

White Paper

Understanding the Business Value of Modern Tape in Production and Protection Storage

By Mark Peters, Senior Storage Analyst; and Jason Buffington, Senior Data Protection Analyst

September 2015

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

Contents

Introduction	3
Understanding the Business Problem and Operational Drivers.....	4
Economics of Storage Management	4
Economics of Data Protection Storage	4
How Modern Tape Addresses the Business Needs of Modern Storage.....	5
Raw Costs of Tape Systems	5
The TCO of Tape	6
Further Economic Impacts and Advantages of Tape	7
How Modern Tape Is Utilized for Modern Data Protection and Preservation	7
Why Tape Is Still Synonymous with Compliance	8
The Bigger Truth	9

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 from time to time. 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.

Introduction

In today’s economy, the reality is that IT is driven by the demands of the business. In ESG’s [2015 IT Spending Intentions Survey](#), reducing costs and improving regulatory compliance were both among the five most commonly cited business initiatives expected to drive technology spending over the course of 2015 (see Figure 1).¹

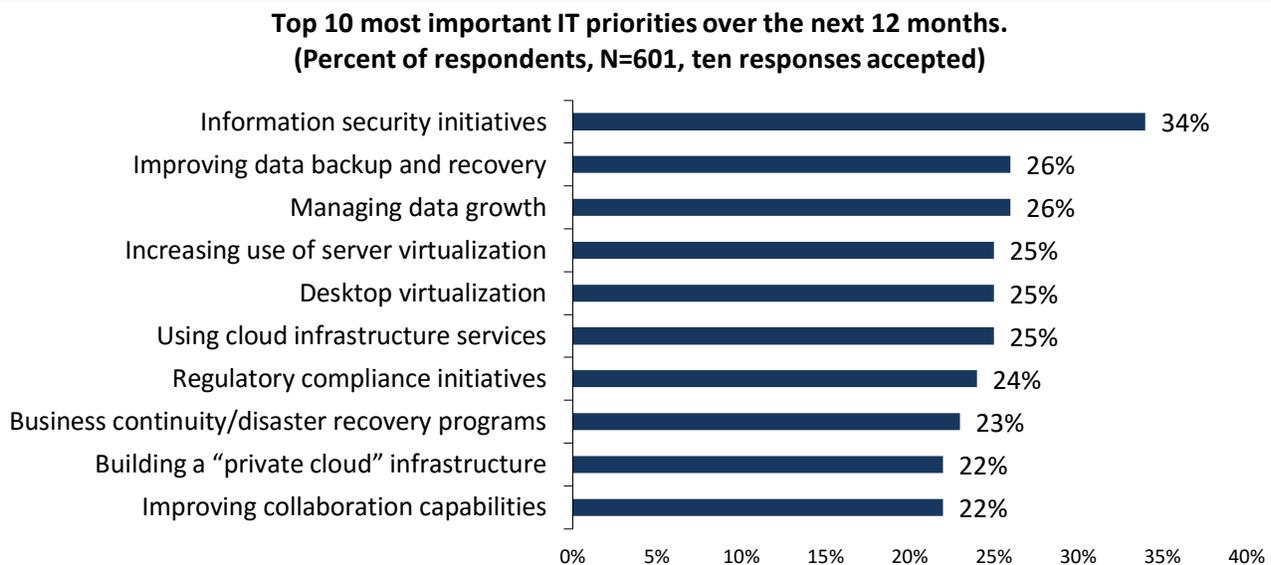
Figure 1. Top Five Business Initiatives Driving Technology Spending in 2015



Source: Enterprise Strategy Group, 2015.

From an IT perspective, organizations’ top priorities often involve storage and data protection (see Figure 2).² It’s a challenge for them, however: Although organizations have always struggled to manage costs, their average production and protection storage *infrastructures* are growing at rates beyond those of their storage and data protection *budgets*.

Figure 2. Top Ten IT Spending Priorities in 2015



Source: Enterprise Strategy Group, 2015.

¹ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.

² *ibid.*

Said another way, while controlling costs may be a primary strategy driver and improving compliance often equates to storing more data for longer periods of time, the exacerbation of the status quo through the ever-growing scale of storage in both production and protection scenarios is unsustainable without a more strategic approach to managing storage.

Understanding the Business Problem and Operational Drivers

To quantify the challenges, it is important to consider the economics of both primary storage and secondary/tertiary protection storage.

Economics of Storage Management

The economics of storage management are often clouded by discussions of “necessary performance” or “prerequisite speed” and “crucial capacity.” In reality, the only factor that actually drives the link between business problems, operational drivers, and storage choices is economics. The economic tail wags the storage dog, not the other way around. The easy proof of this potentially inflammatory statement is to consider a world in which all storage types cost exactly the same—clearly there would not be any HDDs or tape used in that scenario because there would be no reason to do so when solid-state costs the same. The varying prices, TCO models, and ROI impacts of different applications (... and users, and businesses ...) are what drive storage system decisions. Period. The workload and business parameters will—or at least *should* (if they are pragmatically optimizing their environment)—drive users to pragmatically pick storage that is at an overall price that is in line with the application, user, and/or business need. For some users and uses, that can only mean server-based flash. For others, it can perfectly easily (and sensibly) mean a sophisticated, integrated modern tape system.

Economics of Data Protection Storage

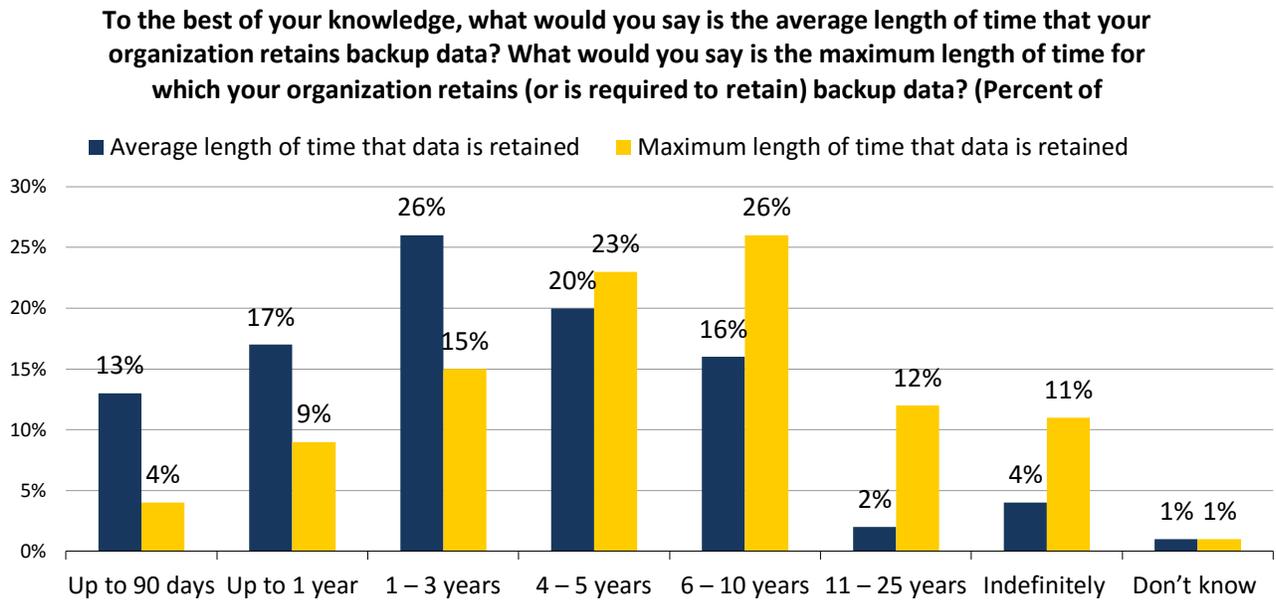
ESG research uncovered that regulatory compliance can be seen as both a business initiative (see Figure 1) and an IT priority (see Figure 2). And business requirements for regulatory compliance often equate to retaining more data for longer. In fact, under the growing pressure to be more agile, more copies of data are being created as well as stored for longer periods of time. According to ESG’s most recent data protection research, organizations report creating—on average—approximately 11 copies of data for various protection and non-protection purposes.³ And although deduplication can help, many of those copies are spread out over retention windows that prohibit a disk-only approach to data retention (see Figure 3).⁴

On average, organizations report creating roughly 11 copies of data for various protection and non-protection purposes.

³ Source: ESG Research Report, *Trends in Data Protection Modernization*, to be published.

⁴ Ibid.

Figure 3. Average and Maximum Data Retention Time Periods



Source: Enterprise Strategy Group, 2015.

How Modern Tape Addresses the Business Needs of Modern Storage

The preceding sections have established that the economic value proposition of tape is well suited to addressing one of the foremost challenges of modern IT: balancing increasing data scale with increasing pressure to constrain costs. Furthermore, the business value proposition of tape has particular suitability for long-term, low-activity, or dormant data and data protection uses. But of course, every storage vendor will claim “good economic value” for its products—whether those are the highest level of flash systems or a plain old filing cabinet. How and why is tape economically attractive?

Raw Costs of Tape Systems

The essential value of tape is that it has always provided the lowest cost per gigabyte/terabyte/petabyte of storage. And that matters: As previously mentioned, cost reduction was among the top business initiatives expected to affect 2015 technology spending, which has been a consistent trend over the last five years.

Tape, therefore, clearly matters in the overall storage equation. Indeed, in our increasingly converged, virtualized, and software-defined world (the sum of which is more fluid data and easier, more automated data movement), it is easy to actually see the role of tape expanding. With data volumes and retention needs outstripping budgets, and with the increasing abilities of software to intelligently and dynamically manage data placement, tape will retain a role wherever high volumes of less active data need to be dealt with economically...whether on-premises or in the cloud.

While entry costs are certainly associated with a tape environment (drives and invariably automated tape libraries), there are entry costs (controllers and invariably software) associated with disk systems, too. The exact starting cost (and, therefore, the precise per-X-byte cost) is always going to vary and be hard to determine across users, geographies, and specific user situations, but asserting that tape will almost always be cheaper is safe. And after the entry costs are accounted for, then the expansion costs are clearly in tape’s favor. As cartridge are added, the raw costs can easily be under one cent per GB. For comparison, cloud storage—even at around or under one cent per GB—is priced *per month*, so the difference is clear even before allowing for additional network and other costs that tend to come with cloud services. Indeed, it is easy to see why cloud providers love tape!

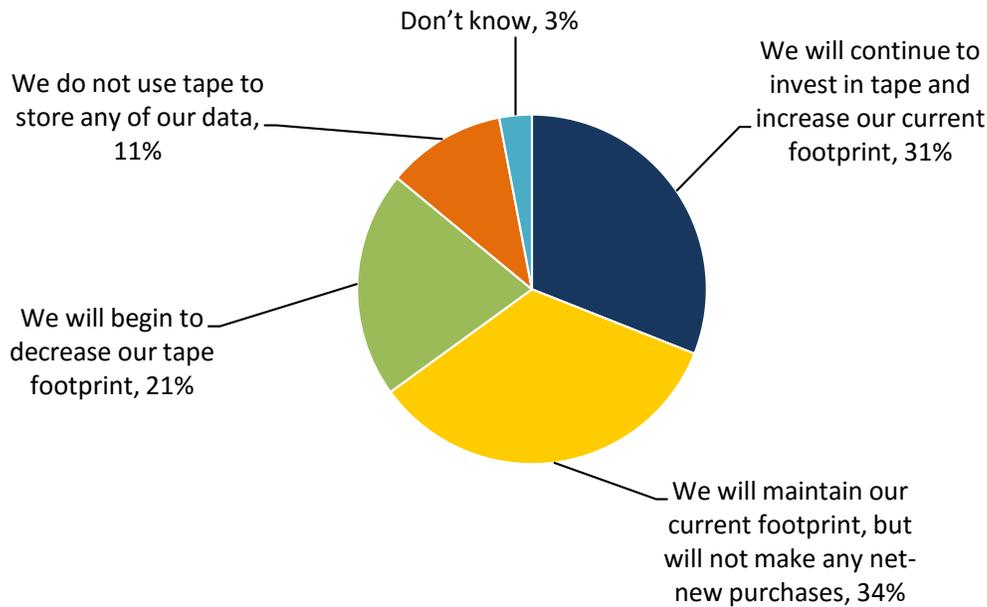
Overall then, tape storage costs per GB/TB/PB will actually decrease as the installed capacity increases and the hardware is amortized over a greater capacity; the same is not true (certainly not to the same extent) with disk-

based installations. Today’s huge tape capacities (10TB uncompressed on the latest IBM TS1150) keep the CapEx low, while the extended longevity of the media (30 years is the standard promise), together with the cross-generation compatibility of tape drives and media, mean that the investment does not need to be repeated as often as with disks.

All told, ESG’s latest data storage-focused research suggests that most organizations understand the value proposition of tape, with roughly two-thirds reporting that they plan to either maintain (34%) or grow (31%) their tape footprint (see Figure 4).⁵

Figure 4. Planned Organizational Tape Strategies

**What is your organization’s forward looking strategy with regards to tape technology?
(Percent of respondents, N=373)**



Source: Enterprise Strategy Group, 2015.

The TCO of Tape

From a strictly CapEx measurement standpoint, tape-based storage systems cost only a fraction of what comparable hard disk storage systems cost, but it is in OpEx terms that tape truly shines.⁶ Because tape cartridges in a library slot or on a shelf do not consume energy when they are not in use, tape uses the least energy (and is therefore the greenest) of all storage technologies and can be used to increase storage capacity with little, if any, increase in energy consumption.

But the savings in TCO go beyond just electricity; the longevity of the media, together with an ability to reuse it, contribute to constraining costs. Furthermore, the longevity leads to easier management and lower installation costs. Additionally, tape systems do not need to be on the same three-year refresh cycle that is found so often across much of IT. The strong raw-reliability of contemporary tape systems such as IBM’s may even translate to a reduction in the number of support-desk calls to deal with and pay for.

⁵ Source: ESG Research Study, *2015 General Storage Trends Survey*, conducted in May 2015.

⁶ Source: ESG’s TCO study (*A Comparative TCO Study: VTLs and Physical Tape*, February 2011) compared an LTO-5 tape library system with a low-cost SATA disk system for backup using deduplication (best case for disk) and showed that disk deduplication has a 2-4X higher TCO than the tape system for backup in several use cases over a five-year period. For archive data, the TCO for disk is considerably higher. ESG is also currently conducting new in-depth “Economic Value Validation” based research in this area.

What might the TCO advantage of tape actually be? While there are myriad “your mileage may vary” issues to consider, IBM claims that the overall TCO of tape is some 6X better than disk. Even taking a decent size “grain of salt” with that number, it can be seen that the TCO advantage for tape over disk is considerable.⁷ Indeed, one of the most significant contributors to that TCO advantage is—no surprise here—electrical power consumption; it is not hard to construct situations where tape has a two orders of magnitude advantage over disk in this regard, which is clearly of especially high value to users that operate in either high-cost-per-kW areas and/or where power availability is actually constrained.

The attractive TCO of tape has been helping it to expand from its historical role—primarily as a backup solution—to a much broader set of workloads, including archive, fixed content, unstructured data, disaster recovery, business continuity, and long-term data retention. A burgeoning global list of government, compliance, and legal regulations now describes not only the way data should be managed and protected, but also the length of time that data is required to be stored. For many applications and files, the lifetime for data preservation has become essentially infinite, further increasing the size, security needs, and management requirements of the digital archive. Meeting such huge and invariably less-active storage requirements with disk would be financially prohibitive for most businesses, while the technological and TCO progress in tape media, drives, and libraries has once again positioned tape as the credible and optimal high-capacity storage solution for the foreseeable future.

Further Economic Impacts and Advantages of Tape

Because economic considerations are a part of all IT decisions, overall economic value matters. As already mentioned, ESG’s latest spending intentions research shows that IT priorities and initiatives driving technology spending are led by a triumvirate of needs for data security, improved backup and recovery, and cost constraint.⁸ Tape can actually address all three. The demands for this threesome are invariably non-negotiable, but meeting them in aggregate is simply not attainable by any other storage media. No one is suggesting that tape should become the basis of all storage provisioning—database processing, for instance, or VDI are going to stay on disks and flash—but it can be an economically attractive and sensible underpinning to everything.

Tape brings more economic value to the IT table. For instance, its fast streaming can mean a fast recovery (getting applications or a whole business up and running again), while its sheer mobility is an efficient and inexpensive way to cloud-seed or cloud-burst. In another very practical implication, the great storage density that tape can support (especially when combining IBM’s TS1150 drives and media in its expansive TS4500 library⁹) can lead to lower data center space costs. And for a final compelling insight into the value of tape, look no further than the many cloud providers that are embracing it for its economic value—not just for the competitive per-TB-stored pricing it enables them to offer (or their achievement of better margins), but also because tape (via such functions as encryption and its “island” status) can avert potentially nasty economic results—precluding, or at least ameliorating, exposure from system glitches. Putting it colloquially, tape can save your bacon ... and that will save money for the business as well as for IT.

How Modern Tape Is Utilized for Modern Data Protection and Preservation

Data protection and preservation are synonymous with tape for good reasons; and modern data protection still requires modern tape to this day. As seen in Figure 5, tape is utilized as part of the primary backup processes for nearly half (49%) of all surveyed IT organizations, a number that is higher among enterprise organizations.

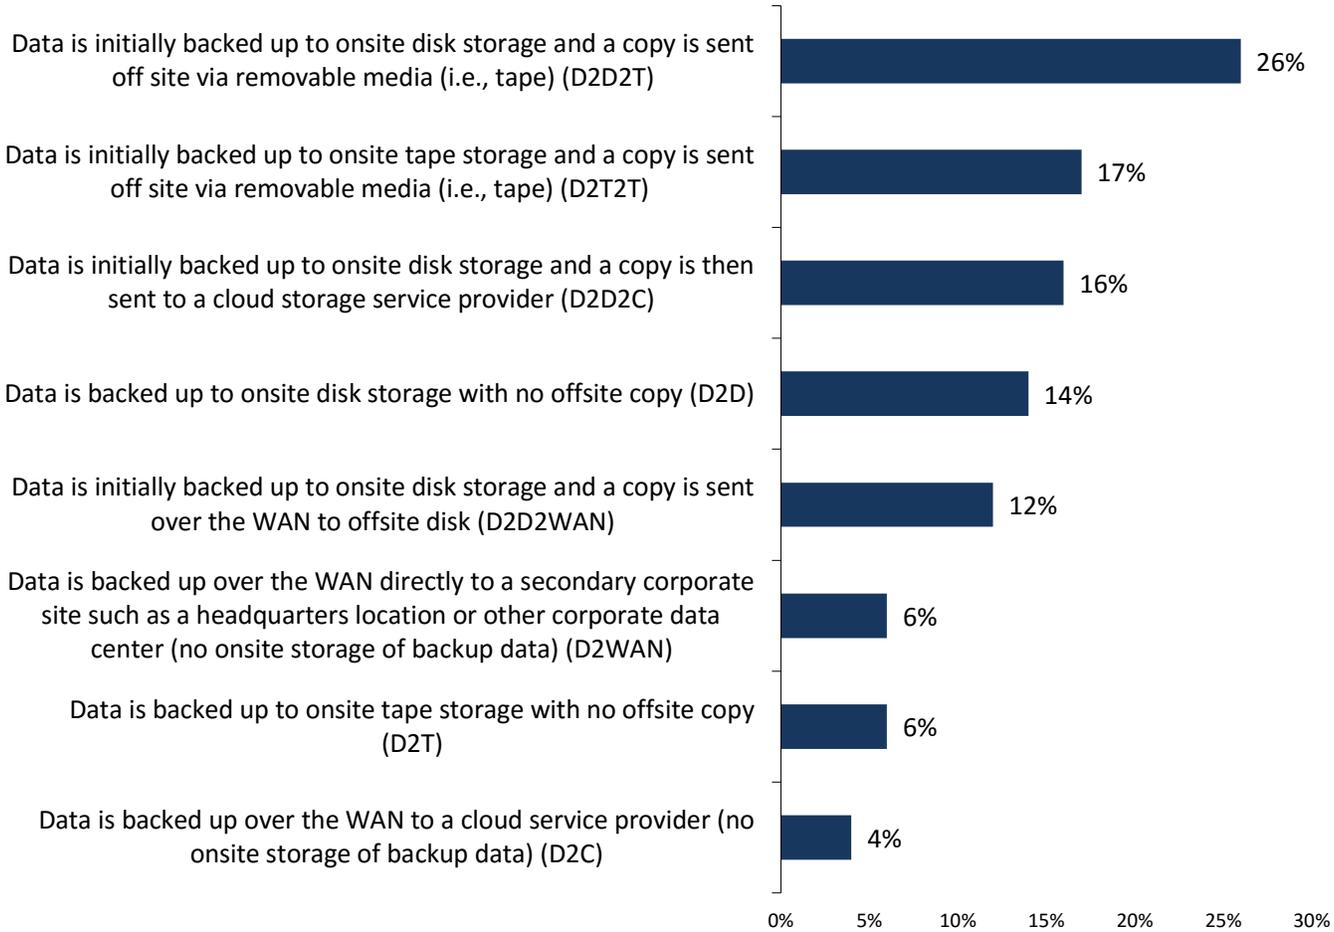
⁷ There is no particular need to doubt IBM’s numbers, since it is a systems vendor that is just as motivated to sell disk as tape. IBM states that the power consumption per TB for tape is over 100X less than that for disk, while its Spectrum Archive can reduce the storage TCO by 90% compared with disk.

⁸ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.

⁹ A single library system can contain 180PB of capacity, and multiples thereof as additional systems are linked.

Figure 5. Manner in Which the Data Backup Process Is Generally Managed

Thinking about your organization’s environment today, which of the following **best describes** how the data backup process is generally managed? (Percent of respondents, N=375)



Source: Enterprise Strategy Group, 2015.

In fact, according to Figure 5, tape is the exclusive media for nearly one in four (23%) organizations’ primary backup topology. And while there may be many reasons, the earlier ESG findings on long-term retention having to do with compliance are among the most compelling reasons to include tape within your strategy, and are also why you explicitly cannot replace tape with disk- or cloud-based storage.

Why Tape Is Still Synonymous with Compliance

One of the key facets of many compliance mandates is the presumption of the defensible authenticity of the restoration media. Said another way, no matter how long ago the data was originally stored, auditors, litigators, and other restore-invokers must be confident that the data has been stored in a pristine state and has been unmodified since the original storage date.

One of the key facets of many compliance mandates is the presumption of the defensible authenticity of the restoration media.

The challenge with disk-based systems is that they are limited in their assurance of pristineness by how the backup/archive/data-management software wrote it. More importantly, running the same protection storage (disk) array for ten or more years is seen as impractical by many. Instead, cloud storage is often considered as an alternative complement

to disk, in deference to tape. But most cloud solutions falter in ensuring the sanctity of the data while still providing agile manageability. Consider the following scenario using tape:

- For the first few years after writing data to tape, a courier service would take the tapes to a warehouse.
- For the next few years, instead of using a courier service, the organization simply stored the tapes at an alternate corporate location.
- In more recent years, the organization has chosen to use a full-service tape management service that can also index the data, maintain working drives, etc.

At any point along the way, the data on the tape is pristine and defensibly authentic from the original date that it was stored on that particular cartridge.

Now consider a similar dynamic of venue changes using cloud storage:

- For the first few years, the data was written to a very popular cloud storage company.
- That cloud provider's costs kept rising higher than competitors, so the organization wanted to switch.
- Later, the organization decided to move the data again for geo-political boundary reasons.

In a cloud model, the data likely retains all of its versions courtesy of the backup software, as long as the original cloud provider is maintained. However, most mechanisms have no ability to pull the entire data set out of a cloud provider so that the data is portable and pristine after three to four years. Instead, some organizations would have to leave their oldest data in the original cloud, even though newer data goes to the less-expensive cloud, in order to maintain its pristineness.

In these scenarios, the most likely way for a cloud provider to ensure pristine and portable data for truly long-term retention is to use—you guessed it—tape.

The Bigger Truth

The most important realization about “modern tape” is that it is in fact “modern.”

From a purely cost-per-TB storage consumption standpoint, tape makes sense—and dollars. But it can only do that because the performance and reliability of modern tape formats have greatly surpassed the poor reputations of the tapes of previous decades. With the continually growing capacities of modern tape cartridges, the actual number of media cartridges can decline as the storage efficiency per cartridge increases, thereby further reducing operational management costs for the cartridges themselves. Of course, high-capacity libraries and auto-loaders further reduce the manual handling costs of tape, so performance, portability, and pristineness become the real decision makers.

While disk systems don't have longevity, cloud services are not yet able to ensure the authenticity of the original data without forcing subscribers into a non-negotiable vendor lock-in to preserve the defensibility of the restoration of data over the entire length of retention that is mandated, regardless of the duration. It is for these reasons, among others, that so many organizations continue to embrace tape as a strategic aspect of their storage management and data preservation strategy.



Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com