How to Store & Protect an Exabyte

White Paper | Spectra Logic

Table of Contents

About Spectra Logic

Spectra Logic modernizes IT infrastructures to preserve, protect, and defend data from days to decades, whether on-premises, in a single cloud, across multiple clouds, or in all locations simultaneously. Our cost-effective solutions enable organizations to manage, migrate, and store long-term data efficiently, ranging from terabytes to exabytes, with features that make them resilient against ransomware.

To learn more, visit www.SpectraLogic.com

Elephant in the Data Center?

Although you've probably seen calculations on how much data is stored in an exabyte, starting a white paper on exabyte storage is hard without setting the stage for just how much data that is. In the mid-1980s, a company called "Exabyte" was started. Exabyte Corporation introduced an 8mm tape drive and cartridge for data storage, which held an amazing 2.2 gigabytes. Although the company's name was lofty, it would have taken over 450 million of those tapes to hold an exabyte.

What does an exabyte of data look like in layperson's terms? A single edition of The King James Bible comprises over 3.1 million letters. If we decided to go digital, we could store more than 320 billion Bibles-worth of text in an exabyte. Of course, a modern-day analogy would be more like digital photographs. At 3 megabytes per photo, if we printed an exabyte-worth of photos and laid them side-by-side, they would stretch more than 48,000,000 miles — enough to wrap around Earth's equator roughly 2,000 times.

No matter how excited we are about Bibles or selfies, we will probably never need 320 billion of one or 48 million miles of the other. So, do we even need to discuss storing, managing, and protecting an exabyte of data?

Most of us are familiar with the phrase, "Calling out the elephant in the room." The "elephant" is a "mostly known" but "mostly ignored" problem or issue that no one wants to acknowledge. Is exabyte storage the new "elephant" in the data center?

In mid-March of 2020, data backup provider Backblaze announced that it stores an exabyte of user data. What made the announcement unusual is not the amount of data but the fact that they announced how much data they manage. Large data centers are notoriously secretive about how much data they store or how they store it. Companies like Google, Amazon, and Microsoft rarely disclose their data quantities, but it's been theorized for years that they all hold tens of exabytes of data.

High-Performance Computing (HPC) environments often work with hundreds of petabytes of data.

When dealing with hundreds of petabytes, getting to an exabyte takes counting to 10. And that's precisely what is happening in data centers across the globe. It's not just social media giants, cloud providers, and HPC data centers. Increasingly more organizations are working with hundreds of petabytes and are rapidly on their way to a public announcement (or not) of storing an exabyte or more.

There's never been a better time to discuss how to store, manage, and protect an exabyte — whether you are there already or plan to be in the next five years.

What Would an Exabyte Look Like in an Archive Tape Library?

To bring the discussion closer to home, an exabyte is one quintillion bytes or a 1 with 18 zeros behind it — 1,000,000,000,000,000,000 bytes.

In the past, storage products were rated in terms of compressed capacity, but storage software either pre-compressed or encrypted most content or data. The Spectra® TFinity® Plus tape library is the world's largest-capacity storage system. The TFinity tape library can hold over 1.6 EB of data based on LTO-10 native capacity or over 4.2 EB based on compressed capacity. This is an industry first as it is the only tape library in the world capable of holding more than 1.6 exabytes of uncompressed information — or as Spectra Logic calls it — ExaScale storage.

A 45-frame TFinity tape library with 168 LTO-10 tape drives holds 56,400 LTO-10 tape cartridges. At 30 terabytes of native data per tape, this equals 1.692 EB of uncompressed data in a single tape library. Physically, this is a large library. It measures 109 feet long and 3.6 feet deep, meaning it will occupy 391 square feet of real estate. Although the footprint is hefty, the data density is the best available: a TFinity of this configuration offers 4,263 TB per square foot, which is highly compact for use in a modern data center.

As no data center “starts” with an exabyte of data, the TFinity library can be configured as a 3-frame unit and scaled up. This white paper examines the concept of storing an exabyte, presenting various configurations, sizes, and associated costs.

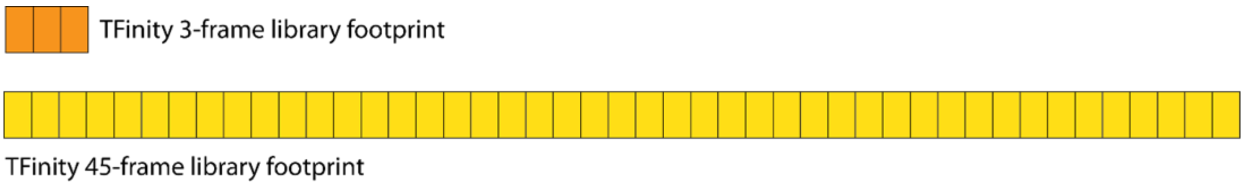

Max capacity in tapes/drives

How Does an Exabyte in an Archive Tape Library Perform?

Several factors must be considered when evaluating the performance of a tape library. Examining the tape library alone — the number of tape drives, the speed of the robotics, and the internal library operating system will determine the maximum performance capabilities. The TFinity library can be configured with up to 168 LTO-10 tape drives. The internal real estate of the tape library balances the number of tape drives and tape cartridges. For example, with 48 tape drives, the 45-frame TFinity is capable of 19.2 GB/s at the rated native throughput of 400 MB/s per drive. This means over 69 TB of data per hour can be read and written to the storage system. Over 605 PB can be transferred into or out of the library in one year. If more throughput were needed, those figures could be increased 3.5 times using 168 tape drives.

While disk performance is often touted as better than tape, it’s important to remember that an exabyte of data under management is typically archived data versus primary, transactional data. Hard disk drives (HDDs) certainly have the performance edge on “time to the first bit of data,” but they would fare terribly against tape in an exabyte-size archive. Using 16 TB HDDs would require a bank of 62,500 HDDs to create an exabyte archive system. The acquisition price of the HDDs alone would eliminate them as a contender. Still, when considering ongoing costs (such as power and cooling), tape has a clear advantage for long-term archiving.

In our tape library example above, with all 48 drives reading and writing at full speed and the robotics moving at their highest velocity, the library will still only use 4,300 watts of power. In comparison, a typical disk storage system with 40 to 60 hard drives (3.5-inch) will consume between 1,200 and 1,500 watts of power with all drives and the motherboard running. With 16 TB hard drives, each disk storage system will hold 1 PB. Four of these disk systems will hold 4 PB and consume the same amount of power as the tape library, which stores 250 times more data.

Floor life is also a major concern when discussing the archiving of an exabyte of data. The floor life of a disk is typically only three years. The floor life of a tape library is 10 to 12 years, and the shelf life of tape media is 30 years.

While all forms of storage have their place, tape is the clear winner in long-term archive projects.

Tiering for Exabytes

Organizations have large amounts of data critical to their businesses, and those data sets are multiplying. Often, these organizations store all active and inactive data in expensive primary-tier storage, which is intended for active data. However, over 80 percent of the data is typically inactive and stored on an inefficient tier, resulting in millions of dollars in annual costs. As organizations approach exabyte storage, they must reexamine data storage methods for data tiering. Older tools, such as Hierarchical Storage Management (HSM), continue to be used in very large tape archives, while modern tools like Amazon S3 Glacier for public and private clouds, as well as Storage Lifecycle Management, are emerging for new, large-scale storage implementations, as well as lower-end solutions.

The Traditional Storage Paradigm

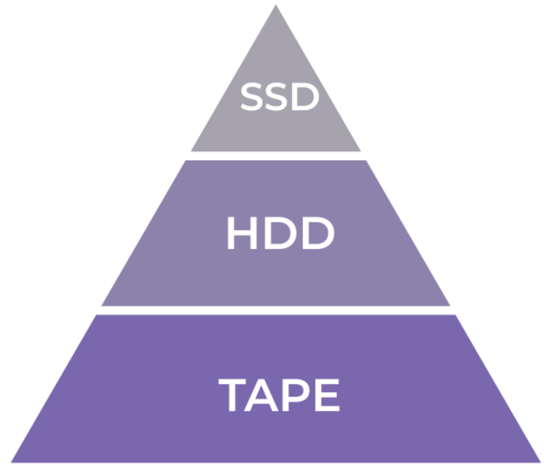

Traditional Storage Paradigm

The storage pyramid is one of the most widely applied storage models in our industry. It is usually represented as three tiers but could depict any number of tiers or combinations of storage technology entities, including cache, RAM, SSD/flash, FC/SCSI disk, SAS/SATA disk, tape, optical, and more. It makes the critical observation that the top of the pyramid is the most responsive, costliest, least dense, and smallest amount of storage in the ecosystem. All those attributes flip as data proceeds down the pyramid. The lowest level of the pyramid is the least responsive, the least costly, and the densest, typically accounting for the largest amount of storage in the ecosystem.

While the basic concept of the storage pyramid is as relevant today as it was 30 years ago, it is a model that doesn't address the newer challenges of modern storage, especially as we approach an exabyte.

With the introduction of the public cloud and object storage technologies, the hierarchical nature of the traditional paradigm has become less effective. This is especially the case when different storage technologies are used in similar roles — SSDs and HDDs in the top tier, disk and tape in backup, and tape and cloud for disaster recovery (DR) and offsite storage. The roles of these technologies may be similar. Still, there are granular differences that enable them to meet the demands of individual data centers and significantly offset costs if those differences can be accounted for.

Likewise, today's storage model must consider the advent of new storage formats. As object storage grows, historical models do not answer many questions. Does object storage apply to a single tier, or do we see block, file, and object storage being used across multiple tiers and intermixed?

A Two-Tiered Storage Model

Groundbreaking storage management software has enabled a modern storage paradigm based on a two-tiered storage model. Rather than focusing exclusively on the storage medium, this model is based on the data or digital content that is being stored. We start by classifying data into two categories — "active," meaning it's being edited, processed, or changed somehow, and "inactive," which refers to everything else. This results in a Primary Tier for the active data and an Archive Tier for inactive data.

The Primary Tier stores all active data and is typically composed of flash, NVMe, or solid-state storage. Organizations can significantly reduce their storage footprint by moving inactive data from the Primary Tier to the Archive Tier. This allows administrators to configure this tier more effectively by utilizing a combination of high-speed storage mediums to achieve the required performance for workflows associated with highly active data.

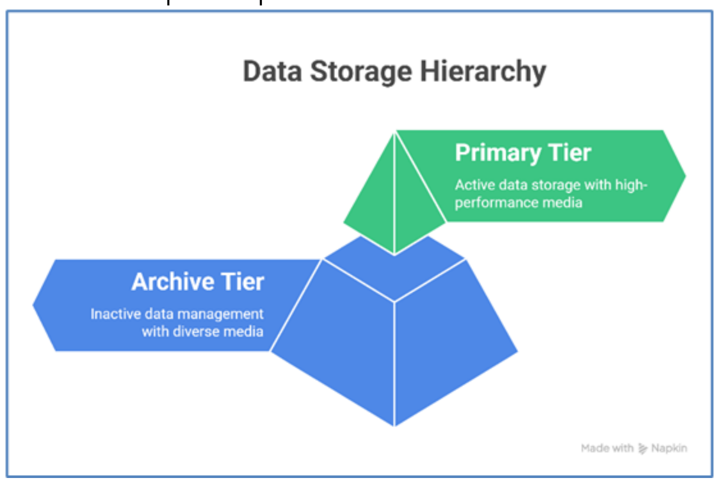

The Archive Tier is dedicated to inactive data and is designed to keep multiple copies of data on various storage mediums, including NAS, object storage disk, cloud, and tape. While the data is not considered "active" on the Archive Tier, quite a bit is happening at this level. The Archive Tier is used for secondary storage, distribution, multiple copies (a responsive copy and a disaster recovery copy), backup, archive, project archive, and traditional disaster recovery. The Archive Tier is the area in which the exabyte archive should reside. As mentioned above, there are multiple mediums in the Archive Tier. When it comes to archiving an exabyte, or even one percent of an exabyte, tape has numerous advantages.

Protecting an Exabyte

Archives are created for several reasons. While sometimes overlooked in value, they are the very essence of both our history and our future. In the words of author and poet William Feather, “The wisdom of the wise and the experience of the ages is preserved into perpetuity by a nation’s proverbs, fables, folk sayings, and quotations.” We will later examine how archives are utilized in the Use Cases section of this white paper. Regardless of the intended use, archives are designed to last and endure for generations to come.

Therefore, data reliability is paramount in digital archives. Individual tape drives have surpassed the reliability of hard disks' bit error rates. LTO-10 tape boasts a bit error rate of 1×10^{20}, while IBM TS1170 tape technology is at 1×10^{20}.

Tape Medium	Hard Error Rate
LTO-8	1 x 10E19 bits
LTO-7	1 x 10E19 bits
TS1155	1 x 10E20 bits

Undetectable bit error is another aspect of data reliability and an even greater one on which to focus. Unlike bit error rates, an “undetectable” bit error is often called “bit rot,” meaning data corruption. Tape has a tremendous advantage over disk in this category as well. Tape boasts an undetected error rate of a single bit for every 1.6×10^{33} bits it reads. To put that into perspective: If you had one million tape drives and one million disk drives running simultaneously, you would get a single undetectable bit error on tape once every 22.5 billion years. In comparison, you would have 1,577 undetected bad sectors every single year with disk.

Hard disk unreliability is mitigated by technology such as RAID. Tape reliability is significantly increased by creating multiple copies. The best-practice recommendation to have two copies on tape could be considered mirroring data. While multiple copies of data on different mediums are always recommended, tape cartridges are the most cost-effective way to ensure data availability at half a penny per gigabyte.

STaRR — A Compelling Approach

Why keep an exabyte of archive data around? The primary reason is so organizations can meet compliance and regulatory requirements. The other purpose is that between one and five percent of disaster recoveries come

from this archive, and it is almost impossible to know what percent of that one to five percent of that exabyte of data will need to be recovered in response to a request. The organization is forced to keep all the data. The capacity requirements are inversely proportional to the number of recovery requests.

While the cost of being unable to restore from an archive is tremendous, a well-planned archive is rarely accessed for restores. This introduces an opportunity for extreme protection at minimal cost.

> **STaRR or 'Store Twice and Read Rarely' is a compelling way to use tape for increased reliability, security, and availability.**

Keeping multiple copies on tape eliminates the risk of data loss due to tape failure. Modern tape cartridges have an annual failure rate of one in 100,000. Low as it is, it could result in a failure rate of about one cartridge every two years in our exabyte-scale library. Implementing a two-copies-on-tape policy reduces the annual odds of data loss due to tape failure to one in ten billion. With each additional copy on tape, the odds decrease even further. And again, the financial burden is one of the lowest in data protection available.

The Spectra TFinity offers enough tape slots and storage capacity to house single and multiple copies of data on-site.

Another advantage of tape is its "removable" nature. Sending one copy off-site provides geographic separation, further ensuring business continuity in the event of a natural disaster, such as a flood, fire, earthquake, or other similar event.

Historically, disaster recovery referred only to those types of disasters. Today, all organizations must be mindful of cyberattacks. Ransomware has become the leading form of cyberattack and the most damaging for organizations that lack a strong disaster recovery strategy.

Storing a disaster recovery copy of data in the cloud creates geographic separation but doesn't eliminate the possibility of ransomware encryption or data destruction. Where there is an electronic path to data, there is also vulnerability. A true disaster recovery archive must be both off-site and offline. Tape offers a true "air-gap" between data and the cybercriminals' attempts to extort it. Even if the backup or archive server is encrypted, and they are often the first targets of attack, tape can rebuild data from scratch. It is not an easy or quick task, but it is a fail-stop measure that assures an organization can recover without paying millions to reclaim its stolen data.

User-Based Encryption

Data breaches, in which data is copied or stolen, continue to be a significant threat. Whether cyber-espionage or identity theft, data breaches are extremely costly in soft and hard costs. An organization's reputation and stock price (if public) are significantly damaged after such breaches. These breaches are a significant liability, both financially and legally. Various industry, federal, state, and world organizational laws mandate how certain information is handled.

Just as cybercriminals use encryption to prevent organizations from being able to access their data, organizations can use encryption to ensure that criminals cannot see or reproduce their data. The Spectra TFinity library offers multiple encryption solutions from which to choose. Spectra Logic is the only library manufacturer that provides fully embedded encryption key management within the tape library. This avoids the cost and complication of additional servers, applications, support agreements, etc. The Spectra LumOS library management software offers two versions of encryption key management: BlueScale Standard Encryption Key Management and BlueScale Professional Encryption Key Management. Standard Encryption is free to all libraries as a function of the Spectra LumOS library management software.

For more complex environments, Spectra also offers externally managed encryption solutions. Our Spectra Key Lifecycle Manager (SKLM) is often selected when multiple libraries or data centers are managed from one area. This involves a server or virtual machine (VM) for hosting, but allows for additional feature sets. Our SKLM solution is FIPS, KMIP, and IKEv2-SCSI compliant. Additional features in SKLM are more extensive audit trails, key grouping, assigning a single key per tape if desired, separate key states, and other features that target policy-based management.

Spectra libraries also support the Hewlett Packard Enterprise (HPE) encryption solution, Enterprise Secure Key Manager (ESKM).

Making Exabyte Tape Archives Easy

New tape technologies are more reliable and offer greater capacity, speed, and affordability per gigabyte than ever before; however, tape must integrate into the modern data center to be viable for storing "ExaScale" archives.

Multiple Tape Technologies

LTO and IBM TS Tape Technology are the only viable choices for today's exabyte-scale archives. However, Oracle's departure from manufacturing the Oracle® T10000 tape drive complicates this situation. Hundreds of thousands of those tape cartridges are still in use, and there is no viable way to continue supporting that technology.

Spectra TFinity libraries are the only ones supporting LTO, IBM® TS11XX, and the Oracle® T10000 tape technology. As made clear, an exabyte of data comprises thousands of tapes. The TFinity makes it simple to import all existing tape, regardless of whether it's LTO, IBM® TS, or OracleT10000. The TFinity is easily partitioned to allow these very different technologies to coexist. Data can be migrated across mediums in the background while current operations are underway.

Spectra TFinity Tape Libraries support three tape technologies.

Multiple Architectures

Historically, tape architecture was limited to traditional file system storage, often resulting in difficulties with scalability and workflow. Spectra offers File System, Object Storage, and even Transparent Access architectures to make tape extremely versatile in the age of ExaScale archives. File system storage has been a mainstay of the storage industry for decades. While that may sound like dated technology, it is still the predominant form of storage and is a proven architecture for many ExaScale archives. File system storage is commonly found in traditional backup and recovery, as well as disaster recovery software applications.

Object storage is a method of structuring data storage similar to a file system, but without the traditional file systems' limitations on growth and performance. Object storage was a mandate for most cloud storage, where trillions of files needed to be stored in a single namespace. File systems could not scale to that level.

Objects are like files containing a single file, document, picture, etc. However, they differ in their use of a unique object ID for tracking and organization. File systems include not only data but also the metadata, physical location, access rights and other information pertinent to that piece of information. This is very useful for transactional, high-edit data, but becomes an unnecessary burden in long-term archive situations.

Object storage is ideal for ExaScale archives; however, until recently, object storage management was not available for tape. Spectra has changed that scenario.

Software running in conjunction with a tape library accomplishes transparent access to tape. This comes in two broad forms: HSM software and Data Storage Management software. Both approaches offer more seamless access to tape, further expanding the workflows that can be accomplished with tape. Spectra tape libraries are designed to support these architectures, making them ideal for archives of any size and scale. However, the best archive solutions are in the multi-petabyte to exabyte size range.

Broad Application Support

Spectra has the broadest range of open systems application support available. Often, a single library is used for multiple applications. As libraries become larger, this aspect becomes increasingly important, representing a significant opportunity to increase efficiency and reduce costs.

Spectra differentiates itself from other tape approaches to large archives because it supports many applications that haven't historically supported tape.

Below is a sample of the numerous applications that Spectra libraries support.

It's challenging to categorize software applications into a single category. We will break this discussion into the following categories for exploration purposes:

1. Traditional Backup &Recovery and Disaster Recovery
2. Enterprise HSM Archive
3. Modern Data Management Software Archive
4. Object Storage Archive

Traditional Backup & Recovery and Disaster Recovery

In this white paper, “backup” refers to making a “copy” of the original data while leaving the original data in place. Active or transactional data should remain on primary storage for greater access speed. The term “archive” refers to the process of moving original data from primary storage to another location, often referred to as migration. This approach is typically used for migrating less active data to a more appropriate storage tier, such as the Archive Tier of storage mentioned earlier.

The backup applications discussed below often have some archiving capabilities but focus more heavily on data copying than data migration.

Backup and recovery are undoubtedly the largest category of software support we will cover. In the past, tape was the mainstay of backup and recovery. At the lower end of this market, higher-density disk drives and public cloud services have reduced the presence of tape due to the ease of recovery. As we discuss the higher end of this market, meaning greater amounts of data, tape still plays a vital role. However, tape is increasingly playing a larger role in archives. In exabyte archives, tape is the leading technology of choice.

Some major backup application players are Cohesity, Commvault, Dell/EMC, IBM, Rubrik, and Veeam. There are many more, but these are the names usually found in the top vendors for “application-based” backups.

All the above solutions, and dozens more, are supported by Spectra libraries. It is common for tape libraries to be “partitioned” to appear as multiple virtual libraries. In this manner, an exabyte-size TFinity tape library could serve as the repository for various applications and workflows. A single organization could use one application to backup Windows-based systems, a second application to backup Linux/Unix-based systems, and even a third application for archiving — all going to a single TFinity library. As mentioned, Backup and Archiving serve different purposes, so this approach is common.

Enterprise Hierarchical Storage Management Archive

Today's HSM market has fewer vendors than the backup market. We often see the names IBM HPSS, HPE DMF, and Versity VFM. HSMs try to map nearline storage to appear online, including tape storage. This is tricky given the long latency of nearline storage. Applications will time out if they don't receive the requested data within the expected timeframe. To accomplish this, HSMs typically use stub files along with filter drivers. The stub file looks like the original file and often contains the beginning of the original file. The stub file can respond to the read request while the HSM retrieves the remaining file and returns it. This is no small feat.

HSMs become part of the file system, usually have kernel code, are very dependent on the operating system (OS), must be upgraded with the OS, and do not allow the application or the user to identify if a file has been moved.

For the above reasons, HSMs don't play well in all environments. They do, however, play exceptionally well in High Performance Computing (HPC) environments where the operating systems — like Lustre, GPFS, etc. — are more "time-out tolerant" or "HSM-aware." Successful HSMs such as those listed are highly effective. They are also expensive, complex, and require a lot of resources, but when you need them, nothing else will do.

Modern Data Management Software Archive

Today's solutions in the data management software category differ from enterprise HSMs because they require less budget, headcount, and infrastructure, and sit well outside the file system. Some exceptions exist, but this categorization generally applies to modern Storage Management Software.

These packages are less complex and compatible with a broader range of applications and use cases. This "class" or "group" of applications has many feature sets. Some of the more well-known names are Arcitecta MediaFlux, StrongBox StrongLink, Komprise, and Spectra StorCycle software.

Symbolic links and/or HTML links are more likely to be used to find the moved data than stub files. These links work quite differently from stub files and from each other.

The user can redirect a read to where the file has been moved by leaving a symbolic link in place of the original file. This works great when moving infrequently accessed data off primary storage (high-speed disk or SSD) to a lower tier of storage like network-attached storage (NAS) disk. Most applications can tolerate a slight increase in latency. However, this methodology does not work well with tape or low-response-level cloud. That is where the HTML links come in.

Spectra created HTML links to support storage mediums with longer latency, such as tape or low-response cloud levels. This is a unique feature of Spectra StorCycle software. When an HTML link is left in place of the migrated file, the user is presented with an HTML page which states that the file has been archived, gives information about when, where, and how this was done, and allows the user to start a restore from tape or a recall from cloud without having to contact IT.

Again, the feature sets of each solution vary greatly. All of them support the Spectra TFinity, and any of them could be used with Spectra TFinity libraries to create and manage an ExaScale archive. When evaluating the appropriate data management solution for any site, individual use cases, IT initiatives, budget, and desired management style are considerations.

Some commonalities in this group include the fact that they all focus on data protection and archiving. While enterprise HSM focuses primarily on data tiering, data management software also incorporates the concept of "multiple copies in multiple locations" and typically works quite well with object storage.

This can be an effective form of data protection. Having copies checksummed and stored on disk, tape, or in the cloud (or any combination thereof) helps ensure that a valid copy of the data can be retrieved when needed. Likewise, administrators can direct a file to be moved to multiple locations and retrieved from whichever location is most efficient.

Another commonality is that these software packages are all integrated into the Spectra BlackPearl® Converged Storage Solution. This enables the unique ability to bring object storage management to tape, representing a significant advancement in creating and managing ExaScale archives.

Object Storage for Tape

Object storage has been one of the most exciting developments in archiving over the past two decades. The concept emerged in the mid-to-late 1990s and has undergone rapid development since the early 2000s.

Object storage is best described as an architecture that manages data as objects, unlike other storage architectures, such as file systems, which manage data in a hierarchical file structure. Instead of embedding each new piece of data in the file hierarchy, along with other files, directories, subdirectories, folders, etc., object storage stores each piece of data independently using a unique object ID.

This creates a flat structure (as opposed to a hierarchical one) that allows for exceptional scalability, independent metadata addition, and tiering. Interfaces can be directly programmable by the application, namespaces can span multiple instances of physical hardware, and most importantly, data management functions like replication — allowing data distribution to occur at a much more granular, object-level.

For these reasons, object storage has become the de facto standard for storing information in the cloud. Only recently has object storage for tape been introduced, bringing the best of both worlds together — simple, long-term data storage management with low-cost, long-lived tape automation. This is key for ExaScale tape archives.

Spectra offers object storage via the BlackPearl Converged Storage Solution. BlackPearl is an intelligent object storage gateway and data management system that allows files to be written to tape, NAS, or cloud as objects. The BlackPearl approach to object storage for tape is unique in that files are managed as objects, yet they are written to tape using the non-proprietary Linear Tape File System (LTFS) format. Users can leverage the benefits of scaling, fast indexing, and searching features in object storage. Simultaneously, tapes are generated under the LTFS format, allowing for complete interoperability with other systems. This ensures that data is accessible in the future, with or without a Spectra solution, providing a non-proprietary archive.

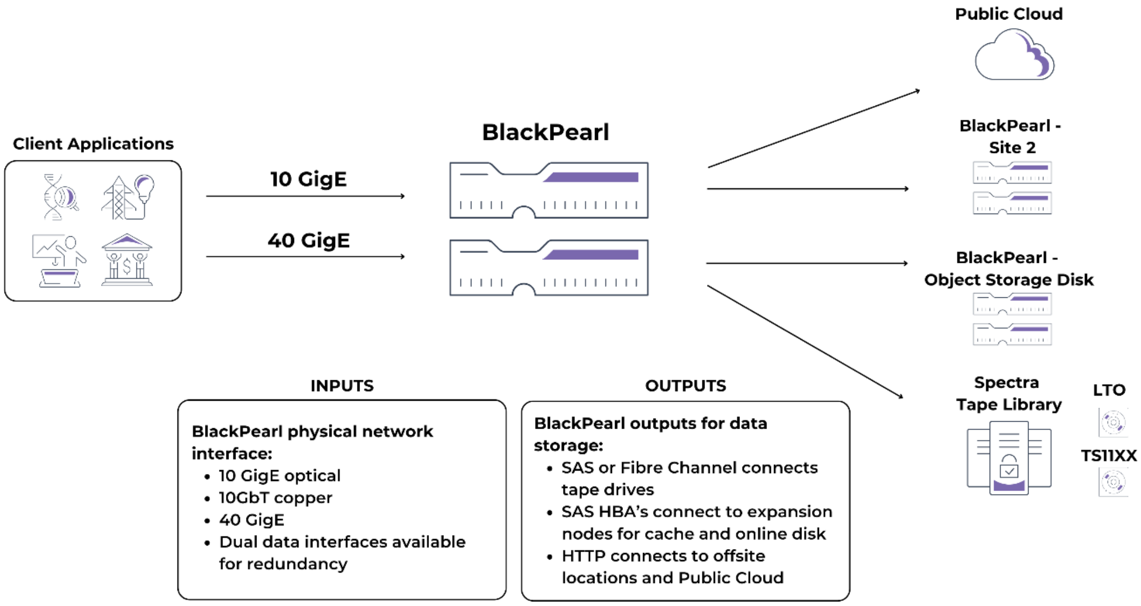

Data can be archived via the BlackPearl object storage solution through several paths or workflows.

Spectra provides a simple, published API for BlackPearl Object Storage Gateway, enabling software manufacturers and end-users to interface with BlackPearl. Commercial applications, such as those listed under data management software above, offer support to simplify writing to tape in an object storage environment.

Many applications in the media and entertainment industry have also integrated support for BlackPearl. These edit suites have not previously offered direct tape integration, requiring additional, expensive software to support archiving. For those integrated with BlackPearl, it's now as simple as selecting the "archive" button.

Over 40 commercially available applications are now integrated into BlackPearl. It's so simple that users can even write their own interface to customize how they archive or backup data.

Spectra BlackPearl can also migrate data from proprietary solutions into open, non-proprietary solutions. In the media and entertainment market, BlackPearl software will import data from the Front Porch Digital DIVA solution and the SGL FlashNet solution, allowing users to replace older, expensive, proprietary archives with a simple, open, modern design at significantly lower cost.

Furthermore, BlackPearl utilizes the Hypertext Transfer Protocol (HTTP). HTTP is the underlying protocol used by the Internet. It defines how messages are formatted, transmitted, and can work with delays in response. This is an ideal protocol for moving, retrieving, and sharing data with higher-latency storage, such as tape and cloud. APIs based on the HTTP protocol are now commonly referred to as REST or RESTful interfaces.

Broad application support is critical when building a multi-petabyte or exabyte archive. Spectra TFinity libraries offer the largest selection of supported vendors. From commercial backup applications to HSMs to data management solutions to custom interfaces, TFinity libraries are designed for archives in any environment imaginable. The TFinity LumOS library management software assures simple integration as new applications become available.

Use Cases for an Exabyte Archive

It's all too easy to become mired down in the "bits and bytes" of constructing a large archive and lose sight of the value it can bring to organizations worldwide. For example, NASA ran the following headline in a news article: "Earth-Size, Habitable Zone Planet Found Hidden in Early NASA Kepler Data."

NASA launched the Kepler space telescope in early 2009 to discover Earth-size planets orbiting other stars in or near "habitable" zones. In other words, the goal was to discover planets where humans can live. The telescope was retired in late 2018, seven years ago. So why is it still making headlines?

A group of scientists from the University of Warwick developed a machine learning algorithm to sift through old NASA data, which contains thousands of potential planet candidates. This new technology has identified a planet very similar to Earth in terms of size and temperature, which is theorized to support liquid water. Is this the beginning of a science fiction movie? Hardly. This is the result of applying new technology to old data. And new worlds are emerging.

That is the power of archives. As a result, archives are growing larger as more data is generated and more organizations establish their retention policies to "forever."

Not all archives are used for such pioneering discoveries. From scientific research to manufacturing, from modeling to weather predictions, entertainment to world history, medicine to self-driving cars — digital archives help us understand our world, keep us safe, get us where we need to be, and even entertain us when we get there. The following use cases exemplify how organizations create archives, breaking every previous record for size and demand.

BioPharma Research

Historically, drug research has depended on traditional, randomized clinical trials to determine the effectiveness of new and existing drugs. As one might imagine, this research is highly time-consuming and expensive. In many cases, this process can delay drugs from reaching the market by years.

The rapid development of data analytics and AI now allows drug companies to study what is being termed “real-world evidence.” That’s any evidence of a drug’s effectiveness gathered outside the traditional clinical trial setting. But where does that data come from?

The last two decades have seen tremendous growth in the digitization of patient records. Combine that with insurance records, diagnostic and genetic testing, fitness wearables, IoT sensors, and social media. Data is being extracted from all these areas in an unprecedented effort to speed the safe development and delivery of life-saving drugs.

Privacy is a primary concern when using such data, but by anonymizing this information, suddenly millions of people from every walk of life can contribute to drug research. This isn’t predicted to replace clinical studies, but it has already rapidly advanced new medicines, given much greater insight into existing medicines and offered greater “individualized” medical treatment.

Breast cancer is not common in men; only about 1% of breast cancer cases occur in men throughout the U.S. That still accounts for 2,000 cases per year in the U.S. alone. Ibrance, developed by Pfizer, is known to be an effective medicine in controlling breast cancer. When the initial clinical trials were conducted, they did not include male participants. Therefore, Ibrance was not approved by insurance companies for male patients. Conducting the study on men would have required an additional three to five years. Looking at the averages, that’s another 6,000 to 10,000 men who would not have access to this drug during those trials.

Pfizer examined the original study results obtained from women and supplemented those results with data generated by male users (real-world evidence) who had paid for their own medication. Within 12 months, Pfizer demonstrated to the U.S. Food and Drug Administration that Ibrance was safe and effective for both men and women.

Additionally, by significantly lowering the cost of entry into a drug market, pharmaceutical companies can develop drugs in less lucrative markets, such as treatments for diseases that disproportionately strike poorer nations.

This all depends on access to tremendous amounts of data. Genetic testing is another area that has provided incalculable information for such research. Genetic testing generates approximately 300 GB of data per person. It would take approximately 99 exabytes to store the genetic testing data of the entire U.S. population.

Regardless of the amount of data biopharma research companies have access to, more information will lead to faster and more effective solutions for disease prevention, control, and cure. In this context, a single exabyte archive seems small.

Large Video Archives

Media and entertainment consumers consistently crave new content. As broadcasters and producers continue to capture the attention of viewers with TV shows, movies, and sporting events, the need for high-capacity storage will only increase. Many outlets currently utilize high-resolution 4K cameras, which contribute to this massive storage demand. Depending on the number of cameras used to shoot, image resolution, hours recorded, frames per second, image bit depth, and compressed type used, shooting in 4K quickly leads to pressure on data storage. For example, a finished two-hour film shot at 4K with 172,000 frames at 24 fps creates roughly 50 MB of data per frame — approximately 10 TB for just the finished film.

When you examine the ratio of raw footage hours to final footage (typically 6:1 or 7:1), it is approximately 60 TB to 70 TB of data created for each two-hour film.

One can see how a network producing a slew of reality shows could exceed an exabyte of total storage, given the sheer number of hours of raw footage required to create one hour of finished content. The ratio is unheard of, ranging from 100 to 500 raw hours to produce a mere one hour of final video. High-end ratios tout a staggering 1,000 raw to 1 finished ratio. When these shows air weekly, we're talking about massive amounts of data being created!

As if the data creation associated with 4K wasn't enough, 8K grows in popularity, and the tremendous amounts of data created with it are staggering. Shooting at such high resolution allows broadcasters to show "finger-tip" finishes to their audience, creating a highly realistic experience for viewers. In fact, portions of the Olympics are currently shot in 8K. All of that comes at a price. Combine 8K resolution with higher-density and higher-frame rates, and you have the recipe (perfect storm) that would cause any content creator to breach the realm of an exabyte-scale archive. Compared to shooting in 4K, 8K films nearly quadruple the data created when shooting. You can expect around 200 MB per frame for 8K DPX, meaning a finished film will be approximately 35 TB. Then, we consider the raw footage to final footage ratio, and we're talking about 245 TB of data.

Considering the exponential storage requirements of shooting in 8K resolution, it is easy to see how quickly a broadcaster can easily reach an exabyte archive.

Media and entertainment organizations must plan for larger archives that can store all their invaluable raw and finished footage to stay ahead of the changing times and bring new content to consumers.

Research Data and Onsite "Glacier-like" Archives

Spectra works with a large university that supports the research efforts of over 50 different groups within the university. Based on the performance of the storage, they offer various service-level agreements (SLAs). Each research group is billed for the storage it uses based on the SLA it selects.

The university has standardized on three storage performance levels:

- A high-speed, solid-state disk-based tier
- A medium-speed, disk-based tier
- A NAS-based tier targeted for an archive of projects

There are multiple challenges for both the university and the individual researchers that the university would like to overcome.

While NAS is the lowest-cost repository in the current storage model, archiving large amounts of fixed content becomes extremely expensive. Researchers have asked for a lower-cost solution for archiving. The data center acts as a "cloud" provider but lacks the last storage level, often referred to as "glacier-like" due to its similarity to AWS Glacier storage. The university would like to introduce an ExaScale exabyte-sized archive based on tape. Still, they cannot introduce "rule-based" file movement across the storage infrastructure to such a tape archive.

The university actively encourages researchers to move data off the high-speed Primary Tier with a bill-back system that offsets their costs. However, when the university runs out of high-speed storage, it doesn't always have the funds to acquire more before the offsets come in, which can hamper research efforts.

The researchers would like to move their data to a low-cost storage tier, but they face a challenge in accomplishing this. The data can be human-generated, application-generated, or machine-generated, and has been accumulating for years. The researchers have access to the data but cannot identify what is actively used, what needs to be archived, and what data is orphaned. And if the data is manually moved, how can it be accessed after the researcher leaves the university? Research grants often require data storage for periods longer than the researchers who work for the lab or university are involved.

A data management software application offered the perfect solution for both parties, enabling a tape-based archive of extreme size. The university selected the Spectra Logic StorCycle solution, which provides a seamless view across all the storage it manages. Both the university and the researchers have access to all the data managed.

Researchers can use the scanning capabilities of StorCycle software to identify and target inactive data sets for archiving. StorCycle project archives will ensure this situation doesn't recur. Project archives enable users to identify and archive all files and directories associated with a project as a single group. This can be done immediately after a large project is completed.

Archived data sets can be tagged with additional information to identify key aspects of the project, such as grants associated with the project, researchers involved, and project names. This metadata can easily be searched at any point in the future. Likewise, StorCycle software produces a manifest for each project archive, which can be accessed as a digital file. The manifest shows exactly what was moved, where the data originated, where it was moved to, and when. It can be digitally displayed by clicking on the finished project archive and stored with other files in the project. It does not require a query into the database and can be worked into existing workflows.

With several hundred petabytes of data currently under management, the university can deploy an affordable, glacier-like tape storage archive capable of growing to an exabyte size.

Machine Learning and Autonomous Driving

"Deep Blue…."

"Watson…."

"AlphaZero….."

These names evoke great fear or inspiration from the world's top chess players. These are the names of the supercomputers and associated machine-learning programs that have beaten the highest-ranked chess champions. And therefore, these names have become synonymous with the concept and power of machine learning.

Whereas Deep Blue and Watson required hundreds of hours of programming and teaching, AlphaZero made headlines in late 2017 for its ability to self-learn. AlphaZero was programmed with nothing more than the rules of chess. Nine hours later, it had played itself millions of times, learned from the experience, and was able to beat or draw not only the best human players, but also all prior computer champions.

This self-play learns from its success and failure and feeds the output into a neural network.

Large quantities of existing information (such as exabyte-scale archives) are unnecessary for machine learning. But that all depends on what the machine is expected to do. AlphaZero plays games — there are two players, complete information is available to each player, and specific rules must be obeyed. AlphaZero is incapable of making an illegal move. Nor would it be able to play another person or computer who "cheats."

In a zero-sum game — one winner and one loser — with no hidden information or elements, AlphaZero is unbeatable. But if you need a ride to the grocery store… you might want to look for another solution.

Autonomous or self-driving cars based on neural networks must be able to handle a much different situation than a chess game.

In daily real-life scenarios, one would expect two people to go to the grocery store simultaneously and both arrive home safely with their groceries — a win/win scenario. But many unknowns or "hidden information" are floating around out there. Weather, cyclists, animals, poor drivers, drunk drivers — the list of unknowns is too large to cover.

The development of autonomous vehicles such as Waymo involves building extensive neural networks (which reside in the car) that are pre-programmed to process camera, LIDAR, and radar data to spot other moving vehicles, pedestrians, traffic signs and lights, animals, bicycles, motorcycles, objects in the road, and even traffic cones.

The neural network's configuration is created and tested with "training data." This training data is collected from video cameras built into vehicles already in operation with human drivers. The more training data a manufacturer holds and processes, the safer the autonomous vehicle will be. Depending on the level of connectedness, autonomous automobiles are expected to generate up to 5 TB of data per car per day. This training data will grow into exabytes of storage.

This is AI in pure form, and it won't stop here. In the future, numerous applications of AI will drive the need for more training data. Autonomous cars fall under the larger Internet of Things (IoT) category, which is the subject of our next discussion.

Internet of Things

The Internet of Things (IoT) refers to all sensors and cameras that collect and share data. IoT devices may include traffic sensors, weather sensors, factory sensors, sensors on an oil pipeline, video surveillance cameras, your smart doorbell, smartwatches, your home's cooling and heating thermostat, your car, and even your toaster.

The advisory firm International Data Corporation (IDC) forecasts that the number of IoT devices will grow to over 40 billion units by 2025. The idea is that millions of sensors could gather data and send it to a central location for processing. The applications for such data are endless — better traffic control, weather forecasts, manufacturing efficiencies, GPFS, health monitoring, etc. There are civilian as well as military applications. As diverse as IoT devices are, they have one thing in common: they generate data.

And generate data, they will. IDC further forecasts that IoT devices will generate 79.4 zettabytes of data by 2025.

Not all of this data will be saved, but if 0.0000125% of it is saved, that alone would be equivalent to one exabyte of data. Ten to 30% or more of this data could be utilized for future insights and improvements in processes, products, manufacturing, security, and overall quality of life.

The IoT has become so pervasive that a new acronym is now being used — IoE, the Internet of Everything. When we have access to data, we have unprecedented insight into the past and the future. Those with access to the most data hold the greatest insight.

Medical, BioTech, and Genomics

We've already covered the discussion of biopharma and big data — drug development is being reduced by years required to reach the market, bringing more affordable, safer drugs to market more rapidly. However, that is just one area of healthcare that is being revolutionized by access to large amounts of data. Computer systems are also being increasingly used in the diagnosis and treatment of diseases.

In October 2019, a study in the Lancet Journal of Digital Health compared the performance of "deep learning" AI and healthcare professionals in detecting diseases based on medical imaging (X-rays, CT scans, MRIs, etc.). The study looked at 14 studies published between January 2012 and June 2019. On average, AI was slightly better at diagnosing disease than human healthcare professionals (87% correct vs. 86% correct, respectively).

AI is even more effective in certain types of cancer screening and detection. In lung cancer, Google has shown that the neural networks of AI can be trained to detect the presence of cancerous cells earlier and more rapidly than radiologists. When reading mammograms to detect pre-cancerous and cancerous cells, computers are currently three to 10 times more accurate in identifying cancer than radiologists.

A notable difference is that AI can continuously train itself on any new data it receives. It never tires, and it never retires! All it needs is more data to become even more accurate and efficient. Combining AI with the "Internet of Medical Things" — home sensors worn by patients, for example, enables health information to be gathered and processed in real time, and changes to the patient's lifestyle may be suggested.

These systems are typically implemented as neural networks. The more training data an entity has, the more accurate its predictions will be. Many more AI-based healthcare systems are expected to be created and implemented in the future, thereby improving disease detection.

Biological research is another medical field that relies on and generates massive amounts of data. Cryoelectron microscopy is a method of understanding underlying cellular mechanics. A single cell may be frozen at cryogenic temperatures, sliced into layers, and scanned by an electron microscope. This method enables samples to be examined without the use of dyes or fixatives. Since samples remain in their native state, scientists uncover new information about viruses and protein complexes at the molecular level. As many as 500 layers may be sliced from a single cell, producing over 200 TB in a single examination.

Genomic simulation has also opened a new door to medical understanding. Genetic testing and the vast amounts of data it creates allow individuals to understand their risks for certain diseases and use further technology (and data creation) to have the best chance of detecting and surviving what was, a decade ago, considered unsurvivable. We have reached a point where medical regimens are tailored to individuals based on their unique genetic profiles.

These activities generate petabytes to exabytes of data storage.

HPC Applications

The world of High-Performance Computing (HPC) relies on supercomputers which are used for a wide range of computationally intensive tasks in various fields, including quantum mechanics, weather forecasting, climate research, oil and gas exploration, molecular modeling, and physical simulations (such as simulations of the early moments of the universe, airplane and spacecraft aerodynamics, automobile design, the detonation of nuclear weapons, and nuclear fusion).

As one might expect, High-Performance Computing generates and requires enormous data sets. First, computational test scenarios manipulate a digital environment or situation. Data is then run through these test scenarios, and the output of the tests or experiments is collected. A single experiment can create hundreds of terabytes up to multiple petabytes. Likewise, a single experiment may take months or even years to run and cost millions.

This creates a challenge for any experimental computing facility. Experiments are often repeated and reexamined. All inputs and outputs — digital data related to a project — must be preserved. HPC sites rival data repositories in any vertical market and are second in size only to Internet operations.

Although data sets are regularly shared among engineers, scientists, and developers who work in the HPC world, their workflows are more "manual" than those of general IT. Once the output of an experiment or modeling session is finalized, it remains unchanged, so it can be moved to a data repository that requires lower access and, hopefully, much greater density and energy efficiency for storage.

Government national labs are a large part of the HPC community. They create and store vast amounts of data for both civilian and military use. In the 2020s, each of these government labs is expected to grow in capacity to one or more exabytes. The tape-based ExaScale archive is already deployed at most HPC sites.
The "TOP 500 list" of supercomputers shows the size and capacities of the world's top sites.

Weather Data

Weather prediction utilizes sophisticated modeling that runs on supercomputers, which are continually being improved. Weather data consists of data collected from thousands to tens of thousands of sensors. This also fits into the category of the "Internet of Things." These sensors may include ground-based sensors, satellite-based sensors, seawater temperature sensors, aircraft-collected data, and occasionally weather balloons. This data is stored forever, along with the resultant forecasts. Weather scientists can tell if their models are working by evaluating the gaps and differences between the actual outcome and what their model predicted for that time, whether currently or in the past.

Therefore, weather data archives are enormous. Most large national weather forecasters are expected to exceed one exabyte by 2025. More than 25 major worldwide organizations exist, including the National Weather Service (NWS), Meteo France, the British MET, the European Center for Medium Range Forecasts, the Korean Meteorological Association, etc.

As sensors become more prevalent and higher in resolution, they will collect more data. As supercomputers continue to grow in computational power, we can expect to see a corresponding increase in the storage required for archives.

High Energy Physics

When the desire is to push the limits of human understanding to answer some of life's most complex questions, the research's amount of time and complexity demand the most secure and scalable storage available. CERN, often called the European Laboratory for Particle Physics, uses some of the world's largest and most complex scientific instruments to study the fundamental particles of matter, quite literally discovering the "God Particle," the missing cornerstone in our knowledge of nature.

They do this with the LHC (Large Hadron Collider), the world's largest and most powerful particle accelerator.

The LHC is 27 km (16.8 miles) of superconducting magnets in the shape of a ring, which sends high-energy particle beams close to the speed of light until they collide, to figure out how the universe was formed. This research means the CERN Data Centre produces an astronomical one petabyte of data per day, pushing their CERN Advance Storage System (CASTOR) to a massive 400PB of data stored on tape.

According to CERN, this is equivalent to 2000 years of 24/7 HD video recording.

Currently, the LHC is shut down while getting a significant performance upgrade. This upgrade will keep it shut down for two years, just long enough to upgrade their data archival software and tape system to handle the much higher data volume.

Once the LHC resumes its next run, data creation is expected to double from 2027 to 2029. This means they will store an additional 600PB or more after running in 2029. This will put CERN close to, if not over, 1 exabyte of data on tape at the end of 2029. The LHC will then shut down again, and upgrades to many of the sensors, magnets, and testing devices in the collider will be made. When it resumes, we can expect a fourfold increase in the current level of data stored. This means over 1.6 exabytes of additional data must be archived to tape during run four.

The new CERN Tape Archive (CTA) software stores 400 petabytes of data and ALL new data created. It is being designed to handle these massive amounts of data. The data produced at CERN is extremely valuable and must be preserved for future generations of physicists, making tape the ideal archive storage technology to use. It is also shared worldwide. CERN has transferred 830 PB of data and 1.1 billion files to other HEPIX research organizations worldwide. This allows other physicists to conduct research, and it also means that the data is archived geographically with multiple copies, ensuring it can be safely kept forever.

CTA can store an exabyte in native capacity, and CERN is rapidly approaching the need for an exabyte of storage. They are counting on the industry to continue innovating and storing multiple exabytes in the coming decade to preserve them for future generations. With the roadmap for LTO tape, it appears to be a partnership that will stand the test of time.

Petroleum Reservoir Storage and Modeling

Quite a bit has been said about High-Performance Computing. One of the first commercial applications for HPC was seismic exploration for minerals and petroleum. What ultrasound is to doctors, seismic vibrations are to oil drillers. Seismic exploration measures the spread of waves through the Earth. Data explorers can predict material composition deep underground by measuring speed, reflection, and refraction on or near the Earth's surface; Seismic data have been used for mineral exploration for decades.

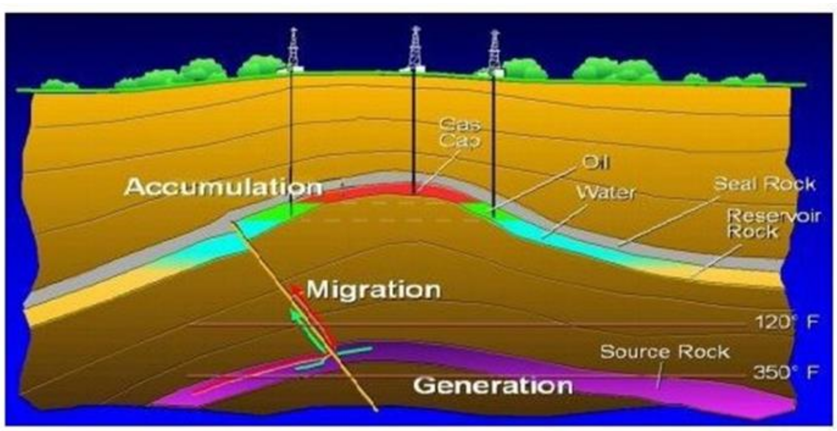

Simplified visualization of hydrocarbon traps recognized in producing oil and gas fields.

(Source: http://www.geologyin.com/2014/12/hydrocarbon-traps.html)

Still, it was far less effective before the invention of high-performance computing, which could generate and analyze the tremendous amounts of data required for efficacy.

Thanks to HPC, the era of drilling to determine what were dry wells and gushers is now as dated as the black-and-white footage that documents it.

Advances in HPC and data analysis have all but eliminated the ambiguity of finding oil, and new developments in extraction, such as horizontal drilling, enable drillers to physically access oil once it has been located with precision using seismic data.

The critical point is that, although technologies for reaching and extracting oil and natural gas deposits have evolved, they still rely on seismic information gathered decades ago on unreachable deposits.

Geologists, therefore, have not had to spend millions of dollars to resurvey and perform new seismic explorations. Regardless of its age, geological data remains as accurate today as it was when it was first recorded and can be reanalyzed using the latest and most accurate methods. This data remains invaluable as new technology emerges; those approaches will be modeled against existing seismic exploration data.

The data generated from seismic exploration falls into the ExaScale category. It continues to be stored on tape, in large tape libraries.

Hybrid Cloud for Archives?

One of the most significant growth areas in our industry is hybrid cloud. "Hybrid Cloud" is a term that's been thrown about ever since the introduction of the public cloud. In this discussion, a hybrid cloud refers to the ability to store data both on-premises and in the cloud. Some data may reside exclusively on-premises, such as highly confidential information or large amounts of data that are too extensive to store in the cloud indefinitely due to monthly charges and egress fees. Other data may reside exclusively in the cloud, such as transitive data, which is unnecessary once a given calculation or computing service is completed. Having copies of the same data stored simultaneously in the cloud and on-premises may be desired.

The ultimate definition of hybrid cloud allows data management to occur from a "single pane of glass" or control mechanism that provides a universal view of the data, regardless of where that data physically resides.

Both public cloud and on-premises storage provide certain advantages and challenges. If done correctly, a hybrid cloud offers the possibility of achieving the best of both worlds. Let's start by comparing cloud storage and on-premises storage. Cost and performance will not bode as well for cloud storage of an exabyte as they do for on-premises storage. That isn't necessarily a strike against cloud storage as much as it is a reason for a hybrid approach of combining on-premises and public cloud.

Before discussing hybrid cloud, it's essential to understand the opportunities/challenges with "cloud only" or "on-premises only" models.

What Does it Cost to Store an Exabyte?

Tape Only

Creating an exabyte archive on tape requires capital and operational expense. The costs are more heavily weighted on the front end due to the purchase of a tape library, tapes, and other related costs. The "pay as you go" model the cloud offers can be replicated with the Spectra TFinity tape library design. As mentioned earlier, the TFinity can start with as few as three frames, one tape drive and 50 enabled tape slots. End users can fully populate existing frames in the field. Once a frame is filled, additional frames can be added for up to 45 frames to store and protect an exabyte of uncompressed data.

We began this white paper by providing a brief description of what an exabyte archive would look like in a TFinity tape library. We'll now show what it looks like, including costs, to grow an extensive archive into a full exabyte archive.

The organizations we're discussing work in a remarkably high data-capture environment. They could be involved with Artificial Intelligence, the Internet of Things, large video archives, HPC computing on an ExaScale supercomputer, or capturing machine data from a Cryogenic Electron Microscope, Square Kilometer Array, or another high-data-output machine.

These environments can easily generate up to 200PB of data per year, reaching an exabyte in five years. To accommodate this, we'll configure an 8-frame TFinity tape library with 48 LTO-10 full-height FC tape drives, two Spectra 2400-OSW Optical SAS Switches (each capable of connecting to 40 tape drives), one year of support, and 50PB of media on day one. The starting cost of this complete configuration, installation and support will be $1.40 million.

Over the next five years, media will be purchased quarterly, at a rate of 50PB per quarter, to take advantage of the regular price reductions offered with LTO media. At the end of the first year, the organization will spend an additional $1.13 million for an incremental 150 PB capacity, resulting in a total one-year spend of $2.53 million and a 200 PB capacity. The starting cost of 200 PB in the AWS Deep Glacier Cloud would be approximately $2.44 million – a bit less expensive for year one.

For the next four years, no additional robotics or tape drives will be needed. Additional media expansion storage frames, media, and yearly support are ongoing costs.

In year two, five media expansion frames are added, bringing the total library size to 13 frames. An additional 50 PB of LTO-10 media will be purchased quarterly, and an extra year of support will be paid, bringing the total additional spend for year two to $1.56 million.

In year three, five more media expansion frames are added, again at a rate of 50 PB of media per quarter and with one year of support, resulting in a yearly spend of $1.40 million.

Year four brings another five media expansion frames, 50PB of media quarterly, and a year of support for $1.23 million. The decreasing cost of media is evident in the lower year-over-year cost for the same capacity added to the library.

The final year of use will add the last five media expansion frames, the final 200 PB of media, and one more year of support, for a total of $1.06 million that year.

This brings the total spend to $7.78 million for the five years, and there is now 1 exabyte of data stored in this environment. As the table below shows, the total cost per Gigabyte of data stored is $0.0078/GB.

One Exabyte	Year 1	Year 2	Year 3	Year 4	Year 5	$/GB
Capacity	200 PB	400 PB	600 PB	800 PB	1,000 PB	
Day 1 Cost	$1,397,267					
Yearly Incremental Cost	$1,133,333	$1,562,492	$1,397,441	$1,228,618	$1,059,795	
Total Accured Cost	$2,530,601	$4,093,092	$5,490,533	$6,719,151	$7,778,946	$0.0078/GB

Cloud Only

The storage cost is always a consideration and increases as data and retention time grow. Cloud often leads with the low cost of storage, usually shown as a “per Gigabyte / per month” charge. AWS Glacier Deep Archive offers storage as low as $0.00099/GB (just under one-tenth of one cent) per GB/month. It’s important to note that this is a “per month” charge, so comparing it in that form to the “total cost” per GB as noted above with tape is a bit challenging.

Storing 200 PB of data for one year at $0.00099/GB/month (the AWS published price per GB/month for Glacier Deep Archive in North America) would cost over $2.4 million annually. As our archive grows by 200 PB per year until it reaches the exabyte mark, calculations for storage alone are straightforward.

The total “storage” expense for an archive growing to 1 exabyte over five years would be slightly over $36 million, or $0.036/GB.

Cloud Archive	Year 1	Year 2	Year 3	Year 4	Year 5	$/GB
Capacity	200 PB	400 PB	600 PB	800 PB	1 Exabyte	
Starting Cost	$2,424,000					
Yearly Incremental Cost		$4,848,000	$7,272,000	$9,696,000	$12,120,000	
Total Accured Cost	$2,424,000	$7,272,000	$14,544,000	$24,240,000	$36,360,000	$0.036/GB

The word “storage” cost is emphasized above because that cost does not include charges or egress fees for accessing data, various levels of storage required to “land” or “stage” data, etc. Nor does it include the cost of bandwidth selected to move the data to and from the cloud - more on that in the sections below.

AWS also offers a seemingly inexpensive cost per Gigabyte to restore data. There are various levels, but the most affordable and least responsive option is $ 0.0025/GB to restore – approximately 2.5 times the amount charged to store the data. Again, with no internal movement charges being calculated, it would cost $2,500 per petabyte to bring back the data. Even a very small read or recovery rate becomes significant at exabyte scale. At 5% per year, read/recovery charges would be $125,000. At 5% per month, those annual charges would be $1.5 million. It would cost $2.5 million to restore the entire exabyte archive. There would be no additional charges if this were restored from a TFinity tape library. Depending on access, this could easily save over a million dollars a year and more than $2 million for disaster recovery or other complete recall.

Applying the cloud pricing model to tape automation (cents per GB) provides more interesting insights. Looking at a “per Gigabyte” cost for tape automation, the total five-year archive, spanning from 200PB to 1EB, would cost $0.008 (under 1 cent) per Gigabyte.

Over the five years, the total cost for cloud storage would be $ 0.036 per GB, making it 4.6 times more expensive to create this archive on the cloud versus tape.

Archives for 10 or More Years

The above calculations are based on a five-year model to reach an exabyte because, as stated earlier, most organizations are on their way to an exabyte rather than already being there. Most archives are designed to last decades or even indefinitely.

What happens to costs after those first five years? Supporting this size tape system will cost a little over $112,000 per year, so the total spend over 10 years will be $8.34 million. This slightly increases the tape archive cost from $ 0.0078/GB to $ 0.0083/GB.

Tape: One Exabyte	Year 6	Year 7	Year 8	Year 9	Year 10	$/GB
Capacity	1,000 PB	1,000 PB	1,000 PB	1,000 PB	1,000 PB	
Five-Year Cost	$7,778,946					
Yearly Incremental Cost	$112,844	$112,844	$112,844	$112,844	$112,844	
Total Accured Cost	$7,891,789	$8,004,633	$8,1171477	$8,230,320	$8,343,164	$0.0083/GB

Cloud storage will incur monthly charges of $12.1 million per year for the exabyte held in archive in years 6 through 10. That would put the total cost of the cloud archive at over $96 million – significantly raising the cost of storage from $.036/GB to $.097/GB -- over 11 times the cost of the same archive on tape! Whereas tape is less than a penny per GB ($.0083), cloud storage costs around 10 cents per GB ($.097).

Cloud	Year 6	Year 7	Year 8	Year 9	Year 10	$/GB
Capacity	1 Exabyte	1 Exabyte	1 Exabyte	1 Exabyte	1 Exabyte	
Starting Cost	$36,360,000					
Yearly Incremental Cost	$12,120,000	$12,120,000	$12,120,000	$12,120,000	$12,120,000	
Total Accured Cost	$48,480,000	$60,600,000	$72,720,000	$84,840,000	$96,960,000	$0.097/GB

Furthermore, the tape archive holds no outside costs to restore data. This eliminates the shock of unplanned expenses when you least need to deal with them – during an unexpected data loss, cyberattack, or natural disaster. Archives are also used to gain new insights into current problems or challenges by utilizing previously collected data.

The Spectra Cube can be expanded from 50 to 1,670 LTO slots and up to 16 full-height tape drives. With LTO-10, a capacity of over 50.1 PB (native).

Having archives immediately available without additional, excessive costs makes future research more viable and increases the value of the archive.

As mentioned, having a second copy on tape virtually assures data availability. In a high-capacity tape system, approximately three-quarters of the total cost of the environment is attributed to media, and one-quarter is allocated to tape library hardware and system support. In this scenario, making a second copy to store outside the library would cost over $5.88 million over the same five-year period using the same media purchasing policy and pricing. Storing two copies in the cloud doubles the $12.1 million for each year a second copy of data is kept.

Pricing Models for Those “On Their Way” to an Exabyte

Similar pricing models can also be used for smaller systems. Below are tables showing the cost for organizations planning to archive 50, 250, or 500 PB. In designing a long-term tape archive, the items that account for cost will be the same as mentioned above—the tape library, support, expansion when needed, and media. Depending on the initial amount of data, Spectra recommends different models of Spectra tape libraries.

Since Spectra tape libraries can scale up to the next size library, even starting with the smallest Spectra library can eventually be grown into a TFinity, while protecting your existing investment.

The Spectra Cube tape library makes sense from a capacity and cost perspective for environments that may need to store 50 PB over five years.

A Spectra Cube with 12 LTO-10 drives would be used to create a 50 PB archive over five years. For the next four years, only media and yearly support will be purchased. It is capable of achieving a total throughput of 4.8 GB/s. The costs for the Spectra Cube example are listed in the table below.

Capacity (PB)	Year 1 – Cost	Year 2 – Cost	Year 3 – Cost	Year 4 – Cost	Year 5 – Cost	Total Cost	$/GB
50	$326,388	$100,568	$91,724	$82,880	$73,164	$674,723	$0.0135

For the other archives – 250PB and 500PB – the Spectra TFinity library is recommended.

Capacity (PB)	Year 1 – Cost	Year 2 – Cost	Year 3 – Cost	Year 4 – Cost	Year 5 – Cost	Total Cost	$/GB
250	$910,315	$423,939	$365,703	$323,184	$297,639	$2,320,781	$0.0093
500	$1,690,864	$821,292	$720,537	$652,473	$551,718	$4,436,884	$0.0089

In a 250PB archive stored over five years, a 3-frame, 24-drive TFinity library with one Spectra OSW-2400 Optical SAS Switch is recommended.

The unit will grow by 50 PB per year, with media purchases of 12.5 PB quarterly. Throughout the five years, the TFinity will be expanded to nine frames. The costs for these two examples are shown in the table below.

For an organization that will need 500 PB stored over five years, a similar configuration of TFinity would be used. The library would initially operate as a five-frame system, featuring 48 LTO-10 tape drives and two Spectra OSW-2400 Optical SAS Switches. 25PB of media will be added quarterly throughout the five years, and, as with all these financial projections, yearly support is also included. In year two, three additional media expansion frames will be added. In year three, another two media expansion frames will be added. Year four requires three additional media expansion frames. In year five, the final two media expansion frames will be added for a completed TFinity with 15 frames total, holding 500 PB of uncompressed data.

With this library size, it costs $0.0089 per GB to store 500 PB of data. A second copy of the media will only add another $2.94 million, further driving down the cost per GB.

Capacity (PB)	Year 1 – Cost	Year 2 – Cost	Year 3 – Cost	Year 4 – Cost	Year 5 – Cost	Total Cost	$/GB
250	$910,315	$423,939	$365,703	$323,184	$297,639	$2,320,781	$0.0093
500	$1,690,864	$821,292	$720,537	$652,473	$551,718	$4,436,884	$0.0089

For a quick comparison, a consolidated table showing costs for all configurations is listed below:

Capacity (PB)	Year 1 – Cost	Year 2 – Cost	Year 3 – Cost	Year 4 – Cost	Year 5 – Cost	Total Cost	$/GB
50	$326,388	$100,568	$91,724	$82,880	$73,164	$674,723	$0.0135
250	$910,315	$423,939	$365,703	$323,184	$297,639	$2,320,781	$0.0093
500	$1,690,864	$821,292	$720,537	$652,473	$551,718	$4,436,884	$0.0089
1,000	$2,530,601	$1,562,492	$1,397,441	$1,228,618	$1,059,795	$7,778,946	$0.0078

What Does it Take to Move an Exabyte of Data?

Another consideration is the performance or bandwidth required of an organization's wide area network (WAN) to move data to the cloud. Cost will vary tremendously depending on the geographic location and type of service. Likewise, an organization's use of a WAN will extend beyond cloud access alone. While it's challenging to determine the average bandwidth cost for transferring data to the cloud, evaluating the bandwidth required to archive an exabyte to the cloud is straightforward.

Connection	Speed	TB Moved Per Day	Time to Move 1 PB	Time to Move 1 EB
Gigabit Ethernet	1 GBs	10.8	92 days	252 years
OC48	2.5 GBs	27	37 days	101 years
OC96	4.976 GBs	53.74	18 days	49 years
OC192	9.600 GBs	103.68	9.6 days	26 years

The egress charges paid to AWS for data retrieval do not include the pipes that customers will use for restoration, nor do they guarantee how long it will take to receive the data back. The service level agreement (SLA) only specifies the timeframe during which the customer can access their data to initiate the restoration.

Various types of internet connections offer varying speeds. This chart illustrates the performance of dedicated connections for the specified connection type, as well as the time required to transfer terabytes, petabytes, and exabytes of data at full dedicated performance.

The bandwidth numbers above are based on a best-case scenario, assuming 100 percent of available bandwidth is utilized. Therefore, any other organizational use must be suspended to achieve these numbers. Even restoring a single petabyte, 1/1000th of an exabyte, could pose serious challenges to an organization's SLA.

A tape-based, on-premises archive offers rapid access and restoration times. Access time to data can be reduced from hours to minutes. More importantly, at 400MB/s uncompressed, LTO-10 restoration offers comparably rapid restoration. The table below shows the speed of a single LTO-10 and the aggregate speed of the 48 LTO-10 drives used in our exabyte archive example above. The OC192 WAN speed is repeated for reference.

	Speed	TB Moved Per Day	Time to Move 1 PB	Time to Move 1 EB
Single LTO-9 Drive (Uncompressed)	400 MBs	34.56	34,056 days	252 years
48 LTO-9 Drives (Uncompressed)	19.2 GBs (aggregate)	1,659	14.3 hours	1.6 years
OC192	9.600 GBs	103.68	9.6 days	26 years

Even the fastest internet connections pale compared to the aggregate speed of LTO-10 tape drives. The aggregate speed of 48 LTO-10 drives is over 16 times faster than an OC-192 connection. Remember that up to 168 drives can be configured in a TFinity, resulting in an aggregate performance of over 67 GB/s.

Cost and performance will prevent the archiving of an exabyte of data in the cloud. But these numbers are essential to consider. Few organizations "start" with an exabyte of data, but if you look at the current data growth combined with longer retention rates, many organizations have already reached this milestone.

The Promise of Hybrid Cloud Revisited

This section started with a brief description of Hybrid Cloud. As noted above, cost and performance limitations hinder the archiving of exabytes of data to the cloud. Yet, the cloud offers services that an on-premises archive might lack. The cloud can be beneficial for distributing or sharing data. The cloud transcoding services ensure shared data can be consumed on multiple, disparate platforms.

The goal of data users and storage manufacturers alike should be to maximize the benefits of both public cloud and on-premises storage. How can we maintain low-resolution copies of content in the cloud for editing, while keeping the full-resolution copies on-premises for lower-cost storage of the higher-resolution files? How can we stream research data to tape for indefinite retention on-premises (creating petabytes and exabytes of data), yet also have accessibility via the cloud? At the same time, it's still relevant to researchers across the globe.

Arguments over "end-point" storage solutions – disk vs. tape, public cloud vs. private cloud, file vs. object – have consumed too much of the storage conversation and have deterred organizations from being able to focus on the real point behind storage – meeting the desired organizational goals that the information/data/content is used for.

If it's true that we will all use the cloud in some fashion, and we believe that it is possible, we need a more consolidated approach to working within, as well as through and even around, the cloud. Would it be possible to access data via the cloud, yet store the data on-premises? Applications are underway, and we believe it will be a short time before these challenges are met.

Moving towards a more "geographically independent" storage model will allow users to focus on using the most appropriate storage once again for a given role without creating independent storage silos which kill efficiencies and waste budget dollars.

In this manner, exabyte-scale archives will become something to aspire to due to their ability to support our research, manufacturing, discovery, medicine, industry, and overall quality of life. Stay tuned!

LTO into the Future

In May 2025, IBM and the LTO Consortium announced the LTO-10 full-height drive, which matches their latest Enterprise drive, the IBM® TS1170, in uncompressed throughput. More significantly, the LTO roadmap is published through LTO-14.

This indicates that the LTO technology in our exabyte-capacity library could increase four times to 8 exabytes over the 10-year period we discussed. As LTO-1 has grown to LTO-10, the growth to LTO-14 is readily possible with tweaks and improvements to the existing technology, portending a bright future.

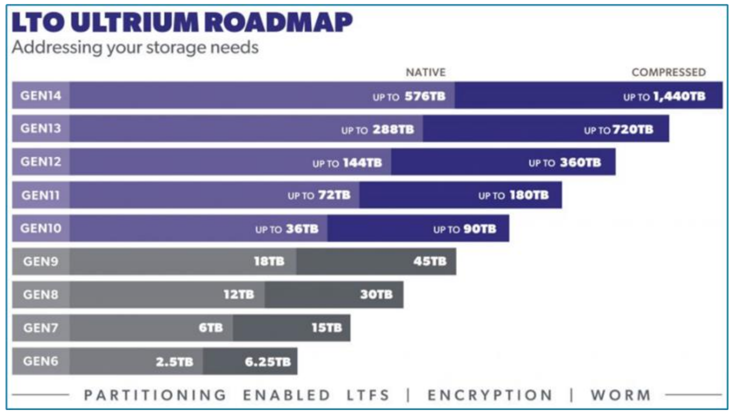

The current LTO technology roadmap is valid through May 27, 2025. An updated roadmap reflecting the next generation of LTO advancements will be published following the industry announcement on that date.

Summary

In more than 40 years of creating, manufacturing and delivering storage solutions, Spectra Logic has seen many of the industry's "ups and downs." Often, one new technology threatens an older technology or requires a yet-to-be-invented technology to reach its promise. It's usually tricky to determine where the market is going or how best to balance organizational mandates with technology and budget.

The one constant in storage is the need for more. That doesn't mean that budgets must be broken or organizational goals must be missed.

Using the right technology, ExaScale archives can be achieved at reasonable costs, often with a very high return on investment. Hopefully, this white paper demonstrates that the TFinity Plus archives are designed to support organizations' pursuits, whether in science, time-to-market, health, prosperity, quality of life, education, or entertainment. We have the technology today to build what we will need tomorrow. Such is the nature of the archive.

Toll Free: 800-833-1132 · 303-449-6400 · 6285 Lookout Road · Boulder, CO 80301 USA
V2-060625