White Paper

Using Tape to Optimize Data Protection Costs and Mitigate the Risk of Ransomware for Data-Centric Organizations

Sponsored by: The LTO Program

Phil Goodwin
April 2018

IDC OPINION

IDC forecasts that 60% of organizations will have developed a digital transformation strategy and will be in the process of implementing it by 2020. Leading organizations are learning to become data centric — a key element of digital transformation. Data-centric organizations are those that seek competitive advantage through real-time, data-driven decision making. Data analytics is a key part of being data centric, but it requires data to be constantly available and in the right place at the right time. As data becomes more and more valuable to organizations, it becomes an increasingly attractive target for theft or malware attacks such as ransomware and distributed denial of service (DDoS). These threats are becoming more frequent, challenging IT organizations to protect themselves from threats, as well as traditional data loss threats including failure of hardware and software, human error, and natural disasters.

The dynamic nature of evolving threats requires IT organizations to deploy dynamic data protection capabilities, technologies, and strategies. IDC believes that best practice organizations should maintain and update a threat analysis as well as match data protection technologies and processes to the specific threats. The result is a matrix of data protection used to create an end-to-end continuum of protection against all foreseeable events. Thus data protection should be regarded as a strategic pillar of an organization's data-driven advantage.

Another advantage of matching specific threats to technologies is defeating threats at the lowest possible cost. Broad-brush approaches to data protection can be successful but may come at a significant cost. Tape is recognized as the most cost-effective means of storing data for long-term retention. Recently, tape has shown itself to be an effective means of providing an "air gap" between live data and protected data. 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.

Other areas where tape continues to prove itself as the strongest, most cost-effective, technology include archive-intensive applications and data streaming applications in industries such as banking, medical, IoT, oil and gas, and media and entertainment. Importantly, as organizations look to also add cloud to their data protection solutions, many cloud providers themselves use tape to deliver low-cost, long-term data archive-as-a-service solutions.

IDC recommends the following as data protection best practices:

- 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.

SITUATION OVERVIEW

It is axiomatic that 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 160ZB by 2025, with a 30% compound annual growth rate (CAGR). Figure 1 illustrates this growth in data.

FIGURE 1

Global Datasphere Forecast, 2010–2025

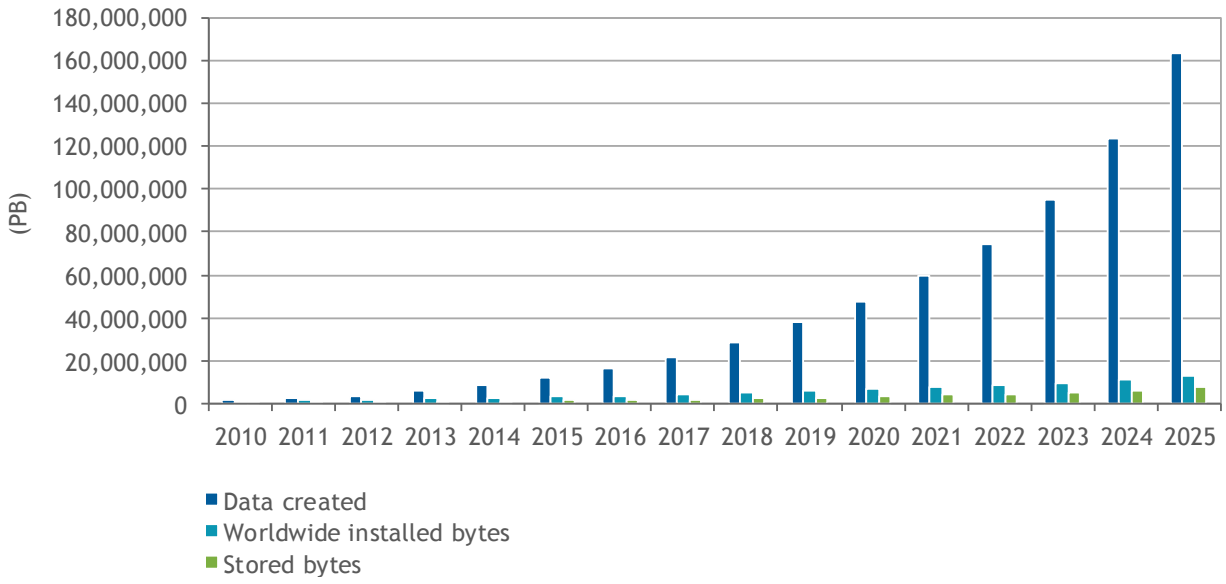

Source: IDC Data Age 2025 Study

Readers should note that the estimate of 160ZB is the cumulative total amount of data generated, both transient and captured data. Figure 1 also shows about 13ZB of installed storage capacity and 7ZB of actual stored data (both cumulative capacities). Worldwide installed bytes refer to the total data storage capacity deployed, including traditional disk drives, cell phones, and other devices that include storage. Stored bytes refer to the amount of storage retained on all devices, but it is the most critical data to be captured and stored. We estimate that approximately 50% of stored bytes is business data. Of that, 60% is neither frequently accessed nor does it require rapid access, what can be referred to as "cold." Therefore, we believe that approximately 2.25ZB of data is eligible for the lowest cost media available.

This proliferation of data causes several challenges for IT organizations.

As the datasphere expands, vast quantities of data must be protected and any data directly accessible is vulnerable to attack or theft. If ransomware attackers can destroy all means of self-recovery, then they are in a position to extort maximum ransom.

These same quantities of data must be recoverable, regardless of loss scenario. Unrecoverable data is a cardinal sin of data protection, yet our research shows that nearly 25% of organizations have suffered unrecoverable data loss within the past three years.

Data replication technologies, such as snapshots, mirrors/clones, remote replication, and object storage, are designed to protect against system or site-related failures and to deliver low RPO/RTO. However, data corruption and malware can be unwittingly replicated, leading to unrecoverable data loss from malicious attacks, especially internal attacks such as those perpetrated by disgruntled employees.

Cloud is not always the answer as its use does not inherently introduce an air gap into the data protection stream. Moreover, attempting to recover large quantities of data from the cloud may be unacceptably long and/or very costly.

One key driver for long-term data retention, and therefore archiving, is data governance requirements. This is especially true in heavily regulated industries, such as financial services and healthcare, but also can apply to almost every organization. In some cases, data must be retained for years or even decades. In other cases, the rules are sufficiently vague that organizations feel compelled to retain data for long periods rather than risk regulatory penalties or legal jeopardy. And many organizations may not be concerned with regulatory compliance but instead retain data long term for business reasons. In industries such as oil and gas, media and entertainment, and scientific research, the data is of critical importance to the business and has value for many years, if not forever. However, this long-term retention can create a jeopardy of its own. Data that is subject to HIPAA, PII, GDPR, or other privacy regulations, or is of proprietary value, must be stored in such a manner that it cannot be accessed by anyone who is not authorized to do so. Unfortunately, the reality is that any data stored online is potentially accessible to someone skilled enough and determined enough to access it. Data encryption is a critical method to protecting data, but encryption keys may be available to internal threats. The only true way to absolutely protect data is to keep it securely offline until it is needed. Securely stored or vaulted tape is the lowest cost method to do this.

The challenge for IT organizations is never allowing data to get into an unrecoverable state while ensuring that data is protected from unauthorized disclosure. This requires a coordinated solution of multiple data protection techniques for specific threats and service levels. For example, snapshots provide very low RPO/RTO for some kinds of data loss, but not system failure, site disasters, or ransomware. Data replication, including remote replication, can provide data survivability for site or system outages but cannot necessarily handle internal threats or ransomware.

The trap for IT organizations is the mistaken quest to have fully automated data protection through automated replication. While much of the data protection process can be fully automated, an air gap must be included to protect from propagated data corruption and ransomware. This air gap is accomplished by a deliberate halt in the data stream, which may involve a manual step of permitting data to be moved from one media to another. In addition, physically separated copies of data (i.e., data sets) ensure that malware does not corrupt prior backup versions. This is possible with replication

technologies but may be very expensive to implement. In contrast, tape operations inherently, cost-effectively introduce this gap, such that at some point, a recoverable image is available.

Disk technology has more recently bifurcated between solid state disk (flash) and low-cost, low-performance disk. Some organizations use this low-cost disk for archival purposes. Even though it is lower in cost than other SSD/HDD technologies, it is inherently more expensive than tape and requires some type of environmental support in terms of floor space, power, and cooling. Data stored on tape, in contrast, has the lowest cost per gigabyte and does not require power or cooling inherently, although the hardware needed to read the tapes obviously does. However, large quantities of data can be stored on tape for a very little cost for years to come.

FUTURE OUTLOOK

Considering LTO Tape

Throughout the modern computing era, tape has been the cornerstone of effective data protection strategies. Although various tape formats and technologies exist, the predominant tape platform for open systems is linear tape open (LTO), supported by a variety of vendors that cooperate on standards, development, and implementation. Originally released in 2000, LTO technology is now in its eighth generation, LTO-8. This latest generation features a 12TB cartridge capacity (30TB compressed) and transfer rates of 360MBps native (750MBps compressed). Transfer rates are especially important when organizations are planning and sizing recovery efforts and comparing recovery from tape to recovery from cloud. LTO technology is supported by every major manufacturer of tape automation systems, and its specification road map has been published through LTO-12, meaning that it has a defined and stable future.

Some organizations have attempted to eliminate tape from their data protection environment, believing that it cannot deliver the RPO or RTO required by the business. However, such thinking ignores the range of threats and recovery scenarios where tape is an important option. Snapshots and replications do offer better RPO/RTO in many cases but may not protect from malware or other threats. Cloud replicas offer offsite storage but may have extremely long recovery times for large data volumes. Each technology has its place, and tape continues to fill a need around large data volume recovery and malware protection. Thus IT leaders must recognize that the use of tape is not an "either/or" for data protection, but rather an "and" technology for specific threats and uses at a cost per gigabyte unmatched by disk technologies. Interleaving these technologies creates a multitier data protection strategy to meet the spectrum of data loss threats at the optimal SLA.

The Linear Tape File System (LTFS) is an important part of the LTO specification. LTFS provides a simple and intuitive means for accessing data files on tape using a native OS browser. This means that files can be found and viewed on tape in a manner very similar to finding data on a disk volume. LTFS was adopted by the Storage Networking Industry Association (SNIA) and is now an industry standard (ISO/IEC 20919:2016). Thus LTFS is a nonproprietary tape format file system that makes the tape self-describing, which helps improve archive management of files with no application dependencies. LTFS permits data sharing across platforms, including Linux, Mac, and Windows.

Also LTFS combined with flash and tape storage (flape) can effectively work together to provide seamless storage solutions with excellent performance, high capacity, and very competitive low costs. Emerging workloads and use cases that require fast transactional processing of data that will remain unaltered will be especially good candidates for flape solutions.

A concern of some organizations is that tape media can be stolen or lost and, therefore, read by the wrong parties. While this is a possibility, stealing tapes is not the preferred method used by cybercriminals and requires infrastructure not common for personal use. Moreover, LTO drives use AES256-GCM encryption, the current state of the art technology for commercial encryption. It is authenticated encryption that achieves very high speeds and low latency in hardware. Thus LTO tapes are highly secure and virtually impossible to read without both the right hardware and encryption keys.

Another benefit of LTO tape is that it supports WORM (write once, read many) technology. WORM technology ensures that data cannot be changed, making it secure against both malware and internal threats. This provides compliance assurance for data that requires it, but can also serve as a fail-safe copy of data that may be useful for normal data stores, even if compliance is not specifically required.

CHALLENGES/OPPORTUNITIES

Tape is a solution that has been around for years — sometimes resulting in an inaccurate impression that newer technologies are inherently better. Some organizations are attempting to eliminate tape from their backup infrastructure, believing that it is a cost- or time-saving move. Tape may require some manual handling, but this manual handling has the benefit of creating an inherent air gap in the data stream to help halt the propagation of malware and prevent attack. Further, tape automation solutions minimize the need for manual effort and some solutions can even be designed where tape cartridges do not have to be handled at all.

Organizations often replace tape with replication products and cloud repositories. While these solutions have their place in a comprehensive data protection scheme, they are inherently more expensive than tape, especially as it relates to the cost of retrieval and accessing data; the higher product and environmental costs may offset any labor savings and may not provide the same level of assured recovery that tape can.

CONCLUSION

Data protection needs and techniques have evolved significantly since tape was the primary protection media. Snapshots, mirrors, synchronous/asynchronous, and remote replication as well as the cloud all play important parts in a comprehensive data protection strategy. Each has their own strengths and weaknesses that match up to specific data loss threats.

Nevertheless, tape continues to address certain data protection needs with its capabilities and low cost that other technologies cannot match. Chief among these capabilities is a bulwark against the rising threat of ransomware and other malware designed to infect replicated data from end to end. The air gap nature of tape operations, along with its ability to maintain separated data images that cannot be corrupted, allows tape to serve as a fail-safe method against ransomware. Moreover, if a company's entire data environment were to be compromised, tape can restore data at speeds that cannot be matched by replication from the cloud.

Tape is receiving a resurgence of interest as many leading IT organizations recognize the unique features of the technology with the added benefit of very low cost of ownership. Moreover, cloud providers are using tape as part of their infrastructure to offer low-cost, reliable archive solutions.

IDC recommends that IT organizations view tape as an "and" technology rather than an "or" technology in their data protection strategies, especially when looking at a multitier data protection

solution. Tape has specific strengths that match well to specific data threats at price points far lower than comparable methods. LTO tape has become the open system standard for commercial applications, is supported by numerous vendors, has recognized industry standard technology, and has a published road map for the next four generations. Tape, specifically LTO technology, remains as the last line of defense against malicious software and still plays an important part of any comprehensive data protection strategy.

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