Open Systems SnapVault® 2.2
Installation and Administration Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: http://www.netapp.com

Part number 215-01627_A0
November 2005

Copyright and trademark information

Copyright information

Copyright © 1994–2005 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, which is copyrighted by Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

   This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER

IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software contains materials from third parties licensed to Network Appliance Inc. which is sublicensed, and not sold, and title to such material is not passed to the end user. All rights reserved by the licensors. You shall not sublicense or permit timesharing, rental, facility management or service bureau usage of the Software.

Portions developed by the Apache Software Foundation (http://www.apache.org/). Copyright © 1999 The Apache Software Foundation.

Portions Copyright © 1995–1998, Jean-loup Gailly and Mark Adler

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

Portions Copyright © 1995 University of Southern California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

Portions of this product are derived from version 2.4.11 of the libxml2 library, which is copyrighted by the World Wide Web Consortium.

Network Appliance modified the libxml2 software on December 6, 2001, to enable it to compile cleanly on Windows, Solaris, and Linux. The changes have been sent to the maintainers of libxml2. The unmodified libxml2 software can be downloaded from http://www.xmlsoft.org/.

Copyright © 1994–2002 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. http://www.w3.org/Consortium/Legal/

Software derived from copyrighted material of the World Wide Web Consortium is subject to the following license and disclaimer:

Permission to use, copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications, that you make:

The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.

Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [$date-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. http://www.w3.org/Consortium/Legal/.

Notice of any changes or modifications to the W3C files, including the date changes were made.

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS

FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

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

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

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

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

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp–the Network Appliance Company, DataFabric, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. Data ONTAP, gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexVol, FPolicy, HyperSAN, InfoFabric, LockVault, Manage ONTAP, NOW, NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapMigrator, SnapSuite, SnapValidator, SohoFiler, vFiler, VFM, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the United States. Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, and SpinServer are registered trademarks of Spinnaker Networks, LLC in the United States and/or other countries. SpinAV, SpinManager, SpinMirror, SpinRestore, SpinShot, and SpinStor are trademarks of Spinnaker Networks, LLC in the United States and/or other countries.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

Preface ix

Chapter 1 **Introduction to Open Systems SnapVault** 1

About Open Systems SnapVault 2

Open Systems SnapVault Features 8

Chapter 2 **Installing the Open Systems SnapVault Software**. 13

Before you install 14

Installing Open Systems SnapVault on Windows platforms 18

Installing Open Systems SnapVault on UNIX and Linux platforms 24

Verifying the installation 31

Upgrading the Open Systems SnapVault agent to version 2.2 32

Uninstalling the Open Systems SnapVault 2.2 agent. 35

Unattended installation and upgrade 37

Chapter 3 **Configuring Open Systems SnapVault**. 43

Configuration interfaces 44

Understanding the Configurator utility interface. 45

Understanding the svsetstanza command 47

Running the Configurator utility 50

Confirming that services are running. 52

Modifying SnapVault parameters 53

Generating and stopping files for debugging. 56

Setting block-level incremental backup options 60

Configuring backup exclusion lists. 62

Configuring open file backup for Windows 65

Chapter 4 **Setting up backups and restores** 69

Setting up SnapVault backup on Open Systems platforms. 70

Configuring the SnapVault secondary storage system 71
Creating an initial baseline copy 72
Scheduling SnapVault update backups. 73

Restoring data. 74
Restoring a directory or a file 75
Restoring an entire primary storage system 78
Restoring files to a primary storage system from tape 79

Chapter 5 **Open Systems SnapVault Management** 81

Locating status and problem reports 82

Backing and restoring the Open Systems SnapVault database. 83

Backing up and restoring Windows System State data. 88
Adding System State data backup 91
Restoring System State data 92
Using System State backup to rebuild a primary system 94

Deleting and recreating Open Systems SnapVault relationships. 97

Migrating a relationship between two secondary systems 98

Migrating between two volumes on one secondary system 101

Setting up a tertiary system for a relationship 104

Reusing a deleted or renamed directory name 105

Resynchronizing restored or broken relationships 107

Retrying failed transfers 109

Encrypted File System (EFS) encrypted file backup and restore. 110

Chapter 6 **Open Systems SnapVault space estimator**. 111

Appendix A **The Logical Replication (LREP) tool** 121

Appendix B **Open Systems SnapVault Error Messages**. 125

Preface

About this guide

This guide describes how to use the Open Systems SnapVault® software product on platforms such as such as Windows®, UNIX®, and Linux.

Audience

This guide is for system administrators who are familiar with operating systems such as UNIX, Windows 2000, and Windows Server™ 2003 and storage systems running Data ONTAP™ software. It also assumes that you are familiar with how to configure these systems and how the NFS, CIFS, and HTTP protocols are used for file sharing or transfers. This guide doesn't cover basic system or network administration topics; it emphasizes the characteristics of the Open Systems SnapVault product and how it is used with the NetApp® storage systems.

Terminology

When the term "filer" is used, the information applies to all Network Appliance™ filer models.

Command conventions

You can enter filer commands on the system console or from any client that can obtain access to the filer using a Telnet session. In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ, depending on your version of UNIX.

Formatting conventions

The following table lists different character formats used in this guide to set off special information.

Formatting convention	Type of information
Italic type	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the `fctest` *adaptername* command, you enter the characters "fctest" followed by the actual name of the adapter. ◆ Book titles in cross-references.

Formatting convention	Type of information
`Monospaced` font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ The contents of files.
`Bold monospaced` font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Keyboard conventions

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide…	What it means…
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
Enter	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.
type	Used to mean pressing one or more keys on the keyboard.
enter	Used to mean pressing one or more keys and then pressing the Enter key.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Caution

A caution contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

Introduction to Open Systems SnapVault 1

About this chapter

This chapter introduces you to the NetApp Open Systems SnapVault software and describes how to find more information about the software and the related technologies.

Topics in this chapter

This chapter contains the following topics:

- "About Open Systems SnapVault" on page 2
- "Open Systems SnapVault Features" on page 8

About Open Systems SnapVault

What Open Systems SnapVault is

Open Systems SnapVault is a disk-to-disk data protection solution that takes advantage of the NetApp SnapVault technology to protect data residing on the following platforms:

- AIX®
- HP-UX®
- IRIX®
- Red Hat® Linux
- Windows
- Solaris™

Note

For a list of currently supported versions of the above platforms, see "Requirements for primary storage systems" on page 14.

Components of Open Systems SnapVault environment

A typical Open Systems SnapVault environment has the following three components, as shown in the following illustration:

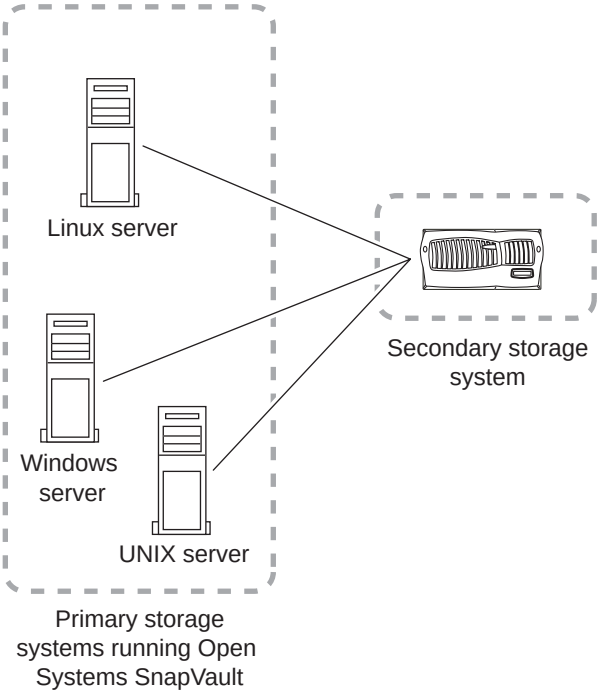

1. The primary system—the system from which you are going to back up data.

   For a list of currently supported primary systems, see "Requirements for primary storage systems" on page 14.

2. The Open Systems SnapVault agent—the software that you purchased from Network Appliance. This software is installed on the primary system.

3. The secondary system—the Network Appliance storage system to which you are going to back up data from the primary system.

   For information about supported secondary systems, see "Requirements for secondary storage systems" on page 16.

How Open Systems SnapVault works

You identify the directories or file system on the primary system from which to back up data and the qtrees on the secondary system to which the data will be backed up.

For the first backup, the secondary storage system requests an initial baseline transfer of the identified directories (file systems) from the primary system. This transfer establishes a SnapVault relationship between the Open Systems platform directories and the SnapVault secondary qtrees.

Subsequent transfers can either be initiated manually or be configured for automatic, unattended transfer (scheduled). You specify the schedules on the secondary system using the commands available in Data ONTAP or using an NDMP-based management application such as DataFabric Manager. Depending on the parameters configured for an Open Systems SnapVault environment, whole files or only the changed blocks are transferred in the subsequent transfers.

For each set of scheduled data transfers, SnapVault creates a set of incremental snapshots that capture the changes to the secondary qtrees after each transfer.

For each set of snapshots, the SnapVault secondary storage system saves the number of secondary storage snapshots you specify and assigns each snapshot in the set a version number (0 for most current, 1 for second most recent, and so on).

If directory or file data needs to be restored to the primary system, SnapVault retrieves the data from one of the retained snapshots and transfers the data back to the primary system that requests it.

For information about SnapVault and how it works to back up data, see the Data ONTAP *Data Protection Online Backup and Recovery Guide*.

Differences between Open Systems SnapVault and Data ONTAP backup and restore operations

Although Open Systems SnapVault is used in a manner similar to how Data ONTAP is used to back up and restore data from NetApp filers, the following differences exist:

- With Open Systems SnapVault you can back up and restore to any directory, subdirectory, or volume, whereas with Data ONTAP you can back up or restore only to qtrees using SnapVault.
- With Open Systems SnapVault, you can restore a single file.

Administration interfaces for Open Systems SnapVault

You can configure the Open Systems SnapVault parameters on your primary storage system in the following two ways:

- Using the `Configurator` utility graphical-user interface (GUI) for configuring Open Systems SnapVault, as shown in the following example.

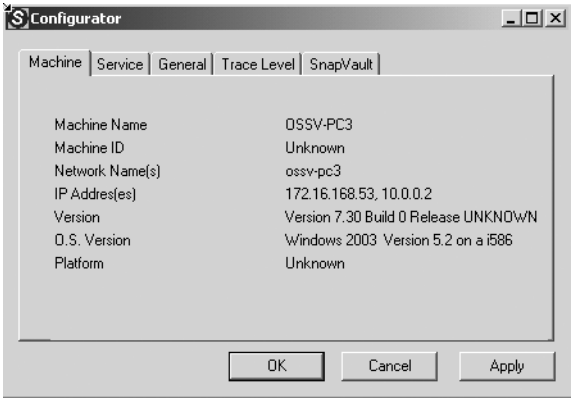

  For more information about the `Configurator` utility, see "Configuration interfaces" on page 45.

- Using the command-line utility

  The `svsetstanza` command is available to configure Open Systems SnapVault parameters.

  For more information about the `svsetstanza` command, see "Understanding the svsetstanza command" on page 47.

On the secondary storage system, you can configure SnapVault and start backups from the command line.

Sample commands

Sample backup command: To back up the C: drive of a Windows primary system using Open Systems SnapVault, from the command line of the SnapVault secondary system, enter the following command:

```
snapvault start -S winserver:C:\ /vol/volname/winserver_C
```

winserver is the name of the Windows primary system , C:\ is the Windows system's C: drive, and */vol/volname/winserver_C* is the path on the secondary system to the SnapVault qtree for the backup.

Sample restore command: To restore a subdirectory from the backup created in the `start` command above, enter the following command from a Windows prompt:

```
install_dir\bin\snapvault restore -S
secondary:/vol/volname/winserver_C/projfiles C:\restored
```

install_dir is the path to the `snapvault` command on the Windows systems. *secondary* is the name of the secondary system. For example, “C:\Program Files\netapp\snapvault\bin\snapvault.exe”.

Note

The example *install_dir* path is enclosed in double quotes (“ ”) because it includes spaces in the path name.

/vol/volname/winserver_C/ is the path of the qtree on the secondary system that stores the backed-up data. *projfiles* is the name of the subdirectory in the qtree to be restored. The subdirectory is restored to the Windows system directory C:\restored.

For an example of a single file restore, see “Using the snapvault restore command” on page 76.

Central management of Open Systems SnapVault agents

Open Systems SnapVault can be managed from a variety of management applications. These applications communicate with the Open Systems SnapVault clients and the NetApp storage systems over a TCP/IP network. Backup schedules, retention policies, backup control, and monitoring is centrally configured on these applications.

The applications that you can use to manage Open Systems SnapVault are as follows.

Vendor	Product	URL
Network Appliance™	DataFabric® Manager 2.2 or later	http://now.netapp.com
BakBone® Software	NetVault™ 7.x or later software with a SnapVault Manager plug-in module	http://www.bakbone.com
Syncsort	Backup Express 2.2 or later	http://www.syncsort.com
CommVault®	QiNetix™	http://www.commvault.com

The following illustration shows a typical Open Systems SnapVault setup using DataFabric Manager.

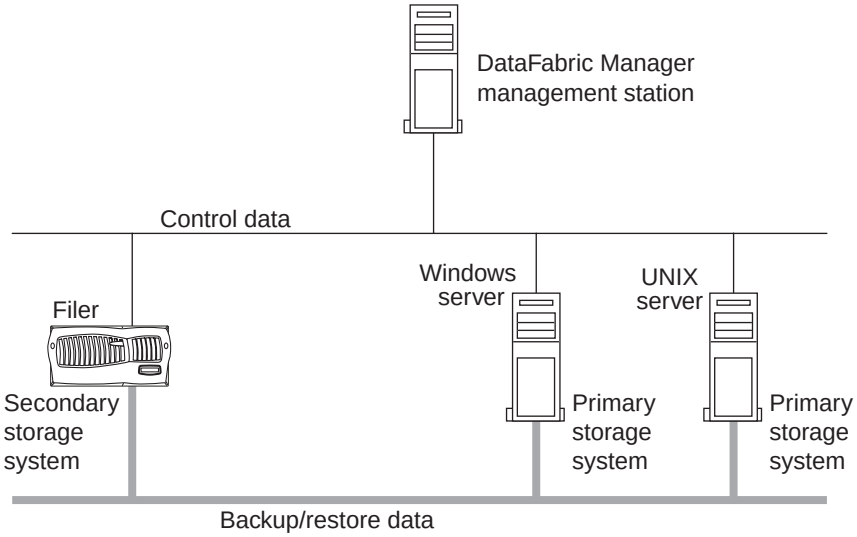

For more information

You can obtain additional information about Open Systems SnapVault and related technologies from the following:

- For SnapVault information and commands, see the following documents:
  - Data ONTAP *Data Protection Online Backup and Recovery Guide* at http://now.netapp.com.
  - Network Appliance Technical Report TR-3234—Leveraging Network Appliance™ SnapVault™ for Heterogeneous Environments, at http://www.netapp.com.
  - Network Appliance Technical Report TR-3252—Enabling Rapid Recovery with SnapVault, at http://www.netapp.com.
- For the Open Systems SnapVault Quick Reference Card or the Frequently Asked Questions (FAQs), see the Data ONTAP *Best Practices Guide* on the NetApp on the Web™ (NOW) site.

Open Systems SnapVault Features

Available features

The Open Systems SnapVault software provides many advanced features, a few of which are as follows:

- "Backing up open files"
- "Excluding specific files or directories from backup"
- "Setting block-level incremental transfer"
- "Backing up and restoring the Open Systems SnapVault database"
- "Backing up and restoring Windows System State"
- "Checkpoints for restart of transfers"
- "Encrypted File System (EFS) files"
- "Determining free space using space estimator"
- "Resynchronizing a broken relationship"

Backing up open files

If files on an Open Systems platform directory are open when a scheduled SnapVault copy takes place, those files are skipped for the current transfer. Changes to those files are backed up during the next scheduled transfer. However, if you have configured your system for open file backup, Open Systems SnapVault can back up open files using one of the following options:

- Open File Manager (OFM)—Used for Windows NT and Windows 2000 platforms

  OFM is automatically installed on the primary system at the same time the Open Systems SnapVault software is installed; however, you do need to have a license to use the OFM feature as described in "License requirements" on page 16.

- VSS snapshot—Used for Windows 2003 platforms

  The VSS snapshot functionality is integrated with the Open Systems SnapVault software as a standard feature and does not require a license to use.

For more information, see "Configuring open file backup for Windows" on page 65.

Excluding specific files or directories from backup

Backup exclusion lists are used by Open Systems SnapVault agents to exclude specified files and directories from backups. Open Systems SnapVault agents support two types of exclusion lists:

- File exclusion lists— A file or directory is excluded if the file name or any path element matches a file exclusion entry in the list.
- Path exclusion lists—If a path exclusion entry specifies a directory, that directory and its files and subdirectories are excluded.

For more information, see "Configuring backup exclusion lists" on page 62.

Setting block-level incremental transfer

A block-level incremental (BLI) backup recognizes that a file has changed based on a timestamp and checksum algorithm. It also determines exactly which blocks in the file have changed, and then backs up only those blocks to the Open Systems SnapVault secondary storage systems during backup transfers. Because only a small percentage of an application's data changes between periodic backups, incremental backups provide an efficient solution to protecting your data.

Note

BLI does not work on EFS-encrypted files on Windows systems.

Typically, incremental backups reduce the amount of time required to back up data, and minimize the resources required to perform backups when compared to baseline or full backups. This technique significantly reduces the amount of data that needs to be transferred to backup storage, as well as the amount of data that must be stored on backup storage disk.

Changed blocks are recognized based on checksum values calculated on 4 KB blocks of file data and stored in an internal database by the Open Systems SnapVault primary agent. This technique works well if an application modifies the file by appending changes to the end of the file. However, applications such as Microsoft Word, Microsoft Excel, and Microsoft PowerPoint (referred to as name-based applications) modify files by inserting new data blocks in the file and rewriting all subsequent data blocks in the file to new positions in the file. Because the modified file is considered new, a backup of all the rewritten blocks and a recalculation of checksum would be required. However, Open Systems SnapVault agents work around this issue by recognizing files by names in addition to identifying the file by the file-system location.

Unlike the corresponding secondary storage system option, this work around is enabled by default. To enable or disable the workaround, see "Enabling or disabling BLI backups for certain name-based applications" on page 61.

For information on how to configure or change BLI settings in Open Systems SnapVault, see "Setting block-level incremental backup options" on page 60.

Backing up and restoring the Open Systems SnapVault database

The Open Systems SnapVault database consists of a set of files that contain information about the Open Systems SnapVault relationship between a primary and a secondary system.

If the Open Systems SnapVault database becomes corrupt or gets out-of-sync with the secondary system, data transfers between the primary and secondary systems cannot continue. If you do not have a way to restore the database, you are forced to initiate a baseline transfer from the primary to the secondary system. However, if you maintain a backup copy of the database, you can restore the database for the relationship and continue with subsequent data transfers with minimal downtime and without the need to perform a baseline transfer.

For information about the Open Systems SnapVault database and how to back up and restore the database, see "Backing and restoring the Open Systems SnapVault database" on page 83.

Backing up and restoring Windows System State

You can back up and restore Windows System State data using Open Systems SnapVault. Doing so can be useful if, for example, an Active Directory entry is accidentally deleted. You can also use Open Systems SnapVault System State data backup in conjunction with complete file system backups as part of a disaster recovery plan.

For more information, see "Backing up and restoring Windows System State data" on page 88.

Checkpoints for restart of transfers

When an Open Systems SnapVault backup process fails, checkpoint restart support allows the backup transfer (baseline or an update) to be resumed from a known good point in the backup stream.

Checkpoints are recorded by the primary storage system when certain predetermined conditions or periodic intervals are met. Checkpoints are taken either at five-minute intervals or at the end of a file, whichever happens later. The primary storage system records the checkpoints and sends them to the secondary storage system.

By default, the number of times a failed transfer will be retried is set to two; however, you can change it. For more information, see "Changing the number of retry attempts made for failed transfers" on page 109.

Encrypted File System (EFS) files

You can back up and recover EFS encrypted files on Microsoft Windows systems; however, EFS encrypted files cannot be backed up using block-level incremental backup. Any time an EFS encrypted file is modified, Open Systems SnapVault will back up the entire EFS encrypted file.

Open Systems SnapVault is capable of restoring up to five EFS encrypted files concurrently.

For more information, see "Encrypted File System (EFS) encrypted file backup and restore" on page 110.

Determining free space using space estimator

The Open Systems SnapVault space estimator is a utility available with the Open Systems SnapVault product that enables you to ascertain if there is sufficient disk space available on the Open Systems SnapVault primary to perform a backup. If this utility is run on a system that does not have the Open Systems SnapVault product installed, it can also provide recommendations on where to install the product and its database and temporary files based on the currently available free space.

For more information, see "Open Systems SnapVault space estimator" on page 111.

Resynchronizing a broken relationship

You can resynchronize a broken or out-of-sync SnapVault relationship between a primary and a secondary and continue incremental data transfers as usual. Doing so eliminates the need to reinitialize the relationship, which involves a lengthy baseline transfer between the primary and the secondary systems.

For more information, see “Resynchronizing restored or broken relationships” on page 107.

Installing the Open Systems SnapVault Software 2

About this chapter

This chapter describes how to install and upgrade the Open Systems SnapVault agent on various platforms.

Topics in this chapter

This chapter discusses the following topics:

- "Before you install" on page 14
- "Installing Open Systems SnapVault on Windows platforms" on page 18
- "Installing Open Systems SnapVault on UNIX and Linux platforms" on page 24
- "Verifying the installation" on page 31
- "Upgrading the Open Systems SnapVault agent to version 2.2" on page 32
- "Uninstalling the Open Systems SnapVault 2.2 agent" on page 35
- "Unattended installation and upgrade" on page 37

Before you install

Before you install

Make sure that your site meets the minimum requirements for the primary and the secondary storage systems, and that you have the correct licenses for both systems before installing the Open Systems SnapVault agent. Also, make sure that you read the information in "Limitations" on page 16 and *Open Systems SnapVault 2.2 Release Notes* for any late-breaking information.

Requirements for primary storage systems

The requirements for the primary storage systems can be categorized as follows:

- Type of system
- Memory and port requirements
- Disk requirements
- File system requirements

Type of system: You can install the Open Systems SnapVault 2.2 agent on the following systems:

32-bit systems:

- Windows Server 2003 Standard Edition and Enterprise Edition on x86-compatible platforms.
- Windows 2000 Server and Advanced Server with Service Pack 4
- Windows NT 4 Server Standard Edition, and Enterprise Edition with Service Pack 6
- Solaris versions 2.6, 7, 8, and 9 on SPARC
- AIX 4.3 and AIX 5L versions 5.1 and 5.2 on PowerPC and POWER processor based systems

Note

The IBM Cluster 1600 and RS/6000 Scalable Parallel (SP) platforms are not supported.

- Red Hat® Linux versions 7.1, 7.2, 7.3, 8.0 and 9; Advanced Server 2.1, and Red Hat Enterprise Linux AS 3.0, ES 3.0, and WS 3.0
- SGI IRIX versions 6.5.x and later

64-bit systems:

- HP-UX versions 11 and 11i on PA-RISC

  The Open Systems SnapVault software running on 64-bit HP-UX systems is 32-bit.

Note

64-bit AMD and Intel processors are not supported by Open Systems SnapVault.

Memory and port requirements: Your primary system must have the following characteristics:

- A minimum of 512 MB of memory
- 100Base-T or Gigabit Ethernet (GbE) network connectivity is recommended for best performance
- TCP port 10566 (SnapVault) must be free
- (Optional) For NDMP-based management applications such as DataFabric Manager, TCP port 10000

  If port 10000 is already in use, you can choose another port using the `Configurator` utility during or after installation.

Storage (disk) requirements: You can use the Open Systems SnapVault space estimator utility, available starting with Open Systems SnapVault 2.2, to obtain recommendations on where to install the Open Systems SnapVault product, its database, and temporary files based on the currently available free space on your system. For more information about the space estimator utility, see "Open Systems SnapVault space estimator" on page 111.

In addition, you can use the following guidelines to determine the amount of free disk space your installation will require:

- Open Systems SnapVault requires temporary disk space for normal operations. Temporary space of 425 MB is needed for every two million files of 20 KB each during baseline transfer. This number goes up to 601 MB during update transfer if the same number of files are updated.
- If Open File Manager (OFM) is used, at least 15 percent free disk space in the volume being backed up.
- The Open Systems SnapVault built-in database requires dedicated storage on your primary storage system. The database disk space requirements depend on the number and average size of file and number of directories.

File system requirements: Open Systems SnapVault software supports the following file systems.

Platform	File systems	ACLs supported on
Windows	NTFS	NTFS
AIX	JFS1	JFS1
HP-UX	VxFS (JFS), HFS	HFS
IRIX	XFS	XFS
Linux	EXT2, EXT3	None
Solaris	UFS	UFS

Requirements for secondary storage systems

To use systems installed with Open Systems SnapVault software, the SnapVault secondary storage system must be running Data ONTAP 6.5.4 or later.

License requirements

The Open Systems SnapVault licenses are installed on the *secondary* system to which they are backed up. You will need to install the following licenses on the secondary system:

- The SnapVault secondary (`sv_ontap_sec`) license
- The SnapVault primary licenses for the platforms you want to back up to the secondary system. The following is a list of the primary licenses:
  - `sv_windows_pri`—For Windows systems
  - `sv_unix_pri`—For UNIX-based systems
  - `sv_linux_pri`—For Linux-based systems
- Optionally, you can install the `sv_windows_ofm_pri` license if you want to use OFM for Windows systems

To obtain the licenses for Open Systems SnapVault, contact your NetApp representative.

Limitations

Review the following limitations and the issues listed in the release notes of this Open Systems SnapVault release before you begin using it to back up data.

Open Systems SnapVault software *does not support* the following:

- Restore of UNIX sockets
- Backup and restore of any primary storage system quota database
- Microsoft 64-bit Windows
- The following operations without root access on UNIX or administrator privileges on Windows primary storage systems:
  - Installing and configuring the Open Systems SnapVault agent
  - Data restoration
- FAT and HPFS file systems on Windows primary storage systems
- Remote NFS or CIFS file systems that have been mounted on or mapped to UNIX or Windows primary storage systems
- NFS v4 access control list (ACL) information functionality
- Any clustered systems
- Multiple file systems in a single backup transfer

  Open Systems SnapVault does not cross mount points. You must specify each local mount point as a separate backup if you are backing up multiple file systems. On Windows systems this means mount points that are not mapped to drive letters cannot be backed up.
- Scripts created to back up databases such as Oracle, SQL, or Sybase

  However, Open Systems SnapVault can back up file systems that contain these databases if the database is in the proper backup mode.

Installing Open Systems SnapVault on Windows platforms

Installing a Windows agent from CD-ROM

To install the Open Systems SnapVault Windows NT, Windows 2000, or Windows 2003 agent on your supported Windows primary storage system from CD-ROM, complete the following steps.

Note

For Windows NT and Windows 2000, St. Bernard's OFM module requires that you reboot the system after installation. You will be unable to use OFM module's Open File Backup functionality until you reboot the system.

Step	Action
1	Log in to your Windows primary storage system with Administrator privileges.
2	Insert the CD-ROM labeled *Open Systems SnapVault Agent* into the primary storage system's CD-ROM drive.
3	From the command line or Windows Explorer, navigate to the CD-ROM directory specific to your platform: ◆ For Windows NT, go to \NT ◆ For Windows 2000, go to \WIN2K ◆ For Windows 2003, go to \WIN2K3 Depending on your system configuration, the files and directories mentioned might appear in lowercase characters. To navigate to the correct directory, use uppercase characters for the directory name (for example, `cd \WIN2K`).
4	After navigating to the appropriate directory for your platform, see "Using the Installation Wizard" to complete the installation.

Installing a Windows agent from NOW

To download the Open Systems SnapVault Windows agent module from the NOW™ site and install the agent on your supported Windows platform, complete the following steps.

Step	Action
1	Log in to your primary storage system with Administrator privileges.
2	Go to the Download Software page of the NOW site at http://now.netapp.com/NOW/cgi-bin/software/ and follow the directions to download the appropriate Open Systems SnapVault package for your platform.
3	Unzip the downloaded package into a temporary directory on your Windows platform.
4	Navigate to the temporary directory where you unzipped the files for your system.
5	Follow the instructions in "Using the Installation Wizard" to complete the installation.

Using the Installation Wizard

To use the Installation Wizard to install the Open Systems SnapVault agent on your system, complete the following steps.

Step	Action
1	Locate and double-click the setup.exe file. **Result:** The Open Systems SnapVault Setup Wizard is launched.
2	Follow the instructions on the screen. Click the "I Agree" radio button to accept the license agreement and click Next. **Result:** The Choose NDMP Username/Password window appears. This username and password is used to communicate with an NDMP-based application such as DataFabric Manager that is used for the central management of Open Systems SnapVault agents.

Step	Action
3	**Optional:** Enter your user name and password, retype your password to confirm it, and click Next. **Result:** The NDMP Listen Port window appears. This port number is used to communicate with an NDMP-based application such as DataFabric Manager that is used for the central management of Open Systems SnapVault agents. Or, simply click Next.
4	Enter the NDMP listening port in the NDMP Listen Port window when prompted, or accept the default listening port setting of 10000, and click Next. **Note** If another application uses 10000 as its listening port, choose a port number greater than 10000 that is currently not being used. 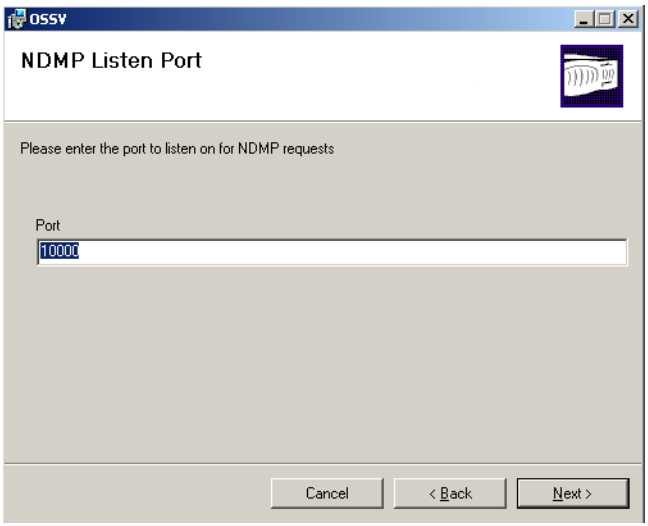

Step	Action
5	When the “Allowed Secondary Names” window appears and you are prompted for allowed secondary storage system names, enter one or more host names or IP addresses of the SnapVault secondary storage system or systems that you want to back up this primary storage system. 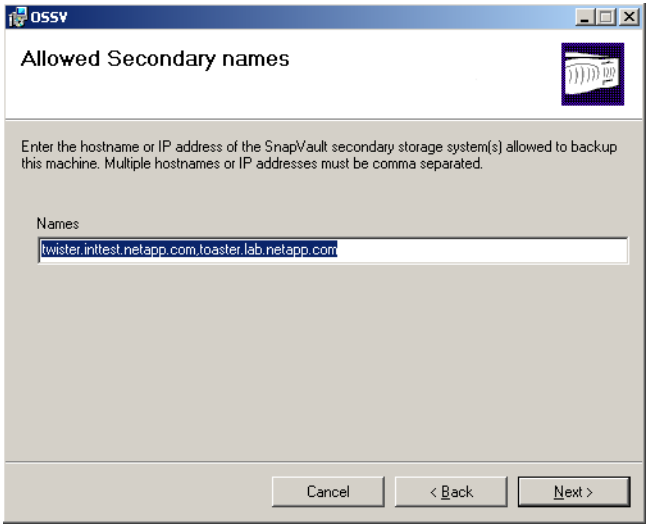 **Note** If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas. Only the secondary storage systems named in this field are accepted by the primary system as valid backup systems.

Step	Action	
6	When the Select Installation Folder window prompts you for the installation directory, click Next to accept the default location, or enter your own path to a different folder and click Next. 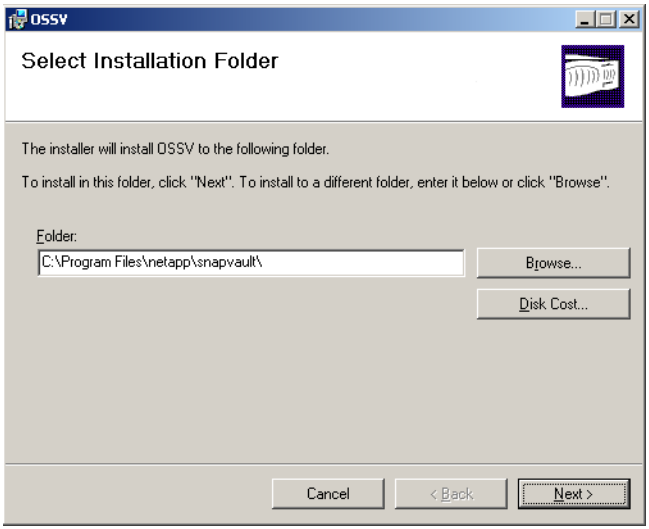 **Note** If the Open Systems SnapVault agent was previously installed and then uninstalled, files might still be resident in the installation directory and a dialog box appears. To empty the directory, click Yes in the dialog box before you click Next. Click Disk Cost to view the available disk space. **Result:** The Confirm Installation window appears to inform you that the installer is ready to install Open Systems SnapVault on your computer.	
7	Click Next to start the installation process and wait until the Installation Complete window appears, then click Close.	
8	**For Windows NT or 2000...**	**For Windows 2003...**
	Reboot the system when prompted if you plan to use OFM; otherwise, go to Step 9.	Go to step 9.

Step	Action
9	After you finish installing the Open Systems SnapVault agent, follow the procedures in the SnapVault chapter of the Data ONTAP *Data Protection Guide* to configure the SnapVault secondary storage system for Open Systems backup.

Installing Open Systems SnapVault on UNIX and Linux platforms

Before you proceed with the installation

Make sure you have read the "Requirements for primary storage systems" on page 14 to ensure that you are installing the Open Systems SnapVault software on one of the *supported* UNIX or Linux platforms.

Installing the Solaris agent from CD-ROM

To install the Open Systems SnapVault Solaris agent on your supported Solaris primary storage system from CD-ROM, complete the following steps.

Step	Action
1	Log in to your primary storage system as root.
2	Insert the CD-ROM labeled *Open Systems SnapVault Agent* into the primary storage system's CD-ROM drive.
3	Navigate to the Solaris pkgadd package on the CD-ROM. The Solaris pkgadd utility transfers the contents of a software package from the distribution medium or directory to install it onto the system.
4	Run the Solaris pkgadd utility using the following command: **`pkgadd -d`** ***`path_to_package`***`/ossv ossv` *path_to_package* is the full path to where the Open Systems SnapVault package is located; for example, /export/home/packages/ossv ossv. **Result:** The package installation script is initiated and asks you a series of questions prior to and during the installation process.
5	Go to "Using the Solaris pkgadd utility to install the Solaris agent" to continue the Solaris package installation.

Installing the Solaris agent from NOW

To download the Open Systems SnapVault Solaris installation package from the NOW site and install it on a primary storage system, complete the following steps.

Step	Action
1	Log in to your primary storage system as root.
2	Go to the Download Software page of the NOW site at http://now.netapp.com/NOW/cgi-bin/software/ and follow the directions to download the Open Systems SnapVault Solaris agent package.
3	After you download the Open Systems SnapVault package, unpack and start the installation program as follows: ◆ Uncompress and untar the downloaded pkgadd package. ◆ Run the Solaris pkgadd utility using the following command: **`pkgadd -d`** ***`path_to_package`***`/ossv ossv` *path_to_package* is the full path to where the Open Systems SnapVault package is located; for example, /export/home/packages/ossv ossv. The Solaris pkgadd utility transfers the contents of a software package from the distribution medium or directory to install it onto the system. **Result:** The package installation script starts and asks you a series of questions prior to and during the installation process.
4	Go to "Using the Solaris pkgadd utility to install the Solaris agent" to continue the Solaris package installation.

Using the Solaris pkgadd utility to install the Solaris agent

To install the Open Systems SnapVault Solaris agent using the pkgadd utility, complete the following steps.

Step	Action
1	Answer yes (y), no (n), or display (d) when asked if you read and agree to the terms of the license. If you answer no (n), you cannot install the package.

Step	Action
2	When prompted, enter the path where you want the SnapVault directory to be created. The default directory is /usr/snapvault. If the directory you entered already exists, you are warned that the current contents of that directory will be destroyed if you continue and then prompted whether you want to continue. Enter yes (y) or no (n).
3	Enter your user name to connect to the target system using the NDMP protocol. This username and password is used to communicate with an NDMP-based application such as DataFabric Manager that is used for the central management of Open Systems SnapVault agents.
4	Enter and confirm the password to connect to the system.
5	Enter the NDMP listener port. The default listener port is 10000. This port number is used to communicate with an NDMP-based application such as DataFabric Manager that is used for the central management of Open Systems SnapVault agents. **Note** If another application uses 10000 as its listening port, choose a port number greater than 10000 that is currently not being used.
6	Enter the host names or IP addresses of the SnapVault secondary storage systems that are allowed to perform backups from this primary storage systems. **Note** If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas. Only the secondary storage systems named in this field are accepted by the primary system as valid backup systems.
7	A series of installation scripts is executed. If successful, the following message appears: `Installation of OSSV was successful.`

Installing the HP-UX, AIX, IRIX, or Linux agent from CD-ROM

To install the Open Systems SnapVault HP-UX, AIX, IRIX, or Linux agent on your supported UNIX primary storage system from CD-ROM, complete the following steps.

Step	Action
1	Log in to your primary storage system as root.
2	Insert the CD-ROM labeled *Open Systems SnapVault Agent* into the primary storage system's CD-ROM drive.
3	From the command line, navigate to the CD-ROM directory specific to your platform: ◆ For HP-UX, go to /HPUX ◆ For AIX, go to /AIX ◆ For IRIX, go to /IRIX ◆ For Linux, go to /LINUX Depending on your system configuration and mount options, the files and directories mentioned might appear in lowercase characters.
4	From the command line, enter the following command using all uppercase characters: `./install`
5	When the `install` program prompts for the installation directory, press Enter to accept the default, or enter your own path. The default value, /usr, installs the Open Systems SnapVault software in /usr/snapvault.
6	**If you plan to manage SnapVault backup through...** / **Then...** Your Data ONTAP command line / Press Enter. A commercial NDMP application, such as DataFabric Manager / When the `install` program prompts you for the user name, password, and NDMP listening port, specify a user name and password authorized through that application.

Step	Action
7	When the `install` program prompts you for Allowed Systems, enter one or more host names or IP addresses of the SnapVault secondary storage system or systems to which you want to back up this primary storage system. **Note** If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas but no spaces.
8	Look for a message similar to the following, which is displayed to indicate a successful installation: `Installation completed successfully.`
9	After you finish installing the Open Systems SnapVault agent, follow the procedures in the SnapVault chapter of the Data ONTAP *Data Protection Guide* to configure the SnapVault secondary storage system for Open Systems SnapVault backup.

Installing the HP-UX, AIX, IRIX, or Linux agent from NOW

To download the Open Systems SnapVault HP-UX, AIX, IRIX, or Linux agent from the NOW site and install it, complete the following steps.

Step	Action
1	Log in to your primary storage system as root.
2	Go to the Download Software page of the NOW site at http://now.netapp.com/NOW/cgi-bin/software/ and follow the directions to download the appropriate Open Systems SnapVault package.

Step	Action	
3	**If...**	**Then...**
	You are installing Open Systems SnapVault on a Linux platform	Enter the following commands after downloading the Open Systems SnapVault package: **`gunzip`** ***`package_name`*** **`tar -xvf`** ***`tar_file_name`***

	You are installing Open Systems SnapVault on any other UNIX platform	Enter the following commands after downloading the Open Systems SnapVault package:
		`uncompress` ***`package_name`***
		`tar -xvf` ***`tar_file_name`***
		For example, on an AIX primary storage system, uncompress and untar the downloaded ossv_aix_v2.2.tar.Z package using the following commands:
		`uncompress ossv_aix_v2.2.tar.Z`
		`tar -xvf ossv_aix_v2.2.tar`

4	Navigate to the directory where the `untar` program placed the files and enter the following command: **`./install`**	
5	When the `install` program prompts for the installation directory, press Enter to accept the default, or enter your own path. The default value, /usr, installs the Open Systems SnapVault software in /usr/snapvault.	

Step	Action	
6	**If you plan to manage SnapVault backup through...**	**Then...**
	Your Data ONTAP command line	Press Enter.
	A commercial NDMP application such as DataFabric Manager	When the `install` program prompts you for the user name, password, and NDMP listening port, specify a user name and password authorized through that application.
7	When the `install` program prompts you for Allowed Systems, enter one or more host names or IP addresses of the SnapVault secondary storage system or systems that you want to back up this primary storage system. **Note** If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas but no spaces.	
8	Look for a message similar to the following, which is displayed to indicate a successful installation: `Installation completed successfully.`	
9	After you finish installing the Open Systems SnapVault agent, follow the procedures in the SnapVault chapter of the Data ONTAP *Data Protection Guide* to configure the SnapVault secondary storage system for Open Systems backup.	

Verifying the installation

Verifying the installation

To ensure that the installation was completed correctly and that the primary and secondary storage systems can back up data, complete the following steps.

Step	Action
1	Verify connectivity to the secondary storage system using the following command: **`ping secondary_system`** *secondary_system* is either the name or the IP address of the secondary storage system.
2	If security was enabled during installation, ensure that the secondary storage system specified is accessible from the primary storage systems, using the procedure described in "Enabling and disabling security" on page 53.
3	Ensure that the volumes you need were created on the secondary storage system before you attempt to back up data to them. **Note** The volumes should not have qtrees already created.
4	If there are firewalls between the primary storage system and the secondary storage system, ensure that TCP port 10000 and 10566 are allowed using the `telnet` command. TCP port 10566 is used by the SnapVault process and TCP port 10000 is used by NDMP applications.
5	Navigate to *install_dir*/bin and run `svinstallcheck` to verify successful installation and to make sure that the services are running properly. *install_dir* is the location where you installed the Open Systems SnapVault agent.

Upgrading the Open Systems SnapVault agent to version 2.2

Prerequisites

You must understand the following prerequisites before upgrading to Open Systems SnapVault agent 2.2:

- You cannot upgrade from a beta version to a release version. If you have a beta version installed, remove it following the appropriate procedure in the section "Uninstalling the Open Systems SnapVault 2.2 agent" on page 35.
- If an Open Systems SnapVault 1.2 agent is upgraded to 2.2, the Enable Unix ACLs tab appears in the `Configurator` utility. You can safely ignore this tab as it has no effect on agent functionality. When upgrading to an Open Systems SnapVault 2.2 agent from either a 1.0 or 1.1 agent, this tab does not appear in the `Configurator` utility.
- If you are upgrading from Open Systems SnapVault 2.0 or Open Systems SnapVault 2.1 and have been using BLI updates, or if you do not know whether you have been using BLI updates, review NetApp Bug 137685 for information about a problem in some versions of Data ONTAP that can cause loss of data in Open Systems SnapVault 2.0 BLI updates. Also, see Knowledge Base article KB4207 for information on a workaround for this problem with Open Systems SnapVault 2.1.

Upgrading the Open Systems SnapVault agent to version 2.2

When upgrading the Open Systems SnapVault agent from version 1.x to version 2.2: Perform the following steps.

Step	Action
1	**1.** Launch the SnapVault `Configurator` utility. See "Running the Configurator utility" on page 50 for instructions. **2.** Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services. **3.** Close the `Configurator` utility. **Note** You can use the `svpmgr shutdown` command as an alternative.

Step	Action
2	**Important:** To ensure that your Open Systems SnapVault database is not accidentally removed during installation, move the database directory to a new location or back it up as discussed in "Backing and restoring the Open Systems SnapVault database" on page 83.
3	Using the command line interface of the Open Systems SnapVault primary system, navigate to the *snapvault/util* directory on your system. ❖ On Windows systems, the default path is: C:\Program Files\snapvault\util ❖ On UNIX systems, the default path is: /usr/snapvault/util
4	Enter the following command: `uninstall`
5	After the `uninstall` command is complete, remove the *snapvault* directory and its contents.
6	Install the Open Systems SnapVault 2.2 agent using the procedure described in "Installing Open Systems SnapVault on Windows platforms" on page 18 or "Installing Open Systems SnapVault on UNIX and Linux platforms" on page 24. **Result:** The agent automatically starts Open Systems SnapVault services after installation.
7	**1.** Launch the SnapVault `Configurator` utility. **2.** Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services. **3.** Close the `Configurator` utility. **Note** You can use the `svpmgr shutdown` command as an alternative.
8	Copy the database directory that you saved in Step 2 to the location of the database created by the new installation.

Step	Action
9	1. Launch the SnapVault `Configurator` utility. 2. Click the Service tab, and then click Start Service to start Open Systems SnapVault services. **Note** You can use the `svpmgr startup` command as an alternative.

When upgrading from any other version of Open Systems SnapVault except version 1.x to version 2.2: Perform the following step.

Step	Action
1	Install the Open Systems SnapVault 2.2 agent using the procedure described in "Installing Open Systems SnapVault on Windows platforms" on page 18 or "Installing Open Systems SnapVault on UNIX and Linux platforms" on page 24. **Result:** The agent automatically starts Open Systems SnapVault services after installation.

Uninstalling the Open Systems SnapVault 2.2 agent

Uninstalling the 2.2 agent on Windows

To uninstall the Open Systems SnapVault 2.2 agent on the Windows platform, complete the following steps.

Step	Action
1	Click Start > Control Panel.
2	Double-click “Add or Remove Programs”.
3	Choose the Open Systems SnapVault agent in the list of programs and click Remove. **Result:** Windows uninstalls the Open Systems SnapVault agent and responds with a message to inform you that the uninstall is successful.

Uninstalling the 2.2 agent on Solaris

To uninstall the Open Systems SnapVault 2.2 agent on a Solaris platform, complete the following steps.

Step	Action
1	Log in to your primary storage system as root.
2	Run the pkgrm utility: `pkgrm ossv`

3	Answer yes (y) when asked whether you want to remove this package. The script responds with the following text:
	`## Removing installed package instance <ossv>`
	`This package contains scripts which will be executed with super-user permission during the process of removing this package.`

Step	Action
4	Answer yes (y) when asked whether you want to continue with the removal of this package. **Result:** If removal is successful, the script responds with the following text: `Removal of <ossv> was successful.`

Uninstalling 2.2 agent on HP-UX, AIX, IRIX, or Linux

To uninstall the Open Systems SnapVault 2.2 agent on the HP-UX, AIX, IRIX, or Linux platform, complete the following steps.

Step	Action
1	Using the command-line interface of the Open Systems SnapVault primary storage system, navigate to the location of the agent (the default location is /usr/snapvault/util).
2	Enter the following command: **`uninstall`**
3	If the snapvault directory still appears after running `uninstall`, remove the directory manually.

Unattended installation and upgrade

What unattended installation and upgrade is

Unattended installation and upgrade method enables you to install or upgrade Open Systems SnapVault software on a primary storage system with minimal user intervention. This technique is most useful for environments with large number of primary storage systems. By using this technique, you can set installation variables non-interactively, and in most cases, you do not need to reboot the system after the installation or upgrade has completed successfully.

Note

The unattended installation method does not provide batch installation of several clients at the same time. However, you can perform remote batch installation on Windows clients as described in "Remote batch installation of the Open Systems SnapVault agent on Windows" on page 41.

Supported platforms and Open Systems SnapVault versions

All platforms on which Open Systems SnapVault is currently supported can be upgraded or installed using this technique. See "Requirements for primary storage systems" on page 14 for information on supported platforms.

Unattended installation is only supported on systems running Open Systems SnapVault 2.x.

Process of unattended installation

To perform an unattended installation of Open Systems SnapVault on a primary system, an installation script and other supporting files are required. A utility called `svconfigpackager` is available in the Open Systems SnapVault software. When run on a primary system running Open Systems SnapVault, the utility saves the current configuration settings to a file. In addition, this utility can create an installation script that in conjunction with the configuration settings file and other files can be used to perform unattended installations or upgrades.

Guidelines to follow

You must understand the following guidelines before proceeding with the unattended installation and upgrade procedure:

- Installation script and other files created by the `svconfigpackager` utility on an operating system cannot be used for running an unattended installation on a different operating system—that is, if you created an installation script on a

Windows 2003 system, you cannot use it to perform an unattended installation on a Windows 2000 system. Similarly, an installation script created for a Solaris system cannot be used to perform an unattended installation on an HP-UX system.

- The following configuration settings cannot be changed when performing an unattended upgrade:
  - Installation path
  - Database directory
  - Trace directory
  - Temporary directory
- Open Systems SnapVault 1.x agents cannot be upgraded using the unattended installation technique.

Preparing for an unattended installation or an upgrade

You must generate an installation script and other files necessary to perform the unattended installation or an upgrade. To generate the installation script and the files, complete the following steps.

Step	Action
1	On a primary system whose configuration settings you want to use for other installations, do the following: **1.** Launch the SnapVault `Configurator` utility. See "Running the Configurator utility" on page 50 for instructions. **2.** Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services.
2	Configure all parameters you want your new installation to have using the `Configurator` utility.
3	Close the `Configurator` utility.
4	Navigate to the *install_dir*/bin directory of your primary system.

Step	Action	
5	**If...**	**Then...**
	You only want to save the configuration settings to a file	Enter the following command: **`svconfigpackager filename`** *filename* is the name of the configuration settings file. **Example:** To create a configuration settings file called svconfig.in, enter the following command: `svconfigpackager svconfig.in`
	You want to save the configuration settings to a file and also create an installation script for new installs	Enter the following command: **`svconfigpackager -i path_name filename`** *path_name* is the directory where the configuration settings file and the installation script will be placed. *filename* is the name of the configuration settings file. **Example:** To create a configuration settings file called svconfig.in and save the installation script in the /usr/snapvault directory, enter the following command: `svconfigpackager -i /usr/snapvault svconfig.in`
	You want to save the configuration settings to a file and create an installation script that will not overwrite the existing configuration values on a system (as in case of an upgrade)	Enter the following command: **`svconfigpackager -h -i path_name filename`** *path_name* is the directory where the configuration settings file and the installation script will be placed. *filename* is the name of the configuration settings file.

Step	Action

6	Once the files in the previous step have been created, you will see a message at the command prompt that lists all the files that have been placed in the directory you specified.
	
	Example: The following is an example of what you will see on a Solaris primary system:
	`The following files have been placed in ‘/usr/snapvault’:`
	`'mypackage.in' (Configuration Package)`
	`'unattinstall.sh' (Unattended install shell script)`
	`'InstallResponseFile' (Unattended install response file)`
	`'InstallAdminFile' (Solaris ‘pkgadd’ Admin file)`
	
	Example 2: The following is an example of what you will see on a Windows primary system:
	`The following files have been placed in 'C:/Program Files/netapp/snapvault':`
	`'mypackage.in' (Configuration Package)`
	`'unattinstall.bat'  (Unattended install batch file)`
	
	Result: You are done preparing the installation script and the configuration files you will require to perform an unattended installation.

Performing an unattended installation or upgrade

To perform an unattended installation or upgrade on a system for which you have generated an installation script and other files necessary, complete the following steps.

Step	Action
1	Download and uncompress the Open Systems SnapVault package from the NOW site or copy it from the CD-ROM to the installation directory on the primary system that you want to install or upgrade.
2	Copy the installation script and other files that were generated in Step 6 of “Preparing for an unattended installation or an upgrade” on page 38 to the directory on the primary system where you uncompressed the installation package.

Step	Action	
3	**If...**	**Then...**
	Your primary system is a Windows machine	Enter the following command to start the unattended install: **`unattinstall.bat`** If the installation does not succeed, the error messages are logged at the following location: ◆ For Windows 2000 and 2003 systems—%SystemRoot%\Documents and Settings*Current User*\Local Settings\Temp ◆ For Windows NT systems—%SystemRoot%\WINNT\TEMP
	Your primary system is a Unix machine	Enter the following command to start the unattended install: **`./unattinstall.sh`** If the installation does not succeed, the error messages are logged in the /tmp directory.
4	A utility `svinstallcheck` is automatically run after the installation is complete to verify successful installation and to make sure that the services are running properly. If any errors are found, they are logged in the snapvault log files in *install_dir*/etc directory.	

Remote batch installation of the Open Systems SnapVault agent on Windows

Network Appliance does not provide a method to batch install Open Systems SnapVault agents. However, it is possible to remotely batch install the Open Systems SnapVault agent onto Windows clients.

The remote batch installation method is based on the Windows domain and active directory. In active directory, you can establish a policy to push the Open Systems SnapVault agent onto a number of clients within the domain. When those clients are rebooted, the agent is installed.

For information about remote batch installation in the Windows environment, see http://www.microsoft.com/.

Configuring Open Systems SnapVault 3

About this chapter

This chapter describes how to modify Open Systems SnapVault parameter settings. You can use either the graphical-user interface (GUI) utility called `Configurator` or the `svsetstanza` command on the command-line interface (CLI) on the system on which you installed the Open Systems SnapVault agent.

Topics in this chapter

This chapter describes the following procedures that you can perform using the SnapVault `Configurator` utility:

- "Configuration interfaces" on page 45
- "Running the Configurator utility" on page 50
- "Confirming that services are running" on page 52
- "Modifying SnapVault parameters" on page 53
- "Generating and stopping files for debugging" on page 56
- "Setting block-level incremental backup options" on page 60
- "Configuring backup exclusion lists" on page 62
- "Configuring open file backup for Windows" on page 65

Configuration interfaces

Available configuration interfaces

You can configure or modify Open Systems SnapVault parameters using either of the following methods:

- The `Configurator` utility—a GUI-based interface
- The `svsetstanza` command—a command-line interface

Topics in this section

The following topics are covered in this section:

- "Understanding the Configurator utility interface" on page 45
- "Understanding the svsetstanza command" on page 47

Understanding the Configurator utility interface

About the Configurator utility interface

The `Configurator` utility interface is the graphical user interface used from the Open Systems Snapvault agent on the primary storage system to configure and manage Open Systems SnapVault environment options.

This section describes all the tabs available in the `Configurator` utility GUI and their purpose.

Components of the Configurator utility GUI

The `Configurator` utility GUI consists of five tabs, as shown in the following example:

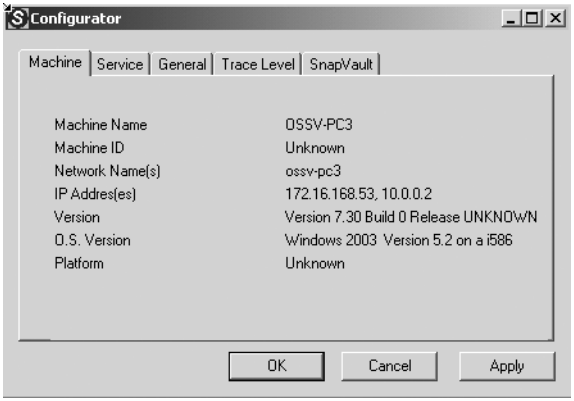

The Machine tab: The Machine tab displays information about the version of Open Systems SnapVault software, the primary storage system machine information such as the IP address and OS version.

The Service tab: The Service tab enables you to start and stop the Open Systems SnapVault service.

The General tab: The General tab enables you to generate debugging files by first selecting “Generate debugging files”, then modifying the default log output settings for various Open Systems SnapVault processes in the Trace Level tab.

You can also modify default directory locations using this tab. The SnapVault log files are located in the *install_dir*/etc directory. For more information about the SnapVault log files, see “Locating status and problem reports” on page 82.

The Trace Level tab: The Trace Level tab enables you to modify the default logging output for the various Open Systems SnapVault processes.

The SnapVault tab: The SnapVault tab enables you to modify multiple parameters such as block-level increment level (BLI), Open File Manager (OFM) parameters), NDMP parameters (for central management of Open Systems SnapVault agents), VSS parameters, and security settings.

Understanding the svsetstanza command

About the svsetstanza command

The `svsetstanza` command is a command-line utility that enables you to configure or modify Open Systems SnapVault parameters. You can use this utility for configuration purposes instead of using the GUI-based `Configurator` utility. The `svsetstanza` command is located in the *install_dir*/snapvault/util directory.

What svsetstanza command changes

The parameters that the `svsetstanza` command can configure or change are in various configuration files located in the *install_dir*/snapvault/config directory.

The following is an example of the configure.cfg file.

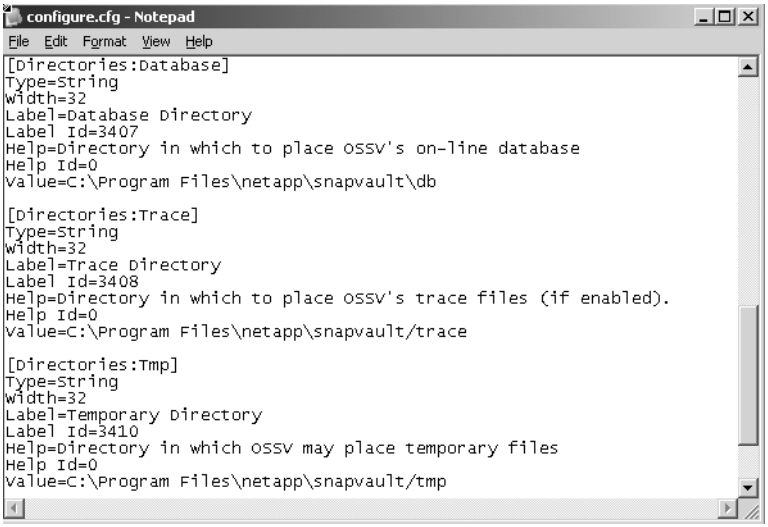

Syntax of the svsetstanza command

The following is the syntax of the `svsetstanza` command:

```
svsetstanza directory file category title keyword value
{asvaluelist=TRUE | FALSE} [replaced_value]
```

directory is the Open Systems SnapVault directory that contains the file in which the value to be changed is present. In most cases, the directory will be *install_dir*/snapvault/config.

file is the Open Systems SnapVault configuration file to be changed. For example, in the previous example, configure.cfg is the file.

category is the section to be changed in the configuration file. For example, in the previous example, Directories is a category.

title is the title of the section to be changed. For example, in the previous example, Trace is a title.

keyword is the parameter to be changed. For example, in the previous example, Value is a keyword.

value is the new value for the parameter.

`asvaluelist` specifies whether the value of the parameter to be changed is a list. Use TRUE if the value is a list, otherwise use FALSE.

replaced_value specifies the value in the list to replace. For example, if a QSM Access List specifies “f840, f880”, it can be changed to “f840, f740”

Notes about variable values: None of the values you specify for the variables in the `svsetstanza` command are case-sensitive. The values with spaces in between must be enclosed in double quotes (“ ”).

Examples

Example1: The trace directory of an Open Systems SnapVault installation needs to be changed from C:\Program Files\netapp\snapvault\trace to D:\Trace.

```
[Directories:Trace]
Type=String
Width=32
Label=Trace Directory
Label Id=3408
Help=Directory in which to place OSSV's trace files (if enabled).
Help Id=0
Value=C:\Program Files\netapp\snapvault\trace
```

Use the following command to accomplish the change:

```
svsetstanza config configure.cfg Directories Trace Value D:\Trace
FALSE
```

Example 2: To turn off the BLI settings, enter the following command:

```
svsetstanza config snapvault.cfg Configuration Checksums Value OFF
FALSE
```

Example 3: To replace f880 with f740 in the QSM Access List, enter the following command:

```
svsetstanza.exe config snapvault.cfg QSM “Access List” Value f740
asvaluelist=TRUE f880
```

Running the Configurator utility

Running the Configurator utility

To run the `Configurator` utility, complete the following steps.

Step	Action
1	Launch the `Configurator` utility using one of the following methods: ◆ Click Start > Programs > OSSV > OSSV Configurator on your Windows primary system. ◆ Run the following command on your UNIX primary system: **`$INSTALL_DIR/bin/svconfigurator`** The INSTALL_DIR is usually /usr/snapvault.

Step	Action
2	Click the appropriate `Configurator` utility tab to change the settings. The most commonly modified Open Systems SnapVault parameters are accessed by clicking the SnapVault tab. 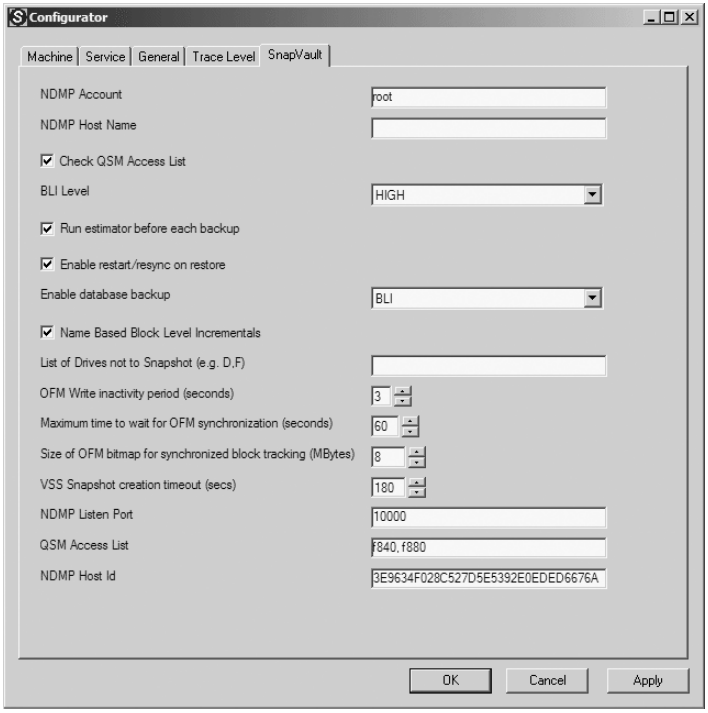

Confirming that services are running

Confirming that services are running

To confirm that Open Systems SnapVault services are running, complete the following steps.

Step	Action
1	Click the Service tab.
2	Verify that Current State Running is shown.
3	If services are not running, click Start Service.
4	Click OK.

Modifying SnapVault parameters

Parameters you can modify

You can modify the basic Open Systems SnapVault parameters from the SnapVault tab of the `Configurator` utility. This section describes how to perform the following tasks:

- Enabling and disabling security
- Modifying the QSM access list
- Modifying the NDMP settings

Enabling and disabling security

As a security measure, Open Systems SnapVault uses an access list to determine the secondary storage systems to which the primary storage system has permission to back up data. To enable or disable security, complete the following steps.

Step	Action
1	Click the SnapVault tab.
2	Click the Check QSM Access List box to enable or disable security.
3	Click OK.

Modifying the QSM access list

You can change the secondary storage systems to which the primary storage system backs up data by modifying the Qtree SnapMirror™ (QSM) access list. To modify the QSM access list, complete the following steps.

Step	Action
1	Click the SnapVault tab.
2	Add to, replace, or delete IP addresses or host names of the secondary storage systems to which you want to back up data.
3	Click OK.

Modifying the NDMP settings

You can modify the following NDMP settings:

- NDMP Listen Port

  Indicates the port on which the primary storage system listens for NDMP connections. By default, this port is set to 10000.

- NDMP Account

  Indicates the value used for NDMP authentication if NDMP-based management applications are used to manage the Open Systems SnapVault agent.

The following settings can be modified, but Network Appliance recommends against changing the default values unless asked by NetApp technical support:

- NDMP Host Name

  Indicates the name of the server (on which the NDMP-based management application exists) to which the Open Systems SnapVault agent connects. This field is set internally by the Open Systems SnapVault agent.

- NDMP Host ID

  Indicates a unique identifier that is automatically filled by the Open Systems SnapVault agent. This identifier is used by NDMP-based management applications to identify the primary storage system. This field must *not* be modified.

Modifying the NDMP Account setting: To modify the NDMP Account setting, complete the following steps.

Step	Action
1	Click the SnapVault tab.
2	Modify the NDMP Account in the window.
3	Click OK. **Note** To change the password, use the command line. Navigate to the *install_dir*\bin directory and run the `svpassword` command to change the password.

Modifying NDMP Listen Port setting: To reassign the NDMP Listen Port setting to another unused port number, complete the following steps.

Step	Action
1	Click the SnapVault tab.
2	In the NDMP Listen Port box, enter an unused TCP port number.
3	Click OK.
4	Stop and restart services after assigning the NDMP listen port.

Generating and stopping files for debugging

About generating debugging files

You can generate debugging files to help troubleshoot problems. However, you should practice the following cautions when generating these files:

- Because the debugging files grow large quickly and there is a performance impact, you should only generate these files when advised by NetApp Technical Support.
- Because the debugging files can be very large in size, make sure you delete the debugging files from your system after you have sent them to NetApp Technical Support.
- Make sure you turn off the generation of these files after you have sent the files to NetApp Technical Support.

Generating debugging files

To generate debugging files, complete the following steps.

Step	Action
1	Click the General tab, shown in the following window. 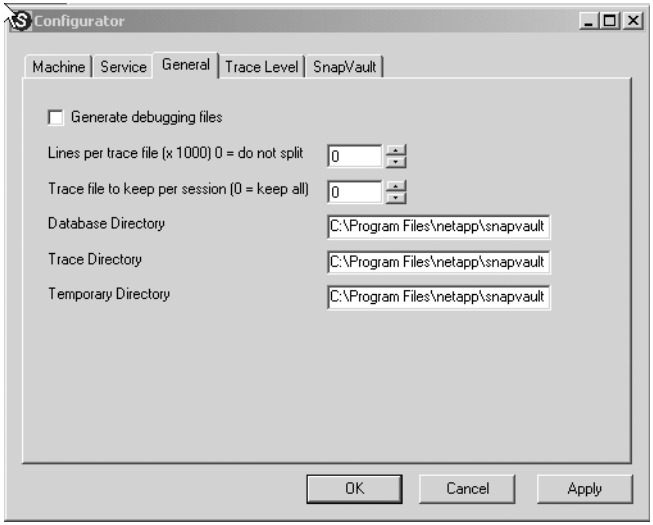

Step	Action
2	Select the “Generate debugging files” check box. **Note** To turn off the generation of debugging files after you are done troubleshooting, unselect the “Generate debugging files” check box.
3	Click the Trace Level tab.
4	From the pull-down menu of the service for which you want to generate debugging information, select from one of the four trace levels—Always, Verbose, Libnormal, and Libverbose. For example, for QSM Server, select VERBOSE, as shown in the following window. 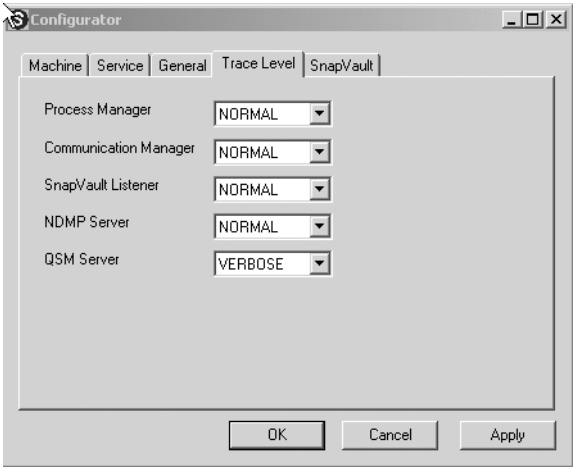 **Note** Make sure you set the trace level back to Normal after you are done troubleshooting.
5	Click Apply.

Step	Action
6	**1.** Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services. Wait until the Current State displays as Stopped. **2.** Click Start Service. Wait until the Current State displays as Running. **3.** Click OK. **Note** Instead of performing the above steps through the `Configurator` utility, you can use the `svpmgr shutdown` and `svpmgr startup` commands on the command line.

Stopping the generation of debugging files

To stop the generation of debugging files, complete the following steps.

Step	Action
1	Click the General tab, shown in the following window. 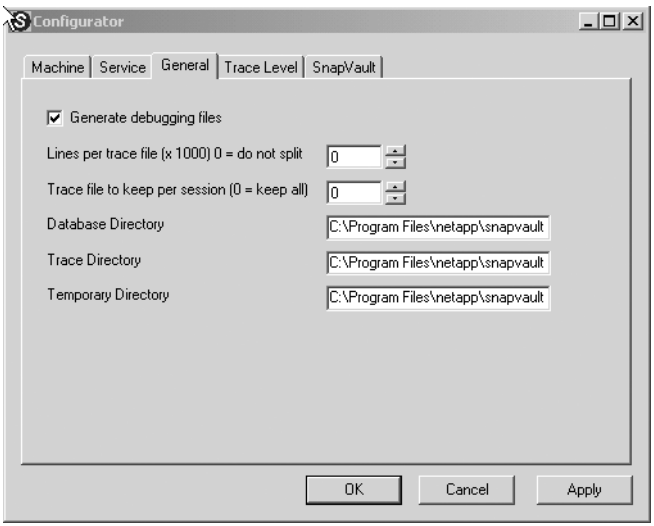

Step	Action
2	Unselect the “Generate debugging files” check box.
3	Click the Trace Level tab.
4	Select NORMAL from the pull-down menu of the service for which you want to stop generating debugging information. For example, for QSM Server, select NORMAL, as shown in the following window. 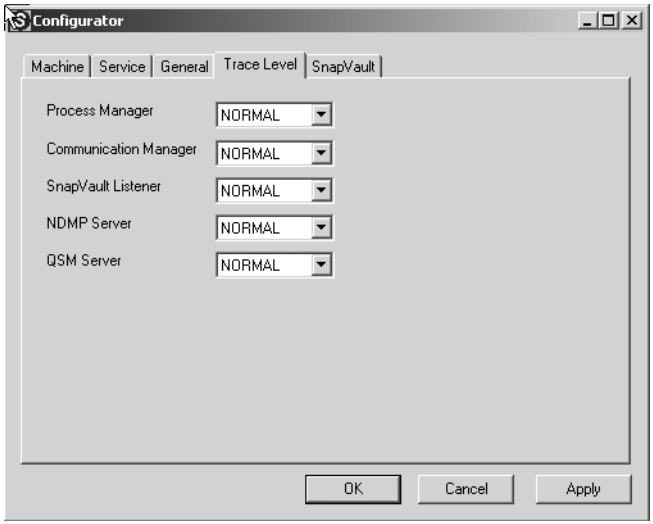
5	Click Apply.
6	Click the Service tab, click Stop Service, click Start Service, and then click OK.

Setting block-level incremental backup options

Setting up checksum computation for block-level incremental (BLI) backups

By default, the BLI level is set to High. However, you can choose the level at which you want to compute BLI checksums or you can disable checksum computations entirely.

To choose the level of BLI checksum computation or to disable it, complete the following steps.

Step	Action
1	Click the SnapVault tab, as shown in the following window. The block-level incremental checksum computation options are displayed in the BLI Level drop-down menu. 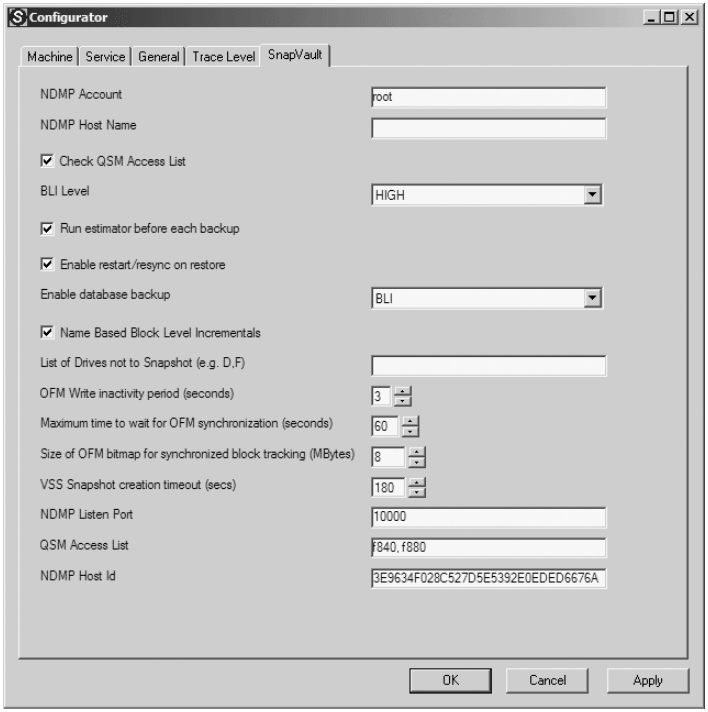

Step	Action
2	Choose the block-level (BLI Level) option you want using the drop-down menu. ◆ HIGH—computes checksums during the initial baseline backup and updates **Note** The HIGH mode might slow down the initial backup transfer. ◆ LOW—computes checksums only during updates ◆ OFF—disables block-level incremental backups altogether

Enabling or disabling BLI backups for certain name-based applications

Name-based block-level incremental backups support for certain applications, such as Microsoft's productivity applications, are enabled by default. To enable or disable this block-level incremental backup feature, complete the following steps.

Step	Action
1	Click the SnapVault tab.
2	To enable or disable block-level incremental backup, select or clear the Name Based Block Level Incrementals option.

Configuring backup exclusion lists

What backup exclusion lists are

Backup exclusion lists are used by Open Systems SnapVault agents to exclude specified files and directories from backups. Open Systems SnapVault agents support two types of exclusion lists:

- File exclusion lists
- Path exclusion lists

File exclusion lists

File exclusion list entries consist of single path elements. A file or directory is excluded if the file name or any path element matches a file exclusion entry in the list.

The file exclusion list is in the file *install_dir*/etc/file-exclude.txt.

Path exclusion lists

Path exclusion list entries consist of complete file system paths to either a directory or a file. If a path exclusion entry specifies a directory, that directory and its files and subdirectories are excluded. A directory path entry must end with a slash (/) on UNIX, or a slash (/) or back slash (\) on Windows.

The path exclusion list is in the file *install_dir*/etc/path-exclude.txt.

Configuring exclusion lists

On Windows systems, exclusion list files are Unicode text files. On UNIX systems, exclusion list files are multibyte text files. Each entry is on its own line. wildcard and special characters are supported.

The following table includes the complete list of wildcard and special characters.

Character	Meaning
/ or \	Path delimiters for Windows You cannot use ! to escape path delimiters.
/	Path delimiter for UNIX You cannot use ! to escape path delimiters.

Character	Meaning
!	Escape character. Use to escape a following special character. However, you cannot escape a path delimiter.
#	If the first character of a line, the line is a comment.
!n	New line
!r	Carriage return
!t	Tab
!f	Form feed
!b	Backspace
!!	A single !
*	Matches any number of characters, including none. Does not cross path boundaries.
?	Matches any one character.

Examples of exclusion list

The following is an example of a file exclusion list:

```
# The following excludes any files ending with .tmp.
*.tmp
# The following excludes any directories with a path
# element ending in .tmp.
*.tmp/
# The following excludes Fred, Frad, and so forth, but not
# Freed.
Fr?d
# Exclusion lists on Unix systems are case sensitive,
# but not on Windows systems. So on a Unix system,
# fred would not be excluded, but fred would be excluded on
# Windows systems.
```

The following is an example of a path exclusion list:

```
# The following excludes a file named tmp in the root directory
/tmp
# (On Windows this includes the root directory for all drives.)
# The following excludes a directory named tmp in the root
# directory.
/tmp/
```

```
# The following excludes directories in /home, but not files.
/home/*/
# On a Windows system, the following excludes files on the
# C drive, in the My Documents folder, with delete in their
# names. Directories with delete in the path element are not
# excluded. On a Unix system, the entry is invalid and is
# ignored.
C:\My Documents\*delete*
```

Configuring open file backup for Windows

About backing up open files

Open Systems SnapVault backups up open files using one of the following options:

- OFM—Used for Windows NT and Windows 2000 platforms

  OFM is automatically installed in your system at the same time the agent is installed; however, you do need to have a license to use the OFM feature as described in "License requirements" on page 16.
- VSS snapshot—Used for Windows 2003 platforms

  The VSS snapshot functionality is integrated with the Open Systems SnapVault agent as a standard feature and does not require a license to use.

About configuring OFM and VSS settings

Under normal conditions, you will not need to change the OFM and VSS settings. These settings need to be changed only for troubleshooting. Therefore, do not change these settings unless directed by NetApp Technical Support.

Configuring OFM

OFM waits for a set time when write activity is quiet and the system is in a safe state to initiate the backup of open files. OFM listens continuously for a period of write inactivity until it is ready to initiate a synchronize for backup or until the preset synchronization timeout period has expired.

OFM requires a quiescent period of process inactivity to synchronize and take a snapshot of the target drive. Without this quiescent period, OFM will not initiate a snapshot. An OFM failure can result in a blank backup.

Note

OFM supports one active snapshot per disk volume letter. For example, if you try to use OFM to back up C:\temp and C:\test, one of the two backups will fail. However, if multiple volumes (or drives) are backed up simultaneously, backups will proceed, but the drives will be backed up sequentially.

If an active snapshot is encountered during an attempted backup, OFM fails.

You need to perform the following operations to use the OFM feature:

- License OFM on the secondary system
- Set OFM parameters

Licensing OFM: To license OFM on the secondary system, complete the following step.

Step	Action
1	On the secondary system, enter the following command: `license add code` *code* is the license code provided to you by Network Appliance.

Setting OFM parameters: To set OFM parameters, complete the following steps.

Step	Action
1	Click the SnapVault tab, as shown in the following window. 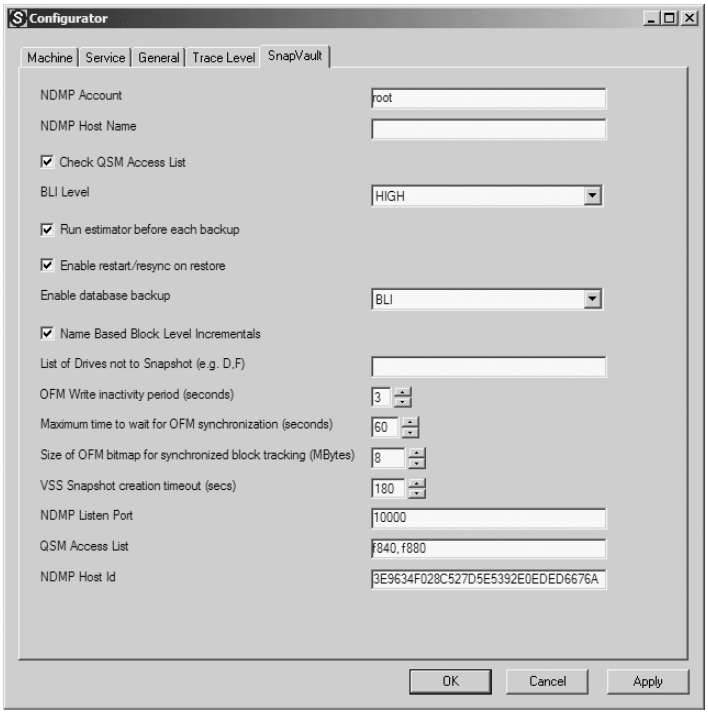

Step	Action
2	Click the up and down arrows to change the “OFM Write inactivity period (seconds)” value. You can select between 1 and 60 seconds. The default value is 3 seconds. **Result:** OFM waits this amount of time, in seconds, for the write inactivity required.
3	Click the up and down arrows to change the “Maximum time to wait for OFM synchronization (seconds)” value. You can set the time between 1 and 60 seconds. The default value is 60 seconds. The synchronization timeout is the maximum wait time that is allowed for OFM to find a period of write inactivity before the wait is ended.
4	Click the up and down arrows to change the “Size of OFM bitmap for synchronized block tracking (MBytes)” value. You can set the size between 1 and 50 MB. The default value is 8 MB. **Result:** This sets the maximum amount of memory in MB that OFM uses to track changed blocks during the backup process.

Disabling open file backup: You can disable open file backup on the secondary storage system.

To disable open file backup on the secondary storage system, complete the following step.

Step	Action
1	On the secondary system, enter the following command: `snapvault modify -o back_up_open_files off`

Setting the VSS snapshot timeout parameter on Windows 2003

Certain conditions must be met before the Open Systems SnapVault agent can acquire a VSS snapshot. You can set the amount of time (snapshot timeout) that the agent waits until it retries a VSS snapshot if the conditions are not right at the time. Setting this parameter avoids unacceptably long waiting periods.

To set the VSS snapshot timeout parameter on Windows 2003, complete the following steps.

Step	Action
1	Click the SnapVault tab.
2	Click the up and down arrows to change the “VSS Snapshot creation timeout (secs)” value. You can set a time between 1 and 180 seconds. The default value is 180 seconds.
3	Click OK.

Setting up backups and restores 4

About this chapter

This chapter describes the basic SnapVault backup and restore functions that you perform using the Open Systems SnapVault software.

Topics in this chapter

This chapter covers the following topics:

- "Setting up SnapVault backup on Open Systems platforms" on page 70
- "Restoring data" on page 74

Setting up SnapVault backup on Open Systems platforms

About configuring SnapVault backup

Setting up SnapVault backup on Open Systems platforms means preparing the systems and SnapVault secondary storage systems to fulfill their backup tasks in the following order:

Note

You must have Open Systems SnapVault licenses for the Open Systems platform and the secondary storage system to use SnapVault.

1. On Open Systems platforms, install the Open Systems SnapVault agent and configure it for backups by the desired SnapVault secondary storage system. See Chapter 2, "Installing the Open Systems SnapVault Software," on page 13 for more information.
2. On the SnapVault secondary storage system, use the filer console commands to license and enable SnapVault, specify the Open Systems platforms to back up. See "Configuring the SnapVault secondary storage system" on page 71.
3. On the SnapVault secondary storage system, start the baseline transfer. See "Creating an initial baseline copy" on page 72.
4. On the SnapVault secondary storage system, schedule times for drives, directories, or subdirectories to be backed up to secondary storage. See "Scheduling SnapVault update backups" on page 73.

Configuring the SnapVault secondary storage system

SnapVault secondary storage system requirement

The SnapVault secondary storage system must be running Data ONTAP 6.5.4 or later to support backup of systems installed with Open Systems SnapVault.

Configuring the SnapVault secondary storage system

To configure the SnapVault secondary storage system to support Open Systems platform SnapVault backup, complete the following steps.

Step	Description
1	License the SnapVault secondary storage system. In the filer console of the SnapVault secondary system, enter the following command: **`license add sv_secondary_license`**
2	License the SnapVault primary storage system. In the filer console of the SnapVault secondary system, enter the following command: **`license add ossv_primary_license`**
3	Enable SnapVault. In the secondary storage system filer console, enter the following command: **`options snapvault.enable on`** For more information see the section on enabling SnapVault of the Data ONTAP *Data Protection Online Backup and Recovery Guide*.
4	Specify the names of the primary storage systems to back up. Enter the following command: **`options snapvault.access host=snapvault_primary1, snapvault_primary2 ...`** **Example:** **`options snapvault.access host=melzhost,samzhost,budzhost`** For more information see the section on enabling SnapVault of the Data ONTAP *Data Protection Online Backup and Recovery Guide*.

Creating an initial baseline copy

Creating a baseline copy

To create an initial baseline copy on the secondary storage system, complete the following step.

Step	Action
1	For each Open Systems platform directory to be backed up to the SnapVault secondary storage system, execute an initial baseline copy from the primary to secondary storage system. ◆ Specify the fully qualified path to the Open Systems host directory that you want to back up. Use the `-S` prefix to indicate the source path. ◆ Even though the Open Systems platform directory to be backed up has no qtree, you *still* need to specify a host and path to the qtree where you will back up this data on the SnapVault secondary storage system. Enter the following command: **`snapvault start -S`** ***`prim_host`***`:`***`dirpath sec_host`***`:/vol/`***`sec_vol/sec_tree`*** **Example 1 (Windows):** **`snapvault start -S melzhost:c:\melzdir sv_secondary:/vol/sv_vol/tree_melz`** **`snapvault start -S samzhost:c:\samzdir sv_secondary:/vol/sv_vol/tree_samz`** **`snapvault start -S budzhost:c:\budzdir sv_secondary:/vol/sv_vol/tree_budz`** **Example 2 (UNIX):** **`snapvault start -S melzhost:/usr/melzdir sv_secondary:/vol/sv_vol/tree_melz`** **`snapvault start -S samzhost:/usr/samzdir sv_secondary:/vol/sv_vol/tree_samz`** **`snapvault start -S budzhost:/usr/budzdir sv_secondary:/vol/sv_vol/tree_bu`**

Scheduling SnapVault update backups

About scheduling SnapVault updates

Open Systems SnapVault supports a maximum of 16 simultaneous transfers from a primary system.

You should plan your backup schedules such that 16 or fewer transfers occur at the same time from the same primary system. You can use the `snapvault status` command on the secondary system to check the number of simultaneous transfers occurring from a primary system.

Scheduling SnapVault update backups

To schedule when Open Systems SnapVault updates backups of drives, directories, or subdirectories, complete the following step.

Step	Action
1	Use the `snapvault snap sched` command to schedule update copying of new or modified data on all Open Systems platform directories that are backed up to qtrees in SnapVault secondary storage. Specify the name of the secondary storage volume containing the secondary qtrees, a snapshot basename (for example: “sv_hourly” or “sv_nightly”), the number of SnapVault snapshots to store on the secondary storage system, and the days and hours to execute. For example: **`snapvault snap sched -x vol1 sv_weekly 1@sat@19`** **`snapvault snap sched -x vol1 sv_nightly 2@mon-fri@19`** **`snapvault snap sched -x vol1 sv_hourly 11@mon-fri@7-18`** **Note** You must use the `-x` parameter in the above command. This parameter causes SnapVault to copy new or modified files from the Open Systems platform directories to their associated qtrees on the secondary storage system. If you do not use the `-x` parameter, the default parameter `-c` is used, which creates snapshots of file systems locally. After all the secondary qtrees on the specified volume have been updated, SnapVault creates a snapshot of this volume for archiving.

Restoring data

When to restore data

In event of data loss or corruption on a primary storage qtree, use the `snapvault restore` command to restore the affected qtree to its state at the time of its last SnapVault snapshot.

Topics in this section

The following topics are covered in this section:

- "Restoring a directory or a file" on page 75
- "Restoring an entire primary storage system" on page 78
- "Restoring files to a primary storage system from tape" on page 79

Restoring a directory or a file

Methods for restoring a directory or a file

In the event of data loss or corruption on the Open Systems platform, the administrator can use one of three different methods for restoring a directory or file:

- Copy files from the secondary storage system to the primary storage system.
- Use the `snapvault restore` command.
- Use the DataFabric Manager restore wizard.

Copying files

You can copy files from the secondary storage system to the NFS or CIFS primary storage system if you want to restore something other than an entire qtree—that is, a single file, a small group of files, or a few directories.

You might want to share the SnapVault destination on the secondary storage system with all the primary storage systems all the time. In this way, end users can perform self-service restores without requiring a backup administrator's assistance.

Note

Some Windows NT and UNIX attributes are not preserved using this method—notably, Windows NT sparse files, Windows NT EFS Encrypted Data, and UNIX access control lists (ACLs).

Copying files to NFS primary storage systems: To restore data by copying files back to an NFS primary storage system, complete the following steps.

Step	Action
1	Mount the backed-up qtree from the SnapVault secondary storage system to the NFS primary storage system.
2	Use the UNIX `cp` command, or an equivalent command, to copy the desired files from the backup to the directory in which you want them.

Copying files to CIFS primary storage systems: To restore data by copying files back to a CIFS primary storage system, complete the following steps.

Step	Action
1	Create a share from the backed-up qtree on the SnapVault secondary storage system to the CIFS primary storage system.
2	Drag and drop the desired files from the backup to the directory in which you want them.

Using the snapvault restore command

You can use the `snapvault restore` command to restore a directory or file on the Open Systems platform to its state at the time of one of its SnapVault snapshots.

To use the `snapvault restore` command, complete the following steps.

Step	Description
1	Navigate to the *install_dir*/bin on your Open Systems platform whose data you want to restore.

Step	Description

2	Enter the `snapvault restore` command and specify the following:
	◆ The secondary storage system host and the path to the secondary qtree, directory, and file that you want to restore from
	◆ The name of the snapshot that you want to restore from (for example, sv_weekly.0, sv_weekly.1, or sv_weekly.2)
	◆ The path on the primary storage system to the directory or file that you want to restore to
	Example 1 (Windows):
	`snapvault restore -s sv_daily.0 -S myvault:/vol/sv_vol/melzdir/evidence.doc a:\melzdir\evidence.doc`
	Example 2 (UNIX system):
	`snapvault restore -s sv_daily.0 -S myvault:/vol/sv_vol/melzdir/evidence.doc /usr/melzdir/evidence.doc`
	Note
	Enter the entire command as a single line.
	Make sure that you do not specify a trailing slash (\ or /) character at the end of the path name in the above command; otherwise, the `snapvault restore` command will fail.

Using the DataFabric Manager restore wizard

The DataFabric Manager restore wizard leads you through the entire restore process. See the *DataFabric Manager Administration Guide* for details.

Restoring an entire primary storage system

Restoring a primary storage system

You can restore an entire NFS or CIFS primary storage system from a SnapVault secondary storage system, but the restore cannot be to a primary storage system that has a blank hard disk. There must be an operating system on the disk.

To restore an entire primary storage system, complete the following steps.

Step	Action
1	Reinstall the operating system on the primary storage system.
2	Reformat the file system as the original file system was formatted.
3	Install the SnapVault agent. See Chapter 2, "Installing the Open Systems SnapVault Software," on page 13.
4	(Optional) Restore the Windows System State data of the primary storage system you are restoring if you backed up its Windows System State data. See "Restoring System State data" on page 92 for more information.
5	Restore the backed-up directories using the `snapvault restore` command. See "Using the snapvault restore command" on page 76 for details.

Restoring files to a primary storage system from tape

About restoring from tape

The process of restoring from tape to a primary storage system involves first restoring the data from the tape to a secondary storage system and then restoring from that secondary storage system to the primary storage system using Open Systems SnapVault.

Restoring from tape

To perform a SnapVault restore to an NFS or CIFS primary storage system from tape, complete the following steps.

Note

Some Windows NT and UNIX attributes are not preserved using this method; notably, Windows NT sparse files, Windows NT EFS Encrypted Data, and UNIX access control lists (ACLs).

Step	Action
1	Mount the tape that has the files that need to be restored.
2	Use the `restore` command to restore from the tape to the SnapVault secondary storage system. See the *Tape Backup and Recovery Guide* for details.
3	Copy the files from the SnapVault secondary storage system to the NFS or CIFS primary storage system. See "Restoring a directory or a file" on page 75 for details.

Open Systems SnapVault Management 5

About this chapter

Most of the Open Systems SnapVault management tasks you might perform are the same as the ones you perform for any other SnapVault relationship. For general SnapVault management tasks, see the Data ONTAP *Data Protection Online Backup and Recovery Guide*. This chapter describes procedures specific to Open Systems SnapVault management.

Topics in this chapter

This chapter describes the following management procedures that you can perform:

- "Locating status and problem reports" on page 82
- "Backing and restoring the Open Systems SnapVault database" on page 83
- "Backing up and restoring Windows System State data" on page 88
- "Deleting and recreating Open Systems SnapVault relationships" on page 97
- "Migrating a relationship between two secondary systems" on page 98
- "Migrating between two volumes on one secondary system" on page 101
- "Setting up a tertiary system for a relationship" on page 104
- "Reusing a deleted or renamed directory name" on page 105
- "Resynchronizing restored or broken relationships" on page 107
- "Retrying failed transfers" on page 109
- "Encrypted File System (EFS) encrypted file backup and restore" on page 110

Locating status and problem reports

Where to find status and problem reports

You can find primary storage system operational status and problem reports in the log files called snapvault.*yyyymmdd*, where *yyyy* is the year, *mm* is the month, and *dd* is the date when the file was created. A new file is created daily at midnight.

The log files are located in the following directory:

`install_dir/snapvault/etc`

install_dir is the directory on the primary storage system in which you installed the Open Systems SnapVault agent. On Windows systems, the default location for *install_dir* is the c:\Program Files directory. On UNIX systems, the default location for *install_dir* is the /usr directory.

Note

If the *install_dir* path includes spaces in the pathname, you must enclose the path in double quotes (“ ”); for example, “C:\Program Files\netapp\snapvault\bin\snapvault.exe”.

You can find secondary storage system reports in the /etc/log/snapmirror file in the root volume.

Backing and restoring the Open Systems SnapVault database

About the Open Systems SnapVault database

The Open Systems SnapVault database consists of a set of files that contain information about the Open Systems SnapVault relationship between a primary and a secondary system. A unique set containing the following files is maintained for each relationship:

- History file
- BLI checksums file (if BLI is enabled)
- ACL file
- (Optional) Checkpoint file (if a transfer failed)

Naming convention for the database files

The following convention is used to name the files comprising the Open Systems SnapVault database.

File	Naming convention
History	*xx* *xx* is a unique integer
ACL	*xx*-ACLs
BLI checksums	*xx*-checksums

Example: For an Open Systems SnapVault relationship with BLI enabled, the following files exist in the Open Systems SnapVault database:

```
D:\Program Files\netapp\snapvault\db\QsmDatabase\Files\
QtreeHistory\0000\00>dir

11/08/05  02:32  4,676  01              History file
11/08/05  02:32  1,996  01-ACLs         ACL file
11/08/05  02:32  48     01-checksums    BLI checksums file
```

Need to backup and restore the Open Systems SnapVault database

If the Open Systems SnapVault database becomes corrupt or gets out-of-sync with the secondary system, data transfers between the primary and secondary systems cannot continue. If you do not have a way to restore the database, you are forced to initiate a baseline transfer from the primary to the secondary

system. However, if you maintain a backup copy of the database, you can restore the database for the relationship and continue with subsequent data transfers with minimal downtime and without the need to perform a baseline transfer.

Note

The method discussed in this section is the only way to restore an Open Systems SnapVault database. The database is not restored when you restore the whole backup or individual directories or files.

How the backup functionality works

By default, backup of the Open Systems SnapVault database (the history file and its corresponding BLI checksums file) occurs automatically every time data is transferred from a primary system to a secondary system. A compressed file of the database is created and transferred to the secondary system during each data transfer.

The ACL file is not included in the backup because this file is recreated at each transfer; therefore, there is no benefit to backing up this file.

Any checkpoint files, if existing, are also not included in the backup.

After the compressed file has been transferred over to the secondary system, it is deleted from the primary system.

At the secondary system, the compressed file is placed in the root of the qtree where backup files for an Open Systems SnapVault relationship are located.

Characteristics of the database backup file

The compressed database file is located in the *install_dir*/tmp directory and is named .OSSV_DATABASE_BACKUP.

Note

The dot (.) in the database file name indicates that it is a hidden file. But the file can be viewed using a CIFS or an NFS client.

If a file named .OSSV_DATABASE_BACKUP exists in the root of the qtree for a SnapVault relationship, the second and subsequent files to be created follow the naming convention .OSSV_DATABASE_BACKUP_*x*, where *x* is an integer used to uniquely identify each file.

The file name .OSSV_DATABASE_BACKUP is *not* case sensitive.

How the restore functionality works

The database file can be restored by using the `snapvault restore` command, however, you must include the file name .OSSV_DATABASE_BACKUP in the command. See "Restoring the agent database" on page 86 for the steps to follow for restoring the database file.

After the database file has been restored, Open Systems SnapVault software decompresses it automatically and places the decompressed files where Open Systems SnapVault database files are located for the relationship. Data transfers can be performed from this point on.

Note

If any data updates occur between the time a database file is backed up and the database is restored from it, they cause the secondary system to get out-of-sync with the primary system; therefore, subsequent data updates cannot be continued. In such a case, you must first resynchronize the relationship with the `snapvault start -r` command. After the resynchronization has completed, data updates can be performed as usual. For more information on resynchronizing a relationship, see "Resynchronizing restored or broken relationships" on page 107.

Backing up the agent database

By default, backup of the Open Systems SnapVault database (the history file and its corresponding BLI checksums file) occurs automatically every time data is transferred from a primary system to a secondary system.

To change the database backup option, complete the following steps.

Step	Action
1	In the Configurator utility, click the SnapVault tab.

Step	Action	
2	**If...**	**Then...**
	You want to back up the history file and its corresponding BLI checksums file	Select “BLI” from the “Enable database backup” drop-down list. **Note** This option is selected by default.
	You want to back up only the history file	Select “DB only” from the “Enable database backup” drop-down list.
	You want to disable the database backup functionality	Select None from the “Enable database backup” drop-down list.

Restoring the agent database

To restore the Open Systems SnapVault agent database, complete the following steps.

Step	Action
1	Using the command-line interface of the Open Systems SnapVault primary storage system, navigate to the snapvault/bin directory. ◆ On Windows systems, the default path is c:\Program Files\snapvault\bin ◆ On UNIX systems, the default path is /usr/snapvault/bin

Step	Action
2	Enter the following command: **`snapvault restore -S`** ***`secondary_system`*****`:`** ***`pathname`*** **`/.OSSV_DATABASE_BACKUP`** ***`pri_pathname`*** *secondary_system* is the secondary storage system. *pathname* is the path where the .OSSV_DATABASE_BACKUP file is located on the secondary system. *pri_pathname* is the path on the primary storage system to which the database will be restored. **Note** Make sure that you do not specify a trailing slash (\ or /) character at the end of the path name in the above command; otherwise, the `snapvault restore` command will fail. **Example:** To restore the database from a secondary system called f840 to a directory \temp\database on the primary system, enter the following command: **`D:\Program Files\netapp\snapvault\bin>snapvault restore –S f840:/vol/vol0/rel5/.OSSV_DATABASE_BACKUP D:\temp\database`**

Disabling database backup

To disable backing up a database, complete the following steps.

Step	Action
1	In the Configurator utility, click the SnapVault tab.
2	Select None from the “Enable database backup” drop-down list.

Backing up and restoring Windows System State data

What System State data is

Depending on configuration, Windows 2000 and Windows 2003 systems have some or all of the following System State data:

- Registry
- COM+ Class Registration database
- System files, including the boot files
- Certificate Services database
- IIS Metadirectory
- System files that are under Windows File Protection
- Performance counters

Additionally, the System State data on domain controllers includes Active Directory and SYSVOL data.

Why you back up and restore System State data

You can add backups of Windows System State data to existing Open Systems SnapVault backup schedules and use the backups to restore a system to an earlier state. This can be useful if, for example, an Active Directory entry is accidentally deleted. You can also use Open Systems SnapVault System State data backup in conjunction with complete file system backups as part of a disaster recovery plan.

Possible issues

If the Windows System State data includes registry and domain information, you might experience some issues when backing up and restoring data. For example, if you restore registry data from one system to a different system, the restored registry entries might not be the correct entries for the new system. In such a case, performance might be degraded or the system might not be functional. See the Microsoft Knowledge Base for information on such issues.

Primary, authoritative, and non-authoritative restores

Two components of the Windows System State data involve functions normally coordinated over multiple systems: Active Directory and SYSVOL (or the File Replication Service). Inherent to these distributed system is their interaction with Active Directory and SYSVOL functions on other systems in a domain/forest.

The Active Directory on a restored system can have its authority option set to authoritative or non-authoritative. The SYSVOL authority option can be authoritative, non-authoritative, or primary.

Windows does not restrict any combinations of Active Directory and SYSVOL authority options. For example, a system can have an authoritative Active Directory and a non-authoritative SYSVOL. However, Microsoft recommends the following when System State data is restored:

- A primary restore should only be used when all domain controllers have been lost and the domain is being completely recreated.
- Because of the way they interact, Active Directory and SYSVOL should have matching authority.

These two points result in three combinations of restore scenario, SYSVOL authority, and Active Directory authority as shown in the following table:

Scenario	**SYSVOL Authority**	**Active Directory Authority**
Recreating the first or only domain controller	primary	authoritative
Performing a non-authoritative restore of System State data or as a part of disaster recovery	non-authoritative	non-authoritative
Performing an authoritative restore of System State data or as a part of disaster recovery	authoritative	authoritative

Active Directory restores performed by Open Systems SnapVault can only be non-authoritative, but you can use the Windows utility `ntdsutil` to change the authority. SYSVOL restores performed by Open Systems SnapVault can be specified as primary, authoritative, or non-authoritative by using the `snapvault restore` command secondary path options `SystemStatePrimary`, `SystemStateAuthoritative`, or `SystemState`, respectively.

Disaster Recovery Planning

As part of your disaster recovery planning, consider the Active Directory and SYSVOL authority status of each system in the domain and use it to determine in what order to restore systems, and what authority settings to use when they are restored. Take into account the following:

- A non-authoritative system cannot become the domain controller until the file replication service permits. This means a non-authoritative system cannot become the domain controller until it contacts another system with a working file replication service.
- A restored domain controller that is assigned certain FSMO (Flexible Single Master Operation) Active Directory roles cannot function as the domain controller until it has replicated with another domain controller.

  For example, if the restored domain controller has the RID (relative ID) Master role, it must contact another domain controller with which it is set to replicate before the restored domain controller will function as the domain controller.

Also review the Microsoft Knowledge Base for additional issues.

For detailed information

The following sections discuss ways to create and use System State data backups:

- "Adding System State data backup" on page 91
- "Restoring System State data" on page 92
- "Using System State backup to rebuild a primary system" on page 94

Adding System State data backup

Initiating and starting a System State data backup

To initiate a System State data backup, and to add System State data backup to an existing Open Systems SnapVault backup schedule, complete the following step.

Note

On Windows 2000 systems, before starting System State data backups, make sure the system has recent Windows Service Packs installed. You can have problems with restored systems if the backed-up system was running an older version of the operating system. For information on which service pack your system must be running, see "Requirements for primary storage systems" on page 14.

Step	Action
1	In the filer console of the SnapVault secondary system, enter the following command: **`snapvault start -S`** ***`prim_host`*****`:SystemState`** ***`sec_host`***`:/vol/`***`sec_vol/sec_tree`*** **Note** The keyword `SystemState` is case independent. **Example:** **`snapvault start -S melzhost:SystemState sv_secondary:/vol/sv_vol/tree_melz`**

Note

System State data backups of domain controllers are only valid for the configured tombstone lifetime setting for the enterprise. The default tombstone lifetime is 60 days.

Boot files and system files are backed up even when they are on different volumes.

Subsequent backups use block incremental backups.

Restoring System State data

Restoring System State data from a backup

Use this procedure to restore the System State data, unless all domain controllers are being recreated, and this is the first domain controller to be restored. In that case, use the procedure "Restoring System State data from a backup and marking it primary" on page 93.

To restore the System State data from a backup, complete the following steps.

Step	Action
1	If the machine is a domain controller, reboot it into Directory Services Restore Mode. (You can enter the Directory Services Restore Mode by holding down the Ctrl key when the machine is booting, and then pressing the F8 key at the startup menu.)
2	On the primary, enter the following command: **`snapvault restore -S sec_host:/vol/sec_vol/sec_tree SystemState`** **Note** The keyword `SystemState` is case independent. Make sure that you do not specify a trailing slash (\ or /) character at the end of the path name in the above command; otherwise, the `snapvault restore` command will fail. **Example:** **`snapvault restore -S sv_secondary:/vol/sv_vol/tree_melz SystemState`**
3	Restored Active Directory information is marked as non-authoritative. When a domain controller with non-authoritative entries reconnects to the domain, replication services update those entries with authoritative values. To mark any or all of the restored entries as authoritative, use the Microsoft ntdsutil tool. Otherwise skip to the next step.
4	Reboot.

Restoring System State data from a backup and marking it primary

Use this procedure to restore the System State data for the first domain controller after all domain controllers have been lost. Do not use this procedure if there are still functioning domain controllers. To restore the System State data from a backup, complete the following steps.

Step	Action
1	If the machine is a domain controller, reboot it into Directory Services Restore Mode. (You can enter the Directory Services Restore Mode by holding down the Ctrl key when the machine is booting, and then pressing the F8 key at the startup menu.)
2	In the primary system, enter the following command: **`snapvault restore -S sec_host:/vol/sec_vol/sec_tree SystemStatePrimary`** **Note** The keyword `SystemStatePrimary` is case independent. Make sure that you do not specify a trailing slash (\ or /) character at the end of the path name in the above command; otherwise, the `snapvault restore` command will fail. **Example:** **`snapvault restore -S sv_secondary:/vol/sv_vol/tree_melz SystemStatePrimary`**
3	Reboot.

Using System State backup to rebuild a primary system

About using System State backup

You can use System State data backups to rebuild a primary system in case of a disaster. The process of backing up and, if needed, rebuilding a primary system involves the following:

- Backing up system drive and the Windows
- Creating a Windows System State data backup to a secondary system
- Rebuilding a primary system from the Windows System State data backup

Creating a backup to rebuild a primary system

To create a backup that you can use to rebuild a primary system, complete the following steps.

Note

This procedure is designed to back up the operating system and its state. The procedure does not reliably back up all application data. Review your application documentation for any steps necessary to back up the application and application data.

Open Systems SnapVault files are not backed up because they are modified during the backup.

Step	Action
1	Back up the entire system drive.
2	If a computer utility partition exists, and is accessible as part of a file system, back up the partition.

Step	Action
3	In the filer console of the SnapVault secondary system, enter the following command: **`snapvault start -S`** ***`prim_host`*****`:SystemState`** ***`sec_host`***`:/vol/`***`sec_vol/sec_tree`*** **Note** The keyword `SystemState` is case independent. **Example:** **`snapvault start -S melzhost:SystemState sv_secondary:/vol/sv_vol/tree_melz`**

Rebuilding a primary system from the System State backup

Guidelines for rebuilding a primary system: You must keep the following guidelines in mind before proceeding with the procedure:

- Using a System State data backup on a system that is not effectively a duplicate of the original system disables the new system.
- The following procedure lists several parameters that are required to be the same (or in the case of disk drives, same or larger), but there are also areas of potential differences, such as video bus types (for example AGP bus versus PCI video bus), that might cause problems depending on the specifics of the systems.

To rebuild a primary system using the backup you created in "Creating a backup to rebuild a primary system" on page 94, complete the following steps.

Step	Action
1	Install and configure the base operating system. Note the following restrictions: ◆ Use the same operating system and service packs that were on the original system. ◆ Use the same machine name as on the original system. ◆ Use the same drive letter mappings. ◆ Make sure each drive is at least as large as the corresponding drive was when the backup was made. ◆ Format each drive with the same file system type and version as on the original system. ◆ Make sure the hardware configuration is identical to the original. Do not perform any operating system configuration tasks beyond those needed to satisfy these restrictions.
2	Install and configure Open Systems SnapVault in the same location it was installed on the original machine.
3	If the machine is a domain controller, boot into Directory Services Restore Mode.
4	Restore the system drive.
5	Restore the System State data. Choose between the procedures "Restoring System State data from a backup" on page 92 and "Restoring System State data from a backup and marking it primary" on page 93 based on the points about authoritative versus non-authoritative and primary versus non-primary restores discussed as part of the procedures.
6	Reboot.
7	Reinstall any applications not restored by steps 1 through 6.
8	Restore any application data not restored by steps 1 through 6.

Deleting and recreating Open Systems SnapVault relationships

About deleting and recreating a relationship

The process of deleting an Open Systems SnapVault relationship requires that you delete an Open Systems SnapVault relationship on the secondary storage system and also release the relationship on the primary system to free up the primary directory for future backups.

If you use the `snapvault stop` command on the secondary to delete an Open Systems SnapVault relationship and try to recreate the relationship without releasing the relationship on the primary system, the attempts fails with an error message similar to this:

date and time `[worker_thread_162:error]: snapvault: destination transfer from` *source file* `to` *destination qtree*`: the qtree is not the source for the snapmirror destination`

`Transfer aborted: the qtree is not the source for the snapmirror destination.`

Deleting and recreating a relationship

Follow these general steps to delete and recreate Open Systems SnapVault relationships.

Step	Action
1	On the secondary system console enter this command: **`snapvault stop`** ***`secondary_path`***
2	On the primary system console enter this command: **`snapvault release`** ***`primary_path`*** **`[ secondary: ]`** ***`secondary_path`***
3	On the secondary system console enter this command: **`snapvault start -S`** ***`primary_path`*** **`[ secondary: ]`** ***`secondary_path`***

Migrating a relationship between two secondary systems

Migrating data from one secondary to another

Before migrating data from one secondary to another secondary, review the basics of setting up SnapVault transfers (access permission, licensing, and correct volume language for example) in the Data ONTAP *Data Protection Online Backup and Recovery Guide*.

To migrate a volume that contains SnapVault destination qtrees from one secondary system to another secondary system without having to perform another baseline transfer, complete the following steps.

Step	Action
1	Ensure that you have Open Systems SnapVault baselines of the directory you are migrating. **Example:** In this procedure, assume a baseline of the bno:C:\500MB directory was backed up to r200-old:/vol/old_vol/bno_C_500MB.

2	Using SnapMirror, replicate the volume from the present secondary system to a volume on the new secondary system. See the SnapMirror chapter in the Data ONTAP *Data Protection Online Backup and Recovery Guide* for details about creating volume replicating SnapMirror relationships.
	
	Example: To replicate the old_vol volume from the r200-old secondary system to the new_vol volume on the r200-new secondary system, complete the following steps on the new secondary system (r200-new):
	
	1. Create the new_vol volume.
	`r200-new> vol create new_vol 3`
	
	2. Mark the new_vol volume restricted.
	`r200-new> vol restrict new_vol`
	
	3. Transfer the old_vol volume to the new_vol volume.
	`r200-new> snapmirror initialize -S r200-old:old_vol new_vol`

Step	Action
3	Quiesce and break the SnapMirror relationship between the old secondary system and the new secondary system. **Example:** To quiesce and break the SnapMirror relationship between r200-old and r200-new, complete the following steps on r200-new: 1. **`snapmirror quiesce new_vol`** 2. **`snapmirror break new_vol`**

4	Check the SnapMirror status and SnapVault status on the new secondary. SnapMirror status should be `Broken-off`. SnapVault status should be `Snapvaulted` to the new volume on the new secondary system.
	
	Example: Perform the following steps from r200-new.
	
	1. **`snapmirror status`**
	`Source  Destination  State`
	`r200-old:old_vol  r200-new:new_vol  Broken-off`
	
	2. **`snapvault status`**
	`Source  Destination  State`
	`bno:C:\500MB  r200-new:/vol/new_vol/bno_C_500MB  Snapvaulted`

5	Confirm that SnapVault configuration information is not present on the new secondary system by using the `snapvault status -c` command. **Example:** Perform the following step from r200-new. **`snapvault status -c`** `Snapvault secondary is ON.`

6	Add SnapVault configuration information to the registry on the new secondary system using the `snapvault start` command.
	
	Note
	This does not start a new baseline, it updates the registry.
	
	Example: Perform the following step from r200-new.
	
	`snapvault start -S bno:C:\500MB r200-new:/vol/new_vol/bno_C_500MB`
	`Snapvault configuration for the qtree has been set.`
	`Qtree /vol/new_vol/bno_C_500MB is already a replica.`

Step	Action

7	Confirm that SnapVault configuration information is present on the new secondary system using the `snapvault status -c` command.
	
	Example: Perform the following step from r200-new.
	
	`snapvault status -c`
	`Snapvault secondary is ON.`
	`/vol/new_vol/bno_C_500MB source=bno:C:\500MB`

8	Test the new SnapVault relationship by manually updating r200-new.
	
	If you are using the command line to manage your environment, continue to the next step; otherwise, you are finished.
	
	Example: Perform the following step from r200-new.
	
	`snapvault update r200-new:/vol/new_vol/bno_C_500MB`
	`Transfer started.`
	`Monitor progress with 'snapvault status' or the snapmirror log.`

9	Recreate any schedules used on the old secondary system to the new secondary system and ensure access permissions are in place.

Migrating between two volumes on one secondary system

Migrating from one volume to another on a secondary system

To migrate a volume that contains SnapVault destination qtrees to another volume on the same secondary system without having to perform another baseline transfer, complete the following steps.

Step	Action
1	Ensure that you have Open Systems SnapVault baselines of the directory you are migrating. **Example:** In this procedure, assume a baseline of the bno:C:\500MB directory was backed up to r200:/vol/old_vol/bno_C_500MB.

2	Using SnapMirror, replicate the volume from the present volume on the secondary system to a new volume. See the SnapMirror chapter in the Data ONTAP *Data Protection Online Backup and Recovery Guide* for details about creating volume replicating SnapMirror relationships.
	
	Example: To replicate the old_vol volume on the r200 secondary system to the new_vol volume, complete the following steps on the secondary system (r200):
	
	1. Create the new_vol volume.
	`vol create new_vol 3`
	
	2. Mark the new_vol volume restricted.
	`vol restrict new_vol`
	
	3. Transfer the old_vol volume to the new_vol volume.
	`snapmirror initialize -S r200:old_vol new_vol`

Step	Action
3	Quiesce and break the SnapMirror relationship between the old volume and the new volume. **Example:** To quiesce and break the SnapMirror relationship between old_vol and new_vol, complete the following steps on r200: 1. **`snapmirror quiesce new_vol`** 2. **`snapmirror break new_vol`**

4	Check the SnapMirror status and SnapVault status of the new volume. SnapMirror status should be `Broken-off`. SnapVault status should be `Snapvaulted` to the new volume.
	
	Example: Perform the following steps from r200.
	
	1. **`snapmirror status`**
	`Source  Destination  State`
	`r200:old_vol  r200:new_vol  Broken-off`
	
	2. **`snapvault status`**
	`Source  Destination  State`
	`bno:C:\500MB  r200:/vol/new_vol/bno_C_500MB  Snapvaulted`

5	Confirm that SnapVault configuration information is not present for the new volume using the `snapvault status -c` command. **Example:** Perform the following step from r200. **`snapvault status -c`** `Snapvault secondary is ON.`

6	Add SnapVault configuration information to the registry on the new volume using the `snapvault start` command.
	
	Note
	This does not start a new baseline; it updates the registry.
	
	Example: Perform the following step from r200.
	
	`snapvault start -S bno:C:\500MB r200:/vol/new_vol/bno_C_500MB`
	`Snapvault configuration for the qtree has been set.`
	`Qtree /vol/new_vol/bno_C_500MB is already a replica.`

Step	Action

7	Confirm that SnapVault configuration information is present on the new volume using the `snapvault status -c` command.
	
	Example: Perform the following step from r200.
	
	`snapvault status -c`
	`Snapvault secondary is ON.`
	`/vol/new_vol/bno_C_500MB source=bno:C:\5000MB`

8	Test the new SnapVault relationship by manually updating new_vol.
	
	If you are using the command line to manage your environment, continue to the next step; otherwise, you are finished.
	
	Example: Perform the following step from r200.
	
	`snapvault update r200:/vol/new_vol/bno_C_500MB`
	`Transfer started.`
	`Monitor progress with 'snapvault status' or the snapmirror log.`

9	Recreate any schedules used on the old volume to the new volume and ensure access permissions are in place.

Setting up a tertiary system for a relationship

Need for a tertiary system

You can protect against disasters on the SnapVault secondary system by using the SnapMirror® feature. The configuration involves setting up SnapMirror relationships from the volumes on your SnapVault secondary system to volumes on a remote (tertiary) Data ONTAP system, as shown in the following illustration. SnapMirror thus provides an exact replica of the SnapVault secondary data on the tertiary system.

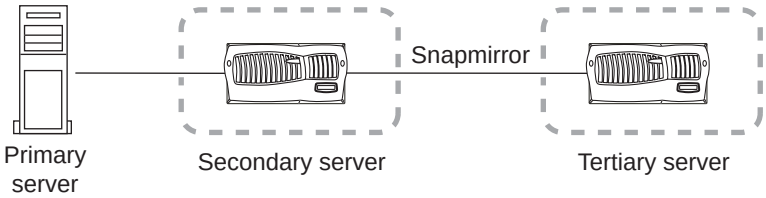

Additionally, softlock support in Open Systems SnapVault enables you to continue SnapVault replication relationships between the original SnapVault primary system and the tertiary system, without any initial baseline transfers. For instance, if your SnapVault secondary system becomes unusable because of a disaster, you can manually redirect your next SnapVault transfers to the tertiary system instead of the old SnapVault secondary system. Effectively, your tertiary system becomes the new SnapVault secondary system, and your SnapVault transfers continue using the latest “snapshot” common to both the primary and the tertiary systems.

“Snapshot” is in quotes in the preceding sentence because Open Systems SnapVault primary systems do not literally keep Snapshot™ copies. They keep the equivalent state in a database that allows the Open Systems SnapVault primary system to perform many operations equivalent to those of a filer primary system.

For more information about using SnapMirror, see the Data ONTAP *Data Protection Online Backup and Recovery Guide*.

Reusing a deleted or renamed directory name

When to reuse a directory name

If you delete or rename a directory that is a source of a SnapVault relationship and create another directory with the same name, the next SnapVault update transfer will fail with the following error messages:

```
Root Inode has changed
Failed to generate update inode values
```

Perform the procedures discussed in this section to delete or recreate the Open Systems SnapVault relationship.

Reusing a deleted or renamed directory name

If the SnapVault source directory on the primary system was intentionally deleted and backups are not necessary: Complete the following steps.

Step	Action
1	Delete the backed-up directory on the secondary system.
2	Release the SnapVault relationship from the Open Systems SnapVault primary system using the `snapvault release` command.

If the SnapVault source directory on the primary system was intentionally renamed: Complete the following steps.

Step	Action
1	Create a new SnapVault relationship for the renamed directory and perform a baseline transfer for the new relationship.
2	When previous backups of the original directory are no longer needed, delete the original SnapVault relationship and release that relationship from the Open Systems SnapVault primary system.

If the SnapVault source directory on the primary system was erroneously renamed: Complete the following steps.

Step	Action
1	Change the name of the directory back to its original name.
2	Continue performing update backup transfers as you did before the erroneous renaming of the directory.

Resynchronizing restored or broken relationships

How systems get unsynchronized

Systems in an Open Systems SnapVault relationship are considered synchronized as long as a common snapshot exists between the primary and secondary systems. A common snapshot is necessary for incremental backups to continue successfully between the primary and secondary systems. Once the common snapshot is lost, incremental backups start failing and the systems get unsynchronized.

The systems in an Open Systems SnapVault relationship can become unsynchronized under the following conditions:

- An older version of the Open Systems SnapVault database is restored on the primary after its database is corrupted.
- Data is restored using the `snapvault restore` command.

  The `snapvault restore` command is used to restore primary data to the state at the time of one of its SnapVault snapshots.
- The state of a destination qtree in a SnapVault relationship is changed to read-writable.

  Even if the contents of the qtree were not modified, if the state of a secondary qtree is changed from read-only to read-write, you must resynchronize the SnapVault relationship between the primary and secondary systems so that the secondary qtree becomes read-only and the incremental transfers can continue.

  Caution

  If the contents of the qtree were modified, all data written to it will be lost upon resynchronizing.

Need for resynchronization

Prior to Open Systems SnapVault 2.2, the only way to resynchronize a SnapVault relationship between a primary and a secondary was by reinitializing the relationship. Reinitializing a relationship involves a lengthy baseline transfer between the primary and the secondary systems, which is not desirable in most cases.

Starting with Open Systems SnapVault 2.2, you can use the `snapvault start -r` command to resynchronize a relationship without having to reinitialize it.

Resynchronizing a relationship

To resynchronize a restored or broken Open Systems SnapVault relationship, complete the following step.

Step	Action
1	On the secondary system, enter the following command: **`snapvault start -r -S`** ***`prim_host`***`:`***`dirpath`*** ***`sec_host`***`:/vol/`***`sec_vol/sec_tree`*** **Example 1 (Windows):** **`snapvault start -r -S melzhost:c:\melzdir sv_secondary:/vol/sv_vol/tree_melz`** **Example 2 (UNIX):** **`snapvault start -r -S melzhost:/usr/melzdir sv_secondary:/vol/sv_vol/tree_melz`**

Retrying failed transfers

About retrying failed transfers

If a transfer stops because of a retryable error such as a temporary network outage, the transfer is automatically retried after 60-second wait. The number of retries is determined by the tries option for the relationship. The default tries value is two, which means that a stopped transfer is retried once (the failed initial try counts as the first try).

You can change the value for the number of times a failed transfer is retried.

Changing the number of retry attempts made for failed transfers

To change the value for the number of retry attempts made for failed transfers, complete the following step on the secondary storage system.

Step	Action
1	Enter the following command: **`snapvault modify -t`** ***`n sec_qtree`*** *n* is the number that specifies the number of retries. **Note** The failed attempt counts as the first try. Therefore, if you set *n* to 5, four attempts are made after the failed try. *sec_qtree* is the qtree on the secondary system where data is being backed up.

Encrypted File System (EFS) encrypted file backup and restore

About EFS encrypted file backup and restore

Open Systems SnapVault is capable of backing up and restoring EFS encrypted files automatically as long as the requirements listed in the following section are met.

You cannot use block-level incremental backup to back up EFS encrypted files. Any time an EFS encrypted file is modified, Open Systems SnapVault backs up the entire EFS encrypted file.

Requirement for EFS encrypted file backup and restore

The following are the requirements for backing up and restoring EFS encrypted files in an Open Systems SnapVault relationship:

- A version of Data ONTAP that has a fix for NetApp Bug 139696. See the following URL to determine the versions of Data ONTAP that have a fix for this bug. You can review the text of bug 139696 on the NOW site by entering the bug number 139696 in the Bugs Online Quick Search Enter Bug ID(s) field or entering in your browser the following URL: http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=139696.

  Note

  If the version of Data ONTAP on your secondary system does not support Open Systems SnapVault backups of EFS encrypted files, Open Systems SnapVault skips those files and creates a log entry (in the snapvault log files, located in the *install_dir*/etc directory) listing the skipped files.

- A sufficient amount of free space in the target Windows volume

  Replacing existing EFS files with restored EFS encrypted files requires the Open Systems SnapVault agent to create a temporary file that is equivalent to the size of the EFS encrypted file that is being replaced. Open Systems SnapVault is capable of restoring up to five files concurrently. This means that a restore of EFS encrypted files may require free space in the target Windows volume that is equal to or greater than the sum of the size of the five largest EFS encrypted files in target volume.

Open Systems SnapVault space estimator 6

What Open Systems SnapVault space estimator is

The Open Systems SnapVault space estimator is a utility available with the Open Systems SnapVault product that enables you to ascertain if there is sufficient disk space available on the Open Systems SnapVault primary to perform a backup. If this utility is run on a system that does not have the Open Systems SnapVault product installed, it can also provide recommendations on where to install the product and its database and temporary files based on the currently available free space.

How space estimates are done

The space estimator bases its calculation on many factors, such as the number of files, size of files, length of directory names, exclusion lists, ACLs (access control lists), and data streams on a system. The space estimator scans the backup path of a system to obtain the values for some of these factors such as number of files, while the values for other factors such as ACLs and data streams are obtained from the values specified in a configuration file.

Degree of accuracy in the reported space estimates

Although the space estimator provides fairly precise results for the amount of free disk space on a system, some of the values taken into consideration for calculation are obtained from a configuration file and not scanned from the system on which the space estimator is run. Therefore, these values are not an absolute reflection of the disk space consumption on a system and might introduce a small degree of inaccuracy in results reported.

Ways to use the space estimator

The Open Systems SnapVault space estimator can be used in two ways:

- Built-in

  In this case, the space estimator runs in the background, at the start of a transfer, and reports whether there is sufficient space to back up based on the current Open Systems SnapVault configuration. The reported results are written to a log file called snapvault in the *install_dir*/etc directory.

  Even if insufficient space is found, a backup operation is *not* aborted by default. However, you can set an option to abort operations. See "Failing a backup if insufficient disk space is found" on page 119.

- Standalone

  In this case, the space estimator is installed as a standalone application on a system that might or might not have an existing Open Systems SnapVault installation.

If Open Systems SnapVault is installed on the system, the space estimator uses the current Open Systems SnapVault configuration to determine disk space.

If Open Systems SnapVault is not installed on the system, the estimator provides recommendations on the directories where Open Systems SnapVault and its database and temporary files can be installed.

The following is a sample output of the estimator report in which Open Systems SnapVault is not installed on a system:

```
Scanning system volumes...
Volume 'A:\' type Removable  Free Space 0%
Volume 'C:\' type Normal NTFS Free Space 2%
Volume 'D:\' type CDRom  Free Space 0%
Volume 'E:\' type Normal FAT32 Free Space 99%
Volume 'G:\' type Normal NTFS Free Space 55%

Analyzing space requirements...
'E:\' is suitable for 'Database requirements'
'E:\' is suitable for 'Temporary space requirements'
Estimator has found sufficient space for backup
```

Configuration files required for space estimator operation

The space estimator requires the following two files to estimate free disk space on a system:

- A configuration file called estimator.cfg

  This file contains user-defined options that are taken into consideration when estimating free disk space. See "Contents of the estimator.cfg file" on page 113 for more information.

- The path and file exclusion list files

  See "Configuring backup exclusion lists" on page 62 for more information.

On a system on which Open Systems SnapVault is installed, the estimator.cfg file is present in the snapvault/config directory and the exclusion list files are present in the *install_dir*/etc directory by default. If any of the files is missing, the space estimator uses the predefined default values for the missing information.

On a system on which Open Systems SnapVault is not installed—that is, you are using the standalone space estimator—you must create an estimator.cfg file (a sample of which is provided in the following section) and the two backup exclusion list files, mentioned previously in this section, in the directory in which the standalone space estimator is run.

Contents of the estimator.cfg file: The following is a sample estimator.cfg file:

```
# Sample Estimator configuration file.

# In standalone estimator mode, this file must be placed in a
#'config' directory within the current working directory.

# (where the estimator is being executed)

#

[ESTIMATOR:BLI enabled]

value = TRUE

[ESTIMATOR:Hist Data enabled]

value = FALSE

[ESTIMATOR:VSS/OFM required % disk space]

value = 15

[ESTIMATOR:Average number of streams per entity]

value = 0

[ESTIMATOR:Average stream size (KB)]

value = 1

[ESTIMATOR:Average stream name length]

value = 6

[ESTIMATOR:Average ACL size]

value = 200

[ESTIMATOR:Average OSSV installaton size (KB)]

value = 12000
```

The following table describes the fields in the estimator.cfg file and the default value associated with each field.

Field	Default Value	Description
BLI enabled	TRUE	◆ This field is ignored during the built-in operation as this information is obtained from the existing Open Systems SnapVault configuration on the system where the space estimator is running. ◆ During the standalone operation, this field enables the space estimator to include checksum sizes in its calculations.
Hist Data enabled	FALSE	◆ This field is ignored during the built-in operation as this information is obtained from the existing Open Systems SnapVault configuration on the system where the space estimator is running. ◆ During standalone operation, this field enables the space estimator to include Redundant Array of Independent Disks (RAID) checksum sizes and ACL sizes in its calculations.
VSS/OFM required % disk space	15	This field specifies the percentage of disk space that the space estimator will *not* take into consideration when determining the amount of free disk space for a backup operation. This space is assumed to be reserved for taking snapshots on the detected drives and therefore not available for backup operation.

Field	Default Value	Description
Average number of streams per entity	0	The space estimator does not examine alternate data streams for files and directories on a system. Instead it uses the value configured for this field when calculating free disk space. The value for this field applies to both files and directories.
Average stream size (KB)	1 KB	This field specifies the average size of a data stream if alternate data streams are present on a system. This field applies to all files and directories.
Average stream name length	6 characters	This field specifies the length of an alternate-data-stream name. The length of a stream name affects the Open Systems SnapVault database size, therefore, this field value is taken into consideration when calculating free disk space.
Average ACL size	200 bytes	This field is used only if the space estimator is being run on a platform that supports ACLs. If this field is set, the space estimator assumes an ACL for every file and directory within the backup.
Average OSSV installation size (KB)	12000 KB (12 MB)	This field specifies the amount of disk space consumed by Open Systems SnapVault installation without the Open Systems SnapVault database. This value is used by the space estimator when the `-i` option is specified in standalone mode.

Logs where space estimator information is written

When space estimator is run in standalone mode on a system, it displays information about all drives scanned and free space available on the console of the system. The following is a sample of the space estimator console output when run in standalone mode:

```
Scanning system volumes...
```

```
Volume 'A:\' type Removable  Free Space 0%
Volume 'C:\' type Normal NTFS Free Space 2%
Volume 'D:\' type CDRom  Free Space 0%
Volume 'E:\' type Normal FAT32 Free Space 99%
Volume 'G:\' type Normal NTFS Free Space 55%
Analyzing space requirements...
Estimator has found sufficient space for backup
```

If the space estimator is run with the debug trace in standalone mode, this information is also written to a debug trace file. Therefore, the console output provided in the previous sample is also written to a trace file.

If the space estimator is run in the built-in mode on a system, it does not display information on the console but writes the information to the snapvault log files and the trace file (if enabled). The snapvault log files contain all the information that is displayed in standalone mode except the drive scan results. For more information about the snapvault log files, see "Locating status and problem reports" on page 82).

The following is a sample of the information that will be written to a snapvault log file for the example discussed in this section:

```
Estimator has found sufficient space for backup
```

Installing the space estimator

The built-in space estimator is automatically installed when the Open Systems SnapVault software is installed on a system.

To install the standalone space estimator, complete the following step.

Step	Action	
1	**If...**	**Then...**
	The Open Systems SnapVault software is installed on your primary storage system	**1.** Go to the *install_dir*/bin directory and locate the file called svestimator. You can run the standalone space estimator from this directory or move the file to another location of your choice.
	The Open Systems SnapVault software is not installed on your primary storage system	**1.** Follow the instructions in Chapter 2, "Installing the Open Systems SnapVault Software," on page 13 to download the installation package from the NOW site or locate the installation directory specific to your platform on the CD-ROM you use to install Open Systems SnapVault agent. **2.** Uncompress or unzip the installation package, if needed. **3.** Locate the svestimator file in the installfiles directory.

Running the space estimator

In built-in mode: The space estimator in built-in mode is enabled by default and runs at the start of each transfer. Therefore, no action is required on your part.

In standalone mode: To run the space estimator in standalone mode, complete the following step.

Step	Action
1	Enter the following command at the command prompt of your primary system where you installed the space estimator: **`svestimator [-o -i -d] root_backup_path`** `-o` is required when the space estimator is run on a system on which the Open Systems SnapVault software is installed. Doing so ensures that the existing Open Systems SnapVault configuration values are used for calculating disk space. Note that the `-o` option does not take into consideration the disk space consumed by the Open Systems SnapVault installation. `-i` is required when you want the space estimator to take into consideration the amount of disk space consumed by the Open Systems SnapVault installation on the primary system. `-d` is required if you want to the debug trace information to be written to a log. If Open Systems SnapVault is installed on the system, the debug trace called svestimator.txt is written to the *install_dir*/trace directory. If Open Systems SnapVault is not installed, a directory called trace is created in the current directory and the trace is written to it. **Note** You must enable the “Generate debugging files” option in the General tab of the `Configurator` utility before using the `-d` option of the `svestimator` command. If the “Generate debugging files” option is not enabled, the log file will not be written to the trace directory. *root_backup_path* specifies the directory you want to back up; for example, C:\MyData\MyDocs. You can specify more than one path. If you do specify more than one path, the space estimator calculation takes into account the combined size of all backups.

Disabling the built-in space estimator

To disable the built-in space estimator, complete the following step.

Step	Action
1	In the SnapVault tab of the `Configurator` utility, uncheck the “Run estimator before each backup” option.

Failing a backup if insufficient disk space is found

By default, backups in built-in mode are not aborted even if insufficient space is found by the space estimator. However, you can configure the space estimator to fail backups. To do so, complete the following step.

Step	Action

1	In the *install_dir*/config directory, set the value of the following entry in the snapvault.cfg file to True:
	`[Configuration:Estimator can fail backup]`
	`Value=True`

The Logical Replication (LREP) tool A

What the Logical Replication (LREP) tool does

The Logical Replication (LREP) tool enables you to write the initial transfer (baseline) data from an Open Systems SnapVault primary system to a portable device or a system such as a NetApp FAS 250. You can then ship the portable device to the location where the secondary system is (data center), move the data from the portable device to the secondary system, and modify the Open Systems SnapVault relationship between the primary and secondary storage systems.

Need for the LREP tool

The baseline transfers can have a crippling effect on the limited bandwidth Wide Area Network (WAN) connections of remote/branch offices. Using the LREP tool, no network bandwidth is used.

As shown in the following illustration, the baseline data is transferred from the primary system to the secondary system using LREP. After the baseline data has been transferred over to the secondary system, the relationship between the primary and secondary storage system is established so incremental transfers can occur directly.

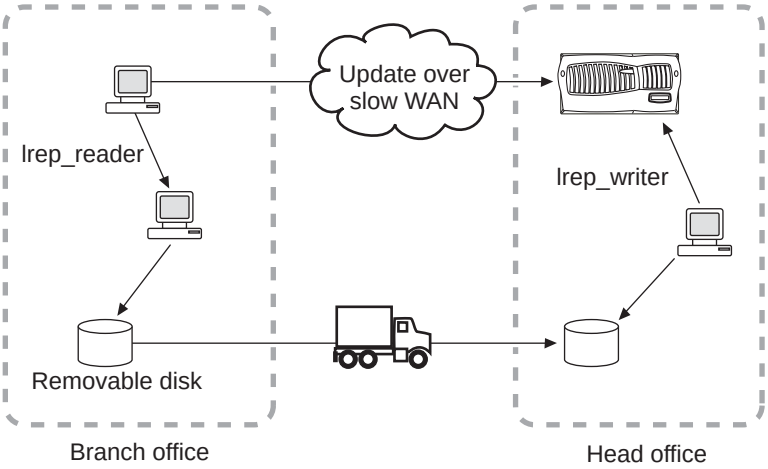

LREP supported platforms

The LREP tool can be run on Windows XP, Windows 2003, and Windows 2000 systems. However, you can use the LREP tool to write data from any of the currently supported Open System SnapVault platforms—for example, you can write data from an AIX system to a NetApp FAS 250 using LREP that is running on a Windows XP system.

Downloading LREP

The binary file for the LREP tool is packaged with the Open Systems SnapVault software. You can access the binary file from the same location where you uncompressed the Open Systems SnapVault software package on your primary system.

Additionally, you can download the LREP tool from the NOW site in one of the following two ways:

- Go to the Download Software page at http://now.netapp.com/NOW/cgi-bin/software/ and follow directions to go to the Download page of Open Systems SnapVault package for your platform. The LREP tool binary file is located on that Download page.
- Go to http://now.netapp.com/eservice/toolchest page and locate the LREP tool.

Installing the LREP tool

Because the LREP tool is a standalone binary file, it does not need to be installed and can be run from the system the binary file resides on.

How the LREP tool works

The LREP tool consists of two utilities:

- `lrep_reader`—Used at the remote office to write data from the Open Systems SnapVault primary to the portable device.
- `lrep_writer`—Used at the location of secondary storage system to write data from the portable device to the secondary system.

  If the `lrep_writer` utility is installed on a machine on which Open Systems SnapVault software is installed, you must stop the Open Systems SnapVault service before using the utility.

Example

A Windows server WIN1 is located in a remote office of an enterprise. A volume `vol1` of a secondary system, SEC1 (located in the data center), will be used to back up data from WIN1. Because the remote office has a small bandwidth WAN connection to the data center, the system administrator uses the LREP tool for the baseline transfer.

Using the `lrep_reader`, the baseline transfer is done from WIN1 to a Zip drive (drive letter E) connected to a Windows XP client, XP1. The Zip drive is shipped to the data center where a second Windows XP client, XP2, uses the `lrep_writer` to write data from the Zip drive to SEC1.

At the remote office: Enter the following command from the directory on XP1 in which the LREP tool is installed:

C:\>**lrep_reader -p snapvault_start -f SEC1 -q /vol/vol1/backup -o E:\test@0 WIN1:D:\dataset**

`-p snapvault_start`—SnapVault is used.

`-f SEC1`—SEC1 is the final destination.

`-q /vol/vol1/backup`—Full path on the final destination.

`-o E:\test@0`—The Zip drive (E:), the name of the file (test) that will be created on the drive, the number of files (@0) of size 2 GB that will be created. The number zero indicates that an infinite number of files can be created. You can span multiple drives using this feature. If your portable drive is smaller than the total amount of data that needs to be stored, you can connect multiple drives to store all the data. For example, if your portable drive is 8 GB and your data is 12 GB, you can connect two drives (E:\ and F:\) and enter the following in place of the `-o` option in the previous command:

-o E:\test@4 -o F:\test@0

`WIN1:D:\dataset`—The source that you want to back up.

At the data center: After the Zip drive arrives at the data center, connect it to the second Windows XP client, XP2, and enter the following command:

Note

Make sure that Open Systems SnapVault software is not running on the machine on which you are going to run the `lrep_writer` command.

C:\>**lrep_writer -p snapvault_start E:\test**

Now, start the data transfer from SEC1 by entering the following command:

SEC1> **snapvault start -S XP2:E:\test /vol/vol1/backup**

Modify the relationship on SEC1 by entering the following commands:

SEC1> **snapvault modify -S WIN1:D:\dataset /vol/vol1/backup**

To make sure that a relationship between WIN1 and SEC1 has been established, force an incremental backup and check status thereafter by entering the following commands:

SEC1> **snapvault update /vol/vol1/backup**

SEC1> **snapvault status**

Open Systems SnapVault Error Messages B

List of error messages

The following table lists the frequently encountered Open Systems SnapVault error messages, their causes, and actions you can take to remedy those errors.

Note

The suggested actions in the “Actions to take” column of the following table assume that you are running Open Systems SnapVault 2.2 on your primary storage system.

Error Message	Cause	Action to take
<SRC> <DEST> Abort (could not read from socket)	This message can occur because of one of the following problems: ◆ The `snapvault abort` command was issued via the command line ◆ The Open Systems SnapVault system or the secondary system has failed or rebooted ◆ A network error was encountered	Check the following logs for more information: ◆ The SnapMirror log on the secondary system ◆ The SnapVault log on the primary system
<SRC> <DEST> Abort (replication destination failed to store entry in inode map)	Usually caused by a Data ONTAP or Open Systems SnapVault bug.	Make sure you are running the latest recommended Data ONTAP release on your secondary system.

Error Message	Cause	Action to take
<SRC> <DEST> Abort (replication destination does not have a directory that the source has modified)	Usually caused by a Data ONTAP or Open Systems SnapVault bug.	Make sure you are running latest recommended Open Systems SnapVault release on your primary system and Data ONTAP release on your secondary system.
snapvault: destination transfer from <SRC> to <DEST> : the qtree is not the source for the snapmirror destination Transfer aborted: the qtree is not the source for the snapmirror destination.	This error is often caused by mistyping. Sometimes this error also occurs if a database move fails. (If the Open Systems SnapVault database is moved without following the correct procedure, it is possible for a transfer to occur to the old database, thus causing the relationship to be out-of-sync.)	Ensure that you follow the correct procedure to move the Open Systems SnapVault database.

Error Message	Cause	Action to take
can't connect to ndmp server <SERVER>. (port 10000)	Usually occurs when TCP/IP port 10000 (required by Open Systems SnapVault) is being used by another process.	◆ Stop the process using port 10000 OR ◆ Change the TCP/IP port Open Systems SnapVault uses for NDMP connections by changing the NDMP Port value in the Open Systems SnapVault Configurator (svconfigurator). For information on how to change the port, see "Modifying NDMP Listen Port setting" on page 55.
Open Systems SnapVault encountered a network error while reading (writing) data.	This error can be caused by network problems or the secondary system aborting the transfer.	Check the log files on secondary for more information.
Transfer aborted: destination qtree is not coalesced	The secondary qtree is still in a transitioning state.	Make sure the destination qtree is in an IDLE state for an update operation to be successful.

Error Message	Cause	Action to take
Current Transfer Error: replication destination cannot find a file for which the source sent data SnapMirror destination transfer from <SRC> to <DEST> : replication destination cannot find a file for which the source sent data.	Usually caused by a Data ONTAP or Open Systems SnapVault bug.	Make sure you are running the latest recommended Open Systems SnapVault release on your primary system and Data ONTAP release on your secondary system.
snapmirror: Message from Read Socket : Connection reset by peer snapvault: destination transfer from <SRC> to <DEST> : source volume is offline, is restricted, or does not exist.	This is often caused when Open Systems SnapVault is not running on the host machine at the time of backup. It may also be caused by mistyping or a very busy network.	Check the SnapMirror log on the secondary system and the SnapVault log on the primary system for errors and retry the transfer.
Data Ontap changed data in <X> data chunk(s) in <DEST> which may have been missing or incorrect.	This error might be caused when Data ONTAP detects that incomplete Open Systems SnapVault block-level incremental backups have occurred.	See bugs 137685 and 140930 at http://now.netapp.com for more information.
destination transfer from <SRC> to <DEST> : qtree snapmirror destination found a mismatch between a directory entry and its inode information	Usually caused by a Data ONTAP or Open Systems SnapVault bug.	Make sure you are running latest recommended Open Systems SnapVault release on your primary system and Data ONTAP release on your secondary system.

Error Message	Cause	Action to take
replication destination failed to store entry in inode map	Usually caused by a Data ONTAP or Open Systems SnapVault bug.	Make sure you are running latest recommended Open Systems SnapVault release on your primary system and Data ONTAP release on your secondary system.
OFM Volume cannot be sync'ed - volume error. (1)	Caused by OFM not being able to create a snapshot of the drive which holds the data for the relationship. (OFM requires a quiescent period of file system inactivity to perform a snapshot.)	Edit the following settings in the Open Systems SnapVault Configurator (svconfigurator) and retry the transfer: ◆ Set “Write inactivity period (seconds)” from the default of 5 to 2. ◆ Set “Synchronization timeout (seconds)” from the default of 60 to 120.
Cannot backup 'c:\' Failed to create volume snapshot	Caused by OFM not being able to create a snapshot of the drive which holds the data for the relationship. (OFM requires a quiescent period of file system inactivity to perform a snapshot.)	Edit the following settings in the Open Systems SnapVault Configurator (svconfigurator) and retry the transfer: ◆ Set “Write inactivity period (seconds)” from the default of 5 to 2. ◆ Set “Synchronization timeout (seconds)” from the default of 60 to 120.
Unexpected close getting QSM data	The Open Systems SnapVault services have been either started or stopped.	If the Open Systems SnapVault services were not started or stopped intentionally, check the SnapMirror log on the secondary system and the SnapVault log on the primary system for more information

Error Message	Cause	Action to take
Failed to write hybrid history record	Open Systems SnapVault has failed to write to its database. This can be caused by a lack of disk space or a disk error.	Check the amount of disk space remaining on the partitions containing the Open Systems SnapVault installation, Open Systems SnapVault database, and the Open Systems SnapVault temporary directory. Check the SnapMirror log on the secondary system and the SnapVault log on the primary system for errors
snapvault: destination transfer from <SRC> to <DEST> : incompatible snapmirror versions on systems	Due to an ordering error in the destination-side error handling, it is possible for SnapVault and qtree SnapMirror transfers to fail, but produce an error message that does not correspond to the error reported on the source.	See bug 147982 at http://now.netapp.com for more information.
Unexpected read select while no data pending	This message can occur because of one of the following problems: ◆ The `snapvault abort` command was issued via the command line ◆ Open Systems SnapVault machine failing/rebooted; network error.	Check the SnapMirror log on the secondary system and the SnapVault log on the primary for errors and then retry the transfer.

Error Message	Cause	Action to take
transfer from <SRC> to <DEST> : request denied by source filer; check access permissions on source	The secondary system has denied the request from Open Systems SnapVault to perform the operation.	On the Open Systems SnapVault machine, check whether the QSM Access List in the Open Systems SnapVault Configurator contains the IP address or network name of the secondary system. For more information, see "Modifying the QSM access list" on page 53.
Current Transfer Error: unable to translate Unicode pathname, please check volume language configuration	Unicode pathname that cannot be translated abort Open Systems SnapVault backup.	You can work around this problem by changing the volume language to a character set that uses UTF-8; for example, C.UTF-8. See bug 133965 at http://now.netapp.com for more information.

Symbols

.OSSV_DATABASE_BACKUP file 84

Numerics

32-bit systems 14
64-bit systems 15, 16

A

aborted transfer, retrying 109
access list (QSM) 53
ACL file 83
administration interfaces 5
agent, Open Systems SnapVault 3
Allowed Secondary Names window 21
Always setting (for generating debugging information) 57
applications for managing Open Systems SnapVault 6

B

backup
    database, how to 85
    disable database 87
    EFS-encrypted files 11, 110
    open files 8
    Open Systems SnapVault database 10
    scheduling 73
    System State data 91
    using Open Systems SnapVault for 79
    Windows System State 10
backup exclusion lists 9, 62
backup file, database (OSSV_DATABASE_BACKUP) 84
baseline copy, creating 72
batch installation 41
BLI
    checksums file 83
    described 9
    for EFS-encrypted files 9
    for name-based applications 9
    level, setting 60
    name-based 61
block-level incremental. *See* BLI
broken relationship
    restoring 107
    resuming 11
built-in space estimator
    described 111
    disabling 119
    enabling 117
    installing 117

C

calculation of free space 111
checkpoint
    file 83
    restart 11
checksum computation, BLI checksum 60
command-line (CLI) utility for administration 47
command-line utility (svsetstanza) 5, 47
commands
    lrep_reader 122
    lrep_writer 122, 123
    restore 79
    snapvault modify 109
    snapvault release 97, 105
    snapvault restore 6, 75, 85, 87
    snapvault snap sched 73
    snapvault start 5, 72, 91, 97, 107
    snapvault stop 97
    svconfigpackager 37
    svestimator 118
    svinstallcheck 31
    svpmgr shutdown 32
    svpmgr startup 34
    svsetstanza 5, 47
configuration file
    configure.cfg 47
    estimator.cfg 112
Configurator
    General tab, purpose of 45
    Machine tab, purpose of 45
    running 50
    Service tab, purpose of 45
    SnapVault parameter, modifying 53
    SnapVault tab, purpose of 46

Trace tab, purpose of 46
utility, described 45
configure.cfg file 47
configuring, Open Systems SnapVault 5
copying files to restore 75

D

Data ONTAP version, supported 16
database
backup 85
disabling backup 87
disk space requirements for 15
files, ACL 83
files, backup file (OSSV_DATABASE_BACKUP) 84
files, BLI checksums 83
files, checkpoint 83
files, described 83
files, history 83
Open Systems SnapVault 10, 83
DataFabric Manager, using to manage Open Systems SnapVault 6
debugging files
generation of 56
stop generating 58
determining free space 111

E

EFS-encrypted files
about 11
BLI and 9
error messages 125
estimating, free space 11
estimator.cfg file
contents of 113
described 112
exclusion lists 9
EXT2 file system 16
EXT3 file system 16

F

failed transfers, retrying 109
file exclusion list
described 9
for space estimator 112
location of 62
file systems, supported 16
free space estimator. *See* space estimator

G

General tab (Configurator) 45
Generate debugging files option 56
GUI interface for administration 5, 45

H

HFS file system 16
history file 83

I

initial backup (baseline) 72
installation
Allowed Secondary Names window 21
batch 41
built-in space estimator 117
limitations 16
on HP-UX, AIX, IRIX or Linux 27
on Solaris 24
on UNIX and Linux 24
on Windows 18
requirements, primary systems 14
script for unattended install 37
Select Installation Folder window 22
Setup wizard 19
standalone space estimator 117
unattended, about 37
upgrading. *See* upgrade
interfaces
for administering 5
Setup wizard 19

J

JFS file system 16

L

Libnormal setting 57
Libverbose setting 57
licenses
    for Open Systems SnapVault 16
    OFM 65
    sv_linux_pri 16
    sv_ontap_sec 16
    sv_unix_pri 16
    sv_windows_pri 16
limitations of Open Systems SnapVault 16
listen port, NDMP 20
lists, exclusion 9
log files
    SnapVault 82
    space estimator 116
    trace 116
log files, SnapVault 45
Logical Replication (LREP) tool, described 121
LREP tool
    described 121
    downloading 122
    installation 122
    supported platforms 121
lrep_reader utility 122
lrep_writer utility 122, 123

M

Machine tab (Configurator) 45
management application, Open Systems SnapVault 6
messages, error 125
migrate a volume 98

N

name-based
    applications, about 9
    block-level incremental, setting 61
NDMP Account field 54
NDMP Host ID field 54
NDMP Host Name field 54
NDMP Listen Port field 20, 54
NDMP-based management applications 6
NTFS file system 16

O

OFM
    about 8, 65
    configuring 65
    disabling 67
    parameters, setting 66
    St. Bernard's module 18
open file manager. *See* OFM
open files, backing up 8
Open Systems SnapVault
    about 2
    database backup 10
    database disk space requirements 15
    how it works 3
    interfaces 5
    licenses 16
    operation 3
    software 3
    systems, primary and secondary 3
OSSV. *See* Open Systems SnapVault
OSSV_DATABASE_BACKUP file 84

P

parameters modifying, SnapVault 53
path exclusion list
    about 9
    described 62
    for space estimator 112
    location of 62
port
    10000 15, 20, 31, 127
    10566 15, 31
    listen, NDMP 20
preinstallation requirements 14
prerequisites
    for installation 14
    for upgrading 32
primary system
    about 3
    administration 5
    controlling secondary system access 53
    installation requirements 14

restoring 78
restoring data 76
svrestore command 76

Q

Qtree SnapMirror (QSM) access list 53

R

relationship
broken 107
deleting 97
migration 98
resynchronizing 107
remote installation
and upgrade 37
batch 41
reports, problem 82
requirements (installation), for primary systems 14
restart backup from a checkpoint 11
restore
a complete system 78
broken relationship 107
by copying files 75
EFS-encrypted files 110
from tape 79
methods for 75
System State data 92
Windows Systems State 10
restore command 79
resume broken relationship 11
resynchronize a relationship 11, 107
retry failed transfer 109

S

schedule backups 73
script for unattended install 37
secondary system
about 3
access control 53
administration 5
configuring for use with open systems 71
enabling SnapVault 71
installation requirements 16
migrating relationship between 97, 98
security, enabling or disabling 53
Select Installation Folder window 22
Service tab (Configurator) 45
Setup Wizard 19
SnapMirror, using for protecting 104
SnapVault
configuring the secondary storage system 71
log files 45, 82
primary system parameters, modifying 53
snapvault modify command 109
snapvault release command 97, 105
snapvault restore command 6, 75, 85, 87
snapvault snap sched command 73
snapvault start command 5, 72, 91, 97, 107
snapvault stop command 97
SnapVault tab (Configurator) 46
snapvault.cfg file (for space estimator), modifying 119
softlock support 104
space estimator
about 11
built-in
described 111
disabling 119
installing 117
calculates, how it 111
console output 116
described 111
enabling 117
estimator.cfg file 112
log files 116
path and file exclusion lists 112
standalone 111, 117
utility 15
St. Bernard’s OFM module 18
standalone space estimator
described 111
installation 117
status logs 82
sv_linux_pri license 16
sv_ontap_sec license 16
sv_unix_pri license 16
sv_windows_pri license 16
svconfigpackager command (utility) 37

svestimator command 118
svinstallcheck command 31
svpmgr shutdown command 32
svpmgr startup command 34
svsetstanza command 5, 47
system restore, primary 78
System State data
    backup 91
    restore 92

T

tape, restore from 79
tertiary system, data protection with 104
tombstone lifetime setting 91
trace file 116
Trace Level tab (Configurator) 46
troubleshooting 125
troubleshooting, generating debugging files for 56

U

UFS file system 16
unattended installation 37
uninstall, how to 35
unsupported configurations 16
upgrade
    from versions 1.x
    from versions 2.x 34
    prerequisites for 32
    unattended 37

V

Verbose setting 57
volume migration 98
VSS snapshot
    about 8, 65
    setting timeout for 67
VxFS file system 16

W

web interface for administration 5, 45
Windows System State 10
wizard, setup 19

X

XFS file system 16