WHITE PAPER
October 2006

New Life for Dead Storage Technologies

Consumer-Driven Optical Storage in the Data Center

Jonathan W. Buckley

Executive Summary

By all accounts, optical technology has been a near failure as an enterprise storage technology over the last 15 years. While optical technology has become the default removable storage medium on the desktop (CD and DVD), in enterprise storage the evidence of failure is unmistakable. According to IDC[1], the total worldwide shipments of all optical automation equipment in 2005 amounted to a mere $69M USD. In contrast, the external spinning disk factory exit values amounted to an estimated $16B during the same period. Meanwhile, the data tape market amounted to an estimated $4.8B . After a decade and a half optical technologies in the enterprise storage market account for only a fraction of 1% of enterprise storage hardware spending.

For a traditional optical-based enterprise storage supplier, these are tough facts to acknowledge. Yet they are undeniable. There will be a dramatically different future, however, for certain optical storage subsystems in the next several years . The change will result from a convergence of new business requirements. These include new and different market demands for archival storage and new technical approaches for employing optical technologies which overcome previous limitations. This paper will cover the following:

- Assert that there are four main reasons for the historical failure of optical technologies to take a meaningful hold in the enterprise, including reliability, cost, performance, and proprietary vendor/technology lock in.
- Examine the changes in enterprise storage demand, which can be uniquely addressed by optical-based systems given the current technological outlook.
- Discuss how new storage architectures are overcoming the four main issues holding back optical in the enterprise.
- Address how consumer-driven technologies, such as Digital Versatile Disk (DVD) and its next-generation, blue laser counterparts, rather than "enterprise optical," are leading to a coming resurgence of optical technology in the enterprise. These consumer-driven technologies are being transformed in the same way Just a Bunch of Disks (JBOD) was transformed in the 1990's by Network Attached Storage (NAS) vendors with the employment of Redundant Array of Inexpensive Disks (RAID).

1 IDC Market Analysis. "Worldwide Optical Automation 2006-2010 Forecast and Analysis". Wolfgang Schlichting. April 2006, IDC#201141, Volume 1

Why Old Optical Failed

The problem for optical storage technologies in the enterprise over the last 15 years is that the design criteria and expectations applied to it were not suited to its nature as a subsystem component. The original promise of optical media was a lower storage cost compared to magnetic disks at only a slight discount in performance. Early efforts focused on using optical disks to lower the costs of the overall storage subsystem. To achieve this result, optical disks became an integral part of the operating system storage infrastructure for general-purpose primary storage. The approach in the storage software solution was to make optical devices behave as much as possible like the magnetic disk.

Optical technology, however, could not keep up in the market. In the early 1990's a magneto-optical (MO) disk was multiple orders of magnitude less expensive than spinning disk and yet enjoyed roughly similar seek times. Additionally, this was a time when an optical disk could hold two to three times more than each individual hard drive. By 1995, the hard disks had climbed to ~2 gigabytes (GB) in capacity from just tens of megabytes (MB) a few years prior while MO was stuck at 640MB per disk. DVD was introduced in 1997 with 4.7GB per disc, but the capacity and performance race was short lived, and was lost decisively.

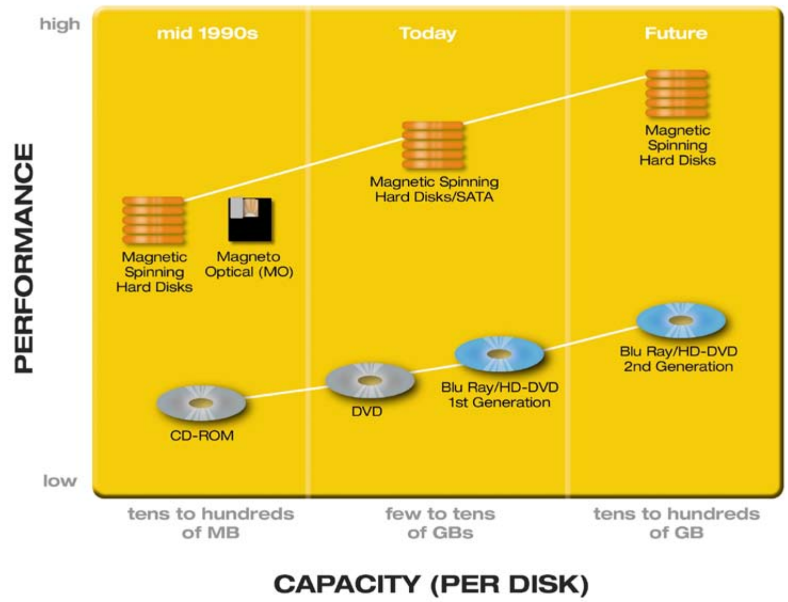

As we write this document, new optical technologies are starting to hit the market with impressive capacity gains. However, it must be noted that at the cutting edge, a 50GB per disc (Blu-ray disc) is going up against the newest hard drives at 500GB per 3.5 inch disk drive, and the seek times on magnetic disks are one-fourth to one-sixth that of optical. Today's low-end serial ATA (SATA) disk drives are about $3-4 per GB. This is a 1000x reduction over a 15-year period. Meanwhile, optical media costs in an automated changer have dropped as well to about $2-3 per GB, a 242x reduction. Space savings and speed alone for disk seem to call the end of optical storage technologies for all but a removable and portable storage medium. The fact is that optical cannot catch up with magnetic drives in terms of density and performance. But it doesn't need to in order to be relevant. There is a new calling for storage in the enterprise that demands a design criteria particularly well suited to optical technology.

The Emerging New Market Demand for the Online Archive

The increase in magnetic disk drive densities over the last decade has been nothing short of stunning. In most recent years, magnetic disks have "slowed" to a yearly average capacity gain of about 25-30% CAGR[1]. As incredible as that seems, given the now common use of multimedia, graphics and the like, some estimates have pegged the average enterprise data growth in the last few years at about 60% CAGR. With this difference alone it did not take long for the growth in number of network filers, not just disk replacement, to hit new records. We have finally outstripped our

1 http://www.almaden.ibm.com/StorageSystems/Advanced_Storage_Systems/Advanced_RAID/

ability to "pack data" with the ability to generate it. Good macro evidence of this fact is that despite a precipitous fall in the price per byte of disk, we continue to reach new highs in disk capacity spending. Recently compounding this natural growth effect of data overall is the emergence of legal data retention requirements, a phenomenon that is relatively new to all but a few market verticals.

A plethora of regulations, including the well publicized Sarbanes-Oxley and HIPAA regulations, have now pushed certain corporations to retain very long-term data. Nearly 30 new and different regulations mandate long-term data retention. Much of this general excess data retention load could be effectively handled by shuttling off aging, but required, data to tape, and sending it off site and out of the data center.

Capacity of Reference Information Worldwide, 2001-2006

Worldwide, Digitized Reference Information created & stored will surpass Non-Reference Information by the end of 2006

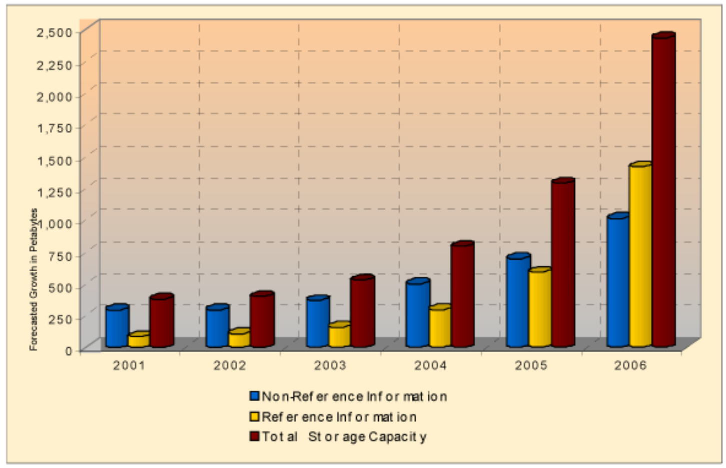

Source: Enterprise Strategy Group, 2002

Data Retention Regulations*

	SEC 17a	NASD 3010	HIPAA	21 CFR 11	Sarbanes Oxley
Public Sector			x		
Public Companies	x		x		x
Financial Services	x	x	x		x
Healthcare			x		x
Life Sciences			x	x	x
Other			x		x
Archive Duration	3 years to 6 years after account closing	3 years to 6 years after account closing	2 years to 21 years (in some cases, for the patient's lifetime)	2 years to 36 years	4 years to 7 years

Yet another new twist in the data retention saga will take effect on December 1, 2006, when a notable amendment to the Federal Rules of Civil Procedure, Rule 26 (b), takes effect. The regulation will likely force the timely access to electronically stored information that is found in "reasonably accessible" sources. This will create an inflection point in the data storage world, since corporations are not only being forced to retain data for very long periods of time but also to make it accessible if it is commercially viable. Naturally, tape sent offsite is not a particularly cost effective means of meeting this compliance requirement if it is needed.

Rule 26 (b)(2)(B) states that:

"A party need not provide discovery of electronically stored information from sources that the party identifies as not reasonably accessible because of undue burden or cost."

This rule will likely cause a bifurcation in the use of tape in the enterprise. Most companies use tape as a default archive medium in addition to its intended and designed purpose, which is to back up and restore critical data. In fact, it is estimated that 99.9% of the world's ~30 Exabytes (10^{18} bytes) of electronically archived data resides on tape today.   After December 2006, if companies wish to exclude data on tape in legal discovery processes, they will need to clearly define the intended purpose of tape in their environment and be consistent with that policy, in order to avoid highly public fines. Either tape will be a strict backup medium and thus not be "reasonably accessible because of undue burden or cost" in the discovery process, or it will remain as the default medium for true archive (not just backup) and the data on tape will be included in the process. If the company chooses to label tape as an “unreasonably burdensome or costly medium” and thus off limits to the discovery process, then it must either migrate its tape archives to an online medium or expunge the archive data and then adhere to a strict policy of not retaining long-term data completely. A company cannot claim that, on the one hand tape is cost effective for archiving its corporate data, and then, on the other hand that it is not cost effective enough to include that data in the discovery process. Retrieving files from tape is neither quick and inexpensive nor reliable over long periods of time. But the prospect of consistently deleting data doesn’t seem to be a winning idea in an information economy either. This one legal change alone can bolster the adoption of an “online archive,” whereby access to all data, not just current data, is facilitated without undue burden on the company.

Online Access to Archive Data Emerging as A Revenue Opportunity

If the “stick” in data retention is a deadly combination of retention and accessibility mandates, the “Long Tail” of data is certainly the “carrot” for businesses[2]. As summarized in a Wikipedia entry, the phrase The Long Tail, as a proper noun, was first coined by Chris Anderson. He argued that products that are in low demand or have low sales volume can collectively make up a market share that rivals or exceeds the relatively few current bestsellers and blockbusters, if the store or distribution channel is large enough. The Long Tail is a potential market and, as examples such as Amazon.com and Netflix illustrate, the distribution and sales channel opportunities created by the Internet often enable businesses to tap into that market successfully.

In the online content world, with the marginal “shelf space” closer to the cost of zero, we are seeing new and unexpected statistics in purchasing patterns. Where it was once thought that new versus older content would be demanded in a classic Pareto distribution model of the top 20% constitutes 80% of the demand, we are finding that the aggregate of the content that never would have made it on the retail store shelves, in the online world, outsells the top 20% consistently. Chris Anderson makes this point very well in his book.

The internal corporate use of legacy data for modeling, referencing, or building new strategies can be expected to follow the same “usage patterns” as Anderson’s Long Tail of consumer content. For example, the very research for this paper required the access to industry reports from the early 1990s. This is not an easy set of information to find when the primary storage “retail store shelves” of the corporation long ago expunged such information either by hitting the delete button or banishing it off to a remote tape vault, which is as good as deleted. With access to data in an archive, perhaps previously thought to be useless internal data, both “consumers” of data within a corporation and external consumers of such data as music files will show just how valuable such “written off” files can be.

The carrot and the stick incentives for online, readily accessible archived data is no doubt reason the Enterprise Storage Group (ESG) is calling for a 98% CAGR in archival storage—data storage that is online but relatively infrequently accessed—over the next decade. In essence, the “push effect” of regulation together with the “pull effect” of revenue on the long tail of data will forever change the market demands and dynamics for long-term data.

---

2 Beginning in a series of speeches in early 2004 and culminating with the publication of a Wired magazine article in October 2004, Anderson described the effects of the long tail on current and future business models. Anderson later extended it into the book The Long Tail: Why the Future of Business is Selling Less of More (2006).

Changing Economics in Storage Design Approaches

Clearly there is an emerging requirement for very long-term, online storage for content which is both fixed (unchanging) and also relatively infrequently accessed when compared to data in process.

At the highest level, tape and disk enterprise storage systems have been and are designed for a functional relationship between cost/performance (speed, scale, availability, reliability) and capacity (or density in design). For the purpose of making the point in this paper, let's represent this design criteria as Device$=ƒ(Performance, Capacity). This criterion is an optimization for speed in writes, re-writes and writes again in a scalable multi-client shared environment. Tape is designed specifically for bulk writes and bulk offload of data in a backup and restore scenario rather than for a collection of many small files written and retrieves randomly. Being magnetic in nature, both disk and tape media have never been intended for long-term storage. In fact, they are generally thought to have a lifespan of 3-5 years. More specifically, the tapes themselves (since this is a removable system) have a best practices requirement for re-writes every 5 years or so in order to avoid data decay even if they are kept strictly at 68° F and 40% RH as most tape manufacturers demand.

Online archive storage requires very different design optimization criteria. Remember that in theory, fixed content need not be written more than once, as it is fixed and not changing. Also, the nature of an archive is that it is unpredictably, and relatively infrequently, accessed on a file-by-file basis. Much as Amazon reduced the cost of "shelf space" compared to bookstores by locating warehouses in inexpensive areas and leveraging the Web as its shelf space, in order to offer a "long tail" of books, so must an online data storage system, in order to hold a "long tail" of data efficiently for the enterprise. The design criteria for an online archival system must be different then its shorter-term, high-performance relatives of the disk-based systems. An archival system must stress total cost of ownership over extended periods as a function of longevity and accessibility. By accessibility, this is the requirement that the archive actually be online at all times or it cannot meet the "long tail" requirement or enable data discovery cost-effectively. For the purposes of this paper, let's represent this design criteria as TCO $=ƒ(Longevity, Accessibility).

Maybe an Opening for Optical's Second Act?

If optical storage technology in the enterprise clearly failed to gain meaningful traction in the last decade and a half by the numbers, what could possibly start to change its course now, so late in the game? To answer that question it is important to recall the previous design criteria for all storage up to now. As was iterated above, Device $=ƒ(Performance, Capacity) was the guiding principle in optical storage system design. This has been the basis of design for all networked storage and was what the market expected even of optical alternatives in the early 1990s. It turns out that write-once optical is, by its very physical nature, very well suited to meet the emerging design criteria of TCO $=ƒ(Longevity, Accessibility), the emerging online archival storage design criteria.

Total costs over the archive period include such factors as power, maintenance, and the cost of data migration over time in order to preserve its integrity. For instance, disk has a practical life expectancy of 3-5 years and tape requires a migration to new media every five years or so. However, write-once optical media requires no such migration for many decades. By the demands of online archive alone, tape should not even be considered when addressing this enterprise need because it can't meet criteria #1: Accessibility. Yes, there are nearline tape systems on the market, but given their cost and design focus, they are all aimed at the backup and restore process. Given their size and capacity, the fact that files need to be sequentially searched rather than randomly searched eliminate these systems for practical consideration.

Offline tape comparison is included here simply because it is the default medium for data archive 99.9% of the time, by current analysts' accounts.

How Much Data (TB) 100

Retention Period (Yrs.) 20

Total Cost of Ownership

Type	Hardware	Software	Media	IT Costs	Floor Space	Power	Data Migrations	Total
DVD Optical	$152,610	$228,910	$47,060	$183,980	$15,160	$49,000	$0	$676,720
Disk CAS	$934,820	$2,504,460	$0	$644,530	$48,380	$735,800	$2,804,460	$7,672,450
Tape	$554,260	$203,120	$29,000	$225,930	$29,670	$128,800	$1,749,780	$2,920,560

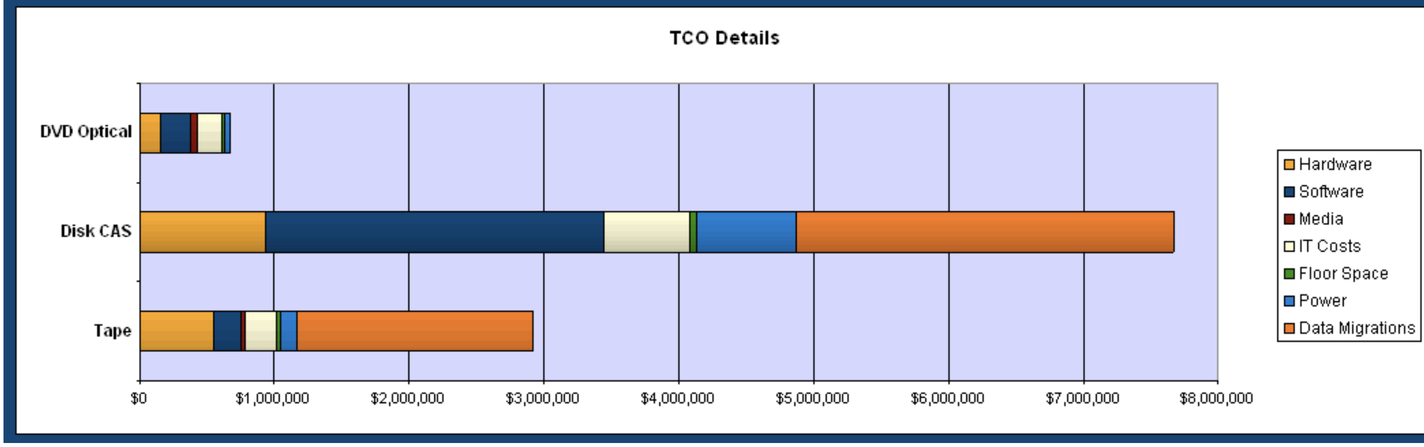

Designing Networked Storage for the "Long Tail" with Magnetic Disks Alone

Today's magnetic spinning disk, network attached storage systems are designed and built for extremely fast data writes and rewrites. The physical parts in the system, often spinning up to 10,000 revolutions per minute, have an expected system replacement timeline of roughly three years. Statistically, approximately 5% of enterprise class disks will fail each year across a large sample. However, it must be noted that since disk failure is highly correlated with environmental conditions and with manufacturing batches, disk failure within a single or co-located systems can be much higher than the 5% overage. In terms of the reliability of such systems, the disk error rate per drive and system has increased logically as disk density has increased. IBM Almaden Research[3] writes: "Disk drive capacities have been increasing at a rapid pace over the years—between 60 percent and 100 percent per year. (Since 2004, the growth rate has slowed down to around 25-35 percent per year). However, the probability of uncorrectable read errors has been relatively constant: around 1 in 10^{15} bits. As a result, with increasing capacities, the probability of a data loss due to an uncorrectable media error has been increasing to the point where it is now a significant factor."

The issue of uncorrectable read errors is not new and is mitigated by the employment of error correction approaches, in addition to RAID applications. In smaller-scale systems, RAID 5 can be effective against uncorrectable read errors which can lead to spindle failure; however, due to the correlation with disk batches and environmental effects on failure, RAID 6 is the most effective in protection from this type of issues. RAID 6 is not as popular as RAID 5 due to the write penalties to the system.

3 http://www.almaden.ibm.com/StorageSystems/Advanced_Storage_Systems/Advanced_RAID/

In order to keep data on spinning disk systems for long periods of time as an online archive, consider the math behind the storage system designed to keep the data intact. Consider a RAID 5 configuration with 500GB drives aimed at archiving 1.5 TB of data. Also known as striping with parity, RAID 5 reserves the space of one drive for parity allowing the system to experience a drive failure without the loss of data. Achieving 1.5TB of usable capacity with 500GB drives and RAID 5 would require four drives or 2TB of raw capacity. Throw in a hot spare and the system configuration will require 2.5TB of raw capacity. Add offsite replication/mirroring and the configuration will require a minimum of 10 drives or 5TB of raw capacity to achieve 1.5TB of usable capacity.

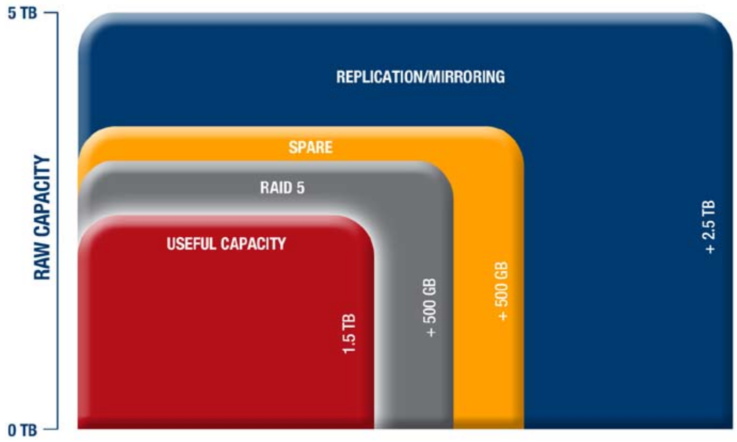

Every disk is spinning at 7,200 rpm or greater in the event the infrequently accessed archive data is needed, and each is drawing power, requiring cooling and needing replacement every three years. A disk subsystem designed to write, rewrite and deal with high traffic is not a good fit for holding data that is intended to be written only once and is infrequently accessed. This approach seems akin to keeping a Ferrari on idle in the driveway for the next three years in the event you need to run an errand. This may seem like a rational approach to highly demanded, mission-critical data, which requires very fast writes and rewrites with very high availability. But it does not seem like a rational approach for archival data.

To meet long tail requirements, Longevity and Accessibility are key design considerations for reducing the TCO of a storage system. In addition to keeping fixed data for ten, twenty or thirty years, "ready access" merely needs to be in seconds, not milliseconds.

With WORM-based optical in the subsystem, the overhead in raw capacity is greater reduced because the physical storage platter and the drive used for reading the media are separate, and thus the physical storage platter is not susceptible to the same uncorrectable read errors after the data is written and verified, to keep the data online and available. With optical-based parity techniques and error avoidance approaches employed, the raw storage to usable storage ratio is within 10% versus 100%-200%. Factor in an equivalent remote mirror copy for disaster recovery protection purposes, there is a greatly multiplied raw to usable capacity difference. This is just one reason that the comparison of low-end, enterprise SATA-based systems marketed for online archive at ~$3,000 per terabyte is an unfair comparison at face value to online archive systems with optical back end components at ~$4,000 per terabyte. The required terabyte count for the SATA based system must be much higher to make an "apple-to-apple" comparison.

Scale Challenges in Online Archive for Device $=ƒ(Performance,Capacity) Approach

The storage market started shifting in the first half of the 2000s in large part because of an inherent lack of physical and financial scalability in network-attached storage (NAS) systems for all the long tail reasons outlined above. There are certain physical scale issues an enterprise will reach, and they will eventually require a tiering of the storage network. Reliance on disk for primary storage and tape media for backup (i.e. "flat" storage architecture) is insufficient to keep pace with a growing volume of data. The first physical "wall" one hits is that of backup time. Generally, backup time for an incremental scheme using traditional methods is about four hours for each terabyte of data. This explains in great part the large rise in the number of startups since 2000 addressing backup times using approaches such as fast disk as the backup medium, software to compress the time and space for the process to take place, and continuous data protection, the newest entrant in technology. These new technologies definitely help with easing the backup time pressure. With the growing storage industry acceptance that ~60% of content on spinning disk is fixed content as established earlier in this paper, many companies are beginning to ask, "Why are we backing up data that is not being changed in the first place?" This has given rise to the growth of "data classification" software companies in the last two years. Data classification, among other things, identifies fixed content and is the first step to move from a flat to a tiered storage architecture. Data in a tiered storage architecture can be moved to and stored in its appropriate "container," in accordance with its value, purpose or use.

If the first "wall" one hits attempting to scale network storage into the hundreds of terabytes or even a petabyte is the backup window, then the second barrier is the power-cooling-and-maintenance overhead. As of this writing, it is estimated that the cost to power and cool a SATA array for five years is approaching 50% of the total original purchase price of the system[4]. "A typical 10,000-square-foot data center consumes enough juice to turn on more than 8,000 60-watt light bulbs (sic). That amount of electricity is six to ten times the power needed to operate a typical office building at peak demand, according to scientists at Lawrence Berkeley National Laboratory. Given that most data centers run 24/7, the companies that own them could end up paying millions of dollars this year just to keep their computers turned on." [5]

Manageability is an equally important consideration, in spite of declining hard disk device costs. For example, hard disks should be replaced every three years, which means that data migration needs to occur every three years as well. In the meantime, instead of being archived, fixed content gets stored together with primary active data, taking up expensive disk space and hampering the performance of production servers. Adding to the problem, fixed content gets caught up in the back-up cycle, where it gets written and rewritten to tape, not to mention the time it takes to retrieve data. In essence, the issue is one of people-costs, not simply dollars-per-gigabyte for hardware.

All of this begs the question: Why is fixed content data consuming a storage technology aimed at speed when the requirement is for a long-term, low-cost data store? The answer is simple up to this point. There have been no practical solutions other than disk systems for the enterprise that still wants online usability of its archives.

The Emergence of the Specially Designed Archival Tier of Storage

Clearly, the answer to all the physical and economical scale issues with keeping the growing stores of fixed content online and accessible is to move it to a storage tier that is specifically designed for this purpose. In this model, primary active data and archives each occupy their own tier of storage. The primary data would sit on a system optimized for

---

4 The Clipper Group, Explorer™, "Tape and Disk Costs – What It Really Costs to Power the Devices" June 4, 2006. Report #TCG2006046L

5 http://www.cio.com/archive/041506/energy.html "Powering Down", Susannah Patton, CIO Magazine. April 15, 2006

Device \$=ƒ(Performance, Capacity) and the archive tier optimized for TCO \$=ƒ(Longevity, Accessibility). Within each tier, the storage subsystem is suited to the nature of the content to be stored.

The time for an archival tier of storage is now. The reasons are two-fold. First, storage software technologies have recently advanced to a stage where they can support a tiered storage strategy, without being a science fair project. Information Lifecycle Management (ILM) software suites and data classification/movement tools are examples of such technologies. They are available from both the world's largest software companies and small independent software vendors alike. Prior to their appearance, organizations faced a "double jeopardy" in which they were unable to adopt a tiered storage model because the software technology to enable tiering was not yet in place.

Second, in the past few years, the broader storage market has started to work through the haze of the difference between backup and archive. Given that most storage infrastructures were flat, if space on disk was recaptured by deleting older data, the thought was that if the files were ever needed again, they could be found on the periodic full backup tape catalogues. While this is currently true across the broader horizontal of corporate data, demand for an archival class of storage has existed in specific market verticals that have had long-term, fixed-content retention requirements. One such industry vertical is medical which was the first industry in the U.S. to be hit with data retention laws. Another example is government, which has always had a "save everything forever" approach.

While backup and archive have been well defined in these two verticals, the broader horizontal of corporate data storage has not faced the pressure to draw the distinction—until recently. The most economical approach was to run a flat storage infrastructure wherein data was divided into online/backed up and archive. The data that was online was backed up to tape, and backup tapes were periodically sent off site as an "archive". This is how "archive" and "offline" became synonymous in the data world. As for what was kept online, when space on disk ran short, two options existed: hit the "delete" button or, more commonly, add more disk space.

Early implementations of an archival tier of storage have included lower-cost, SATA-based solutions. Those early adopters are now finding that cheaper disk does not equal the right solution for archive data as was demonstrated earlier in this paper. An archival tier needs to be truly built from the ground up with a design criteria of TCO \$=ƒ(Longevity, Accessibility).

Optical Attributes For Long Term Data

Optical storage technologies, at their very root, have been designed for very long term storage of data that is written only once. First of all, optical-based digital versatile disc (DVD) and compact disc (CD) have become the de facto standard for archival storage in essence (admittedly in the consumer sector), having more than proven themselves as a storage medium in the living room, in the car, and on the desktop. Consider the following trends in media distribution in the last number of years:

Recordable/Re-Recordable Optical Disc Replication: Worldwide
(Millions of Units)

	2002	2003	2004
Total CD-R, CD-RW	7,425	9,220	10,480
Total DVD-R, DVD-RW*	205	910	1,935
Total Worldwide Recordable/ Re-Recordable Replication	**7,630**	**10,130**	**12,415**

*Includes DVD-RW, DVD+RW, DVD-RAM, DVD-R, and DVD+R

This data is provided as part of a strategic alliance between IRMA and Understanding & Solutions

Note that movies or music are perfect examples of archived digital data, albeit in the consumer domain. Both are long term fixed content, which have a consumer expectation of “keep forever”.

Second, optical-based DVDs and CDs are a highly reliable medium for archival storage. Some popular press reports and editorials have called into question the long-term reliability of optical (particularly consumer optical) technologies. They have written that “laser rot”, de-lamination, sunlight-facilitated oxidation, and other media/data decay causes limit the reliability of the optical medium to less than ten years.[6]

It is very easy to see why this may be the perception. Depending on what quality of media is used, the perception could very well be true. The author’s company tests tens of thousands of pieces of media a year across brands and models, finding wide deviations in media quality and extrapolated life expectancy. More importantly, according to a report by the National Institute of Standards and Technology, “It is demonstrated here the CD-R and DVD-R media can be very stable (sample S4 for CD-R and sample D2 for DVD-R). Results suggest that these media types will ensure data is available for several tens of years and therefore may be suitable for archival uses.”

Longevity of DVDs and CDs Can Vary, But There Is “Archival Quality”

Consider a test reported in October 2004 in The Journal of Research of the National Institute of Standards and Technology. “Stability Comparison of Recordable Optical Discs-A Study of Error Rates in Harsh Conditions” reported on tests whereby a temperature, humidity and light chamber accelerated environmental stress on DVD and CD media. The goal is not to reiterate the entire study here, but the graphs of the results data are very telling:

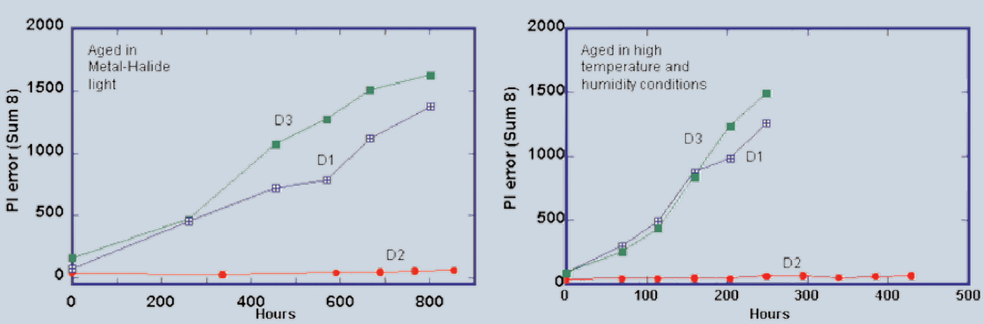

Fig. 5. PI (Sum 8) increase in DVD-R when exposed to (A) M-H and (B) extreme temperature/humidity.

In the graphs above, the “redline” sample D2 for DVD-R media showed remarkable resiliency to environmental accelerated conditions, while samples from other brands of media did not.

So Why DVD and Not UDO or Other Optical Media?

To develop a system with the goal of long-term archiving, the underlying physical platform where the data is held must provide long-term stability, thus eliminating the risk of data loss and the high maintenance needed for short-term data media including magnetic technologies. We have made mention of DVD and CD formats, but what about other options in the market? Do they exhibit the same Longevity characteristics as Ultra Density Optical (UDO) and Professional Disc for Data (PDD)?

---

6 Example: http://msn.pcworld.com/article/id,124312-page,1/article.html “Do Burned CDs Have A Short Life Span?”

Such physical media may be able to retain readable data in decades to come. But when it comes to archival storage, one must consider whether there will be a way to read the data decades from now. For those familiar with the changes in tape formats over the years, it is a common quandary of having the data but not the player. In archival storage it is important to maximize the chances of being able to read the data in many years. DVD as a subsystem component, though not as dense as other formats, provides that maximum probability of access by adhering to the broadest distributed format and form factor in the history of storage technology. When designing for TCO $=ƒ(Longevity, Accessibility), the function of Longevity and Accessibility over Performance or Capacity makes the demand that the system use components with the best chance of being available in 20 years. The components should be truly ubiquitous, not just physically able to hold the data without decay. The greater the ubiquity of the technology standard, the greater the chance one will be able to find it in decades.

Optical data standards and media formats have matured. In earlier times, proprietary formats were widespread with each optical vendor making media and devices that were incompatible with each other. There were some attempts to reach a single standard. For example, the 5 ¼ inch MO disk became available from a handful of vendors. But in general the industry never came to a single format that was backward compatible at either the hardware or software level.

This all changed with the emergence of DVD. As a derivative of CD-ROM technology, DVD is largely driven by the demands of consumer electronics. Available from a multitude of vendors, DVD has reached a true commodity status. The media costs are determined by market volumes and the manufacturing costs are low. Today optical storage technology is absolutely dominated by DVD. Another major development often overlooked is the emergence of a single standard for the Universal Disk Format. As the established data standard, UDF ensures media portability among platforms and applications. The standard eliminates the single-vendor lock-in threat for hardware or software.

If Longevity makes the technology ideal for long-term archiving of important data, then the backward compatibility of DVD data is highly significant. CDs written in 1982 are readable in today's CD and DVD players 24 years later. Likewise, DVDs written in 1996 are readable in today's DVD players, ten years later. We anticipate that DVD data written today will be readable in future HD-DVD and Blu Ray players. This preserves the DVD format for the future, an important capability as DVDs have become the ideal storage media for long-term archival of important documents.

With millions of DVD drives already in circulation, the ubiquity of DVD can offer the archive storage administrator the peace of mind of knowing that the ability to read the standard data for decades into the future is viable for the first time in technology history.

Emergence of Network Storage and Appliances

To this point we have covered all the reasons why optical at its individual component level, particularly DVD optical, is an optimal medium for long term archive given its capacity and longevity. However, we must also address the business complexity of incorporating optical into a storage environment. For all intents and purposes, there are four reasons that have thus far led to the failure of optical in the enterprise:

1. Cost

2. Proprietary nature of medium
3. Performance
4. Reliability

Overcoming Cost and Proprietary Issues

Specifically, when it came to the issue of cost, optical systems traditionally required a longer term total cost of ownership argument in order to justify their first-time cost, together with the requisite optical expertise needed in the enterprise to administer the systems. This causes two problems in the buying process, which we believe have held optical subsystems back. First, CIOs and their staffs tend to have expected tenures, based on historical data of inside 3-4 years. It is nearly impossible to build a business case that makes sense only after executives and employees are gone. Second, addressing the cost issue of long periods of time begets the trouble with the proprietary issues. The fact is the longer the period of payback for a system, the higher the perceived risk of incorporating it. With enhanced timelines to justify the payback of the technology it calls to question its useful life, which in this case is, in part, the ability to read the optical media. When choosing proprietary optical formats given the track record of optical companies, these perceived risks could not have been higher.

Standardizing on consumer-driven, ubiquitous standards such as DVD precipitously drops the upfront systems costs. It also reduces the risk that the format will be unreadable in the years to come, opening up opportunities for customers and vendors alike, while putting more burden on the system provider to create a system-level enterprise-grade product.

Overcoming Performance and Reliability Issues

For all practical purposes, network attached storage (NAS) did not exist in the enterprise environments of the early 1990s, at least in any scale. Network Appliance was a pioneering startup at the time going public in the first half of 1995 on a total revenue stream in the low $40 million range. The concept of a dedicated, single-purpose appliance was alien to most IT managers, primarily due to the high cost of server hardware. Inexpensive Intel-based commodity systems simply did not have the performance capabilities to operate in the enterprise environments and were not off-the-shelf disks (“Inexpensive Disks”).

The advent of NAS essentially introduced the concept of a dedicated, special purpose computer aimed at optimizing storage pooled from multiple disks as a network attached volume. The duty of the system head end is to “virtualize” the complexity of managing the underlying collection of “JBOD” or Just A Bunch of Disks. An individual disk may not be fast enough, reliable enough, large enough, but a collection of disks, completely and logically unified in different pools can address these issues. It is beyond the scope of this paper to delve into Redundant Array of Independent/Inexpensive Disks (RAID) but this approach demonstrated that the problems of speed, size, performance, and reliability issues of relatively inexpensive disks acting independently could be effectively addressed and that storage as a “network utility” could be accomplished inexpensively.

The proliferation of standard network file system protocols enabled easy integration of the appliance architectures into the enterprise storage infrastructure. At the same time, open source technologies reached the level of maturity enabling them to be competitive in enterprise computing. In particular, they accelerated the cost-effective development of the embedded software for storage appliances.

These NAS dynamics have an impact on how optical storage is being implemented today with DVD and will open up a new possibility for optical as an archival medium.   Just like NAS, a dedicated appliance can now control multiple DVD auto-changer libraries. It is integrated into the enterprise storage infrastructure as a network attached storage device and is available to all the enterprise via standard storage protocols (CIFS and NFS). All the details of media management and specifically of DVD I/O are hidden by the software running in the appliance. Speed, size of the incremental media pieces, and reliability of individual media are all addressed in exactly the same fashion NAS addressed the challenges of using JBOD the decade prior. As NAS and RAID did in the decade prior with magnetic disks, these same technological approaches have been shown to very effectively eliminate the issues of ease of use, performance and reliability that individual optical media simply cannot address outside of this system approach.

Archival Storage Tier Based on Optical Technology

The reasons for the coming rebirth of optical technology, specifically DVD and their blue laser counterparts, are now becoming evident. As discussed in this paper, the reasons are partly marketplace dynamics and partly a new and more appropriate design criteria for optical systems. The maturing of standards and formats, changed economics, and other factors have collectively transformed the outlook for optical in the enterprise. Storage architecture and design approaches can leverage optical technology for its optimal use:

TCO $=ƒ(Longevity,Accessibility)

Whereas previous applications of the technology were geared to general-purpose storage, they now support the requirements of long-term archiving. Previously the guiding principle was device price/performance/capacity, but the requirements changed. Today, the driving factor is TCO/longevity/accessibility in addition to performance/capacity. To support the new archiving requirements, enterprise storage needs to offer the best of both worlds: the longevity of optical and the accessibility of network-attached storage. The benefits are clear: An organization can cost-efficiently store vital digital assets for extended periods of time with the assurance that the information will remain in an unaltered form, while maintaining the ability to quickly retrieve the data when needed.

Jonathan W. Buckley
October 11, 2006