



AI Redesigning Storage Architectures

Fred Moore, President
Horison Information Strategies
Horison.com



LIBRARY LIBRARY
OF CONGRESS

Washington DC
March 9-10

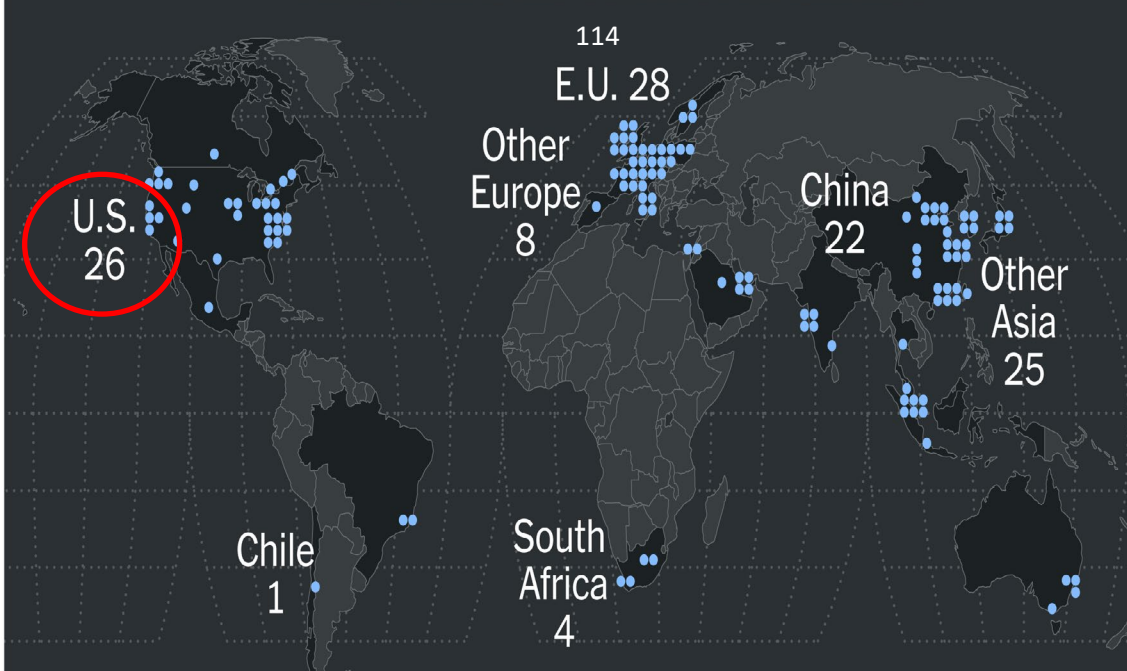
Designing Storage Architectures
for
Digital Collections



AI Fueling Data Center Re-design

AI Is No Longer Just an App – It's a Primary Electrical Load

Where A.I. Data Centers Are Located



- WW digital divide – nations with compute power and those that don't.
- US data centers to use 4.4% of electrical output and 12% of total consumption by 2028.
- AI uses more energy than Ireland, Denmark, Finland, New Zealand and Chile.

- [Stargate data center](#) – a 1,200 acre \$500B joint venture between OpenAI, Oracle and Softbank in Abilene Texas.
- Could consume > 900 MW/h., >450,000 GPUs.
- Fueled by natural gas, wind, solar.
- Closed loop water system uses 1 million gallons of water (low for its size).
- 5 more Stargate data centers planned.
- AI and EV charging are the two fastest growing electric loads.

AI Triggers a VMF (Vertical Market Failure)

Storage Infrastructure Falls Behind Market Demands

- A VMF occurs when the underlying infrastructure fails to meet the demands of the market it serves.
- Current HDD and tape development roadmap progress has fallen behind projected storage specifications and increasing demand.
- VMF has created mounting storage supply shortages and higher prices.
- The Zettabyte scale + AI fueled storage market pressures existing storage architectures and changes access patterns.
- Tape Drive (1) - IBM is the only developer/supplier controlling the entire tape ecosystem specifications.
- Tape Media (2) - Fujifilm and Sony are the only LTO suppliers.
- HDD (3) - Seagate, Toshiba and WD are the only remaining HDD suppliers.
- Tape Library (4) - HPE, IBM, Quantum and Spectra are the primary large-scale library suppliers.
- R/W heads (Tape - WD), (HDD – TDK, WD) are primary suppliers.



- VMF will demand *Redesigning Storage Architectures and data centers – a VMR?*
- *Can Current Storage Technologies Overcome Their Challenges?*
- *Will a New Novel Technology or Solution Ever Arrive?*

(A new tape format, Ceramics, Photonics, DNA, Glass, or the Space Belt)?



AI demands outstrip storage capabilities and supply (2026).

Storage Hits the VMF Wall

VMF (*Failure*)


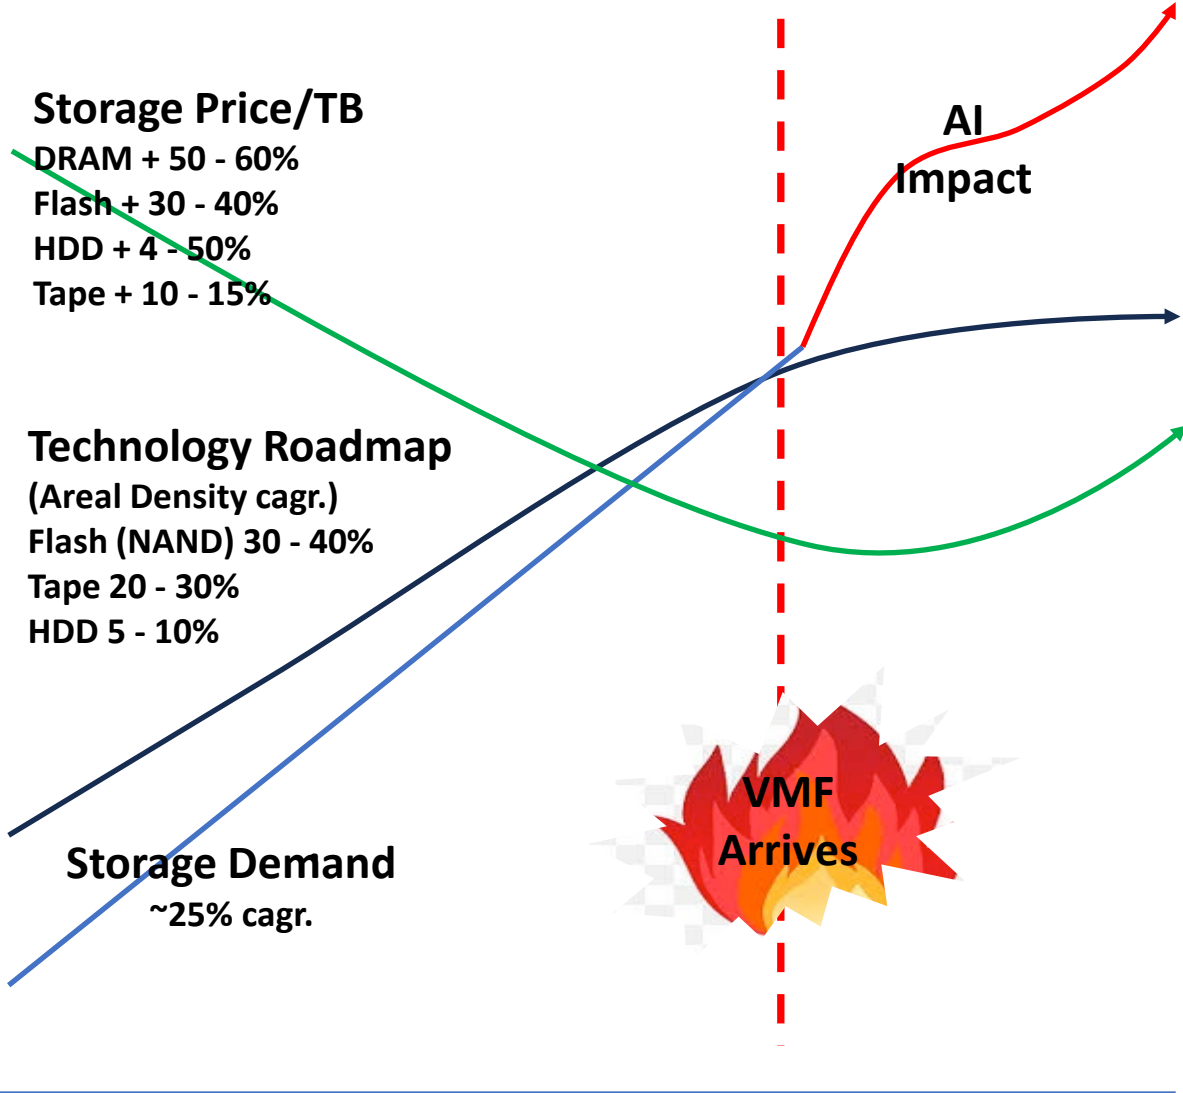
VMR (*Recovery*)?

Compute is cheaper
Storage is cheaper
Engineering is cheaper
Cost optimization: FinOps

4 Data Fabric
 Metadata Maturity, Automation and Operationalization

Cloud Lakehouses

Unified Data Access and Management



AI stresses entire storage technology stack & supply chain

Compute & storage prices higher as demand exceeds supply.

Will storage R&D and roadmaps accelerate?

Higher storage prices = higher revenues. Capital returns over technology investment?

Exponential stock prices for memory, SSD companies (>200%).

Wafer fab cost is excessive (~\$20B)
 Approaching technology limits

Energy costs soaring

Can the VMF Become a VMR?
 Vertical Market Recovery

2020

2025

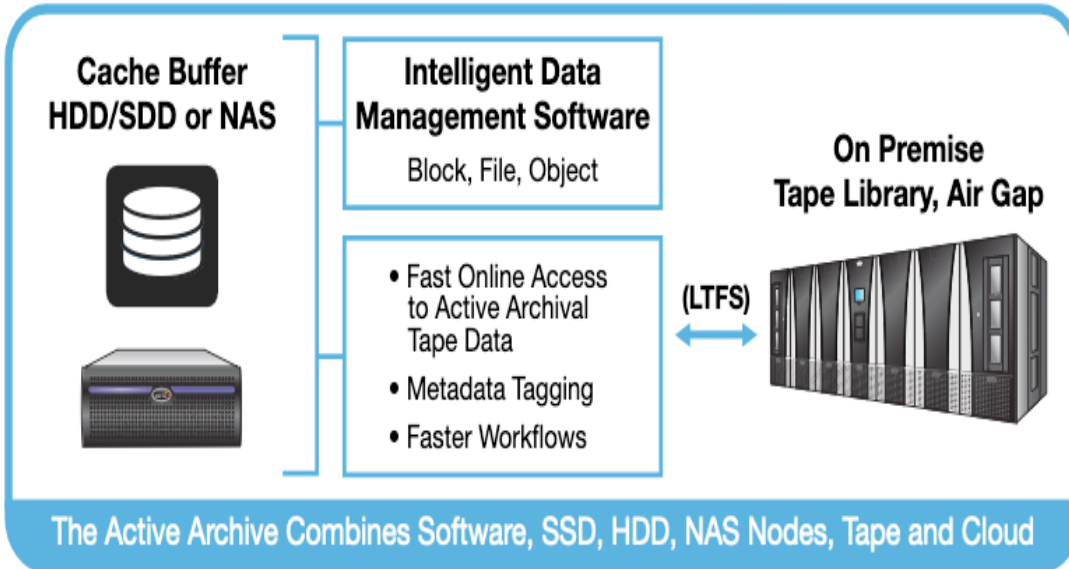
2026

2030+

Active and Dynamic Archives

Approximately 80% of all Data is Low Activity, Archival or Cold

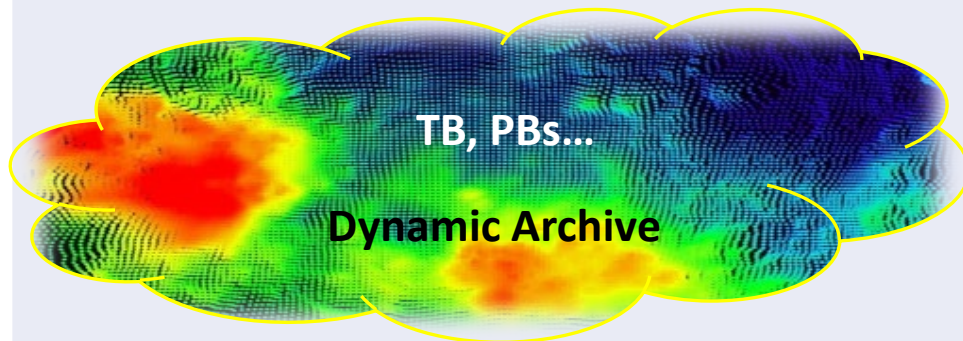
The Active Archive



Secondary Storage Function

A storage hardware system that stores long-term, persistent data (often written once) in online storage allowing brief **bursts of higher - performance access** without permanently migrating that data into expensive primary storage.

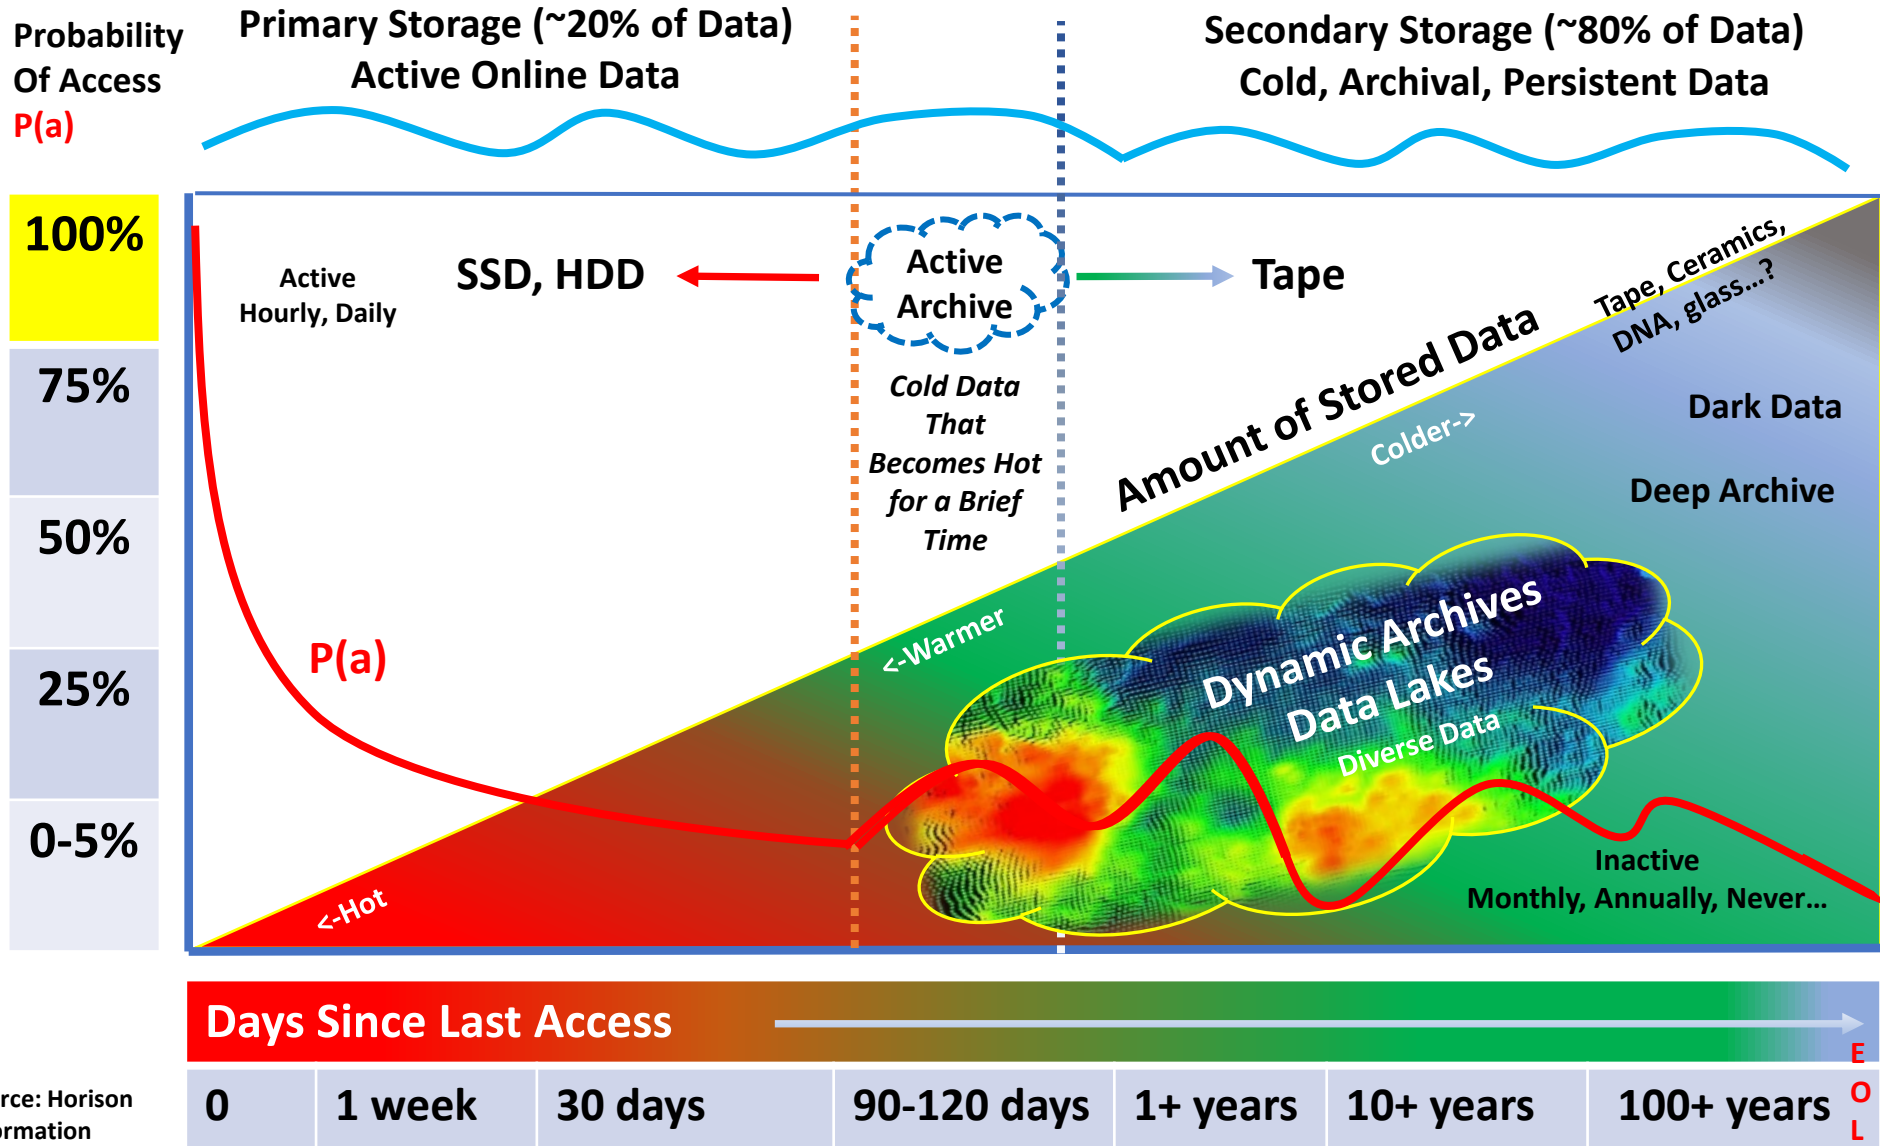
The Dynamic Archive



A data access pattern where data is **not permanently cold**. Instead, diverse archived data can dynamically **move back and forth between cold, warm, and hot states** as its value and use change—often driven by AI, analytics, compliance, or unexpected variable access needs.

New Data Lifecycle and Access Patterns (After AI)

From Hot – to Cold - to Hot - to Warm – to Cold... (The Dynamic Archive Arrives)



The New Value of Data

Lifecycle Profile Factors

Most data typically becomes **cold/archival** in ~90 - 120 days.

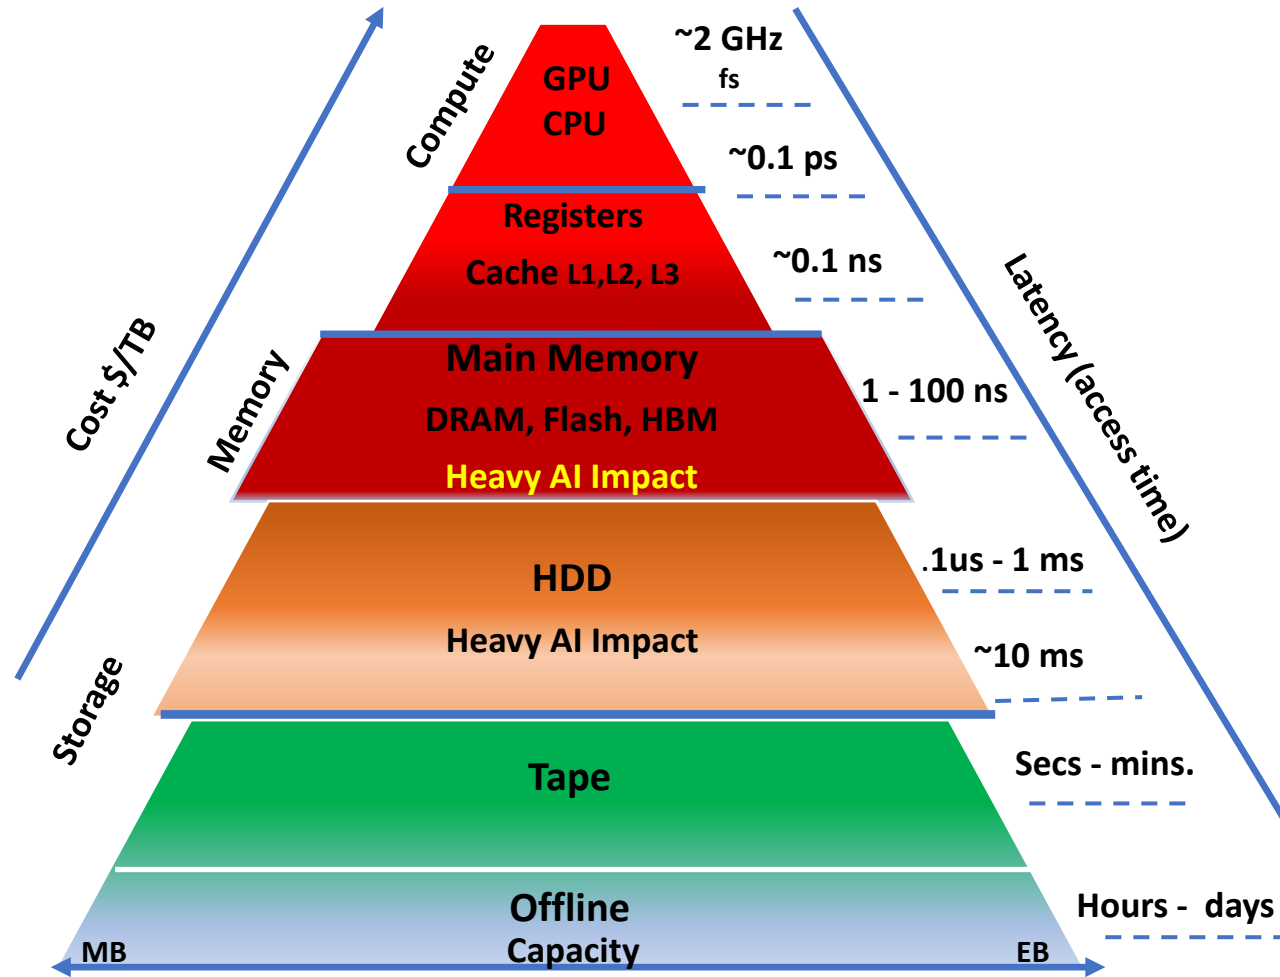
(Dynamic Archive Arrives)
The probability of access $P(a)$ becomes variable with AI.

Cold/Archival data can become **hot** and stay **hot** for variable time periods as the value of data changes.

Archival retention can be >100 years to ∞

AI Unites Compute, Memory and Storage Tiers

Intensive AI, Compute, Memory and Storage Demands *Transforming* the Pyramid



- AI training is compute, memory and storage intensive.
- AI occupies memory for significant longer time
 - Google search ~.5 MB of DRAM
 - AI chat-prompt 1 TB of DRAM and secs. or mins.
 - ms. for a search engine request
 - Memory and storage demands outstrip supply – prices rise.

- Flash and HDDs store AI data sets for fast access online to training models and are stressed by AI demand.
- Tape provides source data for AI to harvest, primary backup and archive target. The *Dynamic Archive Arrives*.


- By 2030 the AI sector could consume up to 1,000 TWh/year, 15-25% cagr.
- Inference consumes 80%+ of all AI energy.
- Global AI demand 4 – 7 billion km³ of water in 2027.
- Is Your Data AI Ready (metadata, storage device, backup)?

• **No Data = No AI!**

➔ **Will AI end the Golden Era of Software Programming?**



AI is a Material Driver of Grid Stress and Water Scarcity



Secondary (Archival) Storage Model – 2026 and Beyond

New Tiers Target the Dynamic Secondary Storage Stack

By YE 2025

~175 ZB
Created



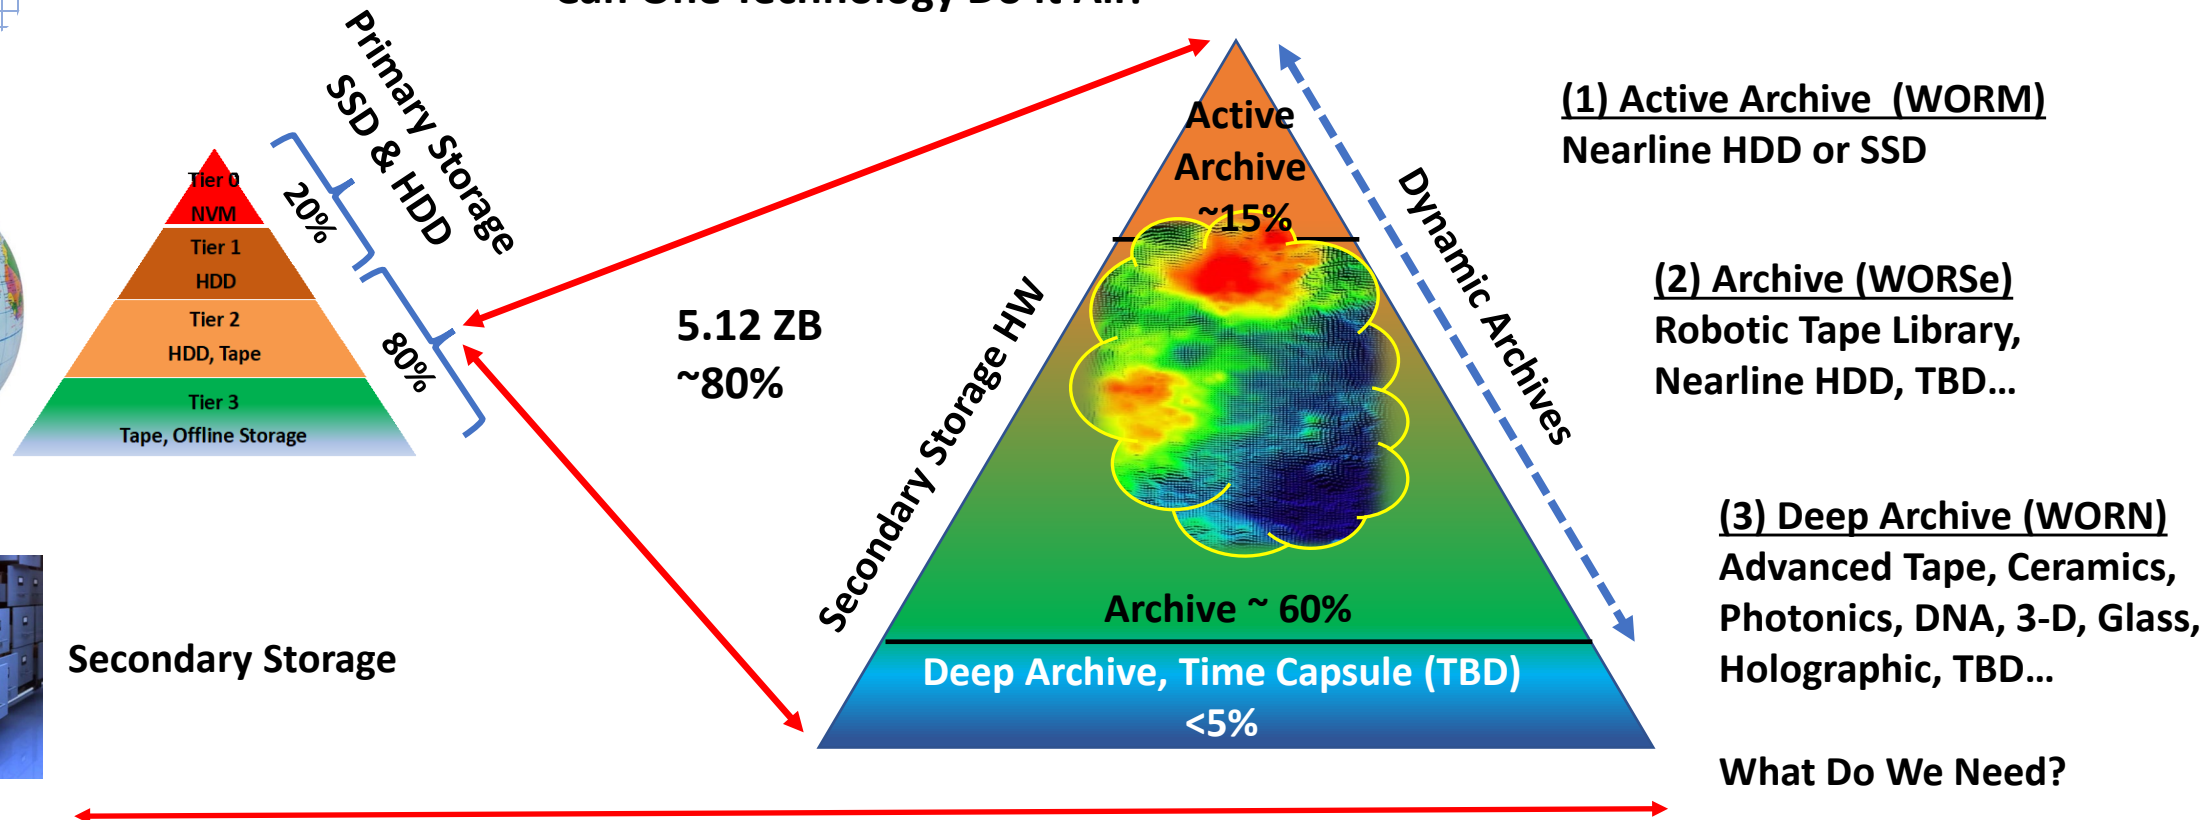
~5.12 ZB
Archival
~80%



- Secondary Storage Has Three (3) Distinct Tiers
- Data Moves Between Tiers
- Can One Technology Do it All?

Archival Data is Mostly Write Once

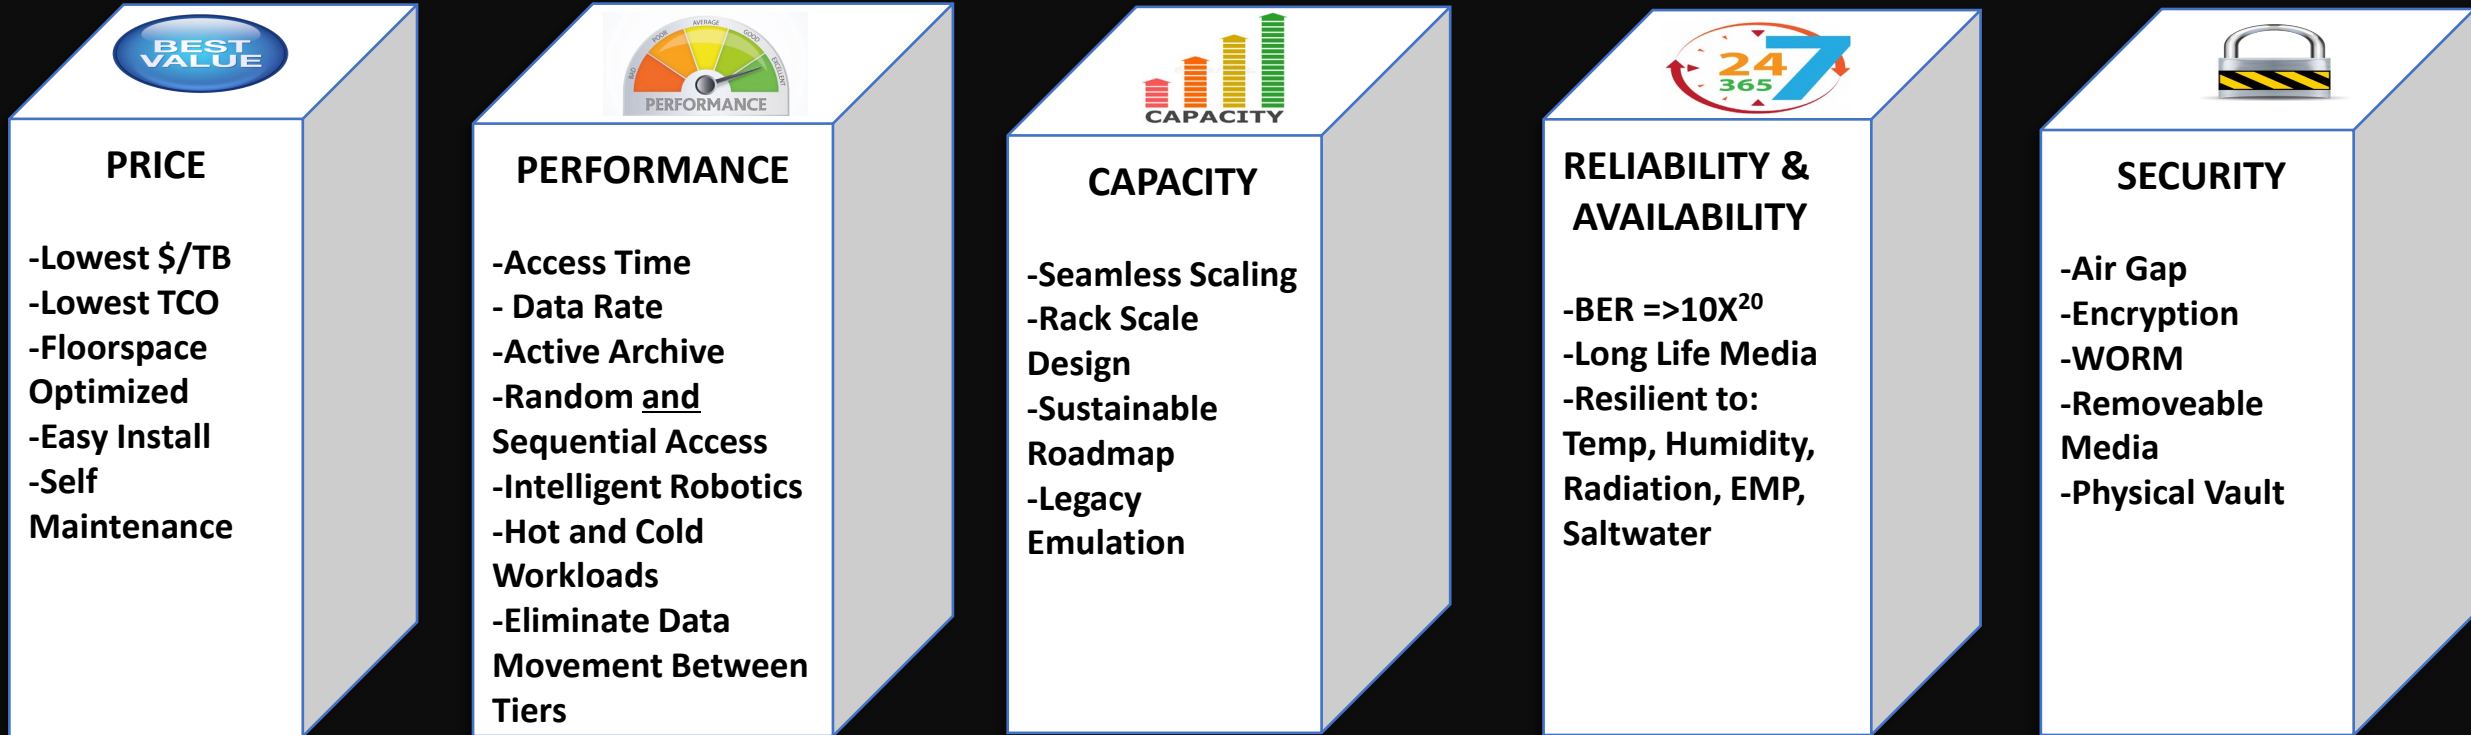
- WORM – Write Once, Read Many
- WORSe -- Write Once, Read Seldom
- WORN – Write Once, Read Never



Re- Designing Storage Architectures – *What Will Be Required?*

Meeting New Secondary Storage Demands Beyond 2025 ?

Critical Architecture Design Requirements



THE RISE of SUSTAINABILITY


Low CO₂e, Power, Water and Electronic Waste Generated, Reliable Energy Sources

By 2025 ~6.4 ZBs of Secondary Storage Capacity Installed

Source: Horison Information Strategies, LLC

Will Any New Emerging Storage Technologies Arrive?

A New Storage Frontier Needs to Unfold – But What and When?



New Tape Format?




DNA
Molecular



Celestial Archives




Lunar and Terrestrial
Storage Systems,
Constellations of Satellites




Quantum



TBD ?




Photonics



Glass nano-layers
Ceramics



Liquid Memory



3D-refraction



Nanostructure