



Storage Industry Update 2015

Fred Moore President

www.horison.com



HORISON
Information Strategies

Did You Know?

There are over 3 billion Internet users – WW pop. 7.3B

Over 90% of all Internet traffic is unwanted material

Mobile devices and laptops now account for 28% of all corporate data

Smartphone shipments reach 338 million globally for 2Q 2015

Est. ~2.13 billion social media users by 2016

2016 model automobiles will average 14 million lines of code – autonomous?

Over 75% of all data is now generated by individuals...

However 80% will become the liability of a large data center

Less than 5% is ever analyzed or touched again – dark archive

Very little data is actually encrypted! (Target, Home Depot, Wiki Leaks, Anthem...)

SEDs (Self-encrypting Drives – HDD and Flash) are here – but seldom used... ???

Most data reaches archival status in less than 30 days

One-fifth of corporations store more than 1 petabyte of data

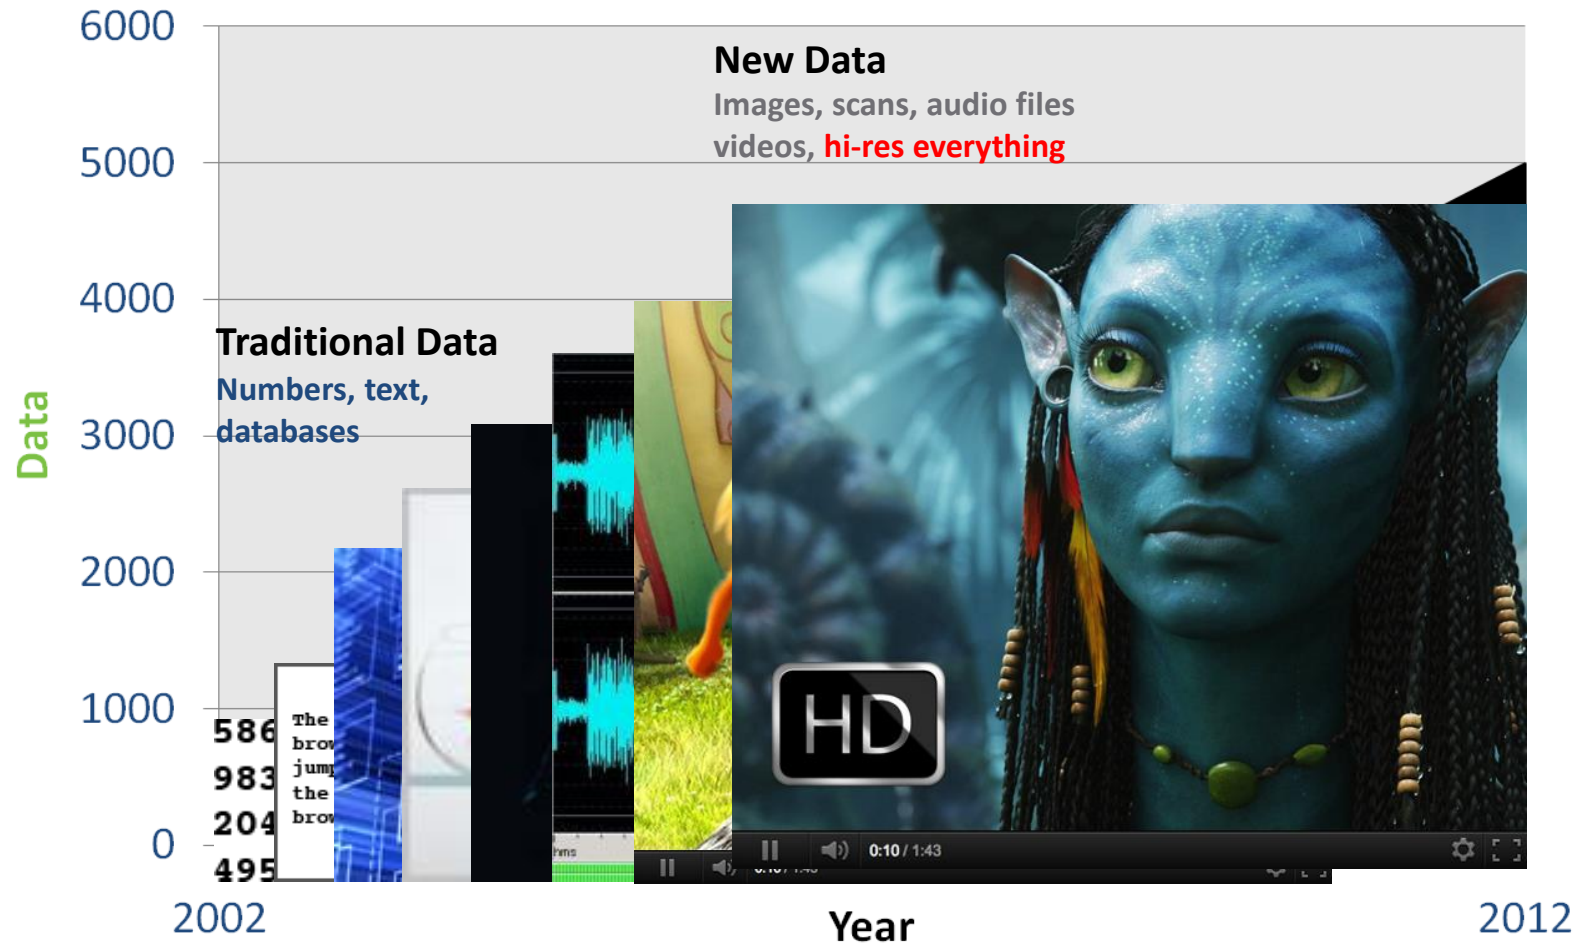
The average open systems disk utilization is <50% - *But* >80% for mainframes

Modern tape cartridges have over 2x the capacity as disk drives

Modern tape is much less expensive and more reliable than disk


Storage and I/O speed steadily falling behind computing speed

The Era of Colossal Content Arrives



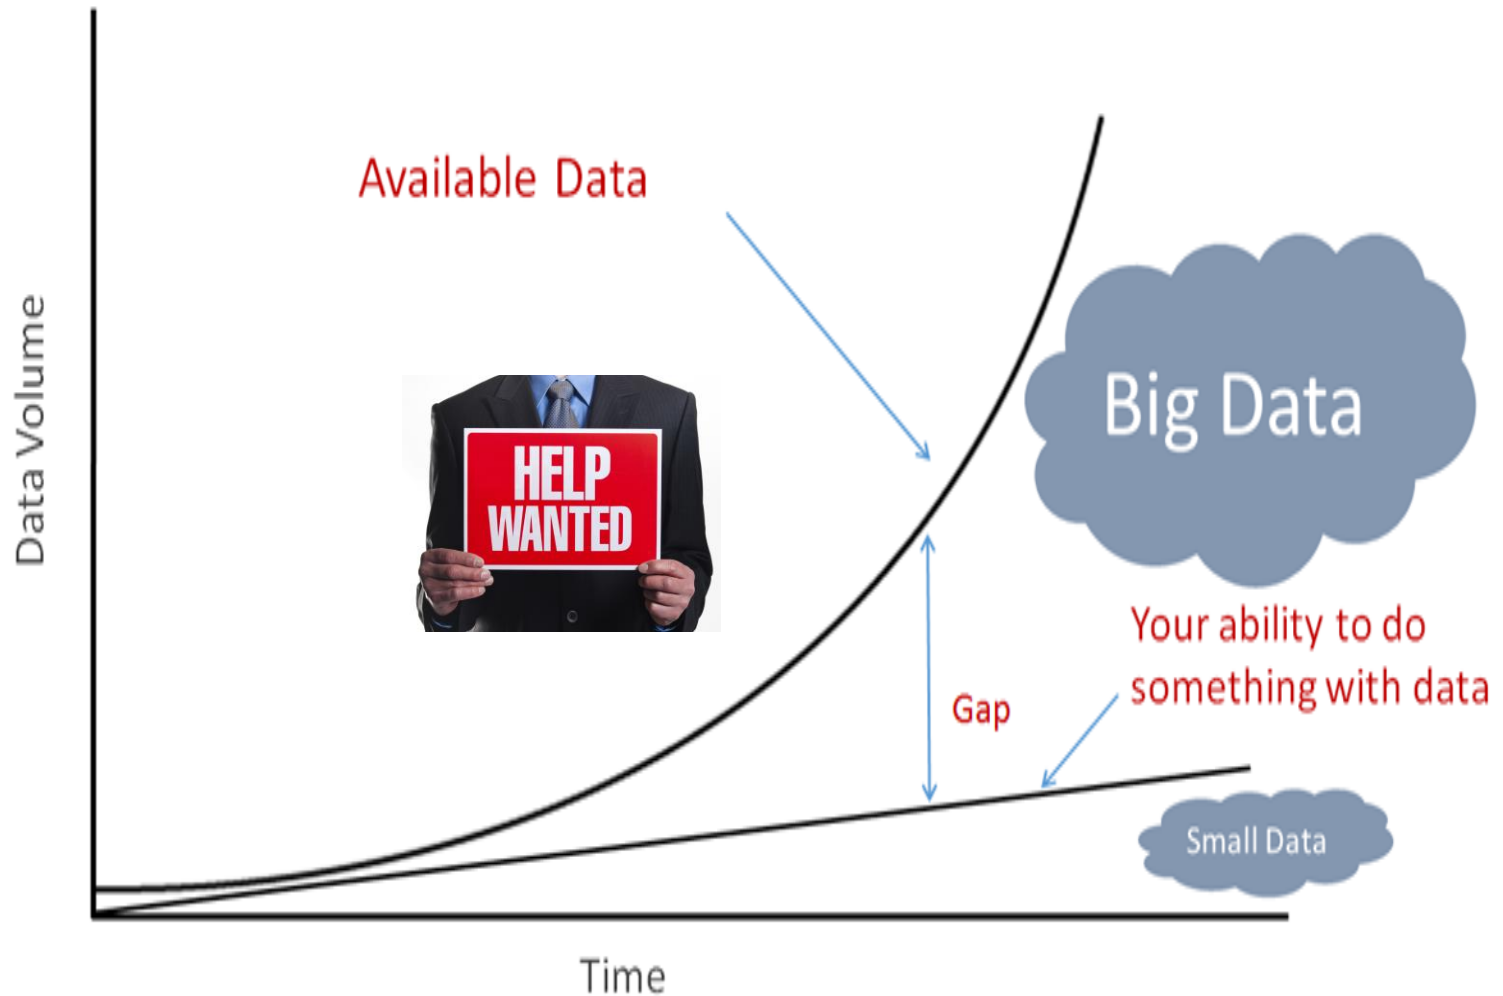
- Cagr. will approach 100% requiring Exascale storage systems with storage tiering
- New drivers are social, mobile, cloud, and unstructured archives with unlimited analytics
- Completely new data management methods required - (Hadoop, HPSS, GPFS, Dispersal, RAIT....)
- Will change the rules of data protection, HA, and DR

The Digital Age....




Digital Data Value...

Data Explosion



The Tiered Storage Hierarchy

Optimal Storage Infrastructure



Availability Index % Primary Technology Used

Tier 0 – 99.999%

Flash & DRAM SSDs

Tier 1 – 99.999%

Enterprise Disk Arrays FC, SAS, RAID, Mirrors, Replication

Tier 2 – 99.99%

Midrange Disk Arrays SATA, VTLs, Dedup, VSM

Tier 3 – 99.999%

Automated Tape Libraries, Remote Tape Vaults – **Energy Efficient Storage**
HSM/HPSS Systems Are Critical to Effectively Manage Storage

Tier 0

Solid State Disk Scenario - 2015

Storage Class Memory - Several SSD and Hybrid Implementations



HDD/SSD 2.5" HDD



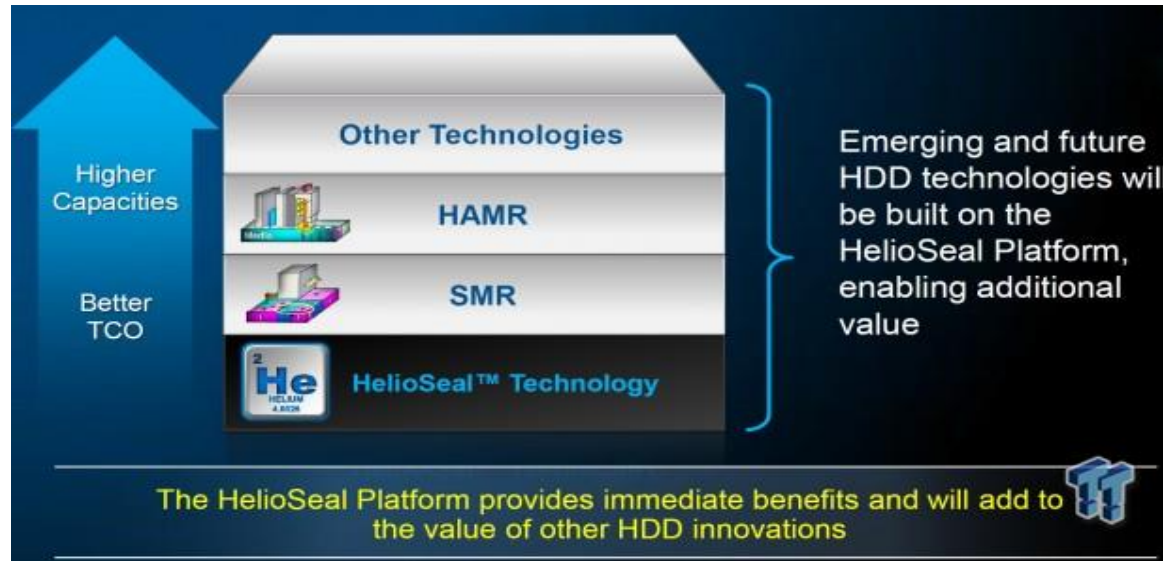
2.5" SSD 4 TB

- First SSD in 1978 (STK 4305 @ \$8800/MB - DRAM)
- Currently 128 Flash Suppliers, 10 Companies Get >90% of ~\$11.3B Revenue in 2015
- All-flash Arrays (AFAs) and Hybrid Flash Arrays (HFAs) - Showing Explosive Growth
- SSD Capacity Now @ 4TB (SAS)
- IOPS Intensive Applications, Databases, OLTP, HPC Burst Buffer
- Non-volatile, Portable, Low Power (1/3 of HDD)
- No Moving Parts, High Reliability - BER 1×10^{17}
- Read Access Times: .2 ms Approx. ~50x Faster Access Than HDD
- Flash Success Impacting HDD Sales at All Levels
- Tier 0 is All About Performance!

Tier 1

Tier 2

Disk Storage Scenario - 2015



Helium-filled HDDs Contain **More Platters.**

The One Shown Here has **7** Rather Than the Usual **3**


- Only 3 HDD Suppliers Remaining - Seagate (49%), WD (42%), Toshiba (9%)
- ~\$32B HDD Revenue From 564M HDDs shipped in 2014
- HDD Capacity Growing 20-30% - but Areal Density Growth Slowing
- Current Maximum HDD Capacity at 8 TB (SAS)
- Low HDD Utilization (<50%) Increases End-user Costs
- Future HDD Performance Gains Are Minimal
- Longer RAID Rebuild Times - Dispersed Storage, Erasure Coding, to Replace RAID?
- Reliability (BER) Has Fallen Behind Tape and Flash Shingled Magnetic Recording (SMR) Helps Address Physical Limits (track overlap)
- TCO Ranges from 2-15x Higher Than Tape
- Remember - HDDs Can Address **All** Data Types and Requirements

HDD Challenge - Capacity Performance Conflict

The Measure of HDD Performance Capability

Access Density =
I/O's per second
per gigabyte

IOPs @10ms	HDD Cap. GB	Access Density
100	1,000	.1
100	2,000	.05
100	4,000	.025
100	8,000	.0125



- HDD Performance Not Scaling With Capacity Growth or Server Speed
- Future HDD Performance Gains are Minimal
- SMR Adds Up to 25% Capacity but Writes Can Be Slower
- Access Density Degrades Response Time (More Actuator Arm Contention)
- Access Density Will Continue to Decline as HDD Capacity Increases
- RAID Rebuild Times Increase (n days) as HDD Capacity Increases
- Creating More Demand For “High - Performance and Hybrid Disk” Category – SSD (Tier 0)

Optical Disc Scenario - 2015

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

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
- Read data rate only 17.1 MB/sec

Reliability BER 1×10^{13} – 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.



Archival Disc Roadmap			
Capacity	300GB	500GB	1TB
Signal Processing Technology		High Linear Density (Inter Symbol Interference Cancellation Technology)	High Linear Density (Multi Level Recording Technology)
Basic Specification	Narrow Track Pitch (Crosstalk Cancellation Technology)		
	Double-Sided Disc Technology $\lambda=405\text{nm}$, NA=0.85, Layer Structure: 3Layers/side		

Tier 3

Tape Storage Scenario – 2015

LTFS

Linear Tape File System

Linear Tape File System (LTFS) is a self-describing magnetic tape format that defines the organization of data on the tape. Supporting LTO-5 and later Ultrium generations, data tapes written in the LTFS allow for direct access to file content data and file metadata. Accessing LTFS formatted tapes is similar to accessing files stored on other storage media like disk or removable flash drives. This format presents a standard file system view of the data stored on the tape media, allowing for users to access the data directly from tape.



- Over 85% of Tape Drive Shipments are LTO (>100,000 PB – 100EB - of LTO shipped)
- Tape Drive Reliability (BER), Data Rate and Capacity Has Surpassed Disk
- Tape Capacities at 10 TB Native, >25 TB Compressed - **Tape Has Highest Capacity**
- Tape Data Rates at 360 MB/sec. Native
- LTFS Provides a Universal, Open File System for Tape
- Disk **Gaining** Backup Applications From Tape via Deduplication (HDD)
- Disk **Losing** Archive Applications to Tape – Economics, Reliability, Media life
- More Than 70% of All Digital Data Classified as Tier 3 (Archive, Fixed Content)
- Over 500 Exabytes Stored on Tape
- Tape Vendors Slow to Expand Market Awareness on Technology Advancements

Unstructured Data Fueling Tape Demand

To Approach 90% of all Data With 40-60% CAGR

WHAT ARE WE STORING?

DATA HITS US IN TWO FORMS

Structured Data



Databases: Microsoft SQL Server, IBM DB2, Oracle Database 12c



XML data: coded text based files found on the Web, Internet and intranets



Data warehouses: Microsoft Analytics Platform System, Oracle Exadata Database Machine, Teradata Aster Discovery Platform



Enterprise systems: Salesforce CRM, Microsoft Dynamics, Oracle's PeopleSoft Enterprise

Unstructured Data

Much Unstructured Data Reaches Archival Status In ~30 days

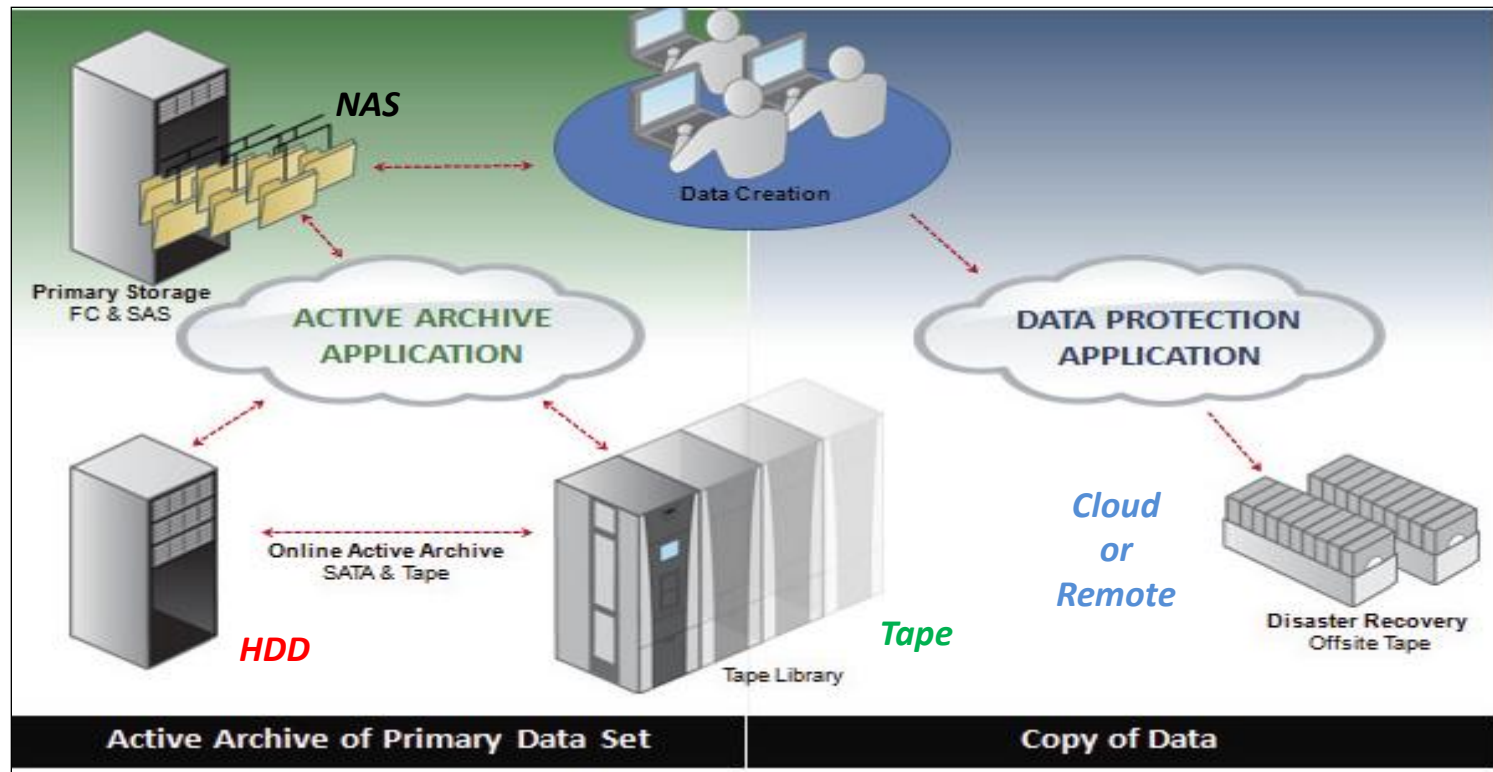
Documents (Word)
Presentations (PowerPoint)
Spreadsheets (Excel)
Email messages
Chat logs
Images
Video
Audio

This will be 90 percent of all data created over the next decade.



The Active Archive

Redefining the Archive Experience



The Active Archive solution combines the simplicity *and performance* of disk with the economics of tape in a highly scalable solution.

Active Archives takes advantage of open LTFS tape standards using modern tape.

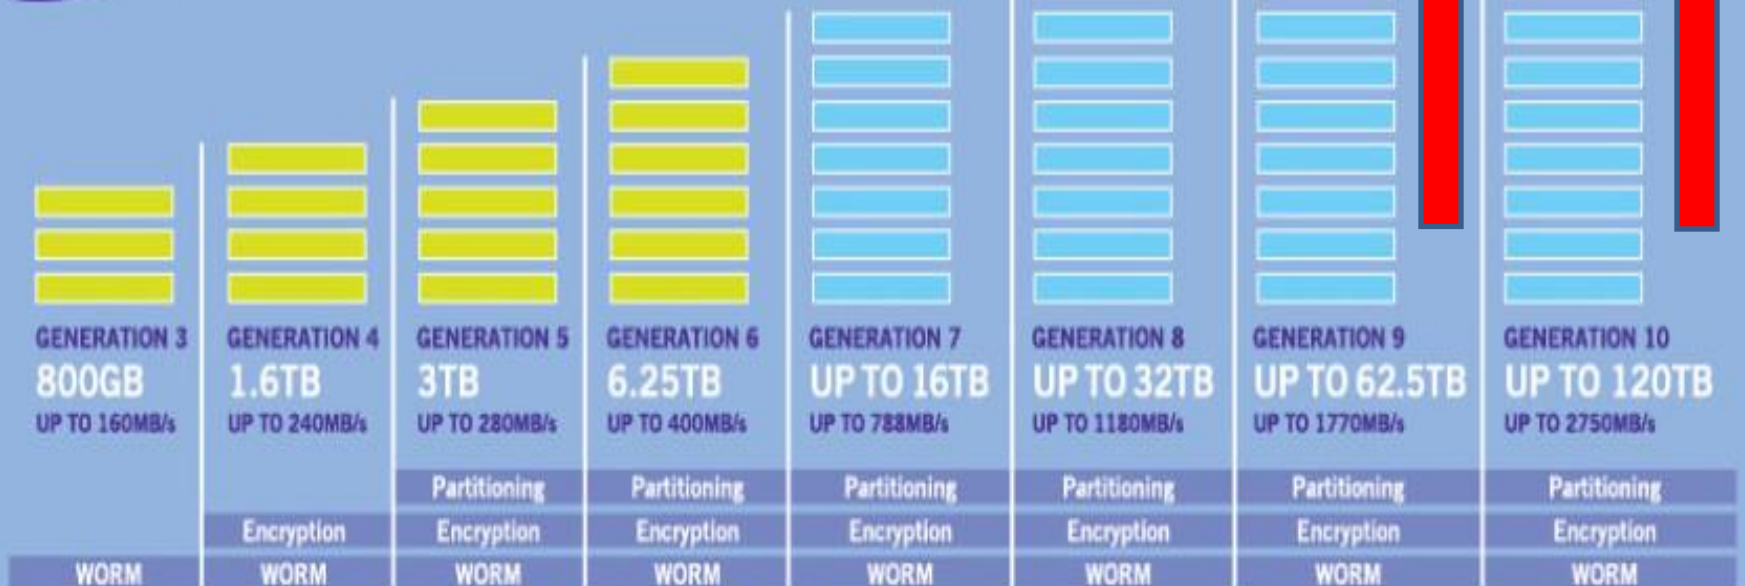
An Active Archive can scale to billions of files and holds more frequently accessed archival data.

The archive of the future will combine HDD, tape, LTFS, data management and cloud services (optionally).

Source: Horison, Inc.

LTO Roadmap – Extended

Gen 9 and Gen 10 Added Sept. 10, 2014



LTO ULTRIUM ROADMAP ADDRESSING YOUR STORAGE NEEDS

Note: Compressed capacities for generations 1-5 assume 2:1 compression. Compressed capacities for generations 6-10 assume 2.5:1 compression (achieved with larger compression history buffer).

Source: The LTO Program. The LTO Ultrium roadmap is subject to change without notice and represents goals and objectives only.

Linear Tape-Open, LTO, the LTO logo, Ultrium, and the Ultrium logo are registered trademarks of HP, IBM and Quantum in the US and other countries.



New Tape *Performance* Patents Granted

IBM Assigned 21 Tape Patents in 2015

- **Writing Multiple Files Simultaneously to Tape # ([9,021,196](#))**

Writing the data from at least two data sources to regions in a first wrap of tape on a data-source basis in a first predetermined order.

- **Reordering Access to Reduce Total Seek Time on Tape Media # ([9,021,175](#))**

A host application sends a list of UDSs to a target tape drive and requests the tape drive to reorder the list to provide better performance.

- **Data Storage Tape with Random Access Data # ([9,019,654](#))**

The data storage tape can be arranged in a concertina format. In such a format, the data storage tape can include a plurality of tape segments.

- **High Performance Cartridge Format # ([9,047,879](#))**

The data storage system is configured to use at least two different track width formats, where which of the at least two different track width formats used by the system during reading and/or writing.

- **Solid State Storage Media Cartridge # ([9,060,414](#))**

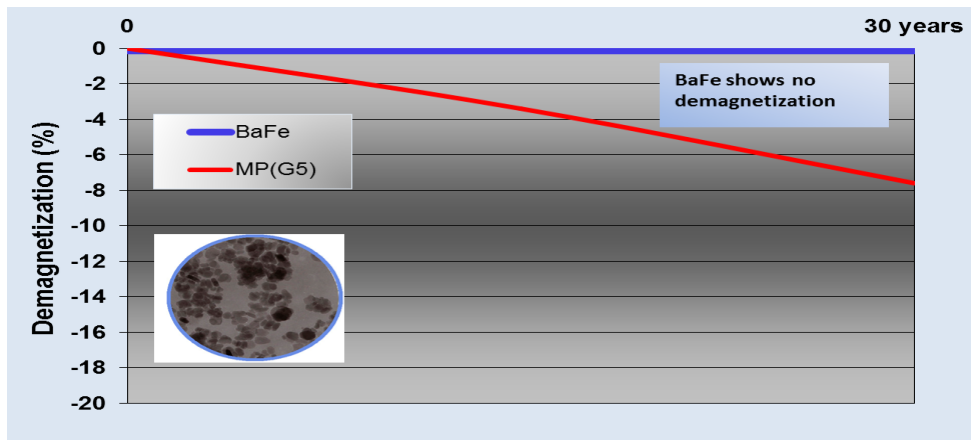
A tape cartridge adapted to house solid state storage where the cartridge may include one or more solid state storage elements held within the reel.

Tape Media Has Longest Archival Life


> 30 Years

- Oxidation is major cause of MP (Magnetic Particle) deterioration with possible data loss.
- However, BaFe is already oxidized yielding a much longer life compared to MP tape.
- MP shows slight degradation in magnetic signal over 30 years, although not detrimental to read/write performance. [Note: average HDD lasts 4.1 years]
- In Fujifilm's experiments, BaFe withstands realistic storage environment simulations and proves its reliability and stability over more than a **30 year** time period.
- BaFe Used for T10000, TS11xx and LTO Media.

Demagnetization



Media Life



Source: Fujifilm

Tape Capacity Trajectory

Fujifilm and IBM Redefine the Limits of Tape Media

Capacity 110,000,000x greater compared with IBM's first tape drive product in 1952


	2006	2010	2014	2015
Aerial Density (bits per sq inch)	6.67 Billion	29.5 Billion	85.9 Billion	123 Billion
Cartridge Capacity Lab	8 Terabytes	35 Terabytes	154 Terabytes	220 Terabytes
Number of Books Stored	8 Million	35 Million	154 Million	220 Million
Track Width (micrometers)	1.5	0.45	0.177	0.140
Linear Density (bits per inch)	400'000	518'000	600'000	680'000
Tape Material	Barium Ferrite - BaFe			
Tape Thickness (micrometers)	6.1	5.9	4.3	4.3
Tape Length (meters)	890	917	1255	1255

Source: Fujifilm and IBM

Areal Density Growth Favors Tape

Annual Areal Density Growth Rate Scenarios

- HDD – 20% to 25% – Transition to New Technology, Sensor Output, Lithography
- NAND Flash – 25% to 30% – Lithography and Endurance
- TAPE – 40% to 80% -- No Lithography Issues, Mechanical Realities



Device Reliability Levels

BER

Published Values for BER – Bit Error Rate
Metric - Hard Read Errors per Bits Read

T10K and TS11xx
 1×10^{19} bits



Flash SSD
 1×10^{17} bits

LTO-5, 6, 7
 1×10^{17} bits



Enterprise FC/SAS
 1×10^{16} bits

Enterprise SATA
 1×10^{15} bits



Desktop SATA
 1×10^{14} bits



Tape Nearing
 1×10^{20} bits

- When the HDD fails, a 4 TB RAID disk drive will take >2 days to rebuild, a 30TB RAID disk drive will take nearly a month to rebuild.
- Encryption, deduplication and compression techniques increase the likelihood that the loss of a single bit will make all the data on a drive unrecoverable.
- The ERA of Modern Tape Has Arrived!

Source: Horison, Inc.

Energy and Infrastructure Expense

Do You Need This – To Support This?



Office Building
3-5 watts/sq. ft.
\$60-90/sq. ft.




Data Center
150 watts/sq. ft.
\$700-950/sq. ft.



Infrastructure – External Consumption	Avg.
Chillers, cooling, air-conditioning, pumps,	24%
Uninterruptible power supply	8%
Air movement, circulation, fans etc.	10%
Misc. lighting, security, surveillance, appliances	3%
Total	45%
IT Hardware – Internal Consumption	
Servers	>30%
Disk drives and control units	12%
Tape drives, robotic libraries	<3%
Network gear, SAN switches and other devices...	10%
Total	55%

TCO Comparisons – Disk and Tape

Total Cost of Ownership



The 5-Year TCO for Disk Ranges 2-15x Higher Than Tape for Backup and Archive Applications – the Gap is Widening

Tape Value Proposition

C Customer
V Value
P Proposition

Component	Remarks
Price	Tape Has the Lowest Acquisition Price - \$/GB
TCO	HDD's TCO is 2-15x Higher Than Tape (For Backup/Dedup and Archive)
Performance	HDD & SSD Much Better for IOPs, Tape Better for Data Rate HDD & SSD for Random and Sequential Access, Tape Sequential Access
Capacity	Tape Cartridge Capacity @10TB (25TB compressed), HDD @8TB (<50%)
Reliability	Tape Reliability (BER) Higher Than HDD for LTO and Enterprise Tape
Energy Consumption	Tape Uses Much Less Energy Than HDDs and SSDs
Portability	Tape Media Easily Portable, HDDs Difficult to Move Safely
Encryption	All Tape Drives offer Encryption, Option on Some HDDs But Seldom Used
Media Life	>30 Years for all Modern Tape, Avg. ~4.1 Years for HDDs
Media Conversion	Tape Supports Current and 2 Prior Versions, HDD Requires Replacement
Recording Limits?	No Foreseen Limits for Tape, HDDs Facing Areal Density Limits



Things Are Moving So Fast...

Even the Future is Obsolete

Storage Landscape in 2020....



- The Flash/SSD and Tape Centric Data Center Arrives as Amount of Data Stored on SSD *and* Tape Will Grow Faster Than Data Stored on Disk - Disk Limitations Mount
- Robotics Proliferate as Automated Tape Libraries Deliver Exascale Architectures
- Expect Faster Access Times With Integrated Flash-HSM-LTFS Active Archive Solutions
- Mobile (And Wearable) Devices Will Generate Over 70% of All Digital Data by 2020
- Clouds Will *Not* Take Over the IT World - but Will Continue to Grow (Esp. Private)
- Surveillance Will Drive Huge, Multi-media Storage and Archive Requirements
- Hacking Will Become the Biggest Threat to *Data* Security – Not Reliability
- Drones Will Become the Biggest Threat to *Physical* Security
- Energy Management Will Become Even More Critical as Building New Data Centers Becomes Cost Prohibitive - Favors Large Tape Storage Systems
- Tier 3 (Tape) Data Will Grow Faster Than Other Tiers as Archival Requirements Soar
- Storage Capacities Will Force Users to Optimize the Storage Hierarchy (HPSS Key)
- **Strategy Shift – “Built to Last” Gives Way to “Built to Change”**

