



NetworkAppliance®

The evolution of storage.™

VFM™ (Virtual File Manager) 5.6

Advanced Configuration and Techniques

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

Copyright © 2003-2005 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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

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

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

This product includes software developed by the Apache Software Foundation.

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

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

Trademark information

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

NuView Incorporated, 738 Highway 6 South, Suite 850, Houston, Texas 77079
(www.nuview.com).

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

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

Table of Contents

PREFACE	III
CHAPTER 1 INTRODUCTION.....	1
CONFIGURING PERMISSIONS.....	1
<i>To configure access permissions to the VFM server:</i>	<i>1</i>
<i>To configure system access permissions on domain-based roots and shared folders:</i>	<i>2</i>
<i>To configure function-level access permissions:</i>	<i>2</i>
Security permissions.....	3
CHANGE THE VFM SERVER CREDENTIALS.....	3
CHANGE THE LISTENING PORT OF THE VFM SERVER	4
CHAPTER 2 MANAGE REPLICATION AGENTS	5
PREREQUISITES	5
UNIX.....	5
Windows.....	6
HOW TO DEPLOY AND MANAGE REPLICATION AGENTS.....	6
Distributed replication.....	7
Manually deploy a replication agent.....	11
Stop and restart an agent.....	11
Upgrade an agent	12
Uninstall a replication agent.....	12
Change the credentials of an agent	13
Change the listening port of an agent.....	13
Purge an agent.....	13
Install a UNIX replication agent	14
Specify default replication agents.....	18
Specify a proxy machine	18
View a list of replication agents	19
CHAPTER 3 MANAGE MONITORING AGENTS.....	21
HOW TO INSTALL AND MANAGE MONITORING AGENTS	21
Install a monitoring agent.....	22
Assign a monitoring agent to a policy	22
Back up a monitoring agent.....	22
Change the credentials of a monitoring agent.....	23
View the status of a monitoring agent.....	23
CHAPTER 4 RUN BATCH FILES	25
RUNNING BATCH FILES	25
Batch files	26
CHAPTER 5 APPLICATION SERVER AVAILABILITY	33
CONFIGURE PRIMARY AND STANDBY SERVERS	34
Manage primary and standby VFM servers	36
Configure a standby server.....	38
Configure synchronization.....	38
Fail over to standby server	39
Fail back to the primary server	39
Synchronize servers immediately.....	40
Change the synchronization schedule.....	40
View information about servers.....	41
Remove relationship between primary and standby servers.....	41
Remove a standby server	42

Advanced Configuration and Techniques

<i>Remove NETBIOS alias from a server</i>	42
EXPORT/IMPORT CONFIGURATION FILES	42
<i>Export VFM configuration files</i>	44
<i>Import VFM configuration files</i>	45
<i>Resume running VFM on the primary machine</i>	46
<i>Resume running VFM on the backup machine</i>	46
CHAPTER 6 USE SSH SHELL CREDENTIALS	47
HOW TO CONFIGURE VFM TO USE SSH	48
<i>Generate a key-pair</i>	51
<i>Copy the public key to a UNIX machine</i>	52
<i>Copy the public key to a filer using NFS</i>	53
<i>Copy the public key to a filer using CIFS</i>	54
<i>Cache the keys</i>	54
<i>Install and configure OpenSSH on Solaris</i>	55
<i>Install and configure OpenSSH on RedHat</i>	57
<i>Set default remote shell credentials</i>	58
CHAPTER 7 USE RSH SHELL CREDENTIALS.....	59
HOW TO ENABLE RSH AS ROOT	59
<i>Enable RSH as root on a Linux system</i>	59
<i>Enable RSH as root on a Solaris system</i>	61
APPENDIX	63
GLOSSARY	65
INDEX	69

Preface

What the guide covers

This guide describes advanced topics, such as how to configure permissions on the VFM™ (Virtual File Manager) server and manage replication and monitoring agents.

The guide is divided into the following sections.

- **Chapter 1 Introduction** – Explains how to evaluate your storage configuration when you install VFM and how to set permissions on the VFM server.
- **Chapter 2 Manage Replication Agents** – Explains how to deploy and manage replication agents on Windows and non-Windows systems.
- **Chapter 3 Manage Monitoring Agents** – Explains how to install and manage monitoring agents.
- **Chapter 4 Run Batch Files** – Explains how to use batch files to customize policies and reports. The chapter provides a list of batch files supplied by VFM that you can use as examples to create scripts.
- **Chapter 5 Application Server Availability** – Explains how to make the VFM server highly available.
- **Chapter 6 Use SSH Shell Credentials** – Explains how to configure VFM to use SSH.
- **Chapter 7 Use RSH Shell Credentials** – Explains how to enable RSH as root on Linux and Solaris systems.
- **Appendix** – Provides a list and description of the VFM guides.
- **Glossary** – Provides definitions of terms used in this guide.

Audience

This guide is for network administrators who are familiar with Windows NT®, Windows® 2000, Windows Server™ 2003, and DFS. It does not cover basic system or network administration topics.

Typographic conventions

- **Bold text** - Items on a dialog box or page on the VFM console that you are instructed to click.
- *<text in italics>* - Substitute the text within the greater than and less than symbols with the value requested.

Using the online help

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

Advanced Configuration and Techniques

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

- In the table of contents, a new topic is represented as a page with a red asterisk , and an entire section of new topics is represented as a book with a red asterisk .
- Definitions of new terms are available wherever they appear in help topics. The first time a new term appears in a help topic, it appears as a [green hyperlink](#).
- Context help contains detailed information about a dialog box or view, and might include links to related topics. Press **F1** or click  **Help**.

Support

Contact NetApp technical support, 24 x 7.

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

Support phone numbers:

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

Notes:

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

Enter new license key

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

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

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

Chapter 1 Introduction

This guide includes advanced topics, such as setting permissions on the VFM server, managing replication and monitoring agents, running batch files in storage policies, and making the VFM server highly available.

Topics in this chapter are arranged in the following sections:

- **Configuring permissions** explains how to configure access permissions on the VFM server, domain-based roots and shared folders. It also explains how to restrict access at the operation or function level.
- **Change the VFM server credentials** explains how you can update the credentials of the server when you change the password on the account where the VFM server is installed.
- **Change the listening port of the VFM server** explains how you can change the listening port normally used by the VFM server.

Configuring permissions

- Permissions settings, such as lack of administrator permission on the hosts supporting roots or the domains containing domain-based roots, can cause problems when you attempt to administer these roots.
- The account running the VFM client must have permissions to write to the *All Users\Application Data\NuView\StorageX\GuiConsole* directory tree. Users whose accounts lack write permissions to these directories will experience errors when attempting to run VFM reports or set VFM system options.
- Lack of sufficient privileges can prevent policies from functioning correctly.
- Make sure that you have administrator permission for any servers hosting stand-alone roots and domains and servers hosting domain-based roots before you attempt to administer them using either VFM or native tools.
- These problems might be manifested as an inability to create, modify, or delete roots or links in the namespace to which the account running VFM lacks appropriate permissions.

To configure access permissions to the VFM server:

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

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

To configure system access permissions on domain-based roots and shared folders:

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

To configure function-level access permissions:

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

- All policies
- A policy category, such as Disaster Recovery Policies
- A policy type, such as Client Recovery Policies, Filer Recovery Policies, and Server Recovery Policies
- To configure application-wide permissions, click **Options** on the Tools menu, click **System Options**, and then click **Set Security**.
- To configure permissions for a DFS root, right-click the root and click **Properties**, and then click the **Application Security** tab.
- To configure permissions for all policies, right-click  **Admin View** and click **Application Security**.
- To configure permissions for a policy category or type, right-click the item and click **Application Security**.

Notes:

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

Security permissions

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

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

Change the VFM server credentials

During the initial installation process of VFM, you specify the account and password for the VFM server or monitoring agent. This account information is used when these servers run, as well as when they deploy replication agents to transfer data. When you change the password on the account where the VFM server is installed, you must update the credentials of the server. If you do not update credentials, the server ceases to function correctly because Windows security prevents the service from authenticating with the system.

Notes:

- When you change the credentials of the VFM server, you must restart the server service so that the changes can take place. Perform the following actions:
 1. Close the VFM client.
 2. Right-click **My Computer** on the Windows desktop and click **Manage**.
 3. Double-click **Services and Applications**.
 4. Double-click **Services**.
 5. Scroll down to the StorageX Server, right-click the server and click **Stop**.
 6. When the server stops, right-click the server and click **Start**.
- After updating the credentials of the VFM server, you must update the password of all the replication deployed by the server. See "Change credentials of an agent."

Advanced Configuration and Techniques

1. On the **Tools** menu, click **Agent Management**.
2. Select the VFM server and click **Set Credentials**.
3. Type the new password and click **OK**.
4. Click the action in the bottom pane and click **Commit Changes**.

Change the listening port of the VFM server

Normally, VFM uses the TCP port 6001 for the VFM server and monitoring agents and port 6002 for replication agents. You have the option of changing these ports. Make sure that the port you select enables the server to communicate through firewalls with all the replication agents the server deployed and the VFM client .

1. On the **Tools** menu, click **Agent Management**.
2. Right-click the VFM server and click **Set Port**.
3. Type the new TCP port and click **OK**.
4. Select the action to change the port in the bottom pane and click **Commit Changes**.

Chapter 2 Manage replication agents

This chapter explains how you can manage Windows and UNIX replication agents that VFM uses to transfer CIFS and NFS data.

Prerequisites

UNIX

General requirements

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

Solaris 8

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

- SUNWcsl
- SUNWlibms
- SUNWcsu
- SUNWloc
- SUNWgzip

RedHat 9

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

- glibc-2.3.2-27.9.7
- bash-2.05b-20.1
- coreutils-4.5.3-19.0.2

Advanced Configuration and Techniques

- glibc-common-2.3.2-27.9.7
- glibc-2.3.2-27.9.7
- gzip-1.3.3-9
- grep-2.5.1-7.8
- procps-2.0.11-6
- Redhat package manager version 4 or greater.

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

RedHat Enterprise 3

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

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

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

Windows

- The Windows machine where you deploy a replication agent must be running Windows 2000, Windows Server 2003, Windows XP file systems, or Windows NT 4 systems where MSXML3 is installed.
- The account used by VFM to deploy a Windows replication agent must have administrator privileges on the machine that receives the replication agent.

How to deploy and manage replication agents

The following topics are included in this section:

- See "Distributed replication" for an explanation and illustration about how distributed replication works.
- See "Manually deploy a replication agent" for instructions about how to deploy replication agents to transfer data.
- See "Start or stop an agent" for instructions about how to start, stop, and pause a replication agent.

- See "Upgrade an agent" for instructions about how to upgrade a replication to a new version.
- See "Uninstall an agent" for instructions about how to uninstall a replication agent.
- See "Purge an agent" for instructions about how to remove a replication or monitoring agent from the VFM database.
- See "Change the credentials of an agent" for instructions about how to change the credentials of a replication or monitoring agent and the VFM server.
- See "Change the listening port of an agent" for instructions about how to change the listening port that VFM normally uses.
- See "Install a replication agent" for instructions about how to install a replication agent on a UNIX machine.
- See "Specify default replication agents" for instructions about how to specify default replication agents for Windows and UNIX machines to use when VFM cannot deploy a replication agent onto the source or destination.
- See "Specify a proxy machine" for instructions about how to specify a proxy machine.
- See "Change service password" for instructions about how to update the account information for the VFM server and monitoring agents used for future deployment of replication agents.
- See "View a list of replication agents" for instructions about how to view a list of all the replication agents deployed.


Distributed replication

The following list illustrates the algorithm used by VFM to deploy replication agents.


- Windows replication agents are deployed to machines running Windows 2000, Windows Server 2003, Windows XP file systems, or Windows NT 4 systems where MSXML3 is installed.
- UNIX replication agents are deployed to machines running Solaris 8, RedHat 9, or RedHat Enterprise 3 operating systems.

Advanced Configuration and Techniques


1. When you move data, the replication agent is deployed to the machine hosting the destination target.




2. If VFM cannot deploy the agent to this machine, it attempts to deploy the agent to the machine hosting the source target.




3. If VFM cannot deploy the agent to either the destination or source target, it attempts to use the agent on the proxy machine of the destination target, if one is configured.



4. If VFM cannot deploy the agent to the proxy machine of the destination target, it attempts to use the agent on the proxy machine of the source target, if one is configured.



5. If VFM cannot use any of the agents listed above, VFM uses the replication agent on the VFM server or the monitoring agent to transfer data.



- You can set the data transfer rate limits for machines that participate in data transfers by using the properties dialog box for the machine in the VFM Physical View.
- The replication agent is deployed as part of VFM policies and file differences reports.
- The account used by VFM to deploy a replication agent needs to have administrator privileges (Windows replication agents) or root privileges (UNIX replication agents) on the machine that receives the replication agent. Windows replication agents are deployed through the machine's ADMIN\$ share and are installed as a service. Both of these operations require administrator privileges on the system where they take place.
- If you change the password on the account where the VFM service is installed, the Service account password and the password of all deployed replication agents and monitoring agents must be updated. **If you do not update passwords, already deployed**

Advanced Configuration and Techniques

agents cannot perform replication and monitoring. See "Change the credentials of an agent" for instructions about how to update credentials.

- When a replication operation uses NTFS Change Journal, the replication agent is deployed to the machine hosting the source target. NTFS Change Journal is only applicable to Windows file systems.
- (CIFS only) When data is replicated using differential replication, VFM deploys a replication agent to both the source and destination. If VFM cannot deploy a replication agent to either the source or destination, you must designate a proxy for that location.
- (CIFS only) If the machine you want to deploy a replication agent is running Windows XP SP 2, you must perform the following actions:
 1. Disable the firewall on the machine you want to deploy the replication agent.
 2. Restart the machine.

Note: The following tables list the network ports used by DFS and VFM.

Network Ports Used by DFS

Service Name	UDP	TCP
NetBIOS Datagram Service	138	
NetBIOS Session Service		139
LDAP Server	389	389
Remote Procedure Call (RPC)		135
Server Message Block (SMB)	445	445

Network Ports Used by VFM

Service Name	UDP	TCP
StorageX Server		6001
Replication Agent		6002

Reasons to use a proxy machine for a replication agent

VFM enables you to specify a proxy machine for a replication agent that is used during data transfer when a replication agent cannot be deployed on the source or destination.

- Use a proxy machine to avoid transferring data across a WAN.
- When the VFM or is not on the same LAN as the source and destination targets and you cannot deploy the replication agent to the destination or the source target, assign a proxy machine. If the replication agent on the VFM server was used in this scenario, the data

transfer path would be from the source target to the VFM server and from the VFM server to the destination target. By deploying a proxy replication agent, you can avoid transferring the data across the WAN during the transfer operation.

Manually deploy a replication agent

VFM automatically deploys Windows and UNIX replication agents to manage data transfers. You can manually deploy agents. See "Manage replication agents" for system requirements.

- You must deploy a replication agent on the source or destination when transferring data.
- VFM enables you to specify default proxy replication agents. When VFM cannot deploy an agent to the source or destination, it uses the proxy replication agent you specified. See "Specify default replication agents" for instructions about how to set these options.

To deploy a replication agent:

1. On the **Tools** menu, click **Agent Management**.
2. Click **Deploy**.
3. In the **Target** field, type the name or UNC path of the machine where you want to deploy a replication agent. VFM uses this field to also determine whether to deploy a Windows or UNIX replication agent.
4. Type the credentials that VFM will use to communicate with the replication agent.
5. Optionally type a different port. VFM normally uses port 6002 to communicate with replication agents.
6. Click **OK**. The action appears in the bottom pane.
7. Select the action to deploy the agent and click **Commit Changes**.

Note: If you change the listening port, make sure that the port you select enables the VFM server to communicate through any firewalls with the VFM client and replication agents the server deployed.

Stop and restart an agent

- When you **stop** a replication agent while it is transferring data, the data transfer terminates and starts again at the next scheduled time. Restarting the agent does not **restart** the data transfer.
- When you **pause** a replication agent while it is transferring data, the data transfer is suspended and **resumes** when the agent starts again.

To stop or pause the agent:

1. On the **Tools** menu, click **Agent Management**.

Advanced Configuration and Techniques

2. Right-click a replication agent and click **Stop Agent** or **Pause Agent**.
3. Select the action in the bottom pane and click **Commit Changes**.

To restart or resume the agent

1. Right-click the replication agent and perform one of the following actions:
 - If you stopped the agent, click **Restart Agent**.
 - If you paused the agent, click **Resume Agent**.
2. Select the action in the bottom pane and click **Commit Changes**.

Upgrade an agent

When you upgrade to a new version of VFM, you must upgrade all deployed replication agents and monitoring agents you installed. If the replication agents are not upgraded, the policies that use the agents to perform data transfer will be disabled.

1. Open VFM after you upgrade to a new version.
2. On the **Tools** menu, click **Agent Management**.
3. Select the replication agents and click **Upgrade**.
4. Select the actions to upgrade the agents in the bottom pane and click **Commit Changes**.

Notes:

- You must use the VFM media to upgrade installed monitoring agents.
- If you get the error message "Network Resource Not Found" when you try to upgrade the replication agent installed on the same machine as the VFM server, see "General Problems" in the Troubleshooting topic.

Uninstall a replication agent

When VFM or a monitoring agent is uninstalled, replication agents deployed by the VFM server or monitoring agent are not deleted. VFM enables you to uninstall one or multiple agents.

1. On the **Tools** menu, click **Agent Management**.
2. Select one or multiple replication agents and click **Uninstall**.
3. Select one of the following options:
 - Uninstall with existing credentials
 - Uninstall with the following credentials. Use this option when the current credentials do not have the security permissions to uninstall the agent.
4. Click **OK**.
5. Select the actions to uninstall the agents in the bottom pane and click **Commit Changes**.

Change the credentials of an agent

When you change the credentials of the VFM server, you must update the credentials of all the replication agents that the server deployed.

Note: If you installed a monitoring agent and you change the credentials of the agent, you must update the credentials of all the replication agents that the monitoring agent deployed.

1. On the **Tools** menu, click **Agent Management**.
2. Select all the replication agents deployed by the VFM server and click **Set Credentials**.
3. Change the credentials to match the credentials of the VFM server.
4. Select the action in the bottom pane and click **Commit Changes**.

Change the listening port of an agent

Normally, VFM uses the TCP port 6002 for replication agents and TCP port 6001 for the VFM server. You have the option of changing the ports. Make sure that the port you select enables the replication agent to communicate with the VFM server through any firewalls.

1. On the **Tools** menu, click **Agent Management**.
2. Perform one of the following actions:
 - To update the port for multiple agents, select the agents and click **Set Listening Port**.
 - To update the port for one agent, right-click the agent and click **Set Port**.
3. Type the new port and click **OK**.
4. Select the action to change the port in the bottom pane and click **Commit Changes**.

Purge an agent

Remove a replication or monitoring agent from the VFM database. Use this option when you cannot uninstall an agent because the machine hosting the agent does not exist or was renamed.

1. On the **Tools** menu, click **Agent Management**.
2. Select the agent you want to remove from the list and click **Purge**.
3. Select the action to purge the agent in the bottom pane and click **Commit Changes**. The agent is removed from the database and from the list of agents.

Install a UNIX replication agent

VFM automatically deploys a replication agent on the source or destination when transferring NFS data. If VFM cannot deploy an agent on either the source or destination, deploy or install an agent on a proxy machine. The UNIX machine where you install a replication agent must be running RedHat® 9, RedHat Enterprise 3, or Solaris® 8 operating system.

Notes:

- You must log in with root privileges on the machine where you install the replication agent.
- NFS 3 or later must be installed on the source and destination. See the illustration of NFS data transfer at the end of this topic.
- The UNIX operating systems are case sensitive.
- The file system where the replication agent will be installed must be mounted read-write. If you attempt to install the replication agent on a read-only file system, the following error message appears:
 - On Solaris:
 - pkgadd: ERROR: Unable to make directory </usr/local/URA>
 - On Linux:
 - error: unpacking of archive failed on file /usr/local/ura/bin:
 - cpio: mkdir failed - No such file or directory
- See "Manually deploy a replication agent" for instruction about how to deploy Windows and UNIX replication agents using VFM.

Prerequisites

Before you install a UNIX replication agent, review the list of requirements at the beginning of this chapter.

RedHat 9 and RedHat Enterprise 3

Use these instructions to install and manage replication agents on RedHat systems.

To install a replication agent:

1. To get a copy of the installation package, perform one of the following actions:
 - In the VFM installation directory, click the *Agents\UNIX Replication Agents\Linux\RedHat* folder.
 - Extract the installation package from the VFMInstall.zip or VFMInstall.exe file.
2. Copy the installation package `vmra-5.5-<build number>.i386.rpm`, where *<build number>* is the current build of the package, to the UNIX machine that will host the replication agent.
3. **View details of the installation package.** In these instructions, build number "1" is used. To view details of the installation package without installing, type the following command:

```
rpm -qilp vmra-5.5-1.i386.rpm
```

4. **Install the replication agent.** You must use the option `--prefix` to specify the full directory path where the replication agent will be installed. This determines where the binaries and log files of the replication agent are kept. In the following example, the binaries are installed in `/usr/local/URA/bin` and the logs and registry are installed in `/usr/local/URA/logs`.

```
rpm -ihv --prefix=/usr/local/URA vfmra-5.5-1.i386.rpm
```

5. **Configure the replication agent.** After you install the replication agent, you must configure it. The following example shows how to configure the agent.

```
/etc/init.d/ura configure
```

You are prompted to enter the following information:

- The TCP port where the replication agent and the VFM server listen.
- The name of your VFM servers.

Notes:

- You can type the name or IP address of the VFM server.
- You must update the replication agent whenever the information you entered changes. Perform the following actions:
 1. Stop the replication agent
 2. Configure the agent
 3. Start the agent again.

6. **Start the replication agent.** Type the following command:

```
/etc/init.d/ura start
```

To verify that the replication agent is running, type the following command:

```
pgrep -fl Repl
```

To stop a replication agent:

Type the following command to stop the replication agent:

```
/etc/init.d/ura stop
```

To remove a replication agent:

1. Stop the replication agent.
2. Type the following command to remove the replication agent:

```
rpm -evv vfmra-5.5-1
```

3. Type the following command to remove the directory `/usr/local/URA`. This directory contains the configuration files created by the agent.

```
rm -Rf /usr/local/URA
```

Advanced Configuration and Techniques

To upgrade a replication agent:

To install a new version of the package, you must uninstall the old package and then install the new package:

1. Complete Steps 1 and 2 under "To remove a replication agent."
2. Follow the instructions under "To install a replication agent."

To show all the packages installed, type the following command: `rpm -qa`

Solaris

VFM deploys UNIX replication agents to the directory `/usr/local/URA`. The following instructions show how to install a replication agent on this directory. To install a replication agent to another directory, see "Installing a UNIX replication to another location."

"The package name is `vmra-5.5-<build number>.pkg`, where `<build number>` is the current build of the package. In these instructions, build number "223" is used.

To install a replication agent:

1. To get a copy of the installation package, perform one of the following actions:
 - In the VFM installation directory, click the Agents\UNIX Replication Agents\Solaris folder.
 - Extract the installation package from the VFMInstall.zip or VFMInstall.exe file.
2. Copy the installation package to the UNIX machine that will host the replication agent.
3. **Unzip the package file.** At a shell prompt, type the following command:

```
gunzip vmra-5.5-223.pkg.gz
```

4. **Install the package.** Type the following command:

```
pkgadd -d ./vmra-5.5-223.pkg
```

The replication agent is installed in the `/usr/local/URA` directory.

5. **Configure the replication agent.** After you install the replication agent, you must configure it. The following example shows how to configure the agent.

```
/etc/init.d/ura configure
```

You are prompted to enter the following information:

- The TCP port where the replication agent and the VFM server listen
- The name of your VFM servers.

Notes:

- You can type the name or IP address of the VFM server.
- You must update the replication agent whenever the information you entered changes. Perform the following actions:
 1. Stop the replication agent
 2. Configure the agent
 3. Start the agent again.

6. **Start the replication agent.** Type the following command:

```
/etc/init.d/ura start
```

To verify that the replication agent is running, type the following command:

```
ps -ef|grep Repl
```

To view packages currently installed:

- To view all packages, type the following command: `pkginfo`
- To view detailed information about a package, type the following command:

```
pkginfo -l NTAPura
```

To search for the replication agent package:

Type the following command: `pkginfo | grep NTAP`

To stop the replication agent:

Type the following command: `/etc/init.d/ura stop`

To remove a replication agent:

1. Stop the replication agent.
2. Type the following command to remove the replication agent:

```
pkgrm NTAPura
```

3. Type the following command to remove the directory `/usr/local/URA`. This directory contains the configuration files created by the agent.

```
rm -Rf /usr/local/URA
```

To upgrade a replication agent:

To install a new version of the package, you must uninstall the old package and then install the new package.

1. Complete Steps 1 and 2 under "To remove a replication agent."
2. Follow the instructions under "To install a replication agent."

Install a UNIX replication agent to another directory on Solaris

VFM automatically deploys UNIX replication agents to the directory `/usr/local/URA`. You can install a replication agent to another directory.

Note: If you install a UNIX replication agent to another directory and VFM upgrades the replication agent, the new replication agent is installed in the default directory `/usr/local/URA`.

1. Complete Steps 1 through 3 under "To install a replication agent."
2. **Install the package.** Type the following command:

Advanced Configuration and Techniques

```
# pkgadd -a none -d ./vmra-5.5-223.pkg
```

You will see the following prompt:

Enter path to package base directory [?,q]

3. **Type the location of the directory**, for example:
 /opt/URA
4. Continue with Step 5 under "To install a replication agent."

Specify default replication agents

You can specify default proxy replication agents for policies that transfer data. VFM enables you to specify a replication agent for Windows and UNIX. When VFM cannot deploy a Windows or UNIX agent onto the source or destination, VFM uses the default agents you specify.

1. On the **Tools** menu, click **Options**.
2. Select **Systems Options**.
3. Under "Default Replication Agents," specify a machine hosting a Windows or UNIX replication agent or a machine that you want VFM to deploy a replication agent during data transfer.
4. Click **OK**.

Specify a proxy machine

Normally, VFM deploys Windows and UNIX replication agents on machines participating in policies that transfer data. You can specify a proxy machine to host the replication agent instead of having the agent deployed to a machine participating in the policy.

Note: If the machine you want to deploy a replication agent is running Windows XP SP 2, you must perform the following actions:

1. Disable the firewall on the machine you want to deploy the replication agent.
2. Restart the machine.

1. In the Physical View tree, right-click a machine and click **Properties**.
2. Click the **Agent** tab.
3. Perform one of the following actions:
 1. Select **Use default proxy**, to use the machine you specified as the default system option.
 2. Select **Use this specified proxy**, to select a machine you want VFM to deploy a replication agent. Type the path of the machine in the appropriate text box or click  **Browse** to locate the machine in your network.
4. Click **OK**.

View a list of replication agents

To view a list of deployed Windows and UNIX replications agents, perform one of the following actions:

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

Chapter 3 Manage monitoring agents

This chapter explains why you should consider distributing monitoring agents and provides instruction about how to install and manage the agents on Windows machines.

Distributed monitoring

- Consider distributing monitoring agents based on your network topology. You can select another monitoring agent that you installed at another location and configure it to work with your VFM server.
- It is recommended that monitoring agents have the same network connectivity as clients for the link targets they monitor.

Assigning monitoring agents to policies

VFM enables you to assign monitoring agents to run policies.

- Uninstalling a VFM monitoring agent disables any policies assigned to that agent.
- Reinstalling the agent on a given host does not automatically reenables policies using the agent on that host. To reenables these policies, you must assign the newly installed agent to run the policies by reselecting the desired agent to be used by the policy through the properties dialog box for the policies.

How to install and manage monitoring agents

VFM enables you to install and manage distributed monitoring agent.

- See "Install a monitoring agent" for instruction about how to install an agent on a machine.
- See "Assign a monitoring agent to a policy" for instruction about how you can assign a monitoring agent to run a policy.
- See "Back up a monitoring agent" for instructions about how to specify where you want to store the snapshot of the status of a monitoring agent and the policies assigned to the agent.
- See "Change the credentials of a monitoring agent" for instructions about how to update credentials.
- See "Purge an agent" in *Chapter 2 Managing Replication Agents* for instructions about how to remove monitoring agents from the VFM database.
- See "View the status of a monitoring agent" for instructions about how to view information about a deployed agent.

Install a monitoring agent

The system requirements for installing and running a monitoring agent are the same as those for VFM. Before you install a monitoring agent, on a server, review the VFM system requirements.

Use the VFM media to install a monitoring agent on a server.

Notes:

- To specify a monitoring agent for a policy, see "Assign a monitoring agent to a policy."
- Before you uninstall a monitoring agent, make sure that policies assigned to the agent are assigned to another monitoring agent or the VFM server because the policies are disabled after the monitoring agent is uninstalled.

Assign a monitoring agent to a policy

The VFM server runs policies. You can select another location to run a policy.

1. Perform one of the following actions:
 - In the Admin View, right-click a policy and click **Properties**.
 - Select a policy in the Admin View to display a **Properties** tab in the right pane.
Click .
2. Click **General**.
3. On the **Use the Monitor on** list, select the host name of the monitoring agent where monitoring is performed. The list displays the host names of the VFM server and all the VFM monitoring agents that are installed and configured to communicate with the VFM server.
3. Click **OK**.

Back up a monitoring agent

When you schedule to back up the application, VFM saves the configuration files of the application and all installed monitoring agents. The configuration files of a monitoring agent are saved in folder under the location you specify for saving the application configuration files. You can specify a different location to store snapshots of the state of a monitoring agent and the policies assigned to the agent.

1. On the **Tools** menu, click **Agent Management**.
2. Select a monitoring agent and click **Set Save Location**.
3. In the **Save location** field, type the path where you want to save the snapshot of the monitoring agent, or click **Browse** to find a location. This should be a location that is different than the location you specified to save the application.
4. Click **OK**.
5. Click the action in the bottom pane and click **Commit Changes**.

Note: For instructions about how to automatically back up the application, see "Export VFM configuration files".

Change the credentials of a monitoring agent

When you change the password on the account where the VFM server or a monitoring agent is installed, you must update the credentials of the VFM server or the monitoring agent. If you do not update credentials, the VFM server or monitoring agent ceases to function correctly because Windows security prevents the service from authentication with the system.

Notes:

- When you upgrade to a newer version of VFM, you must upgrade all installed monitoring agents and deployed replication agents.
- You specify the credentials for the VFM server or monitoring agent during the initial installation process of the product. This account information is used when these servers run, as well as when they deploy replication agents to transfer data.
- When you change the credentials of the VFM server or a monitoring agent, you must update the credentials of the replication agents already deployed by the VFM server or monitoring agent. If you do not update passwords, you cannot perform replication and monitoring on already deployed agents.

1. On the **Tools** menu, click **Agent Management**.
2. Select the VFM server or monitoring agent whose credentials you want to update and click **Set Credentials**.
3. Type the new account name and password and click **OK**.
4. Click the action in the bottom pane and click **Commit Changes**.

View the status of a monitoring agent

After a monitoring agent is installed, you can view status information about the agent. To view system status, perform one of the following actions:

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

Chapter 4 Run batch files

VFM enables you to configure batch files to execute before and after you run a storage policy, run a report, or take a snapshot of the application.

- For information you need to know before you run batch files in VFM, see "Running batch files."
- To view detailed information about the script files you can run in VFM, see "Batch files."

Running batch files

- The batch files are run by VFM on the machine from which the targets of a policy are being monitored by VFM, whether by the VFM server or by a VFM monitoring agent to which the policy is assigned. Batch files related to reports are run by the VFM server.
- VFM uses the directory where a batch file is located as the default/current directory where the command is executed. If a batch file is specified through a UNC path, the TMP directory of the server executing the policy is used as the default/current directory.
- VFM checks the return value from the batch script that is run before failover. If the batch script returns failure, VFM does not continue the failover process. This allows customized checking to be performed before the actual failover takes place.
- If the UNC path to the batch file contains spaces, you must bracket the entire path in quotation marks. You must follow the same rules that apply when running commands from a Windows command prompt.
- VFM provides example scripts in the installation directory under the Examples\Scripts folder. These examples rely on the presence of additional third-party software, such as an RSH client and other utilities, to function. See "Batch files" for information about the scripts.
- Some antivirus programs display a message every time a script is invoked, inquiring whether you want to enable the script to run. VFM server components execute scripts for snapshots taken of the application and batch files configured to run with VFM storage policies. When scripts are invoked by VFM server components, the checks of the antivirus program interfere with the operation of the server. You should configure your antivirus program to enable scripts used by VFM to run unimpeded.
- To set application-wide default settings, click **Options** on the **Tools** menu. Changes affect only new policies. Existing policies are not affected.
- Test the scripts to verify that they are functioning as expected. You must do this by specifying small test directories for the source and destination of the synchronization task and configuring the scripts to write diagnostic output to a text file for review. After the scripts function correctly in the test configuration, you can configure them for use in "real" policies, and you can reduce the amount of diagnostic information logged by the scripts to match the requirements of your environment.

Batch files

Batch files enable you to customize policies, such as send an email, post an event in the Windows event log, or send an SNMP trap in response to a failure. VFM provides some example scripts in the installation directory under the **Examples\Scripts** folder.

The following table provides detailed information about scripts you can run in VFM.

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
Application snapshot	Pre-snapshot script	"Run batch before snapshot" field on the Application Snapshot Options page. Click Options on the Tools menu.	<first user argument>	0 => success all others => error
Application snapshot	Post-snapshot script	"Run batch after snapshot" field on the Application Snapshot Options page. Click Options on the Tools menu.	<first user argument>	0 => success all others => error
Archival migration policies	Pre-scan script	"Run batch before scan" field on the General tab of the policy properties.	1. Policy name 2. Local path to a file to which script can output a list of tasks to be added. The local path is relative to the VFM server or monitoring agent that invokes the script. 3. <first user argument>	0 => success all others => error
Archival migration policies	Post-scan script	"Run batch after scan" field on the General tab of the policy properties.	1. Policy name 2. <first user argument>	0 => success all others => error
Archival migration policies	Criteria script	"Confirmed by user selection criteria script" field on the Migration Criteria tab of the policy properties.	1. Policy name 2. UNC of path to evaluate 3. <first user argument>	0 => path matches criteria 1 => path does not match criteria all others => error
Archival migration policies	Task pre-batch script	"1. Run batch file" field on the Actions tab of the policy properties.	1. UNC path of source 2. UNC path of destination 3. <first user argument>	0 => success all others => error

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
Archival migration policies	Task post-batch script	"9. Run batch file" field on the Actions tab of the policy properties.	- UNC path of source - UNC path of destination - copy=[pass fail skip] - updatedfs=[pass fail skip] - delete=[pass fail skip] <first user argument>	0 => success all others => error
Client disaster recovery	Pre-replication script	"Run Batch Before Replication" field on the Replication Script tab of the policy properties.	- UNC path of link (for Replication Policies, this is the destination link) - UNC path of source - UNC path of destination <first user argument>	0 => success all others => error
Client disaster recovery	Post-replication script	"Run Batch After Replication" field on the Replication Script tab of the policy properties.	- UNC path of link (for Replication Policies, this is the destination link) - UNC path of source - UNC path of destination - Result of the replication: Pass - Operation succeeded Fail - Operation encountered errors Cancel - Operation cancelled by the user Expired - Operation was not completed in the scheduled time window - The time that the it took to perform the operation - Number of files copied - Number of bytes copied - Number of files already in sync - Number of bytes already in sync <first user argument>	0 => success all others => error
Client disaster recovery policies	Pre-failover script	"Run batch before failover" field on the DR Script Settings tab of the policy properties.	- Policy name [not used, will be empty] - UNC path of the link being failed over - Invocation type (/failover or /failback) in case of failback, not provided in case of active target	0 => success all others => error

Advanced Configuration and Techniques

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
			change 5. [first user argument]	
Client disaster recovery policies	Post-failover script	"Run batch after failover" field on the DR Script Settings tab of the policy properties.	1. Policy name 2. UNC path of active target 3. UNC path of the link being failed over 4. Invocation type (/failover or /failback) in case of failback, not provided in case of active target change 5. <first user argument>	0 => success all others => error
Client disaster recovery policies	User decision script	"Run User Decision Script" field on the DR Script Settings tab of the policy properties	1. Policy name 2. Link name 3. Online target 4. Offline targets	0 => success all others => error
Filer disaster recovery policies	Pre-failover script See Client disaster recovery policies script.	See Client disaster recovery policies script.	See Client disaster recovery policies script.	See Client disaster recovery policies script.
Filer disaster recovery policies	Post-failover script See Client disaster recovery policies script.	See Client disaster recovery policies script.	See Client disaster recovery policies script.	See Client disaster recovery policies script.
Filer disaster recovery policies	User decision script	"Run User Decision Script" field on the DR Script Settings tab of the policy properties	1. Policy name 2. Filer name 3. Volume name 4. <first user argument>	0 => success all others => error
Migration policies	Migration task creation script	"Run task creation batch" field on the General tab of the policy properties.	1. Policy name 2. Local path to a file to which script can output a list of tasks to be added. The local path is relative to the VFM server or monitoring agent that invokes the script. 3. <first user argument>	0 => success all others => error
Migration policies	Task pre-batch script See Archival migration policies script.	See Archival migration policies script.	See Archival migration policies script.	See Archival migration policies script.

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
Migration policies	Task post-batch script See Archival migration policies script.	See Archival migration policies script.	See Archival migration policies script.	See Archival migration policies script.
Multiprotocol namespace policies	Pre-sync script	"Run Batch Before Sync" field on the Scripts tab of the policy properties.	1. Policy name 2. Automount files location 3. <i><first user argument></i>	0 => success all others => error
Multiprotocol namespace policies	Post-sync script	"Run Batch After Sync" field on the Scripts tab of the policy properties.	1. Policy name 2. Automount files location 3. <i><first user argument></i>	0 => success all others => error
Replication policies	Pre-replication script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Replication policies	Post-replication script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Report group	Pre-data gathering	"Run batch before gathering data" field on Scripts of the Report Group properties.	1. Policy name 2. <i><first user argument></i>	0 => success all others => error
Report group	Post-data gathering See Pre-data gathering script	"Run batch after gathering data" field on Scripts of the Report Group properties.	See Pre-data gathering script	See Pre-data gathering script
Report publishing option	Before updating index file	"Run batch before updating index file" field on the Tools>Options>Reporting dialog box.	The path to the Index.xml file that VFM creates when publishing reports.	0 => success all others => error
Report publishing option	After updating index file See Before updating index file script	"Run batch after updating index file" field on the Tools>Options>Reporting dialog box.	See Before updating index file script	See Before updating index file script

Advanced Configuration and Techniques

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
Report wizard	Pass report output to script	"Script" field on the Report Output dialog box of the Reports wizard.	- The path to a TMP file containing the report in the selected format. It is expected that the batch file will copy the contents of the file to a new location for subsequent processing. VFM deletes the temporary file after the report output batch file runs. - The data collection process shows the following states: OK - All the data required for the report was successfully collected. Warning - Not all the data required for the report was collected. The information passed to the batch file might not be complete. You need to refresh the associated report group to force collection of the data. Error - Problems occurred while collecting the data for the report. The information passed to the batch file might not be complete. <i><first user argument></i>	0 => success all others => error
Server disaster recovery policies	Pre-replication script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Server disaster recovery policies	Post-replication	See Client disaster recovery	See Client disaster recovery script	See Client disaster

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
	script See Client disaster recovery script	script		recovery script
Server disaster recovery policies	Pre-failover script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Server disaster recovery policies	Post-failover script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Server disaster recovery policies	User decision script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Share Creation wizard	Pre-creation batch file	"Run pre-creation batch file" field when you create duplicate share	1. UNC path of source share 2. UNC path of destination share 3. Local path of the destination share	0 => success all others => error
Share Creation wizard	Post-creation batch file	"Run post-creation batch file" field when you create duplicate share	1. UNC path of source share 2. UNC path of destination share 3. Local path of the destination share	0 => success all others => error
Snapshot scheduling policies	Pre-snapshot script	"Run batch before snapshot" field on the VFM Settings tab of the policy properties.	1. Policy name 2. Filer name 3. Volume name 4. <first user argument>	0 => success all others => error
Snapshot scheduling policies	Post-snapshot script	"Run batch after snapshot" field on the VFM Settings tab of the policy properties.	1. Policy name 2. Filer name 3. Volume name 4. <first user argument>	0 => success all others => error
Storage load balancing policies	Task pre-batch script See Archival migration policies script	"Run Batch File Before Migration" field on the Script tab of the policy properties.	See Archival migration policies script	See Archival migration policies script

Advanced Configuration and Techniques

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
Storage load balancing policies	Task post-batch script See Archival migration policies script	"Run Batch File After Migration" field on the Script tab of the policy properties.	See Archival migration policies script	See Archival migration policies script

Chapter 5 Application server availability

This chapter explains how to make your VFM server highly available. It discusses the methods provided by VFM.

Planning for VFM availability

When you install VFM, consider the following aspects of your storage configuration:

- A convenient location for the VFM console software.
- A reliable/highly available location for the VFM server software.
- How large your VFM configuration is.
- Whether you have storage policies that you want to run from locations that are different from where the VFM server is installed.
- Where the DFS roots that are managed by VFM are located.

Selecting a method

- **Maintain primary and standby VFM servers**

VFM enables you to configure primary and standby VFM servers. The command-line utility enables you to perform fail over and fail back operations and view information about the servers.

- When the primary VFM server is not available, you can execute a command to activate the standby server. This can be scripted when it is appropriate for your environment. This enables you to maintain a hot site.
- The failover operation offers an option to add a NETBIOS alias to the standby server, allowing it to answer NETBIOS requests in place of the failed server.
- For instructions, see "Configure primary and standby servers."

- **Schedule snapshots of the state of the VFM server**

VFM enables you to schedule snapshots of the state of the VFM server.

- When the VFM server is not available, you can import the latest snapshot on another machine where VFM is installed and have it take over as the VFM server.
- Snapshots can be used to restore the VFM server to a known state after unintended changes or disk failure on the system where the server is installed.
- For instructions, see "Export/Import VFM configuration files."

- **Manually export and import VFM server configuration files**

VFM enables you to manually export VFM configuration files.

- When the VFM server and VFM console are not available, you can use the export utility to save the VFM configuration information.
- You can import the configuration files on another machine where VFM is installed and have it take over as the VFM server.
- For instructions, see "Export/Import VFM configuration files."

Notes:

- You must install the VFM server on the standby machine with the same account that you configured the VFM server on the primary machine.
- You might want to create a namespace backup policy to save backups of the logical structure of all the DFS roots in your namespace. This enables you to recreate the DFS namespace if a DFS root server is lost along with the VFM server machine.
- An additional component of any availability plan is maintaining backups of the actual data contained on the servers participating in the DFS namespace.
- VFM does not support VFM servers active on multiple servers managing policies referencing a common set of DFS resources.


Configure primary and standby servers

Use this feature when you want to create a standby server that you can activate when the primary VFM server or a monitoring agent is not available. This section describes how a standby server is configured to stand by for the VFM server or a monitoring agent, how the standby server can take over when the primary server is not available, and how the original primary server can take over when it becomes available again.

Throughout this section, the term *primary server* is used to describe an active VFM server or monitoring agent. The term *standby server* is used to describe a machine that is standing by (not active).

Primary and standby server configuration options

How you configure your primary and standby VFM servers in your network depends on your environment. The following graphic illustrates two network configuration options.



- In **Option 1**, one primary server (the VFM server) is configured to synchronize configuration files with one of the standby servers, and the other two primary servers (monitoring agents) are configured to synchronize their configuration files with the other standby server.
- In **Option 2**, each primary server is configured to synchronize configuration files with a standby server.

Note: It is not recommended that you configure the same standby server for the VFM server and a monitoring agent. The VFM server and a monitoring agent cannot run on the same machine. If both server and agent become unavailable at the same time, you can only activate one of the components if they are sharing the same standby server.

How it works

- Install a monitoring agent to use as a standby server.
- Configure a monitoring agent as the standby server.
- Enable synchronization between the primary and standby server and schedule how often to perform data synchronization.
- When the primary server is not available, manually fail over to the standby server.
- When the primary server becomes available, manually fail back to the primary server.

VFM also enables you to manually synchronize the primary and standby servers, change the schedule, cancel the relationship between the primary and standby servers, change a standby server back to a monitoring agent, and get server state.

Configure primary and standby servers checklist

For a step-by-step explanation for configuring and managing VFM primary and standby servers see "Manage primary and standby servers."

Manage primary and standby VFM servers

The following table illustrates the VFM server availability operations. Use it as a quick reference guide to configure VFM standby servers, fail over to a standby server, and recover back to the original primary server. In this example, *ServerA* is the primary VFM server and *ServerB* is a server you configure as a standby for *ServerA*.

Step	ServerA	ServerB	Comments
1	Install the VFM server For installation instructions, see the VFM Getting Started Guide or Release Notes.	Install a monitoring agent. For instructions, see "Install a monitoring agent."	Installation on ServerB should use either the same credentials used to install ServerA or a set of credentials with the same privileges as those used when installing ServerA.
2		Configure as a standby server. For instructions, see "Configure a standby server."	
3	Add <i>ServerB</i> as a standby to ServerA (the primary server) and enable data synchronization. For instructions, see "Configure synchronization."		
4	Schedule synchronization frequency. For instructions, see "Change the synchronization schedule."		This is an optional step if the default schedule (hourly synchronizations) is not acceptable.
5	Host system panics		
6		Take over as ServerA	

Step	ServerA	ServerB	Comments
		(including NETBIOS alias). For instructions, see "Fail over to a standby server."	
7	Reboot off the network and configure as standby. For instructions, to configure as a standby see "Configure a standby server."		Steps 7 through 14 provide instructions about how to bring the initial primary server back online.
8		Release NETBIOS alias.	
9	Attach to network		
10		Add ServerA as a standby to ServerB (the primary server) and enable data synchronization. For instructions, see "Configure synchronization."	
11		Synchronize configuration files to ServerA. For instructions, see "Synchronize servers immediately."	
12	Take over as primary server For instructions, see "Fail over to a standby server"		Policies start running on ServerA and policies stop running on ServerB.
13		Configure as a standby server. For instructions, see "Configure a standby server."	
14	Add <i>ServerB</i> as a standby to ServerA (the primary server) and enable data		If ServerB was also standing by for other systems, you need to

Advanced Configuration and Techniques

Step	ServerA	ServerB	Comments
	synchronization. For instructions, see "Configure synchronization."		repeat the operation for each system for which ServerB is acting as a standby.

Note: To specify a server port different from the default, you must type the server name in the following format:

[ServerName]:[Port]

This example only applies to the server you are connecting to. Any other proxy calls between primary servers, standby servers, or monitoring agents uses the ports configured in VFM.

Configure a standby server

Configure a monitoring agent as a standby server.

Notes:

- A primary server can be the VFM server or a monitoring agent assigned to manage VFM policies.
- You must configure a standby server with the same account that you configured the primary server.

1. Install a monitoring agent. See "Deploy a monitoring agent" for instructions.
2. On the machine where the monitoring agent is installed or the primary server is running, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
3. Type **Standby** [*Server to Connect*], where [*Server to connect*] is the name of the monitoring agent you installed in Step 1 and is required only when you execute this command on a primary server.
4. Press ENTER. The monitoring agent is configured as a standby server.

Configure synchronization

Configure the relationship between primary and standby servers to enable synchronization between the servers. VFM enables you to configure a standby server for one or multiple primary servers.

1. On the machine where the primary server is running, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
2. Type **StandbyForMaster** <StandbyServer> <MasterName> [, MasterName], where <StandbyServer> is the name of a standby server, <MasterName> is the name of a primary server, and [,MasterName] is the name of another primary server. VFM enables you to configure one standby server for multiple

primary servers.

Note: See "Primary and standby server configuration options" for examples about how to configure the servers in your network.

3. Press ENTER. Data synchronization between the primary servers and the standby server is performed once every 60 minutes. You can change the synchronization frequency. See the topic "Change the synchronization schedule."

Fail over to standby server

Manually fail over to the standby server when a primary server becomes unavailable.

1. On the machine where the standby server is running, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
2. Type **Takeover <StandbyName> <MasterName> [+Alias <AliasName>]**, where *<StandbyName>* is the name of the standby server that will take over as the primary server, *<MasterName>* is the name of the primary server, and **[+Alias <AliasName>]** is the option to add a NETBIOS alias name to the standby server.
3. Press ENTER. The standby server assumes the identity of the primary server.

Note: To remove a NETBIOS alias name from a primary or standby server, see the topic "Remove NETBIOS alias."

Fail back to the primary server

Fail back to the primary server when it becomes available again.

In this example, *ServerA* is a primary server that became unavailable. You executed a failover and made *ServerB* the standby server for *ServerA*, the active server. *ServerA* is now available and you want to make *ServerA* the primary server again.

Step	Instructions	Command
1	Reboot <i>ServerA</i> off the network and configure it to be a standby server. Note: For detailed instructions, see "Configure a standby server."	On <i>ServerA</i> , run the following command: Standby ServerA
2	Remove the NETBIOS alias from <i>ServerB</i> . Notes: • If a standby is activated and takes over the NETBIOS alias of the unavailable primary server, attempting to boot the primary server after a failback encounters difficulties with NETBIOS name conflict if the alias is not removed from the standby prior to booting the primary server with the same name. • For detailed instructions, see "Remove	On <i>ServerB</i> , run the following command: RemoveAlias aliasname

Advanced Configuration and Techniques

Step	Instructions	Command
	NETBIOS alias from a server."	
3	Attach <i>ServerA</i> to the network and configure the relationship between the servers to enable data synchronization between <i>ServerB</i> and <i>ServerA</i> . Note: For detailed instructions, see "Configure synchronization."	On <i>ServerB</i> , run the following command: StandbyForMaster <i>ServerA ServerB</i>
4	Manually synchronize the configuration files in <i>ServerB</i> with <i>ServerA</i> . Note: For detailed instructions, see "Synchronize servers immediately."	On <i>ServerB</i> , run the following command: SyncToStandby <i>ServerB</i>
4	Make <i>ServerB</i> a standby server. Note: For detailed instructions, see "Configure a standby server."	On <i>ServerB</i> , run the following command: Standby <i>ServerB</i>
6	Make <i>ServerA</i> the active server. Note: For detailed instructions, see "Failover to a standby server."	On <i>ServerA</i> , run the following command: Takeover <i>ServerB</i> <i>ServerA</i>
7	Configure the relationship between the servers to enable data synchronization between <i>ServerA</i> and <i>ServerB</i> . Note: For detailed instructions, see "Configure synchronization."	On <i>ServerA</i> , run the following command: StandbyForMaster <i>ServerB ServerA</i>

Synchronize servers immediately

VFM enables you to manually synchronize the data between the primary and standby servers.

1. On the machine where the primary server is running, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
2. Type **SyncToStandby** *<MasterName>*, where *<MasterName>* is the name of the primary server.
3. Press ENTER. The configuration files on the primary server are synchronized with the configuration files on the standby servers.

Change the synchronization schedule

VFM enables you to change how often you want to synchronize configuration data between the primary and standby servers.

1. On the machine where the primary server is running, start a command prompt and change directory to the VFM installation **Tools\StandBy** subdirectory.
2. Type **SetExportFrequency** *<MasterName>* *<Frequency In Minutes>*,

where *<MasterName>* is the name of the primary server and *<Frequency In Minutes>* is how often, in minutes, you want to synchronize data. By default, the servers are synchronized once every 60 minutes.

3. Press ENTER.

View information about servers

VFM enables you to list all the standby servers for a specified primary server, view information about file synchronizations to a standby server, and view the state of a primary or standby server.

To list all standby servers:

1. On the machine where the primary server is running, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
2. Type **ListStandbyServers** *<MasterName>*, where *<MasterName>* is the name of the primary server.
3. Press ENTER. A list shows all the standby servers for the primary server.

To view information about file synchronizations:

1. On a standby server, start a command prompt and change to the VFM installation **Tools** subdirectory.
2. Type **ListStandbyTransfers** *<Standby>*, where *<Standby>* is the name of the standby server on which you want to view a list of configuration files received from the primary server.
3. Press ENTER. A list shows all the configuration files transferred to the standby server.

To view the state of a server:

1. On a primary or standby server, start a command prompt and change to the VFM installation **Tools** subdirectory.
2. Type **ShowState** [*Server to Connect*], where [*Server to Connect*] is the name of the server on which you want to get the current state. Use this option when the server on which you are running the command and the server on which you want to get the current state are not the same.
3. Press ENTER.

Remove relationship between primary and standby servers

Remove the relationship between a primary server and its standby server. You can use this command when a server is a standby machine for multiple primary machines and you do not want to remove the relationship between the standby and all the primary machines.

1. On the machine where the primary server is running, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
2. Type **StopStandingbyForMaster** *<StandbyServer>* *<MasterName>*, where *<StandbyServer>* is the name of the standby server and *<MasterName>* is the name of the primary server.

Advanced Configuration and Techniques

3. Press ENTER. Synchronization between the two servers stops.

Remove a standby server

Change a standby server back to a monitoring agent.

Note: Before you execute the following command, make sure that you removed the relationship between the primary and standby servers (see "Remove relationship between primary and standby servers" for instructions.)

1. On the machine where the standby server is running, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
2. Type **StopStandby** [*Server to Connect*], where [*Server to Connect*] is the name of the standby server you want to change back to a monitoring agent.
3. Press ENTER. The machine is changed back to a monitoring agent and the VFM server is restarted.

Remove NETBIOS alias from a server

Remove a NETBIOS alias from a primary or standby server.

1. On a primary or standby server, start a command prompt and change to the VFM installation **Tools\StandBy** subdirectory.
2. Type **RemoveAlias** <aliasname>, where <aliasname> is the name of the NETBIOS alias on the server.
3. Press ENTER.

Notes:

- It is important to remove the NETBIOS alias of a standby server that has taken over the alias of the primary server before you attempt to reboot the primary server. If you fail to release the alias, you encounter problems when you attempt to attach the primary server to the network.
- To add a NETBIOS alias name to a standby server, see the topic "Fail over to a standby server."

Export/Import configuration files

Use this feature to save VFM configuration files to a known location. The export contains the following information:

- The contents of your VFM tree view, including administrative policies, roots selected for display in the Logical View, Physical View contents, and reports.
- Any backup files you created of the logical structure of DFS roots in your namespace.
- The state of the VFM server.

Planning for VFM availability

The export/import feature enables you to maintain a backup VFM server if the primary VFM machine fails.

- You can import the VFM configuration files into a VFM installation on another machine that you can bring up if the primary VFM machine becomes unavailable.
- When the primary machine becomes available, you can export the configuration files from the backup machine and import the files back to the primary VFM machine.

Notes:

- When exporting data, you might want to create a namespace backup policy to save backups of the logical structure of all the DFS roots in your namespace. This enables you to re-create the DFS namespace if a DFS root server is lost along with the VFM server machine.
- When exporting VFM information, you must save copies of any batch files you configured to run as part of your policies. You must make these files available to the backup server before it is activated. If you install the batch files in a different location on the backup server from where they were on the primary server, you must edit your policies to reference the batch files in their new locations.
- Some antivirus programs might interfere with the ability of VFM to execute scripts. See "Running batch files in policies" for more information.

Restore VFM configuration to a known state

You can use the exported files to restore the VFM configuration on your primary machine to a known state, if necessary.

How it works

- Choose to automate or manually back up VFM configuration files:
 - VFM enables you to schedule automatic snapshots of the VFM configuration information to a removal media or the same server on which you install the standby copy of the VFM server.
 - VFM enables you to manually take a snapshot of the VFM configuration information any time you want.
 - VFM command-line utilities enable you to manually export your VFM configuration information when the VFM console is not available.
- If the machine where VFM is installed becomes unavailable, you install VFM on another machine and import the VFM configuration files to the backup machine.
- The backup machine takes ownership of all monitoring agents in the imported data.
- When the primary VFM machine becomes available, you can perform one of the following actions:
 - Manually take a snapshot of the VFM configuration information on the backup VFM machine and specify the primary VFM machine as the location where the

Advanced Configuration and Techniques

snapshot is saved. Use a command-line utility to import the files to the primary VFM machine.

- Use the same command-line utilities to export the VFM configuration files from the backup machine and import them back to the primary VFM machine.

Notes:

- You must stop all VFM processes, including any remote clients accessing the VFM server, before you can perform an import operation.
- UNIX namespace policies in VFM 3.0 are not imported to version 4.0. In VFM 4.0, UNIX namespace policies have been replaced with multiprotocol namespace policies. After you import your VFM 3.0 configuration files into version 4.0, you can create multiprotocol namespaces policies to replace any 3.0 UNIX policies.

Export/Import checklist

To export VFM configuration files and import the files to VFM on a backup machine, follow the basic steps listed below:

1. Export configuration files.
2. Import configuration files.

Export VFM configuration files

Back up your VFM server configuration files so that you can continue to manage your namespace if the machine where the VFM server is installed becomes unavailable. Before you export your configuration files, read "Export/import considerations."

To automate backups of VFM configuration files:

1. On the **Tools** menu, click **Options**.
2. Under System Options, click **Application Snapshot Options**.
3. Select **Automatically snapshot the application**.
4. Type the UNC path where you want to save the configuration files, or click **Browse** to locate the machine in your network.
5. Type the number of snapshots to retain.
6. Optionally type the UNC path of the batch files that you want to run before and after the snapshot process.
7. Type how frequently you want to back up the files.
8. Click **OK**.

To manually backup VFM configuration files:

Use the following method when you have access to the VFM client and the VFM server is running.

1. On the Tools menu, click Take Application snapshot now.
2. Type the UNC path where you want to save the configuration files, or click Browse to locate the machine in your network.
3. Type the name of the folder where the files are saved.
4. Click **OK**.

Use the following method when you do not have access to the VFM client or the VFM server is not running.

1. Close the VFM console.
2. On the machine where the VFM server is running, start a command prompt and directory to the VFM installation **Tools\Export** subdirectory.
3. Type **export** <directory name>, where <directory name> is the path to the directory where you want to save the VFM configuration files.
4. Press ENTER.

Notes:

- If the destination directory already exists, it must be empty.
- If the destination directory does not exist, it is created during the export process.
- Store the exported files in an appropriate location where you can access them if the machine hosting the VFM server fails.
- The location selected for the application snapshot must be writable by the account under which MSDE is running and the account under which the VFM server is installed. By default, MSDE runs as LocalSystem. Share permissions and directory permissions configured on the snapshot save location must allow both LocalSystem and the account under which the VFM server is installed to write to the directory.

Import VFM configuration files

If the machine where the VFM server is installed becomes unavailable, you can perform a manual failover to a backup machine. The following instructions assume that you previously exported your VFM configuration files using the utility provided by VFM. You must stop all VFM processes, including any remote clients accessing the VFM server, before you can perform an import operation.

1. Install VFM on a backup machine.
2. Close the VFM console.
3. On the machine where the VFM server is installed, start a command prompt and change to the VFM installation **Tools\Export** subdirectory.
4. Type **import** <directory name>, where <directory name> is the path to the directory where the backup files are located.

Advanced Configuration and Techniques

5. Press ENTER.
6. Start the VFM console and validate that the VFM configuration is correct.

Resume running VFM on the primary machine

When you are ready to resume running VFM on the primary machine, perform the following instructions. These instructions assume that you are running the VFM server on a backup machine and that you used the utility provided by VFM to import the VFM configuration files to the backup machine.

1. Perform the following actions on the backup VFM server machine:
 - Run the export operations.
 - On the machine where the VFM server is installed, start a command prompt and change to the VFM installation **Tools\Export** subdirectory.
 - Type **server disable** and press ENTER to stop and disable the VFM server on the backup machine. The primary and backup VFM servers **cannot be active simultaneously**.
2. Perform the following actions on the primary VFM machine:
 - Ensure that VFM is operational.
 - Start with Step 2 of "Import VFM configuration files."

Note: The path you specify for the import operation should be the location of the files that you exported from the backup machine.

Resume running VFM on the backup machine

Follow these instructions for subsequent failovers to the same backup machine.

Assumptions

- VFM server is stopped and disabled on the backup machine.
- You previously exported your VFM configuration files using the utility provided by VFM.
- You followed the instructions in "Resume running VFM on primary machine" to resume operation on the VFM primary machine.

To perform subsequent failovers on the same VFM backup machine:

1. Start a command prompt on the backup machine and change to the VFM installation **Tools\Export** subdirectory.
2. Type **server enable** and press ENTER to enable and start the VFM server on the backup machine.
3. Start with Step 4 of "Import VFM configuration files."

Chapter 6 Use SSH shell credentials

VFM executes shell commands to configure and obtain data from NetApp filers and other non-Windows systems, such as Solaris and Linux systems. VFM normally uses RSH to execute these commands. Where the environment does not support the use of RSH, you can configure VFM to use SSH communications in place of RSH.

Prerequisites

Before you configure VFM to use SSH, make sure that the VFM client and server and machines known to VFM that are involved in data transfer have the following items installed and running:

- The following programs from PuTTY suite 0.55 or later must be installed on the VFM machine:
 - Plink.exe
 - Puttygen.exe
 - pscp.exe

Note: If VFM client, server, and monitoring agent are installed on different machines, you must install the programs on all three machines.

- A UNIX machine must be running OpenSSH. See "Install and configure OpenSSH on Solaris" or "Install and configure OpenSSH on RedHat" for instructions about how to install OpenSSH on Solaris 8, RedHat 9, and RedHat Enterprise 3.
- A NetApp filer must have the SecureAdmin™ package installed. To verify that a filer has this package, look for the directory named /etc/sshd.
- When VFM uses SSH to communicate with UNIX machines and filers, SSH must be configured without a password or passphrase.

Note: Data ONTAP 6.3 does not support passwordless access through SSH.

References

To read more about SSH, OpenSSH, and PuTTY, see the following list of recommended references.

- **Covers UNIX and Windows implementation of SSH:**
Barrett, David J., and Richard Silverman. The Secure Shell: The Definitive Guide. Sebastopol: O'Reilly Media, Inc., 2001.
- **Information about SSH and PuTTY for Windows:**
http://www.jfitz.com/tips/ssh_for_windows.html
- **Information about OpenSSH:**
<http://www.openssh.org>
- **Web site where you can get a copy of PuTTY:**
<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

How to configure VFM to use SSH

The following checklist outlines the steps to configure VFM to use SSH. Follow the steps for the configuration that best describes your environment.

Important

After you complete the steps listed below, you must provide VFM with the remote SSH shell credentials to enable VFM to use plink.exe in ssh mode for remote shell operations and use pscp.exe to deploy replication agents.

- See “Set default remote SSH shell credentials” for instructions about how to set application-wide default SSH shell credentials for all new filers and UNIX machines added to the Physical View.
- See “Set remote SSH shell credentials for a filer” in the VFM [Manage Physical Resources Guide](#).
- See “Set remote SSH shell credentials for a UNIX machine” in the VFM [Manage Physical Resources Guide](#).

	Environment	Step	Instructions
1	The VFM client and server are installed on the same machine and you want to configure VFM to communicate with UNIX machines. The VFM server runs all your policies.	1. Generate a key-pair on the machine where the VFM client and server are installed.	1. "Generate a key-pair"
		2. Copy the public key to each UNIX machine known to VFM.	2. "Copy the public key to a UNIX machine"
		3. Cache the keys onto each UNIX machine known to VFM.	3. "Cache the keys"
2	The VFM client and server are installed on separate machines and you want to configure VFM to communicate with UNIX machines. You installed monitoring agents to run some of your policies.	1. Generate a key-pair on the VFM client machine.	1. "Generate a key-pair"
		2. Generate a key-pair on the VFM server machine.	2. "Generate a key-pair"
		3. Generate a key-pair on each machine you installed a monitoring agent.	3. "Generate a key-pair"

Environment		Step	Instructions
		4. Copy the public key generated on the VFM client to the UNIX machines known to VFM.	4. "Copy the public key to a UNIX machine"
		5. Copy the public key generated on the VFM server to the UNIX machines known to VFM.	5. "Copy the public key to a UNIX machine"
		6. Copy the public key generated on each VFM monitoring agent to the UNIX machines known to VFM.	6. "Copy the public key to a UNIX machine"
		7. Cache the keys on the VFM client onto the UNIX machines known to VFM.	7. "Cache the keys"
		8. Cache the keys on the VFM server onto the UNIX machines known to VFM.	8. "Cache the keys"
		9. Cache the keys on each VFM monitoring agent onto the UNIX machines known to VFM.	9. "Cache the keys"
3	The same VFM configuration as Environment 1, but you want to configure VFM to communicate with NetApp filers that are running CIFS.	1. Generate a key-pair on the machine where the VFM client and server are installed.	1. "Generate a key-pair"
		2. Copy the public key to each filer that is running	2. "Copy the public key to a filer using CIFS"

Advanced Configuration and Techniques


	Environment	Step	Instructions
		CIFS.	
		3. Cache the keys onto the CIFS filers known to VFM.	3. "Cache the keys"
4	The same VFM configuration as Environment 2, but you want to configure VFM to communicate with filers that are running NFS.	1. Generate a key-pair on the VFM client machine.	1. "Generate a key-pair"
		2. Generate a key-pair on the VFM server machine.	2. "Generate a key-pair"
		3. Generate a key-pair on each VFM monitoring agent machine.	3. "Generate a key-pair"
		4. Copy the public key generated on the VFM client to the NFS filers known to VFM.	4. "Copy the public key to a filer using NFS"
		5. Copy the public key generated on the VFM server to the NFS filers known to VFM.	5. "Copy the public key to a filer using NFS"
		6. Copy the public key generated on each VFM monitoring agent to the NFS filers known to VFM.	6. "Copy the public key to a filer using NFS"
		7. Cache the keys on the VFM client onto the NFS filers known to VFM.	7. "Cache the keys"

	Environment	Step	Instructions
		8. Cache the keys on the VFM server onto the NFS filers known to VFM.	8. "Cache the keys"
		9. Cache the keys on each VFM monitoring agent onto the NFS filers known to VFM.	9. "Cache the keys"

Generate a key-pair


When the VFM client and server are on different machines, you must generate key-pairs on both machines. If monitoring agents are running policies that transfer data, you must also generate key-pairs on the machines where the agents are installed.

1. On the machine where the VFM client or server is running, start a command prompt and type the path of the puttygen.exe file and press ENTER.



2. Select SSH1 (RSA) as the type of key to generate.
3. Click **Generate**. An SSH1 RSA key with 1024 bits is generated.

Advanced Configuration and Techniques



4. Click **Save Private Key** and specify a directory and filename. Remember where this file is saved and the name of the file. You will have to provide this information to VFM. See "Set default remote shell credentials."
5. When the message "Are you sure you want to save this key without a passphrase to protect it?" appears, click **Yes**.
6. Select the public key that appears in the "Public key for pasting into authorized_keys file" field, right-click, and click **Copy**.
7. Continue with "Copy the public key to a UNIX machine", "Copy the public key to a filer using NFS" or "Copy the public key to a filer using CIFS."

Copy the public key to a UNIX machine

After you generate the key-pairs, you copy the public key to a UNIX machine.

1. Login to the UNIX machine from your Windows machine.
2. Create a ".ssh" directory in the home directory of root. Execute the following commands:
 - To create the ".ssh" directory type the command: `mkdir .ssh`
 - To make the directory your current directory, type the command: `cd .ssh`
3. Create a file named "authorized_keys."

Note: If you generated a SSH2 (RSA) or SSH2 (DSA) key, the key can be stored in the "authorized_keys" or "authorized_keys2" file. Check with your administrator to see which file is being used in your environment. Wherever the filename "authorized_keys" appears in the instructions, substitute the filename "authorized_keys2."

4. Open the authorized_keys file and paste the public key at the end of the file. Make sure that the public key appears in one line. The key should not be separated by line breaks.
5. Save and close the file.
6. Make sure that the following directories have the appropriate permissions:

Directory	Permission
home	700 or 755
.ssh	700
authorized_keys	700 or 755

7. Check the sshd configuration of the UNIX machine to verify that public-key authentication is enabled. Look for the line:

```
PubkeyAuthentication yes
```

The line can begin with "#".

Note: The configuration file is found in the following directories:

- On RedHat 9 and RedHat Enterprise 3, the file is in /etc/ssh/sshd_config.
- On Solaris 8, the file is in /usr/local/etc/sshd_config.

8. If you changed the configuration, restart the SSH daemon.
9. Continue with the instructions to "Cache the keys."

Copy the public key to a filer using NFS

After you generate the key-pairs, you copy the public key to a NetApp filer using NFS.

1. Login to the UNIX machine from your Windows machine.
2. On a command prompt, type the command:

```
mount -t nfs <filer name>:/ /mnt
```

Filer name is the name of the filer you want to copy the public key.
3. Make sshd your current directory. Type the command:

```
cd /mnt/etc/sshd
```
4. Perform one of the following actions:
 - If the root directory does not exist, create it and then make it your current directory. Type the command:

```
mkdir root
```
 - If the root directory exists, make it your current directory. Type the command:

```
cd root
```
5. Create a directory named .ssh. Type the command:

```
mkdir .ssh
```
6. Create a file named "authorized_keys."
7. Open the authorized_keys file and paste the public key at the end of the file. Make sure that the public key appears in one line. The key should not be separated by line breaks.
8. Save and close the file.
9. Continue with the instructions to "Cache the keys."

Copy the public key to a filer using CIFS

After you generate the key-pairs, you copy the public key to a NetApp filer using CIFS.

1. In Windows Explorer, type the fully qualified path of the filer.
2. Open the `\\<filer name>\etc$\sshd\root\.ssh` directory. Where `<filer name>` is the name of the filer you want to copy the public key.
3. Open the `authorized_keys` file and paste the public key at the end of the file. Make sure that the public key appears in one line. The key should not be separated by line breaks.
4. Save and close the file.
5. Continue with the instructions to "Cache the keys."

Cache the keys

After you copy the public key to the UNIX machine or filer, you need to run the "plink.exe" program on the VFM client and server machines. This program caches the keys. The program must be executed before you can use VFM to transfer data.

Note: If monitoring agents are assigned to run policies that transfer data, you must also run the "plink.exe" program on the machines where the monitoring agents are installed.

1. On the machine where you generated the key-pairs, start a command prompt and type the path of the "plink.exe" file. Type the following command:

```
Plink <host> -l<users>
```

- **Host** is the fully qualified name of the UNIX machine or filer (for example, `secure.mmd.com`).
- **Users** is the account name used to login to the UNIX machine or filer.

Notes:

- The "plink.exe" program must be executed using the account that runs VFM.
- You are prompted to add the key to the cache of the account that runs VFM.

2. Type the command `plink` again. For `<host>` type the NetBIOS name of the UNIX machine or filer (for example, `secure`).
3. Verify that you are not prompted again to add the key to the cache. Complete Step 1 and 2 again.

Install and configure OpenSSH on Solaris

Before you install OpenSSH on a UNIX machine, you must install a patch for /dev/random/. Contact Sun or go to the web site <http://sunsolve.sun.com> to get the latest version of patch 112438-03.

1. Download the patch and install it as indicated in the instructions.
2. Verify that /kernel/drv/random, /dev/urandom, and /dev/random exist on the UNIX machine.
3. Reboot the machine.

Note: If the files listed in Step 2 do not exist, execute the following commands:

```
ln -s /kernel/drv/random /dev/random
ln -s /kernel/drv/random /dev/urandom
```

Step 1 Get the OpenSSH packages for SPARC

1. Get the following packages from sunfreeware:
 - openssl-3.9pl-sol8-sparc-local.gz
 - openssl-0.9.7d-sol8-sparc-local.gz
 - zlib-1.2.1-sol8-sparc-local.gz
 - libgcc-3.3-sol8-sparc-local.gz
2. Copy the packages to the UNIX machine.

Step 2 Install the packages

1. Change to the directory where you copied the packages.
2. Type the following commands:


```
gunzip openssl-3.9pl-sol8-sparc-local.gz
gunzip openssl-0.9.7d-sol8-sparc-local.gz
gunzip zlib-1.2.1-sol8-sparc-local.gz
gunzip libgcc-3.3-sol8-sparc-local.gz
pkgadd -d openssl-3.8.1p1-sol8-sparc-local
pkgadd -d openssl-0.0.7d-sol8-sparc-local
pkgadd -d zlib-1.2.1-sol8-sparc-local
pkgadd -d libgcc-3.3-sol8-sparc-local
```

Notes:

- ssh is installed in /usr/local/bin.
- sshd is installed in /usr/local/sbin.
- Make sure that /usr/local/bin and /usr/local/sbin are in your PATH environment variables.

Step 3 Set up the sshd user and the /var/empty directory

Execute the following commands as root:

```
mkdir /var/empty
chown root:sys /var/empty
chmod 755 /var/empty
groupadd sshd
useradd -g sshd -c 'sshd privsep' -d /var/empty -s
/bin/false sshd
```

Note: The directory /var/empty should not contain any files.

Step 4 Install ssh and sshd

Each machine that you want to communicate with SSH must have an sshd daemon running.

1. Run the following three lines to create the key information for the server machine. Make sure that you have /usr/local/bin and /usr/local/sbin in your PATH. As root, type:

```
ssh-keygen -t rsa -f /usr/local/etc/ssh_host_key -N ""
ssh-keygen -t dsa -f /usr/local/etc/ssh_host_dsa_key -N
""
ssh-keygen -t rsa -f /usr/local/etc/ssh_host_rsa_key -N
""
```

Note: It may take a few minutes to complete the commands.

2. Create a script named "sshd" to start the sshd daemon. Place the script in the /etc/init.d directory.

```
#!/sbin/sh
#
#
[ ! -d /usr/local/sbin ] && exit
case "$1" in
'start')
    /usr/local/sbin/sshd
    ;;
'stop')
    /usr/bin/pkill -x -u 0 'sshd'
    ;;
```

```

*)
    echo "usage: $0 { start | stop }"
    exit 1
;;
esac
exit 0

```

3. Run the following commands as root:

```

chown root /etc/init.d/sshd
chgrp sys /etc/init.d/sshd
chmod 555 /etc/init.d/sshd
ln -s /etc/init.d/sshd/etc/rc2.d/S98sshd

```

Step 5 Start the process

The final step is to start the sshd daemon. Execute the following command as root:

```
/etc/rc2.d/S98sshd start
```

Note: To stop the sshd daemon, type the following command:

```
/etc/rc2.d/S98sshd stop
```

Install and configure OpenSSH on RedHat

Use the following instructions to install and configure OpenSSH on RedHat 9 and RedHat Enterprise 3.

Step 1 Get the latest OpenSSH server packages

1. Copy the appropriate packages for RedHat 9 or RedHat Enterprise 3 to the UNIX machine.
2. Change to the directory where you copied the packages.

Step 2 Install the OpenSSH server packages

Type the command `Rpm -ivh <package filename>`, where `<package filename>` is the name of the package you downloaded.

Step 3 Start the SSH service

1. Configure RedHat to start the SSH service automatically when the system reboots. Type the command:

```
chkconfig --level 35 sshd on
```

Advanced Configuration and Techniques

2. Start the SSH service. Type the command:
`/etc/init.d/ssh start`

Set default remote shell credentials

VFM executes shell commands to configure and obtain data from NetApp filers and other non-Windows systems, such as UNIX and Linux systems. VFM normally uses RSH to execute these commands. Where the environment does not support the use of RSH, you can use SSH communications in place of RSH.

Important

Before you can complete the information requested under SSH Settings, you must configure VFM to use SSH. See "Use SSH shell credentials" for instructions about how to perform the following tasks:

- Install and configure OpenSSH.
- Generate a key-pair.
- Copy the public key to UNIX machines or filers.
- Cache the keys.

1. On the **Tools** menu, click **Options**.
2. Click **Remote Shell**.
3. Select a default shell type from the list.
4. Perform one of the following actions:
 - If you select **RSH** as the default shell, type the location of the rsh and rcp executables and the user name and password that VFM will use to communicate with filers and UNIX machines. When you leave the user name and password blank, VFM uses "root" as the user name with no password.
 - If you select **SSH** as the default shell, type the location of the plink and scp executables and the private key and then type the user name that VFM will use to communicate with filers and UNIX machines. When VFM uses SSH to communicate with UNIX machines and filers, SSH is configured without a password or passphrase.

Chapter 7 Use RSH shell credentials

Normally, VFM communicates with non-Windows machines using RSH. RSH must be enabled as root on Linux and Solaris systems.

How to enable RSH as root

- See "Enable RSH as root on a Linux system" for instructions about how to enable RSH on a Linux system.
- See "Enable RSH as root on a Solaris system" for instructions about how to enable RSH on a Solaris system.

Enable RSH as root on a Linux system

For VFM to communicate with Linux systems using RSH, the systems must enable RSH as root.

1. Make sure that the rshd server is installed. Type the following command to verify that the rshd server is installed:

```
rpm -qa | grep rsh
```

After executing the command, look for a package name, such as "rsh-server-0.17-14." If you do not have this package, you can get it from the RedHat media.

2. Edit */etc/xinetd.d/rsh* and find the line:

```
disable = yes  
and change it to:  
disable = no
```

3. Restart Xinetd. Type the following commands:

```
/etc/init.d/xinetd stop  
/etc/init.d/xinetd start
```

4. Edit */etc/hosts.equiv* or create it if it does not exist. Add the following line:

```
+ +
```

Make sure that there are no leading spaces and that there is only one space between the

Advanced Configuration and Techniques

two plus symbols. This format is important. Instead of the above line you can also add:

```
<machine1.npb.com> +<user1>
```

where <machine1.npb.com> is the name of the machine where the VFM server is installed and <user1> is the account used by the VFM server. This enables the VFM server to access the Linux system using RSH.

Note: The /etc/hosts.equiv file must have permission 600 and must be owned by root.

5. Edit /etc/hosts.allow or create it if it does not exist. Add the following line at the end of the file and make sure that there are no leading spaces:

```
ALL: ALL
```

6. Make sure that /etc/hosts.deny is empty, or has only lines which begin with '#'. These are comments and are ignored. If there are existing entries in hosts.deny/hosts.allow, check with your sysadmin about allowing connection to rsh from specific hosts.
7. Edit /etc/pam.d/rsh and change the following line:

```
auth          required    pam_rhosts_auth.so

to

auth          sufficient   pam_rhosts_auth.so
```

8. Add entries for rsh, rlogin and rexec in /etc/securetty as shown in the following excerpt:

```
.
tty10
tty11
rsh
rlogin
rexec
```

9. Check whether the kernel firewall is blocking incoming or outgoing packets. Type the following command:

```
Iptables -list
```

If rules are listed, the firewall might be blocking connection attempts. Ask the firewall administrator to allow RSH ports through the firewall.

10. Verify RSH from another machine. Type the following command:

```
rsh -l root <your host here> ls -la
```

If you see the directory listing, RSH as root is configured correctly.

Enable RSH as root on a Solaris system

For VFM to communicate with Solaris systems using RSH, the systems must enable RSH as root.

1. Ensure that the SUNWcsu package is installed. The following command should display information about the package if it is installed:

```
pkginfo -l SUNWcsu
```

2. Ensure that the r-services (for example, RSH, REXEC, and RLOGIN) are enabled in */etc/inetd.conf*.
3. Ensure that the */etc/default/login* file has *CONSOLE* variable commented out.
4. Edit the file *.rhosts* in the roots home directory or create it if it does not exist. Add the following line:

```
+ +
```

Make sure that there are no leading spaces and that there is only one space between the two plus symbols. This format is important. Instead of the above line you can also add

```
<machine1.npb.com> +<user1>
```

where *<machine1.npb.com>* is the name of the machine where the VFM server is installed and *<user1>* is the account used by the VFM server. This enables the VFM server to access the Solaris system using RSH. Make sure that the file has permission 600 and is owned by root.

5. Verify RSH from another machine. Type the following command:

```
rsh -l root <your host here> ls -la
```

If you see the directory listing, RSH as root is configured correctly.

Appendix

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

No.	Title	File name	Contents
1.	Advanced Configuration and Techniques	AdvancedConfig.pdf	• Explains how to use VFM to configure security permissions • Explains how to manage replication and monitoring agents • Explains how to run batch files in policies and reports • Explains how to make the VFM server highly available
2.	Data Lifecycle Automation Policies Guide	DataLifecyclePolicies.pdf	Explains how to create and manage archival migration policies to optimize storage.
3.	Data Movement Policies Guide	DataMovePolicies.pdf	• Explains how to transfer CIFS and NFS data by creating migration and replication policies • Explains how to transfer data between NetApp filers by creating SnapMirror policies • Explains how to load balance Windows servers by creating storage load balancing policies
4.	Disaster Recovery Policies Guide	DRPolicies.pdf	Explains how to make data highly available by creating client, filer, and server disaster recovery policies
5.	Getting Started Guide	GettingStarted.pdf	A tutorial to get you started using VFM. It explains how to install VFM and steps you through the process of creating a logical namespace, making the namespace highly available, making data highly available, transferring CIFS and NFS data, creating multiprotocol namespaces, virtualizing user home directories, and backing up data and the namespace.
6.	Manage Physical	PhysicalResources.pdf	Explains how to use VFM to Work

Advanced Configuration and Techniques

No.	Title	File name	Contents
	Resources Guide		with NetApp filers, set replication properties, and work with Windows and UNIX machines
7.	Namespace Guide	Namespace.pdf	Explains how to create a DFS namespace (roots, links, and link targets), namespace creation, permission based, namespace availability, multiprotocol namespace, home folder, and namespace backup policies
8.	Reports Guide	Reports.pdf	Explains how to create and publish reports
9.	Snapshot Policies Guide	SnapshotPolicies.pdf	Explains how to use VFM to manage snapshots on NetApp filers using snapshot access and snapshot scheduling policies
10.	Troubleshooting Guide	Troubleshooting.pdf	• Includes diagnostic information, such as troubleshooting, event IDs, migration, replication, and status indicators • Provides examples of file exclusions and inclusions

Glossary

A

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

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

Alternate volume: SEE Replica.

C

Child node: SEE DFS link.

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

Client: The user interface component of the application.

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

D

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

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

DFS branch: SEE Shared folder.

DFS leaf: SEE Shared folder.

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

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

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

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

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

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

Advanced Configuration and Techniques

DFS volume: SEE Shared folder.

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

E

Export: SEE NFS export

F

Fault-tolerant DFS: SEE Domain-based DFS root

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

FRS: Microsoft File Replication Service.

G

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

I

Inter-DFS link: SEE DFS link.

J

Junction point: SEE DFS link.

K

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

L

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

M

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

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

Monitoring agent: The component that can be installed to evaluate storage policies from a

different machine than the one where the application server is installed.

N

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

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

P

Pool: A group of storage locations.

Primary server: An active application server or monitoring agent.

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

Q

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

R

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

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

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

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

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

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

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

Advanced Configuration and Techniques

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

S

Secondary target: SEE Standby target.

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

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

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

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

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

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

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

T

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

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

U

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

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

W

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

Index

A

Account password	3
Application	
configure permissions	1
export configuration files	44, 46
import configuration files	45, 46
Application server	
change synchronization frequency	40
configure a standby server	38
configure primary and standby servers	36
configure relationship with standby	38
fail back	39
fail over to standby	39
manual synchronization	40
remove relationship with standby server	41

B

Back up configuration files	44
Batch files	26

C

Change server password	3
Configuration files	44, 45, 46
Credentials	13

D

Data synchronization	38
Data transfer	14
Deploy a replication agent	
requirements	11

E

Export configuration files	44, 46
----------------------------------	--------

F

Failback	39, 46
Failover	39, 45, 46
Fault tolerance	44

I

Import configuration files	45, 46
----------------------------------	--------

K

Key-pair	
generate	51
Key-pair	51

L

Linux	59
List of replication agents	19
Listening ports	4, 13

M

Monitoring agents	
assign to a policy	22
back up an agent	22
change credentials	23
configure as standby server	38
install an agent	22
save snapshot of agent to new location	22
uninstall an agent	22
view status	23

N

NETBIOS alias name	39, 42
Network communication	4, 13
Network ports	4, 13
NFS data transfer	14

O

OpenSSH	55, 57
---------------	--------

P

Password	3
Permissions	1
Ports	13, 36
Primary server	
remove relationship with standby server	41
Primary server	36
Private key	
cache keys	54
Private key	54
Proxy machine	7
Public key	
cache keys	54
copy to filer	53
copy to UNIX machine	52

R

RedHat	14, 57
Remove replication agent from database	13
Replication agent	
change credentials	13
change listening port	13
deploy an agent	11
deployment of	7
install an agent	14
pause an agent	11
purge from database	13
restart an agent	11
resume an agent	11
specify a proxy machine	18
specify proxy machine	7
stop an agent	11

Advanced Configuration and Techniques

uninstall an agent	12, 14
upgrade an agent	12, 14
Replication agent	11
Replication agents	
list of UNIX replication agents	19
list of Windows replication agents	19
Specify default proxy agent	18
Restart a replication agent	11
RSH	58, 59

S

Scripts	26
Security	1
Server password	3
Servers	
current state of server	41
Servers	41
Set permissions	1
Shell credentials	58
Solaris	14, 55, 59
SSH	51, 52, 53, 54, 58

Standby server	
change back to monitoring agent	42
view list of	41
Standby server	36, 38
Stop a replication agent	11
synchronize primary and standby servers	
change frequency	40
view information about	41
synchronize primary and standby servers	40

U

Uninstall a replication agent	12
UNIX	14, 18
UNIX replication agents	19
Upgrade a replication agent	12

W

Windows	18
Windows replication agents	19