

White Paper

Tape and Cloud: Solving Storage Problems in the Zettabyte Era of Data

Sponsored by: LTO Consortium

Phil Goodwin
June 2019

INTRODUCTION

IDC research estimates that corporate data will continue to grow at a 40-50% compound annual growth rate (CAGR), doubling every two to three years. In this regard, IT organizations can expect to deal with approximately 7.5ZB of data by 2025, up from 2.6ZB in 2018. At the same time, the majority of data growth (80-90%) is associated with unstructured data. While unstructured data is nothing new, there are new sources that are now generating it. Unstructured data is increasingly generated from social media, search engine queries, real-time streaming, and Internet of Things (IoT) sensors. Data management is also becoming more complex with additional regulatory requirements, long-term retention needs, and "right to be forgotten" laws in the United States and the European Union. As unstructured data becomes more valuable, it is consistently a target of theft and malware attacks, such as ransomware and distributed denial of service (DDoS). It is imperative to keep an archival copy of data offline, creating an "air gap" to protect it against the threat of cyberattacks.

The combination of data growth and more stringent data management requirements strains the infrastructure and staff of IT organizations. While IDC forecasts that IT budgets will increase by 2-3% in 2019, IT staff resources will remain flat. The result will be that as data continues to grow exponentially, it will require businesses to find more efficient methods of managing it. Many organizations fall into one of two categories: those that are trying to maximize prior investment in existing technology that is overdue for a refresh and those that seek "silver bullets" of new technology that may meet specific point solution requirements but cannot address the breadth of data management issues.

As part of digital transformation, enterprise storage strategies have moved from risk management to saving unstructured data in a cost-effective manner so it may be analyzed and repurposed for competitive advantage, essentially providing access to the right data at the right time. Data-driven organizations strive to provide more accurate information to decision makers faster than their competition in order to gain an edge in the marketplace. Simply put, the availability of data and ability to access data are cornerstones of business success.

Cloud has emerged as a key technology as many organizations attempt to become "cloud first," meaning that all new initiatives center around cloud deployments, most often as a combination of on-premise (private cloud) and public cloud, known as hybrid cloud. Cloud is a technology, and like most technologies, it has some significant benefits but also limitations. Therefore, it is not a silver bullet to address all data growth and management problems. Some organizations have attempted to replace tape with cloud, but rather than being substitutable technologies, they really are complementary.

Cloud can offer significant benefits around application and IT agility, while tape can solve many problems around data management, governance, and survivability.

IDC recommends the following best practices:

- Evaluate technologies based on the strengths of the technology and match that technology with the desired business outcome.
- Consider cloud where on-demand resources are required.
- Consider tape for helping tame large-scale data storage and retention requirements.

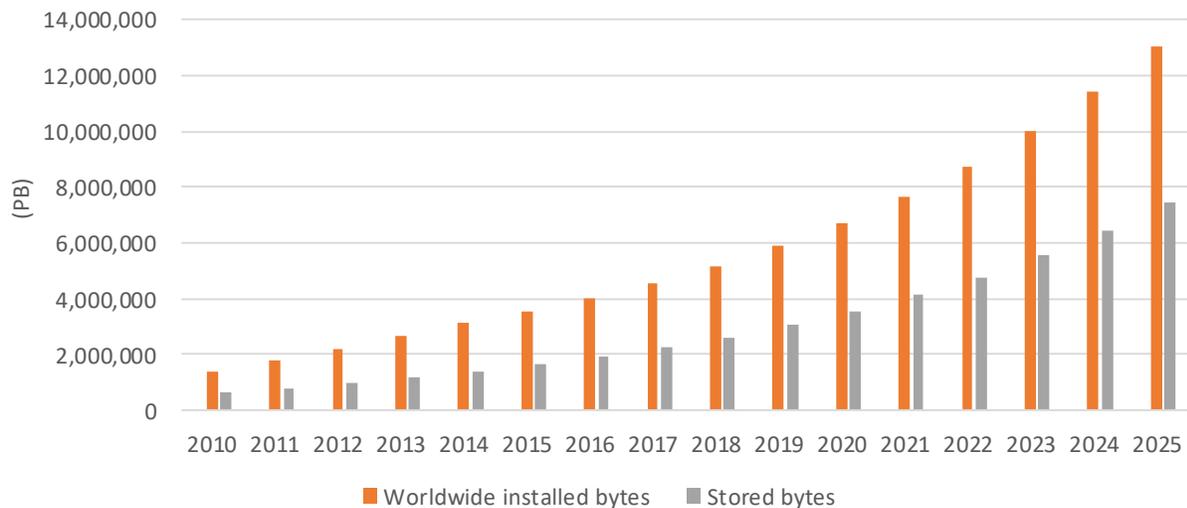
SITUATION ANALYSIS

IDC estimates that as many applications will be deployed in the next five years as have been deployed in total in the previous 40 years. While application deployments include on-premise and cloud configurations, a hybrid cloud and multicloud model will be dominant. Because of the diverse application deployment models, this may result in inconsistent data management, governance, and retention. Cloud-based SaaS apps will rarely have service-level policies in line with corporate requirements. While many organizations like the benefits of cloud-based apps, they also want to copy data on-premise for data assurance, governance, and long-term retention.

IDC's study of the global datasphere illustrates the enormous growth of data that will be impacting IT organizations (see Figure 1).

FIGURE 1

Global Datasphere Forecast



Source: IDC, 2019

By 2025, the world will be generating approximately 7.5ZB of data annually that must be stored and managed by IT. Stored bytes are cumulative, and these numbers are most relevant to IT organizations because they will need to care for and manage the data. In this regard, we expect IT organizations to deal with the increased volume of data in 2025, up from 2.6ZB in 2018.

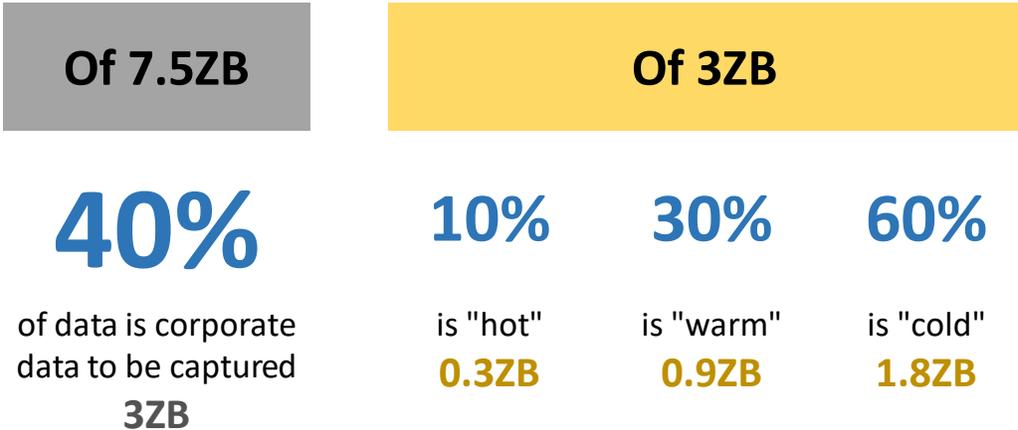
To dissect these numbers further, we can make some assumptions about the data. We can assume that 40% of stored data is commercially related. Of that 40%:

- 10% is "hot," meaning the expectation of retrieval/access is within 10 days.
- 30% is "warm," meaning the expectation of retrieval/access is within 30 days.
- 60% is "cold," meaning that the expectation of retrieval/access is more than 30 days.

Figure 2 graphically illustrates these assumptions.

FIGURE 2

Hot, Warm, and Cold Data Assumptions



Source: IDC, 2019

Based on these assumptions, IT organizations must be prepared to manage 1.8ZB (1.8 trillion gigabytes) of cold data, or medium-term to long-term data storage. Actual impact may vary by industry or organizational profile, so readers are invited to apply their own assumptions against the actual amount of data in their organization.

It is the 60% of "cold" storage that can be addressed by either tape or cloud. Although there is a trend toward all-flash arrays in the datacenter, it is simply not necessary to house infrequently accessed, but still essential, data on the most expensive, low-latency disk. However, it is essential to keep this data safe for when it is needed. Although all data will not always be accessed, in most cases, the users cannot know which data they will need, so they must keep it all. This cold storage could be a candidate for either cloud storage or on-premise tape, but recovery time objectives/recovery point objectives (RTOs/RPOs) along with mounting monthly charges and other hidden fees must be taken into consideration when evaluating storage solutions.

As corporate data stores expand, increasingly vast quantities must be protected considering that any data directly accessible via the web is vulnerable to attack or theft. Ransomware attackers succeed only when they can destroy all means of self-recovery. Once an organization is hijacked of its ability to recover its own data, then cybercriminals are able to extort maximum ransom. Though being unable to recover data is a cardinal sin of data protection, our research shows that nearly 25% of organizations have suffered such an event within the past three years.

Data replication technologies, such as snapshots, mirrors/clones, and remote replication, can protect against system failures and deliver low RPOs/RTOs. However, that data is still vulnerable to corruption and malware through the unwitting replication of it, leading to unrecoverable data loss from malicious attacks, especially internal attacks by disgruntled employees. IT organizations often chase a fully automated data protection scheme through replication technologies, but in doing so, they must also ensure data corruption or ransomware is prevented by creating an air gap – a break in the data stream. This can be accomplished manually by moving data from one storage medium to another or by physically separating copies of data (i.e., data sets) to ensure one does not infect previous backups. While possible through replication technologies, this manual approach can be arduous and expensive. Tape technology, in contrast, natively introduces an air gap and does so cost effectively, ensuring there is always a recoverable master copy of data.

The current model for nearline storage is a combination of solid state disk (flash) and low-cost, low-performance disk. Some organizations use low-cost disk for archival purposes, and while it is less expensive than SSD and performance HDD technologies, it is inherently more expensive than tape and has a higher total cost of ownership associated with power and cooling.

While cloud storage is one solution for long-term retention of infrequently accessed data, it does not inherently include an air gap as part of the data protection continuum. Moreover, attempting to recover large quantities of data from the cloud may be unacceptably long and/or very costly. Beyond the fixed monthly storage cost, cloud storage vendors know enterprises periodically need to access archived data for analysis and profit from that by charging fees to export data on a per-gigabyte basis. In fact, charges associated with periodic data retrieval can dwarf ongoing primary storage costs.

Tape archival systems, in contrast, do not require power or cooling other than the hardware to read the tape cartridges. Even considering the capital costs for hardware and software, floor space, energy consumption, and administration, LTO tape technology still offers a far lower total cost of ownership than other archival methods.

FUTURE OUTLOOK

Considering Tape

Magnetic tape has long been a keystone technology of the datacenter. Tape offers many benefits, including impressively high restore rates based on any specific configurations for very large volumes of data. LTO technology has become the de facto tape format standard worldwide, and tape systems can integrate with multiple systems onsite or offsite, including many on-premise/private cloud solutions.

Current LTO technology eighth-generation standards include up to 30TB of data per cartridge (compressed) and up to 750MBps data transfer throughput per tape drive (compressed). By comparison, Gigabit Ethernet (GbE) offers 125MBps transfer rate. Data transfer rates are crucial when comparing RPOs/RTOs based on either cloud or on-premise tape assets. Because tape drives can

restore data concurrently, total data restore is limited only by the configuration of drives. While replication technologies do offer better RPOs/RTOs in most cases, they do not inherently protect data from malware or corruption.

In addition, Linear Tape File System (LTFS), an important part of the LTO technology specification, makes tape easy to manage through a native OS browser and offers partitioning and indexing of content. One partition holds the content, and the other holds that content's index, making the tape self-describing. Further, tape archive systems have the ability to integrate with multiple solutions onsite or offsite, including many on-premise/private cloud solutions. LTFS is cross-platform compatible, allowing data to be shared across Linux, Mac, and Windows, and can also be combined with flash and tape storage (referred to as *flape*) to act as a management layer to handle workflows in a tiered storage architecture with competitive low costs. Emerging workloads, such as big data and analytics, that require fast transactional processing of data that is expected to remain unaltered will be especially good candidates for flape solutions.

As previously mentioned, cloud services do provide offsite archive storage, but they suffer from long recovery times for large volumes of data. Tape, therefore, should be considered a complementary technology to cloud storage while offering a lower TCO that SSD/HDD cannot match. Creating an optimal SLA without breaking the bank requires a multitier data protection strategy to address both data loss/recovery and cyberthreats.

Moving tape offsite can ensure data survivability, because LTO's tape life is guaranteed for 30 years, and it requires little to no management while providing an inherent air gap that prevents malware infection. While some have voiced concern over the possibility of tape cartridges being stolen, this is far from a preferred method for data thieves who would not only need a proper hardware setup but also face LTO technology's native ability to use AES 256 encryption.

While cloud can also be used for long-term data retention through archive as a service that places data on low-cost storage, the GB/month charges add up over time. Tape, on the other hand, can be stored at little more than the environmental cost of the storage room, which makes a significant difference when data retention is required for years. For example, coldline storage on at least one major cloud vendor costs \$0.007/GB/month, \$7,000/PB/month, or \$84,000/PB/year. While cloud services can offer relatively fast retrieval of smaller amounts of data, large data restores across networks can be very slow and, as mentioned previously, cloud egress costs add up quickly and can be significant. In addition, data stored online is potentially accessible to skilled cybercriminals. While data encryption can be used to further protect data in the cloud, those keys may be open to internal threats. Securing data offline until it is needed on vaulted tape is the only way of thwarting unauthorized disclosure at the lowest cost.

IT organizations will find that cloud and tape are very much complementary technologies. The top use cases for each technology is as follows:

- Consider cloud for:
 - Applications where assured time to first byte is important
 - Applications where online data analysis and searching may be required
 - Applications where the data restore/access volume is relatively small (i.e., individual files, tables, or smaller databases)

- Consider tape when:
 - A data restore is likely to involve very large volumes of data
 - Data retention is more than three years
 - Low cost is an important criterion

Even for a small organization, moving data rapidly into an air gap or archive infrastructure on a cloud service requires a significant investment in dedicated internet connections. In many cases, this cost outweighs the benefit of cloud as a potential lower-cost storage option.

CONCLUSION

Far from being "either/or" technologies, cloud and tape are very complementary. In fact, their strengths and weaknesses match up very nicely. Tape is the ideal supplement for cloud storage with a low-cost tape archive for added data protection; it is no surprise then that cloud providers have embraced tape to complement their services.

IDC believes there is a role for both technologies in most organizations. That role will be tied to the application requirements. Organizations that reflexively dismiss tape are missing an opportunity to optimize the data management capabilities of their organizations, especially on a cost-effectiveness basis.

Tape is expected to remain significantly less expensive than hard disk storage for years to come. Tape is ideal for production IT environments because it is able to reduce costs through new software-defined storage and flash-based architectures that create a tiered data storage architecture suitable for both archive and emerging workloads requiring fast transactional processing. Tape has an overall lower TCO and is essentially a fixed cost that can reside in the operating expense section of budgets, making accounting relatively simple and reducing the impact on the overall IT capital budget. Tape requires far less management and can be vaulted for decades offsite, ensuring validated data is available through an external disaster recovery strategy. At the same time, by using tape onsite, organizations can protect data and thwart cybercrime with the air gap, which is created by using a storage medium that is offline until data is needed.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2019 IDC. Reproduction without written permission is completely forbidden.

