

THE LANDSCAPE OF DATA STORAGE

Cost Advantage per Terabyte of Data

Page | 1

Sponsored by:


Hie Electronics

Author:

Kyle Hodges

Published:

July 2011



Overview

This white paper provides a brief synopsis, advantages, and disadvantages of four different types of media: Hard Disk Drive, Solid-State Drive, Magnetic Tape, and Optical. These media types are evaluated and assessed by their cost per terabyte, capacity, data integrity, and the associated risk in owning that type of media. This paper will focus just on the individual media types and not the associated equipment. Other points of observation are noted to provide a more holistic approach to choosing the correct media for your company's needs.

THE LANDSCAPE OF DATA STORAGE

Hard Disk Drive Solutions

<i>Media Cost per Terabyte:</i>	<i>\$110</i>
<i>Advantages:</i>	<i>High Capacity + Performance</i>
<i>Disadvantages:</i>	<i>Power + Heat</i>
<i>Max Capacity:</i>	<i>3 TB</i>
<i>Archive Life:</i>	<i>3-5 Years</i>
<i>TCO:</i>	<i>High</i>

Hard Disk Drive (HDD) technology has advanced rapidly in the past several years, providing higher performance and capacity at significantly reduced prices. Hard drives are primarily designed for capacity, performance, and short term availability. For the most part, hard drives are constantly spinning, constantly ready to retrieve the requested data at a moment's notice. The constant spinning, vibrations, and harmonics created by the myriad of mechanical parts in the HDD generate heat. Power and cooling requirements are the resulting operational costs involved with hard drive solutions. There is also a need for Tiered Data Architecture for storage using HDD.

A main concern is the reliability of data stored on the hard drives. In order to provide the required levels of reliability, hard disks are paired into a Redundant Array of Independent Disks (RAID) configuration for increased performance using striping or for data redundancy using mirroring and parity. Mirroring is accomplished by striping copies of the data across the RAID to avoid loss of data when one or more drives fail. Striping provides defect management using a single drive or portions of several different drives as error correction buffers. Various RAID configurations can be implemented according to the degree of performance or data protection required. Data managers realize the dire consequences of losing important data with a hard drive failure, and they understand that it is not a question of if the hard drives will fail, but when. In order to protect their data, additional hard drives are being added to the equation, further increasing the workload of data managers and the overall cost of operations. Mirroring an entire RAID array with another offsite RAID array for disaster recovery can effectively more than double the operating cost to manage the data.

Carnegie Mellon University performed a detailed study of Google Inc.'s data storage architecture in February 2007 documenting annual failure rates among their multi-petabytes of hard drive arrays. Two petabytes (2,000TB) of usable storage would require approximately 9,600 250GB Hard Disk Drives (HDD) for initial implementation. With a 20% RAID overhead included for redundancy and assuming a 40% data growth rate, the total hard drive requirement will increase to almost 37,000 drives by the fifth year of operation, representing a total of 7.6TB of data. The Carnegie Mellon study found that in the second, third, and fourth years of operation, hard drives will fail at the rate of approximately 8% per year. In the fifth year this number jumps dramatically to 24.3% and in year six up to over 50%,

THE LANDSCAPE OF DATA STORAGE

requiring the replacement of nearly five HDD per day in year five, and 10 per day in year six. (Gibson & Schroeder, 2007)

In general, HDD are not designed for long term availability and accessibility. Due to how the HDD operates and the components inside them, the usable lifecycle or MTBF of a HDD is 3-5 years. An assumption can be made that after year 4, an HDD will become obsolete and too risky to keep active as the fear of losing data will accelerate. The amount of data stored will steadily grow and the purchase of more hard disks will mirror that. The rate of disk replacement will thus become exponential to ensure data retention.

Solid-State Drive Solutions

Media Cost per Terabyte: \$2,860

Advantages: High Capacity + Long Term Energy Passive + Size

Disadvantages: Cost + Reliability + Energy Use

Max Capacity: 2 TB

Archive Life: 5-7 Years

TCO: Very High

Solid-State Drives (SSD) are different from most styles of storage in that they use solid-state flash memory semiconductors. Due to the nature of the storage medium, the retention of data is held in a relatively non-volatile random-access memory chip with no movable parts. This allows for very low access time and higher immunity to physical shock. The low-latency of the SSD is the greatest advantage, making a very high rate of access when reading or writing data. The rate of error is very low as well; SSD has less vulnerability to failure than HDD. However, an increased number of write cycles on a single cell will decrease the retention and increase the risk of failure of that cell structure entirely. (Gibson S. , 2007)

Power to SSD is not needed to retain data, which is essential for long-term data retention. By using a NAND-based flash memory, a type of solid-state Random-Access Memory (RAM), speed of access can be up to 50 times faster than HDD. Having the same interface as HDD, compatibility is not a problem with most systems. Because SSD hardware is so compact, there is a pump required to help push the needed electricity throughout the circuit. This can cause a buildup of electrical pressure from excessive writing to the SSD. The more information is processed and energy is pumped into the SSD, the more a heat sensitive environment is caused. SSD energy consumption can actually be higher than HDD given the platform used. This is more apparent when continuously using the SDD at capacity; thus, requiring constant energy use. (Roos & Schmid, 2008)

THE LANDSCAPE OF DATA STORAGE

Degradation of an SSD is referred to as seasoning. This becomes more likely when the SSD is at capacity. Seasoning leads to a more pronounced latency, especially during write processes. The lifespan of an SSD depends on how many times it has to be written to. (Tokar, 2010) The measure of failure is known as the MTBF, which manufacturers claim to be a lifespan of 30 years. Being so new to the datacenter storage scene, verification is not available. There is a fear that MTBF is more than likely falsified, similar to HDD portraying life in optimal conditions with minimal input-output. However, there are reports of data retention to range from a minimum of 1 year, to an average of 3 years, and a maximum of 10 years. The industry standard for this media is to maintain a schedule of data migration to occur every 5-7 years. The conclusion is, the more writes and erases, the shorter the lifespan and the smaller the data retention due to degradation of the oxide barrier. (Western Digital, 2009)

Magnetic Tape Solutions

<i>Media Cost per Terabyte:</i>	<i>\$30</i>
<i>Advantages:</i>	<i>Low Cost + Energy Passive</i>
<i>Disadvantages:</i>	<i>Environment Conditions + Reliability</i>
<i>Max Capacity:</i>	<i>3TB</i>
<i>Archive Life:</i>	<i>5-7 Years</i>
<i>TCO:</i>	<i>Medium</i>

Tape based backup and archive has been the backbone of many hierarchical data management policies over the past several decades. Tape media's low cost, inert retention, density, and high sustained data rates have helped magnetic tape manufacturers maintain a commanding hold on the data backup and archival industries.

Magnetic media is very sensitive to environmental conditions, especially to outside magnetic fields such as those created by a magnet or an electromagnetic pulse (EMP) device. Tape media requires a controlled environment with strictly controlled temperature and humidity levels to prevent premature degradation. Tapes access data serially, potentially creating an input-output bottleneck when processing several simultaneous restoration requests. Magnetic tapes are prone to jam, bind, and tear, resulting in a need to replace the damaged tape, forcing data managers to generally create redundant backup tapes. Reliability has increased with each new generation of tapes, but the possibility of data corruption during long term storage has led many to seek out a more reliable solution.

Magnetic particles known as pigment make up the storage medium and will start to degrade shortly after being manufactured. This degradation is known as "sticktion" or the Sticky Shed Syndrome. The Library of Congress has claimed that the bindings in magnetic tapes are "inherently instable" and

THE LANDSCAPE OF DATA STORAGE

must be assumed to have a finite lifespan. The useful active lifespan is noted to be assumed at 5-7 years. (Van Zelst, 2008)

Recent legal proceedings have shown the cost of non-compliance or lost archive data to be increasingly higher; this increase is tipping the scales in favor of higher reliability data retention solutions. This is especially so considering storage of legal documents, which must be held for a minimum of 7 years due to federal retention requirements. (IRS, 2008)

Blu-ray Optical Solutions

<i>Media Cost per Terabyte:</i>	<i>\$110</i>
<i>Advantages:</i>	<i>Retention + Environment</i>
<i>Disadvantages:</i>	<i>Low Capacity + Bitrate</i>
<i>Max Capacity:</i>	<i>128 GB</i>
<i>Archive Life:</i>	<i>50-100 Years</i>
<i>TCO:</i>	<i>Very Low</i>

Optical solutions have only recently achieved a storage density that allows it to truly be a viable option for long term data retention. Blu-ray is ideal for many different reasons; it is optimal in long term data storage, backwards compatible, and low cost. Blu-ray media is extremely resilient and guards data against harm from temperature fluctuations, moisture, water damage, power interruption or spiking, magnetic variations, critical system failures and aging. Blu-ray eliminates the need for a Tiered Data Architecture by having very little environmental requirements. The personal modifications for security on individual discs are quite vast. With Blu-ray having the option to be Write-Once Read Many (WORM) media, it is perfect for maintaining sensitive information for a long period of time without the fear of data being altered. Another option for Blu-ray is the re-writable disc which allows for this optical media to be able to compete against the other media types.

The current capacity of Blu-ray is 128GB of storage per disc. Research is paving the way to even more optical improvements; a recent press release introduced the new Micro-Holographic Disc. This next generation optical storage is technologically advanced to get capacity of 500GB while still having the size and speed of Blu-ray. This technology still works like a Blu-ray disc, which has an orbital read/write media that takes time to process data, causing data access latency. To combat this, Blu-ray uses RAM to achieve a quick response. (GE, 2011)

Blu-ray media has been subjected to accelerated manufacturer lifecycle testing of 80°C/85% relative temperature and humidity media testing. Manufacturer testing has shown that write-once Blu-ray media can maintain the data integrity for up to 100; re-writable Blu-ray media has been rated to

THE LANDSCAPE OF DATA STORAGE

reliably store its data for up to 50 or more years. A distinct advantage is the ability to store data without worries of environmental effect. The Blu-ray disc can be fully submerged in water, dried off, and its data is still accessible. The Blu-ray burn process automatically performs bit level verification of data at the time of write as an added assurance that the data has been successfully written to the disc.

Blu-ray is not susceptible to electro-magnetic pulse (EMP) like HDD and magnetic tape and is not susceptible to the electron separation that makes SSD unreliable. There is no contact between the optical pickup (OPU) and the disc, allowing Blu-ray to be read an infinite amount of times, setting it apart from all other media. (Panasonic, 2011) A rewritable Blu-ray disc can also be rewritten at least 10,000 times. (Bennett, 2011) Blu-ray is also a passive storage media, significantly reducing power consumption, which is a large benefit with the rising trend of cost for electricity. More importantly, the environmental conditions needed to implement Blu-ray media into a datacenter does not require the substantial environmental controls that all the other media recommend. The total cost of ownership using this media will, therefore, dramatically drop.

Conclusion

In a perfect world there would be no latency effect, power consumption of the equipment and environmental controls would be at an acceptable rate, and capacity of raw data per unit would be as high as possible regardless of media type. Most importantly the data retention would not be in question; also, there would be no concern of any legal or ethical requirements. However, the realistic world today gives us multiple choices, each with drawbacks. The dangers could lead to a major disadvantage in the long run or even short run. As noted in this paper, data retention is a major concern and selecting the right storage platform is a very important choice in the route to achieving your company's goals. The chance of losing data is dependent upon the media and an inability to archive data will be detrimental and possibly cause legal consequences. Choosing a reliable long-term media is the only solution.




A data centers' storage medium is a primary concern of the IT department. IT has transitioned from a role of technology implementation and into a strategic area of interest to any company. Cutting cost and minimizing risk are part of any good business tactic, and making sure the IT department is reducing costs in the long run and ensuring reliable data redundancy for the future is becoming a big part of this effort. The next step is to move IT towards power efficient equipment, while not forgetting that data should be on a media solution that is secure, consistent, economical, and advantageous to data retention. The choice that suits all of the criteria will be the Blu-ray disc. Figure A is a comparison chart that reviews this decision while Figure B will show the steps to data migration based on the useful life of the media type.

THE LANDSCAPE OF DATA STORAGE

Figure A: Overview of Media Solutions

	\$/TB	RAW Capacity	Energy Use	Max Archive Life	Reliability	TCO	Latency
HDD	Low	High	High	Very Low	Low	High	Low
SDD	Very High	Medium	Low	Low	Low	Very High	Very Low
Tape	Very Low	High	Medium	Medium	Medium	Medium	High
Blu-ray	Low	Low	Very Low	Very High	High	Very Low	Low

Figure B: Blu-ray Disc Advantages as an Archive Media

	2010-	2020-	2030-	2040-	
 BD BD Media (UDF format) BD Drive					Because it is a de facto standard, the supply of BD drives and parts can be expected to continue for a long time.
 Tape Tape Media Tape Drive					The drive becomes obsolete in about 10 years, resulting in the need for a format change.
 HDD					Since it is assumed that HDDs will be replaced, data migration is frequently required.

Source: Panasonic

Logical Deduction of Long Term Media Solution



THE LANDSCAPE OF DATA STORAGE

Works Cited

- 42U. (2011). *PUE and DCiE Data Center Efficiency Measurement and Benchmarking*. Retrieved June 15, 2011, from 42U: <http://www.42u.com/measurement/pue-dcie.htm>
- Acadia Optronics LLC. (2008). *High Performance Storage*. Retrieved July 07, 2011, from Acadia Optronics: <http://www.acadiaoptronics.com/highperformance.php>
- Bennett, H. (2011). *Disc Longevity, Handling, Storage and Disposal*. Retrieved July 21, 2011, from Hugh's News: <http://www.hughnews.ca/faqs/authoritative-blu-ray-disc-bd-faq/13-disc-longevity-handling-storage-and-disposal>
- Cassella, D. (2009, November 30). *Digital Trends*. Retrieved July 19, 2011, from Experts Say Blu-ray Players are Here to Stay: <http://www.digitaltrends.com/home-theater/experts-say-sony-blu-ray-players-dvds-are-here-to-stay/>
- DataVision. (2011). *Intel 160GB X25M MLC SSD*. Retrieved July 22, 2011, from DataVision: <http://www.datavis.com/cgi-bin/product.cgi?prfnbr=559768&site=GOGL>
- GE. (2011, July 21). *GE Micro-Holographic Disc at 500G*. Retrieved July 27, 2011, from Storage Newsletter: <http://www.storagenewsletter.com/news/optical/ge-micro-holographic>
- Gibson, G., & Schroeder, B. (2007). Disk failures in the real world:. In C. S. Department (Ed.), *5th USENIX Conference on File and Storage Technologies* (pp. 1-16). San Jose, CA: Carnegie Mellon University.
- Gibson, S. (2007, December 13). Security Now, Transcript of Episode #122. (L. Laporte, Interviewer)
- HP. (2011). *HP*. Retrieved July 22, 2011, from HP Storage MSL2024 1 LTO-5 Ultrium 3280 FC Tape Library: <http://h71016.www7.hp.com/MiddleFrame.asp?page=config&ProductLineId=450&FamilyId=2297&BaselId=32979&oi=E9CED&BEID=19701&SBLID=>
- IRS. (2008, April 8). *How long should I keep records?* Retrieved July 22, 2011, from Internal Revenue Service: <http://www.irs.gov/businesses/small/article/0,,id=98513,00.html>
- Panasonic. (2011). *Blu-ray Disc Global*. Retrieved July 22, 2011, from Panasonic: http://panasonic.net/avc/blu-ray_disc/biz_ideal_media.html#operation
- Panasonic. (2011). *Longevity*. Retrieved July 22, 2011, from Panasonic: http://panasonic.net/avc/blu-ray_disc/biz_ideal_media.html
- Perpetual Storage Inc. (2011, February 15). *Storing data on archived back-up tapes may keep data out of court*. Retrieved July 19, 2011, from Perpetual Storage Inc: <http://perpetualstorage.com/storing-data-on-archived-back-up-tapes-may-keep-data-out-of-court-100061>
- Provantage. (2011). *Seagate ST33000650NS*. Retrieved July 22, 2011, from Provantage: <http://www.provantage.com/seagate-st33000650ns~7SEGS282.htm>
- Roos, A., & Schmid, P. (2008, June 27). *The SSD Power Consumption Hoax*. Retrieved July 20, 2011, from Tom's Hardware: <http://www.tomshardware.com/reviews/ssd-hdd-battery,1955.html>
- Run Tech Media. (2011). *10 Pack Optical Quantum 6X 50GB BD-R DL Blu-Ray Blank Disc*. Retrieved July 22, 2011, from Pronto: http://www.pronto.com/product/10-pack-optical-quantum-6x-p_1923974944
- Scaramella, J. (June 2010). *Improving Power and Cooling Efficiency in the Datacenter*. Framingham, MA: IDC.

THE LANDSCAPE OF DATA STORAGE

Seagate. (2011, March). *Constellation ES.2*. Retrieved July 22, 2011, from Seagate Data Sheet:
http://www.seagate.com/docs/pdf/en-GB/datasheet/disc/ds_constellation_es_2.pdf

Steege, P. (2008, November 11). *Four reasons SSD fits in the Enterprise first*. Retrieved July 19, 2011, from Seagate:
<http://storageeffect.media.seagate.com/2008/11/storage-effect/four-reasons-ssd-fits-in-the-enterprise-first/>

Tokar, L. (2010, October 14). *Trimming SSD Performance Degradation*. Retrieved July 21, 2011, from The SSD Review:
<http://thesdreview.com/ssd-guides/optimization-guides/ssd-performance-loss-and-its-solution/>

Van Zelst, L. (2008, March). *Library of Congress*. Retrieved July 20, 2011, from Magnetic Tape "Sticky Shed" Research: Characterization, Diagnosis, and Treatment : http://www.loc.gov/preservation/scientists/projects/sticky_shed.html

Ward, D. (2011, May 02). *Audio Video Revolution*. Retrieved July 19, 2011, from Analysts Say 2011 Will be Strong for Blu-ray and 3D : <http://www.avrev.com/home-theater-news/blu-ray-software-news/analysts-say-2011-will-be-strong-for-blu-ray-and-3d-2.html>

Western Digital. (2009). *NAND Evolution and its Effects on Solid State Drive (SSD) Useable Life*. Retrieved July 21, 2011, from Western Digital: http://www.wdc.com/WDPProducts/SSD/whitepapers/en/NAND_Evolution_0812.pdf

Images Cited

HDD Photo: (Acadia Optronics LLC, 2008)

Tape Photo: (Perpetual Storage Inc, 2011)

SSD Photo: (Steege, 2008)

Blu-ray Photo: (Ward, 2011), (Cassella, 2009)

Figure B: (Panasonic, 2011)

Notes of Importance

- All costs and capacity were found by using RAW data.
- Information about HDD is based on the Seagate Constellation ES.2 ST3300650NS 3 TB Internal Hard Drive
 - Cost per TB was found by taking the \$0.1029/GB and multiplying out by 1024 and rounded up to the nearest ten. Prices based on July 2011. (Seagate, 2011) (Provantage, 2011)
- Information about tape is based on the LTO-5 HP MSL2024.
 - Capacity is based on the max of 3TB per system and cost per TB was found by taking the \$0.0274/GB and then multiplied by 1024 and rounded up to the nearest ten. Prices based on July 2011. (HP, 2011)
- Information on the SSD is based on the more inexpensive version by choosing Multi-Level Cell and on the Intel X25M SATA Solid State Drive
 - Cost per TB was found by taking \$2.79/GB and multiplying out by 1024 and rounded up to the nearest ten. Prices based on July 2011. (DataVision, 2011) Voltage is 5.25 max, active power consumption is 0.15 watts, idle is 0.06 w.
- Information on the Blu-ray is based on Quantum 6X WORM 50GB Blu-ray discs
 - Cost per TB was found by taking \$0.1039/GB and multiplying out by 1024 and rounded up to the nearest ten. Prices based on July 2010. (Run Tech Media, 2011)