



THE CHALLENGE OF DATA GROWTH AND THE PROMISE OF LTO-8 TAPE

BY

JON TOIGO, CHAIRMAN

DATA MANAGEMENT INSTITUTE

WWW.DATAINSTITUTE.ORG

10 OCTOBER 2017

Summary

With the announcement of its support for Generation 8 of the Linear Tape Open Ultrium Tape Format, LTO-8, IBM has confirmed its commitment to and support of magnetic tape storage as part of the total storage technology mix. IBM has announced a new LTO-8 tape drive, the TS2280, and its support across a range of IBM autoloaders, libraries and storage enclosures. It is also extending the support of IBM Spectrum Archive to IBM LTO 8 solutions. This move makes sense, especially given the accelerating rate of data growth that is being witnessed both by cloud service providers and large enterprises. Driven by ongoing analog-to-digital data conversion, increasing use of mobile commerce, and by analytics processing and the internet of things, data is expected to top 160 zettabytes by 2025. Storing that much data will require the integration of multiple storage layers and the intelligent placement of data on storage based on business, regulatory and operational considerations.

Introduction

On October 10, 2017, IBM announced its support for the latest generation of Linear Tape Open Ultrium tape technology, LTO-8, by previewing its LTO-8 tape drive, the TS2280, and a plan to integrate the drive into existing tape autoloader, tape library and multimedia storage enclosures. IBM also announced that its IBM Spectrum Archive (Linear Tape File System or LTFS) software would be updated to support LTO-8 products.

This announcement came well ahead of the formal announcement by the Linear Tape Open Technology Provider Companies (TPC) of the availability of LTO-8 technology media. When finalized by the LTO consortium members, the media will deliver

- 12TB native cartridge capacity and 30 TB compressed capacity
- 360 Mbps native and 900 Mbps compressed transfer rates
- Increased bit density over any prior LTO generations
- More tracks than previous generation LTO with same tape length
- Read/write compatibility with LTO-7
- Archival life of 30+ years

Even more compelling than the technical attributes of LTO-8 is the strong business value case that can be made around the technology. The combination of its cost-containment features, its contributions to risk reduction, and its enablement of improved productivity for information processing deliver outstanding business value. This three-part business value narrative is required to obtain management buy-in and support for any technology initiative within contemporary business organizations.

This paper articulates the continuing case for tape technology and illuminates the rationale for including tape in a robust and well-managed data storage infrastructure.

THE TECHNOLOGY CASE

Before jumping to the business value case for tape, it is useful to re-visit the technical capabilities of the tape storage medium that have sustained its productive use almost since the dawn on modern computing. Tape was among the first types of magnetic media introduced for data storage. From its first implementation by IBM in 1952, magnetic tape has undergone steady evolution and innovation.

Over the years, tape has progressed as a technology from a medium supporting analog recording to one adept at the recording and playback of digital data, from a technology with limited areal density design to hundreds of billions of bits per square inch densities today, and from storage capacities of 224 KB to 15 TB per cartridge (uncompressed). All along, tape has delivered one of the best data throughput rates, featuring extremely limited jitter characteristics that make it a perfect medium for streaming data such as video, genomic data, telemetry, and other “long block” data objects and files.

The durability of tape and its bit error rate characteristics are among the highest of all storage media. Contemporary tape boast bit error rates that are far superior to those of disk, reducing the likelihood of silent corruption of data at the time that data is recorded to that of a statistical anomaly. The rated

lifespan of contemporary tape media is 30 years, though the pace of technology change in tape typically sees users upgrading their drives and media about every five to seven years.

Key to understanding the benefits of tape is how the technology fits together with other storage elements to create a tiered storage infrastructure. In most analyses, tape provides the foundation – or long term retention media – for long term data storage; archive is tape’s “killer app.” Data may be written to and read from flash memory or hard disk-based storage while its usage characteristics remain “hot” (frequent access and frequent update are being made to the data by applications or users), but as data “cools” over time, migrating it to tape makes more sense for several reasons:

- Tape provides the most inexpensive capacity for storing less frequently accessed/updated data, even data used in long term trend analyses and similar Big Data operations. Moving older data to tape frees up space on more expensive hot storage media.
- Data stored on tape is protected by an “air gap.” Data on tape is not connected to servers, networks or applications until it is requested and mounted. Because of this, the data is insulated from most computer malware and viruses, including ransomware attacks. [More and more companies are turning to tape to protect their data assets from hacking events.](#) Modern tape technology also delivers comprehensive and pervasive encryption services to ensure the privacy of data placed on the medium.
- Tape technology uses less floor-space, less environmental conditioning and less energy than other storage media. It is considered “greener” than other storage media and has been demonstrated to reduce the cost of archival storage to a level well below that of disk or flash based archival repositories
- Tape is more efficient than networks as a transport for large quantities of data. Moving 10 TB across a T1/DS1 service will take over one year at rated speeds; assuming that a business can access and afford one, an OC-192 wide area network service will move the same amount of data over distance in approximately 2.25 hours. For bulk transfers of data, no network services can beat tape storage and transport – which has recently garnered the name “cloud seeding” by leading cloud storage proponents.

According to Horison Information Strategies, a well-disciplined tiered storage infrastructure will leverage a tape tier to store between 40 and 60% of all data. Without the use of tape in the storage infrastructure, storing a petabyte of data will cost between \$1 and \$1.5 million; with tape in the mix, costs can be as low as \$587,000.

Going forward, the capacity, durability, performance and cost per TB metrics of tape only show signs of improving. IBM, working in conjunction with tape media makers such as Fujifilm Recording Media and Sony Corporation, has already [demonstrated significant areal density advances for tape media in the laboratory](#). In August 2017, tests with different media coatings yielded areal densities for tape that exceed current capabilities by 20x. These results lend credence to the view, espoused by IBM thought leaders and others, that the runway ahead for tape capacity growth exceeds that of solid state, magnetic disk and optical disc – combined!

BUSINESS VALUE CASE

The technical capabilities of tape are well documented, but may not resonate fully with non-technical business managers. They are less interested with the number of tracks per inch or the transfer rate of

the latest technology than they are with how the technology will help contain cost, reduce risk and improve productivity. Savvy IT practitioners know that technology initiatives need a business-savvy narrative to help convince those who hold the purse strings to allocate budget where it is needed.

In the case of tape, it doesn't help that so many in the storage industry have repeatedly prophesized the end of tape technology. The punchline in the IBM development labs is that tape just celebrated its 65th year of service as a data storage technology, and its 64th year of being incorrectly cast as deceased or soon to be. The truth is quite different, however.

Tape, or at least IBM Tape, is enjoying a renaissance at present, with increased sales of tape-based storage capacity today than in any prior year. The technology has repeatedly profited by "borrowing" technology from other storage types (such as GMR read/write heads, perpendicular magnetic recording, and PRML algorithms from the disk industry) to increase capacity and performance characteristics and has benefited tremendously from recent ease-of-use improvements, including the delivery by IBM of the Linear Tape File System (LTFS). LTFS is a self-describing file system for tape that enables the portability of tape-based data from system to system without the need for proprietary file system or object models: LTFS formatted tape is the equivalent of a USB key for massive quantities of data.

In short, tape has consistently defied the rumors of its own demise and demonstrated its value within contemporary data storage infrastructure. It can deliver demonstrable business value in each of the three categories that business leaders consider: cost containment, risk reduction and improved productivity.

Cost-Containment

Depending on the analyst one consults, storage infrastructure is currently costing organizations between 33 and 75 percent of their IT hardware budget annually. While smaller and medium-sized firms have in recent years outsourced some storage to cloud service providers, this shift has simply driven the acquisition of more storage capacity, including capacity on tape media, to the cloud. For this reason, it is not surprising that cloud vendors and large enterprise data centers have become the proving ground for storage technology and for storage strategies aimed at cost-containment. Tape is not necessarily competing with the cloud – it is enabling the cloud!

Storage costs are a function not only of acquisition costs, but also of intelligent management of infrastructure resources, services and management. Any Total Cost of Ownership analysis will look at the cost to acquire and deploy storage resources and also at the cost to operate storage to derive the greatest value from the investment.

Year over year, the demand for more storage capacity has been accelerating. Leading analysts projected annual growth rates in capacity demand ranging from 40% to 650% (the latter being demand rates in heavily virtualized server environments). At this rate, large enterprises and cloud computing organizations, the proverbial tip of the spear in IT trends, are looking at a requirement to host nearly 60 zettabytes of new data by 2020 and 160 zettabytes by 2025.

Storing this much new data will be impossible using flash and disk storage alone. One recent analysis made by a leading cloud architect placed the total annualized output of the flash storage industry, in terms of media capacity and factoring in all anticipated improvements to the capacity of the components and devices that the industry produces over the next three years, at only 500 exabytes. The comparable total output of the hard disk industry, factoring in current improvement roadmaps, will

be about 780 exabytes per year by 2020. Taken together, less than two zettabytes of capacity will be delivered by the leading random access media producers in time to meet the hosting demands of upwards of 60 zettabytes of new data. This will create a significant storage gap that can only be filled by tape.

Organizations have been attempting to forestall the storage gap in their environments for at least the past decade. Data reduction and compression methods have been introduced to squeeze more data into the same amount of storage space by reducing replication and eliminating waste. Additionally, organizations have been moving from RAID-based data protection to older error-correction coding methods that do not require multiple copies of data to consume more storage space. However, even these approaches need to be viewed pragmatically as tactical, rather than strategic, solutions to the coming storage gap.

Resolving the storage gap will require two things: (1) tiered storage, including tape, to maximize available storage capacity and (2) improved data management, probably to include the application of cognitive data management algorithms to storage, to ensure that the right data is deployed to the right storage based on access characteristics, business value and cost to store. IBM developed many of the foundational models for storage tiering and is working with energy to develop the cognitive data management capabilities that will make the best use of storage from a cost-containment perspective.

Since it is impossible to bend the storage capacity demand curve, it is necessary to drive inefficiency out of storage wherever possible, and to optimize the allocation of storage resources and services to data based on business-facing policies for data governance. This, in the final analysis, is the only strategic way to contain storage costs, and tape must play a role.

IBM thought leaders have taken to describing the infrastructure model using the analogy of a Pittsburgh steak: with hot data on the edges (data stored on disk and flash close to the applications that use them) and a “cold center” (tape storage used to store data that needs to be retained despite low re-reference rates). Making this metaphorical infrastructure happen will be even more important as companies extend their storage infrastructure across multiple sites, including company-owned data centers and multiple cloud service providers.

Bottom line: IT professionals will need to leverage the cost-efficiencies of tape-based storage and hierarchical/tiered storage management if they have any hopes of “bending the storage cost curve.” The latter is an objective that senior management wants IT management to pursue.

Risk Reduction

The second component of a business value narrative required by the front office to justify tape storage (or any other IT initiative) is risk reduction. Risk is conceived not only as the risk of downtime, a common interpretation of the expression in most data centers, but also as investment risk or reputational risk. Senior management does not want investments to be made in technology that is suddenly obsolete and fails to return its investment. Similarly, they do not want to end up on the front pages of a financial news outlet beneath an embarrassing headline regarding their mismanagement of data. These many dimensions of risk can provide a rich forum for discussing the benefits of tape.

Clearly, given its history as a backup media, tape has a strong story to tell in downtime avoidance. The only way to protect data is to make a copy and store the copy at a location sufficiently distant from the original to prevent it from falling prey to the same disaster potential (aka, flood, hurricane, earthquake,

nuclear or chemical event, etc.). One needs to look no further than the recent weather events in North America to understand the importance of data separation at great distance. Tape has long been used as a portable medium for backup data, providing the “air gap” discussed above to keep data safe from malware, viruses, and hackers, as well as the panoply of natural and man-made disasters.

Tape is a very resilient storage medium with a hard error rate expressed as $1.00E+19$ -- or one error in every 11,102.02 petabytes of recorded data. Enterprise hard disk, by contrast offers $1.00E+16$ or one error in every 11.10 petabytes. Put simply, the uncorrectable bit error rate of tape media is 10,000 times less than the BER of hard disk.

Steady improvements in both tape media and drive technology and in robotics and enclosure technologies have bolstered the resiliency of tape well beyond other media, making it a destination for roughly 70% of the backup data in the world. Tape is also the preferred destination for most digital archives because of its low bit error rate and high durability of about 30 years.

Tape also reduces risks of data corruption, deletion, or unauthorized disclosure by implementing both an “air gap” (separating data on tape from on-line data stores) that can prevent virus infections and ransomware attacks, and pervasive encryption to prevent unauthorized access and disclosure. Tape encryption was pioneered by IBM and today offers risk reduction value without performance penalties.

However, risk reduction isn't just about resilient or encrypted media. Tape, particularly tape formats from the Linear Tape Open consortium, offer investment protection, as well. LTO tape is about to move into its eighth generation. LTO-8 will be able to read/write LTO-7 media. This backward compatibility enables tape to provide a comfortable hedge against early obsolescence and enables firms to add to their tape storage infrastructure with the latest technology without requiring a wholesale re-recording of all existing recorded tape media.

Finally, and perhaps most importantly, tape's risk reduction value extends to data management and compliance with legal and regulatory restrictions and mandates around information governance. Most regulatory and legal mandates around data storage have to do with either the retention period required for certain data or with the privacy and protection requirements for certain data assets over time. An unprotected data loss or disclosure of data, or the destruction of data before the end of its mandated retention period, can lead to the adverse consequences that business managers seek to avoid – whether an embarrassing headline or a shareholder or customer legal action.

Used in a tiered storage system, tape storage establishes a reliable, low-cost, long-term archival storage platform on which data can be hosted securely for up to 30 years. Numerous studies testify to the lower cost of this data storage archive when compared to either flash or disk-based archiving platforms, and leading cloud vendors, in acknowledgement of these factors, are beginning to deploy tape in a big way as an archival storage service. Consumers seeking risk reduction should be sure to consider tape archiving for their own infrastructure and should seek comparable services and platforms from the cloud service providers that they elect to use.

Improved productivity

The third part of a comprehensive business value case for tape storage is related to productivity. Productivity is, again, a word with different meanings. Clearly, LTO-8 tape delivers faster access and faster throughput than its predecessors, so backups and other tape read/write workloads are more efficient and performant.

In addition to the speeds-and-feeds improvements delivered by LTO-8, the increased productivity enabled by the technology is actually derived from the proper application of the right storage resources and services to data based on business policy. By itself, no storage technology ensures productivity gains for an organization. However, tape can certainly help businesses to achieve other goals, including more agile response to market opportunities and improved efficiencies in workload processing.

Tape is too often pigeon-holed as archival media. While it certainly provides an excellent platform for archive, planners need to think about what they mean by archive to unlock the real value of tape. Archive is thought to be the right fit for tape because it is written and read sequentially, rather than randomly. Random reads and writes are required by databases and other high performance workloads where tape is seldom used. However, other workloads, such as streaming services (streaming movies or videos from a central repository) are able to leverage tape to great effect.

For example, a media streaming company stores exabytes of video clips that can be requested by users. To play a file from disk or flash, the request is simply transferred to the storage device which moves its random access components to the start of the requested file and begins streaming the data. This instantaneous service sounds like the best fit for the data and the application, but in fact a sequential medium like tape may be better. With tape, the request to stream a specific file may require the selection of a tape from a shelf of the library, the actuation of a robotic “picker” to grab the selected cartridge and to place it into a tape drive, then up to 40 milliseconds in an LTO drive to access the start point of the selected file. In short, depending on the number of tapes and robots and drives, starting the streaming of the selected file might require up to two minutes. What some vendors of movie streaming services have done is to monetize this “delay” to the start of the file by playing an advertisement from flash or disk while the requested object loads. The data is in an archive for low cost retention and ease of playback, but the monetization of the file access delay has produced another line of revenue for the business.

In another case, some motion picture production companies have chosen to record their daily shoots to LTO tape rather than to flash. Doing so enables them to use physical mails and delivery services to forward their daily shoots to central distribution or editing services. Not only is this more dependable than sending large quantities of data across shared wide area network connections (assuming they are available at shooting locations), they also are able to save insurance costs by using tape.

Other companies have leveraged the native encryption that is made possible by current tape vendors, including IBM, to pursue business in locations previously thought to be too laden with security risks to provide business friendly markets. By encrypting data collected at branch office operations or work sites, then moving data via encrypted tape, to the primary data center or cloud, organizations are able to enter into markets that were once closed to them.

For its own part, tape technology has become far more user friendly. Enclosures and libraries provide touch-free or “black box” repository services requiring a minimum of operator training or skills development. Innovations like IBM’s Linear Tape File System (LTFS) have reduced the need for complex data protection or data archiving software suites by enabling a simple file copy operation for recording data stored on other storage media directly to tape. LTFS tapes carry their own indices and are “self-describing,” so getting to the data on an LTFS-formatted LTO (or enterprise) tape cartridge is as simple as inserting the tape into another library or autoloader capable of reading LTFS-formatted LTO media. So, from an operational perspective, today’s tape technology has never been easier to use, to provision, or operate.

CONCLUSION

LTO-8 technology will build on over 65 years of tape innovation and development from vendors such as IBM. In addition to providing the core drive technology, IBM works with its business partners and its customers to ensure that tape adapts to fit the technical and the business requirements that challenge data storage today.

Going forward, higher capacities, higher throughput rates, greater areal densities and continued improvements in consumability and ease of use promise to keep tape an important component of strategic storage architecture. As Mark Twain once said: “The news of his death was greatly exaggerated.” The same can be said for Magnetic Tape!