

In Search of the Long-Term Archiving Solution — Tape Delivers Significant TCO Advantage over Disk

Analysts: David Reine and Mike Kahn

Introduction

Forever is a very long time. In today's hyper-paced society, we have a lot of trouble dealing with today and the short-term demands of our lives. Few people in their twenties are thinking about their lives three or four decades from now. If they are wise, then they might be putting money into a retirement plan or college fund for their kids. Beyond that, they live in the present — *doing what seems to make the most sense — right now*. However, ask them if they have photos or videos that they want to keep forever and, suddenly, forever has a more tangible meaning — a number of decades, for sure.

This bulletin is about preserving valuable digital data for a long time, usually somewhere between a decade and a lifetime, and possibly forever. Our focus is on the compounding quantities of commercial and institutional digital data, which are deemed so important that someone in the enterprise says: *We need to keep this forever!*, which means for longer than can be imagined in normal, 21st-Century human terms. **The question is how to do that, reasonably and in the most cost-effective way, thinking not just about what to do now and in the next few years, but what to do now to prepare for the ensuing decades of continuing data growth, technology change, and increasing long-term preservation requirements.**

The problem is very real but is different than keeping data for a required retention period, say emails or office documents for a seven-year statute-of-limitation period. In many cases, such data is retained *out of fear*, because it has to be kept to meet statutory or other requirements. Retention is seen as a liability (and a reaction to potentially severe penalties) that must be managed to limit exposure, as opposed to retention of an asset so valued that it must be preserved for a long time. There are many digital assets that might need to be kept for a long time. A good example is medical records, especially medical images. Depending on the jurisdiction, medical records may need to be retained for as long as the life of the patient plus several decades. This fits within our concept of *forever*.

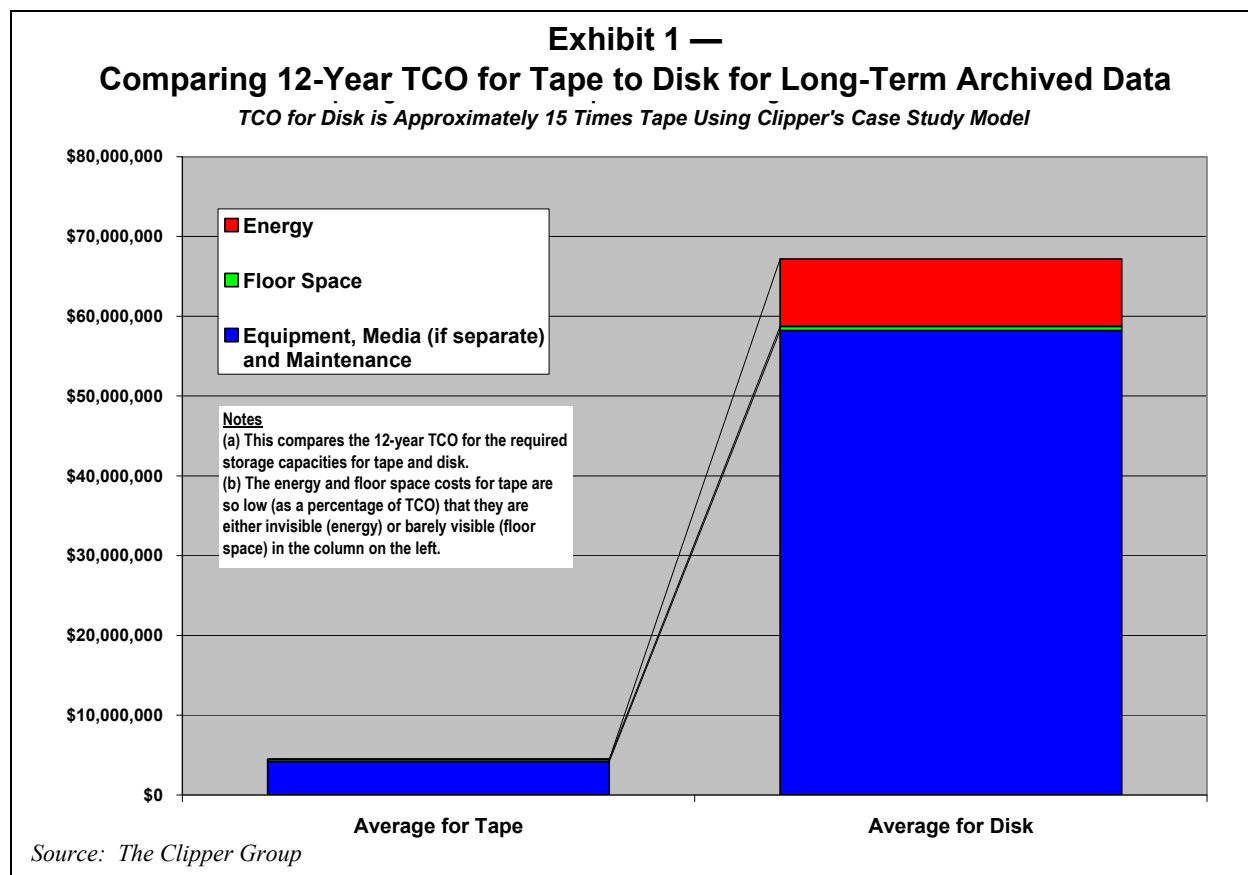
This bulletin documents the results of a major Total Cost of Ownership (TCO) study funded independently by The Clipper Group, Inc., the publisher of this report. It focuses on the underlying media on which data is kept for now and for the foreseeable future. We have designed it to be a generic study, one that is not applicable to just a single problem area or industry. We made many assumptions to simplify our model and to address the many unknowns that *forever* represents and these are described herein. **This study is a comparison of disk versus tape as the predominant storage media for archiving, or the long-term preservation of data.**

In summary, we have documented 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. Read on to learn how we came to these conclusions, including the important assumptions and details that framed our TCO analyses.

Major Conclusions at a Glance

For long-term archiving of digital data:

- The average disk-based solution costs 15 times the average tape-based solution
- The cost of energy alone for the average disk-based solution exceeds the entire TCO for the average tape-based solution



Management Summary

In this section, we will give you the summary answer to the question: ***What is the best media for the long-term storage of digital data, taking into consideration the back-end costs associated with storing it?***

- **The bottom line answer is that tape is always much less expensive than disk and always uses much less energy, when measured on a per petabyte basis.**

However, if the time required to retrieve any of your vast, long-term archived data needs to be a

few seconds or less¹, then archiving to disk may be best but will be much more costly. **In fact, the best solution is really a blend of disk and tape, but – for most uses – we believe that the vast majority of archived data should reside on tape.** Unfortunately, for those readers seeking a quick answer, there is a lot of study methodology, data, equipment details, and economic assumptions (along with more than a few caveats) that will take many pages to explain. If you want to understand how we did what we did, please read the entire report.

So, how big is the economic advantage for tape?

- **As Exhibit 1 (above) shows, over the 12-year span of the study's model, using our assumptions, the cost of disk is approximately 15 times that of tape.**

This represents many tens of millions of dollars in our case study scenario. It should be noted here that all pricing comparisons were calculated based upon configuration data and list prices² provided

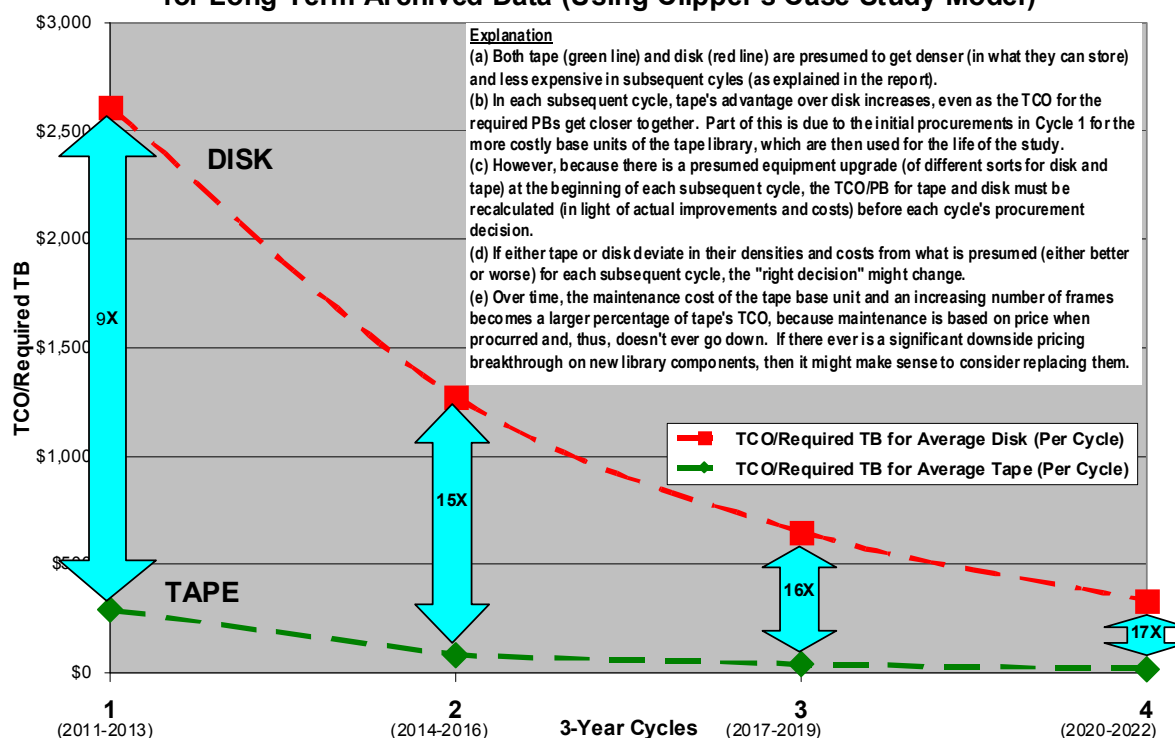
¹ The concept of time will be discussed later, as will the caching assumptions.

² While you can get a valid sense of street pricing on consumer electronics by surfing the Internet, this is not true for high-end IT procurements. Thus, we have chosen to use list pricing as the

IN THIS ISSUE

➤ Introduction	1
➤ Management Summary	2
➤ TCO Methodologies	5
➤ The Long-Term Archiving Challenge	9
➤ Assumptions for an Enterprise's Long-Term Archive	14
➤ Conclusion	17
➤ Appendix A – Assumptions	19

Exhibit 2 —

TCO/Three-Year Cycle (for TBs Required in Each Cycle)
for Long-Term Archived Data (Using Clipper's Case Study Model)

Source: The Clipper Group

by the vendors. Our disk comparison point began with 2TB disk assembled into fully-configured arrays, which can be considered Tier-2 storage. Our tape comparison point began with *Ultrium* generation-5 (LTO-5)³ tape cartridges and drives and housed in a fully-configured, enterprise-class *Automated Tape Library (ATL)*. Because we found that the cost of disk is about 15 times that of tape, we fully expect to see more discounting on the disk acquisitions than for tape⁴.

- **Nonetheless, in every dimension, the TCO of the tape solution was found to be less expensive than the TCO of the disk solution for long-term data retention, especially for energy consumption, where disk**

consumes 238 times as much energy as tape under assumptions that lean toward favoring disk.

As shown in Exhibit 2, above, the Total Cost of Ownership (TCO) for each cycle goes down significantly for both disk and tape as we progress through the study period. This is to be expected, as we presume that disk and tape densities continue to improve.⁵ However, do note that the TCO ratio of disk to tape continues to increase, mostly because there is so much more data to be archived in each ensuing period and because disk costs predominantly are linear while tape costs are not.

In a nutshell, what did we do? We compared low-cost, high-capacity disk solutions against high-end ATLs using LTO tape. We started with current generations of disk and tape technologies and made future generation assumptions based on

best available common reference point. Additionally, if we chose to use the same discount across the board, the ratios would have been the same.

³ When we say *LTO-5*, we are referring to *Ultrium generation-5*, etc. We will use the shorter "LTO-x" notation, where "x" refers to the generation of LTO tape, because it is more compact.

⁴ We have endeavored to make the configurations and pricing equivalent for all vendors in the same class of product. We are confident that they are equivalent within plus/minus 1% (in dollar terms). Given that discounts should be expected for multi-million dollar procurements, the potentially small configuration differences can be considered as insignificant to your procurement decisions.

⁵ Be aware, however, that there is a lot of discussion in the storage community regarding whether rotating disk drives will be able to continue to double the amount of data stored in a single drive or have the response time necessary for tomorrow's processors. If disk is unable to continue to double its capacity every three-or-so years, as we have presumed, then disk will be even more costly than discussed in the management summary, further reinforcing our recommendation to use tape for long-term archiving.

Exhibit 3 —**A Quick Look at Our Assumptions**

These are explained in *Appendix A*.

#1 – Usable Capacity

- Assumption – Both tape and disk are filled to 85% of their usable capacity.
- *Model Bias – Favors disk, potentially significantly*

#2 – Costs of Transition

- Assumption – Disk and tape both have negligible costs for transition between generations.
- *Model Bias – Favors disk, somewhat*

#3 – Energy

- Assumption – Disk and tape consume energy at their maximum operating level.
- *Model Bias – Favors disk, somewhat*

#4 – Compression

- Assumption – Data is uncompressible.
- *Model Bias – Favors disk, potentially significantly*

#5 – Maintenance

- Assumption – All hardware is maintained 24 by 7 with a 4-hour response time.
- *Model Bias – Favors disk, somewhat*

#6 – Floor Space

- Assumption – All hardware requires only the specified floor space, including space needed to provide access to all functional panels and components.
- *Model Bias – Favors disk, somewhat*

#7 – Excess Capacities

- Assumption – Disk and tape solutions are bought in optimal capacities (i.e., the Largest Building Block or LBB, as described in the body of the report. Additionally, tape cartridges are procured in an economic order quantity of 500.
- *Model Bias – Favors disk, somewhat*

#8 – Three-Year Life Cycle

- Assumption – Disk solutions are replaced completely at the end of each three-year cycle, while tape drives (but not the library) are replaced similarly on a three-year cycle.
- *Model Bias – Favors disk, somewhat to potentially more*

#9 – Acquisition Costs

- Assumption – Configurations and terms are fundamentally equivalent for all vendors for the same class of product, and priced at list by the vendors.
- *Model Bias – Model Has Small Bias Favoring Disk*

#10 – No Personnel Costs

- Assumption – No IT personnel costs are included in the model.
- *Model Bias – Model Has No Bias to Small Bias Favoring Disk*

published road maps and historical extrapolations. See Exhibit 3, to the left, for a summary of our assumptions as they apply to our analyses.

What were the most important assumptions? First, we assumed that an *Archiving Appliance*⁶ controls the indexing and placement of data on the long-term storage media, whether disk or tape. **Since the same appliance is used in both cases, we did not include its cost, or the cost of the disk that serves as its operating cache⁷, as part of the TCO for disk or tape.**

Second, we made important assumptions about the nature of the data being archived. We assumed that it was *blob-like* digital data that could not be compressed or deduplicated⁸. One example is scientific data streams, such as soundings from geophysical explorations, astrophysical data from radio telescopes, and detailed medical images.⁹

- **This kind of data is either too expensive to reacquire or cannot be reacquired at all, so it is presumed to be very valuable and, thus, may need to be kept for a long time, measured in decades, at least.**
- **This type of data is also the most challenging for archival purposes because it usually does not benefit much from deduplication or compression technologies and, therefore, represents a straight comparison between disk and tape as the media of choice for long-term data retention.¹⁰**

Third, we assumed that a three-year procurement cycle exists for both disk and tape and have based our 12-year model on four three-year cycles, as shown earlier in Exhibit 2.

- **We believe that this is sufficient to draw a long-term conclusion but, more importantly, it clearly calls for tape to be the media of choice for long-term archiving**

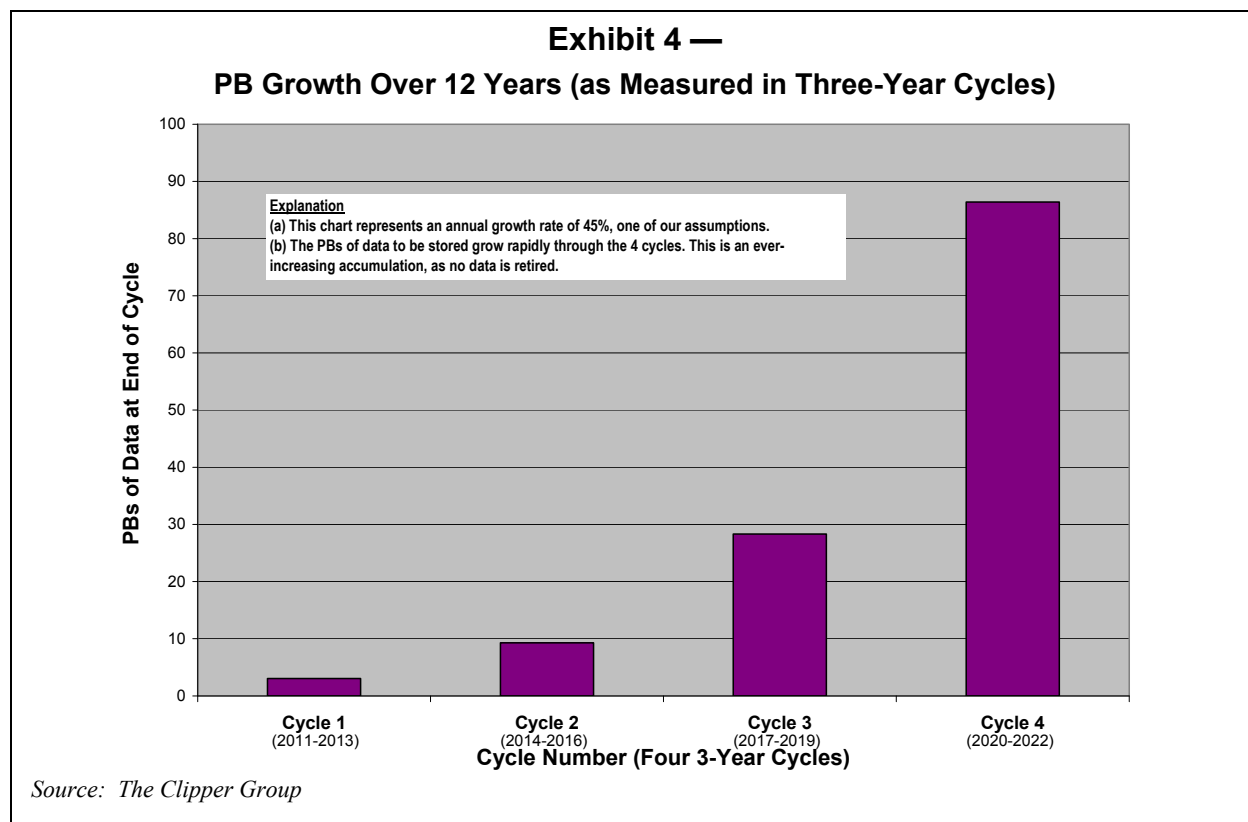
⁶ When used with leading capitalization, “Archiving Appliance” refers to a collection of hardware and software that manages the data being preserved. More details appear later in this report.

⁷ Depending on the application and budget, the operating cache might hold from a couple of months to several years of the most-recently used and/or more-likely-to-be-used data.

⁸ This lack of deduplication is not as big a deal as you might think, as will be explained later.

⁹ There are many other examples, as described later.

¹⁰ If your data is mostly office documents, heavily duplicated and/or capable of being compressed significantly, adjustments will have to be made to our model, which focuses on irreducible digital data. Do recognize that the adjustments go both ways, since it costs extra to deduplicate and/or compress disk files and objects, while LTO tape has compression built-in at no extra charge. (In addition, you will have to retain and maintain the necessary versions of the data deduplication and decompression software for future retrieval of your archived data.)



for the next three years, the only period for which current decision-making is required.

- For the first cycle (assumed to start in 2011), the TCO ratio of disk to tape is 9:1, less than for the 12-year period because some ATL acquisition costs have been frontloaded at the beginning of the first cycle.

Lastly, we tried to be fair in all of our assumptions¹¹, bending over backward to reduce disks' presumed higher costs.

- *In the end, such bias toward disk really didn't make any difference to our conclusion that tape is significantly less expensive than disk for long-term archiving, since the TCO differential was so great. Read on for some history, the assumptions, and the details.*¹²

We consider our assumptions to be very conservative. However, you need to compare your situation and IT procurement practices to our business case. (See Appendix A, for a full discussion of our assumptions.)

TCO Methodologies

At this point in the paper, we had a chicken-or-the-egg decision to make. We went with the egg (describing the TCO methodology and results) and delayed the chicken (explaining the business case and further assumptions and details) until later.

Sources for Specifications and Costs

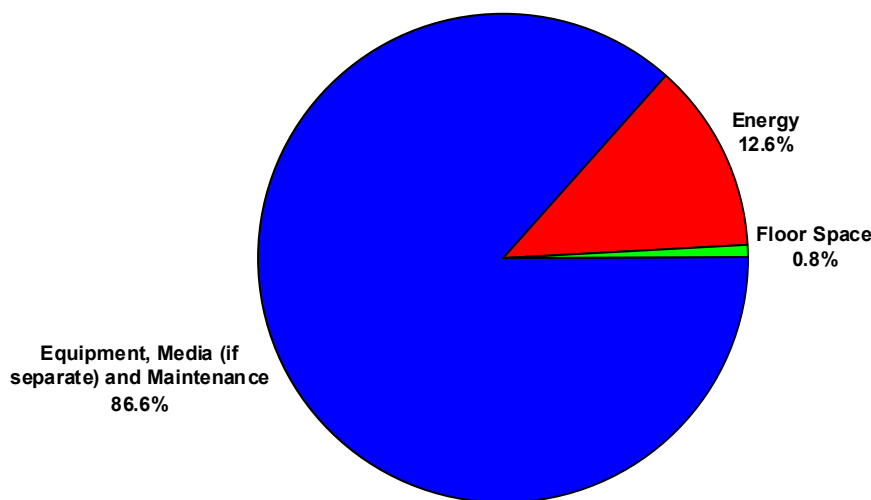
We have collected data from multiple providers of low-cost, high-capacity disk array solutions¹³, which today are using 2TB SATA disk drives as the technology of preference for this application. Each disk vendor provided us with its largest building block, which, as would be expected, is differently sized for each model in the study. For tape, we collected data from multiple providers of enterprise-class ATLs with LTO drives. For comparison purposes, we took tape and disk configurations of differing capacities and densities (in each LBB, or Largest Building Block) and took an average of their costs per PB within each type.

- **Thus, when we speak of TCO for disk or tape (or the cost components within each), we are referring to the average cost for**

¹¹ These are described and explained in detail in the sections that follow, and in Exhibit 3, which is on the previous page.

¹² If you want to skip the details, this is a good place to stop.

¹³ Three or more vendors provided data for disk and for tape. As part of our promise not to disclose their identities or their confidential data, we will leave the exact number unstated.

Exhibit 5 — Distribution of the 12-Year TCO for Disk by Budget Category**TCO = \$67,193,181***Source: The Clipper Group*

the disk solutions or the tape solutions that were included in our study.

Determining the TCO for Disks

Because we are starting out with a 2TB disk building block, we believe that it is mandatory to configure the disk array with a RAID-6 architecture to ensure against a second disk failure in any given RAID group. We have also configured one spare drive per 14-16 disks in a RAID group to facilitate the rebuilding of any failed drive.

Factoring in both RAID parity drives and hot spares, about 87.5% of the LBB's raw disk capacity is available for storage. As stated earlier, we are assuming a utilization factor of 85% of available capacity for all storage, leaving about 74.4% of the raw disk capacity usable for archiving data.

In order to meet the 12-year pattern of data growth at 45% per year, as shown in Exhibit 4¹⁴, the Archival Storage needs to grow from 3PBs for Cycle 1 to more than 86PBs for Cycle 4.¹⁵ The

average TCO for the disk solutions over the 12-year study period is more than \$67 Million.

Cycle 1:	\$2612/PB	\$7,162,915 TCO
Cycle 2:	\$1275/PB	\$11,851,642 TCO
Cycle 3:	\$653/PB	\$18,498,122 TCO
Cycle 4:	\$334/PB	\$28,880,442 TCO
Total:		\$67,193,081 TCO

The costs are dominated by the hardware and its maintenance, as shown in Exhibit 5, above. Not surprisingly, it costs more to store a PB of data on disk in the first cycle than in ensuing cycles, because we assume that the procurement costs per PB will decline in each cycle. Also, the cost of maintenance is higher for disk than for tape, because disk drives are more likely to fail than tape drives and so many more of them will be deployed over the study period.

Maintenance Costs for Disks

Maintenance policies and warranties vary by vendor and degree of bundling. Many of the disk arrays come with three years of "standard maintenance" that may vary in coverage. Thus, we have not broken out maintenance costs for disks (or tapes) because of the incongruent warranties and base-rate service times. In our cost accounting, we have normalized all of the maintenance coverages and made cost adjustments to bring each disk solution to a 24x7x4-hour response time for three years.

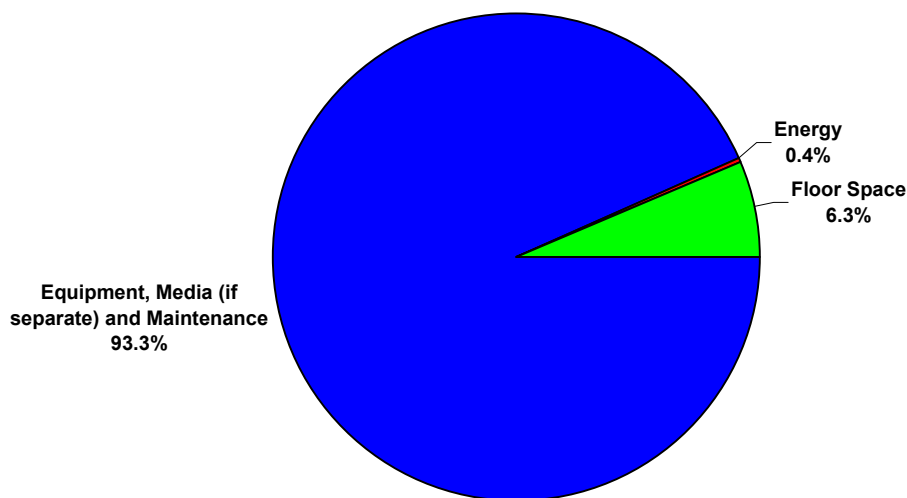
¹⁴ And discussed further later in this report. Be aware that a higher data growth rate will mean vastly more data will be accumulating and the economic advantage of tape over disk will widen.

¹⁵ We have used the following acquisition protocol. All procurements are made at the beginning of the first year of each cycle. Additionally, we are procuring only LBBs, so there is some capacity included that is beyond the actual data requirements for the end of the cycle. This results in an *overprocurement* or maybe it should be described as procurement by means other than *just-in-time*. However, since we are not calculating the time-value of money (as explained later), this does not affect the TCO. It also lessens the concerns about assuming that the drives and cartridges are filled to 85% of capacity, which is beyond what most vendors would recommend for disk. With the

Archiving Appliance doing the writing, and the extra capacity acquired by rounding up to the next full LBB, the maximum capacity required likely will not be reached before the very end of the cycle.

Exhibit 6 — Distribution of the 12-Year TCO for Tape by Budget Category

TCO = \$4,458,228



Source: The Clipper Group

Cost of Energy for Disk Arrays

With a presumed cost of \$0.15 per KWH right now¹⁶, with a presumed annual increase of 5% for the second through twelfth years (and under all of our other model assumptions), the following KWHs (an average, including power and cooling) and projected costs will be consumed during the 12-years for disk.

Cycle 1:	795,116 KWH =	\$375,692
Cycle 2:	1,170,932 KWH =	\$640,474
Cycle 3:	1,800,945 KWH =	\$1,140,350
Cycle 4:	2,761,164 KWH =	<u>\$2,023,940</u>
Total:	19,584,469 KWH =	\$4,180,456

With an all-disk archived storage solution (in one site only), the disks will consume almost 20 GWH of power over 12 years. That's a very big number to consider, since a Gigawatt Hour equals a million Kilowatt Hours. However, as you will see shortly, **the total energy cost for disk approaches the TCO (all costs) for tape. That is mind-boggling.** So don't minimize what is represented by that modest slice of the pie for energy in Exhibit 5.

Floor Space for Disk Array

Assuming a rate of \$200 per square foot of data center space¹⁷, the average disk solution will consume the following square feet and costs over the four cycles:

Cycle 1:	326 SQFT =	\$65,183
Cycle 2:	479 SQFT =	\$95,706
Cycle 3:	723 SQFT =	\$144,628
Cycle 4:	1120 SQFT =	<u>\$224,072</u>
Total:		\$529,589

Over the study period, floor space represents just .8% of the TCO for disk.

Determining the TCO for Tape Libraries

As with disk, we have collected data from multiple providers of large-scale automated tape libraries, using Ultrium (LTO) tape as the technology of choice, starting with the current open state-of-the-art tape, LTO-5. Each vendor provided us with standard configurations for their tape offerings; each of which, as would be expected, is different from the other offerings. We used configurations of differing frame densities, and modeled accordingly, to configure and cost out each offering, and then took an average of the costs for comparison purposes. (See Exhibit 6 above.)

One base frame has been configured for each tape solution, with a variable number of expansion frames required, depending upon the capacity provided by each vendor to satisfy the number of cartridges required for each cycle. The base frame, and the internal robotics, will remain constant throughout the course of the study period, with expansion frames being added in the first year of each cycle. We have configured each frame with its maximum number of slots to simplify the expansion process for each cycle, and to simplify the TCO study for us. Based upon the capacities de-

¹⁶ To be described later.

¹⁷ The cost per square foot per year for data center spaces varies significantly by geography, location, and market. However, even at \$100 or \$300, the impact of the cost of floor space on the TCO is minimal, in relation to the equipment and maintenance.

rived earlier in this report, we have configured sufficient frames to support the minimum number of slots (and cartridges) and drives required for each cycle:

Cycle 1: 2,903 Slots & 3 LTO-5 Drives

Cycle 2: 6,863 Slots & 3 LTO-6 Drives

Cycle 3: 10,823 Slots & 4 LTO-7 Drives

Cycle 4: 18,743 Slots & 8 LTO-8 Drives

We have assumed that each new drive and each new generation cartridge will become available at the same cost as the preceding generation. We have configured enough drives to handle the calculated writing requirements through the ATL, plus one spare tape drive per ATL to protect the data center in the event of a drive failure, and to be able to handle spikes in performance requirements. If your retrieval or writing needs require significantly more drives than we have specified (but without increasing the amount of data to be archived), they likely will increase tape's TCO by no more than 20% and probably much less.¹⁸

The average TCO per PB stored and total TCO on tape for each of the Cycles are:

Cycle 1: \$290/PB \$885,533 TCO

Cycle 2: \$83/PB \$769,620 TCO

Cycle 3: \$40/PB \$1,142,786 TCO

Cycle 4: \$19/PB \$1,660,288 TCO

Total: \$4,458,228 TCO

As might be expected, energy and floor space requirements and costs were small when compared to the costs of the ATL and cartridges. However, the cost of media becomes a significant element in the TCO equation and is included as part of the total cost of equipment (TCE), which also includes maintenance charges needed to meet our 24x7x4-hour requirements.

Maintenance Cost of Tape Library

As with disk, maintenance policies vary from vendor to vendor. Drives that have been replaced with the next generation are retained as a luxury but without maintenance. When they fail or there is no room for them in the ATL, they are no longer used. All calculations to determine the number of drives needed assume that only *the latest generation of drives would be used to meet the writing volume requirements*. The maintenance cost is buried in the TCE amount and not listed separately because the vendors' warranty and maintenance terms and conditions and pricing mechanisms were not uniform.

¹⁸ Even if you doubled the cost of equipment, maintenance, and energy (everything but the cost of the cartridges), tape still would be one-tenth the cost of disk. That doubling would pay for a lot of additional drives and robotics, etc.

Looking at Tape vs. Disk — and the Origins of this Study

For many years now, the naysayers in our industry have been predicting the death of magnetic tape in the data center. For many reasons, such as *recovery time objectives* (RTO), length of the backup window, and the reduced cost of Tier-2 disk, many solution providers have been creating data protection and archiving solutions based upon a *disk-to-disk* (D2D) architecture. Conversely, tape library vendors have long espoused the financial advantages of including tape in a *disk-to-disk-to-tape* (D2D2T) scenario, on a cost per terabyte basis, not to mention the significant TCO advantages resulting from lower administrative costs, reduced floor space requirements, and lower energy costs to both run the IT infrastructure and cool the data center. In an attempt to quantify the differences between running data protection in a D2D environment versus running the same applications with a D2D2T architecture, Clipper published (in February 2008) and updated (in October of that year) an extensive study of the differences in the TCO for the preservation of data in a backup and recovery environment. (See the issue of *Clipper Notes* dated October 21, 2008, entitled *Disk and Tape Square Off Again – Tape Remains King of the Hill with LTO-4*, available at <http://www.clipper.com/research/TCG2008056.pdf>.)

The results of this study were stunning, not necessarily in the conclusion that tape was less expensive than disk as a long-term solution, but in the extent of that difference. In 2008, in comparing disk subsystems with 750GB SATA-2 drives against tape libraries with LTO-4 drives and cartridges, we concluded that the TCO ratio of disk to tape *in a backup environment* was 23:1. When we looked exclusively at the energy consumption between these two environments, we saw an amazing cost ratio of 290:1. These results not only drew the attention then of those within the enterprise data center community, a community experiencing a great deal of pain trying to manage an IT budget in an era of fiscal panic, but to this date continue to draw much attention, as seen by the report's continued distribution and downloading.

With that study in mind, this time we set out to determine the impact that more than two years have made on the technical and economic differences between tape and disk. However, different from our 2008 study where we focused on backup, this time we are focusing on an archival environment. A new backup study is planned for the first half of 2011 and is expected to include data deduplication.

Cost of Tape Media

In order to ensure the best pricing for media, we have established an acquisition pattern to obtain cartridges in lots of 500, with the acquisitions taking place annually. We have established a price of \$90 per cartridge from a reputable media source, which anecdotal evidence suggests is high. The number of tapes needed and ordered is:

Cycle 1: 2391 Needed 2500 Procured
 Cycle 2: 2296 Needed 2500 Procured
 Cycle 3: 3500 Needed 3500 Procured
 Cycle 4: 5634 Needed 6000 Procured

Even though most tape cartridges are certified for a 30-year life, we assume that the data center will migrate (in Cycle 4) their LTO-5 cartridges (from Cycle 1) to LTO-8 when that technology becomes available, rather than retaining older tape drives within the infrastructure for read-only purposes. This results in a consolidation of slots during Cycle 4 (as fewer cartridges and slots are needed when using the much denser LTO-8 cartridges). However, there is a significant increase in the number of cartridges for that cycle because of the migration. The cost of the cartridges is buried in the total cost of the equipment.

Cost of Energy for Tape Library

With a presumed cost of \$0.15 per KWH right now¹⁹ with an annual increase of 5% for the second through twelfth years and our model assumptions, the following KWHs (an average, including power and cooling) and costs will be consumed during the 12-years for disk.

Cycle 1: 19,745 KWH = \$3,110
 Cycle 2: 19,920 KWH = \$3,632
 Cycle 3: 21,304 KWH = \$4,497
 Cycle 4: 25,965 KWH = \$6,344
 Total: 86,934 KWH = \$17,583

While the ATL and drives do consume more energy than is required to make your morning coffee and toast, there is no comparison to the massive energy required to drive the disk solution (as described on page 7). If we now compare this to the energy consumed by the disk solution, we see a ratio of 238:1. Try justifying that for infrequently-accessed archived data.

Floor Space for Tape Library

Each vendor has a slightly different approach to frame placement and panel access requirements in the data center. We adjusted for that and used the same \$200 per square foot as we used for disk and assumed that the extra space for additional frames is created (added) in the cycle when it is

needed. Here are the averages for tape for floor space.

Cycle 1: 210 SQFT = \$42,093
 Cycle 2: 289 SQFT = \$57,710
 Cycle 3: 392 SQFT = \$78,485
 Cycle 4: 516 SQFT = \$103,101
 Total: \$281,389

Floor space for tape costs a little more than half of what disk costs (\$529,589...from page 7).

The Long-Term Archiving Challenge²⁰

Data Requirements

This is a study based on plausible requirements, amalgamated from numerous enterprise sources.

- **Our goal is to come up with some answers to the cost of disk versus tape question for long-term archiving.**

We assume that there will be a lot of data to store. While we assume that the existing amount of real data at the beginning of the study period is a petabyte, this is an arbitrary assumption. We assume that 1PB is large enough to define our problem as *large scale*. The results would have been similar had we started with 2 PBs or 10 PBs of data. We feel confident that we have started with a sufficiently-large-enough amount of data to justify the largest-capacity disk and tape solutions. Thus, we believe that we have some *near steady-state* costs (represented on a per petabyte basis), which can be extrapolated to larger collections of data.²¹ Thus, it should be easy for the reader to compare his or her actual large data size against the 1 PB in this study, to get a reasonable multiplier for the costs and other factors.

Beginning with a basic archive of 1PB and a growth factor of 45% annually²², we get a requirement for 1.45PB of archive storage by the end of year 1, with rather explosive growth to an accumulation of over 86PB of archive storage by the end of year 12. (As shown in Exhibit 4, earlier, which represents the amount of data at the end of each three-year cycle.) From where we sit at the end of

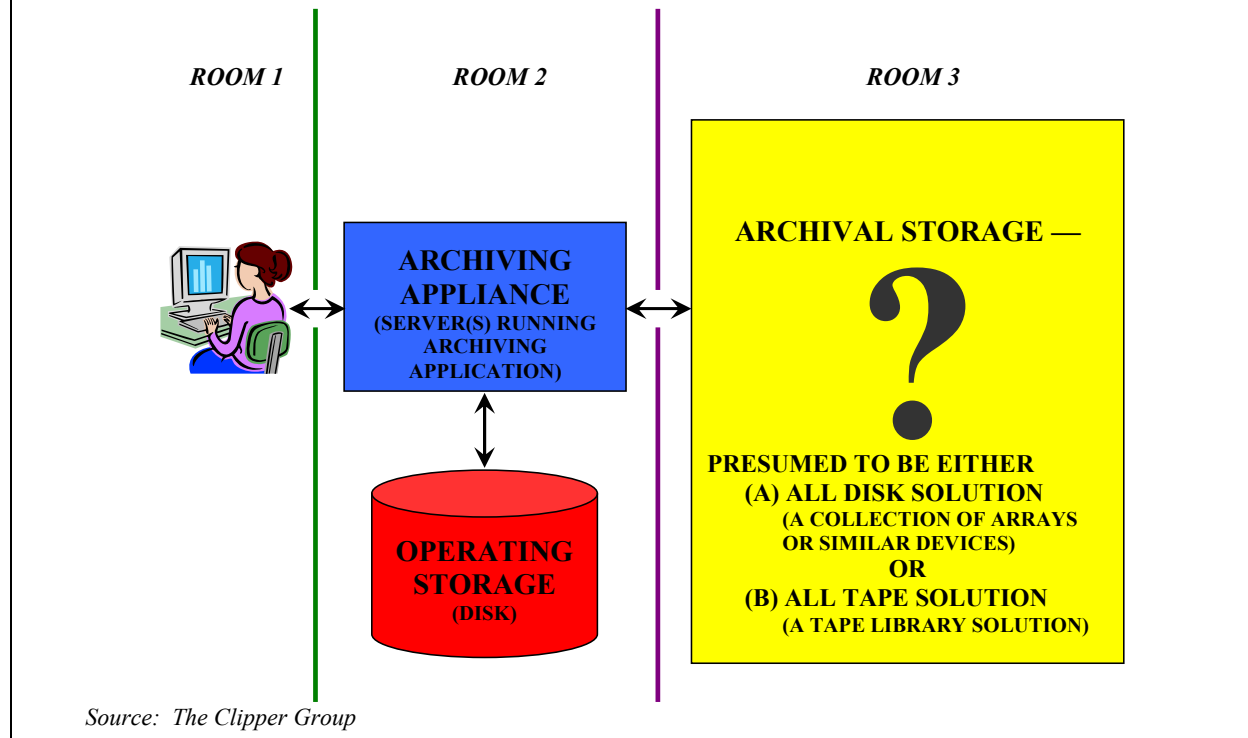
²⁰ This section is the “chicken” (the explanation of why we did what we did) that has followed the “egg” (the summary of findings and conclusions). If you are not interested in the details, go directly to the Conclusion on Page 17.

²¹ Near steady state comes from our use of large, full-capacity (or close to full) building blocks for both disk and tape. We call these *LBBs*, for *Large Building Blocks*. If you have much more data than we have modeled, you still will be building your solution using similar LBBs.

²² We often see annual growth rates of 40 to 50 percent for large enterprise averages. So, we picked 45%. This is not very scientific and well short of the extreme where data growth exceeds 100% per year, but large enough to show the long-term effect of significant growth.

¹⁹ As discussed later.

Exhibit 7 — Dividing the Problem and Focusing on the Right Part



2010 and for the purposes of this study, this is big enough. Nonetheless, recognize that the bulk of the new data is stored in the later periods, so the later periods are when this study gets very interesting and where the costs really begin to accumulate. However, in the real world, you have to deal with the first cycle (i.e., the present) before the others.

- **Thus, what looks good for the next three years probably should be the primary basis for your current decision-making.**

Defining the Storage Challenge

Imagine an archiving process, where there are three rooms that are separated by walls, through which only networking cables pass. (See Exhibit 7, at the top of this page.) The rooms might be separated by many miles or only a few feet. A user sits at her desk in what we abstractly call *Room 1* and retrieves data (think *a file*, for simplicity) that has been stored by an *Archiving Appliance* in a prior action. She has no idea where the data was archived and can only judge the success of the archiving solution by three very important criteria:

1. *Whether it stored what it had been given,*
2. *Whether she gets back what was stored previously, and*
3. *Whether this happens in a timely manner.*

Somewhere beyond the user's view (*Room 1*), an Archiving Appliance sits in *Room 2* and man-

ages all that is necessary to store the data and accomplish these three criteria. Its vendor and brand are not important to this study but certain generalizations can be made about its characteristics.

The Archiving Appliance is presumed to have its own *Operating Storage*²³, whether this is built into the appliance or sits somewhere nearby (presumably in *Room 2*). The Operating Storage serves as an *archival cache* and, depending on the requirements, may hold data for days, weeks, months, or even years, in order to meet the users' service level requirements. Eventually, based on aging, lack of use, or other policies, the data needs to be archived to the least expensive storage available, for possible later use, or for legal or other reasons. In any case, it is presumed that the data is important, i.e., sufficiently valuable that it must be preserved for a long time, beyond the 12 years of the study period.

To preserve it, the Archiving Appliance moves it from the Operating Storage (disk) to the *Archival Storage*²⁴ (low-cost disk or tape, or a blend), which is situated separately, in this example in a third place called *Room 3*. The Archiving Appliance,

²³ When used with capitalized first letters, "Operating Storage" refers to the disk cache that sits in Room 2 with the Archiving Appliance.

²⁴ When used with capitalized first letters, "Archival Storage" refers to the tape libraries or disk arrays that sit in Room 3 and contain the data being preserved at the lowest-cost tier.

like the user, can only judge the adequacy of the Archival Storage in terms of (1) whether it stores what it is given, (2) whether it is able to retrieve it without loss, and (3) whether this happens in a timely manner.

Let's assume that the Archiving Appliance is dependable, of high quality, and, thus, always stores and retrieves, or caches, properly, regardless of where it is storing the data to be preserved (e.g., on disk or tape). Thus, the only remaining critical question, or success factor, is *whether this happens in a timely manner*.

Some applications demand that data be retrieved instantly, such as police records, especially when a pursuit is involved. Other applications are price sensitive, due to limited budgets or huge volumes of archived data, and its users are willing to wait (to varying degrees) for the data to be made available (at a lower cost). These other applications might be seasonal or occasional and its users may be willing to wait longer for the data to become available, especially if its use was not anticipated. In the end, for many applications and users, there is a balancing act between *time-to-availability of the needed data* and the *TCO of storing the archived data* in Room 3.

Typically, the more data that you have, the more likely you are to be looking for ways to reduce the *unit cost* per stored PB, i.e., to archive and retrieve it.

- **This study focuses on those enterprises that archive very large amounts of data for a long time and are looking for ways to keep the costs down.**

Defining the Time Frame

We started this study thinking about *data that has a very long useful life*. In discussions with enterprises and storage vendors, retention periods of both 50 and 100 years were mentioned. Since most of us won't be around in 50 or 100 years, we colloquially began to talk about *forever*, which we defined as longer than we can imagine or beyond which we can plan reasonably.²⁵

There are many kinds of data that need to be kept for a long time. However, for this study we needed to quantify the data in some meaningful way. As researchers often do, we made preliminary assumptions (as shown in Exhibit 3 and Appendix A and also discussed in the body of this report) and decided to see whether we reached some meaningful results. Fortunately, we did – but before we discuss the results further, we need to

explain the methods and rationales that we used.

Trends in Storage Technology

First and foremost, storage technology continues to evolve. We have historical data on the evolution of spinning magnetic media, both disk and tape, although not a lot of historical data for Solid-State Drives (SSDs). It is difficult enough to predict with accuracy what might be the prevailing storage technologies and alternatives ten or fifteen years from now, much less how they will evolve over the next 50 or 100 years. Thus, while we presume for this study that the data will need to be retained for a long time, it seemed unreasonable to try to drive a model based on what we know today and extrapolate it for the next 50 or 100 years. Similarly, we did not assume that today's data would remain on its original media forever.

- **This study is not about the physical durability of storage solutions but about finding the lowest total cost for practically keeping data for a long time.**

So, what do we know today and what can we reasonably predict about disk and tape? Let's start with disk. Disk capacities for 3.5 inch, 7200 RPM lower-cost, high-capacity SATA drives have been doubling in capacity about every three years²⁶. By and large, the doubled capacity of each new generation has come with little or no increase in price. 1TB SATA drives had been widely deployed throughout data centers around the world, but 2TB SATA drives now are the procurement choice for low-cost storage. 3TB SATA drives are just beginning to appear in the consumer space but we consider this an interim step to the next generation, 4TB enterprise-class SATA drives. However, there is no reliable roadmap on which to base expectations for the availability of 4TB drives, and most forecasts made several years ago already have been proven wrong. In addition, the next generation of low-cost drives might be SAS²⁷ and may be housed in a 2.5" enclosure. Whatever the future may bring, we have to start with what is being offered today and all of our disk vendors chose to offer disk solutions based on 2TB SATA drives (without spin down for energy conservation).

- **So, for purposes of this study, today's low-cost disk is a 2TB SATA device.**

As shown in Exhibit 8 at the top of the next page, we presume that SATA drives continue to double

²⁵ In conversations with enterprise archiving managers, they often mentioned *forever* before we did.

²⁶ One might argue that it is really longer than that for widespread adoption by enterprises, but three years is as good a supposition as something shorter or longer is. See additional discussion in Appendix A.

²⁷ SAS=Serial Attached SCSI.

Exhibit 8 — Presumed Tape and Disk Capacities

Study Cycle	Presumed Years	LTO/Ultrium Generation	Tape Cartridge Capacity (Uncompressed) – in TBs	SATA Drive Capacity (Uncompressed) – in TBs
Cycle 1	2011-2013	LTO-5	1.5	2
Cycle 2	2014-2016	LTO-6	3.2	4
Cycle 3	2017-2019	LTO-7	6.4	8
Cycle 4	2020-2022	LTO-8	12.8	16

Source: LTO Program and The Clipper Group

in capacity every three years.²⁸

Storage TCO and energy consumption still remain among the most significant concerns to data center executives around the globe. The focus is shifting from storing everything to reducing what is being stored and for how long. There was a lot of buzz surrounding data deduplication in 2008 and, where appropriate, that is a cost-saving measure.²⁹ Other new technologies, such as self-encrypting disk drives, lend more credence to the capability of securing data on disk while not affecting application performance. Both of these are significant innovations but do not make a lot of sense for the scenario that we have described for long-term archiving of irreducible digital data held behind an Archiving Appliance.

Tape is different, in several ways. For our purposes, we are focusing on *LTO* tape, which has become the *de facto* open standard with drives and cartridges supported by many suppliers. Earlier this year, LTO technology advanced to its latest generation, *LTO-5 tape*, with dual partitioning to enable the creation of disk-like file systems on commodity tape.³⁰ With a native capacity of 1.5 TBs, LTO-5 tape has almost twice the capacity of LTO-4 tape, increasing from 800 GBs to 1.5 TBs – *uncompressed*.³¹ At the same time, LTO-5 technology has continued LTO tape's capabilities in the areas of dynamic WORM, compression, and encryption.³² All of the leading tape vendors are now shipping LTO-5 products.

- **So, it is safe to declare that – for open tape – we now are in the LTO-5 era.**

What is even more impressive for the continuing presence of enterprise tape in the data center is the extensive roadmap through *LTO-8* that

the *LTO Program*³³ has rolled out for increased capacity and improved throughput over the next decade or so.³⁴ (See again Exhibit 8, above.) So, we feel that it is safe for us to assume that each of the next three generations of LTO tape will double in capacity, as described in the LTO roadmap.

- **Based upon the generation jumping that we have seen over the last decade, we feel comfortable presuming that each new generation (of disk and tape) will come at 3-year intervals.**

Now we have some ground rules for projecting the future of disk and tape over the next decade, a period during which we do not expect to see SSDs replace the lower tiers of rotating disk storage, i.e., we expect SATA disks to cost less than SSDs on a TCO basis for many years. If we are wrong, then we would expect the price/capacity curve of SATA devices to continue, more or less, even if 7200 RPM SATA drives are replaced by SSDs.³⁵

- **However, many in the storage community often raise the concern that rotating disks cannot continue to double in capacity every three or so years. They may be right, in which case the costs for storing data archives on disk will grow much faster than we show in this report.**

However, they might be wrong, in which case our model would be closer to the truth. Some disk vendors made the same forecast for tape a decade ago, openly declaring the death of tape, but LTO roadmaps, both past and current, are proof that tape is alive and well. We will attempt to remove all of

²⁸ More on this assumption to follow.

²⁹ Due to the nature of the archive data being stored, data deduplication will not be considered in this study. However, we will include it in our follow-up study for the long-term storage of backup/recovery data, to be published in 1H11.

³⁰ However, the dual partitioning and the *Long-Term File System (LTFS)* are not used in our scenario.

³¹ More on compression to follow.

³² None of these capabilities are used, either. See Appendix A.

³³ The LTO Program was formed in 1997 and three companies – HP, IBM and Quantum – jointly oversee the development and roadmap of Linear Tape-Open (LTO) technology. LTO technology allows users to have multiple sources of product and media. The "open" nature of LTO technology also provides a means of enabling compatibility between different vendors' offerings.

³⁴ See *The Clipper Group Navigator* dated January 29, 2010, entitled *LTO Program Announces Next Gen Tape – LTO-5 Raises the Bar for Tier-3 Storage*, available at <http://www.clipper.com/research/TCG2010002.pdf>.

³⁵ This assumption is not without risk. Today, the technological and economic roadmaps for tape seem far more certain than for disk drives or SSDs.

the fear, uncertainty, and doubt (FUD) relative to the capabilities and longevity of magnetic tape as a medium for archiving. Thus, we have reached another simplification.

- **Both disk and tape are presumed to be improving every three years (or so), with the capacities doubling, while the price per disk drive or tape cartridge remains about the same as the previous generation.**

What's fair for one should be fair for the other, at least until better information is available.³⁶

Back to Forever

We have meandered a long way from defining *forever*. Driven by the LTO roadmap, we presume that LTO-6, LTO-7 and LTO-8 generations will come in years four, seven, and ten. Because of the common three-year cycle that we attribute to both disk and tape, we defined our study with a 12-year analysis period, broken down into four generations that we call *cycles*, with the capacities of tape and disk doubling at the beginning of each cycle. We felt that going any longer to quantify the TCO of tape versus disk would be fraught with too many indefensible assumptions. If the analyses do not deliver conclusive results by the twelfth year, then our study would have failed to achieve its primary objective. Fortunately, the answer was clear even after just three years of modeling.

While 12 years is not forever, it is long enough to draw some reasonable conclusions on the costs of keeping data for a long time, hence, the reason this study focuses on the TCO of long-term archiving. Think of this approach in different terms – in the next 12 years, you need to do something to store your valuable data.

- **If you need to act in 2011, then you must choose between today's tape and disk or a combination thereof. If something new appears or if tape or disk fails to keep pace as we presume, then you can make a reasonable change at the beginning of the second through fourth cycles.**

Fortunately, the Archiving Appliance will make these transitions largely painless and transparent, since it handles the transfers from one generation of technology to the next, even when transferring to a different storage technology.

What about the Costs of Data Protection?

Of course, there is more to consider. What needs to be done to mitigate the potential for a site-wide disaster or forced shutdown? There are three major choices from which to choose, plus a plethora of variations.

1. **A completely duplicated disaster recovery site sits far enough away to provide services (in short order) after the primary site fails.** In this case, the Archiving Appliance in the primary location would be linked to, and synchronized with, another Archiving Appliance at the remote location. There are three Archival Storage possibilities.

- (a) Tape is the archiving storage medium in the primary location and tape also is used at the remote site.
- (b) Disk is the archiving storage medium in the primary location and disk also is used at the remote site.
- (c) Disk is the archiving storage medium in the primary location but tape is used at the remote site.

The key characteristic of all of these alternatives is that the Archiving Appliance is responsible for keeping the remote site up-to-date.

2. **A peer-to-peer back up strategy is in place for the Operating Storage and the Archival Storage.** In this case, each storage solution is responsible for replicating itself to the remote location, although this also might be achieved by writing to the local and remote devices concurrently (i.e., this is a networking solution that also delivers replication).

3. **A traditional backup/recovery solution is deployed on the Operating Storage and the Archival Storage,** with the backup data stored remotely.

In a future Clipper study (targeted for the first half of 2011) comparing tape to disk for backup and recovery scenarios, we will investigate this subject more thoroughly. However, in this long-term archiving study, we focus on the scenarios listed in the first alternative above (a completely-duplicated disaster recovery site). We do this because we are most interested in comparing the alternatives economically and Scenario #1 allows us to do that simply.

First, our economic analyses are focused on Room #3 where the Archival Storage resides. Our cost analyses do not include the contents of Room #2, since that is the same regardless of what storage sits in Room #3. Thus, our decision to duplicate what is in Room #2 at a remote location (including

³⁶ As said several times in this report, procurement for the next three years of archiving is based on the current generation of technology and this cycle is far more important than any later cycle because it is real, here, and now. Later decisions will be based on data that is not available today.

whatever extra software and networking costs this will bring) does not affect our Room #3 analyses.

Second, with the costs of Room #2 assumed away (for the purposes of our comparison of tape to disk), the three sub-scenarios (1a, 1b, and 1c, from above) are all economic derivatives of the cost analyses that we did for Room #3 at the primary site.

(1a) If the tape solution is the same at both sites, then the cost of the Archival Storage at the remote site can be presumed to be the same as the costs at the primary site.

(1b) If the disk solution is the same at both sites, then the cost of the Archival Storage at the remote side can be presumed to be the same as the costs at the primary site³⁷.

(1c) If disk is used at the primary site for Archival Storage and tape is used at the remote site, then the two-site solution is a combination of the cost for disk at one site and tape at the other.

These are calculations that you can do yourself from the cost data that we have calculated for the primary data center Room #3.³⁸ The simplicity of doubling the cost of all-tape or all-disk Archival Solutions is easy to comprehend. If the cost of all-tape is less than all-disk, then the mixed mode remote scenario's cost will be between the all-tape and all-disk. However, if all-disk is less expensive, then there is no cost reason to use tape at the remote location, although, there may be data protection reasons to consider tape.

Assumptions for an Enterprise's Long-Term Archiving

Energy Assumptions

- **It must be noted here that the availability of, and access to, energy is not getting better, and the cost of that energy is not going down.**

Many data centers have already reached a maximum load when it comes to energy consumption. In reality, for many locations there is no more energy available to come through the existing power grid. If data center managers need to deploy addi-

tional storage devices within the IT infrastructure to preserve the mission- and business-critical data that has become the lifeblood of the enterprise, something else has to be removed! Moreover, the cost of that energy is going in the same direction as the cost of gasoline at your neighborhood gas station – *UP!* Can your IT budget withstand the hit that will be caused by the *next* energy crisis? The cost of energy is a significant factor in that budget³⁹, as the cost per KWH (Kilowatt Hour) continues to rise, with no end to price increases in sight.

Trying to fix a unit cost for energy consumption is difficult as access to energy and the cost of that energy varies wildly across the country, from the high rates in urban centers in the northeast to the lower rates found in some parts of our country. In an attempt to determine the energy cost of the various options in this study, we will use an average urban rate, based upon *Electric Power Monthly Report*, as stated by the U.S. Energy Information Administration in their report released on November 15, 2010⁴⁰. The rate we have chosen is \$0.15 per KWH (somewhat above the national average, to reflect a more familiar rate to those folks in California and the Northeast), and assumed an annual increase of 5%.⁴¹ As you will see, while there is a big difference in energy consumption between disk and tape solutions, it still is a small economic factor in the TCO of each. The issue has more to do with the economics of not needing to acquire another data center, or moving to a larger one, issues that we do not consider, but may be important to you.

So, if your data center is located on Wall Street in NYC, be prepared to pay more, a lot more for power, *if you can get it*. If your data center is in Idaho, you will pay a lot less. Don't even ask, if your data center is in Hawaii. You will have to factor in your local rates as appropriate for your region. We have also assumed the industry norm of a dollar to cool the environment for every dollar spent to power the data center's IT infrastructure.

The Data to Be Preserved

For the purposes of this study, we have decided to look at the enterprise requirements for data

³⁷ Since our goal is to keep the costs down, we presume that the least-costly disk storage already is being used at the primary site, so that the opportunity to use even lower-cost disk at the remote site does not exist.

³⁸ In reality, there are many fail-safe and fail-soft alternatives that might be deployed and the number of data centers involved may be more than two (as required by law for certain industries or because of industry-wide best practices). So, we will leave to you the arithmetic calculations for the scenario that you want to (or must) deploy.

³⁹ Some data centers do not pay for their power from their budget but it rolls up into another budget line. Regardless, it is a growing and major concern to most enterprises.

⁴⁰ See http://www.eia.doe.gov/electricity/epm/table5_6_a.html from the U.S. Energy Information Administration report released November 15, 2010, for data measured in August 2010.

⁴¹ None of the cost of energy assumptions is critical to the outcome, since the calculations were done in KWHs. Thus, the ratio will be the same regardless of the cost per KWH.

storage with an Archiving Appliance containing a sizeable, integrated operating storage serving as an archive cache, and a separate, high-capacity archive storage located either within the data center or at some remote location. We assume that this external archive currently consists of 1PB of unique files and binary objects, which need to be retained forever; 12 years for the purpose of this study. We also needed to pick a plausible rate of data growth (45%), as explained earlier. Examples of this data usually include, but are not limited to:

- Medical imaging
- Video production and security recordings
- Astrophysical recordings, such as those from radio telescopes
- Entertainment files and production data, including 3D digital movies
- Scanned libraries of books
- Gas and Oil exploration soundings
- Genetic modeling, databases and results

These tend to be uncompressible and unique *blobs*, data types that do not lend themselves to data deduplication.

- **Because of the uniqueness of these kinds of data in our hypothetical archive, data deduplication is not a factor unless the enterprise tries to store the same file additional times.**⁴²

The data to be archived is assumed to range in size from 1MB to 500 GBs⁴³ and needs to be indexed, archived, and retained for many decades, if not forever. Recently created, or recently accessed, files need to be retained within the Archiving Appliance's Operating Storage for prompt retrieval, and need to be available on a 24x7x365 basis without human intervention or delay. The Archival Storage may be a Tier-2 disk array or magnetic tape in an automated tape library (ATL). This data center decision will be based on multiple factors, primarily the retrieval speed at which the application demands for retrieval and the relative TCO of the solution. As the size of the archive continues to grow, the relative acquisition costs and

TCO factors become more significant, as well.

Procurement Assumptions

Procurements will be made at the beginning of each three-year cycle, with the capacity needed for the end of the cycle being acquired at the beginning of the cycle. Yes, it will cost more to acquire all of the capacity up-front (at the beginning of the first cycle); however, the incremental costs are minor relative to the overall TCO, and it eliminates the cost and/or disruptions related to multiple installations during operational periods. In all procurement decisions, the focus is on maintaining the lowest TCO, not on keeping up with the latest technology. We have taken into account that older technologies cost more to maintain because that cost usually is a fixed percentage of the acquisition price, not the price of the newest generation.⁴⁴

It is assumed that each new generation of disk and tape will hold twice the capacity of the previous generation; however, we have assumed that disks and tape drives will continue to operate within the same energy envelope as the previous generation, as has been the case for the last several generations. The Archiving Appliance will see all disk storage as available LUNs (i.e., in block mode) and will treat tape cartridges as serial storage. All storage will be maintained by its supplier on a 24x7x365 basis, with a four-hour response time⁴⁵.

As noted earlier, whether the tape or disk solution is delivered in a cloud infrastructure or a traditional one, the relative costs and issues of tape versus disk remain the same. Our analysis focuses on the medium and not the delivery mechanism of that medium. This study is not about the ownership or delivery mechanism. Whether the enterprise houses archives in its own data centers or chooses a third party to do so, the service provider has the same disk and tape alternatives from which to choose.

Data Management Assumptions

All files are assumed to be managed centrally by the archiving software or Archiving Appliance. The cost of this software or appliance, and the maintenance costs in support of them, will be the same for the same volume of data, regardless of whether the archive data resides on disk or tape.

⁴² We assume that if deduplication is to be done, it would be done in the Archiving Appliance or in a separate special appliance that also sits in Room 2. What is passed into Room 3 would be a single instance, regardless of whether it is stored on disk or tape. Thus, our presented average costs of storing the single instances in Room 3 would not be affected by the choice of media (disk or tape), whether because the stored data is unique or because it has been deduplicated before reaching Room 3. Do consider, however, that if deduplication is done on the preserved data, the deduplication technology (whether hardware or software) must be maintained forever so that the retrieved data can be put back together again.

⁴³ However, object size is not a factor in our calculations.

⁴⁴ We have assumed or adjusted all acquisition costs to include maintenance for three-years (the entire period) of 24 by 7 service with a 4-hour response time.

⁴⁵ Maybe we should have specified next day service instead of 4-hour service. We chose the latter because the data is important and if a disk array or tape frame were down, data would be inaccessible from the primary site. Additionally, if the vendor does not have a four-hour service plan, we have selected whatever offering is closest and shorter.

Therefore, it will be excluded from the financial equation, as will all networking and personnel costs, even though it can be assumed that disk will require more networking than tape. We have also assumed that the administrative costs of both solutions will be about the same, as the Archiving Appliance will manage all storage automatically. If this assumption seems unreasonable in your data center operations, you may want to adjust your costs accordingly.

Additional Disk Assumptions

We have seen a steady growth in both the capacity and throughput of disk media over the past ten years and feel it reasonable to assume a doubling of capacity every three years for the foreseeable future. We assume that the next-generation drives after 2TB will be available at no higher price than the prior generation. This is not to say that a disk array has a useful life of three years, simply that the cost of acquiring new arrays usually will be less than the cost of maintaining older technologies. All disk hardware for the three year cycle will be acquired at the beginning of the cycle to eliminate any issue of upgrade during the period. No residual value is presumed for the older arrays, as any value will be spent in cleansing the data for formal disposal, but they also can be used by another application, if you don't mind running without a maintenance contract.

Since it is not a question of *if*, but *when* a disk will fail, we have chosen to implement a RAID architecture to protect the data on disk. Currently, with 1TB disk devices deployed, data centers typically choose at least a RAID-5 technology. With a minimum of 2TBs per drive, and the length of time required to rebuild that large a drive, along with the risk of a second drive failure, we have opted to deploy our disk arrays with a RAID-6 architecture, increasing the overhead in using a disk-to-disk (D2D) architecture, but lowering the risk level significantly. We have also configured one spare drive per 14 to 16 drives, and this varies by product.

Additional Tape Assumptions

We will initially populate our ATLs with 1.5TB LTO-5 tape drives. The original ATL will remain in place for the duration of the study; however, tape drives will be upgraded with each new generation, which we presume will be at the beginning of each of the cycles. Some cartridge slots in the ATL may need to be sacrificed to make room for the deployment of the tape drives. We realize that not all tape solutions house their drives in the same manner; therefore, we have paid close

attention to this detail and made the appropriate adjustments, including attention to the write speed of each drive generation. Expansion frames are deployed as required, with the densest frames or drive-less frames being procured to reduce the number of frames and floor space required. All slots are activated upon deployment rather than on an incremental (capacity-on-demand) basis. Even though this accelerates spending for tape, it greatly simplifies our model and doesn't affect the TCO within a cycle.

With an established technological roadmap, we have made the assumption, based upon the LTO technology track record, that we will see LTO-6 in three years, with new versions to follow in three-year intervals. We will populate the ATL with new drives and cartridges as the technology advances in order to maintain the lowest TCO, which is achieved by writing as densely as the available technology will permit.

All drives and cartridges needed for the next three-year cycle will be procured in the first year that they are presumed to be available (i.e., at the beginning of a cycle), and it is assumed that all drives and cartridges will be priced at the same level as the previous generation. The LTO Program has established a policy whereby each new generation of tape drive will be able to read/write the previous generation and read the generation before that. With that in mind, we will copy and consolidate LTO-5 cartridges onto LTO-8 media, when LTO-8 is expected to become available in year 10, thus keeping the slot and frame count to a manageable level, as well as reducing the number of older drives that must be retained. Other than that, all cartridges will remain in their original format and on their original cartridge during the study period.⁴⁶ All cartridges will remain within the ATL to eliminate human intervention (i.e., none will be held outside the library).

The data center will need sufficient tape drives to write new archive data to tape and to copy older data from prior generations. We have assumed that drives will be dedicated to writing no more than ten hours per day, therefore available for retrieval of archive data and ongoing archive maintenance at least 14 hours per day. We have assumed an $n+1$ acquisition plan for all new LTO tape drives; an additional drive will be acquired at the beginning of each cycle to serve as a spare drive and to be available for unanticipated spikes in demand for retrieval and writing. We have retained as many of each generation of older drives for as long as this

⁴⁶ Unless the media has to be replaced due to defects.

makes sense and as long as there is room in the library without extra cost, primarily to be used to retrieve older data. However, these drives will remain without a service contract as newer LTO drives can read back two generations, and because all reading and writing requirements will be satisfied with only the latest generation of drives. Thus, these older drives truly are extra spares for unanticipated uses.

Economic Assumptions

We have chosen to ignore the time-value of money in this study.⁴⁷ Here's why.

1. Data growth over the 12-year period is significant. (As shown earlier in Exhibit 4.) Most of the data to be archived arrives in the second half of the study period.
2. If we were buying most of our equipment and incurring most of our operating expenses in year 1 (as one might when buying a railcar), then it might have made sense to consider the time-value of money. However, by breaking the procurement into four "natural" cycles, we compare costs within each cycle and the time value is not very significant.
3. This is especially true when interest rates are very low, as they are now. Unfortunately, we cannot predict when this will change and to what degree.

Conclusion

In the course of this study, we have reviewed the required configurations from a variety of Tier-1 vendors for both disk arrays and tape libraries and taken the average of the systems studied for the comparison. As a result of this comparison, we show that tape is the better value in terms of total cost of ownership for the long-term preservation of irreducible images and binary data, if you can wait for several seconds to several minute for retrieval. In summary:

- **In terms of TCO under our scenario, the cost to implement a disk solution for long-term archiving is over \$67M, about 15 times the cost to deploy a tape solution of \$4.5M.**
- **A very large expense in deploying a tape archiving solution might be, for some, an operational expense – the cost of media. At \$1.9M over 12 years, the cost of tape cartridges represents about 42% of the TCO for tape.**

⁴⁷ In fact, we did consider it, but it did not affect any of the conclusions drawn, so we removed it to make our model more understandable.

- **As we saw in 2008, the cost of energy continues to be a significant factor in the operational costs of maintaining a disk subsystem for long-term storage. In fact, the cost to run and cool the data center for a long-term Archival Storage on disk is more than \$4M. This represents a ratio of 238:1 in comparison to the energy required for a tape archive, about \$18K, even when we calculate that figure on a worst-case basis with all drives operating 24 hours per day, an extremely unlikely scenario. This is consistent with the ratio of 290:1 from our 2008 backup study.**

Ask Yourself

For decades, because of its cost, most enterprise data centers have used tape as the low-cost media of choice. Now, with any number of D2D solutions dotting the landscape, the data center staff is faced with a potentially budget-busting question:

- ***Does your data center need to spend tens of millions of dollars for sub-second response time for archive retrieval requests or is a response time of up to several minutes sufficient to satisfy the service level agreements that exist with your user community, at 1/15th of the cost?***

Restated in terms of this TCO study:

- ***Is instantaneous response worth over \$60M over 12 years?***

If all of your archive retrieval processes demand one-second response time (or less) for millions of queries during which expensive people are just sitting and waiting to get the data (or worse, as when people are at risk), then the answer probably is "yes, spend the money."

- ***If this doesn't make sense, even just some of the time, then an ATL may be the best investment you will ever make.***

Thanks for investing your time in this study and thanks to the vendors that provided the data. We have attempted to provide you with a fair and balanced scenario and analyses. Now, you need to consider what is best for your data center and your enterprise.

- **We think that tape makes economic sense for long-term archiving, especially if the Archiving Appliance's Operating Storage can satisfy most of your immediate requests.**



Appendix A — Explanations of Model Variables, Assumptions and Bias

#1 – USABLE CAPACITY – *Model Bias Favors Disk...Potentially Significantly*

- **Assumption** – Both tape and disk are filled to 85% of their usable capacity.
- **Explanation** – For disk, this is a generous assumption, since most disk solutions tend to become “filled” at a much lower percentage. For tape being written in serial order (as it comes in, until there is no more room), 85% tends to be an austere limitation. If disk is 10% too high and tape is 10% too low, then there is a significant (20%) bias that favors disk.

#2 – COSTS OF TRANSITION – *Model Bias Favors Disk...Somewhat*

- **Assumption** – Disk and tape both have negligible costs for transition between generations.
- **Explanation** – Because the Archiving Appliance manages data migration from one generation of solution to the next, neither will be as traumatic as migrating a disk array, LUN by LUN, and dealing with the associated addressing issues. However, we assume that each generation of disk is replaced by the next every three years, while older tape cartridges are consolidated and rewritten every nine years. Physically managing the migration of even thousands of tape cartridges is a much less demanding physical chore than installing new disk solutions and, after the data has been migrated and cleansed, removing the old disk solution. This requires staff time, extra space, extra energy, etc., none of which we have included in the disk costs.

#3 – ENERGY – *Model Bias Favors Disk...Somewhat*

- **Assumption** – Disk and tape consume energy at their maximum operating level.
- **Explanation** – For disks, this makes sense, since the media is spinning all of the time even when not being accessed. None of the vendors' disk solutions use any form of spin-down (quiescing). However, for tape, this is an extreme assumption, as it implies that all tape drives are spinning, reading, or writing, all of the time. We have calculated the time needed to write the volumes of new data to tape, but assume that this must be done in a 10-hour window and without using the spare drive, leaving the other 14 hours and the spare drive for retrieval. This assumption means that the energy required for tape may be much too high, but the tape solutions use so little energy that the bias against tape is not significant.

#4 – COMPRESSION – *Model Bias Favors Disk...Potentially Significantly*

- **Assumption** – Data is uncompressible.
- **Explanation** – We know that even with blobs, some lossless compression is possible. Anecdotal evidence suggests that digital data might be compressible to 2-to-2.5 times. A compression capability is built into LTO-5 tape (at no extra charge) but not into any of the disk solutions. Thus, the tape capacity required by the study's model might be 250% higher than necessary, which provides a significant bias toward disk.

#5 – MAINTENANCE – *Model Bias Favors Disk...Somewhat*

- **Assumption** – All hardware is maintained 24 by 7 with a 4-hour response time.
- **Explanation** – By assuming that the disks are RAID-6 configured with a spare for approximately every 16 drives, we feel that the data is protected (in its archived location) and operations can continue uninterrupted while rebuilding takes place and the bad drive(s) are replaced. The same is true for tape, as there is always a spare tape drive of the latest generation. However, the spare tape drive actually can be used operationally (for writing and for reading) while the spare disk drives cannot. Additionally, we have assumed that tape drives purchased in a prior generation will be kept installed (but not maintained) as long as there is room for them within the tape library. Thus, there is always at least one extra tape drive and potentially several extra tape drives, generally increasing the throughput potential, for use when needed. This assumption is a bias against tape, which favors disk.

#6 – FLOORSPACE – *Model Bias Favors Disk...Somewhat*

- **Assumption** – All hardware requires only the specified floor space, including space needed to provide access to all functional panels and components. We assume that floor space in the data center is valuable and available. We make decisions in each cycle that inherently minimize the floor space used.
- **Explanation** – We have considered all specified installation requirements. Racks full of spinning drives generate significant heat (unlike tape libraries and drives) and additional space may be needed to balance the placement of the racks of disks. For example, there may be an airflow plan for hot aisles and cold aisles, potentially increasing the square footage required for disks. Since we did not consider this, this assumption probably favors disk.

#7 – EXCESS CAPACITIES – *Model Bias Favors Disk...Somewhat*

- **Assumption** – Disk and tape solutions are bought in optimal capacities (i.e., the *Largest Building Block* or *LBB*, as described in the body of the report.) Additionally, tape cartridges are procured in an economic order quantity of 500.
- **Explanation** – At the end of each 3-year cycle, we presume that disk solutions are replaced and nothing is carried over to the next cycle. However, for tape we make similar assumptions for the tape drives, but we retain the old drives for use as spares. In addition, because tapes are bought in lots of 500, there almost always are virgin tapes left over, which we assume will not be used. In reality, both the older tape drives and the extra cartridges likely will be used, but we do not estimate this potential cost-lowering contribution to TCO that, in the end, increases the TCO for tape and thus favors disk.

#8 – THREE-YEAR LIFE CYCLE – *Model Bias Favors Disk...Somewhat to Potentially More*

- **Assumption** – Disk solutions are replaced completely at the end of each three-year cycle, while tape drives (but not the library) are replaced similarly on a three-year cycle. We assume that there is no salvage value for used components.
- **Explanation** – We assume that both disk and tape technologies are doubling in capacity every three years. There are two additional factors for disk solutions. First, many question whether disk can keep doubling without hitting some difficult technological barriers. If it can't, the TCO advantage for tape goes up significantly, because the densities promised for LTO -8 have already been demonstrated in the laboratory and we believe that they are operationally achievable in about 9 years. Second, we assume that disk solutions need to be replaced after three years because the cost of maintaining them is very high, with out-of-warranty maintenance contracts based on the original cost of the disk solution, which will be double the cost of buying a new one for the same capacity.⁴⁸ If disk vendors change their out-of-warranty maintenance costs, then it may make more sense to keep disks longer, which might lower the TCO for disk. However, we doubt that this will happen. However, if LTO-8 were able to read back three generations (instead of the two generations done currently), there could be a significant TCO reduction for tape solutions, as no cartridges would need to be replaced during the 12-year study (fewer drives but more racks would be needed). By ignoring both of these factors, there is a significant TCO bias in favor of disks. For tape, we assume that old drives are retained as spares for as long as they work and only when there is space available for them within the frames. Nine-year old cartridges are assumed to have no value.

However, if the cycle is less than three years for disk (or tape) and the arrays were replaced every, say, two years, then the total acquisition costs would be higher, since maybe six generations of disk (or tape drives) would be required. In reality, that two-year replacement cycle is an unlikely real-world scenario; it just would mean that arrays (and tape drives and cartridges) would be bought on an as-needed basis, with the additional densities of next-generation disks (or tapes) reducing the costs of solutions, somewhat, per TB, because fewer drives would be needed.

#9 – ACQUISITION COSTS – *Model Has No Bias*

- **Assumption** – We have endeavored to make the configurations and pricing equivalent for all vendors in the same class of product.
- **Explanation** – There are some small variations among the like offerings. We have adjusted for these variations as best as we could. We are confident that they are equivalent within plus or minus 1% (in dollar terms). Given that discounts should be expected for multi-million dollar procurements, the differences can be considered as insignificant to your procurement decisions. We do not think that these variations affect the averaged TCO significantly.

#10 – PEOPLE COSTS – *Model Has No Bias to Small Bias Favoring Disk*

- **Assumption** – No IT personnel costs are included in the model.
- **Explanation** – Since the Archiving Appliance sits in front of the Archival Storage (disk or tape), most of the “work” is done by the Archiving Appliance. There are generational upgrades for both disk and tape, which for disk are forklift upgrades (wholesale swapping of old arrays for new ones). Thus, there may be a little more IT staff work being done for disks, which results in a small bias to disk.

⁴⁸ For disk, the math is simple. It costs more to maintain a 3-year old disk system than to replace it, given the other assumptions that we have made including doubling of capacity in the next generation for no increase in price. If you have an array that costs \$100K to procure (an arbitrary round number for this explanation), including the first three-years of maintenance, and it costs 20% per year thereafter to maintain it, it will cost \$60K to maintain it for the next three years. While that sounds like a no-brainer decision in favor of keeping the old array for another three years, remember that you will get *twice* as much storage if you buy new. Thus, you need to spend \$120K to maintain two old array(s) compared to spending \$100K to buy a new one with doubled capacity. Furthermore, it costs something to cleanse and dispose of old arrays and that is presumed to offset any salvage value.

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 Authors

Mike Kahn is Managing Director and a cofounder of The Clipper Group. Mr. Kahn is a veteran of the computer industry, having spent more than four decades working on information technology, spending the last 18 years at Clipper. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis, and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation - Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; at Power Factor Corporation, a Boston-based electronics startup – President; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; at Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, at Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan, where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- *Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com or via phone at (781) 235-0085 Ext. 121. (Please dial “121” when you hear the automated attendant.)*

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, 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.