Comparison of Optical versus Tape for Archive

September 11, 2014

Instrumental, Inc.
1450 Energy Park Drive
Suite 375
St. Paul, MN 55108
651-280-4800
http://www.instrumental.com

Contents

Table of Tables

Notice

Product names mentioned in this document may be trademarks and/or registered trademarks of their respective companies and are the property of these companies. Information referenced in footnotes is accurate as of the date of this paper, but may be changed by the respective vendor or author at a later date.

Executive Summary

Tape and optical storage have historically addressed different parts of the archive market, where tape addresses the high end of archive and optical the very low end of the archive market. With the demise of the majority of the optical market with most laptop and desktop computers no longer having optical drives, the market for optical is rapidly shrinking. This shrinking market is in contrast to other developments in optical storage including the fact that there has been a significant increase in optical density with Blu-ray discs and at least one optical vendor, Hitachi, is stating that they have a path to holographic optical storage. As a result of these developments, the optical vendors see an opportunity to open a new optical archive market, but even with the improved density, tape still has significant advantages.

Current Optical Disk Archive Market

As the traditional audio/visual markets for optical storage have become saturated and online streaming services are replacing discs, optical disc manufacturers are turning toward the long-term archive market to boost sales and have settled on specific Blu-ray makes and models.

- Blu-ray Disc Recordable (BD-R) media provides 25 GB to 50 GB of capacity, in a Write-Once, Read-Many (WORM) format.
- BD-XL discs are multi-layered (usually 3 to 4 layers) discs that can provide between 100 GB and 128 GB of capacity per disc in a WORM format.
- BD-DSD discs are new, two-sided, multi-layer discs and can provide 200 GB of capacity per disc in a WORM format.
- MDisc BD-R is a specialized product from Millenniata that increases reliability by etching actual pits into a solid surface rather than changing a dye or metal.

While optical media, such as Blu-ray discs, may be resistant to environmental factors such as humidity, dust, and temperature, they are generally less reliable than magnetic media like hard drives or tape when considering hard error rates[1] and especially silent data corruption[2] due to the consumer technology interfaces used (SATA and USB), both of which are significantly more susceptible to silent data corruption than are SAS or fibre channel (FC). While it would be possible to provide SAS or fibre channel interfaces to BD, this is unlikely and could not happen in the near future as it would require significant and expensive changes to the drives and to the formats of the discs. In addition, it is assumed that current trends in backward compatibility of optical drive technologies will mean that optical drives of the future will be able to read discs written today; however, new advances may not guarantee such backward compatibility.

As a result of these two factors, it is generally assumed that current Blu-ray media will be used as off-line, low availability storage for archives that have long lifespans.

---

[1] A permanent, unrecoverable media error such as a disk read error. Contrast with soft error which is recoverable

[2] Refers to errors in computer data that occur during writing, reading, storage, transmission, or processing, which introduce unintended changes to the original data which is not detected by the system

Challenges for Optical Storage in Archiving

To be recognized as a viable archive storage option, Blu-ray and other optical media face several challenges.

First, Blu-ray has very limited read and write speeds for archiving. The maximum read and write speeds for Blu-ray discs and drives is 16x base Constant Angular Velocity (CAV). With a base "1x" CAV that results in about 4.5 MB/s , a maximum of 16x CAV translates into approximately 72 MB/s maximum for outer cylinders, but the disc must spin at lower linear velocity on parts of the data groove closer to the center.[3] In addition, read and write rates for actual archive appliances are much slower at between 2x and 6x, or 9 MB/s to 27 MB/s[4]. When compared to tape drives that run at about 160 MB/s to 250 MB/s consistently, optical drives have extremely slow transfer rates.

Second, even advanced Blu-ray discs continue to lack the raw storage capacity of other media options. Blu-ray discs being developed now have considerably less capacity than tape. The new BD-DSD standard has only 200 GB. More advanced developments include Sony and Panasonic's Archival Disc, which is projected to have a capacity of about 300 GB to 1 TB in the future[5], and Pioneer's next generation high-capacity disc, which is projected to have 512 GB per disc[6]. This is impressive for Blu-ray technology, but it is not comparable to tape, which can currently store 2.5 TB of data per tape cartridge for LTO tape and more for enterprise formats (4 TB for IBM, 8.5 TB for Oracle).

Third, Blu-ray and optical storage lack scalability or variety of options. Storage infrastructure for Blu-ray and other optical media archives is still relatively new and only a few viable appliances exist as configurations develop. Smaller archive options use archival Blu-ray products and focus on a drive and a set of discs that data is stored on. The discs are then stored separately from the drive, similar to traditional practices for consumer optical media. However, this is usually only designed to store low numbers of TB of data, at most, and is unsuitable for large scale enterprise archive solutions. More advanced storage options are generally stand-alone appliances from third-party vendors using what vendors call enterprise-grade Blu-ray discs and drives from the primary manufacturers like Sony, Panasonic, Pioneer, and Hitachi. These "Jukebox" appliances can store up to hundreds of TB in modular libraries, although it is not specified which discs are used[7]. Some of these options can come with a few modules or compartments of hard drives to handle online storage and connection to clients while the Blu-ray discs serve as near-line or off-line archive storage. Hitachi-LG provides one of the few actual enterprise-scale appliances for optical storage with its HL100 optical archiving library system[8]. The system is built around library units that have 50 TB of raw capacity. The whole system is scalable up to 9 libraries per

---

[3] Blu-ray Disc; http://en.wikipedia.org/wiki/Blu-ray_Disc

[4] Blu-ray FAQ; http://www.blu-ray.com/faq/#bluray_speed

[5] "Archival Disc" standard formulated for professional-use next-generation optical discs; http://www.sony.net/SonyInfo/News/Press/201403/14-0310E/index.html

[6] Pioneer Jointly Develops Next-Generation High-Capacity Optical Disc for Data Archiving with Single-Sided Capacity of 256GB; http://pioneer.jp/press-e/index/1758

[7] Review of the Current Jukebox Library Technology; http://www.imakenews.com/kin2/e_article001411628.cfm?x=b11,0,w

[8] HLDS Launches Blu-ray Optical Archiving Library System; http://www.hughsnews.ca/hlds-launches-blu-ray-optical-archiving-library-system-0046566

42U rack for up to 450 TB of raw capacity. The claimed goal of Hitachi-LG's system, however, is to only provide reliable cold storage that can last up to (as yet unproven) 50 years with new optical media replacement discs, redundant robotics and advanced error correction coding[9]. It is still unsuitable for more active applications, especially with large files, as Blu-ray performance is very low. Although using it as a form of cold storage is a viable option, in which the disc is written to and then left off-line, more active archives present a problem for Blu-ray.

Fourth, reliability tends to present a problem when considering Blu-ray as an archiving solution. Since the last standard for optical storage only requires reliability comparable to computer connections, most estimates place Blu-ray hard error rates at around 1E12[10], which means that on average one in every one trillion bits written to or read from the device will be in error. This means that an error will occur on the Blu-ray for every TiB transferred (see Table 1). In addition, the USB or SATA connections used also do not provide good error correction in the channel, which can lead to silent data corruption. In comparison, tapes use fibre channel or SAS and can provide hard error rates of 1E17 for LTO, or 1E19 or greater for enterprise tapes. The only publicly published exception here is Hitachi-LG's optical archive library, which incorporates added measures, error correction and error detection to achieve a theoretical hard error rate of 8.1E23[11]. It is believed that other specialized archive solutions use similar error correction code (ECC) methods to achieve reliability higher than the rated media reliability. The following is the silent data corruption rate for various technologies using the assumption that the channel is operating within its design specification with no errors. The table shows the number of silent data corruptions per year for the channel if data is moving at the specified rate for 24 hours per day x 365 days per year. For example, for FC and SAS, it is expected that there will be an average of 2.7 silent data corruptions per year at a transfer rate of 10 TB per second.

		Sustain Transfer Rate Per Second for a Year						
Notes	**SDC Rate**	**0.5 GB/sec**	**1 GB/sec**	**10 GB/sec**	**100 GB/sec**	**1 TB/sec**	**10 TB/sec**	**100 TB/sec**
Detection with T10 PI	10E28	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	10E27	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	10E26	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	10E25	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	10E24	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	10E23	0.0	0.0	0.0	0.0	0.0	0.0	0.3
	10E22	0.0	0.0	0.0	0.0	0.0	0.3	2.7
FC/SAS	10E21	0.0	0.0	0.0	0.0	0.3	2.7	27.1
	10E20	0.0	0.0	0.0	0.3	2.7	27.1	270.9
	10E19	0.0	0.0	0.3	2.7	27.1	270.9	2708.9
	10E18	0.1	0.3	2.7	27.1	270.9	2708.9	27089.2
SATA	10E17	1.4	2.7	27.1	270.9	2708.9	27089.2	270892.2
	10E16	13.5	27.1	270.9	2708.9	27089.2	270892.2	2708921.8
	10E15	135.4	270.9	2708.9	27089.2	270892.2	2708921.8	27089217.7
	10E14	135.4	2708.9	27089.2	270892.2	2708921.8	27089217.7	270892177.3
	10E13	135.4	27089.2	270892.2	2708921.8	27089217.7	270892177.3	2708921772.9
	10E12	135.4	270892.2	2708921.8	27089217.7	270892177.3	2708921772.9	27089217729.3
USB est	10E11	135.4	2708921.8	27089217.7	270892177.3	2708921772.9	27089217729.3	270892177293.3

Table 1 - Silent Data Corruption Rate for Various Interfaces

[9] Handling errors when reading data over the next 50 years; http://www.hitachi.com/rd/portal/story/bd_archive/04.html
[10] UJ-235A Blu-ray Combo; http://remedio.tv/diversos/pctita_arquivos/PanasonicSlimBluRayBurner.htm
[11] Introduction of Optical Archiving Library System for Long-term Data Retention; http://hlds.co.kr/v2/HLDS_BD_140212_V1.pdf

Magnetic Tape Archive Options vs. Optical

The following sections examine the differences between Linear Tape Open (LTO), Enterprise tapes and optical and why enterprise tapes should be utilized for large archives.

LTO and Enterprise Tape vs. Optical

There are advantages of tape over both optical disc and hard disk due to the higher reliability of tape for archival data. There are also significant differences between enterprise tape and LTO tape for archival solutions. Large archives require a good understanding of not just the short-term technology and cost issues, but the benefits of a long-term analysis of technology trends, and present and future value costs.

As part of this process, some of the technology issues are examined when making the determination of what technology provides the best value for larger archives. There are three different tape media choices for large archives in the market today. They are:

1. IBM TS1140 enterprise tape
2. Oracle T1000D enterprise tape
3. LTO tape from multiple suppliers

LTO and enterprise tapes have different characteristics and benefits. It should be noted that IBM enterprise tape has been around for almost 30 years while LTO has been a standard for almost 15 years. Likewise, the Oracle tape technology, derived from the long-standing STK products, has a long and rich history. Technology refresh and media reuse are a significant issue for the long-term cost of any archive. As an example, the support for the hardware interface to the tape drive such as the fibre channel interface has about a 7-10 year life before the technology reaches end of life (EOL), although fibre channel is designed to be backward compatible. In a critical archive, any technology should not be reaching EOL and migration to new technology must be completed before EOL. Although the tape media might be rated to last 30 years, there are a large number of issues in other parts of the system and the operational environment that reduce the 30 year number to less than a decade. By comparison, Blu-ray discs may last 50 years, but changes to increase capacity, such as new Archival Disc and holographic technologies, mean that Blu-ray may not have backward compatibility in drive technology and will likewise have practical lifespans shortened to about a decade. This is because the drives for these new technologies will have difficulty reading BD-R, BD-XL, or BD-DSD discs. If, however, vendors choose to stay with Blu-ray, backward compatibility has been proven to last at least 30 years. However, Blu-ray will still have limitations on performance and capacity well into the future.

Media

The preferred tape media is Barium Ferrite (BaFe). Both IBM and Oracle utilize BaFe in their enterprise tapes. The LTO Consortium has also committed to BaFe media, but they have yet to standardize on BaFe media, which is currently optional at LTO-6 and is expected to be the standard for LTO-7 and beyond.

Storage Capacity

Storage capacities of enterprise tapes are significantly higher than those for current LTO tapes, which impacts the costs for a tape archive. Future capacity increases for the IBM TS1150 and the Oracle T10000D follow-on are expected to further increase this disparity. Tape capacities, which range from 1.5 TB to 8.5 TB, far surpass Blu-ray's 25 GB to 128 GB capacity standards. One other key point is that tape media offers much more capability for density improvements than that provided by BD. In fact just recently Fujifilm and IBM announced a new record in areal density of 85.9 billion bits per square inch on linear magnetic tape. This equates to a potential of 154 TB on a standard LTO data cartridge, 62 times greater than current 2.5 TB LTO-6. Contrast this to BD-XL and BD-DSD capacity growth potential which is currently limited to a 10X improvement while the technical feasibility of this is questionable for the following reasons:

1. Laser wavelength – Blu-ray discs use a laser with a wavelength of 405 nm. This means that pits, and data, must remain at a width readable by this laser. An individual pit, or bit, requires a space 130 nm wide by 150 nm long.
2. Disc size restrictions – Standard Blu-ray discs have a set diameter of 12 cm, and any significant change means they cannot fit into optical drives.
3. Layer interference – Based on the previous restrictions, the current solution is to increase the number of data layers available on a disc. However, each new layer adds some level of interference as the laser has to travel farther, focus at different levels, and even read through layers of data. This limits the number of layers practical for a Blu-ray to about ten layers before problems occur in reading and writing data. For example, the new BD-DSD has tried to increase capacity of the BD-XL standard, but could only do so by making the disc double-sided instead of adding layers.

Performance

Performance of the IBM and Oracle enterprise tapes is similar at about 240 MiB/s, but both are significantly higher than the performance for LTO tapes at 153 MiB/s. Further, LTO-6 is still nearly ten times faster than archival Blu-ray options, which read data at about 17 MiB/s, at best.

Error Rates

In addition to performance, capacity and cost differences, there are also reliability differences between LTO and enterprise tape technology; however, both are significantly higher than most Blu-ray options. The hard error rate, which defines how many bits of data can be read before a read fails, for enterprise tape storage is at least 2 orders of magnitude better than that for even enterprise disk storage. The table below displays the hard error rates of different types of storage devices. As can be seen, reliability increases from optical storage options such as CDROM, Blu-ray, and DVD to SATA consumer hard disk drives toward enterprise tape. It takes about 0.01 pebibyte (1024*1024*1024*1024*1024 bytes = PiB) of data for consumer-level SATA disks to experience errors while tape drives such as LTO and enterprise tape can transfer multiple PiB before errors occur. For reliability, the best solution is enterprise tape with a hard error rate of 1E19 or greater, which translates into moving over 1,000 PiB of data before an error occurs in the drive. This is 1,000 times better than the best hard disks, 100 times better than LTO tape, and 10,000,000 times better than standard Blu-ray. Despite this vast improvement, with archives of multiple PiB, even enterprise tape will experience errors. It should be noted that the higher reliability of enterprise tape does not address the silent data corruption possible on a communications channel. This is especially important considering that Blu-ray typically uses USB 3 or SATA interfaces, which are much less reliable than the SAS or fibre channels used by tape.

One other point is that the errors as described in the table below are at the hardware level. For certain applications, such as the Hitachi Blu-ray disc archive, it is possible that Hitachi rewrote the driver to do forward error correction, and in that way, improved the reliability of the archive and Blu-ray discs up to an error rate of 8E23.

Device	Hard Error Rate (1 bit in error in this number of bits moved)	PiB Equivalent Data Moved Before Error
Data / Communications Channel / CD / Blu-ray	1.0E+12	0.0001
DVD	1.0E+15	0.11
Consumer SATA	1.0E+14	0.01
Enterprise SATA	1.0E+15	0.11
Enterprise SAS/FC	1.0E+16	1.11
LTO	1.0E+17	11.10
Oracle T10000A/B/C/D	1.0E+19	1110.22
IBM TS1140	1.0E+20	11102.23

Table 2 - Hard Error Rates of Different Media

Conclusions

While the optical market today for archive is growing, it is small with mostly traditional consumer electronics companies backing it. The following table is a comparison between tape and optical technologies:

Storage Media	TS1140	T10000D	LTO-6	Current Blu-ray (HLDS HL100)	Current Blu-ray (DISC ArXtor 7000)	Current Blu-ray (Open Source Storage Blu-ray Archive)
Introduction	June 2011	September 2013	February 2013	April 2014	November 2012	?
Capacity in 1000s	4TB	8.5TB	2.5TB	100GB	50GB	50GB
Interface	FC-8	FC-16/FCoE 40 Gb	6 Gb/s SAS / FC-8	SATA	SATA	SATA
Hard Error Rate	1.00E+20	1.00E+19	1.00E+17	8.10E+23	~1.00E12?	~1.00E12?
Performance						
Read Data Rate (MiB/s)	238	246	153	17.16	8.58	8.58
Write Data Rate (MiB/s)	238	246	153	8.58	8.58	8.58
Pick Time (s) in Automation (Robot/Library)	5	11	5	8[i]	6	8
Load Time (s) for Media in Device	16	13	16	6[i]	6	6
Average File Access (s)	47	50	50	80[iii]	66[ii]	66[iv]
Time (hr) to Read 1GiB	0.020	0.022	0.022	0.043	0.055	0.055
Time (hr) to Read 1TiB	1.24	1.20	1.92	17.00	33.97	33.97
Time (hr) to Read 1PiB	1253.2	1212.5	1949.4	17381.2	34762.4	34762.4
Time (hr) to Write 1GiB	0.020	0.022	0.022	0.059	0.055	0.055
Time (hr) to Write 1TiB	1.24	1.20	1.92	33.97	33.97	33.97
Time (hr) to Write 1PiB	1253.2	1212.5	1949.4	34762.5	34762.4	34762.4
Environmentals						
Drive Operation	60°F to 90°F (+16°C to +32°C) Humidity 20% to 80% non-condensing	60°F to 90°F (+16°C to +32°C) Humidity 20% to 80% non-condensing	50°F to 104°F (+10°C to +40°C) Humidity 10% to 80% non-condensing	50°F to 95°F (+10°C to +35°C) Humidity 10% to 80% non-condensing	50°F to 104°F (+10°C to +40°C) Humidity 10% to 75% non-condensing	50°F to 95°F (+10°C to +35°C) Humidity 5% to 95% non-condensing
Long-term Media Storage	61°F to 77°F (+16°C to +25°C) Humidity 20% to 80% non-condensing	59°F to 79°F (+15°C to +26°C) Humidity 10% to 95% non-condensing	61°F to 95°F (+16°C to +35°C) Humidity 20% to 80% non-condensing	-40°F to 116°F (-40°C to +47°C) Humidity 10% to 80% non-condensing	-40°F to 116°F (-40°C to +47°C) Humidity 10% to 75% non-condensing	-40°F to 116°F (-40°C to +47°C) Humidity 5% to 95% non-condensing

Notes:

[i] Pick and load time are unavailable but assumed to be similar to other Blu-ray models.

[ii] While the vendor publishes a number of 0.17, it assumes access from a hard disk cache. The average file access time (66 s) for similar models is used.

[iii] This average file access time is Hitachi's published typical access time (80 s) (see http://hlds.co.jp/v2/HLDS_BD_140212_V1.pdf) but it seems excessive in Instrumental's opinion.

[iv] Average file access time is unavailable but assumed to be similar to other Blu-ray models.

Table 3 - Tape and Optical Comparison

This table demonstrates several advantages for tape technology in archiving. These advantages of tape technology include the following:

1. Performance - Even the slower LTO drives have a significant advantage over optical discs, particularly when writing data.
2. Reliability - Enterprise tapes in particular are several orders of magnitude more reliable than optical discs. The only instance where optical discs could match tape is Hitachi. However, this technology is new, untested in production, and requires non-standard error correction coding.
3. Capacity - The latest enterprise tape has almost a 200x higher capacity than standard optical discs. Even considering the new BD-DSD standard, current enterprise tape still has as much as a 42x higher capacity.
4. HSM software - There are numerous HSM packages with a long history of tape support, but for optical technology, the selection of packages is limited and not heavily used from what our research showed.
5. Interface technology - Tape drives utilize more reliable fibre channel or FC over Ethernet (FCoE), or possibly SAS interfaces for LTO tape drives, instead of the USB or SATA interfaces used by optical drives (see Table 2 for a comparison of hard error rates).

Clearly tape has many advantages compared to optical, but what is clear from this chart is that optical also has some advantages over tape. These include:

1. Tape pick and position time are much slower than optical, which means that for small files optical has an advantage.
2. Slightly better power usage during hibernation as optical drives can be placed in cold storage if they are not accessed, requiring no power.

This comparison shows that tape has a longer history and better developed archive support and architecture, provides a wider range of products, is a more reliable platform for active archives, offers greater scalability, and has significantly higher performance than Blu-ray and optical archive products available today or in the foreseeable future. While optical storage may provide some slight edges in power usage for cold storage and in reading small files, the current media and hardware cannot provide sufficient performance, reliability, availability, or scalability for large or enterprise-level data archives without significant internal support and customization. This means that tape is still the best option for data archives.

Appendix A - Acronyms

BaFe - Barium Ferrite
BD - Blu-ray Disc
BD-R - Blu-ray Disc Recordable
CAV - Constant Angular Velocity
CD - Compact Disc
DVD - Digital Video Disc
ECC - Error Correcting Code (Reed Solomon being 1 example)
EOL - End Of Life
FC - Fibre Channel
FCoE - Fibre Channel over Ethernet
GB - Gigabytes = 10^9 bytes = 1000 MB
GiB - Gibibyte = 2^{30} (1,073,741,824) bytes = 1024 MiB
HDD - Hard Disk Drive
HSM - Hierarchical Storage Management
IB - InfiniBand
INCITS - International Committee for Information Technology Standards
LTFS - Linear Tape File System
LTFS-EE - LTFS-Enterprise Edition
LTO - Linear Tape Open
MB - Megabytes = 10^6 bytes = 1000 KB
Mb/s - Megabits per second
MB/s - Megabytes per second
MiB - Mebibyte = 2^{20} (1,048,576) bytes
PB - Petabytes = 10^{15} bytes = 1000 TB
PiB - Pebibyte = 2^{50} (1,125,899,906,842,624) bytes = 1024 TiB
OS - Operating System
SAS - Serial Attached SCSI
SATA - Serial Advanced Technology Attachment
SCSI - Small Computer System Interface
TB - Terabytes = 10^{12} bytes = 1000 GB
TiB - Tebibyte = 2^{40} (1,099,511,627,776) bytes = 1024 GiB
T10 - Technical Committee of INCITS
T10 DIF - T10 Data Integrity Field
T10 PI - T10 Protection Information
USB - Universal Serial Bus
WORM - Write-Once Read-Many