



Technical Report

SnapVault Best Practices Guide

Jeremy Merrill, Amol Chitre, Remington Svarcas, Chris Blackwood, NetApp
February 2010 | TR-3487 | Version 2.0

ABSTRACT

This document is intended to serve as a deployment guide for architecting and deploying SnapVault® in a customer environment. As always, refer to the latest technical publications on the NOW™ (NetApp® on the Web) site for specific updates on processes; Data ONTAP® command syntax; and the latest requirements, issues, and limitations. This document is intended for field personnel who require assistance in architecting and deploying a SnapVault solution. For further information on Open Systems SnapVault, refer to the “OSSV Best Practices Guide” (TR-3466).

TABLE OF CONTENTS

1	INTRODUCTION	4
1.1	INTENDED AUDIENCE	4
1.2	PURPOSE.....	4
1.3	PREREQUISITES AND ASSUMPTIONS	4
1.4	BUSINESS APPLICATIONS	4
2	DETERMINING DATA PROTECTION REQUIREMENTS	4
2.1	THREAT MODELS	5
2.2	USAGE PATTERNS	6
2.3	RESTORE GRANULARITY	7
2.4	RETENTION PERIODS	7
2.5	MEDIA COSTS	7
2.6	LEGAL REQUIREMENTS	8
2.7	EXISTING BACKUP SCHEDULES AND POLICIES	8
3	SNAPVAULT OVERVIEW	9
3.1	HOW SNAPVAULT WORKS.....	10
3.2	BENEFITS OF SNAPVAULT	11
3.3	SNAPVAULT VERSUS SNAPMIRROR: WHAT'S THE DIFFERENCE?	12
3.4	SNAPVAULT MANAGEMENT OPTIONS.....	13
4	CONFIGURING SNAPVAULT	14
5	PROTECTING THE SNAPVAULT SECONDARY	18
6	KNOWN SNAPVAULT BEHAVIORS	18
6.1	TRANSFER OVERHEAD	18
6.2	COMBINING SNAPMIRROR AND SNAPVAULT	19
6.3	PREVENTING SNAPVAULT SNAPSHOT COPIES FROM CYCLING	21
6.4	QUIESCING A SLOW TRANSFER	22
6.5	SINGLE FILE RESTORE.....	22
6.6	INCREMENTAL RESTORE.....	23
6.7	TRADITIONAL VOLUMES VERSUS FLEXIBLE VOLUMES	23
6.8	SIZING VOLUMES ON THE SECONDARY	23
6.9	CONCURRENT TRANSFERS.....	23
6.10	PERFORMANCE EFFECT ON PRIMARY DURING TRANSFER	24
6.11	QUEUEING YOUR TRANSFERS	24
6.12	SNAPVAULT WITHIN A CLUSTERED SYSTEM	24
6.13	SNAPVAULT WITHIN A SINGLE SYSTEM.....	24
6.14	SNAPVAULT AND DEDUPLICATION ON FAS	25
6.15	SNAPVAULT AND VFILER (MULTISTORE) SUPPORT	25

6.16	NDMP MANAGEMENT APPLICATIONS AND DATA ONTAP 7.3 CHANGES	26
7	BEST PRACTICES AND RECOMMENDATIONS	26
7.1	GENERAL BEST PRACTICES	26
7.2	COMMON MISCONFIGURATIONS	27
8	CONCLUSION	29
9	ADDITIONAL RESOURCES	29
	APPENDIX A: LREP DEMO: SEEDING THE SECONDARY USING LREP_READER AND LREP_WRITER WITH SNAPVAULT	29
A.1	AT THE REMOTE OFFICE.....	29
A.2	AT THE DATA CENTER	30
	APPENDIX B: SNAPVAULT AND SNAPMIRROR BUNDLE.....	30
	APPENDIX C: TROUBLESHOOTING SNAPVAULT ERRORS	31
	APPENDIX D: DETERMINING THE RATE OF CHANGE FOR A VOLUME.....	32
	VERSION HISTORY	33

1 INTRODUCTION

This technical report is designed for storage administrators and architects who are already familiar with SnapVault software and are considering deployments for production environments.

1.1 INTENDED AUDIENCE

This document is intended for system administrators, backup administrators, and IT managers who want to benefit from all of the advantages of SnapVault and provide the highest level of protection for their data. It includes a brief overview of SnapVault basics to establish baseline knowledge. After the overview, the guide discusses features, best practices, and deployment of SnapVault.

1.2 PURPOSE

The purpose of this paper is to present a guide for implementing SnapVault technology, addressing step-by-step configuration examples and providing recommendations to assist the reader in designing an optimal SnapVault solution.

1.3 PREREQUISITES AND ASSUMPTIONS

This guide makes the following assumptions:

- The reader has general knowledge of NetApp platforms and products, particularly in the area of data protection.
- The reader has general knowledge of disaster recovery (DR) solutions.

Note: This report is based on features available in Data ONTAP 6.5 and later.

1.4 BUSINESS APPLICATIONS

SnapVault software from NetApp is a reliable and economical way to protect enterprise data and has many significant advantages over traditional backup methods. Although SnapVault can be deployed in configurations designed to emulate the legacy backup methods it replaces, you can realize the full value of the solution by making a significant shift in the way you think about backup and recovery.

The first section includes an overview of SnapVault, focusing on the differences between SnapVault and traditional backup applications. In particular, it covers some of the special benefits that are unique to SnapVault.

Although this document focuses on SnapVault when used to back up data from systems running Data ONTAP, many of the concepts discussed apply equally well to SnapVault backups in heterogeneous storage environments using Open Systems SnapVault software. For more information on using SnapVault to back up data from other operating systems such as Microsoft® Windows® and Sun™ Solaris™, see TR-3466, [“OSSV Best Practices Guide.”](#)

This document complements the “Data ONTAP Data Protection Online Backup and Recovery Guide,” which provides important information, such as detailed procedures for day-to-day operational tasks.

2 DETERMINING DATA PROTECTION REQUIREMENTS

The first step in designing a backup environment is to determine your data protection requirements. There are several questions you need to answer:

- What threats or problems are you protecting your data against?
- What do your users want out of a backup and recovery infrastructure?
- How often do you expect to restore single files or small groups of files?

- How often do you expect to restore entire data sets?
- When a restore is requested, how quickly does it need to be performed? This is known as your recovery time objective (RTO).
- How old is the “most recent backup” allowed to be at any given time? This is known as the recovery point objective (RPO). It is a measure of how much data (expressed in units of time) would be lost if the source data set were destroyed just prior to the next backup; this requirement determines the frequency of backups. In SnapVault, this is measured as “lag time.”
- How frequently do you expect to restore very old data?
- How long should data backups be kept?
- Where is the data located? Is it on NetApp equipment or on another vendor’s storage?

The following sections help you organize your thoughts and determine your requirements; however, knowledge of your data and users is the best tool to help you in this process. If you feel you need to know more about your data or users, you should consider interviewing a sample of your user community to learn more about their backup and recovery needs.

2.1 THREAT MODELS

A variety of threats could alter, destroy, or otherwise interfere with use of your data. In fact, there are so many threats that without unlimited resources it is impossible to defend against all of them. Consider which threats are most likely to occur and which threats would cause the most damage if realized. Your threat model is a concise, detailed list of the threats for which you should prepare.

A threat model might be part of a service-level agreement with your users; it can be viewed as a promise that says, “If any of these bad things occur, our backup and recovery system will protect your data.” You can also develop a threat model to assist in backup planning without including it in your formal service-level agreements.

You need to determine how to mitigate each threat in your threat model. For example, local Snapshot™ copies on a storage system might protect against a user error that deletes a file or group of files, but would not protect against a fire that burns down the building containing the storage system. A synchronous replication system that provides an exact duplicate of a data set at a remote location can protect against the fire, but might not protect against user error if the user action is replicated to the remote site.

There are some broad categories of threat that you should always consider.

• Data Integrity Threats

Some threats cause unintended, unauthorized, accidental, or malicious modification of data. In these cases, any backup copy of the data is acceptable regardless of location. Either a local Snapshot copy on the same storage system or a remote copy of the data on another system serves equally well. There are only two requirements: make the backup (that causes the data integrity problem) before the event, and the backup copy must not be subject to the same threat.

For example, if you are protecting against the possibility that a normal user might accidentally delete a file, either a local Snapshot copy or a backup copy created by backup software and stored on the same disk would provide good protection. In contrast, if you are protecting against an angry user who might deliberately delete the file, a backup copy on the same disk would not be good enough because the user could delete the backup copy as well as the original copy. A local Snapshot copy on the file system would provide enough protection, because Snapshot copies are read-only. If the threat model included a rogue system administrator who might destroy the whole volume (including the Snapshot copies), the situation resembles a media failure threat. SnapVault could be used in this case to make backups on a remote system.

Note: In many cases, replication solutions (which protect against most other types of threats) do not protect against data integrity threats, because the undesirable changes or deletions could be automatically replicated to the backup copy.

- **Media Failures**

Some threats cause or are caused by errors in the storage media, such as:

- Failure of a single disk
- Failure of multiple disks at once
- Bad or unreadable sectors on a disk

Data ONTAP protects data against failures of individual disks or sectors; however, a multiple disk failure of two disks without RAID-DP® or three disks with RAID-DP could still cause data loss if SyncMirror® is not being used. Snapshot images are not sufficient to protect against this category of threat, because they are stored on the same media as the original data.

To fully protect against media failures, a backup copy of the data should be created on separate storage media, either a traditional backup by sending data to tape devices or a more efficient method such as SnapVault.

Using SyncMirror to maintain multiple copies of a volume provides a high level of protection against most types of media failure.

- **Site-Level Disasters**

To protect against some kinds of media failure threats, the backup media must be located some distance away from the source media.

For example, a threat such as a fire or flood might destroy all of the storage media in a building. To protect your data from the threat, the backup media must be preserved as well as the primary storage media.

In traditional tape-based solutions, it is common to ship backup media off-site and store them remotely. Making duplicate copies of the backup data allows one copy to be kept locally for restore purposes, while the other is shipped off-site.

SnapVault provides several superior backup options. SnapVault is directed to a remote SnapVault secondary, or to multiple SnapVault secondaries. The SnapVault secondary can also be backed up to tape and the tapes shipped to a remote location. Adding SnapMirror® to the configuration protects the secondary volumes by mirroring them to another storage system. This configuration also protects you from a site-level disaster.

2.2 USAGE PATTERNS

When planning a backup and recovery system, you should keep in mind the usage patterns of both users and applications. The duty cycle of an application server (when it is busy, when it is idle) influences backup schedules, and the frequency of access or change in a data set can guide you in choosing a backup retention policy.

One key point to remember: when in doubt, restore requirements are more important than backup requirements. Apart from any performance impact a backup might have on a production environment, users are typically not very sensitive about when backups occur or how long they take. However, there is a critical difference to business operations between a restore that takes five minutes and one that takes an hour.

Think about how frequently your users request restores of files from backup media. When they do make such requests, are they asking for the most recent copy, or do they require data from a specific date? Do your users more often request restores of single files and small groups of files, or do they frequently require restores of an entire data set, such as a qtree? Individual file or directory restores are more common in home directory, source code, and engineering data environments. Whole data set and qtree restores are more common in database and application server environments.

2.3 RESTORE GRANULARITY

Although most restores are performed from the most recent backup copy of the data, some situations might require an older copy. For example, suppose that a data corruption problem (caused by a virus, a software bug, or user error) occurs on a Monday afternoon and is not noticed until Wednesday morning. A restore from the Tuesday evening backup would not be acceptable because the backup copy of the data contains the same errors as the current version of the data. In this case, the user wants to restore from the most recent backup prior to the corruption. With SnapVault, you have the ability to schedule backups as frequently as once an hour; the question that arises is how long to keep each hourly backup. If a user requests a restore from three weeks ago, is it important to provide them with a choice between the backup performed at 3 p.m. and the backup performed at 4 p.m.? If so, backup media to retain each hourly backup is quite costly. If not, determine what granularity is required to accommodate user needs.

Note: *Restore granularity* is actually the same concept as RPO (recovery point objective), but relates to restore of data from something other than the most recent backup.

2.4 RETENTION PERIODS

At a certain age, any given piece of user data becomes useless for production needs. The data reaches a point at which it is easier to correct the corruption manually than to reenter the newer data. For example, think of a user's e-mail in-box. In many cases, it would be more desirable to invest a substantial amount of work to remove a virus or correct a data format problem than to recover from a week-old backup copy, due to the inherent data loss and value of the new data received during the week. However, restoring the whole mailbox from a backup made an hour ago would often be acceptable.

For most data sets, you should be able to determine a "maximum age of likely restore," the oldest data you expect a user to request from backup media. This is based on the usefulness of old data for the specific application and the speed with which users or applications are likely to notice bad data. If a problem is noticed quickly, then a restore from a recent copy is more likely. If a problem is not noticed for many weeks, it is quite likely that a weeks-old backup will be required.

Aside from production use, it is sometimes necessary to retain old data for archival or reference purposes. For example, a company might need to restore old source code to determine when a particular bug was introduced into a code line, or an accounting database from several years ago might need to be reviewed to track down a financial inconsistency. In these cases there are usually specific points in time at which the data needs to be preserved, such as at a software release or the end of a business quarter. How long each backup must be retained is less certain.

Many companies have a fixed retention policy for all data sets. However, the cost savings realized by customizing retention policies for each data set are usually worth the time and effort expended to develop them. Furthermore, understanding the differences between user needs and archival or reference needs allows you to provide different service levels for different points in time, potentially saving substantial amounts of money.

2.5 MEDIA COSTS

Keeping many backup copies of a data set for a long time can consume a lot of backup media. Although the cost for any particular piece of media might be low, it is never insignificant when considered in bulk.

SnapVault uses the NetApp WAFL[®] (Write Anywhere File Layout) file system and Snapshot technology to make efficient use of disk space. With Snapshot copies, only changed blocks are stored on disk after the initial backup. SnapVault consumes less space than traditional backup applications that require a full backup and possibly several incrementals that store changed files instead of changed blocks.

Despite this efficient use of space, some companies might have backup retention policies or requirements that would consume too much disk space over time. In these cases it is best to use disk-based backups for most restores and tape for long-term archival. You can accomplish this simply by using dump or an NDMP-enabled backup application to make occasional backups of the SnapVault secondary to tape.

Note: When you dump Open Systems SnapVault data to tape, the dump does not capture the extended attributes, such as encrypted data, sparse files, or UNIX[®] ACLs. Because these attributes are not native to WAFL, SnapVault stores them in hidden metadata directories, which are not transferred during the dump process.

Note: You cannot use these tapes to reestablish a SnapVault relationship.

SnapMirror to tape is an alternative to using dump or NDMP-enabled backup applications. The only way to capture the extended attributes and use the tapes to reestablish the SnapVault relationship is to use SnapMirror to tape to write the Snapshot copies to tape for recovery.

Note: SnapMirror to tape increases the number of Snapshot copies retained because SnapMirror leaves Snapshot copies around to allow future updates.

2.6 LEGAL REQUIREMENTS

In some industries, and with some data sets, there are legal requirements that specify how frequently backups must be performed and/or how long backup copies of data must be kept. Check with your company's legal department to determine whether any such requirements exist.

In addition to specific backup requirements, you can add LockVault™ to provide a solution for environments where data must be retained for a specific period of time. LockVault allows compliance with various regulations, such as SEC17a-4 and Sarbanes-Oxley.

Note: LockVault uses the SnapLock® license. You do not have to purchase a new license.

2.7 EXISTING BACKUP SCHEDULES AND POLICIES

It is useful to present backup schedules, granularity, and retention in the form of a data restore service-level agreement table. For example, a common backup and recovery environment implemented using standard enterprise backup software along with a tape library might look like Table 1.

Table 1) Existing backup schedule.

Data Set	Age of Requested Data	RPO/Granularity	RTO	Backup Method	Protected From...
Individual home directories	13 days or less	1 day	2 hours	Nightly incremental tape backup, stored on-site	Integrity threats, media errors
	Up to a month	1 week	2 hours	Weekly full tape backup, copies on-site and off-site	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2 days	Monthly full tape backup, stored off-site	Integrity threats, media errors, site disasters
Production database	13 days or less	1 day	12 hours	Nightly full tape backup stored on-site	Integrity threats, media errors
	Up to a month	1 week	12 hours	Weekly full tape backup, copies on-site and off-site	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2.5 days	Monthly full tape backup, stored off-site	Integrity threats, media errors, site disasters

After developing a schedule and policies for use with SnapVault, you might come up with a table such as this.

Table 2) Adjusted the SnapVault backup schedule.

Data Set	Age of Requested Data	RPO/Granularity	RTO	Backup Method	Protected From...
Individual home directories	Less than 1 day	1 hour	5 minutes	Local Snapshot copies + off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 7 days	1 day	5 minutes	Local Snapshot copies + off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 13 days	1 day	20 minutes	Off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 3 months	1 week	20 minutes	Off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2 days	Monthly full tape backup, stored off-site	Integrity threats, media errors, site disasters
Production database	1 day or less	1 hour	30 minutes	Hot backup to Snapshot and off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 7 days	1 day	30 minutes	Hot backup to Snapshot and off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 13 days	1 day	2 hours	Off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 3 months	1 week	2 hours	Off-site SnapVault	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2.5 days	Monthly full tape backup, stored off-site	Integrity threats, media errors, site disasters

After you determine the new backup schedules and retention policies, compare them with the legacy schedule. You will see that using SnapVault provides a huge improvement in the service level provided. In the legacy configuration, incremental backups were performed once a day, with full backups once a week. The fastest restore took up to two hours and required the intervention of a backup operator or system administrator. In the SnapVault configuration, incremental backups are performed once an hour. Because SnapVault uses Snapshot technology, each incremental backup can be used as if it were a full backup, and most restores can be performed in minutes or less by end users, without the need for backup operator intervention.

Note: In this scenario, the individual home directories are relatively small (each has a quota of 3GB), and the production database is 50GB.

3 SNAPVAULT OVERVIEW

The following figure shows a simple SnapVault architecture.

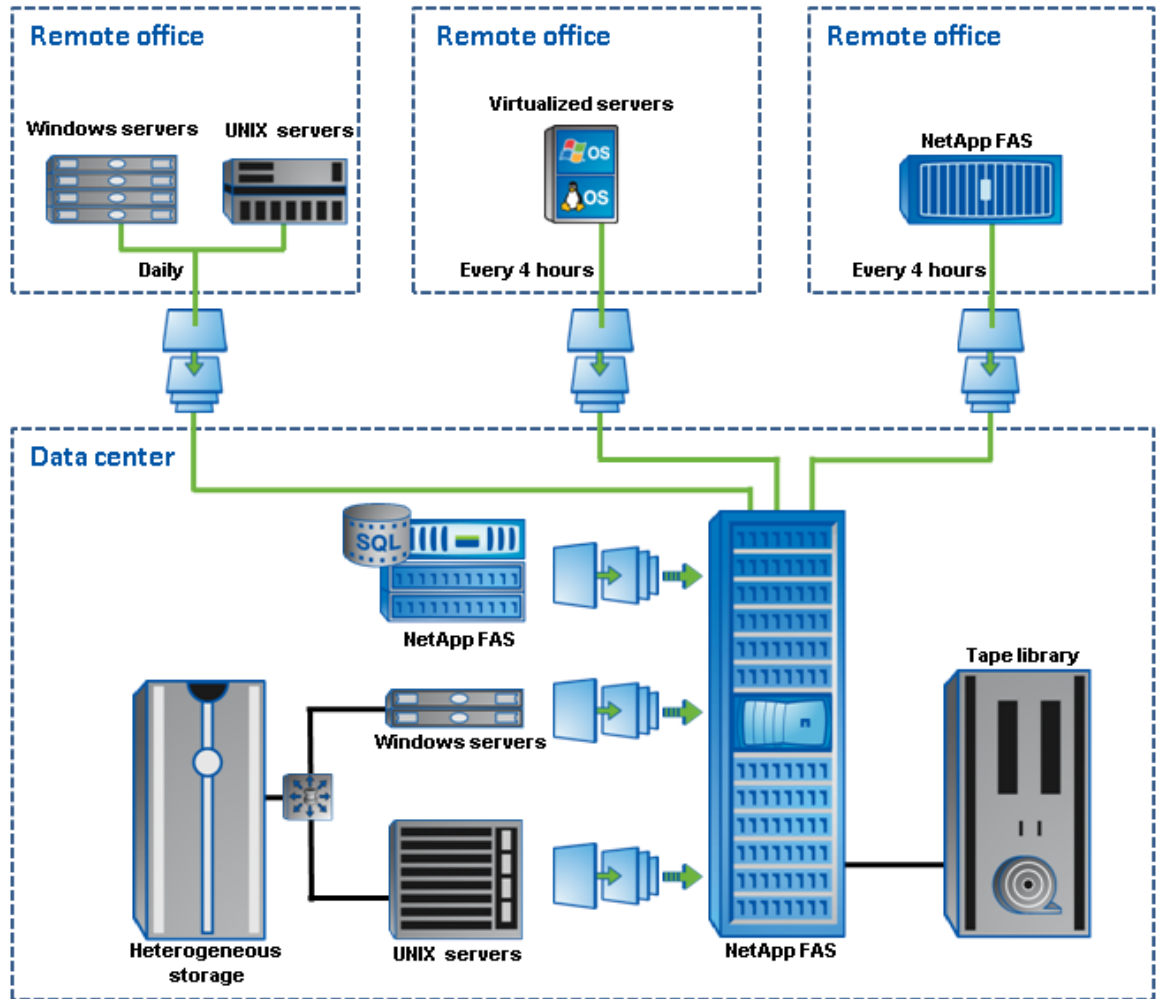


Figure 1) Simple SnapVault implementation.

In this diagram, at both the data center and the remote office, both open systems (with heterogeneous storage) and NetApp storage are backed up to a NearStore[®] system. After the data is on the NearStore system, SnapMirror is used to mirror the data to a remote data center.

3.1 HOW SNAPVAULT WORKS

SnapVault protects data on a SnapVault primary system (called a SnapVault client in earlier releases) by maintaining a number of read-only versions of that data on a SnapVault secondary system (called a SnapVault server in earlier releases) and the SnapVault primary. The SnapVault secondary is always a data storage system running Data ONTAP, such as a NearStore system or a FAS system.

First, a complete copy of the data set is pulled across the network to the SnapVault secondary. This initial, or baseline, transfer might take some time to complete, because it duplicates the entire source data set on the secondary, much like a level-zero backup to tape. Each subsequent backup transfers only the data blocks that have changed since the previous backup.

When the initial full backup is performed, the SnapVault secondary stores the data in a WAFL file system and creates a Snapshot image of the volume for the data that is to be backed up. A Snapshot copy is a read-only, point-in-time version of a data set. SnapVault creates a new Snapshot copy with every transfer, allowing retention of a large number of copies according to a schedule configured by the backup administrator. Each copy consumes an amount of disk space proportional to the differences between it and the previous copy.

Note: If the baseline transfer is a large amount of data, LREP is an option to help seed the secondary storage system. (For more information on LREP, see Appendix A.)

For example, if SnapVault backed up a 100GB data set for the first time, it would consume 100GB of disk space on the SnapVault secondary. Over the course of several hours, users change 10GB of data (not files) on the primary file system. When the next SnapVault backup occurs, SnapVault writes the 10GB of changes to the SnapVault secondary and creates a new Snapshot copy. At this point, the SnapVault secondary contains two Snapshot copies. One contains an image of the file system as it appeared when the baseline backup occurred, and the other contains an image of the file system as it appeared when the incremental backup occurred. The copies consume a combined total of 110GB of space on the SnapVault secondary.

SNAPSHOT COPIES, VOLUMES, AND QTREES

A quota tree, or qtree, is a logical unit used to allocate storage. The system administrator sets the size of a qtree and the amount of data that can be stored in it, but it can never exceed the size of the volume that contains it.

The smallest granularity for SnapVault is a qtree; each qtree can contain different application data, have different users, and meet different scheduling needs. However, the SnapVault Snapshot creations and schedules of a SnapVault transfer by destination volume. Because the scheduling is on a volume level, when you create volumes on the secondary storage system, be sure to group like qtrees (qtrees that have the similar change rates and identical transfer schedules) into the same destination volume.

A volume is a logical storage unit composed of a number of RAID groups. The space available within a volume is limited by the size and number of disks used to build the volume. A Snapshot copy is a read-only, point-in-time version of an entire volume. It contains images of all the qtrees within the volume.

When you start protecting a qtree using the `snapvault start` command, a Snapshot copy is created on the volume that contains the qtree you want to back up. The SnapVault primary reads the image of the qtree from this copy and transfers it to the SnapVault secondary storage system.

Each time a SnapVault incremental backup occurs, the SnapVault primary compares the previous copy with the current copy and determines which data blocks changed and need to be sent to the SnapVault secondary system. The SnapVault secondary system writes these data blocks to its version of the qtree. When all qtrees in the secondary volume have been updated, a Snapshot copy is created to capture and retain the current state of all the qtrees. After this Snapshot copy is created, it is visible for restoring data.

This mechanism effectively combines data from multiple Snapshot copies on multiple primaries into a single copy on the SnapVault secondary system. However, remember that SnapVault does not transfer Snapshot copies; it transfers only selected data from within Snapshot copies.

3.2 BENEFITS OF SNAPVAULT

The following section discusses the benefits of using SnapVault in a production environment for data protection.

INCREMENTAL BACKUPS FOREVER

A full backup copies the entire data set to a backup medium, which is tape in traditional backup applications, or a NearStore system when using SnapVault. An incremental backup copies only the changes in a data set. Because incremental backups take less time and consume less network bandwidth and backup media, they are less expensive. Of course, because an incremental backup contains only the changes to a data set, at least one full backup is required for an incremental backup to be useful.

Traditional backup schedules involve a full backup once per week or once per month and incremental backups each day. Full backups are done so frequently for two reasons:

- **Reliability:** Because a full backup is required to restore from an incremental backup, failure to restore the full backup due to media error or other causes renders all of the incremental backups useless in restoring the entire data set. Tapes used in traditional backup applications are offline storage; you cannot be sure that the data on the tape is readable without placing the tape in a drive and reading from it. Even if each piece of tape is individually read back and verified after being written, it could still fail after being verified, but before being restored.

This problem is usually solved by creating full backups more frequently and by duplicating backup tapes. Duplication of backup tapes serves several purposes, including providing an off-site copy of the backup and providing a second copy in case one copy is bad. However, for certain types of problems it is possible that the bad data will be copied to both sets of tapes.

- **Speed of recovery:** To restore a full data set, a full backup must be restored first and possibly one or more incremental backups. If full backups are performed weekly and incremental backups daily, restores typically involve a level-zero restore and up to six incremental restores. If you perform fewer full backups and more incrementals, restoring a full data set would take considerably longer.

SnapVault addresses both of these issues. It produces backup reliability by storing the backups on disk in a WAFL file system. They are protected by RAID, block checksums, and periodic disk scrubs, just like all other data on a NetApp storage system. Restores are simple because each incremental backup is represented by a Snapshot copy, which is a point-in-time copy of the entire data set, and is restored with a single operation.

For these reasons, only the incremental changes to a data set ever need to be backed up once the initial baseline copy is complete. This reduces load on the source, network bandwidth consumption, and overall media costs.

SELF-SERVICE RESTORES

One of the unique benefits of SnapVault is that users do not require special software or privileges to perform a restore of their own data. Users who want to restore their own data can do so without the intervention of a system administrator, saving time and money. When trying to restore from a SnapVault secondary, connectivity to the secondary must be in place.

Restoring a file from a SnapVault backup is simple. Just as the original file was accessed using an NFS mount or CIFS share, the SnapVault secondary can be configured with NFS exports and CIFS shares. As long as the destination qtrees are accessible to the users, restoring data from the SnapVault secondary is as simple as copying from a local Snapshot copy.

Users can restore an entire data set the same way, assuming that the appropriate access rights are in place. However, SnapVault provides a simple interface to restore a data set from a selected Snapshot copy using the `snapvault restore` command on the SnapVault primary. Starting with Data ONTAP 7.3, there is also the option to perform an incremental restore, where only the blocks required are transferred to the primary system. Details of syntax and procedures for performing such a restore are found in the “Data ONTAP Data Protection Online Backup and Recovery Guide.”

Note: When you use `snapvault restore`, the command prompt returns only when the restore is complete. If the restore needs to be cancelled, press Ctrl-C.

CONSISTENT SECURITY

A common statement in the computer security community is that backups are “a reliable way to violate file permissions at a distance.” With most common backup methods, the backup copy of the data is stored in a format that is usable by anyone with a copy of the appropriate backup software. The backup software can implement access controls, but they cannot be the same as the access controls on the original files.

SnapVault stores backup copies of the data in a WAFL file system, which replicates all of the file permissions and access control lists held by the original data. Users who are not authorized to access a file on the original file system are not authorized to access the backup copies of that file. This allows the self-service restores described earlier to be performed safely.

3.3 SNAPVAULT VERSUS SNAPMIRROR: WHAT’S THE DIFFERENCE?

The following list describes some of the key differences between SnapVault software and the qtree-based SnapMirror feature.

- SnapMirror software uses the same software and licensing on the source appliance and the destination server. SnapVault software has SnapVault primary systems and SnapVault secondary systems, which provide different functionality. The SnapVault primaries are the sources for data that is to be backed up. The SnapVault secondary is the destination for these backups.

Note: As of Data ONTAP 7.2.1, SnapVault primary and SnapVault secondary can be installed on different heads of the same cluster. Data ONTAP 7.3 supports installing both the primary and secondary on a standalone system.

- SnapVault destinations are typically read-only. Unlike SnapMirror destinations, they cannot be made into read-write copies of the data. This means that backup copies of data stored on the SnapVault server can be trusted to be true, unmodified versions of the original data.

Note: A SnapVault destination can be made into read-write with the SnapMirror and SnapVault bundle. For more information, see Appendix B.

- SnapMirror transfers can be scheduled every few minutes; SnapVault transfers can be scheduled at most once per hour.
- Multiple qtrees within the same source volume consume one Snapshot copy each (on the source system) when qtree-based SnapMirror software is used, but consume only one Snapshot copy total when SnapVault software is used.
- The SnapMirror software deletes SnapMirror Snapshot copies when they are no longer needed for replication purposes. The copies are retained or deleted on a specified schedule.
- SnapMirror relationships can be reversed, allowing the source to be resynchronized with changes made at the destination. SnapVault provides the ability to transfer data from the secondary to the primary *only* for restore purposes. The direction of replication cannot be reversed.
- SnapMirror can be used to replicate data only between NetApp storage systems running Data ONTAP. SnapVault can be used to back up both NetApp and open systems primary storage, although the secondary storage system must be a FAS system or a NearStore system.

3.4 SNAPVAULT MANAGEMENT OPTIONS

This section documents the applications that are available for managing SnapVault relationships, transfer schedules, and restores.

DATA ONTAP CLI

One method of managing SnapVault relationships and their transfer schedules uses the Data ONTAP CLI to create SnapVault schedules, manage relationships, and perform updates and restores. In addition, you can abort transfers, stop the relationship, and check the status. The command set differs, depending on whether you are on the primary or the secondary storage system. Figure 2 shows the commands that can go with the snapvault CLI command on the primary storage system.

```
f825-rtp01*> snapvault
The following commands are available; for more information
type "snapvault help <command>"
abort          help          restore        status
destinations   release        snap
```

Figure 2) SnapVault options, primary.

Figure 3 shows the options that go with the snapvault command on the secondary storage system.

```

r100-rtp01*> snapvault
The following commands are available; for more information
type "snapvault help <command>"
abort          help          snap          stop
convert        modify        start         update
destinations   release       status

```

Figure 3) SnapVault options, secondary.

For more help with the snapvault command, type snapvault help <command>. This document focuses on configuring SnapVault using the Data ONTAP CLI.

PROTECTION MANAGER

Unlike homegrown scripts and competitors' products, only Protection Manager takes full advantage of the NetApp APIs and industry standards to deliver a full suite of storage management capabilities for enterprise storage and content delivery infrastructures. Protection Manager enables customers to simplify management and increase the success of backup and recovery operations by providing easy-to-use policies and global monitoring of data protection operations.

NetApp Protection Manager provides higher data protection by eliminating the manual errors associated with configuring data protection in a dynamic environment. Automation and policy-based management reduce the possibility of user errors. This provides a higher level of assurance that the data is protected and hence available.

Protection Manager enables users to consolidate the secondary storage resources into resource pools. These consolidated resource pools are then allocated to several data sets, thus making sure of maximum use of secondary storage resources.

For more information on using Protection Manager, see TR-3560, "[NetApp Protection Manager: Scenario Setup](#)."

COMMVAULT

The CommVault Simpana suite, based on the CommVault Common Technology Engine, provides data protection by managing data throughout its lifecycle using integrated backup and recovery, migration, archiving, replication, and storage management. For more information, visit the CommVault Web site.

Note: In addition to managing SnapVault, CommVault can also manage Open Systems SnapVault.

SYNCSORT

Syncsort Backup Express has been certified for NetApp Data ONTAP and is currently in collaborative development on Data ONTAP 7.0. Fully integrated Open Systems SnapVault management is available with Backup Express. Backup Express includes complete support for NetApp SnapVault, including Open Systems SnapVault management for Windows, Linux®, and UNIX.

Note: In addition to managing SnapVault, Syncsort can also manage Open Systems SnapVault. Syncsort provides its own agent similar to Open Systems SnapVault.

4 CONFIGURING SNAPVAULT

This section provides step-by-step procedures for configuring SnapVault and examples of configurations.

STEP ONE: DETERMINE THE OVERALL BACKUP SCHEDULE

Determine the overall backup schedule that you want to implement. This document uses the schedule shown in Table 2, in section 2.7.

The following examples assume that you are configuring backups for a single FAS system named `fas3050-pri`, using a single NearStore system named `fas3070-sec`. The home directories are in a qtree

on fas3050-pri called /vol/vol1/users; the database is on fas3050-pri in the volume called /vol/oracle and is not in a qtree.

STEP TWO: SCHEDULE SNAPSHOT COPIES ON THE SNAPVAULT PRIMARIES

The following steps occur on the SnapVault primary, fas3050-pri.

1. License SnapVault and enable it.

```
fas3050-pri> license add ABCDEFG
fas3050-pri> options snapvault.enable on
fas3050-pri> options snapvault.access host=fas3070-sec
```

2. Turn off the normal Snapshot schedules, which will be replaced by SnapVault Snapshot schedules.

```
fas3050-pri> snap sched vol1 0 0 0
fas3050-pri> snap sched oracle 0 0 0
```

3. Set up schedules for the home directory hourly Snapshot copies.

```
fas3050-pri> snapvault snap sched vol1 sv_hourly 22@0-22
```

This schedule creates a Snapshot copy every hour, except for 11 p.m. It keeps nearly a full day of hourly copies and, combined with the daily or weekly backups at 11 p.m., makes copies from the most recent 23 hours always available.

4. Set up schedules for the home directory daily Snapshot copies.

```
fas3050-pri> snapvault snap sched vol1 sv_daily 7@23
```

This schedule creates a Snapshot copy once each night at 11 p.m. and retains the seven most recent copies.

The schedules created in steps 3 and 4 give 22 hourly and seven daily Snapshot copies on the source to recover from before needing to access any copies on the secondary. This enables more rapid restores. However, it is not necessary to retain a large number of copies on the primary; higher retention levels are configured on the secondary.

STEP THREE: SCHEDULE SNAPSHOT COPIES ON THE SNAPVAULT SECONDARY

The following steps occur on the SnapVault secondary, fas3070-sec.

1. License SnapVault and enable it.

```
fas3070-sec> license add HIJKLMN
fas3070-sec> options snapvault.enable on
fas3070-sec> options snapvault.access host=fas3050-pri
```

2. Create a FlexVol® volume for use as a SnapVault destination.

```
fas3070-sec> aggr create sv_flex 10
fas3070-sec> vol create vault sv_flex 100g
```

The size of the volume should be determined by how much data you need to store and other site-specific requirements, such as the number of Snapshot copies to retain and the rate of change for the data on the primary FAS system.

Depending on site requirements, you might want to create several different SnapVault destination volumes. You might find it easiest to use different destination volumes for data sets with different schedules and Snapshot copy retention needs.

3. Optional (recommended): Set the Snapshot reserve to zero on the SnapVault destination volume.

```
fas3070-sec> snap reserve vault 0
```

Due to the nature of backups using SnapVault, a destination volume that has been in use for a significant amount of time often has four or five times as many blocks allocated to Snapshot copies as it does to the active file system. Because this is the reverse of a normal production environment, many users find that it is easier to keep track of available disk space on the SnapVault secondary if SnapReserve is effectively turned off.

4. Turn off the normal Snapshot schedules, which will be replaced by SnapVault Snapshot schedules.

```
fas3070-sec> snap sched vault 0 0 0
```

5. Set up schedules for the hourly backups.

```
fas3070-sec> snapvault snap sched -x vault sv_hourly 4@0-22
```

This schedule checks all primary qtrees backed up to the vault volume once per hour for a new Snapshot copy called `sv_hourly.0`. If it finds such a copy, it updates the SnapVault qtrees with new data from the primary and then creates a Snapshot copy on the destination volume, called `sv_hourly.0`.

Note that you are keeping only the four most recent hourly Snapshot copies on the SnapVault secondary. A user who wants to recover from a backup made within the past day has 23 backups to choose from on the primary FAS system and has no need to restore from the SnapVault secondary. Keeping four hourly Snapshot copies on the secondary merely lets you have at least the most recent four backups in the event of a major problem affecting the primary system.

Note: If you do not use the `-x` option, the secondary does not contact the primary and transfer the Snapshot copy. Only a Snapshot copy of the destination volume is created.

6. Set up schedules for the daily backups.

```
fas3070-sec> snapvault snap sched -x vault sv_daily 12@23@sun-fri
```

This schedule checks all primary qtrees backed up to the vault volume once each day at 11 p.m. (except on Saturdays) for a new Snapshot copy called `sv_daily.0`. If it finds such a copy, it updates the SnapVault qtrees with new data from the primary and then creates a Snapshot copy on the destination volume, called `sv_daily.0`.

In this example, you maintain the most recent 12 daily backups, which, combined with the most recent two weekly backups (see step 7), slightly exceeds the requirements shown in Table 2, in section 2.7.

7. Set up schedules for the weekly backups.

```
fas3070-sec> snapvault snap sched vault sv_weekly 13@23@sat
```

This schedule creates a Snapshot copy of the vault volume at 11 p.m. each Saturday for a new Snapshot copy called `sv_weekly.0`. There is no need to create the weekly schedule on the primary. Because you have all the data on the secondary for this Snapshot copy, you will simply create and retain the weekly copies only on the secondary.

In this example, you maintain the most recent 13 weekly backups, for a full three months of online backups.

STEP FOUR: PERFORM THE INITIAL BASELINE TRANSFER

At this point, you have configured schedules on both the primary and secondary systems, and SnapVault is enabled and running. However, SnapVault does not yet know which qtrees to back up or where to store them on the secondary. Snapshot copies will be created on the primary, but no data will be transferred to the secondary.

To provide SnapVault with this information, use the `snapvault start` command on the secondary:

```
fas3070-sec> snapvault start -S fas3050-pri:/vol/vol1/users  
/vol/vault/fas3050-pri_users  
  
fas3070-sec> snapvault start -S fas3050-pri:/vol/oracle/-  
/vol/vault/oracle
```

If you later create another qtree called `otherusers` in the `vol1` volume on `fas3050-pri`, it can be completely configured for backups with a single command:

```
fas3070-sec> snapvault start -S fas3050-pri:/vol/vol1/otherusers
/vol/vault/fas3050-pri_otherusers
```

No additional steps are needed because the Snapshot schedules are already configured on both primary and secondary for that volume.

SPECIAL CASE: DATABASE AND APPLICATION SERVER BACKUPS

Simply scheduling a Snapshot copy on a database volume might not create a safe, consistent image of the database. Most databases, such as Oracle® and DB2 databases, can be backed up while they continue to run and provide service, but they must first be put into a special hot backup mode. Other databases need to be quiesced (which means that they momentarily stop providing service), and some need to be shut down completely, enabling a cold backup.

In any of these cases, you must take certain actions before and after the Snapshot copy is created on the database volume. These are the same steps that you should take for any other backup method, so your database administrators probably already have scripts that perform these functions. Although you could set up SnapVault Snapshot schedules on such a volume and simply coordinate the appropriate database actions by synchronizing the clocks on the storage systems and database server, it is easier to detect potential problems if the database backup script creates the Snapshot copies using the `snapvault snap create` command.

In this example, you want to create a consistent image of the database every four hours, keeping the most recent day's worth of Snapshot copies (six copies), and you want to retain one version per day for a week. On the SnapVault secondary, you will keep even more versions. The first step is to tell SnapVault the names of the Snapshot copies to use and how many copies to keep. No schedule should be specified, because all Snapshot creations will be controlled by the database backup script.

```
fas3050-pri> snapvault snap sched oracle sv_hourly 5@-
```

This schedule creates a Snapshot copy called `sv_hourly` and retains the most recent five copies, but does not specify when to create the copies.

```
fas3050-pri> snapvault snap sched oracle sv_daily 1@-
```

This schedule creates a Snapshot copy called `sv_daily` and retains only the most recent copy. It does not specify when to create the copy.

After this has been done, you must write the database backup script. In most cases, the script has the following structure:

```
[first commands to put the database into backup mode]
rsh fas3050-pri snapvault snap create oracle sv_hourly
[end with commands to take the database out of backup mode]
```

You would then use a scheduling application (such as `cron` on UNIX systems or the Windows Task Scheduler program) to create an `sv_hourly` Snapshot copy each day at every hour other than at 11 p.m. A single `sv_daily` copy would be created each day at 11 p.m., except on Saturday evenings, when an `sv_weekly` copy would be created instead.

In most cases, it is entirely practical to run such a database backup script every hour because the database needs to be in backup mode for only a few seconds while the script creates the Snapshot copy.

SPECIAL CASE: BACKUP OF FCP OR ISCSI LUNS

Backing up logical units (LUNs) used by Fibre Channel Protocol (FCP) or iSCSI hosts presents the same issues as backing up databases. You should take steps to make sure that the Snapshot copies created represent consistent versions of the user data.

If the LUN is being used as raw storage for a database system, then the steps to be taken are *exactly* the same as described in sections 4.1 through 4.4.

If the LUN is being used as storage for a file system, such as UFS, NTFS, or VxFS, the steps to take depend on the file system. Some file systems have commands or APIs to synchronize and quiesce the file

system, while others might require that the file system be unmounted or disconnected prior to creating the Snapshot copy. In some cases, certain logging file systems might not require any action at all, but this is rare.

In addition to the backup steps for the file system, it is important to take any steps required by applications that use the file system as well.

Finally, if you are backing up LUNs using SnapVault, consider turning space reservations on for the SnapVault secondary volume. Enabling space reservation allows writes to the LUN if the amount of data to be retained is greater than the available space in the LUN. For example, if you have a 10GB LUN on the primary FAS system and rewrite all 10GB, the next SnapVault transfer sends all 10GB. The SnapVault transfer does not fail because it uses the 10GB space reservation to complete those writes. SnapVault can't delete the 10GB that was overwritten because it's still required for the previous Snapshot copy.

SCHEDULING TAPE BACKUPS OF SNAPVAULT SECONDARY

Using an NDMP-enabled backup application, a set of scripts, or manual commands, you can dump the data to tape. As an example, the most recent weekly backup (sv_weekly.0) might be dumped to tape at the beginning of each month and sent to an off-site storage facility. This procedure means that an off-site copy to tape is available and that the latest weekly Snapshot copy contains all relevant data.

In planning this step, note that the previous backup procedures kept two years of monthly backups (24 sets of tapes) and one month of weekly backups (five sets of tapes), stored at the off-site tape storage vendor. You might want to reduce expenses by renegotiating with the vendor to store fewer tapes, or you might take the opportunity to store more than two years of monthly backups off-site.

5 PROTECTING THE SNAPVAULT SECONDARY

Although SnapVault is incredibly effective in protecting the data stored on primary storage systems, some sites might also want to take measures to protect against disasters that affect the SnapVault secondary itself.

In a SnapVault environment, the loss or failure of a SnapVault secondary does not affect primary systems any more than does the loss or failure of a tape library in a traditional backup environment. In fact, some data protection continues, because the loss of a SnapVault secondary does not interrupt the process of creating Snapshot copies on the primary systems.

You could simply configure a replacement system in response to a lost or failed SnapVault secondary. This requires restarting backups from each primary qtree, including a complete baseline transfer of each qtree. If the SnapVault secondary is located on the same network as the primaries, this might not be a problem. You can perform periodic backups of the SnapVault secondary to tape with an NDMP-enabled backup application to preserve long-term archive copies of data.

One of the best options is to protect the SnapVault secondary with SnapMirror technology. Simply use volume-based mirroring to copy all of the SnapVault destination volumes (including all Snapshot copies) to another SnapVault secondary at a remote site. If the original SnapVault secondary fails, the extra SnapVault secondary can continue to back up the SnapVault primaries. One other option is to create periodic backups of the SnapVault secondary using the SnapMirror store command to copy the entire volume (including all Snapshot copies) to tape.

6 KNOWN SNAPVAULT BEHAVIORS

The following section will discuss known SnapVault behaviors, which the user should be aware of before implementing SnapVault.

6.1 TRANSFER OVERHEAD

For every transferred inode, the SnapVault primary sends a 4kB header. Also, all changed data is rounded up to 4kB. Thus, a 1-byte file is much more expensive than a 0-byte file. When a file is created, deleted, or renamed, that changes a directory, causing a 4kB header transfer for that directory. If a file or directory is larger than 2MB, an additional 4kB header is transferred for every 2MB.

In addition to the inodes, the SnapVault primary transfers all the changed ACLs for a volume. Unlike all other inodes, ACLs are not associated with a qtree. This increases the number of files or directories that can

share an ACL, but can use extra network bandwidth on qtree SnapMirror. Given the overhead with ACLs, this also causes the baseline transfer to consume more space on the secondary storage system.

6.2 COMBINING SNAPMIRROR AND SNAPVAULT

You can use SnapVault to protect a volume SnapMirror destination (as shown in Figure 4), but there are some things that need to be taken into consideration. Schedules must be managed such that SnapMirror and SnapVault do not run at the same time for a given volume. Otherwise, replication jobs will fail. This is not a concern when using SnapMirror to protect a SnapVault destination.



Figure 4) Protecting a volume SnapMirror destination with SnapVault.

Prior to Data ONTAP 7.3.2, when SnapVault is used to protect a volume SnapMirror (VSM) destination, SnapVault ignores any specified Snapshot copy as part of the snapvault update command (or SnapVault schedule) and uses the most recent VSM-created Snapshot copy.

Starting with Data ONTAP 7.3.2 SnapVault can update from a specific Snapshot copy that resides on a volume SnapMirror destination. In order to accomplish this, the Snapshot copy to be specified needs to be preserved on the volume SnapMirror primary system. Preserving the Snapshot copy on the volume SnapMirror source adds an “acs” soft lock to the Snapshot copy, preventing autodelete from removing the Snapshot copy (when the autodelete commitment level is set to “try”). Preserving SnapVault Snapshot copies can be achieved in two ways.

Note: This functionality requires Data ONTAP 7.3.2 or later on each of the systems.

SNAPVAULT SCHEDULE

Starting with Data ONTAP 7.3.2 SnapVault Snapshot copies on a SnapVault primary system will automatically get preserved/unpreserved whenever a SnapVault schedule is used within Data ONTAP.

If the volume containing the preserved Snapshot copy is replicated to the volume SnapMirror destination system (Figure 4), SnapVault can update from that specific Snapshot copy. The SnapVault update can be scheduled from the SnapVault destination system.

For example, on the volume SnapMirror primary system (system_A):

```
system_A> snapvault snap sched vol1 snapA 4@0-23@mon-fri
```

On the SnapVault secondary system (system_C):

```
system_C> snapvault snap sched -x vol1 snapA 30@0-23@mon-fri
```

The SnapVault baseline (relationship) would need to be established between system_B and system_C.

In this example, SnapVault will update from snapA. In order to accomplish this, system_A and system_B both require a SnapVault primary license. A SnapVault secondary license is required on system_C.

MANUAL PRESERVE

A specific Snapshot copy can be manually preserved on the volume SnapMirror primary system with the `snapvault snap preserve` command. Once this Snapshot copy is preserved and replicated to the

volume SnapMirror destination system, it can then be used in a SnapVault operation. The syntax for the command is as follows:

```
snapvault snap preserve vol_name snap_name [tag_name]
```

Once the Snapshot copy is preserved and replicated with volume SnapMirror, the `snapvault update` command on the SnapVault destination system can be directed to this Snapshot copy using the “-s” flag.

For example, on the volume SnapMirror primary system (system_A):

```
system_A> snapvault snap preserve vol1 snapA
```

On the volume SnapMirror destination system (system_B), update the VSM destination volume, which will transfer the volume containing “snapA” to the SnapMirror destination system:

```
system_B> snapmirror update vol1
```

Finally, on the SnapVault destination system (system_C), data from “snapA” can now be transferred to the SnapVault secondary system:

```
system_C> snapvault update -s snapA /vol/vol1/qt1
```

This method does not require a SnapVault license on the volume SnapMirror primary system (system_A).

Preserved Snapshot copies can be listed with the following command:

```
snapvault snap preservations <vol> [<snapname>]
```

REMOVING PRESERVED SNAPSHOT COPIES

When it is necessary to manually remove a Snapshot copy from the volume SnapMirror primary system of a cascaded relationship (as shown in Figure 4), special consideration must be taken so that a SnapVault base Snapshot copy is not inadvertently deleted. In the event that the SnapVault base Snapshot copy is removed, the next volume SnapMirror update will fail. Figure 5 illustrates this decision process.

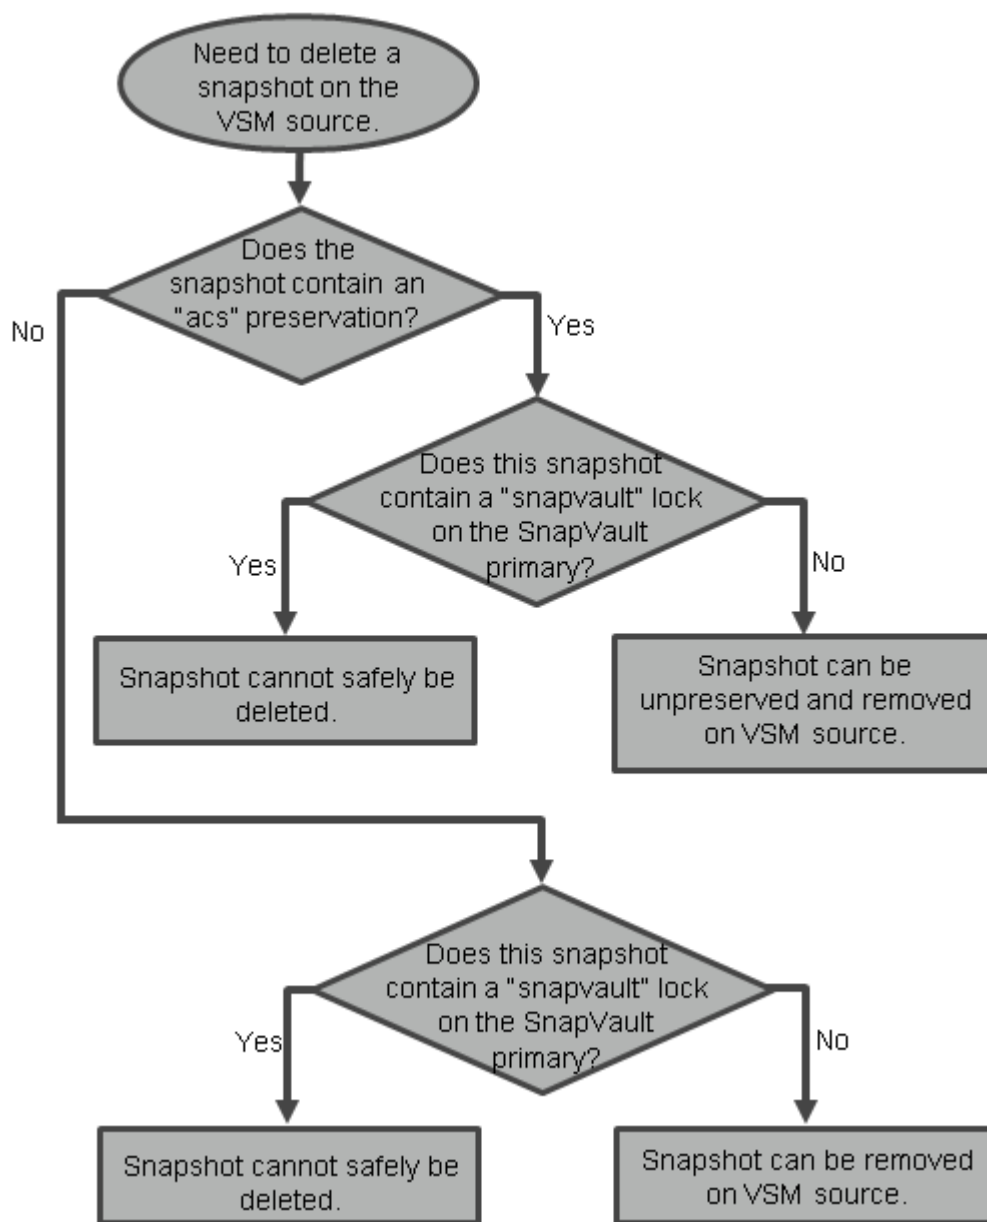


Figure 5) Removing Snapshot copies from the volume SnapMirror primary system.

Snapshot copies can be unpreserved with the following command:

```
snapvault snap unpreserve <vol> <snapname>
```

6.3 PREVENTING SNAPVAULT SNAPSHOT COPIES FROM CYCLING

Starting with Data ONTAP 7.3.2, SnapVault can be configured to prevent Snapshot copies from rolling off when the number of Snapshot copies reaches the retention level set by the schedule. This is useful in situations when an administrator wants to be able to manually intervene prior to any Snapshot copies being rolled off. For example, a SnapVault schedule might be set such that it retains 50 SnapVault Snapshot copies. With SnapVault Snapshot preservation enabled, SnapVault is prevented from creating a new Snapshot copy on the destination system after 50 Snapshot copies is reached (SnapVault updates will fail). The administrator could then take manual steps to archive and delete some of the Snapshot copies. After

the Snapshot copies are removed, SnapVault is able to continue normally. This feature also includes the ability to send EMS and SNMP alerts once the number of Snapshot copies reaches a specified number.

There are two ways to implement this feature. It can be enabled on a per schedule basis, or it can be enabled globally.

For example, to enable this on a per schedule basis on the SnapVault destination system:

```
snapvault snap sched -x -o preserve=on,warn=5 voll sv_daily 50@0-23
```

To enable this on a global basis:

```
options snapvault.preservesnap on
```

Note: individual preservations set on a schedule basis will override the global setting.

6.4 QUIESCING A SLOW TRANSFER

Because SnapVault transfers and schedules are based on the volume, it is important to group qtrees with the same characteristics into the same volume. Obviously, there will be instances where a qtree has an abnormal rate of change, which can't be avoided.

What needs to be avoided is grouping into a volume qtrees that don't have similar characteristics. For example, suppose that you have a volume (/vol/vault) that has 16 qtrees (qtree1 through qtree16). Assume that each qtree has to transfer 1GB worth of changed data, except for qtree4, which has 10GB worth of changed data. This volume is scheduled to complete only one daily transfer, at 11 p.m.

Given this scenario, qtree4 holds up the SnapVault transfer because SnapVault cannot create a Snapshot copy of the destination volume. When you run the snapvault status command on the secondary system all completed qtrees show a status of quiescing. The one qtree that is still being transferred shows a status of transferring and displays the amount of data that has transferred. The other 15 qtrees in the volume do not have an available Snapshot copy until the last qtree in the destination volume has completed its transfer. If there is a slow link between the primary and the secondary system, the 10GB of changed data can take a long time to transfer. This would clearly be a flaw in the layout of the schedule and qtrees to the secondary volume. Figure 6 shows an example of a SnapVault transfer with qtrees in a quiescing state.

r100-rtp01:/vol/sv_dest/qtree1_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/qtree2_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/qtree3_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/qtree4_dest	Snapvaulted	01:05:16	Transferring (248 MB done)
r100-rtp01:/vol/sv_dest/qtree5_dest	Snapvaulted	01:05:16	Quiescing

Figure 6) Example of a transfer in a quiescing state.

Notice that in the example qtree4 is still transferring while all other qtrees are in a quiescing state. It would be a good idea to monitor qtree4 in this SnapVault transfer to see if it continues to cause the other qtrees to be in a quiescing state. The change rate of qtree 4 might not be similar to the other qtrees on the destination volume, and it would make more sense to move this qtree to another volume.

6.5 SINGLE FILE RESTORE

When it is necessary to restore a single file, you cannot use the snapvault restore command. The snapvault restore command allows you to restore the entire qtree contents back to the original primary qtree. After you have restored the entire contents of the qtree, you can choose either to resume the scheduled SnapVault backups (snapvault start -r) or to cancel the SnapVault relationship and the corresponding backups (snapvault release).

For single file restores, use the ndmcopy command in Data ONTAP or Protection Manager (if available) or use CIFS/NFS and copy the file from the Snapshot copy to the correct location.

6.6 INCREMENTAL RESTORE

Before Data ONTAP 7.3, a restore operation had to be performed to a qtree that did not exist on the source. Starting with Data ONTAP 7.3, you can restore to an existing qtree, and only the blocks required to recover to the specified point in time (Snapshot copy) are transferred and stored on the primary system. To use this functionality, both the primary and secondary systems must be running Data ONTAP 7.3 or later. Details of the syntax and procedures for performing such a restore are found in the “Data ONTAP Data Protection Online Backup and Recovery Guide.”

6.7 TRADITIONAL VOLUMES VERSUS FLEXIBLE VOLUMES

When you are setting up the secondary volumes, it's a good idea to use flexible volumes to maximize performance. This allows resizing the volumes as needed, making it easier to retain more Snapshot copies if necessary. In addition, it allows the user to reduce the size of the volume if the number of Snapshot copies that need to be retained changes. The configuration of the secondary volume is independent of the primary, so if the source volumes on the primary are traditional volumes, you can still choose to have the destination volumes be flexible volumes.

In addition to the resizing feature of flexible volumes, FlexClone® and SnapMirror can also be used to make a copy of the SnapVault destination that is writable. FlexClone volumes are a point-in-time copy of the parent volume (SnapVault destination). Changes made to the parent volume after the FlexClone volume is created are not reflected in the FlexClone volume.

6.8 SIZING VOLUMES ON THE SECONDARY

The sizing of volumes on the secondary can vary based on the RTO, RPO, and granularity required, plus the rate of change for the source volume. In addition to the rate of change on the source volumes and/or qtrees, you must consider performance and tape backup factors. Because the rate of change can fluctuate, you should determine the average rate of change for the qtrees and then group like qtrees into the same destination volume. The ability to manipulate the size of a flexible volume makes it an ideal volume type for the SnapVault destination. If the rate of change, retention requirements, or size of the primary changes, you can adjust the size of the destination volume.

Grouping the qtrees by the desired Snapshot schedule and then adding together the disk space requirements for each group of qtrees determines the overall amount of space required by each group. If this results in volume sizes larger than desired (or larger than supported by Data ONTAP) the groups should be split into smaller ones.

Also available in Data ONTAP is the snap delta command. This command reports the rate of change between Snapshot copies. The command compares all copies in a volume, or just the copies specified. Although snap delta can be used to help determine the rate of change for sizing the secondary volume, the future work load should also be considered.

6.9 CONCURRENT TRANSFERS

There is a maximum number of concurrent replication operations for each NetApp system. A storage system might not reach the maximum number of simultaneous replication operations for the following reasons:

- Storage system resources, such as CPU usage, memory, disk bandwidth, or network bandwidth, are taken away from SnapMirror or SnapVault operations.
- Each storage system in a cluster has a maximum number of simultaneous replication operations. If a failover occurs, the surviving storage system cannot process more than the maximum number of simultaneous replication operations specified for that storage system. These can be operations that were scheduled for the surviving storage system, the failover storage system, or both.

Note: Take this limitation into consideration when you are planning SnapMirror or SnapVault replications using clusters.

NetApp systems with a NearStore license are optimized as a destination for QSM and SnapVault replication operations. Replication operations of which the NearStore system is the QSM source, SnapVault source, volume SnapMirror (VSM) source, or VSM destination count *twice* against the maximum number.

For details on the maximum concurrent streams, see the “Data ONTAP Data Protection Online Backup and Recovery Guide” for your version of Data ONTAP.

6.10 PERFORMANCE EFFECT ON PRIMARY DURING TRANSFER

Because a SnapVault transfer is a pull operation, resource usage on the secondary is expected. Remember that a SnapVault transfer also requires resource usage on the primary. This is important because you want to make sure that you don't negatively affect the primary storage system for a SnapVault transfer when setting up SnapVault schedules. Many factors affect how many resources on the primary are used. For this example, suppose that you have two data sets, both 10GB in size. The first data set, dataset1, has approximately a million small files, and the second data set, dataset2, has five files, all 2GB in size. During the baseline transfer, dataset1 requires more CPU usage on the primary or requires a longer transfer time than dataset2. For SnapVault, maximum throughput is generally limited by CPU and disk I/O consumption at the destination.

6.11 QUEUING YOUR TRANSFERS

When scheduling transfers, you must take into consideration the size of the transfer and group like qtrees into the same destination volume. Because scheduling is volume based, not qtree based, poor scheduling causes many issues. There is a limit on the number of concurrent streams supported by the platform you are running. For the list of such limits, refer to the “Data ONTAP Data Protection Online Backup and Recovery Guide.” If you schedule more than the allowed number of concurrent streams, the remaining qtrees to be transferred are queued. However, there is a limit to the number of qtrees that you can queue. You can schedule up to 1,024 transfers with Data ONTAP 7.3 (for both SnapMirror and SnapVault). Any queued transfers in addition to 1,024 are not scheduled for transfer, causing backups to be lost.

Prior to Data ONTAP 7.3, the maximum number of concurrent SnapVault targets supported by a storage system was equal to the maximum number of concurrent SnapVault transfers possible for the system. A SnapVault target is a process that controls the creation of a scheduled SnapVault Snapshot copy on a SnapVault destination volume. There will be a SnapVault target for each SnapVault destination volume that has qtrees being updated.

There is a maximum number of concurrent SnapVault targets for each platform. Only the qtrees in those volumes can be updated concurrently. If the number of SnapVault targets exceeds the limit, the number of concurrent SnapVault transfers might be affected. Despite the maximum number of concurrent SnapVault targets, you can configure SnapVault relationships in as many volumes as required. However, only the qtrees in the limited number of volumes can be updated. See the release notes for the most up-to-date limits.

Note: SnapVault transfers are scheduled based on the volume and not the qtree. Therefore, if a destination volume has 32 qtrees, all 32 qtrees are transferred when the schedule is run.

6.12 SNAPVAULT WITHIN A CLUSTERED SYSTEM

SnapVault in Data ONTAP 7.2.1 includes the ability to use SnapVault within a clustered system. You can therefore install a SnapVault primary license on one head (or controller) of a clustered system and a SnapVault secondary license on the other one. Another type of configuration enabled by this new functionality includes bidirectional backup between two different clustered systems. This feature enables customers to use SnapVault within a cluster from FC drives to SATA drives in the same system. In the event that a cluster fails over, the SnapVault transfers will continue to run, but the maximum number of concurrent transfers is the same as a single head.

6.13 SNAPVAULT WITHIN A SINGLE SYSTEM

SnapVault in Data ONTAP 7.3 includes the ability to use SnapVault within a standalone system. You can therefore install a SnapVault primary and SnapVault secondary on a single controller. This functionality lets you use SnapVault to send the data from FC drives to lower-cost ATA drives and provide local recovery and retention on a single controller. In addition, you could also use two storage systems to act as a SnapVault destination for the other system, enabling bidirectional SnapVault transfers between two storage systems. One limitation is that a SnapVault volume cannot contain both primary and secondary qtrees.

6.14 SNAPVAULT AND DEDUPLICATION ON FAS

Starting with Data ONTAP 7.3, SnapVault and FAS deduplication are integrated to work together on the SnapVault destination system. After an update transfer completes for all qtrees in the target, a “base” Snapshot copy is created. An archive Snapshot copy is also created, which is the Snapshot copy used by the retention policy. After the Snapshot copies are in place, SnapVault then calls FAS deduplication to start. After FAS deduplication completes successfully, SnapVault creates another Snapshot copy and moves the previous archive Snapshot copy to the new archive Snapshot copy. If dedupe fails/aborts, the archive Snapshot copy is not moved and the duplicate blocks will remain locked in the Snapshot copy until that Snapshot copy is recycled.

FAS deduplication might be implemented and scheduled on the SnapVault primary system, but SnapVault and FAS deduplication are not integrated; the schedules are independent of each other. Because SnapVault is replicated at the qtree level, deduplication is not maintained during the transfer.

For more information on FAS deduplication, see [TR-3505, “NetApp Deduplication for FAS Deployment and Implementation Guide.”](#)

6.15 SNAPVAULT AND VFILER (MULTISTORE) SUPPORT

SnapVault can be implemented in a vFiler™ environment, but there are some rules to follow.

SnapVault primary volumes can be owned by the default vFiler unit (vfiler0) or a nondefault vFiler unit. In addition, SnapVault primary volumes can be managed through the default vFiler unit regardless of the ownership. Primary volumes owned by a nondefault vFiler unit can be managed through the default vFiler unit or the nondefault vFiler unit. Table 3 illustrates this.

Table 3) SnapVault primary vFiler support.

Management of SnapVault Primary Volume	Ownership of SnapVault Primary Volume	
	Default vFiler Unit (vfiler0)	Nondefault vFiler Unit
Default vFiler unit (vfiler0)	Yes	Yes
Nondefault vFiler unit	No	Yes

SnapVault secondary volumes can be owned by the default vFiler unit (vfiler0) or a nondefault vFiler unit, but they can only be managed through the default vFiler unit. Table 4 illustrates this.

Table 4) SnapVault secondary vFiler support.

Management of SnapVault Secondary Volume	Ownership of SnapVault Secondary Volume	
	Default vFiler Unit (vfiler0)	Nondefault vFiler Unit
Default vFiler unit (vfiler0)	Yes	Yes
Nondefault vFiler unit	No	No

Note: Protection Manager supports the management of SnapVault relationships for volumes through the default vFiler unit (vfiler0) context only. When using Protection Manager, the following limitations apply for SnapVault relationships involving nondefault vFiler units.

- You can only view SnapVault relationships configured through the default vFiler unit (vfiler0). You cannot view any SnapVault relationships configured through nondefault vFiler units.
- You can configure new SnapVault relationships for a volume only through the default vFiler unit (vfiler0), even if the volume belongs to a nondefault vFiler unit.

6.16 NDMP MANAGEMENT APPLICATIONS AND DATA ONTAP 7.3 CHANGES

Various NDMP-based management applications (Protection Manager, Syncsort, CommVault, Bakbone) provide the ability to monitor and manage your SnapVault and Open System SnapVault transfers. Customers using both an NDMP management application and Data ONTAP 7.3 will not benefit from the increased concurrent streams for Data ONTAP 7.3.

7 BEST PRACTICES AND RECOMMENDATIONS

The following sections discuss best practices and recommendations for implementing SnapVault. This information will be useful when planning the SnapVault deployment.

7.1 GENERAL BEST PRACTICES

You need to be aware of a variety of best practices to perform a successful SnapVault deployment. The sections below discuss in detail some of the general best practices that you should follow.

MONITORING LOGS

When a SnapVault transfer fails, review the log file to determine the problem. All operations (both primary and secondary) are logged to the `/etc/log/snapmirror` log. This log contains various messages that could affect the scheduled transfers. You can see the amount of time a SnapVault session was in the quiescing state or whether it tried to roll back to the last good Snapshot copy in the event of a failed transfer.

SCHEDULING GUIDELINES

Before setting up the SnapVault schedule, you should first gather the following information:

- What is the maximum size to which this qtree is expected to grow?
- What is the estimated rate of change for this qtree in megabytes per day?
- How many days of Snapshot copies should be maintained on the destination volume?

A primary consideration for grouping qtrees within destination volumes is the number of days that Snapshot copies will be retained on the destination volume. The available space on the destination volume is then the secondary criterion.

Another thing to remember when setting up SnapVault schedules is that you need to disable all scheduled Snapshot copies that are invoked by `snap sched` on both the primary and the secondary systems. You should also keep all the Snapshot copy names the same on all primary systems, regardless of the volume, as in the following example:

Snapshot Name	Snapshot Frequency
sv_hourly	Hourly
sv_daily	Daily
sv_weekly	Weekly

In addition to knowing which Snapshot copy should be used, this practice helps to determine the transfer schedule of the specific Snapshot copy.

When scheduling the Snapshot copies, make sure that you add up *all* qtrees in every volume. When adding new schedules for volumes, be sure to take into account the existing schedule. Also, when scheduling, be aware of how many Snapshot copies will be retained for the volume, including copies for SnapVault and SnapMirror.

PRIMARY SNAPSHOT COPY RETENTION

When planning your SnapVault transfer schedule, keep in mind that you can also retain SnapVault Snapshot copies at the primary. Keeping hourly copies on the secondary might not be required, so this would be an ideal situation for primary copy retention. The schedule created in section 4 kept more hourly Snapshot

copies on the primary than on the secondary. The reasoning is that if you need to go back just one or even 10 hours, that copy should be maintained locally. This helps keep the amount of restore time lower than restoring from the secondary. It also helps reduce the number of transfers from the primary to the secondary and makes it easier to maintain a complex schedule. Given this scenario, you would have hourly copies on the primary, but would perhaps transfer only four hourly copies (one every six hours).

CHANGING THE “TRIES” COUNT

The `-t` option ('tries') of the `snapvault` command sets the number of times that updates for the `qtree` should be tried before giving up. The default is two. When the secondary starts creating a Snapshot copy, it first updates the `qtrees` in the volume (assuming that the `-x` option was set on the Snapshot schedule). If the update fails for a reason such as a temporary network outage, the secondary tries the update again one minute later. The `-t` option specifies how many times the secondary should try before giving up and creating a new Snapshot copy with data from all the other `qtrees`. If set to 0, the secondary does not update the `qtree` at all. This is one way to temporarily disable updates to a `qtree`.

If you leave this option at the default, the first attempt to update the secondary counts as the first try. In that case, SnapVault attempts only one more time to update the destination before failing. If there are potential network issues, you should increase the number of tries for the transfer. If the tries count needs to be modified after the relationship has been set up, use the `snapvault modify` command. This is useful when there is a planned network outage.

PRIMARY DATA LAYOUT

With the introduction of FlexVol volumes, there are some alternative ways to lay out data on the SnapVault primary system, which can handle small files and millions of files. If the SnapVault primary system contains millions of files, then using a single FlexVol volume in place of each `qtree`, or even creating just one `qtree` in each volume, helps boost SnapVault performance. This minimizes the amount of scan time before data is sent during the SnapVault transfer. When performing the baseline, the `snapvault start` command is still used, but `“-”` is used in place of the source `qtree` name. The `“-”` signifies that SnapVault backs up all data in the volume that does not reside in a `qtree`. If `qtrees` also exist in that volume, create a separate SnapVault relationship for those `qtrees`.

Note: The `nonqtree` part of the primary storage system volume can be replicated only to a `qtree` SnapVault secondary storage system. The data must be restored to a `qtree` on the primary storage system, but it *cannot* be restored as `nonqtree` data.

If the Data ONTAP CLI is used to perform restores, it's recommended to use one `qtree` inside of the volume. Using this configuration allows restores to function like any other SnapVault restore.

7.2 COMMON MISCONFIGURATIONS

This section examines some common misconfigurations that a user might encounter with SnapVault. You should consider these possible problems during the planning phase in order to achieve a successful SnapVault deployment.

TIME ZONES, CLOCKS, AND LAG TIME

One thing to consider when scheduling is that the SnapVault operations are initiated by the clock on the storage system. For example, on the primary, the Snapshot copies are scheduled by using the `snapvault snap sched` command. When this time is reached, the primary storage system creates its copy. On the secondary, you use the `snapvault snap sched -x` command (`-x` tells the secondary to contact the primary for the Snapshot data) to schedule the SnapVault transfer. This can pose a huge problem with lag times if the clocks are skewed.

The following example shows the output from the `snapvault status` command. Here is the output from the primary storage system:

Source	Destination	State	Lag	Status
f825-rtp01:/vol/sv_vol/qtree1	r100-rtp01:/vol/sv_dest/qtree1_dest	Source	00:02:22	Idle
f825-rtp01:/vol/sv_vol/qtree2	r100-rtp01:/vol/sv_dest/qtree2_dest	Source	00:02:15	Idle
f825-rtp01:/vol/sv_vol/qtree3	r100-rtp01:/vol/sv_dest/qtree3_dest	Source	00:02:08	Idle
f825-rtp01:/vol/sv_vol/qtree4	r100-rtp01:/vol/sv_dest/qtree4_dest	Source	00:01:58	Idle
f825-rtp01:/vol/sv_vol/qtree5	r100-rtp01:/vol/sv_dest/qtree5_dest	Source	00:01:49	Idle

Figure 7) SnapVault status on primary.

And here is the output from the secondary:

f825-rtp01:/vol/sv_vol/qtree1	r100-rtp01:/vol/sv_dest/qtree1_dest	Snapvaulted	-00:03:43	Idle
f825-rtp01:/vol/sv_vol/qtree2	r100-rtp01:/vol/sv_dest/qtree2_dest	Snapvaulted	-00:03:50	Idle
f825-rtp01:/vol/sv_vol/qtree3	r100-rtp01:/vol/sv_dest/qtree3_dest	Snapvaulted	-00:03:57	Idle
f825-rtp01:/vol/sv_vol/qtree4	r100-rtp01:/vol/sv_dest/qtree4_dest	Snapvaulted	-00:04:07	Idle
f825-rtp01:/vol/sv_vol/qtree5	r100-rtp01:/vol/sv_dest/qtree5_dest	Snapvaulted	-00:04:16	Idle

Figure 8) SnapVault status on secondary.

As you can see, the secondary storage system has a negative lag time, because the clock on the primary storage system is ahead of the secondary. In this case, it is only a matter of a couple of minutes, but it could be worse. If the secondary is ahead of the primary, the issue could be even larger. Suppose that the secondary is ahead of the primary by 15 minutes. At 11 p.m., the secondary is scheduled to get the data for the daily Snapshot copy. In this case, when it's 11 p.m. on the secondary, it's only 10:45 on the primary, so the primary hasn't yet created the sv_daily.0 Snapshot copy. This gives a lag time of 23:45 on the secondary, and you are now exposed to potentially losing a day's worth of data. To avoid this problem, be sure to verify that the clocks are in sync with the SnapVault schedule.

Clocks and scheduling also come into play when the primary and secondary are in different time zones.

When the primary and secondary are in different times zones, it is important to remember that the schedules are based on the local clock. Given this scenario, assume that there are two storage systems, one on the East Coast and one on the West Coast (a three-hour difference in time zones). You must make sure that the schedules coincide with the time zone difference, or you will end up with either negative lag times or lag times greater than expected based on the schedule.

MANAGING THE NUMBER OF SNAPSHOT COPIES

With Data ONTAP 6.4 and later, each volume on the SnapVault secondary system can have up to 255 Snapshot copies. SnapVault software requires the use of four Snapshot copies (regardless of the number of qtrees or data sets being backed up), leaving 251 copies for scheduled or manual Snapshot copy creation. In most cases, fewer than 251 copies are maintained due to limitations on available disk space. It is recommended that you do not attempt to retain more than 250 total Snapshot copies of a volume. With improper scheduling, this limit can quickly be reached on the secondary because SnapVault creates a Snapshot copy of the volume after every transfer. Again, it's important to make sure that the qtrees within a SnapVault destination have the same characteristics to avoid reaching the 250-copy limit.

VOLUME TO QTREE SNAPVAULT

When issuing the snapvault start command, you are not required to specify a qtree name for the source; however, this practice is not recommended. This type of relationship increases the performance of the SnapVault transfer, but it also increases the time it takes to perform a backup. Since you must specify a qtree for the SnapVault destination, an entire volume then resides in a qtree on the destination. When it's time for the restore using the Data ONTAP CLI, the entire contents of the qtree, which contains all the data from the source volume, are restored to a qtree on the SnapVault primary system. Once the data is restored, you must then manually copy the data back to the appropriate location.

8 CONCLUSION

SnapVault software can be configured and deployed with a minimum amount of time and planning to duplicate the capabilities of legacy backup solutions while still providing several unique advantages. With careful preparation and investigation of user needs, SnapVault can deliver data protection, backup, and recovery capabilities orders of magnitude beyond those available with traditional solutions.

9 ADDITIONAL RESOURCES

- “Data ONTAP Data Protection Online Backup and Recovery Guide,” available on [NOW](#)
- Data protection portal at www.netapp.com/solutions/data_protection.html
- NearStore product information at www.netapp.com/products/nearstore/
- SnapVault product overview at www.netapp.com/solutions/data_protection-br.html
- “Data Protection Strategies for Network Appliance Storage Systems” at www.netapp.com/library/tr/3066.pdf
- “Best Practices Guide for Tape with NearStore Appliances” at www.netapp.com/library/tr/3149.pdf

APPENDIX A: LREP DEMO: SEEDING THE SECONDARY USING LREP_READER AND LREP_WRITER WITH SNAPVAULT

This example customer has a secondary system named `r200` in its data center in Raleigh, North Carolina. `vol1` is the secondary volume. There is a small Windows server named `nt1` in the company's Smithfield, North Carolina, office. A second Windows machine named `nt2` at the data center functions as the `lrep writer`. A Zip drive, drive letter `E`, is moved between `client1` and `client2`.

A.1 AT THE REMOTE OFFICE

First, unpackage `lrep_reader` on the remote server, `nt1`. Navigate to the directory that contains the `lrep` executable and enter the following command:

```
Client1:> lrep_reader -p snapvault_start -O -f Secondary -q  
/vol/dstvol/dstqtree -o /Primary/vol1/lrep_dump/lrep_srcqtree@0  
Primary:/vol/srcvol/srcqtree
```

Examining one argument at a time:

-p `snapvault_start` = use SnapVault protocol

-O = disable OSSV

-f `Secondary` = the final destination

-q `/vol/dstvol/dstqtree` = the full path on the final destination

-o `/Primary/vol1/lrep_dump/lrep_secqtree@0` = the location where LREP writes the data, a name for the file that is created, @number of 2GB files (0=infinite) * number of 2GB files created. This feature allows you to span multiple drives.

Primary:/vol/srcvol/srcqtree = the source you want to mirror

If your portable drive is small, say, 8GB, and your data is 12GB, and you have the option of connecting two portable drives at `E:\` and `F:\`, then you could use the following argument:

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

Note: /Primary/vol1/lrep_dump/lrep_secqtree@4 means to create a maximum of four 2GB files, so it directs lrep_reader to store the first 8GB in E:\.

Note: /Primary/vol1/lrep_dump/lrep_secqtree@0 (0 means unlimited) means to create all files until the end of the stream in /vol/lrep_dump/lrep_secqtree.

A.2 AT THE DATA CENTER

Now move the lrep images to the data center and lrep_writer host, client2, and start lrep_writer:

```
Client2:> lrep_writer -p snapvault_start  
/Secondary/vol2/lrep_dump/lrep_srcqtree
```

Note: Open Systems SnapVault cannot be installed on the lrep_writer machine due to contention for TCP port 10566.

Now start the transfer from Secondary (the secondary system):

```
Secondary> snapvault start -S <IP address of  
client2>:/vol/srcvol/srcqtree /vol/dstvol/dstqtree
```

APPENDIX B: SNAPVAULT AND SNAPMIRROR BUNDLE

SnapVault does not currently have the ability to create a writable destination on the secondary system. However, you can use SnapMirror to convert the SnapVault destination to a SnapMirror destination, making it a typical SnapMirror destination that can be quiesced and broken.

REQUIREMENTS

The minimum version of Data ONTAP is 6.4 or 6.5.

LICENSING

- a) Primary systems: SnapVault primary license
- b) Secondary systems: SnapVault or SnapMirror bundle license

Note: In order to propagate any changes made on the secondary back to the primary, the SnapMirror license must be on the primary storage system.

CONVERTING AND MAKING THE SECONDARY READ/WRITE

Perform the following steps to convert an Open Systems SnapVault or SnapVault secondary backup destination to a usable/writable destination, typically for disaster recovery (DR) situations.

1. Secondary: Turn SnapMirror and SnapVault off.
2. Secondary: Switch to privileged mode (`priv set diag`).
3. Secondary: Convert SnapVault qtree to SnapMirror qtree (`snapmirror convert <sec_qtree_path>`).
4. Secondary: Turn SnapMirror on.
5. Secondary: Quiesce the qtree.
6. Secondary: Break the mirror, making it writable.

7. Secondary: Turn SnapVault on.

REESTABLISHING THE RELATIONSHIP

The following steps apply only to storage-system-to-storage-system SnapVault. Because Open Systems SnapVault does not consist of a primary running Data ONTAP, these steps are not used in an Open Systems SnapVault relationship.

To reestablish the storage-system-to-storage-system SnapVault relationship, there are two scenarios.

Scenario 1: Preserve all the changes made to the secondary during the DR period.

1. Primary: Resync the primary qtree (`snapmirror resync <pri_qtree_path>`).
2. Primary: Quiesce the qtree (`snapmirror quiesce <pri_qtree_path>`).
3. Primary: Break the mirror, making it writable.
4. Secondary: Resync the secondary qtree (`snapmirror resync <sec_qtree_path>`).
5. Secondary: Turn SnapMirror and SnapVault off.
6. Secondary: Convert SnapMirror qtree to SnapVault qtree (`snapvault convert <sec_qtree_path>`).
7. Secondary: Turn SnapVault and SnapMirror on.

Scenario 2: Discard all the changes made to the secondary during the DR period.

1. Secondary: Resync the secondary qtree (`snapmirror resync <sec_qtree_path>`).
2. Secondary: Turn SnapMirror and SnapVault off.
3. Secondary: Convert SnapMirror qtree to SnapVault qtree (`snapvault convert <sec_qtree_path>`).
4. Secondary: Turn SnapVault and SnapMirror on.

Storage-system-to-storage-system SnapVault can now update the qtree as if no changes had occurred.

APPENDIX C: TROUBLESHOOTING SNAPVAULT ERRORS

It is important to check the logs on both the primary and secondary when troubleshooting errors with SnapVault. The errors are located in `/etc/logs/snapmirror` on both the primary and secondary storage systems. Here are some of the common errors encountered when running SnapVault displayed either on the console or in the log file.

source contains no new data; suspending transfer to destination

The Snapshot copies on the primary do not contain any new data, so no data is transferred.

destination requested Snapshot that does not exist on the source

The SnapVault secondary has initiated a transfer, but the Snapshot copy doesn't exist on the source. Either the `snapvault` command was entered incorrectly or the Snapshot copy was deleted on the primary.

request denied by source filer; check access permissions on source

To resolve this error, check options `snapvault.access` on the primary. You might see this issue if a new secondary is being configured or if the hostname or IP address of the secondary has changed.

snapvault is not licensed

The license `sv_ontap_pri` or `sv_ontap_sec` is not on the storage system. Input the license key to unlock the `snapvault` commands.

Transfer aborted: service not enabled on the source

This error appears when a SnapVault secondary contacts the primary for the transfer. If there is a SnapVault license on the primary, verify that SnapVault is on with the options `snapvault.enable` command.

snapvault: request while snapvault client not licensed on this filer

This error is displayed on the console of the primary and means that a secondary has requested a SnapVault transfer, but is not currently licensed on the primary. Check the licensing on the primary and the command syntax on the secondary.

APPENDIX D: DETERMINING THE RATE OF CHANGE FOR A VOLUME

The amount of disk space required for a SnapVault destination volume depends on a variety of factors, the most important of which is the rate of change for data in the source volume and/or qtrees.

The backup schedule and the Snapshot schedule on the destination volume both affect disk usage on the destination volume, and rate of change on the source volume is not likely to be constant. It is a good idea to provide a buffer of additional storage capacity above that which seems to be required to accommodate future changes in end-user or application behavior.

If at all possible, estimate the rate of change on source volumes and qtrees based on the historical size of SnapVault data transfers.

When planning for SnapVault deployments, it might be useful to make estimates based on the historical size of incremental tape backups. The network bandwidth used for transferring data between the SnapVault primary and the SnapVault secondary systems should be about the same size as an incremental backup to tape, while the actual amount of disk space used will generally be significantly less.

There are two ways to determine the rate of change for a volume. With Data ONTAP 7.0 and later, use the `snap delta` command to display the rate of change between Snapshot copies. For more information on `snap delta` and how to read the output, see the man page for the `snap` command.

The second way to determine the rate of change for a volume is a manual process. If real historical data is not available, and you are running a version earlier than Data ONTAP 7.0, the easiest way to estimate the rate of data change on a volume is to adjust the volume Snapshot schedule temporarily and use the `df` command to get statistics on disk space usage. Assuming that the source volume is called `sourcevol`, the procedure is as follows:

1. Select an hour during the day when the rate of change for data on the source volume is expected to be at or near peak.
2. Use telnet or a serial console to connect to the source storage system.
3. Type `snap sched sourcevol` to get the current Snapshot schedule configuration. Take note of this information so that you can undo the changes made as part of this procedure.
4. Confirm with the system administrator, application owners, and/or users that turning off Snapshot for a one-hour period is acceptable.
5. Type `snap sched sourcevol 0 0 0` to disable automatic Snapshot copies on the source volume.
6. Type `df /vol/sourcevol ; snap create sourcevol mysnap`.
7. Reading the `df` output, look at the Used column on the line beginning with `/vol/sourcevol/.Snapshot`. Take note of this number.
8. Take note of the number in the Used column on the line beginning with `/vol/sourcevol` (the line without `.Snapshot`).
9. Wait one hour.
10. Type `df /vol/sourcevol` and again take note of the Used column of the `/vol/sourcevol/.Snapshot` line.

11. Again take note of the Used column of the /vol/sourcevol line (the line without .Snapshot).
12. Subtract the number obtained in step 7 from the number obtained in step 10. The result is the number of kilobytes of data that changed on the source volume during the one-hour period.
13. Subtract the number obtained in step 8 from the number obtained in step 11. The result is the number of kilobytes of data created on the source volume during the one-hour period. If this number is less than zero, treat it as zero.
14. Add the numbers obtained in steps 12 and 13 to find the estimated hourly rate of change for the source volume. This is the final result.
15. Type `snap delete sourcevol mysnap` to delete the temporary Snapshot copy created in step 6.
16. Use the `snap sched` command and the information from step 3 to reset the Snapshot schedule to its original values.

VERSION HISTORY

Version 2.0	February 2010
	Added version history
	Updated for Data ONTAP 7.3.2
	Updated Figure 1
	Section 6.2 is now "Combining SnapMirror and SnapVault"
	Expanded section 6.2 to include SnapVault preservation
	Added section 6.3, "Preventing SnapVault Snapshot Copies from Cycling"
	Added section 6.15, "SnapVault and vFiler (MultiStore) Support"

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein must be used solely in connection with the NetApp products discussed in this document.