HORISON
Information Strategies

STORAGE OUTLOOK

FUJIFILM
Value from Innovation

Fujifilm's 8th Annual
Global IT Executive Summit
SEPTEMBER 21-24, 2016 • HYATT REGENCY BOSTON HARBOR HOTEL, BOSTON
Exploring the new world of storage!

Fred Moore
President
www.horison.com

A New World of Information Growth

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

Nearly 20% of corporations store more than 1 PB.

2016 model automobiles average 14 million lines of code.
Autonomous will use much more.

1.1 trillion photos in 2015 - 16% cagr.
Equals 1.1 roundtrips to the Sun.

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

Over 75% of all data generated by individuals. However 80% will become the liability of a large data centers.

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

USA 324M

China 1.4B

India 1.3B

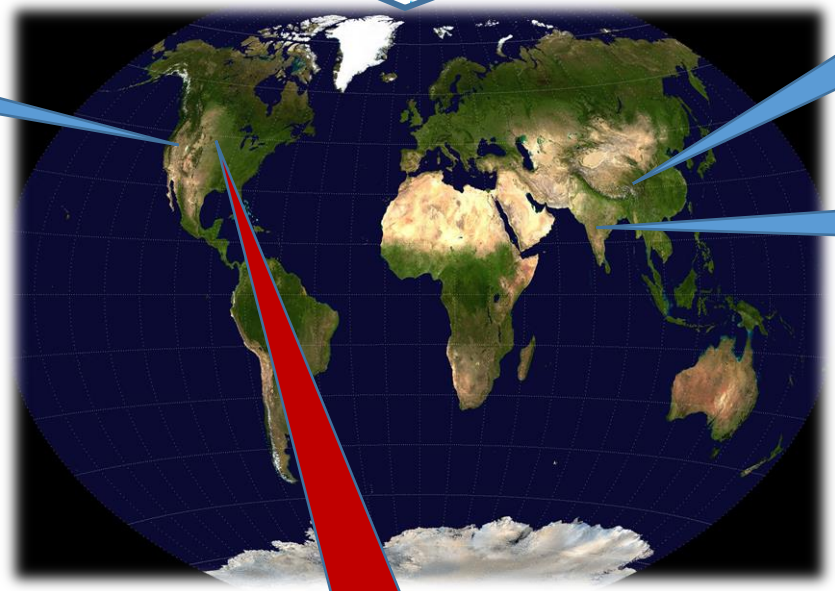

Data centers use ~2% of all US electricity

5.75 M new servers installed every year.

Average lifespan
Server - 3 years.
HDD ~ 4 years.
Tape Drive ~ 8 years.

WW quarterly HDD shipments exceeded 150 M in 2013.
Now are < 100 M.

Over 90% of all Internet traffic is unwanted material!

Mobile devices now account for 28% of all corporate data.

IoT expected to exceed 25B connected devices by 2020, 35% more data.

The human genome takes 600GB to digitize.

Data Growth and Storage Requirements

Are Not the Same After Data Reduction…

2016 Storage Scenario

- Digital Universe was 4.4 ZB in 2013 -> 44 ZB in 2020
- (*doubling every two years*)
- In 2020, 44 ZB equals 5.2 TB per Person Worldwide
- Approx. 42% of all Data is Duplicated (at least once)
- Approx. 33% of Data is Compressed (2x)
- Approx. 5% of Data is Transient (temporary)
- Approx. 75% of all Data is Unstructured -hard to navigate
- Annual global internet traffic is nearing 2 ZB threshold
- Video content is approximately 80% of all global consumer internet traffic.

Note: Amount of Data Created ≠ Amount Data Stored

Data Classification by Business Value

Mission Critical!	15%
Business Critical/Vital	20%
Sensitive	25%
Non-critical	40%

Not All Data is Created Equal!

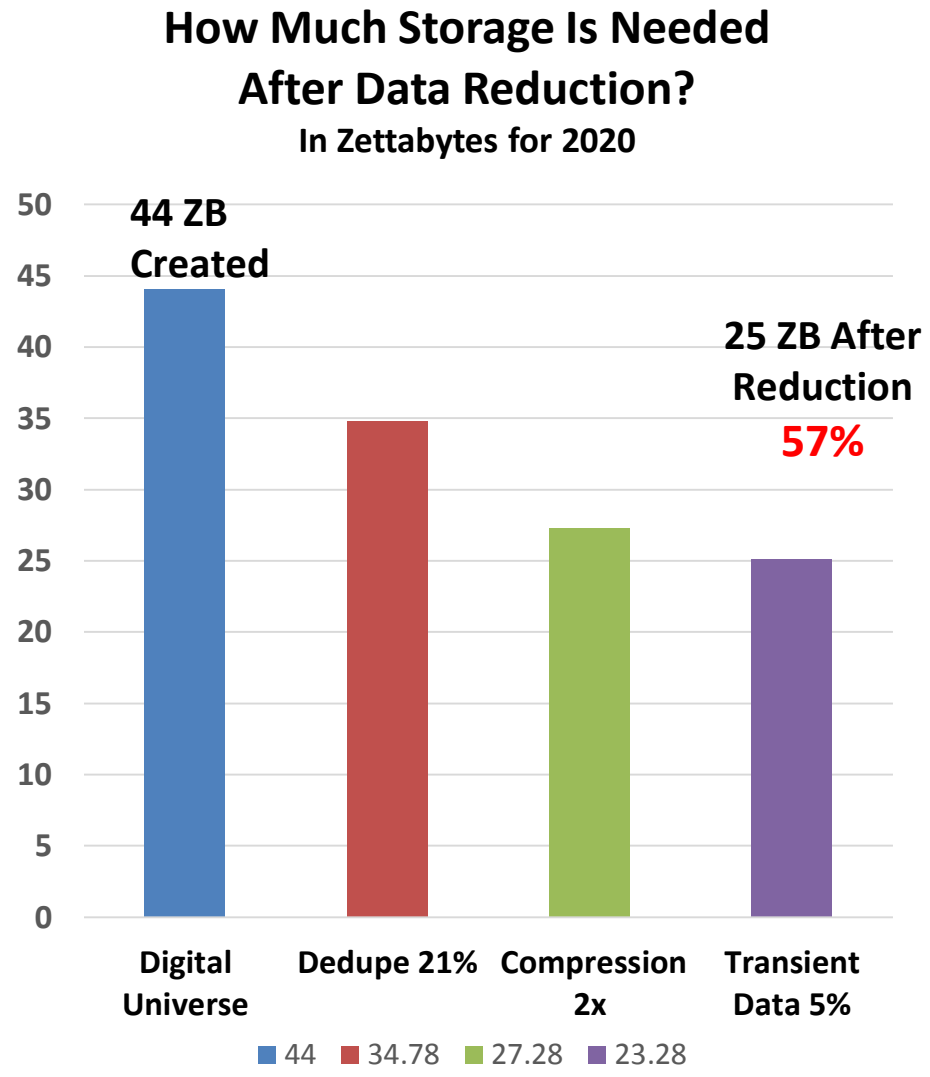

Source: EMC IDC http://www.emc.com/leadership/digital-universe/index.htm, Horison, Inc.

HPC Pushing Exascale Requirements

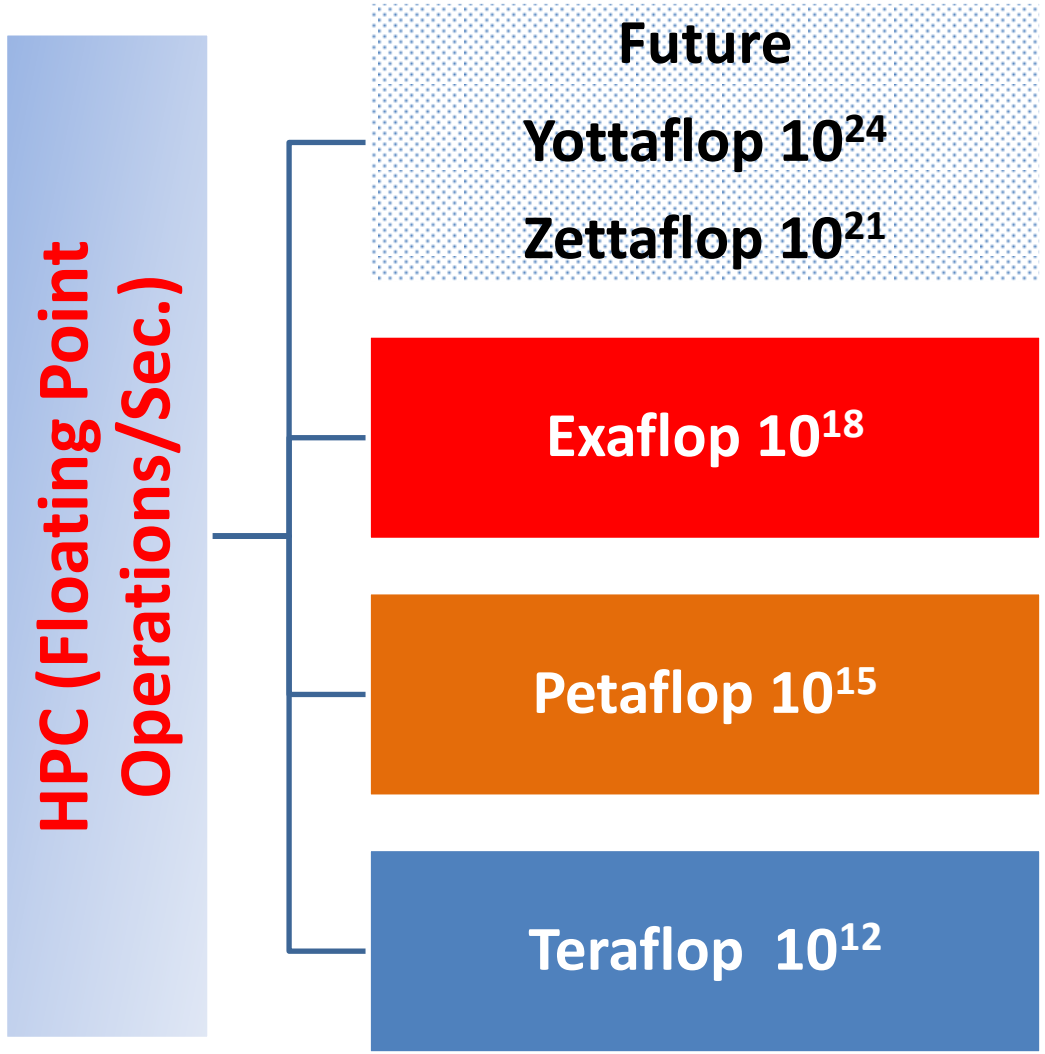

- HPC currently at Petascale compute levels.
- Note: Largest HPC today is 33 Petaflops.
- Exascale parameters
  - Compute speed – > 1 ExaFlop
  - Storage - systems that scale > 1 EB
  - Data transfer - > 100 TB/sec.

- Exascale capacity tape libraries have arrived.
- Faster Solid State Storage needed.
- Data availability requirement has surpassed 99.999%.
- HPC TCO mounting – energy, facilities, security.

- Oak Ridge National Lab's (ORNL) Summit super-computer expected to hit 200 petaflops in 2018.

Exascale requirements will arrive - regardless of the technology capabilities available.

Source: Horison, Inc.

Hyperscale Data Centers Arrive - in a BIG Way

Trend to fewer but *much* larger data centers

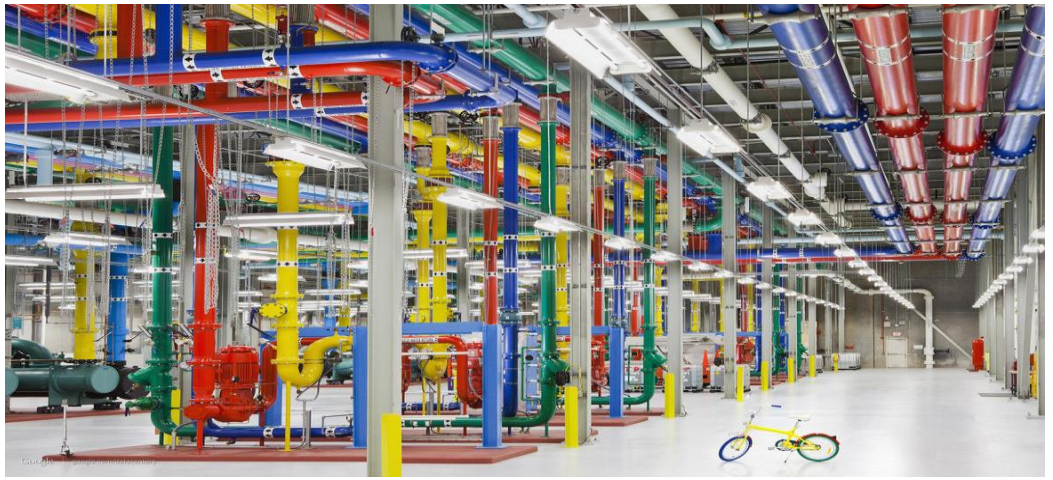

- **Hyperscale is an enormous distributed computing environment – scaling from PBs to EBs.**
- **Exponential increase in volume of data and workload.**
- **Servers and DAS are the basic unit - data is widely spread.**
- **Not built with redundant components – if a failure - workload moves to another server.**
- **Hyperscale storage typically serve millions of users with a few applications.**
- **Hyperscale storage has a minimal feature set and may lack redundancy.**
- **Amazon Web Services, Microsoft, IBM, and Google collectively control more than half of the WW cloud infrastructure service market.**
- **Unprecedented energy consumption.**

Source: Horison, Inc.

IT Energy Consumption to Accelerate

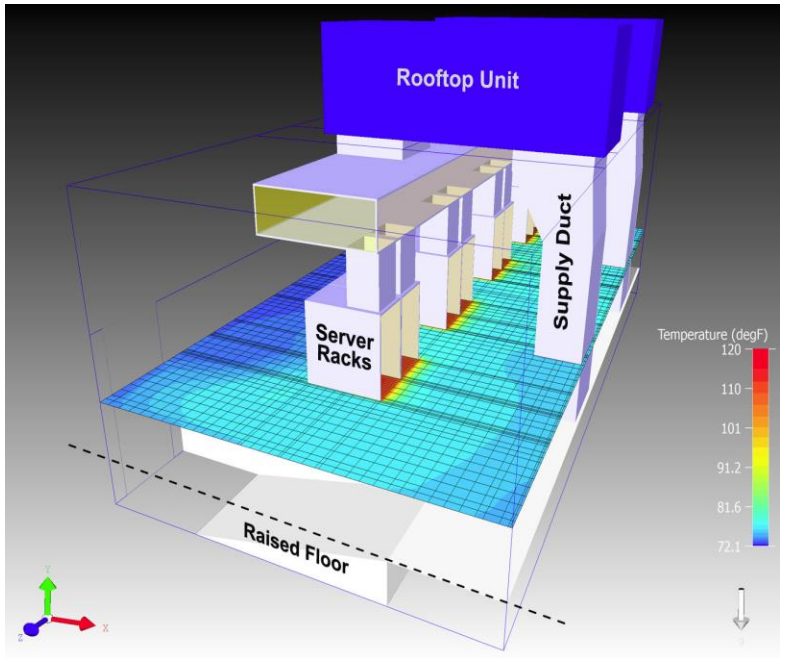

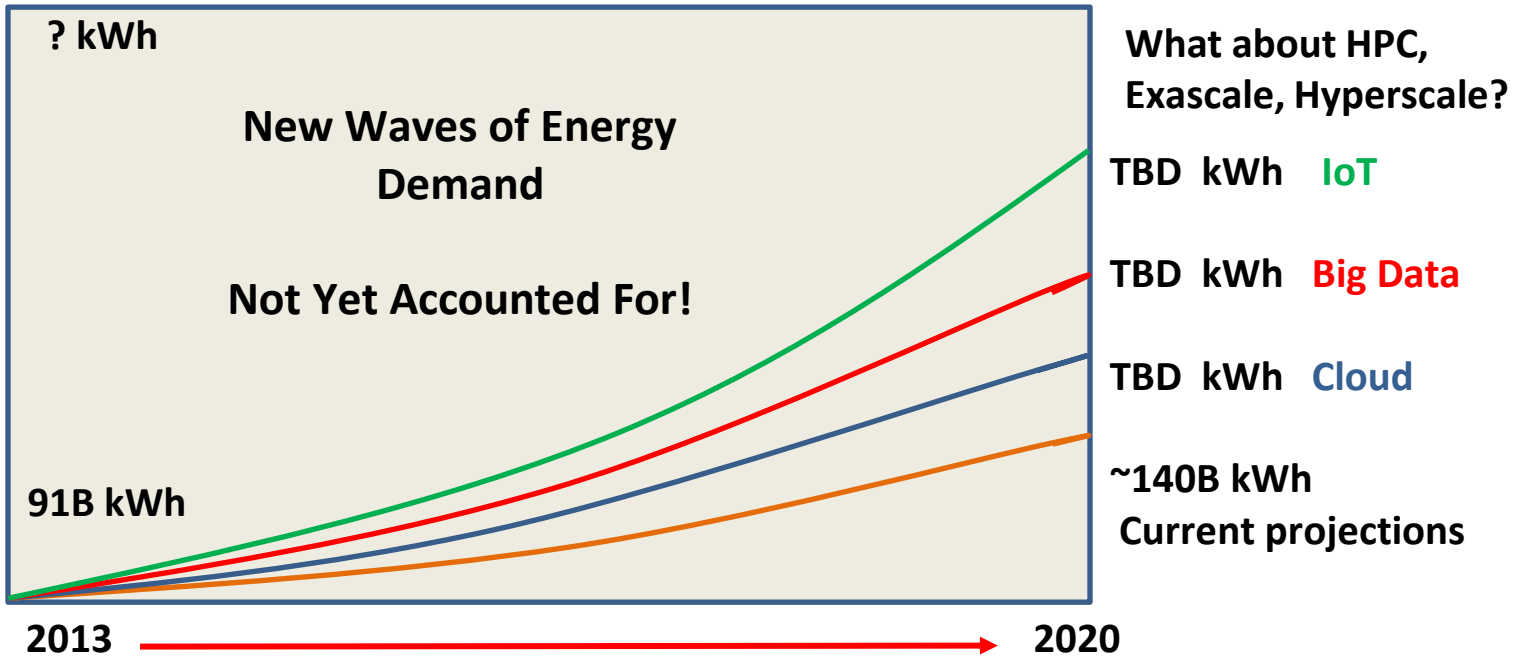

- Several New Waves of Storage Demand Will Increase Data Center Energy Consumption.
- In 2013, U.S. data centers consumed an estimated 91 billion kilowatt-hours of electricity, equivalent to the annual output of 34 large (500-megawatt) coal-fired power plants.
- Data center electricity consumption is projected to increase to roughly 140 billion kWh annually by 2020 - based on current trends
- But not including IoT, Big Data and Cloud demand.

Source: Horison, Inc.

The Tiered Storage Hierarchy

Data Aging Profile	
Age in Days	Probability of Re-use
1	70 - 80%
3	40 - 60%
7	20 - 25%
30	1 - 5%
90+	Near 0

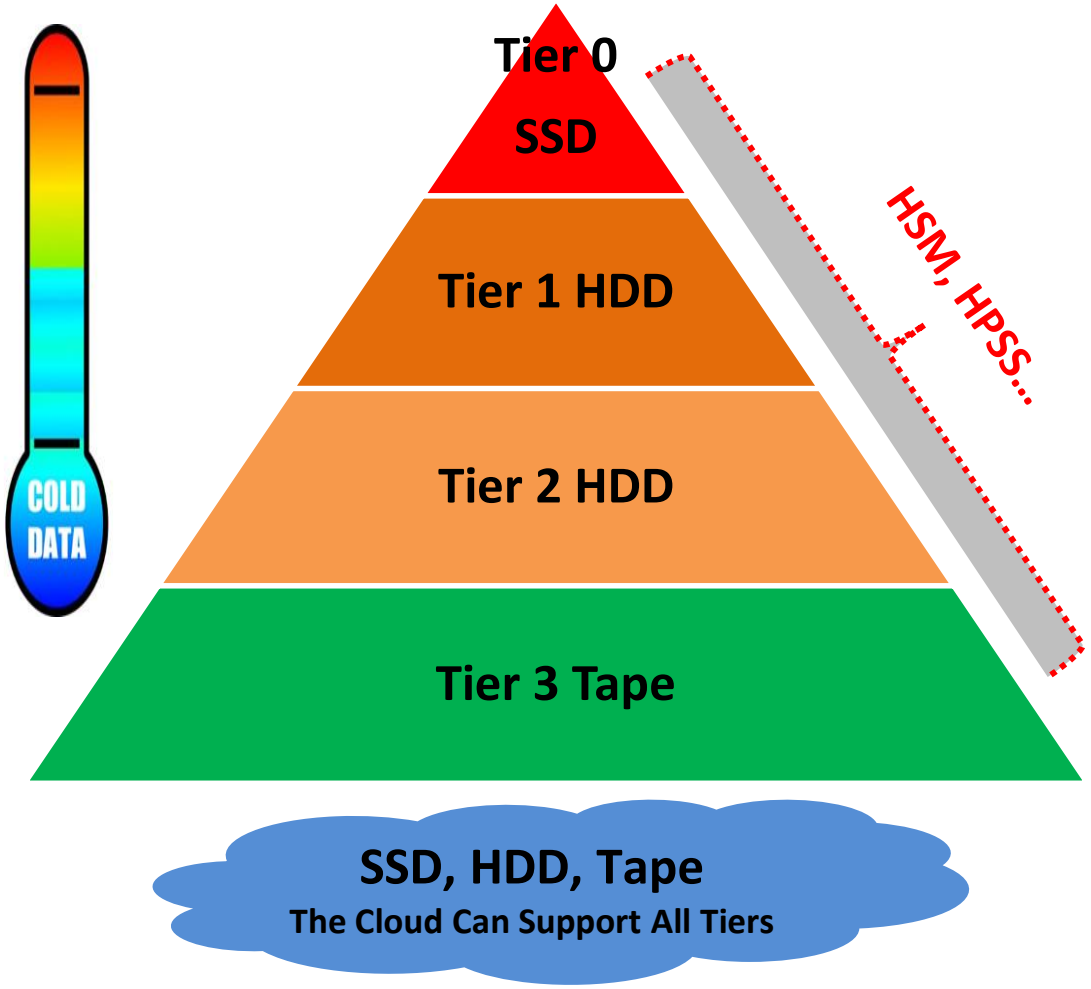

Avg. Data By Tier	Key Applications
Tier 0 5%	Very Hi- Performance Apps, Critical Data, OLTP
Tier 1 15%	Mission-critical, OLTP, Revenue Generating, Hi-performance Apps
Tier 2 20%	Backup/recovery, Vital and Sensitive Data, Moderate Perf., Test and Development, and File Services, Big Data
Tier 3 60%	Archive, Backup, Recovery, Long-term Retention, Big Data, DR, Compliance
Cloud	Archive-as-a-Service (AaaS), DR/BR-as-a-Service (DRaaS), BR-as-a-Service (BRaaS)

Source: Horison, Inc.

Cost Reduction Using Tiered Storage

Optimize Data Placement - Move Archival Data Off of Disk

Example For 1 PB of Storage	Cost/GB ASP	% Alloc. One HDD Tier	Total Cost	% Alloc. Two HDD Tiers	Total Cost	% Alloc. Four Tiers	Total Cost
Tier 0 SSD	$5.00	0		0	0	5% 50 TB	$250K
Tier 1 HDD	$1.50	100% 1,000 TB	$1.5M	50% 500 TB	$750K	15% 150 TB	$225K
Tier 2 HDD	$0.50	0	0	50% 500 TB	$250K	20% 200 TB	$100K
Tier 3 Tape	$0.02	0	0	0	0	60% 600 TB	$12K
Totals			$1.5M		$1.0M	100%	$587K

Allocation Percentages Use Industry Average Data Distribution per Tier
Pricing Uses ASP (Average Selling Price Per GB), Not List Price. Prices will vary greatly.
Using the Tape Tier provides greatest TCO advantage!

Source: Horison Inc.

Solid State Disk (SSD) Scenario

Storage Class Memory - Several SSD and Hybrid Implementations

Hybrid Flash Array

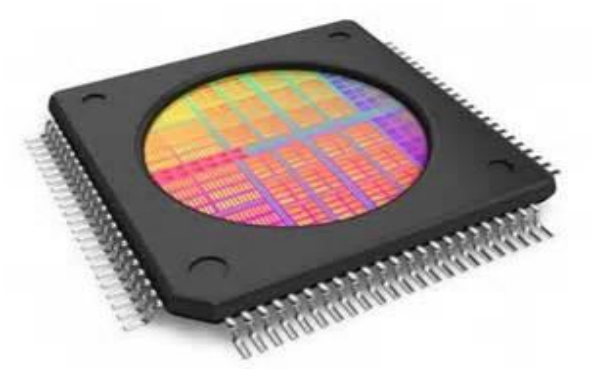

Flash Memory

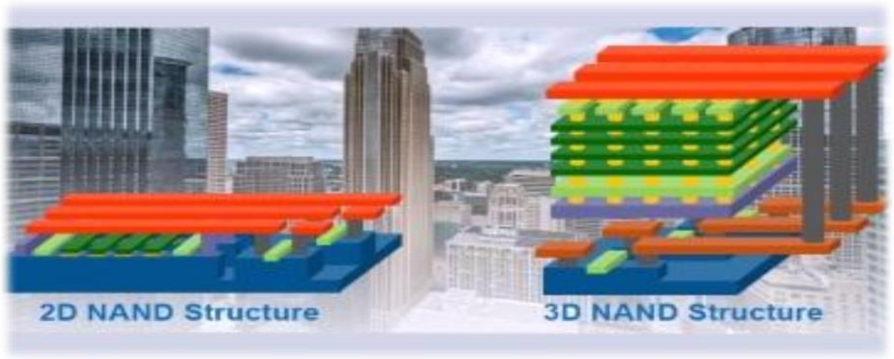

3D NAND Flash *is next…*

- First SSD in 1978 (STK 4305 @ $8800/MB - DRAM)
- 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 16% of WW total HDD capacity shipped in 2015
- 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 Impacting HDD Sales From Enterprise to Desktop
- 3D NAND on the Horizon - Stacks Flash Cells Vertically in 32 Layers to Achieve a 3x capacity Increase.

Source: Horison, Inc.

Disk (HDD) Storage Scenario

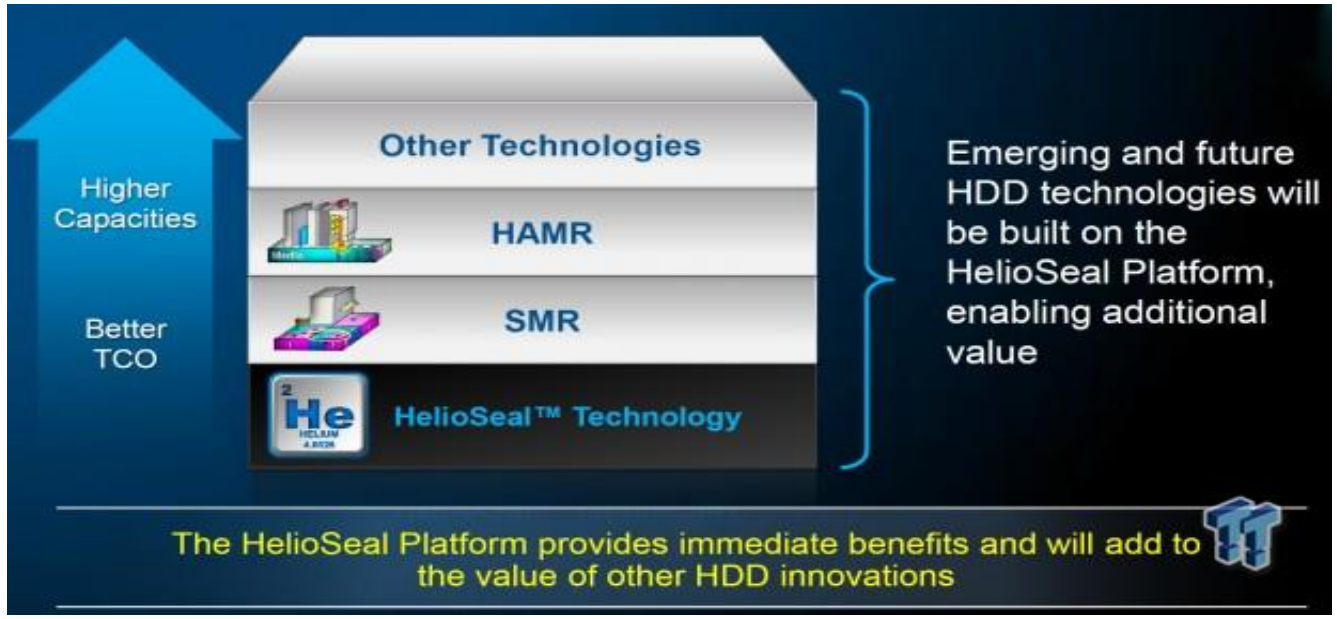

Helioseal is a registered trademark of Western Digital's HGST

Helium-filled HDDs Contain More Platters.

The Platter above has 7 Rather Than the Usual 3.

- **HDD Capacity Growing 20-30% - but Areal Density (~1.2 Tb/sq. in^2) Growth Rate is Slowing**
- **Current Maximum HDD Capacity at 10 TB (SAS)**
- **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 as Flash Usage Accelerates**

Source: Horison, Inc.

New Disk Concepts Are Arriving

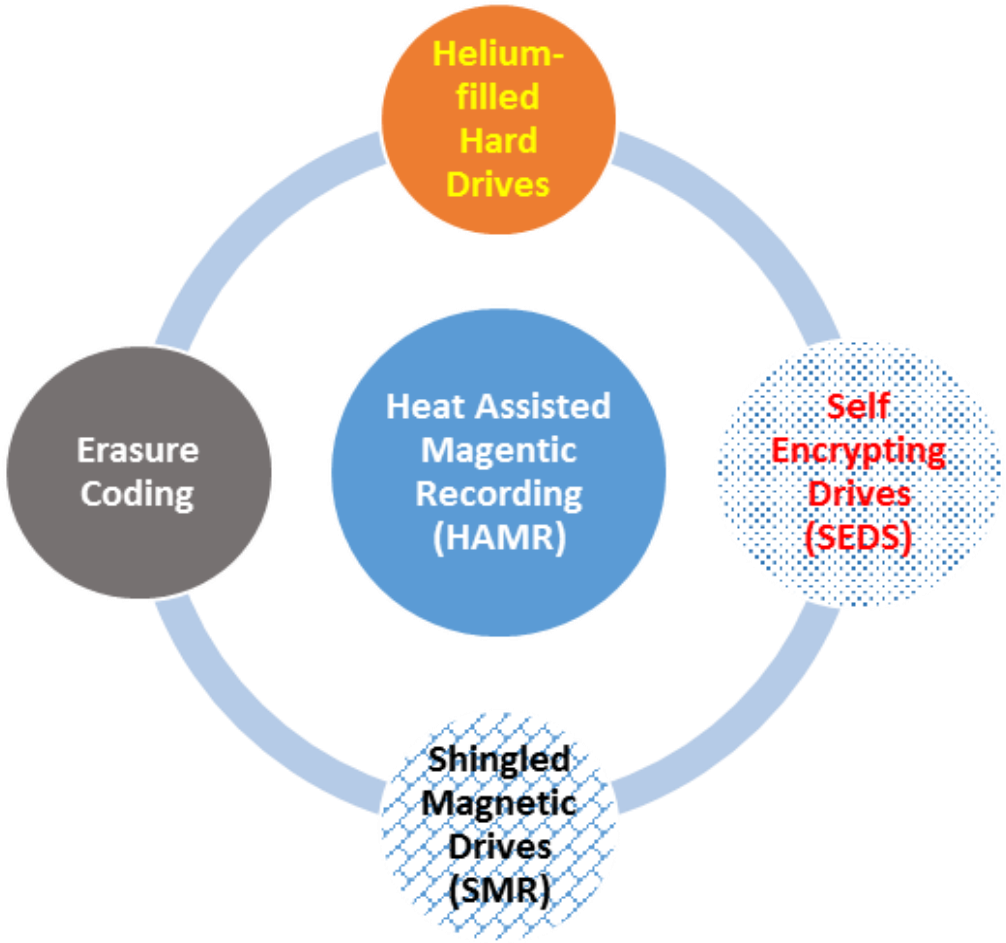

Helium Filled 	**Helium-filled drives use less power to spin disks (which spin easier thanks to less resistance), they run cooler, and can stack more disks in the same space.**
SEDs 	**SEDs use disk drive controllerhardware to encrypt all data written to the magnetic media and decrypts all the data read from the media automatically. Minimal usage.**
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.**
Erasure Coding 	**Data is broken into encoded fragments and stored across different locations, such as HDDs, storage nodes or distant geographical locations.**
HAMR 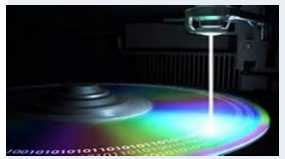	**HAMR uses a small laser to heat the part of the disk surface that is being written to remove the super-paramagnetic effect. Note: presently under development.**

Source: Horison, Inc.

Data Profile

What's On Your Disk?

Disk Utilization Profile

System Overhead, RAID, Control Fields	5%
Orphaned (Inactive or Unknown Files)	5%
Allocated - Used (Live Data)	40%
Allocated - Unused (Gas)	20%
Inert – Free Space (Available for Use)	30%

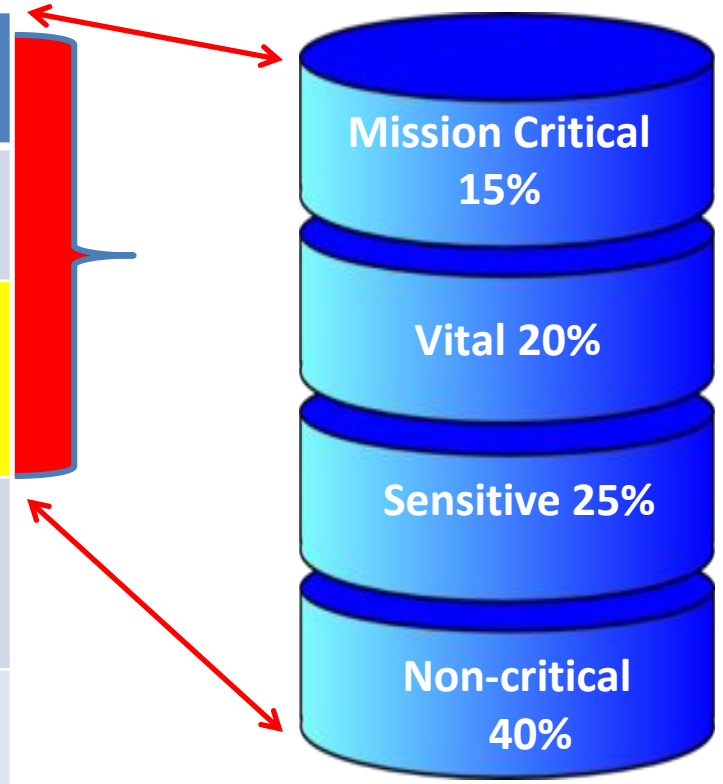

Classification of Data by Value

Sensitive data can take up to several hours for recovery without causing major operational impact.

Non-critical data is not critical for immediate business survival but is often retained for secure archives and Big Data analysis.

Note: Average Disk Allocation Levels for Open Systems

Source: Horison, Inc.

Magnetic Tape Storage Scenario

Barium Ferrite

- Over 85% of Tape Drive Shipments are LTO (>100,000 PB – *100 EB* - of LTO shipped)
- BaFe Driving Media Progress – Lifespan at 30 Years or More
- Tape Has 200x Times More Surface Area to record on Than HDD
- Tape Drive Reliability (BER), Data Rate and Capacity Has ***Surpassed Disk***
- Tape Capacities at 10 TB Native, >25 TB Compressed Areal Density > 6 Gb in^2
- Tape Data Rates at 360 MB/sec. Native Streaming Mode – RAIT is on the Way
- Disk Gaining Backup Applications From Tape via Deduplication (HDD)
- Disk Losing Archive Applications to Tape – Economics, Reliability, Media life
- More Than 60% of All Digital Data Classified as Tier 3 (Archive, Fixed Content) – fastest growing
- Clouds Embracing Tape Solutions for Better TCO and More Secure Archival Services

Source: Horison, Inc.

Magnetic Tape Future Projections

PARAMETER	2017	2019	2021	2023	Change rate
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

RAIT – Has Its Time Arrived?

Redundant Arrays of Independent Tape - Multiply Tape Data Rate

RAIT 5

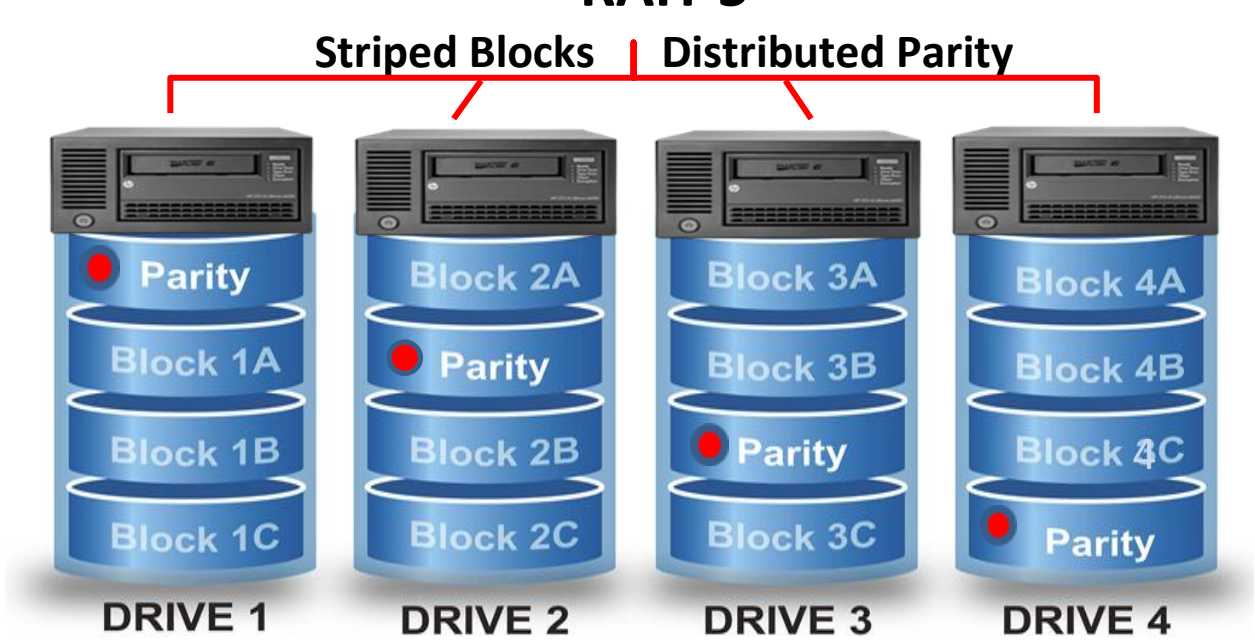

Parity Rotated Across all Cartridges in Stripe

RAIT 0 Striping

Striped Blocks – No Parity

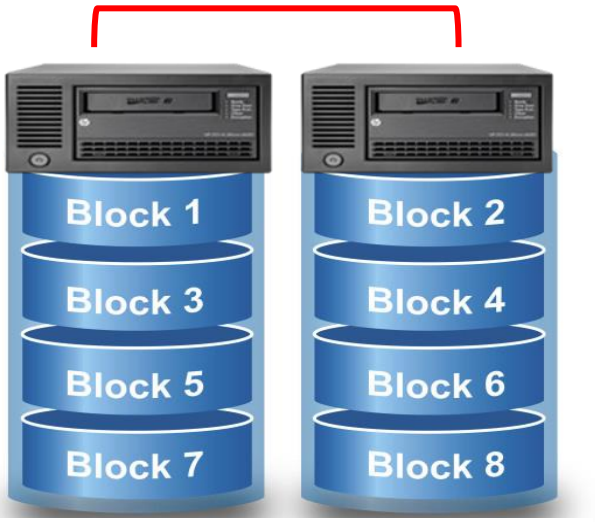

RAIT 5 and RAIT 0 are tape striping architectures
Software segments files (blocks) and distributes them across *n* drives in stripe (*n* = 2, 4, 8, 12, 16..).
Throughput - optimizes higher tape data transfer – RAIT 5 Ex: effective transfer rate increases up to 4x.
Availability - fault tolerance provides much 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.

A Glimpse Into the Future

Key Tape Patents Granted 2015 - present

Oracle	(9,311,044)	• Supporting Efficient Buffer Usage With Single External Memory Interface.
	(9,275,666)	• Rotatable Recording Head Actuator For Correcting Angular Error In Tape Drives.
	(9,244,961)	• Concurrent Access For Hierarchical Data Storage.
	(9,195,402)	• Target And Initiator Mode Configuration Of Tape Drives For Data Transfer Between Source And Destination Tape Drives.
IBM	(9,021,196)	• Writing Multiple Files Simultaneously To Tape Media.
	(9,275,678)	• Primary Storage Media With Associated Secondary Storage Media For Efficient Data Management.
	(9,060,414)	• Solid State (Tape) Storage Media.
	(9,285,996)	• Tape Drive Buffer Utilization For Improving Tape Drive Efficiency Using A Processor.
	(9,047,879)	• High Performance Cartridge Format.
	(9,019,654)	• Data Storage Tape With Random Access Data.
	(9,021,175)	• Re-ordering Access To Reduce Total Seek Time On Tape Media.
Spectra	(9,336,815)	• Dynamic Write Once Read Many (WORM) Tape Cartridge Designation.
Fujifilm	(9,401,171)	• Granted a total of 10 patents related to magnetic particles and method of manufacturing for Fujifilm recording media.

Off-Line Storage

Optical Disc Scenario - 2016

Blu-ray Optical Disc (BD) has set the standard for high definition picture and audio quality - popular in the home and car – but not in the data center.

Market Shares – CD 35%, DVD 51.5%, BD 13.5%.

Specifications (Blu-ray)
- **Single layer Blu-ray discs up to 25GB/layer.**
- **Dual layer discs up to 50GB/layer.**

- **New ULTRA HD Blu-ray discs to hold 66GB and 100GBs.**
- **300GB to 1TB multi layer discs planned – TBD.**
- **One LTO-7 native cartridge can contain 120 Blu-ray discs.**
- **Read data rate only 17.1 MB/sec.**
- **Reliability BER $1x10^{12}$ – far below all magnetic devices.**

WORM format is more reliable than RW.

Bottom line: Optical disc has fallen far behind magnetic storage in capacity, performance and reliability.
Not cost-effective for data center usage.

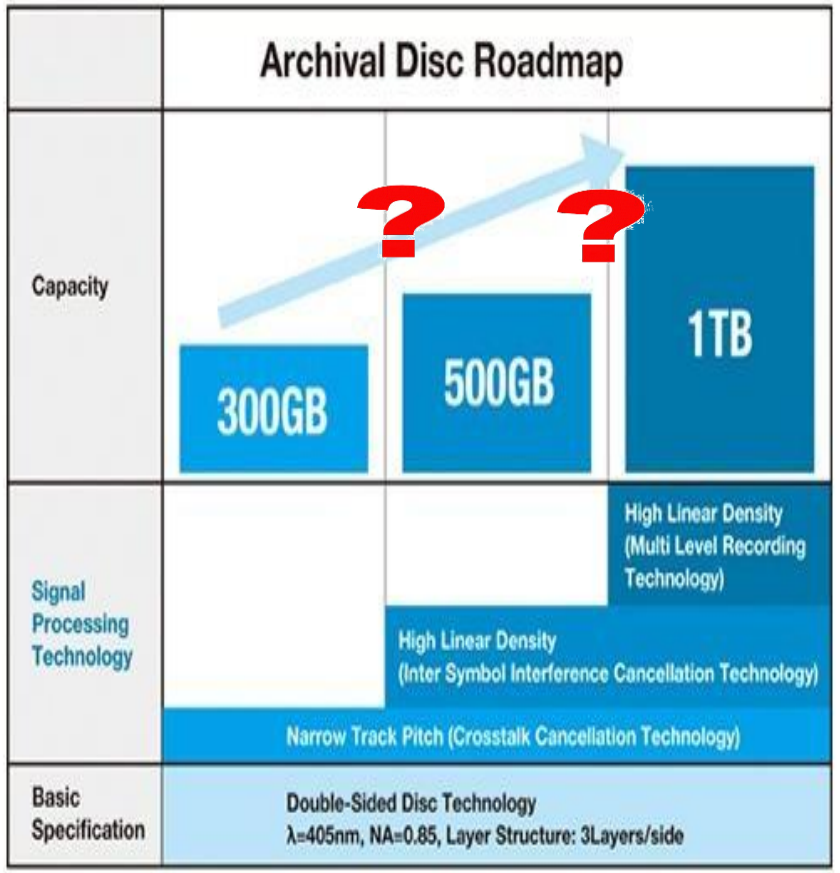

Storage Reliability Levels

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, and 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	> 30 years

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 – Number of Bits in Error per Total Bits Read

Tape is Now the Most Reliable Storage Technology Available

Source: Horison, Inc.

Backup and Archive are Not the Same!

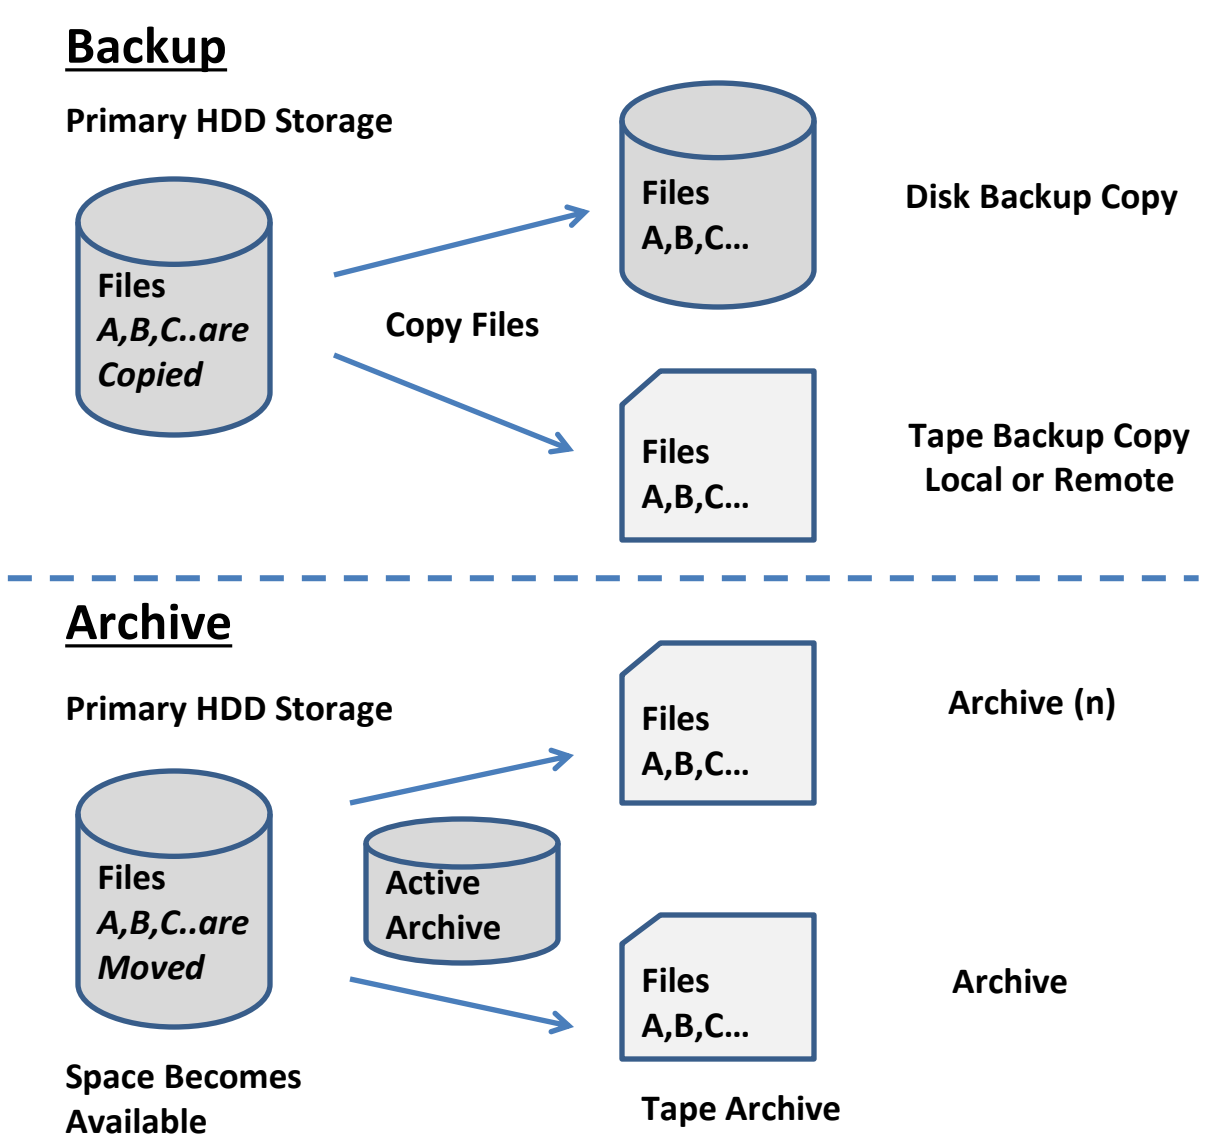

Backup: (Copy data)

Back up - or the process of backing up - is making copies of data which may be used to *restore or recover* the original after a data loss event. Short-term storage.

Archive: (Move data)

An archive moves data to a new location and refers to data specifically selected for long-term retention.

Archives are usually data that is not actively used and was moved from its initial location to a more cost-effective long-term storage location.

Note: Archive data should have a second (backup) copy.

Source: Horison, Inc.

Archive Retention Requirements

Signals Need for Advanced Long-term Archival Solutions

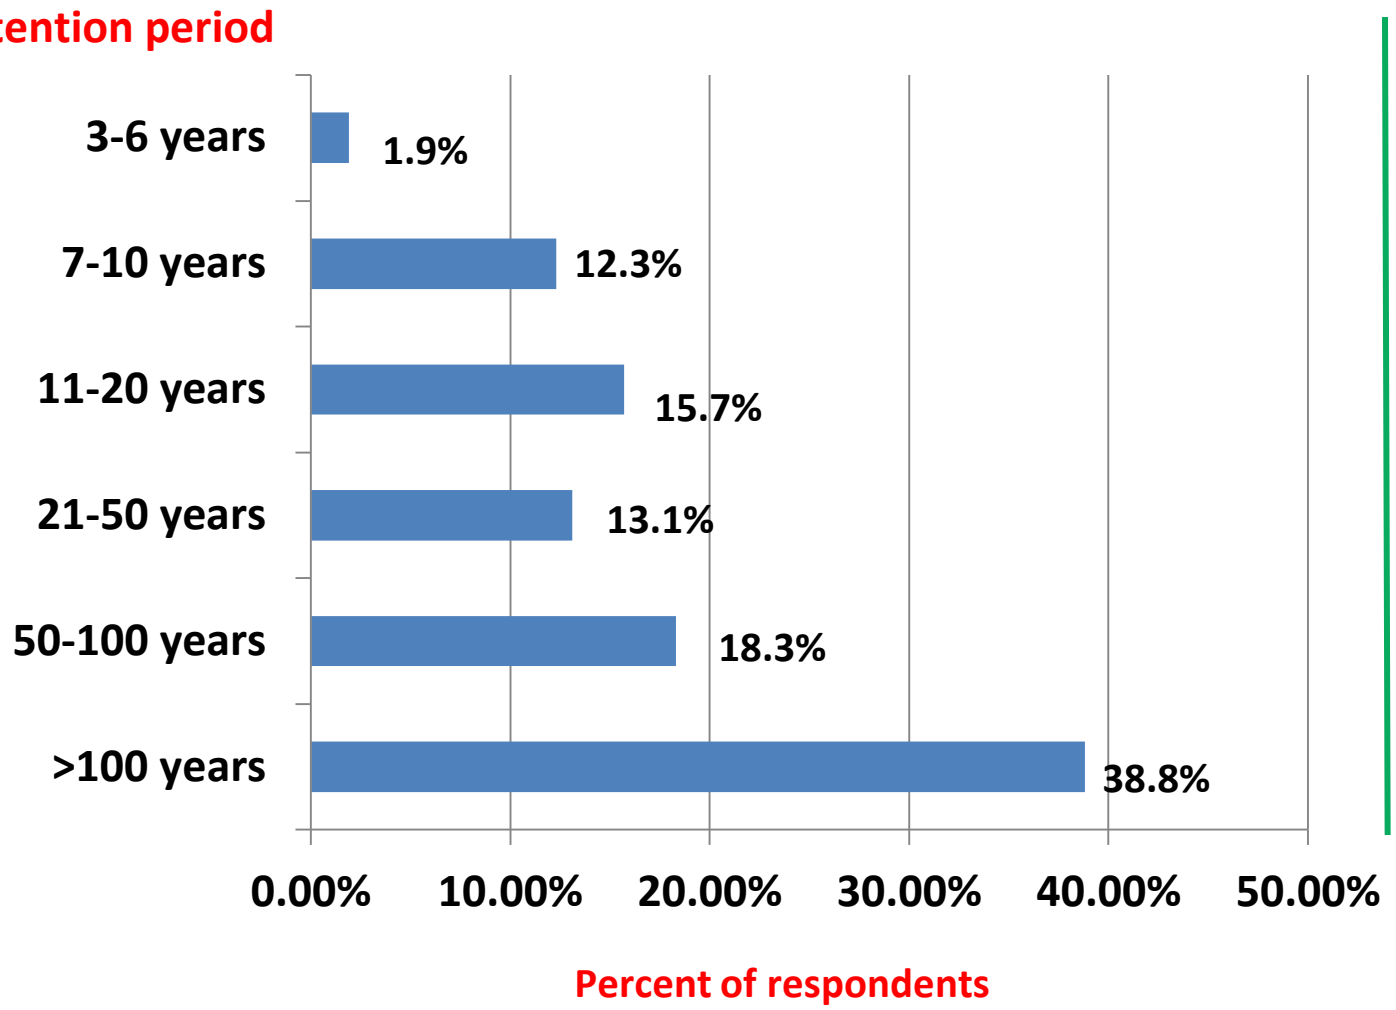

- Long-term Archival Storage Requirements Increasing
- 20+ Years Retention Required by 70% of respondents
- 50+ Years Retention Required by 57% of Respondents
- Key Driving factors
  - Compliance
  - Legal risk
  - Business risk
  - Security risk

Source: 2013 Storage Developer Conference SNIA LTR TWG

Optimizing Archive Storage

Compare Archival Capabilities – Onsite and Cloud

Archive Function	Tape	Disk
Technology Refresh Cycle for Archive	✓ Yes, 30 years or more on all new media and ~ years for drives.	~4-5 years for most HDDs before upgrade or replacement.
Reliability	✓ Tape BER has surpassed disk.	BER not improving as fast as tape.
Portability	✓ Yes, media completely removable and easily transported.	Disks are difficult to remove and to safely transport.
Inactive Data Does Not Consume Energy - Green	✓ Yes, this is becoming a goal for most data centers. "If the data isn't being used, it shouldn't consume energy".	Rarely for disk, except in the case of "spin-up, spin-down" disks. HDDs consume much more energy than tape.
Security	✓ Encryption and WORM available on all midrange and enterprise tape drives.	✓SED available on selected disk products, PCs and personal appliances. Not yet widely used.
Scalability	✓ Tape scales in both capacity and throughput/performance.	Disk only scaling in capacity - minimal performance gains.
Total Cost of Ownership (TCO)	✓ Heavily favors tape for backup (2-4:1) and archive (15:1).	Higher TCO, more frequent conversions and upgrades. High energy costs.
Storage Admin. Capability	✓ Can manage PBs of data ($1x10^{18}$)	Can manage TBs of data ($1x10^{15}$)
Cloud Archiving	✓ Cloud data can be physically ingested using tape where large data quantities prohibit the use of network transmission. ✓ Tape makes cloud storage services more cost effective and secure. Ideal archive solution.	Not cost effective for archiving. The more Cloud data you move - the more you pay! Most enterprise organizations' "cloud" storage infrastructures are actually "hybrid" clouds.

Source: Horison Inc.

Cloud Performance Model

Scenario – Cloud times transferring a 20 Gigabyte file over a Public Internet connection with bandwidth of 1, 10, and 40 Gigabit/sec. (Gbps) Ethernet - xGigE.

Public Cloud

Many organizations, resource pooling

The biggest Cloud challenge is limited Internet bandwidth

- **20 GB file size packetized for TCP/IP = 224 billion bits (+ 40% framing).**
- **Add 10 - 15% to size *if* encryption is used.**
- **Typical storage price = $.0275 - $.03/GB/mon.**
- **Typical XFR prices $.02 - $.09/GB.**
- **Or $20 - $90/TB transferred.**
- **Download speeds are typically faster than upload speeds, and reads are faster than writes.**
- **The more data you move – the more you pay!**
- **Data rate of 1 Gig E = 1 billion bits/sec.**
- **224 billion bits @ 1 billion bits/sec (1 Gig E) = 224 seconds.**

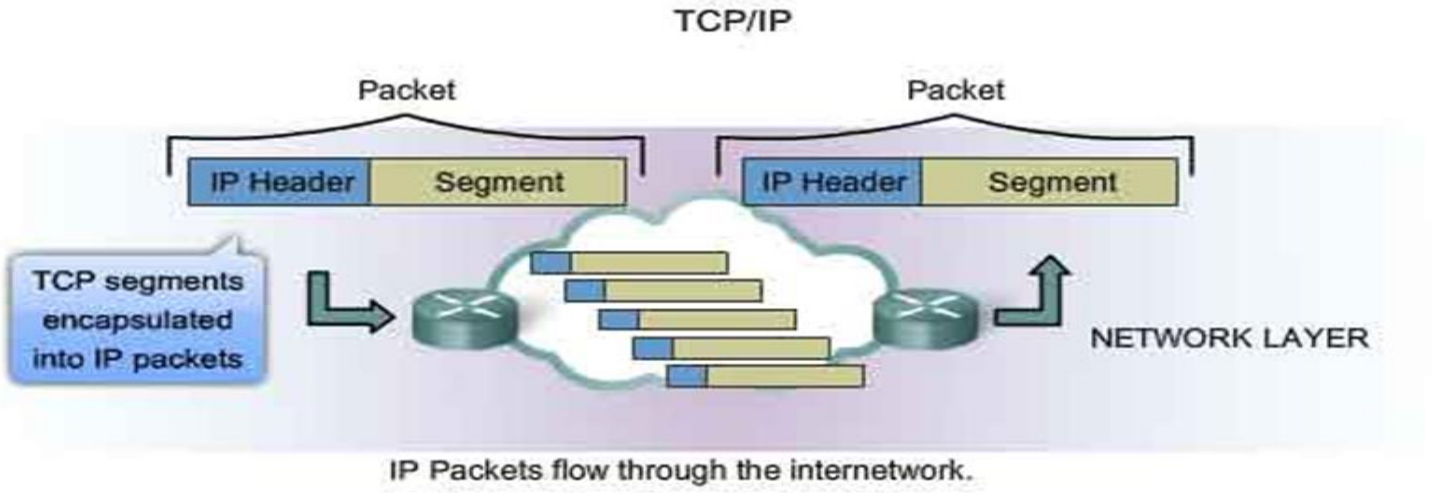

IP Packets flow through the internetwork.

- Connectionless - No connection is established before sending data packets.
- Best Effort (unreliable) - No overhead is used to guarantee packet delivery.
- Media Independent - Operates independently of the medium carrying the data.

Tape Becomes a Key Cloud Component

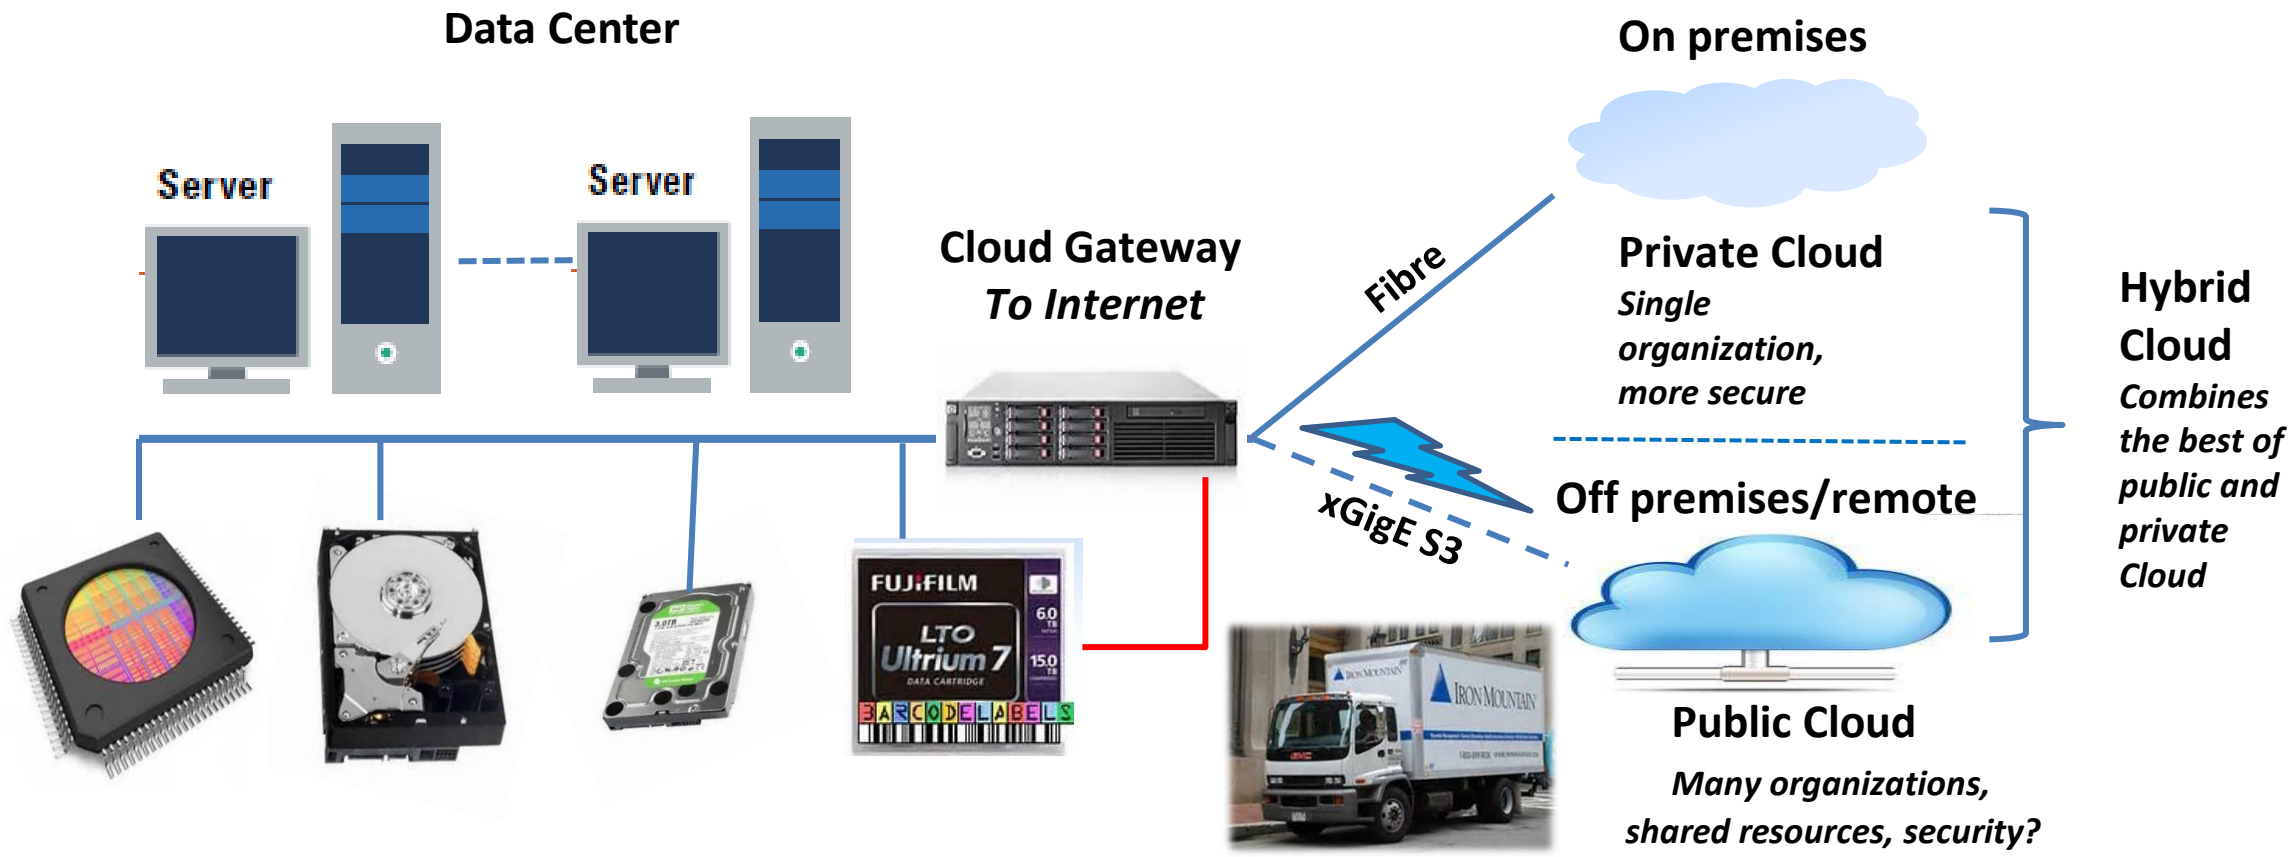

Tier 0	Tier 1	Tier 2	Tier 3	--- To Internet ---	Cloud Storage
Hi-activity			*Moderate activity*	*Lower activity/long-term archives/cold data*	

Source: Horison, Inc.

Storage Squeeze Play – HDD Caught in Middle

The Storage Landscape is Shifting

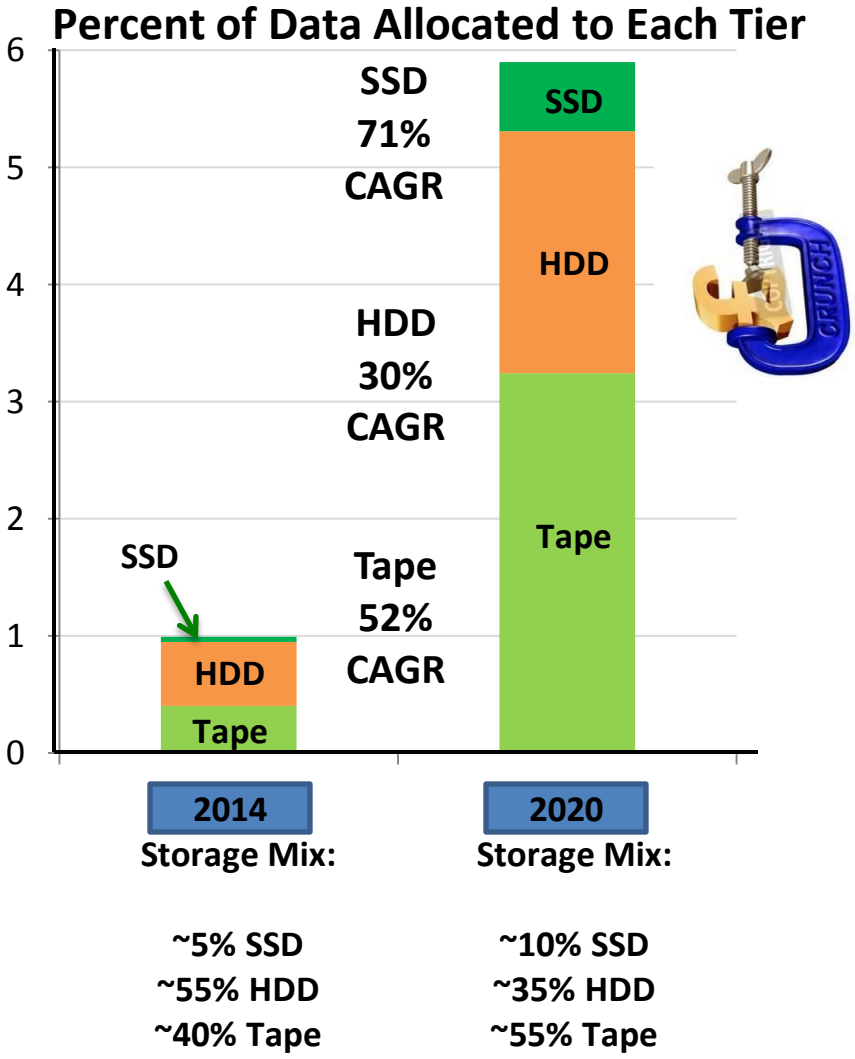

HDD Challenges are Mounting

- Further Disk Performance Gains Minimal
- Disk Re-build Times Excessive (*n* * days)
- Disk Capacity Gains Facing Limits
- Disk Adding Platters to Increase Capacity
- Disk TCO Higher Than Tape (4-15x)
- Poor Utilization Requires More HDDs
- Disk Data at Rest is Main Target for Hackers
- SEDs Usage Remains Low
- Tape Reliability Has Surpassed Disk
- Tape is Much Greener Than Disk
- Tape Media Life Now 30 Years or More

Old way: Keep adding more disk
New way: Optimize using SSD, disk and tape

Source: Horison, Inc.

Remember …
Things Are Changing So Fast…

Even the Future is Obsolete