

A Solutions Approach for Low TCO: What are the Costs to Store Data Long Term on Tape and Disk?

David Reine

Director, Enterprise Systems

david.reine@clipper.com


781-235-0085 Ext. 123

The Clipper Group, Inc.

Technology Acquisition Consultants and Strategic Advisors

Navigating Information Technology Horizons™ since 1993


www.clipper.com



Management Summary


- **Our Objectives**

- *Determine the TCO Benefits of Tape vs. Disk in long-term storage for a mid-sized data center*
- *Compare the costs of long-term SATA storage in a disk archive vs. the long-term costs of an LTO-4 Tape library archive using brand name products*
- *Selection of 750GB disk drives and 800GB tape cartridges*




Management Summary

- **Our Methodology**
 - *We chose backup archiving as a common long-term storage activity, as every data center does it*
 - *Backup and archive are ideal environments to determine long-term effects of a very large and growing storage requirement*
 - *As the environment grows, incremental costs (tape or disk) become fairly constant in terms of cost/TB*
- **Our Finding – Cost ratio of disk to tape is 23:1**




Backup/Recovery Challenges

- Optimize performance vs. cost on a \$/TB basis
- Determine gating factors: Backup window, library slots, etc.
- Complete the backup within a fixed time period (backup window)
- Short-term recovery must be completed in a timely fashion to avoid an out-of-service condition




Case Assumptions

- **Size of Initial Storage = 50TB**
- **Annual Growth Rate = 50%**
- **Five-year storage total is > 6.5PB**
- **Total Cost of Ownership factors**
 - *Acquisition Cost (at list prices)*
 - *Power/Cooling Energy*
 - *Floor Space*
 - *Maintenance*
 - *Media*
- **Daily incremental and full weekly backups for most recent 13 weeks always go to disk for both the disk and tape scenarios**
- **Quarterly backups archived thereafter go to the disk or tape archive for this TCO analysis**




Disk Parameters

- All disk arrays will be F.C. and fully configured to simplify acquisition and deployment costs at \$10.6K/TB
- All arrays will have SATA drives
- All arrays will be replaced after the three year warranty to take advantage of higher capacities and refreshed warranty
- 750 GB SATA devices used for the disk array (capacity over 350TB)
- We used same capacity utilization ratio (85%) for disk as tape even though some studies position tape as 10% to 15% higher
- Two RAID 5 arrays per drawer to eliminate need for accelerated maintenance
- One hot spare per drawer




Tape Parameters

- Initial purchase of a tape library large enough (>6000 slots) to support all cartridges nearline for the 5-year study to expedite recovery
- LTO-4 drives and cartridges used in the tape library
- LTO-4 media costs \$12.6K in lots of 100
- Reserve adjacent data center floor space for year 5 capacity requirement even though some of it will be empty for years
- Tape drives added as required to satisfy backup window




Tape Parameters – continued

- Configured n+1 drives to ensure device availability and ability to backup storage (backup window) within 7 days, in background
- Tape compression ratio assumed as 2:1 (data dependent)
- Adjusted write throughput downward by 15% to account for possible network inefficiencies
- No urgency for recovery after 13 weeks
- Did not utilize LTO-4's capability for WORM or encryption in this scenario
- Value of LTO-4 features varies for each data center
- Did not consider LTO-5




Energy Parameters

- Energy availability is not, in itself, a limitation
- Data center consumes matching kilowatts of energy for cooling as for powering equipment
- Energy costs (average urban rate) assigned as \$.12/KWH
- Energy rate held as a constant for term of study
- Higher costs favor tape due to higher energy demand from disk




Floor Space Parameters

- **Disk Archive**
 - *One fully configured array requires 19.5 sq. ft. at a cost of \$300/sq.ft./year*
 - *Floor space doubled to allow for aisles*
 - *Adds \$590,000 to 5-year TCO*
- **Tape Archive**
 - *Each library frame occupies 10.2 sq.ft.*
 - *Floor space doubled to allow for aisles*
 - *Adds \$30K/year for library floor space*




Accounting Parameters

- **TCO for 5 years including**
 - *Hardware acquisition*
 - *Media (tape) costs*
 - *Cost of energy*
 - *Cost of floor space*
- **Did not include cost of backup software, backup servers, network operation or time-value of money**
- **Model includes only costs allocated to each year of the 5-year study period**
- **The cost of the 13-week interim disk storage buffer was excluded in the cost analysis as it is common to both scenarios**




Study Findings

- 1st 3 years consume over 2.5PB of disk in the disk archive
- Each disk array consumes 11,000 Watts/hr (\$11.5K/yr) plus equal amount for cooling
- 5-year cost of energy for disk archive is \$420,000 at \$.12/KWH
- Tape archive energy cost for 5 years is \$1,300
- Energy and floor space represent only 4% of the total cost
- 5-year TCO for Disk Archive is \$51.73M
- TCO for disk archive solution is 23X the cost of the tape archive solution
- Energy cost for the disk archive solution is 290X the tape configuration




Conclusions

- Blended environment provides best mix of high-performance in short-term with lower TCO for long-term storage
- Costs associated with a large disk archive solution are probably too great for an organization to absorb for long-term storage
- Tape continues to be the most economical solution for long-term storage
- When using a Data Deduplication VTL with a 20:1 reduction factor – the estimated TCO is still 5 times greater than the tape solution but might add significantly to acquisition cost
- Your situation and costs may vary



Other Considerations for a Blended Environment

- **Tape is removable**
 - *Not susceptible to intentional or unintentional corruption*
- **Tape is portable**
 - *Economical method to have a copy of critical data off-site or out of region for data protection*
- **Tape is fast**
 - *LTO-4 can backup up to 864GB per hour (2:1 compressed)*
- **Tape is reliable**
 - *Over 2M LTO drives and over 80M LTO cartridges shipped*
- **Tape can address data security**
 - *LTO-4 has hardware data encryption*
- **Tape can address compliance**
 - *LTO-4 has WORM capability*




To Find Out More

- This webinar is based on a paper written by today's presenters and published on 2/13/2008
 - See URL at bottom of this page
 - We have not presented all of the details included therein
- About 600 Clipper bulletins are available at www.clipper.com/publications.htm
- Q&A

Clipper Notes™

VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE
FOR THE INFORMATION TECHNOLOGY BUYER

Published Since 2007



Navigating Information Technology Horizons™

Report #TCG2008009

February 13, 2008

Disk and Tape Square Off Again — Tape Remains King of the Hill with LTO-4

Analysts: David Reine and Mike Kahn

Management Summary

Progress happens in many ways. At times, progress can be *evolutionary*, as technology advances one-step at a time, as we have seen in the development of the microprocessor, from Intel's 16-bit 8086 in 1978 to the 64-bit multi-core CPUs of 2007. It can also be *revolutionary*, as with the invention of the cell phone that literally turned the business world upside-down with mobile communications. Sometimes, we can have a revolutionary, new technology competing with the evolution of an older technology for the hearts, minds, and wallets of the business community. The newer technology usually has a lot of buzz about it, creating an excitement to be the first to implement, rather than rolling out the latest iteration of the legacy solution. Both sides make claims as to the functional advantages of their solution. Both sides also readily admit that there are additional costs associated with the new technology; replacing a legacy system is expensive. The question that persists, however, is "how much?" The CIO may be anxious to push forward with the newer technology, but the CFO may want to apply traditional measuring sticks to the acquisition. The enterprise needs to determine the total cost of ownership (TCO) of both solutions and the impact on the enterprise bottom line before making a choice.

This very situation exists in today's mid-sized data center where new technologies are competing with tape, *again*, for the favors of an I.T. staff struggling to support a rapidly expanding storage base while protecting the assets of the enterprise. The purpose of this bulletin is to review these new technologies against the value of retaining, or upgrading, your current tape environment. To do this, we will try to establish a framework of comparison – using a backup/recovery and archive scenario – between the latest disk technologies and the latest iteration of tape storage, LTO-4 technology.

We chose backup and archive for two reasons. First, every business does it, or at least should! Second, it is an excellent example of a very large and growing storage requirement, from which many conclusions can be drawn. For example, once you have one or more disk arrays deployed in the solution, adding another will not change the cost/TB by very much, except for the underlying reductions in cost per terabyte that come with time. For a tape library example, if you assume that the enterprise initially acquires a configuration with maximum tape automation and minimal number of tape drives (i.e., more than enough room for later growth), then incremental costs for additional drives and media are largely linear. While this might front load, somewhat, the total costs of ownership for tape, it does show a very favorable payback over disk for long-term storage. Our finding is that for long-term storage over our five-year study period, the cost of disk is about 23 times that of tape, while the cost of energy for disk is about 290 times that of tape. Read on, for the details.

The Cost Ratio for a Terabyte Stored Long-Term on SATA Disk versus LTO-4 Tape is about

23:1

For energy cost, it is about

290:1

Read this bulletin to understand what this really means.

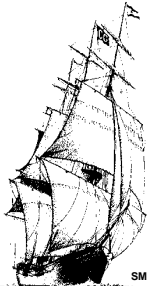
IN THIS ISSUE

- Assumptions for a Mid-Sized Enterprise's Long-Term Backup Scenario.....2
- Fundamental Building Blocks3
- The Disk Array Solution4
- The Tape Solution5
- Conclusion.....8

The Clipper Group, Inc. - Technology Acquisition Consultants • Strategic Advisors

888 Worcester Street • Suite 140 • Wellesley, Massachusetts 02482 • U.S.A. • 781-235-0085 • 781-235-5454 FAX

Visit Clipper at www.clipper.com • Send comments to editor@clipper.com



<http://www.clipper.com/research/TCG2008056.pdf>