Data Protection Strategies in IBM System Storage N Series

IBM data protection solutions in detail

Knowing your data

Business issues affecting data protection

Alex Osuna
Sakina Fakhruddin
Keith Knudson
Bobby Oommen

Redbooks

ibm.com/redbooks

International Technical Support Organization

Data Protection Strategies in IBM System Storage N Series

June 2008

SG24-7591-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (June 2008)

This edition applies to Data ONTAP Version 7.1 or later

Contents

Preface

Systems fail, users accidentally delete files, natural disasters occur, and mistakes happen. Businesses are losing critical data. One of the most important questions IT management must ask is, "What is my data recovery plan?" IBM® System Storage™ N Series provides a variety of choices for data protection and recovery. This IBM Redbook® publication addresses many available options, and recommends solutions for protecting data using the IBM System Storage N Series.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.

Alex Osuna is a project leader at the International Technical Support Organization, Tucson Center. He writes extensively and teaches IBM classes worldwide on all areas of storage. Before joining the ITSO 3 years ago, Alex was a Principal Systems Engineer with the Tivoli® Western Region. He has 30 years in the I/T industry focused mainly on storage. He olds certifications from IBM, Microsoft® and Redhat.

Sakina Fakhruddin is a independent I/T consultant she was a Software Engineer at IBM India Systems and Technology Lab, Pune. Sakina has been working with the IBM Storage management team on development of "Storage Configuration Manager" (SCM), mainly focussing on IBM BladeCenter® SAS Connectivity Module in BladeCenter-S. She holds a Bachelor of Engineering Degree in Computer Science from Cummins College of Engineering, Pune University.

Keith Knudson is a independent contractor in the I/T industry from the State of Texas in the USA. He holds several Operating system certifications and is well versed in storage including Network Attached Storage.

Bobby Oommen is an Alliance engineer for IBM DB2® at NetApp® Inc RTP, North Carolina. He holds a Bachelors degree in Mechanical Engineering from University of Dharwad, Karnataka(India). He has been in involved in IT industry over 13 years, in areas of database programming, Dpropr, Qrepl, modeling, designing, administration, performance tuning and storage integeration. He is an IBM DB2 Certified DBA.

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@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

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 is 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 illustrate 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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at:

http://www.ibm.com/legal/copytrade.shtml

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

Redbooks (logo) ®
BladeCenter®
DB2®
ESCON®
FICON®
FileNet®
IBM®
Lotus®
Redbooks®
System p™
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, LockVault, WAFL, SyncMirror, SnapVault, SnapRestore, SnapMirror, SnapLock, FlexVol, FlexClone, FilerView, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

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

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

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries 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.

1

Introduction to data protection in IBM System Storage N Series

Data protection strategies in the IBM System Storage N™ Series is discussed in this chapter. Data protection requirements vary with business needs. This chapter gives an overview of the need for data protection, and what IBM offers.

This chapter introduces the following:

- Data protection requirements for businesses
- IBM solutions

1.1 Introduction

Systems fail, users accidentally delete files (Figure 1-1 on page 2), natural disasters occur, and mistakes happen. Businesses lose critical data. One of the most important questions IT management must ask is, "What is my data recovery plan?" The N Series provides a variety of choices for data protection and recovery. This book addresses available options, and recommends solutions for protecting data on the N Series .

IBM System Storage N Series data protection strategies looks at data protection requirements, technology solutions, and performance issues to provide recommendations. The goal of this book is to help prepare for future challenges of protecting your system against system downtime. Costs of unplanned outages are increasing, along with requirements for continuous data access.

Disk storage capacity on servers are increasing at an alarming rate. Traditional data protection mechanisms are stretched to the limit. The lack of a backup window in many enterprises further escalates the problem. Storage capacity is scaling faster than tape cartridge capacity and tape bandwidth. Protecting multiterabyte storage systems with tape media can far exceed reasonable windows for both backup, and restore. These trends suggest that you must consider additional methods complementary to tape backup.

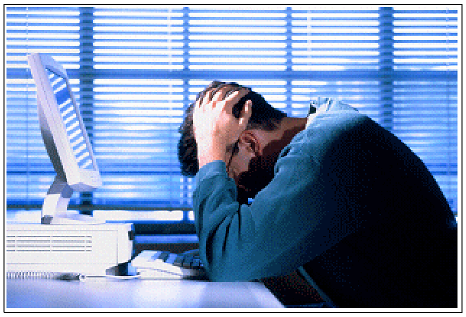

Figure 1-1 Disaster

1.2 Data protection requirements

Data is everywhere, from your mobile computer, to network servers, emails, and USB devices.Individuals must protect their data. In businesses, the need expands. Business success depends on the ability to maintain a resilient infrastructure. There are many factors to consider as you define data protection requirements. Typically businesses have five data protection requirements:

1. Fast, user-initiated recovery of accidentally deleted files

- We all experience situations where we accidentally delete important files. It is critical to have a mechanism that ensures retrieval of accidentally deleted files quickly, and with little effort. Files must be backed up frequently, making the recovered files as close as possible to the versions lost. Ideally, users perform online recovery operations by copying an earlier version of an accidentally deleted file from disk to their home directory. This capability frees the system administrator from having to restore a file from tape.

2. Tape archival of file systems, or current unreferenced data for future use

- Archived data provides a complete, self-consistent replica of a collection of data. Many businesses require that archived tapes are stored at an off-site location, for recovering a file , or file system many years later. For example, a software development firm must reconstruct an early version of their software product to fix an error for a customer. This requires an archive-to-tape scheme from earlier systems.

3. Minimized backup and recovery windows

- A *window* is the maximum amount of interruption time available for an application during a backup, or recovery operation. It is important to have a minimized backup, and recovery window that ensures minimal impact on users. While restore is infrequent, backups are constant.

4. Fast recovery from natural, or human-caused disasters

- What are the costs per minute, hour, or day of a business shutdown resulting from a disaster? How long can the company afford to not actively do business? Realistically calculate the costs of lost user productivity, missed business opportunity, and system administrator data management time. Typically, the recovery, or restore *window* is short. A data protection mechanism is required to meet these short restore *windows*. Some businesses require a server-mirroring solution, while others require a geographically separate site that is available to go online within minutes of disaster.

5. Data protection for compliance regulation purposes, and referencing

Figure 1-2 on page 3 illustrates the use of SnapShot for daily online backups, and SnapVault® technology, for backing up remote systems to a nearline storage subsystem.

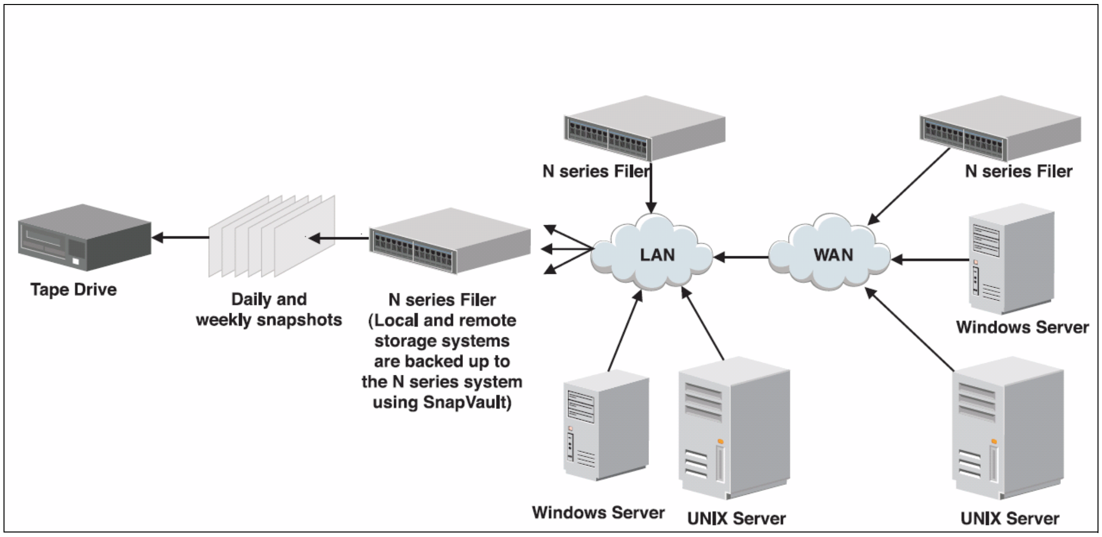
N series Filer
N series Filer
Tape Drive
Daily and weekly snapshots
N series Filer (Local and remote storage systems are backed up to the N series system using SnapVault)
LAN
WAN
Windows Server
Windows Server
UNIX Server
UNIX Server

Figure 1-2 IBM N Series with SnapVault and SnapShot configuration

1.3 IBM N Series solutions

The IBM N Series provides a unique set of solutions to address:

- SnapShot technology:
  - Daily online backups
  - Ensures short backup windows
  - Online backup maintenance for instantaneous access to previous versions of data
- SnapRestore® software:
  - Near-instantaneous recovery of files, or entire file systems to an earlier state
  - Recovery of an individual file to an entire file system
- SnapMirror® software:
  - Automated data replication
  - Fast and flexible enterprise solution for, mirroring, or replicating data over LAN, WAN, or Fibre Channel networks
- SnapVault software:
  - Disk-based backup and recovery of the IBM N Series, as well as open systems.
  - Improved and efficient data backup and recovery.
  - Provides a centralized backup system, reducing backup cost
- MetroCluster
  - Disaster recovery over WAN
  - Integrated high-availability, and busine ss continuance solution that leverages industry-proven technologies
  - Expands the comprehensive N Series portfolio capabilities of high-availability, and disaster recovery solutions that offer failover, data replication, and backup
- Third-party data protection products with NDMP support for archiving data to tape:
  - NDMP is an open protocol used to control data backup, and recovery communications between primary and secondary storage devices, such as storage systems, and tape devices in a heterogeneous network environment.
- SnapLock® and LockVault™:
  - Provides compliance regulations
  - Organizations face highly strict regulations for records retention that requires rigorous best practices to ensure accurate data retrieval at any time. The SnapLock function helps manage the permanence, accuracy, integrity, and security of data by storing business records in an inalterable form, and allowing for their rapid online accessibility for long periods of time.
- Native dump and restore:
  - Backup and restore from tape
  - Native dump and restore commands bundled in the Data ONTAP® software. These form the foundation of a scalable tape backup strategy.
- VTL Solution
  - Faster backups, by emulating tape devices on disk storage

2

Knowing your data

Knowing the datas role when choosing your data protection strategy is discussed in this chapter. The first step in choosing a data protection plan is to identify the critical data in the disaster recovery scenario. Major factors governing data protection strategy are discussed in this chapter.

2.1 Data classification

Several key factors drive data protection strategies:

- Acceptable recovery windows for business-critical data
- How up to date is the restored dynamic data
- Total data set size, volume size, and quota tree(qtree)
- Number and file size
- Directory structure
- Data types and compression

2.1.1 Business critical data

An essential step in determining a companys data protection strategy is to identify the priority of the data in the enterprise, and how quickly the data must be recovered in case of a disaster. You must also decide how up to date your recovered files must be in the event of a recovery operation. Categorizing data by how critical it is to the business allows system administrators to design flexible data protection strategies around restoration requirements. Figure 2-1 on page 6 shows the data classification model[1] hierarchy of critical data. The value of the information dictates the technology you use to protect it.

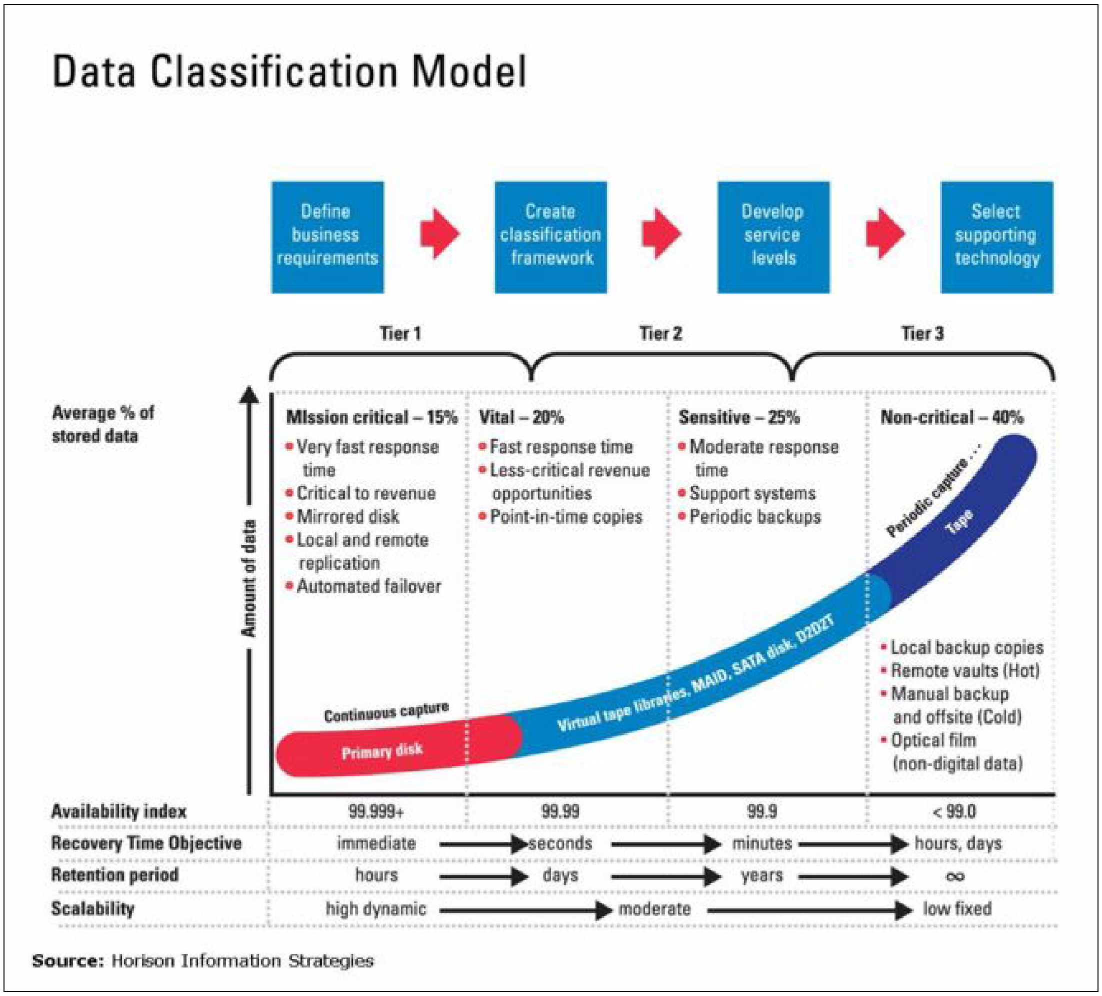

Availability index	99.999+	99.99	99.9	< 99.0
Recovery Time Objective	immediate →	seconds →	minutes →	hours, days
Retention period	hours →	days →	years →	∞
Scalability	high dynamic →	moderate →		low fixed

Source: Horison Information Strategies

Figure 2-1 Data classification model

[1] Source: Horison Information Strategies

Mission-critical data Used in key business processes, or customer facing applications. Critical data accounts for as much as 15 percent of all data stored online, and typically has very fast response time requirements. Snapshot™ copies are maintained for quick recovery.

Vital data Virtual data is important, but does not require instant recovery for business processes to continue operating.

Sensitive data Recovery takes several minutes to several hours without causing an impact to major operations, or business. With sensitive data, alternative sources exist for accessing, or reconstructing the data in case of lost data.

Non-critical data Represents approximately 40 percent of all data stored online making it the largest classification category. Lost, corrupted, or damaged non-critical data is reconstructed with minimal effort, and acceptable recovery times range from hours, to several days because this data is not essential for business survival. Most non-critical data is backed up to lower-cost storage solutions. Tape backups are the most popular choice.

We recommend that our customers optimize data protection by organizing their file systems into multi-volumes, and qtrees. For example, by isolating critical data into its own volume, or qtree you are able to:

- Mirror business-critical data, so it is available for immediate disaster recovery
- Create frequent Snapshot copies so that recent versions of files are available online
- Limit the size of volumes, or qtrees so tape backups, and restores take less time

Separating critical data for backup and restore purposes increases overall backup and restore rates for all data. Backing up critical data separately reduces the backup frequency for data that is less critical.

2.1.2 Dynamic data

According to the computer science definition, dynamic data is information that is changed asynchronously, as further updates to the information become available. Though it is similar to business-critical data, rapidly changing data warrants a different protection strategy. For example, dynamic data sets might require more frequent backups to capture frequent changes. Isolating dynamic data in its own volume, or qtree allows a business to design a protection strategy specifically around it.

For example, incremental backups of rapidly changing data takes longer than backups of more static data because there is more changed data to back up. If data is in its own volume, run frequent incremental backups. Sometimes, changing archival data is placed in a separate volume that gets occasional incremental backups, and full backups only at widely spaced intervals.

The amount of change occurring with a data set, also affects the use of SnapMirror, and SnapVault technology. IT managers can schedule incremental transfers more frequently to ensure mirrored data, and disk-based backups are up-to-date.

Total data set size

Knowing the total data set size and the size of each volume and qtree lets system administrators estimate tape backup and restore windows and SnapMirror or SnapVault data transfer times and then decide whether the estimated windows are adequate for the business.

Table 1 on page 8 calculates tape backup and restore windows using average backup and restore rates for a single LTO tape drive.

Table 1 Sample backup and restore windows

Backup time	500 GB volume	1.4 TB volume
Backup time @ 60 GB/hr	8.3 hours	23.3 hours
Restore time @ 40 GB/hr	12.5 hours	35 hours

If tape backup or restore times are unacceptable, use faster tape drives if possible. Or administrators can consider dividing large volumes into smaller volumes, or qtrees. For example, if a 1.4 TB volume is divided into four qtrees, you can backup each qtree to a separate tape drive, or you can perform separate full backups on four different nights.

For large volumes with long restore windows, consider using SnapMirror, or SnapVault software for fast disaster recovery. Volume sizes over 1.4 TB, or total data set sizes greater than 4 TB might exceed the natural performance limitations of SCSI, or Fibre Channel and tape, and might therefore be good candidates for a SnapMirror, or SnapVault solution.

Number and size of files

For a given volume size, large numbers of small files take longer to process than small numbers of large files. More files mean a bigger directory structure to process. In addition, there is a fixed amount of overhead associated with each file being backed up, regardless of the size of the file. Knowing the number and size of files gives administrators insight into the tape backup, or restore performance they can expect.

Directory structure

Directory structure also affects the performance of a backup and recovery solution. Both large directories and deep directories decrease performance. Large directories increase the complexity of memory management when each directory is processed. The entire contents of a small directory can be loaded into memory at once, whereas only parts of a large directory can be in memory at one time due to memory size restrictions. Memory management and frequent reading of the directory from disk introduce overhead.

Deep directory structure generally means higher file-to-directory ratio. To contain a certain number of files, more directories are needed in a deep directory structure than in a flat one. More directories translate to more work when backing up the files.At restore time, a child file, or directory is not created unless its parent directory is already created. This child-parent dependency prevents file, and directory creation from being parallel.

3

SnapShots

SnapShot technology is a data protection mechanism, offered by the IBM System Storage N Series, is discussed in this chapter. SnapShot is a unique feature that allows administrators to maintain read-only versions of online file systems.

3.1 SnapShots for data protection

Various data vendors provide SnapShot technology at different levels of the storage stack. SnapShot solutions are provided both at hardware, and software layers. These solutions are now becoming prevalent to perform data protection. IBM System Storage N Series provides SnapShots at file system level. File system based SnapShots provide a better level of recovery granularity than hardware SnapShots.

The WAFL® file system supports SnapShot. SnapShot copies, a bundled component of Data ONTAP software, allow users to recover accidentally damaged or deleted data by copying a desired file from a SnapShot directory. Versions of Data ONTAP currently supports 255 SnapShot copies per volume. Figure 3-1 illustrates the features of SnapShots.

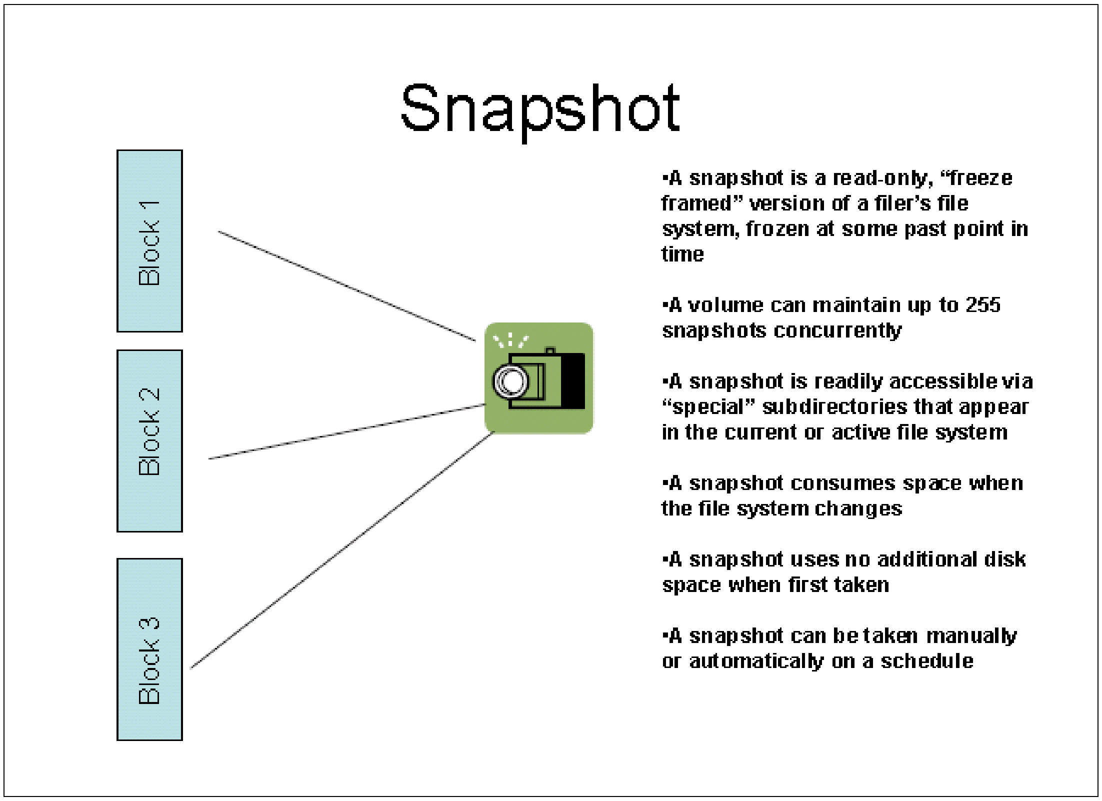

Snapshot

- A snapshot is a read-only, “freeze framed” version of a filer’s file system, frozen at some past point in time
- A volume can maintain up to 255 snapshots concurrently
- A snapshot is readily accessible via “special” subdirectories that appear in the current or active file system
- A snapshot consumes space when the file system changes
- A snapshot uses no additional disk space when first taken
- A snapshot can be taken manually or automatically on a schedule

Figure 3-1 SnapShot features

SnapShot copies augment and simplify an overall enterprise data protection strategy. SnapShot copies can serve as daily backups for users to recover their own files. The system administrator defines a schedule (every minute, hourly, nightly, or weekly SnapShot copies), and defines how long to retain each SnapShot copy. By creating SnapShot copies throughout the day (every three hours from 8 a.m. to 8 p.m., for example), a system administrator can guarantee that recent file versions are available for recovery. A business can choose to create nightly SnapShot copies, and retain them for one week, instead of nightly incremental backups to tape.

SnapShot technology form the core of the IBM System Storage N Series data protection solutions. Tape backup, SnapRestore, SnapMirror, and SnapVault, all described later in this book, also use the SnapShot feature. The dump command, and NDMP-compliant products read backup data directly from a SnapShot copy, this eliminates taking the file system offline, or dealing with open file conflicts. SnapRestore software allows the administrator to almost instantaneously revert a file, or an entire file system. The SnapMirror automated replication software uses SnapShot copies to provide asynchronous mirroring. SnapVault software uses SnapShot copies to provide disk-to-disk block-level incremental backup and archive capabilities to IBM storage systems.

Certain businesses choose to perform daily tape, or SnapVault backups, as well as create daily SnapShot copies. For example, in the unlikely occurrence of a concurrent failure of three disks in a single RAID-DP™ RAID group, data might not be recoverable from online SnapShot copies. For very valuable data, or longer-term storage, a business might decide to perform nightly tape backups, or SnapVault updates.

3.2 How SnapShot works

Some users assume that SnapShot copies incur significant disk space penalties, because each SnapShot copy appears as though it is a read-only copy of the file system. However, in reality, SnapShot copies usually only require a small disk space premium. SnapShot copies are maintained as pointers to disk blocks containing data. When the WAFL file system creates a SnapShot copy, it makes a copy of the set of pointers from the active file system, but does not copy data blocks. These pointers are consistently-sized segments within the file system referencing the data. In UNIX® terms, these pointers are called *inodes.* As the active file system changes, SnapShot copies continue to point to deleted or changed disk blocks, holding these blocks from the file system's free space, thereby using disk space, which can considered an overhead.

Understanding that the WAFL file system is a tree of blocks rooted by the root inode is the key to understanding SnapShots. To create a virtual copy of this tree of blocks, WAFL duplicates the root inode (Figure 3-2 on page 12).

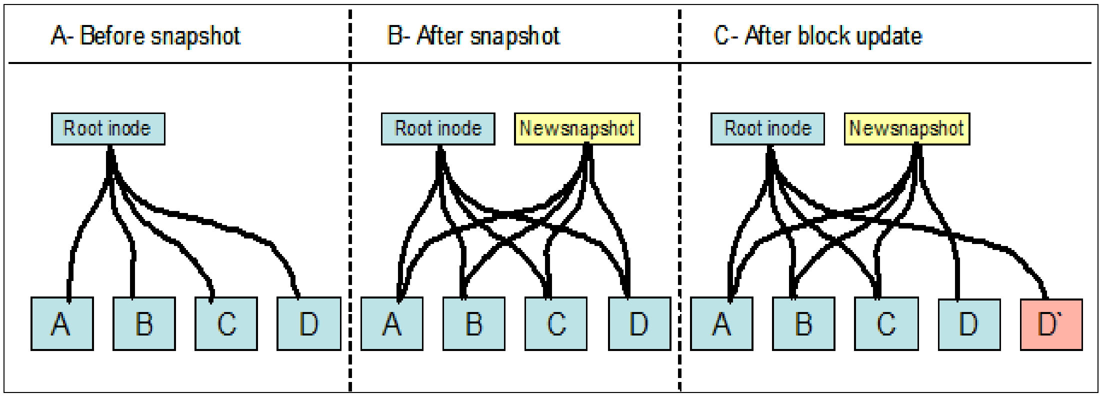

Figure 3-2 WAFL creates SnapShots by duplicating root inode

WAFL creates SnapShots by duplicating root inode. The new inode becomes the root inode of the SnapShot, which references the same data in the active file system. Therefore, a brand new SnapShot consumes no more space than the space required for an inode.

Column C in Figure 3-2 shows what happens when a user modifies data block D. WAFL writes the new data to block D on disk, and changes the active file system to point to the new block.The SnapShot still references the original block D, which is unmodified on disk.

By duplicating only the root inode WAFL ensures that disk I/O is reduced and SnapShots use minimal disk space. Also this creates SnapShots very quickly.

3.3 SnapShot example

SnapShots are created using a command line, or filer view. In this section, we illustrate the creation of SnapShots using a command line.

1. Create the SnapShot using the **`snap create`** command (see Example 3-1 on page 13)

Example 3-1 SnapShot creation

```
snap create vol1 vol1_snapshot
```

2. Use the **`snap list`** command to verify your SnapShot (see Example 3-2 on page 13)

Example 3-2 snap verification

```
Nseries> snap list
Volume vol_0
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  1% ( 1%)    0% ( 0%)  Nov 08 12:00  hourly.0

Volume vol1
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  Nov 08 12:36 vol1_snapshot
```

3.4 Summary

IBM System Storage N Series with SnapShot technology offers unique benefits. Point-in-time copy technology is available from a variety of data storage vendors, but not all technologies are created equal. The IBM System Storage N Series, with point-in-time SnapShot capability, offers important advantages. SnapShot technology helps deliver data stability, scalability, recoverability and performance capabilities.The IBM System Storage N Series leverages SnapShot technology as the foundation for developing a range of data protection solutions. These solutions incorporate, and extend the advantages of SnapShot technology to support advanced enterprise data protection.

In summary, SnapShot copies provide:

- User-initiated recovery of accidentally deleted files
- Replacement for nightly incremental backups to tape
- Ability to save data more frequently than incremental backups to tape
- A consistent copy of the file system for dump, and NDMP to use when creating tape backups
- Underlying technology for SnapRestore, SnapMirror, and SnapVault software, as described in the following sections

4

FlexVol and FlexClone

IBM System Storage N Series with *FlexClone®* and *FlexVol®* technologies provide new opportunities for organizations to work with the challenges of increased overhead, management costs, and data protection. Using *FlexVol* and *FlexClone* technology as data protection techniques are discussed in this chapter.

4.1 Overview of FlexVol and FlexClone

IBM System Storage N Series with *FlexClone* and *FlexVol* technologies help deliver storage virtualization solutions that support sharing and pooling enterprise data. Using *FlexVol* technology, system administrators can dynamically assign storage space to a user from the available pool of storage resources, based on the users space requirements. This flexibility helps your organization:

- Simplify operations
- Improve utilization and efficiency
- Make changes quickly and seamlessly

FlexClone technology generates nearly instantaneous replicas of data sets and storage volumes that require no additional storage space. Each cloned volume is a transparent virtual copy is used for enterprise operations.

Data protection has gained significant advantage with the introduction of *FlexVol* technology. Data ia placed on *FlexVol* volumes for performance, effective usage of disks, and data protection.

FlexVol volumes are created on top of a large pool of disks called an *aggregate*. There can be more than one *aggregate,* if required. *FlexVol* volumes are striped across every disk in the *aggregate*, and have their own attributes, and are independent of each other. For example, they can have their own SnapShot schedule, or their own replication schedule. *FlexClones* are created for a *FlexVol,* or *FlexClone*. See Figure 4-1 on page 17.

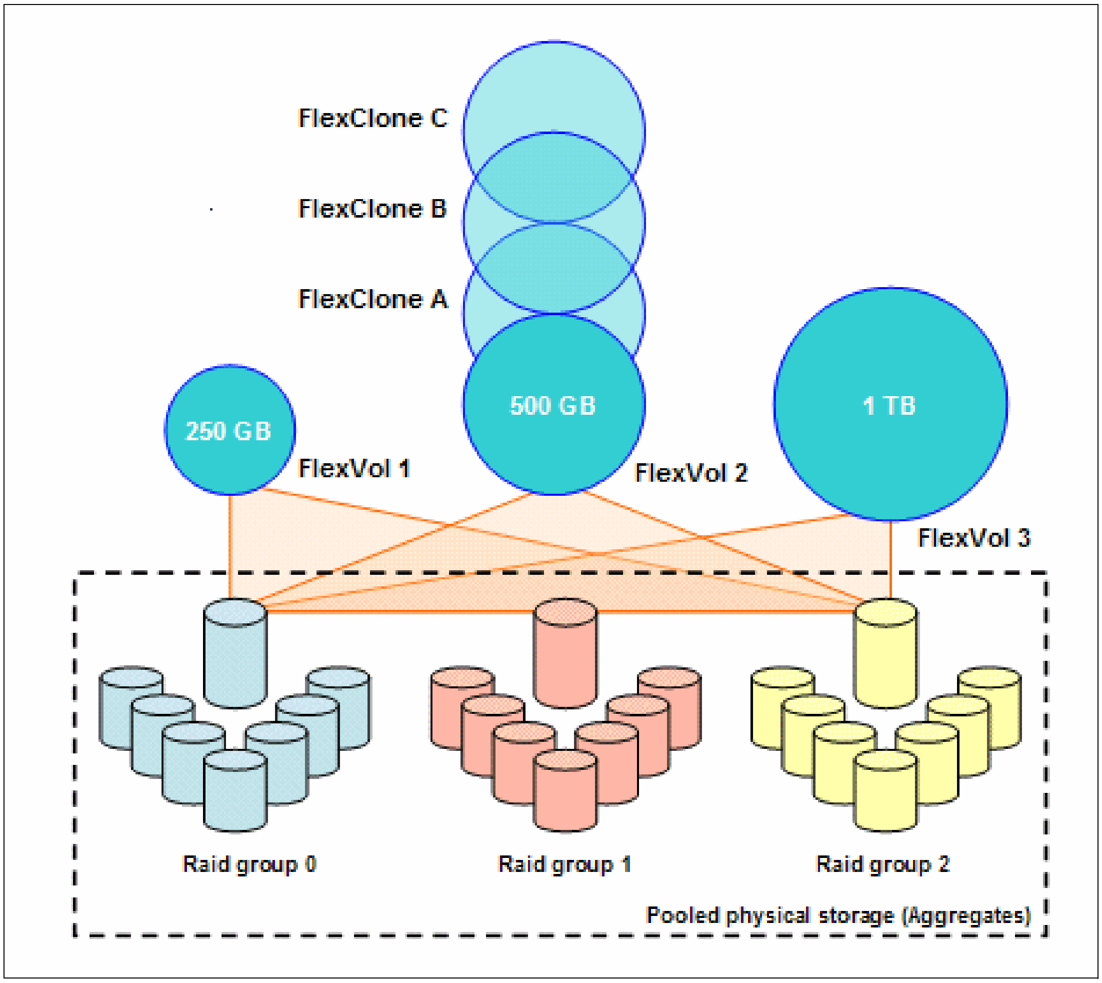

Figure 4-1 Overview of FlexVol and FlexClone

4.2 FlexVol for data protection

FlexVol technology is a ground breaking technology that comes embedded with DATA ONTAP software. *FlexVol*s are independent of the underlying physical storage. These are the logical entities that are sized, resized, managed, and moved independently of the underlying storage.

Volumes remain the primary unit of data management. Flexible volumes refer to logical entities, not (directly) to physical storage, and are transparent to the administer.

Flexible volumes are no longer bound by the limitations of the disks on which they reside. A *FlexVol* volume is simply a pool of storage that is sized based on how much data you store in it, rather than, on what your disk size dictates. A *FlexVol* volume is shrunk, or increased without any downtime. Flexible volumes have all the spindles in the *aggregate* available to them at all times. For I/O-bound applications, flexible volumes run much faster than equivalent-sized traditional volumes.

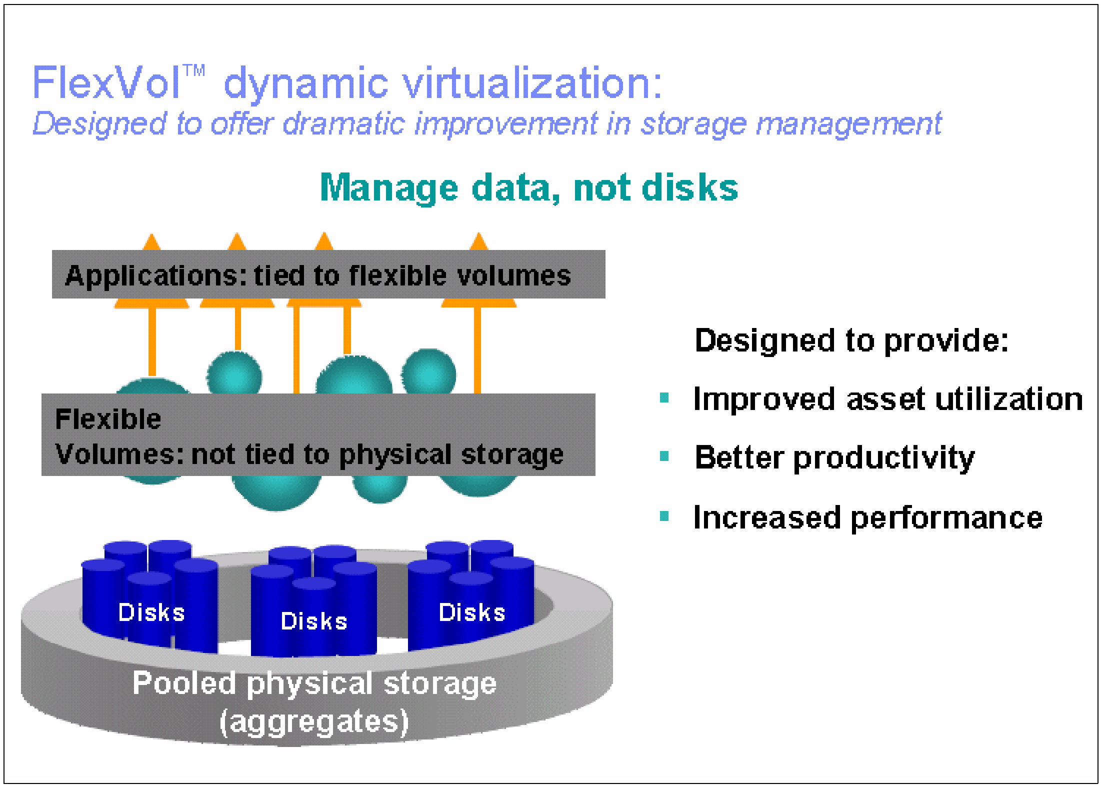

FlexVol™ dynamic virtualization:

Designed to offer dramatic improvement in storage management

Manage data, not disks

Applications: tied to flexible volumes
Flexible
Volumes: not tied to physical storage
Disks
Disks
Disks
Pooled physical storage
(aggregates)

Designed to provide:

- **Improved asset utilization**
- **Better productivity**
- **Increased performance**

Figure 4-2 An aggregate: a pool of many disks from which space is allocated to volumes

The size of *FlexVol* volumes are increased, or decreased quickly. Space that is allocated to *FlexVol* but not used, is quickly taken away, and reallocated to another *FlexVol* volume that needs it. The size of the *aggregate*(s) is also increased quickly.You can clone *FlexVol* volumes using *FlexClone* technology. A *FlexClone* volume represents a space efficient point-in-time copy (read/write), of the parent *FlexVol* volume, but can also turn into a fully independent *FlexVol* volume.

4.2.1 FlexVol creation

In the IBM N Series, *FlexVol* is created using a command line or *FilerView®*. In this section we discuss the creation of *FlexVol* using the command line.

1. From the command line enter the **vol create** command (see Example 4-1 on page 19).

Example 4-1 Vol create

```
vol create vol1 -l en_US aggr 10g
```

2. Check the status of your newly created volume with the **vol status** command use the **-v** option for detail.

Example 4-2 Check volume status

```
itsotuc1> vol status -v vol1
         Volume State      Status            Options
      vol1 online     raid4, flex  nosnap=off,
                                            nosnapdir=off,
                                            minra=off,
                                            no_atime_update=off,
                                            nvfail=off,
                                           ignore_inconsistent=off,
                                             snapmirrored=off,
                                             create_ucode=off,
                                             convert_ucode=off,
                                             maxdirsize=10470,
                                             schedsnapname=ordinal,
                                             fs_size_fixed=off,
                                             guarantee=volume,
                                            svo_enable=off,
                                             svo_checksum=off,
                                             svo_allow_rman=off,
                                             svo_reject_errors=off,
                                             no_i2p=off,
                                             fractional_reserve=100,
                                             extent=off,
                                             try_first=volume_grow
                Containing aggregate: 'aggr'

                Plex /aggr/plex0: online, normal, active
                    RAID group /aggr/plex0/rg0: normal
```

4.3 FlexClone for data protection

A *FlexClone* volume is a writable point-in-time image of a *FlexVol* volume, or another *Flexclone* volume, as illustrated in Figure 4-3 on page 20.

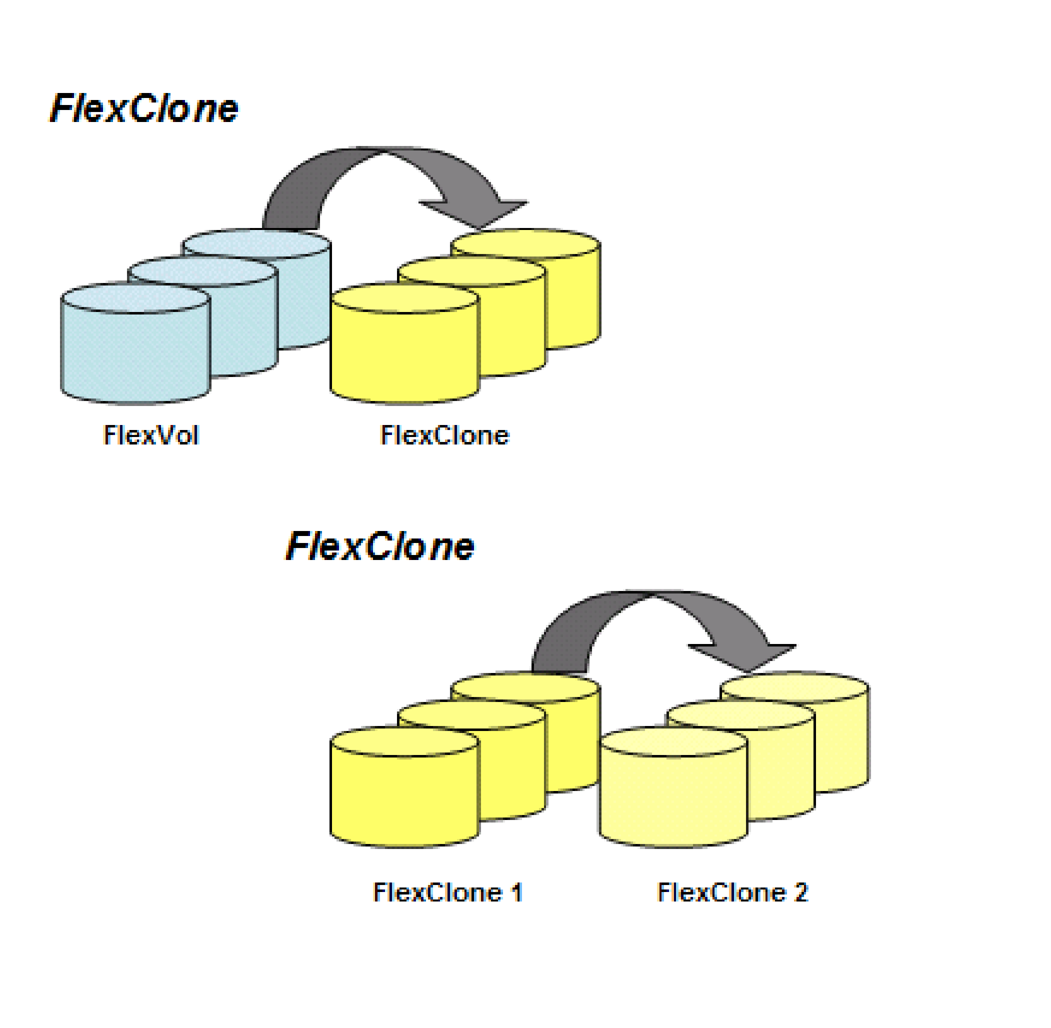

Figure 4-3 FlexClone created by cloning a FlexVol or FlexClone

FlexClone volumes have all the capabilities of a regular flexible volume, including growing, shrinking, and being the source for SnapShot copies, or the source for another clone.

Think of a *FlexClone* volume as a transparent writable layer in front of a SnapShot copy (Figure 4-4 on page 21). A *FlexClone* volume is writable, so it needs physical space to store the data that is written to the clone. It uses the same mechanism used by SnapShot copies to get available blocks from the *aggregate*. A Snapshot copy links to existing data that is overwritten in the parent, a *FlexClone* volume stores the data written to it on disk (using WAFL), and then links to the new data. The disk space associated with the SnapShot copy, and *FlexClone* is separate from the data in the parent *FlexVol* volume.

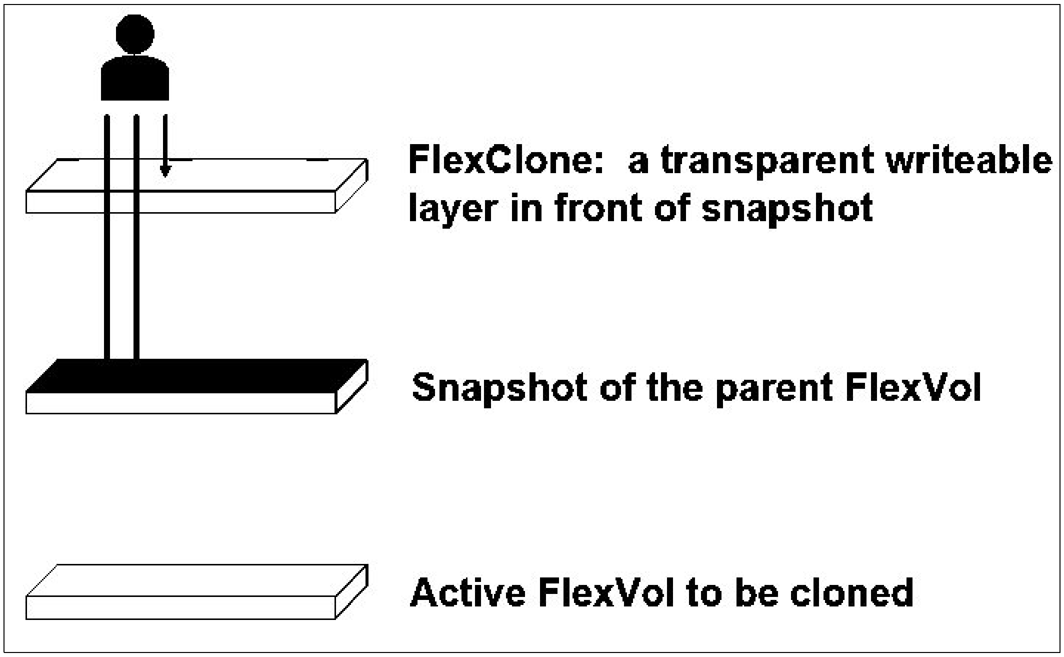

FlexClone: a transparent writeable layer in front of snapshot

Snapshot of the parent FlexVol

Active FlexVol to be cloned

Figure 4-4 A FlexClone volume as a transparent writable layer in front of a SnapShot

A SnapShot copy links to existing data that is overwritten in the parent. In contrast, a *FlexClone* volume stores the data written to it on disk (using WAFL), and then links to the new data (see Figure 4-5 on page 21). The disk space associated with the SnapShot copy, and *FlexClone* is separate from the data in the parent *FlexVol* volume.

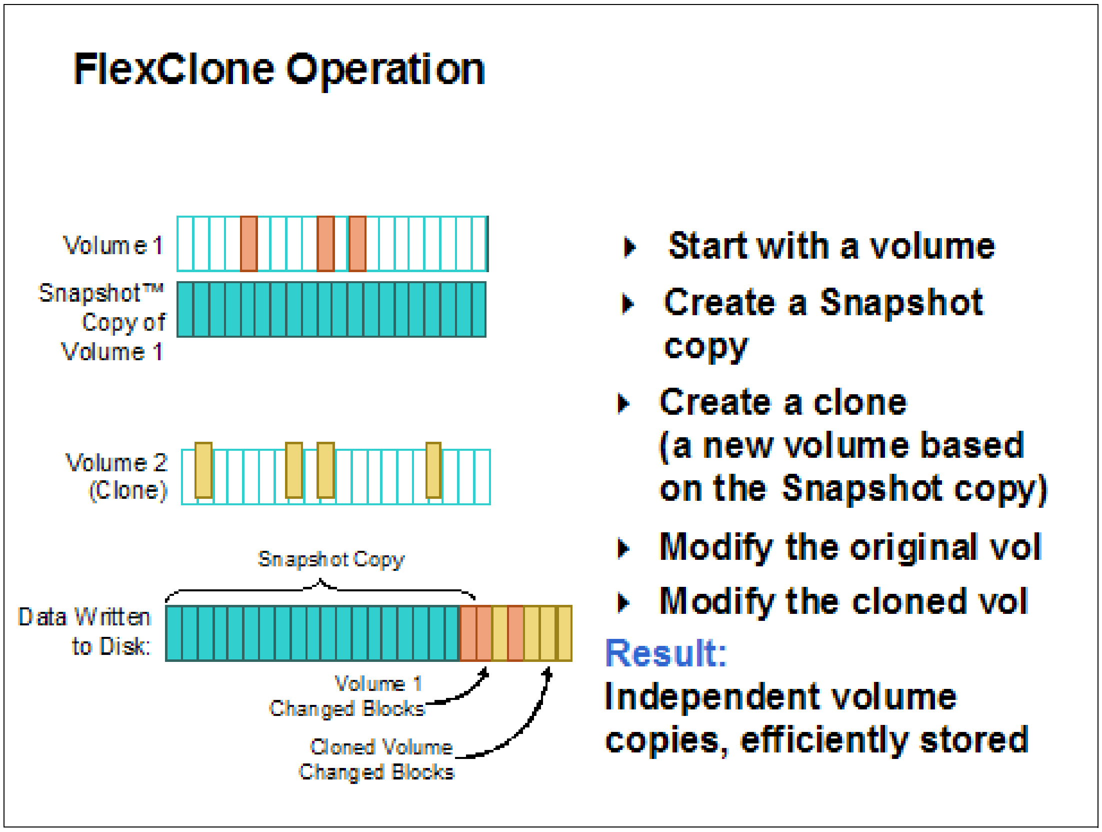

FlexClone Operation

- Start with a volume
- Create a Snapshot copy
- Create a clone (a new volume based on the Snapshot copy)
- Modify the original vol
- Modify the cloned vol

Result:
Independent volume copies, efficiently stored

Figure 4-5 FlexClone operation

When a *FlexClone* volume is first created, it needs to know the parent *FlexVol* volume and a SnapShot copy of the parent to use as its base. The SnapShot copy can already exist, or it is created automatically as part of the cloning process. The *FlexClone* volume gets a copy of the SnapShot copy metadata, and then updates its metadata as the clone volume is created.

The parent *FlexVol* volume can change independently of the *FlexClone* volume because the Snapshot copy is there to keep track of changes, and prevent the original parents blocks from being reused while the SnapShot copy exists. The same SnapShot copy is read-only, and is efficiently reused as the base for multiple *FlexClone* volumes.

Space is used efficiently, since the only new disk space used is either associated with the small amounts of metadata, updates, or additions to either the parent *FlexVol* volume, or the *FlexClone* volume.

FlexClone volumes appear to the storage administrator just like a *FlexVol* volume; that is, they look like a regular volume and have all of the same properties and capabilities.

FlexClone volumes enable administrators to access the destination mirror created by the IBM N Series SnapMirror product. Previously, it was necessary to break the mirror in order to make any changes to the destination copy. With *FlexClone*, an administrator can clone a SnapShot copy held in the mirror, and make it available for both reading, and writing at the remote site, while allowing the mirror facility to continue running unaffected.

4.3.1 FlexClone example

1. Run the status command on your volume to get the current status of the volume, and to identify the *aggregate* that it belongs to (Example 4-3 on page 22).

Example 4-3 Vol status command

```
itsotuc1> vol status -v vol1
         Volume State      Status            Options
      vol1 online     raid4, flex  nosnap=off,
                                            nosnapdir=off,
                                            minra=off,
                                            no_atime_update=off,
                                            nvfail=off,
                                           ignore_inconsistent=off,
                                             snapmirrored=off,
                                             create_ucode=off,
                                             convert_ucode=off,
                                             maxdirsize=10470,
                                             schedsnapname=ordinal,
                                             fs_size_fixed=off,
                                             guarantee=volume,
                                            svo_enable=off,
                                             svo_checksum=off,
                                             svo_allow_rman=off,
                                             svo_reject_errors=off,
                                             no_i2p=off,
                                             fractional_reserve=100,
                                             extent=off,
                                             try_first=volume_grow
                Containing aggregate: 'aggr'

                Plex /aggr/plex0: online, normal, active
                    RAID group /aggr/plex0/rg0: normal
```

2. Use the `df` command to check the available disk space on your volume (Example 4-4 on page 23).

Example 4-4 df command

```
itsotuc1> df -g
Filesystem          total       used      avail capacity  Mounted on
/vol/vol_0/   24GB        0GB       23GB        1%  /vol/vol_0/
/vol/vol_0/.snapshot       6GB        0GB        5GB       0%
/vol/vol_0/.snapshot
/vol/vol1/            8GB        0GB        7GB       0%  /vol/your_name/
/vol/vol1/.snapshot   2GB        0GB        2GB       0%  /vol/vol1/.snapshot
```

3. Again, use the df command to check the available disk space on the *aggregate* (Example 4-5 on page 23).

Example 4-5 df command

```
itsotuc1> df -Ag
Aggregate                total       used      avail capacity
vol1 56GB        0GB       56GB      0%
vol1/.snapshot        2GB        0GB        2GB       0%
Vfilers                  170GB        0GB     170GB       0%
Vfilers/.snapshot          8GB        0GB        8GB       0%
aggr_0                   56GB       30GB       26GB      53%
aggr_0/.snapshot           2GB        0GB        2GB       9%
aggr 56GB       10GB       46GB      18%
aggr/.snapshot            2GB       0GB        2GB       0%
```

4. Now we are ready to clone the existing volume (Example 4-6 on page 23).

Example 4-6 clone create command

```
itsotuc1> vol clone create vol1_clone -b vol1
Thu Nov  8 14:16:09 MST [itsotuc1: wafl.snaprestore.revert:notice]: Reverting
volume vol1_clone to a previous snapshot.
Creation of clone volume 'vol1_clone' has completed.
```

5. The `snap list` command shows one new volume named *vol1_clone*. It also lists the new SnapShot copy named *clone_vol1_clone.1*. Your *flexClone* creation is now successful (Example 4-7 on page 24).

Example 4-7 Successful FlexClone creation

```
itsotuc1> snap list
Volume vol_0
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  Nov 08 12:00  hourly.0
  1% ( 0%)    0% ( 0%)  Nov 08 08:00  hourly.1
  1% ( 0%)    0% ( 0%)  Nov 08 00:00  nightly.0
  2% ( 1%)    0% ( 0%)  Nov 07 20:00  hourly.2

Volume vol1
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
 29% (29%)    0% ( 0%)  Nov 08 14:16  clone_vol1_clone.1 (busy,vclone)

Volume vol1_clone
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  8% ( 8%)    0% ( 0%)  Nov 08 14:16  clone_vol1_clone.1
```

4.4 Summary

FlexVol technology decouples direct connections between volumes and their associated physical disks, vastly increasing flexibility, and storage efficiency. A new entity termed an *aggregate* provides the connection between the logical flexible volume and the underlying physical storage, and isolates the volume from this connection. A flexible volume is now able to stripe all of its data across the entire *aggregate*, improving performance for small volumes. They also help you to organize logical unit numbers (LUN)s according to host, application, or function.

Storage administrators have access to greater flexibility and performance. Flexible volumes provide unparalleled levels of storage virtualization, enabling IT staff to economically manage and protect enterprise data without compromise. *FlexClone* volumes are one of the many powerful features that make this possible, providing instantaneous writable volume copies that use only as much storage as necessary to hold new data.

FlexClone volumes enable, and simplify many operations. Application testing provides less risk, less stress, and higher service levels. Use *FlexClone* volumes to try out changes on clone volumes, and upgrade under tight maintenance windows by simply swapping tested *FlexClone* volumes for the originals. Data mining, and parallel processing benefit by using multiple writable *FlexClone* volumes from a single data set, all without using additional physical storage to hold the updates.

FlexClone volumes are used as online backup and disaster recovery volumes, immediately resuming read-write operations if a problem occurs. System deployment becomes much easier by cloning template volumes for testing, and rollout. IT operations benefit from multiple copies of production system that are used for testing and development, and refreshed as needed to mirror the live data.

5

SnapMirror

SnapMirror is a software feature that allows replication of a dataset between the N Series System Storage over a network for backup, or disaster recovery purposes. This chapter discusses SnapMirror as a data protection solution.

5.1 The need for SnapMirror

SnapMirror software provides a fast, flexible enterprise solution for mirroring, or replicating data over local, or wide area networks. SnapMirror is used for:

- Disaster recovery
- Remote enterprise-wide online backup
- Data replication for local read-only access at a remote site
- Application testing on a dedicated read-only mirror
- Data migration between IBM storage systems

SnapMirror technology is a key component of enterprise data protection strategies. If a disaster occurs at a source site, businesses can access mission-critical data from a mirror on another N Series System Storage , ensuring uninterrupted operation (see Figure 5-1 on page 28). Enterprise tape backups are made from SnapMirror, not a production system, reducing CPU load on the production system.

The N Series can be located virtually any distance from the source. It can be in the same building, or on the other side of the world, as long as the interconnecting network has the necessary bandwidth to carry the replication traffic that is generated.

SnapMirror software provides fast recovery in a disaster situation compared to restoring a file system, or qtree from tape. To assist customers in determining the economic impact to their company, resulting from total system downtime, you must realistically calculate costs per day, of business shutdown resulting from a disaster. As the cost of downtime rises for an organization, enterprises cannot afford to operate without a disaster recovery solution in place. SnapMirror allows organizations to quickly and easily implement an economical, and reliable disaster recovery solution.

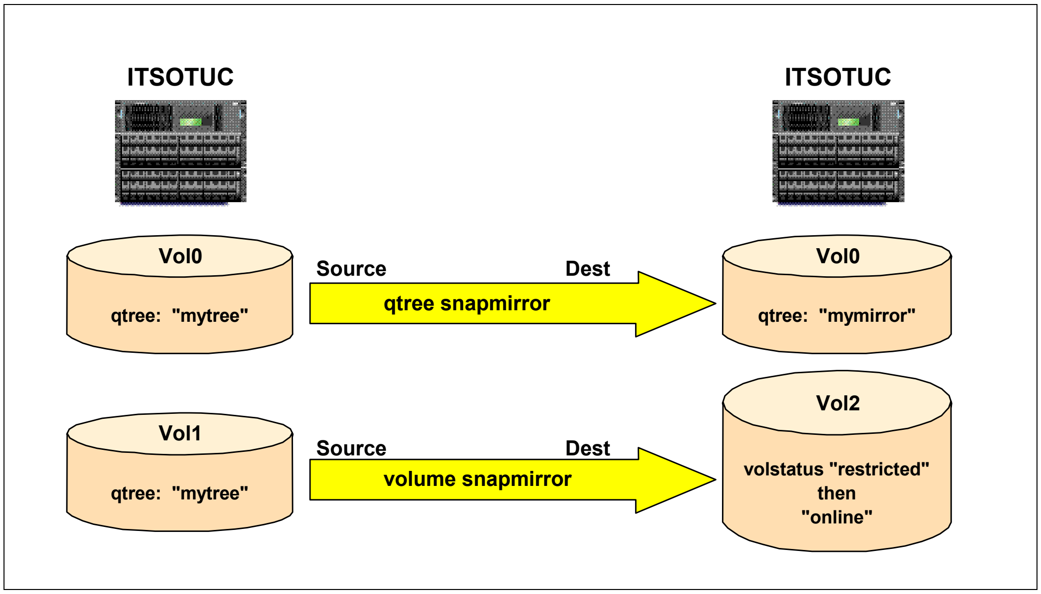

Figure 5-1 SnapMirror

Some environments require off-site storage, or off-site archiving. Tape archiving is costly. When a tape device is attached to a SnapMirror destination, data is moved to tape periodically. SnapMirror is also used for backup consolidation, and offloading tape backup overhead from production servers. This facilitates centralized backup operations, reducing backup administrative requirements at remote locations. Our solution dramatically reduces overhead from stressful backup operations, caused by small backup windows on production storage systems. Because backup operations are not occurring on the production systems, small backup windows are not as important.

SnapMirror technology leverages the WAFL SnapShot capability to create, and update a copy of a source volume, or qtree on the N Series. The mirror copy is accessible to users in read-only mode. SnapMirror software makes a baseline transfer of the data (comparable to a full backup for tape backups). The initial transfer is accomplished through a network connection, or by restoring a tape on the destination. SnapMirror then updates the mirror by replicating only new, or changed data blocks. Mirror copies are consistent because SnapMirror software operates on consistent SnapShot copies.

System administrators specify the intervals, and times SnapMirror SnapShot copies are created. Determining this schedule depends upon how much the data changes during the day, how up-to-date the mirror needs to be, CPU usage on the source, and available network bandwidth.

SnapMirror operates in three modes:

- Asynchronous mode:
  - In asynchronous mode, SnapMirror performs incremental, block-based replication as frequently as once per minute. Performance impact on the source N Series is minimal, as long as the system is configured with sufficient CPU and disk I/O resources. Because asynchronous replication is periodic, SnapMirror is able to consolidate writes and conserve network bandwidth. There is minimal impact on write throughput and write latency. As soon as data is written to the NVRAM of the source N Series , applications using this data are free to continue processing, without waiting for the data to reach a destination system. Updates take place in the background, so the application does not experience any additional transaction latency.
- Synchronous:
  - Certain environments have very strict uptime requirements. All data that is written to one site must be mirrored to a remote site, or system synchronously. SnapMirror in synchronous mode sends updates from the source to the destination as they occur, rather than according to a predetermined schedule. This guarantees that data written on the N Series source is protected on the destination, even if the entire source system fails. In synchronous mode, SnapMirror immediately replicates all data written to the source file system. This guarantees zero data loss in the event of a failure, but can have a significant performance impact. It is not necessary, or appropriate for all applications.
- Semi-synchronous:
  - Semi-synchronous mode minimizes data loss in a disaster, while also minimizing how much the replication impacts the performance of the source system. Unlike asynchronous mode, which replicates volumes, or qtrees, synchronous, and semi-synchronous modes work only with volumes. Complete details are found in chapter 18 of the *IBM System Storage N Series Redbook SG24-7129-00*.

5.2 SnapMirror creation

Using the N Series, you can create a SnapMirror by using the command line interface, or by using *FilerView*. In this section, we illustrate the creation of SnapMirror using the command line.

Example 5-1 on page 30 shows you how to setup SnapMirror between two separate N Series nodes. In the example, we consider two N Series storage systems: *itsotuc1* and *itsotuc2*

1. Telnet to *itsotuc1* and *itsotuc2*
2. Verify that the snapmirror license is installed on both the systems. If not, install the Snapmirror license (Example 5-1).

Example 5-1 Install snapmirror license

```
itsotuc1>license add PRO23UB
A snapmirror site license has been installed.
   snapmirror enabled
```

3. In *Itsotuc1*:

- Create a source volume (Example 5-2 on page 30), for example *vol_primary* on the storage, and confirm its online status with the **`vol status`** command.

Example 5-2 Create volume and check the status

```
itsotuc1>vol create vol_primary aggr1 25m
Creation of volume 'vol_primary' with size 25m on containing aggregate
'aggr1' has completed
itsotuc1> vol status vol_primary
Volume State Status Options
vol_primary online raid_dp, flex  create_ucode=on
            convert_ucode=on
Containing aggregate: 'aggr1'
```

4. In *itsotuc2*:

- Create a destination volume, for example *vol_secondary* on the storage and confirm its online status with the **`vol status`** command as in the above example.

Note: It is required that the destination volume used for mirroring is larger than, or equal to the source volume. Before proceeding with snap mirroring, verify this.

5. Verify that the volume, *vol_secondary* is of the same size or larger than the volume on the other filer, *vol_primary* using the **vol status -b** command. (*see* Example 5-3 on page 31)

Example 5-3 Verify that the source and destination volumes are of acceptable size

```
itsotuc1> vol status -b vol_primary
Volume             Block Size (bytes)  Vol Size (blocks)   FS Size (blocks)
------            ------------------  ------------------  ----------------
vol1_prim          4096                2621440             2621440
itsotuc2> vol status -b vol_secondary
Volume             Block Size (bytes)  Vol Size (blocks)   FS Size (blocks)
------             ------------------  ------------------  ----------------
vol1_primary       4096                2621440             2621440
```

6. In itsotuc2, restrict the destination volume using **vol restrict** command. (See Example 5-4 on page 31)

Example 5-4 Restrict the destination volume

```
itsotuc2> vol restrict vol_secondary
Volume 'vol_secondary' is now restricted
```

7. In both the nodes, specify the destination hosts that are allowed to access the source. In Example 5-5 on page 31, we allowed cross mirroring by defining the partner nodes in each of their snapmirror access options using the **snapmirror.access option** to set and verify the current configuration and specify the hosts allowed.

Note: This step is essential even when the SnapMirror is created using GUI/FilerView

Example 5-5 Specify access to snap mirror relationship

```
itsotuc1> options snapmirror.access host=itsotuc2
itsotuc1> options snapmirror.access
snapmirror.access            host=itsotuc2

itsotuc2> options snapmirror.access host=itsotuc1
itsotuc2> options snapmirror.access
snapmirror.access            host=itsotuc1
```

8. Start snapmirror on both the systems, using the **snapmirror on** command.
9. In both storage systems, create CIFS shares for source and destination volumes:

- In *itsotuc1*:
  - **itsotuc1> cifs shares -add srcvol /vol/vol_primary**
- In *itsotuc2*:
  - **itsotuc2> cifs shares -add destvol /vol/vol_secondary**

10. Access SnapMirror source volume *srcvol* share, and create new files, or directories (Figure 5-2 on page 32) using the following sign on specification:

- Share Login: **administrator**
- Password: **<CIFS password you configured during setup>**

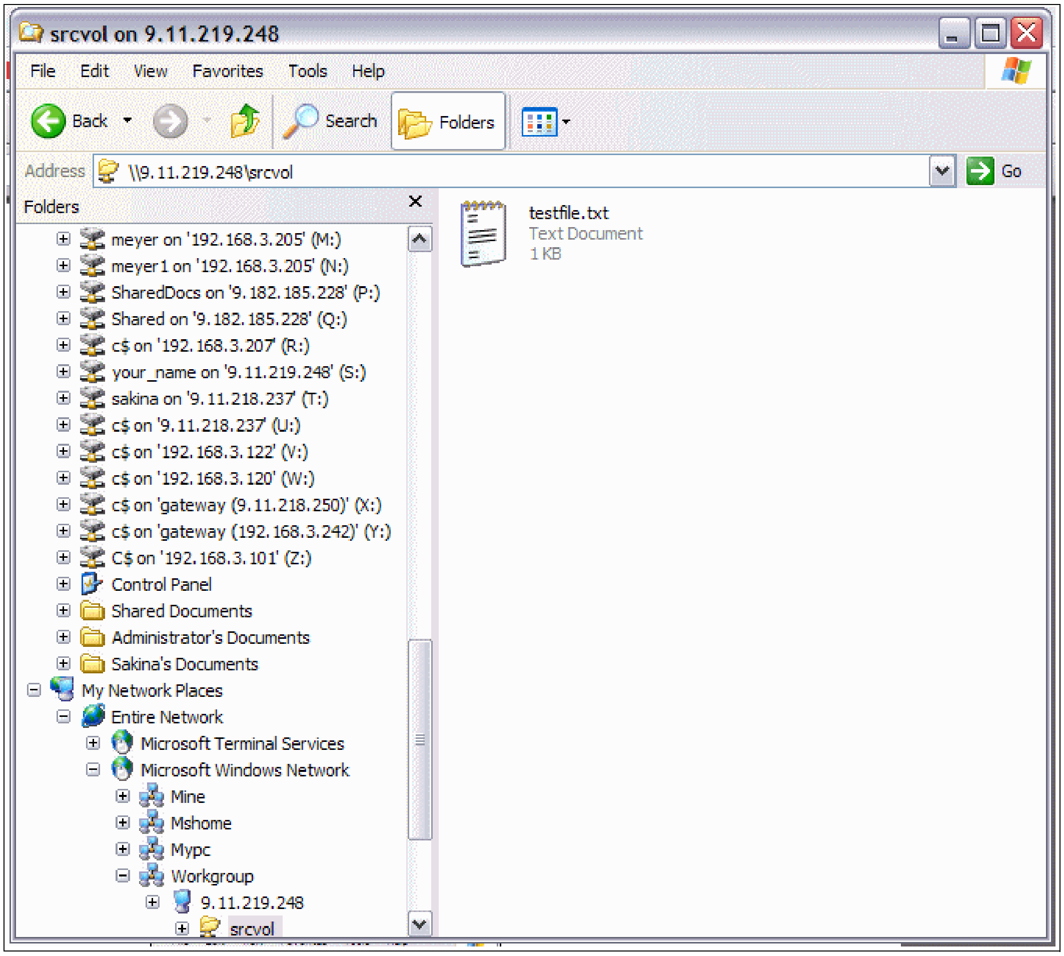
srcvol on 9.11.219.248
File
Edit
View
Favorites
Tools
Help
Back
Search
Folders
Address
\\9.11.219.248\srcvol
Go
Folders
meyer on '192.168.3.205' (M:)
meyer1 on '192.168.3.205' (N:)
SharedDocs on '9.182.185.228' (P:)
Shared on '9.182.185.228' (Q:)
c$ on '192.168.3.207' (R:)
your_name on '9.11.219.248' (S:)
sakina on '9.11.218.237' (T:)
c$ on '9.11.218.237' (U:)
c$ on '192.168.3.122' (V:)
c$ on '192.168.3.120' (W:)
c$ on 'gateway (9.11.218.250)' (X:)
c$ on 'gateway (192.168.3.242)' (Y:)
C$ on '192.168.3.101' (Z:)
Control Panel
Shared Documents
Administrator's Documents
Sakina's Documents
My Network Places
Entire Network
Microsoft Terminal Services
Microsoft Windows Network
Mine
Mshome
Mypc
Workgroup
9.11.219.248
srcvol
testfile.txt
Text Document
1 KB

Figure 5-2 Source volume directory

11. In *itsotuc2*, start the first SnapMirror baseline transfer from *itsotuc1* to *itsotuc2* using the **`snapmirror initialize`** command (Example 5-6 on page 32).

Example 5-6 Initialize snap mirroring on destination volume

```
itsotuc2> snapmirror initialize-S itsotuc1:vol_primary itsotuc2:vol_secondary
Transfer started.
Monitor progress with 'snapmirror status' or the snapmirror log
```

12. Enter the **`snapmirror status`** command to check the transfer status (Example 5-7 on page 32).

Example 5-7 Check transfer status

```
itsotuc2> snapmirror status
Snapmirror is on.
Source               Destination            State          Lag        Status
itsotuc1:vol_primary itsotuc2:vol_secondary Uninitialized - Transferring
itsotuc2> snapmirror status
Snapmirror is on.
Source               Destination            State          Lag        Status
itsotuc1:vol_primary itsotuc2:vol_secondary Snapmirrored   00:00:26   Idle
```

13. Access the destination volume, try to edit, or delete the test file. You get a warning similar to Figure 5-3 on page 33.

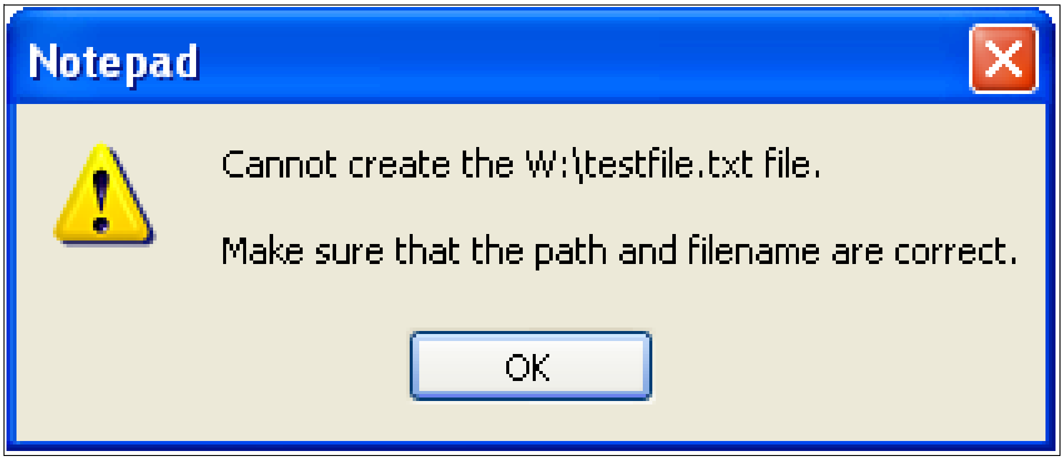

Cannot create the W:\testfile.txt file.

Make sure that the path and filename are correct.

Figure 5-3 writing to restricted volume

> **Note:** With SnapMirror active, you cannot edit or delete the files in the destination volume.

5.2.1 Breaking the SnapMirror relationship

As we discuss SnapMirror technology as a data protection technique, it is essential to know how SnapMirror helps in a crisis.

As shown above, destination volume cannot be accessed when snapmirror is active. When catastrophe happens and the source SnapMirror storage fails, you have to break the mirrors and make the destination SnapMirror storage read-write so that the business can continue on the destination storage.

Use the following procedures to make the destination volume active:

1. Turn the SnapMirror off on both the storage systems using **`snapmirror off,`** and check the status using **`snapmirror status.`** (See Example 5-8 on page 33).

Example 5-8 Turn SnapMirror off

```
itsotuc2> snapmirror off
itsotuc2> snapmirror status
Snapmirror is off.
Source             Destination          State         Lag        Status
itsotuc1:vol_primary  itsotuc2:vol_secondary  Snapmirrored   00:05:26   Idle
```

2. Break the mirror of *itsotuc1:vol_primary* on *itsotuc2:vol_secondary* (See Example 5-9 on page 34).

Example 5-9 Break the SnapMirror

```
itsotuc2>snapmirror break itsotuc2:vol_secondary
snapmirror break: Destination vol_secondary is now writable.
Volume size is being retained for potential snapmirror resync. To grow the volume,
and not expect to resync, set vol option fs_size_fixed to off.
```

After the source storage is back to normal **`snapmirror resync`** and **`snapmirror update`** command should be used to resynchronize broken mirrors and to re-establish the SnapMirror relationship between source and destination volumes.

5.3 Summary

In summary, SnapMirror technology in the IBM System Storage N Series is used for the following:

- Data replication to a local, or remote site:
  - Fast data replication and failover-can help reduce downtime in case of a failure at the primary site
- Fast recovery from disaster (no lengthy restores from tape are required)
- Replicating data to remote systems
- Creating writable:
  - *FlexClone* copies of real-time data (useful in application testing and development)
- Migrating data between IBM storage systems

6

MetroCluster

MetroCluster is another data protection mechanism. MetroCluster is an integrated, high-availability, and business-continuance solution for clustering two N Series N5000, or N7000 storage systems located miles apart.

6.1 Overview of MetroCluster

In the event of catastophe, organizations look for a sound solution to facilitate rapid recovery. But current plans, often with added high-availability requirements, are costly, and difficult to implement and administer. They are also complicated by a storage infrastructure that includes data centers at sites located miles apart. The data at the remote site must be kept up-to-date so that it can serve as the principal data in case of a disaster at the primary site. MetroCluster functionality of is illustrated in Figure 6-1 on page 36.

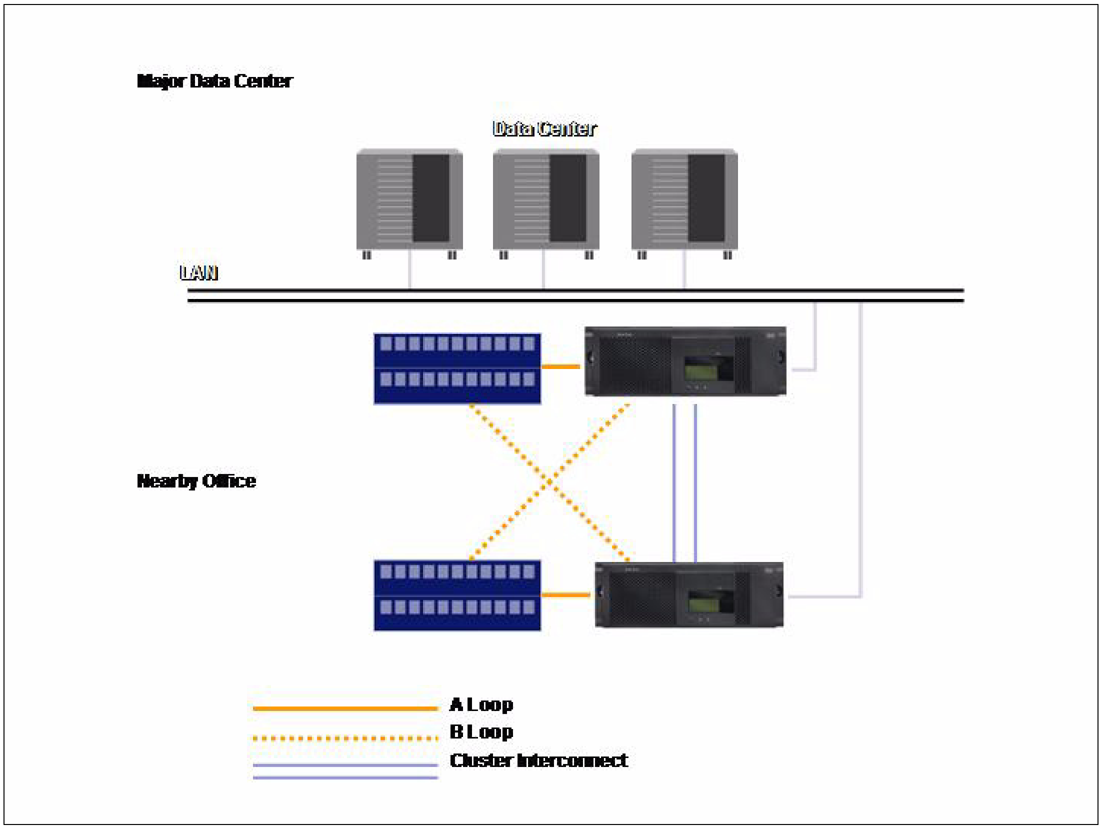

Figure 6-1 MetroCluster

MetroCluster is a cost-efficient high availability and disaster recovery solution. It provides an enterprise solution for high availability over WAN. MetroClusters extend the cluster failover capabilities from primary to a remote site. This makes it an important component of enterprise data protection strategies. Using MetroCluster, if a disaster occurs on a primary site, businesses can continue to run, and access data from the remote site.

MetroCluster is designed to provide mission-critical applications with redundant storage services, for site-specific disasters, by tolerating site-specific disasters with minimal interruption to mission-critical applications, and zero data loss by synchronously mirroring data between two sites.

6.2 MetroCluster implementation

There are two ways to implement MetroCluster according to business needs, and the distances required:

- Stretch MetroCluster
- Fabric MetroCluster

6.2.1 Stretch MetroCluster

In Stretch MetroCluster, the controllers and expansion shelves are attached to Fibre Channel switches, and the switches have GBICs to communicate across the WAN to one another. SyncMirror® is built into the MetroCluster so that every write is written to two separate expansion units to two separate *aggregate* groups. Stretch MetroCluster is available on N5000 and N7000 series.

In the Gateway Interoperability matrix, Stretch MetroCluster is available on IBM N Series Gateway models.

Imagine a disk subsystem like the IBM DS4800, separating the controllers miles apart, and maintaining the two disk groups to ensure failover, instead of having two IBM DS4800 disk subsystems, and synchronously mirroring between the two. Figure 6-2 on page 37 illustrates Stretch MetroCluster.

Note: Stretch Cluster requires MetroCluster and SyncMirror license.

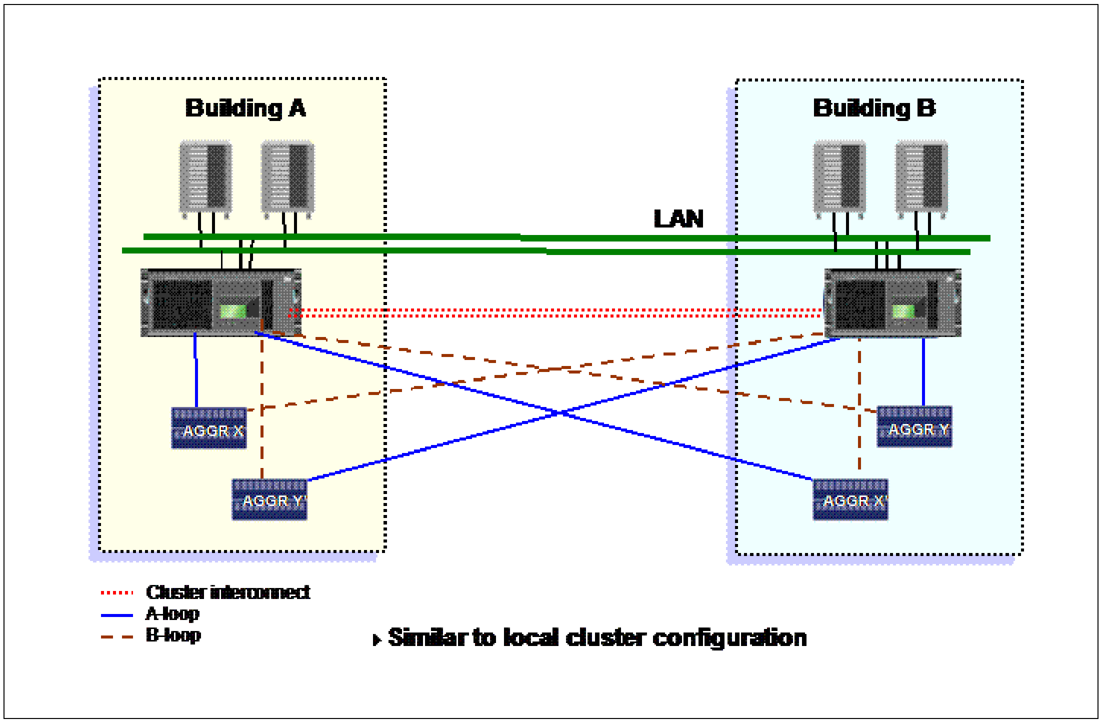

Figure 6-2 Stretch MetroCluster

Note: Stretch MetroCluster provides a disaster recovery option at distances up to 500 meters between each IBM N Series system

6.2.2 Fabric MetroCluster

Fabric MetroCluster provides disaster recovery of up to 100 kms distance. Figure 6-3 on page 38 illustrates Fabric MetroCluster (MetroCluster > 500 m).

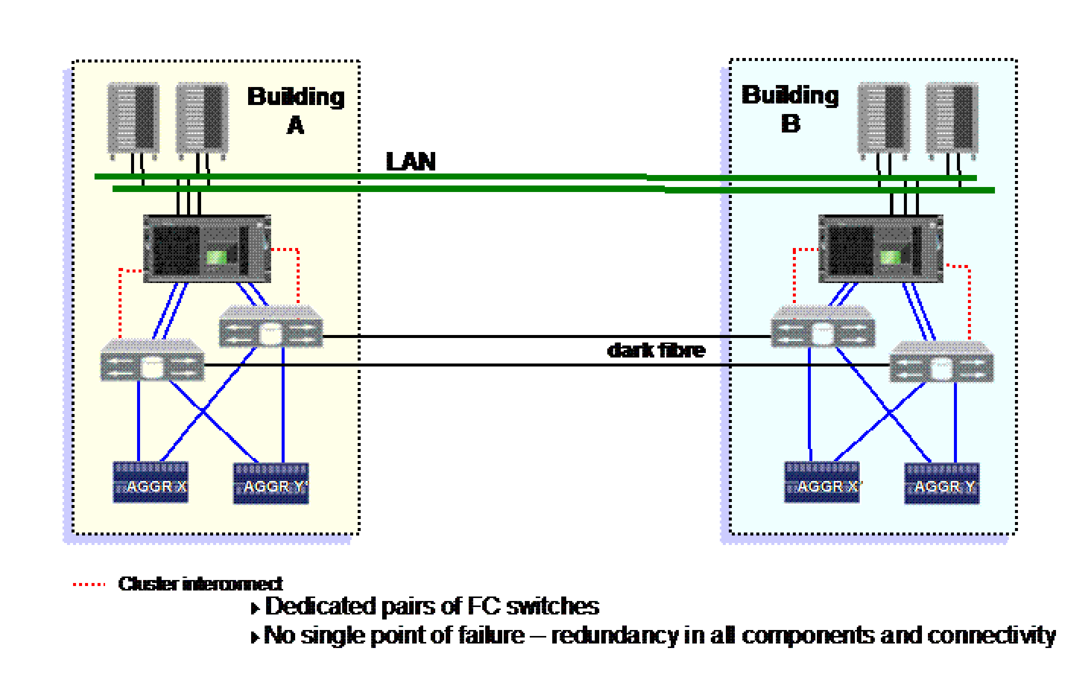
Building A
LAN
Building B
dark fibre
AGGR X
AGGR Y'
AGGR X'
AGGR Y
Cluster interconnect
Dedicated pairs of FC switches
No single point of failure – redundancy in all components and connectivity

Figure 6-3 Fabric MetroCluster

- Fabric MetroCluster:
  - Available on N5000 and N7000 series
  - Supported on N7000 series with Data ONTAP version 7.2.3 or higher
  - Uses switches for longer distance disaster recovery solution
  - Enables the use of IBM 2005-16B switches for longer-distance disaster recovery solutions
  - Requires the use of a dual-port HBA (#1006), and a cluster interconnect adapter (#1018)

Using dense wave division multiplexing switches

Dense wave division multiplexing (DWDM) is a method of multiplexing multiple channels of fiber optic-based protocols, such as ESCON®, Fibre Channel, FICON®, and Gbit Ethernet, onto physical cable by assigning different wavelengths of light (colors) to each channel, and fanning it back out at the receiving end. The major players in the enterprise class DWDM marketplace are: Nortel Networks, Cisco (ONS 15540), and Lucent.

DWDMs are data link layer 2 tools. The typical DWDM machine does not perform any switching, routing, or protocol conversion.

Figure 6-4 on page 39 shows a fabric MetroCluster installation using DWDM switches.

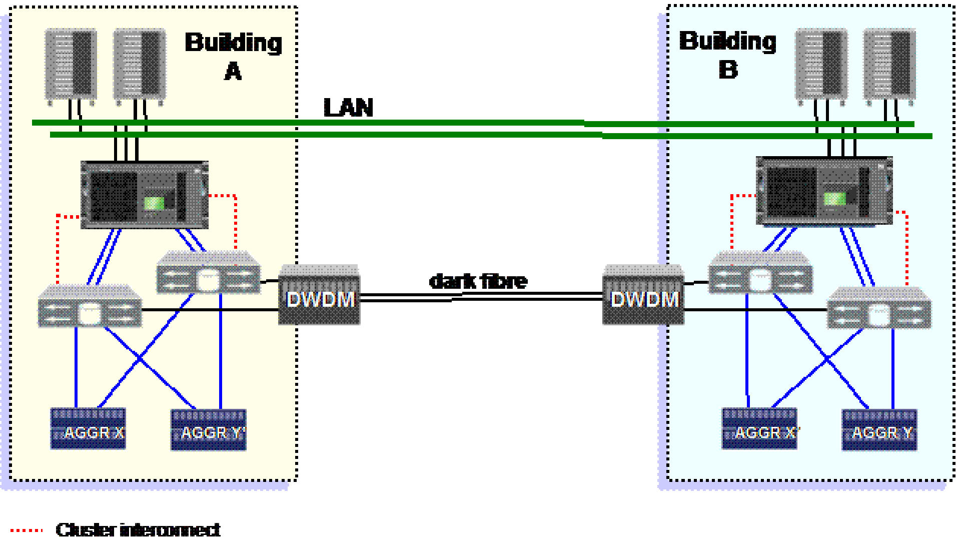

- Allows MC to plug into your existing Fibre
- All DWDM equipped certified with Brocade supported

Figure 6-4 DWDM environment

6.3 Benefits of MetroCluster

Using MetroCluster in your enterprise provides the following benefits:

- Designed to be a simple-to-administer solution that extends failover capability from within a data center to a remote site
- Designed to provide replication of data from the primary site to a remote site, helping keep data at the remote site current
- Combining failover and data replication aids in disaster recovery by helping to prevent the loss of data in less time than is otherwise possible
- Extends Clustered failover capabilities from a primary site to a remote site
- Replicates data from the primary site to the remote site to ensure that data there is completely up-to-date and available
- If Site A goes down, MetroCluster allows you to rapidly resume operations at a remote site minutes after a disaster as illustrated in Figure 6-5 on page 40

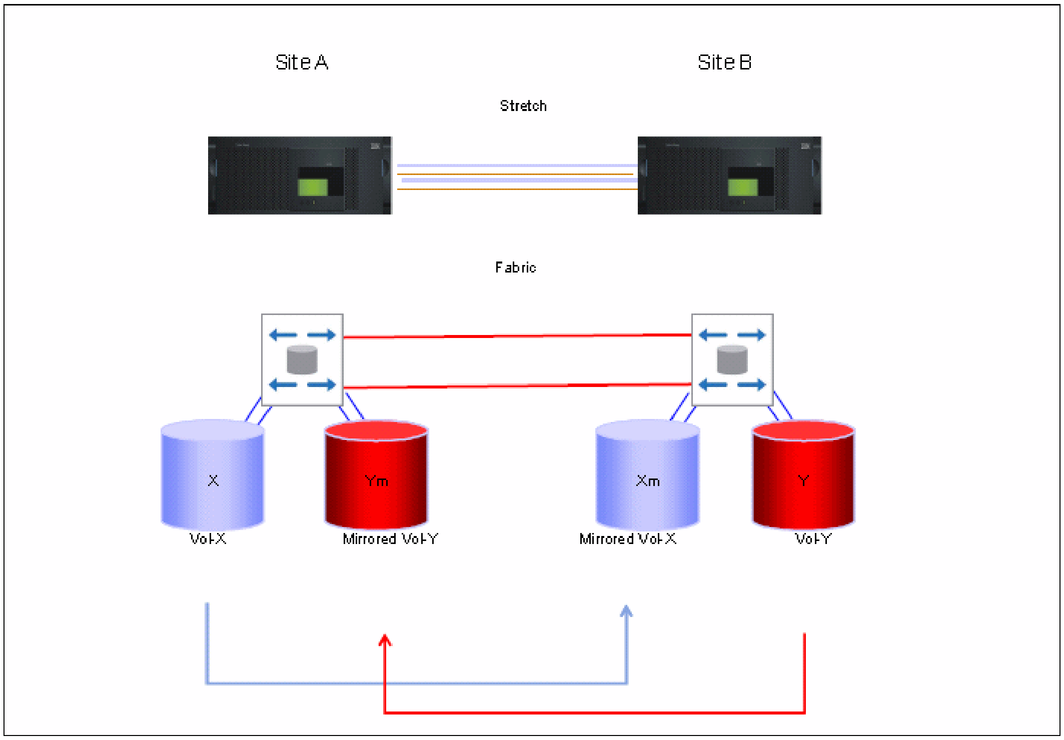

Figure 6-5 Disaster recovery using MetroCluster

7

SnapVault

SnapVault technology in the N Series is discussed in this chapter. SnapVault provides a centralized, disk-based backup solution for multiple IBM systems.

This chapter includes the following

- Overview of SnapVault technology
- How SnapVault works
- SnapVault configuration using CLI
- Benefits of SnapVault

7.1 Overview of SnapVault

SnapVault software from IBM N Series is a reliable and economical way to protect enterprise data. It offers many significant advantages over traditional backup methods. Although SnapVault is deployed in configurations designed to emulate the legacy backup methods it replaces, the full value of the solution is realized only by making a significant shift in the way you think about backup and recovery. SnapVault renders many common backup policies, and schedules obsolete.

Figure 7-1 on page 42 shows the basic operation of the SnapVault architecture with regular, and SnapLock volumes depicted. Open systems protocols are supported with IBM System Storage N Series. The connection is through WAN or LAN.

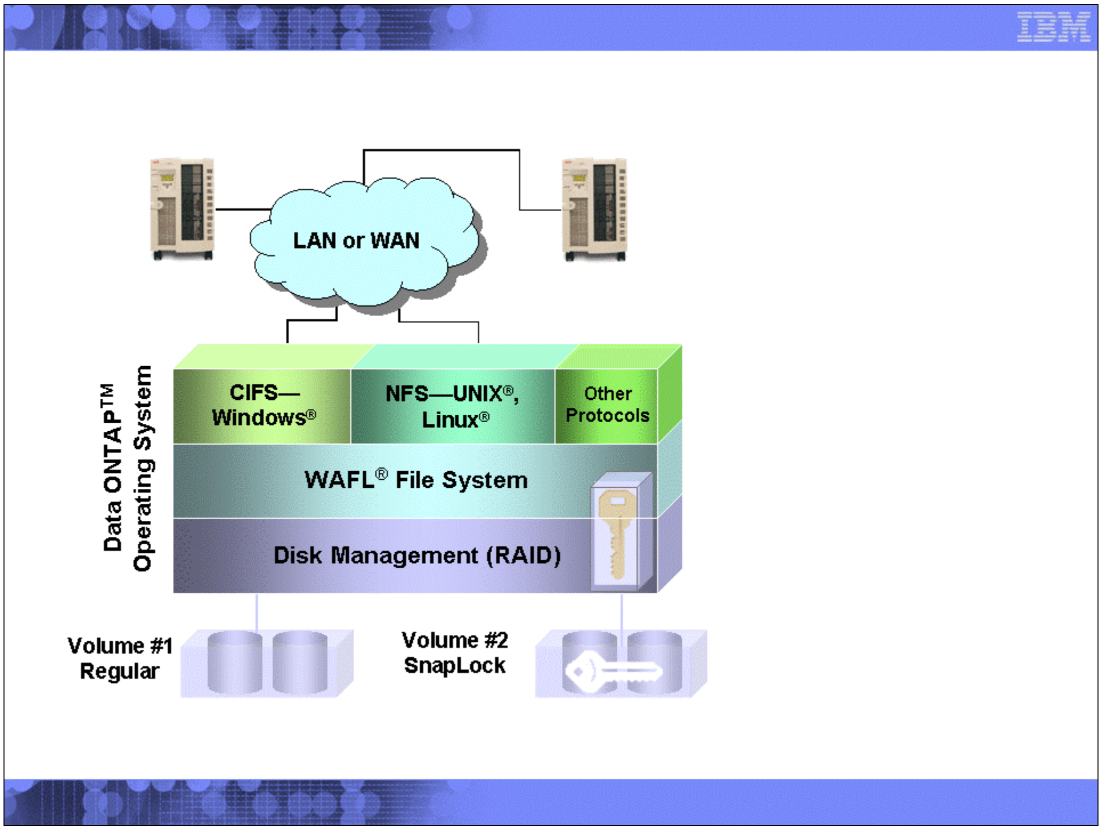

Figure 7-1 Basic SnapVault operation

7.2 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 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 can 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 are performed, the SnapVault secondary stores the data in a WAFL file system and creates a SnapShot image of the volume for the data being 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 consumes 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 (Figure 7-2 on page 43).

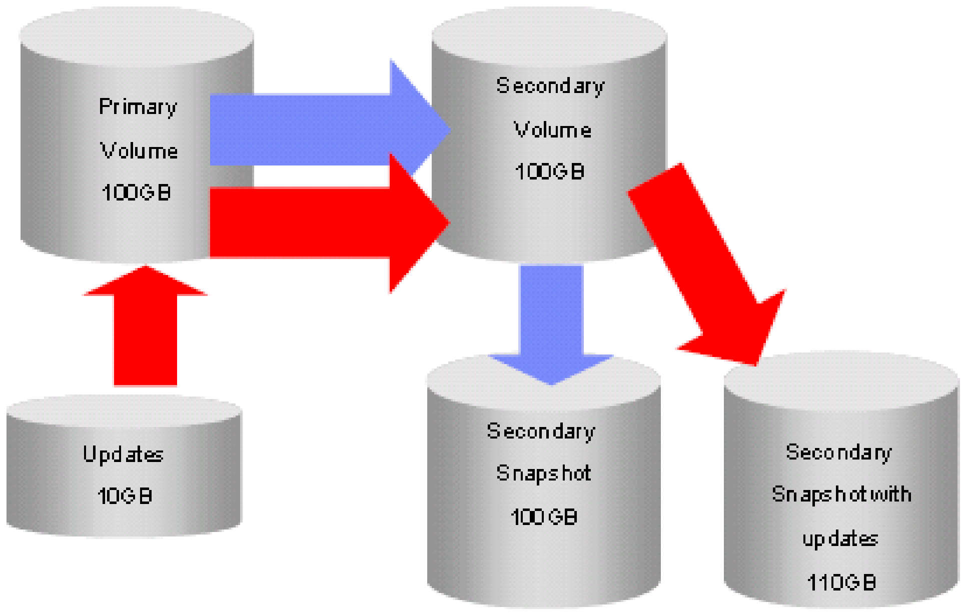

Figure 7-2 SnapVault to secondary

7.3 SnapVault example

With the N Series, SnapVault can be configured using the command line interface.

The following example how to setup SnapVault between two separate N Series nodes. In the example, we consider two N Series storage systems: *itsotuc1* and i*tsotuc2*.

The home directories are in a qtree on *itsotuc1,* called /vol/vol1/users, the database is on *itsotuc1,* in the volume called /vol/oracle:

1. Telnet to *itsotuc1* and *itsotuc2*
2. License SnapVault
3. Enable SnapVault on both the systems. See Example 7-1 on page 44.

Example 7-1 Install SnapVault license

```
itsotuc1>license add ABCDEFG
itsotuc1> options snapvault.enable on
itsotuc1> options snapvault.access host=itsotuc2

itsotuc2> license add HIJKLMN
itsotuc2> options snapvault.enable on
itsotuc2> options snapvault.access host=itsotuc1
```

4.Schedule SnapShot copies on the SnapVault primary, *itsotuc1*. Refer Example 7-2 on page 44.

Example 7-2 Schedule SnapShot copies on SnapVault primary

```
a.
itsotuc1>snap sched vol1 0 0 0
itsotuc1>snap sched oracle 0 0 0
b.
itsotuc1>snapvault snap sched vol1 sv_hourly 22@0-22
c.
itsotuc1>snapvault snap sched vol1 sv_daily 7@23
```

a. Turn off the normal SnapShot schedules, which is replaced by SnapVault SnapShot schedules

b. Set up schedules for the home directory hourly SnapShot copies

   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.

c. Set up schedules for the home directory daily SnapShot copies

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

5. Schedule SnapShot copies on the SnapVault secondary, itsotuc2. Refer Example 7-3 on page 45.

Example 7-3 Schedule SnapShot copies on SnapVault secondary

```
a.
itsotuc2> aggr create sv_flex 10
itsotuc2> vol create vault sv_flex 100g
b.
itsotuc2>snap sched vault 0 0 0
c.
itsotuc2>snapvault snap sched -x vault sv_hourly 4@0-22
d.
itsotuc2>snapvault snap sched -x vault sv_daily 12@23@sun-fri
e.
itsotuc2>snapvault snap sched vault sv_weekly 13@23@sat
```

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

b. Turn off the normal SnapShot schedules, which is replaced by SnapVault SnapShot schedules

c. Set up schedules for the hourly backups

   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 a copy, updates the SnapVault qtrees with new data from the primary, and then takes a SnapShot copy on the destination volume, called *sv_hourly.0*.

d. Set up schedules for the daily backups

   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 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*.

e. Set up schedules for the weekly backups

   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, create and retain the weekly copies on the secondary only.

6. 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 know which qtrees to back up, or where to store them on the secondary. SnapShot copies are taken on the primary, but no data is transferred to the secondary.

   To provide SnapVault with this information, use the `snapvault start` command on the secondary (See Example 7-4 on page 45).

Example 7-4 Perform baseline transfer

```
itsotuc2> snapvault start -S itsotuc1:/vol/vol1/users /vol/vault/itsotuc1_users
itsotuc2> snapvault start -S itsotuc1:/vol/oracle/ /vol/vault/oracle
```

7.4 Benefits of SnapVault

When it comes to backup, SnapVault has several benefits, and options to protect your data, and aid high availability.

7.4.1 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. 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. Reliability and speed of recovery are the main reasons why full backups are done so frequently. 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. And also, if a full backup is done frequently, it ensures lower restore time since only few incremental backups are being restored.

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 are backed up once the initial baseline copy is complete. This reduces the source load, network bandwidth consumption, and overall media costs.

7.4.2 Self service restores

One of the unique benefits of SnapVault is that users do not require special software, or privileges to perform a restore of their own data. Users who 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.

7.4.3 Consistent security

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 perform safely.

8

SnapLock

This chapter discusses the SnapLock software in the N Series. The SnapLock function is designed to deliver high performance, and high-security data function to a disk-based nearline, and primary IBM System Storage N Series.

This chapter contains the following:

- Overview of SnapLock
- How SnapLock works
- Benefits of SnapLock

8.1 Overview of SnapLock

SnapLock is an advanced storage solution that provides an alternative to traditional optical (write-once-read-many) WORM storage systems for non-rewritable data. SnapLock is a license-based, open-protocol feature that works with application software to administer non-rewritable storage of data.

The SnapLock function helps manage the permanence, accuracy, integrity, and security of data by storing business records in an inalterable form, and providing fast online accessibility for long periods of time.

Storage platform requirements (for compliance and business needs) can be broadly divided into three areas:

- Security and confidentiality
- Flexibility
- Data permanence

A complete compliance solution must deliver in all three areas, SnapLock and ONTAP does, as illustrated in Figure 8-1 on page 48

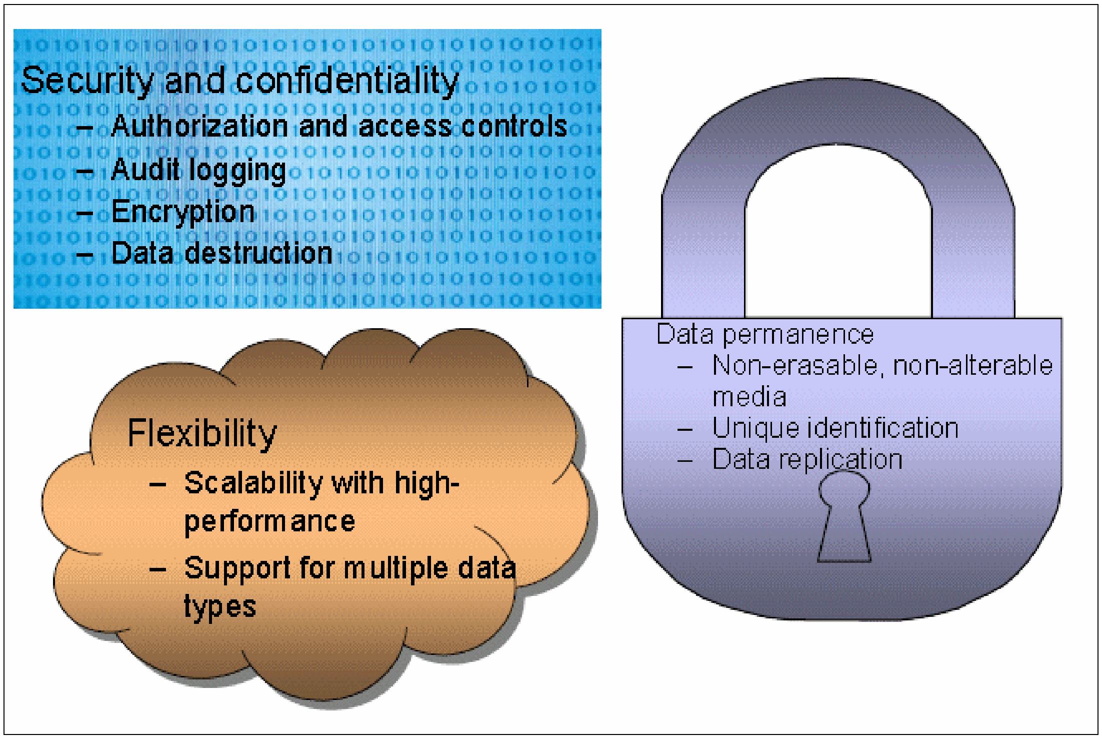
Security and confidentiality
– Authorization and access controls
– Audit logging
– Encryption
– Data destruction
Flexibility
– Scalability with high-performance
– Support for multiple data types
Data permanence
– Non-erasable, non-alterable media
– Unique identification
– Data replication

Figure 8-1 Storage platform requirements

SnapLock is available in two forms:

- SnapLock compliance
- SnapLock enterprise

8.1.1 SnapLock compliance

SnapLock compliance is designed to assist organizations in implementing a comprehensive archival solution for meeting the Security Exchange Commission (SEC), or governmental regulations for data retention. Records and files committed to WORM storage on a SnapLock.

Compliance volume cannot be altered or deleted before the expiration of their retention period. Moreover, a SnapLock compliance volume *cannot* be deleted until all data stored on it has passed its retention period, and deleted by the archival application, or some other process.

8.1.2 SnapLock enterprise

SnapLock enterprise is geared towards assisting organizations with meeting self-regulated, and best practice guidelines for protecting digital assets with WORM type data storage. Data stored as WORM on a SnapLock enterprise volume is protected from alteration or modification with one main difference from SnapLock compliance:

- Data being stored is not for regulatory compliance, a SnapLock enterprise volume *can* be deleted, including the data it contains, by an administrator.

A comparison of SnapLock compliance and SnapLock enterprise is illustrated in Figure 8-2 on page 49.

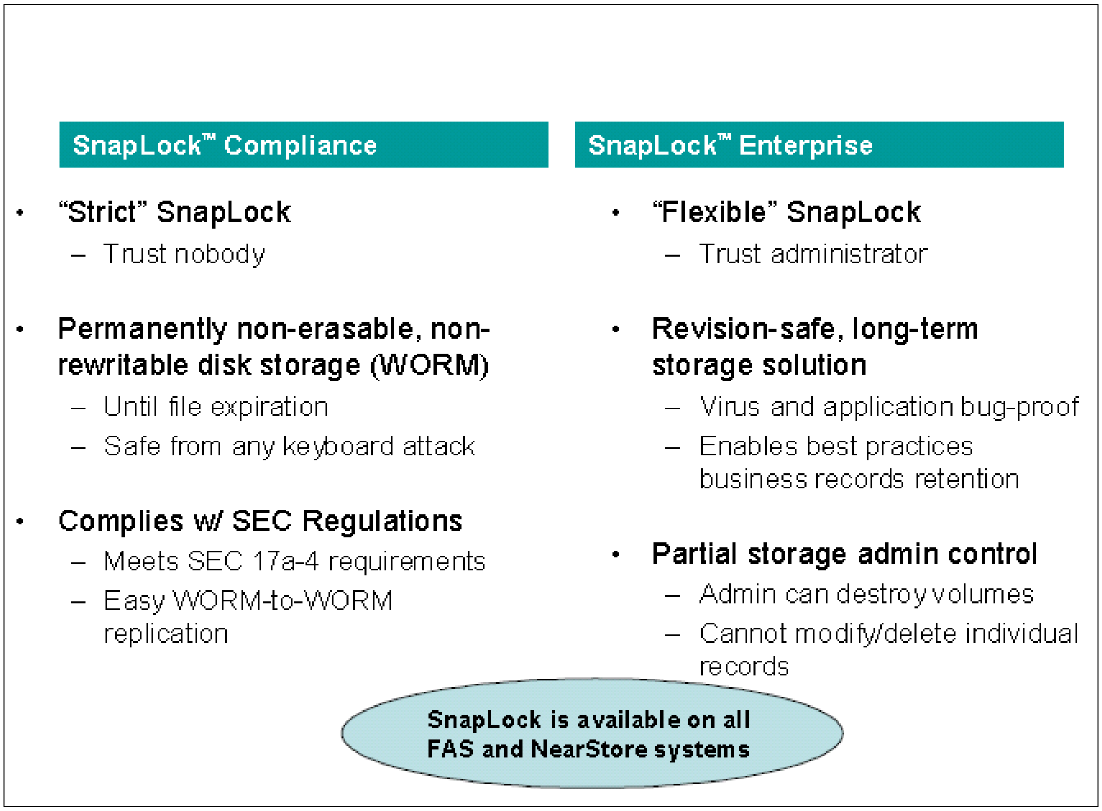

SnapLock™ Compliance

- **"Strict" SnapLock**
  - Trust nobody
- **Permanently non-erasable, non-rewritable disk storage (WORM)**
  - Until file expiration
  - Safe from any keyboard attack
- **Complies w/ SEC Regulations**
  - Meets SEC 17a-4 requirements
  - Easy WORM-to-WORM replication

SnapLock™ Enterprise

- **"Flexible" SnapLock**
  - Trust administrator
- **Revision-safe, long-term storage solution**
  - Virus and application bug-proof
  - Enables best practices business records retention
- **Partial storage admin control**
  - Admin can destroy volumes
  - Cannot modify/delete individual records

SnapLock is available on all FAS and NearStore systems

Figure 8-2 Comparison of SnapLock compliance and Snaplock enterprise

8.2 Implementing SnapLock

SnapLock compliance and SnapLock enterprise are extensions to the Data ONTAP operating system (see Figure 8-3 on page 50), which has a ten-year proven track record in online data storage.

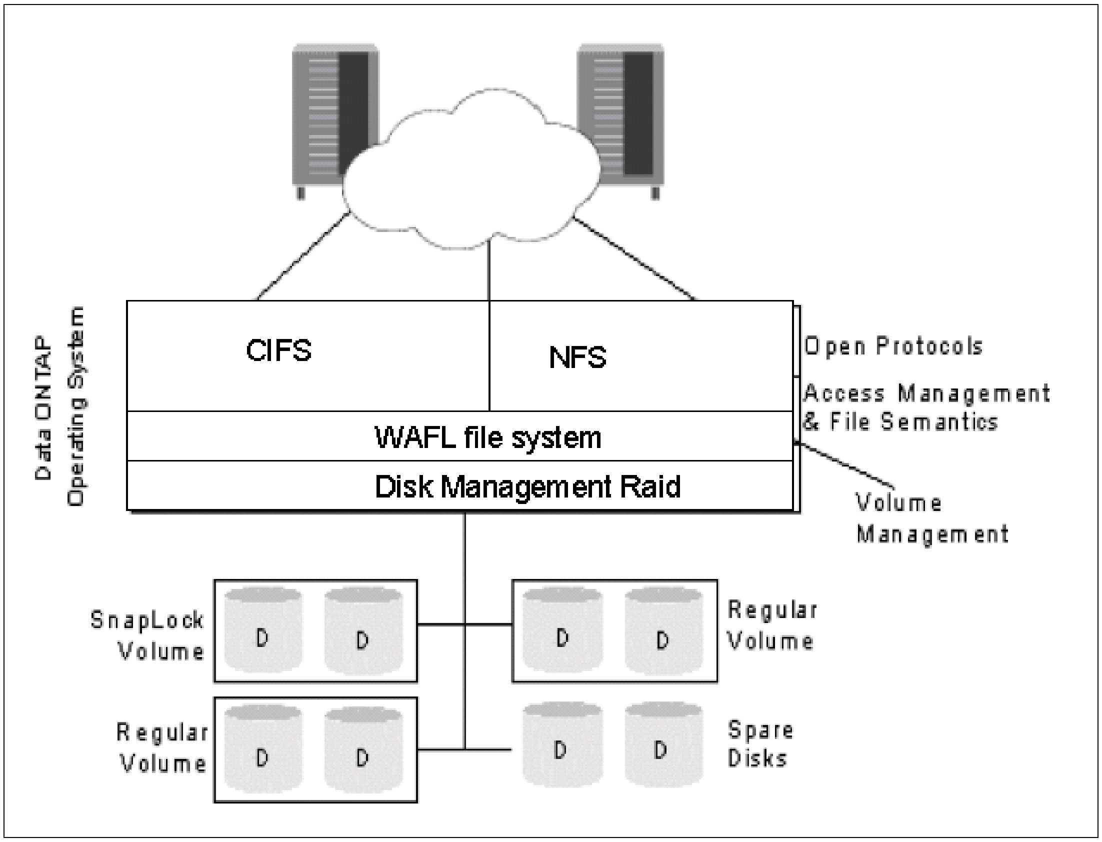

Figure 8-3 SnapLock integrated with ONTAP

Data ONTAP provides a complete infrastructure for storage, including RAID protection for data, a suite of tools and products to promote high data availability, and open protocol connectivity for data access. The same hallmarks of Data ONTAP (such as ease of deployment, ease of management, and ease of administration) also apply to both SnapLock software products.

8.3 Benefits of SnapLock

- WORM and security capabilities enable best business practices

  The use of SnapLock addresses issues faced by growing business requirements for WORM data storage, and alleviates issues inherent with traditional WORM storage solutions. SnapLock allows companies to implement the data permanence functionality of traditional WORM storage in an easier-to-manage, faster-access, lower-cost magnetic disk-based solution.

- Maximize return on investment (ROI), and total cost of ownership (TCO)

  As technology has improved, the evolution of WORM data storage (which started with paper and microfiche and progressed to optical) has now arrived at a new best-of-breed solution: IBM N Series configured with SnapLock Compliance, and SnapLock Enterprise software for high levels of data integrity, retention, and low TCO. Storing, and accessing massive amounts of information about economical, high-capacity, and easily expandable N Series storage solutions ensures the maximum return on your IT infrastructure investment.

9

LockVault

LockVault software integrates SnapLock and SnapVault technologies to create a solution specifically designed to help businesses address regulatory compliance requirements for unstructured data. In addition, the N Series with LockVault offers backup, and disaster recovery support for comprehensive, and integrated unstructured data compliance, and protection.

The following topics are discussed in this chapter:

- Compliance requirements
- N Series compliance offerings
- Uses and benefits of LockVault

9.1 Overview of LockVault

LockVault is an unparalleled solution offered by IBM System Storage N Series. It extends the regulated compliance for unstructured data. LockVault is the result of tight integration between the N Series SnapLock, and SnapVault product lines. LockVault combines SnapLock and SnapVault features into a single, unified solution. Figure 9-1 on page 54 shows an example of LockVault backup, recovery, and regulatory compliance.

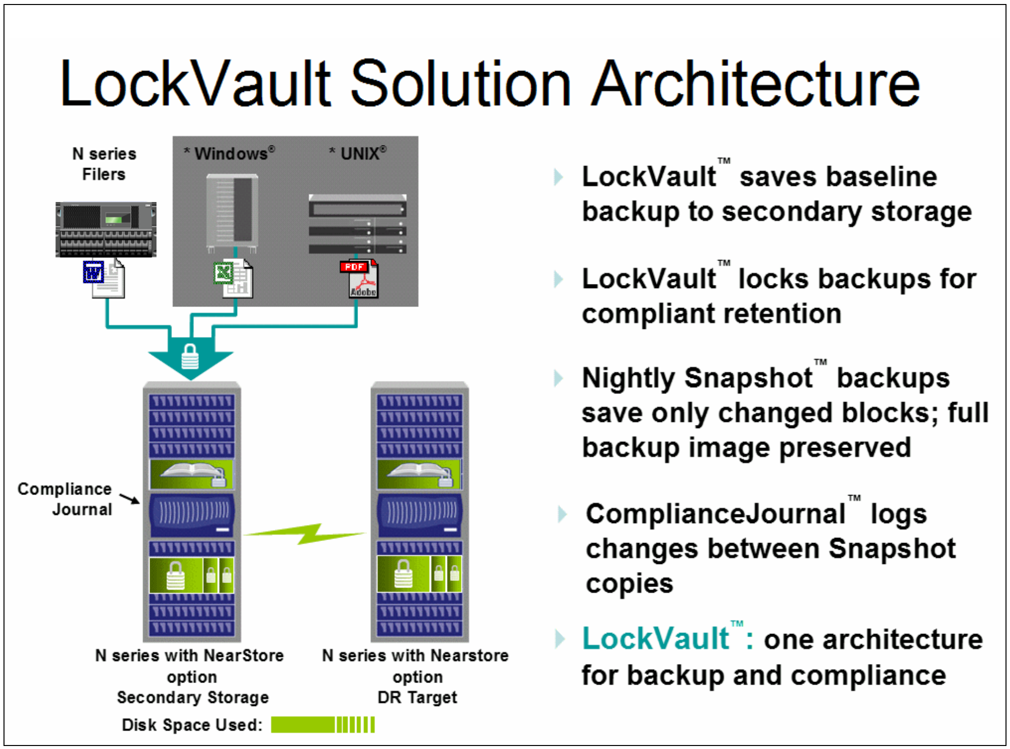

LockVault Solution Architecture

- LockVault™ saves baseline backup to secondary storage
- LockVault™ locks backups for compliant retention
- Nightly Snapshot™ backups save only changed blocks; full backup image preserved
- ComplianceJournal™ logs changes between Snapshot copies
- LockVault™: one architecture for backup and compliance

N series Filers
* Windows®
* UNIX®
Compliance Journal
N series with NearStore option Secondary Storage
N series with Nearstore option DR Target
Disk Space Used:

Figure 9-1 LockVault helps to unify backup and compliance

LockVault delivers a capacity-efficient regulatory solution by:

- Making backups compliant (one copy of data serves two purposes)
- Saving only block-level incremental changes
- Storage-efficient (block incremental) daily SnapShot copies backed up to secondary storage (using SnapVault technology), and protected against modification or deletion, until a specified retention date (using SnapLock technology)
- Integrating with Open Systems SnapVault (OSSV)
- Creating a compliance solution for open systems without compliant storage

9.2 LockVault uses and benefits

LockVault usage offers the following benefits:

- LockVault combines the features of SnapLock and SnapVault into a single, unified solution, providing unification of backup, recovery, and regulatory compliance.
- The initial focus on the financial and insurance industries mitigates risk. LockVault eliminates the need to rely on manual ,or policy-based methods of identifying and isolating records subject to regulatory compliance rules.
- LockVault delivers fast access for search and discover. Nightly compliant archives of the enterprise are available online for instant search, retrieval, or restore.
- LockVault minimizes storage capacity consumption. The block-level incremental scheme uses less than 1/20 of the capacity consumed by a traditional tape backup scheme over a one-year period.
- LockVault simplifies infrastructure deployment and management. A single, unified platform that is easy to deploy and manage can handle requirements of unstructured, structured, and semi-structured data.
- LockVault imparts flexibility. Open protocols-archival avoids the complexity and performance penalties of API-based solutions and assures true protection against obsolescence or vendor lock-in.
- SnapMirror can also be used to make duplicate copies of the LockVault backup-compliant images because SnapMirror supports WORM-to-WORM remote replication
- SnapVault is designed primarily for unstructured data compliance

9.2.1 Compliance drivers and requirements

To mitigate legal risk, companies want solutions that make it easy, and cost effective to enforce corporate wide retention policies, ensure immutability for records that are retained (as a non-repudiation tool), and have a way to dispose of them securely when they expire. There are many regulations worldwide which dictate the way businesses have to store information. In addition, enterprises store information securely to protect their intellectual property, and defend themselves against litigation. When we look at the way all these regulations and internal corporate governance requirements impact data storage, the requirements fall into two major categories:

1. Data permanence
2. Privacy and security

Some of the compliance drivers and requirements are listed in Figure 9-2 on page 56.

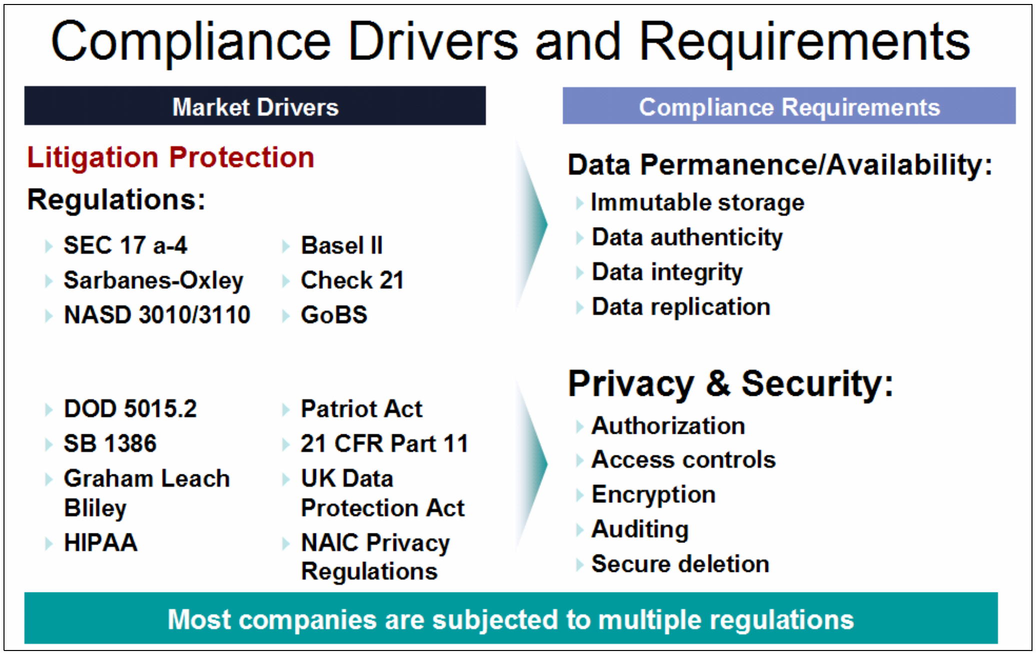

Compliance Drivers and Requirements

Market Drivers

Litigation Protection

Regulations:

- SEC 17 a-4
- Sarbanes-Oxley
- NASD 3010/3110
- Basel II
- Check 21
- GoBS
- DOD 5015.2
- SB 1386
- Graham Leach Bliley
- HIPAA
- Patriot Act
- 21 CFR Part 11
- UK Data Protection Act
- NAIC Privacy Regulations

Compliance Requirements

Data Permanence/Availability:

- Immutable storage
- Data authenticity
- Data integrity
- Data replication

Privacy & Security:

- Authorization
- Access controls
- Encryption
- Auditing
- Secure deletion

Most companies are subjected to multiple regulations

Figure 9-2 Compliance drivers and requirements

Gap in todays compliance solutions

When you review the solutions available for data permanence, privacy, and security across all data types in todays market, the most significant void exists in the area of data permanence for unstructured data (see Figure 9-3 on page 56).

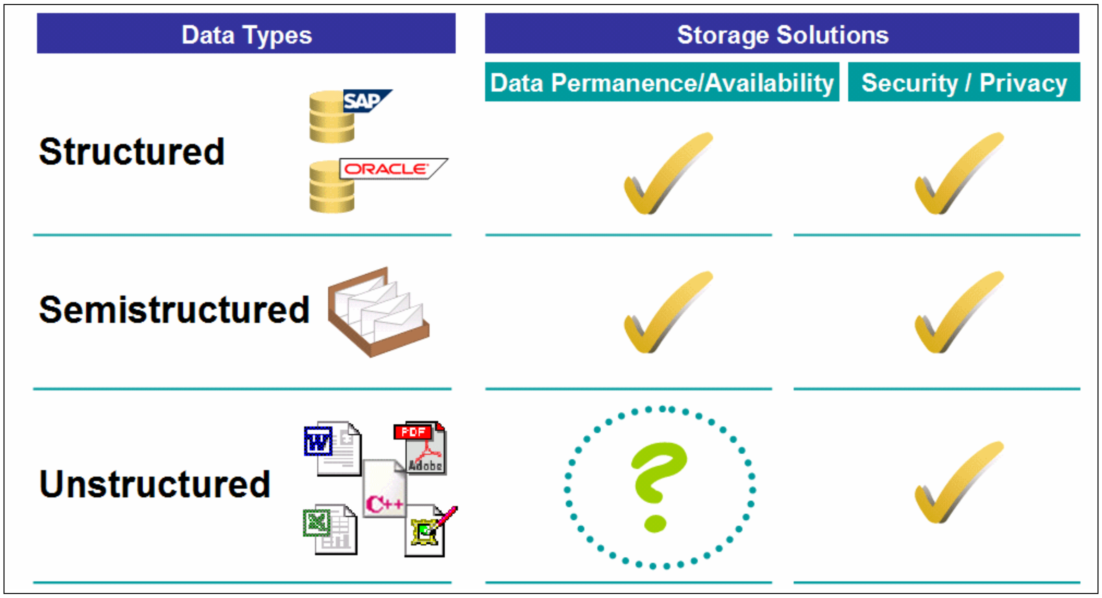

Data Types	Storage Solutions	
	Data Permanence/Availability	Security / Privacy
Structured	✓	✓
Semistructured	✓	✓
Unstructured	?	✓

SAP
ORACLE
W
PDF
Adobe
C++

Figure 9-3 Gap in today's data compliance solutions

- Database archival applications combined with WORM storage provide a good solution for structured data permanence. Examples of database archival vendors include Princeton, Softech, and Outerbay. These applications when combined with WORM storage solutions like SnapLock solve the data permanence needs with structured data. Similarly, email archival vendors like KVS, and IXOS, and document archival vendors like FileNet® and Documentum, in conjunction with SnapLock solves semi-structured data permanence.
- Privacy and security features available in applications, databases, operating systems, and directory infrastructures like Active Directory® reasonably address privacy and security issues with all forms of data.

While the bulk of enterprise data (over 50% by most estimates) is unstructured, there are no practical solutions to addressing data permanence of unstructured data. Unstructured data is strewn enterprise-wide in home directories, file systems, web servers, application servers, etc. It is becoming increasingly clear to several large enterprises, specifically in financial services vertical, that portions of these data can *ALSO* fall under regulatory purview.

9.2.2 N Series compliance offerings

The IBM N Series compliance architecture eliminates the need to add yet another silo of storage for compliance. The IBM N Series flexible storage architecture allows customers to meet all enterprise data management needs: primary storage, backup and recovery, disaster recovery, reference storage, and compliance storage.

LockVault covers all data types, including structured data like databases, semi-structured data, such as mail, and unstructured data, such as presentations, spreadsheets, and documents. Using a combination of SnapLock and LockVault, and the security features of the N Series enables you to meet your compliance requirements, see Figure 9-4 on page 57.

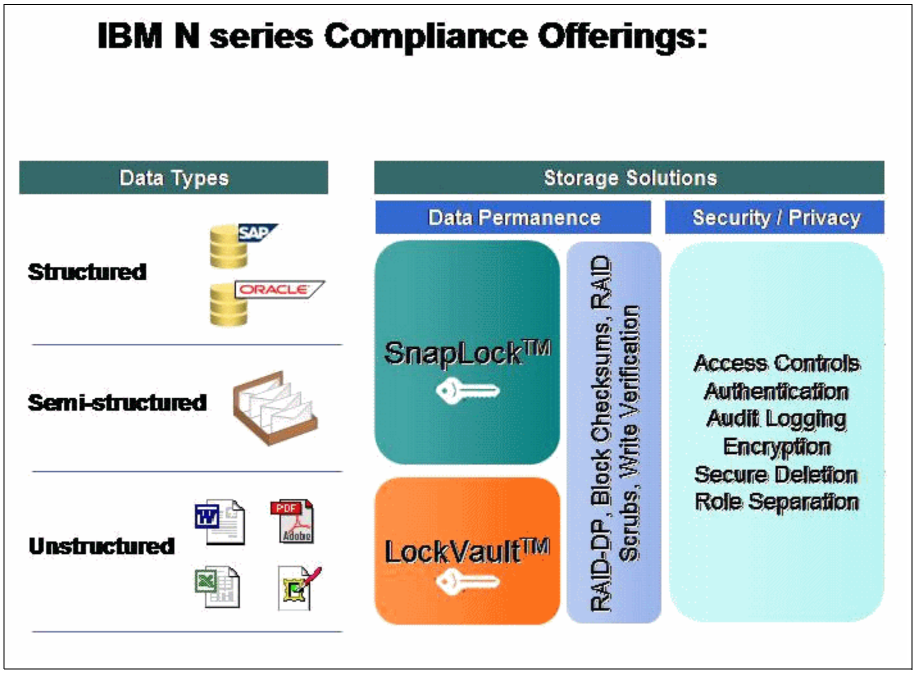

Figure 9-4 Compliance offerings

9.3 Summary

LockVault is designed for retaining large amounts of unstructured data:

- Documents
- Project files
- Home directories

LockVault is built upon SnapLock and SnapVault products. With LockVault, retention periods are set on SnapShot automatically after a SnapVault transfer takes place. LockVault integrates with OSSV as well, creating a compliance solution for open systems without compliant storage.

10

SnapRestore

SnapRestore software is used for data protection. Various data protection mechanisms are discussed in the previous chapters, this chapter discuses the importance of having an effective way to restore your data. SnapRestore technology helps to quickly recover the data, ranging in size from a single file, to a multiterabyte volume.

10.1 Overview of SnapRestore

Achieving efficient data recovery is absolutely necessary for enterprises. Disaster strikes at any time. With the traditional data recovery technologies, this cost hours in productivity, and cost. The N Series, with SnapRestore technology, recovers data fast and easy.

Using SnapRestore, reverting a file, or volume is much faster than restoring files from tape, or copying files from a SnapShot copy to the active file system.

> **Note:** SnapRestore requires a license code. You *must* license SnapRestore before using it.

SnapRestore leverages the SnapShot feature of Data ONTAP software by restoring a file, the entire file system, or LUN to an earlier state. It is used to recover a damaged or deleted file, or to recover from a corrupted database, application, or damaged file system. Example 10-1 on page 60 shows the SnapRestore syntax.

Example 10-1 Syntax of SnapRestore

```
snap restore -s <Snapshot name> <flexvol name>
```

The system administrator can restore a file, or the entire file system, LUN, or entire volume, from any existing SnapShot copy. Without rebooting, the restored file, volume, file system, or LUN is available for full production use, having returned to the precise state that existed when the selected SnapShot copy was created. From a single home directory, to a huge production database, SnapRestore does the job in seconds, regardless of the size of the file, or volume.

10.2 SnapRestore operation

After you select a SnapShot for reversion, the filer restores the volume, or file to contain the same data, and timestamps as it did when the SnapShot was taken. As mentioned, all data that existed before the reversion is overwritten.

> **Important:** You cannot undo a SnapRestore reversion to change the volume back to the state it was in prior to the reversion.

10.2.1 What SnapRestore reverts

SnapRestore has the ability to revert to various stages of data life cycle. It only reverts file contents. It does *not* revert attributes of a volume, such as the SnapShot schedule, volume option settings, RAID group size, and maximum number of files per volume.

Using SnapRestore, you can restore the entire volume, or LUN, or you can perform a single file SnapRestore on a LUN. Volume level SnapRestore can be undesirable since all the data in the filer volume is reverted to the SnapShot state.

However, option settings applicable to the entire filer can be reverted. This is because the option settings are stored in a registry in the /etc directory on the root volume. If you revert the root volume, the registry is reverted to the version in use at the SnapShot creation time.

You can revert a volume to a SnapShot taken when the filer was running a different Data ONTAP version, as well. However, this can cause problems, because of version incompatibilities.

Important: You cannot revert a volume to recover a deleted SnapShot.

If you delete the hourly two SnapShot, and revert the volume to the hourly one SnapShot, you cannot find the hourly two SnapShot after the reversion. Although the hourly two SnapShot existed at the creation time of the hourly one SnapShot, SnapRestore cannot revert the contents of the hourly two SnapShot because you already deleted it.

Note: After you revert a volume to a specific SnapShot, you lose SnapShots that are more recent than the SnapShot used for the volume reversion.

As another example, after you revert the volume to the hourly zero SnapShot, you no longer have access to more recent SnapShots, such as the hourly one SnapShot. This is because at the creation time of the hourly one SnapShot, the hourly zero SnapShot does not exist.

10.2.2 Applying SnapRestore

SnapRestore is a data recovery facility available with the N Series. SnapRestore can be applied in the following scenarios:

- Disaster recovery
- Database corruption recovery
- Application testing, such as a development environment using large data files

If a client application corrupts data files in a volume, you can revert the volume to a SnapShot taken before the data corruption.

Prerequisites for applying SnapRestore

- SnapRestore must be licensed on the system. See Example 10-2 on page 61.

Example 10-2 Licensing SnapRestore

```
itso>license add ABCDEFG
```

Note: *ABCDEFG* is the license code purchased

- There must be at least one SnapShot copy on the system that you select to revert
- The volume to be reverted must be online and must not be a mirror used for data replication
- The LUN must be unmounted before using SnapRestore to revert the volume containing the LUN, or single file SnapRestore of the LUN. For a single file SnapRestore, the LUN must also be offline.

Examples of applying SnapRestore

- A messaging application or database application stores user data in one or two files that grow to several hundreds of GB in a volume. If this application corrupts the files, you can revert the volume to a SnapShot taken before the data corruption. Or, if a single file is corrupt, you can revert only the specific file, as illustrated in Example 10-3 on page 62.

  Example 10-3 Restoring a database file

  ```
  itsotuc> snap restore -s dbdata_snap.1 dbdata
  ```

- You can revert a volume or file in a test environment to its original state after each test. Example 10-4 illustrates the use of SnapRestore to restore a volume. The command used here is

  ```
  snap restore -t vol -s <snapshot name> <volume name>
  ```

  Example 10-4 Reverting a file in a test environment

  ```
  itsotuc> snap restore -t vol -s nightly.0 /vol/vol1

  itsotuc> WARNING! This will restore a volume from a snapshot into the active
  filesystem.  If the volume already exists in the active filesystem, it is
  overwritten with the contents from the snapshot.
  Are you sure you want to do this? y
  You have selected file /vol/vol1, snapshot nightly.0
  Proceed with restore? y
  ```

10.3 Summary

SnapRestore software reverts a file or an entire file system to an earlier stored SnapShot copy within seconds, with minimal downtime. SnapRestore is used to recover from a corrupted database, application, or damaged file system. The advantages are summarized below:

- It maximizes availability by reverting an entire volume in less than one second, regardless of size.
- SnapRestore greatly reduces the dependency on tape making data recovery a faster process.
- Also useful in a testing environment that requires frequent returns to baseline state.
- Essential for database or messaging environments (e.g., DB2, Oracle®, Lotus®, Exchange) prone to application errors, virus attacks, etc.

11

Summary of recommendations

Recommendations for data protection using the N Series are discussed in this chapter. This IBM redbooks publication has introduced various data protection strategies offered by the N Series. However,data protection is identified as a process that involves identifying different data types, choosing an appropriate data protection strategy, and implementing the strategy effectively. This chapter summarizes the salient steps involved in data protection.

11.1 Choosing a backup software technology

Organizations protect data so that they can restore it later. Data recovery falls into three categories:

1. Recovery of accidentally deleted files
2. Long-term, single, or multiple file recovery from archived data
3. Recovery of a file system after a disaster

As illustrated in Figure 11-1 on page 64, the N Series addresses increasing data protection needs for application recovery, synchronous replication for DR, or continuous replication for business continuance solutions. IBM provides a range of solutions based on customer requirements.

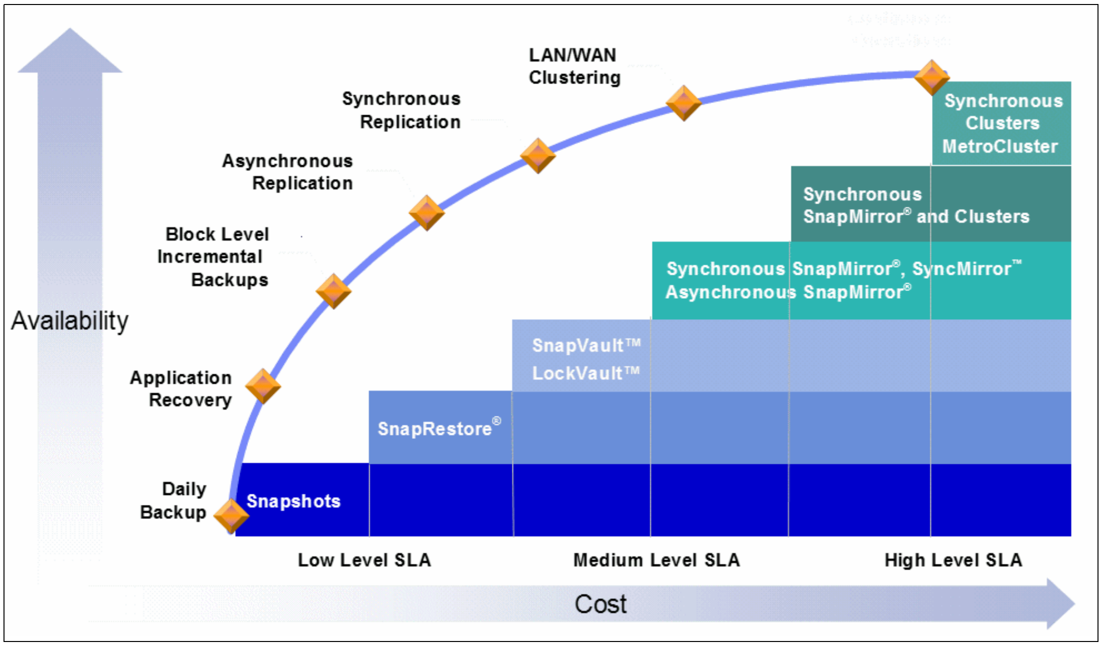

Figure 11-1 Data protection scales to meet customer needs

SnapShot technology for online recovery from user errors

Figure 11-1 on page 64 shows that SnapShots form the starting point for data protection. The integrate feature of data ONTAP microkernel SnapShot technology, provides fast and easy recovery of accidentally deleted files. By scheduling SnapShots throughout the day, the system administrator is guaranteed that recent files are available for recovery. Users can easily copy data from a SnapShot directory to their own directories. Because SnapShot technology is designed within the system, and the overhead is minimal.

FlexVol

FlexVol technology delivers true storage virtualization solutions that lower overhead and capital expenses, reduce disruption and risk, and provide the flexibility to adapt quickly and easily to the dynamic needs of the enterprise. *FlexVol* technology pools storage resources automatically, and enables you to create multiple flexible volumes on a large pool of disks.

FlexClone

FlexClone technology enables true cloniing, instant replication of data volumes and data sets, without requiring additional storage space. Each cloned volume is a transparent, virtual copy for essential enterprise operations, such as:

- Testing and bug fixing
- Platform and upgrade checks
- Multiple simulations against large data sets
- Remote office testing and staging
- Market-specific product variations

SnapVault for disk-to-disk backup and reduced impact on operations

SnapVault software provides a fast and efficient, disk-to-disk backup and recovery for IBM storage systems, and open system platforms. Block-level incremental changes are stored on the destination, but are usable as full backups. Space utilization is minimal, and system restore is fast and easy to perform. NDMP-based tape backup is used to protect data on the destination, offloading all overhead from tape-based data protection from primary storage systems.

SnapMirror for disaster recovery

For disaster recovery, we recommend you use *SnapMirror* to mirror data to remote locations. If you have serious file system damage, the system administrator turns the mirrored file system into the active file system. Even if a disaster occurs at one site, the data remains safe and online at the remote location.

MetroCluster

MetroCluster extends failover capability from the data center to sites at remote locations. It also replicates the data from the primary site to the remote site. Replication ensures that your data is current. The combination of failover, and data replication ensures recovery from disaster—with no loss of data—in minutes, rather than hours, or days. The built-in simplicity of MetroCluster allows for near continuous storage service operations during disasters with little to no administrator intervention.

SnapLock

SnapLock enables compliance with regulatory, and best-practices records-retention requirements by allowing the creation of nonrewritable, nonerasable *WORM* volumes on IBM near-line, and N Series storage systems, thereby preventing critical files from being altered or deleted until a specified retention date.

LockVault

LockVault allows customers to make permanent *WORM* backups of their unstructured data by copying, and storing only the unique blocks written since the most recent incremental backup. Every block variation is protected, so no data falls through the cracks. LockVault also automatically provides online compliant backups of unstructured data at multiple points in time.

11.2 Configuring near-line storage as a backup storage system

To gain full advantage of IBM data protection technologies and methodologies, you can configure a near-line storage device as a backup storage system. Figure 11-2 on page 66 shows IBM N Series System Storage, *Thomas and Abe*, and open systems servers, *Franklin and Jack*, replicate their data to four discrete volumes on the near-line storage device. *George*, an NDMP-compliant backup solution runs local simultaneous tape backups from near-line storage, to a high-capacity, multidrive tape library. The steps are as follows:

1. Run a SnapVault baseline transfer for each of the four systems
2. Run incremental updates to the volumes according to the schedule that fits your needs
3. Backup the four volumes on *George*, in parallel, to multiple tape drives in the tape library

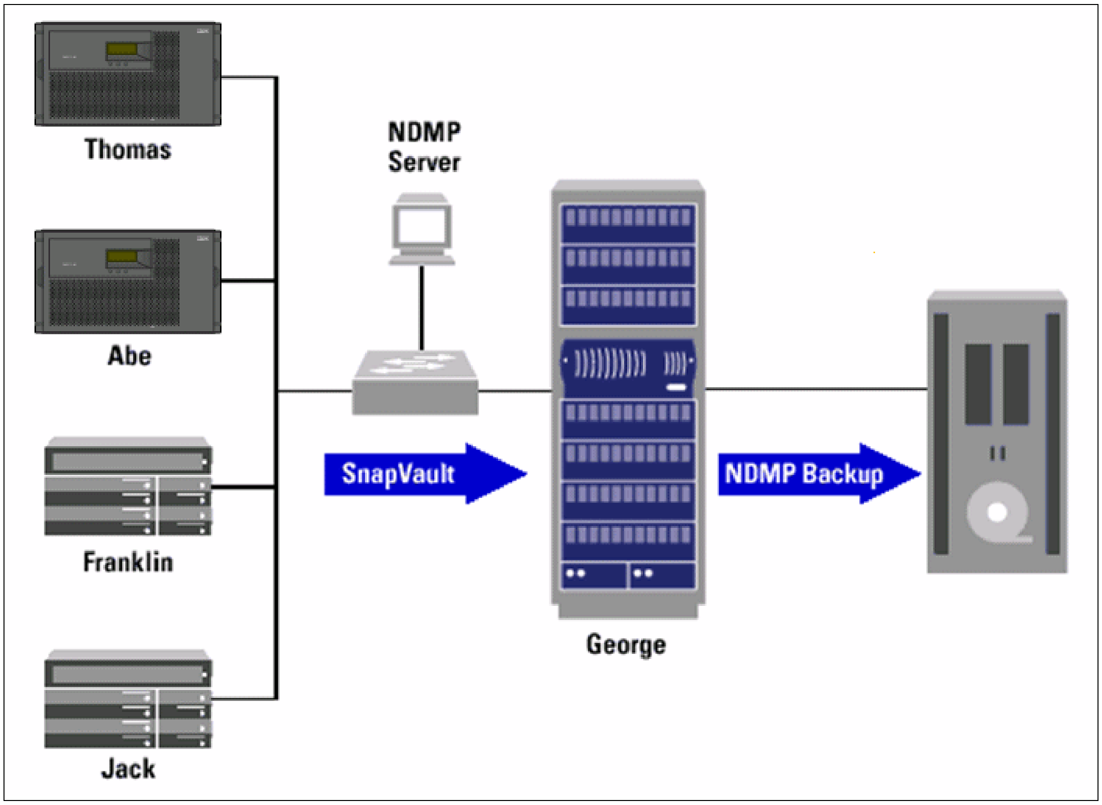

Figure 11-2 Configuring a backup storage system

Steps 2 and 3 can overlap. For example, you can run incremental SnapVault transfers on *Thomas and Abe*, while backing up *Franklin and Jack* volumes to tape, creating a pipeline schedule to fit your environment.

This configuration offers the following benefits:

- Allows for tape backups with no backup window
- Reduces the load on application systems, because they are not running local tape backup sessions
- Meets archiving needs
- Allows load balancing with read-only access of data on near-line storage
- Fits into existing NDMP-based, third-party infrastructures

11.3 Organizing file system data into volumes and qtrees

The N Series supports multiple volumes on the same storage system, and multiple qtrees within each volume. Backup performance is optimized by setting up volumes, and qtrees correctly. Be sure to use multiple volumes to ensure that the volume size meets your recovery time estimates. Within volumes, configure qtrees. The reasons follow:

- A full dump works on an entire volume.
  - A full dump of a very large volume takes longer than a full dump of a smaller volume. You can minimize backup windows by dividing your system into smaller volumes and doing full dumps of different volumes on different nights.
- Large volumes take longer to restore.
  - Match volume size to an acceptable time for a full disaster recovery from tape. For example, if your restore window for any volume is eight hours, and you estimate restore rates to be 40 GB per hour, use the following formula to calculate your largest volume:
    - 40 GB/hr. * 8 hours = 320GB volume

Note: If a 12-hour restore window is acceptable, the volume size can go up to 480 GB

Double-parity RAID (RAID-DP), is a unique technology that provides protection against data loss equivalent to RAID1 (mirroring) without the impact on usable capacity (see Figure 11-3 on page 67). RAID-DP groups are configured large enough to match RAID4 parity overhead (cost impact). Performance is fundamentally the same as RAID4. RAID-DP is available at no cost, or special hardware requirements.

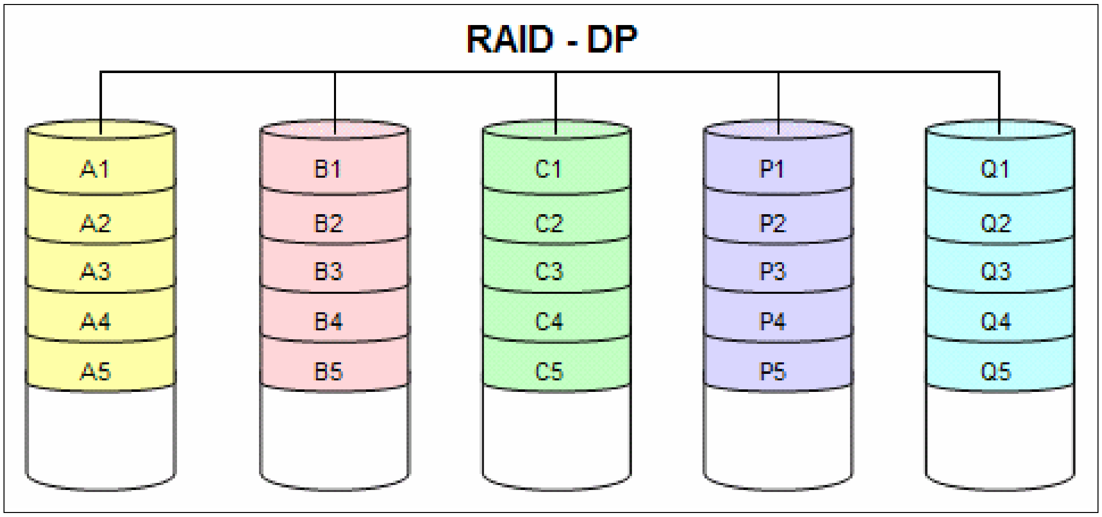

Figure 11-3 RAID double parity implementation

Traditional single-parity RAID technology offers protection from a single failed disk drive. The caveat is that no other disk fail, or that uncorrectable bit errors do not occur during a read operation while reconstruction of the failed disk is still in progress. If either secondary event occurs during reconstruction, then some, or all data contained in the RAID array, or volume is lost. With modern larger disk media, the likelihood of an uncorrectable bit error is fairly high, because disk capacities have increased, but bit error rates are the same. The traditional single-parity RAID to protect data is stretched past its limits.

11.4 Classifying data by value and change rate

We recommend that you classify data by its value, and by how dynamic it is. Isolating different types of data into different volumes, or qtrees provides flexibility in backup policies, for both tape backups, and SnapShot copies. Important, and dynamic data in one volume warrant frequent full dumps to tape. Place archival data in a separate volume, for occasional incremental backups, and full dumps at widely spaced intervals.

Multiple volumes, or qtrees also allow you to configure different SnapShot schedules for different kinds of data. For example, you might have home directories, and a source code repository on the same IBM storage system. Enabling frequent SnapShot copies for home directories so that users can always restore their own files saves time, and system administrator resources. Disable SnapShot copies for the source code repository, SnapShot copies are not practical, because it changes so frequently.

11.5 Summary

Risks of data loss include deleted files, system crashes, application crashes, viruses, and natural disasters. Because of these risks, a carefully thought out data protection plan is imperative. Putting a data protection plan in place involves:

- Identifying business-critical data
- Defining requirements:
  - Protection from accidentally deleted files
  - Archiving data for future use
  - Reducing backup and restore windows
  - Recovering from a disaster
- Analyzing your data:
  - How dynamic is it
  - Size of data set
  - Number and size of files
  - Directory structure
  - Data type
- Choosing data protection software and hardware solutions that address your requirements
- Configuring these technologies to ensure that performance meets your needs:

IBM System Storage N Series provides a unique set of solutions:

- Built-in SnapShot technology:
  - Enable online SnapShot copies throughout the day
  - Provides online file recovery
  - Provides a potential replacement for time-consuming incremental tape backups
- SnapRestore software:
  - Provides near-instantaneous recovery of individual files, or entire volumes
  - Minimizes downtime to recover from virus infections, or database corruption
- NDMP gives IBM users a variety of choices among high-performance tape backup and restore vendors
- SnapMirror technology:
  - Enables organizations to replicate data volumes at high speeds over a network
  - Provides an up-to-date mirrored volume
  - Ensures uninterrupted operation in case of disaster
- SnapVault software:
  - Enables efficient disk-to-disk backup and recovery of both IBM storage systems, and open systems storage

The best data protection strategies often combine these solutions, each of which solves a different piece of the backup problem.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 72. Note that some of the documents referenced here may be available in softcopy only.

- *IBM System Storage N Series MetroCluster*, REDP-4259
- *Using IBM System Storage N Series SnapDrive for Lotus Domino for Windows*,REDP-4287
- *IBM System Storage N Series MetroCluster Planning Guide*, REDP-4243
- *N Series SnapManager with Microsoft SQL*, REDP-4174

Other publications

These publications are also relevant as further information sources:

- *IBM System Storage N Series Data ONTAP 7.2 System Administration Guide,* GC26-7974
- *IBM System Storage N Series Data ONTAP 7.2 File Access and Protocols Management Guide, GC26-7965*
- *IBM System Storage N Series Data ONTAP 7.2 Data Protection Online Backup and Recovery Guide,* GC26-7967
- *IBM System Storage N Series Data ONTAP 7.2 Data Protection Tape Backup and Recovery Guid,* GC26-7968

Online resources

These Web sites are also relevant as further information sources:

- Support for Data ONTAP

  https://www-304.ibm.com/systems/support/myview/supportsite.wss/supportresources?taskind=7&brandind=5000029&familyind=5329797&typeind=0&modelind=0&osind=0
- Support for Network attached storage (NAS) & iSCSI

  https://www-304.ibm.com/systems/support/supportsite.wss/mainselect?brandind=5000029&familyind=0&continue.x=20&continue.y=4&oldbrand=5000029&oldfamily=0&oldtype=0&taskind=1&psid=bm
- Support for Open Systems SnapVault

  https://www-304.ibm.com/systems/support/supportsite.wss/supportresources?brandind=5000029&familyind=5329835&taskind=1

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads:

ibm.com/support

IBM Global Services:

ibm.com/services

Index

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5” spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Redbooks

Data Protection Strategies in IBM System Storage N Series

(1.5” spine)
1.5”<-> 1.998”
789 <->1051 pages

Redbooks

Data Protection Strategies in IBM System Storage N Series

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

Redbooks

Data Protection Strategies in IBM System Storage N Series

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

IBM Redbooks

Data Protection Strategies in IBM System Storage N Series

(0.2”spine)
0.17”<->0.473”
90<->249 pages

(0.1”spine)
0.1”<->0.169”
53<->89 pages

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Data Protection Strategies in IBM System Storage N Series

(2.5” spine)
2.5”<->nnn.n”
1315<-> nnnn pages

Data Protection Strategies in IBM System Storage N Series

(2.0” spine)
2.0” <-> 2.498”
1052 <-> 1314 pages

Data Protection Strategies in IBM System Storage N Series

IBM data protection solutions in detail

Knowing your data

Business issues affecting data protection

Systems fail, users accidentally delete files, natural disasters occur, and mistakes happen. Businesses are losing critical data. One of the most important questions IT management must ask is, "What is my data recovery plan?" IBM® System Storage™ N Series provides a variety of choices for data protection and recovery. This IBM Redbook® publication addresses many available options, and recommends solutions for protecting data using the IBM System Storage N Series.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7591-00 ISBN 0738431036