

# **Capacity Management in the *Real* World**

**Palindrome Corporation**

**A Seagate Software Company**

April 20, 1995



## Introduction

Palindrome™ Corporation has been developing backup and storage management software tailored specifically for the unique demands of local area networks. Since it first started business in 1989, Palindrome's unique fault tolerant technology has focused on simplifying the complex tasks of storage management through superior automation and smarter operation.

In August of 1994, Seagate Technology purchased Palindrome to build a substantial presence in the data management software market. Seagate, like thousands of users, industry analysts and the press, saw that Palindrome's technology was capable of changing the way people think about backup and storage management.

Beyond backup, the rapid growth of LAN-based data is stressing many organizations' abilities to keep up. Whether it is adding new discs fast enough, locating and retrieving prior versions of critical documents, or simply getting backups done in time -- the need for capacity management is becoming essential for many larger or more complex LANs.

Organizations have found traditional capacity management solutions -- adding new disc drives, setting user quotas for disc space, compressing or deleting files -- only produce short-term relief. Many are now looking to hierarchical storage management (HSM) to solve this problem. While HSM may offer an effective means of controlling on-line storage growth, major investments in hardware and software make most HSM implementations impractical for all but the largest networks.

Palindrome™ Storage Manager advances HSM's usability by providing a step-by-step implementation of HSM functionality, fully integrated with backup and restore. Specific functionality included in Palindrome Storage Manager is:

- Backup and restore
- Migration
- Easy data retrieval
- Archiving
- Hierarchical storage management

Palindrome Storage Manager integrates the award winning backup and restore capabilities of Palindrome™ Backup Director for unparalleled data protection. And, since Storage Manager is designed around Palindrome's unique SafeStore™ technology, it's built to thrive in the "real world" by preventing, detecting, and recovering from system, user, hardware and media failures. (*See companion document "Backup in the Real World"*)

The first section of this document compares Storage Manager's step-by-step approach with manual strategies. The second section compares Storage Manager's fully integrated HSM features with higher cost stand-alone HSM systems.

We believe the following examples will illustrate how Palindrome Storage Manager provides a cost-effective solution that addresses "real world" capacity management issues better than *any* alternative solution available.

# Have you *really* thought about what happens when...

**you add additional hard discs to an already overflowing server?**

## **Situation:**

You have been here before. Volumes are running out of disc space, users can't even save files to disc at times. Network broadcasts can be heard across the office shouting "Volume out of space!" Your boss wants to know what are you going to do about this problem!

## **The Conventional Approach:**

You call your network vendor and get quotes on disc drives, controllers and additional server RAM. You schedule network downtime at night or on weekends so you can come in and install the new hardware. In the mean time you scramble for a way to move or delete files to keep the network going for a few more days. You take the quote to your boss for approval and the first thing he says is, "Didn't we just do this a few months ago?"

## **Palindrome Storage Manager - the *practical* solution:**

Palindrome Storage Manager provides a permanent and cost effective solution to capacity management problems. Using the same hardware you already use for backup, Storage Manager provides flexible tools that allow you to easily identify and prioritize inactive network data and manually or automatically move it from network volumes to near-line or off-line storage. Off-site protection is ensured and on-site availability to users easy and secure.

Palindrome's Storage Analyst tool can profile your data usage characteristics and prepare a printed profile of future storage growth, projected capacity needs, impact on backup and recovery strategy and an ROI analysis of possible solutions. Storage Manager provides a long term solution to the problem that will reduce storage hardware costs, management and user time and effort and improve the performance of both network backup and disaster recovery.

**How long can you afford to let *this* continue?**

# Have you *really* thought about what happens when...

**you add hard discs and now your backups take too long!**

## **Situation:**

O.K., so you bought more hard-disc again. Those new discs are already filling up, your nightly backups don't fit on a single tape any more, and they take almost twice as long.

## **The Conventional Approach:**

You give your boss a quote for a new larger tape drive. The first thing he asks is, "if backups are taking twice as long, **how long would it take for us to recover** from a server crash?"

Maybe you need to look at *what* exactly is stored on these discs?

## **Palindrome Storage Manager - The *practical* solution:**

Storage Manager controls on-line storage growth by identifying unused files and moving them to other types of storage such as optical or tape. By removing this inactive data from your hard discs, nightly backups and disaster recovery have less work to do and complete more quickly. In the event of a disaster, users get back to work faster!

If over time you accumulated 20GB of total data and only 20% is used on a regular basis, it would be 5 times faster to restore the 4GB of data that users really need than having users wait around while you restore the entire 20GB of data, most of which they don't really need. Storage Manager would enable you to keep only that active data on hard-disk and shorten your backup and recovery time.

**What would an extra hour of downtime cost your company?**

# Have you *really* thought about what happens when...

**you manually delete files from hard disc?**

## **Situation:**

You have analyzed your disc storage and determined that over 80% of your data hasn't been accessed in over a month. Of that, nearly half hasn't been accessed in over a year! No wonder you are continually running out of disc space! Now you need to determine which files to remove and how to do it.

## **The Conventional Approach:**

You could manually search or use utilities such as NDIR to search for files of a certain age but what do you do then? You can't just delete them because even though there may be thousands of files that will *never* need to be accessed, there will certainly be hundreds that *will* be needed. The cost of losing those files or having to recreate them would be significantly higher than just keeping them on disc in the first place...they need to be permanently protected with both on and off-site storage. You could copy them first to special backup tapes but how will you maintain catalogues of these "archived" files? What will happen if someone leaves this tape in the drive, could it be accidentally overwritten?

## **Palindrome Storage Manager - The *practical* solution:**

Palindrome Storage Manager comes with powerful default migration rules that can be easily modified and new rules added down to the directory, wildcard and file level. Groups of files can be excluded from migration or given unique aging criteria and then prioritized by age, size or weighed criteria. You might, for example, be able to free enough disc capacity by migrating 3 large files rather than 300 smaller files?

Because Storage Manager provides *integrated* backup and capacity management, you can be sure that all files are protected with multiple permanent copies (including off-site) prior to being eligible for migration. Storage Manager makes user access to any version of any file easy and secure or even automatic for the user *and* administrator.

**Are willing to spend this much time and take these risks?**

# Have you *really* thought about what happens when...

**you try to use a backup system for long-term capacity management?**

## **Situation:**

Manually deleting or "archiving" files is too complicated and time consuming and you can't afford the risk of losing user data. You decide to look at automated grooming features available in many LAN backup software packages.

## **The Conventional Approach:**

Many popular backup applications have added disc grooming features and some even claim to ensure multiple copies, off-site storage and prevent accidental overwrite. The question is for how long will your data be safe? For example, one very popular LAN backup package makes 3 copies on "monthly" backup tapes prior to grooming but then asks for those tapes to be recycled a year later. If you follow its instructions you will NOT be able to retrieve groomed files after one year! Because these features were "tacked on" to existing product rather than designed into the product architecture, these solutions offer very limited rules for selecting files, often do not offer manual migration, do not leave placeholder files on disc and migrate EVERYTHING that meets migration criteria even if that means all the data on your disc and do so whether you need to free space or not!

## **Palindrome Storage Manager - The *practical* solution:**

Capacity management features such as archiving and migration were designed as an integral part of the Palindrome SafeStore technology from the start. Palindrome Storage Manager safely ensures that a configurable number of copies have been made on *permanent* sections of media prior to migrating files. One copy would be made on designated on-site "near-line" media while other copies would be made on designated off-site media to ensure protection.

Disc capacity thresholds can be set and actual disc capacity monitored on a continual basis. When a "high water mark is exceeded", Storage Manager can notify you through the SmartAlert™ system (Alert Palette, MHS, NetWare SEND or SNMP) or can be set to automatically migrate files according to the rules and priorities you have set until the volume reaches a low water mark. Only enough is migrated to maintain the desired amount of free space on disc while using the maximum capacity of your hard disc to "cache" the most frequently accessed files.

**Are you using the right tool for the job?**



# Have you *really* thought about what happens when...

**users *need* to access data that has been deleted, groomed, or migrated?**

## **Situation:**

Through deletion or grooming you have removed data that existed in a users directory. The user remembers that it existed and needs it immediately. What do they do?

## **Conventional Methods:**

Chances are the user picks up the phone and calls you, leaves you voice mail, waits, pages you, waits, and then leaves you another not so pleasant voice mail again. When you drop what you're doing and call the user back, they explain what they are looking for, where they think it was, when they think it was there and you go off and try to find it. Finding it alone is **not** an easy task with most products. Because of the database tracking used by most products designed for backup, databases grow very large very quickly (often to 20% of your total disc capacity by themselves!) and are typically purged of information more than a few months old as a manufacturers recommendation to control growth. This means you will likely have to mount old tapes one at a time, and physically search their contents until you think you have what you need.

## **Palindrome Storage Manager - The *practical* solution:**

Storage Manager makes user access easy or even automatic. Placeholder or "phantom" files may be left on disc allowing users to see *all* of there data even if it has been migrated. If enabled, users can simply access any migrated file in any application to initiate automatic recall of that file. With an on-screen pop-up, the user has full notification and control over the recall process.

In addition, Storage Manager's simple "file manager" interface gives users a secure method to search and access any file or file version that has been backed up, deleted or migrated without making a request to an operator. A single file history database and media library contains **all** available data but allows user access only within existing file-system security. Typical file history databases consume only 1-2% of the total managed volume capacity even after many years of operation. Robotic AutoLoader or jukebox devices *may* be added at any time to completely automate the media loading process for even greater automation.

**Do you have the time for this?**

## **Section II - Storage Manager vs. Other HSM Solutions**

The first section of this document contrasted conventional capacity management approaches with the flexible step-by-step approach provided by Palindrome Storage Manager. This section will compare the benefits provided with Storage Manager's fully-automated HSM features with those provided by more expensive and complicated HSM systems.

By looking at the benefits provided by either solution you will clearly see that Palindrome Storage Manager provides all of the benefits of high-end HSM solutions at a lower cost, that's easier to manage and without the compatibility concerns of having two separate systems managing different portions of your data.

# Have you *really* thought about what happens when...

you implement an automated HSM solution?

## **Situation:**

You realized the need to find a permanent, cost-effective and automated way of managing LAN storage growth. You need a solution that will automatically monitor disc capacity levels, migrate files automatically and retrieve files automatically when users access them.

## **The Competition:**

If you purchase a stand-alone HSM package, you will probably spend \$20,000 to \$50,000 on the initial software and hardware, plus you probably will need to buy more software each time you add additional servers or devices. This doesn't include the cost of network backup software and hardware. You may have to purchase a dedicated HSM server and hard-discs or at least purchased additional backup devices just to protect your HSM near-line storage.

HSM only changes the *location* of data and in virtually all cases, near-line storage is NOT being backed up along with your hard-discs (if it was you would lose one of the major benefits of HSM...reduced backup and recovery time). This means that you need separate hardware and procedures to backup and protect the data stored on near-line storage in the event fire, flood, theft, etc. These procedures must be carefully synchronized with those protecting your on-line data including the media rotation and off-site storage.

## **Palindrome Storage Manager...the *practical* approach:**

Palindrome Storage Manager provides a step by step approach to solving capacity management problems. A full range of capacity management benefits including HSM are included right in the product along with powerful backup and recovery. You can begin with on-demand migration using only a single tape drive shared with backup and be sure that migrated files are protected, including off-site storage, prior to being migrated.

Because all storage management functions are controlled and managed from a single console with shared configuration, databases and a common media library and rotation, overall administrator time is reduced and simplified. Compatibility concerns are non-existent because a single system is responsible for managing and protecting all data.

**Didn't you buy this to *simplify* management and save money?**

# Have you *really* thought about what happens when...

**you try to do a full backup of data migrated by an HSM system?**

## **Situation:**

You have added fully-automated HSM to your network. This system begins removing inactive data from disc and replacing it with placeholders. The question is, what happens when you try to back up these files?

## **The Competition:**

Most HSM systems make provisions to allow backup applications to access placeholder files without causing the file to be recalled. In the best of cases, these systems work by "filtering" recall requests from a particular user id used only by the backup application, a particular NLM name or a particular user connection. Others actually require that the HSM system shut itself down at a particular time while the backup application is running.

This *usually* works fine until something changes and the HSM system is not updated to reflect the change. A new version of backup software uses different NLMs or user id's, backups run longer and continue even once the HSM system reactivates itself. These types of changes may cause a nightly backup process to begin recalling *each and every file* from near-line storage which may very quickly fill up the hard disc.

## **Palindrome Storage Manager - The *integrated* solution:**

Because Palindrome Storage Manager is a completely integrated system, backup and HSM operations are not only guaranteed to coexist under all circumstances, but they actually work together to provide the best efficiency. For example, a single nightly directory scan will not only backup on-line files but pre-stage inactive data to near-line storage reducing the workload on the network. A single database and media library ensures that all files are adequately protected with backup and off-site storage prior to being removed from disc. All files and versions stored on either near-line or backup media are easily accessible from a single screen so there is no wondering about which files are stored where or how to access them. *No two separate systems can make these claims, even if developed by the same company.*

**Is this worth the risk?**

# Have you *really* thought about what happens when...

**you need to recover *all* of your data after a disaster?**

## **Situation:**

Your facility was destroyed in a fire and you need to recover your data from off-site media as quickly as possible.

## **The Competition:**

With separate HSM systems you need to have off-site backups of not only your on-line data (which contains active files and placeholders) but all near-line data as well. You need to regenerate the network operating system and install appropriate backup and HSM software and hardware which may include an optical jukebox, dedicated HSM server and dedicated hard discs.

Each set of off-site backups must be in sync to the day to avoid having data falling into a "black hole". If your backup is restored as of yesterday and HSM restored as of one week ago, any files migrated in the last week will be restored to disc as placeholders but will not correspond to data stored by the HSM system. This data may be permanently lost.

## **Palindrome Storage Manager - The *integrated* solution:**

Palindrome Storage Manager greatly reduces the effort in restoring a full server and providing users the fastest possible access to both on-line and near-line data. Server recovery diskettes can be produced to manage the entire process for you -- no need to completely reinstall NetWare or Storage Manager to recover your server. Access to *all* data is provided in through a single recovery procedure.

Because Storage Manager completely integrates backup and capacity management, a single off-site media library contains all data required to recover on-line hard discs and near-line data. As a matter of fact, you can even manually or automatically recall files from off-site tape backups even if the files were originally migrated to optical media...Storage Manager knows about every copy and version of every file...you don't need to have exact replacement equipment to provide access to near-line data. Even a single tape drive will do!

If you do have duplicate optical or tape hardware, you can benefit from Storage Manager's media duplication facility to make exact image duplicates of all your near-line media, simply load these discs or tapes into a new device and provide immediate access!

## **What would *you* want in a crisis situation?**

# Have you *really* thought about what happens when...

**you fill your near-line storage with migrated data?**

## **Situation:**

Your on-line storage growth is under control and you have been migrating inactive files for some time now. All of a sudden, there is no more room in near-line storage. Do you need to add more? Can you add more?

## **The Competition:**

Different products vary greatly on the options but virtually all require you either to add more storage, and often pay for software support for new devices, or destroy existing data to make room for more. Even worse, although you can destroy near-line media and remove placeholders from disk, your on-line backup system will retain copies of these placeholders on tape for some time and let you restore placeholders to disc that no longer correspond to data available from the HSM system.

## **Palindrome Storage Manager - The *practical* solution:**

Palindrome Storage Manager supports true off-line storage with a sophisticated queuing system shared by the Palindrome File Manager user interface. When near-line storage reaches capacity, you have the option of adding more near-line storage without additional software cost or simply removing media from the drive, autoloader or jukebox and replacing it with blank media to store new data.

Users will still see all data in it's original location and can access any file as usual. If users request a file that is off-line they will receive a friendly notification that the file they requested could not be retrieved without an operator and notification will be sent through Storage Manager's SmartAlert system to the Alert Palette, SNMP, NetWare SEND or MHS E-Mail to designated administrators or tape operators. The job is held in a queue for processing at the operators convenience. The operator simply mounts the requested media and the recall or restore operation will process and the user may be notified when it is complete.

This sophisticated handling of off-line storage and support for "daisy chaining" of devices allows for virtually infinite storage growth.

**Do you want your future growth limited?**

## **Conclusion**

Palindrome's software beats the competition in critical "real world" scenarios with higher performance, lower cost and better reliability. Unique fault tolerant technology provides the intelligence you need to effectively manage and protect your critical data in the real world. That's why thousands of users are buying Palindrome software to protect their businesses' critical data.

Palindrome's technology is changing the way people think about backup and capacity management and is the choice of those in the know -- like, Novell, Hewlett-Packard, Hitachi, and many others.