NEXSAN

MAKING CENTS OF TAPE VS. DISK
THE COMPLETE GUIDE ON WHAT TO CHOOSE AND WHY

Since the cost comparison between tape and disk is obviously far more complex than the media itself, this brief outlines all the considerations necessary for a more complete understanding of the benefit and true cost comparison between tape and disk to help IT professionals choose and justify their backup environment.

WHITE PAPER

NEXSAN

VOLUME OF DATA
PAST AND FUTURE

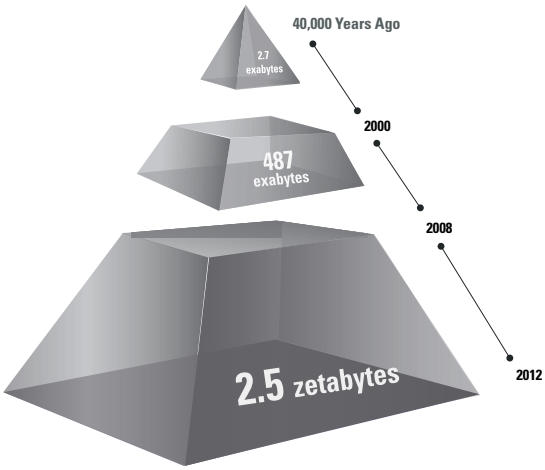

The rate of data growth simply exceeds the capability to store and adequately protect data on aging storage media, which is why magnetic tape is in a serious state of atrophy.

Perhaps the most interesting thing about the tape verses disk argument… is that there is still an argument. Most storage mediums transition from old to new without even so much as a fuss as the benefit of new media is clearly understood. However, it seems that tape verses disk continues to be the anomaly. Whereas, many IT professionals see a clear and distinct winner, others do not.

On one hand, there is the high speed, highly reliable, random access, online benefit of disk. On the other hand, there is the low speed, low reliability, linear access, offline nature of tape. One is new, fast and flashy. The other is old, slow and cumbersome. The winner is clear, right?

Although the choice may seem clear when looking at the benefit comparison, IT professionals who choose tape for their backup environment will end up citing a couple things like, "tape is still fast enough to meet their window," or, "their organization can handle extended periods of downtime while waiting on a restore." However, the most commonly used answer for a tape deployment over disk is that the sheer expense of tape is simply, "cheaper." With the performance, management and reliability benefit clearly belonging to disk, the outstanding issue seems to be a "perceived" cost issue.

When making a direct cost comparison of media, it is true that the cost-per-byte is cheaper for tape than it is for disk. For some IT professionals, that's where they draw the line and make a decision. For them, the cost of media is the race, and tape is the winner. The problem, however, has to do with the fact that the race really isn't about the cost of media; it's the associated cost of several other factors: downtime, reliability, management, availability, data growth and the cost of the backup system itself. In other words, it's about the big picture.

It should be noted that some IT professionals have circumvented the whole tape vs. disk decision dilemma and have implemented tiered solutions that use both in concert, otherwise known as Disk-to-Disk-to-Tape (D2D2T). With this approach, IT professionals are writing directly to a disk array for their backup then passing it on to tape for deep archiving and off-site portability. With this approach, organizations are leveraging the many benefits of online disk storage while maintaining the portability and long term retention aspects they are used to receiving with tape.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com

"

The best case scenario for tape reliability is still a data center's worst operational risk.

NEXSAN

TAPE VS. DISK

The Big Picture

1. Reliability
2. Performance
3. Management
4. Data Availability
5. Power Efficiency

Since the cost comparison between tape and disk is obviously far more complex than the media itself, this brief outlines all the considerations necessary for a more complete understanding of the benefit and true cost comparison between tape and disk to help IT professionals choose and justify their backup environment. Whereas the decision was more controversial just a few short years ago, it has never been clearer and easier to understand than now with new technologies, capacities and market prices.

TAPE VS. DISK – RELIABILITY

The perceived value of tape is to offer adequate data protection at low cost. Gartner estimates that one in ten recovery images on tape is unrecoverable from the large enterprise point of view. Figures are much worse for those without an enterprise class library and transports. Gartner has gone on to say that data up to one year old has a 10-15% failure rate, and the failure rate of data five or more years old is 40-45%.

Other studies have revealed that much of this goes unnoticed as Gartner and Storage Magazine reported that 34% of companies, who backup their data to tape, never test their backups. They went on to say that 77% of those companies, who did test their backups, found backup failures.

If the above failure rates with tape are the best case scenario for a data center, that means it is still far more than an organization can afford, especially with 7 out 10 small firms going out of business from a major data loss[1]. The best case scenario for tape reliability is still a data center's worst operational risk. Companies try to mitigate that risk by depending on multiple copies of the same information on different tapes that came from their full and incremental backups. Although the cost of media might be cheaper for tape than disk, after one considers the number of copies necessary for tape attempting to achieve acceptable levels of reliability, the cost-per-byte protected FAR exceeds disk.

That's why cost comparisons shouldn't revolve around "bytes stored" but rather "bytes protected." Tape only outperforms disk on outright media costs when organizations accept the associated reliability risk and don't retain multiple copies.

[1] Source: Boston Computing Network, Data Loss Statistics

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com

77% organizations who tested their tape backup and found failures.

RECOVERY TIME
TAPE VS. DISK

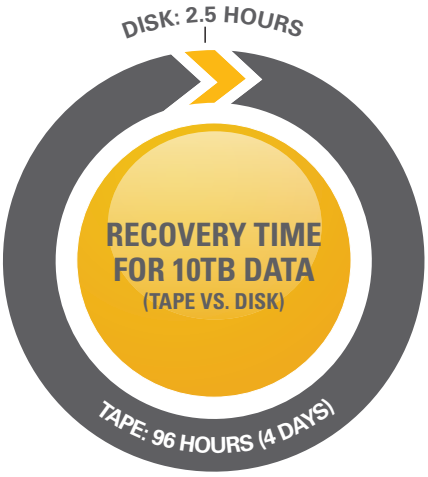

Disk, unlike tape, has a multitude of reliability and protection elements that are built-in and commonly used like RAID and automated error checking. There is no such thing as RAID with tape. If one tape out of a backup job group fails, the whole restore collapses. By utilizing RAID 6, organizations are protected against the rarest of occurrences like dual drive failures. With tape, there is no built-in redundancy.

"Reliability versus cost" is one of the key determinants to "disk versus tape." Unless a user has the latest high-end tape transport and library targeted at the enterprise, serious reliability issues are to be expected notwithstanding the complexity of tape management and risk from handling. Additionally, users need many transports and libraries to stand a chance of getting nightly backups done on time. Then there is the media. Anyone using tape understands that the cost of the media is the big expense. Every couple of years, as new transports are announced, users discover the old media will no longer work, and all of it has to be replaced. The process of replacement is not only expensive, it is disruptive.

When taking into consideration the reliability exposures and the necessary retention of multiple copies for a single byte, the cost-per-byte associated with tape is far more expensive than the base measurement. And that still doesn't take into consideration what has always been seen as the necessary evil — performance, networks, resource conflicts, scheduling, media management and more. Using disk as a library offers a flexible, ultra-reliable, high-performance and operationally efficient solution. Features from backup software vendors have made backup to disk the logical choice for simple and flexible backup and recovery.

When considering applications like VMware, Exchange and SharePoint, protection, recoverability and performance are key. While tape is still used, it is rare to see it used exclusively today. The benefits of Disk-to-Disk (D2D) are too great, which is why at least 70% of all backups are written to disk first[1].

[1] Source: ESG White Paper – Disk Backup and EMC: Addressing Today's Business Problems 2008

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com

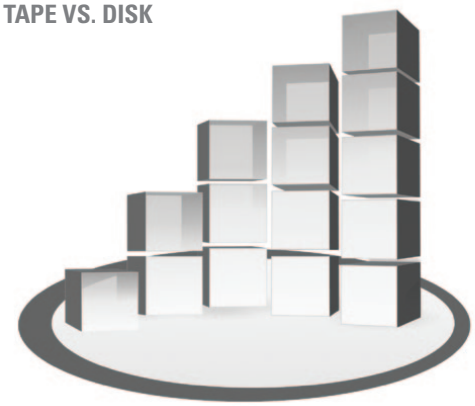

6,300 LTO2 TAPES

To protect:
42TB data/year

2TB DEDUPLICATED DISK STORAGE

To protect:
42TB data/year

TAPE VS. DISK – PERFORMANCE

For many organizations, the only cost that really matters is the business cost of downtime. The fundamental question to ask when looking at tape or disk is, "What are your recovery objectives?" After all, it's not really about the backup, it's about the restore. It has to work, and it has to work on time.

Tape is not a random access medium. Backed up data must be accessed as it was written to tape. Backup software will gather data from multiple job streams, typically 15 or more, then interleave the data into a super-block, which is then sequentially written to tape across multiple cartridges.

To recover something as small as a single file, ALL the data must be read from the super-block by stripping away 14 out of 15 records to get the one record needed. Multiply that by all the super-blocks associated with the sequential file that is often spread across 10 to 30 or more tapes, and one can understand just how slow and arduous this process can be. Aside from being slow, it is prone to failure from media and transport reliability issues that have no level of redundancy whatsoever. If there is a single "read error," the entire backup is lost.

Protection objectives are measured as Recovery Point Objectives (the amount of data at risk) and Recovery Time Objectives (the amount of downtime you can tolerate) are a concern with tape. Tape can easily take nearly four days 23 hours to recover 10TB for an LTO-4 as compared to 2.5 hours for a high performance disk array used as a protection library. What is the cost of downtime for you?

When using disk as a protection library, the problem is solved. Backup software such as Veritas's Netbackup can index all the data as it is stored directly to disk. If there is a need to recover a single sub-object to VMware, an email message to Exchange or a SharePoint document, users can recall them individually with a simple "point and click." Due to the nature of random access recovery, performance is unparalleled. Use the NDMP protocol to write directly to disk for easy configuration and management of network based backups. With NDMP, network congestion is minimized because the data path and control path are separated.

With disk used as a protection library, backup can occur locally — from file servers direct to disk — while management can occur from a central location. Operation is simple because it is indexed by the backup application directly to disk. The decreased infrastructure complexity makes everything easier and far more operationally efficient. With disk, backup is faster to restore and much easier to manage.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com

"10 degrees of temperature change can change the life of a tape by 10 yrs or more.

TAPE VS. DISK – MANAGEMENT

Backup to tape has always been an administrative challenge with the amount of manual intervention needed to perform backups. Tape backup must be closely supervised, equipment needs to be regularly maintained, heads have to be cleaned, tapes must be loaded, replaced, labeled and transported. While multiple tapes and monitoring are required for a single backup to tape, backup to disk is a completely automated procedure — "set and forget."

Backup to tape typically uses a Grandfather-Father-Son (GFS) managed retention plan. The backup schedule will generally include daily incremental backups and weekly full backups. Every TB of primary disk causes 25TB to be written to tape for protection.

The cost to implement, maintain and manage this level of protection can be overwhelming. Here's one example to illustrate the capital expense: if an organization were backing up 42TB of primary disk[1], they would need 788 LTO-4 tapes over the course of a year. This assumes an 80% efficiency usage for each cartridge. At $38 per cartridge, the cost is $27,563. Using GFS the cost of storing 25 copies of the data would rise to $137,812. By comparison the cost of a second 42TB array as a backup target is in the range of $45,000.

A restore of a single user or application can easily require loading and reading 10 to 30 cartridges or more. Finding the right cartridges, and having each one of them work without failure, is a major concern. The manpower required to manage a tape library is FAR beyond the manpower needed to manage disk as a protection library.

A tape library is typically a serialized resource. Backup jobs are scheduled by priority; resources are switched and allocated to a job. When that job completes, resources are switched again, and the process goes on. The associated monitoring and administration of complex processes creates heavy bandwidth on the IT department and easily leads to operational failures.

Using disk for a protection library, backups are routed through a centralized backup infrastructure while allowing users to share resources among multiple servers, simultaneously, whether it is on a SAN or through the network by way of iSCSI — no monitoring, no switching, no hassles. Backup jobs run simultaneously, avoiding the imposed requirement from tape whereby users must wait for the backup job to complete, switch resources, then begin again. With a disk array, multiple streams can run at the same time. Users can also easily collect or move data offsite on a WAN for geographically protected data.

[1] Disk Capacity Assumptions: Generally backup utilities to disk have the ability to compress and to deduplicate. Compression, as an example, would require an initial full backup. With a 2:1 compression efficiency, the initial backup will take 21 TBs. Change data rates average about 10% per week, after compression that adds 4 TBs, which becomes 2 TBs compressed each additional week. That would consume 27 TBs used for the initial month or 38 TBs for 3 months, which then would be overwritten round robin. If you instead use deduplication, also common for Backup applications, you will typically get a 20:1 efficiency that would reduce the capacity required to 2.1 TBs for the initial backup and another 1.1 TBs for the remainder or the 90 days for a total of 3.2 TBs. The example conservatively uses a target of 42TBs for financial comparison versus tape, offering considerable headroom for growth.

> Cost comparisons shouldn't revolve around "bytes stored" but rather "bytes protected."

For even easier management via a storage reduction strategy, users can leverage deduplication for an estimated 20x reduction and gain significant improvements to backup and restore performance. With tape requiring 6,300 cartridges to protect 42 TB over the course of a year, a deduplication disk storage system would need only 2 TB. And with backup data reduced to its raw essentials, data is even more easily transferred over a network to a disk system at a disaster recovery site.

Users can even use disk as a VTL if so desired. This takes full advantage of disk as a protection library and allows users to write to tape on the back-end if data portability is needed — the one remaining useful function of tape for the small to medium enterprise.

TAPE VS. DISK – AVAILABILITY

It is well understood that magnetic tape degrades over time. Temperature and humidity have a dramatic impact on shelf live. Ten degrees of temperature change can change the life of a tape by ten years or more. If an administrator stores tape in a room without a raised floor, there is a great danger temperature and humidity changes will accelerate the effects of thermal decay which, in turn, will destroy data in as little as five years.

The Library of Congress and the National Media Lab recommends, "for data having permanent value, storage areas should be kept at a constant 45 to 50° F or colder (do not store magnetic tapes below 46° F as it may cause lubrication separation from the tape binder) and 20 to 30% Relative Humidity (RH) for magnetic tapes (open reel and cassette) and 45 to 50% RH for all others. Environmental conditions must not fluctuate more that ±5° F or ±5% RH over a 24 hour period. Tape should be stored in dark areas except when being accessed, being sure to keep recordings away from UV sources (unshielded fluorescent tubes and sunlight)."

Widely fluctuating temperature, or RH, severely shortens the life span of all tape. This is one of the main reasons why tape is only viable for the large enterprise that can afford a library large enough to maintain tape on raised floor handled exclusively by a robot.

The design of the cartridge and the transport are critical to tape reliability as well. The enterprise class transports used today are in the 400,000-hour range. A well-managed cartridge (correctly controlled temperature and humidity), that is also a stagnant cartridge (i.e. a cartridge that has not been used), has a shelf life of around 20 years.

NEXSAN

SMALL TO LARGE

Nexsan has the disk library for all your backup and archiving needs:

- SATABoy delivers 14 disks in 3U and is perfect for protecting small application environments with advanced backup software
- SATABeast is Nexsan's high density offering with 42 disks in 4U. Modules can scale to whatever capacity and performance level is needed
- E-60 with 60 drives, can be used for even larger backup demands
- For the ultimate in simplicity with full-featured storage services like snapshots, virtualization and more, look at the Nexsan iSeries IP SAN
- Use Assureon for the ultimate in online archive. Assureon is a self-managed archive that guarantees conformance to regulatory objectives for privacy, record integrity, and data longevity. Assureon reduces the amount of complexity, infrastructure and cost necessary to manage a protection archive over and above any other solution on the market today.

Considering a shelf life of 20 years, at least 6 generations of change would have evolved in transports. Without the transport that wrote the cartridge along with the application software, operating system, computer hardware, operations manuals, ample spare parts and the recorded media itself, data cannot be retrieved. Even with all of those moving parts in harmony and perfect environmental conditions, chances of getting all that data back are still bleak when considering the reliability statistics earlier cited.

If anything goes wrong with any of the cartridges used for backup, there is no redundancy, which means an organization is unable to retrieve their data. IT organizations deal with this by re-mastering data onto new transports and new media with every generation they change — a very expensive and laborious process.

The mechanism for reading and writing tape is FAR more complicated than disk. With disk, there is a flat, stable surface that spins without flexing in a hermetically-sealed and contaminate-free enclosure. Beyond the disk itself, disk arrays offer complete data redundancy with RAID technology and 99.999% availability via hot-swappable components, redundant controllers, power supplies, etc.

TAPE VS. DISK – POWER EFFICIENCY

Tape has long been considered the most power efficient media since it can be taken offline. However disk has made huge advances in power efficiency with spin-down technology, like Nexsan's AutoMAID, that enables the cost-efficient, long-term retention of files by placing disks into progressively deeper states of sleep while offering near instantaneous response. With advanced power savings, Nexsan disk arrays allow the power and simplicity of disk backup with greatly reduced power consumption.

With an easy-to-use power configuration manager, users can create policies for desired power savings after user-defined periods of idle time. When idle thresholds are met, AutoMAID progressively reduces disk drive power consumption.

The first I/O request will wake the array up to full power. Once the array is awake, it performs at 100% performance until enough idle time has passed to satisfy the energy savings policy, which then places the array into increasingly deep levels of sleep. All of this happens automatically while providing great response performance — "speed with green."

"

34% companies who backup their data to tape and never test their backups.

CONCLUSION

Although the cost-per-byte stored on tape media is less than disk, it is an isolated figure that gives a very incomplete look at a much larger picture. Typical "Grandfather-Father-Son" management produces about 25 to 1 more copies on tape than disk. That alone makes tape much more expensive. The choice is even clearer when adding the soft cost to manage tape, the risk of data loss and downtime, performance limitations and the inconvenience of data that is offline. Protection, performance, reliability, management and cost all favor disk storage. And with Nexsan's exclusive AutoMAID power intelligence, online retention of infrequently accessed data is justified, making Nexsan the best choice for disk-to-disk backup.

From the early 1950s until the late 1990s, the volume of data made sense for tape technology. But with the explosion of the digital universe, tape can't reasonably sustain the role it once held. For most organizations, that threshold has already been reached as they can't even backup all their data within the necessary window, let alone restore data fast enough to meet business requirements.

As the pioneer of disk-to-disk backup, Nexsan was the first to understand and deliver the benefit of low-cost disk for the backup environment. As such, Nexsan's unique position in the marketplace has been delivering unparalleled value and leadership to enterprises of every size for over ten years.

For a free consultation on how your organization can benefit from a Nexsan disk backup solution, call **866.4.NEXSAN** or visit **www.nexsan.com**.

ABOUT NEXSAN

Nexsan® is a leading provider of innovative data storage systems with over 10,000 customers worldwide. Nexsan's pioneering hybrid storage systems combine solid-state technologies, spinning disk storage and advanced software to deliver radical new levels of performance and capacity at lower cost. The company's advanced technologies enable organizations to optimize traditional, virtual and cloud computing environments for increased productivity and business agility. With more than 28,000 systems deployed since 1999, the company delivers its data storage systems through a worldwide network of solution providers, VARs and system integrators. Nexsan is based in Thousand Oaks, Calif. For more information, visit **www.nexsan.com.**

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | **p.** 866.4.NEXSAN **f.** 805.418.2799 | www.nexsan.com