the event log.

Sample batch scripts are provided in the example directory of the installation. You can perform these activities before or after a failover. You can use the Run

Batch Before Failover script to even determine whether a failover is to be performed. For example, you can have different failover thresholds during the weekdays and the weekends, and you can script these advanced checks.

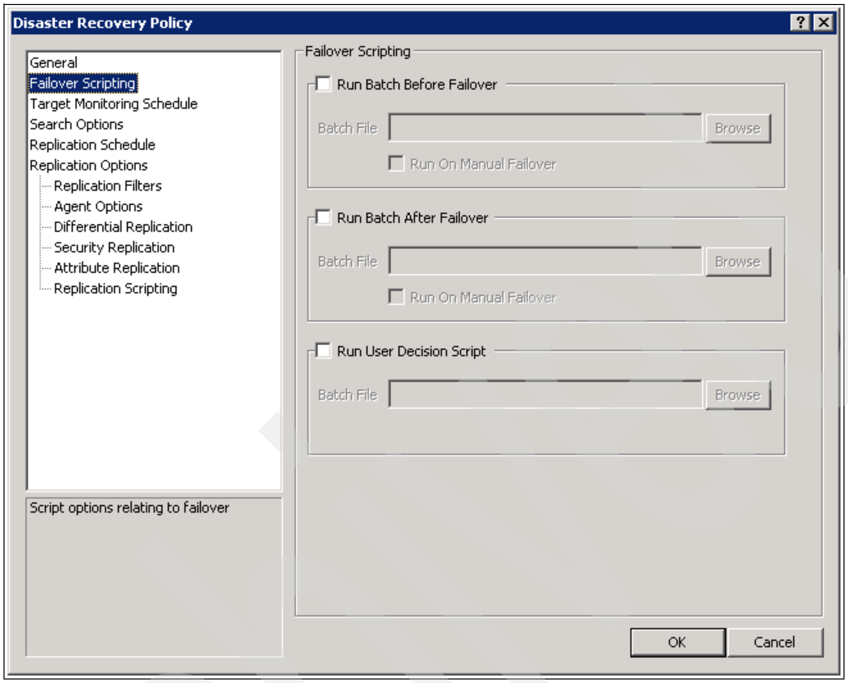

Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Script options relating to failover

Failover Scripting

□ Run Batch Before Failover

Batch File

□ Run On Manual Failover

□ Run Batch After Failover

Batch File

□ Run On Manual Failover

□ Run User Decision Script

Batch File

Figure 11-5 Client recovery policy: Failover Scripting panel

It is important to understand how VFM decides whether a share that is a target of a DFS link in a namespace is unavailable. VFM does not ping the machine hosting the share. It checks to see if it can enumerate the contents of the share. If VFM can enumerate the share, it means that the users can also get to the data. In essence, it simulates what the user does based on the assumption that the machine can respond to a ping, but the server service hosting the shares might be unavailable. A ping indicates only that the network stack on the primary server is working.

If you are an advanced administrator and you do not like the way VFM checks for data availability, you can check using the User Decision script. This script replaces VFM enumerating the share.

11.3.3 Target Monitoring Schedule

This panel provides a way to adjust the frequency of monitoring the data on the primary target of the Client DR policy (see Figure 11-6).

It specifies how often VFM monitors each link target in the policy to determine whether that link can be failed over to an alternate target. If the primary link target is unavailable, the failover occurs based on whether manual or automatic failover options are set. If the primary link target is unavailable, a red X is displayed on the VFM console corresponding to the policy and the DFS links to be failed. If the VFM console is not open, your only indication is through the execution of the Run Batch Before Failover script.

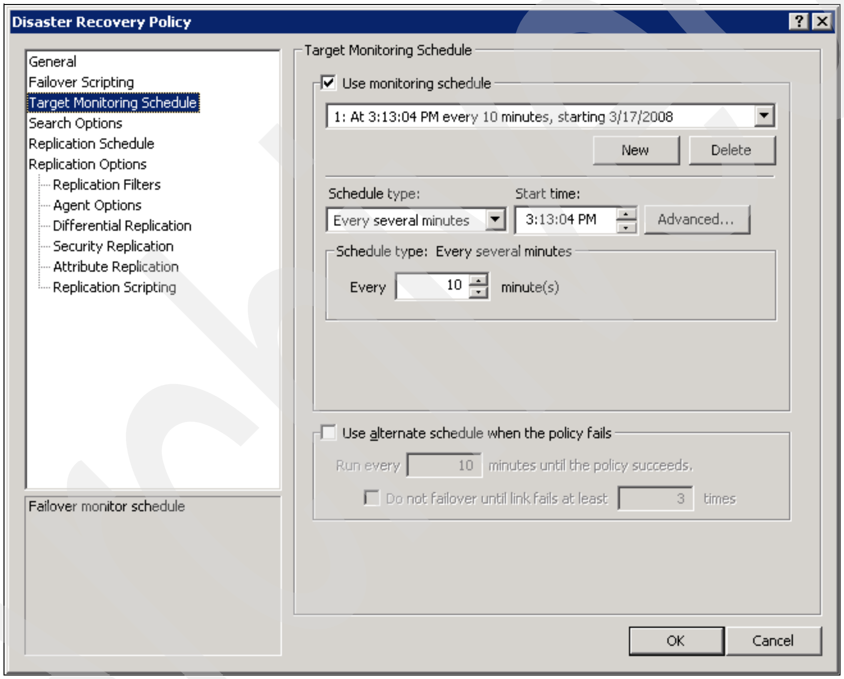
Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Failover monitor schedule

Target Monitoring Schedule

☑ Use monitoring schedule

1: At 3:13:04 PM every 10 minutes, starting 3/17/2008

Schedule type: Every several minutes

Start time: 3:13:04 PM

Schedule type: Every several minutes

Every 10 minute(s)

☐ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

☐ Do not failover until link fails at least 3 times

Figure 11-6 Client recovery policy: Target Monitoring Schedule panel

You must closely control the target monitoring schedule because it determines how often the VFM server monitors the primary server for availability.

If you set the target monitoring schedule to a very low value, you can be unnecessarily generating a lot of network traffic because the VFM server monitors the primary server for availability.

If you set this value too high, you can delay detecting the loss of the primary server. Your phone might ring before your policy notifies you that the primary server is unavailable.

It is best to consult your users, the recovery time objectives (RTO), and the recovery point objectives (RPO) before deciding on an appropriate value. In most instances, a target-monitoring interval value of 5 or 10 minutes is adequate for any policy. You can choose to have different polices with different monitoring intervals.

Note: If a link belongs to one DR policy, it cannot belong to another DR policy.

For the purposes of this example, specify a interval value of one minute for the target-monitoring schedule as shown in Figure 11-7. Use this value only for demonstration purposes because monitoring every minute is too frequent in production situations.

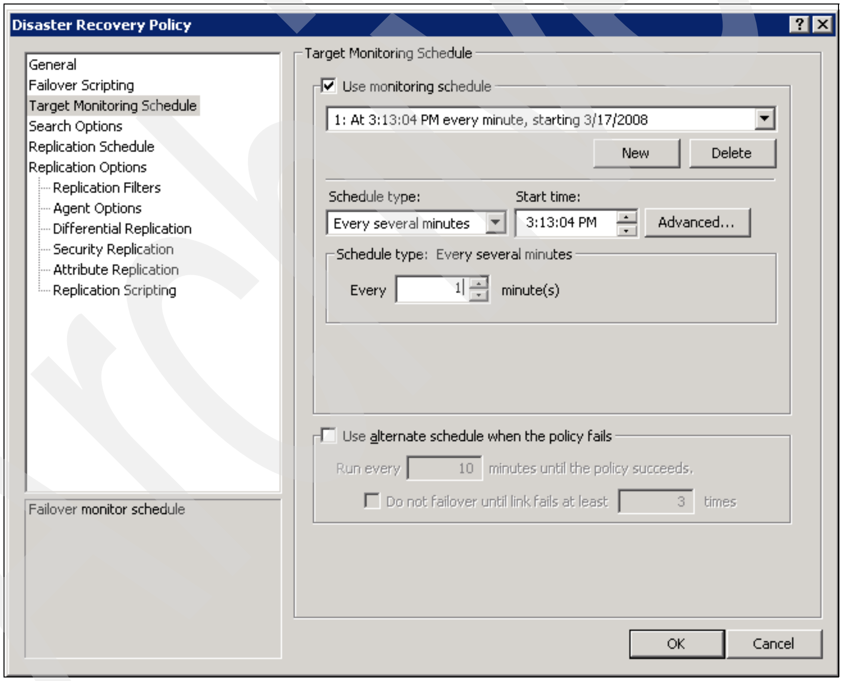

☑ Use monitoring schedule

1: At 3:13:04 PM every minute, starting 3/17/2008

Schedule type: Every several minutes

Start time: 3:13:04 PM

Schedule type: Every several minutes

Every 1 minute(s)

☐ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

☐ Do not failover until link fails at least 3 times

Figure 11-7 Disaster Recovery Policy: scheduling

11.3.4 Search options

The search options are advanced techniques that control the automatic addition and removal of links from a client DR policy. We recommend you use these options when tighter controls of policies are required.

The following options are available from the Search Options panel (see Figure 11-8 on page 401):

- **Automatically set the first link target online**

  This option provides a way to set which link target is the primary target and which link targets are considered the secondary link targets. To understand this, you need to understand DFS links, link targets, multitargeted DFS links, and the online and offline status of link targets. See 5.1.1, “Distributed file system overview” on page 118.

- **Automatically set the first link target online**

  This option provides a way for you to help VFM determine which link target to set online and which link targets to take offline. VFM has no way of knowing which server is the primary and which server is the secondary server, so you have to assist it in this instance.

For the purposes of our example, select the **Automatically set the first link target online** check box to be ready to adjust the link targets appropriately when the link is added to the policy. Keep the rest of the check boxes at their default values (see Figure 11-8 on page 401).

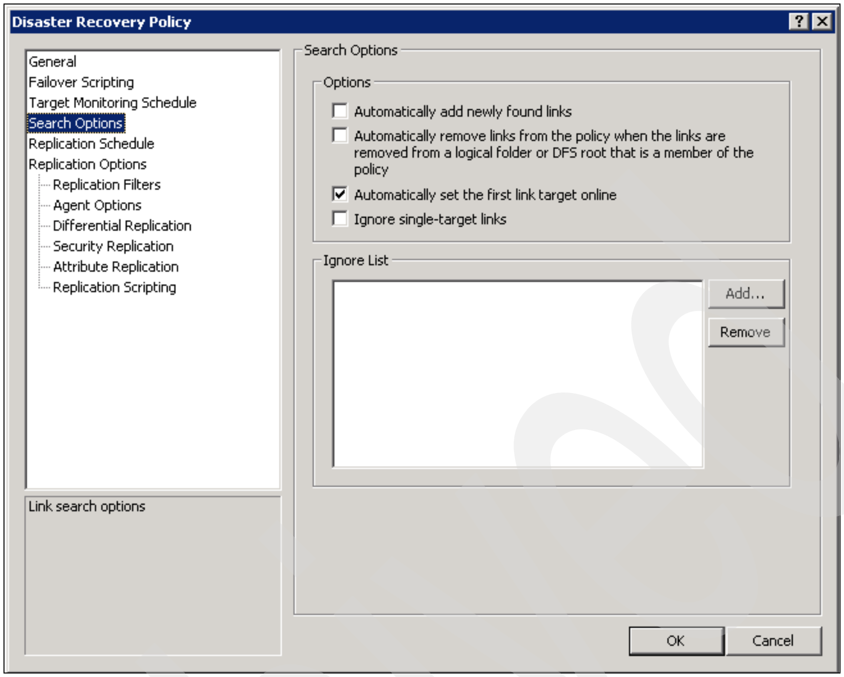

☐ Automatically add newly found links

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

☑ Automatically set the first link target online

☐ Ignore single-target links

Figure 11-8 Disaster recovery policy: Search Options panel

11.3.5 Replication schedule

A replication schedule in a client DR policy determines the frequency of replication between the primary online link target and all secondary offline link targets.

In some instances, it is satisfactory to replicate data between the primary and secondary targets once a night. In other instances, you might need to replicate the data every four or eight hours. The frequency of replication for any policy is always based on the business need and the business value of the data.

> **Important:** Do not schedule the replication to occur more frequently than the time required to complete a run of the replication. For example, if a replication is expected to take one hour to complete, do not schedule the replication to run every 30 minutes. If a replication is running and the next scheduled run of the replication elapses, the replication waits until the current run of the policy finishes.

The VFM console does not have to be open when the data is being replicated. The replication is initiated by the VFM server. The policy is run based on the time settings and schedule of the VFM server, which might or might not be in the same time zone of the source or destination machines. Take care when setting the run times of a policy and take into account the source and destination time zones and where the VFM server is located.

Replication always takes place from the first server listed in the Client DR policy to all servers listed below that server. Cascading of replication does not occur. Replication always follows a one-to-many topology from the primary server to all other servers.

Note: The replication schedule is different than the monitoring schedule. The replication schedule deals with how frequently you copy data from the source to the DR sites; the monitoring schedule is a measure of how frequently you check to see whether the primary link target is available.

For the purposes of our example, set the replication interval for 5 minutes so that changes are constantly copied (see Figure 11-9). VFM does an initial replication before the DR policy is active.

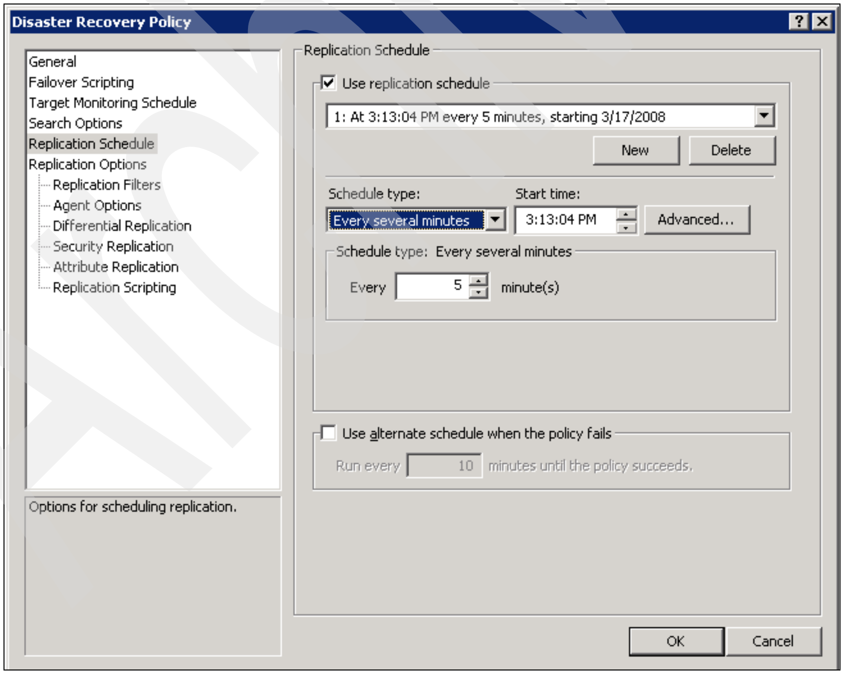

Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options for scheduling replication.

Replication Schedule

☑ Use replication schedule

1: At 3:13:04 PM every 5 minutes, starting 3/17/2008

Schedule type: Every several minutes

Start time: 3:13:04 PM

Schedule type: Every several minutes

Every 5 minute(s)

□ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

Figure 11-9 Disaster Recovery Policy: Replication Schedule panel

Note: In production environments, never set the replication interval at this low value. In our example, we use these low values to quickly demonstrate exactly what you see on the VFM client.

11.3.6 Replication options

Replication options enable you to provide control of data movement when VFM performs the replication. The following options are available (see Figure 11-10 on page 407):

- **Include subfolders**
- **Delete orphaned files in destination folders**

  This option controls what happens if extra files are on the destination, that is, files that do not correspond to any files on the source. If you select to delete the orphaned files, the Replication Agent makes the destination a mirror image of the source.
- **Allow loss of additional file streams (CIFS only)**

  In some instances, the destination file system does not support alternate data streams. Some Microsoft Windows applications store data in alternate data streams. Macintosh-created data files make extensive use of streams (called *resource forks*). In most cases, you might or might not know whether alternate data streams are used. If VFM cannot copy a file because the destination does not support streams, the only way to copy the file is to select the **Allow loss of additional file streams** check box.
- **Enable Copy-in-place**

  Many optimizations in a data movement policy are a compromise between speed and safety. The Copy-in-place option in the Replication Options can help a replication policy finish faster with a slight security risk. Copy-in-place comes into effect during a incremental copy. Assume the following scenario for a replication policy that moves files from one machine to another:

  – A file xyz.doc is on the source.

  – The xyz.doc file is copied to the destination.

  – The file is modified on the source.

  – Because the xyz.doc file on the source is newer than the destination, xyz.doc on the source must replace the file on the destination.

If Copy-in-place is not checked, the following events occur:

- The file xyz.doc is copied to the destination as a temporary file.
- When the copy is complete, the file xyz.doc is deleted on the destination, and the newly copied temporary file is renamed as xyz.doc.

If Copy-in-place is checked, the following events occur:

- The file xyz.doc is truncated to 0 bytes on the destination.
- The file from the source is copied to the truncated file.

Checking Copy-in-place is faster because no rename operations are required. Copy-in-place is preferred when the source and destination are on the same LAN or connected by a high-speed network link.

If the source and destination are connected by a low-speed link or the link is not reliable, clear Copy-in-place and enable the safe rename technique to work. This way, if the network link breaks when the copy is in progress, you are not left with a zero-length file on the destination.

- **Use NTFS Change Journal (CIFS only)**

  The NTFS change journal is a “file” that maintains a list of accessed, changed, created, and deleted files.

  VFM can optimize the search of changed source files on Microsoft Windows machines by integrating with the NTFS change journal. If the Replication Agent scans the source NTFS change journal, it can detect the changed, deleted, and newly created files in a fraction of the time compared to doing a full file scan of TBs of data. This way, a replication resynchronization proceeds much faster.

  The change journal file can wrap, thus resulting in lost changes. The maximum change journal size on the volume is based on space you allocate. If the change journal wraps, VFM resorts to a full scan of the source to determine the changed files.

  The NTFS change journal is only available when the source is a Microsoft Windows machine. In all other cases, this option is silently ignored.

- **Use snapshot data**

  During any replication policy run, one or more files can be open on the source. This is usually the case for PST files and Access databases where the user has an exclusive lock on the file. If VFM can get a read lock on the file, it copies the file. Depending on how the policy is set up, VFM can either skip the opened file or copy the file as best as it can from the snapshot.

  VFM can integrate with Volume Shadow Service (VSS) on Microsoft Windows 2003 or N series snapshots.

If you expect files to be open, VFM can integrate with snapshots. VFM can take a temporary snapshot of the source volume and copy all the files from the source snapshot. Once the files are copied, the source snapshot is automatically deleted. If the source is not Microsoft Windows 2003 or N series, use the Retry failed file option.

Files not copied are logged as errors and are reported in the manifest based on the settings of the policy. If one or more files are not copied successfully, the policy is reported as a unsuccessful run. This is because VFM chooses to warn you that the policy did not copy all files completely.

- **Copy files**

  When a policy runs (either based on the replication schedule or a manual run), VFM examines the source and destination to determine which files must be copied. VFM uses the file name, the date, the attributes, and the file size to check whether two files are identical. Based on the policy settings, if two files are identical, no data is transferred. If the files on the source and destination are different, whether a file is copied is controlled according to this setting.

  The options available from the Copy files drop-down list box if the destination file is different or missing and if the destination is older or missing. At times, the file on the destination is newer than the source. You must take care when setting the replication options for any policy.

- **Retry failed file opens**

  This option also determines how VFM proceeds when open files are on the source of the replication. This option can make the Replication Agents retry the open files for a number of times, with the specified interval between each try.

- **Abort if no file is sucessfully opened in**

  If the Replication Agent cannot open any files for a certain time, the replication fails. With this option, you can set how much time elapses before the replication aborts.

- **Event details**

  A replication policy can store a list of the files that failed during replication or the list of all files moved during the run of a policy.

  We recommend you use the Only list files that encounter errors option, only when you require a full list of replicated files.

If you select to list all files on large replications, you can overload the Replication Agent and VFM server by transferring all that data. We do not recommend passing the file list of full migrations with millions of files because doing so generates a significant amount of network traffic. The manifest file is stored on the machine running the Replication Agent and is transferred each time it is requested.

If files cannot be copied between the source and the target, the policy is considered to complete with errors. Any open files that are not copied result in an error. If an alternate schedule exists, the policy then runs based on the alternate schedule. If the VFM server cannot communicate with the Replication Agent performing the data movement, the VFM server forces the policy into an error state.

For the purposes of our example, select the following options (see Figure 11-10 on page 407):

- **Include subfolders**
- **Delete orphaned files in destination folders**
- **Enable Copy-in-place**
- **Event details** (only list files that encounter errors)

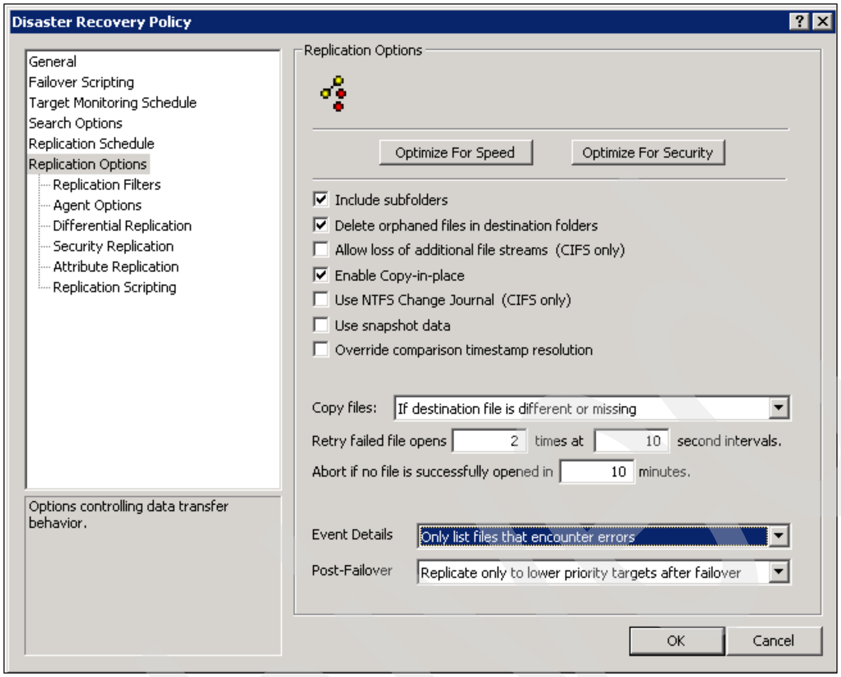

Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options controlling data transfer behavior.

Replication Options

☑ Include subfolders
☑ Delete orphaned files in destination folders
☐ Allow loss of additional file streams (CIFS only)
☑ Enable Copy-in-place
☐ Use NTFS Change Journal (CIFS only)
☐ Use snapshot data
☐ Override comparison timestamp resolution

Copy files: If destination file is different or missing

Retry failed file opens 2 times at 10 second intervals.

Abort if no file is successfully opened in 10 minutes.

Event Details: Only list files that encounter errors

Post-Failover: Replicate only to lower priority targets after failover

Figure 11-10 Disaster Recovery Policy: Replication Options panel

11.3.7 Replication filters

This option provides a more flexible way of controlling the files to be copied or excluded in a replication (see Figure 11-11 on page 408). At times, you can choose not to copy temp files. Or you might want to copy files, but you want to exclude MP3 files from being placed on the destination N series. Or you have a large share, and you want to exclude two top-level folders from being copied. Instead of specifying each folder to copy, you can choose to specify the folders to exclude. Not all the data has to be on the DR site. The subset of the data that must be available during a failover is replicated, and this option provides fine-grained control of the replication.

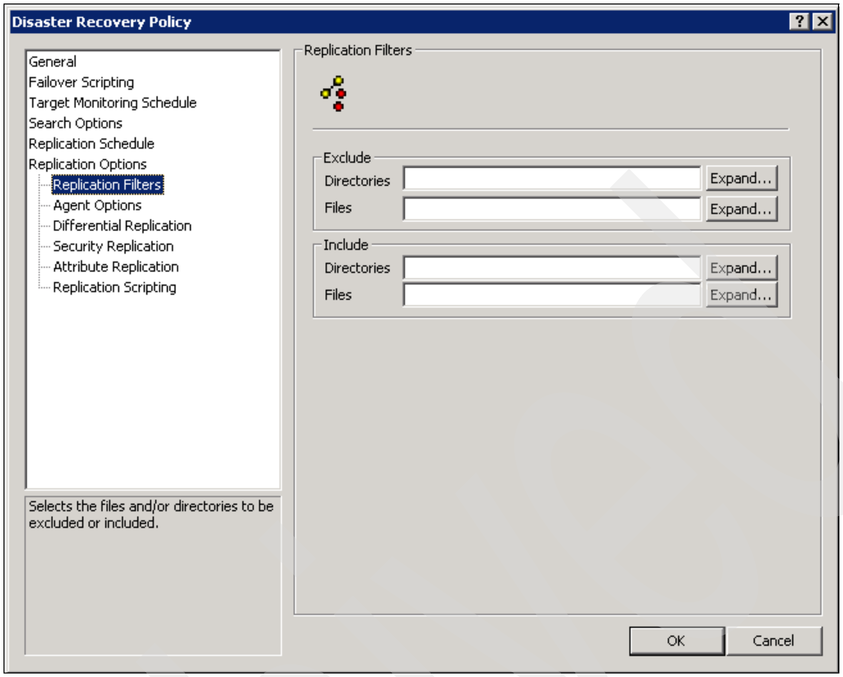

Selects the files and/or directories to be excluded or included.

Figure 11-11 Disaster Recovery Policy: Replication Filters panel

11.3.8 Agent options

By default, VFM uses the Replication Agent on the destination to pull data from the source. This method is preferred because the agent on the destination knows how much data it can handle based on the load on the destination machine.

The available options are as follows and are shown Figure 11-12 on page 410):

- **On the source side**

  When the Replication Agent cannot run on the destination (for example, the destination is an N series storage system), you can use this option to force the agent to run on the source and push data to the destination.

- **On the destination side**

  This is the default behavior when you select VFM selects. This method is preferred because the agent on the destination knows how much data it can handle based on the load on the destination machine. Using this option, you can force this behavior, or just make it clear that it is being used.

- **On this machine or corresponding proxy**

  When Replication Agents cannot be deployed on the source or destination, you must use a Replication Agent on another machine, which acts as a proxy and makes the replication possible.

  One typical scenario for enabling this option is when you want to migrate data from one N series to another. Because Data ONTAP does not allow agents to be installed, you require a proxy. If you choose VFM selects, VFM can automatically select a machine to act as the proxy, but it does not consider the geographical distance, link cost, or bandwidth because it does not have this information.

  If Replication Agents cannot be deployed on the source or the destination, manually select a machine to act as a proxy as a best practice.

  Remember that you do not have to manually install Replication Agents. Based on the policy settings, the Replication Agent is automatically deployed to the specified machine, regardless of whether it is a source, destination, or a proxy.

- **In this agent group**

  Agent grouping provides a way to share the load between multiple Replication Agents. When the replication job runs, VFM sends the policy to the next available agent in the group. A round robin technique is used to determine which Replication Agent in the agent group performs the policy. Using Agent Groups prevents a Replication Agent from being overloaded while sharing the load across multiple agents to perform replications.

- **VFM selects**

  This default method transfers the responsibility to choose where to run the agent to VFM.

  This option first tries the destination; if it fails, it tries the source. If it cannot use either source or destination, it uses the Replication Agent on the machine where the VFM server installed. Check the on this machine or corresponding proxy check box as a best practice when you require a proxy.

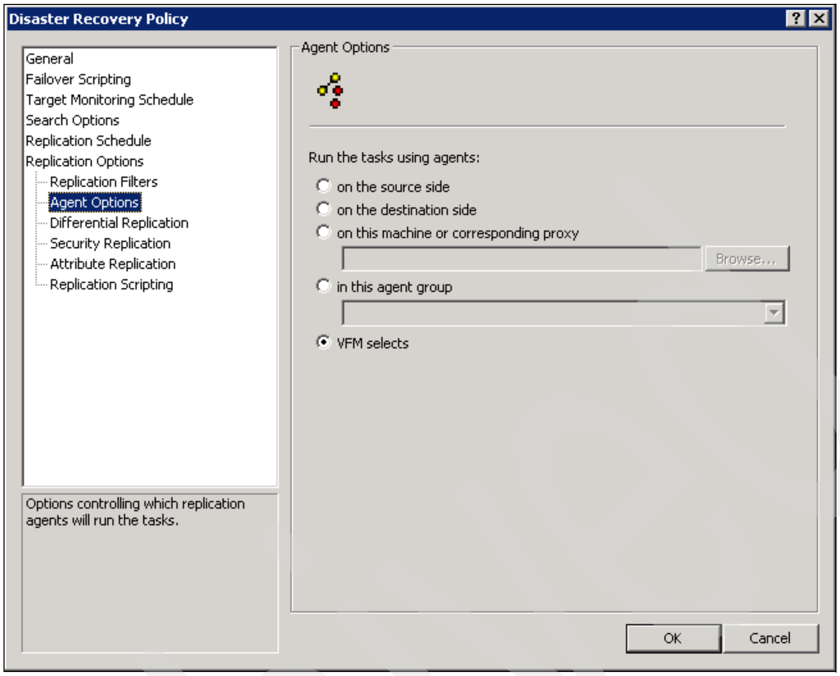

Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options controlling which replication agents will run the tasks.

Agent Options

Run the tasks using agents:

○ on the source side
○ on the destination side
○ on this machine or corresponding proxy
○ in this agent group
● VFM selects

Figure 11-12 Disaster Recovery Policy: Agent Options panel

Remember that Replication Agents run under the credentials of the service account that VFM runs under. This service account must have administrator rights on the source and destination machines.

11.3.9 Differential replication

VFM provides a feature using byte level replication (BLR) to control data movement across slow links. BLR (see Figure 11-13 on page 411) minimizes the data sent across the wire, especially with slow network links, by sending only the changed bytes.

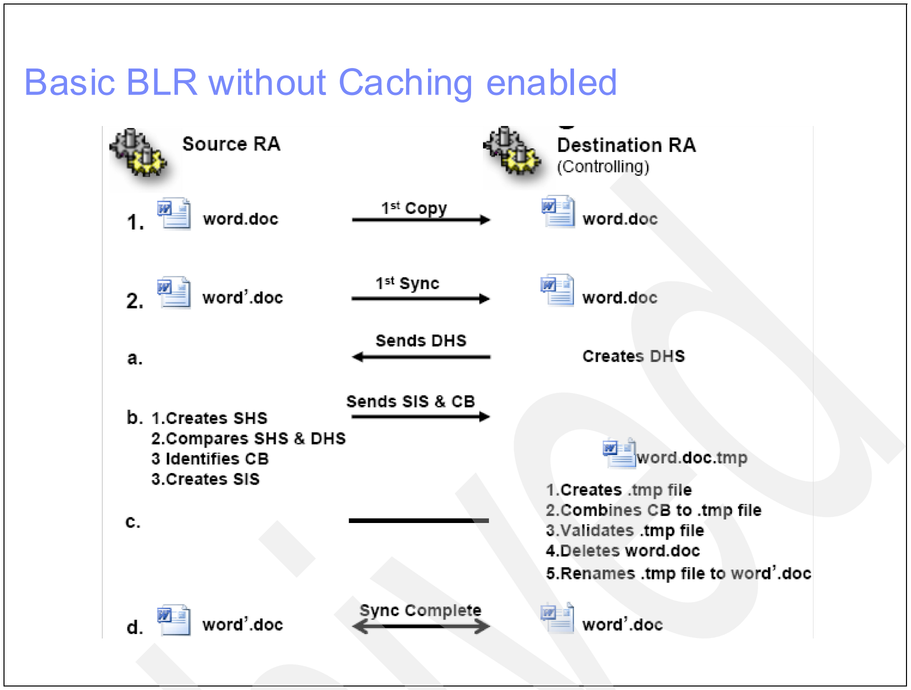

Basic BLR without Caching enabled

Source RA

Destination RA
(Controlling)

1. word.doc — 1st Copy → word.doc

2. word'.doc — 1st Sync → word.doc

a. ← Sends DHS — Creates DHS

b. 1.Creates SHS
2.Compares SHS & DHS
3 Identifies CB
3.Creates SIS

Sends SIS & CB →

word.doc.tmp

c. 1.Creates .tmp file
2.Combines CB to .tmp file
3.Validates .tmp file
4.Deletes word.doc
5.Renames .tmp file to word'.doc

d. word'.doc ↔ Sync Complete ↔ word'.doc

Figure 11-13 byte level replication

Suppose you have a 10 MB PowerPoint file copied from the source to the destination share. Assume that the user at the source changes one or two slides in the presentation. Instead of sending the 10 MB changed file again from the source to the destination, VFM can intelligently send the changes across the wire.

While the time required in computing the changed bytes in extremely large files can be significant, network usage is minimized. You sacrifice the extra CPU cycles to figure out the changed bytes so that the limited resource of network bandwidth is optimized. BLR requires two Replication Agents to work together to enable a successful replication.

Use BLR only across slow links. In LAN speed networks, the overhead of BLR computation is more than the benefit it provides. BLR is also not effective for graphic files (such as JPG and BMP) where a single change between images is scattered through all parts of the file.

BLR works when the destination Replication Agent computes a rolling hash set across all the blocks of the file at the destination. The destination Replication Agent sends the rolling hash set as a file signature to the source Replication Agent. This results in a small amount of data transferred from destination to source. The source Replication Agent computes the rolling hash set of the changed file at the source. When different hash blocks are encountered, the source Replication Agent computes an instruction set to send from the source to the destination. The instruction set is used to insert or remove data blocks from the file at the destination to make it look like the source. Using BLR over slow WAN links results in an 80% savings in network traffic of slightly changed data files.

For the purposes of this example, do not enable differential replication.

11.3.10 Replication security

Every replication requires the files be copied accurately and completely. VFM provides several mechanisms to control the copying of files. Where possible, VFM signals errors when files cannot be copied completely or correctly.

Some Active Directory deployments use local groups to secure the access to file resources on the old server. In NT4 domains, Microsoft recommends the use of local groups. A local group is only valid on the source system on which it is defined. If the security on a file is controlled by the permissions of a local groups and that file is copied to a destination machine, the security is lost. This means that the users who had access to the files on the old server might not have access to the replicated copy.

VFM provides a way to handle this situation as shown in Figure 11-14 on page 413. VFM can create a local group on the destination with the same name as the local group on the source. All domain users and domain groups that are members of the old local group on the source server are automatically added to the newly created local group of the destination server. This way, users that had access to the data through the old local group on the source server have access to the data on the new server through the newly created local group.

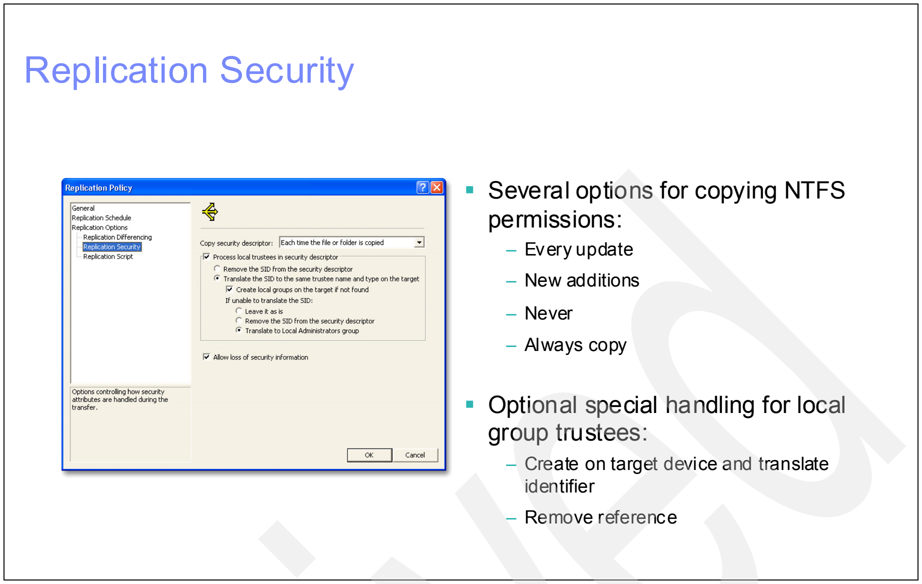
Replication Security

Copy security descriptor: Each time the file or folder is copied

☑ Process local trustees in security descriptor

- ○ Remove the SID from the security descriptor
- ● Translate the SID to the same trustee name and type on the target
  - ☑ Create local groups on the target if not found
  - If unable to translate the SID:
    - ○ Leave it as is
    - ○ Remove the SID from the security descriptor
    - ● Translate to Local Administrators group

☑ Allow loss of security information

Options controlling how security attributes are handled during the transfer.

- Several options for copying NTFS permissions:
  - Every update
  - New additions
  - Never
  - Always copy
- Optional special handling for local group trustees:
  - Create on target device and translate identifier
  - Remove reference

Figure 11-14 Replication security

At times, you care only about replicating the data from one location to another without regard to the security of the data on the new system. For example, if you replicate data from a Microsoft Windows server to a UNIX-style qtree on the new N series, it is advisable to check the Allow loss of security information check box.

If the source machine or the destination machine is a Domain Controller, deselect the Allow loss of security information check box to process local groups (technically called *local trustees*). This is because the Domain Controller does not support local groups.

For the purposes of our example, keep the default check boxes and radio buttons selected (see Figure 11-15).

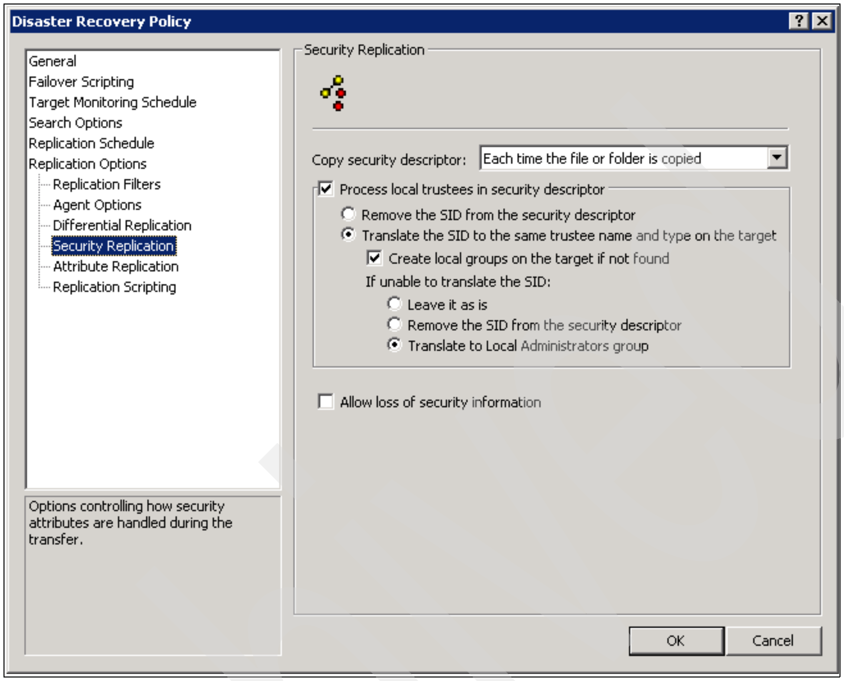
Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options controlling how security attributes are handled during the transfer.

Security Replication

Copy security descriptor: Each time the file or folder is copied

☑ Process local trustees in security descriptor

○ Remove the SID from the security descriptor

● Translate the SID to the same trustee name and type on the target

☑ Create local groups on the target if not found

If unable to translate the SID:

○ Leave it as is

○ Remove the SID from the security descriptor

● Translate to Local Administrators group

□ Allow loss of security information

Figure 11-15 Disaster Recovery Policy: Security Replication panel

11.3.11 Attribute replication

In a replication policy, you control the attributes used to compare two files to determine whether they are identical or different. By default, VFM compares the names, size, dates, and attributes. You can override this by checking or clearing the check box for attribute replication. You can also specify how the attribute is set on the destination file after it has been copied. For example, you can clear the archive attribute of a copied file so that the file is backed up on the replication target.

When copying files, you can preserve the last access time of the file on the source. Even though VFM copies the source file, it resets the last accessed time on the source file, which is useful if HSM or other reporting software is based on the last access time. Overhead is involved on the Replication Agent in preserving the last access time on the source. So if this feature is not required, uncheck the Preserve last access time on source check box. Because the DR policy is replicating data from one machine to the next, we recommend you check this box when dealing with production machines (see Figure 11-16).

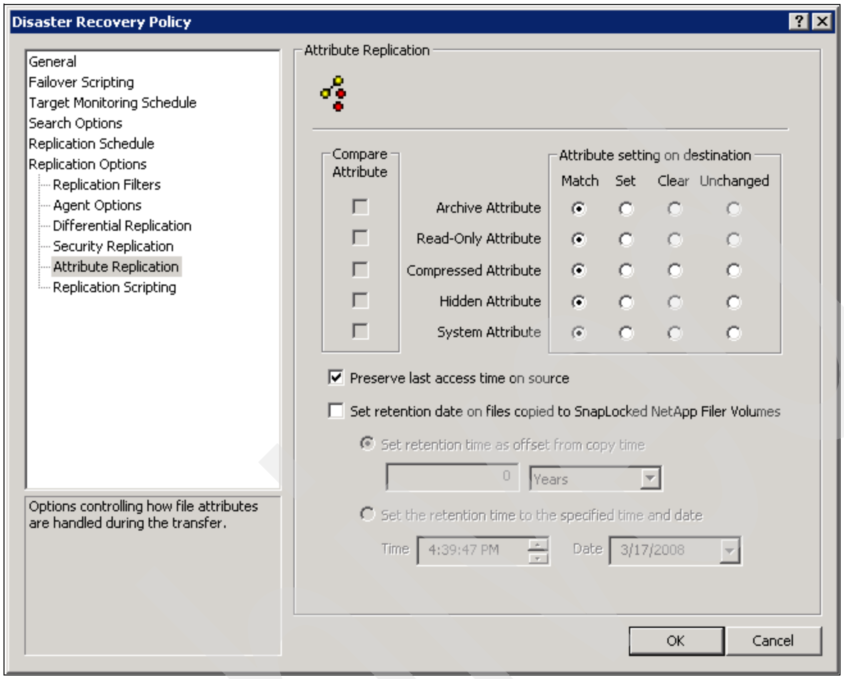

Disaster Recovery Policy

- General
- Failover Scripting
- Target Monitoring Schedule
- Search Options
- Replication Schedule
- Replication Options
  - Replication Filters
  - Agent Options
  - Differential Replication
  - Security Replication
  - Attribute Replication
  - Replication Scripting

Options controlling how file attributes are handled during the transfer.

Attribute Replication

Compare Attribute		Attribute setting on destination: Match	Set	Clear	Unchanged
☐	Archive Attribute	●	○	○	○
☐	Read-Only Attribute	●	○	○	○
☐	Compressed Attribute	●	○	○	○
☐	Hidden Attribute	●	○	○	○
☐	System Attribute	●	○	○	○

☑ Preserve last access time on source

☐ Set retention date on files copied to SnapLocked NetApp Filer Volumes

● Set retention time as offset from copy time

0 Years

○ Set the retention time to the specified time and date

Time 4:39:47 PM Date 3/17/2008

Figure 11-16 Disaster Recovery Policy: Attribute Replication panel

For the purposes of this example, check the **Preserve last access time on source** check box as shown in Figure 11-16. Then click **OK** to create the policy.

After the policy is created, the policy is listed as shown in Figure 11-17.

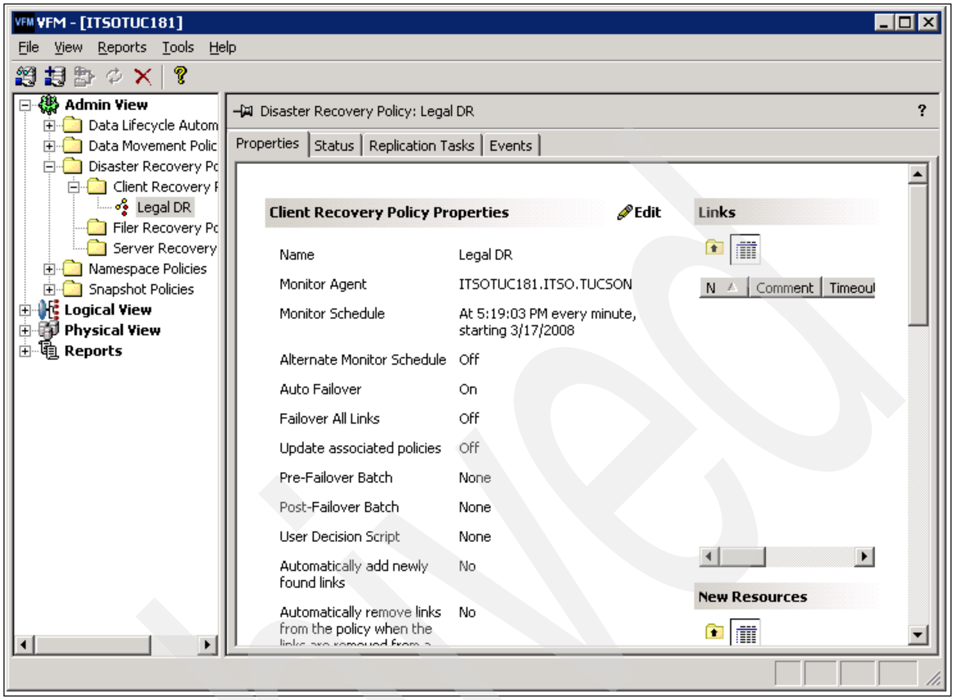

Client Recovery Policy Properties

Name	Legal DR
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Monitor Schedule	At 5:19:03 PM every minute, starting 3/17/2008
Alternate Monitor Schedule	Off
Auto Failover	On
Failover All Links	Off
Update associated policies	Off
Pre-Failover Batch	None
Post-Failover Batch	None
User Decision Script	None
Automatically add newly found links	No
Automatically remove links from the policy when the links are removed from a	No

Figure 11-17 Disaster Recovery Policy: listing new created policy

11.3.12 Adding a DFS link to a DR policy

For the DR policy to be complete, it requires a DFS link, but you do not add it during policy creation itself. To add a DFS link to the DR policy, you can either drag-and-drop the Legal DFS link from the logical view to the policy. Or you can perform the following steps:

1. Right-click the created DR policy.
2. Click **Add Link to policy** (see Figure 11-18 on page 417).

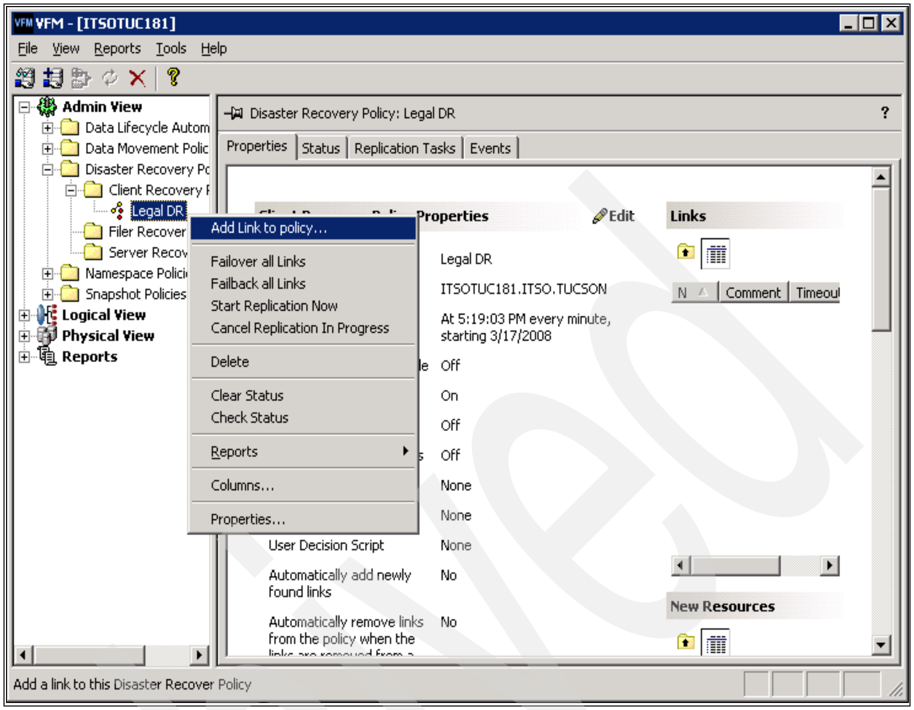

Legal DR

ITSOTUC181.ITSO.TUCSON

At 5:19:03 PM every minute, starting 3/17/2008

Off

On

Off

Off

None

None

User Decision Script None

Automatically add newly found links No

Automatically remove links from the policy when the No

Figure 11-18 Add Link to policy menu

3. Click the **Legal** link (see Figure 11-19).

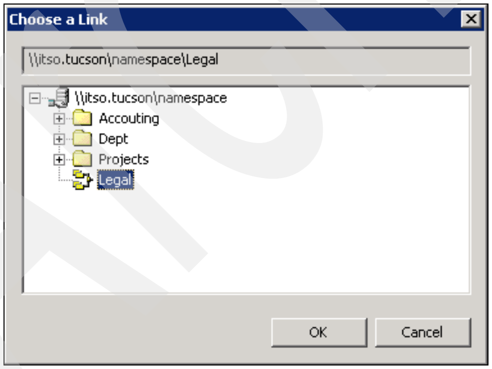

Figure 11-19 Choosing Legal link to add to policy

Note: You must add only multitargeted DFS links to the policy. If you add a single-targeted DFS link, VFM runs the wizard that creates and selects a second DR or failover target.

4. At this point, you must communicate to the VFM which target on the multitargeted DFS link you wish to make the primary. Click **Next** (see Figure 11-20).

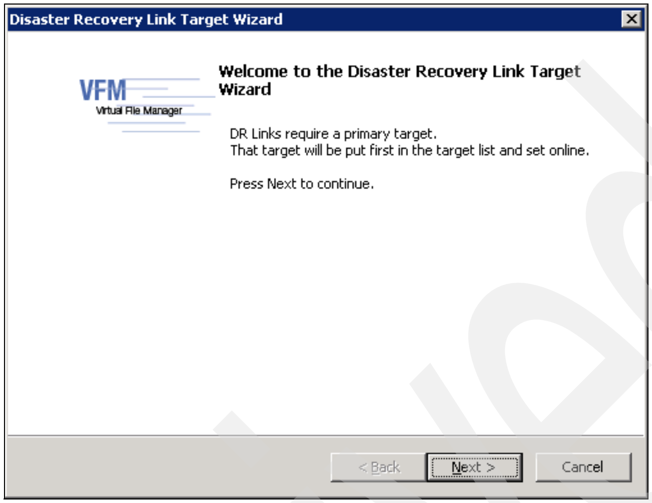

Disaster Recovery Link Target Wizard

Welcome to the Disaster Recovery Link Target Wizard

DR Links require a primary target.
That target will be put first in the target list and set online.

Press Next to continue.

Figure 11-20 Disaster Recovery Link Target Wizard

5. Select your primary target (see Figure 11-21). All targets on the DFS link are displayed.

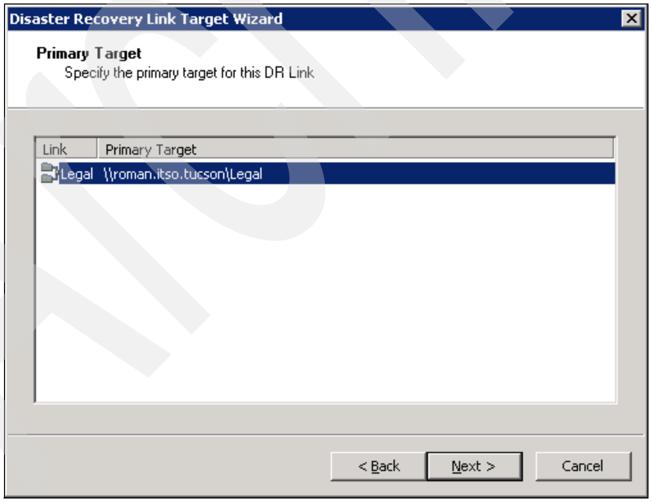

Disaster Recovery Link Target Wizard

Primary Target
Specify the primary target for this DR Link

Link	Primary Target
Legal	\\roman.itso.tucson\Legal

Figure 11-21 Disaster Recovery Link Target Wizard: selecting primary target

6. Click **Finish** (see Figure 11-22).

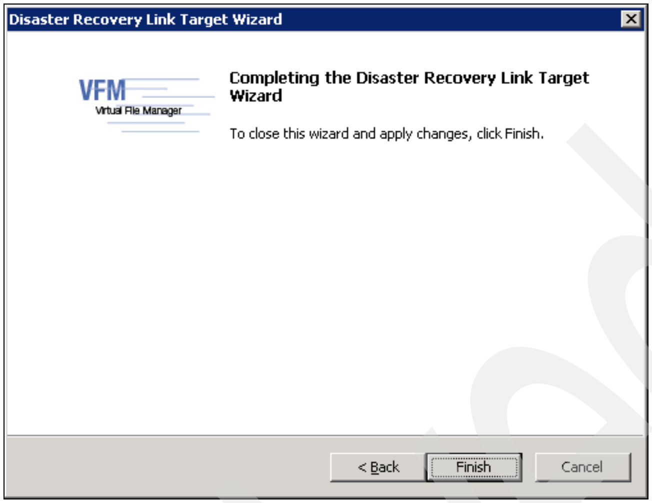
Disaster Recovery Link Target Wizard

Completing the Disaster Recovery Link Target Wizard

To close this wizard and apply changes, click Finish.

Figure 11-22 Disaster Recovery Link Target Wizard: finish

The DFS link is added to the policy and the primary target is selected.

After adding the DFS link, you might see a yellow warning sign adjacent to the policy (see Figure 11-23 on page 420). You have to wait for the first replication and monitoring to take place for the warning icon to disappear (see Figure 11-24 on page 421). This could take a few minutes based on the monitoring schedule.

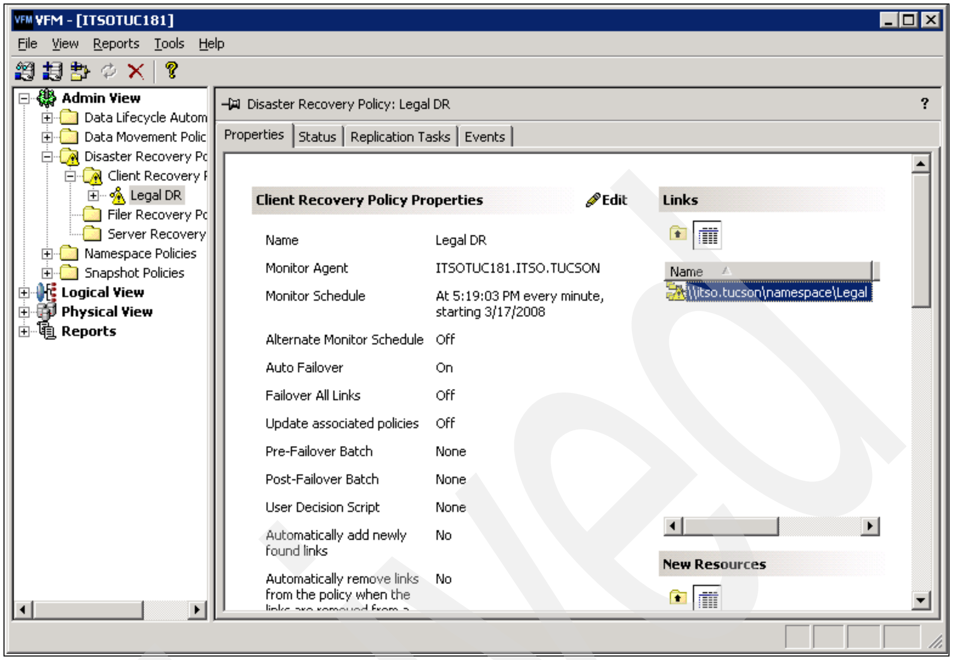

Client Recovery Policy Properties

Name	Legal DR
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Monitor Schedule	At 5:19:03 PM every minute, starting 3/17/2008
Alternate Monitor Schedule	Off
Auto Failover	On
Failover All Links	Off
Update associated policies	Off
Pre-Failover Batch	None
Post-Failover Batch	None
User Decision Script	None
Automatically add newly found links	No
Automatically remove links from the policy when the links are removed from a	No

Links

Name
\\itso.tucson\namespace\Legal

Figure 11-23 DR policy with new DFS link

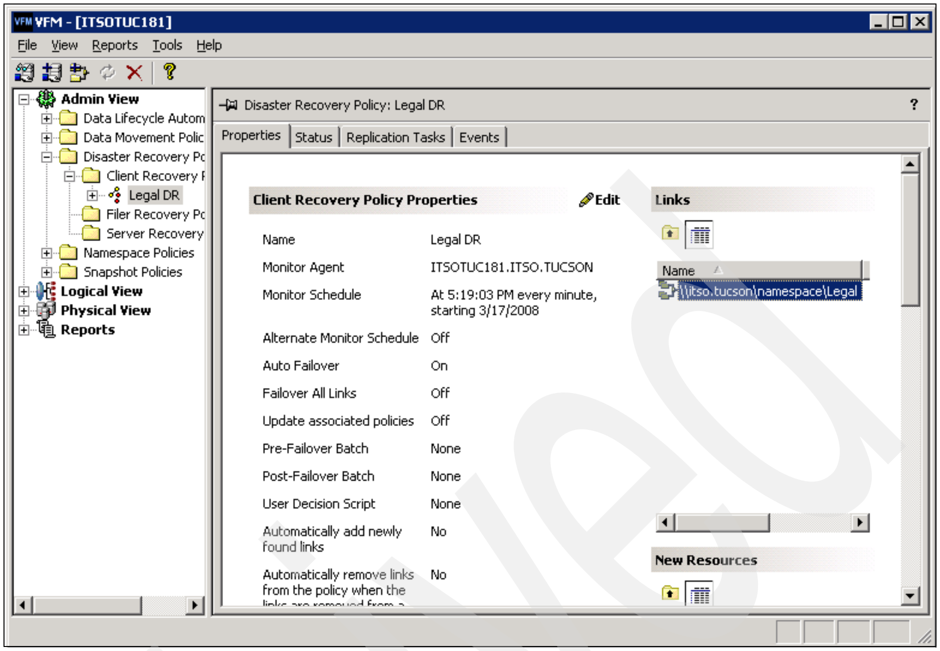

Client Recovery Policy Properties

Name	Legal DR
Monitor Agent	ITSOTUC181.ITSO.TUCSON
Monitor Schedule	At 5:19:03 PM every minute, starting 3/17/2008
Alternate Monitor Schedule	Off
Auto Failover	On
Failover All Links	Off
Update associated policies	Off
Pre-Failover Batch	None
Post-Failover Batch	None
User Decision Script	None
Automatically add newly found links	No
Automatically remove links from the policy when the links are removed from a	No

Links: \\itso.tucson\namespace\Legal

Figure 11-24 DR policy with DFS links after with no warning displayed

Notice on the DFS Link Properties panel that the check boxes adjacent to the link targets are converted to radio buttons (see Figure 11-25 on page 422).

The active (primary) target is the target with the radio button enabled. The offline (DR) target is the target with the radio button cleared. You can adjust the active and DR target using the yellow arrow icons at the right side of the window.

Always make sure that:

- The active or primary target is online.
- The active or primary target is listed at the top of the list.

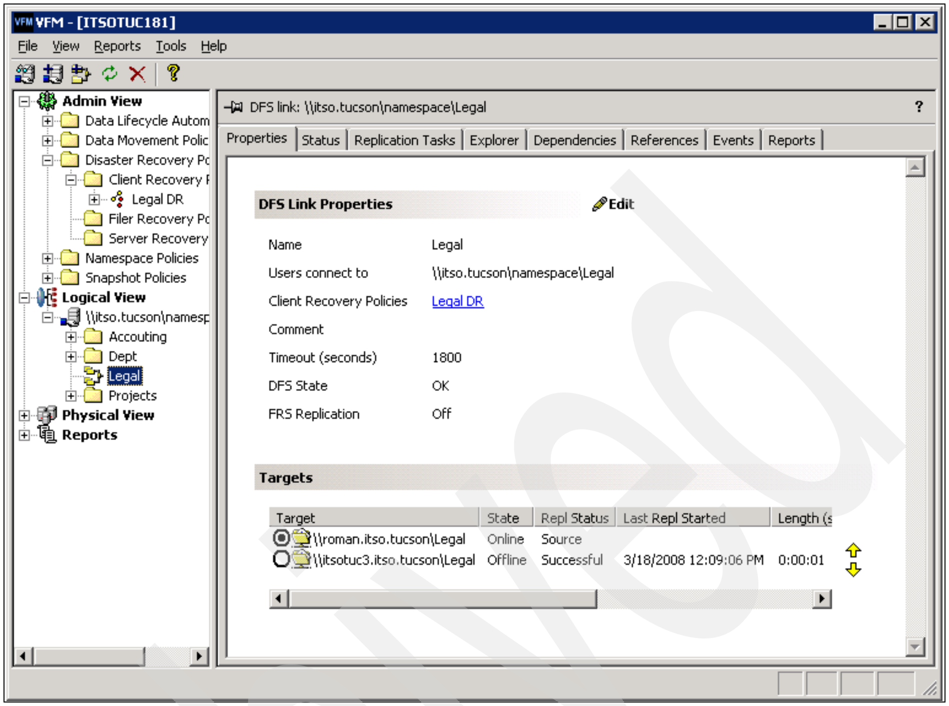

DFS Link Properties

Name	Legal
Users connect to	\\itso.tucson\namespace\Legal
Client Recovery Policies	Legal DR
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

	Target	State	Repl Status	Last Repl Started	Length (s
●	\\roman.itso.tucson\Legal	Online	Source		
○	\\itsotuc3.itso.tucson\Legal	Offline	Successful	3/18/2008 12:09:06 PM	0:00:01

Figure 11-25 DFS link in a DR policy

Note: A single client DR policy can control the failover of one or more DFS links, so more links can be added to the same client DR policy by repeating the process.

To continue with our example, make sure the warning icon on the link in the policy is no longer displayed. At that point, you can simulate a failover.

You can do this by taking the primary share offline. Go to the server hosting the primary share and unshare that share. Because you have set the client DR policy to monitor every minute, it takes about one to two minutes before VFM detects that the primary target is unavailable.

Based on the settings of the policy, VFM fails over the Legal DFS link to the next available destination target. You can check the DFS link properties, indicating that the failover occurred. At this point, the second target is online and is selected (see Figure 11-26).

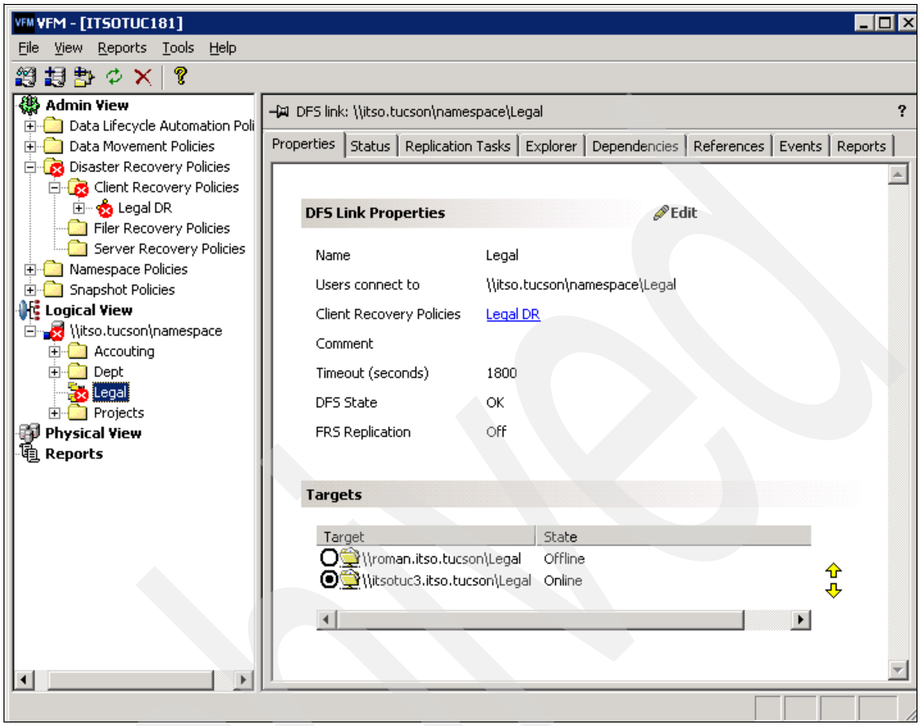

DFS Link Properties

Name	Legal
Users connect to	\\itso.tucson\namespace\Legal
Client Recovery Policies	Legal DR
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
○ \\roman.itso.tucson\Legal	Offline
● \\itsotuc3.itso.tucson\Legal	Online

Figure 11-26 DFS link showing a failover

You can also check the status on a user machine. Open Windows Explorer and access \\domain\ns\legal. The window displays the contents of Legal. At this time, the contents are visible from the DR server on the failover site.

11.3.13 Controlling failback

VFM can fail over users manually or automatically from one server to another, but failback is *always* manual and must be carefully controlled. Follow these steps:

1. Wait for a time when no one is accessing the data on the DR site.
2. The data must be resynchronized with the new primary. You can do that manually, as shown on Figure 11-27, or wait for the next scheduled replication.

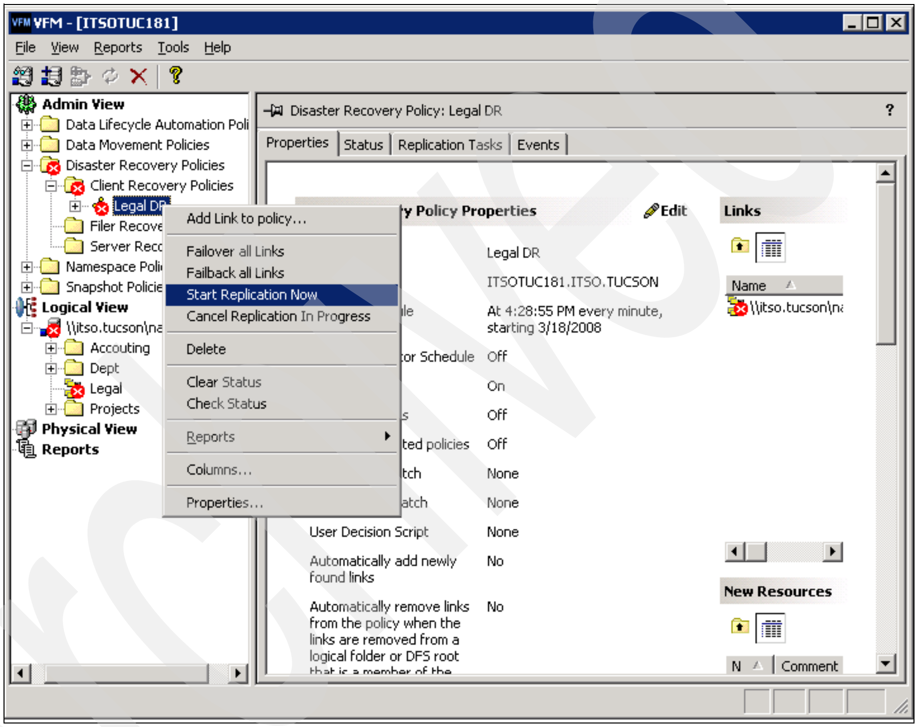

Legal DR

ITSOTUC181.ITSO.TUCSON

le At 4:28:55 PM every minute, starting 3/18/2008

or Schedule Off

On

s Off

ted policies Off

tch None

atch None

User Decision Script None

Automatically add newly found links No

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

Figure 11-27 Start replication manually after disaster

3. You can update the namespace to point back to the newly recovered primary location. To do so, follow these steps:

   a. Right-click the DFS Link and click **Failback** (see Figure 11-28).

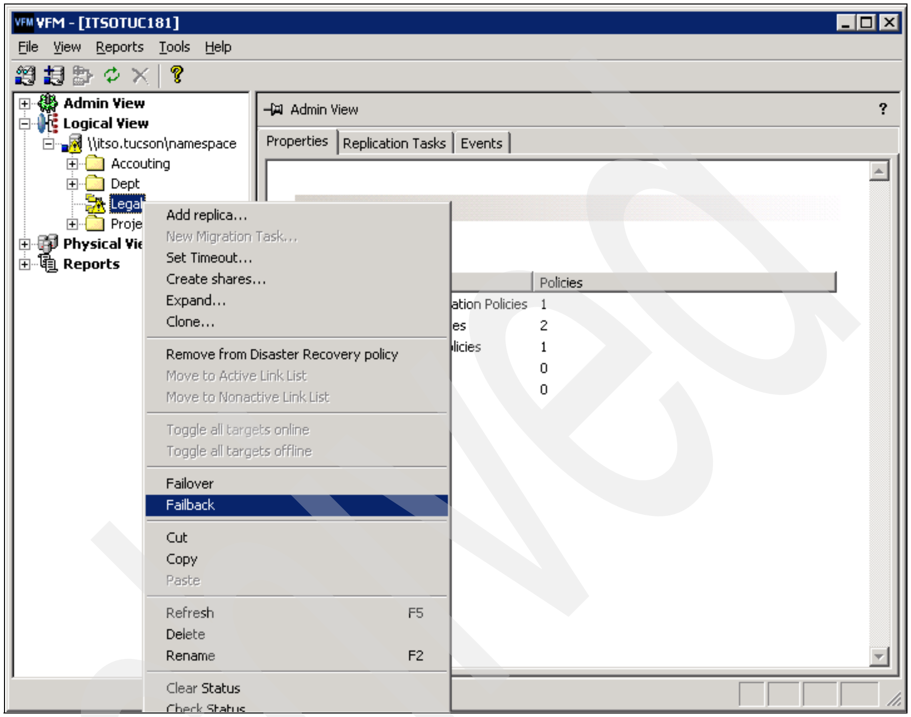

Figure 11-28 Failback to primary target

   b. Click the **Change** button to change only the namespace (see Figure 11-29).

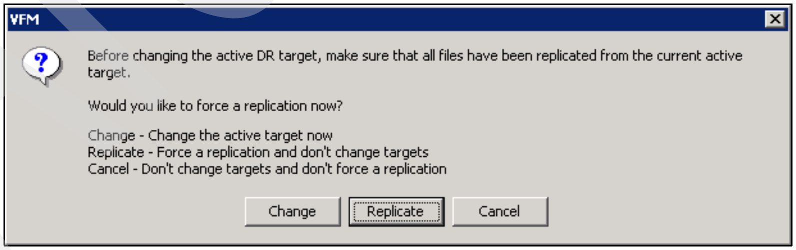
VFM

Before changing the active DR target, make sure that all files have been replicated from the current active target.

Would you like to force a replication now?

Change - Change the active target now
Replicate - Force a replication and don't change targets
Cancel - Don't change targets and don't force a replication

Figure 11-29 Failback dialog

Choosing Replicate at this point starts a replication in the same way as shown earlier.

Be sure to change only the namespace back to the primary target after the replication is finished.

11.4 Setting a filer recovery policy

A filer recovery policy provides a flexible and powerful DR solution for protecting data residing on N series machines. If the data residing on N series is replicated to another N series machine at the DR site using Volume SnapMirror or qtree SnapMirror, the filer recovery policy automates the setup, control, and maintenance of the failover policy.

In this section, we use this scenario as an example: You have a DFS namespace with a set of DFS links that point to the data on a N series called itsotuc3. You have an asynchronous SnapMirror relationship that replicates data between itsotuc3:/vol1 and itsotuc4:/vol1. The SnapMirror relationship was set up external to VFM. (However, the SnapMirror can be set up using VFM.)

If the primary N series storage system fails, the filer recovery policy can perform a manual failover or an automatic failover, based on the settings of the policy. All the shares on the original source machine referenced by the namespace are created and made available on the destination machine automatically. All the namespace links are updated to point to the DR N series volume and the user continues to access the data as before. This provides a high level of DR and BC of the data because users are not disrupted by the loss of the primary N series machine.

Before starting this exercise, make sure the following functions are in place:

- The N series that is part of the filer recovery policy are members of the same Active Directory domain as the VFM server.
- DNS is working correctly for name resolution between the VFM server and the N series that are part of the policy (see 3.2.2, “DNS settings” on page 88).

- The VFM service account has admin rights on the N series storage systems.

  The VFM service account has all necessary rights (see the steps outlined in 3.4, “Creating a VFM service account” on page 93).

- Both N series storage systems are already in the Physical View on the VFM client.

  In this example, we use the N series called itsotuc3 and itsotuc4 (see Figure 11-30).

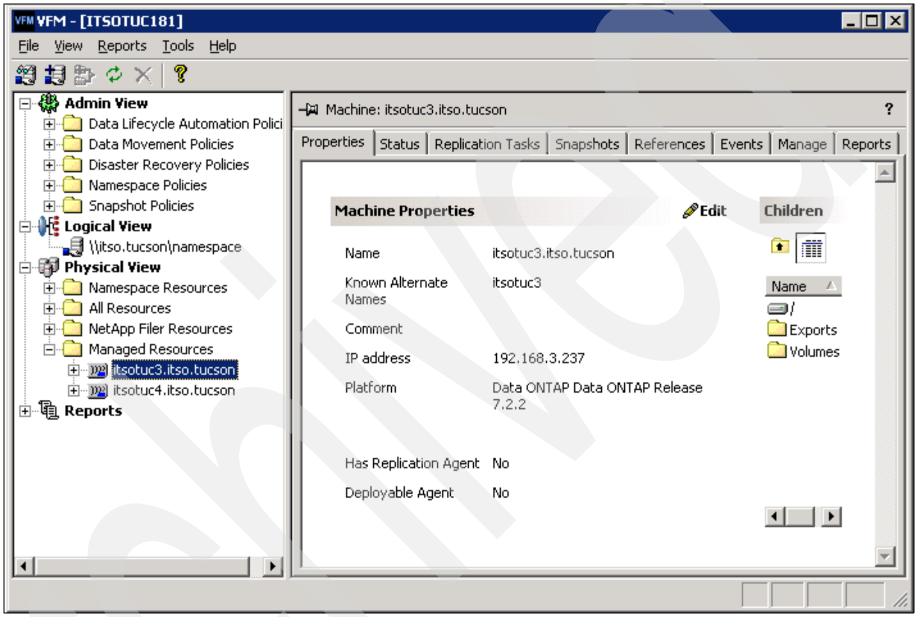

Name	itsotuc3.itso.tucson
Known Alternate Names	itsotuc3
Comment	
IP address	192.168.3.237
Platform	Data ONTAP Data ONTAP Release 7.2.2
Has Replication Agent	No
Deployable Agent	No

Figure 11-30 Managed resources with two N series

- N series storage systems are correctly identified as filers. You can verify this by viewing the icon on the VFM client (see Figure 11-31).

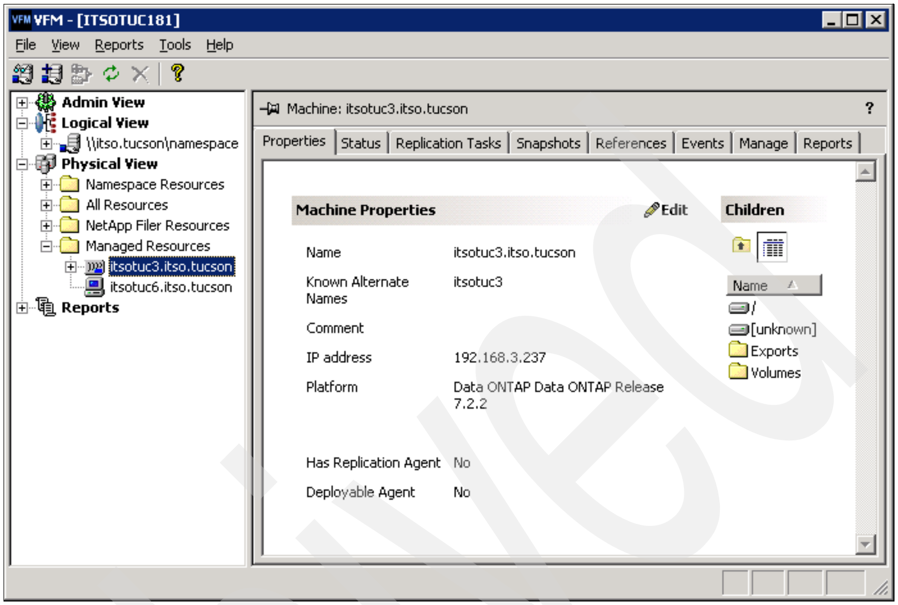

Name	itsotuc3.itso.tucson
Known Alternate Names	itsotuc3
Comment	
IP address	192.168.3.237
Platform	Data ONTAP Data ONTAP Release 7.2.2
Has Replication Agent	No
Deployable Agent	No

Figure 11-31 Physical view showing a filer

- RSH credentials are set up for the N series.
- No red X marks the N series icon in the Physical View.
- SnapMirror licenses are loaded for each N series.

After checking all the items on the preceding list, you run the wizard to set up the filer DR policy. Follow these steps:

1. On the VFM client, go to **Admin View** → **Disaster Recovery Policies** → **Filer Recovery Policy** (see Figure 11-32).

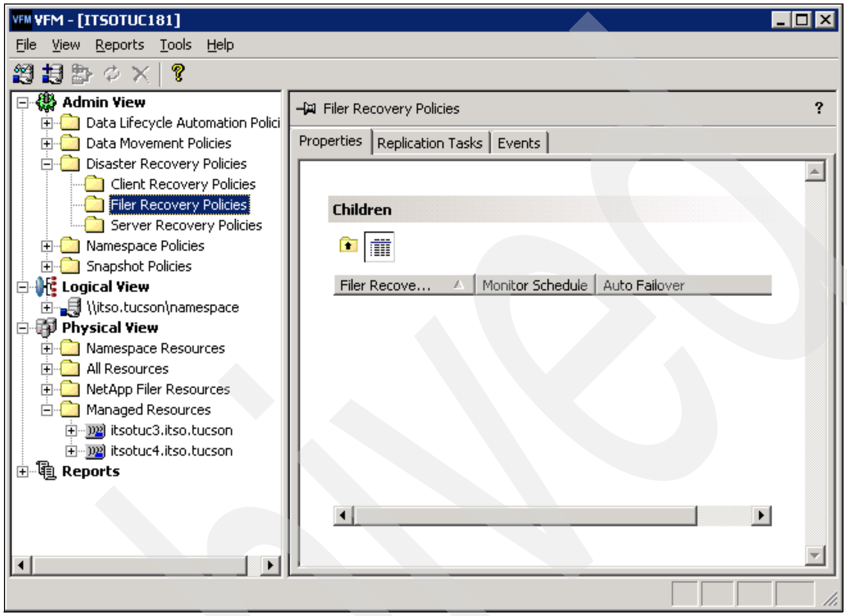

Figure 11-32 Filer Recovery Policies panel

2. Then right-click **Filer Recovery Policies** (see Figure 11-33) and **New Filer Recovery Policy**.

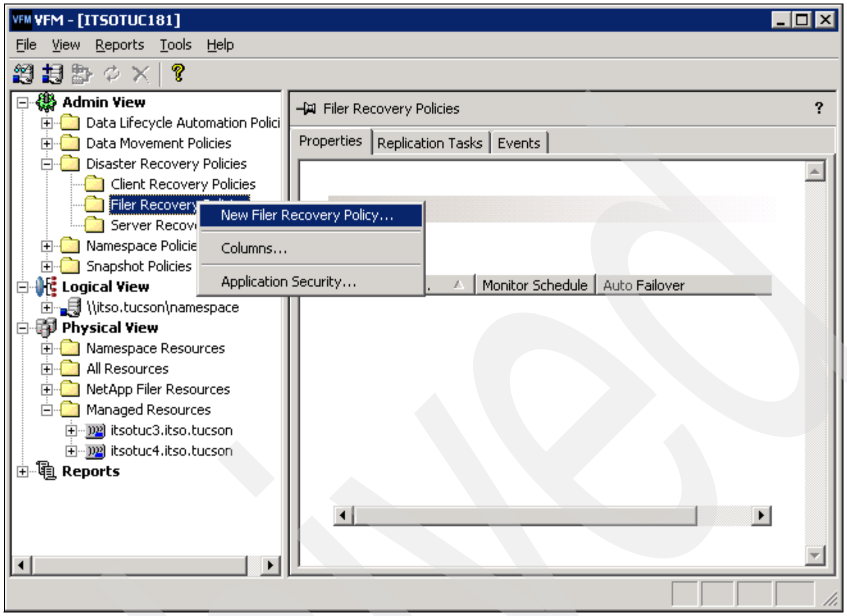

Figure 11-33 New Filer Recovery Policy

3. A welcome window (see Figure 11-34) is displayed. Click **Next**.

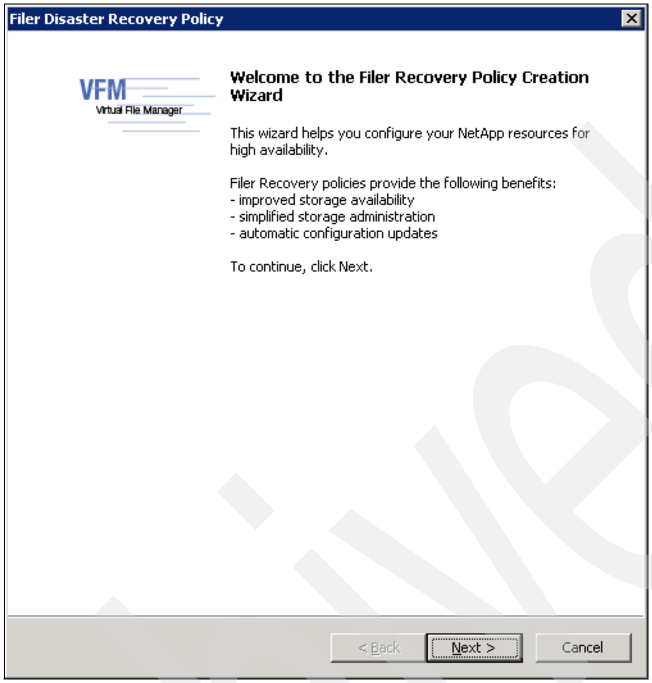
Welcome to the Filer Recovery Policy Creation Wizard

This wizard helps you configure your NetApp resources for high availability.

Filer Recovery policies provide the following benefits:
- improved storage availability
- simplified storage administration
- automatic configuration updates

To continue, click Next.

Figure 11-34 Filer Recovery Policy Creation Wizard

In the sections that follow, we describe each option for setting up a filer recovery policy creator.

11.4.1 Filer Disaster Recovery Options

The Filer Disaster Recovery Options window enables you to configure the filer recovery policy (see Figure 11-35 on page 432). In this window, you specify the name of the policy. You specify the destination filer in the DR location (itsotuc4) and optionally the name of the volume on the destination filer that is a target of the SnapMirror relationship.

For purposes of this example, specify only the filer in the DR site. The following options are available on the Filer Disaster Recovery Options window:

- **Break SnapMirror on failover**

  This check box enables the SnapMirror relationship between the primary and the secondary to be broken when a failover happens. If the check box is not selected, the failover occurs, and the target is still in read-only mode. If the check box is selected, the failover is associated with the target, becoming read-write.

For the purposes of this example, select the **Break SnapMirror on failover** check box.

- **Update associated policies on failover**

  This check box is an advanced feature of the product that enables multiple policies to work together. For the purposes of this example, do not select this check box.

- **Automatically failover**

  This check box enables you to set a manual failover or an automatic failover in case of DR. In production environments, you do not want to check this box because a false failover can occur if the primary server is briefly unavailable.

  For this example, you check the box to initiate an automatic failover (see Figure 11-35).

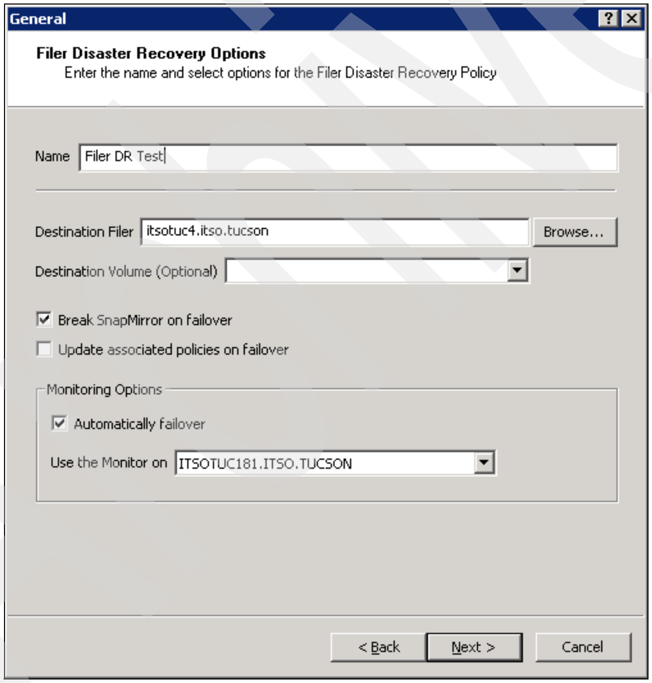

Filer Disaster Recovery Options

Enter the name and select options for the Filer Disaster Recovery Policy

Name: Filer DR Test

Destination Filer: itsotuc4.itso.tucson

Destination Volume (Optional):

☑ Break SnapMirror on failover

□ Update associated policies on failover

Monitoring Options

☑ Automatically failover

Use the Monitor on: ITSOTUC181.ITSO.TUCSON

Figure 11-35 Filer recovery policy options

11.4.2 Filer Disaster Recovery Roots

In the Filer Disaster Recovery Roots window, you can select the namespaces to monitor. VFM monitors these namespaces to determine whether the links that are part of them are available. For each link, it ensures the existence of a DR target to the secondary. In addition, if a new share is created on the primary volume and this share is targeted by a namespace link, this share is created on the DR volume.

For this example, specify the DFS roots to monitor and click **Next** (see Figure 11-36).

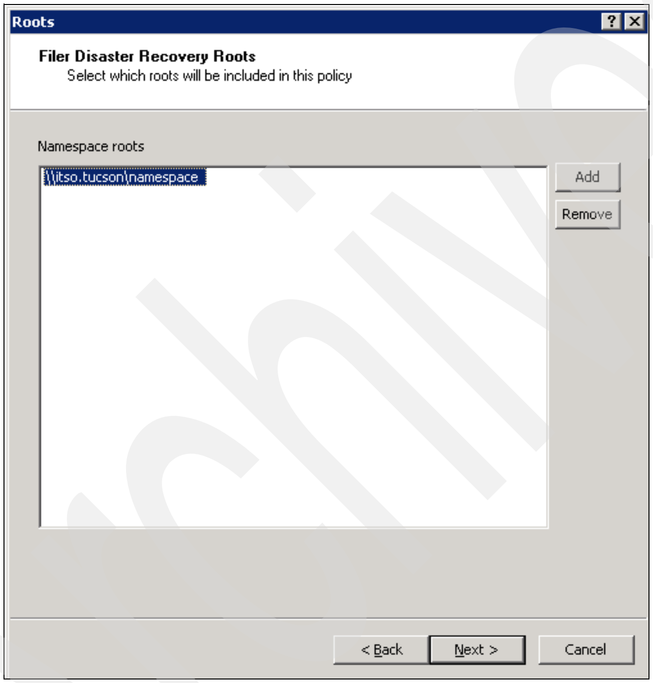

Roots

Filer Disaster Recovery Roots

Select which roots will be included in this policy

Namespace roots

\\itso.tucson\namespace

Figure 11-36 Namespaces to monitor

11.4.3 Specifying volumes and qtrees to be made highly available

Follow these steps to specify the volumes and qtrees to be made highly available (see Figure 11-37):

1. Click **Add**.

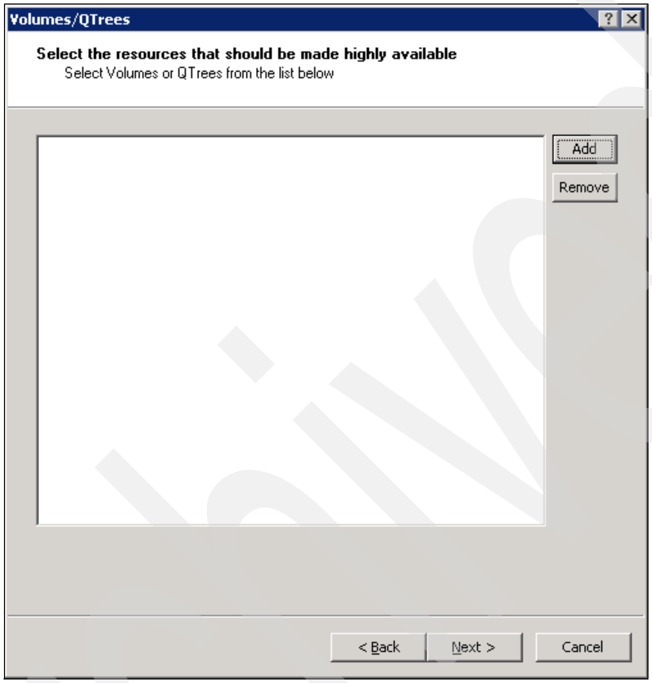
Volumes/QTrees
?
Select the resources that should be made highly available
Select Volumes or QTrees from the list below
Add
Remove
< Back
Next >
Cancel

Figure 11-37 Volumes or qtree

2. Choose the N series and a volume on it (see Figure 11-38), and click **OK**.

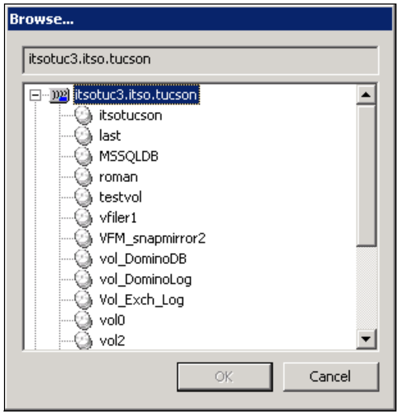

Figure 11-38   Browsing volumes

3. The chosen volume is displayed in the list (see Figure 11-39). Click **Next**.

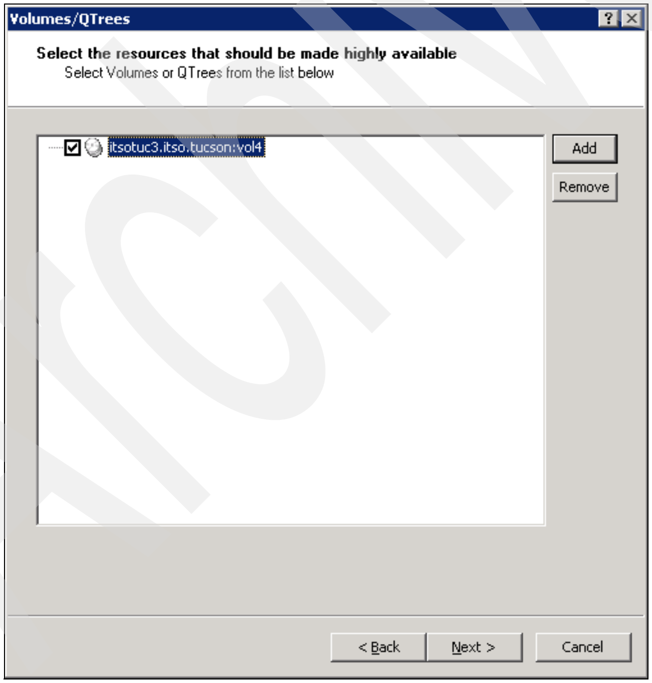
Volumes/QTrees
Select the resources that should be made highly available
Select Volumes or QTrees from the list below
☑ itsotuc3.itso.tucson:vol4
Add
Remove
< Back
Next >
Cancel

Figure 11-39   Dialog showing chosen volume

A drop-down list box is displayed. It contains the existing SnapMirror relations for the chosen volume or the option to create one (see Figure 11-40).

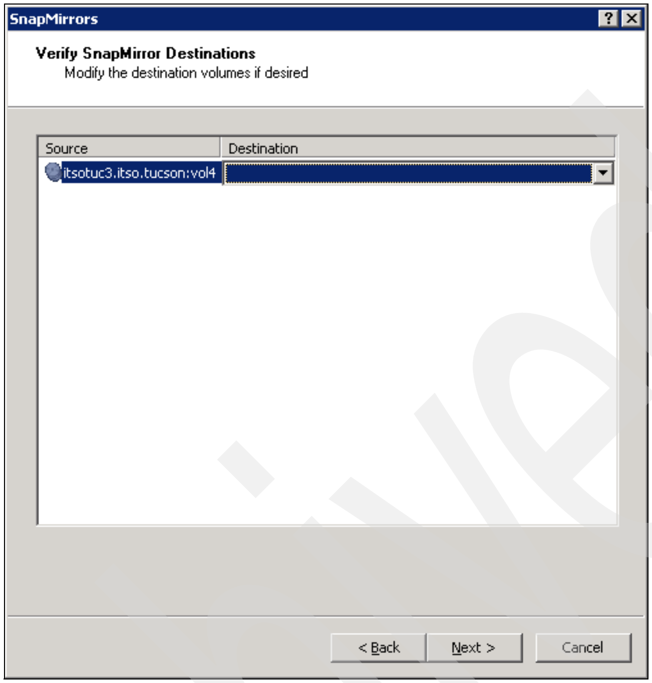
SnapMirrors
Verify SnapMirror Destinations
Modify the destination volumes if desired
Source
Destination
itsotuc3.itso.tucson:vol4
< Back
Next >
Cancel

Figure 11-40 SnapMirror destinations

11.4.4 Adding a SnapMirror

In this section, we demonstrate how to create a new SnapMirror for the chosen volume. Follow these steps:

1. Click the **Destination** list and then click **Add new SnapMirror** (see Figure 11-41).

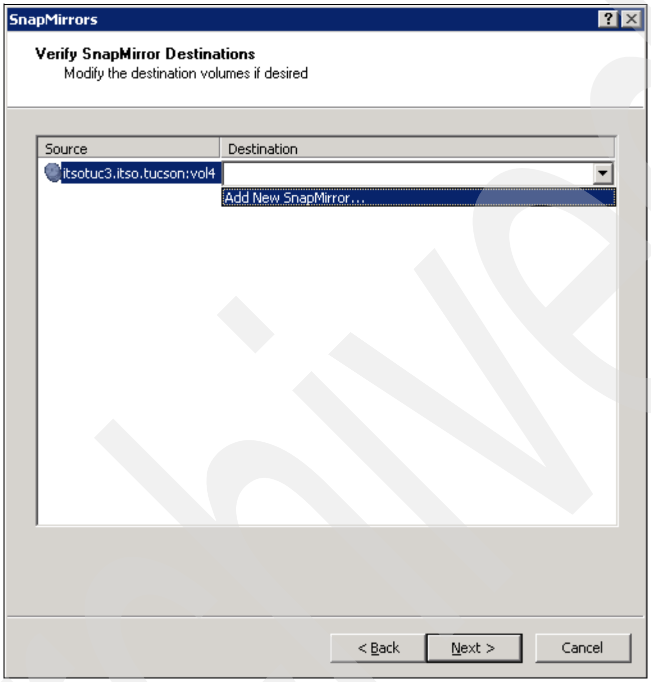

SnapMirrors

Verify SnapMirror Destinations

Modify the destination volumes if desired

Source	Destination
itsotuc3.itso.tucson:vol4	Add New SnapMirror...

Figure 11-41 Adding a new SnapMirror

2. Choose the destination volume (see Figure 11-42).

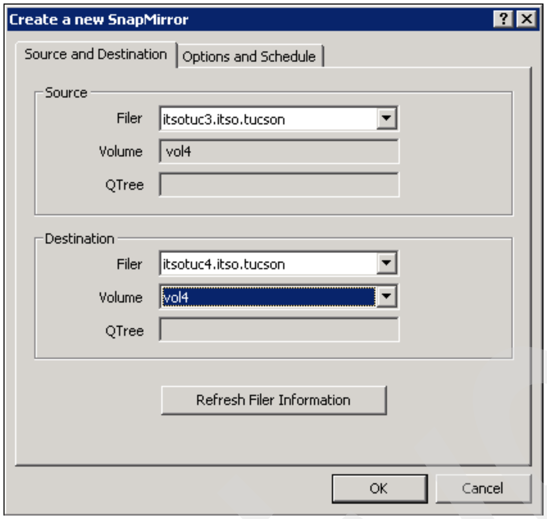

Source

Filer: itsotuc3.itso.tucson

Volume: vol4

QTree:

Destination

Filer: itsotuc4.itso.tucson

Volume: vol4

QTree:

Figure 11-42 Creating a SnapMirror

3. At this point, the SnapMirror is set between the two volumes on the two N series (see Figure 11-43). Click **Next**.

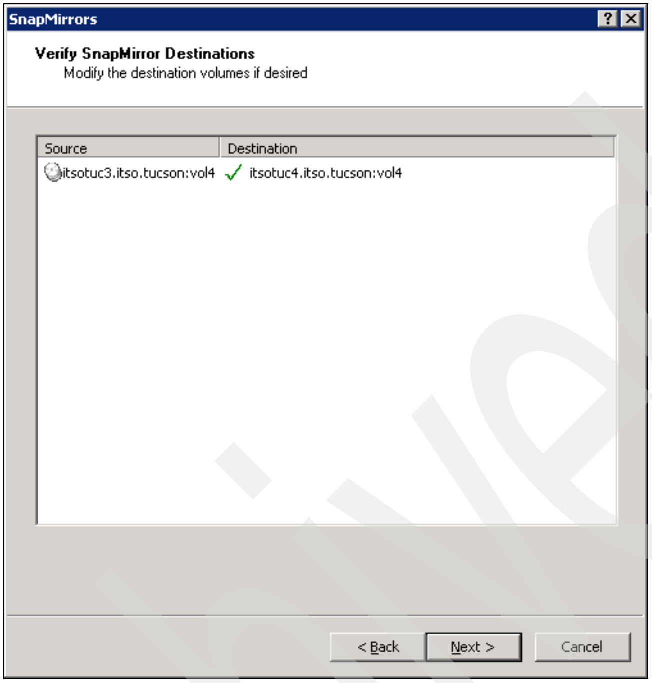

Verify SnapMirror Destinations

Modify the destination volumes if desired

Source	Destination
itsotuc3.itso.tucson:vol4	✓ itsotuc4.itso.tucson:vol4

Figure 11-43 SnapMirror destination set

11.4.5 Find New Resources option

VFM can be set to create new objects on the destination bases of the objects on the source. For every object created on the source, VFM can create it on the destination (see Figure 11-44).

- **Find new links**

  This option directs VFM to create the new destination share for the source share when a namespace link points to the source share.

- **Find new qtrees**

  This option directs VFM to search the volume to find any newly created qtrees. If a new qtree is found, it creates the appropriate qtree SnapMirror relation for that source.

For the purposes of our example, select both options.

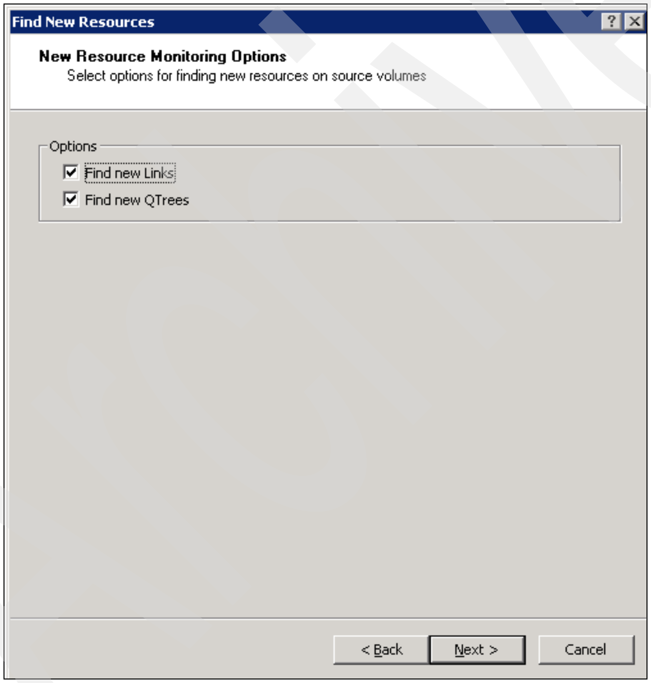

Find New Resources

New Resource Monitoring Options

Select options for finding new resources on source volumes

Options

☑ Find new Links

☑ Find new QTrees

Figure 11-44 Find New Resources window

11.4.6 Resource monitoring schedule

In the New Resource Monitoring Schedule window, you can specify how often the source volume is scanned to look for new resources (such as shares and qtrees). You can set this option to scan once every hour (see Figure 11-45).

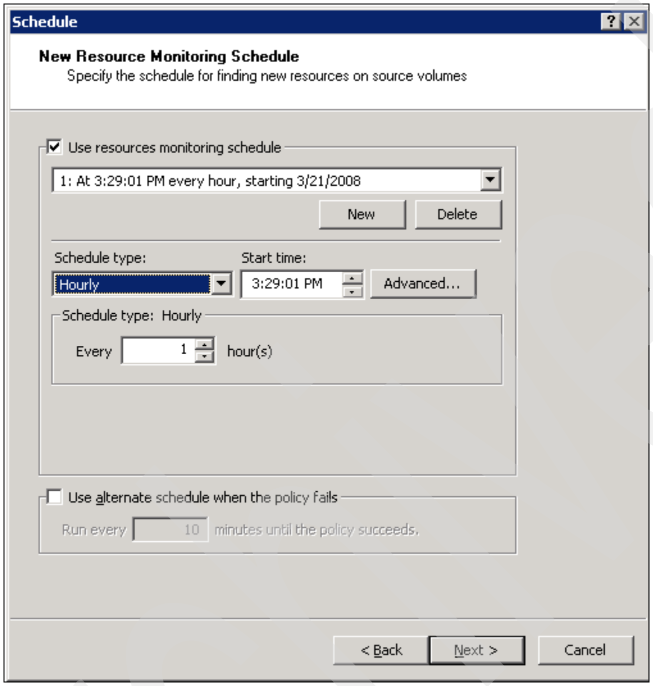

Schedule

New Resource Monitoring Schedule

Specify the schedule for finding new resources on source volumes

☑ Use resources monitoring schedule

1: At 3:29:01 PM every hour, starting 3/21/2008

Schedule type: Hourly

Start time: 3:29:01 PM

Schedule type: Hourly

Every 1 hour(s)

☐ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

Figure 11-45 New Resource Monitoring Schedule window

11.4.7 Filer Disaster Recovery Schedule

In the Filer Disaster Recovery Schedule window, you can specify how often the source is scanned for availability (see Figure 11-46 on page 442). The target monitoring schedule must be closely controlled because it determines how often the VFM server monitors the primary server for availability.

If you set this option to a low value, you might be unnecessarily generating a lot of network traffic as the VFM server monitors the primary server for availability. If you set this value too high, you might delay detecting the loss of the primary server. It is best to consult your users, your recovery time objectives (RTO), and recovery point objectives (RPO) before deciding on an appropriate value. In most instances, a target-monitoring interval value of 5 or 10 minutes is adequate for any policy.

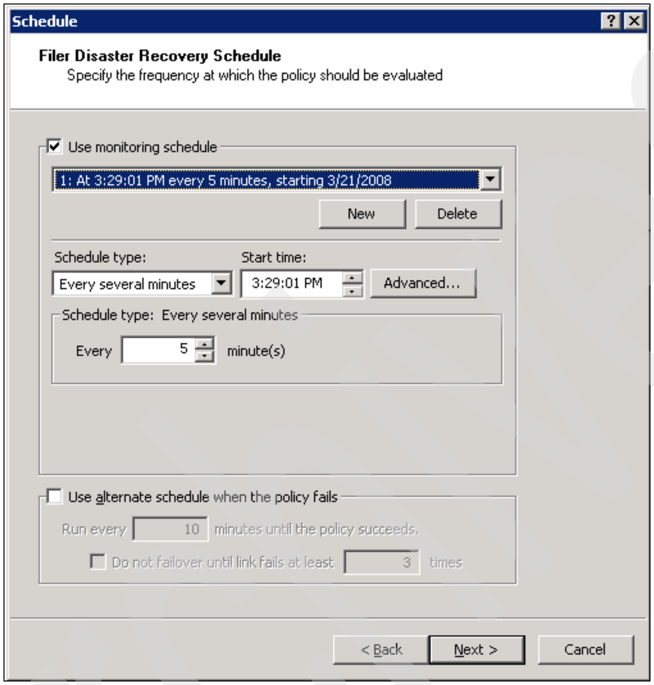

Schedule

Filer Disaster Recovery Schedule
Specify the frequency at which the policy should be evaluated

☑ Use monitoring schedule

1: At 3:29:01 PM every 5 minutes, starting 3/21/2008

Schedule type: Every several minutes

Start time: 3:29:01 PM

Schedule type: Every several minutes

Every 5 minute(s)

□ Use alternate schedule when the policy fails

Run every 10 minutes until the policy succeeds.

□ Do not failover until link fails at least 3 times

Figure 11-46 Availability monitoring schedule

11.4.8 DR Script Settings

In the DR Script Settings window, you can set scripts to run before and after a failover. You can even change the decision to failover (see Figure 11-47). For the purposes of our example, click **Next**.

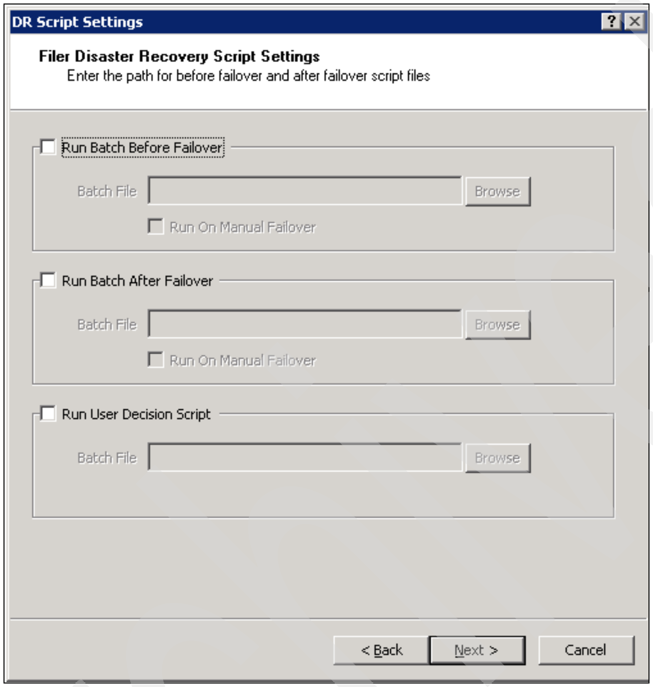

DR Script Settings

Filer Disaster Recovery Script Settings

Enter the path for before failover and after failover script files

□ Run Batch Before Failover

Batch File [ ] Browse

□ Run On Manual Failover

□ Run Batch After Failover

Batch File [ ] Browse

□ Run On Manual Failover

□ Run User Decision Script

Batch File [ ] Browse

< Back | Next > | Cancel

Figure 11-47 Script settings

At this point, you have reached the end of the Filer Disaster Recovery Wizard (see Figure 11-48). Click **Finish** to actually create the policy.

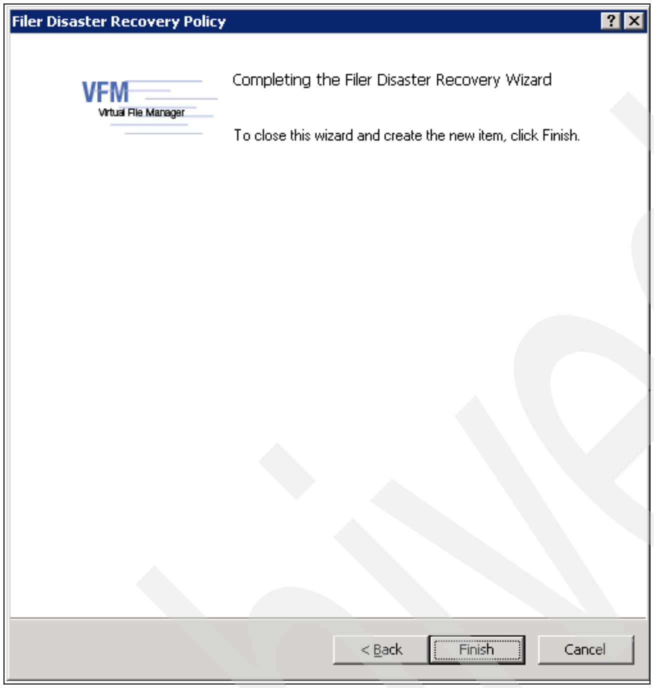

Completing the Filer Disaster Recovery Wizard

To close this wizard and create the new item, click Finish.

Figure 11-48 Filer DR wizard: finish

11.4.9 Testing the example

Follow these steps to test our example:

1. Share the vol4 DFS link in both N series.
2. Create a DFS link pointing to the primary target only. (See Chapter 5, “DFS namespace” on page 117.)

   Figure 11-49 displays the DFS link vol4 pointing to itsotuc3.

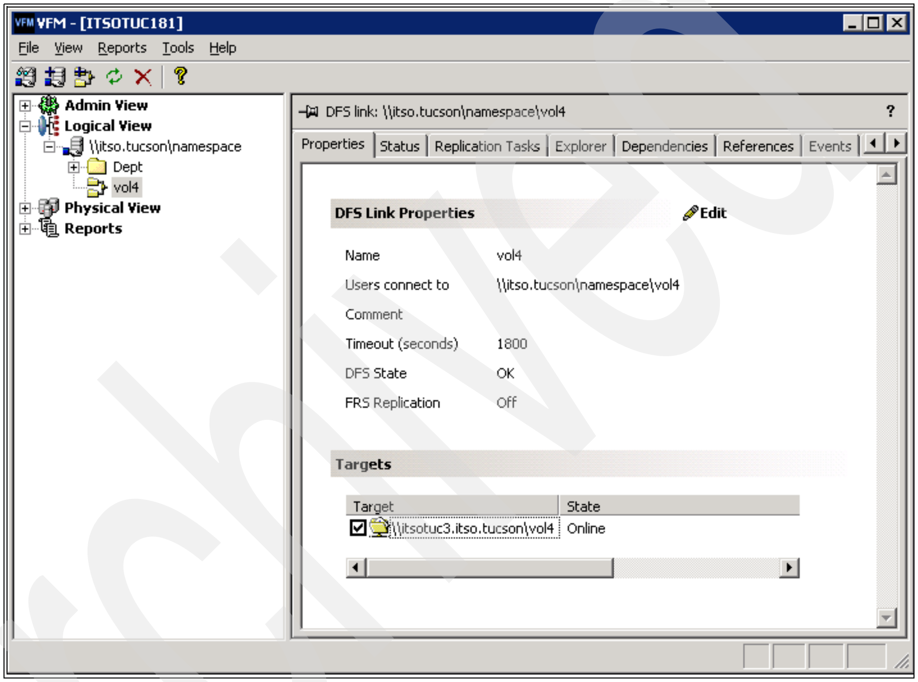

DFS link: \\itso.tucson\namespace\vol4

DFS Link Properties

Name	vol4
Users connect to	\\itso.tucson\namespace\vol4
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
☑ \\itsotuc3.itso.tucson\vol4	Online

Figure 11-49 DFS link targeting vol4 on itsotuc3

Wait a few minutes, and notice that VFM has automatically found the share on the secondary N series, adding it as a target on the DFS link (see Figure 11-50).

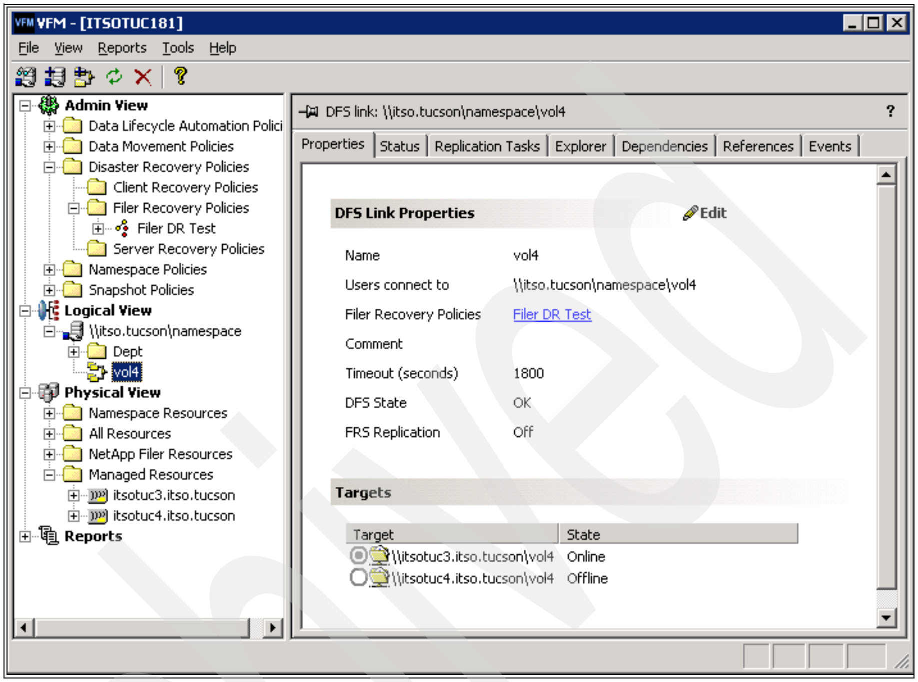

DFS Link Properties

Name	vol4
Users connect to	\\itso.tucson\namespace\vol4
Filer Recovery Policies	Filer DR Test
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
● \\itsotuc3.itso.tucson\vol4	Online
○ \\itsotuc4.itso.tucson\vol4	Offline

Figure 11-50 DFS link showing both targets for DR

To check that VFM can fail over, you can shut down the primary N series. Or if you cannot shut it down, you can manually failover by following these steps:

1. On the VFM client, go to **Admin View** → **Disaster Recovery Policies** → **Filer Recovery Policy**.
2. Then right-click the **Filer DR Test** policy.

3. Click **Failover All** (see Figure 11-51).

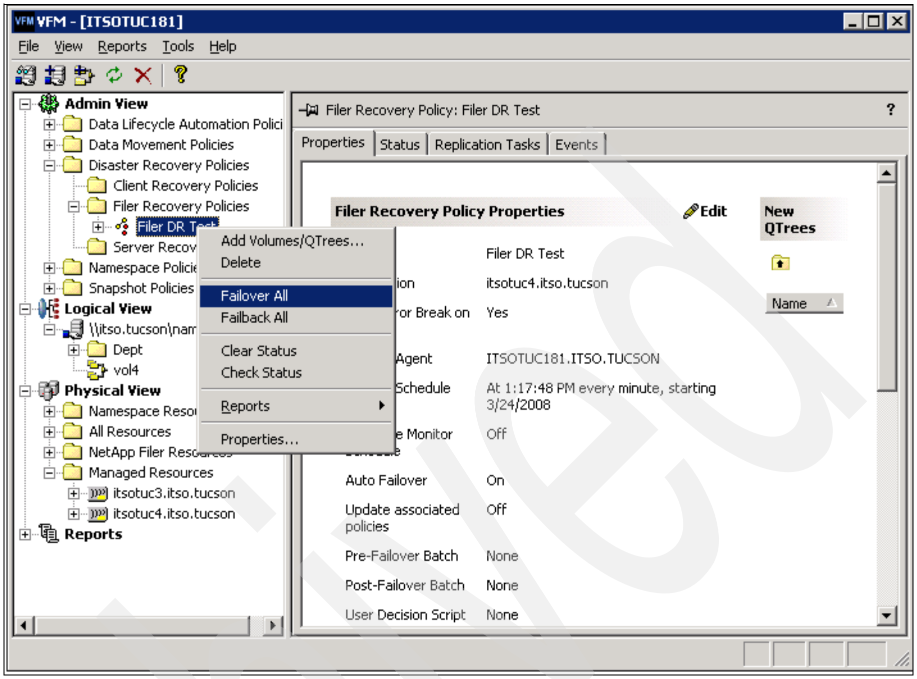

	Filer DR Test
ion	itsotuc4.itso.tucson
or Break on	Yes
Agent	ITSOTUC181.ITSO.TUCSON
Schedule	At 1:17:48 PM every minute, starting 3/24/2008
e Monitor	Off
Auto Failover	On
Update associated policies	Off
Pre-Failover Batch	None
Post-Failover Batch	None
User Decision Script	None

Figure 11-51 Failover All menu

Whether you shut down the primary N series or you perform a manual failover, VFM switches the online target on the DFS link (see Figure 11-52).

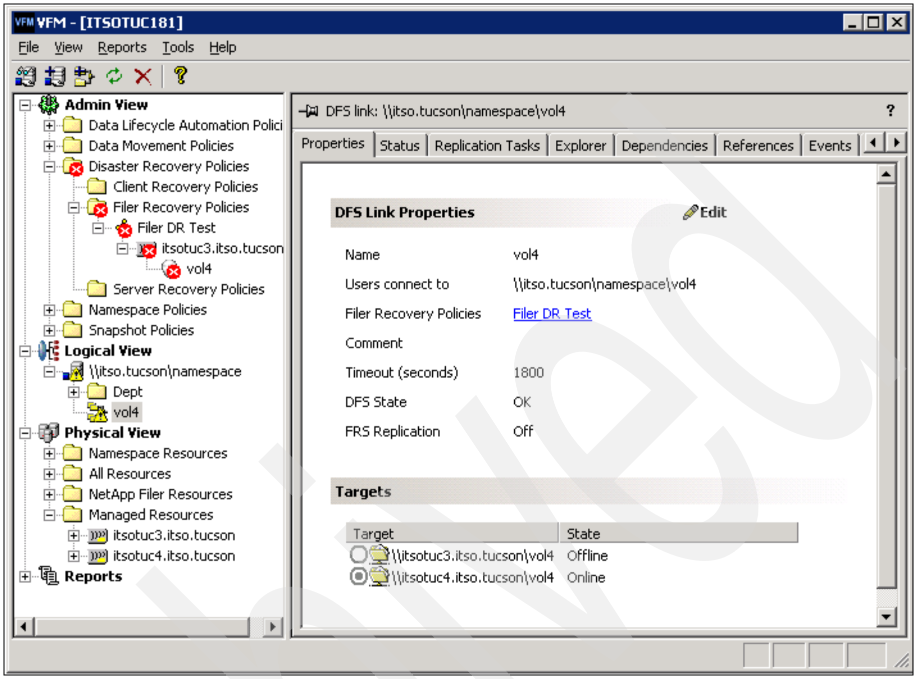

DFS Link Properties

Name	vol4
Users connect to	\\itso.tucson\namespace\vol4
Filer Recovery Policies	Filer DR Test
Comment	
Timeout (seconds)	1800
DFS State	OK
FRS Replication	Off

Targets

Target	State
○ \\itsotuc3.itso.tucson\vol4	Offline
● \\itsotuc4.itso.tucson\vol4	Online

Figure 11-52 DFS link after failover

Then VFM breaks the SnapMirror (see Figure 11-53).

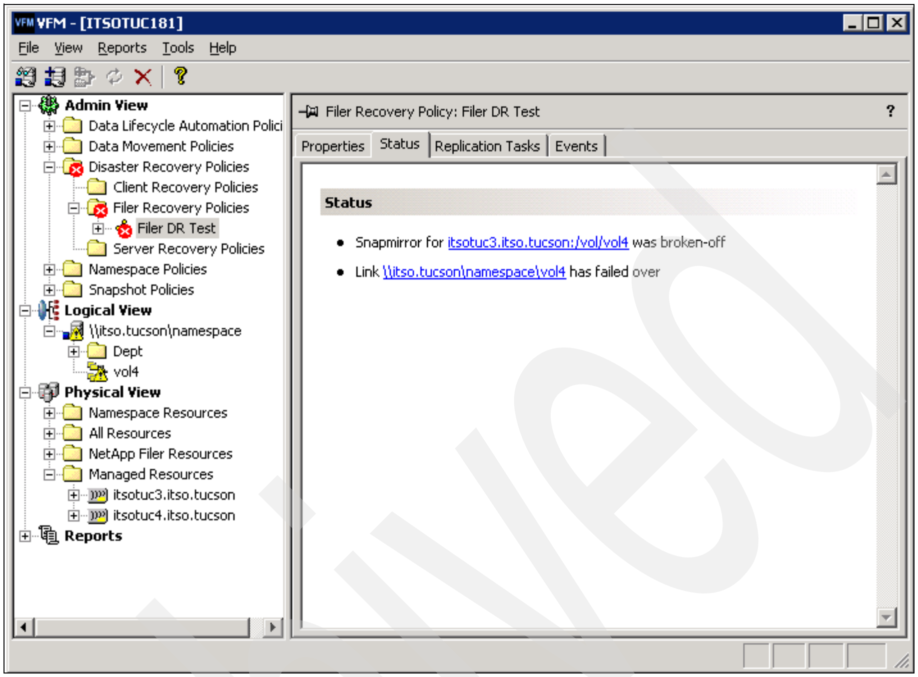

- Snapmirror for itsotuc3.itso.tucson:/vol/vol4 was broken-off
- Link \\itso.tucson\namespace\vol4 has failed over

Figure 11-53 DR policy status after failover

12

Adding IBM N series storage to VFM

This chapter briefly describes how to incorporate the IBM N series with VFM to obtain information and manage several of VFM features.

12.1 Integrating IBM N series with VFM

No real differences exist between *adding* a regular server and an N series. The differences are apparent only after they have been added. In the sections that follow, we show the steps to add a N series to VFM.

12.1.1 Creating a Managed Resources folder

VFM automatically updates the Namespace Resources and All Resources folders on Physical View (see Figure 12-1), and it enables you to create as many folders as required to organize your servers, making it easier to manage them.

To manually add a machine to VFM, you must create a folder first and add the machines to your folder. In this section, we create a Managed Resources folder. Follow these steps:

1. Go to the VFM client (in **Start** → **Programs** → **IBM** → **Virtual File Manager** → **VFM Client**).
2. Right-click **Physical View** and then click **Add Folder**, as shown in Figure 12-1.

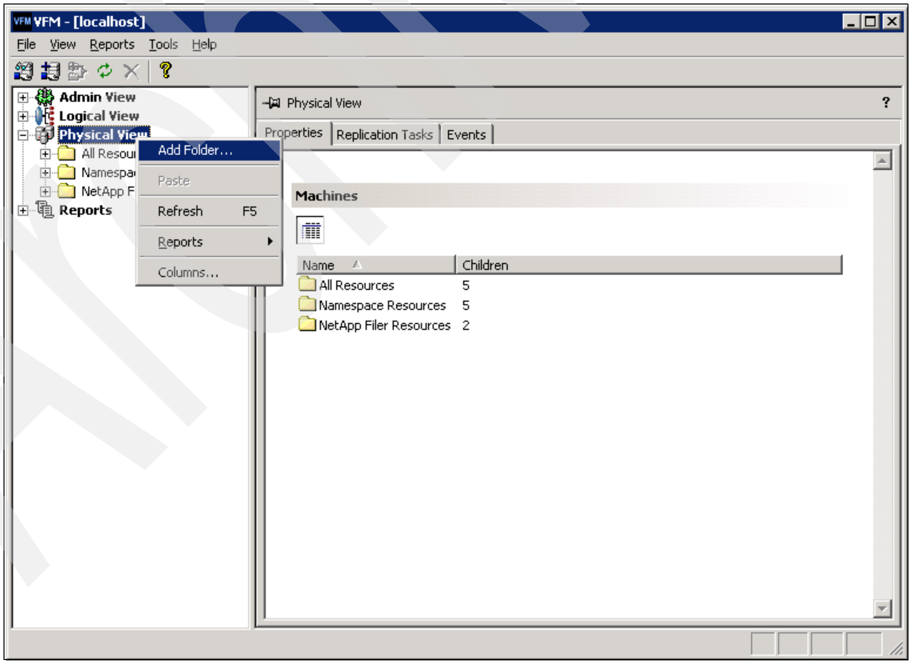

Name	Children
All Resources	5
Namespace Resources	5
NetApp Filer Resources	2

Figure 12-1 Physical View on VFM client

3. The Add Folder window is displayed, and you can set the folder name. Use the name Managed Resources, as shown in Figure 12-2.

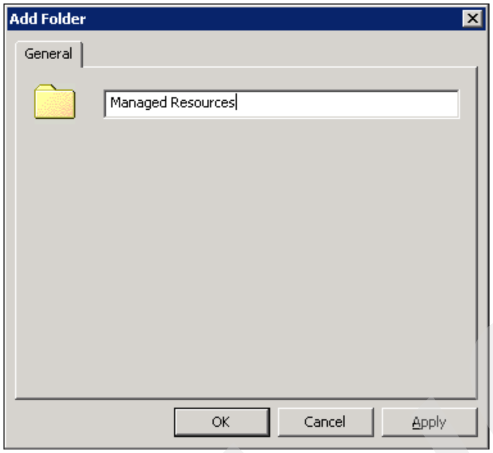

Managed Resources

Figure 12-2 Adding Managed Resources folder

The Managed Resources folder is available as shown in Figure 12-3.

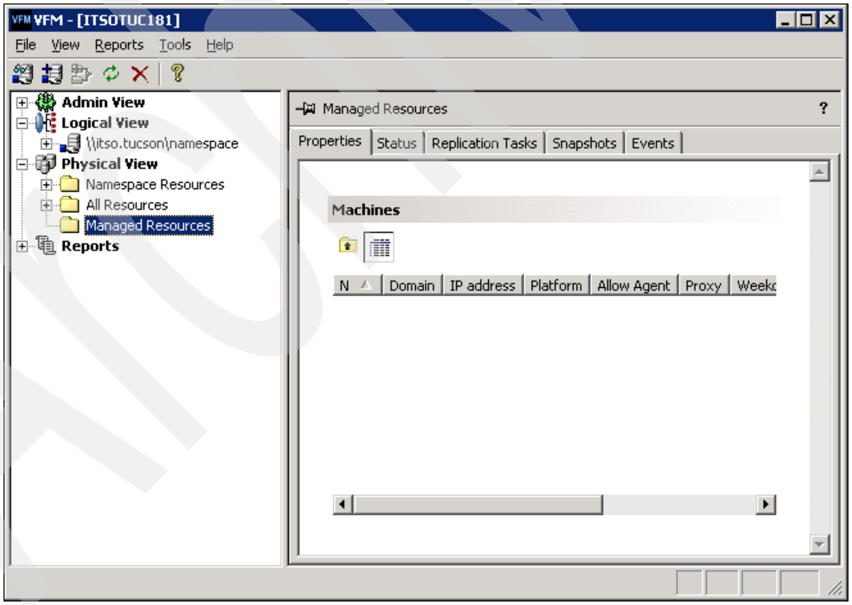

Figure 12-3 Physical View with Managed Resources folder

12.1.2 Adding the IBM N series

After creating the Managed Resources folder, we can add machines to it.

Important: You must be sure that DNS name resolution is working correctly between the VFM server and the N series (see 3.2.2, “DNS settings” on page 88.

In this section, we add the itsotuc3.itso.tucson N series storage system. Follow these steps:

1. Right-click the **Managed Resources** folder (see Figure 12-4).

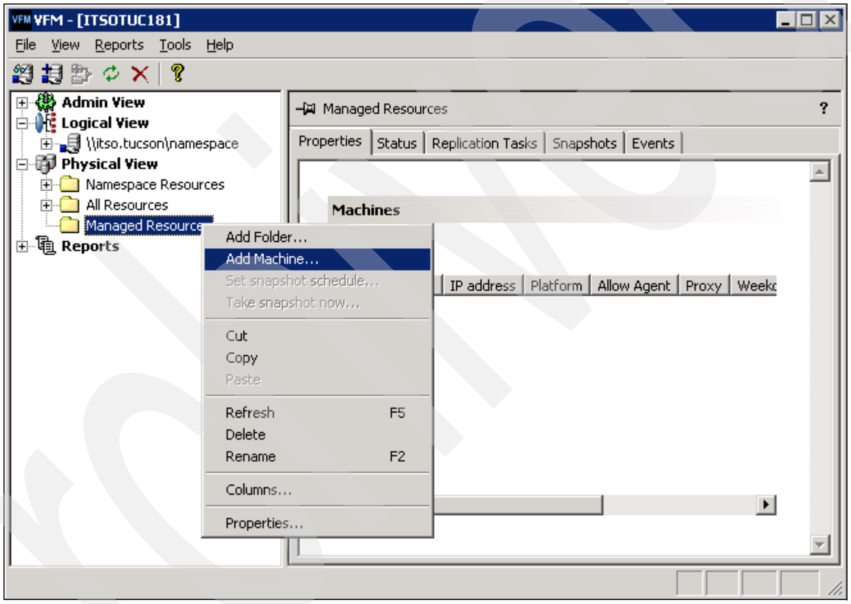

Figure 12-4 Add Machine menu

2. Enter the host name of the N series you wish to add (see Figure 12-5).

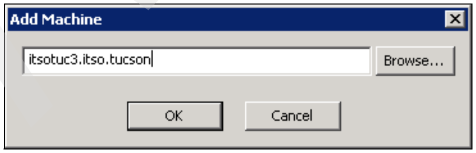

itsotuc3.itso.tucson

Figure 12-5 Adding itsotuc3

The N series is listed in the Managed Resources folder (see Figure 12-6). The red X adjacent to the N series indicates VFM does not have the information required to access it using RSH.

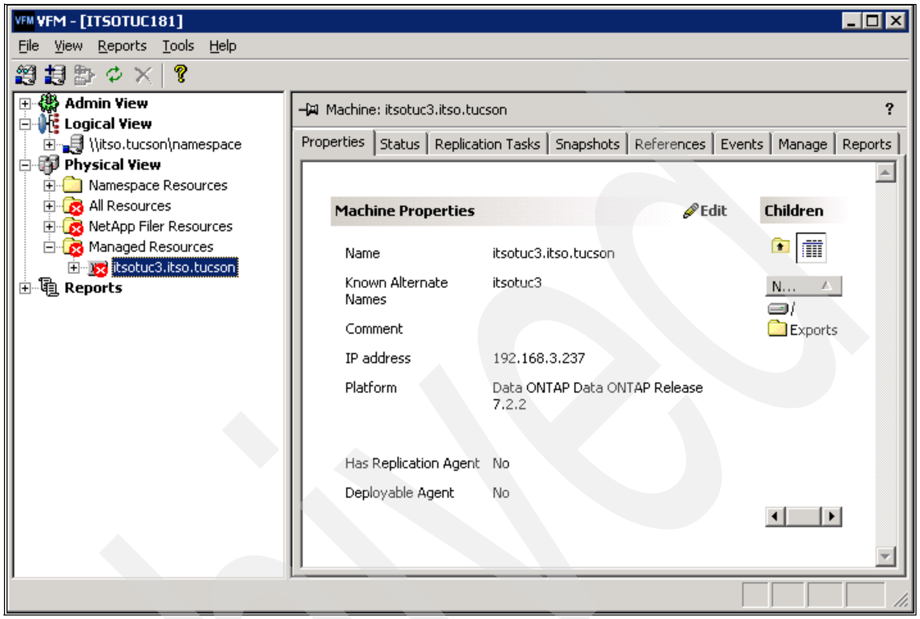

Name	itsotuc3.itso.tucson
Known Alternate Names	itsotuc3
Comment	
IP address	192.168.3.237
Platform	Data ONTAP Data ONTAP Release 7.2.2
Has Replication Agent	No
Deployable Agent	No

Figure 12-6 Managed Resources with itsotuc3

Note: VFM always displays this error when you first add an N series. We demonstrate how to fix it in 12.3, “Setting RSH information about VFM” on page 459.

12.2 Setting VFM to recognize a filer

At times VFM identifies N series storage systems as regular Microsoft Windows servers and does not enable you to perform any N series-specific operations. You can fix this situation by following these steps:

1. Right-click the machine name (see Figure 12-7).

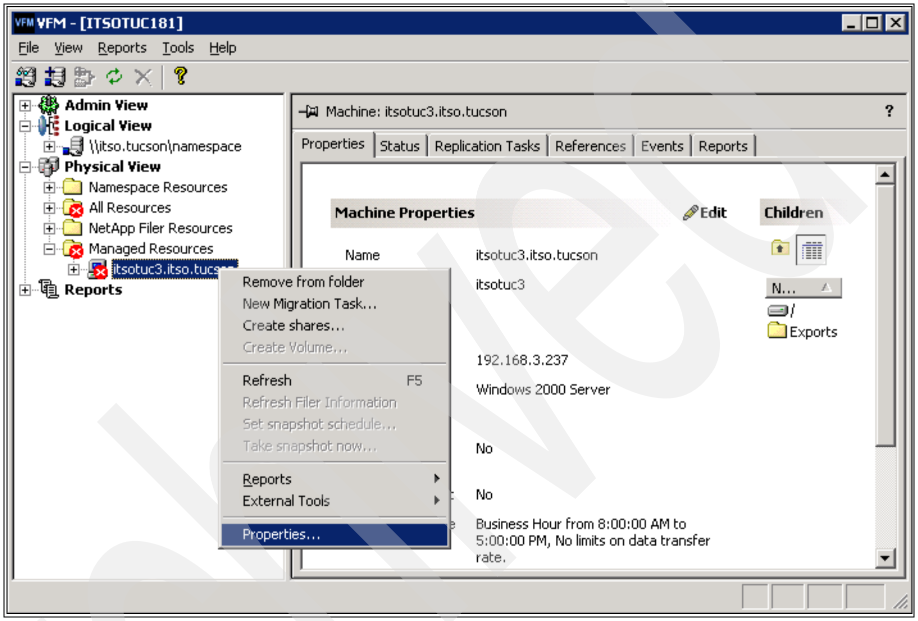

Name itsotuc3.itso.tucson

itsotuc3

192.168.3.237

Windows 2000 Server

No

No

Business Hour from 8:00:00 AM to 5:00:00 PM, No limits on data transfer rate.

Figure 12-7 Machine right-click menu

The Machine Identity field defaults to Auto-detect machine identity as shown in Figure 12-8.

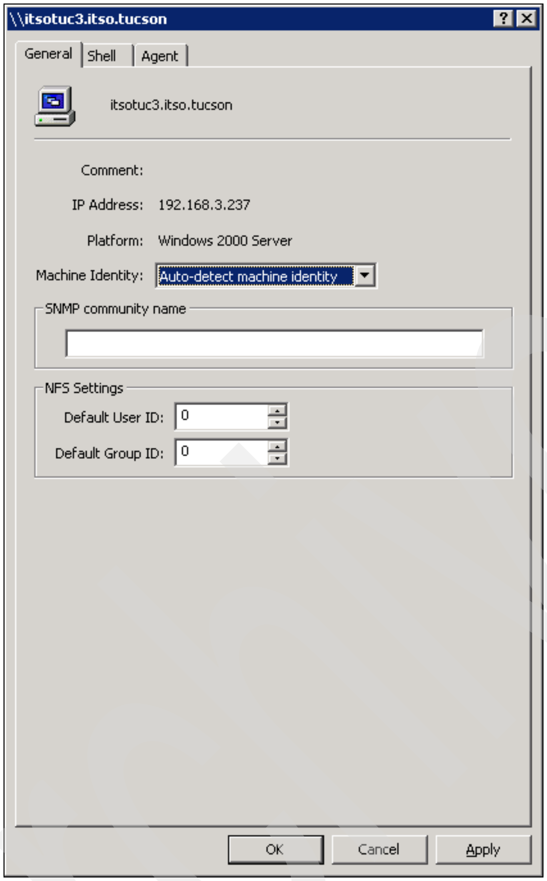

\\itsotuc3.itso.tucson

itsotuc3.itso.tucson

Comment:

IP Address: 192.168.3.237

Platform: Windows 2000 Server

Machine Identity: Auto-detect machine identity

SNMP community name

NFS Settings

Default User ID: 0

Default Group ID: 0

Figure 12-8 Properties showing Auto-detect

2. Change the option in the Machine Identity list box to NetApp® filer (see Figure 12-9).

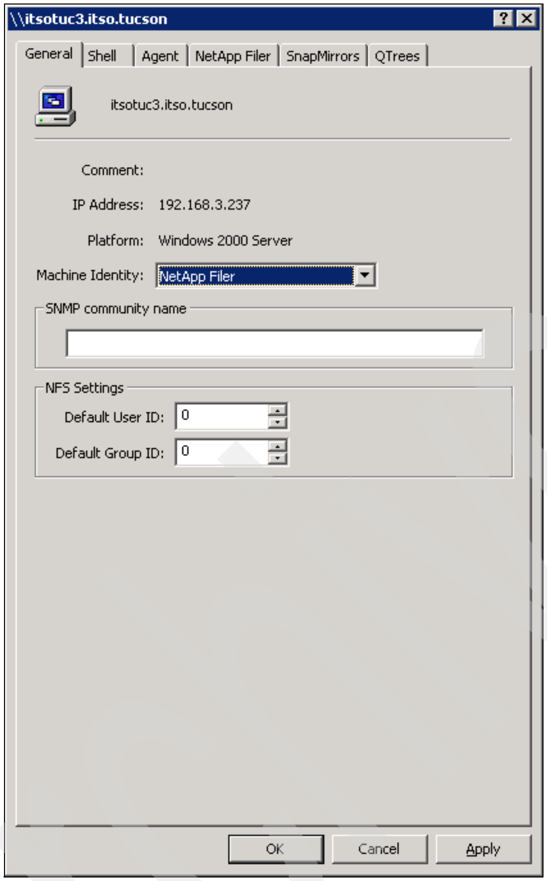

itsotuc3.itso.tucson

Comment:

IP Address: 192.168.3.237

Platform: Windows 2000 Server

Machine Identity: NetApp Filer

SNMP community name

NFS Settings

Default User ID: 0

Default Group ID: 0

Figure 12-9 Properties with NetApp filer set

At this point, VFM recognizes that this machine is a filer, and the corresponding options are displayed (see Figure 12-10).

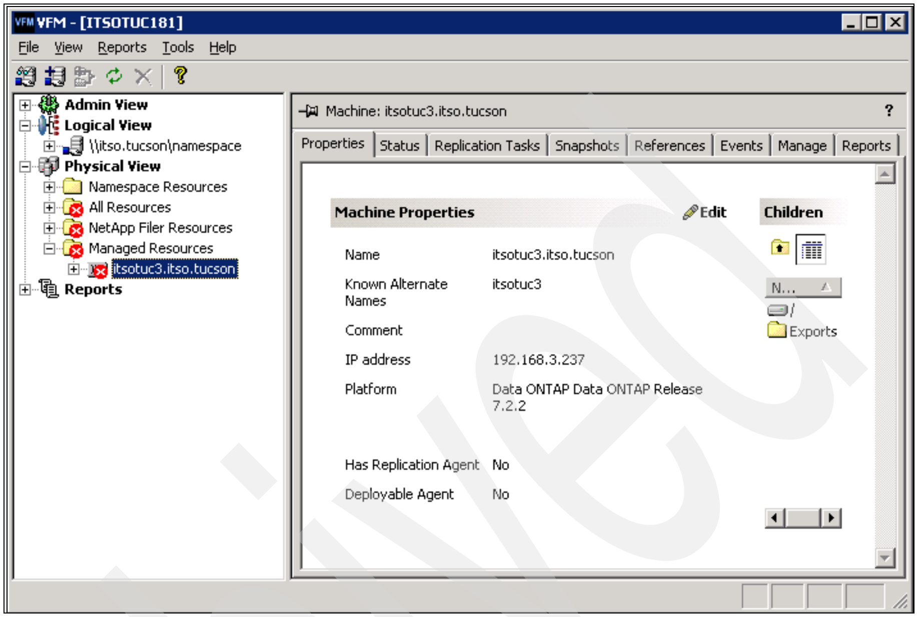

Name	itsotuc3.itso.tucson
Known Alternate Names	itsotuc3
Comment	
IP address	192.168.3.237
Platform	Data ONTAP Data ONTAP Release 7.2.2
Has Replication Agent	No
Deployable Agent	No

Figure 12-10 Filer correctly identified

12.3 Setting RSH information about VFM

VFM needs RSH access to the N series storage systems to provide detailed information and to control it. You must manually enter this information for each filer you want to control with VFM.

Note: This procedure is not required if you use only the filer's CIFS shares, without having to control anything on the filer.

Follow these steps to set the RSH information to your N series:

1. Right-click the machine name (see Figure 12-11).

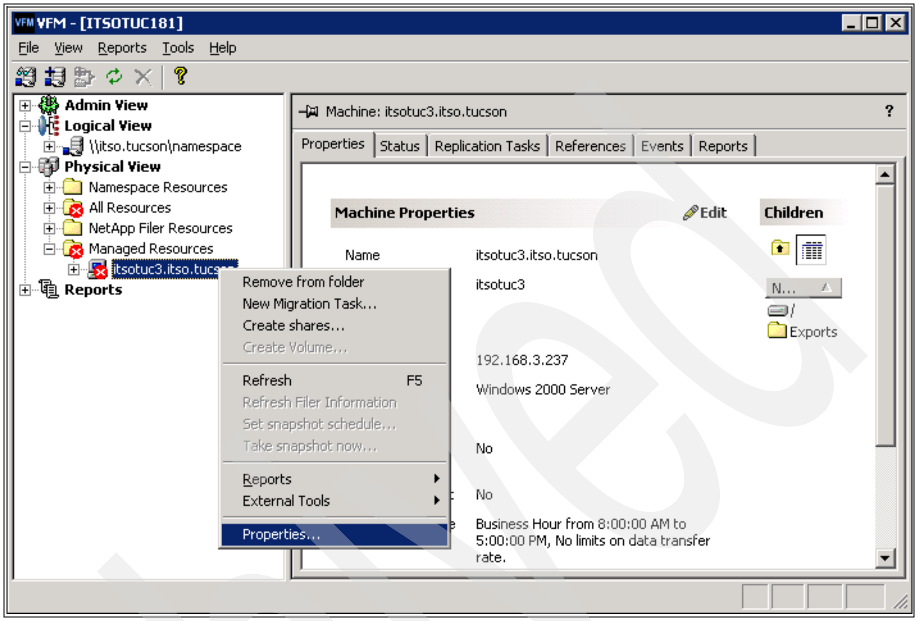

Name itsotuc3.itso.tucson

itsotuc3

192.168.3.237

Windows 2000 Server

No

No

Business Hour from 8:00:00 AM to 5:00:00 PM, No limits on data transfer rate.

Figure 12-11 Machine menu

2. Click the **Shell** tab (see Figure 12-12).

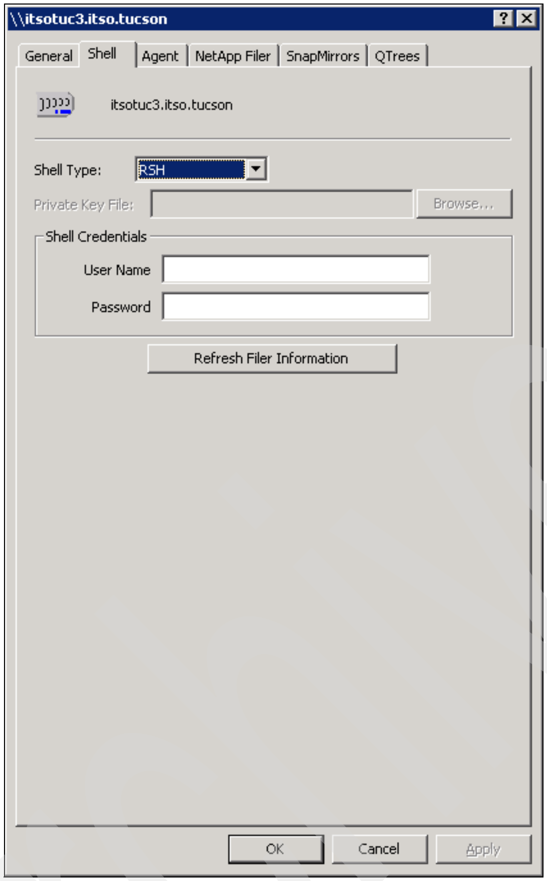

\\itsotuc3.itso.tucson

itsotuc3.itso.tucson

Shell Type: RSH

Private Key File:

Shell Credentials

User Name

Password

Figure 12-12 Machine properties: Shell tab

3. Make sure the **Shell Type** list box is set to RSH.

4. Enter the user name and password to access the N series (see Figure 12-13).

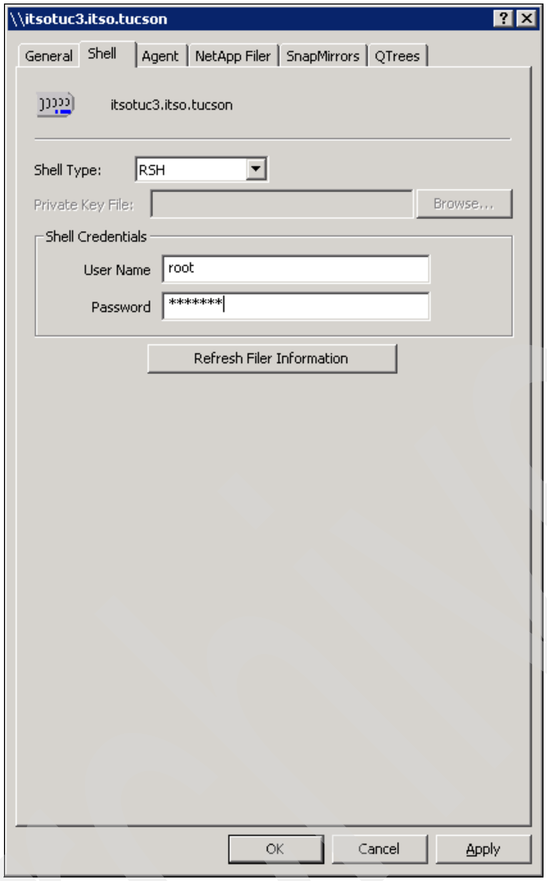

\\itsotuc3.itso.tucson

itsotuc3.itso.tucson

Shell Type: RSH

Private Key File:

Shell Credentials

User Name: root

Password: *******

Figure 12-13 Shell tab with root information

5. Click **OK**.

   After a few seconds, VFM updates the filer information and then removes the error sign from the machine (see Figure 12-14).

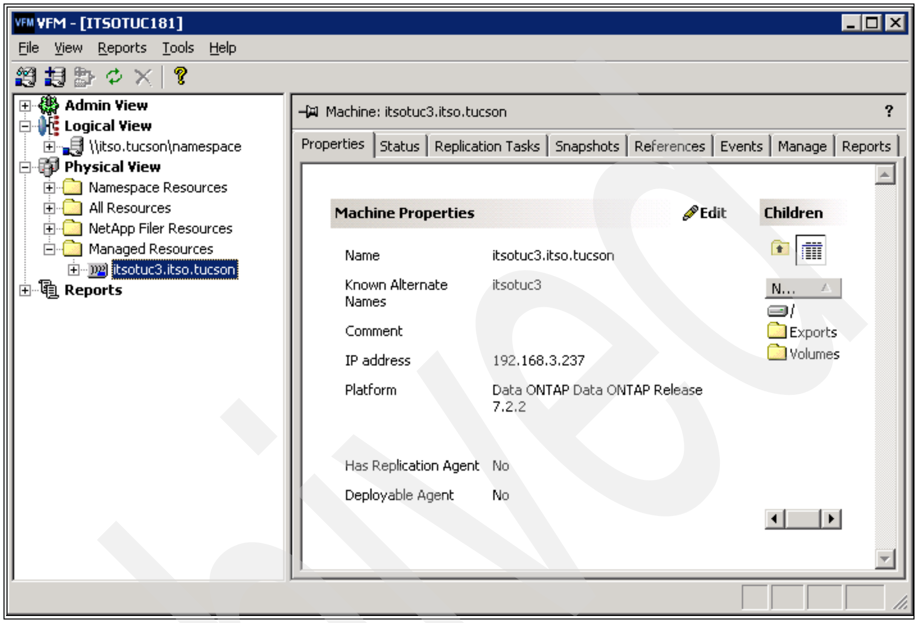

Name	itsotuc3.itso.tucson
Known Alternate Names	itsotuc3
Comment	
IP address	192.168.3.237
Platform	Data ONTAP Data ONTAP Release 7.2.2
Has Replication Agent	No
Deployable Agent	No

Figure 12-14 Filer with RSH information

12.4 Creating volumes

You can use VFM to create volumes on N series storage systems. In this section, we show the steps for doing so:

1. Add the N series storage system you wish to use to the Managed Resources folder (see 12.1, “Integrating IBM N series with VFM” on page 452 for instructions).
2. On the VFM client, choose **Physical View** → **Managed Resources**.
3. Right-click the N series you wish to use. In this example, we use itsotuc3.itso.tucson.

4. Click **Create Volume** (see Figure 12-15).

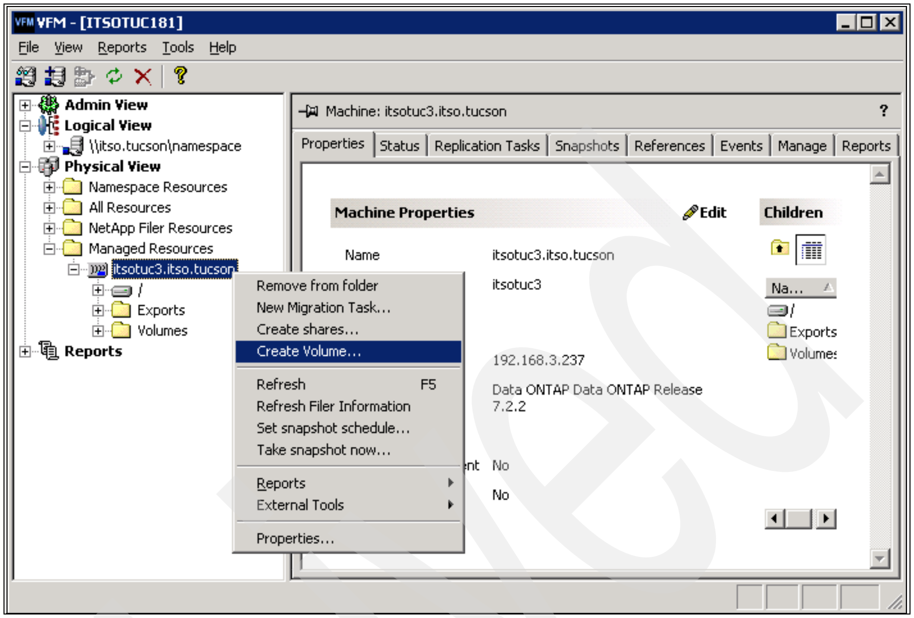

Name itsotuc3.itso.tucson

itsotuc3

192.168.3.237

Data ONTAP Data ONTAP Release 7.2.2

nt No

No

Figure 12-15 Machine menu: Create Volume

5. At this point, you can set the basic properties of the new volume and designate the disks to be used (see Figure 12-16). You can manually select the disks, or just choose the number of disks and disk size.

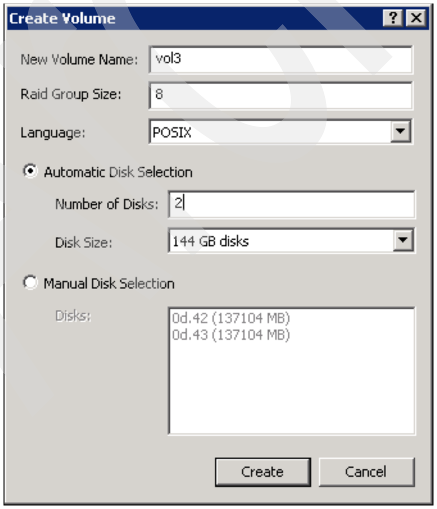

New Volume Name: vol3

Raid Group Size: 8

Language: POSIX

● Automatic Disk Selection

Number of Disks: 2

Disk Size: 144 GB disks

○ Manual Disk Selection

Disks:
0d.42 (137104 MB)
0d.43 (137104 MB)

Figure 12-16 Create Volume dialog

6. Click **Create**.

The volume is set and is being created.

7. You can check this by expanding the Volumes folder and clicking the just created volume (see Figure 12-17). The Volume State displays the creation status.

   Depending on the size of the disks you are using, the Volume State can remain in this status for some time.

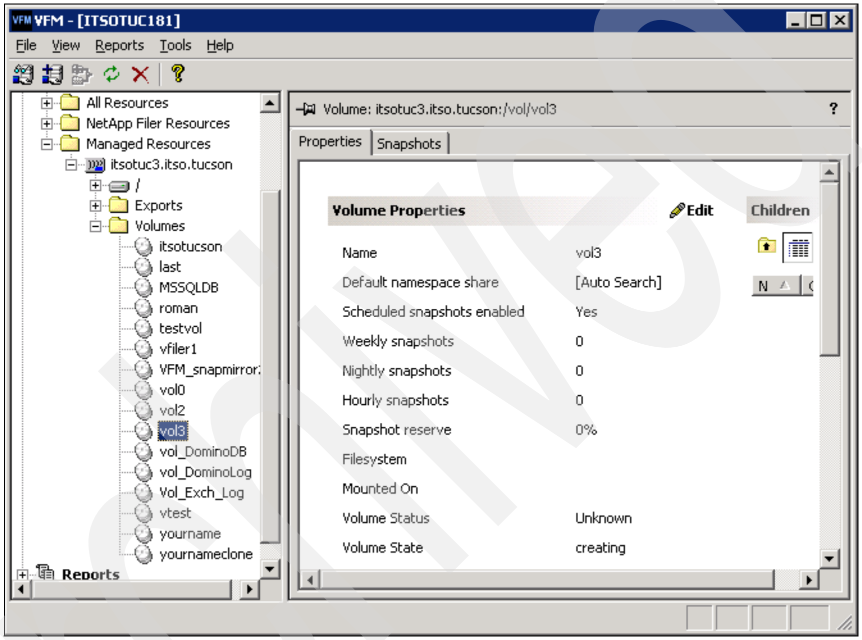

Name	vol3
Default namespace share	[Auto Search]
Scheduled snapshots enabled	Yes
Weekly snapshots	0
Nightly snapshots	0
Hourly snapshots	0
Snapshot reserve	0%
Filesystem	
Mounted On	
Volume Status	Unknown
Volume State	creating

Figure 12-17 Vol3 created

When creation is complete, the Volume State changes to online (see Figure 12-18), and the volume is ready to use.

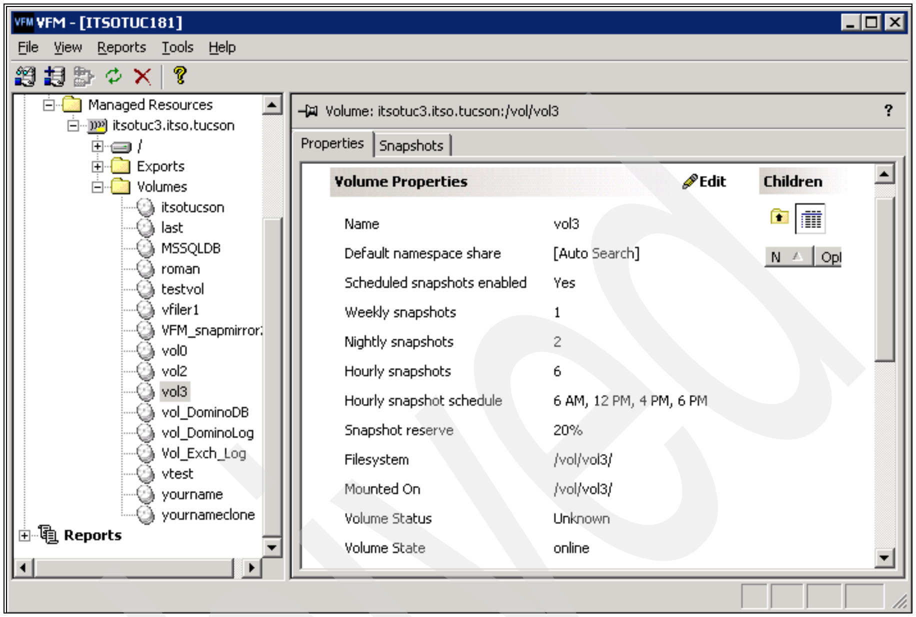

Name	vol3
Default namespace share	[Auto Search]
Scheduled snapshots enabled	Yes
Weekly snapshots	1
Nightly snapshots	2
Hourly snapshots	6
Hourly snapshot schedule	6 AM, 12 PM, 4 PM, 6 PM
Snapshot reserve	20%
Filesystem	/vol/vol3/
Mounted On	/vol/vol3/
Volume Status	Unknown
Volume State	online

Figure 12-18 Vol3 online

13

Reporting

This chapter provides an overview of the IBM Virtual File Manager reporting tool as well as examples demonstrating how to create reports. The following topics are discussed:

- VFM reports
- Types of VFM reports
  - Administrative reports
  - Logical namespace reports
  - Physical resources report
  - File system reports
  - N series reports
- Publishing reports
- Patch files
- Creating reports with VFM
  - Share permissions report
  - File differences report
  - File ages report

13.1 VFM reports

The VFM reporting tool provides a means of creating an inventory of assets or reports.

VFM provides over 30 preconfigured reports that provide valuable information about network data and storage resources at both the physical and the logical level. For example, you can determine file usage and access patterns for files by department by running a report at the logical level, or by storage resource by running a report at the physical level (Figure 13-1).

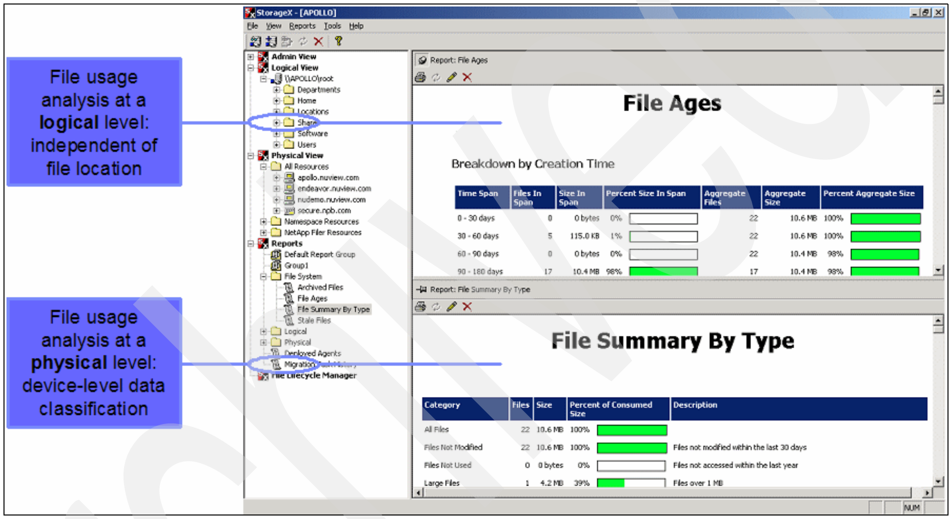

File usage analysis at a **logical** level: independent of file location

File usage analysis at a **physical** level: device-level data classification

File Ages

Breakdown by Creation Time

Time Span	Files In Span	Size In Span	Percent Size In Span		Aggregate Files	Aggregate Size	Percent Aggregate Size	
0 - 30 days	0	0 bytes	0%		22	10.6 MB	100%	🟩
30 - 60 days	5	115.0 KB	1%		22	10.6 MB	100%	🟩
60 - 90 days	0	0 bytes	0%		22	10.4 MB	98%	🟩
90 - 180 days	17	10.4 MB	98%	🟩	17	10.4 MB	98%	🟩

File Summary By Type

Category	Files	Size	Percent of Consumed Size		Description
All Files	22	10.6 MB	100%	🟩	
Files Not Modified	22	10.6 MB	100%	🟩	Files not modified within the last 30 days
Files Not Used	0	0 bytes	0%		Files not accessed within the last year
Large Files	1	4.2 MB	39%	🟩	Files over 1 MB

Figure 13-1 File usage analysis at logical and physical levels

You can use the flexible reporting capability of VFM in a number of ways to streamline storage operations and reduce storage costs. A share properties report indicates the capacity utilization levels of all storage resources within the global namespace to identify data migration targets and rebalance storage capacities. The Stale Files report identifies files that have not been accessed for a selected period of time and are candidates for movement to a secondary, or an archived, storage (Figure 13-2).

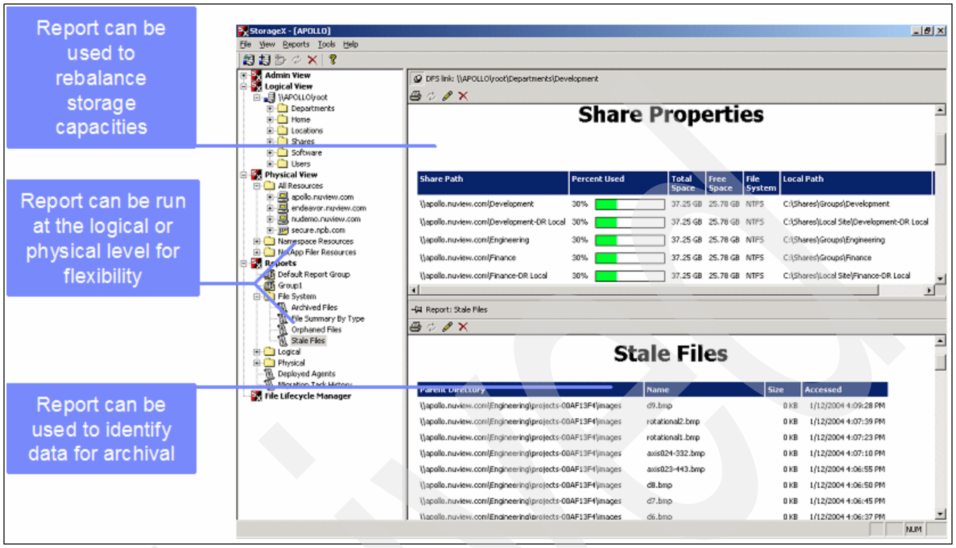

DFS link: \\APOLLO\root\Departments\Development

Share Properties

Share Path	Percent Used	Total Space	Free Space	File System	Local Path
\\apollo.nuview.com\Development	30%	37.25 GB	25.78 GB	NTFS	C:\Shares\Groups\Development
\\apollo.nuview.com\Development-DR Local	30%	37.25 GB	25.78 GB	NTFS	C:\Shares\Local Site\Development-DR Local
\\apollo.nuview.com\Engineering	30%	37.25 GB	25.78 GB	NTFS	C:\Shares\Groups\Engineering
\\apollo.nuview.com\Finance	30%	37.25 GB	25.78 GB	NTFS	C:\Shares\Groups\Finance
\\apollo.nuview.com\Finance-DR Local	30%	37.25 GB	25.78 GB	NTFS	C:\Shares\Local Site\Finance-DR Local

Report: Stale Files

Stale Files

Parent Directory	Name	Size	Accessed
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	d9.bmp	0 KB	1/12/2004 4:09:28 PM
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	rotational2.bmp	0 KB	1/12/2004 4:07:39 PM
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	rotational1.bmp	0 KB	1/12/2004 4:07:23 PM
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	axis024-332.bmp	0 KB	1/12/2004 4:07:10 PM
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	axis023-443.bmp	0 KB	1/12/2004 4:06:55 PM
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	d8.bmp	0 KB	1/12/2004 4:06:50 PM
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	d7.bmp	0 KB	1/12/2004 4:06:45 PM
\\apollo.nuview.com\Engineering\projects-00AF13F4\images	d6.bmp	0 KB	1/12/2004 4:06:37 PM

Figure 13-2 Stale Files report

The benefits of VFM reports are summarized as follows:

- Integrates over 30 reports for logical storage, physical storage, storage appliances, policies, and so on.
- Provides agent or agentless reporting.
- Reports on individual servers or across the enterprise.
- Enables the reuse of collected data for all reports.
- Enables proactive decision making.

13.2 Types of VFM reports

VFM enables you to create reports and publish them to a specified location, to process report data, and to archive reports. VFM provides the following two categories of reports:

- Reports generated from data periodically gathered by report source groups. These reports appear both under the report source groups to which they are assigned and under the Generated Reports folder.
- Reports generated from recorded information, such as events. These reports are not scheduled and appear only under the Generated Reports folder.

Within these two categories, VFM can generate the following five types of reports:

- Administrative reports
- Logical namespace reports
- Physical resources report
- File system reports
- N series reports

A description of the reports in each category is provided in the sections that follow.

13.2.1 Administrative reports

The three types of administrative reports are as follows:

- Disaster recovery policies status

  Enables you to review links contained in the selected disaster recovery policies and check storage and replication status for each target of a policy.
- Policy references

  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.

13.2.2 Logical namespace reports

The types of logical namespace reports are as follows:

- 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 the *BLOB size*).

> **Note:** Gathering BLOB size information can take a long time, especially for standalone roots that contain a large number of links.

- Physical references

  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.

- Policy references

  Enables you to view the resources referenced in policies. 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. The Replica properties reports option also enables you to view links that are not highly available because they have only one target.

13.2.3 Physical resources reports

This section describes the types of physical resources reports (see Figure 13-3).

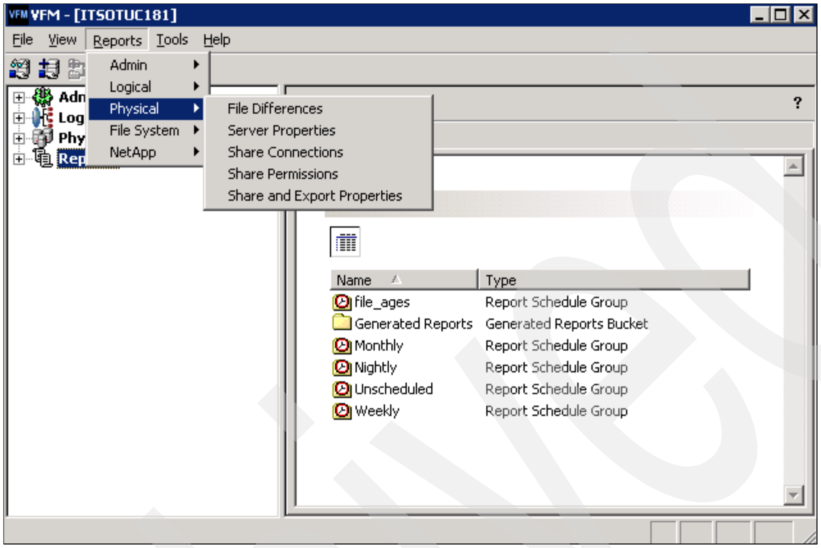

Name	Type
file_ages	Report Schedule Group
Generated Reports	Generated Reports Bucket
Monthly	Report Schedule Group
Nightly	Report Schedule Group
Unscheduled	Report Schedule Group
Weekly	Report Schedule Group

Figure 13-3 Types of physical resources reports

The physical resources reports are:

- File Differences

  Provides an easy way to compare the differences that might exist in the contents of network shares.

- Physical References (Physical View selection)

  Enables you to view the DFS links or roots that reference the targets in the namespace.

- Policy References (Physical View selection)

  Enables you to view the resources referenced in policies. Using the Report wizard, you can filter the report by items you select in the Physical View.

- 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.

13.2.4 File system reports

The types of file system reports are shown in Figure 13-4 and are described in this section.

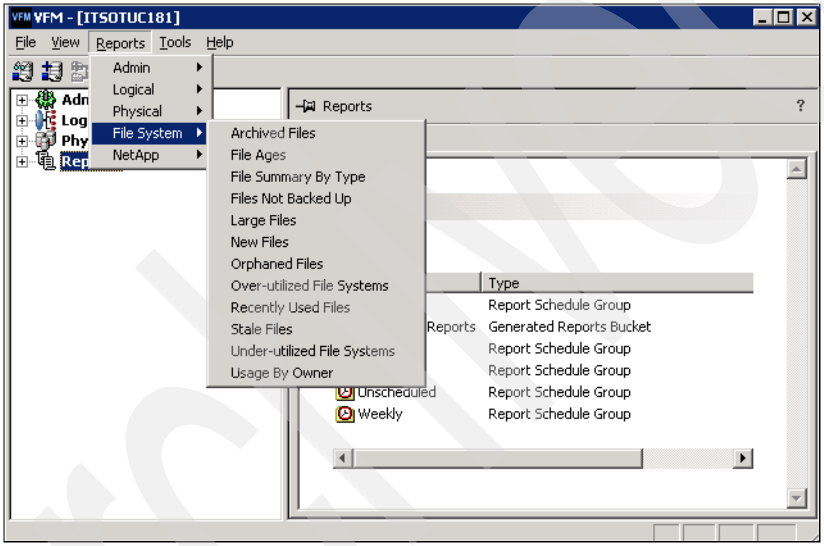

	Type
	Report Schedule Group
Reports	Generated Reports Bucket
	Report Schedule Group
	Report Schedule Group
Unscheduled	Report Schedule Group
Weekly	Report Schedule Group

Figure 13-4 Types of File System reports

VFM enables you to create these file system reports:

- Archived Files

  Files with no archive attribute set.
- File Ages

  Files arranged by age and file size for a specified location.
- File Summary By Type

  The numbers of files and space consumed based on the types of the files.
- Files Not Backed Up

  Files modified within a specified time and whose archive attribute is set.

- Large Files

  Files larger than a specified size in the selected locations. You can exclude files from the report that have been accessed recently.

- New Files

  Files created within a specified time.

- Orphaned Files

  Files whose owner is unknown in your environment.

- Over-utilized File Systems

  File systems with small amounts of free space remaining.

- Recently Used Files

  Files accessed within a specified time.

- Stale Files

  Files not accessed within a specified time.

- Under-utilized File Systems

  File systems containing a large amount of free space.

- Usage By Owner

  The amount of space consumed by files owned by each user.

13.2.5 N series reports

The 10 types of logical namespace reports are as follows:

- Filer CIFS Configuration

  Review information about the CIFS set up on a N series.

- Filer CIFS Statistics

  Review CIFS protocol usage information by N series. The report contains information about the physical filer and all vFilers hosted on the filer.

- Filer File System Status

  Review statistics about each volume on a N series.

- Filer Properties
  - View the filers participating in your namespace.
  - Review 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

  Review CIFS and NFS protocol operations for a N series.

- Filer Qtrees by Volume

  Review information about available qtrees by volume for a N series.

- Filer RAID

  Review RAID group configuration information by volume for a N series.

- Filer Snap Mirror Status

  Review statistics about Snap Mirror status on selected N series storage systems.

- Filer Snapshot

  Review information about available snapshots by volume for a N series.

- Filer vFiler Status

  Review statistics about each vFiler on selected N series.

13.3 Publishing reports

VFM reports can be published in a folder specified by the user. To publish VFM reports, report publishing must be enabled. Reports are published in HTML.

Specify a location for the published reports, and VFM generates an index file that is automatically copied to the same location as the published report. You must configure the global report publishing settings before publishing a report.

The steps for configuring a global report publishing location are as follows:

1. On the VFM main menu, click **Tools** and then **Options** as shown in Figure 13-5.

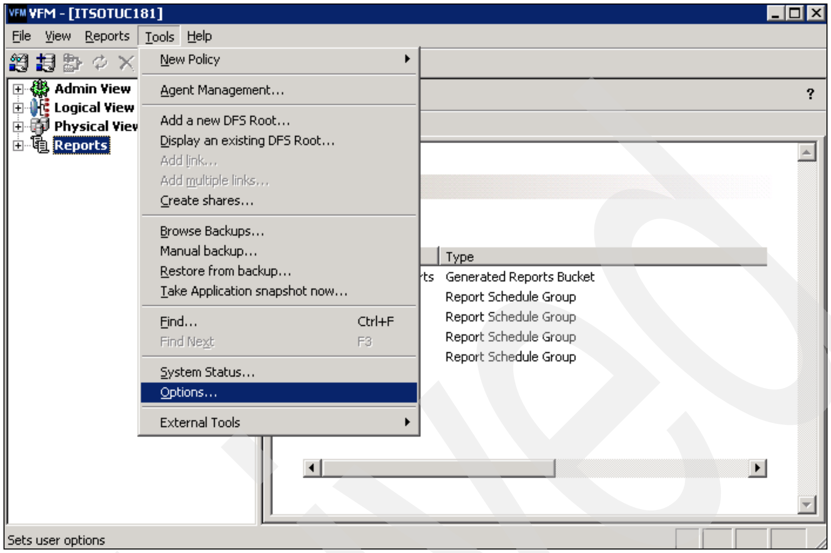

Figure 13-5 Configuring Publishing reports

2. Click **Reporting** under System Options on the console as shown in Figure 13-6.

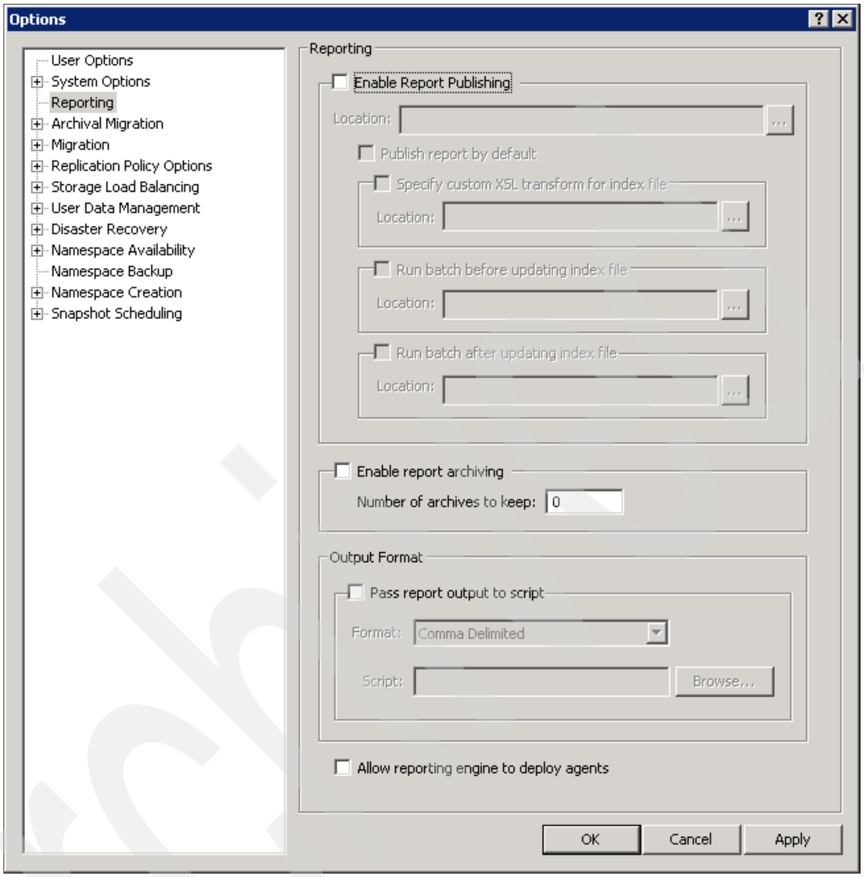

Options

Reporting

☐ Enable Report Publishing

Location: ...

☐ Publish report by default

☐ Specify custom XSL transform for index file

Location: ...

☐ Run batch before updating index file

Location: ...

☐ Run batch after updating index file

Location: ...

☐ Enable report archiving

Number of archives to keep: 0

Output Format

☐ Pass report output to script

Format: Comma Delimited

Script: Browse...

☐ Allow reporting engine to deploy agents

Figure 13-6 Enable Report Publishing

3. Click the **Enable Report Publishing** check box and enter or browse to the path of a folder where you intend the reports to be published (in this example, the path is \\itsotuc181.itso.tucson\VFM_reports). Select the **Publish report**

by default check box to set report publishing as the default report output for all new reports created, as shown in Figure 13-7.

.

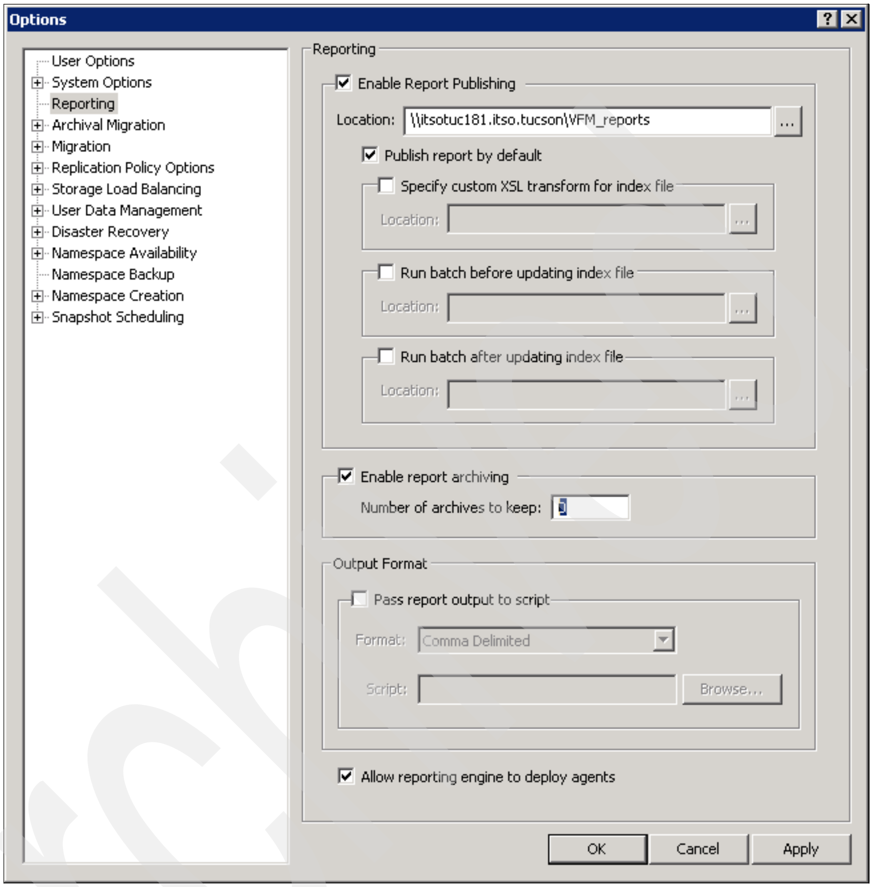

Reporting

☑ Enable Report Publishing

Location: \\itsotuc181.itso.tucson\VFM_reports

☑ Publish report by default

□ Specify custom XSL transform for index file

Location:

□ Run batch before updating index file

Location:

□ Run batch after updating index file

Location:

☑ Enable report archiving

Number of archives to keep: 0

Output Format

□ Pass report output to script

Format: Comma Delimited

Script:

Browse...

☑ Allow reporting engine to deploy agents

Figure 13-7 Setting Publishing reports location

The other selections on the Reporting panel are optional and are briefly described here:

– Specify custom XSL transform for index file

  Enter 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.

– Run batch before updating index file

  Enter the name of a batch file, or click **Browse** to find a batch file, to run before the index file is updated (refer to 13.4, “Batch files” on page 479 for details).

- Run batch after updating index file

  Enter the name of a batch file, or click Browse to find a batch file, to run after the index file is updated (refer to 13.4, “Batch files” on page 479 for details).

- Enable report archiving

  After checking this check box, specify the number of copies of a report you want to archive. You can specify this option for all reports or a single report.

- Output Format

  Select the **Pass report output to script** check box, specify a format, and browse for a batch file (see 13.4, “Batch files” on page 479 for details).

- Allow reporting engine to deploy agents

  Generate reports from agents (see 3.1, “Understanding VFM components” on page 84).

4. Click **OK** to apply your changes and exit the window.

13.4 Batch files

Batch files enable you to customize policies, such as sending an e-mail, posting an event in the Microsoft Windows event log, or sending an SNMP trap in response to a failure.

The batch files are run by VFM on the machine from which VFM monitors the targets of a policy, 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 fails, VFM does not continue the failover process. This process enables customized checking to be performed before the actual failover takes place.

Scripts specified in Run Batch Before Replication run before 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)

VFM provides example scripts in the installation directory under the Examples\Scripts folder. To function, these examples rely on the presence of additional third-party software, such as an RSH client or other utilities.

13.5 Creating reports with VFM

VFM provides the same report wizard steps for the creation of all categories and types of reports. The sections that follow describe examples showing how to create a Share Permissions report, File Differences report, and File Ages report. The first two are part of the Physical resources report category, and the latter one is part of the File System report category.

13.5.1 Creating Share Permissions report

In this example, a Share Permissions report is generated out of all the shares under the DFS namespace. A report for the following folders is generated:

- Namespace
- Dept
- Marketing and Sales (these last two are linked shares)

Follow these steps to create a Share Permissions report:

1. On the VFM main window, right-click the DFS namespace (in this example, \\itso.tucson\namespace) and choose **Reports** → **Physical** → **Share Permissions** as shown in Figure 13-8 on page 481.

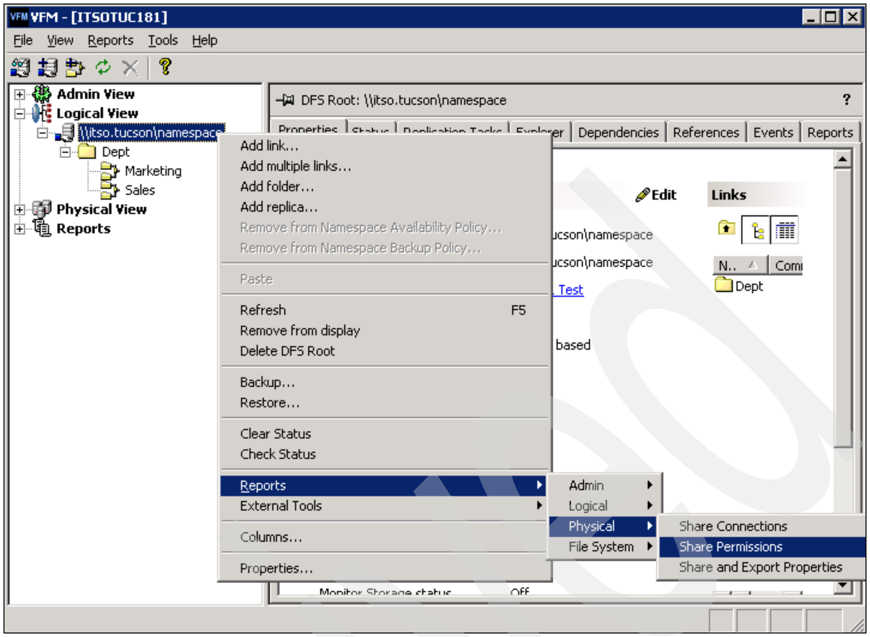

Figure 13-8 Creating a Share Permissions report

2. On the Report Wizard window, select the **Publish report** check box (refer to 13.3, “Publishing reports” on page 475 for Publish report details). Click **Next** as shown in Figure 13-9.

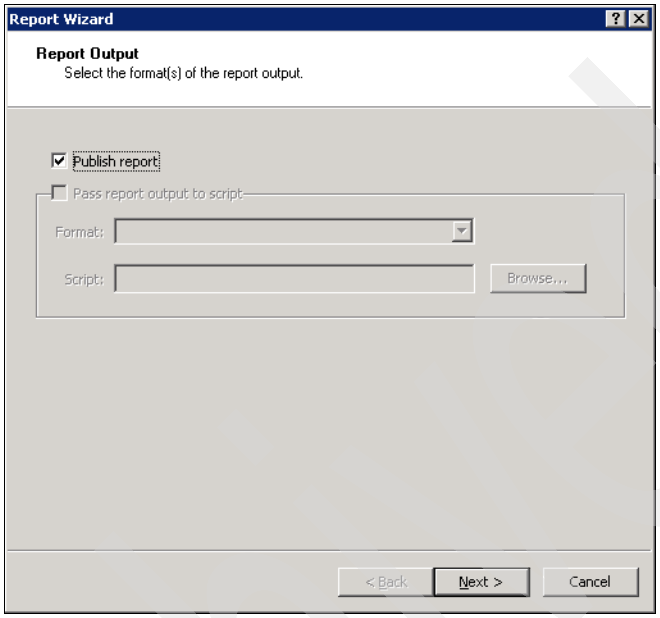

Report Output

Select the format(s) of the report output.

☑ Publish report

☐ Pass report output to script

Format:

Script:

Figure 13-9 Report Wizard

On the Report Group panel, four default report groups (Monthly, Nightly, Unscheduled, and Weekly) are shown (Figure 13-10).

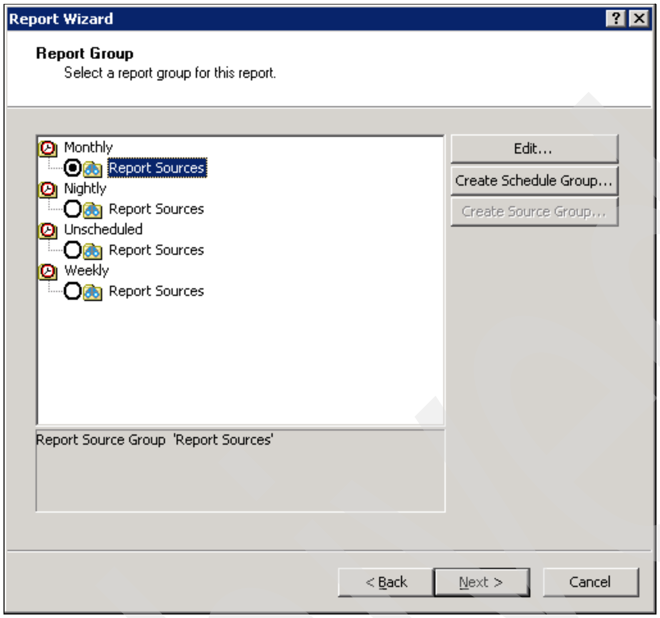

Report Wizard

Report Group

Select a report group for this report.

- Monthly
  - ● Report Sources
- Nightly
  - ○ Report Sources
- Unscheduled
  - ○ Report Sources
- Weekly
  - ○ Report Sources

Report Source Group 'Report Sources'

Figure 13-10 Report Group panel

3. In this example, we create another group. After clicking **Create Schedule Group** button, the window shown in Figure 13-11 is displayed.

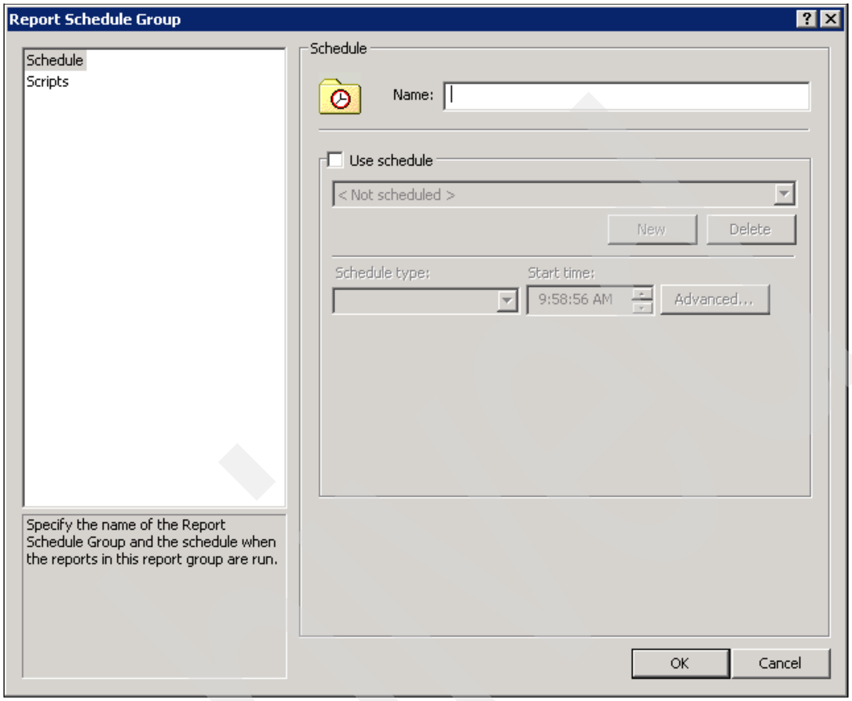

Specify the name of the Report Schedule Group and the schedule when the reports in this report group are run.

Figure 13-11 Schedule panel

4. Enter a name that best applies (in this example, the name is share_permissions), and select the **Use schedule** check box.

5. The **New** button is displayed. Click **New** to set the Schedule type and Start time (in this example, the schedule type is Every several minutes and Every 5 minutes) as shown in Figure 13-12.

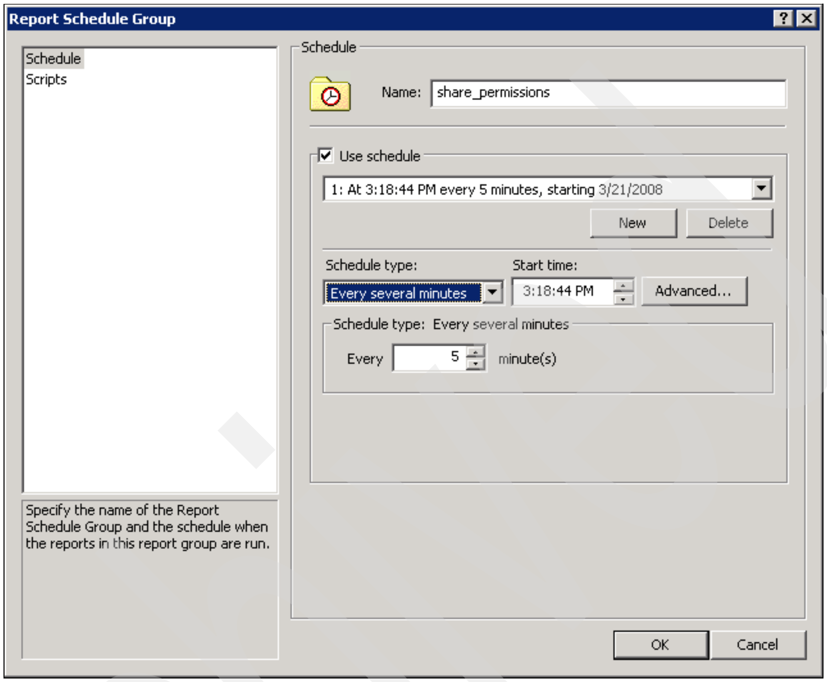

Report Schedule Group

Schedule
Scripts

Specify the name of the Report Schedule Group and the schedule when the reports in this report group are run.

Schedule

Name: share_permissions

☑ Use schedule

1: At 3:18:44 PM every 5 minutes, starting 3/21/2008

New Delete

Schedule type: Every several minutes

Start time: 3:18:44 PM

Advanced...

Schedule type: Every several minutes

Every 5 minute(s)

OK Cancel

Figure 13-12 Setting report group name and schedule

6. Click **OK**. The newly created group is displayed (in this example, share_permissions). Select it and click **Next** (see Figure 13-13).

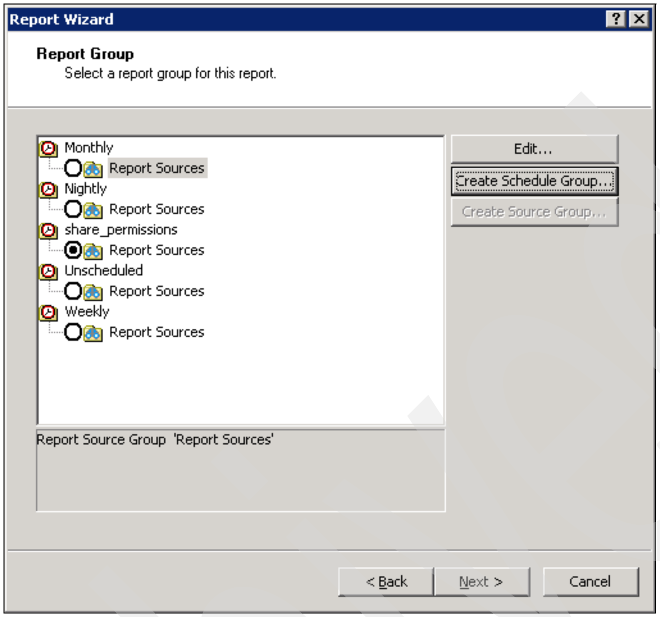

Report Wizard

Report Group

Select a report group for this report.

Report Source Group 'Report Sources'

Figure 13-13 New report group added

7. A VFM pop-up message is displayed (see Figure 13-14), which prompts you to add the context items. Click **Yes** and then **OK**.

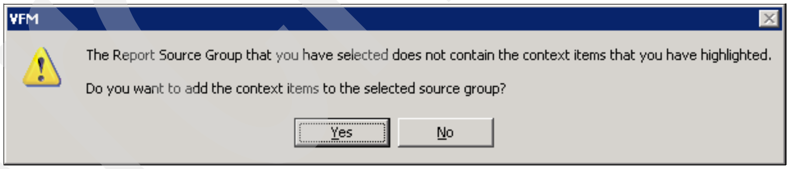

VFM

The Report Source Group that you have selected does not contain the context items that you have highlighted.

Do you want to add the context items to the selected source group?

Figure 13-14 VFM pop-up message

8. On the Report name panel, enter the name that best applies (in this example, share_permissions), and click **Enable reporting archive** if appropriate (in this example, enter 1 in the **Number of archives to keep** field). Click **Finish** as shown in Figure 13-15.

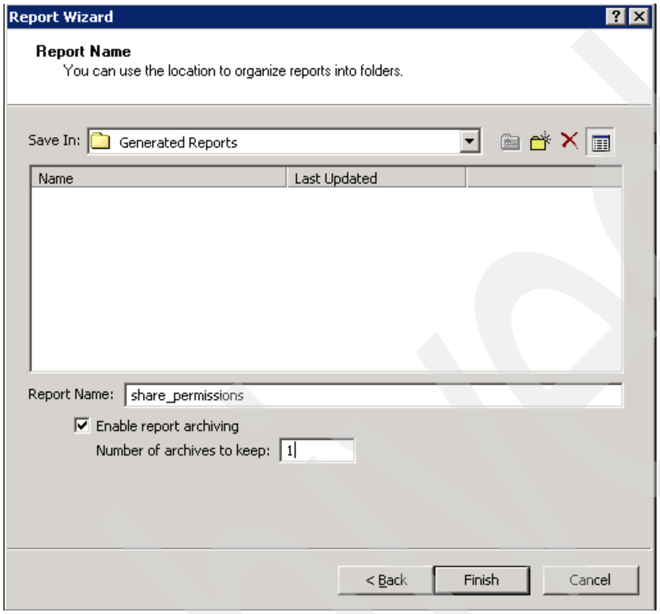

Report Name: share_permissions

☑ Enable report archiving

Number of archives to keep: 1

Figure 13-15 Report Name panel

9. To visualize the report, on the VFM main window expand **Reports** → **share_permissions** and click **Share_Permission**s as shown in Figure 13-16.

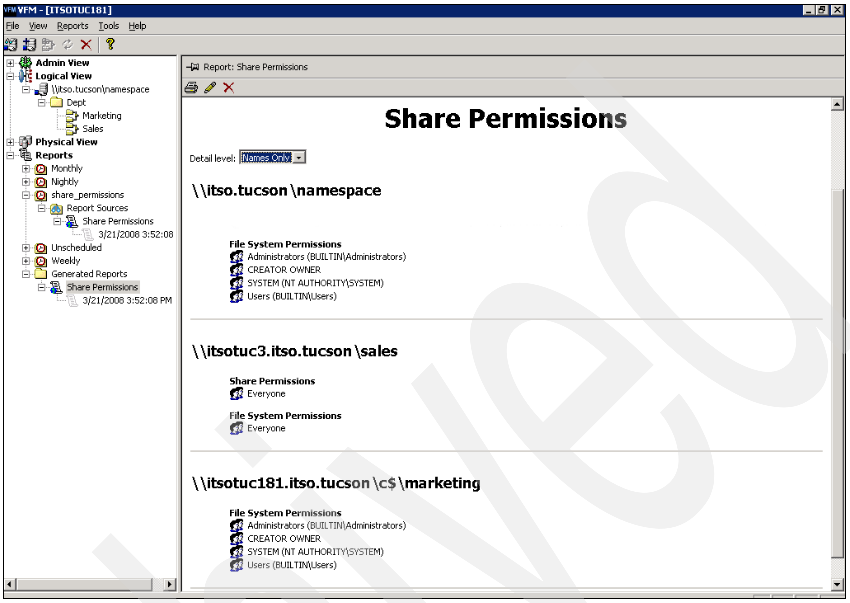

Share Permissions

Detail level: Names Only

\\itso.tucson\namespace

File System Permissions
Administrators (BUILTIN\Administrators)
CREATOR OWNER
SYSTEM (NT AUTHORITY\SYSTEM)
Users (BUILTIN\Users)

\\itsotuc3.itso.tucson\sales

Share Permissions
Everyone

File System Permissions
Everyone

\\itsotuc181.itso.tucson\c$\marketing

File System Permissions
Administrators (BUILTIN\Administrators)
CREATOR OWNER
SYSTEM (NT AUTHORITY\SYSTEM)
Users (BUILTIN\Users)

Figure 13-16 Share Permissions report visualization

A detail-level option with two options for visualization (summary and advanced) is shown in Figure 13-17.

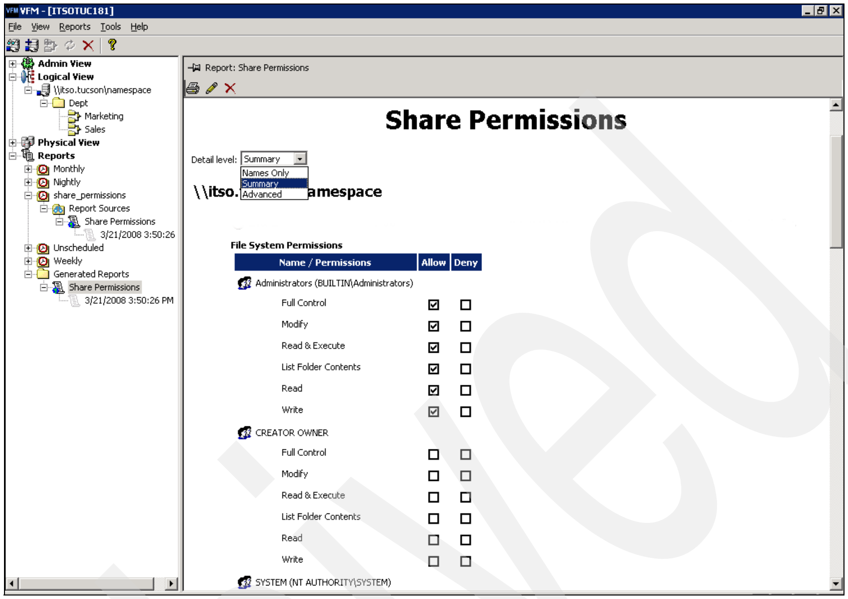

Share Permissions

Detail level: Summary

\\itso.tucson\namespace

File System Permissions

Name / Permissions	Allow	Deny
Administrators (BUILTIN\Administrators)		
Full Control	☑	□
Modify	☑	□
Read & Execute	☑	□
List Folder Contents	☑	□
Read	☑	□
Write	☑	□
CREATOR OWNER		
Full Control	□	□
Modify	□	□
Read & Execute	□	□
List Folder Contents	□	□
Read	□	□
Write	□	□
SYSTEM (NT AUTHORITY\SYSTEM)		

Figure 13-17 Summary-level report window

10. Move the scroll bar to visualize all the report contents.

13.5.2 Creating File Differences report

In this example, two folders labeled Marketing are compared. The source folder is located on a different server than the destination folder. In this example, the source Marketing folder is located in \\itsotuc181.itso.tucson\c$\Marketing, and the destination folder is located in \\itsotuc4.itso.tucson\c$\Marketing. Each folder has a different size and different contents.

To create a File Differences report, refer to the following steps:

1. On the VFM main window, expand the DFS namespace to the Marketing link (in this example, a link to the Marketing share was added under the Dept folder; refer to Chapter 5, “DFS namespace” on page 117 for details on adding links). Right-click **Marketing** and choose **Reports** → **Physical** → **File Differences** as shown in Figure 13-18 on page 490.

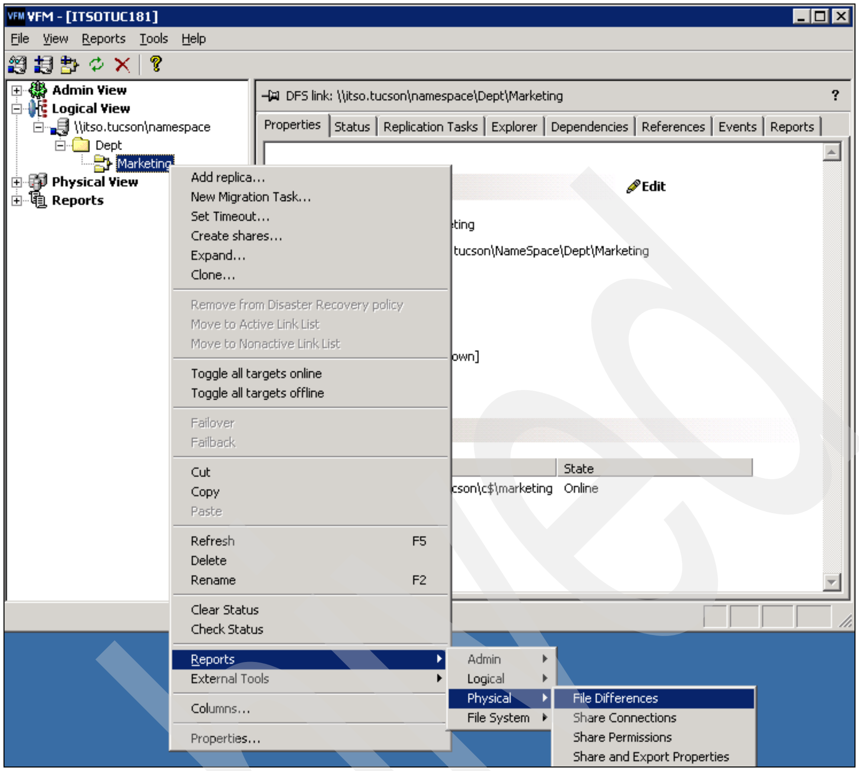

Figure 13-18 Creating a Filer Differences report

After you click **Filer Differences**, the Report Wizard window is displayed (see Figure 13-19).

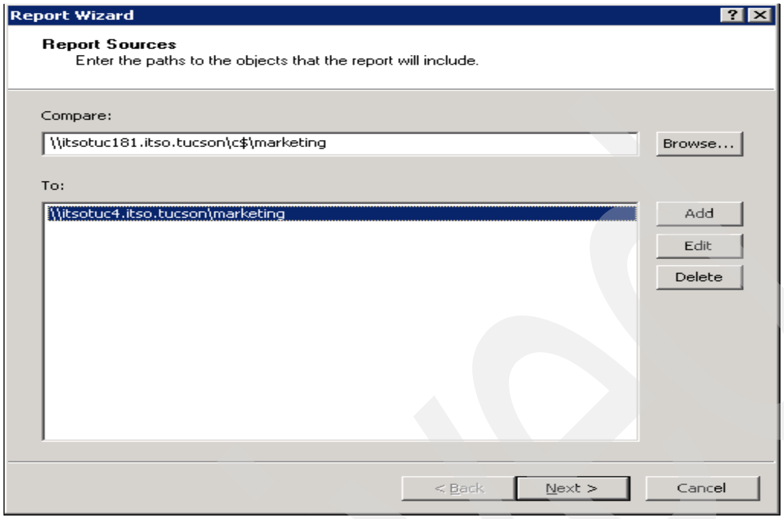

Compare: \\itsotuc181.itso.tucson\c$\marketing

To: \\itsotuc4.itso.tucson\marketing

Figure 13-19 Report Sources panel

2. In the Compare field, enter the name that best applies or browse the source path for the name. Click **Add** to enter the name of or browse to the destination folder. (In this example, the source folder is \\itsotuc181.itso.tucson\c$\Marketing, and the destination is \\itsotuc4.itso.tucson\c$\Marketing.)

3. Then click **Next**. The Report Output panel is displayed (see Figure 13-20).

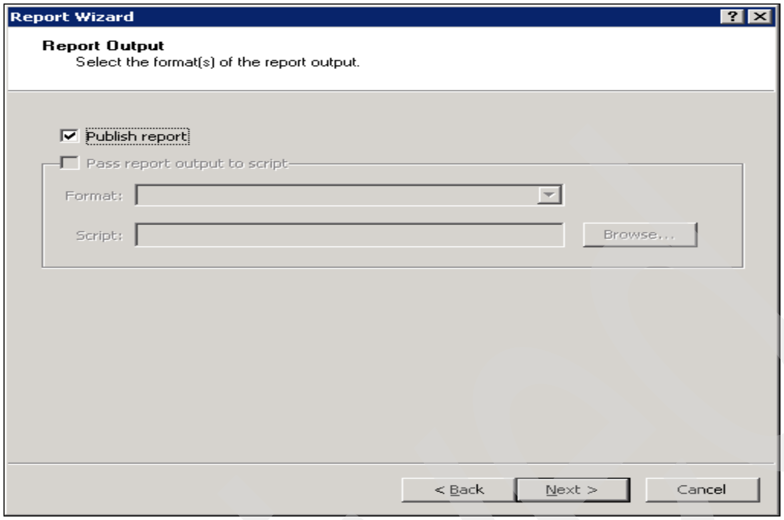

Report Output

Select the format(s) of the report output.

☑ Publish report

□ Pass report output to script

Format:

Script:

Figure 13-20 Report Output panel

4. Select **Publish report** and click **Next** (refer to 13.3, “Publishing reports” on page 475 for Publishing report details). The Report Group panel is displayed (see Figure 13-21).

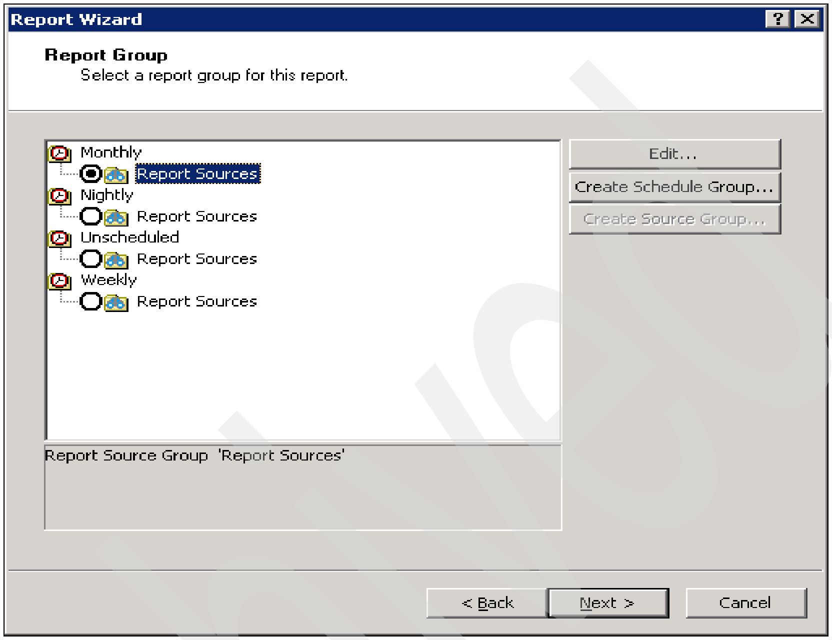

Report Wizard

Report Group

Select a report group for this report.

- Monthly
  - ● Report Sources
- Nightly
  - ○ Report Sources
- Unscheduled
  - ○ Report Sources
- Weekly
  - ○ Report Sources

Report Source Group 'Report Sources'

Figure 13-21 Report Group panel

5. The Schedule panel displays four default groups (Monthly, Nightly, Unscheduled, and Weekly). In this example, we create another group. After clicking the **Create Schedule Group** button, the window shown in Figure 13-22 is displayed.

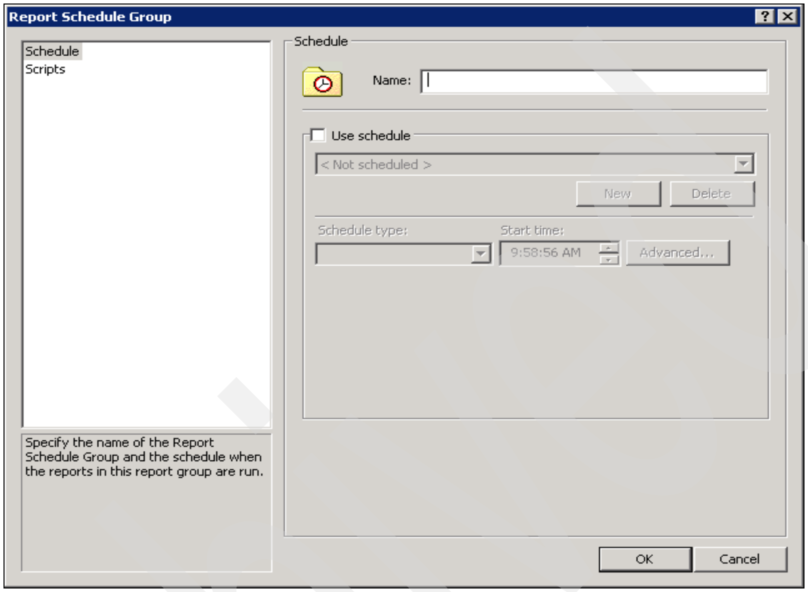

Name:

□ Use schedule

< Not scheduled >

New Delete

Schedule type:

Start time: 9:58:56 AM

Advanced...

Specify the name of the Report Schedule Group and the schedule when the reports in this report group are run.

Figure 13-22 Schedule panel

6. Enter a name that best applies (in this example, we enter the name file_differences) and select **Use schedule**.

7. The **New** button is displayed. Click **New** to set the Schedule type and Start time (in this example, the Schedule type is Every several minutes and Every 1 minutes) as shown in Figure 13-23.

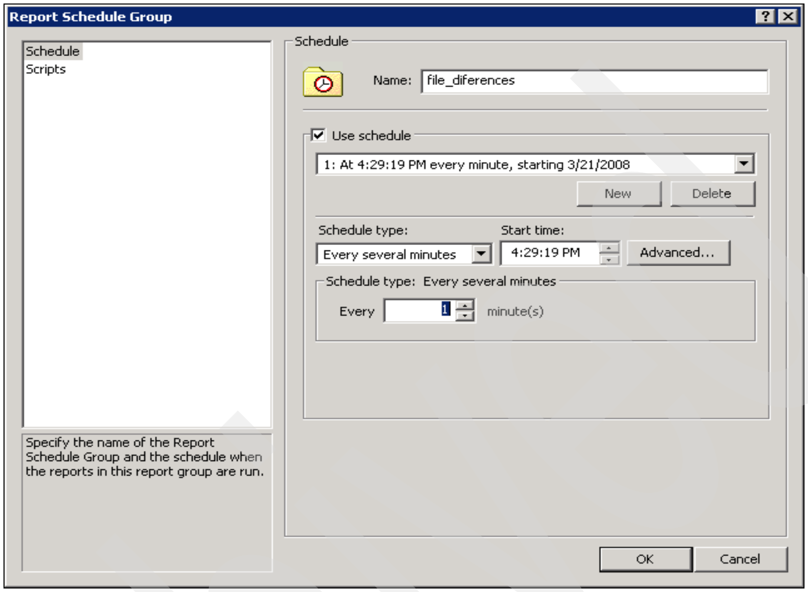

Report Schedule Group

Schedule
Scripts

Specify the name of the Report Schedule Group and the schedule when the reports in this report group are run.

Schedule

Name: file_diferences

☑ Use schedule

1: At 4:29:19 PM every minute, starting 3/21/2008

Schedule type: Every several minutes

Start time: 4:29:19 PM

Schedule type: Every several minutes

Every 1 minute(s)

Figure 13-23 Setting report group name and schedule

8. Click **OK**. The previous window is displayed with the newly created group as shown in Figure 13-24.

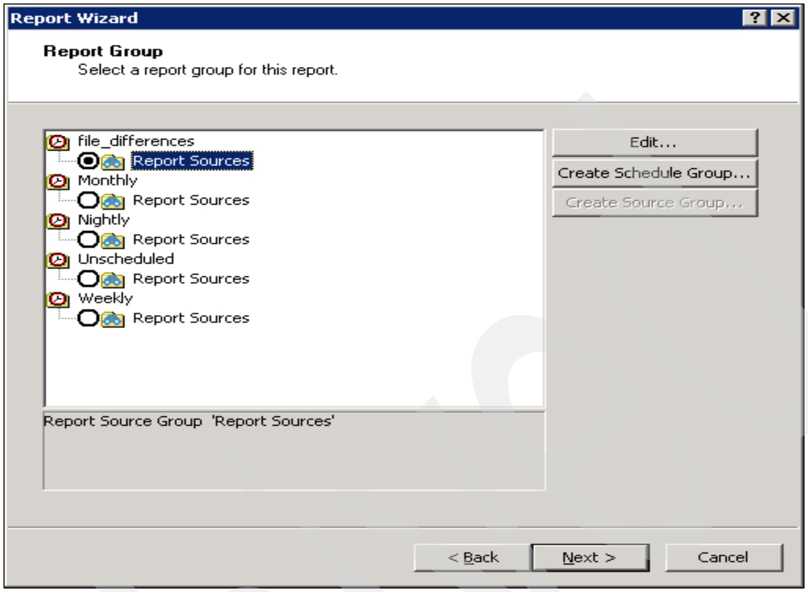

Figure 13-24 New report group added

9. Click **Next**. A VFM pop-up message is displayed (see Figure 13-25) prompting you to add the context items. Click **Yes** and then **OK**.

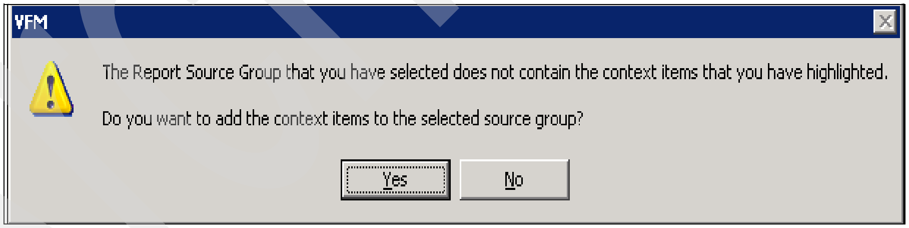
VFM

The Report Source Group that you have selected does not contain the context items that you have highlighted.

Do you want to add the context items to the selected source group?

Figure 13-25 VFM pop-up message

10. On the Report Name panel, enter the name that best applies (in this example, file_differences), and select **Enable report archiving** if appropriate. (In this example, we set the Number of archives to keep at 1.) Click **Finish** (see Figure 13-26).

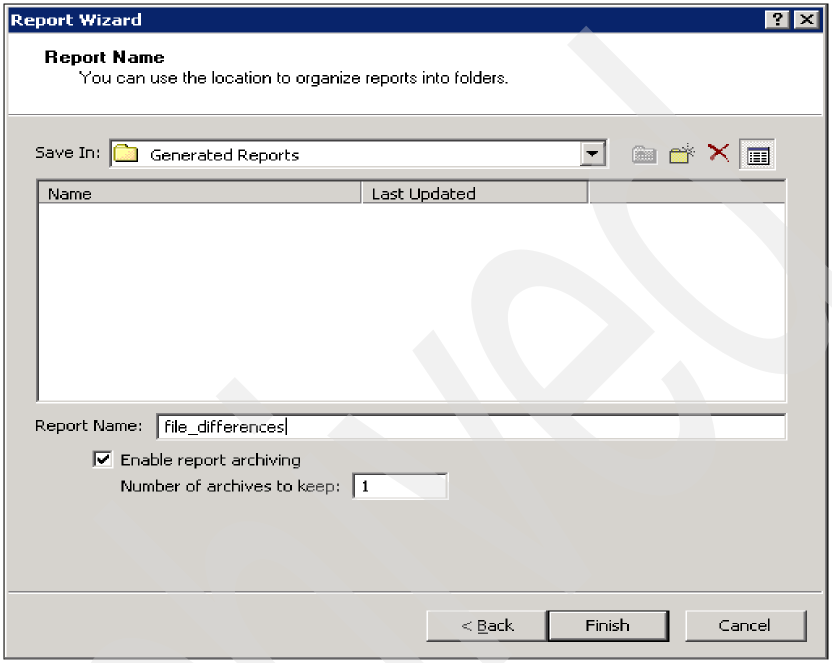

Report Wizard

Report Name

You can use the location to organize reports into folders.

Save In: Generated Reports

Name	Last Updated

Report Name: file_differences

☑ Enable report archiving

Number of archives to keep: 1

Figure 13-26 Report Name panel

11. To visualize the report, access the VFM main window and expand **Reports** → **file_differences**, then click **File Differences** (see Figure 13-27).

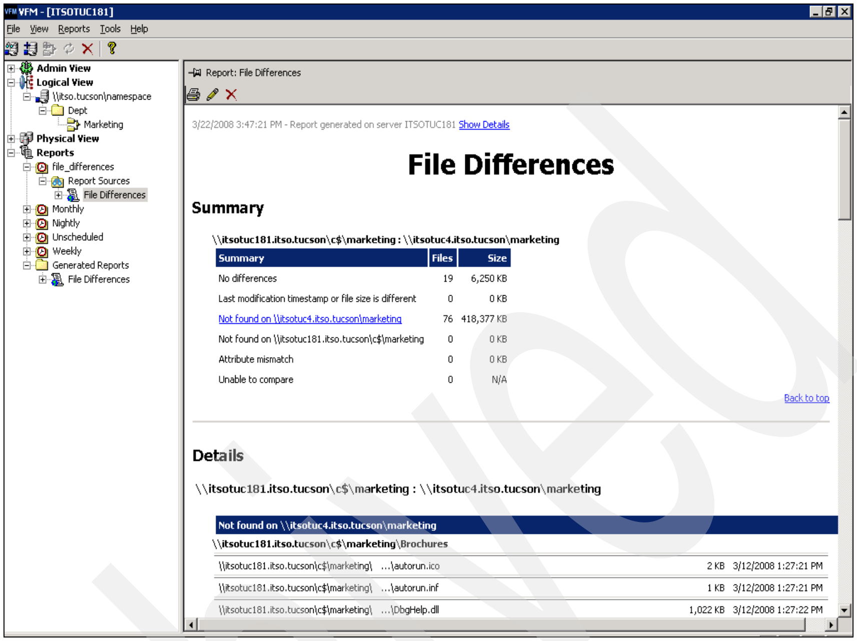

3/22/2008 3:47:21 PM - Report generated on server ITSOTUC181 Show Details

File Differences

Summary

\\itsotuc181.itso.tucson\c$\marketing : \\itsotuc4.itso.tucson\marketing

Summary	Files	Size
No differences	19	6,250 KB
Last modification timestamp or file size is different	0	0 KB
Not found on \\itsotuc4.itso.tucson\marketing	76	418,377 KB
Not found on \\itsotuc181.itso.tucson\c$\marketing	0	0 KB
Attribute mismatch	0	0 KB
Unable to compare	0	N/A

Back to top

Details

\\itsotuc181.itso.tucson\c$\marketing : \\itsotuc4.itso.tucson\marketing

Not found on \\itsotuc4.itso.tucson\marketing

\\itsotuc181.itso.tucson\c$\marketing\Brochures

\\itsotuc181.itso.tucson\c$\marketing\	...\autorun.ico	2 KB	3/12/2008 1:27:21 PM
\\itsotuc181.itso.tucson\c$\marketing\	...\autorun.inf	1 KB	3/12/2008 1:27:21 PM
\\itsotuc181.itso.tucson\c$\marketing\	...\DbgHelp.dll	1,022 KB	3/12/2008 1:27:22 PM

Figure 13-27 File Differences report window

12. The Summary field lists items comparing the two folders. Move the scroll bar to view report details about the differences between the two folders.

    All the contents in the source folder (\\itsotuc181.itso.tucson\c$\Marketing) are migrated to the destination folder (\\itsotuc4.itso.tucson\c$\Marketing) so that the two folders have the same size and number of files. The results are displayed in Figure 13-28.

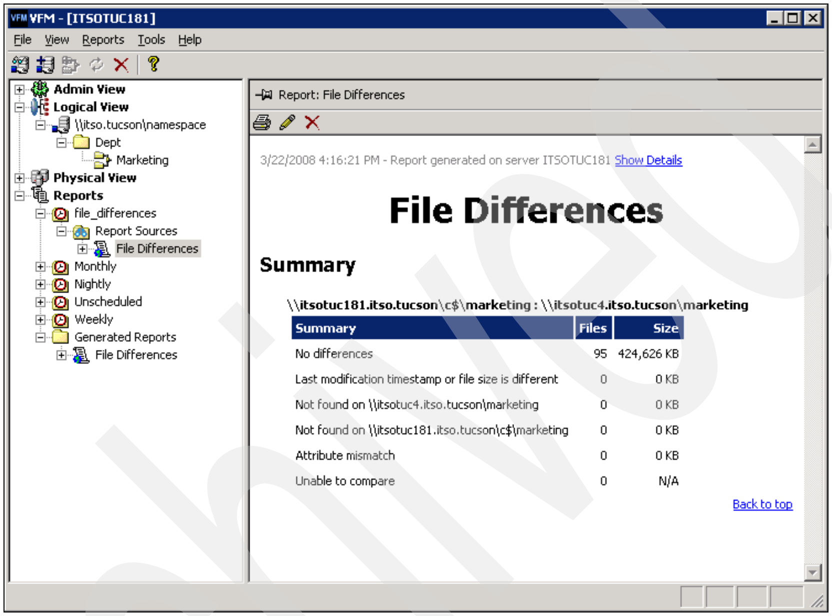

3/22/2008 4:16:21 PM - Report generated on server ITSOTUC181 Show Details

File Differences

Summary

\\itsotuc181.itso.tucson\c$\marketing : \\itsotuc4.itso.tucson\marketing

Summary	Files	Size
No differences	95	424,626 KB
Last modification timestamp or file size is different	0	0 KB
Not found on \\itsotuc4.itso.tucson\marketing	0	0 KB
Not found on \\itsotuc181.itso.tucson\c$\marketing	0	0 KB
Attribute mismatch	0	0 KB
Unable to compare	0	N/A

Back to top

Figure 13-28 File differences report displaying folders of the same size

13.5.3 Creating File Ages report

In this example, file information from two shares, Marketing and Sales, under the DFS namespace. To create a File Ages report, follow these steps:

1. On the VFM main window, expand the **DFS namespace to Dept** folder. (In this example, we add the Dept folder as well as the links. (Refer to Chapter 5, “DFS namespace” on page 117 for details on how to add folders and links.)
2. Right-click **Dept** and choose **Reports** → **File System** → **File Ages** as shown in Figure 13-29 on page 500.

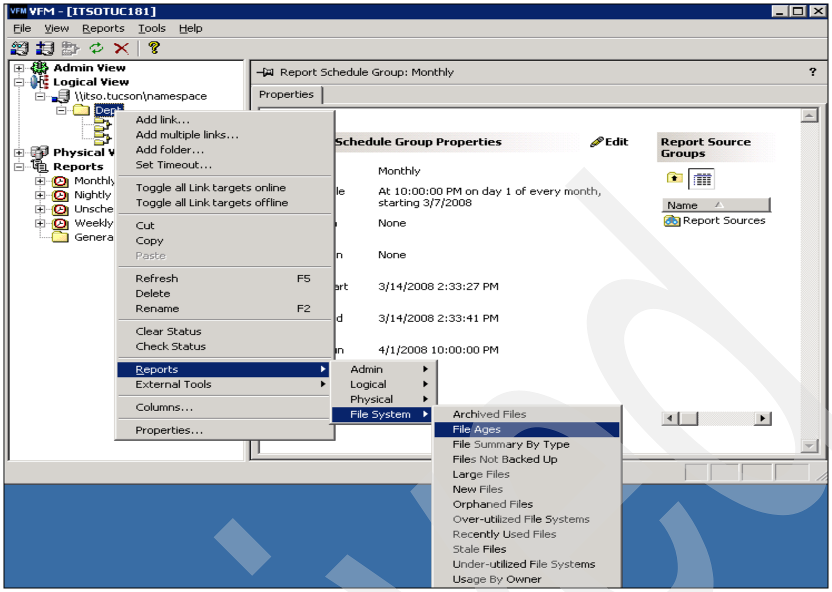

Monthly

le At 10:00:00 PM on day 1 of every month, starting 3/7/2008

None

n None

rt 3/14/2008 2:33:27 PM

d 3/14/2008 2:33:41 PM

n 4/1/2008 10:00:00 PM

Figure 13-29 Creating a File Ages report

3. After clicking **File Ages**, the Report Wizard window is shown (see Figure 13-30).

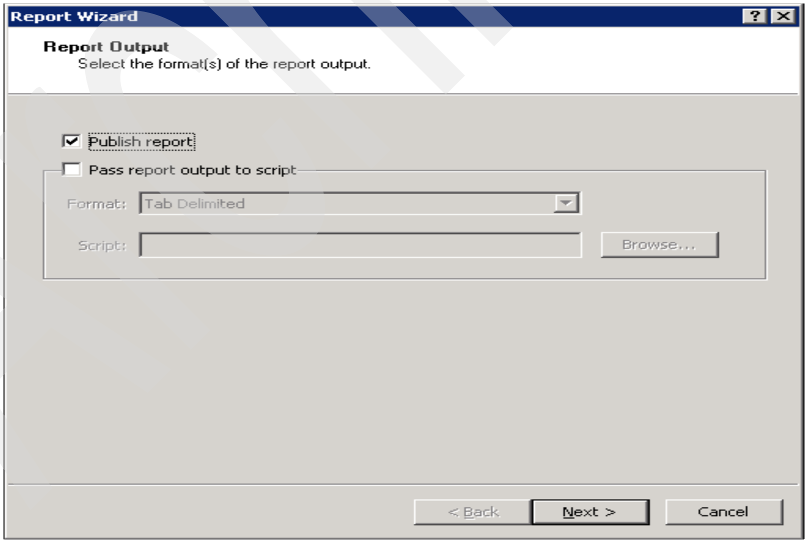

Report Output
Select the format(s) of the report output.

☑ Publish report

□ Pass report output to script

Format: Tab Delimited

Script:

Figure 13-30 Report Output panel

4. Select **Publish report** and click **Next** (refer to 13.3, “Publishing reports” on page 475 for publishing report details). The Report Group panel is displayed (see Figure 13-31).

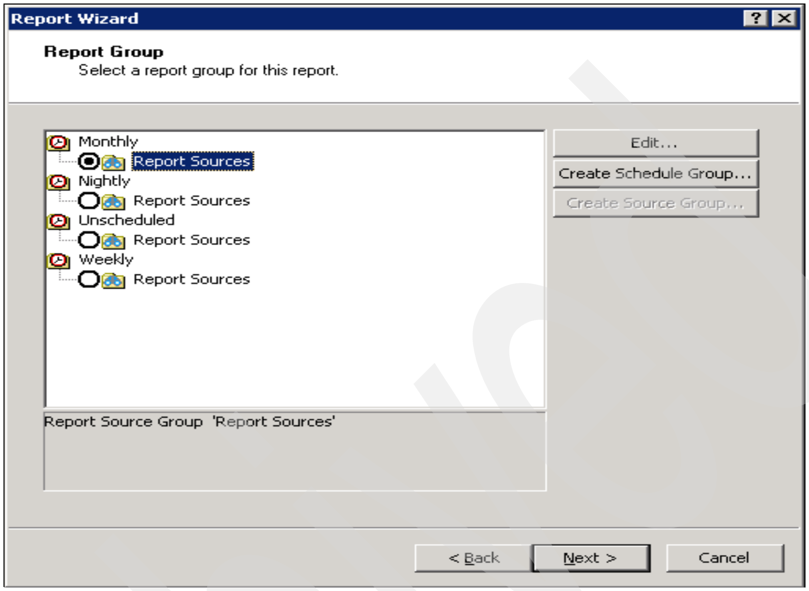

Report Wizard

Report Group

Select a report group for this report.

- Monthly
  - ● Report Sources
- Nightly
  - ○ Report Sources
- Unscheduled
  - ○ Report Sources
- Weekly
  - ○ Report Sources

Report Source Group 'Report Sources'

Figure 13-31 Report Group panel

5. The Report Group panel displays four default groups (Monthly, Nightly, Unscheduled, and Weekly). In our example, we create another group. After you click the **Create Schedule Group** button, the window shown in Figure 13-32 is displayed.

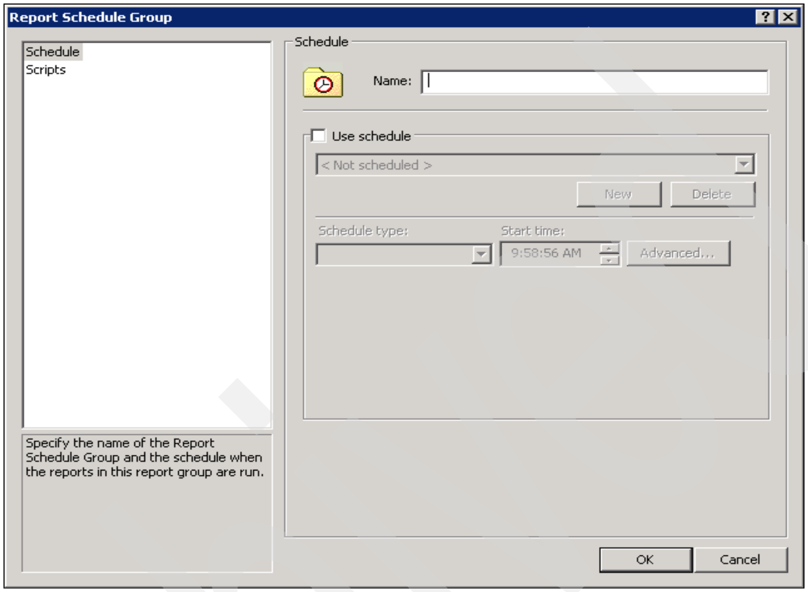

Specify the name of the Report Schedule Group and the schedule when the reports in this report group are run.

Figure 13-32 Report Schedule Group window

6. Enter a name that best applies (in our example, we enter the name file_ages), and select **Use schedule**.

7. The New button is displayed. Click **New** to set the Schedule type and Start time (in this example, the Schedule type is Every several minutes and Every 1 minutes) as shown in Figure 13-33.

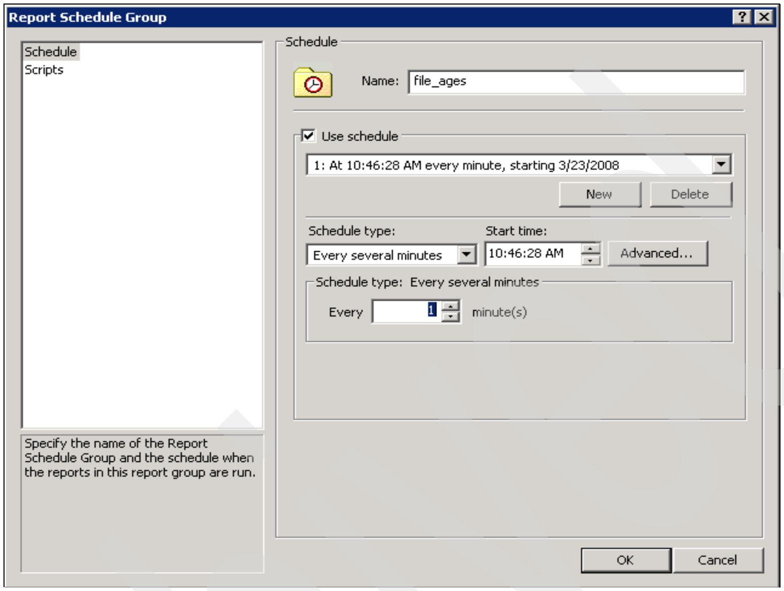

Name: file_ages

☑ Use schedule

1: At 10:46:28 AM every minute, starting 3/23/2008

Schedule type: Every several minutes

Start time: 10:46:28 AM

Schedule type: Every several minutes

Every 1 minute(s)

Specify the name of the Report Schedule Group and the schedule when the reports in this report group are run.

Figure 13-33 Setting group name and schedule

8. Click **OK** when you are finished. The previous window is displayed with the newly created group as shown in Figure 13-34.

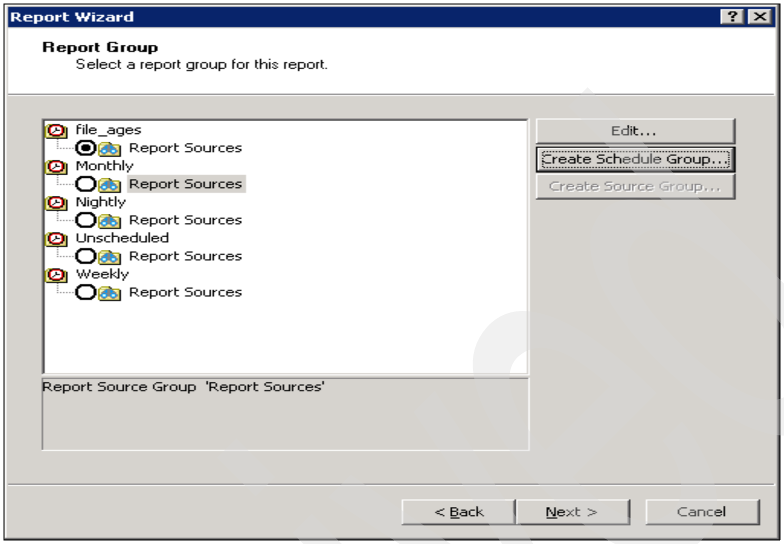

Report Group
Select a report group for this report.

- file_ages
  - ● Report Sources
- Monthly
  - ○ Report Sources
- Nightly
  - ○ Report Sources
- Unscheduled
  - ○ Report Sources
- Weekly
  - ○ Report Sources

Report Source Group 'Report Sources'

Figure 13-34 New group report added

9. Click **Next**. A VFM pop-up message is displayed (see Figure 13-35), prompting you to add the context items. Click **Yes** and then **OK**.

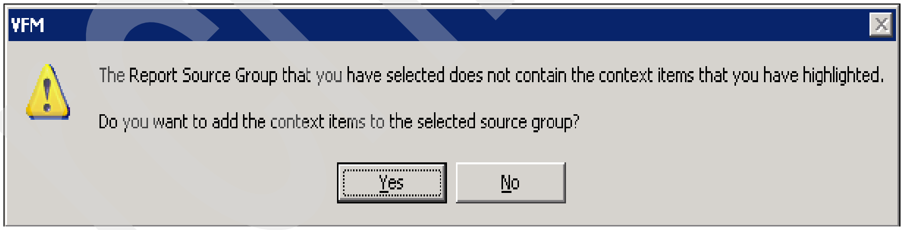

The Report Source Group that you have selected does not contain the context items that you have highlighted.

Do you want to add the context items to the selected source group?

Figure 13-35 VFM pop-up message

10. On the Report Name panel, enter the name that best applies (in this example, the name we enter is file_ages). Click **Enable report archiving** if appropriate and click **Finish** (see Figure 13-36).

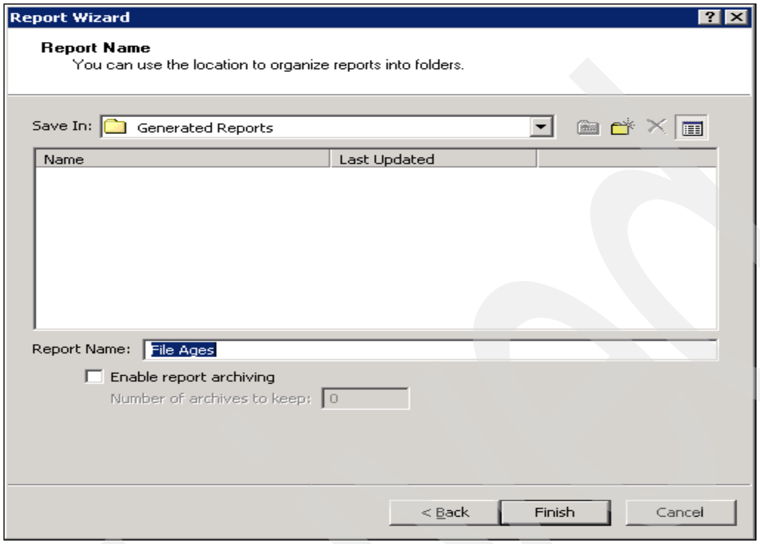

Report Name

You can use the location to organize reports into folders.

Save In: Generated Reports

Name	Last Updated

Report Name: File Ages

☐ Enable report archiving

Number of archives to keep: 0

Figure 13-36 Report Name panel

11. To visualize the report, access the VFM main window and expand **Reports** → **file_ages** and click **File_Ages** as shown in Figure 13-37.

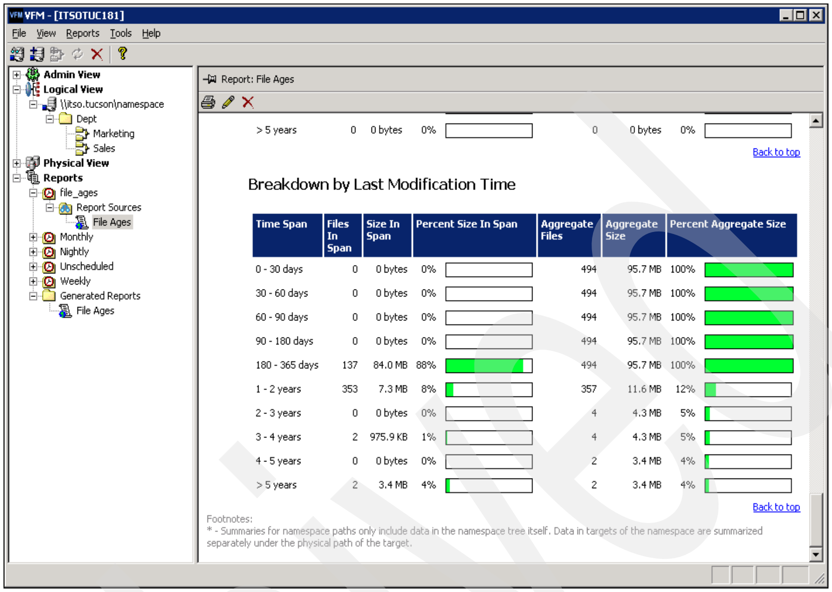

> 5 years	0	0 bytes	0%	0	0 bytes	0%

Breakdown by Last Modification Time

Time Span	Files In Span	Size In Span	Percent Size In Span	Aggregate Files	Aggregate Size	Percent Aggregate Size
0 - 30 days	0	0 bytes	0%	494	95.7 MB	100%
30 - 60 days	0	0 bytes	0%	494	95.7 MB	100%
60 - 90 days	0	0 bytes	0%	494	95.7 MB	100%
90 - 180 days	0	0 bytes	0%	494	95.7 MB	100%
180 - 365 days	137	84.0 MB	88%	494	95.7 MB	100%
1 - 2 years	353	7.3 MB	8%	357	11.6 MB	12%
2 - 3 years	0	0 bytes	0%	4	4.3 MB	5%
3 - 4 years	2	975.9 KB	1%	4	4.3 MB	5%
4 - 5 years	0	0 bytes	0%	2	3.4 MB	4%
> 5 years	2	3.4 MB	4%	2	3.4 MB	4%

Footnotes:
* - Summaries for namespace paths only include data in the namespace tree itself. Data in targets of the namespace are summarized separately under the physical path of the target.

Figure 13-37 File ages report visualization panel

12. Move the scroll bar to view all the report contents.

14

VFM best practices and troubleshooting

This chapter provides troubleshooting information, including best practices when using VFM.

14.1 Troubleshooting

The more common problems you can encounter and their resolutions are provided in this section.

14.1.1 General problems

This section provides information to enable you to resolve general problems.

- You can take the following actions when the response from VFM is slow or it 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 Microsoft Windows tools, such as Windows Explorer.
  - Check which level of status monitoring is enabled (root, link, storage) and the frequency of monitoring taking place.
  - Enabling the Show drives in physical views option causes VFM to gather additional information and can slow its operation, especially when target machines are not reachable from the VFM console.
- If you experience `access denied` errors while setting VFM system options, check the permissions on the account running the VFM client. This account must have permissions to write to the directory tree All Users\Application Data\IBM\VFM\GuiConsole. Write permission for these directories is typically granted to members of the Local Administrators group and the local power users group.
- If you cannot drag-and-drop items out of Microsoft Windows Explorer onto VFM, be sure the items you want to drag to VFM are 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\).
- If you encounter errors creating logical folders in your namespace, check that the permissions on the DFS root share grant the account running the VFM console the ability to make changes to the share.

  Microsoft 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 permission to write to the DFS root share to create logical directories in the namespace.

- If you changed your DFS configuration but client machines still access the old locations, be aware that clients cache referrals for a period of time. You can flush the client cache on Microsoft Windows 2000 and later by performing one of the following actions:
  - Use the DFS tab of the Windows Explorer Property dialog box to clear the history.
  - Run dfsutil with the appropriate options on the client machine.
  - You must reboot clients running older operating systems, such as Microsoft NT 4, Windows 98, or Windows 95, to clear the cache.
- If the connection was refused by the VFM Server service, consider these suggestions:
  - The VFM Server service might have been stopped through the Microsoft 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.
- If the VFM client reports an `access denied` error while attempting to connect to the VFM server, check or obtain access to the VFM server for your VFS service account.

14.1.2 DFS root problems

This section suggests resolutions to DFS root problems.

- When experiencing DFS root problems, you might not see your roots in the VFM browse dialog box. You might see an `Unable to display the DFS Root` <*host_name*>, a `the network resource was not found`, or an `Unavailable` status in the VFM client, and the root error status messages. If so, here are actions you can take to resolve this issue:
  - Check whether network browsing is working correctly in Microsoft Windows Explorer and correct any browsing problems.
  - Enter the root name directly rather than browsing for it.
  - Check whether the DFS service is started on the machines hosting the roots.

- You can take these actions if you cannot create a DFS root using VFM or you cannot create links on a root, and if you encounter one of the following messages:
  - 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`

  Perform the following checks:

  - Network connectivity by pinging the server the root is created on from the VFM machine.
  - That the account running VFM has appropriate privileges for the domain and servers where the domain-based root is being created or for the server where the standalone root is being created.
  - Whether forward and reverse DNS lookups are accurate for the machine on which you want to create the root.
  - Whether you are trying to use a share to create the root that contains a root.
  - Whether you are trying to use a folder to create the root that is a subdirectory of a share that contains a root.
  - Whether the DFS service is started on the machine where you want to create the root.
  - Whether the DFS service is started on the machine hosting the root where you want to create a link.
- If you cannot add links to a DFS root, perform the following actions:
  - 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 Microsoft Windows Server 2003 systems.
- If you have problems creating a domain-based root with replicas; if after the domain-based root is created, it appears as a single replica root; or if the

status of the root is displayed as inaccessible, these considerations might apply or the suggested actions might resolve the issue:

  - 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 have to refer to the domain-based root by its fully qualified domain name when accessing it if 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 Microsoft Windows Server 2003 systems.

- The error message `The request is not supported` when attempting to create a DFS root typically occurs when attempting to create a second DFS root on Microsoft Windows 2000 or Windows Server 2003 Standard Edition server. Windows 2000 and Windows Server 2003 Standard Edition support only one DFS root per computer. Windows Server 2003 Enterprise Server and Windows Server 2003 Datacenter support multiple DFS roots per server.

14.1.3 Namespace problems

This section provides information about resolving namespace issues.

- If 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.
- If 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 DFS root (domain roots only).

  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.
- If links were not created when you ran the policy, perform the following checks:

  - 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.

  - Whether the NTFS permissions for a folder hosting a share to which you want the policy to create link contain <*host_name*>\users. Remove the <*host_name*>\users group from the NTFS permissions on the folder and run the policy.
  - 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.

14.1.4 Replication and migration problems

This section provides information to help you resolve replication and migration problems.

- If VFM reports an error when attempting to copy files when you can successfully copy them using Windows Explorer, these considerations and actions apply:
  - 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 Microsoft 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 Microsoft Windows 2000 NTFS to NT4 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. NT4 NTFS does not support sparse files.
  - Check network connectivity between machines hosting the Replication Agent and those hosting the VFM server or Monitoring Agent.
- If VFM Server service or Replication Agent service fail to start, perform the following checks:
  - The Microsoft Windows NT event log for errors related to startup failure.
  - The Service Control Manager to see whether the VFM server or Replication Agent was stopped or the startup of either was disabled.
  - Whether the account used by the service was removed from the domain.

  - Whether the password was changed on the account used by the service.
  - The account and password information used by the service. Validate that this information is correct.
  - Whether the account used by the service has the Logon as service privilege enabled.
  - Whether the VFM server or Replication Agent is installed on the system.

> **Note:** When you install VFM, you are prompted for the account and password for the VFM server. The server caches password information using Microsoft 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.

- If the deployment of Replication Agent fails, perform the following checks:
  - Whether the system being deployed to is up and reachable. You can do so by pinging it as well as using the net view command to view the shares that it publishes in the network.
  - Whether domain trust issues are preventing deployment.
  - Whether network routing and infrastructure issues are preventing network connectivity in both directions between the system hosting the server or the Monitoring Agent and the system that hosts the Replication Agent.
  - 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.
  - Whether the target system is running one of the following platforms: Microsoft Windows 2000, Windows Server 2003, or Windows XP. Agents cannot be pushed to other platforms.

    > **Note:** If the machine on which you want to deploy a Replication Agent is running Microsoft Windows XP SP 2, you must disable the firewall temporarily on the machine on which you want to deploy the Replication Agent and restart the machine.

  - Whether the normal admin$ shares are published on the destination system. Try viewing the admin$ share for the target system using Microsoft Windows Explorer.

- 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 so by using the Microsoft Windows “run as . . .” facility to run Windows Explorer and browse the admin$ share on the server.
- Whether sufficient space is on the destination system to receive the Replication Agent. To be installed, the Replication Agent requires approximately 700 KB.
- Whether the account under which the VFM Server service is running has privileges to install and run as a service on the destination machine.
- Whether any firewalls exist 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). (See Chapter 3, “Preparation” on page 83.)
- Whether any process is already listening on the port the VFM server and Replication Agent services use and thus are preventing VFM from operating.
- 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.
- Whether the system running on the machine where you want to deploy a Replication Agent was upgraded from Microsoft 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.
- Whether the intended destination is a Microsoft Cluster Server cluster. If so, additional steps are required to configure the Replication Agent to run on the cluster. Contact support for this configuration information.
- 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.
- 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.
- Whether you are running a version of VFM that is later than version 5.5. If so, you cannot deploy Replication Agents to Windows NT 4 systems.
- Whether the Replication Agent can communicate with the machine hosting the VFM server using the NETBIOS host name.

- If VFM cannot communicate with Replication Agents, perform the following checks:
  - Whether any firewalls between the system hosting the VFM server or VFM Monitoring Agent services and the intended destination system for the Replication Agents are blocking the port where the Replication Agent service listens (TCP port 6002). (See Chapter 3, “Preparation” on page 83.)
  - 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 the firewall allowing communication.
  - Whether any process is already listening on the port the VFM server and Replication Agent services use, thus preventing VFM from operating.
  - Whether the system hosting the Replication Agent can ping the system hosting the VFM server using the server name.
  - Whether the Replication Agent can communicate with the machine hosting the VFM server using the NETBIOS host name.
- If VFM does not display complete information about shares hosted on the filer, depending on the N series implementation, VFM might not support the detailed level of information provided by Microsoft Windows systems. VFM displays as much information as can be obtained about storage targets.
- If invalid security descriptor messages are reported during migration operations, the following actions or possibilities apply:
  - 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 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 Microsoft 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 N series, check the filer’s /etc/messages file for any error messages indicating problems communicating with its domain controller on the N series.
  - The files might be 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.

- If replication or migration operations produce `Disk is full` or `The system cannot find the file specified` error messages, perform the following checks:
  - 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 Microsoft Windows systems or qtree quotas in filers.
  - On N series storage system, check the /etc/messages file of the N series 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 large number of files.
  - If the destination is hosted on a 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 through 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 filer, check whether the destination volume has sufficient nodes available to accommodate the files being transferred.

- If 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`

  then:

- 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 through Microsoft 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 the Local Administrators group of the source or destination server.
- If replication or migration operations produce `Cannot delete source/destination`, `It is being used by another process`, or `Cannot open source/destination` error messages, perform the following checks:
  - 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 Microsoft Outlook® data files such as PSTs and database files used by Microsoft Access, SQLServer, FoxPro, or MS® Exchange.
  - 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 might lead to the errors cited earlier.
  - On N series, 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.
- If replication or 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`

You can diagnose permission problems by checking or verifying whether the following conditions are true for your environment:

  - The VFM Server service account is in the Local Administrators group of the source server.
  - The VFM Server service account is in the Local Administrators group of the destination server.
  - The VFM Server service account is not explicitly denied share permissions on the source or destination share.
  - If the source or destination is a filer, verify that the system can communicate with domain controllers. You can do so by using the command **testdc**.
  - The filer's /etc/messages file contains errors related to communication with domain controllers.
  - The account used by the Replication Agent is a member of the Local Administrators and backup operators.

- If VFM reports `access denied` errors when attempting to transfer files to a filer and vscan is in use on the filer, then:
  - 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.
  - Check whether either of these conditions might be causing problems for virus scanning software: changes to the account under which the virus scanning software runs or network connectivity problems between the filer and the virus scanning host.
  - Check whether files and directories on the source and destination have different permission settings. If so perform the following checks:
    - Inherited permissions on the source and destination directory trees.
    - 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 ACLed destination directories. After VFM has completed its work in a directory tree, the permissions are set to the actual permissions intended for the destination. If the transfer is interrupted, it is possible that extra permissions are present on destination directories because VFM was prevented from completing the permission-setting phase of the transfer.
    - Whether local trustee processing is enabled for the policy. Right-click the policy, click **Properties** and click **Replication Security**.

> **Note:** Local trustee processing is not applicable when the source of a transfer is a Domain Controller. Domain Controllers do not host local SIDs.

- If you have problems transferring data using the NFS protocol, then:
  - 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 on to the machine where the UNIX replication is installed and ping the NFS server.
    - In a mixed Microsoft Windows and UNIX environment, always use the fully qualified DNS name.
  - Log on 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 displays the rights you expect.
  - Check the exports file on the NFS server you are trying to contact. This file is found in the following directories:
    - Linux: /etc/exports
    - Solaris: /etc/dfs/dfstab
    - IBM: <*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, the VFM machine 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 “VFMmachine(rw)”.

      > **Note:** Other rights can be within the parentheses; however, it is important that “rw” (read/write) access is granted.

    - Solaris and IBM: Entries are separated by commas. Search for the entry “rw”, “rw=VFMmachine”, or “root=VFMmachine”.

      > **Note:** Make sure that “ro” or “ro=VFMmachine” 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 reports data transfer failures.

  To suppress these error reports, perform the following actions

  i. Right-click the policy that contains the UNIX machine or filer. and click **Properties**.

  ii. Click **Security Replication**.

  iii. Select **Allow loss of security information**.

  iv. 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 IBM: When the VFM machine is not given explicit root access, the entry “anon=0” must be included in the exports file.

  > **Note:** After changing the exports file, run the commands: **exportfs -ua** and **exportfs -a**.

- – 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 N series, 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:

  i. Verify that the UNIX user is mapped to an NTFS user in the file /etc/usermap.cfg.

  ii. Select **Allow loss of security information** in the properties of the policy. This suppresses error messages reported by the UNIX Replication Agent.

- ► If 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.

- If replication or migration operations produce the following error messages:
  - `Error copying source.`
  - `There is not enough space on the disk.`

  Check whether sufficient space is on the destination.

14.1.5 Report problems

If you experience errors attempting to run VFM reports, be sure the account running the VFM client has permissions to write to the directory tree All Users\Application Data\IBM\VFM\VFM Client. Write permission for these directories is typically granted to members of the Local Administrators group and the local power users group.

If reports are slow to generate, perform these actions:

- Check the items under "14.1.1, "General problems" on page 508".
- 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.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 524. Note that some of the documents referenced here might be available in softcopy only.

- *IBM N Series Storage Systems in a Microsoft Windows Environment*, REDP-4083
- *Multiprotocol Data Access with IBM System Storage N series*, REDP-4176
- *Data Protection Strategies in IBM System Storage N Series*, SG24-7591
- *IBM System Storage N series File System Design for an NFS File Server*, REDP-4086

Other publications

These publications are also relevant as further information sources:

- *IBM System Storage N series Data ONTAP 7.3 System Administration Guide*, GC52-1279-02:

  http://www-01.ibm.com/support/docview.wss?uid=ssg1S7002470
- *IBM System Storage N series Data ONTAP 7.3 Network Management Guide*, GC52-1280-02:

  http://www-01.ibm.com/support/docview.wss?uid=ssg1S7002471
- *IBM System Storage N series Data ONTAP 7.3 Software Setup Guide*, GC27-2206:

  http://www-01.ibm.com/support/docview.wss?uid=ssg1S7002465

Online resources

These Web sites are also relevant as further information sources:

- IBM System Storage N series Data ONTAP 7.3RC1 Filer Publication Matrix

  `http://www-1.ibm.com/support/docview.wss?rs=1147&uid=ssg1S7002181`

- Support for Data ONTAP

  `https://www-304.ibm.com/systems/support/myview/supportsite.wss/supportresources?taskind=7&brandind=5000029&familyind=5329797&typeind=0&modelind=0&osind=0`

- Support for Network attached storage (NAS) & iSCSI

  `https://www-304.ibm.com/systems/support/supportsite.wss/mainselect?brandind=5000029&familyind=0&oldbrand=5000029&oldfamily=0&oldtype=0&taskind=1&psid=bm&continue.x=13&continue.y=8`

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Redbooks

IBM System Storage and Virtual File Manager

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM System Storage and Virtual File Manager

Life cycle management with VFM and N series

Intelligent data movement with VFM and N series

Data replication using VFM and N series

This IBM Redbooks publication describes how IBM System Storage N series Virtual File Manager (VFM) logically aggregates user file data distributed across heterogeneous environments and provides administrators with tools and policies to automate data management.

VFM is designed to provide data management functionality for server and storage consolidation, migration, remote office data management, and disaster recovery while avoiding disruption to users. VFM provides this functionality through automated policy-based data management leveraging a global namespace.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7597-00 ISBN 0738431575