

Sony Electronics Inc.

1 Sony Drive

Park Ridge, NJ 07656

(800) 829-7669

www.Sony.com/Datasystems

Sony's SAIT: The Atlas Holding the Enterprise's Tape World

Preface

The continued growth in storage demand is putting additional stress on the storage infrastructure in general — and on tape systems in particular. One reason is that growth of new applications and storage on disk has a multiplier effect on tape. Backing up an extra gigabyte (GB) of disk storage does not just result in a one-to-one increase in tape storage, but rather a one-to-many increase owing to keeping multiple tape generations of the same data.

Moreover, when data ages its frequency of access goes down, and it effectively changes to read-only, fixed-content data. Although some of this data is so old and of so little value that it can be safely archived at a remote site, companies tend to retain most data in a form and a location where it can be restored in a reasonable time. This data still has value, as in when it is used to analyze customer behavior or to review regulatory compliance. When low frequency of access, low cost, and high capacity all play a role, tape is likely to act as a data repository.

The need to store data on tape, therefore, seems insatiable. The need to push tape storage to ever-higher capacities and performance is key for many IT organizations. Although the road map for open systems-based tape technologies indicates the requirements will be met in the future, only Sony's Super Advanced Information Technology (SAIT) meets that need now.

SAIT technology illustrates a high-end tape technology that meets performance and capacity challenges for the tape-devouring requirements of many enterprises. This Aberdeen *Profile* examines SAIT technology, which builds upon Sony's well-established AIT recording technology to deliver dramatically higher capacity for a standard half-inch-form-factor tape cartridge. SAIT technology delivers very high capacity and performance while enabling IT organizations to maximize their limited budget dollars.

Executive Summary

Aberdeen research shows that storage consolidation and improving the backup/restore process are still high on the agenda for IT organizations. As the move to storage networking continues, many smaller libraries are consolidated into one larger capacity library. With half of a terabyte (TB) native — and more than one TB for compressed data — on a single data cartridge, SAIT technology can deliver a large capacity tape library with fewer slots than is otherwise needed. That means operational savings in managing fewer pieces of tape media as well — and where space is tight, a smaller footprint can save big money.

In a time of shrinking backup windows and the high number of backup jobs that fail to run to completion, tape drive performance counts. Performance is not only a matter of rated drive speed — at 30 MB/s, SAIT drives compare favorably with linear tape technologies. SAIT's helical scan technology does not suffer from the stop/start problems that may afflict linear tape technologies that move tape at high speed. The result is an improved and more efficient backup process.

Sony introduced SAIT in order to meet capacity and performance requirements for the high-end, open systems, enterprise-class tape automation technology market. With SAIT, Sony widened its AIT tape from 8 mm (about one-third of an inch) to one-half of an inch. An SAIT media cartridge stores more than twice the native capacity in the same space as linear tape technologies by using the same media length and cartridge size of half-inch linear tape technologies.

Sony can supply each of the three elements that make up a tape system — the media in the form of the SAIT tape cartridge, the SAIT tape drive, and the robotic tape automation system — in the CSM PetaSite library series. Matsushita Kotobuki Electronics Industries, Ltd. (MKE) is a second manufacturing source for tape drives, and Matsushita Electric Industrial Co., Ltd. (MEI) — which is well-known for its Panasonic-brand products — is an alternative source for SAIT media. Leading open system tape automation vendors have already announced support for tape automation solutions that contain SAIT technology. IT organizations benefit from having alternative sources for tape libraries and a second source for tape drives and tape media.

Sony uses Advanced Metal Evaporated (AME) tape technology, which has an expected shelf life of about 30 years or more. Sony intends that future generations of SAIT drives will be backward compatible with the current generation of SAIT drives. IT organizations can, therefore, protect their tape media investment technology without the risk of having to migrate the data in the future.

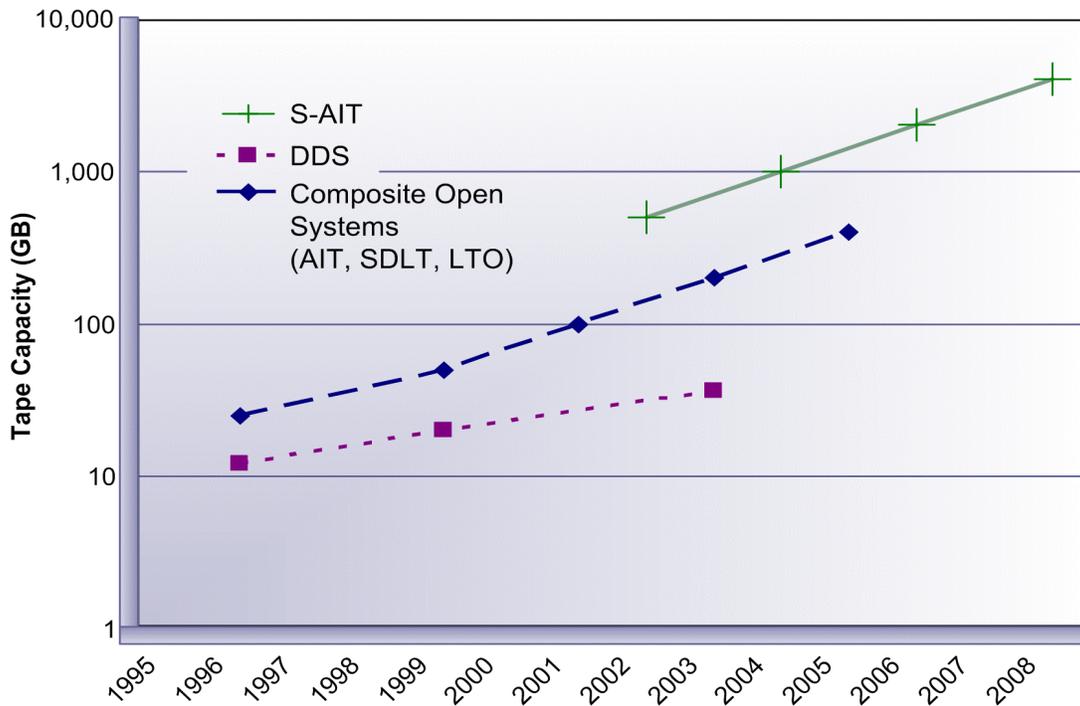
SAIT technology is useful for meeting the criteria of an IT organization for consolidating (space-wise as well as in the number of tape libraries), for storing bulk data that no longer needs to reside on high-performance disk but needs to be available, and for providing cost savings and operational benefits of managing fewer pieces of media.

Deep Density: Key to the SAIT Architecture

The dramatic improvement in the areal density in Gb/cm² for both disk and tape is likely to continue in the foreseeable future. Winding 600 meters of tape into a standard cartridge yields a constant volume of tape whose volumetric density will increase proportionately to areal density (Figure 1). That volumetric density is deep — in the sense of large in quantity — density, as the native (uncompressed) capacity will reach one TB within a couple of years. Deep density brings about both operational (fewer pieces of tape media to manage) and cost-savings benefits (by lowering the overall cost of a tape solution for the same capacity).

As a helical scan recording tape technology, SAIT technology has an inherent — and, therefore, sustainable — deep-density advantage over linear recording tape technologies. Helical scan recording can record many more data tracks than linear, longitudinal recording. Helical scan technology uses a stable spinning drum with read/write heads recording data on slowly moving tape. By contrast, linear recording technology moves tape very quickly past fixed-position read/write heads. Slower moving tape does not have to compensate for the degree of skew tracking problems of faster moving tape; therefore, helical scan technology can more accurately record tracks closer together — thus, a deeper density.

Figure 1: Tape Capacity Trends



Source: Sony and Aberdeen Group, September 2003

Overall, the deep density of SAIT technology yields operational benefits through the simple expedient of managing fewer pieces of media, meaning less chance for error in tracking and handling the media.

The deep density of SAIT media also delivers improved cost savings. The media itself has a lower cost per GB. Moreover, deep density means that more TBs can be stored in the same physical space. The use of SAIT media results in a lower cost per GB for large tape automation solutions (i.e., consisting of a tape library with a number of tape drives and a considerable number of tape media slots). This is because the tape automation takes up less physical floor space (dollars-per-square-foot savings) and requires fewer tape slots (smaller cost tape automation).

Even Deeper Density through Compression

Unlike disk, tape typically contains compressed data. Sony states that its Adaptive Lossless Data Compression (ALDC) can achieve up to a 2.6: 1 data compression ratio. The 500-GB native capacity could translate into as much as 1.3 TB compressed, and the native speed of 30 MB/s would be 78 MB/s compressed. Compression is valuable in keeping tape-solution cost per GB low, relative to disk solutions.

Balancing Performance and Deep Density

Even though the native performance speed is relatively fast compared with, say, AIT, the density has increased even faster. Although the apparent result is that the time needed to restore a full tape is typically longer, the sheer volume of data restored for a given time period is far greater.

Moreover, in most cases, the time needed to recover a full tape is not the key factor. Only time-sensitive production data, such as an online transaction processing system, needs to be recovered quickly — and, typically, that is not the majority of data at a site. If the quickly recoverable data is a single file or a database that needs to be restored sequentially, then SAIT's performance is the key, as the overall restore time would be less regardless of how much data would be restored.

In the extreme case of a bare metal restore at a new site after a major disaster, the management process would be very complex and, having fewer pieces of media to manage, would probably be an overall advantage. Bare metal restores are going to take considerable time anyway, and minimizing the possibility of media management errors would seem to outweigh the longer times needed to read a full tape.

No Shoe Shines Necessary

Getting a backup process to run smoothly can be very frustrating to an IT administrator. If there is an overall data-rate mismatch between what the backup server is trying to deliver and what a tape drive, even with buffering, can handle smoothly, the result can be a lot of stop/start action that is reminiscent of a shoe shine. After a tape drive has stopped recording, it has to backup and accelerate forward to the last recorded position and continue recording.

“Shoe shining” can significantly slow the overall performance of a tape drive from its rated speed. Moreover, shoe shining is hard on the tape media and the tape drive mechanism. The impact of shoe shining is very limited with SAIT because helical scan tape has a very low tape velocity (about one inch per second) compared with typical linear recording tape technologies (more than 100 inches per second). With helical scan technology, the read/write heads move; with linear recording technology, the read/write heads are fixed, and the tape media travels at a high speed over the heads. Thus, SAIT requires minimal repositioning of the tape and, therefore, little performance and wear degradation caused by shoe shining.

Shoe shining is not the fault of either backup software or the tape drives. The problem is that pushing data through a backup server over a network at a preset speed is difficult to do — if it is even possible. One solution is to put a disk-based system in front of a tape automation system for the sole purpose of acting as a cache to buffer speed changes and prevent shoe shining. By utilizing helical scan recording, SAIT significantly reduces the shoe-shining “disease” for the system administrator, who now does not have to spend extra money for a disk array to treat the symptoms.

Getting in Compliance

As enterprise customers determine their electronic-records-compliance strategy, one clear result is that compliant records will be stored in a nonrewritable and nonerasable format. Disk and tape systems that comply with those conditions would meet the standards to house compliant data records. Sony already offers a compliant product with its Write-Once, Read-Many (WORM) capability for AIT-2 and AIT-3, and intends to extend that capability to SAIT.

Each SAIT tape cartridge contains a Memory-in-Cassette (MIC) flash memory chip that contains media information and statistics, such as serial number and media usage. Each MIC chip for a WORM tape cartridge contains an unalterable, factory-written code that designates the media as write-once. The detection of the special MIC encoding by an SAIT drive forces the drive firmware to write data only to unrecorded portions of the tape media, satisfying the nonrewritable condition. As SAIT tape drive’s special firmware also prevents data overwriting on WORM-specified media, it satisfies the nonerasable condition as well.

Having a WORM option for a particular tape technology will quickly become a necessity, not just a “nice to have.” Executive management is increasingly interested in reducing personal and litigation exposure to ward off the risk of noncompliance.

All SAIT drives will be “multifunction” and will perform standard tape operations or WORM-required operations, depending on the media type used with the SAIT drive. IT does not have to make a future investment in WORM-enabled tape drives and has the flexibility to use the same drive for any purpose.

The CSM PetaSite Library Series Illustrates SAIT Technology in Action

When talking about quantity of storage, only recently has the unit of measurement gone from gigabytes to terabytes. Now, the discussion also covers petabytes (PBs) — *three* orders of magnitude greater than TBs. Sony recognizes the trend in this direction with its PetaSite tape libraries.

The SAIT CSM PetaSite libraries come in three *native* capacity sizes — the CSM60 PetaSite systems, from 10 TB to 30 TB of capacity; the CSM100 PetaSite series system, from 42 TB to 54 TB; and the CSM200 PetaSite system that scales from 96 TB to 1.5 PB. The basic storage system is a standard 19-inch rack base footprint that contains the library robotics, as well as tape drives and tape cartridge slots. Larger systems contain extension consoles that consist of a cartridge-only console, a drive and cartridge console, and/or a redundant robot console (to ensure no single robotics point of failure).

A CSM PetaSite library with SAIT fits nicely into the current roles that tape plays, as well as in the future roles of tape. For many enterprises, data protection will continue to be tape's primary role. However, as mid-line (i.e., cost-effective, ATA-based) disks play a greater role as the first line of defense in recovering from local data outages, tape automation will serve as a second line of defense, as well as a source for generating tapes for off-site tape rotations and archiving. A CSM PetaSite library will be particularly effective in staging from disk to tape. This staging will assume a non-production-disrupting role, as the staging can be done continuously in the background. The high performance and capacity of SAIT drives in a CSM PetaSite Library require fewer drives and fewer tape cartridges than other open systems tape technologies can accommodate.

However, tape in general — and a CSM PetaSite library, in particular — can play a more active role as well. Fixed content — data that is effectively read-only — actually consumes the majority of today's data and will consume an even larger percentage tomorrow. This includes structured data, such as closed transactions and data warehousing information; semi-structured data, such as e-mails and their attachments; and unstructured data, such as medical images and videos. Much of this data has a very low frequency of access requirements and sub-minute response times are adequate, all of which are characteristics of duty cycle and responsiveness that can be easily be met by a CSM PetaSite library.

Fixed-content data tends to be bulky in the sense of individually large files (e.g., medical images) or as the sum of a large number of small records or files (e.g., telemetry data). These characteristics play to the capacity strengths of a CSM PetaSite library.

Data has to be "online" but "online" and "disk" are not equivalent. Online does not refer to response time (as long-term users of the Web can attest), but whether or not a business user (primarily in contrast with an IT administrator for data protection) can retrieve and use data for a business application. Sony enables a CSM PetaSite library for this purpose through the use of its PetaServe Hierarchical Stor-

age Management (HSM) software, which migrates data from disk to tape transparently to the user and supplies files back to the business user quickly on-demand. Unlike the original concept of HSM, which was primarily time-based migration, PetaServe HSM software delivers policy-based migration, including frequency of access over a period of time, size of file, and type of file, in addition to time.

Whether or not the need is to get compliance data, an MRI for a patient visit, a really out-of-date presentation that may contain useful information for a new marketing campaign, or a long-term customer buying pattern analysis, keeping that data on a CSM PetaSite library is more cost-effective than keeping it on disk.

The Investment-Protecting SAIT Road Map

Sony has made each generation of AIT drives compatible with previous generations, and it is continuing the same strategy with SAIT drives. The expected shelf life of AIT and SAIT media of 30 or more years is also critical and is especially important for compliance data using WORM, in which writing to a new piece of media would have to be conducted under strict controls. With the Sony products, the media would not be rewritten very frequently (one can never say never) because of both backward capacity and long shelf life.

Sony also plans an alternating strategy in which the latest generation of AIT is introduced ahead of the next SAIT generation. A large portion of the technology needed for a new generation of SAIT will already be in the field so that the transition to the new SAIT generation should not be that difficult.

Each new generation of SAIT drives and media should have about twice the capacity and performance of the prior generation (Figure 1). Therefore, a native TB-per-tape cartridge is not that far from reality and not just as a laboratory demonstration. Organizations that recognize the need to plan their storage architectures for at least the next three years can feel comfortable about putting the SAIT road map into their storage infrastructure equation.

Aberdeen Conclusions

SAIT is, in reality, the Atlas of tape storage, with each tape cartridge capable of 500-GB native. That fits with the unceasing IT trend to fit more into the same size form factor, be it servers or storage. Deep tape density translates into greater cost-effectiveness through more (capacity) through less (fewer number of slots and footprint for a tape library). Moreover, physical costs savings are augmented via operational savings — reducing the IT administrator's burden. That burden is more a function of how many objects (i.e., pieces of tape media) have to be managed than of overall quantity to be managed (i.e., TBs of data on tape).

Building on an already widely used data tape technology (AIT) has minimized the design considerations for the new tape (SAIT); therefore, the new tape technology is more like an old friend than a new one. The increase in native capacity over other open systems tape technologies is dramatic. Although the drive speed is

quite good, enterprises may not want to take the time to restore a full tape at full speed. For applications that stress the performance of SAIT rather than its capacity, a shorter length tape would be welcome. IT could mix shorter tapes for applications with a high service-level requirement for fast restoration. That should be the minority of the data in a data center. The rest could be on longer tapes.

SAIT shows the future direction of tape — storage-demand absorbing capacity, backup process improving performance, and the flexibility to also support compliance data, all contained in a form factor that can be utilized with standard tape automation systems. Aberdeen advises that organizations that are evaluating high-end, open systems tape technologies to thoroughly examine SAIT.

To provide us with your feedback on this research, please go to www.aberdeen.com/feedback.

*Aberdeen Group, Inc.
260 Franklin Street
Boston, Massachusetts
02110-3112
USA*

*Telephone: 617 723 7890
Fax: 617 723 7897
www.aberdeen.com*

*© 2003 Aberdeen Group, Inc.
All rights reserved
September 2003*

Aberdeen Group is a computer and communications research and consulting organization closely monitoring enterprise-user needs, technological changes, and market developments.

Based on a comprehensive analytical framework, Aberdeen provides fresh insights into the future of computing and networking and the implications for users and the industry.

Aberdeen Group performs projects for a select group of domestic and international clients requiring strategic and tactical advice and hard answers on how to manage computer and communications technology. This document is the result of research performed by Aberdeen Group, which was underwritten by Sony Electronics, Inc. Aberdeen Group believes its findings are objective and represent the best analysis available at the time of publication.