



Tape The Digital Curator of the Information Age

Fred Moore, President
Horison, Inc.
www.horison.com

Introduction

Since the first successful drive first appeared in 1952, the magnetic tape industry continues to evolve. Markets are shifting as disk slowly encroaches on tape's traditional backup/recovery market while tape is positioning itself to address the exploding tier 3 applications such as fixed content, compliance and archive. In recent years tape has been relegated to the mainframe and the middle to high end of the SMB (Small to Medium Business) markets with decreasing penetration below this level due to availability of small diameter disk drives and flash memory alternatives.

StorageTek, acquired by Sun Microsystems in 2005 and again when Oracle acquired Sun Microsystems in 2010, was clearly an innovator and what set StorageTek apart were its market leading tape innovations. StorageTek introduced the Nearline automated tape library in 1988 using a new, circular robotic architecture for the first time. This bold move essentially saved the company as StorageTek found many new applications for tape beyond traditional backup and recovery. In 1998, IBM and StorageTek popularized the first Integrated VTL (Virtual Tape Library) solutions bringing a new, improved economic model to the mainframe tape industry. The Nearline library was the de-facto standard library for 15 years and many are still in operation today. Tape was seldom seen as an innovative technology, but these innovations often proved the exception to that rule and StorageTek and IBM consistently lead the industry in tape innovation.

Tape Industry Overview

Annual tape industry revenues worldwide total over \$3B including all drives, robotic libraries and media compared to a \$25B disk drive industry in 2009. Overall awareness and vendor marketing efforts for tape have been minimal in recent years impacting customer perceptions of the future of tape. As a result, many customers' perceptions of tape have become outdated and are not current with the latest developments and improvements in the tape industry. These developments include much longer media life, improved drive reliability, higher duty cycles and much faster data rates than any previous tape drives. Tape cartridge capacities have exceeded those of disk drives. Oracle's StorageTek T10000 and IBM's TS1130 enterprise-class cartridges along with

the popular LTO cartridges in the midrange-class range up to 1.5 terabytes native (and to 3 terabytes compressed 2:1) capacities while the largest disk drive contains 2 terabytes. When using tape, you actually get to use a higher percentage of the total cartridge capacity than on a disk as disk inherently requires significantly more overhead. This is especially true for RAID (Redundant Array of Independent Disks) implementations. Tape is steadily positioning itself to address the many new high capacity, high growth tier 3 storage opportunities as over 65% of the world's digital data is optimally suited for tier 3 storage. Tier 3 applications are rapidly growing storage requirements at nearly 60% compounded annually and it is this tier that will derive the most benefit from future tape developments and new architectures.

Becoming the “Digital Curator”

Given these improvements, it's now time for the tape industry to seize the sizeable tier 3 storage opportunities that await it and position itself in the 21st century as “The Digital Curator of the Information Age.” A curator means "officer in charge of a museum, library, or a collection.” What does it take to become the “Digital Curator of the Information Age? To most effectively be the Digital Curator, the following capabilities must be met.

Tier 3 - Archive Capability	Tape	Disk
Long-life media	Yes, 15-30 years on all new media.	~4-5 years for most HDDs before upgrade or replacement
Portability	Yes, media completely removable and easily transported.	Disks are difficult to remove and to safely transport.
Move data to remote location for DR with or without electricity	Yes, can move remotely with or without electricity. Natural disasters can force physical media movement.	Difficult to move disk data to remote location for DR without electricity.
Inactive data does not consume energy	Yes, this is a becoming a goal for most data centers. “If the data is not being used, it shouldn't consume energy”.	Rarely for disk, except in the case of MAID or “spin-up, spin-down” disks.
Encryption for highest security level	Yes, encryption capability available on essentially all tape drives. Highest demand with mainframe and mission critical tape applications.	Quickly becoming available on selected disk products, PCs and personal appliances.

Can tape become the Digital Curator of the Information Age? Clearly from the above table the answer is yes from a technological perspective. Let's look more closely at tier 3 storage requirements and see how tape solutions can optimize this rapidly growing storage requirement.

Tiered Storage Began With Mainframes, Nearline and DF/HSM

In 1988, the IBM mainframe operating system introduced DFSMS (Data Facility Storage Management Subsystem), which became a highly effective policy engine for managing mainframe storage resources requiring users to get to know their data and to better understand its true business value. The DFSMS concept enabled mainframe businesses to have access to the right data at the right place at the right time while optimizing the storage hierarchy for cost and performance.

There are several components included with SMS but the established DF/HSM (Data Facility - Hierarchical Storage Management) functionality was the single most important function for tiered storage, enabling a management capability for data throughout its lifetime. HSM proved to be the catalyst, enabling businesses to finally address the storage dilemma of matching policy based data attributes with the most cost-effective technologies in the storage hierarchy. The combination of DF/HSM and the StorageTek robotic library were the genesis of the tiered storage concept, enabling easy policy-based migration of inactive data from costly disk to lower cost tape storage – as a result the concept of tiered storage was born.

The space management capabilities of DF/HSM quickly began to optimize storage hierarchy investments for large enterprises. Obsolete data are deleted and inactive data are moved, or migrated, to lower-cost tiers of storage, such as low-cost disk or automated tape libraries for long-term data retention. Migrating the data to lower cost storage tiers allows the data to be automatically retrieved, or recalled, should the data be required at some future time, yet frees space on more costly disk devices for more active data sets. As a result of widespread use of DF/HSM concepts, robotic tape libraries became a new tier of storage between online disk and offline tape. This new level of storage was called Nearline.


To implement a data lifecycle management strategy from a technology perspective, the de-facto standard three-tiered storage hierarchy model emerged as the preferred and optimal choice. These tiers include primary storage, always disk-based for highly active, mission critical and for most customer-facing revenue generating applications. Secondary storage included midrange disk arrays, integrated virtual tape for enterprise systems or low cost SATA-based disk systems for data that has a lower activity level. The third tier, long-term archival storage, remains the realm of magnetic tape and automated tape libraries. Optical storage is used heavily in the entertainment industry and with personal appliances but it is no longer a data center storage technology. Long term storage applications such as archive, compliance and fixed content apps are now the largest growth market for tape as the economics of tape storage for these applications are more favorable than other technologies. All data is not created equal and the value of data can change throughout its lifetime. For lifetime data management, “It doesn't matter if the data is ever used; it does matter if the data is there and can be accessed.”

A true tiered storage concept for non-mainframe systems has been much slower to develop. Without a continuously running DF/HSM-like function for Unix, Windows or

Linux servers, tape remained primarily a backup solution for these systems while low activity files frequently accumulated on disk. That will begin to change as tiered storage is poised to take off in non-mainframe markets where the total amount of installed disk capacity usually is greater than on a mainframe system.

Businesses often calculate *availability* indexes for key applications in terms of “the number of 9s.” The estimated average costs of system failures would be nearly fatal to some companies and can exceed \$3 million per hour of downtime in certain industries. A storage subsystem or server that is 99 percent available may seem highly available, but will actually be unavailable 5,000 minutes per year. Availability figures range from approximately 99 percent for Windows-based servers to over 99.999-plus percent or five-9s for the latest enterprise z/Series mainframe servers. Revenue lost per hour of outage reflects on the criticality of the IT function to a particular business. When implementing tiered storage, it’s important to match application availability requirements with the technology that can meet those requirements.

The Tiered Storage Hierarchy



Source: Horison Inc.

Integrated Virtual Tape Libraries are a Key to Optimizing the Storage Hierarchy

Virtual tape concepts for mainframe computers were pioneered and popularized in 1998 by IBM and StorageTek as the effective utilization of tape cartridges was historically low. An integrated VTL is the optimal integration of tier 2 and tier 3 storage solutions for mainframe storage applications. The integrated VTL combines disk arrays as a front-end to an automated tape library. The disk storage serves as a cache or buffer for more active data for the larger-capacity and lower-cost tape library. The device images presented to the operating system appear as multiple tape drives rather than physical disk drives, therefore “virtualizing” the disk by making it look like something other than it really is. Integrated VTLs store multiple virtual tape volumes on a single physical tape cartridge, commonly enabling cartridge utilization levels to reach 80 percent or more. Effective utilization of storage assets becomes more important as cartridge capacities steadily increase.

Pre-established policies such as file size, capacity utilized and frequency-of-usage patterns determine when the data is moved directly to and from the disk buffer and the automated library without going through the server – using zero additional server processing overhead. In this case, the policy-based functionality resides outboard of the host application server, directly controlling bi-directional data movement between disk and tape storage enabling direct data transfer between both Tier 2 storage and Tier 3 storage to combine fast backup and recovery on Tier 2 disk with Tier 3 tape archiving capabilities.

The tape to disk data ratio on mainframes is higher than this ratio on non-mainframe systems. This is due to the aggressive mainframe HSM software that continually sweeps less active data from more expensive disk to less expensive tape. Tape is the least expensive medium to store data long term in the data center, is environmentally very friendly, is easily transportable with or without electricity, and actually has the highest bandwidth (if physically transported on the highway).

Integrated VTLs make traditional tape solutions much more cost-effective and appealing by:

- Reducing the number of tape cartridges by increasing cartridge utilization.
- Reducing the number of tape drives by presenting multiple virtual tape drive images which can be utilized by many backup jobs during the backup window.
- Improving tape drive and library reliability by satisfying many tape requests from the disk buffer and avoiding access to the tape subsystem. Virtual tape is optimized by ensuring that, while the probability of recall is reasonably high, the data is still in the disk buffer. The outcome is fewer physical tape mounts, and therefore less wear and tear on tape drives and media.
- Lowering the overall costs of the tape environment by reducing hardware since fewer tape drives and cartridges are needed. Fewer drives also lower energy consumption.

- Improving IO performance as many requests are satisfied by the disk buffer avoiding a physical tape mount – larger disk buffers make performance even better.

An integrated VTL is the optimal mainframe solution for both backup/recovery and tier 3 applications such as fixed content, compliance and archival data as it integrates tier 2 and tier 3 storage levels into a single architecture.

For non-mainframe systems, the market opportunity for automated tape libraries is rapidly increasing as a result of the explosion of tier 3 data and the requirement for managing, storing and protecting data throughout its lifecycle. This makes managing and moving data through tiers of storage increasingly important and is more effectively addressed through the adoption of tiered storage strategies that use enabling software such as Oracle's SAM-QFS (Storage Archive Manager - Quick File System). QFS is an open source file system that is tightly integrated with SAM software extending the QFS file system transparently to manage archival storage. SAM-QFS integrates tier 2 and tier 3 and it is this software that makes Oracle's tiered storage become an "active" tiered storage solution.

Tiered Storage Future Considerations

Tape users and vendors will need to develop strategies to more closely integrate and manage tier 2 and tier 3 storage capabilities - what might they be considering to capture the huge tier 3 storage opportunity? Tiered storage software can take advantage of increasing HDD (Hard Disk Drive) capacities for larger buffers and coupled with continual capacity increases in tape cartridges to solidify the position of tape as the Digital Curator of the Information Age and setting the stage for multi-petabyte (PB) architectures that will soon be required. For example, imagine a very large disk buffer with sub-tiers of tape storage behind it. One tier potentially could be very low cost disk; the second tier might be a high-performance fast access tape drive, with the 3rd tier being a very high capacity terabyte plus tape drives.

What Applications are Optimal for Tape?

While some backup applications are moving to disk in non-mainframe systems, long-term storage applications like digital archives, fixed content/multi-media and compliance are ideally and economically better suited for tape storage. A growing list of government and legal regulations worldwide now describe the way data must be managed and stored throughout its lifetime. Features such as encryption and WORM (Write-Once-Read-Many) are required for many of these applications. The amount of data at the back end of the data life cycle is now growing, not shrinking, as was the case a decade ago. For some data, the storage preservation requirement will become infinite. For the majority of digital data, the generalized axiom of "90 days on disk and 90 years on tape" refers to lifetime data management and to tape as the curator.

Data retention policies are now being based on data value and legal requirements in addition to reference activity, mandating that a universal, standard data classification scheme must emerge. For some data files, the duration of storage preservation

requirements has become infinite such as digital videos of most major sporting events which will never be deleted. Most of the fixed content, compliance and archive data that makes up tier 3 is unstructured data in the form of file storage. The social networking wave has become a major generator of unstructured digital data. Some of the primary drivers of tier 3 storage are listed below.

Primary Application Drivers of Tier 3 Storage Demand

Digital Assets	Rich Media (3D, Multi-dimensional)
- E-mail and long term archives, ~80% is stored in attachments	- Digital Audio & Video (iPods, MP3, Digital Surveillance, YouTube, Motion and Streaming)
- Database Archives	- Medical Images (3D MRIs, CAT Scans)
- Compliance & Litigation	- Geophysical, Geospatial, GIS, Seismic, Google Earth etc.
- Medical Files/Images	Fixed Content and Archival Futures...
- Insurance Claims	XAM, a New Fixed Content Data Management Interface Standard - Using Metadata
- Web Content, static images	Automated Tiered Storage for Tape and Disk (HSM-like policy-based software)
- Social Networking Explosion, Wikis	Linear Tape File System (LTFS) for enhanced tape archive storage and retrieval for open systems and mainframes
- Cloud Applications	
- Digital Photos	
- Document Imagery (Printed Materials, Books, Contracts)	

Tape Means Dark Green Storage

As the cost of power continues to increase, the favorable impact of using tape storage on the IT budget will only grow. Rising electricity costs, scheduled power outages, limited floor space, growing data volumes, and longer data retention periods for low-activity data will ultimately affect all businesses. Average data-center energy costs are growing at 15-20% per year or more and are expected to double by 2012. Unlike storage providers, energy providers have shown little interest in lowering their rates! Compounding this dynamic is the fact that power density is going up for most IT equipment, at a rate of 20% to 30% per year. This has the following domino effects; 1) much more power needs to be supplied to each square foot of a data center 2) more power is required to cool hotter equipment and 3) more heat extraction equipment needs to be supplied to each square foot of a data center. The maximum heat density that can be air-cooled in a data center is approximately 10,000 watts/sq ft, yet most data centers have been designed for power densities of less than 1,000W/sq ft. The bottom line is that IT is approaching the limits of power distribution in many data centers, which is forcing organizations to begin exploring new cooling techniques such as water-cooled racks or in some cases, building

another data center. Building another data center is normally a last resort and is very expensive.

Average Electrical Power usage inside Data Centers

IT equipment as a percentage of IT electrical consumption

Servers 48%

Storage (disk) 30%

Storage (tape) 6%

Networks 12%

Other 4%

Source: Horison, Inc. and estimates from various industry

Average electrical consumption within a data center (not including external HVAC) indicates that tape is the greenest storage technology and typically uses about one-fifth of the electricity that disk does. Utility companies in certain areas are restricting the amount of power some businesses can use. Focusing on disk storage can decrease its 30% electrical consumption figure considerably by shifting less-active data to tape storage and this has clear advantages when it comes to green initiatives in the data center.

The growing energy crisis will continue to drive more aggressive use of environmentally conservative, green storage solutions, particularly favoring tape and tiered storage. Developing an energy strategy is becoming mandatory and several businesses have created a new position of “Energy Systems Manager” to ensure that floor layouts and airflow are properly designed and adapt to ongoing equipment changes in the data center while minimizing inefficiencies.

Best Practice for Tape

For organizations that need fast recovery for large data files, tape remains a key backup technology that is unlikely to go soon. The heavy use of tape in mainframe and enterprise environments continues to prove this point. The promise that bandwidth would replace trucks for moving large amounts of data has not come true. The growth in data that needs to be backed up exceeds the growth rate of bandwidth, and will continue to for the foreseeable future. This gives tape a financial advantage over disk as an off-site backup technology that is compounded by the lower energy cost and overall TCO of tape versus disk over the lifetime of the stored data. Tape also has a much longer storage life than disk, eliminating the need and costs to move data to new media as frequently during its lifetime. Organizations that need a copy of data in WORSE mode,” Write Once Read Seldom if Ever,” can benefit from automated tape economics that presently show a 1/3 to 1/10 \$/GB advantage over disk technologies depending on the configuration.

Encryption implementation is becoming a critical feature for tape to protect digital archives for indefinite periods of time and most tape drives offer encryption. The tape drive should be crypto-ready and be designed to simplify data security and mitigate risk by providing device-level encryption while requiring no infrastructure changes or decryption software. The encryption solution should be easy for IT staff to implement,

paying particular attention to selecting a bullet-proof key management system. Device-level encryption provides the easiest way to decrypt data over time. Data can be compressed prior to encryption so you can optimize storage capacity. With tape-based encryption, you add an extra measure of protection when tape cartridges are stored offsite or if they should be lost.

Best Practices for Effective Tape Usage

Select the optimal data for tape storage – fixed content, compliance, archive and backup.
Use automation whenever possible from small-scale autoloaders to robotic libraries.
Implement integrated virtual tape and tiered storage concepts to reduce CAPEX & OPEX while improving performance and availability.
Take advantage of tape’s much longer media life and drive reliability improvements.
Consider encryption for tape and mobile media including PCs and personal appliances.
Clearly understand the vendor’s future commitment to tape development and innovation.

Tape Economics

Economics continue to favor tape storage over other storage products for both an acquisition cost and ongoing operating expense perspective. The table below highlights the average selling price ranges for various storage devices. Prices vary widely as the amount of storage, account status with a vendor and strategic value of a potential customer can influence actual pricing.

Subsystem category	ASP (Capex) Range \$/GB	Notes
Solid State Disk (DRAM) FC, SCSI (Flash SSD/HDD)	\$300-\$500 \$50-100 (current ASP est.)	ASP range based on capacity and high-availability features
Enterprise disk FICON, ESCON, SCSI, FC, SAS	\$7-20	Includes controller, cache & drives. Add-on storage modules with no controller and cache are lower
Midrange disk SCSI, FC, SAS, Virtual Tape, NAS	\$1-8	Includes controller, cache & drives. Add-on storage modules with no controller and cache are lower
Optical disc library (Blu-laser) SCSI/WORM	\$5-20	Includes drives, media and library
Economy disk SATA, MAID, JBOD	\$1-5	Price range varies widely based on capacity
Automated tape library	\$.20-\$2.00 (non mainframe) <\$.20 (mainframe, varies by configuration and cartridge to drive ratio)	Includes tape drives, media and library and uses a 2:1 compression to determine library capacity

Source: Horison Information Strategies (April 2010 update)

A tape rebirth may have been signaled. HP, IBM Corporation and Quantum Corporation, the three primary providers for the Linear Tape-Open (LTO) Program, released survey

results that strongly suggest that storage customers that use a disk-only infrastructure are now looking at tape storage as part of a tiered storage infrastructure to support backup and the growing archive requirement. Over two thirds of businesses surveyed said they were looking to add tape storage back into their overall network infrastructure and of those respondents; over 80-percent plan to add tape storage solutions within the next 12 months. This doesn't sound like tape is going away – does it?

Positioning Tape Growth for the Future

It is clear that future application needs for tape are in place. Tape is one of the most efficient ways to address long-term data protection while delivering cost-efficiency and energy savings to businesses. The economics of tape for many applications are compelling from an acquisition cost and cost of ownership perspective. The opportunities for tape success and sustainability are now evident, its time for the tape industry to aggressively raise the awareness of these important capabilities.

Opportunity Areas for Tape Growth	Comments
Tier 3 storage applications such as fixed content, compliance, archive and social networks are exploding.	This market is generating true storage demand at ~60% annually and is the fastest growing storage segment. Over 60% of data center data and much unstructured data are tier 3 candidates.
Green initiatives and energy consumption favor tape. Ultimately, data that isn't accessed shouldn't consume energy	Tape storage is the greenest of storage solutions; tape vendors will more aggressively market these advantages and facilitate customers to easily move inactive data to-and-from disk to more eco-friendly tape storage..
Today's tape has significantly improved media life and drive reliability compared to mid 1990s levels.	Customer perceptions of tape are outdated and left over from the late 1990s. The new tape products have 8-10x improved MTBF, duty cycle and media life capability compared to those of 15 years ago positioning tape as the optimal archive solution.
Economics and pricing of tape are favorable compared to disk.	Tape is less expensive per GB to acquire than disk and uses less energy than disk – potentially zero energy. These metrics must stay below disk prices for tape to remain viable. Tape technology roadmaps indicate this will happen.
Successful tape vendors need to escalate marketing activity to build mindshare and improve awareness of tape improvements	Websites, trade journals, storage events, user's groups presently have decreasing focus on tape compared to disk solutions.
Removable media has advantages.	The promise that bandwidth would replace trucks for moving large amounts of data has not come true because the amount of data that needs to be backed up exceeds the growth rate of affordable bandwidth – and is expected to for the foreseeable future.

Electrical outage versatility	Tapes can be moved during electrical outages while disk drives almost always require electricity to move data.
A new vision for tape is needed...	Tape needs to offer more than increased capacity and data rates in future roadmaps. Look for tape vendors to develop exciting new tape capabilities such as LTFS.
Tape Requires Fewer Device and Media Conversions Than Disk.	Tape drives have longer useful life than disk drives as current tape media is rated at 15-30 years, greatly reducing the need to move data to new media as frequently.
Tape Has a Lower CAPEX and OPEX Than Disk.	Not only is the CAPEX (\$/GB) cheaper for automated tape than disk, the ongoing costs are lower as ultra-high capacity cartridges, simpler backup processes, and lower energy costs keep down expenses.

Conclusion – Tape is Extending Its Role to be the Digital Curator

Tape isn't going away as its role is in transition, expanding from a pure backup solution to that of a premier long-term storage technology and archive. In the last few years a few vendors have tried to advance integrated tape library technology (as opposed to pure disk "virtual tape" solutions) beyond the mainframe market. Integrated VTLs have demonstrated for mainframes significant Capex and Opex benefits for optimizing the cost/benefit ratio of long-term storage administration, making the biggest question in tape's future not its viability, but how quickly the tape vendors will begin to drive new technology capabilities forward.

In difficult economic times, the value and cost savings of using a tiered storage strategy with more economical tape as the preferred tier 3 technology choices couldn't be greater. Tape vendors who want to seize the mantle of "Digital Curator" will jump on this opportunity to stimulate new and growing interest in tape storage solutions. Remote storage, data vaults, lower operating expenses and the ability to move data when electricity isn't available all favor tape. The 21st century data explosion is here – and tape is well positioned to become The Digital Curator of the Information Age.