
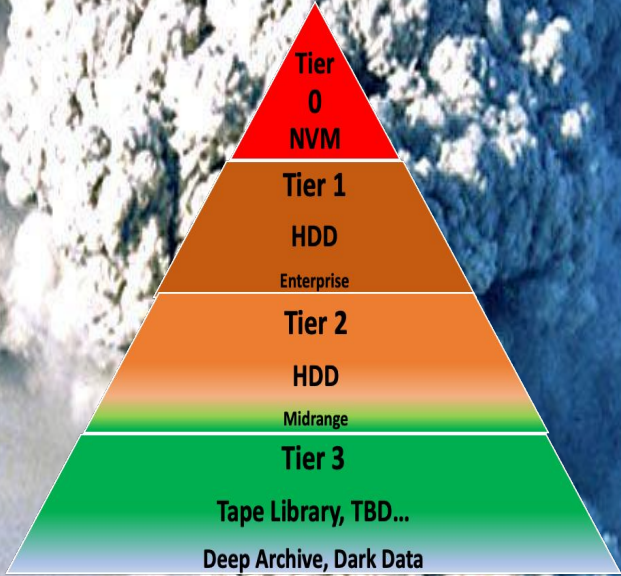


Storage Outlook for the Zettabyte Era And the *Rise* of Secondary Storage



June 29, 2023, 2-3pm MDT



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

~175 ZBs to be Generated
75% Generated Outside the Data Center
~11.7 ZB to be Stored

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



Fueling the Shift to Secondary Storage

Key Trends - 2025 and Beyond

DEFINITION

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.

- ~8.4ZB Stored on SSD, HDD and Tape by 2025 (cagr. 25-30%)
- ~80% of All Data Stored is Archival (Seldom Accessed, Cold, Dark)
- Over 80% of Data Created is Unstructured
- The Archival Copy is Usually the Only Copy of Data
- Retention Periods Over 100 Years are Common
- AI, ML, Big Data Analytics, Virtual Reality, Augmented Reality, Gaming and Robotics Fill and Activate Archives
- Data Protection Critical as Cybercrime Reaches \$23.84 T by 2027. [Statista](#)
- The Active Archive Becomes a De-facto Standard Tier
- Expect a New Secondary Storage and Deep Archive Tier to Arrive
- The race to Zero (\$0.00/TB) continues.


Hyperscale Data Centers Reshape IT Landscape

Heavy Influence on 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 growth.
- Space, full automation, self healing, infinite bandwidth, seamless (PBs to EBs) scalability are critical.
- HSDC secondary storage architecture implements erasure coding, RAIL, geo-spreading.
- HSDC demand curves surpassing all storage limits – fueling a VMF?

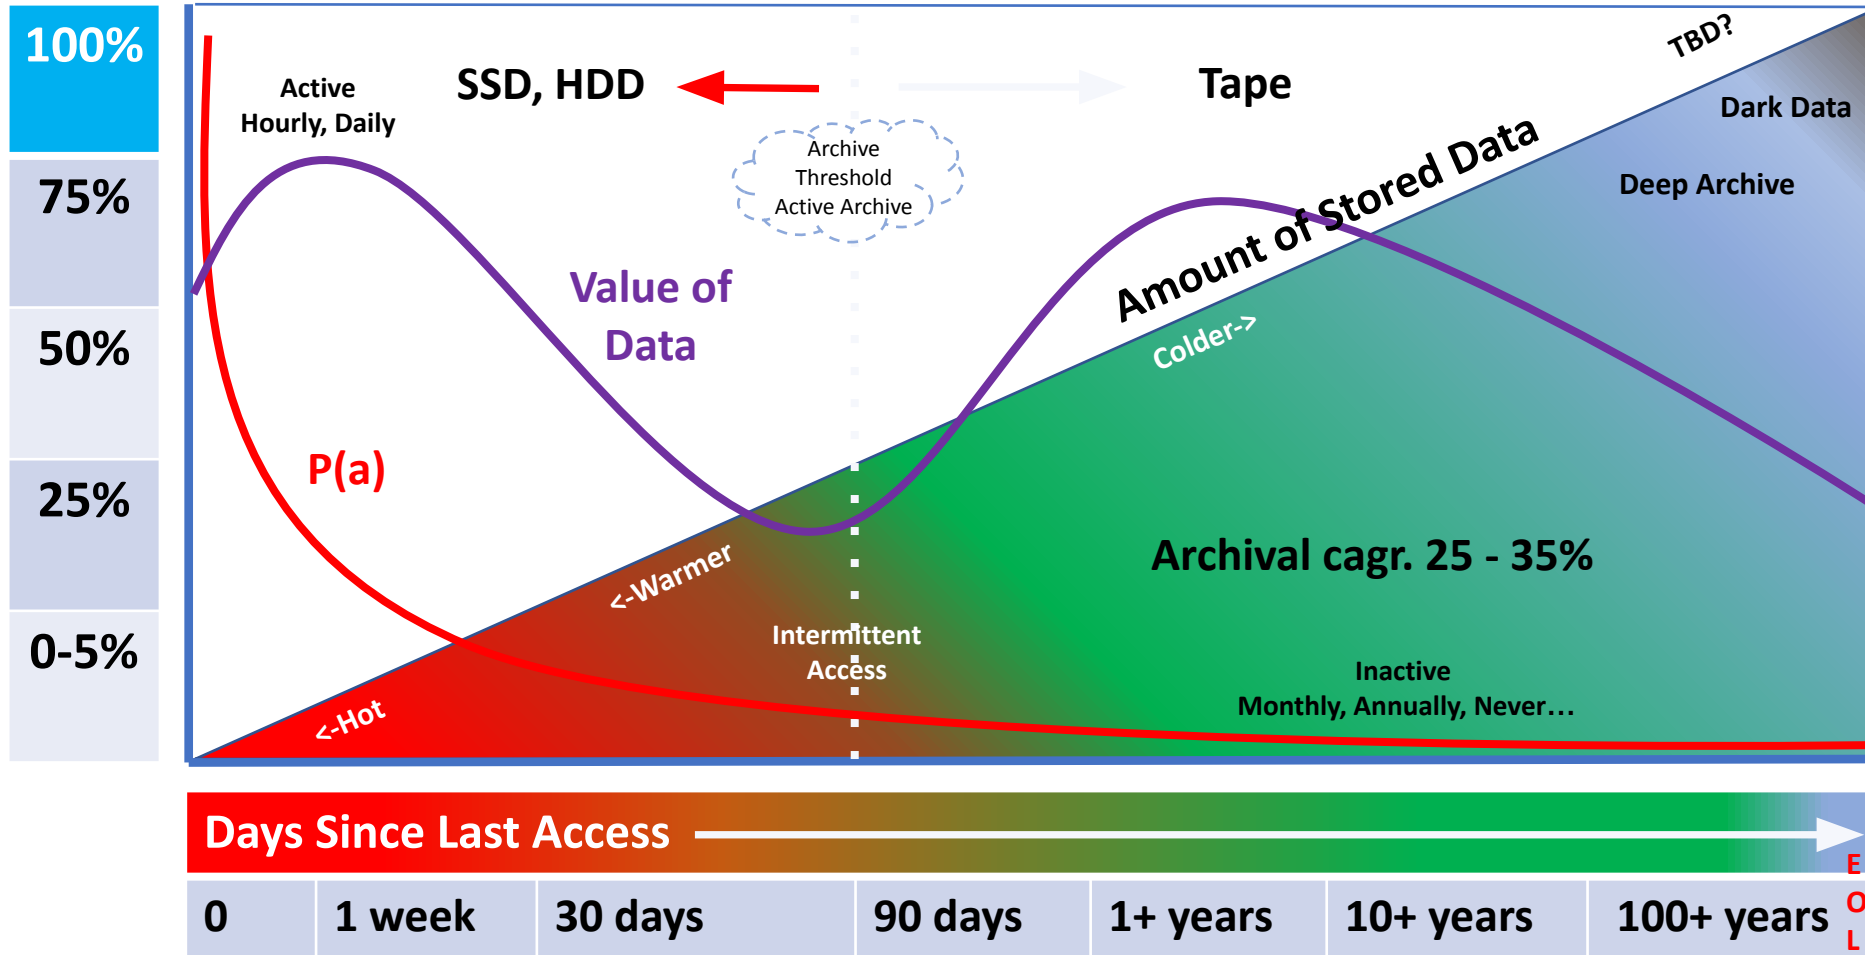
Digital Data Lifecycle

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

Archival data piling up faster than it is analyzed

Seldom Backed Up (1 copy)

The **value of data** can vary over time

Archival retention can be >100 years to ∞

The Tiered Storage Model

By 2025 a New Secondary Storage Paradigm Begins to Emerge – HDD, Tape or ?


Enterprise Data
By 2025

Source: TrendFocus

~8.4 ZB Stored

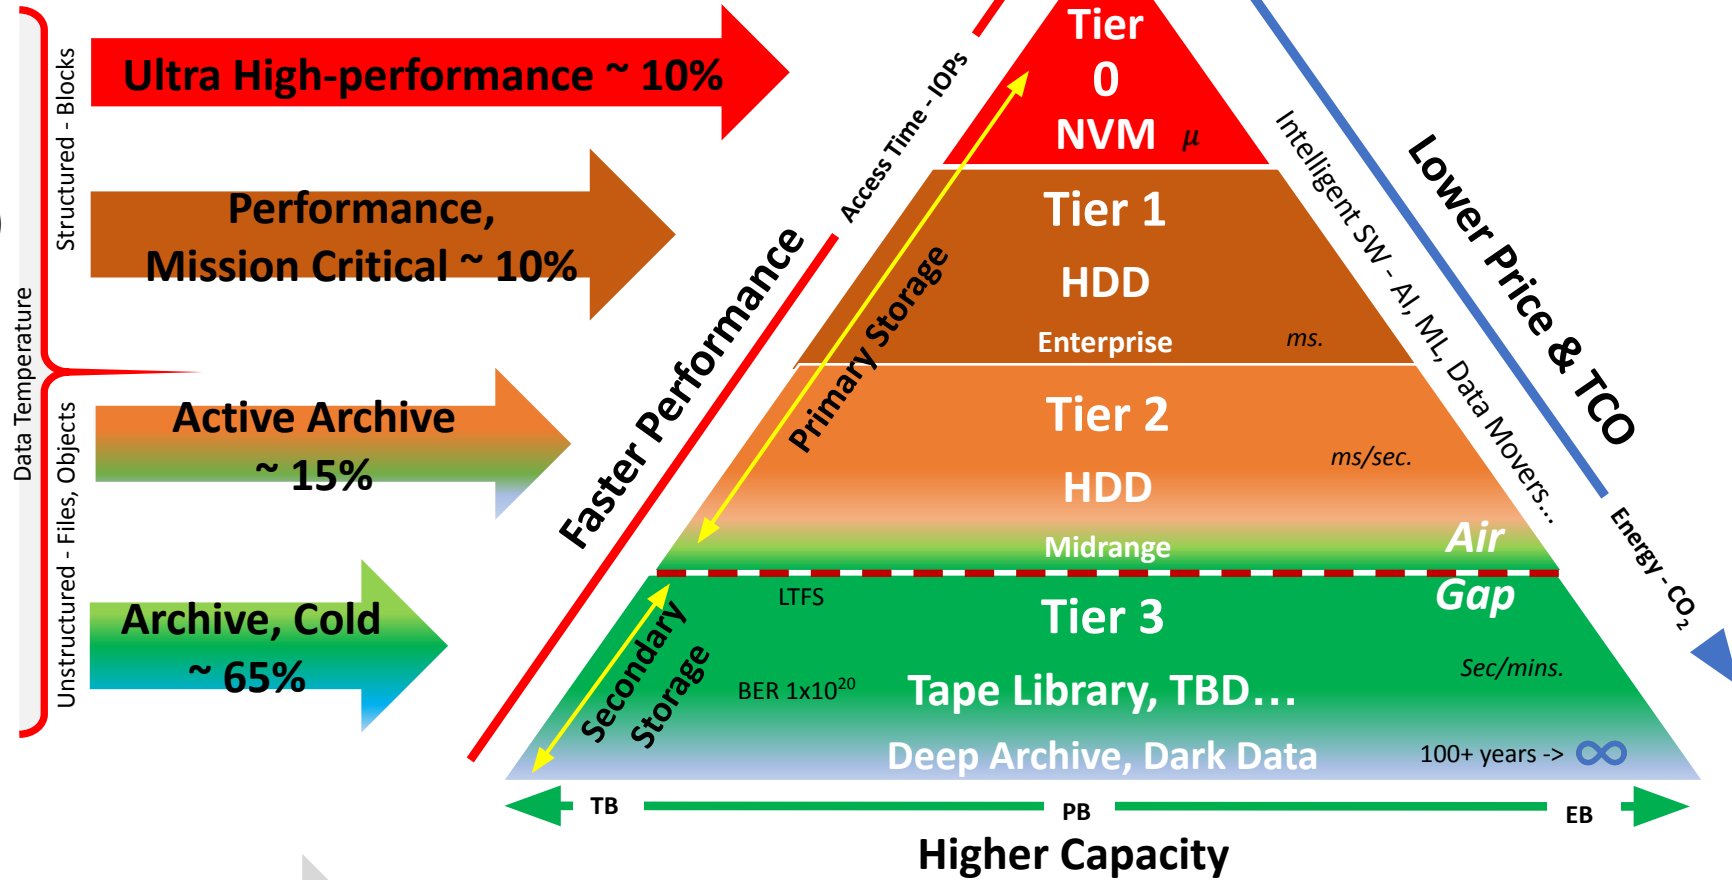
~94%
of data
generated is
transient

Intelligent
Software (AI, ML)



Legacy, Unknown, Orphaned Data

Optimized Data Classification Stored by Tier
% of All Data



Tertiary
Storage



Source: Horison Information Strategies

After 2025 a New Secondary Storage Model Begins to Emerge

New Tiers Target the Archival Avalanche

By 2025

- ~8.4 ZB total enterprise data stored
- >60% stored in HSDCs
- ~20% of stored data is active
- ~80% (~6.7 ZB) of stored data is cold/archival
- Active archive becoming a standard tier
- Majority stored on the *wrong tier*

Immutable Data Properties

Archival Data is Mostly Write Once*

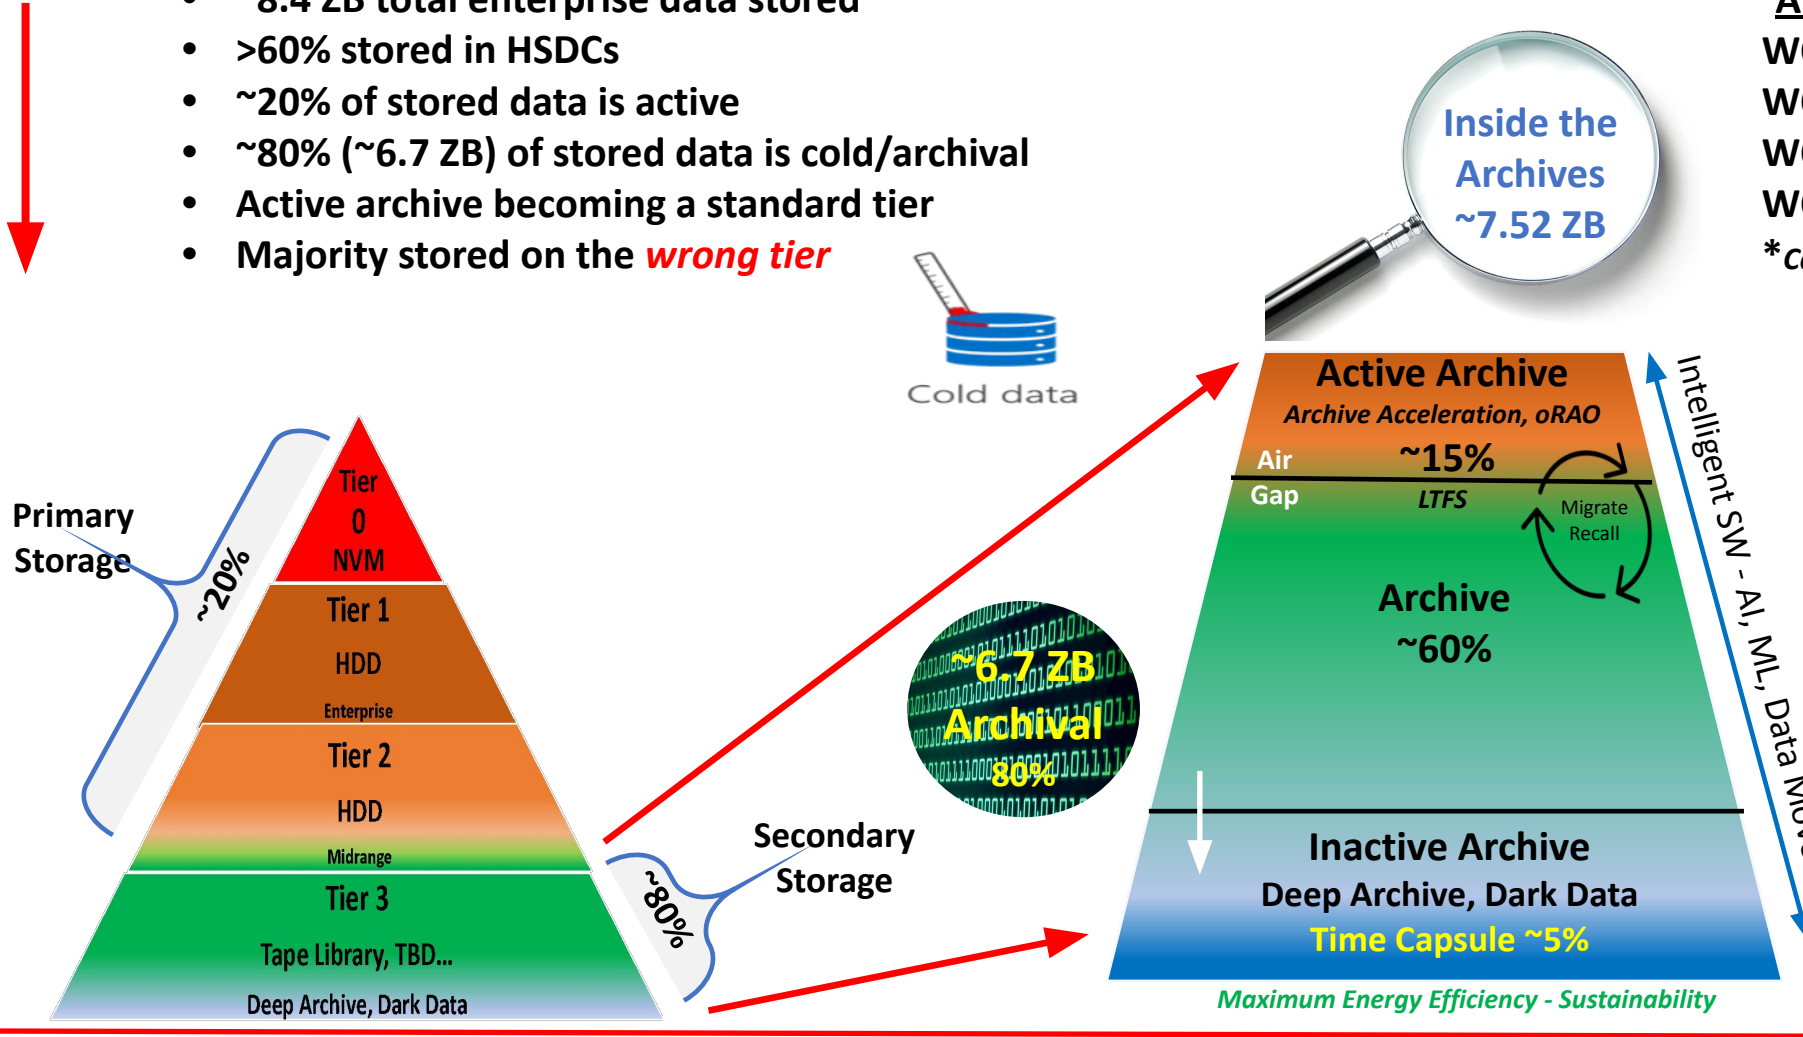
WORM – Write Once, Read Many

WORSE -- Write Once, Read Seldom

WORN – Write Once, Read Never

WORF -- Write Once, Read Forever

*Can't be deleted, modified or overwritten



Optimal Long-term Solution

Nearline HDD - Active Archive

Flash SSD – Instant Archive

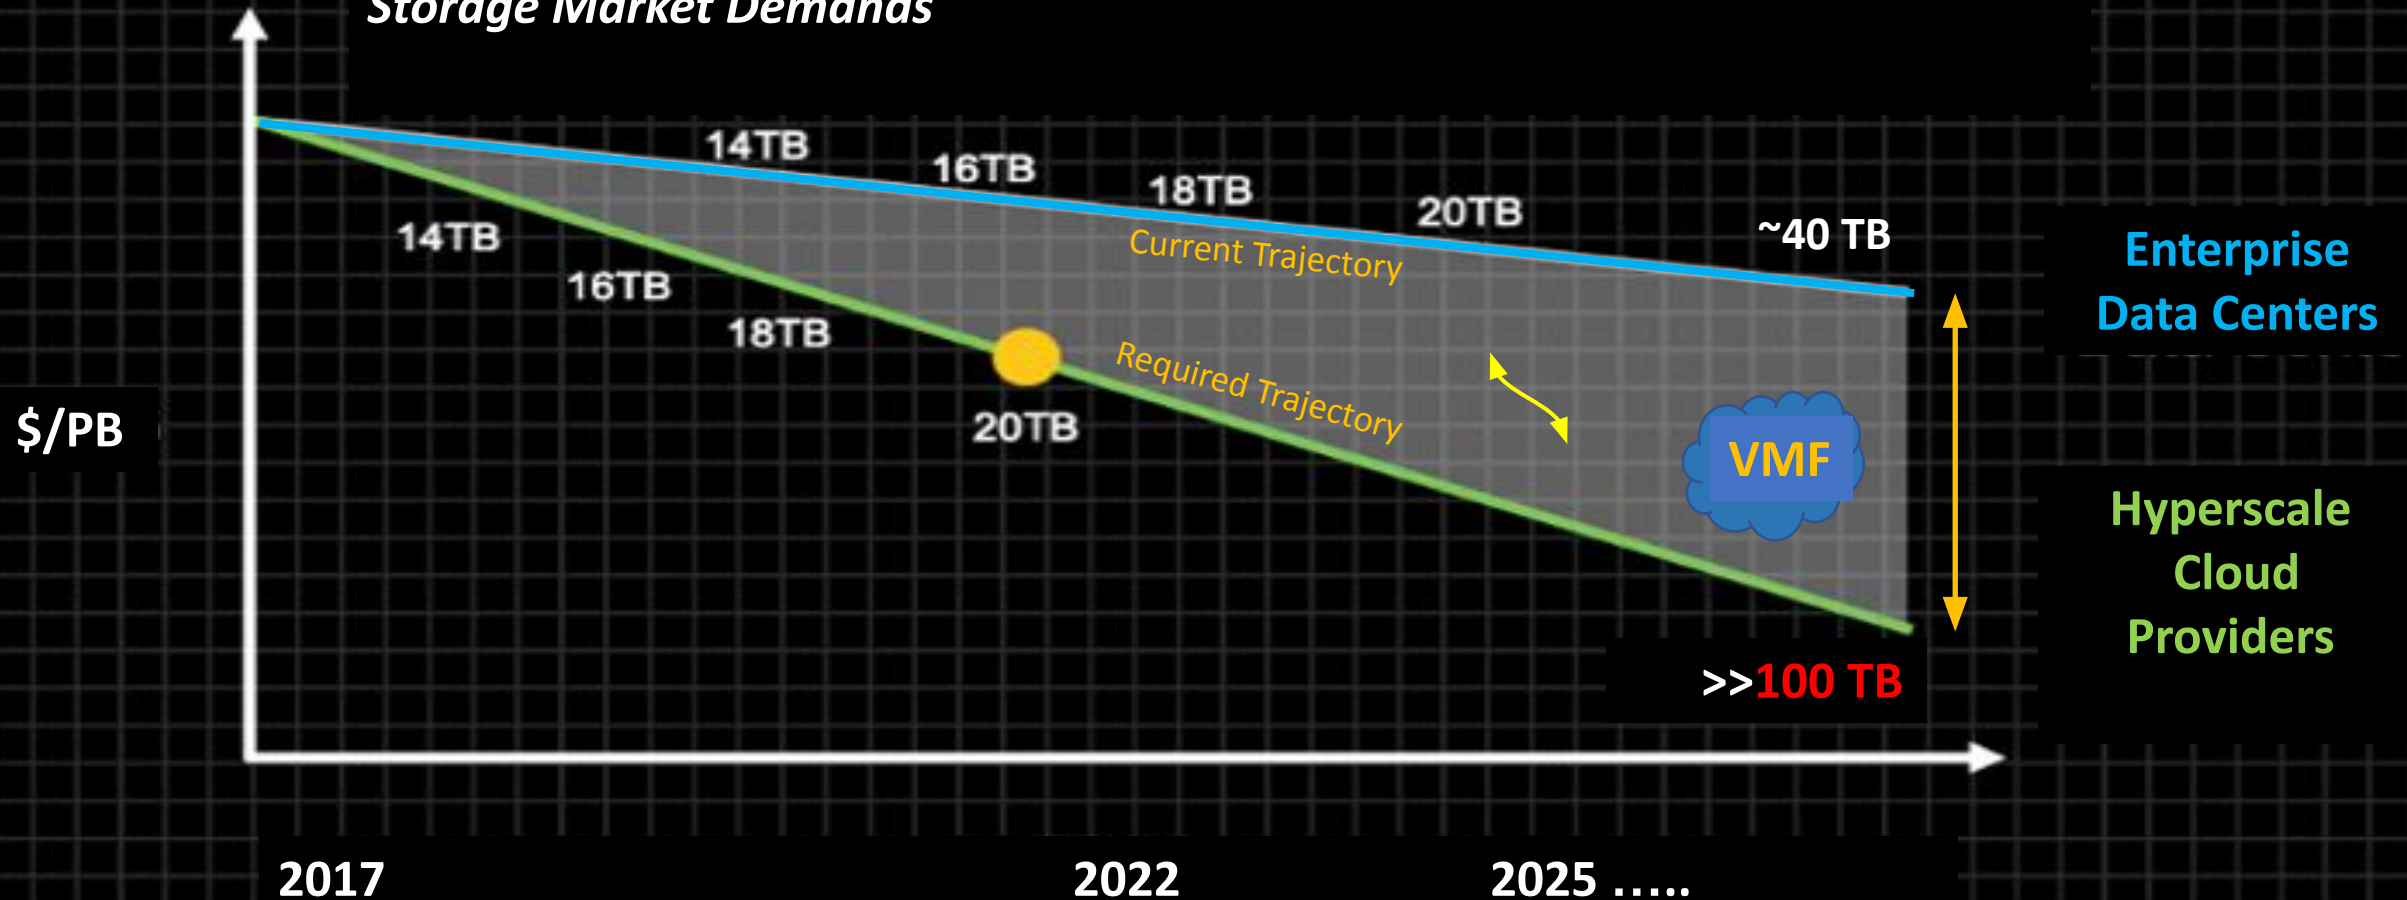
TBD - Optimized Active Archive

Tape Library (PB, EB...)

Tape, Photonics, Glass, 3-D Holographic, DNA, VMF....?

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



VMF Risk Factors

To the Current Secondary Storage Model

- The Zettabyte scale WW secondary storage market has become the exclusive domain of few suppliers.
- IBM is the only tape drive developer/supplier controlling tape ecosystem specifications.
- Fujifilm and Sony are the only tape media suppliers.
- HPE, IBM, Quantum and Spectra are the primary large-scale tape and library suppliers.
- Seagate, Toshiba and WD are the only 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 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 a VMF be Mitigated by Private Sector Market Participants Through Risk-share Agreements (an Alliance) Between Customers and Suppliers?*
- *Can Current Storage Technologies Overcome Their Challenges?*
- *Will a New Novel Technology Solution Arrive?*

Storage Prices Tumble, Inventories Rise as Race to Zero (\$0.00/TB) Accelerates in 2023

TODAY'S TOP NEWS March 16, 2023

Revenue from Enterprise SSDs Totaled \$3.8 Billion for 4Q22, Dropping by 27.4% Q/Q

Due to slumping demand and widening decline in SSD contract prices

TODAY'S TOP NEWS March 15, 2023

Pessimistic View of Near-Term Storage Demand in Asia

One industry that remains in free fall is NAND.

TODAY'S TOP NEWS Feb. 1, 2023

WD: Fiscal 2Q23 Financial Results

[[comments](#)] | Hideous quarter like for Seagate, with revenue at \$3.1 billion, down 17% Q/Q and 36% Y/Y, and net loss

SSD Prices Have Fallen 15 to 30 Percent Since January 2023

By Avram Piltch published 15 days ago. Declining Flash memory prices have lead to great savings.

2TB WD SSD Now Only 5 Cents per GB: Real Deals (The Race to zero \$/TB) March 6, 2023

By Stewart Bendle published 12 days ago. A great price on 2TB of Gen 4 SSD storage.

News March 23, 2023

SSDs are super cheap right now

A downturn in the NAND storage market means great SSD deals for consumers.

WW NAND Flash Revenue Down 25% Q/Q in 4Q22 in **Flash crash**. 8.1% Expected in 1Q23 as ASP Drops Further. The global NAND flash market has been facing a demand headwind since 2H22. In response, the supply chain has been scrambling to clear out inventory, driving down NAND flash contract prices by 20-25%. Enterprise SSD took the brunt of the fall with prices plummeting 23-28%.

Global DRAM Revenue Fell by More Than 30% Q/Q in 4Q22 to \$12.3 Billion

As suppliers made large price concessions to drive shipments

This is a Press Release edited by StorageNewsletter.com on March 29, 2023 at 2:02 pm

This report, published on March 2, 2023

Micron May Have Had Largest Ever Loss in Q2

Company expected to report negative gross margins for the first time in 15 years

By Nathan Reiff Published March 28, 2023

TODAY'S TOP NEWS. April 14, 2023

HDD Shipments Dip 35% Y/Y to 34 Million in 1Q23

On slight seasonal reductions

Cloud Projects Keep Being Postponed Amid Economic

Uncertainty April 23, 2023

HDDs: Unstoppable Decline in Sales. In this respect, it can already be estimated that there will only be a two-part market in the future, since the Toshiba brand will disappear completely. 4/20/2023

TODAY'S TOP NEWS April 21, 2023 Seagate: Fiscal 3Q23 Financial Results

| Sales down 34% Y/Y and further decreasing next quarter, and much higher net loss

WD: Fiscal 3Q23 Financial Results. May 9, 2023

Sales down 36% Y/Y and 10% Q/Q, bad period following morose HDD and SSD markets.

TODAY'S TOP NEWS

NAND and SSD Market Continue to Struggle in 1Q23

Unit shipments dropped 15% Q/Q to 64 million, exabytes down by 13.2%, to 57.

TODAY'S TOP NEWS. 6/8/2023

WW NAND Flash Revenue Falls 16.1% in 1Q23 at \$8.6 Billion

As demand continues to decline.

WW Enterprise SSD Revenue Falling by Nearly Half Following Slump in Volume and Price. 6/14/2023

Data Protection Strategies

Backup is Important – Recovery is Everything
Protect *All* Data In The Security Ecosystem

The 3 – 2 – 1 – 1 Backup Strategy

Backup, Recovery and DR for Primary (active) Storage


3	2	1	1
Copies of Data	Different Media Types	Copy Offsite ... Cloud, Vault	Air Gap Copy Offline Media
Cyclic Process	(SSD, HDD, Tape) Technology Diversity	(SSD, HDD, Tape) Geographic Redundancy	(Tape) No Electronic Connection

The 2 – 2 – 1 Archive Strategy

Backup, Recovery and DR for Secondary (archival) Storage

2	2	1
Copies of Data	Geographic Locations	Air Gap Copy Offline Media
Periodic Process	<ul style="list-style-type: none"> On-Premise, Cloud, or Vault Replicated or Erasure Coded Synchronous or Asynchronous Geographic Redundancy	(Tape) No Electronic Connection

Immutable Archive Copy – WORM, Encryption

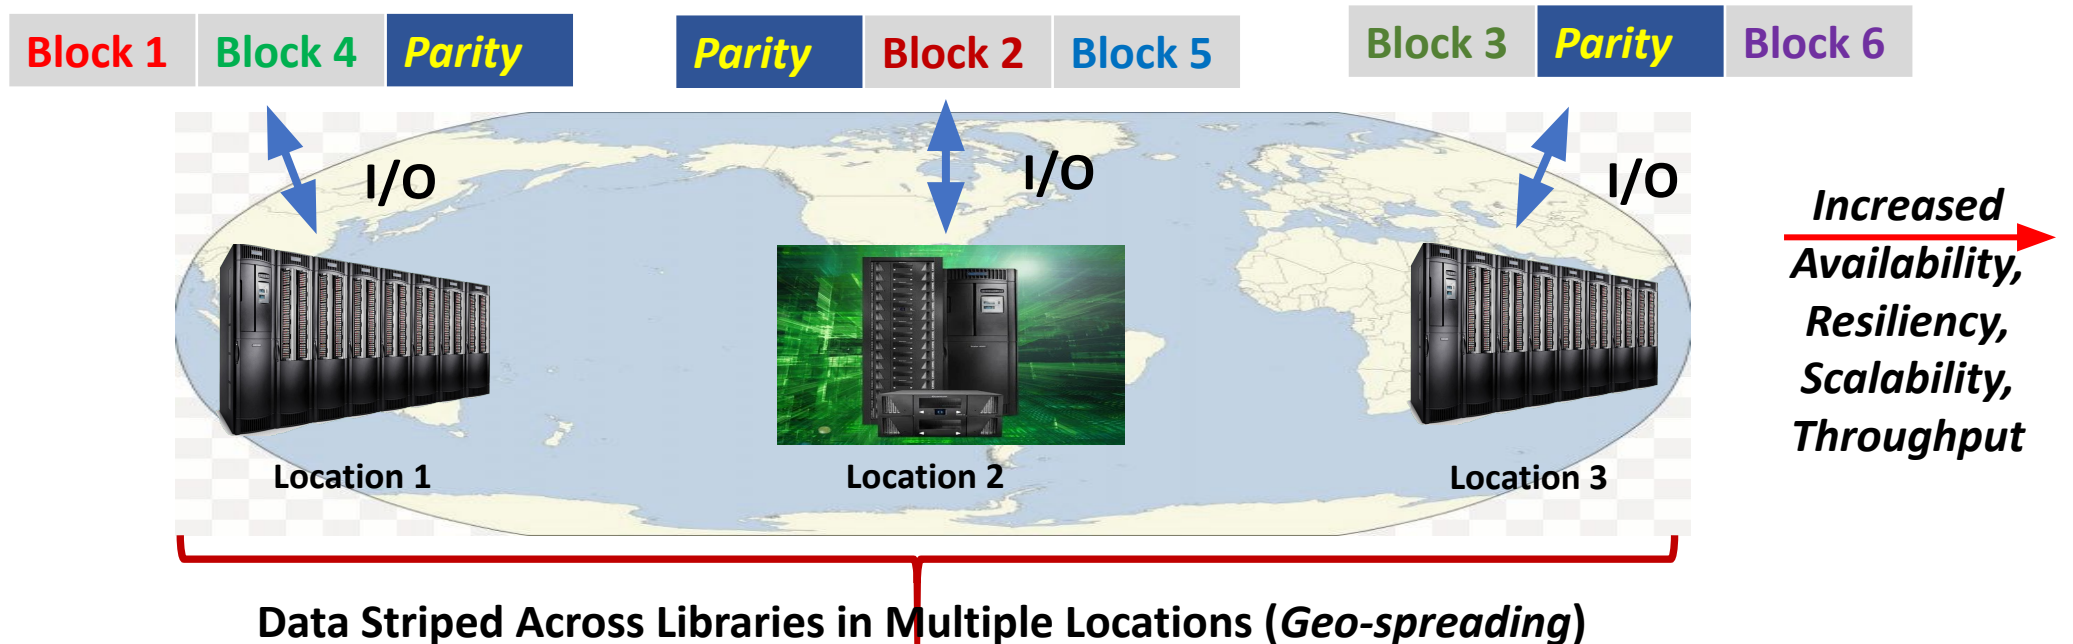


RAIL - Redundant Array of Independent Libraries

Erasure Coding and Geo-spreading 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 Cloud Object Archives

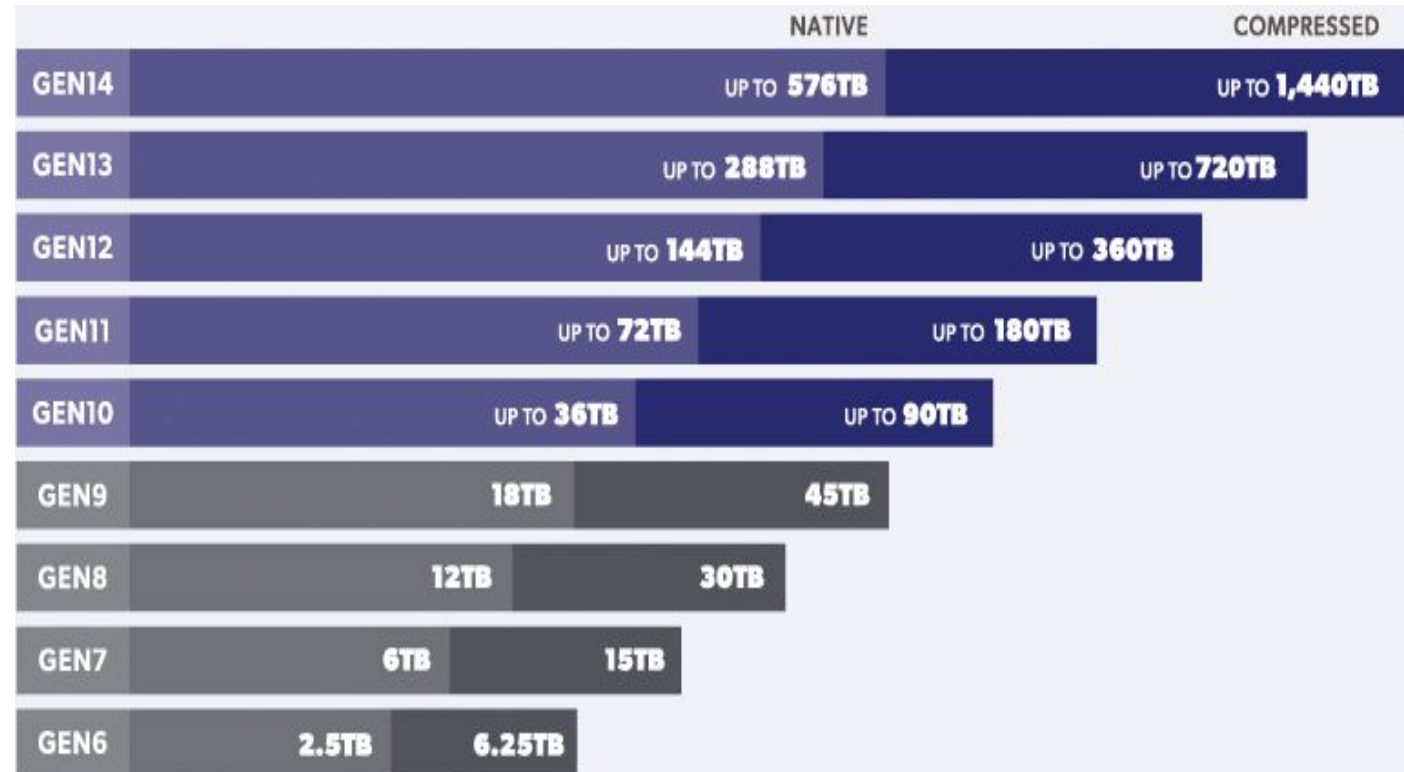
Tape Recording Technology Roadmap

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 (Ohkoshi Laboratory, **The University of Tokyo** / Nakajima Laboratory, **Osaka University** / Recording Media Research Laboratories, **FUJIFILM Corporation**) commissioned by NEDO of METI.



PARTITIONING ENABLED LTFS | ENCRYPTION | WORM

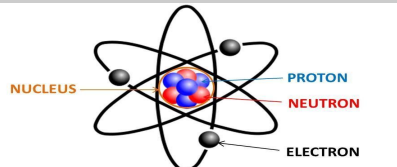
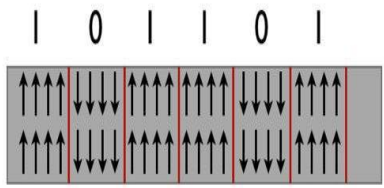
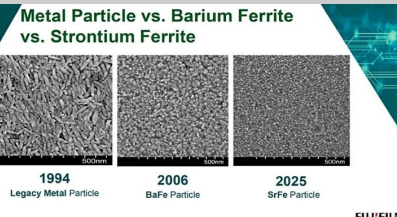

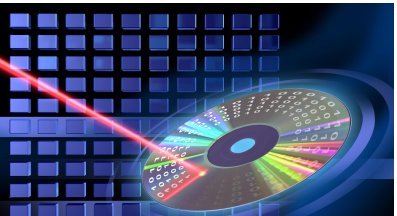
NOTE: Compressed capacity for generation 5 assumes 2:1 compression. Compressed capacities for generations 6-14 assume 2.5:1 compression (achieved with larger compression history buffer).

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.

From Here to Where...?

Current Technology	Recording Technique	Roadmap Capability	Future Developments	Challenges
<p>NVM (SRAM, DRAM, MRAM, NAND Flash, PCM, 3D-Xpoint)</p>	<p>Electronic Charge</p> 	<p>Aggressive development, multiple technologies emerging, CXL, NVMe(oF)</p>	<p>Multi-layer 3D stacking (500+), faster garbage collection, new tiers?</p>	<p>Price</p>
<p>HDD</p>	<p>Magnetic Field</p> 	<p>Performance limited and capacity growth slowing</p>	<p>HAMR, MAMR, (? Tb/in²), more-platters (9-11), zones, 2-4 actuators, bit patterned, ordered granular, cold HDD</p>	<p>Access density (IOPs), TCO, high energy consumption, \$/TB/watt, CO₂</p>
<p>Tape</p>	<p>Magnetic Field</p> 	<p>Well defined and sustainable capacity growth, 580 TB demo, high patent activity</p>	<p>Strontium Ferrite (SrFe), Epsilon Ferrite (ϵ-Fe₂O₃), TMR, deep archive, RAIL, erasure coding, Geo-spreading, fixed tape</p>	<p>Access time, No consumer market, The race to \$0/TB...</p> 
<p>Optical Disc</p>	<p>Reflective Spot</p> 	<p>Minimal progress compared to magnetics, <u>Not</u> presently a data center technology</p>	<p>Photonic (fluorescent) multi-layer recording has most potential for optics, EMP proof media</p>	<p>Price, performance, capacity, reliability, throughput, slow learning curve</p>

The Optimal Secondary Storage Strategy Currently Favors Tape

Tape Function	Benefits Analysis <i>Tape Re-enters Growth Phase</i>
Price/TCO	Tape Has the Lowest Acquisition Price \$/TB, Lowest TCO.
Energy, CO ₂ Sustainability	Tape Uses Much Less Energy and Has Much Lower Carbon Footprint Than HDDs (~85% Lower).
Performance (Access time)	Much Faster Access Times - Active Archives, Fastest Data Rates, Smarter and Faster Robotics, RAIT, RAIL, New Time to 1 st Byte Features (oRAO, TAOS). Re-writable Media.
Capacity	LTO-9 Cartridge Capacity @18 TB (45 TB compressed) with 400 MB/sec Data Rate. Smart Exabyte ⁺ Capacity Libraries are Available. Lab Demos Reach 580 TBs per Cartridge.
Scalability	Tape Easily Scales Capacity (PBs to EBs) by Adding Media/Racks Without Adding Energy Consumption, HDDs Scale Capacity by Adding Drives and Adding Energy Consumption.
Portability	Tape Media Easily Portable in Case of Disaster, HDDs More Difficult to Physically Move.
Cybersecurity	Tape Air Gap, WORM and Encryption Options Protect Against Malware Attacks, Provide Immutability.
Durability/Media	LTO Reliability BER (1x10 ²⁰) Surpassed HDDs (1x10 ¹⁶), Media Life >50 Years for all Modern (BaFe) Tape.
Recording Limits	HDDs Facing Areal Density and Performance (IOPs) Limits. Tape Has a Well-Defined Roadmap (LTO-14).
Open Standards	SW (S3 API) Support for Tape Object Storage. LTO and LTFS Provide Open Standard File Interface, APIs.
Tape and Cloud Ecosystem	Tape Interfaces Seamlessly With Clouds Using Industry Standard APIs to Enable Hybrid Clouds. Native Cloud Applications Can Write To and Read From Tape.

Will New Emerging Technologies Arrive?

A New Storage Frontier Unfolds – But What and When?

Future Success
or VMF...?

580+ TB Tape

Photonics

DNA

Glass

3-D Holographic

TBD...

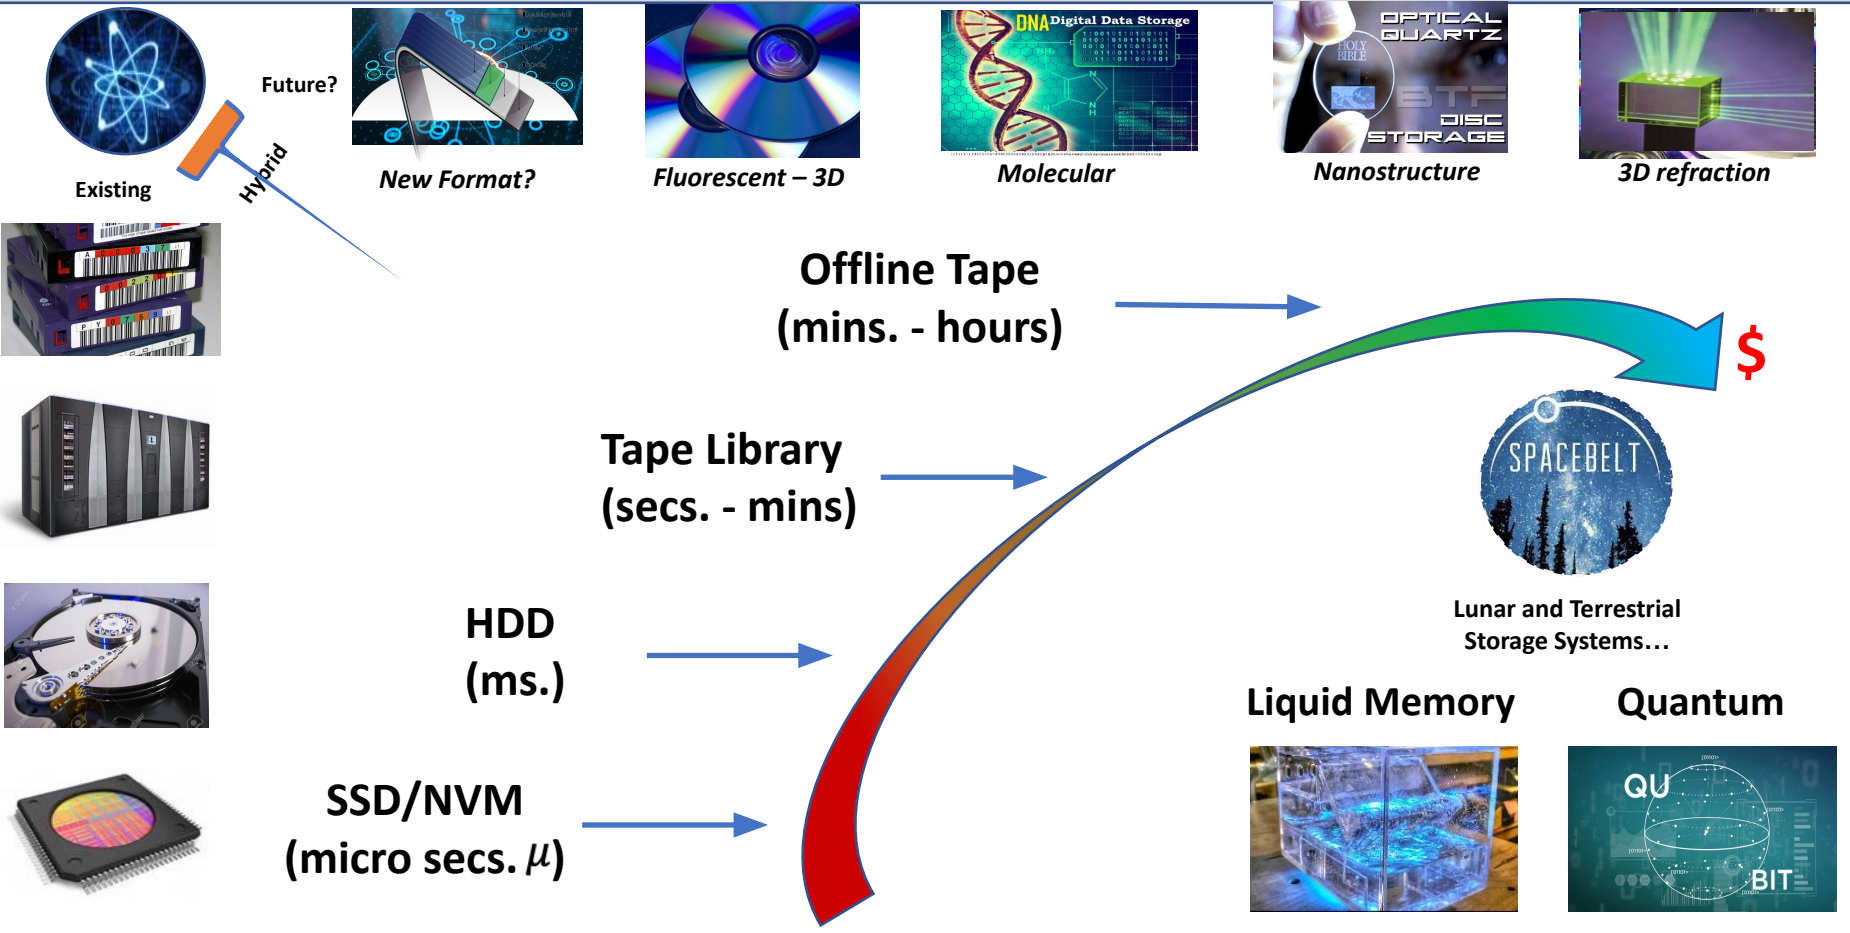
Access Time
Performance

100 + years

Data
Preservation
Time

~5 years

Fastest



Highest

Price per Unit of Capacity

Lowest

The Zettabyte Era and Beyond

Every 3-D System is Bounded by Space, Speed and Time



Saganbyte
 1×10^{33}
 Galactic Data



>186,282 m/s
 FTL

Geopbyte
 1×10^{30}

Brontobyte...
 1×10^{27}

Yottabyte



...
 1×10^{24}

Zettabyte

...
 1×10^{21}



Zettabyte Era and Beyond



You Are Here



Artificial Intelligence

Hyper Cloud

- AI, ML -> Singularity
- Quantum in the Cloud
- Arrays of Atoms
- Each Atom is a Qubit
- 1 Million+ Qubits in a Thumbnail
- Massive Parallelism, Data Fabrics
- HyperScale Compute and Storage
- New Storage Paradigm



SEISMIC SHIFTS



Hydrogen Cells

Energy

SOLAR Electricity POWER WATER
 OIL HYDRO SUNLIGHT TIDAL ENERGY GAS Wind Petroleum PIPELINE
 Hydrogen Geothermal energy

REMEMBER

**Things are Changing so Fast
 Even the Future is Obsolete**