Storage Outlook

Welcome to the Zettabyte Era

Unprecedented Demand
Driving
Unparalleled Solutions

Fujifilm 12th Annual Global IT Executive Summit

June 22-25, 2022

San Diego

HORISON Information Strategies

Fred Moore
President
Horison.com

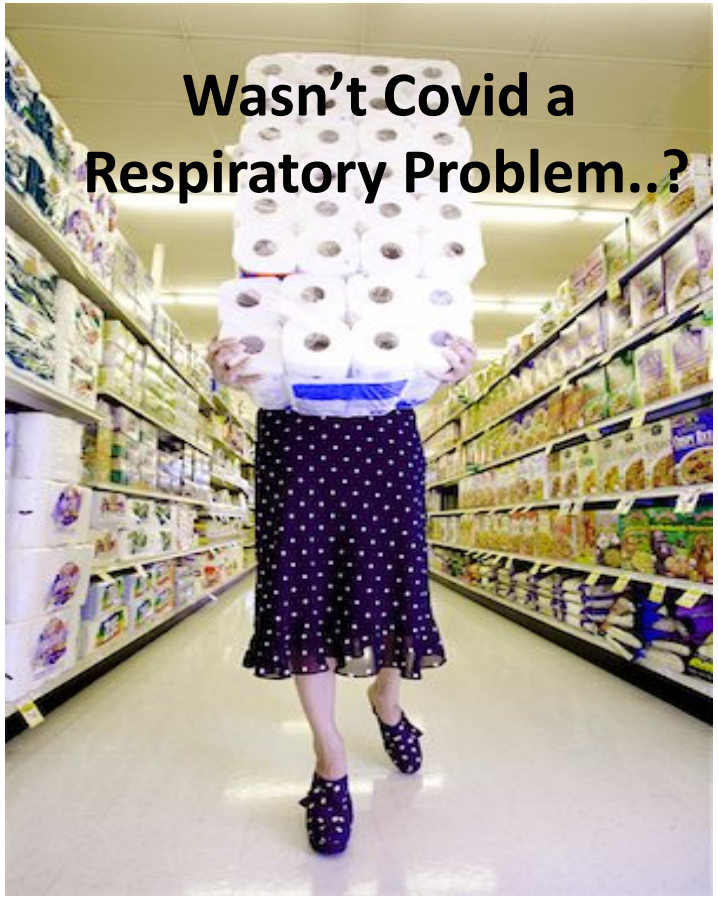

The Official Flag of Covid 19, 20, 21, 22, ...

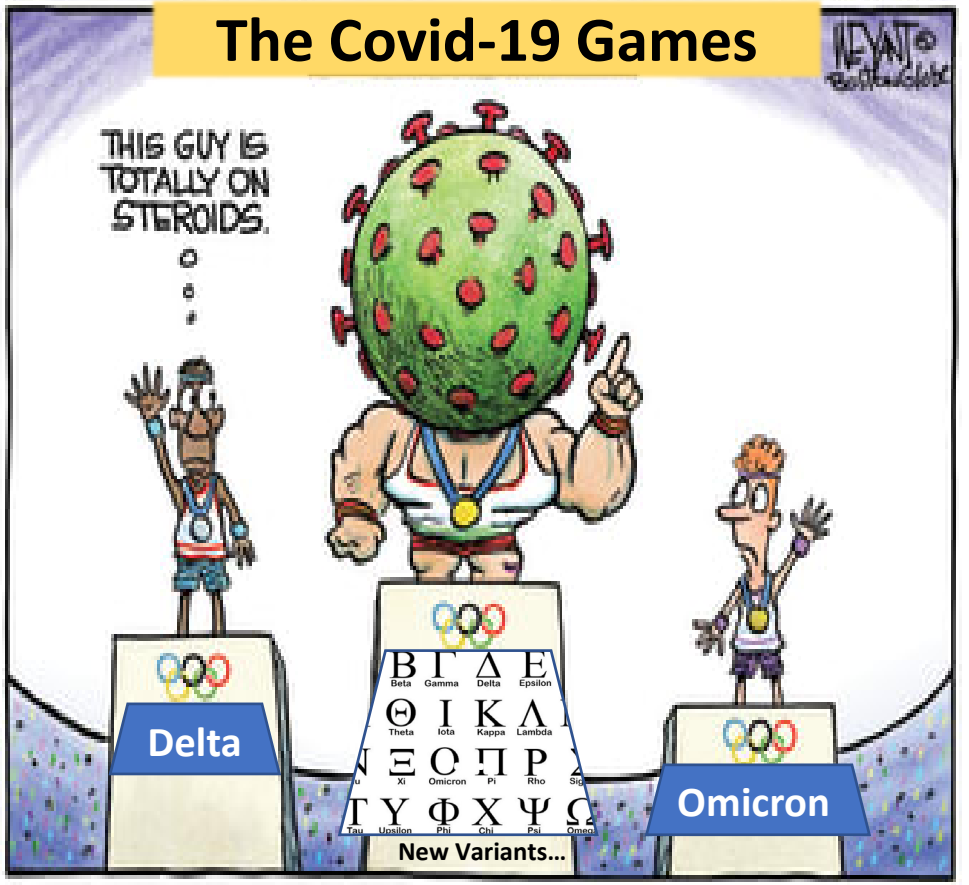

Before We Examine the Zettabyte Era – When Will the Covid Era End?

The Zettabyte Era – How Big Is It?

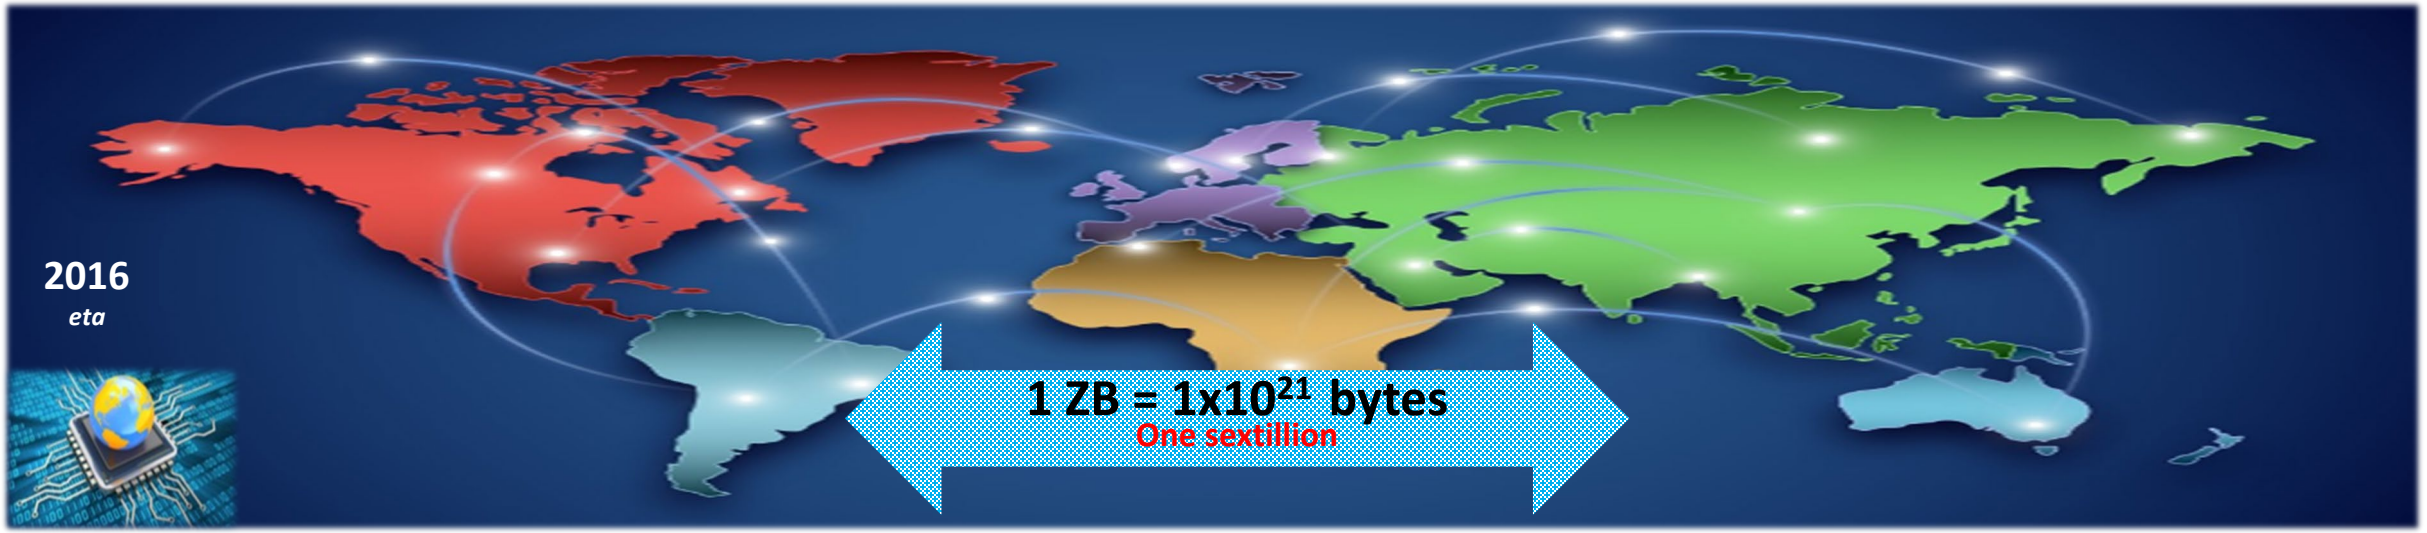

- By 2025, 175 ZBs projected to be created, ~11.7 ZB will be stored.
- 1 ZB equals 66.7 years of the Large Hadron Collider's experimental data.
- 1 ZB would store over 7.5 trillion MP3 songs.
- 1 ZB = 57.5 billion 32 GB Apple iPads or 250 billion DVDs.
- 1 ZB = 125 million years of 1 - hour TV shows.
- 3.3 ZBs of global IP traffic generated in 2021 (2/3 from wireless, 15% from Netflix, 11.4% from YouTube).
- 1 ZB would fill 55.36 million LTO-9 (18 TB) cartridges or 50 million 20 TB HDDs.

Source: Numerous industry sources, Horison, Inc.

The Digital Universe

Archival Pileup Should Exceed 9.0 ZB by 2025

Digital Data Lifecycle

When Does Data Become Archival?

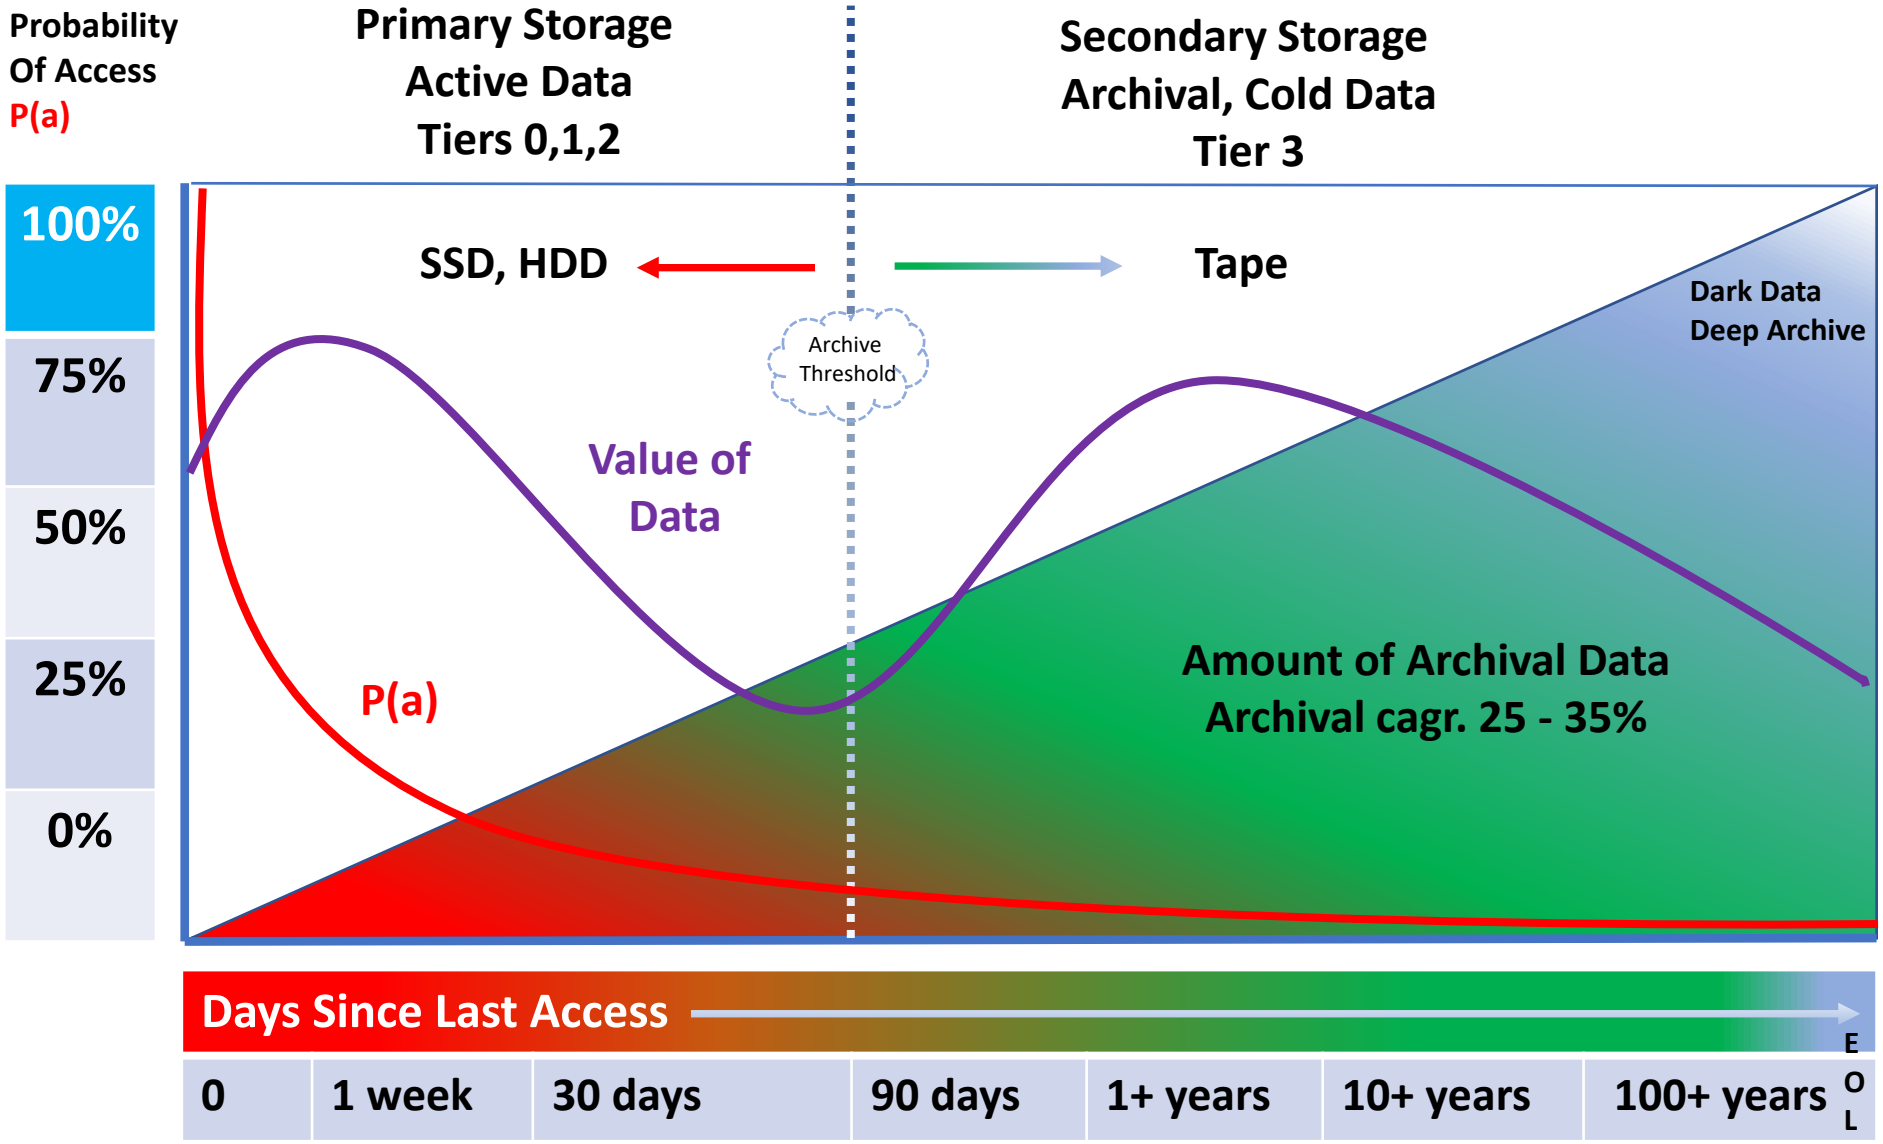

Key Factors

Probability of access P(a) declines as data ages

Data typically becomes archival in ~90-120 days

Archival data piling up faster than it is analyzed

The value of data can vary over time

Active archives support growing AI/ML Ops

Archival retention can be >100 years to ∞

Source: Horison, Inc.

The Tiered Storage Model

By 2025 a New Secondary Storage Paradigm Begins to Emerge

By 2025
Source: IDC

~175 ZB Created

~11.7 ZB Stored

~9.3 ZB (80%) Archival

Intelligent Storage Software (AI, ML)

Optimal Data Classification by Tier
% of All Data

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

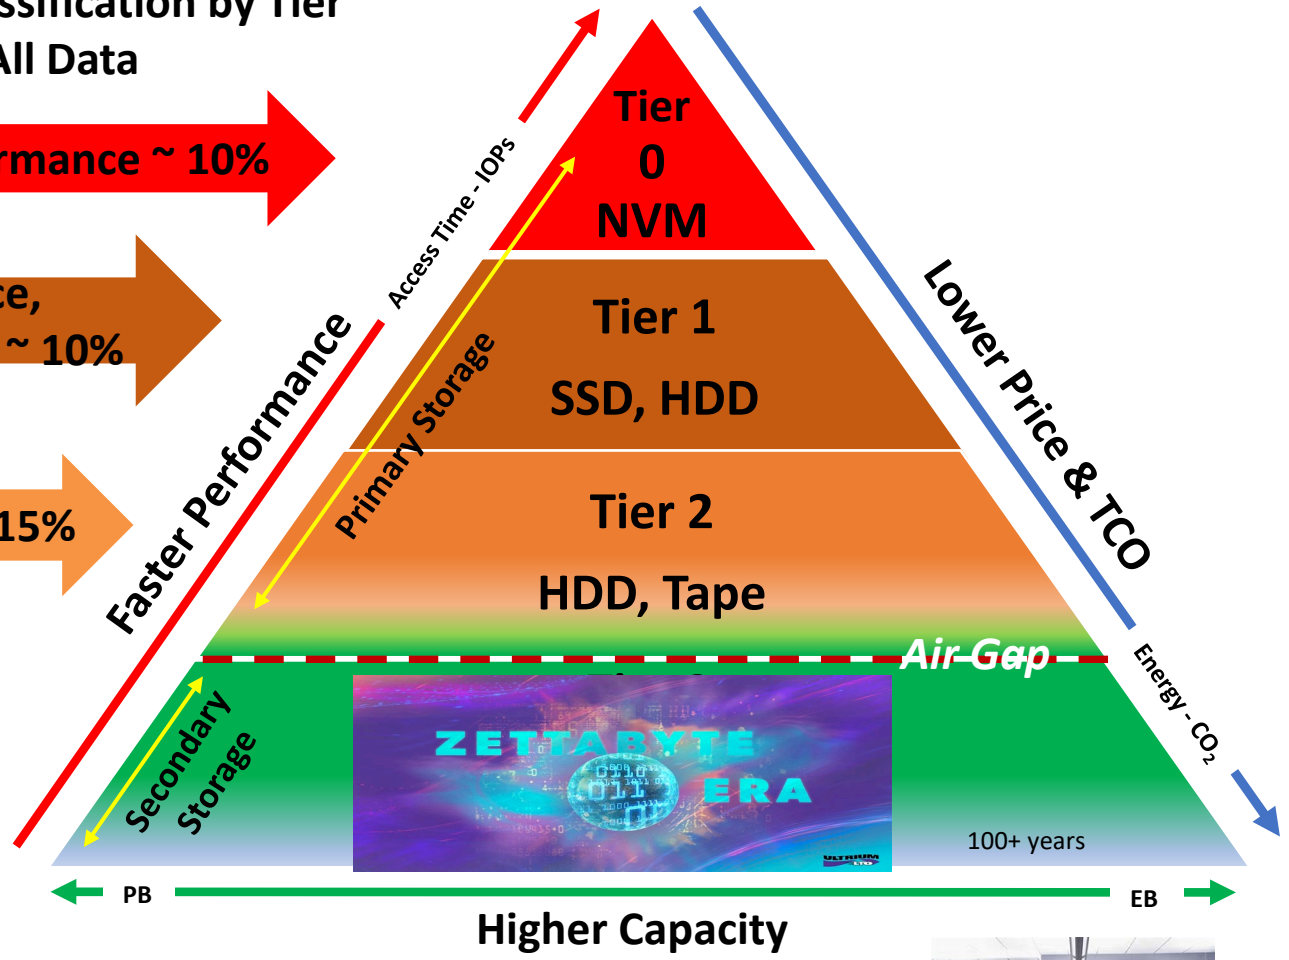

Hybrid Cloud, Air Gap, Vault

The Cloud Uses All Storage Tiers SSD, HDD, Tape

Source: Horison, Inc.

Analyzing the Archives – 2025

Three Secondary Storage Tiers Arrive to Address 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 becomes a standard
- Much stored on the *wrong tier*

Archival Data is Mostly Write Once

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

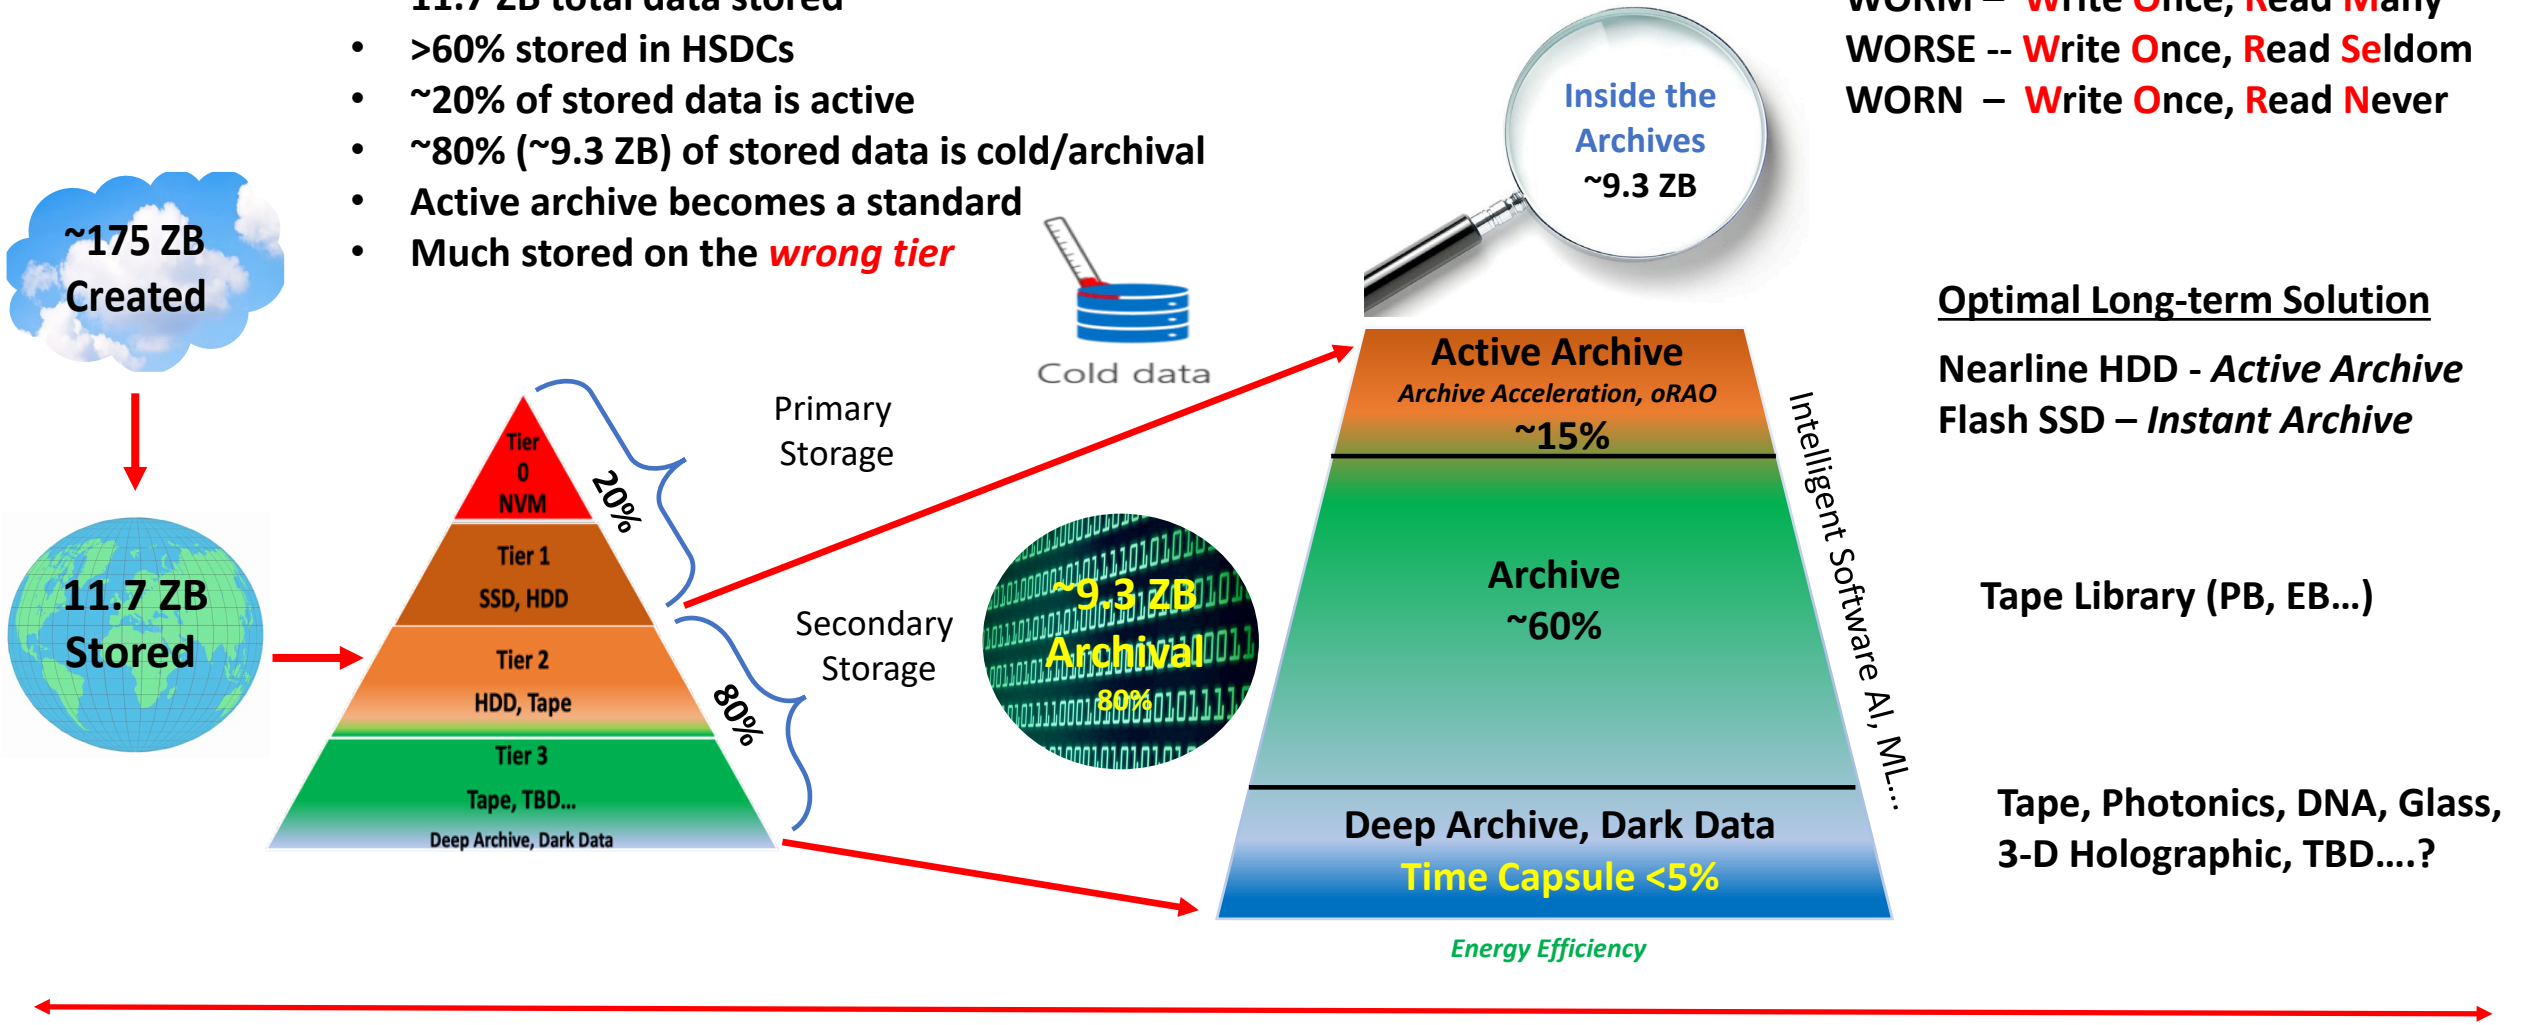

Optimal Long-term Solution

Nearline HDD - *Active Archive*
Flash SSD – *Instant Archive*

Tape Library (PB, EB…)

Tape, Photonics, DNA, Glass, 3-D Holographic, TBD….?

Source: Horison, Inc.

Storage Formats Key for Data Classification

Object Storage Becoming Primary Format for Massive Archives

Format

Structured Data (Blocks)

Well organized

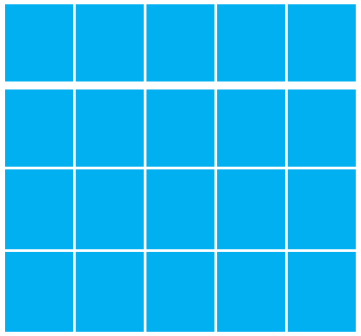

~20% of all data

- Easily searchable metadata
- OLTP, Data bases, DWHS
- OS, Mission critical data
- Hi-performance apps (IOPs)

Primary Technology

HDD and SSD

Semi-structured Data (Objects)

Loosely organized

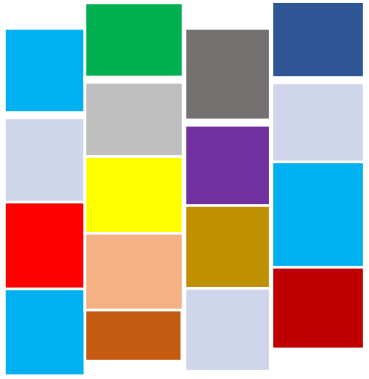

~80% of all data

- Customizable metadata, tags
- No scaling limits
- Object storage supports S3 API
- Fast sequential throughput – ideal for streaming tape
- Surveillance, IoT
- **Optimal choice for archives**

HDD

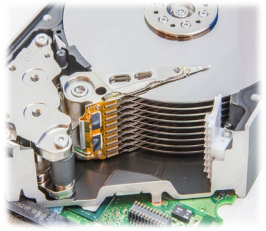

Tape

Unstructured Data (Files)

Not organized

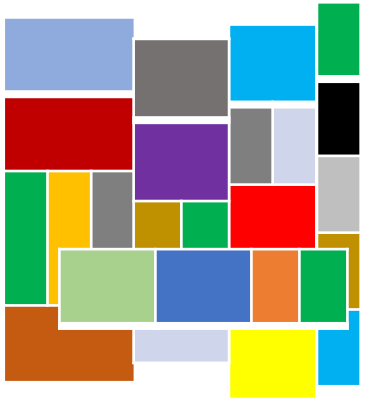

- Limited metadata and scaling
- Big data, archives, BUR, DR
- Images, text, PDF, video
- Mobile data, scientific

Cloud

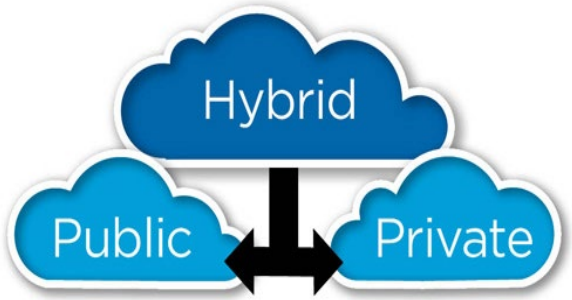

Source: Horison, Inc.

Hyperscale Data Centers Reshape IT Landscape

- A Hyperscale Data Center (HSDC) is a massive infrastructure – often over 100,000 ft^2, largest is > 10.763 million ft^2 (= 132.9 soccer fields).
- Use self healing redundant components, geo-spreading - compute and storage intensive.
- Full automation and seamless scalability (PBs to EBs) are critical.
- Extreme energy consumption, carbon footprint and sustainability challenges mount.
- Tape usage increasing and ***will be critical*** to enable growth and control infrastructure costs.
- HSDCs pushing all technology limits.

Source: Horison, Inc.

Hybrid Clouds Become Integral To Data Center Operations

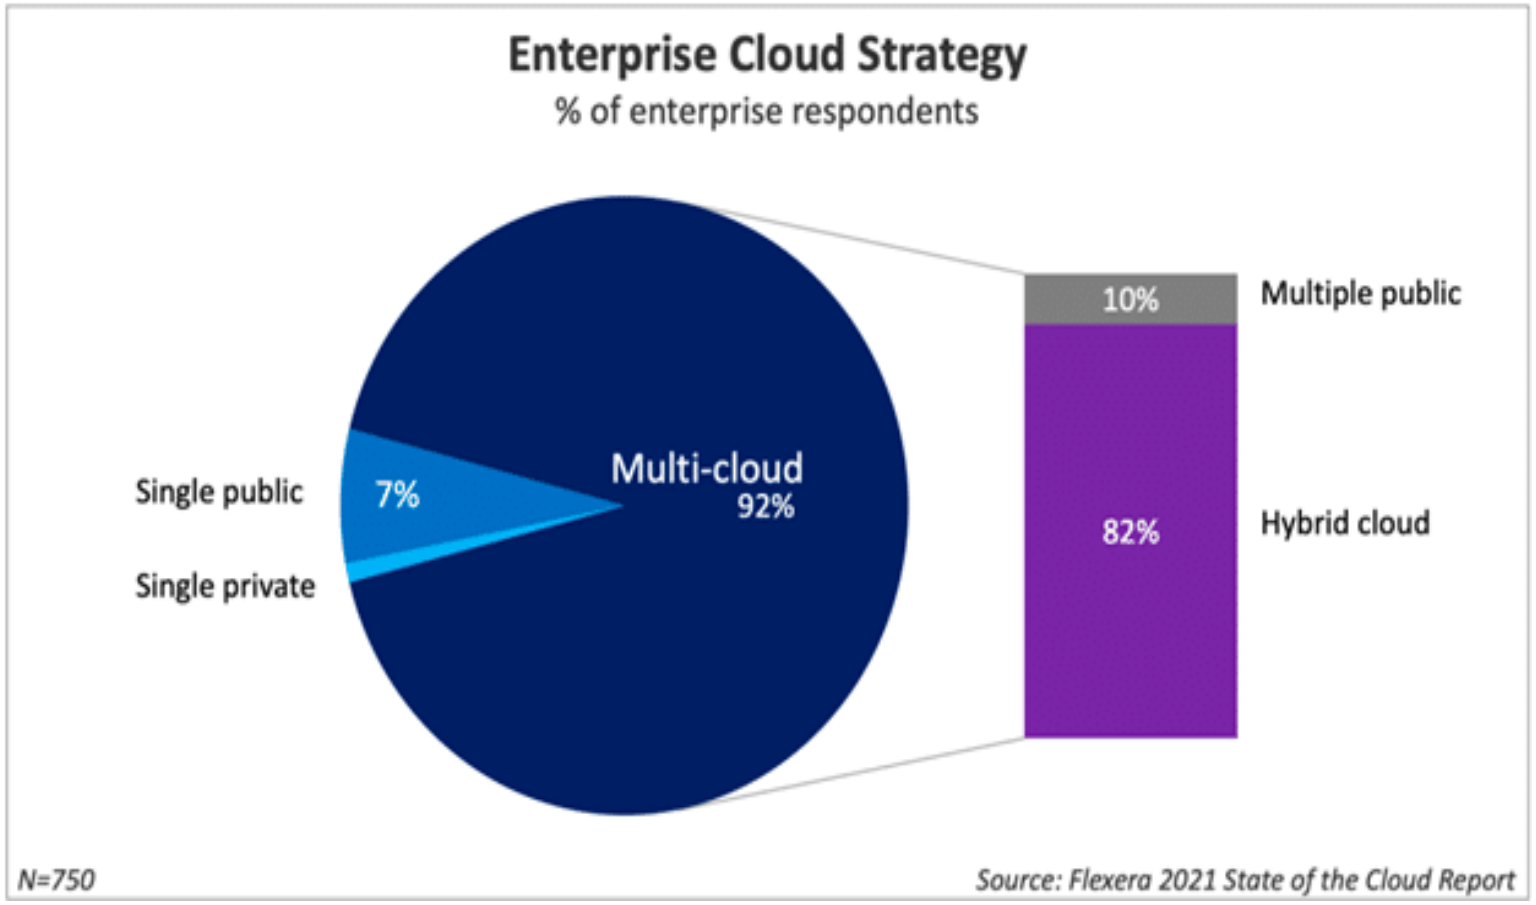

Hybrid Cloud Storage Includes

- Private cloud (on-premise) for OLTP, performance, security, data control, no hidden costs
- Public cloud for faster scalability, lower capital costs

STaaS (Storage as a Service) - Backup, archives, DR, big data are key hybrid apps.

Cloud storage tiering (SSDs, HDDs, tape) optimizes private and public clouds.

Hybrid clouds supports the high-availability 3-2-1-1 strategy for backup and archive data. 1 - one air gap copy

Multi-cloud is the use of multiple cloud computing and storage services in a single heterogeneous architecture.

The Cloud, Fog and Edge

Computing and Storage at the Logical Extremes of a Network

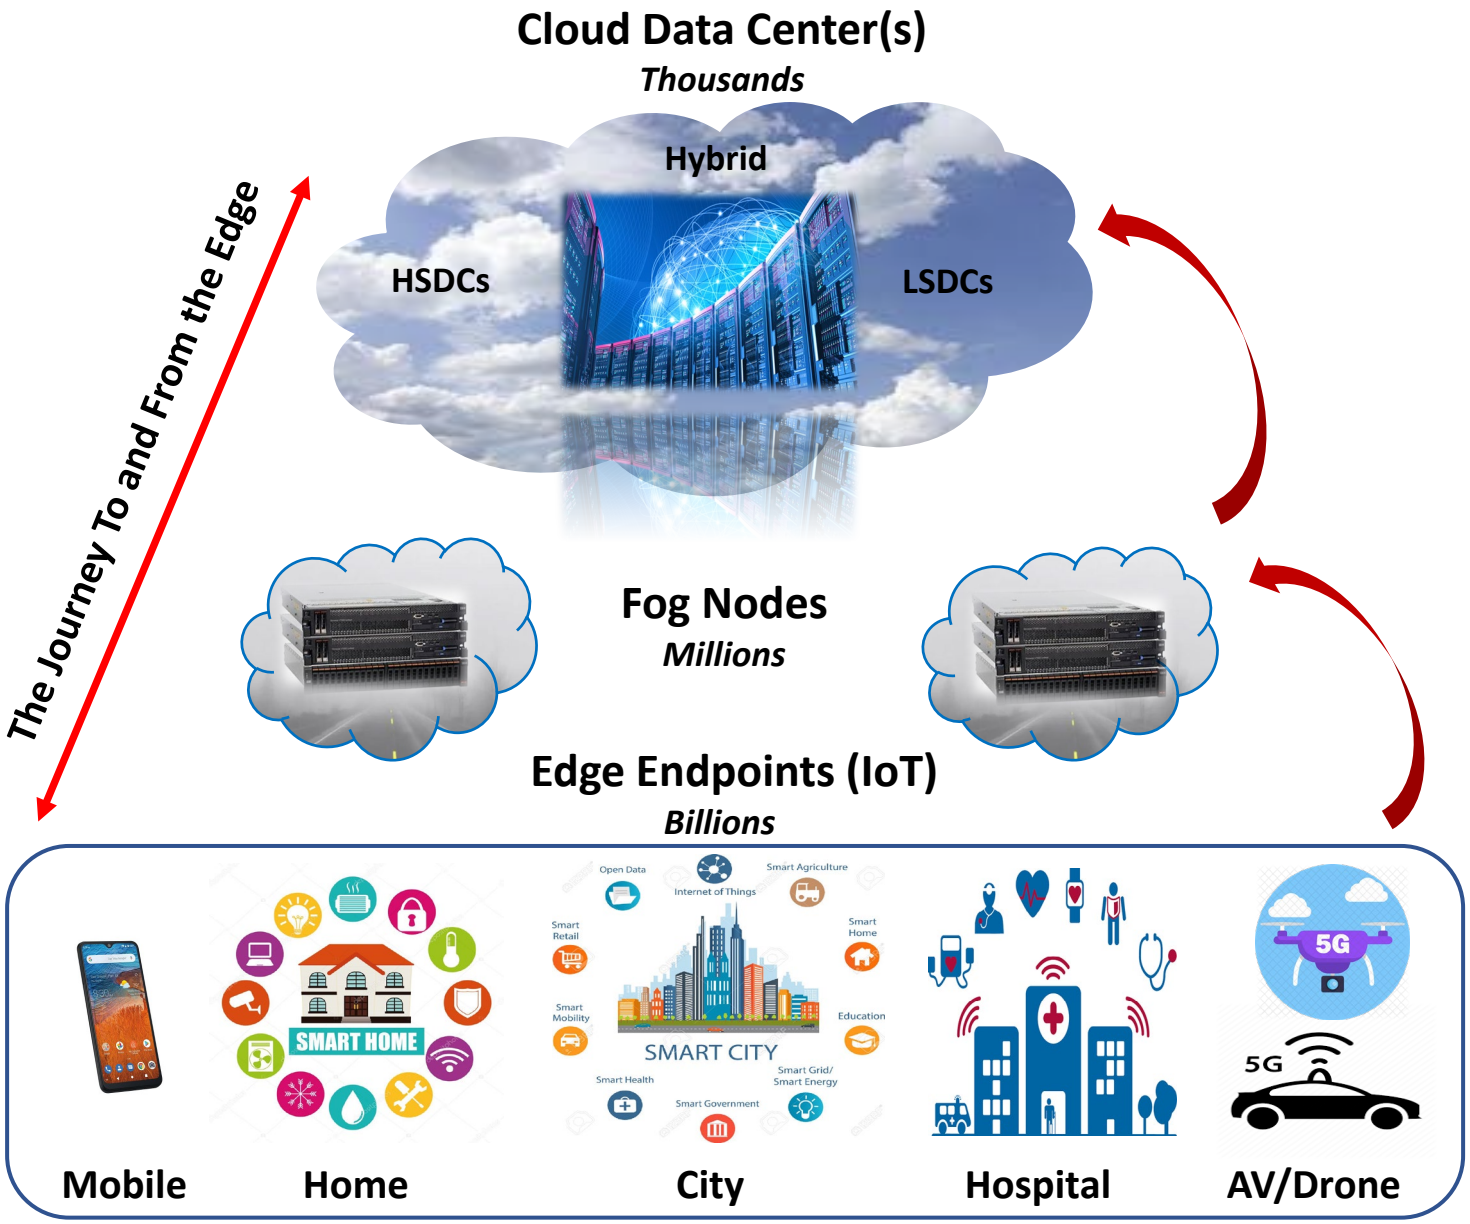

DMZ

The Cloud Ecosystem

The Cloud (core)
- Big Data Processing
- Data Warehouses
- Data Lakes (*oceans..*)
- Key Role for Tape

The Fog
- Data Reduction
- Data Aggregation
- Real Time Processing
- Local Processing
- Quick Analysis
- Avoid Data Floods

The Edge (IoT, 5G)
- Sensors, end points
- Smart Vehicles, Smart Devices, Drones
- Geo-fencing
- Connected Systems
- Security concerns↑
- Key Role for Flash

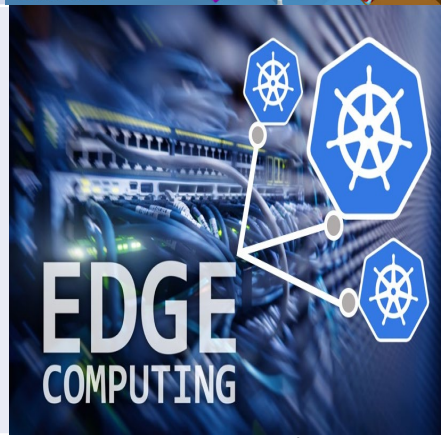

Source: Horison, Inc.

Cybercrime Scenario 2022

Tape Air Gap Plays A Key Role In The Cyber Security Ecosystem

The 3 – 2 – 1 – 1 Data Protection Strategy

For Backup and Archival Data

3	2	1	1
Copies of Data	Different Media Types (SSD, HDD, Tape)	Copy Offsite … Error-free (SSD, HDD, Tape)	Air Gap Copy Offline (Tape)

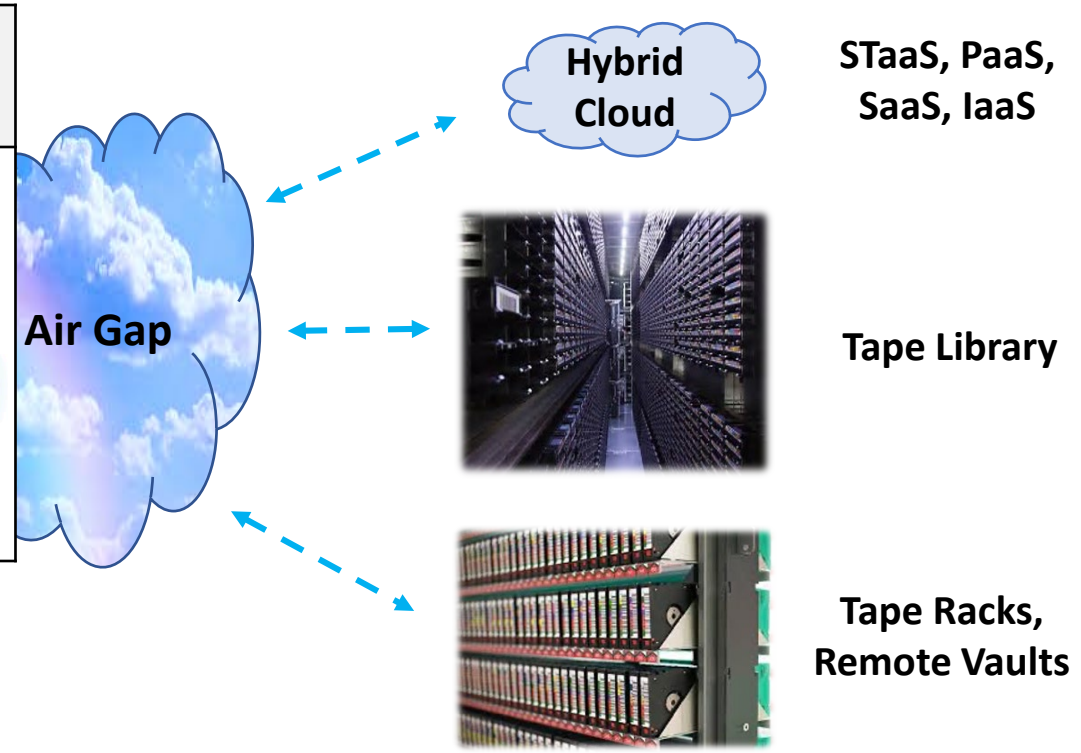

- Zettabyte Era, IoT, Edge, and Shadow Assets Expand the Cybercrime Attack Surface.
- Ransomware attacks damage from rose from $11.5 B in 2019 to $20 B in 2021.
- Highest ransom fee was $70 M, 82% defaulted to paying the ransom.
- Cloud cyber attacks account for 20% of all cyber attacks.
- WW ~3.5 million unfilled cybersecurity jobs in 2021. Cybersecurity Ventures.
- Protection solutions *lagging* for emerging technologies like 5G, the edge, IoT, crypto, AI and ML – ***all TBD***.

Source: Horison Information Strategies

How Secure is the Energy Grid?

Without Electricity There is No IT Industry

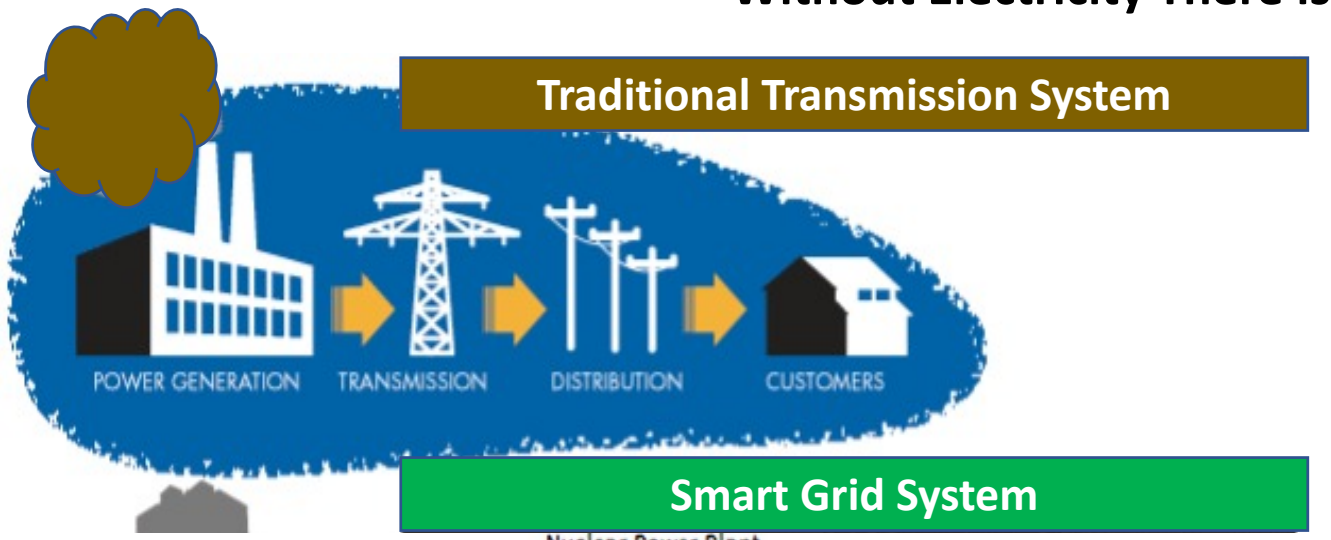

- Assume everything on the energy grid can be hacked.
- Malware can open circuit breakers, stopping electricity flow creating a sustained power outage.
- Power grid strategy transitions to a *smart grid* to be more responsive to changing power needs.
- Smart grids must have self-healing and predictive, capabilities, multiple sources.
- AI and ML are key in building the smart and secure energy grids of the future.
- EVs and Bitcoin increase grid demands.
- Electrical grid critical to national defense.

Source: Horison, Inc.

Data Center Energy Consumption by Component

Data Center Monthly Energy Costs

Ex: PUE = 65/35 = 1.86

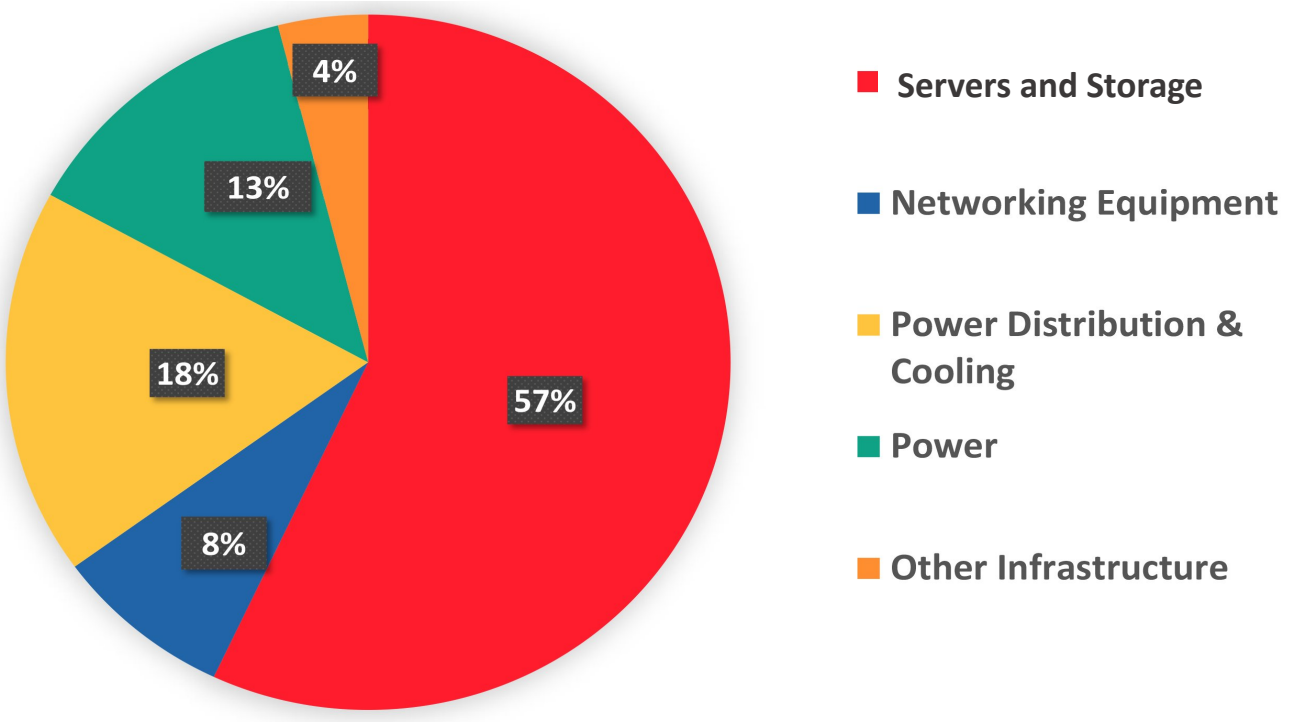

- **Data centers consume nearly 3% of world's power generation**
- **SSD wattage ~6W/module**
- **HDD wattage ~8.6W/drive**
- **Tape wattage negligible**
- **1 Bitcoin transaction – 2,258.49 kWh (10 minutes/coin)**
- **100,000 Visa transactions – 148.63 kWh**

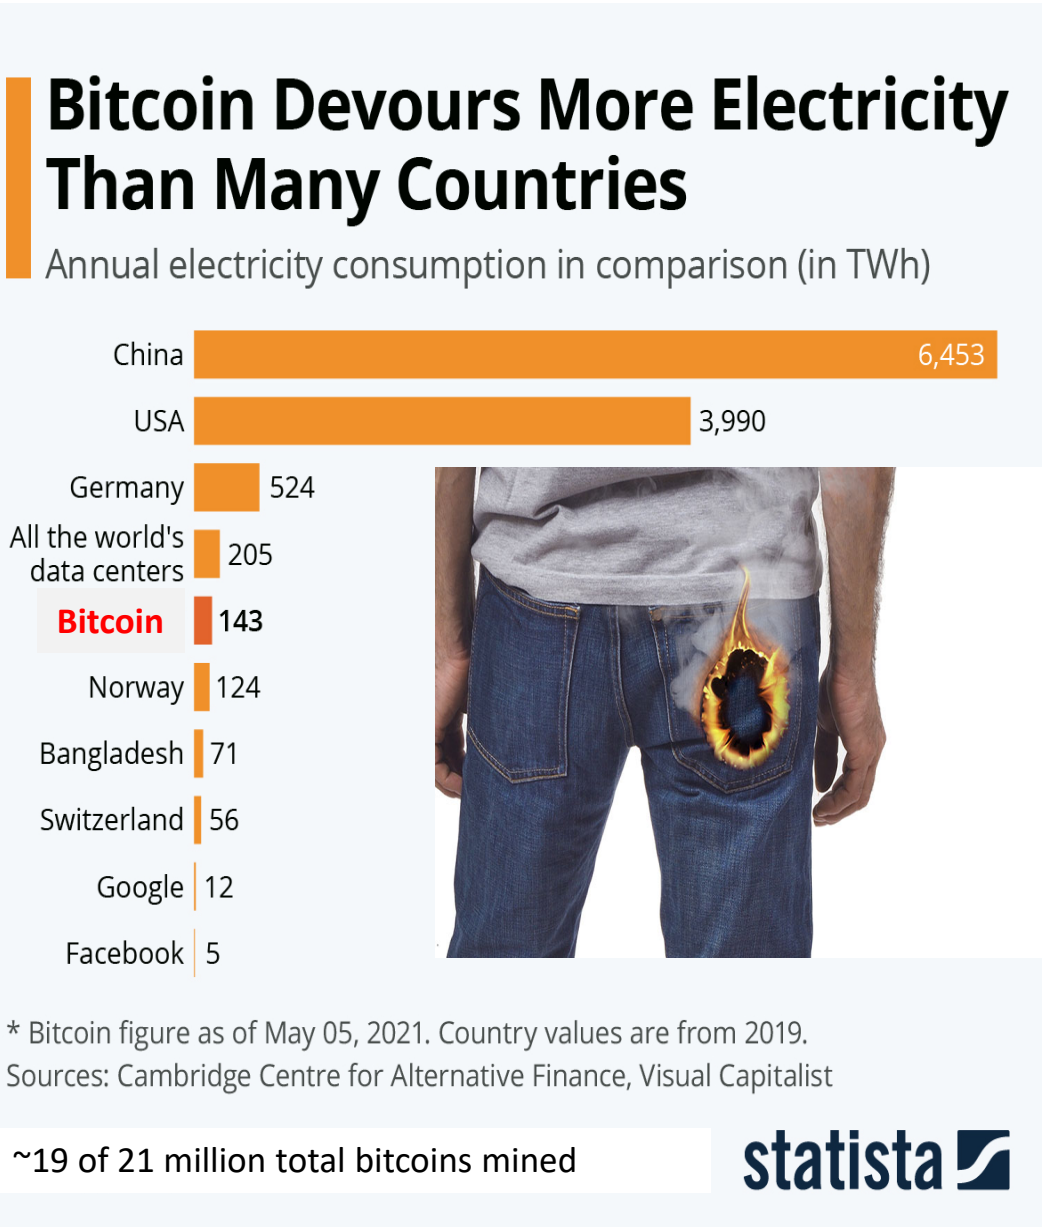

Source: https://www.firetrace.com/fire-protection-blog/three-levels-of-data-center-fire-protection

Storage Outlook - Current Technology Scenario

Technology	Recording Technique	Roadmap Capability	Future Developments	Challenges
NVM (MRAM, DRAM, NAND Flash, PCM, 3D-Xpoint)	Electronic Charge 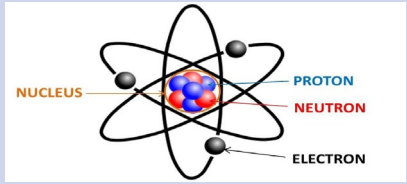	Aggressive development, multiple emerging technologies, CXL, NVMe(oF)	Multi-layer 3D stacking (256^{+}), faster garbage collection, new tiers?	Price
HDD	Magnetic Field 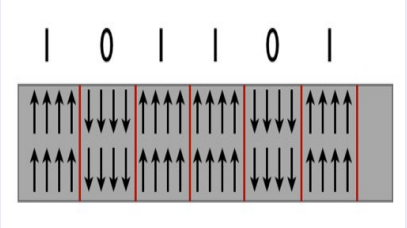	Performance and capacity growth challenges	HAMR, MAMR, (? Tb/in^2), multi-platters (9-11), zones, 2-4 actuators, ordered granular, bit patterned, cold HDD?	Access density (IOPs), TCO, high energy consumption, \$/TB/watt, CO_2
Tape	Magnetic Field 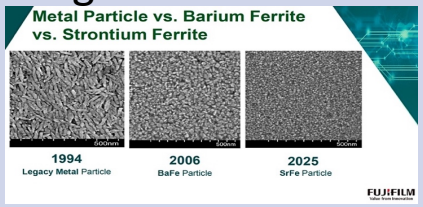	Well defined and sustainable capacity growth, 580 TB demo, high patent activity	Strontium Ferrite (SrFe), Epsilon Ferrite (ε-Fe2O3), deep archive, RAIL, erasure coding, Geo-spreading	Access time, customer awareness, *The race to \$0/TB…*
Optical Disc	Reflective Spot 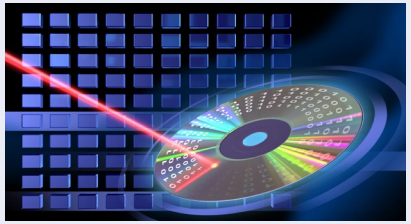	Slow progress compared to magnetics, <u>Not</u> presently a data center technology	Photonic (fluorescent) recording has most potential for optics, multi-layer, EMP proof media	Price, performance, capacity, reliability, throughput, slow learning curve

Source: Horison, Inc.

Tape is Best Positioned for the Zettabyte Era

Function	Benefits Summary
Price/TCO	Tape Has the Lowest Acquisition Price $/TB, Lowest TCO.
Performance (Access time)	Much Improved Access Times- Active Archives, Fastest Data Rates, RAIT, Smarter and Faster Robotics, RAIL, New Time to 1st Byte Features (oRAO, TAOS).
Capacity	LTO-9 Cartridge Capacity @18 TB (45 TB compressed) with 400 MB/sec Data Rate. Exabyte Capacity Libraries are Available. Tape Lab Demonstrations Reach 580 TBs.
Scalability	Tape Easily Scales Capacity (PBs to EBs) by Adding Media *Without Adding* Energy Consumption, HDDs Add Capacity by Adding Drives and Increasing Energy Consumption.
Energy, CO_2 Sustainability	Tape Uses Much Less Energy and Has Much Lower Carbon Footprint Than HDDs (~85% Lower).
Portability	Tape Media Easily Portable in Case of Disaster, HDDs Difficult to Physically Move.
Cybersecurity	Tape Air Gap, WORM and Encryption Defend Against Malware Attacks, Provide Immutability.
Durability/Media	LTO Reliability BER ($1x10^{19}$) Has 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.
Open Standards	LTO and LTFS Provide Open Standard File Interface and APIs. SW (S3 API) Supports Tape Object Storage.
Tape and Cloud Ecosystem	Tape Interfaces Seamlessly With Clouds Using Industry Standard API's. Native Cloud Applications Can Write *To a*nd Read *From* Tape.

Source: Horison, Information Strategies 6/2022

Zettabyte Era Fueling the Secondary Storage Disruption

A New Storage Frontier Unfolds

Future Success or VMF…? | 580+ TB Tape | DNA | Photonics | Glass | 3-D Holographic | *TBD…*

New Format?

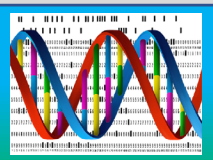

Fluorescent

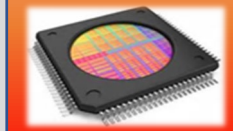

Offline Tape
(mins. - hours)

Tape Library
(secs. - mins)

HDD
(ms.)

SSD/NVM
(micro secs. μ)

$

Access Time Performance

Fastest

Data Preservation Time

100 + years

~5 years

Highest | Price per Unit of Capacity | Lowest

Source: Horison Inc.

The Zettabyte Era

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

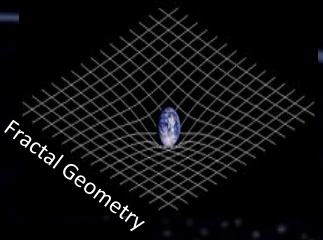

Galactic Data

Geopbyte
$1x10^{30}$

Brontobyte…
$1x10^{27}$

Yottabyte…
$1x10^{24}$

Zettabyte…
$1x10^{21}$

Exit…

>186,282 m/s
FTL

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
- ZettaScale+ Storage and Compute
- New Storage Paradigm

You Are Here

Artificial Intelligence

SEISMIC SHIFTS

Zettascape and Beyond

REMEMBER

Things are Changing so Fast
Even the Future is Obsolete

Source: Horison Inc.,
Fred Moore President