

The evolving role of disk and tape in the data center



Best Practices for backup and long-term data retention

Written by Debbie Beech

Introduction

Few organizations today would deny their absolute dependency on both their computer systems and on the information they store. Exponential data growth continues, and despite the current economic climate, the growth rate of storage requirements has not abated. Trends such as increased user-created content coupled with mounting pressures from regulatory compliance, are causing us to store more information over longer periods of time.

Simply adding more storage capacity to keep up with this data growth is no longer an acceptable strategy with organizations facing ever tighter constraints in budget, physical floor space, power, and management resources.

Against this backdrop, storage is continuing to evolve and diversify to keep pace with data growth. Technologies such as virtualization, deduplication, and solid state disk are adding to the range of options available. While much is being made of the potential opportunities opened up by disk-based solutions for data protection, as opposed to traditional tape, the question for the majority of businesses is not about whether to choose disk or tape, but about how to use both technologies to best effect in modern tiered-storage architecture. Put simply, decisions on storage should be made in the context of a balanced equation which includes the business need they meet and the cost involved to maximize the return on asset (ROA) of each storage device.

This whitepaper describes five best practices for backup and long-term data protection and explores the aspects of disk and tape solutions that, when combined into a disk-to-disk-to-tape solution, provide the most robust answer to backup and archival for the data center.

Increasing demands on storage

Ziff-Davis Enterprise Storage Buyers survey, December 2008:

- 89% of respondents will maintain or increase storage purchasing in 2009
- 90% expect storage volumes to increase in 2009
- 57% expect storage volume to increase by more than 25%

What causes data loss?

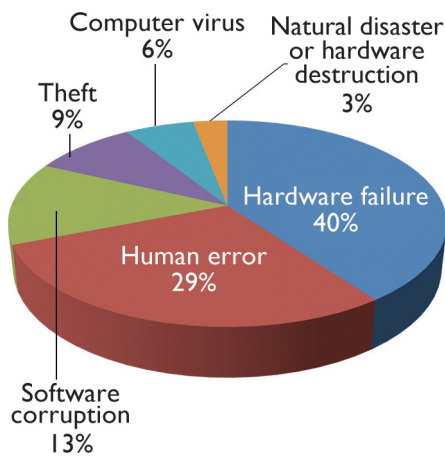


Figure 1. Source: Dr. David M Smith, Pepperdine University, 2003

Unstoppable data growth

Organizations of all sizes are grappling with the need to keep larger and larger volumes of ‘business critical information’, as well as ‘other’ information, for longer and longer periods of time. A recent InformationWeek Analytics/Byte and Switch Storage Management Survey¹ reported that over 50% of organizations surveyed were actively managing between 1TB and 99TB’s of data, capacities that would have been unthinkable just a few years ago.

The reasons behind this extraordinary data growth are many and varied, but are largely concerned with:

1. Business requirements to hold more detailed information on products, services and customers that will shape business strategy, drive growth and increase service levels (business applications and data warehousing).
2. User-driven growth in unstructured content, through the increasing use of images, video and audio alongside more traditional content.
3. The need to mitigate risk influenced by both internal policies and compliance with externally imposed regulations, which necessitates the storing of data for longer periods of time.

At the end of the day, IT Managers are less concerned with the causes of this phenomenal data growth as to what to do with all the data they need to store. This is especially taxing in the area of backup and archival storage, where company data is in effect ‘stockpiled’. The challenge for organizations becomes balancing out the need to minimize storage requirements and budgets with the business imperatives of compliance and protecting company data from the many threats that can result in catastrophic data loss (see figure 1.)

Keeping pace: storage evolution

The evolution of storage in terms of capacity, access speed, and cost has been evident since early computers employed disk drives the size of washing machines, and tape reels as large as bicycle wheels. Until recently there were few architectural choices to be made and consequently both disk and tape systems were deployed in parallel. However, over the last few years significant developments in disk technology have resulted in substantial reductions in unit price. Furthermore, technology advances such as storage virtualization, data deduplication and low-bandwidth data replication bring new opportunities. By expanding the use of disk-based storage in a data protection role, one that was formerly purely the domain of tape-based solutions, enhancements can be made to SLAs through improved data recovery objectives.

The first tape drives were introduced almost sixty years ago, and since then tape technology has also continued to evolve in parallel with advances in disk technology. Storing up to 1.6TBs of compressed data, the latest LTO-4 Ultrium tape drives hold 8 times the capacity of the first LTO tape drive launched in 2000. In addition, long before it was offered on disk, tape technology offers Write Once Read Many (WORM) capability as well as hardware-based encryption.

Which best described your current backup and recovery infrastructure?

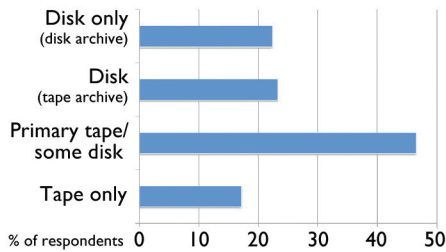


Figure 2. InforStor Annual Reader Survey
Published March 2009

“The low cost and data protection portability of tape storage remains key for our business resilience, now and for the foreseeable future.”

Dick Cosby,
System Administrator, EDPS

Despite the evolution in tape, as disk-based backup solutions have gained in momentum attempts have been made to promote the idea that tape is “dead”, or at best a legacy solution. However, published research supports a very different picture, with only a minority of organizations adopting a disk-only approach to the storage of data. The InfoStorⁱⁱ Annual Reader Survey published in March 2009, indicates that tape is still a valuable component of the storage hierarchy, with three-quarters of those surveyed stating that they used tape either exclusively or as a primary element of their backup and recovery infrastructure (as shown in figure 2). Additional research by Fleishman Hillardⁱⁱⁱ reported that over two-thirds of respondents currently using disk-only backup were looking to add tape back into their storage infrastructure, while 40% of managers in tape only environments are forecasting the increased use of tape.

It would seem that tape technology continues to offer an optimum balance of price and performance for storing data with lower access requirements. Additionally, companies looking to store vast amounts of streaming data (e.g. databases, digital imaging, multi-media etc.), can also benefit from tape's high performance backup speeds. With native transfer rates of up to 120MB/s, the LTO-4 tape drives stream up to 864GB of 2:1 compressed data per hour, easily equaling or surpassing the backup speed of the majority of today's disk drives in data streaming applications.

The considerable demand for tape is further evidenced by the fact that tape drives and tape libraries are offered and promoted as part of the storage ecosystem by all major system manufacturers. In fact, the branded tape drive market was worth \$3.47bn in 2008 with shipments of over one million drives across all market sectors. Furthermore, major manufacturers including HP, IBM, Quantum and others, continue to invest in further development of successful tape technologies such as LTO and DAT.

The key to understanding the best use of disk and/or tape for data protection lies in recognizing their place in the storage hierarchy. The following best practices for data backup and longer-term data retention provide a guide to choosing storage strategies and technologies to optimize your organizations data protection.

Best practices in backup and long term data retention

Best practice I: Be efficient - store wisely

Not all data is created equal - a consistent theme from industry experts and the first step in controlling storage costs is the need for data lifecycle management. That is, it is helpful to consider the value of the different types of data that an organization must store, and how that value changes over time. While keeping all data close at hand on high speed disks might seem ideal for access purposes, in reality to do so could be prohibitively expensive in terms

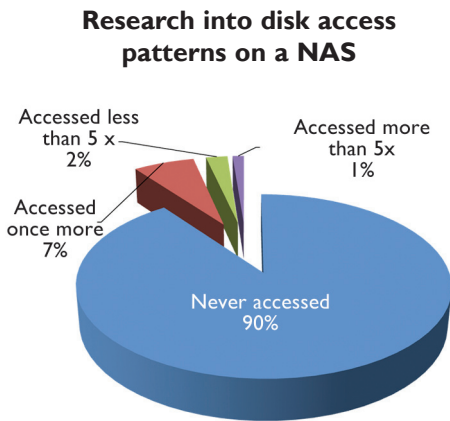


Figure 3: Research by University of California, Santa Cruz

“Employing a tiered storage approach is just patently good sense; it has resulted in a reduction to both operational and capital expense.”

Roger Smale,
Western Power Distribution

Example of tiered storage

A production system that requires 24 x 7 uptime:

- Production data held in primary disk storage
- Disk mirroring provides a copy of all production data in a 24 hour window
- Backup to tape every 24 hours from the mirrored disks (does not impact system performance)
- Tapes sent offsite for archive and disaster recovery purposes

of both hardware purchases and the cost of power, cooling and physical space, especially when compared with tape storage. In other words, why fly the data ‘first class’ when ‘coach’ will do!

Reliable estimates have shown that more than 95 percent of data stored is rarely accessed beyond 90 days after it was created. A recent study by the University of California at Santa Cruz^{iv} showed that 90% of data stored to NAS was never accessed again, and another 6.5% of the data was only accessed once more. The challenge is to identify the 95% of data that is probably and unlikely to be accessed again, and store it to a more cost-effective medium, such as tape.

Best practice 2: Implement tiered storage architectures

Tiered storage architectures - also known as hierarchical storage management (HSM), provide a way of matching the business value of the data to the most cost-effective form of storage. At its most basic, it would seem appropriate to put the highest value, time-critical information on storage media that can be easily accessed with minimal time to access data, and to archive little-used information onto low-cost storage media with a proven shelf-life yet acceptable access time. For example, the use of a secondary disk system or virtual tape library meets high performance backup requirements. By blending with physical tape for low recall backups and archival, the IT manager also benefits from offline, offsite, low cost, low energy data protection and disaster recovery.

To determine the tier of storage required for each type of data, it is useful to consider the following service requirements:

- **Recovery Time Objective (RTO)** - how quickly you need to get this type of data back
- **Recovery Point Objective (RPO)** - how recent the data must be in order to minimize impact to your business - minutes, hours or a few days

As well as RPO and RTO, the third criterion is, of course, budget. This needs to include both initial purchase cost and other factors that add to the total cost of ownership of both the equipment, and the information it stores. With these factors in mind, the storage scenarios concerned with backup and long term data retention include:

Archiving - Email and older versions of files that are in an archive state (non-changing) can clog storage devices; removing them and saving them to a lower tier of storage (SATA disk or tape, depending on the RTO) can both free up valuable “productive” storage space and reduce the costs of the overall storage environment. Key considerations when choosing to archive storage include:

- Time to access archived data (RTO)
- Cost of long-term storage - media cost, energy costs, cost to transport and hold data offsite

Consequence of non-compliance:

In February 2006, a U.S. investment bank offered to pay \$15 million to resolve an investigation by U.S. regulators into the bank's failure to retain email messages. Email took center stage in a \$1.58 billion judgment against the company in a case that focused on the firm's inability to produce email documents.

Analysts predict data loss spike:

According to analysts KPMG, 47.8 million people were affected by data loss incidents in the three month period between September and November 2008. This statistic has lead KPMG to predict that 2009 will see unprecented numbers of data breach incidents, doubling the levels of 2008 and impacting more than 190 million people.

Security tops list:

InformationWeek Analytics Storage Management Survey, 51% list data loss/data security as their top concern when it comes to storage.

- Portability and amount of physical space required
- Shelf-life of the media used
- Level of data security needed

Data retention for compliance and e-discovery (deep archiving) - separate from archival of more unstructured, infrequently used data is the need to retain information for compliance and business governance reasons. There are now over 10,000 regulations in place throughout the world that require records to be held for certain periods of time., for example; Sarbanes Oxley (US), European Union Data Protection Act (Europe), Electronic Ledger Storage Law (Japan) AIPA (Italy), HIPPA (USA) to name but a few. Companies that do not comply face hefty financial penalties, bad PR and even imprisonment for key board members. The same considerations as those for choosing active archiving solutions may be used for longer term data retention and compliance.

Data backup and restore - ensuring the timely restoration of data following a user error, system failure or other occurrence. Critical decisions to determining which storage technology to choose include:

- RTO and RPO
- Physical data storage and security, especially for data in transit

Business continuity and disaster recovery - in the event of a significant system failure due to malicious act or natural disaster, what provision needs to be in place to get the business back up and running? Considerations include

- RPO and RTO
- Determine off-site storage locations, it is vital to remove data from the same threats posed to the data center Using different tiers of storage to meet these various backup and data retention requirement can significantly enhance storage efficiency and reduce costs.

Best practice 3: Mitigate risk

Don't put your data at risk - there are a number of best practices concerned with minimizing the risk to your data from threats previously described in this paper.

Have multiple levels of protection - it is advisable to hold at least 3 copies of data in different locations, including one of these stored in a remote region for disaster recovery purposes in the case of fire, flood, earthquake or criminal attack. This latter event impacted a UK flight simulator site as described in the left column at the top of page six.

It should also be noted that having data copies held on different forms of storage media, for example a mixture of disk and tape, can also avoid system or media process issues.

When sabotage strikes

May 2009: A flight simulator site in the UK was “destroyed” by malicious hackers who took down the site’s two servers.

Unfortunately the company had dutifully backed up the servers to each other and had not established an external backup system. There was no fall-back position when both servers were attacked.

Illustrating the need for system isolation

January 2009: Blogging service Journal-Space was completely wiped out after the drives that housed their entire database were accidentally overwritten.

The problem was that their backups weren’t actually backups at all. The servers were set up with a mirrored RAID system so that if the primary drive should fail the secondary drive would be used to recover the primary. This on its own is a risky business, as it only protects you from the failure of one drive. In the case of Journal-Space, the drive didn’t fail but instead the data was overwritten/erased on one drive leading the other drive to follow suite and clear itself. A data recovery team was unable to retrieve the database.

If a tape backup had been made the night before, the data would have been offline and therefore protected from the replication error.

Another method in off-site data protection is the use of low bandwidth replication (LBR) - particularly relevant with virtual tape libraries. This technology offers the benefits of automating the transfer of data offsite, using deduplication to transmit only changed data and hence reducing the amount of bandwidth required to effectively replicate virtual tape cartridges between sites. However, LBR does not protect against viruses, sabotage and human/system error. A corruption at the source site could be replicated to the target site. For this reason multiple iterations of the backup data is recommended along with an original copy of the data to physical tape. Older iterations should be moved to tape to free up the more expensive disk space.

Keep at least one copy offline (system isolation) - avoids intentional or unintentional corruption. Some IT departments are looking at disk only solutions for backup, restore, archive and disaster recovery. This may include replicating data from disk to disk which can help protect it, however if there is a virus, sabotage or system error, then the replicated data may also be corrupted. The data was not isolated from the system. On the other hand, a backup copy held on tape, offline away from the system, protects against intentional or unintentional corruption, physical disk failure and provides a logically “clean” copy of the data. (See sidebar story “illustrating the need for system isolation”)

The low cost per Gbyte of tape storage (approximately 4 cents per Gbyte for an LTO-4 cartridge^v), means that multiple iterations of the data can be inexpensively stored on tape in the event a point in time copy is required. Removable disk solutions may also be implemented for removing data off-site; however these are available in much lower capacity points than many tapes and considerably more expensive in terms of cost per Gbyte of storage (approximately \$1.45 per Gbyte for a removable disk cartridge^{vi}).

Protect data at rest and in transit - every month we hear of new horror stories about loss or theft of sensitive data. With mounting legislation around the world, protecting data at rest is becoming a new business imperative and as a result data encryption is a hot topic. Data encryption can be performed using software, dedicated encryption devices and also using device-based encryption. A number of tape drives feature hardware based data encryption that delivers a low-cost, scalable solution that does not impact on system performance.

Implement the right technology mix - hardware and software failure is an inescapable fact of life, and in reality both disk drives and tape drives do fail on occasion. However, neither technology is inherently “unreliable”.

A number of myths have grown up around tape reliability. In reality, the vast majority of failures that created this myth have been attributed to other sources: human administration errors, software or system issues. In the rare event of a tape error, the data can typically be recovered. In the event of a disk error, the data is typically unrecoverable. Today’s tape drives include technology features that help overcome these issues to provide outstanding

“We utilize a disk-to-disk-to-tape strategy and send encrypted LTO-4 tapes offsite. This gives us the best performance, low cost, compliance, disaster recovery and security.”

Mack Kigada

Technical Architect, Servers & Storage
Providence Health and Services

data integrity and reliability. Let's take LTO technology as an example. Advances in the coating of tape film, read-after-write data verification and powerful error correction codes provide confidence in the integrity of data stored on tape. These robust tape cartridges are coupled with drive technology that features simpler tape paths and servo tracking systems to promote error-free tape handling. Consequently, LTO drives are specified with an impressive mean-time-between-failure rate (MTBF) of 250,000 hours at 100% duty cycle, that's 700% more than the MTBF of tape technologies created a decade ago.

To enhance reliability, disk drive manufacturers developed RAID, and some would point to this to suggest that disk is more reliable than tape. However, RAID is more about “uptime” than it is about data protection. If a drive fails, you simply keep running from the other drive(s) in the array, if a spare drive is available. However, if you have a virus that destroys data, if files are accidentally, or deliberately erased, should you suffer from a catastrophic OS failure that requires reformatting and reinstalling, then RAID does not offer protection. Only an offline backup is going to really protect you from losing data.

When it comes to minimizing risk, not one but all of the above factors should be taken into account. It's not just about the reliability of the technology you choose or the security of your location, but about the overall strategy for holding multiple copies on different media, online and offline, secured and protected.

Best practice 4: Consider TCO

“The economics of tape are pretty compelling, especially when you figure in the cost of power.”

Greg Schulz, Storage I/O Group

More than acquisition cost - In meeting today's SLAs, IT Managers need to consider all aspects of the value of a solution, not only with regard to backup window and recovery times, but also the total ongoing cost of delivering the service.

While disk invariably provides for faster data access, the introduction of lower cost disk such as SATA drives has caused some people to believe that disk is a cheaper option than tape. However, the acquisition cost of a data protection solution is only a small part of the total cost of ownership and therefore any evaluation of storage technologies for backup and longer term data retention should include the following factors:

Energy costs - the fact remains that tape systems will always require less electricity than disk systems with similar capacity. In general, disks must be kept spinning whether data is being accessed or not. Heat emission and the need for cooling add to the energy requirements for disk systems. With tape on the other hand, energy is required to write the data onto the tape media and from that point on it sits on a shelf, or in a tape library slot, where it requires no further energy apart from the occasional instance when the data may need to be retrieved.

In a data archiving TCO study^{vii}, The Clipper Group examined the total cost of ownership over a five year period for the long-term storage of data

“Storing infrequently-accessed data on disk is equivalent to keeping your car running constantly in the driveway - it wastes energy and it costs money.”

The Clipper Group 2007

in a tiered storage archiving environment. The analysis compared a disk-to-disk solution (D2D) to a solution consisting of a mixture of disk and tape (D2D2T). After factoring in acquisition costs of equipment, media, electricity costs and data center floor space, The Clipper Group found that the total cost of archiving solution based on SATA disk was about 23 times more expensive than an LTO-4 tape library archiving solution. When comparing the energy costs for the competing approaches, the energy costs for the disk-only solution jumped to about 290 times that of the tape solution.

As previously mentioned, data deduplication technologies for disk can allow a significant reduction in data backup volumes over a period of time that allows organizations to store more backups to disk, or store backups on disk for longer periods of time without requiring additional disk purchases. However, when the same Clipper Group Study described above considered the use of data deduplication with disk, the TCO was still estimated to be 5 times more costly than the tape solution. Furthermore, best practice would suggest that having at least one full copy of critical data in a ‘non-deduped’ state is advisable in case the disk or data deduplication system pointers become corrupt. Tape provides a low cost, portable solution that eliminates the requirement for a deduplication appliance at the disaster recovery site.

Physical storage space - with floor space in computer rooms and offices at a premium, the high capacity and small footprint of tape provides a cost advantage that's often overlooked, particularly in longer-term archival and offsite storage scenarios.

Management - the people costs to manage a storage solution should also be considered within the TCO. With the advent of advanced backup/archive management software and physical or virtual automated libraries, the resource costs required for managing and maintaining disk or tape based data protection solutions are roughly the same.

Scalability - as data volumes grow, scalability of storage becomes an important factor. How much does it cost to expand your storage system, will it add to management overheads and complexity? In particular, solutions such as Tape Libraries are infinitely capacity scalable; just add in another cartridge within the library.

Calculating the TCO of different technologies for your backup and archival scenarios should also include, hardware maintenance, software acquisition, software support and maintenance after warranty and a look at the need for management software and disk controllers, in addition to the cost of data migration. Typically migration to a new generation of disk system needs to occur every 3 to 5 years, whereas tape technology has a longer cycle of 7 years or more. Generally the cost of migration effort with tape is much less for tape than for disk. All of these factors have an impact on the true cost of your solution.

Best practice 5: Ensure you can restore

This is the bottom line - when all is said and done, the proof point of any data protection solution is in the ability to dependably restore files, or entire systems, when the need arises. At this point, you don't want to discover that something has gone wrong. It is therefore important to consider the following factors:

Regular checks - with all archived data, regardless of technology used for archive storage, frequent testing of restore capability is recommended, if not essential.

Shelf life - if you are planning on keeping data for a number of years in deep archival for governance or compliance reasons, then you need to ensure that the storage medium selected has sufficient expected shelf-life. In general, tape offers between 4 and 6 times the life expectancy of disk, with media manufacturers specifying up to 15 years for DAT and up to 30 years for LTO tape media.

Efficient restores - while the random access of disk supports fast file access for restore, for batch restores of large amounts of data, tape may prove equally efficient. Powerful indexing through software can also facilitate individual file restore from tape.

Disk may be necessary, but is not sufficient - new disk technologies such as disk mirroring and replication have helped to expand the use of disk into a primary backup medium. However, should something happen to the original data due to error, system failure or virus attack, data loss may be propagated through mirroring or replication to the backup copy. Used as a complementary technology in a disk-to-disk-to-tape solution, tape provides a secure, offline copy of the data and the ultimate recovery mechanism.

Recommendations for a balanced approach

In the data center, when it comes to disk and tape it is not a question of 'either-or'. The real discussion should focus on 'where' and 'how' to best use the available storage technologies as part of a comprehensive, tiered storage architecture while meeting the needs of RPO, RTO, budget and the best practices listed above.

For mission critical data that requires the fastest RTO and RPO, primary disk mirroring may be necessary. Backups should be made to tape storage for offline, offsite data protection and long term archive.

For daily short term backups users can implement a disk to disk solution such as a virtual tape library (VTL) to allow for fast recall of data. The data should be copied to tape for off-line, offsite long term data protection and to lower storage costs. As the data becomes infrequently accessed on disk it can then be removed to free up the disk space for new data. Many disk to disk solutions offer data reduction technologies such as data deduplication to help reduce disk storage requirements.

These technologies are briefly described below and then compared in the following table:

- Primary Disk Mirroring - a method in which mission critical data is written to two duplicate disks simultaneously. If one of the disks were to fail, the system can instantly switch to the other disk without any loss of data or service.
- Virtual Tape Libraries (VTLs) or Disk to Disk - backing up from primary disk storage to a secondary backup disk. A common method is a VTL where data is backed up to a disk that emulates many features of a tape drive/library, but with the random access recall speed advantage of disk.
- Standalone Tape Drives or Automated Tape Libraries - tape based backup and restore for backup, archival and data protection.
- Disk to disk to tape solutions (D2D2T) - starts out backing up to disk, then offloading data to tape for longer term data retention.

Storage technology	Advantages	Restrictions	Suggested applications
Primary Disk Mirroring	<ul style="list-style-type: none"> • Random access for faster back-up and restore than tape • More frequent RPO 	<ul style="list-style-type: none"> • Most expensive • High energy consumption • Not portable to offsite • Online (suffers from same risk exposure as primary storage) • Requires maintenance (e.g. RAID sets etc.) 	<ul style="list-style-type: none"> • When RTO is in seconds to minutes and RPO is immediate to minutes or hours dependent on backup policy/cycle • First tier backup from primary disk (D2D) • Mission critical data
Virtual Tape Libraries (VTL) or Disk to Disk	<ul style="list-style-type: none"> • Faster backup and single file restore than tape solutions • Can offer data deduplication and replication • Automation • Emulates tape - can simplify implementation and minimize changes 	<ul style="list-style-type: none"> • Higher cost and energy consumption than tape • Media not physically portable to offsite • Appear more complex to use than tape 	<ul style="list-style-type: none"> • When RTO is in minutes and RPO is specified in hours dependent on backup policy/cycle • First tier backup from primary disk (D2D)
Tape Drives and Tape Libraries	<ul style="list-style-type: none"> • Lower TCO • Offline data protection • Automation (Libraries) • Energy efficient • Long shelf life • Portable • High capacity/small footprint • Fast streaming of batch data • Easily scalable • WORM (LTO-3 & 4) • LTO-4 hardware encryption 	<ul style="list-style-type: none"> • Slower restore for single files 	<ul style="list-style-type: none"> • RTO is typically more minutes than for VTL and RPO is specified in hours dependent on backup policy/cycle • Second tier backup from disk or VTL • First tier backup in cost conscious environments • Active and deep archival • Offsite disaster recovery
Disk to disk to tape solutions (D2D2T)	<ul style="list-style-type: none"> • Combines the data access speed of disk for backup and the economies and data protection attributes of tape for longer term data retention 	<ul style="list-style-type: none"> • Need to set policies and procedures to manage migration of data from disk to tape 	<ul style="list-style-type: none"> • Covers all data protection and archiving scenarios • Most Backup & Archive software support

In summary, primary disk mirroring and high-performance disk-based backup in a disk-to-disk or VTL solution will be needed in environments requiring the most stringent RPO and RTO. Tape will complement these D2D solutions; migrating data from disk backup to tape backup provides secure, offline data protection and data archival to mitigate risk. This Disk-to-disk-to-tape (D2D2T) architecture combines the data access speeds of disk with the low TCO of tape for best effect. Applications that do not require stringent RPO's and RTO's can be backed up directly from primary disk to tape.

In addition to low-cost backup, tape is ideal for active archival of data such as human resources records, emails, or last year's patient's records, for example, with low access frequency. With a substantially longer shelf life, tape is also the ideal medium for archiving of compliance data. It is also used as the ultimate defense against collapse in disaster recovery solutions.

Conclusion

The roles of disk and tape have evolved in the data center and are complementary in fulfilling best practice strategies. In many environments, tape is no longer the primary backup target for mission-critical applications that have stringent RPO and RTO objectives. Yet, it continues to maintain its role as the primary backup target for applications with less stringent requirements and for data protection and TCO. Tape has a prominent role in disk-to-disk-to-tape strategies. It is a cost effective media for migrating older backups from more expensive disk-based targets. Tape should also be viewed as the low-cost, dependable, data protection media to store archived data for long periods of time.

About the author

Debbie Beech is Director of Sylvatica Consultants. Based in the UK, she draws upon experience gained from over 15 years of senior management positions in HP's worldwide storage business. For the last 8 years, Debbie's focus has been to deliver practical marketing advice and support to a wide range of IT companies.

Reach Debbie Beech via email at debbie@sylvatica.co.uk or telephone: +44 (0) 1454 616533

Sylvatica MC Ltd. reserves all rights regarding the use of its trademarks. All other trademarks belong to their respective owners.

Reproduction of this document or any sections within it is limited to Sylvatica and individual member companies of the LTO Program. To request advance permission to reproduce, please email Debbie Beech or geoff.mordock@fleishman.com

References

- i. InformationWeek Analytics / Byte and Switch Storage Management Survey of 328 Business Technology Professionals, November 2008
- ii. Infostor: Annual backup and recovery trends survey results 2009 - www.backupreview.info/2009/03/29/reader-survey-reveals-backup-and-recovery-trends-2
- iii. Fleishman-Hillard survey of 200 network administrators technical specialists in mid-size to larger companies across the US conducted in calendar Q4 2008
- iv. 2008 study by University of California looking into disk access patterns for the network of a large local business over 22TB of disk-based data
- v. Based on current average internet purchase price (May 2009), of a single LTO-4 cartridge holding up to 1.6TBs with 2:1 data compression
- vi. Based on current average internet purchase price (May 2009), of a single RDX cartridge holding 160GB of data
- vii Source: "Disk and Tape Square Off Again" The Clipper Group at: www.ultrium.com/News/paper-form.php?section=2&subsec=paper