Applied Engineering Science Inc.
www.appliedengineeringscience.com
PO Box 696 Eastlake, CO 80614
(303) 920-8180

Protecting Digital Assets for the Long Term

By Joe Jurneke

Abstract

This paper discusses the various aspects of archive operations and risks to data stored in the digital domain. The most probable technologies for long-term archive reliability based on modern technology will be discussed.

With the massive migration of information from the historical "analog" mechanisms toward the ever-growing digital domain, one question remains paramount in the minds of everyone – "how long will my data survive, and can I get it back?" The content herein is based on published papers showcasing technologies, performance, and limitations of specific technologies

Competing Technologies

Today there are four competing technologies for archiving digital data. This includes magnetic tape, magnetic disk, optical disk, and solid-state non-volatile ram (flash, or equivalent). Two are competitively priced for storage of large digital databases ($/GB capacity) – magnetic disk and magnetic tape.

Optical disk is lower in capacity, significantly higher in $/GB, and significantly slower in transfer rate than either disk or tape. Solid-state memory is prohibitively expensive for archive applications. Therefore, discussions will be limited to disk and tape technologies.

There are four fundamental applications for magnetic disk and tape – online data storage, nearline data storage, backup, and offline archive applications.

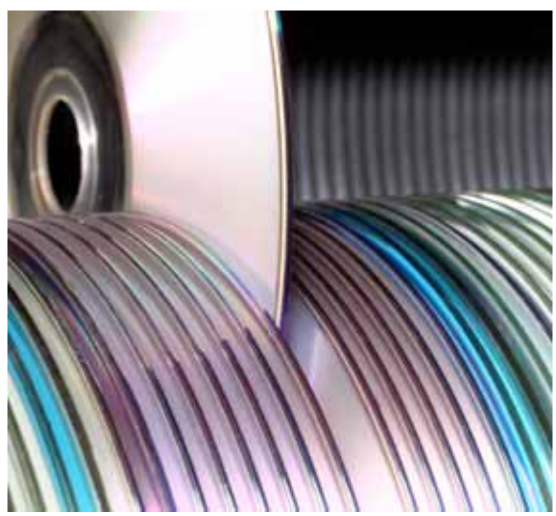

In online storage applications, storage drives (disk drives and tape drives) are typically powered on and left operational for many hours without interruption. Nearline data storage involves the requirement to access data less frequently than online, and because of the sheer quantity of data, can comprise both spinning disk and library accessed tape or sometimes just lower cost disk arrays.

As backup applications are more and more frequently understood to be specifically for disaster recovery situations, solutions for backup have been migrating from tape to spinning disk in recent years.

Archive applications in particular comprise solutions where data is written to media for the purpose of long term preservation with the intent of reading at a future point in time. Archived media is often stored on the shelf, but can be kept online or nearline as well.

Applications Demanding Maximum Reliability

The most demanding application for reliability is long-term archive. Customers expect to recover data for many years after creation of the archived data set. Loss of data is excruciatingly painful, and in many cases intolerable.

Video assets are a good example of the need for long-term archive. In this environment, additional copies of the data may not be available. Hence a loss of data is a permanent loss of assets. Maintaining multiple copies may help alleviate this problem, but in certain circumstances may also provide for a false sense of security.

Verifying the integrity of data and data quality is paramount to creating an archive quality data set. Verification may involve checking the integrity of data on media after some period of storage time, which as discussed later is more important for disk drives than for tape. Frequently archivists are more

Archive Environments

Many organizations have inventories of thousands to tens of thousands of volumes of data.

Examples of some of the largest archive environments are:

1. Video Tape Libraries – TV stations, broadcast networks
2. Movie and entertainment libraries – asset creation
3. Animators, special effects laboratories
4. Museums
5. Internal Revenue Service
6. Social Security Administration
7. Library of Congress
8. Fortune 500 corporations
9. Federal Government archives
10. National Archives
11. Space Missions
12. Geophysical
13. AT&T Long lines Division
14. NSA/CIA/Homeland Security
15. Medical Records

focused on verification that the data was written correctly to begin with. The two technologies under examination here differ greatly in this respect:

- In the case of tape, data is verified as it is being written (read-verified-writes) with no penalty in transfer rate.
- In the case of disk, a read after write pass is required to verify integrity of the data, thereby reducing the transfer rate by at least a factor of 2.

Archive Failure Mechanisms

Several factors affect long-term stability of magnetic recordings. Storage and operating environment play a critical role in long-term stability of recorded data. Temperature and humidity should be controlled to specifications outlined by recommendations from the manufacturers of the magnetic media. This applies to both disk and tape systems.

The National Media Lab (NML) in 1995 published a report on archiving magnetic media. NML recommended maintaining a temperature of 20 Deg Celsius at a relative humidity of 40%.[1] Since that time, areal densities (amount of data stored per square inch of magnetic surface) have increased many-fold. The environmental considerations today are more important than they were 15 years ago.[2]

In fact AMIA (Association of Moving Image Archivists) in 2007 recommended a storage environment of 7 Deg Celsius and a relative humidity of 25%.[2]

Mechanical handling can play a crucial role

Disk/Tape ECC Characteristics

Disk defect management is typically based on eliminating defects by retiring bad sectors at the time of manufacture.

When a customer records data to disk, the drive knows which sectors are bad, which are good. The ECC is intended to provide adequate coverage for a small amount of operational degradation over time, or from defects missed in the manufacturing process.

Tape does not have to manage defects. With the ability to read the data at time of recording, the tape drive knows if the data is intact. Should the data be found either bad, or marginal, tape drives will automatically re-record data in the next good block on tape to ensure reliability.

in loss of reliability. Drop a disk drive on the floor, and the probability of retrieving data is significantly reduced. Based on published prices from data recovery firms, the cost of recovering data from a disk drive can exceed $4,000, after rebuilding the drive. Success in recovery is not guaranteed.

Drop a tape cartridge on the floor, and there is a probability that some edge damage may take place. With tape, in a worst-case drop event, a majority of the data can be recovered, if not all data. Nearly every tape manufacturer states that their tape cartridges pass rigorous drop testing evaluations.

Both technologies are sensitive to severe shock, but tape has been demonstrated to be much more robust in this regard.

Error Correction Code (ECC) organization between disk and tape are quite different. Disk organizes it's ECC intermixed with the data it protects in the same track/sector. Tape organizes its ECC multi-dimensionally. Tape spreads the ECC syndromes out over the surface, down track, and across multiple tracks thereby obtaining significant protection from damage and defects. Additionally, in many cases, the data itself is spread multi-dimensionally to improve robustness against defects and damage. This is why it is easier to recover data from tape errors than from hard disk errors.

Assuming that one can navigate around the environmental and the mechanical shock issues, what other concerns remain?

Disk and tape drives are complex mechanisms. Both are very sophisticated devices, with a tremendous amount of technology content. When archiving to disk, the entire disk drive is archived, electronics, media, heads, actuators, spin motors, etc. In contrast, in an archived tape environment, only the tape volume is archived whereas the drive mechanism with its heads and motors is not – much less hardware to degrade over time. When it comes to ultimate reliability, in the aggregate less complexity results in higher reliability.

Hard Disk Drive Failure Rates

Hardware reliability is broken down into three key time periods – infant mortality (early life failure), useful life, and wear-out (end of life). Utilizing disk drives for archive applications especially when stored individually on the shelf makes concerns over infant mortality failure rates more justified and the results of such failures all the more damaging.

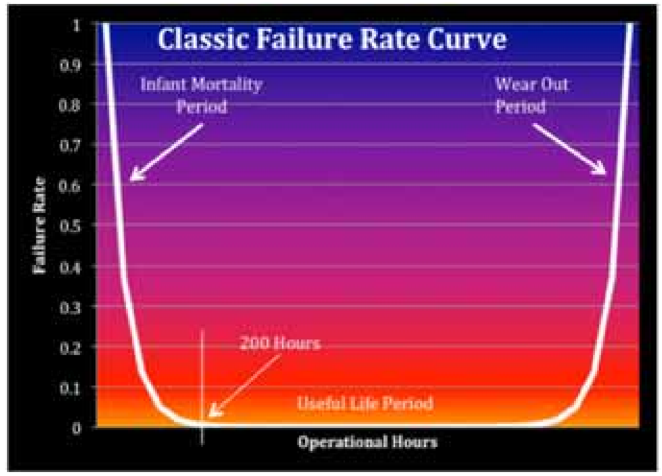

Based on the published data from Google[17], CMU[6] and others[5, 7], disk suffers from a significant early failure rate (infant mortality) problem. In addition, as capacity grows, the failure rates tend to increase proportional to capacity. The aforementioned papers are based on large drive populations in excess of 100,000 units. From the perspective of archive applications, infant mortality (early failure) failure of disk drives is most disturbing. Since most "shelf storage" offline archive operations will involve new drives, these volumes are particularly vulnerable to infant mortality failures.

Drive manufacturers typically specify MTBF (Mean Time Between Failure) as a measure of reliability. Most MTBF computations remove "No defect found" and "Cannot Replicate Problem" from the failure rate estimates. Literature suggests that this may result in an over-optimistic estimation of MTBF by 30-40%.[6, 18] MTBF may not be an appropriate metric for use in determining archive performance.

From the customer's perspective, any drive removed from service for any reason constitutes a failure. With many hours of time tied up in creating the recorded volume, failure for any reason becomes expensive and time consuming.

While the specified 1 million hour MTBF looks impressive, many drives fail within 90 days. In fact, according to the Google data, a drive is 39 times more likely to fail within 60 days than during the useful life period.[17]

Some reports suggest that disk drives prior to being placed in service should be burned in for at least 3 months, or 2000 hours.[3, 4]

Data also suggests that the reliability of “high end” drives is no better than consumer grade hardware.[18]

The CMU[6] and Google[17] data suggests that failures are affected by power on hours, operational activity, vintage, operating environment, and supplier. The CMU reports the infant mortality concerns persist for up to 200 days of operation before transitioning into useful life, with a peak in failure events at about 20 days.

This infant mortality data raises concern over utilizing shelf stored, offline disk as a reliable archive medium. A prudent archivist might consider making a duplicate copy of an asset. Even this approach has its risks.

When organizations purchase hardware, they would typically procure a fairly large quantity of drives at one time. Drive failures are not necessarily random. Many factors affect “vintage” concerns, such as changes in processes, suppliers, personnel changes in assembly, new equipment introduced into manufacturing operations, etc. The new devices are likely coming from the same manufacturing lot, all produced at about the same time, using parts from the same manufacturing lots and processes.

Failure rates and failure mechanisms for this specific lot of material would be expected to be quite similar as reported in the literature[17]. If one drive fails early, it might be reasonable to assume that other drives of the same vintage might also fail early. Vintage issues may imply that use of multiple copies as a strategy for improving reliability may not be appropriate. If a given vintage fails early, there is a distinct possibility that other drives from the same manufacturer’s vintage may also fail early leaving the integrity of the archive volume in question.

Thermal Stability of Magnetic Particles

Magnetic particles, whether tape or disk, are subject to degradation over time. How quickly they degrade is of concern to the archivist.

Temperature has a significant impact on the stability of magnetized particles. Increased temperature reduces the expected lifetime as reported in the NML[1] and AMIA[2] documents.

There are two competing forces at work in magnetized media. The blocking force, that which establishes and attempts to maintain the magnetization state, is driven by the magnetic properties of media and by the particle (or grain) volume. The destabilizing force is driven by temperature[24]. When the blocking force and the destabilizing force are equal, the particle can change magnetic state. Particle magnetization tends to decay exponentially over time. The decay occurs from a fully magnetized condition to a near zero remnant condition.

To quantify these phenomena, a figure of merit has been established for thermal stability. This figure of merit is referred to as "KuV/kT." Ku is the particle anisotropy constant, V is the particle volume, k is Boltzman's constant, and T is temperature in degrees Kelvin. As long as this ratio is greater than 65, a 30 year life can be expected due to thermal stability. The effect of this figure of merit is exponential. A life of 5 years has a KuV/kt of about 45.

Thermal Stability Concerns

Thermal stability is a critical parameter in the longevity of signals recorded on magnetic media.

Lower thermal stability of the media can result in premature signal fade thus reducing a media's archival life. This is why some experts recommend spinning up disk drives periodically to check for signal loss and rewrite data that is fading.[20, 21]

Conversely, magnetic tape's higher thermal stability figures suggest that it will exhibit a significantly longer shelf life.

Literature suggest that magnetic disk (modern perpendicular oriented media, or patterned media) supports a KuV/kT coefficient of 45-60.[12, 13, 14] Magnetic tape supports a ratio much greater than 100 for conventional MP particles.[14] Hence, tape is much more stable thermally from an archive perspective.

Long Term Storage on Tape

There are still customers that regularly utilize open reel 7 and 9 track NRZI as well as 1600 bpi PE/6250 GCR digital data tapes created from the 1950's through early 1980's. Tape drive repair and maintenance services and refurbished tape are available from several sources. At least one company is providing a device for use in conversion to more modern standards.Archive life for tape is heavily dependent upon environment. By following the archive environment recommendations, obtaining 30-year life is entirely feasible with modern technologies. All mid-range and enterprise class tape devices for the past 40 years has driven archive performance as a major consideration in design. Data storage media utilized in LTO format devices based on metal particle formulations is tested to and promoted as having a 30 year life expectancy by media manufacturers such as Fujifilm, Maxell, Sony & TDK.

Recent announcements from FujiFIlm[8, 9] IBM, and Oracle have demonstrated at least 3 future generations of tape product feasibility utilizing new state of the art Barium Ferrite tape for digital data storage. Along with those demonstrations, the archive life of 30 years has been reported for Barium Ferrite. This claim has been verified by tape drive manufacturers.[8]

Tape plays a crucial role in every aspect of mid-range and enterprise data integrity. There are many different tape formats in the marketplace. Focus consideration on the most popular backed by strong companies, e.g. LTO technology backed by HP, IBM & Quantum. – LTO currently accounts for over 90% of midrange tape drive sales. Since LTO's inception ten years ago over 3.5 Million tape drives and over 150 Million tape cartridges have been sold.[10, 11, 19, 22, 23]

Summary Recommendations

A true archive technology (for 100 years or more) does not exist today for digital content. The task at hand is to choose a medium that will provide the most protection possible for the longest possible period.

Disk was never intended as an archive medium historically. It has always been intended for direct access, online data storage. Burning in disk for a thousand hours, or 3 months, as recommended by one reference[4] prior to using it is very expensive and time consuming.

Tape on the other hand has always been intended for backup and long-term archive applications. As an archive media, it does not suffer from infant mortality or electro-mechanical failure mechanisms. In addition tape media has higher thermal stability, which translates into being able to read data from shelf stored media for longer periods of time compared to disk.

In comparing the intrinsic reliability of a tape cartridge versus archiving an entire disk drive, simplicity should always win out – the tape should be more reliable from a hardware reliability perspective.

There are many different tape technologies in the marketplace. Some are superior to others. Choose the technology and the supplier wisely. Formats that are experiencing growth in market share are usually good candidates. Technologies with shrinking market share should be avoided. LTO technology, backed by HP, IBM and Quantum, commands over 90% share of the midrange data tape drive market.

Failing tape drives are usually not as catastrophic as failing disk drives. Tape products are easier to repair. In making repairs,

Cautions for Hard Disk Drive Archiving

The conclusions drawn herein would suggest that archiving to hard disk drives is not a recommended practice. That said, many users are doing so anyway. Thus they should take into consideration the best practices that can be distilled from the reports consulted in this paper.

Infant mortality statistics suggest that hard disk drives should always be burned in before being used for archive – any burn-in period will help, the longer the better with some sources recommending 2000 hours.

Thermal stability characteristics of hard disk drive media suggest that they be stored in a cool dry location, and that increases in temperature can lead to increased likelihood of data loss.

Physical protection for hard disk drives is rarely emphasized as USB and firewire enclosures populate closet shelves in growing numbers. Placing these enclosures in locations where they aren't easily knocked over and handling them carefully can't be undervalued in archive applications.

Periodic refresh of data on disk by spinning up each drive and reading each bit is a recommended practice by independent sources in the industry.[20, 21]

your data is not put at risk, as is the case with disk drive repair, especially failures involving the interface to the Head/Disk assembly.

A number of particular issues have been brought to light in this paper when considering archival solutions for data. Infant mortality is an often-neglected issue when archiving to hard disk drives and can significantly contribute to data loss scenarios in shelf-stored disk drives used for archival purposes. Further, lower coefficients of thermal stability of hard disk drive media as compared to data tape, can result in data loss over time. In summary, magnetic tape is a superior archival media for offline, shelf stored content.

Bibliography

[1] Dr. John W.C. Van Bogart "Magnetic Tape Storage and Handling" June 1995, National Media Laboratory

[2] Association of Moving Image Archivists, 2007

[3] Bianca Schroeder Garth A. Gibson Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? 5th USENIX Conference on File and Storage Technologies, San Jose, CA, Feb. 14-16, 2007

[4] http://itmanagement.earthweb.com/datbus/article.php/3843706/How-to-Reduce-Hard-Drive-Crashes.htm

[5] Feng-Bin Sun, PhD, Shaoang Zhang, PhD, Does Hard Disk Drive Failure Rate Enter Steady-State After One Year? Reliability and Maintainability Symposium, 2007, RAMS '07 Annual pp 356 – 361

[6] Bianca Schroeder Garth A. Gibson Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? 5th USENIX Conference on File and Storage Technologies, San Jose, CA, Feb. 14-16, 2007

[7] Sandeep Shah, Jon G. Elerath, Reliability Analysis of Disk Drive Failure Mechanisms Reliability and Maintainability Symposium, 2005. Proceedings. Annual pp226 – 231

[8] http://www.fujifilm.com/news/n100910.html "Long-term Archivability and Stability of Fujifilm magnetic tape using Barrium Ferrite (BaFe) particle."

[9] M. L. Watson, R. A. Beard, S. M. Keintz, and T. W. Feebeck, "Investigation of Thermal Demagnetization Effects in Data Recorded on Advanced Barium Ferrite Recording Media" IEEE Transaction on Magnetics, vol 44, No. 11, Nov 2008.

[10] Fred Moore, Horison, Inc.Tape: New Game. New Rules - Tape Re-architects for 21st Century Data Explosion

[11] http://www.clipper.com/research/TCG2010054.pdf In Search of the Long-Term Archiving Solution —Tape Delivers Significant TCO Advantage over Disk Analysts: David Reine and Mike Kahn

[12] Wei Choong Poh., N. Piramanayagam, and T. Liew Magnetic Properties and Corrosion Resistance Studies on Hybrid Magnetic Overcoats for Perpendicular Recording Media IEEE TRANSACTIONS ON MAGNETICS, VOL. 46, NO. 4, APRIL 2010 1069 pp 1069-1076

[13] Wei Choong Poh., N. Piramanayagam, and T. Liew Magnetic Properties and Corrosion Resistance Studies on Hybrid Magnetic Overcoats for Perpendicular Recording Media IEEE TRANSACTIONS ON MAGNETICS, VOL. 46, NO. 4, APRIL 2010 pp 1069-1076

[14] M. V. Lubarda, S. Li, B. Livshitz, E. E. Fullerton, and V. Lomakin Reversal in Bit Patterned Media With Vertical and Lateral Exchange Center for Magnetic Recording Research, Department of Electrical and Computer Engineering, University of California, 2010 San Diego, CA 92093 USA

[15] Long Term Archivability and Stability of Fujifilm Magnetic Tape Using Barium-Ferrite (BaFe) Particle White paper

[16] M. V. Lubarda, S. Li, B. Livshitz, E. E. Fullerton, and V. Lomakin Reversal in Bit Patterned Media With Vertical and Lateral Exchange, IEEE TRANSACTIONS ON MAGNETICS, VOL. 47, NO. 1, JANUARY 2011 pp 18-25

[17] Failure Trends in a Large Disk Drive Population Eduardo Pinheiro, Wolf-Dietrich Weber and Luiz Andr´e Barroso, Google Inc. 1600 Amphitheatre Pkwy Mountain View, CA 94043, 2007

[18] Jon G. Elerath, Network Appliance, Inc., Sunnyvale Sandeep Shah, Network Appliance, Inc., Sunnyvale "Server Class Disk Drives: How Reliable Are They?" RAMS 2004

[19] Back-Up Tape Technology: 2011, Santa Clara Consulting Group

[20] Jordan, Larry: Hard Disk Warning!, Larry's Monthly Final Cut Studio Newsletter, September 2008, http://www.larryjordan.biz/articles/lj_hard_disk_warning.html

[21] Jordan, Larry: A Free Utility to Refresh Hard Drives, Larry's Monthly Final Cut Studio Newsletter, December 2010, http://www.larryjordan.biz/articles/lj_refresh_hard_drives.html

[22]LTO Ultrium Consortium http://www.ultrium.com/index.html[23]http://davesimpsonsstorageblog.blogspot.com/2011/02/lto-crushes-other-tape-formats.html

[24] Darin J. Sleiter and Michael P. Sharrock "Magnetic Time Effects in Small Metallic Particles for Recording: Orientation Dependence" IEEE TRANSACTIONS ON MAGNETICS, VOL. 40, NO. 4, JULY 2004 2413

About the Author

Mr. Joe Jurneke has 38 years of engineering experience in storage peripherals, 12 years in Disk drive development, 26 in tape systems. He has contributed both as an individual technical contributor, and as a technology manager/Director, with 28 years of management experience. He has played a critical role in the development of 10 enterprise level world class storage products, a multitude of mid range products, and has spent 8 years driving head and media technology for tape products providing advanced technology 2 to 5 years ahead of the product requirements. Mr. Jurneke holds as inventor/co-inventor 9 US patents, 4 others pending. He also has received 2 inventor recognition awards, and the StorageTek Chairman's quality award 1996.

Applied Engineering Science, Inc

Po Box 696 Eastlake, CO 80614
(303) 920-8180
Info@appliedengineeringscience.com
www.appliedengineeringscience.com