

A ROADMAP FOR COST CONTAINMENT IN DATA ARCHIVING

Kenneth Singer, Scott Ijaz and Steven Santamaria

EXECUTIVE SUMMARY

The burgeoning growth of digital data has triggered new assessments of data storage, for both the most active storage tier and the long-term archive tier. On the most active tiers, speed and performance are prime, while exploding demand for archive tiers makes low-cost most important. A common misperception is that ‘archiving’ implies low maintenance and passivity, like keeping a book on a shelf. In the case of digital archiving, that assumption misses the mark as data centers migrate previously archived data to new media every five years through a process known as remastering. This requires careful planning, accurate budgeting, and detailed execution in order to prevent data corruption and financial turmoil.

Data center managers constantly seek out cost containment opportunities. One expense, media acquisition cost, deserves closer scrutiny. From the early 1990s to 2008, media acquisition costs remained steady because media suppliers were able to continuously increase the areal density, resulting in sizable annual unit cost reductions in equilibrium with rapidly growing stored volume. However, hard disk drives (“HDD”), are nearing the end of their areal density roadmap, now with a meager 15% annual growth. Consequently, demand for storage is outpacing unit cost reduction. Organizations are spending more on media purchases and deleting valuable data.

In the wake of a changing landscape and accelerating data growth, data centers need to adopt new technologies whose features are more aligned to the evolving market environment: long media lifetime, backward compatible drives, and technology with robust capacity roadmaps. Such technologies will require less remastering and lower costs in the face of rapid growth and long-term retention. Optical discs best meet these criteria because they can last up to 100 years and optical drives have backward read compatibility. The industry, however, has been slow to adopt optical because of the historic low capacity implying high up-front cost per disc, even though the total cost of ownership is more favorable.

The newest generation of optical technology, Folio Photonics’ DataFilm Disc™ (“DFD™”), will enter into the market with the lowest cost and highest capacity per unit, with backward drive compatibility, a robust roadmap and long shelf life. Folio’s DFD™’s 100-year remastering cycles are far superior to competitors’ 5 years cycles, with significant savings. *Our model indicates that Folio’s DFD™ will upend the data archiving industry by lowering annual media acquisition costs over incumbent technology by up to a factor of 50 in the next decade.*

Introduction

The production of digital data is growing at an enormous rate with the total amount of stored data from 2018 to 2025 expected to reach 7.5 zettabytes, up from the recent 2.6 zettabytes. Further, 60% or about 4.5 zettabytes will constitute long-term archival storage.¹ The rapid growth of this segment is driven not only by increased production, but also by the growth of long-term storage requirements due to government regulations, business and legal requirements, and emerging technologies such as blockchain, machine-learning, and the internet of things. In addition, the rapid usefulness and growth in big data science drives the need to retain data for future analysis as machine and deep learning algorithms are more useful with larger datasets. Entertainment media, medicine, surveillance, finance, and government all require long retention times...from many decades to a century and beyond. Many of these applications require not only long-term retention, but also easy availability. Thus, active archiving is expected to become a greater fraction of data stored for the long term.

The present upfront cost per gigabyte (“cost/GB”) of the various archiving media, while important, is not adequate to describe the economics of archival storage, nor to properly compare the various archiving technologies. While existing comparisons of the total cost of ownership (“TCO”) of the various archiving technologies^{2,3} address the long-term costs, these analyses do not have the granularity to inform shorter-term budgeting. In this white paper, we present a model for annualizing the cost of acquisition for archiving media as a tool for comparing various technologies and informing approaches to cost containment.

Our analysis brings the impact of remastering to the forefront, that is, periodic migration of archived data to new media, as well as the capacity/cost roadmap. As we will illustrate, the ratio of the Data Retention Lifetime to the Remastering Interval (“DRL/RI”) is an important driver of annual acquisition volume for archival storage. When the impact of remastering on the acquisition volume is combined with the year over year media cost reduction rate (“CRR”) due to increased storage density, the differences between the various technologies come into sharp focus. Further, our annualized model serves to present these costs in a way that corresponds to planning cycles of IT professionals, namely the annual present acquisition of archival media and associated costs. We conclude that media acquisition is a dominant differentiating factor among total costs of ownership for the various

¹ Horison Information Strategies (horison.com)_

² https://web.stanford.edu/group/lockss/resources/2012-09_The_Economics_of_Long-Term_Digital_Storage.pdf

³ Gupta, et al. DOI: [10.1109/MASCOTS.2014.39](https://doi.org/10.1109/MASCOTS.2014.39)

media technologies. Additionally, long remastering intervals and backward compatibility of drives provides a strong advantage for optical archiving media. Further, combining this advantage with a low-cost, high-capacity optical medium with a strong capacity and cost reduction roadmap such as Folio Photonics' DataFilm Disc™ (DFD™), optical archiving technology is compelling.

Archival Media

The hardware for long term archival storage is currently dominated by magnetic media, hard disk drives ("HDD") and, to a lesser extent, magnetic tape. The economic viability of data storage on magnetic media in light of the explosive growth of storage needs has been made possible because of Kryder's law,⁴ the data storage analog of Moore's law for semiconductor devices. Here, the storage capacity per unit (areal density) has grown exponentially by a combination of both scientific and engineering advances while the production cost per media unit (single tape cartridge or hard disk) has been flat or declining. As a consequence, the media cost/GB has dropped exponentially. From the early 1990s to 2008, data storage costs were manageable because the two key drivers – cost per GB and data growth – maintained equilibrium. However, the slowing in Kryder's law for incumbent magnetic media demonstrates that future annual cost will grow faster. This fact, coupled with the explosion of archival data, indicates that the past equilibrium will end, resulting in unsustainable archiving costs and the unhappy prospect of discarding valuable data.

Table I summarizes some important properties comparing HDD, tape, and optical (Sony-Panasonic's Archival Disc™) storage. All three are secure for errors and have reasonably similar transfer rates. HDD and optical have the access advantages of random-access media. Tape and optical have much lower energy consumption costs since HDD is often kept spinning. Optical also has an advantage in environments lacking temperature and humidity control resulting in further cost and energy savings. The main disadvantage of optical is the high up-front media cost of discs arising from

⁴ Walter, Chip (August 2005). "[Kryder's Law](#)". *Scientific American*.

Table I: Properties of incumbent archival media

Feature	Tape	HDD	Archival Disc™
Native capacity (TB)	2.5-12	4-10	0.3
Transfer rate (MB/s)	<400	<210	<360*
Time to first bit (s) (maximum)	140	millisecond	40
Error rate	<10 ⁻²⁰	<10 ⁻¹⁷	<10 ⁻⁴¹
Refresh interval** (yr)	5-10	5	<100
Power consumption (wh/GB)	0.2	5.3	0.5

*System level parallel

**Typical; tape drives backward compatible for 5 years

Some data from Horison Information Strategies and Panasonic Freeze-Ray

historic low capacity. Folio's DFD™ brings the capacity and costs in line with tape, while much lower cost than HDD. This, along with the random access, robustness, and durability of write-once-read-many media ("WORM"), make a strong case for Folio Photonics' high capacity medium.

Hard disk drives

HDD are currently the principal storage media for archiving since ready access to data from its short time-to-first-bit and random-access capability are strong positive factors. However, its short media lifetime and limited roadmap requires a re-thinking of the economics in light of the explosion of archival data and requirements for long retention. Additionally, the favorable trade-off between decreasing cost per GB and increasing capacity is being challenged by the recent slowing of Kryder's law for HDD technology. Figure 1 below depicts the current HDD archiving pain point, where capacity growth far exceeds annual cost reduction.

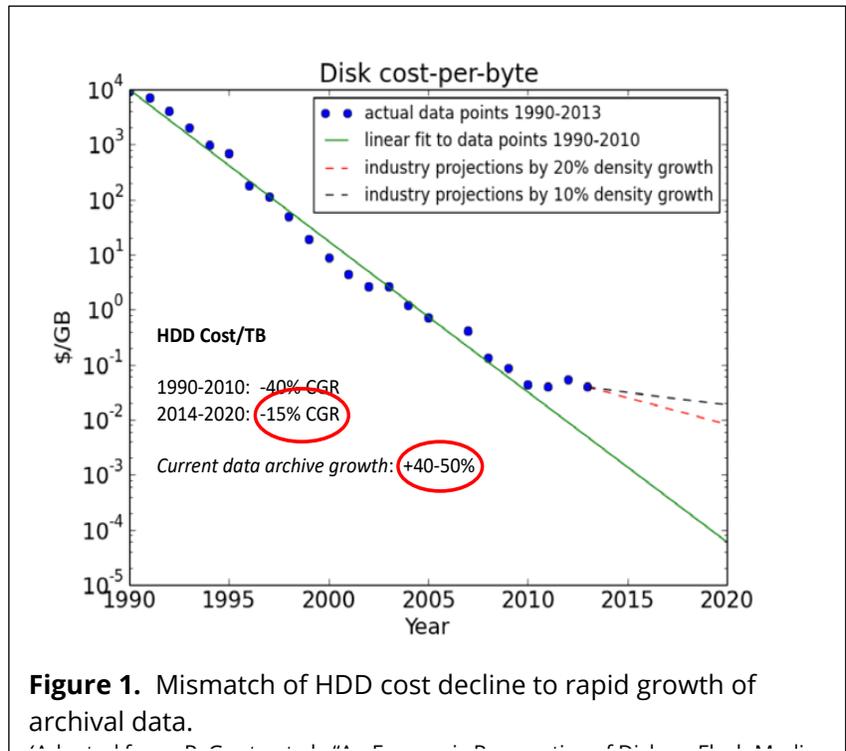
It has been pointed out that Kryder's law reflects a series of ever-flattening S-curves associated with a series of technological and scientific advances.⁵ As HDD approaches the fundamental areal density limit associated with magnetic domains,⁶ the growth in areal density has slowed from over 100% per year around 2000 to the current 15%. Current approaches to extend the growth of areal density toward the limit are mostly focused on overcoming the energy barrier to perpendicular writing as the bit size decreases. Additional sources of energy such as laser heating ("HAMR") and additional

⁵ https://www.theregister.co.uk/2014/11/10/kryders_law_of_ever_cheaper_storage_disproven/

⁶ Richter, H.J., *J. Magnetism and Magnetic Materials* **321** (2009) 467-476

magnetic fields (“MAMR”) are emerging to extend the roadmap. In the future, pre-patterned surfaces for locating bits (“PBM”) might be introduced, but at higher cost arising from the challenges of implementing the exquisite timing needed. Nonetheless, Kryder’s areal growth rate will continue to decline for HDD as the physical limit is approached in the coming years. Recently, the continued increase in capacity and decrease in cost per GB has mostly been driven by including multiple disks in a single package.

The reliability of HDD technology depends both on the retention lifetime of the media surfaces and on the reliability of the drive mechanics and electronics since HDD integrate the media with the drive. As a result, it has been a common practice to remaster the archived data every 3-5 years. This short migration interval has also been driven by the rapid increase in capacity and decrease in cost driven by Kryder’s law as new

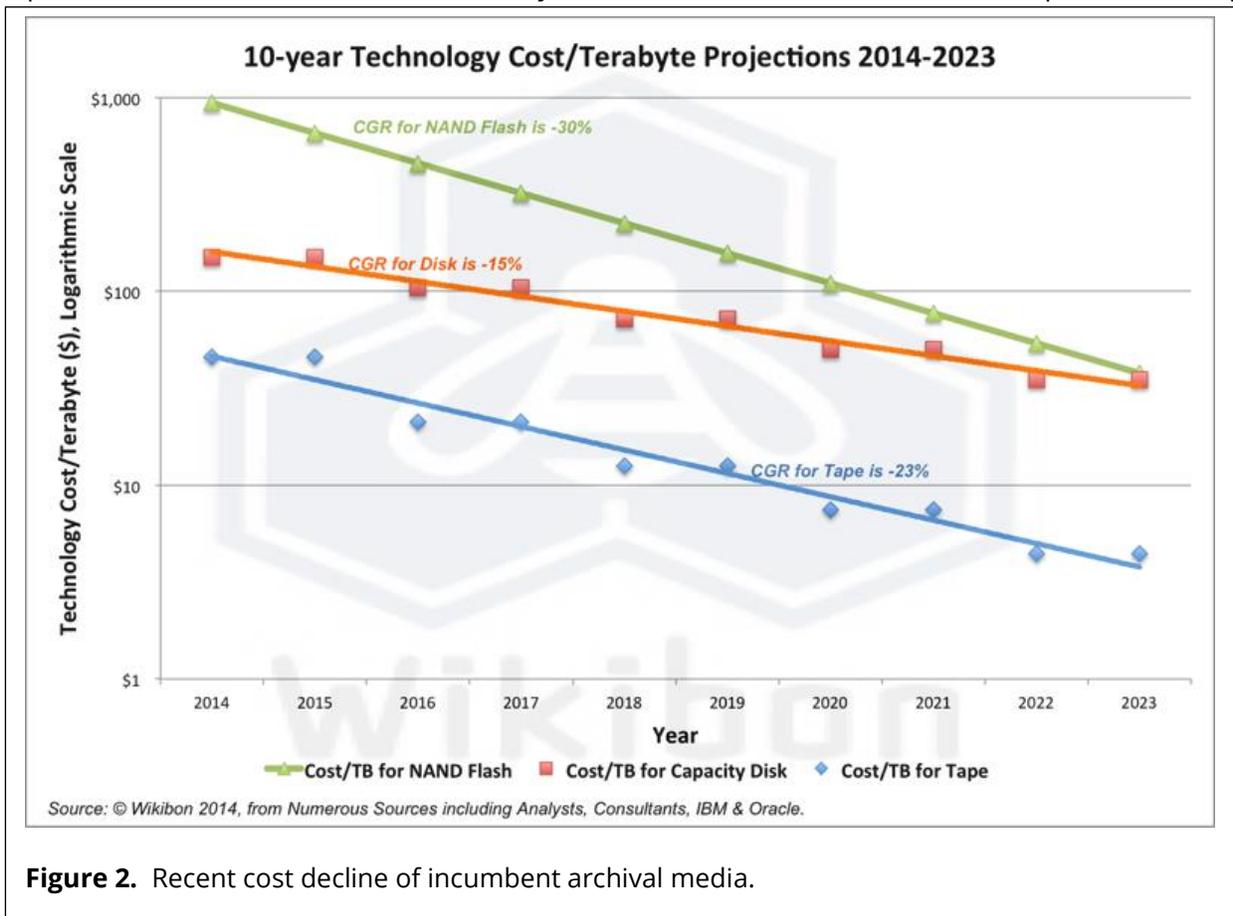


generations of HDD emerge. While it might be possible to extend HDD remastering intervals, the degradation of HDD due to the number of accesses may make this a difficult calculation. As data retention requirements increase, these short remastering intervals for HDD are a significant factor in the costs of this technology for archiving data. Decades of retention will require a substantial number of remastering events.

Magnetic Tape

Tape is a growing, though small segment of the storage of archival data, mainly due to the low cost per GB of the media (see Figure 2), currently several times lower than HDD. Tape is also several years and generations behind HDD in its areal density (see Figure 1) and thus has a longer lifetime on Kryder's curve by taking advantage of advances in areal density previously made by HDD. In addition, the cartridge capacity is large due to the large surface area over its long length...now near a kilometer. However, the long tape format implies a long-time-to-first-bit as well as challenges in data efficiency schemes such as de-duplication due to the need to locate complete files close together. This lack of random-access is the principal drawback of tape and is a significant issue for active archiving.

While the shelf life of tape is long, up to several decades under ideal conditions, its useful life is limited more by the number of accesses, storage environment, and most importantly, by the lack of backward compatibility of read/write drives. For LTO tape, backward compatibility of 1-2 generations is expected. Given the history of version introductions every two or three years since the introduction of LTO-1 in 2000, this implies an effective media lifetime of approximately 5-7 years as tapes need to be rewritten to be readable by the new drives. Unless old drives are kept on hand, only



a 5-year interval is sustainable for the long term. For enterprises where data growth is sufficiently low and companies are willing and to keep old drives on hand, tapes are sometimes kept for 10 years.

Solid state drives

Solid state drives (SSD) have been making major inroads to the HDD market for personal computing and high-performance enterprise storage tiers. SSD's lightning speed, lack of mechanical components, and compactness are highly desirable in these applications. Still, the cost per GB is considerably higher than HDD, and the storage density roadmap will not continue to expand quickly. SSD is not widely considered for archiving, but its somewhat longer useful life could make its TCO competitive with hard drive. However, the useful life is highly dependent on the number of accesses. When seldom accessed, the useful life can be as long as 10 years, but the uncertainty in the number of accesses usually results in considerable shorter remaster planning cycles. It is common to remaster SSD every 5 years.

Optical discs

Libraries of optical discs for archival data storage were recently introduced. These libraries are based on Blu-Ray® (BR®) and Sony-Panasonic's Archival Disc™. As random-access media, optical discs retain the advantages of data efficiency techniques. In disc library appliances, the time-to-first-bit is limited by the library robotics, averaging well below one minute. These libraries generally use system level paralleling schemes to increase the data rate up to and beyond that of HDD and tape, making it a good candidate for active archiving. As a WORM medium and with the potential for off-line storage, integrity is enhanced. The advertised lifetime of the Archival Disc™ is 100 years in ideal conditions and over 50 years in environments that are not environmentally controlled. As WORM media, the lifetime is less dependent on the number of accesses. They are robust against moisture and even water immersion, and are impervious to electromagnetic pulses, making them safer in the face of natural and man-made disasters. In addition, the low power consumption of optical libraries greatly reduces the TCO and environmental footprint. The long life and robust nature of optical discs implies that the discs will not require remastering even with data retention of 100 years. This, as will be demonstrated below, is a major factor in the economics of data archiving.

On the downside, BR® and Archival Disc™ platters have relatively low capacity, currently 200-300 GB, and, as such, the cost per GB is considerably higher than both tape and HDD. The current Archival Disc™ roadmap to 1 TB is likely not enough to assure economy in the future. The capacity is

limited by the areal density and the number of layers. Currently, work on the incumbent discs is mainly focused on increasing linear density since only 3-4 layers per side is economically manufacturable using the additive BR type manufacturing technology. The current Archival Disc™ is 300 GB on a two-sided, three-layer per side format, and BD-XL disc (enterprise BR technology) is currently 200 GB on 4 layers.

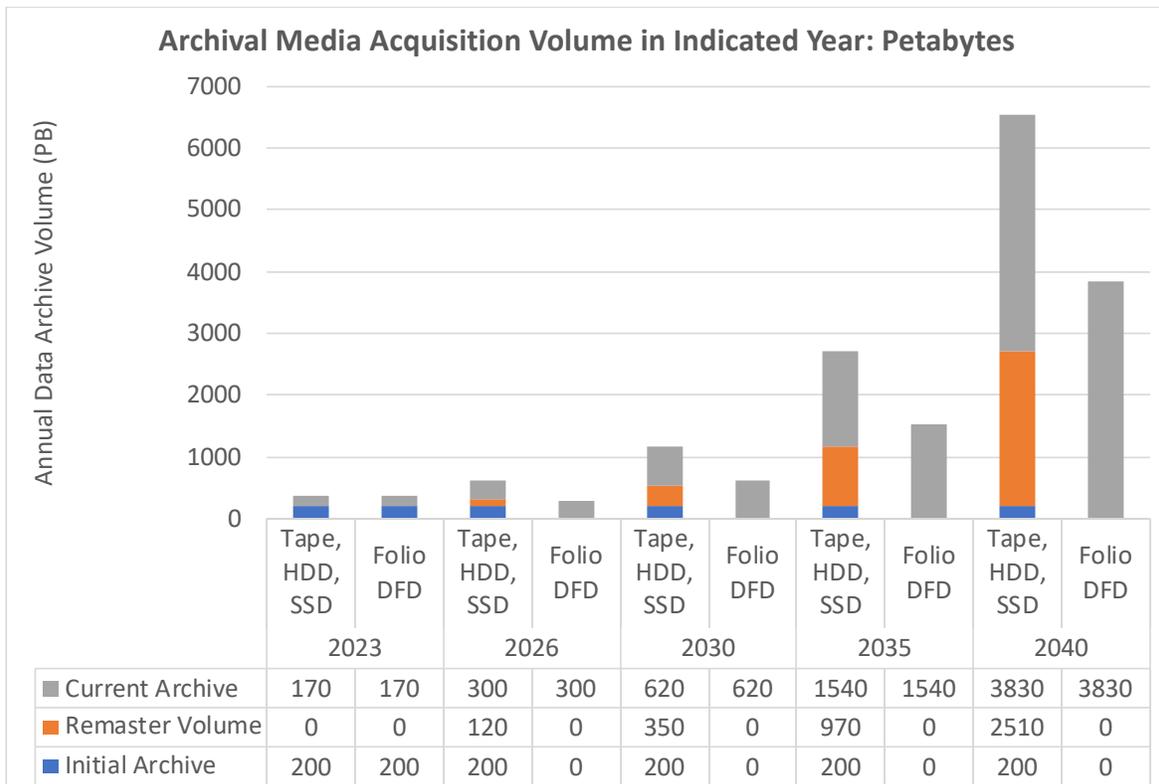
The ability to manufacture additional layers at low cost to bring the cost per GB of optical disc at or below that of tape would be disruptive. Folio's DFD™ is based on a widely-used, low-cost, roll-to-roll polymer co-extrusion method that produces all the layers at once, eliminating the yield problems associated with additive manufacturing processes. Currently, a 16-layer fluorescent disc is under development.⁷ This disc will be priced at or below that of the effective cost of tape per GB. The disc and drive are based on BR laser technology. The drive will be backward read compatible along its roadmap extending beyond 10 TB, and it is expected to maintain backward compatibility far into the future. Thus, the low cost per GB and 50-100 year remastering cycle suggests a disruptive data archiving technology. As will be shown below, the low TCO along with annual savings will be compelling.

Annualized Archival Media Acquisition Model

We present our model for the quantity and cost of archival data storage media on an annual basis starting in 2020 in Appendix A. The model is presented along with a table of the input and output parameters applicable to Figures 3 and 4. Model inputs include the retention period for the data, remastering interval, the present size and future growth rate of the archive volume, the cost reduction rate ("CRR"), and the size of the past archive that will require migrating in 2020 through 2025. We have adapted the data in Figure 2 to estimate the cost reduction rate of the incumbent media (CRR = -CGR). The CRR of HDD and SSD have been reduced somewhat to reflect expected decreases in the future as the growth in the number of layers of SSD slows and HDD approaches the physical areal density limit. Results of our annualized media acquisition model are shown in Figures 3 and 4 assuming a 20% year

⁷ www.foliophotonics.com

Figure 3. Amount of archival media in petabytes (100PB in 2020) to be acquired in the indicated year. Tape, HDD and SSD have the same volume due to their common remastering interval of 5 years. The remastering interval of Folio's DFD™ is 50-100 years. The growth rate of the annual archive is 20%.



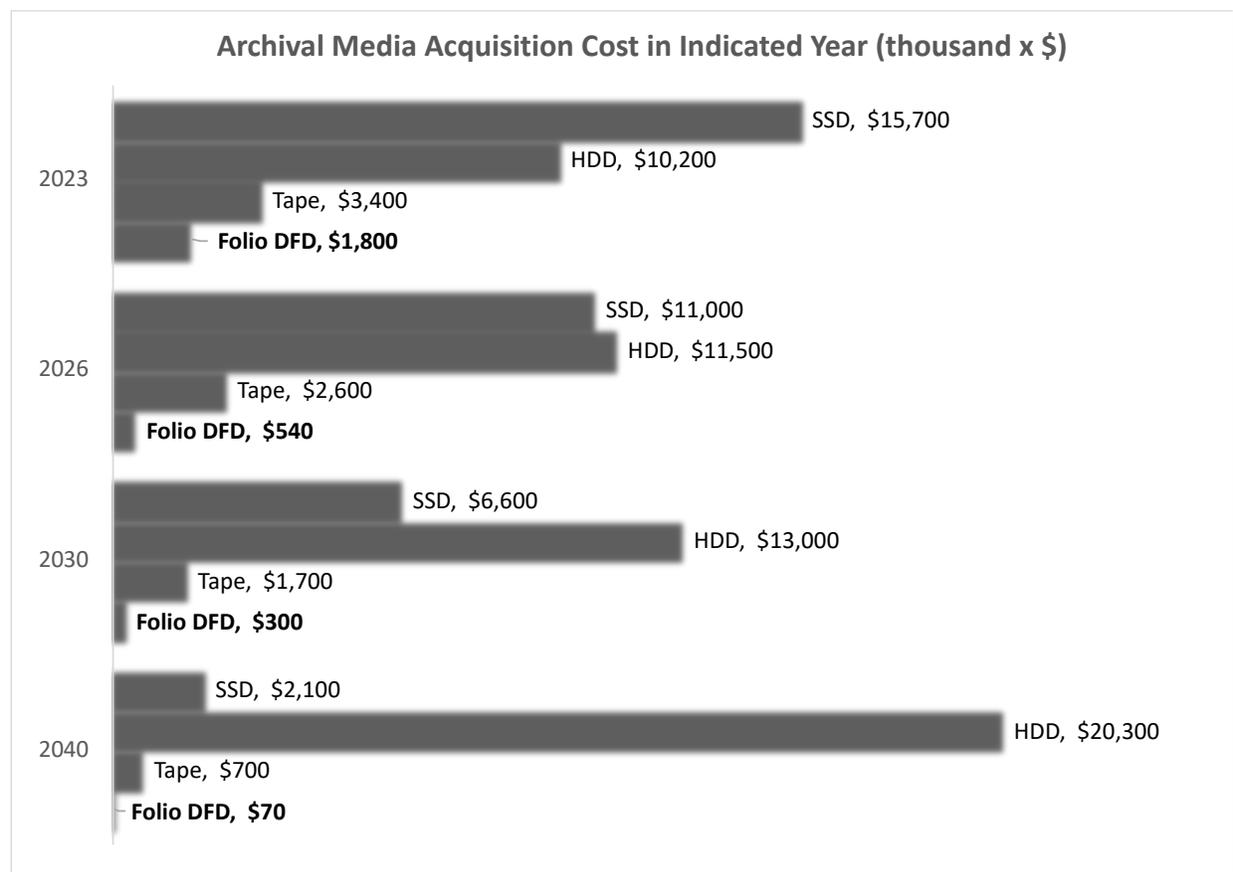
over year growth of new archival data. We note that the cost of tape includes refresh of the drives required every two generations of tape.

Figure 3 indicates that approximately 40-50% of the annual volume purchased will be devoted to remastering. That fraction, however, depends on the year-over-year growth of new archival data, and it is found that a 5% year-over-year new archival data growth requires 75% of annual volume devoted to remastering, while a 50% year-over-year growth rate requires that 15-20% of the acquisition is devoted to remastering. This, of course, is a significant fraction of annual purchases.

Figure 4 is based on a calculation that includes the cost/TB and its CRR. It is apparent from the figure that from a cost perspective, HDD and SSD are not competitive with Tape and Folio's DFD™. In a few years, HDD acquisition costs exceed a factor of 10 more than DFD™. Even for magnetic tape, an ever-widening gap with DFD™ is apparent due to both the remastering volume and differential CRR. Table B-1 in Appendix B presents additional results for various year-over-year archive growth rates. Here, CRR becomes a very strong differentiator among the storage media choices. Table B-2 in Appendix B presents a comparison of Tape and DFD™ where both are assumed to have CRR of 0.25. In this case, we extend the refresh interval to 7.5 years to reflect that some segments may retain tape that long despite the obsolete drives. In this case, the differences are attributable to remastering volume and the costs for tape drive refresh.

We note that, in addition to the cost of media (and associated drives) for remastering, it is estimated that an additional 17% of the cost of the remaster storage arrays is required for the

Figure 4. Cost of the archival media corresponding to the acquisition volumes of Figure 3 to be acquired in the indicated year. The remastering interval of Folio's DFD™ is 50-100 years. The growth rate of the annual archive is 20%.



remastering process.⁸ This amount includes personnel costs, software and other contributors. This factor would make the benefits of low-cost, long-lived optical storage even more compelling than depicted in Figure 4.

In contrast to previous publications that have addressed cost of archiving data via the TCO in terms of all contributing costs, we have focused solely on the acquisition of media on an annualized basis in order to reflect typical purchase decisions. We find that media acquisition is a major differentiator among the various storage technologies due to the accumulation of remastered data when the retention lifetime greatly exceeds the remastering interval (DRL/RI large). In addition, the values of CRR among the various technologies amplifies the volume differential leading to enormous cost differences. Energy consumption is often also quoted as a differentiating factor between optical and tape on one hand, and HDD on the other. This differential depends on the fraction of time that HDD are kept spinning as well as on the environmental conditions that the storage technology can tolerate. As WORM media, optical discs are more tolerant of elevated temperature and humidity, and since environmental control can be a large cost factor, it is possible that this is an additional advantage for optical media. Even without these additional factors, our cost analysis on media acquisition alone strongly favors Folio's DFD™ -based optical archive.

We conclude that data remastering, backward compatibility of media drives and especially CRR have a significant impact on the costs of archival media acquisition. Further, our results indicate that a high capacity, low-cost optical solution, such as Folio's DFD™ would be a game changer. The case is more compelling when coupled with the energy efficiency, ability to store under harsh environmental conditions, savings from associated costs of remastering, and imperviousness to water immersion and electromagnetic pulse. Our results also indicate that the current dominance of HDD in archiving is seriously mismatched to this application from an economic perspective, although HDD manufacturers are trying to enlarge this market as SSD takes over the higher performance tiers.

⁸ Source: Wikibon 2014

Notes:

Our model and its inputs for annualized ownership costs are given in Appendix A. The intention of this model here is to demonstrate the impact of remastering and the cost reduction rate (CRR= - CGR) on media acquisition economics.

1. The quoted CRR may be somewhat less than indicated in Figure 2 to account for potential reductions over the period. We expect that HDD CRR will decrease the most as HDD approaches the physical limit and SSD CRR due to maturity of the technology.
2. The remastering interval of 5 years is typical for Tape, HDD and SSD. Note that the lifetime of Tape is longer, however, the lack of backward compatibility of the drive beyond two generations stimulates a remastering interval of 5-10 years depending on their capacity to retain obsolete drives.
3. We have assumed that the data stored prior to 2020 is migrated in equal increments over the remastering interval of the incumbent media on which it was stored.
4. Volume and costs are based on an existing prior archive of 1000 PB and a 2020 new archive of 100 PB.
5. We add \$70 to the cost of single tape cartridges due to the cost of the drive, assuming an average drive cost of \$7000 and a tape to drive ratio of 100 in a tape library. This results in a cost of approximately \$20/TB for LTO-7 and LTO-8. (BDXL optical disc drives cost about \$200; therefore, we ignore drives costs for optical drives.)
6. Backward compatibility of optical drives (so far demonstrated in the field for almost 40 years) allows the remastering interval to correspond to the media lifetime, 100 years for the Archival Disc™.
7. We assume that DFD™ enters the market in 2020 with a 16-layer disc with capacity 0.75 TB and cost of \$10/disc.

Appendix A: *The Annualized Archival Media Acquisition Model*

Model parameters and variables:

Inputs:

u = year index (0 for the present year)

L_D = required lifetime of archived data (years)

L_M = years between migration (effective media lifetime) (years)

p = fractional increase of new archive data year to year

C_u = media cost per TB (\$/TB) in year u

$L_M(0)$ = Lifetime of archiving media prior to $u=0$

$V(0)$ = volume of archived data prior to $u=0$ (TB)

Outputs:

V_u = amount of new data needing archiving in year u (TB): $V_u = V_0(1 + p)^u$

F_u = media purchase in year u devoted to migration: $F_u = \sum_{n=1}^{INT(\frac{L_D}{L_M})} (1 + p)^{-nL_M}$

T_u = archival capacity to purchase in year u (TB)

P_u = total cost of archiving media in year u (includes new and migrating data) \$

V_{old} = Annual volume of media required for migrating fraction of $V(0)$

$$P_u = C_u T_u$$

Where, for u up to the lifetime of the data,

$$T_u = V_0(1 + p)^u \left(1 + \sum_{n=1}^{INT(\frac{u}{L_M})} (1 + p)^{-nL_M} \right) + \frac{V(0)}{L_M(0)} \text{rect} \left\{ \frac{u - INT(\frac{u}{L_M})L_M}{L_M(0)} - \frac{1}{2} \right\} = V_u(1 + F_u) + V_{old}$$

where INT is the greatest integer (step) function and

$$\text{rect}\{x\} = \begin{cases} 0 & \text{for } |x| > \frac{1}{2} \\ \frac{1}{2} & \text{for } |x| = \frac{1}{2} \\ 1 & \text{for } |x| < \frac{1}{2} \end{cases}$$

Model Inputs

Data Lifetime > 20 years

Year 2020 Archive = 100 PB

Prior-to-Year 2020 Archive = 1000 PB

	HDD	SSD	Tape	DFD™
Remaster Interval (yr)	5	5	5	>50
Archive Growth Fraction	0.2	0.2	0.2	0.2
Media CRR	0.12	0.25	0.23	0.28
Year 2020 Cost (\$/TB)	40	100	20*	13

*includes \$70 per cartridge for drive replacement