TAPE ROADMAP

INSIC REPORT

2019

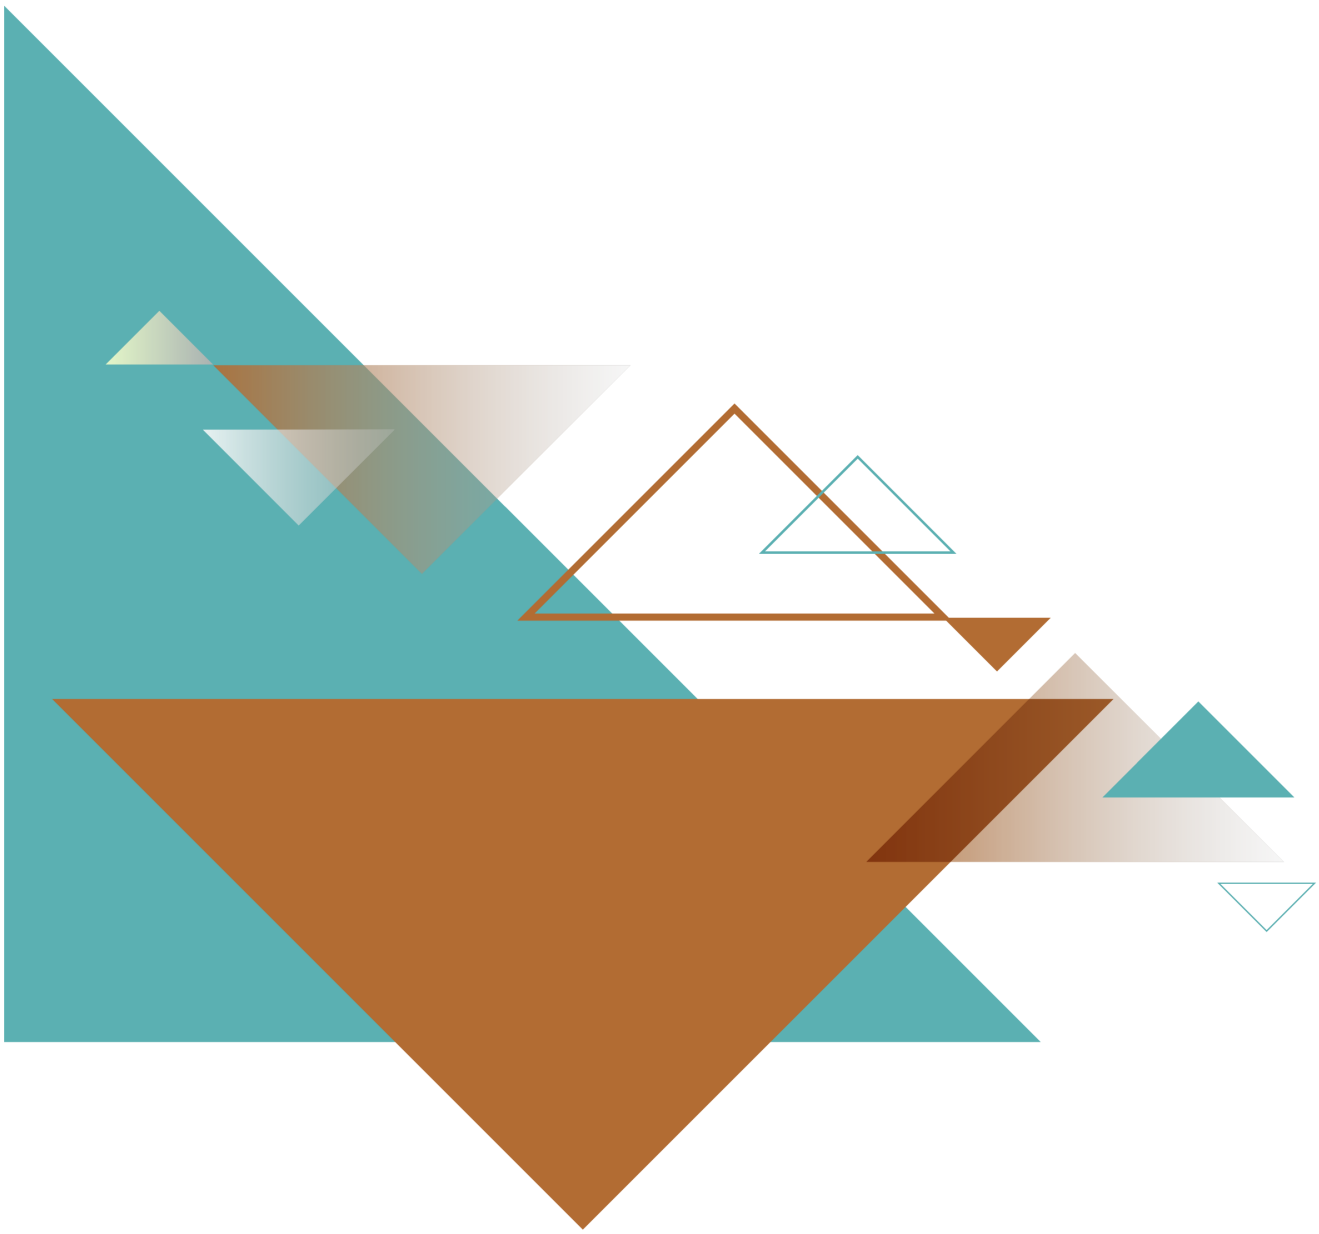

1.0 | 2019 APPLICATIONS & SYSTEMS TAPE ROADMAP

Team Director

Dave Woito, INSIC

Contributors

Carlos Sandoval, IBM
Julio Hernandez, IBM
Sandra Peláez, IBM
Mark Hill, IBM
Turguy Goker, Quantum
Molly Presley, Quantum
Ryan Duffy, Quantum
Richard Bickers, Hewlett Packard Enterprise
Carl Hoerger, Hewlett Packard Enterprise
Laura Loredo, Hewlett Packard Enterprise
Tom Wultich, Oracle
Robert Raymond, Oracle
Christine Rogers, Oracle
Jon Benson, Spectra Logic
David Feller, Spectra Logic
Susan Merriman, Spectra Logic
Matt Ninesling, Spectra Logic

Quoted Industry Experts

Tom Coughlin, Coughlin Associates
Phil Goodwin, IDC
Jason Hick, Los Alamos National Lab
Fred Moore, Horison
Mark Peters, Enterprise Strategy Group
Dave Thomson, QStar Technologies

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Contents

1.0 | 2019 APPLICATIONS & SYSTEMS TAPE ROADMAP 3
1.1 | EXECUTIVE SUMMARY 6
1.1.1 | Notable quotes from INSIC interviews 7
1.2 | MARKET TRENDS 8
1.2.1 | Green Data Center 8
1.2.3 | Cold Data Storage 10
1.2.4 | Big Data / Analytics 10
1.2.5 | Video Content 11
1.2.6 | Cloud 12
1.2.7 | Longer Retention Times 14
1.2.8 | Ransomware 15
1.2.9 | GDPR (Achieve privacy on Tape) 16
1.3 | TECHNOLOGY TRENDS 17
1.3.1 | Competing Technologies 17
1.3.2 | Emerging Technologies 18
1.3.3 | Tiered Storage with LTFS 19
1.3.4 | Object Interface in Open Source 21
1.3.5 | Tape Encryption and Quantum Computing 22
1.3.6 | Erasure Coding & Tail Latency 23
1.3.7 | Software Defined Storage 27
1.3.8 | Host interface 27
1.4 | COST & ENERGY CONSUMPTION ADVANTAGES (TCO) 28
1.4.1 | Introduction 28
1.4.2 | ESG TCO Study 28
1.4.3 | TCO Calculators 30
1.4.4 | Modeling Methodology 31
1.4.5 | Energy Consumption of Storage 31
1.4.6 | Trends in Storage Costs 32
1.4.7 | TCO Conclusions 33
1.5 | THROUGHPUT AND RELIABILITY 33
1.5.1 | Data Rate 33
1.5.2 | Durability of the Data 35
1.5.3 | Media Wear Out 44

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.5.4 | Roadmap Benefits from the Changing Usage Model 45
1.5.5 | Not Compromising Time to First Byte 47
1.6 | CHALLENGES FOR TAPE ADOPTION 48
1.7 | INDUSTRY USE CASES FOR TAPE 49
1.7.1 | Big Data / Analytics 49
1.7.2 | Archive 50
1.7.2.1 | Data Migration 51
1.7.3 | Backup and Recovery 53
1.7.4 | Cold Storage 54
1.7.5 | HSM / Tiering 55
1.7.6 | Bulk Transfer 56
1.7.7 | Cloud 57
1.7.8 | Disaster Recovery / Air Gap Protection 59
1.8 | INDUSTRIES WHERE TAPE IS USED 60
1.8.1 | Overview 60
1.8.2 | HPC/Scientific 61
1.8.3 | Cloud 61
1.8.4 | Media & Entertainment 62
1.8.5 | Healthcare 62
1.8.6 | Video Surveillance 63
1.8.7 | Traditional IT 64
1.8.8 | Oil & Gas/Seismic 64
1.9 | CONCLUSIONS AND RECOMMENDATIONS 65
1.10 | ACKNOWLEDGEMENT 67
1.11 | REFERENCES 69

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.1 | EXECUTIVE SUMMARY

The Information Storage Industry Consortium (INSIC) 2019 roadmap contains quite a few updates 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 2019 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. In other words, advancing tape technology in the future faces much less technical risk than disk technology.
- There is a general shift of tape use from short-term backup towards long-term archive. Tape is playing a key role in some of the largest, high-growth archive use cases, which are highlighted in this report.
- This roadmap now incorporates technology advances such as tape tension control and tension effects on Tape Dimensional Stability. This new capability is very important to enabling higher tape capacities.
- For the first time, the roadmap includes a projected scaling rate for the Uncorrectable Bit Error Rate of tape. Providing this information allows for easier comparison to other storage technologies and highlights a very significant advantage over disk technology.
- At the end of year 2029, the tape roadmap still has a bit aspect ratio of about 2 times where disk is today, providing significant margin for future advancement of tape technology.

The inherent, fundamental advantages of tape continue to be its low acquisition cost, extremely low power consumption and cooling requirements, excellent footprint, scalability, and of course, reliability. Other attributes that can be significant are tape's offline data protection and portability. The dramatic increase in malware / ransomware in the market has driven more demand for offline storage, including tape, which provides significant protection from cyberattacks. Tape's future is driven by the needs of multiple markets and use cases where these attributes are important. In short, tape is best for applications that do not require fast access, so in essence, the opposite end of the spectrum of flash applications.

Tape has made significant progress in multiple, very large, high-growth archive markets. While these markets have very different applications that create and access the data, they all present very large-scale, high-growth storage requirements where tape provides a huge advantage over any other technology today or in the foreseeable future. Solutions typically involve tiered storage that includes flash and/or disk along with tape, but tape is the key technology of choice for vast amounts of storage for each of these markets as evidenced in this report.

Much continues to be written about the other end of the spectrum where high IOPs and cost per IOP is king, and where flash adoption is growing rapidly and displacing disk. However, the

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

archive market segments identified above present some of the largest and fastest growing markets for any storage technology. Tape is well positioned for each of those markets today and for the future.

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/category/blogbytes) for timely information.

As mentioned above, interviews with analysts and experts were conducted as research for this roadmap. Some notable quotes are listed below and provide an overview of key points in the INSIC 2019 roadmap.

1.1.1 | Notable quotes from INSIC interviews

"Tape is the most affordable option for multi-PB storage capacity."
– Jason Hick, Los Alamos National Laboratory

"In the future, I believe tape volumes will dramatically increase because of the price/performance/capacity/reliability and roadmap advantages."
– Fred Moore, Horison Information Strategies

"Tape's offline capability gives us a layer of protection that the other storage options do not have."
– Jason Hick, Los Alamos National Laboratory

"Five years out, tape is still going to be quite relevant. There is going to be more data than ever that needs to be stored, and they'll be looking for cost-effective ways of doing that."
– Tom Coughlin, Coughlin Associates

"Storing data in the cloud is inexpensive; moving data in and out of the cloud is very expensive. Not so with tape."
– Fred Moore, Horison Information Strategies

"Most higher capacity users prefer on-premise archives."
– Dave Thomson, QStar Technologies

"Tape backup operations is becoming a part-time job due to the degree of automation available. Disk management has its own amount of labor associated with it."
– Phil Goodwin, IDC

"Most of the issues with tape have to do with perception rather than reality. The challenge is really fiction, not function."
– Mark Peters, ESG

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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]

IT consumes about 7 percent of global electricity, forecasted to rise to 13 percent by 2030.[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. For a discussion on power consumption, please see: https://www.lto.org/wp-content/uploads/2014/06/Data-Storage-Technology-Tape-and-Sustainability-Sept-2017.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 the Enterprise Strategy Group concluded that a disk archive is an order of magnitude more than the power of a tape-based solution. Tape's footprint is also reduced, minimizing the square footage required. The same study showed a footprint savings with tape of about 10 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, IBM and Google collectively control more than half of the WW

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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.
- 2016 model automobiles averaged 14 million lines of code. Autonomous automobiles may reach 100 million.
- 1.2 trillion photos in 2017 – 9 percent CAGR (compound annual growth rate).
- Surveillance = 2.5 EB/day in 2019.
- The IoT is forecasted to connect over 25 billion "things" as Internet nodes by 2020.
- IoT has 6.4 million developers.
- World's largest data center in Chicago is 1.1M sq. ft.
- By 2020, One third of all data will pass through the cloud.
- Average Server Lifespan - 3 years; Hard Disk Drive (HDD) 4-5 years, Tape Drive 7-10 years.

According to IDC the amount of data being stored is exploding, and it has been and will continue to do so for years to come. IDC estimates that the global datasphere will exceed 175 ZB by 2025, with a 30 percent CAGR.[4]

Figure 1 illustrates this growth in data.

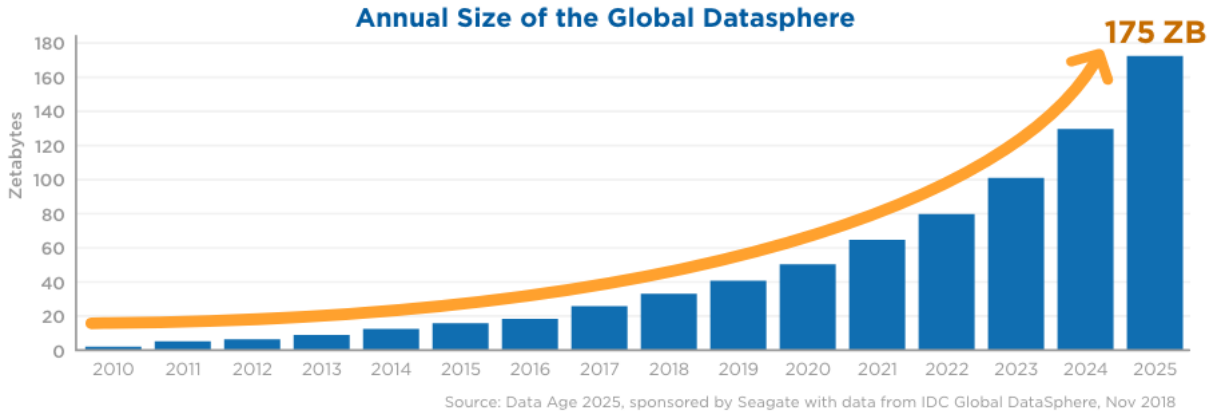

Figure 1 | Annual size of the global datasphere

Source: IDC

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Figure 2 shows a reference for understanding data measurements in this exponential data growth environment.

Figure 2 | Reference for data measurements

Source: Moore, F. (2017, October). Storage Outlook. Retrieved from https://www.lto.org/wp-content/uploads/2018/10/LTO-MGM-63-Fred-Moore.pdf?rs_file_key=13488642735c23d7ec65fb9608114877

1.2.3 | Cold Data Storage

As discussed in a previous section, 60 percent of the total datasphere is neither frequently accessed nor does it require rapid access. This data can be referred to as "cold." Therefore, studies propose that approximately 2.25 ZB of data is eligible for the lowest cost media available, and IDC believes that cold storage in the future will be an integrated strategy incorporating both tape and cloud.

Given the focus on cost and the de-emphasis of access time, both onsite tape and cloud archive tiers are excellent solutions for cold data storage. Google's Nearline and Microsoft Azure are two prominent cloud offerings targeting cold data storage.

Cold data storage is ultra-low-cost storage for data that is very rarely (if ever) accessed. Simply put, cost is paramount and access time is not as important. For example, data held for purely legal compliance reasons falls into this category. Much of the data captured and stored by national labs, such as the National Energy Research Scientific Computing Center (NERSC), also fall into this category. Researchers Wayne Hurlbert and Damian Hazen stated the following regarding the case for archiving and long-term data retention. "We need to preserve this data; it's not infrequent for researchers to need to go back to their (old) data sets, often to use modern techniques to reanalyze. The archive allows us to safeguard irreplaceable data generated over decades; the data represents millions of dollars of investment in computational hardware and immense time, effort, and scientific results from researchers around the world."[5]

1.2.4 | Big Data / Analytics

Information has always been important to society, but the amount of data, both structured and unstructured, now collected, shared, analyzed, stored and leveraged is becoming so massive it has to be managed using new methodologies. This is big data and our increasing ability to make use of it is profoundly changing the world in which we live.

As said previously, IDC forecasts that the global datasphere will grow to 175 zettabytes by 2025, more than 10 times the amount of data created just a few years ago.[6]

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Emerging trends, such as IoT applications, predictive analytics, sensor networks, artificial intelligence, smart cities, automated customer service, and connected cars will create even more data that will deliver business insights and benefit mankind.

The applications for Big Data appear to be limitless in today's data-driven world. For example, healthcare analytics has the potential to save lives, reduce treatment costs, predict epidemics and improve the quality of life for much of humanity. Big Data in law enforcement enables a precinct in California to check a nationwide database for fingerprints, DNA and other forensic information captured in New York. Big Data analytics can also analyze past criminal behaviors, locations and crimes to generate 'hotspot' mapping. A large California Police Department, for example, collects and analyzes crime data that indicates if another illegal offense will occur in the same area, enabling police to proactively beef up security in that zone to thwart crimes.

With a seismic sea wave of data overwhelming IT infrastructures, and volumes swelling year over year, many organizations are re-examining their data storage strategies. Cloud is one platform for storing and analyzing Big Data. But, because much of Big Data does not need to be accessed instantly, tape is regarded as the most economical way to preserve this information for the long term. Big Data 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 Content

The media and entertainment ecosystem is changing rapidly, with enormous amounts of content being created by television and film as well as by disruptive trends, such as virtual reality, augmented reality, live streaming, social media, video games, global sports broadcasts, live entertainment and video streaming services. Greater camera resolutions, such as 4k and 8k formats, smart TVs, smart phones, tablets, laptops and all types of digital devices are all contributing to the tectonic shift. This growth in content, especially video, is showing no signs of slowing down. For example, more than 400 hours of video content is uploaded to YouTube every minute with 1.5 billion regular monthly users. And, the average person will spend 84 minutes a day watching videos online by 2020, according to Zenith's annual Online Video Forecast Report.[7]

In general video format terms, 4K is four times bigger than High Definition (HD) but the data size will be 12 times larger in comparison between (XDCAM 422 HD 50Mbps.) If the same multiplier is applied between 4K and 8K, 8K from HD will be calculated 144 times in data size. FLAPE (combined flash and tape) infrastructure is key in this environment, as flash is the most suitable storage tier supporting 4K and 8K workloads. At the same time, data can be offloaded to tape, moving it from a higher cost storage tier to the lowest cost archiving tier for long-term storage.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

The advent of 3-D technology fosters more storage requirements, and future plans include customer-selected point of view, which requires that the content from multiple cameras can be accessible, and 360-degree video takes this even further.

This explosion in digital content is being created and leveraged by broadcasters, post-production houses, sports organizations, houses of worship, universities, research organizations, and social media users, among many others. Organizations find themselves needing to preserve this content for future usage (such as when news outlets and sports organizations replay historical footage on an event anniversary or during a game). Many organizations require a hybrid solution that allows them to use cloud services to meet certain operational needs, such as quickly distributing working files across studios in different locations, while leveraging tape technology to affordably store massive amounts of content locally and mitigate the rising costs of cloud. With new technologies such as asset management application systems, tape has become more accessible as an efficient and economical tier of content storage. These applications allow users to easily access both partial and full files from storage. With tape fully integrated, users can directly access content on tape from the asset management GUI, which delivers substantial time and cost savings.

Video surveillance is another large growth market. Many industries deploy video surveillance for a variety of reasons today, including greater security, activity monitoring, evidence collection and legal protection. The introduction of body-worn cameras for law enforcement as well as high-profile mass shootings have heightened awareness regarding the value of video footage. At many medical marijuana dispensaries, video surveillance cameras are mandatory. Moreover, many U.S. cities have begun to stipulate that video surveillance footage be kept for a minimum of 30 days, but more frequently to be kept as long as two to five years. While video resolutions have advanced considerably, the focus on how to retain the growing amount of footage has lagged. As surveillance video assets multiply and cameras with higher resolutions require more storage capacity, new tiered storage strategies will be necessary for organizations to meet budgetary and operational directives.

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-premise 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 the industry is still in the early phases of cloud adoption. The majority of organizations going to the cloud are combining some form of cloud with on-premise compute/storage for the most effective, economical and secure data protection.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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 is growing interest in 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.

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 fairly 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 Gbps 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-8 tape technology is 3 times the performance of a dedicated 1 Gbps connection. A single LTO-8 tape drive, at native transfer rate of 360MBs, would take the 1 PB restoration from 3 months down to one month. A bank of 10 LTO-8 drives would take the restoration down to three days. That’s using native transfer rates. Compression rates vary, but could add up to an additional 2.5 times in performance over those numbers. The LTO-8 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. And cost is another reason “Cloud Storage as a Service” doesn’t make sense for large data sets being created outside of the cloud.

Data recovery from the cloud is based on total capacity of files being stored, total number of files being stored, retrieval and access costs, transfer and transfer acceleration costs and replication costs. Depending on the amount of data to be retrieved, costs can climb sharply. Restoration costs are often 4 times to 10 times the cost of storing the data. While monthly charges per GB to store data can be incredibly low, those numbers are often deceptive when stated as a “monthly per GB” charge. At .004 ¢ per GB/month, the 1 PB data set referenced above will cost over $700,000 to store for 15 years not including any restoration charges or bandwidth charges.[8]

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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-premise 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-premise operations.

1.2.7 | Longer Retention Times

Drivers for longer data retention include scientific research, healthcare, litigation, public safety and national security. 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 health care 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).

The growing use of analytics is driving the need to keep more data longer. In order to gain a competitive advantage, businesses are reviewing historical trends. Use of tape allows for storing data for a longer period of time, giving better accuracy and longer-term views on models and trends.

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.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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 also at the file level with industry standard hash codes incorporated into the archive storage software. Tape also 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 | Ransomware

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 2017, the FBI's internet Crime Complaint Center received 1783 ransomware complaints, which cost victims over $2.3 million. This number does not factor in the hundreds of organizations that paid ransoms without reporting it for privacy purposes. In its annual Internet Crime Report,[10] the bureau's Internet Crime Complaint Center reviewed 301,580 complaints in 2017, and calculated that losses from various cybercrime schemes exceeded $1.41 billion in the U.S. that year.

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.[11]

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.2.9 | GDPR (Achieve privacy on Tape)

The General Data Protection Regulation (GDPR) is legislation developed to strengthen and unify data privacy protection for natural persons within the European Union. It can impact any business, regardless of location, that processes European individual personal data. Non-compliance could mean fines of up to the greater of €20 million Euros or 4 percent of the company's annual worldwide turnover for a corporate group.

The primary objective of the GDPR is to provide rules relating to the protection of natural persons with regard to the processing of personal data and rules relating to the free movement of personal data. This objective comes at a time of unprecedented creation of digital information with more and more of this personal data being processed than ever before.

The key principles of GDPR relating to the processing of personal data are lawfulness, fairness, transparency, accuracy, integrity, confidentiality and accountability, as well as ensuring purpose and storage are limited and data minimized. The definition of what constitutes personal data has also been considerably widened in the GDPR to include IP addresses; genetic, health and biometric data; as well as data types such as Global Positioning System (GPS); media access control (MAC) addresses; unique mobile device identifiers (UDID); and International Mobile Equipment IDs (IMEI).

In essence, the GDPR means that safeguards for personal data will have to be designed into the very fabric of how personal data is captured, managed, used and stored during its lifecycle. While this might appear challenging, for many organizations GDPR is actually an opportunity to gain better insight into where personal data is gathered, used and stored in order to assess the robustness of how it is protected.

GDPR is explicitly "technology neutral" so that the protection of individuals' rights in respect to their personal data is not dependent on the techniques used. In other words, the GDPR does not recommend any single storage technology as a means of compliance. It is up to organizations to determine that the technology they use to process and store personal data complies with the GDPR. There is no "one size fits all" storage platform or technique that will enable companies to simply tick a box marked "GDPR compliant." Some storage technologies fit certain aspects of GDPR better than others. And no two businesses are alike. Enterprises will probably need to evaluate their options and use a mix of solutions depending on which aspect of compliance they are trying to address.

Encryption—a key feature of LTO technology—is one of only a few techniques specifically mentioned by the GDPR in the context of data protection and security. And when it comes to encryption, tape is a highly efficient and effective data protection solution.

Another core benefit of tape technology relevant to the security of processing considerations is the fact that tape is essentially an offline storage medium. Once removed from a tape drive or library host, cartridges can be placed in a controlled and secure facility to reduce the risk of data loss or disruption caused by cyberattack, malware or other hostile intent. This core feature of

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

LTO tape technology may be useful for companies facing the challenges of Article 32, in particular:

- The ability to ensure the ongoing confidentiality, integrity, availability and resilience of processing systems and services;
- The ability to restore the availability and access to personal data in a timely manner in the event of a physical or technical incident;

Deployment of tape, therefore, may greatly increase the resilience of an organization to recover personal data in the event of a physical or technical incident and lessen the potential risk of GDPR non-compliance.

Read the full whitepaper on "How tape can support a GDPR compliant data processing system" here.

1.3 | TECHNOLOGY TRENDS

1.3.1 | Competing Technologies

Since the 2015 INSIC roadmap, there have been some changes in the storage industry. Solid State Disk (SSD) has pushed further with 3D skyscrapers being built to hold more cells for flash NAND, as well as further exploration of more voltage levels in the same cell so that bit density can increase. Additional investment is being made in manufacturing fabrication to increase the amount of flash NAND that can be made every year, and even HDD companies are partnering or merging with flash suppliers. SSD continues to be the growth portion of the storage market.[12]

In the HDD world, Seagate R&D engineers and scientists continue to work toward making the practical application of HAMR (Heat Assisted Magnetic Recording) a marketplace reality. As of this roadmap, HAMR HDDs are not out yet, but Seagate continues to invest significant resources to make it happen.

Seagate also has a team of people who are looking to apply their HDD products and technology in new ways. Seagate released an ethernet-attached HDD and continues to look for new ways of integrating their technology in places in the market. Besides the ethernet attachment, Seagate also presented high-level visions at the Open Compute Project for Storage.[13]

A big shift in HDD technology came in October of 2017 when Western Digital (now merged with HGST) announced they were stopping work on HAMR and instead working on Microwave Assisted Magnetic Records (MAMR).[14] MAMR creates a spin torque oscillator in the disk head magnetic field, allowing it to be more finely controlled and allowing for smaller bits / higher bit density. Western Digital announced their intention of having MAMR HDDs available to the marketplace in 2019.[15]

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Western Digital believes that MAMR can carry the HDD industry, or at least their HDD technology, through 2025 with a 40 TB HDD. This represents a 15 percent CAGR in areal density, which is respectable but slower than the projected tape density.

The net summary as that no big change is expected in base technology. SSD will continue to take market share and high profit at the high end of the storage market, while disk is being squeezed between flash and tape. Google made a similar statement about high capacity / nearline HDDs continuing to be in the marketplace for the foreseeable future at the FAST 2016 conference.[16] As one of the largest cloud vendors in the world, Google drives a lot of demand and revenue for storage suppliers. As long as HDD costs remain below SSDs, the scale of cloud will require an HDD tier.

1.3.2 | Emerging Technologies

There are new technologies being developed in research labs around the world. IBM demonstrated atomic-level control of storage and spelled out “IBM” with atoms[17] as well as a nanomechanical data storage device.[18] There is investigation being made into DNA storage. There also is research being done on 5D crystal technology (so-called “Superman memory crystals,” inspired by the American movie released in the 1980s), resulting in a storage disc being put in the Tesla Roadstar car launched by Elon Musk into space.[19]

Although great work is being done, there does not appear to be any new technologies capable of getting to a production phase and generally available in the next decade. DNA storage is the encoding and manipulation of DNA to store information. Although incredible densities are possible, including getting to EBs per cubic inch, there are also many challenges before this technology is ready for a typical consumer.[20] In today’s marketplace a typical DNA reader runs at 1 Mbit/sec, has 95 percent accuracy, and is ~$2K/reader (using Nanopore readers). Writing information into DNA is done at a rate of ~1bit/sec and is possibly WORM only. Costs are failing because of the large demand in the medical industry to study DNA, but there is a long way to go before DNA storage reaches a competitive cost. The DNA Storage research is in its early infancy.

5D crystal technology is a multi-state optical WORM technology, potentially for permanent records. According to the University of Southampton, the format has a unique 5-dimensional method of storing data.[21] Ultra-fast laser pulses encode data in 3 spatial dimensions and have changes in polarity and intensity of light reflections in a quartz substrate, providing the 5 dimensions. According to the University of Southampton, “This may sound like science fiction, but it’s basically a really fancy optical illusion. In this case, the 5 dimensions inside of the discs are the size and orientation in relation to the 3-dimensional position of the nanostructures. The concept of being 5-dimensional means that one disc has several different images depending on the angle that one views it from, and the magnification of the microscope used to view it. Basically, each disc has multiple layers of micro and macro level images.”[22] Unfortunately, femtosecond lasers are big and expensive today, precluding the use in a high volume storage

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

device. Although some university and research organizations are validating the viability of this technology, no storage company has come forward with products based on this technology.

What is closer to the market is Storage Class Memories (SCMs). These are technologies between flash NAND and DRAM.[23] There are several players looking at MRAM and other SCMs.[24] These have the potential to squeeze the flash industry from the top of they become more viable. However, they are all far more expensive than tape, which is the focus of this roadmap. While emerging technologies need to be watched, there is not another low-cost storage technology coming out soon to challenge the tape industry.

1.3.3 | Tiered Storage with LTFS

LTFS is a self-describing format adopted by the Storage Networking Industry Association (SNIA). The LTFS Technical Work Group (TWG) includes members from the storage industry. IBM, Oracle, HPE and Quantum are working together toward the adoption of an open LTFS industry standard recognized by the American National Standards Institute (ANSI) and International Organization for Standardization (ISO).

LTFS represents a significant step forward in moving tape storage away from its reputation as complex and difficult to use. Unlike disk and flash storage, tape has previously required device-specific software to read, write and manage data stored on its cartridges in libraries. But LTFS sets a new standard for ease of use and portability for open systems tape storage. Users can run any application designed for disk files against tape data without concern for the fact that the data is physically stored on tape.

LTFS uses the file system's format and operating system (OS) resources to graphically display the contents of a tape cartridge in the OS user interface format, typically a folder/tree structure. The metadata of each cartridge, once mounted, is cached in server memory. Metadata operations, such as browse directory and filename search do not require tape movement. Similar to using a USB drive or memory stick, LTFS enables users to share data across platforms by simply dragging and dropping files. Access to data is faster and easier as users just load a tape into the drive and mount it into the file system, and files become visible as if they were on a disk. Because access is not through application software, such as backup software, cartridges formatted with LTFS can be exchanged more easily between users working in different operating systems, using different software, and in different locations. Users' operational agility and efficiency is increased because LTFS doesn't require knowledge of tape.

Additionally, a tape formatted with LTFS represents a key component for the Cyber Resiliency framework, providing the necessary software intelligence to enable portable and self-describing content which optimizes the appropriate retrieving of the information located offsite, whether in vault or in a different location. Thus, LTFS can help deliver "air gap" protection.

The value of LTFS includes:

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

- Self-describing tape format addresses tape archive requirements, eliminating the need for proprietary tape management software.
- Makes tape look and work similar to other removable media (e.g., thumb drives), enabling simpler data sharing among users.
- Open format whereby a user can access to the files directly without the application that wrote the data.

With LTFS:

- Tape can be used without making any changes to file-system-based applications
- Tiering from flash or disk to tape is enabled, e.g., FLAPE concept
- Long-term retention or permanent storage can be placed on tape cartridges in an open format for future read back
- TCO can be lowered by over 90 percent using tape for tier 2 and tier 3 storage
- Energy consumption can be 1,500 times less than a comparable storage disk array

LTFS has expanded its use cases. The large-scale enterprise data solutions from companies such as Fujifilm, IBM, Oracle, Quantum and Spectra Logic are offering customers advanced solutions that not only reduce storage costs, but also seamlessly integrate LTFS tape into unstructured data workflows. Other areas where LTFS continues to prove itself include archive-intensive and data streaming applications in industries such as banking, medical, IoT, hybrid cloud, cloud service providers (CSP), Oil and Gas and analytics, as well as media and entertainment. Also, as organizations look to add cloud to their data protection solutions, many cloud providers themselves use tape as a means to deliver low-cost, long-term data archive-as-a-service solutions.

LTFS combined with flash and tape storage (FLAPE) can effectively work together to provide seamless storage solutions with excellent performance, high capacity, and at very competitive costs. Emerging workloads and use cases that require fast transactional processing of data that will remain unaltered are especially good candidates for a FLAPE solution. Today's workload information is rapidly getting colder, with much of that in a "deep-freeze" state, never to be touched. Storage professionals must manage the retention of data; "defensible deletion" is often not a good enough option. LTFS provides a method of changing the face of data infrastructures. Automated tiering managed through policies is now fully in the vision of the IT storage community. Reducing data in the primary repository (i.e., tier 0 and tier 1) enables storage savings and reduces backup and recovery workloads for Big Data environments.

Managing the expectation of end users can be the storage professional's biggest challenge. LTFS can be part of an overall solution that sets expectations on both retrieval time and cost of storing data.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

The LTFS format is transportable and easily integrated with file system-based solutions. The metadata index on LTFS-formatted media can be read and parsed into large-scale file systems without ever touching the data on tape, with no external database or format changes -- just an XML parsing and “symbolic link” creation in the file system.

For further reference on how LTFS works, visit the LTO consortium website which explains this information in an easy way. https://www.lto.org/technology/ltfs/

Also, other resources describe the value of LTFS:

https://horison-groflyllc.netdna-ssl.com/wp-content/uploads/2018/09/Fujifilm-LTFS-Redefines-the-Future-for-Tape-2018.pdf

According to a Gartner report from 2018 titled, “Predicts 2018: Storage,” through 2022, data center tape will remain actively in use in more than 50% of large enterprises.[25]. We expect LTFS solutions to be a key player in reducing storage costs.

“Five years out, tape is still going to be relevant. There is going to more data than ever that needs to be stored, and they’ll be looking for cost-effective ways of doing that.”

– *Tom Coughlin, Coughlin Associates.*

1.3.4 | Object Interface in Open Source

There are several proprietary object solutions for tape provided by different companies. Some explicitly state that tape is the back-end storage device and are used in private cloud installations, while other public clouds include tape as a storage tier inside their data centers. For example, Google and Microsoft have talked about tape usage in their companies.

Progress has been made on creating open source projects for object interfaces to Tape technologies. OpenStack Swift is a good example of an open source project designed for object interfaces. IBM, along with BDT and others have worked to release a connector for high latency media storage devices (devices that can take a long time to satisfy a read or write command).[26] This component, called SwiftHLM, plugs into the OpenStack SWIFT stack and adds capabilities for tape, optical, or any high latency media.

It is difficult to size the adoption of open source in the industry. In some cases, the open source code is used as an example and a starting point and software engineers inside a company will create their own optimized software stack to take advantage of the tape system. In other cases, the open source may be adopted with the support structure being handled internal to the company with help from the open source community. This structure limits the types of companies and types of data that can be put into a system, vs commercial solutions.

Open source projects can be important kick starters for companies and industries to understand capabilities and strategic approaches, and allows free expertise to be shared.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.3.5 | Tape Encryption and Quantum Computing

Quantum Computers have 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.[27] Although there are small quantum computers available on the internet, it is clear that quantum computers are not yet as powerful as conventional computers.

However, there are teams of researchers and developers working on quantum computers. Many applications exist that will benefit from this new way of computing, and a lot of investment from companies and governments is pouring into this field.

One of the claims of quantum computing is the ease with which it can break today's encryption. Quantum computing allows for quicker factoring and discrete logarithms of the keys used at the heart of encryption codes. The projection is that current asymmetrical encryption schemes such as RSA and error correction code (ECC) will become insecure in the near future. The symmetrical encryption schemes such as AES will last longer.[28]

Recent reviews of AES in particular have shown it resistant to Shor's algorithm. Reviewers also used Grover's algorithm to analyze the strength of AES 256. AES-256 in a quantum computing world is theoretically equivalent in operations to AES-128 with today's classical computers.[29] In practice, Grover's algorithm is less powerful than generally assumed.[30] The number of bits needed to add to keep the same key strength is lower. However, the trend of reduced security for current encryption algorithms is clear.

Today's LTO tape drives use GCM-AES-256 encryption. The tape solution encryption techniques are estimated to be secure over the next decade. However, this cannot last forever. As quantum computers get more and more powerful, existing encryption algorithms will get easier to break.

Some alternatives must be considered by tape storage suppliers over the next decade.[31] One alternative is to push the encryption and security requirements higher up the storage stack into the host server that sends the data. However, once data is encrypted, it cannot subsequently be compressed or deduped. If a tape drive receives encrypted data it cannot use hardware compression which can compress cleartext data at line speed, without a performance penalty. All modern tape drives have compression enabled by default and do not require a server to expend any compute cycles or do anything special. Compression done at the host requires computational resource which can significantly affect performance – history has shown this to be a non-optimal solution. Therefore, INSIC expects that tape suppliers will incorporate new Quantum Safe Cryptography (QSC) encryption schemes in tape storage devices over the next 10 years.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.3.6 | Erasure Coding & Tail Latency

Increasingly, cloud suppliers are adopting tape storage devices into cloud infrastructure[32]. As data continues to grow, adoption of a low-cost storage technology is more and more attractive. However, a tape storage system has different behaviors than an HDD or SSD storage system. Because tape is a removable media it requires different infrastructure and different ways of integrating into the cloud software stack. Some examples include erasure coding in a cloud; dealing with tail latency; and handling of file sizes. These highlight the differences in technologies and what can be done.

Cloud suppliers have many software engineers who provide value-add to the infrastructure. Assuming that hardware is a commodity and can be bought at low prices, software can be modified to provide high availability and add reliability to hardware infrastructure. Examples of commodities in clouds include HDDs; processors in servers; and memory in servers. Modifying software to tolerate hardware failures is a requirement when clouds get to a large enough scale. If software is required to tolerate hardware failures, then perhaps hardware requirements can be eased, making them even lower cost. Finding the right balance between software tolerance and base hardware reliability is a task that the storage and cloud industry have been researching for several years.

Google presented their view on how data is stored on HDD and SSD[33]. Google states that both technologies will be used in cloud data centers for at least the next decade, so it's important to understand how these technologies are used before looking at a different technology like tape storage. Google describes their "collection view" of how data is distributed among geographies, data centers and storage devices in their paper on "Disks for Data Centers"[34]. This is shown graphically in a presentation at FAST 2016 USENIX conference on FILE and Storage Technologies[35].

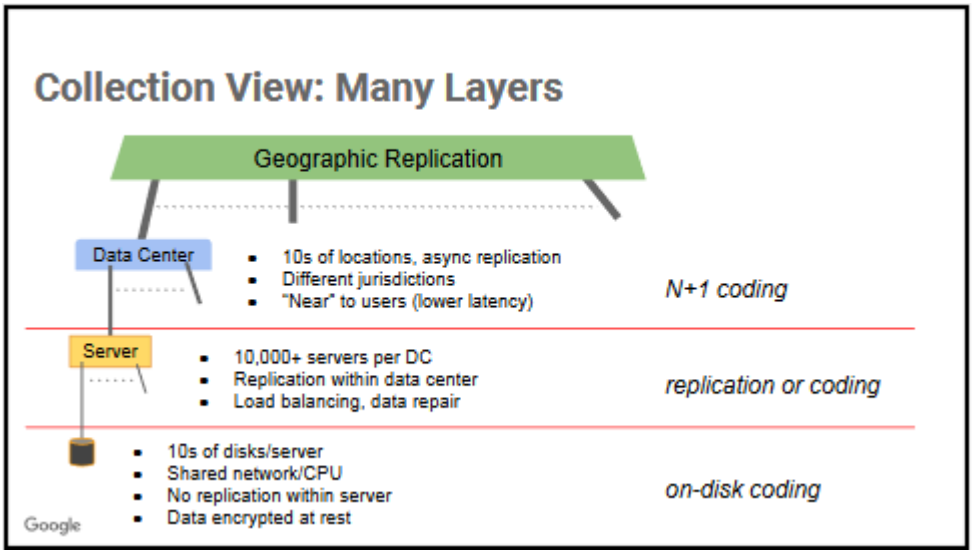

Figure 4 | Distribution of Google data by geography, data center and storage device
Source: Brewer, E. (2016, February 23). Disks and their Cloudy Future. Retrieved from https://www.usenix.org/sites/default/files/conference/protected-files/fast16_slides_brewer.pdf

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Google is not alone. Netflix uses Amazon for their cloud services. After a serious incident at Amazon, Netflix changed their strategy on distributing data and workload to use redundant data in multiple regions in the US.[36]

To take advantage of distributed data cross multiple data centers and regions, erasure coding techniques are used. Erasure coding spreads data across multiple storage devices in a way that increases availability and reliability. Extra data is added to allow checking for lost bits and reconstruction of the data when lost bits are detected. This is similar to RAID techniques used widely with disk but is more flexible in its applications. Examples for cloud applications include spreading data over geographies or data centers, while RAID is typically limited to disks contained in one enclosure.

Erasure coding techniques have been in cloud infrastructure for many years, as shown in Microsoft, Google and Backblaze papers [37] [38] [39] This well proven technique is successful with HDD and SSD devices. Applying this same technique to a removable media such as tape requires more logistics. In order to access data, a tape cartridge must be loaded into a tape drive. If data is spread across multiple tape cartridges, each must be loaded into a tape drive at the same time. If data is spread across multiple libraries, potentially in multiple geographies, then multiple robots must load the cartridges in the tape drives. At smaller scales where RAID on Tape (or RAIT) has been tried, data is spread within the same library and a small number of tape library robots must load the cartridges. This causes some performance problems. In larger scale, where multiple tape libraries and robots load cartridges, additional factors such as latency across geographies and mount queue handling must be considered.

Tape infrastructure such as library robots and tape drives can be considered part of the connectivity on a cloud. The data resides on the tape cartridge, while data passes through the tape drive and is carried around by the tape library robot. In this sense, tape infrastructure is similar to the network connectivity in a cloud.

There have been several studies on how networks impact clouds. In a paper from Microsoft Research and Purdue University called “When the Network crumbles…” [36] reliability of the network is reviewed and shown to be a major factor impacting service availability. This paper talks about networks inside of data centers and WANs connecting to the data centers. Figure 5 shows that connectivity losses are the greatest factor impacting service.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

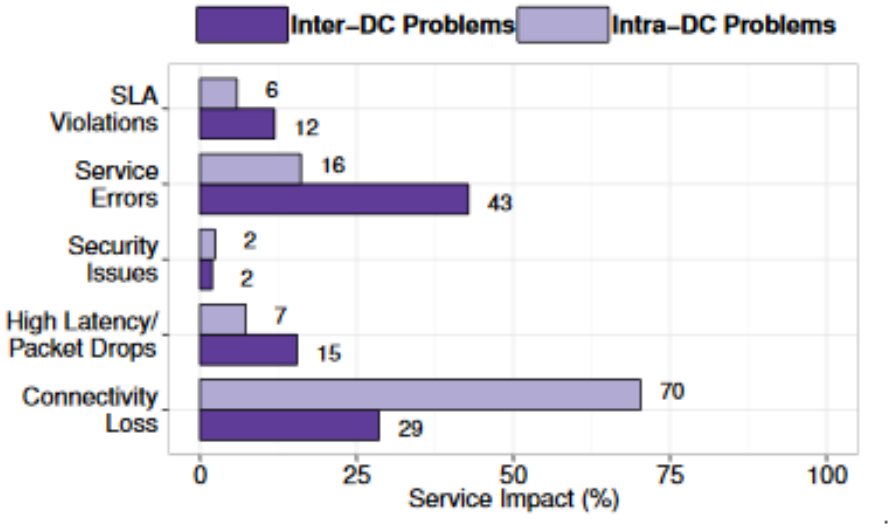

Figure 5 | Top 5 categories of service impact in a cloud provider caused by network failures

Source: Rahul Potharaju of Purdue University, Navendu Jain of Microsoft Research. "When the Network Crumbles: An Empirical Study of Cloud Network Failures and their Impact on Services"; Oct 2013.

This paper also states that the strategy of cloud providers is to build networks based on a scale-out design using inexpensive commodity hardware, similar to the storage strategy in clouds.

The paper describes several layers in a network for a cloud, including inter-data center links, access servers, aggregation switches (AGG) and top-of-rack switches (ToR). One key finding in the paper is that "redundancy is more effective at the AGG-ToR level (an improvement of 28.7 percent) …" As stated earlier, in a tape storage system the tape cartridge is the commodity device and the tape hardware is infrastructure and more analogous to a network in the cloud. In a tape storage system, the tape robotics can be analogous to a top-of-rack switch as it is connected to all the tape drives. If this analogy holds true to tape robotics, then a statement could be made that redundant robotics is a key requirement for service level agreements (SLA) in a cloud system that includes tape.

This makes sense. What would ethernet be without failover and redundant paths? Similarly, redundant robotics and connectivity through tape drives is a requirement for a reliable and highly available cloud that incorporates tape technology. The commodity hardware in a tape storage is the tape cartridge. A tape library contains many tape cartridges and much more data than a single HDD. A tape drive is a pipeline for reading and writing data but does not contain data itself. A tape library and tape drives in the library is more analogous to connectivity such as ethernet. Data is fully encapsulated on a tape cartridge; data can be spread through erasure coding or RAIT on to multiple tapes. The initial thinking might be to include tape drives or even entire tape libraries as an encapsulated device, but this is not the same as an HDD or SSD.

One aspect of redundancy for availability is cost. There is almost always a cost tradeoff for features like availability. Traditionally cost is captured in a TCO study. TCO in a "collection view" as described by Google is dominated by the cost of HDDs[33]. As tape is more widely adopted

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

into cloud infrastructure, we can say it is the same for tape: TCO is dominated by the cost of the tape cartridge.

One can make an analogy of HDD as removable media as well, at least from a service point of view. When a HDD fails and data is not available, it is replaced. Since data is spread across multiple HDDs, the data is rebuilt and put on the new HDD (or a spare HDD in the pool). Similarly, if a tape cartridge fails it can be replaced. Since data is spread across multiple tape cartridges, the data can be rebuilt on a replacement tape cartridge.

Another implication of utilizing erasure coding in a cloud deals with file sizes. An important attribute of storage systems is performance. Hard disk and flash storage have fast access times to the first byte of data in a file. When large files are used, tape systems have fast access times to the last byte of data in a file. Software using tape storage, including erasure coding schemes, can take advantage of this characteristic to improve overall performance in cases where access to a complete file or dataset is important. When spreading data across several regions or storage devices in a data center, or even across tape drives in a single tape library, a large amount of data increases efficiency. If smaller file sizes are used, then erasure coding can be done but is inefficient. An optimized system would aggregate smaller file sizes into a larger aggregation of data and use that as the base for storing into a storage device.

Besides erasure coding, another key issue in clouds is tail latency, as described by Google. Tail latency is defined as slow reads in a statistical distribution of reads. If 100 read cycles are measured, there will be a distribution of read times, with a tail representing longer access times. If data is distributed and 100 reads are required to access the data, then tail latency becomes significant. Erasure coding schemes allow for some relief from the tail latency problem. In an erasure code, not all data has to be accessible. If data is split into 7 chunks, and only 5 of 7 chunks are needed to recreate the entire data set, then some tolerance of slower storage devices or missing data is allowed. If there is a storage device or infrastructure causing a slow read, but all other chunks are available, then data can be reconstructed, and the slow device can be sent a command to cancel or abort the read command.

The strategy to handle tail latency with HDD or SSD by using erasure coding can also be used in tape storage. If a particular tape system is slow to respond, but other systems are not, then the tail latency of a particular read command can be masked. However, the implementation must be different. Cancelling a read command to a removable media storage system such as tape has other implications than cancelling a read command to an HDD. This leads into vendor-specific details on handling of tapes in tape structures, as well as dealing with mount queues in software, so is beyond the scope of the INSIC roadmap.

Several aspects of using a removable storage technology in cloud infrastructure has been explored. There is much more that can be explored. This is not too dissimilar to the adoption of SSD into a cloud infrastructure, where handling of erasure cycles must be accommodated. With any different technology, there are different characteristics and behaviors. To gain the benefit of the technology, the software stack must be modified to take advantage of the benefits of the new storage technology while also mitigating any new costs. With tape being the lowest cost

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

storage technology for the foreseeable future, investment in optimized cloud software stacks will be a focus for the next several years.

1.3.7 | Software Defined Storage

Over the last few years, Software-Defined Storage (SDS) has transitioned from marketing buzz to real products available in the market. Several companies have now released SDS versions of storage devices, allowing storage administrators to run software on their choice of hardware. INSIC views this as a good trend for the tape industry.

As software is virtualized from the hardware, additional changes in behavior and timings must be tolerated for a robust solution. The changes are required to tolerate different hardware, as well as take advantage of the different storage technologies. This gives greater capability to handle different storage behavior. In addition, SDS solutions generally handle tiered layers of storage technologies. If there is a set of storage allocated for high performance, medium performance, and lower performance, then the ability to handle latency capabilities are designed into the software, more naturally supporting tape solutions. LTFS is one example of an SDS solution that is being incorporated into other software stacks to enable tape to be used as part of a larger solution.

Software Defined Storage also allows more flexibility in host interfaces. One example is the work being done by multiple companies with OpenStack Swift to enable a cloud/object interface based on LTFS. Another example is Oracle's announcement of the incorporation of OpenStack Swift into their HSM software, or Quantum's StorNext object interfaces, or IBM's Spectrum Scale solution with an object interface.[40]

According to the “Software Defined Storage for Dummies” e-book publication from IBM[41], another way to think about SDS is to imagine it as a layer of intelligence inserted at strategic points in the data path between applications and their data. With smart software setting up residence between the relatively passive physical storage media such as flash, disk or tape and the rest of the IT infrastructure, a wide world of possibilities suddenly opens.

SDS today represents the layer(s) of intelligence and functionality deployed between application hosts and the physical storage device.

SDS is just code. It can be downloaded and deployed when you want, as you like, as long as you can accommodate its modest processing and storage requirements. This flexibility is a key reason it is becoming more widely adopted in the storage industry.

1.3.8 | Host interface

Tape has traditionally followed disk technologies for its interfaces, which is why tape is predominately attached through FC or SAS interfaces. Both FC and SAS interfaces have robust roadmaps being followed by the storage industry. Other interfaces such as Infiniband have

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

found their place in niche markets, but have been more successful with server-to-server connections than server-to-storage connections.

As cloud and web companies continue to grow, there has been more and more interest in ethernet attachment to tape. iSCSI and FCoE have been around for a while in the disk controller market, and SANs have allowed switching from FC to FCoE to enable different solutions. However, SAN responsibility and LAN/WAN responsibilities have been split in many companies, and having an independent SAN for storage is becoming a larger proportional burden. As a result, a growing trend is ethernet attachment to storage. Whether iSCSI and FCoE continue to dominate, or other interfaces such as iWarp or RoCE take over, is yet to be seen.

1.4 | COST & ENERGY CONSUMPTION ADVANTAGES (TCO)

1.4.1 | Introduction

With cost as a driving force behind IT investments, tape storage can offer significant cost advantages compared to long-term storage on disk or in the cloud for many applications. This, coupled with the technical attributes of long shelf-life and extremely high data reliability, make tape an excellent fit for long-term data storage.

In this section we review the Total Cost of Ownership (TCO) and energy consumption of tape storage compared to disk-based and cloud-based storage. We begin with the results of recently published TCO studies and a tool to calculate TCO for various scenarios. Next we discuss the factors affecting TCO modeling. We then consider the changes since 2016 and the impact of the changing underlying technologies to provide some insight into likely future trends.

1.4.2 | ESG TCO Study

The Enterprise Strategy Group (ESG) performed a detailed TCO comparison between tape, disk and cloud storage in 2016 with an update in 2018. The 2016 study compared two scenarios for long-term data storage:

- **LTO tape scenario:** In this scenario, the customer is using a properly sized, high-end tape library, housing tape drives and LTO-6 cartridges at the beginning of the time horizon, and transitions to LTO-7 after five years. The configuration also includes servers to run LTFS directory services and software costs for LTFS itself.

- **PMO scenario**: In this scenario, the customer is using a comparable set of hardware and software components based on tier-2 disk solutions, including properly sized disk arrays, nearline disks (accounting for RAID overhead), and archive software.

In addition, a hybrid disk-cloud scenario was considered in which 20 percent of the data was stored in the cloud.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

The paper details the methodology which incorporated eight factors affecting cost: *hardware, software, media, infrastructure (i.e., power costs and rack space opportunity costs), maintenance charges, staff costs, data migration costs, and data access costs. The sum of these categories equals the total cost of ownership (TCO) of each solution.*

The 2016 ESG study was updated in 2018 to include advances in tape (LTO-8) as well as disk technology. It also included an all-cloud storage scenario for comparison. Results of the 2016 study are summarized in Figure 6. The study found an all-tape system to cost only 14 percent relative to an all-disk solution.[42]

	All Disk	All Tape	80% disk 20% cloud
Total 10-year cost	**$16.8M**	**$2.4M**	**$13.9M**
Cost relative to disk	**100%**	**14%**	**83%**

Figure 6 | 2016 ESG Study
Source: INSIC

The 2018 update found similar relative savings of a tape-based system compared to disk. In addition, compared to an all-cloud system, tape was only 34 percent of the total cost of ownership, as shown in Figure 7.[43]

	All Disk	All Tape	All Cloud
Total 10-year cost	**$15.2M**	**$2.1M**	**$6.3M**
Cost relative to disk	**100%**	**14%**	**41%**
Cost relative to cloud	**241%**	**34%**	**100%**

Figure 7 | 2018 ESG Study
Source: INSIC

Figure 8 shows the same comparison graphically.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

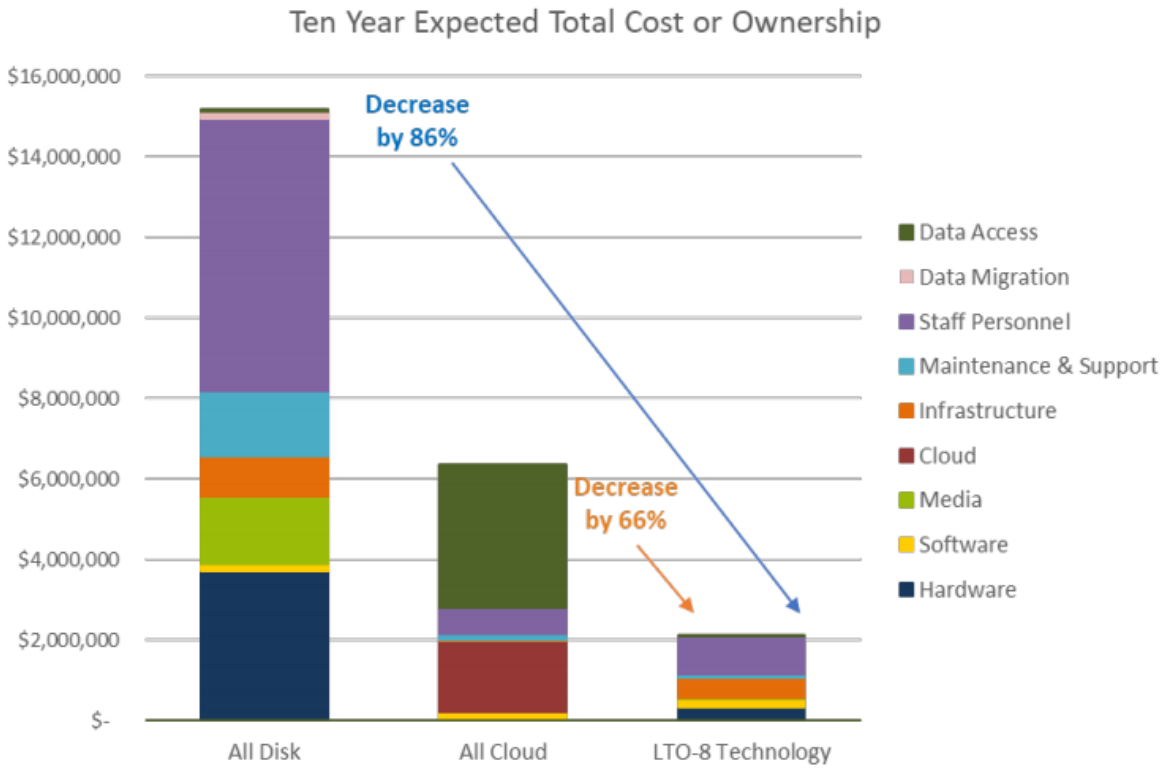

Figure 8 | TCO comparison of disk vs cloud vs tape
Source: Arcilla, A.(2018, August) Quantifying the Economic Benefits of LTO-8 Technology. Retrieved from https://www.lto.org/wp-content/uploads/2018/08/ESG-Economic-Validation-Summary.pdf

1.4.3 | TCO Calculators

The LTO Program has published an online TCO calculator. This calculator allows the user to enter relative cost and performance parameters to compare TCO estimates between tape, disk and cloud-based systems. It includes entries for annual data growth as well as how much data is retrieved. This allows a user to compare systems that most closely represent their application. Here's the link to the tool:

https://www.lto.org/resources/tcotool/

Fujifilm has also developed a TCO calculator for the same purpose.

https://page.dternity.net/TCO.html

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.4.4 | Modeling Methodology

As with any modeling methodology, the results can vary widely depending on the factors considered in the model and the underlying assumptions. Primary costs considered in TCO modeling include:

- Initial cost of storage system and media
- Annual floor space expense
- Energy expense (primary and secondary, e.g., cooling)
- Service and support costs

Important assumptions in cost modeling include:

- Equipment and media life
- Data storage life
- Data compression
- Data redundancy
- Equipment discounting over list price
- Data growth rates
- Data migration strategies
- Monthly amount of data retrieved
- System bandwidth

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, disks have a much higher failure rate than tapes so this must be accounted for in differing redundancy levels. Also, disk systems tend to last significantly shorter than tape systems, leading to higher replacement costs. These costs are included in the service charges for cloud systems; however, cloud also charges for data retrieval. And for all systems, a data migration strategy must be considered. Several of these factors are discussed in more detail in the ESG Studies.

In addition, financial decisions are usually constrained by capital and expense budgets, which often dictate whether a more heavily capitalized or expensed system is preferred.

1.4.5 | Energy Consumption of Storage

The 2015 Clipper Group study[44] included comparisons between tape and disk for energy consumption. In this study, disk systems were assumed to power-up the drives continuously and tape systems also powered drives continuously. Energy consumption for each was then doubled to account for cooling. The study found the disk system to consume 140 times the energy of the tape system. In fact, the energy cost for the disk system was more than 60 percent of the total cost of ownership of the tape system. Although recent advances in disk spin-down technology have reduced this difference, a detailed comparison to tape is not known to the authors.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

In 2018, ESG reviewed and updated the study for LTO-8 drives confirming that the differences in energy consumption between disk and tape continue to be significant with current technologies.[45]

Also, a study by EchoLeaf of the global energy needs for storing the world's data found just how much power is needed. It finds that HDD storage consumes roughly 2.5 percent of the electricity produced worldwide even though 'a great amount of stored data will only be accessed infrequently'.

With the low power consumption of tape, it concludes that 'simply transferring just 33 percent of stored data to tape could reduce global electricity consumption by 1 percent per year'![46]

1.4.6 | Trends in Storage Costs

Historical studies can show interesting trends in storage costs over time. In addition to the ESG study, The Clipper Group performed a similar study in 2015. Although the studies use different data grow rates and a variety of other factors, they may be compared by looking at the ratio of the TCO between disk and tape solutions; this normalizes many of the differences between the studies.

	2015 Study	2016 Study	2018 study
TCO Ratio Disk:Tape	**6.2**	**7.0**	**7.2**

Figure 9 | Studies comparing TCO between Disk and Tape
Source: INSIC

While the multiplier between disk and tape systems has dropped significantly between 2013 and 2015, it is important to note that several factors are at play, with the following being the most significant.

- The update cycles of tape and disk systems are independent. Depending on when the study data was gathered, one may have leapt ahead more recently than the other. In the July 2015 study, LTO-6 had been available for several years and LTO-7 was just on the horizon, while many of the disk solutions had leapt to 8 TB drives.

- The 2013 study focused on block storage costs, whereas the later studies included servers and software capable of delivering access to files on tape in a library, which were not included in the 2013 study.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Nonetheless, it is clear that of tape solutions cost significantly less than either disk or cloud solutions. This cost difference represents many millions of dollars in all of 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.

"Tape is the most affordable option for multi-petabyte storage capacity."

- Jason Hick, Los Alamos National Labs

1.4.7 | TCO Conclusions

Tape systems continue to cost significantly less than disk systems. Although the gap may have narrowed over the past five years, the recent TCO study estimated costs over the next decade showing disk systems still costing seven times that of a tape system. The most recent study comparing tape to a cloud system show the cloud system to cost more than 5 times the equivalent to tape system. In addition, technology challenges for disk systems will likely slow their rate of cost declines.

1.5 | THROUGHPUT AND RELIABILITY

1.5.1 | Data Rate

From an architectural point of view there are two major differences between tape and disk systems that result in substantial differences in transfer rates: multi-channel recoding and multi-bump head technologies. Both of these features play a key role in achieving very high transfer rates with tape systems.

In archiving applications, the verification process is one of the key requirements to ensure data integrity. Multi-bump tape heads enable tape systems to read the just-written data, performing on-the-fly verification without compromising the transfer rate. However, with disk systems, the host must request a full read process, resulting in an overall reduction of the transfer rate. Checksums used in the disk system can only inform the user that the data read is not valid; they don't prevent invalid data from being written to the disk. Read-after-write verification is the only way to prevent bad data from being stored on the disk. With tape systems, this is done in real time without any impact to the transfer rates; disk systems require a second pass at the time data is written to verify that data, which results in a significant performance degradation.

Multi-channel head technology is another enabler of faster transfer rates in tape drives. Disk drives use a single active channel. Figure 10 shows the roadmap for transfer rate growth for both tape and disk systems. CAGR for tape is 22 percent and 10 percent for HDD.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

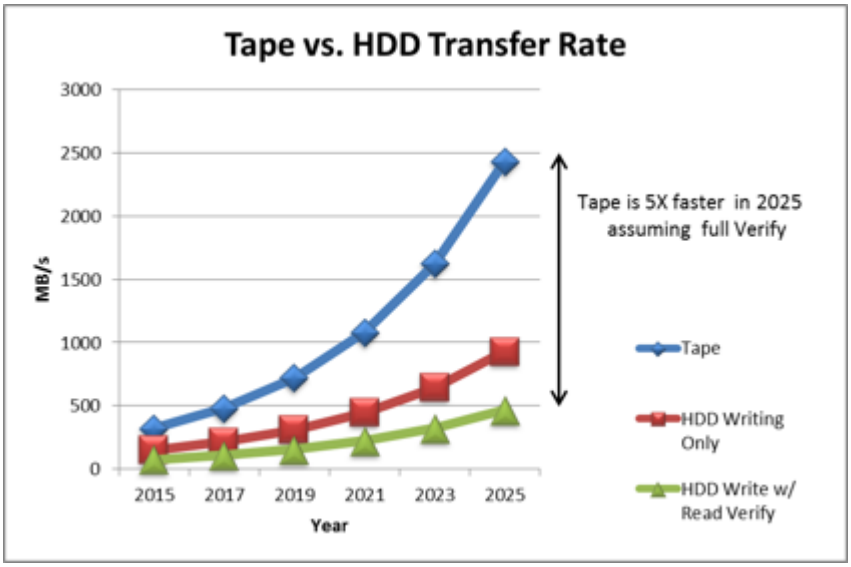

Figure 10 | Tape vs. HDD Transfer Rate
Source: INSIC

Unlike disk drives, most current tape systems use single cartridges rather than spreading data over multiple cartridges, such as in a RAIT system. Future tape applications can also benefit from erasure code-based RAIT type data spreading technologies to further shorten the archiving time. This concept potentially can lessen the need for replication while reducing the overall archiving times, improving data reliability, robustness to lost cartridges and more efficient overheads. There are already some usage examples of parity-based RAID system using tape (RAIT), such as Google backing up Gmail along with exabytes of data. In this application, Google used five LTO-6 tapes with a 4 data plus 1 parity policy.[47]

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Figure 11 shows projected data rates of Google's RAIT technology compared to current copy-based method and a potential Erasure Code concept applied to the capacities in 2025

Projected Archiving Times Study in 2025 per Roadmap			
	Single Tape	Google RAIT	EC Tape
Tape Capacity	248 TB		
Tape Transfer Rate (per tape drive)	2435 MB/s		
Number of Copies	2	None	None
Parity Policy	None	4 data + Parity	8 data + 2 parity
Overhead	100%	25%	20%
Number of cartridges can be lost or damaged with data still be recovered	1	1	2
Time archive 248 TB data	28 hrs.	7 hrs.	3.5 hrs.

Figure 11 | Projected Archiving Times
Source: INSIC

The results in this simple "what if?" study show 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

Data reliability with tapes can be described by four main parts:

- **Media life**
- **Wear**, which is number of passes across the magnetic head

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

- The **data durability** which is a function of random errors defined by BER specification
- **Correlated** or **nonrandom errors** including events such as drop pins, lost tapes, etc., based on AFR numbers

Media life is specified to be a minimum of 30 years if stored in approved archival environmental conditions:

- Temperature range 16 °C to 25 °C
- Relative humidity range: 20% to 50%

The number of passes for the media to reach 'end of life' is shown below for previous INSIC and current INSIC tape roadmaps. Based on the number of tracks and the number of head channels, the maximum number of full volume usage (FVU) passes will reduce as shown in Figure 12. The data in this table can be found in the INSIC-2015 technology roadmap. FVU stands for full volume usage and it's a function of # of passes to EOL and # of passes per full volume write or read.

Tape	TB	# of passes to write tape	# pass to EOL	# FVU
INSIC-2015	8	140	27200	194
INSIC- 2019	25	313	32200	102
	49	483	34560	71
	96	497	37093	74
	188	768	39812	51
	369	1187	42730	35
	723	1375	45862	33

Figure 12 | Capacities vs full volume usage per cartridge
Source: INSIC

As capacities increase, the maximum full volume usage per cartridge will decrease. Data reliability must take this into account by adapting the read, write and tape access profiles.

Given that the media is stored in specified environmental conditions to ensure its 30-year life and that the maximum number of passes must be below the limits for each generation,

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

the data durability can be determined based on the protection methods used such as replication and erasure coding.

One of the key requirements for archiving is the preservation of data integrity and the durability of the written data while at rest. For example, the majority of its life cycle while it is being stored.

- Data integrity is the assurance that the host records are correctly written to the media, despite potential errors occurring anywhere along the data path -- including the media itself (e.g., disk or tape).

- Durability is defined as the probability that a data object will be intact and available for access after one year; for instance, 99.999 percent probability, or five 'nines' of durability, meaning that for an archive of 10,000 objects one might lose one object every 10 years or so.

Preservation of data integrity and durability are based on a complex series of CRC checks and balances along with replication, ECC (including RAID) to ensure that the data is correctly written to the media and available with the advertised "durability nines" for the life of the archive.

Durability is a function of file size, bit error rate (BER); annual failure rate (AFR); protection methods, such as replication or erasure coding; and required support functions, such as mean time to repair a failed device or data. BER plays a key role especially as file size increases. Due to its dual layer orthogonal Reed Solomon ECC with multi-dimensional interleave format, tape has superior BER specifications compared to HDD. This is shown in Figure 13.

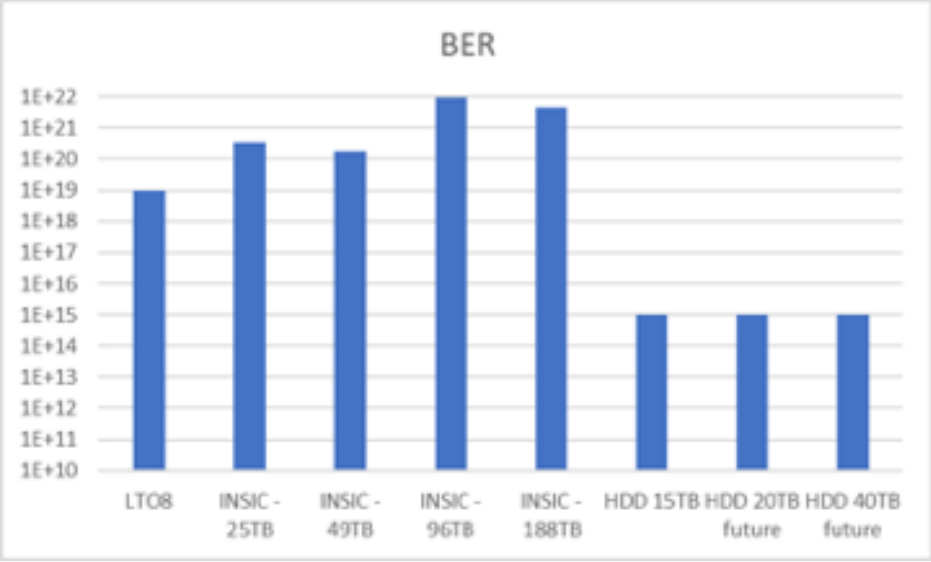

Figure 13 | Tape and HDD BER Roadmap
Source: INSIC

Figure 13 shows the durability estimation for HDD and tape-based on current products and future roadmap operating points. The durability is estimated based on the BER, file size and

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

AFR numbers where AFR is annual failure rate. Number of NINES can be calculated using BER and size of a data for both Tape and HDD based on their specifications. The assumption is that a BER event which is a hard-read case will result in the loss of data set or file on the storage device. See equations shown in Figure 14 to calculate NINES. Tape BER and capacities can be found in the INSIC-2019 Technology Roadmap white paper. For HDD, we used BER value of 1E-15 which is enterprise quality performance. Figure 14 shows three different durability nines data:

- The blue legend data shows the 'durability nines' computed assuming files are less than a data set for tape (typically a data set is 5 MB) or less than a sector for HDD (4K sector size) where the system can recover the rest of the data. These are typically small burst or random media errors. However, most common errors are nonrandom and correlated, therefore, this durability estimation may be too optimistic of a number. So, a file size-based BER is more relevant to typical real-world error cases.

- The orange legend data shows the file size 'durability nines' based on a BER event that may result in a loss of the file. The analysis shows the nines based on an average file size of 100 MB. Note that as file size increases the durability nines will decrease.

- The gray color legend data shows the combination of file-based BER and typical AFR numbers associated with tape- and disk-based systems. For tape, the analysis uses the AFR estimate of 0.1 percent tape and 5 percent for the disk.[48]

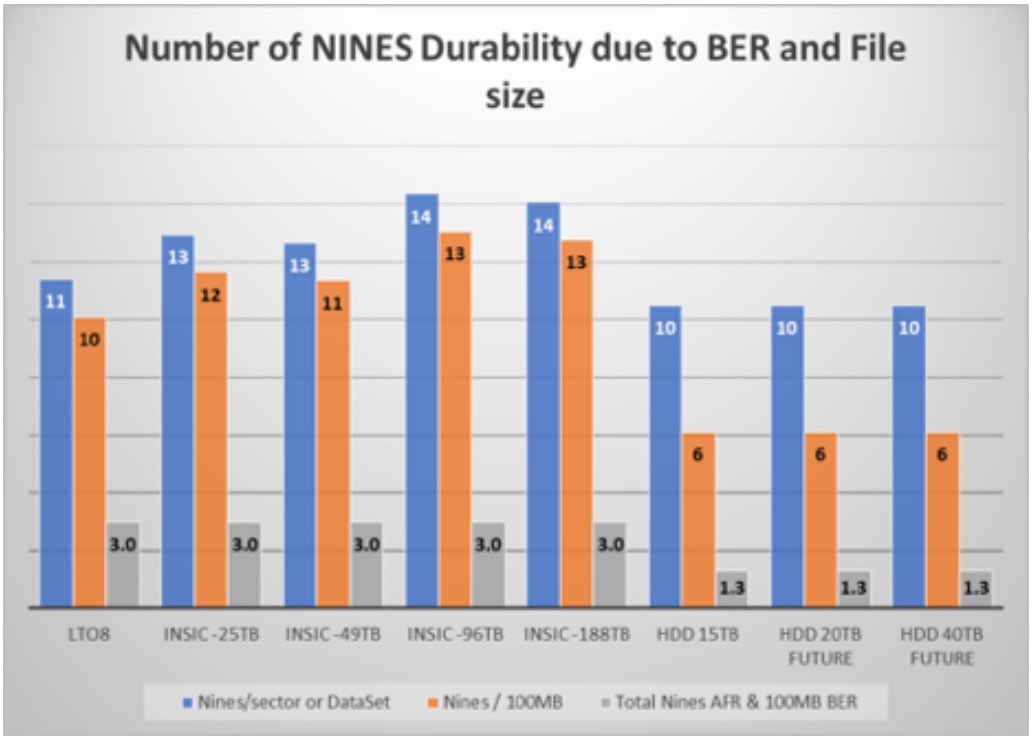

Figure 14 | Tape and HDD roadmap data durability
Source: INSIC

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Figure 14 shows the expected loss of data as a function of durability nines for both tape- and disk-based systems. Here, the analysis assumes a 1 PB data set where the average file size is 100 MB. Calculation of # of objects lost per year based on a selected # of NINES is shown below:

$$\#\ of\ ObjectsLost\ per\ year = (1 - \%\ Reliability) * (\#\ of\ Ave\ Objects\ stored\ in\ a\ Year)$$

$$\text{Where}\ \ \%Reliability = 1 - \frac{1}{10^{NINES}}$$

Using 100 MB file size, 1 PB dataset (10 million the average number of files) and the gray legend data from Figure 14 (the nines for 100 MB file and AFR), the tape might lose less than 10,000 files in a year, whereas the HDD loss might be up to 500,000 files. For archival applications, based on this relative comparison, it can be seen that tape is 50 times more durable than disk systems as long as it is maintained per its recommended storage and operating conditions.

BER is calculated assuming all the errors are random and uncorrelated. For correlated and non-random error events, AFR can be used which is based on the MTBF number from field data. For HDD, there are several publications that describe AFR in a data center application. Manufacturing specs for HDD typically define AFR as 0.88 percent (or one million hours of MTBF). Studies by Google and others show that field-based AFR for HDD can be as high as 4-6 percent which is substantially more than the manufacturing predicted data of 0.88 percent. In this paper, we use 5 percent for HDD AFR.

Tape AFR is more complex due to its removability. A tape that fails in one drive may work in a different drive or under different environmental or operating conditions including using a cleaning process to recover the data. The analysis focuses on the actual repeatable tape errors rather than the tape and drive combination errors. These are typically loss of a cartridge, dropped pins, cut tape, accidentally erased media or spatially worn media where the internal ECC format cannot help. Tape and drive non-repeatable errors are like the HDD, 4-5 percent but with removable media the actual repeatable tape errors as described above are significantly less.

In NERSC's report on a massive data migration process to upgrade old tape cartridges to a newer system, 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. 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. In our analysis we used the full failed cartridges instead of partial reads, which resulted in an AFR of less than 0.1 percent.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

The ability to recover user data based on removable media technology by trying different drives or conditions may increase latency but it will lower AFR numbers substantially compared to HDD, resulting in a higher durability archival solution. AFR for device failures dominate the HDD- based archiving system where even three copies fail to provide five nines of durability. With tape systems, the removability of a media cartridge from a drive which has failed still enables the user to recover the data by simply using another drive or even a newer model that has backward read capability. This feature allows us to use the AFR for the media failures rather than the AFR of the drive -- a significant benefit of removable media.

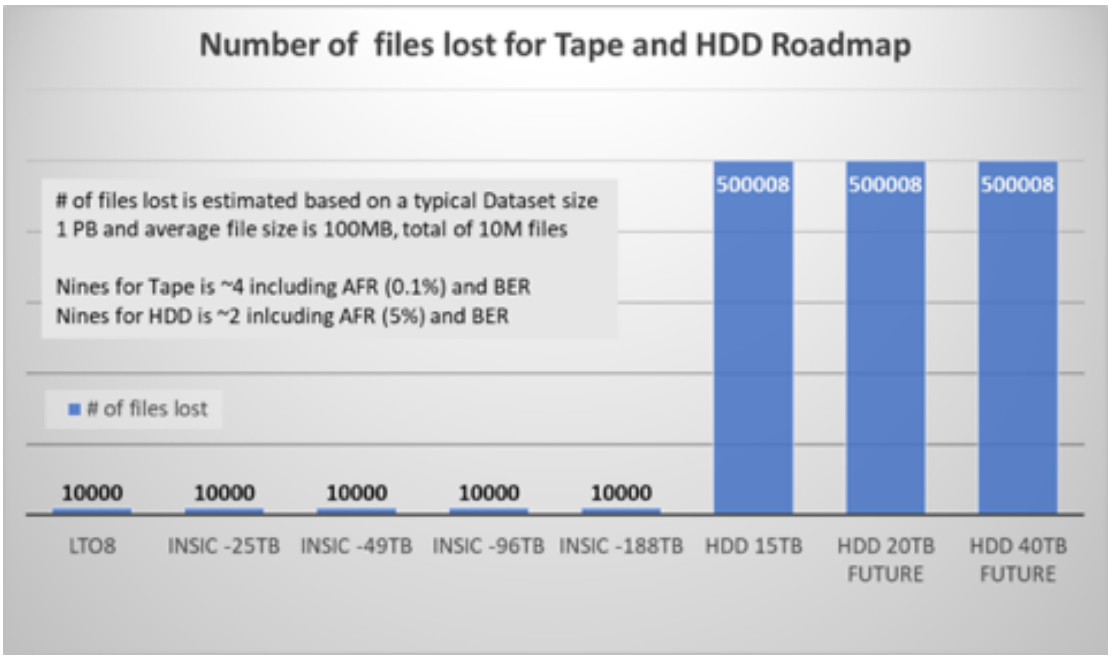

Figure 15 | Durability and Data Loss Relationship for Tape and HDD
Source: INSIC

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

	Units	SATA Disk	LTO-8	Notes
# of Copies for Archiving	#	3	3	Assume replication as durability control
Device Capacity (TB)	TB	12	12	Similar capacities
AFR (%) Media based on NERSC publication[49]	%	NA	0.1%	
AFR (%) Drive[50]	%	4.8%	4%	
BER specified		1.E-15	1.E-19	BER based on storage device technology
Probability of BER Failure	%	9.6%	0.00096%	=BER*8*Capacity(Bytes)
Probability of Check Sum Errors[50]	%	0.66%	~0%	Multi-Dimensional Interleaved orthogonal C1-C2 Tape format
Probability of Latency sector Errors[50]	%	9.5%	~0%	Tape has Read after Write with Error detection; all user data is written and checked
Reliability of Device is when No Error (BER, AFR, Latent Error and Check Sum error) Reliability = 1-(1-AFR)*(1-BER)*(1-CheckSumError)*(1-SilentError) $$NINES = \log_{10} \frac{1}{1 - Reliability}$$ This assumes an error even such as BER will result in loss of entire capacity of the storage device				
Archived Data Durability # of Nines	#	2	9	See description above

Figure 16 | Durability and Data Integrity Analysis Results
Source: INSIC

Another feature of tape is the superior end-to-end data integrity, compared to HDD-based systems in archival applications. Disks are prone to silent write errors such as bit rot, torn writes or misdirected writes, and the only way to detect these is to do a full verification which is a very costly process that basically results in a 50 percent reduction of the overall transfer rate. A paper presented at USENIX by Andrew Krioukov[51] describes this problem and addresses a set of complex CRC checks and mirroring to be able to identify issues later when the data is

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

requested to be read back. However, unless a full verify is done, the silent errors will be stored forever on the disk for the duration of the archiving until it's detected, at which time it's too late and the data has already been lost.

Tape has built-in verification in its hardware architecture that implements a read-while-writing process that checks written data as it is put on the tape. The end-to-end data integrity is implemented by CRC-protected Host records, and by the tape drive doing a read-while-write verification of the actual written data, using a combination of CRC checks and two-layer ECC protection based on spatial interleaving. Together, these result in less than one Silent Write Error in 10^{30} bits transferred.

The data durability as shown in Figure 16 are both based on single device (tape or HDD) reliability including BER and AFR without any additional protection such as erasure coding or replication. Figure 16 shows the impact of replication for both Tape and HDD but again based on single device performance including BER and AFR estimates. Overall durability based on AFR, file size and BER (see Figure 16) can be substantially improved by using replication or erasure coding as shown in Figure 17. Tape's key durability benefits compared to HDD can be summarized as:

- Tape has 30-year life compared to HDD's maximum five years
- Tape has more than five orders of magnitude lower BER compared to HDD
- Tape has 50 times lower AFR compared to HDD.

Given that tapes are stored and operated in the recommended environmental conditions and their usage doesn't exceed the End of Life media wear numbers, tape offers extremely high durability with replication or erasure coded archival storage protection policies compared to HDD as shown in Figure 17. The durability calculations for erasure-coded tapes are based on Markov Model. The following references describe the modeling data durability for storage devices where erasure codes are used.

Notes on reliability calculations by IBM Research:
http://domino.watson.ibm.com/library/cyberdig.nsf/papers/BD559022A190D41C85257212006CEC11/$File/rj10391.pdf

In this example we used a typical 12/4 policy (8 data tapes and 4 parity). We used 0.1% AFR rate and associated BER numbers as the INSIC-2019 Tape Technology Roadmap describes. Repairs were set yearly for tapes. For HDD we used typical repair times which are based on capacity and transfer rates which result in the Markov Model repair periods to be substantially different between the devices. HDD BER values were kept at 1E-15 and AFR numbers were 5%.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

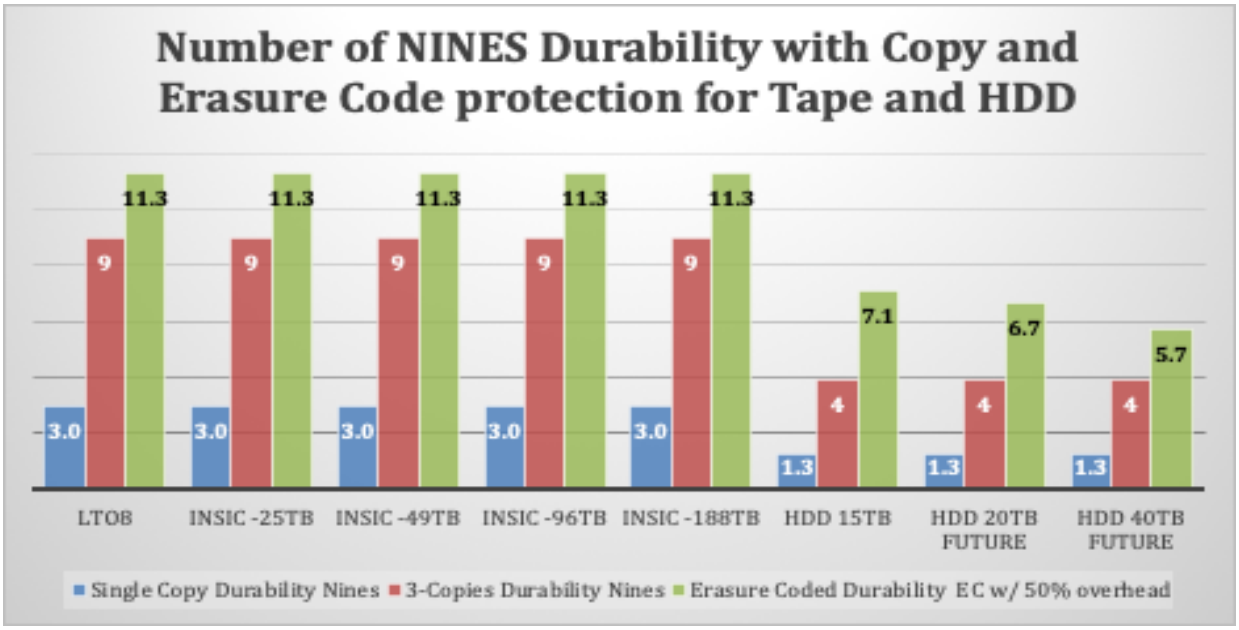

Figure 17 | Durability with replication and erasure coding
Source: INSIC

If we use the 'durability nines' as shown in Figure 17 with a 1 PB dataset, by simple translation of the 'nines' to AFR, we can estimate the potential data loss in megabytes for each case as it is shown in the figure below. This simple estimate demonstrates the power of tape due to its substantially strong BER and its removability where the system can use different drives and conditions to re-read tapes whereas HDD systems are not able to do this.

As a result, tape can offer much lower AFR numbers than HDD-based systems especially when combined with replication or erasure coding methods. Tape-based archival system offers extremely durable solutions.

Amount of Data loss (MB) with each Policy based on 1PB dataset				
Tape	TB	1-Copy	3-Copy	EC
LTO-8	12	1000000	1	0.005
INSIC	25	1000000	1	0.005
	49	1000000	1	0.005
	96	1000000	1	0.005
	188	1000000	1	0.005
HDD	15	50000760	125006	79
	20	50000760	125006	200
	40	50000760	125006	1995

Figure 18 | Amount of data loss (MB) with each policy-based 1 PB dataset
Source: INSIC

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.5.3 | Media Wear Out

Tape wear is a complex function of media magnetic and tribological properties, environmental conditions, and the drive usage model by the host 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 BaFe media under environmental conditions. Based on studies, media companies estimate the archive life of the Barium Ferrite tape to be beyond 30 years. However, the number of passes across the head and the guiding system, along with the usage frequency will impact this extremely long archival life. The INSIC 2015 technology roadmap paper shows maximum number of passes to EOL where EOL is end of life for each generation of tape. Based on the operating points for each generation, a full volume read or write will require different number of passes. The estimates were 27200 passes with 140 passes per full volume operation as shown in Figure 12. Also, the usage frequency of the media (as long as one does not exceed the maximum number of passes count) was shown to play a key role in determining the archival life. It is estimated that performing one full volume data transfer per month reduces the life of the cartridge down to 25 years.[52]

Another variable in the archival life of the BaFe media is the storage environmental conditions. This is normally based on estimation using the Arrhenius equation, which is a formula for the temperature dependence of chemical reaction rates. Based on this we can also estimate the archival life of a media if it was stored in different environmental conditions, including monthly usage method. Figure 19 shows the archival life for two usage patterns as a function of storage temperature. At 35 degrees C storage temperature, the media archival life will be in excess of 10 years even when it's used every month.

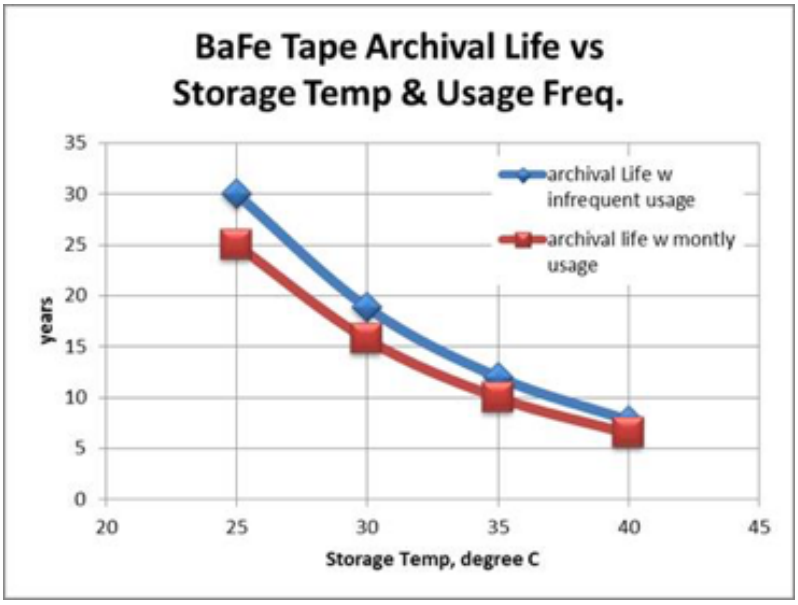

Figure 19 | BaFe Archival Life
Source: INSIC

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

However, in order to obtain this remarkable length of life, it is critical that the tape utilization by the user application software also is considered. This 30-year life period is an archival life time in which the media is stored and the data is to be retrieved within the 30-year period. However, when a tape cartridge is kept in a library and used frequently to transfer data to and from the drive, the life will be de-rated based on the usage profile (e.g., how often the tape is used, the number of passes across the magnetic head and tape guiding, how well the tape is kept streaming by the host such that repositions are kept at minimum), and the environmental conditions in which the tape is used.

Tape capacity, number of tracks and number of data channels all play a role in translating the maximum pass count to a full volume usage. Figure 20 shows what this means for the user in terms of data size and life. We use total data transferred per cartridge to EOL as a metric to describe this complex media wear function.

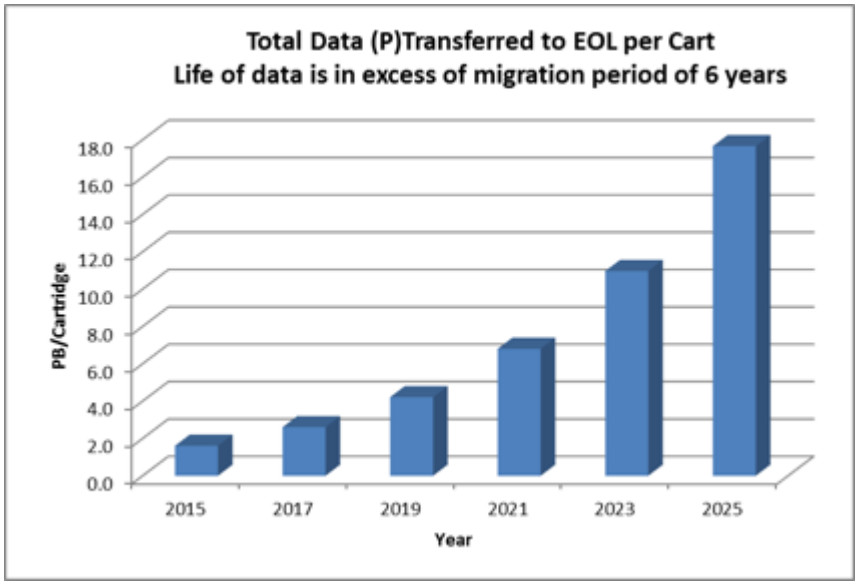

Figure 20 | Total data transferred per cartridge to End of Life
Source: INSIC

The chart shows how much data the user can transfer per cartridge over the generations where the data is guaranteed to last beyond the data migration period (which is currently every six years if the cartridge is used not more than once per month). The improvements in media durability along with increases in capacity and number of data channels result in an exponentially increasing amount of data transferred to EOL, where the life of the data is guaranteed beyond the migration period of six years.

1.5.4 | Roadmap Benefits from the Changing Usage Model

The usage model for tape is changing. For many years it was the primary backup media and tapes would be used and re-used many times over as they were cycled round the full and

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

incremental backups lifecycle. In this usage model it is important that the tape is able to support many full volume passes over its expected life.

The new usage model for tape is archive -- storing vast amounts of data for very long periods of time. In this case, the data is written once and only occasionally accessed. This allows the requirements on full volume passes to be reduced and used for the benefit of the roadmap.

With fewer full volume passes, the requirements on the tape coating are less constrained though durability is always a factor. It can be thinner and smoother enabling significant increases to the bit density and therefore overall capacity. In this way, the changing usage model of tape is benefiting the roadmap.

Figure 21 shows how the reduced full volume passes enable the significant increases in capacity over time.

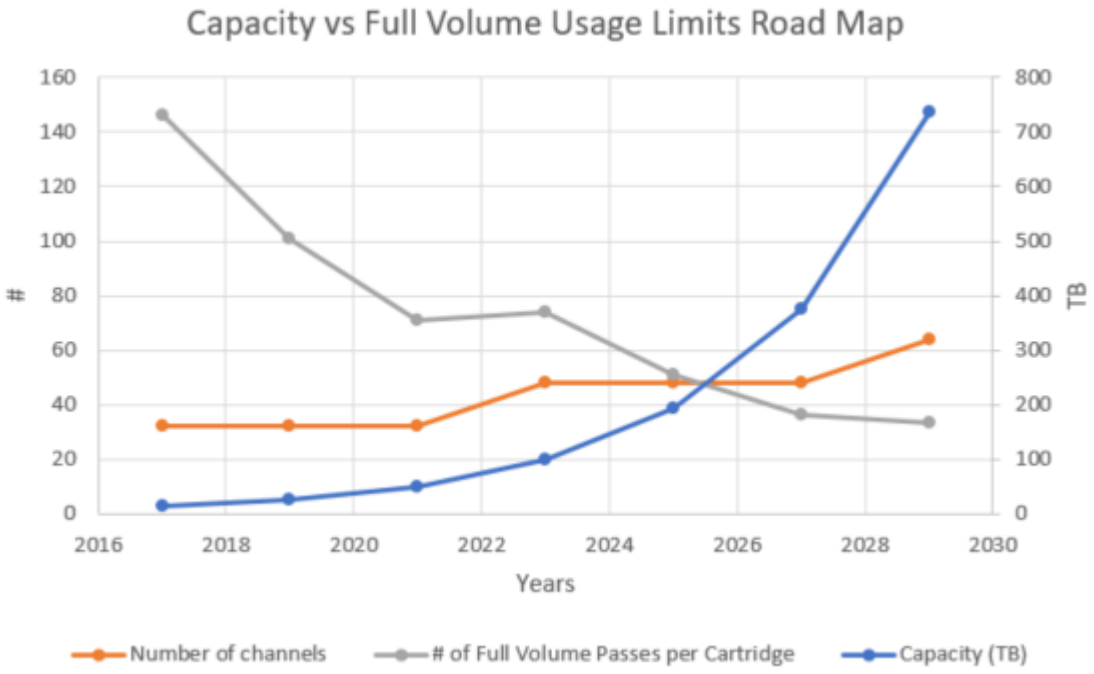

Figure 21 | Capacity versus full volume passes trade-off
Source: INSIC

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.5.5 | Not Compromising Time to First Byte

As the capacity of each tape increases over the course of the roadmap, the time taken to fill a tape increases too. There are benefits to this as it reduces the tape handling (both manual and automated) but the key metric for accessing tape is the time it takes to retrieve the first byte. This must be kept close to current times for the usage models to remain popular.

The main drivers for access time are tape speed and the length of the tape and these are not changing over time in the same way as capacity. The 10 times capacity increases are mostly achieved through bit density improvements so access times remain about the same over time.

Figure 22 shows how the time to first byte increases slightly over the course of the roadmap while the capacity increases by 10 times.

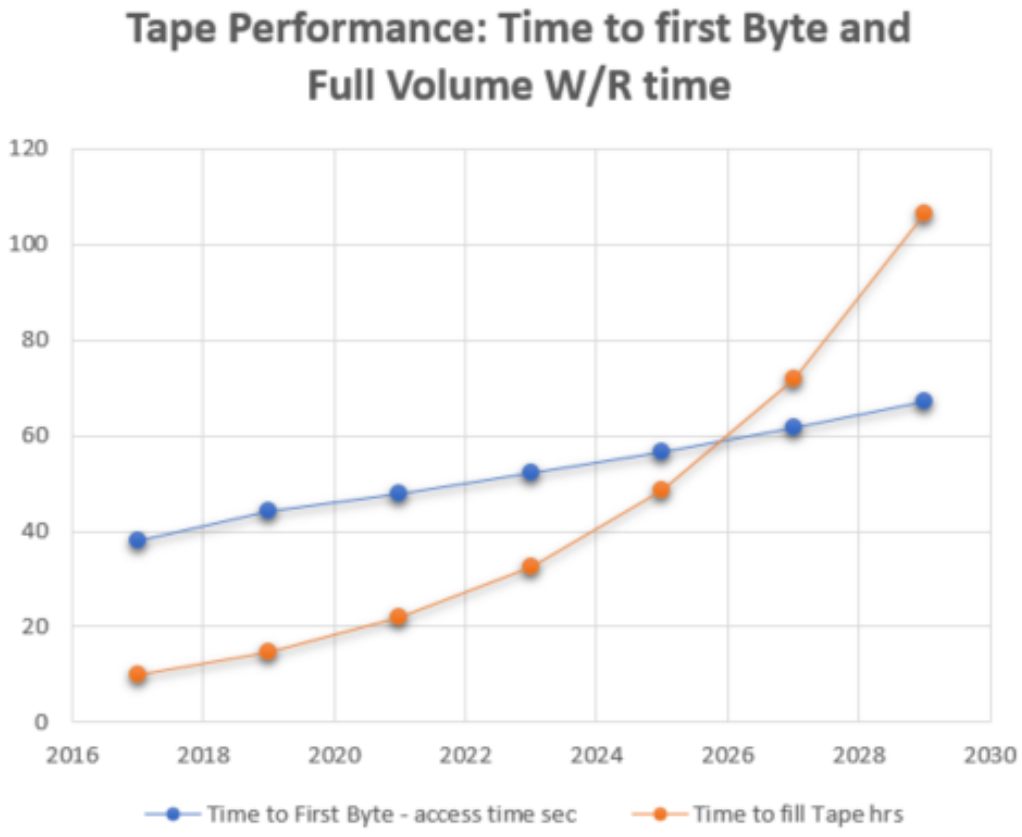

Figure 22 | Time to first byte over time
Source: INSIC

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.6 | CHALLENGES FOR TAPE ADOPTION

While tape has long enjoyed an acquisition cost advantage compared to magnetic disk and optical disk, Mark Peters, an analyst with ESG in an interview with the INSIC panel, stated, “Most of the issues with tape have to do with perception rather than reality. The challenge is really fiction, not function.” One of the primary perception problems tape has faced is ease of use**.** Specifically, tape has been perceived as more difficult to use than disk, requiring special skill sets, complex processes and more manual intervention.

Tape’s challenges in its ease of use stem from two main factors: The first is a direct result of tape’s removable media architecture. Tape derives a significant cost advantage to scale its capacity via its removable media (i.e., the capacity of a tape storage system can be increased by adding cartridges with media, whereas disk requires additional heads, electronics, power supplies, etc., to scale). However, the removable media comes at a cost of support complexity. When a problem is encountered (e.g., failure on a read or write), it is not necessarily obvious if the problem relates to bad media or a bad drive, whereas disk’s lack of removable media results in a clearer identification of a bad FRU (field replaceable unit) that can be replaced. In the past, this was handled with either reviewing service log files by the vendor (covered under warranty or service contract), or manual queries by the user. To overcome this, tape vendors have put a significant effort into analytics that report to the customer as well as to the vendor’s support service. This analytics capability typically identifies bad drives or media as well as drives or media that are starting to perform at a suboptimal level and should be replaced or have diagnostics run on them. These advances in the area of analytics have made it much simpler for a customer to have confidence in their tape storage system and its support.

The second factor making tape harder to use has to do with its interface to the host computer. Disk systems present either a block or file interface to the host. Operating systems and application software are written to support one or both of these interfaces. Tape (as a sequential access system) has required a different driver and also needs the application to understand that the time to access a file may be a minute rather than milliseconds. Many applications, having been written for disk, 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 method is to provide a file system on tape (LTFS) that makes each cartridge appear to an application with a list of files that can be independently accessed just as they would be if they were on disk. This method makes tape as accessible as thumb drives. Users can point and click to drag-and-drop files to and from tape as easy as they move files to and from disk or flash. This means that many customers who required IT assistance for specialized skills to use tape in the past no longer need to do so; they can easily use tape themselves now.

Another method to address 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

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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.

"Tape archive operations is becoming a part time job due to the degree of automation available."

- Phil Goodwin, IDC.

1.7 | INDUSTRY USE CASES FOR TAPE

1.7.1 | Big Data / Analytics

Before the era of big data, enterprises "only" had to worry about the data and systems within their own data center. However, the transformations over the past decades mandated and created a range of additional data sources, both inside and outside an organization, for enterprises to consider.

But, what is Big Data? Big data is currently one of the most critical emerging technologies; Big data is the term used to describe the recent explosion of different types of data from diverse sources. In the market, "Big Data" is characterized by moving applications and data processing closer to the data being processed. The premise is that with so much information, it's more efficient to move the work to the data instead of moving the data to the work. Much of the work involves analyzing existing information to derive new data or insights that represent business, scientific, or technical value.

But even with a scale-out approach to co-locating the analytics and processing work with the data, with expansive amounts of data, not all data is active all the time (or need be active all the time). Analytics or processing work may be batch-oriented, focused on certain projects or sub-areas of data, such that only relevant data need be online and available for the work. All the remaining data may remain accessible, but may tolerate some delay as relevant data is staged in for processing, while other data is staged out when no longer immediately needed.

Keeping all data immediately available and accessible in a Big Data environment can be cost- and space-prohibitive, especially if only a subset of data at a time is actually used for analytics or processing. As such, tape finds a good opportunity to participate within this use case. Data that is not immediately needed for processing can be staged to tape to free up resources for use by data relevant for the current analytics. 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.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Tremendous progress has been made in tape storage technology since the announcement of the first tape drive. Tape can serve as a cost-effective “extended primary” tier for the Big Data/Analytics use case, reserving the more limited data storage resources for the analytics data “working set” instead. As described in earlier sections, the lower cost of tape can also allow a larger amount of flash to be used for analytics, increasing the value of the overall solution.

1.7.2 | Archive

Archive can be defined as a permanent (indefinite or per some time period) record of fixed content. In other words, some information that is being kept per business, scientific, legal, etc. requirements for a specified period of time, and the data is not changing. In many cases, most of the data is infrequently accessed, so the storage selected is often secondary storage instead of primary storage as the access performance is not paramount compared to the overall cost, migration considerations, and data integrity capabilities.[53]

Given the focus on cost and data integrity, tape has been chosen for many of the world’s largest archives for many years.

Some of the archive categories include preservation (discussed below) as well as the following that are covered elsewhere in this paper: rich media, high performance computing and cloud.

Preservation includes records retention, both for corporations (e.g.; corporate filings, human resources records, intellectual property, etc.) and the public sector, including the world’s libraries, which are digitizing their information and offering it online. Tape’s built-in data integrity, low cost, and long shelf life, along with solutions that incorporate file fixity and open standards, has made it a staple in the leading digital archives around the world. The U.S. Library of Congress has posted some information about their solution here: https://www.smpte.org/media-archiving-library-congress

Data integrity is paramount for most digital archivists, whether for preservation, scientific or legal reasons, and this is an area that differentiates tape from other storage technologies.

End-to-end data integrity begins upon ingest of data, and carries protection through all movement of the data to recording it on the final storage medium. End-to-end data protection features common in tape solutions ensure that all data received is in fact recorded correctly. The CRCs are recorded with the data so that the same integrity validation exists for data retrieval (i.e., the host is able to validate that the data retrieved has not been corrupted). In addition to data error detection, tape’s advanced error code correction capabilities are able to correct detected errors and return the original data intact. The T10 Protection Information standard enables end-to-end data integrity in systems built with tape technology. In addition, some tape-based solutions offer the ability to periodically verify the integrity of data stored on tape through automated scanning processes, with options for corrective actions. Additional protection in these solutions also includes file fixity, whereby the archiving application can apply a standard hash code for an entire file, thereby ensuring that the file has not been altered in any way.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

In addition to the data integrity capabilities of tape, the offline security noted in the data protection section of this paper applies to archiving as well.

"Tape's offline capability gives us a layer of protection that other storage options do not."

- Jason Hick, LANL.

Digital archives are typically built with a tiered storage architecture, with flash and/or disk for active data and as a storage buffer, and tape representing the primary, large-scale storage tier. As stated above, tape's low cost coupled with block and file integrity provided by these solutions has made it the architecture of choice for digital archivists.

"With tape, we are confident in our ability to get data off tape that is 15 years old."

- Jason Hick, LANL.

1.7.2.1 | Data Migration

Archiving data requires very large volumes of data to be preserved indefinitely or at least for many years. As a result, that data will likely outlive whatever technology it was initially stored on so must be migrated from time to time until it is no longer needed.

Data migration can be expensive and time consuming but the backwards compatibility features of tape drives can be used to significantly reduce those costs and the whole process can be used positively to check and refresh the data while also validating new tape libraries before they are commissioned for full production use.

A typical disk system is seldom used much longer than five years, but tape can offer customers the option of delaying migrations to much longer times because of its backward read capabilities and because the libraries support many generations of tape drive. For example, given a tape roadmap of two-and-a-half years/generation, two generations of drives manufactured at any given time (i.e., five-year product availability), and five-plus years of support past manufacture, a tape can still be read 10 to15 years after being written. Delaying migrations can reduce the number of migrations needed and therefore reduce cost, albeit at the expense of needing to keep the older, lower capacity media for longer.

To illustrate, below are two different strategies for migrating data using LTO tape drives as an example. An archivist can opt for either extreme or a balance that works best for them:

- **Minimum migration:** the user does as little migration as possible by exploiting the ability of tape drives to read two generations back and /or keeping older technology drives in the library to extend that back even further. As long as there is a tape drive available that can read the older technology tapes, then it is not essential to migrate the data at that time. As libraries are rarely fully populated with drives, it allows older drives to be easily retained for both accessing older technology tapes and for redundancy. Even the older drives are compatible with the newer libraries. This dramatically reduces the amount of data that

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

needs to be migrated, although it doesn't take advantage of the benefits of following the 'minimum media' strategy outlined next. Figure 23 depicts an example of a minimum migration strategy showing an example where **no data needs to be migrated** at all because the combination of drives available can read back all the tapes in the library, just at the expense of needing to retain a larger number of the older tapes. A key point is that although no data has been migrated, it is **all available online**. Notice the **linear** increase to the number of tapes needed despite the **exponential** amount of data being recorded.

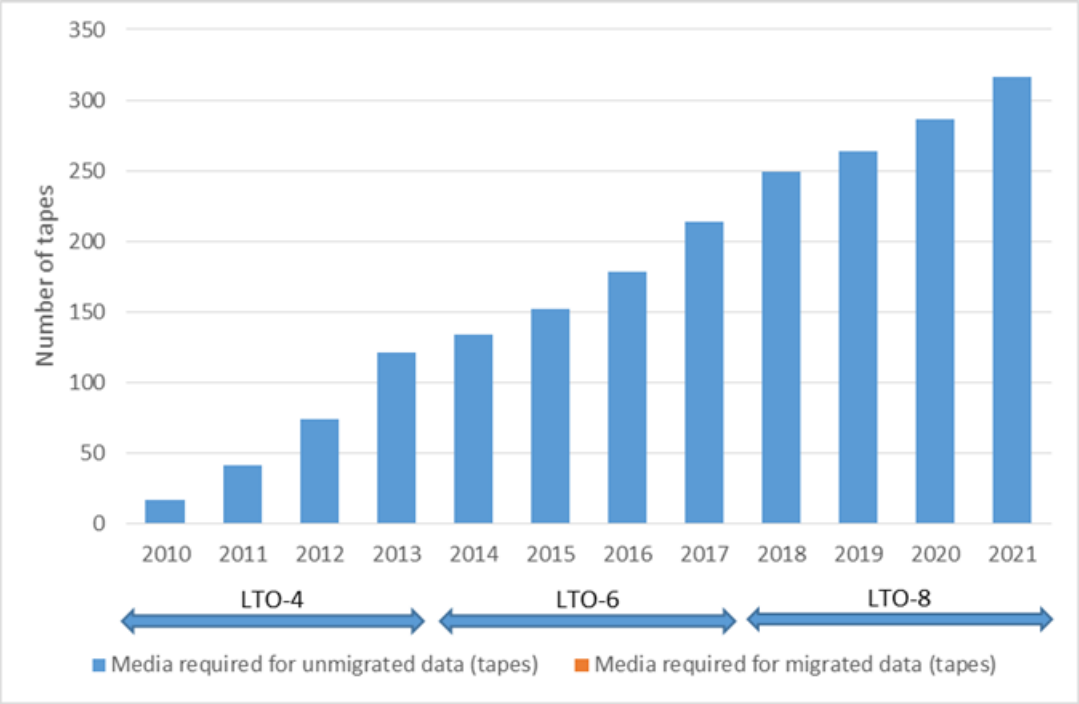

Figure 23 | Minimum Migration Strategy
Source: INSIC

- **Minimum media**: the user migrates data from older technology tapes as soon as newer technology is available. This has the advantage of reducing the number of cartridges needed quite dramatically. A rough calculation based on the amount of new data doubling every two years and new technology available every two to three years with doubled capacity shows that only around 20 percent of the total storage capacity would be taken up by migrated data even if everything was kept online. This would potentially allow the *entire* user's archived data to be available online for a low additional cost which dramatically simplifies the access model. This approach also has the advantage of refreshing the data every few years thereby ensuring confidence that it is in good condition throughout its valued life. Figure 24 depicts an example of a minimum media migration strategy showing a much lower number of tapes needed, although requiring more frequent data migration. In this case the number of tapes **reduces** over time despite the exponential increase of the volume of data recorded.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

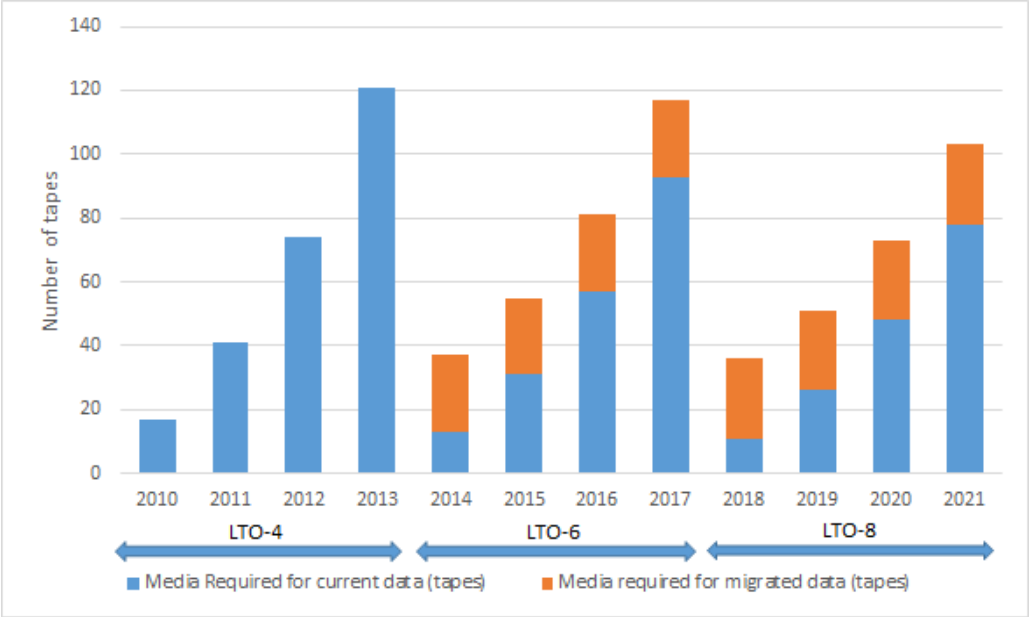

Figure 24 | Minimum Media Strategy
Source: INSIC

It works very well to align this migration process with the introduction of newly acquired, next-generation hardware, providing an excellent way of acceptance testing the new system at the same time. The migration can be made to the new technology at a time the hardware is not yet being used in production thereby allowing full use of all the drives, significantly reducing the time taken. Once complete, the new hardware will have to be fully verified and can be handed over for production use with confidence.

1.7.3 | Backup and Recovery

Data protection solutions, as opposed to archive discussed above, can be defined as solutions designed to restore content that has been corrupted or lost. This includes both business continuance, which includes day-to-day restores of files while continuing business operations, and disaster recovery, which includes massive restores of arrays or even entire data centers.

Tape's role in data protection for disaster recovery is discussed later in this paper. Business continuance includes technologies such as mirroring to enable quick fail-over, ideally transparent to mission-critical applications, and backup/restore, where files that have been corrupted or lost are restored.

The backup and recovery portion of data protection has seen a rapid adoption of purpose-built backup appliances, typically incorporating deduplication technology. IDC estimates this market to be about $7B.[54]. However, the backup and recovery market is not served exclusively by purpose-built backup appliances. Other solutions incorporate continuous data protection. Cloud providers also offer backup as a service.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Tape's role in backup and recovery often includes providing the low-cost tier for longer term backups. As noted in the archive use case, many customers use backups for a dual purpose of both data protection and archival storage. Most major backup applications support disk and tape and allow the customer to set policies for the migration to achieve this objective automatically. Some new continuous data protection solutions also support tape directly for long-term retention, again using policies to automate the process of migration. For customers who don't have a digital archive solution, using the data protection solution to achieve their archive objects is attractive, with tape providing the low-cost, long-term data tier.

Tape is also used in backup and recovery for data that cannot be deduplicated. For example, compressed formats for video, or data that has been compressed and/or encrypted at the source is ideally backed up to tape. Purpose-built backup appliances rely on significant deduplication ratios to reduce their costs, and without the ability to deduplicate the data, they can be cost prohibitive relative to tape.

In addition to the cost consideration, tape provides a secure offline solution for backups. In fact, best practice data protection is often described as the rule of 3-2-1 rule: at least 3 copies of the data, on at least 2 different technologies, and at least 1 copy that is offline and offsite. A recent FBI report highlighted the need for offline copies of data to withstand the growing ransomware issue: "Ensure backups are not connected to the computers and networks they are backing up. Examples might include securing backups in the cloud or physically storing them offline."[55]

So in summary, tape continues to be used in backup and recovery, both for long-term backup retention and also for data that cannot be deduplicated by purpose-built backup appliances. Its offline architecture provides a security advantage not found in other storage technologies.[56]

1.7.4 | Cold Storage

As mentioned earlier in this report, IT consumes about 7% of global electricity, forecasted to rise to 13% by 2030.[57] This alarming increase in power consumption has driven the rise in cold storage solutions whereby the solution is optimized for minimal power consumption at the expense of access time and performance. Electricity costs around the world range from about $.08/kWh in China and India to $0.20 in the UK and $0.35 in Germany. The cost of powering and cooling IT at scale has caused many large-scale storage customers, including cloud, to seriously consider cold storage.

Cold storage architectures are built to minimize power consumption, reserving power for ingest of new data and drawing minimal power for storage of data, typically at a powered off state. It is understood and accepted that there is a significant delay (e.g., minutes to hours) in access to data because of the reduction in power and cost. The cold storage tier is for "cold" data, typically fixed content data that is very rarely accessed. Higher performance storage is used for the "hot" data, but ideally that is a very small percentage of the total storage.

Tape is well suited for cold storage as the removable tape cartridges do not consume power and the digital libraries consume very little power. Tape drive technologies have reduced their

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

idle power as well. Typically, tape is used in a tiered storage architecture inside a cold storage solution, where flash or disk is used at the ingest point as a buffer for tape, and flash or disk is used for the hot data requirements in the overall solution. The solutions typically have an open standard interface, such as OpenStack Swift, NFS, CIFS, etc., so that applications don't need to be modified to use the cold storage solution.

A storage architecture that exploits data access patterns can yield a much more efficient data center solution for both product cost and power and cooling. Such an architecture places fixed content, inactive data in a cold storage solution and reserves the expensive and power consuming storage for the "hot" active data that requires it.

1.7.5 | HSM / Tiering

Different storage technologies have existed in the data center for years as the price/performance characteristics have varied widely. Flash, disk, optical disk, tape, etc., provide a menu of very different choices to a solution architect. Moreover, technologies have all changed over time, making the storage landscape a dynamic place that can offer improvements in performance and cost.

What is required for this dynamic storage landscape is an abstraction layer that hides the underlying storage technologies from the applications. In other words, the applications should be able to perform their tasks without having to understand how or where the data is stored, i.e., the storage technologies are transparent to the application. 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. This is where HSM (Hierarchical Storage Management) or storage tiering comes into play.

Data can be written (ingested) into a high performance tier, which may be very expensive, but offers excellent performance and random access. Policies can be set to migrate that data and/or make copies to less expensive storage tiers, which can include cloud and offsite storage. This can be done for archival purposes, for example, keeping corporate financial and human resources records, or for data protection, such as securing additional copies that are offsite and offline. Data is retrieved from the highest performing tier that has a copy of the data.

Tape is a great fit for the lowest cost tier as it also brings high reliability and low power and cooling requirements. A flash and/or disk tier above it may be the tier that performs most of the reads/recalls, and tape is used for the older, cold data that is seldom recalled. It can also be used as the extra copy in an HSM system for disaster recovery.

"A way users minimize the time to first bye of data is by adding a disk buffer in front of tape."

- Fred Moore, Horison Strategies.

An HSM system can be changed dynamically (tuned), adding in new storage technologies as they emerge, including cloud tiers. Policies automate the process, so the storage administrator

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

can change its performance by changing policies and/or adding storage to the different tiers as required. Most HSM systems keep the metadata and active data on flash and put the inactive data on tape. Most make multiple copies to provide for data protection, with tape often chosen as an offline tier that can also be offsite.

Modern HSM systems also offer “self-healing” where data that has been found to be corrupted (detected by systematic audits or checking the CRC of read data) can be restored from a good copy inside the system, all without the involvement of the application.

1.7.6 | 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. However, in situations when an amount of data to transmit in a timely fashion exceeds the available capabilities of traditional data communication methods, the “bulk” transfer of data via tapes is a viable alternative. To overcome expense, throughput, bandwidth, and/or latency limitations of data communication approaches, data is instead ‘transmitted’ via tape volumes, shipped via private or commercial carriers.

With the advent of LTFS, data portability via tape has been made easier, through use of familiar file-based access and a standardized format. Building upon this simplified and standardized approach, in 2013 the SNIA LTFS Technical Working Group began work to standardize how LTFS could be used for bulk transfer. The LTFS Bulk Transfer Standard specifies workflows for envisioned use cases, XML schema for transfer requests and responses, some guidance for error handling, and considerations for security. (Details of the standard cover packaging of files and objects onto LTFS volumes for transport between destinations, defining how files and objects are merged into a destination name space, and transfer verification and atomicity.)[58]

Part of the premise for using tape for bulk transfer of data is that networks are expensive and still limited in bandwidth and latency. And with the emergence of more network-based storage, such as cloud storage, transfers of large volumes of data are not always feasible within desired amounts of time. By way of example, the SNIA LTFS Technical Working Group estimated that transferring 2 PB of data to a cloud destination over an OC-12 link would take 370 days, at a cost of $61,700,[59] whereas using tape it would take 23 days, and cost $89,200. As another example, they estimated that transferring 10 TB of data over a 10Mbit/sec link would take 92.6 days, at a cost of $1,852, whereas using tape it would take four days, and cost $468 (all estimates based on 2013 values). The spectrum of time savings versus cost savings varies by use case, but the differences are compelling enough to consider tape-based bulk transfer as an option.

- The standard itself considers five variations of the bulk transfer use case:
- Transfer a large volume of data to or from a remote location
- Update or synchronize a subset of data stored at a remote location

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

- Transfer a large volume of data to a second enterprise, or public/ private cloud (e.g., cloud seeding)
- Transfer a large volume of data from a public/private cloud
- Transfer a large volume of data from one public/private cloud to a second public/private cloud [58]

In each of these variations, the common theme is having a large volume of data that exceeds the practical capabilities of the available network resources. Tape is able to provide a high-volume transfer of data very cost effectively, but possibly of greater importance is the faster time in which it can do so (leveraging existing global shipping resources). As a result, tape-based bulk data transfer can be a valuable tool within a data management environment.

For more information on Bulk Transfer, read:

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

1.7.7 | Cloud

Applications for the cloud generate massive data that quickly becomes historical data that can be stored for decades. Adding tape to a cloud infrastructure can eliminate network threats with offline storage, reduce costs, drive business value and create competitive advantages.[60]

Although some cloud service providers (CSP) initially established their services based on disk storage only, it has not taken long for major CSPs to recognize a critical need for less expensive storage in these infrastructures. CSPs are in a constant battle of competitive pricing, driving the cost to the consumer to the lowest possible price while still being able to maintain the operations. In some cases like the devastating closure of Nirvanix [61], have shown that this is a tough market to survive in, with considerable risk to the customer data in some cases. Tape's role in CSPs is two-fold: tape provides a very economical tier as part of the CSP's data protection for its storage and tape enables a low-cost archival storage offering via its very low acquisition and operational costs.

It is important to note that cloud serves as a centralization of smaller infrastructures, taking advantage of the benefits of storage at scale. As such, cloud plays to tape's strengths. Tape's advantages grow at scale, where its low acquisition cost, extremely low power and cooling cost, etc., give it enormous advantages over an architecture built exclusively with disk. CSPs are seeing the benefits of tiering data based on analytic principles of activity and access patterns to lower cost storage. Tape is serving as the lowest cost tier for second and third copies to reduce disk usage and power and as full resident copies of data not being accessed.

CSPs deal with data storage at scales barely imaginable by most organizations and as a result have entertained the other considerations of storage, such as power and especially bit error rate. Power is a fairly simple reasoning; tape drives take little power and tape media take no

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

power. Industry averages of power consumption for tape drives and automation can be readily derived. In addition, according to ESG, the expected 10-year TCO for the LTO-8 data retention solution is 86 percent lower than that of an all-disk solution.[62]

Bit error rate at scale is another factor that is often taken lightly until the data approaches hundreds or thousands of petabytes. To reduce costs, CSPs use bulk purchases of commodity-based disk drives. If a bit error is expected every 110 TB, the management of rewriting data to ensure availability and durability of data consumes massive compute resources. In the instance of tape, the bit error is expected on the order of an exabyte.

The intense competition in the CSP market will drive tiered storage architectures that incorporate tape, primarily because of the cost advantages that tape provides, in addition to the other advantages listed above. One forward-looking analysis by Aaron Ogus, Storage Architect for Microsoft, compared different technologies.[63] Cloud is certainly one of the high growth markets for the future and tape has a major role for cloud storage.

Tape and cloud have made enhancements in capacity, durability and data access capabilities while improving its economic benefits. Therefore, tape technology is perfect to complement cloud service providers and also for organizations that are looking for a cheap, reliable and durable solution to protect their data from cyberattacks.

Organizations have started to rethink using tape as part of their tiered data protection strategies. Many now opt to use all three types of storage - cloud, disk and tape to accomplish different goals. Some see tape as a good offsite backup/recovery option to keep data secure and offline from primary digital networks, while others see it as a good long-term archive option.

Tape technology is the best option for offline storage as data is not connected to the network most of the time. And, given the rising number of cyberattacks and the importance of data protection, moving away from a mixed storage strategy may not be the best. For example, according to a report from security firm RedLock, hackers broke into Tesla's public cloud environment, gaining access to non-public data and stealing compute resources within the company's cloud environment for cryptojacking.[64]

Information stored on tape is almost invulnerable to hackers because of its offline nature. According to TechTarget: “In this case, tape even beats the cloud, which is an increasingly common platform for backup, but which is susceptible to cyberthreats because of its online nature.”[65]

All backup technologies have pros and cons. That is why it important to use a tiered strategy to meet organization’s unique needs. Tape and cloud can address the most critical backup, recovery and archiving requirements, allowing real efficiencies and cost savings.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.7.8 | Disaster Recovery / Air Gap Protection

Tape is still a key element in disaster recovery plans. Disaster recovery (DR) includes the processes, habits, plans and best practices that give an organization the best chance of salvaging critical IT systems and data after a disaster.[66]

DR ensures that data can be restored after any corruption or loss; it generally refers to creating, transporting and preserving copies of production data at alternate locations to protect from disasters affecting the primary copy of data at a primary location. Disaster recovery capability is often further characterized by recovery time objectives (RTO) and recovery point objectives (RPO).

RTO describes the goal for how long it takes to recover from a disaster, and RPO describes the goal for the point in time to which to recover (i.e., how granular a data recovery can be made in terms of days, hours, minutes, seconds, etc.). For business continuance purposes, RTO can be quite small (for example, minutes to hours) and RPO very fine-grained (for example, no loss of business continuity). For other uses, such as recovery of non-business-critical applications or repositories, RTO may be larger (hours to days) and RPO course-grained (hours).

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 is often used, and those snapshots are copied to tape and protected offsite. Alternatively, second copies of operational backups are written to tape, and those copies are sent offsite.

The primary goal of a DR solution is to have copies of data at locations different than, and far enough away from, a site to be protected, but which can be made available again within the RTO. Recovery may not be to the original primary site; it may be at the site with the DR data copies, or at an alternate site, to where the DR data copies need to be sent. The portability and high capacity of tape provides flexibility in the DR recovery scenarios, since the DR data copies can typically be transferred quickly and efficiently via the tape media directly, rather than relying on network infrastructure, which may also have been affected by the originating incident. With tape, all resources needed to recover may be flexibly co-located or relocated as needed to facilitate recovery.

A key requirement for DR is that the data is isolated from the online copies. Best practice, as stated earlier, is summarized as the 3-2-1 rule: at least 3 copies of data, on at least 2 different technologies, and at least 1 copy that is offsite and offline. The last aspect is critical. Architectures where all the data is online are susceptible to software bugs, human errors and sabotage, and enable the data protection copy to be deleted along with the primary data.

Ideally, events that would require use of the DR data copies would never occur, yet that data needs to be available to provide business resiliency. Retaining that data on storage that requires ongoing power and cooling is an expensive proposition, relative to the frequency of use. Tape provides a very cost-effective way to protect DR data copies, since the data can be stored with minimal power and cooling costs.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

In fact, IDC recommends the following as data protection best practices.[67]

- Adopt a continuum of data protection technologies matched to the data protection loss threat and service-level requirements.
- Thwart ransomware risks by ensuring an "air gap" between live data and recovery data.
- Control costs by selecting the lowest-cost technology that neutralizes the threat of data loss and meets the business' availability requirements.[67]

According to Phil Goodwin Research Director within IDC, “Recently, tape has shown itself to be an effective means of providing an "air gap" between live data and protected data.[67] This air gap is essential to thwarting more sophisticated ransomware and malware that attempts to corrupt live, backup, and archive data simultaneously. While some are attempting to write off tape, it has proven to be an important building block of a complete modern data protection plan all at the industry's lowest cost per gigabyte.”

A robust, overall data management plan includes contingencies for disaster recovery, and tape lends itself well to optimize that part of a plan.

1.8 | INDUSTRIES WHERE TAPE IS USED

1.8.1 | Overview

The long-term trend of storing more and more data for longer periods of time shows no signs of slowing. Companies are managing storage capacities that would have been unthinkable a few years ago. Simply throwing more disk at the problem doesn’t work because of cost, power and cooling, and floor space. Budgets remain tight for people and equipment.

Data growth continues to be driven by a variety of factors including Compliance regulations, such as the Sarbanes-Oxley Act in the United States, internal retention policies, business requirements and the growth of unstructured content through the increasing use of imaging. In many industries, digitization has seen a move from analog imaging techniques to digital imaging and also a continued increase to higher quality digital resolutions. There has also been a growing need for long-term digital archiving of all assets.

The requirements for digital archives include ingesting, searching and retrieving data; mobility, accessibility and integrity of that data for long periods of time, typically measured in years. It can also be very difficult to know what to save and what to discard, or how valuable an asset will be in the future. Tape has been integrated into solutions that provide these key requirements, including standard interfaces such as NFS, CIFS, OpenStack Swift, etc. and given its low total cost of ownership, is key to organizations being able to meet their archival needs.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.8.2 | HPC/Scientific

HPC (High Performance Computing) applications tend to be used in the field of computational science for a wide range of computational-intensive tasks, including quantum mechanics, weather forecasting, climate research, seismic studies, oil and gas exploration, molecular modeling, physical simulation, etc. These applications handle large volumes of data over short periods of time running on Supercomputers with a high level of computational capacity.

Data created through these applications requires careful storage and long-term archiving. This data represents significant investments and in many cases experiments or scenarios that are very difficult and prohibitively expensive to replicate. New advances in technologies, algorithms and discoveries could provide new ways of analyzing this data without having to recreate the original study. Also, new theories and models can be validated and refined by comparing against historical data. Hyperion Research projects that the total storage for HPC will exceed $7.5B by 2022.[68]

Tape technology is typically used in these industries in a tiered storage solution. Often a flash or disk tier provides a buffer for data on its way into the tape archive, optimizing the tape archive, and also as an area for analysis of data that is recalled, such as when applying modeling. Tape's use in HPC as nearline storage and as a data archiving solution, scaling out as the storage requirements grow, provides low cost of ownership, long shelf life, and high performance and capacity in a small footprint.

"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.

1.8.3 | Cloud

Cloud provides a variety of use cases for which tape is being used as an integral part of the storage hierarchy, including data protection and archive. Overall, Mordor Intelligence projects that the global cloud storage market will reach $101.6B in 2023.[69]

Cloud offerings provide a host of benefits to customers. Tape's role is typically part of the cloud provider's data protection architecture. The consolidation of many customers' or departments' data into a private or public cloud presents the kind of large-scale data protection problem for which tape's high throughput, high reliability, high capacity and low cost are ideal. Often the data protection architecture includes a disk tier with deduplication and a tape tier for the older, longer-term backups and also as an offline storage tier that is the last line of defense against human error or an attack. Tape's ability to provide transparent encryption without sacrificing performance is typically part of this data protection architecture, providing another level of security from attacks.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Cloud Infrastructure as a Service presents another use case for which tape is ideal. Providers such as Amazon, Microsoft, Google, and Oracle, among others, are offering a very low cost, deep archive tier. The service level agreements for this type of storage are typically measured in hours and the competitive landscape is centered on price and reliability. For these offerings, tape is typically used as the 'final' storage technology, often behind a flash or disk tier that provides a "landing zone" for data on ingest. Tape provides both very low CAPEX (low cost per PB) and also very low OPEX (extremely low energy and low management cost per PB) that have enabled these cloud archive tiers.

1.8.4 | Media & Entertainment

Some of the main factors driving storage capacity and management growth in the media and entertainment industry are the requirements for creation of new content with higher resolution formats such as 4K and 8K. A single 6K camera produces 85 TB in 24 hours and new offerings are emerging that offer the consumer the option to select their point of view, increasing the number of cameras required. Other factors impacting storage demands are the digital distribution, conversion and presentation as well as archiving of new content. Total media and entertainment storage revenue is forecasted to reach $12.8B in 2023 with about 34 percent of the storage being archival and preservation.[70]

Studios are not just challenged by the constant increase in capacity per production, but also they see an increase in the number of productions per year and the addition of more digital effects put into each production.

Tape technology provides a low-cost, scalable and easy-to use-solution from capture through distribution of assets. In production, tape technology is used to move data on set or on location directly to tape for direct asset archive or to share content. LTFS is an open standard that is common for this use case. In Post Production, content stored on tape can be brought online for editing; work can then be stored back on tape once complete. The SMPTE-defined AXF format and the LTFS format are archive standards for this use case. Open specification, cross-platform interchangeability, reliability and robustness make tape technology an ideal media for distribution. A single tape cartridge can hold hours of content for broadcast distribution. Digital archiving is a key step in asset preservation for future re-purposing and re-monetization. Long-term data retention, reliability, cost-effectiveness and low maintenance make tape technology the ideal medium to preserve entire projects.

Tape is uniquely placed to protect and retain digital content in the media and entertainment industry, being low cost, scalable, durable, and reliable. Furthermore, with open standards such as LTFS and AXF, data is easy to use, access and share across the industry.

1.8.5 | Healthcare

In the healthcare industry, the moves to electronic health data sets and from film imaging to digital imaging are well underway. The exponential growth of medical images and the reticence

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

to delete patient data, combined with legal requirements like the Health Insurance Portability and Accountability Act (HIPAA), and also the need for historical trend analysis and research are all contributing to the overall challenge of efficiently and cost effectively storing data as well as maintaining its access in the future.

The main concerns in healthcare regarding archiving and management of medical images, patient records and administrative data include: how to utilize its value to improve quality of care, how to share all this information, how to secure it and how to maintain and access it in the long-term.

A tiered storage approach is typically the optimal strategy, with the most recent data stored online in hard drives, and data moved to tape as it ages, where it is kept on an affordable, secure, reliable and easy-to-use media. Tape with LTFS simplifies the health data lifecycle by easily integrating with any custom or out-of-the-box applications used by departments across the practice. Tape with LTFS allows clinicians and administrators to quickly retrieve and share medical data that is typically only stored for a short time. These records can be shared without workflow impact to any other department in the practice, and without conflict with any vendor-specific medical data and image software already in use. Tape also simplifies the storage of large medical images that are not easily deduplicated. When it comes to transmitting large healthcare data files such as x-rays, tape doesn't consume resources that affect users and individual files are readily accessible.[71]

1.8.6 | 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 systems used for competitive advantage. For example, video surveillance is now being used to analyze customer behavior. The quality of the images has been increasing – High Definition video and 30 frames-per-second are now standard for many applications – and 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. The average data generated daily by new surveillance cameras shipped globally is estimated to be over 2.5 exabytes in 2019.[72]

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.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.8.7 | Traditional IT

Traditional IT uses tape in two broad categories. The traditional role of tape in data protection has shifted to a tiered solution that incorporates disk and tape. Disk is typically used for fast restores of files, often incorporating deduplication technology to minimize cost, and also as a staging area before moving the data to tape. Tape is used for two key aspects of a data protection solution: The long-term retention of older backups and an offline copy that provides an additional degree of protection from human error, sabotage, cyberattacks and ransomware, and environmental threats. In short, tape's low cost and offline characteristics make it ideal for disaster recovery. Most major data protection applications manage the movement of data across the disk and tape tiers according to policies, providing an automated solution with minimal manual activity. Other solutions incorporate continuous data protection to provide a Recovery Point Objective and Recovery Time Objective of near zero. These solutions also often support staging data to tape, based upon policies for the same reasons mentioned above. Some customers ship tapes offsite as part of their disaster recovery architecture, while other companies maintain their own tape vault at a second or third site.

Traditional IT incorporates tape for archiving a variety of corporate records, ranging from email and office documents to training videos and more. In these solutions, tape is one of the storage tiers incorporated in a system that presents a standard interface, such as NFS, CIFS, OpenStack Swift, etc. This allows the applications to move files to and from the archive without knowing the actual storage technology used. The key benefit to this approach is that the applications work with the system as they would with a typical NAS or cloud solution, and the archive solution can migrate data across the storage tiers, based on policy, in the background. Archival solutions also migrate data to private or public clouds and can migrate from one generation of technology to another, all without any changes to the applications that wrote the data.

HTF projects that the NAS market size is growing at an estimated CAGR of over 20 percent and will grow to $45.2B by 2023.[73]

1.8.8 | Oil & Gas/Seismic

Data volumes are exceeding 10 TB per day for a single oil well.[74] Given the high value of the data as well as the high cost to collect it, along with significant research and development to improve the modeling forecast, this is a "true" archive market with a very high proportion of the data collected saved as an archive.

Some industry trends that are driving additional storage requirements are higher resolution capture methods, such as wide azimuth or multi-azimuth. A 3D survey yields 10 times the storage requirements compared to a 2D survey format. The use cases involve both the initial data acquisition as well as the computational systems running advanced modeling algorithms, and long-term archiving.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

The role of tape in this market starts with data acquisition, where the data from sensors is often streamed to tape. Tape's excellent streaming performance and portability make it ideal for this application. Petabytes of data on tape are easily transported to a data center where they are then copied as part of the data protection solution that provides for disaster recovery. The data can be stored in a digital archive on tape because of its economic advantage over other storage tiers. Data is then retrieved from tape to disk for modeling software to analyze. The movement of the data from tape to disk is automated by middleware that has been integrated with the analysis software. Like the markets discussed above, this middleware typically presents a NAS image to make it easy to integrate with the application software.

1.9 | 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 = 10^{21} 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 particular attributes that must be optimized in a given 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 also 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.

"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

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

(by revenue) going forward for tape."

- Phil Goodwin, IDC.

Tape has a number of 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 a number of 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, in order 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 cost 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.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

1.10 | 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.

In memory of Jon Toigo who passed away Feb.12, 2019, and left his unique stamp on the storage industry with a focus on long-term data retention and disaster recovery. More importantly, he focused on the needs of the end users who required help understanding the technology and the channel partners who implemented it. He will be dearly missed.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

About The Information Storage Industry Consortium

The Information Storage Industry Consortium (INSIC) is the research consortium for the worldwide information storage industry. INSIC membership consists of corporations, universities and government organizations with common interests in the field of digital information storage. Corporate membership includes major information storage product manufacturers and companies from the storage industry infrastructure. INSIC has its headquarters in San Diego and was originally incorporated in April 1991 as NSIC, a non-profit mutual benefit corporation.

For further information or to contact INSIC please go to http://www.insic.org/

Copyright ©2019 INSIC. All rights reserved worldwide. All other trademarks and registered trademarks are property of their respective owners. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, without prior written permission. All opinions in this paper are those of INSIC and based on information from company material, industry reports, news outlets and interviews.

1.11 | REFERENCES

[1] UNFCCC. (2016, August 12). ICT Sector Helping to Tackle Climate Change. Retrieved from https://unfccc.int/news/ict-sector-helping-to-tackle-climate-change

[2] Sadler, R. (2017, January 31). Video demand drives up global CO2 emissions. Retrieved from https://climatenewsnetwork.net/video-demand-drives-global-co2-emissions/

[3] Buffington, J. (2016, February). Analyzing the Economic Value of LTO Tape for Long-term Data Retention. Retrieved from https://page.dternity.https://www.lto.org/wp-content/uploads/2014/06/ESG-WP-LTO-EVV-Feb_2016.pdfnet/TCO.html

[4] Reinsel, D. (2018, November). The Digitization of the World from Edge to Core. Retrieved from https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

[5] Inside HPC (2019, February 11). Moving Mountains of Data at NERSC. Retrieved from https://insidehpc.com/2019/02/moving-mountains-of-data-at-nersc/

[6] Reinsel, D. (2018, November). The Digitization of the World From Edge to Core. Retrieved from https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

[7] Zenith The ROI agency (2018, July 16). Online viewing to exceed an hour a day in 2018, Retrieved from https://www.zenithmedia.com/online-video-viewing-to-exceed-an-hour-a-day-in-2018/

[8] Ultrium LTO (n.d.). Are you aware of the hidden costs of storing data in the cloud?. Retrieved from https://www.lto.org/resources/tcotool/assumptions/

[9] Swan, E. (2018, April 23). Document Retention Requirements Per Industry. Retrieved from https://www.efilecabinet.com/document-retention-requirements-for-various-industries/

[10] Scott, S. (2017, October 12). Internet Crime Report. Retrieved from https://pdf.ic3.gov/2017_IC3Report.pdf

[11] Moore, F. (n.d.). The Tape Air Gap. Retrieved from https://horison-groflyllc.netdna-ssl.com/wp-content/uploads/2017/10/Tape-Air-Gap.pdf

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

[12] Decad, G., & Fontana, R. (2017, July). A Look at Cloud Storage Component Technologies Trends and Future Projections. Retrieved from http://ibmsystemsmag.com/mainframe/storage/support/cloud-trends-projections

[13] Wikipedia contributors (2019, February 20). Storage - OpenCompute. Retrieved from https://www.opencompute.org/wiki/Storage

[14] Western Digital (2017, October 11). Western Digital Unveils Next-Generation Technology To Preserve And Access The Next Decade Of Big Data. Retrieved from https://www.westerndigital.com/company/newsroom/press-releases/2017/2017-10-11-western-digital-unveils-next-generation-technology-to-preserve-and-access-the-next-decade-of-big-data

[15] Fink, M. (n.d.). Innovation at Western Digital: The Data Revolution Begins Here. Retrieved from https://www.westerndigital.com/company/innovations

[16] USENIX (n.d.). USENIX-Slides brewer. Retrieved from https://www.usenix.org/legacy/sites/default/files/conference/protected-files/fast16_slides_brewer

[17] Wikipedia contributors. (2019, March 25). IBM (atoms) - Wikipedia. Retrieved from https://en.wikipedia.org/wiki/IBM_(atoms)

[18] Strachowski, N. Nanomechanical data storage device from ZRL. Retrieved from https:/www.zurich.ibm.com/news/05/millipede.html

[19] University of Southhampton (2018, February 7). 5D storage crystal joins Tesla Roadster on incredible space journey. Retrieved from https://www.southampton.ac.uk/news/2018/02/spacex-5d-crystal.page

[20] Marchon, B. (2017, June 9). DNA Storage. Retrieved from https://wdc.app.box.com/s/pmmvf8r70jnod5w3fiwrv1ccidfarbgy

[21] University of Southampton. (2016, February 18). Eternal 5D data storage could record the history of humankind. Retrieved from https://www.southampton.ac.uk/news/2016/02/5d-data-storage-update.page

[22] Wikipedia contributors. (2019b, April 14). 5D optical data storage - Wikipedia. Retrieved from https://en.wikipedia.org/wiki/5D_optical_data_storage

[23] Hewlett Packard Enterprise. (2017, May 3). Memory-Driven Computing Explained. Retrieved from https://news.hpe.com/memory-driven-computing-explained/

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

[24] Wellise, C. (2018, June 5). Energy efficient technologies for the digital economy. Retrieved from https://www.hpe.com/us/en/insights/articles/energy-efficient-technologies-for-the-digital-economy-1806.html

[25] Gartner, Predicts 2018: Storage, (February 2018). Retrieved from https://www.gartner.com/doc/3860375/predicts--storage

[26] Wiki OpenStack contributors. (n.d.). Swift/HighLatencyMedia - OpenStack. Retrieved from https://wiki.openstack.org/wiki/Swift/HighLatencyMedia

[27] 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

[28] Bernstein, D. (2018, August 18). Introduction to post-quantum cryptograph. Retrieved from http://pqcrypto.org/www.springer.com/cda/content/document/cda_downloaddocument/9783540887010-c1.pdf

[29] 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

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

[31] Mavroeidis, V. (n.d.). The Impact of Quantum Computing on Present Cryptograph. Retrieved from https://arxiv.org/pdf/1804.00200.pdf

[32] Decker, S. (2018, October 17). Bloomberg - The Future of the Cloud Depends on Magnetic Tape. Retrieved from https://www.bloomberg.com/news/articles/2018-10-17/the-future-of-the-cloud-depends-on-magnetic-tape

[33] Brewer, E., & Lawrence, Y. (2016, February 23). Disks for Data Centers. Retrieved from https://storage.googleapis.com/pub-tools-public-publication-data/pdf/44830.pdf

[34] Brewer, E., & Lawrence, Y. (2016, February 23). Disks for Data Centers. Retrieved from https://storage.googleapis.com/pub-tools-public-publication-data/pdf/44830.pdf

[35] Brewer, E. (2016, February 23). Disks and their Cloudy Future. Retrieved from https://www.usenix.org/sites/default/files/conference/protected-files/fast16_slides_brewer.pdf

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

[36] Rahul, P., & Navendu, J. (2013, October 3). When the Network Crumbles: An Empirical Study of Cloud Network Failures and their Impact on Service. Retrieved from https://54e57bc8-a-62cb3a1a-s-sites.googlegroups.com/site/2013socc/home/program/a15-potharaju.pdf?attachauth=ANoY7crEU9xw2LxKYBBak-JwHJToILLmFfy204eE4UVRDgrr9vpnWafSnJR4RkxbRSA_2fCpzb2m1x_v3g2ty70B9w9TxqxMb79DVm8Tc-43g1z63qTh1SGmaHeL6znKp-p8N63QzgKmESPDHIrOwH2SuUZjK8HNFf9PpnbqzWBPWd-_CO3YNzjsUlmHHLnaoHNISMvdr9eMZAI2CMZLIAaW27uUZGMDzd9DR6ftNQE92JjNbm9j02s%3D&attredirects=1

[37] Huang, C., Simitci, H., Xu, Y., Ogus, A., Calder, B., Gopalan, P., . . . Yekhanin, S. (November, 2008). Erasure Coding in Windows Azure Storage. Retrieved from https://www.cs.princeton.edu/courses/archive/spring13/cos598C/atc12-final181.pdf

also presented at USENIX 2012 annual technical conference

Huang, C. (2012, June 12). Erasure Coding in Windows Azure Storage | USENIX. Retrieved from https://www.usenix.org/conference/atc12/technical-sessions/presentation/huang

[38] Beach, B. (2018, February 16). Backblaze Open Sources Reed-Solomon Erasure Coding Source Code. Retrieved from https://www.backblaze.com/blog/reed-solomon/

[39] Amazon. (n.d.). Product Details. Retrieved from https://aws.amazon.com/vmware/details/

[40] SNIA. (n.d.). Software Defined Storage | SNIA. Retrieved from https://www.snia.org/sds

[41] Saul, C. (2019). Software-Defined Storage for Dummies, Second Edition. Retrieved from https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=TSM03044USEN

[42] Arcilla, A. (2018, August). Quantifying the Economic Benefits of LTO-8 Technology. Retrieved from https://www.lto.org/wp-content/uploads/2018/08/ESG-Economic-Validation-Summary.pdf

[43] Buffington, J. (2016, February). Analyzing the Economic Value of LTO Tape for Long-term Data Retention. Retrieved from https://www.lto.org/wp-content/uploads/2014/06/ESG-WP-LTO-EVV-Feb_2016.pdf

[44] Reine, D. (2015, July 15). The clipper group calculator. Retrieved from http://www.clipper.com/research/TCG2015006.pdf

[45] Arcilla, A. (2018, August). Quantifying the Economic Benefits of LTO-8 Technology. Retrieved from https://www.lto.org/wp-content/uploads/2018/08/ESG-Economic-Validation-Summary.pdf

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

[46] Bonner, F. (n.d.). Data Storage Technology, Tape, and Sustainability. Retrieved from https://www.lto.org/wp-content/uploads/2014/06/Data-Storage-Technology-Tape-and-Sustainability-Sept-2017.pdf

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

[48] Schroeder, B. (n.d.). Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? Retrieved from http://www.cs.toronto.edu/~bianca/papers/fast07.pdf

[49] The Tape AFR was based on the report by National Energy Research Scientific Computing Center NERSC; https://www.nersc.gov/assets/pubs_presos/AAA-Case-Study-NERSC-FINAL2-6-12.pdf

[50] AFR estimates were based on the study done by the University of Illinois under a grant by NetApp corp. and a presentation at USENIX; https://www.usenix.org/legacy/events/fast08/tech/jiang.html

https://www.usenix.org/legacy/events/fast08/tech/full_papers/bairavasundaram/bairavasundaram.pdf Parity Lost and Parity Regained, a study done by the University of Wisconsin-Madison under a grant from NetApp; https://www.usenix.org/legacy/events/fast08/tech/krioukov.html

[51] Krioukov, A. (n.d.). 6th USENIX Conference on File and Storage Technologies (FAST '08). Retrieved from https://www.usenix.org/legacy/events/fast08/tech/krioukov.html

[52] Fuji technical paper on tape archival life; Retrieved from: http://fujistore.hu/wp-content/uploads/2014/11/Macomp-technical-hotline-for-storage-tapes-November-2014.pdf

[53] Moore, F. (2018). Archival data storage. Retrieved from https://horison-groflyllc.netdna-ssl.com/wp-content/uploads/2018/09/Fujifilm-Archival-Upheaval-2018-Final.pdf

[54] IDC. (2018). Worldwide Data Protection and Recovery Software Forecast, 2018â□“2022. Retrieved from https://www.idc.com/getdoc.jsp?containerId=US44079518

[55] Federal Bureau of investigation. (2016, September 15). Internet Crime Complaint Center (IC3) | Ransomware Victims Urged to Report Infections to Federal Law Enforcement. Retrieved from https://www.ic3.gov/media/2016/160915.aspx

[56] SNIA. (2017, November). Data Protection Best Practices Technical White Paper | SNIA. Retrieved from https://www.snia.org/education/whitepapers/dataprotection-wp?utm_source=SNIA+Email+List

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

[57] Climate news network. (2017, January 31). Video demand drives up global CO2 emissions. Retrieved from https://climatenewsnetwork.net/video-demand-drives-global-co2-emissions/

[58] SNIA LTFS Bulk Transfer Standard. https://www.snia.org/sites/default/files/technical_work/LTFS/LTFS_Bulk_Transfer_1_0.pdf

[59] SNIA Applications of LTFS for Cloud Storage Use Cases. http://snia.org/sites/default/files/CloudTapeUseCases_v1.0.pdf

[60] IBM Employees. (2018). WHY IBM TAPE FOR CLOUD? Retrieved from https://www.ibm.com/downloads/cas/MALZYQNK

[61] ServerWatch (2013) Nirvanix Shut-Down Sends Shockwaves through the Cloud Services Industry. http://www.serverwatch.com/server-trends/nirvanix-shut-down-sends-shockwaves-through-the-cloud-services-industry.html

[62] Arcilla, A. (2018, August). Quantifying the Economic Benefits of LTO-8 Technology. Retrieved from https://www.lto.org/wp-content/uploads/2018/08/ESG-Economic-Validation-Summary.pdf

[63] Zettabyte data storage [Data at Scale 10] Retrieved from https://youtu.be/4hoc7tjwtu8

[64] DeNisco, A. (2018, February 20). Tesla public cloud environment hacked, attackers accessed 'non-public' company data. Retrieved from https://www.techrepublic.com/article/tesla-public-cloud-environment-hacked-attackers-accessed-non-public-company-data/

[65]Crocetti, P. (n.d.). What is magnetic tape storage? - Definition from WhatIs.com. Retrieved from https://searchdatabackup.techtarget.com/definition/magnetic-tape

[66] TechTarget. (n.d). The role of tape in disaster recovery. Retrieved from https://searchdisasterrecovery.techtarget.com/feature/The-role-of-tape-in-disaster-recovery

[67] Goodwin, P. (2018, April). Using Tape to Optimize Data Protection Costs and Mitigate the Risk of Ransomware for Data-Centric Organizations. Retrieved from https://www.lto.org/wp-content/uploads/2018/04/US43710518.pdf

[68] SlideShare User. (2018, May 2). HPC Market Update from Hyperion Research. Retrieved from https://www.slideshare.net/insideHPC/hpc-market-update-from-hyperion-research

[69] Mordor Intelligence. (2018). Cloud Storage Market. Retrieved from https://www.mordorintelligence.com/industry-reports/cloud-storage-market

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

[70] Coughlin, T. (2018, August). 2018 Digital Storage for Media and Entertainment Report. Retrieved from https://tomcoughlin.com/wp-content/uploads/2018/08/ME-Storage-Report-2018-Brochure-082118.pdf

[71] Ultrium LTO. (2018, May). AFFORDABLE, EFFICIENT, COMPLIANT HEALTH DATA STORAGE THROUGH LTO WITH LTFS. Retrieved from https://www.lto.org/wp-content/uploads/2018/02/LTO18_0570_ValueProp_Healthcare_2018.pdf

[72] IHS Markit. (2014, June). Caught On Tape, Now Keep It Secure. Retrieved from https://www.lto.org/wp-content/uploads/2014/06/IHS-Markit-Whitepaper-Caught-on-Tape-Keep-it-Secure-highres.pdf

[73] MarketWatch. (2018, November 5). Network-Attached Storage (NAS) market size will grow upto USD 45.2 Billion by 2023, at an estimated CAGR of 20.2%. Retrieved from https://www.marketwatch.com/press-release/network-attached-storage-nas-market-size-will-grow-upto-usd-452-billion-by-2023-at-an-estimated-cagr-of-202-2018-11-05

[74] Amery, N. (2018, August 21). Data: Is the oil and gas industry's most valuable resource being overlooked? Retrieved from https://ihsmarkit.com/research-analysis/Data-is-oil-gas-industry.html

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited