

Open Systems SnapVault® 2.3

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 215-02042_A0
October 2006

Copyright and trademark information

Copyright information

Copyright © 1994–2006 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.

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

THIS SOFTWARE IS PROVIDED BY NETWORK APPLIANCE “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 NETWORK APPLIANCE 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.

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, Data ONTAP, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the U.S.A. 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, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapDrive, SnapFilter, SnapMigrator, SnapSuite, SnapValidator, SohoFiler, SpinAV, SpinManager, SpinMirror, SpinRestore, SpinShot, SpinStor, 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 U.S.A.

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	vii
Chapter 1	Introduction to Open Systems SnapVault	1
	About Open Systems SnapVault	2
	Open Systems SnapVault features	9
Chapter 2	Installing the Open Systems SnapVault Software.	13
	Prerequisites	14
	Installing Open Systems SnapVault on Windows platforms	19
	Installing Open Systems SnapVault on UNIX and Linux platforms	26
	Verifying the installation	34
	Upgrading to Open Systems SnapVault 2.3	35
	Uninstalling the Open Systems SnapVault 2.3 agent.	38
	Unattended installation and upgrade	41
Chapter 3	Configuring Open Systems SnapVault.	47
	Configuration interfaces	48
	Understanding the Configurator utility interface.	49
	Understanding the svsetstanza command	51
	Running the Configurator utility	54
	Confirming that services are running.	56
	Modifying Open Systems SnapVault parameters	57
	Enabling and disabling debugging	60
	Setting block-level incremental backup options	64
	Configuring backup exclusion lists.	66
	Configuring open file backup for Windows	69
Chapter 4	Setting Up Backups and Restores	73
	Setting up SnapVault backup on Open Systems platforms.	74

	Configuring the SnapVault secondary storage system	75
	Creating an initial baseline copy	76
	Scheduling SnapVault update backups.	77
	Restoring data.	78
	Restoring a directory or a file	79
	Restoring an entire primary storage system	82
	Restoring files to a primary storage system from tape.	83
Chapter 5	Open Systems SnapVault Management	85
	Locating status and problem reports	86
	Backing up and restoring the Open Systems SnapVault database	87
	Backing up and restoring Windows System State data.	91
	Adding System State data backup	94
	Restoring System State data	95
	Using System State backup to rebuild a primary system	97
	Deleting and re-creating Open Systems SnapVault relationships	100
	Migrating a relationship between two secondary systems	101
	Migrating between two volumes on one secondary system	104
	Setting up a tertiary system for a relationship	107
	Reusing a deleted or renamed directory name	109
	Resynchronizing restored or broken relationships	111
	Retrying failed transfers	113
	Encrypted File System (EFS) file backup and restore	114
Chapter 6	Open Systems SnapVault Space Estimator	115
Appendix A	The Logical Replication (LREP) Tool	125
Appendix B	Open Systems SnapVault Error Messages.	129
	Index	139

Preface

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

Audience This guide is for system administrators who are familiar with operating systems such as UNIX, Windows™ 2000 Server, 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 does not 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 storage system commands on the system console or from any client that can obtain access to the storage system by using a Telnet session. In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ from what you see, depending on the version of UNIX that you are using.

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

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fc test adaptername</code> command, you enter the characters “fc test” 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.
<i>Enter</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.
<i>enter</i>	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.

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 9

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®
- ◆ Linux®
- ◆ Windows®
- ◆ Solaris™

Note

For a list of currently supported versions of these platforms, see “[Requirements for primary storage systems](#)” on page 14.

Components of the 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 SnapVault secondary storage systems](#)” on page 17.

How Open Systems SnapVault works

You identify the directories or file systems 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 be either initiated manually or configured for automatic, unattended transfer (scheduled). You specify the schedules on the secondary system by using the commands available in Data ONTAP or by 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 Snapshot™ copies that capture the changes to the secondary qtrees after each transfer.

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

If directory or file data needs to be restored to the primary system, SnapVault retrieves the data from the specified Snapshot copy 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 storage systems, 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 by 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 the 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 49.

- ◆ 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 51.

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 by 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 80.

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

Vendor	Product	URL
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>.
 - ❖ Network Appliance Technical Report TR-3466—Open Systems SnapVault Best Practices Guide, 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 the space estimator](#)”
- ◆ “[Resynchronizing a broken relationship](#)”

Backing up open files

Typically, if files on an Open Systems platform directory are open when a scheduled SnapVault transfer takes place, they are not backed up until the next scheduled transfer. However, you can configure the system for open file backup, during which Open Systems SnapVault uses one of the following options:

- ◆ Open File Manager (OFM)—Used for 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 17.
- ◆ VSS Snapshot copy—Used for Windows 2003 platforms
The VSS Snapshot copy 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 69.

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 66.

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 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 the backup storage, as well as the amount of data that must be stored on the 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 workaround is enabled by default. To enable or disable the workaround, see [“Enabling or disabling BLI backups for certain name-based applications”](#) on page 65.

For information on how to configure or change BLI settings in Open Systems SnapVault, see [“Setting block-level incremental backup options”](#) on page 64.

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 system 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 system 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 up and restoring the Open Systems SnapVault database”](#) on page 87.

Backing up and restoring Windows System State

You can back up and restore Windows System State data by 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 91.

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 transfer will be tried is set to two, giving only one retry; however, you can change it. For more information, see [“Changing the number of retry attempts made for failed transfers”](#) on page 113.

Encrypted File System (EFS) files

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

For more information, see [“Encrypted File System \(EFS\) file backup and restore”](#) on page 114.

Determining free space using the 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 system 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 115.

Resynchronizing a broken relationship

You can resynchronize a broken or out-of-sync SnapVault relationship between a primary system and a secondary system, 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 secondary systems.

For more information, see [“Resynchronizing restored or broken relationships”](#) on page 111.

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:

- ◆ [“Prerequisites”](#) on page 14
- ◆ [“Installing Open Systems SnapVault on Windows platforms”](#) on page 19
- ◆ [“Installing Open Systems SnapVault on UNIX and Linux platforms”](#) on page 26
- ◆ [“Verifying the installation”](#) on page 34
- ◆ [“Upgrading to Open Systems SnapVault 2.3”](#) on page 35
- ◆ [“Uninstalling the Open Systems SnapVault 2.3 agent”](#) on page 38
- ◆ [“Unattended installation and upgrade”](#) on page 41

Prerequisites

Before you install

Ensure that the site meets the minimum requirements for the primary and the secondary storage systems, and that you have the correct licenses for both, before installing the Open Systems SnapVault agent. Also, ensure that you read the information in “[Limitations](#)” on page 18 and *Open Systems SnapVault 2.3 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

Type of system: You can install the Open Systems SnapVault 2.3 agent on the following systems.

Operating system	Software versions	File systems	ACLs supported
Windows 2000 and 2003 on x86 and x86-64/EM64T compatible hardware	Windows 2000 Server and Advanced Server for x86 with Service Pack 4	NTFS	Yes
	Windows Server 2003 Standard Edition and Enterprise Edition for x86 Note _____ “Windows 2003” in this document refers to Windows Server 2003 and Windows Storage Server. However, manageability is not supported on Windows Storage Server. _____	NTFS	Yes

Operating system	Software versions	File systems	ACLs supported
Linux on x86 and x86-64/EM64T compatible hardware	Red Hat® Linux 9 for x86	ext2 and ext3	No
	Red Hat Enterprise Linux 3.0 ES/AS/WS for x86 and x86-64/EM64T	ext2 and ext3	Yes, support for ext3 only
	Red Hat Enterprise Linux 4.0 AS/ES/WS for x86 and x86-64/EM64T	ext2 and ext3	
	SUSE® Linux Enterprise Server 8 for x86 and x86-64/EM64T	ext2, ext3, JFS, XFS, and ReiserFS	Yes, supported only on 32-bit kernel
	SUSE Linux Enterprise Server 9 for x86 and x86-64/EM64T	ext2, ext3, JFS, XFS, and ReiserFS	Yes
Solaris on UltraSPARC systems	Solaris 9	UFS, VxFS 3.5, and VxFS 4.0	Yes
	Solaris 10	UFS, VxFS 4.0, and VxFS 4.1	Yes
AIX on PowerPC® and POWER processor-based systems	5L versions 5.1, 5.2, and 5.3 Note _____The IBM Cluster 1600 and RS/6000 Scalable Parallel (SP) platforms are not supported. _____	JFS1 and JFS2	Yes (AIXC type ACLs on JFS1 and JFS2)
HP-UX on PA-RISC	HP-UX 11.11	HFS and JFS	Yes
	HP-UX 11.23	HFS and JFS	Yes

Note

Open Systems SnapVault 2.3 is supported on both 32-bit and 64-bit kernels on those UNIX operating systems that have the capability of booting into both these kernels. Open Systems SnapVault binaries are 32 bit in both the cases. For Windows, Open Systems SnapVault is available only on 32-bit operating systems.

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

- ◆ A minimum of 512 MB of memory
- ◆ 100Base-T or Gigabit Ethernet (GbE) network connectivity, for best performance
- ◆ Available TCP port 10566 (SnapVault)
- ◆ 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.
- ◆ For NetApp Host Agent, available HTTP port 4092 and HTTPS port 4093

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 the system. For more information about the space estimator utility, see [“Open Systems SnapVault Space Estimator”](#) on page 115.

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

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

Requirements for SnapVault secondary storage systems

To use systems installed with the Open Systems SnapVault software, the SnapVault secondary storage system must be running Data ONTAP 6.5.6 or later.

Supportability matrix of Open Systems SnapVault with NetApp Host Agent

Open Systems SnapVault 2.3 release includes NetApp Host Agent 2.3.1, providing better manageability through DataFabric Manager. During installation, Open Systems SnapVault installs the NetApp Host Agent and copies the interface libraries at the location on which NetApp Host Agent is installed. The platform support for Open Systems SnapVault with NetApp Host Agent is as follows.

Platform	Versions supported
Linux on x86	◆ Red Hat Enterprise Linux 3.0 ES (update 3 or later) ◆ Red Hat Enterprise Linux 3.0 AS (update 3 or later) ◆ SUSE Linux Enterprise Server 9
Windows on x86 and x86-64/EM64T	◆ Windows 2000 Server ◆ Windows Server 2003

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 Open Systems SnapVault primary licenses for the platforms you want to back up to the secondary storage 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
- ◆ For softlocks, SnapMirror licenses are required for the secondary and tertiary 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:

- ◆ Backup and 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 mountpoints.
You must specify each local mountpoint as a separate backup if you are backing up multiple file systems.
 - ❖ Open Systems SnapVault does not back up remote points and special mountpoints.
For example, on Windows systems, this means mountpoints that are 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.
- ◆ Open Systems SnapVault does not support ACLs on SUSE Linux Enterprise Server 8.0 (x86-64 compatible platform).

Installing Open Systems SnapVault on Windows platforms

Installing the Windows agent from CD-ROM

To install Open Systems SnapVault on the Windows primary storage system from the CD-ROM, complete the following steps.

Note

For 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 the Windows primary storage system with Administrator privileges.
2	Insert the CD-ROM labeled <i>Open Systems SnapVault Agent</i> 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 the platform: ◆ For Windows 2000, go to \WIN2K ◆ For Windows 2003, go to \WIN2K3 Depending on the 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, <code>cd \WIN2K</code>).
4	After navigating to the appropriate directory for the operating system, see “Using the Installation Wizard” to complete the installation.

Installing the Windows agent from NOW

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


Step	Action
1	Log in to the primary storage system with Administrator privileges.


Step	Action
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 the platform.
3	Unzip the downloaded package into a temporary directory on the Windows platform.
4	Navigate to the temporary directory where you unzipped the files for the system.
5	Follow the instructions in “Using the Installation Wizard” to complete the installation.

Using the Installation Wizard

To install the Open Systems SnapVault agent on the system using the Installation Wizard, 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 Next. Result: The Open Systems SnapVault Setup License Agreement window appears.
3	Click “I Agree” to accept the license agreement and click Next. Result: The Choose NDMP user name/password window appears. This user name and password are 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
4	Optional: Either type your user name and password, retype your password to confirm it, and click Next, or just 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.
5	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
6	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
7	When the Select Installation Folder window prompts you for the installation directory, either accept the default location, or enter your own path. Click Disk Cost to view the available disk space, and click Next. The screenshot shows a Windows-style dialog box titled 'OSSV' with the subtitle 'Select Installation Folder'. Inside, it says 'The installer will install OSSV to the following folder.' and 'To install in this folder, click "Next!". To install to a different folder, enter it below or click "Browse".'. There is a text field labeled 'Folder:' containing 'C:\Program Files\netapp\snapvault\' and a 'Browse...' button. Below that is a 'Disk Cost...' button. At the bottom are 'Cancel', '< Back', and 'Next >' buttons. 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, select the “Yes” radio button before you click Next. Result: The Install NetApp Host Agent window appears.

Step	Action
8	When the Install NetApp Host Agent window prompts you to choose whether to install the NetApp Host Agent, select Yes or No. If you do not want the NetApp Host Agent installed on the default location, select an alternate location on which to install the NetApp Host Agent, and click Next. The Confirm Installation window appears.  Alternatively, if you select No, NetApp Host Agent will not be installed. To install NetApp Host Agent later, complete these steps: - Using the command-line interface (CLI), go to the Open Systems SnapVault install location. - Run these commands: <pre>cd manageability InstallHostAgent.exe (for Windows)</pre>
9	Click Next to start the installation process and wait until the Installation Complete window appears with the following message: <pre>OSSV has been successfully installed. Click "Close" to exit.</pre>

Step	Action	
10	For Windows 2000 Server...	For Windows Server 2003/ Windows Storage Server...
	Reboot the system when prompted if you plan to use OFM; otherwise, go to Step 11.	Go to step 11.
11	After you finish installing the Open Systems SnapVault agent, follow the procedures in the SnapVault chapter of the Data ONTAP <i>Data Protection Guide</i> 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

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 platform.

Installing the Solaris agent from CD-ROM

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

Step	Action
1	Log in to the primary storage system as root.
2	Insert the CD-ROM labeled <i>Open Systems SnapVault Agent</i> 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 by using the following command: pkgadd -d <i>path_to_package</i>/ossv ossv <i>path_to_package</i> 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 the 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 <i>path_to_package</i>/ossv ossv <i>path_to_package</i> is the full path to the Open Systems SnapVault package—for example, /export/home/packages/ossv ossv. The Solaris pkgadd utility installs the package 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 user name 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, or Linux agent from the CD-ROM

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

Step	Action
1	Log in to the primary storage system as root.
2	Insert the CD-ROM labeled <i>Open Systems SnapVault Agent</i> into the primary storage system's CD-ROM drive.
3	From the command line, navigate to the CD-ROM directory specific to the platform: ◆ For HP-UX, go to /HPUX ◆ For AIX, go to /AIX ◆ For Linux, go to /LINUX Depending on the system configuration and mount options, the files and directories mentioned might appear in lowercase characters.
4	From the command line, enter the following command: <code>./install</code>
5	When the <code>install</code> program prompts you to read and agree to the terms of the license, answer y (yes), n (no), or d (display). If you answer n (no), the Open Systems SnapVault agent will not be installed.
6	When the <code>install</code> 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	
7	If you plan to manage Open Systems SnapVault backup through...	Then...
	The Data ONTAP command line	Press Enter.
	A commercial NDMP application, such as DataFabric Manager	When the <code>install</code> program prompts you for the user name, password, and NDMP listening port, specify a user name and password authorized through that application.
8	When the <code>install</code> 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, not spaces.	
9	In the Linux environment, when the <code>install</code> program prompts you to choose whether to install the NetApp Host Agent, enter y (yes) or n (no). If you answer n (no), NetApp Host Agent will not be installed. To install NetApp Host Agent later, complete these steps: a. Using the command-line interface (CLI), go to the Open Systems SnapVault installation location. b. Run these commands: <pre>cd manageability ./InstallHostAgent.sh (for Linux)</pre>	

Step	Action
10	Look for a message similar to the following, which is displayed to indicate a successful installation: <pre>NetApp Host Agent is installed at /opt/NTAPagent Stopping the NTAPagent daemon Starting the NTAPagent daemon Installation completed successfully.</pre>
11	After you finish installing the Open Systems SnapVault agent, follow the procedures in the SnapVault chapter of the <i>Data ONTAP Data Protection Guide</i> to configure the SnapVault secondary storage system for Open Systems SnapVault backup.

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

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

Step	Action
1	Log in to the 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: <pre>gunzip package_name tar -xvf tar_file_name</pre>
	You are installing Open Systems SnapVault on any other UNIX platform	Enter the following commands after downloading the Open Systems SnapVault package: <pre>uncompress package_name tar -xvf tar_file_name</pre> For example, on an AIX primary storage system, uncompress and untar the downloaded ossv_aix_v2.3.tar.Z package using the following commands: <pre>uncompress ossv_aix_v2.3.tar.Z tar -xvf ossv_aix_v2.3.tar</pre>
4	Navigate to the directory where the untar program placed the files and enter the following command: <pre>./install</pre>	
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 Open Systems SnapVault backup through...	Then...
	The Data ONTAP command line	Press Enter.
	A commercial NDMP application such as DataFabric Manager	When the <code>install</code> 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 <code>install</code> 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 <i>Data ONTAP Data Protection Guide</i> 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: <code>ping secondary_system</code> <i>secondary_system</i> is either the name or the IP address of the secondary storage system.
2	If security was enabled on the primary storage system during installation, ensure that the secondary storage system specified can access the primary storage system, using the procedure described in “Enabling and disabling security” on page 57.
3	Ensure that the volumes you need were created on the secondary storage system before you attempt to back up data to them.
4	If there are firewalls between the primary storage system and the secondary storage system, ensure that TCP ports mentioned in “Memory and port requirements:” on page 16 are kept open.
5	Navigate to <i>install_dir/bin</i> and run <code>svinstallcheck</code> to verify successful installation and to make sure that the services are running properly. <i>install_dir</i> is the location where you installed the Open Systems SnapVault agent.

Upgrading to Open Systems SnapVault 2.3

Prerequisites

Open Systems SnapVault 2.1 or Open Systems SnapVault 2.2 must be installed on the primary storage system.

Upgrading to Open Systems SnapVault 2.3

When upgrading Open Systems SnapVault 2.1 or 2.2 to 2.3, complete the following steps.

Step	Action
1	Stop Open Systems SnapVault services, as follows: a. Launch the Open Systems SnapVault Configurator utility. For instructions, see “Running the Configurator utility” on page 54. b. Click the Service tab. c. Click Stop Service to stop Open Systems SnapVault services. d. Close the Configurator utility. Note Alternatively, you can use either the <code>svpmgr shutdown</code> or <code>snapvault service stop</code> command.
2	Important: To ensure that the existing 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 up and restoring the Open Systems SnapVault database” on page 87.

Step	Action
3	Install the Open Systems SnapVault 2.3 agent using the procedure described in “Installing Open Systems SnapVault on Windows platforms” on page 19 or “Installing Open Systems SnapVault on UNIX and Linux platforms” on page 26. Note _____ During installation, when prompted whether you want to upgrade, select “Yes” (for Windows) or enter “Y” (for other supported operating systems) to continue with the upgrade. Result: The Open Systems SnapVault agent automatically starts Open Systems SnapVault services after installation.
4	a. Launch the Open Systems SnapVault Configurator utility. b. Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services. c. Close the Configurator utility. Note _____ Alternatively, you can use the <code>svpmgr shutdown</code> command.
5	Copy the database directory that you saved in Step 2 to the location of the database created by the new installation.
6	a. Launch the Open Systems SnapVault Configurator utility. b. Click the Service tab, and then click Start Service to start Open Systems SnapVault services. Note _____ Alternatively, you can also use the <code>svpmgr startup</code> command.

Support for ACLs: Open Systems SnapVault 2.3 supports ACLs on Linux and HP-UX JFS. For existing Open Systems SnapVault relationships (that is, for versions earlier than 2.3), after upgrade to Open Systems SnapVault 2.3, ACLs are backed up only for changed or newly added files.

To avoid such a partial backup of ACLs after upgrading to Open Systems SnapVault 2.3, add the following stanza to snapvault.cfg:

```
[QSM:EAs Updated]  
Value=FALSE
```

This causes Open Systems SnapVault to take the following actions, based on the BLI level:

- ❖ If BLI is set to HIGH, ACLs for all files in the relationship are sent to the secondary storage system during the next transfer.
- ❖ If BLI is set to OFF, the entire data set including ACLs are sent to the secondary storage system during the next transfer.

After the first update transfer is complete, either remove the stanza from snapvault.cfg or set the value to TRUE.

Uninstalling the Open Systems SnapVault 2.3 agent

Uninstalling the 2.3 agent on Windows

To uninstall the Open Systems SnapVault 2.3 agent on the Windows platform, complete the following steps.

Step	Action
1	Click Start > Control Panel.
2	Double-click Add or Remove Programs.
3	Select OSSV in the list of programs and click Remove. The Welcome to the OSSV Removal Wizard appears. 
4	To uninstall NetApp Host Agent, select Yes. Note— If you want to uninstall NetApp Host Agent later, select No. You can uninstall NetApp Host Agent later from the Add/Remove Programs Wizard.

Step	Action
5	Click Finish to remove Open Systems SnapVault from the computer. Result: Windows uninstalls the Open Systems SnapVault agent and displays the following message to inform you that the uninstallation is successful: OSSV has been successfully removed. Click "Close" to exit.

Uninstalling the 2.3 agent on Solaris

To uninstall the Open Systems SnapVault 2.3 agent on a Solaris platform, complete the following steps.

Step	Action
1	Log in to the primary storage system as root.
2	Run the pkgrm utility: pkgrm ossv
3	Enter Yes (y) when asked whether you want to remove this package. The script responds with the following text: <pre>## Removing installed package instance <ossv> This package contains scripts which will be executed with super-user permission during the process of removing this package.</pre>
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: <pre>Removal of <ossv> was successful.</pre>

Uninstalling the 2.3 agent on HP-UX, AIX, or Linux

To uninstall the Open Systems SnapVault 2.3 agent on the HP-UX, AIX, or Linux platform, complete the following steps.

Step	Action
1	Enter the following command: <install_dir>/util/uninstall
2	During uninstallation on Linux, when prompted whether to uninstall NetApp Host Agent, enter y (yes). This prompt appears only if NetApp Host Agent was installed using Open System SnapVault.
3	If the <code>snapvault</code> directory still appears after running <code>uninstall</code> , 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 45.

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

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: a. Launch the Open Systems SnapVault Configurator utility. For instructions, see “Running the Configurator utility” on page 54. b. Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services.
2	Configure all parameters you want the new installation to have, using the Configurator utility.
3	Close the Configurator utility.
4	Navigate to the <i>install_dir/bin</i> directory of the primary system.

Step	Action	
5	If...	Then...
	You only want to save the configuration settings to a file. For details on the installation script, see “Unattended installation script” on page 46.	Enter the following command: <pre>svconfigpackager filename</pre> <i>filename</i> is the name of the configuration settings file. Example: To create a configuration settings file called <code>svconfig.in</code>, enter the following command: <pre>svconfigpackager svconfig.in</pre>
	You want to save the configuration settings to a file and also create an installation script for installers The installation script will be <i>unattinstall.bat</i> (for Windows) and <i>unattinstall.sh</i> (for UNIX). The configuration settings file, the install script, and the response file (for UNIX) will be located at the install root directory for Open Systems SnapVault.	Enter the following command: <pre>svconfigpackager [-h -i "<installation path>"] <package name></pre> <i>installation path</i> is the Open Systems SnapVault installation directory for the unattended installs. <i>package name</i> is the name of the configuration settings file. Example: To create a configuration settings file called <code>svconfig.in</code> and save the installation script in the <code>/usr/snapvault</code> directory, enter the following command: <pre>svconfigpackager -i /usr/snapvault svconfig.in</pre>
	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: <pre>svconfigpackager -h -i path_name filename</pre> <i>path_name</i> is the directory where the configuration settings file and the installation script will be placed.

Step	Action
6	After the files in the previous step are created, you see a message at the command prompt that lists all the files placed in the directory you specified. Example 1: The following is displayed on a Solaris primary system: <pre>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)</pre> Example 2: The following is displayed on a Windows primary system: <pre>The following files have been placed in 'C:/Program Files/netapp/snapvault': 'mypackage.in' (Configuration Package) 'unattinstall.bat' (Unattended install batch file)</pre> Result: You are done preparing the installation script and the configuration files that will be required 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 “ Performing an unattended installation or upgrade ” on page 44 to the directory on the primary system where you uncompressed the installation package.

Step	Action	
3	If...	Then...
	The primary system is a Windows machine	Enter the following command to start the unattended installation or upgrade: 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
	The primary system is a UNIX machine	Enter the following command to start the unattended installation or upgrade: ./unattinstall.sh If the installation does not succeed, the error messages are logged in the /tmp directory.
4	On Windows, run the <code>svinstallcheck</code> utility to verify the successful installation and to make sure that the services are running properly. On UNIX, <code>svinstallcheck</code> is automatically run after the installation. If any errors are found, they are logged in the <code>snapvault</code> log files in the <code>install_dir/etc</code> 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 the 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 Open Systems SnapVault agent is installed.

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

Unattended installation script

Run the `svconfigpackager` utility located in `install_dir/bin` location to create a configuration settings file and an installation script for unattended installations or upgrades. The `svconfigpackager` utility prompts you to answer the following queries:

- ◆ Do you accept the agreement (Y/N)?
- ◆ Do you want to install NetApp Host Agent for managing OSSV (Y/N)?
- ◆ Do you want to upgrade any existing NetApp Host Agent for managing OSSV (Y/N)?

After the installation script is successfully created, the following message is displayed:

```
Operation completed successfully
The following files have been placed in 'D:/Program
Files/netapp/snapvault':
'svconf.in' (Configuration Package)
'unattinstall.bat' (Unattended install batch file)
```

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 49
- ◆ [“Running the Configurator utility”](#) on page 54
- ◆ [“Confirming that services are running”](#) on page 56
- ◆ [“Modifying Open Systems SnapVault parameters”](#) on page 57
- ◆ [“Enabling and disabling debugging”](#) on page 60
- ◆ [“Setting block-level incremental backup options”](#) on page 64
- ◆ [“Configuring backup exclusion lists”](#) on page 66
- ◆ [“Configuring open file backup for Windows”](#) on page 69

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 49
- ◆ “[Understanding the svsetstanza command](#)” on page 51

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 interface 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 Open Systems SnapVault log files, see “[Locating status and problem reports](#)” on page 86.

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.

List of configuration files

The parameters that the different tabs of the *Configurator* utility interface can configure or change are included in the following configuration files. These configuration are located in the *install_dir/snapvault/config* directory.

File name	Description
configure.cfg	The General tab of the Open Systems SnapVault <i>Configurator</i> interface represents the values in this file.
estimator.cfg	The <i>svestimator</i> utility uses the values specified in this file to arrive at a better estimate of space requirements for Open Systems SnapVault installation and data transfers.
programs.cfg	The Trace Level tab of the Open Systems SnapVault <i>Configurator</i> interface represents the values in this file.
snapvault.cfg	The SnapVault tab of the Open Systems SnapVault <i>Configurator</i> interface represents the values in this file.

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. For more information on the configuration files, see [“List of configuration files”](#) on page 50.

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 Qtree SnapMirror Access List specifies “f840, f880”, it can be changed to “f840, f740”.

Notes about variable values: The values that you specify for the variables in the *svsetstanza* command are not case-sensitive. Values that contain spaces must be enclosed in double quotes (“ ”).

Examples

Example1: In the following example, imagine that 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 Qtree SnapMirror Access List, enter the following command:

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

Note


The svsetstanza command does not validate the given values. It just writes the values to the .cfg file.

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 the Windows primary system. ◆ Run the following command on the UNIX primary system: <code>\$INSTALL_DIR/bin/svconfigurator</code> 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.  The screenshot shows the 'Configurator' window with the 'SnapVault' tab selected. The settings are as follows: <table border="1"><thead><tr><th>Parameter</th><th>Value</th></tr></thead><tbody><tr><td>NDMP Account</td><td>root</td></tr><tr><td>NDMP Host Name</td><td></td></tr><tr><td>Check QSM Access List</td><td>☒</td></tr><tr><td>BLI Level</td><td>HIGH</td></tr><tr><td>Run estimator before each backup</td><td>☒</td></tr><tr><td>Enable restart/resync on restore</td><td>☒</td></tr><tr><td>Enable database backup</td><td>BLI</td></tr><tr><td>Name Based Block Level Incrementals</td><td>☒</td></tr><tr><td>List of Drives not to Snapshot (e.g. D,F)</td><td></td></tr><tr><td>OFM Write inactivity period (seconds)</td><td>3</td></tr><tr><td>Maximum time to wait for OFM synchronization (seconds)</td><td>60</td></tr><tr><td>Size of OFM bitmap for synchronized block tracking (MBytes)</td><td>8</td></tr><tr><td>VSS Snapshot creation timeout (secs)</td><td>180</td></tr><tr><td>NDMP Listen Port</td><td>10000</td></tr><tr><td>QSM Access List</td><td>840, f880</td></tr><tr><td>NDMP Host Id</td><td>3E9634F028C527D5E5392E0EDED6676A</td></tr></tbody></table> Buttons: OK, Cancel, Apply	Parameter	Value	NDMP Account	root	NDMP Host Name		Check QSM Access List	☒	BLI Level	HIGH	Run estimator before each backup	☒	Enable restart/resync on restore	☒	Enable database backup	BLI	Name Based Block Level Incrementals	☒	List of Drives not to Snapshot (e.g. D,F)		OFM Write inactivity period (seconds)	3	Maximum time to wait for OFM synchronization (seconds)	60	Size of OFM bitmap for synchronized block tracking (MBytes)	8	VSS Snapshot creation timeout (secs)	180	NDMP Listen Port	10000	QSM Access List	840, f880	NDMP Host Id	3E9634F028C527D5E5392E0EDED6676A
Parameter	Value																																		
NDMP Account	root																																		
NDMP Host Name																																			
Check QSM Access List	☒																																		
BLI Level	HIGH																																		
Run estimator before each backup	☒																																		
Enable restart/resync on restore	☒																																		
Enable database backup	BLI																																		
Name Based Block Level Incrementals	☒																																		
List of Drives not to Snapshot (e.g. D,F)																																			
OFM Write inactivity period (seconds)	3																																		
Maximum time to wait for OFM synchronization (seconds)	60																																		
Size of OFM bitmap for synchronized block tracking (MBytes)	8																																		
VSS Snapshot creation timeout (secs)	180																																		
NDMP Listen Port	10000																																		
QSM Access List	840, f880																																		
NDMP Host Id	3E9634F028C527D5E5392E0EDED6676A																																		

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 displayed.
3	If services are not running, click Start Service.
4	Click OK.

Modifying Open Systems 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 system 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	Select or clear the Check QSM Access List check 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 Qtree SnapMirror access list, complete the following steps.

Step	Action
1	Click the SnapVault tab.
2	Add, 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 NetApp recommends against changing the default values unless technical support asks you:

- ◆ 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 <i>install_dir</i> \bin directory and run the <i>svpassword</i> 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.

Enabling and disabling debugging


About generating debug files

Open Systems SnapVault can generate debug files to help troubleshoot problems. When generating these files, ensure the following:

- ◆ Enable debugging only when advised by technical support, because the debug files grow quickly and affect performance.
- ◆ Disable the generation of these files after you have sent a batch to technical support.
- ◆ Delete the debug files from the system after you have sent them to technical support, to minimize the aforementioned performance impact.

Enabling the generation of debug files

To enable the generation of debug 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 disable the generation of debug files after you are done troubleshooting, clear the “Generate debugging files” check box. _____
3	Click the Trace Level tab.
4	From the drop-down list of the service for which you want to generate debug information, select one of the five trace levels—ALWAYS, NORMAL, VERBOSE, LIBNORMAL, or 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	a. Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services. Wait until the Current State displays as Stopped. b. Click Start Service. Wait until the Current State displays as Running. c. Click OK. Note _____ Instead of performing these steps through the Configurator utility, you can use the <code>svpmgr shutdown</code> and <code>svpmgr startup</code> commands on the command line. _____

Disabling the generation of debug files

To disable the generation of debug files, complete the following steps.

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


Step	Action
2	Clear the “Generate debugging files” check box.
3	Click the Trace Level tab.
4	Select NORMAL from the drop-down list of the service for which you want to stop generating debug 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 backup is set to HIGH. However, you can select the level at which you want to compute BLI backup checksums or you can disable checksum computations entirely.

To select 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 list. The screenshot shows a 'Configurator' window with the 'SnapVault' tab selected. The window contains various configuration fields and checkboxes. The 'BLI Level' dropdown menu is set to 'HIGH'. Other visible options include 'Check QSM Access List', 'Run estimator before each backup', 'Enable restart/resync on restore', 'Enable database backup' (set to 'BLI'), 'Name Based Block Level Incrementals', and several numeric fields for OFM Write inactivity period, Maximum time to wait for OFM synchronization, Size of OFM bitmap, and VSS Snapshot creation timeout. The bottom of the window has 'OK', 'Cancel', and 'Apply' buttons. Machine Service General Trace Level SnapVault NDMP Account: root NDMP Host Name: ☒ Check QSM Access List BLI Level: HIGH ☒ Run estimator before each backup ☒ Enable restart/resync on restore Enable database backup: BLI ☒ Name Based Block Level Incrementals List of Drives not to Snapshot (e.g. D,F): OFM Write inactivity period (seconds): 3 Maximum time to wait for OFM synchronization (seconds): 60 Size of OFM bitmap for synchronized block tracking (MBytes): 8 VSS Snapshot creation timeout (secs): 180 NDMP Listen Port: 10000 QSM Access List: f840, f880 NDMP Host Id: 3E9634F028C527D5E5392E0EDED6676A OK Cancel Apply

Step	Action
2	Select the BLI Level backup option that you want, by using the drop-down list. ◆ 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

Support for name-based block-level incremental backups 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 other 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 lists

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 backs up open files using one of the following options:

- ◆ OFM—Used for Windows 2000 platforms

OFM is automatically installed in the system at the same time the Open Systems SnapVault agent is installed; however, you do need to have a license to use the OFM feature as described in “[License requirements](#)” on page 17.

- ◆ VSS Snapshot copy—Used for Windows 2003 platforms

The VSS Snapshot copy 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 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 period of process inactivity to synchronize and make a Snapshot copy of the target disk drive. Without this quiescent period, OFM will not initiate a Snapshot copy. An OFM failure can result in a blank backup.

Note

OFM supports one active Snapshot copy 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. If multiple volumes (or drives) are backed up simultaneously, backups will proceed, but the drives will be backed up sequentially.

If an active Snapshot copy 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 <i>code</i> 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.  The screenshot shows the 'Configurator' window with the 'SnapVault' tab selected. The settings are as follows: - NDMP Account: root - NDMP Host Name: (empty) ☒ Check QSM Access List - BLI Level: HIGH ☒ Run estimator before each backup ☒ Enable restart/resync on restore - Enable database backup: BLI ☒ Name Based Block Level Incrementals - List of Drives not to Snapshot (e.g. D,F): (empty) - OFM Write inactivity period (seconds): 3 - Maximum time to wait for OFM synchronization (seconds): 60 - Size of OFM bitmap for synchronized block tracking (MBytes): 8 - VSS Snapshot creation timeout (secs): 180 - NDMP Listen Port: 10000 - QSM Access List: f840, f880 - NDMP Host Id: 3E9634F028C527D5E5392E0EDED6676A Buttons at the bottom: OK, Cancel, Apply.

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: <pre>snapvault modify -o back_up_open_files off</pre>

Setting the VSS Snapshot copy timeout parameter on Windows Server 2003

Certain conditions must be met before the Open Systems SnapVault agent can acquire a VSS Snapshot copy. You can set the amount of time (Snapshot timeout) that the agent waits until it retries a VSS Snapshot copy 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.

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 74
- ◆ “[Restoring data](#)” on page 78

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. For more information, see Chapter 2, “[Installing the Open Systems SnapVault Software](#),” on page 13.
2. On the SnapVault secondary storage system, use the storage system console commands to license and enable SnapVault, and specify the Open Systems platforms to back up. See “[Configuring the SnapVault secondary storage system](#)” on page 75.
3. On the SnapVault secondary storage system, start the baseline transfer. See “[Creating an initial baseline copy](#)” on page 76.
4. On the SnapVault secondary storage system, schedule times for drives, directories, or subdirectories to be backed up to the secondary storage. See “[Scheduling SnapVault update backups](#)” on page 77.

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 the Open Systems platform SnapVault backup, complete the following steps.

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

Setting up SnapVault backup on Open Systems platforms

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 <i>still</i> 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: <pre>snapvault start -S prim_host:dirpath sec_host:/vol/sec_vol/sec_tree</pre> Example 1 (Windows): <pre>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</pre> Example 2 (UNIX): <pre>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</pre>

Setting up SnapVault backup on Open Systems platforms

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 <code>snapvault snap sched</code> 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 Snapshot copies to store on the secondary storage system, and the days and hours to execute. Example: <pre>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</pre> Note You must use the <code>-x</code> 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 <code>-x</code> parameter, the default parameter <code>-c</code> is used, which creates Snapshot copies of file systems locally. After all the secondary qtrees on the specified volume have been updated, SnapVault creates a Snapshot copy of this volume for archival.

Restoring data

When to restore data

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

Topics in this section

The following topics are covered in this section:

- ◆ [“Restoring a directory or a file”](#) on page 79
- ◆ [“Restoring an entire primary storage system”](#) on page 82
- ◆ [“Restoring files to a primary storage system from tape”](#) on page 83

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 primary storage system using NFS or CIFS 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 restore operations without requiring a backup administrator’s assistance.

Note

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

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

Step	Action
1	Mount the backed-up qtree from the SnapVault secondary storage system to the primary storage system, using NFS.
2	Use the UNIX <code>cp</code> 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 primary storage system using CIFS, complete the following steps.

Step	Action
1	Create a share from the backed-up qtree on the SnapVault secondary storage system to the primary storage system using CIFS.
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 Snapshot copies.

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

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

Step	Description
2	Enter the <code>snapvault restore</code> 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 copy that you want to restore from (for example, <code>sv_weekly.0</code>, <code>sv_weekly.1</code>, or <code>sv_weekly.2</code>) ◆ The path on the primary storage system to the directory or file that you want to restore to Example 1 - Single file restore (Windows): <code>snapvault restore -s sv_daily.0 -S myvault:/vol/sv_vol/melzdir/evidence.doc a:\melzdir\evidence_restore.doc</code> Example 2 - Single file restore (UNIX system): <code>snapvault restore -s sv_daily.0 -S myvault:/vol/sv_vol/melzdir/evidence.doc /usr/melzdir/evidence_restore.doc</code> Note Enter the entire command as a single line. Ensure that you do not specify a slash (\ or /) character at the end of the path name in the above command; otherwise, the <code>snapvault restore</code> command will fail. Example 1 - Single directory restore (Windows system): <code>snapvault restore -s sv_daily.0 -S myvault:/vol/sv_vol/melzdir/dir1 a:\melzdir\dir1_restore</code> Example 2 - Single directory restore (UNIX system): <code>snapvault restore -s sv_daily.0 -S myvault:/vol/sv_vol/melzdir/dir1 /usr/melzdir/dir1_restore</code>

Using the DataFabric Manager restore wizard

The DataFabric Manager restore wizard leads you through the entire restore process. For details, see the *DataFabric Manager Administration Guide*.

Restoring an entire primary storage system

Restoring a primary storage system

You can restore an entire primary storage system from a SnapVault secondary storage system using NFS or CIFS, 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 Open Systems SnapVault agent. See Chapter 2, “ Installing the Open Systems SnapVault Software ,” on page 13.
4	(Optional) If you backed up the Windows System State data of the primary storage system, restore its Windows System State data if you are restoring. For more information, see “ Restoring System State data ” on page 95.
5	Restore the backed-up directories using the <code>snapvault restore</code> command. For details, see “ Using the snapvault restore command ” on page 80.

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 a primary storage system from tape, using NFS or CIFS, 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 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 <code>restore</code> command to restore from the tape to the SnapVault secondary storage system. For details, see the <i>Tape Backup and Recovery Guide</i> .
3	Copy the files from the SnapVault secondary storage system to the primary storage system using NFS or CIFS. For details, see “Restoring a directory or a file” on page 79.

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 86
- ◆ [“Backing up and restoring the Open Systems SnapVault database”](#) on page 87
- ◆ [“Backing up and restoring Windows System State data”](#) on page 91
- ◆ [“Deleting and re-creating Open Systems SnapVault relationships”](#) on page 100
- ◆ [“Migrating a relationship between two secondary systems”](#) on page 101
- ◆ [“Migrating between two volumes on one secondary system”](#) on page 104
- ◆ [“Setting up a tertiary system for a relationship”](#) on page 107
- ◆ [“Reusing a deleted or renamed directory name”](#) on page 109
- ◆ [“Resynchronizing restored or broken relationships”](#) on page 111
- ◆ [“Retrying failed transfers”](#) on page 113
- ◆ [“Encrypted File System \(EFS\) file backup and restore”](#) on page 114

Locating status and problem reports

Where to find status and problem reports

All log files are located in the following directory:

install_dir/snapvault/etc

install_dir is the directory on the primary storage system on 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 path name, 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.

You can find the operational status and problem reports of the primary storage system in the log files called "snapvault". A new file is created daily at midnight or as soon after midnight as the first subsequent activity on the system takes place; the existing file is not archived until a new one is created. The following message is logged:

"Previous snapvault log file is archived to: *install_dir*/etc/snapvault.yyyymmdd"

Note

The current file is called "snapvault" and it has no extension. However, the archived files have the .yyyymmdd extension, where yyyy is the year, mm is the month, and dd is the date when the file was created.

Backing up 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)
- ◆ Checkpoint file (if a backup transfer had failed with a checkpoint)

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	<i>xx</i> <i>xx</i> is a unique integer
BLI checksums	<i>xx</i> -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   48      01-checksums     BLI checksums file
```

Need to back up 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 must 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.

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.

Note

Any checkpoint files, are not included in the backup. Also any softlock state in the database are not backed up.

Characteristics of the database backup file

The database files for an Open Systems SnapVault relationship on the primary system are compressed and backed up as a file named .OSSV_DATABASE_BACKUP.

If a file named .OSSV_DATABASE_BACKUP already exists in the directory being backed up, 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.

Note

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 90 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 111.

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.	
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\netapp\snapvault\bin. ◆ On UNIX systems, the default path is /usr/snapvault/bin.
2	Enter the following command: <pre>snapvault restore -S <i>secondary_system:pathname</i> /<i>.OSSV_DATABASE_BACKUP pri_pathname</i></pre> <i>secondary_system</i> is the secondary storage system. <i>pathname</i> is the path where the .OSSV_DATABASE_BACKUP file is located on the secondary system. <i>pri_pathname</i> 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: <pre>D:\Program Files\netapp\snapvault\bin>snapvault restore -S f840:/vol/vol0/re15/.OSSV_DATABASE_BACKUP D:\temp\database</pre>

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 re-created.
- ◆ 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
Re-creating 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 Flexible Single Master Operation (FSMO). 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 94
- ◆ [“Restoring System State data”](#) on page 95
- ◆ [“Using System State backup to rebuild a primary system”](#) on page 97

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 storage system console of the SnapVault secondary system, enter the following command: <pre>snapvault start -S prim_host:SystemState sec_host:/vol/sec_vol/sec_tree</pre> Note— The keyword <code>SystemState</code> is case independent. Example: <pre>snapvault start -S melzhost:SystemState sv_secondary:/vol/sv_vol/tree_melz</pre>

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 re-created 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 96.

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: <pre>snapvault restore -S sec_host:/vol/sec_vol/sec_tree SystemState</pre>Note The keyword <code>SystemState</code> 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 <code>snapvault restore</code> command will fail. Example: <pre>snapvault restore -S sv_secondary:/vol/sv_vol/tree_melz SystemState</pre>
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 <code>ntdsutil</code> 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: <pre>snapvault restore -S sec_host:/vol/sec_vol/sec_tree SystemStatePrimary</pre> 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: <pre>snapvault restore -S sv_secondary:/vol/sv_vol/tree_melz SystemStatePrimary</pre>
3	Reboot.

Backing up and restoring Windows System State data

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 storage system console of the SnapVault secondary system, enter the following command: <pre>snapvault start -S prim_host:SystemState sec_host:/vol/sec_vol/sec_tree</pre> Note_____ The keyword <code>SystemState</code> is case independent. _____ Example: <pre>snapvault start -S melzhost:SystemState sv_secondary:/vol/sv_vol/tree_melz</pre>

Rebuilding a primary system from the System State backup

Keep the following guidelines in mind before proceeding:

- ◆ Using a System State data backup on a system that is not effectively a duplicate of the original system disables the new system.
- ◆ When you rebuild a new primary system, many of its characteristics must be the same as the original primary system from which you perform the system state backup. In the case of disk drives, the new primary system should be either the same size or larger than the original. Those characteristics that might be different on the new primary might cause problems, depending on the specifics of the system. One potential difference could be the type of video bus—for example, an AGP bus versus a PCI video bus.

To rebuild a primary system using the backup you created in [“Creating a backup to rebuild a primary system”](#) on page 97, 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 that 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 that 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 95 and “Restoring System State data from a backup and marking it primary” on page 96 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 re-creating Open Systems SnapVault relationships

About deleting and re-creating a relationship

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

If you use the `snapvault stop` command on the secondary system to delete an Open Systems SnapVault relationship and try to re-create the relationship without releasing the relationship on the primary system, the attempt 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 re-creating a relationship

Follow these general steps to delete and re-create Open Systems SnapVault relationships.

Step	Action
1	On the secondary system console, enter this command: snapvault stop <i>secondary_path</i>
2	On the primary system console, enter this command: snapvault release <i>primary_path</i> [<i>secondary:</i>] <i>secondary_path</i>
3	On the secondary system console, enter this command: snapvault start -S <i>primary_path</i> [<i>secondary:</i>] <i>secondary_path</i>

Migrating a relationship between two secondary systems

Migrating data from one secondary system to another

Before migrating data from one secondary system to another, review the basics of setting up SnapVault transfers (for example, access permission, licensing, and correct volume language) 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 <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i> 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): a. Create the new_vol volume. <code>r200-new> vol create new_vol 3</code> b. Mark the new_vol volume restricted. <code>r200-new> vol restrict new_vol</code> c. Transfer the old_vol volume to the new_vol volume. <code>r200-new> snapmirror initialize -S r200-old:old_vol new_vol</code>

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: a. <code>snapmirror quiesce new_vol</code> b. <code>snapmirror break new_vol</code>												
4	Check the SnapMirror status and SnapVault status on the new secondary system. SnapMirror status should be <code>Broken-off</code>. SnapVault status should be <code>Snapvaulted</code> on the new volume on the new secondary system. Example: Perform the following steps from r200-new: a. <code>snapmirror status</code> <table><tr><td>Source</td><td>Destination</td><td>State</td></tr><tr><td>r200-old:old_vol</td><td>r200-new:new_vol</td><td>Broken-off</td></tr></table> b. <code>snapvault status</code> <table><tr><td>Source</td><td>Destination</td><td>State</td></tr><tr><td>bno:C:\500MB</td><td>r200-new:/vol/new_vol/bno_C_500MB</td><td>Snapvaulted</td></tr></table>	Source	Destination	State	r200-old:old_vol	r200-new:new_vol	Broken-off	Source	Destination	State	bno:C:\500MB	r200-new:/vol/new_vol/bno_C_500MB	Snapvaulted
Source	Destination	State											
r200-old:old_vol	r200-new:new_vol	Broken-off											
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 <code>snapvault status -c</code> command. Example: Perform the following step from r200-new: <code>snapvault status -c</code> Snapvault secondary is ON.												
6	Add SnapVault configuration information to the registry on the new secondary system using the <code>snapvault start</code> command. Note_____ This does not start a new baseline, it updates the registry. _____ Example: Perform the following step from r200-new: <code>snapvault start -S bno:C:\500MB r200-new:/vol/new_vol/bno_C_500MB</code> 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 <code>snapvault status -c</code> command. Example: Perform the following step from <code>r200-new</code>: <pre> snapvault status -c Snapvault secondary is ON. /vol/new_vol/bno_C_500MB source=bno:C:\500MB </pre>
8	Test the new SnapVault relationship by manually updating <code>r200-new</code>. 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 <code>r200-new</code>: <pre> snapvault update r200-new:/vol/new_vol/bno_C_500MB Transfer started. Monitor progress with 'snapvault status' or the snapmirror log. </pre>
9	Re-create any schedules used on the old secondary system to the new secondary system and ensure that 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 <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i> 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): a. Create the new_vol volume. <code>vol create new_vol 3</code> b. Mark the new_vol volume restricted. <code>vol restrict new_vol</code> c. Transfer the old_vol volume to the new_vol volume. <code>snapmirror initialize -S r200:old_vol new_vol</code>


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: a. snapmirror quiesce new_vol b. 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. a. snapmirror status <table><tr><td>Source</td><td>Destination</td><td>State</td></tr><tr><td>r200:old_vol</td><td>r200:new_vol</td><td>Broken-off</td></tr></table> b. snapvault status <table><tr><td>Source</td><td>Destination</td><td>State</td></tr><tr><td>bno:C:\500MB</td><td>r200:/vol/new_vol/bno_C_500MB</td><td>Snapvaulted</td></tr></table>	Source	Destination	State	r200:old_vol	r200:new_vol	Broken-off	Source	Destination	State	bno:C:\500MB	r200:/vol/new_vol/bno_C_500MB	Snapvaulted
Source	Destination	State											
r200:old_vol	r200:new_vol	Broken-off											
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 <code>snapvault status -c</code> command. Example: Perform the following step from r200. <pre> snapvault status -c Snapvault secondary is ON. /vol/new_vol/bno_C_500MB source=bno:C:\5000MB </pre>
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. <pre> snapvault update r200:/vol/new_vol/bno_C_500MB Transfer started. Monitor progress with 'snapvault status' or the snapmirror log. </pre>
9	Re-create any schedules used on the old volume to the new volume and ensure that access permissions are in place.

Setting up a tertiary system for a relationship

Need for a tertiary system

You can protect the SnapVault secondary system from disasters 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.



Also, the 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. A softlock is a request to the primary system to retain the context for re-running a transfer. In Network Appliance terminology, a softlock is a reference to a Snapshot™ copy. It is a destination system's request to the source system to not delete a particular Snapshot copy.

In an open system's context, the destination system is requesting that the source system keep enough storage space free to be able to run a "SnapVault update" from that particular point in time. For instance, if your SnapVault secondary system becomes unusable because of a disaster, you can manually redirect the next SnapVault transfers to the tertiary system instead of the old SnapVault secondary system. Effectively, the tertiary system becomes the new SnapVault secondary system, and the SnapVault transfers continue using the most recent Snapshot copy common to both the primary and tertiary systems.

After the secondary system comes up, resynchronize it using the tertiary system. Then, you can configure the secondary system in any one of the following ways:

- ◆ primary -> tertiary -> secondary

In this scenario, you manually release the SnapVault relationship between the primary and secondary systems from the secondary system. The original tertiary system then takes the place of the secondary system, and vice versa.

- ◆ primary -> secondary -> tertiary

To revert to the previous relationship, in which the original secondary system is retained, manually release the previous SnapMirror relationship (that is, between tertiary and secondary systems) and create a new relationship between the original secondary and tertiary systems.

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 re-create 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 <code>snapvault release</code> 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 copy exists between the primary and secondary systems. A common Snapshot copy is necessary for incremental backups to continue successfully between the primary and secondary systems. If the common Snapshot copy 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:

- ◆ Data is restored using the `snapvault restore` command. The `snapvault restore` command is used to restore primary data to its state at the time of one of its SnapVault Snapshot copies.
In this condition, if you restore the data from the qtree to another location on the primary system, and then perform subsequent incremental backups from the restored location to the same qtree on the secondary system, you must resynchronize the relationship.
- ◆ An older version of the Open Systems SnapVault database is restored on the primary system after the primary database is corrupted.
- ◆ The state of a destination qtree in a SnapVault relationship is changed to read-write.

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 are modified before resynchronization, all data written to this qtree is lost upon resynchronization.

Need for resynchronization

Prior to Open Systems SnapVault 2.2, the only way to resynchronize a SnapVault relationship between primary and secondary systems 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 SnapVault relationship, perform the following steps.

Step	Action
1	Log in to the secondary storage system.
2	On the console, enter the following command: <pre>snapvault start -r -S <i>prim_host:dirpath</i> <i>/vol/sec_vol/sec_tree</i></pre> Example 1 (Windows): snapvault start -r -S melzhost:C:\melzdir /vol/sv_vol/tree_melz Example 2 (UNIX): snapvault start -r -S melzhost:/usr/melzdir /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: <pre>snapvault modify -t <i>n</i> <i>sec_qtree</i></pre> <i>n</i> is the number that specifies the number of retries. Note The failed attempt counts as the first try. Therefore, if you set <i>n</i> to 5, four attempts are made after the failed try. <i>sec_qtree</i> is the qtree on the secondary system where data is being backed up.

Encrypted File System (EFS) file backup and restore

About EFS file backup and restore

Open Systems SnapVault is capable of backing up and restoring EFS 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 files. Any time an EFS file is modified, Open Systems SnapVault backs up the entire EFS file.

Requirements for EFS file backup and restore

The following are the requirements for backing up and restoring EFS 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 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 files requires the Open Systems SnapVault agent to create a temporary file that is equivalent to the size of the EFS file that is being replaced. Open Systems SnapVault is capable of restoring up to five files concurrently. This means that a restore of EFS files might 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 files in the target volume.

What Open Systems SnapVault space estimator is

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

How space estimates are made

The space estimator bases its calculation on many factors, such as the number of files, size of files, length of directory names, exclusion lists, access control lists (ACLs), 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 124.

- ◆ Stand-alone

In this case, the space estimator is installed as a stand-alone application on a system that might not have an existing Open Systems SnapVault installation. For more details, see “[Configuration files required for space estimator](#)”

[operation](#)” on page 117. Run the svestimator executable available on the Open Systems SnapVault 2.3 CD-ROM.

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

The following is a sample command and subsequent output of the estimator report in which Open Systems SnapVault is installed on a system:

```
C:\>svestimator.exe -o C:\ E:\
```

```
Scanning system volumes...
Volume 'C:\' type Normal NTFS Free Space 52%
Volume 'D:\' type CDROM Free Space 0%
Volume 'E:\' type Normal NTFS Free Space 47%
```

```
Examining 'C:\'...
```

```
Estimated space requirements so far:
Database: 57.00 MB
Temp: 110.00 MB
```

```
Examining 'E:\'...
```

```
Estimated space requirements so far:
Database: 207.00 MB
Temp: 298.00 MB
```

```
Analyzing space requirements...
Estimator has found sufficient space for backup
```

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

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

```
C:\>svestimator.exe -i C:\ E:\
```

```
Scanning system volumes...
Volume 'C:\' type Normal NTFS Free Space 52%
Volume 'D:\' type CDROM Free Space 0%
Volume 'E:\' type Normal NTFS Free Space 47%
```

```

Examining 'C:\'...

Estimated space requirements so far:
Installation: 12.00 MB
Database: 57.00 MB
Temp: 110.00 MB

Examining 'E:\'...

Estimated space requirements so far:
Installation: 12.00 MB
Database: 207.00 MB
Temp: 298.00 MB

Analyzing space requirements...
'C:\' is suitable for 'Installation requirements'
'C:\' is suitable for 'Database requirements'
'C:\' 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 the 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. For more information, see [“Contents of the estimator.cfg file”](#) on page 118.
- ◆ The path and file exclusion list files
For more information, see [“Configuring backup exclusion lists”](#) on page 66.

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 stand-alone 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 stand-alone space estimator is run.

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

```
# Sample Estimator configuration file.

# In stand-alone 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 on which the space estimator is running. ◆ During the stand-alone operation, this field enables the space estimator utility 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 utility is running. ◆ During the stand-alone 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 <i>not</i> take into consideration when determining the amount of free disk space for a backup operation. This space is assumed to be reserved for making Snapshot copies on the detected drives and therefore not available for backup operation.
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.

Field	Default Value	Description
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 Open Systems SnapVault 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 <code>-i</code> option is specified in stand-alone mode.

Logs to which the space estimator information is written

When the space estimator is run in stand-alone 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 stand-alone 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 stand-alone 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 stand-alone mode except the drive scan results. For more information about the snapvault log files, see “[Locating status and problem reports](#)” on page 86).

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 stand-alone space estimator, complete the following step.

Step	Action	
1	If...	Then...
	The Open Systems SnapVault software is installed on your primary storage system	a. Go to the <i>install_dir/bin</i> directory and locate the file called svestimator. You can run the stand-alone 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	a. 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. b. Uncompress or unzip the installation package, if needed. c. 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 stand-alone mode: To run the space estimator in stand-alone 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: <pre>svestimator [-o -i -d] root_backup_path</pre> -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 the debug trace information to be written to a log. If Open Systems SnapVault is installed on the system, the debug trace called <code>svestimator.txt</code> is written to the <code>install_dir/trace</code> directory. If Open Systems SnapVault is not installed, a directory called <code>trace</code> 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 <code>Configurator</code> utility before using the -d option of the <code>svestimator</code> command. If the “Generate debugging files” option is not enabled, the log file will not be written to the trace directory. <i>root_backup_path</i> specifies the directory you want to back up—for example, <code>C:\MyData\MyDocs</code>. 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 space estimator

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

Step	Action
1	In the SnapVault tab of the Configurator utility, clear 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 <i>install_dir</i>/config directory, set the value of the following entry in the snapvault.cfg file to True: <pre>[Configuration:Estimator can fail backup] Value=True</pre>


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 same location as the secondary system (data center), transfer the data from the portable device to the secondary storage 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 systems is established for incremental transfers to occur directly.



LREP supported platforms

The LREP tool can be run on all supported operating systems. 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 2000 system.

Versioning

LREP 1.1 is packaged with Open Systems SnapVault 2.3. To check the version number, run the following command from <INSTALL_DIR>\bin\:

```
C:\>lrep_reader -v
```

Downloading the LREP tool

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

Also, 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 binary file of the LREP tool is located on that Download page.
- ◆ Go to <http://now.netapp.com/eservice/toolchest> page and locate the LREP tool.

Note

Before you download LREP 1.1, delete any existing earlier versions.

Installing the LREP tool

Because the LREP tool is a stand-alone binary file, it does not need to be installed and can be run from the system on which the binary file is located. If Open Systems SnapVault is installed, the following files are available:

- ◆ For Windows:
 <INSTALL_DIR>\bin\lrep_reader.exe
 <INSTALL_DIR>\bin\lrep_writer.exe
- ◆ For UNIX:
 \$INSTALL_DIR/bin/lrep_reader
 \$INSTALL_DIR/bin/lrep_writer

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.

Note

If the `lrep_writer` utility needs to be run on a machine on which the Open Systems SnapVault software is installed, 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 client, XP1. The Zip drive is shipped to the data center where a second Windows 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.exe -f SEC1 -q /vol/vol1/backup -o E:\test@0  
WIN1:D:\dataset
```

`-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 client, XP2, and enter the following command:

```
C:\>lrep_writer.exe 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 command:

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

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

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

```
SEC1> snapvault status
```

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 these errors.

Note

The suggested actions in the following table assume that you are running Open Systems SnapVault 2.3 on the primary storage system.

Error number	Error message	Cause	Action to take
	Cannot connect to the 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 that Open Systems SnapVault uses for NDMP connections by changing the NDMP port value in the Open Systems SnapVault Configurator tool (svconfigurator). For information on how to change the port, see “Modifying NDMP Listen Port setting” on page 59.
	Open Systems SnapVault encountered a network error while reading (writing) data.	Can be caused by network problems or the secondary storage system aborting the transfer.	For more information, check the log files on the secondary storage system.

Error number	Error message	Cause	Action to take
	OFM volume cannot be synchronized - volume error. (1)	This occurs when OFM is not able to create a Snapshot copy of the drive that holds the data for the relationship. (OFM requires a period of file system inactivity to create a Snapshot copy.)	Edit the following settings in the Open Systems SnapVault Configurator (svconfigurator) and retry the transfer: ◆ Set the “Write inactivity period (seconds)” from the default of 5 to 2. ◆ Set “Synchronization timeout (seconds)” from the default of 60 to 120.
	Cannot back up 'c:\' Failed to create a volume snapshot	This occurs when OFM is not able to create a Snapshot copy of the drive that holds the data for the relationship. (OFM requires a period of file system inactivity to create a Snapshot copy.)	Edit the following settings in the Open Systems SnapVault Configurator (svconfigurator) and retry the transfer: ◆ Set the “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 number	Error message	Cause	Action to take
	Failed to write a hybrid history record	Open Systems SnapVault has failed to write to its database. This error can occur either due to 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, the 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.
	Unexpected read select while no data pending	This error can occur because of one of the following problems: ◆ The <code>snapvault abort</code> command was issued using the command line. ◆ The Open Systems SnapVault machine failed/rebooted. ◆ There was a network error.	Check the SnapMirror log on the secondary system and the SnapVault log on the primary system for errors, and then retry the transfer.
	Transfer from <SRC> to <DEST>: request denied by the source filer; check access permission on the source filer.	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 Configurator utility contains the IP address or network name of the secondary system. For more information, see “Modifying the QSM access list” on page 57.

Error number	Error message	Cause	Action to take
	Current Transfer Error: unable to translate Unicode path name, please check volume language configuration.	A Unicode path name that cannot be translated aborts the 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. For more information, see bug 133965 at http://now.netapp.com .
2007	Unable to process the softlock data.	Open Systems SnapVault could not process the softlock due to an internal problem.	Retry the transfer. If you see the problem again, contact technical support.
2008	Unable to restore the file.	There might be no space left on the disk or you might not have adequate permission.	Check the availability of disk space and ensure that you have the adequate permission to write to the file.
2009	Unable to set an attribute on the file.	During a restore operation, there might be no space left on the disk or you might not have adequate permission.	Check the availability of disk space and ensure that you have the adequate permission to write to the file.
3001	Failed to open the file.	During backup or restore operations, the file is in use by some other process or you might not have adequate permission.	Retry after closing other processes and ensure that you have the adequate permission.

Error number	Error message	Cause	Action to take
3002	Failed to seek in file.	During backup or restore operations, the file is in use by some other process or you might not have adequate permission.	Retry after closing other processes and ensure that you have the adequate permission.
3003	Failed to read the file.	During backup or restore operations, the file is in use by some other process or you might not have adequate permission.	Retry after closing other processes and ensure that you have the adequate permission to read from the file.
3004	Failed to write to the file.	There might be no space left on the disk, or you might not have adequate permission.	Check the availability of disk space and ensure that you have the adequate permission to write to the file.
3005	Failed to get or set the file information	During backup or restore operations, the file is in use by another process or you might not have adequate permission.	Ensure that the file is not being used by any other process and retry the operation later.
3006	Failed to read from the Open Systems SnapVault database.	This message can occur because of one of the following conditions: ◆ The database is in use by some other process. ◆ You might not have adequate permission. ◆ The database is corrupted.	Retry after closing other processes and ensure that you have the adequate permission. You might need to reset the relationship to its baseline if the database is corrupted.

Error number	Error message	Cause	Action to take
3007	Failed to write to the Open Systems SnapVault database.	There might be no space left on the disk or you might not have adequate permission.	Check the availability of disk space and ensure that you have the adequate permission to write to the database.
3008	Failed to open the Open Systems SnapVault database.	The database is in use by some other process or you might not have adequate permission.	Retry after closing other processes and ensure that you have the adequate permission to open the database file on the primary system.
3009	An internal processing error has occurred.	An unexpected situation was encountered.	Ensure that you are running the most recent version of the storage system and Open Systems SnapVault. If you see the problem again, contact technical support.
3010	An invalid path was specified. A possible attempt to update an empty directory.	This occurs when an empty directory is backed up.	Ensure that the directory is not empty.
3012	Snapshot copy failure	The VSS or OFM shadow copy fails. The Snapshot copy has encountered a failure in the primary system.	In the case of Windows 2000, check for OFM failure or refer to the troubleshooting section of the Microsoft Volume Shadow Copy Service on Windows 2003.

Error number	Error message	Cause	Action to take
3013	Insufficient disk space to perform the operation.	During a backup operation or, more likely, during a restore operation, there might not be sufficient space left on the disk.	Check the availability of disk space on the primary system before performing a backup or restore operation.
3014	An error is encountered while processing checkpoint information.	This occurs during checkpoint read or write.	Ensure that you are running the supported versions of the secondary storage system and Open Systems SnapVault. If that does not solve the problem, contact technical support.
3016	A network error has occurred.	The network socket was closed unexpectedly or the transfer was aborted by the user.	Verify network connectivity between the Open Systems SnapVault primary system and the secondary storage system.
	<SRC> <DEST> Abort (could not read from socket)	This message can occur because of one of the following problems: ◆ The <code>snapvault abort</code> 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 that you are running the most recent recommended version of Data ONTAP on the secondary system.

Error number	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 that you are running the recommended versions of Open Systems SnapVault on the primary system and Data ONTAP on the 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 transferred using incorrect procedure, it is possible that an old database gets transferred, thus causing the relationship to be out-of-sync.)	Ensure that you follow the correct procedure to transfer the Open Systems SnapVault database.
	Transfer aborted: destination qtree is not coalesced	The secondary qtree is still in a transition state.	Ensure that the destination qtree is in an IDLE state before attempting a transfer.

Error number	Error message	Cause	Action to take
	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.	Ensure that you are running the latest recommended versions of Open Systems SnapVault on the primary system and Data ONTAP on the 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 might 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.	For more information, see bugs 137685 and 140930 at http://now.netapp.com .

Error number	Error message	Cause	Action to take
	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 that you are running the most recent recommended versions of Open Systems SnapVault on the primary system and Data ONTAP on the secondary system.
	Replication destination failed to store an entry in the inode map	Usually caused by a Data ONTAP or Open Systems SnapVault bug.	Make sure that you are running the most recent recommended versions of Open Systems SnapVault on the primary system and Data ONTAP on the secondary system.
	SnapVault: destination transfer from <SRC> to <DEST>: incompatible SnapMirror versions on systems	Due to an ordering error at the destination, it is possible for SnapVault and qtree SnapMirror transfers to fail and display an error message that does not correspond to the error reported on the source side.	For more information, see bug 147982 at http://now.netapp.com .

Index

Numerics

32-bit systems 16

64-bit systems 18

A

aborted transfer, retrying 113

access list (QSM) 57

ACL file 87

Active Directory 11, 46, 91, 92

Allowed Secondary Names window 22

Always setting (for generating debugging information) 61

applications for managing Open Systems SnapVault 6

B

backing up

- disable database 90

- EFS files 12, 114

- open files 9

- using Open Systems SnapVault for 83

- Windows System State 11

backup exclusion lists 10, 66

backup file, database

(OSSV_DATABASE_BACKUP) 88

backups

- EFS files 12

- scheduling 77

- System State data 94

baseline copy, creating 76

batch installation 45

BLI backups

- checksums file 87

- described 10

- for EFS files 10

- for name-based applications 10

- level, setting 64

- name-based 65

broken relationship

- restoring 111

- resuming 12

built-in space estimator

- description 115

- disabling 124

- enabling 122

- installing 122

C

calculation of free space 115

checkpoint

- file 87

- restart 12

checksum computation, BLI checksum 64

command-line (CLI) utility for administration 51

command-line utility (svsetstanza) 5, 51

commands

- lrep_reader 126

- lrep_writer 126, 127

- restore 83

- snapvault modify 113

- snapvault release 100, 109

- snapvault restore 6, 79, 88, 90

- snapvault snap sched 77

- snapvault start 5, 76, 94, 100, 112

- snapvault stop 100

- svconfigpackager 41

- svestimator 123

- svinstallcheck 34

- svpmgr shutdown 35

- svpmgr startup 36

- svsetstanza 5, 51

configuration files

- configure.cfg 51

- estimator.cfg 117

Configurator utility

- description 49

- General tab, purpose of 49

- Machine tab, purpose of 49

- running 54

- Service tab, purpose of 49
- SnapVault parameter, modifying 57
- SnapVault tab, purpose of 50
- Trace tab, purpose of 50
- configure.cfg file 51
- copying files to restore 79

D

- Data ONTAP version, supported 17
- database files
 - ACL 87
 - backup file (OSSV_DATABASE_BACKUP)
 - 88
 - BLI checksums 87
 - checkpoint 87
 - described 87
 - history 87
- databases
 - backup 89
 - disabling backup 90
 - disk space requirements for 16
 - Open Systems SnapVault 87
- DataFabric Manager, using to manage Open Systems SnapVault 6
- debug files
 - generation of 60
 - stop generating 62
- determining free space 115

E

- EFS files
 - about 12
 - BLI and 10
- error messages 129
- estimating free space 12
- estimator.cfg file
 - contents of 118
 - described 117
- exclusion lists 10

F

- failed transfers, retrying 113

- file exclusion lists
 - described 10
 - for space estimator 117
 - location of 66
- free space estimator. *See* space estimator

G

- General tab (Configurator) 49
- Generate debugging files option 60
- GUI interface for administration 5, 49

H

- history file 87

I

- initial backup (baseline) 76
- installation
 - Allowed Secondary Names window 22
 - batch 45
 - built-in space estimator 122
 - limitations 18
 - on HP-UX, AIX, or Linux 29
 - on Solaris 26
 - on UNIX and Linux 26
 - on Windows 19
 - requirements, primary systems 14
 - script for unattended install 41
 - Select Installation Folder window 23
 - Setup wizard 20
 - stand-alone space estimator 122
 - unattended, about 41
 - upgrading. *See* upgrade
- interfaces
 - Setup wizard 20

L

- Libnormal setting 61
- Libverbose setting 61
- licenses 17
- limitations of Open Systems SnapVault 18
- listen port, NDMP 21
- lists, exclusion 10

- log files
 - SnapVault 86
 - space estimator 121
 - trace 121
- log files, SnapVault 50
- Logical Replication (LREP) tool 125
- LREP tool
 - description 125
 - downloading 126
 - installation 126
 - supported platforms 125
- lrep_reader utility 126
- lrep_writer utility 126, 127

M

- Machine tab (Configurator utility) 49
- management application, Open Systems SnapVault 6
- messages, error 129
- migrating a volume 101

N

- name-based
 - applications, about 10
 - block-level incremental backup, setting 65
- NDMP Account field 58
- NDMP Host ID field 58
- NDMP Host Name field 58
- NDMP Listen Port field 21, 58
- NDMP-based management applications 6
- NetApp Host Agent
 - installing 24
 - supportability matrix 17
 - uninstalling 38, 40
- Normal setting 61

O

- OFM
 - about 9, 69
 - configuring 69
 - disabling 71
 - parameters, setting 70
 - St. Bernard's module 19

- OFM licenses 69
- Open File Manager. *See* OFM
- open files, backing up 9
- Open Systems SnapVault 2
- OSSV. *See* Open Systems SnapVault
- OSSV_DATABASE_BACKUP file 88

P

- parameters modifying, SnapVault 57
- path exclusion lists
 - about 10
 - described 66
 - for space estimator 117
 - location of 66
- port
 - listen, NDMP 21
- ports
 - 10000 16, 21, 34, 129
 - 10566 16
 - 4092 16
 - 4093 16
- preinstallation requirements 14
- prerequisites
 - for installation 14
- primary system
 - about 3
 - administration 5
 - controlling secondary system access 57
 - installation requirements 14
 - restoring 82
 - restoring data 80
 - svrestore command 80

Q

- Qtree SnapMirror (QSM) access list 57

R

- relationships
 - broken 111
 - deleting 100
 - migration 101
 - resynchronizing 111

- remote installation
 - and upgrade 41
 - batch 45
- reports, problem 86
- requirements (installation), for primary systems 14
- restart backup from a checkpoint 12
- restore command 83
- restoring
 - a complete system 82
 - broken relationship 111
 - by copying files 79
 - EFS files 114
 - from tape 83
 - methods for 79
 - System State data 95
 - Windows Systems State 11
- resuming broken relationships 12
- resynchronizing a relationship 12, 111
- retrying a failed transfer 113

S

- scheduling backups 77
- script for unattended install 41
- secondary system
 - about 3
 - access control 57
 - administration 5
 - configuring for use with open systems 75
 - enabling SnapVault 75
 - installation requirements 17
 - migrating relationship between 100, 101
- security, enabling or disabling 57
- Select Installation Folder window 23
- Service tab (Configurator utility) 49
- Setup Wizard 20
- SnapMirror, using for protecting 107
- SnapVault
 - configuring the secondary storage system 75
 - log files 50, 86
 - primary system parameters, modifying 57
- snapvault modify command 113
- snapvault release command 100, 109
- snapvault restore command 6, 79, 88, 90
- snapvault snap sched command 77

- snapvault start command 5, 76, 94, 100, 112
- snapvault stop command 100
- SnapVault tab (Configurator utility) 50
- snapvault.cfg file (for space estimator), modifying 124
- softlock support 107
- space estimator
 - about 12
 - built-in
 - description 115
 - disabling 124
 - installing 122
 - calculates, how it 115
 - console output 121
 - described 115
 - enabling 122
 - estimator.cfg file 117
 - log files 121
 - path and file exclusion lists 117
 - stand-alone 115, 122
 - utility 16
- St. Bernard's OFM module 19
- stand-alone space estimator
 - described 115
 - installation 122
- status logs 86
- sv_linux_pri license 17
- sv_ontap_sec license 17
- sv_unix_pri license 17
- sv_windows_pri license 17
- svconfigpackager command (utility) 41
- svestimator command 123
- svinstallcheck command 34
- svpmgr shutdown command 35
- svpmgr startup command 36
- svsetstanza command 5, 51
- system restore, primary 82
- System State data
 - backup 94
 - restore 95

T

- tape, restoring from 83
- tertiary system, data protection with 107

- tombstone lifetime setting 94
- trace file 121
- Trace Level tab (Configurator utility) 50
- troubleshooting 129
- troubleshooting, generating debug files for 60

U

- unattended installation 41
- uninstall, how to 38
- unsupported configurations 18
- upgrade
 - from versions 1.x
 - unattended 41

V

- Verbose setting 61
- volume migration 101
- VSS Snapshot copy
 - about 9, 69
 - setting timeout for 71

W

- Web interface for administration 5, 49
- Windows System State 11
- wizard, setup 20

