

WHITE PAPER

Comparison of Tape Backups With Terremark Backup and Restore

EXECUTIVE SUMMARY

The need for reliable data backup has never been greater. Recent studies have shown that 50% of companies that experience total data loss are out of business within one year and 90% within two years. Data backup is an insurance policy that no business can afford to be without.

Tape backup systems have been the predominant method of data backup for 40 years. Like most 40-year-old technologies, tape is outmoded for data backup and restore. According to the research firm The Yankee Group, 40% of tape backups fail while industry analyst Baroudi Bloor reports that 50% of tape restores fail. Thus, disk-based backup systems, especially off-site, online backup system such as the Terremark Backup and Restore service, is increasing. According to the Gardner Group, by 2008, expenditures on disk-based backup will be greater than expenditures on tape backup.

This white paper explores the reasons why tape is being replaced for ongoing data backup operations by online, offsite backup systems such as Terremark Backup and Restore.

TAPE IS UNRELIABLE

When a single byte of information is written onto tape, a thin piece of iron oxide-coated plastic ribbon is stretched, pulled, prodded and magnetically imprinted. Repeat this process a million times in a year and it is apparent why in a recent Storage Magazine survey, over 60% of IT professionals experience tape failure at least once per week and 25% experience tape failure at least twice a week. The problem is exasperated as tapes are handled, dropped, and exposed to adverse environmental conditions.

Recently, a technician for the State of Alaska accidentally deleted all pertinent data for a \$38 billion oil-funded account. As the backup tapes were unreadable, the state was forced to spend over \$200,000 recreating and recovering the data. While this was a well-publicized case, the scenario is repeated over and over across the corporate landscape.

Getting data across the tape heads and onto the tape media is a time consuming process, resulting in tape backup jobs not finishing during their backup windows. Consequently, many businesses have stopped performing data integrity checks on their backup jobs because that increases the time to completion by 20-30%. A recent study by the research firm IDC showed that less than 40% of businesses actually test their backups on a regular basis.

Terremark Backup and Restore uses highly reliable RAID-5 storage for its data storage. If RAID-5 experiences a rare hardware issue, the system pinpoints errors and allows technicians to replace bad drives without data corruption.

Further, the Terremark Backup and Restore system continuously runs a self-healing process that checks for data errors on the disk and repairs bad clusters on the fly. Terremark's customers can instantly initiate a separate data validation process or test restores at any time to ensure the integrity of their backup data. With tape, this would mean hours, if not days, of retrieving tapes, mounting them, indexing them and restoring them.

Writing data to disks is inherently much faster than writing to tape. Due to data reduction schemes such as compression, common file elimination and delta-blocking (where only those portions of a file that change are stored), Terremark Backup and Restore writes less data than tape to complete backup jobs well within allotted windows.

TAPE BACKUP IS COSTLIER

For many IT professionals, it comes as a shock that tape is more expensive than disk for data backup. After all, a single tape costs \$15 and holds 1+ TB of data (uncompressed) while a 1 TB hard disk costs around \$90. However, the Total Cost of Ownership (TCO) for using tape is much higher than for Terremark Backup and Restore because a TCO analysis takes into account the following factors:

1. Hardware costs
2. Software costs
3. Media costs
4. Human factors costs
5. Offsite storage costs

Taking into account the above factors, the total cost of ownership for Terremark Backup and Restore is usually less than for tape over a given 3 or 4 year timeframe (the time over which tape hardware costs are amortized).

Terremark can easily perform a full TCO analysis for any corporation upon request that will show how much money will be saved using Terremark Backup and Restore. One Terremark customer, a government contractor with 2 TB of data under management, saves \$14,000 per month by switching to Terremark from tape. Each of these costs factors is discussed in detail in Appendix A.

TAPE PROVIDES INCOMPLETE COVERAGE

Many businesses have IT assets beyond the reach of traditional tape infrastructure. Business travelers with laptops perform work at customer sites, on planes and in hotels. This data is vulnerable until the laptop is back in the corporate LAN (or VPN), which sometimes can be days or weeks. Furthermore, if the user doesn't copy the data to a file server, it will most likely never be backed up. Why is this important? According to Gardner, up to 20% of laptops suffer hardware failure within the first three years. Further, the FBI reports that nearly \$7 million worth of laptops were stolen in 2004. However, IDC reports that less than 40% of the laptop data is backed up.

Terremark Backup and Restore allows each laptop to independently back up to the Terremark Data Center using any available Internet connection (corporate LAN, home DSL/Cable or hotel broadband) without user intervention. While daily backup operations are decentralized, management can be centralized, if required, via reporting functions that provide the IT manager with insight across all laptops in a corporate enterprise.

Remote offices present another challenge for centralized tape backup systems. Businesses either forego remote office backups or install smaller tape backup units at each location. These tape backup systems lead to increased costs across all five cost dimensions mentioned earlier. Logistics also become a nightmare as non-IT personnel at the remote sites start swapping tapes and checking backup logs.

Terremark solves this problem by installing a gateway at each site and letting it perform all backups in an automated fashion. Terremark even monitors the jobs, although a central management console at a central site (e.g. headquarters) can also be used by the IT manager.

TAPE MAKES YOUR DATA VULNERABLE

Tape backups only reduce your DR exposure only when they are carried offsite. Whether the tape is being carried by an employee or by a professional tape handling server, there is a certain amount of time that the data is vulnerable until the tape is physically in a protected facility.

If the backup tapes fall into the wrong hands, a company's sensitive data can be compromised as recent studies have shown that only 7% of companies encrypt the data. Tapes can and do get lost in transit to offsite locations. Recent examples of backup tapes being lost in transit to or from an offsite facility include:

- Bank of America: 1,200,000 customers
- Ameritrade Holdings: 200,000 current and past customers
- Citigroup: 3,900,000 customers
- MasterCard: 13,900,000 customers
- Time Warner: 600,000 customers

For many organizations, recent legislation has made it mandatory that all backup data that leaves the organization's site be encrypted. Two examples are those organizations subject to HIPPA (health care records) and those subject to PCI (credit card records). Using tapes is simply not an option anymore for many organizations.

Terremark Backup and Restore encrypts all data before it leaves the gateway appliance that is within the corporate enterprise (i.e. behind the firewall). While the latest version of all data is stored locally on this gateway, all data is encrypted and automatically copied to the data vault in Virginia and then replicated to Dallas. Terremark encryption can be as secure as AES256, the government standard for the encryption of classified data.

The software agents for tape backup software also are a major security risk. While Veritas and other manufacturers have issued code patches, business-software maker Quallsys reports that administrators take an average of 62 days

to implement fixes, a long time to be vulnerable. Terremark Backup and Restore addresses this issue by not deploying agents.

TAPE DECREASES YOUR REVENUE STREAM

A company's IT staff is a very valuable resource. IT professionals are highly trained, highly motivated people who usually have not gone to school to swap tapes and check endless backup tape logs. After a while, these highly trained people begin to make mistakes during the backup process. According to Storage Magazine, human error was responsible for tape backup errors at least 50% of the time.

Most business want to free up their IT professionals to do what they do best: help support the daily operations of end users. By using a "set it and forget it" automated system such as Terremark Backup and Restore, IT personnel gain back free time to help end users and implement any long overdue time- and money-saving projects.

Companies also see the use of Terremark Backup and Restore as a competitive edge when bidding on highly-competitive government and corporate contracts. By promising their customers that project data will be kept safe and secure, businesses in many verticals are finding that Terremark Backup and Restore is a great proposal-stuffer. One government contractor has recently implemented a policy of including Terremark Backup and Restore on all current and future proposals as a method of guarantee data integrity to their clients.

CONCLUSION

The use of tape is being shifted to the narrow corridor of long term archiving while day to day backup operations are being migrated from tape-based initiatives to online, offsite systems. Terremark Backup and Restore has been leading the way in this arena for many years. This paper has addressed several areas of concern relative to tape backup and how Terremark addresses those issues. The final conclusion is that tape is more expensive and doesn't protect data as well as Terremark Backup and Restore. Contact Terremark today for begin your free trial and find out how much easier data backup can be.

APPENDIX A: DETAILS OF COST FACTORS

Hardware Costs

Tape backup systems can run into the hundreds of thousands of dollars for the initial capital outlay. Beyond that, maintenance is typically around 20% of the initial cost that is recurring every year for the life of the hardware. Scalability also becomes a concern because once a company's data amount grows past the hardware's ability to back it up, the company is faced with another large capital expenditure. Conversely, a company can over-commit resources to anticipate future data growth but then is wasting hardware capability until such time that the data has grown to fully utilize the hardware.

With Terremark, there are no hardware or hardware maintenance charges to pay. Scalability is never an issue as the Terremark data center uses the latest SAN and NAS capabilities to provide virtually unlimited data storage.

An often overlooked source of hardware costs involves the use of legacy systems. Most companies have data retention policies that span many years, meaning that data that was written to tape several years ago must be accessible. However, the hardware that was used several years ago to write the tape was probably replaced or upgraded at least once. After each upgrade, the old tape systems must be maintained in case data needs to be restored from the older tapes. Maintaining multiple tape systems is an expense that is not incurred with Terremark Backup and Restore. Terremark updates its hardware periodically with no impact on the customer.

Software Costs

Software costs for tape backup begin with the cost of the tape backup software. Then, there is the purchase of software licenses: core tape system, all mail, database and file servers, and every end user desktop or laptop. Enterprise class software will typically cost a SMB thousands of dollars and a large company tens of thousands. Finally, there is the maintenance charge, 20% recurring every year.

There are two additional hidden costs of software licensing. First, if the IT department loses the piece of paper that lists the software license codes, that license needs to be paid for again. Second, the company needs to pay for everything again in 2-4 years when Microsoft releases a new OS and drops support for the old one.

All in all, it's a cycle of recurring and non-recurring costs that never ends until the switch is made to an online, offsite backup service. Terremark Backup and Restore does not have any licensing costs and makes the software freely available to all critical and non-critical components of a company's network: file, database and mail servers, laptops and desktops.

Media Costs

While the cost of physical media is coming down, the sheer number of tapes that a company has to purchase in a given year adds up in a hurry. Typically, a company will have one 6-tape set for dailies (used every week for 3 months), one 5-tape set for weeklies, and one 12-tape set for monthlies. The number of tapes multiplies quickly if older, lower capacity technology is used and multiple tapes are required for each backup.

As an example, consider a business with 600 GB of data with 120 GB (20%) changing on a nightly basis. With each DLT tape holding 40 GB (uncompressed), the number of tapes in a given year can be calculated as follows:

1. Dailies: 3 tapes daily x 6 days = 18 tapes. 4 sets per year = 72 tapes
2. Weeklies: 15 tapes weekly x 5 weeks = 75. 4 sets per year = 300 tapes
3. Monthlies: 15 tapes monthly x 12 months = 180 tapes

Total media cost = 552 tapes x \$40 per tape = \$22,080

Human Factors Costs

Human factors involve all costs with purchasing, installation, configuring and maintaining a tape-based infrastructure. Specific costs items include:

1. Investigating the vast array of tape options. Performing a Google search on the term “tape backup system” results in 3-5 ‘above the results’ hits, 8-10 ‘right side’ hits, and thousands of search hits.
2. Training and configuration. Usually, this involves expensive on-site consultants that charge by the hour for training, implementation and certification.
3. Managing the backups. IT personnel must perform the mundane tasks of swapping tapes, checking backup logs, and reacting to errors. Many businesses cannot devote the necessary time to the latter 2 tasks and therefore, subject their backup data to failure due to overlooked errors.
4. Managing off-site media storage. This involves either managing a vendor or physically carrying tape to a safe, secure location.
5. Supporting restore requests. When an end user requests data to be restored, the proper tape must be determined, retrieved from an offsite location, mounted, indexed and searched before the first file can be restored. Usually, this sequence of steps takes many hours, sometimes days if the tape is coming from the IT director’s garage.

Many businesses underestimate the costs of human factors because an IT engineer’s salary is usually only 50% of the true cost to the company. The company still has to pay for the employee’s benefits, time spent out of the office, overhead, etc. Therefore, an engineer whose salary is \$40,000 per year (approximately \$20/hr) actually costs the company \$80,000 per year (approximately \$40/hr).

Off-site Storage Costs

Engaging a tape handling vendor to transport tapes to a secure off-site location is an expensive, non-deterministic process. It is non-deterministic because each restore request involves a fee and a company cannot predict how many restores will be required.

Specific cost items include:

1. Monthly maintenance fee. This is a flat fee that is charged before the first tape is picked up and stored.
2. Basic pickup and storage charges. These are usually on a per-tape basis and increase as the number of tapes increase to meet a company’s growing data needs.
3. Basic retrieval charge. The cost to locate and deliver one or more tapes, usually within 24 hours
4. Emergency retrieval charge. The cost to locate and deliver tapes with a faster turnaround time (e.g. 4 hours).