HORISON
Information Strategies

STORAGE OUTLOOK

Fujifilm's 9th Annual Global IT Executive Summit

September 13-16, 2017
Motif Seattle Hotel

Fred Moore
President
Horison Information Strategies
Horison.com

A New World of Information Growth

Digital Universe now doubling every two years.
Less than 5% of data is *ever* analyzed or touched again.

Nearly 20% of corporations store more than 1 PB.

2016 model automobiles averaged 14 million lines of code.
Autonomous may reach 100 million.

1.2 trillion photos in 2017 – 9% cagr.

Surveillance = 2.5 EB/day in 2019.

The IoT should connect over 25 billion "things" as Internet nodes by 2020.

Now has 6.4 million developers.

$11 trillion market by 2025 est.

7.4B WW Population
3.6B Internet Users
6.8B Mobile Phone Users

USA 324M

China 1.4B

India 1.3B

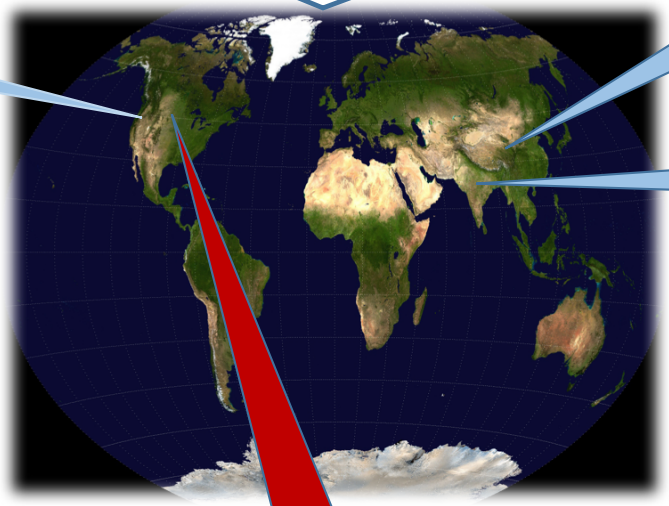

Data centers use ~2% of all US electricity

World's largest data center in Chicago is 1.1M sq. ft.

By 2020, 1/3 of all data will pass through the cloud.

Average lifespan
Server - 3 years.
HDD 4-5 years.
Tape Drive 7-10 years.

Over 4,000 Cyber-attacks every day.
Annual Cyber-crime cost at $400B.

Over 90% of all Internet traffic is unwanted material.

The Hyperscale Effect Re-shaping the IT Industry.

The human genome takes 600 GB to digitize.

Megabytes Gigabytes Terabytes Petabytes Exabytes Zettabytes Yottabytes ($1x10^{24}$) Xenotabytes ($1x10^{27}$) Shilentobytes ($1x10^{30}$).....

Hyperscale Data Centers Arrive - in a BIG Way

Trend Toward Fewer but *Much* Larger Data Centers

- **A Hyperscale Data Center is an enormous distributed computing environment – scaling from PBs to EBs.**
- **Massive infrastructure - over 400,000 ft^2, largest is >1.1 million ft^2 (= 19 football fields).**
- **Scale compute and storage independently.**
- **Designed with redundant components – if a failure - workload moves to another server.**
- **Hyperscale storage typically serves millions of users with a few applications.**
- **Amazon Web Services, Microsoft, IBM, and Google collectively control more than half of the WW cloud infrastructure service market.**
- **Unprecedented energy consumption – requires advanced cooling systems, redundant power - and lots of tape.**
- **Tape will be critical to sustainable hyperscale costs and growth!**

Source: Horison, Inc.

Hyperscale Heatwave - IT Energy Consumption Accelerates

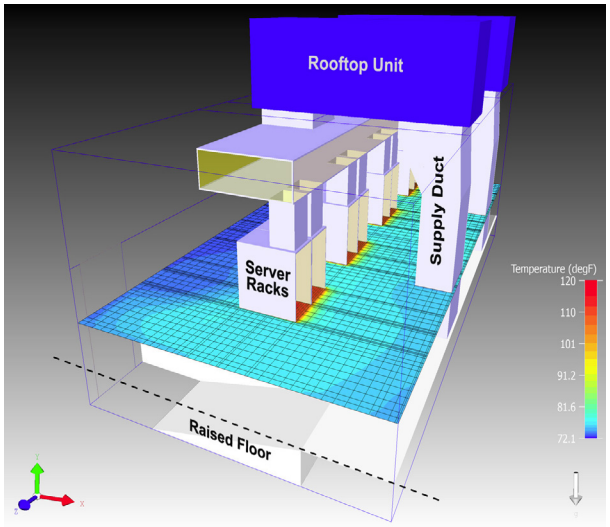

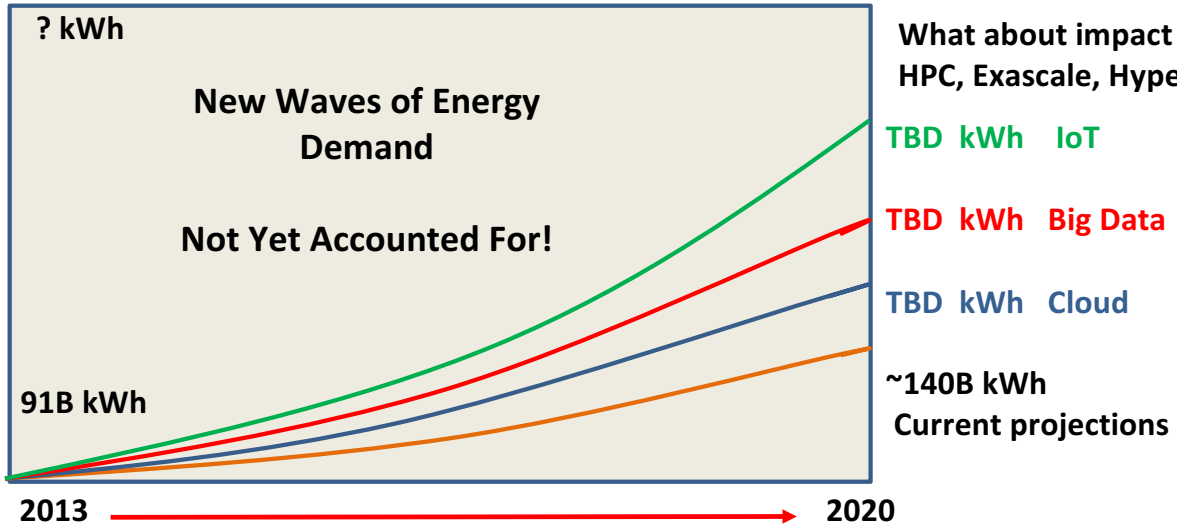

- Several New Waves of Storage Demand and Cloud Will Increase Data Center Energy Consumption.
- The World's Data Centers Now Consume Almost as Much Energy as Spain.
- In 2013, U.S. data centers consumed an estimated 91 billion kilowatt-hours of electricity, equivalent to the annual output of 34 large (500-megawatt) coal-fired power plants and 2% of all US electricity.
- Data center electricity consumption is projected to increase to roughly 140 billion kWh annually by 2020 - based on current trends.
- Does not include IoT, Big Data and Cloud demand.

Source: Horison, Inc.

New Data Center and Storage Energy Innovations

Geography Plays a Key Role

Microsoft Underwater Data Center

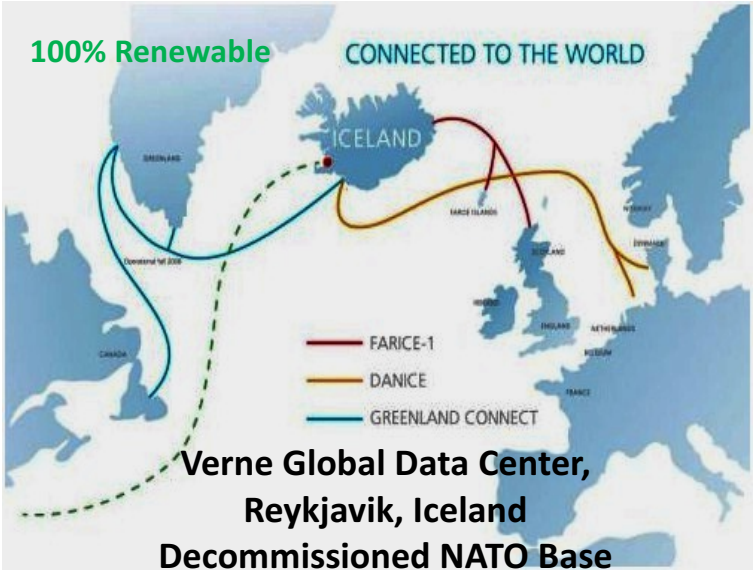

Verne Global Data Center, Reykjavik, Iceland
Decommissioned NATO Base

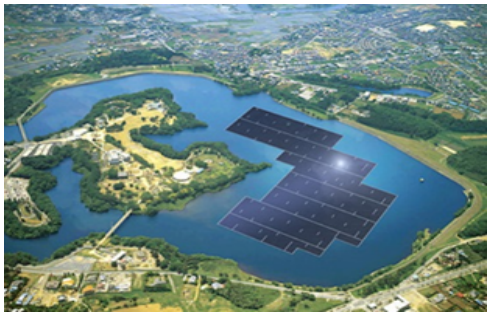

Floating Solar Panels
Yamakura Dam, Japan

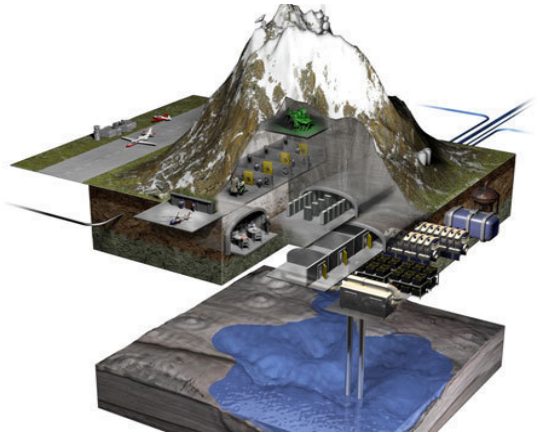

Swiss Fort Knox is Concealed in a Nuke-proof Bunker

Solid Granite Data Vault Wasatch Mtns. Utah

Bahnhof Modular Data Center, Kista, Sweden

The Global Datasphere

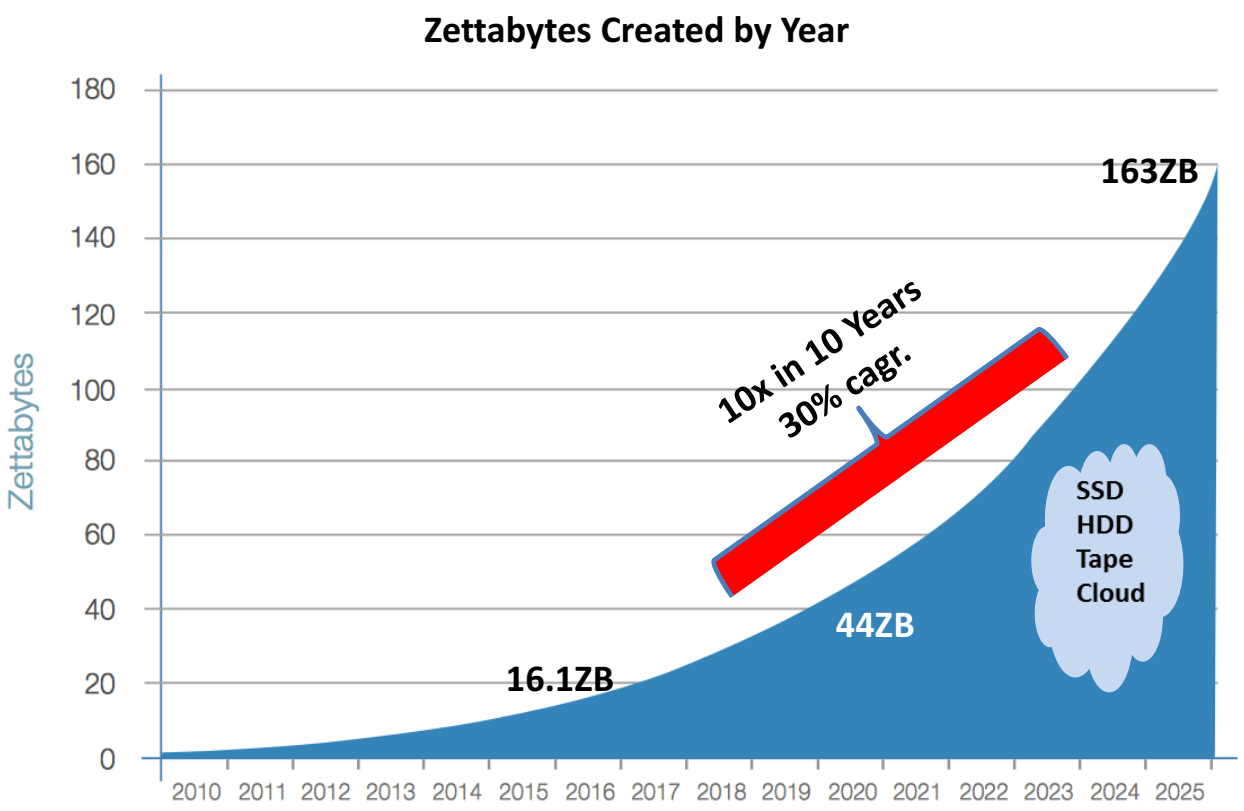

Data Classification by Business Value

Mission Critical!	15%
Business Critical/Vital	20%
Sensitive	25%
Archive, Long-term	40%

Not All Data is Created Equal!

Cloud Plays Major Role

The percentage of data that is processed, stored, or delivered by public cloud datacenters will nearly double to 26% from 2016 to 2025.

About 60% of all data best suited for tape.

The 163ZB projected in 2025 is only the beginning!

Source: Storage Newsletter 04/2017

Digital Universe by Data Class and Tier

Percent of WW Digital Data by Storage Tier In 2025.

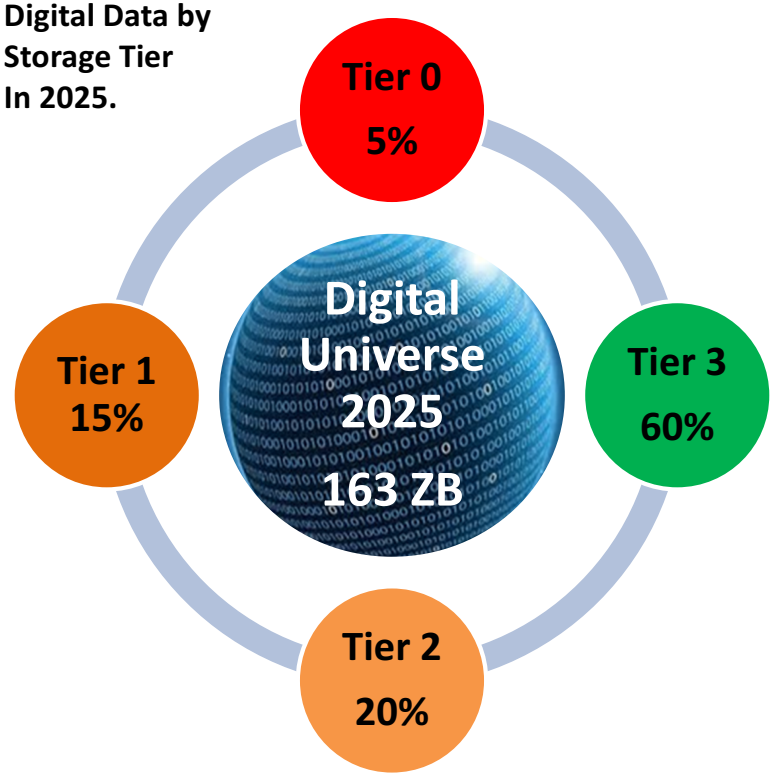

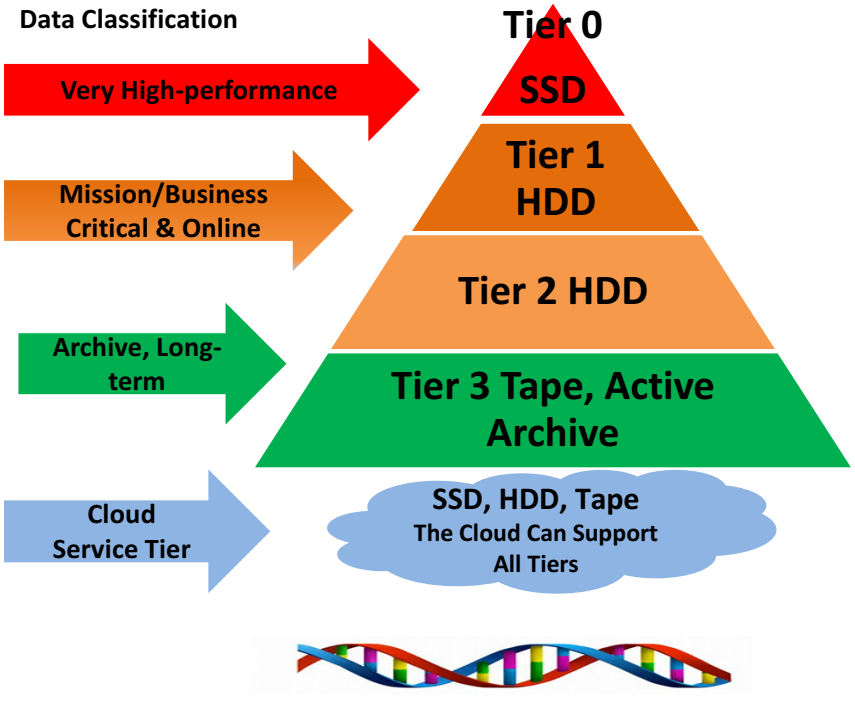

At Least 60% of All Data is Archival

Source: Horison, Inc. The Digital Universe IDC

The Data Classification Model

Tier 0 and Tier 1	Tier 2	Tier 3 (Archives – TB, PB,EB, ZBs..
Primary Storage	**Secondary Storage**	**Long-term Retention**
SSD, Enterprise HDD, Mission-critical, OLTP, DBs, Hi-perf. apps	Midrange HDD, Backup/recovery, Moderate Perf. Apps, Test and Development, Big Data	Tape, Offsite vaults, Archives, Permanent Copies, Big Data, Compliance, Video, Master Copies, RIM, Cold Data…

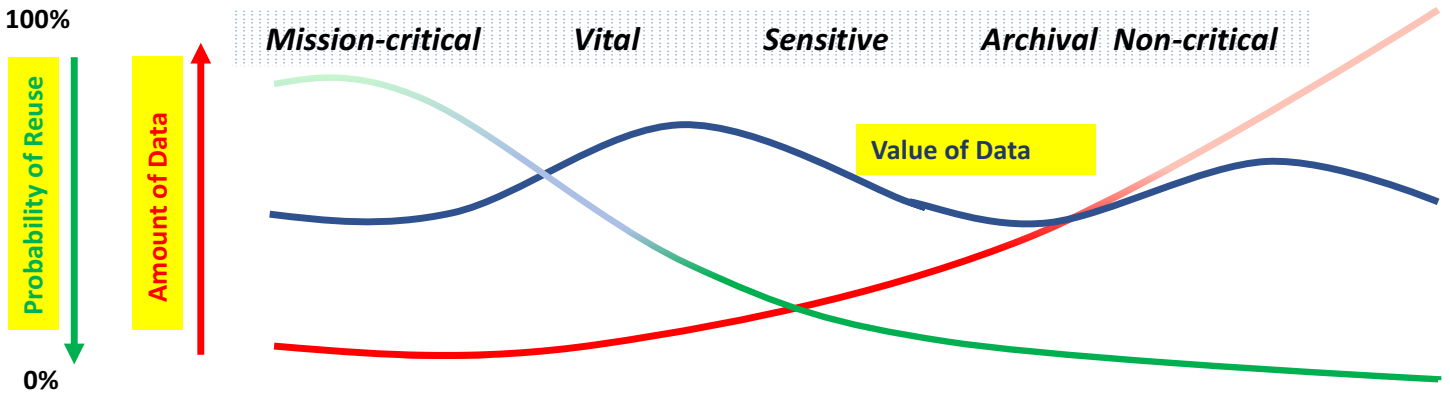

Days Since Creation	0 days	30+ days	90+ days	1+ years	Permanent and Infinite Storage
Response Time	ms.	ms. -secs.	Minutes	Hours	Hours-days
Components, Challenges	Policy engine, Tiered storage, Availability, BUR, Data Security		Data movers, Active archive	Re-mastering, Enormous Data Migrations	

Creation Security/encryption Long-term Storage Destruction

Source: Horison, Inc.

Solid State Disk (SSD) Scenario

Storage Class Memory Evolving Fast!

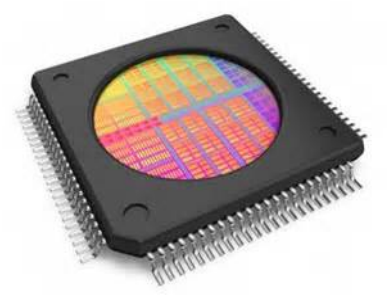

2D Flash Memory

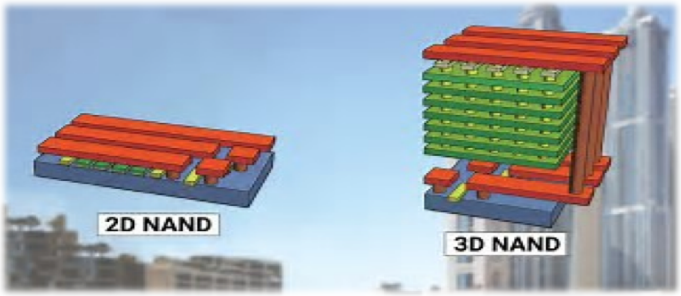

3D Flash *is next…*

Followed by 3D XPoint

- **Tier 0 is All About Performance - IOPS Intensive Apps, Databases, OLTP, HPC Burst Buffer.**
- **All-flash Arrays (AFAs) and Hybrid Flash Arrays (HFAs) - Showing Explosive Growth in all Markets.**
- **SSD capacity shipped was 19% of WW total HDD capacity shipped in 2016.**
- **Non-volatile, Low Power (~1/3 of HDD).**
- **Read Access Times: .2 ms Approx. ~50x Faster Access Than HDD (Some Write Fatigue Limits).**
- **Flash Success Heavily Impacting HDD Sales From Enterprise to Desktop.**
- **NVMe Optimized for Enterprise and Client SSDs.**
- **3D NAND arriving - Stacks Flash Cells Vertically in Multi-layers. Up to 512 layers planned.**
- **3D Xpoint to Break Storage and Memory Barrier – NAND Density and DRAM Speed.**

Source: Horison, Inc.

HDD (Disk) Storage Scenario

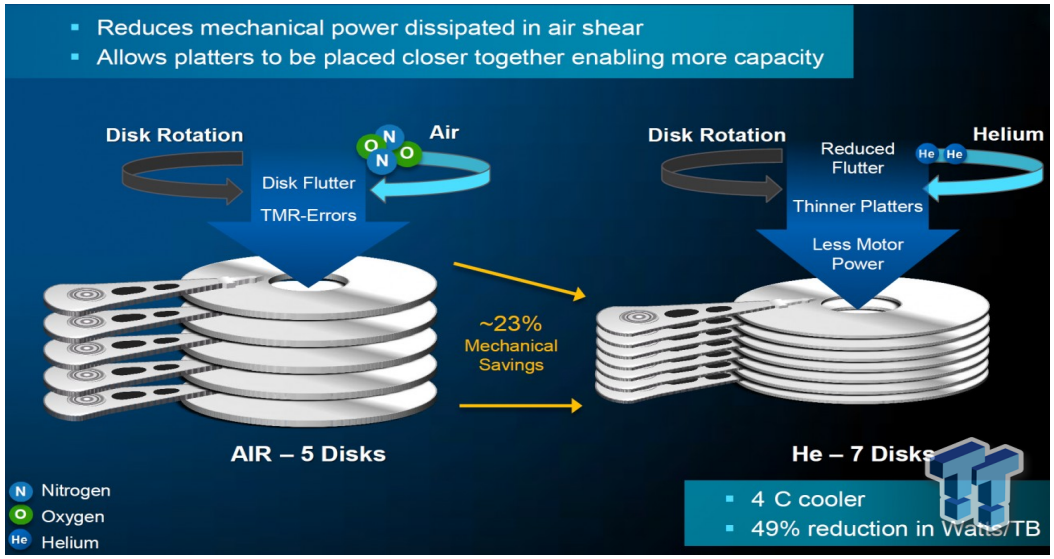

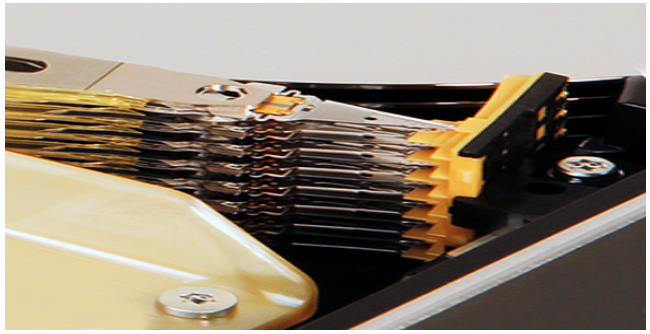

Helium-filled HDDs Contain More Platters.

The Drive above has 7 Platters Rather Than the Usual 3, 4 or 5.

- **HDD Capacity Growing 20-30% - but Areal Density (~1.2 Tb/sq. in^2).**
- **Capacity Growth Rate is Slowing – Add More Platters, Helium to Increase.**
- **Current Maximum HDD Capacity at 12 TB or 14 TB SMR – using 8 Platters.**
- **Drive Performance is Not Improving.**
- **Low HDD Utilization (<50%) Increases End-user Costs.**
- **RAID Rebuild Times Can Take Several Days, Erasure Coding to replace RAID?**
- **Reliability (BER) Has Fallen Behind Tape.**
- **Remember - HDDs Can Address All Data Types and Requirements.**
- **HDD Shipments Declining (-37.5% since 2013) as Flash Increases.**

Source: Horison, Inc.

New Disk Concepts Are Arriving

But None Address Performance…

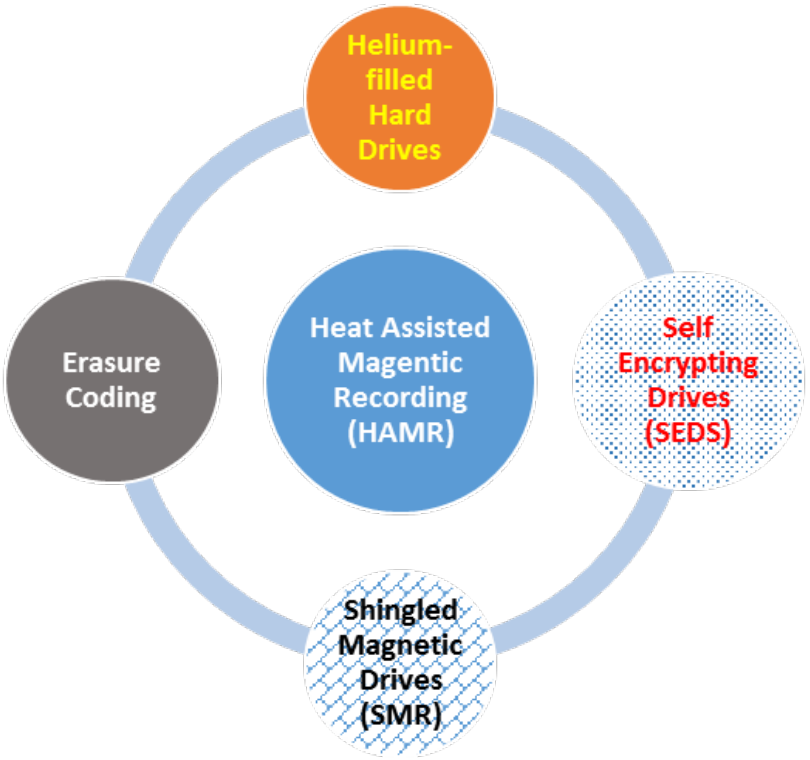

Helium Filled 	Helium-filled drives use less power to spin disks (which spin easier thanks to less resistance), they run cooler, and can stack more disks in the same space.
SEDs 	SEDs use disk drive controllerhardware to encrypt all data written to the magnetic media and decrypts all the data read from the media automatically. **Minimal market acceptance.**
Shingled HDD 	SMR achieves higher areal densities by overlapping tracks, analogous to shingles on a roof, allowing more data to be written to the same space. **Slow writes.**
Erasure Coding 	Data is broken into encoded fragments and stored across different locations, such as HDDs, storage nodes or distant geographical locations. **To Replace RAID?**
HAMR 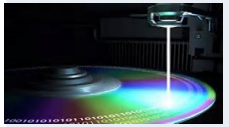	HAMR uses a small laser to heat the part of the disk surface that is being written to remove the super-paramagnetic effect. **Several issues – still under development.**

Source: Horison, Inc.

Magnetic Tape Storage Scenario

LTFS
Linear Tape File System

- Over 85% of Tape Drive Shipments are LTO (96,000 PB – *96 EB* - of LTO shipped in 2016).
- Tape Media Lifespan at 30 Years or More.
- Tape Areal Density Demos Indicate Cartridge Capacities Over 200 TB are Achievable.
- Tape Drive Reliability (BER), Data Rate and Capacity Has ***Surpassed Disk.***
- Tape Capacities at 15TB Native, 45 TB Compressed, Areal Density > 9 Gb in^{2}.
- Tape Data Rates Reach 360 MB/sec. Native Streaming Mode – RAIT is on the Way.
- HDD TCO 6-25x Higher Than Tape.
- More Than 60% of All Digital Data Classified as Tier 3 (Archive, Fixed Content) – Fastest Growth Area.
- Clouds Embracing Tape Solutions for Better TCO and More Secure Archival Services.

Source: Horison, Inc.

Storage Eclipse

SSD and Tape Capturing Data from Traditional HDD Market

Further HDD Performance Gains are Minimal

HDD Re-build Times Excessive (several days)

HDD Capacity Growth Facing Density Limits

HDD Adding Platters to Increase Capacity

HDD TCO Higher Than Tape (4-15x)

SSD Capturing High Performance Data

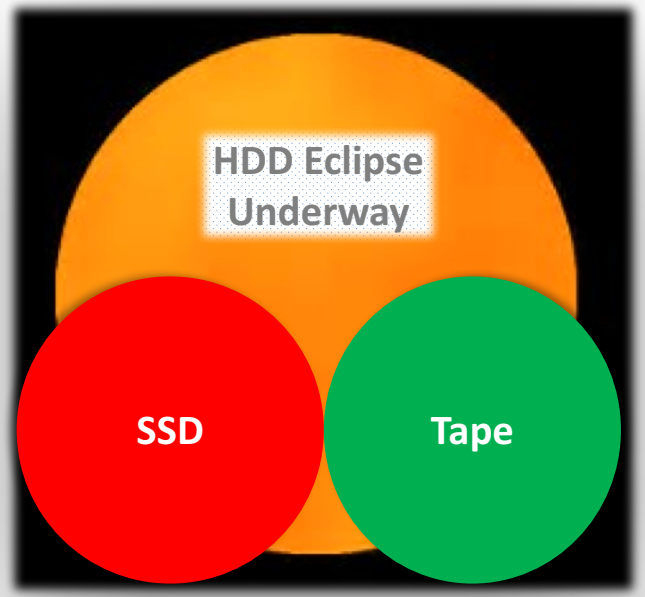

Poor Utilization (<50%) Requires More HDDs

HDD Data at Rest is Main Target for Hackers

Tape Reliability Has Surpassed HDD

Tape is Much Greener Than HDD

Tape Media Life Now 30 Years or More

Tape Capturing Archive and Permanent Data

Source: Horison, Inc.

HDD Units Shipment Forecast

Projecting a 17% Decline 2016 – 2020

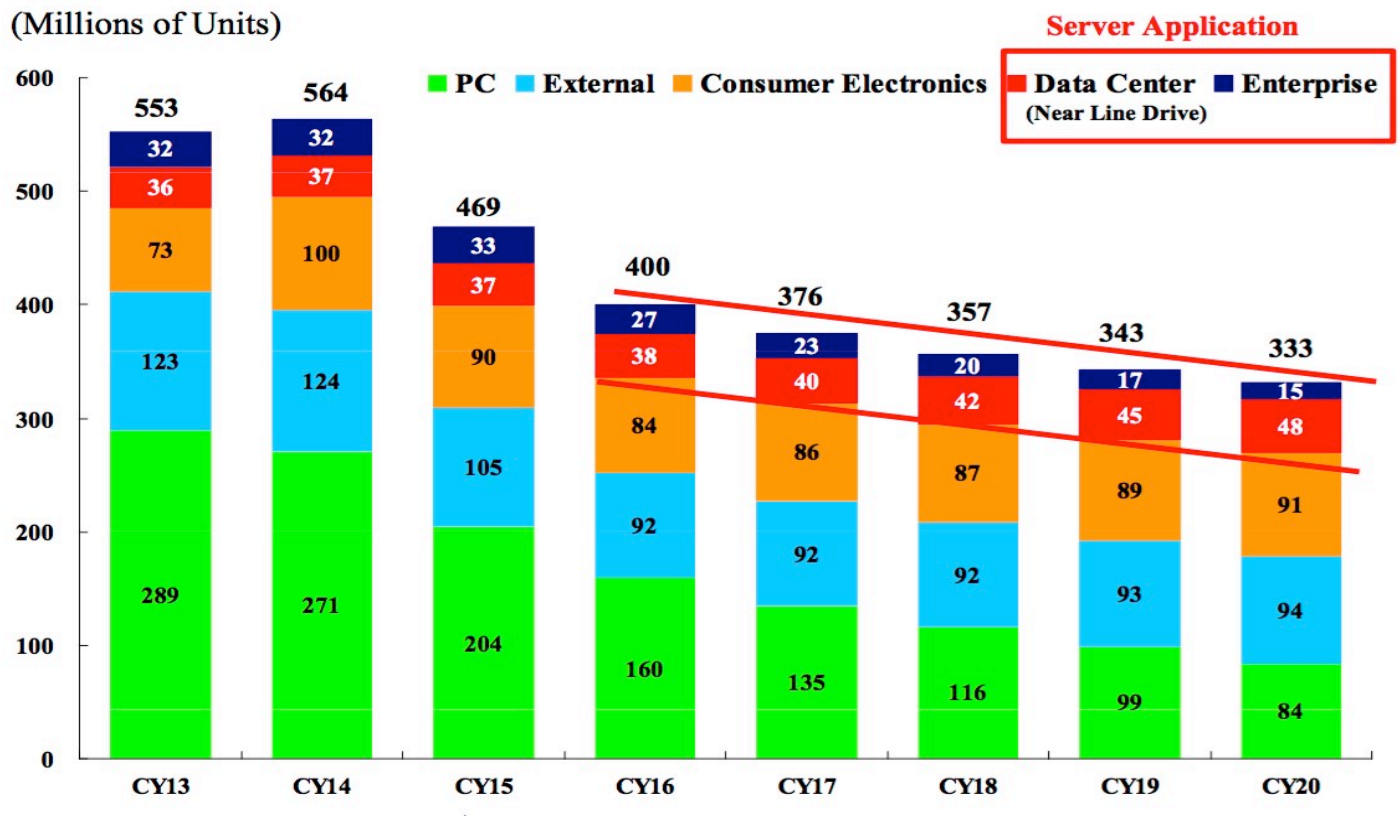

Key Factors

- **Declining PC demand for HDDs**
- **Strong Flash Demand in All Categories**
- **Preliminary 1Q17 HDD Results: 98.5 to 100 Million Units Shipped**

Source: Nidec, Trendfocus

Storage Reliability Levels

Bit Error Rate	Device	Media Lifespan Average
$1x10^{19}$	Enterprise tape – T10K, TS11xx, LTO-7 (Approaching $1x10^{20}$)	> 30 years
$1x10^{17}$	LTO tape – LTO 5, 6 Flash SSD	> 30 years 5-10 years
$1x10^{16}$	Fibre Channel and SAS HDD	3-5 years
$1x10^{15}$	Enterprise SATA	3-5 years
$1x10^{14}$	Desktop SATA	3-5 years
$1x10^{12}$	Blu-ray Optical Disc	> 50 years
$1x10^{10}$	TCP/IP (correctable after *n* retries with degradation)	NA

Published Values for BER – Bit Error Rate
BER Metric for HDD and SSD – Number of Sectors in Error per Total Bits Read
BER Metric for Tape and Optical – Number of Bits in Error per Total Bits Read

Tape is Now the Most Reliable Storage Technology Available

Source: Supplier Data, Horison, Inc.

Tape Subsystem Performance

Tape Library Mount Time	Tape Drive Load Time	File Access Time	Total Time to 1st Byte of File	Sequential Access Within File
4 - 10 secs	11 sec	10 - 100 secs	25 – 110 secs	*What About Data Rates*?

Robotic Mount and Load Time
Library Mount Time = 4 – 10 secs.
Drive Load Time = 11 secs.
File Access Time = 10 – 100 secs.

Drive type	Format	Max. Native Capacity	Sustained Data Rate	Throughput Density
LTO-7	Linear	6 TB	1.080 TB/sec	5.55
TS1155	Linear	15 TB	1.296 TB/sec	11.57

Source: Horison, Inc.

Magnetic Tape Future Projections

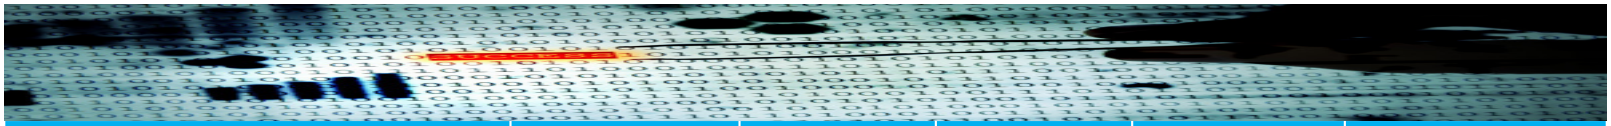

PARAMETER	2017	2019	2021	2023	Change rate (Learning curve)
Capacity (TB/cart)	16	32	64	128	41.00 %/year
Total data rate (MB/sec)	480.2	720.6	1,081.4	1,622.7	22.50 %/year
FC Roadmap (MB/sec)	6,400	12,800	12,800	25,600	
Recordable length (meters)	1,051	1,141	1,238	1,343	3.9% winding reserve
Track density (TPI)	15,652	22,868	33,543	49,372	21.17 %/year
Linear bit density (KFCI)	581	703	850	1,029	10.00 %/year
Total data tracks	6,639	9,856	14,660	21,842	21.95 %/year
Areal density (Gbits/inch^2)	9.09	16.07	28.52	50.80	33.28 %/year
Tape speed (meters/sec)	5.8	6.2	6.7	7.2	3.61 %/year
Time to fill a tape (minutes)	552	731	969	1,284	15.10 %/year
# Passes to end of media life	29,194	31,333	33,630	34,840	3.6 %/year
Tape width (mm)	12.65	12.65	12.65	12.65	No change

Tape Data Rates Surpass HDD

Tape Data Rates to Exceed HDD by 5X

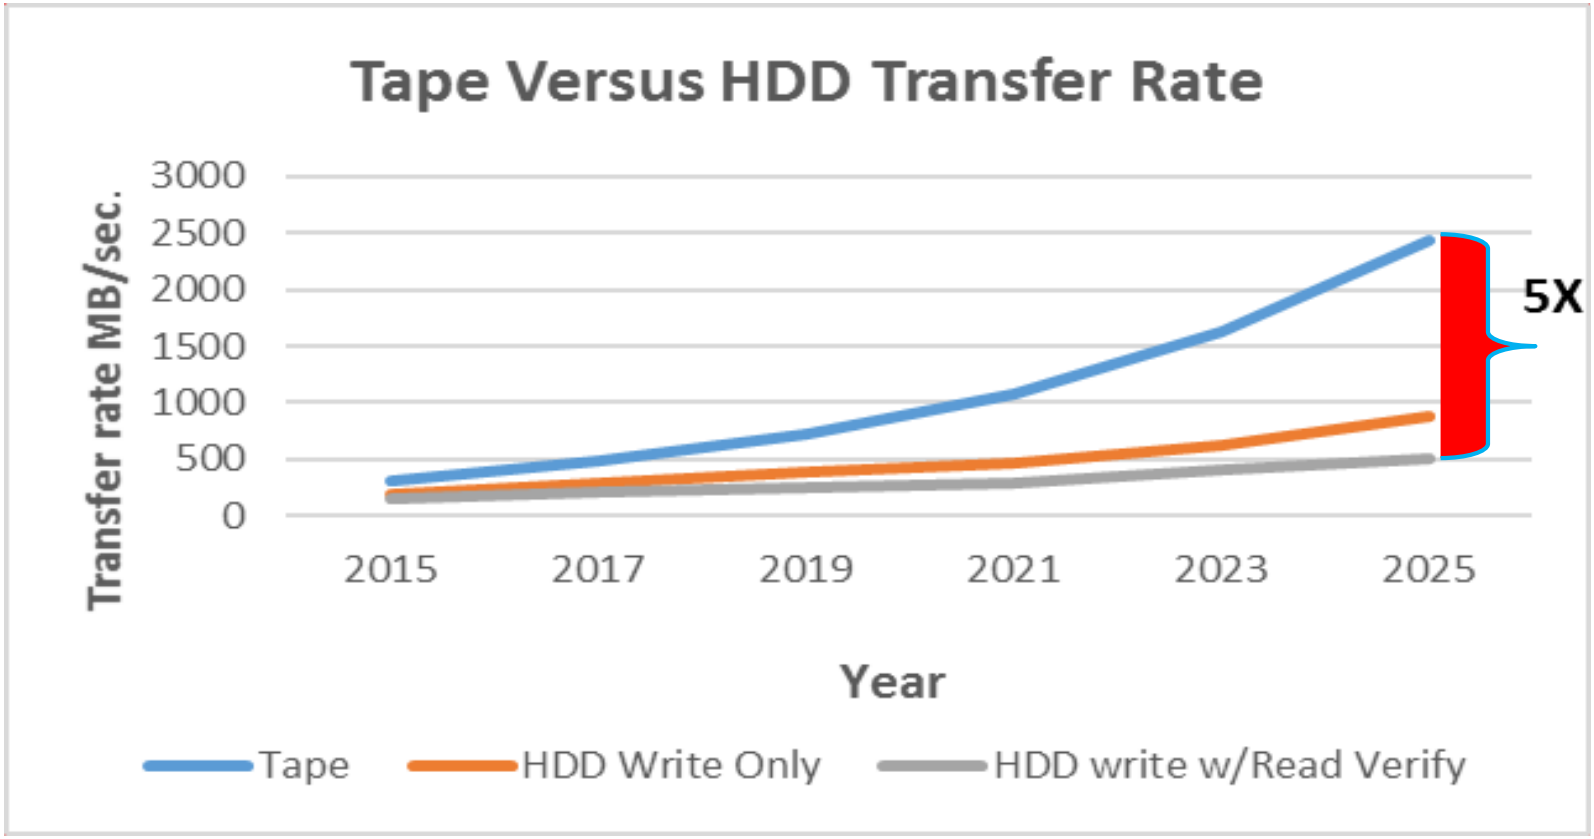

Source: TSC State of the Tape Industry Memo 2016

RAIT Is Gaining Momentum

Redundant Arrays of Independent Tape – Making Tape Data Rates Soar

RAIT 5

Striped Blocks | Distributed Parity

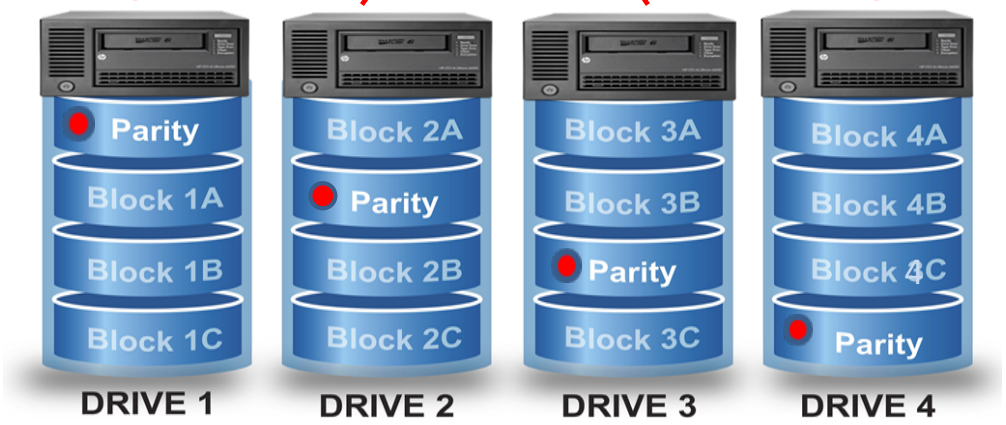

Parity Rotated Across all Cartridges in 4-drive Stripe

RAIT 0 Striping

Striped Blocks – No Parity

Tape Drives

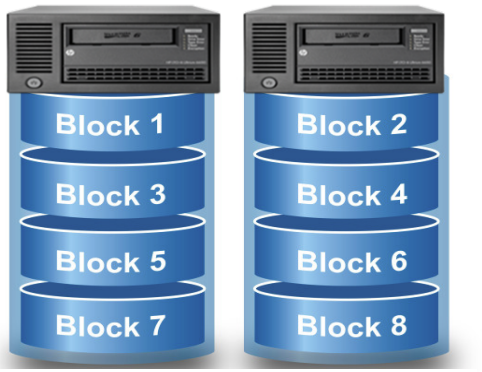

RAIT 5 and RAIT 0 are highest performance tape striping architectures

- Software segments files (blocks) and distributes them across *n* drives in stripe (*n* = 2, 4, 8..).
- Throughput - optimizes higher tape data transfer – RAIT 5 Ex: effective transfer rate increases up to 4x.
- Availability - fault tolerance provides higher availability – RAIT 5 Ex: 1 drive failure in 4 is permitted.
- Capacity efficiency = 80% with RAIT 5, 100% with RAIT 0.

Source: Horison, Inc.

RAIT Provides Higher Bandwidth and Higher Availability

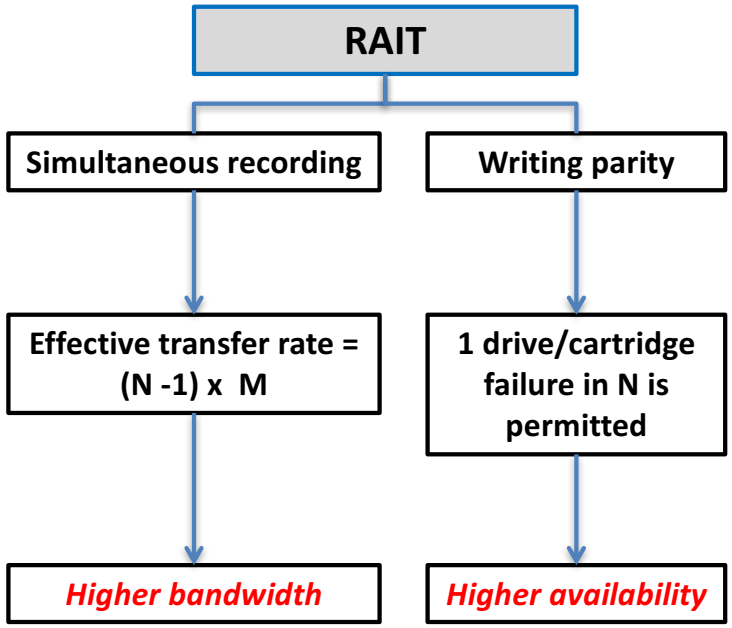

RAIT Effective Transfer Rate MB/sec.

Drive Type	M	N=3 E=67%	N=4 E=75%	N=5 E=80%
LTO-6	160	320	480	640
LTO-7	300	600	900	1,200
TS1155	360	720	1,080	1,440

Effective Transfer Rate = (N-1) x M
E: Efficiency factor for RAIT group. E = (N-1)/N
M: Transfer rate of single drive
N: The number of drives in a RAIT group. Min. N =3, typically N <= 5

RAIT Proof of Concept

FUJIFILM

Amanda backup software with RAIT 4 to write data to tape, with commodity PC server, Linux OS.

3 LTO-6 drives used, 160 MB/sec. 256K blocks, no compression, file size ~19GB.

Expected transfer rate = (3-1) x 160MB/sec.
= 320 MB/sec.

Results: Observed transfer rate was 320 MB/sec.
Backup time was 56 seconds.

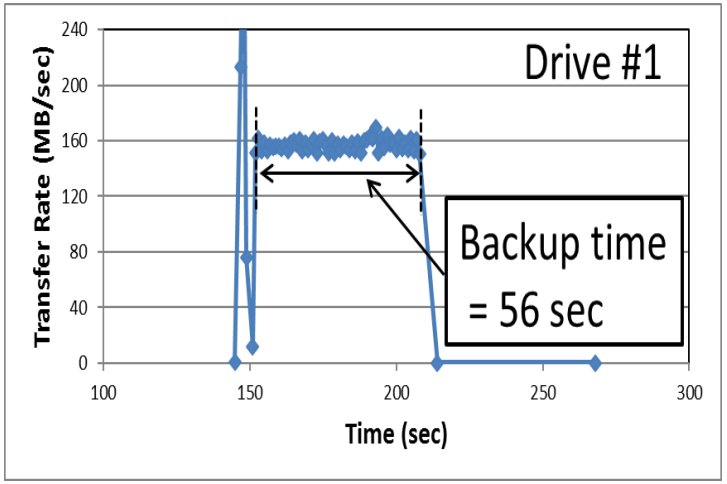

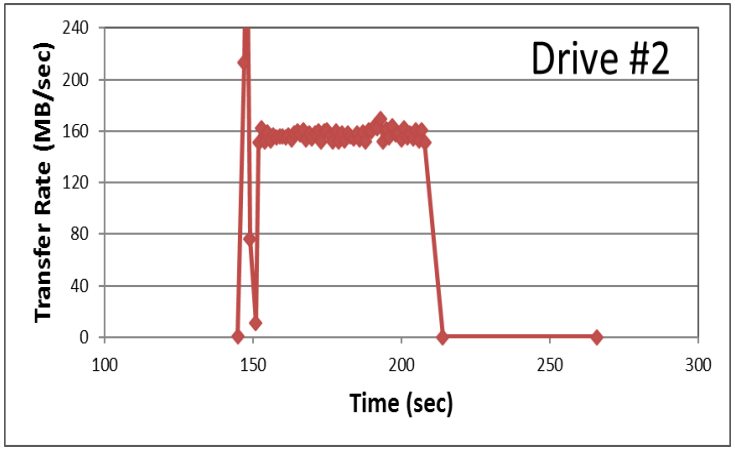

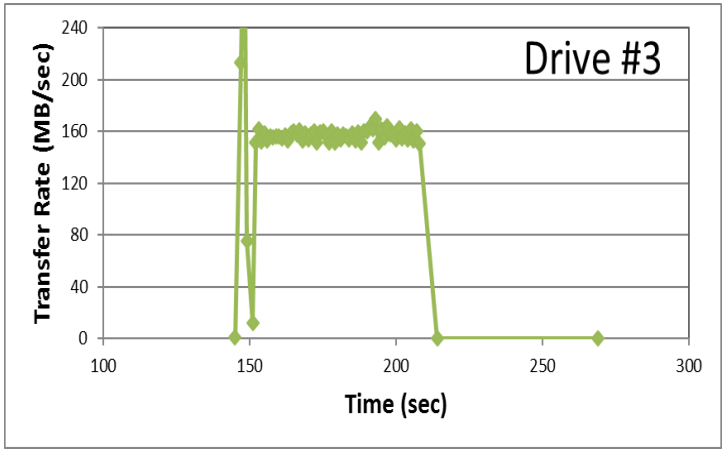

Source: FUJIFILM Recording Media USA., Inc. by Yuichi Kurihashi, Manager, Engineering/Technical

Recommended Access Order

Faster Tape Access/Recall times

Tape ordered recalls – recommended access order

- Enterprise tape drives now support Recommended Access Ordering (RAO)
- Multiple tape recalls are properly ordered by the tape drive to reduce recall time
- Tests show that RAO improves multiple file recalls by 30% to 60%
  - This SAME example illustrates 2:06 of tape movement without tape I/O
- **The tape drive includes a processor and memory storing a file location table for the tape cartridge.**

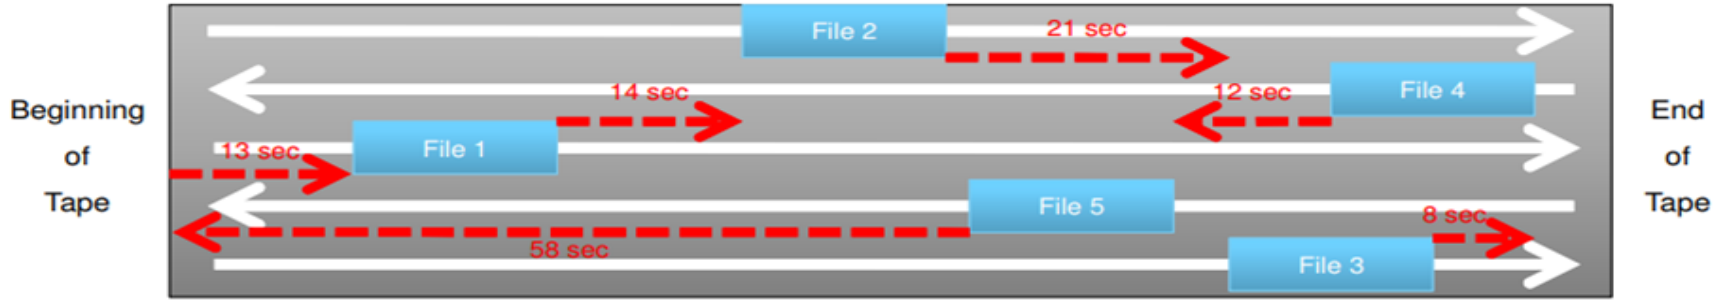

When Does Data Reach Archival Status?

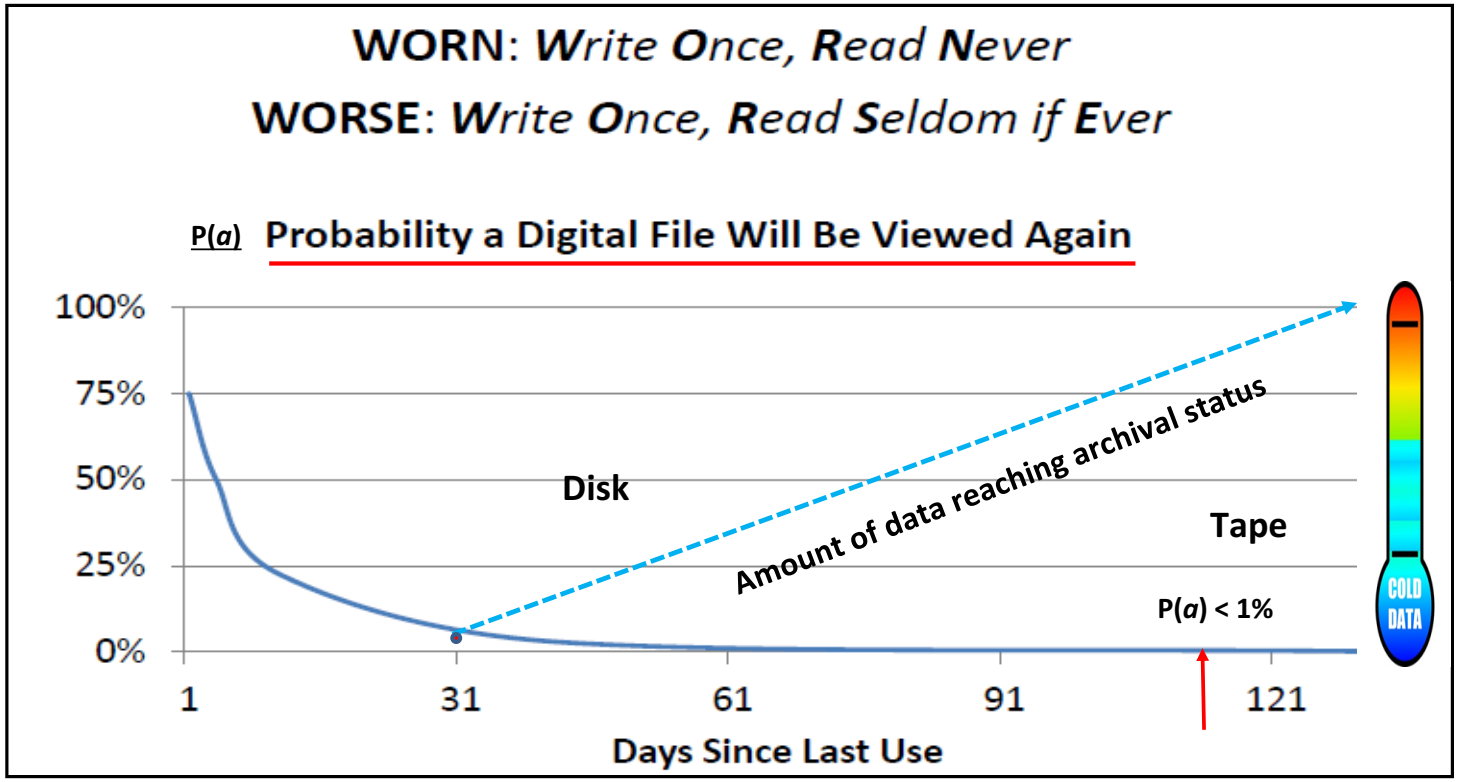

Source: Horison, Inc.

Deep Archive - a storage location for data that will probably not be accessed again, but must be kept as the final copy for preservation.

aka

Cold Storage - the retention of inactive data that an organization rarely, if ever, expects to access for years or perhaps decades.

The Active Archive

Combines HDD and Tape to Improve Archival Access Times

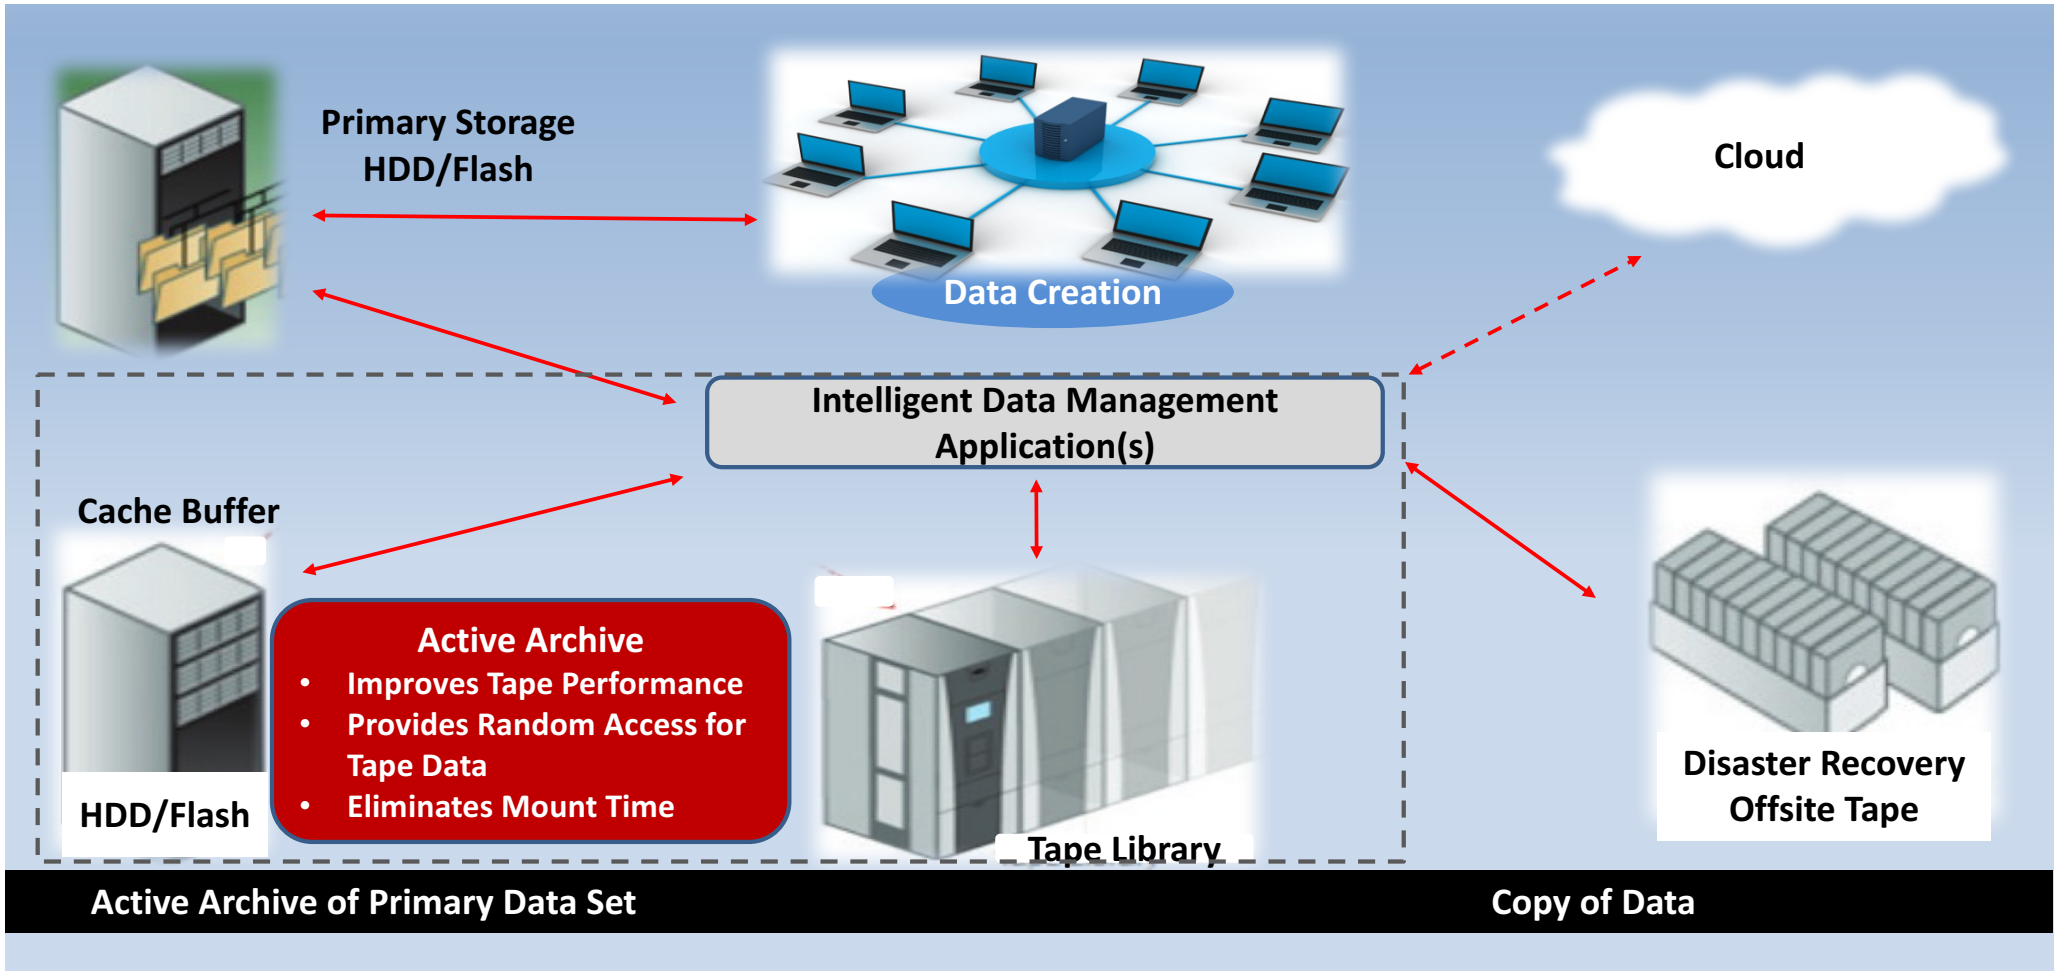

Tape Performance Enhancements

Active Archive 	**Provides HDD-like Access Time to 1st Byte and Random Access for 60-90% of Accesses (hit ratio).** **Increasing in value as digital archives swell.**
Faster Data Rates 	**Tape Data Rates Expected to be 5X Greater than HDD by 2025.**
LTFS 	**Drag and Drop Feature Provides Disk-like Access to Files.**
RAIT 	**Striping Multiplies Tape Drive Data Rates and Improves Availability for High-performance File Transfers.**
RAO 	**Improves File Recall Times (Time to 1st Byte) by 30-60% by Building Recommended Access Order List.**

Source: Horison, Inc.

The Growing Need for Data Protection

87% of Data Stored to Require Protection by 2025

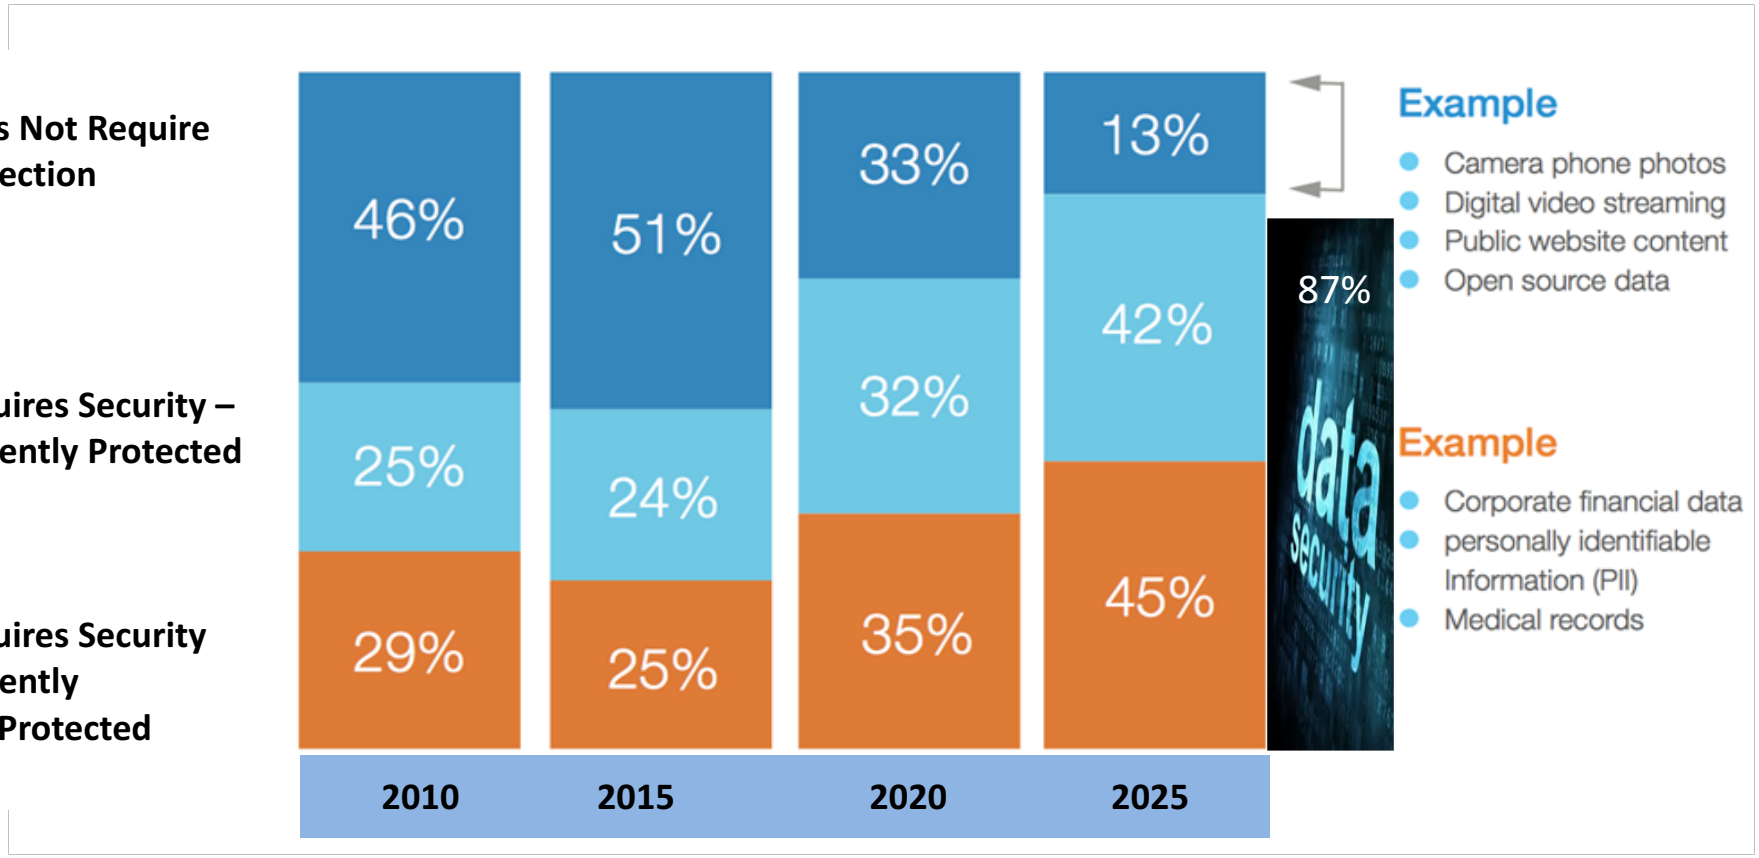

Source: Storage Newsletter April 7, 2017

Calling All Hackers!!

DEF CON 25 is July 27-30, 2017 at Caesars Palace in Las Vegas!

$260 USD for all four days! Cash only at the door, there is no pre-registration.

Contests, awards, prizes!

Over 16,000 attendees in 2017.

And FBI and NSA agents looking to recruit potential and up-and-coming cyber-culture hackers.

The Tape Air Gap = Data Security

- The Average Total Cost Of a Data Breach Was $3.79 Million In 2016.
- Cybercrime Will Become a $2.1 Trillion Problem By 2019!
- Estimated 4,000 ransomware attacks occurring each day.
- Tape Air Gap Prevents Unauthorized Electronic Access – Data Security.
- Tape is Going Greener and Deeper.

Online Tape Access

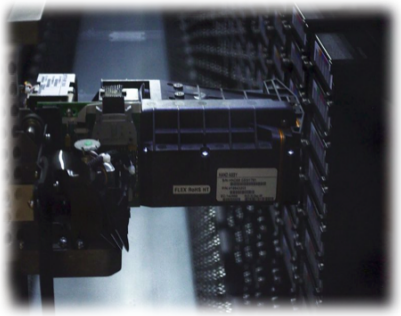

Manual

Robotic

AIR GAP

Brute-force Hack Attack	Catfish	Drive-by Download	Ghosting	Hash Busters
Keylogger	Malvertising	Man-in-the-middle attack	Pharming	Phishing
Ransomware	Scareware	Skimming	Smishing	Spear-fishing
Spoofing	Spyware	Vishing	Whaling	WikiLeaks… Human leaks

Tape's Role in the Cloud is Expanding

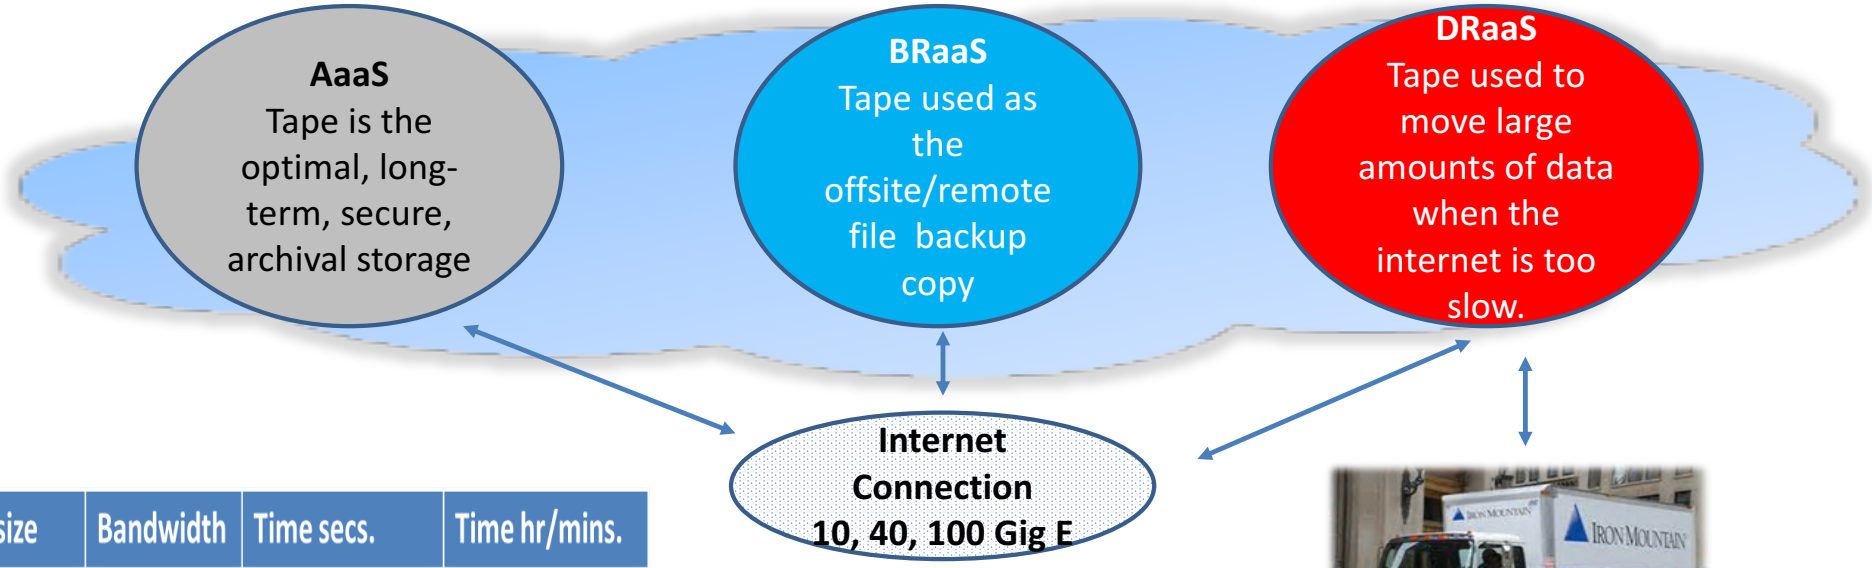

File size	Bandwidth	Time secs.	Time hr/mins.
1 TB 10 TB	10 Gig E	1,120 sec 11,200 sec	18 min 36 sec 3 hr 6 min 40 sec
1 TB 10 TB	40 Gig E	280 sec 2,800 sec	4 min 40 sec 46 min 40 sec
1 TB 10 TB	100 Gig E	112 sec 1,120 sec	1 min 52 sec 18 min 40 sec

Tape Providing Lower Cost, Highly Secure, and Unlimited Capacity Scaling for Service Providers.

State of Tape Industry - 2017

Function	Remarks Source: Horison, Inc.
Price/TCO	Tape Has the Lowest Acquisition Price $/GB, HDD TCO 4-15x Higher Than tape.
Performance	• Active Archive Addresses 1) Tape "Random" Access 2) Time to First Byte. • Tape Data Rates @360MB/sec. and Expected to be 5X Greater Than HDD by 2025. • RAIT Provides High Data Rate/Throughput Multiplier. • RAO Improves Recall and Time to First Byte. Needs LTO Support.
Capacity	Tape Cartridge Capacity Max. @15TB (45TB compressed) with 200X More Area, HDD @14TB. Lab Demos Demonstrate Native Cartridge Capabilities to Reach 100s of TBs.
Scalability	Tape Adds capacity by Adding Media, HDDs Add capacity by Adding Drives.
Reliability	Tape (BER) $1x10^{19}$ HDD (BER) $1x10^{16}$ Tape is 1000x Higher Than HDD!
Energy Usage	Tape Uses Much Less Energy Than HDDs, Can Move Tape Data w/o Electricity.
Portability	Tape Media Easily Portable, HDDs Difficult to Move.
Encryption/WORM	All Tape Drives offer Encryption, Option on Some HDDs But Seldom Used.
Cyber Security	Tape Air Gap Prevents Cybercrime Attacks.
Media Life	>30 Years for all Modern Tape, Avg. ~4-5 Years for HDDs.
Media Conversion	Tape Supports Current and 2 Prior Versions, HDD Requires Drive Replacement.
Recording Limits	No Foreseen Limits for Tape, HDDs Facing Areal Density Limits.
Cloud	Improves Cloud Reliability and Security, Lowers Storage Costs, Unlimited Capacity Scaling.

Remember …

Things Are Changing So Fast…

Even the Future is Obsolete

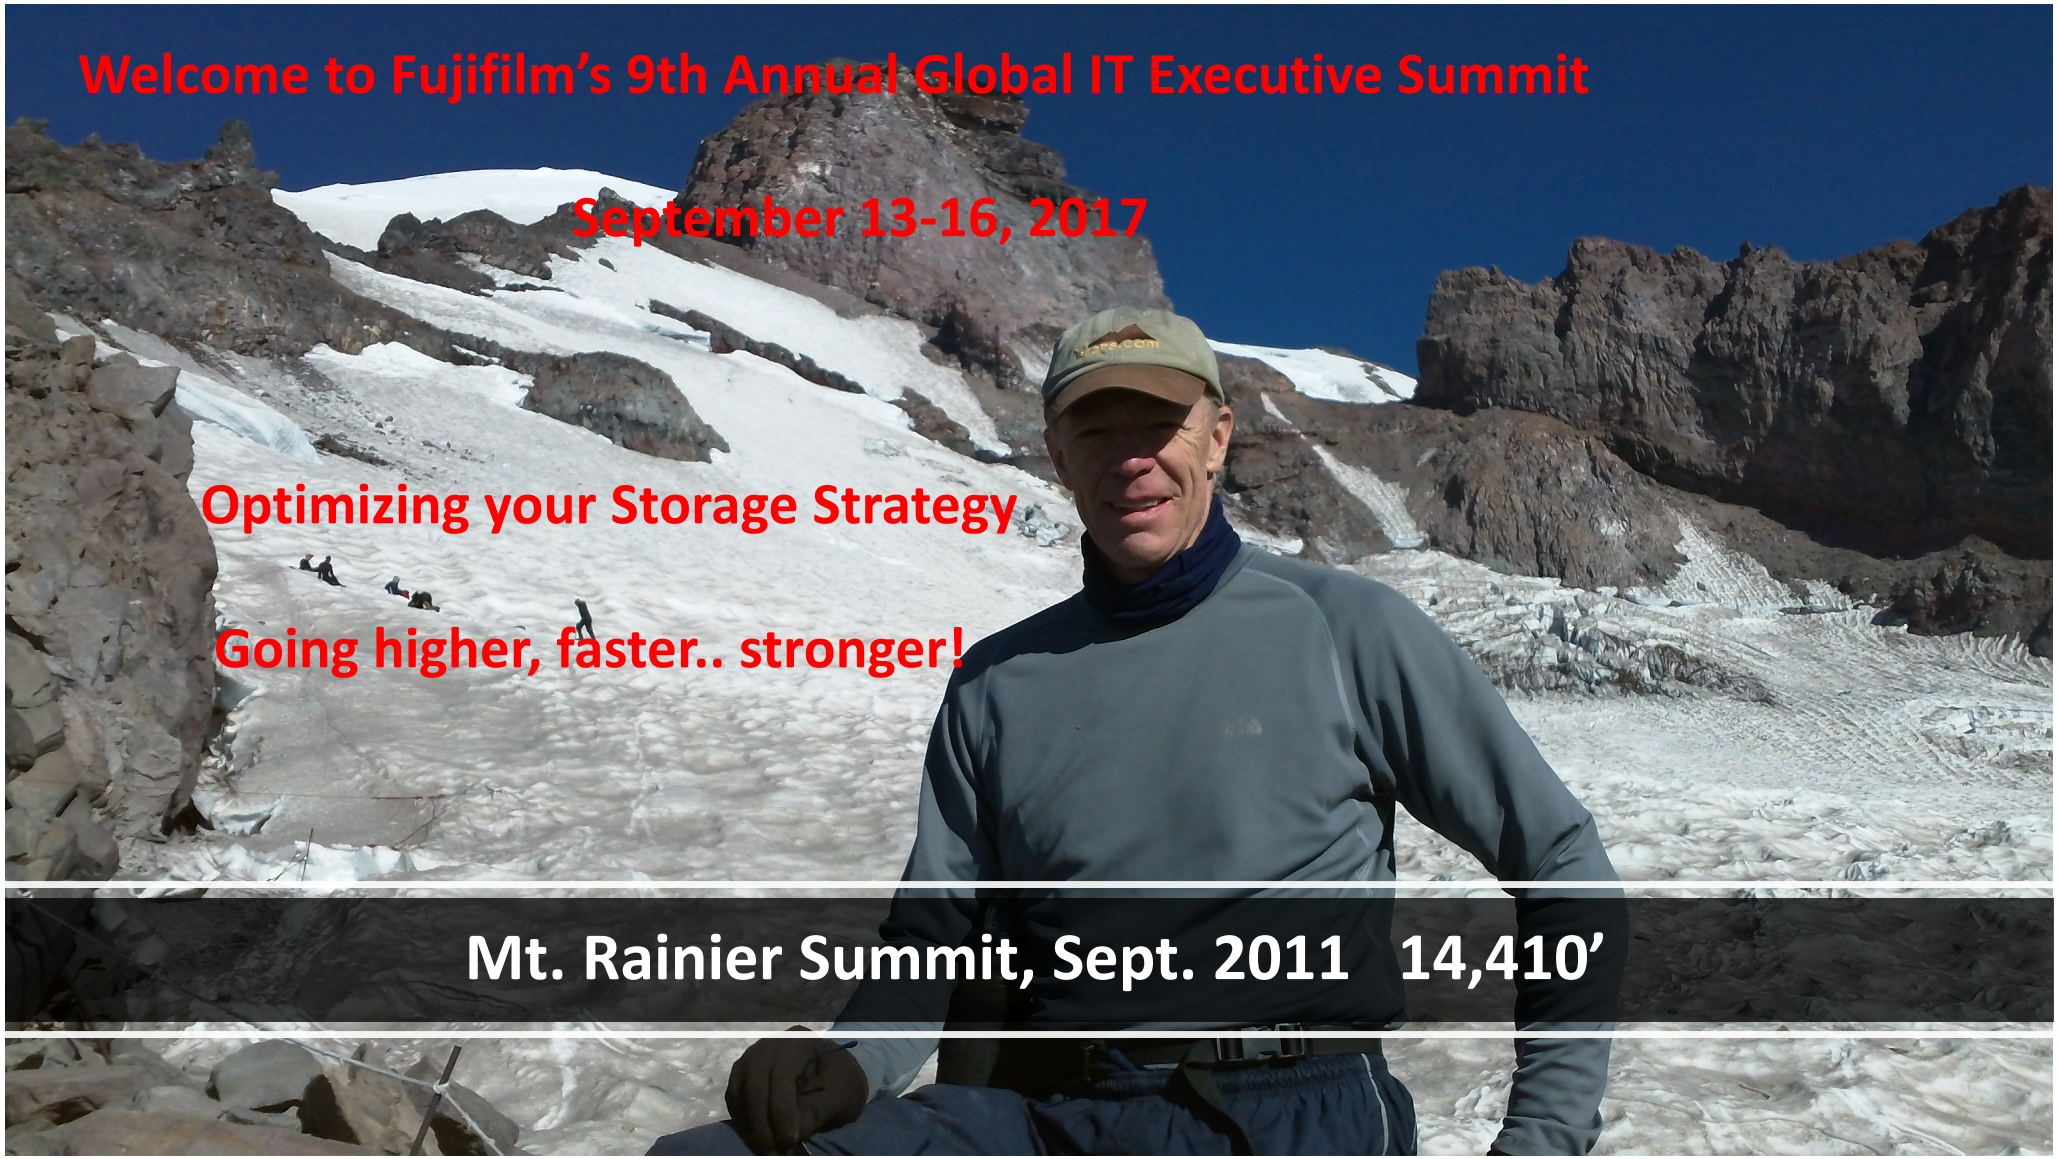
Welcome to Fujifilm's 9th Annual Global IT Executive Summit
September 13-16, 2017
Optimizing your Storage Strategy
Going higher, faster.. stronger!
Mt. Rainier Summit, Sept. 2011 14,410'