Large-scale Storage

Tape and Archives
For the Zettabyte Era

LIBRARY OF CONGRESS

Designing Storage Architectures

Washington D.C.

March 27-28th, 2023

HORISON Information Strategies

Fred Moore
President
Horison Information Strategies
Horison.com

The Zettabyte Era Takes Off

2016 eta
Zettabyte era
One sextillion

1 ZB = 1,000 EBs = 1x10^{21} bytes

~2043 eta
Yottabyte era
One septillion

One Zettabyte Equivalents

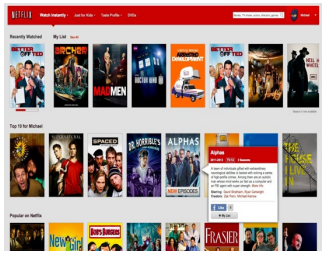

Watching the entire Netflix catalog 3 million times

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)

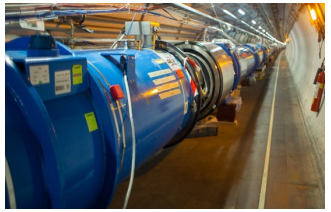

66.7 years of the Large Hadron Collider's experimental data

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, Horison Inc.

Digital Transformation Fuels Secondary Storage

Archival Data Pileup Could Exceed 9.0 ZB by 2025

By 2025…

Up to 11.7 ZB Stored

~80% (~9.3 ZB) of all Data is Archival!

The Archive Copy is Usually the Only Copy of Data

~27.1 B IoT Devices

Deeper Archives
Retention Periods Over 100 Years are Common

AI, ML, Deep Learning Harvest Archives

Active Archive Becomes a De-facto Standard

Expect a New Sustainable Archive Tier/Technology *Or VMF*

Cybercrime Damage to Reach $10.5 T by 2025.
Cybersecurity Ventures

Software Defined "Everything"

~338 Billion Lines of New Software

~30 Million Software Developers

Source: Horison, Inc.

Backup and Archive are Different Processes

Backup is a Copy Process – Multiple Copies of Source Data

Servers and Applications

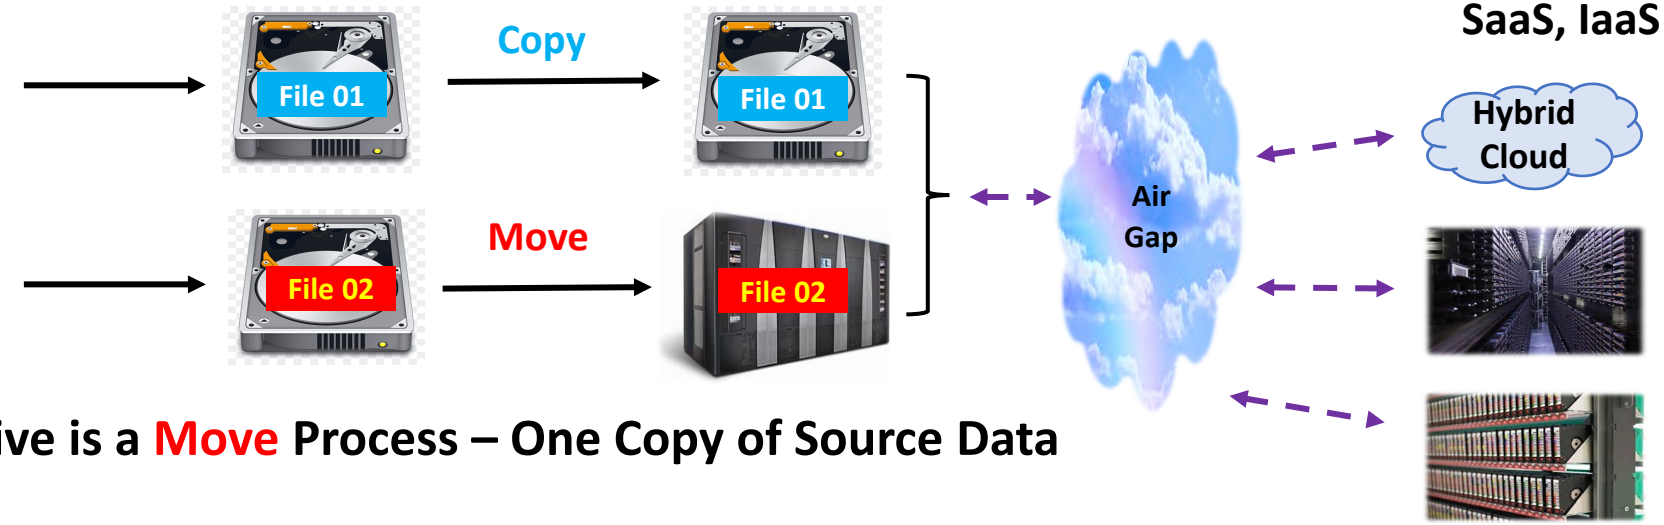

Archive is a Move Process – One Copy of Source Data

Active Archive – Provides *Fast* Access to Archives

Backup	Archive	Active Archive
Copies data for protection and fast recovery (RTO), source data left in place, potential for redundant data, often stored as blocks	Moves infrequently used data to more cost-effective storage, frees space on source devices	A combined solution of intelligent software, SSD, HDD, and tape library systems
Restores files to desired point in time in event of data loss Backup is important – recovery is everything!	Retrieves files for future reference and analysis ⚠ Archive copy is usually the only copy	Provides fast file and object level access to higher activity (dynamic) archival data
Short retention duration ~1-120 days Cyclic process, overwrites itself at end of retention time	Periodic process, forever growing, seldom overwritten - WORM, encryption, air gap	AI, ML and Big Data analytics increase archive activity – for blocks, files and objects

Source: Horison, Inc.

Digital Data Lifecycle

When Does Data Become Archival?

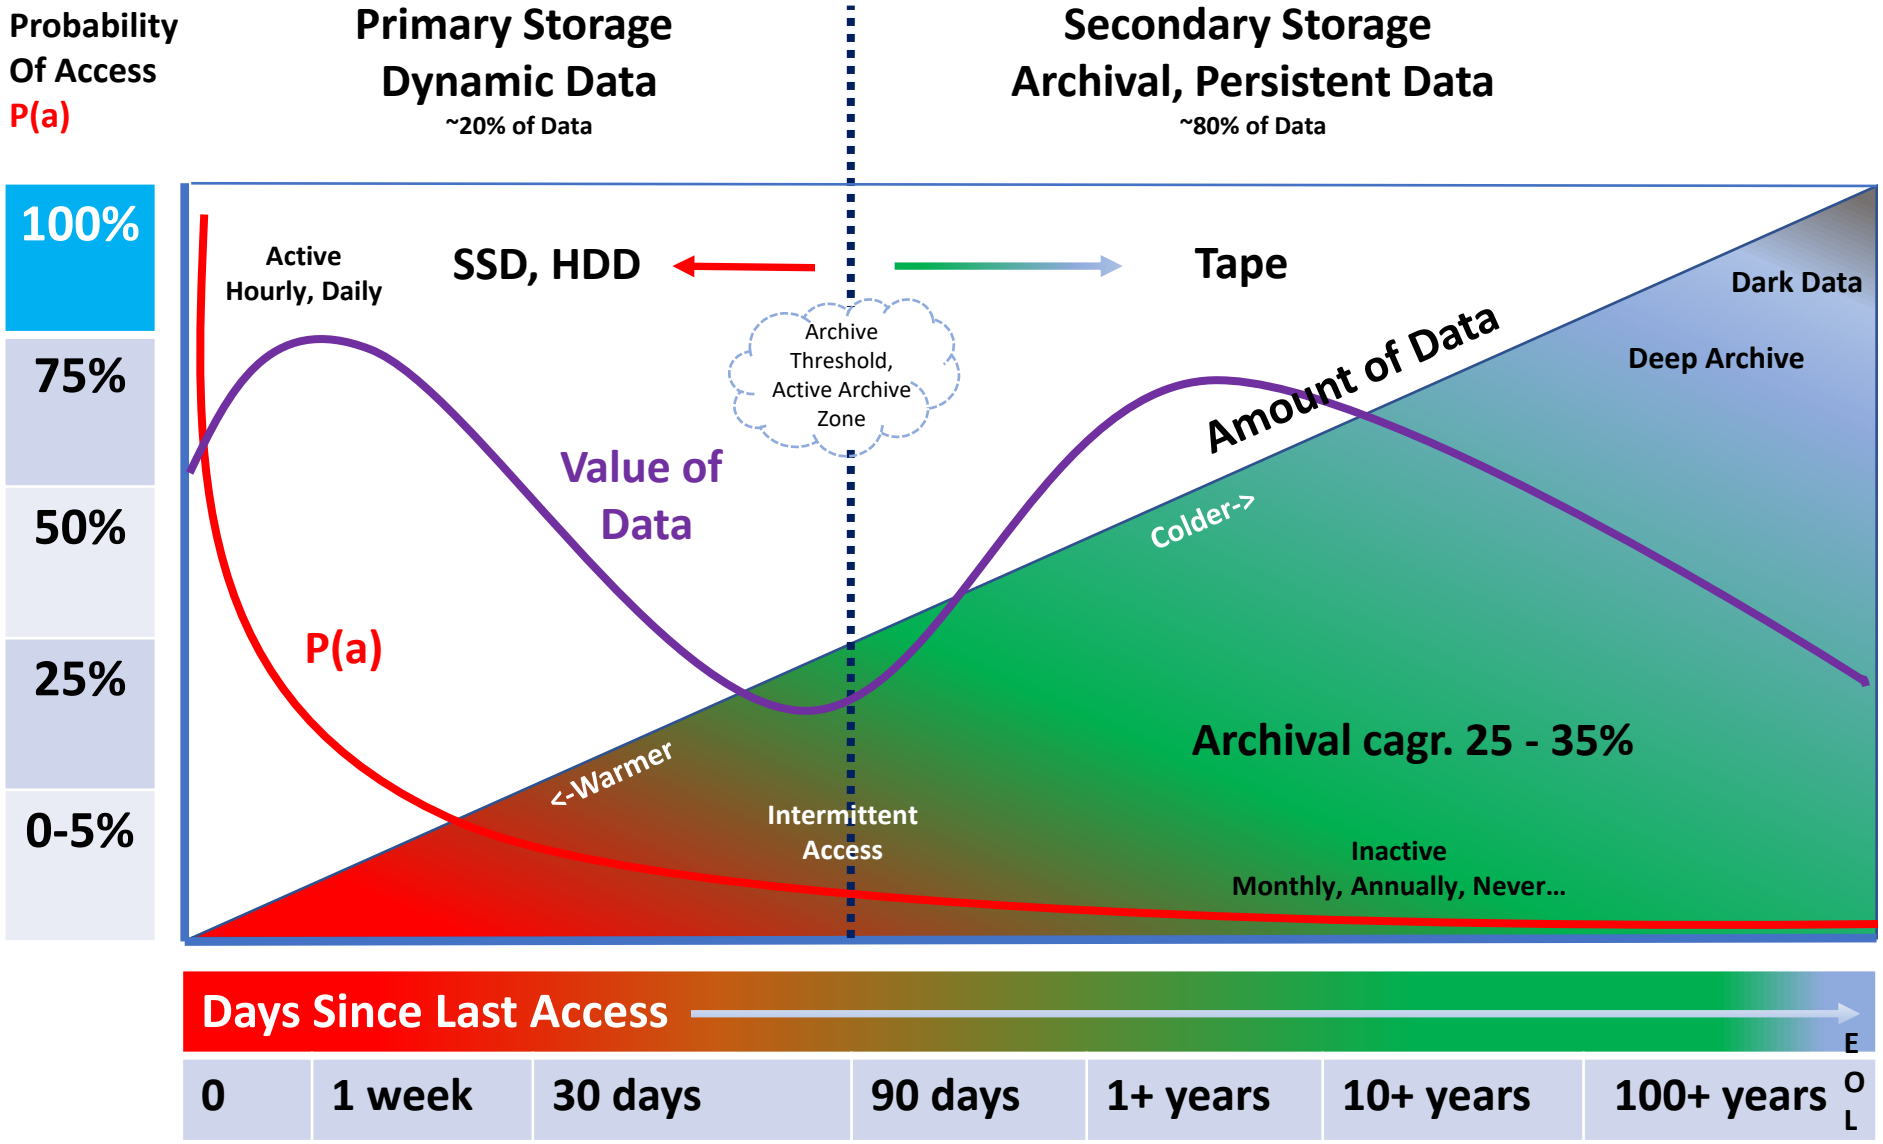

Lifecycle Profile Factors

Probability of access P(a) declines as data ages

The value of data can vary over time

Data typically becomes archival in ~90-120 days

Archival data piling up faster than it is analyzed

Seldom Backed Up (1 copy)

Archival retention can be >100 years to ∞

Source: Horison, Inc.

The Tiered Storage Model

Zettabyte Era Pushes Limits for Secondary Storage

By 2025

Source: IDC

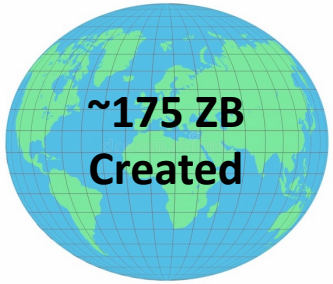

~94% transient data

Intelligent Storage Software (AI, ML)

~11.7 ZB Stored

Legacy, Unknown, Orphaned Data

Optimal Data Classification by Tier
% of All Data

Data Temperature

Blocks - Structured

Files, Objects - Unstructured

- Ultra High-performance ~ 10%
- Performance, Mission Critical ~ 10%
- Active Archive ~ 15%
- Archive, Cold ~ 65%

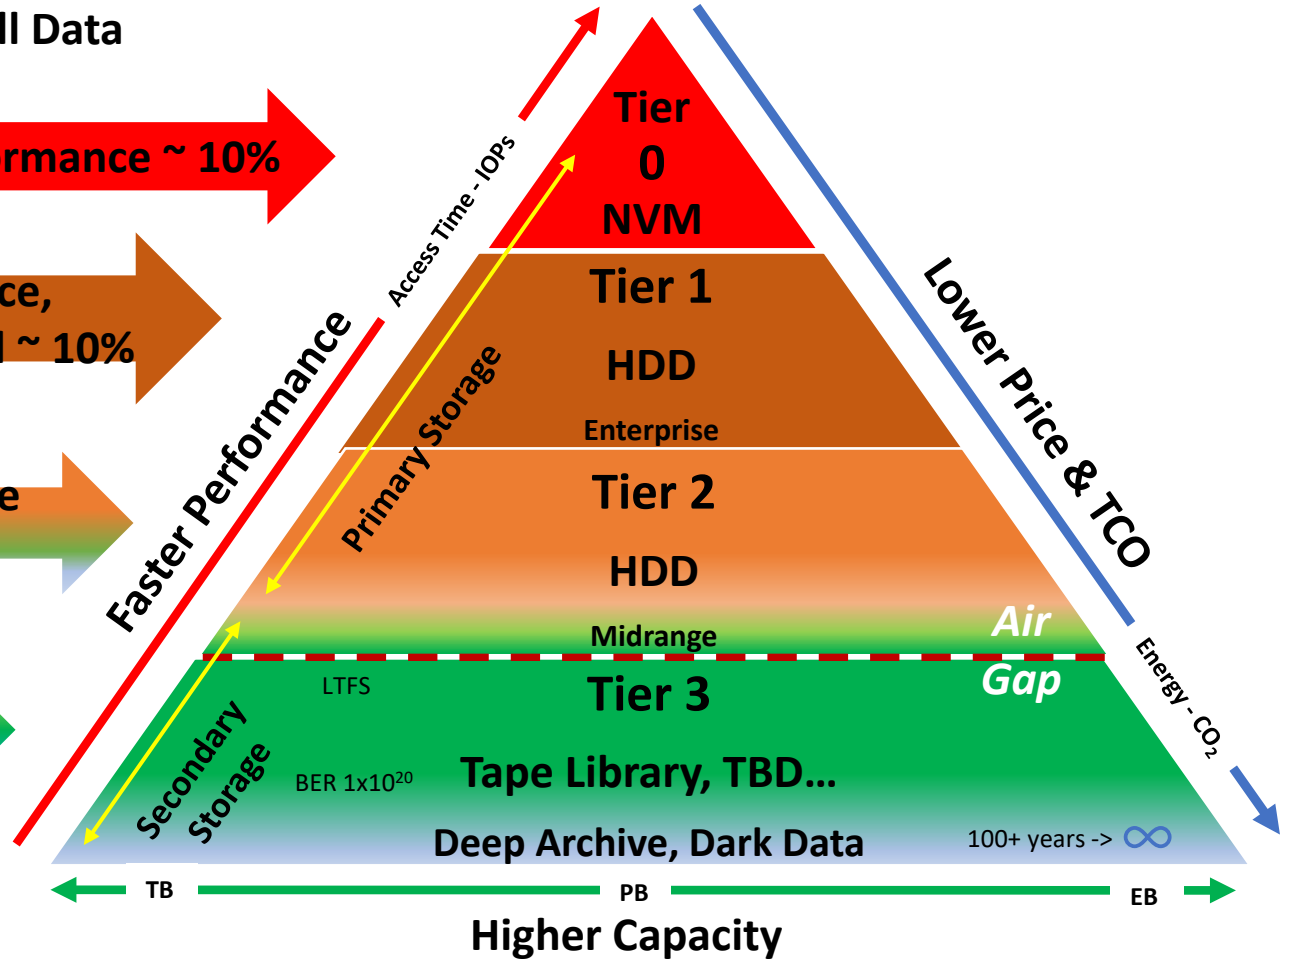

Hybrid Cloud, Air Gap, Vault

The Cloud Uses All Storage Tiers SSD, HDD, Tape

Source: Horison, Inc.

After 2025 a New Secondary Storage Architecture Begins to Emerge

Distinct Tiers Target the Archival Avalanche

By 2025

- ~11.7 ZB total data stored
- >60% stored in HSDCs
- ~20% of stored data is active
- ~80% (~9.3 ZB) of stored data is cold/archival
- Active archive quickly becoming a standard tier
- Majority stored on the *wrong tier*

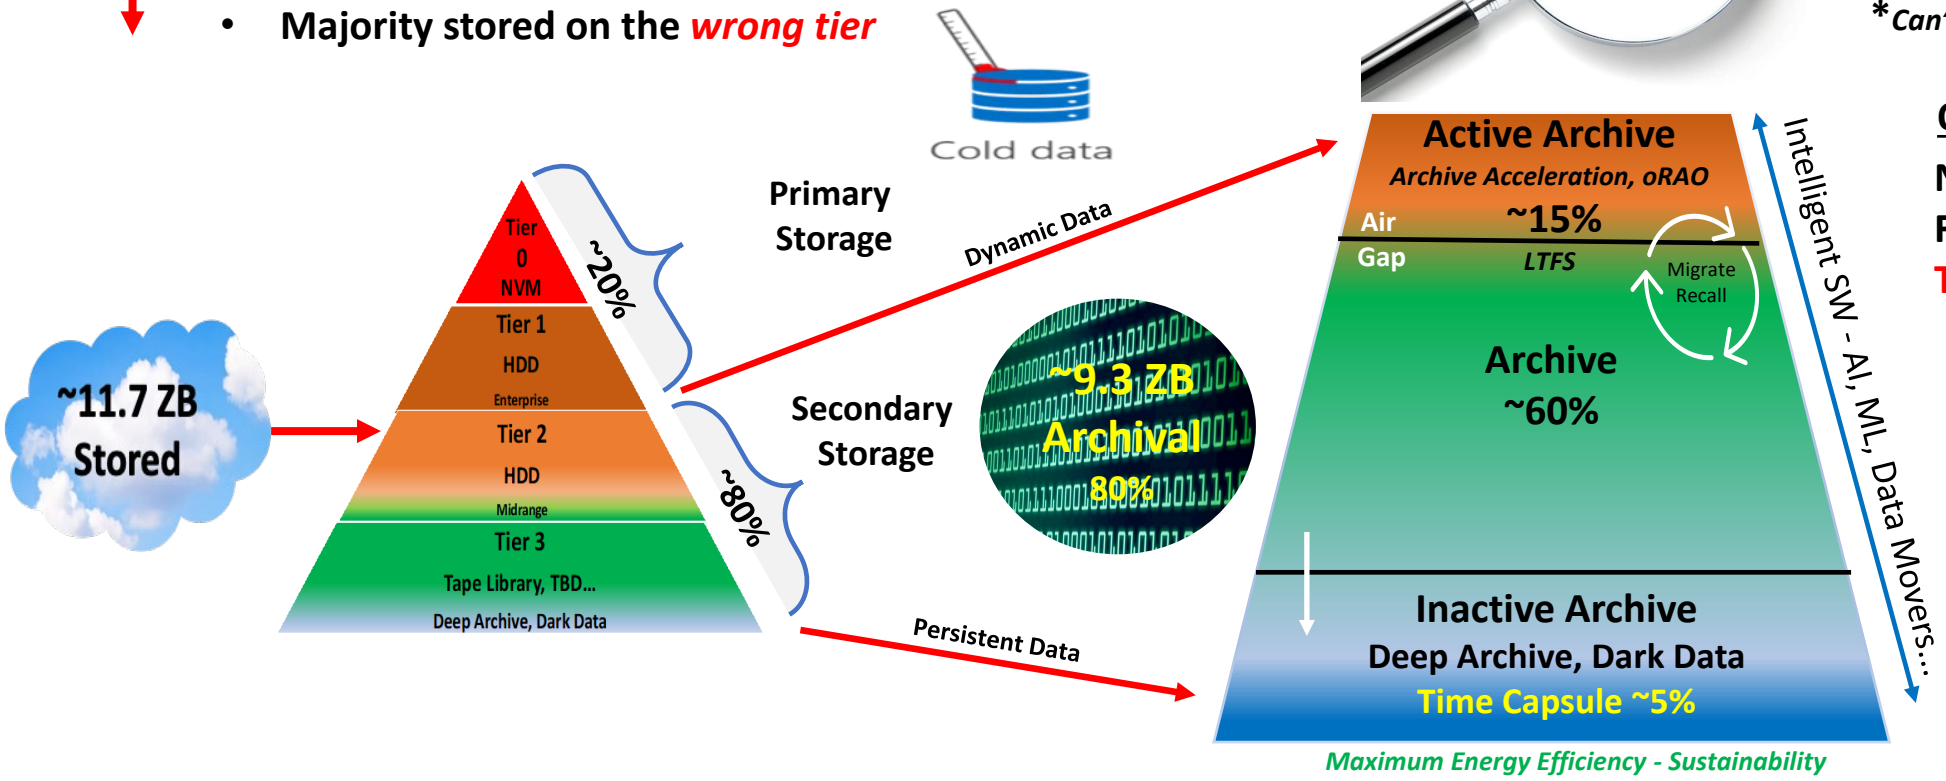

Immutable Data Properties

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, DNA, Glass, 3-D Holographic, TBD….?

Source: Horison, Inc.

Sustainability and TCO Comparison for Secondary Storage

Heavily Favors Tape Over Disk for Archives

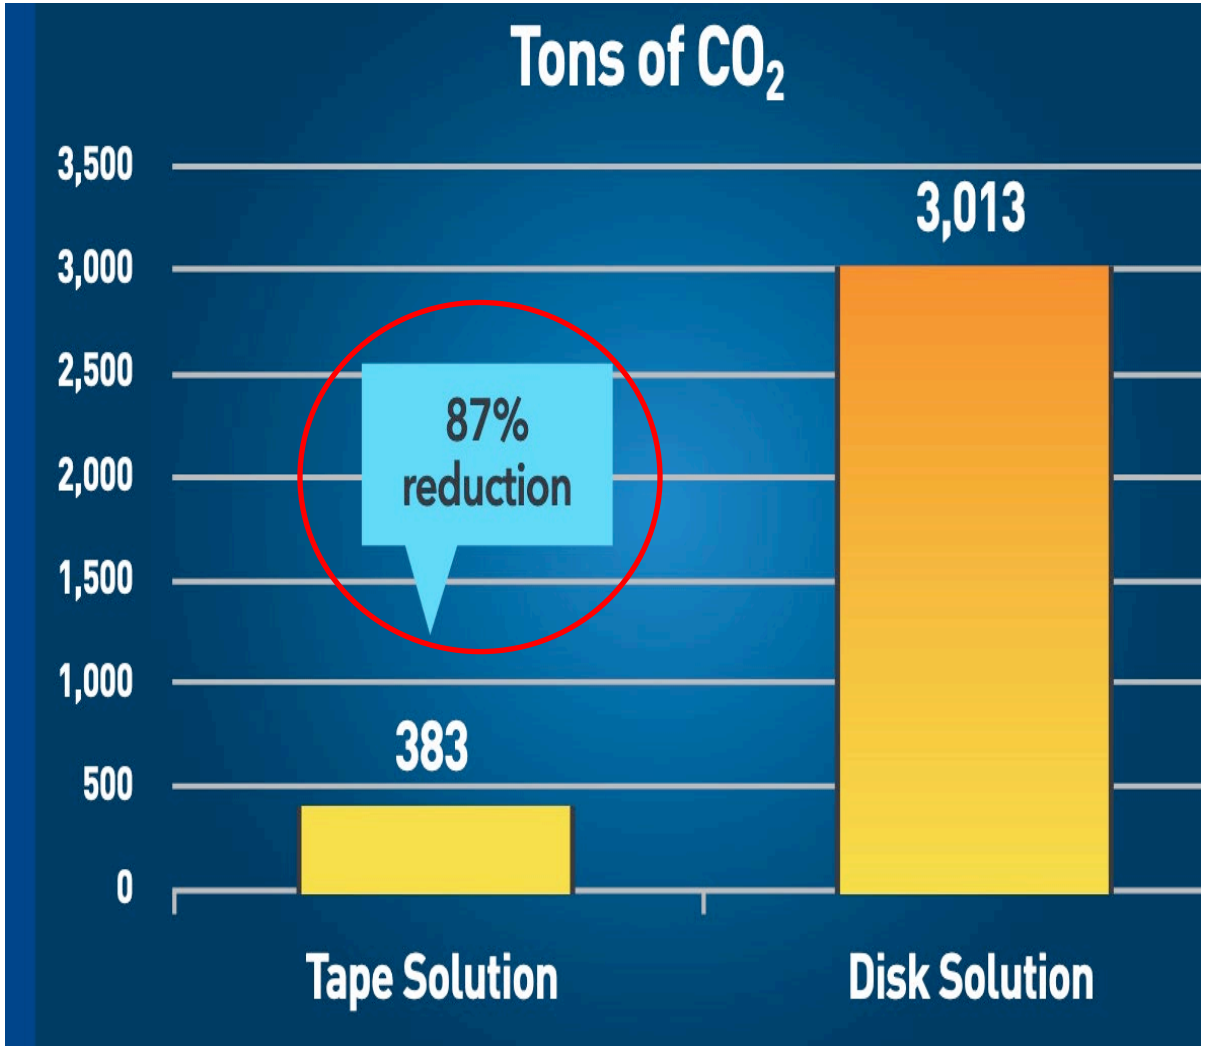

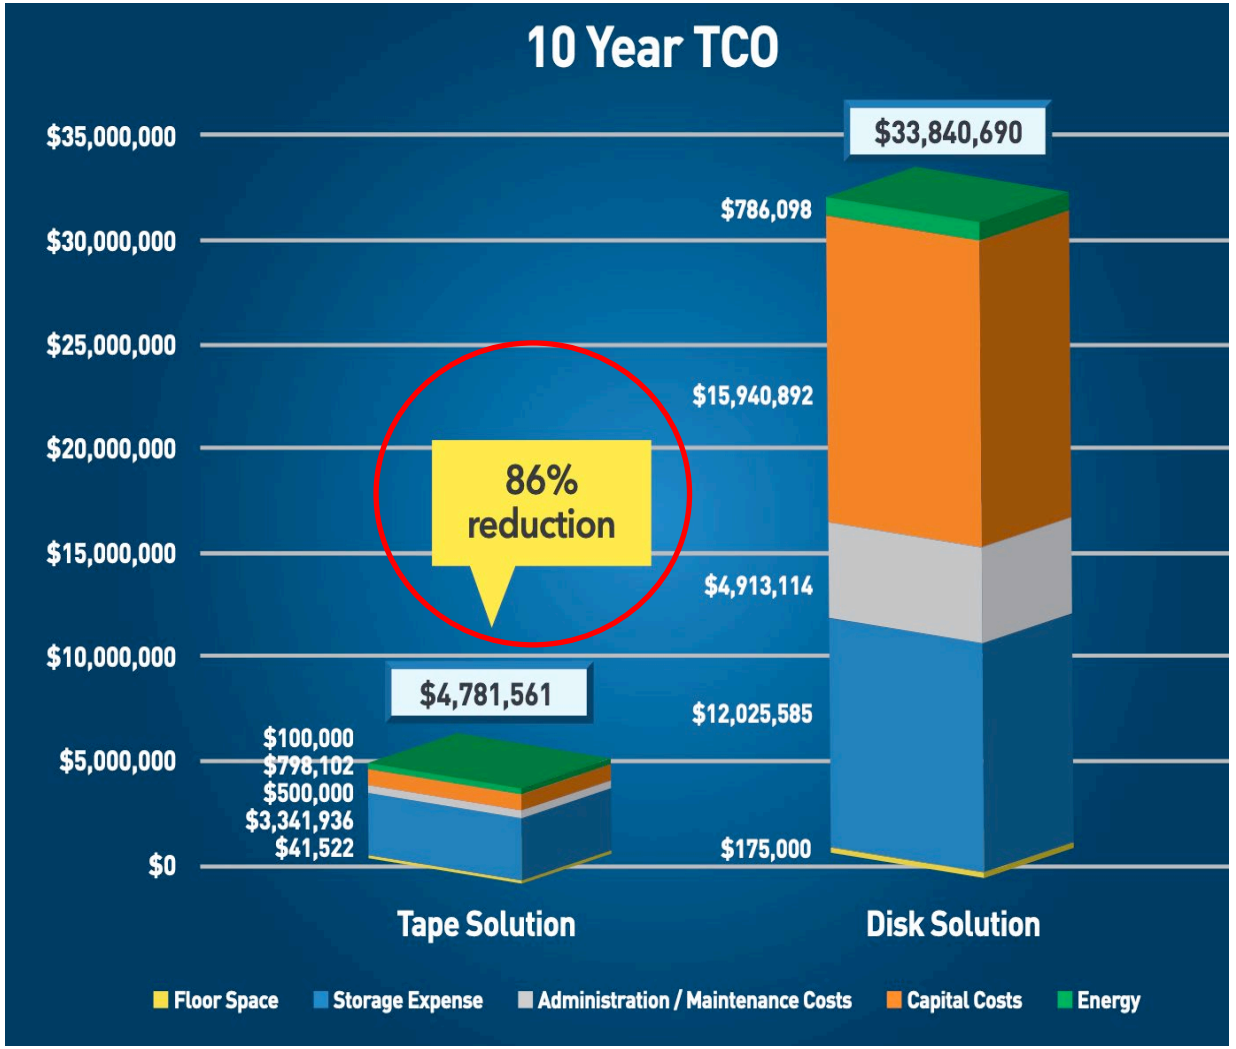

- Ten-year TCO and CO_2 Reduction for 10 PB Growing at 35% Annually

Source: Brad Johns Consulting

From Here to Where…?

Current Technology 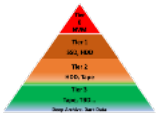	Recording Technique	Roadmap Capability	Future Developments	Challenges
NVM (SRAM, DRAM, MRAM, NAND Flash, PCM, 3D-Xpoint)	Electronic Charge 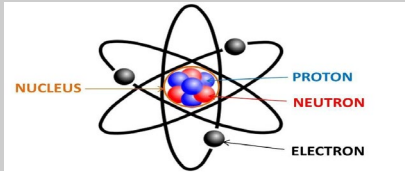	Aggressive development, multiple technologies emerging, CXL, NVMe(oF)	Multi-layer 3D stacking (500^{+}), faster garbage collection, new tiers?	Price
HDD	Magnetic Field 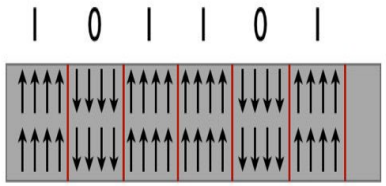	Performance limited and capacity growth slowing	HAMR, MAMR, (? Tb/in^2), multi-platters (9-11), zones, 2-4 actuators, bit patterned, ordered granular, cold HDD	Access density (IOPs), TCO, high energy consumption, \$/TB/watt, CO_2
Tape	Magnetic Field 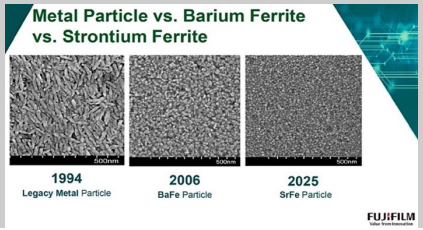	Well defined and sustainable capacity growth, 580 TB demo, high patent activity	Strontium Ferrite (SrFe), Epsilon Ferrite (ε-Fe2O3), TMR, deep archive, RAIL, erasure coding, Geo-spreading, fixed tape	Access time, No consumer market The race to \$0/TB…
Optical Disc	Reflective Spot 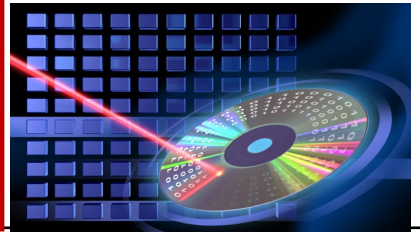	Slow progress compared to magnetics, Not presently a data center technology	Photonic (fluorescent) multi-layer recording has most potential for optics, EMP proof media	Price, performance, capacity, reliability, throughput, slow learning curve

Source: Horison, Inc.

Future Developments for Secondary Storage

A New Storage Frontier Unfolds – But What and When…?

Future Success or VMF…?

580⁺ TB Tape | **Photonics** | **DNA** | **Glass** | **3-D Holographic** | ***TBD…***

WORM ONLY

100 + years

Future

Existing

Hybrid

New Tape Format?

Fluorescent – 3D

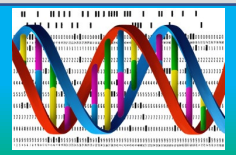

Molecular

Nanostructure

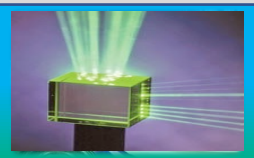

Refraction – 3D

Offline Tape
(mins. - hours)

Access Time Performance

Tape Library
(secs. - mins)

HDD
(ms)

SSD/NVM
(micro secs. μ)

Fastest

Data Preservation Time

Terrestrial

Liquid Memory

Quantum

~5 years

Highest — **Price per Unit of Capacity** — **Lowest**

Source: Horison Inc.

The Optimal Secondary Storage Strategy Currently Favors Tape

Tape Function	Benefits Summary *Tape Re-enters Growth Phase*
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 Improved Access Times - Active Archive, Fastest Data Rates, Smarter and Faster Robotics, RAIT, RAIL, New Time to 1st Byte Features (oRAO, TAOS), Re-writable.
Capacity	LTO-9 Cartridge Capacity @18 TB (45 TB compressed) with 400 MB/sec Data Rate. Smart Zone 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	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 >30 Years for all Modern Tape.
Recording Limits	HDDs Facing Areal Density and Performance (IOPs) Limits. Tape Has a Well-Defined Roadmap (pace..).
Open Standards	LTO and LTFS Provide Open Standard File Interface, APIs. SW (S3 API) Support for Tape Object Storage.
Tape and Cloud Ecosystem	Tape Interfaces Seamlessly With Clouds Using Industry Standard API's. Native Cloud Applications Can Write To and Read From Tape. Hot and Cold clouds.

Source: Horison Inc. 3/2023

The Zettabyte Era and Beyond

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

Saganbyte
$1x10^{33}$

Galactic Data

Geopbyte
$1x10^{30}$

Brontobyte…
$1x10^{27}$

Yottabyte…
$1x10^{24}$

Zettabyte…
$1x10^{21}$

ERA's

Exit…

>186,282 m/s
FTL

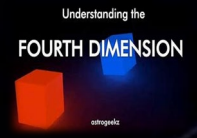

Spacebelt

Zettabyte Era and Beyond

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

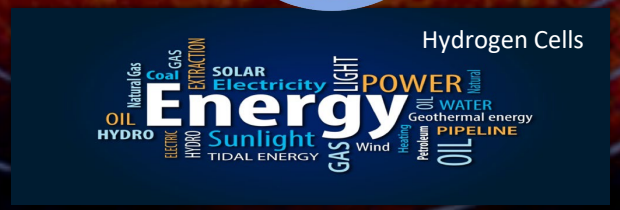

REMEMBER

Things are Changing so Fast
Even the Future is Obsolete

Fractal Geometry

You Are Here

Artificial Intelligence

SEISMIC SHIFTS

Source: Horison Inc.,
Fred Moore President