



**Alex Osuna
Jose Hanchi
Darrin Chapman
Amol Chitre
Jeremy Merrill
Remington Svarcas**

IBM System Storage N series SnapVault Best Practices Guide

Introduction

This IBM® Redpaper guides you in implementing SnapVault® technology. It provides step-by-step configuration examples, and gives recommendations to help you design an optimal SnapVault solution. This document is intended for field personnel who require assistance in deploying and architecting a SnapVault solution; it is useful for system administrators, backup administrators, and IT managers who want to benefit from all of the advantages of SnapVault, and also provide the highest level of protection for their data.

For further information about this topic, you can refer to the latest publications at the following site:

www.redbooks.ibm.com

For specific updates about processes, Data ONTAP® command syntax, and current requirements, issues, and limitations, you can refer to the following site:

<http://www.ibm.com/support>

Business applications of SnapVault

SnapVault software from IBM N series is a reliable and economical way to protect enterprise data, and it offers 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. SnapVault is so useful that it renders many common backup policies and schedules obsolete.

This paper first provides 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 (see Figure 1).

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.

Note: This publication does not replace the Data ONTAP document *Data Protection Guide*, GA32-0522-01. That guide provides important information that is not covered here, including detailed procedures for day-to-day operational tasks.

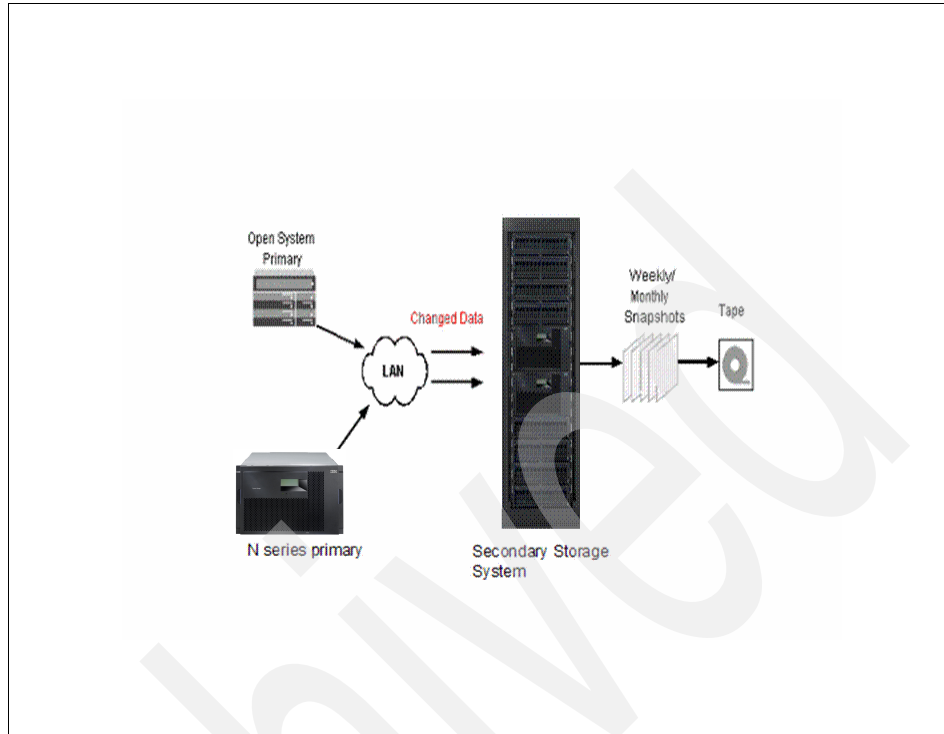


Figure 1 Basic SnapVault

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 IBM System Storage™ N series, or on a different vendor's storage?

The following sections help organize concepts and determine your requirements. However, knowledge of your data and users is the best guide to help in this process. If you feel you do not know enough about your data or users, consider interviewing a sample of users to learn more about their backup and recovery needs.

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, having local Snapshot™ copies on a storage system may protect against a user error that deletes a file or group of files, but it 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.

Figure 2 depicts how each Snapshot is represented by a list of pointers to the datablocks modified or added during the period the Snapshot is covering.

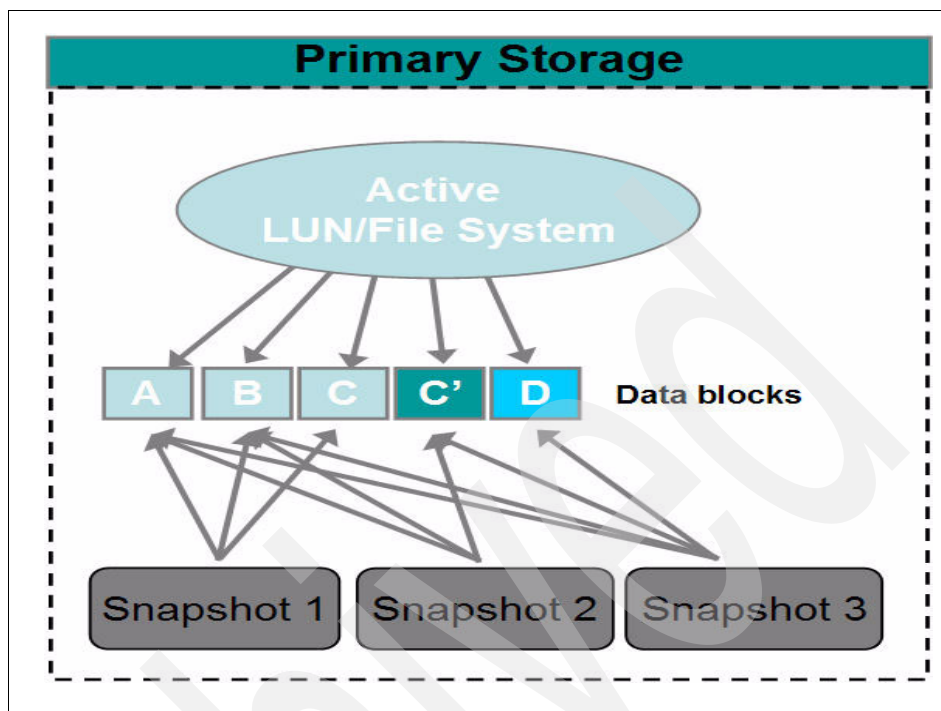


Figure 2 Snapshot detail

There are 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; a local Snapshot copy on the same storage system or a remote copy of the data on another system serve equally well. There are only two requirements:

- The backup must be made prior to the event that causes the data integrity problem.
- 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 may be caused by errors in the storage media, such as:

- Failure of a single disk
- Failure of multiple disks at once
- Corrupted 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™ (see Figure 3) 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.

Double Parity RAID: RAID-DP™

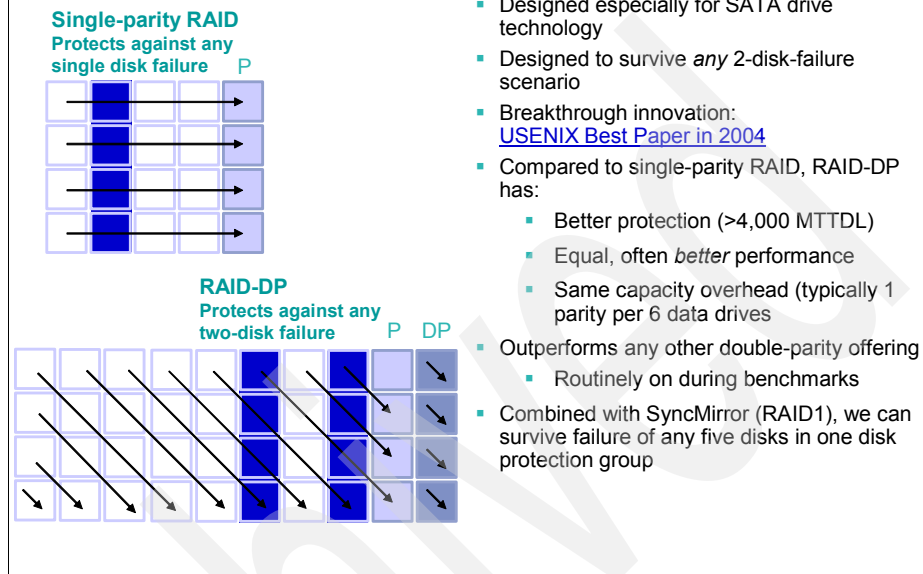


Figure 3 RAID-DP

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 offsite 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 offsite.

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 (see Figure 4). This configuration also protects you from a site-level disaster.

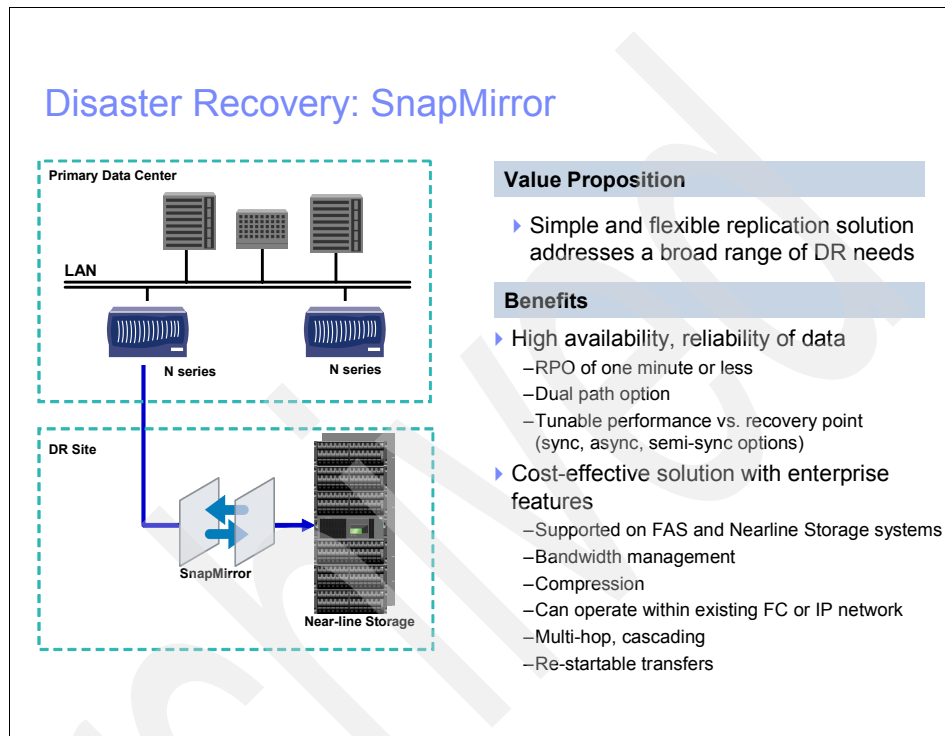


Figure 4 SnapMirror illustrated

Usage patterns

When planning a backup and recovery system, 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.

A key point to remember is this: 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.

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 every hour (see Figure 5).

The other consideration to keep in mind is how long to keep each 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 costly. In essence, it is important to determine what granularity is required to accommodate user needs.

Note: Restore granularity is the same concept as recovery point objective (RPO), but Snapshot technology extends this idea to provide the ability of restoring data from something other than the most recent backup.

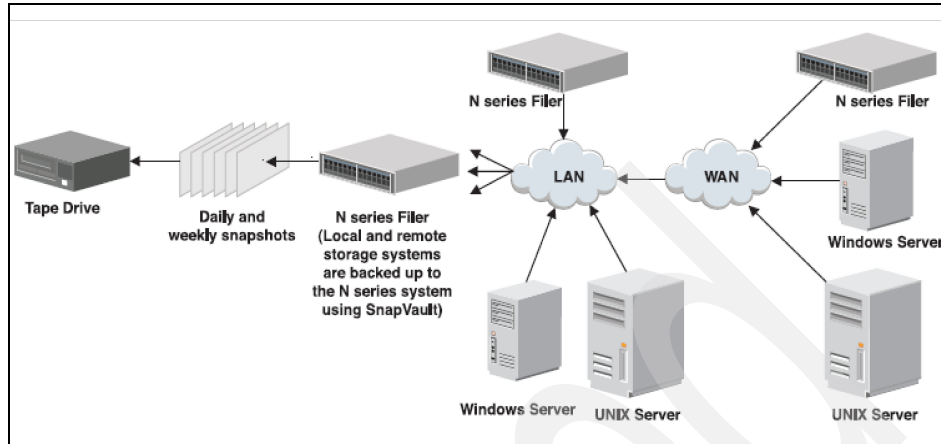


Figure 5 SnapVault scheduling

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” (that is, 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 corrupted 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 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.

Media costs

Keeping many backup copies of a data set for a long time can consume a significant amount 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 IBM Storage System N series WAFL® (Write Anywhere File Layout) file system (see Figure 6) 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.

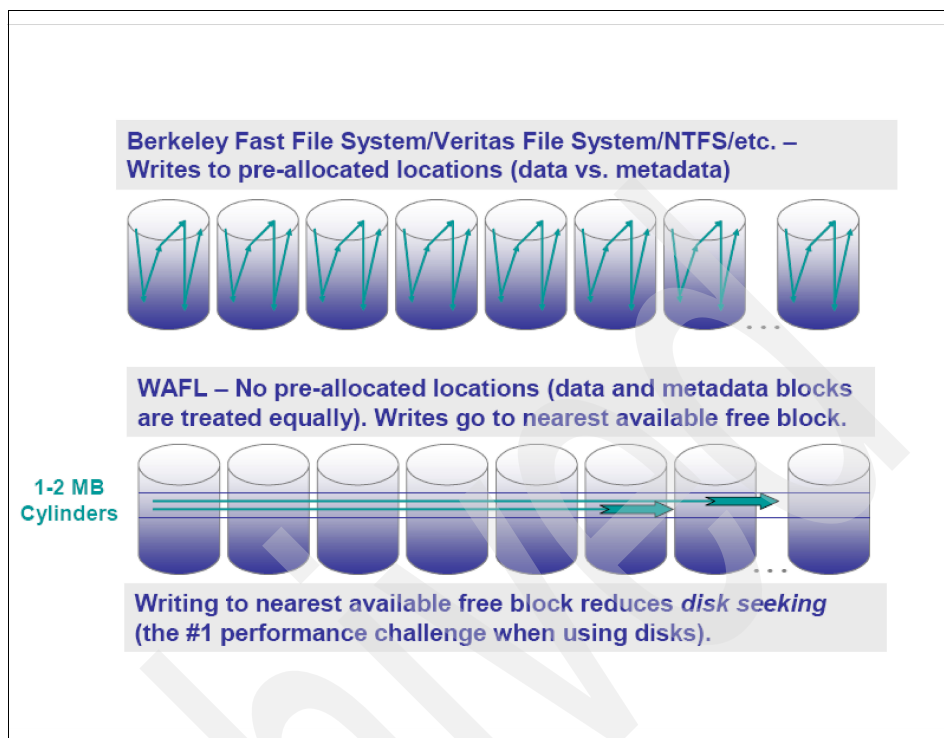


Figure 6 WAFL

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.

Legal requirements

In some industries, and with some data sets, there are legal requirements that specify how frequently backups must be performed, 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 is based on the SnapLock license.

Existing backup schedules and policies

It is useful to present backup schedules, granularity and retention in the form of a Data Restore Service Level Agreement table. For example, a common backup and recovery environment implemented using standard enterprise backup software along with a tape library might look like Table 1.

Table 1 Existing backup schedule

Dataset	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 onsite	Integrity threats, media errors
	Up to a month	1 week	2 hours	Weekly full tape backup, copies onsite and offsite	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2 days	Monthly full tape backup, stored offsite	Integrity threats, media errors, site disasters
Production database	13 days or fewer	1 day	12 hours	Nightly full tape backup stored onsite	Integrity threats, media errors
	Up to a month	1 week	12 hours	Weekly full tape backup, copies onsite and offsite	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2.5 days	Monthly full tape backup, stored offsite	Integrity threats, media errors, site disasters

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

Table 2 Adjusted SnapVault backup schedule

Dataset	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 + offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 7 days	1 day	5 minutes	Local Snapshot copies + offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 13 days	1 day	20 minutes	Offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 3 months	1 week	20 minutes	Offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2 days	Monthly full tape backup, stored offsite	Integrity threats, media errors, site disasters
Production database	1 day or less	1 hour	30 minutes	Hot backup to Snapshot and offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 7 days	1 day	30 minutes	Hot Backup to Snapshot and offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 13 days	1 day	2 hours	Offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 3 months	1 week	2 hours	Offsite SnapVault	Integrity threats, media errors, site disasters
	Up to 2 years	1 month	2.5 days	Monthly full tape backup, stored offsite	Integrity threats, media errors, site disasters

After the new backup schedules and retention policies have been determined, compare them with the existing schedule. You will clearly see that utilizing SnapVault is a huge improvement to the service level provided. In the existing 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.

By contrast, 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.

SnapVault overview

Figure 7 shows a simple SnapVault architecture with both the data center and the remote office. Both open systems (with heterogeneous storage) and IBM System Storage N series are backed up to a near-line system. After the data is on the near-line system, SnapMirror is used to mirror the data to a remote data center.

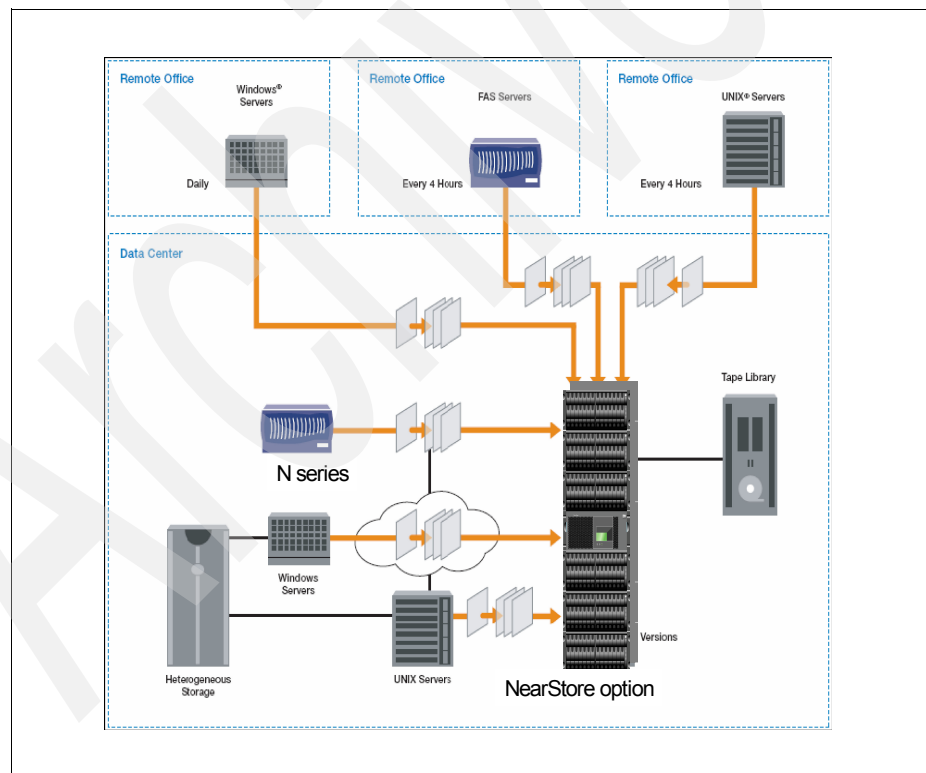


Figure 7 Simple SnapVault implementation

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.

First, a complete copy of the data set is sent across the network to the SnapVault secondary. This initial, or *baseline*, transfer may take a long 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.

For example, if SnapVault backed up a 100 GB data set for the first time, it would consume 100 GB of disk space on the SnapVault secondary. Over the course of several hours, users change 10 GB of data on the primary file system. When the next SnapVault backup occurs, SnapVault writes the 10 GB 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 110 GB of space on the SnapVault secondary (see Figure 8).

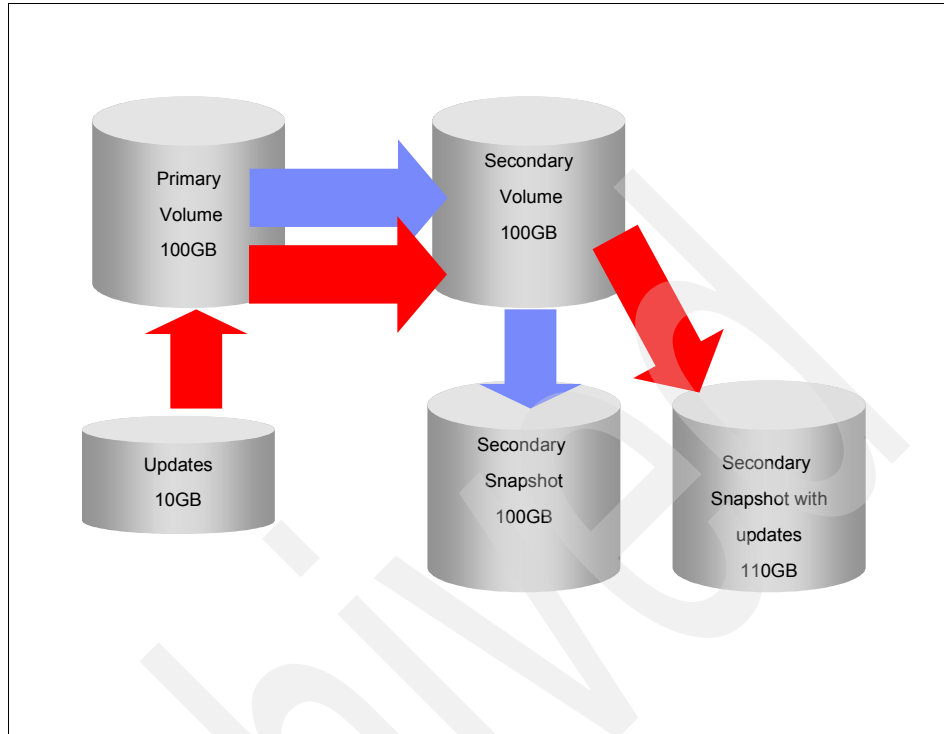


Figure 8 SnapVault to secondary

Note: If the baseline transfer is a large amount of data, you can use LREP as indicated in Chapter 6 of the IBM Redbooks® publication *IBM System Storage N series SnapMirror*, SG24-7260.

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 similar change rates and identical transfer schedules) into the same destination volume.

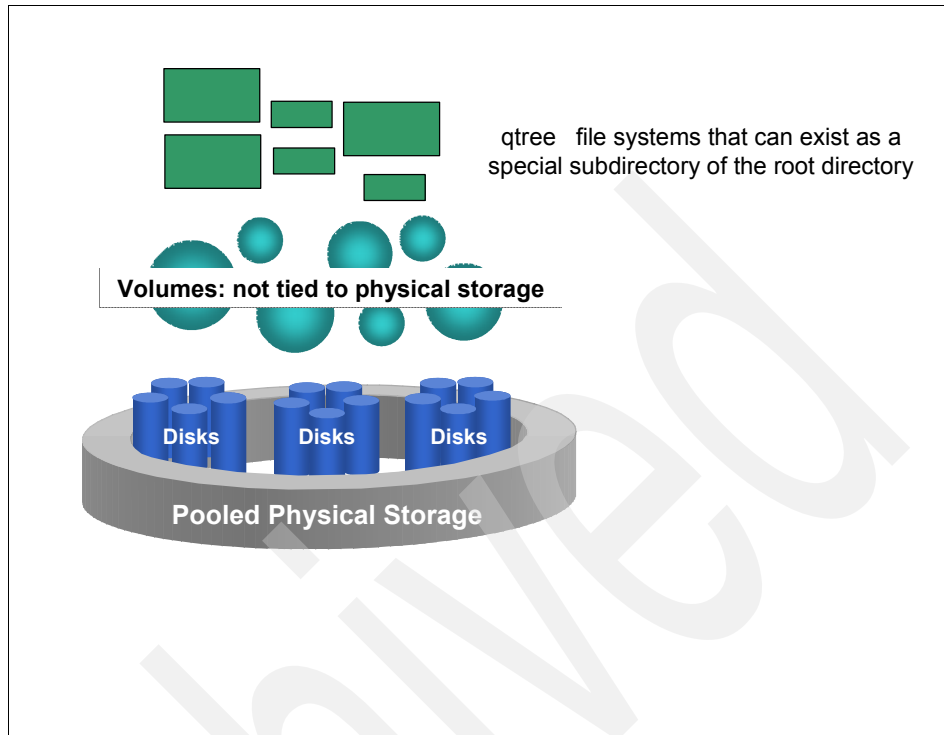


Figure 9 A qtrees

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 qtrees using the **snapvault start** command, a Snapshot copy is created on the volume that contains the qtrees you want to back up. The SnapVault primary reads the image of the qtrees 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 qtrees. 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. After 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.

Benefits of SnapVault

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

Incremental backups forever

A *full backup* copies the entire data set to a backup medium, which is tape in traditional backup applications, or an IBM System Storage N series near-line 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 offsite copy of the backup and providing a second copy in case one copy is corrupted. However, for certain types of problems it is possible that the corrupted 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. Backups are protected by RAID, block checksums and periodic disk scrubs, just like all other data on an IBM System Storage N series. Restores are simple because each incremental backup is represented by a Snapshot copy, which is a point-in-time copy of the entire data set, and is restored with a single operation.

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

Self-service restores

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

Restoring a file from a SnapVault backup is simple. Just as the original file was accessed 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.

Example 1 SnapVault restore syntax

```
itstocl> snapvault restore
usage:
snapvault restore [-f] [-s <snapname>] [-k <n>] -S
<secondary_filer>:<secondary_
path> [<primary_filer>:<primary_path>
```

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.

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.

Differences between SnapVault and SnapMirror

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 A: SnapVault/SnapMirror bundle” on page 47.

- ▶ 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 IBM System Storage N series running Data ONTAP. SnapVault can be used to back up both IBM N series and Open Systems primary storage, although the secondary storage system must be an IBM N series system or a near-line system.

SnapVault management options

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

Data ONTAP CLI

One method of managing SnapVault relationships and their transfer schedules 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.

Example 2 shows the commands that can go with the **snapvault** CLI command on the primary.

Example 2 SnapVault syntax

```
itsotuc3> snapvault
The following commands are available; for more information
type "snapvault help <command>"
abort          help          restore        status
destinations   release        snap
```

Example 3 shows the options that go with the **snapvault** command on the secondary.

Example 3 SnapVault secondary syntax

```
itsotuc4> snapvault
The following commands are available; for more information
type "snapvault help <command>"
abort          modify        start          stop
destinations   release        status         update
help           snap
```

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 **ndmcopy** command.

Operations Manager

Unlike homegrown scripts and competitor products, only Operations Manager takes full advantage of the Data ONTAP APIs and industry standards to deliver a full suite of storage management capabilities for enterprise storage and content delivery infrastructures.

To enable Operations Manager of SnapVault and Open Systems SnapVault, the business continuance option must be purchased and added to the Operations Manager installation. Operations Manager utilizes NDMP to access the primary and secondary systems. TCP port 10000 must be open if firewalls exist. If multiple interfaces exist on the Operations Manager server, NDMP-preferred interfaces can be utilized.

Operations Manager 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 *DataFabric Manager 3.4.1 Installation and Upgrade Guide*, GC26-7892.

When relationships already exist on the storage systems, Operations Manager allows the user to import these relationships for management, avoiding another costly full baseline transfer. The Operations Manager backup manager (see Figure 10) recognizes that the relationship exists when it is selected for backup.

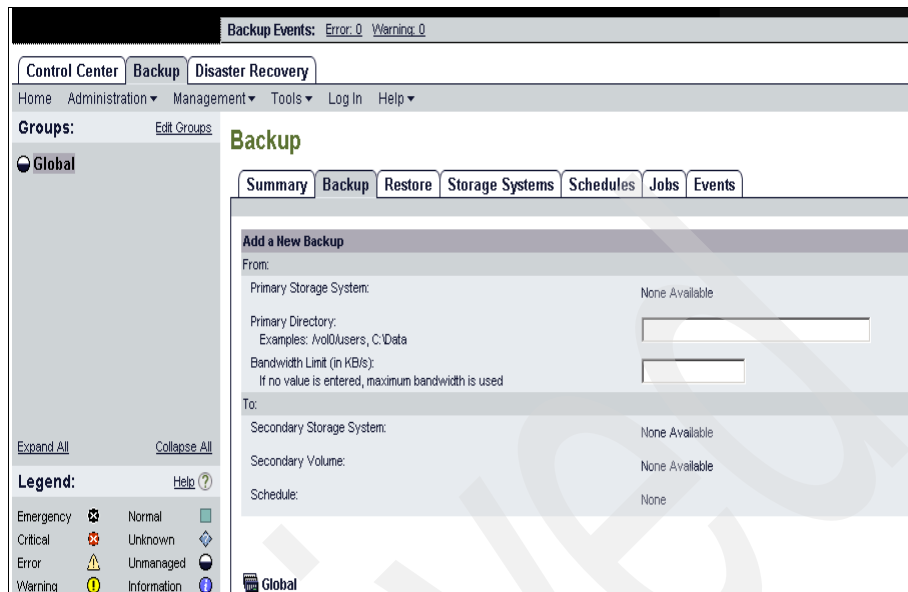


Figure 10 Operations Manager backup Administration

All scheduling can be performed in Operations Manager. Multiple schedules can exist, and retention policies can be placed on these schedules. If a schedule already exists from using the Data ONTAP command line and Operations Manager 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 Operations Manager, there is no need to create a qtree name; Operations Manager creates a unique qtree name. It is important to be aware of this when specific naming schemes are in place. Operations Manager is an excellent tool for data restoration. The ability to browse Operations Manager created backups on the secondary makes restoration much simpler.

Beginning with DFM 3.2 and continuing with Operations Manager, you have the functionality to create prescripts 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 Operations Manager 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, and so on)
- Privileges needed to run the script

By default, the scripts are installed in “script-plugins” in the root of the Operations Manager installation. The scripts can be run manually via Operations Manager or by a schedule (see Figure 11). These scripts can be useful to put a database into hot backup mode before an Open Systems SnapVault transfer, and then release it from hot backup mode upon completion.

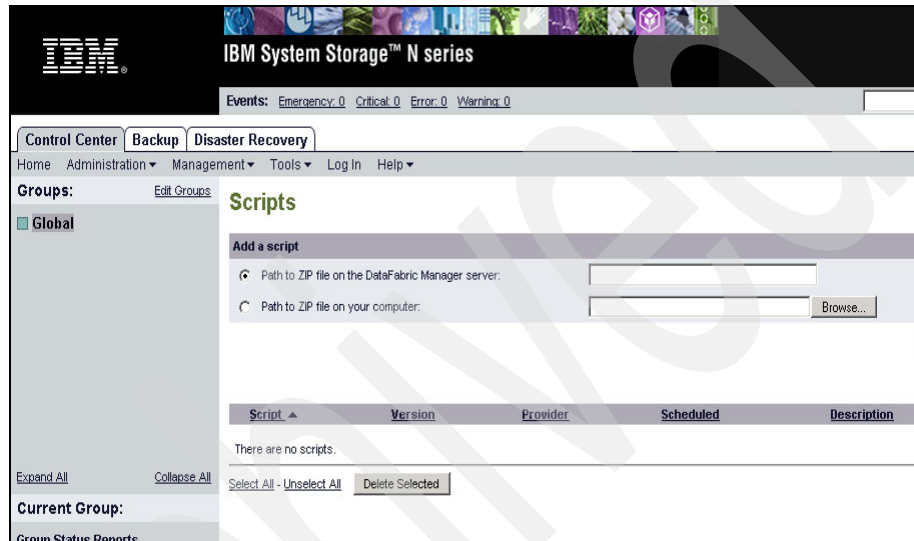


Figure 11 Script Management

For more details about Operations Manager and Operations Manager prescripts and postscripts, refer to the Operations Manager documentation on the IBM support Web site.

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.

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 (OSSV).

SyncSort

SyncSort Backup Express has been certified for IBM System Storage N series 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 IBM N series 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.

Configuring SnapVault

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

Step 1 - Determine the overall backup schedule

Determine the overall backup schedule you want to implement. This document uses the schedule shown in Table 2 on page 14.

The following examples assume that you are configuring backups for a single IBM N series system named itsotuc-pri, using a single near-line system named itsotuc-sec.

The home directories are in a qtree on itsotuc-pri called /vol/vol1/users; the database is on itsotuc-pri in the volume called /vol/oracle, and is not in a qtree.

Step 2 - Schedule Snapshot copies on the SnapVault primaries

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

1. License SnapVault and enable it.

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

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

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

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

```
itsotuc-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.

```
itsotuc-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.

Step 3 - Schedule Snapshot copies on the SnapVault secondary

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

1. License SnapVault and enable it.

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

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

```
itsotuc-sec> aggr create sv_flex 10  
itsotuc-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 N series storage 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.

```
itsotuc-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.

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

5. Set up schedules for the hourly backups.

```
itsotuc-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 N series storage 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 do not use the -x option, the secondary does not contact the primary to transfer the Snapshot copy. A Snapshot copy of the destination volume is merely created.

6. Set up schedules for the daily backups.

```
itsotuc-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 on page 14.

7. Set up schedules for the weekly backups.

```
itsotuc-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.

Step 4 - 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:

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

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

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

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

Special case - Database and Application Server backups

Simply scheduling a Snapshot copy on a database volume 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.

```
itsotuc-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.

```
itsotuc-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.

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

```
[ first commands to put the database into backup mode ]  
rsh itsotuc-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.

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 “Step 1 - Determine the overall backup schedule” to “Special case - Database and Application Server backups”.

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 case of an event where the amount of the data needed to be retained is greater than the available space in the LUN (see Example 4.) For example, if you have a 10 GB LUN on the primary N series storage system and rewrite all 10 GB, the next SnapVault transfer sends all 10 GB. The SnapVault transfer does not fail because it utilizes the 10 GB space reservation to complete those writes. SnapVault cannot delete the 10 GB that was overwritten because it is still required for the previous Snapshot copy.

Example 4 space reservation syntax

```
qtree level: qtree reservation qtree_path [enable|disable]
file level: file reservation file_name [enable|disable]
LUN level: lun set reservation lun_path [enable|disable]
```

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 offsite storage facility. This ensures that an offsite 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 offsite 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 offsite.

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 (see Figure 12). 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.

Another 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.

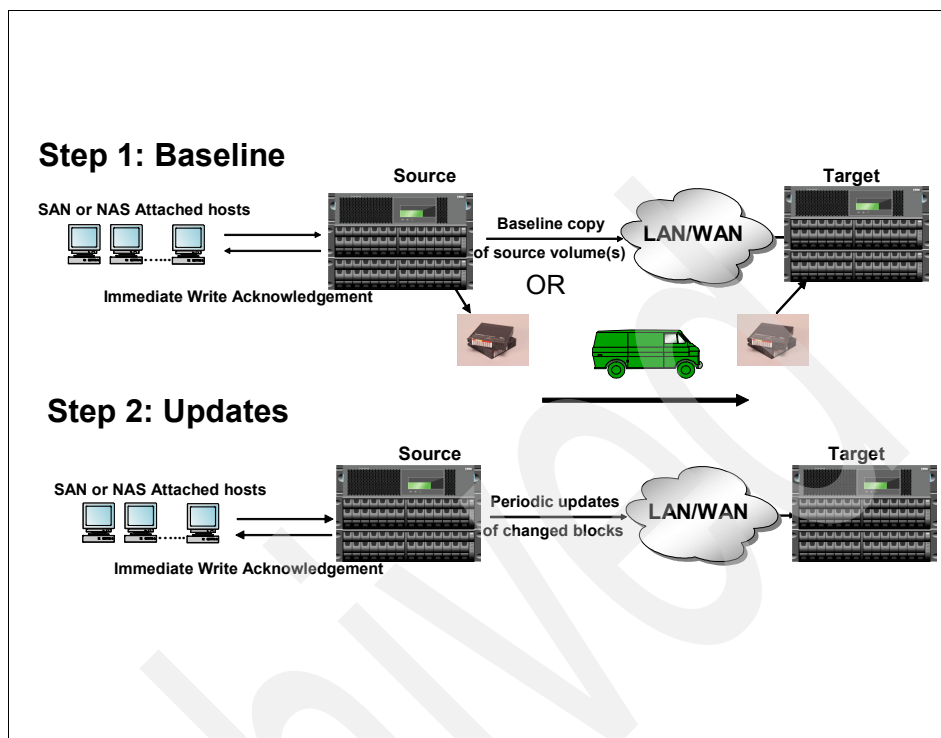


Figure 12 SnapMirror operations

Known SnapVault behaviors

The following section discusses SnapVault behaviors that users should be aware of prior to implementing SnapVault.

Transfer overhead

For every transferred inode, the SnapVault primary sends a 4 KB header. Also, all changed data is rounded up to 4 KB. 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 4 KB header transfer for that directory. If a file or directory is larger than 2 MB, an additional 4 KB header is transferred for every 2 MB.

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.

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 stopped. (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.

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 cannot be avoided.

What needs to be avoided is grouping in a volume qtrees that do not 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 1 GB worth of changed data, except for qtree4, which has 10 GB 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 will not 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 qtree that is still being transferred, qtree4, shows a status of Transferring and the amount of data already 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 10 GB of changed data to take a while to transfer. This would clearly be a flaw in the layout of the schedule and qtrees to the secondary volume. Figure 13 shows an example of a SnapVault transfer with qtrees in a Quiescing state.

r100-rtp01:/vol/sv_dest/mtree1_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/mtree2_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/mtree3_dest	Snapvaulted	01:05:16	Quiescing
r100-rtp01:/vol/sv_dest/mtree4_dest	Snapvaulted	01:05:16	Transferring (248 MB done)
r100-rtp01:/vol/sv_dest/mtree5_dest	Snapvaulted	01:05:16	Quiescing

Figure 13 Example of a transfer in a Quiescing state

Notice that in the example, mtree4 is still transferring while all other mtrees are in a Quiescing state. It would be a good idea to monitor mtree4 in this SnapVault transfer to see if it continues to cause the other mtrees to be in a Quiescing state. The change rate of mtree4 is not similar to the other mtrees in the destination volume, and it would make more sense to move mtree4 to another volume.

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 mtree contents back to the original primary mtree. After you have restored the entire contents of the mtree, you can choose either to resume the scheduled SnapVault backups (**snapvault start -r**) or to cancel the SnapVault relationship and the corresponding backups (**snapvault release**).

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

Traditional volumes versus flexible volumes

When you are setting up the secondary volumes, to maximize performance it is 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.

Sizing volumes on the secondary

The sizing of volumes on the secondary can vary based on the RTO, RPO and granularity required, in addition to the rate of change of the source volume and qtrees. Moreover, 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 in question so as to group similar qtrees into the same destination volume. Flexible volumes are the ideal type of volume for SnapVault destinations: if the rate of change, retention requirements, or size of the primary changes, the size of the destination volume can be adjusted.

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), then 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 workload also needs to be considered.

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, which is the maximum number allowed. Always try to spread the transfers out to avoid scheduling too many at any given time. For more information about scheduling transfers, see “Queuing your transfers” on page 41.

Table 3 Concurrent transfer limits

	Number of simultaneous transfers allowed FC drives	Number of simultaneous transfers allowed ATA drives	
IBM N series Storage System	Data ONTAP > 7.0.1	Data ONTAP > 7.0.1	Notes
N3700	8	4	
N5200	16	8	1
N5500	16	8	1
N7600	24	12	1, 2
N7800	32	16	1, 2

Table notes:

1. Storage systems that use ATA drives.

A storage system that has any ATA drives, other than near-line 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. The N7600 and N7800 require a minimum of Data ONTAP 7.2.

Near-line 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 near-line system is the QSM source, SnapVault source, Volume SnapMirror (VSM) source, or VSM destination count *twice* against the maximum number.

For example, suppose that you have a near-line 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 near-line system has 30 replications running concurrently:

$$20 \text{ transfers (SnapVault)} + 5 \times 2 \text{ transfers (VSM)}$$

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.

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

NearStore option

To enable customers to utilize the IBM Storage System N series as a secondary storage system, a new software license option, called the `nearstore_option`, has been introduced. This license option can be installed only on the IBM Storage Systems N5200/N5500. This option is supported on Data ONTAP 7.1 and later versions.

This license option provides increased concurrent streams when IBM Storage Systems N5200/N5500 are used as destinations for SnapMirror® or 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 N5x00 and N7x00 are described in Table 4.

Table 4 Concurrent streams per IBM N series storage system

IBM System Storage System	Number of concurrent streams
N5200	16
N5300	16

IBM System Storage System	Number of concurrent streams
N5500	16
N7600	24
N7800	32
N5200 with ATA drives	8
N5500 with ATA drives	8
N7600 with ATA drives	12
N7800 with ATA drives	16

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 N5x00 system with FC drives, if you have 2 QSM sources, 6 QSM destinations and 3 VSM sources concurrently running, you can start only up to 5 more SnapVault destinations. Note that a replication operation within the same storage system is considered 2 concurrent streams.

The maximum stream count shown is *per controller*. In a clustered N5200 configuration with FC drives and both controllers active, each controller has a maximum limit of 16 streams. After 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 stopped.

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

Table 5 describes the maximum streams for each replication operation.

Table 5 Near-line option concurrent transfers

Operation	N5200 maximum streams	N5300 maximum streams	N5500 maximum streams	N7600 maximum streams	N7800 maximum streams
SnapVault source	16	16	16	24	32
SnapVault destination	32	64	64	96	128

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

If you are using the N5500 with near-line personality as a qtree SnapMirror (QSM)/SnapVault destination alone, you can have up to 64 *concurrent* streams. Without the nearstore_option license on the same N5500, you are limited to 16 concurrent streams with FC drives and 8 with ATA drives. This is precisely a scenario where a N5500 with near-line personality would be beneficial to customers.

For a clustered configuration consider the maximum values of stream count shown in Table 5 are *per controller*. For example, in a clustered N5200 configuration where both controllers are active, each controller has the maximum limit of 32 streams for QSM destinations. After a controller is taken over by the other controller, a maximum of 32 streams is available to the entire system. During takeover and giveback, all transfers running on either the controller being taken over or given back are stopped.

Performance impact on primary during transfer

Given that a SnapVault transfer is a pull operation, resource utilization will occur 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 do not 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 10 GB in size. The first dataset, dataset1, has approximately a million small files, and the second dataset, dataset2, has five files, all 2 GB 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.

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 “Concurrent transfers” on page 36.

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 beyond 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.

SnapVault within a clustered system

SnapVault in Data ONTAP 7.2.1 and above 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 bidirectional backup between two different clustered systems. This feature will enable customers to SnapVault within a cluster from FC drives to SATA drives in the same system.

In the event that a cluster fails over, the SnapVault transfers will continue to run, but the maximum number of concurrent transfers is the same as a single head (see Figure 14).

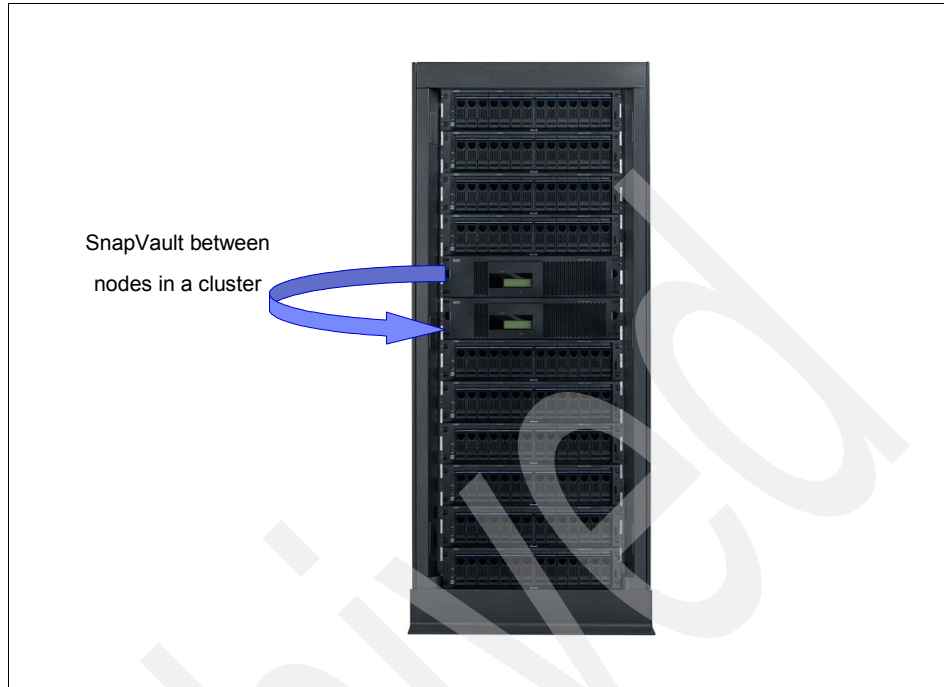


Figure 14 SnapVault between Nodes

Note: you may not install both a primary and a secondary license on the same controller or head. This means SnapVault may be done between cluster nodes.

Best practices and recommendations

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

General best practices

There are many best practices to be aware of in order to ensure a successful SnapVault deployment; some of the general best practices are explained here.

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. In it 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.

Scheduling guidelines

When setting up the SnapVault schedule, 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 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 shown in Table 6.

Table 6 Snapshot names and frequency

Snapshot name	Snapshot frequency
sv_hourly	Hourly
sv_daily	Daily
sv_weekly	Weekly

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

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

Primary Snapshot copy retention

When planning your SnapVault transfer schedule, keep in mind that you can also retain SnapVault Snapshot copies at the primary. It may not be a requirement to

keep hourly copies on the secondary, but this would be an ideal situation for primary copy retention.

In the schedule created in “Configuring SnapVault” on page 26, it was decided 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 reduces 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 perhaps transfer only four hourly copies (one every six hours).

Changing the “tries” count

The `-t` option (tries) of the **snapvault** command sets the number of times that updates for the qtree should be tried before giving up; the default is 2. When the secondary starts creating a Snapshot copy, it first updates the qtrees in the volume (provided that the `-x` option was set on the Snapshot schedule). If the update fails for some reason (such as a temporary network outage), then 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. When 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 once more 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 (Example 5). This is useful when there is a planned network outage.

Example 5 SnapVault modify syntax

```
itstotuc4*> snapvault modify
usage:
snapvault modify [-k <kbs>] [-t <n>] [-o <options>] [-S
<primary_system>:]<primary_path>] [<secondary_filer>:]<secondary_path>
      where <options> is
<opt_name>=<opt_value>[[,<opt_name>=<opt_value>]...]
```

Primary data layout

With FlexVol volumes there are 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 a dash (-) 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 is recommended to use one qtree inside the volume. Using this configuration allows restores to function like any other SnapVault restore.

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

Common misconfigurations

The following sections describe common misconfigurations that a user may encounter with SnapVault. These are issues you should consider prior to the deployment in order to achieve a successful SnapVault deployment.

Time zones, clocks, and lag time

A scheduling consideration 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. Figure 15 on page 46 shows 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 15 SnapVault status on primary

Figure 16 shows the output from the secondary.

/vol/sv_vol/qtrees1	r100-rtp01:/vol/sv_dest/qtrees1_dest	Snapvaulted	-00:03
/vol/sv_vol/qtrees2	r100-rtp01:/vol/sv_dest/qtrees2_dest	Snapvaulted	-00:03
/vol/sv_vol/qtrees3	r100-rtp01:/vol/sv_dest/qtrees3_dest	Snapvaulted	-00:03
/vol/sv_vol/qtrees4	r100-rtp01:/vol/sv_dest/qtrees4_dest	Snapvaulted	-00:04
/vol/sv_vol/qtrees5	r100-rtp01:/vol/sv_dest/qtrees5_dest	Snapvaulted	-00:04

Figure 16 SnapVault status on secondary

As you can see, the secondary storage system has a negative lag time because the clock on the primary storage system is ahead of the secondary. In this case, it is only a matter of a couple of minutes, but it could be worse. If the secondary is ahead of the primary, the issue could be even larger.

Suppose that the secondary is ahead of the primary by 15 minutes. At 11 p.m., the secondary is scheduled to get the data for the daily Snapshot copy. In this case, when it is 11 p.m. on the secondary, it is only 10:45 p.m. on the primary, so the primary has not yet created the sv_daily.0 Snapshot copy. This gives a lag time of 23:45 on the secondary, exposing you to potentially losing a day's worth of data. Make sure you verify that primary and secondary 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. In this case it is important to remember that schedules are based on the local clock. Given this scenario, assume that there are two storage systems, one on the U.S. East Coast and one on the U.S. West Coast. You must ensure that schedules account for the three-hour difference, otherwise you will either have negative lag times or lag times greater than what is expected based on the schedule.

Managing the number of Snapshot copies

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 is important to make sure that the qtrees within a SnapVault destination have the same characteristics in order to avoid reaching the 250 copy limit.

Volume to Qtree SnapVault

When issuing the **snapvault start** command, you are not required to specify a qtree name for the source; this should be avoided. This type of relationship increases the performance of the SnapVault transfer, and increases the amount of time it takes to perform a backup.

Because you must specify a qtree for the SnapVault destination, an entire volume now resides in a qtree on the destination. In the event of a restore via 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. Subsequently, you must copy manually the data back to the appropriate location.

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 with orders of magnitude beyond those available with traditional solutions.

Additional resources

- ▶ *IBM System Storage N series SnapMirror, SG24-7260*
- ▶ *IBM System Storage N series Data Protection Online Backup and Recovery Guide, GA32-0522*
- ▶ *DataFabric Manager 3.4.1 Installation and Upgrade Guide, GC26-7892*

Appendix A: 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.

Licensing

1. Primary systems: SnapVault primary license
2. 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 does not 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 B: 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 does not 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 host name 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 C: 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 or qtrees.

The backup schedule and the Snapshot schedule on the destination volume affect disk usage on the destination volume. Considering that the rate of change on the source volume is not likely to be constant, it is useful to provide a buffer of additional storage capacity above the requirement, 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.

To determine the rate of change for a volume, use the **snap delta** command to display the rate of change between Snapshot copies. For more information about **snap delta** and how to read the output, refer to the man page for the **snap** command.

The team that wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Tucson, Arizona.

Alex Osuna is a Project Leader at the International Technical Support Organization, Poughkeepsie Center. He has more than 28 years of experience in the IT industry, and has spent 19 years specializing in the storage area. Alex holds 10 certifications from IBM, Microsoft® and Red Hat. Before joining the ITSO, he worked as a System Engineer for Tivoli®.

Jose Hanchi is an IT Specialist with Techline providing pre-sales technical support to IBM sellers and Business Partners in Latin America on disk systems and SAN infrastructure. He has been with IBM since 2004 and previously worked with the ESS/DS8000/DS6000 field defect triage team primarily on open systems analysis.

Thanks to the following Network Appliance™ Corporation personnel for their contributions to this project:

Darrin Chapman

Amol Chitre

Jeremy Merrill

Remington Svarcas

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

© Copyright International Business Machines Corporation 2007. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.



Send us your comments in one of the following ways:


- ▶ Use the online **Contact us** review redbook form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbook@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099, 2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



Redpaper

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo) ®
DB2®

IBM®
Redbooks®

System Storage™
Tivoli®

The following terms are trademarks of other companies:

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Snapshot, RAID-DP, FlexVol, Network Appliance, WAFL, SyncMirror, SnapVault, SnapMirror, NearStore, DataFabric, Data ONTAP, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

Solaris, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.