

Open Systems SnapVault® 3.0 **Installation and Administration Guide**

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

Part number: 215-05190_A0
August 2010

Copyright and trademark information

Copyright information

Copyright © 1994-2010 NetApp, Inc. All rights reserved. Printed in the U.S.A.

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

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP 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.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. 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 NetApp.

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 NetApp logo; the Network Appliance logo; Bycast; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StorageGRID; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, 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.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Introduction to Open Systems SnapVault	1
	About Open Systems SnapVault	2
	Open Systems SnapVault features	9
Chapter 2	Installing the Open Systems SnapVault Software.	25
	Prerequisites	26
	Installing Open Systems SnapVault on Windows platforms	32
	Installing Open Systems SnapVault on UNIX and Linux platforms	36
	Verifying the installation	41
	Upgrading to Open Systems SnapVault 3.0	42
	Uninstalling Open Systems SnapVault	46
	Unattended installation and upgrade	48
Chapter 3	Configuring Open Systems SnapVault.	57
	Configuration interfaces	58
	Understanding the Configurator utility interface.	59
	Understanding the svsetstanza command	61
	Running the Configurator utility	64
	Confirming that services are running	65
	Modifying Open Systems SnapVault parameters	66
	Enabling and disabling debugging	70
	Setting block-level incremental backup options	73
	Configuring backup exclusion lists	75
	Configuring open file backup for Windows	78
	Configuration for preserving Snapshot copies	80
	Configuration for DataFabric Manager restore to non-ASCII path	81
	Primary storage system reporting through AutoSupport	82

Chapter 4	Microsoft Cluster Services Support	83
	Microsoft Cluster Services Support on Open Systems SnapVault	84
	Setting up and configuring a two-node cluster.	88
	Protection Manager support for Microsoft Cluster.	90
	Uninstalling Open Systems SnapVault in an MSCS environment.	93
Chapter 5	Perform Backup and Restore.	95
	Perform SnapVault backup on Open Systems platforms.	96
	Configuring the SnapVault secondary storage system.	97
	Creating an initial baseline copy	99
	Scheduling SnapVault update backups.	100
	Backing up empty source directories.	101
	Perform SnapVault restore on Open Systems platform	102
	Restoring a directory or a file	103
	Restoring an entire primary storage system	107
	Restoring files to a primary storage system from tape.	108
	Volume mountpoint data backup and restore.	109
Chapter 6	Microsoft SQL Server Backup and Restore.	111
	Configuring backup and restore of Microsoft SQL databases	117
	Viewing SQL Server database from the command-line interface	124
	Backing up and restoring Microsoft SQL Server databases	128
	Backing up using the command-line interface.	129
	Restoring using the command-line interface.	131
	Backing up Microsoft SQL Server database using Protection Manager.	132
	Restoring Microsoft SQL Server database using Protection Manager.	135
Chapter 7	Open Systems SnapVault Management	137
	Locating status and problem reports	138
	Backing up and restoring the Open Systems SnapVault database	140
	Backing up and restoring Windows System State data.	145
	Adding System State data backup	149
	Restoring System State data	150
	Using System State data backup to rebuild a primary storage system.	152
	Deleting and re-creating Open Systems SnapVault relationships	155
	Migrating a relationship between two secondary storage systems	156

	Migrating between two volumes on one secondary storage system	159
	Setting up a tertiary system for a relationship	162
	Reusing a deleted or renamed primary backup root directory name	164
	Reusing a renamed Open Systems SnapVault primary host name	166
	Renaming a SnapVault secondary volume	167
	Resynchronizing restored or broken relationships	169
	Retrying failed transfers	172
	Encrypted File System (EFS) file backup and restore	174
Chapter 8	Changelog Minifilter Driver	175
	Changelog filter driver management	179
Chapter 9	Open Systems SnapVault Space Estimator	187
Chapter 10	Open Systems SnapVault solution for VMware.	197
	VMware terminology	198
	Overview of Virtualization and VMware ESX.	200
	Open Systems SnapVault on ESX server.	201
	Installing Open Systems SnapVault 3.0 on ESX server	203
	Configuration of Open Systems SnapVault on ESX server	204
	Backup and restore of virtual machines	211
	Open Systems SnapVault support for VMotion	216
Appendix A	The OSSVINFO Tool	217
Appendix B	Open Systems SnapVault Error Messages.	219
	Index	251

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, Linux, Solaris®, 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	<p>When the term “storage system” is used, the information applies to all NetApp storage system models.</p> <p>In this document, the term “storage system” is also referred to as “filer”.</p> <p>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.</p>
Command conventions	You can enter storage system commands on the system console or from any client that can obtain access to the storage system by using a Telnet session. In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ from what you see, depending on the version of UNIX that you are using.
Formatting conventions	The following table lists different character formats used in this guide to set off special information.

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

Keyboard conventions

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

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

Special messages

This guide contains special messages that are described as follows:

Note

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

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 that resides on the following platforms:

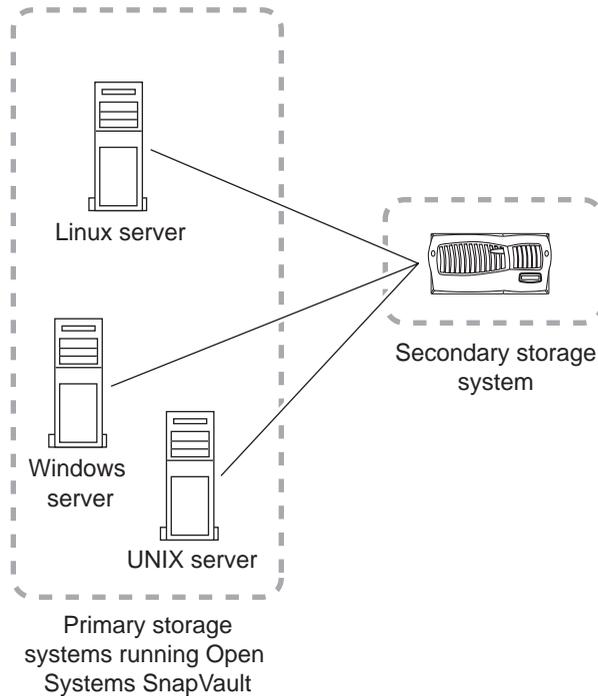
- ◆ Microsoft Windows
- ◆ Red Hat® Enterprise Linux
- ◆ Novell® SUSE® Linux Enterprise Server
- ◆ Sun Solaris™
- ◆ IBM AIX®
- ◆ HP-UX®
- ◆ VMware® ESX

Note

For a list of currently supported versions of these platforms, see “[Requirements for primary storage systems](#)” on page 26.

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 26.
2. The Open Systems SnapVault agent—the software that is installed on the primary storage system
3. The secondary storage system—the NetApp 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 29.

How Open Systems SnapVault works

You identify the directories or file systems and the qtrees on the primary and secondary storage system to back up data.

For the first backup, the secondary storage system requests an initial baseline of the primary data. This transfer establishes a SnapVault relationship between the primary data and the SnapVault secondary qtrees.

The subsequent backups can be scheduled or performed manually. 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, you can transfer whole files or only the changed blocks 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 retains the number of secondary storage Snapshot copies you specify. The SnapVault secondary storage system assigns each Snapshot copy in the set a version number, beginning with 0 (zero) for the most recent and so on.

If you want to restore a directory or file data to the primary system, SnapVault retrieves the data from the specified Snapshot copy and transfers the data to the primary storage system.

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

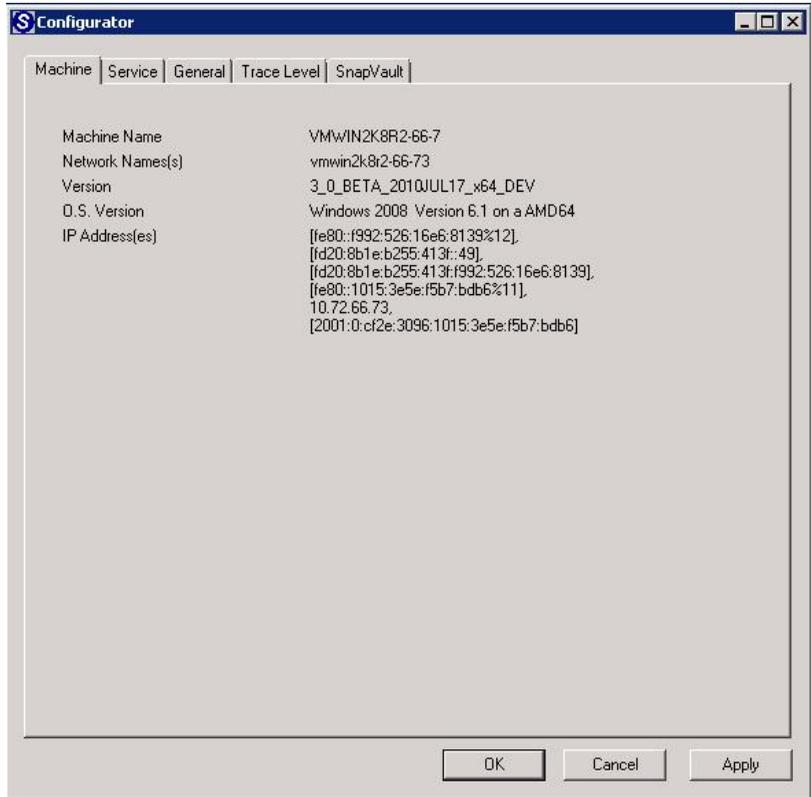
Although Open Systems SnapVault backs up and restores data in a manner similar to how SnapVault backs up and restores data from NetApp storage systems, the following differences exist:

- ◆ With Open Systems SnapVault, you can back up and restore qtrees, volumes, and 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.

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 59.

- ◆ 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 61.

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

Sample commands

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

```
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, enter the following command from the Open Systems SnapVault primary system:

```
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.

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 to be restored. You can restore the subdirectory to the Windows system directory C:\restored.

For an example of a single file restore, see [“Using the snapvault restore command”](#) on page 104.

Central management of Open Systems SnapVault agents

You can manage Open Systems SnapVault 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.

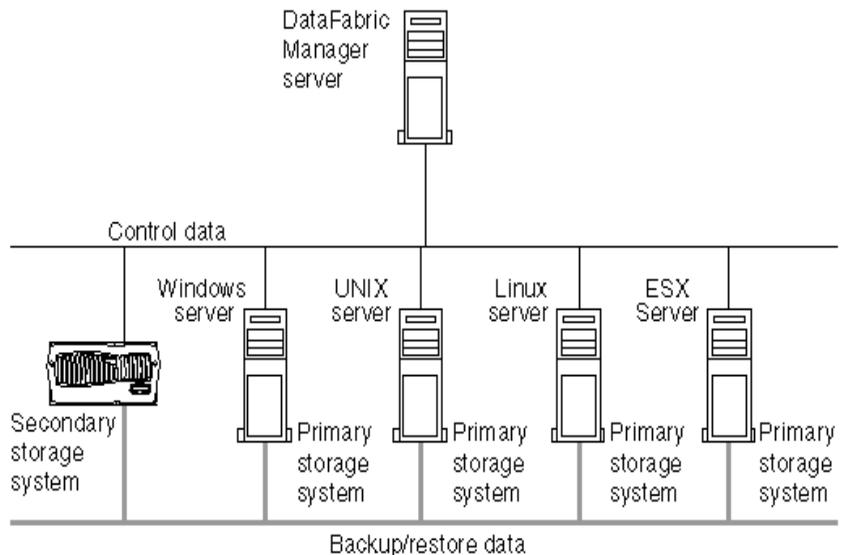
Note

The management application uses NDMP for communication.

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

Vendor	URL
NetApp - Protection Manager	www.netapp.com
BakBone Software® NetVault®	www.bakbone.com
Syncsort® Backup Express®	www.syncsort.com
CommVault® Simpana	www.commvault.com

The following illustration shows a typical Open Systems SnapVault setup that uses the Protection Manager component within 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. You can use Operations Manager to monitor, alert, and report on the storage and NetApp storage system infrastructure.

Protection Manager provides policy-based management for Open Systems SnapVault. Protection Manager controls the start and stop services of Open Systems SnapVault with the help of NetApp Host Agent. You must install NetApp Host Agent on the same primary system where Open Systems SnapVault is running.

Additional information

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

- ❖ Data ONTAP *Data Protection Online Backup and Recovery Guide* at <http://now.netapp.com>
- ❖ NetApp Technical Report TR-3234—Leveraging NetApp SnapVault for Heterogeneous Environments, at <http://www.netapp.com>
- ❖ NetApp Technical Report TR-3252—Enabling Rapid Recovery with SnapVault, at <http://www.netapp.com>
- ❖ NetApp Technical Report TR-3466—*Open Systems SnapVault Best Practices Guide*, at <http://www.netapp.com>
- ❖ NetApp Technical Report TR-3653—*Open Systems SnapVault Best Practices Guide for Protecting Virtual Infrastructure*, at <http://www.netapp.com>.

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”](#) on page 9
- ◆ [“Excluding specific files or directories from backup”](#) on page 10
- ◆ [“Setting block-level incremental transfer”](#) on page 10
- ◆ [“Backing up and restoring the Open Systems SnapVault database”](#) on page 11
- ◆ [“Backing up and restoring Windows System State”](#) on page 12
- ◆ [“Checkpoints for restart of transfers”](#) on page 12
- ◆ [“Encrypted File System \(EFS\) files”](#) on page 13
- ◆ [“Determining free space using the space estimator”](#) on page 13
- ◆ [“Resynchronizing a broken relationship”](#) on page 13
- ◆ [“Compression on primary and secondary storage systems”](#) on page 13
- ◆ [“Microsoft Cluster Support”](#) on page 17
- ◆ [“Microsoft SQL Server database backup and restore”](#) on page 17
- ◆ [“Changelog minifilter driver for faster incremental backup”](#) on page 17
- ◆ [“Volume mountpoint backup and restore”](#) on page 17
- ◆ [“IPv6 support”](#) on page 18
- ◆ [“Client-based bandwidth throttling”](#) on page 20

Backing up open files

Volume Shadow Snapshot copy Service (VSS) Snapshot copy—Used for Windows 2003 and later platforms.

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

For more information, see [“Configuration for preserving Snapshot copies”](#) on page 80.

Common Snapshot Management: Prior to Open Systems SnapVault 2.6 release, whenever the transfer of files failed, the VSS Snapshot copies were deleted on the primary system and a new Snapshot copy was created when the

transfer restarted. In case of common snapshot management, Open Systems SnapVault 2.6 and later has the ability to retain old Snapshot copies and use the copies when the transfer is restarted.

Excluding specific files or directories from backup

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

- ◆ File exclusion lists—You can exclude a file or directory if the file name or any path element matches a file exclusion entry in the list.
- ◆ File system exclusion list—A file system is excluded if the file system type matches the file system exclusion entry.
- ◆ Path exclusion lists—If a path exclusion entry specifies a directory, you can exclude the directory and its files and subdirectories.

For more information, see “[Configuring backup exclusion lists](#)” on page 75.

By default, Open Systems SnapVault ignores the following files:

- ◆ pagefile.sys
- ◆ hibernate.sys

Additionally, Open Systems SnapVault does not back up the files under the following registry key:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\BackupRestore\FilesNotToBackup.

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. 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. 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 74.

For more information about how to configure or change BLI settings in Open Systems SnapVault, see [“Setting block-level incremental backup options”](#) on page 73.

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. It is backed up with every backup transfer.

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 should 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. 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 140.

Backing up and restoring Windows System State

You can back up and restore Windows System State data by using Open Systems SnapVault. This is useful when, for example, an Active Directory entry is accidentally deleted. You can also use Open Systems SnapVault System State data backup along 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 specified 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 145.

Checkpoints for restart of transfers

When an Open Systems SnapVault backup process fails, checkpoint restart support allows the backup (baseline or update) to resume 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 retry value 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 172.

Block-level checkpoints: Prior to Open Systems SnapVault 2.6 release, checkpoints were allowed only at the end of files. If you transferred large number of files, the checkpoint mechanism was not useful, 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 and later support 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 300 seconds (5 minutes).

For more information about how to configure block-level checkpoints, see [“Configuring the checkpoint interval”](#) on page 172.

Encrypted File System (EFS) files

You can back up and recover EFS files on Windows systems; however, EFS files cannot be backed up using block-level incremental backup. If you modify an EFS file, Open Systems SnapVault backs up the entire EFS file.

For more information, see [“Encrypted File System \(EFS\) file backup and restore”](#) on page 174.

Determining free space using the space estimator

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

For more information, see [“Open Systems SnapVault Space Estimator”](#) on page 187.

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. This 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 169.

Compression on primary and secondary storage systems

The compression feature of Open Systems SnapVault enables data compression over the network. This feature helps to optimize the bandwidth for Open Systems SnapVault data transfers.

Data ONTAP 7.3 and later support bandwidth optimization in Open Systems SnapVault using the compression feature. The compression of data reduces the amount of data sent over the network, thereby optimizing the bandwidth and enabling efficient backup and restore transfers. This section describes the compression of the Open Systems SnapVault data stream.

For more information about how to set up backup relationships by using different compression options on the secondary storage system, see the *Data ONTAP 7.3 Release Notes*.

Note

Starting with Data ONTAP 7.3.1, you can run the following command on the secondary storage system to enable the compression feature for all Open Systems SnapVault relationships:

```
options snapvault.ossv.compression on
```

By default, the compression feature is enabled on the primary system where Open Systems SnapVault is installed.

On the primary storage system: The following configuration options for the compression feature are available on the Open Systems SnapVault primary storage system.

You can use the `svsetstanza` utility to set the following options:

- ◆ To disable compression for all relationships on the primary storage system:

```
svsetstanza.exe config snapvault.cfg QSM
```

```
"EnableCompression" Value FALSE FALSE
```

You can view the value in the `snapvault.cfg` file.

```
[QSM:EnableCompression]
```

```
Value=FALSE
```

The default value is TRUE.

- ◆ To set the compression level when compression is enabled:

```
svsetstanza.exe config snapvault.cfg QSM
```

```
"CompressionLevel" Value MEDIUM FALSE
```

You can view the value in the `snapvault.cfg` file.

```
[QSM:CompressionLevel]
```

```
Value=MEDIUM
```

You can set the flag to LOW, MEDIUM, or HIGH. The default value is MEDIUM.

If you set the value to LOW then the percentage of data that is compressed is low.

If you set the value to HIGH, the percentage of data that is compressed is high depending on the type of data. For example, if the data is a text file the compressed ratio is high, and if the data is an image file then the compressed ratio is less.

To set compression at a lower priority to minimize the CPU usage, when compression is enabled:

```
svsetstanza.exe config snapvault.cfg QSM
```

```
"CompressionLowPriority" Value TRUE FALSE
```

You can view the value in the snapvault.cfg file.

```
[QSM:CompressionLowPriority]
```

```
Value=TRUE
```

The default value is FALSE.

If the value is set to TRUE, it maximizes the CPU usage, thus, increasing the compression time.

To enable compression for restore transfers, use the `-c` option with the `snapvault restore` command on the primary system.

Step	Action
1	<p>On the console of the primary system, enter the following command:</p> <pre>snapvault restore -c -S <secondary>:<qtree> <primary path></pre> <p><i>qtree</i> is the qtree on the secondary storage system where the data is backed up.</p> <p><i>secondary</i> is the path where the data is backed up in the secondary storage system.</p> <p><i>primary path</i> is the path where you want to restore the data.</p>

You can use the `snapvault status -l` command to see compression statistics such as the compression ratio for baseline, update, and restore transfers. For more information about the compression ratio, see the SnapVault log files. The SnapVault log files contain information about the actual size of the data and the compressed data set.

Example:

On the secondary storage system, run the following command to start the backup with compression enabled for a new Open Systems SnapVault relationship:

```
snapvault start -o compression=on -s 10.73.44.32:c:\test
/vol/vol10/test21
```

When the backup is in progress, you can view the compression ratio on the secondary storage system. When the backup is over, you can view the percentage of the data compressed on the primary storage system.

Result: When the status of the data is transferring, the compression ratio is displayed on the secondary storage system.

```
Source:                10.73.44.32:c:\test
Destination:           f3050-230-53:/vol/vol10/test21
Status:                Transferring
Progress:              2088 KB
Compression Ratio:     1.9 : 1
State:                 Uninitialized
Lag:                   -
Mirror Timestamp:     -
Base Snapshot:         -
Current Transfer Type: Initialize
Current Transfer Error: -
Contents:              Transitioning
Last Transfer Type:    -
Last Transfer Size:    -
Last Transfer Duration: -
Last Transfer From:    -
```

When the backup transfer is over and the status of the data is Idle, you can see the percentage of data that is being compressed on the primary storage system.

```
Source:                10.73.44.32:c:\test
Destination:           f3050-230-53:/vol/vol10/test21
Status:                Idle
State:                 Source
Lag:                   00:01:48
Mirror Timestamp:     -
Base Snapshot:         -
Current Transfer Type: -
Contents:              -
Last Transfer Type:    -
Last Transfer From:    -
Last Transfer Size:    1108 KB (Compression=48%)
```

```
Last Transfer Duration:      00:00:06
Total files to transfer:     2
Total files transferred:    2
Current File Size:          -
Current File Progress:      -
Current File Name:          -
Transfer Error ID:          -
Transfer Error Message:     -
```

Note

When compression is enabled on the primary and the secondary storage systems, it activates the compression feature in an Open Systems SnapVault relationship.

If you want to disable the compression feature on a particular primary storage system, set the [QSM:EnableCompression] value to FALSE on the primary storage system.

Microsoft Cluster Support

You can configure Open Systems SnapVault on a two-node Microsoft Cluster Services (MSCS) for backing up data from the cluster nodes. You can also use Protection Manager functionality of DataFabric Manager for backing up and restoring data from cluster nodes. For more information, see “[Microsoft Cluster Services Support on Open Systems SnapVault](#)” on page 84.

Microsoft SQL Server database backup and restore

Open Systems SnapVault enables you to back up and restore Microsoft SQL Server database. You can back up and restore SQL server databases using Open Systems SnapVault command-line utility and Protection Manager functionality of DataFabric Manager. For more information, see “[Microsoft SQL Server Backup and Restore](#)” on page 111.

Changelog minifilter driver for faster incremental backup

Changelog minifilter driver enables faster incremental backups after initial baseline transfer. You can use this feature for application data backup and file system data backup. Open Systems SnapVault allows you to enable or disable the Changelog minifilter driver. For more information, see “[Changelog Minifilter Driver](#)” on page 175.

Volume mountpoint backup and restore

Open Systems SnapVault enables you to back up and restore volume mountpoints on Windows platforms. For more information, see “[Volume mountpoint data backup and restore](#)” on page 109.

IPv6 support

Open Systems SnapVault 3.0 supports both Internet Protocol version 6 (IPv6) and IPv4. In IPv6, the IP address size is 128 bits, which is larger than the IPv4 address size of 32 bits. This larger address space provides expanded routing, security, and addressing capabilities.

You can use Open Systems SnapVault 3.0 in a network environment with either IPv6 or IPv4 only, or in a mixed network environment with both IPv4 and IPv6. Open Systems SnapVault 3.0 can communicate with all IPv6-enabled platforms.

You can use IPv6 and IPv4 addresses for backup and restore purposes. The command-line interface of Open Systems SnapVault and Logical Replication (LREP) tool support IPv6 addresses for backup and restore purposes.

In the command-line interface, you must enclose all IPv6 addresses in parentheses during a backup and restore process. However, you can optionally enclose IPv4 addresses and host names in parentheses during the restore process.

Example of an IPv6 address:

```
snapvault restore -S [0ffa::88fe:3456:7654:AA34]:/vol/vol10/test  
c:\test
```

Note

To use IPv6 for backup and restore purposes, the secondary storage system must be running Data ONTAP 7.3.3 or later in the 7.3 release family.

The output of the `svinstallcheck.exe` utility displays both IPv4 and IPv6 connections with NDMP and `qtree SnapMirror`. For communicating with external interfaces, such as `qtree SnapMirror` and NDMP interfaces, Open Systems SnapVault 3.0 uses both IPv4 and IPv6.

Supported platforms: Open Systems SnapVault supports IPv6 on the following platforms:

- ◆ Microsoft Windows Server 2003, 2008, and 2008 R2
- ◆ Red Hat® Enterprise Linux
- ◆ Novell® SUSE® Linux Enterprise Server
- ◆ Sun® Solaris
- ◆ IBM AIX
- ◆ HP-UX

Note

Backup and restore over IPv6 is possible only if IPv6 is enabled on both primary and secondary storage systems.

Supported connection modes for restore: A connection mode specifies how a host name resolves to IP addresses during restore.

The following are the different supported connection modes for restore:

- ◆ The use of *inet* as the connection mode for IPv4
 - ❖ When you specify a host name with *inet* as the connection mode, the host name resolves to a list of IPv4 address. If the host name cannot resolve with an IPv4 address, then an error is displayed.
- ◆ The use of *inet6* as the connection mode for IPv6
 - ❖ When you specify a host name with *inet6* as the connection mode, the host name resolves to a list of IPv6 address. If the host name cannot resolve with an IPv6 address, then an error is displayed.
- ◆ The use of *unspec* as the connection mode for not specifying any particular connection modes
 - ❖ When you specify a host name with *unspec* as the connection mode, the host name attempts to resolve with all possible addresses (IPv6 and IPv4). Open Systems SnapVault attempts to connect to the addresses until the connection is established.

You can select the connection mode during restore process.

To specify the connection mode during a restore process, you must use the `-m` option followed by the connection mode.

Example:

```
snapvault restore -m inet6 -S nebula:/vol/vol0/test c:\test
```

Note

If you do not specify the preferred connection mode, Open Systems SnapVault uses the *unspec* mode.

Setting IPv6 as preferred type: Open Systems SnapVault 3.0 uses IPv6 or IPv4 for communicating with different internal modules of Open Systems SnapVault. The internal communication among Open Systems SnapVault features depends on the flag value you set for `[Network:Prefer IPV6 for Messaging]` in the `configure.cfg` file.

- ◆ If you set the flag value as `TRUE`, Open Systems SnapVault prefers IPv6 over IPv4. In case IPv6 is not available, then IPv4 is used for internal communication.
- ◆ If you set the flag value as `FALSE`, Open Systems SnapVault prefers IPv4 over IPv6. In case IPv4 is not available, then IPv6 is used for internal communication in Open Systems SnapVault 3.0.

Note

After changing the flag, you must stop and start the Open Systems SnapVault services.

Setting IPv6 as the preferred address type: The *configure.cfg* file includes a new flag [*Network:Prefer IPV6 for Messaging*] that enables you to select IPv6 for internal communication. By default, the flag value is FALSE.

To set IPv6 as the preferred type, complete the following steps:

Step	Action
1	Navigate to the <i>install_dir/snapvault/config</i> directory.
2	In the config directory, open the <i>configure.cfg</i> file.
3	Find for the [<i>Network:Prefer IPV6 for Messaging</i>] flag in the file.
4	Set the value = TRUE.
5	Save and close the file.

Client-based bandwidth throttling

This client-based throttling feature in the Open Systems SnapVault 3.0 release enables you to manage network bandwidth during a SnapVault backup process from a primary storage system to a secondary storage system. It enables you to schedule bandwidth throttling and modify the schedule for the transfers that are currently active. Any changes in the bandwidth throttle apply to the transfers that are currently active and for future transfers. This feature is available in all Open Systems SnapVault supported platforms. You can allocate the required bandwidth for Open Systems SnapVault backups after analyzing the bandwidth usage by other applications in the network. The throttle value must be in KBps (kilobytes per second).

The bandwidth throttling configuration file: Open Systems SnapVault includes a new configuration file called *wan.cfg* for scheduling bandwidth throttling. The *wan.cfg* file enables you to schedule bandwidth throttling, specify a time interval for checking the schedule, and modify the schedule according to your requirements. The *wan.cfg* file is installed in the *install_dir/snapvault/config* directory during the Open Systems SnapVault 3.0 installation.

The *wan.cfg* file has the following parameters:

- ❖ THROTTLE:Schedule
- ❖ THROTTLE:CheckingInterval

For `CheckingInterval`, you can set the value to zero or any value from 60 to 86400 seconds. The default value for the checking interval is 900 sec. You must enter the value in seconds. If the value is zero, the client-based throttling schedule is disabled. If the value is set to the default value or a value in the permissible range, the schedule is enabled.

Note

You must restart the Open Systems SnapVault services after you disable and reenable the client-based bandwidth throttling schedule.

How client-based bandwidth throttling works: Open Systems SnapVault checks the configuration file for bandwidth throttle details before beginning the backup. The configuration file provides the schedule and interval of the throttling. If the interval is zero, there is no client-based throttle for the current transfer.

If Open Systems SnapVault finds the schedule entries in the configuration file, it reads the entries and follows the throttling schedule at the time specified for transfers. If no throttling details are present in the configuration file for the current time data transfer, then Open Systems SnapVault uses the throttle value defined by you using the `-k` option.

However, if you use the SnapVault feature (the `-k` option) also for throttling, then Open Systems SnapVault takes the lowest value of the two values.

For example, if your network usage is high from 9 hours to 20 hours and low from 21 hours to 7 hours, then you can use the `wan.cfg` file to schedule the bandwidth throttle depending on the usage. From 9 hours to 18 hours the network usage is high, you can set the bandwidth throttle as 500 KBps. From 21 hours to 7 hours the network usage is low, you can set the bandwidth throttle as 1000 KBps. The checking interval is 900 seconds by default. For a period that is outside the schedule, Open Systems SnapVault uses the throttle value defined by you using the `-k` option.

If Open Systems SnapVault does not find any schedule entries, it checks for changes to the throttle schedule every 15 minutes and continues data transfers with the default throttle.

All active transfers share the bandwidth throttle. For example, if the schedule throttle value is 2 KBps and if there are two active transfers in the client, each transfer should take 1 KBps.

Scheduling bandwidth throttle: When you schedule the bandwidth throttle, you should follow these guidelines:

- ◆ A comma should be used to separate the throttle entries in the *wan.cfg* file.
- ◆ The throttle entry can be set for any day in a week or a range of days of a week. If you do not specify a day, the throttle entry is used for all the days of in a week.
- ◆ The throttle entry must have a time range in the 24-hour format.
- ◆ The time range should be preceded an @ sign.
- ◆ The throttle value should be preceded by the pound (#) sign.
- ◆ The throttle value should end with an exclamation (!) mark, except for the last throttle value.
- ◆ The Checking interval must be in seconds.

To schedule the bandwidth throttle, complete the following steps:

Step	Action
1	Navigate to the <i>install_dir/snapvault/config</i> directory.
2	In the config directory, open the <i>wan.cfg</i> file.
3	Type the throttle schedule values. Note _____ You should follow the guidelines for scheduling the throttle. _____
4	Type the throttle checking interval.
5	Save and close the file.

Examples:

The following is an example for throttle schedule:

[THROTTLE:Schedule]

Value=mon-thu@9-18#100!,fri@9-18#150!,18-21#200

This schedule translates to:

- ◆ Schedule throttle value of 100 KBps from Monday to Thursday from 9 a.m. to 6 p.m.
- ◆ Schedule throttle value of 150 KBps on Friday from 9 a.m. to 6 p.m.
- ◆ Schedule throttle value of 200 KBps on all days from 6 p.m. to 9 p.m.

[THROTTLE:CheckingInterval]

Value= 900

Open Systems SnapVault checks the bandwidth throttling schedule every 900 sec.

The following is an example for throttle schedule from one particular day to another day in a week:

[THROTTLE:Schedule]

Value=mon-sun@9-18#100

This schedule translates to:

Schedules throttle value of 100 KBps from Monday to Sunday from 9 a.m. to 6 p.m.

[THROTTLE:CheckingInterval]

Value= 900

The following is an example for throttle schedule for some specific days in a week:

[THROTTLE:Schedule]

Value=sat-mon@9-18#100

This schedule translates to:

Schedules throttle value of 100 KBps from Saturday to Monday from 9 a.m. to 6 p.m.

[THROTTLE:CheckingInterval]

Value= 900

The following is an example for overlapping bandwidth throttle schedules:

[THROTTLE:Schedule]

Value=sat-mon@9-10#100!,sat-mon#9-18#200

This schedule translates to:

Schedules throttle value of 100 KBps from Saturday to Monday 9 a.m. to 10 p.m.

Schedules throttle value of 200 KBps from Saturday to Monday 10 a.m. to 6 p.m.

[THROTTLE:CheckingInterval]

Value= 900

In the preceding example, the second throttle value (sat-mon#9-18#200) is overlapped by the first throttle value (sat-mon9-10#100) from 9 a.m. to 10 a.m. from Saturday to Monday. So, Open Systems SnapVault uses the lowest of the two bandwidth throttle values from 9 a.m. to 10 a.m. In the example, the lowest throttle value is 100 KBps.

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 26
- ◆ [“Installing Open Systems SnapVault on Windows platforms”](#) on page 32
- ◆ [“Installing Open Systems SnapVault on UNIX and Linux platforms”](#) on page 36
- ◆ [“Verifying the installation”](#) on page 41
- ◆ [“Upgrading to Open Systems SnapVault 3.0”](#) on page 42
- ◆ [“Uninstalling Open Systems SnapVault”](#) on page 46
- ◆ [“Unattended installation and upgrade”](#) on page 48

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. Also, ensure that you read the information in “[Limitations](#)” on page 30 and the latest information in the *Open Systems SnapVault 3.0 Release Notes*.

Requirements for primary storage systems

You can categorize the requirements for the primary storage systems as follows:

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

Type of system: You can install the Open Systems SnapVault 3.0 agent on the following operating systems:

Operating system	Software versions	File systems	ACLs supported
Windows 2003, 2003 R2, 2008 on x86 and x86-64/EM64T compatible hardware	<p>Windows Server 2003 Standard Edition, Windows Server 2003 Standard x64 Edition, Windows Server 2003 Enterprise Edition, and Windows Server 2003 Enterprise x64 Edition.</p> <hr/> <p>Note— “Windows 2003” in this document refers to Windows Server 2003 and Windows Storage Server.</p> <hr/> <p>Windows Server 2008 Standard Edition, Windows Server 2008 Enterprise Edition, and Windows Server 2008 Datacenter Edition.</p>	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 Red Hat Linux 5.4	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 and 11	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
Note _____ Open Systems SnapVault 3.0 does not support VxFS 3.5, and VxFS 4.0 file systems for Solaris.			
Solaris x86	Solaris 10	UFS	Yes
AIX® on PowerPC® and POWER® processor-based systems	5L versions 5.1, 5.2, 5.3, and 6.1 Note _____ Open Systems SnapVault 3.0 does not support IBM® Cluster 1600 and RS/6000 Scalable Parallel (SP) platforms.	JFS1 and JFS2	Yes (AIXC type ACLs on JFS1 and JFS2)
HP-UX® on PA-RISC®	HP-UX 11.23	HFS and JFS	Yes
	HP-UX 11.31	HFS and JFS	Yes
VMware® ESX	ESX 3.0, ESX 3.0.1, ESX 3.0.2 and ESX 3.5	vStorage VMFS and NFS	NA

Note _____
Open Systems SnapVault 3.0 supports 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 3.0 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 requirements:

- ◆ A minimum of 128 MB memory for Windows Server 2003

- ◆ A minimum of 512 MB memory for Windows Server 2008
- ◆ A minimum of 256 MB memory for Linux and ESX server
- ◆ A minimum of 512 MB 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 select another port that uses the Configurator utility during or after installation.

Storage (disk) requirements: You can use the 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 187.

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

- ◆ Open Systems SnapVault requires temporary disk space for normal operations. You require a temporary disk space of 425-MB for every two million files of 20-KB each during baseline transfer. This number increases to 601 MB during an update transfer if you update the same number of files.
- ◆ 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, in addition to 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 7.1 or later.

Supportability matrix of Open Systems SnapVault with NetApp Host Agent

Starting with Open Systems SnapVault 2.6.1, NetApp Host Agent is no longer packaged along with Open Systems SnapVault. You must ensure that the NetApp Host Agent version installed on the primary system is correct. For more information about the NetApp Host Agent versions supported with different platforms, see http://now.netapp.com/NOW/knowledge/docs/olio/guides/ossv/OSSV_Support_matrix.shtml.

Go to the Download Software page of the NOW site at <http://now.netapp.com/NOW/cgi-bin/software> and follow the directions to download the appropriate NetApp Host Agent.

License requirements

The Open Systems SnapVault licenses are installed on the secondary storage system to which they are backed up. You must 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_vi_pri`—For VMware ESX

For softlocks, you require SnapMirror licenses for the secondary and tertiary systems. For more information about softlocks, see “[Setting up a tertiary system for a relationship](#)” on page 162.

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 3.0 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

- ◆ Multiple file systems in a single 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. However, it allows you to back up volume mountpoints on Windows platforms.
- ◆ Resynchronizing restored subdirectories and single files
- ◆ Any new system features in Windows Server 2008 that are not present in Windows Server 2003, except for firewall settings, TCP/IP changes, and bit-locker features.

Installing Open Systems SnapVault on Windows platforms

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 http://now.netapp.com/NOW/cgi-bin/software/ and follow the directions to download the appropriate Open Systems SnapVault package for the platform.
3	Decompress the downloaded package into a temporary directory on the Windows platform.
4	Navigate to the temporary directory where you decompressed the files for the system.
5	Follow the instructions in “ Using the installation wizard ” on page 32 to complete the installation.

Using the installation wizard

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

Step	Action
1	Locate and double-click the Setup.exe file. Result: The Open Systems SnapVault Setup Wizard is launched.
2	Follow the instructions on the screen. Click Next. Result: The Open Systems SnapVault Setup License Agreement window appears.

Step	Action
3	<p>To accept the license agreement, click I Agree, and click Next.</p> <p>Result: The Choose NDMP user name and password window appears.</p> <p>This user name and password are used to communicate with an NDMP-based application, such as you use DataFabric Manager for the central management of Open Systems SnapVault agents.</p>
4	<p>Optional: Perform any of the following actions:</p> <ul style="list-style-type: none"> ◆ Type your user name and password, retype your password to confirm it and click Next. ◆ Directly, click Next. <p>Result: 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"> ◆ Enter the NDMP listening port in the NDMP Listen Port window when prompted. Click Next. ◆ Accept the default listening port setting of 10000, and click Next. <p>Note_____</p> <p>If another application uses 10000 as its listening port, choose an unused port number greater than 10000.</p> <p>_____</p>

Step	Action
6	<p>The Allowed Secondary Names window appears, enter the 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>Note_____</p> <p>If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas(,).</p> <p>_____</p> <p>The primary system accepts the secondary storage systems that are listed in this field as valid backup systems.</p>
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>Note_____</p> <p>If the Open Systems SnapVault agent is uninstalled, files might still be resident in the installation directory; a message appears, asking if you want to empty the directory. To empty the directory, select the Yes option before you click Next.</p> <p>_____</p>

Step	Action	
8	Open Systems SnapVault displays a message.	
	If NetApp Host Agent is...	Then...
	Not installed on the system	<p>Open Systems SnapVault displays a message to install NetApp Host Agent to manage Open Systems SnapVault through DataFabric Manager. The message displays the path to the installation file, which you can use to install the Open Systems SnapVault plug-ins. The path to the installation file is <i>install_dir\manageability\InstallHostAgentPlugins.exe</i>.</p> <p>For more information about installing NetApp Host Agent, see “Supportability matrix of Open Systems SnapVault with NetApp Host Agent” on page 29.</p>
Already installed on the system	Open Systems SnapVault displays a message that the Open Systems SnapVault plug-ins will be copied.	
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. Click “Close” to exit.</p>	
10	<p>After you install the Open Systems SnapVault agent, follow the procedures in the SnapVault chapter of the <i>Data ONTAP Data Protection Guide Online Backup and Recovery 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 26 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 NOW

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

Step	Action
1	Log in to the primary storage system as root.
2	Go to the Download Software page of the NOW site at http://now.netapp.com/NOW/cgi-bin/software/ and follow the directions to download the Open Systems SnapVault Solaris agent package.
3	<p>After you download the Open Systems SnapVault package, unpack and start the installation program:</p> <ul style="list-style-type: none">◆ Uncompress and untar the downloaded pkgadd package.◆ Use the following command to run the Solaris pkgadd utility: <pre>pkgadd -d <i>path_to_package</i>/ossv -a <i>path_to_package</i>/ossv ossv</pre><i>path_to_package</i> is the full path to the Open Systems SnapVault package—for example, /export/home/packages/ossv ossv. <p>The Solaris pkgadd utility installs the package on the system.</p> <p>Result: The package installation script starts and asks you a series of questions before and during the installation process.</p>
4	Go to “ Installing the Solaris agent by using the Solaris pkgadd utility ” on page 37 to continue the Solaris package installation.

Installing the Solaris agent by using the Solaris pkgadd utility

Use the pkgadd utility to install the Open Systems SnapVault Solaris agent:

Step	Action	
<p>1</p>	<p>If you read and agree to the terms of the license, answer yes (y), no (n), or display (d) when asked.</p>	
	<p>If you answer...</p>	<p>Then...</p>
	<p>yes (y)</p>	<p>Go to the next step of installation.</p>
	<p>display (d)</p>	<p>The license is displayed.</p>
	<p>no (n)</p>	<p>You cannot install the Open Systems SnapVault package.</p>
<p>2</p>	<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 get a warning message 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>	
<p>3</p>	<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>	
<p>4</p>	<p>Enter and confirm the password to connect to the system.</p>	
<p>5</p>	<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>Note _____ If another application uses 10000 as its listening port, choose an unused port number greater than 10000.</p> <p>_____</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>Note _____ If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas (.).</p> <hr/> <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> <p>If the installation is 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 NOW

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

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

Step	Action	
3	If...	Then...
	You are installing Open Systems SnapVault on a Linux platform	Enter the following commands after downloading the Open Systems SnapVault package: <pre>gunzip package_name tar -xvf tar_file_name</pre>
	You are installing Open Systems SnapVault on any other UNIX platform	Enter the following commands after downloading the Open Systems SnapVault package: <pre>uncompress package_name tar -xvf tar_file_name</pre> <p>For example, on an AIX primary storage system, uncompress and untar the downloaded ossv_aix_v3.0.tar.Z package using the following commands:</p> <pre>uncompress ossv_aix_v3.0.tar.Z tar -xvf ossv_aix_v3.0.tar</pre>
4	Navigate to the directory where the untar operation placed the files and enter the following command: <pre>./install</pre>	
5	When the install program prompts for the installation directory, press Enter to accept the default, or enter your own path. The default value, /usr, installs the Open Systems SnapVault software in /usr/snapvault.	

Step	Action	
6	If...	Then...
	You plan to manage Open Systems SnapVault backup through the Data ONTAP command-line interface	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>Note _____ If you specify multiple SnapVault secondary storage systems, separate the host names or IP addresses with commas(,), but no spaces.</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 install the Open Systems SnapVault agent, follow the procedures in the SnapVault chapter of the Data ONTAP <i>Data Protection Guide Online Backup and Recovery Guide</i> to configure the SnapVault secondary storage system for Open Systems backup.</p>	

Verifying the installation

Verifying the installation

To verify that the installation is completed correctly and that the primary and secondary storage systems can back up data, complete the following steps:

Step	Action
1	Use the following command to verify connectivity to the secondary storage system: <code>ping secondary_system</code> <i>secondary_system</i> is either the name or the IP address of the secondary storage system.
2	If security is enabled on the primary storage system during installation, ensure that the secondary storage system specified can access the primary storage system, using the procedure described in “ Enabling and disabling security ” on page 66.
3	Create the volumes you need on the secondary storage system before you attempt to back up data to them.
4	If there are firewalls between the primary and the secondary storage system, ensure that the TCP ports are open. For more information, see “ Memory and port requirements: ” on page 28.
5	Navigate to <i>install_dir</i> /bin and run <code>svinstallcheck</code> to verify successful installation and to ensure that the services are running. <i>install_dir</i> is the location where you installed the Open Systems SnapVault agent.

Upgrading to Open Systems SnapVault 3.0

Prerequisites

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

Note

Windows Server 2003 and 2008 on x86-64/EM64T do not support 32-bit Open Systems SnapVault installation.

Upgrading to Open Systems SnapVault 3.0

To upgrade to Open Systems SnapVault 3.0, 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">a. Launch the Open Systems SnapVault Configurator utility. For instructions, see “Running the Configurator utility” on page 64.b. Click the Service tab.c. Click Stop Service to stop Open Systems SnapVault services.d. Close the Configurator utility. <p>Note— Alternatively, you can use either the <code>svpmgr shutdown</code> or <code>snapvault service stop</code> command.</p>
2	<p>Move the database directory to a new location or back it up as discussed in “Backing up and restoring the Open Systems SnapVault database” on page 140 to ensure that during installation the Open Systems SnapVault database is not removed accidentally.</p>

Step	Action
3	<p>Install the Open Systems SnapVault 3.0 agent using the procedure described in “Installing Open Systems SnapVault on Windows platforms” on page 32, or “Installing Open Systems SnapVault on UNIX and Linux platforms” on page 36.</p> <p>Note _____ During installation, when prompted whether you want to upgrade, select Yes (for Windows) or enter Y (for other supported operating systems) to continue with the upgrade.</p> <p>_____</p> <p>Result: The Open Systems SnapVault agent automatically starts Open Systems SnapVault services after installation.</p>
4	<p>a. Launch the Open Systems SnapVault Configurator utility.</p> <p>b. Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services.</p> <p>c. Close the Configurator utility.</p> <p>Note _____ Alternatively, you can use the <code>svpmgr shutdown</code> command.</p> <p>_____</p>
5	<p>Copy the database directory that you saved in Step 2 to the location of the database created by the new installation.</p>
6	<p>a. Launch the Open Systems SnapVault Configurator utility.</p> <p>b. Click the Service tab, and then click Start Service to start Open Systems SnapVault services.</p> <p>Note _____ Alternatively, you can also use the <code>svpmgr start-up</code> command.</p> <p>_____</p>

Support for ACLs: Open Systems SnapVault 2.3 supports ACLs on Linux and HP-UX JFS. For Open Systems SnapVault relationships (that is, for versions earlier than 2.3), after the upgrade to Open Systems SnapVault 3.0, ACLs are backed up only for changed or newly added files. Add the following stanza to `snapvault.cfg`, to avoid a partial backup of ACLs after upgrading to Open Systems SnapVault 3.0:

[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.

Upgrading from Windows Server 2003 with Open Systems SnapVault installed to Windows Server 2008

To upgrade Windows Server 2003 with Open Systems SnapVault installed to Windows Server 2008, complete the following steps:

Step	Action
1	Upgrade the existing Open Systems SnapVault version to Open Systems SnapVault 3.0 on Windows Server 2003.
2	Upgrade the operating system from Windows Server 2003 to Windows Server 2008.

You can use an alternative method to upgrade Windows Server 2003 with Open Systems SnapVault installed to Windows Server 2008.

Step	Action
1	Upgrade the operating system from Windows Server 2003 to Windows Server 2008. Result: Open Systems SnapVault service is in an unknown state.
2	Upgrade the existing Open Systems SnapVault version to Open Systems SnapVault 3.0. Result: Open Systems SnapVault service is working.

Open Systems SnapVault does not support the following after the upgrade from Windows Server 2003 to Windows Server 2008:

- ◆ Update and restore of system drive, for example C:\, which was backed up before the operating system upgrade.
- ◆ Update and restore of system state.

Uninstalling Open Systems SnapVault

Uninstalling the Open Systems SnapVault agent on Windows

To uninstall the Open Systems SnapVault 3.0 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 from the list of programs and click Remove. The Welcome to the OSSV Removal Wizard window appears. Note Open Systems SnapVault does not remove any preinstalled NetApp Host Agent. It restarts the NetApp Host Agent service during uninstallation to remove its plug-ins from NetApp Host Agent.
4	Click Finish to remove Open Systems SnapVault from the computer. Result: Windows uninstalls the Open Systems SnapVault agent and displays the following message to inform you that the uninstallation is successful: OSSV has been successfully removed. Click "Close" to exit.

Uninstalling the Open Systems SnapVault agent on Solaris

To uninstall the Open Systems SnapVault 3.0 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: <code>pkgrm ossv</code>

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 <ossv></pre> <p>Execute the scripts in this package contains scripts with superuser 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>Result: If removal is successful, the script responds with text similar to the following:</p> <pre>Removal of <ossv> was successful.</pre>

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

To uninstall the Open Systems SnapVault 3.0 agent on the HP-UX, AIX, or Linux platform, complete the following steps:

Step	Action
1	<p>Enter the following command:</p> <pre>install_dir/util/uninstall</pre>
2	<p>If the SnapVault directory still appears after running <code>uninstall</code>, remove the directory manually.</p>

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 noninteractively, and usually, you do not need to reboot the system after the installation or upgrade has been 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 54.

Supported platforms and Open Systems SnapVault versions

You can upgrade or install all Open Systems SnapVault supported platforms using this method. See [“Requirements for primary storage systems”](#) on page 26 for information about 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, you require an installation script and other supporting files. 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, along 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 different platform. Similarly, an installation script created for a Solaris system cannot be used to perform an unattended installation on an HP-UX system.

- ◆ You cannot change the following configuration settings 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 55. 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"> a. Launch the Open Systems SnapVault Configurator utility. For instructions, see “Running the Configurator utility” on page 64. b. Click the Service tab, and then click Stop Service to stop Open Systems SnapVault services.
2	Configure all parameters you want the new installation to have, using the Configurator utility.
3	Close the Configurator utility.
4	Navigate to the <i>install_dir/bin</i> directory of the primary storage system.

Step	Action	
5	If...	Then...
	<p>You only want to save the configuration settings to a file</p>	<p>Enter the following command:</p> <pre>svconfigpackager filename</pre> <p><i>filename</i> is the name of the configuration settings file.</p> <p>Example: 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>Note _____</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) are located at the install root directory for Open Systems SnapVault.</p> <p>_____</p>	<p>Enter the following command:</p> <pre>svconfigpackager [-h -i installation path] package name</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>Example: To create a configuration settings file called <code>svconfig.in</code> and save the installation script in the <code>/usr/snapvault</code> directory, enter the following command:</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>svconfigpackager -h -i path_name filename</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>Example 1: 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>Example 2: 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>Result: 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"> ◆ Download and uncompress the Open Systems SnapVault package from the NOW site. ◆ Copy the Open Systems SnapVault package to the installation directory on the primary storage system where you want to install or upgrade.

Step	Action	
2	Copy the installation script and other generated files in Step 6 of “ Performing an unattended installation or upgrade ” on page 51 to the directory on the primary storage system where you decompressed the installation package.	
3	If...	Then...
	The primary storage system is a Windows machine	<p>Enter the following command to start the unattended installation or upgrade:</p> <pre>unattinstall.bat</pre> <p>If the installation does not succeed, you can find the error messages at the following location:</p> <p>For Windows 2003 systems— %SystemRoot%\Documents and Settings\<i>Current User</i>\Local Settings\Temp</p>
	The primary storage system is a UNIX machine	<p>Enter the following command to start the unattended installation or upgrade:</p> <pre>./unattinstall.sh</pre> <p>If the installation does not succeed, you can find the error messages in the /tmp directory.</p>
4	On Windows, run the svinstallcheck utility to verify the successful installation and to ensure that the services are running. On UNIX, svinstallcheck runs automatically after the installation. If you find any errors see the SnapVault log files in the <i>install_dir/etc</i> directory.	

Performing an unattended installation or upgrade of NetApp Host Agent on Windows

To perform an unattended installation or upgrade of NetApp Host Agent with Open Systems SnapVault on Windows, complete the following steps:

Step	Action
1	Run the svconfigpackager utility to generate the unattended install package for Windows.
2	Copy the generated files to the unzipped Open Systems SnapVault package directory.
3	Download the NetApp Host Agent installer (a single program file) from http://now.netapp.com/NOW/cgi-bin/software/ , to the Open Systems SnapVault package directory.
4	<p>Edit the unattinstall.bat file to add a new command so that it precedes the OSSV install command. This command can vary for NetApp Host Agent installation or upgrade.</p> <ul style="list-style-type: none"> ◆ To install and upgrade NetApp Host Agent, add the following command: <pre>start /wait agentsetup-<NHA version>-win32.exe /s /v/qn</pre> ◆ Add the following command to change the installation path: <pre>start /wait agentsetup-<NHA version>-win32.exe /s /v"/qn INSTALLDIR="<the actual path>"</pre>
5	<p>The new script looks like:</p> <pre>start /wait agentsetup-<NHA version>-win32.exe /s /v/qn msiexec /i ossv.msi /qn targetdir="c:\Program Files\netapp\snapvault" db_dir="c:\Program Files\netapp\snapvault\db" trace_dir="c:\Program Files\netapp\snapvault\trace" tmp_dir="c:\Program Files\netapp\snapvault\tmp" reboot=ReallySuppress UNATTENDED_INSTALL=1 HONOR_EXISTING_CONFIG=1 CONFIG_FILE=ossv</pre> <p>Result: This procedure ensures that NetApp Host Agent is installed before Open Systems SnapVault is installed or upgraded. The Open Systems SnapVault plug-ins are also installed with NetApp Host Agent.</p>

Performing an unattended installation or upgrade of NetApp Host Agent on Linux

To perform an unattended installation or upgrade of NetApp Host Agent with Open Systems SnapVault on Linux, complete the following steps:

Step	Action
1	Run the <code>svconfigpackager</code> utility to generate the unattended install package for Linux.
2	Copy the generated files to the Open Systems SnapVault package directory.
3	Download the NetApp Host Agent installer (<code>agentsetup-<NHA version>-linux.bin</code>) from http://now.netapp.com/NOW/cgi-bin/software/ , to the Open Systems SnapVault package directory.
4	Modify the <code>unattinstall.sh</code> script to include the following line: <code>./agentsetup-<NHA version>-linux.bin</code>
5	The new unattended install script is similar to the following: <code>./agentsetup-<NHA version>-linux.bin</code> <code>./install.sh InstallResponseFile</code>

Remote batch installation of the Open Systems SnapVault agent on Windows

NetApp 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. Reboot those clients to install the Open Systems SnapVault agent.

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)?

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 command-line interface 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 59
- ◆ [“Running the Configurator utility”](#) on page 64
- ◆ [“Confirming that services are running”](#) on page 65
- ◆ [“Modifying Open Systems SnapVault parameters”](#) on page 66
- ◆ [“Enabling and disabling debugging”](#) on page 70
- ◆ [“Setting block-level incremental backup options”](#) on page 73
- ◆ [“Configuring backup exclusion lists”](#) on page 75
- ◆ [“Configuring open file backup for Windows”](#) on page 78
- ◆ [“Configuration for preserving Snapshot copies”](#) on page 80
- ◆ [“Configuration for DataFabric Manager restore to non-ASCII path”](#) on page 81
- ◆ [“Primary storage system reporting through AutoSupport”](#) on page 82

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

This section covers the following topics:

- ◆ [“Understanding the Configurator utility interface”](#) on page 59
- ◆ [“Understanding the svsetstanza command”](#) on page 61

Understanding the Configurator utility interface

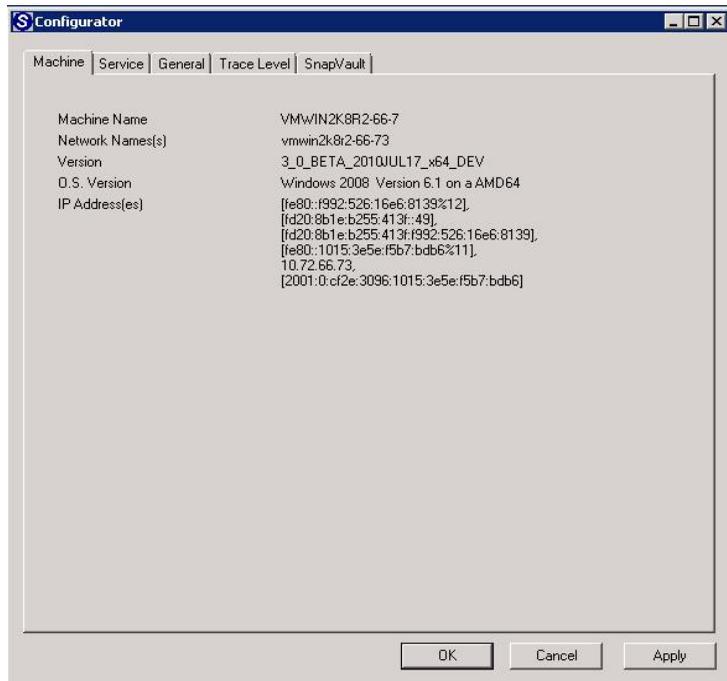
About the Configurator utility interface

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 Open Systems SnapVault software version and the primary system information, such as the IPv4 and IPv6 network addresses and operating system 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. You can find the SnapVault log files in the *install_dir/etc* directory. For more information about the Open Systems SnapVault log files, see “[Locating status and problem reports](#)” on page 138.

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), parameters, NDMP parameters (for central management of Open Systems SnapVault agents), VSS parameters, and security settings.

List of configuration files

The following configuration files include the parameters that the different tabs of the Configurator utility interface can configure or change. You can find these configuration files 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

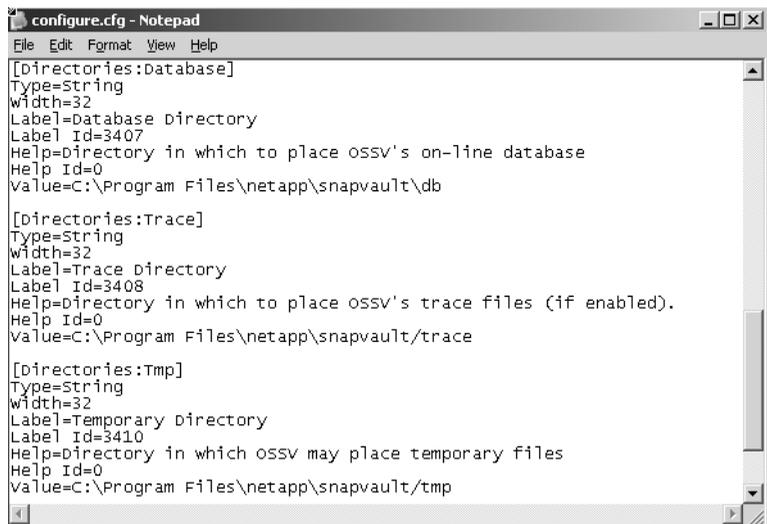
About the svsetstanza command

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. You can find the svsetstanza command in the *install_dir/snapvault/util* directory.

What the svsetstanza command changes

You can find the parameters that the svsetstanza command can configure or change in the *install_dir/snapvault/config* directory. For more information about the configuration files, see [“List of configuration files”](#) on page 60.

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. Usually, the directory is *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, 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 specified values. It 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">◆ Click Start > Programs > OSSV > OSSV Configurator on the Windows primary storage system.◆ Run the following command on the UNIX primary storage system: <code>\$INSTALL_DIR/bin/svconfigurator</code> The default location for <code>INSTALL_DIR</code> is <code>/usr/snapvault</code>.
2	<p>Click the appropriate Configurator utility tab to change the settings.</p> <p>Click the SnapVault tab to access the most commonly modified Open Systems SnapVault parameters.</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”](#)
- ◆ [“Modifying the qtree SnapMirror access list”](#)
- ◆ [“Modifying the NDMP settings”](#)
- ◆ [“Enabling and disabling Windows EventLog”](#)

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 the 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>Note</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.

You can modify the following settings (but you should not change the default values unless technical support asks you to change them):

- ◆ 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 the Open Systems SnapVault agent fills automatically. This identifier is used by NDMP-based management

applications to identify the primary storage system. You must *not* modify this field.

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

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

Modifying 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.

Step	Action
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. Note _____ 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.

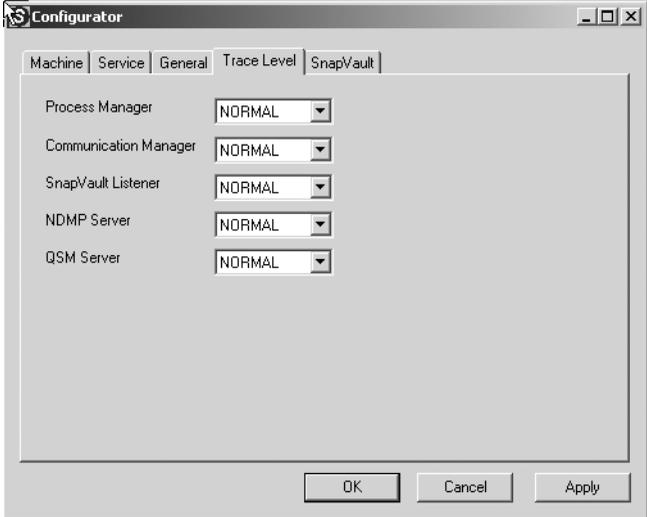
Step	Action
1	In the Configurator GUI, click the General tab.
2	Select the “Generate debugging files” check box. Note _____ To disable the generation of debug files after you are done troubleshooting, clear the “Generate debugging files” check box. _____
3	Click the Trace Level tab.
4	From the drop-down list of the service for which you want to generate debug information, select one of the five trace levels—ALWAYS, NORMAL, VERBOSE, LIBNORMAL, or LIBVERBOSE. Note _____ Ensure that you set the trace level back to NORMAL after you are done troubleshooting. _____
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>Note _____ Instead of performing these steps through the Configurator utility, you can use the <code>svpmgr shutdown</code> and <code>svpmgr startup</code> commands on the command-line interface. _____</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.

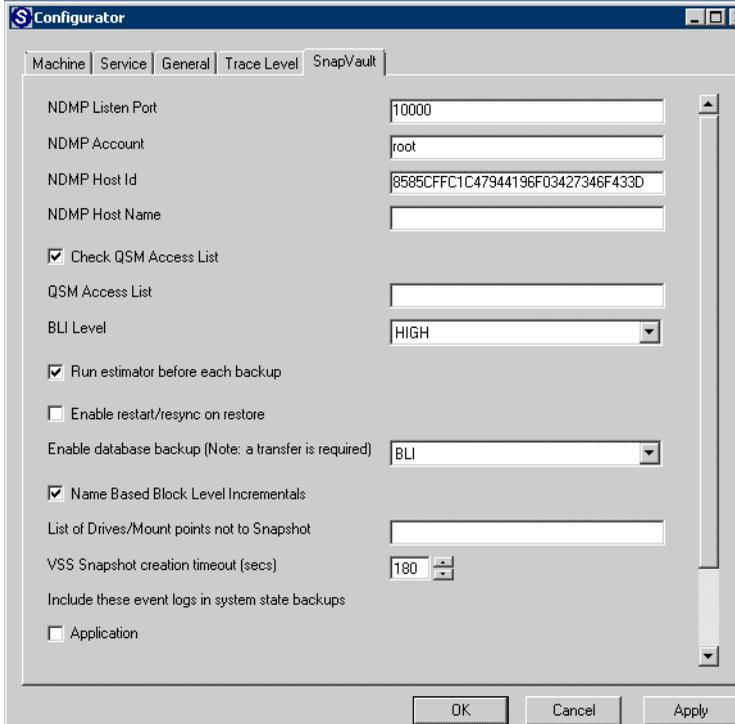
Step	Action
4	<p>Select NORMAL from the drop-down list of the service for which you want to stop generating debug information.</p> <p>For example, for qtree SnapMirror Server, select NORMAL, as shown in the following image.</p> 
5	Click Apply.
6	Click the Service tab, click Stop Service, click Start 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 BLI Level drop-down list the block-level incremental checksum computation options.</p>  <p>The screenshot shows the 'Configurator' window with the 'SnapVault' tab selected. The 'BLI Level' dropdown menu is open, showing 'HIGH' as the selected option. Other visible settings include NDMP Listen Port (10000), NDMP Account (root), NDMP Host Id (8585CFFC1C47944196F03427346F433D), and 'Run estimator before each backup' checked.</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"> ◆ HIGH—computes checksums during the initial baseline backup and updates <p>Note _____ The HIGH mode might slow down the initial backup process.</p> <hr/> <ul style="list-style-type: none"> ◆ LOW—computes checksums only during updates ◆ OFF—disables block-level incremental backups altogether

Enabling or disabling BLI backups for certain name-based applications

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

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

Configuring backup exclusion lists

What backup exclusion lists are

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

- ◆ File exclusion lists
- ◆ File system 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*.

File system exclusion lists

File system exclusion list entries consist of file system types. A file system is excluded if the file system type matches the file system exclusion entry.

Note

Windows and AIX platforms do not support File system exclusion.

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

You can specify one file system exclusion entry per line of the file. The exclusion entry should not have any wildcards or delimiters. It should be a complete file system type name, for example, ext2, ext3, and so on.

Note

The root (/) is always included even if the root file system is excluded.

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.
!	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 file system exclusion list:

```
guest192-224:/usr/snapvault/etc # vi file-system-exclude.txt
# file-system-exclude.txt file for SnapVault Linux file system
exclusions
# Exclusion file system entries must be in ascii format
# Exclusion file system entries are case sensitive
# Specify one exclusion expression only per line of file
# Wildcards are not supported
# All spaces in an exclusion entry are significant
xfs
jfs
```

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 the VSS Snapshot copy feature on Windows 2003 and later. Open Systems SnapVault agent includes the VSS Snapshot copy functionality as a standard feature and does not require a license.

About configuring VSS settings

The VSS settings need to be changed only for troubleshooting. Therefore, do not change these settings unless directed by technical support.

Backing up open files

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

The system must meet certain conditions 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 platforms, complete the following steps.

Step	Action
1	Click the SnapVault tab.

Step	Action
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 on the Open Systems SnapVault primary storage system 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 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 end the transfer.

- ❖ When the value is set to **TRUE**, the transfer aborts.
- ❖ When the value is set to **FALSE**, the transfer continues with a new Snapshot.

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 to restore backed-up data to a non-ASCII path.

For Open Systems SnapVault 2.6.1, the default value is `TRUE`.

Note

The value is `FALSE` for releases prior to Open Systems SnapVault 2.6.1. If you have upgraded to Open Systems SnapVault 2.6.1, you have to set the value to `TRUE`.

Primary storage system reporting through AutoSupport

Starting with Open Systems SnapVault 3.0, the primary storage system sends system information to the secondary storage system during backup. If AutoSupport feature is enabled on the secondary storage system, it sends the information to NetApp technical support on a weekly basis. By default, the option is enabled on your primary storage system. You can disable or enable this option from the *snapvault.cfg* file. To disable this option, open the *snapvault.cfg* file and change the *EnableASUP* flag value to FALSE.

Open Systems SnapVault primary system sends the following information through AutoSupport messages to technical support:

- ◆ The operating system type and version of the primary storage system on which Open Systems SnapVault is installed
- ◆ File system type when backing up file system path
- ◆ Open Systems SnapVault version
- ◆ Total storage present on the primary system
- ◆ Application-specific details:
 - ❖ Application name and version
 - ❖ Total application data stored on the primary system
 - ❖ Total number of application instances and total databases

About this chapter This chapter describes how to configure Open Systems SnapVault on a two-node (active/active and active/passive) Microsoft Cluster Services (MSCS) and the procedures to configure cluster services Support using Protection Manager.

Topics in this chapter

This chapter covers the following topics:

- ◆ [“Microsoft Cluster Services Support on Open Systems SnapVault”](#) on page 84
- ◆ [“How Open Systems SnapVault works in an MSCS environment”](#) on page 85
- ◆ [“Migrating a stand-alone Windows node with Open Systems SnapVault to a cluster”](#) on page 86
- ◆ [“Enabling cluster support”](#) on page 86
- ◆ [“Disabling cluster support”](#) on page 87
- ◆ [“Setting up and configuring Open Systems SnapVault on a two-node cluster”](#) on page 88
- ◆ [“Setting up and configuring a two-node cluster”](#) on page 88
- ◆ [“Protection Manager support for Microsoft Cluster”](#) on page 90
- ◆ [“Uninstalling Open Systems SnapVault in an MSCS environment”](#) on page 93

Microsoft Cluster Services Support on Open Systems SnapVault

Overview

Starting with Open Systems SnapVault 3.0, you can back up and restore local and cluster file system and application data from the nodes in a Microsoft Cluster Services (MSCS) and Windows Server 2008 Failover Clustering solution. You can install Open Systems SnapVault on two-node cluster configurations in active/active or active/passive mode. Open Systems SnapVault supports all its features on MSCS.

During cluster failover and failback or administrative disruption, you do not have to run the baseline transfer again.

A new utility called **svcluster** is available to configure Open Systems SnapVault on MSCS. You should run svcluster utility on both the nodes to enable cluster support on Open Systems SnapVault.

You can use DataFabric Manager's Protection Manager for backup and restore purposes.

The following Windows platforms supported on both on 32-bit and 64-bit OS type:

- ❖ Windows Server 2003
- ❖ Windows Server 2008
- ❖ Windows Server 2003 R2
- ❖ Windows Server 2008 R2

Open Systems SnapVault supports Microsoft Cluster Services on virtual machines and functions in the same way that it does on physical systems.

Open Systems SnapVault 3.0 supports the virtual machines of the following virtualization software applications for MSCS:

- ❖ VMware® ESX 3.5
- ❖ VMware ESX 4.0
- ❖ Hyper-V™ virtual machine

How Open Systems SnapVault works in an MSCS environment

To provide Open Systems SnapVault support in a Microsoft Cluster environment, the Open Systems SnapVault database is shared for each backup relationship with the node, and a separate custom cluster resource is created for gracefully aborting transfers during failover.

Distributed Open Systems SnapVault database: The Open Systems SnapVault database resides in the `\ossvdb` directory at the root of the volume for all the relationships configured on that volume. For example, if `F:\data1` and `F:\data2` are part of backup in two separate relationships, then Open Systems SnapVault creates a database directory `F:\ossvdb`. The information about both the backup relationships is stored in the `F:\ossvdb` directory. Therefore, if the volume `F:` is moved within the cluster, the relevant database information that is stored in the `ossvdb` is also moved with the volume and enables Open Systems SnapVault export the configurable database path. However, the configured database path stores only information about System State relationships. For all other relationships (including relationships from local drives), the database is stored on the drive from which data is being backed up.

Custom cluster resource: You should run `svcluster` utility on both the nodes to enable cluster support on Open Systems SnapVault. After enabling cluster support, a new custom cluster resource called *OSSVResourceType* is created. The `OSSVResourceType` resource is used to gracefully abort all transfers that are in progress during the cluster failover process. It sends a notification to Open Systems SnapVault when the disk resources are moved and restarts the Open Systems SnapVault services after aborting all the backup transfers that are in progress.

The "`OSSVResourceType`" also ensures that the Open Systems SnapVault database is in a consistent state before the disk resources are unmounted.

Open Systems SnapVault configuration in a cluster: You should consider the following points when configuring Open Systems SnapVault in a cluster:

- ◆ Both nodes in a cluster must have the same configuration when configuring Open Systems SnapVault.
- ◆ The disk resources are mounted with the same drive letter on both the nodes of a cluster.
- ◆ The mountpoints are at the same location on both the cluster nodes. The mountpoint location for clustered volume mount should also be a cluster resource so that when the volume mountpoint is moved, its mountpoint location is moved with it.

Migrating a stand-alone Windows node with Open Systems SnapVault to a cluster

You can migrate a stand-alone Windows node that has Open Systems SnapVault 3.0 installed on it to a cluster.

After the migration, you can use Open Systems SnapVault in two different scenarios.

- ◆ You can continue to use the node for backup and recovery purposes without making cluster-related changes in Open Systems SnapVault. Although the node is part of a cluster, Open Systems SnapVault backs up the node as if it is a standalone node.
- ◆ You can make cluster-related changes in Open Systems SnapVault, but you must rebaseline all backup relationships, except for the System state.

Enabling cluster support

Open Systems SnapVault provides a utility called `svcluster`. This utility enables you to configure Open Systems SnapVault for MSCS support. You can enable or disable the cluster support by using this utility. After enabling cluster support, the Open Systems SnapVault services restart and create the `OSSVResourceType`. When you disable cluster support on Open Systems SnapVault restarts the Open Systems SnapVault services and removes the `OSSVResourceType`.

This utility is installed in the `cluster/mscs` sub-directory during the Open Systems SnapVault 3.0 installation for Windows.

Note

You must enable the cluster support on both nodes of a cluster.

To enable cluster support, complete the following steps:

Step	Action
1	In the primary system management console, navigate to the <code>mscs</code> directory. <code>C:\Program Files\NetApp\snapvault\cluster\mscs></code>
2	In the console, enter the <code>svcluster enable</code> command. <code>C:\Program Files\NetApp\snapvault\cluster\mscs>svcluster enable</code>
3	Repeat the steps 1 and 2 on the other node of the cluster. You can start using Open Systems SnapVault on MSCS.

Disabling cluster support

To disable cluster support on a node in a cluster environment, complete the following steps:

Note

You should remove the dependency of the OSSVResourceType on all the disk resources before disabling cluster support.

Step	Action
1	In the primary system management console, navigate to the mscs directory. <code>C:\Program Files\NetApp\snapvault\cluster\mscs></code>
2	Enter the <code>svcluster disable</code> command. <code>C:\Program Files\NetApp\snapvault\cluster\mscs>svcluster disable</code>
3	Repeat steps 1 and 2 on the other node of the cluster.

Microsoft Cluster Services Support on Open Systems SnapVault
Setting up and configuring a two-node cluster

Setting up and configuring Open Systems SnapVault on a two-node cluster

To set up and configure Open Systems SnapVault on a Microsoft two-node cluster, complete the following steps:

Note _____
 To perform this procedure, you must have Cluster Administrator privileges.

Step	Action
1	Set up and configure Microsoft cluster. For instructions about setting up and configuring MSCS, see the Microsoft documentation.
2	Configure the cluster resource groups. A resource group should contain disk resources for backup purpose.
3	Install Open Systems SnapVault on both the nodes. For the installation instructions, see “ Installing the Open Systems SnapVault Software ”. Note _____ The install path for Open Systems SnapVault must be the same for both the nodes.
4	Configure Open Systems SnapVault as per your requirements.
5	Run <i>svcluster</i> utility to enable the cluster support. For enabling the cluster support, see “ Enabling cluster support ”.
6	From the Cluster Administrator user interface, for each cluster resource group add a new resource type of "OSSVResourceType". <ul style="list-style-type: none"> ❖ For file-systems backup, make the resource dependent on all the disk resources. ❖ For Microsoft SQL backups, make the resource dependent on the SQL Agent resource.

**Backing up using
command-line
interface in an
MSCS environment**

If you are backing up data using the command-line interface, then the cluster resource should be backed up using the cluster IP address and the local drive should be backed up using the local IP address.

For more information about backup and restore, see “[Perform Backup and Restore](#)” on page 95.

Microsoft Cluster Services Support on Open Systems SnapVault Protection Manager support for Microsoft Cluster

Configuring Open Systems SnapVault in Protection Manager for a Microsoft Cluster environment

You can configure Open Systems SnapVault in Protection Manager for MSCS environments in a two-node active/active or active/passive cluster.

An Open Systems SnapVault client in a Microsoft Cluster environment has multiple IP addresses. You can add these multiple IP addresses as hosts in Protection Manager and use them for backup and restore purposes. For example, if you have a two-node active/active Microsoft SQL Server cluster. Each of the cluster nodes have their physical IP addresses in addition to a virtual IP address for each of the resource groups used for failover. In the case where there are two resources groups between the two nodes, there are a total of four IP addresses that need to be added to Protection Manager. The physical IP addresses are used to protect local data such as data in the C:\ and the System State, whereas the virtual IP addresses are used to protect SQL Server or other file system data owned by the resource group.

Prerequisites:

- ◆ You should install Open Systems SnapVault on each node and enable clustering support before performing this task. For more information, see [“Setting up and configuring Open Systems SnapVault on a two-node cluster”](#) on page 88.
- ◆ You should disable Host Agent Discovery option from Operations Manager user interface. You can also disable from command-line utility using this command `dfm option set discoverAgents=no`.
- ◆ You should disable NDMP Host Discovery option from Operations Manager user interface. You can also disable from the command-line utility using this command `dfbm option set discoverNdmp=no`.

To configure Open Systems SnapVault in Protection Manager in an active/passive or active/active two-node cluster environment, complete the following steps:

Step	Action
1	Log in to Protection Manager using NetApp Management Console.
2	From the navigation pane, click Hosts > OSSV .

Step	Action
3	Click Add to start the Add OSSV Host wizard. Enter the host name or IP address of the first node in the cluster (the physical host) and complete the wizard.
4	Click Add to start the Add OSSV Host wizard. Enter the host name or IP address of the second node in the cluster (the physical host) and complete the wizard. Verify that the both the nodes have been included in the hosts list in the OSSV Hosts window. You might need to refresh the window before you can view the new hosts in the host list.
5	Access the OSSV Configurator on the first node and click the SnapVault tab.
6	Change the NDMP Host Name to match the host name of the virtual interface that you want to add as an Open Systems SnapVault host in Protection Manager.
7	Change the last four characters of the NDMP Host ID. NDMP Host IDs must be unique. Duplicate NDMP Host IDs prevents Protection Manager from adding the host entries.
8	Click Apply .
9	Access Protection Manager and add the virtual host name as an Open Systems SnapVault host.
10	Repeat Step 5 through Step 9 to add remaining virtual host names of both the cluster nodes.

For more information about Protection Manager, see the *Provisioning Manager and Protection Manager Printable Help and NetApp Management Console Online Help*.

Postrequisites:

- ◆ Enable Host Agent Discovery option from the Operations Manager user interface. You can also enable from command-line utility using this command `dfm option set discoverAgents=yes`

- ◆ Enable NDMP Host Discovery option from Operations Manager. You can also enable from command-line utility using this command `dfbm option set discoverNdmp=yes`.

Uninstalling Open Systems SnapVault in an MSCS environment

To uninstall Open Systems SnapVault in an MSCS environment, complete the following steps:

Step	Action
1	In the Cluster Administrator user interface, delete all OSSVResourceType resources from all the cluster resource groups.
2	Disable cluster support on cluster nodes. For more information, see “Disabling cluster support” on page 87.
3	Uninstall Open Systems SnapVault from both the nodes. For more information, see “Uninstalling the Open Systems SnapVault agent on Windows” on page 46.
4	Delete the Open Systems SnapVault database directories (<i>ossvdb</i>) from all local and cluster disk drives. Note _____ The ossvdb directory resides in the root of a disk drive. _____

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 96
- ◆ [“Perform SnapVault restore on Open Systems platform”](#) on page 102

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 that you have completed the following steps in the following order:

Note

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

1. On Open Systems platforms, install the Open Systems SnapVault agent and configure it for backups by the desired SnapVault secondary storage system. For more information, see Chapter 2, “[Installing the Open Systems SnapVault Software](#),” on page 25.
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 97.
3. On the SnapVault secondary storage system, start the baseline transfer. See “[Creating an initial baseline copy](#)” on page 99.
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 100.

Topics in this section

This section covers the following topics:

- ◆ “[Configuring the SnapVault secondary storage system](#)” on page 97
- ◆ “[Creating an initial baseline copy](#)” on page 99
- ◆ “[Scheduling SnapVault update backups](#)” on page 100
- ◆ “[Backing up empty source directories](#)” on page 101

Configuring the SnapVault secondary storage system

SnapVault secondary storage system requirement

The SnapVault secondary storage system must be running Data ONTAP 7.1 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 7.1.2, 7.2, or later.

Note

If your secondary storage system is part of a Data ONTAP active/active configuration or an HA pair, then you must add the SnapVault secondary license on the both nodes.

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>Example:</p> <pre>license add sv_ontap_sec</pre> <p>Note</p> <p>Ensure that you select the correct storage system or NearStore® secondary license.</p>

Step	Description
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>Example 1 (Windows):</p> <pre>license add sv_windows_pri</pre> <p>Example 2 (UNIX):</p> <pre>license add sv_unix_pri</pre>
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 <i>Data ONTAP 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>Example:</p> <pre>options snapvault.access host=melzhost,samzhost,budzhost</pre> <p>For more information, see the section on enabling SnapVault in the <i>Data ONTAP 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">◆ 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.◆ 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. <p>Enter the following command:</p> <pre>snapvault start -S prim_host:dirpath sec_host:/vol/sec_vol/sec_tree</pre> <p>Example 1 (Windows):</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>Example 2 (UNIX):</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>

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 <code>g</code> 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>Example:</p> <pre>snapvault snap sched -x voll sv_weekly 1@sat@19 snapvault snap sched -x voll sv_nightly 2@mon-fri@19 snapvault snap sched -x voll sv_hourly 11@mon-fri@7-18</pre> <p>Note</p> <p>You must use the <code>-x</code> parameter in the preceding 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 updating all the secondary qtrees on the specified volume, SnapVault creates a Snapshot copy of this volume for archival.</p>

Backing up empty source directories

Prior to Open Systems SnapVault 2.6.1, 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: [QSM:BackupEmptyMount] value=TRUE
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. You can also use Protection Manager for restoring the data.

Topics in this section

This section covers the following topics:

- ◆ [“Restoring a directory or a file”](#) on page 103
- ◆ [“Restoring an entire primary storage system”](#) on page 107
- ◆ [“Restoring files to a primary storage system from tape”](#) on page 108

Restoring a directory or a file

Methods for restoring a directory or a file

If there is a 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, you can perform restore operations without requiring a backup administrator’s assistance.

Note

This method do not preserve some Windows and UNIX attributes—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 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"> ◆ The secondary storage system host and the path to the secondary qtree, directory, and file that you want to restore from. ◆ 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>). ◆ 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. ◆ The path on the primary storage system to the directory or file that you want to restore to. <p data-bbox="494 644 1126 673">Example 1—Single file restore (Windows system):</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">Example 2—Single file restore (UNIX system):</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">Note</p> <hr data-bbox="494 946 1224 949"/> <p data-bbox="494 963 1224 1085">Enter the entire command as a single line. Ensure that you do not specify a slash (<code>\</code> or <code>/</code>) 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 1097 1224 1100"/> <p data-bbox="494 1133 1204 1163">Example 1—Single directory restore (Windows system):</p> <pre data-bbox="494 1168 1220 1220">snapvault restore -s sv_daily.0 -k 10 -S myvault:/vol/sv_vol/melzdir/dir1 a:\melzdir\dir1_restore</pre> <p data-bbox="494 1246 1150 1275">Example 2—Single directory restore (UNIX system):</p> <pre data-bbox="494 1281 1013 1362">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 format of the file system.
3	Install the Open Systems SnapVault agent. See Chapter 2, “ Installing the Open Systems SnapVault Software ,” on page 25.
4	Optional: If you backed up the Windows System State data of the primary storage system, restore its Windows System State data. For more information, see “ Restoring System State data ” on page 150.
5	Restore the backed-up directories using the <code>snapvault restore</code> command. For details, see “ Using the snapvault restore command ” on page 104.

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

The following method does not preserve some Windows NT and UNIX attributes, notably Windows NT sparse files, Windows NT EFS data, and UNIX ACLs.

Step	Action
1	Mount the tape that has the restored files.
2	Use the <code>restore</code> command to restore from the tape to the SnapVault secondary storage system. For details, see the <i>Tape Backup and Recovery Guide</i> .
3	Copy the files from the SnapVault secondary storage system to the primary storage system using NFS or CIFS. For details, see “Restoring a directory or a file” on page 103.

Volume mountpoint data backup and restore

Overview

Starting with Open Systems SnapVault 3.0, you can back up and restore NTFS file system mounted folders. A volume mountpoint is an association between a volume and a directory on another volume.

The following Windows platforms support mountpoint backup and restore operations:

- ❖ Windows Server 2003
- ❖ Windows Server 2008
- ❖ Windows Server 2003 R2
- ❖ Windows Server 2008 R2

Open Systems SnapVault 3.0 enables you to back up data inside a volume mountpoint.

You can restore data to a volume mountpoint or to a folder inside a volume mountpoint.

During system state backup, if an NTFS volume is mounted as a drive letter and also as multiple volume mountpoint, then the disk quota for that volume is backed up multiple times.

When you run the `svinstallcheck.exe` utility, it displays all the mountpoints along with the drive letters that are appropriate for backup. The `-f` option of the `svinstallcheck.exe` scans and displays all the unsupported reparse points.

Protection Manager for backup and restore

You can use Protection Manager for backup and restore of the mountpoints. Protection Manager lists all the mountpoints.

The Open Systems SnapVault collects all mountpoints and drives associated with a volume and makes a list, which is then sent to Protection Manager. For example, if a volume is mounted as `F:\` and also as `C:\mnt`, then Open Systems SnapVault sends the `F:\` and `C:\mnt` to Protection Manager.

Backing up and restoring data in a volume mountpoint

You can back up a complete volume mountpoint or you can back up a folder inside a volume mountpoint.

Example: Backing up a volume mountpoint `C:\mnt` mountpoint.

```
snapvault start -s windows-machine:c:\mnt /vol/vol0/mnt_backup
```

Example: Backing up a folder inside a mountpoint. C:\mnt is a mountpoint and *data* is a folder.

```
snapvault start -S windows-machine:c:\mnt\data  
/vol/vol1/mnt_data_backup
```

You can restore data to a mountpoint or to a folder inside the mountpoint.

During restore, Open Systems SnapVault validates the mountpoint and restores the mountpoint. The restore path should be a valid mountpoint, it should not contain any reparse point, and the mountpoint should be an NTFS file system. If the validation fails, the restore operation is aborted.

Example: Restoring a folder to into a mountpoint.

```
snapvault restore -S storage system :/vol/vol1/mnt_backup/data  
c:\mnt\data
```

Example: Restoring the volume and the volume mountpoint. First restore the volume, and then restore the mountpoint.

```
snapvault restore -S filer:/vol/vol1/c_drive c:\  
snapvault restore -S filer:/vol/vol1/c_mnt_mount_point c:\mnt
```

You should create a mountpoint manually before restoring the data to a mountpoint.

When you want to restore both the volume mountpoint and the volume in which the mountpoint exists, you must restore the volume first followed by the volume mountpoint. For example, if C:\ is a volume and C:\mnt is the mountpoint, you must restore C:\ first and then c:\mnt.

For more information about a backup and restore process, see [“Perform Backup and Restore”](#) on page 95.

About this chapter This chapter describes how to back up and restore Microsoft SQL Server databases using Open Systems SnapVault and Protection Manager. It describes the configuration procedures for backing up various SQL Server database files.

Topics in this chapter This chapter covers the following topics:

- ◆ “[Overview](#)” on page 111
- ◆ “[How Open Systems SnapVault backs up an SQL server database](#)” on page 112
- ◆ “[How Open Systems SnapVault restores SQL server database](#)” on page 115
- ◆ “[Configuring backup and restore of Microsoft SQL databases](#)” on page 117
- ◆ “[Backing up using the command-line interface](#)” on page 129
- ◆ “[Restoring using the command-line interface](#)” on page 131
- ◆ “[Backing up Microsoft SQL Server database using Protection Manager](#)” on page 132
- ◆ “[Restoring Microsoft SQL Server database using Protection Manager](#)” on page 135

Overview Starting with Open Systems SnapVault 3.0, you can back up and restore the Microsoft SQL Server™ application database. The support of Microsoft SQL Server enables backup administrators to perform scheduled backups of a specific SQL database and enables the administrator to restore and recover any specific database when required.

The following Microsoft SQL Server versions are supported:

- ◆ Microsoft SQL Server version 2005
- ◆ Microsoft SQL Server version 2008

You can use Protection Manager 3.8 and if you are not using Protection Manager for back up and restore, the Data ONTAP command-line interface on the secondary storage system and Open Systems SnapVault command-line interface on the primary system can be used for the backup and restore of Microsoft SQL Server database. The *svapp* command-line utility enables you to view the list of databases on the primary system.

A new plug-in called `mssql` enables you to back up and restore the Microsoft SQL Server database. The plug-in is a DLL file (`ossv_mssql.dll`) installed along with Open Systems SnapVault 3.0 on all supported Windows platforms in this path `<install path>\apps\mssql`. The plug-in is also installed during an unattended installation of Open Systems SnapVault 3.0.

You can back up and restore the full database and transaction logs from a single node or from either of the two nodes in a two-node Microsoft Cluster. To back up from a two-node Microsoft Cluster, both the nodes must be running Open Systems SnapVault 3.0 version or later.

The `svinstallcheck` command lists the Microsoft SQL Server paths that are suitable for backup and the `svestimator` tool estimates the space required for Microsoft SQL Server database backup.

Configuration files: The two configuration files `ossv-mssql.cfg` and `ossv-mssql-local-Tlog-DBs.cfg` enable you to set the behavior of the Microsoft SQL Server database per your requirement.

How Open Systems SnapVault backs up an SQL server database

The SQL Server database backup can be performed using Data ONTAP command-line interface on the secondary storage system and Protection Manager.

The `svapp` utility enables you to view all the SQL server database backup relationships and their file system paths from the Open Systems SnapVault command-line interface.

In Protection Manager, after adding the OSSV as a host you can see the `app:mssql` folder in NetApp Management Console. The `app:mssql` appears as one of the folders of the primary storage system. When you click the `app:mssql` folder, then all the MSSQL instances are displayed. Under each instance, you can see the database files. You can create datasets to back up MSSQL database. The `app:mssql` folder lists the SQL server database in the form of instances and each instance has full database folders and log folders.

Note

You should not select the `apps:mssql` folder and instances for creating datasets, select only the individual databases.

Full database backup: You must create a separate dataset for databases, set a separate schedule and protection policy for each dataset to back up. When the backup starts, Protection Manager sends the Microsoft SQL Server datasets path

in the form of *app:mssql:\<instance name>\<db name>* to **SnapVault or OSSV**. Open Systems SnapVault receives the paths and with the help of the mssql plug-in it discovers the actual file system path.

In an MSCS environment, the mssql plug-in discovers the Open Systems SnapVault database location. The Open Systems SnapVault database is created on a volume drive that has the SQL Server master database. The master database location helps the plug-in to find the actual file system path.

Any change in the master database location causes subsequent backups to fail. If you are moving the master database, you should also move the Open Systems SnapVault database to the same location. For more information about Open Systems SnapVault database, see “[Distributed Open Systems SnapVault database](#)” on page 85.

Open Systems SnapVault requests the Changelog minifilter driver to monitor the file system path.

The mssql plug-in takes the SQL writer aware VSS snapshot copy of the database. The plug-in also provides actual database and log file names that need to be backed up. Open Systems SnapVault performs a complete backup of the files if it is a initial backup. It uses minifilter driver change logs or BLI checksums to perform incremental backup.

Transaction log backup: When the transaction backup starts, Protection Manager sends the Microsoft SQL Server datasets path in the form of *app:mssql:\<instance name>\<db name:Tlog>* to SnapVault or OSSV. Open Systems SnapVault receives the paths and with the help of the mssql plug-in, it discovers the actual file system path. The plug-in deletes the backup copy of transaction log file if a full DB backup occurs after the last transaction log backup.

Note

You must create a separate dataset for transaction log databases, set a separate schedule, and protection policy for each transaction logs dataset to back up.

For backing up transaction logs, the mssql plug-in uses the custom-generated script called T-SQL and take the VSS snapshot. The plug-in provides the actual transaction log file names to be backed up.

The role of Changelog minifilter driver in transaction log backup is the same as that in case of the full database backup. For more information, see “[Configuring full database recovery behavior](#)” on page 118.

Local transaction log backup: Local transaction logs are unique in Open Systems SnapVault 3.0. The local transaction log enables you to back up the data changes between two scheduled SnapVault backups. Local transaction logs are stored in the primary storage system. You can set the local transaction logs backup interval in the *ossv_mssql.cfg* file by setting the *MSSQL:Local TLog Backup Interval* flag. The backup of these logs enables you to restore to a particular point-in-time state of the data.

For example, assume that you schedule hourly transaction log backups at 8 a.m., 9 a.m., and so on, and local transaction logs backup every 15 minutes. Your local transaction logs backups will happen at 8:15 a.m., 8:30 a.m., 8:45 a.m., and at 9:00 a.m. Your entire local transaction log backups that occurred at a particular hour are part of the hourly transaction log backup at 9:00 a.m.

If there is some data corruption and you want to restore to the database state at 8:45 a.m., you can restore it to the database state at 8:45 a.m. because you have taken local transaction logs backups. If you do not have the local transaction logs, you can only restore it to the database state at 8:00 a.m. For more information, see “[Setting up local transaction log backup interval](#)” on page 120.

Transaction log truncation: Open Systems SnapVault truncates the transaction log file if the database that is backed up uses the full recovery model or Bulk-logged recovery model. If the truncation is not applied, the transaction log file can use up all the space on the volume and can make the database nonoperational.

The log is normally truncated during the full database backup. The log file created during the full database backup is not backed up. You can choose not to truncate the transaction log file by setting the option *MSSQL:Transaction Log Truncate* to *FALSE* in the configuration file.

Note

The scheduled transaction log backups or the transaction log truncation do not result in a reduction in the file size. Truncation only helps in removing the committed transactions from the transaction log file and freeing up space in the transaction log file. However, the logical file size of the Transaction Log remains the same.

You can manually truncate the Transaction Log file if the Transaction Log file is huge and it is not possible for you to wait until the next scheduled backup. However, you should not to truncate the Transaction Log files manually because it breaks the log sequence and the Transaction Log backups cannot be used until the next full database backup. For more information, see [“Configuring the truncation behavior of transaction logs”](#) on page 122.

How Open Systems SnapVault restores SQL server database

The Open Systems SnapVault command-line interface enable you to restore a backed up an SQL Server database to the original location or to an alternate location in the same instance on your primary storage system. However, using Protection Manager you can only restore to the original location.

When you select the database to restore in Protection Manager or run the snapvault restore command from the command-line utility, the restore database path is provided to Open Systems SnapVault as `app:mssql:\<inst name>\<db name>`. Open Systems SnapVault discovers the restore locations with the help of the mssql plug-in. The mssql plug-in performs a tail-log backup to save the current transaction logs from being overwritten and puts the SQL database in Restore mode.

Open Systems SnapVault restores all the paths provided by the mssql plug-in as subfolders. After the restore is complete, Open Systems SnapVault puts the database into operational mode, based on the configuration setting in the `ossv-mssql.cfg` configuration file. If the *Recover After DB Restore* value is true in the configuration file, then the database is functional immediately after the restore. If the value is false, then you should manually make the database operational. For more information, see [“Configuring full database recovery behavior”](#) on page 118.

Restore SQL Server database with an alternate name: You can restore an SQL server database with an alternate database name in the same SQL instance only by using the Open Systems SnapVault command-line interface. To restore, you should ensure that the SQL instance to which you are restoring does not have a database name same as your alternate database name.

For example, if you specify the alternate database name is Test1_DB and if the SQL instance to which you are restoring has already a database by the same name (Test1_DB), then Open Systems SnapVault does not allow you to restore. You must provide a unique name.

Transaction Log files restore: When you select the transaction log to restore in Protection Manager or run the snapvault restore command from the command-line utility, the restore path is provided to Open Systems SnapVault as app:mssql:\<inst name>\<db name:Tlog>. Open Systems SnapVault discovers the restore locations with the help of the mssql plug-in. The mssql plug-in puts the SQL database in Restore mode.

Open Systems SnapVault restores all the paths provided by the mssql plug-in as subfolders. It uses SQL scripts for restoring. After the transaction log restore is complete, the mssql plug-in runs custom generated MSSQL scripts and puts the database into operational mode, based on the configuration setting in the ossv-mssql.cfg configuration file. If the [MSSQL:Recover After TLog Restore] value is true in the configuration file, then the database is functional immediately after the restore. If the value is false, then you should manually make the database operational. For more information, see “[Configuring transaction log recovery behavior](#)” on page 119.

Note

You must restore a full database before restoring the transaction logs.

Configuring backup and restore of Microsoft SQL databases

About configuration files

Open Systems SnapVault provides the *ossv_mssql.cfg* and *ossv-mssql-local-tlog-DBs.cfg* configuration files for managing the Microsoft SQL Server backup and recovery process. Using these configuration files, you can set the behavior of full database backup and recovery, the behavior of the transaction logs and local transaction logs. You can also set the truncation behavior of transaction logs.

Topics in this section

This section covers the following topics:

- ◆ [“Configuring full database recovery behavior”](#) on page 118
- ◆ [“Configuring transaction log recovery behavior”](#) on page 119
- ◆ [“Configuring local transaction log behavior”](#) on page 119
- ◆ [“Setting up local transaction log backup interval”](#) on page 120
- ◆ [“Adding a database list for backing local transaction logs”](#) on page 121
- ◆ [“Configuring the truncation behavior of transaction logs”](#) on page 122
- ◆ [“Specifying a directory path for transaction log files”](#) on page 122

Configuring full database recovery behavior

To set the full database recovery behavior, complete the following steps:

Step	Action	
1	Navigate to the <i>install_dir/snapvault/config</i> directory.	
2	In the config directory, open the <i>ossv_mssql.cfg</i> file in a notepad or WordPad.	
3	Depending on your requirement, set a value for the <i>MSSQL:Recover After DB Restore</i> flag:	
	If..	Then..
4	You want the database to be operational after the database restore	Set the value = TRUE
	You want to make the additional recovery before making the database operational	Set the value = FALSE
	By default, the value is FALSE.	
	Note _____ For transaction log restore, you must set the value as FALSE. _____	
5	Save and close the file.	

Configuring transaction log recovery behavior

To set the transaction log recovery behavior, complete the following steps:

Step	Action	
1	Navigate to the <i>install_dir/snapvault/config</i> directory.	
2	In the config directory, open the <i>ossv_mssql.cfg</i> file in a notepad or WordPad.	
3	Depending on your requirement, set a value for the <i>MSSQL:Recover After TLOG Restore</i> flag:	
	If..	Then..
4	You want the database to be operational after the transaction log restore	Set the value = TRUE
	You want to make additional recovery before making the database operational	Set the value = FALSE
	The default the value is TRUE.	
5	Save and close the file.	

Configuring local transaction log behavior

To set the local transaction log recovery behavior, complete the following steps:

Step	Action	
1	Navigate to the <i>install_dir/snapvault/config</i> directory.	
2	In the config directory, open the <i>ossv_mssql.cfg</i> file in a notepad or WordPad.	
3	Depending on your requirement, set a value for the <i>MSSQL:Restore Local TLog</i> flag:	
	If..	Then..

Step	Action	
4	You want to restore the local transaction logs along with the transaction logs	Set the value = TRUE
	You do not want to restore the local transaction logs along with the transaction logs and want to restore the local transaction logs manually	Set the value = FALSE
	By default, the value is FALSE.	
5	Save and close the file.	
	Note _____ The local transaction log behavior is also applicable to the tail-log backups. _____	

Setting up local transaction log backup interval

You can set the time interval for taking backups of the transaction log files locally. This configuration is applicable only to the database files that are listed in the *ossv-mssql-local-Tlog-DBs.cfg* file.

To set up local transaction log backup interval, complete the following steps:

Step	Action
1	Navigate to the <i>install_dir/snapvault/config</i> directory.
2	In the config directory, open the <i>ossv_mssql.cfg</i> file in a notepad or WordPad.
3	Depending on your requirement, set a value for the <i>MSSQL:Local TLog Backup Interval</i> flag.

Step	Action
4	<p>Set the value between 5 and 55 minutes.</p> <p>Any other value you enter is considered as zero.</p> <p>Note_____</p> <p>If you are changing the value from zero to any valid value, you must restart the Open Systems SnapVault service.</p> <p>_____</p>
5	Save and close the file.

Adding a database list for backing local transaction logs

The `ossv-mssql-local-Tlog-DBs.cfg` configuration file enables you to add a list of MSSQL Server databases. From this list you can back up the local transaction logs. By default, the file is empty. You can add the database in either of the following formats:

- ❖ Instance:DBName
- ❖ Instance

Note_____

The instance and database name can be a non-ASCII name.

Step	Action
1	Navigate to the <code>install_dir/snapvault/config</code> directory.
2	In the config directory, open the <code>ossv-mssql-local-Tlog-DBs.cfg</code> file in a notepad or WordPad.
3	<p>In the <code>ossv-mssql-local-Tlog-DBs.cfg</code> file add list of the SQL server database names.</p> <p>Example: To include local Tlogs of MyDB2 and MyDB4 in Instance1 and MyDB5 in Instance2 during backup, add the database names in the <code>ossv-mssql-local-Tlog-DBs.cfg</code> file as shown below.</p> <pre>Instance1:MyDB2 Instance1:MyDB4 Instance2:MyDB5</pre>

Step	Action
4	Save and close the file.

Configuring the truncation behavior of transaction logs

You can configure Open Systems SnapVault to truncate the transaction logs during a full database backup or transaction log backup. By default the value is TRUE, and Open Systems SnapVault truncates transaction logs during backup.

Step	Action	
1	Navigate to the <i>install_dir/snapvault/config</i> directory.	
2	In the config directory, open the <i>ossv_mssql.cfg</i> file in a notepad or WordPad.	
3	Depending on your requirement, set a value for the <i>MSSQL:TLog Truncate</i> flag:	
4	If..	Then..
	You want to truncate transaction log files	Set the value = TRUE
	You do not want to truncate the transaction log files	Set the value = FALSE
	By default, the value is TRUE.	
5	Save and close the file.	

Specifying a directory path for transaction log files

You can specify a different directory path for saving transaction log files other than the default directory.

To specify a different directory, complete the following steps:

Step	Action
1	Navigate to the <i>install_dir/snapvault/config</i> directory.

Step	Action
2	In the config directory, open the ossv_mssql.cfg file in a notepad or WordPad.
3	Find for the MSSQL:TLog Backup Director flag in the file.
4	Enter a directory path for saving transaction log files. Note _____The path should be an ASCII path._____
5	Save and close the file.

Viewing SQL Server database from the command-line interface

About the svapp command-line utility

The svapp.exe command-line utility enables you to view the Microsoft SQL server database instances on your primary storage system. You can use this utility if you are not using Protection Manager for Microsoft SQL Server database backup and recovery. The svapp utility provides details about the existing SnapVault backup relationships and helps you to identify the database files that need to be backed up.

This section covers the following topics:

- ◆ [“Commands to view Microsoft SQL server database when using the svapp utility”](#) on page 124
- ◆ [“Viewing all the SQL server database instances”](#) on page 125
- ◆ [“Viewing an SQL server database list in a particular instance”](#) on page 126

Commands to view Microsoft SQL server database when using the svapp utility

The following table lists the commands that help to view the Microsoft SQL Server database details if you are using the svapp utility.

Use this <i>svapp</i> utility command...	To display information about...
<code>svapp list</code>	Supported applications.
<code>svapp list -path app_name</code>	Components of an application path.
<code>svapp list -verbose app_name</code>	Details of the application, database, and log file location for each database.
<code>svapp list -recursive</code>	All subcomponents of an application or application component.

Viewing all the SQL server database instances

To view all the Microsoft SQL Server database instances, complete the following steps:

Step	Action
1	In the primary system management console, navigate to the bin directory. <code>C:\Program Files\NetApp\snapvault\bin></code>
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\bin>svapp list mssql</code> Example: <code>C:\Program Files\NetApp\snapvault\bin>svapp list mssql</code> Instance ----- SQLInstance1 SQLInstance2 SQLInstance3

Excluding SQL Server files during complete drive backup

You can exclude the SQL Server files that are backed up during file system files backup or complete drive backup. This exclusion ensures that the SQL server files are not backed up twice. During the backup of file system, the SQL Server database files are backed up as application-inconsistent files. Therefore, it is a best practice to exclude the SQL Server files from backup during backup of file system or complete drive backup and avoid duplication of SQL Server files backup.

Identify the SQL Server database files path using the command-line utility and add the paths to the file-exclude.txt. For more information, see [“Configuring backup exclusion lists”](#) on page 75.

To find the SQL Server database file paths, complete the following steps:

Step	Action																
1	In the primary system management console, navigate to the bin directory. <code>C:\Program Files\NetApp\snapvault\bin></code>																
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\bin>svapp list -verbose mssql</code>																
<p>Example:</p> <pre>Instance : SQLInstance1 Version : 9.00.1399.06 Language : English Clustered : FALSE Microsoft SQL Server Enterprise Edition (64-bit)</pre> <table border="1"> <thead> <tr> <th>Database\TLog</th> <th>Backup Path</th> <th>Protected</th> <th>FS Paths</th> </tr> </thead> <tbody> <tr> <td>DB1</td> <td>app:mssql:SQLInstance1:DB1</td> <td>Yes</td> <td>c:\mssql\db1.mdf c:\mssql\db1.ldf</td> </tr> <tr> <td>DB1 TLog</td> <td>app:mssql:SQLInstance1:DB1:TLog</td> <td>Yes</td> <td>c:\mssql \backedupTlogs \MSSQLSERVER\db 1\db1.trn</td> </tr> <tr> <td>DB2</td> <td>app:mssql:SQLInstance1:DB2</td> <td>No</td> <td>c:\mssql\db2.mdf</td> </tr> </tbody> </table>		Database\TLog	Backup Path	Protected	FS Paths	DB1	app:mssql:SQLInstance1:DB1	Yes	c:\mssql\db1.mdf c:\mssql\db1.ldf	DB1 TLog	app:mssql:SQLInstance1:DB1:TLog	Yes	c:\mssql \backedupTlogs \MSSQLSERVER\db 1\db1.trn	DB2	app:mssql:SQLInstance1:DB2	No	c:\mssql\db2.mdf
Database\TLog	Backup Path	Protected	FS Paths														
DB1	app:mssql:SQLInstance1:DB1	Yes	c:\mssql\db1.mdf c:\mssql\db1.ldf														
DB1 TLog	app:mssql:SQLInstance1:DB1:TLog	Yes	c:\mssql \backedupTlogs \MSSQLSERVER\db 1\db1.trn														
DB2	app:mssql:SQLInstance1:DB2	No	c:\mssql\db2.mdf														

Viewing an SQL server database list in a particular instance

To view the Microsoft SQL Server database and log files of a particular path, complete the following steps:

Step	Action
1	In the primary system management console, navigate to the bin directory. <code>C:\Program Files\NetApp\snapvault\bin></code>

Step	Action
2	In the console, enter the following command: <pre>C:\Program Files\NetApp\snapvault\bin>svapp list mssql -path</pre>
<p>Example:</p> <pre>C:\Program Files\netapp\snapvault\bin>svapp list mssql:SQLInstance1 -path Instance: SQLInstance1 Database\TLog Backup Path Protected FS Paths ----- DB1 app:mssql:SQLInstance1:DB1 Yes c:\mssql\db1.mdf c:\mssql\db1.ldf DB1 TLog app:mssql:SQLInstance1:DB1:TLog Yes c:\mssql \backedupTlogs \MSSQLSERVER\db 1\db1.trn DB2 app:mssql:SQLInstance1:DB2 No c:\mssql\db2.mdf</pre>	

Backing up and restoring Microsoft SQL Server databases

Overview

You can back up and restore Microsoft SQL Server database by using the command-line interface from your secondary storage system or by using the Protection Manager.

This section covers the following topics:

- ◆ [“Backing up using the command-line interface”](#) on page 129
- ◆ [“Restoring using the command-line interface”](#) on page 131
- ◆ [“Backing up Microsoft SQL Server database using Protection Manager”](#) on page 132
- ◆ [“Restoring Microsoft SQL Server database using Protection Manager”](#) on page 135

Backing up and restoring Microsoft SQL Server database

Backing up using the command-line interface

Backup process

To back up Microsoft SQL Server database from the command-line interface you should perform tasks from the primary and secondary storage systems. From the primary storage system you should select the database to be backed up. The *svapp* utility helps you to view the details. After you select the database that needs to be backed up, you can initiate backup process from the secondary storage system.

To back up Microsoft SQL Server database from command-line interface, complete the following steps:

Step	Action																				
1	In the primary system management console, navigate to the bin directory. <code>C:\Program Files\NetApp\snapvault\bin></code>																				
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\bin>svapp list -path</code>																				
<p>Example:</p> <pre>C:\Program Files\netapp\snapvault\bin>svapp list mssql:SQLInstance1 -path Instance: SQLInstance1</pre> <table border="1"> <thead> <tr> <th>Database\TLog</th> <th>Backup Path</th> <th>Protected</th> <th>FS Paths</th> </tr> <tr> <th>-----</th> <th>-----</th> <th>-----</th> <th>-----</th> </tr> </thead> <tbody> <tr> <td>DB1</td> <td>app:mssql:SQLInstance1:DB1</td> <td>Yes</td> <td>c:\mssql\db1.mdf c:\mssql\db1.ldf</td> </tr> <tr> <td>DB1 TLog</td> <td>app:mssql:SQLInstance1:DB1:TLog</td> <td>Yes</td> <td>c:\mssql \backedupTlogs \MSSQLSERVER\db 1\db1.trn</td> </tr> <tr> <td>DB2</td> <td>app:mssql:SQLInstance1:DB2</td> <td>No</td> <td>c:\mssql\db2.mdf</td> </tr> </tbody> </table>		Database\TLog	Backup Path	Protected	FS Paths	-----	-----	-----	-----	DB1	app:mssql:SQLInstance1:DB1	Yes	c:\mssql\db1.mdf c:\mssql\db1.ldf	DB1 TLog	app:mssql:SQLInstance1:DB1:TLog	Yes	c:\mssql \backedupTlogs \MSSQLSERVER\db 1\db1.trn	DB2	app:mssql:SQLInstance1:DB2	No	c:\mssql\db2.mdf
Database\TLog	Backup Path	Protected	FS Paths																		
-----	-----	-----	-----																		
DB1	app:mssql:SQLInstance1:DB1	Yes	c:\mssql\db1.mdf c:\mssql\db1.ldf																		
DB1 TLog	app:mssql:SQLInstance1:DB1:TLog	Yes	c:\mssql \backedupTlogs \MSSQLSERVER\db 1\db1.trn																		
DB2	app:mssql:SQLInstance1:DB2	No	c:\mssql\db2.mdf																		
3	Make a note of the database that you want to back up from the list.																				

Step	Action
4	<p>In the secondary storage system console, enter the following command:</p> <pre> snapvault start -S primary_host:app:mssql:instance: database /vol/sec_vol/sec_tree </pre> <p>Example:</p> <pre> snapvault start-S primary_host:app:mssql:SQLInstance1: DB2 /vol/sqlvol/sql dbs </pre> <p>After the initial baseline transfer, you can schedule the backups.</p> <p>For information about scheduling backups, see “Scheduling SnapVault update backups” on page 100.</p>

Restoring using the command-line interface

Restore process

You should consider the following points before restoring the MSSQL databases:

- ◆ You can only restore a database to the original location from where you backed it up previously.
- ◆ You must restore a full database before restoring the transaction logs.
- ◆ Ensure that the primary and secondary systems names resolved.
- ◆ You must set the *MSSQL:Recover After DB Restore* flag option and other flags in the *ossv_mssql.cfg* file, based on your requirements.

Step	Action
1	In the primary system management console, navigate to the bin directory. <code>C:\Program Files\NetApp\snapvault\bin></code>
2	To restore a database, enter the following command: <code>snapvault restore -S <secondary host>:<qtree name> app:mssql:inst:database</code>

Backing up Microsoft SQL Server database using Protection Manager

Prerequisites

You should meet the following prerequisites to back up and restore Microsoft SQL server databases using Protection Manager:

- ◆ Open Systems SnapVault 3.0 is installed.
- ◆ Microsoft SQL server 2005 or Microsoft SQL server 2008 is installed on the Open Systems SnapVault primary storage system.
- ◆ DataFabric Manager 3.8 or later is installed.
- ◆ Protection Manager license is enabled.
- ◆ NetApp Management Console is installed on the primary storage system
- ◆ Windows Server 2003 or Windows Server 2008 Server added as the Open Systems SnapVault host.
- ◆ NetApp Host Agent is installed on the system to manage Open Systems SnapVault using DataFabric Manager.
- ◆ A NetApp secondary storage system is added as a storage host.
- ◆ Resource pool is added from the secondary storage system.

For information about adding hosts and creating a resource pool, see the *Provisioning Manager and Protection Manager Printable Help and NetApp Management Console Online Help*.

Backing up Microsoft SQL database using Protection Manager involves the following tasks:

- ◆ Creating datasets
- ◆ Assigning protection policy
- ◆ Scheduling the backups

Backing up using Protection Manager

To back up the Microsoft SQL database using Protection Manager, complete the following steps:

Step	Action
1	Log in to Protection Manager using NetApp Management Console.
2	From the navigation pane, click Data > Datasets > Overview .
3	Click Add to start the Add Dataset wizard.
4	Enter a name for your dataset and click Next .
5	In the Available Resources property sheet, select the host system from the resource pool trees display.
6	In the resources pool tree, select <i>app:mssql</i> .
7	Under the <i>app:mssql</i> resource, select the database to backup.
8	Click “>” to add the resources to the dataset and click Next . Note _____ The maximum number of datasets you can add is 50.
9	Complete the steps in the wizard to create a dataset. You have to assign a protection policy to the dataset.

Assigning a protection policy

To assign a protection policy for the dataset you created, complete the following steps:

Step	Action
1	From the navigation pane, click Data > Datasets > Overview .
2	Select a dataset and click Protection Policy to start the Dataset Policy Change wizard.
3	In the Welcome to Protection Policy wizard, click Next .
4	In the Protection Policy property sheet, select the <i>Remote backups only</i> protection policy and click Next .

Step	Action
5	In the Modify New Node Resources property sheet, select <i>Provision and attach Resources using policy</i> or <i>Assign resources manually</i> .
6	In the Destination Node Resources property sheet, select the destination resources for backup and click “>” to add to the Resources in this node pane.
7	<p>Click Next and complete the steps in the wizard to assign a protection policy.</p> <p>Protection Manager creates a relationship and performs a baseline backup of the database to the secondary storage system.</p> <p>You have to create a schedule for the backups. For information about scheduling the backups, see the <i>NetApp Management Console</i> online help.</p>

Restoring Microsoft SQL Server database using Protection Manager

Restoring database

You should consider the following points before restoring database:

- ◆ You can only restore database to the original location from where you backed it up previously.
- ◆ You must restore a full database before restoring the transaction logs.
- ◆ You must set the *MSSQL:Recover After DB Restore* flag option and other flags in the *ossv_mssql.cfg* file, based on your requirements.

To restore Microsoft SQL database using Protection Manager, complete the following steps:

Step	Action
1	Log in to Protection Manager using NetApp Management Console.
2	From the navigation pane, click Data > Datasets > Overview .
3	From the list of datasets, select the dataset that you want to restore.
4	Click Restore to start the Restore wizard.
5	Complete the steps in the Restore wizard.

About this chapter Most of the Open Systems SnapVault management tasks you perform are similar to 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 138
- ◆ [“Backing up and restoring the Open Systems SnapVault database”](#) on page 140
- ◆ [“Backing up and restoring Windows System State data”](#) on page 145
- ◆ [“Deleting and re-creating Open Systems SnapVault relationships”](#) on page 155
- ◆ [“Migrating a relationship between two secondary storage systems”](#) on page 156
- ◆ [“Migrating between two volumes on one secondary storage system”](#) on page 159
- ◆ [“Setting up a tertiary system for a relationship”](#) on page 162
- ◆ [“Reusing a deleted or renamed primary backup root directory name”](#) on page 164
- ◆ [“Reusing a renamed Open Systems SnapVault primary host name”](#) on page 166
- ◆ [“Renaming a SnapVault secondary volume”](#) on page 167
- ◆ [“Resynchronizing restored or broken relationships”](#) on page 169
- ◆ [“Retrying failed transfers”](#) on page 172
- ◆ [“Encrypted File System \(EFS\) file backup and restore”](#) on page 174

Locating status and problem reports

Where to find status and problem reports

You can find all the log files 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

`snapvault` is the current file, 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.

Deleting the old snapvault logs

The `snapvault` log files consume a lot of space with time. Therefore, the old log files need to be deleted to make sufficient space for the new log files.

To keep the number of `snapvault` log files you can set the value in the General tab of the Configurator GUI.

If the number of log files to keep is changed in the `configure.cfg`, you should restart the Open Systems SnapVault services.

If you set the value to 0 (zero) in the *old snapvault logs to keep*, it means that it will keep all old snapvault log files.

If the value is set to 1 (one), it keeps the current snapvault log file and the recent old log file.

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. Each relationship maintains a unique set containing the following files:

- ◆ History file
- ◆ BLI checksums file (if BLI is enabled)
- ◆ Checkpoint file (if a backup process 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	.xx .xx is a unique integer
BLI checksums	.xx-checksums

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

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

11/08/05  02:32  4,676  01                History file
11/08/05  02:32   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 is transferred to the secondary storage system, the primary storage system deletes the file.

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

Backup process does not include checkpoint files. Also Open Systems SnapVault does not backup a softlock state in the database.

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

You can restore the database file 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 143 for the steps to restore the database file.

After restoring the database file, Open Systems SnapVault software decompresses it automatically and places the decompressed files where Open Systems SnapVault database files are located for the relationship. You can perform the data transfers from this point onward.

Note

If any data updates occur between the time a database file is backed up or restored, they cause the secondary storage system to get out-of-sync with the primary storage system; therefore, subsequent data updates cannot continue. In such a case, you must first resynchronize the relationship with the `snapvault start -r` command. After the resynchronization has completed, perform the data updates as usual. For more information about resynchronizing a relationship, see [“Resynchronizing restored or broken relationships”](#) on page 169.

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

Restoring the agent database

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

Step	Action
1	Using the command-line of the Open Systems SnapVault primary storage system, navigate to the snapvault/bin directory. <ul style="list-style-type: none"> ◆ On Windows systems, the default path is C:\Program Files\netapp\snapvault\bin. ◆ On UNIX systems, the default path is /usr/snapvault/bin.

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 you restore the database.</p> <hr/> <p>Note</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> <hr/> <p>Example: 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>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 the configuration, Windows 2003 systems have some or all 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 68.

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 deleted accidentally. You can also use Open Systems SnapVault System State data backup along with complete file system backups as part of a disaster recovery plan.

Taking backup to restore the records and clearing the record helps to maintain the Windows EventLog. Clear the EventLog before it reaches its maximum size, otherwise, 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 about such issues.

Event logs should be backed up separately from any other system files. During a system backup, the event log files are copied and therefore unusable. If you attempt to open a backed up or copied event log file 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 149
- ◆ [“Restoring System State data”](#) on page 150
- ◆ [“Using System State data backup to rebuild a primary storage system”](#) on page 152

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.

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>Note</p> <hr/> <p>The keyword SystemState is case-independent.</p> <hr/> <p>Example:</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 151.

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>Note_____</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>_____</p> <p>Example:</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>Note_____</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>Example:</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">snapvault start -S prim_host:SystemState sec_host:/vol/sec_vol/sec_tree</pre> <p>Note</p> <hr/> <p>The keyword SystemState is case-independent.</p> <hr/> <p>Example:</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 152, 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"> ◆ Use the same operating system and service packs that were on the original system. ◆ Use the same machine name as on the original system. ◆ Use the same drive letter mappings. ◆ Make sure that each drive is at least as large as the corresponding drive was when the backup was made. ◆ Format each drive with the same file system type and version as on the original system. ◆ Make sure that the hardware configuration is identical to the original. <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 “Restoring System State data from a backup” on page 150 and “Restoring System State data from a backup and marking it primary” on page 151 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: <pre>snapvault stop <i>secondary_path</i></pre>
2	On the primary storage system console, enter the following command: <pre>snapvault release <i>primary_path</i> [<i>secondary:</i>] <i>secondary_path</i></pre>
3	On the secondary storage system console, enter the following command: <pre>snapvault start -s <i>primary_path</i> [<i>secondary:</i>] <i>secondary_path</i></pre>

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>Example: 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>Example: 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">Create the new_vol volume: <pre>r200-new> vol create new_vol 3</pre>Mark the new_vol volume as restricted: <pre>r200-new> vol restrict new_vol</pre>Transfer the old_vol volume to the new_vol volume: <pre>r200-new> snapmirror initialize -S r200-old:old_vol new_vol</pre>

Step	Action												
3	<p>Quiesce and break the SnapMirror relationship between the old secondary storage system and the new secondary storage system.</p> <p>Example: 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"> a. <code>snapmirror quiesce new_vol</code> b. <code>snapmirror break new_vol</code> 												
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>Example: Perform the following steps from r200-new:</p> <ul style="list-style-type: none"> a. <code>snapmirror status</code> <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"> b. <code>snapvault status</code> <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>Example: 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>Note This does not start a new baseline, it updates the registry.</p> <hr/> <p>Example: Perform the following step from r200-new: <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>Example: Perform the following step from r200-new: <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 r200-new.</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>Example: Perform the following step from r200-new: <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>Example: 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>Example: 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">a. Create the new_vol volume. <code>vol create new_vol 3</code>b. Mark the new_vol volume as restricted. <code>vol restrict new_vol</code>c. Transfer the old_vol volume to the new_vol volume. <code>snapmirror initialize -S r200:old_vol new_vol</code>

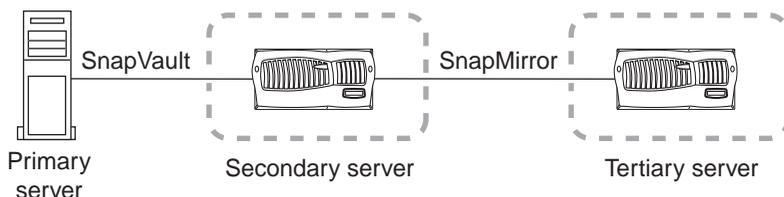
Step	Action												
3	<p>Quiesce and break the SnapMirror relationship between the old volume and the new volume.</p> <p>Example: To quiesce and break the SnapMirror relationship between old_vol and new_vol, run the following commands on r200:</p> <ul style="list-style-type: none"> a. snapmirror quiesce new_vol b. snapmirror break new_vol 												
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>Example: Perform the following steps from r200.</p> <ul style="list-style-type: none"> a. snapmirror status <table border="0" data-bbox="490 737 1142 798"> <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"> b. snapvault status <table border="0" data-bbox="490 850 1229 911"> <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 snapvault status -c command.</p> <p>Example: 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 <code>snapvault start</code> command.</p> <hr/> <p>Note This does not start a new baseline; it updates the registry.</p> <hr/> <p>Example: 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 <code>snapvault status -c</code> command.</p> <p>Example: 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 <code>new_vol</code>.</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>Example: 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 NetApp 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.

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

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

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

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: <code>snapvault modify -s new_source_hostname:path destination_filer:path</code>
2	Run the <code>snapvault update</code> command and the new host name is displayed in the SnapVault status: <code>snapvault update destination_filer:path</code>

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: <code>vol rename <i>oldvolname</i> <i>newvolname</i></code>
2	Enter the following commands to verify the changes: <code>snapvault status</code> snapvault status displays the new path. <code>snapvault status -c</code> snapvault status -c does not display the new path.
3	Enter the following command: <code>snapvault start -S <i>primary_filer:primary_qtree</i> <i>secondary_filer:secondary_qtree</i></code> 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: <code>snapvault status -c</code> snapvault status -c now displays the new path.
5	Enter the following command to verify whether the change was successful: <code>snapvault update <i>secondary_qtree</i></code>

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>snapvault stop /vol/oldvolname/secondary_qtree</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 7.1.2, 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 1002 374">snapvault start -r -S prim_host:dirpath /vol/sec_vol/sec_tree</pre> <p data-bbox="494 401 780 427">Example 1 (Windows):</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">Example 2 (UNIX):</p> <pre data-bbox="494 548 1052 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>Note</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 2.6.1 and later, 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.

About this chapter

This chapter provides information about the Open Systems SnapVault Changelog minifilter driver, such as how it works, its limitations, how to install and uninstall the minifilter driver, and other related information.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Overview”](#) on page 175
- ◆ [“How the Changelog minifilter driver works”](#) on page 176
- ◆ [“Limitations”](#) on page 177
- ◆ [“How filter driver works with other features”](#) on page 178
- ◆ [“Datasets favorable for the Changelog filter driver”](#) on page 178
- ◆ [“The Changelog minifilter driver configuration files”](#) on page 178
- ◆ [“Values of the configuration”](#) on page 178
- ◆ [“Changelog filter driver management”](#) on page 179
- ◆ [“Verifying the Changelog minifilter driver install status”](#) on page 179
- ◆ [“Verifying the Changelog minifilter load status”](#) on page 179
- ◆ [“Enabling or disabling the Changelog minifilter driver for applications data”](#) on page 179
- ◆ [“Enabling or disabling the Changelog minifilter for file system data”](#) on page 180
- ◆ [“Viewing the log files count”](#) on page 181
- ◆ [“Setting limits on the log files count”](#) on page 181
- ◆ [“Loading the Changelog minifilter driver”](#) on page 183
- ◆ [“Unloading the Changelog minifilter driver”](#) on page 183
- ◆ [“Uninstalling the Changelog minifilter driver”](#) on page 184
- ◆ [“Troubleshooting”](#) on page 184

Overview

The Changelog minifilter driver is a new feature in Open Systems SnapVault 3.0, which enables you to perform faster incremental backups after initial baseline transfer. The minifilter driver is a file system filter driver developed by using the Microsoft minifilter model. It is available on all Open Systems SnapVault supported Windows platforms.

After the baseline transfer, the Changelog minifilter driver monitors changes in the files data that are part of the Open Systems SnapVault backup. The information about the modified blocks is logged in a log file. The log file with changed blocks information helps Open Systems SnapVault to perform faster incremental backups. The Changelog minifilter driver also works in a Microsoft Cluster environment.

The following Windows platforms are supported for the Changelog minifilter driver:

- ◆ Windows Server 2003 SP1
- ◆ Windows Server 2003 R2
- ◆ Windows Storage Server 2003
- ◆ Windows Server 2008
- ◆ Windows Server 2008 R2

The Changelog minifilter driver is installed along with Open Systems SnapVault installation. You can also choose to install the driver later.

For incremental backups of application data, Open Systems SnapVault by default uses the Changelog minifilter driver. For incremental backups of file system data, Open Systems SnapVault does not use the Changelog minifilter driver. It uses BLI for these backups. However, you can make configuration changes and use the Changelog minifilter driver for file system data backups also.

Note

You can make configuration changes and control the usage of the Changelog minifilter driver for incremental backups of application or file system data.

How the Changelog minifilter driver works

Open Systems SnapVault uses the following method to identify the changed blocks when the Changelog minifilter driver is not being used.

- ◆ Open Systems SnapVault scans the backup folder to identify the changed files.
- ◆ It performs a checksum of every 4 KB of the file and compares with the checksum of the previous backup and identifies the changed blocks of the file.
- ◆ If only a few blocks in a file are changed, the checksum is computed on the complete file to identify the few changed blocks. This method is time consuming for large files.

The Changelog minifilter driver method of identifying the changed blocks enables faster incremental backups.

Open Systems SnapVault uses the following method to identify the changed blocks when the Changelog minifilter driver is being used.

- ◆ Open Systems SnapVault scans the backup folder to identify the changed files.
- ◆ It receives changed blocks information of all the files from the Changelog minifilter driver.
- ◆ It performs a checksum of only the modified blocks and updates the checksum data.
- ◆ The checksum data can be used if Open Systems SnapVault has to revert to traditional BLI checksum based backups for some uses cases, such as after system reboot or after failover/failback in a cluster environment.
- ◆ The Changelog minifilter driver integrates with the VSS Snapshot copy mechanism and identifies the changes to ensure backup data is consistent.

Note

You can stop using the traditional checksum method by disabling BLI.

Example: A 20-MB file is backed up to a secondary storage system. After backup is performed, there is a 1-MB data change in the file on the primary system. The Changelog minifilter driver monitors the data changes and logs the changes in the driver log file. For the next incremental back up, Open Systems SnapVault performs checksum only for the changed 1-MB data and thus enables faster incremental backups.

Limitations

The following are the Changelog minifilter driver limitations:

- ◆ If the system is booted in the safe mode, the minifilter driver does not load and track the changes.
- ◆ After primary system reboot, Open Systems SnapVault uses the traditional checksum for the first backup and for subsequent backups the Changelog minifilter driver.
- ◆ In case of resynchronization after restore, the minifilter driver continues to monitor the original backup relationship till it is released on the primary system.
- ◆ Minifilter driver does not track changes to hard links, sparse files, and encrypted and compressed files. However, Open Systems SnapVault continues to back up these files using the traditional checksum method.

How filter driver works with other features

System state: The Changelog minifilter driver is not used for backing up the System State data.

Compression: You can use the Changelog minifilter driver and compression feature together for backup.

Restore : The Changelog minifilter is not used for during restore process.

Datasets favorable for the Changelog filter driver

The Changelog minifilter driver enables faster incremental backups for file sizes greater than 1 MB.

When using the Changelog minifilter driver for backing up a large number of small files, it does not cut down the incremental backup time.

The Changelog minifilter driver configuration files

The Changelog minifilter driver configuration files are installed in the Windows Registry.

The Changelog filter configuration has the following options:

- ◆ Log count
- ◆ Log path

You can change the status of these configuration files using the FilterCfg command-line utility.

Values of the configuration

Log count: You can set the maximum number of log files to retain. The default number of log files to retain is three.

Log path : The default path is *C:\Program Files\NetApp\snapvault\changelog* for the log files. You can change the path for the log files.

Changelog filter driver management

Verifying the Changelog minifilter driver install status

To verify that the Changelog minifilter driver is installed, complete the following steps:

Step	Action
1	In the primary system management console navigate to the fltdrvr directory. <code>C:\Program Files\NetApp\snapvault\fltdrvr</code>
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\fltdrvr\filtercfg installstatus</code>

Verifying the Changelog minifilter load status

To verify the Changelog minifilter driver load status, complete the following steps:

Step	Action
1	In the primary system management console navigate to the fltdrvr directory. <code>C:\Program Files\NetApp\snapvault\fltdrvr</code>
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\fltdrvr\filtercfg filterstatus</code>

Enabling or disabling the Changelog minifilter driver for applications data

To enable or disable the Changelog minifilter driver for backing up application data, complete the following steps:

Step	Action
1	Navigate to the <code>install_dir/snapvault/config</code> directory.

Step	Action	
2	In the config directory, open the <i>Snapvault.cfg</i> file in a notepad or WordPad.	
3	Depending on your requirement, set a value for the <i>OSSV:UseChangelogsForApps</i> flag:	
	If..	Then..
4	You want to use the Changelog minifilter driver for application data back up	Set the value = TRUE
	You do not want to use Changelog minifilter driver for application data back up	Set the value = FALSE
	By default, the value is TRUE.	
5	Save and close the file.	

Enabling or disabling the Changelog minifilter for file system data

To enable or disable the Changelog minifilter driver for backing up file system data, complete the following steps:

Step	Action	
1	Navigate to the <i>install_dir/snapvault/config</i> directory.	
2	In the config directory, open the <i>Snapvault.cfg</i> file in a notepad or WordPad.	
3	Depending on your requirement, set a value for the <i>OSSV:UseChangelogsForFileSystems</i> flag:	
	If..	Then..

Step	Action	
4	You want to use the Changelog minifilter driver for file system data back up	Set the value = TRUE
	You do not want to use the Changelog minifilter driver for file system data back up	Set the value = FALSE
	By default, the value is FALSE.	
5	Save and close the file.	

Viewing the log files count

To view the total number of log files, complete the following steps:

Step	Action
1	In the primary system management console navigate to the fltdrvr directory. <code>C:\Program Files\NetApp\snapvault\fltdrvr</code>
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\fltdrvr\filtercfg get-config-integer /stanza <stanza name></code> Example: <code>C:\Program Files\NetApp\snapvault\fltdrvr\filtercfg get-config-integer /stanza Log Count</code>

Setting limits on the log files count

The default number of log files is three. You can change this count.

To set limit on the number of log files, complete the following steps:

Step	Action
1	In the primary system management console navigate to the fltdrvr directory. <code>C:\Program Files\NetApp\snapvault\fltdrvr</code>

Step	Action
2	<p>In the console, enter the following command:</p> <pre>C:\Program Files\NetApp\snapvault\fltldr\filtercfg set-config-integer /stanza <stanza name> /value <value></pre> <p>Example:</p> <pre>C:\Program Files\NetApp\snapvault\fltldr\filtercfg set-config-integer /stanza Log Count /value 5</pre>

Viewing the log file path

To view the path where the log files are saved, complete the following steps:

Step	Action
1	<p>In the primary system management console navigate to the fltdrvr directory.</p> <pre>C:\Program Files\NetApp\snapvault\fltldr</pre>
2	<p>In the console, enter the following command:</p> <pre>C:\Program Files\NetApp\snapvault\fltldr\filtercfg get-config-string /stanza <stanza name></pre> <p>Example:</p> <pre>C:\Program Files\NetApp\snapvault\fltldr\filtercfg get-config-string /stanza Log Path</pre>

Modifying the log files path

If you want move the log files from the default location to another location, complete the following steps:

1	<p>In the primary system management console navigate to the fltdrvr directory.</p> <pre>C:\Program Files\NetApp\snapvault\fltldr</pre>
---	--

2	<p>In the console, enter the following command:</p> <pre>C:\Program Files\NetApp\snapvault\fltldr\filtercfg set-config-string /stanza <stanza name> /value <value></pre> <p>Example:</p> <pre>C:\Program Files\NetApp\snapvault\fltldr\filtercfg set-config-string /stanza Log Path /value D:\log files</pre>
----------	---

Loading the Changelog minifilter driver

If the Changelog minifilter driver is unloaded manually, you can reload the driver.

To load the Changelog minifilter driver, complete the following steps:

Step	Action
1	<p>In the primary system management console navigate to the fltldr directory.</p> <pre>C:\Program Files\NetApp\snapvault\fltldr</pre>
2	<p>In the console, enter the following command:</p> <pre>C:\Program Files\NetApp\snapvault\fltldr\filtercfg load</pre>

Unloading the Changelog minifilter driver

You can unload the Changelog minifilter driver for troubleshooting or when you are requested by NetApp technical support.

Caution

You should not unload or uninstall the Changelog minifilter driver if you are stopping the Open Systems SnapVault services for some reason. Unloading the Changelog minifilter driver results in revert to BLI checksum and impacts incremental backup time, especially if the backup includes large files.

To unload the Changelog minifilter driver, complete the following steps:

Step	Action
1	In the primary system management console navigate to the fltdrvr directory. <code>C:\Program Files\NetApp\snapvault\fltdrvr</code>
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\fltdrvr\filtercfg unload</code>

Uninstalling the Changelog minifilter driver

To uninstall the Changelog minifilter driver, complete the following steps:

Step	Action
1	In the primary system management console navigate to the fltdrvr directory. <code>C:\Program Files\NetApp\snapvault\fltdrvr</code>
2	In the console, enter the following command: <code>C:\Program Files\NetApp\snapvault\fltdrvr\filtercfg uninstall /path <inf path></code>

Troubleshooting

For troubleshooting issues, you should collect the dump file.

- ◆ If the system crashes, you should collect the dump file from %SystemRoot%\MEMORY.DMP.
- ◆ If the system does not respond, you should use the keyboard shortcut keys to collect the dump file from %SystemRoot%\MEMORY.DMP.

Setting up system for creating dump file

To set up system for creating dump file, complete the following steps:

Note

You should ensure that the paging file is on the boot drive.

Step	Action
1	Click Start > Control Panel .
2	Double-click the System icon.
3	Click Advanced .
4	Click Settings under Startup and Recovery section.
5	In the Startup and Recovery window, select Kernel memory dump from the drop-down list of the Write debugging information section.
6	Verify that dump file location shows the default %SystemRoot%\MEMORY.DMP.
7	Select the Overwrite any existing file check box.
8	Click OK .

Set up system for generating dump file using the keyboard

You should set up the keyboard shortcut keys to generate a dump file when the system does not respond.

Follow the instructions in the KB 244139 available in the Microsoft page.

Guidelines for the paging file size

As per the information provided in the Microsoft's KB 307973 the required size of the paging file depends on the amount of RAM in your computer.

The maximum amount of available space for a kernel memory dump on a 32-bit system is 2 GB plus 16 MB; on a 64-bit system, the maximum amount of available space for a kernel memory dump is the RAM size plus 128 MB.

The following table contains guidelines for the size of the paging file:

RAM size	Paging file size
256 MB–1,373 MB	1.5 times the RAM size
1,374 MB or greater	32-bit system: 2 GB plus 16 MB 64-bit system: RAM size plus 128 MB

Enabling the minifilter driver traces

You should enable the minifilter driver traces only when you are requested to do so for troubleshooting.

To enable the Changelog minifilter driver traces, complete the following steps:

Step	Action
1	Download the tracelog tool from the Microsoft download page.
2	<p>Run the following command to start the trace session:</p> <pre>tracelog -start changelog -guid changelog.ctl -f changelog.etl -flags traceflag -level tracelevel</pre> <p><i>traceflag</i> identifies a subsystem in the filter driver.</p> <p><i>tracelevel</i> sets the verbosity level of the traces.</p> <p>Note _____ The actual trace flags and trace level differ according to the scenario and are provided by NetApp technical support.</p> _____
3	Reproduce the issue or scenario to be traced.
4	<p>Run the following command to stop the trace session after the issue is reproduced:</p> <pre>tracelog -stop changelog</pre> <p>The tracelog tool generates an output file Changelog.etl which has the traces in binary format.</p> <p>Send the files to NetApp Technical Support.</p>

What the Open Systems SnapVault space estimator is

The Open Systems SnapVault space estimator enables you to find out if there is enough disk space available on the Open Systems SnapVault primary storage system to perform a backup. When you run this utility on a system that does not have Open Systems SnapVault installed, it provides recommendations on where to install Open Systems SnapVault, its database, and the 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, volume mountpoints, 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. You can obtain the values for other factors, such as ACLs and data streams, from the values that are specified in a configuration file.

Degree of accuracy in the reported space estimates

Although the space estimator provides 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 enough space to back up based on the current Open Systems SnapVault configuration. You can find the results in the snapvault log file, 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 end operations. See [“Failing a backup if insufficient disk space is found”](#) on page 196.

- ◆ 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 189.

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 190.
- ◆ The path, file system, and file exclusion list files
For more information, see [“Configuring backup exclusion lists”](#) on page 75.

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"> ◆ Ignore this field during the built-in operation as this information is obtained from the existing Open Systems SnapVault configuration on the system on which the space estimator is running. ◆ During the stand-alone operation, this field enables the space estimator utility to include checksum sizes in its calculations.
Hist Data enabled	FALSE	<ul style="list-style-type: none"> ◆ Ignore this field during the built-in operation as this information is obtained from the existing Open Systems SnapVault configuration on the system where the space estimator utility is running. ◆ During the stand-alone operation, this field enables the space estimator to include Redundant Array of Independent Disks (RAID) checksum sizes and ACL sizes in its calculations.
VSS required % disk space	15	This field specifies the percentage of disk space that the space estimator does <i>not</i> consider when determining the amount of free disk space for a backup operation. Assume this space 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, consider this field value when calculating free disk space.
Average ACL size	200 bytes	If the space estimator is run on a platform that supports ACLs then this field is used. 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. Space estimator uses this value when you specify the <code>-i</code> option 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 trace file also contains the console output provided in the previous sample.

If the space estimator is run in the built-in mode on a system, it does not display information about 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 138.

The following example is available in a snapvault log file:

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

Installing the space estimator

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

To install the stand-alone space estimator, complete the following step.

Step	Action	
1	If...	Then...
	The Open Systems SnapVault software is installed on your primary storage system	<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"> a. Follow the instructions in Chapter 2, “Installing the Open Systems SnapVault Software,” on page 25 to download the installation package from the NOW site. b. Uncompress the installation package, if needed. c. Locate the <i>svestimator</i> file in the <i>installfiles</i> directory.

Running the space estimator

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

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>svestimator [-o -i -d] root_backup_path</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. The <code>-o</code> option does not consider the disk space consumed by the Open Systems SnapVault installation.</p> <p><code>-i</code> is required when you want the space estimator to consider 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>Note</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 3.0 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 198
- ◆ [“Overview of Virtualization and VMware ESX”](#) on page 200
- ◆ [“Open Systems SnapVault on ESX server”](#) on page 201
- ◆ [“Installing Open Systems SnapVault 3.0 on ESX server”](#) on page 203
- ◆ [“Configuration of Open Systems SnapVault on ESX server”](#) on page 204
- ◆ [“Backup and restore of virtual machines”](#) on page 211
- ◆ [“Open Systems SnapVault support for VMotion”](#) on page 216

VMware terminology

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

Virtual machine	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).
VMware ESX	VMware ESX is a virtualization software that abstracts processor, memory, storage, and networking resources into multiple virtual machines.
VMware vCenter Server	VMware vCenter Server (formerly VirtualCenter) is a management software to manage a group of ESX hosts and the associated virtual machines. vCenter server is the backend and virtual infrastructure client is a user interface.
VMotion	VMotion is a feature that enables you to move running virtual machines from one ESX server to another without interrupting service. VMware vCenter Server activates the VMotion. The vCenter Server centrally coordinates all VMotion activities.
Universal Unique Identifier (UUID)	<p>The UUID is a 128-bit value that is used for unique identification of virtual machines.</p> <p>Example:</p> <pre>00 11 22 33 44 55 66 77-88 99 aa bb cc dd ee ff</pre>
VMware vStorage VMFS	VMware vStorage VMFS (formerly VMFS) is a file system that is used by ESX servers to store virtual machine files. Each virtual machine represents a collection of files under the vStorage VMFS volume.

Virtual Machine Disk (VMDK)

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 Red Hat Enterprise Linux 3.0 operating system used as a management interface to the ESX server.

Overview of Virtualization and VMware ESX

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

VMware ESX 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 3.0 in the service console of ESX server

Before Open Systems SnapVault 3.0, Open Systems SnapVault supported backup within individual virtual machines, with the following limitations:

- ◆ Managing backup agents inside each virtual machine was time-consuming.
- ◆ Backing up the data within a virtual machine, instead of the whole virtual machine.

To address the limitations of running Open Systems SnapVault within each virtual machines, Open Systems SnapVault 3.0 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 registered virtual machine 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 Systems SnapVault is installed:

NDMP port (default value is 10000)

FILESERVER port-10555.

QSMSEVER port-10566.

Installing Open Systems SnapVault 3.0 on ESX server

Installing Open Systems SnapVault in the service console of ESX server

The procedure to install Open Systems SnapVault 3.0 in the service console of ESX server is similar to the procedure to install Open Systems SnapVault 3.0 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 vCenter Server 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 vCenter Server 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 NOW”](#) on page 38.

Configuration of Open Systems SnapVault on 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 is 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>	Configure the key name of the option
<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	Displays all matching options with their values. In this case, the list of possible values is not displayed.
<i>option_key_prefix</i> is specified and it matches an option exactly	Displays the option, its value, and its possible values.
You specify <i>option_value</i>	Verify the option value before it is set. If you do not verify, 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 NORMAL VERBOSE LIBNORMAL LIBVERBOSE]	trace.level.proc_mgr	Trace level of Process Manager
Communication Manager:Trace Level	String [ALWAYS NORMAL VERBOSE LIBNORMAL LIBVERBOSE]	trace.level.comm_mgr	Trace level of Communications Manager
SnapVault Listener:Trace Level	String [ALWAYS NORMAL VERBOSE LIBNORMAL LIBVERBOSE]	trace.level.svlistener	Trace level of svlistener

Stanza name /default value	Value type/range	Options key	Notes
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

Stanza name /default value	Value type/range	Options key	Notes
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 vCenter Server 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 vCenter Server is listening. The default value is 80.
VmWare:HTTPS_PORT	Range (0,65535)	vmware.port.https	The HTTPS port on which vCenter Server is listening. The default value is 443.

Stanza name /default value	Value type/range	Options key	Notes
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>

Stanza name /default value	Value type/range	Options key	Notes
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 represents 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.

Once the transfer completes Open Systems SnapVault deletes the virtual machine Snapshot copy.

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, organize the different files that comprise the virtual machine 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
DISK_x_y	Holds disk files for the hard disk at the virtual device node SCSI (x:y)

Directory	Content
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 that is 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 vCenter Server.

-u *user name* is the user name of the ESX host or the vCenter Server.

-p *password* is the password of the ESX host or the vCenter 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 96.

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:Qtree  
app:vmware:UUID:cfg=current
```

The preceding syntax restores only the vmdk files. 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 qtree 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. Restore only the disk files, not the configuration files.

- ◆ 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 preceding syntax, backed up files 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/vml app:vmware:503f7bac-c758-3401-5613-8482ed7f3451:cfg=original
```

Filer1 is the secondary host name.

/vol/vol0/vml 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/vml /vmfs/nfs/vml
```

Filer1 is the secondary host name.

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

/vmfs/nfs/vml 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, trace files, and bandwidth throttle schedules to this output directory if either the `-q` (for Windows only) or the `-all` (for all platforms) option is specified.

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 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 qtreedid]  
[-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 3.0 is packaged with Open Systems SnapVault 3.0. To check the version number, run the following command:

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

Note

OSSVINFO.exe runs only on Windows systems that have the 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 3.0 on the primary storage system.

Primary system error messages: The following error messages are displayed on the primary system 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"> ◆ Stop the process using port 10000 OR ◆ Change the TCP/IP port that Open Systems SnapVault uses for NDMP connections by changing the NDMP port value in the Open Systems SnapVault Configurator tool (svconfigurator). <p>For information on how to change the port, see “Modifying the NDMP Listen Port setting” on page 68.</p>
	Open Systems SnapVault encountered a network error while reading (writing) data.	Can be caused by network problems or the secondary storage system aborting the transfer.	For more information, check the log files on the secondary storage system.

Error number	Error message	Cause	Solution
	OFM volume cannot be synchronized - volume error. (1)	<p>This error 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"> ◆ Set the “Write inactivity period (seconds)” from the default of 5 to 2. ◆ Set “Synchronization timeout (seconds)” from the default of 60 to 120.
	Cannot back up 'c:\' Failed to create a volume snapshot	<p>This error 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"> ◆ Set the “Write inactivity period (seconds)” from the default of 5 to 2. ◆ Set “Synchronization timeout (seconds)” from the default of 60 to 120.
	Failed to write a hybrid history record	<p>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>	<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>

Error number	Error message	Cause	Solution
	Unexpected read select while no data pending	<p>This error can occur because of the following reasons:</p> <ul style="list-style-type: none"> ◆ The <code>snapvault abort</code> command was issued using the command line on the SnapVault secondary storage system. ◆ There was a network error. 	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.
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.

Error number	Error message	Cause	Solution
2755	Open Systems SnapVault upgrade (2.2 to 2.6) fails	<p>The installation was done in one of the following ways:</p> <ol style="list-style-type: none"> 1. Remote desktop or terminal server is used to connect to the Windows 2000 machine. 2. The installation happens from a mapped network drive. 	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.
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.

Error number	Error message	Cause	Solution
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.	<p>This message can occur because of one of the following conditions:</p> <ul style="list-style-type: none"> ◆ The database is in use by some other process. ◆ You might not have adequate permission. ◆ The database is corrupted. 	Retry after closing other processes and ensure that you have the adequate permission. You might need to reset the relationship to its baseline if the database is corrupted.
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.

Error number	Error message	Cause	Solution
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 error 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.
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.

Error number	Error message	Cause	Solution
3014	An error is encountered while processing checkpoint information.	This error 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.
	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"> ◆ Third party software like antivirus scanners consume a lot of paged pool memory. You could disable such softwares and retry the restore operation. ◆ 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.

Error number	Error message	Cause	Solution
	<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 <code>netstat -an</code> 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 open the history file.	This error occurs in one of the following conditions: <ul style="list-style-type: none"> ◆ The history file is deleted. ◆ The history file is corrupted. ◆ The <code>\tmp</code> directory is missing on the system. 	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.

Error number	Error message	Cause	Solution
	<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"> ◆ Timeout (30000 milliseconds) waiting for the Microsoft Software Shadow Copy Provider service to connect. ◆ The service did not respond to the start or control request in a timely fashion. 	<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"> ◆ Increase the retry count. ◆ Schedule the backups such that multiple drives are not backed up simultaneously.

Error number	Error message	Cause	Solution
	Invalid qtree or Snapshot requested	<p>This error occurs due to the following reasons:</p> <ol style="list-style-type: none"> 1. When the secondary storage system requests for a transfer of a qtree that is not present in the primary storage system. 2. When a secondary storage system requests an updated transfer of a qtree from a Snapshot copy that is not recognized by the primary storage system. <p>Note_____</p> <p>The primary storage system and the secondary storage systems need to maintain a common Snapshot copy for Open Systems SnapVault incremental updated transfers.</p> <hr/> <ol style="list-style-type: none"> 3. A restore transfer is initiated by specifying a destination path on the Open Systems SnapVault primary storage system. 	<p>If the error is because of the first or second reason, ensure that you do not delete any files accidentally present in the Open Systems SnapVault db directory. If the error occurs, restore Open Systems SnapVault db and resync the relationship. For more information about the backup and restore of the db directory, see “Backing up and restoring the Open Systems SnapVault database” on page 140 and “Resynchronizing restored or broken relationships” on page 169.</p> <p>If the error is because of the third reason, the restore operation, you cannot specify a destination path that is involved in another relationship. You can specify a different destination path for the restore, or break the relationship which is using the same primary path.</p>
	<p><path_name> is in the SnapVault temporary file directory</p> <p><path_name> is in the SnapVault database directory</p> <p><path_name> is in the SnapVault trace directory</p>	<p>These errors occur if you are trying to back up the SnapVault db, tmp, or trace directories or their subdirectories as the root of the qtree.</p>	<p>You should not back up the SnapVault db, tmp, or trace directories or their subdirectories as the root of the qtree.</p>

Error number	Error message	Cause	Solution
	Access Denied: Secondary < <i>secondary_name</i> > is not on the primary access list	Open Systems SnapVault allows backup requests from only those secondary storage systems that are present in the QSM Access List. This error occurs if you try to back up from a different secondary storage system.	Add your secondary storage system name to the QSM Access List or clear the Check QSM Access List check box in the svconfigurator GUI.
	Failed to output Database Backup File	Open Systems SnapVault fails to transfer the db file due to some internal errors.	Enable tracing and collect VERBOSE traces for qtree SnapMirror server and contact technical support. As a temporary arrangement, you can disable the Open Systems SnapVault database backup. For more information about database backup, see “Backing up and restoring the Open Systems SnapVault database” on page 140.
	Estimator reports insufficient disk space to complete the operation	Open Systems SnapVault requires temporary space to perform the backup operation. A built-in estimator runs prior to the transfer to verify that sufficient free space exists for the backup to happen. This error occurs when the estimator detects insufficient space.	Run the stand-alone svestimator tool provided in the snapvault\bin folder to check the approximate space requirements for the transfer. Clean up some disk space accordingly and retry the transfer.

Error number	Error message	Cause	Solution
	Have not been supplied a drive name	<p>If the backup source path or the restore destination path is specified without a valid drive letter name, this error occurs.</p> <p>Note_____</p> <p>This error occurs only on Windows.</p> <p>_____</p>	Ensure that the backup source path and the restore destination path starts with a valid Windows drive letter, such as C:\, D:\ , and so on.
	<path_name> includes a reparse point	This error occurs if you try to back up or restore to Windows reparse point.	You should not back up or restore to Windows reparse point.

Error number	Error message	Cause	Solution
	Unexpected error on QSM connection	This message is displayed when Open Systems SnapVault detects connectivity issues while trying to receive any data from the secondary storage system.	This might happen due to connection failures or even transient network issues. Ensure that the network connection between the primary and secondary storage systems is working properly. In case of transient network errors, Open Systems SnapVault backup transfers have a checkpoint restart mechanism to restart and continue the transfer. (You should have higher SnapVault retry counts if the network connectivity is poor). However, if it happens for an Open Systems SnapVault restore transfer, it has to be initiated all over again. If the network connectivity is fine, contact technical support and provide the Open Systems SnapVault side traces and secondary SnapMirror logs.

Error number	Error message	Cause	Solution
	Unexpected send error	This message is displayed when Open Systems SnapVault detects connectivity issues while trying to send data to the secondary storage system.	This might happen due to connection failures or even transient network issues. Ensure that the network connection between the primary and secondary storage systems is working properly. In case of transient network errors, Open Systems SnapVault backup transfers have a checkpoint restart mechanism to restart and continue the transfer. (You should have higher SnapVault retry counts if the network connectivity is poor). However, if it happens for an Open Systems SnapVault restore transfer, it has to be initiated all over again. If the network connectivity is fine, contact technical support and provide the Open Systems SnapVault side traces and secondary SnapMirror logs.

Error number	Error message	Cause	Solution
	Unexpected close getting QSM data	This error occurs when Open Systems SnapVault is waiting to receive some data from the secondary storage system but detects that the socket is closed at the other end. This typically happens when the secondary storage system terminates the connection because of unsuccessful negotiation for the backup transfer. For example, during negotiation, if the secondary storage system detects that it does not have valid SnapVault licenses to receive data from the primary storage system, it terminates the connection, while the primary storage system waits to receive an acknowledgement of the negotiation.	Check the SnapMirror logs on the secondary storage system for the transfer to terminate. Contact technical support if the reason is not clear.
	Failed to bind QSM port (10566) to TCP socket	Open Systems SnapVault listens for incoming backup requests on TCP port 10566. This message appears when Open Systems SnapVault detects that this port is not free or is being used by another application.	Use the <code>netstat -a</code> command to check whether any other application is using this port. Free this port for Open Systems SnapVault to use.

Error number	Error message	Cause	Solution
	Failed to create root directory	This message appears when Open Systems SnapVault fails to create the root directory for the restore operation.	<p>Ensure that Open Systems SnapVault has the necessary permissions to create directories in the restore path. Also ensure that you have sufficient disk space for the restore operation. Collect the QSM verbose traces and contact technical support.</p> <p>Note_____</p> <p>On Windows, Open Systems SnapVault runs as a local system account and on UNIX systems as root account.</p>
	<p>Directory in the wrong phase</p> <p>Restore to <file_path> could not be started: cannot initialize an existing qtree</p>	A restore transfer is initiated by specifying a destination path on the Open Systems SnapVault primary storage system, which is bound to another Open Systems SnapVault relationship.	For the restore operation, you must not specify a destination path that is bound to another relationship. Specify a different destination path for the restore, or break the relationship which is using the same primary path.
	Error building restart or resync files restart or resync will not be available	This message is displayed when Open Systems SnapVault is not able to generate the data needed to do a resync of the restored relationship. This might be mainly due to any internal errors while generating the resync-related data.	<p>Contact technical support after collecting the QSM verbose traces.</p> <p>Note_____</p> <p>If you see this error while doing the restore, the subsequent resync operation fails.</p>

Error number	Error message	Cause	Solution
	Failing resync on break. Configured version is 'n' Filer version is 'm'	This message is displayed when Open Systems SnapVault detects an incompatible qtree SnapMirror version for the resync operation.	Resync is supported only in qtree SnapMirror version 11 and later. Ensure that you have the right version of Data ONTAP. For more information about the Data ONTAP versions, see https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb12138
	Error operation(s) failed < VSS Error ID>	This message is displayed when Open Systems SnapVault encounters VSS-related issues. VSS Snapshot copy is used for backup transfers on the Windows 2003 platform.	For more information about the VSS operation errors, see the Windows event logs. For further troubleshooting, see “ Configuration options for Microsoft Volume Shadow copy Services (VSS) in Open Systems SnapVault ” on page 248 and the MSDN documentation on < VSS Error ID>.
	Cannot restart with different request options	This error occurs when the Open Systems SnapVault restart transfer is requested with options that do not match with the original set of options. Open Systems SnapVault transfer options are specified using -o <name:value> in the <code>snapvault start</code> or <code>snapvault update</code> command.	Ensure that you use the same set of options that was used for the original transfer and retry the transfer.

Error number	Error message	Cause	Solution
	Requested checkpoint restart is not available	<p>During a checkpoint restart, the secondary storage system requests a checkpoint number to restart the transfer. The primary storage system stores a list of files corresponding to each of the checkpoint numbers. If the requested checkpoint file is not available, this error is displayed. This error occurs when the <code>snapvault abort -h</code> command is used from the primary side to abort a transfer and a retry happens for that transfer.</p>	<p>Baseline transfers cannot proceed without this file, and hence might need a rebaseline transfer. Update transfers roll back to the previous Snapshot copy after the specified number of retries. The fresh update transfer works fine.</p>
	missing libsv.dll	<p>This error can occur if the target directory path was specified using forward slashes (/) instead of back slashes (\) in the <code>targetdir</code> path variable.</p> <p>Before Open Systems SnapVault 2.6.1 release, this error occurs due to an unattended installation.</p>	<ol style="list-style-type: none"> 1. Verify that the unattended batch file, <code>unattinstall.bat</code>, was generated using <code>svconfigpackager.exe</code>. 2. Determine if the target directory path is correct (that is, the target directory path is specified using forward slashes (/) in the target directory path variable).

The following warning messages are displayed on the primary system:

Error number	Error message	Cause	Solution
	Failed to get file information for <file_name> <error_string>	This warning message appears when GetFileInformationByHandle Windows API fails.	Try to manually access the file and ensure that there are no issues. See the Microsoft documentation on the <error_string> that is logged. During backup, Open Systems SnapVault skips the file with no information and moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update.
	Failed to get status of <file_name> <error_string>	This warning message appears when the stat operation on the file fails in a UNIX environment.	Try to manually access the file and ensure that there are no issues. See the Microsoft documentation on the <error_string> that is logged. During backup, Open Systems SnapVault skips the file with no information and moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update.

Error number	Error message	Cause	Solution
	<p>Unable to read stream in file <file_name> Not backing up</p> <p>Failed to read stream in file <file_name> Not backing up</p>	<p>This warning message appears when BackupRead Windows API fails.</p>	<p>Try to manually access the file and ensure that there are no issues. This particular file skips the backup process, and Open Systems SnapVault moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update. For 'file in flux' errors, you can close the applications using the file and retry the transfer. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.</p>
	<p>Failed to read stream bytes in file <file_name> Not backing up</p>	<p>This warning message appears when BackupRead Windows API fails. Read the requested number of bytes from a file.</p>	<p>Try to manually access the file and ensure that there are no issues. During backup, Open Systems SnapVault skips the file with no information and moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update. For 'file in flux' errors, you can close the applications using the file and retry the transfer. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.</p>

Error number	Error message	Cause	Solution
	Failed to read stream name in file <file_name> Not backing up	This warning message appears when BackupRead Windows API fails. Read the name of the given stream in the file.	Try to manually access the file and ensure that there are no issues. During backup, Open Systems SnapVault skips the file with no information and moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update. For 'file in flux' errors, you can close the applications using the file and retry the transfer. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.
	File is in flux, or cannot be opened: So skipping it.	This warning message appears when Open Systems SnapVault detects that the file is in flux (currently being modified by another application) or cannot be opened when restarting from a block-level checkpoint in that file.	Try to manually access the file and ensure that there are no issues. During backup, Open Systems SnapVault skips the file with no information and moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update. For 'file in flux' errors, you can close the applications using the file and retry the transfer. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.

Error number	Error message	Cause	Solution
	File <file_name> is in flux, not backed up.	This warning message appears when Open Systems SnapVault detects that the file it is trying to back up is in flux (currently being modified by another application).	Try to manually access the file and ensure that there are no issues. During backup, Open Systems SnapVault skips the file with no information and moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update. For 'file in flux' errors, you can close the applications using the file and retry the transfer. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.
	Failed to open <file_name> <error_string>	This warning message appears when Open Systems SnapVault is not able to open the file it is trying to back up.	Try to manually access the file and ensure that there are no issues. For more information about the <error_string>, see the Microsoft documentation or UNIX manual pages. This particular file skips the backup process, and Open Systems SnapVault moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update.

Error number	Error message	Cause	Solution
	File < <i>file_name</i> > has changed	<p>At the start of the transfer, Open Systems SnapVault scans the entire data set and collects information about the files and directories in the data set. It revisits each of the files to transfer the file contents. This error occurs when Open Systems SnapVault revisits the file and notices any discrepancies between the current file information and the information that is collected during scanning. The error occurs when Open Systems SnapVault notices:</p> <ul style="list-style-type: none"> ◆ Changes in the file index of the file ◆ The file has been converted into a directory 	<p>During backup, Open Systems SnapVault skips the file with no information and moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.</p>
	File < <i>file_name</i> > has been modified	<p>This warning message appears when Open Systems SnapVault detects that the file has been modified (based on the last modified time) when restarting from a block-level checkpoint in that file. When Open Systems SnapVault has to restart from a block-level checkpoint inside a file, it has to ensure that the file has not been modified since the time that the transfer was aborted. This is to ensure file data consistency.</p>	<p>This particular file skips the backup process, and Open Systems SnapVault moves on to the next file. Open Systems SnapVault automatically tries to process and transfer the skipped file in the subsequent update. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.</p>

Error number	Error message	Cause	Solution
	Unable to scan directory <directory_name> <error_string>	This message occurs when Open Systems SnapVault is unable to open any directories present in the given data set. The reason for the failure is given by a call to GetLastError and the same is logged by Open Systems SnapVault. The reasons might be because of insufficient permissions to access the directory, the directory is in use by another process, or an incorrect path.	Try to manually access the path provided in the error message. For more information about the <error_string>, see the Microsoft documentation. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.
	Error scanning directory <directory_name>	This warning message appears when FindNextFile Windows API fails or the opendir or readdir system call on UNIX fails during the directory scan.	Try to manually browse through the directory logged in the warning message. Collect the QSM verbose traces and contact technical support. The use of Snapshot copies (OFM or VSS) for transfers reduces the possibility of such errors.

Secondary storage system error messages

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 source system has denied the request from the destination 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 “Modifying the qtree SnapMirror access list” on page 66.</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 http://now.netapp.com.</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"> ◆ The <code>snapvault abort</code> command has been issued through the command line on the primary system. ◆ The Open Systems SnapVault primary system has failed or rebooted. ◆ A network error has been encountered. 	<p>Check the following logs for more information:</p> <ul style="list-style-type: none"> ◆ The SnapMirror log on the secondary system ◆ The SnapVault log on the primary system

Error number	Error message	Cause	Solution
	<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.

Error number	Error message	Cause	Solution
	<p>SnapMirror: Message from Read Socket : Connection reset by peer</p> <p>SnapVault: destination transfer from <SRC> to <DEST> : 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 <X> data chunk(s) in <DEST> 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 http://now.netapp.com</p>
	<p>Destination transfer from <SRC> to <DEST> : 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><SRC> <DEST> 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>

Error number	Error message	Cause	Solution
	SnapVault: destination transfer from <SRC> to <DEST>: incompatible SnapMirror versions on systems	Due to an ordering error at the destination, it is possible for SnapVault and qtree SnapMirror transfers to fail and display an error message that does not correspond to the error reported on the source side.	For more information, see bug 147982 at http://now.netapp.com .
	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"> ◆ The history file is deleted. ◆ The history file is corrupted. 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.

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/Mount points 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.

If the mountpoint name is not found in the *List of drives/Mountpoints not to Snapshot* list, it takes the Snapshot copy of the volume that is mounted at the mountpoint and backs up the data from the Snapshot copy to the secondary storage system. If the mountpoint name is in the List of drives/mountpoints not to snapshot, then Snapshot of the volume is not taken and the backup is performed from the live file system.

[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.

Error messages	Description
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 28

64-bit systems 30

A

aborted transfer, retrying 172

ACL file 140

Active Directory 12, 54, 145, 146

Allowed Secondary Names window 34

Always setting (for generating debugging information) 70

applications for managing Open Systems

SnapVault 6

AutoSupport

 disabling 82

 enabling 82

B

backing up

 disable database 144

 EFS files 13, 174

 using Open Systems SnapVault for 108

 Windows System State 12

backup

 syntax to backup virtual machines 212

backup exclusion lists 10, 75

backup file, database

(OSSV_DATABASE_BACKUP) 141

backups

 EFS files 13

 scheduling 100

 System State data 149

baseline copy, creating 99

batch installation 54

BLI backups

 checksums file 140

 described 10

 for EFS files 10

 for name-based applications 11

 level, setting 73

 name-based 74

broken relationship

 restoring 170

 resuming 13

built-in space estimator

 description 187

 disabling 196

 enabling 194

 installing 194

C

calculation of free space 187

Changelog minifilter driver 175

 about 175

changelog minifilter driver

 configuration files 178

 datasets favorable 178

 disabling for file system data 180

 enabling for applications data 179

 how 176

 limitations 177

 load status 179

 loading 183

 log file path 182

 management 179

 modifying log files path 182

 setting limit on log files 181

 troubleshooting 184

 uninstalling 184

 unloading 183

 viewing install status 179

 viewing log files count 181

 with other features 178

checkpoint

 file 140

 restart 12

checksum computation, BLI checksum 73

Client-based bandwidth throttling

 about 20

 configuration file 20

command-line utility (svsetstanza) 5, 61

- command-line utility for administration 61
- commands
 - restore 108
 - snapvault modify 172
 - snapvault release 155, 164
 - snapvault restore 6, 103, 142, 144
 - snapvault snap sched 100
 - snapvault start 6, 99, 149, 155, 170
 - snapvault stop 155
 - svconfigpackager 48
 - svestimator 195
 - svinstallcheck 41
 - svpmgr shutdown 42
 - svpmgr startup 43
 - svsetstanza 5, 61
- Compression 13
 - primary storage system 14
- configuration files
 - configure.cfg 61
 - estimator.cfg 189
- Configurator utility
 - description 59
 - General tab, purpose of 60
 - Machine tab, purpose of 59
 - running 64
 - Service tab, purpose of 60
 - SnapVault parameter, modifying 66
 - SnapVault tab, purpose of 60
 - Trace tab, purpose of 60
- configure.cfg 20
- configure.cfg file 61
- configuring SQL Server behavior
 - add database list for local transaction logs 121
 - full database recovery 118
 - how 117
 - local transaction log 119
 - local transaction log backup interval 120
 - transaction log recovery 119
 - transaction logs truncation 122
- copying files to restore 103

D

- Data ONTAP version
 - compression 14

- Data ONTAP version, supported 29
- database files
 - ACL 140
 - backup file (OSSV_DATABASE_BACKUP) 141
 - BLI checksums 140
 - checkpoint 140
 - described 140
 - history 140
- databases
 - backup 142
 - disabling backup 144
 - disk space requirements for 29
 - Open Systems SnapVault 140
- DataFabric Manager, using to manage Open Systems SnapVault 6
- debug files
 - generation of 70
 - stop generating 71
- determining free space 187
- disable 78
- Disabling 78
- disabling 82

E

- EFS files
 - about 13
 - BLI and 10
- error messages 219
- estimating free space 13
- estimator.cfg file
 - contents of 190
 - described 189
- ESX Server
 - definition 198
 - Overview 200
- ESX server 200
- exclusion lists 10

F

- failed transfers, retrying 172
- file exclusion lists
 - described 10
 - for space estimator 189

- location of 75
- file system exclusion
 - definition 75
 - exclude file system 10
- file system exclusion, path 75
- free space estimator. *See* space estimator

G

- General tab (Configurator) 60
- Generate debugging files option 70
- GUI interface for administration 4, 59

H

- history file 140

I

- initial backup (baseline) 99
- installation
 - Allowed Secondary Names window 34
 - batch 54
 - built-in space estimator 194
 - limitations 30
 - on Solaris 36
 - on UNIX and Linux 36
 - requirements, primary systems 26
 - script for unattended install 48
 - Select Installation Folder window 34
 - Setup wizard 32
 - stand-alone space estimator 194
 - unattended, about 48
- interfaces
 - Setup wizard 32
- IPv6 support
 - about 18
 - connection modes 19
 - inet6 19
 - platforms 18
 - setting IPv6 as preferred type 19
 - unspec 19

L

- Libnormal setting 70

- Libverbose setting 70
- licenses 30
- limitations of Open Systems SnapVault 30
- listen port, NDMP 33
- lists, exclusion 10
- log files
 - SnapVault 138
 - space estimator 193
 - trace 193
- log files, SnapVault 60

M

- Machine tab (Configurator utility) 59
- management application, Open Systems SnapVault 6
- messages, error 219
- Microsoft Cluster Services Support 83
 - about 83
 - configuration 85
 - custom cluster resource 85
 - Disabling cluster support 87
 - distributed Open Systems SnapVault database 85
 - enabling cluster support 86
 - how it works 85
 - migration from stand-alone 86
 - OSSVResourceType 85
 - overview 84
 - Protection Manager support 90
 - setting up two-node cluster 88
- Microsoft Cluster Support
 - on Open Systems SnapVault 17
- Microsoft SQL Server Backup and Restore 111
- migrating a volume 156
- MSCS 84

N

- name-based
 - applications, about 11
 - block-level incremental backup, setting 74
- NDMP Account field 67
- NDMP Host ID field 67
- NDMP Host Name field 67
- NDMP Listen Port field 33, 67

NDMP-based management applications 6
NetApp Host Agent
 install 35
20
Normal setting 70

O

open 78
Open Systems SnapVault 2
OSSV. *See* Open Systems SnapVault
OSSV_DATABASE_BACKUP file 141

P

parameters modifying, SnapVault 66
path exclusion lists
 about 10
 described 75
 for space estimator 189
 location of 75
port
 listen, NDMP 33
Port requirements
 FILESERVER 10555 202
 NDMP 10000 202
 QSMSEVER 10566 202
ports
 10000 29, 33, 41, 219
 10566 29
preinstallation requirements 26
prerequisites
 for installation 26
primary storage system
 about 3
primary storage system reporting through
AutoSupport 82
primary system
 administration 4
 controlling secondary system access 66
 installation requirements 26
 restoring 107
 restoring data 104
 svrestore command 104

R

relationships
 broken 170
 deleting 155
 migration 156
 resynchronizing 170
remote installation
 and upgrade 48
 batch 54
reports, problem 138
requirements (installation), for primary systems 26
restart backup from a checkpoint 12
Restore 213
restore command 108
Restore of virtual machine
 restore example 213
 three methods to restore 213
restoring 140
 a complete system 107
 broken relationship 170
 by copying files 103
 EFS files 174
 from tape 108
 methods for 103
 System State data 150
 Windows Systems State 12
resuming broken relationships 13
resynchronizing a relationship 13, 170
retrying a failed transfer 172

S

scheduling backups 100
script for unattended install 48
secondary system
 about 3
 access control 66
 administration 5
 configuring for use with open systems 97
 enabling SnapVault 98
 installation requirements 29
 migrating relationship between 155, 156
security, enabling or disabling 66
Select Installation Folder window 34
Service console

- definition 199
- Service tab (Configurator utility) 60
- Setup Wizard 32
- SnapMirror, using for protecting 162
- SnapVault
 - configuring the secondary storage system 97
 - log files 60, 138
 - primary system parameters, modifying 66
- snapvault modify command 172
- snapvault release command 155, 164
- snapvault restore command 6, 103, 142, 144
- snapvault snap sched command 100
- snapvault start command 6, 99, 149, 155, 170
- snapvault stop command 155
- SnapVault tab (Configurator utility) 60
- snapvault.cfg file (for space estimator), modifying 196
- softlock support 162
- space estimator
 - about 13
 - built-in
 - description 187
 - disabling 196
 - installing 194
 - calculates, how it 187
 - console output 193
 - described 187
 - enabling 194
 - estimator.cfg file 189
 - log files 193
 - path and file exclusion lists 189
 - stand-alone 187, 194
 - utility 29
- SQL Server backup and restore
 - backing up using cli 128, 129
 - backing up using Protection Manager 132
 - configuration files 112, 117
 - configuring 117
 - full database backup 112
 - how 112
 - local transaction log 114
 - mssql plug-in 112
 - mssql-local-Tlog-DBs.cfg 112
 - ossv_mssql.dll 112
 - ossv-mssql.cfg 112

- overview 111
- restore how 115
- restore with an alternate name 115
- restoring using cli 131
- restoring using Protection Manager 135
- supported SQL versions 111
- transaction log backup 113
- transaction log restore 116
- transaction log truncation 114
- stand-alone space estimator
 - described 187
 - installation 194
- status logs 138
- sv_linux_pri license 30
- sv_ontap_sec license 30
- sv_unix_pri license 30
- sv_windows_pri license 30
- svapp
 - excluding SQL Server files 125
 - list of commands 124
 - viewing database list 126
 - viewing SQL instances 125
 - viewing SQL Server database 124
- svcluster 84
- svconfigpackager command (utility) 48
- svestimator command 195
- svinstallcheck command 41
- svpmgr shutdown command 42
- svpmgr startup command 43
- svsetstanza command 5, 61
- system restore, primary 107
- System State data
 - backup 149
 - restore 150

T

- tape, restoring from 108
- tertiary system, data protection with 162
- tombstone lifetime setting 149
- trace file 193
- Trace Level tab (Configurator utility) 60
- troubleshooting 219
- troubleshooting, generating debug files for 70

U

- unattended installation 48
- uninstall, how to 46
- unsupported configurations 30
- upgrade
 - unattended 48
- UUID
 - definition 198

V

- Verbose setting 70
- Virtual machine
 - definition 198
- virtual machine
 - process to backup a virtual machine 211
 - virtual machine backup 212
- Virtualization
 - overview 200
- VMDK
 - definition 199
- VMotion 216

- definition 198
- vmware license 30
- VMX
 - definition 199
- volume migration 156
- volume mountpoint
 - backing up 109
 - backup and restore 109
 - overview 109
 - Protection Manager for backup and restore 109
 - restoring 109
- volume mountpoint data backup and restore 109
- VSS Snapshot copy
 - about 9
 - setting timeout for 78

W

- Web interface for administration 4, 59
- Windows System State 12
- wizard, setup 32