

Backup in the *Real* World

Palindrome Corporation
A Seagate Software Company

Introduction

Palindrome Corporation has been developing backup and storage management software tailored specifically for the unique demands of local area networks. Since first starting 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.

LAN backup products deployed in the field run into an astonishing array of system, user, hardware and media failures, most of which are rarely seen in the lab when the product is being developed. The single most important variable in real-world reliability is how well the backup software prevents, detects or recovers from these errors. A well-designed and implemented LAN backup application can prevent many user errors by limiting dependency on user intervention, and by clearly guiding the user through operations. System, media and hardware failures can be overcome if the software "understands" concepts of redundant protection, monitors and alerts the user of a pending failure, evaluates the extent of the failure, and can take corrective action.

Palindrome's products utilize fault tolerant technology to accommodate these real world situations. Today, Palindrome customers are finding real benefits in this technology through a reduction in overall storage management costs and increased reliability. Simply put, Palindrome software is designed to keep working when things break - you could say it has been *built for the real world*.

We believe the following examples will illustrate these points and clearly show you why Palindrome's fault tolerant technology is changing the way people think about backup.

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

the backup operator leaves the wrong tape in the drive for tonight's backup?

Situation:

Today is Friday. Tape X was used for an incremental backup last night. Tape Y is expected for a full backup tonight. The backup operator either forgets to change tapes (leaving tape X in the tape drive), or inadvertently puts tape Z in the drive. Either way, the operator discovers the mistake Monday night and puts tape Y in the tape drive.

The Competition:

Most other backup products expect to do a full backup on tape Y. When the software sees tape X in the drive, it posts an error in the log and stops. No backup happens (a server crash over the weekend would lose all data created or modified on Friday). When the correct tape (tape Y) is placed in the drive on Monday you may or may not get the full backup you expected and needed. Some products will skip the full backup entirely and continue with the regularly scheduled operation for that night.

Palindrome Backup Director - The Right Way:

Palindrome Backup Director expects to do a full backup on Friday night but notices the incorrect tape in the drive. With tape X, it does one more incremental (or differential) backup operation, and posts an error to the log that a tape rotation is past due. With tape Y in the drive on Monday, it does the scheduled full backup (to "catch up"), and adjusts the rotation schedule to accommodate. Either way, when tape Y comes back, the delayed full backup is done.

Are you comfortable with how your solution handles this situation?

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

your full backup tape is damaged mid-week?

Situation:

You do full backups every Friday night, differential or incremental backups the other nights. Through an accident, last Friday's full backup tape is damaged during the day Tuesday. As a result, this week's differentials/incrementals won't be sufficient to do a full restore.

The Competition:

With competing products, all the burden is left on the operator to ensure they are fully protected. The rotation schedule (if one exists) must be suspended, and a full backup must be performed to another piece of media. This is required to ensure proper volume recovery since it is based on the last full backup plus any incremental or differential tapes. Competing products assume that the operator understands all the implications of the lost tape, and will take manual corrective action to ensure adequate protection again.

Palindrome Backup Director - The Right Way:

With Palindrome Backup Director, the operator simply instructs the software to "forget" the damaged tape. Backup Director then analyzes what data was on that tape via its powerful on-line database, and the next regularly scheduled backup operation will automatically ensure that these files are reprotected. Other than telling the software to forget the damaged tape, the operator doesn't need to do a thing. Simply tell Backup Director what happened, and it assures full protection.

Are you comfortable with how your solution handles this situation?

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

your backup device and/or media fail during a backup?

Situation:

A large server has four DAT drives, and is backing up 10 volumes on 3 servers, along with 20 workstation disks. During the night, DAT drive #2 fails with a jammed tape.

The Competition:

With competing products that allow parallel operations, volumes are statistically mapped to backup devices. All volumes that were to be protected by the failed drive (other than those completed before the failure) will not be backed up. If the drive actually failed, the operator will need to reconfigure their system in the morning to handle the next backup, or replace the drive before the next backup.

You should also note that some competing products do not support the parallel use of multiple drives at all and when the single device fails there is no way for the software to perform the necessary backup.

Palindrome Backup Director - The Right Way:

Because of Backup Director's dynamic parallelism, all remaining backup sessions (for the 30 total volumes) are automatically routed to the surviving three drives, without any human intervention, including the partially completed job. The backup completes successfully and the administrator is notified of the drive failure.

Are you comfortable with how your solution handles this situation?

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

the amount of data on your network changes and the impact it has on backup performance?

Situation:

Four DAT drives are installed on one large server backing up its volumes, those of two other servers, and 20 workstations. Over the following six months, the amount of data on these volumes changes, the size of the volumes change, and night by night, the speed of data transfer from them varies. Will all the volumes be backed up by morning as the network changes and grows?

The Competition:

With competing products, the manager must assign volumes to specific devices at installation, taking into account the amount of data on each volume, and the speed of data transfer from that volume. For example, the manager would have to assume a different transfer rate for the volumes on a local server versus volumes on remote servers. If done well, the first backup will keep most of the drives busy. As time goes on and data distribution changes, some of the tape drives will finish early in the night, while others will still be running well into the morning, requiring frequent reevaluation and reconfiguration of the drive assignments. If one volume or server is unusually slow one night, other volumes assigned to the same tape drive will be waiting for backup well into the morning, even though the other three drives were idle much of the night.

Palindrome Backup Director - The Right Way:

Backup Director's dynamic parallelism provides totally automatic load balancing. Jobs are automatically assigned to devices as the previous jobs are completed. Changes in data transfer rate, or long-term changes in volume size, are automatically distributed among the devices, keeping all devices active until there are no jobs remaining.

Are you comfortable with how your solution handles this situation?

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

you have to restore an entire server or volume?

Situation:

You have a five (5) gigabyte server that has crashed and has to be completely recovered from tape. Three (3) gigabytes haven't changed since the last full backup, two (2) gigabytes have changed and need to be restored from the differential tape.

The Competition:

With competing products, all data from the full backup tape is restored, wasting considerable time restoring 2 gigabytes of data that is obsolete. Two gigabytes are then restored from the differential tape. A total of 7 gigabytes are transferred, consuming substantial time in a pressure-filled full restore. The situation is even more dramatic with incremental backups, where some files are restored as many as five times when only one copy is needed.

Palindrome Backup Director - The Right Way:

Since Backup Director approaches recovery completely from a file orientation it begins the restore from the full backup tape, restoring **ONLY** the 3 gigabytes of data required, not wasting time on the 2 gigabytes that is obsolete. The remaining 2 gigabytes are restored from the differential tape.

Are you comfortable with how your solution handles this situation?

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

a full server or volume recovery is interrupted by a server crash or other event?

Situation:

A conflict in one of your servers caused a crash, corrupting the server's file system in the process. A full restore of 5 GB is necessary. Two hours and four gigabytes into the restore, a configuration error in the new server causes an abend. The data isn't lost, but the restore process is interrupted.

The Competition:

With competing products, the manager restarts the restore, overwriting all the files that are already on disk. All the time spent before the abend was wasted.

If your restore operation requires a tape that was not immediately available, competing products require you to stop the restore and restart it later - starting from the beginning. With these products, tapes must be used in the exact order requested by the software - even if another tape is available that has data that needs to be restored.

Palindrome Backup Director - The Right Way:

When the restore is restarted, Backup Director scans the server to see what data is already restored, and resumes the restore with the final gigabyte of data. The only time you lose is the time it took to reboot the server after the abend.

Similarly, if your restore operation required a tape that was not immediately available, Backup Director allows you to use an alternative tape that contains data to be restored or provides an easy mechanism to stop the restore and restart it at the exact point you left off when the tape is available.

Are you comfortable with how your solution handles this situation?

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

you need to restore a file but you don't have immediate access to all your tapes?

Situation:

While working at a station on the network (not in the server room), the system manager needs to restore a deleted file. Using the backup software, he selects a file and submits the job for automatic restore. The file to be restored has been backed up on tapes W, X, Y and Z.

The Competition:

With competing products, he sees that there is a copy on tapes W, X, Y and Z. He selects the copy on tape W. The server engine receives the request, finds that tape W isn't in the autoloader, and stops. Five minutes later, the request is dropped. The identical file on tape X is not used.

Unfortunately, the user didn't know which tapes were in the drive. This is important since competing products represent specific file copies on specific media in their user interfaces. That is, you must select the exact file COPY you want to restore.

Palindrome Backup Director - The Right Way:

With Backup Director, he sees that there are four copies of the file VERSION he needs. He tags that version, and submits the job for processing. The server-based engine receives the request. Tapes W and Y are off site, tape X is in the autoloader, and tape Z is on the shelf. Backup Director automatically mounts tape X and restores the file - notifying the user upon completion.

Are you comfortable with how your solution handles this situation?

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

you try to simplify your workload by allowing users to restore their own files?

Situation:

An autoloader has been installed to minimize operator intervention in backups and restores. This makes it possible for end users to restore not only yesterday's file versions, but a good selection of previous versions. To provide better service to the users, and to reduce the burden on the IS group, users are given the ability to automatically restore their own files through the backup application, without IS intervention.

The Competition:

With some competing products, the user must be given rights to the entire backup database, allowing every user on the network to restore any file to their directories, totally bypassing network security. Since this blatant disregard of security is not acceptable in most organizations, all restore requests must be routed through the network manager. Other products only allow users to restore data from jobs they actually defined.

Palindrome Backup Director - The Right Way:

With Backup Director, each user gets a simple utility, very similar in operation to the File Manager. The user sees a simple directory tree, showing files on disk and deleted files that are available on tape. Each user sees only those files that he has rights to see under the network security system. After tagging the file version he wants restored, the user can monitor the restore requests progress, and can even request e-mail notification when the restore completes.

Are you comfortable with how your solution handles this situation?

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

you request to have a file restored but the tape is not available?

Situation:

A user or the system manager has requested a file restore, but the file does not exist on the media in the drive or autoloader.

The Competition:

With competing products, the job is suspended, and a request for the required media is displayed on the server. If the correct media is not mounted within five minutes, the job is discarded.

Palindrome Backup Director - The Right Way:

With Backup Director, the job is processed, and placed on "server hold". The user sees that the job was suspended. An alert is sent to the network manager, so that she knows a media change is needed. When the network manager is able to respond, the queued jobs are handled with manual media mounts. As each restore happens, the originating user receives an e-mail telling them that the file has been restored.

Are you comfortable with how your solution handles this situation?

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

you utilize products that do not comply with industry standards to protect your data?

Situation:

You've been diligently backing up your data for the past year. Your company has just been purchased by a company that utilizes a different backup product. They have a corporate policy that requires interchange of backup media for disaster recovery purposes.

The Competition:

Most competing products continue to utilize proprietary tape formats with their products. Their position is that the use of standards will allow competing products to enter their captive markets. Even those that pledge future support of industry standard tape formats will do so only as a secondary format or will only support the ability to read industry standard tapes. In either case, without complete support for these standards, they are locking their customers into their proprietary formats.

Palindrome Backup Director - The Right Way:

Palindrome's products fully support the System Independent Data Format (SIDF) standard which is being standardized by the international standards organization ECMA as ECMA-208. Palindrome supports SIDF as its native tape format, ensuring that all data written by Palindrome's NLM products are written in an industry standard format.

Palindrome's full support of SIDF ensures its customers of long-term data retrieval. This is becoming increasingly important as more and more customers begin to utilize archiving and hierarchical storage management solutions (Palindrome's archiving and HSM products also utilize SIDF).

Are you comfortable with how your solution handles this situation?

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

you recycle a weekly backup tape from a month ago?

Situation:

Most backup strategies rely on periodic full system backups plus series of incremental or differential backups to protect server data.

The Competition:

With competing products, these weekly full system backups are usually recycled and erased in a matter of weeks requiring all requiring each and every file on the network to be copied back to tape to complete a new full backup. Chances are that 70-80% of the data just copied was already on the tape before it was overwritten and hadn't changed since that tape was written the last time.

Some companies are forced to preserve each and every weekly and/or monthly full backup tape in order to insure that records of major projects are available on tape with a variety of historical versions. This causes the media library to become very large and tape catalogs on disk to grow very quickly to the point that many backup software developers recommend or even *default* to purging records from the catalog after a few months. It is not uncommon with many competing products for tape catalogs to consume up to 20% of the total disk space being protected. 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:

Through the implementation of Dynamic Baselining™, Storage Manager creates and maintains a baseline of stable network data automatically. During backups, Storage Manager recognizes these archived files, backing up only those files that are active and changing. This means backups can be as much as three times faster than competing products and need much less media to complete. And, most importantly, the full backup sessions mandated by every other backup system become strictly optional -- allowing one system to protect and manage a much larger portion of the network.

Because Storage Manager Archives all stable data, the system creates and manages an on-line catalog of historical file versions. Storage Manager makes it easy for any user to retrieve not only the most recent version of a file, but any version catalogued by the system -- simply and automatically, right from their workstation. Typical file history databases consume only 1-2% of the total managed volume capacity even after many years of operation so you can keep records of every file and every version on tape on-line and accessible at all times.

Conclusion

Palindrome's software beats the competition in critical *real* world scenarios with higher performance, lower cost and better reliability. Our 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 is the choice of those in the know -- like, Novell, Hewlett-Packard, Hitachi, and many others.
