

# **Open Systems SnapVault® 2.6 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](mailto:doccomments@netapp.com)  
Information Web: <http://www.netapp.com>

Part number 215-03715\_A0  
March 2008

# Copyright and trademark information

---

## Copyright information

Copyright © 1994–2008 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.A. 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, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, 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, FlexShare, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot,

SpinStor, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

IBM, the IBM logo, AIX, and System Storage are trademarks and/or registered trademarks of International Business Machines Corporation.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. 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

---

	<b>Preface</b> . . . . .	ix
<b>Chapter 1</b>	<b>Introduction to Open Systems SnapVault</b> . . . . .	1
	About Open Systems SnapVault . . . . .	2
	Open Systems SnapVault features . . . . .	9
<b>Chapter 2</b>	<b>Installing the Open Systems SnapVault Software.</b> . . . . .	15
	Prerequisites . . . . .	16
	Installing Open Systems SnapVault on Windows platforms . . . . .	22
	Installing Open Systems SnapVault on UNIX and Linux platforms . . . . .	27
	Verifying the installation . . . . .	35
	Upgrading to Open Systems SnapVault 2.6 . . . . .	36
	Uninstalling Open Systems SnapVault . . . . .	39
	Unattended installation and upgrade . . . . .	42
<b>Chapter 3</b>	<b>Configuring Open Systems SnapVault.</b> . . . . .	49
	Configuration interfaces . . . . .	50
	Understanding the Configurator utility interface. . . . .	51
	Understanding the svsetstanza command . . . . .	53
	Running the Configurator utility . . . . .	56
	Confirming that services are running . . . . .	57
	Modifying Open Systems SnapVault parameters . . . . .	58
	Enabling and disabling debugging . . . . .	62
	Setting block-level incremental backup options . . . . .	66
	Configuring backup exclusion lists . . . . .	68
	Configuring open file backup for Windows . . . . .	71
	Configuration for preserving Snapshot copies . . . . .	74
	Configuration for DataFabric Manager restore to non-ASCII path . . . . .	75

<b>Chapter 4</b>	<b>Perform Backup and Restore</b> . . . . .	77
	Perform SnapVault backup on Open Systems platforms . . . . .	78
	Configuring the SnapVault secondary storage system . . . . .	79
	Creating an initial baseline copy . . . . .	81
	Scheduling SnapVault update backups . . . . .	82
	Backing up empty source directories . . . . .	83
	Perform SnapVault restore on Open Systems platform . . . . .	84
	Restoring a directory or a file . . . . .	85
	Restoring an entire primary storage system . . . . .	89
	Restoring files to a primary storage system from tape . . . . .	90
<b>Chapter 5</b>	<b>Open Systems SnapVault Management</b> . . . . .	91
	Locating status and problem reports . . . . .	92
	Backing up and restoring the Open Systems SnapVault database . . . . .	93
	Backing up and restoring Windows System State data . . . . .	98
	Adding System State data backup . . . . .	102
	Restoring System State data . . . . .	103
	Using System State data backup to rebuild a primary storage system . . . . .	105
	Deleting and re-creating Open Systems SnapVault relationships . . . . .	108
	Migrating a relationship between two secondary storage systems . . . . .	109
	Migrating between two volumes on one secondary storage system . . . . .	112
	Setting up a tertiary system for a relationship . . . . .	115
	Reusing a deleted or renamed primary backup root directory name . . . . .	117
	Reusing a renamed Open Systems SnapVault primary host name . . . . .	119
	Renaming a SnapVault secondary volume . . . . .	120
	Resynchronizing restored or broken relationships . . . . .	122
	Retrying failed transfers . . . . .	125
	Encrypted File System (EFS) file backup and restore . . . . .	127
<b>Chapter 6</b>	<b>Open Systems SnapVault Space Estimator</b> . . . . .	129
<b>Chapter 7</b>	<b>Open Systems SnapVault solution for VMware</b> . . . . .	139
	VMware terminology . . . . .	140
	Overview of Virtualization and VMware ESX Server . . . . .	142

	Open Systems SnapVault on ESX Server . . . . .	143
	Installing Open Systems SnapVault 2.6 on ESX Server . . . . .	145
	Configuration of Open Systems SnapVault for ESX Server . . . . .	146
	Backup and restore of virtual machines . . . . .	153
	Open Systems SnapVault support for VMotion . . . . .	158
<b>Appendix A</b>	<b>The OSSVINFO Tool . . . . .</b>	<b>159</b>
<b>Appendix B</b>	<b>Open Systems SnapVault Error Messages. . . . .</b>	<b>161</b>
	<b>Index . . . . .</b>	<b>179</b>



# Preface

---

- About this guide** This guide describes how to use the Open Systems SnapVault® software product on Windows®, UNIX®, and Linux® platforms.
- 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 “storage system” is used, the information applies to all Network Appliance™ storage system models.
- In this document, the term “storage system” is also referred to as “filer”.
- ACL (access control list)—A list of access control entries (ACEs) that provide information about the users and groups that are allowed access to an object.
- 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.

<b>Formatting convention</b>	<b>Type of information</b>
<i>Italic type</i>	<ul style="list-style-type: none"> <li>◆ Words or characters that require special attention.</li> <li>◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fcstest <i>adaptername</i></code> command, you enter the characters “fcstest” followed by the actual name of the adapter.</li> <li>◆ Book titles in cross-references.</li> </ul>
Monospaced font	<ul style="list-style-type: none"> <li>◆ Command and daemon names.</li> <li>◆ Information displayed on the system console or other computer monitors.</li> <li>◆ The contents of files.</li> </ul>
<b>Bold monospaced font</b>	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.

<b>What is in this guide...</b>	<b>What it means...</b>
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.

---

---

### **Attention**

An attention notice 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

---

## Open Systems SnapVault overview

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:

- ◆ Microsoft Windows
- ◆ Red Hat® Enterprise Linux
- ◆ Novell® SUSE® Linux Enterprise Server
- ◆ Sun Solaris™
- ◆ IBM AIX®
- ◆ HP HP-UX®
- ◆ VMware® ESX Server

---

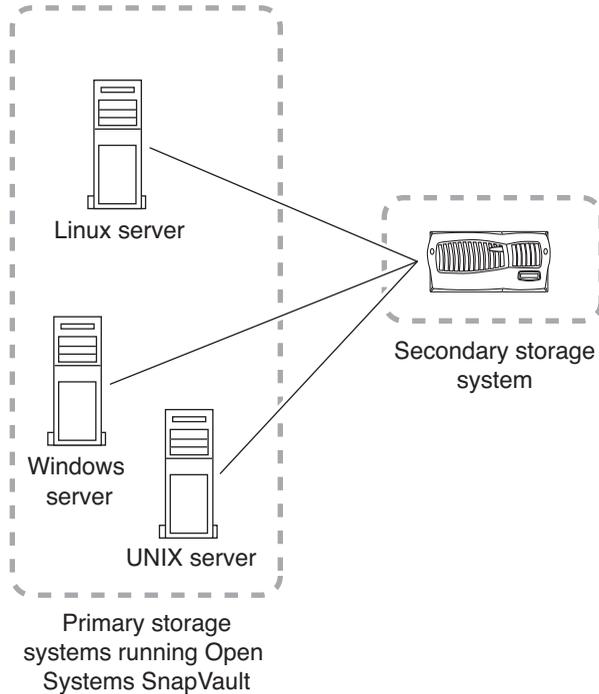
### Note

For a list of currently supported versions of these platforms, see “[Requirements for primary storage systems](#)” on page 16.

---

## Components of the Open Systems SnapVault environment

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



1. The primary storage system—the system from which you are going to back up data  
For a list of currently supported primary storage systems, see [“Requirements for primary storage systems”](#) on page 16.
2. The Open Systems SnapVault agent—the software that you purchased from Network Appliance  
This software is installed on the primary storage system.
3. The secondary storage system—the Network Appliance storage system to which you are going to back up data from the primary storage system  
For information about supported secondary storage systems, see [“Requirements for SnapVault secondary storage systems”](#) on page 19.

## How Open Systems SnapVault works

You identify the directories or file systems on the primary storage system from which to back up data, and the qtrees on the secondary storage 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 storage system. This transfer establishes a SnapVault relationship between the Open Systems platform directories and the SnapVault secondary qtrees.

You can initiate subsequent transfers either manually or configure for automatic, unattended transfer (scheduled). You specify the schedules on the secondary storage system by using the commands available in Data ONTAP, or by using an optional 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, Open Systems 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, beginning with 0 (zero) for the most recent and so on.

If directory or file data needs to be restored to the primary storage system, SnapVault retrieves the data from the specified Snapshot copy and transfers the data back to the primary storage 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 SnapVault backup and restore operations

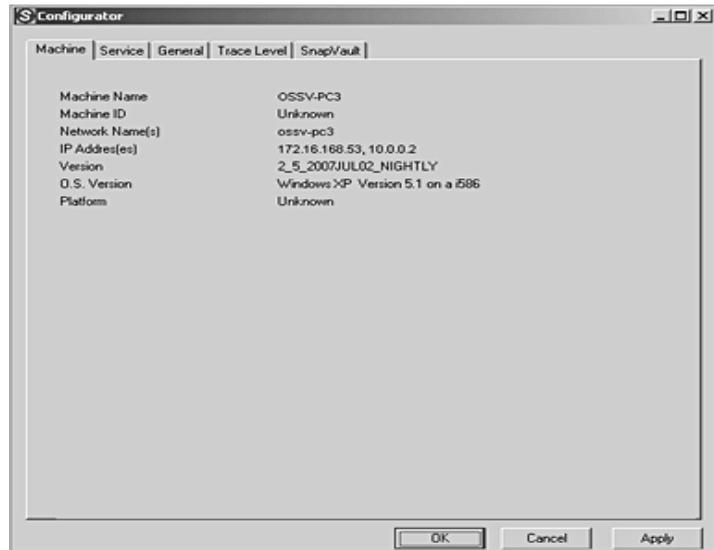
Although Open Systems SnapVault backs up and restores data in a manner similar to how SnapVault backs up and restores data from NetApp storage systems, the following differences exist:

- ◆ With Open Systems SnapVault, you can back up and restore qtrees, volume, and any contents of directories from a non-NetApp primary storage system (for example, Windows or UNIX) to a NetApp secondary storage system (Data ONTAP), whereas with SnapVault, you can back up or restore data only from a NetApp primary storage system to a NetApp secondary storage system.
- ◆ With Open Systems SnapVault, you can restore a single file using NFS or CIFS.

## Administration interfaces for Open Systems SnapVault

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

- ◆ Use 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 51.

- ◆ Use 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 53.

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 storage system by using Open Systems SnapVault, from the command line of the SnapVault secondary storage system, enter the following command:

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

*winserver* is the name of the Windows primary storage system.

C:\ is the Windows system's C: drive.

*/vol/volname/winserver\_C* is the path on the secondary storage system to the SnapVault qtree for the backup.

**Sample restore command:** To restore a subdirectory from the backup created in the `start` command in the preceding sample command, 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 storage 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 storage 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 86.

## 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.

---

**Note**

NDMP is used by the management application.

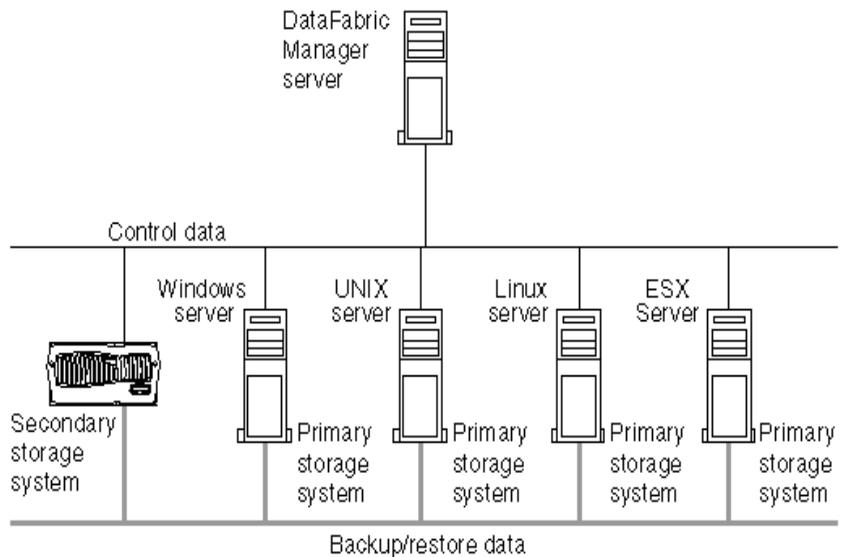
---

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

Vendor	URL
Network Appliance	<a href="http://www.netapp.com">www.netapp.com</a>

Vendor	URL
BakBone Software® NetVault®	<a href="http://www.bakbone.com">www.bakbone.com</a>
Syncsort® Backup Express®	<a href="http://www.syncsort.com">www.syncsort.com</a>
CommVault® Galaxy	<a href="http://www.commvault.com">www.commvault.com</a>

The following illustration shows a typical Open Systems SnapVault setup using the Protection Manager component within the DataFabric Manager.



DataFabric Manager provides infrastructure services, such as discovery, monitoring, role-based access control, auditing, logging for products in the NetApp storage and data suites. DataFabric Manager software runs on a separate workstation or server. It does not run on the storage systems.

Operations Manager is the Web interface of DataFabric Manager. It is used for day-to-day monitoring, alerting, and reporting on storage and appliance infrastructure.

Protection Manager provides policy-based data protection using NetApp storage systems that have SnapVault or SnapMirror® licenses.

## 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 NOW™ 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)—Optionally used for Windows 2000 platforms  
OFM is an optionally licensed feature for protecting open files on Windows Server 2000. It is automatically installed on the primary storage 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 20.
- ◆ Volume Shadow Snapshot copy Service (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 71.

**Common Snapshot Management:** In releases prior to Open Systems SnapVault 2.6, whenever the transfer of files failed, the Snapshot copies were deleted and a new Snapshot copy was created during the transfer restart.

However, Open Systems SnapVault 2.6 has the ability to retain old Snapshot copies and to use these copies subsequently during transfer restarts. This process is termed as common Snapshot management.

A Snapshot copy is a point-in-time copy of a file system that facilitates data recovery. You can access Snapshot copies to recover from accidental deletions. Whenever there is a transfer of a data set or files from a primary storage system to a secondary storage system, Snapshot copies the files and keeps a backup. Thus, if the connection between the primary storage system and the secondary storage system aborts or the transfer of files restart, the data is retained from the Snapshot copy.

For more information, see “[Configuration for preserving Snapshot copies](#)” on page 74.

### **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 68.

By default, the following files are not backed up by Open Systems SnapVault:

- ◆ pagefile.sys
- ◆ hibernate.sys

Additionally, files under the registry key `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\BackupRestore\FilesNotToBackup` are also not backed up.

### **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, compared to baseline or full backups.

**Recognizing files by using the Open Systems SnapVault primary agent:** 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. As 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 67.

For information on how to configure or change BLI settings in Open Systems SnapVault, see [“Setting block-level incremental backup options”](#) on page 66.

## **Backing up and restoring the Open Systems SnapVault database**

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

If the Open Systems SnapVault database becomes corrupt or gets out-of-sync with the secondary storage system, data transfers between the primary and secondary storage 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 storage system to the secondary storage 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 93.

## Backing up and restoring Windows System State

You can back up and restore Windows System State data by using Open Systems SnapVault. This can be useful when, 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.

**Windows EventLog support:** With Open Systems SnapVault support for Windows EventLog, you can maintain the records of all the events that occur in the system. The EventLog files record information about all that is happening in a system at any given point in time. It is necessary to record the events to help you carry out tasks, for example, troubleshooting problems or capacity planning.

For more information, see [“Backing up and restoring Windows System State data”](#) on page 98.

## Checkpoints for restart of transfers

When an Open Systems SnapVault backup process fails, checkpoint restart support allows the backup transfer (baseline or 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. 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 125.

**Block-level checkpoints:** Prior to Open Systems SnapVault 2.6 release, checkpoints were allowed only at the end of files. The checkpoint mechanism was not useful if large files were transferred, as checkpoint can be taken only after the end of the file transfer. This resulted in sending all the file data again, if there was a transfer failure.

For such transfers, Open Systems SnapVault 2.6 supports the following improvements in the checkpoint mechanism:

- ◆ Enabling checkpoints at block levels inside files—This improvement is useful, especially when the data set contains large files (greater than 100 MB). Checkpoints are allowed inside files, therefore you can restart the transfer even from the middle of a file.
- ◆ Configuring checkpoint intervals—You can configure a checkpoint interval. The default value is set to 300 seconds (5 minutes).

For more information on configuration of block-level checkpoints, see [“Configuring the checkpoint interval”](#) on page 125.

## **Encrypted File System (EFS) files**

You can back up and recover EFS files on 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 127.

## **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 storage 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 129.

## **Resynchronizing a broken relationship**

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

For more information, see [“Resynchronizing restored or broken relationships”](#) on page 122.



**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 16
- ◆ [“Installing Open Systems SnapVault on Windows platforms”](#) on page 22
- ◆ [“Installing Open Systems SnapVault on UNIX and Linux platforms”](#) on page 27
- ◆ [“Verifying the installation”](#) on page 35
- ◆ [“Upgrading to Open Systems SnapVault 2.6”](#) on page 36
- ◆ [“Uninstalling Open Systems SnapVault”](#) on page 39
- ◆ [“Unattended installation and upgrade”](#) on page 42

# 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 20 and *Open Systems SnapVault 2.6 Release Notes* for the latest 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 2.6 agent on the following systems.

<b>Operating system</b>	<b>Software versions</b>	<b>File systems</b>	<b>ACLs supported</b>
Windows 2000 and 2003 on x86 and x86-64/EM64T compatible hardware	Windows 2000 Server and Advanced Server for x86.	NTFS	Yes
	Windows Server 2003 Standard Edition, Windows Server 2003 Standard x64 Edition, Windows Server 2003 Enterprise Edition, and Windows Server 2003, Enterprise x64 Edition.  <b>Note</b> _____“Windows 2003” in this document refers to Windows Server 2003 and Windows Storage Server. However, NetApp Host Agent is not supported on Windows Storage Server.	NTFS	Yes
Linux on x86 and x86-64/EM64T compatible hardware	Red Hat Enterprise Linux 4.0 ES/AS/WS	ext2 and ext3	Yes, support for ext3 only
	Red Hat Enterprise Linux 5.0 AS/ES/WS	ext2 and ext3	
	SUSE Linux Enterprise Server 9	ext2, ext3, JFS, XFS, and ReiserFS	Yes, supported only on 32-bit kernel
	SUSE Linux Enterprise Server 10	ext2, ext3, JFS, XFS, and ReiserFS	Yes

Operating system	Software versions	File systems	ACLs supported
Solaris on UltraSPARC systems	Solaris 9	UFS	Yes
	Solaris 10	UFS	Yes
<p><b>Note</b>_____</p> <p>VxFS 3.5, and VxFS 4.0 file systems for Solaris are not supported by Open Systems SnapVault 2.6.</p> <p>_____</p>			
AIX on PowerPC® and POWER® processor-based systems	5L versions 5.1, 5.2, and 5.3	JFS1 and JFS2	Yes (AIXC type ACLs on JFS1 and JFS2)
<p><b>Note</b>_____</p> <p>The IBM® Cluster 1600 and RS/6000 Scalable Parallel (SP) platforms are not supported.</p> <p>_____</p>			
HP-UX on PA-RISC®	HP-UX 11.23	HFS and JFS	Yes
	HP-UX 11.31	HFS and JFS	Yes
ESX Server	ESX Server 3.0, ESX Server 3.0.1, ESX Server 3.0.2	VMFS	NA

**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 2.6 binaries are available as 32-bit (in 32-bit systems) and 64-bit (in 64-bit systems).

\_\_\_\_\_

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

- ◆ A minimum of 128 MB of memory for Windows
- ◆ A minimum of 256 MB of memory for Linux and ESX Server
- ◆ A minimum of 512 MB of memory for AIX, HP-UX, Solaris

- ◆ 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 on the 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 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 129.

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 7.0 or later.

### **Supportability matrix of Open Systems SnapVault with NetApp Host Agent**

The Open Systems SnapVault 2.6 release includes NetApp Host Agent 2.3.1, providing better manageability through DataFabric Manager. During installation, Open Systems SnapVault installs 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	<ul style="list-style-type: none"> <li>◆ SUSE Linux Enterprise Server 9</li> <li>◆ Red Hat Enterprise Linux 4.0</li> </ul>
Windows on x86 and x86-64/EM64T	<ul style="list-style-type: none"> <li>◆ Windows 2000 Server</li> <li>◆ Windows 2003 Server (32-bit and 64-bit)</li> </ul>

## License requirements

The Open Systems SnapVault licenses are installed on the secondary storage system to which they are backed up. You will need to install the following licenses on the secondary storage 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 systems
  - ❖ `sv_linux_pri`—For Linux systems
  - ❖ `sv_vmware_pri`—For VMware ESX Server
- ◆ Optionally, you can install the `sv_windows_ofm_pri` license if you want to use OFM for Windows Server 2000. For more information about OFM, see [“Backing up open files”](#) on page 9.

For softlocks, SnapMirror licenses are required for the secondary and tertiary systems. For more information on softlocks, see [“Setting up a tertiary system for a relationship”](#) on page 115.

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 the Open Systems SnapVault 2.6 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

- ◆ 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
- ◆ Any application cluster software such as Microsoft Cluster Server
- ◆ 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.
- ◆ Resynchronizing restored subdirectories and single files
- ◆ Network interface card (NIC) teaming or trunking

# Installing Open Systems SnapVault on Windows platforms

---

## Installing the Windows agent from the CD-ROM

To install Open Systems SnapVault on a 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	<p>From the command line or Windows explorer, navigate to the CD-ROM directory specific to the platform:</p> <ul style="list-style-type: none"><li>◆ For Windows 2000, go to \WIN2K</li><li>◆ For Windows 2003 x32, go to \WIN2K3x32</li><li>◆ For Windows 2003 x64, go to \WIN2K3x64</li></ul> <p>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>).</p>
4	After navigating to the appropriate directory for the operating system, see " <a href="#">Using the installation wizard</a> " to complete the installation.

## Installing the Windows agent from NOW

To download 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.
2	Go to the Download Software page of the NOW site at <a href="http://now.netapp.com/NOW/cgi-bin/software/">http://now.netapp.com/NOW/cgi-bin/software/</a> and follow the directions to download the appropriate Open Systems SnapVault package for the platform.
3	Uncompress the downloaded package into a temporary directory on the Windows platform.
4	Navigate to the temporary directory where you uncompressed the files for the system.
5	Follow the instructions in “ <a href="#">Using the installation wizard</a> ” to complete the installation.

## Using the installation wizard

To install the Open Systems SnapVault agent on a Windows system using the installation wizard, complete the following steps.

Step	Action
1	Locate and double-click the Setup.exe file. <b>Result:</b> The Open Systems SnapVault Setup Wizard is launched.
2	Follow the instructions on the screen. Click Next. <b>Result:</b> The Open Systems SnapVault Setup License Agreement window appears.
3	To accept the license agreement, click I Agree, and click Next. <b>Result:</b> The Choose NDMP user name and 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	<p><b>Optional:</b> Perform any of the following actions:</p> <ul style="list-style-type: none"> <li>◆ Type your user name and password, retype your password to confirm it and click Next.</li> <li>◆ Directly, click Next.</li> </ul> <p><b>Result:</b> The NDMP Listen Port window appears.</p> <p>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.</p>
5	<p>Perform any of the following actions:</p> <ul style="list-style-type: none"> <li>◆ Enter the NDMP listening port in the NDMP Listen Port window when prompted. Click Next.</li> <li>◆ Accept the default listening port setting of 10000, and click Next.</li> </ul> <p><b>Note</b>_____</p> <p>If another application uses 10000 as its listening port, choose a port number greater than 10000 that is currently not being used.</p> <p>_____</p>
6	<p>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 on the primary storage system.</p> <p><b>Note</b>_____</p> <p>If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas(,).</p> <p>_____</p> <p>Only the secondary storage systems named in this field are accepted by the primary storage system as valid backup systems.</p>

Step	Action						
7	<p>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.</p> <p><b>Note</b>_____</p> <p>If the Open Systems SnapVault agent was previously installed and then uninstalled, files might still be resident in the installation directory; a dialog box appears, asking if you want to empty the directory. To empty the directory, select the Yes radio button before you click Next.</p> <p>_____</p> <p><b>Result:</b> The Install NetApp Host Agent window appears.</p>						
8	<p>The Install NetApp Host Agent window prompts you to choose install NetApp Host Agent; select Yes or No.</p> <table border="1" data-bbox="473 760 1243 1097"> <thead> <tr> <th data-bbox="473 760 835 819">If...</th> <th data-bbox="835 760 1243 819">Then...</th> </tr> </thead> <tbody> <tr> <td data-bbox="473 819 835 1010">You select Yes, and you do not want the NetApp Host Agent to be installed on the default location</td> <td data-bbox="835 819 1243 1010">Select an alternate location on which you want to install the NetApp Host Agent, and click Next. The Confirm Installation window appears.</td> </tr> <tr> <td data-bbox="473 1010 835 1097">You select No</td> <td data-bbox="835 1010 1243 1097">NetApp Host Agent will not be installed.</td> </tr> </tbody> </table> <p>To install NetApp Host Agent later, complete the following steps:</p> <ol style="list-style-type: none"> <li data-bbox="534 1166 1189 1229">a. Using the command-line interface (CLI), go to the Open Systems SnapVault install location.</li> <li data-bbox="534 1253 1189 1357">b. Run the following commands: <ul style="list-style-type: none"> <li data-bbox="584 1305 795 1328"><b>cd manageability</b></li> <li data-bbox="584 1333 1024 1357"><b>InstallHostAgent.exe (for Windows)</b></li> </ul> </li> </ol>	If...	Then...	You select Yes, and you do not want the NetApp Host Agent to be installed on the default location	Select an alternate location on which you want to install the NetApp Host Agent, and click Next. The Confirm Installation window appears.	You select No	NetApp Host Agent will not be installed.
If...	Then...						
You select Yes, and you do not want the NetApp Host Agent to be installed on the default location	Select an alternate location on which you want to install the NetApp Host Agent, and click Next. The Confirm Installation window appears.						
You select No	NetApp Host Agent will not be installed.						

Step	Action	
9	<p>Click Next to start the installation process and wait until the Installation Complete window appears with a message similar to the following:</p> <p>OSSV has been successfully installed.</p> <p>Click "Close" to exit.</p>	
10	<b>For Windows 2000 Server...</b>	<b>For Windows Server 2003/ Windows Storage Server...</b>
	Reboot the system when prompted if you plan to use OFM; otherwise, go to <a href="#">Step 11</a> .	Go to <a href="#">Step 11</a> .
11	<p>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.</p>	

# Installing Open Systems SnapVault on UNIX and Linux platforms

---

## Before you proceed with the installation

Read “[Requirements for primary storage systems](#)” on page 16 to ensure that you are installing the Open Systems SnapVault software on one of the *supported* UNIX or Linux platforms.

## Installing the Solaris agent from the 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 on the system.
4	Run the Solaris pkgadd utility by using the following command:  <code>pkgadd -d <i>path_to_package/ossv</i> -a <i>path_to_package/ossv ossv</i></code>  <i>path_to_package</i> is the full path to where the Open Systems SnapVault package is located; for example, <code>/export/home/packages/ossv ossv</code> .  <b>Result:</b> The package installation script is initiated and asks you a series of questions prior to and during the installation process.
5	Go to “ <a href="#">Using the Solaris pkgadd utility to install the Solaris agent</a> ” 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 <a href="http://now.netapp.com/NOW/cgi-bin/software/">http://now.netapp.com/NOW/cgi-bin/software/</a> and follow the directions to download the Open Systems SnapVault Solaris agent package.
3	<p>After you download the Open Systems SnapVault package, unpack and start the installation program:</p> <ul style="list-style-type: none"><li>◆ Uncompress and untar the downloaded pkgadd package.</li><li>◆ Run the Solaris pkgadd utility using the following command: <code>pkgadd -d <i>path_to_package/ossv</i> -a <i>path_to_package/ossv</i></code> <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 on the system.</li></ul> <p><b>Result:</b> The package installation script starts and asks you a series of questions prior to and during the installation process.</p>
4	Go to “ <a href="#">Using the Solaris pkgadd utility to install the Solaris agent</a> ” 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	If you read and agree to the terms of the license, answer yes (y), no (n), or display (d) when asked.	
	<b>If you answer...</b>	<b>Then...</b>
	yes (y)	Go to the next step of installation.
	display (d)	The license is displayed.
	no (n)	You cannot install the Open System SnapVault package.
2	<p>When prompted, enter the path where you want the SnapVault directory to be created. The default location is /usr/snapvault.</p> <p>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).</p>	
3	<p>Enter your user name to connect to the target system using the NDMP protocol.</p> <p>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.</p>	
4	Enter and confirm the password to connect to the system.	
5	<p>Enter the NDMP listener port. The default listener port is 10000.</p> <p>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.</p> <p><b>Note</b> _____  If another application uses 10000 as its listening port, choose a port number greater than 10000 that is currently not being used.  _____</p>	

Step	Action
6	<p>Enter the host names or IP addresses of the SnapVault secondary storage systems that are allowed to perform backups from the primary storage system.</p> <p><b>Note</b> _____            If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas(.).</p> <p>_____</p> <p>Only the secondary storage systems named in this field are accepted by the primary storage system as valid backup systems. A series of installation scripts is executed.</p>
7	<p>If successful, a message similar to the following appears:</p> <pre>Installation of OSSV was successful.</pre>

### 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	<p>From the command line, navigate to the CD-ROM directory specific to the platform:</p> <ul style="list-style-type: none"> <li>◆ For HP-UX, go to /HPUX</li> <li>◆ For AIX, go to /AIX</li> <li>◆ For Linux, go to /LINUX</li> </ul> <p>Depending on the system configuration and mount options, the files and directories mentioned might appear in lowercase characters.</p>
4	<p>From the command line, enter the following command:</p> <pre>./install</pre>

<b>Step</b>	<b>Action</b>	
<b>5</b>	When the install program prompts you to read and agree to the terms of the license, answer y (yes), n (no), or d (display).	
	<b>If you answer...</b>	<b>Then...</b>
	yes (y)	Go to the next step of installation.
	display (d)	The license is displayed.
<b>6</b>	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.	
	<b>7</b>	<b>If...</b>
	You plan to manage Open Systems SnapVault backup through the Data ONTAP command line	Press Enter.
	You plan to manage Open Systems SnapVault backup through a commercial NDMP application, such as DataFabric Manager	When the install program prompts you for the user name, password, and NDMP listening port, specify a user name and password authorized through that application.
<b>8</b>	When the install program prompts you for Allowed Systems, enter one or more host names or IP addresses of the SnapVault secondary storage system or systems to which you want to back up the primary storage system.	
	<p><b>Note</b> _____</p> <p>If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas(,) not spaces.</p> <p>_____</p>	

Step	Action	
9	In the Linux environment, when the install program prompts you to choose whether to install NetApp Host Agent, enter y (yes) or n (no).	
	<b>If you answer...</b>	<b>Then...</b>
	n (no)	NetApp Host Agent will not be installed.
	y (yes)	Complete the following steps to install NetApp Host Agent: <ul style="list-style-type: none"> <li>a. Using the command-line interface (CLI), go to the Open Systems SnapVault installation location.</li> <li>b. Run the following commands:               <pre>cd manageability ./InstallHostAgent.sh (for Linux)</pre> </li> </ul>
10	A message similar to the following 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 <a href="http://now.netapp.com/NOW/cgi-bin/software/">http://now.netapp.com/NOW/cgi-bin/software/</a> and follow the directions to download the appropriate Open Systems SnapVault package.	
3	<b>If...</b>	<b>Then...</b>
	You are installing Open Systems SnapVault on a Linux platform	Enter the following commands after downloading the Open Systems SnapVault package:  <b><i>gunzip package_name</i></b> <b><i>tar -xvf tar_file_name</i></b>
	You are installing Open Systems SnapVault on any other UNIX platform	Enter the following commands after downloading the Open Systems SnapVault package:  <b><i>uncompress package_name</i></b> <b><i>tar -xvf tar_file_name</i></b>  For example, on an AIX primary storage system, uncompress and untar the downloaded ossv_aix_v2.6.tar.Z package using the following commands:  <b><i>uncompress ossv_aix_v2.6.tar.Z</i></b> <b><i>tar -xvf ossv_aix_v2.6.tar</i></b>
4	Navigate to the directory where the untar operation placed the files and enter the following command:  <b><i>./install</i></b>	
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	<b>If...</b>	<b>Then...</b>
	You plan to manage Open Systems SnapVault backup through the Data ONTAP command line	Press Enter.
	You plan to manage Open Systems SnapVault backup through a commercial NDMP application, such as DataFabric Manager	When the install program prompts you for the user name, password, and NDMP listening port, specify a user name and password authorized through that application.
7	<p>When the install program prompts you for Allowed Systems, enter one or more host names or IP addresses of the SnapVault secondary storage system or systems to which you want to back up the primary storage system.</p> <p><b>Note</b> _____            If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas(,), but no spaces.</p> <p>_____</p>	
8	<p>A message similar to the following is displayed to indicate a successful installation:</p> <pre>Installation completed successfully.</pre>	
9	<p>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.</p>	

# 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: <b>ping secondary_system</b> <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 “ <a href="#">Enabling and disabling security</a> ” on page 58.
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 “ <a href="#">Memory and port requirements:</a> ” on page 18 are kept open.
5	Navigate to <i>install_dir/bin</i> and run <b>svinstallcheck</b> 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 2.6

---

## Prerequisites

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

---

### Note

Installing 32-bit Open Systems SnapVault on Windows 2000 and 2003 on x86 and x86-64/EM64T is not supported.

---

## Upgrading to Open Systems SnapVault 2.6

When you upgrade Open Systems SnapVault 2.2 or 2.3 to 2.6, complete the following steps.

Step	Action
1	<p>Stop Open Systems SnapVault services by performing the following actions:</p> <ul style="list-style-type: none"><li>a. Launch the Open Systems SnapVault Configurator utility. For instructions, see <a href="#">“Running the Configurator utility”</a> on page 56.</li><li>b. Click the Service tab.</li><li>c. Click Stop Service to stop Open Systems SnapVault services.</li><li>d. Close the Configurator utility.</li></ul> <p><b>Note</b>— Alternatively, you can use either the <code>svpmgr shutdown</code> or <code>snapvault service stop</code> command.</p>
2	<p>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 <a href="#">“Backing up and restoring the Open Systems SnapVault database”</a> on page 93.</p>

Step	Action
3	<p>Install the Open Systems SnapVault 2.6 agent using the procedure described in “<a href="#">Installing Open Systems SnapVault on Windows platforms</a>” on page 22, or “<a href="#">Installing Open Systems SnapVault on UNIX and Linux platforms</a>” on page 27.</p> <p><b>Note</b> _____  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.</p> <p>_____</p> <p><b>Result:</b> The Open Systems SnapVault agent automatically starts Open Systems SnapVault services after installation.</p>
4	<p><b>a.</b> Launch the Open Systems SnapVault Configurator utility.</p> <p><b>b.</b> Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services.</p> <p><b>c.</b> Close the Configurator utility.</p> <p><b>Note</b> _____  Alternatively, you can use the <code>svpmgr shutdown</code> command.</p> <p>_____</p>
5	<p>Copy the database directory that you saved in <a href="#">Step 2</a> to the location of the database created by the new installation.</p>
6	<p><b>a.</b> Launch the Open Systems SnapVault Configurator utility.</p> <p><b>b.</b> Click the Service tab, and then click Start Service to start Open Systems SnapVault services.</p> <p><b>Note</b> _____  Alternatively, you can also use the <code>svpmgr startup</code> command.</p> <p>_____</p>

**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 the upgrade to Open Systems SnapVault 2.6, ACLs are backed up only for changed or newly added files.

Add the following stanza to snapvault.cfg, to avoid such a partial backup of ACLs after upgrading to Open Systems SnapVault 2.6:

```
[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 Open Systems SnapVault

## Uninstalling the Open Systems SnapVault agent on Windows

To uninstall the Open Systems SnapVault 2.6 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 window appears.
4	To uninstall NetApp Host Agent, select Yes.  <b>Note</b> ——— If you want to uninstall NetApp Host Agent later, select No. You can uninstall NetApp Host Agent later from Add/Remove Programs Wizard. —————
5	Click Finish to remove Open Systems SnapVault from the computer.  <b>Result:</b> 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 Open Systems SnapVault agent on Solaris

To uninstall the Open Systems SnapVault 2.6 agent on a Solaris platform, complete the following steps.

Step	Action
1	Log in to the primary storage system as root.
2	Run the <code>pkgrm</code> command:  <b><code>pkgrm ossv</code></b>

Step	Action
3	<p>Enter Yes (y) when asked whether you want to remove this package. The script responds with text similar to the following:</p> <pre>## Removing installed package instance &lt;ossv&gt;</pre> <p>This package contains scripts which will be executed with super-user permission during the process of removing this package.</p>
4	<p>Answer Yes (y) when asked whether you want to continue with the removal of this package.</p> <p><b>Result:</b> If removal is successful, the script responds with text similar to the following:</p> <pre>Removal of &lt;ossv&gt; was successful.</pre>

## Uninstalling the Open Systems SnapVault agent on HP-UX, AIX, or Linux

To uninstall the Open Systems SnapVault 2.6 agent on the HP-UX, AIX, or Linux platform, complete the following steps.

Step	Action
1	Enter the following command: <b>install_dir/util/uninstall</b>
2	During uninstillation 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.  NetApp Host Agent can also be installed independently.
3	If the SnapVault directory still appears after running <code>uninstall</code> , remove the directory manually.

# Unattended installation and upgrade

---

## What unattended installation and upgrade is

The 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 method, 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 47.

---

## Supported platforms and Open Systems SnapVault versions

All platforms on which Open Systems SnapVault is currently supported can be upgraded or installed using this method. See [“Requirements for primary storage systems”](#) on page 16 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 storage 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 storage 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:

- ◆ The 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. For details on the installation script, see “[Unattended installation script](#)” on page 47. To generate the installation script and the files, complete the following steps.

Step	Action
1	On a primary storage system whose configuration settings you want to use for other installations, do the following: <ul style="list-style-type: none"> <li>a. Launch the Open Systems SnapVault Configurator utility. For instructions, see “<a href="#">Running the Configurator utility</a>” on page 56.</li> <li>b. Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services.</li> </ul>
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 storage system.

Step	Action	
5	<b>If...</b>	<b>Then...</b>
	<p>You only want to save the configuration settings to a file</p>	<p>Enter the following command:</p> <pre><b>svconfigpackager filename</b></pre> <p><i>filename</i> is the name of the configuration settings file.</p> <p><b>Example:</b> To create a configuration settings file called <code>svconfig.in</code>, enter the following command:</p> <pre>svconfigpackager svconfig.in</pre>
	<p>You want to save the configuration settings to a file and also create an installation script for installers</p> <p><b>Note</b> _____</p> <p>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.</p> <p>_____</p>	<p>Enter the following command:</p> <pre><b>svconfigpackager [-h -i installation path] package name</b></pre> <p><i>installation path</i> is the Open Systems SnapVault installation directory for the unattended installs.</p> <p><i>package name</i> is the name of the configuration settings file.</p> <p><b>Example:</b> 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:</p> <pre>svconfigpackager -i /usr/snapvault svconfig.in</pre>
<p>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)</p>	<p>Enter the following command:</p> <pre><b>svconfigpackager -h -i path_name filename</b></pre> <p><i>path_name</i> is the directory where the configuration settings file and the installation script will be placed.</p>	

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

## 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 necessary files, complete the following steps.

Step	Action
1	<p>Perform any of the following actions:</p> <ul style="list-style-type: none"> <li>◆ Download and uncompress the Open Systems SnapVault package from the NOW site.</li> <li>◆ Copy the Open Systems SnapVault package to the installation directory on the primary storage system where you want to install or upgrade.</li> </ul>

Step	Action	
2	Copy the installation script and other files that were generated in <a href="#">Step 6</a> of “ <a href="#">Performing an unattended installation or upgrade</a> ” on page 45 to the directory on the primary storage system where you uncompressed the installation package.	
3	<b>If...</b>	<b>Then...</b>
	The primary storage system is a Windows machine	Enter the following command to start the unattended installation or upgrade: <b>unattinstall.bat</b> 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\ <i>Current User</i> \Local Settings\Temp
	The primary storage system is a UNIX machine	Enter the following command to start the unattended installation or upgrade: <b>./unattinstall.sh</b> If the installation does not succeed, the error messages are logged in the /tmp directory.
4	On Windows, run the svinstallcheck utility to verify the successful installation and to make sure that the services are running properly. On UNIX, svinstallcheck is automatically run after the installation. If any errors are found, they are logged in the SnapVault log files in the <i>install_dir/etc</i> 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 on 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 the `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, a message similar to the following 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 GUI utility called Configurator or the `svsetstanza` command on the 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 51
- ◆ [“Running the Configurator utility”](#) on page 56
- ◆ [“Confirming that services are running”](#) on page 57
- ◆ [“Modifying Open Systems SnapVault parameters”](#) on page 58
- ◆ [“Enabling and disabling debugging”](#) on page 62
- ◆ [“Setting block-level incremental backup options”](#) on page 66
- ◆ [“Configuring backup exclusion lists”](#) on page 68
- ◆ [“Configuring open file backup for Windows”](#) on page 71
- ◆ [“Configuration for preserving Snapshot copies”](#) on page 74
- ◆ [“Configuration for DataFabric Manager restore to non-ASCII path”](#) on page 75

# 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 51
- ◆ [“Understanding the svsetstanza command”](#) on page 53

## Understanding the Configurator utility interface

---

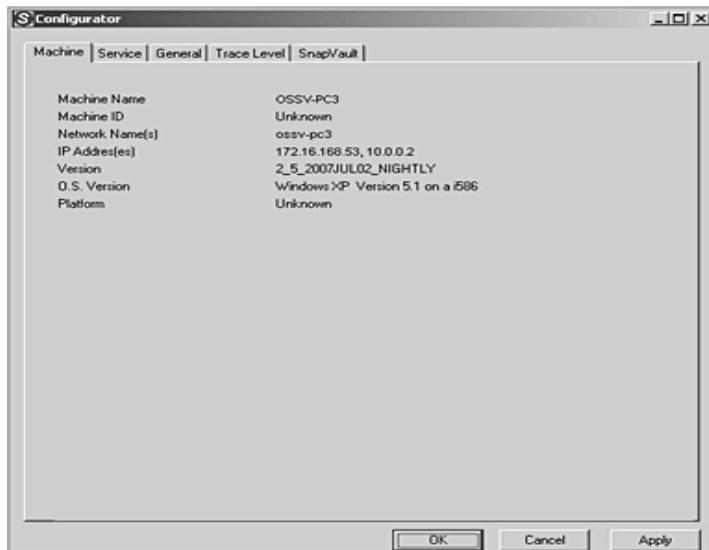
### About the Configurator utility interface

The Configurator utility interface is the GUI 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 and 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 92.

**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 files are located in the *install\_dir/snapvault/config* directory.

File name	Description
configure.cfg	The General tab of the Open Systems SnapVault Configurator interface represents the values in this file.
estimator.cfg	The svestimator 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 Configurator interface represents the values in this file.
snapvault.cfg	The SnapVault tab of the Open Systems SnapVault Configurator interface represents the values in this file.

## Understanding the svsetstanza command

---

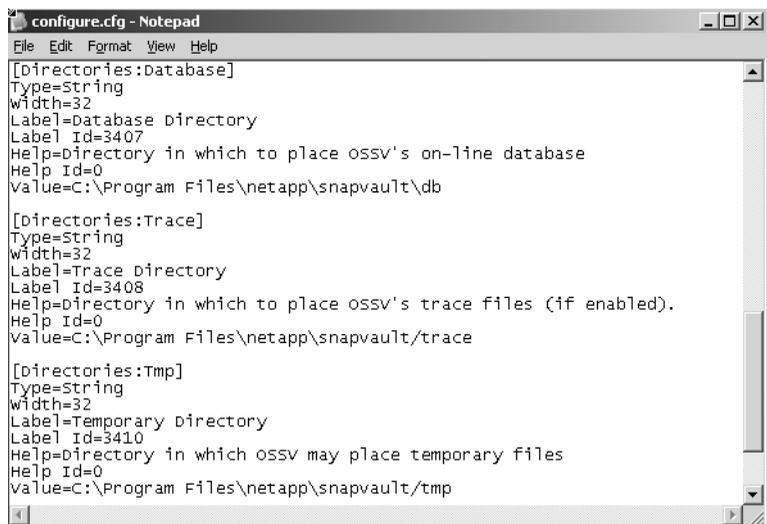
### About the svsetstanza command

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 Configurator utility. The `svsetstanza` command is located in the `install_dir/snapvault/util` directory.

### What the 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 52.

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



```
configure.cfg - Notepad
File Edit Format View Help
[Directories:Database]
Type=String
width=32
Label=Database Directory
Label Id=3407
Help=Directory in which to place OSSV's on-line database
Help Id=0
Value=C:\Program Files\netapp\snapvault\db

[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

[Directories:Tmp]
Type=String
width=32
Label=Temporary Directory
Label Id=3410
Help=Directory in which OSSV may place temporary files
Help Id=0
Value=C:\Program Files\netapp\snapvault/tmp
```

### 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. In the preceding example (image), *configure.cfg* is the file.

*category* is the section to be changed in the configuration file. In the preceding example (image), *Directories* is a category.

*title* is the title of the section to be changed. In the preceding example (image), *Trace* is a title.

*keyword* is the parameter to be changed. In the preceding example (image), *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”.

---

**Note**

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, it is assumed 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	<p>Launch the Configurator utility using one of the following methods:</p> <ul style="list-style-type: none"><li>◆ Click Start &gt; Programs &gt; OSSV &gt; OSSV Configurator on the Windows primary storage system.</li><li>◆ Run the following command on the UNIX primary storage system: <b><code>\$INSTALL_DIR/bin/svconfigurator</code></b> The default location for <code>INSTALL_DIR</code> is <code>/usr/snapvault</code>.</li></ul>
2	<p>Click the appropriate Configurator utility tab to change the settings.</p> <p>The most commonly modified Open Systems SnapVault parameters are accessed by clicking the SnapVault tab.</p>

## Confirming that services are running

---

### Confirming that services are running

To confirm that the 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 the 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](#)” on page 58
- ◆ “[Modifying the Qtree SnapMirror™ access list](#)” on page 58
- ◆ “[Modifying the NDMP settings](#)” on page 59
- ◆ “[Enabling and disabling Windows EventLog](#)” on page 60

## 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 Qtree SnapMirror Access List check box to enable security; clear it to disable security.
3	Click OK.

## Modifying the Qtree SnapMirror™ access list

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

Step	Action
1	Click the SnapVault tab.

Step	Action
2	<p>Add, replace, or delete IP addresses or host names of the secondary storage systems to which you want to back up data.</p> <p>To add two secondary storage system values, enter the secondary storage values separated by a comma.</p> <p>For example, f3070-202-170, r200-192-196</p> <p>The snapvault.cfg file appears as follows:</p> <pre data-bbox="494 494 960 633">[QSM:Access List]   Type=String   Label=QSM Access List   Label Id=108000149   value=f3070-202-170!,r200-192-196</pre> <p><b>Note</b></p> <p>Do not edit the snapvault.cfg file.</p>
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.  <b>Note</b> 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 the 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 Windows EventLog

You can enable or disable support for Windows EventLog as part of the System State backup. The EventLog options are as follows:

- ◆ Application EventLog  
Indicates backup of application logs only
- ◆ Security EventLog  
Indicates backup of security logs only
- ◆ System EventLog  
Indicates backup of system logs only

**Modifying Windows EventLog:** To enable or disable the support for Windows EventLog, complete the following steps.

<b>Step</b>	<b>Action</b>
1	Click the SnapVault tab.
2	Select or clear the application, security, or system event logs check box to enable or disable support for Windows EventLog.  <b>Note</b> _____ You can select one or more EventLog check boxes at a time. _____
3	Click OK.

# 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 impact on performance.

## Enabling the generation of debug files

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

Step	Action
1	In the Configurator GUI, click the General tab.
2	Select the “Generate debugging files” check box.  <b>Note</b> _____ 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.

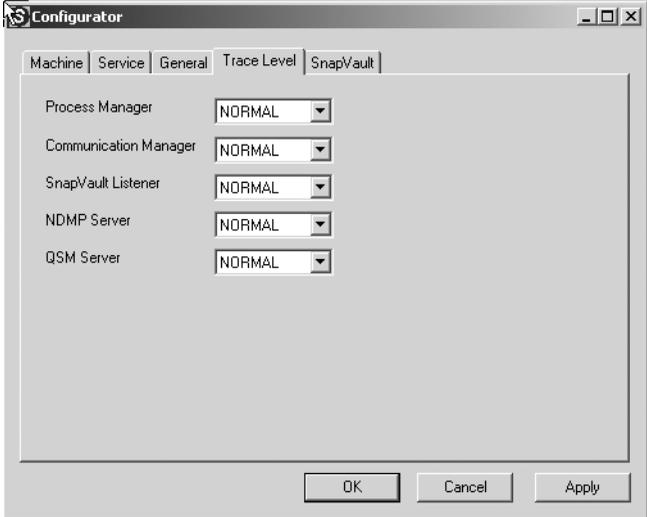
Step	Action
4	<p>From the drop-down list of the service for which you want to generate debug information, select one of the five trace levels—ALWAYS, NORMAL, VERBOSE, LIBNORMAL, or LIBVERBOSE.</p> <p>For example, for Qtree SnapMirror Server, select VERBOSE, as shown in the following image.</p> <div data-bbox="575 491 1147 958" data-label="Image"> <p>The screenshot shows a window titled 'Configurator' with several tabs: Machine, Service, General, Trace Level, and SnapVault. The 'Trace Level' tab is active. It contains a list of services with corresponding trace level dropdown menus:</p> <ul style="list-style-type: none"> <li>Process Manager: NORMAL</li> <li>Communication Manager: NORMAL</li> <li>SnapVault Listener: NORMAL</li> <li>NDMP Server: NORMAL</li> <li>QSM Server: VERBOSE</li> </ul> <p>At the bottom of the window are three buttons: OK, Cancel, and Apply.</p> </div> <p><b>Note</b> Make sure you set the trace level back to NORMAL after you are done troubleshooting.</p>
5	Click Apply.

Step	Action
6	<p>a. Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services.</p> <p>Wait until the Current State displays as Stopped.</p> <p>b. Click Start Service.</p> <p>Wait until the Current State displays as Running.</p> <p>c. Click OK.</p> <p><b>Note</b> _____  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.  _____</p>

### Disabling the generation of debug files

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

Step	Action
1	In the Configurator GUI, click the General tab.
2	Clear the “Generate debugging files” check box.
3	Click the Trace Level tab.

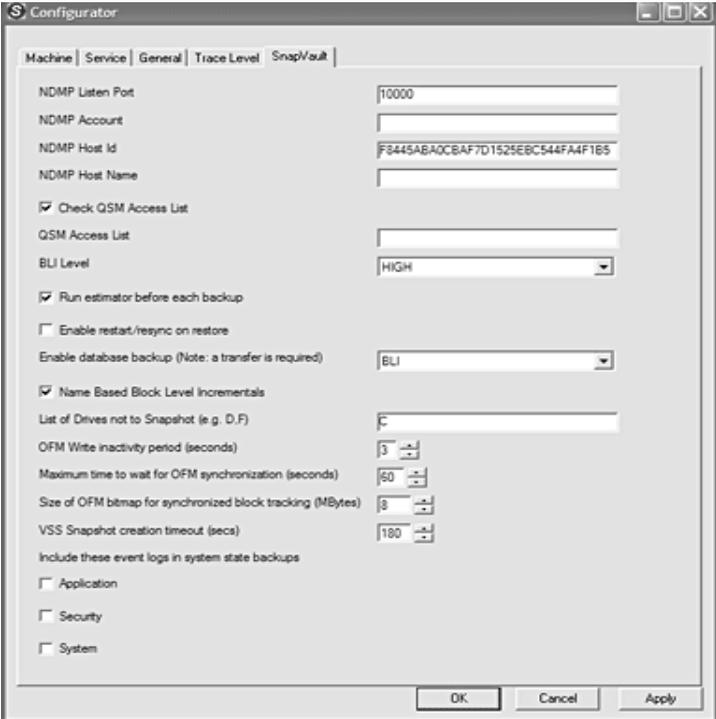
Step	Action
4	<p>Select NORMAL from the drop-down list of the service for which you want to stop generating debug information.</p> <p>For example, for Qtree SnapMirror Server, select NORMAL, as shown in the following image.</p>  <p>The screenshot shows a window titled 'Configurator' with a tabbed interface. The 'SnapVault' tab is selected. Under the 'Trace Level' section, there are five entries, each with a dropdown menu set to 'NORMAL':</p> <ul style="list-style-type: none"> <li>Process Manager: NORMAL</li> <li>Communication Manager: NORMAL</li> <li>SnapVault Listener: NORMAL</li> <li>NDMP Server: NORMAL</li> <li>QSM Server: NORMAL</li> </ul> <p>At the bottom of the dialog are three buttons: 'OK', 'Cancel', and 'Apply'.</p>
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 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	<p>Click the SnapVault tab. The block-level incremental checksum computation options are displayed in the BLI Level drop-down list.</p>  <p>The screenshot shows the 'Configurator' window with the 'SnapVault' tab selected. The 'BLI Level' dropdown menu is set to 'HIGH'. Other visible options include 'NDMP Listen Port' (10000), 'NDMP Account', 'NDMP Host Id' (F3445ABA0C8AF7D1525EBC544F4F165), 'NDMP Host Name', 'Check QSM Access List' (checked), 'QSM Access List', 'Run estimator before each backup' (checked), 'Enable restart/resync on restore' (unchecked), 'Enable database backup (Note: a transfer is required)' (BLI), 'Name Based Block Level Incrementals' (checked), 'List of Drives not to Snapshot (e.g. D.F)' (C), '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), and 'Include these event logs in system state backups' (Application, Security, System).</p>

Step	Action
2	<p>Select the BLI Level backup option that you want, by using the drop-down list.</p> <ul style="list-style-type: none"> <li>◆ HIGH—computes checksums during the initial baseline backup and updates</li> </ul> <p><b>Note</b> _____ The HIGH mode might slow down the initial backup transfer.</p> <hr/> <ul style="list-style-type: none"> <li>◆ LOW—computes checksums only during updates</li> <li>◆ OFF—disables block-level incremental backups altogether</li> </ul>

### 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 20.
- ◆ 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.

## Backing up open files

OFM waits for a set time when there is no write activity 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 storage system
- ◆ Set OFM parameters

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

Step	Action
1	On the secondary storage system, enter the following command: <b>license add code</b> <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	In the Configurator GUI, click the SnapVault tab.
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.  <b>Result:</b> 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.  <b>Result:</b> 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: <code>snapvault modify -o back_up_open_files off</code>

### Setting the VSS Snapshot copy timeout parameter

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.

# Configuration for preserving Snapshot copies

---

## About Common Snapshot Management

Common Snapshot management ensures that the same Snapshot copy is used for backup.

There are two possible configurations for common Snapshot management:

- ◆ **MaxCPRestartWaitTime**  
It is the maximum waiting time a Snapshot copy is retained after the transfer failure. If the transfer of files restarts after the maximum waiting time, the Snapshot copy is lost and needs to be created again. The default waiting time is 10 minutes.
- ◆ **FailCPRestartOnNewSnapshot**  
If the corresponding Snapshot copy is not available during restart of a transfer due to a system restart, elapsed time, or Open Systems SnapVault restart, either you allow the transfer to continue using a new Snapshot copy or abort the transfer.
  - ❖ When the value is set to **TRUE**, the transfer is aborted.
  - ❖ When the value is set to **FALSE**, a new Snapshot copy is created and the transfer continues.

# Configuration for DataFabric Manager restore to non-ASCII path

---

## Restoring backed-up data to a primary storage system

Using DataFabric Manager, the configuration flag [**NDMP: ForceUTF8Encoding**] has to be TRUE in the snapvault.cfg file in order to restore backed-up data to a non-ASCII path.

For Open Systems SnapVault 2.6, the default value is TRUE.

---

### Note

The value is FALSE for releases prior to Open Systems SnapVault 2.6. If you have upgraded to Open Systems SnapVault 2.6, you have to set the value to TRUE.

---



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

**Topics in this chapter** This chapter covers the following topics:

- ◆ [“Perform SnapVault backup on Open Systems platforms”](#) on page 78
- ◆ [“Perform SnapVault restore on Open Systems platform”](#) on page 84

# Perform SnapVault backup on Open Systems platforms

---

## Before configuring SnapVault backup

To set up SnapVault backup on the Open Systems platform, you must prepare the systems and SnapVault secondary storage systems to fulfill their backup tasks. Ensure you have completed the following steps in the order given.

---

### 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 15.
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 79.
3. On the SnapVault secondary storage system, start the baseline transfer. See “[Creating an initial baseline copy](#)” on page 81.
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 82.

## Topics in this section

The following topics are covered in this section:

- ◆ “[Configuring the SnapVault secondary storage system](#)” on page 79
- ◆ “[Creating an initial baseline copy](#)” on page 81
- ◆ “[Scheduling SnapVault update backups](#)” on page 82
- ◆ “[Backing up empty source directories](#)” on page 83

## 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.

To support resync after restore, the SnapVault secondary storage system must be running Data ONTAP 7.0.5, 7.1.2, 7.2, or later.

### 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	<p>License the SnapVault secondary storage system. In the storage system console of the SnapVault secondary storage system, enter the following command:</p> <pre>license add sv_secondary_license</pre> <p><b>Example:</b></p> <pre>license add sv_ontap_sec</pre> <p><b>Note</b> _____ Ensure that you select the correct storage system or NearStore® secondary license.</p> _____
2	<p>License the SnapVault primary storage system. In the storage system console of the SnapVault secondary storage system, enter the following command:</p> <pre>license add ossv_primary_license</pre> <p><b>Example 1 (Windows):</b></p> <pre>license add sv_windows_pri</pre> <p><b>Example 2 (UNIX):</b></p> <pre>license add sv_unix_pri</pre>

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

## 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	<p>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.</p> <ul style="list-style-type: none"><li>◆ Specify the fully qualified path to the Open Systems host directory that you want to back up. Use the <code>-S</code> prefix to indicate the source path.</li><li>◆ Even though the Open Systems platform directory to be backed up has no <code>qtree</code>, you <i>still</i> need to specify a host and path to the <code>qtree</code> where you will back up this data on the SnapVault secondary storage system.</li></ul> <p>Enter the following command:</p> <pre><b>snapvault start -S prim_host:dirpath sec_host:/vol/sec_vol/sec_tree</b></pre> <p><b>Example 1 (Windows):</b></p> <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> <p><b>Example 2 (UNIX):</b></p> <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>

Perform 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 storage system.

You should plan your backup schedules such that 16 or fewer transfers occur at the same time from the same primary storage system. You can use the `snapvault status` command on the secondary storage system to check the number of simultaneous transfers occurring from a primary storage 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	<p>Use the <code>snapvault snap sched</code> command to schedule the updated copying of new or modified data on all Open Systems platform directories that are backed up to qtrees in SnapVault secondary storage.</p> <p>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.</p> <p><b>Example:</b></p> <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> <p><b>Note</b></p> <p>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.</p> <p>After all the secondary qtrees on the specified volume have been updated, SnapVault creates a Snapshot copy of this volume for archival.</p>

Perform SnapVault backup on Open Systems platforms

## Backing up empty source directories

---

Prior to Open Systems SnapVault 2.6, SnapVault updates of empty source directories used to fail with the error “could not read from socket” on the secondary storage system when the directory on the primary storage system was empty.

This error indicated that the transfer failed because the primary storage system closed the TCP socket. To determine the cause of this failure, view the SnapVault log file at *Install\_Path*\etc on the primary storage system. For Windows, the typical installation path is C:\Program Files\netapp\snapvault.

The SnapVault log displays the following error from the failed update:

```
2007/01/17 08:56:44: ERROR      : C:\backup dest-
filer:/vol/ossv/win_C_backup Possible attempt to update empty
mount, aborting. Set config option BackupEmptyMount to override
2007/01/17 08:56:44: ERROR      : C:\backup dest-
filer:/vol/ossv/win_C_backup Failed to generate update inode values
```

To allow an Open Systems SnapVault agent to back up empty primary paths, complete the following steps.

Step	Description
1	Modify the snapvault.cfg file at <i>Install_Path</i> \config on the primary storage system. Open this file using WordPad or Notepad and add the following entry to the bottom of the file:  <b>[QSM:BackupEmptyMount]</b> <b>value=TRUE</b>
2	Save and close this file.
3	Stop and start the Open Systems SnapVault service on the primary storage system. The backups will now succeed.

# Perform SnapVault restore on Open Systems platform

---

## 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 85
- ◆ [“Restoring an entire primary storage system”](#) on page 89
- ◆ [“Restoring files to a primary storage system from tape”](#) on page 90

## 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 Operations Manager (the DataFabric Manager user interface).

### 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 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	<p data-bbox="494 239 1220 265">Enter the <code>snapvault restore</code> command and specify the following:</p> <ul data-bbox="494 279 1224 618" style="list-style-type: none"> <li>◆ The secondary storage system host and the path to the secondary qtree, directory, and file that you want to restore from.</li> <li>◆ The <code>-s</code> option sets 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>).</li> <li>◆ The <code>-k</code> option sets the maximum speed at which data is transferred in kilobytes per second. If this option is not set, the storage system transmits data as fast as it can.</li> <li>◆ The path on the primary storage system to the directory or file that you want to restore to.</li> </ul> <p data-bbox="494 644 1123 673"><b>Example 1 - Single file restore (Windows system):</b></p> <pre data-bbox="494 678 1013 760">snapvault restore -s sv_daily.0 -k 10 -S myvault:/vol/sv_vol/melzdir/evidence.doc a:\melzdir\evidence_restore.doc</pre> <p data-bbox="494 786 1072 815"><b>Example 2 - Single file restore (UNIX system):</b></p> <pre data-bbox="494 821 1013 902">snapvault restore -s sv_daily.0 -k 10 -S myvault:/vol/sv_vol/melzdir/evidence.doc /usr/melzdir/evidence_restore.doc</pre> <p data-bbox="494 928 556 958"><b>Note</b></p> <hr data-bbox="494 946 1224 949"/> <p data-bbox="494 963 942 989">Enter the entire command as a single line.</p> <p data-bbox="494 1025 1220 1116">Ensure that you do not specify a slash (\ or /) character at the end of the path name in that command; otherwise, the <code>snapvault restore</code> command will fail.</p> <hr data-bbox="494 1128 1224 1131"/> <p data-bbox="494 1164 1201 1194"><b>Example 1 - Single directory restore (Windows system):</b></p> <pre data-bbox="494 1199 1220 1251">snapvault restore -s sv_daily.0 -k 10 -S myvault:/vol/sv_vol/melzdir/dir1 a:\melzdir\dir1_restore</pre> <p data-bbox="494 1277 1147 1307"><b>Example 2 - Single directory restore (UNIX system):</b></p> <pre data-bbox="494 1312 1013 1394">snapvault restore -s sv_daily.0 -k 10 -S myvault:/vol/sv_vol/melzdir/dir1 /usr/melzdir/dir1_restore</pre>

## **Using the Operations Manager restore wizard**

The Operations Manager restore wizard leads you through the entire restore process. For details, see the Operations 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, “ <a href="#">Installing the Open Systems SnapVault Software</a> ,” on page 15.
4	<b>Optional:</b> If you backed up the Windows System State data of the primary storage system, restore its Windows System State data. For more information, see “ <a href="#">Restoring System State data</a> ” on page 103.
5	Restore the backed-up directories using the <code>snapvault restore</code> command. For details, see “ <a href="#">Using the snapvault restore command</a> ” on page 86.

## 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 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 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 <a href="#">“Restoring a directory or a file”</a> on page 85.

## About this chapter

Most of the Open Systems SnapVault management tasks you 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 using Open Systems SnapVault:

- ◆ [“Locating status and problem reports”](#) on page 92
- ◆ [“Backing up and restoring the Open Systems SnapVault database”](#) on page 93
- ◆ [“Backing up and restoring Windows System State data”](#) on page 98
- ◆ [“Deleting and re-creating Open Systems SnapVault relationships”](#) on page 108
- ◆ [“Migrating a relationship between two secondary storage systems”](#) on page 109
- ◆ [“Migrating between two volumes on one secondary storage system”](#) on page 112
- ◆ [“Setting up a tertiary system for a relationship”](#) on page 115
- ◆ [“Reusing a deleted or renamed primary backup root directory name”](#) on page 117
- ◆ [“Reusing a renamed Open Systems SnapVault primary host name”](#) on page 119
- ◆ [“Renaming a SnapVault secondary volume”](#) on page 120
- ◆ [“Resynchronizing restored or broken relationships”](#) on page 122
- ◆ [“Retrying failed transfers”](#) on page 125
- ◆ [“Encrypted File System \(EFS\) file backup and restore”](#) on page 127

## 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 .yyymmdd 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 storage 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 storage system, data transfers between the primary and secondary storage systems cannot continue. If you do not have a way to restore the database, you must initiate a baseline transfer from the primary storage system to the secondary storage 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 storage system to a secondary storage system. A compressed file of the database is created and transferred to the secondary storage system during each data transfer.

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

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

---

**Note**

No checkpoint files are included in the backup. Also softlock state in the database is not backed up.

---

## Characteristics of the database backup file

The database files for an Open Systems SnapVault relationship on the primary storage 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 96 for the steps to restore 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 onward.

---

### 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 storage system to get out-of-sync with the primary storage 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 122.

---

## 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 storage system to a secondary storage system.

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

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

Step	Action	
2	<b>If...</b>	<b>Then...</b>
	You want to back up the history file and its corresponding BLI checksums file	Select “BLI” from the “Enable database backup” drop-down list.  <b>Note</b> _____ 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 CLI of the Open Systems SnapVault primary storage system, navigate to the snapvault/bin directory. <ul style="list-style-type: none"> <li>◆ On Windows systems, the default path is C:\Program Files\netappsnapvault\bin.</li> <li>◆ On UNIX systems, the default path is /usr/snapvault/bin.</li> </ul>

Step	Action
2	<p>Enter the following command:</p> <pre>snapvault restore -S <i>secondary_system:pathname</i> /.OSSV_DATABASE_BACKUP <i>pri_pathname</i></pre> <p><i>secondary_system</i> is the secondary storage system.</p> <p><i>pathname</i> is the path where the .OSSV_DATABASE_BACKUP file is located on the secondary storage system.</p> <p><i>pri_pathname</i> is the path on the primary storage system to which the database will be restored.</p> <hr/> <p><b>Note</b></p> <p>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.</p> <hr/> <p><b>Example:</b> To restore the database from a secondary storage system called f840 to a directory \temp\database on the primary storage system, enter the following command:</p> <pre>D:\Program Files\netapp\snapvault\bin&gt;snapvault restore - S f840:/vol/vol0/rel5/.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.

By default, System State data backup does not include EventLog as it is not a part of Microsoft definition of System State. For more details, see “[Enabling and disabling Windows EventLog](#)” on page 60.

## 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 when, 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.

Taking back up to restore the records and clearing the record helps to maintain the Windows EventLog. If the EventLog is not cleared before it reaches its maximum size, it either stops recording any new events or starts overwriting older events.

---

### Note

When you clear an event log, the operating system does not delete the previous event log file. Instead, Windows creates a new 64-KB log file that replaces the old log file. Before you clear an event log, create a backup of that log.

---

## 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.

Event logs should be backed up separately from any other system files. In a regular system backup, the event log files are copied and therefore are unusable. If you attempt to open an event log file that has been copied or backed up by using any means, other than the Event Log Backup Application Programming Interface, you receive an error message stating that the event log file is corrupt. This error message is the result of a unique characteristic of event log files.

## 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 systems is their interaction with Active Directory and SYSVOL functions on other systems in a domain or 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 the 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

Scenario	SYSVOL authority	Active Directory authority
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 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 102
- ◆ [“Restoring System State data”](#) on page 103
- ◆ [“Using System State data backup to rebuild a primary storage system”](#) on page 105

## Adding System State data backup

---

### Initiating and starting a System State data backup

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

**Note**

---

On Windows 2000 systems, before starting System State data backups, make sure the system has recent Windows Service Packs installed. You can have problems with restored systems if the backed-up system is running an older version of the operating system. For information on which service pack your system is running, see [“Requirements for primary storage systems”](#) on page 16.

---

Step	Action
1	<p>In the storage system console of the SnapVault secondary storage system, enter the following command:</p> <pre>snapvault start -S prim_host:SystemState sec_host:/vol/sec_vol/sec_tree</pre> <p><b>Note</b></p> <p>The keyword <code>SystemState</code> is case-independent.</p> <p><b>Example:</b></p> <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 104.

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 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	<p>On the primary storage system, enter the following command:</p> <pre>snapvault restore -S sec_host:/vol/sec_vol/sec_tree SystemState</pre> <p><b>Note</b></p> <p>The keyword <code>SystemState</code> is case-independent.</p> <p>Ensure that you do not specify a trailing slash ( \ or / ) character at the end of the path name in the preceding command; otherwise, the <code>snapvault restore</code> command will fail.</p> <p><b>Example:</b></p> <pre>snapvault restore -S sv_secondary:/vol/sv_vol/tree_melz SystemState</pre>
3	<p>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.</p> <p>To mark any or all of the restored entries as authoritative, use the Microsoft <code>ntdsutil</code> tool. Otherwise skip to the next step.</p>
4	Reboot the system.

## 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	<p>In the primary storage system, enter the following command:</p> <pre>snapvault restore -S sec_host:/vol/sec_vol/sec_tree SystemStatePrimary</pre> <p><b>Note</b>_____</p> <p>The keyword SystemStatePrimary is case-independent.</p> <p>Make sure that you do not specify a trailing slash ( \ or /) character at the end of the path name in the preceding command; otherwise, the snapvault restore command will fail.</p> <p>_____</p> <p><b>Example:</b></p> <pre>snapvault restore -S sv_secondary:/vol/sv_vol/tree_melz SystemStatePrimary</pre>
3	Reboot the system.

## Using System State data backup to rebuild a primary storage system

---

### About using System State backup

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

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

### Creating a backup to rebuild a primary storage system

To create a backup that you can use to rebuild a primary storage 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	<p>In the storage system console of the SnapVault secondary storage system, enter the following command:</p> <pre data-bbox="494 314 1013 366"><b>snapvault start -S prim_host:SystemState sec_host:/vol/sec_vol/sec_tree</b></pre> <p><b>Note</b></p> <hr/> <p>The keyword <code>SystemState</code> is case-independent.</p> <hr/> <p><b>Example:</b></p> <pre data-bbox="494 545 999 597">snapvault start -S melzhost:SystemState sv_secondary:/vol/sv_vol/tree_melz</pre>

### Rebuilding a primary storage 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 storage system, many of its characteristics must be the same as the original primary storage system from which you perform the system state backup. In case of disk drives, the new primary storage system should either be of the same size or larger than the original. Characteristics that are different on the new primary storage can cause problems, depending on the specifics of the system. One potential difference can be the type of video bus—for example, an AGP bus is different from a PCI video bus.

To rebuild a primary storage system using the backup you created in “[Creating a backup to rebuild a primary storage system](#)” on page 105, complete the following steps.

Step	Action
1	<p>Install and configure the base operating system.</p> <p>Note the following restrictions:</p> <ul style="list-style-type: none"> <li>◆ Use the same operating system and service packs that were on the original system.</li> <li>◆ Use the same machine name as on the original system.</li> <li>◆ Use the same drive letter mappings.</li> <li>◆ Make sure that each drive is at least as large as the corresponding drive was when the backup was made.</li> <li>◆ Format each drive with the same file system type and version as on the original system.</li> <li>◆ Make sure that the hardware configuration is identical to the original.</li> </ul> <p>Do not perform any operating system configuration tasks beyond those needed to satisfy these restrictions.</p>
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	<p>Restore the System State data.</p> <p>Choose between the procedures “<a href="#">Restoring System State data from a backup</a>” on page 103 and “<a href="#">Restoring System State data from a backup and marking it primary</a>” on page 104 based on the points about authoritative versus non-authoritative and primary versus non-primary restores discussed as part of the procedures.</p>
6	Reboot the system.
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 from the secondary storage system and also release the relationship from the primary storage system, to free the primary directory for future backups.

If you use the `snapvault stop` command on the secondary storage system to delete an Open Systems SnapVault relationship and try to re-create the relationship without releasing the relationship on the primary storage system, the attempt fails with an error message similar to the following:

```
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 storage system console, enter the following command: <code>snapvault stop <i>secondary_path</i></code>
2	On the primary storage system console, enter the following command: <code>snapvault release <i>primary_path</i> [ <i>secondary:</i> ] <i>secondary_path</i></code>
3	On the secondary storage system console, enter the following command: <code>snapvault start -S <i>primary_path</i> [ <i>secondary:</i> ] <i>secondary_path</i></code>

# Migrating a relationship between two secondary storage systems

## Migrating data from one secondary storage system to another

Before migrating data from one secondary storage 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 storage system to another secondary storage system without having to perform another baseline transfer, complete the following steps.

Step	Action
1	<p>Ensure that you have Open Systems SnapVault baselines of the directory you are migrating.</p> <p><b>Example:</b> In this procedure, assume a baseline of the bno:C:\500MB directory was backed up to r200-old:/vol/old_vol/bno_C_500 MB.</p>
2	<p>Using SnapMirror, replicate the volume from the present secondary storage system to a volume on the new secondary storage system. For details about creating volume-replicating SnapMirror relationships, see the SnapMirror chapter in the Data ONTAP <i>Data Protection Online Backup and Recovery Guide</i>.</p> <p><b>Example:</b> To replicate the old_vol volume from the r200-old secondary storage system to the new_vol volume on the r200-new secondary storage system, complete the following steps on the new secondary storage system (r200-new):</p> <ol style="list-style-type: none"><li>Create the new_vol volume: <pre>r200-new&gt; vol create new_vol 3</pre></li><li>Mark the new_vol volume as restricted: <pre>r200-new&gt; vol restrict new_vol</pre></li><li>Transfer the old_vol volume to the new_vol volume: <pre>r200-new&gt; snapmirror initialize -S r200-old:old_vol new_vol</pre></li></ol>

Step	Action												
3	<p>Quiesce and break the SnapMirror relationship between the old secondary storage system and the new secondary storage system.</p> <p><b>Example:</b> To quiesce and break the SnapMirror relationship between r200-old and r200-new, run the following commands on r200-new.</p> <ul style="list-style-type: none"> <li>a. <code>snapmirror quiesce new_vol</code></li> <li>b. <code>snapmirror break new_vol</code></li> </ul>												
4	<p>Check the SnapMirror status and SnapVault status on the new secondary storage system. The SnapMirror status should be <code>Broken-off</code>. The SnapVault status should be <code>Snapvaulted</code> on the new volume on the new secondary storage system.</p> <p><b>Example:</b> Perform the following steps from r200-new:</p> <ul style="list-style-type: none"> <li>a. <code>snapmirror status</code></li> </ul> <table border="0" data-bbox="494 777 1142 829"> <thead> <tr> <th>Source</th> <th>Destination</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>r200-old:old_vol</td> <td>r200-new:new_vol</td> <td>Broken-off</td> </tr> </tbody> </table> <ul style="list-style-type: none"> <li>b. <code>snapvault status</code></li> </ul> <table border="0" data-bbox="494 888 1236 972"> <thead> <tr> <th>Source</th> <th>Destination</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>bno:C:\500MB r200-</td> <td>new:/vol/new_vol/bno_C_500MB</td> <td>Snapvaulted</td> </tr> </tbody> </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	<p>Confirm that SnapVault configuration information is not present on the new secondary storage system, by using the <code>snapvault status -c</code> command.</p> <p><b>Example:</b> Perform the following step from r200-new:</p> <pre> snapvault status -c Snapvault secondary is ON.</pre>												

Step	Action
6	<p>Add SnapVault configuration information to the registry on the new secondary storage system using the <code>snapvault start</code> command.</p> <hr/> <p><b>Note</b> This does not start a new baseline, it updates the registry.</p> <hr/> <p><b>Example:</b> Perform the following step from <code>r200-new</code>:  <pre> snapvault start -S bno:C:\500MB r200-new:/vol/new_vol/bno_C_500MB Snapvault configuration for the qtree has been set. Qtree /vol/new_vol/bno_C_500MB is already a replica.</pre> </p>
7	<p>Confirm that SnapVault configuration information is present on the new secondary storage system, by using the <code>snapvault status -c</code> command.</p> <p><b>Example:</b> 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> </p>
8	<p>Test the new SnapVault relationship by manually updating <code>r200-new</code>.</p> <p>If you are using the command line to manage your environment, continue to the next step; otherwise, migration of data is complete between two secondary storage systems.</p> <p><b>Example:</b> 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> </p>
9	<p>Re-create any schedules used on the old secondary storage system to the new secondary storage system, and ensure that access permissions are in place.</p>

# Migrating between two volumes on one secondary storage system

---

## Migrating from one volume to another on a secondary storage system

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

Step	Action
1	<p>Ensure that you have Open Systems SnapVault baselines of the directory you are migrating.</p> <p><b>Example:</b> In this procedure, assume a baseline of the bno:C:\500MB directory was backed up to r200:/vol/old_vol/bno_C_500MB.</p>
2	<p>Using SnapMirror, replicate the volume from the present volume on the secondary storage system to a new volume. For details about creating volume-replicating SnapMirror relationships, see the SnapMirror chapter in the <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i>.</p> <p><b>Example:</b> To replicate the old_vol volume on the r200 secondary storage system to the new_vol volume, complete the following steps on the secondary storage system (r200):</p> <ul style="list-style-type: none"><li>a. Create the new_vol volume. vol create new_vol 3</li><li>b. Mark the new_vol volume as restricted. vol restrict new_vol</li><li>c. Transfer the old_vol volume to the new_vol volume. snapmirror initialize -S r200:old_vol new_vol</li></ul>

Step	Action												
3	<p>Quiesce and break the SnapMirror relationship between the old volume and the new volume.</p> <p><b>Example:</b> To quiesce and break the SnapMirror relationship between old_vol and new_vol, run the following commands on r200:</p> <ul style="list-style-type: none"> <li>a. snapmirror quiesce new_vol</li> <li>b. snapmirror break new_vol</li> </ul>												
4	<p>Check the SnapMirror status and SnapVault status of the new volume. The SnapMirror status should be Broken-off. The SnapVault status should be Snapvaulted to the new volume.</p> <p><b>Example:</b> Perform the following steps from r200.</p> <ul style="list-style-type: none"> <li>a. snapmirror status</li> </ul> <table border="0" data-bbox="494 743 1142 795"> <thead> <tr> <th>Source</th> <th>Destination</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>r200:old_vol</td> <td>r200:new_vol</td> <td>Broken-off</td> </tr> </tbody> </table> <ul style="list-style-type: none"> <li>b. snapvault status</li> </ul> <table border="0" data-bbox="494 855 1233 907"> <thead> <tr> <th>Source</th> <th>Destination</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>bno:C:\500MB</td> <td>r200:/vol/new_vol/bno_C_500MB</td> <td>Snapvaulted</td> </tr> </tbody> </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	<p>Confirm that the SnapVault configuration information is not present for the new volume using the <b>snapvault status -c</b> command.</p> <p><b>Example:</b> Perform the following step from r200.</p> <pre>snapvault status -c</pre> <p>Snapvault secondary is ON.</p>												

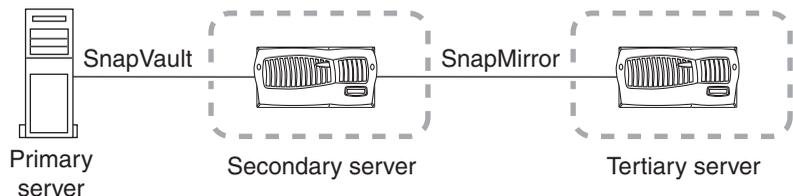
Step	Action
6	<p>Add SnapVault configuration information to the registry on the new volume using the <b>snapvault start</b> command.</p> <hr/> <p><b>Note</b> This does not start a new baseline; it updates the registry.</p> <hr/> <p><b>Example:</b> Perform the following step from r200.</p> <pre> snapvault 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.</pre>
7	<p>Confirm that SnapVault configuration information is present on the new volume using the <b>snapvault status -c</b> command.</p> <p><b>Example:</b> Perform the following step from r200.</p> <pre> snapvault status -c  Snapvault secondary is ON.  /vol/new_vol/bno_C_500MB source=bno:C:\5000MB</pre>
8	<p>Test the new SnapVault relationship by manually updating new_vol.</p> <p>If you are using the command line to manage your environment, continue to the next step; otherwise, migration between two volumes on one secondary storage system is complete.</p> <p><b>Example:</b> Perform the following step from r200.</p> <pre> snapvault update r200:/vol/new_vol/bno_C_500MB  Transfer started.  Monitor progress with 'snapvault status' or the snapmirror log.</pre>
9	<p>Re-create any schedules used on the old volume to the new volume and ensure that access permissions are in place.</p>

## Setting up a tertiary system for a relationship

---

### Need for a tertiary system

You can protect the SnapVault secondary storage system from disasters by using the SnapMirror feature. The configuration involves setting up SnapMirror relationships from the volumes on your SnapVault secondary storage system to volumes on a remote (tertiary) Data ONTAP system, as shown in the following illustration. SnapMirror therefore 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 storage system and the tertiary system, without any initial baseline transfers. A softlock is a request to the primary storage 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 storage system becomes unusable because of a disaster, you can manually redirect the subsequent SnapVault transfers to the tertiary system instead of the old SnapVault secondary storage system. Effectively, the tertiary system becomes the new SnapVault secondary storage system, and the SnapVault transfers continue using the most recent Snapshot copy common to both the primary and tertiary storage systems.

## Configuration of secondary storage system using a tertiary system

After the secondary storage system comes up, resynchronize it using the tertiary system. Then, you can configure the secondary storage 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 storage systems from the secondary storage system. The original tertiary storage system then takes the place of the secondary storage system, and the secondary storage system takes place of the original tertiary system.

- ◆ primary -> secondary -> tertiary

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

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

## Reusing a deleted or renamed primary backup root directory name

---

### When to reuse a directory name

If you delete 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
```

To reuse the existing relationship, complete the following steps.

Step	Action
1	Enable the <code>FixRootInodeChanges</code> flag in <code>snapvault.cfg</code> : <code>[QSM:FixRootInodeChanges]</code> <code>Value=TRUE</code>
2	Run the <code>snapvault update</code> command and the update completes successfully: <code>snapvault update destination_filer:path</code>

### Reusing a deleted or renamed directory name

If the SnapVault source directory on the primary storage system was intentionally deleted, and backups are not necessary, complete the following steps to reuse the deleted directory's name.

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

If the SnapVault source directory on the primary storage system was intentionally renamed, complete the following steps to reuse the renamed directory's name.

<b>Step</b>	<b>Action</b>
<b>1</b>	Create a new SnapVault relationship for the renamed directory and perform a baseline transfer for the new relationship.
<b>2</b>	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 storage system.

If the SnapVault source directory on the primary storage system was erroneously renamed, complete the following steps.

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

# Reusing a renamed Open Systems SnapVault primary host name

---

## Procedure to update relationship on secondary storage system

If the Open Systems SnapVault primary storage system is renamed, you can update the SnapVault relationship without performing a new baseline transfer. Complete the following steps.

Step	Action
1	Update the SnapVault relationship on the SnapVault secondary storage system to reflect the new primary host name, by using the following command on the secondary storage system: <b>snapvault modify -S new_source_hostname:path destination_filer:path</b>
2	Run the <code>snapvault update</code> command and the new host name is displayed in the SnapVault status: <b>snapvault update destination_filer:path</b>

## Renaming a SnapVault secondary volume

---

Generally, Open Systems SnapVault updates do not work after a volume is renamed. However, you can rename a secondary volume without a new baseline transfer.

To rename a secondary volume, complete the following steps.

Step	Action
1	Enter the following command:  <b>vol rename oldvolname newvolname</b>
2	Enter the following commands to verify the changes:  <b>snapvault status</b>  snapvault status displays the new path.  <b>snapvault status -c</b>  snapvault status -c does not display the new path.
3	Enter the following command:  <b>snapvault start -S primary_filer:primary_qtree secondary_filer:secondary_qtree</b>  A message similar to the following appears: Snapvault configuration for the qtree has been set. Qtree /vol/newvolname/secondary_qtree is already a replica.
4	Enter the following command:  <b>snapvault status -c</b>  snapvault status -c now displays the new path.
5	Enter the following command to verify whether the change was successful:  <b>snapvault update secondary_qtree</b>

Step	Action
6	<p>The output of <code>snapvault status -c</code> also contains entries that refer to the old volume name in addition to the new volume name.</p> <p>Enter the following command to remove these entries:</p> <pre><b>snapvault stop /vol/oldvolname/secondary_qtree</b></pre> <p>A message similar to the following appears:</p> <pre>Snapvault configuration for the qtree has been deleted. Could not delete qtree: destination qtree does not exist</pre> <p>The output reflects that the configuration information is deleted and the qtree does not exist on the disk because the volume name is changed.</p>

# 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 storage systems. A common Snapshot copy is necessary for incremental backups to continue successfully between the primary and secondary storage 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 creation of one of its SnapVault Snapshot copies.

In this condition, if you want to restore the data from the qtree to another location on the primary storage system, and then perform subsequent incremental backups from the restored location to the same qtree on the secondary storage system, you must resynchronize the relationship.

- ◆ An older version of the Open Systems SnapVault database is restored on the primary storage 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 storage systems so that the secondary qtree becomes read-only and the incremental transfers can continue.

- ◆ Restoring the secondary qtree to different primary storage system  
If a primary storage system having the Open Systems SnapVault relations, gets crashed or corrupted, you can restore the secondary qtree to a different primary storage system. Further, proceed with resynchronizing this relationship.

### Attention

---

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 storage systems was by re-initializing the relationship. Re-initializing a relationship involves a lengthy baseline transfer between the primary and the secondary storage 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.

## Resync after restore

In releases prior to Open Systems SnapVault 2.2, this feature was not available. Resync after restore or break allows you to resynchronize a relationship without requiring a new baseline transfer.

Before restoring, enable the *Enable restart/resync on restore* check box, in the Configurator GUI, to make the resync work.

## Resynchronizing a relationship for different Data ONTAP versions

Resynchronizing a relationship with different Data ONTAP versions is based on the value of the stanza `[QSM: Resync version]` in the `snapvault.cfg` file.

To support Data ONTAP 7.2 or later the value is 11. It is the default value.

To support Data ONTAP 7.0.5 and 7.1.2, the value has to be changed as follows:

```
[QSM: Resync version]
```

```
Value=6
```

To resynchronize a SnapVault relationship, perform the following steps.

Step	Action
1	Log in to the secondary storage system.

Step	Action
2	<p data-bbox="494 239 1189 300">On the console, enter the following command to resynchronize a SnapVault relationship:</p> <pre data-bbox="494 322 1005 374">snapvault start -r -S prim_host:dirpath /vol/sec_vol/sec_tree</pre> <p data-bbox="494 401 780 427"><b>Example 1 (Windows):</b></p> <pre data-bbox="494 435 1026 487">snapvault start -r -S melzhost:C:\melzdir /vol/sv_vol/tree_melz</pre> <p data-bbox="494 513 727 539"><b>Example 2 (UNIX):</b></p> <pre data-bbox="494 548 1053 600">snapvault start -r -S melzhost:/usr/melzdir /vol/sv_vol/tree_melz</pre>

## Retrying failed transfers

---

### About retrying failed transfers

If a transfer stops because of an error, such as a temporary network outage, the transfer is automatically retried after a 60-second wait. The number of retries is determined by the value of the `tries` option for the relationship. The default `tries` value is 2, 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 number of retry attempts made for failed transfers, complete the following step on the secondary storage system.

Step	Action
1	<p>Enter the following command:</p> <pre>snapvault modify -t n sec_qtree</pre> <p><i>n</i> is the number that specifies the number of retries.</p> <hr/> <p><b>Note</b></p> <p>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.</p> <hr/> <p><i>sec_qtree</i> is the <code>qtree</code> on the secondary storage system where data is being backed up.</p>

### Configuring the checkpoint interval

In Open Systems SnapVault 2.6, you can configure the checkpoint interval by executing the following command:

```
[QSM:Checkpoint Interval]  
Value = 300
```

The default value is 300 seconds (5 minutes).

**Note**

---

60 seconds is the minimum checkpoint interval and any value less than 60 seconds is considered as 60 seconds.

---

# 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 ID 139696 on the NOW site by entering the bug number 139696 in the Bugs Online > Quick Search > Enter Bug ID(s) field or the following URL in your browser:  
<http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=139696>

---

### Note

If the version of Data ONTAP on your secondary storage 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 will 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 the Open Systems SnapVault space estimator is

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

## 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, ACLs, and data streams on a system. The space estimator scans the backup path of a system to obtain the values for some factors, such as number of files. 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 results are written to a log file called `snapvault`, in the `install_dir/etc` directory.

Even if space is insufficient, 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 138.

- ◆ 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 131. Run the `svestimator` executable available on the Open Systems SnapVault 2.6 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.

**Example:**

The following is an example of a command and the 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.

**Example:**

The following is an example of a command and the 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 [“Example”](#) on page 132.
- ◆ The path and file exclusion list files  
For more information, see [“Configuring backup exclusion lists”](#) on page 68.

On a system on which Open Systems SnapVault is installed, the `estimator.cfg` file is present in the `snapvault` or `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 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 (an example follows) and the two backup exclusion list files (mentioned previously in this section) in the directory in which the stand-alone space estimator is run.

**Example:**

The following is an example of an 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	<ul style="list-style-type: none"> <li>◆ 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.</li> <li>◆ During the stand-alone operation, this field enables the space estimator utility to include checksum sizes in its calculations.</li> </ul>
Hist Data enabled	FALSE	<ul style="list-style-type: none"> <li>◆ 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.</li> <li>◆ 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.</li> </ul>
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 the 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 an example 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 92).

The following is an example of the information that will be written to a snapvault log file for the example discussed in the above 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	<b>If...</b>	<b>Then...</b>
	The Open Systems SnapVault software is installed on your primary storage system	<p>Go to the <i>install_dir/bin</i> directory and locate the file called <i>svestimator</i>.</p> <p>You can run the stand-alone space estimator from this directory or move the file to another location of your choice.</p>
	The Open Systems SnapVault software is not installed on your primary storage system	<ol style="list-style-type: none"> <li><b>a.</b> Follow the instructions in Chapter 2, “<a href="#">Installing the Open Systems SnapVault Software</a>,” on page 15 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 the Open Systems SnapVault agent.</li> <li><b>b.</b> Uncompress the installation package, if needed.</li> <li><b>c.</b> Locate the <i>svestimator</i> file in the <i>installfiles</i> directory.</li> </ol>

## 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	<p>Enter the following command at the command prompt of your primary storage system where you installed the space estimator:</p> <pre><b>svestimator [-o -i -d] root_backup_path</b></pre> <p><code>-o</code> 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 <code>-o</code> option does not take into consideration the disk space consumed by the Open Systems SnapVault installation.</p> <p><code>-i</code> 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 storage system.</p> <p><code>-d</code> 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.</p> <hr/> <p><b>Note</b></p> <p>You must enable the “Generate debugging files” option in the General tab of the Configurator utility before using the <code>-d</code> 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.</p> <hr/> <p><code>root_backup_path</code> 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.</p>

## 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/config</i> directory, set the value of the following entry in the <i>snapvault.cfg</i> file to True:  [Configuration:Estimator can fail backup] Value=True

## About this chapter

This chapter describes the use of Open Systems SnapVault as part of a virtual environment based on VMware ESX Server. It describes the ESX Server architecture and file storage, Open Systems SnapVault 2.6 backup and restore of virtual machines, and installation and configuration procedures.

## Topics in this chapter

This chapter contains the following topics:

- ◆ [“VMware terminology”](#) on page 140
- ◆ [“Overview of Virtualization and VMware ESX Server”](#) on page 142
- ◆ [“Configuration of Open Systems SnapVault for ESX Server”](#) on page 146
- ◆ [“Backup and restore of virtual machines”](#) on page 153
- ◆ [“Open Systems SnapVault support for VMotion”](#) on page 158

# VMware terminology

---

The following are definitions of VMware specific terms you will come across in this chapter. For more information, see <http://vmware.com/>.

<b>Virtual machine</b>	A virtual machine is a tightly isolated software container that can run its own operating system and applications as if it were a physical computer. A virtual machine behaves like a physical computer and contains its own CPU, RAM hard disk, and network interface card (NIC).
<b>ESX Server</b>	VMware ESX Server is a virtualization software that abstracts processor, memory, storage, and networking resources into multiple virtual machines.
<b>VirtualCenter</b>	VMware VirtualCenter is a management software to manage a group of ESX hosts and the associated virtual machines. VirtualCenter server is the backend and virtual infrastructure client is a user interface.
<b>VMotion</b>	VMotion is a feature that enables you to move running virtual machines from one ESX Server to another without interrupting service. VMotion is activated by the VirtualCenter Server. The VirtualCenter Server centrally coordinates all VMotion activities.
<b>Universal Unique Identifier (UUID)</b>	<p>The UUID is a 128-bit value used for unique identification of virtual machines.</p> <p><b>Example:</b></p> <pre>00 11 22 33 44 55 66 77-88 99 aa bb cc dd ee ff</pre>
<b>Virtual Machine File System (VMFS)</b>	VMFS is a file system used by ESX Servers to store virtual machine files. Each virtual machine is represented as a collection of files under the VMFS volume.
<b>Virtual Machine Disk (VMDK)</b>	VMDK file is a file representation for the hard disk of the virtual machine.

**Virtual Machine Extensions (VMX)**

A VMX file is the primary configuration file for a virtual machine.

**Service console**

The service console is a RHEL 3.0 operating system used as a management interface to the ESX Server.

# Overview of Virtualization and VMware ESX Server

---

## Virtualization overview

According to VMware, “Virtualization is technology that enables multiple operating systems and multiple applications to run on a single computer simultaneously. Essentially, it allows one computer to do the job of many, thus greatly increasing the usefulness and flexibility of your hardware.” For more information, see <http://vmware.com/>.

## Overview of ESX Server

ESX Server is a *bare metal* hypervisor that partitions physical servers into multiple virtual machines. Each virtual machine represents a complete system, with processors, memory, networking, storage, and Basic Input Output System (BIOS) code.

Multiple virtual machines can share physical resources and run concurrently on the same server.

Operating systems and applications can run unmodified in virtual machines.

# Open Systems SnapVault on ESX Server

---

## Advantages of running Open Systems SnapVault 2.6 in the service console of ESX Server

Prior to Open Systems SnapVault 2.6, Open Systems SnapVault supported backup within individual virtual machines, with the following limitations:

- ◆ Managing backup agents inside each virtual machine was time-consuming.
- ◆ The data within a virtual machine was backed up, and not the whole virtual machine.

To address the limitations of running Open Systems SnapVault within each virtual machines, Open Systems SnapVault 2.6 runs in the service console of an ESX Server. This offers you the following advantages:

- ◆ Each virtual machine is visible to Open Systems SnapVault as a set of files. Therefore, a single Open Systems SnapVault agent can be used to back up and restore multiple virtual machines.
- ◆ Disaster recovery is possible as the entire virtual machine is backed up.
- ◆ Because Open Systems SnapVault supports update transfers, you can maintain multiple Snapshot copies on the secondary storage system, which enables you to restore a virtual machine to any of its previous states.

## Limitations

Open Systems SnapVault *does not support* the following:

- ◆ If a virtual machine has physical Raw Device Mapping (RDM), then Open Systems SnapVault cannot back up the virtual machine.
- ◆ If a virtual machine has virtual Raw Device Mapping (RDM), the disk is excluded from backup.
- ◆ Open Systems SnapVault backup fails for virtual machine having existing Snapshot copy.
- ◆ Checkpoint restart transfers after the system reboots or SnapVault service restarts.
- ◆ Initiating baseline transfer of a virtual machine that is registered with an ESX Server other than the ESX Server of the Open Systems SnapVault primary system.
- ◆ Resync after restore
- ◆ File system backup
- ◆ Update transfers after changing the UUID of the virtual machine.

**Note**

---

Ensure that the following TCP ports are open before Open System SnapVault is installed:

NDMP port (default value is 10000)

FILESERVER port-10555.

QSMSEVER port-10566.

---

# Installing Open Systems SnapVault 2.6 on ESX Server

---

## Installing Open Systems SnapVault in the service console of ESX Server

The procedure to install Open Systems SnapVault 2.6 in the service console of ESX Server is similar to the procedure to install Open Systems SnapVault 2.6 on the Linux platform.

However, when you run the **install** script, you will have to make a few additional inputs, such as the following:

Enter the Host Name or IP address of the Virtual Center Host [localhost] :

You can enter the host name or the IP address of the VirtualCenter host if ESX Server is not running independently.

Enter the User Name to connect to the Virtual Center Host :  
Please enter the password to connect to the Virtual Center Host :  
Confirm password:

Enter the user name and password for the VirtualCenter host.

Should HTTPS be used to connect to the Virtual Center Host?  
If you specify n, HTTP will be used (y n) [y] :

HTTPS is the default value.

For more information about installing Open Systems SnapVault on the Linux platform, see [“Installing the HP-UX, AIX, or Linux agent from the CD-ROM”](#) on page 30.

# Configuration of Open Systems SnapVault for ESX Server

---

## Command-line interface utility for Open Systems SnapVault on ESX Server

The Configurator utility GUI (svconfigurator) cannot run in the service console of ESX Server because of the non availability of the x-libraries. Therefore, a new command-line interface utility, svconfig, has been introduced.

The command-line interface provides a way to configure all the options that are available through the svconfigurator GUI. Additionally, the VMware options incorporated for ESX Server are also available through the command-line interface.

Use the svconfig utility as follows:

```
svconfig [option_key [option_value]]
```

```
svconfig [option_key_prefix]
```

Option	Description
<i>option_key</i>	Key name of the option that needs to be configured
<i>option_value</i>	Value of the option key

If...	Then...
<i>option_value</i> is missing	The value of <i>option_key</i> is displayed along with the list of all possible values for this option.
<i>option_key_prefix</i> is specified	All matching options with their values are displayed. In this case, the list of possible values is not displayed.
<i>option_key_prefix</i> is specified and it matches an option exactly	The option, its value, and its possible values are displayed.

If...	Then...
You specify <i>option_value</i>	That option value is verified before it is set. If it is not verified, the configuration is not modified.

The options output is in the following format:

```
option_key tab_character option_value
tab_character Possible Values:value 1, value 2, value 3 ...
```

### Example 1:

This following command prints all the options that start with VMware.

```
# svconfig vmware
vmware.vchost                localhost
vmware.username              root
vmware.password
vmware.https.enable          true
vmware.poweroff_before_ss    false
vmware.backup_powered_down_vm true
```

### Example 2:

The following command gives the value of the `vmware.https.enable` option. Because it is also the only option, it also lists the possible values.

```
# svconfig vmware.https.enable
vmware.https.enable          true
Possible values : true, false
```

### Example 3:

This command sets the preceding option to false.

```
# svconfig vmware.https.enable false
Option vmware.https.enable set to false
```

## Supported options:

The following table specifies all the supported options.

Stanza name /default value	Value type/range	Options key	Notes
Directories:Database	Path	path.db	Database path
Directories:Trace	Path	path.trace	Trace path
Directories:Tmp	Path	path.tmp	Temporary objects path
Trace:Trace to File	Boolean	trace.enable	Enable debug trace
Trace:Lines per File	Range (0, 32000) * 1000	trace.lines_per_file	Number of trace lines per file. Zero means do not split files.
Trace:Files to Keep	Range (0, 1000)	trace.files_to_keep	Number of trace files to keep. Zero means keep all files.
Process Manager:Trace Level	String [ALWAYS NORM AL VERBOSE  LIBNORMAL LI BVERBOSE]	trace.level.proc_mgr	Trace level of Process Manager
Communication Manager:Trace Level	String [ALWAYS NORM AL VERBOSE  LIBNORMAL LI BVERBOSE]	trace.level.comm_mgr	Trace level of Communications Manager
SnapVault Listener:Trace Level	String [ALWAYS NORM AL VERBOSE  LIBNORMAL LI BVERBOSE]	trace.level.svlistener	Trace level of svlistener

<b>Stanza name /default value</b>	<b>Value type/range</b>	<b>Options key</b>	<b>Notes</b>
NDMP Server:Trace Level	String [ALWAYS NORMAL VERBOSE]  LIBNORMAL LIBVERBOSE]	trace.level.ndmp_server	Trace level of NDMP server
QSM Server:Trace Level	String [ALWAYS NORMAL VERBOSE]  LIBNORMAL LIBVERBOSE]	trace.level.qsm_server	Trace level of Qtree SnapMirror server
NDMP:Listen Port	Range (0,65535)	ndmp.port	NDMP listen port
NDMP:Account	String:256	ndmp.account	NDMP account name
NDMP:Password	String:256	ndmp.password	The NDMP password in encoded form
NDMP:Host Id	String:256	ndmp.hostid	NDMP host ID
NDMP:Host Name	String:256	ndmp.hostname	NDMP host name
QSM:Check Access List	Boolean	qsm.accesslist.enable	Enables or disables host access list (white list)
QSM:Access List	String:2048	qsm.accesslist.hosts	Qtree SnapMirror access list of comma separated hosts
QSM:Generate Verify Checksums	Boolean	config.cf.enable	
Configuration:Check sums	String [HIGH LOW OFF]	config.bli	BLI level
Configuration:Run Estimator	Boolean	config.estimator.enable	Run estimator before each backup

<b>Stanza name /default value</b>	<b>Value type/range</b>	<b>Options key</b>	<b>Notes</b>
QSM:Enable Restart	Boolean	config.resync.enable	Enable restart or resync on restore
QSM:Backup Database	String [BLI DB only NONE]	config.db_backup	Enable database backup
VMware:VCHost	String:2048	vmware.vchost	Host name or IP address of the VCHost. If VirtualCenter is not available, this can also be specified as localhost.
VMware:Username	String:256	vmware.username	User name used to connect to the VCHost.
VMware>Password	String:256	vmware.password	Password to connect to VCHost.
VMware:HTTPS	Boolean	vmware.https.enable	Mode of connection to VCHost. The two possible values are <i>Yes</i> or <i>No</i> . If the value is <i>No</i> , HTTP is used.  The default value is HTTPS.
VmWare:HTTP_PORT	Range (0,65535)	vmware.port.http	The HTTP port on which VirtualCenter Server is listening. The default value is 80.
VmWare:HTTPS_PORT	Range (0,65535)	vmware.port.https	The HTTPS port on which VirtualCenter Server is listening. The default value is 443.

<b>Stanza name /default value</b>	<b>Value type/range</b>	<b>Options key</b>	<b>Notes</b>
VMware:PowerOff Before SS	Boolean	vmware.poweroff_before_ ss	<p>The two possible values are TRUE or FALSE.</p> <p>If there are multiple disks in different modes, VMware recommends that the virtual machine be powered down before taking a Snapshot copy. Hence, if this value is TRUE, Open Systems SnapVault turns off the virtual machine, takes a Snapshot copy and turns on the virtual machine again to continue the transfer.</p> <p>If the value is FALSE, the virtual machine (with multiple disks in different mode) is not backed up.</p> <p>The default value is FALSE.</p>

<b>Stanza name /default value</b>	<b>Value type/range</b>	<b>Options key</b>	<b>Notes</b>
VMware:Backup powered down VM	Boolean	vmware.backup_powered_ down_vm	<p>The two possible values are TRUE or FALSE.</p> <p>The default value is FALSE and hence powered down virtual machines will not be backed up by default. When the backup of a virtual machine is in progress, do not modify or change the configuration settings for the virtual machine (like adding or deleting disks).</p> <p>You can make the configuration changes by powering off the virtual machine and Open Systems SnapVault will skip backing up that virtual machine.</p> <p>If this value is set to TRUE, then powered off virtual machines are backed up.</p>

# Backup and restore of virtual machines

---

This section describes the backup and restore of virtual machines using the Open Systems SnapVault agent.

## Backing up a virtual machine from the primary storage system

Each virtual machine is represented as a collection of files under the VMFS volume. To backup a virtual machine, Open Systems SnapVault backs up the following files:

- ◆ **vmx:** Virtual machine definition file which contains the references to all the components of the virtual machine (CPU, memory disk, and so on)
- ◆ **.vmdk:** The disk descriptor and disk data files.
- ◆ **.nvram:** File containing the BIOS configuration of the virtual machine.
- ◆ **.log:** log files contain the activities of the virtual machine.

When the virtual machine is running, its vmdk files are locked and not accessible for external applications; however, taking virtual machine Snapshot would freeze and release the lock on the vmdk files allowing Open Systems SnapVault to backup these files. As long as the virtual machine Snapshot copy is active, changes to the virtual machine are tracked in a delta file.

Open Systems SnapVault deletes the virtual machine Snapshot copy once the transfer is completed.

Open Systems SnapVault backups and restores work integrated with the VMware environment and uses VMware infrastructure SDK to communicate with VMware environment for activities like creating Snapshot copies, registering, and powering on the virtual machine.

**Directory structure of the secondary storage system:** After a virtual machine is backed up, different files that comprise the virtual machine are organized according to the following directory structure in the secondary qtree.

Directory	Content
CONFIG	Holds the VMX file and NVRAM file
LOGS	Holds all log files

Directory	Content
DISK_x_y	Holds disk files for the hard disk at the virtual device node SCSI (x:y)
OSSV__APP__ CONFIG__	Open Systems SnapVault configuration file for a virtual machine

In addition, OSSV\_DATABASE\_BACKUP is backed up at the root of the secondary qtree.

## Creating an initial baseline copy

To initiate a backup of the virtual machine from the primary storage system, run the following command:

```
snapvault start -S esx-server:app:vmware:uuid secondary qtree
```

*esx-server* is the source ESX Server.

*app:vmware* is a keyword to specify that the backup request is for a virtual machine while Open Systems SnapVault is installed in an ESX Server.

*uuid* is the universal unique identifier of 128-bit value used for identification of the individual virtual machines.

Example:

```
00 11 22 33 44 55 66 77-88 99 aa bb cc dd ee ff
```

*secondary qtree* provides the path for the secondary qtree where the data files are to be backed up.

### Example to backup a virtual machine:

```
snapvault start -S esx1:app:vmware:503f7bac-c758-3401-5613-8482ed7f3451 /vol/vol0/vm1
```

*esx1* is the host name of the ESX Server.

*503f7bac-c758-3401-5613-8482ed7f3451* is the uuid of the virtual machine to be backed up.

*/vol/vol0/vm1* is the secondary qtree to backup data.

To find the UUID of a particular virtual machine, run the following command on the ESX Server console:

```
vcbVmName -h esx host or VC server -u username -p password -s Any
```

-h *esx host or VC Server* is the host name or IP address of the ESX host or the VirtualCenter server.

-u *user name* is the user name of the ESX host or the VirtualCenter server.

-p *password* is the password of the ESX host or the VirtualCenter server.

Any lists the details of all virtual machines.

**Example:** To get the UUIDs of the virtual machines on the local host ESX Server, run the following command:

```
vcbVmName -h localhost -u root -p ossvbtc211 -s Any
```

For more information about the backup process, see “[Perform SnapVault backup on Open Systems platforms](#)” on page 78.

## Methods to restore a virtual machine

There are three methods to perform the restore operation of a virtual machine. You can either restore only virtual machine data (only vmdk files), or restore virtual machine data along with all the backed up virtual machine configuration (vmdk along with vmx, nvrAm and log files)

- ◆ To restore only the virtual machine data, enter the following command:

```
snapvault restore -S Secondary host:Qtrees  
app:vmware:UUID:cfg=current
```

With the above syntax, only vmdk files are restored. The configuration files (vmx, nvrAm and log files) will not be restored and the virtual machine continues to use existing configuration. Before doing the restore, ensure that the virtual machine already exists and registered.

After restore is complete, the virtual machine will be powered on automatically.

### Example to restore only the virtual machine data:

```
snapvault restore -S filer1:/vol/vol0/vm1 app:vmware:503f7bac-  
c758-3401-5613-8482ed7f3451:cfg=current
```

Filer1 is the secondary host name.

/vol/vol0/vm1 is the secondary Qtrees from where the data is restored.

app:vmware:503f7bac-c758-3401-5613-8482ed7f3451 is the virtual machine to be restored.

cfg=current means the restore is based on current configuration on the ESX Server. Only disk files are restored, configuration files are not restored.

- ◆ To restore the virtual machine data along with all the backed up configuration, enter the following command:

```
snapvault restore -S Secondary host:Qtree  
app:vmware:UUID:cfg=original
```

With the above syntax, all the files that were backed up are restored. Hence this will overwrite the current configuration settings of the virtual machine if it is already registered.

### Example to restore virtual machine data and configuration:

```
snapvault restore -S filer1:/vol/vol0/vm1 app:vmware:503f7bac-  
c758-3401-5613-8482ed7f3451:cfg=original
```

Filer1 is the secondary host name.

/vol/vol0/vm1 is the secondary qtree from where the data is restored.

app:vmware:503f7bac-c758-3401-5613-8482ed7f3451 is the virtual machine to be restored.

cfg=original means restore is based on the original backed up configuration. All the files backed up are restored. Hence this will overwrite the current configuration settings if exists.

### Note

---

Select the configuration options depending on whether you want to restore only the virtual machine data (only vmdk files) or virtual machine data along with the configuration (vmx, nvram and log files). The default is to restore the virtual machine data along with the backed up configuration (*cfg=original*)

If the virtual machine is destroyed, you must restore the virtual machine using *cfg=original*. A restore using *cfg=original* will restore the virtual machine whether it exists or not. If the virtual machine exists, it will be overwritten.

---

- ◆ Another type of restore is similar to the normal Open Systems SnapVault qtree restore. Enter the following command:

```
snapvault restore -S Secondary system:Qtree data-store path on  
the primary
```

When you run this command, virtual machine files are restored to the primary path specified in the `restore` command with no additional functionality of integration with VMware environment.

### Example:

```
snapvault restore -S filer1:/vol/vol0/vm1 /vmfs/nfs/vm1
```

Filer1 is the secondary host name.

/vol/vol0/vm1 is the secondary qtree from where the data is restored.

/vmfs/nfs/vm1 is the location on primary ESX Server where data needs to be restored.

---

**Note**

It is important to restore the virtual machine to one of the available data stores. If you restore the virtual machines to any of the local paths (non-VMFS) on ESX Server, these paths cannot be used as virtual machine repositories. For both backup and restore operations, you need to know the UUID for the virtual machine.

---

## Open Systems SnapVault support for VMotion

---

VMotion is the process of live migration of virtual machines from one ESX host to another.

Open Systems SnapVault supports the backup of virtual machine during and after VMotion. Overlapping VMotion does not affect the backup, and block-level incremental (BLI) based updates continue to work even after the virtual machine is moved to another ESX host by VMotion.

The prerequisite for VMotion support is that Open Systems SnapVault should be installed and running on both the source and target ESX hosts that are involved in VMotion and port 10555 should be open on the target ESX host.

## What the OSSVINFO tool does

OSSVINFO is a data collection tool that collects Open Systems SnapVault-related information from primary and secondary storage systems. It writes this data to a text file in a specific format to the output directory. Also, it collects the ChangeLog and trace files to this output directory if either the `-q` (for Windows only) or `-all` (for all platforms) option is given.

## OSSVINFO-supported platforms

There are two versions of OSSVINFO that are available: an executable file for Windows, and a Perl script and an executable file for UNIX.

- ◆ *OSSVINFO.exe* runs on Windows 2000 and Windows 2003, on which Open Systems SnapVault is installed.
- ◆ *OSSVINFO.pl* runs on Solaris, Linux, HP-UX, and AIX, on which Open Systems SnapVault is installed.

## How the OSSVINFO tool works

**For Windows:** The following are the list of OSSVINFO commands for Windows and their description.

The following command displays the list of relationships in the primary storage system:

```
OSSVINFO.exe -list
```

The following command retrieves the trace files and ChangeLog files to the output directory in addition to the other information that is collected in this directory:

```
OSSVINFO.exe [ -s secondary ] [ -l username:password ] [-q qtreeid]  
[-all] Output_Dir
```

---

### Note

In Windows, `-q` and `-all` commands stop the Open Systems SnapVault service and collect the trace files and ChangeLog files. After the files are collected, Open Systems SnapVault service is restarted automatically.

---

`-s secondary`—name of the secondary storage system

`-l username:password`—user name and password of the secondary storage system

`-q qtreeid`—to retrieve the trace files and the ChangeLog files of a particular qtree

`-all`—to retrieve the trace files and ChangeLog files of all the qtrees in the primary storage system

*Output\_Dir*—name of the directory where the output is stored

**Example:**

```
OSSVINFO.exe -s sv_secondary -l username:password -q qtreeid] -all  
Output_Dir
```

The following command displays the version of the installed OSSVINFO tool. OSSVINFO 2.0 is packaged with Open Systems SnapVault 2.6. To check the version number, run the following command:

```
INSTALL_DIR\bin\OSSVINFO.exe -version
```

**Note**

---

OSSVINFO.exe can be run only on Windows systems that have Open Systems SnapVault agent installed.

---

**For UNIX:** The following are the list of OSSVINFO commands for UNIX and their description.

The following command retrieves all the trace files and ChangeLog files to the output directory in addition to the other information that is collected in this directory:

```
OSSVINFO.pl [ -s secondary ] [ -l username:password ] [-all]  
Output_Dir
```

`-s secondary` is the name of the secondary storage system.

`-l username:password` is the user name and password of the secondary storage system.

`-all` retrieves the trace files and ChangeLog files of all the qtrees in the primary storage system.

*Output\_Dir* is the name of the directory where the output is stored.

**Example:**

```
OSSVINFO.pl [ -s secondary ] [ -l username:password ] [-all]  
Output_Dir
```

The following command displays the version of the installed OSSVINFO tool:

```
- OSSVINFO.pl -version
```

## List of error messages

The following table lists the frequently encountered Open Systems SnapVault error messages, their causes, and their solutions.

### Note

The solutions provided in the following table assume that you are running Open Systems SnapVault 2.6 on the primary storage system.

The following error messages are displayed on primary systems running Open Systems SnapVault.

Error number	Error message	Cause	Solution
	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.	<ul style="list-style-type: none"> <li>◆ Stop the process using port 10000 OR</li> <li>◆ 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 <a href="#">“Modifying the NDMP Listen Port setting”</a> on page 60.</li> </ul>
	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.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	OFM volume cannot be synchronized - volume error. (1)	<p>This occurs when OFM is not able to create a Snapshot copy of the drive that holds the data for the relationship.</p> <p>(OFM requires a period of file system inactivity to create a Snapshot copy.)</p>	<p>Edit the following settings in the Open Systems SnapVault Configurator (svconfigurator) and retry the transfer:</p> <ul style="list-style-type: none"> <li>◆ Set the “Write inactivity period (seconds)” from the default of 5 to 2.</li> <li>◆ Set “Synchronization timeout (seconds)” from the default of 60 to 120.</li> </ul>
	Cannot back up 'c:\' Failed to create a volume snapshot	<p>This occurs when OFM is not able to create a Snapshot copy of the drive that holds the data for the relationship.</p> <p>(OFM requires a period of file system inactivity to create a Snapshot copy.)</p>	<p>Edit the following settings in the Open Systems SnapVault Configurator (svconfigurator) and retry the transfer:</p> <ul style="list-style-type: none"> <li>◆ Set the “Write inactivity period (seconds)” from the default of 5 to 2.</li> <li>◆ Set “Synchronization timeout (seconds)” from the default of 60 to 120.</li> </ul>
	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.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	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.	<p>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.</p> <p>Check the SnapMirror log on the secondary system and the SnapVault log on the primary system for errors.</p>
	Unexpected read select while no data pending	<p>This error can occur because of the following reasons:</p> <ul style="list-style-type: none"> <li>◆ The snapvault abort command was issued using the command line on the SnapVault secondary storage system.</li> <li>◆ There was a network error.</li> </ul>	Check the SnapMirror log, on the secondary storage system and the SnapVault log, on the primary storage system for errors, and then retry the transfer.
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.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
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.
2755	Open System SnapVault upgrade (2.2 to 2.6) fails	The installation was done in one of the following ways: <ul style="list-style-type: none"> <li>1. Remote desktop or terminal server is used to connect to the Windows 2000 machine.</li> <li>2. The installation happens from a mapped network drive.</li> </ul>	This problem is due to a limitation in the Windows 2000 system and has been removed in later Windows versions.
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.
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.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
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: <ul style="list-style-type: none"> <li>◆ The database is in use by some other process.</li> <li>◆ You might not have adequate permission.</li> <li>◆ The database is corrupted.</li> </ul>	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.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
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 storage 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.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
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.
	Root Inode has changed	Open Systems SnapVault update fails with this error message if the base directory is deleted and another directory with the same name is created.	If the SnapVault relationship anchor directory was intentionally deleted from the Open Systems SnapVault primary storage system and it no longer needs to be backed up, delete the corresponding SnapVault relationship on the SnapVault secondary storage system and release the relationship from the primary storage system.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	Insufficient system resources exist to complete the requested service	This error occurs due to insufficient disk space or insufficient memory.	<p>Try one of the following:</p> <ul style="list-style-type: none"> <li>◆ Third party software like antivirus scanners consume a lot of paged pool memory. You could disable such softwares and retry the restore operation.</li> <li>◆ Increase the paged pool size. Set the registry value as follows:  HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\PagedPoolSize to 0xFFFFFFFF (the maximum possible)  Reboot the system and retry the restore operation.</li> </ul>

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	<p>Root Inode has changed</p> <p>Failed to generate update inode values</p>	<p>Open Systems SnapVault update fails with this error message if the base directory is renamed and another directory with the same name is created.</p>	<p>If the SnapVault relationship anchor directory was intentionally renamed, create a new SnapVault relationship for the renamed anchor directory and perform a baseline transfer for this new relationship. When previous backups of the directory with the original name are no longer needed, delete the SnapVault relationship for the anchor directory with the original name, and release the relationship from the Open Systems SnapVault primary storage system.</p> <p>If the SnapVault relationship anchor directory was erroneously renamed, rename it back to its original name and continue performing update backup transfers.</p>

Error number	Error message	Cause	Solution
	OSSV fails to install, 'postinstall.sh' script exits with error code	The script 'postinstall.sh', which is run as part of the installation, exits with this error. The output from the script may be found in the file <code>..\local Settings\temp\postinstall.sh.o utput6</code>	Set the system variables <i>temp</i> and <i>tmp</i> to point to the standard path <code>C:\temp</code> , then rerun the Open Systems SnapVault setup.  Also, check if port 10000 on the client machine is occupied by another application. To do so, enter <b>netstat -an</b> command. Port 10000 is used for NDMP by default. If it is taken by another application, change the port during the Open Systems SnapVault setup. Also ensure that no firewall is blocking the port that is used by Open Systems SnapVault.
	Unique seeding fails - can't get hostname	During an installation of Open Systems SnapVault on either the UNIX or Linux server, the post installation script fails with this error because the host name is not set properly.	Verify the host name setting on the UNIX or Linux server using the <code>hostname</code> command. If the host name is incorrect or not set, then correct the name before installing Open Systems SnapVault.
	Failed to output Database Backup File	An Open Systems SnapVault transfer displays this error message if the <code>\tmp</code> directory is missing in the system.	To complete the update to Open Systems SnapVault, re-create the <code>\tmp</code> directory.
	Failed to open the history file.	This error occurs in one of the following conditions: <ul style="list-style-type: none"> <li>◆ The history file is deleted.</li> <li>◆ The history file is corrupted.</li> <li>◆ The <code>\tmp</code> directory is missing on the system.</li> </ul>	Check whether the history file is deleted or corrupted. If yes, contact technical support. Check whether the <code>\tmp</code> directory exists on the system.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	<p>Open Systems SnapVault backups may fail when multiple drives are backed up simultaneously.</p> <p>When this problem occurs, the following error messages are seen in Windows system eventlog:</p> <ul style="list-style-type: none"> <li>◆ Timeout (30000 milliseconds) waiting for the Microsoft Software Shadow Copy Provider service to connect.</li> <li>◆ The service did not respond to the start or control request in a timely fashion.</li> </ul>	<p>For some unknown reason, the Microsoft Shadow Copy Provider service fails to start. Hence, the Volume ShadowCopy Service fails to create a Snapshot copy.</p>	<ul style="list-style-type: none"> <li>◆ Increase the retry count.</li> <li>◆ Schedule the backups such that multiple drives are not backed up simultaneously.</li> </ul>

The following error messages are displayed on the SnapVault secondary storage systems.

Error number	Error message	Cause	Solution
	Transfer from <SRC> to <DEST>: request denied by the source storage system; check access permission on the source storage system.	The secondary system has denied the request from Open Systems SnapVault to perform the operation.	<p>On the Open Systems SnapVault agent, check whether the QSM Access List in the Configurator utility contains the IP address or network name of the secondary system.</p> <p>For more information, see “<a href="#">Modifying the Qtree SnapMirror™ access list</a>” on page 58.</p>
	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.	<p>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 <a href="http://now.netapp.com">http://now.netapp.com</a>.</p>
	<SRC> <DEST> Abort (could not read from socket)	<p>This message can occur because of one of the following problems on the primary storage system:</p> <ul style="list-style-type: none"> <li>◆ The <code>snapvault abort</code> command has been issued through the command line on the primary system.</li> <li>◆ The Open Systems SnapVault primary system has failed or rebooted.</li> <li>◆ A network error has been encountered.</li> </ul>	<p>Check the following logs for more information:</p> <ul style="list-style-type: none"> <li>◆ The SnapMirror log on the secondary system</li> <li>◆ The SnapVault log on the primary system</li> </ul>

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	<SRC> <DEST> Abort (replication destination does not have a directory that the source has modified)	This error is 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 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, 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 the IDLE state before attempting a transfer.
	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.

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	<p>SnapMirror: Message from Read Socket : Connection reset by peer</p> <p>SnapVault: destination transfer from &lt;SRC&gt; to &lt;DEST&gt; : source volume is offline, is restricted, or does not exist.</p>	<p>This is often caused when Open Systems SnapVault is not running on the host machine at the time of backup. It can also be caused by incorrect typing or a very busy network.</p>	<p>Check the SnapMirror log on the secondary system and the SnapVault log on the primary system for errors and retry the transfer.</p>
	<p>Data Ontap changed data in &lt;X&gt; data chunk(s) in &lt;DEST&gt; which may have been missing or incorrect.</p>	<p>This error can be caused when Data ONTAP detects that incomplete Open Systems SnapVault block-level incremental backups have occurred.</p>	<p>For more information, see bugs 137685 and 140930 at <a href="http://now.netapp.com">http://now.netapp.com</a></p>
	<p>Destination transfer from &lt;SRC&gt; to &lt;DEST&gt; : qtree snapmirror destination found a mismatch between a directory entry and its inode information</p>	<p>Usually caused by a Data ONTAP or Open Systems SnapVault bug.</p>	<p>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.</p>
	<p>&lt;SRC&gt; &lt;DEST&gt; Abort (replication destination failed to store entry in inode map)</p>	<p>Usually caused by a Data ONTAP or Open Systems SnapVault bug.</p>	<p>Ensure that you are running the most recent recommended versions of Open Systems SnapVault on the primary system and Data ONTAP on the secondary system.</p>

<b>Error number</b>	<b>Error message</b>	<b>Cause</b>	<b>Solution</b>
	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 <a href="http://now.netapp.com">http://now.netapp.com</a> .
	service not enabled on the source	The Open Systems SnapVault primary system does not permit backing up of data from a mapped drive on a primary storage system.	This operation is currently not supported.
	Source qtree is not accessible	During deployment and reconfiguration, this error is displayed when starting an Open Systems SnapVault task.	Ensure that the service account (the account that is running the Open Systems SnapVault service on the client) matches the credentials on the secondary storage system or within DataFabric Manager, if present. Access is granted both ways to and from the secondary storage system as well as through DataFabric Manager, if present.
	Base snapshot for transfer no longer exists on the source	This error occurs in one of the following conditions: <ul style="list-style-type: none"> <li>◆ The history file is deleted.</li> <li>◆ The history file is corrupted.</li> </ul> The \tmp directory is missing on the system.	Check whether the history file is deleted or corrupted. If yes, contact technical support. Check whether the \tmp directory exists on the system.
	Source qtree does not exist	The directory that you are trying to backup does not exist on the primary system.	Check if the directory exists on the primary system.

Error number	Error message	Cause	Solution
	missing libsv.dll	<p>This error can occur if the target directory path was specified using forward slashes (/) instead of back slashes (\) in the targetdir path variable.</p> <p>Prior to the Open Systems SnapVault 2.6 release, this error occurs due to unattended installation.</p>	<ol style="list-style-type: none"> <li>1. Verify if the unattended batch file, unattinstall.bat, was generated using svconfigpackager.exe.</li> <li>2. Determine if the target directory path is correct (that is, the target directory path is specified using forward slashes (/) in the targetdir path variable).</li> </ol>

**Configuration options for Microsoft Volume Shadow copy Services (VSS) in Open Systems SnapVault**

**VSS Snapshot creation timeout (secs):**

Using this configuration, you can set the amount of time (Snapshot timeout) that Open Systems SnapVault waits until it retries a VSS Snapshot copy in case of transient errors. The default value is 180 seconds (the maximum value). The minimum value is one second.

**List of drives not to Snapshot:**

Using this configuration, you can prevent Open Systems SnapVault from taking VSS Snapshot copy and use the live file system. This can be used for volumes that are not supported by VSS and when there are some unsolvable problems with Snapshot creation.

**[VSS:UseSystemProvider]:**

Using this configuration, you can force Open Systems SnapVault to use the Microsoft software system provider when the default provider is changed and cannot support Open Systems SnapVault's Snapshot calls. The default value is FALSE.

**VSS error messages:**

Open Systems SnapVault handles the error messages that are returned by VSS. Open Systems SnapVault aborts the transfer on some error messages and retries the transfer for some error messages. The errors on which Open Systems SnapVault retries the transfer are transient errors. The following errors are transient errors:

Error messages	Description
VSS_E_SNAPSHOT_SET_IN_PROGRESS	The creation of a shadow copy is in progress, and only one shadow copy creation operation can be in progress at one time.
VSS_E_FLUSH_WRITES_TIMEOUT	The system was unable to flush I/O writes.
VSS_E_HOLD_WRITES_TIMEOUT	The system was unable to hold I/O writes.
VSS_E_PROVIDER_VETO	The provider was unable to perform the request at this time.
VSS_E_UNEXPECTED_PROVIDER_ERROR	The provider returned an unexpected error code.

When any of the above errors occur, Open Systems SnapVault sleeps for a second and retries creation of Snapshot copy till the time specified by the VSS Snapshot copy creation timeout configuration option.

Other errors like VSS\_E\_VOLUME\_NOT\_SUPPORTED are non-transient and Open Systems SnapVault aborts the transfer immediately.



# Index

---

## Numerics

- 32-bit systems 18
- 64-bit systems 20

## A

- aborted transfer, retrying 125
- ACL file 93
- Active Directory 12, 47, 98, 99
- Allowed Secondary Names window 24
- Always setting (for generating debugging information) 63
- applications for managing Open Systems SnapVault 6

## B

- backing up
  - disable database 97
  - EFS files 13, 127
  - open files 9
  - using Open Systems SnapVault for 90
  - Windows System State 12
- backup
  - syntax to backup virtual machines 154
- backup exclusion lists 10, 68
- backup file, database (OSSV\_DATABASE\_BACKUP) 94
- backups
  - EFS files 13
  - scheduling 82
  - System State data 102
- baseline copy, creating 81
- batch installation 47
- BLI backups
  - checksums file 93
  - described 10
  - for EFS files 11
  - for name-based applications 11
  - level, setting 66
  - name-based 67
- broken relationship

- restoring 123
- resuming 13
- built-in space estimator
  - description 129
  - disabling 138
  - enabling 136
  - installing 136

## C

- calculation of free space 129
- checkpoint
  - file 93
  - restart 12
- checksum computation, BLI checksum 66
- command-line utility (svsetstanza) 5, 53
- command-line utility for administration 53
- commands
  - restore 90
  - snapvault modify 125
  - snapvault release 108, 117
  - snapvault restore 6, 85, 95, 97
  - snapvault snap sched 82
  - snapvault start 5, 81, 102, 108, 123
  - snapvault stop 108
  - svconfigpackager 42
  - svestimator 137
  - svinstallcheck 35
  - svpmgr shutdown 36
  - svpmgr startup 37
  - svsetstanza 5, 53
- configuration files
  - configure.cfg 53
  - estimator.cfg 131
- Configurator utility
  - description 51
  - General tab, purpose of 52
  - Machine tab, purpose of 51
  - running 56
  - Service tab, purpose of 51
  - SnapVault parameter, modifying 58
  - SnapVault tab, purpose of 52

- Trace tab, purpose of 52
- configure.cfg file 53
- copying files to restore 85

## D

- Data ONTAP version, supported 19
- database files
  - ACL 93
  - backup file (OSSV\_DATABASE\_BACKUP)
    - 94
  - BLI checksums 93
  - checkpoint 93
  - described 93
  - history 93
- databases
  - backup 95
  - disabling backup 97
  - disk space requirements for 19
  - Open Systems SnapVault 93
- DataFabric Manager, using to manage Open Systems SnapVault 6
- debug files
  - generation of 62
  - stop generating 64
- determining free space 129

## E

- EFS files
  - about 13
  - BLI and 11
- error messages 161
- estimating free space 13
- estimator.cfg file
  - contents of 132
  - described 131
- ESX Server 142
  - definition 140
  - Overview 142
- exclusion lists 10

## F

- failed transfers, retrying 125
- file exclusion lists

- described 10
- for space estimator 131
- location of 68
- free space estimator. *See* space estimator

## G

- General tab (Configurator) 52
- Generate debugging files option 62
- GUI interface for administration 5, 51

## H

- history file 93

## I

- initial backup (baseline) 81
- installation
  - Allowed Secondary Names window 24
  - batch 47
  - built-in space estimator 136
  - limitations 20
  - on HP-UX, AIX, or Linux 30
  - on Solaris 27
  - on UNIX and Linux 27
  - on Windows 22
  - requirements, primary systems 16
  - script for unattended install 42
  - Select Installation Folder window 25
  - Setup wizard 23
  - stand-alone space estimator 136
  - unattended, about 42
  - upgrading. *See* upgrade
- interfaces
  - Setup wizard 23

## L

- Libnormal setting 63
- Libverbose setting 63
- licenses 20
- limitations of Open Systems SnapVault 20
- listen port, NDMP 24
- lists, exclusion 10
- log files

- SnapVault 92
  - space estimator 135
  - trace 135
- log files, SnapVault 52

## M

- Machine tab (Configurator utility) 51
- management application, Open Systems SnapVault 6
- messages, error 161
- migrating a volume 109

## N

- name-based
  - applications, about 11
  - block-level incremental backup, setting 67
- NDMP Account field 59
- NDMP Host ID field 59
- NDMP Host Name field 59
- NDMP Listen Port field 24, 59
- NDMP-based management applications 6
- NetApp Host Agent
  - installing 25
  - supportability matrix 19
  - uninstalling 39, 41
- Normal setting 63

## O

- OFM
  - about 9, 71
  - configuring 71
  - disabling 72
  - parameters, setting 72
  - St. Bernard's module 22
- OFM licenses 71
- Open File Manager. *See* OFM
- open files, backing up 9
- Open Systems SnapVault 2
- OSSV. *See* Open Systems SnapVault
- OSSV\_DATABASE\_BACKUP file 94

## P

- parameters modifying, SnapVault 58
- path exclusion lists
  - about 10
  - described 68
  - for space estimator 131
  - location of 68
- port
  - listen, NDMP 24
- Port requirements
  - FILESERVER 10555 144
  - NDMP 10000 144
  - QSMSEVER 10566 144
- ports
  - 10000 19, 24, 35, 161
  - 10566 19
  - 4092 19
  - 4093 19
- preinstallation requirements 16
- prerequisites
  - for installation 16
- primary storage system
  - about 3
- primary system
  - administration 5
  - controlling secondary system access 58
  - installation requirements 16
  - restoring 89
  - restoring data 86
  - svrestore command 86

## R

- relationships
  - broken 123
  - deleting 108
  - migration 109
  - resynchronizing 123
- remote installation
  - and upgrade 42
  - batch 47
- reports, problem 92
- requirements (installation), for primary systems 16
- restart backup from a checkpoint 12
- Restore 155

- restore command 90
  - Restore of virtual machine
    - restore example 155
    - three methods to restore 155
  - restoring 93
    - a complete system 89
    - broken relationship 123
    - by copying files 85
    - EFS files 127
    - from tape 90
    - methods for 85
    - System State data 103
    - Windows Systems State 12
  - resuming broken relationships 13
  - resynchronizing a relationship 13, 123
  - retrying a failed transfer 125
- S**
- scheduling backups 82
  - script for unattended install 42
  - secondary system
    - about 3
    - access control 58
    - administration 5
    - configuring for use with open systems 79
    - enabling SnapVault 80
    - installation requirements 19
    - migrating relationship between 108, 109
  - security, enabling or disabling 58
  - Select Installation Folder window 25
  - Service console
    - definition 141
  - Service tab (Configurator utility) 51
  - Setup Wizard 23
  - SnapMirror, using for protecting 115
  - SnapVault
    - configuring the secondary storage system 79
    - log files 52, 92
    - primary system parameters, modifying 58
  - snapvault modify command 125
  - snapvault release command 108, 117
  - snapvault restore command 6, 85, 95, 97
  - snapvault snap sched command 82
  - snapvault start command 5, 81, 102, 108, 123
  - snapvault stop command 108
  - SnapVault tab (Configurator utility) 52
  - snapvault.cfg file (for space estimator), modifying 138
  - softlock support 115
  - space estimator
    - about 13
    - built-in
      - description 129
      - disabling 138
      - installing 136
    - calculates, how it 129
    - console output 135
    - described 129
    - enabling 136
    - estimator.cfg file 131
    - log files 135
    - path and file exclusion lists 131
    - stand-alone 129, 136
    - utility 19
  - St. Bernard's OFM module 22
  - stand-alone space estimator
    - described 129
    - installation 136
  - status logs 92
  - sv\_linux\_pri license 20
  - sv\_ontap\_sec license 20
  - sv\_unix\_pri license 20
  - sv\_vmware\_pri license 20
  - sv\_windows\_pri license 20
  - svconfigpackager command (utility) 42
  - svestimator command 137
  - svinstallcheck command 35
  - svpmgr shutdown command 36
  - svpmgr startup command 37
  - svsetstanza command 5, 53
  - system restore, primary 89
  - System State data
    - backup 102
    - restore 103
- T**
- tape, restoring from 90
  - tertiary system, data protection with 115

- tombstone lifetime setting 102
- trace file 135
- Trace Level tab (Configurator utility) 52
- troubleshooting 161
- troubleshooting, generating debug files for 62

## U

- unattended installation 42
- uninstall, how to 39
- unsupported configurations 20
- upgrade
  - from versions 1.x
  - unattended 42
- UUID
  - definition 140

## V

- Verbose setting 63
- Virtual machine
  - definition 140

- virtual machine
  - process to backup a virtual machine 153
  - virtual machine backup 154
- Virtualization
  - overview 142
- VMDK
  - definition 140
- VMotion 158
  - definition 140
- VMX
  - definition 141
- volume migration 109
- VSS Snapshot copy
  - about 9, 71
  - setting timeout for 73

## W

- Web interface for administration 5, 51
- Windows System State 12
- wizard, setup 23

