



SnapVault® Best Practices Guide

Jeremy Merrill, Darrin Chapman, Amol Chitre, Remington Svarcas
Network Appliance, Inc.

July 2006 | TR-3487

Abstract

This document is intended to serve as a deployment guide for architecting and deploying SnapVault in a customer environment. As always, please 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 deploying and architecting a SnapVault solution. For further information on Open Systems SnapVault, please refer to the [“OSSV Best Practices Guide” \(TR 3466\)](#).

Table of Contents

1.0 Introduction.....	6
1.1 Intended Audience	6
1.2 Purpose	6
1.3 Prerequisites and Assumptions.....	6
1.4 Business Applications	6
2.0 Determining Data Protection Requirements.....	6
2.1 Threat Models.....	7
2.2 Usage Patterns.....	9
2.3 Restore Granularity	9
2.5 Media Costs.....	10
2.6 Legal Requirements	11
2.7 Existing Backup Schedules and Policies	11
3.0 SnapVault Overview	13
3.1 How SnapVault Works	13
3.1.1 Snapshots, Volumes, and Qtrees	14
3.2 Benefits of SnapVault.....	15
3.2.1 Incremental Backups Forever	15
3.2.2 Self-Service Restores.....	16
3.2.3 Consistent Security	16
3.3 SnapVault Versus SnapMirror: What's the Difference?	16
3.4 SnapVault Management Options	17
3.4.1 Data ONTAP CLI.....	17
3.4.2 Data Fabric Manager.....	18
3.4.3 Symantec NetBackup.....	19
3.4.4 CommVault.....	19

3.4.5 SyncSort.....	19
4.0 Configuring SnapVault.....	19
4.1 Step One: Determine the Overall Backup Schedule	20
4.2 Step Two: Schedule Snapshot Copies on the SnapVault Primaries.....	20
4.3 Step Three: Schedule Snapshot Copies on the SnapVault Secondary.....	20
4.4 Step Four: Perform the Initial Baseline Transfer	22
4.5 Special Case: Database and Application Server Backups	22
4.6 Special Case: Backup of FCP or iSCSI LUNs	23
4.7 Scheduling Tape Backups of SnapVault Secondary.....	24
5.0 Protecting the SnapVault Secondary	24
6.0 Known SnapVault Behaviors.....	24
6.1 Transfer Overhead	25
6.2 SnapMirror-SnapVault Interlock	25
6.3 Quiescing with Slow Transfer	25
6.4 Single File Restore	26
6.5 Traditional Volumes Versus Flexible Volumes.....	26
6.6 Sizing Volumes on the Secondary	26
6.7 Concurrent Transfers	28
6.7.1 NearStore Option	29
6.8 Performance Impact on Primary During Transfer	31
6.9 Queuing Your Transfers.....	31
6.10 SnapVault within a Clustered System	32
7.0 Best Practices and Recommendations	32
7.1 General Best Practices.....	32
7.1.1 Monitoring Logs.....	32
7.1.2 Scheduling Guidelines.....	32
7.1.3 Primary Snapshot Copy Retention.....	33

7.1.4 Changing the “Tries” Count.....	33
7.1.5 Primary Data Layout.....	33
7.2 Common Misconfigurations.....	34
7.2.1 Time Zones, Clocks, and Lag Time	34
7.2.2 Managing the Number of Snapshot Copies	35
7.2.3 Volume to Qtree SnapVault	35
8.0 Conclusion.....	36
9.0 Additional Resources.....	36
10.0 Terms and Acronyms	36
Appendix A: LREP Demo: Seeding the Secondary Using Irep_reader and Irep_writer with SnapVault.....	40
At the Remote Office	40
At the Data Center	40
Appendix B: SnapVault/SnapMirror Bundle	41
Appendix C: Troubleshooting SnapVault Errors	42
Appendix D: Determining the Rate of Change for a Volume.....	43
Revision History.....	45

List of Figures

Figure 1) Simple SnapVault implementation.	13
Figure 2) SnapVault options, primary.	17
Figure 3) SnapVault options, secondary.	18
Figure 4) Example of a transfer in a Quiescing state.	26
Figure 5) SnapVault status on primary.	34
Figure 6) SnapVault status on secondary.	35

List of Tables

Table 1) Existing backup schedule.	12
Table 2) Adjusted SnapVault backup schedule.	12
Table 3) Concurrent transfer limits.	28
Table 4) Concurrent streams per storage system model.	30
Table 5) NearStore option concurrent transfers.	31

1.0 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 guide is intended to be used by field personnel responsible for architecting and deploying successful SnapVault solutions. A brief overview of SnapVault basics is presented in order to establish baseline knowledge. After the overview, this guide discusses features, best practices, and finally the 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 Network Appliance™ 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 Network Appliance 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, the full value of the solution can be realized only by making a significant shift in the way you think about backup and recovery. To be blunt, SnapVault is so good that it makes many common backup policies and schedules obsolete.

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. 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 the [OSSV Best Practices Guide \(TR 3466\)](#).

This document is not a replacement for the Data ONTAP *Data Protection Guide*, which provides important information not covered here, including detailed procedures for day-to-day operational tasks.

2.0 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 backups of the data 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 do not know enough 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 that should be defended against.

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 happen, 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 may 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 may protect against the fire, but may not protect against the user error if the user action is replicated to the remote site.

There are some broad categories of threat that should always be considered.

- **Data Integrity Threats**

Some threats cause unintended, unauthorized, accidental, or malicious modification of the 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: the backup must be made prior to the event that causes the data integrity problem, 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. On the other hand, 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. A local Snapshot copy on the file system would provide enough protection, however, 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 that 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 not be destroyed at the same time 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 may require an older copy. For example, suppose that a data corruption problem (caused by a virus, 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 capability 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 will be 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.

Completely 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 to 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 may 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 may 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 doesn't capture the extended attributes, such as encrypted data, sparse files, or UNIX® ACLs. Since 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 utilize the tapes for reestablishing the SnapVault relationship is to utilize SnapMirror to tape to write the Snapshot copies to tape for recovery. Note

that 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: In order to utilize LockVault, you are not required to purchase a new license; it's based on the SnapLock 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 the following table.

Data Set	Age of Requested Data	RPO / Granularity	RTO	Backup Method	Protected From...
Individual Home Directories	13 Days or Fewer	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 Fewer	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
--	---------------	---------	----------	---	---

Table 1) Existing backup schedule.

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

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

Table 2) Adjusted SnapVault backup schedule.

Once the new backup schedules and retention policies have been determined, compare them with the legacy schedule. You will clearly see that utilizing SnapVault is a huge improvement to 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 utilizes Snapshot technology, each incremental backup is usable 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 in size (each has a quota of 3GB), and the production database is 50GB.

3.0 SnapVault Overview

The following figure shows a simple SnapVault architecture.

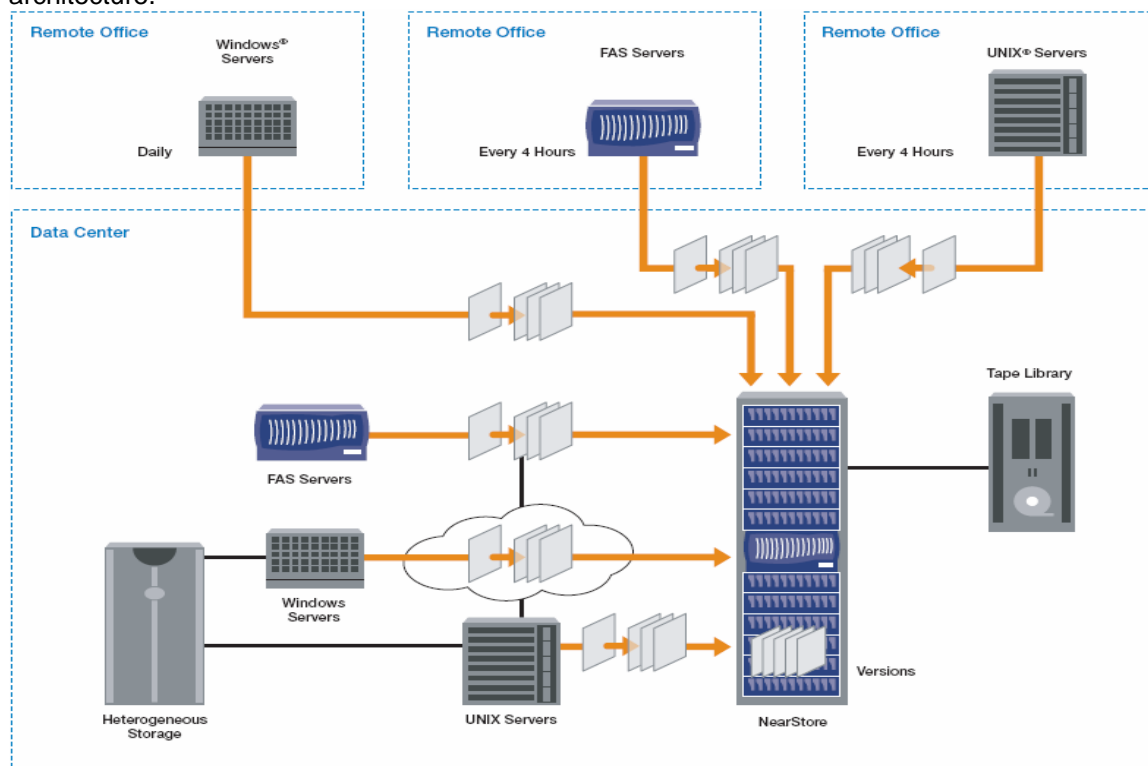


Figure 1) Simple SnapVault implementation.

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. Once 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 may take some time to complete, because it is duplicating 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, and allows 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. (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 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.

3.1.1 Snapshots, 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 have different scheduling needs. However, the SnapVault Snapshot creations and schedules of a SnapVault transfer per volume. Because the scheduling is on a volume level, when you create volumes on the secondary, 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.

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. The SnapVault secondary writes these data blocks to its version of the qtree. When all qtrees in the secondary volume have been updated, a Snapshot copy is taken to capture and retain the current state of all the qtrees. Once this copy has been 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. However, it is important to remember that SnapVault does not transfer Snapshot copies; it only transfers selected data from within copies.

3.2 Benefits of SnapVault

The following section will discuss the benefits of utilizing SnapVault in a production environment for data protection.

3.2.1 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 in order 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. There are two reasons why full backups are done so frequently:

- **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 when 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 taking 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 simply be copied to both sets of tapes.

- **Speed of recovery:** In order 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 ensures 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

3.2.2 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 via 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 an entire data set from a selected Snapshot copy using the `snapvault restore` command on the SnapVault primary. Details of syntax and procedures for performing such a restore are found in the *Data ONTAP Data Protection Guide*.

Note: When you use `snapvault restore`, the command prompt does not return until the restore has completed. If the restore needs to be cancelled, press Ctrl-C.

3.2.3 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; access controls can be implemented by the backup software, 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. Installing both the Primary and Secondary on a standalone system is not yet supported.

- SnapVault destinations are typically read-only. Unlike SnapMirror destinations, they cannot be made into read-write copies of the data. This ensures 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/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.
- SnapMirror Snapshot copies are deleted by the SnapMirror software 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 capability 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.

3.4.1 Data ONTAP CLI

One method of managing SnapVault relationships and their transfer schedules is from the Data ONTAP CLI, from which you can 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. Figure 2 shows the commands that can go with the `snapvault` CLI command on the primary.

```
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.

```
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.

Note: You cannot perform a single file restore with the `snapvault restore` command. For single file restores, you can either mount the NFS mount or CIFS share and use copy and paste; use DFM; or use the `ndmpcopy` command.

3.4.2 Data Fabric Manager

Unlike homegrown scripts and competitors' products, only Data Fabric® Manager takes full advantage of the Network Appliance APIs and industry standards to deliver a full suite of storage management capabilities for enterprise storage and content delivery infrastructures.

To enable DFM management of SnapVault and Open Systems SnapVault, the business continuance option must be purchased and added to the DFM installation. DFM utilizes NDMP to access the primary and secondary systems. TCP port 10000 must be open if firewalls exist. If multiple interfaces exist on the DFM server, NDMP-preferred interfaces can be utilized. DFM can be installed on the following types of servers: Windows 2000, Windows 2003, Solaris 8, Linux® workstation, and Linux server. For the latest installation requirements, be sure to check the *Data Fabric Manager Installation and Upgrade Guide*.

When relationships already exist on the storage systems, DFM allows the user to import these relationships for management, avoiding another costly full baseline transfer. The DFM backup manager recognizes that the relationship exists when it is selected for backup.

All scheduling can be performed in DFM. Multiple schedules can exist, and retention policies can be placed on these schedules. If a schedule already exists from the Data ONTAP command line and DFM is introduced later, limit the scheduling mechanism to only one of the two options. Two separate scheduling mechanisms can cause confusion and may interfere with other backups. The preexisting Data ONTAP schedules cannot be imported, unlike the relationships.

When creating a new relationship in DFM, there is no need to create a qtree name; DFM creates a unique qtree name. It's important to be aware of this when specific naming schemes are in place. DFM is an excellent tool for data restoration. The ability to browse DFM-created backups on the secondary makes restoration much simpler.

Beginning with DFM 3.2, you have the functionality to create pre- and postscripts. These scripts are useful to put a database into hot backup mode before the backup and then release it from hot backup mode upon completion. The scripts must be PERL scripts, and PERL 5.6 or later must be installed on the DFM server. The scripts are installed using a Zip file that contains:

- The script
- An XML file named `package.xml`, which includes:

- Packaging Information (file name, script version, etc.)
- Privileges needed to run the script

By default, the scripts are installed in “script-plugins” in the root of the DFM installation. The scripts can be run manually via DFM or by a schedule. These scripts can be useful to put a database into hot backup mode before an OSSV transfer and then release it from hot backup mode upon completion.

For more details on DFM and DFM pre- and postscripts, refer to the DFM documentation on NOW™ (NetApp on the Web).

3.4.3 Symantec NetBackup

Symantec® NetBackup™ Enterprise Server 6.0 software, combined with the NetBackup Advanced Client option, provides fully integrated support for SnapVault. The NetBackup Administration Console is used to configure, control, and manage SnapVault disk-to-disk backup and recovery operations such as:

- Creation and management of SnapVault relationships between SnapVault primary and SnapVault secondary storage platforms
- Scheduling of Snapshot copies
- Scheduling of SnapVault transfers
- Support for individual file, subdirectory, and entire qtree recoveries
- Oracle® database backup and recovery

Another way to restore data from the Snapshot copies created by NetBackup is to CIFS or NFS mount the share and then drag and drop the files that are needed.

3.4.4 CommVault

The CommVault QiNetix suite, based on the CommVault Common Technology Engine, provides data protection: managing data throughout its lifecycle via 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.

3.4.5 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 2.2. Backup Express includes complete support for NetApp SnapVault, including OSSV management for Windows, Linux, and UNIX.

Note: In addition to managing SnapVault, SyncSort can also manage Open Systems SnapVault. SyncSort provides its own OSSV-like agent.

4.0 Configuring SnapVault

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

4.1 Step One: Determine the Overall Backup Schedule

Determine the overall backup schedule 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 `r200-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.

4.2 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=r200-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 takes 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., ensures that copies from the most recent 23 hours are 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 takes 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 7 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.

4.3 Step Three: Schedule Snapshot Copies on the SnapVault Secondary

The following steps occur on the SnapVault secondary, `r200-sec`.

1. License SnapVault and enable it.

```
r200-sec> license add HIJKLMN
r200-sec> options snapvault.enable on
r200-sec> options snapvault.access host=fas3050-pri
```

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

```
r200-sec> aggr create sv_flex 10
r200-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 may want to create several different SnapVault destination volumes. You may find it easiest to use different destination volumes for data sets with different schedules and Snapshot copy retention needs.

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

```
r200-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.

```
r200-sec> snap sched vault 0 0 0
```

5. Set up schedules for the hourly backups.

```
r200-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 takes 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 ensures that you have at least the most recent four backups in the event of a major problem affecting the primary system.

Note: If you don't use the `-x` option, the secondary does not contact the primary and transfer the Snapshot copy. A Snapshot copy of the destination volume is merely created.

6. Set up schedules for the daily backups.

```
r200-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 takes 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 2 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.

```
r200-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 on the secondary only.

In this example, you maintain the most recent 13 weekly backups, for a full 3 months of online backups.

4.4 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 taken 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:

```
r200-sec> snapvault start -S fas3050-pri:/vol/vol1/users  
/vol/vault/fas3050-pri_users
```

```
r200-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:

```
r200-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.

4.5 Special Case: Database and Application Server Backups

Simply scheduling a Snapshot copy on a database volume may not create a safe, consistent image of the database. Most databases, such as Oracle and DB2, 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 take 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 takes a Snapshot copy called `sv_hourly`, and retains the most recent five copies, but does not specify when to take the copies.

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

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

Once 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 take an `sv_hourly` Snapshot copy each day at every hour other than at 11 p.m. A single `sv_daily` copy would be taken each day at 11 p.m., except on Saturday evenings, when a `sv_weekly` copy would be taken 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.

4.6 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 ensure that the Snapshot copies taken 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 may require that the file system be unmounted or disconnected prior to taking the Snapshot copy. In some cases, certain logging file systems may 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 via SnapVault, consider turning space reservations on for the SnapVault secondary volume. Enabling space reservation allows writes to the LUN in an event where the amount of the data needed 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 utilizes 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.

4.7 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) may be dumped to tape at the beginning of each month and sent to an off-site storage facility. This ensures 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 (5 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.

5.0 Protecting the SnapVault Secondary

Although SnapVault is incredibly effective at protecting the data stored on primary storage systems, some sites may 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 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 may 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 take periodic backups of the SnapVault secondary using the SnapMirror `store` command to copy the entire volume (including all Snapshot copies) to tape.

6.0 Known SnapVault Behaviors

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

6.2 SnapMirror-SnapVault Interlock

When you are using SnapVault in combination with Volume SnapMirror, it is important to understand their relationship with Snapshot. You cannot utilize SnapVault to protect a Volume SnapMirror destination, because SnapVault and SnapMirror both use the same Snapshot copies; they cannot run simultaneously. Schedules must be managed to accommodate the interlock that keeps SnapVault and SnapMirror from stepping on each other. If a SnapMirror session is transferring data, and SnapVault begins, the current SnapMirror transfer is aborted. This issue does not affect utilizing SnapMirror to protect a SnapVault destination. The only way to accomplish this configuration would be to suspend SnapMirror transfers until the SnapVault transfers are complete.

6.3 Quiescing with 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 won't be able to take a Snapshot copy of the destination volume. When a `snapvault status` command is executed on the secondary, all completed qtrees show a status of Quiescing. The one qtree that is still being transferred shows a status of Transferring, and shows 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, this can cause the 10GB of changed data to take quite a while to transfer. This would clearly be a flaw in the layout of the schedule and qtrees to the secondary volume. Figure 4 shows an example of a SnapVault transfer with qtrees in a Quiescing state.

r100-rtp01:/vol/sv_dest/qtrees1_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/qtrees2_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/qtrees3_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/qtrees4_dest	Snapvaulted	01:05:16	Transferring (248 MB done)
r100-rtp01:/vol/sv_dest/qtrees5_dest	Snapvaulted	01:05:16	Quiescing

Figure 4) Example of a transfer in a Quiescing state.

Notice that in the example qtrees4 is still transferring while all other qtrees are in a Quiescing state. It would be a good idea to monitor qtrees4 in this SnapVault transfer to see if it continues to cause the other qtrees to be in a Quiescing state. The change rate of qtrees 4 may not be similar to the others in the destination volume, and it would make more sense to move this qtrees to another volume.

6.4 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 qtrees contents back to the original primary qtrees. Once you have restored the entire contents of the qtrees, 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 `ndmpcopy` command in Data ONTAP or Data Fabric Manager (if available); or use CIFS/NFS and copy the file from the Snapshot copy to the correct location.

6.5 Traditional Volumes Versus Flexible Volumes

When you are setting up the secondary volumes, to maximize performance, it's a good idea to utilize flexible volumes. 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

6.6 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 also needs to be considered.

6.7 Concurrent Transfers

Table 3 shows the maximum number of concurrent transfers allowed for specific storage systems. When taking into account the maximum amount of transfers, remember that each qtree is considered one transfer. If the volume `/vol/vol1/sv_dest` has 16 qtrees and the SnapVault transfer for that volume begins, that is 16 transfers, the maximum number allowed. Always try to spread the transfers out to avoid scheduling too many at one time. For more information on scheduling transfers, see section 6.9.

Model	Number of simultaneous transfers allowed FC drives		Number of simultaneous transfers allowed ATA drives		Notes
	Data ONTAP ≤ 7.0.1	Data ONTAP > 7.0.1	Data ONTAP ≤ 7.0.1	Data ONTAP > 7.0.1	
	F85/F87	4	4	-	
F720/F740/F760	4	4	-	-	
F810/F820/F825	8	8	-	-	
F840/F880	16	16	-	-	
FAS250	4	4	-	-	
FAS270	8	8	-	-	
FAS920	8	8	4	4	1
FAS940/FAS960 /FAS980	16	16	8	8	1
FAS3020	8	16	4	8	1
FAS3050	16	16	8	8	1
FAS6030		24		12	1,3
FAS6070		32		16	1,3
GF270	8	8	8	8	
GF825	8	8	8	8	
GF940/GF960/ GF980	16	16	16	16	
R100	128	128	128	128	1, 2
R150	128	128	128	128	1, 2
R200	128	128	128	128	1, 2

Table 3) Concurrent transfer limits.

Table Notes:

1. **Storage systems that use ATA drives:** A storage system that has any ATA drives, other than NearStore systems, has half the maximum number of simultaneous replication operations that the same storage system has using FC drives. Table 3 lists the maximum number of simultaneous replication operations for storage systems that can use FC drives and ATA drives.
 2. **NearStore systems are optimized as a destination for QSM and SnapVault replication operations.** Therefore, the number of simultaneous replication operations shown in the table represents the total number of simultaneous QSM or SnapVault replication operations of which the storage system is the destination. 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.

Example: Suppose that you have a NearStore system that is the destination of 20 SnapVault relationships and the source of 5 VSM relationships to another storage system. If all the replications occur simultaneously, the NearStore system has 30 replications running concurrently—20 transfers (SnapVault) + 5*2 transfers (VSM).
 3. **The FAS6030 and FAS6070 require a minimum of Data ONTAP 7.2.**
- **Factors that might reduce the maximum number of operations:** 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 the maximum number of simultaneous replication operations listed in Table 3. 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 failed over storage system, or both. For example, each FAS960 in a cluster can run a maximum of 16 simultaneous replication operations. If one FAS960 fails over to the other, it still has a maximum of 16 operations, which can be operations that were scheduled by the surviving FAS960, the failed FAS960, or both.

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

6.7.1 NearStore Option

Background

To enable customers to utilize the FAS storage system as a secondary storage system, a new software license option, called the NearStore Personality option (`nearstore_option`), has been introduced. This license option can be installed only on the FAS3020/3050 systems. This option is supported on Data ONTAP 7.1 and later versions. This license option provides increased concurrent streams when FAS3020/3050 storage systems are used as destinations for

SnapMirror/SnapVault transfers and to enable SnapVault for NetBackup. This license option should not be installed on these storage systems that handle primary application workloads.

Concurrent Replication Limits

The default Data ONTAP behavior *without* the **nearstore_option** license is to maintain a fixed upper limit for concurrent SnapMirror and SnapVault transfers based on the type of disks the storage system has attached. Without the license installed, the total concurrent SnapMirror/SnapVault transfer limits for the FAS30x0 and FAS60x0 are described in Table 4.

Filer Model	Number of Concurrent Streams
FAS3020	16
FAS3050	16
FAS6030	24
FAS6070	32
FAS3020 with ATA drives	8
FAS3050 with ATA drives	8
FAS6030 with ATA drives	12
FAS6070 with ATA drives	16

Table 4) Concurrent streams per storage system model.

The values listed are the combined total of all source *and* destination transfers that can be concurrently run on the given platform. For example, on a system with FC drives, if you have two QSM sources, six QSM destinations, and three VSM sources concurrently running, you can start up only five more SnapVault destinations. Note that a replication operation within the same storage system is considered two concurrent streams.

In a clustered configuration, the maximum stream count shown is *per controller*. So in a clustered FAS3020 configuration with FC drives where both controllers are active, each controller has a maximum limit of 16 streams. Once a controller is taken over by the other controller, a maximum of 16 streams is available to the two controllers. During takeover and giveback, all transfers running on either the controller being taken over or given back are aborted.

Once the nearstore_option license is installed, the storage system switches to NearStore Personality. When the storage systems take on the NearStore personality, Data ONTAP limits the maximum concurrent transfers based on the type of the replication operation. The NearStore Personality also removes the restriction of 8 concurrent streams for ATA drives.

Table 5 describes the maximum streams for each replication operation.

Operation	FAS3020 Maximum Streams	FAS3050 Maximum Streams	FAS6030 Maximum Streams	FAS6070 Maximum Streams
SnapVault Source	16	16	24	32
SnapVault Destination	32	64	96	128

Table 5) NearStore option concurrent transfers.

Note: The streams shown in the table are *not* cumulative.

If you are using the FAS3050 with NearStore Personality as a QSM/SnapVault destination alone, you can have up to 64 *concurrent* streams. Without the `nearstore_option` license on the same FAS3050, you are limited to 16 concurrent streams with FC drives and 8 with ATA drives. This is precisely a scenario where a FAS3050 with NearStore Personality would be beneficial to customers.

In a clustered configuration, all the values of maximum stream count shown in Table 5 are *per controller*. So in a clustered FAS3020 configuration where both controllers are active, each controller has the maximum limit of 32 streams for a QSM destination. Once a controller is taken over by the other controller, a maximum of 32 streams is available to the two controllers. During takeover and giveback, all transfers running on either the controller being taken over or given back are aborted.

6.8 Performance Impact on Primary During Transfer

Given that a SnapVault transfer is a pull operation, there is expected resource utilization on the secondary. Keep in mind that a SnapVault transfer also requires resource utilization on the primary. This is important because you want to make sure that when setting up SnapVault schedules, you don't negatively affect the primary storage system for a SnapVault transfer. There are many factors that affect how many resources on the primary are used. For this example, suppose that you have two datasets, both 10GB in size. The first dataset, `dataset1`, has approximately a million small files, and the second dataset, `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.

6.9 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 limitation on the number of concurrent streams supported by the platform you are running. For the list of limitations, see section 6.7, "Concurrent Transfers." 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 only 600 transfers (this is a total of both SnapMirror and SnapVault). Any queued transfers past 600 are lost, and not scheduled for a transfer, causing backups to be lost.

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

6.10 SnapVault within a Clustered System

SnapVault in Data ONTAP 7.2.1 includes the ability to SnapVault within a clustered system. What this means is that you can 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 that is enabled by this new functionality includes bi-directional backup between two different clustered systems. This feature will enable customers to SnapVault within a clustered 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.

Note: you may not install both a Primary and a Secondary license on the same controller or head.

7.0 Best Practices and Recommendations

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

7.1 General Best Practices

There are many best practices a user needs to be aware of to ensure a successful SnapVault deployment. The sections below will go into detail about some of the general best practices that should be followed.

7.1.1 Monitoring Logs

In most cases, when a SnapVault transfer fails, the problem can be determined by reviewing the log file. 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. Here you can see the amount of time a SnapVault session may have been in the quiescing state, or whether it tried to roll back to the last good Snapshot copy in the event of a failed transfer.

7.1.2 Scheduling Guidelines

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

1. What is the maximum size to which this qtree is expected to grow?
2. What is the estimated rate of change for this qtree in megabytes per day?
3. 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. Also, a best practice is to keep all the Snapshot names the same on all primary systems, regardless of the volume, as in the following example:

Snapshot Name	Snapshot Frequency
<code>sv_hourly</code>	Hourly
<code>sv_daily</code>	Daily
<code>sv_weekly</code>	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 are going to be retained for the volume, including copies for SnapVault and SnapMirror.

7.1.3 Primary Snapshot Copy Retention

When planning your SnapVault transfer schedule, keep in mind that you can also retain SnapVault Snapshot copies at the primary. It may not be a requirement to keep hourly copies on the secondary, so this would be an ideal situation for primary copy retention. In the schedule created in section 4, it was determined to keep more hourly Snapshot copies on the primary than on the secondary. The reasoning is that if you need to go back just 1 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).

7.1.4 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 2. 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, it is recommended that you 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.

7.1.5 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, is advantageous for SnapVault

performance. This minimizes the amount of scan time prior to data being 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, a separate SnapVault relationship must be created for those qtrees.

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

Note: With current versions of Data ONTAP, the maximum number of FlexVol volumes per system is 200.

7.2 Common Misconfigurations

The following are some common misconfigurations that a user may encounter with SnapVault. These are things to be sure to consider prior to the deployment in order to achieve a successful SnapVault deployment.

7.2.1 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 the time for the copy to be created 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/qtrees1	r100-rtp01:/vol/sv_dest/qtrees1_dest	Source	00:02:22	Idle
f825-rtp01:/vol/sv_vol/qtrees2	r100-rtp01:/vol/sv_dest/qtrees2_dest	Source	00:02:15	Idle
f825-rtp01:/vol/sv_vol/qtrees3	r100-rtp01:/vol/sv_dest/qtrees3_dest	Source	00:02:08	Idle
f825-rtp01:/vol/sv_vol/qtrees4	r100-rtp01:/vol/sv_dest/qtrees4_dest	Source	00:01:58	Idle
f825-rtp01:/vol/sv_vol/qtrees5	r100-rtp01:/vol/sv_dest/qtrees5_dest	Source	00:01:49	Idle

Figure 5) SnapVault status on primary.

And here is the output from the secondary.

f825-rtp01:/vol/sv_vol/qtrees1	r100-rtp01:/vol/sv_dest/qtrees1_dest	Snapvaulted	-00:03:43	Idle
f825-rtp01:/vol/sv_vol/qtrees2	r100-rtp01:/vol/sv_dest/qtrees2_dest	Snapvaulted	-00:03:50	Idle
f825-rtp01:/vol/sv_vol/qtrees3	r100-rtp01:/vol/sv_dest/qtrees3_dest	Snapvaulted	-00:03:57	Idle
f825-rtp01:/vol/sv_vol/qtrees4	r100-rtp01:/vol/sv_dest/qtrees4_dest	Snapvaulted	-00:04:07	Idle
f825-rtp01:/vol/sv_vol/qtrees5	r100-rtp01:/vol/sv_dest/qtrees5_dest	Snapvaulted	-00:04:16	Idle

Figure 6) SnapVault status on secondary.

As you can see, the secondary storage system has a negative lag time. This is caused by the fact that the clock on the primary storage system is ahead of the secondary. In this case, it's 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 days' worth of data. To fix this issue, be sure to verify that the clocks are in sync along 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 what is expected based on the schedule.

7.2.2 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 4 Snapshot copies (regardless of the number of qtrees or datasets being backed up), leaving 251 copies for scheduled or manual Snapshot 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 takes 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.

7.2.3 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; however, it also increases the amount of 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 via the Data ONTAP CLI, the entire contents of the qtree, which contains all the data from the source volume, is 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.

8.0 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 some advance preparation and investigation of user needs, SnapVault can deliver data protection, backup, and recovery capabilities orders of magnitude beyond those available with traditional solutions.

9.0 Additional Resources

- *Data ONTAP Data Protection Guide*, available on [NOW](#)
- Data protection portal at http://www.netapp.com/solutions/data_protection.html
- NearStore product information at <http://www.netapp.com/products/nearstore/>
- SnapVault product overview at http://www.netapp.com/solutions/data_protection-br.html
- *Data Protection Strategies for Network Appliance Filers* at <http://www.netapp.com/library/tr/3066.pdf>
- *Best Practices Guide for Tape with NearStore Appliances* at <http://www.netapp.com/library/tr/3149.pdf>

10.0 Terms and Acronyms

ATA

Advanced Technology Attachment, 1994 ANSI standard commonly referred to as AT Attachment for Disk Drives. Originally known as IDE drives, over the years, IDE drives evolved from low-cost/low-performance disk drives to higher-performance/higher-reliability drives to meet the challenging requirements of today's storage requirements.

Backup Catalog

A database hosted on the backup server. It contains metadata about the backed up data. A file catalog might maintain the mapping between backed up files and volumes. A media catalog would maintain the mapping between backup volumes and tape cartridges. A backup catalog is literally the brains of an enterprise storage network and must be carefully protected. Also see backup index.

Backup Index

Some backup applications use this term for a backup catalog. Functionally they are equivalent.

Backup Server

At the heart of the traditional backup network, a backup server hosts a backup index or catalog, directs clients to media servers, and also launches backups. The backup is configured on this server as well, so here is where one would configure backup schedules, add more backup clients, set backup launch times, etc.

Control Application

Issues SnapVault commands to primary and secondary storage using NDMP v4 SnapVault Management extensions.

Differential Backups

During a differential backup, the backup application scans the file system or volume, determines all the changes that have occurred since the last full backup, and then writes those changes to the backup destination.

Full Backups

During a full backup, the entire file system or volume is copied to a backup destination. Some backup applications also refer to this as a “level 0” backup. Archival backup is another term that is used to mean the same thing.

Incremental Backups

During an incremental backup, the backup application scans the file system or volume, determines all the changes that have occurred since the last incremental backup, and then writes those changes to the backup destination.

Interleaving

This is the technique by which backup applications write multiple backup streams from various clients to a single backup volume or tape cartridge. Interleaving keeps the tape drive streaming and thereby speeds backup. Interleaved backup volumes are difficult to recover from, since the backup streams are scattered all over the backup volume. Also see multiplexing.

Media Server

An open systems server attached to a backup destination such as a tape library, tape drive, or a disk array. A backup server directs clients to send backup data to the media server. Also see storage node.

Multiplexing

Another term for interleaving.

NearStore

A product line of Network Appliance designed for the nearline storage market as an alternative collection point for backing up various NetApp systems as well as third-party systems. NearStore systems may be utilized for other purposes as well, for example one volume on a NearStore may be used for backup, while another might be used for archiving old software projects.

Open Systems SnapVault Agent

Software module installed on Open Systems which enables communication to a Network Appliance storage system to back up/restore over a TCP/IP network.

Primary Directory

A directory within a primary storage system that is configured for backup to a particular secondary volume. Primary directories on Network Appliance primary storage systems must be qtrees or the non-qtrees area of a volume. Primary directories on primary storage systems other than Network Appliance systems can be simply files or directories in a file system on an open systems computer.

Note: The non-qtrees area of a volume can only be restored to a qtrees within the primary storage system. It cannot be restored back to the non-qtrees area of the volume.

Primary System

A computer system to be protected by a secondary storage system using SnapVault through a configured SnapVault relationship. Each primary storage system has a list of secondary storage systems authorized to back up directories from the primary storage system. Primary systems are typically UNIX storage, Windows storage, or NetApp storage.

Qtree

A special subdirectory in a file system that acts as a virtual file system with special attributes, primarily quotas and permissions.

Recovery Point Objective (RPO)

Recovery point objective describes the age of the data you want the ability to restore in the event of a data loss.

Recovery Time Objective (RTO)

The recovery time objective is the time needed to restore data or recover from a disaster or, saying it another way, how long you can afford to be without your applications.

Secondary System

A system providing data protection for one or more primary storage systems through a configured SnapVault relationship. In the usual configuration, a secondary storage system is a NearStore system, but it may also be a Network Appliance filer. Each secondary storage system has a list of authorized primary storage systems that are allowed to restore data from the secondary storage system. These systems are typically NearStore products.

Secondary Volume

A volume within a secondary storage system providing data protection for one or more primary directories. Each secondary volume contains various qtrees obtained from primary qtrees or primary directories. Each secondary volume has a backup schedule based on backup interval and retention period.

SnapMirror

A product that performs automated file system replication of a data set onto the same or a separate disk or NetApp storage.

Snapshot

A Data ONTAP feature that creates an online, read-only copy of the entire file system, called a Snapshot copy that protects against accidental deletions or modifications of files without duplicating file contents.

SnapVault

Data ONTAP software option that backs up selected NetApp storage and open systems data sets over IP to a central online repository on the network.

SnapVault Incremental Transfer

A follow-up backup to the secondary storage system that contains only the changes to the primary storage data between the current and last transfer actions.

SnapVault Relationships

A persistent configuration that defines the primary directory and secondary volume for a SnapMirror software based data transfer.

SnapVault Snapshot Copies

The backup images that SnapVault creates at intervals on its secondary storage systems. SnapVault Snapshot copies capture the state of primary qtree and directory data on each primary system. This data is transferred to secondary qtrees on the SnapVault secondary system, which creates and maintains versions of the combined data for long-term storage and possible restore operations.

Storage Node

Some backup applications use this term for a media server.

Tape Library

A device that contains one or more tape drives, storage for tape cartridges, and a mechanism such as a robotic arm that moves the tape cartridges in and out of the drive(s).

Update

An incremental backup of a SnapVault qtree or OSSV file system

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`.

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:

```
-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`.

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: OSSV 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/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

Minimum version of Data ONTAP is 6.4 or 6.5.

Licensing

- a) Primary systems: SnapVault primary license
- b) Secondary systems: SnapVault/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 OSSV 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 OSSV doesn't consist of a primary running Data ONTAP, these steps are not used in an OSSV 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 may 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 may 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.

Revision History

Date	Name	Description
12/2006	Jeremy Merrill	Updated for 7.2.1
07/2006	Jeremy Merrill	Creation



© 2006 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, DataFabric, Data ONTAP, NearStore, SnapMirror, SnapVault, SyncMirror, and WAFL are registered trademarks and Network Appliance, FlexVol, LockVault, NOW, RAID-DP, and Snapshot are trademarks of Network Appliance, Inc. in the U.S. and other countries. Linux is a registered trademark of Linus Torvalds. Microsoft and Windows are registered trademarks of Microsoft Corporation. Oracle is a registered trademark of Oracle Corporation. Solaris and Sun are trademarks of Sun Microsystems, Inc. Symantec is a registered trademark and NetBackup is a trademark of Symantec Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.