



**EMC RepliStor for Microsoft Windows**  
Version 6.1

**ADMINISTRATOR'S GUIDE**

**P/N 300-002-464**  
**REV A01**

**EMC Corporation**  
*Corporate Headquarters:*  
Hopkinton, MA 01748-9103  
1-508-435-1000  
[www.emc.com](http://www.emc.com)

Copyright © 2003 - 2005 EMC Corporation. All rights reserved.

Published November, 2005

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

### **Trademark Information**

EMC<sup>2</sup>, EMC, EMC ControlCenter, AlphaStor, ApplicationXtender, Catalog Solution, Celerra, CentraStar, CLARAlert, CLARiiON, ClientPak, Connectrix, Co-StandbyServer, Dantz, Direct Matrix Architecture, DiskXtender, Documentum, EmailXtender, EmailXtract, HighRoad, Legato, Legato NetWorker, Navisphere, OpenScale, PowerPath, RepliStor, ResourcePak, Retrospect, Smarts, SnapshotServer, SnapView/IP, SRDF, Symmetrix, TimeFinder, VisualSAN, VSAM Assist, Xtender, Xtender Solutions, and where information lives are registered trademarks and EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, EMC Storage Administrator, Access Logix, ArchiveXtender, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CLARevent, Codebook Correlation Technology, Common Information Model, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, DiskXtender 2000, EDM, E-Lab, EmailXaminer, Engenuity, eRoom, FarPoint, FLARE, FullTime, Global File Virtualization, Graphic Visualization, InfoMover, Invista, MirrorView, NetWin, NetWorker, OnAlert, Powerlink, PowerSnap, Rainfinity, RecoverPoint, RepliCare, SafeLine, SAN Advisor, SAN Copy, SAN Manager, SDMS, SnapImage, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, UltraPoint, Viewlets, VisualSRM, and WebXtender are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

<b>Preface</b> .....	xv
<b>Chapter 1 Introduction</b>	
RepliStor Software Overview.....	1-2
RepliStor Software Components .....	1-3
RepliStor Replicating Functions .....	1-4
Failover Capabilities.....	1-4
Alias Failover.....	1-5
VSS-compliant Shadow Copies .....	1-7
<b>Chapter 2 Getting Started</b>	
Before You Begin.....	2-2
Usage Considerations .....	2-2
Common Mistakes to Avoid .....	2-2
Active Directory Support.....	2-4
Starting the RepliStor Client.....	2-6
The RepliStor Client Elements.....	2-7
Starting a RepliStor Site .....	2-12
Attaching to Remote RepliStor Sites .....	2-14
Creating Specifications.....	2-15
Global Exclude Specifications.....	2-16
File/Directory Specifications .....	2-18
Share Specifications.....	2-28
Registry Specifications .....	2-30
Macintosh Specifications .....	2-36
Creating Specifications for Circular Mirroring.....	2-37
Synchronizing Specifications .....	2-39
Synchronizing Specifications .....	2-39

	Specifying an E-mail Address for Synchronization Notifications .....	2-41
	Detaching from a Site .....	2-42
<b>Chapter 3</b>	<b>Configuring, Taking, and Managing Shadow Copies</b>	
	Overview .....	3-2
	What is VSS? .....	3-3
	Allocating Shadow Storage .....	3-5
	Configuring Shadow Copies .....	3-7
	Taking Shadow Copies .....	3-10
	Scheduling Shadow Copies .....	3-12
	Microsoft Cluster Server (MSCS) Support .....	3-14
	Managing Shadow Copies .....	3-15
	Canceling a Shadow Copy in Progress .....	3-16
	Troubleshooting Shadow Copy Creation .....	3-18
<b>Chapter 4</b>	<b>Configuring Options</b>	
	Configuring Client Options .....	4-2
	Configuring RepliStor Options .....	4-3
	Configuring the Updating and Forwarding Options .....	4-5
	Configuring Updating and Forwarding .....	4-5
	Configuring Throttling Options .....	4-6
	Configuring Throttling Options .....	4-7
<b>Chapter 5</b>	<b>Configuring Failover</b>	
	Alias Failover .....	5-2
	Failover Setup Roadmap .....	5-4
	Configuring Failover on the Source System .....	5-5
	Configure a Failover .....	5-5
	Using Scripts .....	5-10
	Manually Forcing a Failover with Aliases .....	5-13
	Configuring RepliStor Software with Two NIC Adapters .....	5-14
	Configuring Manual Alias Activation .....	5-15
	Activating the Alias on the Target System after a Failover .....	5-16
	Configuring Failover across a WAN .....	5-17
	Failover Agent for WAN Failover .....	5-17
	Configuring Failover for Macintosh Volumes .....	5-20

<b>Chapter 6</b>	<b>Administering RepliStor Software</b>	
	Using the Windows Performance Monitor.....	6-2
	Sizing the Kernel Cache .....	6-3
	Using the Performance Tab in the Options Dialog Box .....	6-5
	Using the Processes Tab in the Options Dialog Box.....	6-7
	Administering RepliStor Sites .....	6-10
	Checking Site Status .....	6-10
	Correcting Site Problems .....	6-19
	Targeting a Remote Share Hosted by a NAS Device.....	6-23
	Registering a UNC Path.....	6-23
	Enabling or Disabling a UNC Path .....	6-27
	Modifying the Properties of a UNC Path .....	6-28
	Deleting a UNC Path .....	6-28
	Administering Specifications .....	6-30
	Modifying Specifications .....	6-30
	Deleting Specifications .....	6-31
	Enabling and Disabling Specifications .....	6-31
	Synchronizing Specifications .....	6-32
	Managing Specifications for Microsoft Cluster Server.....	6-36
	Turning Processes On and Off.....	6-41
	Stopping and Starting File Mirroring and Forwarding....	6-41
	Pausing File Updates.....	6-43
	Stopping and Starting File Attribute Updates .....	6-44
	Working with the Message Log.....	6-45
	Viewing the Message Log.....	6-46
	Marking Messages as Read .....	6-47
	Purging the Message Log .....	6-48
	Closing Application Files.....	6-48
	Using Tape Backups with RepliStor Software .....	6-49
	Reporting/Simulating Bandwidth	
	Between Source and Target.....	6-50
	Overview .....	6-50
	Reporting.....	6-51
	Simulation .....	6-54
<b>Chapter 7</b>	<b>Recovering Data</b>	
	Alias Failover Recovery.....	7-2
	Restoring the Original Configuration .....	7-2
	Setting the Alias to Fail Over to the Original Source.....	7-3

**Chapter 8 Commands**

Using RepliStor Commands.....	8-2
Syntax Conventions.....	8-3
Command Descriptions .....	8-4
aliasadd .....	8-4
aliasop .....	8-5
blockedsites .....	8-6
checkpoint.....	8-6
close .....	8-7
configexport/configimport.....	8-8
deletedatadir .....	8-8
forwarding.....	8-8
isidle .....	8-8
mirroring.....	8-9
modifyspec .....	8-9
numblockedfiles .....	8-10
pause.....	8-11
pauseattributes.....	8-11
report/simulate .....	8-11
siteforwarding.....	8-12
ShadowCopy.....	8-13
ShadowUtil.....	8-14
spec .....	8-15
spec/ globalexclude .....	8-18
spec/registry .....	8-19
spec/shares.....	8-21
specexport.....	8-21
specimport .....	8-22
specstatus.....	8-22
start .....	8-23
stop .....	8-23
sync.....	8-23
throttle.....	8-25
unblockfiles .....	8-27

**Appendix A RepliStor Components**

Kernel Driver .....	A-2
RepliStor Service .....	A-2
RepliStor Control Service .....	A-3
RepliStor Client .....	A-3
Performance Monitor Support .....	A-3
SNMP Support .....	A-4
Microsoft Cluster Server Support .....	A-4
Help .....	A-4
Documentation .....	A-5

**Appendix B RepliStor Exchange Support**

Overview .....	B-2
Installing RepliStor Exchange Support .....	B-3
Considerations When Using Failover .....	B-9

**Appendix C Security**

Security Levels .....	C-2
Connection .....	C-2
Integrity .....	C-2
Encrypted .....	C-3
Setting the Default Security Level for All Future RepliStor Sites .....	C-3
Setting the Security Level for a Single RepliStor Site .....	C-3
Encrypted File Support .....	C-4
Exporting Encrypted File Certificates from the Source to the Target .....	C-4
Setting Client Security .....	C-8
Windows NT LAN Manager and Kerberos .....	C-9
NTLM .....	C-9
Kerberos.....	C-9
Account Settings .....	C-10
Disallowing a System Account .....	C-11
Assigning RepliStor Administrative Permissions.....	C-11
Checklist for Increasing the Security Level .....	C-12

**Appendix D Files in the Data Directory**

Files in the Data Directory ..... D-2

**Appendix E Utilities**

Using the regutil Utility ..... E-2

Using the Set Active Computer Name Utility ..... E-3

**Glossary** ..... g-1

**Index**..... i-1



2-1	Site List in Windows without Active Directory .....	2-4
2-2	Site List in Windows with Active Directory .....	2-4
2-3	RepliStor Client Window .....	2-6
2-4	RepliStor Toolbar .....	2-7
2-5	Customizing the Toolbar .....	2-7
2-6	RepliStor Workspace .....	2-8
2-7	Displaying Information on an Item in the Directory Tree .....	2-9
2-8	RepliStor Traffic Light Icons .....	2-10
2-9	RepliStor Site Started .....	2-12
2-10	Global Exclude Window .....	2-17
2-11	Global Exclude Specification List .....	2-18
2-12	Select Source Window .....	2-19
2-13	Select Target Window .....	2-20
2-14	Target Path Drop-Down List .....	2-21
2-15	Delete Directory Drop-Down List .....	2-22
2-16	Options Dialog Box .....	2-22
2-17	VSS Parameters Dialog Box .....	2-24
2-18	Synchronization Dialog Box .....	2-27
2-19	Mirror Shares Dialog Box .....	2-29
2-20	Select Source Window .....	2-31
2-21	Viewing and Selecting Subkeys .....	2-32
2-22	Select Target Window .....	2-32
2-23	Options Dialog Box .....	2-33
2-24	Synchronization Dialog Box .....	2-34
2-25	Alias Maintenance Dialog Box .....	2-35
2-26	Add Computer Alias Dialog Box .....	2-35
2-27	Sync Options Dialog Box .....	2-39
2-28	Viewing Synchronization Status .....	2-40
2-29	Defining E-mail Addresses for Notification .....	2-41
3-1	VSS Framework Overview .....	3-3

3-2	Vssadmin command .....	3-5
3-3	VSS Parameters Tab in the Specifications Dialog Box .....	3-7
3-4	Taking a Shadow Copy .....	3-11
3-5	Shadow Copy Confirmation Dialog Box .....	3-11
3-6	Scheduling Shadow Copies .....	3-12
3-7	Shadow Copy in Progress .....	3-13
3-8	Shadow Copy Sets in the Tree Pane of the RepliStor Client Window .....	3-14
3-9	Target Shadow Copy Right-Click Menu .....	3-15
3-10	Cancelling a Shadow Copy in Progress .....	3-17
3-11	Vssadmin list writers command .....	3-19
4-1	Client Options Dialog Box .....	4-2
4-2	RepliStor Options Dialog Box .....	4-3
4-3	RepliStor Options Account Tab .....	4-4
4-4	Schedule Options Dialog Box .....	4-5
4-5	Throttling Options .....	4-7
5-1	Using Alias Failover .....	5-3
5-2	Alias Maintenance Dialog Box .....	5-5
5-3	Add Computer Alias Dialog Box — Alias Tab .....	5-6
5-4	Add Computer Alias — IP Addresses Tab .....	5-7
5-5	Add Computer Alias — Services Tab .....	5-8
5-6	Add Computer Alias — Target Tab .....	5-9
5-7	Alias Maintenance Dialog Box with New Alias and Target Listed .....	5-9
5-8	Options Dialog Box—Directories Tab .....	5-12
5-9	Script Repository Dialog Box .....	5-12
5-10	Activating Aliases on the Target System .....	5-16
5-11	Failover Tab on the Options Dialog Box .....	5-18
6-1	RepliStor Performance Monitor Window .....	6-2
6-2	Performance Tab in the Options Dialog Box .....	6-5
6-3	Processes Tab in the Options Dialog Box .....	6-7
6-4	Add Process Dialog Box .....	6-9
6-5	Performance Information in the Client Window .....	6-10
6-6	Status of Sites in the Sites Pane of the Client Window .....	6-11
6-7	Status of the Source and Target Servers in the Tree Pane of the Client Window .....	6-12
6-8	RepliStor Monitor Traffic Lights .....	6-13
6-9	Viewing Target Site Information .....	6-15
6-10	Viewing Specification Information .....	6-16
6-11	Viewing Synchronization Status .....	6-17
6-12	Pending Synchronizations .....	6-18
6-13	Checking File Update Status .....	6-19
6-14	Checking Source Status from the Target .....	6-19
6-15	Options — Log Tab .....	6-22
6-16	Specifications Dialog Box .....	6-24

6-17	Select Target Dialog Box .....	6-25
6-18	Network Attached Storage Admin Dialog Box .....	6-25
6-19	Network Attached Storage Properties Dialog Box .....	6-26
6-20	Modifying Specifications .....	6-30
6-21	Sync Options Dialog Box .....	6-32
6-22	Viewing Pending Synchronizations .....	6-35
6-23	Cluster Administrator .....	6-37
6-24	New Resource Dialog Box .....	6-37
6-25	Bringing the Resource Online .....	6-39
6-26	Site Forwarding Option .....	6-43
6-27	Marking Messages as Read .....	6-47
6-28	Simulation/Reporting Feature in Maintenance Menu .....	6-50
6-29	Bandwidth Simulation/Reporting Dialog Box .....	6-51
6-30	Source Report Example .....	6-52
6-31	Target Report Example .....	6-53
6-32	Reporting in Progress Message in Lower Right Status Bar .....	6-55
B-1	Introduction Window .....	B-3
B-2	License Agreement Window .....	B-4
B-3	Exchange Parameters — Alias Window .....	B-4
B-4	Exchange Parameters - VSS Window .....	B-5
B-5	Ready to Install the Program Window .....	B-7
B-6	Installing RepliStor Exchange Support .....	B-8
B-7	Setup Complete Window .....	B-8
C-1	Microsoft Management Console .....	C-5
C-2	Certificate Export Wizard .....	C-6
C-3	Client Options — Comm Tab .....	C-8
C-4	User/Groups Permissions .....	C-11



1-1	Situations in which to Use Aliases .....	1-6
2-1	Description of the Four Panes in the RepliStor Workspace .....	2-8
2-2	Keyboard Shortcuts .....	2-11
2-3	Specification Types .....	2-15
3-1	Description of Target Shadow Copy Right-Click Menu Options .....	3-16
5-1	Environment Variables Set When Running Executables .....	5-10
6-1	Performance Tab Field Descriptions .....	6-6
6-2	Traffic Light and Computer Icon Condition in the Site Pane .....	6-11
6-3	Traffic Light and Computer Icon Condition in the Tree Pane .....	6-13
6-4	RepliStor Monitor Buttons .....	6-14
6-5	Target Site Status Information in the List Pane .....	6-15
6-6	Target Site Specification Information in the List Pane .....	6-16
6-7	Specification Synchronization Status Information in the List Pane .....	6-18
6-8	Troubleshooting Using the Monitor Bar .....	6-20
6-9	Pending Synchronization Messages .....	6-35
6-10	Message Log Severity Levels .....	6-46
6-11	Source Report Column Headings .....	6-52
6-12	Target Report Column Headings .....	6-54
8-1	Syntax Conventions .....	8-3
8-2	aliasadd Required Parameters .....	8-4
8-3	aliasadd Optional Parameters .....	8-4
8-4	aliasop Optional Parameters .....	8-6
8-5	checkpoint Optional Parameters .....	8-7
8-6	modifyspec Required Parameters .....	8-9
8-7	modifyspec Optional Parameters .....	8-10
8-8	report Optional Parameters .....	8-12
8-9	spec Required Parameters .....	8-16
8-10	spec Optional Parameters .....	8-16
8-11	spec/globalexclude Parameters .....	8-19
8-12	spec/register Required Parameters .....	8-20

8-13	spec/registry Optional Parameters .....	8-20
8-14	spec/shares Parameters .....	8-21
8-15	specstatus Parameters .....	8-22
8-16	sync Optional Parameters .....	8-24
8-17	throttle Optional Parameters .....	8-26
D-1	Files in the Data Directory .....	D-2
E-1	regutil Parameters .....	E-2

---

*As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this EMC RepliStor for Microsoft Windows Version 6.1 Administrator's Guide may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.*

*This EMC RepliStor for Microsoft Windows Version 6.1 Administrator's Guide provides information on how to configure and manage the EMC RepliStor software.*

*Install the RepliStor software on the appropriate server and clients before using the information presented in this guide. Refer to the EMC RepliStor for Microsoft Windows Version 6.1 Installation Guide for installation instructions.*

*Post-release information is contained in the Release Notes for this product. This document is available online at the Legato website ([www.legato.com](http://www.legato.com)). Refer to the website periodically to view the latest Release Notes.*

*If a product does not function properly or does not function as described in this EMC RepliStor for Microsoft Windows Version 6.1 Administrator's Guide, please contact your EMC representative.*

**Audience**

This information in this guide is intended for use by system administrators responsible for installing software and maintaining the servers and clients on a network.

**Organization**

Here is an overview of where information is located in this guide.

Chapter 1, *Introduction*, introduces the EMC RepliStor software.

Chapter 2, *Getting Started*, explains how to get started using RepliStor software and how to create specifications for replicating data.

Chapter 3, *Configuring, Taking, and Managing Shadow Copies*, describes how to configure, take, and manage consistent point-in-time data in the form of VSS shadow copies.

Chapter 4, *Configuring Options*, provides instructions for configuring options for RepliStor software. These options determine how the Client window appears and how RepliStor software operates.

Chapter 5, *Configuring Failover*, describes how to configure RepliStor failover options.

Chapter 6, *Administering RepliStor Software*, describes the administrative tasks you can perform once RepliStor software is installed and configured.

Chapter 7, *Recovering Data*, describes how to use RepliStor software to recover replicated data when a system fails and a failover occurs.

Chapter 8, *Commands*, describes RepliStor commands that you can run from a command prompt.

Appendix A, *RepliStor Components*, describes the various utilities that can be used with the RepliStor software.

Appendix B, *RepliStor Exchange Support*, provides information on installing the RepliStor Exchange Support utility. This utility should be installed if you are planning to take VSS-compliant shadow copies of Exchange Server 2003 databases.

Appendix C, *Security*, discusses security issues related to RepliStor software, including available security facilities, ramifications, and trade-offs.

Appendix D, *Files in the Data Directory*, describes the various RepliStor components including the file type, default installation directory, filename, and functions.

Appendix E, *Utilities*, provides information about files that reside in the RepliStor data directory. The data directory is the main location where data on the target system is stored.

The *Glossary* contains terms related to the RepliStor software.



## Related Documentation

The following documents, available on the EMC Legato website at <http://web1.legato.com/cgi-bin/catalog?sf=Releases>, provide additional information about the RepliStor software:

- ◆ *EMC RepliStor for Microsoft Windows Version 6.1 Installation Guide*, P/N 300-002-463 Rev A01
- ◆ *EMC RepliStor for Microsoft Windows Version 6.1 Release Notes*, P/N 300-002-461 Rev A01

An extensive archive of product documentation is available at <http://web1.legato.com/cgi-bin/catalog?sf=Releases>. Most of the documents are in Adobe Acrobat Portable Document Format (PDF), and can be viewed by downloading and installing the Adobe Acrobat Reader. The Reader is available in the `/viewers/acroread` directory on the Legato Documentation Suite CD-ROM, or directly from Adobe at [www.adobe.com](http://www.adobe.com). To install and use the Reader on the preferred platform, refer to the instructions in the CD-ROM's `/viewers/acroread/readme.txt` file or at the Adobe website.

## Conventions Used in This Guide

This document uses the following typographic conventions and symbols to make information easier to access and understand.

---

A note presents information that is important, but not hazard-related.

---



---

AVANT GARDE

Keystrokes

---

**Palatino,  
bold**

- ◆ Dialog box, button, icon, and menu items in text
- ◆ Selections you can make from the user interface, including buttons, icons, options, and field names

---

*Palatino,  
italic*

- ◆ New terms or unique word usage in text
  - ◆ Command line arguments when used in text
  - ◆ Book titles
-

---

*Courier, italic* Arguments used in examples of command line syntax.

---

Courier System prompts and displays and specific filenames or complete paths. For example:  
 working root directory [/user/emc]:  
 c:\Program Files\EMC\Symapi\db

---

**Courier, bold** ♦ User entry. For example:  
**sympoll -p**  
 ♦ Options in command line syntax

---

### Where to Get Help

Legato offers a variety of methods; including electronic, telephone, and fax support to obtain company, product, and technical information. For questions about technical support, call your local sales office or service provider.

### General Information

The Legato website provides most of the information you might need. Technical bulletins and binary patches are also accessible on the Legato FTP site. For specific sales or training needs, e-mail or call Legato.

EMC Legato Service or Resource	Technical Bulletins	Binary Patches	Company and Product Information	Training Programs
www.legato.com	Yes	Yes	Yes	Yes
ftp.legato.com	Yes	Yes		
Legato Sales (650) 210-7000 Option 1 sales@legato.com			Yes	
Legato Education Services (650) 842-9357 training@legato.com				Yes

### Technical Support

The [www.legato.com/support](http://www.legato.com/support) website provides contact information, software patches, technical documentation, and information about available support programs.

You can contact technical support at:

**Worldwide Toll Free:** 877-534-2867<sup>a</sup>

**Worldwide:** +1-905-315-4788

- a. For international dialing, excluding Canada and Guam, enter your AT&T Direct Access Number first, wait for the prompt, then dial 877-534-2867.

You can also e-mail technical support at [support@legato.com](mailto:support@legato.com).

Customers with an active support agreement have access to Legato's integrated product knowledge base. Help with Legato software issues is also available through Legato Technical Support.

Customers without an active support agreement can contact Legato Technical Support services for per-update/per-incident support or Support Sales and Renewals to purchase annual software update subscriptions. You can contact Support Sales and Renewals by using one of the following contact methods:

Contact Method	Americas and Asia	Pacific Europe, Middle East, Africa
Phone Number	(650) 812-6000, option 2	+31 23 554-8870
Email	supportsales@legato.com	supportsalesEMEA@legato.com

## Licensing and Registration

To license and register this product, go to the Legato licensing website. To change contact information, transfer licenses, or for questions about licensing, use one of the following contact methods:

Licensing and Registration	Contact
EMC Legato licensing website	<a href="http://legato.com/support/licensing">http://legato.com/support/licensing</a>
Telephone number	(650) 812 6000 (option 3, option 2) <sup>a</sup> +31 23 554 8881 <sup>b</sup>
Fax number	(650) 745-1477 <sup>a</sup> +31 23 554 8808 <sup>b</sup>
E-mail	<a href="mailto:licensing@legato.com">licensing@legato.com</a> <a href="mailto:licensingemea@legato.com">licensingemea@legato.com</a>

- a. Contact information for the Americas, Asia, and the Pacific.  
b. Contact information for Europe, the Middle East, and Africa.

## Your Comments

Comments and suggestions about software features, the installation procedure, and documentation are welcome. Please send suggestions and comments to [feedback@legato.com](mailto:feedback@legato.com). Receipt of all e-mail correspondence is confirmed. Although every request cannot be

responded to personally, all comments and suggestions are considered during product design.

Help improve the documentation by completing a brief survey. Visit the Legato website at [www.legato.com](http://www.legato.com), navigate to the documentation page, and click on the survey link.

---

This chapter introduces the EMC RepliStor software. This chapter includes the following sections:

- ◆ RepliStor Software Overview ..... 1-2
- ◆ RepliStor Software Components ..... 1-3
- ◆ RepliStor Replicating Functions ..... 1-4
- ◆ Failover Capabilities ..... 1-4
- ◆ VSS-compliant Shadow Copies ..... 1-7

## RepliStor Software Overview

EMC® RepliStor® software provides real time data replication and increases the availability and reliability of Windows servers without the use of proprietary or specialized hardware. RepliStor software operates over both local area network (LAN) and wide area network (WAN) connections, and allows remote administration and installation. RepliStor software also operates across Windows domains; you can administer RepliStor software across domains from any RepliStor client in the network.

RepliStor software replicates files, directories, shares, and registry keys from a *source system* (the computer that has the data that needs to be protected), to a *target system* (the computer to which the data is replicated). After an initial synchronization of the data, RepliStor software replicates any changes made to the files (or keys), transfers those changes to the target system, and updates the files on the target system.

Before you begin replicating, you must create a *specification*, that tells RepliStor software which files, directories, Registry keys, and shares to replicate. For information about creating specifications, refer to *Creating Specifications* on page 2-15.

After you create specifications, you must *synchronize* the data from the source to the target system. Synchronization ensures that the replicated data on the target system exactly matches the original data on the source system before the replicating process begins. For more information on synchronization, refer to *Synchronizing Specifications* on page 2-39 and page 6-32.

## RepliStor Software Components

The RepliStor software provides the following components:

- ◆ **RepliStor client** — The product's graphical interface. Use the RepliStor client to configure RepliStor software, start the RepliStor server, and administer the RepliStor solution. You can also use the client to administer RepliStor software remotely.
- ◆ **RepliStor server** — Provides capabilities for data replication. When a Windows server runs the RepliStor server, it is known as a RepliStor *site*.
- ◆ **RepliStor driver** — Works with the RepliStor server to provide data replication.
- ◆ **RepliStor Control service** — Provides functions such as remote starting and stopping of a RepliStor server.
- ◆ **RepliStor installation program** — Installs and removes RepliStor software on supported Windows platforms.
- ◆ **SNMP Agent Extension dynamic link library (DLL)** — Extends the standard Windows Simple Network Management Protocol (SNMP) service to allow RepliStor software to send messages as SNMP events, enabling the RepliStor software to work with system management software.
- ◆ **Performance Monitor DLL** — Provides performance statistics on RepliStor software in the Windows Performance Monitor.

For more information about RepliStor components, refer to Appendix C, *Security*.

## RepliStor Replicating Functions

RepliStor software uses the following functions to replicate data in real time:

- ◆ *Mirroring* — Captures changes in data at the source system.
- ◆ *Forwarding* — Sends changes in data from the source system to the target system.
- ◆ *Updating* — Applies changes to the files on the target system.

Mirroring and forwarding must be enabled on the source system, and updating must be enabled on the target system for RepliStor software to operate normally. For more information about RepliStor software processes, refer to *Turning Processes On and Off* on page 6-41.

## Failover Capabilities

During normal operation, RepliStor software sends a *heartbeat* signal from the source system to the target system to let the target know the source is available and functioning normally. When a Windows server fails, RepliStor software performs a *failover* in which the processing and identity of that server are transferred to another server. Failover can occur transparently, without interrupting the processing of any client systems that may be attached to the failed server.

The failover option that RepliStor software provides is known as *Alias*. Alias failover takes advantage of Windows features that allow a target system to take over for one or more source systems while maintaining its own identity and processing.

You can configure how often the heartbeat signal is sent from the source system and how often the target system expects to receive it. If the target system does not receive the heartbeat signal within the expected period of time, then RepliStor software begins the failover process. First, RepliStor software attempts to PING the source system (with a standard IP/ICMP PING) to make sure the source system has actually failed.



RepliStor software allows you to configure the following when a failover occurs:

- ◆ Services started on the target system.
- ◆ The IP addresses and subnet masks forwarded from the source system to the target system.
- ◆ Commands or batch files to execute on the target system, both before and after the failover.
- ◆ Systems in the network that are notified of the failover.
- ◆ The time it takes RepliStor software to fail over. You can do this by changing the value in the **I'm Alive Recv (secs)** field in the **Time Limits** tab of the **Options** dialog box.

---

For information on configuring failover, refer to Chapter 5, *Configuring Failover*.

---

---

## Alias Failover

Alias failover is the only failover method supported in RepliStor Version 6.1. In RepliStor versions prior to Version 6.0, there was a choice for failover - *Automatic Switch Over* (ASO) and *Super Automatic Switch Over* (Super ASO). However, these failover methods were part of Windows NT, and Windows NT is not supported in this release.

The Alias failover option allows a target system to take over automatically for a source system when the source fails. Configuring the source and target options using Alias failover allows you to determine when and how the failover occurs.

With RepliStor software, you can define an *alias* and associated IP addresses for the source system for use in an Alias failover. On a failover, the alias information is sent to the target system to allow clients to connect using the alias when the actual source system is unavailable. You can also elect to add the alias to domain name system (DNS) servers.

With aliases, you can specify a particular target system for the source system to use on a failover. In addition, aliases are particularly useful in the situations described in Table 1-1 on page 1-6.

**Table 1-1 Situations in which to Use Aliases**

Situation	Description
Failover between two Windows 2000 (and later) domain controllers with Active Directory configured.	Aliases allow a failover to occur without machine name conflicts in Active Directory caused when the actual source machine name is added to the target system on a failure. With aliases, the source machine name is never sent to the target system, only the alias is sent.
Failover across a WAN, where IP addresses cannot be sent across different subnets.	For failover across a WAN, you must have a DNS server configured. Define one or more aliases that clients can continue to connect to on a failure and specify that the aliases be added to the DNS server. However, do not define associated IP addresses.
Failover recovery with minimal client connection disruption.	Clients connect to the alias on the source system before failover and continue to connect to the alias on the target system after failover. Specifications can be sent automatically to the target system and file changes continue to be tracked. When the source system is restored, remove the source machine name from the target system Added Names list, and all aliases and specifications are automatically removed from the target system. The aliases are added back to the source system. Re-enable and synchronize specifications on the source system, and client connections to the alias are interrupted only as long as it takes to transfer the alias.

For information on configuring failover options, refer to Chapter 5, *Configuring Failover*.

## VSS-compliant Shadow Copies

RepliStor allows you to create and manage Exchange Server 2003 and file system shadow copy backups via the Volume Shadow Copy Service (VSS) framework provided by Microsoft.

VSS is a framework that integrates user applications, backup applications, and storage systems to produce consistent point-in-time data in the form of VSS shadow copies. Specifically in RepliStor, VSS enables the integration of Exchange and RepliStor along with third-party hardware to provide shadow copy functionality.

Shadow copy backups of the Exchange Server 2003 database or file systems are created on the target system in case of source system failure. The shadow copy types supported on the target include CLARiiON® SnapView™ snapshots, Windows system shadow copies, and custom scripts.

---

The VSS-compliant shadow copy functionality is only supported on Windows 2003 systems.

---

Shadow copies allow you to obtain a consistent, restartable copy of the Exchange database or file system. This copy is available on the target side. If you want to perform a recovery, take the copy from the target and either use it locally or copy it back to the source server and start it.

---

Before creating Exchange shadow copies, you must install the Exchange Support utility, as described in Appendix B, *RepliStor Exchange Support*.

---

For more information on shadow copies, refer to Chapter 3, *Configuring, Taking, and Managing Shadow Copies*.



---

This chapter explains how to get started using RepliStor software and how to create specifications for replicating data.

This chapter includes the following sections:

- ◆ Before You Begin ..... 2-2
- ◆ Active Directory Support ..... 2-4
- ◆ Starting the RepliStor Client ..... 2-6
- ◆ Attaching to Remote RepliStor Sites ..... 2-14
- ◆ Creating Specifications ..... 2-15
- ◆ Creating Specifications for Circular Mirroring ..... 2-37
- ◆ Synchronizing Specifications ..... 2-39
- ◆ Detaching from a Site ..... 2-42

---

## Before You Begin

Before using RepliStor, read through this section to familiarize yourself with certain aspects of the RepliStor software to avoid problems during replication.

---

### Usage Considerations

- ◆ By default all target files are made read-only.
- ◆ The Delete Orphans functionality deletes all files on the target that do not exist on the source. The **Delete Orphans** option is found on the **Synchronization** dialog box (refer to Figure 2-18 on page 2-27). The **Delete Orphans** option has no effect on real-time mirroring. If a file is deleted on the source, it will be deleted on the target whether the **Delete Orphans** option is set or not.
- ◆ Circular mirroring should be used only in very carefully controlled situations. If you are considering using circular mirroring, contact EMC Customer Support.

---

For more information on circular mirroring, refer to *Creating Specifications for Circular Mirroring* on page 2-37.

---

### Common Mistakes to Avoid

- ◆ Never replicate to system files that are in use by the operating system, or that are specific to the target operating system.
- ◆ Never have an application open the replicated files on the target.
- ◆ Before replicating to any specific target directory, make sure that replicating to the target directory will not cause any problems when this directory is overwritten. For example, if you attempt to replicate from the C:\Documents and Settings directory to the same directory on the target, you will overwrite that directory on the target. This would be a bad user practice.

Certain folders such as Program Files, Documents and Settings, WINNT, Windows or other system-generated folders should be very carefully evaluated to make sure they:

- need to be replicated
- are not machine-specific

- will not cause an application or the operating system to become un-usable.

Since there is no way to identify each and every host-specific file, it is usually easier to replicate only those files which are verified to be required on the target and that will not cause any issues with the target's processes.

- ◆ In general, any file or folder that is flagged as a hidden, read-only, or system file should not be replicated.

## Active Directory Support

RepliStor software provides Windows 2000 (and later) Active Directory support. Some of the RepliStor dialog boxes and windows in this document show the product configured without an Active Directory, while others show it configured with an Active Directory. As Figure 2-1 and Figure 2-2 show, the main difference is that there is no **Get List** button in the **Site List** window with an Active Directory configured, because in that environment the list of all servers in the domain is immediately available to any client or server.




Figure 2-1 Site List in Windows without Active Directory

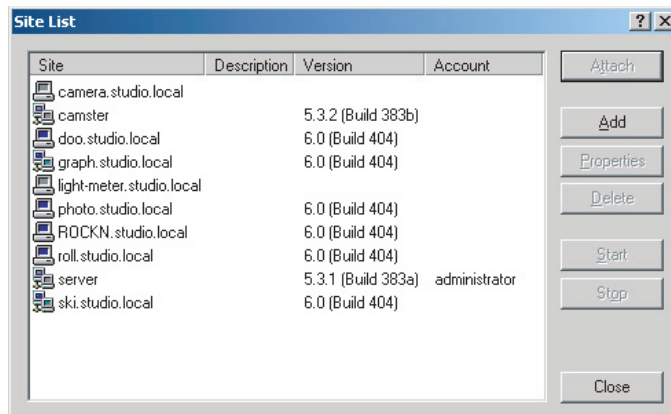


Figure 2-2 Site List in Windows with Active Directory



When installing RepliStor software, you can optionally register it in Active Directory. This has the following benefits:

- ◆ There is a common site list for all RepliStor installations in the domain. If any server makes a change to the site list, *all* servers see that change immediately.
- ◆ There is a common users list for all installations in the domain (on the **Users** tab in the **Options** dialog box).
- ◆ The site list is available immediately. Manually adding a site is only necessary if the site is outside the domain.

---

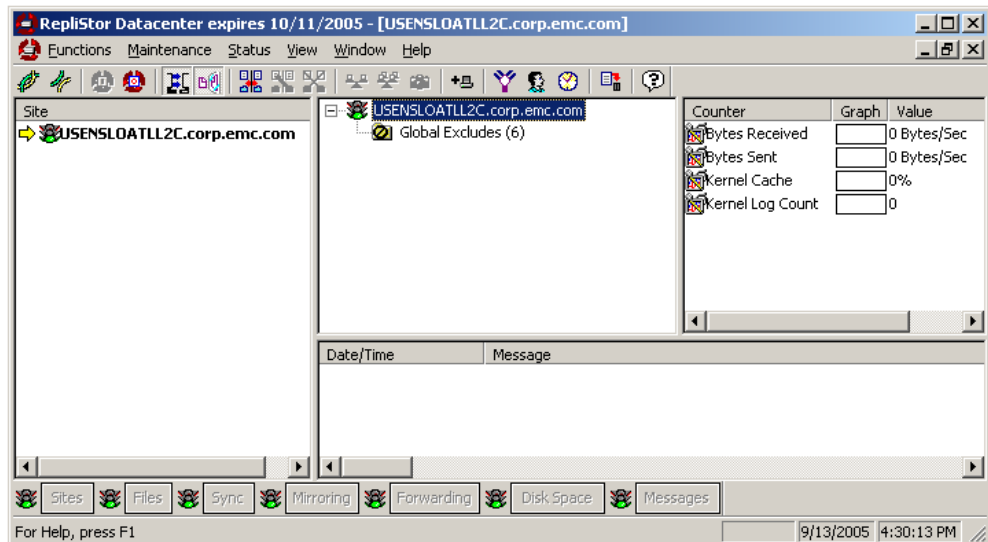
To register RepliStor software in Active Directory, the user performing the installation must be in the **Domain Admins** group. Once RepliStor software is installed, a domain administrator can assign administration rights to any other user or group.

---

## Starting the RepliStor Client

To start the RepliStor client, select **Start, Programs, RepliStor** from the Windows desktop.

When you initially use the product, the local *RepliStor site* is attached automatically, and the RepliStor client window opens, as shown in Figure 2-3. A RepliStor site is a Windows system, typically a server, running the RepliStor server.



**Figure 2-3** RepliStor Client Window

**Important:** To see an option's description in any of the dialog boxes, right-click the option's name, and then click **What's This?**

## The RepliStor Client Elements

The following sections describe key elements in the RepliStor client window and how to use them.

### The Toolbar

The toolbar, shown in Figure 2-4, contains icons that allow you to quickly execute many of the commands available on the RepliStor menus. To see a description of a toolbar button, move the pointer over the button.



Figure 2-4 RepliStor Toolbar

### How to Customize the Toolbar

To customize the RepliStor client window toolbar:

1. Select **Modify Tool Bar** from the **View** menu.

The **Customize Toolbar** dialog box appears, as shown in Figure 2-5.

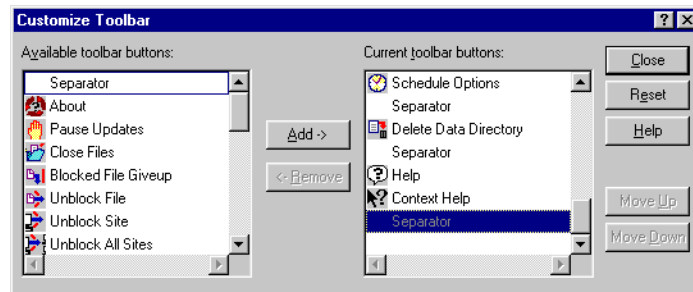


Figure 2-5 Customizing the Toolbar

2. To include a button in the toolbar, click the name of the button in the **Available Toolbar Buttons** list, and then click **Add**. You can also double-click the button to add it.
3. To remove a button from the toolbar, click the name of the button in the **Current Toolbar Buttons** list, and then click **Remove**.
4. To change the order of the buttons as they appear in the toolbar, click the name of the button in the **Current Toolbar Buttons** list that you want to move, and then click **Move Up** or **Move Down**.

5. To return the buttons to their original positions, click **Reset**.  
The buttons in the toolbar change immediately as you add, remove, and move them.
6. When customization is complete, click **Close**.

## The Workspace

The workspace is divided into four main sections: site pane, tree pane, list pane, and message pane. How these sections are arranged depends on the **Display** settings you choose in the **Client Options** dialog box (select **Client Options** from the **Maintenance** menu). For information about configuring client options, refer to *Configuring Client Options* on page 4-2. Figure 2-6 shows the default RepliStor workspace. Table 2-1 describes the four panes in the RepliStor workspace.

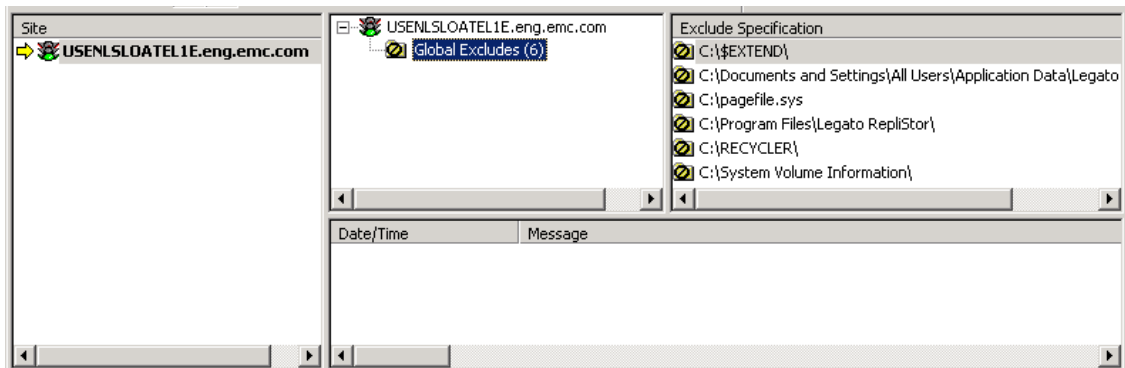


Figure 2-6 RepliStor Workspace

Table 2-1 Description of the Four Panes in the RepliStor Workspace

Pane	Description
<b>Site pane</b>	The left pane provides a list of sites you can attach.
<b>Tree pane</b>	The upper-middle pane provides a hierarchical directory information tree on the attached site. For example, you may see a list of <b>Global Exclude</b> specifications, file specifications, and aliases defined for the RepliStor site.
<b>List pane</b>	The upper-right pane provides a detailed list of information about the highlighted item; for example, aliases or <b>Global Exclude</b> specifications defined.
<b>Message pane</b>	The lower pane provides messages pertaining to the attached site.

RepliStor software displays one workspace for each attached site. If you attach to another site by clicking **Attach** from the **Functions** menu, another workspace opens in the RepliStor client window. If you double-click a site in the site pane, you detach from one server and attach to the other. You can verify if you are attached to multiple sites by clicking **Window**; the list of windows appears at the bottom of the menu.

### How to Display Information in the Directory Tree

To display information from the directory tree:

1. Click an item in the directory tree; for example, **Aliases**.
2. Right-click the item in the directory tree to display a submenu of available functions, as shown in Figure 2-7.




Figure 2-7 Displaying Information on an Item in the Directory Tree

### How to Change the Split in the Workspace

You can change the split position of the workspace by selecting **Split** from the **View** menu and then moving the sizing arrows. You can also change the split from vertical to horizontal, as follows:

1. Select **Client Options** from the **Maintenance** menu.
2. Clear the **Vertical Split** option, and then click **OK**.

The split change appears the next time you open a new window.

### How to Close the Message Pane

You can close the message pane and have messages accessible via an icon as follows:

1. Select **Client Options** from the **Maintenance** menu.
2. Clear the **Separate Message Pane** option, and then click **OK**.

The next time you open a new window, the message pane does not appear.

## Using Traffic Lights to Determine Status

Traffic light icons are used in the workspace to display the status of RepliStor sites, as shown in Figure 2-8.



Figure 2-8 RepliStor Traffic Light Icons

### Traffic Lights in the Site Pane

- ◆ A traffic light with no lights on (that is, a solid black traffic light not showing either a green, yellow, or red light) means the site is not running, but RepliStor software is installed and the control service is running at that site.
- ◆ A dimmed light (that is, a traffic light that is grayed out) means no RepliStor process is running on the site or the site does not exist.
- ◆ A red light means one or more of the following is true:
  - at least one unread *Severe* log message is at the site
  - the Disk Space Monitor is on and the amount of free space has dropped below the Stop level
  - at least one target site is blocked
  - at least one file is blocked
- ◆ A yellow light means one or more of the following is true:
  - at least one unread *Warning* log message is at the site
  - the Disk Space Monitor is on and the amount of free space has dropped below the Notify level
- ◆ A green light means all systems are functioning normally.

For more information about using traffic lights, refer to *Using Traffic Lights and Computer Icons in the Site Pane of the Client Window* on page 6-11.

### Traffic Lights in the Tree Pane

- ◆ A traffic light with no lights on (that is, a black traffic light not showing either a green, yellow, or red light) means the RepliStor server is not running at the attached site.
- ◆ A green light means all systems are functioning normally.

- ◆ A yellow light means systems are functioning and currently working, but may have encountered problems that require attention.
- ◆ A red light means systems have stopped and are not functioning.

For more information about using traffic lights, refer to *Using Traffic Lights and Computer Icons in the Tree Pane of the Client Window* on page 6-12.

### Traffic Lights in the Monitor Bar

The monitor bar lists various RepliStor functions and indicates their status using a traffic light. A green light means systems are functioning; a yellow light means systems are functioning and working, but require attention; a red light means systems have stopped and are not functioning.

### The Monitor Bar

To make the monitor larger:

1. Select **Client Options** from the **Maintenance** menu.
2. In the **Client Options** dialog box, select the **Large Monitor Bar** option, and then click **OK**.

For more information about the monitor bar and how to use it, refer to *Using the Monitor Bar* on page 6-13.

### Keyboard Shortcuts

The RepliStor client supports keyboard shortcuts to make using the interface easier.

Table 2-2 lists the keyboard shortcuts available with the RepliStor client.

**Table 2-2 Keyboard Shortcuts**

Keyboard Shortcut	Description
F5	Refreshes all panes in the workspace. <hr/> This would only be necessary to use in the event where systems are not functioning properly. <hr/>
F6	Displays the next pane, where the pane is the directory tree of the workspace.
SHIFT+ F6	Displays the previous pane.
CTRL + F6	Displays the next window in the workspace.
SHIFT + CTRL + F6	Displays the previous window in the workspace.

Table 2-2 Keyboard Shortcuts (continued)

Keyboard Shortcut	Description
SPACE	In the <b>Select Target</b> tab of the <b>Specification</b> dialog box, the SPACE key will check or uncheck a site entry.
TAB	In the <b>Select Target</b> tab of the <b>Specification</b> dialog box, the TAB key is used to make the drop-down list on the selected site entry appear. A path may be directly entered or the ALT down arrow key can be pressed to drop down the list.

## Starting a RepliStor Site

Once you open the RepliStor client, do one of the following to start a RepliStor site.

If the site is listed in the left pane:

1. Select the site you want to start from the list.
2. Select **Start Server** from the **Functions** menu.

The RepliStor site starts and its icon displays a green light, as shown in Figure 2-9.

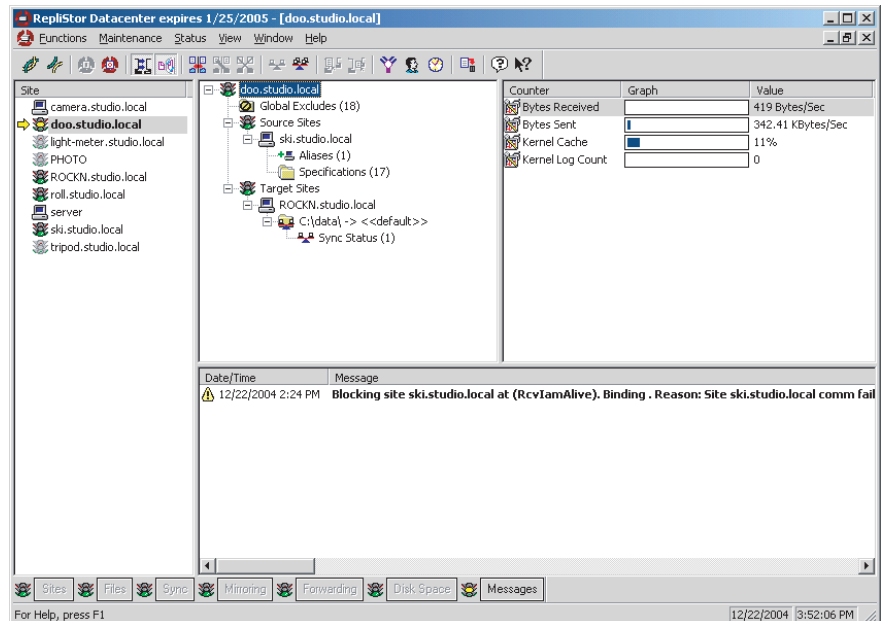


Figure 2-9 RepliStor Site Started



If the site is not listed in the left pane:

1. Select **Attach** from the **Functions** menu to open the **Site List** dialog box (refer to Figure 2-1 and Figure 2-2 on page 2-4).
2. Select the site you want to start, and then click **Start**. To attach to the site, click **Attach**.

The **Site List** dialog box closes and the RepliStor site icon displays a green light, as shown in Figure 2-9 on page 2-12.

RepliStor software provides some default **Global Exclude** specifications. These specifications indicate files to exclude when replicating from the source system to the target system. For information about creating specifications, refer to *Creating Specifications* on page 2-15.

## Attaching to Remote RepliStor Sites

Before configuring or administering a remote RepliStor site, you must attach to and start that site.

**Important:** To connect to a remote RepliStor server, you must have an account on the remote system in the **Windows Administrators** group.

To attach to a remote RepliStor site:

1. Start the RepliStor client.
2. If the remote site is listed in the left pane, double-click it to attach to it.
3. If the remote site is not listed in the left pane, select **Attach** from the **Functions** menu.

The **Site List** dialog box opens, as shown in Figure 2-1 and Figure 2-2 on page 2-4. If the desired site is not listed, click **Add**, and in the **Site Properties** dialog box specify the site or, on a system configured without an Active Directory, click **Get List** to search the network for all target sites. This may take a few minutes.

4. (Optional) Right-click the remote site in the left pane and select **Properties** to set the properties and the communication security level for the remote site in the **Site Properties** dialog box. For more information about setting security levels, refer to *Security Levels* on page C-2 in Appendix B, *RepliStor Exchange Support*.
5. Make sure the RepliStor server is running on the remote site by verifying that its icon displays a green light. If the RepliStor server is not running on the remote site, start it by selecting **Start Server** from the **Functions** menu or by right-clicking the remote site in the left pane and selecting **Start**.

## Creating Specifications

After attaching to the local and remote sites, you must create specifications that tell RepliStor software which data to replicate (or exclude from replicating). A specification identifies the path and file information that RepliStor software requires to replicate data from a source system to a target system.

Table 2-3 describes the kinds of specifications you can configure with the RepliStor software.

**Table 2-3 Specification Types**

Specification Type	How to Create Specification
<b>Global Exclude</b>	Identify files, directories, and shared directories that you do not want replicated from the source system. A <b>Global Exclude</b> specification takes precedence over all other specifications created after it. For instructions on how to configure <b>Global Exclude</b> specifications, refer to <i>Global Exclude Specifications</i> on page 2-16.
<b>File/Directory</b>	Identify the files, directories, and shared directories to replicate to a target system. For instructions on how to create <b>File/Directory</b> specifications, refer to <i>File/Directory Specifications</i> on page 2-18.
<b>Share</b>	Replicate all of the shares on the source system to the target system. For instructions on how to create <b>Share</b> specifications, refer to <i>Share Specifications</i> on page 2-28.
<b>Registry</b>	Replicate specific keys within the <code>HKEY_LOCAL_MACHINE</code> Windows Registry hive. For instructions on how to create <b>Registry</b> specifications, refer to <i>Registry Specifications</i> on page 2-30.
<b>Macintosh</b>	Replicate Macintosh volumes residing on a Windows Server. For instructions on how to create <b>Macintosh</b> specifications, refer to <i>Macintosh Specifications</i> on page 2-36.  A <b>Macintosh</b> specification is not a separate type of specification. A <b>File</b> specification is created for Macintosh data and when that is completed, it is called a <b>Macintosh</b> specification.

Keep the following in mind when creating specifications:

- ◆ Create **Global Exclude** specifications first since they take precedence over all other specifications.
- ◆ Do not replicate the following types of files:
  - Virtual drives (same as *redirected drives*)
  - RepliStor files
  - Windows system files
- ◆ RepliStor replicates compressed files on a source system as compressed files on the target system.

If the source or target system uses NTFS file systems, there are special requirements for the user ID in addition to those described in the *EMC RepliStor for Microsoft Windows Version 6.1 Installation Guide*:

- ◆ The user ID used to add a specification must have Full Control permission on the *source* system for the files, directories, and shares to be replicated.
- ◆ The user ID used to add or modify a specification must have Full Control access on the *target* system to the files, directories, and shares to be replicated.

The account under which RepliStor software is run, typically the System account, must have Full Control access to files on both the source and target systems.

---

## Global Exclude Specifications

**Global Exclude** specifications identify files, directories, and shared directories that you do not want replicated from the source system. The RepliStor software comes with the following **Global Exclude** specifications:

- ◆ One for each swap file (C:\pagefile.sys)
- ◆ RepliStor program directory (usually C:\<Program FilesDirectory>\Legato RepliStor)
- ◆ RepliStor data directory (C:\Documents and Settings\All Users\Application Data\Legato RepliStor Data)

---

For a list of files that reside in the RepliStor data directory, refer to Appendix D, *Files in the Data Directory*.

---

- ◆ All System Volume Information directories on all drives.
- ◆ All RECYCLER directories (or RECYCLED directories on FAT volumes)
- ◆ All \$EXTEND directories on NTFS file systems.

**Important:** Configure **Global Exclude** specifications before configuring file/directory or share specifications. If you configure file/directory or share specifications first, RepliStor software performs the initial synchronization, and any subsequent **Global Exclude** specifications do *not* take effect retroactively.

## Creating a Global Exclude Specification

To create a **Global Exclude** specification:

1. Start the RepliStor client and attach to the source system for which you are creating the specification.
2. From the **Maintenance** menu, select **Add Specification, Global Exclude**.

The **Global Exclude** window opens, as shown in Figure 2-10.

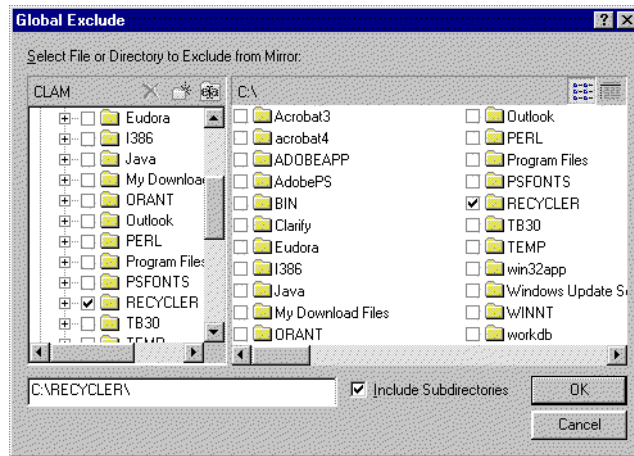


Figure 2-10 Global Exclude Window

3. Select a drive or a single file or directory to exclude from replication.

When the **Include Subdirectories** checkbox is selected, all subdirectories in the selected drive or directory are also excluded.

The left pane in the **Global Exclude** window lists the drives found on the source system. The right pane lists the contents of the selected drive.

The files to exclude appear in the text box in the lower left portion of the window. You can also enter a filepath in this text box. For example, to exclude all files of type `.tmp` in the `WINWORD` directory, enter the following:

```
C:\WINWORD\*.TMP
```

4. Click **OK**.

The specification is added to the **Exclude Specification** list in the RepliStor client window. To see the list, click the **Global Excludes** folder in the tree pane (the upper-middle pane) of the client window, as shown in Figure 2-11.

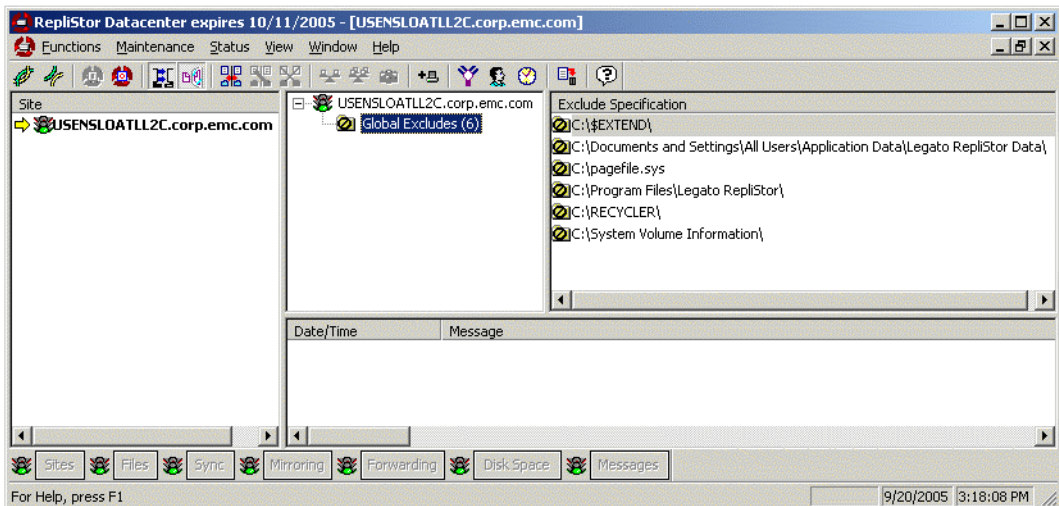


Figure 2-11 Global Exclude Specification List

## File/Directory Specifications

**File/Directory** specifications identify the files, directories, and shared directories to replicate to a target system.

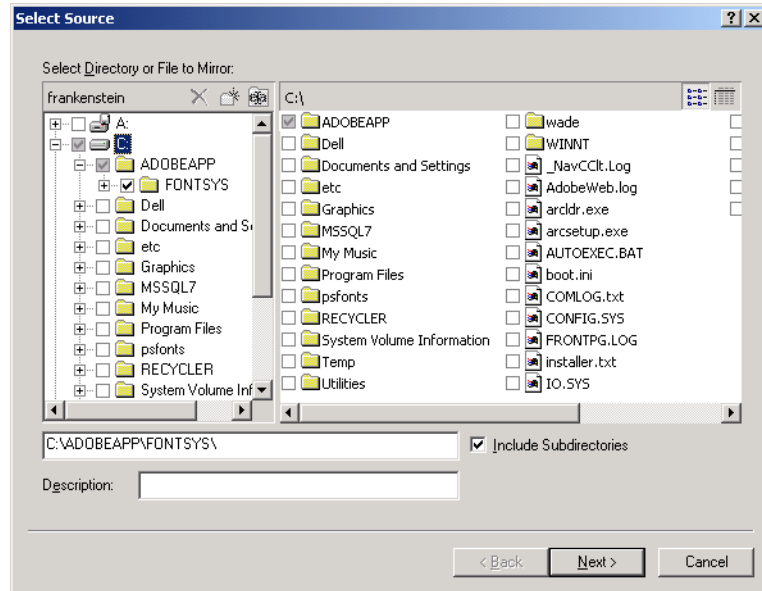
Once you have configured a **File/Directory** specification, you must synchronize the specification in order for RepliStor software to replicate the selected files. For detailed instructions on synchronizing, refer to *Synchronizing Specifications* on page 2-39. Also refer to page 6-32 for more information about synchronizing specifications.

### Creating a File/Directory Specification

To create a **File/Directory** specification:

1. Start the RepliStor client and attach to the source system for which you are creating the specification.
2. From the **Maintenance** menu, select **Add Specification, File/Directory**.

The **Select Source** window opens, as shown in Figure 2-12.



**Figure 2-12** Select Source Window

3. Select a drive, or select a single file or directory on a drive, to include in the specification.

When the **Include Subdirectories** checkbox is selected, all subdirectories in the selected drive or directory are also included in the specification.

The files that you select for replication appear in the text box in the lower left portion of the window. You can also enter a filepath in this text box. For example, to include all files of type `.doc` in the `WINWORD` directory, enter the following:

```
C : \WINWORD\* .DOC
```

4. Click **Next**.

The **Select Target** window opens, as shown in Figure 2-13 on page 2-20.




Figure 2-13 Select Target Window

5. Select all of the target sites to use for this specification. If the desired target site is not listed, click **Sites** to open the **Site Select** dialog box.

The **Select Target** tab allows you to select the target server from the list of known RepliStor servers.

Select **Show All Sites** to show the sites in your site list. This is normally selected if you need to add additional target sites to the specification. This is the default when creating a new specification.

Select **Show Selected Sites Only** to only show the sites that are targets to the specification. This is the default when editing an existing specification.

All known RepliStor sites are selectable. If a site is not selectable, either the RepliStor server is not running at that site, or it is not recognized as an active site in the network.



6. When you select a target site, a drop-down list appears in the **Target Path** column, as shown in Figure 2-14. Select one of the following target path options:
  - **Same As Source** — RepliStor automatically replicates the data selected in the **Select Source** tab to the target machine using exactly the same path as the source.
  - **Default Target Directory** — RepliStor automatically replicates the data selected in the **Select Source** tab to the default directory you selected on the target machine on the **Maintenance, Options, Directories** tab.
  - **Browse** — RepliStor replicates the data selected in the **Select Source** tab to the folder that you select in the **Select Target Path** dialog box. If **Browse** is selected, do one of the following: select the folder on the target system where you want to replicate the data or type the file path where you wish to replicate in the field below the main window.





Figure 2-14 Target Path Drop-Down List

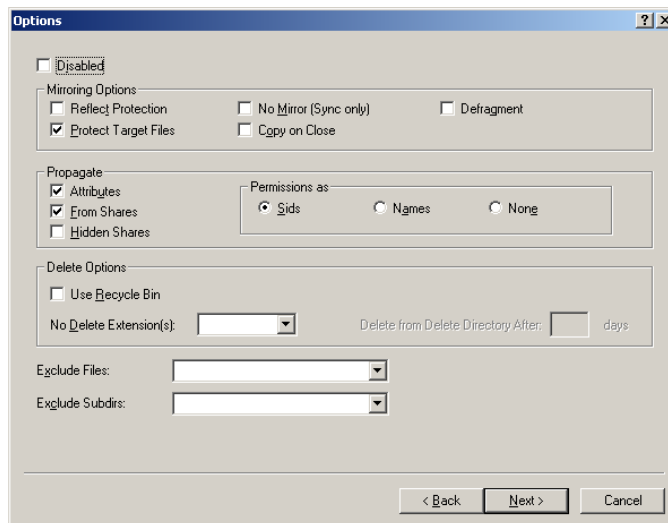
(Optional) The **Delete Directory** option is used on the target system to store files that have been deleted on the source. Click in the **Delete Directory** column and a drop-down list appears, as

shown in Figure 2-15. Select **Browse** to open the **Select Target Path** dialog box where you can select the target directory, or type the target path.



**Figure 2-15 Delete Directory Drop-Down List**

7. Click **Next**. The **Options** dialog box opens, as shown in Figure 2-16.



**Figure 2-16 Options Dialog Box**

8. Select the desired options. Refer to the online help for a description of each option.

---

If the **Disabled** option is selected in the **Options** dialog box, the synchronization options are disabled and the specification will not be synchronized.

---

9. To exclude files of a particular type from the specification, enter the file type you want to exclude in the **Exclude Files** box. For example:

**\*.tmp**

To specify more than one file type, enter a semicolon (;) between each type. For example:

**\*.tmp;\*.dot;\*.exe**

You can also exclude subdirectories in this way in the **Exclude Subdirs** box. When excluding subdirectories of a particular type from the specification, you may use the special characters \* and ?, where \* matches any number of any characters and ? matches any single character. Multiple subdirectories can be separated by a semicolon (;).

10. Click **Next**.

If your system is running on a Windows 2003 platform, the **VSS Parameters** dialog box appears (refer to step 11).

If your system is not running on a Windows 2003 platform, the **Synchronization** dialog box appears (refer to step 12 on page 2-27).

11. The **VSS Parameters** dialog box allows you to configure a VSS-compliant remote shadow copy for a particular specification on Windows 2003 systems (Figure 2-17 on page 2-24).

RepliStor Version 6.1 now allows you to create and manage Exchange Server 2003 and file system shadow copy backups via the Volume Shadow Copy Service (VSS) framework provided by Microsoft. Shadow copy backups (consistent point-in-time data copies) of the Exchange Server 2003 database or file system are created on the target system in case of source system failure. The shadow copy types supported on the target include CLARiON SnapView shadow copies, Windows system shadow copies, and custom scripts.

Shadow copies allow you to obtain a consistent, restartable copy of the Exchange database or file system. This copy is available on the target side. If you want to perform a recovery, take the copy from the target and either use it locally or copy it back to the source server and start it.

Before creating Exchange shadow copies, you must install the Exchange Support utility, as described in Appendix B, *RepliStor Exchange Support*.

For more information on how VSS works, refer to *What is VSS?* on page 3-3.

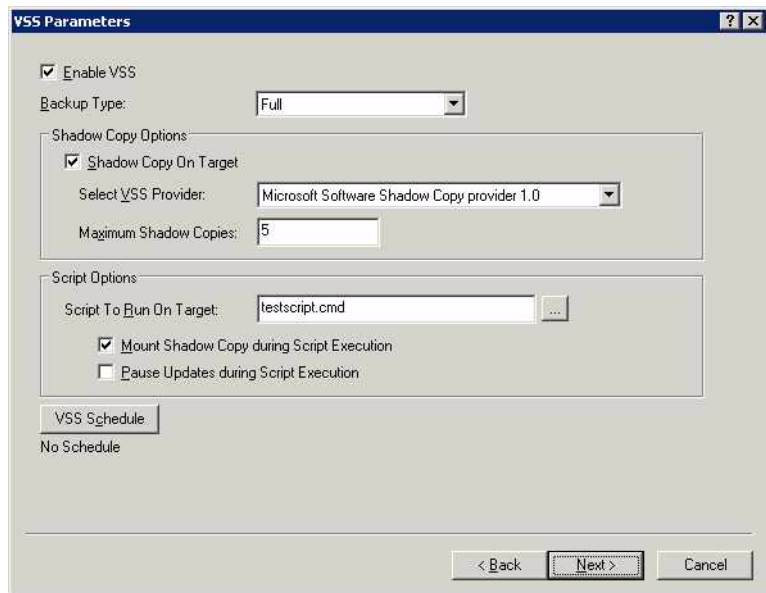


Figure 2-17 VSS Parameters Dialog Box

When the **Enable VSS** checkbox in the **VSS Parameters** dialog box is enabled, RepliStor takes a shadow copy of the drives associated with this specification. Enabling this option means you can request to take a shadow copy at any time on a particular specification.

The other parameters in the **VSS Parameters** dialog box are only relevant if the **Enable VSS** checkbox is enabled.

- ◆ **Backup Type** — This option specifies the type of shadow copy (backup type) that is performed on the source. This option is application/writer-dependent; in other words, it depends on

what you want your specific application to do. For example, if you are taking a shadow copy of an Exchange database, and select **Full**, the log files will be truncated. If you are taking an Exchange shadow copy and selected **Copy** for the backup type, the log files would not be truncated. The backup types are the following:

- **Full** — All files (regardless of whether they have been marked as backed up), are saved. Each file's backup history is updated to reflect that it was backed up.
- **Copy** — Files are copied to a backup medium regardless of the state of each file's backup history, and the backup history is not updated.
- ◆ **Shadow copy on Target** — Taking a target shadow copy is optional. The user may not require a shadow copy at the target and may instead run a script that for example, starts a backup or only copies files. When this option is checked, a *persistent* shadow copy is created on the target; the shadow copy will exist until you delete it.
- ◆ **Select VSS Provider** — The target system is contacted and a list of all VSS providers on that system is presented to the user. This option allows you to select the VSS provider to use when taking a shadow copy on the target such as CLARiiON SnapView shadow copies or Windows system shadow copies.

---

RepliStor does not support reverting shadow copies made using the CLARiiON provider.

---

- ◆ **Maximum Shadow copies** — This option determines the maximum number of shadow copies you want to keep on the target prior to creating the shadow copy on the target. Old shadow copies are automatically deleted if the total number of shadow copy sets exceeds the user-defined maximum.

For example, if you enter 5 maximum shadow copies, this means that RepliStor will delete all shadow copies on the target except for the latest five before taking the target shadow copy. If you enter zero in this field, this means that you want all shadow copies deleted on the target before executing the current shadow copy.

- ◆ **Script To Run On Target** — This option allows you to select a script to run on the target after the shadow copy is on the target. For example, you can run `eseutil.exe` on the target to make sure the target shadow copy is valid or a backup can be performed.
- ◆ **Mount Shadow Copy** — When this option is checked, the target shadow copy is mounted to a volume to become visible to the local host during execution of the script. It is unmounted when the script terminates. The following environment variables are set while executing the script: `SHADOW_SET_PATH` and `SHADOW_SET_ID`. The `SHADOW_SET_PATH` variable is set to the path where the shadow copy set has been mounted. `SHADOW_SET_ID` is the GUID of the shadow copy set.
- ◆ **Pause Update during Script Execution** — Updates are always paused during the target shadow copy operations to ensure that the shadow copy taken on the target is identical to the state of the source. Check this option if you need updates to remain paused during execution of the script. This option may be used if no shadow copy was taken on the target and the script operation is relatively short.
- ◆ **VSS Schedule** — Click this button if you want to schedule shadow copies for specification(s). Refer to *What is VSS?* on page 3-3 for more information on scheduling shadow copies.

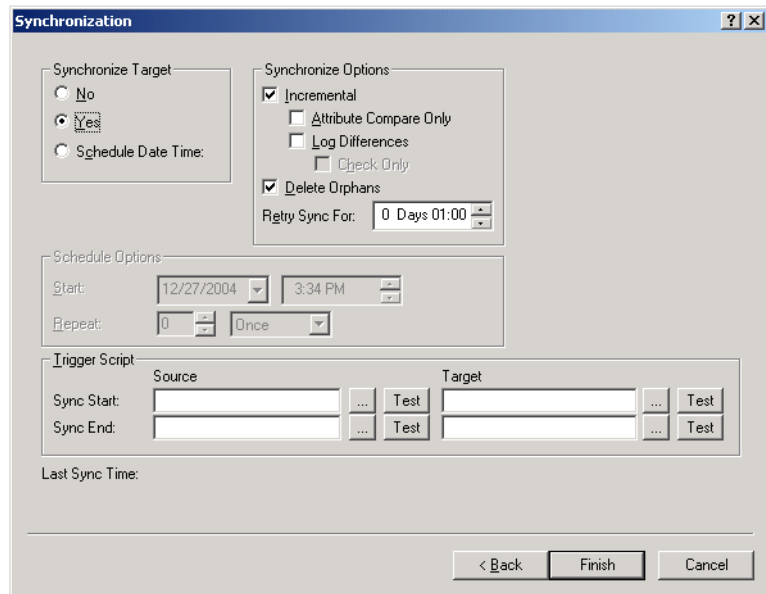
After selecting the shadow copy configuration options in the **VSS Parameters** dialog box, select **Next**.

---

For information on taking shadow copies, refer to *Taking Shadow Copies* on page 3-10.

---

12. The **Synchronization** dialog box opens, as shown in Figure 2-18



**Figure 2-18 Synchronization Dialog Box**

13. Set the desired synchronization options. Refer to the online help for a description of each option.
14. Click **Yes** in the **Synchronize Target** section of the dialog box to make sure the new specification is copied from the source system to the target system.
15. Click **Finish**.

RepliStor software synchronizes the specification and adds it to the RepliStor client window under the appropriate target site.

---

## Share Specifications

A share refers to a shared directory or folder. Shares can be replicated through a **File/Directory** specification, but if you want to replicate all shares from a drive, follow the steps in this section.

If you are creating a **Share** specification, then you must maintain identical file structures on both the source and target systems. (For example, you cannot create a specification that replicates `C:\SHARE\` on the source system to `D:\BKSHARE\` on the target system.)

If you do not have identical file structures on both systems, then you must create a **File/Directory** specification to replicate that file structure from the source system to the target system and select the **From Shares** option in the **Options** dialog box to replicate and translate the shares. If you are replicating multiple shares located in several locations on the source system, replicate the entire drive from the source system to the target system.

Keep the following in mind when creating a share specification:

- ◆ You cannot create a share specification if you are running RepliStor software in a Microsoft Cluster Server (MSCS) environment.
- ◆ Once you have configured a share specification, you must synchronize the specification to replicate the shares to the target. The files within the share are not replicated unless a file specification is also created. For detailed instructions on synchronizing, refer to *Synchronizing Specifications* on page 2-39. Also refer to page 6-32 for more information about synchronizing specifications.
- ◆ You cannot replicate network printer shares.

### Creating a Share Specification

To create a share specification:

1. Start the RepliStor client and attach to the source system for which you are creating the specification.
2. From the **Maintenance** menu, select **Add Specification, Share**. The **Mirror Shares** dialog box opens, as shown in Figure 2-19 on page 2-29.






Figure 2-19 Mirror Shares Dialog Box

3. In the **Target** text box, enter or select the name of the target site. If the desired target site is not listed, click **Sites** to open the **Site List** window.

4. To exclude shares from replicating, enter the share name to exclude. For example:

**Temp**

To specify more than one share name, use a semicolon (;) between each name. Use an asterisk (\*) as a wildcard character to specify groups of shares. For example:

**Share1; Share2; User\***

5. To synchronize the specification when it is created, click **Synchronize**.

6. To replicate any hidden shares on the source system, click **Hidden Shares**.

7. To replicate the permissions associated with the shares, click **Propagates Permissions**.

8. Click **OK**.

RepliStor software adds the specification to the RepliStor client window under the appropriate target site.

---

## Registry Specifications

When an application is installed on a system, configuration information is generally placed in the `HKEY_LOCAL_MACHINE` Registry hive. You can create specifications to replicate specific keys within the `HKEY_LOCAL_MACHINE` Windows Registry hive. This configuration information may be static, where the information does not change after installation, or it may be dynamic, where the application updates its Registry key during normal use. Thus, to fail over to the target an application with dynamically updated configuration information, and have that application run as it did on the source, you may need to replicate the updates to its Registry key.

---

Some applications also update the `HKEY_USERS` Registry hive with per-user information. RepliStor software cannot replicate these hives.

---

**Important:** You must be familiar with the exact information a specific application places in the Registry in order to know which keys to replicate. In addition, do not replicate machine-specific information from a source to target key, since this can cause serious problems on the target.

When creating a Registry specification, you can specify that permissions associated with an application's key be replicated from the source system to the target system. If you do replicate permissions, the target system must be within the source system's domain, or the target keys will be inaccessible. Note that the permissions are based on Security Identifiers (SIDs), not on text account names, which is an option with file specifications. A SID has no meaning outside of its domain.

If you do not replicate permissions to the target, the target uses the parent key's default permissions.

As a rule, only replicate small increments of the Registry—for example, only the key for a critical application. In addition, to avoid writing over configuration information in the target Registry, you may want to stage the replication and specify a different key path on the target system. For more information, refer to *Staging Registry Replication* on page 2-34.

## Creating a Registry Specification

To create a Registry specification:

1. Start the RepliStor client and attach to the source system for which you are creating the specification.
2. From the **Maintenance** menu, select **Add Specification, Registry**.

The **Select Source** window opens, as shown in Figure 2-20.

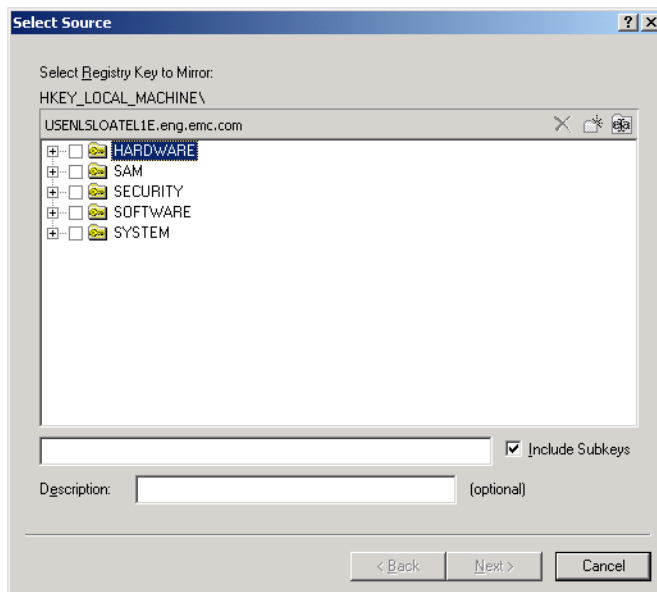
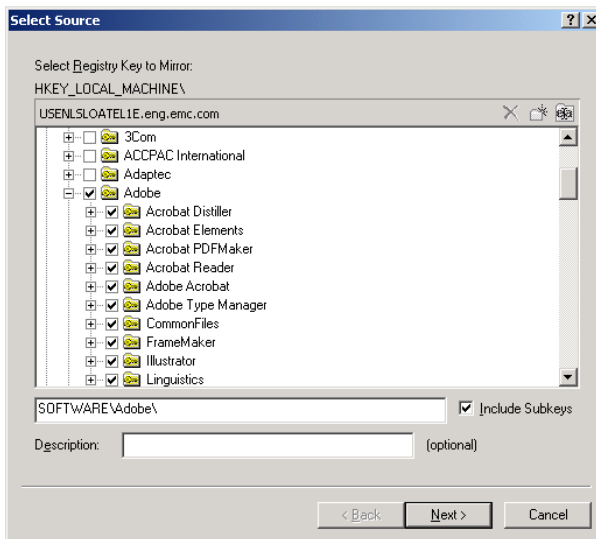


Figure 2-20 Select Source Window

Only the keys in the `HKEY_LOCAL_MACHINE` hive are listed.

3. Click the plus sign to the left of the keys to view a list of subkeys, as shown in Figure 2-21 on page 2-32.

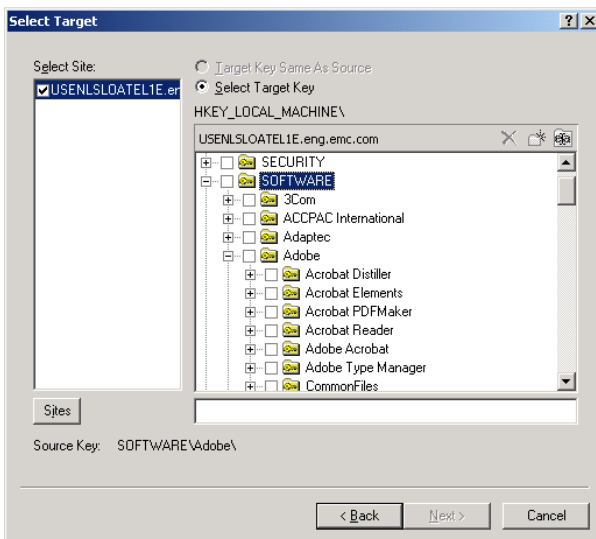
**Important:** Do not select an entire key, such as `Software`. Select only a specific application's key. For example, `Software\Company Name\App Name`.



**Figure 2-21 Viewing and Selecting Subkeys**

4. Select Next.

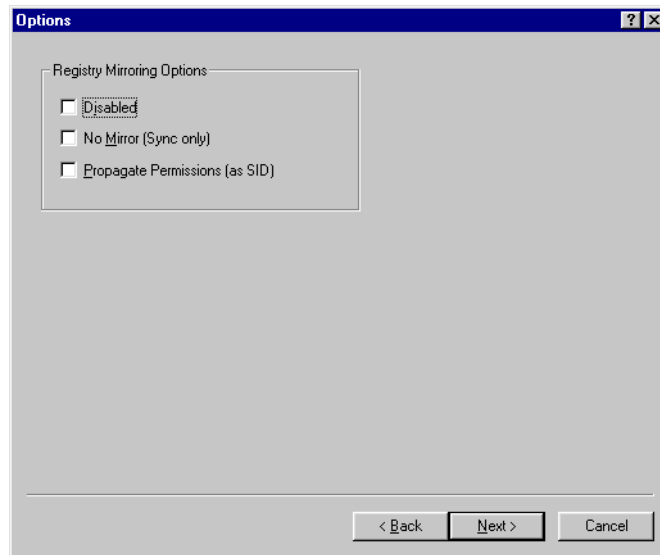
The **Select Target** window opens, as shown in Figure 2-22.



**Figure 2-22 Select Target Window**

5. Select the site and target key where you want RepliStor software to replicate the source key information. To stage the Registry replication, specify a different Registry key on the target system than on the source system. For more information about staging, refer to *Staging Registry Replication* on page 2-34.
6. Select **Next**.

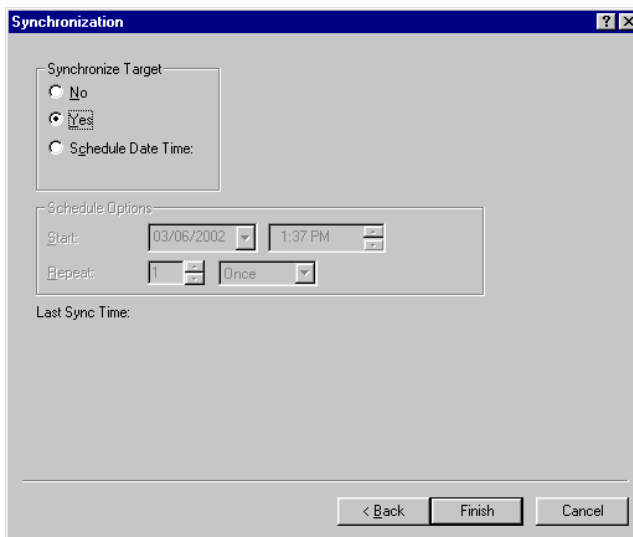
The **Options** dialog box opens, as shown in Figure 2-23.



**Figure 2-23 Options Dialog Box**

7. Set the Registry mirroring options as needed. Refer to the online help for a description of each option.
8. Click **Next**.

The **Synchronization** dialog box opens, as shown in Figure 2-24 on page 2-34.



**Figure 2-24 Synchronization Dialog Box**

9. Select the synchronization options as needed. Refer to the online help for a description of each option.
10. Click **Finish**.

RepliStor software adds the specification to the RepliStor client window under the appropriate target site.

## Staging Registry Replication

Certain applications may have Registry keys you do not want to replicate to the target system in real time. Instead, you may want to update the target Registry only at the time of failover.

To update the target Registry only during a failover:

1. Create a Registry specification, specifying a target Registry key path different than the source system. All Registry updates will go to this *staging area* in the target Registry. Make sure the staging area is not a subkey of the final destination key.
2. Create a script that uses the `regutil` utility to copy the keys and values from the staging area to the destination key area in the target system's Registry. For more information about the `regutil` utility and its syntax, refer to *Using the regutil Utility* on page E-2 in Appendix E, *Utilities*. Be sure to specify the `/y` option to delete all orphaned Registry keys.

- From the **Maintenance** menu, select **Alias**. The **Alias Maintenance** dialog box opens, as shown in Figure 2-25.

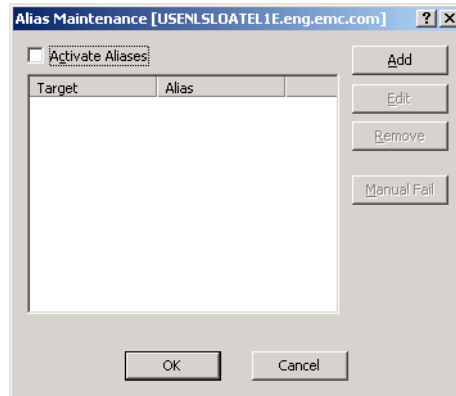


Figure 2-25 Alias Maintenance Dialog Box

- Click **Add** in the **Alias Maintenance** dialog box. The **Add Computer Alias** dialog box opens (Figure 2-26).

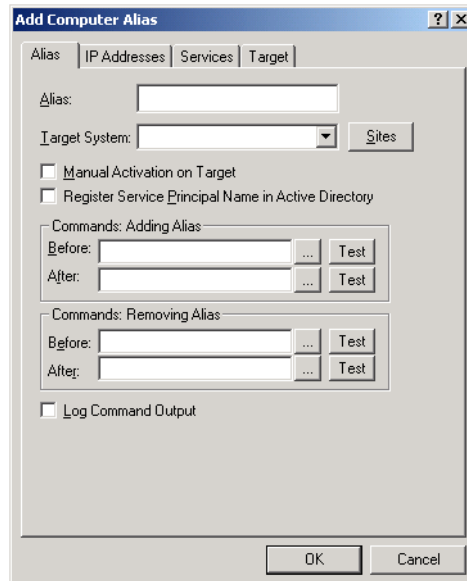


Figure 2-26 Add Computer Alias Dialog Box

5. In the appropriate **Before:** or **After:** text boxes beneath **Commands: Adding Alias**, enter the name of the script created in step 2. When a failover occurs, this script copies the updates from the staging Registry key to the destination key.

---

## Macintosh Specifications

Macintosh clients can access files on a Windows server just as they can in a Macintosh environment. With RepliStor software, you can create a Macintosh specification to replicate the Macintosh volumes residing on a Windows Server.

### Creating a Macintosh Specification

To create a Macintosh specification:

1. Install Macintosh services on both the source and target systems.
2. Make sure the names of the Macintosh volumes on the source system match the volume names on the target system.
3. Create a specification that replicates the root of the Macintosh volume (or above) and includes subdirectories.
4. Stop the File Server For Macintosh service on the target system.

For instructions on configuring failover for Macintosh volumes, refer to *Configuring Failover for Macintosh Volumes* on page 5-20.



## Creating Specifications for Circular Mirroring

Circular mirroring refers to configuring two servers to replicate the same data to each other:

```
Server A: c:\data -> Server B: c:\data
```

And a specification on Server B:

```
Server B: c:\data -> Server A: c:\data
```

In order for this to work, the **Reflect Protection** option must be set on both specifications and the **Protect Target Files** option must *not* be set on both specifications.

On the surface, this configuration seems to be the perfect way to keep two directories identical even if changes are made at either server. The reality is somewhat different. This configuration is supported, but in a limited way. The following restrictions apply:

- ◆ A file cannot be modified on both servers at roughly the same time or even opened at the same time.
- ◆ On Windows 2000, the files being replicated should not be executable files (that is, .exe or .dll files).
- ◆ No files are memory mapped.
- ◆ No files are encrypted.

In addition, there is a performance penalty when selecting the **Reflect Protect** option. The router functions in the service can be multithreaded, which means connections from multiple source systems can be processed simultaneously (or from a single source system if a synchronization is occurring with multiple target connections set). When setting the **Reflect Protect** option, the router functions revert to being single threaded.

If a file is modified at approximately the same time on both servers, both files will likely become corrupt. RepliStor cannot guard against this situation (since **Protect Target Files** needs to be off). RepliStor would not notice that corruption occurred, so it cannot log any message.

If an executable file is copied into a circular mirrored specification, it may not mirror, or it starts mirroring data endlessly, filling up the kernel cache and then the disk (with OC\$nnnnn files). The kernel in Windows 2000 has been *enhanced* to optimize handling of executable files. The executable file will generally be memory mapped and

RepliStor will not be able to know how that file has been updated, so the **Reflect Protect** function is not guaranteed to always work.

### Example of Where Circular Mirroring Would Not Be Feasible

A user has two SQL servers and wants to set up a circular mirror so that:

- ◆ both databases stay in sync all of the time;
- ◆ users can update either database and the changes will be sent to the other server.

This configuration is not possible for the following reasons:

- ◆ If the file is in use on the target, RepliStor will not be able to apply the updates.
- ◆ In circular mirroring, the same file on both the source and target systems cannot be modified at roughly the same time. This example configuration would violate this principle because the SQL database files on both servers are always open.
- ◆ SQL is not designed for its database file to be modified while it has it open. SQL will open the files in *exclusive* mode to prevent this. RepliStor honors this. If RepliStor did not honor this, and RepliStor were to slip these file changes in, SQL would not notice these updates and the data cached by SQL would not properly reflect the contents of the file. The result would probably be a corrupt file.
- ◆ RepliStor is an *asynchronous* replication product. This means if a change occurs on the source, it is made on the target at some point. The application does not wait for the change to be made on both servers before moving on to the next file operation. Even if the above problems were solved, this alone would prevent this configuration from working since SQL needs to lock records before making updates to prevent two processes from making the same update. Even if a lock translates to a file modification that will be replicated, it won't be able to negotiate with the target since it is a one-way, asynchronous mirror.

## Synchronizing Specifications

Before RepliStor software can replicate data successfully, an exact copy of the data from the source system must reside on the target system. To do this, you must synchronize the specifications.

**Important:** If you accept the defaults for synchronization when creating specifications, RepliStor software automatically does the initial synchronization.

RepliStor allows you to synchronize open files (that is, files in use). This is especially useful when you are replicating databases that hold files open continuously, such as SQL Server or Exchange databases. You must synchronize open files (for example, Exchange database files) when you create a specification so that RepliStor can later mirror changes to that file.

## Synchronizing Specifications

To synchronize specifications:

1. In the RepliStor client window, select the specifications to synchronize.
2. From the **Maintenance** menu, select **Synchronize** or **Synchronize All**. The **Sync Options** dialog box opens, as shown in Figure 2-27.




Figure 2-27 Sync Options Dialog Box

3. Select the appropriate synchronization options. Refer to the online help for a description of each option.

For a description of *full*, *incremental*, and *partial* synchronization, refer to *Synchronizing Specifications* on page 6-32.

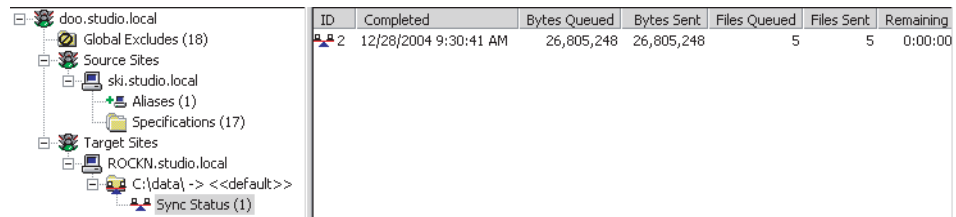
To perform a full synchronization (that is, copy the entire file from the source system to the target system), clear the **Incremental** checkbox.

To perform a partial synchronization (that is, copy a single file or subdirectory from the source system to the target system), click the **Sub-Path** button.

Click **Sync Now** to synchronize now, or click **Save** if you selected a scheduled date and time.

If you clicked **Sync Now**, click **Yes** in response to the confirmation message. If you clicked **Save**, click **OK** in response to the confirmation message.

4. To view the status of the synchronization, click **Sync Status** in the left pane of the RepliStor client window, as shown in Figure 2-28.



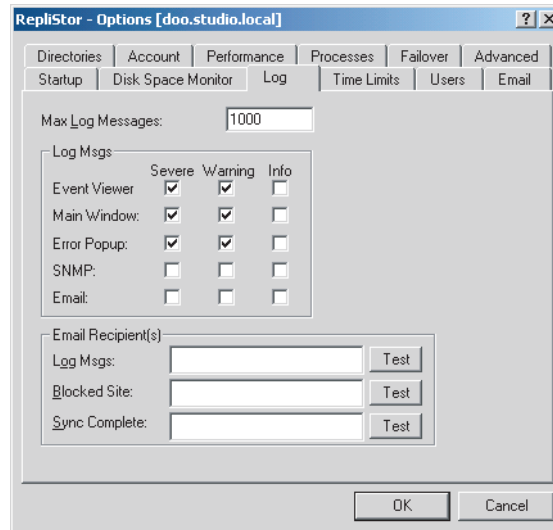
ID	Completed	Bytes Queued	Bytes Sent	Files Queued	Files Sent	Remaining
2	12/28/2004 9:30:41 AM	26,805,248	26,805,248	5	5	0:00:00

Figure 2-28 Viewing Synchronization Status

## Specifying an E-mail Address for Synchronization Notifications

To specify an e-mail address for synchronization notifications:

1. From the **Maintenance** menu, select **Options**.
2. In the **Options** dialog box, click the **Log** tab, as shown in Figure 2-29.



**Figure 2-29** Defining E-mail Addresses for Notification

3. In the **Sync Complete** text box, enter an e-mail address, and then click **OK**.

## Detaching from a Site

To detach from a RepliStor site:

1. Make sure the site you want to detach is the active site in the RepliStor client window. To do this, select the site from the **Window** menu.
2. From the **Functions** menu, select **Detach**.
3. Click **Yes** in response to the confirmation message.

## Configuring, Taking, and Managing Shadow Copies

RepliStor provides shadow copy backups for Microsoft Exchange 2003 databases and file systems. This chapter describes how to configure, take, and manage shadow copies.

This chapter includes the following sections:

- ◆ Overview ..... 3-2
- ◆ What is VSS? ..... 3-3
- ◆ Allocating Shadow Storage ..... 3-5
- ◆ Configuring Shadow Copies ..... 3-7
- ◆ Taking Shadow Copies ..... 3-10
- ◆ Managing Shadow Copies ..... 3-15
- ◆ Troubleshooting Shadow Copy Creation ..... 3-18

---

## Overview

RepliStor allows you to create and manage Exchange Server 2003 and file system shadow copy backups via the Volume Shadow Copy Service (VSS) framework provided by Microsoft.

VSS is a framework that integrates user applications, backup applications, and storage systems to produce consistent point-in-time data in the form of VSS shadow copies. Specifically in RepliStor, VSS enables the integration of Exchange and RepliStor along with third-party hardware to provide shadow copy functionality.

Shadow copy backups of the Exchange Server 2003 database or file systems are created on the target system in case of source system failure. The shadow copy types supported on the target include CLARiiON SnapView snapshots, Windows system shadow copies, and custom scripts.

---

The VSS-compliant shadow copy functionality is only supported on Windows 2003 systems.

Shadow copies allow you to obtain a consistent, restartable copy of the Exchange database or file system. This copy is available on the target side. If you want to perform a recovery, take the copy from the target and either use it locally or copy it back to the source server and start it.

---

Before creating Exchange shadow copies, you must install the Exchange Support utility, as described in Appendix B, *RepliStor Exchange Support*.



## What is VSS?

VSS is a framework that integrates user applications, backup applications, and storage systems to produce consistent point-in-time data in the form of VSS shadow copies. Specifically in RepliStor Version 6.1, VSS enables the integration of Exchange and RepliStor along with third-party hardware to provide shadow copy functionality.

There are three main VSS components: *Requestors*, *Writers*, and *Providers*. VSS acts as the coordinator between the activities of all the VSS Providers, Writers, and Requestors in the creation and use of shadow copies. The components of VSS are described below, and their relationship in the VSS architecture is shown in Figure 3-1.

- ◆ **The VSS Requestor** — is a replication application; it requests a shadow copy. RepliStor software is a VSS Requestor. The RepliStor Requestor has interfaces to the RepliStor console to associate it with a specification.
- ◆ **The VSS Writer** — is the application-specific logic needed in the shadow copy creation and restore/recovery process. The VSS Writer is provided by Exchange Server 2003 and other applications.
- ◆ **The VSS Provider** — is third-party hardware control software that actually creates the shadow copy.

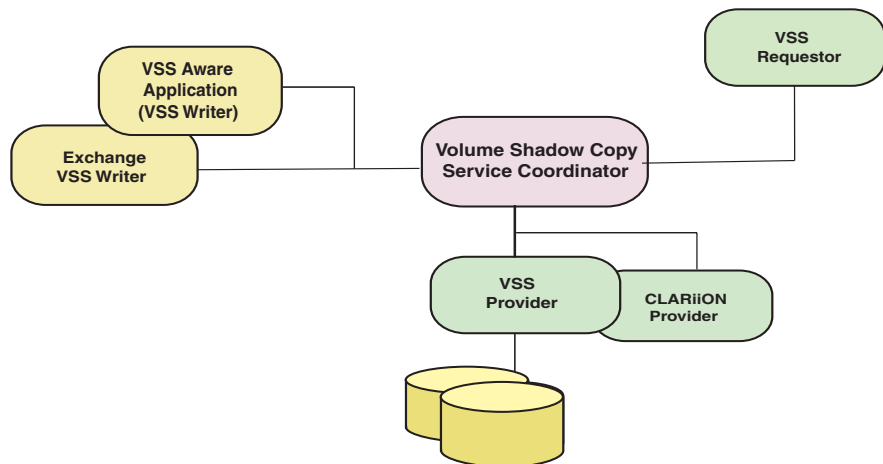


Figure 3-1 VSS Framework Overview

In general, the RepliStor/VSS sequence of events that occurs in the creation of a shadow copy on the target is the following:

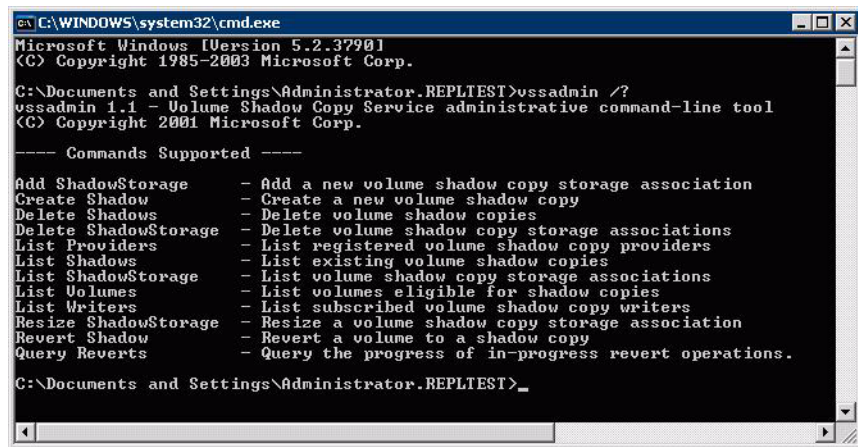
1. On the source side, the RepliStor Requestor (initiated via the RepliStor GUI) requests a VSS shadow copy.
2. VSS then asks the Exchange VSS Writer to freeze and flush all its I/O.
3. When this is completed, VSS asks the System Provider to take a shadow copy.
4. The RepliStor kernel interfaces with the VSS System Provider to determine when the I/O is that of a VSS shadow copy. This is the consistent data that is denoted by a *marker* which is passed to the target RepliStor server.
5. When the kernel *marker* is received at the target side, this indicates that the replicated data on the target is at the same state as the source at the time the data was frozen. This data is the consistent data for Exchange.
6. The target RepliStor server then optionally uses the RepliStor VSS Requestor to initiate a shadow copy on the target.
7. Once the shadow copy is on the target, the shadow copy can be mounted (if configured) and a user-supplied script can be run (if configured).

## Allocating Shadow Storage

Prior to configuring any shadow copies, you will need to allocate shadow storage.

The `vssadmin` CLI is the interface to the Microsoft System Provider. This command line will be used to allocate storage for shadow copies and view the status of the Exchange writer when troubleshooting shadow copy creation problems.

In a DOS window type `vssadmin /?` for a list of commands. Refer to Figure 3-2.



```
C:\WINDOWS\system32\cmd.exe
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

C:\Documents and Settings\Administrator.REPLTEST>vssadmin /?
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line tool
(C) Copyright 2001 Microsoft Corp.

---- Commands Supported ----
Add ShadowStorage      - Add a new volume shadow copy storage association
Create Shadow          - Create a new volume shadow copy
Delete Shadows         - Delete volume shadow copies
Delete ShadowStorage  - Delete volume shadow copy storage associations
List Providers        - List registered volume shadow copy providers
List Shadows           - List existing volume shadow copies
List ShadowStorage    - List volume shadow copy storage associations
List Volumes          - List volumes eligible for shadow copies
List Writers           - List subscribed volume shadow copy writers
Resize ShadowStorage  - Resize a volume shadow copy storage association
Revert Shadow         - Revert a volume to a shadow copy
Query Reverts         - Query the progress of in-progress revert operations.

C:\Documents and Settings\Administrator.REPLTEST>_
```

Figure 3-2 Vssadmin command

It is important to be sure that sufficient shadow storage is being allocated on the target server as this is where the permanent shadow copies will be stored. The amount of storage necessary will vary depending on the size of the Exchange Storage Group(s), amount of data change during the retention period and the number of shadow copies which are being maintained before recycling.

Depending on the available resources, allocating multiple drives to increase performance and eliminate a single point of failure would be optimal.

Once the storage has been presented, formatted and mounted on the target server, use the `vssadmin add shadowstorage` command to allocate shadow copy space. In the example below `F:` is the target drive on the RepliStor target server and `S:` is the shadow copy storage location:

```
vssadmin add shadowstorage /for=F: /on=S:
```

Additionally, allocating some storage for the temporary shadow copies that are created on the source can be helpful. This shadow storage can be minimal, as the shadow copies are deleted automatically by a VSS cleanup routine.

## Configuring Shadow Copies

If your system is running on a Windows 2003 platform, you can configure shadow copies for a specification via the **VSS Parameters** dialog box. This dialog box can be accessed in one of two ways:

- ◆ when you add a specification (select **Add Specification** from the **Maintenance** menu), the **VSS Parameters** dialog box displays after the **Options** dialog box.

For more information on configuring shadow copies via the **VSS Parameters** dialog box, refer to step 11 on page 2-23).

- ◆ when you right-click a specification in the tree pane, select **Modify** and then click the **VSS Parameters** tab in the **Specifications** dialog box (Figure 3-3).

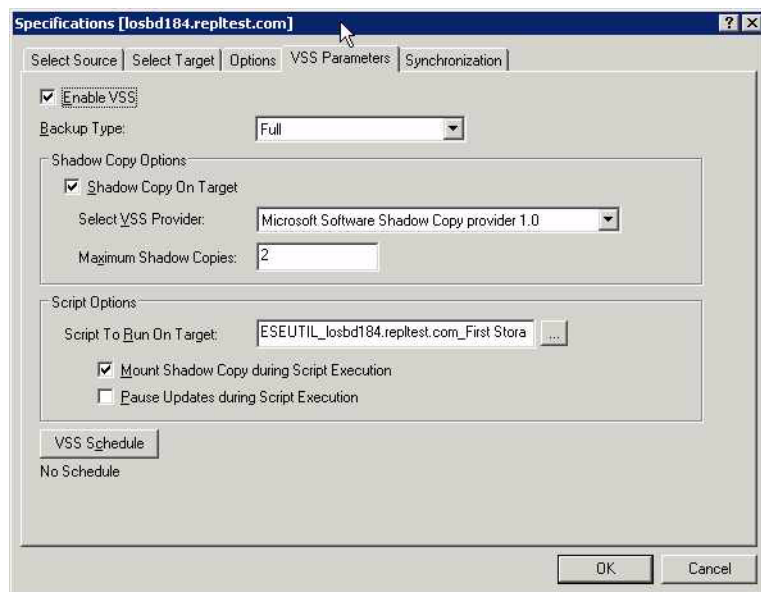


Figure 3-3 VSS Parameters Tab in the Specifications Dialog Box

When the **Enable VSS** checkbox in the **VSS Parameters** tab is enabled, RepliStor takes a shadow copy of the drives associated with this specification. Enabling this option means you can request to take a shadow copy at any time on a particular specification.

The other parameters in the **VSS Parameters** dialog box are only relevant if the **Enable VSS** checkbox is enabled.

- ◆ **Backup Type** — This option specifies the type of shadow copy (backup type) that is performed on the source. This option is application/writer-dependent; in other words, it depends on what you want your specific application to do. For example, if you are taking a shadow copy of an Exchange database, and select **Full**, the log files will be truncated. If you are taking an Exchange shadow copy and selected **Copy** for the backup type, the log files would not be truncated. The backup types are the following:
  - **Full** — All files (regardless of whether they have been marked as backed up), are saved. Each file's backup history is updated to reflect that it was backed up.
  - **Copy** — Files are copied to a backup medium regardless of the state of each file's backup history, and the backup history is not updated.
- ◆ **Shadow Copy on Target** — Taking a target shadow copy is optional. The user may not require a shadow copy at the target and may instead run a script that for example, starts a backup or only copies files. When this option is checked, a *persistent* shadow copy is created on the target; the shadow copy will exist until you delete it.
- ◆ **Select VSS Provider** — The target system is contacted and a list of all VSS providers on that system is presented to the user. This option allows you to select the VSS provider to use when taking a shadow copy on the target such as CLARiiON SnapView shadow copies or Windows system shadow copies.

---

RepliStor does not support reverting shadow copies made using the CLARiiON provider.

---

- ◆ **Maximum Shadow Copies** — This option determines the maximum number of shadow copies you want to keep on the target. Prior to creating the shadow copy on the target, old shadow copies are automatically deleted if the total number of shadow copy sets exceeds the user-defined maximum.

For example, if you enter 5 maximum shadow copies, this means that RepliStor will delete all shadow copies on the target except for the last five. Before taking the target shadow copy, if this field is empty, it will never delete a shadow copy before taking another one.


- ◆ **Script To Run On Target** — This option allows you to select a script to run on the target after the shadow copy is created. For example, you can run `eseutil.exe` on the target to make sure the target shadow copy is valid or a backup can be performed.
- ◆ **Mount Shadow Copy** — When this option is checked, the target shadow copy is mounted to a volume to become visible to the local host during execution of the script. It is unmounted when the script terminates. The following environment variables are set while executing the script: `SHADOW_SET_PATH` and `SHADOW_SET_ID`. The `SHADOW_SET_PATH` variable is set to the path where the shadow copy set has been mounted. `SHADOW_SET_ID` is the GUID of the shadow copy set.
- ◆ **Pause Update during Script Execution** — Updates are always paused during the target shadow copy operations to ensure that the shadow copy taken on the target is identical to the state of the source. Check this option if you need updates to remain paused during execution of the script. This option may also be used if no shadow copy was taken on the target and the script operation is relatively short.
- ◆ **VSS Schedule** — Click this button if you want to schedule shadow copies for specification(s). Refer to *Scheduling Shadow Copies* on page 3-12 for more information on scheduling shadow copies.

## Taking Shadow Copies

A shadow copy is associated with one or more specifications. Typically, a group specification is created that specifies all the related data. For example, a group specification may specify an Exchange Storage Group.

**Important:** If you are going to create a shadow copy of an Exchange Server 2003 database, before doing so, you must install the Exchange Support utility, as described in Appendix B, *RepliStor Exchange Support*.

You can take a shadow copy of one specification or one group specification, or multiple specifications or group specifications by doing either of the following:

- ◆ right-click the specification(s) in the tree pane of the RepliStor client and select **Shadow Copy** (as shown in Figure 3-4) or select the shadow copy button  from the toolbar.
- ◆ manually issue the `ShadowCopy` CLI command.

---

For more information on taking shadow copies from the CLI, refer to *ShadowCopy* on page 8-13.

---



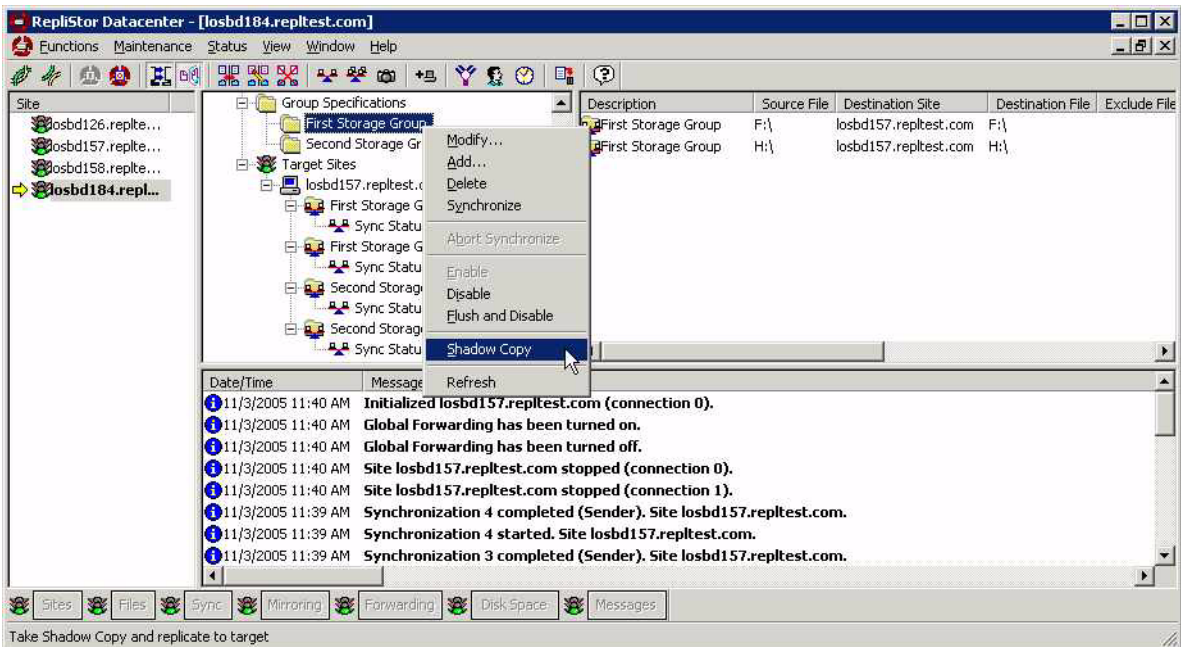


Figure 3-4 Taking a Shadow Copy

After you select **Shadow Copy**, the confirmation message shown in Figure 3-5 appears.




Figure 3-5 Shadow Copy Confirmation Dialog Box

In the **Confirm Shadow Copy** dialog box, you can select the **Backup Type**. (refer to *Configuring Shadow Copies* on page 3-7 for more information on this option) and you can select the **Schedule** option to schedule shadow copies for specification(s).

## Scheduling Shadow Copies

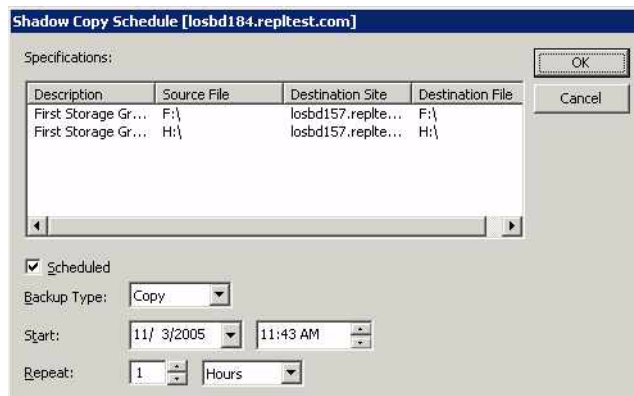
When you click **Schedule** in the confirmation dialog box shown in Figure 3-5, the **Shadow Copy Schedule** dialog box displays, as shown in Figure 3-6. You can set a repeating schedule to take shadow copies by specifying a starting time and repeat intervals (for example, minutes, hours, days) in the **Shadow Copy Schedule** dialog box.

When setting up shadow copy schedules, you cannot enter an end time or distinguish between weekdays and weekends. In addition, using the GUI, you cannot set two separate group specifications to shadow at the same time. If you want to set two separate group specifications to shadow simultaneously, you must use the CLI.



### CAUTION

**Microsoft Exchange routine maintenance performs a series of operations including defragmentation. Maintenance occurs by default between 1-5 am (configurable). It is highly recommended not to perform shadow copies during this period to avoid any additional overhead that may be caused by resource contention during this period.**



**Figure 3-6 Scheduling Shadow Copies**

After setting the shadow copy schedule (if desired), click **OK** in the **Shadow Copy Schedule** dialog box, and then click **OK** in the confirmation dialog box shown in Figure 3-5.

While the shadow copy is in progress, a **Shadow Copy** icon will appear in the upper part of the tree pane, and the Shadow Copy in Progress status will display in the right list pane, as shown in Figure 3-7.

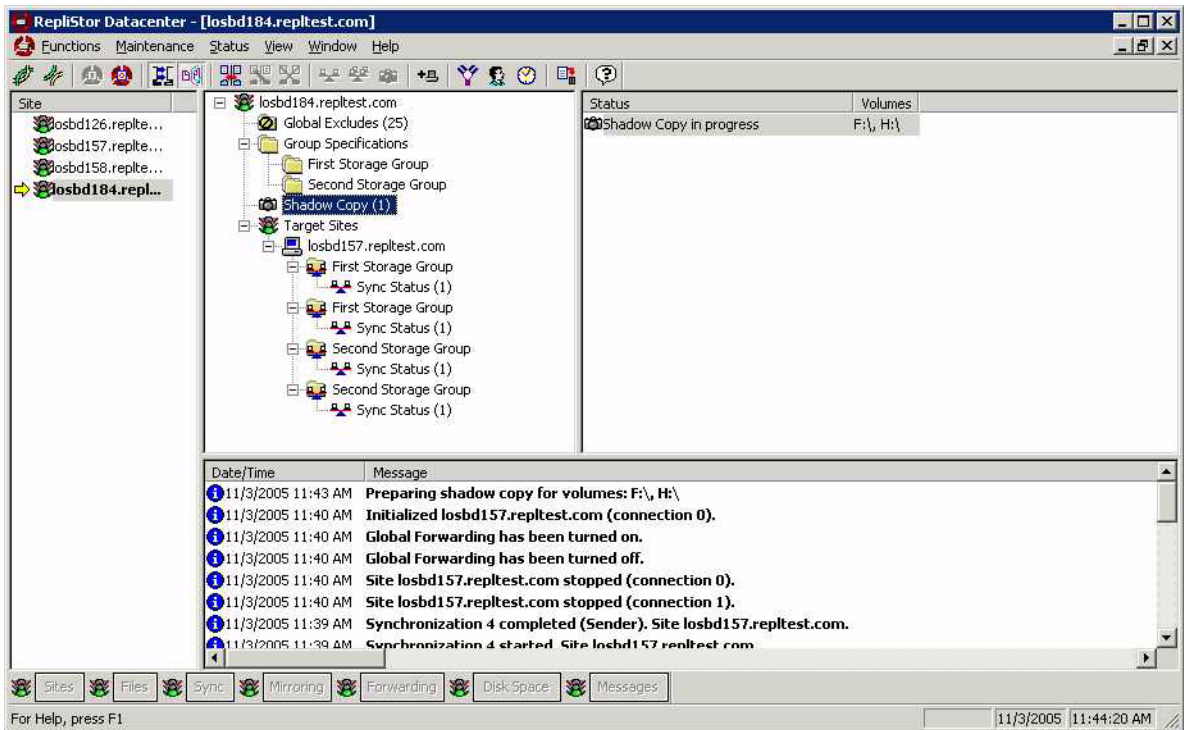


Figure 3-7 Shadow Copy in Progress

It is important to note that during the shadow copy process, a temporary shadow copy is created on the source system. Once a shadow copy is initiated, the source shadow copy shows up in the tree pane as the **Shadow Copy** icon, which is highlighted in gray in the above Figure 3-7. A source shadow copy has the following properties:

- ◆ It is not complete until the target shadow copy is complete and successful, and the target script (if any) has run and returns a SUCCESS error code.
- ◆ It is immediately marked *complete* if the target is not configured to take a shadow copy and there is no script to run.

- ◆ It is always marked as auto-delete and always uses the system VSS provider.


When the shadow copy has completed on the target, the target shadow copy set is designated by the **Shadow Copies Sets** icon  in the tree pane, and the temporary source **Shadow Copy** icon disappears from the upper part of the tree pane, as shown in Figure 3-7. The target shadow copy sets are listed under the source site that they are associated with, but they reside on the target site.



Figure 3-8 Shadow Copy Sets in the Tree Pane of the RepliStor Client Window

## Microsoft Cluster Server (MSCS) Support

RepliStor supports data replication from an MSCS cluster. If a node failure happens before the shadow copy is sent to the target, the shadow copy will not be created on the target. All future scheduled shadow copies will continue on the new node, and any shadow copy schedules will also resume on the new node.

## Managing Shadow Copies

Shadow copies are managed through the tree pane of the RepliStor client window, as shown in Figure 3-7 on page 3-13.

To manage target shadow copy sets, you can either right-click the target shadow copy set in the tree pane or in the list pane. By right-clicking on a target shadow copy set, you can mount, unmount, delete, revert, or list the properties for a particular shadow copy set. Figure 3-9 shows the target shadow copy right-click menu, and Table 3-1 on page 3-16 describes the right-click menu options.

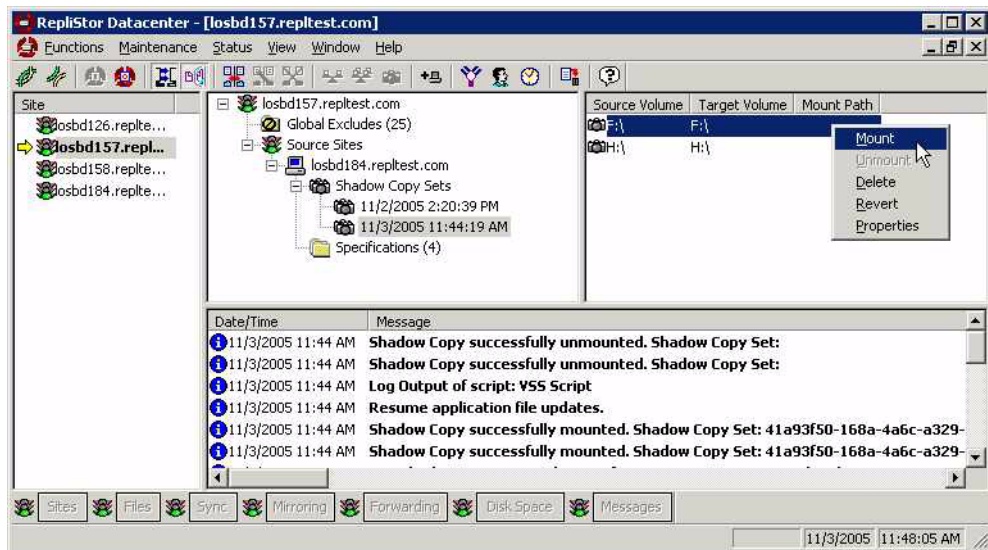


Figure 3-9 Target Shadow Copy Right-Click Menu

Table 3-1 Description of Target Shadow Copy Right-Click Menu Options

Right-Click Menu Option	Description
Mount	Mounts the selected shadow copy to an empty directory specified by the user.
Unmount	Unmounts the selected shadow copy from the directory.
Delete	Deletes the selected shadow copy.
Revert	Reverts the volume to the state of the selected shadow copy. The supported VSS providers for this option are the Microsoft system provider (if Windows 2003 SP1 is used).  VSS-compliant Service Pack 1 is required for the revert functionality. Microsoft recommends a hotfix, KB891957, that alleviates problems arising from depleted paged pool memory. Refer to KB 891957 on the Microsoft website, <a href="http://support.microsoft.com">http://support.microsoft.com</a>
Properties	Displays all the information known about this shadow copy. This should include: <ul style="list-style-type: none"> <li>• Creation time</li> <li>• Shadow Copy set GUID</li> <li>• Shadow Copy GUID</li> <li>• Source server of the shadow copy</li> <li>• Source volume</li> <li>• Target volume</li> <li>• Mount path, if any</li> <li>• CLARiiON shadow copy information (if CLARiiON provider)</li> <li>• XML recovery document path</li> </ul>

### Canceling a Shadow Copy in Progress

You may want to cancel a shadow copy in progress, if for example, your target system becomes unavailable. You can cancel a shadow copy by right-clicking the shadow copy in progress in the right list pane and selecting **Cancel**, as shown in Figure 3-10. When you cancel a shadow copy, a Warning log message is created indicating that the shadow copy was cancelled by the user.

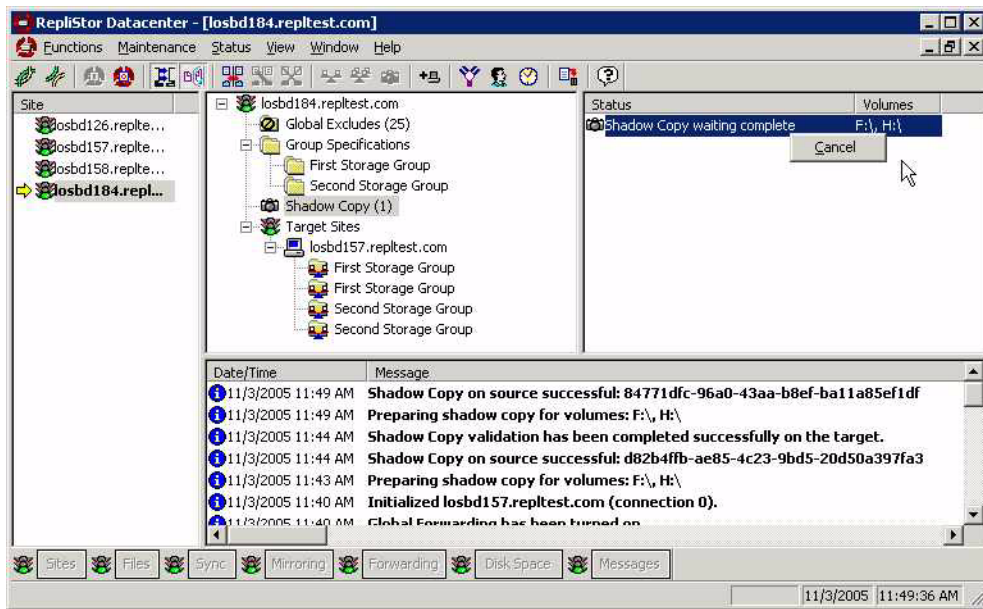


Figure 3-10 Cancelling a Shadow Copy in Progress

## Troubleshooting Shadow Copy Creation

Exchange shadow copy creation failures are logged in both the RepliStor message logs and the Application Event Log. Microsoft logs all application specific information in the Event Log which consists of the Application Event log and the System log. Thus, any given message may be logged in either the Application Event log or the System log.

When troubleshooting shadow copy creation, consider the following:

- ◆ Check the RepliStor messages log. To see RepliStor messages in the message pane, right click and enable Info, optionally Debug messages can be enabled.
- ◆ To enable debug messages edit the following registry key and set the value to 1:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\  
RepliStorControl\Parameters\Enable Debug Log
```

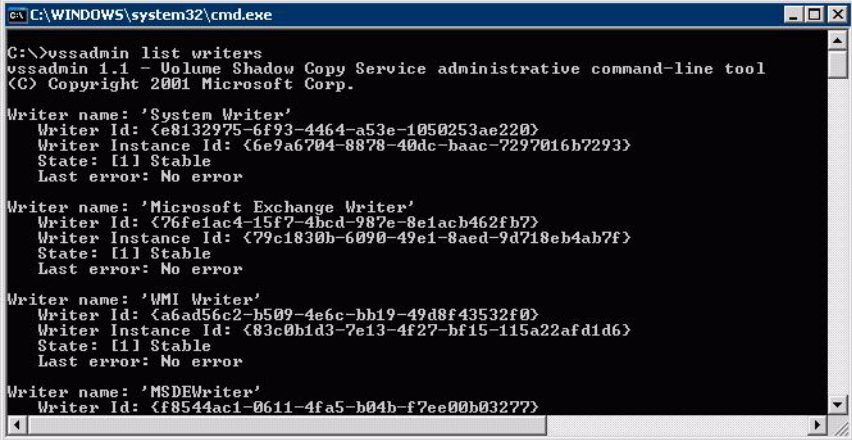
- ◆ After setting the key, restart the RepliStor server and control services.

In addition to the logs, querying the status of the writer can also determine its current state. For example the writer may not have been properly started or a backup may have been in progress. If the event log is backed up, the VSS admin commands can give better results.

- ◆ Issue the following command from a DOS window to view the state (refer to Figure 3-11):

```
Vssadmin list writers
```





```
C:\WINDOWS\system32\cmd.exe
C:\>vssadmin list writers
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line tool
(C) Copyright 2001 Microsoft Corp.

Writer name: 'System Writer'
  Writer Id: {e8132975-6f93-4464-a53e-1050253ae220}
  Writer Instance Id: {6e9a6704-8878-40dc-baac-7297016b7293}
  State: [1] Stable
  Last error: No error

Writer name: 'Microsoft Exchange Writer'
  Writer Id: {76fe1ac4-15f7-4bcd-987e-8e1acb462fb7}
  Writer Instance Id: {79c1830b-6090-49e1-8aed-9d718eb4ab7f}
  State: [1] Stable
  Last error: No error

Writer name: 'WMI Writer'
  Writer Id: {a6ad56c2-b509-4e6c-bb19-49d8f43532f0}
  Writer Instance Id: {83c0bid3-7e13-4f27-bf15-115a22afd1d6}
  State: [1] Stable
  Last error: No error

Writer name: 'MSDEWriter'
  Writer Id: {f8544ac1-0611-4fa5-b04b-f7ee00b03277}
```

Figure 3-11 Vssadmin list writers command

If any error is returned, restart the Volume Shadow Copy Service and the Exchange Information Store Service from the Services Applet to reset the writer state.



This chapter provides instructions for configuring options for RepliStor software. These options determine how the RepliStor client window appears and how RepliStor software operates. If you choose not to configure options, RepliStor software uses the default settings.

This chapter contains the following sections:

- ◆ Configuring Client Options..... 4-2
- ◆ Configuring RepliStor Options ..... 4-3
- ◆ Configuring the Updating and Forwarding Options ..... 4-5
- ◆ Configuring Throttling Options..... 4-6

## Configuring Client Options

You can configure client options to change the look and feel of the RepliStor client window and to enable TCP/IP communication.

To configure RepliStor client window options:

1. Start the RepliStor client and attach to the site for which you are configuring client options.
2. Select **Client Options** from the **Maintenance** menu.

The **Client Options** dialog box opens, as shown in Figure 4-1.

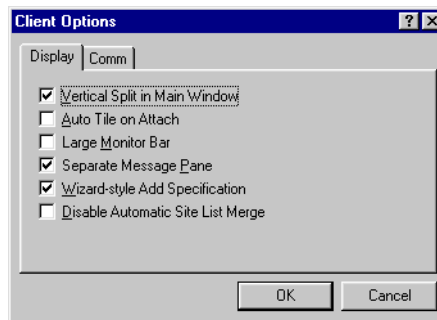


Figure 4-1 Client Options Dialog Box

3. Select the appropriate options on the **Display** and **Comm** tabs. Refer to the online help for a description of each option. For information on configuring client options, refer to *The RepliStor Client Elements* on page 2-7.
4. Click **OK** to save the options.

## Configuring RepliStor Options

To configure RepliStor options that affect the product's operation:

1. Start the RepliStor client and attach to the site for which you are configuring options.
2. Select **Options** from the **Maintenance** menu.

The **Options** dialog box opens, as shown in Figure 4-2.

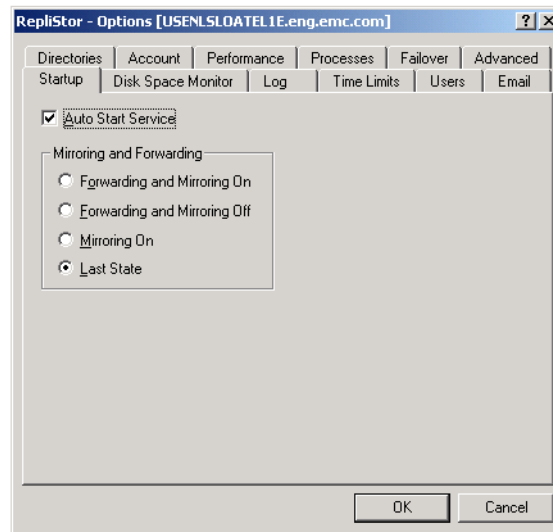
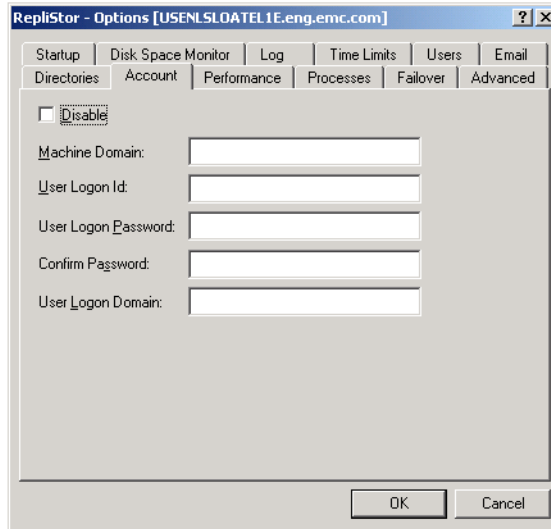


Figure 4-2 RepliStor Options Dialog Box

3. Use the tabs to access the options you want to configure. Refer to the online help for a description of each option.

Fill in the **Account** tab, as shown in Figure 4-3, if you want to modify dynamic DNS entries during an Alias failover. The DNS server is configured for secure updates only.



**Figure 4-3** RepliStor Options Account Tab

The **Account** tab is used to provide domain administrator credentials that may be required in an Alias failover.

4. Click **OK** to save the options.

## Configuring the Updating and Forwarding Options

RepliStor software allows you to schedule updating and forwarding. *Updating* is a function on the target system during which updates from the source are applied to data files on the target system. *Forwarding* is a function on the source system during which changes to data are forwarded from the source to the target system. Forwarding and updating are always on by default.

You can schedule *forwarding* or *site forwarding*. Turning off forwarding not only discontinues forwarding, but also discontinues the heartbeat signal sent from the source system to the target system.

### Configuring Updating and Forwarding

To configure updating and forwarding options:

1. Start the RepliStor client and attach to the site for which you are configuring updating and forwarding options.
2. Select **Schedule Options** from the **Maintenance** menu. The **Schedule Options** dialog box opens, as shown in Figure 4-4.

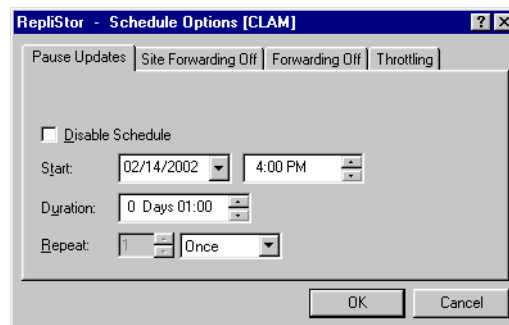


Figure 4-4 Schedule Options Dialog Box

3. Use the **Pause Updates**, **Site Forwarding Off**, and **Forwarding Off** tabs to configure updating and forwarding options. Refer to the online help for a description of each option.

If the **Disable Schedule** option is selected, clear it to enable you to set the options.

4. Click **OK** to save the options.

## Configuring Throttling Options

Throttling is an optional feature that can give you greater control over network bandwidth. With *throttling*, you can determine the maximum data rate, in kilobytes per second, for your replicated data. By setting this rate, you ensure the amount of replicated data transferred across the network from the current source system will never exceed the maximum data rate, as long as throttling is enabled.

Throttling can be permanently enabled, or turned on and off on a schedule. You can also configure the throttling rate for all target sites, each target site individually, or no target sites.

RepliStor software also can protect against poorly designed throttle settings. Since it is possible for replicated data to accumulate faster than the throttle rate, the files on the source system could grow until all disk space is exhausted. To avoid this, enable the RepliStor **Disk Space Monitor** options.

Make sure the **Automatically Recover** checkbox is checked at the bottom of the **Disk Space Monitor** tab. If the captured file operations cannot be sent to the target as fast as they are being captured, they will start to queue on the source. If the kernel cache overflows, it will create `OC$nnnnn` files in the RepliStor data directory to hold the captured file operations. If the disk becomes full (as set by the Disk Space Monitor), then replication stops. If the **Automatically Recover** checkbox is not checked, you must do the following:

1. Delete the data directory
2. Enable mirroring
3. Resynchronize all specifications

If the **Automatically Recover** checkbox is checked, RepliStor will automatically perform the above steps.

When the **Automatically Recover** option is enabled, the throttle rate is automatically disabled if the free disk space goes below the Warning level. It is re-enabled automatically when the free disk space goes above the Warning level.



## Configuring Throttling Options

To configure throttling options:

1. Start the RepliStor client and attach to the site for which you are configuring throttling options.
2. Select **Schedule Options** from the **Maintenance** menu.
3. In the **Schedule Options** dialog box, click the **Throttling** tab.
4. Clear the **Disable Schedule** option to enable the throttling options, as shown in Figure 4-5.

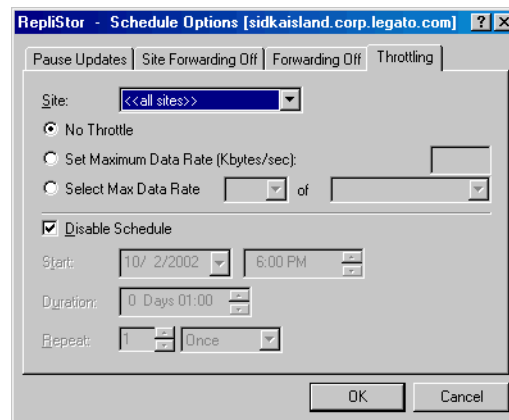


Figure 4-5 Throttling Options

5. Set throttling options as desired. Refer to the online help for a description of each option.
6. Click **OK** to save the options.



You can configure RepliStor software to fail over to a target system when the source system fails. This chapter describes how to configure RepliStor failover options.

This chapter contains the following sections:

- ◆ Alias Failover ..... 5-2
- ◆ Failover Setup Roadmap ..... 5-4
- ◆ Configuring Failover on the Source System ..... 5-5
- ◆ Configuring RepliStor Software with Two NIC Adapters ..... 5-14
- ◆ Configuring Manual Alias Activation ..... 5-15
- ◆ Configuring Failover across a WAN ..... 5-17
- ◆ Configuring Failover for Macintosh Volumes ..... 5-20

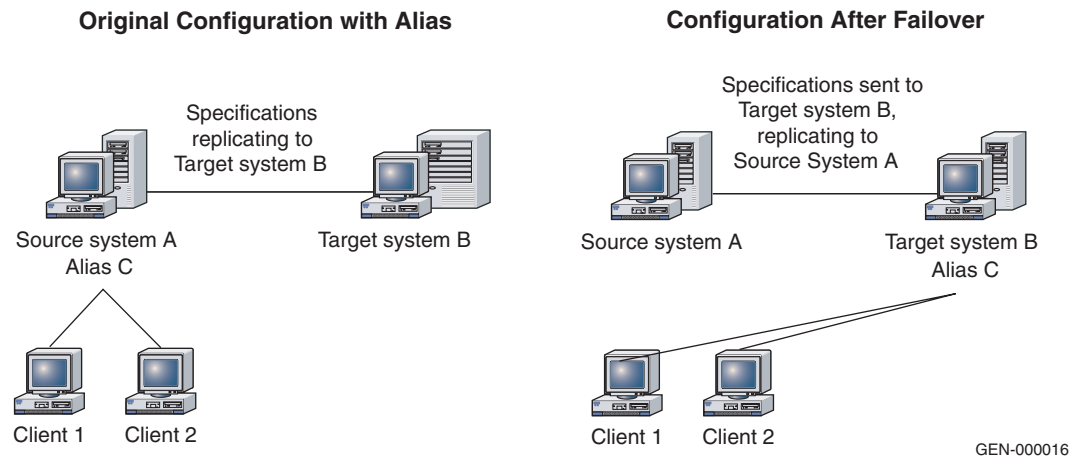
## Alias Failover

RepliStor software provides the Alias failover option, which allows a target system to take over automatically for a source system when the source fails. Configuring the source and target options using Alias failover allows you to determine when and how the failover occurs.

With Alias failover, you can define any number of source system aliases to use during failover, along with IP address and DNS server information.

Aliases allow you to fail over between two Windows 2000 or later servers. They also make it easier to recover from a failover with minimal disruption of client connections. With an alias, the source system can be brought up without affecting client connections to the alias on the target system. This allows the source system to be repaired and brought up, and the data synchronized back to the old source without having to bring down the target system and disconnect any clients using the target system.

Figure 5-1 on page 5-3 shows that clients connected to the alias on the source before failover are still connected to the alias on the target after a failover. Specifications can be sent automatically to the target and continue to replicate back to the source system. When the source is restored, simply remove the source machine name from the target, and all aliases and specifications are removed from the target automatically. The aliases are added back to the source. You then re-enable and synchronize the specifications on the source; client connections are disrupted for only the few seconds it takes to bring up the alias.



**Figure 5-1 Using Alias Failover**

Alias failover enables you to:

- ◆ Fail over multiple aliases or virtual machine names.
- ◆ Associate a list of IP addresses with each alias, and optionally register the IP addresses with the dynamic DNS server. The aliases are added to the target system at failover.
- ◆ Fail over between two Windows 2000 or later servers hosting Active Directory.
- ◆ Recover from a failover faster and easier, with minimal disruption of client connections.
- ◆ Implement WAN failover, where IP addresses cannot be forwarded if source and target systems are in different subnets.
- ◆ Specify a target system to take over for the source when defining the alias. No other target is sent the alias. If a target is specified with Alias failover and any of the source systems fail, that target is sent the source name.

## Failover Setup Roadmap

The following roadmap provides an overview of the tasks involved in setting up failover.

**Task 1:** Start the RepliStor client and attach to the source system. Make sure the RepliStor server is running.

**Task 2:** Configure failover options for the source system:

- ◆ For Alias failover in a Windows 2000 or later domain controller environment or for failover across a WAN, define one or more source system aliases. You also can specify IP addresses (except for WANs) and optional DNS information to be used during failover.
- ◆ In an environment where there are too many variables to allow an automatic failover, such as a WAN failover, you may want to define the alias for Manual Activation on the target system. For more information about *manual alias activation*, refer to *Configuring Manual Alias Activation* on page 5-15.
- ◆ Decide which, if any, services need to be started when the alias is activated on the source and target.
- ◆ Determine if the specifications should be transferred to the target during a failover.
- ◆ Decide if the specifications should automatically synchronize.
- ◆ Establish who will be notified when a failover has occurred.

For more detailed information about source system options, refer to *Configuring Failover on the Source System* on page 5-5.

## Configuring Failover on the Source System

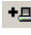
Use the steps below to configure Alias failover options for the source system; you must be attached to the source system.

### Configure a Failover

To configure a failover:

1. Fill in the **Account** tab with domain credential information if you want to create a Service Principal Name (SPN) for the alias. To enter domain credential information:
  - a. In the RepliStor client window, select **Options** from the **Maintenance** menu.
  - b. On the **Account** tab, enter the domain credentials.

You do not have to enter the domain credentials before configuring a failover if you are not creating a SPN for the alias.

2. Select **Alias** from the **Maintenance** menu or select the Alias Maintenance button  from the toolbar. The **Alias Maintenance** dialog box opens, as shown in Figure 5-2.

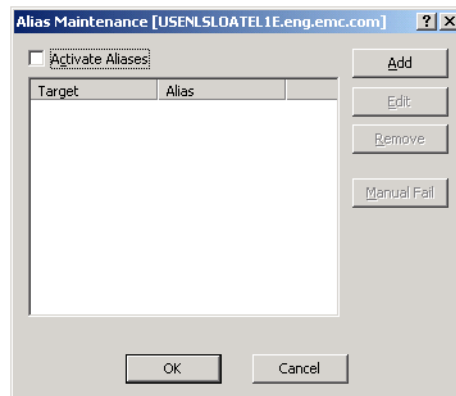
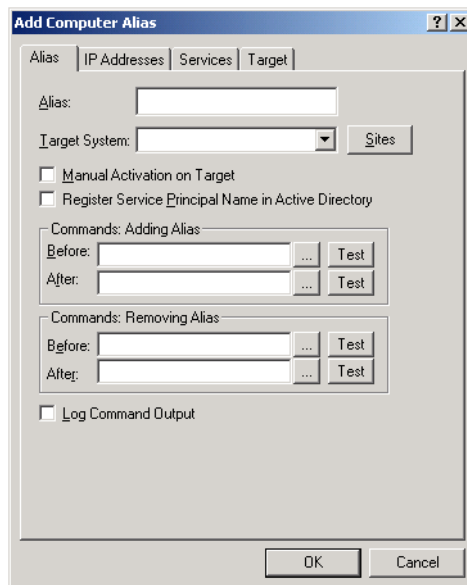


Figure 5-2 Alias Maintenance Dialog Box

3. Use the **Alias Maintenance** dialog box to define aliases by clicking the **Add** button.

The **Add Computer Alias** dialog box opens, as shown in Figure 5-3.



**Figure 5-3 Add Computer Alias Dialog Box — Alias Tab**

For more information on writing failover scripts, refer to *Using Scripts* on page 5-10.

4. Enter alias information as needed on the **Alias** tab.
5. On the **IP Addresses** tab, as shown in Figure 5-4 on page 5-7, enter information about one or more IP addresses to associate with the source system alias. This IP address information is sent to the target on a failover. You can optionally specify that the alias be registered with a DNS server.



**Important:** If you are setting up a failover across a WAN, specify DNS server information only, not IP addresses.





Figure 5-4 Add Computer Alias — IP Addresses Tab

6. Click the **Services** tab to specify services to be started or stopped on the system hosting the alias (either the source or target system) in the event of a failover, as shown in Figure 5-5.



**Figure 5-5 Add Computer Alias — Services Tab**

7. Click the **Target** tab to specify whether you want to transfer specifications to the target or autosynchronize the target (Figure 5-6 on page 5-9).

Click **Transfer Specifications to Target** if you want the RepliStor specifications on the source system to be re-created on the target system during a failover. If you check this box, the specifications are re-created on the target system in a mirror image of how they were on the source. For example, `D:\data` on the source was replicated to `E:\backup\data` on the target. During a failover, a specification will be created from `E:\backup\data` to the original source `D:\data`. If you do not select **Transfer Specifications to Target**, the specifications configured on the source system are lost and must be reconfigured.

Click **Auto Synchronize** if you want to do an incremental synchronization of data when the source system is available again.

Configure the **Notifications** options to specify who RepliStor will notify when the failover occurs.

You can test notification by clicking **Test** in the **Notifications** group box.

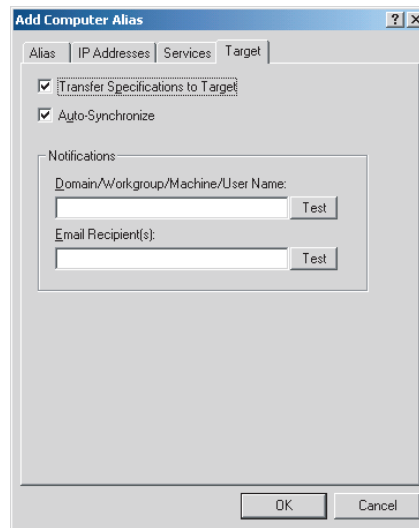


Figure 5-6 Add Computer Alias — Target Tab

Click **OK**. The **Alias Maintenance** dialog box reappears listing the new alias and its target, as shown in Figure 5-7.

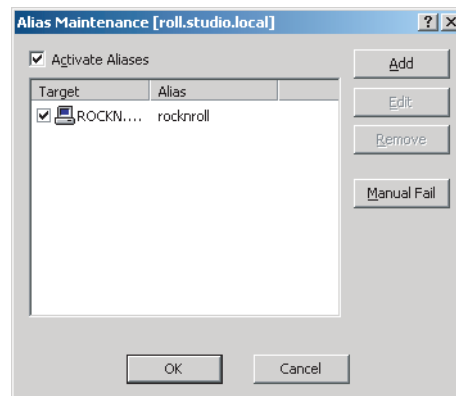


Figure 5-7 Alias Maintenance Dialog Box with New Alias and Target Listed

8. Make sure the checkbox next to the alias is selected.
9. Select the **Activate Aliases** checkbox, and then click **OK**.

## Using Scripts

A script can be any executable or batch file. Before RepliStor software runs a script, it does the following:

- ◆ Modifies the `PATH` environment variable by inserting the RepliStor install directory and default script directory paths in front of the variable.
- ◆ Sets the current directory as the RepliStor default script directory.
- ◆ Sets the environment variables shown in Table 5-1.

As long as the scripts are kept in the default script directory, it is not necessary to enter path information in the **Script Directory** text box in the **Options** dialog box (on the **Directories** tab). You also can create a *script repository* for centralized management of all scripts for a set of servers. For more information about script repositories, refer to *Script Repository* on page 5-11.

**Table 5-1 Environment Variables Set When Running Executables**

Variable	Description	Example
ALIAS	If the script is associated with the activation or deactivation of an alias, the variable contains the alias name.	aliasname
DNSCOMPUTERNAME	Fully qualified DNS name of local computer.	name1.company.com
DOMAINCOMPUTERNAME	NetBIOS name of domain controller. To make this available, configure the <b>Machine Domain</b> attribute in the <b>Options</b> dialog box (on the <b>Account</b> tab).	DCSERVER
SOURCESITE	Fully qualified DNS name of the computer that failed and caused the failover.	source.company.com
SOURCESITENETBIOS	NetBIOS of computer that failed and caused the failover.	SOURCE
TARGETSITE	Fully qualified DNS name of the computer that is the target of the failover (alias scripts only).	target.company.com
TARGETSITENETBIOS	NetBIOS of computer that is the target of the failover (alias scripts only).	TARGET

Table 5-1 Environment Variables Set When Running Executables (continued)

Variable	Description	Example
UserDNSComputerDN	Fully qualified distinguished name of local computer.	CN=NAME1 CN=Computers DC=company, DC=com
UserDNSDomainDN	Fully qualified distinguished name of domain.	DC=company, DC=com

If a script requires any features of the command shell (`cmd.exe`), specify this in the script command line. For example, running a batch file (`.bat`) may require features of the command shell, so the command line would be entered as:

```
cmd /c batchfile.bat %ALIAS%
```

If a script is required to run when an alias is created or removed, then that script file must reside on both the source and target systems, since the alias may exist on either machine. Use a script repository so that only one script directory is necessary for any number of servers.

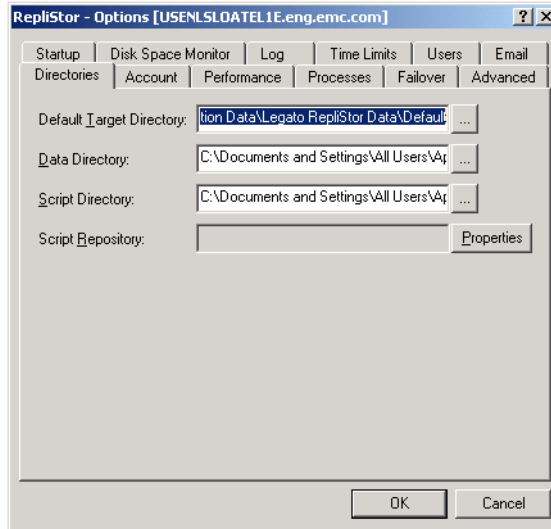
### Script Repository

RepliStor software supports the creation of a script repository, which allows a script to be centrally located and administered. A script repository is a directory on a file server accessible to all RepliStor servers using a universal naming convention (UNC) path. If a script repository is defined, then before executing any script, RepliStor software copies the script from the repository to the local script directory. After script execution, RepliStor software deletes the script file from the local directory.

Using a script repository requires no changes to the way scripts are written or specified. The only exception is if a script requires additional script files not specified on the command line. RepliStor software automatically parses the command line and determines the script file that needs to be copied from the repository. If any additional files need to be copied, specify the files by enclosing them in parentheses.

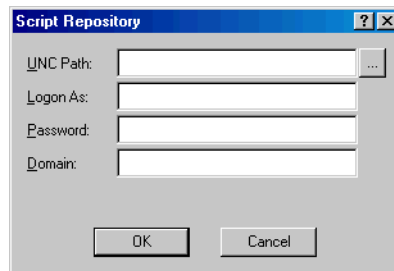
**To create a script repository:**

1. In the RepliStor client window, select **Options** from the **Maintenance** menu.
2. In the **Options** dialog box, click the **Directories** tab, as shown in Figure 5-8.



**Figure 5-8 Options Dialog Box—Directories Tab**

3. Click **Properties**. The **Script Repository** dialog box opens, as shown in Figure 5-9.



**Figure 5-9 Script Repository Dialog Box**

4. In the **UNC Path** text box, enter the UNC path to the script repository using the following format:  

```
\\computer\share\dir1\dir2
```
5. In the **Logon As** text box, enter the logon ID used to access the script repository.
6. In the **Password** text box, enter the password used to access the script repository.
7. In the **Domain** text box, enter the domain associated with the script repository. If the computer is not in a domain, enter the computer name.
8. Click **OK**. The path to the script repository appears in the **Script Repository** list on the **Directories** tab.

### Example: Using a Script Repository

Enter the following command at the command line to instruct RepliStor software to copy `masterscript.pl`, `file1.pl`, and `file2.pl` from the script repository to the local directory:

```
(file1.pl;file2.pl)cmd /c masterscript.pl parameter
```

The script/executable files must include the file extension. For example, `.pl`, `.exe`, `.bat`.

### Manually Forcing a Failover with Aliases

To test the configuration, you can manually force a failover. To manually force a failover if you have aliases defined:

1. In the RepliStor client window on the source system, select **Alias** from the **Maintenance** menu.
2. In the **Alias Maintenance** dialog box, make sure one or more aliases have been activated with the **Activate Aliases** option.
3. Click **Manual Fail**.

RepliStor software sends all selected aliases to their respective target systems. It also removes the selected aliases and associated IP addresses from the source system.

Alternatively, select **Aliases** on the tree pane, right-click on any alias, and select **Manual Fail**.

## Configuring RepliStor Software with Two NIC Adapters

You can configure RepliStor software to use two NIC adapters, which enables you to use:

- ◆ One LAN interface for all file mirroring traffic. This guarantees available bandwidth for RepliStor software and does not interfere with other LAN traffic.
- ◆ Both LAN interfaces to determine if the source computer has failed.
- ◆ If the connection to the target fails over one LAN interface, it will attempt to reestablish it over all LAN interfaces.

This section provides an example of how to specify such a configuration. Assume you have the following scenario:

- ◆ A source computer named `SRC`.
- ◆ A target computer named `TARGET`.
- ◆ `SRC` and `TARGET` are connected on the public LAN, along with the other computers in the environment. This LAN contains the DNS server that is used to resolve `SRC` and `TARGET` to their IP addresses.
- ◆ A dedicated LAN uses the second NIC on both `SRC` and `TARGET` and is connected to no other computer.
- ◆ The IP address of the dedicated LAN is not kept in the DNS server or host file.
- ◆ `SRC`'s dedicated LAN IP address is `192.168.0.1`
- ◆ `TARGET`'s dedicated LAN IP address is `192.168.0.2`

Configure RepliStor software so that all mirroring traffic goes over the dedicated LAN. If `SRC` fails, then `TARGET` should ping `SRC` on both the dedicated and public LAN.

1. In the RepliStor client window, select **Attach** from the **Functions** menu.
2. In the **Site List** window, click **Add**.
3. In the **Site Properties** dialog box, create two site entries for the dedicated LAN interface on both `SRC` and `TARGET`.

For the first entry, use the following information:

- **Site:** `TARGET_DED`



- **Description:** TARGET over dedicated LAN from SRC
- **Account, Domain, and Password:** Enter as needed (usually only required if SRC is in a different domain)
- Select **Use IP Address** and enter 192.168.0.2

For the second entry, use the following information:

- **Site:** SRC\_DED
  - **Description:** SRC over dedicated LAN from TARGET
  - **Account, Domain, and Password:** Enter as needed (usually only required if SRC is in a different domain from TARGET)
  - Select **Use IP Address** and enter 192.168.0.1
4. Create one or more specifications on SRC using TARGET\_DED as the target system.
  5. Create one or more aliases on SRC using TARGET\_DED as the target system.

---

## Configuring Manual Alias Activation

You may have a situation where there are too many variables to allow an automatic failover between systems; for example, across a WAN. In this case, you can configure Alias failover with aliases that require you to manually activate the aliases on the target system before the failover actually occurs.

### To Configure Manual Alias Activation:

1. Attach to the source system and the target system.

#### On the Source System:

2. In the RepliStor client window, select **Options** from the **Maintenance** menu.
3. In the RepliStor **Options** dialog box, select the **Account** tab, as shown in Figure 4-3 on page 4-4. Specify domain credentials for the source system for which you are defining an alias.
4. Select **Alias** from the **Maintenance** menu.
5. In the **Alias Maintenance** dialog box, select **Add** to add a new alias for the source system.

6. In the **Add Computer Alias** dialog box, specify attributes as needed, and then select **Manual Activation on Target** (refer to Figure 5-3 on page 5-6).
7. Click the **IP Addresses** tab to optionally specify the following, and then click **OK**:
  - DNS information for adding the alias to DNS servers in the environment.
  - IP and subnet information to send to the target on a failover.
 The new alias and its target appear in the alias list in the **Alias Maintenance** dialog box.
8. Make sure the checkbox next to the alias is selected, select **Activate Aliases**, and then click **OK**. The new alias appears in the directory tree listed under the source system name.
9. Configure additional target options as needed, and then click **OK**. Refer to the online help for a description of each option.

On a failover, RepliStor software sends the alias to the target system, but does not actually create the alias or IP addresses on this system until you manually activate it.

### Activating the Alias on the Target System after a Failover

To activate the alias on the target system after a failover:

1. In the tree pane of the RepliStor client window on the target system, select **Aliases**.
2. Right-click on the alias name in the list pane and select **Activate** (Figure 5-10).




Figure 5-10 Activating Aliases on the Target System

---

## Configuring Failover across a WAN

Generally, it is not possible to perform a failover across a WAN because IP addresses cannot be forwarded across different subnets. However, in a Windows 2000 or later environment with dynamic DNS server configured, you can set up a failover between a source and target system over a WAN.

### How to Configure an Alias Failover across a WAN

1. On the source system, define an alias for the source. For detailed instructions on defining an alias, refer to *Configure a Failover* on page 5-5.
2. In the **Add Computer Alias** dialog box, as shown in Figure 5-3 on page 5-6, add alias and target information.
3. Select the **IP Addresses** tab, as shown in Figure 5-4 on page 5-7.
4. Select **Add to DNS Servers**.
5. Enter the **DNS Domain** and **Time To Live** information, and then click **OK**.

The **Time to Live** value determines how quickly clients see any changes to the DNS record on the server. On a failover, the DNS entry is modified to refer to the alias sent to the target.

*Important:*

Do not specify any IP addresses to associate with the alias.

6. Specify the remaining failover options for the source and target as needed. Refer to the online help for a description of each option.

---

## Failover Agent for WAN Failover

When failing over between two servers across a WAN, there are special considerations. There are more failure modes that can occur besides a source system failure. For example, a WAN link or VPN connection may fail. A typical failover sequence is:

1. The source system fails and heartbeats stop to the target.
2. The target notices the heartbeats have stopped and pings the source.
3. If the ping is successful, the failover is aborted.
4. If the ping is not successful, the failover occurs.

In the above sequence, if the WAN had failed and not the source server, the target would have taken over for the source and there would be two systems performing the functions of the source, with possibly some clients connected to the source and others connected to the target. This is called a *split brain*<sup>1</sup> condition.

To prevent a split brain condition from occurring, you can configure a Failover Agent. A Failover Agent is a third server that is local to the source system, preferably with at least two LAN connections to the source. A Failover Agent must have RepliStor running on it (either a full installation or just the Administrative client).

On the source system, enter the name of the Failover Agent in the **Failover** tab on the **Options** dialog box (Figure 5-11).

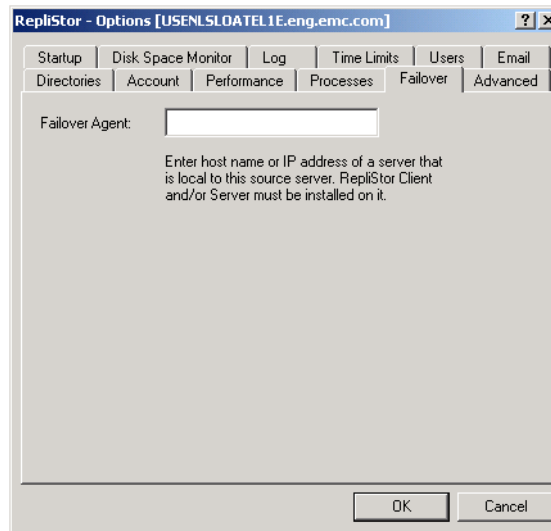


Figure 5-11 Failover Tab on the Options Dialog Box

1. A total communication failure, while both nodes remain operational, is referred to as a *split brain* condition and is a potential cause of logical data corruption. For example, if both sides assume that the other is dead and begin processing new transactions against their copy of the data, two separate and unreconcilable copies of the data can be created.

With the above configuration complete, a WAN failover will now work like this:

1. The source system fails and heartbeats stop to the target.
2. The target notices the heartbeats have stopped and pings the source
3. If the ping is successful, the failover is aborted.
4. If the ping is not successful, the failover continues.
5. The target sees there is a Failover Agent configured on the source. It attempts to connect to the agent. It will then instruct the agent to ping the source:
  - If the target can connect to the agent, and the agent can ping the source, the failover is aborted.
  - If the target can connect to the agent, and the agent cannot ping the source, the failover continues.
  - If the target cannot connect to the agent, the failover is converted to a **Manual Activate** failover.

## Configuring Failover for Macintosh Volumes

Configure specifications for replicating Macintosh volumes as described in *Macintosh Specifications* on page 2-36. Configure failover for Macintosh volumes as follows:

1. On both source and target systems, stop the File Server for Macintosh service and set it to **Manual**.
2. In the RepliStor client window on the source system, select **Alias** from the **Maintenance** menu.
3. In the **Alias Maintenance** dialog box, select **Add** to add a new alias for the source system.
4. In the **Add Computer Alias** dialog box, give the alias a unique network name and specify a target system to take over for the source. Specify additional attributes as needed.
5. Click the **IP Addresses** tab to optionally add an IP address to which the Macintosh clients can connect.
6. Click the **Services** tab.
7. Add **File Server for Macintosh** to the **Services to Start** box, and then click **OK**.

The new alias and its target appear in the alias list in the **Alias Maintenance** dialog box.

8. Make sure the checkbox next to the alias is selected, select **Activate Aliases**, and then click **OK**. The new alias appears in the directory tree listed under the source system name.

---

This chapter describes the administrative tasks you can perform once RepliStor software is installed and configured.

This chapter contains the following sections:

- ◆ Using the Windows Performance Monitor ..... 6-2
- ◆ Using the Performance Tab in the Options Dialog Box..... 6-5
- ◆ Using the Processes Tab in the Options Dialog Box..... 6-7
- ◆ Administering RepliStor Sites ..... 6-10
- ◆ Targeting a Remote Share Hosted by a NAS Device ..... 6-23
- ◆ Administering Specifications ..... 6-30
- ◆ Turning Processes On and Off..... 6-41
- ◆ Working with the Message Log ..... 6-45
- ◆ Closing Application Files..... 6-48
- ◆ Using Tape Backups with RepliStor Software ..... 6-49
- ◆ Reporting/Simulating Bandwidth Between Source and Target6-50

## Using the Windows Performance Monitor

You can use the Windows Performance Monitor to view statistics on the performance of RepliStor software. To start the Performance Monitor:

1. In the RepliStor client window, select **Performance Monitor** from the **Help** menu.

The **Performance** window opens, with counters for the RepliStor server object already selected and active, as shown in Figure 6-1.

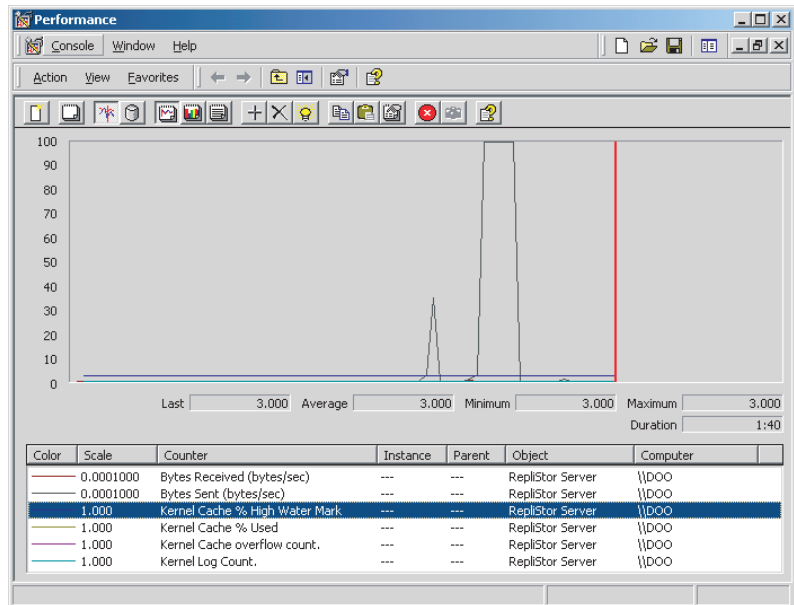


Figure 6-1 RepliStor Performance Monitor Window

2. To monitor kernel counters for RepliStor sites and specifications while replicating data, click the plus (+) button.
3. In the **Add Counters** dialog box select one of the following from the **Performance object** list:
  - **RepliStor Server**
  - **RepliStor Site**
  - **RepliStor Specification**
  - **RepliStor Target**



Site and Specification objects only exist if you have at least one file specification.

4. In the **Counter** list, select a counter. To see the counter's description, click **Explain**.
5. Select an instance.
6. Click **Add**.

---

## Sizing the Kernel Cache

The kernel cache is a fixed amount of shared memory used for queuing data going from the source to the target. When you install the RepliStor software, it looks at the physical memory and pre-configures an appropriate setting for the system. However, you can modify this parameter, if needed, in the RepliStor **Options** dialog box (on the **Advanced** tab).

It is important to size this memory segment correctly. If it is too small, output from memory operations may overflow to disk (that is, OC\$nnnnn files in the RepliStor data directory), hurting performance. If the memory segment is too large, it reduces the amount of memory available to other applications.

Use the following Performance Monitor counters to determine the correct size of the kernel cache:

- ◆ **Kernel Cache % Used**
- ◆ **Kernel Cache % High Water Mark**
- ◆ **Kernel Cache Overflow Count**
- ◆ **Kernel Log Count**
- ◆ **Kernel Log Count High Water Mark**

## Determining the Kernel Cache Size

To determine the kernel cache size:

1. Use RepliStor software with the default kernel cache size.
2. In the Performance Monitor, note the value of the **Kernel Cache Overflow Count**.
3. Use RepliStor software over a time period that typifies normal operations, including high-traffic times.
4. Check the **Kernel Cache Overflow Count**. If it has increased, the kernel cache is probably too small.

5. Check the **Kernel Log Count High Water Mark** to determine how much it overflowed. Each log holds approximately 1 MB of data, so increase the kernel cache by that many megabytes, plus an additional percentage.
6. If the **Kernel Cache Overflow Count** did not increase, check the **Kernel Cache % High Water Mark**. If it is greater than 70 percent filled, then the kernel cache is sized correctly. If it is less than that, you can reduce its size. However, do not set the kernel cache less than 4 MB.

If the specifications include a large number of files, performing a synchronization may overflow the kernel cache. However, if the cache is sized correctly for normal operations, you do not need to increase it to accommodate the synchronization. This is because performance is not affected if the cache overflows during a synchronization, and it may not be possible to increase the kernel cache size enough to prevent this overflow.

## Using the Performance Tab in the Options Dialog Box

The **Performance** tab is accessed by selecting **Options** from the **Maintenance** menu. The **Performance** tab allows you to configure options to improve system performance (Figure 6-2). There are three preconfigured settings that can be selected on the **Performance** tab:

- ◆ **TCP/IP Buffer Size**
- ◆ **Full Sync Buffer Size**
- ◆ **Sync Target Connections**

You can select the default values for these options or you can modify these settings manually.

Table 6-1 on page 6-6 describes the various options on the **Performance** tab.

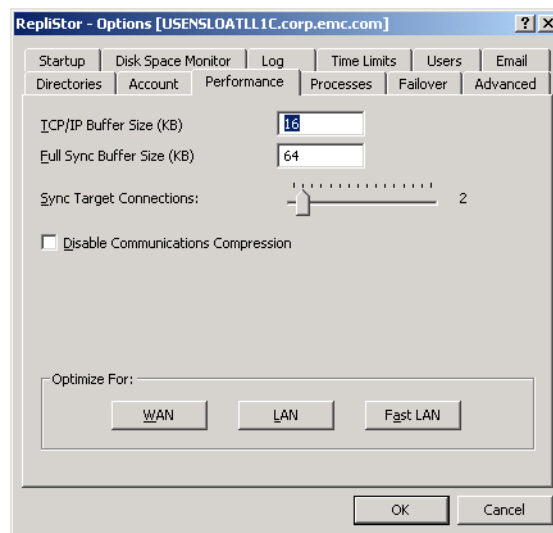


Figure 6-2 Performance Tab in the Options Dialog Box

Table 6-1 Performance Tab Field Descriptions

Field	Description
<b>TCP/IP Buffer Size</b>	The TCP/IP buffer controls the sizes of the send and receive buffers within RepliStor and the TCP/IP driver. A higher TCP/IP Buffer Size may provide better performance for high-speed networks but uses greater system memory. The default is 8 K.
<b>Full Sync Buffer Size</b>	The amount of data read and sent to the target in a single unit during a full synchronization. A higher number will have less overhead. Fast networks could benefit from a bigger buffer.
<b>Sync Target Connections</b>	Sets the number of connections the source server will connect to the target server during a sync operation. This setting can increase the performance of a sync operation, especially over fast LANs. This setting has no effect for real time mirroring. When the sync operation finishes, the number of connections will revert to 1 after a time-out period.
<b>Disable Communications Compression</b>	When checked, all outgoing connections will not be compressed. This option should be used if the network bandwidth is such that compression would not be beneficial and if conserving CPU activity is important.
<b>Optimize for WAN</b>	Sample settings for a low-speed connection.
<b>Optimize for LAN</b>	Sample settings for a 10/100 MByte LAN.
<b>Optimize for Fast LAN</b>	Sample settings for Gig LAN.

## Using the Processes Tab in the Options Dialog Box

The **Processes** tab on the **Options** dialog box allows you to ignore operations by process. The **Processes** tab functions when set on both the source and target server (Figure 6-3).

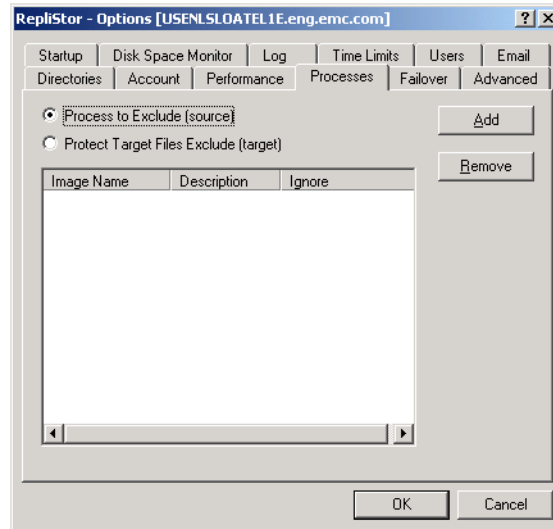


Figure 6-3 Processes Tab in the Options Dialog Box

### Source

A process list can be created that will cause RepliStor to ignore operations based on the process that is performing it. This is typically used in the following situations:

- ◆ **Antivirus scans** — An antivirus program will typically make attribute changes to a file during the scan. When each individual file has finished being scanned, the attributes are set back to the initial state. RepliStor will replicate these files, but this is an inefficient use of time and resources since the files are never changed in any way.
- ◆ **Backup programs** — A backup program will set the *archive* bit for each backed-up file. This causes RepliStor to replicate this activity even though it may not be required on the target.

- ◆ **DiskXtender 2000™** — DiskXtender 2000 may migrate a file from the source disk to an offline media. The operations used to perform this migration should not be replicated, since it may result in a truncation of that file's data on the target.

When adding a process to the list, you can choose to ignore all operations, or just ignore attribute operations. If you want to ignore just attribute operations, select the **Ignore Only Attribute Operations** checkbox in the **Add Process** dialog box (refer to Figure 6-5 on page 6-10). If you leave this checkbox unchecked, all operations will be ignored. In the above examples, the antivirus and backup programs would check the **Ignore Only Attribute Operations** checkbox option, and the DiskXtender 2000 would not check this option, since it would need to exclude all operations.

This list is global across all specifications.

### Target

On the target side, the process list is as an exception list for the **Protect Target Files** function. For example, assume that you have DiskXtender 2000 on the target, and DiskXtender 2000 needs read/write access to the target files in order to migrate them to and from offline storage. You can still enable **Protect Target Files**; but make sure to put DiskXtender 2000 in the list to allow only DiskXtender 2000 access to the files.

### Adding a Process Exclude

Clicking the **Add** button brings up the **Add Process** dialog box (Figure 6-5 on page 6-10). There are three ways to add a process:

- ◆ Select it from one of the currently running processes.
- ◆ If it is not currently running, enter the executable name of the application.
- ◆ Select it from a list of predefined applications.

An optional description may be entered.

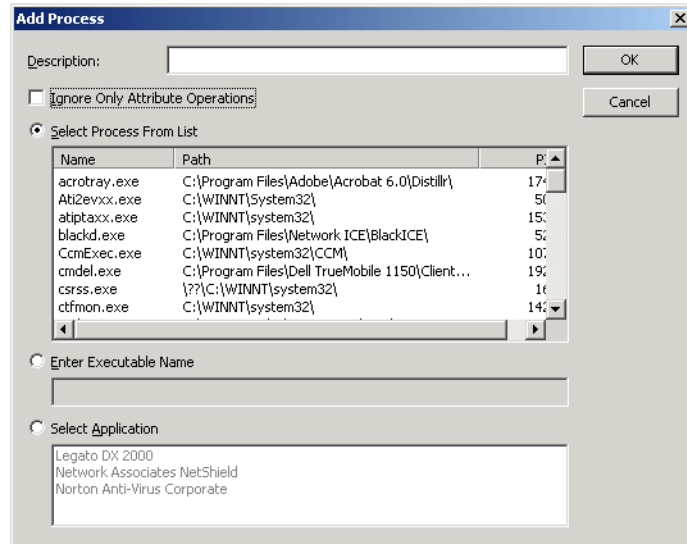


Figure 6-4 Add Process Dialog Box

## Administering RepliStor Sites

A RepliStor site is a Windows system, typically a server, running the RepliStor server. (For information on attaching to RepliStor sites, refer to *Attaching to Remote RepliStor Sites* on page 2-14.)

### Checking Site Status

To make sure RepliStor software is functioning normally, check the status of an attached site and its target sites. This section describes several methods for checking site status.

### Using Performance Information in the Client Window

The RepliStor client window provides a subset of the same performance information about the source server that is available in the Performance Monitor window, as shown in Figure 6-5.

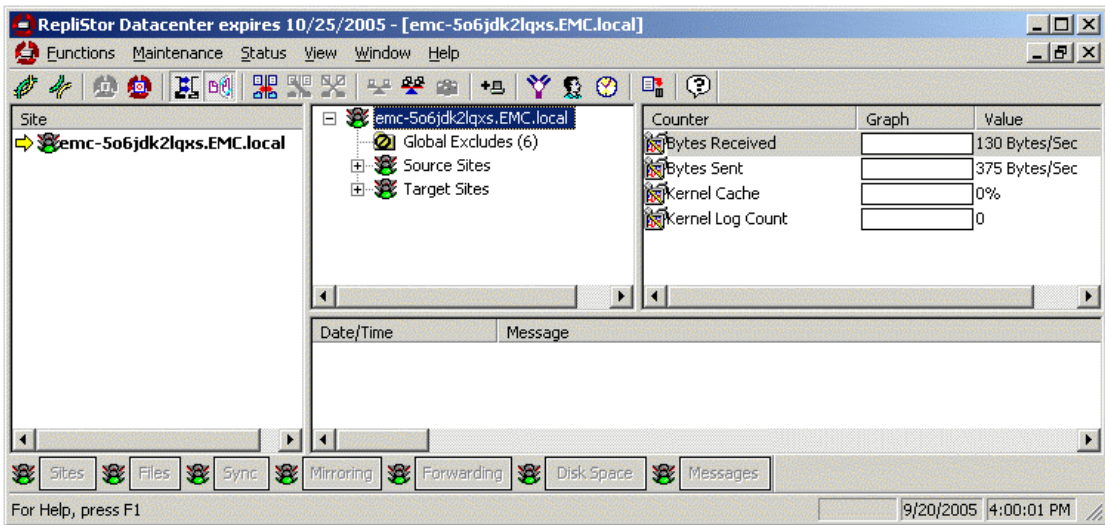



Figure 6-5 Performance Information in the Client Window



## Using Traffic Lights and Computer Icons in the Site Pane of the Client Window

The site pane, or the left pane of the RepliStor client window, shows the status of the sites to which you can attach. The site status is indicated by traffic lights and computer icons, as shown in Figure 6-6.



**Figure 6-6** Status of Sites in the Sites Pane of the Client Window

Table 6-2 provides information about each condition that the traffic light and computer icons can be in for each site in the site pane:

**Table 6-2** Traffic Light and Computer Icon Condition in the Site Pane

Icon	Condition	Status
Site Traffic Light	Off (no lights, solid black icon)	The site is not running, but RepliStor software is installed and the control service is running on that site.
	Dimmed (grayed out)	No RepliStor process is running on the site, or the site does not exist.
	Red Light	One or more of the following is true: <ul style="list-style-type: none"> <li>At least one unread <i>Severe</i> log message is at the site.</li> <li>The Disk Space Monitor is on and the amount of free space has dropped below the Stop level.</li> <li>At least one target site is blocked.</li> <li>At least one file is blocked.</li> </ul>
	Yellow Light	One or more of the following is true: <ul style="list-style-type: none"> <li>At least one unread <i>Warning</i> log message is at the site.</li> <li>The Disk Space Monitor is on and the amount of free space has dropped below the Notify level.</li> </ul>
	Green Light	All RepliStor systems are functioning normally at the site.

Table 6-2 Traffic Light and Computer Icon Condition in the Site Pane (continued)

Icon	Condition	Status
Computer Icon	Blue	The site is running a previous version of RepliStor. <i>Note:</i> The status of the site (that is, red, yellow, or green) cannot be obtained for sites running a previous version of RepliStor.

Note that the red, yellow, and green traffic light conditions of the sites in the sites pane of the RepliStor client window do not necessarily correspond to the red, yellow, and green traffic light conditions of the sites listed in the tree pane (refer to *Using Traffic Lights and Computer Icons in the Tree Pane of the Client Window* below). In the sites pane, the traffic light colors provide general information about the sites; in the tree pane, the traffic light colors provide more details on the status of the source and target servers.

### Using Traffic Lights and Computer Icons in the Tree Pane of the Client Window

The tree pane, or the upper-middle pane of the RepliStor client window, displays a hierarchical tree of information that shows the status of the source server and each of its target sites. The status of the server is indicated by traffic lights and computer icons, as shown in Figure 6-7. When the server is running, the traffic light is green.

Note that the traffic light status in the tree pane is cumulative as you move up the directory tree. The top traffic light displays the cumulative status of the traffic lights below it. For example, if there is a Warning log message at one site (that is, a yellow traffic light) and a blocked site (that is, a red traffic light), then the top traffic light will be red.

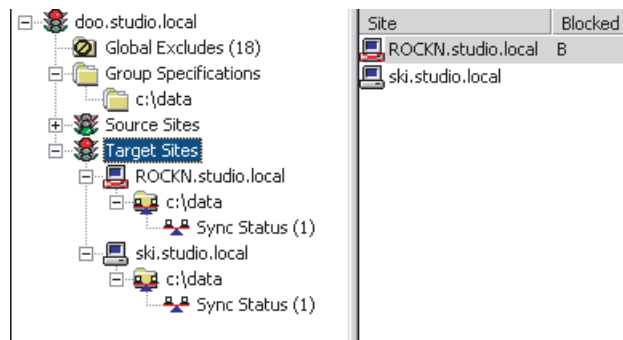


Figure 6-7 Status of the Source and Target Servers in the Tree Pane of the Client Window

Table 6-3 provides information about each condition that the traffic light and computer icons can be in for the source and target servers in the tree pane:

**Table 6-3 Traffic Light and Computer Icon Condition in the Tree Pane**

Icon	Condition	Status
Site Traffic Light	Off (no light, black icon)	The RepliStor server is not running at the attached site.
	Red Light	One or more target sites are blocked.
	Yellow Light	One or more of the following is true: <ul style="list-style-type: none"> <li>The site has at least one blocked file.</li> <li>One or more target sites have synchronization pending.</li> <li>One or more specifications are disabled.</li> </ul>
	Green Light	All RepliStor systems are functioning normally at the attached site and at all target sites.
Computer Icon	No X	The target site is functioning normally.
	Red X through CPU	The target site is blocked.
	Red X through monitor	The RepliStor client cannot connect to the site because of a communications problem.
	Red X through both CPU and monitor	The target site is blocked and has a communication problem.

**Using the Monitor Bar** You can use the monitor bar to check the status of RepliStor sites and to troubleshoot site problems. The monitor bar displays both the status of the site in the current window and the status of all sites to which the RepliStor client is currently attached.

To increase the monitor bar size:

1. Select **Client Options** from the **Maintenance** menu.
2. In the **Client Options** dialog box, select the **Display** tab.
3. Select the **Large Monitor Bar** option, and then click **OK**.

The monitor bar displays a series of traffic lights and buttons, as shown in Figure 6-8.



**Figure 6-8 RepliStor Monitor Traffic Lights**

Table 6-4 provides descriptions for the buttons on the monitor bar.

**Table 6-4 RepliStor Monitor Buttons**

Button	Condition of Traffic Light	Status
<b>Sites</b>	Off (no light)	No specification is defined, or the RepliStor client is not attached to the server at a site.
	Red Light	At least one attached site is blocked, or communication cannot be established between the RepliStor client and the server at a site.
	Yellow Light	Synchronization is pending for at least one attached site and/or one or more specifications are disabled.
	Green Light	No blocks or pending synchronization at any attached site.
<b>Files</b>	Off (no light)	The RepliStor client is not attached to the server at a site.
	Red Light	A file at a site is blocked, or any file referenced in a specification at a site is blocked.
	Green Light	No blocked files.
<b>Sync</b>	Off (no light)	The RepliStor client is not attached to the server at a site.
	Red Light	At least one pending synchronization.
	Yellow Light	At least one active synchronization that is not complete.
	Green Light	No pending or active synchronization.
<b>Mirroring</b>	Off (no light)	The RepliStor client is not attached to the server at a site.
	Red Light	At least one site has mirroring turned off.
	Green Light	Mirroring is on.
<b>Forwarding</b>	Off (no light)	The RepliStor client is not attached to the server at a site.
	Red Light	At least one site has forwarding turned off.
	Green Light	Forwarding is on.
<b>Disk Space</b>	Off (no light)	The RepliStor client is not attached to the server at a site, or disk space monitoring is disabled at all sites.
	Red Light	The amount of free disk space is less than the configured Stop Size.
	Yellow Light	The amount of free disk space is less than the configured Warning Size.
	Green Light	The amount of free disk space is more than the configured Warning Size.

Table 6-4 RepliStor Monitor Buttons (continued)

Button	Condition of Traffic Light	Status
Messages	Off (no light)	The RepliStor client is not attached to the server at a site.
	Red Light	The message log contains at least one severe message that is unread.
	Yellow Light	The message log contains at least one warning message that is unread.
	Green Light	The message log contains no unread messages that are severe or warning messages.

## Finding Target Site Information

To check the status of each target site defined for the source site, click **Target Sites** in the tree pane. RepliStor software lists the names of the target sites, along with detailed information in the right list pane, as shown in Figure 6-9.

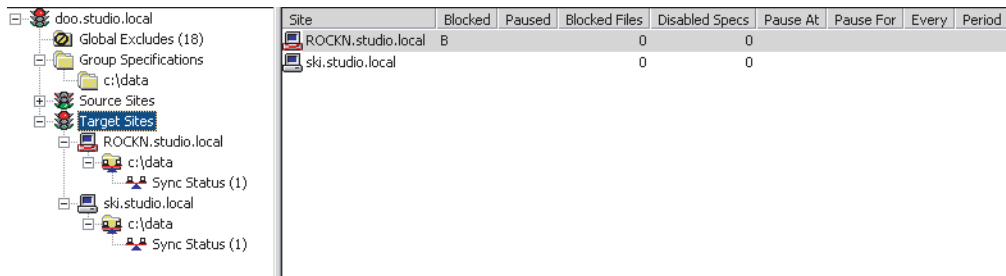


Figure 6-9 Viewing Target Site Information

Table 6-5 describes the status information provided in the right list pane for each target site.

Table 6-5 Target Site Status Information in the List Pane

List Pane Column Heading	Target Site Status Information
Site	Name of the target site.
Blocked	Information on the target site, which can be B for a blocked site or C for a communications error.
Paused	Pause updates has been enabled on the target site.

Table 6-5 Target Site Status Information in the List Pane (continued)

List Pane Column Heading	Target Site Status Information
<b>Blocked Files</b>	Number of blocked files at the target site.
<b>Disabled Specs</b>	Number of specifications that are disabled.
<b>Pause At, Pause For, Every, and Period</b>	Information on scheduled pauses in updating at the target site.

If a site is blocked, the **Target Sites** light is red, and the site's computer icon displays one or more red Xs.

### Finding Specification Information

To see all specifications configured for a target site, click the site's name in the tree pane. Detailed specification information appears in the right list pane, as shown in Figure 6-10.

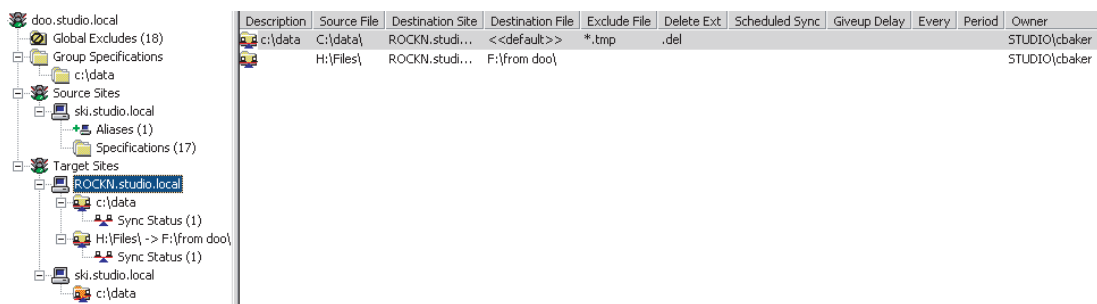


Figure 6-10 Viewing Specification Information

Table 6-6 describes the specification information provided in the right list pane for each target site.



Table 6-6 Target Site Specification Information in the List Pane

List Pane Column Heading	Target Site Specification Information
<b>Description</b>	Description of the specification.
<b>Source File</b>	Complete path name of the source specification.
<b>Destination Site</b>	Name of the target site to which data is mirrored.
<b>Destination File</b>	Complete path name of the target specification.
<b>Exclude File</b>	Specification of files excluded from mirroring.

Table 6-6 Target Site Specification Information in the List Pane (continued)

List Pane Column Heading	Target Site Specification Information
Delete Ext	Extensions used to store files on the target system that have been deleted from the source.
Scheduled Sync	Scheduled date and time that the specification will be synchronized.
Giveup Delay	Date and time when RepliStor software stops trying to synchronize the specification after synchronization is unsuccessful.
Every and Period	Period of time when RepliStor will resynchronize a specification.
Owner	Username of the person who created the specification.

## Checking Synchronization Status

When a specification is synchronized, a seesaw icon appears directly below the specification in the directory tree pane, as shown in Figure 6-11. When the seesaw icon is level , the synchronization is complete. When the seesaw icon is tilted , the synchronization is active and not complete.




Figure 6-11 Viewing Synchronization Status

If the seesaw is tilted, the synchronization is active and not complete. If the seesaw is straight across (as shown), the synchronization is complete. For a complete list and description of specification and synchronization icons, right-click the icon, then click **What's This?**

To see detailed information on the synchronization status, click the seesaw. An entry for each attempted synchronization on the specification appears in the right pane.

Table 6-7 describes the **Sync Status** information provided in the right list pane for each specification.


**Table 6-7 Specification Synchronization Status Information in the List Pane**

List Pane Column Heading	Target Site Specification Information
<b>ID</b>	Unique number for the synchronization.
<b>Completed</b>	Date and time synchronization completed. If a synchronization has not completed, its progress is shown by a graphical bar.
<b>Bytes Queued</b>	Total number of data bytes to be synchronized.
<b>Bytes Sent</b>	Number of data bytes actually sent.
<b>Files Queued</b>	Number of file contained in the specification.
<b>Files Sent</b>	Number of files copied from the source to the target system.
<b>Remaining</b>	Approximate time remaining to complete the synchronization.

To clean up the RepliStor client window, you can delete the **Sync Status** icons after all synchronizations are complete. Right-click the **Sync Status** icon and select **Delete All**. All entries for each completed synchronization are deleted. If all synchronizations are complete, the **Sync Status** icon is also deleted.

You also can delete completed synchronization entries from the right pane. To do so, select one or more completed entries, then right-click and select **Delete** or **Delete All**. Entries cannot be deleted for synchronizations that are not complete.

If a synchronization is stopped or cannot occur because of an error, and you have scheduled synchronization with **Retry Sync For**, then it is *pending*. A pending icon appears directly beneath the specification, as shown in Figure 6-12. For more information about pending synchronization, refer to *Reissuing Pending Synchronizations* on page 6-34.



**Figure 6-12 Pending Synchronizations**



## Checking File Update Status

If a file is blocked and cannot be updated, a **Blocked Files** icon appears below the specification for the file, as shown in Figure 6-13. If the attached site is a target for one or more of the specifications, any blocked files on the target appear directly below the name of the attached site at the top of the tree.



Figure 6-13 Checking File Update Status

## Checking Source Status from the Target System

You can use the RepliStor client on the target system to view information about the target's source sites, specifications, and aliases, if defined. You also can view which source systems are set to forward IP addresses and the addresses, as shown in Figure 6-14.




Figure 6-14 Checking Source Status from the Target

## Correcting Site Problems

This section provides instructions for:

- ◆ *Troubleshooting Site Problems Using the Monitor Bar* on page 6-20
- ◆ *Correcting Blocked Sites and Communication Errors* on page 6-21
- ◆ *Correcting Blocked Files* on page 6-22

## Troubleshooting Site Problems Using the Monitor Bar

When a traffic light in the monitor bar displays a red or yellow light, you can use the associated monitor button to find the problem. Table 6-8 explains how to troubleshoot site problems using the monitor bar.

This troubleshooting section applies only to the sites you have attached to by selecting **Attach** from the **Functions** menu. If you click **View**, you can see all the attached sites at the bottom of the menu.

**Table 6-8 Troubleshooting Using the Monitor Bar**

Button	Troubleshooting Action
<b>Sites</b>	When the light is red, click <b>Sites</b> to find the next site that is blocked or has a pending synchronization. When the light is yellow, click <b>Sites</b> to find the next site with a pending synchronization. The <b>Target Sites</b> entry in the hierarchical tree is highlighted, and so is the name of the offending site in the detailed list in the right pane.
<b>Files</b>	Click <b>Files</b> to find the next blocked file, either at the current site or the next site. For example: <ol style="list-style-type: none"> <li>Click <b>Files</b>. The <b>Blocked Files</b> header for the site appears.</li> <li>Click <b>Files</b> again. The <b>Blocked Files</b> header under the specification for the first blocked file is highlighted.</li> <li>Click <b>Files</b> again to see the specification for the next blocked file. When no blocked files remain, the workspace for the next attached site with a blocked file appears with the <b>Blocked Files</b> header for the site highlighted.</li> </ol>
<b>Sync</b>	When the light is red, click <b>Sync</b> to find the next pending or active synchronization, either at the current site or the next site. Either the pending or seesaw icon directly below the specification is highlighted. When the light is yellow, click <b>Sync</b> to find the next active synchronization, either at the current site or the next site. The seesaw icon directly below the specification is highlighted.
<b>Mirroring</b>	Click <b>Mirroring</b> to find the next site for which mirroring is disabled. The site name is highlighted.
<b>Forwarding</b>	Click <b>Forwarding</b> to find the next site for which forwarding is disabled.
<b>Disk Space</b>	When the light is red, click <b>Disk Space</b> to find the next site with remaining disk space less than the configured Stop Size or Warning Size. When the light is yellow, click <b>Disk Space</b> to find the next site with remaining disk space less than the configured Warning Size.
<b>Messages</b>	When the light is red, click <b>Messages</b> to find the next unread severe or warning message, either at the current site or the next site. When the light is yellow, click <b>Messages</b> to find the next warning message, either at the current site or the next site. The newest messages are listed at the top of the message log. When you click <b>Messages</b> , the message log is searched from bottom to top (that is, starting with the oldest message).

## Correcting Blocked Sites and Communication Errors

A blocked site displays a **B** status code next to the name of the site in the right pane of the workspace. This means the RepliStor server has lost its connection to the target server. A status code of **C** indicates a communication error, which means the RepliStor client has lost communication with that server.

### To correct a blocked site:

If the monitor bar indicates a site is blocked, correct the problem as follows:

1. In the RepliStor client window, click **Target Sites** in the directory tree.
2. To immediately force RepliStor software to reconnect, right-click the name of the blocked target site and select **Unblock**.

RepliStor software periodically tries to reconnect with a blocked site. If there are multiple NIC interfaces to the target site, each will be used. When the connection is established successfully, all updates stored for the target machine are sent. If RepliStor software successfully unblocks the site, the **B** status code disappears. If **B** still appears, then the site remains blocked.

In some cases, the source connection to the target machine is lost for an indefinite period. Rather than allow file changes to accumulate on the source machine, it may be better to synchronize all specifications going to that site. A synchronization to a blocked specification automatically defers until the site becomes unblocked. In the meantime, the site is treated as if it were disabled and file changes are not queued to be sent to the site.

You also can use the RepliStor `blockedsites` command to display the number of blocked sites in the environment. In addition, you can specify an e-mail address to notify when a site becomes blocked through the RepliStor **Options** dialog box (on the **Log** tab), as shown in Figure 6-15 on page 6-22.

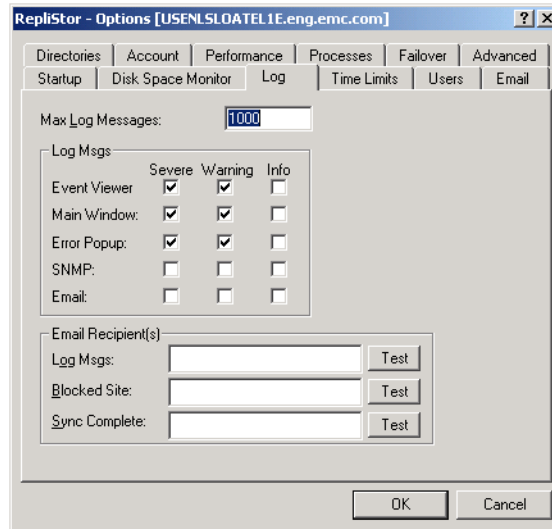


Figure 6-15 Options — Log Tab

## Correcting Blocked Files

When a target system cannot apply updates to its local copy of a file being replicated, the file is blocked. A file may be blocked for the following reasons:

- ◆ It is being used by another application or accessed by a user.
- ◆ The disk is full and the update requested to the file is attempting to increase the file's size.

When a file is blocked, RepliStor software generates a log message with enough detail to determine the cause of the block.

RepliStor software automatically checks at one-minute intervals (configurable on the **Time Limits** tab of the **Options** dialog box) whether updates can be applied, and if so, unblocks the files and applies the updates.

You also can unblock a file by right-clicking the file and selecting **Unblock** from the menu. If you select **Giveup**, RepliStor software discards all updates for the blocked file. This unblocks the file, but because updates are lost, the specification is no longer synchronized.

In addition, you can use the `numblockedfiles` command to display the total number of blocked files or from a specific site. Use the `unblockfiles` command to attempt to immediately correct all blocked files without waiting for the one-minute interval. For information on these commands, refer to Chapter 7, *Recovering Data*.

---

## Targeting a Remote Share Hosted by a NAS Device

With RepliStor software, you can target a remote share hosted by a Network Attached Storage (NAS) device. Targeting a NAS device is identical to normal RepliStor operation, except that the UNC path to the share needs to be registered with RepliStor software. Once registered, it can be used as any local drive.

---

### Registering a UNC Path

Before targeting a remote share hosted by a NAS device, you must register the UNC path.

**To register a UNC path:**

1. In the RepliStor client window, select the specification for which you want to register a UNC path.
2. Select **Modify Specification** from the **Maintenance** menu.

The **Specifications** dialog box opens, as shown in Figure 6-16 on page 6-24.

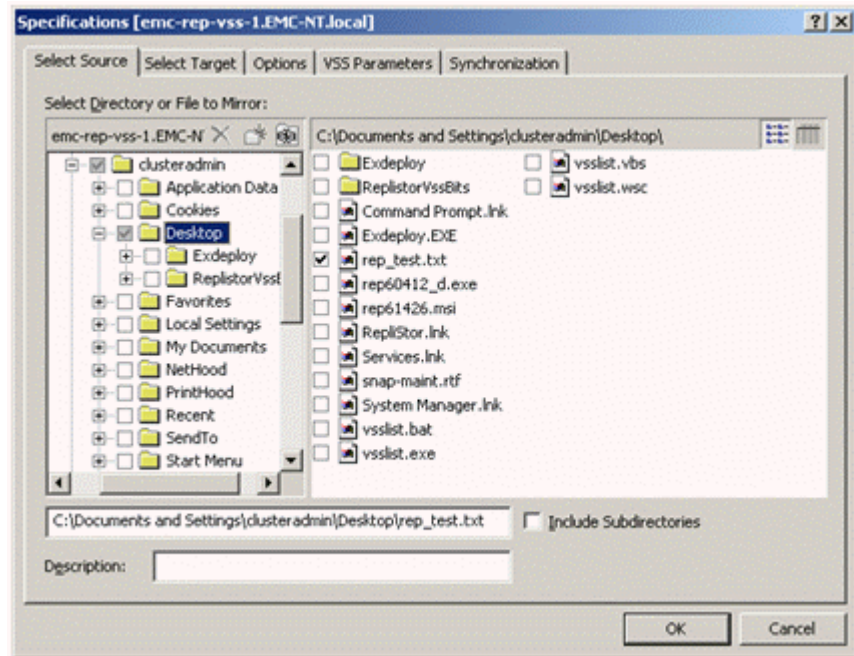
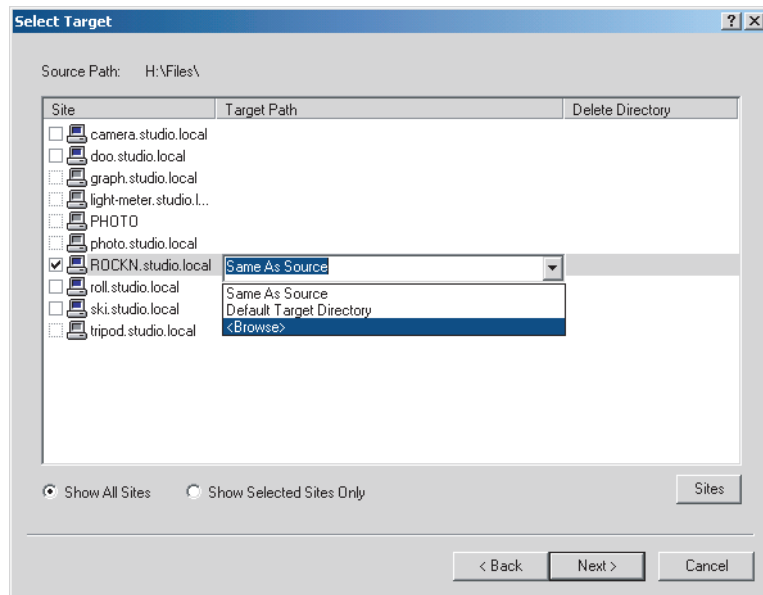



Figure 6-16 Specifications Dialog Box

3. Click the **Select Target** tab.
4. Click the **Target Path** of the target site.  
A drop-down list appears to the right of the target name.
5. From the drop-down list, select **Browse**.

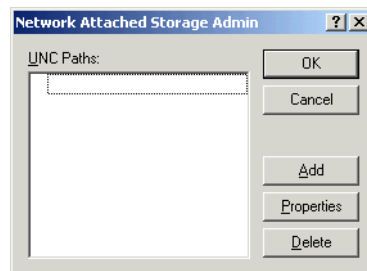
The **Select Target** dialog box appears (Figure 6-17).



**Figure 6-17 Select Target Dialog Box**

6. Click the  button.

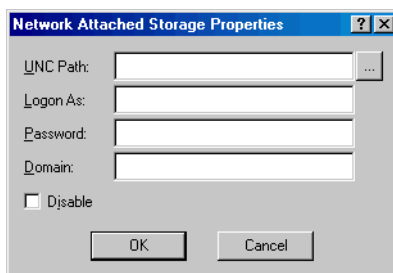
The **Network Attached Storage Admin** dialog box opens, as shown in Figure 6-18. The **UNC Paths** window in the **Network Attached Storage Admin** dialog box lists all currently registered UNC paths.



**Figure 6-18 Network Attached Storage Admin Dialog Box**

7. To add a path to the box, click **Add**.

The **Network Attached Storage Properties** dialog box opens, as shown in Figure 6-19.



**Figure 6-19 Network Attached Storage Properties Dialog Box**

8. In the **UNC Path** text box, enter the UNC path to the NAS share, using the following format: `\\computer\share`. Or, click **Browse** on the right to locate the NAS share.
9. In the **Logon As** text box, enter the logon ID used to access the NAS share.
10. In the **Password** text box, enter the password used to access the NAS share.
11. In the **Domain** text box, enter the domain associated with the NAS share.
12. (Optional) Click the **Disable** checkbox to register and add the UNC path to the **UNC Paths** box, but disable it.

A UNC path can be enabled later by clicking the checkbox next to the path in the **UNC Paths** box in the **Network Attached Storage Admin** dialog box. For more information, refer to *Enabling or Disabling a UNC Path*.

---

The **Protect Target Files** option in the **Options** dialog box cannot be enabled if the target specifies a UNC path. (Refer to Figure 2-16 on page 2-22.)

---

13. Click **OK**. RepliStor software attaches to all enabled UNC paths. Enabled paths appear as target *drives* on the **Select Target** tab and can be used as any other drive.

---

Each UNC path must be a path to a NAS device. If it is not, an error message appears and that path is disabled.

---



---

## Enabling or Disabling a UNC Path

If you did not enable a UNC path during the UNC path registration process, you must enable it before you can use it. Likewise, if you registered and enabled a UNC path that you no longer need, you can disable it. If you may need a UNC path in the future, consider disabling it instead of deleting it. Once you delete a UNC path, you must reregister it if you want to use it again.

### To enable or disable a UNC path:

1. Select the specification associated with the UNC path you want to enable or disable.
2. Select **Modify Specification** from the **Maintenance** menu.

The **Specifications** dialog box opens, as shown in Figure 6-16 on page 6-24.

3. Click the **Select Target** tab.
4. Click the **Target Path** of the target site.

A drop-down list appears to the right of the target name.

5. From the drop-down list, select **Browse**.

The **Select Target Path** dialog box appears.

6. Click the  button.

The **Network Attached Storage Admin** dialog box opens, as shown in Figure 6-18 on page 6-25.

7. The **UNC Paths** window lists all currently registered UNC paths. A check mark next to a path indicates it is enabled. Click the checkbox next to any path to enable or disable it. (If you are enabling a path, a check mark appears. If you are disabling a path, the check mark disappears.)

---

The **Protect Target Files** option in the **Options** dialog box cannot be enabled if the target specifies a UNC path.

---

8. Click **OK**.

---

## Modifying the Properties of a UNC Path

### To modify the properties of a UNC path:

1. Select the specification associated with the UNC path you want to modify.
2. Select **Modify Specification** from the **Maintenance** menu.

The **Specifications** dialog box opens, as shown in Figure 6-16 on page 6-24.

3. Click the **Select Target** tab.
4. Click the **Target Path** of the target site.

A drop-down list appears to the right of the target name.

5. From the drop-down list, select **Browse**.

The **Select Target Path** dialog box appears.

6. Click the  button.

The **Network Attached Storage Admin** dialog box opens, as shown in Figure 6-18 on page 6-25.

7. Select the UNC path to modify and click **Properties**.

The **Network Attached Storage Properties** dialog box opens, as shown in Figure 6-19 on page 6-26.

8. Modify the attributes in the dialog box as desired, then click **OK**.

---

## Deleting a UNC Path

---


If you may need a UNC path in the future, consider disabling it instead of deleting it. Once you delete a UNC path, you must reregister it if you want to use it again. For instructions on how to disable a UNC path, refer to *Enabling or Disabling a UNC Path* on page 6-27.

---

### To delete a UNC path:

1. Select the specification associated with the UNC path you want to delete.
2. Select **Modify Specification** from the **Maintenance** menu.

The **Specifications** dialog box opens, as shown in Figure 6-16 on page 6-24.

3. Click the **Select Target** tab.
4. Click the **Target Path** of the target site.  
A drop-down list appears to the right of the target name.
5. From the drop-down list, select **Browse**.  
The **Select Target Path** dialog box appears.
6. Click the  button.  
The **Network Attached Storage Admin** dialog box opens, as shown in Figure 6-18 on page 6-25.
7. Select the UNC path to delete and click **Delete**.  
The path is removed from the **UNC Paths** box.

## Administering Specifications

This section provides instructions for working with specifications. For information on creating specifications, refer to *Creating Specifications* on page 2-15.

### Modifying Specifications

To modify specifications for data that users have changed, you must resynchronize the specifications immediately after modification. Do not modify specifications waiting for blocked files or sites to become unblocked.

#### To modify a specification:

1. In the RepliStor client window, select the specification to modify:
  - For file or share specifications, right-click the specification to modify under **Target Sites**, as shown in Figure 6-20.

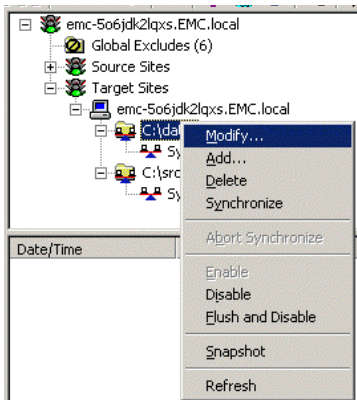


Figure 6-20 Modifying Specifications

- For **Global Exclude** specifications, click **Global Exclude** under the attached site, then select the specification to modify.
2. Right-click the specification and select **Modify**. The **Specifications** dialog box opens, as shown in Figure 6-16 on page 6-24.
3. In the **Specification** dialog box, modify the specification as needed, and then click **OK**.

4. When prompted for confirmation, do one of the following:
  - To add a new specification without changing the original specification, click **Add**.
  - To save the modifications to the specification, click **OK**.

---

## Deleting Specifications

You can delete a specification when you no longer need to replicate the files, directories, or shares it specifies.

### To delete a specification:

1. In the RepliStor client window, select the specification you want to delete:
  - For a file or share specification, click the specification to delete under **Target Sites**.
  - For a **Global Exclude** specification, click **Global Exclude** under the attached site, and then select the specification.
2. Press **DELETE**.
3. When prompted for confirmation, click **OK**.

---

## Enabling and Disabling Specifications

When you disable a specification, you stop it from synchronizing, replicating, or forwarding.

To disable an existing specification, right-click it in the RepliStor client window, and then select **Disable**. Once the specification is disabled, the specification is dimmed, the **Site** and **Target Sites** traffic lights turn yellow, and the **Sites** traffic light in the monitor bar turns yellow.

To re-enable a disabled specification, right-click it in the RepliStor client window, and then select **Enable**.

## Synchronizing Specifications

Before RepliStor software can replicate data successfully, an exact copy of the source system's data must exist on the target system. To make sure this data is exactly the same, you must synchronize it each time a source or target system fails. You also may need to resynchronize the specifications after turning them off and then back on.

You can synchronize all files or only those files that do not match between the source and target system.

**Important:** Make sure the source files are valid before you synchronize files after a system failure.

For information on how to synchronize specifications, refer to *Synchronizing Specifications* on page 2-39.

## Full, Incremental, and Partial Synchronization

The following sections describe the full, incremental, and partial synchronization options, which are set in the **Sync Options** dialog box.

To get to the **Sync Options** dialog box:

- ◆ In the RepliStor client window, right-click the specification you want to synchronize and select **Synchronize**.
- ◆ The **Sync Options** dialog box opens, as shown in Figure 6-21.




Figure 6-21 Sync Options Dialog Box

### Full Synchronization

In a full synchronization, all files are copied from the source to the target system. All aspects of each file are transferred including data streams, attributes (such as date and file size), and permissions.

For full synchronization, clear the **Incremental** option.

### Incremental Synchronization

If you select the Incremental option only, with no additional options, each file is scanned on both source and target systems. If the target file is not significantly different from that on the source, the source forwards an incremental update to the target. The update includes alternate streams, attributes, and permissions. If the file is significantly different, the entire file is forwarded from the source to the target, similar to a full synchronization.

If you select **Incremental** with the **Attribute Compare Only** option, each file on the source and target is scanned for differences in the last modified date and file size. If these attributes match, the file is skipped. If they do not match, the entire file is forwarded from the source to the target.

If you select **Incremental** with the **Log Differences** option, any differences found are forwarded from the source to the target and a message is sent to the message log. You can use this option to verify that mirroring is set up and working correctly.

The **Check Only** option can be used only with **Log Differences**. With this option, differences found are not forwarded, but are logged only in the message log.

In general, if the communications link is slow, select the **Incremental** synchronization option, without any additional options.

If the communications link is fast, the overhead of scanning the files on both source and target may negate any advantage gained by lowering network traffic during a synchronization. In this case, it may be faster to perform a full synchronization.

### Partial Synchronization

In a partial synchronization, a specific subdirectory or file is copied from the source to the target system. Click **Sub-Path** to specify the subdirectory or file.

A partial synchronization is useful when, for example, an error occurs during a sync operation and a single file or group of files does not get copied successfully. Rather than having to do a complete

resync of all files, which can be time consuming, the sync operation can be limited to copying just the missing files.

### Reissuing Pending Synchronizations

If a specification is configured for synchronization at a set time in the **Retry Sync For** option, and the synchronization cannot occur (because the source files are being accessed by users or applications), it is called a *pending synchronization*. The **Retry Sync For** option in the **Sync Options** dialog box specifies how long RepliStor software should attempt to synchronize the specification before abandoning it.

When a synchronization is pending, RepliStor software displays a pending icon under the specification. Select the icon to view detailed information on the status of the synchronization.

If you are replicating databases that hold files open continuously, such as SQL Server or Exchange databases, this is typically not a problem. To correctly mirror open files, you must synchronize the specification when it is created. This enables RepliStor software to manage the open files correctly.


Pending synchronizations have the following options:

- ◆ Reissue the synchronization of an individual file within the specification.
- ◆ Reissue the synchronization of all files within the specification.
- ◆ Give up the synchronization of an individual file within the specification.
- ◆ Give up the synchronization of all files within the specification.



## Reissuing or Giving Up a Pending Synchronization

To reissue or give up a pending synchronization:

1. In the directory tree, click the **Pending Synchronization** icon  (Figure 6-22). RepliStor software displays the status of all the files whose synchronization is pending for the specification.




Figure 6-22 Viewing Pending Synchronizations

Table 6-9 describes the meaning of the various status messages.

Table 6-9 Pending Synchronization Messages

Message	Description
Status	<ul style="list-style-type: none"> <li>• If status is <i>Pending</i>, RepliStor software is still attempting to synchronize the file.</li> <li>• If status is <i>Give Up</i>, RepliStor software has stopped attempting to synchronize the file because the <b>Give Up Time</b> has elapsed.</li> </ul>
Give Up Time	Shows the time after which RepliStor software stops attempting to synchronize the file.
Last Try Time	Shows the time when RepliStor software last tried to synchronize the file.
From File	Shows the open source file that RepliStor software is attempting to synchronize.
To Site	Shows the system to which RepliStor software is attempting to synchronize.

2. To reissue or give up *one or more* pending synchronization files within a specification:
  - a. In the directory tree, click the **Pending Synchronization** icon.
  - b. Select one or more pending synchronization files in the right-hand pane.

- c. Right-click the files.
  - To change the **Give Up Time** for the selected pending synchronization files, click **Reissue** to open the **Reissue Sync** dialog box, change the **Retry Sync Time**, and then click **OK**.
  - To remove the pending synchronization files for the selected pending synchronization files, click **Giveup**.
3. To reissue or give up *all* of the pending synchronization files within a specification, right-click the **Pending Synchronization** icon in the directory tree and then do one of the following:
  - To change the **Give Up Time** for all of the pending synchronization files for the selected specification, click **Reissue All** to open the **Reissue Sync** dialog box, change the **Retry Sync Time**, and then click **OK**.
  - To remove the pending synchronization files for the selected synchronization, click **Giveup All**.

---

## Managing Specifications for Microsoft Cluster Server

To replicate data from a Microsoft Cluster Server, you must create specifications to mirror data from the shared disk drives so that RepliStor software can continuously replicate data out of the Microsoft Cluster Server environment.

After creating specifications, you must create a RepliStor cluster resource and then associate it with the cluster group whose data you want to replicate. By associating the specification with a cluster group, the specification is managed by Microsoft Cluster Server and moved node-to-node within the resource group.

### Creating File/Directory Specifications

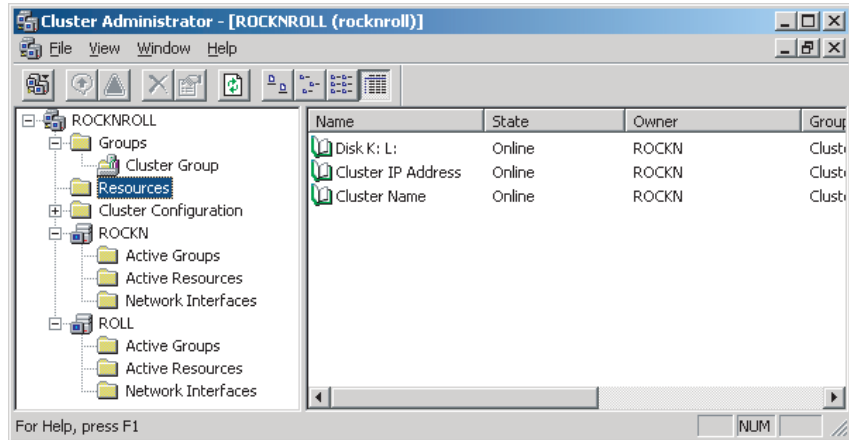
You can add a **File/Directory** specification as a RepliStor resource in Cluster Administrator by creating a **File/Directory** specification that mirrors data from a shared disk drive to a machine outside the Microsoft Cluster Server environment (follow the instructions in *File/Directory Specifications* on page 2-18).

### Creating a RepliStor Cluster Resource

A cluster resource allows Microsoft Cluster Server to manage the specifications you create. After you create the cluster resource, you must associate it with the cluster group whose data you want to replicate.

**To create a RepliStor cluster resource:**

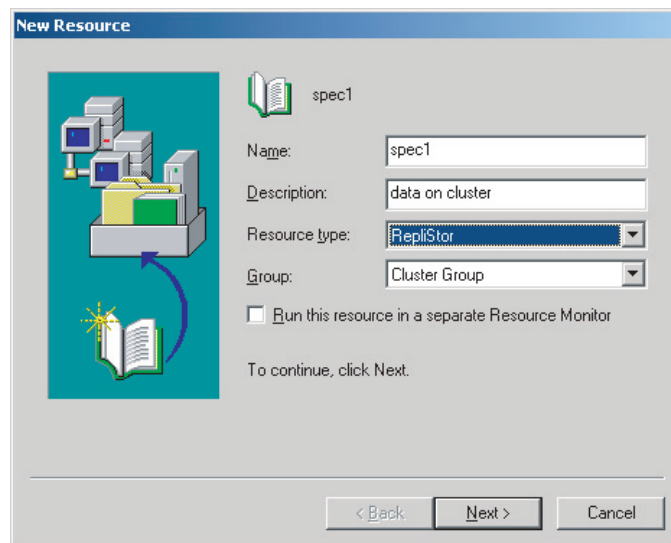
1. Start the Cluster Administrator.
2. In the left pane of the Cluster Administrator, select **Resources**, as shown in Figure 6-23.



**Figure 6-23 Cluster Administrator**

3. Select **File, New, Resource**.

The **New Resource** dialog box opens, as shown in Figure 6-24.



**Figure 6-24 New Resource Dialog Box**

4. Supply the following information:

- **Name:** Enter the resource's name. For example, RepliStor Spec1.
- **Description:** Enter a description.
- **Resource type:** Select RepliStor.
- **Group:** Select the cluster group you want to associate with RepliStor software.

If you have installed the cluster-based application that you want associated with RepliStor software, select that application's cluster group in the **Group** list.

If you have not yet installed the cluster-based application you want associated with RepliStor software, or if you are not sure which cluster-based application you want associated with RepliStor software, select **Cluster Group** in the **Group** list.

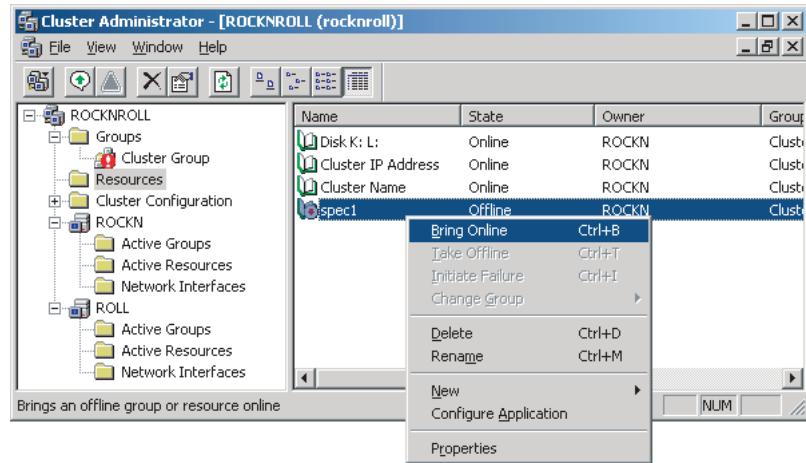
5. Click **Next**. The **Possible Owners** dialog box opens.
6. Click **Next** to allow both nodes in the cluster to potentially take ownership of the RepliStor resource.

The **Dependencies** dialog box opens. At a minimum, make sure the specification is dependent on the shared disk it is replicating.

7. Click **Next**. The **MSClust Parameters** dialog box is displayed.
8. Select the appropriate specification, and click **Finish**.

You can select only one specification per RepliStor resource. If there are several specifications that need to be associated with the **Resource Group**, you can create additional RepliStor resources and add them to the same **Cluster Group**.

9. After the resource is created, it must be brought online. Do this by right-clicking the resource and selecting **Bring Online** (Figure 6-25).



**Figure 6-25** Bringing the Resource Online

## Modifying Specifications Associated with a RepliStor Cluster Resource

You can run the Cluster Administrator on any node in the cluster when creating a RepliStor resource and associating it with a specification, or when editing an existing specification. Previously, you had to run the Cluster Administrator on the node that owned the resource group of which the resource was a member.

You also can modify a specification associated with a RepliStor resource. When you modify a specification, RepliStor automatically updates it in the cluster database.

## Changing the RepliStor Cluster Resource

To change the cluster resource group with which RepliStor software is associated, drag the RepliStor resource icon from its current resource group to its new resource group.

## Replicating Data with Microsoft Cluster Server

Before replicating data from the cluster to a target stand-alone server or a target Microsoft Cluster Server, do the following:

- ◆ Install RepliStor software in a Microsoft Cluster Server environment.

- ◆ Create the RepliStor resource and associate that resource with a cluster group.

When you are replicating out of a cluster environment, keep the following rules in mind:

- ◆ You can only replicate data on a shared disk drive.
- ◆ You can only replicate data located on the same shared disk drives associated with the same cluster group that you associated with RepliStor software.

For example, if you configured SQL Server to be associated with two shared disk drives, drive x and drive z, and you associated RepliStor software with the SQL Server cluster group, you can only replicate data from drive x and drive z.

- ◆ You cannot replicate data from the shared disk drives to one of the local drives in the cluster.

## Turning Processes On and Off

This section provides instructions for turning the data mirroring, forwarding, and updating processes on and off.

- ◆ *Stopping and Starting File Mirroring and Forwarding*, which follows
- ◆ *Pausing File Updates* on page 6-43
- ◆ *Stopping and Starting File Attribute Updates* on page 6-44

### Stopping and Starting File Mirroring and Forwarding

By default, when the RepliStor server starts, mirroring and forwarding are on. You can change the default settings for mirroring and forwarding at startup.

You might want to temporarily stop mirroring and forwarding for the following reasons:

- ◆ To perform routine maintenance, such as tape backup, on the source system
- ◆ To eliminate network traffic from RepliStor software for a period of time

**Important:** If you turn mirroring on or off, you must resynchronize specifications, because any changes in data that occur when mirroring is off are not captured.

Resynchronizing ensures the data on the target system is an exact copy of the source.

### Turning Mirroring On and Off

#### To turn mirroring on and off:

1. Make sure you are attached to the source system on which you want to stop or start mirroring.
2. In the RepliStor client window, select **Mirroring** from the **Functions** menu.
3. If you are turning mirroring off, click **Yes** when the confirmation message appears.

When mirroring is on, a check mark appears next to the **Mirroring** option in the **Functions** menu, the mirroring button is highlighted in the toolbar, and the monitor bar light is green.

When mirroring is off, no check mark appears next to the **Mirroring** option in the **Functions** menu, the mirroring button is dimmed in the toolbar, and the monitor bar light is red.

**Important:** When mirroring is off, any file updates that may occur on the source server are not captured. This means after re-enabling mirroring, you should re-synchronize all specifications.

### Turning Global Forwarding On and Off

Global forwarding affects all target sites. When you turn *global forwarding* on and off, you also turn the system's heartbeat signal on and off.

#### To turn global forwarding on and off:

1. Make sure you are attached to the source system where you want to stop or start forwarding.
2. In the RepliStor client window, select **Forwarding** from the **Functions** menu.
3. When prompted for confirmation, click **OK**.

If you are turning forwarding off, you can click **Schedule** in the confirmation window instead to specify a date, time, interval, and frequency at which to stop forwarding. For information, refer to *Configuring the Updating and Forwarding Options* on page 4-5.

When forwarding is off, no check mark appears next to the **Forwarding** option in the **Functions** menu, the **Forwarding** icon is dimmed in the toolbar, and the monitor bar light is red.

When forwarding is on, a check mark appears next to the **Forwarding** option in the **Functions** menu, the **Forwarding** button is highlighted in the toolbar, and the monitor bar light is green.

### Turning Site Forwarding On and Off

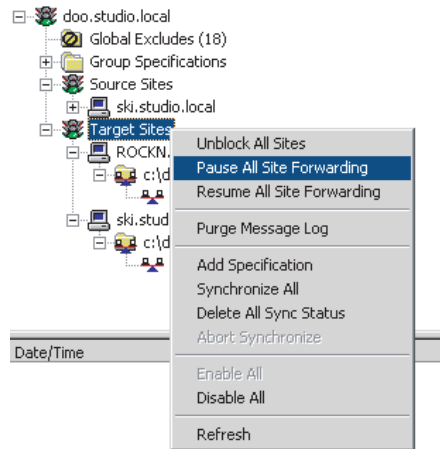
Use site forwarding to control forwarding of changes to a specific site. When you turn site forwarding off, you do not stop the system's heartbeat signal.

#### To turn site forwarding on and off:

1. Make sure you are attached to the source system on which you want to stop or start forwarding.



2. Do *one* of the following in the RepliStor client window:
  - To stop forwarding on all sites, right-click **Target Sites** in the directory tree and select **Pause All Site Forwarding**, as shown in Figure 6-26. To restart forwarding on all sites, select **Resume All Site Forwarding**.



**Figure 6-26 Site Forwarding Option**

- To stop forwarding on an individual site, right-click the name of the site, and select **Pause Site Forwarding**. A check mark appears next to the menu option. To restart forwarding on the site, select **Pause Site Forwarding** (the check mark disappears).
3. When prompted for confirmation, click **Yes**.

## Pausing File Updates

At times you may want to temporarily stop RepliStor software from applying updates to the files on a target machine, for example, when performing a tape backup of the target machine.

The **Pause Updates** option on the **Functions** menu allows you to stop the update process until the tape backup is complete. To pause file updates, you must select **Pause Updates** from the **Functions** menu while you are attached to the target system. During the pause, the target continues to receive and store updates, but they are not applied to the target files until you unselect **Pause Updates** on the source system.

When you perform a tape backup of the target system, make sure the RepliStor data directory is not included. Since RepliStor software continues to receive and store updates, it will continue to write to the data directory from the source system.

If you back up files regularly, you may want to schedule Pause Updates to occur at the same time every day or within some other time period you specify. The updates are applied to the target files when the scheduled pause is completed. For information, refer to *Configuring the Updating and Forwarding Options* on page 4-5.

## Pausing File Updates

### To pause file updates:

1. Make sure you are attached to the target system.
2. In the RepliStor client window, select **Pause Updates** from the **Functions** menu.
3. When prompted for confirmation, click **OK**.

When you pause updates, you can click **Schedule** in the confirmation windows to specify a date, time, interval, and frequency at which to pause updates.

When you pause updates, a check mark appears next to the **Pause Updates** option in the **Functions** menu. When you resume updates, the check mark is removed.

---

## Stopping and Starting File Attribute Updates

RepliStor software allows you to turn on and off the process of updating file attributes on the target system. File attributes are flags that define file types as hidden, read only, archive, or system files. These flags also control the updating of the file times (for create, modify, and access).

### To stop and start file attribute updates:

1. Make sure you are attached to the target system.
2. In the RepliStor client window, select **Attribute Update** from the **Functions** menu.
3. When prompted for confirmation, click **OK**.

When you start updates to file attributes, a check mark appears next to the **Attribute Update** option in the **Functions** menu. When you stop updates, the check mark is removed.

## Working with the Message Log

During normal operation and when errors occur, RepliStor software sends messages with the following information:

- ◆ The date, time, site name, and process affected
- ◆ The severity of the message, which can be:
  - *Severe*: RepliStor software encountered an error that terminated the affected process
  - *Warning*: RepliStor software encountered an error you should address
  - *Information*: RepliStor software is providing you with information on its normal operation
- ◆ A textual description
- ◆ The text of the Windows error, if any

This text usually provides the reason for the error. Most errors are related to Windows system operations, such as writing to a file. For information about Windows messages, refer to the Windows documentation.

RepliStor messages can appear in the following places, depending on the product configuration:

- ◆ Windows Event Log
- ◆ SNMP Event Log
- ◆ Client window
- ◆ Pop-up window
- ◆ E-mail

This section provides instructions for:

- ◆ *Viewing the Message Log*, on page 6-46
- ◆ *Marking Messages as Read* on page 6-47
- ◆ *Purging the Message Log* on page 6-48

## Viewing the Message Log

You can view all messages RepliStor software generates about any RepliStor site and select the severity level of the messages you want to see. Table 6-10 shows the severity levels.

**Table 6-10 Message Log Severity Levels**

Icon	Meaning
Question Mark	RepliStor software is unclear about the types of messages in the message log, because you are not attached to the site or the message log is empty.
Small Blue "i"	There are informational messages in the message log.
Yellow Exclamation Point	There are warning messages in the message log.
Stop Sign	There are severe messages in the message log.

## Viewing the Message Log

To view the message log:

1. Attach to the site on which you want to view the message log.

By default, the message pane in the RepliStor client window displays the message log for the site. If you have changed the default workspace settings to not view the message pane, click **Messages** in the directory tree.

2. To select the severity levels of the messages that appear:
  - a. Select **Messages** from the **Status** menu.

A check mark appears next to the severity levels that currently appear in the client window.

- b. Select the message types you want to see.

3. To view detailed information on any message in the client window:
  - a. Double-click the message in the message pane. A **Detail** window opens containing the details of the selected message.
  - b. Click **Prev** in the **Detail** window to view details on the previous message in the message log.
  - c. Click **Next** to see details on the next message in the message log.

- d. Click **OK** to close the **Details** window and return to the client window.
- e. Click **Copy** to place the detailed message on the Windows Clipboard.

Messages appear in bold in the messages list until you view the details. Once viewed, the message no longer appears in bold.

## Marking Messages as Read

To avoid having to open each message in the log, you can mark them as read:

1. In the RepliStor client window, click the **Messages** icon if the message pane is not open.
2. Click one or more of the messages in bold in the message pane, as shown in Figure 6-27.
3. Right-click the messages and select **Mark as Read** or **Mark All As Read**.

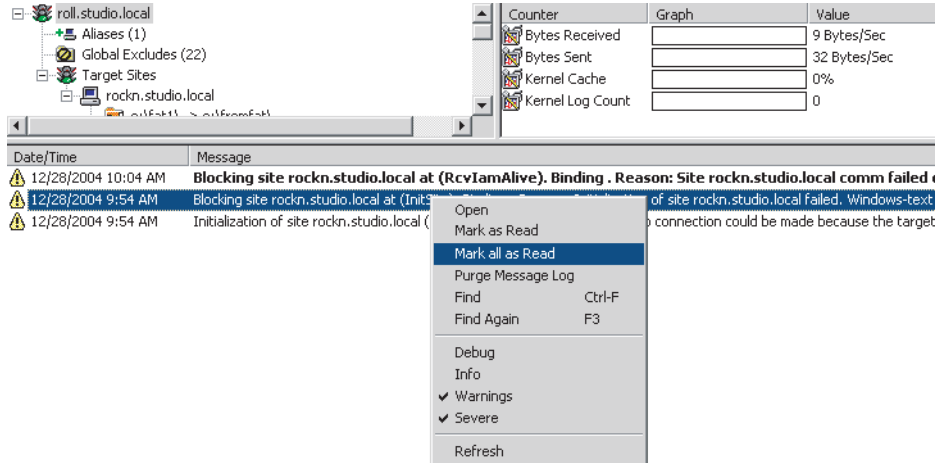


Figure 6-27 Marking Messages as Read

---

## Purging the Message Log

To purge messages from the log once you have viewed them:

1. In the RepliStor client window, select **Status, Messages, Purge Message Log**.
2. When prompted for confirmation, click **Yes**.

---

## Closing Application Files

RepliStor software opens a file on the target system when it replicates a specification. RepliStor software opens files exclusively on the target and essentially keeps all open files on the source also open on the target. RepliStor software closes all open files when the connection between the source and target systems is broken or when you stop RepliStor software on the source machine.

You can also force RepliStor software to close open application files. You may need to do this to perform routine maintenance on a system or to enable RepliStor software to synchronize specifications.

Using this function to close files on the target does not guarantee they will stay closed. A file update on the source reopens the file on the target.

### To close application files:

1. Make sure you are attached to the RepliStor site on which you want to close application files.
2. In the RepliStor client window, select **Close App Files** from the **Functions** menu.
3. When prompted for confirmation, click **Yes**.

## Using Tape Backups with RepliStor Software

Before performing a tape backup on the RepliStor system, do the following:

1. If you are performing a tape backup on a server that is a target system in a RepliStor configuration, make sure **Pause Updates** is selected in the **Functions** menu.

This function can be run from the command prompt. Thus, if the software you are using to perform the tape backup can execute a `Before Backup` and an `After Backup` command, you may want to set up this command.

2. If you are performing a tape backup on a source server in a RepliStor configuration, you may want to exclude the attribute updates that may be caused by the backup application (for example, the setting of the archive bit). Use the **Processes** tab in the **Options** dialog box to exclude attribute updates (refer to *Using the Processes Tab in the Options Dialog Box* on page 6-7).

**Example:** When performing a tape backup on any server in a RepliStor configuration, you must exclude the RepliStor data directory. Successfully backing up this directory and then restoring it may cause RepliStor to corrupt the data (on the target) it is replicating.

## Reporting/Simulating Bandwidth Between Source and Target

### Overview

The simulation/reporting feature is a menu item you can select under **Maintenance** in the RepliStor console (Figure 6-28). This feature allows you to simulate and report data bandwidth between the source and target systems. As shown in Figure 6-28, you can select a site in the left site pane, and then select **Maintenance**, **Simulation/Reporting** to monitor or simulate the traffic between the selected site and the target site. When you do this, the **Bandwidth Simulation/Reporting** dialog box appears (Figure 6-29 on page 6-51).

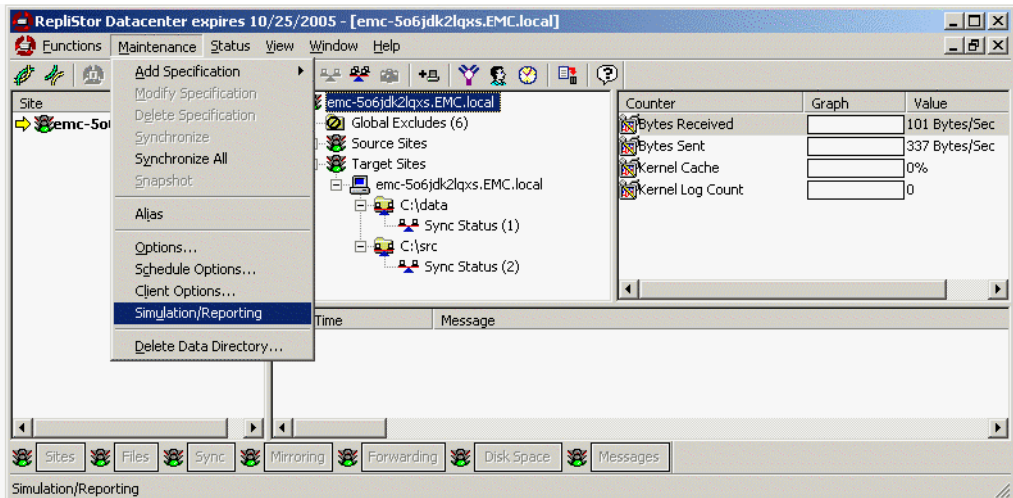
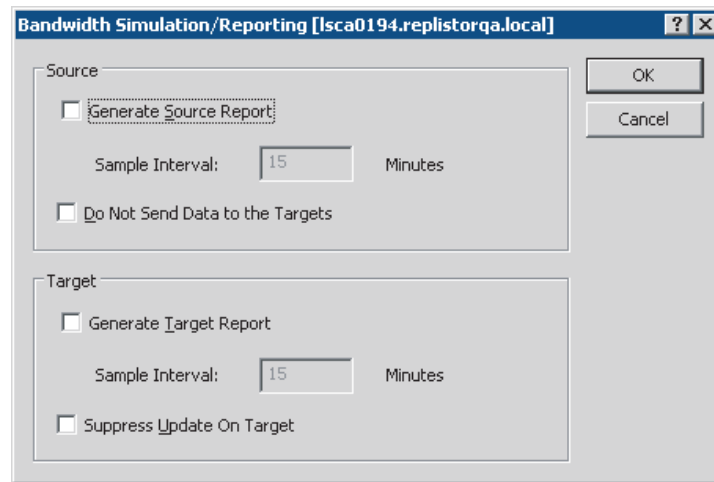


Figure 6-28 Simulation/Reporting Feature in Maintenance Menu





**Figure 6-29 Bandwidth Simulation/Reporting Dialog Box**

As shown in Figure 6-29, you can configure simulation and/or reporting on either the source or the target. For example, you can report on the source system and suppress updates on the target system, and so on. It is important to note that the local RepliStor server can be either a source or a target.

## Reporting

You can specify the sample interval in minutes at which you want to generate the source or target reports (for example, every 5 minutes, every 45 minutes, and so on). You can set the sample interval to generate source or target reports between 1 and 60 minutes. A smaller sample interval increases the accuracy of the bandwidth statistics between the source and the target at the expense of additional processing and larger file sizes.

The source and target reports are generated as text files in a comma-separated variable (csv) file format. Each row in the report is for a particular time period. The reports are stored in the Data directory.

An example of a source report is shown in Figure 6-30 on page 6-52. The column headings displayed in the source report are described in Table 6-11 on page 6-52. The source report file name is of the form: `SourceReport.year.month.day.hour.minute.csv`.

A target report example is shown in Figure 6-31 on page 6-53 and the target report column headings are described in Table 6-12 on page 6-54. The target report file name is of the form:  
 TargetReport.year.month.day.hour.minute.csv.

Time	Bytes Sent	Bandwidth Used, Bytes Sent	Kernel Cache Used	Kernel Cache High Water Mark	Kernel Cache Overflow Count	Kernel Cache Log Count	Kernel Cache Log Count High Water Mark	Number Open Files	rockn.studio.local-Replication Latency
5/17/2005 17:33	123583538	19460695	0	4	0	0	0	298	0
5/17/2005 17:38	140367733	21213232	0	4	0	0	0	455	0
5/17/2005 17:43	136496443	21328115	0	4	0	0	0	405	0
5/17/2005 17:48	123936886	19621642	0	4	0	0	0	379	0
5/17/2005 17:53	134563348	20634448	0	4	0	0	0	416	0
5/17/2005 17:58	134677879	20810912	0	4	0	0	0	520	0
5/17/2005 18:03	131217452	20464782	0	4	0	0	0	619	796
5/17/2005 18:08	123769775	19551314	0	4	0	0	0	330	0
5/17/2005 18:13	134535863	20508324	0	4	0	0	0	436	15
5/17/2005 18:18	142511724	22255441	0	4	0	0	22255441	558	15
5/17/2005 18:23	123858998	19637073	0	4	0	0	0	340	0
5/17/2005 18:28	139735546	21353282	0	4	0	0	0	491	0

Figure 6-30 Source Report Example

Table 6-11 Source Report Column Headings

Source Report Parameter	Description
Time	The date and time at which the source report was generated (in the example shown above in Figure 6-30, source reports were generated every five minutes).
Bytes Sent	The number of application data bytes sent from the source site to the target site since the last time the source report was generated. This number is cumulative in between intervals.
Bandwidth Used, Bytes Sent	The number of application data bytes sent from the source site to the target site, after any compression, since the last sample time. This number is cumulative in between intervals and represents the actual amount of data sent on the network.
Kernel Cache Used	The percent of the kernel cache used by the source site in sending data to the target. The kernel cache is a fixed amount of shared memory used for queuing data going from the source to the target.
Kernel Cache High Water Mark	The highest percent the kernel cache has used since the last time it overflowed and was reset.
Kernel Cache Overflow Count	The number of times the kernel cache has overflowed and began using OC\$ files.

Table 6-11 Source Report Column Headings (continued)

Source Report Parameter	Description
Kernel Cache Log Count	The number of OC\$nnnnn .rdf files that are currently in the data directory. Each log holds approximately 1 MB of data by default.
Kernel Cache Log Count High Water Mark	The highest number of OC\$nnnnn .rdf files that had existed.
Number Open Files	The number of open files that the RepliStor kernel driver is currently tracking.
<target> - Replication Latency	The number of milliseconds that it currently takes between the time the application performs the file operation and when that file operation is performed on the target. This column displays for every current target site.

Time	Bytes Received	Bandwidth Used, Bytes Received	photo.studio.local:0-Target Queue
5/17/2005 17:36	306668	306668	0
5/17/2005 17:41	306964	306964	0
5/17/2005 17:46	303100	303100	0
5/17/2005 17:51	303268	303268	0
5/17/2005 17:56	303212	303212	0
5/17/2005 18:01	303520	303520	0
5/17/2005 18:06	303100	303100	0
5/17/2005 18:11	303100	303100	0
5/17/2005 18:16	303212	303212	0
5/17/2005 18:21	303744	303744	0
5/17/2005 18:26	303044	303044	0
5/17/2005 18:31	303044	303044	0
5/17/2005 18:36	303800	303800	0
5/17/2005 18:41	303156	303156	0
5/17/2005 18:46	308756	308756	0

Figure 6-31 Target Report Example

Table 6-12 Target Report Column Headings

Target Report Parameter	Description
Time	The date and time at which the target report was generated (in the example shown above in Figure 6-31, target reports were generated every five minutes).
Bytes Received	The number of application data bytes received by the target site since the last sample interval. This number is cumulative in between intervals.
Bandwidth Used, Bytes Received	The number of application data bytes received by the target site, after any compression, since the last sample time. This number is cumulative in between intervals and represents the actual amount of data received on the network.
<source>:<connection ID> - Target Queue	Size of the target queue at the sample interval. This value, by default, should never exceed 100. If it is varying between 50 and 100, this indicates that there is a bottleneck on the target side in updating the disk, and it is slowing down the source system from sending data.

## Simulation

RepliStor allows you to simulate data replication where data is not actually sent to the target system. If used in conjunction with the reporting function, the simulation function is useful for determining the volume of data that is to be replicated as well as peak activity times, without loading the network or having a target system capable of handling the data from the source.

If you don't want to use real data for the source configuration, you can simulate the source bandwidth by selecting the **Do Not Send Data to the Targets** checkbox, as shown in Figure 6-29 on page 6-51. Selecting this option will prevent the source from sending data to the targets. If this option is selected, any data that is to be sent to the targets will be lost. If the **Do Not Send Data to the Targets** checkbox is not selected, the data is sent to the target (this is the default).

For the target configuration, selecting the **Suppress Update on Target** checkbox means that the local RepliStor server will not be updated with incoming data (that is, any data received from the source system will be ignored). All incoming data will be lost. This option is recommended if you want to run a simulation where the target system does not have the disk capacity to handle all of the source data. If the **Suppress Update on Target** checkbox is not selected, then data will be written to the local server.

When you select the desired options in the **Bandwidth Simulation/Reporting** dialog box, click **OK**.

If a report is in progress, you will see a **Reporting in Progress** message in the status bar in the lower right corner of the RepliStor client, as shown in Figure 6-32.

If a simulation is in progress (that is, if you selected either or both of the **Do Not Send Data to the Targets** or the **Suppress Update On Target** options), you will see a **Simulation in Progress** message in the lower right corner of the RepliStor client. This message displays to remind you that a simulation is running, as this is easy to forget once you set up the simulation.

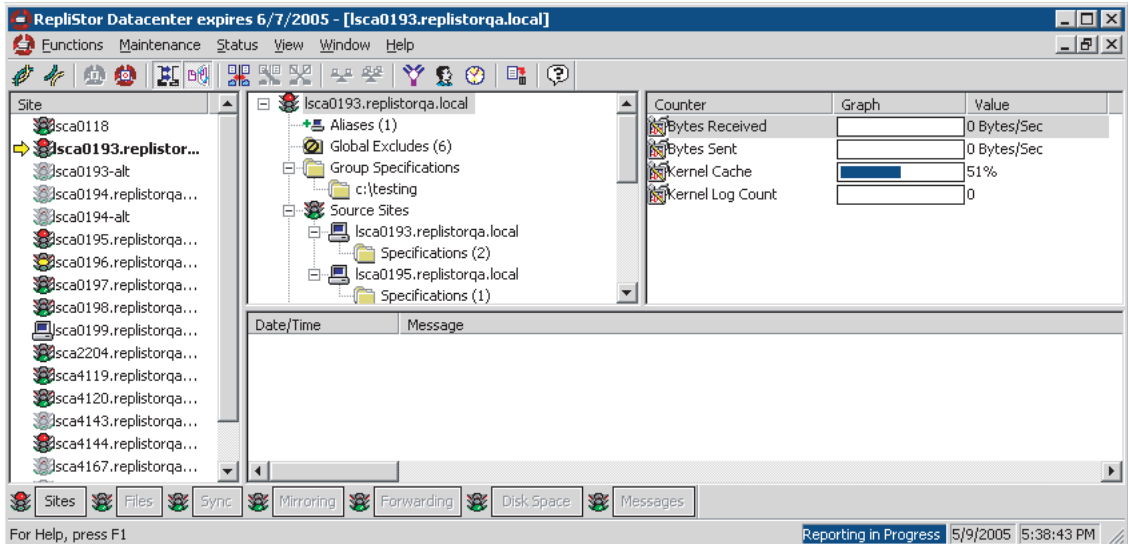


Figure 6-32 Reporting in Progress Message in Lower Right Status Bar



---

This chapter describes how to use RepliStor software to recover replicated data when a system fails and a failover occurs.

This chapter contains the following section:

- ◆ Alias Failover Recovery .....7-2

## Alias Failover Recovery

Aliases allow client connections to continue uninterrupted after a source failure and during recovery. They also help avoid computer name conflicts in the Active Directory that can occur when a failover takes place.

Assume you have the following configuration with aliases defined for use in an Alias failover:

- ◆ Source system A, with an alias C (and associated IP addresses) specified for the source. The alias is configured to use target system B. Specifications will be transferred to the target system during failover.
- ◆ Client connecting to the alias C.
- ◆ Specifications replicating data from source system A to target system B.

When a failover occurs, the alias C and its IP addresses are sent to target system B, and clients still connect to C (now on target B). All specifications replicated to target B are sent to the target and are modified so that they replicate back to source system A, which has a blocked status. When you first restore source system A, all specifications replicated to target B are disabled and alias C is not restored on the source because it was transferred to the target at failover. Clients are still connecting to alias C on target B.

At this point, you can either restore the original configuration or leave the alias on system B and set the alias to fail over to system A. To do the latter, refer to *Setting the Alias to Fail Over to the Original Source* on page 7-3.

---

### Restoring the Original Configuration

To restore system A as the source and return to the original configuration:

1. Make sure the synchronization of all data going from target B to original source A is complete.
2. On the target system, select **Aliases** in the middle pane.
3. Right-click the Alias name in the right pane, and then click **Remove**.



All aliases, along with associated IP addresses, that were sent to the target are automatically removed from the target system and added back to the source system. In addition, any specifications added to the target from the source are removed. Clients are still connected to alias C, which is now on the source system.

4. To re-enable specifications on the source, right-click the specification in the RepliStor client window on the source system, and then select **Enable**.
5. To synchronize specifications, right-click the specification, and then select **Synchronize**.

The specifications return to replicating from source system A to target system B.

---

### Setting the Alias to Fail Over to the Original Source

To leave the alias on system B and set the alias to fail over to system A (essentially switching server roles):

1. In the RepliStor client window on system B, select **Aliases** in the middle pane.
2. Right-click the Alias name in the right pane and click **Activate Aliases** if the Alias is not already active.
3. Right-click the Alias name in the right pane and select **Set Source**.

The alias you selected is now configured to fail over to the original source system A in case system B, the new source, fails.



---

This chapter describes RepliStor commands you can run from a command prompt.

This chapter contains the following sections:

- ◆ Using RepliStor Commands .....8-2
- ◆ Syntax Conventions .....8-3
- ◆ Command Descriptions .....8-4

## Using RepliStor Commands

To use a RepliStor command:

1. Open the **Command Prompt** window.
2. At the prompt, enter one of the following:
  - **rep\_srv cmd** *command* to run RepliStor commands on the local system
  - **rep\_srv cmd:***servername* *command* to run RepliStor commands on a remote system where:
    - *command* is the RepliStor command you want to run
    - *servername* is the name of the server you want to run the command on, either local or remote. If you do not specify a server name, the local server is used.

The command line interface (CLI) commands return either 0 for success or 1 for error. If it returns 1, it will output an error message that describes the error.

## Syntax Conventions

**Important:** ReplStor commands are not case-sensitive.

This chapter uses the conventions described in Table 8-1 to indicate proper command syntax.

**Table 8-1 Syntax Conventions**

Syntax	Convention	Example
Variable parameters	Italics	<code>spec <i>source_path</i> <i>target_site</i> <i>target_path</i></code> where you must supply values for <i>source_path</i> , <i>target_site</i> , and <i>target_path</i> .
Required parameters	Are not enclosed in brackets or parentheses.	<code>spec <i>source_path</i> <i>target_site</i> <i>target_path</i></code> where <i>source_path</i> , <i>target_site</i> , and <i>target_path</i> are required values.
Optional parameters	Enclosed in brackets.	<code>spec/dfs <i>target_site</i> [/delete]</code> where <i>target_site</i> is required, but <code>/delete</code> is optional.
Parameter options with one choice	Enclosed in parentheses and separated by a vertical bar.	<code>forwarding (on off)</code> where you must use either <code>on</code> or <code>off</code> .
Parameter options with multiple choices	Enclosed in parentheses and separated by a comma.	<code>modifyspec <i>source_path</i> <i>target_site</i> ((/enable   /disable),(/mirror   /nomirror))</code> where in addition to specifying the <i>source_path</i> and the <i>target_site</i> , you must specify the following: <code>/enable or /disable</code> and/or <code>/mirror or /nomirror</code>

## Command Descriptions

This section describes the RepliStor commands.

### aliasadd

The `aliasadd` command allows you to create an alias from the command line. The `aliasadd` command uses the following syntax:

```
rep_srv cmd aliasadd alias_name target_site
                {ipaddress_subnet} [optional parameters...]
```

Table 8-2 lists the parameters to use when creating an alias.

**Table 8-2 aliasadd Required Parameters**

Parameter	Description
<code>alias_name</code>	Name of the alias to create. A NetBIOS and DNS name are created from this name, so it must be unique on the network.
<code>target_site</code>	Failover target site. For a failover to occur, at least one specification must exist to replicate data to the target site.

Table 8-3 lists the optional parameters you can use when creating an alias.

**Table 8-3 aliasadd Optional Parameters**

Parameter	Description
<code>ipaddress_subnet</code>	List of IP addresses and subnets that are associated with this alias. When the alias is activated, these IP addresses are created.
<code>/manual</code>	Sets manual activation. If the source system fails, the alias is added to the source, but not automatically activated.
<code>/log</code>	Captures any output created during script execution and places it in the RepliStor message log for informational purposes.
<code>/disable   /enable</code>	Disables or enables the alias when it is created.
<code>/scriptaddbefore <i>script_name</i></code>	Script to run <i>before</i> adding an alias to the source or target system.
<code>/scriptaddafter <i>script_name</i></code>	Script to run <i>after</i> adding an alias to the source or target system.

Table 8-3 `aliasadd` Optional Parameters (continued)

Parameter	Description
<code>/scriptremovebefore script_name</code>	Script to run <i>before</i> removing an alias from the source or target system.
<code>/scriptremoveafter script_name</code>	Script to run <i>after</i> removing an alias from the source or target system.
<code>/addtodns</code>	Adds DNS entries for each IP address associated with the alias when the alias is activated. If the alias is deactivated, the DNS entries are removed. If there are no IP addresses associated with the alias, then a DNS entry is created that associates the alias with the local computer name.
<code>/dnsdomain domain_name</code>	Required parameter if the <code>addtodns</code> parameter is specified. Type the DNS domain. For example: <b>company.com</b> .
<code>/ttl number_of_seconds</code>	Time to live. Enter the number of seconds that the DNS entries are valid. A larger value caches DNS information on the clients for a longer period of time, so DNS traffic is reduced. A smaller value shortens the time it takes for all clients to recognize a change in a DNS entry. The default is 600 seconds.
<code>/servicetostop service_list</code>	List of services to stop when the alias is activated (and started on deactivation). Separate service names with a comma. If there are any spaces in the list of service names, enclose the list in quotes.
<code>/servicetostart service_list</code>	List of services to start when the alias is activated (and stopped on deactivation). Separate service names with a comma. If there are any spaces in the list of service names, enclose the list in quotes.

## aliasop

The `aliasop` command allows you to perform any alias function (for example, activate, fail, or set source). In most cases you must specify the `alias_name` and `target_site` that will perform the operation. Refer to Table 8-4 on page 8-6 for exceptions to this rule.

The `aliasop` command uses the following syntax:

```
rep_srv cmd aliasop alias_name target_site [optional parameters...]
```

Table 8-4 lists the optional parameters you can use with the `aliasop` command.

**Table 8-4 aliasop Optional Parameters**

Parameter	Description
<code>/activate   /deactivate</code>	Activates or deactivates the alias.
<code>/setsource</code>	Turns a target alias into a source alias after a failover. An alias is created on the target that fails over to the original source.
<code>/remove</code>	Removes the alias after a failover and reverts to the previous configuration. If this is a source alias, the alias is deleted. If this is a target alias (after a failover), then the alias is removed from the target and, if the source is available, instructs the source to activate the corresponding alias.
<code>/globalactivate</code>	Activates all enabled aliases. The <code>alias_name</code> and <code>target_site</code> variables are not used with this parameter. The <code>globalactivate</code> parameter is always disabled when RepliStor software starts.
<code>/globaldeactivate</code>	Deactivates all enabled aliases. The <code>alias_name</code> and <code>target_site</code> variables are not used with this parameter.
<code>/fail</code>	Performs a manual failover of all enabled aliases. The <code>alias_name</code> and <code>target_site</code> variables are not used with this parameter.

---

## blockedsites

The `blockedsites` command displays a list of currently blocked sites to standard output. A site is blocked when data cannot be sent from the source to the target because the target is down or the network is not functioning. To list currently blocked sites, enter:

```
rep_srv cmd blockedsites
```

---

## checkpoint

The `checkpoint` command can be used in a consolidated backup configuration to achieve consistent backups at the central site. The following is an example of how to use the Checkpoint function:

1. Quiesce the source application.
2. Run the RepliStor `Checkpoint` command. This will queue a sentinel token for the selected site. Specify a target script and `/PauseUpdates`.



3. Resume the source application.
4. When the queued token reaches the target system, `Updates` will be automatically paused so that any file operations that occurred after the checkpoint will not be applied. The target files will be frozen at the point the `Checkpoint` command occurred.
5. The script (specified in the `Checkpoint` command) is run. This should perform a shadow copy, backup, or whatever operation used to preserve the target files.
6. When finished, resume updates.

This command runs as a CLI function only. The syntax of the command is:

```
> rep_srv cmd CheckPoint <source path> <target site> [/PauseForwarding]
  [/PauseUpdates] [/SourceScript <source Script>]
  [/TargetScript <target Script>]
```

Table 8-5 lists the optional parameters you can use with the checkpoint command.

**Table 8-5 checkpoint Optional Parameters**

Parameter	Description
<code>/PauseForwarding</code>	Automatically pauses forwarding when the sentinel is encountered in the source, so that the target files are frozen at the time the sentinel was written.
<code>/PauseUpdates</code>	Automatically pauses updates on the target when the sentinel is received by the target system, so that the target files are frozen at the time the sentinel was written.
<code>/SourceScript &lt;script&gt;</code>	Runs the script when the source encounters the sentinel.
<code>/TargetScript &lt;script&gt;</code>	Runs the script when the target receives the sentinel.

## close

The `close` command closes application files on the local system or on a specified remote system. For more information about closing application files, refer to *Closing Application Files* on page 6-48.

To close the application files on the local system, enter:

```
rep_srv cmd close
```

To close the application files on remote system Server A, enter:

```
rep_srv cmd:servera close
```

---

## configexport/ configimport

The `ConfigExport` command saves all configuration information, and the `ConfigImport` command restores all configuration information. The syntax of this command is:

```
rep_srv cmd ConfigExport <file>  
rep_srv cmd ConfigImport <file>
```

where:

<file> is a file where the configuration is saved (`ConfigExport`) or restored from (`ConfigImport`).

---

## deletedatadir

The `deletedatadir` command deletes the RepliStor data directory on the local system or on a specified remote system.

To delete the data directory on the local system, enter:

```
rep_srv cmd deletedatadir
```

To delete the data directory on remote system Server A, enter:

```
rep_srv cmd:servera deletedatadir
```

---

## forwarding

The `forwarding` command stops or starts forwarding, and stops or starts the system's heartbeat signal on the local system or on a specified remote system. For information about starting and stopping forwarding from the RepliStor client window, refer to *Turning Processes On and Off* on page 6-41.

To stop forwarding on the local system, enter:

```
rep_srv cmd forwarding off
```

To start forwarding on remote system Server A, enter:

```
rep_srv cmd:servera forwarding on
```

---

## idle

The `idle` command determines if there are file operations queued to be sent to any target by checking the status of the kernel cache and the presence of OC\$n files.

- ◆ If the kernel cache is empty and there are no OC\$n files, this command returns `IDLE`.
- ◆ If either the kernel cache is not empty or at least one OC\$n file exists, it returns `BUSY`. Note that in each case this command returns an error level of 1.

---

## mirroring

The `mirroring` command stops or starts mirroring on the local system or on a specified remote system. For information about starting and stopping mirroring from the RepliStor client window, refer to *Turning Processes On and Off* on page 6-41.

To stop mirroring on the local system, enter:

```
rep_srv cmd mirroring off
```

To start mirroring on remote system Server A, enter:

```
rep_srv cmd:servera mirroring on
```

---

## modifyspec

The `modifyspec` command disables and enables an existing specification, and turns mirroring on or off for the specification.

When you use the `modifyspec` command, in addition to the `source_path` and the `target_site`. Optional parameters include `/disable` or `/enable`, `/mirror` or `/nomirror`, and `/resource` or `/noresource`.

The `modifyspec` command uses the following syntax:

```
rep_srv cmd modifyspec source_path target_site (/disable | /enable), (/mirror | /nomirror), (/resource|/noresource)
```

Table 8-6 lists the parameters to be used with the `modifyspec` command. To modify any aspect of a specification, use the `spec` command. For information about using the `spec` command, refer to *spec* on page 8-15.

**Table 8-6 modifyspec Required Parameters**

Parameter	Description
<code>source_path</code>	The source path of the existing specification.
<code>target_site</code>	The target site to which the existing specification is mirroring data.

Table 8-7 lists the optional parameters that can be used with the `modifyspec` command.

**Table 8-7 modifyspec Optional Parameters**

Parameter	Descriptions
<code>/disable   /enable</code>	Disables or enables the existing specification. You can use one of these parameters.
<code>/mirror   /nomirror</code>	Turns mirroring on or off in an existing specification. You can use one of these parameters.
<code>/resource   /noresource</code>	Sets the flag on or off that indicates the specification is a resource managed by Legato® Cluster or Microsoft Cluster Server. You can use one of these parameters.
<code>/flushdisable [wait &lt;wait-millisecond&gt;</code>	Disables the specification when all current operations for that specification have been forwarded to the target. If you specify a wait time, the command will not return until either the specification has become disabled or the wait time expires, whichever is sooner. If the wait time expires without the specification becoming disabled, it will return a time-out error.

To disable an existing specification that mirrored all data in `C:\DATA` on the local system to Server B, enter:

```
rep_srv cmd modifyspec c:\data serverb /disable
```

To start mirroring for an existing specification that exists on remote system Server A, which mirrors all of the data in `C:\DATA` to Server B, enter:

```
rep_srv cmd:servera modifyspec c:\data serverb /mirror
```

---

## numblockedfiles

The `numblockedfiles` command displays the number of blocked files from a specific site or from all sites to standard output. The `numblockedfiles` command uses the following syntax:

```
rep_srv cmd numblockedfiles sitename
```

To display the number of blocked files on a specific Server B, enter:

```
rep_srv cmd numblockedfiles serverb
```

---

## pause

The `pause` command pauses or resumes the updating of data forwarded from a source machine to the local system or a remote system. For information about pausing and unpausing file updates on the target system from the RepliStor client window, refer to *Turning Processes On and Off* on page 6-41.

To pause data updating on the local system, enter:

```
rep_srv cmd pause on
```

To resume data updating on a remote Server B, enter:

```
rep_srv cmd:serverb pause off
```

---

## pauseattributes

The `pauseattributes` command pauses or resumes the updating of file attributes (for example, read only) from the source machine to the local system or a specified remote system. For information about pausing file updates from the RepliStor client window, refer to *Pausing File Updates* on page 6-43.

To pause the updating of file attributes on the local system, enter:

```
rep_srv cmd pauseattributes on
```

To resume the updating of file attributes on a remote Server A, enter:

```
rep_srv cmd:servera pauseattributes off
```

---

## report/simulate

The `report` command performs the reporting/simulation function, which allows you to report and simulate data bandwidth between the source and target systems.

---

For more information on the reporting and simulation feature, refer to *Reporting/Simulating Bandwidth Between Source and Target* on page 6-50.

The `report` command uses the following syntax:

```
rep_srv cmd report [/SourceReport] [/SourceInterval n]
[/noSourceReport] [/noSendData] [/SendData]
[/TargetReport] [/TargetInterval n]
[/noTargetReport] [/noUpdateData] [/UpdateData]
```

Table 8-8 lists the optional parameters that can be used with the `report` command.

**Table 8-8** `report` Optional Parameters

Parameter	Descriptions
<code>/SourceReport</code>	Generate a source report.
<code>/SourceInterval n</code>	Set interval to take source reports (minutes).
<code>/noSourceReport</code>	Do not generate a source report.
<code>/noSendData</code>	Suppress sending of data to the target system.
<code>/SendData</code>	Enable sending of data to the target system.
<code>/TargetReport</code>	Generate a target report.
<code>/TargetInterval n</code>	Set interval to generate target reports (minutes).
<code>/noTargetReport</code>	Do not generate a target report.
<code>/noUpdateData</code>	Suppress updating of data on the target system.
<code>/UpdateData</code>	Enable updating of data on the target system.

To generate source and target reports every 30 minutes and to run a simulation where no real data is sent from the source to the target system and the target system is not updated with incoming data, enter:

```
rep_srv cmd report /SourceReport /30 /noSendData
/TargetReport /30 /noUpdateData
```

To generate source reports every 10 minutes without running a simulation, where real data is sent to the target and the local RepliStor server is updated with incoming data, enter:

```
rep_srv cmd report /SourceReport /10 /SendData
/noTargetReport /UpdateData
```

---

## siteforwarding

The `siteforwarding` command starts or stops forwarding from the local system or a remote system to the specified site. If no site is specified, forwarding is stopped for all target systems. For information about turning **Site Forwarding** on and off from the RepliStor client window, refer to *Turning Processes On and Off* on page 6-41.

The `siteforwarding` command uses the following syntax:

```
rep_srv cmd siteforwarding off | on [site]
```

This command does not start or stop the system's heartbeat signal, so a failover does not occur if **Site Forwarding** is turned off.

To turn site forwarding to Server B off, enter:

```
rep_srv cmd siteforwarding off serverb
```

To turn site forwarding on for all target sites of remote system Server A, enter:

```
rep_srv cmd:servera siteforwarding on
```

---

## ShadowCopy

The `ShadowCopy` command allows you to take a shadow copy (a consistent point-in-time copy of data) of the specifications listed. Multiple specification descriptions may be listed. If the `/scheduled` option is specified, a shadow copy schedule is created. If the `/noscheduled` option is specified, any existing schedule is deleted.

The `VSSBackupType` option specifies the type of shadow copy (backup type) that is performed on the source. This option is application/writer-dependent; in other words, it depends on what you want your specific application to do. For example, if you are taking a shadow copy of an Exchange database, and select **Full**, the log files will be truncated. If you are taking an Exchange shadow copy and selected **Copy** for the backup type, the log files would not be truncated. The backup types are the following:

- ◆ **Full** — All files (regardless of whether they have been marked as backed up), are saved. Each file's backup history is updated to reflect that it was backed up.
- ◆ **Copy** — Files are copied to a backup medium regardless of the state of each file's backup history, and the backup history is not updated.

The `ShadowCopy` command uses the following syntax:

```
rep_srv cmd ShadowCopy {<Group Spec Description>} [/scheduled <sched time>]  
[/every <num> <time units>] [/noscheduled] [/VSSBackupType <backup type>]
```

To take a shadow copy of group specification `D:\Exchange database` every day at 6:00 AM where the log files are truncated, enter:

```
rep_srv cmd ShadowCopy d:\Exchange database /scheduled
6:00 AM /every 1 day /Full
```

---

## ShadowUtil

The `ShadowUtil` command allows you to delete, mount, or unmount a shadow copy or shadow copy set. In addition, you may revert the target volume to the specified shadow copy or display (to standard output) the properties of a particular shadow copy.

The `ShadowUtil` command uses the following syntax:

```
rep_srv cmd ShadowUtil [/Set] [Delete <GUID>] [Mount <GUID> <path>]
[Unmount <GUID>] [Revert <GUID> [/ForceUnmount]] [Properties [<GUID>]]
```

`/Set` — if specified, the GUID must refer to a shadow copy set, and the operation applies to the entire set.

Only one of the following commands can be on a single command line:

`Delete` — Deletes the specified shadow copy (or set).

`Mount` — Mounts the specified shadow copy (or set) to the specified path. The path must refer to an empty directory. When mounting a shadow copy set, it will create subdirectories for each mounted shadow copy.



`Unmount` — Unmounts the shadow copy(or set).

`Revert` — Reverts the target volume to the specified Shadow Copy. If `/ForceUnmount` is specified, RepliStor will unmount and revert the volume even if it is currently in use by other applications.

---

VSS-compliant Service Pack 1 is required for the revert functionality. Microsoft recommends a hotfix, KB891957, that alleviates problems arising from depleted paged pool memory. Refer to KB 891957 on the Microsoft website, <http://support.microsoft.com>

---

`Properties` —Displays (to standard output) the properties of a particular shadow copy. If a GUID is not supplied, this command will display the properties of all shadow copy.

The shadow copy properties displayed include:

- ◆ Creation time
- ◆ Shadow copy set GUID
- ◆ Shadow copy GUID
- ◆ Source server of the shadow copy
- ◆ Source volume
- ◆ Target volume
- ◆ Mount path, if any
- ◆ CLARiiON shadow copy information (if CLARiiON provider)
- ◆ XML recovery document path

---

## **spec**

The `spec` command allows you to create, edit, or delete a mirror file specification. To synchronize the specification once it is created, use the `sync` command. For information about creating a mirror file specification from the RepliStor client window, refer to *File/Directory Specifications* on page 2-18.

The `spec` command uses the following syntax:

```
rep_srv cmd spec source_path target_site target_path [optional parameters...]
```

Table 8-9 lists the required parameters to use when creating a specification with the `spec` command.

**Table 8-9 spec Required Parameters**

Parameter	Description
<code>source_path</code>	The path to be mirrored from the source to the target.
<code>target_site</code>	The target site to which the files in the <code>source_path</code> will be mirrored from the source site.
<code>target_path</code>	The path on the target specifying the directory or file where the files from the <code>source_path</code> will be mirrored.

Table 8-10 lists the optional parameters available for adding features to a specification.

**Table 8-10 spec Optional Parameters**

Parameter	Description
<code>/enable</code>	Enables a disabled specification.
<code>/exclude spec</code>	Specifies the files or file extensions not to be mirrored to the target system. Multiple entries are separated by semicolons (;). For example: <code>/exclude readme.txt; *.tmp</code>
<code>/subtree</code>	Mirrors all subdirectories in the source path in the mirror file specification.
<code>/nosubtree</code>	Indicates that subdirectories in the source path in the mirror file specification should not be mirrored.
<code>/permissions_sids</code>	Passes file security to the remote file as a security ID. This option is enabled by default.
<code>/permissions_none</code>	File security information is not passed to the mirrored file.
<code>/permissions_names</code>	Passes file security to the remote file as an account name.
<code>/compatible</code>	Mirrors files properly from an NTFS drive to a FAT drive.
<code>/reflectprotection</code>	Does not allow changes from the source to be mirrored back from the target.
<code>/noreflectprotection</code>	Turns off <b>Reflect Protection</b> . This option is enabled by default.

Table 8-10 spec Optional Parameters (continued)

Parameter	Description
/mirror	Ensures the files, directories, and shares within the specification are mirrored to the target system after the specification has been synchronized. This option is enabled by default.
/nomirror	Does not enable the files within the specification to be mirrored from the source system to the target. The specification needs to be synchronized.
/copyonclose	Does not synchronize any of the files in the specification if the file is currently open.
/nocopyonclose	Turns off <b>Copy On Close</b> . This option is enabled by default.
/prop_attributes	Copies the attributes of each source file and directory to its corresponding mirrored file and directory on the target system. This option is enabled by default.
/noprop_attributes	Turns off <b>Propagate Attributes</b> .
/prop_fromshares	Mirrors the additions and deletions of shares to the path.
/noprop_fromshares	Turns off <b>Propagate From Shares</b> . This option is enabled by default.
/nodeleteext <i>ext</i>	One or more extensions that RepliStor software adds to a file on the target system to indicate that this file is a copy of a file that has been deleted from the source system. Multiple entries are separated by semicolons. For example: *.oc1;*.oc2
/deletedir <i>directory</i>	The directory on the target system used to store deleted files.
/delete	Deletes the specified mirror file specification.
/prop_hiddenshare	Propagates hidden shares within the specification. A hidden share is on a name ending in "\$".
/noprop_hiddenshare	Does not propagate hidden shares. This option is enabled by default.
/disable	The specification is created, but is disabled. No mirroring or synchronizing can occur while it is disabled.
/protect	Protects target files from modification. Only read-only access is allowed; read/write access causes an <b>Access Denied</b> error. This option is enabled by default.
/noprotect	Turns off target file protection.

Table 8-10 spec Optional Parameters (continued)

Parameter	Description
<code>/defaulttarget</code>	If specified, the <i>target_path</i> should not be specified. The target path used is defined on the <b>Directories</b> tab in the <b>Options</b> dialog box on the target system.
<code>/nodefaulttarget</code>	Specifies that no default target path should be defined on the <b>Directories</b> tab in the <b>Options</b> dialog box on the target system.
<code>/description <i>description</i></code>	Specifies a description for the specification.
<code>/deletedirage <i>number_of_days_until_delete</i></code>	Specifies the number of days before deleted files are actually deleted on the target. This option is used only if <code>/nodeleteext</code> is specified.
<code>/syncstartscript <i>script_file</i></code>	Specifies a script that will be run when a sync operation is started.
<code>/syncendscript <i>script_file</i></code>	Specifies a script that will be run when a sync operation completes or is aborted.
VSSBackupType	Specifies the type of shadow copy (backup type) that is performed on the source. This option is application/writer-dependent; in other words, it depends on what you want your specific application to do. The backup types are the following: <ul style="list-style-type: none"> <li>• <b>Full</b> — All files (regardless of whether they have been marked as backed up), are saved. Each file's backup history is updated to reflect that it was backed up.</li> <li>• <b>Copy</b> — Files are copied to a backup medium regardless of the state of each file's backup history, and the backup history is not updated.</li> </ul>

To create a mirror file specification mirroring all data in `C:\DATA` on the local system to `C:\DATA1` on Server B, not propagating permissions, and enabling **Copy On Close**, enter:

```
rep_srv cmd spec c:\data serverb c:\data1\ /permissions_none /copyonclose
```

To create the same specification from remote system Server A, enter:

```
rep_srv cmd:servera spec c:\data serverb c:\data1\ /permissions_none /copyonclose
```

---

## spec/ globalexclude

The `spec/globalexclude` command creates a **Global Exclude** specification. For information about creating a **Global Exclude** specification from the RepliStor client window, refer to *Global Exclude Specifications* on page 2-16.

The `spec/globalexclude` command uses the following syntax:

```
rep_srv cmd spec/globalexclude source_path [/subtree] [/delete]
```

Table 8-11 lists the parameters used with the `spec/globalexclude` command.

**Table 8-11 spec/globalexclude Parameters**

Parameter	Description
<code>source_path</code>	The source path to be excluded from mirroring. This parameter is required.
<code>/subtree</code>	Includes all subdirectories in the source path in the <b>Global Exclude</b> specification.
<code>/delete</code>	Deletes the specified <b>Global Exclude</b> specification.

To create a **Global Exclude** specification that excludes the `C:\WINNT` directory and all of its subdirectories, enter:

```
rep_srv cmd spec /globalexclude c:\winnt /subtree
```

To delete a **Global Exclude** specification on remote system Server A that excludes the `C:\WINNT` directory and all of its subdirectories, enter:

```
rep_srv cmd:servera spec /globalexclude c:\winnt /subtree /delete
```

## spec/registry

The `spec/registry` command creates a registry specification. For information about creating a registry specification from the RepliStor client window, refer to *Registry Specifications* on page 2-30.

The `spec/registry` command uses the following syntax:

```
rep_srv cmd spec/registry source_registry_path target_site target_registry_path
[optional parameters...]
```

Table 8-12 lists the parameters to be used with the `spec/registry` command.

**Table 8-12 spec/registry Required Parameters**

Parameter	Description
<code>source_registry_path</code>	The path to be mirrored from the source to the target.
<code>target_site</code>	The target site to mirror the files in the <code>source_registry_path</code> from the source registry.
<code>target_registry_path</code>	The path on the target specifying the registry to mirror the files from the <code>source_registry_path</code> .

Table 8-13 lists optional parameters that can be used with the `spec/registry` command.

**Table 8-13 spec/registry Optional Parameters**

Parameter	Description
<code>/enable</code>	Enables a disabled specification.
<code>/subtree</code>	Mirrors all subdirectories in the source registry path in the registry specification.
<code>/permissions_sids</code>	Passes file security to the target registry as a security ID. This option is enabled by default.
<code>/permissions_none</code>	File security information is not passed to the target.
<code>/mirror</code>	Ensures the registry keys within the specification are mirrored to the target system after the specification is synchronized. This option is enabled by default.
<code>/nomirror</code>	Does not allow the keys within the specification to be mirrored from the source system to the target. The specification needs to be synchronized.
<code>/delete</code>	Deletes the specified registry specification.
<code>/disable</code>	The specification is created, but is disabled. No mirroring or synchronizing can occur while it is disabled.
<code>/description <i>description</i></code>	Specifies a description for the specification.
<code>/syncstartscript <i>script_file</i></code>	Specifies a script to run when a sync operation is started.
<code>/syncendscript <i>script_file</i></code>	Specifies a script to run when a sync operation completes or is aborted.

## spec/shares

The `spec/shares` command creates a mirror shares specification. For information about creating mirror shares specifications from the RepliStor client window, refer to *Share Specifications* on page 2-28.

The `spec/shares` command uses the following syntax:

```
rep_srv cmd spec/shares target_site [optional parameters...]
```

Table 8-14 lists the parameters to use with the `spec/shares` command.

**Table 8-14 spec/shares Parameters**

Parameter	Description
<code>target_site</code>	The target site to mirror the shares from the source site. This parameter is required.
<code>/exclude share</code>	Lists the shares not to mirror to the target system. Separate multiple shares with semicolons (;).
<code>/prop_hiddenshare</code>	Mirrors all hidden shares in the specification.
<code>/noprop_hiddenshare</code>	Turns off <b>Propagate Hidden Shares</b> . This option is enabled by default.
<code>/permissions_sids</code>	Passes file security to the remote file as a security ID. In most cases, you should use this mode when mirroring files.
<code>/permissions_none</code>	Does not pass file security information to the mirrored file.
<code>/delete</code>	Deletes the specified mirror shares specification.

## specexport

The `specexport` command exports all specifications in the current configuration to a file. Once exported, the specifications can be restored using the `specimport` command. For information on using the `specimport` command, refer to *specimport* below.

The `specexport` command uses the following syntax:

```
rep_srv cmd specexport file_name
```

where `file_name` is the path and filename of the file to be created.

---

## specimport

The `specimport` command imports specifications from a file to the current configuration.

The `specimport` command uses the following syntax:

```
rep_srv cmd specimport file_name
```

where *file\_name* is the path and filename of the file containing the specifications to import.

---

The file containing the specifications must have been created using the `specexport` command.

---



---

## specstatus

The `specstatus` command displays the status of a specification to standard output. The `specstatus` command uses the following syntax:

```
rep_srv cmd specstatus source_path target_site
```

Table 8-15 lists the parameters to use with the `specstatus` command.

**Table 8-15 specstatus Parameters**

Parameter	Description
<code>source_path</code>	The source path of the specification.
<code>target_site</code>	The target site of the specification.

This command displays the following status information for the specification:

```
enable|disable mirror|nomirror [resource] (DMR)
```

where `disable mirror resource (DMR)` corresponds to numbers, as follows:

- ◆ D — 0 if enabled; 1 if disabled
- ◆ M — 0 if mirroring is on; 1 if mirroring is off
- ◆ R — 0 if the specification is not a cluster resource; 1 if it is a cluster resource



---

**start**

The `start` command starts the RepliStor server on either the local system or a specified remote system.

To turn on the RepliStor server on the local system, enter:

```
rep_srv cmd start
```

To turn on the RepliStor server on remote system Server A, enter:

```
rep_srv cmd:servera start
```

---

**stop**

The `stop` command stops the RepliStor server on either the local system or a specified remote system.

To turn off the RepliStor server on the local system, enter:

```
rep_srv cmd stop
```

To turn off the RepliStor server on remote system Server A, enter:

```
rep_srv cmd:servera stop
```

---

**sync**

The `sync` command synchronizes one or more specifications from either the local system or a specified remote system. You can also use this command to edit the current synchronization options. If you do not specify any parameters, all of the specifications are synchronized immediately.

The `sync` command uses the following syntax:

```
rep_srv cmd sync [optional parameters...]
```

If a synchronization is requested while the target system is blocked, the synchronization is deferred. While a specification is in a deferred state, it is disabled so that any changes to the file data are not applied. As soon as the target becomes available, the specification is enabled, the synchronization starts, and changes are applied.

If a `sync` command has been issued but not started when a site becomes blocked, the synchronization is terminated and the specification is placed in a deferred sync state.

For more information about synchronization, refer to *Synchronizing Specifications* on page 6-32. For detailed instructions on synchronizing, refer to page 2-39.

Table 8-16 lists the optional parameters available with the `sync` command.

**Table 8-16 sync Optional Parameters**

Parameter	Description
<code>to_site</code>	Synchronizes all specifications mirrored to the specified target site. If not specified, then all specifications are synchronized.
<code>/spec from_file</code>	Only those specifications with a matching source file are synchronized. Enter the entire source path of the specification you want to synchronize. For example: <code>/spec c:\sql\data\</code>
<code>/inc</code>	Performs an incremental synchronization on the selected specifications.
<code>/noinc</code>	Turns off incremental synchronization on the selected specifications.
<code>/deleteorphans</code>	Deletes any files on the target path not found on the source path when the file is synchronized.
<code>/nodeleteorphans</code>	Turns off delete orphans for the selected specification.
<code>/retry elapsed_time</code>	Designates the amount of time in which RepliStor software attempts to resynchronize the specification if the initial synchronization fails.
<code>/noretry</code>	Turns off any <b>Retry Sync</b> value associated with the specification.
<code>/noscheduled</code>	Turns off the schedule options for the selected specification. If the schedule options are turned off, the synchronization for the specification is also turned off (the specification is not synchronized).
<code>/every num time_units</code>	If you specify the <code>/schedule</code> parameter, you can use this parameter to specify how often the specification is to be resynchronized. <ul style="list-style-type: none"> <li>• <code>num</code> must be a number</li> <li>• <code>time_units</code> must be once, seconds, minutes, hours, days, weeks, or months</li> </ul>

Table 8-16 sync Optional Parameters (continued)

Parameter	Description
<code>/scheduled "date time"</code>	<p>Specifies the date and time to synchronize the specification. The <code>/scheduled</code> parameter uses the following syntax (with quotes):</p> <pre><code>/scheduled "year/month/day hour:minutes am/pm"</code></pre> <ul style="list-style-type: none"> <li>• If the date is not specified, the current day is used.</li> <li>• If the year is not specified, the current year is used.</li> <li>• If the month is not specified, the current month is used.</li> <li>• If the time is not specified, 12:00 a.m. is used.</li> <li>• If a.m. or p.m. is not specified, 24-hour format is used.</li> </ul> <p>For example, to set a parameter to begin on December 1, 2002 at 4:30 p.m., enter:</p> <pre><code>/scheduled "2002/12/1 4:30pm"</code></pre> <p>To set the parameter to begin on August 21 of the current year, enter:</p> <pre><code>/scheduled "8/21"</code></pre> <p>To set the parameter to begin at 10:00 a.m. today, enter:</p> <pre><code>/scheduled "10:00"</code></pre>
<code>/logdifferences</code>	Logs differences between source files and target files, but does not synchronize them.
<code>/checkonly</code>	Enables you to check whether the source and target files are synchronized without actually synchronizing them. You must specify <code>/logdifferences</code> to use this parameter.
<code>/sharesonly</code>	Only applies to a mirror file/directory specification with the <b>From Shares</b> option selected. Synchronizes only the associated shares, not the files or directories.
<code>/compareattributes</code>	Compares the time stamp and size of source and target files, and if they are different, copies the entire file from source to target.
<code>/subpath subpath</code>	Specifies a portion of the specification to sync. The <code>subpath</code> can specify a directory within the spec or a single file within the spec.
<code>/subpathsubtree</code>	If specified and if <code>/subpath</code> is specified, it will sync all subdirectories under <code>subpath</code> .

## throttle

The `throttle` command configures throttling on either the local system or the specified remote system. For more information about setting throttling from the client window, refer to *Configuring Throttling Options* on page 4-6. The `throttle` command uses the following syntax:

```
rep_srv cmd throttle [optional parameters...]
```

Table 8-17 on page 8-26 lists the optional parameters that can be used with the `throttle` command.

To set a throttling rate of 200 kb per second on the local system for all target systems from 8:00 a.m. to 6:30 p.m. to occur each day, enter:

```
rep_srv cmd throttle 200 /scheduled 8:00am /for 10:30 /every 1 days
```

**Table 8-17 throttle Optional Parameters**

Parameter	Description
<code>kbytes/sec</code>	The maximum data transfer rate from the source system to the target system. If the rate is 0 or left blank, throttling is disabled.
<code>/site site</code>	The site to which the throttling rate applies. If you do not specify a site, the throttling rate applies to all target sites.
<code>/for "elapsed_time"</code>	<p>If you specify the <code>/scheduled</code> parameter, you must also specify the length of time to enable throttling. The <code>/for</code> parameter uses the following syntax (with quotes):</p> <pre>"days hours:minutes"</pre> <p>For example, to set a parameter to be active for 4 days, 2 hours, and 23 minutes, enter:</p> <pre><b>/for "4 2:23"</b></pre> <p>To set a parameter to be active for 6 days, enter:</p> <pre><b>/for "6"</b></pre>
<code>/every num time_units</code>	<p>If you specify the <code>/scheduled</code> parameter, you must also specify how often to enable throttling.</p> <ul style="list-style-type: none"> <li><code>num</code> must be a number</li> <li><code>time_units</code> must be one of the following values: <code>once</code>, <code>seconds</code>, <code>minutes</code>, <code>hours</code>, <code>days</code>, <code>weeks</code>, <code>months</code></li> </ul>
<code>/scheduled "date time"</code>	<p>The date and time to apply the throttling rate. If you do not specify a schedule, the throttling rate is set immediately. The <code>/scheduled</code> parameter uses the following syntax (with quotes):</p> <pre><code>/scheduled "year/month/day hour:minutes am/pm"</code></pre> <ul style="list-style-type: none"> <li>If the date is not specified, the current day is used.</li> <li>If the year is not specified, the current year is used.</li> <li>If the month is not specified, the current month is used.</li> <li>If the time is not specified, 12:00 a.m. is used.</li> <li>If a.m. or p.m. is not specified, 24-hour format is used.</li> </ul> <p>For example, to set a parameter to begin on December 1, 2002 at 4:30 p.m., enter:</p> <pre><b>/scheduled "2002/12/1 4:30pm"</b></pre> <p>To set the parameter to begin on August 21st of the current year, enter:</p> <pre><b>/scheduled "8/21"</b></pre> <p>To set the parameter to begin at 10:00 a.m. today, enter:</p> <pre><b>/scheduled "10:00"</b></pre>

---

## **unblockfiles**

The `unblockfiles` command immediately attempts to unblock currently blocked files and apply pending file operations.

To display the number of blocked files from a specific site, enter:

```
rep_srv cmd unblockfiles
```



## RepliStor Components

This appendix describes the various RepliStor components including the file type, default installation directory, filename, and functions.

This appendix contains the following sections:

◆ Kernel Driver .....	A-2
◆ RepliStor Service .....	A-2
◆ RepliStor Control Service.....	A-3
◆ RepliStor Client .....	A-3
◆ Performance Monitor Support.....	A-3
◆ SNMP Support .....	A-4
◆ Microsoft Cluster Server Support.....	A-4
◆ Help.....	A-4
◆ Documentation .....	A-5

## Kernel Driver

<b>File Type:</b>	Kernel mode file system filter driver
<b>Default Install Directory:</b>	%SystemRoot%\system32\drivers
<b>File:</b>	replistor.sys
<b>Functions:</b>	<ul style="list-style-type: none"><li>◆ Monitors all file system activity and captures any file operations described in any specification. Operations are queued to the <code>KernelCache</code> or <code>OC\$n\$nnnnn</code> file.</li><li>◆ Enforces the <b>Protect Target Files</b> specification option on the target system.</li></ul>

## RepliStor Service

<b>File Type:</b>	User mode service
<b>Default Install Directory:</b>	%SystemDrive%\Program Files\Legato RepliStor
<b>File:</b>	rep_srv.exe
<b>Functions:</b>	<ul style="list-style-type: none"><li>◆ Detects new entries placed into the <code>KernelCache</code> or <code>OC\$n\$nnnnn</code> file (placed there by the kernel driver) and forwards them to the appropriate target system(s) (forwarding).</li><li>◆ On the target system, accepts file operations from the source system and applies them to the application files (updating).</li><li>◆ Serves as an interface to the RepliStor client.</li><li>◆ Performs scheduling functions.</li><li>◆ Performs replication of shares and registry keys.</li><li>◆ Handles command line functions.</li><li>◆ Queues and performs synchronizations.</li><li>◆ Handles failover functions.</li><li>◆ Maintains the message log.</li><li>◆ Handles routing of SMTP messages.</li></ul>



---

## RepliStor Control Service

<b>File Type:</b>	User mode service
<b>Default Install Directory:</b>	%SystemDrive%\Program Files\Legato RepliStor
<b>File:</b>	CtrlService.exe
<b>Functions:</b>	<ul style="list-style-type: none"><li>◆ Allows the RepliStor client to stop and start remote RepliStor services.</li><li>◆ Performs the <b>Get Sites</b> function if not installed in the Active Directory domain.</li></ul>

---

## RepliStor Client

<b>File Type:</b>	User mode executable
<b>Default Install Directory:</b>	%SystemDrive%\Program Files\Legato RepliStor
<b>File:</b>	rep_clt.exe
<b>Function:</b>	Administrative client for configuration and status display.

---

## Performance Monitor Support

<b>File Type:</b>	Performance Monitor extension DLL
<b>Default Install Directory:</b>	%SystemDrive%\Program Files\Legato RepliStor
<b>File:</b>	repperf.dll
<b>Functions:</b>	<ul style="list-style-type: none"><li>◆ Interfaces RepliStor software to the Windows Performance Monitor functions.</li><li>◆ Real time display of the current operation.</li><li>◆ Allows registering of performance alerts so problems can be automatically monitored and reported.</li></ul>

## SNMP Support

<b>File Type:</b>	SNMP extension DLL
<b>Default Install Directory:</b>	%SystemDrive%\Program Files\Legato RepliStor
<b>File:</b>	SNMPAgnt.dll
<b>Function:</b>	Allows message log entries to be directed as an SNMP trap.

## Microsoft Cluster Server Support

<b>File Type:</b>	SNMP extension DLL
<b>Default Install Directory:</b>	%SystemRoot%\system32
<b>File:</b>	repspec.dll, repspecex.dll, repspec.ini
<b>Functions:</b>	<ul style="list-style-type: none"><li>◆ repspec.dll: Implements the RepliStor resource.</li><li>◆ repspecex.dll: Implements the RepliStor resource administrative extension.</li></ul>

## Help

<b>File Type:</b>	Help files
<b>Default Install Directory:</b>	%SystemDrive%\Program Files\Legato RepliStor
<b>File:</b>	rep_clt.chm
<b>Function:</b>	Provides online and context-sensitive help.

## Documentation

<b>File Type:</b>	Adobe Acrobat PDF files
<b>Default Install Directory:</b>	%SystemDrive%\Program Files\Legato RepliStor \pdf
<b>Files:</b>	replistorag_6.1.pdf, replistorig_6.1.pdf,
<b>Function:</b>	Provides online product documentation: Requires Adobe Acrobat Reader.



**RepliStor Exchange  
Support**

---

---

This appendix provides information on installing the RepliStor Exchange Support utility. This utility should be installed if you are planning to take VSS-compliant shadow copies of Exchange Server 2003 databases.

- ◆ Overview ..... B-2
- ◆ Installing RepliStor Exchange Support..... B-3

## Overview

RepliStor Exchange Support allows you to set Exchange parameters to facilitate creating and managing Exchange Server 2003 shadow copy backups via the Volume Shadow Copy Service (VSS) framework provided by Microsoft. Shadow copies are consistent-point-in-time data copies stored on the target system. Shadow copies allow you to obtain a consistent, restartable copy of the Exchange database.

---

For information on VSS shadow copies, refer to *VSS-compliant Shadow Copies* on page 1-7.

---

The installation of RepliStor Exchange Support is required if you are going to take Exchange Server 2003 shadow copy backups. Before creating Exchange shadow copies, you must install the Exchange Support utility, as described in the following section.

---

The VSS-compliant shadow copy functionality is only supported on Windows 2003 systems.

---

The RepliStor Exchange Support utility performs the following tasks during the install process:

- ◆ Validates that RepliStor Version 6.1 or later is installed.
- ◆ Validates that Exchange 2003 is installed.
- ◆ Installs the latest Exchange 2003 .dll into the RepliStor directory.
- ◆ Prompts the user for failover parameters.
- ◆ Prompt the users for VSS parameters.
- ◆ Performs any configuration required.
- ◆ (Optionally) configures your system for failover.

---

RepliStor Exchange Support can be configured to set up failover. In a failover configuration, you must validate that the source and target paths are identical. This is a requirement for RepliStor Exchange Support to provide failover functionality.

---

In summary, You may use Exchange Support to take shadow copies, provide failover, or both.



### CAUTION

EMC recommends you move all Exchange data off the C:\drive. When VSS shadow copies are taken, if I/O is frozen on the system disk, it is not clear how the system will respond. In addition, you should have a minimum of two drives to separate logs and database; this is a Microsoft best practice recommendation.

## Installing RepliStor Exchange Support

To install RepliStor Exchange from a CD-ROM:

1. Start Windows.  
The user ID to log on to Windows must have administrative privileges, or Domain Admins if you are installing in a domain.
2. Insert the RepliStor Exchange Support CD-ROM into the CD-ROM drive. If the RepliStor Setup Wizard doesn't start automatically, click the `RepliStor.msi` file.
3. Follow the prompts in the Windows Installer program to install the program.

The RepliStor Exchange Support Wizard welcome screen displays as shown in Figure B-1.

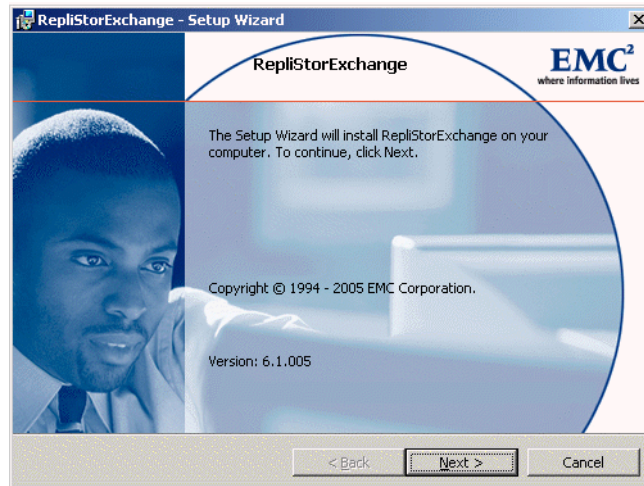


Figure B-1 Introduction Window

- Click **Next** and the **License Agreement** window displays as shown in Figure B-2.

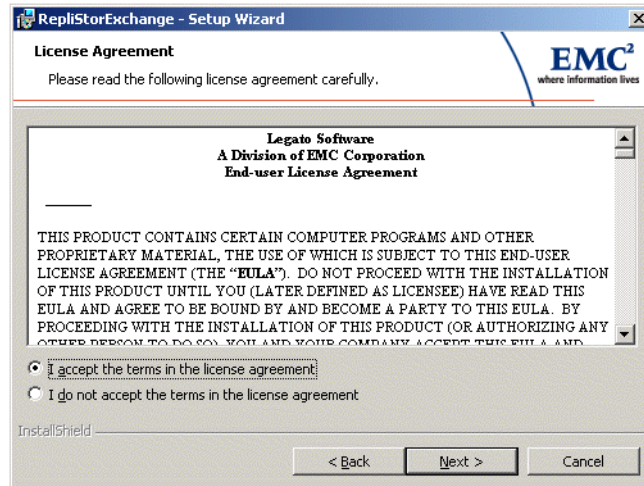


Figure B-2 License Agreement Window

- After accepting the license conditions, click **Next**, and the **Exchange Parameters - Alias** window displays as shown in Figure B-3.

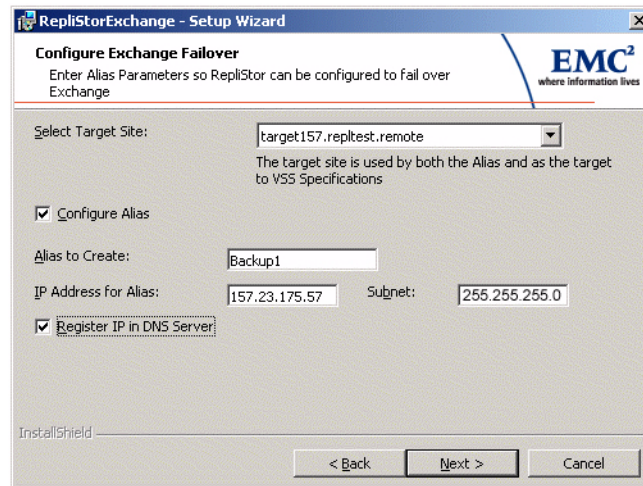
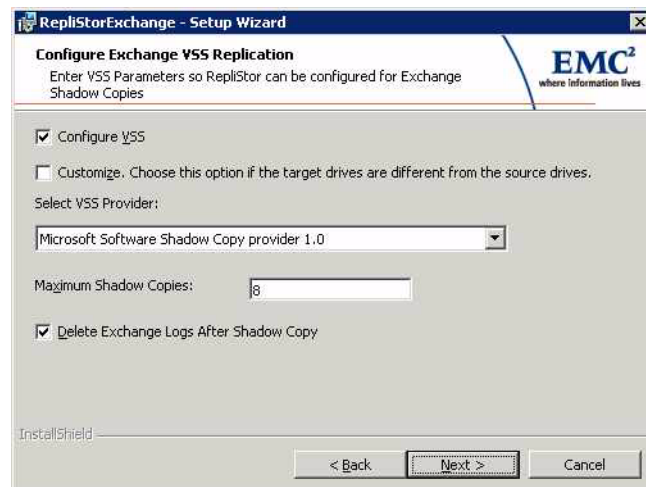


Figure B-3 Exchange Parameters — Alias Window



This window (Figure B-3) allows you to set alias parameters on your Exchange system so that RepliStor can be configured for failover. The configurable alias parameters are described below.

- ◆ **Select Target Site** — This option allows you to select the target site to which you want to replicate the files. This list is pre-filled from the RepliStor site list.
  - ◆ **Configure Alias** — Check this option to configure RepliStor for Exchange failover to the target site.
  - ◆ **Alias to Create** — Enter the alias name that will be created. The alias name should be unique on the network.
  - ◆ **IP Address for Alias** — Enter the optional IP address to be associated with the alias.
  - ◆ **Subnet:** — Enter the appropriate subnet.
  - ◆ **Register IP in DNS Server**— If this option is checked, the alias will be registered in the DNS server when it is activated. This option requires a dynamic DNS server.
6. After you set the alias Exchange parameters, click **Next**. The **Exchange Parameters - VSS** window displays as shown in Figure B-4.



**Figure B-4** Exchange Parameters - VSS Window

This window (Figure B-4) allows you to set VSS parameters on your Exchange system so that RepliStor can be configured correctly to

create and manage Exchange shadow copy backups. The configurable alias parameters are described below:

- ◆ **Configure VSS** — Check this option to configure RepliStor for VSS shadow copies. When checked, RepliStor will configure the VSS parameters in the specifications created. It will also create a script file for each storage group that will run the `eseutil.exe` utility on the shadow copy that was taken.
- ◆ **Customize** — If this is *not* checked, the specifications created will assume that the source and target paths are the same. You should check this option when you want the target path to be different than the source path for a specification.

For example, if you are not configuring failover, you may want to have the target system be the target for several source Exchange servers. In that case, you may want the target path to be different from the source path. Selecting this checkbox will cause the RepliStor Exchange Support Setup Wizard to prompt you for the target path for each specification. It will show you the storage group and the source path. The **Target Path** field will show the source path by default. You can then modify that path.

- ◆ **Select VSS Provider** — The target system is contacted and a list of all VSS providers on that system is presented to the user. This option allows you to select the VSS provider to use when taking a shadow copy on the target such as CLARiiON SnapView snapshots or Windows system shadow copies.

---

RepliStor does not support reverting shadow copies made using the CLARiiON provider.

---

- ◆ **Maximum Shadow Copies** — This option determines the maximum number of shadow copies you want to keep on the target prior to creating the shadow copy on the target. Old shadow copies are automatically deleted if the total number of shadow copy sets exceeds the user-defined maximum.

For example, if you enter 5 maximum shadow copies, this means that RepliStor will delete all shadow copies on the target except for the latest five before taking the target shadow copy. If you enter zero in this field, this means that you want all shadow copies deleted on the target before executing the current shadow copy.

- ◆ **Delete Exchange Logs after Shadow Copy** — If checked, Exchange logs are deleted after a successful shadow copy and RepliStor performs a full backup of the Exchange shadow copy; if unchecked, RepliStor will perform a copy backup of the Exchange shadow copy. In a full backup, all non-Exchange files (regardless of whether they have been marked as backed up), are saved. The backup history of each file is updated to show that it was backed up.

In a copy backup, files are copied to a backup medium regardless of the state of each file's backup history, and the backup history is not updated.

7. After you set the VSS Exchange parameters, click **Next**. The **Ready to Install the Program** window displays (Figure B-5).


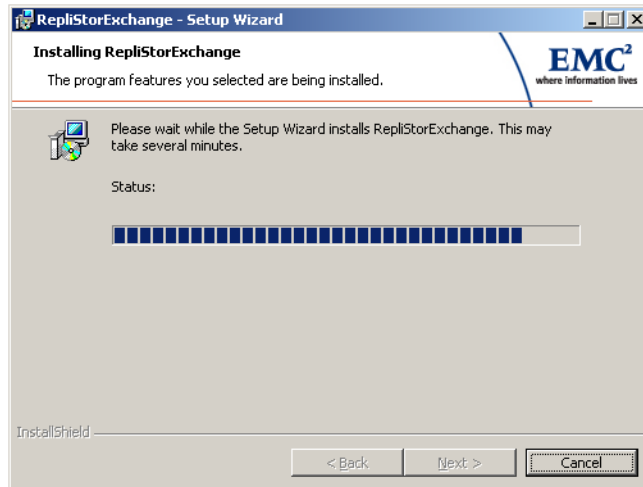


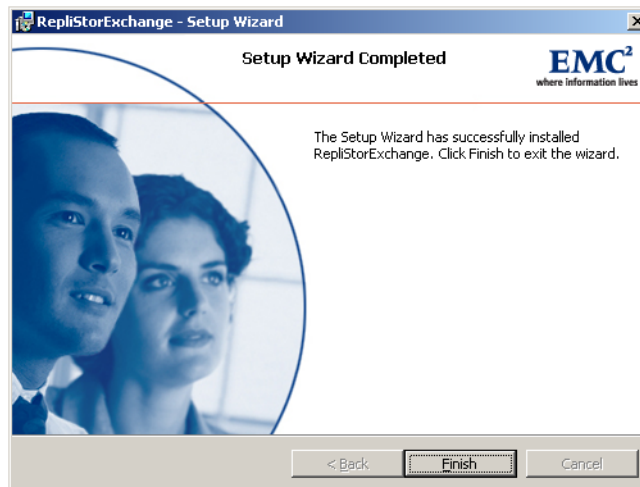
Figure B-5 Ready to Install the Program Window

- Click **Install**, and the Setup Wizard then installs RepliStor Exchange Support as shown in Figure B-6.



**Figure B-6** Installing RepliStor Exchange Support

- Click **Next** and the **Setup Wizard Completed** window displays as shown in Figure B-7. You have successfully installed RepliStor Exchange Support on your system.



**Figure B-7** Setup Complete Window

---

## Considerations When Using Failover

When using failover there are certain things to consider. Once the Exchange Installer completes you should:

1. Manually activate the alias from the Alias Maintenance dialog in the RepliStor client.
2. In the Add Computer Alias dialog, verify the services list as, by default, all Exchange services are added to the Services tab. The services used will vary depending on your Exchange configuration. For example not all Exchange implementation use POP3 or Site Replication Services. In addition any prerequisite services should be added to insure they are started, for example SMTP.



---

This appendix discusses security issues related to RepliStor software, including available security facilities, ramifications, and trade-offs.

This appendix contains the following sections:

- ◆ Security Levels..... C-2
- ◆ Encrypted File Support ..... C-4
- ◆ Setting Client Security ..... C-8
- ◆ Windows NT LAN Manager and Kerberos ..... C-9
- ◆ Account Settings..... C-10
- ◆ Checklist for Increasing the Security Level ..... C-12

---

## Security Levels

RepliStor software offers three levels of security: *connection*, *integrity*, and *encrypted*. The security level is set for *each* site in the **Site List** dialog box, or for *all* sites in the **Client Options** dialog box. When a client or server attempts to connect to a site, the target site properties determine the security level requested. If the target site does not support the requested security level, it reverts to the highest common security level. To verify the security level of all connections, check the **Connection** list (click **Client Connections** or **Site Connections** under the **Status** menu).

---

### Connection

The connection security level authenticates all server-to-server and client-to-server connections. It is the default security level when RepliStor software is installed. The authentication uses a *challenge-response* algorithm that encrypts the account, password, and domain information before passing it between the servers. It also passes the information so that a third party cannot record the exchange and replay it at a later time.

Points to remember about the connection security level include:

- ◆ It is the lowest level of security and, therefore, cannot be disabled.
- ◆ System clock times should be within 30 minutes of each other; otherwise, the server will think the challenge from the client has expired and authentication will fail.

A RepliStor server, by default, accepts a connection by any authenticated user listed in the Administrators group, or from another System account. This behavior can be modified on the **Users** tab of the **Options** dialog box, as shown in Figure 4-2 on page 4-3.

---

### Integrity

The integrity security level adds an additional element of security over the connection level. All messages are signed, which adds a 16-byte signature to each message. The contents of the signature depend on the message contents and the sequence of the messages. While a signature does not add privacy, it guards against tampering and can detect altered, out-of-sequence, and replayed messages.



---

## Encrypted

Encrypted is the highest security level. It performs all the same functions as the connection and integrity levels, plus it adds privacy. All messages are fully encrypted, in addition to being digitally signed. Encryption is appropriate when sensitive data travels over a public network, such as the Internet.

---

## Setting the Default Security Level for All Future RepliStor Sites

To set the default security level for all future RepliStor sites:

1. In the RepliStor client window, select **Client Options** from the **Maintenance** menu.
2. Click the **Comm** tab.
3. Select the security level from the **Default Security** list, and then click **OK**.

---

## Setting the Security Level for a Single RepliStor Site

To set the security level for a single RepliStor site:

1. In the RepliStor client window, select **Attach** from the **Functions** menu.
2. In the **Site List** dialog box, select a site from the **Site** list.
3. Click **Properties**.
4. In the **Site Properties** dialog box, select the security level from the **Security** list.
5. To verify the security level of all connections, select **Client Connections** or **Site Connections** from the **Status** menu to view the **Connection List**.

When a client or server attempts to connect to a site, the target site communication properties are used. If the target site does not support the requested security level, then the highest common communication level is used.

## Encrypted File Support

Keep the following in mind when replicating encrypted files:

- ◆ Files are replicated in a *copy on close* mode. This means RepliStor software does not replicate file operations on an encrypted file as they occur, but waits until the file is closed. As soon as the file is closed, it is sent to the target. Note that *any* change to an encrypted file while it is open, no matter how small, results in the entire file being sent to the target.
- ◆ Certain file types, such as database files, are always open and never closed. These files should not be encrypted, since RepliStor software cannot replicate encrypted files that are always open.
- ◆ Files are transferred to the target in *raw* mode and are not decrypted during the transfer. Thus, you do not need to encrypt the communications line during the transfer.
- ◆ The file is not accessible from the target unless the certificate (including the private key) used to create the file on the source is transferred to the target. You must export these certificates from the source system *before* it is shut down. For instructions on how to export encrypted file certificates, refer to *Exporting Encrypted File Certificates from the Source to the Target* on this page.

---

### Exporting Encrypted File Certificates from the Source to the Target

1. Log on to the source server as the user for whom you want to export the certificate.
2. From the **Start** menu, select **Run**.
3. In the **Open** text box, enter **mmc**, and then click **OK**.

The Microsoft Management Console (MMC) opens, as shown in Figure C-1 on page C-5.

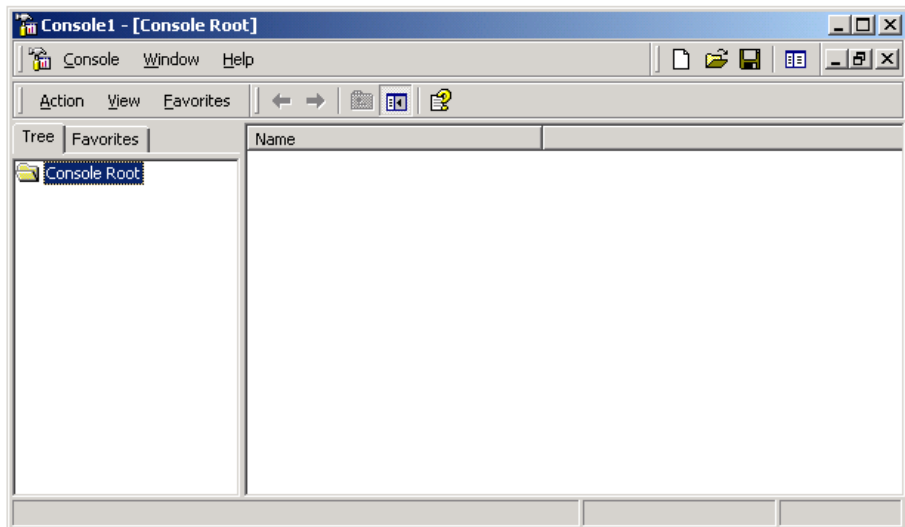


Figure C-1 Microsoft Management Console

4. From the **Console** menu, select **Add/Remove Snap-in**.
5. In the **Add/Remove Snap-in** dialog box, click **Add**.
6. In the **Add Standalone Snap-in** window, select **Certificates**, and then click **Add**.
7. In the **Certificates Snap-in** window, click **My User Account**, and then click **Finish**.
8. Close the **Add Standalone Snap-in** window, and then click **OK** in the **Add/Remove Snap-in** dialog box.
9. In the Tree frame of the **Console Root** window, click the plus sign next to **Certificates – Current User** to expand the list of folders.
10. Click the plus sign next to the **Personal** folder to expand it, and then click the **Certificates** folder.  
The certificates appear in the right pane.
11. In the **Issued To** column, right-click the certificate you want to export, and then select **All Tasks, Export**.

The **Certificate Export Wizard** opens, as shown in Figure C-2.



**Figure C-2** Certificate Export Wizard

12. Click **Next** to continue.
13. Click **Yes, export the private key**, and then click **Next**.
14. Accept the default selections on the **Export File Format** screen, and then click **Next**.
15. Type a password and confirm it, and then click **Next**.
16. Specify a name for the certificate, and then click **Next**.
17. Click **Finish** to export the certificate and close the wizard.
18. Repeat steps 1 through 17 for each user for whom you want to export certificates.
19. Back up all exported certificates to a floppy disk or some other secure location.

### Accessing Encrypted Files on the Target

1. Log on to the target server as the user for whom the certificate was created.
2. From the **Start** menu, select **Run**.
3. In the **Open** text box, enter **mmc**, and then click **OK**.

The Microsoft Management Console (MMC) opens.

4. From the **File** menu, select **Add/Remove Snap-in**.
5. In the **Add/Remove Snap-in** dialog box, click **Add**.
6. In the **Add Standalone Snap-in** window, select **Certificates**, and then click **Add**.
7. In the **Certificates Snap-in** window, click **My User Account**, and then click **Finish**.
8. Close the **Add Standalone Snap-in** window, and then click **OK** in the **Add/Remove Snap-in** dialog box.
9. In the Tree frame of the **Console Root** window, click the plus sign next to **Certificates – Current User** to expand the list of folders.
10. Click the plus sign next to the **Personal** folder to expand it, and then right-click the **Certificates** folder.
11. Select **All Tasks, Import**.  
The **Certificate Import Wizard** opens.
12. Click **Next** to continue. In the **File to Import** window, enter the path and filename of the certificate to import, or click the **Browse** button to search for it, and then click **Next**.
13. Enter a password and select any other necessary options, and then click **Next**.
14. Accept **Personal** as the default certificate store, and then click **Next**.
15. Click **Finish** to import the certificate and close the wizard.
16. Repeat steps 1 through 15 for each user for whom you want to import certificates.

The encrypted files can now be accessed on the target.

## Setting Client Security

Typically, there is no need to encrypt messages between the client and server since file data never travels on that connection. Thus, client connections default to Connection security.

An example of where you may want to encrypt messages passing between the client and the server is event reporting, which may include e-mail addresses.

To set all client-to-server communications to the same security level as server-to-server communications to ensure encryption:

1. In the RepliStor client window, select **Client Options** from the **Maintenance** menu.
2. In the **Client Options** dialog box, click the **Comm** tab, as shown in Figure C-3.

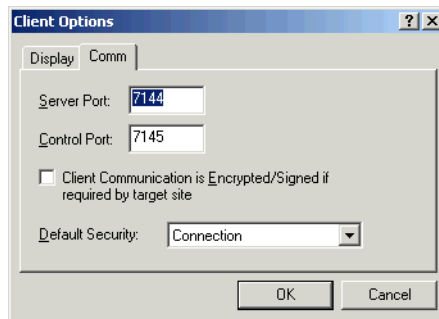


Figure C-3 Client Options — Comm Tab

3. Click **Client Communication Is Encrypted/Signed if required by target site**.
4. Select **Encrypted** from the **Default Security** list, and then click **OK**.

All directory and file displays in RepliStor software are accessed and all file specifications are validated according to the client's permissions. Thus, no client can view or replicate data that the client does not have access to. When you create a specification, the client's account and domain must be able to access the source and target systems and directories.

---

## Windows NT LAN Manager and Kerberos

---

### NTLM

RepliStor software uses the standard Windows NT 4.0 security provider, NT LanManager (NTLM), for authentication, signature, and encryption. NTLM Version 2 provides greater security than earlier versions, so use Version 2 if you require additional security. For more information, refer to Microsoft Knowledge Base article Q147706, *How to Disable LM Authentication on Windows NT*, available on the Microsoft website.

RepliStor also supports an additional security provider, Kerberos.

---

### Kerberos

All connections between the Administrative client and the RepliStor service, and also between the source RepliStor service and the target RepliStor service, need to be authenticated<sup>1</sup>. RepliStor versions prior to Version 6.0 used the NTLM security provider for this function. RepliStor Version 6.1 supports Kerberos in addition to NTLM.

Kerberos provides the following advantages:

- ◆ Both the source and target RepliStor services are authenticated. In NTLM, the target can validate the identity of the source, but the source is not sure it connected to the intended target.
- ◆ A service running in the System context can be authenticated using Kerberos. In NTLM, a process running in the System authenticates as any other process running System in the domain. This can pose a security risk.
- ◆ Authentication takes fewer challenge-response cycles.
- ◆ Kerberos provides a higher level of security over NTLM and NTLM version 2.

---

1. For the purposes of this section, *client* refers to the process initiating the connection and *server* refers to the target of the connection. In RepliStor, the Administrative client is always the client, and the source RepliStor server is the client to the target server.

There are restrictions in using Kerberos. If a system cannot authenticate using Kerberos, it will then attempt to use NTLM. The following conditions must be met to use Kerberos:

- ◆ Both the source and target systems must be in the same Windows 2000 or later domain.
- ◆ The computer name the client is connecting to needs to be the actual name of the target computer. For example, it will fail when connecting to a virtual computer name in a cluster.

To configure a system to accept only Kerberos connections, do the following:

1. Select **Options** from the **Maintenance** menu.
2. Select the **Users** tab and check **Accept Only Kerberos Connections**. Note that when this option is selected, the **Disallow System Account** option is disabled. When using Kerberos, a System account is fully authenticated, so this is not necessary.

To determine how a connection authenticated, select **Client Connections** and/or **Site Connections** from the **Status** menu.

---

## Account Settings

RepliStor software is installed as a service that runs under the System account by default. The System account is not an authenticated account and there is no way for RepliStor software to distinguish one System account from another. For increased security, RepliStor software can be run under an authenticated account instead of under the System account; however, the authenticated account must be a member of the Administrators group.

If all RepliStor servers are running under an authenticated Administrators account, then they should be configured to *not* accept connections from System accounts. All connections must originate from authenticated administrators.



---

## Disallowing a System Account

To configure RepliStor servers to disallow a System account:

1. In the RepliStor client window, select **Options** from the **Maintenance** menu.
2. In the **Options** dialog box, click the **Users** tab.
3. Select **Disallow System Account (Authenticated Access Only)**, and then click **OK**.

---

## Assigning RepliStor Administrative Permissions

RepliStor provides a range of administrative permissions, which enables RepliStor to delegate administration to a variety of users and groups. The administrative permissions are shown in Figure C-4.

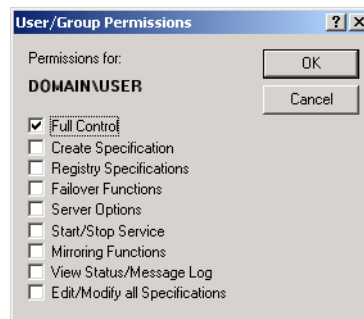


Figure C-4 User/Groups Permissions

The administrative permissions include:

- ◆ **Full Control** — Allow full administration control.
- ◆ **Create Specification** — Allow creation, modification, and deletion of a specification.
- ◆ **Registry Specifications** — Allow creation, modification, and deletion of a registry specification.
- ◆ **Failover Functions** — Allow access to **Alias Maintenance** and **Target Options** dialog boxes. This includes all Alias functions such as create alias, manual fail, set source, and so on.
- ◆ **Server Options** — Allow server options (on **Options** dialog box), schedules, and administrative tasks.

- ◆ **Start/Stop Service** — Allow the user to start or stop a RepliStor service.
- ◆ **Mirroring Functions** — Allow starting and stopping of mirroring, forwarding, updating, and so on.
- ◆ **View Status/Message Log** — Allow log messages to be viewed.
- ◆ **Edit/Modify all Specifications** — Allow the user to view, modify, and delete specifications created by any user. If this permission does not exist, the user can view, modify, and delete only those specifications that were created by the user.

---

All directory/file displays are accessed with respect to the client's permissions. Also, all file specifications are validated with respect to the client's permissions. This means no client can view or replicate data to which it does not have access. When creating a specification, the source and target computers and directories must be accessible by the client's account/domain.

---

---

## Checklist for Increasing the Security Level

When RepliStor software is installed, it defaults to a *connection* security level. To increase the security level, use any of the following methods discussed in this chapter:

- ◆ Configure your system to accept only Kerberos connections, as outlined in *Windows NT LAN Manager and Kerberos* on page C-9.
- ◆ Use Version 2 of the NTLM security provider.
- ◆ Configure RepliStor to run under an account in the Administrators group.
- ◆ In the **Options** dialog box, on the **Users** tab, click **Disallow System Account (Authenticated Access only)**, and then click **OK**.

If there is a possibility of replicating within an unsecure environment (for example, over the Internet or over a campus-wide LAN), then all sites should be set to the integrity or encryption security level. If the RepliStor client accesses sensitive data (such as e-mail addresses), then the client should also be set to encrypt data.

## Files in the Data Directory

This appendix provides information about files that reside in the RepliStor data directory. The data directory is the main location where data on the target system is stored.

***Important:*** Deleting the data directory destroys replicated data, and should be done only when recovering the original environment.

This appendix contains the following section:

- ◆ Files in the Data Directory .....D-2

## Files in the Data Directory

Table D-1 lists the files in the data directory and describes each file's function.

**Table D-1 Files in the Data Directory**

Filename	Description
ActiveFile.rdf	Tracks all files open on the source system. When a failover occurs, ActiveFile.rdf displays a log message that lists all the files open at the time of the failover. This warns the administrator of potential problems with these files.
Block.rdf	Tracks file operations for all blocked target files. Also tracks file operations for all updates if the <b>Pause Updates</b> option is enabled.
KernelCache.rdf	Contains file update, delete, rename, and copy commands recorded by the kernel driver for mirrored files. This is a memory mapped file and the current contents are in the memory segment. The file is written to only during shutdown. The memory is organized as a circular first-in/first-out queue. If a file operation cannot be queued because it is full, the operation creates OC\$nnnnn files.
Msg.rdf	Holds the RepliStor message log.
OC\$nnnnn.rdf	Contains file update, delete, rename, and copy commands recorded by the kernel driver for files being mirrored. These files are created only if the kernel cache is full. When all data in the files has been sent to and acknowledged by all target machines, the files are marked as complete and then deleted.
send.<target_site>.ack.rdf	Contains an entry for each target machine that is to receive file updates. Each entry contains the last offset kernel cache entry or OC\$nnnnn file and is offset into the file that the target machine acknowledged receiving. This information is used on restart, so RepliStor software is aware of which files it has already sent. If a <b>Kernel Cache Corrupt</b> error message appears, the most likely cause is that this file and the kernel cache were not updated together.
send.lid.rdf	Contains the latest load ID for the sender thread in the RepliStor service. Whenever the sender starts or sends a new configuration to a remote site, the value in this file increases by one. This ensures that if RepliStor software shuts down and restarts, all targets will know their configuration is outdated and will request another one.
Send.rdf	Holds all <b>Pending Sync</b> entries. These are created during a synchronization if the sender cannot open a file because it is in use.
Sharedmem.dat.rdf	Memory mapped file that contains all Performance Monitor counters.

Table D-1 Files in the Data Directory (continued)

Filename	Description
<code>Sntp.rdf</code>	Holds all messages queued to be sent as e-mail.
<code>SyncStatus.rdf</code>	Holds all <b>Sync Status</b> entries.
<code>TargetFileLog.2005.01.03.15.50-2005.01.03.17.59.csv</code>	<p>Target file log messages will only be in the data directory if the user has configured target file logging. If the machine is a target, it will list changes it has made in this file. The file that contains the target file log messages consists of many column fields, including:</p> <ul style="list-style-type: none"> <li>• Path</li> <li>• Site</li> <li>• Last Modification</li> <li>• WRITE</li> <li>• TRUNCATE</li> <li>• ROLLLOG</li> <li>• DELETE</li> <li>• RENAME</li> <li>• CONFIG</li> <li>• SYNC_ID</li> <li>• FILE_EXIST</li> <li>• COPY</li> <li>• ADD_SHARE</li> <li>• DEL_SHARE</li> <li>• COPYDONE</li> <li>• CLSAPPFILS</li> <li>• CLOSE</li> <li>• ATTRIBUTES</li> <li>• ACLS</li> <li>• ACLS_DATA</li> <li>• BACKUPWRIT</li> <li>• BACKUPEND</li> <li>• CREATE</li> <li>• ACLS_TEXT</li> <li>• SYNC_COPY</li> <li>• SETCOMPRES</li> <li>• FILEHEADER</li> <li>• SYNCDONE</li> <li>• RESYNCCOPY</li> <li>• RESYNCDONE</li> <li>• REGISTRY</li> <li>• CHECKPOINT</li> </ul>



---

This appendix describes the various utilities that can be used with the RepliStor software.

This appendix contains the following sections:

- ◆ Using the regutil Utility ..... E-2
- ◆ Using the Set Active Computer Name Utility ..... E-3

## Using the regutil Utility

You can use the `regutil` utility to copy registry keys from one machine to another or from one key to another on the same machine. This utility is installed with the product under the *EMC RepliStor* directory.

**Important:** Use this utility to copy keys *only* under the `HKEY_LOCAL_MACHINE` hive in the registry.

The `regutil` utility uses the following command line syntax:

```
regutil src key [/s src] [/d dest] [/k dest key] [/p][/y]
```

Table E-1 lists the parameters available with the `regutil` utility.

**Table E-1 regutil Parameters**

Parameter	Description
<code>src key</code>	Source key without <code>HKEY_LOCAL_MACHINE</code> prefix.
<code>/k dest key</code>	Destination key without <code>HKEY_LOCAL_MACHINE</code> prefix (if omitted, same as <code>src key</code> ).
<code>/s src</code>	Source machine name (if omitted, local machine).
<code>/d dest</code>	Destination machine name (if omitted, same as source).
<code>/p</code>	Copies security descriptors for each key.
<code>/y</code>	Synchronizes the target with the source—delete values and keys on the target that do not match the source.

**Important:** You must be familiar with the information a specific application places in the registry to know which keys to replicate. In addition, do not replicate machine-specific information from a source to a target key as this can cause serious problems on the target.

Use the `regutil` utility with caution.

### Example: Using the regutil Utility to Copy Registry Keys

To copy all keys under `HKLM\Software\Legato Systems` from the local machine to machine A, enter the following:

```
regutil "software\Legato Systems" /d A
```

To copy all keys under `HKLM\Software\abc` to `HKLM\Software\def` on machine B:

```
regutil software\abc /k software\def /s B
```



## Using the Set Active Computer Name Utility

You can use the Set Active Computer Name (`SETACN`) utility in failover scripts to temporarily set the **Active Computer Name** in the Windows registry to the target system name, then set it back to the source system name after the source is restored.

The command line syntax is:

```
setacn [new_active_computer_name]
```

where *new\_active\_computer\_name* replaces the current computer name, which is saved. If you do not specify a name, the active computer name is restored to the saved name.

You might include this utility in a script (specified in the **After:** textbox in the **Commands: Adding Alias** group box in the **Add Computer Alias** dialog box) that is run after a failover occurs. For example, you might want to:

- ◆ Stop the printer spooler that was going to the failed source system.
- ◆ Set the active computer name to another system temporarily.
- ◆ Start the printer spooler (now directed to the new active computer name).
- ◆ Reset the active computer name to the saved name when the source is restored.



---

This glossary contains terms related to EMC RepliStor for Microsoft Windows software. Many of these terms are used in this guide.

### A

- alias** An alias is a virtual name for the source system you can define for use with an Alias failover. The alias is sent to the target in a failure and clients can continue to connect to it while the actual source system is unavailable. When the source system is backed up, clients connect to the alias on the source with no disruption in connection.
- attach** Connecting the RepliStor client to the RepliStor server for configuration and administration.

### B

- blocked** When the target system cannot apply updates to its local copy of the files being replicated, the files are considered blocked.
- When data cannot be sent to the target system because the network connection is broken or the target system is down, a site is considered blocked.

**F**

**failover** The process in which the identity, IP addresses, and services of one system are moved over to another system when the first system fails. Also sometimes called *switchover*.

**file attributes** File attributes are flags that define file types as hidden, read-only, archive, or system. These flags also control the process of updating the file times (for create, modify, and access).

**forwarding** The process by which RepliStor software sends mirrored file updates from the source system to the specified target system.

**G**

**giveup delay** The date and time RepliStor software stops trying to synchronize the specification after synchronization is unsuccessful.

**global exclude** A type of RepliStor specification, specifying those files, directories, or shares that are not included in data mirroring. A global exclude overrides other types of specifications.

**global forwarding** Stops and starts the forwarding process and the system's heartbeat signal.

**GUID** Globally unique identifier, a 128-bit integer that uniquely identifies a particular object class or interface.

**H**

**heartbeat** The *I'm Alive* signal that the source system sends periodically to the target system, letting it know that the source is still running.

**K**

**kernel cache** A fixed amount of shared memory used for queuing data going from the source to the target.

**L**

**license key** The unique alphanumeric string needed to run RepliStor software. You must enter the license key using uppercase letters.

**log files, send and receive** Disk-resident files used by RepliStor software to send and receive data between source and target systems. A log file entry is written for each file update associated with a mirrored file.

## M

**Macintosh-accessible volumes** Volumes on a Windows system that store Macintosh files.

**manual alias activation** The process by which you create an alias for the source system that you manually activate on the target system after failover, rather than allowing RepliStor software to automatically activate it. Select this option when there are too many variables in the environment (for example, across a WAN) to make automatic failover viable.

**message log** A file RepliStor software creates listing all the operational messages that occurred during normal operation.

**mirroring** The process by which RepliStor software captures the updates to a user-specified file in the kernel cache on the source system.

**monitor bar** A part of the RepliStor user interface used to indicate the summary status of all attached RepliStor sites.

**MSCS** Microsoft Cluster Server, which is designed to keep server-based applications available regardless of component failures. RepliStor software is *cluster-aware* (manageable as a cluster resource).

## P

**performance monitor DLL** Application extension DLL used to provide performance statistics in the Windows Performance Monitor, allowing you to go into the Performance Monitor and see statistics on RepliStor performance.

## R

**redirected drive** A mapped drive, not a physical drive, on the local system.

**Reflect Protection** A RepliStor feature that prevents changes coming from the source to be mirrored back to the source from the target. Changes that originate from the target, however, are still updated on the source.

Use the Reflect Protection feature only when the target is mirrored back to the source, in a circular path. When this option is used, it allows circular mirroring to operate correctly.

**RepliStor client** The RepliStor user interface, used to configure and administer the RepliStor product.

**RepliStor device** The component that works with the RepliStor server to provide capabilities for data mirroring.

**RepliStor server** The RepliStor background service that runs on a Windows system to provide capabilities for failover and data mirroring.

**RepliStor site** A Windows system, typically a server running the RepliStor server.

**role reversal** Redefining the *target* machine as the *source* machine and the *source* machine as the *target* machine.

## S

**script repository** A directory on a file server that is accessible by all RepliStor servers using a UNC path. It allows scripts to be centrally located and administered. Before executing a script, if a script repository is defined, RepliStor software copies the script from the repository to the local script directory. After script execution, RepliStor software deletes the script file from the local directory.

**shadow copy** A consistent point-in-time copy of data.

**site forwarding** Stops and starts the forwarding process but does not stop and start the system's heartbeat signal.

**SNMP agent extension DLL** Extends the standard Windows SNMP service to allow RepliStor software to send its messages as SNMP events, enabling RepliStor software to work with systems management software.

**source system** The origination site for mirrored files. In the context of failover, this is the system protected by the failover and whose identity is switched over to a target system if the source fails.

**specification** Information that identifies the files, directories, and shares to be mirrored.

**synchronize** The process by which RepliStor software ensures there is an exact copy of the mirrored data from the source system on the target system. Data must be synchronized before mirroring can start.

## T

**Target File log** A log that lists which files were updated on the target system and the operations that occurred for those files.

**target system** The destination site of mirrored files. In the context of failover, this is the system ready to take over the identity of a source system, should that source system fail.

**throttling** A rate that establishes the maximum level of data to be transferred between a source system and a target system.

## U

**unblocking** The reapplication of file updates to files or the reestablishment of communications between source and target systems.

**UNC** Abbreviation for Universal Naming Convention, a PC format used to specify the location of resources on a LAN. With RepliStor software, UNC paths are used to specify the target data, if the target is a network attached storage (NAS) device. UNC uses the following format: `\\server-name\shared_resource_pathname`.

**updating** The process by which RepliStor software updates mirrored files on the target system with changes received from the source.

## V

**virtual drive** The same as a *redirected drive*.

**Volume Shadow Copy Service (VSS)** A framework for creating shadow copies of files and directories at predefined points in time in Microsoft Windows Server 2003 operating systems.

*See also shadow copy.*

## W

**workspace** A part of the RepliStor client interface that displays information on the status of data mirroring at the attached site.





**A**

- active computer name, setting E-3
- Active Directory, RepliStor site list with 2-4
- administering RepliStor sites 6-10
- Alias failover
  - recovery with aliases 7-2
  - source system 5-5
- aliasadd command 8-4
- aliases
  - associating IP addresses with 5-6
  - manually activating after failover 5-15
  - overview 1-5
  - recovery using 7-2
  - setting to failover to original source 7-3
- aliasop command 8-5
- application files, closing 6-48
- attaching to RepliStor sites 2-6
- availability 1-4

**B**

- blocked
  - files, correcting 6-22
  - sites
    - correcting 6-21
    - e-mail notification 6-21
- blockedsites command 8-6

**C**

- checkpoint command 8-6
- circular mirroring 2-37
- client options 4-2
- Client window

- keyboard shortcuts 2-11
- monitor 2-11
- opening 2-6
- toolbar 2-7
- workspace 2-9
- close command 8-7
- closing application files 6-48
- cluster resource, creating 6-36, 8-10
- commands
  - aliasadd 8-4
  - aliasop 8-5
  - blockedsites 8-6
  - checkpoint 8-6
  - close 8-7
  - configexport 8-8
  - configimport 8-8
  - deletedatadir 8-8
  - forwarding 8-8
  - isidle 8-8
  - mirroring 8-9
  - modifyspec 8-9
  - numblockedfiles 8-10
  - pause 8-11
  - pauseattributes 8-11
  - report 8-11
  - ShadowCopy 8-13
  - ShadowUtil Delete 8-14
  - ShadowUtil Mount 8-14
  - ShadowUtil Properties 8-15
  - ShadowUtil Revert 8-15
  - ShadowUtil Unmount 8-15
  - simulate 8-11
  - siteforwarding 8-12
  - spec 8-15

- spec/globalexclude 8-18
- spec/registry 8-19
- spec/shares 8-21
- specexport 8-21
- specimport 8-22
- specstatus 8-22
- start 8-23
- stop 8-23
- sync 8-23
- throttle 8-25
- unblockfiles 8-27
- common mistakes to avoid 2-2
- communication errors, correcting 6-21
- compressed files 2-15
- configexport command 8-8
- configimport command 8-8
- configuring
  - client options 4-2
  - forwarding options 4-5
  - options 4-3
  - throttling options 4-6
  - updating options 4-5

**D**

- Delete
  - orphans 2-2
- Delete directory 2-21
- deletedatadir command 8-8
- deleting specifications 6-31
- detaching from a site 2-42
- disabling specifications 6-31
- DiskXtender 2000 6-8
- domain credentials for failover 5-5

**E**

- e-mail notification, blocked sites 6-21
- enabling specifications 6-31
- encrypted files
  - accessing C-6
  - exporting certificates from the source to the target C-4
  - importing certificates to the target C-6
- Exchange support B-2
- installing B-3

**F**

- failover
  - agent for WAN failover 5-17
  - alias 1-5, 5-2
  - configuring for WAN 5-17
  - configuring on Macintosh 5-20
  - configuring options on the source system 5-5
  - configuring source system options 5-5
  - manually forcing with aliases 5-13
  - process 1-4
  - tab 5-17
- file
  - attribute updates 6-44
  - blocked 6-20
  - forwarding, turning on and off 6-42
  - mirroring, turning on and off 6-41
  - updates
    - checking status of 6-19
    - pausing 6-43
- file/directory specifications
  - creating 2-18
  - creating for Microsoft Cluster Server 6-36
- forwarding
  - command 8-8
  - global, turning on and off 6-42
  - IP address 1-5
  - options, configuring 4-5
  - site, turning on and off 6-42

**G**

- global exclude specifications
  - creating from command line 8-18
  - creating using Client window 2-16

**H**

- heartbeat signal 6-42, 8-8, 8-13

**I**

- I'm Alive 1-5
  - See also* heartbeat signal
- install
  - Exchange support B-3
- IP address, associating with aliases 5-6
- isidle command 8-8

**K**

Kerberos C-9, C-12  
 kernel cache sizing 6-3  
 keyboard shortcuts 2-11

**L**

LAN 1-2

**M**

Macintosh  
     failover 5-20  
     specifications 2-36  
 manual alias activation 5-15  
 message log  
     marking messages as read 6-47  
     purging 6-48  
     viewing 6-46  
     working with 6-45  
 Microsoft  
     Management Console (MMC) C-4  
     Volume Shadow Copy Service (VSS) 1-7,  
     2-23, 3-2, 3-10, B-2  
 Microsoft Cluster Server  
     associating applications 6-38  
     changing the RepliStor cluster resource 6-39  
     Cluster Group 6-38  
     creating a RepliStor cluster resource 6-36  
     creating global exclude and mirror files  
     specifications for 6-36  
     managing specifications for 6-36  
     replicating data with 6-39  
 mirroring  
     command 8-9  
     turning on and off 6-41  
 mistakes to avoid 2-2  
 modifying specifications 6-30  
 modifyspec command 8-9  
 monitor  
     description of 2-11, 6-13  
     increasing the size of 6-13  
     troubleshooting site problems with 6-20  
     using to check site status 6-13

**N**

NAS device, targeting 6-23

NICs, configuring two adapters 5-14  
 NTFS, special user ID requirements 2-16  
 NTLM C-9, C-12  
 numblockedfiles command 8-10

**O**

open file synchronization 2-39, 6-34  
 options  
     client 4-2  
     forwarding 4-5  
     RepliStor 4-3  
     throttling 4-6  
     updating 4-5

**P**

pause command 8-11  
 pauseattributes command 8-11  
 pausing file updates 6-43  
 pending synchronizations 6-18  
     reviewing and reissuing 6-34  
 Performance Monitor  
     DLL 1-3  
     using with RepliStor 6-2  
 Performance tab 6-5  
 Processes tab 6-7  
 purging the message log 6-48

**R**

redirected drives 2-15  
 Registry keys  
     copying E-2  
     mirroring 2-30  
     specifications 2-30  
 regutil utility E-2  
 RepliStor  
     Active Directory 2-4  
     Client 1-3  
     starting 2-6  
     Client window  
     keyboard shortcuts 2-11  
     monitor 2-11  
     toolbar 2-7  
     cluster resource, changing 6-39  
     components 1-3  
     configuring with two NIC adapters 5-14

- control service 1-3
- driver 1-3
- Exchange support B-2
  - installing B-3
- keyboard shortcuts 2-11
- monitor 2-11
- options 4-3
- Server 1-3
- setup and uninstall programs 1-3
- sites
  - administering 6-10
  - attaching to 2-6
  - checking status of 6-10
  - detaching from 2-42
- toolbar 2-7
- using Performance Monitor with 6-2
- workspace 2-9
- report command 8-11
- reporting bandwidth 6-50

**S**

- scripts
  - environment variables 5-10
  - script repository 5-11
  - using 5-10
- security
  - Kerberos C-9, C-12
  - NTLM C-9, C-12
  - setting client level C-8
  - setting levels C-3
- SETACN E-3
- setting the active computer name E-3
- ShadowCopy
  - command 8-13
  - Delete 8-14
  - Mount 8-14
  - Properties 8-15
  - Revert 8-15
  - Unmount 8-15
- ShadowUtil commands
  - Delete 8-14
  - Mount 8-14
  - Properties 8-15
  - Revert 8-15
  - Unmount 8-15
- shares specifications, creating from CLI 8-21

- Simple Network Management Protocol (SNMP)
  - 1-3
- simulate command 8-11
- simulating bandwidth 6-50
- siteforwarding command 8-12
- sites
  - administering 6-10
  - attaching to 2-6
  - blocked 6-20
  - checking file update status 6-19
  - checking status
    - using the monitor 6-13
    - using traffic lights and computer icons 6-10, 6-12
  - checking synchronization status 6-17
  - detaching from 2-42
  - finding specification information 6-16
  - forwarding, turning on and off 6-42
  - target, finding information about 6-15
  - troubleshooting problems 6-20
- sizing kernel cache 6-3
- snapshot 1-7, 2-23, 3-2, 3-10, B-2
- SNMP Agent Extension DLL 1-3
- source system options 5-5
- spec command 8-15
- spec/globalexclude command 8-18
- spec/registry command 8-19
- spec/shares command 8-21
- specexport command 8-21
- specifications
  - circular mirroring 2-37
  - creating for Microsoft Cluster Server 6-36
  - creating from command line 8-15
  - deleting 6-31
  - displaying status 8-22
  - enabling and disabling 6-31
  - file/directory 2-18
  - finding information about 6-16
  - global exclude, creating 2-16, 8-18
  - managing for Microsoft Cluster Server 6-36
  - mirror files, creating 8-15
  - modifying 6-30, 8-9
  - overview of 1-2
  - Registry 2-30
  - registry 8-19
  - shares, creating 8-21
  - synchronizing 2-39, 6-32

- after recovery 7-3
- specimport command 8-22
- specstatus command 8-22
- start command 8-23
- starting the Client 2-6
- status
  - checking
    - file update 6-19
    - specification information 6-16
    - synchronization 6-17
    - target site information 6-15
  - checking site 6-10
    - using the monitor 6-13
    - using traffic lights and computer icons 6-10, 6-12
- stop command 8-23
- sync command 8-23
- synchronization
  - open file 2-39, 6-34
  - options 6-32
  - pending 6-18, 6-34
  - specifying options 6-32
  - status 6-17
- synchronizing
  - specifications 2-39, 6-32
  - specifications after recovery 7-3
- System account 2-16

**T**

- tape backups 6-49
- target
  - pause updates 2-26, 3-9
- target sites 6-15
- throttle command 8-25
- throttling, configuring 4-6
- To B-3
- toolbar
  - customizing 2-7
  - in Client window 2-7
- traffic lights
  - in workspace 2-10

- on monitor 2-11
- using to check site status 6-10, 6-12

troubleshooting

- blocked sites 6-21
- site problems 6-20
- viewing the message log 6-46

## U

- unblockfiles command 8-27
- UNC path
  - defined g-5
  - deleting 6-28
  - disabling 6-27
  - enabling 6-27
  - modifying 6-28
  - registering 6-23
- updates
  - file attribute, stopping and starting 6-44
  - pausing file 6-43
- updating options 4-5
- usage considerations 2-2
- user ID, special requirements for NTFS 2-16

## V

- virtual drives 2-15
- Volume Shadow Copy Service (VSS) 1-7, 2-23, 3-2, 3-10, B-2

## W

- WAN 1-2
  - failover 1-6, 5-17
- Windows 2000
  - Active Directory 2-6
  - configuring failover 5-2
- workspace
  - changing the split of 2-9
  - closing the message pane 2-9
  - description of 2-9
  - traffic lights 2-10

