

Tape: The Zombie Technology

Tape plays a continuing role in meeting the storage challenge

May 12, 2009

by Mark Ferelli

Pundits who, in spite of all evidence to the contrary, declare that tape is moribund missed a wakeup call at the two-day executive conference sponsored by Fujifilm Recording Media USA, Inc. in early February. Far from a commercial, the total amount of valuable information was, in the words of one data manager, "staggering."

The seminar theme, Into Tomorrow with Tape Storage, Protecting the Data and the Bottom Line, was a broadly based topic that allowed presenters considerable latitude to deliver both plaudits and criticisms of the current generation of tape technology and the applications that work with them. Tape vendors were criticized for weak marketing and a poor performance perception. On the other hand, attendees were quick to identify the continued value of tape as an archival medium and to acknowledge its energy efficiency.

Expert Introduction

The first presenter of the seminar was John Teale, a veteran of IBM's Systems group and one of the early engineers involved with IBM 3480 technology (1983) and subsequent tape-drive developments all the way through LTO-4. He discussed the evolving role of tape in the overall storage environment. His first point was the most significant: data centers will always need an "ultimate insurance policy" protecting critical data. This need will never go away. Teale believes that "disk is the most unreliable storage device ever designed," and in IBM installs where disk is selected for backup, IBM strongly recommends that tape be incorporated behind the disk.

Teale also noted that tape technology is progressing more quickly than disk and is subject to fewer limitations. He cited a joint development program with Fujifilm making headway in the areas of head/tape interface compatibility, head and media magnetics requirements and integration, and the tradeoffs between tape dimensional stability and head/actuator designs required to achieve higher areal densities. One of the practical outcomes has been to prove the capability to store 8.0 TB of data on a single cartridge.

As far as future technologies that would replace tape, Teale doesn't see any for now, commenting that "holographic is and always will be the technology of the future".

Tape TCO

From the analyst firm The Clipper Group, David Reine presented the results of one of their reports. He compared SATA with LTO-4 in an analysis of total cost of ownership (TCO) between disk and tape.

David and his team arrived at remarkable conclusions considering the popular "disk über alles" hype common to the trade press. Even where data deduplication on disk is used, tape offers a TCO advantage of 5:1, with a significant addition of acquisition cost. Where deduplication is not present, the advantage of tape grows to roughly 23:1. Finally, he noted a dramatic 290:1 advantage for tape in terms of energy efficiency.

Instrumental, Inc., a consultancy with considerable Federal government involvement, examined the benefits of tape and how it could be maximized. Nathan Schumann asserted that tape is not yesterday's technology and never will be. He examined the proposition that tape is not needed and concluded that the proposition could not stand up to rigorous analysis. Finally, he stated that MAID (Massive Array of Idle Disks), often identified as an energy-efficient approach to disk array architectures, cannot currently replace tape.

Schumann's analysis demonstrated that tape errors were two orders of magnitude less than disk. In terms of scalability to bandwidth, neither tape nor disk overwhelmed, but tape still had an edge. He criticized mean time to repair for disk as opposed to tape, demonstrating that the problem was exacerbated in SATA disk and compounded by array configurations.

Tape drives, Schumann said, perform best when streaming data in large blocks, therefore managing block size can help optimize performance. He also cautioned that some of the "bad rep" that tape suffers might be due to problems with server memory bandwidth and PCI bus bandwidth. "Blame shifts to tape," Schumann noted, "but what is further up the I/O path could be affecting performance."

A Green Data Center

Greg Schulz of the StorageIO Group consultancy covered the ever-popular issue of the "green-ness" of the data center. Unlike many observers of the industry, Schulz noted that some of the discussion of energy efficiency in the data center is so much hype. More on point, he observed that energy approaches in the data center have evolved. The original emphasis was on energy avoidance: powering down, over-consolidation, reducing the amount of useful work and, in general, decreasing the amount of energy used. Closing the "green gap" has moved from energy avoidance to more energy efficiency, characterized by faster components that need to draw less power, increasing the amount of useful work, and decreasing net energy used.

An industry trend that Schulz called out was the tiering of resources, which results in balanced performance, high availability, the making of best use of capacity and the best use of energy. In the tiered structure, tape remains the most energy-efficient alternative. Schulz pointed to other alternatives, such as IRM (Integrated Recovery Management) and MAID, but neither is as efficient as tape. He referred to tape as a "zombie" technology, since it has often been declared dead but everyone either still uses it or at least should continue to use it in a sensible tiered-storage strategy.

Real World HPC

Sharan Kalwani from General Motors presented a case study based on his own experience struggling with high data turnover in his high-performance computing (HPC) environment. In the design of automobiles, high-performance computers grapple with extreme data turnover on disk and tape backup alike as new designs and simulation results replace the old. Latency between processors and disk or tape peripherals was a major problem in the engineering design group, with 50 percent of processing time taken up by communications.

The solution GM found to handle the latency problem was InfiniBand. InfiniBand is a switched fabric communications link; a bidirectional serial link designed to connect processors with high-speed peripherals. It boasts flexibility in signal rates, and links can be bonded together for additional bandwidth (similar to PCI Express). Sharan also underscored the benefits of tape for mass storage of the huge amount of data generated by HPC applications with the ability to effectively data mine from tape archives.

Best Practices

AT&T Mobility also presented a case study; John Menzik discussed the management of 11 petabytes of monthly backup. Their operation encompasses six domestic data centers in several geographic locations, and only has 11 staff members serving a client base of 4844 users. Currently, they have 107,000 tapes in offsite repository.

The challenges in Menzik's operation are considerable. Data is growing steadily, and the staffing level is low. Support for RMAN, the Oracle retrieval system, is a non-trivial undertaking. Finally, Menzik's department is struggling to give senior management a full understanding of the backup infrastructure and scope.

The bulk of Menzik's solutions revolve around management techniques that he instituted. Aggressive user feedback tools, clear goals and empowerment for his team, strict workflow processes, and a rigid chain of custody policy all contributed to his success. Notably, he had abandoned e-mail in favor of a ticketing system to initiate and pursue projects.

In terms of managing backup, data retention policies were focused largely on AT&T's legal team as the group best positioned to set them.

Menzik has developed a set of standards to judge whether the best handling of tape is being done. For example, he asks whether there is a balance between the recovery point objective and tape ejects. Further, he wants to be sure that tape media is being sent off site completely full. He monitors the speed backups to, and restores from, tape; media handling is considered key to ensure tape and data are readily recoverable. Media that is dropped should be pulled out of the rotation.

The executive believes that tape remains the long-term storage "king" and that tape use will continue in his installation while speed and transfer rates continue to increase. He believes that vendors for backup may make it more expensive to use disk storage for backup, but cautions that tape's cost per gigabyte should remain cheaper than other backup means. Latencies in backup have, in Menzik's experience, been more host issues than problems with backup hardware.

E-Discovery

Ernst & Young representative Heidi Stenberg gave a complete overview of issues in e-discovery, one of the drivers for enterprise data preservation. She discussed e-discovery and best practices issues.

Stenberg started with a discussion of some of the trends in storing information. For example, 93 percent of all information created during 1999 was in a digital form. Instant messaging (IM) is growing as a de facto communication tool; by 2013, IM will be used by 95 percent of all employees. The dependence on e-mail is overwhelming ... Americans send 2.2 billion e-mail messages daily.

With this foundation, she turned to discovery, which is a civil litigation term for identifying, preserving, collecting, and producing information relevant to the matter under dispute. Stenberg noted that the definitions and practices discussed were equally applicable to audits and other kinds of investigations in addition to civil litigation. She said that the challenges of e-discovery were founded in complexity. Data being managed in the enterprise is easily duplicated, easily disseminated, stored in complex backup systems, is easily portable and high in volume, and often residing on archival data tapes. Identification, preservation, and collection are difficult but essential; collection, review, and production are expensive, but failure to produce is far worse in terms of legal costs and court fines.

Even more relevant to IT managers at the conference was the definition of a legal hold, a communication that suspends an enterprise's normal processing of records. These holds impact the data life cycle, and suspension of data deletion can overwhelm systems. Retrieval of legacy data or from legacy systems also takes time and ties up revenue-generating employees.

Stenberg recommended a discovery response program emphasizing the understanding of an enterprise's universe of data and the implementation of a discovery response team. This team would include stakeholders from IT, legal, and business units. She also recommended the construction of an ESI map, an overview of all electronically stored data for use in the "meet and confer" sessions demanded by courts.

Ensuring Reliability

The only presentation by a vendor at the conference was an examination of tape system reliability from David Cerf of Crossroads Systems. He provided a short historical look at Fibre Channel connectivity, and a tape technology booster's reminder that the tape industry needs to be more aggressive in broadcasting tape's real-world advantages.

Cerf carefully examined the reliability issues in tape subsystems, including load/unload latency and the media wear and damage possible in start/stop. He advocated large data transfers and staying above the streaming rate. Cerf cautioned that small transfers with frequent loads/unloads were major time consumers.

In the ideal tape world, Cerf suggested that:

- Tapes be loaded and fully written with large blocks
- Data be transferred at compression rates
- Tapes be unloaded immediately after write operations
- High-quality media and drives be used
- The entire system be monitored
- Written media should be verified regularly

In light of the importance of monitoring, Fujifilm is now offering tape library site surveys to their customers that can help prevent data recovery failures, save time and money, and improve overall data tape performance making use of some of the tools that Cerf and Crossroads Systems offers.

The conference concluded with a lively, moderated panel discussion featuring interactive discussion between the attendees and the conference presenters that covered such topics as the present and future of the tape industry, the importance of actually managing backups, and other issues arising from the presentations.

Overall, it was clear that backups are managed functions that cannot be deferred or neglected. Tape technology is recognized by enterprises large and small as the most cost-effective and energy efficient means to store and restore critical archives. If tape is dead, as some pundits repeat wearily, then it is as Schulz suggests, a zombie technology that keeps rising from the crypt. It's alive!

Mark Ferelli is an independent technology journalist, and commentator. He specializes in mass storage and was editor-in-chief of Computer Technology Review, Storage Management Solutions, and other storage-specific publications.
