



TAPE IT

HP OEM NEWSLETTER - March 2011

Contents

- Tape saves the day for Google's Gmail
- The hot new storage technology for 2011 is... tape?
- LTO-5: Extending the TCO advantage versus disk
- Shipped tape media capacity reaches record levels
- LTO-5 continues to win versus Oracle T10KC tape drive

Tape saves the day for Google's G-mail



Sunday 27th February 2011 saw Google Gmail suffer one of its worst outages in recent history, when thousands of Gmail accounts were deleted due to a software bug caused by a storage software update. Fortunately for Google, they were able to rely upon tape to restore the deleted data.

Ben Treynor, VP Engineering and Site Reliability Czar for Google Gmail, used the official Gmail blog to explain the situation and provided a powerful endorsement for off-line removable tape backup.

“I know what some of you are thinking: how could this happen if we have multiple copies of your data, in multiple data centers...well, in some rare instances software bugs can affect several copies of the data. That's what happened here. Some copies of mail were deleted, and we've been hard at work over the last 30 hours getting it back for the people affected by this issue... To protect your information from these unusual bugs, we also back it up to tape. Since the tapes are offline, they're protected from such software bugs”.

For the full Gmail blog visit:
<http://gmailblog.blogspot.com/2011/02/gmail-back-soon-for-everyone.html>

Google's statement reflects the thoughts of a great many other customers who see that the industry is moving on from rehashing tired old, 'tape is dead' clichés to discussing a bright future for an innovative and still essential storage technology. The roles of disk and tape have evolved and are complementary in fulfilling best practice data protection strategies.

“The hot new storage technology for 2011 is...tape?”

A new research note published by the Data Mobility Group¹ suggests that; “Tape-based storage is more relevant today than ever before” and concludes “DMG encourages all organizations to take a fresh look at tape.”

The Data Mobility Group reports that, contrary to decade-long proclamations from disk-only vendors that tape is obsolete, there is a rising tide of evidence to suggest that tape will become even more important in the years ahead. They point to several underlying factors on which they base their belief:

Trends in data storage are matched by the traditional advantages of tape - including tape’s ability to manage rampant data growth, offer exceptionally low cost-of-ownership, and deliver dependable data archive over extended periods.

Advances in tape technology – DMG refer to the significant increase in tape drive reliability including the ability for tape libraries to automatically monitor status and health, and the ability for many tape drives to proactively notify administrators of potential problems. DMG also points to the massive capacity and throughput of LTO-5, together with future projections outlined in the LTO Ultrium roadmap. They also applaud new security features; WORM and Data Encryption. Finally, the Linear Tape File System (LTFS) is referred to in the report as “truly a breakthrough tape technology”.

A vibrant ecosystem that expands the utility of tape – the traditional advantages of tape and the steady-stream of improvements in tape technology, have fostered a thriving ecosystem of after-market solution providers for the tape market. DMG believes that there is now a wealth of innovative solutions in software and hardware, which leverage tape’s strengths while overcoming some of its unique challenges.

The ideal solution for archive storage – DMG recommend that enterprises seriously consider using tape in their archive tier, stating that “continuous data growth as well as data center energy costs and limitations has made it ever more untenable to keep inactive data on primary disk, but for various reasons even inactive data must remain accessible. An active archive infrastructure is now an absolute necessity for many enterprises”.

A key part of disaster recovery – DMG believe that replication from disk to disk is still too expensive for many organizations and large-scale restore over WAN can be problematic. Therefore, tape will continue to play a key role in backup and disaster recovery solutions, storing off-site tape backups in a disk-to-disk-to-tape approach.

Inherently well suited to some primary storage applications – such as streaming data (audio, video), and data that is generated or processes sequentially (such as genomic sequencing data and other kinds of scientific data capture).



In summary, the report author, Walt Purvis, states:

“Depending on the application and overall environment, tape can be an optimal choice for backup, disaster recovery, archiving, or even primary storage. Tape especially shines for large-scale, long-term archiving. When implementing an active archive strategy, tape has significant advantages over disk, both in cost and functionality. Now that software exists to make tape a first-class online storage target, the cost savings—in particular, the greatly reduced power and cooling expenses—and the benefits of being able to create a portable off-line copy of invaluable data are too great to ignore.”

¹Data Mobility Group, Research Note, March 1 2011: ‘The hot new storage technology for 2011 is...tape?’

LTO-5: Extending the TCO advantage

Proponents of disk-only storage claim that the falling cost of disks, together with data deduplication and data replication technology, has combined to make disk as cost-effective as tape. However, two new analyst studies refute this claim, by demonstrating tape's clear and compelling TCO advantage over disk and disk-based virtual-tape-libraries (VTLs).

The previous two editions of the HP Tape-It newsletter have focused on the emerging role for tape, and LTO-5 Ultrium in particular, as a premier long term storage and archive technology, in addition to a dependable high volume backup solution. To quote one storage analyst "The 21st century data explosion is here – and tape is well positioned to become The Digital Curator of the Information Age."²

The low total-cost-of-ownership (TCO) delivered by tape is one of the compelling reasons that position it as an ideal solution for digital archiving, a fact supported by two new studies:

- **Enterprise Strategy Group (ESG), February 2011: A Comparative TCO Study: VTLs and Physical Tape.**
- **The Clipper Group, December 2010: In search of the long-term archiving solution – tape delivers significant TCO advantage over disk.**

Both studies conclude that tape is the least costly solution by a substantial margin in terms of total-cost-of-ownership for multi-year data retention.

ESG: A comparative TCO study of VTLs and physical tape

This report is focused on the total-cost-of-ownership (TCO) aspects of LTO-5 Ultrium Tape Libraries compared to a selection of deduplication scenarios using a Virtual Tape Library (VTL) over a five-year timeframe. In the ESG model, the TCO of VTL with deduplication ranges from 1.78 to 4.16 times higher than the various implementations of a LTO-5 tape library.

For the full report, please visit www.ultrium.com

The assumptions underpinning the ESG study are as follows:

- The initial amount of disk capacity at the primary site assumption started at 35TBs
 - o With 52% (industry average allocation for non-mainframe systems) allocated to live data
 - o Leaving 18TB of data to be physically backed-up
- A compound annual storage growth rate of 35%
- Assumes list prices with a 30% discount to determine a consistent average selling price
- The costs used for comparison include hardware, software and personnel to manage each device
- The virtual tape library employs data deduplication in order to increase disk storage efficiency by a ratio of 15:1




² Tape: The Digital Curator of the Information Age - Fred Moore, President Horison, Inc.

SCENARIO 1: Deduped VTL compared with an LTO-5 Tape Library

In this first scenario, both the VTL and the LTO-5 Tape Library are assumed to be located within the data center, and both the backup and the data archive are performed locally.

Scenario 1: TCO of Deduped VTL vs. LTO-5 Tape Library



In Scenario 1, the TCO for the VTL with deduplication is 2.46 times greater than the LTO-5 Tape Library in year one. By year five, it is 1.78 times greater than the tape library. The relative difference between the two results is primarily the result of scaling:

- The tape library having been upgraded with expansion modules (each containing 92 additional tape slots with added media) and two additional LTO-5 drives during the period.

- The VTL with deduplication is upgraded once in this period (in year three), to support a higher raw capacity for the data ingest.

The study shows that the LTO-5 Tape Library is less expensive to run than the VTL with deduplication in every single year. Using the Tape Library instead of the VTL in this scenario could result in a total saving of more than \$430,000 over the five year period.

SCENARIO 2: Replicated VTLs compared with an LTO-5 Tape Library

The second scenario compares the TCO of backup followed by remote offsite storage for disaster recovery purposes.


For the VTL solution, two VTLs are employed for replication between the data center and the remote site. The backup data is first deduplicated, and then replicated to the remote

VTL using a wide-area-network (WAN) .

For the LTO-5 Tape Library based solution, tapes are transported offsite using a truck.

Note: in this example, the WAN pricing is based on half of a dedicated T3 line. Truck-based storage services are priced at \$450 per month.

Scenario 2: Replicated VTLs vs. LTO-5 Tape Library with tapes transported off-site by truck



In scenario 2, the TCO for the WAN-based data replication from VTL to VTL is 4.06 times more expensive than the TCO for the LTO-5 Tape Library using a truck to move backup tapes offsite in year one. Over a five-year period, the TCO for the VTL replication solution is 2.78 times greater than the LTO-5 tape library and truck-based offsite transport combination.

It would appear that adding remote tape storage via a truck to transport tapes to a remote storage facility adds relatively little to the TCO for a tape library solution. The tape library solution remains significantly less expensive than a VTL deduplication and replication solution.

“Throughout ESG’s study and scenarios, the LTO-5 Tape option was the least costly backup solution (in terms of TCO) over a five-year period by a substantial margin.”

Mark Peters, Enterprise Strategy Group, 2011



Clipper Group: In search of the long-term archiving solution

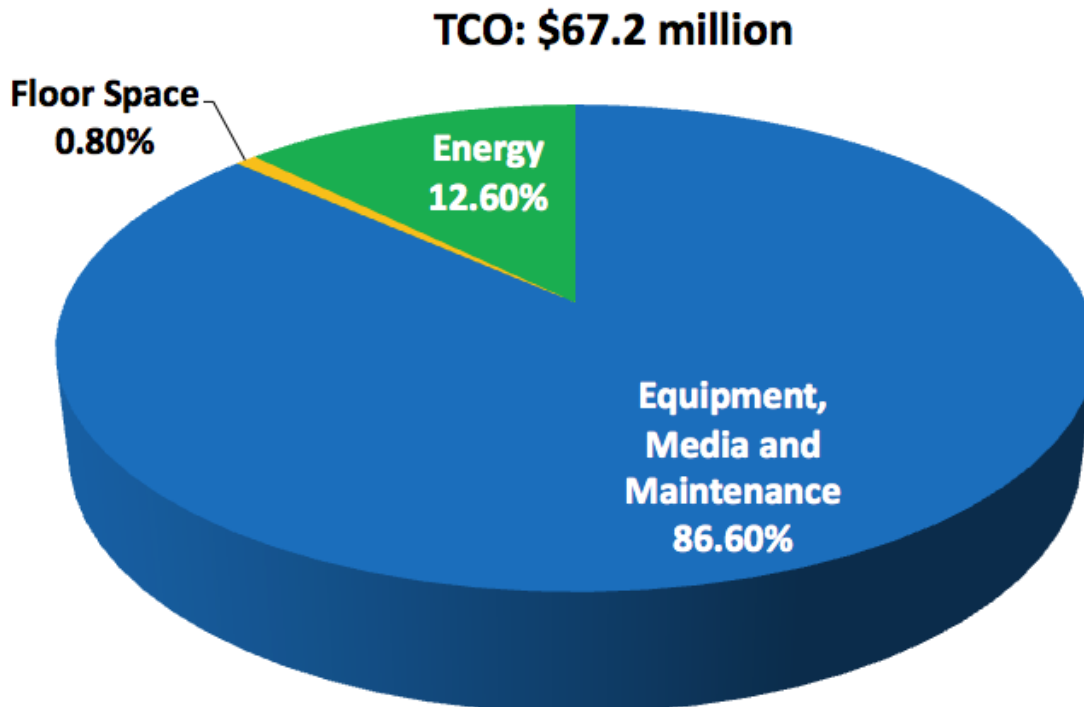
The Clipper Group recently released an independently funded report that compares the total cost of ownership (TCO) of disk and tape as the predominant storage media for long term archiving. The results of this study conclude that disk is more than fifteen times more expensive than tape, based upon vendor-supplied list pricing, and uses 238 times more energy (costing more than the all costs for tape) for an archiving application of large binary files with a 45% annual growth rate, all over a 12-year period.

To read the full report, please visit:
<http://www.clipper.com/research/TCG2010054.pdf>

The assumptions underpinning the Clipper Group study are as follows:

- Starting with 1PB of archive data and 1.45PB of archive storage by the end of year 1, and growing to 86PBs of storage by the end of year 12
- The model used a 2TB SATA disk assembled into full-configured arrays, and an LTO-5 tape library
- Both tape and disk are filled to 85% of their usable capacity
- Disk and tape both have negligible costs for transition between generations
- Disk and tape consume energy at their maximum operating levels
- No IT personnel costs are included in the model
- Both disk and tape are replaced at the end of a three-year cycle
- Assumes that data is not compressible

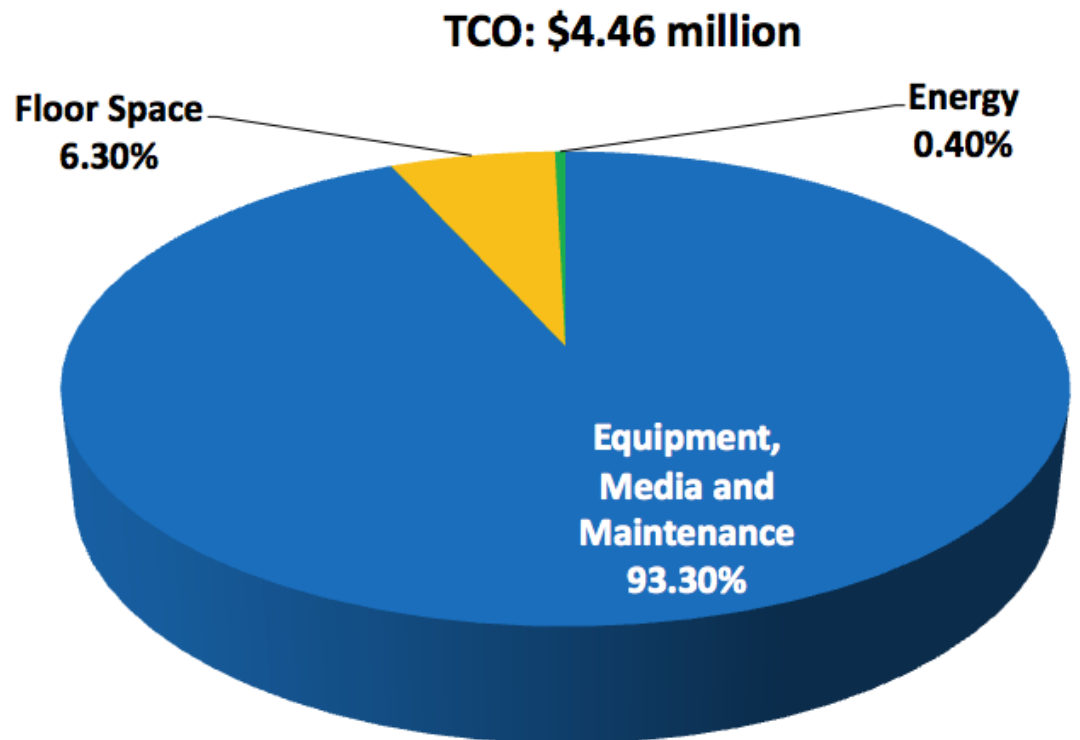
Disk: Costs by Budget Category



Source: The Clipper Group, December 2010

The Clipper Group study shows that the cost of energy is a significant factor in the operational costs of maintaining a disk subsystem for long-term storage. The study concludes that the cost to run and cool a long-term disk archive was more than \$4m. This is compared to tape with energy costs of only \$18,000, even using the unlikely worst case scenario with tape drives operating 24-hours per day.

Tape: Costs by Budget Category



Source: The Clipper Group, December 2010

“In every dimension, the TCO of the tape solutions was found to be less expensive than the TCO of the disk solution for long-term data retention.”

The Clipper Group

In conclusion, while a disk-based solution offers faster data access times than tape, the real question becomes; “does your data center need to spend millions of dollars to achieve sub-second response times for archive retrieval requests? Or is a response time of several minutes sufficient to meet service level agreements at 1/15th of the cost?”³




³ The Clipper Group: In search of long term archiving

Shipped tape media capacity reaches record levels – surpassing shipped external disk capacity by around 23%


The best metric for the continuing value of a technology is market adoption. A further positive proof point for the ongoing robustness of tape market demand can be seen in the increasing amounts of shipped tape media capacity.

In quarter four of calendar year 2010, total shipped tape storage media capacity continued

to surpass (by around 23%) total shipped external disk systems storage capacity. Total tape storage media shipments reached a record level of 4,374 Petabytes in Q4 CY'10, according to The Santa Clara Consulting Group. This compares to 3,537 Petabytes of total external disk systems storage capacity shipping in Q4 CY'10, according to IDC.



Tape media, and disk storage system capacity shipment trends CY08 to CY10



Sources: IDC WW Disk Storage Tracker CQ410– March 2011 and Santa Clara Consulting Group Tape Media Tracker - February 2011

LTO-5 continues to win versus Oracle T10KC tape drive

In January 2011, Oracle announced the Sun-StorageTek T10000C (T10KC) tape drive. At first glance, it might appear from the capacity and performance specifications that the T10KC has more to offer than LTO-5; but as the chart below illustrates, LTO-5 offers superior customer value in the areas of market adoption, investment protection, backwards compatibility, configuration flexibility, enterprise class reliability at an affordable price, TCO and energy efficiency.

At a glance comparison:

	Oracle T10KC	LTO-5 Ultrium
Positioning	Designed for mainframe data protection in an enterprise environment	Wide range of customer environments – direct attach SMB through to mid range and enterprise environments
Capacity (native)	5TB	1.5TB
Performance (native)	240MB/s	140MB/s
Security	Data encryption WORM	Data encryption WORM
Market adoption	Niche player with 1.4% unit share of the overall tape drive market in Q3'10	The world's most popular tape format with 74% unit share of the overall tape market in Q3'10
Format & compatibility	Proprietary format tape technology Single supplier No publically published roadmap Limited backward compatibility	Open standard format tape technology Multiple sources of supply Publically published roadmap Assured compatibility
Configuration	Large form-factor tape drive only, rack-only server support Automation in Oracle SL8500 and Oracle SL3000 enterprise libraries only Interface: 4Gb FC and native FICON Minimal choice of complementary products/services	Choice of half-height and full-height, internal/external/rack-mount/blade server support Automation in autoloaders through mid-range and enterprise libraries from multiple suppliers Choice of native interfaces – SCSI, SAS, 8Gb FC Wide and established range of media suppliers, software vendors and service offerings
Total Cost of Ownership	Very expensive tape drive Media expected to be at least double the price of LTO Less energy-efficient – requires 67W of power	Tape drive less expensive than Oracle T10KC Media expected to be half the price of Oracle T10KC More energy-efficient – requires only 24W for LTO-5

In Summary

When considering the high capacity and performance of the Oracle T10KC Tape Drive, customers must balance out the potential gains with the risk of investing in a proprietary format. LTO Ultrium continues to deliver an open source technology, that's widely adopted throughout the industry today, offering the assurance of supply and a competitive pricing advantage for OEMs and end-user customers alike.