



ESG WHITE PAPER

How Tape Technology Delivers Value in Modern Data-driven Businesses

IBM and Fujifilm

By Vinny Choinski, ESG Senior Analyst; and Christophe Bertrand, Senior Analyst

September 2021

This ESG White Paper was commissioned by IBM and Fujifilm and is distributed under license from ESG.



Contents

Introduction.....	3
New Technology Overview	4
Hyperscalers' Influence on Modern Data Center Architectures	5
Infrastructure Enhancements.....	6
Use Cases	7
The Bigger Truth	9

Introduction

This ESG White Paper focuses on exciting, new advances in tape technology that are now positioning tape for a critical role in effective data protection and retention, as we have now officially entered the zettabyte age of storage. Organizations are challenged with the overwhelming amount of data being created, as well as TCO, energy consumption, cybersecurity threats, scalability, reliability and how to effectively protect all of the company data. Tape, in many cases, is the answer to these challenges and the last line of defense against cyber-attacks. Tape technology can play a major role in any organization’s data protection and storage strategy.

In a recent ESG research survey, ESG asked participants what benefits they realized from using tape technology. The two most common benefits identified were proven reliability (46%) and proven security (44%) (see Figure 1).¹ These benefits represent some of the core values of tape as a part of any data lifecycle management program. Tape has always been viewed as a reliable solution since its creation in 1952, when one reel could hold about 2MB of data, up to today with a single cartridge capacity of 20.0 TB. The ability to isolate tapes offline delivers a strong security component of cyber resiliency when implemented properly. Total cost of ownership, better storage efficiency, and better media longevity when compared with alternatives all ranked very high with ESG research respondents as business benefits of using tape as a data protection and retention technology. The long-term value is extremely high, as the cost for management is low and life expectancy exceedingly long. Tape has been on a steady path of increasing capacity and continues to promise dramatic future capacity increases with the recent announcement by IBM and Fujifilm. Now, the two firms together have demonstrated their ability to meet or exceed capacity roadmaps with their groundbreaking 580TB technology.

Figure 1. Top Seven Benefits of Tape Technology



Source: Enterprise Strategy Group

Although tape is a very mature technology, it has continued to improve and remains a cost-effective way of retaining large amounts of data at less than a penny per gigabyte. It also helps that tape has proven to be readable for many decades and offers the advantage of implementing a physical barrier thanks to the ease of removability and portability not easily achieved by other storage solutions. Furthermore, stealing data stored on tapes would mean actually stealing the physical


¹ Source: ESG Research Report, [Tape’s Place in an Increasingly Cloud-based IT Landscape](#), January 2021.

tapes and having access to the right equipment to read them. Participants in the ESG research survey felt a sense of comfort from using tape and having tape as part of their data protection and archiving strategies.

New Technology Overview

It is estimated that there will be as much as 17 ZB of persistent data needing to be stored by 2025. That's 17 trillion gigabytes. There needs to be a cost-effective way of storing it all. IBM and Fujifilm have recently broken the areal density world record on linear magnetic tape with their new 580TB capacity tape demo, which is an enormous jump in the amount of data capacity previously demonstrated. In comparison, this achievement represents nearly 48 times greater capacity than current cartridges (LTO8) and is enough to store data equivalent to 120,000 DVDs on a single tape. The tape prototype used in the demo has an areal density of 317Gbit/sq in. In order to achieve such a high density, Fujifilm's research teams had to develop a brand-new magnetic particle, employing Strontium Ferrite (SrFe) in the process. The most recent generations of magnetic tapes have relied on Barium Ferrite (BaFe), but SrFe offers the potential for much higher density storage on the same amount of linear tape. Yet, unlike alternative coating methods such as sputtering, SrFe can be produced on the same production lines as BaFe, which will help deliver a significant cost advantage in bringing SrFe to market.

Figure 2. Areal Density Improvements



Source: Enterprise Strategy Group

Figure 2 shows how tape is outpacing disk capacity through areal density improvements over time. As of today, LTO 8 is shipping at an average price of \$79 per cartridge with a native storage capacity of 12TB. This scenario delivers a \$6.58 per TB cost for native data. In comparison, disk today has a native capacity of 18TB for \$683, which is \$37.94 per TB. Currently, there is a very clear price differential between tape and disk. Now, the IBM and Fujifilm demonstration dramatically increases potential capacity to a level that was not expected for many years.

Some of the key highlights include:

- **Increased Density:** 580TB capacity with an aerial density of 317Gbit/sq in.

- **New Process:** Changed from Barium Ferrite (BaFe) to Strontium Ferrite (SrFe) to achieve improved density, resulting in higher capacity. SrFe magnetic particles are 60% smaller than BaFe yet maintain better magnetic signal strength and archival life than BaFe.
- **High Signal to Noise Ratio:** Fujifilm's proprietary NANOCUBIC magnetic tape technology uses SrFe magnetic particles and a disbursement formula that evenly distributes and arranges the ultra-fine SrFe magnetic particles at nanoscale to achieve a high signal-to-noise ratio.
- **Tape Head Advancement:** To take advantage of the use of narrower tracks (56.2nm wide compared to typical 103nm), a new servo technology was created that enables the read-write heads to locate and follow the narrower tracks as the tape moves across the heads.

In conjunction, the IBM team developed a family of new servo-mechanical technologies, including a new servo pattern, pre-recorded in the servo tracks; a new head actuator; and a set of servo controllers. The end result is a very high-capacity tape solution that can read data while moving at a very high rate of speed. This is an important breakthrough since hybrid clouds will rely on magnetic tape for decades to come. Disk drive areal density growth has been challenged as areal densities have reached the superparamagnetic thresholds. Now, even with a lower areal density (317 vs 1,000+) compared to HDD, tape is poised to become the predominant media leveraged to keep pace with the rise in unstructured data backup and archiving needs. According to the Information Storage Industry Consortium (INSIC) trendline for tape areal density growth, tape has a steeper line, at 34% a year, than hard disk drive (HDD) technologies, which has a forecast areal density growth of just 7.6% a year.

The new demonstration from IBM and Fujifilm also represents a major advantage for hyperscalers and other organizations that rely heavily on data storage technologies to operate. Tape, like disk, is storage. But HDD has been unable to demonstrate the type of high-capacity capabilities that the IBM and Fujifilm technology demonstrated. The level of areal density achieved by IBM and Fujifilm demonstrates the ability to deliver capacities that extend far beyond the projected status of not only disk roadmaps, but also the present LTO and enterprise tape roadmaps.

Hyperscalers' Influence on Modern Data Center Architectures

Hyperscalers are major new consumers of tape technologies due to the massive data volumes they need to manage and their need to control costs. A hyperscale data center (HSDC) is one that is substantially larger than a typical enterprise data center. With the rise of hyperscaling, the data center is changing. Tape is now a major component of today's hyperscale environments. Cloud hyperscalers are starting to place tape in the back of their cloud environments because of how economically scalable it is. For example, in these hyperscale environments, cold data is typically stored on tape but may need to be moved when it becomes hot again to a higher performance layer, such as front-end flash, for better accessibility and data delivery from the cloud. Now, artificial intelligence (AI), machine learning (ML), and system automation tools are being leveraged to orchestrate optimal data placement in these large environments. The end-user organizations do not need to concern themselves with what media their data is stored on in an HSDC as long as the agreed upon service level is achieved, and data is accessible when they need it for production or recovery.


Many large enterprises and data centers are experiencing the same challenges as HSDCs, including the need to scale their environments while keeping their costs down. Their growth is being driven by new, innovative technologies and business initiatives, including IT automation, IoT, 5G, big data, gaming, and more. The same strategy used by HSDCs for cost controls and tiered storage is used in these environments. Tape is becoming an integral and growing part of the strategy for cost control and data availability. This holds true for both the enterprise data center and the biggest cloud providers.

Hard disk drives have played a major role in capacity scaling and have been joined by flash SSDs, but the magnitude of hyperscale storage requirements and costs are shifting more focus to the economics of scale. HSDCs are taking advantage of tape technologies to manage their data growth, retention, and availability challenges. For the vast amount of data in an HSDC, most of it doesn't need to be immediately accessible, and tape can be the optimal media to store it on. Future advanced tape architectures, such as the new IBM and Fujifilm 580TB tape, will allow HSDCs to achieve business objectives by providing data protection for critical assets, archive capabilities, easy capacity scaling, the lowest TCO, high reliability, fast throughput, and cybersecurity protection via the air gap. These benefits are expected to become of increasing value to HSDCs and other organizations faced with the challenges of increasing growth and volume of data. For any large-scale data center, adding tape is strategic. Data centers are re-engineering storage strategies to manage their extreme data growth and expanding their use of tape technologies to drive costs down and deliver efficiency.

Infrastructure Enhancements

Every few years, a new LTO tape generation is released as part of a very long-term roadmap. Each time, inevitably, hurdles arise as the tape and drive manufacturers attempt to attain better performance, as well as larger capacities. Considering these historical hurdles, it is remarkable that IBM and Fujifilm have already succeeded in demonstrating the ability to surpass the specifications of LTO Gen 12, which is years away from market introduction; we are now at Gen 8 with Gen 9 having just been released.

Figure 3. Enhancing Infrastructure Lifecycle



Source: Enterprise Strategy Group

For comparison, as seen in Figure 3, a single LTO-8 tape stores 12TB of native data (30TB compressed). The current LTO roadmap extends out to LTO-12 with 144TB of native capacity (360TB compressed). LTO-12 is not expected until 2029. The IBM and Fujifilm demonstration shows the viability of meeting LTO roadmap timeline expectations and the potential to push capacity beyond the projected LTO 12 capabilities to as much as 580TB of capacity. This also highlights that, according to the roadmap, hard drives are no longer keeping pace with the capacity growth of tape.

ESG looked at the real-world impact this would have on existing tape libraries, which can be easily updated with new drive technology to accommodate the new tape formats. The second half of Figure 3 shows an IBM 3U 4300 tape library with the ability to internally manage 40 tapes. The comparison shows that, if this tape drive was at its capacity today at Gen 8, the

maximum data stored on the tapes would be 480TB of native data. Now, compare this with the demonstration by IBM and Fujifilm, and the 40 cartridges would hold 23,200TB of native data. This turns existing systems into high performance solutions with a capacity increase of 48 times over the current configuration.

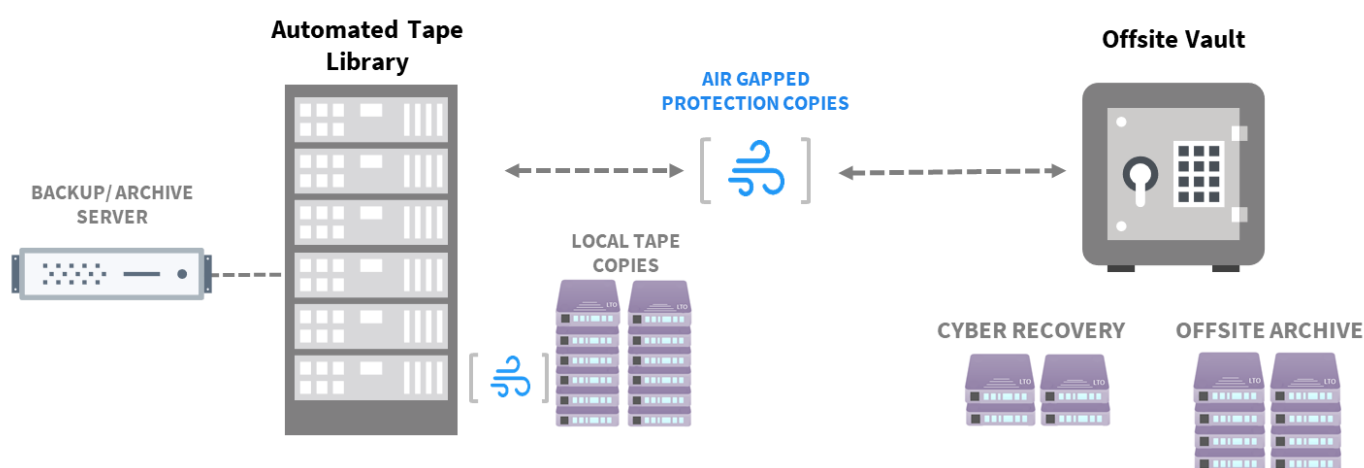
Unlike a typical disk-based system, which requires a substantial upgrade to modify a system for this type of large-scale change, a tape system is non-disruptive and seamless. The only modification is to slide out the drives and replace them with new generations. New LTO tapes can then be used right away.

Another factor to consider is the energy advantage. Data centers are energy consumers. High-density, multi-core data center servers typically use between 500 and 1,200 watts while HDDs use about 6-15 watts per hour, approximately three times more than SSDs. When compared to disk, tape demonstrates sustainability value, especially as tape media capacity grows from 18TB to 580TB while keeping the same manufacturing process and form factor. This ultimately leads to a significant CO2/TB improvement over HDD, which needs to implement energy inefficient processes like heat assist. Reducing the number of servers and moving low-activity data from disk to tape present the greatest data center energy savings opportunities and move organizations toward net-neutral CO2 emissions environments.

Use Cases

There are shifting dynamics in the tape storage world. Today, end-user organizations are challenged by the need to manage large amounts of data in complex environments. Data is growing relentlessly, and now, with the ability to apply advanced analytics to derive competitive advantage, the value of data is increasing, and it must often be retained for longer periods of time. Tape use cases have shifted from small data sets to larger, linear file pulls that are well suited for more strategic tape use cases such as intelligent data lifecycle management where tape is playing an even greater role with unstructured data integration in self-service models. Tape is cost effective with high capacity, but it also has a well-defined and achievable roadmap for advancement well into the future, as demonstrated by IBM and Fujifilm.

Figure 4. Data Protection, Archive, and Resiliency



Source: Enterprise Strategy Group

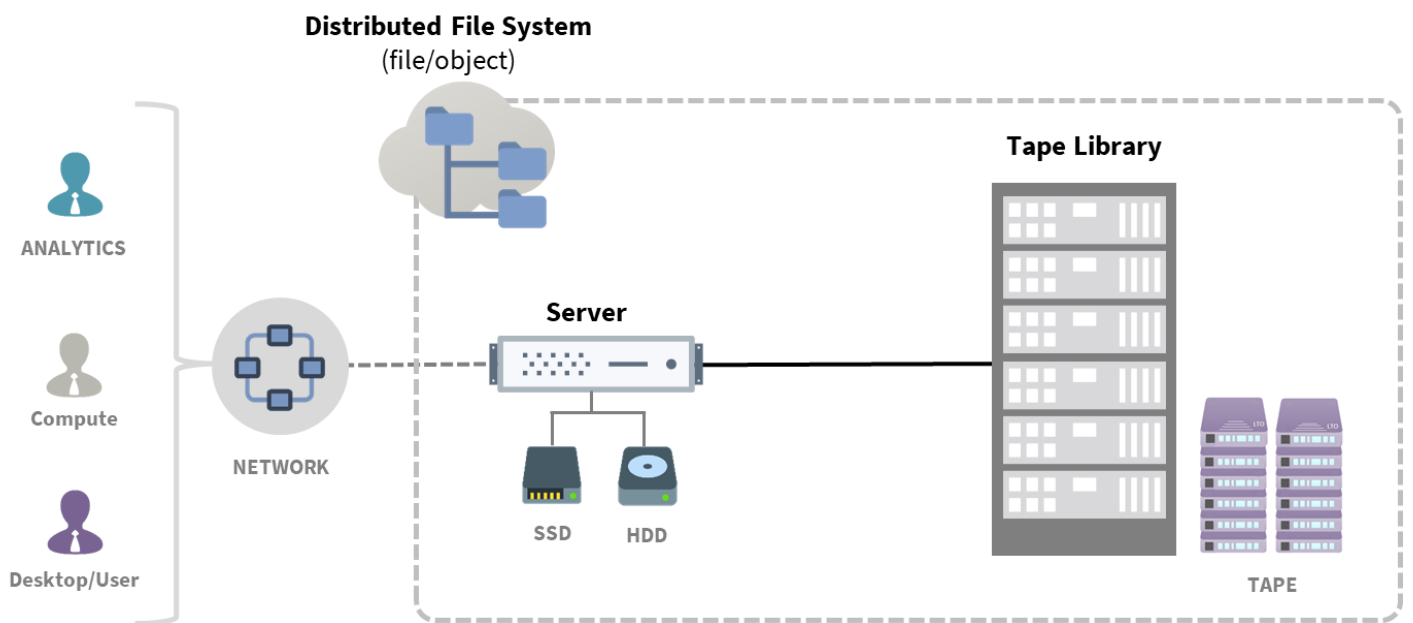
ESG considered the use case of tape for data protection, archiving, and resiliency (see Figure 4). Data is a critical asset of every organization, with backup being the first copy of production data. As data volumes increase, the business case for

backup on tape becomes more compelling because it provides the ability to economically store and recover more data from a single cartridge faster. For data that needs to be kept longer for compliance, regulatory, and corporate governance reasons, tapes are easily moved offsite as data ages. Another key advantage of tape is in cybersecurity. The cybercrime epidemic has become a major problem for many organizations of all sizes. Tape can play a key role in cybercrime prevention and provides WORM (write once read many) and encryption capabilities, providing a secure storage medium for compliance, legal, and protection of valuable files. Tape as an “air gap” solution, shown as a Cyber Recovery Vault in Figure 4, has gained momentum, providing an electronically disconnected copy of data that prevents unauthorized access by cyber-criminals. Disk or flash systems that remain online 7x24 are the primary target, as they are always vulnerable to a cyber-attack.

ESG looked at a second use case of intelligent data management, as shown in Figure 5. In this use case, intelligent data management software automatically moves data from expensive tiers of storage to economy tiers of storage based on user-defined policies. Too many organizations are limited by a paradigm that says they should focus on disk and cloud. They would be wise to leverage the multi-faceted benefits that today’s modern tape technology has to offer.

As more organizations are starting to adopt hyperscale designs across the board, the hyperscalers themselves are now using approaches such as AWS Snowball and Glacier to move data to and from tape in the cloud. As mentioned, ESG has observed a trend in which hyperscalers are using tape in the cloud, front-ended by a much more expensive type of storage, including HDD and SSD. As the trend toward increasing the capacity of tapes grows, the overall TCO decreases for organizations, including the data center footprint, rack space, power consumption, and carbon footprint. In fact, independent studies show that tape consumes 87% less energy than equivalent amounts of disk and reduces CO2 emissions by 87% as well. This same concept of using data at scale with tape can be easily adopted by any size organization to create true economies of scale. There are lessons to be learned from the practices being developed by hyperscalers when it comes to cost reductions, scalability, and even environmental impact.

Figure 5. Intelligent Data Management



Source: Enterprise Strategy Group

The world's data storage needs are increasing exponentially, with HD 4K/8K video, IoT/ICT, hybrid cloud environments, and the proliferation of "big data" analysis. Emerging technologies like smart manufacturing, connected vehicles, and hybrid cloud environments ensure data will continue its growth trajectory. "Cold data," stored long-term and rarely accessed, is estimated to account for as much as 80% of all data stored today and is increasing in its value to businesses. This includes active archiving, backup and recovery, as well as data kept for compliance, regulatory, and corporate governance requirements.

Organizations are investing heavily into IoT and may not realize that 92% of all IoT data is lost or thrown away ephemerally, and they may be missing key indicators that could grow their business. Tape is the only way to retain more of that data to get more value from IoT investments. Intelligent data management also ensures that data is accessible when needed to prevent any business interruptions and to utilize data for active archiving and smart business decisions while also tiering data to lower cost solutions as allowed. The overall goal is to create a balance that meets both operational and financial objectives. Tape can be an integral part of the strategy and easily drive costs down in a small data center environment and exponentially in a large one.

The Bigger Truth

The current limitations on disk highlighted in this paper also indicate that tape capacity is outpacing disk and has the fundamental building blocks to continue well into the future. Some IT professionals may have misconceptions about tape as being more problematic or error-prone than disk. But the truth is that today's highly advanced modern tape is more reliable than disk, especially when organizations leverage robotic automation and reduce physical handling. At this point, with the ability to store so much data on a single cartridge, and the head technology that can easily keep up with the reads/writes, the industry has come to a point where it is possible to store huge amounts of data on smaller, server-sized autoloaders right up to multi-frame automated libraries that can deliver major positive impacts on storage management.

The current tape capacity roadmaps are outpacing disk, and tape is quickly being adopted in many new use cases by hyperscalers and in traditional data centers. Archiving and backup and recovery are two common use cases, but as capacities and performance outpace HDD, the use cases are increasing, and tape is now common in active archiving and a core part of lifecycle management. The demonstration by IBM and Fujifilm further demonstrates the direction of tape and how even more use cases will continue to be supported to change the dynamics of the data center environment in extremely positive ways. While a 580 TB cartridge is a technology demonstration, current generations and those on the immediate horizon are coming to market at the right time and balance of cost, capacity, and performance to meet the many data management challenges. Hyperscalers, concerned with cyber-attacks, are driving down costs and reducing energy waste as they expand on their use of tape as an alternative to disk. With rapidly expanding tape capacities, known reliability, and the ease of air-gapping, tape is a powerful solution that may create a competitive differentiation for organizations that adopt it now. Organizations should also not underestimate the value of tape, especially as an air-gap solution, given the seemingly constant cyber-threat environment organizations need to deal with today.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.