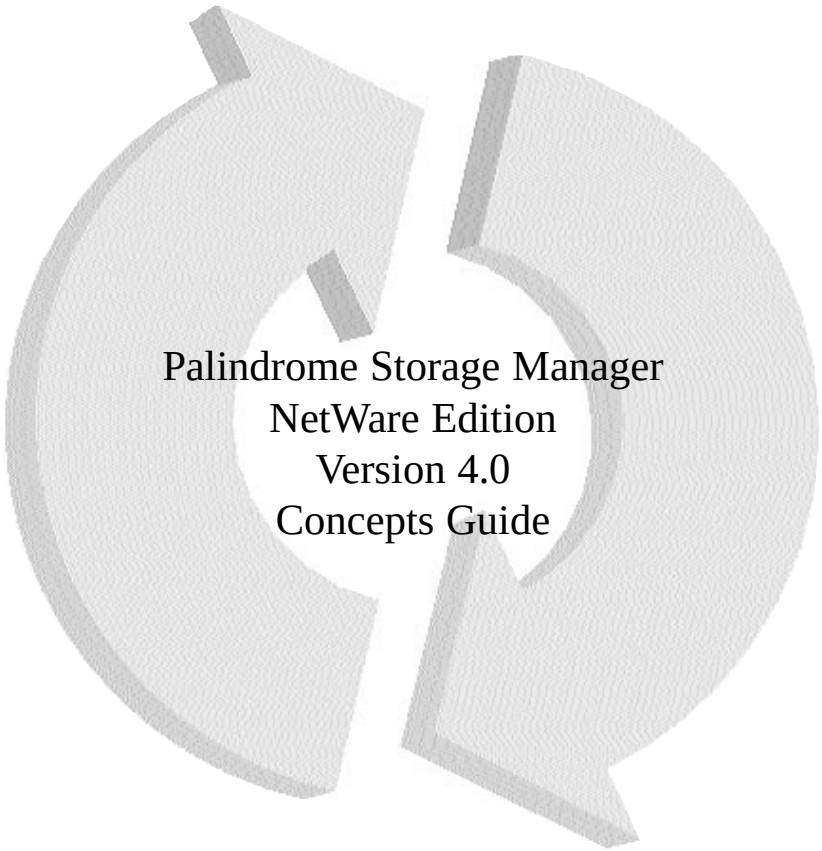

Palindrome Storage Manager
NetWare Edition
Version 4.0
Concepts Guide

Palindrome Corporation
Palindrome Storage Manager V.4.0 (NetWare Edition) Concepts Guide
Manual Number: 0-NWSM-CONCEPT-40b
October 1995

Palindrome Corporation
600 East Diehl Road
Naperville, Illinois 60563
(708) 505-3300

Seagate Software, Ltd.
Hayley House
London Road
Bracknell, Berkshire RG12 2US
England
+44 344 360888

Seagate Software, GmbH
Hanns-Martin-Schleyer-Str. 34
47877 Willich
Germany
+49 2154916350

Table of Contents

Chapter 1

LAN Backup History

Overview

This chapter covers the LAN environment and how its changing hardware, software, and user interfaces technology has affected backup, restore, and secondary storage management software design and function.

LAN Backup Evolution

Like the LANs themselves, the products and strategies used to back up rapidly expanding amounts of data have gone through several distinct phases, and continue to evolve today.

The Pre-Tape Period

Because even a twenty megabyte LAN required about 60 diskettes and considerable hours of diskette shuffling time, many LANs received little or no backup protection.

Tape Backup

As larger LANs became the norm, affordable tape backup devices became available. The shared network disk could be copied to tape weekly or even nightly. Though the backups had no built-in intelligence, the convenience of tape backup gave the LAN manager a method to protect data on a timely basis.

Disk Image Backups

Early tape backup applications simply copied the contents of the disk to tape one cluster at a time—literally creating an image of the disk on tape. In most cases, the server had to shut down for the backup to be executed.

Single File Restores

Although the disk-image backup process did allow fast backups (and equally fast full volume restores), it did not provide a mechanism to recover individual files. Because recovering a volume was an “all or none” process, recovering a single file required several time consuming and dangerous steps:

1. The server would be taken down, and a full backup would be run to preserve the current contents of the server.

2. The desired tape would be mounted, and the old contents of the volume would be restored.

3. The server would then be temporarily booted, and the desired file (or files) would be copied to a workstation.

4. The server would again be taken down, and the backup session from the first step would be restored, bringing the server back to its original state.

5. And finally, the server would be booted again, and the recovered files would be copied from the workstation to their proper location on the server.

First Generation LAN Backup Systems

Benefits:	**Limitations:**
Simple software	Required skilled Network Administrators
High-speed backup	Single-file restores were virtually impossible—required shutting down the network and performing two full restores.

Although disk-image backups were fast and conceptually simple, severe difficulties in recovering files led to the first revolution in LAN backup software.

The First Revolution: Selective, File-by-File Backups

With the increasing storage capacities of LANs, there was a need for a completely different approach to the backup process. Rather than

backing up a volume as a continuous set of disk clusters, backup systems were created that would work with the operating system to allow file-by-file backups. While adding much greater flexibility, this approach added considerable complexity in programming the backup software. File naming conventions, network attributes, network security, and rights issues now came into play.

As LANs grew, it became common for the server's disk capacity to increase beyond the capacity of the tape backup device. This required network manager intervention to complete the backup operation, often late at night and with great inconvenience.

To avoid spending late nights changing backup tapes, a system of "incremental" backups was introduced. Under this strategy, a full backup would be done periodically, usually once a week. On the remaining nights, an incremental backup would be run, backing up only those file that had changed since the previous backup. Using incremental backups reduced the tape-swapping marathon to a once a week ordeal, but did not eliminate the problem. Larger, faster, and more expensive tape drives were the only answer.

Now there was a "trade-off". The simplicity of backing everything up every time was traded for a more flexible but complex incremental backup using less media. However, the restore process was much easier, retrieving a single file was now simply a matter of loading tapes and scanning their directories until the right version was found, and then using the tape software to restore that version. Single file restores were much less difficult, and could be performed without taking the server down.

As networks continued to grow, however, this approach started to show its limits. Although the daily incremental backups could run unattended, servers (or groups of servers) continued to expand beyond the capacities of the tape drives, requiring time consuming, attended weekly backups and adding greatly to the number of media in a tape library. Though selecting a specific file to restore was possible, large media libraries often meant considerable waits for the file to restore. To address this problem, some vendors added disk-based catalogs to track the contents of the tape library. These catalogs required significant

amounts of primary server disk storage storage for their indexes, but made it less time consuming to locate files on tape.

Second Generation LAN Backup Systems

Benefits:	Limitations:
Single file restores much faster than first generation system	Regular full backup required
	Backup integrity highly susceptible to user error
More flexible in selective file backups	Effective implementation required including elaborate planning by the network administrator

Just as a lack of flexibility led to the first revolution in LAN backup software, the limitations of inefficiency, reliability and manageability led to the second revolution.

The Second Revolution: Intelligent Storage Management

The Need for a Change

While the second generation systems provided excellent file restore flexibility, they required a great deal of attention from the network manager. Designing tape rotation schedules, writing scripts, creating backup templates, and training non-technical users to implement the backup strategy safely all combined to burden the network administrator with additional responsibilities.

The Purpose of Backup Systems

The goal of any backup system should be to provide the best protection possible for all network data, as quickly and reliably as possible, with a minimum of media, hardware, and human intervention required. The result of this strategy was the creation of the next generation of backup and archiving systems, intelligent storage management.

Characteristics of an Intelligent Storage Management System

The third generation backup product has various identifying characteristics:

- A detailed, relational database of file history and media tracking
- A rules-driven or goal-oriented user interface
- An intelligent engine to examine file histories and implement network administrator defined rules
- The ability to write to multiple devices, including optical jukeboxes, tape autoloaders, and run concurrent operations on these devices

An intelligent storage management system tracks files by backup date, thus allowing restoration of a specifically dated version. A relational database allows directory paths, file names, and file versions to be stored just once, with each instance of a file on backup media pointing to the commonly shared directory path.

To remain accurate, the appropriate entries must be removed from the databases as backup media are re-used. Intelligent storage management recognizes which file versions are obsolete and erases only those, leaving the permanent file sessions on media intact.

Intelligent management provides a stable, rich variety of file versions to compensate for faulty physical media and user file application or deletion errors. The default settings should guide the network administrator with a well-designed media rotation scheme, but with room for customization.

An intelligent system is designed to make run-time decisions—interpreting the goals defined by the network administrator, in the context of the established file histories, into specific backup operations. This allows the administrator to define precise protection for files or groups of files on the network, without having to constantly revise backup scripts, templates, and schedules, and without accumulating a hodge-podge of special-purpose media.

An intelligent system must use the right media at the right time. This includes never erasing permanent data, asking for new, blank media when appropriate, and above all, compensating for the inevitable user errors.

Since the vast majority of network files are stable—they haven't been modified for months or years—and have been long ago been permanently archived on several media, and there is no need to back them up again unless they change.

Benefits of Intelligent Storage Management Over Second Generation Systems

- Efficient use of resources
- Increased reliability of data protection
- Increased network administrator productivity
  - Allows routine tasks to be safely decentralized
  - Moves focus to higher-level storage management issues

Chapter 2

File Management

Overview

At the heart of intelligent storage management is the concept of file management. This chapter provides an in-depth review of file types and their “life cycles”.

File Life Cycles

For some LAN files, the life cycle is simple and uneventful. A memo may be created, remain on the network for a short period, and then be deleted. For other files, the life cycle is more complex. An application may be installed, used frequently, upgraded several times, and eventually replaced by a similar application from another vendor.

Stable and evolving files

At any point during its life cycle, a file may be categorized as either stable or evolving. An **evolving** file is one that is newly created, or has been recently modified. A **stable** file is one that has existed for a defined period of time without modification.

Evolving files need to be backed up when they change—an "incremental backup" will contain nothing but recently modified files. Evolving files can expect one of three fates: they may continue to be modified; they may stop being modified and become stable; or they may be deleted before they stabilize.

Stable versions of files are candidates for Archiving. In upcoming months or years, the versions of files that are most likely to be needed are the versions that conclude a project or work effort—the final draft of a document, the year-end sales spreadsheet, the production-level circuit design.

At different points in their life cycles, a single file may be created (**evolving**), be modified (**evolving**), and then is no longer modified (becomes **stable**). A spreadsheet that is used for forecasting might be modified frequently while it is being developed, become stable at some point, and then later be changed again (become an evolving file) as company procedures change. Thus, during the life cycle of a single file, various stable versions could exist and a copy of each could be archived to storage media for reference or retrieval at a later date.

Active and Dormant files

In addition to being thought of as evolving or stable, files may also be classified as active or dormant. An **active** file is simply one that has been accessed (created, modified, read, or executed) recently. A **dormant** file is one that has not been accessed over a given period of time (usually measured in months).

All evolving files are considered active, since they were recently modified. Stable files may be either active or dormant. Dormant files are the best candidates for migration (removal from primary storage) when disk space runs low, especially if they've been permanently archived to storage media, thus guaranteeing their restore in the future.

	ACTIVE STATUS	DORMANT STATUS
EVOLVING FILES	Any file that has recently been created or modified has been accessed. Example: Recently created word processing documents	**By definition, a file cannot be both Evolving and Dormant.**
STABLE FILES	Any file not modified but frequently accessed Examples: Most *.EXE and *.COM files.	Any file not modified, not read, and not executed recently. Examples: Old electronic mail and old word processing files.

Types of LAN Files

Almost all files on today's LANs can be categorized as one of five types. There are unique requirements for storage and protection of each file type.

Type 1—Created and never changed.

Examples: Memos, letters, e-mail attachments, application files

Time Line of a "Type 1" File

Overview: Most LANs are comprised of this file type. Type 1 files, such as "*.EXE" and "*.COM" application files, are added to the LAN as stable files shipped from vendors or copied from other locations and may have been created months prior to installation on the LAN. Other Type 1 files, such as e-mail and memos, are created over time but are intended for one-time use and are never modified.

Protection requirements: All files should be protected in the first backup session after they are created. They should be copied to multiple backup media as soon as possible to protect the files from loss due to media failure. Since there is no certainty that this file will remain

unchanged in the near future, it should not yet be added to the permanent archives. It should, however, be copied to multiple media as they are used, to protect against data loss due to media failure.

When files reach the period of stability as defined by the network administrator, the files should be permanently archived. At least one archived copy should be kept off-site.

Storage: All stable active files should be kept on disk for quick access. Stable *dormant* files may be kept on disk, but are most often migrated off to secondary storage devices. Selective “rules” may be set on specific files or file groups to keep them from being removed from disk.

Type 2—Created and periodically changed

Examples: Annual, quarterly or monthly updated output files from databases or accounting systems, etc.

Time Line of a "Type 2" File

Overview: These files are normally inactive until updated in synchronization with normal company business cycles.

Protection Requirements: Prior to reaching their final form as an annual, quarterly, monthly, or weekly report, these files should be copied to multiple media for protection.

When the current version reaches stability, it should be archived to multiple media for future recall (similar to the Type 1 files).

Storage: Type 2 files should remain on primary storage (disk) while they are still evolving and kept on primary storage if regularly accessed after becoming stable. (For example, a sales forecast spreadsheet may be read many times during a year, even if it isn't modified). Type 2 files should only be migrated to media when seldom used for reference.

Type 3—Modified frequently and then forgotten.

Examples: Lengthy documents, complex spreadsheets, software development files.

Time Line of a "Type 3" File

Overview: Type 3 files go through a development period, during which numerous revisions are made. This may involve continuous activity, with one or multiple revisions daily, or may have brief periods

of inactivity—such as a week-long pause while a draft is being reviewed. At some point, a “final” version is created, which will remain current until a replacement document is needed.

Protection requirements: As Type 3 files evolve, they should be protected on multiple media. After a period of time, the early versions of the file may be erased from storage and replaced with more current versions. The availability of a variety of recent file versions for reference is desirable. When the file version reaches stability, it should be added to permanent (non-erasable) archive sessions on multiple media.

Storage: Type 3 files, like all others, need to be on primary storage during their evolutionary phase. After they reach stability, they may be migrated to secondary storage for later retrieval.

Type 4—Constantly updated

Examples: Corporate databases (inventory, accounts receivable, customer lists, invoice records, payroll), e-mail message files, group calendars

Time Line of a "Type 4" File

Overview: Although fewer in number, these files are generally large, and are an integral part of day to day operations. They are characterized by frequent updates.

Protection requirements: Because of their importance, it is vital that the current version of each of these files be protected during each backup session. In addition, to limit losses in the event of data corruption on disk, several previous versions should be stored.

Depending upon the contents of the file, it may be appropriate to selectively add versions of the file to the permanent archive sessions (even if the file has not reached a "stable" condition). Some auditing requirements require keeping permanent copies of ***all*** file versions for future review.

Storage: Because Type 4 files are active by definition, they should **never** be migrated from disk to backup media.

Type 5—Created and deleted after a short time

Examples: Temporary files, text and spreadsheet files, etc.

Time Line of a "Type 5" File

Overview: Often, many Type 5 files become Type 1 files because there was not a regular procedure to review dormant files and delete obsolete ones from disk.

Protection requirements: Type 5 files normally are not needed for long-term restoration purposes, so should be backed up to erasable media configured for a fairly short retention period. If it is known that a specific Type 5 file will be needed for later restoration, it should be permanently archived to multiple media.

Storage: Type 5 files should be kept on primary storage for only a short duration, then deleted.

File Management Requirements

Backup

At the heart of any complete file management system is backup—maintaining extra copies of files (usually on a secondary media) to protect against loss due to disaster, damage, or deletion of the primary copies. A minimal backup system must maintain at least one copy of each file on a secondary storage media.

In addition, a thorough file management strategy must provide some depth for rapidly changing files—several previous versions should be retained (especially from recent days) to allow recovery should the current version be corrupted. No backup strategy is complete without off-site storage of some media—extra backup copies to be used in the event of an on-site disaster that destroys not only the LAN, but the primary media as well.

Archiving

File management must also include control of the archives—permanent copies of important versions of files. This requires periodically updating the archives, making redundant copies of the data, and managing on-site and off-site rotation of archived information.

At one extreme, using write-once media for backups could provide a permanent copy of every version of every file. In most environments, the cost of backup media (and the overhead of tracking its contents) keeps this approach from being practical.

The “art” of archiving is finding a way to preserve several redundant copies of important information, without archiving unimportant data, and without preserving unnecessary redundant copies of any data. This involves determining which version of files should be archived, what

backup media to which they should be archived, how many copies have already been made, and when the archiving should be done.

When done correctly, archiving provides and an extremely high probability that the most important data will always be accessible, despite deletion of the primary copies, media failure, or on-site disasters.

Vault rotation

File management includes careful coordination of the media kept onsite and readily available and media kept offsite, but still available for restoring if needed. Most LAN administrators consider the inconvenience of some file copies being offsite as a necessary "trade-off" between total efficiency and disaster recovery. Using its media copy utility, Storage Manager allows all original backup media to be onsite and a copy of all media to also be offsite.

Vault rotation scheduling is built into Storage Manager backup software and the number of copies on and off site is configurable. The media rotation schedules are designed to keep a rich mixture of files onsite and offsite by making multiple copies of files to different media.

File migration

Because disk storage is both expensive and finite, proper management of this resource is vital to the day-to-day operation of the LAN environment. Storage Manager provides capacity management by safely deleting obsolete or seldom used files automatically. When volume monitoring and automatic migration are turned on, eligible files are automatically deleted whenever volumes reach their configured high water mark (default=90% utilization).

To be eligible for automatic migration, files must:

- meet their migration rule (default=not accessed for 12 weeks)
- be archived on three different media sets (by default)

Also, if a near line set is configured, a copy of a file to be migrated must exist on the near line set.

When files are migrated, a phantom file (a zero-byte placeholder file) is automatically left on disk in place of the migrated file. If using Storage Manager's automatic recall feature, whenever end users try to access files that have been migrated, a restore job is automatically placed in the job queue to restore the file; thereby bypassing the need for end users to contact administrator's when they need a file restored.

For more information on disk capacity management and migration, see Chapter 9.

Chapter 3

Intelligent Storage Management

Overview

This chapter reviews key concepts and defines requirements for Intelligent Storage Management as it applies to LAN backups.

Introduction

Intelligent storage management uses the power of computer software to make storage decisions based upon individual file histories and administrator defined configuration parameters. It is critical for optimal file protection and disaster recovery.

Requirements for Intelligent Storage Management

Goal-oriented control

The most important implication of an intelligent storage management system is its effect on the network administrator's approach to storage management. An intelligent system should perform all routine decision making—which operation should be run next, which files need to be included, etc. This frees the network administrator from being bogged down with operational details (schedules, scripts, batch files, templates) and allows him or her to focus on the larger picture—providing the appropriate type of protection for all the files on the LAN.

This requires a user interface that allows the network administrator to define what level of protection is appropriate for a whole volume, for a group of files, or even for individual files. This capability must apply to all storage devices being protected—no matter how many servers or workstations are involved.

To satisfy the needs of LAN users, the network administrator must have backup and archiving software capable of providing considerable flexibility, both in scheduling and in applying specific criteria (configuration parameters) to files or perhaps entire directories and/or volumes.

Intelligent engine

A goal oriented or rules-based interface for the network administrator creates the need for an intelligent backup engine. Since the backup drive still works at the mechanical level of file sets and directories, the intelligent system must somehow implement the goals defined by the network administrator.

This requires the ability to make run-time decisions on which files should be included in each backup media operation based upon the contents of the media, the operation being performed, the age of the file version being considered, and that file's backup history and protection status.

Recovery

Intelligent backup is not an end in itself; its important feature is the ability to recover files when they are needed. The efficiency of data recovery depends upon the intelligence of the backup scheme and relies upon the database responsible for tracking multiple backup events.

Integration

Backup, **archiving**, and **migration** are interdependent functions. If the current version of a file is fully protected—that is, a sufficient number of permanent archive copies of it exist—there is no need to keep copying it to media. Storage Manager software defaults to a configuration that notes this redundancy and allows for speedier backups once the stable files are fully protected on multiple media.

The need for file migration very often occurs at a crisis moment when a network volume is nearly out of file space. Stopping operations to mount backup media for a manual migration of files from the volume would be unacceptable. An intelligent archiving system should, as part of its normal operations, monitor disk utilization and archive qualified

files to media in preparation for migration. Then migration becomes merely a function of deleting the file from disk. File histories made when the files were archived to media indicate where copies available for restoration are located.

Backup Media and System Management

Because of the mission-critical nature of backup and archiving, network administrators are directly involved in the day-to-day aspects of storage management. Intelligent storage software automates the backup operation so that the operations are controlled by network administrators, but may be initiated by less experienced personnel, thus allowing the administrator to attend to other systems tasks.

Automation considerations should include:

- **Media rotation**—scheduling media changes, and determining what types of backups occur.
- **Off-site rotation**—which media should be off-site at any point in time, to protect against on-site disaster.
- **Message notification**—determining who shall be notified if a backup or restore job fails.
- **Monitor media**—reviewing the reusable space available on the backup media and retiring nearly full media, monitoring high media soft error rates as a possible indication of either pending media or pending hardware failure.

Chapter 4

Product Architecture

Overview

This chapter identifies the components of Storage Manager, and describes their purpose and interactions.

Overall Structure

Storage Manager is composed of the following major components:

Physical Components:	SCSI interface card, SCSI cable, Backup device(s)
Database Components:	System control database, File history database(s)
Program Modules:	Job Server, Job Queue, Backup and Restore Engines, etc.
User Interface:	Graphical User Interface icon selections
Target Service Agents (TSA):	To protect servers, databases, and workstations

SMS Architecture

Storage Manager version for NetWare incorporates Novell's Storage heterogeneous workstation clients, NetWare server volumes (including multiple name spaces), and non-volume resources such as the bindery.

This architecture was developed to make it easier for Storage Manager applications to work with future versions of NetWare and with the file systems NetWare supports. To do this, Storage Manager uses unique Target Service Agents (TSAs) for each supported platform.

Target Service Agents, Targets, and Resources

TSAs (Target Service Agents) are installed on every server and client workstation to be protected by Storage Manager. Server volumes and non-volume resources require that an appropriate TSA be loaded for the server to be protected. Protection of client workstations requires that interrelated TSAs be loaded both at the workstation and on one server within the directory tree. Information from the client workstation is stored by the server TSA and manages communication between the server and workstation. Once loaded, the servers and workstations become the target services.

Volumes on servers and workstations are known as resources. Resources are designated for protection by adding them to the Protected Resource List. A resource may be a volume such as a file system or a non-volume resource such as the bindery, Network Directory Services (NDS), or an SQL database. Screens displaying volumes indicate a defined size, while non-volume resources display as zero-byte size.

Some agents are responsible for a single target service, such as the TSA312 agent responsible for NetWare 3.12. Other agents, such as

Data Interchange

SMS has the capability to transfer data between heterogeneous clients using SMS's System Independent Data Format (SIDF). Storage Manager writes data to backup media using Target Service Agent generated data in SIDF. Prior to Palindrome's SMS 3.1 and Storage Manager versions, data was written in a Palindrome proprietary format.

TSADOS.NLM and WSMAN.NLM , handle multiple targets (TSASMS.COM on both Windows and DOS clients). There may be multiple resources associated with a single target, such as the C: and D: drives of a PC, or the various volumes on a NetWare file server.

Storage Manager Databases

Each installation of Storage Manager is comprised of a unique **single** System Control Database, and **multiple** File History Databases, one File History Database for each volume that has been added to the Protected Resource List. The System Control Database is comprised of two files and each File History Database is comprised of eight files.

The System Control Database

Each System Control Database, comprised of two files (ASDB.PAC and ASNX.PAC), is used to store information for a specific Storage Manager installation. A LAN with sixteen servers being backed up by four Storage Manager installations would have four distinct System Control Databases (one associated with each installation). The System Control Database is stored in the volume and directory specified at the time of installation.

Configuration information stored in the System Control Database includes information regarding:

System Configuration

- The rotation schedule and pattern
- Default restore options
- Setup information such as vaulting, number of archive copies, etc.

Media Management

- The “Library ID” (the “name” of the backup system) forms the base name for all backup media labels within the library.

  For example, if the Library ID=“NEWLIB”, then the first managed media created would be NEWLIB:A:1 indicating the Library ID (NEWLIB):Media letter set (A): and media number (1) within the A: media set.

- The names of the backup media that have been created (see above example)
- The backup, archive, control database, and file history database session names and session sizes for all media
- The status of all media (active or retired)
- Total media capacity, capacity free, bytes of archived data, bytes of backup data, and other statistics per backup media

Media Scheduling

All media scheduling information is stored in the System Control Database.

- The date, time, and status of the last operation
- The date of the next scheduled media rotation

License Serial Numbers

The System Control Database tracks the serial numbers for the base software and any installed options.

Protected Resource List

The Protected Resource List is used to locate, access, and protect both volume and non-volume resources for servers and volumes for client workstations on each LAN installation.

File History Databases

Each volume on the Protected Resource List has a File History Database comprised of eight files. See *Appendix B* for details on each of the File History Database files.

By default, the File History Databases are centrally located (all histories are stored on one volume).

Using Resource Manager, the File History Database(s) can be moved to a different volume or distributed onto the volume each represents.

Advantages of Centralizing Databases

Centralizing File History Databases for all protected resources has advantages.

- Storage Manager can access each resource's file history information more quickly and update it with minimal risk of corruption due to other network traffic.

- Unless the missing resource is the installation volume, the file history database for any resource can be accessed even if the resource is unavailable.

  For example: Vol2 is unavailable for some reason, but data needs to be restored from the last backup operation that occurred for Vol2. The file history database for Vol2 (centrally located on the installation volume) can be accessed and the file tagged for a *redirected* restore to an *available* volume. If necessary, the whole File History Database for Vol2 can be redirected to an available volume and files tagged to restore at this new location.

Because the File History Databases track the details of each file version of all files still on media, the Storage Manager history database was designed to minimize the amount of stored data. The following is information that could be required for a simple text listing of **one** version of **one** file:

Path Name:	Up to 128 bytes
File Name:	11 bytes for DOS filenames, 31 bytes for Macintosh filenames
File Size:	4 bytes
Date Stamp:	4 bytes
Media Name:	11 bytes
Session Name:	8 bytes

A typical entry (not worst case), with a 20-character path name, would require almost 60 bytes. If there were five previous versions (plus the current one), and each had three copies stored on backup media, there would be 18 instances of this file occurring in the File History Database. At 60 bytes each, this file's entry alone could occupy 1080 bytes. Given an average file size of 10 to 20K, this would require allocating five to ten percent of the total disk space to track the contents of the media.

Since much of the path information is redundant, Storage Manager database design eliminates as much of this information as possible. By using relational database technology, histories can be stored in a small fraction of the space that would otherwise be required. In most cases, this is less than one percent of the volume's capacity.

When backup media are re-used, erasable backup file sessions (CPnnnn sessions) are removed from the media and also from associated file history databases.

The User Interface

Storage Manager uses a Graphical User Interface (GUI) to allow customization of configuration parameters, set rules for file backup occurrences and migration from disk, access volumes and file history databases, view data being protected, and provide for custom backup and restore requests for multiple installations from one location.

Customizing Using Rules

Storage Manager's flexibility allows system administrators to customize for individual LAN environments by customizing file "rules" settings. Default rules are set at the time of installation and provide excellent file protection. However, the system administrator may change these rules at the volume level, at directory levels, or even to the individual file level, if desired, to meet specific auditing and regulatory requirements or to conform to company policy.

Rules settings allow custom backup choices:

- Never backup a file (security reasons and temporary files)
- Always backup a file (required extra protection)
- Backup the file only when it changes (efficiency)
- Archive files after a period of dormancy (efficiency, permanency)
- Migrate the file to secondary storage after a selected period with no user access (efficient disk management)

Viewing Files in Native Format

Storage Manager displays files according to the tracking name space on the volume where the file resides.

This can be especially important at restore time. Consider the following example. A Macintosh user installs an application (called "Application") on the server. During the installation process, it creates a number of folders (subdirectories) under the Application subdirectory. These folders are named "Application Templates", "Application Program Files", "Application Examples", "Application Data", and so on.

At restore time, however, difficulties become apparent. Even if the backup system stored the full Macintosh name for the folders (not always the case), the user is presented with a series of directories named "APPLICA1", "APPLICA2", "APPLICA3", "APPLICA4", etc.—the quasi-random truncated names NetWare used for the DOS name space. The user has no way of telling whether "APPLICA2\IMPORTAN", "APPLICA3\IMPORTAN", or "APPLICA4\IMPORTAN" is the one that had been "Application Data:Important Data File". Even within one folder, the truncated and numbered names that appear under DOS may make it almost impossible for the Macintosh user to determine which file to restore.

For any LAN with a significant number of Macintosh users, it is essential that the network administrator be able to view the full path and file names of all AFP (AppleTalk File Protocol) files on backup media. For this reason, pre-SMS versions of Storage Manager displayed the AFP name wherever it exists—all 31 characters, even preserving upper and lower case letters. For files that had no AFP name, the standard DOS name is used.

Resource and File Selection

Resource Selection

The user interface provides powerful tools for locating and selecting (tagging) individual files, directories, trees, volumes, or entire servers for user-designated backup or restore operations.

All protected resources can be accessed through Resource Manager. Options to collapse or expand the existing level, the directory tree, or all levels make viewing the resource efficient, especially if there are many protected resources for the installation.

Options allow easy movement from one volume to another and to directories within the volumes.

Directory Sorting

Within File Manager, Storage Manager displays files in a directory in the same order NetWare would display them—they are unsorted. This is the fastest option, and has a useful side effect. NetWare stores deleted files as a group. A single “End” keystroke positions to the end of the file list, where the deleted files, if any, are displayed.

There are four ways in which the directory listing of files may be displayed:

- Unsorted
- Sorted by extension, based upon the optional three-letter extensions to the file names. This sorting option is the most useful when defining disk management operations—for example, viewing all the “.BAT” batch files within a directory.
- Sorted by basename, starting with the first alpha-numeric character in the file’s name. This is useful when searching for a particular file (or group of similar files) by name.

- *Sorted by date* (last modified date), with the oldest files at the top of the list. Deleted files are inserted in the list based upon the date of the most recent version that exists in the media library.

Pattern Filters

Pattern filters may be used in conjunction with any of the sorting options. This is most useful in reviewing files of a similar type, such as *.EXE, *.COM, or *.TXT as might be created by a particular application.

Even though the currently displayed list may show only files left after a previous filter has been applied, all files are still available for display.

For example: If the “PAY*.*” pattern has been set on, only files meeting that criteria will be display. Resetting the pattern filter to “ACCT*.*” will allow ACCT*.* files to display.

NOTE: Turning off the pattern filter will again display all file patterns at that directory tree level.

Date Filters

An especially powerful tool in the user interface is the Date Filter option. After selection of this option, a starting and ending date (and time to the minute, if desired) range is entered. The directory display will reflect only files whose create/modify date fits within the entered range.

- If the date range is set to end one week ago, the files displayed may appear much differently then the “current” file versions as they were reflected on disk most recently. Any files that did not exist a week ago disappear from the display as they “haven’t been created” from this (week-prior) point of view.

  The date filter, used with pattern filters and directory tree selections, make it possible to quickly specify intricate file restores.

File Selection

Most custom operations involve one or more “selected” files.

The simplest case involves selecting a single file for an operation.. The user locates the file, marks it, then selects the operation to be performed.

A more complex example would be to select an an entire directory tree to be restored to a new location on disk. This would involve locating the directory tree, tagging it (and all included files), and selecting the restore and redirect option.

Job Server and Job Queue

Whenever a user submits a restore or utility operation, that job is placed into the Job Queue.

The Job Server monitors the Job Queue and services job requests by calling on the appropriate engine to perform the task.

In addition, the Job Server monitors a list of scheduled jobs (e.g. such as quarterly backups) and submits them when appropriate.

Jobs in the Job Queue are serviced in the order they were submitted to the queue.

Notification

If a job is not processed for some reason (e.g. no media is mounted), users and system administrators can be automatically notified via NetWare send messages, EMail, or SNMP.

Alert buttons indicate to the user that a situation requires attention. Storage Manager monitors each installation for system messages and conditions that require attention and activates the appropriate alert button. When active, the alert button appears in color; when inactive, the alert buttons are gray. The alert button bar indicates the following installation conditions.

Job Server Inactive—Indicates that the job server is not loaded.

Job(s) Require Attention—An automatic job requires a response to a system message.

Job(s) on Hold—A fatal error occurred while processing.

Check Last Automatic—The last automatic operation failed. For example, a fatal media or device error has occurred.

Check Next Media—The media expected for the next operation has not been mounted.

Check Device—The device is not available for scanning.

Automatic Disabled—The automatic job has been disabled through Job Scheduler and cannot run as scheduled.

Check Resource Monitor—either a migration job is on hold because a resource is in a high state or resource monitor is not accessible.

When selected, each alert button prompts the user with appropriate response options.

Restore Engine

Storage Manager's restore engine is activated when the Operation/Restore option is selected from within **Media**, **Resource**, and **File Manager**.

See Chapter 5 for additional detail as to how Storage Manager determines the location of files on media, the file version to restore, and situations that might cause scheduled file(s) not to restore after scheduling.

Disaster Recovery

Recovering a complete directory of files with associated rights and trustee assignments, restoring a complete volume, and restoring complete servers with possible Network Directory Services tree (NetWare 4.x servers) and security information is a complex process. In some cases where a full volume or server restoration is required, Storage Manager's databases and Novell NetWare installation files may have been lost along with all user application and data files.

System Recovery Procedures

Storage Manager has been designed for efficient restoring of single files, multiple files, complete server volumes, Storage Manager databases, and complete LAN installations. All restore operations make use of the restore engine.

See Chapter 5 for more system recovery detail and *Flow Chart of System Restore Process* on the following page to view decisions and actions taken during a system restore.

Backup and Archiving Engine

Storage Manager's backup engine handles all backup instructions launched as a result of job requests placed into the Job Queue. The backup engine is most commonly invoked to handle automatic backup operations.

On rotation days (when media sets are changed), some files are selected for *archiving* to permanent storage and other files are selected for *backup* to erasable storage.

Backup Procedure

The most common operation for Storage Manager is the automatic backup. When this operation occurs, the backup and archiving engine will examine the System Control Database to review the volumes to be protected and perform the appropriate backup on each item. See Chapter 8 for detailed information on this process.

Unattended Operation

Errors occurring during backups could be user created (forgetting to change backup media on rotation day), network related (file locked by another workstation), or media-related (media failure, media full, etc.). Errors that would normally expect operator intervention, Storage Manager handles with the "best choice" course of action, posts a message to System Messages (the message log), and continues if possible.

The following table shows a few possible situations that might occur during the backup process:

BACKUPS in ATTENDED and UNATTENDED MODES

SITUATION	Attended Mode	Unattended Mode
Unexpected Media Encountered	Prompt User for preferred media: Options include changing to preferred media, quitting, or continuing with existing media	Continue with any eligible media: Perform backup, force rotation, and adjust rotation schedule
End of all Available Eligible Media in Backup Device(s)	Prompt user with possible options to continue (introduce blank to add to media set, insert other eligible media,etc.)	End backup session, log information in Message Log for review and further options to continue backup operation
User presses <ESC>	Discontinue operation after confirming intention to quit	Prompt for password to interrupt operation

Complete Nightly Backups

The intelligence, flexibility, and concurrent backup capabilities of Storage Manager, combined with recent technological advancements in hardware design, allow complete nightly backups to become reality.

- **Intelligent Software:** Unattended mode's capability to react "forgivingly" to unexpected situations or lack of preferred media, coupled with the backup and archiving engines' intelligence and adaptability, dramatically increases the probability of a backup session successfully completing.
- **Flexible Software:** The ability to configure for specialized full, incremental, or differential backups means that the backups can be scheduled to avoid times of peak LAN utilization.
- **Backup Hardware:** Consideration should be given to the use of multiple backup devices, multiple-drive devices, tape autoloaders, or optical jukeboxes as backup devices. As both additional managed and blank media can be contained within these devices in anticipation of media filling, the probability of complete backups during each operation is greatly enhanced.

Concurrent Backup Operations

If more than one backup device has been configured, Storage Manager has the ability to backup more than one volume simultaneously. This ability provides efficiency, as backups take less time and file protection is more immediate. Additional memory is required to allow concurrent operations, but some speed degradation may occur when more than four processes are configured.

To enable concurrent backups, at least two devices must be configured. Devices can be single media storage such as 4mm DAT (digital audio tape) or 8mm drives or multiple media devices such as tape autoloaders (single-drive or multiple-drive units), and optical jukeboxes. They may be used in any combination.

All devices used for concurrent operations should be configured using Device Manager, and there should be no device restrictions for backup operations. All configured devices containing eligible media are then available to handle any backup sessions scheduled by the concurrent scheduling agent.

Interaction with Other System Components

Databases

During automatic operations, Storage Manager queries the System Control Database to determine which resources are included on the Protected Resource List for protection.

Each volume's File History Database is consulted to view the history of each file on the volume. For each file, the backup and archiving engine examines the status of the current version (how many copies exist on media), determines the effective protection rules for that file (which the network administrator can use to define protection criteria for various files), and decides if that file should be included in the current operation.

After each session (archive or backup) has completed, the volume's File History Database is updated to reflect the current status of files protected during that session.

Scheduling Agent

Storage Manager queries data stored in the System Control Database using a scheduling "agent". Information stored in the System Control Database includes the last backup media used, the next scheduled media change (and media expected), and the next operation to be performed.

Based upon this information, the current date and time, and the media that is currently in the drive, Storage Manager determines which actions should be taken. Should a scheduled operation be interrupted abnormally, the Storage Manager Console/Monitor window will display any system errors. The displayed error and any linked errors can be reviewed by accessing System Messages.

Disk Management

With the rapidly expanding addition of new applications to LAN systems, monitoring disk utilization is vital to keeping LANs functional. Due to the high cost of primary disk storage, it is important that only actively used files be stored there.

With Storage Manager's built-in resource monitoring and automatic migration features, inactive files can be automatically migrated from server disk to off-line storage freeing up valuable disk space for active files.

Using Storage Manager recall agents, whenever a user attempts to open a migrated file, the recall agent will automatically submit a restore request to the job queue.

Resource monitoring

Storage Manager provides resource monitoring with PALRMON.NLM which is loaded on the installation server during installation of the Storage Manager software.

PALRMON.NLM monitors the disk utilization of protected volumes and submits a migration job to the Palindrome job queue when volumes exceed their configured high water mark (default=90%).

Resource monitoring can be configured for the entire Protected Resource List or for individual volumes.

Migration

Migration frees disk space by deleting seldom used or less active files from primary disk storage. Only files that already exist on secondary

storage (such as tape or optical) are eligible to be automatically migrated.

Storage Manager's resource monitoring and file rules are the keys to automatic migration: only when a volume reaches its configured high water mark (default=90% utilization) does automatic migration occur and only files that are eligible for migration are migrated (see page 2-11 for more information on file eligibility).

With Storage Manager's phantom file option turned on, a zero-byte file is left on disk and a marker is placed into the File History Database for each file that has been migrated. Since this marker references an actual (permanently archived) file on media, it can be viewed in the file history window, selectively tagged, and restored.

NOTE: Zero-byte phantom files should remain in the directory in which they were created. If they are moved to a different directory, they will be backed up there as a zero-byte file, will not reference an actual data file, and will contain no data if subsequently tagged and restored.

Recall Agents

Storage Manager recall agents allow end users to submit restore requests for migrated files thereby eliminating the need for administrator intervention for end-user restore requests.

The recall agents intercept file open requests of migrated "phantom" files and submit the restore requests to the job queue.

Storage Manager provides recall agents for:

Windows Workstations	(PALWINRC.EXE)
DOS workstations	(PALRECAL.EXE/PALSMRCL.EXE)
NetWare servers	(PALRECAL.NLM)

You do not need to install both workstation and server recall agents. For example, if you have installed the server-based recall agent, you do not have to install the workstation agents.

The workstation agents (except PALSMRCL.EXE) provide a dialog box during file recall so the user can monitor progress of the recall. The server recall agent does not provide a dialog box.

Sample Recall Agent Configurations

Chapter 5

Restoring Resources and Files

Overview

This chapter describes Storage Manager's approach to restoring resources, security information, directories, and files.

Introduction

There are three Storage Manager options that allow selecting and restoring files from media. File Manager and Resource Manager use the File History Database(s) located on disk. Using Media Manager allows directly tagging file(s) within sessions *on media*.

File View

Use File Manager to show a combination of all files currently on server disk, backed up, but later deleted from disk (the copy must still exist on media). Because all files existing on media are readily visible (including all history versions of that filename), it is easy for the user to "tag" a particular file version to restore it to its original location.

Resource View

Use Resource Manager to select an installation, then display the resource directory tree to the level desired by collapsing or *expanding* the tree, and "tag" (mark selected) resources or entire servers to restore files.

Media View

Use Media Manager to physically journal the file history database sessions and files *on media*. This method is useful when the File History Database on disk may be corrupted or the media containing the file to be restored is not a "tracked" media (not referenced within the File History Database).

With this method, files located on media can be tagged to restore files outside of normal querying of the File History Database on disk and

allows viewing the System Independent Data Format (SIDF) formatted media that may have been imported from a different Storage Manager installation.

How Storage Manager Restores Files

Restoring a Single File

Single file restores are generally performed through File Manager. If a specific file version is not tagged for restore, the most recent version on media is restored. File(s) are normally restored to their original location.

Once a file restore job is submitted, Storage Manager queries the File History Database, to determine which backup media (and backup sessions on media) contains the tagged file versions.

If media containing the tagged file version is located in the backup device, Storage Manager will perform the following steps:

1. Mount the media (for tape autoloaders and optical jukebox devices)
2. "Open" the media
3. Seek to the location on media for the designated session (for devices that support high-speed search capability).
4. Read the media index to determine where the session is located.
5. Move through the media to the end of the session containing the needed file and read the session index. Each session index contains a list of files copied to that session. Sessions types are **CP**—Backup, **SV**—Archive, **DC**—System Control Database, and **DH**—File History Database.
6. Read the session index to determine the exact location of the tagged file, and position the media to that point (again using the drive's high-speed capability, if available).
7. Restore the file to disk.

Restoring Multiple Files

Multiple file restore operations are handled similarly to single file restores. File versions to be restored are located and tagged, and a restore operation is submitted to the job queue for execution.

Multiple file restores may require more than one media be made available for the operation to complete successfully. Providing all required media are mounted in a tape autoloader or an optical jukebox device, locating sessions that are on different media happens transparently (the device automatically loads each media as needed). If a single backup device is used, the operator is prompted for additional media as required to restore all files on the tagged file list.

Disk File Overwrite Options

The "Overwrite File Action" parameter set using Configuration Manager determines how files are handled during restore operations. The options are as follows:

- **Prompt**—on conflicts (the default)
- **Overwrite**—with newer file(s) only
- **Never—**overwrite file(s) on disk
- **Always—**overwrite file(s) on disk

System Restore

Depending upon the severity of the situation causing the need to restore files, a number of steps may need to be taken. If a server crash has caused the loss of NetWare operating information, server volume information, and other installed NLM modules, the following steps must be taken before Storage Manager can be used to restore files from media:

1. Novell NetWare must be re-generated
2. Resources must be re-created on the server
3. Storage Manager must be re-installed

Once the above has been accomplished, restoring the system to its prior status may involve restoring some or all of the following from media:

- System Control Database(s)
- File History Database(s)
- NetWare 4.x Directory Services or Bindery information
- Directory structure and trustee assignments (user rights) for one or more volumes.
- The most recent version of all files for a given directory or volume .

The complex nature of full server or volume restores requires intelligent application software, complete database information, and user-definable configuration settings—all built-in functions of Storage Manager.

System Restore Example

To illustrate the system restore process, this section will follow the logic the program uses to recover from a complete server loss—all

volumes are destroyed, and the server must be regenerated from the beginning.

In this example, assume a file server named FS1 had three volumes: SYS:, VOL1:, and DATA:. At the start of the restore, all three volumes are inaccessible and need to be re-created—meaning all data and security information are gone, and the Storage Manager databases on those volumes are unavailable.

Before the restoration can begin, the server must be regenerated—to the point where the volumes exist (without any data), and there is some place (SYS:) to restore the server's previous security system. This process is left completely to the network administrator.

Storage Manager programs must then be re-installed. The original protected volumes do not need to be added during re-installation. When the program recovers the System Control Databases, the protected volumes will be known.

Restoring Storage Manager's Databases

Storage Manager relies upon two different sets of databases (System Control Database and File History Database) to perform a restore. The first step the System Restore Agent takes is to locate and verify the integrity of all Storage Manager databases—the System Control Database and each volume's File History Database. This process will fail immediately, since no System Control Database exists on disk. The program will ask the user if the System Control Database should be recovered from backup media. Each installation has its own unique set of history files, therefore the media used to backup the installation must be used to recover it. The sequence in which the databases are replaced is important and is as follows:

The System Control Database files serve as an "umbrella" over all of the File History Database file "sessions" and therefore the System Control Database must be in place before any actions may take place that will use the File History Database files. During a restore, Storage Manager will immediately search for the System Control Database files, and if not found, will prompt the user to insert the most recently used backup media into the backup device (usually the media still

inside the device) to restore the System Control Database. The System Control Database is copied to media (in its own DCnnnn session) following each archive or backup session. When the System Control Database has been restored, it will have Protected Resource List information that existed at the time of the last backup.

Next the File History Databases need to be restored for each resource. If not available for a resource, Storage Manager will prompt the user to insert media to restore these files. It is important that the File History Databases be from the **same date as the System Control Database files or earlier**. Normally, all resources on the Protected Resource List are backed up nightly, and therefore the latest media would contain a File History Database version that will work correctly with the System Control Database in place.

If the latest File History Database for a volume is not available on the latest media, it will need to be restored from the next most recent media containing the File History Database files. The (DHnnnn) history sessions on media may be viewed for associated volume and backup date by using the Media Manager's Operations/*Journal* option. Once located, the history session is tagged and restored to disk.

NOTE: If the File History Database files to restore are from a software version prior to the installation version currently in place, they will be translated during the restore process.

Locating Files: Before Storage Manager can restore a file(s), it must locate the file(s) on backup media. The File History Database is first searched for the requested file(s) and in what session they are contained. To restore, System Control Database is queried to see what media need to be loaded to restore the files.

Once Storage Manager has restored as many files as possible from the first requested media, it prompts the user for the next media to insert into the restore device. With a tape autoloader or an optical jukebox, this is handled automatically, and prompting occurs only if the media needed are not located in the cartridge handling mechanism of the autoloader or jukebox device.

NOTE: Should all files not be restored (due to open files or a resource that is not available at the time of restore), the user will be prompted to optionally create a file that can be used to restore the remaining files at a later time.

Backup media may be inserted in any order desired, as only the latest version of each file is restored. In general, though, it is most efficient to begin with the last used media, as it contains the latest file version updates.

If the system administrator has advance notice that a volume will be down for a server disk replacement, a Full Backup operation should be performed to insure that all files needed for a restore are contained within one media set. *See Chapter 8 for more information.*

Restoring a Single Volume

When a volume is restored, Storage Manager restores files using the latest File History Database associated with that volume. The following sequence of events takes place:

1. Storage Manager checks the availability and integrity of the associated File History Database
2. If the files are corrupted or not accessible, Storage Manager attempts to restore the latest File History Database from media
3. The volume's directory tree is recreated, including empty directories, and directory user and group trustee rights are restored.
4. File data and file rights information are restored.

NOTE: Files that were deleted before the last media rotation *will not* be restored, but files deleted after the last rotation *will be* restored.

Restoring Multiple Volumes

When multiple volumes are scheduled for restore, such as with a full server restore, each volume is restored individually. *See above information on Restoring a Single Volume.*

Chapter 6

Goal-Oriented Control

Overview

This chapter describes how Storage Manager allows network administrators to use configurable file “rules” to customize the backup software for individual LAN environments.

Introduction

Prior to Intelligent Storage Management, network administrators were forced to manually manage all detailed backup operations. This often left the critical data protection function in the hands of inexperienced personnel. The Storage Manager strategy is to provide a reasonable set of default file rules that govern when, and how often, files are backed up.

During backups, rules are checked against the backup history of each file version as stored in the File History Database. To understand the decision making process that takes place prior to backup of a file, it is important to understand the difference between archiving a file version and making a *backup* of the file version.

Archive file sessions on media are permanent and are never erased from media during automatic operations. Archive sessions include files that are stable—for example, .COM files, .EXE files, quarterly reports, and annual reports.

Backup sessions are written to a reusable portion of media. The sessions contain files that are changing frequently.

Migrated files are stable for a configured period (default rule is 6 weeks) and are copied into archive sessions on media. When files are both stable and also not accessed for long periods of time (default rule is 12 weeks), they are normally part of a disk management process and are deleted. For automatic migration to take place, the file to be migrated must have multiple permanent copies in archive sessions on media and the volume must have reached a predefined utilization level. Once these criteria have been met, the file is deleted from disk and a marker put in its place for reference should a future restore be requested.

Default Rules

Default rules are created at the root of each volume resource when Storage Manager is installed. For NetWare resources, special default rules are set for hidden files, bindery and other non-volume resources, system files within the SYS: volume, and for print queue files. The default setting at the root of the volume is wild card pattern *.*

Default Rules Settings

Backup	Archive	Migrate
Include	After6w	After12w

Rules apply down the directory tree from the root to all subdirectories unless a more specific rule is set somewhere within the subdirectory tree. That more specific rule applies from the point it has been invoked to all lower levels.

Backup / Include

The “Backup/Include” rule instructs Storage Manager to write a backup of this file on media whenever the file changes. Each time media is brought into rotation, all backup sessions are erased from media and related references within the file history database(s) are deleted.

Archive After 6 Weeks

If a file remains stable (no modifications) for a period of six weeks, it is then available for permanent archiving to media. It will be copied to an

archive session on media and not erased when the media comes back into rotation in the future. Archive sessions on media are labeled SVnnnn, with the “nnnn” portion representing a unique numeric session identifier within the file history databases.

Migrate After 12 Weeks

If a file has remained stable for twelve weeks and has not been accessed (no modification, opening for reading, not executed, etc.), it will become available for migration. Though available for automatic migration, it must also:

- Have been fully protected (default is three archive copies on three separate media sets)

  AND

- The utilization level at which migration is configured for each volume must have been reached.

  For example, if the configuration is set for automatic migration when the volume is at 90% utilization, but the volume is only 80% utilized, files will not be migrated.

To check for files eligible for migration, use File Manager.

Scope of Rules

Each rule includes a wild card specification. This wild card may be as general as *.* (all files) or as specific as ONEFILE.TXT (a single file). A rule applies to all files matching the wildcard pattern from the directory tree location where it has been invoked to all lower subdirectory levels. For example, a rule set at the root of the SYS: volume would affect all files on SYS:, but a rule set at the SYS:\SYSTEM directory level would not apply to files residing on the SYS:\PUBLIC directory or its subdirectories.

Precedence of Rules

Because rules can be defined anywhere in the directory tree, it is common for one file to be governed by multiple rules. When this is true, the most specific rule at the same level applies.

The first test in determining rule precedence is where the rule is defined. If a file in \USER\GUEST is covered by two rules—one defined in \USER and one defined at the root—the rule from \USER will take precedence. The rule with the smaller scope (the more specific rule) overrides the more general one.

If two or more rules are defined at the same directory level, the rule with the most specific wild card pattern takes precedence. For example, if rules were defined for *.WK? and *.WK1, all WK1 files would be covered by the WK1 rule—since "1" is more specific than "?". As with the location within the tree, the rule with the smallest scope is considered the most specific, and takes precedence.

See above display with user-defined rules *.WK? and *.WK1. Note that DOCMGR.WK1 is covered by ***.WK1**, the more specific of the two rules.

Chapter 7

Archiving

Overview

This chapter reviews Storage Manager's strategy for archiving files to permanent media storage and achieving "full protection" to avoid unnecessary file backup redundancy.

Definition and Purpose

Archiving is the process of making permanent copies of stable files—files that have not been modified for a specific time. Archiving serves a number of purposes:

- Reduces the time it takes to complete a backup operation
- Reduces the number of media needed within the media library
- Allows disk grooming by migrating files after fully protecting them.

Files archived to media may still exist on disk. These files are available for restore on dates far into the future as they are never removed from media during normal media rotation operations. A file is considered fully protected when it has been archived three times (to three different media sets), and then is no longer backed up.

Storage Manager allows the system administrator to override the full protection criteria and copy fully protected files during any backup operation. This capability is normally only used when a special backup operation is to be performed (for example,prior to replacement of a disk), so that all files will be contained on the same media set for quicker volume restores.

Manual Archiving

Manually archiving has several limitations:

- It depends upon the network administrator or user to perform this function at appropriate intervals.
- It requires an extra full backup, adding to both the time spent on backup operations and to the number of media within the media library.

- It is possible that considerable redundancy of file copies on media will result
- A large number of short-lived temporary files are additionally copied to media that more logically should have only minimal copies and then be erased and replaced during media rotations
- Either the network administrator or a user must determine which media containing selective backups must be retired from rotation so that the files will not be accidentally erased.

Intelligent Archiving

Storage Manager has the built in intelligence to determine if files should be archived and to which media they should be archived without requiring Network Administrator intervention.

What Should be Archived?

Unless laws, regulations, or company policy dictate, only files that are seldom changed should be permanently archived to media.

To perform archiving efficiently, it is important to archive data that may be needed in the future, without wasting media on temporary files and short-lived data.

Storage Manager examines a file's age to determine whether it should be archived. Files that have not been modified in six weeks are considered stable, and worth archiving. At the time of installation, the rule set at the volume's root for the wildcard ***.*** is "Archive After 6 Weeks" (after a file has been stable for six weeks without modification), but this setting can be modified by the network administrator, if desired.

Is this appropriate? If a file version is less than six weeks old, one of two things will happen—either it will be modified or deleted before it stabilizes, or it will reach the age of six weeks. If it is deleted or

modified before reaching stability, it is very likely that the data in it was of a temporary nature, and should not be recorded forever.

On the other hand, if a file reaches stability, it falls in one of two classes—important data that should be archived (considering the extremely low cost per megabyte of backup media, and the high cost of data reconstruction), or it is unimportant information abandoned by the user. If that is the case, archiving is appropriate to facilitate migration of the file when fee disk space runs low.

For these reasons, a file version's age is used to decide whether or not it should be archived. If the network administrator prefers to make occasional permanent copies of some files that change frequently (for example, an accounts receivable database), a rule of "Archive Every 4 Weeks" could be used.

The Archiving Process

With its default rules in effect, Storage Manager will consider a file to be stable when its DOS "modified" data is more than six weeks old. At every rotation, Storage Manager will perform an archive operation on each volume, in addition to the normal backups.

These archive sessions are appended to the front section of backup media, then the backup sessions are appended to the end of the archive area. (If the "Put Archive on Separate Media Than Backup" configuration option has been chosen, archive sessions are written to physically separate media—providing more room for the backup sessions on their own media, but requiring a media change after the archive session. This option is available on all media except the DC6000 systems, where they are required).

The files that are eligible for archiving will be added to the archive areas of several separate backup media. (A default value of three copies is used, and highly recommended). After the three archive copies are made, the file is considered fully protected, and that version of the file is not archived again.

For files that are dormant (not being used), archiving stores permanent copies on inexpensive media, allowing the on-line copy to be migrated in the future, if the network administrator should choose to do so. For files that are in use, archiving provides essential protection of the current data in the file. Should the file be modified in the future (as is common with many quarterly and annual business files), the existing archive copies remain on backup media, and new ones are added when the on-line version stabilizes. Through this action, files that have had multiple stable versions display a rich history in the user interface—with each stable version protected forever, in the permanent section of multiple backup media.

This protection is especially important with executable files. Not only will previous versions of the application be available, but older copies are also available for restoring in the event of a virus attack.

Chapter 8

Backup

Overview

This chapter reviews backup strategy, backup options, and explains Storage Manager's efficient intelligent storage management methods.

Purpose and Strategy

Purpose

Complete daily (and in some cases hourly) data protection, once considered an afterthought, is now considered a necessity due to:

- Large increases in LAN size
- The complexity of interrelated user applications
- Disaster recovery plan requirements
- Company policy requiring periodic LAN “snapshots”
- Duplicate media

Strategy

A backup strategy must maintain at least one copy of each file on a server or client resource. In addition, good backup strategy will:

- Copy each file version to multiple physical media to guards against backup media failure
- Provide an intelligent vault rotation scheme (Disaster recovery planning)
- Provide a selection of previous versions of files that have changed

Intelligent Backup

Storage Manager was designed with an intelligent backup engine, using media rotation schedules, file rules, and intuitive error recovery capabilities to fully protect files.

Commonly, 60 to 70% of LAN disk storage involves stable files that are seldom modified. Storage Manager copies stable files to 3 media sets (by default), then skips these *fully protected* files during subsequent automatic operations (unless otherwise configured). Some of the advantages of intelligent backup are:

- Backups complete more quickly
- Typically, less media are required
- Smaller capacity backup devices can be used

When resources cannot be accessed for backup, Storage Manager intelligently handles the situation and protects as much data as possible without terminating the backup process.

The intelligent nature of Storage Manager greatly simplifies the installation process. No scripts are needed—the user specifies which volumes are to be protected by Storage Manager, modifies the default rules if desired, and Network does the rest. In addition, Storage Manager reacts properly to error situations.

Automated Backup Example: Assume that there are 10 volumes on the Protected Resource List (Vol1 through Vol10). On rotation day, VOL8 was not available at backup time. The system administrator reviewed System Messages (the message log) the following morning and noticed that Vol 8 had not been protected. Vol8 was then made available and the backup job re-submitted to backup files missed the prior night.

The following table illustrates how the volumes' files were treated when Vol 8 was not available and when the operation was re-submitted.

AUTOMATED BACKUP INTELLIGENCE

Original Backup Operation	Re-submitted Backup Operation
Vol 1 through Vol 7, Vol 9, Vol 10 - Full backup of all files not fully protected	Incremental or Differential backup of files modified since prior night's backup
Vol 8 - No backup of any files, volume skipped, message posted to System Messages (message log), operation continued	Full backup of all files not fully protected, including any modified for Vol 8 since prior night's backup

Workstation Backup

To protect client workstation volumes, Storage Manager uses a combination of workstation and server TSAs.

Workstation TSAs

For the backup server and workstation to communicate, proper TSAs to match the client platform must be loaded.

For example, DOS and WINDOWS clients require TSASMS.COM to be loaded at each workstation, but TSADOS.NLM needs to be loaded only once on the LAN server to perform backups. OS/2 clients require TSAOS2.EXE to be loaded at the workstation and TSAPROXY.NLM to be loaded on the LAN.

See Chapter 4 for more on the relationship between server and client TSAs.

Managed Operations

Automatic

Automatic backups are the backbone of file protection strategy and serve to protect all file data on a daily basis. There are three automatic backup options available:

- **Full Backup** copies all files to media from the network except fully protected files or files already on the media. Full backups occur on rotation days.
- **Incremental Backup** copies files to media that have changed since the last backup operation and by default occur between media rotation days.
- **Differential Backup** copies all files modified since the last full backup, which means files that are changed will get copied to the same media redundantly.

Periodically Scheduled

Besides using the default backup schemes, jobs can be scheduled to occur at different intervals, using the Job Scheduler. For example, if a "snapshot" was desired of a volume once a month, that volume could be tagged by selecting *Operation*, *Backup*, *Full*, and *Schedule* when the job is to execute (day, time, hour, and repeat frequency).

Storage Manager's Job Server monitors for scheduled jobs and submits backup requests to the Job Queue at scheduled dates and times.

Custom Operations

The primary advantage of Storage Manager over second generation systems is its ability to intelligently determine which operations are needed, and to perform them automatically—under the direction of the rules defined and configuration options selected by the network administrator.

There may be situations, however, in which the network administrator needs to perform a manual operation beyond the automatic protection. Common examples include copying a large set of databases to backup media for shipment, or performing a full backup of a volume prior to its replacement to allow a single-media restore.

Manual operations may be performed at any time, either from the command line, or through the user interface. With the exception of export operations, these backup sessions are added to the normal backup media library, and are recorded in the databases exactly as an automatic backup would be recorded.

Tracked and Untracked Sessions

Tracked sessions are recorded in the File History Database as any other session and automatically written by Storage Manager.

Untracked sessions are not recorded in the File History Database and should only be used if the media is to be sent off-site, as the media will not be available for restores.

Chapter 9

Managing Primary Storage

Overview

This chapter reviews the Storage Manager's migration feature and how it helps network administrators determine what infrequently accessed files can be removed.

Definition and Purpose

Primary storage refers to the fastest media used for permanent storage of data. On most LANs, this takes the form of hard disks in (or attached to) the servers.

In a NetWare file server, several megabytes of RAM are usually dedicated to caching the most recently used sections of the server's disks. This caching significantly improves the system's overall throughput, because many read requests can be serviced without waiting for the much slower physical disk.

Because RAM chips cost so much more than hard disks (on a dollar per megabyte basis), it usually isn't possible to cache the entire file system. A trade-off between cost and performance is reached. Data that is very frequently used is generally located in RAM, less often used data is generally not cached.

A similar price difference exists between hard disk and tape media. The price of backup media, in dollars per megabyte, is usually well over 100 times more affordable than hard disk storage. On the other hand, data stored on hard disks can be accessed much more quickly than data stored on tape.

One of the challenges of managing primary storage is determining what data should be on primary storage, and what data could be migrated to the less expensive media. As long as there is ample free space on the primary media, there is little need for file migration. When the disk fills, however, the network manager is faced with the choice of buying more primary storage, or moving some data off-line.

On most LANs, as much as 30 to 40% of the data on the server's primary storage has not been accessed in months or years. A great deal of money can be saved by identifying these files, copying them to multiple backup media, and removing them from primary storage.

Manual Migration

Unfortunately, the process of identifying which data could be migrated, and performing the migration safely, can be difficult at best. The key to identifying which files are unused is the "last access" date. This date identifies the last time the file was opened for any purpose—creation, updating, reading, executing, etc. This would include times when the file was opened for backup.

Identifying Eligible Files

Unless the backup system "covers its tracks" and preserves the "real" last access date, it appears that every file was used at the time of the last full backup. Even if the backup system preserves that date, and a tool is available to identify files that have not been used for a period of time, other difficulties emerge. Some files should not be migrated no matter how long they've been unused. (For example, some NLMs in a server that has been running for six months without rebooting, will appear to be dormant. Migrating these files would prevent the server from booting correctly after the next shutdown.)

Even when an appropriate list of files is generated, there is still the task of making multiple permanent copies of these files, in the event they are needed in the future. Having multiple copies on backup media (with at least one copy off-site) assures that the file will be accessible when needed.

Disk Full Conditions

Manually identifying files for migration could be considered an inconvenience, if it weren't for the nature of storage management. All too often, migration is performed because a volume has little or no free space left. If space isn't cleared in a matter of minutes, serious problems can occur. Applications can crash or hang, print jobs can be terminated, and even NetWare's Transaction Tracking System can run out of room. Under these situations, the minutes or hours lost while selecting and migrating files can have serious consequences.

Intelligent Migration

Storage Manager's disk management and intelligent migration capabilities relieve administrator of having to manually track disk usage and manually migrate files.

Storage Manager's built in disk management features include:

- automatic resource monitoring
- determining which files are eligible for migration
- automatically migrating files when volumes reach their configured high water mark
- automatic recall of migrated files

Resource Monitoring

Resource monitoring can be configured for your entire Protected Resource List (using Configuration Manager) or for individual volumes (using Resource Manager).

With resource monitoring and automatic migration enabled, when a volume reaches its configured high water mark (default=90%), a migration operation is submitted to the job queue in an attempt to bring the volume's disk usage down to the low water mark (default=80%)

Migration Rules

Storage Manager's rules allow the administrator to define under what conditions files should be considered eligible for automatic migration. The default rule, migrate after 12 weeks, is in effect when the system is installed. Under this rule, any file that is stable, fully protected

(permanently archived to at least three different media sets), and not accessed for over 12 weeks is eligible for migration.

Different files receive special migration rules upon installation of Storage Manager. The bindery files, for example, include the rule "Migrate Never."

Network administrators can further customize their migration rules. Common examples include "Migrate on Demand" for SYS:\PUBLIC files (migrate only if specific files are manually selected), "Migrate on Demand" for all EXE and COM files, and "Migrate after 52 weeks" for certain directory trees. In addition, the default rule itself can be modified, changing the default eligibility period for an entire volume.

Eligible File Lists

Prior to migrating files, Storage Manager builds a list of files eligible for migration called the prestaged file list. Building the prestaged file list happens at every resource-level migration operation (if no list exists). If an automatic migration operation does not result in a volume reaching its low water mark, a new prestaged file list is built and another migration operation is performed.

If desired, you can configure Storage Manager to periodically build the prestaged list so that the list is always up-to-date at any automatic migration.

Files are added to the list only if all of the following criteria are met:

- The current version of the file is fully protected.
- The file's Migrate Rule allows it to be migrated.
- The file has not been accessed for at least the time period specified in its effective rule.
- The file is on the near line set (if configured)

Because of Storage Manager's continuous archiving process, files will automatically be fully protected before they are old enough to be migrated. Once the list is generated, the files can immediately be removed from the server's primary storage.

Volumes on Hold

If the volume does not get below the high water mark, a message is posted in the error log and the volume is put in an "on hold" state. By simply changing the migration rules or reconfiguring the high and low water marks, the volume can be made eligible for migration operations again.

Automatic Recall

One drawback to migrating files is that users may think a migrated file is unretrievable, and re-create a file rather than simply restoring it from backup media.

To overcome this limitation, Storage Manager allows end users to recall migrated files.

To recall migrated files, you need:

- a recall agent loaded on the workstation or server
- a phantom file left in place of a migrated file

See chapter 4 for more information on the recall agents.

Phantom Files

As its name suggests, a phantom file is an image of the file that was once there. To the user looking at a disk directory, the file looks identical, with the exception of its 0-byte file size.

When Storage Manager migrates a file and creates a phantom file, it truncates it to 0 bytes, and marks it as a phantom file. (Storage

Manager can distinguish between normal 0-byte files and phantom files. In File Manager, a ghost icon appears next to phantom files.)

Leaving a phantom file still requires a directory entry on the server, but it completely frees the data space occupied by the original file.

Near Line Storage

Near Line storage refers to data stored on devices other than your primary disk storage (for example, on an optical jukebox). The benefit of near line storage is realized with automatic recall as migrated files are made accessible to end users as if the files were still on primary disks. Storage Manager provides two parameters that utilize near line storage: near line set and Near Line Device.

The Near Line Set parameter (on the Archive/Migrate tab in Configuration manager), when enabled, creates a unique media set that contains a archive copy of every file eligible for migration. That is, before a file can be migrated, an archive copy of the file must exist on the near line set.

The near line set is your most frequently used media set and is never rotated offsite. For optimal use automatic recall, Palindrome recommends creating the near line set.

The Near Line device parameter (Configured in Device Manager) restricts a device to be used exclusively for your near line set and allows instant recall of migrated files (assuming the entire media set is available in the device). If all the media in your near line set is always in your backup device, you never have to mount media for a recall request. This option is especially useful if you have multiple devices (for example, a tape autoloader and an optical jukebox).

With near line device set to YES, you can preserve your more expensive media (optical) for more critical data (your migrated files) and use your other device (tape autoloader) for all other backup and archiving operations.

Chapter 10

Media Management

Overview

This chapter reviews Storage Manager's automation strategy for managing backup media.

The Media Library

Library IDs

At installation, Storage Manager is given a Library ID that becomes the prefix of the label of all managed media for that library. The Library ID uniquely identifies media for an individual installation. The Library ID may be changed after installation, but there are some considerations:

- All media in the original library are retired (discussed later in this chapter).
- Archived (permanent) file copies that were counted toward a file's full protection are re-archived, if available
- The prior Library ID cannot later be reused.
- The original Library ID's media information remains in the System Control and File History Database(s)
- Additional media are required to build the new Library.

Sets

Storage Manager organizes managed backup media into groups called media sets. Each set has a prefix (Library ID) a letter (A,B,etc.) and a number (2491,1880, etc.). For example, the first media set within the NEWLIB Library ID would be NEWLIB:A:1, the second media set would begin NEWLIB:B:1, and so forth. Storage Manager provides for media sets A through ZZ, though this many is seldom attained. Because additional archive (permanent) sessions are periodically added to media and not erased, the media will eventually fill and a new media will be added to the set. For example, when NEWLIB:A:1 fills, a blank will become NEWLIB:A:2 .

Labels

Storage Manager allows both managed and custom media. During automatic operations, managed media are automatically labeled and are tracked within the File History Database.

Custom media are created for special backup operations, and are not part of the managed rotation media sets. Custom media labeled by the user at backup time may be named “MONDAY”, “MARCH”, “USERNAME”, or etc. By default, sessions on custom media are tracked within the File History Database, but they do not have to be. For example, if a particular media is to be sent off-site, and never used again, the sessions can be designated as untracked to save File History Database space.

Layout

Storage Manager writes Archive sessions at the beginning of media and appends Backup sessions to them. (In addition, an updated System Control Database and File History Database session is written following each Archive and Backup session). *See Appendix A for a complete media layout view.*

Layout of Archive/Backup Sessions on Media

To keep from “trapping” erasable Backup sessions between permanent Archive sessions when automatic rotation occurs, Storage Manager first erases all existing Backup sessions, appends additional Archive sessions, then appends new Backup sessions.

If the “put archives on separate media from backups” option is turned on through Configuration Manager’s Configure/Operations, all Archive sessions will be copied to separate media from Backup sessions. With a single backup device, this would necessitate physically changing media during rotation day operations. With an autoloader device or two backup devices configured, one device can hold only Archive media while the other holds only Backup media, allowing complete backups without intervention.

DC6000 backup devices require archive and backup sessions to be on separate media, while with other devices it is optional.

Sessions

Each Archive, Backup, System Control Database, and File History Database session has a unique identifier composed of letters (SV, CP, DH, or DC) and numbers (*nnnn*).

The following example is typical of session identifiers as written to media during an automatic backup operation:

“SV1”= **Archive** session (permanent, never erased, stable files)
“DH1”= **File History Database** session
“DC1”= **System Control Database** session
“CP2”= **Backup** session (erasable, evolving files)
“DH2”= **File History Database** session
“DC2”= **System Control Database** session

Every time an **Archive** or **Backup** session is created, the unique session number (*nnnn*) is incremented and stored within the System Control Database.

NOTE: System Control Database sessions contain information for **all** media and **all** volumes, not just the last session created for a specific volume.

File History Database sessions contain all history information for the current volume, not just information for the latest session written.

Archive (SVnnnn) sessions:
Archive sessions contain information on LAN data files that are stable. These sessions are never erased from media.

Backup (CPnnnn) sessions:
Backup sessions contain information on LAN data files that are evolving. These sessions are normally erased when the media is called into rotation for an automatic operation.

File History Database (DHnnnn) sessions:
File History Database sessions contain information for all Archive and Backup sessions (on all media within the Library ID) for one *volume*.

System Control Database (DC*nnnn*) sessions:
System Control Database sessions contain information on all media (and associated Archive and Backup sessions) for all resources for one *installation.*

It is important to select the proper installation prior to viewing databases or restoring databases, as each Library ID, and its associated databases, is unique to a specific installation.

If Storage Manager detects that there is no System Control Database, it will attempt to restore the latest version from media. If the latest version is corrupted or the tape containing that version has been destroyed, an earlier version can be restored. Restoring an earlier System Control Database requires that the latest File History Database session, one per volume, be restored for that same prior date. An error will result if the two databases are out of sync.

Media Diagnostics

Intelligent media management begins with intelligent backup software. Storage Manager manages the rotation of media to and from the vault in a manner that assures a diverse set of file versions and copies, and assures that the physical wear is spread over a number of media. See *Chapter 11* for rotation detail.

Media do wear out eventually, so starting with approved, quality media is important. As media wears, software should detect this condition and warn against impending media failure.

Modern backup devices contain error detection circuitry to guarantee that data is correctly recorded on backup media. Tape drives can detect "soft errors" (excessive rewrite tries) and Storage Manager logs this information into System Messages (the message log). Excessive soft errors can indicate physical wear due to the number of passes made on each media, the quality of the media used, and possible dirty heads or head wear on backup devices.

Retensioning Tape

Sometimes otherwise reliable, quality media that is vaulted for a long period may need "loosening up" prior to use. This is called *retensioning* and can be accomplished by running the media through the backup device. This is not necessary with most backup devices and is not possible with 4mm Digital Audio Tape (DAT) drives and 5 gig 8mm drives. However, it is often needed with DC6000 150/250 meg drives.

Retiring Media

As media deteriorates (often a flaking off of the oxide layer), the soft error rate will rise. Sometimes all that is needed to correct high soft

error percentages is to run a cleaning tape through the backup device. Failure to heed excessive soft error warnings can result in media that becomes so worn that neither backup nor restore operations are possible. Continual use of a marginal tape may damage the drive read/write heads, causing damage to other tapes and possible expensive repairs.

While a media is still readable for restores, Storage Manager provides a means to retire that tape from rotation, still tracks it in the file history database, and allows for file restores from that media. This option is available through Media Manager's Operations/*Retire*.

When media are retired, a new blank media may be requested to continue the media set for the next backup operation. For example, if NEWLIB:E:2 was retired, a blank to become NEWLIB:E:3 may be requested.

Forgetting Media

If media becomes unreadable or has been lost and files cannot be restored from it, all references to files that had existed on that media should be removed from the file history databases. Archived files that counted toward full protection should be re-copied onto new archive sessions on other media.

Media Manager's Operations/*Forget* option removes all references to that backup media from the history databases—as if the media never existed. This ensures that the databases continue to accurately reflect what is actually available on the backup media, and does not give the user a false sense of security.

Chapter 11

Media Rotation

Overview

This chapter provides an in-depth look at the Storage Manager's rotation strategies: the Tower of Hanoi Weekly (the default), Tower of Hanoi Daily, and Grandfather, Father, Son.

Rotation Scheduling

Rotation schedules automate the process of moving tapes into rotation from vault storage to provide multiple copies of files in case of onsite disasters. In addition, they provide a diverse mixture of versions of each file for restoration. Complex schedules can be automatically implemented without the need for time-consuming calculations and monitoring by the system administrator.

Grandfather, Father, Son Rotation

Storage Manager's Grandfather, Father, Son (GFS) rotation schedule assigns unique media labels to media sets reserved for Daily, Weekly, and Monthly backups and optionally also provides for Quarterly and Annual sets. The media sets have the following characteristics:

Daily media set:

- Defaults to seven daily media sets, with daily media changes.
- Media are (LIBRARY ID):Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday.
- Is scheduled for a Full Backup operation each day.
- Does not contain any archive sessions.

Weekly media sets:

- Are rotated once each week (default rotation day is Friday.
- Defaults to 5 media sets (LIBRARY ID):WEEK1, WEEK2, WEEK3, WEEK4, and WEEK5.
- Both archive and backup sessions are copied to media.

- Can be removed from rotation to protect backup sessions using GFS option to Retire Weekly Media Set.

Monthly media sets:

- Used once per month.
- By default contain 12 media sets (LIBRARY ID):JANUARY, FEBRUARY, etc. (If configured for less than 12 media sets, they are labeled MONTH1, MONTH2, etc.)
- Default rotation day is *last day of month*, but is configurable.
- Both archive and backup sessions are copied to media.
- Can be removed from rotation after backup using GFS configuration option to *Retire Monthly Media Set*.

Quarterly media sets:

- Can be configured for starting month (default is March).
- Rotate quarterly, beginning with month configured, and at three-month intervals thereafter.
- Both archive and backup sessions are copied to media.
- Media sets are (LIBRARY ID):QUARTER1, QUARTER2, QUARTER3, and QUARTER4

Annual media sets:

- Used once per year.
- Both archive and backup sessions copied to media.
- Default month/day is December 31, but is configurable.
- Media sets are (LIBRARY ID):YEAR1, YEAR2, etc.

NOTE: Media ***sets***, not specific media, have rotation schedules. The exact day of the week that a specific Daily media is used may vary each week, depending upon free space on media, last date used, etc. For example, Daily media *Monday:1* may be used on one Monday and *Monday:3* on Monday of the following week, but all are from the *Monday* media set.

"Tower of Hanoi" Rotation

The "Tower of Hanoi" rotation schedule gets its name from a popular mathematical puzzle that involves moving rings back and forth between three posts. This rotation scheme uses the same method to rotate media as the puzzle uses to rotate rings between posts.

Tower of Hanoi rotation is commonly used for data protection in large mainframe and minicomputer environments and was adopted by Palindrome as one option for PC LAN data protection.

When Storage Manager is installed, the rotation is set to default to *Tower of Hanoi Weekly* rotation and uses five media sets NEWLIB:A through NEWLIB:E, and rotate as follows:

NEWLIB: Media Set	A	B	C	D	E
Weeks Between Rotations	2	4	8	16	32

As some sets come into rotation frequently (NEWLIB:A and :B) and others less frequently (NEWLIB:C, D:, E:), diverse sets of files remain on media at all times, some copies of files are always vaulted offsite, and physical media wear varies in intensity. A maximum of twelve media sets can be in managed rotation.

Weekly and Daily Rotation Schedules

Storage Manager provides the option of choosing either Weekly or Daily Rotation schedules, with configuration options for each. The default at installation is *Tower of Hanoi Weekly* rotation.

The pattern of media sets used is independent of whether the media is changed on a daily or weekly basis. Similarly, the frequency of media changes is independent of the frequency of backup sessions. For example, media could be changed weekly, with backups run several times a day.

Storage Manager allows the user to select weekly or daily rotation. With weekly rotation, the user may specify on which day of the week rotation should occur.

Under weekly rotation, Storage Manager requires a minimum of five media sets, guaranteeing a minimum backup horizon of eight weeks. (Note that archive sets are not erased, and have an infinite horizon. Only backups, containing non-stable versions of files, are recycled.)

Under daily rotation, the schedule is accelerated by a factor of five, because five rotations occur each week, Using only five media sets would result in a very short backup horizon, as little as eight working days. To compensate, Storage Manager requires (as a minimum) that a sixth and seventh media set be added, extending the horizon to 16, then 32 working days. This 6 1/2 week horizon is then similar to the 8 week horizon provided by weekly rotation.

Media Usage

Depending upon the rotation schedule choice, the following maximum number of media sets are allowed:

- Tower of Hanoi Weekly defaults to five media sets but will allow a total of twelve sets.

- Tower of Hanoi Daily defaults to seven media sets but will allow a total of twelve sets.
- Grandfather, Father, Son defaults to 7 Daily, 5 Weekly, and 12 Monthly media sets but up to a total (all types) of 32 media sets are configurable.

The number of media sets required in your library depends upon:

- The number of rotation days (configurable).
- The rotation schedule chosen (GFS, Tower of Hanoi weekly, Tower of Hanoi daily)
- The *Preserve Backups* field configuration (default is 0 days).
- The maximum allowed for each rotation choice (12 for Tower of Hanoi and 32 for Grandfather, Father, Son).

Which Backups Remain on Media?

Archive sessions are permanent, as they are never erased. Backup sessions on media are erased at each rotation, providing configuration has not been set to preserve them for a period of time (*Preserve Backups* configured). When rotating often, Backup sessions on these media have the potential for being erased and replaced as often as every other day. When rotating infrequently, some media in sets seldom come into rotation and may contain file versions that are months old.

For example, if an “E” media set was comprised of media E:1, E:2, and E:3, at some point each of these media will be in rotation and all Backup sessions erased from them. E:1 may have been used a month ago and contain all of the sessions from that point in time, E:2 may have been used last week and all original sessions erased and replaced, and E:3 may have been last used two weeks ago and has not come back into rotation a second time. In all cases, media E:1 through E:3 will contain any Archive (permanent) sessions.

How Long Are Sessions Preserved on Media?

Sessions on media may be either permanent or erasable.

Permanently protected sessions include:

- Archive sessions (are automatically permanent)
- All sessions on retired media (no longer in rotation)
- Sessions created by Custom Operation backups to non-managed media
- Sessions on Grandfather, Father, Son rotation Weekly, Monthly, Quarterly, or Annual media sets if configured to *Retire Media Set* .
- All archive and backup sessions for a complete Library ID when the Library ID is changed, as all media in the current library are automatically retired.

Erasable (temporary) sessions include:

- All Backup sessions (when not protected through optional configuration setting to extend their life cycle)
- Sessions within Grandfather, Father, Son *Daily* media sets
- Sessions on Grandfather, Father, Son Weekly, Monthly, Quarterly, and Annual media when not configured to *Retire Media Set*

NOTE: Due to the intelligence built into Storage Manager rotation schedules, Backup sessions, though normally erased when rotation occurs (the default), are often available for months without configuring to extend their life cycle.

The Scheduling System and Blank Media

After installation, Storage Manager prompts for the insertion of blank media at each rotation until the minimum number of media sets has been created.

During the course of automatic operations, Storage Manager can predict the most logical media to use to allow a complete backup to occur.

Should a media fill, Storage Manager will prompt for the insertion of blank media to continue the backup to include all scheduled sessions. Archive sessions on media are permanent, and append on media as additional file versions qualify. When a media is filled with archive sessions, it will no longer be used in rotation. When the option to put archives on separate media is configured, only media configured for Backup will be considered when determining which Backup media to use.

Use of optical jukeboxes and tape autoloaders can make the insertion of blank media automatic, requiring no human intervention and no delay in completing backups (an inconvenience when users have to wait while the backup operation completes).

Adding Media Sets

The number of media sets in the Library can be increased using Configuration Manager (and Storage Manager will automatically add media sets when necessary). There are a number reasons to add media to the Library:

- To increase the default minimum number of media sets to enrich the number of file versions available for restores.

- To increase the life cycle of each media (number of days between rotations of each existing media set is decreased by half with each media set added).
- The number of archive copies has been configured above three copies (the default).
- The installation has been configured to put archive sessions on separate media (required with DC6000 series backup devices).
- The rotation has been changed from Tower of Hanoi to Grandfather, Father, Son (which automatically retires the existing media and also requires more media sets).
- Quarterly or Annual media sets are configured (default is zero sets) using Grandfather, Father, Son rotation.

Retiring Media

Both sets and individual media in a set may be retired. Reasons for retiring media are:

- The Library ID has changed (automatically retires all media within library). This can be due to management policy (for example, retire all media at end of each fiscal year) or because changing Rotation Schedules has forced a Library ID change.
- Weekly, Monthly, Quarterly, and Annual media (GFS rotation) have been configured to automatically *Retire Media Set* after each backup operation.
- Data on specific media need to be protected from erasure.
- Specific media need to be sent offsite and taken out of rotation.
- Media diagnostics indicate excessive physical wear (media is marginal for backups), but needs to be available for restores.

NOTE: Archive copies of files no longer count toward *full protection.* If the file is still available, additional archive copies of affected files will be re-archived onto other media until full protection is again attained.

The Rotation Process

A rotation occurs whenever Storage Manager performs an automatic backup to a media set other than the one that was used for the last automatic backup. (Remember that "automatic" means Storage Manager decided which operation should be performed. "Automatic" backups can be manually or automatically initiated). During the rotation process the following events occur:

- All Backup sessions are erased (not archived).
- An archive operation is performed on each volume for any files that have qualified for archiving (according to archive rules set for the files).
- A Full Backup is performed on all files.
- Files that have been deleted from disk since the last rotation are noted in the File History Database(s).
- Files eligible for migration are deleted from disk when the disk is nearly full (according to the Migrate Level setting), providing Auto Migrate has been enabled.
- Following each archive and backup session written to media, the File History Database associated with the session's resource is written to media in a DHnnnn session and the System Control Database is written to media in a DCnnnn session.
- The System Control Database and File History Database are updated on disk.

Incremental Backup

Incremental Backups are the default automatic operation on non-rotation days. Incremental Backups copy all files that have

changed since the last backup operation (typically 20 to 30% of all LAN files).

Differential Backup

Differential Backup operations copy all files modified since the last full backup. For weekly rotation with daily media changes, copies of files modified since the last rotation could be included multiple times on the same media.

NOTE: Files copied to media during incremental backups and differential backups do not count toward a file's fully archived protection, as these operations do not include any Archive Sessions.

Chapter 12

Data Integrity

Overview

This chapter provides a review of different methods Storage Manager employs to ensure data integrity when written to backup media.

Introduction

If the primary purpose of a backup system is to safeguard important data, great care must be taken to ensure data integrity. This effort is reflected in every aspect of Storage Manager's design.

Device Failure

Data loss or corruption can occur during the process of writing data to, or reading data from, backup media. All backup devices supported by Storage Manager support multiple levels of error correction and detection. An examination of the error correction capabilities of 8mm tape drives will illustrate some of these capabilities. (These features are shared by all 8mm tape systems supported by Storage Manager.)

Ensuring Proper Recording of Data

On 8mm and 4mm (DAT) tape drives, a method of recording called "helical scanning" is used. (For details, see Appendix A.) In a helical scan system, the tape is partially wrapped around a rotating metal drum. The read and write heads are located on the outside edges of this drum. As the tape slowly moves by the rapidly spinning write heads, a series of diagonal "stripes" of data are recorded on the tape.

Immediately after each block of data is written onto a data stripe by the write head, a read head passes over the stripe, reads the data, and compares it to the original information on disk. If an error is found, that block of data (still available in the tape drive's buffer) is re-written to another stripe a fraction of a second later. The process repeats until the data is written successfully or the high soft errors level causes a fatal tape error and the process terminates.

Ensuring Proper Reading of Data

In addition to storing the block of data, the tape drive also stores a sequence of bits for error correction and detection. This pattern of bits is created by applying a mathematical formula to the data in the block. When the data is read back in, the error correction information can be used to determine if the original data block is accurate. (The data block is read from tape, the formula is applied, and the calculated error correction pattern is compared to the previously stored bit pattern for that data block.) If an error occurs during the bit pattern comparisons, a read error would be posted to System Messages (message log).

If a reasonably small number of bits were damaged or read improperly, the error correction pattern makes it possible for them to be identified and corrected.

These built-in error detection and correction techniques (which occurs in the tape drive while it is reading or writing at full speed) allow incredibly low undetected error rates. Published statistics indicate that the undetected error rates for helical scan drives may be as low as one bit in thousands of gigabytes of stored data.

Soft Errors

"Soft errors", detected and recorded while relocating improperly recorded data to better sections of tape, are invisible to the user. Since the tape drive makes these statistics available to the backup program, Storage Manager monitors them and alerts the user if the errors rise to an unacceptable level.

Soft errors occur during most backups and are not abnormal as tape backup media contain imperfections. Soft errors do not result in lost data; they indicate that the hardware had to take extra steps to read or write the data *successfully*.

If soft errors occur only on one tape, it usually indicates a marginal tape—whether through a manufacturing defect, abuse, normal wear, or simply a need for retensioning.

If soft errors occur on multiple tapes, it usually indicates a drive read or write head in need of cleaning, a bus termination problem, or a dirty air filter in need of replacement. If cleaning does not remedy the problem, the drive may need repair.

By alerting the user to unusual soft error levels, Storage Manager makes it possible to retire and replace defective tapes or repair failing drives before a backup session fails or data is lost.

Media Failure

Even with the great care taken to ensure reliable recording of data, some media failure is inevitable. Portions or all of some media may eventually become unreadable, may be lost, accidentally reformatted, or destroyed. With the long-term nature of data protection, planning for media failures and possible lost data is absolutely essential.

Multiple Media

Storage Manager makes use of "managed redundancy" whenever possible. If a file exists long enough, Storage Manager will ensure its storage on multiple media. When a file is eligible for archiving, Storage Manager makes sure it is permanently stored on several media, at least one of which is always off-site.

Retiring Media

When backup media is the cause of unacceptable soft error levels, it should be taken out of rotation, so that data may still be restored from it. There are a number of ways to retire media from rotation:

- Highlighting the media in Media Manager and selecting the *Retire* option.
- Changing the Library ID
- Adding additional sets to the Library through Configuration options (does not retire media, but brings it into rotation less often so that wear occurs more slowly)

Forgetting Media

Occasionally, a backup media will be lost, destroyed, or damaged so badly it cannot be read. Storage Manager includes a “forget” option that removes all references (to the specific media) from the File History Databases. Blank media are requested to replace any media that has been forgotten. For example, if the E media set contains E:1, E:2, and E:3 and the E:2 media is forgotten, the next media added to the set will become E:4.

Any stable files that are not fully protected after the media is forgotten will then be archived to other media, ensuring that full protection is maintained. Storage Manager intelligently reacts to lost or damaged backup media, and takes all possible steps to maintain complete protection.

Note: Using Media Manager, “Forgotten” media can still be used to restore data. The contents of sessions on the media can be viewed, file(s) highlighted, tagged, and restored.

File Contention

A less obvious danger to data integrity is file contention, which involves a file being used by another station when the backup occurs.

Unshareable Files

Files can be opened either shareably or privately. If a file is opened privately, no other station or program may open it. When Storage Manager encounters an unshareable file, it notes the occurrence in System Messages, and continues the backup.

Shareable Files, Opened for Reading

While performing backups, Storage Manager opens files for reading, specifying “Deny Write”. NetWare will grant access to Storage Manager as long as no other user has opened the file for writing.

When a user is loading an application, for example, the .EXE or .COM file is briefly held open (shareably) for reading. Storage Manager will backup these files, since the file was available and there was no risk of losing data integrity.

Shareable Files, Opened for Writing

Files shareable and open for writing are common with the multi-user databases found in many LAN applications.

When Storage Manager attempts to open these files for reading with the “deny write” option, the request is denied. (Since other stations already have write access to the file, Storage Manager cannot be guaranteed that no changes will be made to this file.)

This presents an interesting dilemma. Should the file be backed up, or should it be skipped?

- **Backup file considerations**:

  Copying the file to media introduces a risk of data integrity.

  Suppose, for example, that this database has several “key fields” on which it is sorted—fields like last name, zip code, account number, etc. and a user is updating a record near the end of the database. After the backup system has read the index area (near the beginning of the database), the user changes the customer’s account number (near the end of the database). The data copied to media will include an old customer account number index reference and a new customer account number data reference. All of the file information will have been copied without corruption, but will be incorrectly cross-referenced. If the file is later restored, serious problems could result.

  If someone searches for the new account number, it might appear the customer doesn’t exist in the index. Also, other databases updated by the indexed account numbers are unknowingly passed incorrect information. By the time the errors are evident, days or weeks may have passed, and recovery may be impossible.

- **Skip file considerations**:

  The conservative approach would be to skip the file during backup, even though it might have been possible to backup the file correctly. This would eliminate the risk of data integrity, but it would not protect the latest data, so the customer’s account information would be incomplete should a restore be necessary before the next backup opportunity.

  A “suspect” version of the file might be better than none at all. For example, if the network manager had been warned of the suspect file situation, the file could have been re-indexed after it was restored.

- **Storage Manager's approach**:

  Storage Manager solves this dilemma by backing up these files whenever possible, and marking them in the File History Database as having "suspect" protection. Storage Manager posts a warning in System Messages so that the system administrator is aware that the file's protection is suspect.

  If a file has been backed up with suspect protection, it can be tagged and backed up again so that a correct (non-suspect) version is on media. Most importantly, the system administrator is alerted that the file version is suspect when it is tagged for restore. If a non-suspect version of the file exists, it can be restored instead. If no other versions exist, at least it is known that corrective action may have to be taken on the suspect file copy after restoration.

User Error

Storage Manager goes to great lengths to both prevent and detect user error, and to ensure that data protection takes place as soon as possible after such errors occur.

Unavailable Volumes

Using the Protected Resource List, Storage Manager attempts to access and backup all files contained in volumes associated with LAN designated LAN resources. This includes client workstation volumes.

It is not necessary that a client workstation have a user logged in for backup to occur, merely that the proper TSA's for the workstation be loaded and that the workstation be turned on. User failure to leave the workstation in this state is the most common reason for lack of protection for user workstation volumes.

Storage Manager posts notes to System Messages regarding any volumes on the Protected Resource List that were not available for backup, thus giving the system administrator the opportunity to bring the workstations online and perform a special custom backup to protect the client data.

After posting the error, the backup process continues to the next resource on the list, backing up all volumes that are available.

NOTE: If a resource is online and physically available, but the proper TSA (Target Service Agent) is not loaded, the resource is skipped.

“Dangerous” Operations

Whenever a user attempts an operation that could result in the loss of data (such as formatting), the user is warned of the consequences and given the option to abort the operation.

In the event that a user formats a managed backup media, erasing its contents, Storage Manager reacts intelligently by removing all references to the media from the System Control and File History Databases and, where possible, re-protecting files that had been in sessions on that media.

Custom and Foreign Media

Occasionally users will insert the media from the wrong Library ID into the backup device. Storage Manager will not copy to media from the wrong Library ID and will alert the operator that the backup operation is being attempted on incorrect media.

This situation usually arises when a Library ID has been recently changed and media from the prior ID not taken to vault or when two backup installations are in close proximity.

Media created by Custom Operations are sometimes inadvertently inserted into the backup device, and also will not be used for managed backup operations.

Error Prevention

Storage Manager has the intelligence to determine whether a full, incremental, or differential backup, and archive or backup copy session should be performed next. In many installations, regular automatic operations are all that need be performed. Errors due to users running the wrong script file or batch file are greatly minimized.

Following all automatic operations, the system administrator may review the System Messages (message log) and the media scheduling screens for all installations and all resources to determine that the backup operations properly protected data for all resources.

If the event of a rotation error during an unattended backup (or during an attended backup where the user chose to override the “incorrect media” prompt), Storage Manager will take the best possible action it can to complete the backup operation. If it is simply a delayed rotation (the most common error), a backup will copy modified files to the existing media and notify the system administrator that incorrect media was used so that the correct rotation media can be inserted.

If the incorrect media is used, Storage Manager will perform its rotation procedure on the available media. (This will remove the obsolete Backup sessions, add Archive sessions, if appropriate, and add new Backup sessions.) It will then adjust the rotation schedule to account for the error. (Using the “bit” analogy, the rotation counter is “fast forwarded” to the next time the media set would have been used.) Compared to second generation systems, which may allow the erasing of archive data and not even record an error, this level of error handling is truly unique.

Appendix A

Media Layout

Overview

This appendix outlines Palindrome's method for organizing sessions written to media in SIDF (System Independent Data Format).

Why Use SIDF?

As companies are merged system administrators often find they are responsible for executing a disaster recovery plan that includes data restores from backup media created by different vendors.

Most vendors continue to utilize proprietary media formats with their backup products. Those that pledge future support of industry standard media formats often do so only as a secondary format, thus locking their customer's data into media written in a proprietary format.

Palindrome's NLM products fully support the System Independent Data Format (SIDF), which is being standardized by the international standards organization ECMA (European Computer Manufacturer's Association) as ECMA-208.

Palindrome supports SIDF as its native media format, writing to media in an industry standard format that ensures a long-term data restore solution.

"Logical" Media Layout

All Palindrome media includes the Media Index, Session Index, and Session Data.

Media Index

The Media Index is best thought of as a table of contents containing a listing of the sessions contained on that media.

The Media Index is composed of a series of fixed-length entries, each of which refers to one session on the media. The media's label (the name of the media, such as "MONDAY:1"), is stored in the description string of the first entry.

In the Media Index, each session's entry contains the following information:

- A "signature" that indicates this is a valid entry. The lack of a signature indicates that there are no more entries on the media.
- A description string, the name of the session (for example,"CP1047)".
- The date and time of the session, the names of the server (or workstation) and volume from which the files were copied, and other information about the session.
- The session's physical addresses on the media.
- The size of the directory portion of the session (the Session Index).
- The size of the actual data section of the session.

Sessions

There are usually many sessions on a media. During automatic operations, sessions are created corresponding to each backup or archiving operation performed for each volume on the Protected Resource List. If an operation requires spanning multiple media, separate sessions are created on each media. (A long backup operation that spans media appears to be one operation to the user, but is stored in separate sessions on each of the spanned media.)

Each session is composed of two sections: a Session Index that describes the contents of the session, and the Session Data containing the data from all files included in the session.

Session Index

The Session Index is composed of a series of variable length records, each of which corresponds to a file or directory on disk. Each entry contains the following information:

- The file's attributes
- The file's "modification date" or "DOS date"
- The size of the file's data within the session.
- System specific data, such as the NetWare "Last Accessed" date, etc.
- The name of the file or subdirectory.

Session Data

The Data Set section of each session holds the actual contents of every file in the session. The files are stored in the same order in which they

are listed in the Session Index. For each file, the following pieces of information are included:

- A data header signature, to verify that this is the start of the file.
- A copy of this file's entry from the Session Index is included for verification.
- The file or directory's full path name.
- The contents of the file.

System Control Database

The System Control Database is backed up during every operation. See *Appendix B for detail on the System Control Database files.*

File History Database(s)

The File History Database(s) are backed up during every operation for each volume. See *Appendix B* for detail on the File History Database(s).

See the following page for the general layout of the Media Index, Session index, and Session Data in System Independent Data Format.

Appendix B

Database Layout

Overview

The most important architectural element of Storage Manager is its compact, flexible, and high performance database structure. This appendix covers the location, contents, function, and interrelation between the files that make up the System Control Database and the File History Database(s).

Introduction

As explained in Chapter 4, there is one System Control Database per Storage Manager installation. There is one File History Database for each volume being protected by that installation of Storage Manager. (If two installations are both protecting one volume, that volume will have two independent File History Databases.)

All Palindrome database files have a .PAC (Palindrome Archive Catalog) extension.

The System Control Database

Overview

The System Control Database is stored in a directory on the volume selected at installation time. Storage Manager may be installed in any directory, but that directory name may not be changed after installation, as the original file and path names are referenced should older File History Databases for this installation need to be restored from media.

The System Control Database is used to store, in one place, all configuration information for all resources for one installation.

ASDB.PAC

(Archive System DataBase) This is the data portion of the System Control Database. It contains system-wide information and summary information about the protected volumes and media Library. The System Control Database contains information about System Controls, Protected Volumes, Media Sets, and Individual Media.

System Control Information

This section contains the following system wide information for one installation:

- The system's current Library ID (the label prefix for all managed backup media)
- The rotation options (daily or weekly, which day of the week)
- Job Scheduler information
- Auto-login user name and password
- Restore options (overwrite, etc.)
- Media rotation information (rotation counter, number of media sets in use)
- Media operation information (last operation, next operation, current operation)
- Last session number assigned
- Average backup size information (per rotation period)

Protected Volume Information

The following information is maintained for each volume:

- Sequence in Protected Resource List
- Server (or machine) and volume name
- Volume size

Media Set Information

The following information is tracked for each media set created by one installation of Storage Manager:

- If this set is active, its position in the rotation schedule

- Number of backup media in this set
- Total sizes of all archives and backups in this set

Individual Media Information

The following information is tracked for each individual media created by one installation of Storage Manager:

- Media set number (to refer to media set info)
- Last operation performed on this media
- Active/Retired flag
- Date of last format, update, and access
- Number and size of archive and backup sessions on media
- Media label (media name, such as NEWLIB:A:1)

ASNX.PAC

(Archive System iNdeX) This file stores indexing information for all "key" (sorted) fields in the System Control Database. This allows specific records to be located quickly, without scanning the entire System Control Database. Although ASDB.PAC and ASNX.PAC are separate files, they should be thought of as being a single, unified database.

The File History Databases

Overview

One File History Database is created and maintained for each volume on the Protected Resource List. This database contains file history for

existing files on that volume and deleted files for which copies still exist on backup media.

If more than one Storage Manager installation is protecting a volume, each will maintain its own File History Database in its own directory.

AVDB.PAC

(Archive Volume DataBase) This file stores the version of Storage Manager, the serial number of the version, and the resource that it protects.

AVPA.PAC

(Archive Volume Path) This file stores the directory names currently on the volume, the names of deleted directories that still have files on backup media, and deleted directories whose child directories still have files on backup media. The full directory names are stored, including the long Macintosh “folder” names for directories created by Macintosh clients.

AVFS.PAC

(Archive Volume File Status) This file contains information about each file on disk or deleted (or migrated) from disk but still on backup media. This includes the file’s status (deleted, migrated, phantom), as well as the number of archive copies. These records are linked to records in AVFH.PAC, the “file history” file.

AVLN.PAC

(Archive Volume Long Name) This file is used to store the full name for files that have long AFP (Macintosh) file names. Files with names that fit within DOS’s “8.3” convention are stored in **AVNX.PAC**.

AVFH.PAC

(Archive Volume File History) This file contains the archive and backup histories of files listed in AVFS.PAC. Each record contains information linking specific instances of files back to the media set,

individual media, and session on which it is stored. In addition, it stores each file's "modified date" (DOS date) and file size.

AVFC.PAC

(Archive Volume File Criteria) This file contains all rules defined for this volume. These rules are used to determine which files should be included in each backup, archive, or migration operation. In addition to storing the wildcard pattern for which the rule applies, each record includes the backup, archive, and migrate criteria, and is indexed into AVPA.PAC, making high speed "run time" decisions possible.

AVFT.PAC

(Archive Volume FileseT) This file tracks information about each session on backup media. It includes the session ID number, the type of operation which generated the session, the number of files in it, and the session's total size.

AVNX.PAC

(Archive Volume iNdeX) This file contains indexes for all key fields in the other File History Database files. These indexes form the relational basis for Storage Manager's high speed database work, making it possible for Storage Manager to rapidly build a "virtual directory" integrating on-disk files, deleted files, and previous versions of files into a single display.

Index

G

I

L

M

N

O

P

Q

R

S

T

V

W