SPECTRA

Digital Data Storage Outlook 2020

2020

800-833-1132 • 303-449-6400 • 6285 Lookout Road • Boulder, CO 80301 USA SPECTRALOGIC.COM

CONTENTS

EXECUTIVE SUMMARY

"Action without vision is only passing time; vision without action is merely daydreaming; but vision with action can change the world." ***- Nelson Mandela***

This fifth annual Data Storage Outlook report from Spectra Logic explores the ways in which the world manages, accesses, uses and preserves its ever-growing data repositories. It also covers the strategies and technologies that are anticipated to make impacts on our future. As Nelson Mandela wisely says in the quote above, action alone and vision alone have no effect, but vision coupled with action can change the world. This combination will be the key to meeting today's data storage and data management challenges and will be vital in driving the technological leaps that will be required by the generations to come.

With that said, the year 2020 started as no other in modern history. The Covid-19 pandemic moved through the world at an astonishing pace, forcing countries to shutter borders, businesses, schools and homes. As the world sheltered in place, many organizations quickly pivoted, becoming almost 100 percent virtual, while other industries were forced to contract to safeguard the health of families, neighbors, colleagues, friends and strangers. As of this writing, the coronavirus war is still being waged, leaving many industries in flux and others poised to mount unrivalled recoveries.

Understanding the drivers that impact us today will help us take the steps necessary to improve our world tomorrow. But, accurately predicting the future, especially in light of today's extraordinary circumstances, can be more challenging than picking the winning horse at the Kentucky Derby. And, unlike a horse race, the stakes are much higher. Data storage and media manufacturers, application developers, and all sizable users of data must anticipate future technologies, applications, use cases and costs because accurate planning today will deliver a much more promising future.

The Size of the Digital Universe

- A 2018 IDC report, commissioned by Seagate Technology, predicts that the Global Datasphere will grow to 175 zettabytes (ZB) of digital data by 2025.[1] While this report took a top-down appraisal of the creation of all digital content, Spectra projects that much of this data will never be stored or will be retained for only a brief time. Data stored for longer retention, furthermore, is frequently compressed. As also noted in the report, the "stored" digital universe is therefore a smaller subset of the entire digital universe as projected by IDC.

The Storage Gap

- While there will be great demand and some constraints in budgets and infrastructure, Spectra's projections show a small likelihood of a constrained supply of storage to meet the needs of the digital universe through 2030. Expected advances in storage technologies, however, need to occur during this timeframe. Lack of advances in a particular technology, such as magnetic disk, will necessitate greater use of other storage mediums such as flash and tape.

Storage Apportionment and Tiering

- Economic concerns will increasingly push infrequently accessed data onto lower cost media tiers. Just as water seeks its own level, data will seek its proper mix of access time and storage cost.
- Spectra envisions a logical two-tier architecture comprised of multiple storage types. We further envision that the first logical tier's storage requirements will be satisfied entirely through solid-state storage technologies while the second tier requirements will be satisfied by magnetic disk and tape deployed as object storage either on-premise or in the cloud.

Highlights from 2020 Report

- 2019 saw the third generation of 3D XPoint technology shipped under the Intel Optane™ product family name. Unlike the previous two generations, that were packaged as NVMe solid-state disk (SSD), this generation resides directly on the memory bus. With performance and cost characteristics that sit between DRAM and SSD, this technology defines an entirely new tier of storage.
- After experiencing 18 months of oversupply of flash in the market, resulting in substantial price reductions, 2020 will see reductions in supply versus demand of pricing such that there will be a 10% to 40% price increase over the year.
- As the magnetic disk manufacturers continue to experience delays in shipments of advanced technology drives, based on heat-assisted magnetic recording, they have extended the technology roadmap for perpendicular magnetic recording as a means to continue increasing drive capacities.
- For the first half of 2019, a legal issue between Fujifilm and Sony prevented Sony from shipping LTO-8 media, which kept the price of LTO-8 media high. This issue was resolved in August, but even prior to that, tape customers had the option of purchasing LTO-7 barium ferrite media that, when initialized in an LTO-8 drive, provided an immediate 50% improvement in tape capacity.

OVERVIEW

In 2019, Spectra reached a historic corporate milestone and celebrated 40 years of business success. Spectra's Founder and CEO Nathan Thompson continues to lead the company in its goal to provide innovative and complete storage solutions for organizations around the world that want to protect and preserve their digital assets indefinitely. Demonstrating continuity and longevity, the firm was an early innovator in robotics, automation and interfaces. Today the company delivers a broad range of data storage and data management solutions to the global market, including disk, object storage, tape and software for storage lifecycle management.

Spectra takes a technology and media-agnostic view of storage, and its strong partner and customer relationships provide a clear glimpse at the company's future direction. As a storage system developer and manufacturer the firm cannot ignore costs and trends in data storage performance and pricing.

This is the fifth edition of an annual overview of trends and predictions in storage technology development and availability. At the time of this writing the world was entering the coronavirus-19 pandemic. The pandemic has resulted in a very high level of uncertainty in all markets with the storage market not being an exception. For this reason, we are predicting lower capacity storage growth for 2020. Given the average capacity of devices of all storage types are increasing this will result in lower volumes of storage devices than in 2019.

THE NEXT STORAGE ARCHITECTURE

Increasing scale, level of collaboration, and diversity of workflows are driving users toward a new model for data storage. The traditional file-based storage interface is well suited to in-progress work but breaks down at web scale. Object storage, on the other hand, is built for scale. Rather than attempting to force all storage into a single model, a sensible combination of both is the best approach.

File vs. Object Storage

File systems are called on to serve many purposes, ranging from scratch storage to long-term archival. Like a jack of all trades, they are a master of none, and storage workflows are exceeding the capabilities of the traditional file system. The file system interface includes a diverse range of capabilities. For example, an application may write to any file at any location. As this capability expanded to network file systems (NFS, SMB), the complexity scaled up as well – for instance, allowing multiple writers to any location within a file.

The capabilities of the file system interface make it excellent for data that is being ingested, processed or transformed. As a user creates content or modifies something they're working on, the application may quickly hop around in its data files and update them accordingly. It must do this with enough performance that the user's creative process is not interrupted, and also with sufficient safety that the user's data will be intact in the event of malfunction. The file system is the user's critical working space.

Object storage is simply another way of saying "the web." From its beginning, the web's HyperText Transfer Protocol (HTTP) was a simple method of sending an object over the public internet, whether that object was a web page, image or dynamically-generated content. Any web browser is a basic "object storage client." HTTP has methods for getting and putting whole objects but lacks the notion of interactive, random I/O.

This simplicity, however, is a powerful enabler for object storage to operate at scale. Every object has a Uniform Resource Identifier (URI) which enables that object to be addressed -- whether it's on a server in the next room or a data logger in the remote wilderness. It doesn't matter if the network topology or storage system is involved, or whether it is traversing multiple caches and firewalls. Objects may be migrated to different storage media or even moved from a company's data center into a public cloud; so long as the URI remains unchanged, users will neither know nor care.

The cloud grew out of the web, so it's no surprise that cloud is primarily based on object storage. The first product of Amazon Web Services (AWS) —pre-dating their compute offerings—was the Simple Storage Service (S3) released in 2006. The S3 protocol is simply HTTP with minimal additions. S3 includes methods for retrieving a range of an object, or sending an object in multiple parts, but in general it maintains a very simple, high-level interface. AWS has released other storage services, including a parallel file system, but S3 remains the backbone of their cloud.

The dramatic contrast between file system and object system capabilities means that the ideal storage interface is both. The data-intensive company should use a combination of systems to fully capitalize on the strengths of each.

Comparison of File vs. Object System Properties

Feature	File System	Object System
Connection	Direct-attach or local network/VPN	VPN or public internet
Standardization	POSIX, Windows	Lacking; AWS S3 popular
Read/Write Mix	Arbitrary read/write	Write-only/read-many
Data Mutability	Update any file in any place	Objects immutable; make new version
App compatibility	Broad	Limited; new applications only
Network/technology independent	No	Yes
Transparent storage class migration	No	Yes
Versioned, auditable	No	Yes

New Storage Tiers

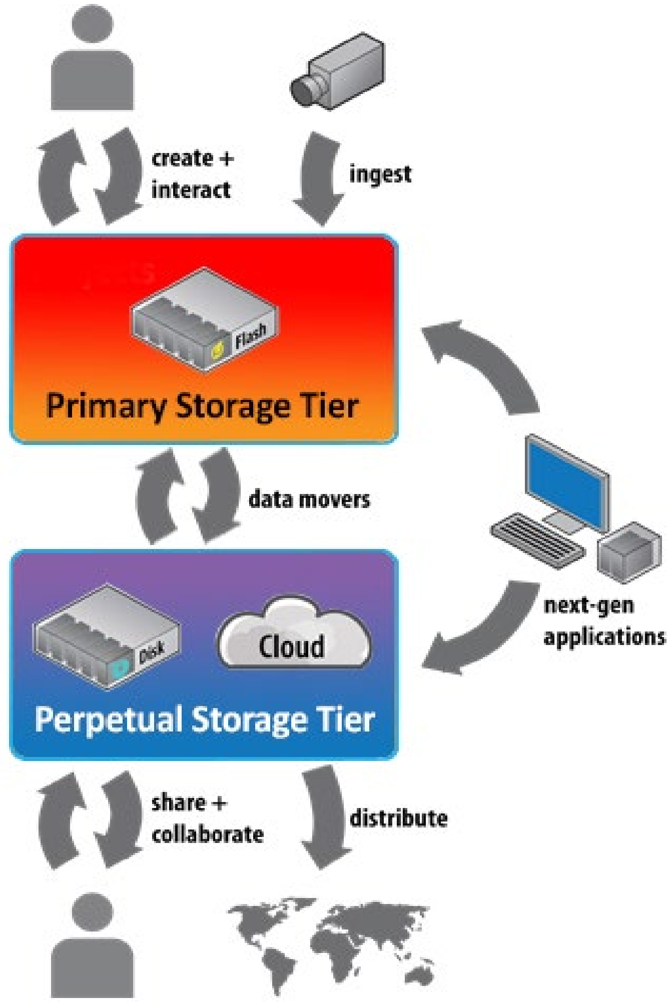

Figure 1: New Storage Tiers

In the past, data storage usage was defined by the technology leveraged to protect data using a pyramid structure, with the top of the pyramid designated for solid-state disk to store 'hot' data,' SATA disk drives used to store 'warm' data and tape used for the bottom of the pyramid to archive 'cold' data. Today, Spectra proposes a new two-tier architecture to replace the aging pyramid model.

The new two-tier paradigm focuses on the usage of the data rather than the technology. The two-tier paradigm combines a Primary or Project Tier where in-progress data resides, which is file-based, and a second or Perpetual Tier where finished and less frequently changed data resides, which is object-based. Data moves seamlessly between the two tiers as data is manipulated, analyzed, shared and protected.

The Project Tier

- **Data ingest, where raw data streams need to be captured rapidly**. For example, a media production may need to capture camera streams, audio streams and timecode simultaneously. Data will be raw, uncompressed, and require extremely high bandwidth. These streams may be stored to separate devices (e.g., flash cards within each camera) or captured on a central system (RAID box or filer).
- **Work-in-progress, where a user may hop around and edit content in any location**. This may include edit-in-place such as some image editing applications, where a user may work across the X/Y image plane and multiple layers. It may also include non-destructive applications, where a change stream is captured but underlying data is never changed. Regardless of technique, the application must respond instantly to user input.
- **Computation scratch space, where the volume of data exceeds RAM and/or checkpoints are saved to stable storage**. Most of it will be discarded after the job is complete; only the results will live on. Storage must have high bandwidth, as time spent waiting for a checkpoint to finish is wasted.

The file system's ability to write to any location within a file is critical for capturing data as it happens. Some applications use the file system interface directly (open a file handle and write to it) while others use software libraries such as SQLite or HDF5 to write structured data in a crash-consistent manner.

But what happens when the user is done with their edits, and the dynamically-changing data becomes static? It moves to the Perpetual Tier.

The Perpetual Tier

- **Project assets that must be shared across a team so they can be the basis for future work**. Video footage leaving production and going into post-production may need to be used by teams of editors, visual effects, audio editing, music scoring, color grading, and more. These teams may be spread across geographic regions and the object store may pre-stage copies in each region. Each source asset will have a globally-resolvable name and data integrity hash code. These are never modified. In some cases, new versions may be created, but the prior versions will be kept as well. The lifetime of raw assets is effectively forever—they are the studio's lifeblood—and they may be migrated across storage technologies many times.
- **Completed work that must be distributed**. Object storage, along with public cloud providers, offer an ideal way to distribute data to end users across the globe. A finished media production may result in a variety of distribution files, along with a descriptive manifest, for example, MPEG-Dash as used by YouTube and Netflix. Because the objects are static, they may be cached in global content delivery networks.
- **Finished computational results to be shared across researchers**. Encryption and access controls, such as those provided in the S3/HTTP protocol, allow for sharing of sensitive data across the public internet. Storage cost may prompt users to apply a cold-storage-only policy to older data, knowing that it can be restored later if needed.

Data moves between the project and perpetual tiers in both directions. Users may migrate from Project Tier to Perpetual Tier once files are complete, but migration may go the other way as well. A visual effects company may start from source files that exist in object storage in a public cloud, staging those to their Project Tier when they start work. Afterward, the finished shots are copied back to the cloud.

Whereas the software applications of the past used file systems only, next-generation applications will support both tiers directly. They will use a file system for their workspace and object storage (including cloud) as the source and destination for more permanent data. Additionally, some software libraries are supporting object storage natively, for example, there is a HDF5 library that can use a S3-compatible object store directly.

Data Movers

Until applications can natively utilize both the project and perpetual tier data movers will be required in order to move data to and from between the two tiers. Based on customers varying requirements will require different types of data movers. Some customers may want the ability to move large amounts of project data over to the Perpetual Tier once a project is completed. This serves two purposes in that it frees up the Project Tier for new projects and it archives the data of the project making it available for future processing. Another set of customers may want to selectively prune the Project Tier of files that have not been accessed for a long period time. This frees up Project Tier storage such that expansion of that storage is required. Another customer may use the

Perpetual Tier as a means to distribute data globally to multiple groups working on the same project. A data mover that allows users to “check out” a project by moving the data from the Perpetual Tier to a local Project Tier. Once changes to the Project Tier are made they can be “checked in” back to the Perpetual Tier, thereby making those changes available to all sites.

STORAGE TECHNOLOGIES

Serving a market sector of more than $50 billion[2], the storage device industry has exhibited constant innovation and improvement. This section discusses current technologies and technical advances occurring in the areas of persistent memory, flash, magnetic disk, magnetic tape, optical disc and future technologies, as well as Spectra’s view of what portion of the stored digital universe each will serve.

Persistent Memory

This year we have added the persistent storage tier that is defined by the two characteristics of persistence across power outages and performance close enough to DRAM to exist on a memory bus. Though various products, such as battery-backed DRAM, have been available for many years, they have always served a niche market. With the release in 2019 of Intel’s Optane™ 128GB and 256GB persistent memory modules, we believe this situation has changed resulting in an entirely new category of storage. The Optane products are the third generation from Intel whereby the first two generations were only released as NVMe SSDs. Optane technology differs from flash in that it has longer wear and lower latency, allowing it to be placed on the memory bus. This technology has the potential to being highly disruptive to the DRAM marketplace. This market of $80 billion is dominated primarily by three players: Samsung, SK Hynix and Micron. Intel does not participate in this market, which means that any disruptive gains it makes is net new business. Below we see a performance comparison of Optane to that of standard DRAM. As can be seen, the DRAM wins every performance category by many multiples. However, performance is not the only factor that matters as Optane technology wins in the categories of cost, density, power consumption and, of course, persistence. For many applications the performance tradeoffs are worthwhile given these other factors.

Latency	Optane DIMM	DRAM
Idle Sequential Read Latency	~170ns	~75ns
Idle Random Read Latency	~320ns	~80ns
Per DIMM Bandwidths		
Sequential Read	~7.6 GB/s	~15 GB/s
Random Read	~2.4 GB/s	~15 GB/s
Sequential Write	~2.3 GB/s	~15 GB/s
Random Write	~0.5 GB/s	~15 GB/s

Two application spaces that can derive almost immediate benefit from this technology are in-memory databases and large shared caches. In-memory databases have become more popular over the last few years as they provide the lowest latency with highest transaction rate as compared to databases running out of other storage mediums. However, these databases need as much DRAM in the system as the size of the database. This is both costly and, in some cases, just not possible given the amount of memory that can be placed into a server. Another issue is that, given the non-persistent nature of DRAM, the database needs to be “checkpointed” on a regular basis -- typically to an SSD. This is required because, if power is lost, the database can be restored to a good state. The Optane technology solves these problems with large formats, such as 256GB, which enables larger databases to be supported along with being persistent so that checkpointing is not required. For example, a database that has been checkpointed may require up to a half hour to be restored, while the same database running on Optane could be brought up in less than half a minute.

Another application that can be easily moved onto the Optane technology is that of cluster-wide caches such as Memcached. These caches hold items in-memory on each node of a cluster such that a vast majority of requests that come into the cluster can be serviced without needing to go to the backend storage. For example, when a user logs into their account in a web application, their home page gets cached and is read one time from the backend. As the user moves around the application, new information that is brought into the application is additionally cached. Optane technology is an excellent fit for this application as its high capacity allows for millions of items to be cached. An additional benefit is that the cache can be restored quickly after a power cycle. Besides these easy-to-convert applications, many new applications will have this technology designed into them.

Given the disruptive nature of this technology, some of the largest DRAM producers are working on alternative solutions. Samsung, for instance, has plans to announce products based on Magnetoresistive random access memory (MRAM) that has many similar properties to that of Optane. SK Hynix is working on technology similar to Optane while Micron co-developed Optane with Intel and therefore has technology based on it.

Flash

The fastest growing technology in the storage market is NAND flash. It has capabilities of durability and speed that find favor in both the consumer and enterprise segments. In the consumer space it has become the de facto technology for digital cameras, smart phones, tablets, laptops and desktop computers. As previously discussed, we predict that the Project Tier will be comprised of solid-state storage technologies.

Previous versions of this paper highlighted the flash vendors transition from planar (i.e., 2D) to 3D-Nand manufacturing. This transition was required in order to provide flash with a roadmap whereby increased capacities could be achieved for many years to come. During the time of this transition, those being the years 2016 through 2017, increases in flash capacity were suppressed resulting in relatively small declines in price. In order to achieve these increases, substantial investment in existing and new flash wafer fabrication facilities as well as time was required for the 3D manufacturing process to achieve good yields.

In fact, this investment, over the past three years, along with what was projected through 2019, is roughly $100 billion. With this transition almost complete, solid-state disk pricing decreased dramatically by almost 50 percent during 2019. It appears that in 2020 this trend will reverse itself and there will be a shortage of flash that will drive prices higher. Though investment is shrinking the price of NAND started rising toward the end of 2019, and it is predicted to rise another 10% to 40% in 2020. This seems to be contrary to the capacity coming online. True, the demand for flash is growing; however, it's not growing faster than it did in prior years. Given that the capacity improvements of the technology are roughly following Moore's law (a doubling each 18 months), it seems that the price increases are being driven by something else. One point of view is that these increases are being driven by the vendors exhibiting greater pricing discipline. All vendors have invested heavily into fabrication facilities that now must show good returns. A race to the bottom in terms of price would make that difficult.

There are five companies that own and run NAND fabrication lines: Samsung, Toshiba/Western Digital, Micron Technology SK Hunix and Yangtze Memory Technologies. According to a roadmap produced by TechInsights, all 2D flash is essentially reaching end of life with no future products planned. Next, every vendor is projecting 128 layer chips arriving to market between mid-2019 to mid-2020. They also are projecting 192 layer parts starting at the end of 2020. Besides adding more layers, there are two other aspects of 3D flash that can provide greater capacity. The first is adding more voltage detection levels inside each cell. With the first flash chips produced, each cell was either charged or not charged; meaning that each represented a single binary bit referred to as single level cell (SLC). This was followed by detection of four levels, with two bits of information per cell referred to as multiple level cell (MLC). Later, a triple level cell (TLC) holding three bits of information per cell was produced. In 2018 the first QLC parts were shipped as consumer technology, with each cell holding four bits of information. Currently QLC is prevalent in consumer quality devices while TLC is used in enterprise. The final method to increase capacity of flash is to decrease the cell size. This results in reducing signal integrity, however, making it more difficult to detect voltage levels while reducing the number of bits that can be detected per cell. Given the flash roadmap, it appears that 19 Nano-meters is as small as the industry plans on producing. Given it is already at 20 nm, it doesn't appear that this method will yield much capacity gain.

According to the roadmap, all vendors have announced availability of 128-layer parts in 2020 except for Yangtze Memory Technologies (YMTC) who is just in the process of delivering their first 64-layer parts. Of these, Samsung has actually delivered the first product to market consisting of single-stack 136-layer part. This is in comparison to some other vendors who plan on using "string stacking" in order to meet the 128-layer goal. String stacking is a technique where multiple chips of some number of single-stacked parts are "glued" together to create a module of the desired number of layers. For example, a vendor could take four 32-layer parts and combine them to create a 128-layer part. Both techniques have their advantages and disadvantages. To do a single stack part exceeding 100 layers takes more upgrade investment in the fabrication facility then adding string stacking to a facility already producing a 64-layer part. Also, the initial yield of the 100-plus layer part is certainly going to be much lower than that of the already established production of the lower layer part. On the other hand, the higher layer level part can win on cost when the overall manufacturing cost is reduced to a level that makes it cheaper than manufacturing multiple parts. The higher layer part also has the capability to expand into the future roadmap more easily.

For example, Samsung has indicated that they will be string stacking three 136-layer parts together to produce a 400-layer part in the next few years. To do so with a 32-layer part would require 12 or 13 parts. There are complex issues with building 100-plus layer parts. For instance, just the alignment of so many layers is problematic. For this reason and others, there are no vendors talking about building past 128-layers in a single-stack part. So, we predict that future capacity gains will be primarily achieved by string stacking parts together.

The question will be what cost advantages will this yield longer term? The other question besides density is what advantages will higher layer string stack parts have over individual parts of smaller layers? For these reasons we are projecting that the price decreases in flash will slow and be more a function of the yield and manufacturing cost improvements that vendors are able to achieve.

According to the flash roadmap, the two technologies that require further explanation are Z-NAND from Samsung and Xpoint from Intel/Micron. Both products offer lower latency (though not necessarily higher bandwidth) than all other flash products on the market. They also offer better wear characteristics making them suitable for caching applications. It is believed that Z-NAND is a combination of single level cell (SLC) flash along with improvements in the architecture of the companion flash controller. It appears that Samsung created this technology to hold off Xpoint from moving up in the flash space. As described previously, XPoint is a completely different technology and is now mainly positioned to compete in the persistent memory segment.

A flash competitor that warrants special attention is the last one on the chart -- Yangtze Memory Technologies (YMTC). It is a China-based company that is supported by the Chinese government as flash technology is considered a key technology area. They plan on producing 60,000 64-layer wafers a month starting in 2020. They will most likely be using string stacking to move to 128-layer in the first half of 2021. Given what has happened with Chinese intrusion into other markets, there is great concern as to the impact in this market. Trendforce’s comment on this subject is probably the most telling when they declared: "YMTC's impact on the future market is inevitable and unstoppable."

Along with flash capacities, strides have also been made in flash controllers. A controller along with some number and capacity of flash chips packaged together comprise a solid-state drive. A new generation of enterprise drives is capable of sequential reading and writing at greater than 3 GB per second. They can support more than 500,000 random small block reads and more than 250,000 random writes. These drives now sell for about $220 per terabyte. Consumer drives do not perform as well and have price points in the less than $110 per terabyte range.

Flash requires the least amount of physical space per capacity of all the storage technologies. Much hype has been made regarding which technology has the largest capacity in a 2.5-inch disk form factor, with some vendors announcing capacities of up to 100TB. Those statements are misleading. The only advantage of such devices is that the cost of the controllers can be amortized over a greater amount of flash and fewer slots are required in a chassis. But, both of these costs are trivial compared to the cost of the flash. The disadvantage of large flash disks is that one controller needs to service a large capacity. A better approach is to maintain the ratio of one controller to a reasonable amount of flash. In order to address this issue, new smaller form factors have been created. These form factors allow for plugging many flash disks into a small chassis rack. These emerging systems provide high capacity, high bandwidth and low latency all housed in a 1U chassis. These new form factors are a break from flash being backward compatible with chassis designed for magnetic disk systems, and more importantly, they represent the end of magnetic disk being considered a primary storage device.

Along with being constrained by the physical interface, flash drives were also constrained by the electrical interface -- either SAS or SATA. These interfaces added latency to data transfers that, with magnetic disk, were “in the noise”, but with flash, became major contributors to overall performance. For this reason, the industry has moved to an interface that directly connects flash drives to the PCIe bus. The NVMe interface can be thought of as the combination of non-volatile memory (NVM) and PCIe (e). Though the specifications for the NVMe have been around for several years, it has only been over the last few years that adoption has started to take off. This lag was primarily caused by the desire to maintain backward compatibility with magnetic disk systems' chassis. As a proof point that NVMe is now the predominant interface for enterprise flash, many of the current generation of enterprise flash drives only support NVMe and do not offer SAS versions. NVMe is following a parallel path and labelling as PCIe. For example, NVMe3 indicates NVMe running over a PCIe Gen 3 bus while NVMe4 indicates NVMe running on a PCIe Gen 4 bus.

NVMe is great for moving data from a processor to flash drives inside a chassis, but in order to fully displace SAS, required the creation of a technology that allowed for a box of flash drives, referred to as a JBOF (just a bunch of flash) to be accessible by one or more controllers. This needed to be done without adding substantial latency on top of the NVMe protocol. The technology developed is referred to as "NVMe over fabrics" (NVMe-oF). The fabric can be PCIe (e), infiniband, SAS or Fibre Channel, but for new systems, it will predominantly be remote direct memory access (RDMA) over converged Ethernet (RoCE). With this latter technology, the physical connection between the controllers and the JBOF is commodity Ethernet. RoCE technology is becoming a commodity both at the chip and HBA level. RoCE technology will find rapid adoption for all interconnections that require high bandwidth and low latency. This includes interconnections between clients and block or file controllers, interconnections between those controllers and the shared storage JBOFs, and the interconnections between cluster members in scale-out storage solutions. Most JBOF chassis run a x4 (four lane) PCIe Gen3 connection to each flash drive. Since each lane is capable of about 1 GB/s, an individual drive is restricted to reading or writing at 4 GB/s. Currently enterprise SSDs are designed with this limitation in mind; however, going forward, with the advent of more performance per flash chip as a result of more layers, and PCIe Gen4 shipping in greater volume, we expect that SSDs exceeding 7 GB/s sequential reading and writing will be available in the 2020-2021 time-frame.

Along with the physical interface changes described above, work is being done on developing new logical interfaces. Samsung, for example, has announced a key/value store interface. This interface is appropriate for object level systems whereby a system can take incoming objects and create some number of fixed size hunks (i.e., values) by either splitting larger objects or packing together smaller ones. These then could be distributed out to a group of SSDs that have this key/value interface thereby reducing the complexity of the application. Another Samsung initiative involves virtualization such that a single SCSI SSD can be split into 64 virtual SSDs. This is useful in virtual processing environments where multiple virtual machines are running with all wanting access to some amount of independent storage from a single SSD. In this case, the hypervisor no longer needs to perform memory mapping between the virtual machine and SSD. Western Digital is taking a different approach by providing a logical interface that allows an application to manage the underlying zones (i.e., flash blocks) of an SSD directly. This effort is called the "Zone Storage Initiative" and applies to all "zone" storage device types which include flash and shingled magnetic recording (SMR) disk. Regardless of the media type, a "zoned" storage device is one in which the storage is broken into equal size areas (i.e., zones) with properties that allow them to be written sequentially without being arbitrarily overwritten.

This interface moves much of the burden of data placement and garbage collection into the application instead of the SSD controller. This allows for better use of the flash capacity and performance at the cost of adding substantial complexity to the application. For this reason, this interface will be limited in use to cloud companies and storage system providers. More on this interface will be discussed in the magnetic disk section of this paper.

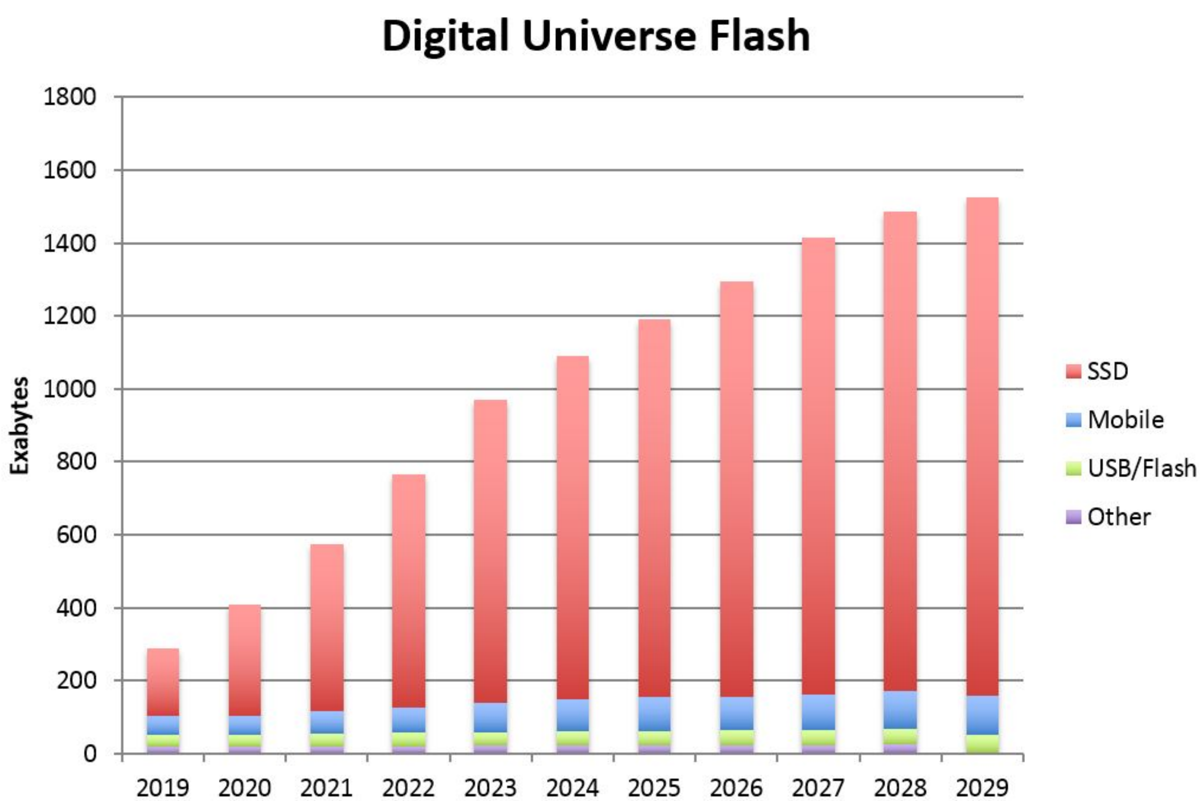

Figure 2: Stored Digital Universe Solid-State Technology

New market segments in gaming and automotive are increasing the demand for flash as they are both essentially brand-new markets for the technology. The enterprise SSD market continues to be a growth segment as IT shops continue to make the move from magnetic to solid-state drives. The lack of performance improvements of magnetic disk will drive cloud companies to spend a greater amount of their storage budgets on flash-based technology. Consumer demands for increased capacities for phones, cameras and laptops is waning and being replaced by demands for lower cost for the same capacities.

Magnetic Disk

For many years now the disk drive industry has had three major providers: Seagate, Western Digital and Toshiba. It is interesting that two of those suppliers, Western Digital and Toshiba, also share flash fabrication facilities and, as such, are not as exposed to disk displacement by flash. However, for Seagate, it is essential that the disk segment continues to be healthy. Looking at the disk drive volume shipments from the last year through 3Q 2019, we see the volumes shipped over the last four quarters to be 328 million compared to 392 million for the prior year's four quarters, or about a 16 percent drop in volume. All consumer and 2.5-inch high performance categories were down. More recently, the major game console manufacturers have introduced their next-generation products that all use NAND flash rather than small hard drives. We expect this will accelerate the demise of this category of disk drives over the next few years. Accepting this, the disk manufacturers have been disinvesting in research and development of these drives as can be seen by the lack of any capacity improvements over several years. The segment that did see year-to-year increases in both capacity and volume shipments is the 3.5-inch nearline drive category. It now comprises over 50% of all disk revenue. Developing a singular product, with a few variations, has allowed the disk companies to focus their resources enabling them to remain profitable even as a good portion of their legacy business erodes.

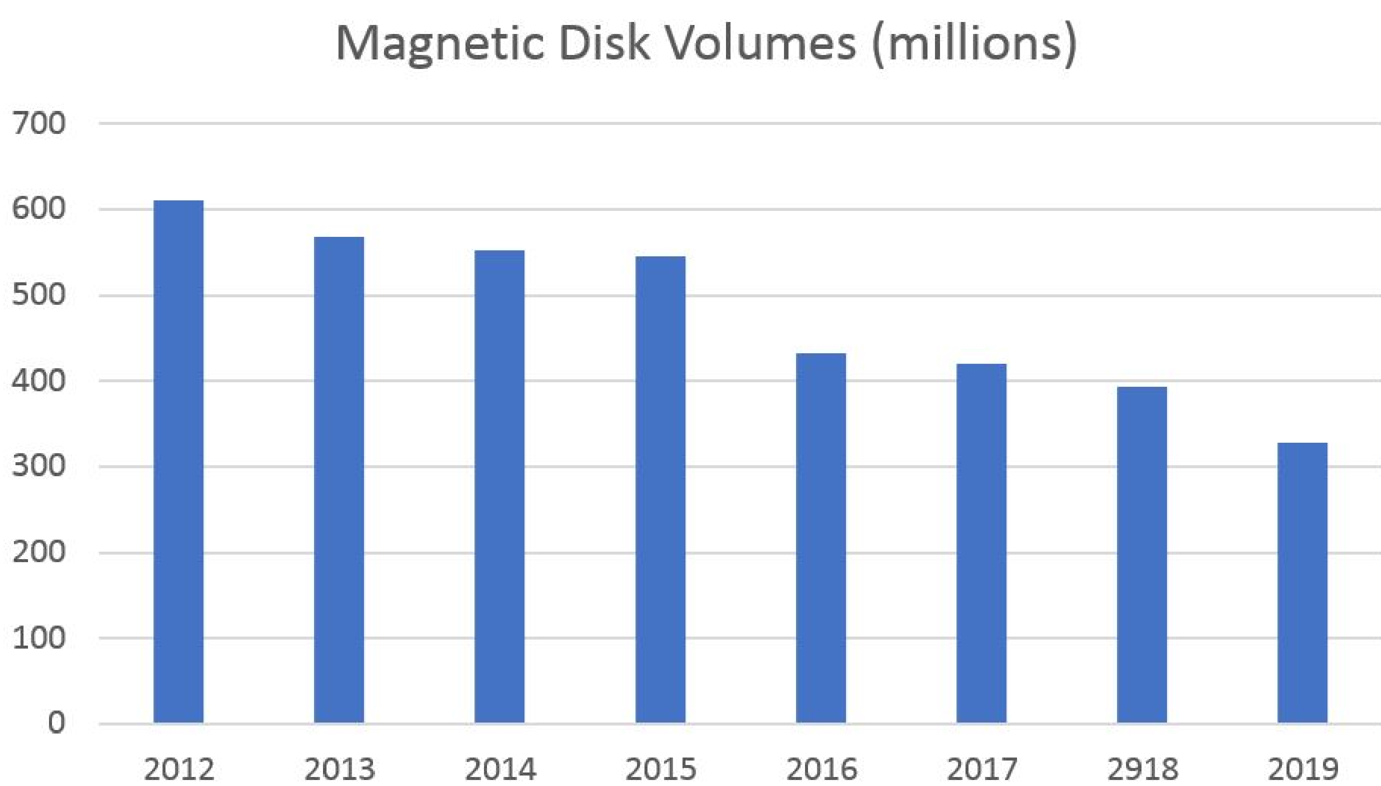

Figure 3: Worldwide Disk Drive Shipments

More and more, the disk industry will be shipping a singular product line, that being high-capacity 3.5-inch nearline drives. These are sold predominantly to large IT shops and cloud providers. Though there will be several versions of these drives, their base technologies are identical allowing all versions to be produced off the same manufacturing line. Variation of these drives will be in the areas of single or dual actuator, shingled or conventional recording, and SAS or SATA interface. In order to sustain that market, their products must maintain current reliability, while at the same time, continue to decrease their per-capacity cost. Protection of market share requires a multiple cost differential over consumer solid-state disk technologies.

Heat Assisted Magnetic Recording (HAMR) increases the areal density of a disk platter by heating the target area with a laser. This heated area is more receptive to a change in magnetic properties (reduced coercivity), allowing a lower and more focused charge to “flip” a smaller bit. The media then immediately cools and regains its high coercivity properties thereby “locking” the bit into place such that it requires a strong magnetic field to reverse it. For this reason, this technology can be thought of as being able to store data for a longer time then traditional disk technology. Microwave Assisted Magnetic Recording (MAMR) uses a microwave field generated from a spin torque oscillator (STO). In this method, the STO located near the write pole of the head generates an electromagnetic field that allows data to be written to the perpendicular magnetic media at a lower magnetic field.

For many years the disk industry has been investing heavily in HAMR and MAMR technology, realizing its importance for the product roadmap. The industry saw samples of disk drives with the HAMR technology in 2018 with promises of production in 2019. This production did not occur and it is now unclear when HAMR drives will be available in the marketplace and what the current issues are with the technology. HAMR technology is being introduced by Seagate, while MAMR technology will be supported by Western Digital. Recently Western Digital and Toshiba both announced that they will be utilizing MAMR technology in the short term to increase capacities with plans to move to HAMR when drives of greater than 24TB are required. This leaves Seagate all in with HAMR and they will need to bring it to market in order to compete from a capacity perspective with MAMR drives coming from its competitors.

Due to the delay of the HAMR technology, the Advanced Storage Technology Consortium (ASTC) has substantially revised its roadmap as shown in Figure 4. The new roadmap shows PMR technology now being extended until the end of 2021 while, at the same time, the start and end of HAMR has been shifted to about two years out. Another notable item is that the CAGR of capacity grown is below the 30% shown in the older roadmap.

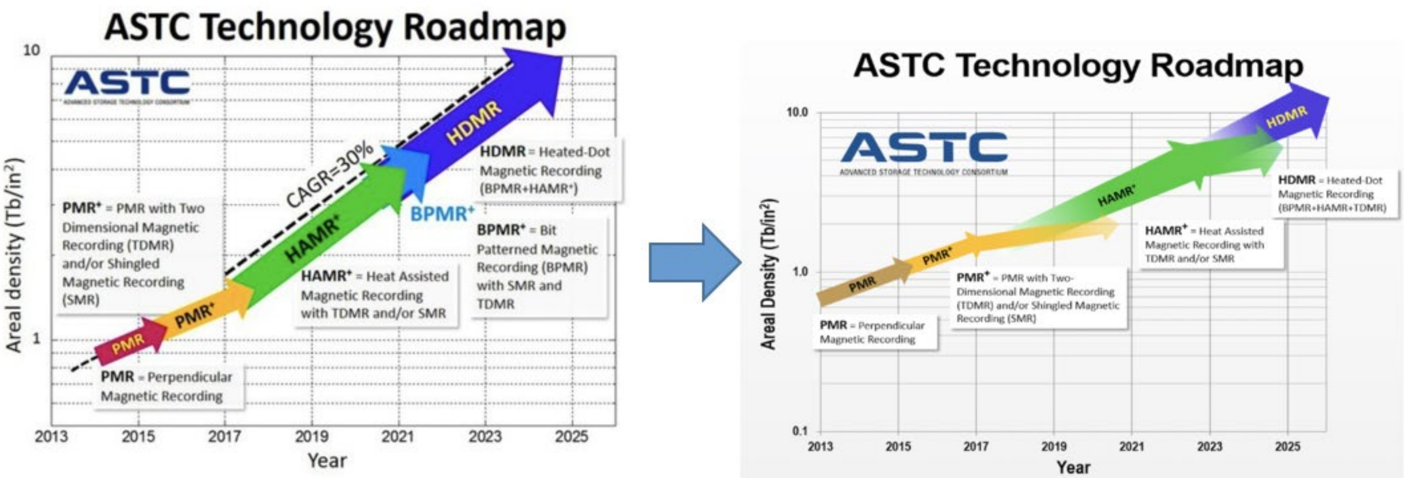

Figure 4: ASTC Technology Roadmap (Left is previous roadmap, right is new roadmap)

Given the simplification of the disk drive roadmaps into a single nearline product line, two paths for controlling these drives are emerging. One path will support the traditional disk SCSI command set thereby satisfying the current storage markets, such as Network Attached Storage (NAS). These drives will be formatted in conventional media recording (CMR) mode which will prevent the need to rewrite disk-based applications. In this mode the disk drive virtualizes the physical placement of the data from the application. The other path will be specifically for cloud companies and is for products that are specifically designed to store fixed content. A drive that supports this interface is known as a host-managed SMR drive, which is essentially the “zoned” block interface discussed earlier in the flash section of this paper. These drives cannot be purchased in the general market as the disk vendors ensure that they are only placed into environments that have been explicitly designed to support them. SMR takes advantage of the fact that the read head of a disk drive is smaller than the write head. This allows for tracks to be written in an overlapping manner as shown in the diagram below. This leads to a capacity increase of up to 20% of the same drive formatted in CMR mode, as the normal track-to-track gaps are eliminated. A side effect is that a track cannot be updated as doing so would overwrite the tracks around it. For this reason, an SMR drive is broken into “zones” whereby each zone is a region on the disk, typically 256MB in length.

Without going into the details of how these zones need to be managed, it is sufficient to say that the host software using these needs to be fully aware of their properties. Besides the capacity advantage, other advantages exist in the areas of improved write/read sequential performance and allowing the host to physically place data into zones that match the performance needs of the data. The performance of a zone corresponds to where it exists on the physical disk. Faster zones are at the outer diameter of the disk while slower zones are at the inner diameter. The performance of the fastest zone to the slowest zone is roughly 2.5 times, which corresponds to the ratio of the circumferences of the disk at the outer edge and the inner hub.

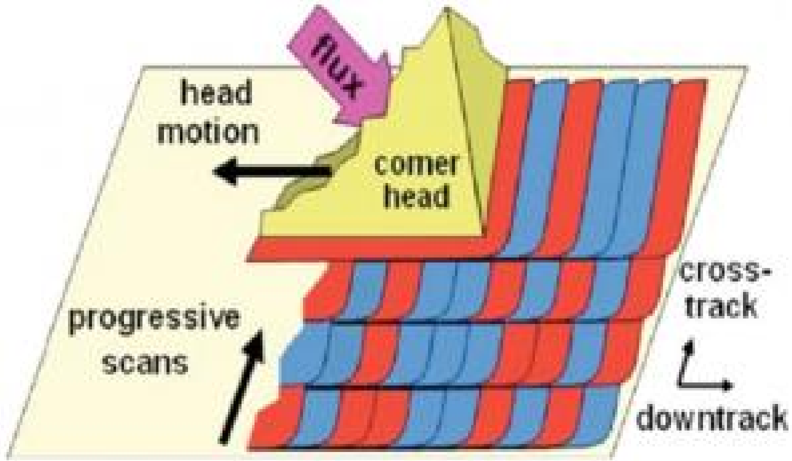

As noted in the architecture section of this paper, we predict that magnetic disk drives predominantly will be used as community storage. Given that community storage will contain data that is not frequently updated, a drive that has high capacity but not a particularly good random I/O performance will be adequate. However, for cloud providers the story is completely different. Cloud providers have two-tier architectures with the first being flash and the second being magnetic disk. They have years of data collected and fully understand the workload patterns of their various applications. Many of these applications have access patterns that are time-based in that the older the data, the less frequently it is recalled. With this information, they can derive how many I/Os will need to be serviced by the flash tier vs. the disk tier. Given the large discrepancy between the I/O performance of flash and disk, it is important that the majority of I/O requests are serviced from the flash tier while it is preferable that the bulk of the data be stored on more cost-effective disk. The lower the I/O rate of the disk tier as a function of capacity (I/Os per TB), the more flash will need to be purchased to avoid backing up I/Os on the disk tier, resulting in time delays to the consumer. This is a chicken and egg problem in that if the disk industry overcomes the technical challenges associated with increasing capacity, then they have to face how to improve I/O performance at the same rate (or better) than the capacity increases.

For example, if a cloud company buys 12TB drives today, and in a year, 24TB drives with the same I/O rate as the 12TB drives become available, then the organization might opt to buy only half the amount of 24TB drives to fulfill its capacity requirements, and buy more flash drives to gain the additional I/Os not provided by the disk tier. As discussed above, one technique that improves I/Os is through strategic placement of SMR zones with different speeds. A second technique being touted by both Seagate and Western Digital is to add additional independent actuators inside the drive. This is a blast from the past in that most disk drives manufactured prior to 1990 had two or more actuators. This automatically doubles the I/O performance of the disk drive; however, the cost of the second actuator could result in a 20 percent to 30 percent price increase. It is to be determined as to whether the cloud companies will see enough benefit to justify the cost. Non-cloud deployment applications that require high I/O will move to flash as it has two orders of magnitude better performance than that of even dual actuator disk drives.

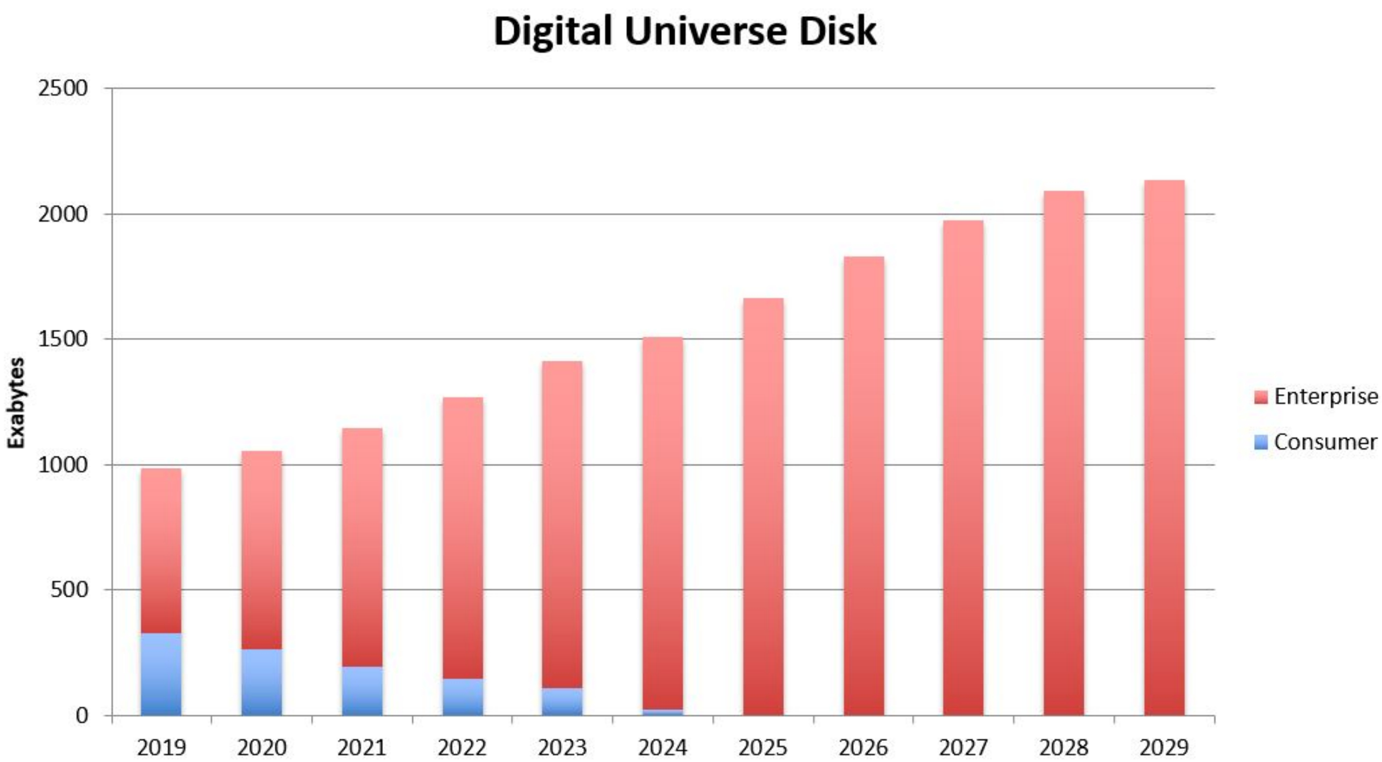

Figure 5: Digital Universe Disk

As seen above, Spectra is predicting a very aggressive decrease in the aggregate shipped capacity of consumer magnetic disk as flash disk takes over that space. Capacity increases in enterprise storage will not maintain a pace that will allow the disk industry to realize volume or revenue gains.

Some reservations are warranted as to the market's ability to deliver advanced technologies and restart the historical cost trends seen in disk for decades. If the industry is unable to cost effectively and reliably deliver on this technology, the intrusion of flash into its space will be greater.

Tape

The digital tape business for backing up primary disk systems has seen year-to-year declines as IT backup has moved to disk- based technology. At the same time, however, the need for tape in the long-term archive market continues to grow. Tape technology is well suited for this space as it provides the benefits of low environmental footprint on both floor space and power; a high level of data integrity over a long period of time; and a much lower cost per gigabyte of storage than any other storage medium.

A fundamental shift is underway whereby the market for small tape systems is being displaced by cloud-based storage solutions. At the same time, large cloud providers are adopting tape -- either as the medium of choice for backing up their data farms or for providing an archive tier of storage to their customers. Cloud providers and large scale-out systems provide high levels of data availability through replication and erasure coding.

These methods have proven successful for storing and returning the data “as is.” However, if during the lifecycle of that data, it becomes corrupted, then these methods simply return the data in its corrupted form. For the tape segment to see large growth, a widespread realization and adoption of “genetic diversity,” defined as multiple copies of digital content stored in diverse locations on different types of media, is required to protect customers’ digital assets. More recently, due to ransomware and other forms of attacks, we are seeing a greater interest in using tape as a last means of defense.

Linear Tape Open (LTO) technology has been and will continue to be the primary tape technology. The LTO consortium assures interoperability for manufacturers of both LTO tape drives and media. In 2018, the eighth generation of this technology was introduced, providing 12TB native (uncompressed) capacity per cartridge. A new seventh generation tape formatted in an eighth-generation tape drive provides a capacity point of 9TB and, as of this writing, has the lowest cost for capacity ratio. Each generation of tape drive has been offered in both a full height and a more cost-effective half-height form factor. As seen in the following table, the LTO consortium is providing a very robust roadmap in terms of future products all the way to LTO-12 at a capacity point of 192TB on a single piece of media. The majority of capacity increase will be gained through adding more tracks across the tape rather than increasing the linear density of the tape. The challenges for realizing this roadmap are multi-fold.

Tape, from a capacity perspective, has a large surface area, which means it has a much lower bit density then that of current disk drives; however, as a removable media, interchange of cartridges between drives requires that the servo systems have enough margin to handle variances in drive writing behaviors. This variance is directly correlated to how precisely the tape can be moved across the tape heads at high speed. A rule of tape drive design is that the longer and heavier the tape path, the better the tape can be positioned. This presents a challenge to the half-height drive as its tape path is shorter and lighter than that of the full-height drive. For this reason, half-height drives may need to be slowed down, resulting in lower bandwidth, or operated at lower capacity points. Another challenge for tape is that as tape capacities have increased, the bandwidth to read and write a tape has improved at a much slower rate. For cloud companies, who measure their performance as a function of capacity, this poses a problem as they are required to purchase a larger number of drives for each successive tape generation. For example, a customer’s tape system requirement might be that, for each petabyte stored, there should be 360 MB/s bandwidth available. For LTO-8, this would be satisfied by using a ratio of 84 cartridges (1000 TB / 12TB per cartridge) per drive. Looking into the future and considering an LTO-11 drive that can transfer at 500 MB/s, this would result in a cartridge to drive ratio of 14 (500/350 x (1000 / 96). Given that the “true” cost of a tape cartridge is the cost of the media plus the cost of the drive divided by the number of cartridges per drive, this will erode some of the cost-per-capacity advantages of tape.

There are three primary methods for improving tape drive performance each posing its own challenges. First the tape cartridge can be run across the tape head faster. As stated previously, the bigger the tape path the better the control, which in turn drives higher drive costs. We believe this method will provide very little uplift of bandwidth in future generations. Another option would be to increase the linear bit density of the tape. This requires a more advanced media formulation similar to the switch from metal particle used on LTO-6 with a linear density of 15,142 bits/mm to the barium ferrite LTO-8 at 20,668 bits/mm respectively. We believe this method to have some amount of potential upside. The last method would be to increase the number of tracks on the tape head from the current 32 to possibly 64. This would result in a more expensive and complicated tape drive, however, would provide the best method for increasing tape drive performance in a substantial way.

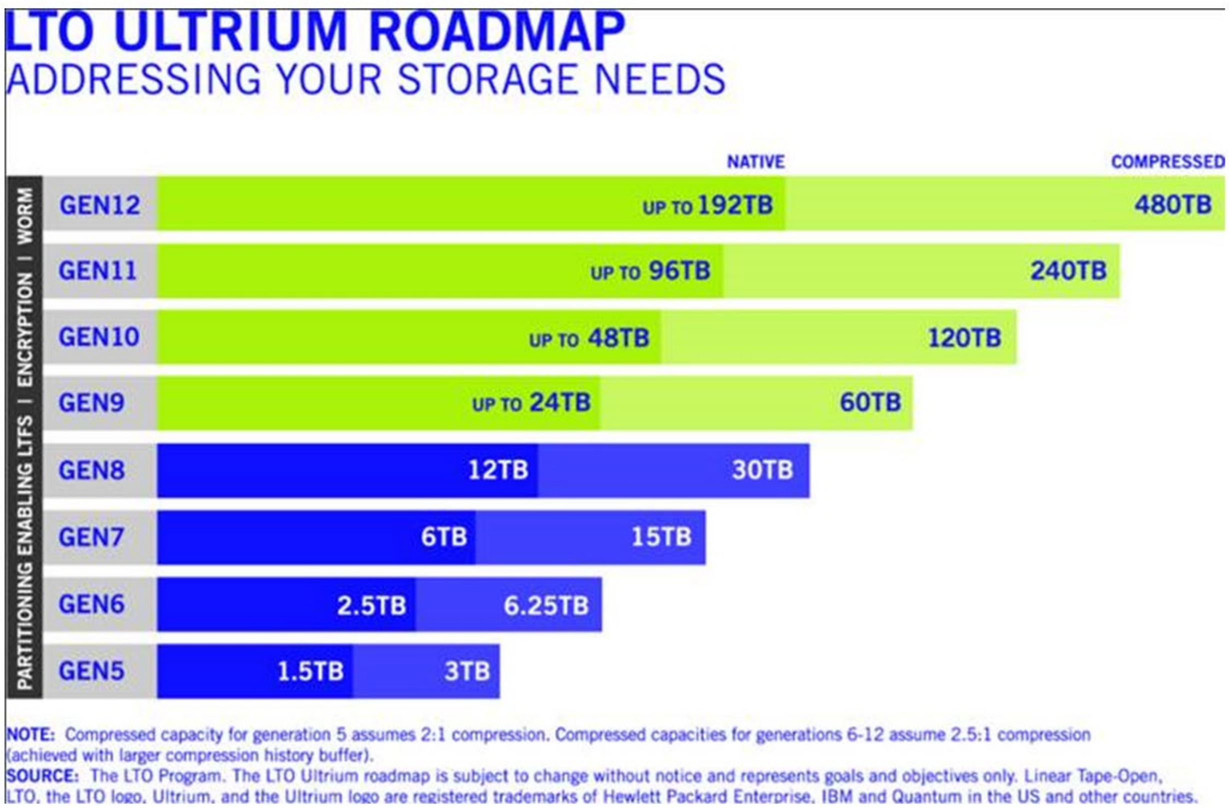

Customers with high duty-cycle requirements can consider using enterprise drives from IBM tape technology. IBM is now shipping the IBM® TS1160 Tape Technology with a native capacity of 20TB. These tape drives use TMR technology (tunneling magnetoresistance) which should allow the capacity to double three more times and forms the basis for future LTO generations. Additionally, these drives are offered with an optional native high-speed 10GE RoCE Ethernet interface.

With Oracle's exit from the tape business complete, IBM now is the only tape drive supplier. Fujifilm and Sony are the market's two tape media suppliers. In 2018 a lawsuit between these two suppliers drove prices higher on LTO-8 media. The legal issue was resolved in August 2019. Similar to other storage technologies, when new generations of tape are introduced, the cost per gigabyte of the technology is priced higher than the older technology on the market. Figure 6 shows the cost per TB of different LTO cartridges over time. As can be seen at the time of LTO-8 introduction, the price point of just below $12 a TB was higher than that of LTO-6 and LTO-7. It became cheaper than LTO-6 and LTO-7 in October 2017 and January 2019 respectively. Even with a sole provider for LTO-8 media, its price declined at about the same rate as previous tape generations. What drove early LTO-8 drive sales was its capability to format an unused LTO-7 tape cartridge that had a capacity of 6TB into an LTO-7 Type M cartridge with a capacity of 9TB. This provided an immediate 50% improvement in capacity without any additional costs. As can be seen, this media has been the lowest cost media over the last two years, and only recently has LTO-8 media been at a close enough price point to prompt customers into moving to the next generation.

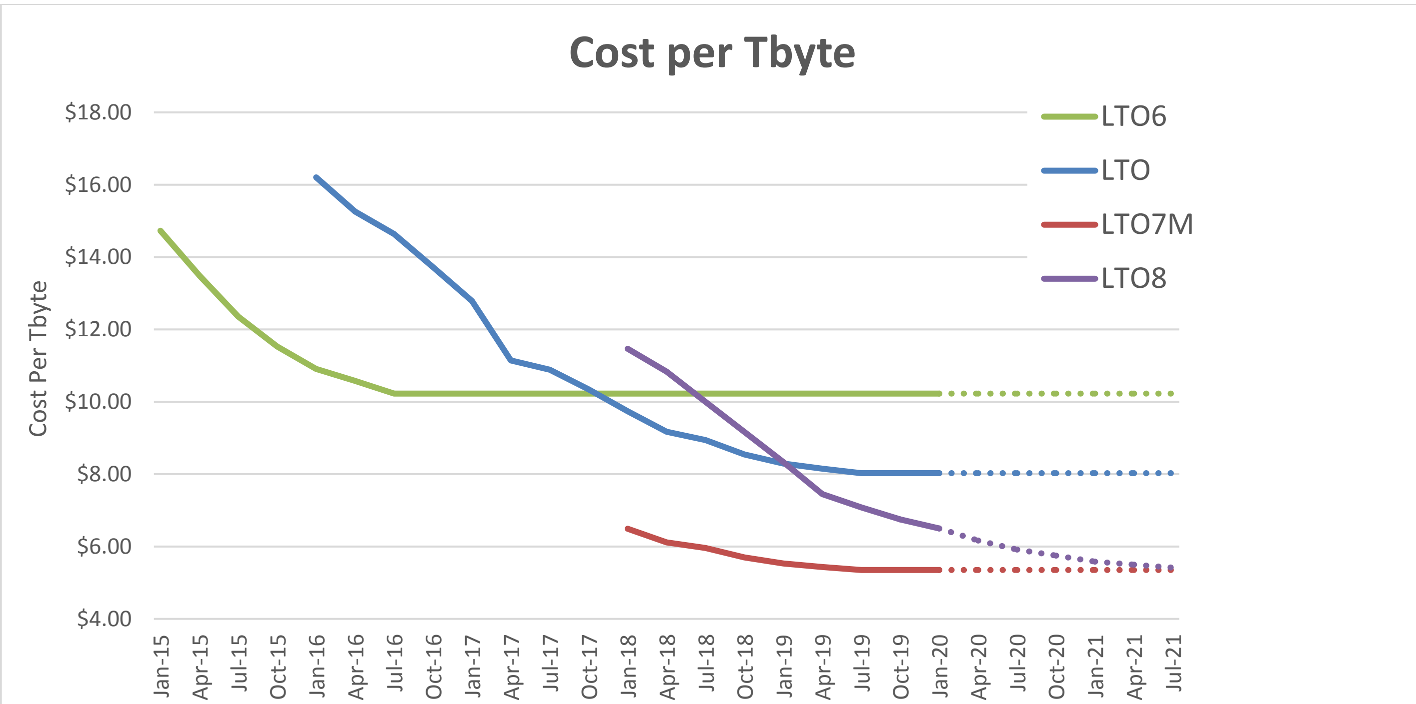

Figure 6: LTO Cost per Tbyte

As a result of some LTO-7 media being formatted into LTO-7 Type M media with a capacity of 9TB, it difficult to estimate the total capacity of tape sold in 2019. Though not reported, we believe tape drive volumes to be fairly consistent at around 100,000. For this reason, we are estimating the market to have shipped a 10% capacity increase in 2019 over 2018.

As small and medium enterprises move their disaster recovery strategies to the cloud, the volume of small tape systems will decrease. That data, however, will be archived on large tape systems installed by cloud manufacturers. This trend will lead to more aggressive adoption of new tape media as the cloud providers have incentive to lower the cost per gigabyte.

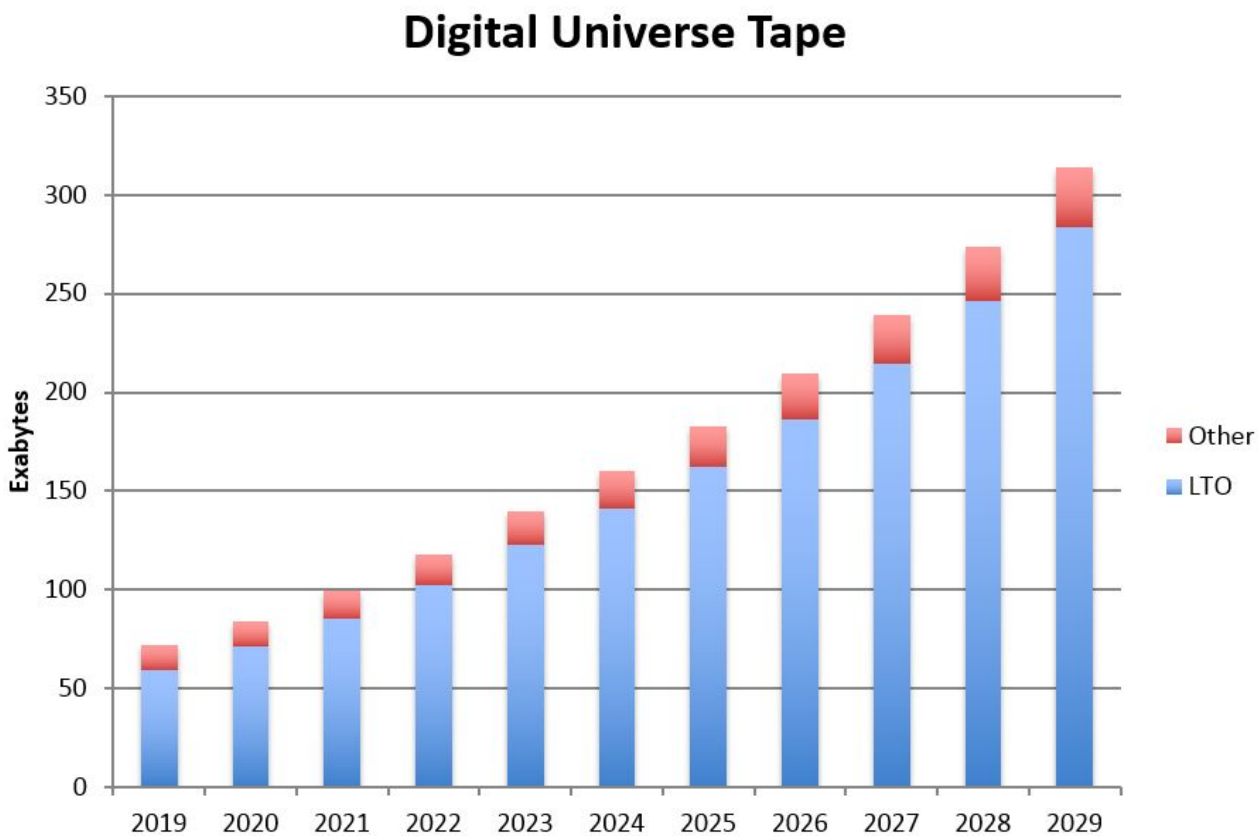

Figure 7: Digital Universe Tape

Cloud providers will mostly adopt LTO, and given their strength in purchasing overall tape technology, this will lead to a greater percentage of LTO shipments versus enterprise tape technology. The challenges for greater tape adoption with cloud providers lies partially in the environmental requirements of tape versus other equipment utilized (e.g., servers, disk, networking). Tighter controls of temperature and humidity are contrary to cloud providers' desire to be "green" by utilizing techniques that save cost, such as using external air. Tape library offerings that solve this problem efficiently without requiring the cloud provider to change their facility plan will find favor.

Optical

In early 2014, Sony Corporation and Panasonic Corporation announced a new optical disc storage medium designed for long-term digital storage. Trademarked "Archival Disc", it will initially be introduced at a 300GB capacity point and will be write-once. An agreement covers the raw unwritten disk such that vendors previously manufacturing DVD will have an opportunity to produce Archival Disc media. Unlike LTO tape, there is no interchange guarantee between the two drives. In other words, a disc written with one vendor's drive may or may not be readable with the others. Even if one vendor's drive is able to read another's, there may be a penalty of lower performance.

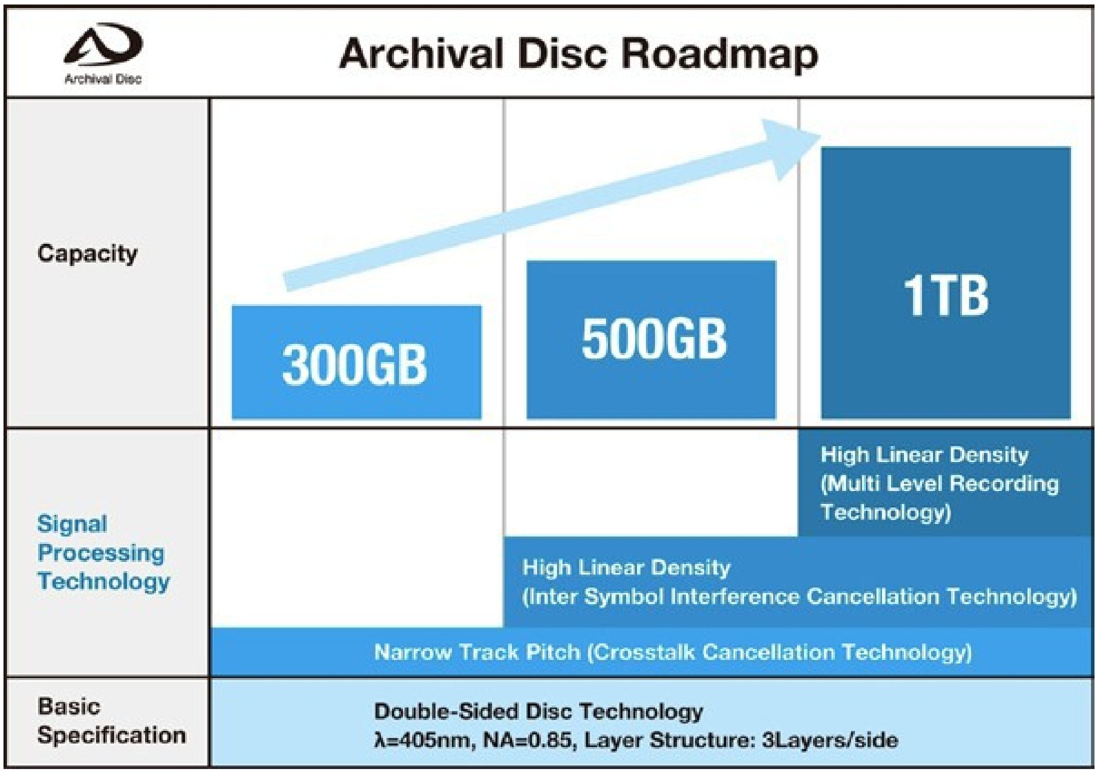

Figure 8: Archival Disc Roadmap
Source: Sony & Panasonic

Sony's solution packages 11 discs into a cartridge that is slightly larger in both width and depth than an LTO tape cartridge. The value proposition of this technology is its longevity and its backward read compatibility. As shown in the following chart, even stored at the extreme temperatures of 50 degrees Centigrade (122 degrees Fahrenheit), the disc has a lifetime of 50 years or more. Error rates are still largely unknown.

Sony and Panasonic are also guaranteeing backward read compatibility of all future drives. Unlike disk and tape, the media will not require migration to newer formats and technologies as they become available. This matches the archive mentality of writing once and storing the media for extended recoverability. Best practices in magnetic disk systems, conversely, where the failure rate of the devices considerably increases over time, indicate that data should be migrated between three to five years. In 2019, Sony announced the second generation 500GB per disc product resulting in a cartridge capacity of 5.5 B.

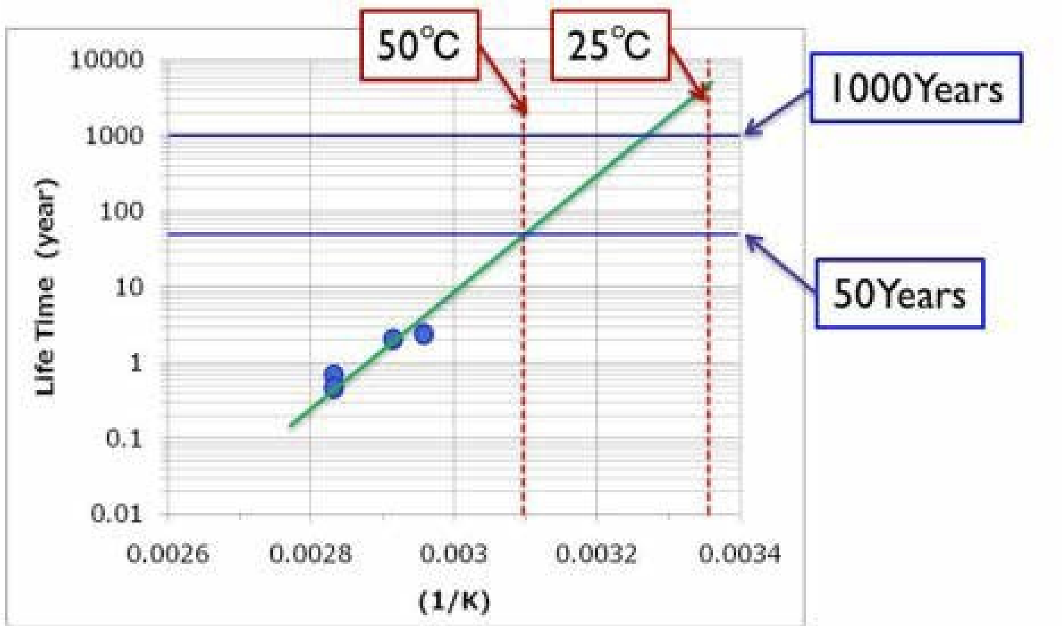

Figure 9: Long-Term Storage Reliability in Archival Disk
Source: Sony and Panasonic

For customers that have definitive long-term (essentially forever) archival requirements, the archive disc will find favor. The size of this particular market segment is small compared with the overall market for archival storage. The ability of this technology to achieve greater market penetration is primarily a function of the pricing of the media and, to some extent, the drives. If the media and drives were cheap enough, this would be an ideal archive media because of its longevity and its less restrictive environmental control requirements. However, this is the fifth year we have reported on this technology and the cost reduction that is required to become competitive has not occurred. Looking at the prices of different substrates that could be considered for archive, we see magnetic disk at less than $25 per TB, tape at less than $10 per TB while optical technology sits at just below $50 per TB. following is a comparison of the first and second generations of the Sony archive disc product. Note that the cost per storage ($/TB) has not substantially changed from one generation to the next. It is presumed that the new generation of media is five layers deep. This assumption is based on the fact that the performance of the product has improved at the same ratio as that of capacity.

Product	3.3 TB (Gen 2)	5.5 TB (Gen 3)
Introduced	2017	2020
Cart Price	$155	~$250
$/TB	$47	$45
Read BW (MB/s)	250	375
Write BW (MB/s)	125	187
Double Sided	YES	YES
Layers Per Side	3	Assuming 5

In order to achieve broader market acceptance, the price of the optical media would need to see a five to ten times cost reduction on a per-TB basis. This could be accomplished by a combination of increasing the capacity of the disc while at the same reducing its manufacturing costs. This type of breakthrough is being pursued by a company that was spun out of Case Western Reserve University in 2012 by the name of "Folio Photonics". The Folio Photonics technology depends on effective use of polymer co-extrusion, a manufacturing method that creates a low-initial cost disc technology. The process uses thin, flexible polymer film that can be cut and laminated to discs, so that 64 extremely thin layers that can be read on hardware are designed for that purpose. It is unclear if and when this technology will reach the market. If it does get to market, the question remains as to whether the cost and capacity points will be attractive enough to create a disruptive change in the marketplace.

Stored Digital Universe Optical Disc

Given the number of manufacturers and the variety of products (such as pre-recorded DVDs, Blue-Ray, etc.), it is difficult to project the stored digital universe for optical disc. To remain consistent with the intent of this paper, Spectra conservatively estimates the storage for this technology at 5EB per year, and that, with the introduction of the Archive Disc, this will grow at a rate of 1EB per year for the next 10 years. This value could change in either direction based on the previous discussion regarding disc pricing.

Future Storage Technologies

Being a $50 billion a year market, the storage industry has and will always continue to attract venture investment in new technologies. Many of these efforts have promised a magnitude of improvement in one or more of the basic attributes of storage, those being cost (per capacity), low latency, high bandwidth, and longevity. To be clear, over the last 20 years, a small portion of the overall venture capital investment has been dedicated to the development of low-level storage devices, with the majority dedicated to the development of storage systems that utilize existing storage devices as part of their solution. These developments align more with the venture capital market in that they are primarily software based and require relatively little capital investment to reach production. Additionally, they are lower risk and have faster time-to-market as they do not involve scientific breakthroughs associated with materials, light or quantum physics phenomenon.

Much of the basic research for advanced development of breakthrough storage devices is university or government funded or is funded by the venture market as purely a proof-of-concept effort. For example, an announcement was made regarding storing data in five dimensions onto a piece of glass or a quartz crystal capable of holding 360TB of data, literally forever. Advanced development efforts continue in attempting to store data into holograms, a technology that, for a long time, has been longer on promises than results.

Developments at the quantum level include storing data through controlling the “spin” of electrons. Though these and other efforts have the potential to revolutionize data storage, it is difficult to believe that any are mature enough at this point in time to significantly impact the digital universe through at least 2026. Historically many storage technologies have shown promise in the prototype phase, but have been unable to make the leap to production products that meet the cost, ruggedness, performance, and most importantly, reliability of the current technologies in the marketplace. Given the advent of cloud providers, the avenue to market for some of these technologies might become easier (see next section).

Cloud Provider Storage Requirements

Over the period of this forecast, cloud providers will consume, from both a volume and revenue perspective, a larger and larger portion of the storage required to support the digital universe. For this reason, storage providers should consider whether or not their products are optimized for these environments. This brings into question almost all previous assumptions of a storage product category. For example, is the 3.5-inch form factor for magnetic disk drives the optimum for this customer base? Is the same level of the cycle redundancy check (CRC) required? Can the device be more tolerant of temperature variation? Can power consumption and the associated heat generated be decreased? Does the logical interface need to be modified in order to allow the provider greater control of where data is physically placed?

Another way to consider the requirements for these providers is to ask the reverse question, which is, ‘What is it that they don’t need?’ Equipment that was designed for IT data centers may have substantial features that add cost and/or complexity to a product that are neither needed or wanted by customers. Additionally, systems that are managed as separate entities do not fit the cloud model because, within these operations, hundreds of identical systems may need to be managed from a central point of control.

For flash, numerous assumptions should be questioned. For example, what is the cloud workload and how does it affect the write life of the device and could this lead to greater capacities being exposed? Similar to disk, questions should be asked regarding the amount and nature of the CRC and the logical interface as well as the best form factor. Better understanding and tailoring of lower power nodes along with the need for refresh should be understood and tailored to meet cloud providers' needs.

Regarding the use of tape technology for the cloud, several questions arise, such as what is the best interface into the tape system. Given that tape management software takes many years to write and perfect, a higher level interface, such as an object level REST interface might be more appropriate for providers that are unwilling to make that software investment. When cloud providers have made that investment, the physical interface to the tape system needs to match their other networking equipment (i.e., Ethernet). Due to the fact that tape has tighter temperature and humidity specifications than other storage technologies, solutions that minimize the impact of this requirement to the cloud provider should be considered. Additionally, there are features provided by tape drives that are not needed, such as backward read compatibility, as systems stay in place until their contents are migrated into a new system. If tape capacities or time to market can be accelerated by dropping backward compatibility it should be seriously considered.

Cloud providers have a unique opportunity to adopt new storage technologies, based on the sheer size of their storage needs and small number of localities, ahead of volume commercialization of these technologies. For example, consider an optical technology whereby the lasers are costly, bulky and prone to misalignment, and the system is sensitive to vibration. If the technology provides enough benefit to a cloud provider, it might be able to install the lasers on a large vibration-isolating table with personnel assigned to keep some systems operational and in alignment. In such a scenario, an automated device might move the optical media in and out of the system. In a similar scenario where the media has to be written in this manner but can be read with a much smaller and less costly device, the media may be, upon completion of the writing process, moved to an automated robotics system that could aid in any future reads to be done.

SIZE OF THE DIGITAL UNIVERSE

The IDC report, published in November 2018 and commissioned by Seagate, predicts the 'global datasphere' will grow by more than 175 zettabytes (ZB) by 2025. This causes many in the industry to wonder whether there will be sufficient media to contain such huge amounts of data.

The Internet of Things (IoT), new devices, new technologies, population growth, and the spread of the digital revolution to a growing middle class all support the idea of explosive, exponential data growth. Yes, 175ZB seem aggressive, but not impossible. The IDC report took a top-down appraisal of the creation of all digital content. Yet, much of this data is never stored or is retained for only a brief time.

For example, the creation of a proposal or slide show will usually generate dozens of revisions -- some checked in to versioning software and some scattered on local disk. Including auto-saved drafts, a copy on the email server, and copies on client machines, there might easily be 100 times the original data which will eventually be archived. A larger project will create even more easily discarded data. Photos or video clips not chosen can be discarded or relegated to the least expensive storage. In addition, data stored for longer retention is frequently compressed, further reducing the amount of storage.

In short, though there might indeed be upwards of 175ZB, when a supply and demand mismatch is encountered, and there are many opportunities to synchronize:

- A substantial part of the data created will be by nature transitory, requiring little or very short retention.
- Storage costs will influence retention and naturally sort valuable data from expendable.
- Long-term storage can be driven to lower-cost tiers. Cost will be a big factor in determining what can be held online for immediate access.
- Flash, magnetic disk, and magnetic tape storage is rewritable, and most storage applications take advantage of this. As an example, when using tape for backups, new backups can be recorded over old versions up to 250 times, essentially recycling the storage media.
- The "long-tail" model will continue to favor current storage – as larger capacity devices are brought online, the cost of storing last year's data becomes less significant. For most companies, all their data from 10 years ago would fit on a single tape today.

Spectra's analysis also differs from the larger projections by omitting certain forms of digital storage such as pre-mastered DVD and Blue-Ray disks.

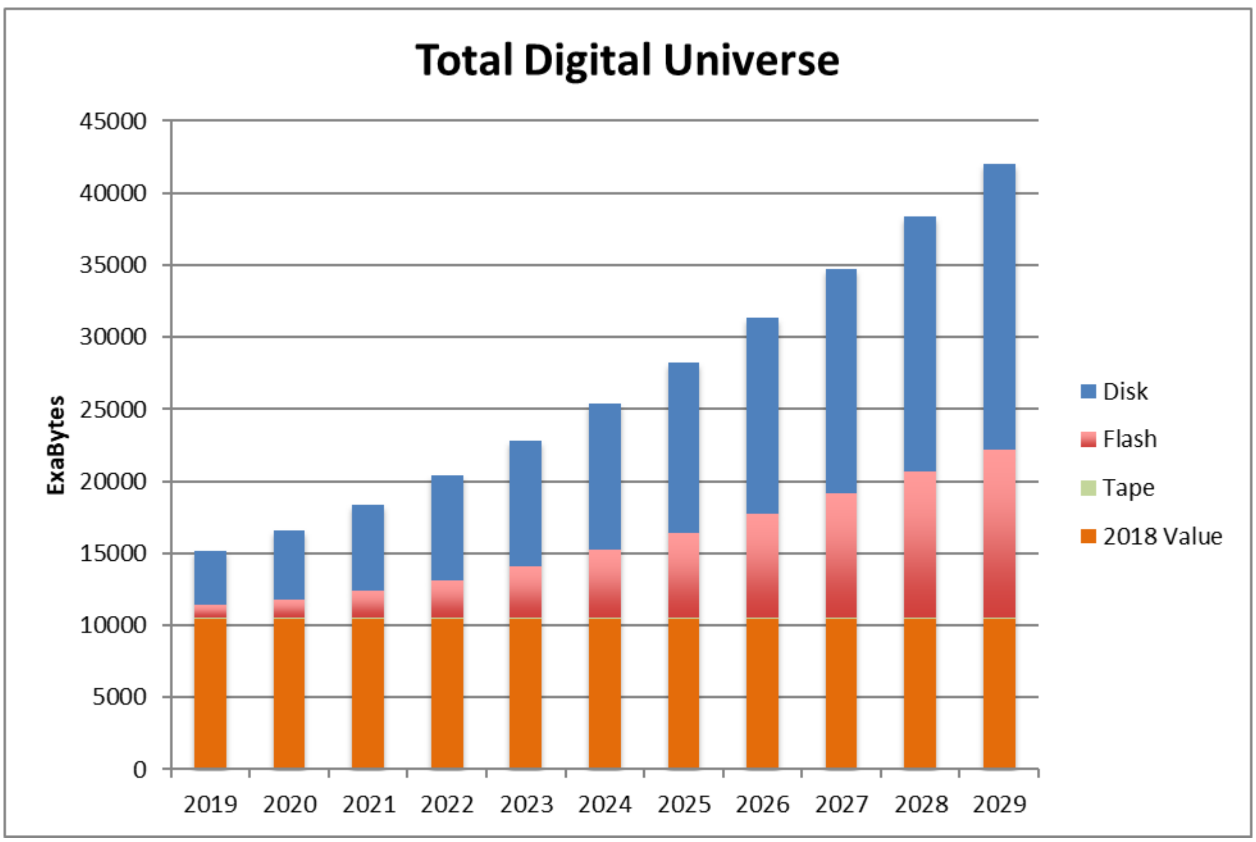

Figure 10: Total Digital Universe

CONCLUSION

We agree with Nelson Mandela's quote at the beginning of the report that states, 'Action without vision is only passing time; vision without action is merely daydreaming; **but vision with action can change the world**.' Today, more than ever, as many sectors of the economy remain unclear due to the coronavirus outbreak, the vision and steps taken today to advance technology and fortify supply chains will further preserve the world's treasury of information for tomorrow.

For the foreseeable future, the storage growth requirements of customers will be fulfilled by storage device providers who will continue to innovate with higher performance and higher capacities to meet increasing demand. As noted in the report, every storage category is exhibiting technology improvements. First, we see memory hosted Xpoint technology becoming the latest high-performance standard for database storage. At the flash layer, 3D fabrication technology is allowing for the creation of density parts while lowering the cost per gigabyte. In the meantime, disk manufacturers are closing in on delivery of HAMR and MAMR technologies that will allow them to initially deliver disk drives of 20TB while also enabling a technology roadmap that could achieve 50TB or greater over the next 10 years. Finally, tape has enough technology headroom that it will achieve storage capacities of 100TB or higher on a single cartridge in the next decade.

Data Storage Dilemma

Given that a singular storage technology has yet to be invented that combines the highest performance at the lowest cost, customers will continue to face the dilemma of what data, at what time, should be stored on which medium. Data that supports a project in progress one day may be suitable for archive once that project is completed. This would thereby lower overall storage costs by freeing up storage capacities for future projects. Software tools that allow customers to identify the usage patterns of their data and then provide for the movement of infrequently accessed data to lower tiers of storage have been available for quite a while; however, these tools have been priced such that most of the benefit of the storage savings are lost. A new generation of tools is required that improves data storage efficiencies while mitigating storage costs.

Designing with the Cloud in Mind

Over the last few years, a new question has arisen for storage administrators, which is 'where' to deploy 'what' storage. More specifically, what data should be placed in the cloud, what data should be located on premise, and what data should be stored in both locations. Each location provides benefits and cost trade-offs. The demand by storage customers to use cloud-based storage prompted many legacy storage providers to 'shoehorn' basic cloud capability into their existing products. Primarily this has consisted of providing customers with the capability of making cloud disaster recovery copies of their on-premise data. This is a pattern that has been seen before such as in the adoption of flash technology into disk arrays. The first generation of storage systems to utilize flash were existing products that were designed before flash storage was available. For this reason, it was typically integrated into these systems as extended cache because that is where it could most easily fit into these existing architectures. Customers gained some benefit, but not the full scope of the technology. Second and third generation solutions were designed with flash in mind and provided tremendous capability to the customer. Over the last few years the solutions have become the hottest segment in the storage system business.

Supporting Complex Workflows

We consider cloud integration by on-premise storage systems to be in this first phase. Next-phase products will be designed from the ground up with the cloud in mind. These products will allow for seamless integration of applications into the storage infrastructure, regardless of storage location -- whether in the cloud and/or in multiple on-premise locations. Complex customer workflows will be supported, through policies set by the customer, that allow data to be automatically moved to the right location(s), to the right storage tiers, at the right time. With this capability, customers will have the freedom to decide which processes they want to run locally and which ones in the cloud – all without having to think about the underlying storage system.

There are many interesting storage ideas being pursued in laboratory settings at different levels of commercialization: storing data in DNA, 3D Ram, (5 dimension optical) hologram storage – plus many that are not yet known. Technology always allows for a singular breakthrough, unimaginable by today's understanding, and this is not to discount that possibility.

Planning for the Future

Spectra's projections do not call for shortages or rising media costs. Due to the coronavirus pandemic there could be short-term supply-side shortages; however, it is unclear at this point whether reduced demand will result in a balanced or unbalanced market. But there are credible risks against expectations of precipitously declining storage costs. Storage is neither free nor negligible and proper designs going forward need to plan for growth and apportion it across different media types, both for safety and economy. Corporations, government entities, cloud providers, research institutions, and curators must continue to plan for data management and preservation today, evaluating data growth against projected costs.

CONTACT US

Spectra has stepped out for its fifth year to make predictions on the data storage industry's future based on what we see today. Think these predictions are too high? Too low? Missing something important? Spectra plans to update and publish this document with new data and new thinking where needed. Please let us know your thoughts.

Spectra Logic:
www.spectralogic.com

APPENDIX NOTES

Footnotes:

[1]Source: IDC Data Age 2025: the Digitization of the World from Edge to Core, November 2018
https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

[2]Source: WW Enterprise Storage Market, IDC website: https://www.idc.com/getdoc.jsp?containerId=prUS45155319

Charts:

Figure 4: Page 17, Source: ASTC https://hexus.net/tech/news/storage/123953-seagates-hdd-roadmap-teases-100tb-drives-2025/

Figure 8: Page 24, Source: Sony and Panasonic https://hexus.net/tech/news/storage/67165-sony-panasonic-create-archival-disc-standard/

Figure 9: Page 25, Source:
https://www.snia.org/sites/default/orig/DSI2015/presentations/ColdStorage/Yasumori_Archival_Disc_Technology-2.pdf (slide 23)

This preliminary calculation is based on objective data and Sony offers no guarantee that media are capable of storing data for 100 years irrespective of the environment.

All unsourced charts in this report were created by Spectra Logic.

About Spectra Logic

Spectra Logic develops data storage solutions that solve the problem of short- and long-term digital preservation for business and technology professionals dealing with exponential data growth. Dedicated solely to storage innovation for more than 40 years, Spectra Logic's uncompromising product and customer focus is proven by the adoption of its solutions by industry leaders in multiple vertical markets globally. Spectra enables affordable, multi-decade data storage and access by creating new methods of managing information in all forms of storage—including archive, backup, cold storage, private cloud and public cloud. To learn more, visit www.SpectraLogic.com.