

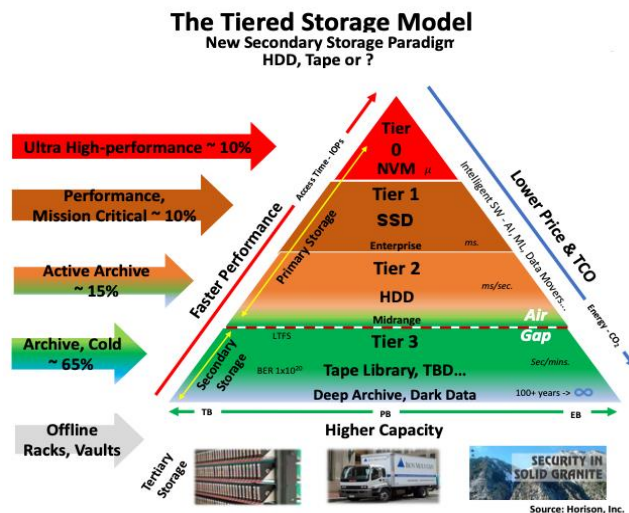


Magnetics Society Joint Chapter

Santa Clara Valley and San Francisco Sections

The Future of Tiered Storage –Tape or HDD or ?

September 14 @ 6:30 pm - 8:00 pm PDT



Fred Moore, President
Horison Information Strategies

Horison.com

The Zettabyte Era Takes Off

2016 eta
Zettabyte era
One sextillion

1 ZB = 1,000 EBs = 1×10^{21} bytes

~2043 eta
Yottabyte era
One septillion

One Zettabyte Equivalents

Watching the entire Netflix catalog 3 million times

Enough to record a video call that's more than 237,000,000 years long

A stack of books from Earth to Pluto 20 times (72 billion miles)

66.7 years of the Large Hadron Collider's experimental data

By 2025

~8.4 ZB to be Stored

>75% Generated Outside the Data Center

55.36 million LTO-9 (18 TB) cartridges, 50 million 20 TB HDDs, 250 billion DVDs

125 million years of 1 hour TV shows, 10 billion 4k movies, ~ 7.5 trillion MP3 songs

~5B internet users, 3.9 ZBs of global IP traffic generated, 82% from video in 2022



Source: Numerous industry sources, Horizon Information Strategies.

New Apps Creating Data at Warp Speed...



Every “tweet”, “swipe left or right”, “accept”, “post”, “tag”, “friend” or “like” creates some amount of stored data.

AI, Social Media and New Apps Activate the Archives

Time to Reach 100 Million Users	
5 days	
ChatGPT	2 months (AI)
WhatsApp	3 yrs, 4 months
YouTube	4 yrs, 1 month
Twitter	5 yrs, 2 months
LinkedIn	7 yrs, 11 months

How Much
Data?

How Much
Storage?

Where is it
Stored?

How Much
Value?

Source: UBS, Yahoo Finance, Horison Information Strategies





Big Data, AI, ML Fuel Shift to Secondary Storage

Key Archive Trends - 2025 and Beyond

Secondary (persistent storage) is designed to keep less critical data on highly economical, secure storage mediums that doesn't need to be accessed as frequently as data in [primary storage](#).

☀ Secondary Storage Is the Largest and Fastest Growing Storage Tier

- “ ~8.4 ZB Stored on SSD, HDD and Tape by 2025 (cagr. 25-30%)
- “ ~80% of All Data Stored (6.72 ZB) is Lower Activity, Archival, Cold
- “ Over 80% of Data Created is Unstructured
- “ AI, ML, Big Data Analytics, Edge/IoT Computing, Virtual Reality, Augmented Reality, Gaming and Robotics are Filling the Archives
- “ The Archival Copy is Usually the Only Copy of Data
- “ Retention Periods Over 100 Years are Common
- “ The Active Archive Becomes a De-facto Standard Tier

The Pillars of Secondary Storage

*Requirements To Meet And Sustain Future Secondary Storage Demands
Can One Technology Meet the Challenge?*



PRICE

- Lowest \$/TB
- Lowest TCO
- Floorspace
- Easy Install
- Self maintenance



PERFORMANCE

- Access Time
- Active Archive
- Random and Sequential Access
- Dynamic archive workloads



CAPACITY

- Seamless Scaling
- Rack Scale Design
- Sustainable Roadmap
- Legacy Emulation
- Open standards



RELIABILITY & AVAILABILITY

- BER =>10X²⁰
- RAID, RAIT, RAIL
- Erasure Coding
- Long Life Media
- Resilient to:
Temp, Humidity,
Radiation, EMP,
Saltwater



THROUGHPUT

- Raw Data Rate
- RAID, RAIT, RAIL



SECURITY FEATURES

- Air Gap
- Encryption
- WORM
- Removeable Media
- Physical Vault
- Cybercrime prevention

THE RISE of SUSTAINABILITY

Optimization Keys: Power Consumed, Water Consumed, CO2e and Electronic Waste Generated, Use Reliable Energy Sources

By 2025 ~6.72 ZBs of Secondary Storage Data Stored

Source: Horison Information Strategies


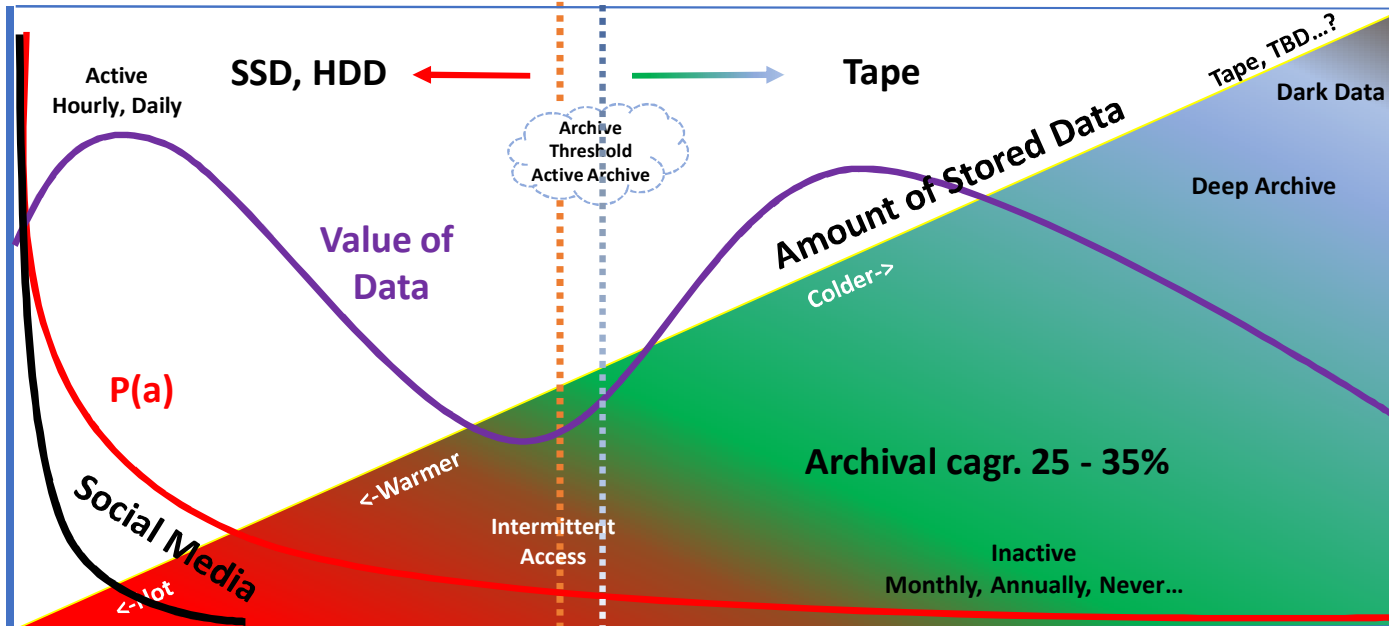

Digital Data Lifecycle

When Does Data Become a Candidate for Secondary Storage?

Probability
Of Access
 $P(a)$

Primary Storage
Active Online Data
~20% of Data

Secondary Storage
Archival, Persistent Data
~80% of Data



Lifecycle Profile Factors

Probability of access $P(a)$ declines as data ages

Most data typically becomes archival in ~90 - 120 days

Seldom Backed Up (1 copy)

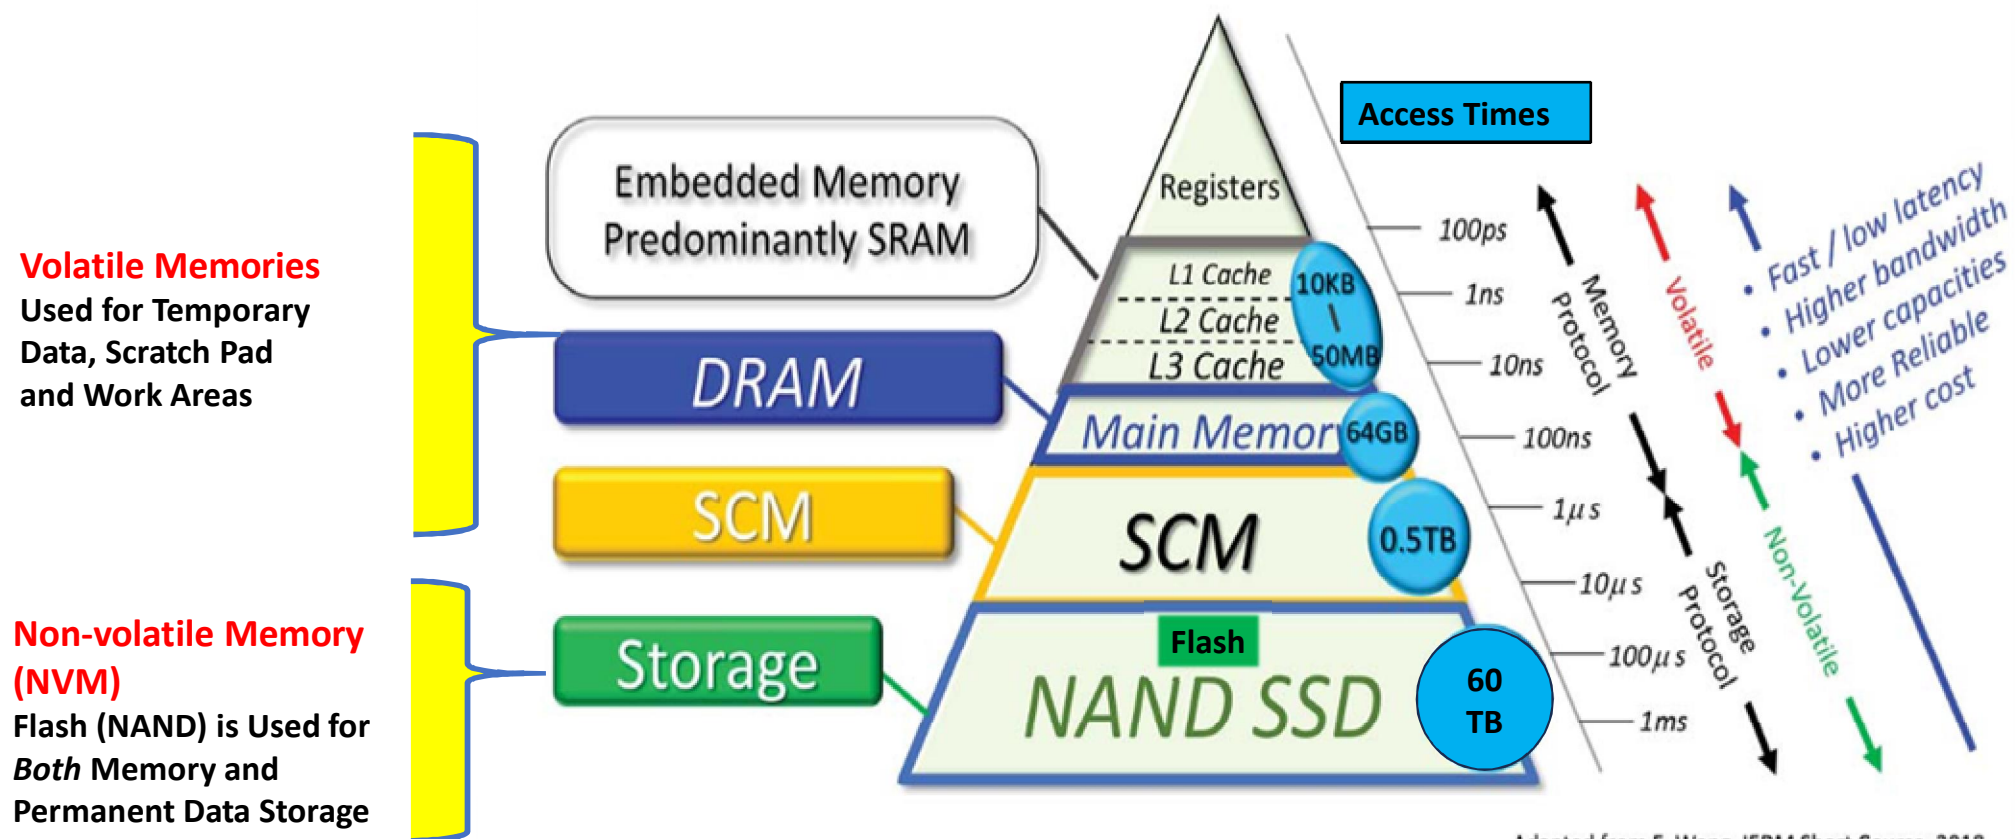
The value of data can vary over time

Archival retention can be >100 years to ∞

Source: Horison Information Strategies

Memory Hierarchies

Flash (NAND) Memory Technology is Also Used for Data Storage



Enterprise Data By 2025

Source: TrendFocus

~8.4 ZB Stored
WW

~6.72 ZB
(80%)
of data
is low
activity

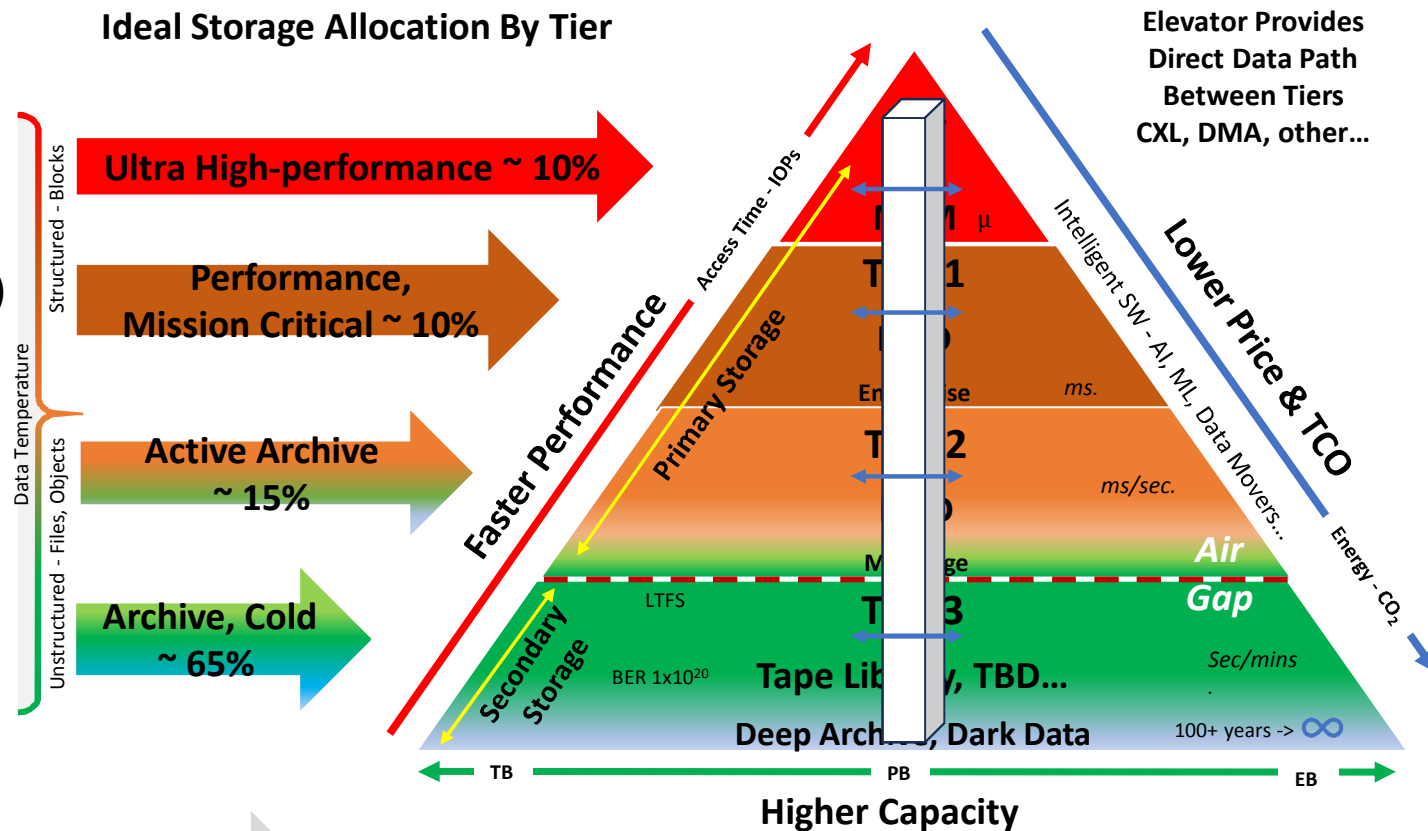
Intelligent
Software (AI, ML)



Legacy, Unknown, Orphaned Data

New Secondary Storage Paradigm Begins to Emerge – HDD, Advanced Tape or ?

Ideal Storage Allocation By Tier



Offline
Racks, Vaults

Tertiary
Storage



Source: Horison Information Strategies.

After 2025 a New Secondary Storage Model Begins to Emerge

New Sub-Tiers Target the Archival Avalanche

By 2025

- ~8.4 ZB total enterprise data stored WW
- ~20% of stored data is active (Primary, online)
- ~80% (~6.7 ZB) of stored data is cold/archival (Secondary)
- Active archive becoming a standard tier
- However, the majority is stored on the **wrong tier**

Immutable Data Properties

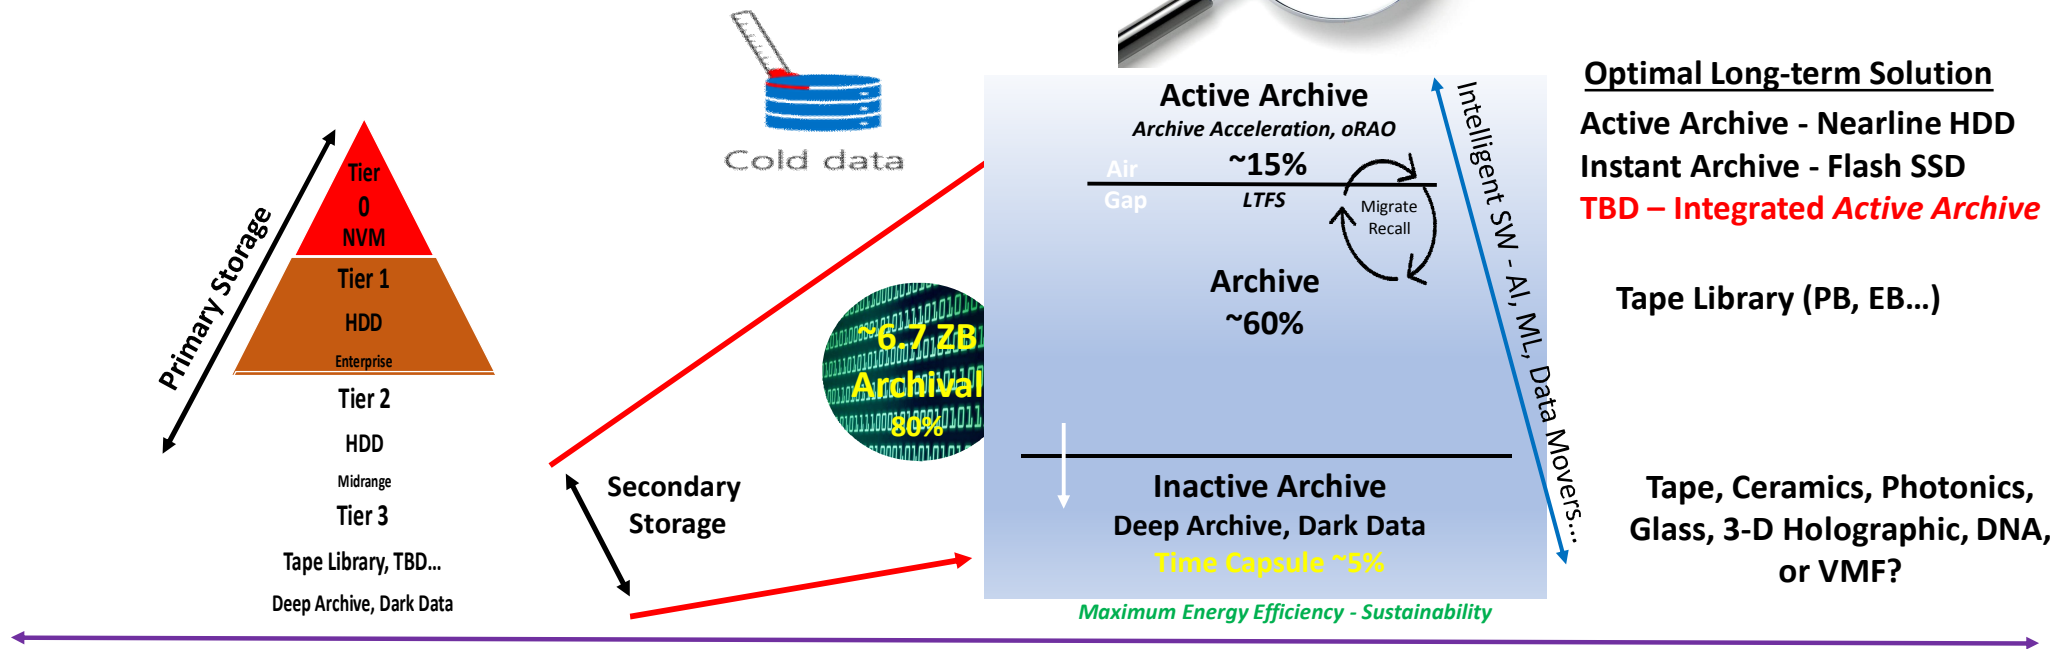
Archival Data is Mostly **Write Once***

WORM – **Write Once**, **Read Many**

WORSE -- **Write Once**, **Read Seldom**

WORN – **Write Once**, **Read Never**



*Can't be deleted, modified or overwritten



Source: Horison Information Strategies

Sustainability and TCO Comparison for Secondary Storage

Data That Isn't Being Used Shouldn't Consume Energy



- “ Ten-year TCO and CO₂ Reduction for 10 PB Growing at 35% Annually
- “ Heavily Favors Tape Over Disk for Archives

Source: Brad Johns Consulting


Hyperscale Data Centers Reshape IT Landscape

Heavy Influence on *FUTURE* Storage Infrastructure



DID YOU KNOW?


- “ Over 700 Hyperscale Data Centers (HSDCs) WW - US 49%, China 15%.
- “ HSDCs contain > 50% of WW storage capacity and servers.
- “ Many over 100,000 ft².
- “ Largest is > 10.763 million ft² = 132.9 soccer fields.
- “ Avg. of ~25,000 servers per support technician per shift, 48U racks.
- “ Avg. server repair time ~2 mins., was one hour.
- “ HSDCs use ~20% of the world's data center electricity, avg. PUE ~1.1.



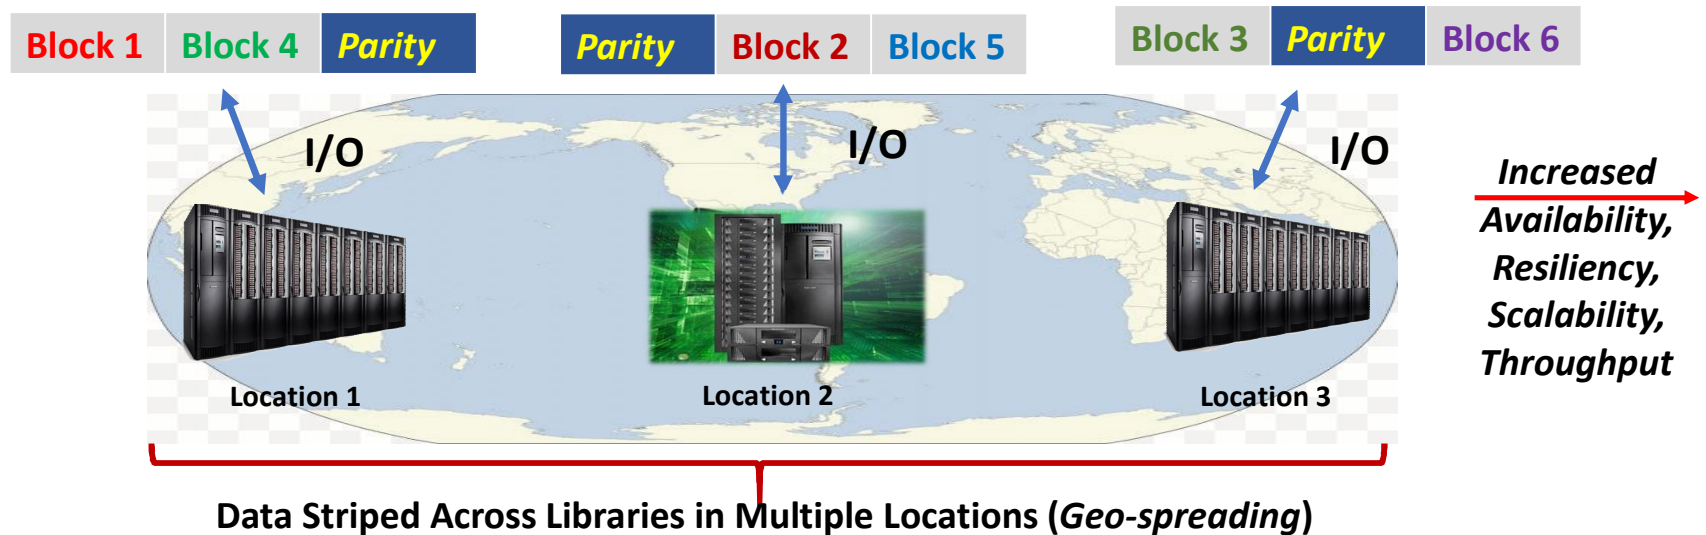
- “ Energy, sustainability and carbon footprint challenges become severe with enormous growth.
- “ HSDC secondary storage architecture includes tape based erasure coding, RAIL, geo-spreading.

RAIL - Redundant Array of Independent Libraries

Geo-spreading and Erasure Coding Enable Unprecedented Resiliency for Secondary Storage



Erasure Coding Splits Files or Objects Into Data Blocks (chunks) and Creates Parity Blocks For Data Recovery

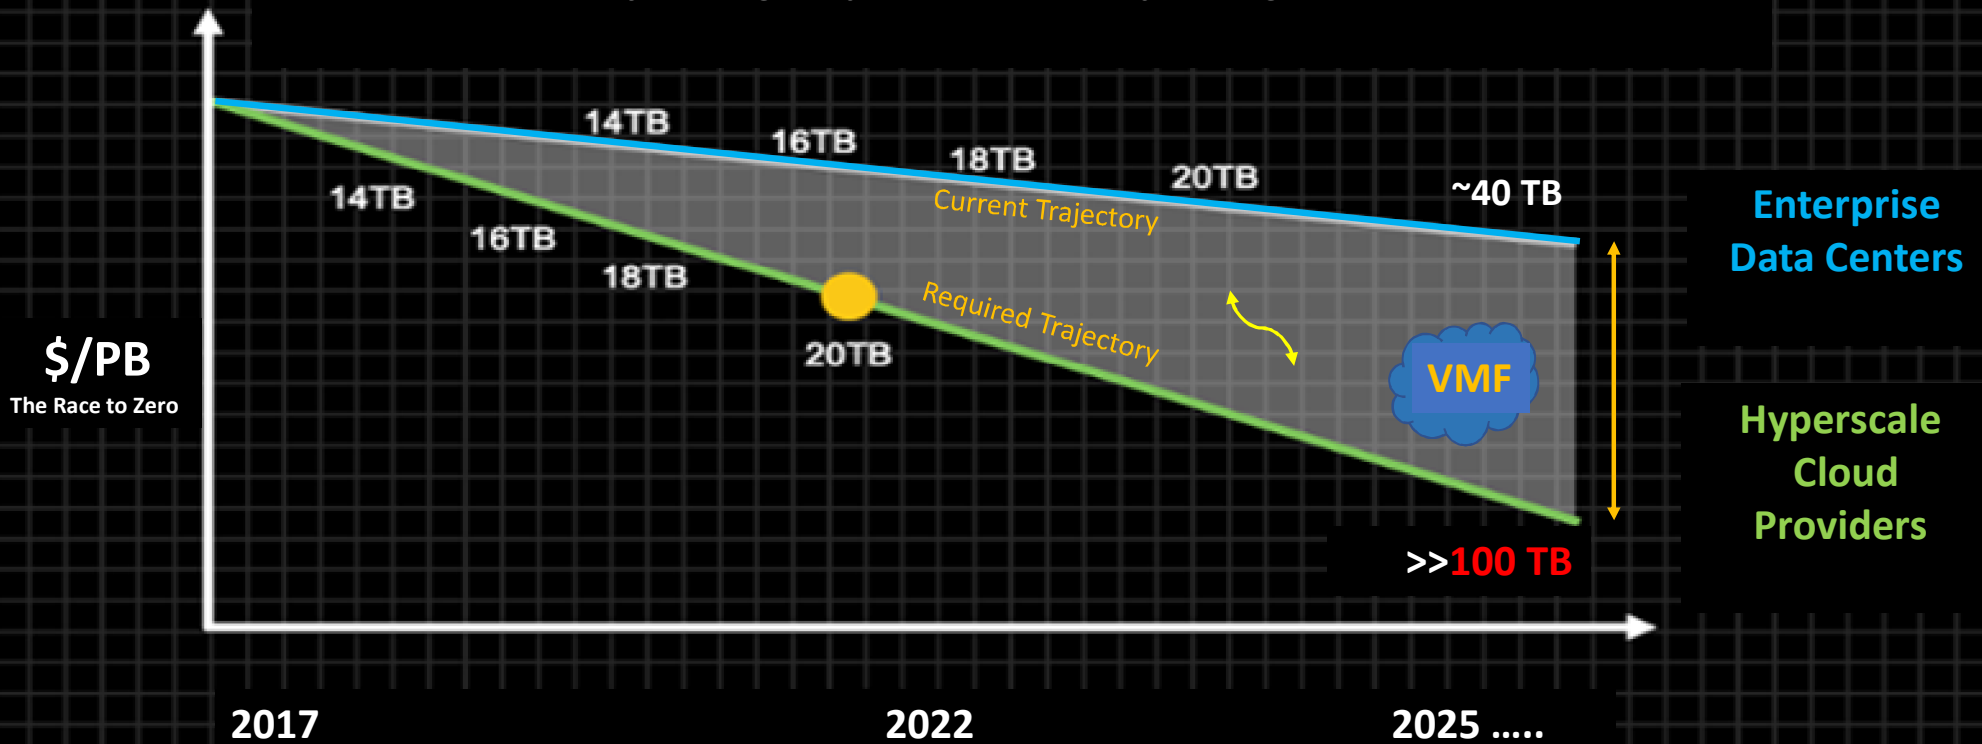


- “ RAIL Provides Parallel Access and Data Transfer to Drives in Multiple Libraries
- “ Erasure Codes Provide ~50% Space Saving Versus HDD RAID
- “ Ideal for Hyperscale and Large-scale and Cloud Object Archives

Source: Horison Information Strategies.

An Impending **VMF** – Can It Be Averted?

➔ A **VMF** (Vertical Market Failure) in Secondary Storage Exists When The Underlying Storage Infrastructure Becomes Insufficient to Address Secondary Storage Market Demands. Secondary Storage represents >80% of all Digital Data.



VMF (Vertical Market Failure) Risk Factors

To the Current Secondary Storage Model


- “ The Zettabyte scale secondary storage market (archive) has become the exclusive domain of few suppliers.
- “ IBM is the only (1) tape drive developer/supplier controlling tape ecosystem specifications.
- “ Fujifilm and Sony are the only (2) LTO tape media suppliers.
- “ HPE, IBM, Quantum and Spectra are the primary large-scale tape and library suppliers.
- “ Seagate, Toshiba and WD are the only (3) remaining HDD suppliers.
- “ The HDD and tape development roadmaps are slowing.
- “ HSDCs leverage their bargaining and buying power to drive down prices impacting vendor margins, R&D.
- “ In the event of a tape (or another energy efficient technology) based VMF, sustainability challenges will become insurmountable for HDDs to address.
- “ The race to zero (\$0.00/TB) places supplier margins, future R&D funding, roadmaps, and innovation at significant risk.
- “ *Can Current Storage Technologies Overcome Their Challenges? **Can Flash Help?***
- “ *Will a New Novel Technology Solution Arrive? (Photonics, Ceramics, DNA, Glass, Holographic, new tape format?)*



Tape Recording Technology Roadmap


Steady Innovation for Drives and Media Outlined

Beyond BaFe and SrFe: Epsilon Ferrite



Epsilon Ferrite enabled by Focused Millimeter Wave-Assisted Magnetic Recording (F-MIMR)*

The present research was supported in part by the Advanced Research Program for Energy and Environmental Technologies/Development of a millimeter wave assisted magnetic recording method for magnetic tapes project by Ohkoshi Laboratory, The University of Tokyo/Nakajima Laboratory Osaka University/Recording Media Research Laboratories. Fujifilm commissioned by NEDO of METI.



PARTITIONING ENABLED LTFS | ENCRYPTION | WORM

NOTE: Compressed capacity for generation 5 assumes 2:1 compression. Compressed capacities for generations 6-14 assume 2.5:1 compression (achieved with larger compression history buffer).

SOURCE: The LTO Program. The LTO Ultrium roadmap is subject to change without notice and represents goals and objectives only. Linear Tape-Open LTO, the LTO logo, Ultrium and the Ultrium logo are registered trademarks of Hewlett Packard Enterprise Company, International Business Machines Corporation and Quantum Corporation in the US and other countries. Please contact your supplier/manufacturer for more information.



Hewlett Packard Enterprise Company, International Business Machines Corporation and Quantum Corporation collaborate and support technology specifications, licensing, and promotions of LTO Ultrium products.

Fujifilm and IBM Announce TS1170 50TB High-Density Tape Storage System

Announced: August 27, 2023 Available August 29, 2023

The IBM TS1170 storage system represents the world's highest cartridge capacity ever announced.

Capacities with 50 TB in a single cartridge reach 150 TB with 3:1 compression.

250% increase in capacity as compared to the previous IBM TS1160 tape drive at 20 TB cartridge capacity.

Improved areal recording density (gb per in²) and a 15% longer tape cartridge enable the capacity increase.

Uses fine hybrid Strontium Ferrite (SrFe) magnetic nanoparticles.

The *sustained* data rate is *unchanged* from TS1160 at 400 MB/sec. *Burst data* rates have reached 1,600 MB/sec.

The IBM TS1170 tape drives consist of two new models:

- 1) the TS1170 Model 70F with a dual-port 16 Gb Fibre Channel interface
- 2) the TS1170 Model 70S with a dual-port 12 Gb SAS interface.

Other features include:

- 1) RAO (Recommended Access Order) which improves recall time and time to first byte (=<83%).
- 2) IBM Storage Archive (LTFS format) for direct, intuitive, and graphical access to data.

Note: This an Enterprise tape announcement. The LTO family is not affected by this announcement.



High-Capacity Tape Drive Roadmap

Tape Continues to Push Capacity Limits

LTO Roadmap



LTO-9 (2021)

18 TB and 400 MB/sec
1 generation backward R/W
BaFe

LTO-10 (2025)

Up to 36 TB and 500 MB/sec
Downward Write/Read
Technology Dependent
BaFe

LTO-11 (202x)

Up to 72 TB and 800 MB/sec
Downward Write/Read
Technology Dependent
BaFe

3592 Roadmap



TS1160 (2018)

20 TB, 400 MB/sec
Media Based Downward
R/W Compatibility
BaFe

TS1170 (8/27/2023)

50 TB* and 400 MB/sec
No Downward Compatibility
SrFe

TS1180 (202x)

Up to 100 TB and 800 MB/sec
No Downward Compatibility
SrFe

All Capacities are Native (non-compressed)

BaFe – Barium Ferrite

SrFe – Strontium Ferrite

* Record Tape Cartridge Capacity

Tape Performance Accelerates

Access Time (Time to First Byte)

Throughput (Data Rate)

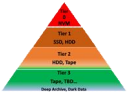
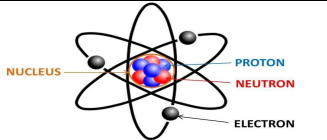
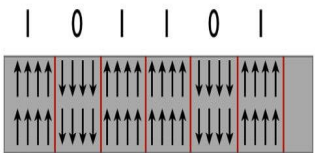
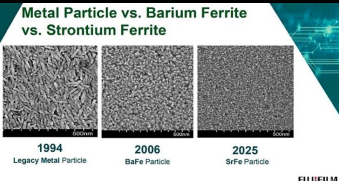
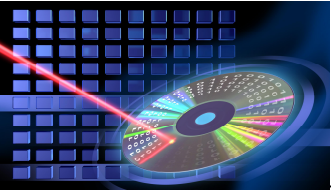
Robotic Mount Time	Drive Load Point Time	Drive/File Access	Total Access Time Range To 1st Byte of File	Throughput/Data Transfer Rate
Locate and Mount Cartridge In Tape Drive	Time to Reach Cartridge Load Point	Time to Locate the File or Object on the Cartridge	Σ Robotic+Load+Access Time	Sequential Data Transfer Rates/Drive LTO-9 400 MB/sec. 24 GB/hour TS1160/70 400 MB/sec. 24 GB/hour
4 - 10 secs	Up to 11 sec	10 - 100 secs	25 – 121 secs	Data Rate and Throughput Improvements

Access Time Improvements for LTO tape

Active Archive (ms)	Provides HDD or SSD-like Cache Access Time to 1 st Byte of Tape Files in ms. (cache hit ratio ~60-90%)	Data Rates Today	Tape Data Rates are 1.5-2.5x Faster Than HDDs.
oRAO – Open Recommended Access Order (Enterprise and LTO-9)	oRAO Orders Tape Requests to Optimize Tape Movement Time to 1 st Byte. Reduces Drive and Media Wear. Up to 73% Improvement.	Faster Data Rates Tomorrow	Projected to be as Much as 3.5X Greater than HDD by 2025.
LTFS - (Enterprise and LTO)	Partitioned Self-describing Open File System to Drag and Drop Files for Faster Access.	RAIT	Striping Data Multiplies Tape Drive Data Rates. Increases Availability With Fault-tolerance.
Faster and Smarter Robotics	Sorting Move Commands and Optimizing Move Sequence-Based on Robot Location. Faster Robotic Movement, Multi-Media Support, Improved Reliability.	RAIL	Increases Availability and Stripes Data Across Drives in Different Libraries and Geographic Locations.

Source: Horison, Inc.

From Here to Where...?

Current Technology	Recording Technique	Roadmap Capability	Future Developments	Challenges
 <p>HDD</p>	 <p>Magnetic Field</p> 	Performance limited and capacity growth slowing	HAMR, MAMR, (? Tb/in ²), more-platters (9-11), zones, 2-4 actuators, bit patterned, ordered granular, cold HDD	Access density (IOPs/TB), TCO, high energy consumption, \$/TB/watt, CO ₂
<p>Optical Disc</p>	 <p>Reflective Spot</p> 	Minimal progress compared to magnetics, <u>Not</u> presently a data center technology	Photonic (fluorescent) multi-layer recording has most potential for optics, EMP proof media	Price, performance, capacity, reliability, throughput, slow learning curve

Source: Horison Information Strategies

The Optimal Secondary Storage Strategy Currently Favors Tape

Tape and HDD Likely to Share the Secondary Storage Load Indefinitely

Function	Benefits Analysis	<i>Tape Repositions for Growth Phase</i>
Energy, CO ₂ Sustainability	Tape Uses Much Less Energy and Has Much Lower Carbon Footprint Than HDDs (~85% Lower).	
Capacity	LTO-9 Cartridge Capacity @18 TB (45 TB compressed) and TS1170 @50TB with 400 MB/sec Data Rate. Smart Exabyte ⁺ Capacity Libraries are Available. Lab Demos Reach 580 TBs (SrFe) per Cartridge.	
Portability	Tape Media Easily Portable in Case of Disaster, HDDs More Difficult to Physically Move.	
Durability/Media	LTO Reliability BER (1×10^{20}) Surpassed HDDs (1×10^{16}), Media Life >30 Years for all Modern Tape.	
Open Standards	SW (S3 API) Support for Tape Object Storage. LTO and LTFS Provide Open Standard File Interface, APIs.	

Will Any New Emerging Technologies Arrive?

A New Storage Frontier Unfolds – But What and When?

*Future Success
or VMF...?*

580+ TB Tape

Photonics

Ceramics

DNA

Glass

3-D Holograph

TBD...

100 + years



Existing



Future



New Format?



Fluorescent – 3D



Matrix


Glass nano-layers



Molecular



Nanostructure



3D refraction




Tape Library
(secs. - mins)



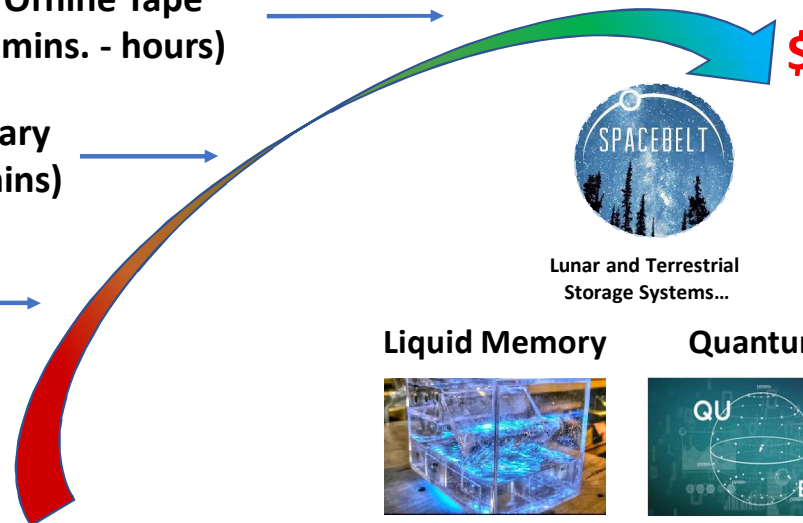
HDD
(ms.)



SSD/NVM
(micro secs. [?])



Offline Tape
(mins. - hours)



Lunar and Terrestrial
Storage Systems...

Liquid Memory



Quantum



Data
Preservation
Time

~5 years

Access Time
Performance

Fastest

Highest

Price per Unit of Capacity

Lowest

Source: Horison Information Strategies.



❖ Horison Information Strategies

- ❖ Home page: **horizon.com**
- ❖ Contact us: **fmoore@horison.com**

❖ IEEE SCV Magnetism

- ❖ Home page: **scvmag.org**
- ❖ Contact us: **secretary@scvmag.org**
- ❖ YouTube channel: **http://channel.scvmag.org**
not http**s**://