

Tape Storage Demand in 2005 and Beyond?



By Fred Moore, published in Computer Technology Review, 2005

Storage strategy conversations centered on magnetic tape normally debate questions such as:

- Can tape cartridge capacity keep pace with disk drive capacity growth?
- Will tape storage continue to be less expensive than disk storage on a per-gigabyte basis?
- Will tape remain the preferred backup and archive technology, given the advancements in disk drives?
- What will drive the future growth for tape?
- Will tape survive in the long term?

Three general classes of tape technologies exist. Enterprise tape solutions offer higher duty cycles, much higher MTBF (Mean Time Between Failure), virtual tape, more scalable, higher performance libraries and higher prices than their counterparts. Midrange tape offers tape automation and scalable libraries and is poised to receive its first virtual tape subsystem offerings. Low-end tape solutions offer a variety of auto-loaders to ease labor costs but have fewer features and less functionality than enterprise and midrange drives.

Magnetic tape industry revenue has been down slightly since the middle of 2000 when the global and economic downturn began, but the financial decline has been much less than that of the disk industry. In 2000, worldwide storage industry revenues were approximately \$31 billion for disk and just under \$5 billion for tape (drives, libraries and media). In 2003, the worldwide disk revenues were \$20 billion and tape totaled approximately \$4.5 billion. Tape library shipments are expected to increase just over 8% through 2009 while revenues are projected to grow over 7% during the same period. What are the primary reasons for these optimistic and rather significant growth forecasts given the negative growth rates since 2000?

Enabling Factor # 1: Tape Cartridge Capacity Accelerates

Though other devices were developed earlier, the first successful tape drive was launched in 1952 and the first successful disk drive appeared in 1956, both delivered from IBM. For over 40 years, the capacity of a round tape reel—and later the successor tape cartridge—was less than the capacity of a disk drive. By the late 1990s, this began to change as tape technology began to experience the most significant technological improvement curve in its history. As a result, tape-cartridge capacity is now growing faster than disk-drive capacity. This is particularly significant because disk areal densities have been increasing on an average of 60% annually since the early 1990s. Disk-drive projections indicate capacities of nearly 5TBs in 2013 and native tape cartridge capacities ranging as high as 10TBs by 2013.

Native capacities for tape cartridges are quickly approaching 1TB. The recently announced LTO Ultrium 3 data cartridge offers 400 gigabytes of native (non-compressed) capacity and has a data rate of 80 megabytes/sec. The Sony SAIT format has already attained the 500-gigabyte native capacity level yielding over 1TB per cartridge in compressed mode. Media life, once a concern, now ranges from 15-30 years for most enterprise and midrange cartridges. Just in case you missed it, things have changed significantly in the tape business.

Bottom line: Expect the capacity of tape cartridges to increase faster than disk drives for at least the next five years. This will continue to lower tape prices and make automated tape storage more appealing for WORSE (write-once read-seldom-if-ever) backup/recovery, archive, compliance, and fixed content applications. These applications represent the single fastest growth area for data storage demand.

Enabling Factor # 2: Tape Pricing Trends Downward

For years, tape was considered to be cheaper than disk storage on a per-gigabyte basis. In the late 1990s, after several years of disk price erosion exceeding 30% annually and relatively few technological enhancements to tape technology, the perception about disk pricing falling below tape pricing became more widespread. In the late 1990s, tape technology began to evolve more quickly than in the past. Historically, multiple tape cartridges or reels had been required to back up the entire capacity of a disk drive. By 2001, a single tape cartridge could back up multiple copies of the largest disk. The typical high-capacity tape cartridges now have native capacities of 200-500 gigabytes and a compressed capacity of 400 gigabytes to over 1TB using a 2-to-1 compression ratio. The largest disk HDA announced today has a capacity of 500 gigabytes but the most popular disk is presently the 80-gigabyte disk.

For typical working automated tape configurations, the price per gigabyte of automated tape ranges from one-fourth to less than one-twentieth the price of an equivalent amount of disk storage. Automated tape consumes less electricity than devices that spin and require cooling constantly further lowering the TCO. Automated tape presently offers the lowest price per gigabyte for midrange and enterprise systems, and technology roadmaps indicate that this is expected to remain so for the foreseeable future. Any viable optical or holographic option still isn't apparent.

Bottom line: Expect the price per gigabyte of automated tape to stay below and even diverge from the price per gigabyte of disk for at least the next 5 years. This will make tape storage a more viable, cost-effective storage solution for archive, compliance and in the traditional backup/recovery markets.

Demand Factor #1: Backup/Recovery

Tape storage has been the most popular choice for backup and recovery applications since the beginning of the IT industry. The use of tape and disk for backup/recovery increased after September 11, 2001. Disk mirroring, copies one and two, is ideally used for mission-critical applications and is the fastest way to return data to operational status from a hardware device or subsystem failure. Restore operations using mirroring occur in just a few seconds or less by simply pointing to a mirrored and executable copy. Mirroring is expensive as it can double disk requirements and should be used for mission critical data—but not for all data. On the average, about 15% of data is classified as mission critical. Mirroring alone is not enough to implement a successful data protection strategy. Hardware problems are not the only problems that storage administrators need to protect against. Mirroring does not help protect against the growing data corruption problems (hackers, worm, virus, intrusion, programmer or software errors) as it instantly spreads the corrupted data to the (n) mirror copies.

The storage industry is now realizing that a truly effective mirroring strategy

should be accompanied by point-in-time copies, also referred to as copies 3 through n, for data that enable a restore to occur from non-corrupted data prior to the time the corruption occurred. The most cost-effective solution for keeping multiple copies of files, volumes or data sets over a period of time is automated tape since the cost of keeping multiple copies of the same data on disk can quickly get out of control. Most backup copies of data on tape are not executable however and data backed up on tape must often be moved to disk before it can be used. It takes time to recover data from tape. Clearly, the usage of disk-based replication is increasing. However, tape backup should accompany disk mirrors to most cost effectively address the growing intrusion and corruption problem. Tape also will continue to accompany the use of disk for backup/recovery as the "T" in the popular D2D2T solutions market.

Bottom line: In the long-term, expect tape to accompany disk in a variety of ways addressing the backup market, possibly in a somewhat different role than in the past, as the primary recovery source for critical data is shifting to disk storage solutions.

Demand Factor # 2: Long Term Data Storage Requirements Soar

The highest growth area for tape storage and the storage market overall will be in managing growing pools of fixed content, compliance and archival data for long periods of time. Compliance and archival data demand is estimated to grow between 20%-70% per year depending on the business and now represents the fastest growing segment of the storage industry. Estimates indicate that the fixed content, compliance and archive hardware and software storage market will exceed \$10 billion in annual revenue by 2006. New retention periods for most compliance, archive and fixed content data can range from 3 years to ∞.

Compliance is driving data security to new levels. Security practices such as encryption and WORM (write-once-read-many) will begin to apply to stored data in addition to data in transit. The recent number of tape vendor announcements supporting the use of WORM for tape devices signals tape's future role will be more widespread in support of compliance and non-alterable storage requirements. One physical copy of compliance data will not be enough to ensure high availability; compliance data will also need to be backed up effectively in different geographical locations potentially doubling many storage estimates.

Bottom line: Compliance, archive and fixed content, the fastest growing segment of the storage industry, favors tape storage over other options as the price per gigabyte and overall TCO for tape is the lowest of any storage technology. Future tape roadmaps project this advantage will continue for the foreseeable future.

Will Tape Survive in the Long Term?

The magnetic tape industry is in its third era since the first successful tape drive appeared in 1952. As the year 2000 approached, growing concerns over the viability of tape were voiced. Since 2000 however, tape cartridge capacities have accelerated and are on an unprecedented growth pace as cartridges are expected to reach 2 TBs within five years. The first 1-TB cartridge was first previewed (but not announced) in 2004. The automated tape library market revenue is forecasted to grow as much as 7% annually through 2007 as backup and recovery issues and new, long-term archival and compliance storage requirements mount for large businesses and the SMB market. Architectural enhancements including embedded disk arrays for caching, tape SANs, WORM

and virtualization provide large-capacity, automated libraries with the capability to meet new data retention requirements with end-user pricing ranging from one-fourth to less than one-twentieth the price per gigabyte of magnetic disk storage.

Bottom line: Having answered many of the challenges it faced prior to 2000, the tape industry has successfully positioned itself for success in the foreseeable future.

www.horison.com