THE CLIPPER GROUP
Explorer™

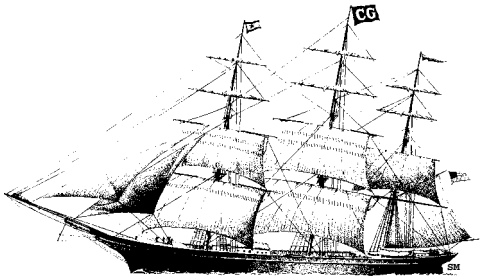

Navigating Information Technology Horizons℠

Published Since 1996 Report #TCG2004040 May 7, 2004

Tape Drive Selection — A How-To Guide

Analyst: David Reine

Management Summary

Tape is an essential medium for the data center. It is as important to the secure operation of the computer room as tires are to the safe operation of an automobile, with as many or more choices. Both are commodity products. Goodyear sells tires that will fit on everything from a Chevrolet to a Lexus. Bridgestone sells tires that fit a Toyota or a Mercedes. Likewise, Quantum sells tape drives that will operate with anything from a Dell PC to an IBM eServer. Exabyte sells tape products that will work with any UNIX box. You can even buy an IBM 3592 drive in an IBM frame to work with a StorageTek library![1] There are a significant number of variables associated with the selection of either. Furthermore, it is not always obvious which variable is the most significant.

The first option that the tire shopper faces is the warranty mileage. Does he want a tire that has a 20,000, 40,000, or 50,000-mile warranty? The first reaction might be to select the tires that will last the longest. However, they will probably cost the most! It is important to understand how many miles the car will drive over the next year. It is important to calculate how many more years that the driver will continue to own this car. If you drive 15,000 miles a year and want to keep the car for two more years, then a tire rated at 40,000 miles could be perfect. Moreover, the shopper also needs to evaluate how the tires will perform, given the specific expected driving conditions. Do you need high-performance tires or standard highway tires? Do you drive in snow? Will the tires hold the road well in slippery conditions? There are many variables in the tire selection process. The same is true for tape.

It is safe to assume that every data center today uses tape. However, for what purpose does the data center use tape? Is it a backup drive? Is it a data processing device? How much data needs to be stored? How long is the backup window? Is the tape drive used in a pure open systems environment or is there a mainframe lurking in the shadows with connection requirements? Moreover, the question that every CIO is asking is, "How much more will my data expand over the coming years"? The requirements placed upon every enterprise by laws such as Sarbanes-Oxley and the Patriot Act have forced these enterprises to double and triple, if not more, their backup and data protection projection.

There are many formats for tape: *SDLT*, *LTO*, *SAIT*, as well as *3592* and *T9940B*. This issue of **The Clipper Group Explorer** will take an in-depth look at the various options that the different formats present to the data center and try to help you in your selection process. If you are looking to review or upgrade your tape environment, please read on.

[1] But see box on page 5.

IN THIS ISSUE

Today's Data Center Environment

Enterprises today are gathering an ever-increasing store of information in order to gain a competitive advantage in their industry and to increase profitability. Some of this information is associated to mission-critical applications running the core of enterprise business. Other data relates to improving customer and vendor relationships. Enterprises retain even more data in order to protect themselves from potential government action. The data center staff must protect all of it in order to assure the financial health of the enterprise.

Legislation, such as Sarbanes-Oxley and the Patriot Act, require enterprises to preserve their financial data for long periods. Health protection acts such as HIPPA place special import not only on the retention of health records, but also on the privacy of such information. In January, we looked at the role that Information Lifecycle Management (ILM) plays in keeping tape alive as a medium in the data center[2]. At that time, we examined the generic use of tape as a backup and recovery medium as well as a data processing medium. We identified some of the reasons for this rise, such as a renewed importance on standard backup-recovery processes, and protection from disasters such as 9/11. We also identified additional topics that we hoped to cover in later bulletins. In this second issue in a series on tape, we examine the different formats and architectures that are available to the Information Technology (IT) community and try to identify which formats may provide superior performance in your specific environment.

New Tape Architectures

The turn into the 21st Century has brought with it a renewed interest in the development of new, improved tape formats to archive data and to insure business continuity. These new formats bring to the market a variety of capacity and throughput characteristics, with SDLT, LTO, and SAIT leading the open format charge toward 1TB cartridges. IBM and STK, meanwhile, continue to champion the performance of their mainframe and open tape offerings, while at the same time, extending them into the open systems fray. Because of this renewed competition for the long-term storage portion of the ILM market, the cost of tape on a $/GB basis has been driven well below the $1.00/GB level. In fact, a Super DLTtape cartridge is available for about $.35/GB, while LTO-2 and SAIT cartridges cost even less.

How does this compare to disk, for those who are looking at local/remote disk options as a permanent back-up/recovery option? **Open systems disk prices have come down significantly, but not as far as tape!** A typical Fibre Channel (FC) disk device, in an array, will cost between $20-$30/GB, depending on whether you select 10K or 15K RPM drives. In a multi-tier ILM environment, you can also select an ATA solution that could reduce that cost to about $6.00/GB, with significantly less functionality. This media cost is still about 5 times that of tape. Furthermore, once your data center has made the initial investment in a tape library, the cost of adding incremental tape storage is so low that it increases the difference between tape and disk.

The decision on which tape format to use is primarily made based on the requirement for capacity and performance, although issues such as reliability and data protection have developed an increasing importance. There are also some instances when compatibility with the past plays a factor in that decision. A large legacy library of DLT cartridges, for example, may be a contributing factor to pursue SDLT in the future. Let's look at some of the new open tape formats that have been introduced since 2000 and how they might fit into your operations.

SDLT

Before 2000, the DLT8000, with a capacity of 40GB and a throughput of 6MB/s, was the market leader for backup and recovery, archiving, disaster recovery, and network storage in the mid-range systems marketplace. This covered UNIX, Linux, and Windows, among other open systems. Quantum took that architecture and enhanced it for better performance and higher capacity to address growing ILM requirements in both the enterprise and departmental/workstation environments. These enhancements took two parallel paths, depending upon system needs and budget. The first path, *Super DLT (SDLT),* is aimed at enterprise requirements of high performance and high capacity, and has a product roadmap spanning four generations. (See Exhibit 1.) The alternative, called *DLT VS*, aims at entry to mid-range server environments, also delivering

[2] See **The Clipper Group Explorer** dated January 29, 2004, entitled *The Role of Tape in Multi-Tiered Storage – Alive, Kicking, and Rolling Along* at http://www.clipper.com/research/TCG2004008.

Exhibit 1 – Super DLT Roadmap

	SDLT320	SDLT600	SDLT1200	SDLT2400
Capacity (GB)	160	300	600	1200
Throughput (MB/s)	16	36	>50	>100
Availability	2002	2004	2005	2007

increased performance and capacity, however, with a reduced cost and better value. With a growth path designed to double capacity every 18-24 months, SDLT appears to be ideally suited to keep pace with expanding enterprise storage requirements, while at the same time protecting investments made by the enterprise in new technology.

With an installed base of two million legacy tape drives, Quantum designed the *SDLT320* to be backward-read compatible with *DLT x000, DLT1,* and *DLT VS80,* using *DLTtape IV* media, ensuring compatibility with legacy tapes in the enterprise library. It also has read-write compatibility with the *SDLT220* using Super DLTtape 1. With the recent introduction of *SDLT600,* targeted for automated solutions, this format includes a native FC interface to improve the throughput and ease the transition of tape libraries into the SAN. It also includes an Ultra160 SCSI LVD interface. The average access time for the 600 is 91 seconds. Fourth generation *SDLT2400* drives will support a throughput of >360GB/hour with TB data cartridges.

SDLT drives feature a file mark index at the beginning of each tape to assist the drive in locating selected files in a high-speed search. They also use an adaptive cache buffering system to monitor the host's data transfer rate in order to match the drive's buffering operation. This helps to prevent unnecessary stop/starts that impair the drive's performance.

SDLT uses media with a single-reel design with stationary read/write heads, rather than the two reels per cartridge found in helical scan devices. Data is recorded on one set of tracks, and then the heads are repositioned for the next track, in the reverse direction.

SDLT drives all run Quantum's new diagnostic software *DLTSage* that provides intelligent maintenance for the tape users. In addition to proactively monitoring drive/media health, DLTSage also aids administrators in determining whether a failure is due to the drive or the media.

LTO

In 1998, HP, IBM, and Certance (formerly known as Seagate RSS) joined forces to create a pair of open-format tape specifications. This was driven by the confusing number of formats and technologies that existed in the tape storage marketplace. Based on the partnership's *Linear Tape-Open (LTO)* technology, and named *Ultrium* and *Accelis*[3], they were designed to bring open standards and new levels of scalability, reliability, and automation to the open systems tape backup market. Like SDLT, LTO uses a mechanically simple linear format that minimizes the number of moving parts. It is enhanced in the areas of hardware data compression, optimized track layouts, and highly efficient error correction codes. LTO maximizes capacity and performance. The partners were unencumbered with the need to maintain compatibility with the past. The Ultrium format is ideally suited for backup, restore, and archiving. These plans provided customers with a single technology roadmap for broad vendor support. Additional vendors may license the specifications of this open tape format.

One of the unique features of LTO is the existence of a cartridge memory, an EEPROM embedded in the Ultrium cartridge. This module contains 4KB of non-volatile memory for storage and retrieval of information about the cartridge and the data, such as cartridge ID and usage history. This permits the drive to know where on tape a record exists without reading a directory or table off the tape. In contrast, SDLT stores its user and error data directly on the magnetic media.

With support for the same, open environments that are available with SDLT, the LTO Ultrium specification introduced the format with 100 GB cartridges (native) and a roadmap covering four generations, to a cartridge capacity of 800GB. (See Exhibit 2, on the next page.) With a projected throughput of 160MB/s,

[3] Accelis was designed for fast access, but has been abandoned by the Consortium.

Exhibit 2 – LTO Product Family

	LTO1	LTO2	LTO3	LTO4
Capacity (GB)	100	200	400	800
Throughput (MB/s)	15	35	40-80	80-160
Availability	2000	2003	2005	TBD

Exhibit 3 – SAIT Product Family

	SAIT1	SAIT2	SAIT3	SAIT4
Capacity (GB)	500	1000	2000	4000
Throughput (MB/s)	30	60	120	240
Availability	2003	2005	2007	2009

generation 4 LTO drives will be able to backup 576GB of data per hour with an average file-access time of 51 seconds.

Introduced in 2003, the *LTO-2* format is offered by all three partners in a variety of environments, including internal, desktop, rack-mount and library. LTO-3 will play leapfrog with SDLT600 in 2005, providing 33% more capacity at 400GB native.

SAIT

For tape to remain cost-effective as a storage solution, growth in capacity must parallel that of disk, while retaining its existing cost/GB advantage. The industry roadmap to achieve that has established a target of 1TB of native storage per cartridge by 2006. We have seen in the exhibits above that Quantum (for SDLT) and the HP, IBM, Certance Consortium (known as the LTO Provider Companies) have both targeted cartridges with 1TB of native capacity in that timeframe. If you consider the doubling of capacity every 18 months to be valid, they may reach it. Some industry planners, however, project the limitations in recording density of today's technologies may limit that capability.

In response, Sony has introduced (in 2003) a new class of tape technology. They have adapted their *Advanced Intelligent Tape (AIT)* with a longer and wider media in a single-reel, half-inch cartridge. Using AIT's helical scan approach, they have a 4X advantage over linear tape technologies in recording density. Helical scan recording uses a very stable rotating drum/head platform that requires a low tape tension and uses a single pass operation. This permits accurate and reliable recording and read-back at higher densities. This has enabled Sony to introduce *SAIT* with an initial format of 500GB of native capacity and a planned capability of 1TB in 2005. Because of the helical scan architecture, SAIT has an MTBF of 500,000 hours with an average head life of 50,000 hours. This compares quite favorably with SDLT and LTO with an MTBF of 250,000 hours and a head life of about 30,000 hours. In addition, the average access time is 70 seconds without load time, 93 seconds with it. They have established a roadmap with a family of four products. (See Exhibit 3, above.) SAIT, therefore, provides us with a third alternative for high-capacity, high-performance storage.

SAIT contains a flash memory chip, a 64K Remote Memory-in-Cassette (R-MIC), embedded into the SAIT media to allow for local storage of important media data and statistics. You can identify media type, serial number, and error recovery data, along with a search map to provide high access to any file on tape. In addition, the short re-cycle time of helical scan recording allows SAIT to outperform linear devices in repositioning performance. Sony has implemented a Write-Once, Read-Many (WORM) architecture using the MIC, satisfying government regulations regarding data preservation. WORM is currently available for OEM SAIT-1 drives. Sony-branded drives will be available soon.

Mainframe Support

SDLT, LTO, and SAIT are all designed for open systems architectures. One other class of tape, however, must be included in this study: tape that is for open and mainframe use[4]. In addition to the backup, recovery, and archiving requirements seen in the open systems solutions

[4] Open systems use SCSI and Fibre Channel for connectivity, while mainframes use ESCON and FICON.

discussed above, mainframes have traditionally used tape in data processing applications. This requires a drive with very fast access times as well as high performance for better productivity. Drives such as the STK *T9940* or the IBM *3592*, typically, have an average access time 70% faster than SDLT320 or LTO1, and 25% faster than LTO2. However, the availability of disk-caching virtual tape products has transferred the responsibility for fast access from the drive to the front-end processing. The requirement remains for those who do not want to invest in virtual tape systems.

Historically, IBM and STK have controlled the storage needs of the mainframe community. Their latest products, the IBM 3592 tape drive[5] and the STK T9940B drive, continue to dominate new spending in this space.

StorageTek T9840/T9940

The natural compliment to the T9840, the T9940 is available in two varieties, 60GB (T9940A) and high-capacity 200GB (T9940B). This is a significant increase in capacity from the T9840. That device had a capacity of 20GB (T9840A) or 40GB (T9840C). With both proprietary and open systems interfaces, the T9940B is an excellent tape storage device for use in collecting, moving, and storing large volumes of mission-critical data, both quickly and reliably. The T9940B uses a 2Gb/sec bandwidth, a 30 MB/sec transfer rate, 200 GB native capacity, high-duty cycle reliability and significantly lowered failure rates. This allows customers to do more work in less time than ever before with a direct attach FICON interface planned for June 2004. This is also available on the T9840C. By using fewer tape drives, it achieves a superior standard of reliability than before. In a standard backup environment, with a 2:1 compression ratio, the IT staff can transfer up to 252GB/hour, reducing the length of the backup window.

For applications that use tape in a data processing mode, the T9840 drive has an average access time of 8 seconds, with the T9840C having Backward Read Compatibility with the T9840A/B. This is significantly faster than the T9940B, which has an average access time of 41 seconds, not including tape load time and thread to ready access. This also compares quite favorably with the 79 seconds quoted for access time for the SDLT600 or the 49 seconds for file access for LTO2 in a pure open systems environment.

The STK T9x40 drives use a proprietary factory written servo system to provide precision head positioning and track following. These drives have sophisticated error control systems comprised of hierarchical error detection and error correction codes, data formatting for error immunity and extensive error recovery procedures. These features provide the error recovery to assure enterprise-class data integrity and data throughput performance. In our age of regulation compliance where the security of the data is essential, the T9940B also comes with a *VolSafe* WORM-storage capability. Using the 200GB 9940 cartridges provides a high-capacity WORM storage solution more robust than those available using optical solutions. VolSafe is a non-erasable, non-rewriteable, cost effective medium for the secure archival of critical data.

IBM 3592 Enterprise Tape Drive

As a replacement for the IBM *3590* tape drive, IBM's *Enterprise Tape Drive 3592* provides high capacity and high performance for storing mission-critical data, with five times the capacity and 2.5 times the data transfer rate of the previous device. By offering significant improvements in both of these categories, the 3592 addresses storage requirements that are normally filled by two different types of drives, those that provide fast access and those that have high capacity. In addition to supporting all of IBM's open systems servers (iSeries, pSeries, and xSeries), the 3592 also provides storage support

Battle of Tape Titans Continues?

StorageTek (STK) states that the IBM 3592 tape drive in a C20 frame has not been tested or approved on the 9310 Powderhorn Library. As a result, STK cannot support this configuration and is refusing to service these systems. IBM states that they have installed over 4,000 3590 drives in STK silos and has successfully installed the new 3592 enterprise drive in a number of STK silos as well. IBM has encouraged STK to formally incorporate the 3592 drive in its library offering, and allow customers to make their purchase decisions based on value. IBM continues to offer the 3592 in the STK library.

[5] See **The Clipper Group Navigator** dated August 31,, 2003, entitled *IBM Delivers Advanced Tape for Open Systems and Mainframes* at http://www.clipper.com/research/TCG2003039.pdf.

for the *zSeries* and older mainframes.

With a native cartridge capability of 300GB, the 3592 provides 50% more capacity than the T9940B. Incorporating the suggested compression ratios endorsed by the manufacturers, a 3592 cartridge claims a compressed capacity of 900GB, while the T9940B cartridge may contain between 400 and 800GB. In terms of throughput, the IBM drive has a performance of 40MB/s, compared to the T9940B measured at 30MB/s. Both of these drives have a 2Gb/s Fibre Channel interface, in addition to the traditional ESCON mainframe connection. The 3592 also has a FICON interface. The average file access for the 3592 with a 300 GB cartridge is 39 seconds. Both the STK and IBM drives may share in an environment with mainframes and open systems servers, thus facilitating a multi-platform consolidation effort to reduce IT costs.

With over 50 years of IBM tape technology behind it, including participation in the partnership that developed the LTO drive, the 3592 tape drive has the high reliability required in mission-critical data gathering environments. By incorporating error-correction code (ECC) and factory-written servo tracks on the tape cartridge for precise head positioning, the 3592 can help increase data integrity. The tape drive also includes redundant power supplies with automatic failover to help prevent outages. In addition, IBM has improved streaming job productivity with its exclusive "Virtual Backhitch Technology."

In a recent announcement, IBM added a new, lower priced 60GB cartridge to the 300GB cartridge already available. Designed for use in environments where fast file access and retrieval is critical, or where data sets are smaller, the new cartridge has a file access time of 11 seconds. This is considerably faster than the 39-second specification for the 300 GB cartridge. The smaller capacity cartridge is also essential when recovery time for multiple data sets is critical. The two cartridge formats provide the data center with capacity options to meet dynamic application needs.

IBM also announced the availability of a pair of new WORM cartridges, 60 and 300 GB, for the 3592 Enterprise Tape Drive. Using the Cartridge Memory (CM), which is a passive, contact-less, silicon storage device that is physically a part of the cartridge, WORM identification flags are written to the cartridges during the manufacturing process. This includes low-level encoding written to the tape media and to the CM. The CM holds data that is specific to that cartridge, the media in that cartridge, and the data on the media. A robust algorithm then uses the encoding to prevent data tampering. The data center can append to the file but not overwrite.

Conclusion

Without a clear understanding of the tape environment in a specific data center, it is impossible to make any specific recommendations as to what is best for any given architecture. We have some general conclusions:

1. In an open systems environment, high throughput is a necessity for a high-volume data center to increase productivity, shrink the backup window, shrink the batch window, decrease recovery time, and improve response time. Typically, high capacity will go hand in hand with the high throughput, depending upon the budgetary constraints. Obviously, if the data center has a large quantity of DLT cartridges on the shelf, then the backward read compatibility of SDLT will certainly have to be examined.
2. Reliability is also a major concern in these data centers. Around-the-clock duty cycles (high reliability, data integrity) save time, reduce business interruptions, increase operations efficiencies, and protect data. These could outweigh the financial advantage of preserving an historic architecture, if another solution proves more reliable. The wide acceptance of LTO tape drives (over 500,000 shipped) with its open standards reliability, high performance and software, along with server and tape library integration, make it a strong candidate for consideration in this arena.
3. If the mainframe is alive and well in your operations center, then your architecture demands the performance, reliability, and functionality of an enterprise tape drive. Your IT staff needs to weigh the benefits of the T9940B and the newly updated 3592 tape drives.

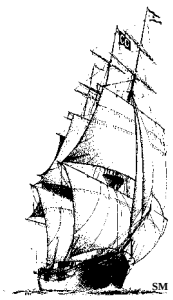

Tape offers many solutions for important enterprise challenges. Check out all that tape can do for you!

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine **is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 23. (Please dial "1-23" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.