INSIC
Information Storage
Industry Consortium

Global Trends, Applications and Use Cases for Tape Adoption

INSIC REPORT | 2024

insic.org

2024 "Global Trends, Applications and Use Cases for Tape Adoption"

Team Director

Shawn Brume - IBM

Contributors

Carlos Sandoval	IBM
Mark Hill	IBM
Turguy Goker	Quantum
Ryan Duffy	Quantum
Robert Biskeborn	Western Digital Corporation
Laura Loredo	Hewlett Packard Enterprise
Andrew Barton	Hewlett Packard Enterprise
Andrew Dodd	Hewlett Packard Enterprise
Matt Ninesling	Spectra Logic
Scott Hamilton	Western Digital Corporation

Quoted Industry Experts

Tom Coughlin	Coughlin Associates
Jim Bain Ph.D.	Carnegie-Mellon
Todd Heer	Lawrence Livermore Laboratory
Steffen Hellmold	Pithos LLC
John Monroe	Furthur Market Research
Philippe Nicolas	Coldago Research
Jay Bartlett	Cozaint
Brad Johns	Brad Johns Consulting
Jason Hick	LANL
Emily LeProust	CEO Twist Bioscience

Contents

1.0 TEAM DIRECTOR AND CONTRIBUTORS 2

1.1 EXECUTIVE SUMMARY 5

1.1.1 Notable quotes from INSIC interviews 6

1.2 MARKET TRENDS 7

1.2.1 Green Data Center 7

1.2.2 Data Explosion 7

1.2.3 Cold Data Storage 8

1.2.4 AI / Data Analytics 8

1.2.5 Video Surveillance 9

1.2.6 Cloud 9

1.2.7 Longer Retention Times 10

1.2.8 Security and Cyber Resiliency 11

1.3 TECHNOLOGY TRENDS 12

1.3.1 Competing Technologies 12

1.3.2 Unstructured Data Integration 13

1.3.3 UBER (uncorrectable bit error rate) 14

1.3.4 RAIL with Erasure Coded Tapes 14

1.3.5 Host interface 15

1.3.6 Post-Quantum Cryptography 16

1.4 COST & ENERGY CONSUMPTION ADVANTAGES 17

1.4.1 TCO Operational Planning 17

1.4.2 Modeling Methodology 17

1.4.3 Energy Consumption 19

1.4.4 Trends in Storage Costs 20

1.4.5 TCO Conclusions 20

1.5 THROUGHPUT AND RELIABILITY 21

1.5.1 Data Rate 21

1.5.2 Durability of the Data 22

1.5.3 Media Wear Out 23

Continued...

Contents

1.6 CHALLENGES FOR TAPE ADOPTION 24

1.7 VERTICAL INDUSTRY USE CASES FOR TAPE 26

1.7.1 Financial Service 26

1.7.2 Life sciences and Healthcare 26

1.7.3 Oil and Gas, Geo physical, and Geological Research 26

1.7.4 Internet of Things and Manufacturing 26

1.7.5 Video Surveillance 26

1.7.6 Retail and Transaction Processing 26

1.7.7 Media, Broadcasting, News & Publishing, Entertainment 27

1.7.8 Scientific Research and High-Performance Computing 27

1.8 BENEFITS OF TAPE ACROSS IT WORKLOADS 28

1.8.1 Big Data / Analytics 28

1.8.2 Archive 28

1.8.3 Object Storage Active Archives and Tape 28

1.8.4 Data Backup and Recovery 29

1.8.5 Tape and Primary Storage 29

1.8.6 Tape and Secondary Storage 29

1.8.7 Tape and Cloud-based Data Protection and Backup 29

1.8.8 Disaster Recovery 29

1.8.9 Ransomware and Cyberthreats 29

1.8.10 Bulk Transfer 30

1.8.11 Data Regulations & Compliance 30

1.8.12 Tape Encryption and WORM immutability 30

1.9 SUSTAINABILITY 31

1.9.1 Tape Lifecycle 31

1.9.2 Data Destruction 31

1.9.3 Media Disposal 32

1.9.4 Carbon impact 32

1.10 CONCLUSIONS AND RECOMMENDATIONS 33

ACKNOWLEDGEMENT 35

REFERENCES 36

1.1

Executive Summary

The Information Storage Industry Consortium (INSIC) 2024 roadmap has been significantly revised over the previous INSIC roadmap. This document was created in a collaborative fashion by many dedicated individuals, and represents their expertise and data gathered from interviews of analysts and experts in the field. At the end of this executive summary, we have included some notable quotes that summarize the themes of the 2024 roadmap.

The Key Points in the Roadmap Include:

- **Compared to hard disk, the current track density of tape is lower by a factor of 50 to 100, providing the greatest potential leverage for advancing tape technology and growing the capacity of tape cartridges.**

- **Data creation and retention forecasts far exceed the production capability of HDD and SSD.** Storage data placement is growing in importance to meet price and demand, tape is the most viable technology for the near future.

- **This roadmap has an increased coverage of carbon impact and energy.**

- **The roadmap has increased coverage of the Uncorrectable Bit Error Rate of tape.** Providing "Durability" comparisons enabling easier comparisons to other storage technologies.

- **2030 tape production capacity capability will exceed 2 zettabytes.**

The fundamental advantages of tape continue to be its low acquisition cost, extremely low power consumption, excellent footprint density, scalability, and of course, reliability. The inherent air-gap design of tape makes it the choice for offline data protection and portability, providing significant protection from cyberattacks. Tape's future is driven by the needs of multiple markets and use cases where these attributes are important.

Tape is the best solution for infrequently accessed data stores. Tape is well positioned for each of those markets today and for the future. Including Software solutions that obfuscate tape from the end user.

This roadmap can only be a snapshot in time. For continuous news on the tape industry, visit places such as the Blogbytes at the LTO consortium website: https://www.lto.org/blogbytes/

Interviews with analysts and experts were conducted as research for this roadmap.

All of the bugs and practicalities of access, scale, manufacturing, and supply chain have been worked out for the tape industry and it's [now] really about advancing the technology. I see some solutions for the problems in the next 5-10 years, whereas pore sequencing for DNA or [similar technical challenges] have a huge number of issues to make it a practical solution.

Jim Bain Ph.D. - Carnegie-Mellonn

1.1.1
Notable Quotes from INSIC Interviews

"I currently do not see tape being displaced for 10-15 years at a minimum. The biggest squeeze is on HDD..."

Todd Heer - Lawrence Livermore Laboratory

"DNA is one of the technologies that is a "future" candidate for deep preservation. Has funding today, but in early stages of commercialization, achieving enabling price points will be critical to realizing deployment at a meaningful scale."

Steffen Hellmold - Hellmold Consulting

"Tape's low energy is a huge positive in the storage industry, giving companies huge saving in cost and carbon impact. HDD and SSD are handling far too much cold data using massive amounts of energy."

John Monroe - Furthur Market Research

"Without tape, the industry will have a problem."

Philippe Nicolas - Coldago Research

1.2
MARKET TRENDS

1.2.1
Green Data Center

Power and cooling requirements and expenses have grown to be a substantial consideration in data center design. Data centers generate up to 2 percent of global CO2 emissions.[1]

Datacenters consume 1.4% of global energy, which is equivalent to the total electricity consumption of Germany and Japan combined. However, in common datacenter geographic locations, such as Ireland, datacenters consume as much as 18% of all energy production.[2]

In addition to the environmental cost and implications of this power consumption, there is a financial burden as well. Green Data Centers are designed from scratch to minimize power consumption and environmental impact. Modernization of data centers has reduced the rate of growth in energy consumption significantly. For a discussion on power consumption, please see:

Read more:
https://www.lto.org/wp-content/uploads/2014/06/Data-Storage-Technology-Tape-and-Sustainability-Sept-2017.pdf

Read more:
https://www.bradjohnsconsulting.com/_files/ugd/8b8555_78e13c6f9b454179af76877d67fbb9db.pdf

Tape has an inherent architectural advantage for inactive data at scale. Using removable media that consumes no power when not in use provides a huge power and cooling savings over other storage technology choices. Automated libraries of tape consume very little power to provide access to these removable tape cartridges on demand.

Tape's power use is significantly less than disk, especially for data that is inactive. One study by Brad Johns Consulting concluded that a disk archive consumes 96 percent more power than a tape-based solution. Tape's carbon footprint is also reduced, the same study showed a CO2e savings with tape of about 97 percent.[3]

1.2.2
Data Explosion

The digital data explosion that commenced with the advent of computing has shown no signs of slowing down. There are multiple drivers of this continued high growth, including conversion from analog storage to digital storage (as seen in industries like media and entertainment, video surveillance, IoT, cloud services, genomics, healthcare, etc.); the high growth rate in consumer-generated digital content (which is driving huge growth in cloud storage); data generated from the internet; and other sources of structured and unstructured data. Another major factor driving this growth is the need to keep digital data longer (in some cases, forever), which is sometimes driven by legal compliance requirements, but also driven by value in the older data, for example, in the case of analytics or scientific research.

The hyperscale effect is reshaping the IT industry as compute and storage are being scaled massively in architectures with redundant components that migrate workloads in case of failures. Hyperscale storage solutions often serve millions of users with a few applications. Amazon Web Services, Microsoft, and Google collectively control more than half of the WW cloud infrastructure service market. This is driving unprecedented energy consumption that requires advanced cooling systems, redundant power and lots of tape.

Some examples of exponential growth are:

- Nearly 20 percent of corporations store more than 1 PB.
- Each Connected car generates 450 TB of data per year.
- 1.7 trillion photos taken in 2023.
- Surveillance = 2.5 EB/day in 2019.
- The IoT is forecast to connect to over 22 billion devices by 2025.
- IoT has 6.4 million developers.
- The Square Kilometer Array will archive over 700 PB per year.
- Average Server Lifespan - 3 years; Hard Disk Drive (HDD) 4-5 years, Tape Drive 7-10 years.

According to Furthur Market Research the amount of stored data is exploding, and it has been and will continue to do so for years to come. Furthur estimates that the amount of stored data will exceed 43 ZB by 2035, as illustrated in figure 1.[4]

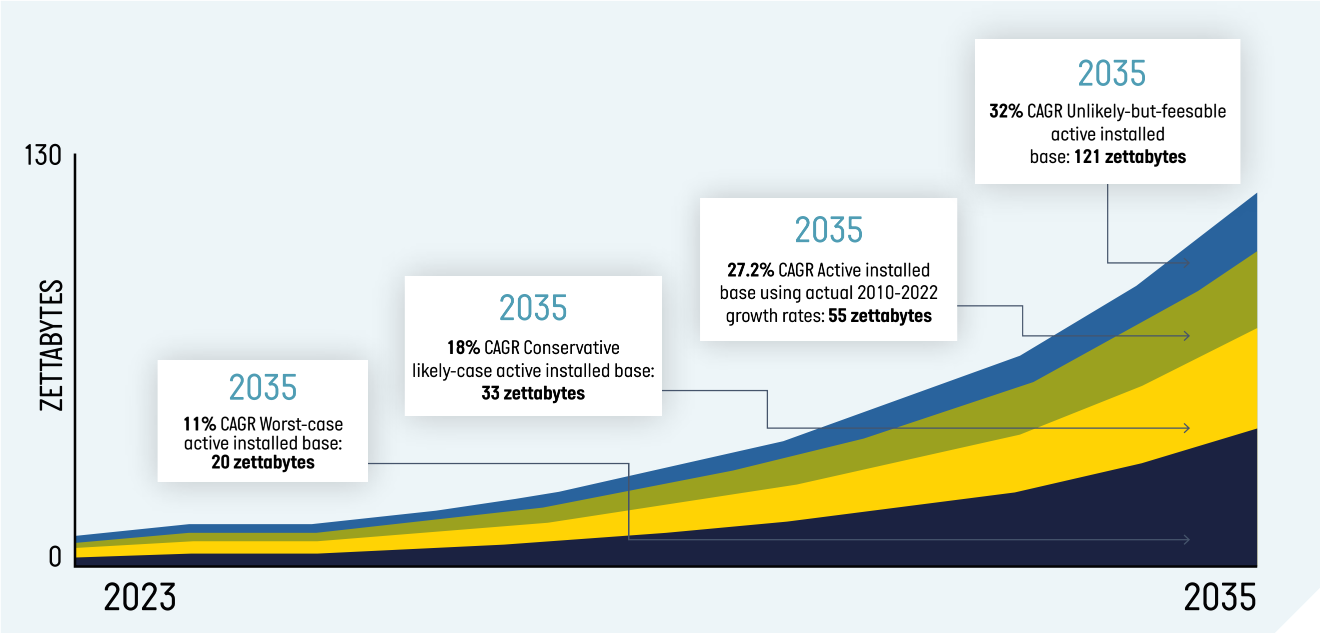

Figure 1. Global Active Storage Installed Base, 2023-2035
Source: Furthur Market Research

Multi-byte Unit Reference

$1x10^6$	$1x10^9$	$1x10^{12}$	$1x10^{15}$	$1x10^{18}$	$1x10^{21}$	$1x10^{24}$	$1x10^{27}$	$1x10^{30}$
Megabyte	Gigabyte	Terabyte	Petabyte	Exabyte	Zettabyte	Yottabyte	Ronnabyte	Quettabyte

Figure 2. Reference for data measurement
Source: INSIC

IDC predicts that by 2027, the Datasphere of data generated annually will reach a staggering 291.7 ZB, the overwhelming majority of which will be Enterprise, rather than consumer data.[6]

Of course, not all of the data that is created needs to be stored, but IDC's adjacent Storagesphere report presents a sobering picture for all those IT professionals tasked with managing and protecting their organization's data. IDC expects the total installed base of storage capacity will roughly double between 2023 and 2027, to reach just over 20 ZB at the end of this period.[6]

1.2.3
Cold Data Storage

Data growth – new data generation – is one part of the equation. The other driver of installed base capacity is the fact that we now live in an era of minimal data deletion. In a recent study by Further Research businesses expressed overwhelming preference for keeping all their data rather than try to selectively curate their data archives.[7]

What is remarkable about the immense data volumes anticipated by the two analysts above is that the majority of that growth is expected to be for cold and unstructured data.

According to Gartner, 80% of all data stored in the future will be unstructured. Between 60% and 70% is Enterprise content that is infrequently accessed but just needs retained for some essential purpose. Furthur Market Research estimates that as much as 35 ZB of cold storage will be needed by 2035.

"We project that stored cold data will grow from 3.1 zettabytes in 2023 to 25.9 zettabytes in 2035. This is an annual growth rate of over 19%."

Brad Johns

1.2.4
AI / Data Analytics

AI is driving virtually every aspect of business, research, and development, multi-petabyte archives are becoming organizational standards. Recent lawsuits over the use of copyrighted materials for the training of AI models, as well as defamation litigation responding to false information generated by AI chatbots, only highlight the need for this training data to be preserved for the long term.

With a seismic sea wave of data overwhelming IT infrastructures, and volumes swelling year over year, many organizations are reexamining their data storage strategies. Cloud is one platform for storing and analyzing these data sets. But, because after a model is completed much of the data does not need to be accessed instantly, tape is regarded as the most economical way to preserve this information for the long term. It can then be restored and used in a future learning model or kept to support possible litigation in the future. AI is here to stay and tape technology combines many key attributes to make it the optimal choice for storing a tsunami of data, including its high capacity, fast performance, low power consumption, 30-year archival life, scalability, extensive roadmap, reliability, and affordability.

1.2.5

Video Surveillance

The move from analog systems to IP cameras has enabled companies to expand the use of their security systems beyond pure monitoring into business intelligence and analytics systems used for competitive advantage. For example, video surveillance and analytics are now being used in retail to monitor and analyze customer behavior; transportation to monitor traffic patterns and delivery optimization; manufacturing to monitor processes and optimize inventory management and production; healthcare for managing access and retrieval of healthcare records; government for public safety, threat detection and avoidance; and the military for observation, planning, and operations.

The quality of the images has continued to increase. While High Definition video and 30 frames-per-second (fps) have generally been a standard, many applications that require the ability to capture faster moving objects coupled with finer details – e.g., license plate and facial recognition – are moving toward 4K resolution and 60 fps. Although specialized software may limit recording to when there is something in the camera's field of vision, data can accumulate very quickly to massive amounts. Even a small surveillance system can generate more data than many companies have in their entire IT environments. Data growth in video surveillance is also driven by legal requirements for longer retention periods, longer hours of operation, and the increased demand for security. According to Mike Jude, IDC research director for Video Surveillance and Vision Applications, "the video surveillance camera market will continue to grow at a double-digit pace." Consequently, the estimates for the average data generated daily by new surveillance cameras shipped globally has grown from over 2.5 exabytes in 2019 to over 5.5 exabytes in 2023.[4]

Tape technology with LTFS or AXF offers an open format that makes data accessible for active archive or long-term retention, while at the same time offering the lowest total cost of ownership with a limitless archive tier.

A solution with a NAS interface that incorporates tape for the archive tier can also provide an ideal solution for this industry since video data is normally generated as large files of streaming data. This solution integrates with existing advanced systems allowing data to be moved to tape in a seamless manner while reducing the cost of storage.

"The center of the universe for video surveillance users is the Video Management Software (VMS) where they look at live and recorded video. If any video surveillance storage solution is not integrated into the VMS, the user acceptance fails. LTO must be part and parcel of the VMS and then users will be open to such a solution."

Jay Bartlett, Cozaint

1.2.6

Cloud

The cloud boom was kicked off when Amazon introduced a suite of cloud-based services, Amazon Web Services, in 2002. Since that time Google, Microsoft and many others have joined the fray. The value proposition for cloud includes access to scalable and reliable IT infrastructures and applications without the capital outlay for on-premises servers, storage, networking equipment and labor.

In the early days of cloud, it sounded too good to be true, and many organizations were given a 'cloud first' mandate. What has occurred today is that after moving to the cloud many organizations are recognizing that there are cost and usage advantages to keep data both in the cloud and on premise. The majority of organizations are now combining some form of cloud with on-premises compute/storage for the most effective, economical and secure data protection.

Why are organizations choosing this hybrid approach? Because "The Cloud" isn't a single "thing" or even service. There are multiple approaches to engaging a remotely owned and managed data center (cloud) to assist in achieving organizational goals. The three cloud offerings most commonly referred to are Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). For digital storage, and specifically tape, these represent large pools of storage at centralized locations for which tape is ideal as a long-term data protection and archive tier. The data is created in the cloud and often stays in the cloud. The storage target is hidden from the user. If that storage target is tape, the user wouldn't know that. They would interact with the cloud without knowledge of the underlying data protection and archive scheme incorporated by the cloud provider. This is an excellent opportunity for tape to be implemented within the cloud by the actual cloud service provider vs. the end user.

There was also a large movement to a fourth form of cloud service often referred to simply as, "cloud storage." This involves moving data created outside of the cloud into the cloud. No compute, application or networking services are provided. The cloud simply offers a remote landing zone for secondary storage in the form of backup, disaster recovery or archive. This is the area where data that is stored in the cloud and used on a regular basis has a much higher long-term cost and where data repatriation back to on-premises storage is starting to happen.

Virtually every smart phone in use "backs up to the cloud." This is not only the perfect example of cloud storage as a service; it's the perfect use case for cloud storage. Storing small amounts of data in the cloud is straightforward. Large data sets, however, are dependent on how much internet bandwidth an organization has. And in almost no cases do organizations have access to 100 percent of their bandwidth to move data to/from the cloud. The bandwidth is needed for all other operations related to WAN activity as well. With a dedicated 1 Gbs connection, common in organizations today, a restoration of 500 terabytes would take over 45 days to achieve. That's with 100 percent of the bandwidth for the entire duration of the 45-day restoration. A single petabyte restoration would require over three months. Depending on the amount of data to be retrieved, a lengthy restore time could negatively affect the business.

By comparison, LTO-9 tape technology is over 3 times the performance of a dedicated 1 Gbs connection. A single LTO-9 tape drive, with a native transfer rate of 400MBs, would take the 1 PB restoration from 3 months down to 30 days. A bank of 10 LTO-9 drives would take the restoration down to three days. The LTO-9 drives require an initial capital outlay, but there's no monthly charge after that as there would be for an oversized WAN connection to increase performance.

Many organizations have determined that a hybrid approach to cloud makes good sense. A combination of cloud and tape can help to achieve best practices in digital preservation for the long-term by providing IT managers with three copies of data, on two different media types and one offsite. Organizations with on-premises tape also have control of their data and the option of abandoning cloud or switching cloud providers rather than paying exorbitant egress fees.

Combined with modern-day object storage interfaces for tape, a hybrid model of public/private cloud storage has become equally enticing to the early promise of public cloud as a replacement for all on-premises operations. In this approach the exact same interface used to access the cloud storage can be used to access the on-premises archive.

1.2.7

Longer Retention Times

Drivers for longer data retention include scientific research, healthcare, litigation, public safety and national security.[8] Privacy is a growing concern and there is litigation that limits retention in certain circumstances, such as The Data Retention and Investigatory Powers Act in the United Kingdom, and the European Union Directive 2006/24/EC. The U.S. has seen changing regulations around retention of cell phone metadata.

Nonetheless, the trend in many industries is to store data for longer periods of time, and this continues to be one of the key drivers for customers to use tape products. **Some examples from the U.S. include:**

- Diagnostic Images (HIPPA regulation): **5 years.**
- Loan applications (Home Mortgage Disclosure Act): **3 years.**
- Insurance claims (Federal Register): **10 years** after termination.
- HR personal records (OSHA regulation): **7 years after termination.**[9]

In general, requirements vary greatly on how long data must be retained. In the healthcare industry there have been mandates and new legal requirements established toward ever longer retention times (e.g., keeping records and images past the life of the patient, or keeping records from birth to a period beyond a patient's life).

Another change in retention times is for litigation reasons, especially in the United States. Some companies have had policies to manage the deletion of data. As the amount of data proliferates, it is increasingly difficult to enforce the managed deletion, or to have data categorized for automatic deletion.

In some cases, legal reasons require companies to keep data for longer periods of time until legal deliberations complete.

For these particular use cases, storing data on tape is a cost-effective way to keep the data. In addition, storing the data in a standardized format such as Linear Tape File System (LTFS), tar or AXF allows for interchange as well as being good preservation practice.

Another key requirement in record retention is data integrity. If used for a health care decision or in litigation, the end user wants to be sure that the data has been preserved accurately. Tape solutions today have built-in, end-to-end data integrity both at the block level with integrated CRC checks, and at the file level with industry standard hash codes incorporated into the archive storage software. Tape can be set up as WORM to provide another layer of protection from manipulation of records.

The combination of tape's low cost, standard formats and data integrity make it a common choice for long-term retention of large amounts of data.

"Tape is ideal for long term, deep archives. Our customers typically transition their data to tape when it is over two years old. We have never lost content."

Dave Thomson, QStar Technologies

1.2.8

Security and Cyber Resiliency

The staggering amount of data now being collected, analyzed, and stored has advanced our world's technological progress. But it has also created opportunities for criminals to exploit the new Digital Age for their own nefarious ambitions. Cybercrime, including Ransomware attacks, are on the rise. According to Interpol, individual hacktivists are being replaced by global cybercriminal networks that have the potential to commit wide-ranging deeds of cybercrime with catastrophic results.

Organizations today are more at risk than ever from cybercrime. One such scheme is Ransomware, which is malware that locks down a screen or network until a requested payment is made to an anonymous recipient using bitcoin. In 2022, the FBI's Internet Crime Complaint Center received 800,944 complaints, with losses exceeding $10.3 billion. This number does not factor in the hundreds of organizations that paid ransoms without reporting it for privacy purposes.

To fight cybercrime, organizations must defend themselves with smart data storage and security strategies. One such strategy that organizations are adopting is to store copies of data on electronically disconnected devices – even when copies are stored on disk or in the cloud. This disconnection creates what is called an 'Air Gap' or separation from the network that prevents attackers from accessing the data. Tape, due to its removability, enables organizations to keep data offline and safe to protect themselves from disruption of service; loss of critical digital assets; and hefty ransom payments.[9]

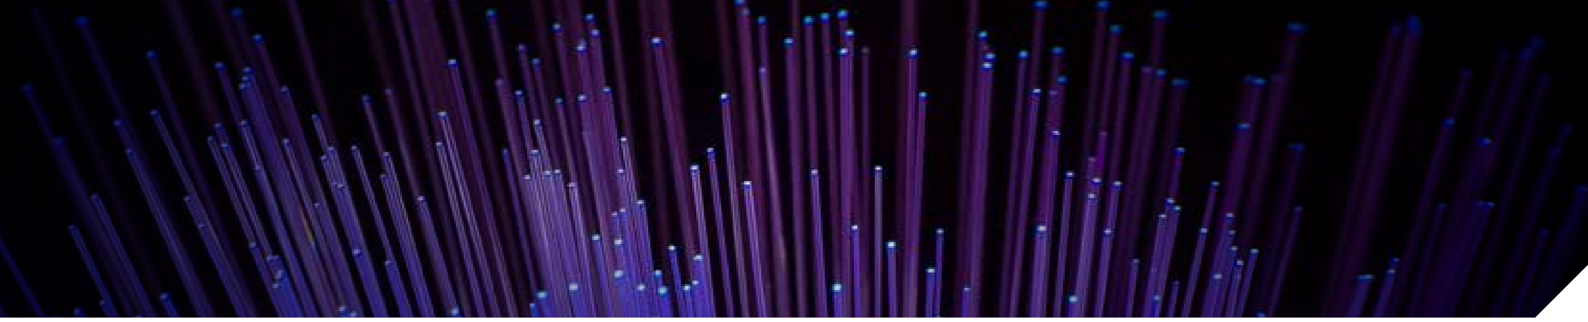

1.3

TECHNOLOGY TRENDS

1.3.1

Competing Technologies

The primary requirements for archival storage are high availability, security, and reliability together with the ability to store massive amounts of cold user data for extended periods of time with a very minimal carbon footprint and low TCO. Magnetic tape has long been the medium of choice for archival storage using multi-copy protection architectures, where tapes are kept in recommended archival environmental conditions reliably using optimized data migration policies for up to thirty years. These architectures typically require well controlled tape handling in addition to environmental requirements such as optimizing operations for streaming mode—where random reads might not be the primary mode of operation—minimizes back hitch or shoeshine-type motions.

Cold storage archives utilize traditional storage media such as HDD and SSD with software driven systems that power down the devices or utilize special HDDs. This schema is designed to reduce energy consumption at the cost of longer time to data access. A SMR HDD drives are often used in these solutions to provide high density at a lower cost of acquisition. These energy conserving systems require special software to manage the infrastructure, and even more special software if using SMR HDD given the linear operation of the HDD drives. Examples of cold storage HDD systems include Dropbox SMR deployment and Microsoft Pelican Project. Both deployments recognize benefits of the systems, yet neither reach the goal of utilizing less energy than simply utilizing existing tape infrastructures for infrequently accessed data.

Read more:
https://dropbox.tech/infrastructure/four-years-of-smr-storage-what-we-love-and-whats-next

Read more:
https://www.microsoft.com/en-us/research/project/project-pelican/

Visual Media are medias that have been purposed to have service life similar to movie film. These devices require a coding method to turn digital data into a visual reference, such as QR codes. Film based storage, physically lasts hundreds of years when the environment is properly controlled. While this is one of the most secure methods of data storage, a drawback to this type of storage for massive digital data repositories is the performance and accessibility of the data. Each visual reference must be visually read from the media, transformed to digital data, and restored as a full data set. Because of the data performance visual storage is best when for visual data and small digital data where data preservation and provenance is more important than cost for performance. An example of this type of storage solution can be found at https://piql.com

Ceramic Nano Media is the latest generation of optical laser storage. CNM is significantly different than traditional laser disk. Digital data is encoded into data matrices, the matrices are written to ceramic nano layers on ultra-thin glass using a laser beam to write the data in nano patterns. Another way to describe this is that small squares of patterned data are placed into larger patterns to form the digital data. The data is both permanent and highly resistant to the environmental factors that stress other forms of media. Some examples are water, heat, fire, electromagnetic pulse, UV light and radiation. Ceramic Nano Media is a visual media, all the data is read using a high-speed high-resolution microscope. Once the data matrices are read, they are then decoded back into digital data. While CNM is moving quickly by utilizing existing market technologies commercialization is in the initial phase of market development. The main developer of the technology Cerabyte has stated a 2024 target of 100MB per second data rate, 1PB per rack and 90 second access time with the goal of reaching an extremely dense target of 100PB per rack, with data accessible within 10 seconds by 2030.

Read more:
https://www.cerabyte.com/how-it-works/

Multi-layer fluorescent film storage stores data in bleached or unbleached fluorescent dots in a multi-layer film[38] that is created by extruding, splitting, folding and re-extruding a polymer film until it has 8 layers. Then a disk-shaped section is cut out and laminated onto a platter. Data is written and read using lasers focused at the right depth in the multi-layer recording medium. The production process is low-cost compared to tape and disks can hold from 500GB to 1 TB of data with 2 TB to 4 TB using 32 layers in prospect. A multi-disk cartridge could hold even more. While this technology is currently available, significant adoption by enterprises has not been observed. This may be due to the increased number of assets that need to be managed along with the relative low data write performance of the media. This technology is developed and being marketed by Folio Photonics.

Read more:
https://foliophotonics.com/

DNA data storage is based on the fundamental building blocks of life. By coding the digital bits to DNA sequence letters data can be stored on synthetically modified material. The synthetic material is then encapsulated in sealed storage pods that resemble small "pill capsules" which can last 50 to 100 years in a reasonably controlled environment. Like visual data once the data is stored it cannot be modified without being destroyed and completely rewritten, making the medium completely immutable. DNA structures are smaller than any other modern bit storage mechanism, as a result it is possible that a Zettabyte of data could fit in a shoebox. The challenge for DNA storage is the coding and decoding process. As of this writing decoding of human DNA can take hours and the industry is working to bring this down to minutes. The cost for HDD is currently ~$100 per TB. According to Twist Bioscience CEO Emily LeProust, currently the challenge for DNA is that it is more expensive than HDD but the cost for DNA is coming down quickly as more investment is made in the coding and decoding processes. This storage option is leveraging the science of DNA not just data storage.

Read more:
https://dnastoragealliance.org

Silica media, otherwise what can be considered glass is a sustainable media that has the possibility to be a long-term storage technology. As with other non-magnetic forms of data storage, silica has a service life of hundreds of years, is low cost, durable and EMF proof. Durability of the glass platters is dependent on zero handling and proper storage; it is still glass. Storage of data in glass is based on the fundamental idea of laser disks at a much higher performance and density level. The current project Silica from Microsoft utilizes a Femtosecond pulse laser to write and read data on glass platters. The laser penetrates the glass and changes the polarization of the write beam to orient the beam as "voxels", each of which can contain multiple bits of data. Unlike traditional storage solutions, the media does not move, the laser moves to stabilize the data at time of write and read. The silica platters enable multiple layers of data writes greatly increasing volumetric data density. Today the raw density estimate calculates to 7 TB contained in a 2mm think DVD size platter. There is complexity in the solution created from using different technologies to write vs read the data. Reading is accomplished using polarization sensitive microscopy. As with many laser technologies, cost of lasers and individual data rate performance is a challenge, a challenge now being addressed by the developers. While INSIC sees this as a viable method of storage for long term data, the technology is most likely to benefit managed storage providers. Certain challenges must be overcome in this technology to become commercially viable, glass is relatively fragile, and systems breakdown meaning the longer-term management of this system must take into account modernization of the larger storage solution, media handling being part of that management.

Read more:
https://www.microsoft.com/en-us/research/project/project-pelican/

As of this writing there are other announcements related to archival storage systems. The primary announcement has been from Huawei using magneto-electric disk storage. However, the announcements have not been accompanied by hard data enabling analysis of the cost, durability, sustainability, and viability of the proposed solution. While all the covered competing technologies have benefits for long-term data retention, there is no technology that is commercially available that provides more value for long term scale data storage than tape. We expect this to be true through the end of the decade.

1.3.2
Unstructured Data Integration

As the amount of data generated continues to grow, retention has become more critical as the need to generate value from the data has also grown. Unstructured data now dominates data generation and retention. The two primary types of unstructured data are Filesystem and Object storage.

According to Gartner

80%

of generated and stored data is unstructured[10]

Each of the data types have advantages and disadvantages, but both have one thing in common – the need for tiered storage of infrequently accessed archive data. Filesystems have supported data tiering in many forms and have supported tape as the low-cost tier for archives for more than a decade. Recently, software and hardware vendors have been focused on integrating tape into object storage solutions. Most of these solutions use the LTFS self-describing tape formats, but have proprietary tiering mechanisms, includes orchestrated command-sets.

Nearly all solutions utilize the standardized S3 interface including the S3 Glacier command set. As 2024 S3 solutions supporting tape are available from Fuji Film, Grau, IBM, Nodeum, Point, Quantum, Qstar, and Spectralogic. The ability of all these vendors to provide the standard S3 API, any software that can utilize the interface can now utilize tape as an archive target for backup and recovery, cyber resilient airgap, and data archives without coding the tape interface.

1.3.3

UBER (uncorrectable bit error rate)

Tape media raw durability is computed based on Bit Error Rate (BER) and represents the quality of data written and retained in an offline archive state. Operational data durability is highly dependent on the infrastructure, environmental conditions, and system architecture of a data operation. End-user operational experience may vary.

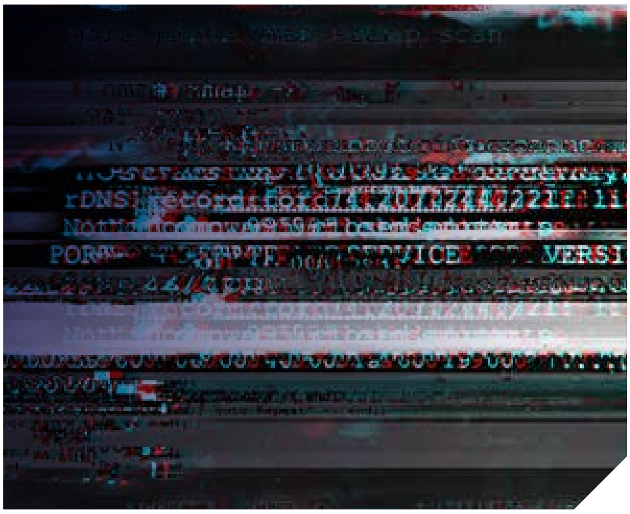

The basis for exceptionally strong UBER performance—which, with new tape technologies, can reach 17-nines of durability[11]—is incredibly resilient and efficient ECC Format and advanced magnetic recording technology. The unique Read-While-Write (RWW) process of tape technology basically ensures that data on tape is written error free because any detected errors that don't comply with ECC format requirements will be rewritten in full streaming mode down the same tape resulting in complete verification of written data durability. As a result, assuming that tapes are written by well-maintained drives with an efficient writing process and in archival environmental conditions which are normally more stringent than modern data center environmental conditions, archival state of tape with data can be defined as 17-nines, which we define as the Archive Durability.

As previously stated, operational durability, which describes the data durability during Read Mode streaming or random, is dependent upon numerous external factors, such as host software, robotics, reading drive quality and as environmental conditions. To maximize the nines for operational durability, protection strategies like replication, erasure coding, and RAIT can be employed.[34]

Read more:
https://asset.fujifilm.com/www/us/files/2020-03/77e3d6ef9afef39c8a5e2d235c75430d/Data_Tape_Care_Handling.pdf

1.3.4

RAIL with Erasure Coded Tapes

RAIL stands for Redundant Array of Independent Libraries when used with Erasure Coded tape technology it enables spreading encoded objects or files across tapes, individual libraries, and over different locations even across the globe such that failure of a site, tape or even a library will not affect data availability and durability.[36] Hypescalers commonly employ this fundamental technology, RAIL with Erasure Coded Tapes, to increase availability across the globe while providing backup for archiving data using a very large number of Operational NINES[12].

To get a better understanding of Erasure Coding with Tapes, one may need to look at the revolution that has taken place in HDD servers starting late 1980's when researchers at the University of California, Berkeley, offered a more automated and efficient way to achieve redundancy and speed improvements called RAID technology. Before the widespread use of RAID, basic data redundancy often involved copying data to multiple hard drives like what current tape systems do. However, eventually, RAID-5 migrated into RAID-6, and then Erasure Code based architectures dominated data center applications. To get a file any size, one must read multiple HDD's like if 12/4 ECC policy is used the read must be from any 8 drives. Because HDDs are never powered off, data is always highly available with very minimum latency. High transfer rates are also achieved because several drives are pooled to read data.

Cartridges containing archived data, on the other hand, are kept offline in a library system and are loaded into drives by local robotics when needed for data operations such as reads. Libraries typically have an approximate cartridge to drive ratio of 120:1, which means that there are more tapes than drives. Because of this, systems using typical erasure-coded architectures like HDD applications may experience longer tail latencies, particularly when dealing with small files because of limited drive availability. However, especially with Active Archive applications, RAIL based Erasure Coded Tape with S3 Object store may offer benefits including aggregated bandwidth across multiple tapes and libraries[12], high operational data durability NINES, and high storage density footprints due to less tape usage. Several S3 interface Erasure Coded Tape Object stores have recently been made available; these new products may help users in selecting from alternatives, including a combination of replication and erasure coding based on their use case requirement. The tables below show some of the key features with benefits and limitations for both methods.

Durability & Overhead	Replication / Copy	RAIL w/ Erasure Code & RAIT
Archival (UBER) Durability	17 NINES	17 NINES
Operational Durability	Higher number of copies are required to increase NINES	Erasure Code policy provides parities to increase NINES
Overhead	Minimum 100% with 2 Copy and 200% with 3 copies	Less than 100% for all Policy Options

Figure 3. Durability and overhead calculations

Configurations	Replication / Copy	RAIL w/ Erasure Code & RAIT
Geo Spread	2 copy approach is ideal for 2 Geo site	Works best with 3 Geo site or more
Drive Count	Can operate with single drive although performance may suffer	Requires more drives than replication

Figure 4. Configuration options

Performance	Replication / Copy	RAIL w/ Erasure Code & RAIT
Capacity Density, PB/ft2	Lower due to higher Tape count	Higher based on Erasure Code Policy
Migration	Must manage migration of all copies; more tapes to be migrated	Higher efficiency migration due to low storage overhead and lower tape counts
Transfer Rate	Limited by single tape xfer rates	Higher Xfer rates due to multiple drives reading in parallel
Tail Latency	Low Tail Latency using single tape	Based on architecture Tail latencies may vary
Small File performance	Need single tape to access	Based on architecture more than one tape may be required

Figure 5. Performance distribution across usage and configurations.
Source: INSIC

1.3.5

Host Interface

Tape has traditionally followed disk technologies for its interfaces, which is why tape is predominately natively attached through Fibre Channel (FC) or Serial Attached SCSI (SAS) interfaces. The tape industry monitors and implements interfaces based on larger market directions. The longevity of tape infrastructure forces the tape industry to be diligent in the observation of trends related to connectivity. The info below is an analysis of the current state of the industry.

FC has a robust roadmap being followed by the storage industry. Despite the frequent calls for integration of SAN with WAN/LAN capability, data storage infrastructure requires extreme performance that often results in segmented infrastructure. This maintains FC dominance in the market, currently 47% of connected capacity. SAS continues to be utilized widely but is under pressure from the NVME protocol over ethernet. Other interfaces such as Infiniband and RoCE have found their place in niche markets and have been most successful with server-to-server connections. USB-C performance is a viable option for consumers and desktops, but has not been adopted by data centers.

Ethernet attached storage continues to have viability and market acceptance, mainly with HDD and SSD storage devices, capturing 11% of attached capacity.[13,14] Ethernet attachment is further enhanced with NVMe-oF. While NVMe-oF has great interest in low latency architectures, most analysts have not seen it as a replacement for SAN. Since NVMe-oF is currently architected for low latency storage, tape attachment would require significant changes, so NVMe-oF is not a consideration as a new interface for tape for the next 5+ years. The tape industry must be vigilant in monitoring the technology as maturity makes it more affordable for high-latency infrastructures like tape.

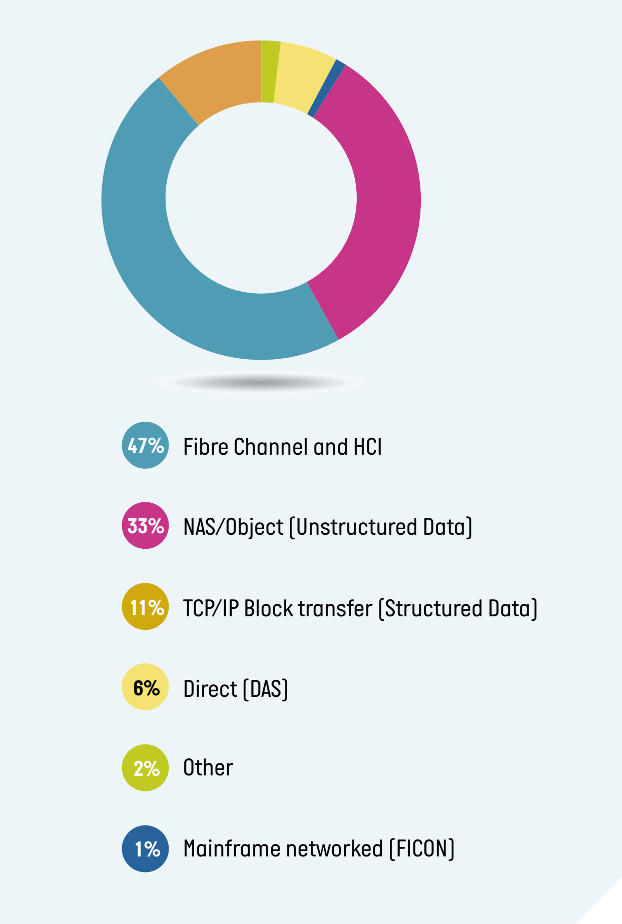

Figure 6. Performance distribution across usage and configurations.Source: IBM 2024 External protocol report

As adoption of various host interfaces shifts in the marketplace, the tape industry should react with interface converters and eventually new native interfaces. For the next 5-10 years, FC and SAS native interfaces should be sufficient.

1.3.6

Post-Quantum Cryptography

Quantum computing has been discussed in the news in the past few years as companies like IBM, Microsoft and Google give access to, or demonstrate quantum computing on, higher and higher numbers of quantum bits.[15] Quantum computing is a highly focused industry making strides quickly to be more powerful than traditional supercomputers.

Quantum computers erode the security of the current encryption algorithms as the number of qubits grows. While implementation of Grover's Algorithm[15] on quantum computers makes many of the current asymmetric cryptography keys less secure or insecure, the symmetrical encryption schemes such as AES will last longer.[16, 17] Urgency to replace existing cryptography depends on the data migration time and shelf-life time set against the threat timeline.

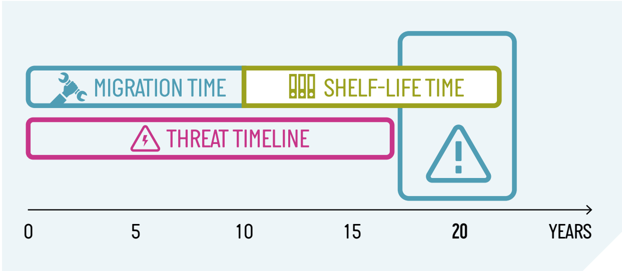

Figure 7. Threat urgency
Source: Global Risk Institute, threat timeline

According to the Global Risk Institute, within the next decade, 55% of expert respondents surveyed indicated that there is a greater than 50% chance RSA-2048 will be breakable within 24 hours by Quantum computers.[19] Today's tape drives use GCM-AES-256 encryption. The tape solution encryption techniques are estimated to be secure over the next decade. However, there are still communication paths that use asymmetrical encryption schemes like RSA and ECC. These paths will be updated by the tape vendors to reduce risk on the threat timeline.

The National Institute of Standards and Technology in the USA has been working with researchers world-wide on new encryption algorithms that are more secure against quantum computers. NIST announced that new algorithms should be finalized and standardized in 2024, allowing for software libraries and implementations to be available afterwards. As these become available, they will be adopted by the tape industry, improving the longevity and security of archives.

Read more:
https://www.nist.gov/news-events/news/2023/08/nist-standardize-encryption-algorithms-can-resist-attack-quantum-computers

1.4

COST & ENERGY CONSUMPTION ADVANTAGES

1.4.1

TCO Operational Planning

While TCO is only one measure of usability of a solution, it is one of the few metrics that can be tracked and placed into IT budget planning. The TCO of tape has been demonstrated across multiple use cases, most common amongst these is active archives and deep archives of data. Tape is most cost effective where infrequently accessed multiple Petabytes of data exist.

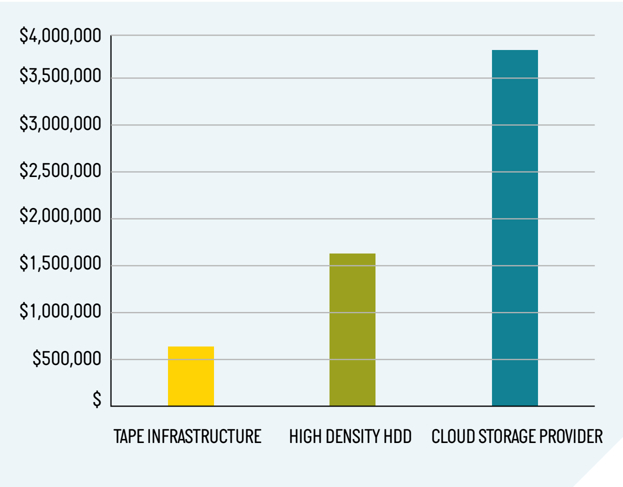

Figure 8. 10 Year Total Cost of Ownership Comparison for Tape, Disk and Cloud
Source: LTO.org

Figure 8 compares 18 TB high-density HDD, and S3 Glacier Deep Archive (bulk recall). Specific data points include: 20 PB stored for 10 years, with no data expansion, 1% data retrieval/egress per month. Tape: Zero compression, no hardware refresh with current roadmap. HDD: Zero deduplication, 72-HDD drawers, 3 expansions per controller, OCP implementation, planned refresh at year 5.

The model yields outcomes that are very similar to other TCO tools. HDD is 2.5x and S3 glacier is 5.9x more expensive than the tape solution. While tape is less expensive than HDD it is not the position of INSIC that all archive data must be solely stored on tape.

In fact, given the importance of archival data for long-term preservation and/or future business monetization, INSIC is proponent of protection methods such as multiple copies, erasure coding, RAIT or combination depending on the needs of the application, which may translate to multiple mediums depending on the data access frequency. The value of HDD is in the frequency of access performance to the active archive data. Operations of hybrid tape and HDD environments, provide users the most optimum experience. Storage data placement and the associated meta data are rapidly becoming the most important aspects to meet both IT budgets and business expectations.[33]

Read more:
https://www.lto.org/tco-calculator

Read more:
https://www.fujifilm.com/us/en/business/data-storage/resources/tco-tool

Read more:
https://aws.amazon.com/s3/glacier/pricing/

1.4.2

Modeling Methodology

As with any modeling methodology, the results can vary widely depending on the factors considered in the model, the underlying assumptions, and implementation of calculations. Nearly all models have input variables that are standardized to improve the usability of the tool.

Primary inputs include:

- Initial data capacity required.
- Annual expected growth of data storage.
- Term of retention comparison.
- Amount of data recall applied to cloud storage.

Secondary input variables include:

- Annual floor space expense.
- Cost of energy.
- Software.
- Starting Media capacity.
- Cloud SLAs.

Once the primary data is entered the most influential part of the TCO studies are the assumptions and calculations used to derive the TCO comparison. Each storage solution has individual component contributing factors influencing the outcome. This table outlines some of the most influential assumptions for each storage solution.

used beyond the manufacturers standard warranty). While these assumption factors are different, they must be applied fairly across the TCO calculations.

These costs are included in the service charges for cloud systems; however, cloud also charges for data retrieval and egress.[7] All credible tools list the factors

Tape	HDD	Cloud
Equipment and media useful life	Equipment and media useful life	Data storage plan
Equipment configuration	Solution configuration	Number of availability zones
Data compression ratio assumed	Deduplication ratio assumed	Location of availability zone
Data migration strategies	Data and equipment refresh	Data retrieval performance
Solution future capacity	Solution future capacity	Dedicated telco charges
Energy consumption	Energy consumption	Monthly amount of data egressed
Thermal energy	Thermal energy	Billing verification
Operational overhead	Operational overhead	Information security costs
Cost of maintenance	Cost of maintenance	

Figure 9. Requisite Assumptions for TCO modeling. Source: INSIC

In comparing tape, disk and cloud systems it is important to make realistic assumptions, many of which are different for the three systems. For example, HDD and tape have different failure rates and those rates are calculated differently for media and other components.[20] HDD infrastructure tend to have product refresh cycles of 4-6 years vs 9-11 years for tape (though service and support need to be considered for any storage system

and assumptions used in the calculations with enough detail to determine the fairness of the TCO comparisons. Although outcomes of a TCO compare may lead to questioning why anyone would pay a given fee, it does not discredit the TCO tool outcome, only the validity of the storage solution for a given use case, often this is the case for deep archives from cloud providers where long -term storage and even minimal data egress are required.

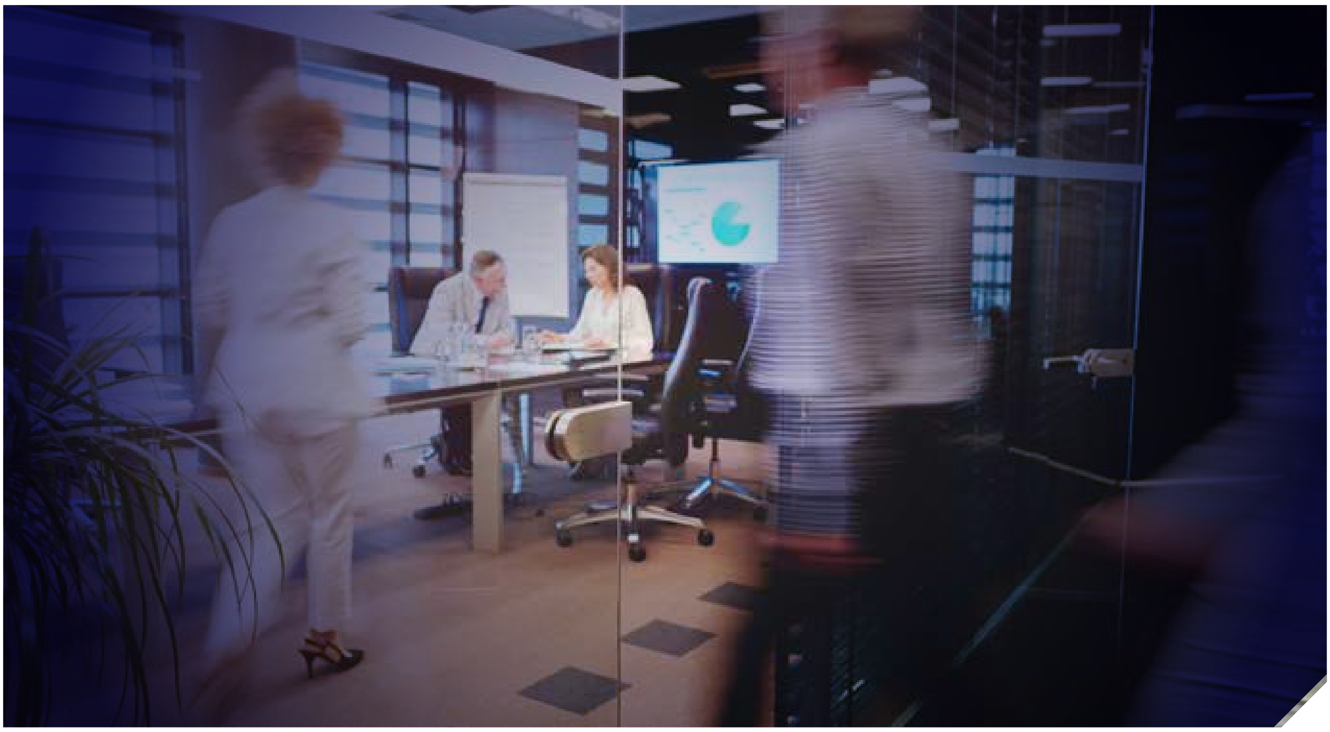

1.4.3

Energy Consumption

According to an article in the Financial Times, the energy consumption of data centers has stretched the electricity grids in Europe to the point that new data centers and existing data center expansions are being paused in some locations. The largest focus in greening data centers is on the server efficiency. However, according to a study by Furthur Market Research data storage power consumption will grow from 17% of data center energy in 2020 to as much as 29% by 2035[4].

As mentioned in the TCO tool section above, the TCO and related energy are all dependent on the configuration of the solution. Tape inherently is low power usage, HDD has a higher per unit and solution consumption, while also offering much greater efficiency in frequency of access. Most TCO tools provide an energy usage comparison specific to the data values entered. Tape consistently uses 85-97% less energy than HDD in archives.[3]

Figure 10 illustrates the total power draw according to global usage of storage technologies. The figure compares SSD, HDD and Active archives including tape between 2020 and 2035.

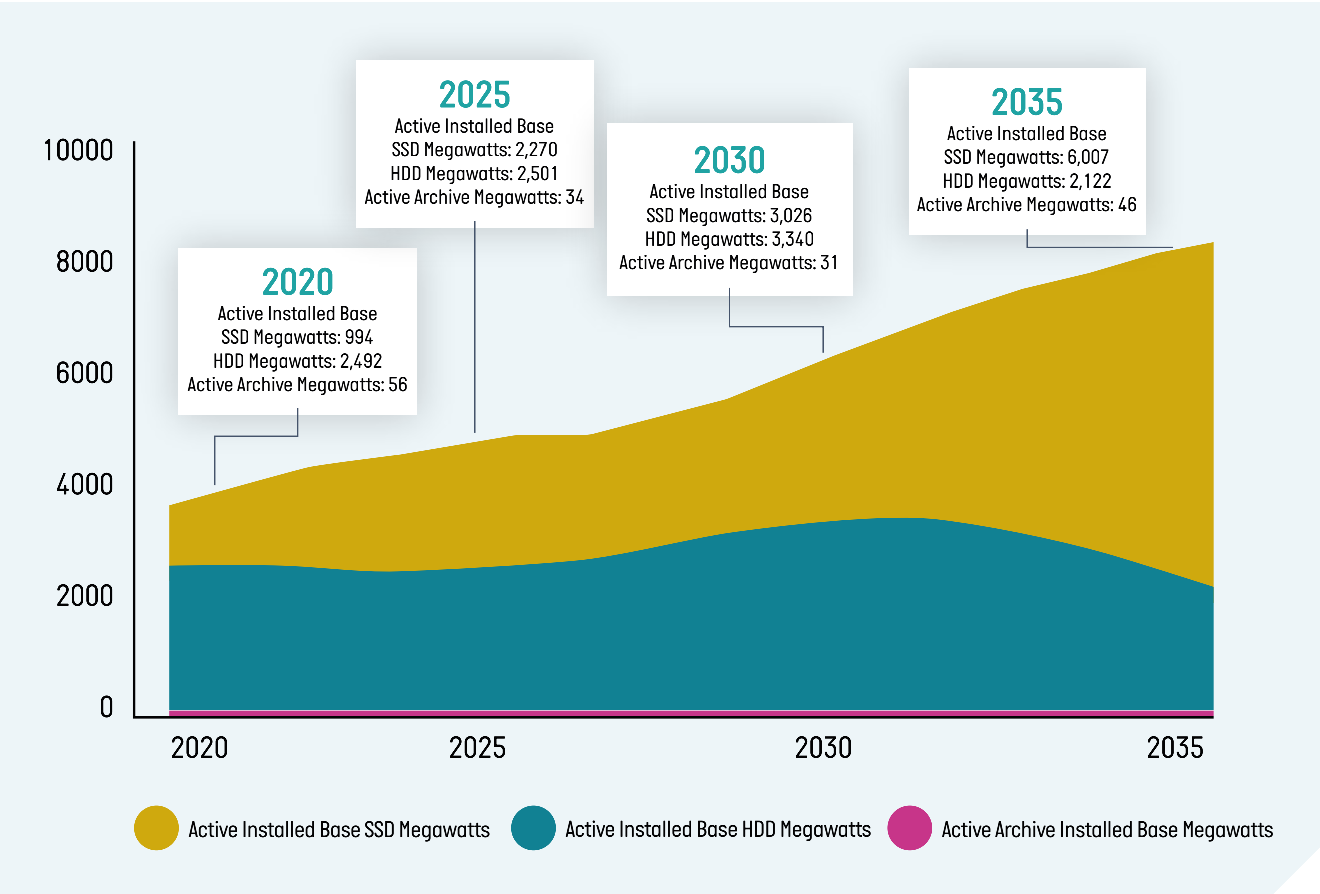

Figure 10. Annual SSD, HDD and total "Active Archive" Power Draw
Source: Furthur market research

1.4.4
Trends in Storage Costs

Historical studies can show interesting trends in storage costs over time. Although each study may use different methodologies, each study is based on available data that yields similar results. Figure 11 below illustrates a log scale of the historical and projected trends comparing the most prominent storage devices; SSD, HDD, and Tape-based active archive systems.

Read more
https://furthurdata.com/wp-content/uploads/2024/02/Sustainable-Preservation-of-Enterprise-Data_V13.pdf

“Tape will have an enduring and widening cost advantage.”
John Monroe, Furthur Market Research

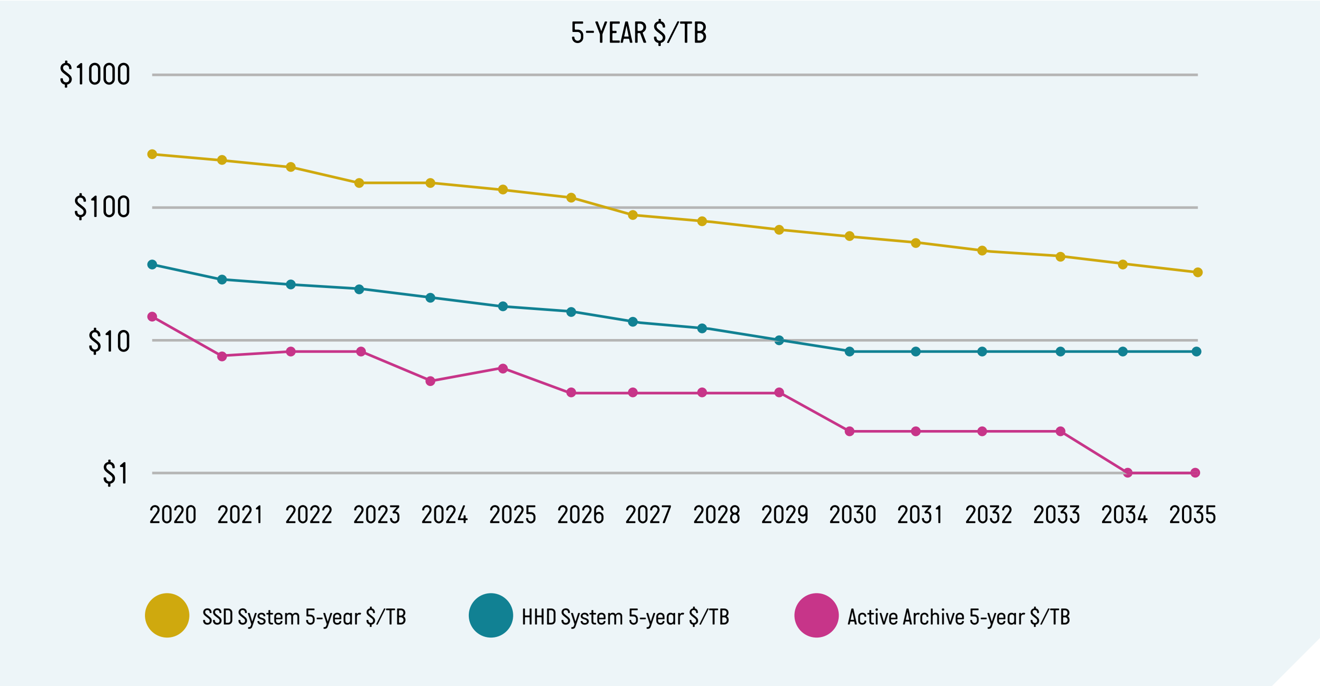

Figure 11. Five-year cost-per-terabyte trends 2020-2035
Source: Furthur Market Research

While the multiplier between disk and tape systems fluctuates over time, the average gap in the TCO is 3.98 times less expensive for tape solutions.

Regardless of the cost gap between HDD and Tape, tape solutions cost significantly less than either disk or cloud solutions, particularly for multi-petabyte environments. This cost difference represents many millions of dollars in all these studies. As the studies emphasize, it is not a question of disk or tape but a question of how much of each is needed to deliver a good archiving solution. Clearly, for infrequently used archive objects (such as files), tape offers significant TCO advantages.

1.4.5
TCO Conclusions

Tape systems continue to cost significantly less than disk systems. The gap between cost for HDD and tape will continue to be variable but more stabilized than observed prior to 2020. The cost of data storage must also take into consideration the amount of energy consumed with long-term data at scale. When frequency of access is not a consideration, tape and HDD are significantly lower in total costs than Cloud storage at scale. Moreover, what is often overlooked in TCO models, beyond the costs of operation shift, are the security considerations and costs associated with moving the data to and from the cloud.

1.5
THROUGHPUT AND RELIABILITY

1.5.1
Data Rate

Data rate of tape drive is a function of linear bit density (Kbit/m), format efficiency (%eff), tape speed (m/s), and number of data channels. Data rate is limited to min and max data rates because tape speed is limited to a range. Min tape speed and number of channels determines min data rate.

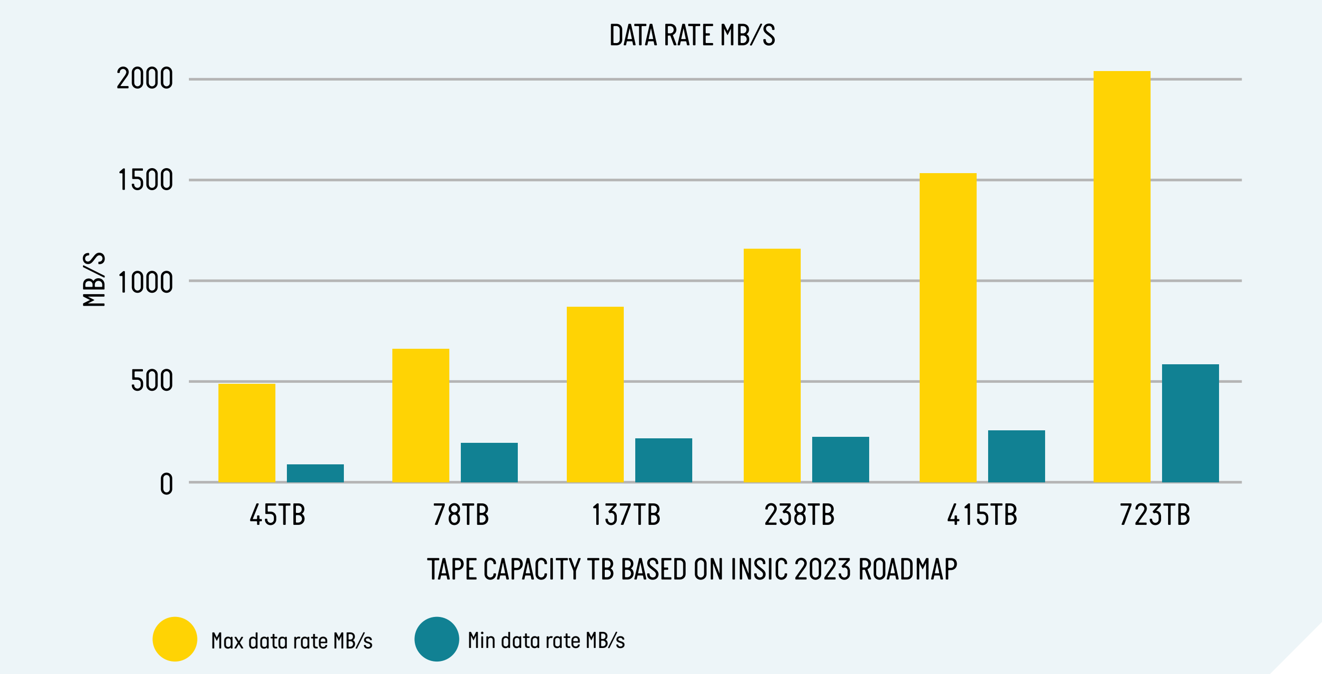

Figure 12. Tape transfer rate
Source: INSIC

With recent LTO-9 drives using 32 channels, the min tape speed is approximately 100 MB/s with max speed approaching 400 MB/s also depending on the form factor (HH or FH). It is important to note that as max data rate increases as number of channels increase, the min data rate also goes up; figure above shows that for capacities 78 to 415 TB the min data rate will be around 200 to 268MB/s due to 64 channel head configurations. This can create challenges for the Host software and hardware architectures to make sure tape drives are streaming at optimum speeds with minimum under or over runs to due lack of buffer size and speed mismatches. This can be an especially important technology when optimizing tape handling for writes and reads with tape systems where min data rates are higher. Larger drive buffers and optimized host server architectures can be ways to address these problems. ROA algorithm can be an effective solution to minimize not only overall time but also minimize the overall back hitch motions.

Current drive technologies offer multiple discrete speeds which drive FW adaptively selects best speed to match host data rates for writing and reading. There are other algorithms that utilize variable speed operation that makes decisions to select optimum tape speed using a variable speed mode.

Figure 12 below is based on current 2023 INSIC technology roadmap showing maximum and minimum data rates per cartridge capacity.

The results in this simple “what if?” study, demonstrates that tape-based archiving applications have much upside potential over the next 10 years as new backup applications come into play. With RAIT or EC tape configurations, backup applications have new challenges to deal with, such as management of multiple tapes to recover a file or object, and access to multiple drives; however as new technologies are adopted, the advantage of tape in terms of low power, low TCO and high durability makes it a strategic archiving platform.

1.5.2

Durability of the Data

Annual Failure Rate (AFR) is a metric frequently used in reliability engineering to estimate the probability of a device failing within a year, which in our application device can be HDD or tape cartridge. Estimation of probability assumes exponentially distributed failures over the flat section of a typical Bathtub Curve referred to as the constant failure rate zone. Also, the exponential probability distribution assumes that the time between failure events is a Poisson process, where events occur continuously and independently at a constant average rate[21, 22].

Based on these statistical assumptions, AFR can be modelled using MTBF by the following relationship:
AFR = 1 – exp (– Annual Operating Hours / MTBF)[22]

The most current HDD study by BackBlaze estimates average AFR for an HDD utilizing 35 distinct models and 4 different manufacturers to be 1.70%[23].

In contrast to HDD systems, tape systems offer a removable medium with extensive interchangeability. Because of its special backward read compatibility, tapes can be read by both current and future generations of drives.

When performing restore or data migration operations, using backward read compatibility can be an effective recovery countermeasure for some challenging TDS issues. By using a next generation of drives with narrower readers, for instance, a tape having certain rare TDS errors from corner case environmental conditions during read vs write mode can be read with optimized margins.

Even though UBER describes the durability of data on tape at rest mode as being extremely high—up to 17-nines—when tapes are loaded into drives for read/write operations, such as data migration, typical tape use-case problems may happen because of the quality of the drives, the environment conditions, robotics problems, user software, and hardware problems for managing streaming tapes impacting AFR numbers. Contrary to HDD usage, tape removability provides more flexibility for data recovery using complex recovery algorithms that take into account variations in drives, backward reads, and environmental factors. For archival applications, the probability of failure of tapes can be significantly lower than those of HDD-based systems especially with copies and erasure-code applications.

As the numbers suggest, HDD and Tape are very reliable mediums for data storage. While many studies are available on the reliability of HDDs, few examples exist for tape. One example that is publicly available is published by NERSC. The NERSC team observed the center's actual tape data failure rate to be less than 0.09 percent using 40,489 cartridges, of which they could not recover data from only 35 of these. This study focused on a major data migration operation to update old tape cartridges to a newer system. Even of the 35 failed cartridges, they were only unable to read data from 178 meters of the 22,065,763 total meters of tape, resulting in an astonishing 0.00009 percent failure rate. The findings flew in the face of conventional wisdom: 99.9991 percent of tapes were 100 percent readable, representing a 0.0009 percent failure rate.[12]

The modeling of AFR for tape systems differs from that of HDD cases due to the removable nature of tape and the introduction of several distinct complex but very effective data recovery algorithms. A simulation-based algorithm that uses a tape-optimized Markov Model to simulate a deep recovery process could be a more beneficial approach for developing a Tape AFR model. This process could involve trying different drives, including next generation drives in backward read mode, head cleaning process, and even different environmental conditions. The Tape based Markov Models can then use this information to estimate the mean time to data loss, which can be translated into durability nines.[24]

It's important to note that unlike HDD systems, data on tape is cold and access rates are much lower. HDD errors are detected quickly so repair process can take place immediately. Data on tape is cold and protected by 17-nines of UBER while at rest given archival storage conditions. After a long period, system decides to use the tape to read, then the tape and drive based failure modes including deep and powerful recovery methods will take place. This is why typical HDD AFR models based on continuous and independently occurring failure modes are not applicable to Tape however simulation-based systems using real data information can offer a more realistic AFR and durability nines estimations.[24]

It is important to note that tape failures for archival applications can be further improved by using erasure code strategies making tape a very durable, green, long-life, and ransomware resilient storage medium for archival applications.

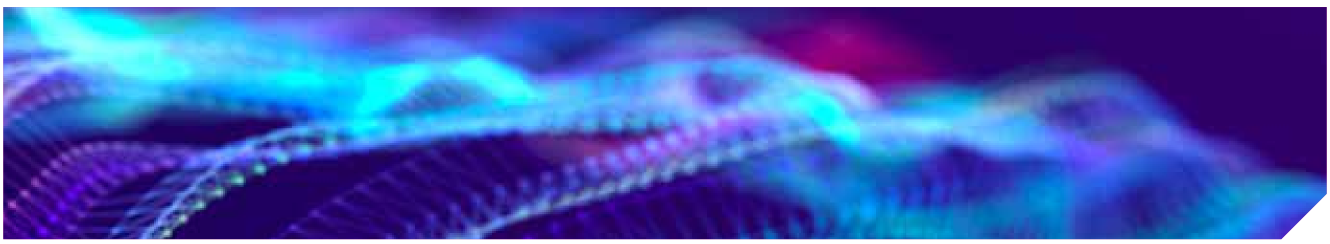

1.5.3
Media Wear Out

Tape wear is a complex function of media magnetic, tribological properties, environmental conditions, drive mechanics, FW algorithms and host software application. Studies have used surface roughness, friction, and stability of organic materials, such as binder polymers and lubricants to establish the magnetic and tribological properties of magnetic media under varying environmental conditions. Media wear can be explained by the number of Full Volume Passes to End of Life where a full volume pass is how many times to media must pass across the head, FWD and REV, to read or write full capacity, a cumulative motion count. The following are key factors effecting its metric:

- **Number of data tracks:** This will increase at 50% each generation to support capacity growth. Higher the number of tracks, lower the media EOL.
- **# of number channels:** Head mechanical design. Current generation of tapes has 32 channels, but INSIC roadmap estimates this to increase to 64 and eventually 128. The higher the number of data channels better the media EOL.
- **Media Tribology:** Continued technological improvements in magnetic media design and manufacturing to push the ceiling in media End of Life pass numbers. This is shown in INSIC 2023 technical roadmap with an expected 7% improvement every generation.
- **Drive Mechanism and Motion, Tension control FW:** Friction, acceleration, deceleration performance, tension control back hitch are all factors that can contribute to media wear if not optimized. Technological improvements in mechanism design and optimized tape motion algorithms will help with extending media EOL Recent RAO type optimized tape motion for improved Random File Reads are perfect examples of drive FW optimization improving media wear.[35]
- **Host Software and Hardware:** Tape is a streaming sequential media and data buffer management can result in drive under and overruns where an underrun is defined as tape drive internal data buffer emptying out causing fall of stream in write. Overrun is same condition in Read Mode where drive does not have space in buffer to fill with data from tape causing falling in streaming motion. Both of these can be due to non-optimum Host and drive synchronization and inefficiencies in Host Software including inability of hist to match drive speeds. Excessive underruns and overruns can further cause higher tape passes which can affect tape EOL.
- **Environmental Conditions:** Environmental conditions such as temperature, humidity, dust, can all affect the wear characteristics of magnetic tape media due to tribology. Extreme temperatures or high humidity levels can accelerate the degradation process, while exposure to dust can introduce contaminants or cause physical damage to the tape.

Overall, the picture below depicts the full volume pass count number of an optimal tape system for each generation of tapes, together with estimated capabilities based on the INSIC 2023 Technology roadmap.

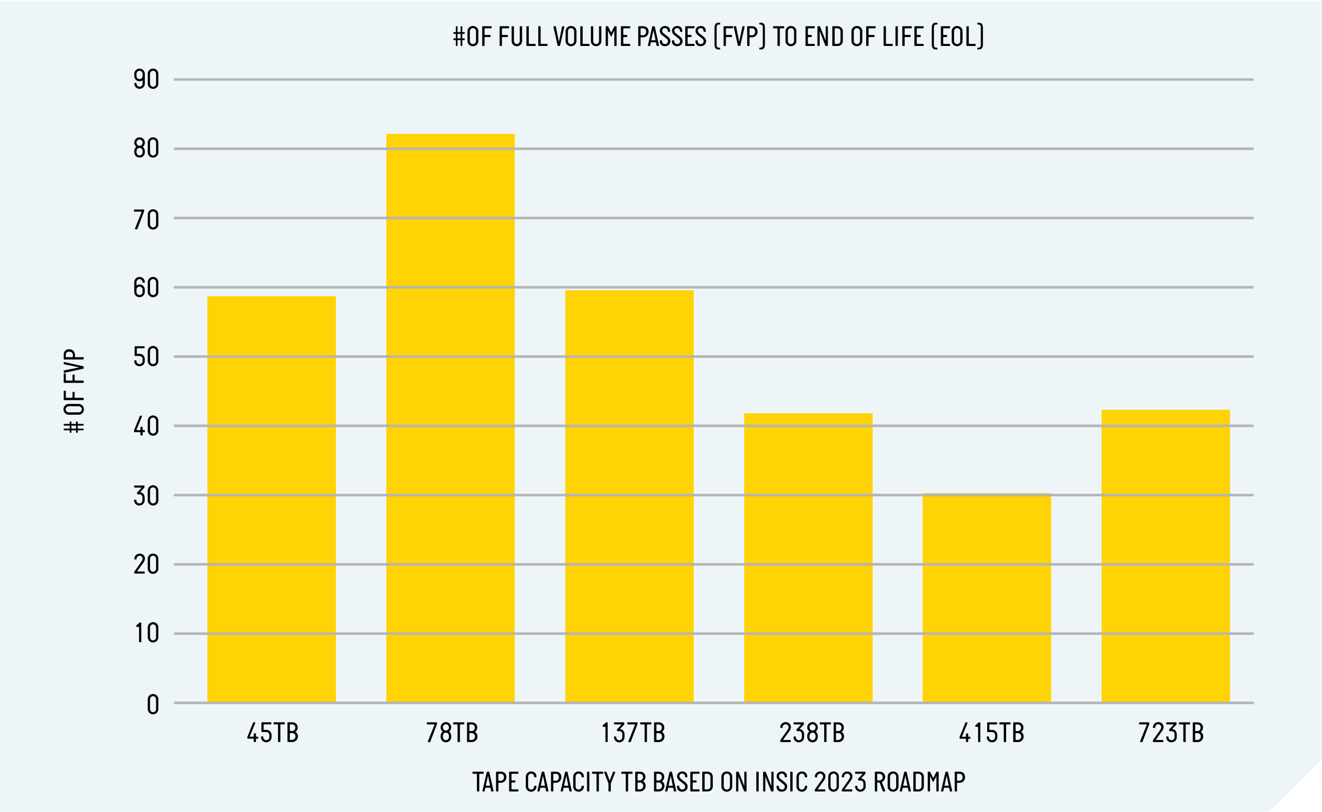

Figure 13. Full volume pass to end-of-life
Source: INSIC

1.6
CHALLENGES FOR TAPE ADOPTION

A persistent perception for tape is that it more difficult to use than disk, requiring special skill sets, complex processes, and more manual intervention. In addition, in some sectors there is general lack of awareness that tape is even used today. This point was made clear during our interview with Jay Bartlett of Cozaint Surveillance:

"With the video surveillance market looking very similar to the Media & Entertainment market, they both got here from different directions. Video surveillance -in the analog days- utilized VHS tape to record camera feeds. Unfortunately, when you mention (LTO) tape to these users, they immediately think of old, outdated technology."

Unlike disk, tape data is stored on removeable cartridges. This brings a layer of support complexity that would not be familiar to non-tape users. For example, when a problem is encountered in tape, it may not be obvious if it is due to bad media or a bad drive. For disk, a bad drive is identified and then simply replaced. To overcome this, tape vendors have created better analytics that can identify drives or media that are beginning to perform at a suboptimal level and issue alerts that units should either be replaced or have diagnostics run. Advances in analytics have made it much simpler for a tape user to have confidence in their tape storage system and its support.

Another difference between disk and tape that relates to ease of use has to do with tape's interface to the host. Disk systems typically present a block or file interface. Operating systems and application software support one or both interfaces. On the other hand, tape is a sequential access system and requires different drivers and needs the application to understand that the time to access a file may be measured in minutes rather than milliseconds. Applications written for disk may have time-outs that will fail a read from a storage device that doesn't respond in a few seconds. Tape vendors have responded to this in two ways. One is providing a file system on tape (LTFS) that enables each cartridge to present a list of files that can be independently accessed just as if they were on disk. With this method, users can drag-and-drop files to and from tape as easily as moving files to and from disk or flash. customers who required IT assistance for specialized skills to use tape in the past may no longer need to do so - they can easily use tape themselves now.

Another method for addressing the accessibility of information on tape is to include tape as a tier behind an abstraction layer that provides a standard file system image to the applications. The abstraction layer is configured with policies to migrate the data to/from tape and disk. This migration is hidden from the applications. To accommodate the read delays, different techniques are available, including keeping stub files on the disk tier to cover the gap in time while waiting for the tape to locate to the start of the file; keeping low resolution versions on disk for imaging applications and the high resolution files on tape; keeping the most likely data to be retrieved on disk and using the tape tier for the rarely read data; and keeping one copy on disk and using tape as a data protection storage tier.

With the software and drivers available today, it is possible for users to record and access their data on tape without them needing to perform any special processes or ask their IT department. In most cases they won't even know their data is stored on tape. Using tape today, then, provides an unbeatable combination of the ease of use of a disk but with all the benefits of tape. For more information see the HSM data tiering section 1.7.5.

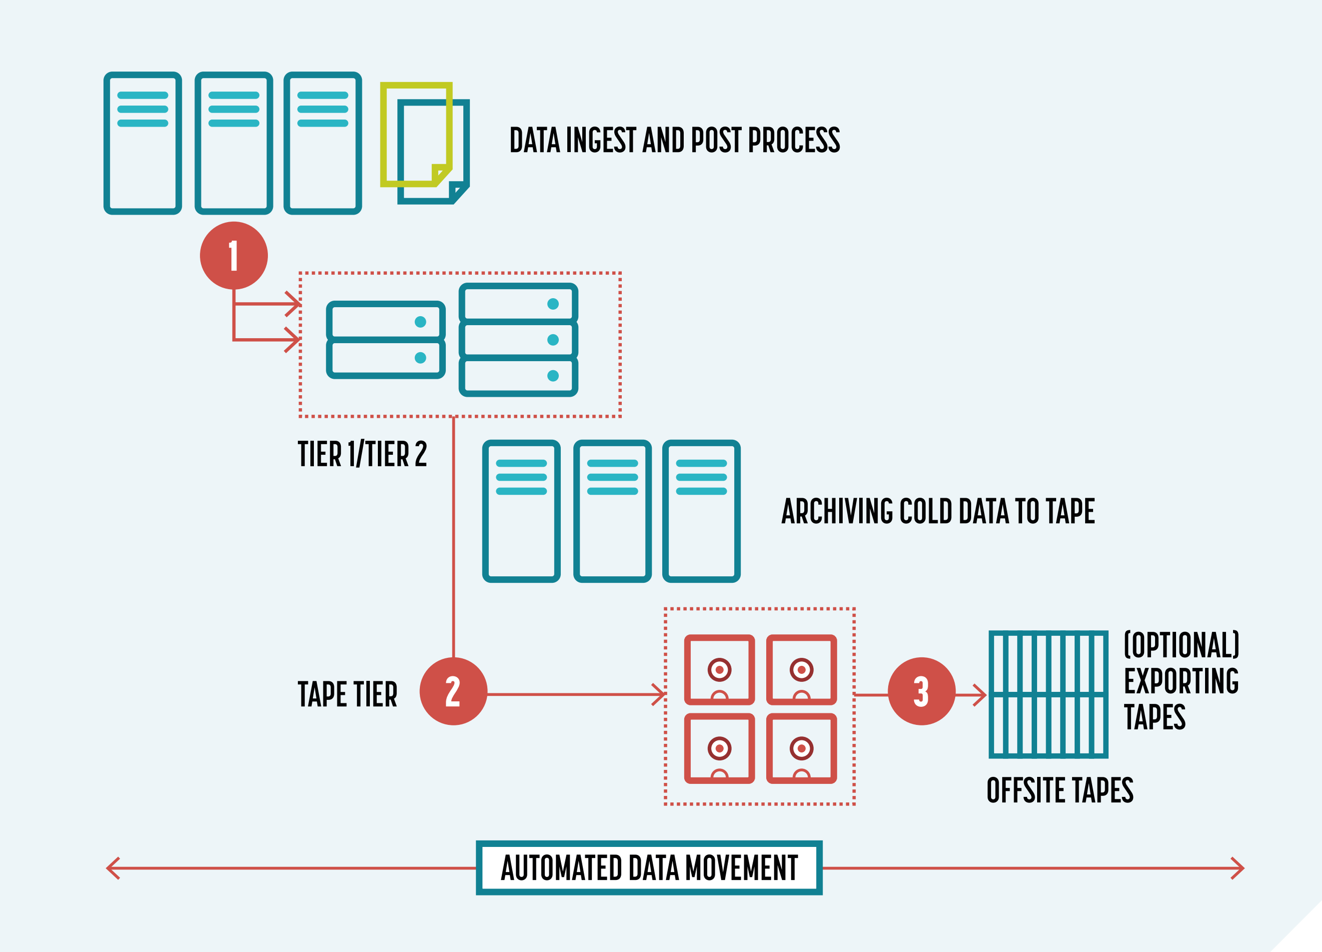

Figure 14. Automated storage data placement
Source: IBM

1.7

VERTICAL INDUSTRY USE CASES FOR TAPE

The use of tape spans various industries, offering secure and cost-effective solutions tailored to their unique needs.

1.7.1

Financial Service

In a data-driven sector – examples could be customer financial records, transaction histories, or compliance documentation – tape ensures secure storage for sensitive financial data, meeting regulatory requirements like GDPR and SOX. It offers a cost-effective solution for archiving large volumes of data and over a retention timeframe that can span decades, it has a much lower total cost of ownership compared to hard drives or solid-state drives, whether on-prem or in the cloud.

1.7.2

Life Sciences and Healthcare

In the rapidly evolving field of life sciences and bioinformatics, the volume of data generated by researchers and organizations is expanding at an unprecedented rate. A single cryogenic electron microscope can capture terabytes of data in only a few hours, and a published breakthrough on cancer research can take anywhere from a few years to a couple of decades. With burgeoning data in genomics and medical records, tape's scalability and security prove invaluable. It addresses storage needs while complying with data security and retention regulations.

1.7.3

Oil and Gas, Geo physical, and Geological Research

While geological data may not have the personal or individual dimension of other sectors, its commercial value in an era of scarce energy resource can often be immeasurable. Data security and regulatory compliance are paramount in these industries, where sensitive geological and drilling data, as well as environmental and safety records, needs to be protected. For vital geological and seismic data, tape ensures long-term preservation, accessibility, and security. Its role in disaster recovery is crucial, given the industry's susceptibility to disruptions.

1.7.4

Internet of Things and Manufacturing

Manufacturing organizations generate vast amounts of data, ranging from production logs and quality/defect control records to machine performance data. Similarly, IoT applications, including sensors and connected machinery, are producing a continuous stream of data related to machine performance, environmental conditions, energy consumption, and product quality, which must be preserved for predictive analysis and decision-making. Tape excels at low cost, long-term data preservation, making it an ideal choice for manufacturing and IoT applications.

1.7.5

Video Surveillance

By 2027, IDC predicts that data generated by computer vision devices, including video surveillance, is expected to make up almost 50% of the total number of exabytes in the Datasphere. Tape's role in storing immense video surveillance data becomes apparent due to its cost-efficiency, longevity, and security features. Its archival capacity surpasses traditional storage methods, addressing regulatory obligations and cost concerns. Tape encryption adds an extra layer of security, vital for regulatory compliance.

1.7.6

Retail and Transaction Processing

Retail and transaction processing operations generate substantial amounts of operational point-in-time data, including sales records, inventory data, and customer information. This data is essential for decision-making, analytics, and reporting and is often subject to stringent compliance and privacy legislation. Tape offers a practical and cost-effective solution for the long-term preservation of all this commercially sensitive material.

1.7.7
Media, Broadcasting, News & Publishing, Entertainment

Tape remains pivotal in preserving vast media content, ensuring its accessibility and integrity over decades. Its integration into content creation workflows and active archive strategies proves valuable for historical footage and digital assets. With a remarkable shelf life that can span decades and an ultra-low UBER for maximum reliability, tape technology ensures the integrity and accessibility of archived content over time. This is invaluable in an industry that depends on historical footage, audio recordings, and digital assets for content licensing, restoration, and compliance with copyright requirements.

1.7.8
Scientific Research and High-Performance Computing

Higher education and governmental research laboratories deploy vast computational resources in their fields of research and as a consequence, require similar proportions of ultra-reliable, highly secure, highly durable storage. Petaflop compute and petabyte storage requirements go hand in hand. Tape caters to the colossal storage demands of institutions dealing with massive data volumes, facilitating long-term, accessible, and cost-effective archiving.

In summary, tape technology, with its scalability, security, and cost-effectiveness, plays a critical role across industries by addressing their unique data storage, security, and regulatory compliance needs.

"At large scale capacity, tape has a very significant cost advantage over disk, although it requires some learning on how to operate and manage. The larger the scale of deployment, the more skeptical I am that labor savings by using disk could overcome the cost savings of tape The gap grows too large, and tape isn't that hard to use."
Jason Hick, LANL.

INDUSTRY	EXAMPLES
HPC / Scientific	Defense / Modeling Weather forecasting Particle physics AI/ML data and model retention
Cloud / Hyperscale	Traditional internal data archiving Archiving-as-a-service
Media & Entertainment	Streaming asset capture (e.g. 8k video) FILM & TV archiving and preservation
Healthcare	Regulatory records, video and imaging content retention
Life Sciences / Genomics	Backing up and protecting large data sets
Video Surveillance	Incident video retention, legal holds, time-based video retention
Traditional IT	Banking Finance Manufacturing
Government / Agencies	Library of Congress Streaming data collection and retention
Oil & Gas/Seismic	Near real-time collection of seismic data

Figure 15. Tape Industry Usage Examples
Source: INSIC

1.8

BENEFITS OF TAPE ACROSS IT WORKLOADS

Tape is used in many different industries to address broad and diverse needs, such as compliance regulations, data retention policies, growth of unstructured content, and cybersecurity. The need for long-term retention of data will continue to expand. In many cases, management of data value is complex, as the future value of an asset may not be well understood. Tape's low total cost of ownership, operational longevity and resilience helps organizations to address these and other concerns.

In most, if not all, cases, the value of tape is its relative low energy consumption and low total cost of ownership, especially compared to disk, for archive storage. There are other significant advantages to using tape, such as airgap to prevent data loss in a ransomware attack, ease of transporting cartridges, high volumetric density, and proven roadmap. These advantages are common across a number of different IT workloads.

1.8.1

Big Data / Analytics

Tape is often overlooked as a storage medium for Big Data analytics because applications and workflows that process and generate enormous data sets are often associated with high performance computing and storage. These applications typically involve complex mathematical calculations, machine learning, and deep learning tasks, which require from parallel processing, low latency, and real-time analytical capabilities.

But tape can support Big Data workflows, particularly in capacity-intensive processes that are focused on handling and managing large volumes of data. These processes are prevalent in batch processing, data warehousing, and long-term storage scenarios. They are characterized by sequential I/O, low sensitivity towards latency and throughput, and the processing of large, contiguous file sets. Further, results of analytics and processing can be staged to tape for long-term storage and retrieval, to avoid having to rerun processing tasks again, or to extend the data within the overall data set.

1.8.2

Archive

In simple terms, a typical active archive solution can integrate different storage technologies (flash, disk, tape), locations (on-prem, co-lo, public cloud) and architectures (block storage, file systems, object storage), so that data is stored in the most appropriate storage-class based on considerations such as ease of access, security, cost and scalability.

The key characteristic of "active" in active archiving is that the archived data is managed by an intelligent software layer that indexes and co-ordinates the placement of data across the different tiers. Crucially, the storage in use should be transparent to the application or user. There should be no need to update the application when data is migrated from a faster storage tier to a lower cost, slower tier, or from an older technology to its replacement. The primary benefit of incorporating an active archive is that organizations can use their primary storage, secondary storage and tape resources more efficiently and reduce the costs associated with long term data storage in an era of minimal data deletion.

1.8.3

Object Storage Active Archives and Tape

For longer term retention, object storage has emerged as a key component of active archive deployments and is worthy of additional consideration. Object storage has traditionally been associated with the cloud because virtually every large cloud vendor today uses object storage (the most common format is Amazon's Simple Storage Service or S3) as the underlying architecture for their infrastructure. However, the unstoppable growth of object data has changed assumptions about the storage technologies used to manage it, wherever it resides.

But although some businesses may still attempt to place all their object archive data on SATA disk-based devices, for the majority, data growth rates are exceeding the ability of traditional HDD technology (and physical rack space within data centers) to keep pace. And all-disk infrastructure is still many times more costly than a mixture of tape and disk.[26]

Faced with these pressures, an increasing number of organizations are attracted to the value proposition of tape technology for object storage. The LTO tape roadmap, for example, projects individual cartridge capacities of 1.4 PB per unit by LTO-14, which more or less matches data growth rate predictions for the coming decade. This gives organizations that use tape for object storage a significant cost of ownership advantage over disk and flash alternatives both in terms of the potential for massive storage density combined with ultra-low cost/GB.

Today, there are multiple different 'front end' enabling technologies that can utilize tape solutions to store data in native object format, (typically S3 Glacier class). This simplifies replicating and tiering data from disk-based object storage to tape.

1.8.4

Data Backup and Recovery

Data protection solutions can be defined as products and services designed to restore content that has been corrupted or lost. This includes both business continuity processes – for example day-to-day restores of lost or damaged files which have minimal or no effect on the wider organization – and disaster recovery, which includes urgent and / or massive restores of mission critical databases, storage arrays or even entire data centers, where the organization's survival is at stake.

1.8.5

Tape and Primary Storage

Tape can play an essential supporting role in both day-to-day and mission critical backup processes. Deploying tape allows organizations to offload data from primary storage to free up space for more valuable information and make most efficient use of their storage resources.

1.8.6

Tape and Secondary Storage

For data restores from snapshots – e.g. point in time representations of data – it is more common to use a 'secondary' storage array, which could be SSD or HDD-based. These arrays and appliances still provide relatively rapid restores of files but without the maximal performance of primary storage. Again, tape can play a supporting role here in offloading older snapshot data from capacity-constrained SSD or HDD arrays onto a less expensive, high capacity storage medium to extend RPO thresholds.

1.8.7

Tape and Cloud-based Data Protection and Backup

For clusters of small files or folders, and for relatively low volumes of data, cloud-based backup can be a more straighforward and practical solution because of its relative ease of recovery. Yet some drawbacks exist when storing large amounts of data in the cloud: the cost of the storage as datasets grow over time; the cost of retrieval, which is normally charged per gigabyte; and the time it takes to move very large datasets across the internet which can so severely impact recovery time objectives as to render the exercise impractical.

1.8.8

Disaster Recovery

Tape plays a role in DR solutions due to its portability, low cost, high reliability, high density, and offline characteristics. To support RPOs, snapshot technology can be used, and those snapshots can be copied to tape and protected offsite. Tape is particularly useful as an extension to secondary storage systems because it permits a much deeper backup history to be preserved as space on the physical array is filled with newer snapshot information. Of course, snapshots are not in themselves infallible, and they rely on the underlying storage system to be available. If the storage system is damaged or destroyed, snapshots might become unavailable or unusable.

A better approach, therefore, is to maintain periodic full backups on tape and keep copies both nearline in a tape library, for rapid accessibility, and offline and offsite in the event of total device or data center compromise.

1.8.9

Ransomware and Cyberthreats

In the ever-evolving landscape of cybersecurity, ransomware attacks have emerged as one of the most pervasive and damaging threats. A recent study found that in 93% of ransomware incidents, the threat actors target connected backup repositories, resulting in 75% of victims losing at least some of their backups during the attack, and more than one-third (39%) of backup repositories being completely lost. In late 2023, a notable and highly publicized ransomware attack on a European cloud provider saw the attackers successfully encrypt all production servers, as well as everything on the primary and secondary backup systems. Most customers lost 100% of their data as a consequence.

The very connectedness that is the strength of disk or cloud-based backup solutions makes them vulnerable in others. What makes tape a key weapon in the fight against ransomware is that it's the only truly offline storage solution that can place Enterprise data behind a physical, disconnected, air gap barrier. It is the final part of the 3-2-1-1 rule which proposes users should maintain three copies of their data, on at least two different media types, with one stored offsite and one stored offline.

1.8.10

Bulk Transfer

Tape has always been a removable and portable storage medium, and has often been used to transport data, typically for storing data offsite for added protection. Tape strong encryption of media protects critical data from loss and malicious intent. Multiple PBs of data can be transported with minimal cost and at speeds unmatched by data center connections.

Read more:
https://www.snia.org/sites/default/files/LTFS_PPT_FINAL.pdf

Read more:
https://www.snia.org/sites/default/files/DavidSlik_LTFS_Bulk_Transfe-R2.pdf

1.8.11

Data Regulations & Compliance

An important part of any data protection and management strategy is ensuring that business information is stored and managed in keeping with regulatory and compliance legislation. Data privacy and compliance regulations vary from country to country, but many share common principles and requirements, especially in the era of the General Data Protection Regulation (GDPR) in the European Union. The penalties for breaching these requirements can be extremely severe and apply both to individuals as well as organizations.

Similarly, a host of industries have strict obligations for maintaining robust data security measures to protect sensitive financial data from unauthorized access, manipulation, or deletion. Examples include: Sarbanes-Oxley Act, Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI DSS), the Basel III international regulatory framework, and finally the Securities and Exchange Commission Rule 17a-4.

1.8.12

Tape Encryption and WORM Immutability

Tape technology strongly assists data governance and compliance objectives. It offers enhanced security because information can easily be encrypted as it is being written to the tape, typically using native hardware-based encryption within the tape device itself. This means there is no need to invest in extra software or separate hardware to get the strongest protection for data. For even greater security, modern tape formats support data immutability via Write Once Read Many (WORM) recording media technology that can store data in a non-rewritable format.

1.9

SUSTAINABILITY

Defining sustainability in the general terms is a broad spectrum of impacts including economic, social, Governance and traditional carbon impacts human life has on earth. Carbon Dioxide is the most prevalent and most impacting carbon to the earth's continuation, thus minor impacting elements are expressed in the 'e'- equivalent. 'e' does not mean another element is less impacting as a measurement of weight or relative harm, only that the total is very low in comparison to CO2. This section focusses on impacts related to the carbon footprint of data storage products in the retention of data, often reflected as CO2e.

1.9.1

Tape Lifecycle

The lifecycle of products in the carbon protocols is best summed up by the greenhouse gas protocol diagram below.

Scope 1 and 2 are well defined and directly accounted to the product in evaluation. Upstream scope 3 are also very well defined by process and material. When combined these 2 lifecycle sections account for the product embedded carbon. In this section we will be discussing the third box pertaining to downstream scope 3 emissions, often accounting for 70-80% of all carbon impacts of a product, where "use" is 90% of this number[27].

The use phase is also a critical influence on end-of-life. Useful life of a product, specifically for products used to retain digital data is critical in the influence of carbon footprint. The longer the useful life of the product the less replacements of the infrastructure are required, lowering carbon impacts.

Tape media has a documented durable life of 30-years.[25] For data that must be retained and remains offline, tape is the only foreseeable solution to reduce all aspects of carbon footprint. Acknowledging the need for data access even for archives, tape as physical managed infrastructure, tape drives, automation, media has a "useful" life of 10-12 years. Compared to the average accepted useful lifecycle of HDD of 3-5 years. Tape's longer lifecycle reduces data management burden, carbon footprint, landfill waste, and recycling burdens. The following sections outline the details of the sustainability of tape.

1.9.2

Data Destruction

INSIC is in full support of the recommendations put forth by IEEE 2883 regarding data destruction methods[28] and NIST regarding data classification and disposition[29]. Due to litigation and risk assessment many large organizations, such as Hyperscale classified companies, have gone to an extreme to ensure data on physical media is not accessible, resulting in massive shredding and incineration of SSD, HDD and Tape media.

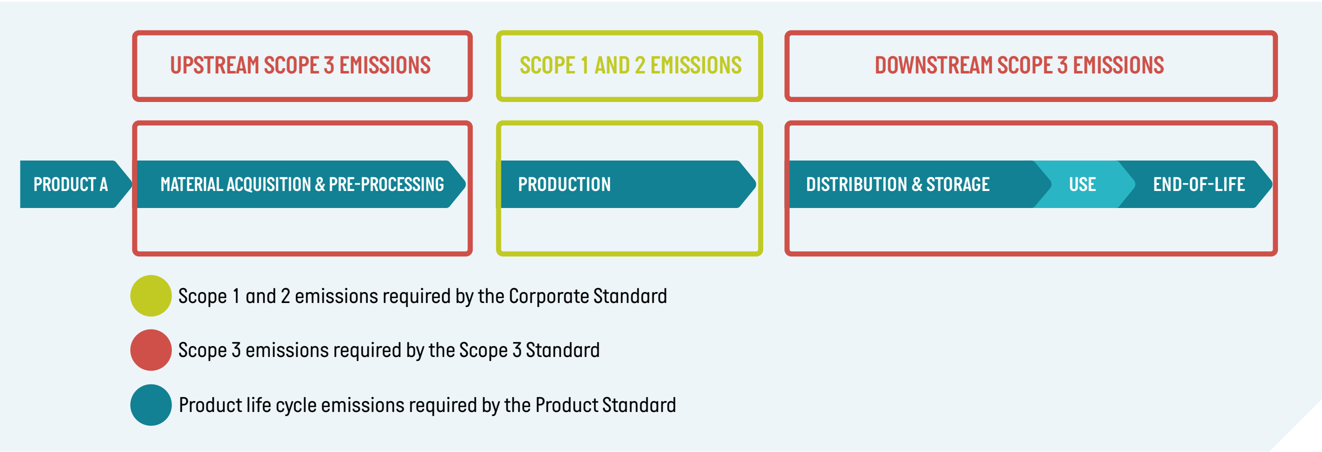

Figure 16. GHG Product life cycle diagram
Source: ghgprotocol.org[27]

INSIC recommends that all data at rest be strong encrypted on the media. Tape native encryption enables this capability with zero impact to performance or data reliability. Once the data is encrypted, deleting the keys associated with the data reasonably eliminates any access to the data, cryptographic erase, with near zero carbon impact for all the data in an infrastructure.

1.9.3
Media Disposal

Media disposal carries objectified risk to data exposure. Proper data destruction mitigates those risk, placing the focus of disposal on reducing e-waste and carbon impact. The goal of all digital media should be to contribute to a circular process as possible. While hyperscale and highly confidential users have focused on shredding and incineration of the media, both these methods have high carbon impacts. Shredding of HDD results in contamination of all the components of the HDD resulting in e-waste. Incineration has the highest impact to the environment. In many cases incineration of digital components results in up to 3 times more CO2e than burn material weight.[30]

98% of tape components are recyclable with no contamination risk or high-energy shredding required, less than 2% by weight results in e-waste. Tape cartridges are 95% landfill safe and are 99% recyclable materials by weight. The cartridge memory is the only e-waste component. The recyclability of tape is dependent on modernized processing of polyester materials.[30]

Improving the benefit of the circular economy is a corporate responsibility, Utilization of products with nature contributions is critical to achieving NetZero goals.

1.9.4
Carbon Impact

Determining and reporting of carbon impact is complicated by the variability of any specific solution for data storage. A result of this variability in the reporting of CO2e is that fundamental assumptions must be made in any comparison. The comparison in this section utilizes 27PB, which is within the upper 35 percentile by data capacity at rest as determined in the recent ESG Research survey commissioned by the Active Archive Alliance[31]. Other assumptions include a 10-year data life, 5-year HDD refresh cycle, HDD Open Compute Project architecture.[32]

Proper storage data placement optimizes performance of solutions and as demonstrated in the image above, has the potential to reduce CO2e and direct energy up to 97%. According to Furthur Market Research and Brad Johns Consulting, data stored on tape that is airgap or long-term deep archive has the potential to reduce CO2e by up to 97% and up to 99% reduction of energy consumption.

While this data is representative of a specific infrastructure, the tape industry has provided specific documentation to enable organizations to measure and report their specific configuration with reasonable accuracy. Moving the 55-65% of non-accessed retained data to tape is a carbon improvement for any organization.

97% CO2E REDUCTION USING TAPE

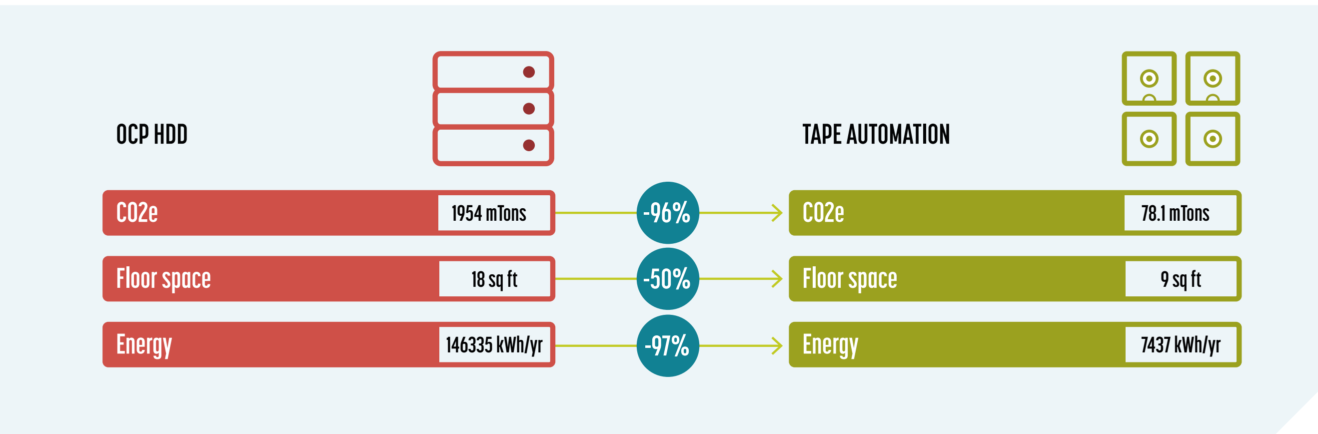

Figure 17. IBM Physical Tape Carbon impact Compare
Source: IBM

1.10

CONCLUSIONS AND RECOMMENDATIONS

A dominant influence on the future market demand for tape storage is the accelerating growth of data that individuals and organizations are generating and retaining. The requirement to preserve such data carries with it the need for the availability of storage technologies with ever improving density, performance, increased capacity, and lower cost. Estimates of the worldwide capacity for archival storage of digital data exceed ten zettabytes (1 zettabyte = 1021 bytes) in the next few years. Therefore, the need to increase available storage capacity will remain with us for the foreseeable future. The specifications for this increased storage range widely, depending on the attributes that must be optimized in each solution (e.g., cost, performance, data reliability, data security, regulatory compliance, disaster recovery, etc.)

The use of tape technology to store data has undergone a transition. Historically, the most significant use of tape has been to provide the backup and restore functions, i.e., to ensure that system hardware or software problems do not lead to a loss of data. Over the past several years, it has become recognized that much of the data collected for the backup and restore functions is needed only for short-term retention, and that the ability to search for and rapidly restore small amounts of data, such as a single office document, is a valuable attribute. These aspects have made it attractive to adopt disk-based solutions for backup and restore, and their use for this purpose has been increasing, while the use of tape for backup and restore has diminished. However, the increase in cyber security requirements has led best practice enterprise data protection to include an offline copy that is often stored offsite, for which tape is ideal. Many users now employ both disk and tape in a tiered solution, whereby most file restores come back from disk for performance reasons, but tape provides the last line of defense against a larger data loss or security breach.

It is now recognized that a very important application of storage is for long-term preservation (or archiving) of data. It is especially in the archival context that the use of tape has remained strong and is projected to grow significantly, offsetting the decline in backup.

- Tape has several attributes which favor it over disk for long-term archival use, including its considerably lower cost (both acquisition and operating costs) and its higher data reliability. An important contributing factor to tape's lower cost is its significantly lower power consumption compared to disk. Energy saving in the data center environment has taken on an importance of its own, beyond strictly cost implications, because data centers are now concerned about the availability of sufficient power for their future operations.

- For the above reasons and more, the use of tape continues to be an important element in most enterprise IT operations. The most common approach for enterprise storage today is not to deliberate whether to use flash or disk or tape, but instead how to best use flash, disk, and tape in an integrated approach for both data protection and archive, taking advantage of what each storage technology offers. Software that presents a standard interface, such as NFS, CIFS, or OpenStack Swift, has enabled tiered storage solutions that are transparent to the applications. There are several industry segments where archival storage on tape has become a strategic part of the operation, of which several examples have been detailed in this report.

- A further ongoing development in the IT environment is the growth of public cloud storage, which continues to accelerate. Cloud storage services are offered with a variety of business models and pricing schemes. Many of these services are highly cost sensitive and/or critically dependent on maintaining security and integrity of the stored data. Tape's fundamental attributes, especially low cost, high data reliability, and transparent encryption, make it ideal for its use in cloud storage offerings.

- We expect that tape will continue to play a strong role in future data storage environments; however, that will require ongoing attention to advancing tape technology and its ease of use. First and foremost, the research and development to support aggressive tape density and capacity advances must continue, for tape to maintain its significant cost advantage over disk. Please refer to section 2 for details on the 10-year roadmap, which shows robust improvements in capacity resulting in lower TCO to maintain that ost advantage. To support increased use of tape for archival storage, advances in tape data organization and associated supporting software, that enable easy access to specific data records in the archive, must continue to be developed. Continuation of software products that migrate data to new generations of hardware have benefited the adoption and use of tape and they must continue. The development of tape drive and media products that extend backward compatibility to two or three generation have also helped reduce the need for frequent migrations and helped tape's adoption for large digital archives and must continue going forward. Finally, although the outlook for tape presented here is very positive, the industry should strive to better educate customers and better publicize the inherent and significant advantages of tape storage.

As cloud and large digital archives, such as HPC and video, continue to grow rapidly, we expect that the declines in the overall tape market are behind us and now project a flat market (by revenue) going forward for tape

Phil Goodwin - IDC

Acknowledgement

The Information Storage Industry Consortium wishes to express its appreciation to the team members from the tape storage industry, and their companies, who came together to gather and discuss the available data, and to craft the messages and the accompanying narrative in this Applications and Systems Roadmap. We would also like to thank the industry analysts and tape storage application specialists who participated with us to share their data and unique perspectives on the future of tape storage.

About

INSIC
Information Storage Industry Consortium

The role of the Information Storage Industry Consortium (INSIC) is to serve as a resource for anyone interested in current and future trends in digital magnetic tape recording technologies and platforms, which are an important component of the data storage landscape. A primary goal is to make this information easily accessible and valuable for existing and prospective tape users. Participants today include corporations involved in creating and producing tape recording hardware, media, systems and infrastructures. INSIC was originally incorporated in April 1991 as NSIC, a nonprofit mutual benefit corporation, and is currently located in Monroe, Virginia, USA.

Read more:
http://www.insic.org/

1.12
REFERENCES

[1] https://unfccc.int/news/ict-sector-helping-to-tackle-climate-change

[2] https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks#

[3] Johns, B. (2012, February). Improving information Technology Sustainability with Modern Tape Storage.

[4] Furthur Market Research, Monroe, J., Johns, B. (2024, February). The Sustainable Preservation of Enterprise Data Report

[5] Veeam "Ransomware Trends Survey", 2023

[6] IDC Global StorageSphere 2023-2027 Forecast, May 2023

[7] Further Research "Storage Management in an Age of Minimal Data Deletion", May 2023

[8] Swan, E. (2018, April 23). Document Retention Requirements Per Industry. Updates from https://www.crownrms.com/insights/your-record-retention-cheat-sheet-a-comparative-look-from-healthcare-to-government/

[9] Moore, F. (n.d.) What is the tape Air Gap? Retrieved from video https://datastorage-na.fujifilm.com/tape-air-gap-provides-defense-against-cybercrime/

[10] Gartner, (2023) Magic Quadrant for Distributed File Systems and Object Storage reference# 376728

[11] UBER Paper: https://www.lto.org/wp-content/uploads/2022/08/LTO-UBER-Technical-Paper-August-2022.pdf

[12] Rise of secondary Storage and Role of Active Archive with discussion on RAIL with Erasure Coded Tapes https://horison.com/publications/active-archive-webinar-the-rise-of-secondary-storage-and-the-role-of-active-archive/

[13] Brume, S (2024, February), IBM 2024 External Protocol Report

[14] IDC, (2020) Worldwide Enterprise Storage Systems Installed Base Forecast, 2020-2024 reference#US46726820

[15] The Economist. (2018, August 18). The race is on to dominate quantum computing. Retrieved from https://www.economist.com/business/2018/08/18/the-race-is-on-to-dominate-quantum-computing

[16] Cisco Systems. (n.d.). Reassessing Grover's Algorithm. Retrieved from https://eprint.iacr.org/2017/811.pdf

[17] ResearchGate (2017, April). The AES-256 Cryptosystem Resists Quantum Attacks. Retrieved from https://www.researchgate.net/publication/316284124_The_AES-256_Cryptosystem_Resists_Quantum_Attacks

[18] Bernstein, D. (2018, August 18). Introduction to post-quantum cryptograph. Originally retrieved from http://hyperelliptic.org/tanja/vortraege/20220825-qsancus.pdf"http://hyperelliptic.org/tanja//vortraege/20220825-qsancus.pdf

[19] Mosca, M., Piani, M. (2022, December) Quantum Threat timeline Report 2022 https://globalriskinstitute.org/publication/2022-quantum-threat-timeline-report/

[20] Pinheiro, E. (n.d.). Failure Trends in a Large Disk Drive Population. Originally retrieved from http://static.googleusercontent.com/media/research.google.com/en//archive/disk_failures.pdf

[21] Wikipedia, pulled from: https://en.wikipedia.org/wiki/Annualized_failure_rate#:~:text=Annualized%20failure%20rate%20(AFR)%20gives,devices%20are%20run%20per%20year

[22] Seagate, pulled from: http://enterprise.media.seagate.com/2010/04/inside-it-storage/diving-into-mtbf-and-afr-storage-reliability-specs-explained/

[23] Backblaze 2023 Hard Drive Reliability Statistics (2024) pulled from: https://www.servethehome.com/backblaze-releases-2023-hard-drive-reliability-stats-for-35-

[24] Arsian,2. Peng, j., Goker, T. (2019, October) A Data-Assisted Reliability Model for Carrier-Assisted Cold Data Storage Systems https://asset.fujifilm.com/www/us/files/2024-04/69ca08096529949d3aa09b45494353af/FF_Care_Handling_2024_Update.pdf

[25] Fuji technical paper on tape archival life; https://asset.fujifilm.com/www/us/files/2024-04/69ca08096529949d3aa09b45494353af/FF_Care_Handling_2024_Update.pdf

[26] Solutions North Consulting, "Data Economics: LTO-9 Technology Optimizing Data Value", (May 2023)

[27] GHG Product Lifecycle Accounting Reporting Standard (2024) https://ghgprotocol.org/sites/default/files/standards/Product-Life-Cycle-Accounting-Reporting-Standard_041613.pdf

[28] IEEE 2883-2022 (2022), IEEE Standard for Sanitizing Storage https://standards.ieee.org/ieee/2883/10277/

[29] NIST SP 800-88 rev 1 (2014, December), Guidelines for Media Sanitization https://csrc.nist.gov/publications/detail/sp/800-88/rev-1/final

[30] Brume, S (2023, November), IBM Physical Tape Media Disposition Report https://ibm.seismic.com/Link/Content/DCbae1c674-42c6-4d42-8203-65e929a6416e

[31] ESG Research_Active Archive Alliance 2023 Intelligent Active Archive Study_Key Findings (January, 2024)

[32] Brume, S (2023, November), IBM Physical Tape Carbon impact Compare 2024 https://www.ibm.com/downloads/cas/YE5WAQ0B

[33] SNIA TWG (2023, October 27), Storage Data Placement TWG Proposal. Originally taken from: https%3a%2f%2fmembers.snia.org%2fdocument%2fdl%2f52007/RK=2/RS=vf5f4m_FSzZr2zlR6c0RqIWTdbA-

[34] RAIT https://hpss-collaboration.org/rait/

[35] RAO and ECC Tape https://tapestorage.org/wp-content/uploads/Tape_Storage_Council_2020_December.pdf

[36] Rise of secondary Storage and Role of Active Archive with discussion on RAIL with Erasure Coded Tapes https://horison.com/publications/active-archive-webinar-the-rise-of-secondary-storage-and-the-role-of-active-archive/

[37] Ultrium LTO (n.d.). Are you aware of the hidden costs of storing data in the cloud? Retrieved from https://lto-calculator.hpe-media02.bitwise.nl/pdf/LTO_TCO_TOOL_USER_GUIDE.pdf?v3

[38] https://blocksandfiles.com/2021/08/31/seeing-the-archive-light-folio-photonics-hopes-to-crack-the-optical-archive-disk-market-wide-open/