NetworkAppliance®

The evolution of storage.™

VFM (Virtual File Manager) v6.0
Advanced Configuration and Techniques

May 29, 2007

Publication Number: 215-03203_A0

Copyright © 2007 Brocade Communications Systems, Inc. All Rights Reserved.
Copyright © 2007 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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

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

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

This product includes software developed by the Apache Software Foundation.

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

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

Trademark information

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

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

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

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

Preface

What the guide covers

This guide describes advanced topics, such as how to configure permissions on the Network Appliance® VFM® 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 Agent Management** – Explains how to deploy and manage replication agents on Windows and non-Windows systems, how to install and manage monitoring agents, and how to manage the VFM server.
- **Chapter 3 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 4 Application Server Availability** – Explains how to make the VFM server highly available.
- **Chapter 5 Use SSH Shell Credentials** – Explains how to configure VFM to use SSH.
- **Chapter 6 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.
- *`code_text_in_italics`* – Variable; replace with actual value.

Using the online help

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

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

- Definitions of new terms are available wherever they appear in help topics. The first time a new term appears in a help topic, it appears as a green

hyperlink.

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

Support

Contact NetApp technical support, 24 x 7.

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

Support phone numbers:

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

Note: To get your product build number and license information, click About **VFM** 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 VFM**.
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**.

Note: After you change the activation key, you must restart the VFM client.

Table of Contents

PREFACE ..... **I**

CHAPTER 1 INTRODUCTION **1**

PERMISSIONS ON ACCOUNTS USED WITH VFM COMPONENTS 1
To configure access permissions to the VFM server *1*
To configure system access permissions on domain-based roots and shared folders *2*
To configure function-level access permissions *2*
Security permissions *3*

CHAPTER 2 AGENT MANAGEMENT **5**

ACCOUNT PERMISSIONS AND PRIVILEGES FOR RUNNING VFM 5
REPLICATION AGENT REQUIREMENTS 6
UNIX replication agents *6*
Windows replication agents *7*
THE DESKTOP 8
Host view *8*
Properties view *8*
Status view *8*
DISTRIBUTED REPLICATION 9
Algorithm used by VFM *9*
Network ports *12*
Reasons to use a proxy machine for a replication agent *13*
View the status of an agent *13*
MANAGE REPLICATION AGENTS 13
Checklist *13*
Manually deploy a replication agent *14*
Set default credentials *15*
Stop and restart an agent *16*
Pause and resume an agent *16*
Upgrade an agent *16*
Create an agent group *17*
Uninstall replication agents *17*
Change the credentials of an agent *18*
Change the listening port of an agent *19*
Install a UNIX replication agent *19*
Purge an agent *25*
MANAGE MONITORING AGENTS 25
Distributed monitoring *25*
Assigning monitoring agents to policies *26*
Checklist *26*
Install a monitoring agent *26*
Change the credentials of a monitoring agent *27*
Change the listening port *27*
Assign a monitoring agent to a policy *28*
Specify a backup location *28*
Purge an agent *29*
MANAGE THE APPLICATION SERVER 29
Checklist *29*
Change the server credentials *29*
Change the listening port *30*
Specify backup location *30*

CHAPTER 3 RUN BATCH FILES **33**

RUNNING BATCH FILES 33
BATCH FILES 35

CHAPTER 4 VFM SERVER AVAILABILITY **43**

Planning for VFM availability *43*
Selecting a method *43*
EXPORT/IMPORT CONFIGURATION FILES 44
Export VFM configuration files *46*
Import VFM configuration files *47*
Resume running VFM on the primary machine *48*
Resume running VFM on the backup machine *48*

CHAPTER 5 USE SSH SHELL CREDENTIALS **49**

Prerequisites *49*
References *49*
Checklist *50*
GENERATE A KEY-PAIR 54
COPY THE PUBLIC KEY TO A UNIX MACHINE 55
COPY THE PUBLIC KEY TO A FILER USING NFS 56
COPY THE PUBLIC KEY TO A FILER USING CIFS 57
CACHE THE KEYS 58
ADD PRIVATE KEYS TO PAGEANT 59
INSTALL AND CONFIGURE OPENSSH ON SOLARIS 59
Step 1 Get the OpenSSH packages for SPARC *60*
Step 2 Install the packages *60*
Step 3 Set up the sshd user and the /var/empty directory *60*
Step 4 Install ssh and sshd *61*
Step 5 Start the process *62*
INSTALL AND CONFIGURE OPENSSH ON REDHAT 62
Step 1 Get the latest OpenSSH server packages *62*
Step 2 Install the OpenSSH server packages *62*
Step 3 Start the SSH service *62*
SET DEFAULT REMOTE SHELL CREDENTIALS 62

CHAPTER 6 USE RSH SHELL CREDENTIALS **65**

ENABLE RSH AS ROOT ON A LINUX SYSTEM 65
ENABLE RSH AS ROOT ON A SOLARIS SYSTEM 66

APPENDIX **69**

VFM GUIDES 69

GLOSSARY **71**

INDEX **75**

Chapter 1 Introduction

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

Permissions on accounts used with VFM components

- 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

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

  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 appropriate privileges can prevent policies from functioning correctly.
- Make sure that you have administrator privileges for any servers hosting stand-alone roots and domains and servers hosting domain-based roots before you attempt to administer them using either VFM client or native tools.
  - **Stand-alone DFS roots -** The account must be a member of local administrators on the system hosting the root and must have permissions to write to the root share on that system.
  - **Domain-based DFS roots -** See the Microsoft DFS FAQ.

  http://www.microsoft.com/windowsserver2003/community/centers/fileservices/dfsfaq.mspx?pf=true
- These problems might be manifested as an inability to create, modify, or delete roots or links in a namespace for which the account running VFM lacks appropriate permissions.

Note: For more information about account permissions and privileges for installing and running VFM, see the VFM release notes

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 creating or deleting a policy. You can configure application-wide, function-level access permissions and configure permissions for a DFS root in the Logical View, application snapshots, and 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**, then click **Set Security**.
- To configure permissions for a DFS root, right-click the root and click **Properties**, then click the **Application Security** tab.
- To configure permissions for all policies, right-click Admin View and click **Application Security**.
- To configure permissions for a policy category or type, right-click the item and click **Application Security**.

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

  `SERVERMACHINE` is the machine hosting the VFM server.

  - The VFM client on `SERVERMACHINE` is used to restrict VFM security access so that only `SERVERMACHINE\Administrators` can perform management operations in VFM.
  - When the VFM client hosted on machine `SECONDCLIENTMACHINE` connects to the VFM server on `SERVERMACHINE` and views application security settings, the display shows `SECONDCLIENTMACHINE\Administrators` and not `SERVERMACHINE\Administrators`.

  This is only a display issue. This happens because the SID of the local administrators group is "well known" as part of the Windows infrastructure. The VFM server running on `SERVERMACHINE` will prevent operations by accounts other than those in the `SERVERMACHINE\Administrators group`.
- For information about function-level permissions, see "Security 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.

Chapter 2 Agent Management

Agent Management enables you to manage the VFM server, monitoring agents, and Windows and UNIX replication agents. The tool enables you to perform the following tasks:

- Manually deploy Windows and UNIX replication agents
- Set default credentials for replication agents
- Change the credentials of an agent or the VFM server
- Create replication agent groups to distribute migration workload
- Upgrade a replication agent
- Change the listening port of an agent or the VFM server
- Specify a backup location for a monitoring agent or the VFM server

Account permissions and privileges for running VFM

- The account used by the VFM server, a replication agent, or a VFM monitoring agent must have the following permissions:
  - **SeBackupPrivilege** ("Back up files and directories") - required to open source directories/files with backup semantics
  - **SeRestorePrivilege** ("Restore files and directories") - required to open destination directories/files with backup semantics
  - **SeSecurityPrivilege** ("Manage auditing and security log") - required to access SACLs (System Access Control Lists)
  - **SeTakeOwnershipPrivilege** ("Take ownership of files or other objects") - required to delete directories/files
- The account used for the VFM server, which is specified during installation, must have the following permissions and privileges:
  - Appropriate permissions to administer the DFS roots that are managed
    - **Standalone roots** - The account must be a member of local administrators on the system hosting the root and must have permissions to write to the root share on that system.
    - **Domain-based roots** - Microsoft Active Directory must already be operational. The account used for VFM must have administrator permissions for both the domain and the server where you create the DFS root. See the Microsoft DFS FAQ

      http://www.microsoft.com/windowsserver2003/community/centers/fileservices/dfsfaq.mspx?pf=true
  - Administrator privileges on any machines hosting either source or destination of data replications
  - "Run as service" privilege
- The account running the VFM client must have permissions to write to the *All Users\Application Data* directory tree. Users whose accounts lack write permissions to

these directories will experience errors when attempting to run VFM reports or set VFM system options.

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

Replication agent requirements

Before you deploy a replication agent, review system requirements.

Important:

- You must use a UNIX replication agent when you transfer NFS data.
- You must use a Windows replication agent when you transfer CIFS data.
- Make sure that the machine hosting the VFM server can be reached by a Windows or UNIX replication agent using its NETBIOS host name. It is not sufficient for the machine hosting the VFM server to be reached only by its fully qualified DNS name.

UNIX replication agents

- The UNIX machine where you deploy or install a replication agent must be running RedHat® Enterprise 4, or Solaris® 10 operating system.
- The account used by VFM to deploy or install a UNIX agent must have root privileges on the machine where you deploy or install the replication agent.
- NFS 3 or later must be installed on the source or destination if a UNIX replication agent is not installed. You do not have to install NFS on the source or destination that is hosting a UNIX replication agent because the replication agent reads the files directly.
- 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 10

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

- SUNWcsl
- SUNWlibms
- SUNWcsu

- SUNWloc
- SUNWgzip

RedHat Enterprise 4

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

- glibc-2.3.4-2
- glibc-common-2.3.4-2
- bash-3.0-19.2
- coreutils-5.2.1-31
- gzip-1.3.3-9
- grep-2.5.1-16
- procps-3.2.3-7
- Redhat package manager version 4 or greater.

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

Windows replication agents

- The Windows machine where you deploy a replication agent must be running Windows 2000, Windows Server 2003, or Windows XP file systems.
- The account used by VFM to deploy a Windows replication agent must have administrator privileges on the machine that receives the replication agent. The account must be able to perform the following actions:
    - Access remote registry
    - Write to Admin$ share
    - Install and run a service

The desktop

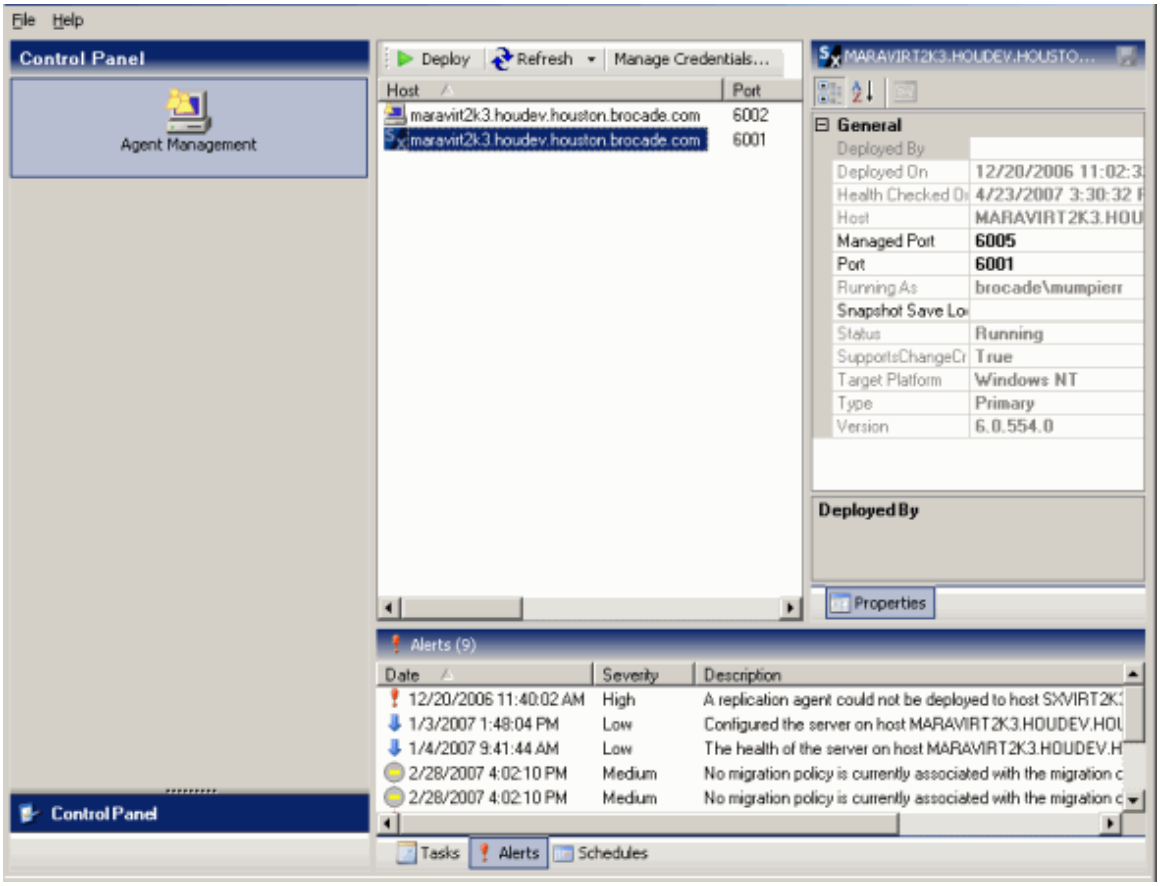

Host view

- Manually deploy a Windows or UNIX replication agent.
- Set default credentials for all Windows and UNIX replication agents that you and the VFM server deploy.
- Right-click the server or an agent to view a context menu.
- Sort the list by platform, credentials, agent group and more.
- Specify how you want to refresh the list.

Properties view

Change the properties of selected replication and monitoring agents and the application server, such as listening port and credentials.

Status view

- View tasks in progress, such as changing credentials of an agent.
- View information about problems and errors that might occur when a task runs.

Distributed replication

This section explains how VFM automatically distributes CIFS and NFS replication agents to transfer data. The section also provides instructions about how you can view the status of an agent.

Algorithm used by VFM

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

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

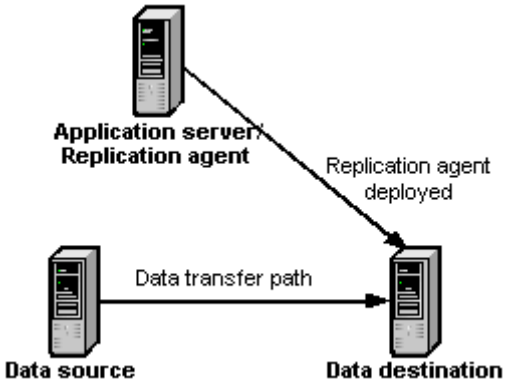

2. If VFM cannot deploy the agent to this machine, it attempts to deploy the agent to the machine hosting the source target.

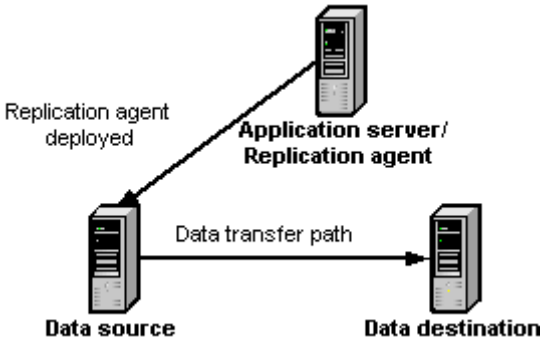

3. If VFM cannot deploy the agent to either the destination or source target, it attempts to use the agent on the proxy machine of the destination target, if one is configured.

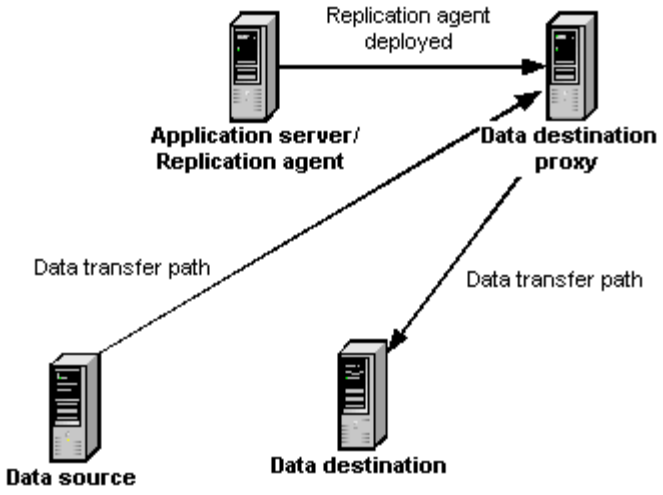

4. If VFM cannot deploy the agent to the proxy machine of the destination target, it attempts to use the agent on the proxy machine of the source target, if one is configured.

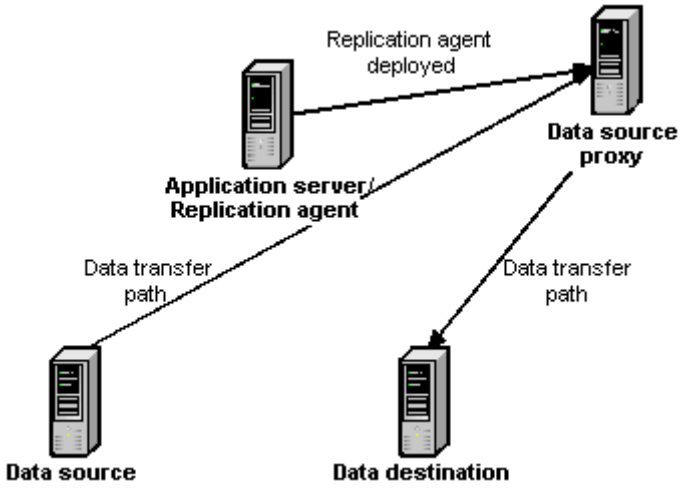

5. If VFM cannot use any of the agents listed above, VFM uses the replication agent on the VFM server or the monitoring agent to transfer data.

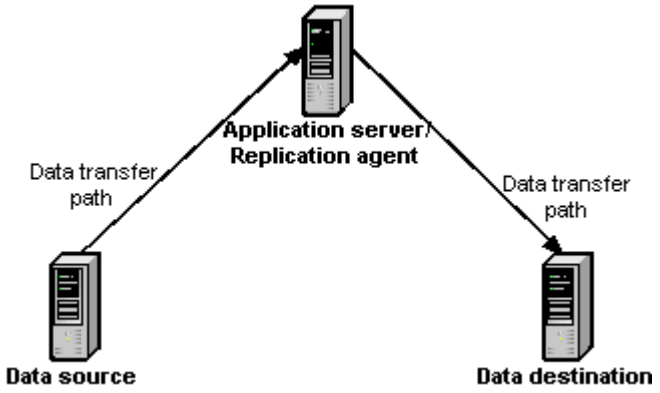

- You can set the data transfer rate limits for machines that participate in data transfers by using the properties dialog box for the machine in the VFM Physical View.
- The replication agent is deployed as part of VFM policies and file differences reports.
- The account used by VFM to deploy a replication 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 agent 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 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 files systems.
- (CIFS only) When data is replicated using differential replication, VFM deploys a replication agent to both the source and destination. If VFM cannot deploy a replication agent to either the source or destination, you must designate a proxy for that location.
- (CIFS only) If the machine 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.

Network ports

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
Tapestry StorageX Server or monitoring agent		6001and 6005
Replication Agent		6002

Network Ports Used on NetApp filers with NFS enabled

Service Name	UDP	TCP
NFS mount daemon (mountd)	602	603
NFS status daemon (statd,statmon)	604	605
NFS lock manager (lockd, nlockmgr)	606	607
NFS quota daemon (quotad, rquotad)	608	

Network Ports Used by other systems

Service Name	UDP	TCP
NFS mount daemon (mountd)	611	612
NFS status daemon (statd,statmon)	613	614
NFS lock manager (lockd, nlockmgr)	615	616
NFS quota daemon (quotad, rquotad)	617	

Note:

- The VFM server and UNIX Replication Agent communicate with storage devices using NFS during NFS migrations.
- The VFM server uses SNMP when identifying Network Appliance filers.
- The VFM server uses RSH or SSH when issuing some commands to UNIX machines and Network Appliance filers.

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.

View the status of an agent

View status information about a replication or monitoring agent or the VFM server.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click a replication or monitoring agent or the VFM server and click **Check Health**.
4. Click the **Alerts** tab to view errors.

Manage replication agents

This section provides instructions about how to manage Windows and UNIX replication agents.

Checklist

- To deploy a replication to a machine to specify, see "Manually deploy a replication agent."
- To set default credentials for all Windows and UNIX replication agents deployed by the VFM server, see "Set default credentials."
- To stop and restart a replication agent, see "Stop and restart an agent."
- To pause and resume a replication agent, see "Pause and resume an agent."
- To upgrade a replication agent to a new version, see "Upgrade an agent."
- To group replication agents to distribute the migration workload, see "Create an agent group."
- To uninstall a replication agent, see "Uninstall replication agents."
- To remove a replication agent from the VFM database, see "Purge an agent."
- To change the credentials of a replication agent, see "Change the credentials of an agent."
- To change the listening port that VFM normally uses, see "Change the listening port of an agent."
- To install a replication agent on a UNIX machine, see "Install a UNIX replication agent."
- To specify application-wide default replication agents for Windows and UNIX machines to use when VFM cannot deploy a replication agent onto the source or destination, see

"Specify default replication agents."

Note: For instructions about how to set replication properties on machines participating in the namespace or a storage policy, see the following topics:

- To specify a Windows or UNIX machine as a proxy to host a replication agent instead of having the agent deployed to a machine participating in a policy, see "Specify a proxy machine."
- To prevent deployment of agents to individual machines, see "Disallow deployment of replication agents."
- To set throttling on a machine hosting a replication agent, see "Set data transfer rate."

Manually deploy a replication agent

VFM automatically deploys Windows and UNIX replication agents to manage data transfers. You can manually deploy agents. Before you deploy an agent, review the system requirements.

- VFM cannot deploy replication agents to NetApp filers.
- A proxy replication agent is used when VFM cannot or you do not want to deploy an agent to the source or destination during data transfer.
    - For instructions about how to set application-wide default proxy Windows and UNIX replication agents, see "Specify default replication agents."
    - For instructions about how to specify a Windows or UNIX machine as a proxy to host a replication agent when you do not want an agent deployed to a machine, see "Specify a proxy machine."

To deploy a replication agent

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, click **Deploy**.
4. In the **Host** box, type the name or UNC path of the machine where you want to deploy a replication agent or click **Browse** to find a machine in your network. VFM determines whether to deploy a Windows or UNIX replication agent.
5. (Optional) Type a different port. VFM normally uses port 6002 to communicate with replication agents.
6. Specify the credential that VFM will use to communicate with the replication agent.
    - Select **The default replication agent user name** when you want to use the default credentials specified by VFM or use the default credentials you specified.

Notes:

Application default credentials

- Deployed Windows replication agents use the VFM server credentials.
- Deployed UNIX replication agents use root with no password.

User specified default credentials

To specify default credentials for Windows and UNIX replication agents that are deployed, see "Set default credentials."

- Select **The user name and password specified below** and type new credentials.

7. Click **OK**.
8. Click the **Alerts** tab to verify that the replication agent was deployed successfully.
9. Click **Refresh** in the Host view to display the new replication agent.

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.

Set default credentials

When you specify default Windows and UNIX replication agent credentials, all replication agents deployed after you change the credentials will use the default credentials. Replication agents that are already deployed are not affected.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, click **Manage Credentials**.
4. Click **Change** for the appropriate system.
5. Type a new user name.
6. Type the password.
7. Type the password again to confirm.
8. Click **OK**.
9. Repeat Steps 4 through 8 to change the credentials for replication agents deployed to another system.
10. Click **OK**.

Note: To change the credentials of deployed replication agents, see "Change the credentials of an agent."

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.

To stop an agent

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click a replication agent and click **Stop**.

To restart the agent

In the Host view, right-click the replication agent and click **Start**.

Pause and resume an agent

When you **pause** a replication agent while it is transferring data, the data transfer is suspended and **resumes** when the agent starts again.

To pause an agent

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click a replication agent and click **Pause**.

To resume the agent

In the Host view, right-click the replication agent and click **Resume**.

Upgrade an agent

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

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click a replication agent and click **Upgrade**.
4. To view the progress of the update, click **Tasks** in the Status view.
5. To view information about errors that might occur, click **Alerts** in the Status view.

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 "Troubleshooting General Problems."

Create an agent group

When you want to distribute a migration workload or increase the availability of agents during migration, assign replication agents to groups and then assign the agent group to a policy.

- When the policy runs, VFM selects a replication agent from the group that is not running a migration operation to run the migration.
- When a machine hosting a replication agent is not available, VFM uses another agent within the group to run the migration.
- When a migration operation contains multiple migration tasks, VFM distributes the workload among the replication agents in the group.
- A replication agent can belong to only one agent group.
- There are no limits to the number of agents that an agent group can contain.

To create a replication agent group

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, click a replication agent .
4. In the Properties view, perform one of the following actions:
   - Type a name for the group in the **Group** text box.
   - Click the down arrow, click **New**, and type a name for the group.
5. (Optional) Type a word or phrase that best describes the group.
6. Click **OK**. The replication agent is assigned to the group and the group name is added to the list.

To assign a replication agent to an existing group

1. In the Host view, click a replication agent.
2. In the Properties view, click the **Group** arrow and select a group from the list.
3. Click **Save**.

To view the replication agents assigned to a group

In the Host view, click the **Group** column to sort the list alphabetically by agent group.

Note: Migration, replication, archival migration, storage load balancing, and data user management policies enable you to select which replication agent or agent group performs data migration. For information about how to select a replication agent or agent group, see "Select a replication agent."

Uninstall replication agents

When VFM or a monitoring agent is uninstalled, replication agents deployed by the VFM server or monitoring agent are not deleted. You have the option of uninstalling replication agents.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.

3. In the Host view, right-click the replication agent and click **Uninstall**.
4. Select one of the following options:
   - **Uninstall using the same credentials that were used to deploy each of the selected replication agents.** Select this option when the credentials used to run the replication agent have the security permissions to uninstall the agent.
   - **Uninstall the selected agents using the following credentials.** Select this option when the credentials used to run the replication agent do not have the security permissions to uninstall the agent.
5. If you selected to uninstall the replication agent using different credentials, type the credentials.
6. Click **OK**.
7. Click the **Alerts** tab to verify that the replication agent was uninstalled successfully.
8. In the Host view, click **Refresh** to remove the replication agent from the list.

Note: If you have problems uninstalling a replication agent, check whether the option "Allow replication agent to be deployed to this host" is cleared in the properties of the machine hosting the agent. For more information, see "Replication and migration problems."

Change the credentials of an agent

Agent Management enables you to change the credentials of deployed Windows and UNIX replication agents.

Notes:

- When deployed Windows replication agents are using the credentials of the VFM server and you change the credentials of the server, you must update the credentials of all the Windows replication agents deployed by the server. If you do not update passwords, you cannot perform replication and monitoring on already deployed agents.
- When you change the default agent credentials for Windows or UNIX replication agents, you must update the credentials of deployed agents that are using the default credentials. If you do not update the credentials, you cannot perform replication and monitoring on the agents.
- If you installed a monitoring agent and you change the credentials of the agent, you must update the credentials of all the replication agents deployed by the monitoring agent.
- You can change the credentials for one agent or you can change the credentials for multiple agents (replication and monitoring) and the VFM server.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click a replication agent and click **Change Credentials**.

Run Using	Description
The default replication agent user name	Run the replication agent using the default replication agent credentials. To specify default replication agent credentials for all Windows and UNIX replication agents that

Run Using	Description
	you deploy, see "Set Default credentials."
One of the credentials in the credentials vault	VFM stores all the credentials that are used or were used to run replication and monitoring agents or the VFM server. Select one of these credentials to run the replication agent.
The credentials specified below	Before you specify new credentials for the replication agent, make sure that the account has the correct permissions and privileges.

Change the listening port of an agent

VFM uses the TCP port 6002 for replication agents and TCP ports 6001 and 6005 for the VFM server. You have the option of changing the ports. Make sure that the port you select enables the replication agent to communicate with the VFM server through firewalls.

1. Click **Control Panel**.
2. In the Control Panel View, click **Agent Management**.
3. Select one of the following options:

Option	Step
To change the listening port of multiple replication agents	1. In the Host view, select the replication agents . 2. In the Properties view, select "Multiple Values" in the **Port** box and type the new port number. 3. Click **Save**.
To change the listening port of a replication agent	1. In the Host view, select a replication agent. 2. In the Properties View, type the new port number. 3. Click **Save** .

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, you can deploy or install an agent on a proxy machine. The UNIX machine where you install a replication agent must be running RedHat Enterprise 4, or Solaris 10 operating system.

Note: For instruction about how to deploy Windows and UNIX replication agents, see "Manually deploy a replication agent."

Prerequisites

Before you install a UNIX replication agent, review the list of requirements.

RedHat Enterprise

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

To....	Then do this.....

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\ ReplicationAgent\Linux* folder.
	• Extract the installation package from the VFMInstall.zip or VFMInstall.exe file.
	2. Copy the installation package vfmra-<*product version number*>-<*product build number*>.i386.rpm, where <product version number> is the current version of VFM and <*product build number*> is the current build of VFM, to the UNIX machine that will host the replication agent.
	3. **View details of the installation package**. In these instructions, product version "99.9" is used and build number "1" is used. **You must substitute these numbers with the current version and build of VFM.** To view details of the installation package without installing, type the following command:
	`rpm -qilp vfmra-99.9-1.i386.rpm`
	4. **Install the replication agent**. Type the command listed below to install the UNIX replication agent This determines where the binaries and log files of the replication agent are kept. In the following example, the binaries are installed in "/opt/URA/bin" and the logs and registry are installed in "/opt/URA/logs".
	`rpm -ihv vfmra-99.9-1.i386.rpm`
	5. **Configure the replication agent**. After you install the replication agent, you must configure it. The following example shows how to configure the agent.
	`/etc/init.d/ura configure`
	You are prompted to enter the following information:
	• The TCP port where the replication agent and the VFM server listen.
	• The name of your VFM servers.
	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.

To....	Then do this.....

	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`

Stop a replication agent	Type the following command to stop the replication agent: `/etc/init.d/ura stop`

Remove a replication agent	1. Stop the replication agent.
	2. Type the following command to remove the replication agent:
	`rpm -evv vfmra-99.9-1`
	3. Type the following command to remove the directory /opt/URA . This directory contains the configuration files created by the agent.
	`rm -Rf /opt/URA`

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 | grep vfm

Solaris

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

The package name is vfmra-<*product version number*>-<*product build number*>.pkg, where <*product version number*> is the current version of VFM and <*product build number*> is the current build of VFM. **In these instructions, product version "99.9" is used and product build number "1" is used. You must substitute these numbers with the current version and build of VFM.**

To...	Then do this...
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\ReplicationAgent\Solaris folder. • Extract the installation package from the VFMInstall.zip or VFMInstall.exe file. 2. Copy the installation package to the UNIX machine that will host the replication agent.

To...	Then do this...

	3. **Unzip the package file**. At a shell prompt, type the following command. **Make sure you substitute "99.9" with the current version of VFM and "1" with the current build of VFM.**
	`gunzip vfmra-99.9-1.pkg.gz`
	4. **Install the package**. Type the following command:
	`pkgadd -d ./vfmra-99.9-1.pkg`
	The replication agent is installed in the */opt/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`

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
Search for the replication agent package	Type the following command: `pkginfo | grep NTAP`
Stop the replication agent	Type the following command: `/etc/init.d/ura stop`
Remove a replication	1. Stop the replication agent. 2. Type the following command to remove the replication agent:

To...	Then do this...

agent	`pkgrm NTAPura`
	
	3. Type the following command to remove the directory /opt/URA. This directory contains the configuration files created by the agent.
	`rm -Rf /opt/URA`

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 `/opt/URA`. You can install a replication agent to another directory. **Note:** If you install a UNIX replication agent to another directory and VFM upgrades the replication agent, the new replication agent is installed in the default directory `/opt/URA`. 1. Complete Steps 1 through 3 under "To install a replication agent." 2. **Install the package**. Type the following command: # pkgadd -a none -d ./vfmra-99.9-1.pkg You will see the following prompt: Enter path to package base directory [?,q] 3. **Type the location of the directory**, for example: /opt/Software/URA 4. Continue with Step 5 under "To install a replication agent."

Notes:

- The replication agent can be installed on the source, destination, or a proxy machine.
- The following examples illustrate NFS data transfer:

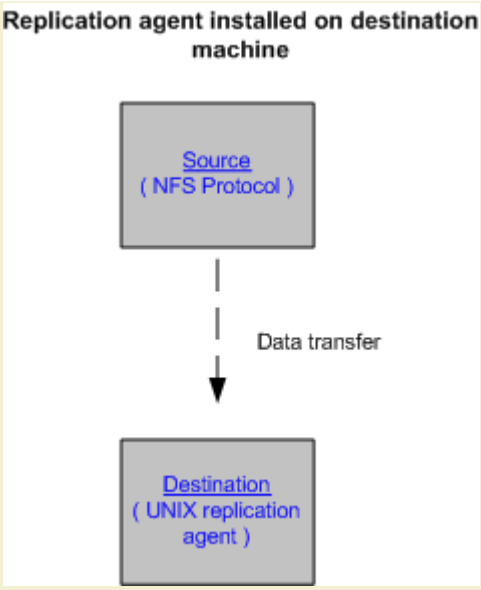
Replication agent installed on destination machine
Source
(NFS Protocol)
Data transfer
Destination
(UNIX replication agent)

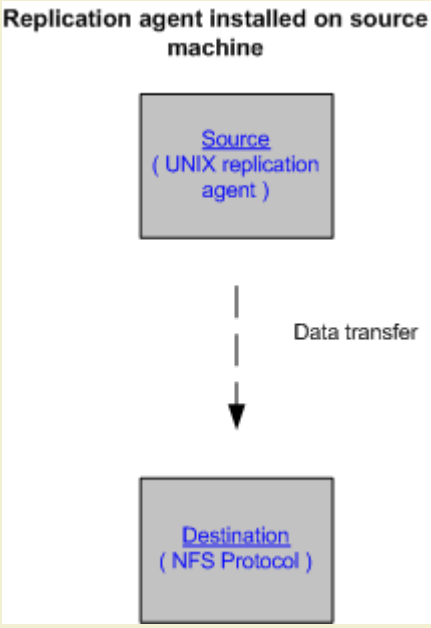
Replication agent installed on source machine
Source
(UNIX replication agent)
Data transfer
Destination
(NFS Protocol)

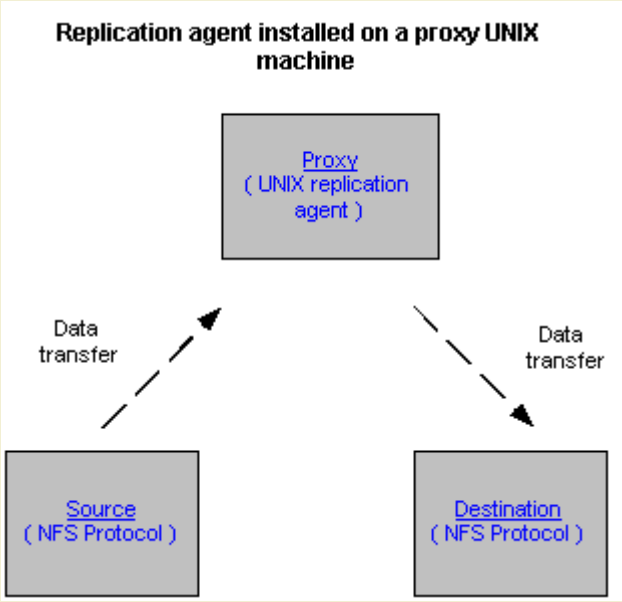

Purge an agent

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. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click the agent you want to remove from the list and click **Purge**. The agent is removed from the database and from the list of agents.

Manage monitoring agents

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

Distributed monitoring

- Consider distributing monitoring agents based on your network topology. You can select another monitoring agent that you installed at another location and configure it to work with your VFM server.
- 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 reenable policies using the agent on that host. To reenable these policies, you must assign the newly installed agent to run the policies by reselecting the desired agent to be used by the policy through the properties dialog box for the policies.

Checklist

VFM enables you to install and manage distributed monitoring agent.

- To install an agent on a machine, see "Install a monitoring agent."
- To update credentials of a monitoring agent, see "Change the credentials of a monitoring agent."
- To change the listening port of a monitoring agent, see "Change the listening port."
- To assign a monitoring agent to run a policy, see "Assign a monitoring agent to a policy."
- To specify where you want to store the configuration files of a monitoring agent and the policies assigned to the agent, see "Specify a backup location."
- To remove a monitoring agent from the VFM database, see "Purge an agent."
- To view information about a deployed agent, see "View the status of an agent."

Install a monitoring agent

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

Use the VFM media to install a monitoring agent.

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.

Change the credentials of a monitoring agent

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

Notes:

- When you upgrade to a newer version of VFM, you must upgrade all installed monitoring agents.
- You specify the credentials for a monitoring agent during the installation process of the agent. This account information is used when the monitoring runs.
- You can change the credentials for one agent or you can change the credentials for multiple agents (replication and monitoring) and the VFM server.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click the agent and click **Change Credentials**.
4. Select one of the following options:

Run Using	Description
One of the credentials in the credentials vault	VFM stores all the credentials that are used or have been used to run replication and monitoring agents or the VFM server. Select one of these credentials to run the monitoring agent.
The credentials specified below	Before you specify new credentials for the monitoring agent, make sure that the account has the correct permissions and privileges.

5. Click **OK**.

Change the listening port

VFM uses the TCP ports 6001 and 6005 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 or monitoring agent to communicate through firewalls with the VFM client and all the replication agents deployed by the server or monitoring agent.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, click the VFM server or a monitoring agent.
4. In the Properties view, type the new TCP port and click **Save** .

Assign a monitoring agent to a policy

The VFM server runs policies. You can select a monitoring agent installed on another machine to run a policy.

1. Perform one of the following actions:
   - In the VFM 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 Edit.
2. Click **General**.
3. On the **Use the Monitor on** list, select the host name of the monitoring agent where monitoring is performed. The list displays the host names of the VFM server and all the VFM monitoring agents that are installed and configured to communicate with the VFM server.
4. Click **OK**.

Note: For information about assigning and unassigning monitoring agents, see "Mange monitoring agents."

Specify a backup location

When you schedule backups of the application, the configuration files of the application and all installed monitoring agents are saved. The configuration files of monitoring agents are saved in folders under the location you specify to save the application configuration files. You can specify a different location to store snapshots of the state of a monitoring agent and the policies assigned to the agent.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, click a monitoring agent to show its properties.
4. In **Snapshot Save location**, type the UNC path of the location where you want to save the configuration files of the monitoring agent or click **Browse** to find a location. This should be a location that is different than the location you specified to save the application.
5. Click **Save**.

Note: For snapshot location requirements, see the Notes for specifying a backup location for the VFM server.

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. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click the agent you want to remove from the list and click **Purge**. The agent is removed from the database and from the list of agents.

Manage the application server

This section provides information about how you can change the VFM server credentials, the listening port of the server, specify a location to back up the server configuration files, and stop and restart the server.

Checklist

- To update the account information for the VFM server, see "Change server credentials."
- To change the listening port for the VFM server, see "Change the server listening port."
- To specify the location where you want to back up the VFM configuration files, see "Specify a backup location."

Change the 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.

Notes:

- When you change the password on the account where the VFM server is installed, you must update the credentials of the server. If you do not update credentials, the server ceases to function correctly because Windows security prevents the service from authenticating with the system.
- After updating the credentials of the server, you must update the password of all the replication deployed by the server. See "Change credentials of an agent."
- You can change the credentials for the VFM server only or you can change the credentials for the server and multiple agents (replication and monitoring).

To change the VFM server credentials

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, right-click the VFM server and click **Change Credentials**.
4. Select one of the following options:

Run Using	Description
One of the credentials in the credentials vault	VFM stores all the credentials that are used or were used to run replication and monitoring agents or the VFM server. Select one of these credentials to run the server.
The credentials specified below	Before you specify new credentials for the VFM server, make sure that the account has the correct permissions and privileges.

5. Click **OK**.

Note: During the process of changing the server credentials, the VFM server service is stopped and restarted.

Change the listening port

VFM uses the TCP ports 6001 and 6005 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 or monitoring agent to communicate through firewalls with the VFM client and all the replication agents deployed by the server or monitoring agent.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, click the VFM server or a monitoring agent.
4. In the Properties view, type the new TCP port and click **Save** .

Specify backup location

Specify the location where you want to save the VFM configuration files.

1. Click **Control Panel**.
2. In the Control Panel view, click **Agent Management**.
3. In the Host view, click the VFM server to show its properties.
4. In **Snapshot Save location**, type the UNC path of the location where you want to save the configuration files of the VFM server or click **Browse** to find a location.
5. Click **Save** .

Notes:

- The location selected for the application snapshot must be writable by the account under which MSDE is running and the account under which the VFM server is installed. By default, MSDE runs as LocalSystem. Share permissions and directory permissions configured on the snapshot save location must allow both LocalSystem and the account under which the VFM server is installed to write to the directory.
- If the destination directory already exists, it must be empty.
- Application snapshots cannot be saved on a mapped network drive.
- Store the exported files in an appropriate location where you can access them if the machine hosting the VFM server fails.
- For instructions about how to schedule backups or manually back up the application, see "Export configuration files".

Chapter 3 Run Batch Files

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 monitoring agent to which the policy is assigned. Batch files related to reports are run by the VFM server.
- VFM uses the directory where a batch file is located as the default/current directory where the command is executed. If a batch file is specified through a UNC path, the TMP directory of the server executing the policy is used as the default/current directory.
- VFM checks the return value from the batch script that is run before failover. If the batch script returns failure, VFM does not continue the failover process. This allows customized checking to be performed before the actual failover takes place.
- If the UNC path to the batch file contains spaces, you must enclose the entire path in quotation marks. You must follow the same rules that apply when running commands from a Windows command prompt.
- Scripts run under the following conditions:
    - The script specified in **Run Batch Before Replication** runs *before* the following migration actions:
        - Scan for potential problems copying the data (initial phase)
        - Perform a baseline copy (initial phase)
        - Copy the data (incremental phase)
        - Do a final replication prior to deleting the source data (final phase)
        - Delete the source data (final phase)
    - The script specified in **Run Batch After Replication** runs after the migration actions listed previously.
    - The script specified in the initial phase runs before the actions in the initial phase run, and the script specified in the final phase runs after the actions in the final phase are completed.
    - When the script specified in **Run Batch Before Replication** or the script **specified in the initial phase** runs without errors, the script specified in **Run Batch After Replication** or the script **specified in the final phase** is executed.
    - When the script specified in **Run Batch Before Replication** or the script **specified in the initial phase** returns an error, the replication is not executed and the script specified in **Run Batch After Replication** or the script **specified in the final phase** is not executed.

- VFM provides example scripts in the installation directory under `Examples\Scripts.. 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 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.
- Review the following list of some common problems you might encounter when working with scripts.
    - The path to the batch file must be wrapped in quotation marks if that path contains spaces. If you browse to the script location, VFM automatically enters the quotation marks. If you enter the path to the batch file manually , you must enter the quotation marks.
    - Batch files run with their working directory set to the directory containing the batch file when they are specified by a local path (for example, `c:\bats\myfile.bat`). When the batch file is referenced by a UNC path, its working directory is the system32 of the server or monitoring agent that runs the batch file.
    - Any file invoked by VFM must be a batch file. It is possible to have a top-level batch file invoke other tools, such as `jscript` files or other executables.
    - The arguments passed to the script are slightly different than expected. It can be very helpful to have your script log into a text file the time it was invoked along with the arguments it received.
    - Multiple copies of a script might invoke simultaneously, such as pre- and post-migration task batch files, when tasks are running at the same time. This can also occur in disaster recovery and replication scripts.
    - VFM uses the exit value of the script to assess whether the script succeeded or failed. Scripts must set the return value correctly to signal either success or failure.
    - If your script is not producing the desired results, you can insert different exit codes throughout the script to determine where exactly the script is exiting. You can also insert logging statements showing progress through the script to also trace execution.

- Scripts must not prompt for user input. There is no console available to display the script prompt at run time. Scripts that prompt for input and the policy executing the script will hang indefinitely and might cause the application to hang. This is true for any prompt, whether in an error path or in a normal execution path in the script.
- VFM *does not validate the path* to the script you specify. It is your responsibility to ensure that the path is accurate and accessible to the VFM server or monitoring agent that invoke the script.

Batch files

- Batch files enable you to customize policies, such as sending an email, posting an event in the Windows event log, or sending an SNMP trap in response to a failure.
- Scripts are executed by the VFM server or monitoring agent to which a policy is assigned. Report-related scripts are run by the VFM server.
- You can find example scripts in the installation directory under the `Examples\Scripts` folder.
- Before you execute a script, review the list of common problems you might encounter when working with scripts.

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.	1. <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.	1. <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 `<first user argument`	0 => success all others => error

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
			>	
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 <first user argument>	0 => path matches criteria 1 => path does not match criteria all others => error
Archival migration policies	Pre-Design phase batch script	"Run batch before migration task" field on the **Task Scripting** tab of the policy properties.	3. Path of source 4. Path of destination 5. *<first user argument>*	0 => success all others => error
Archival migration policies	Post-Finalization phase batch script	"Run batch after migration task" field on the **Task Scripting** tab of the policy properties.	1. Path of source 2. Path of destination 3. copy=[pass|fail|skip] 4. updatedfs=[pass|fail|skip] 5. delete=[pass|fail|skip] 6. *<first user argument>*	0 => success all others => error
Archival migration policies	Pre-replication script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Archival migration policies	Post-replication script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Client disaster recovery	Pre-replication script	"Run Batch Before Replication" field on the **Replication Script** tab of the policy properties.	1. UNC path of link (for Replication Policies, this is the destination link) 2. UNC path of source 3. UNC path of destination 4. *<first user argument>*	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.	1. UNC path of link (for Replication Policies, this is the destination link) 2. UNC path of source 3. UNC path of destination 4. The time that the it took to perform the operation 5. Result of the	0 => success all others => error

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
			replication: • **Pass** - Operation succeeded • **Fail** - Operation encountered errors • **Cancel** - Operation canceled by the user • **Expired** - Operation was not completed in the scheduled time window • Number of files copied • Number of bytes copied • Number of files already in sync • Number of bytes already in sync • *<first user argument>*	
Client disaster recovery policies	Pre-failover script	"Run batch before failover" field on the **DR Script Settings** tab of the policy properties.	1. Policy name 2. [not used, will be empty] 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	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 5. *<first user argument>*	0 => do not fail over 1=>do fail over all others => error; logged in event log.

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
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. Qtree name (when applicable) 5. *<first user argument>*	0 => do not fail over 1=>do fail over all others => error; logged in event log.
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	Pre-Design phase batch script See Archival migration policies script	See Archival migration policies script.	See Archival migration policies script.	See Archival migration policies script.
Migration policies	Post-Finalization phase batch script See Archival migration policies script	See Archival migration policies script.	See Archival migration policies script.	See Archival migration policies script.
Migration policies	Pre-replication script See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script
Migration policies	Post-	See Client disaster	See Client disaster recovery	See Client disaster

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
	replication script See Client disaster recovery script	recovery script	script	recovery 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. *<first user argument>*	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. *<first user argument>*	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 schedule group or report source group	Pre-data gathering	"Run batch before gathering data" field on **Scripts** of the Report Schedule Group or Report Source Group properties.	1. Report source group name 2. *<first user argument>*	0 => success all others => error
Report schedule group or report source group	Post-data gathering See Pre-data gathering script	"Run batch after gathering data" field on **Scripts** of the Report Schedule or Report Source 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
Report wizard	Pass report	"Script" field on	1. The path to a TMP file	0 => success

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
	output to script	the **Report Output** dialog box of the Reports wizard.	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. 2. The data collection process shows the following states: • **OK** - All the data required for the report was successfully collected. • **Warning** - Not all the data required for the report was collected. The information passed to the 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. 3. *<first user argument>*	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 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	Pre-failover script See Client	See Client disaster recovery script	See Client disaster recovery script	See Client disaster recovery script

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
	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	1. The name of the policy 2. The name of the link 3. The online or active link target 4. The offline link targets 5. *<first user argument>*	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
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

Policy/Feature	Type of Script	Where to run script	Arguments	Exit codes
Storage load balancing policies	Pre-Design phase batch script See Archival migration policies script			
Storage load balancing policies	Post-Finalization phase batch script See Archival migration policies script			
Storage load balancing policies	Pre-replication script See Client disaster recovery script			
Storage load balancing policies	Post-replication script See Client disaster recovery script			

Chapter 4 VFM Server Availability

This section 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 is your VFM configuration.
- Whether you have storage policies that you want to run from locations that are different from where the VFM server is installed.
- Where the DFS roots that are managed by VFM are located.

Selecting a method

Select a method for backing up and restoring VFM server configuration files.

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

  You can schedule snapshots of the state of the VFM server.

  - When the VFM server is not available, you can restore the latest snapshot on another machine where VFM is installed and have it take over as the VFM server.
  - 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 back up and restore VFM server configuration files**

  You can manually back up VFM configuration files.

  - When the VFM server and VFM console are not available, you can use the export utility to back up the VFM configuration information.
  - You can restore 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 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 servers active on multiple servers managing policies referencing a common set of DFS resources is not supported.

Export/Import configuration files

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

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

Planning for VFM availability

The backup and restore feature enables you to maintain a backup VFM server if the primary VFM machine fails.

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

Notes:

- When backing up 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 backing up VFM information, you must save copies of any batch files you configured to run as part of your policies. You must make these files available to the backup server before it is activated. If you install the batch files in a different location on the backup server from where they were on the primary server, you must edit your policies to reference the batch files in their new locations.
- 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 server configuration files you backed up 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 back up your VFM configuration information when the VFM console in 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 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 on which the VFM server is installed becomes unavailable. Before you export your configuration files, read "Export/import considerations."

To back up VFM configuration files automatically

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

Notes:

- When you specify a snapshot location for the VFM server in the Agent Management tool, VFM uses the location specified in Agent Management instead of the location you specify when scheduling application snapshots.
- For additional requirements, see the Notes at the end of this topic.

To back up VFM configuration files manually

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 to the location at which you want to save the configuration files, or click **Browse** to locate the machine in your network. You cannot save the snapshot on a mapped network drive.
3. Type the name of the folder where the files are saved.
4. Click **OK**.

Notes:

- The location you type at which to manually back up the VFM configuration files overrides the snapshot location you specify for the VFM server in the Agent Management tool.
- For additional requirements, see the Notes at the end of this topic.

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** subfolder.
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:

- The location selected for the application snapshot must be writable by the account under which MSDE is running and the account under which the VFM server is installed. By default, MSDE runs as LocalSystem. Share permissions and directory permissions configured on the snapshot save location must allow both LocalSystem and the account under which the VFM server is installed to write to the directory.
- The destination directory must exist and must be empty.
- Application snapshots cannot be saved on a mapped network drive.
- Store the exported files in an appropriate location where you can access them if the machine hosting the VFM server fails.
- For problems you might encounter, see "Problem taking snapshots of the application."

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.

Prerequisites

- You must stop all VFM processes, including any remote clients accessing the VFM server, before you can perform an import operation.
- When the SQL service is running on the Local Network Account on the backup machine, perform one of the following tasks:
    - Save the exported snapshot on the local drive of the backup machine before you import the snapshot to VFM.
    - Make sure that the SQL service is running under an account that has full privileges across all domains.

To import VFM configuration files

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

6. Start the VFM console and validate that the VFM configuration is correct.
7. Make sure that policies are assigned to the proper server or monitoring agent.

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:
   a. Run the export operations.
   b. On the machine where the VFM server is installed, start a command prompt and change to the VFM installation `Tools\Export` subfolder.
   c. Type `server disable` and press ENTER to stop and disable the VFM server on the backup machine. The primary and backup VFM servers **cannot be active simultaneously**.
2. Perform the following actions on the primary VFM machine:
   a. Ensure that VFM is operational.
   b. 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` subfolder.
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 5 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 and 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`

  `pageant.exe`

  Important: If the 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 10 and RedHat Enterprise 4.
- A NetApp filer must have the SecureAdmin™ package installed. To verify that a filer has this package, look for the directory named `/etc/sshd`.

  Note: Data ONTAP 7.0 or later includes built-in SecureAdmin functionality with SSL protocol.

- When VFM uses SSH to communicate with UNIX machines and filers, SSH must be configured without a password or passphrase.

  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/

Checklist

The following checklist outlines the steps required 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.

- Set default remote SSH shell credentials for all new filers and UNIX machines added to the Physical View.
- Set remote SSH shell credentials for a filer in the Physical View
- Set remote SSH shell credentials for a UNIX machine in the Physical View

	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 public keys onto each UNIX machine known to VFM.	3. "Cache the keys"
		4. Add the private keys into memory on the VFM machine.	4. "Add private keys to pageant"
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	3. "Generate a key-pair"

	Environment	Step	Instructions
		monitoring agent.	
		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 public keys on the VFM client onto the UNIX machines known to VFM.	7. "Cache the keys"
		8. Cache the public keys on the VFM server onto the UNIX machines know to VFM.	8. "Cache the keys"
		9. Cache the public keys on each VFM monitoring agent onto the UNIX machines known to VFM.	9. "Cache the keys"

	Environment	Step	Instructions
		10. Add the private keys into memory on the VFM client machine.	10. "Add private keys to pageant"
		11. Add the private keys into memory on the VFM server machine.	11. "Add private keys to pageant"
		12. Add the private keys into memory on each VFM monitoring agents.	12. "Add private keys to pageant"
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 CIFS.	2. "Copy the public key to a filer using CIFS"
		3. Cache the public keys onto the CIFS filers known to VFM.	3. "Cache the keys"
		4. Add the private keys into memory on the VFM machine.	4. "Add private keys to pageant"
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"

	Environment	Step	Instructions
		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 public keys on the VFM client onto the NFS filers known to VFM.	7. "Cache the keys"
		8. Cache the public keys on the VFM server onto the NFS filers know to VFM.	8. "Cache the keys"
		9. Cache the public keys on each VFM monitoring agent onto the NFS filers know to VFM.	9. "Cache the keys"
		10. Add the private keys into memory on the VFM client machine.	10. "Add private keys to pageant"

	Environment	Step	Instructions
		11. Add the private keys into memory on the VFM server machine.	11. "Add private keys to pageant"
		12. Add the private keys into memory on each VFM monitoring agents.	12. "Add private keys to pageant"

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.

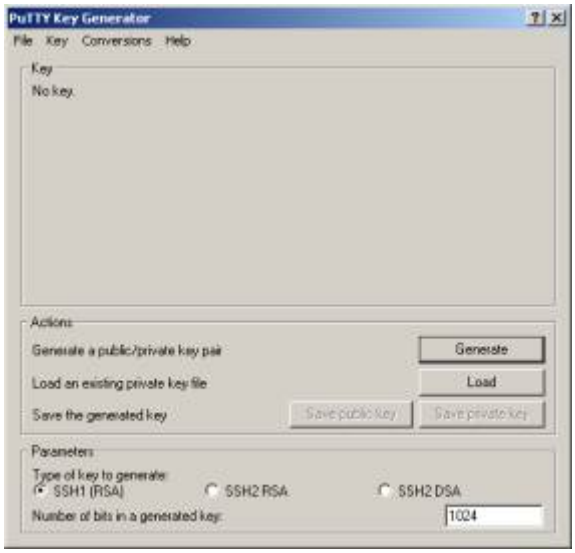

2. Select SSH1 (RSA) as the type of key to generate.
3. Click **Generate**. An SSH1 RSA key with 1024 bits is generated.

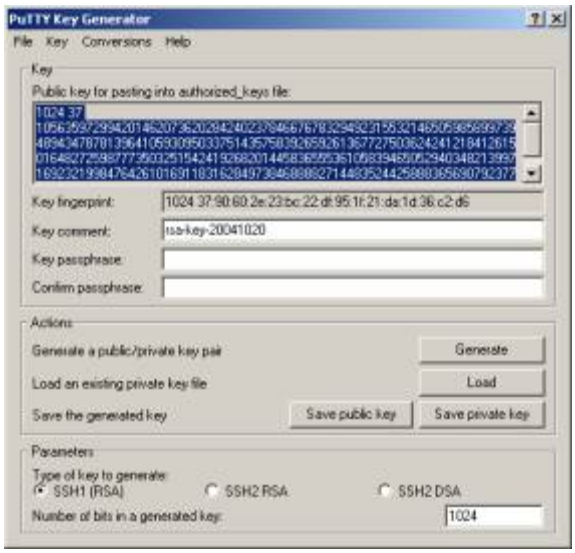

4. Click **Save Private Key** and specify a directory and filename. Remember where this file is saved and the name of the file. You will have to provide this information to VFM. See "Set default remote shell credentials."
5. When the message "Are you sure you want to save this key without a passphrase 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 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:

   Using a root account

   - If the **root directory does not exist**, create it then make it your current directory. Type the command:

     ```
     mkdir root
     ```

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

     ```
     cd root
     ```

Using a non-root account

- If the **directory does not exist**, create it then make it your current directory. Type the command:

  ```
  mkdir <non-root account name>
  ```

  where `<non-root account name>` is the name of the directory you want to create. The name of the directory must be the same as the name of the account you are using to create the directory.

- If the **directory exists**, make it your current directory. Type the command:

  ```
  cd <non-root account name>
  ```

  where `<non-root account name>` is the name of the directory. The name of the directory is the same as the name of the account you are using to configure SSH.

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` directory. Where `<filer name>` is the name of the filer you want to copy the public key.
3. Perform one of the following actions:

   Using a root account

   - If the **root directory does not exist**, create it then make it your current directory. Type the command:

     ```
     mkdir root
     ```

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

     ```
     cd root
     ```

Using a non-root account

- If the **directory does not exist**, create it then make it your current directory. Type the command:

  ```
  mkdir <non-root account name>
  ```

  where `<non-root account name>` is the name of the directory you want to create. The name of the directory must be the same as the name of the account you are using to create the directory.

- If the **directory exists**, make it your current directory. Type the command:

  ```
  cd <non-root account name>
  ```

  where `<non-root account name>` is the name of the directory. The name of the directory is the same as the name of the account you are using to configure SSH.

4. Create a directory name .ssh. Type the command:

   ```
   mkdir .ssh
   ```

5. Create a file named "authorized_keys."
6. Open the authorized_keys file and paste the public key at the end of the file. Make sure that the public key appears in one line. The key should not be separated by line breaks.
7. Save and close the file.
8. 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.

Notes:

- The `plink.exe` program must be executed using the account that runs VFM.
- 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.npb.com`).
   - `users` is the account name used to login to the UNIX machine or filer.

   Note: You are prompted to add the key to the cache of the account that runs VFM.

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

Add private keys to pageant

After you cache the public keys, you must run the "pageant.exe" program on the VFM client and server machines. This program copies the private keys into memory. The private keys are decoded so that VFM can communicate with a filer or UNIX machine without a password.

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

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

   ```
   pageant <path>
   ```

   `path` is the path to the private key and the name of the file containing the private key.

   Note: The "pageant.exe" program must be executed using the account that runs VFM.

2. Verify that you are not prompted for a password. At the command prompt, type the following command:

   ```
   plink <host> -l <users> ls
   ```

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

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

a. Get the following packages from sunfreeware:
   - `openssh-3.9pl –sol8-sparc-local.gz`
   - `openssh-0.9.7d-sol8-sparc-local.gz`
   - `zlib-1.2.1-sol8-sparc-local.gz`
   - `libgcc-3.3-sol8-sparc-local.gz`

b. `Copy the packages to the UNIX machine.`

Step 2 Install the packages

a. Change to the directory where you copied the packages.

b. Type the following commands:

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

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.

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

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

b. 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
```

c. 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 Enterprise 3.

Step 1 Get the latest OpenSSH server packages

a. Copy the appropriate packages for RedHat Enterprise 3 to the UNIX machine.
b. 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

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

```
Chkconfig –level 35 sshd on
```

b. 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.
- Add the private keys into memory

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 6 Use RSH Shell Credentials

Normally, VFM communicates with non-Windows machines using RSH. RSH must be enabled as root on the non-Windows machines.

- 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 Linux 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 the file *.rhosts* in the roots home directory or create it if it does not exist. Add the following line:

   ```
   + +
   ```

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

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

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

   Note: The file *.rhosts* must have permission 600 and must be owned by root.
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 <your host here> -l root "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 Solaris 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 <your host here> -l root "ls -la"
```

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

Appendix

VFM Guides

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*	VFMAdvConfig.pdf	• Explains how to use VFM to configure security permissions • Explains how to manage replication and monitoring agents • Explains how to run batch files in policies and reports • Explains how to make the VFM server highly available
2.	*Create and Manage Namespaces*	VFMNamespace.pdf	Explains how to create a DFS namespace (roots, links, and link targets), how to quickly populate the namespace, make the namespace highly available, create multiprotocol namespaces, virtualize user home directories, and back up the namespace.
3.	*Getting Started*	VFMGettingStarted.pdf	A tutorial to get you started using VFM. It explains how to install VFM and steps you through the process of creating a logical namespace, making the namespace highly available, making data highly available, transferring CIFS and NFS data, creating multiprotocol namespaces, virtualizing user home directories, and backing up data and the namespace.
4.	*Make Data Highly Available*	VFMMakeDataAvailable.pdf	Explains how to make data highly available by creating client, filer, and server disaster recovery policies
5.	*Manage Physical Resources*	VFMPhysicalResources.pdf	Explains how to use VFM to Work with NetApp filers, set replication properties, and work with Windows and UNIX machines

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

Glossary

A

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

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

C

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

Client: The user interface component of the application.

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

D

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

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

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

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

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

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

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

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

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

E

Export: SEE NFS export

F

Fault-tolerant DFS: SEE Domain-based DFS root

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

FRS: Microsoft File Replication Service.

G

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

K

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

L

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

M

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

Migration task: Identifies the source and destination.

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

N

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

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

P

Pool: A group of storage locations.

Primary server: An active application server or monitoring agent.

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

by the DFS root.

Q

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

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

R

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

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

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

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

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

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

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

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

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

S

Secondary target: SEE Standby target.

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

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

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

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

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

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

T

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

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

U

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

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

W

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

Index

A

Application
  configure permissions 1
  export configuration files 46, 48
  import configuration files.....47, 48
Application server
  change credentials 29
  change listening port 27, 30
  specify backup location.....30
Application server 29
Application server 29
Application server 30

B

Back up configuration files 46
Back up location of server.....28, 30
Batch files.....33, 35

C

Configuration files.....46, 47, 48
Credentials.....15, 18, 27, 29

D

Data transfer rate 13
Delete from database 25, 29
Deploy a replication agent.....14

E

Export configuration files.....46, 48

F

Failback 48
Failover.....47, 48
Fault tolerance 46

I

Import configuration files.....47, 48
Install an agent.....19, 26

K

Key-pair
  generate 54
Key-pair.....54

L

Linux 65
Listening ports.....19, 27, 30
Load balancing 17

M

Monitoring agent
  assign to policies 25, 28
  backup location 28
  change credentials 27
  install an agent..... 26
  purge from database 25, 29
  view status..... 13
Monitoring agent 28

N

Network communication 19
Network ports 19
NFS data transfer..... 19

O

OpenSSH 59, 62

P

Pause an agent 16
Permissions..... 1
Policies 25
Ports..... 19
Private key
  cache keys 58
Private key 58
Proxy machine 13
Public key
  cache keys 58
  copy to filer 56
  copy to UNIX machine 55
Public key 58

R

RedHat.....19, 62
Remove from database 25, 29
Replication agent
  availability..... 17
  change credentials 15, 18
  change listening port 19
  deployment of..... 14
  install an agent..... 19
  pause an agent 16
  purge from database 25, 29
  requirements..... 14
  restart an agent 16
  resume an agent..... 16
  stop an agent..... 16
  uninstall an agent..... 17
  upgrade an agent 16
Replication agent 14
Replication agent 15
Replication agent 17
Replication group 17
Restart an agent 16
RSH 62, 65, 66

S

Scripts.....35
Security.....1
Set throttling limits.....13
Shell credentials 62
Solaris.....19, 59, 66
SSH.....54, 55, 56, 57, 58, 62
Stop an agent 16

U

Uninstall an agent.....17, 19, 26
UNIX 19
Upgrade a replication agent.....16, 19