Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE FOR THE INFORMATION TECHNOLOGY BUYER

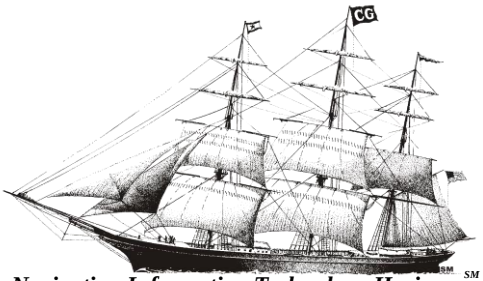

Navigating Information Technology Horizons℠

Published Since 2007 Report #TCG2014015 July 5, 2014

Is Tape the Best Low-Cost Technology for the Preservation of Data?

Analyst: David Reine

Management Summary

Did you know that you have, or are about to have, a new technology in your wallet? Well, the smartcard that has been in use throughout Europe, is coming to America, and the American consumer is about to hop on, or be pulled on to, the bandwagon. If you are in Europe, you have probably been using one for several-to-many years, and you may not have even realized it, as the technology is transparent to the consumer. You use it just like a credit/debit card except for the fact that you need to enter a PIN to complete any transaction, with identification and authentication embedded on a microchip on the smartcard. With the fraudulent use of credit/debit cards a worldwide epidemic, smartcards are the obvious solution, so why not use them? The rollout of this technology in America is a no-brainer, as we have seen with the recent announcements from Wal-Mart's Sam's Club and Target, trying to make their transactions more secure against hacking.

Transparent technology has been present in the data center of every enterprise, as we migrate our open systems IT architecture from one generation to the next without the necessity of a "fork-lift upgrade". As each new generation of IT professionals becomes entrenched in the enterprise data center, we need to take a step back and review the new, and old, issues that face today's data center staff on a regular basis, and how they can best be solved. Some of the issues that must be addressed, in no particular order, are: uncontrolled storage growth, cloud access, demanding analytics, and infrastructure cost. It would be impossible to address all of these issues (and others) in a single issue of Clipper Notes without putting half of the readers to sleep. Therefore, today I will focus on the rampant growth of data, access to it, and the need to protect and preserve that data, securely, for a very long time. In fact, preserve it forever, without having to worry about the technology becoming obsolete or breaking the enterprise IT budget.

One of the most serious problems affecting all enterprises, regardless of size, is data storage. Many enterprises are amassing a tremendous amount of information on a daily basis in an attempt to satisfy customer desires and to get an edge on the competition. These include, but are not limited to, entertainment, video surveillance, medical imaging, and analytics. Some of that data may be needed immediately, while a significant portion must be preserved for future use. There is urgency for accessing some of this data, e.g., for mission- and business-critical data. Urgent data requiring immediate availability must continue to be stored on solid state disk (SSDs), or spinning hard disks (HDDs), because the time to retrieve the data is short and timeliness is the name of the game. However, for archival data (often the primary source data) and extracts of databases (copies of data stored somewhere else) may need to be preserved for long-term use, often sporadic (if at all accessed), usually without the same urgency for instantaneous access.

As all forms of storage media, including SSD, HDD, and the lowest-cost form of storage, tape, continue to expand in capacity and contract in terms of price per GB/TB/PB, we may need to take a step back and gain a new perspective on the direction of storage in our immediate future. With

IN THIS ISSUE

SSDs, now available in TB capacities, carrying a cost per IOPS ratio below that of HDDs, and with the cost per TB also dropping into the range of HDDs, the role of spinning disk for primary storage in the data center may disappear in the not too distant future. In a turn of the table, perhaps it is the disk media industry that is at risk rather than low-cost tape. In fact, for the purposes of mission-critical applications, SSDs may soon overtake HDDs in terms of everyday usage. As for long-term storage, tape is very strong. When compared to disks or SSDs on a total cost of ownership (TCO) basis for archived data, tape is the best answer for enterprise archiving, a situation that should continue for the next decade. Why should you care what media your files are stored on as long as you can retrieve them, transparently, at a low cost in a reasonable and timely manner.

Assuming a continuing drumbeat of capacity and TCO improvements, enterprise data centers will continue to require at least two tiers of storage, one for high performance and the other for high capacity, but which two? *Do we need both SSDs and HDDs for our daily activities? Do we need both disk and tape for the long-term preservation of data?*

Do we need another new, inexpensive storage technology to reduce the TCO for the preservation of long-term data within the IT infrastructure? Or perhaps, is tape still the low-cost technology to meet our needs? All we have to do is re-discover magnetic tape with its new-found features and capabilities to make a transparent transition from old to new, from disk to tape. If you are feeling the pain from the strain being placed on your infrastructure budget, please read on.

Dealing with the Storage Explosion in the Data Center

The IT industry has always been able to meet the storage needs of the data center, however in some cases, at a cost, a steep cost. The evolution of primary storage easily can be demonstrated by reviewing the history of the most common storage medium today: magnetic disks.

The first disk created in the middle of the 20th Century was only 5MB in size and was designed mainly to address the immediate access to random data within the data center. Each succeeding generation of disk was designed to hold more data in an ever-shrinking package, until, by the end of the 20th Century, hundreds of MBs were being stored on a single 5" device, and this was proliferated by the deployment of PCs and open systems servers throughout the enterprise. By 2005 that capacity had grown to hundreds of GBs in 3.5 inch packaging, and by 2007, to 1TB, and that was only the beginning. Today, 4TB HDDs are in common use around the world, with 6TB drives now making their initial appearance, albeit at a very high cost per TB.

Cost for disk storage, however, continues to rise, especially on a TCO basis. The floor space and energy required to keep decades – and petabytes – of information on spinning media, not to mention the administration required, is rapidly becoming cost prohibitive in the aggregate, and there is no indication that data growth is going to abate any time soon. With the availability of SSDs in TB capacities, we may soon be asking about the viability of HDDs in an enterprise data center with extremely high IOPS demands. In order to achieve the required IOPS with disk, data needs to be thinly striped over many more disks than should be necessary from a disk available-capacity standpoint. Additionally, HDDs simply are too slow for the most demanding enterprise data center. This leads to extreme inefficiency and an unnecessarily high TCO.

Over the past decade, the rapid growth of data has created a new set of problems for the IT professionals in the data center. Using SSDs, or even the less expensive HDDs, for access to primary storage continues to make sense due to the need for high application performance and for the immediate availability of critical data to the mission- and business critical applications that drive enterprise profitability. In addition, with increasing demands to access new analytical applications, the requirement for the immediate availability of important information to real-time decision-making makes enterprise storage on SSD or HDD, more important than ever. The deployment of the virtualization of these critical applications throughout the enterprise infrastructure only reinforces that need by magnifying the requirement for more and more random access to data.

Unfortunately, disk does carry multiple burdens for the I.T. staff. An immediate issue always is the acquisition cost of these devices, but the TCO can place an unwarranted burden on the enterprise, especially for long-term storage. Issues with energy consumption, security, portability (through many generations of devices, usually each on a 3-to-5 year basis), and others are causing concern for the continued use of disk for second- and third-tier storage, not to mention the continuing maintenance costs upon the termination of warranty, often forcing the replacement of the HDDs after only three years. (See Exhibit A for,

Exhibit A— Universal Characteristics of Long-Term Storage

- **High capacity** – the ability to hold many PBs.
- **High scalability** – the ability to keep on growing (in terms of the PBs being stored).
- **Sufficiently fast** – it needs to be fast enough for the kinds and volumes of data that are envisioned for long-term preservation; this probably is measured in minutes and not is seconds or fractions thereof.
- **Long-term durability** – the ability for the media to last a long time, measured in decades.
- **Long-term viability** – the ability of the storage infrastructure to last for many years (without a fork-lift upgrade).
- **Very low cost** – where the total cost of ownership per PB is very low (in comparison to other viable storage options).
- **Evolving** – where the TCO per petabyte is getting lower because of improving technologies.
- **Energy efficient** – where the watts per TB stored per year is very low (in comparison to other storage technologies); this is about power and cooling.
- **Space efficient** – where the amount of space required per PB is very low (in comparison to other storage technologies).
- **Reliable** – where stored data is inherently safe, stable, and persistent.
- **Diligent** – where the storage infrastructure increases its reliability by problem detection and error correction before a problem arises.
- **Secure** – where data can be encrypted in transit and at rest.
- **Upgradable** – where transition to denser media or new storage technologies can be accomplished automatically (and without pain).
- **Self-describing** – where a set of data can be independently read and its content able to be re-categorized.
- **Simple** – where the storage solution and architecture are easy to comprehend and use.
- **Sufficient** – where the solution does what is needed (in a straightforward way) without a lot of unneeded baggage.
- **Integratable** – where the solution fits into the data center without a lot of new requirements for networking (protocols), location planning, or special handling.

Source: The Clipper Group

above, a more complete list of the requirements for long-term storage.) Even with advancements in storage capacity, the growth of data almost always far outstrips any gains made from higher device capacity. Increasing the amount of storage in the infrastructure is inevitable. Furthermore, with limited floor space in the data center, the staff must be constantly aware of the density efficiency of the data on the devices in order to minimize the over-provisioning of storage. HDDs simply are too inefficient for high IOPS requirements and too expensive for largely static and lightly used data where sub-second retrieval is not required, especially for the very large volumes of data that now must be archived[1].

This leads to the next problem, limitations on the availability of energy within the data center. There is only so much electricity available for the power to run the disk drives and air conditioning necessary to cool the environment, to ensure the viability of the infrastructure. Having thinly-striped data across many HDDs only magnifies the problem. The quantity of energy available may not even be the most serious energy issue for the data center. The cost of that energy, much like the cost of gas for our automobiles, continues to rise, especially in the more densely populated areas. The cost of energy in cities such as New York, Los Angeles, London, or Paris will put a serious crimp in the bottom line of any budget-conscious CIO.

The preservation and protection of data, one of the most valuable assets of any enterprise is no small task either. With hackers constantly trying to penetrate the security of the enterprise firewall and steal privileged information, as well as potential disgruntled employees looking to cause havoc, the enterprise staff must deploy technologies (such as encryption) to protect enterprise data. With the ever-present danger of data corruption,

[1] See the issue of *The Clipper Group Calculator* dated May 13, 2013, entitled *Revisiting the Search for Long-Term Storage – A TCO Analysis of Tape and Disk*, and available at http://www.clipper.com/research/TCG2013009.pdf.

both by accident and knowingly, the enterprise must institute backup and recovery procedures for both short-term and long-term protection, with long-term protection likely involving in a remote location. With the constant threat of natural disasters hanging over our heads, both figuratively and literally, the enterprise must find the means to reproduce the entire infrastructure, including the data, from a remote location. The cost of this protection and preservation on disk is becoming prohibitive, especially with a doubling of the data to be stored every 12-to-18 months, while the portability of these devices also is problematical.

Obviously, whatever storage architecture is in play, it must have the reliability, availability, and serviceability (RAS) to ensure business continuity for employees, partners, and customers alike in a 24x7x365 environment. Any storage system that loses integrity (and thus cannot provide the data being sought) is worthless. With devices constantly growing in capacity, RAID-6 also becomes a protection necessity, a potentially cost-prohibitive one in terms of the added cost that results.

One storage architecture that has proven itself time and again is magnetic tape. Tape has been used for the long-term preservation of critical enterprise data since the beginning of time – well, since the start of the computer era, for over 60 years. However, quite clearly we are not talking about the same technology that existed in the last century. A new era of tape technology has appeared with higher capacities and more reliability plus security features to protect that data.[2] Not only that, but with single cartridge capacities now available at more than 5TB per cartridge, tape can provide a very low-cost alternative to spinning disks, especially for infrequently-accessed archived data.

Debunking the Myths of Tape

First of all, let us dispel some of the many myths of tape, the first being *tape is not dead but alive and now well endowed.* In fact, it is thriving in the enterprise data center, mainly for long-term storage and archiving. In some instances, such as geophysical applications, it is even used as an initial storage depository.

Next, *tape is not an old technology* without a future. The capacity and features of today's tape are nothing like the tape from 20, 40 or 60 years ago. Today's tape continues to expand capacities per cartridge and add new capabilities via new technologies, including features such as extremely high capacity, WORM, encryption, and the *Linear-Tape File System (LTFS)* for use with standard file systems. This is new technology created for a growth industry.

Third, *tape is reliable* and *90% of tape backups do not fail*. Unfortunately, no one can verify this urban myth, yesterday or today. In fact, the reliability of tape exceeds that of enterprise FC/SAS drives or enterprise SATA, which pulls up the rear.

Fourth, *tape is less expensive* than disk. I do not know where that one came from[3], but anyone who has analyzed storage media knows that the TCO of disk is many times more expensive than tape.[4]

There are other myths which also lack credibility. But rather than rehash them all, let's just say that tape is a proven technology and deserves fair and reasonable consideration for long-term data storage in your enterprise data center.

A New Era of Tape Technology

There are certain qualities that tape has retained over the years, such as low-cost and portability. Today, however, many new features have also been added. Because tape is a sequential medium, accessing a specific piece of data on it could require more than a few seconds as the library automatically mounts the tape cartridge and the tape drive scans down the tape looking for the desired record, even with LTFS. If immediate access is a requirement, then SSD, or even HDD, technology is required.[5]

Because of the many myths about tape, many data centers have gravitated away from tape in favor of disk. In fact, some data centers are staffed with many younger folks who may not even know of the many advantages provided by tape. Recent developments over the past decade have changed that landscape, improving reliabil-

[2] The same can be said for disk. So why is it so hard to believe for tape?

[3] Maybe from an era when tape libraries did not have automation and thus relied on a bevy of operators to feed and manage it.

[4] Clipper has done much analysis and written many reports on the TCO of tape versus disk, and the conclusion always was the same: tape costs significantly less than disk per TB for data held over the long term (as in archiving). See report referenced in footnote #1.

[5] However, remember that this is the time to get to the first byte. If you are accessing a big file or a large database, what usually is important is the time to get the last byte into the server. For large files, this is the more important metric and tape is very competitive here, because it streams faster than RAIDed hard drives.

ity, while at the same time enhancing capacity, performance, security, and facilitating access. In fact, we have been undergoing a technological evolution in what some would have you believe is a dead technology. Let's see if tape can do what users need for the storage and retrieval of archived data. Furthermore, let's see if tape can do it with transparency and the lowest cost. If tape can satisfy these needs, shouldn't you want to consider it?

Capacity

The capacity of a single open systems cartridge (*LTO-6*) at 6.25TBs[6] outstrips the capacity of open systems tape drives and cartridges deployed in today's data center. Enterprise tape drives and cartridges also have been developed by both IBM (*TS1140*) and Oracle (*T10000D*) which have even greater capacity. In fact, the IBM TS1140 cartridge has a native capacity of 4TB (12TB with a compression ratio of 3:1), while the Oracle T10000D, using their maximum capacity feature, has a native capacity of 8.5TB (21.25TB with a compression ratio of 2.5:1). Furthermore, both Fujifilm and Sony have demonstrated new media with enormous head room, to the tune of 154TB (Fujifilm) and 185TB (Sony) per cartridge. The tape capacity story only will get better.

Performance

With a native throughput of up to 160MB per second, an LTO-6 tape drive has sufficient performance for any data center storage application. For high-performance computing, the enterprise drives have a native throughput in excess of 250MB/second. This creates an environment that is ideally suited for the fast streaming of sequential data, exactly what most archives require.

File Access

Linear Tape File System (LTFS) is the format for data recorded on special tape media and the implementation of specific software that uses this data format to provide a file system interface to that data stored. The LTFS format is a self-describing, indexed tape format developed primarily to address tape archiving requirements. It enables applications with the capability to access files on tape in a persistent format as if they were on disk from any connected platform, including servers, PCs, laptops, and smartphones. The data center no longer has to retain proprietary applications written specifically for the long-term preservation of data. It is ideal for retaining large numbers of small files as well as big data applications, such as entertainment, video surveillance, medical imaging, etc. As a result, tape has evolved into an easy-to-use technology. What happens behind the scenes in the data center is transparent to the user or application, which only sees familiar files in directories/folders.

More recently, a new enterprise version of LTFS is now also available to enable the data center to replace disk for tiered storage. This new version enables the virtualization of many LTFS tapes in a larger library environment.

Reliability

Despite myths to the contrary, tape is more reliable than disk. With a published uncorrected bit error rate (UBER) of 1 x $10E^{-17}$ bits, LTO tape has established a clear lead in reliability over enterprise SATA drives at 1 x $10E^{-15}$ bits and enterprise FC/SAS drives at 1 x $10E^{-16}$ bits. Furthermore, disk UBER figures pale in comparison to an enterprise tape benchmark of up to 1 x $10E^{-19}$ bits.

With higher capacity and higher throughput, the enterprise can deploy a tape library with fewer drives and fewer frames to hold the cartridges. As a result, the enterprise can achieve improved reliability and do so with less data center floor space. With advanced, integrated tape management and analytic tools, the data center staff can monitor the health of the entire tape infrastructure, drives, robots, and even the integrity of the media.

Security and Compliance

Tape provides the enterprise with the capability to use WORM media to satisfy many government and industry regulations, including compliance requirements. Data written to WORM media is permanent; it cannot be overwritten or erased by accident or through malicious intent. In addition, modern tape comes with built-in encryption to prevent anyone from reading your critical information without the encryption key. In addition, should a tape cartridge go missing (i.e., should it "fall off a truck"), the media is unreadable and worthless and the enterprise is protected against potential liability.

The "fall off the truck" phenomenon is rare, but it could happen, because tape media is portable. Portability enables the data center to remove media from the data center and store it in a secure, remote location, protecting it from being infected with any of the viruses that abound in a connected world. In the event of a disaster, those cartridges could be recalled and shipped back to the same or another data center for a complete recovery,

[6] This is a compressed capacity and assumes a 2.5 times compression factor.

providing another layer of security and confidence.

Scalability

Tape provides the data center tape library with virtually unlimited scalability, on demand. Libraries are available in any size, from a single media tray that can be rack-mounted, to a multi-frame library with anywhere from hundreds to thousands of cartridge slots capable of supporting many PBs of data. Often, additional slots can be procured when provisioned on demand, in a pay-as-you-grow fashion.

Total Cost of Ownership

At this point, there should be no confusion with regard to the TCO of tape versus disk. The TCO of a disk subsystem has been shown to be up to 26 times the cost of a comparable tape library[7]. The TCO includes the hardware, media, maintenance, energy, and floor space required to support the infrastructure. This continues to be a moving target but, in general, the cost per TB for tape is dropping faster than the cost per TB for disk.

Energy

The cost of energy alone for the average disk-based solution exceeds the entire TCO for the average tape-based solution.[8] Disk consumes about of 105 times as much energy as a tape library (for the same large amount of disk space). The reason is simple: tape media is being stored in cartridge slots within the library (which consumes little energy when idle), and which then waits to be selected for transport to the tape drive (which consume little energy, compared to disks). The entire disk solution, however, is always under power, consuming kilowatts of electricity each hour, assuming (of course) that the data center has sufficient power available.

Floor Space

Today's densest tape libraries can store about 2.2 PBs of uncompressed data in a 10-square-foot (single enclosure) footprint, making it more than twice as dense per square foot as a full-sized rack of high-capacity (4TB) disk drives.

Durability

For those of you who are still uneasy about the durability of tape, rest assured. Tape has been estimated to have a shelf life of 30 years.[9] With tape media management tracking the health of any particular piece of media dynamically, the IT staff can be sure that they will have sufficient notification to replace the media before it is corrupted.

Conclusion

Obviously, the ability of the computer industry to address the capacity growth of primary storage has not been a problem as many enterprises still are doubling the capacity of their disk farms on an annual basis. Preserving that data in the form of backups and archives for long-term storage and use in a cost-effective manner, however, is another story. The expanding consumption of data center floor space and energy has simply made it too expensive to continue to use disk drives as a secondary media for long-term data preservation. In fact, with the developments being made in solid state storage, within five years we may find that the two-tier deployment of storage in the data center consists of SSD and tape. There may not be a need for disk at all.

Tape, with a more than a 60-year history of success in the data center, has evolved to a point where it can be the solution of choice for the preservation of data in every data center. Tape can meet data center requirements in terms of capacity, high performance, security, and reliability, as well as resolving such TCO issues as scalability, floor space, and energy efficiency and, most importantly, cost per TB.

If you need to find a new technology to deliver low-cost storage to the enterprise, tape can reduce the TCO of your IT infrastructure. It remains the low-cost alternative for data stored for the long term. Tape is the obvious solution; so why not use it?

[7] See report mentioned in footnote #1.

[8] See again, the report mentioned in footnote #1.

[9] Many data centers migrate drives and media (cartridges) every two-to-three generations, not because they are concerned about data loss but because the increased densities (capacity on a single cartridge) yields much lower costs per TB stored.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-second year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 40 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 a Senior Contributing Analyst for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, 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. 123. (Please dial "123" 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, The Clipper Group Voyager, Clipper Notes, The Clipper Group Calculator, 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.

Disclosures

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.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

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.