Fujifilm's 7th Annual Global IT Executive Summit

October 7 – 10, 2015

A peek into the future (The future of Tape)

Dr. Mark Lantz
Principle Research Staff Member, Manager Advanced Tape Technologies
IBM Research - Zurich

Excerpts from IBM EDGE May 11-15, 2015

Outline

- Introduction: The role of tape in the era of big data
- The Future of Tape
  - Tape areal density trends and future scaling potential
  - New **world record tape** low cost particulate tape areal density demo of **123 Gb/in^2** (IBM – FujiFilm collaboration)
  - Technologies enabling the 123 Gb/in^2 demo
  - Tape technology roadmap
- What About the Other Storage Technologies?
- Investing in the Future
- Conclusions

The data deluge

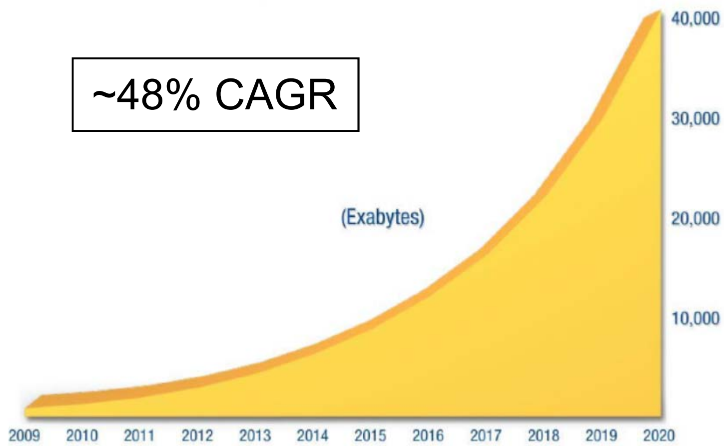

Source: IDC's Digital Universe Study, sponsored by EMC, December 2012

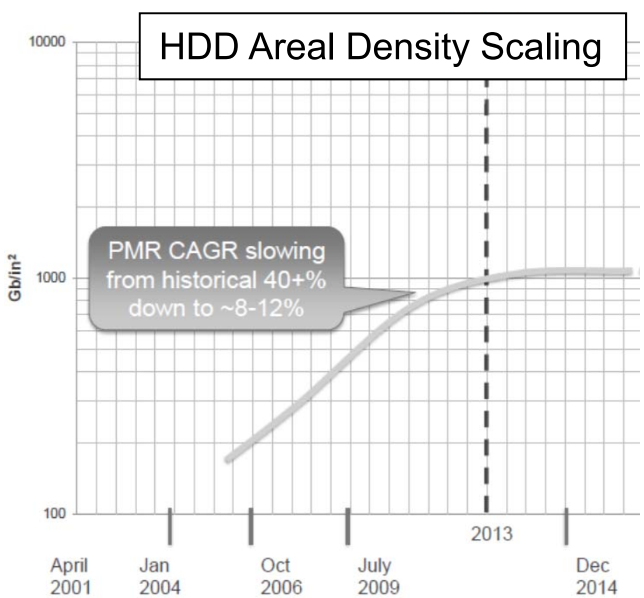

Source: D. Anderson, 2013 IEEE Conf. on Massive Data Storage

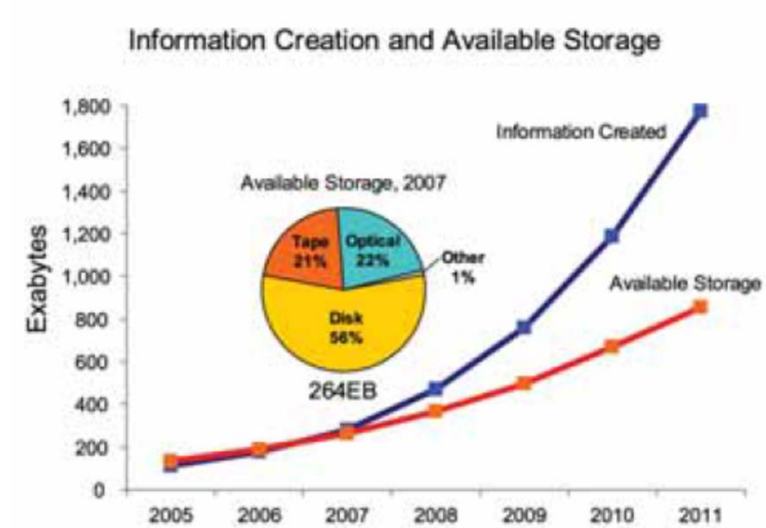

80% of all files created are inactive – no access in at least 3 months!

Tape advantages for long-term storage

- **Very energy efficient**: no power needed once data is recorded
- **Very secure**:
  - Data is inaccessible when cartridge is not mounted
  - Drive level encryption
  - Portable
- **Very long** expected media **lifetime** (30+ years)
- **Very reliable**: Typically no data loss in case of drive failure
- **Main net advantage of tape for archival storage is cost**

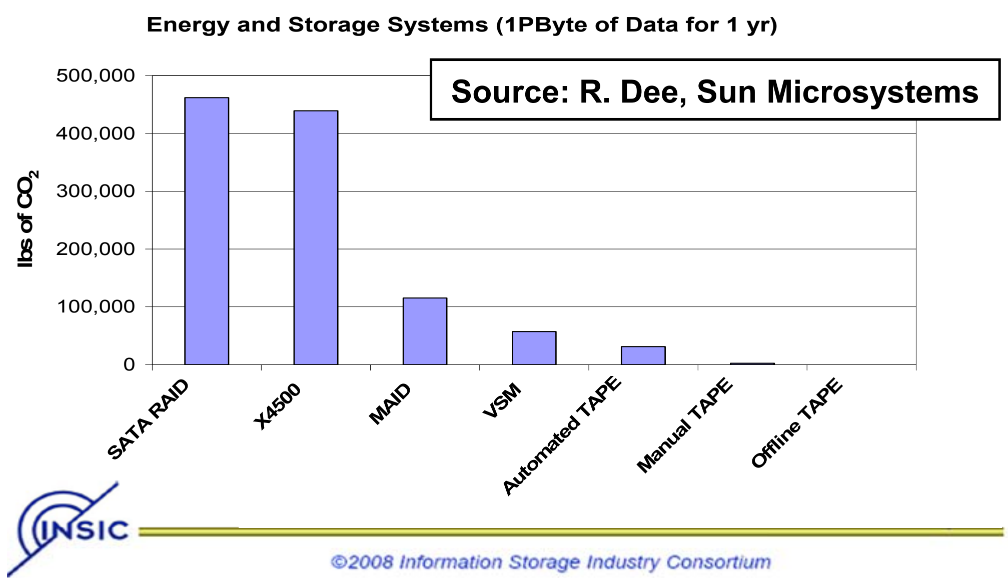

Tape TCO for Long-Term Archiving – Clipper Group

Recent studies from the Clipper Group:

1) Continuing the Search for the Right Mix of Long-Term Storage Infrastructure – A TCO Analysis of Disk and Tape Solutions (15 July 2015) Report # TCG2015006
2) The Impact of LTO-7 on The TCO of Long-Term Storage (15 Sept. 2015) Report #TCG2015008

Investigate 9 year TCO of a 1PB archive that grows to 52 PB (55% CAGR)

Major Finding: 6.7x TCO advantage of LTO Tape of Disk

HDD Areal Density Scaling:

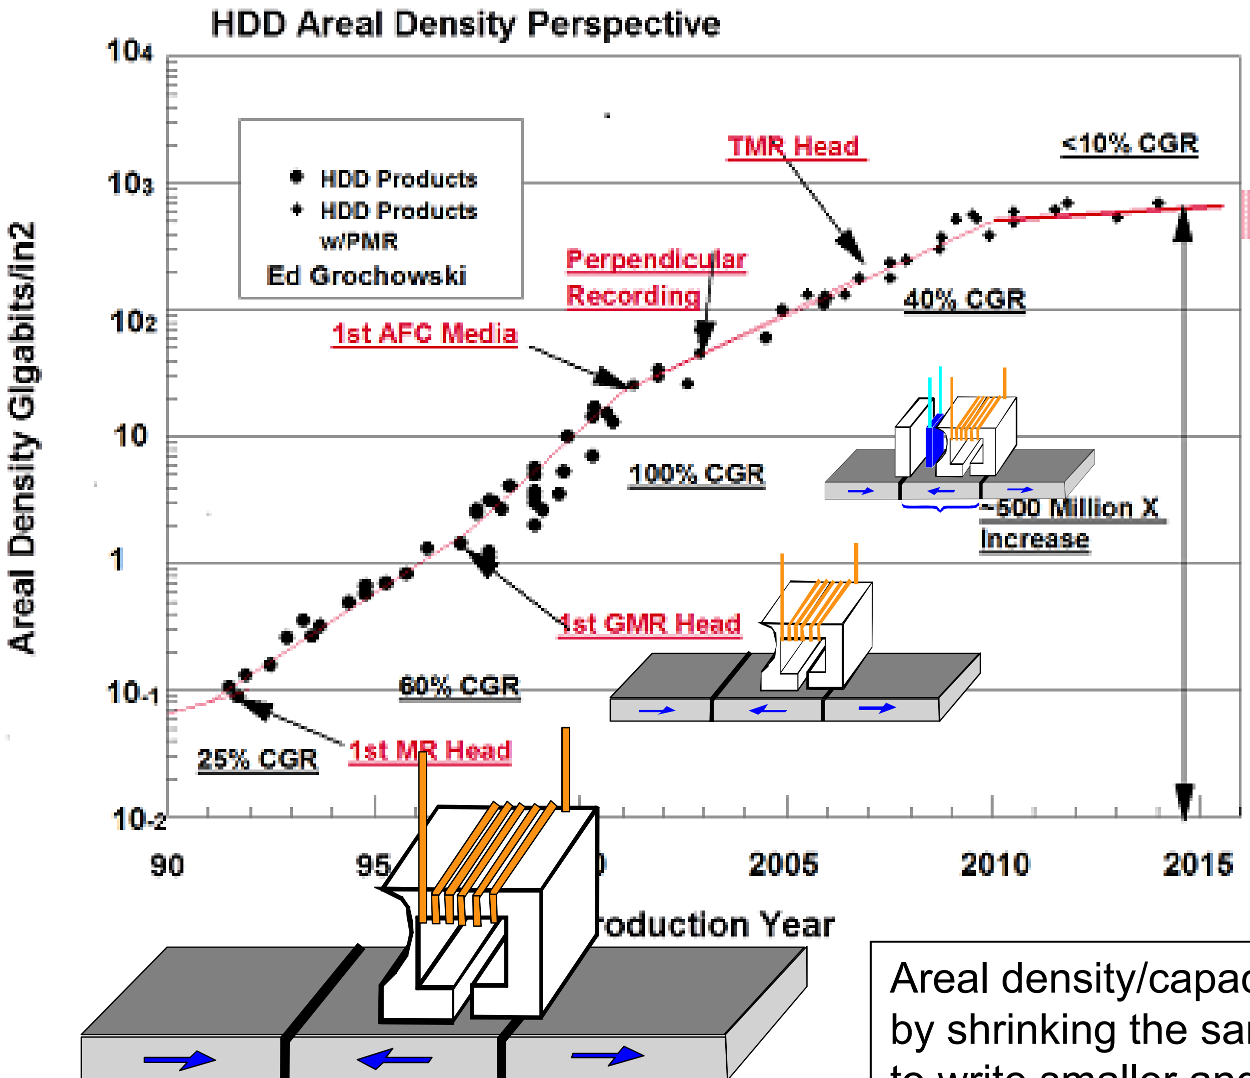

Areal density/capacity scaling achieved by shrinking the same basic technology to write smaller and smaller bits on disk

Ref: http://www.storageacceleration.com/author.asp?section_id=3670&doc_id=274482

Noise and Magnetic Media Structure

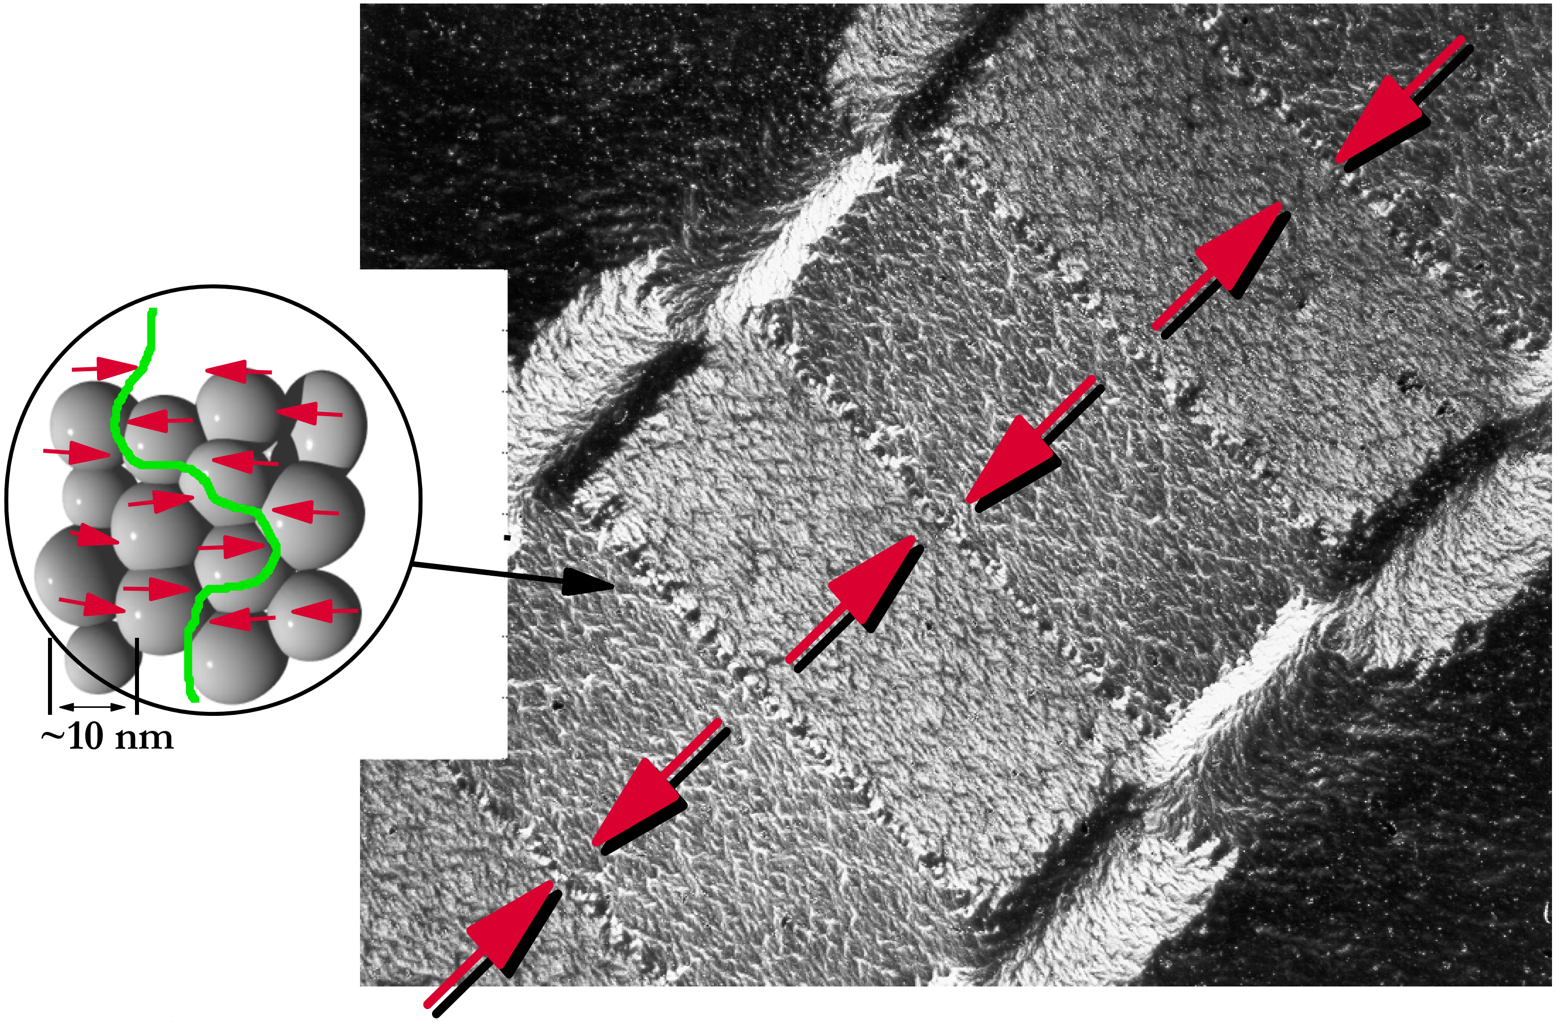

Information is encoded in transition edge. Large grains → media noise
To shrink the size of a bit, we need to shrink the size of the grains

If grains become too small, magnetic state is unstable → superparamagnetic effect

The Superparamagnetic "Limit"

Magnetic Media "Trilemma":

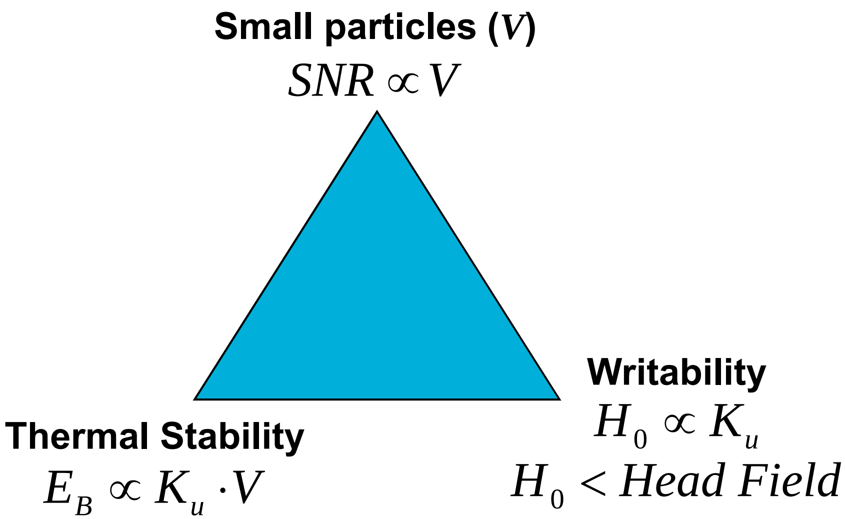

HDD has reached the limit of (known) materials to produce larger write fields.

Technologies to go beyond the superparamagnetic limit:

- Two dimensional magnetic recording (TDMR)
- Heat Assisted Magnetic Recording (HAMR)
- Microwave Assisted Magnetic Recording (MAMR)
- Bit Patterned Media (BPM)

Magnetic tape (r)evolution

Product / Year:	IBM 726 /1952	LTO7 / 2015	TS1150 /2014	**Demo 2015**
Capacity:	2.3 MBytes	6 TBytes	10 TBytes	220 TBytes
Areal Density:	1400 bit/in^2	4.28 Gbit/in^2	6.7 Gbit/in^2	123 Gbit/in^2
Linear Density:	100 bit/in	485 kbit/in	510 kbit/in	680 kbit/in
Track Density:	14 tracks/in	8.83 ktracks/in	13.2 ktracks/in	181 ktracks/in

HDD vs. Tape Areal Density Scaling:

IBM-FujiFilm demonstration of 123 Gb/in2 on BaFe tape

Goal: Demonstrate the feasibility of tape roadmap for the next 10+ years

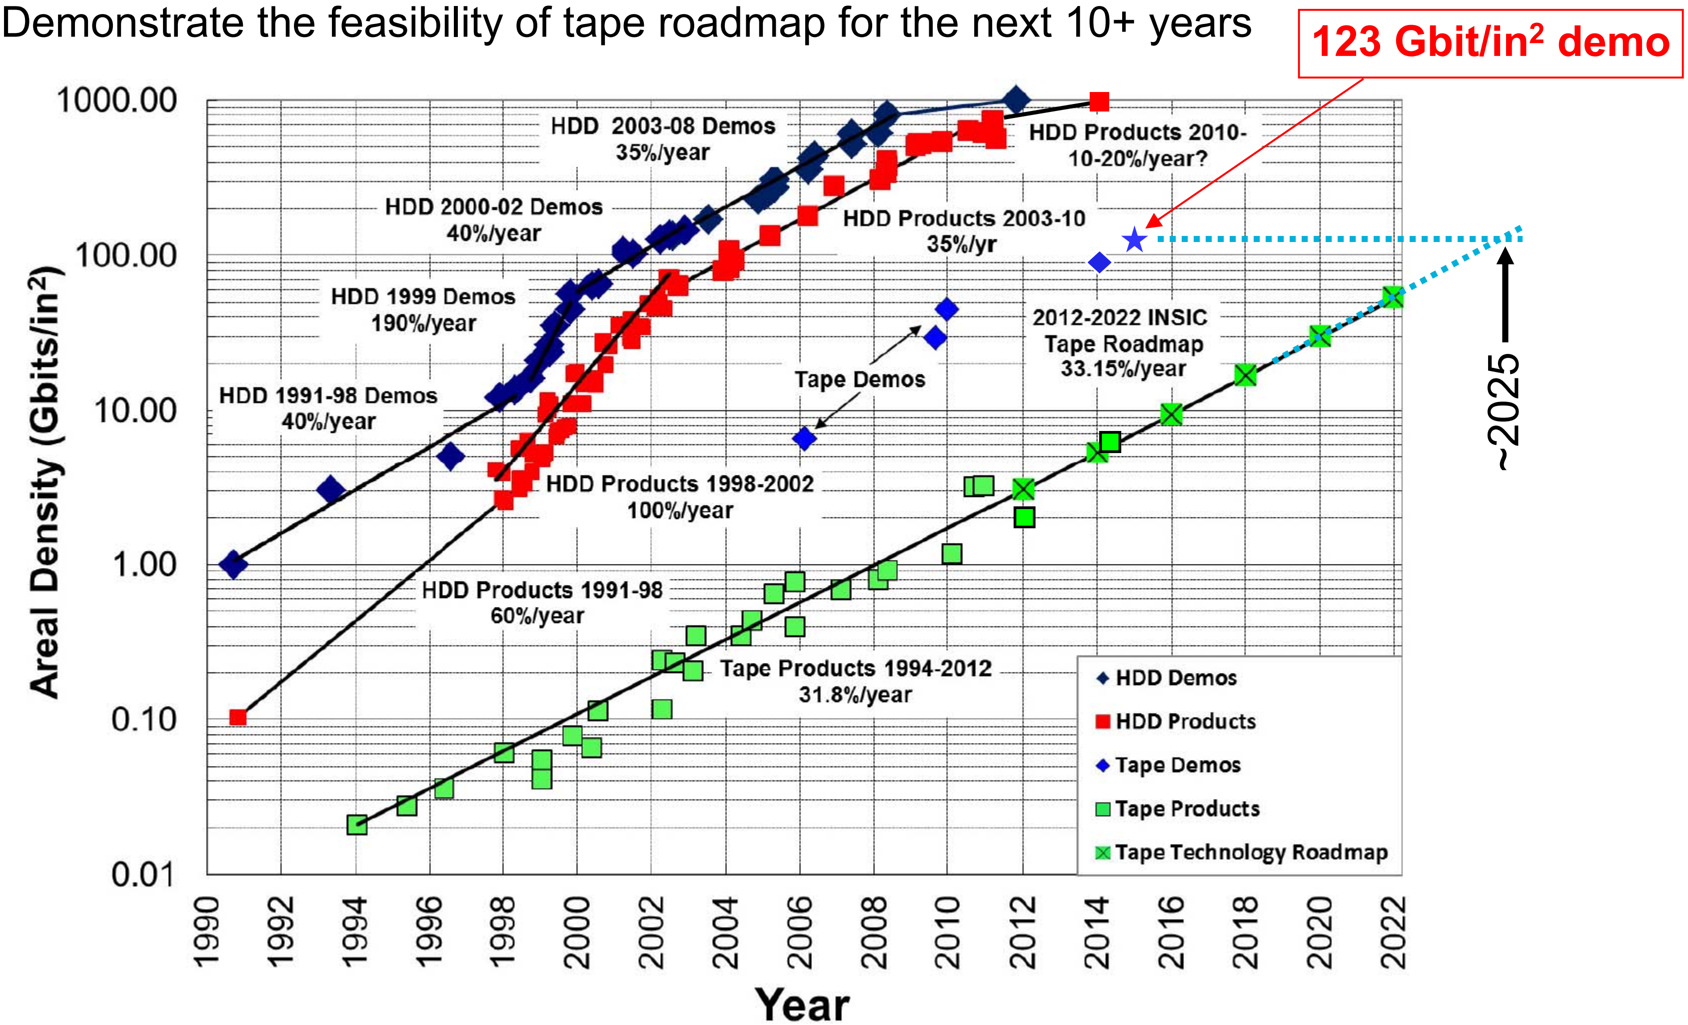

(Source: INSIC 2012-2022 International Magnetic Tape Storage Roadmap)

2015 Storage Bit Cells and Extendibility

- Scaled bit cells:
- Magnified 25x:

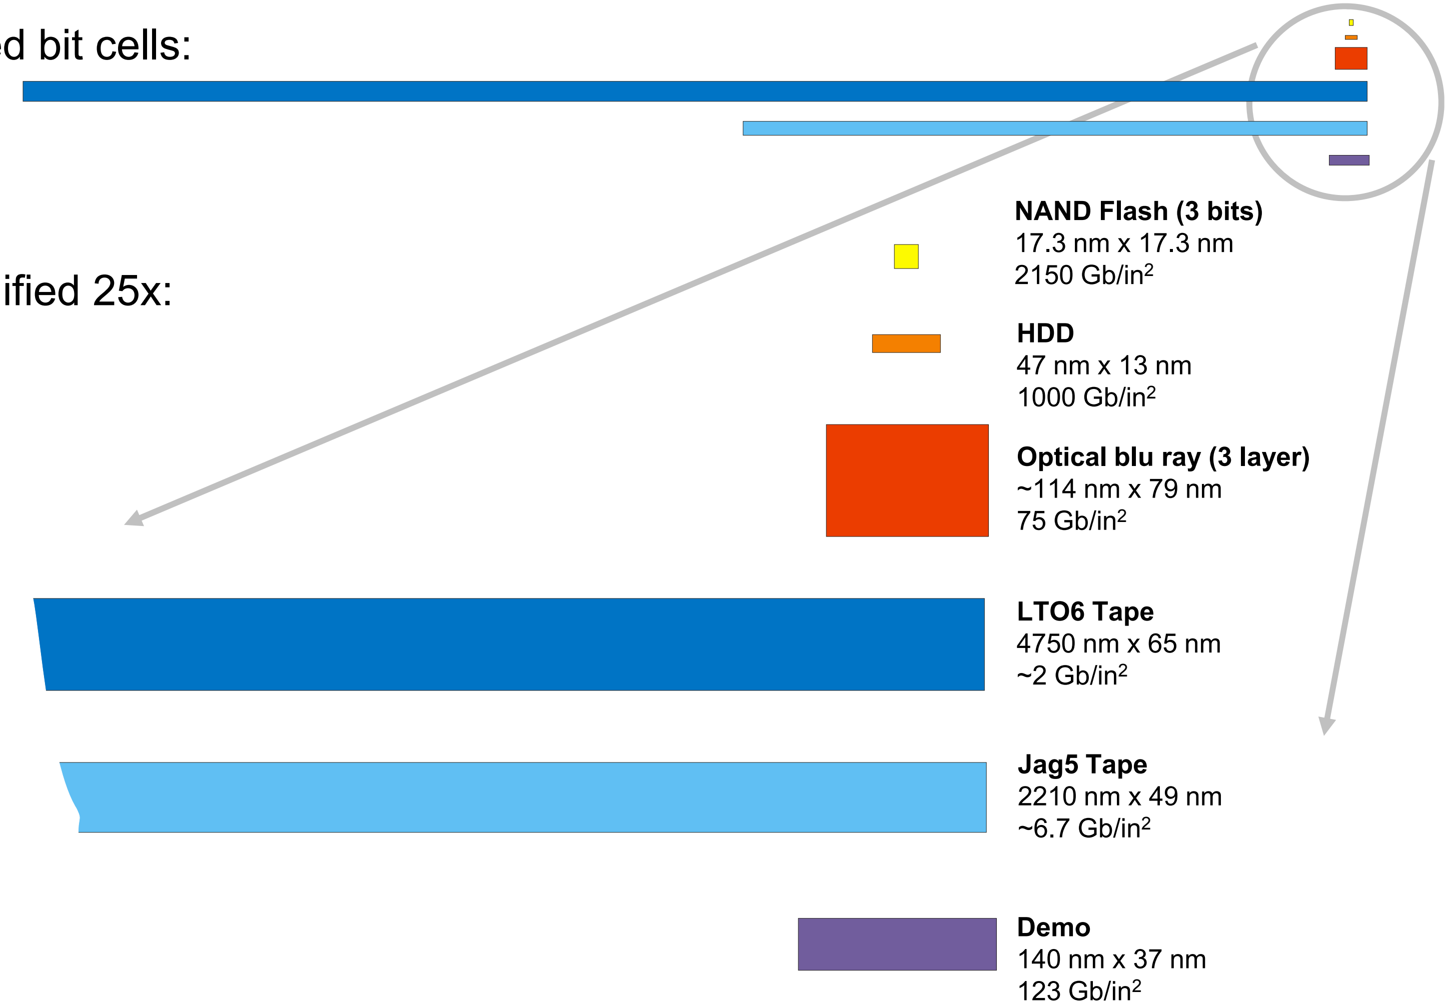

→ Tremendous potential for future scaling of tape track density

→ Key technologies: improved track follow servo control
improved media, reader, data channel

Demo Technologies

Focus on aggressive track density scaling

- Require:
  - dramatic improvement in track following → enables track width reduction
  - reduce reader width from a few microns to 90 nm
- Ultra narrow reader results in a dramatic loss in read back signal that must be compensated for with
  - **improved media technology** → require improved writer technology
  - improved signal processing and coding
  - improved reader technology

Servo pattern design for high areal density demo

Main design goal: nm-scale positioning fidelity

- Increased azimuth angle ⇨ increased resolution
- Increased pattern density ⇨ increased servo bandwidth and resolution

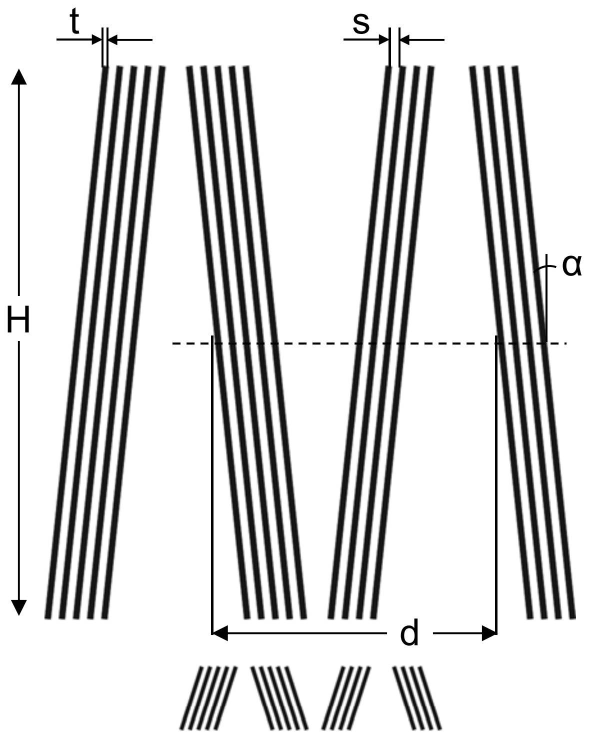

Standard LTO Pattern
H = 186 μm, t = 2.1 μm, s = 5 μm
α = 6°, d = 100 μm

Demo Pattern
H = 23.25 μm, t = 1.0 μm, s = 2.4 μm
α = 24°, d = 52 μm

4x angle
2x rate

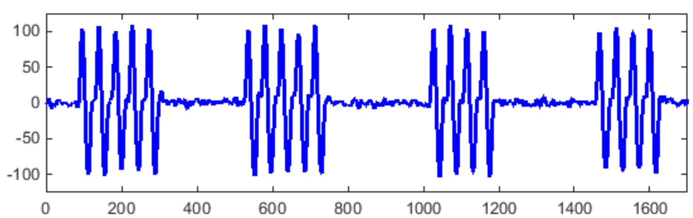

Compatible with Future 16 Data Band Tape Format

Synchronous servo channel

- Servo channel decodes the readback signal from the servo pattern and provides position information to the track follow control system
- Servo channel optimized for p-BaFe → improved resolution
- Optimized servo channel in combination with advanced BaFe media formatted with the 24°demo servo pattern provides **nanoscale position information**

Servo readback signal

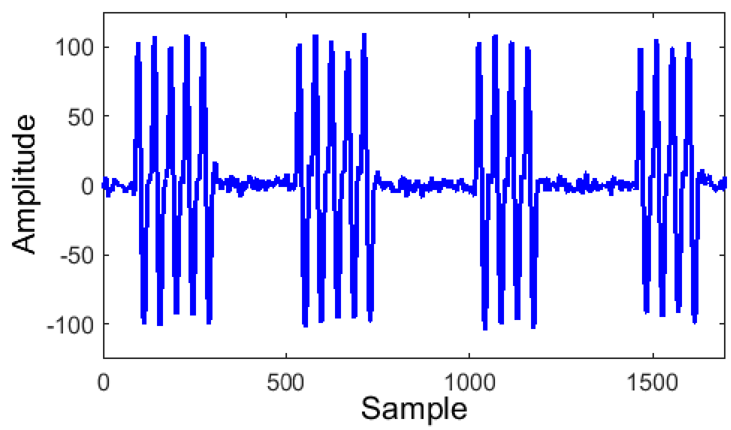

Servo channel

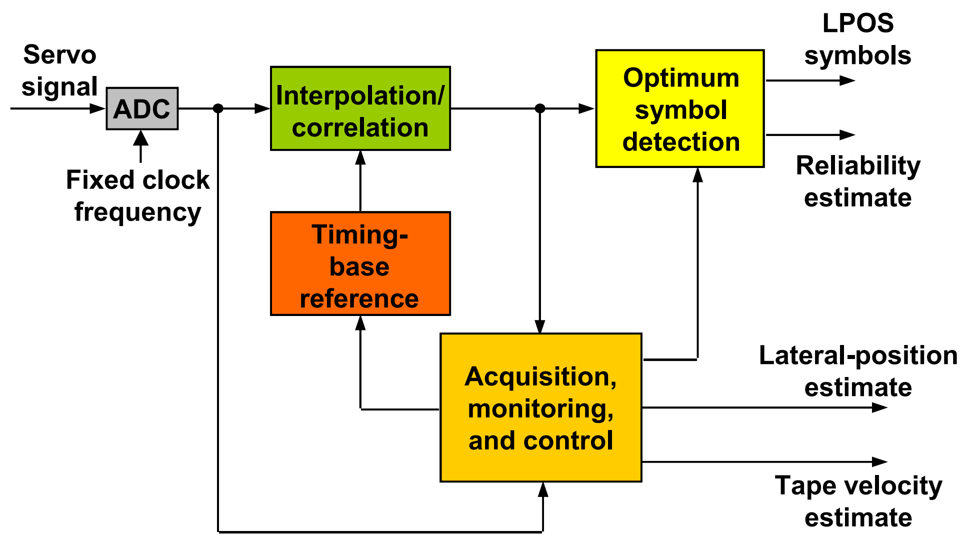

New H∞ track-follow control system

- **Key features**
  - Prototype high bandwidth head actuator
  - A speed dependent model of the system delay is used for control design
  - The tape speed is used as a parameter to select the controller coefficients
  - Disturbance rejection is enhanced at the frequencies of the tape path disturbances

High Bandwidth Actuator

Track-follow control system

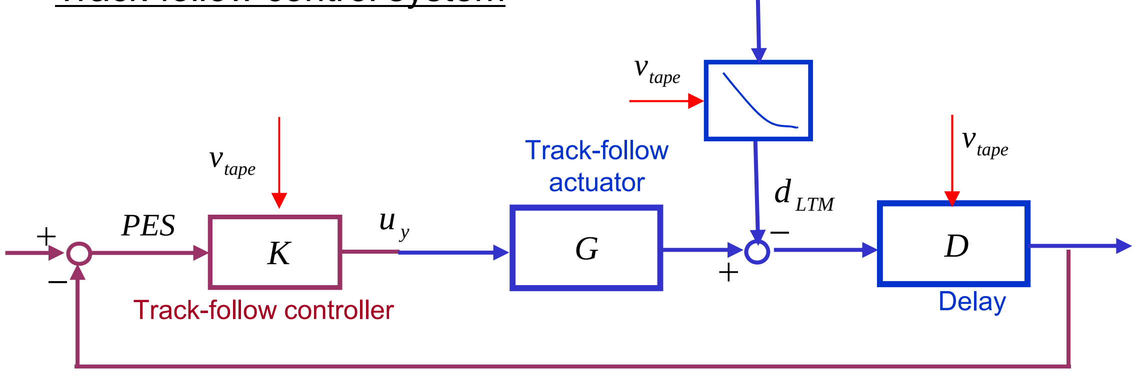

Actuator Response

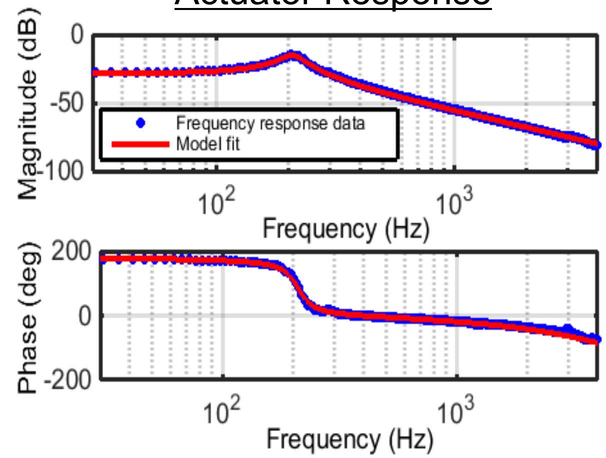

Prototype tape transport & hardware platform

- Precision flangeless tape path with grooved rollers & pressured air bearings to minimize disturbances
- TS1140 electronics card for reel-to-reel control and analog front end
- FPGA Board: System-on-Chip (SoC)
  -> Servo channels
  -> Microprocessor for synchronous track-follow (TF) servo controller

FPGA Board

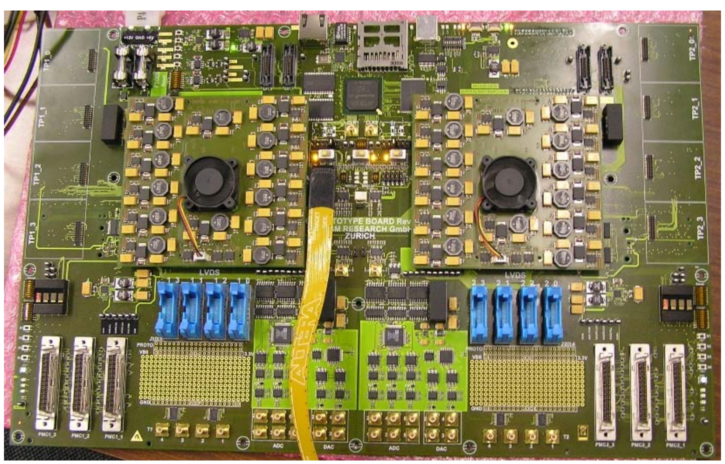

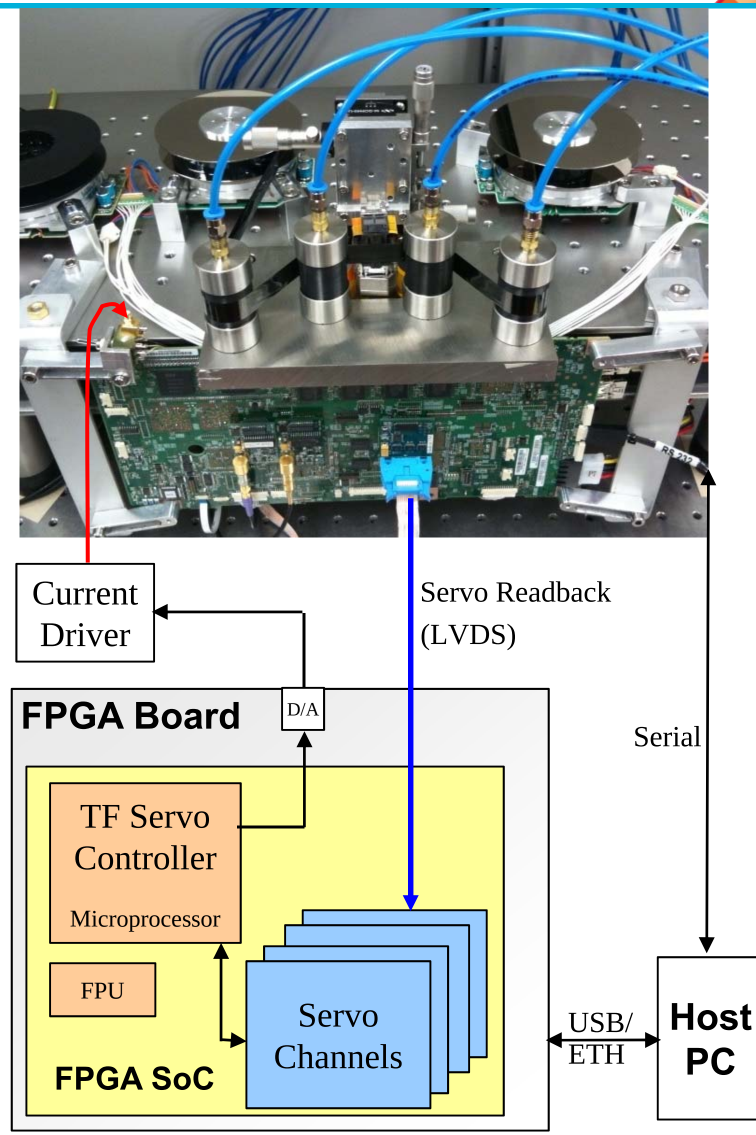

Track-follow performance on BaFe tape

- Track width computation based on measured position error signal: PES (INSIC method)
- σPES = standard deviation of position error signal: measure of track following fidelity
- Track width = 2*√2 * 3*σPES + Reader Width   (Reader Width = 90nm)

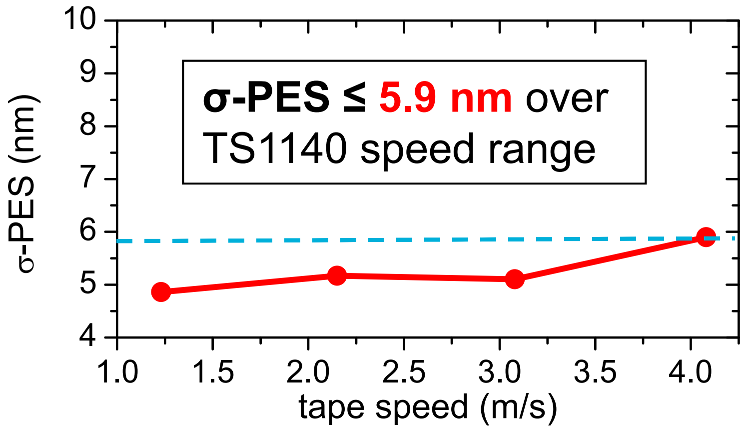

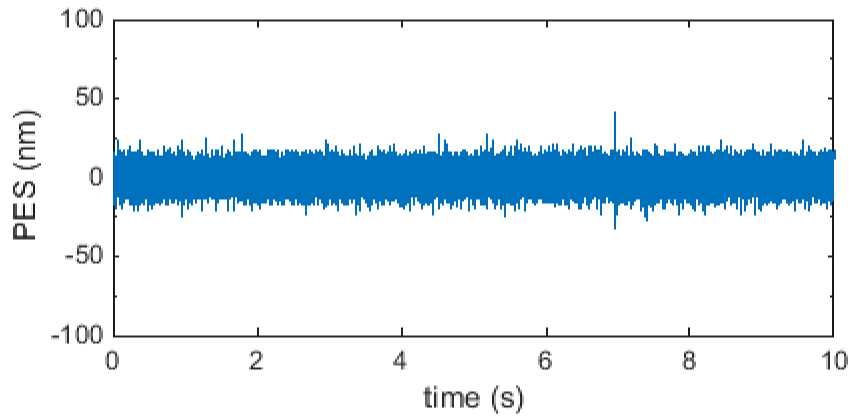

Reader Width = 90nm
σ-PES ≤ 5.9 nm

Track width = 140 nm
Track density = 181 ktpi

Advanced BaFe Media Technology

Key technologies for advanced tape media

1. Fine magnetic particles with high coercivity → archival lifetime
2. Smooth surface
3. Perpendicular orientation of magnetic particles

Metal particle vs. Barium-ferrite particle

	Metal particle (MP)	Barium ferrite (BaFe)
Shape	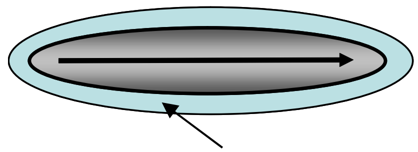 Acicular	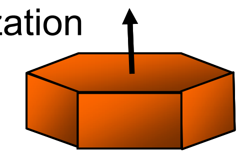 Hexagonal platelet shaped
Origin of magnetic energy	Shape anisotropy	Magneto-crystalline anisotropy
Material	FeCo alloy	$BaO(Fe_2O_3)_6$ Oxide
Passivation layer	Needed	Not needed

- The magnetic properties of BaFe particles are NOT affected by its shape.
- BaFe particles do NOT need an oxide passivation layer because it is an oxide.
- The size of BaFe particles can be reduced while maintaining high coercivity.

TEM image of fine Barium-ferrite particles

FUJIFILM Value from Innovation | Edge2015 #ibmedge

Latest MP
Volume :2850 nm^3
coercivity:189kA/m(2380Oe)

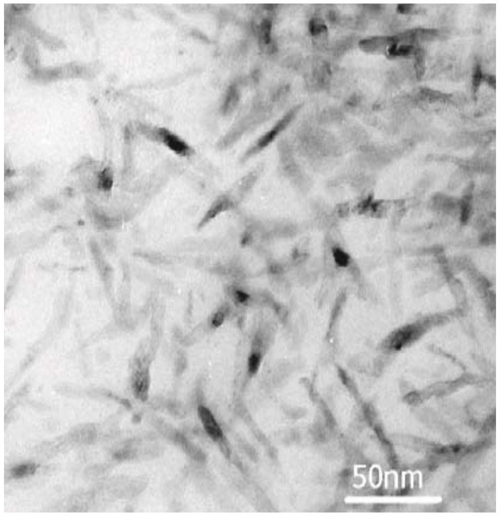

Demo Tape BaFe
Volume : 1600 nm^3
coercivity: 223kA/m(2800Oe)

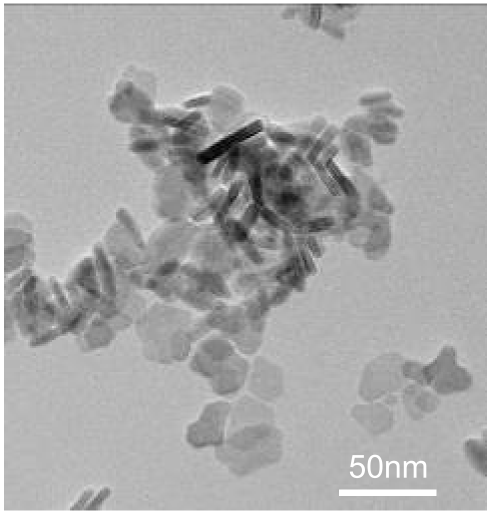

The volume of barium ferrite particle used in the demo tape is 45% smaller than the latest MP, reducing media noise and improving SNR

SEM Image of tape surface

FUJIFILM Value from Innovation

Edge2015 #ibmedge

Latest MP tape

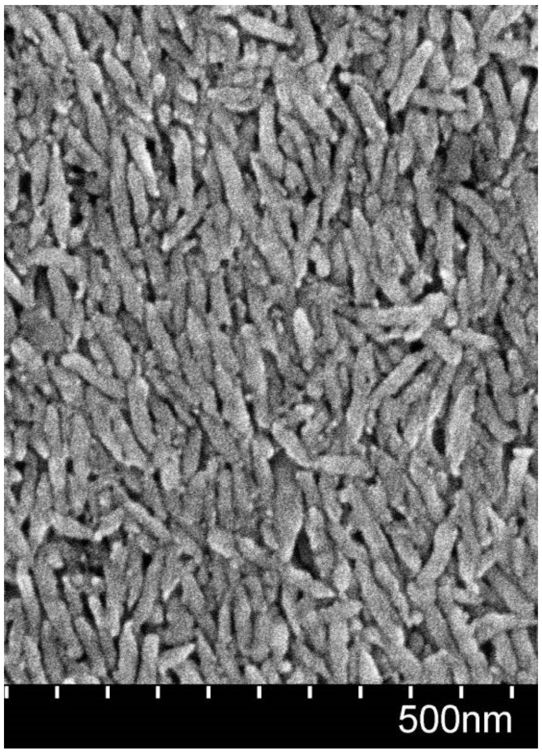

123Gb/in^2 demo tape

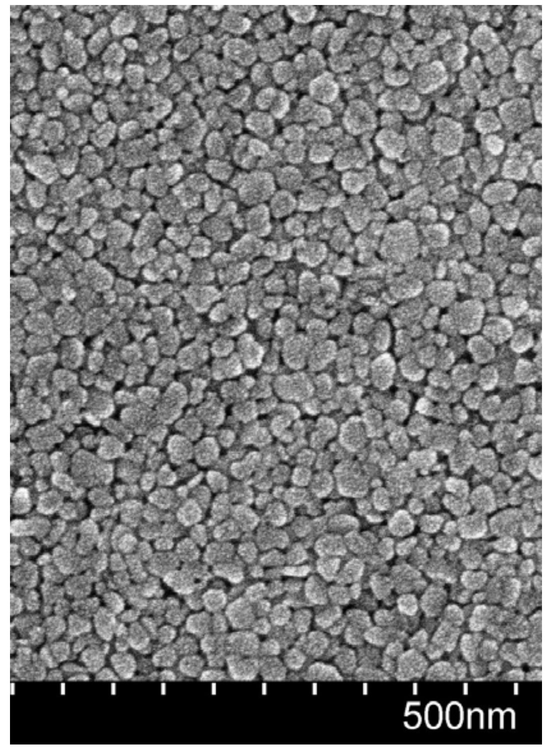

Barium ferrite particles are well isolated and packed with high density.

Particle volume vs. coercivity

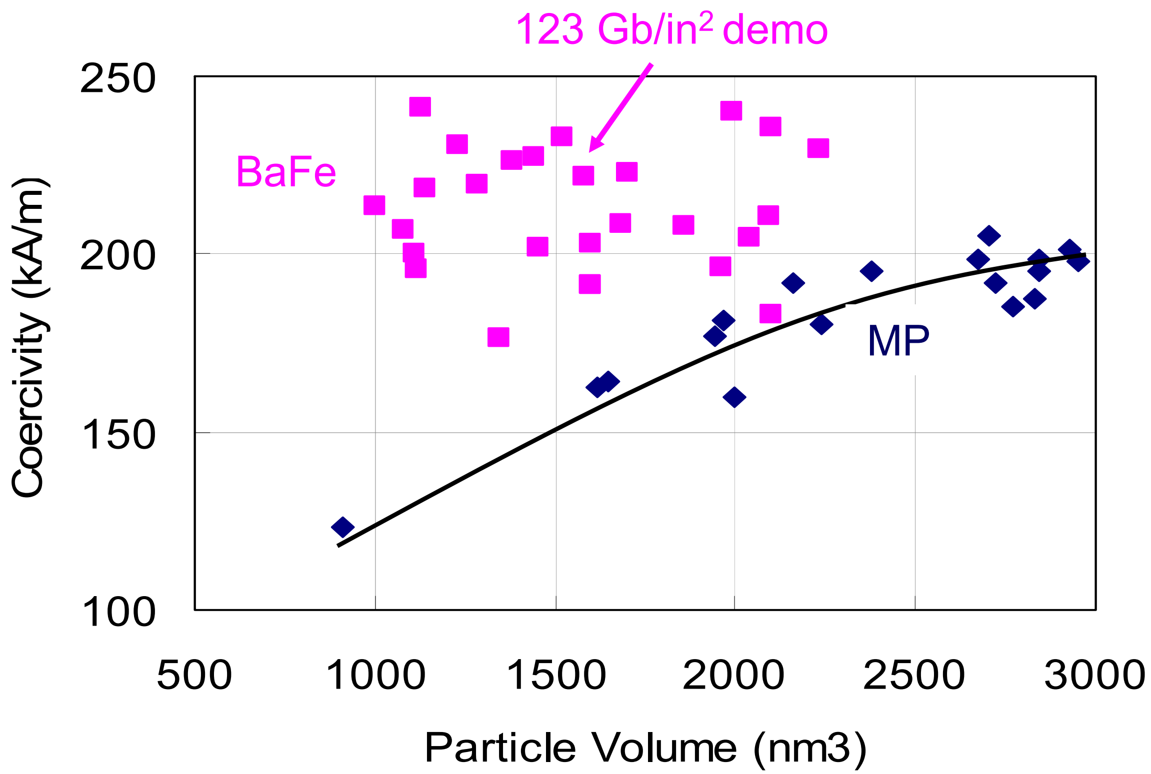

- The coercivity metal particles smaller than 3000nm^3 decreases with size
- The coercivity of barium ferrite particles can be tuned independently of size enabling small particle media with long archival lifetime
- BaFe particles as small as 1000nm^3 have been developed indicating the further scaling potential of BaFe tape

Surface profile

FUJIFILM Value from Innovation | Edge2015 #ibmedge

	Latest MP tape	TS1150 JD tape	123Gb/in^2 Demo tape
Optical interferometry roughness	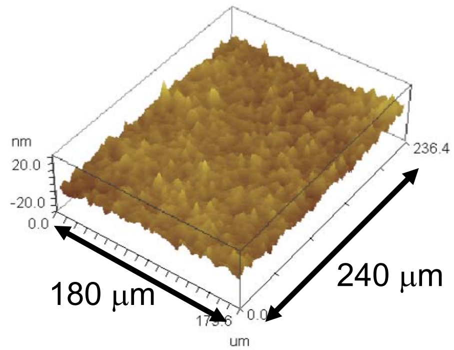 Ra 2.0nm	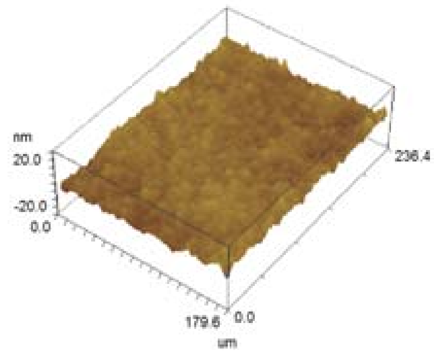 Ra 1.6nm	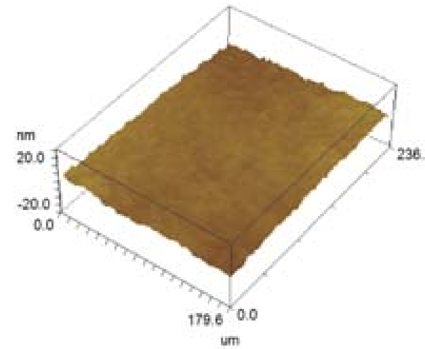 Ra 0.9nm
AFM	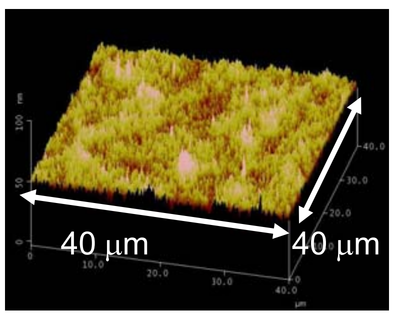 Ra 2.4nm Rz 40nm	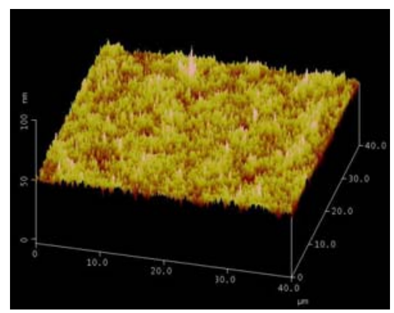 Ra 2.0nm Rz 34nm	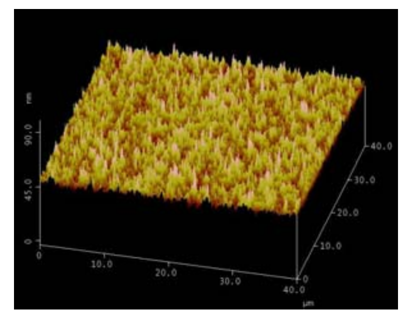 Ra 1.8nm Rz 27nm

Reduced surface roughness of demo tape increases the media SNR

22

Perpendicular orientation

Longitudinal orientation (MP tape)

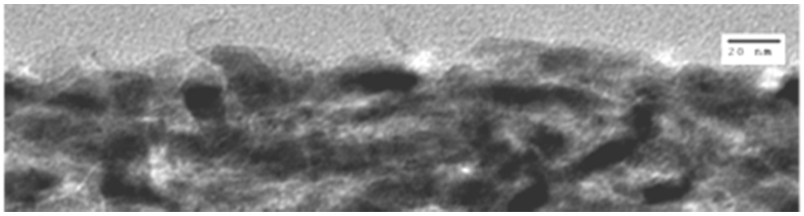

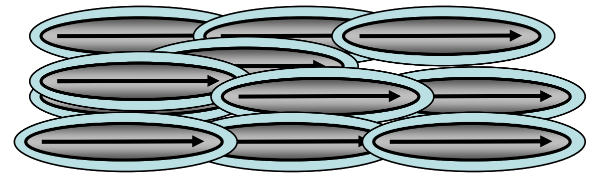

Random orientation (TS1140 JC and TS1150 JD tape)

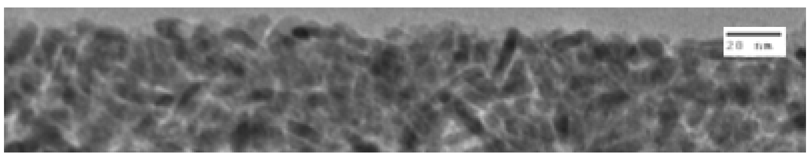

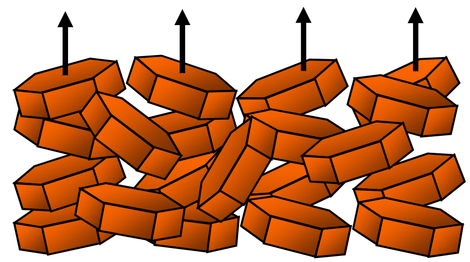

Highly perpendicular orientation (123Gb/in^2 demo tape)

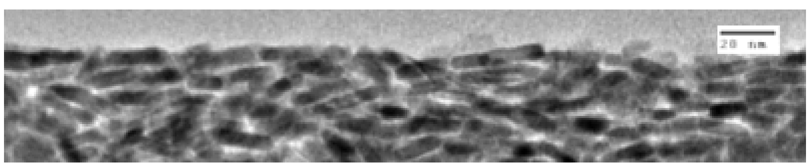

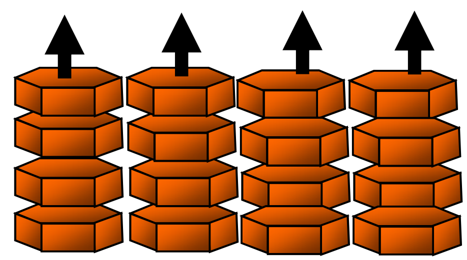

The perpendicular orientation of BaFe particle provides a strong increase in SNR

Read/write performance

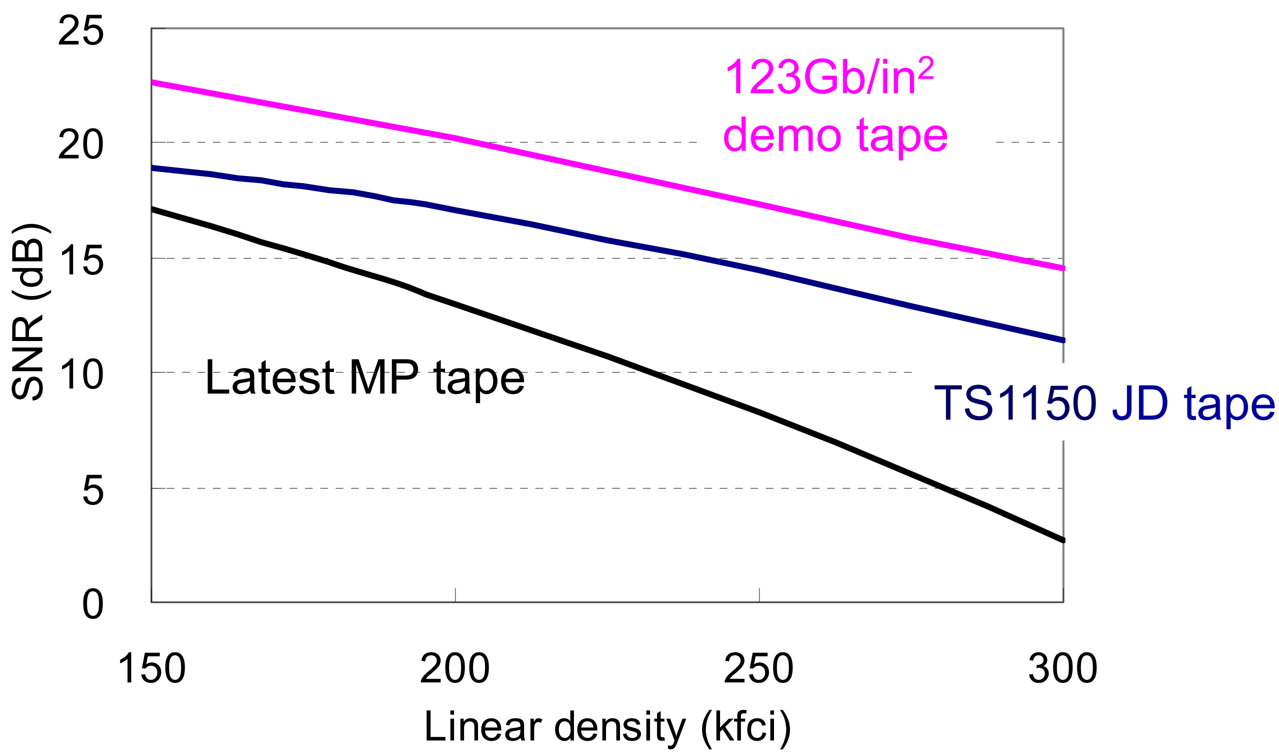

The combination of small particle volume, smooth surface and perpendicular BaFe particle orientation provide a major increase in SNR.

Enhanced Write Field Head Technology

Magnetic Media "Trilemma":

Small particles (V)

$$SNR \propto V$$

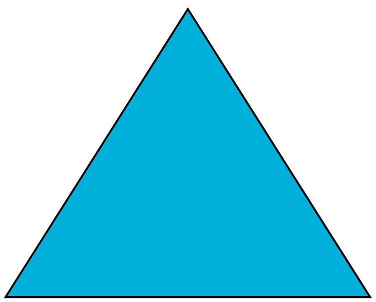

Thermal Stability

$$E_B \propto K_u \cdot V$$

Writability

$$H_0 \propto K_u$$

$$H_0 < Head\ Field$$

IBM developed a **new high moment (HM) layered pole write head** that produces much larger magnetic fields enabling the use of smaller magnetic particles

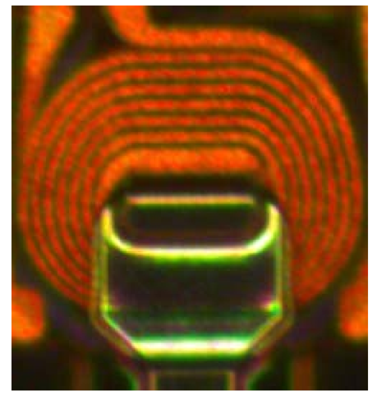

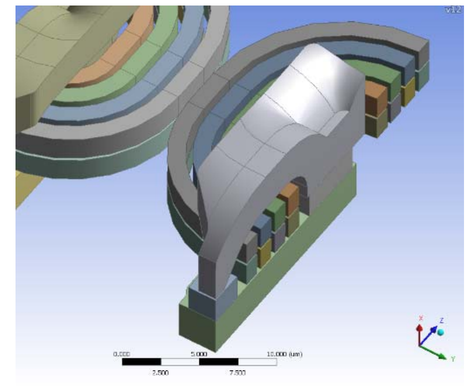

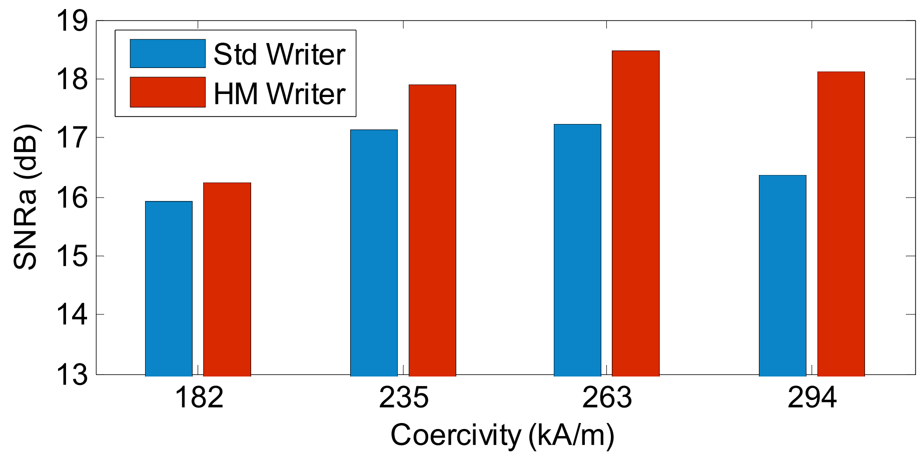

Increasing media coercivity

Iterative decoding

Edge2015 #ibmedge

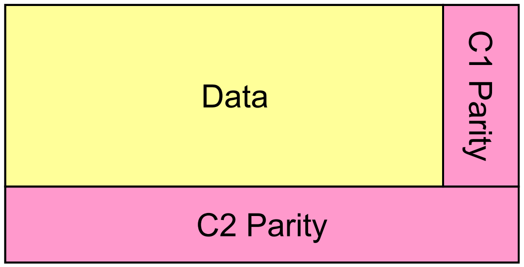

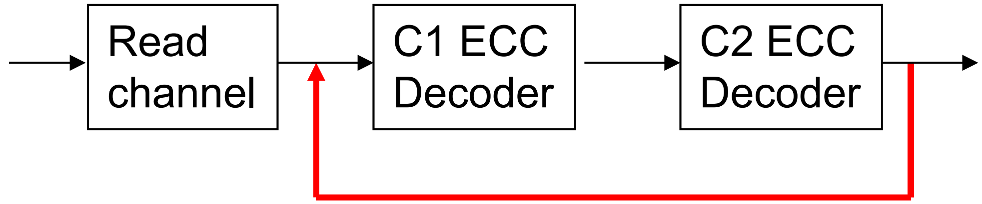

- **A user byte-error rate of 10^{-20}** can be achievable using two C1-C2 iterations with a byte error rate of ≈ $4 \cdot 10^{-2}$ at the output of the detector

- With EPR4 detection $4 \cdot 10^{-2}$ byte error rate ≈ 10^{-2} bit error rate

- Require **SNRa ≈ 10.5 dB** at the input of the detector to achieve a raw bit error rate < 10^{-2} at the output of the detector

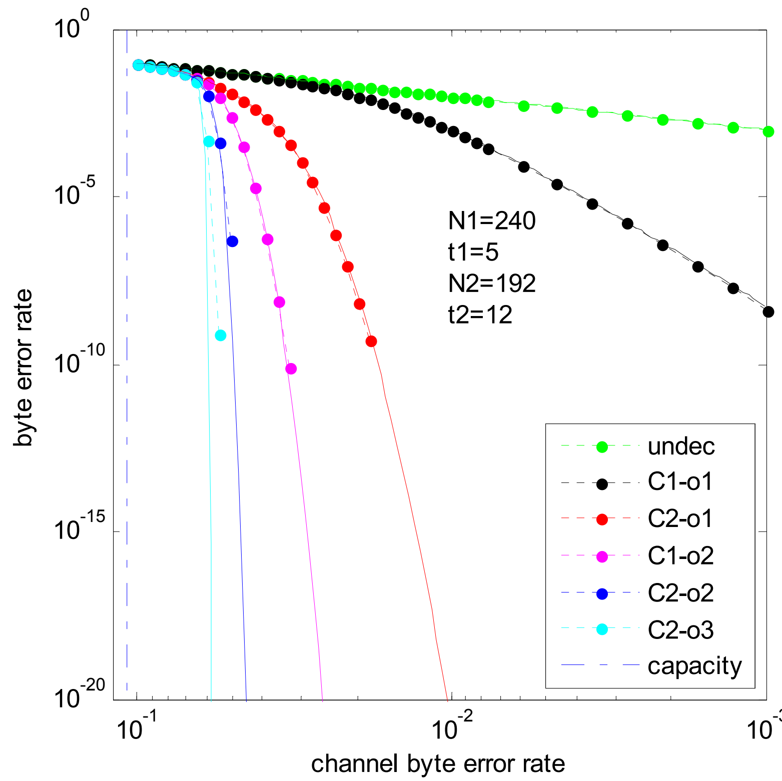

Recording performance of BaFe with High moment writer & 90 nm GMR Reader

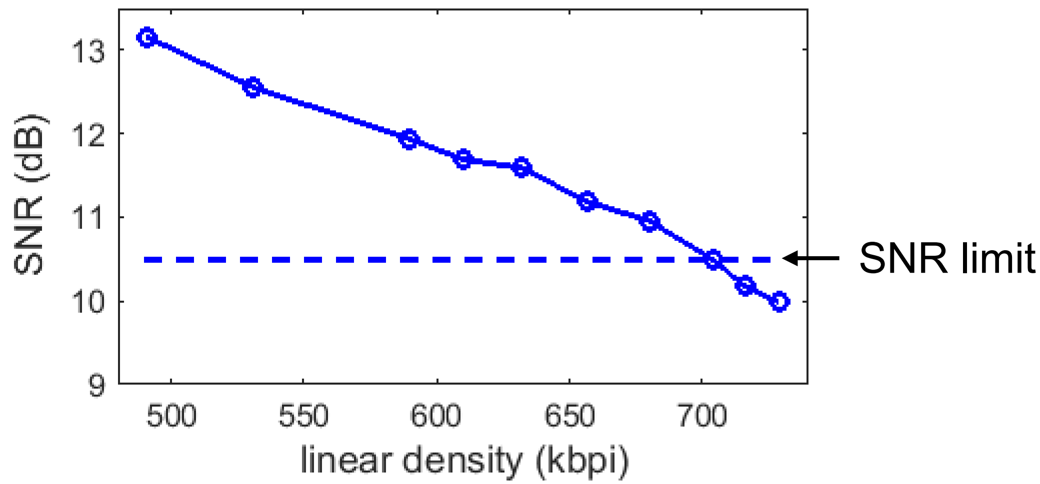

SEM image of GMR reader

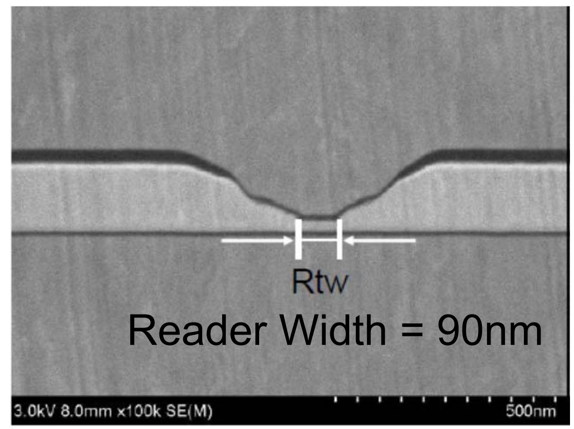

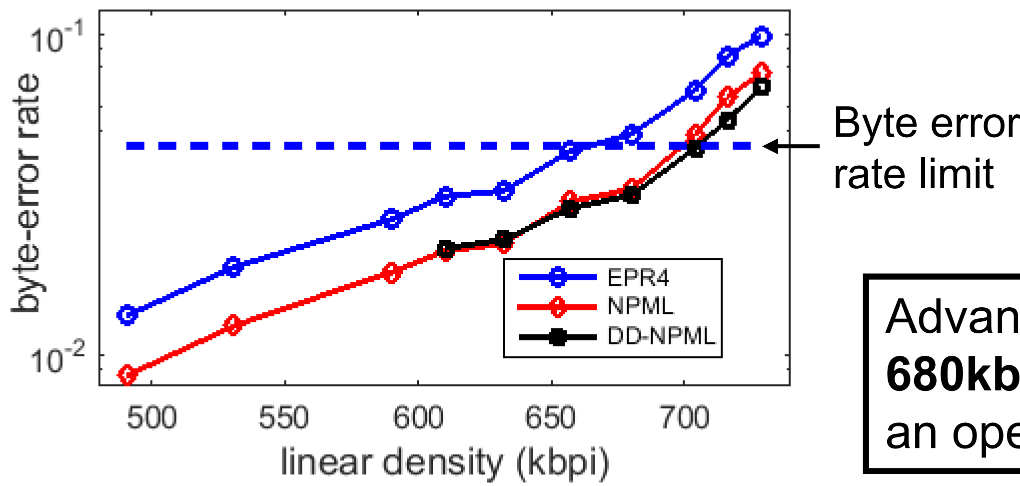

Advanced BaFe supports a linear density of **680kbpi** with a 90nm reader and provides an operating margin of ~ 0.5dB SNR

Summary of demo results

- Advanced Perpendicular BaFe medium
- Linear density = 680 kbpi w/ 90 nm reader (single-channel recording)
- 1-sigma PES = 5.9 nm,
- Track density = 181 ktpi (track width = 140 nm)

Areal recording density : 123 Gb/in^2

61x LTO6 areal density

→ 220 TB cartridge capacity (*)

This demonstration shows that tape technology has the potential for significant capacity increase for years to come!

(*) 220 TB cartridge capacity, assuming LTO6 format overheads and taking into account the 48% increase in tape length enabled by the thinner Aramid tape substrate used

INSIC 2012-2022 Tape Roadmap

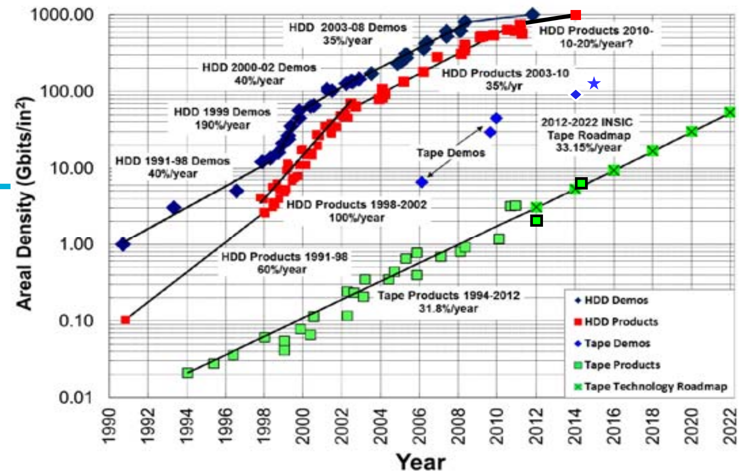

Table 1: 2012 Tape Technology Roadmap Detail.

Parameter/Year	2012	2014	2016	2018	2020	2022		
1. Capacity (TB)	4	8	16	32	64	128	**41.42%**	per year
2. Data rate per channel (MB/sec)	11.2	13.8	17.0	20.9	25.8	31.8	**11.00%**	per year
3. Total data rate (MB/sec)	270.0	405.0	607.4	911.0	1366.5	2049.6	**22.47%**	per year
4. FC Speed Roadmap (MB/sec)*	3200	6400	12800	12800	25600	25600		
5. Number of channels	24	29	36	43	53	64	10.33%	
6. Tape thickness (μm)	**6.00**	**5.53**	**5.10**	**4.70**	**4.33**	**3.99**	**-4.00%**	per year
7. Data capacity reserve	5.0%	5.0%	5.0%	5.0%	5.0%	5.0%		defect reserve
8. Tape length that is recordable (meters)	867	941	1,021	1,107	1,202	1,304	**3.90%**	winding reserve
9. Tape length total (meters)	902	979	1,062	1,152	1,250	1,357	4.17%	
10. Track density (TPI)	6,506	9,773	14,787	22,498	34,393	52,791	23.29%	per year
11. Linear bit density (KFCI)**	467	545	635	741	864	1,008	**8.00%**	per year
12. Areal density (Gbits/inch^2)	3.04	5.32	9.39	16.67	29.73	53.22	33.15%	per year
13. Tape speed (m/sec)	6.4	6.6	6.8	7.0	7.3	7.5	1.63%	
14. Tape width (mm)	**12.65**	**12.65**	**12.65**	**12.65**	**12.65**	**12.65**		
15. ECC and formatting overhead	24.00%	21.84%	19.88%	18.09%	16.47%	14.99%	**-4.60%**	per year
16. Servo track and layout overhead***	18.46%	16.59%	15.02%	13.67%	12.50%	11.49%	**-7.00%**	per year
17. Number of passes to write a tape	110	138	175	222	283	361		
18. Time to fill a tape (minutes)	247	329	439	585	781	1,041	15.47%	
19. Number of passes to end of media life	30,300	33,406	36,830	40,605	44,767	49,356	**5.00%**	per year
20. Number of data tracks	2,642	4,060	6,258	9,673	14,988	23,272	24.31%	per year
21. Bit Aspect Ratio (BAR)	88	67	51	38	29	22	-13.12%	per year

INSIC Roadmap available at: http://www.insic.org/news/2012Roadmap/news_12index.html

Fujifilm’s 7th Annual Global IT Executive Summit

October 7 – 10, 2015

What About the Other Storage Technologies?

Ed Childers
STSM, Manager Tape and LTFS
IBM Storage Development, Tucson, AZ

Excerpts from IBM EDGE May 11-15, 2015

How do you store a bit?

Magnetize something → Tape & HDD

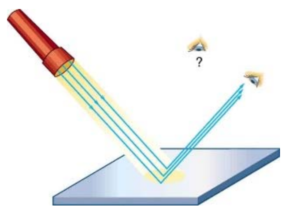

Change optical reflection → BluRay, DVD

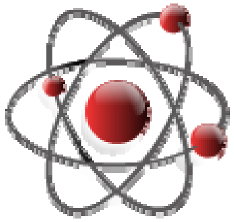

Capture charge → Flash

What's limiting ability to increase areal density?

Tape

- ➢ No limitations in multi-decade time frame
  - ➢ Size of elements >100x than HDD
    - ➢ No lithographic challenges
  - ➢ Magnetic domain stability not an issue
    - ➢ No fundamentals of physics issues

HDD

- ➢ Size of magnetic domains
  - Stability of magnetic domains @ room temp
  - Long awaited superparametric limit is here
- ➢ Lithographic capability to make heads smaller
  - Approaching limits / reaching chip requirements

Flash

- ➢ Captured charge / # of electrons in level
- ➢ Longer term – aspect ratio's of structures
- ➢ Longer term – 3D eventually

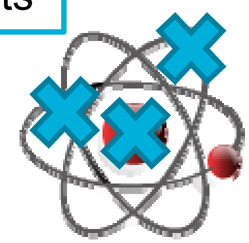

Optical

- ➢ Reducing wavelength of light in semiconductor laser, there is no viable path
  - ➢ ie. nothing after Blu – no roadmap

HDD Areal Density Scaling:

HDD utilizes equipment developed for chip industry Feature size approaching that of silicon

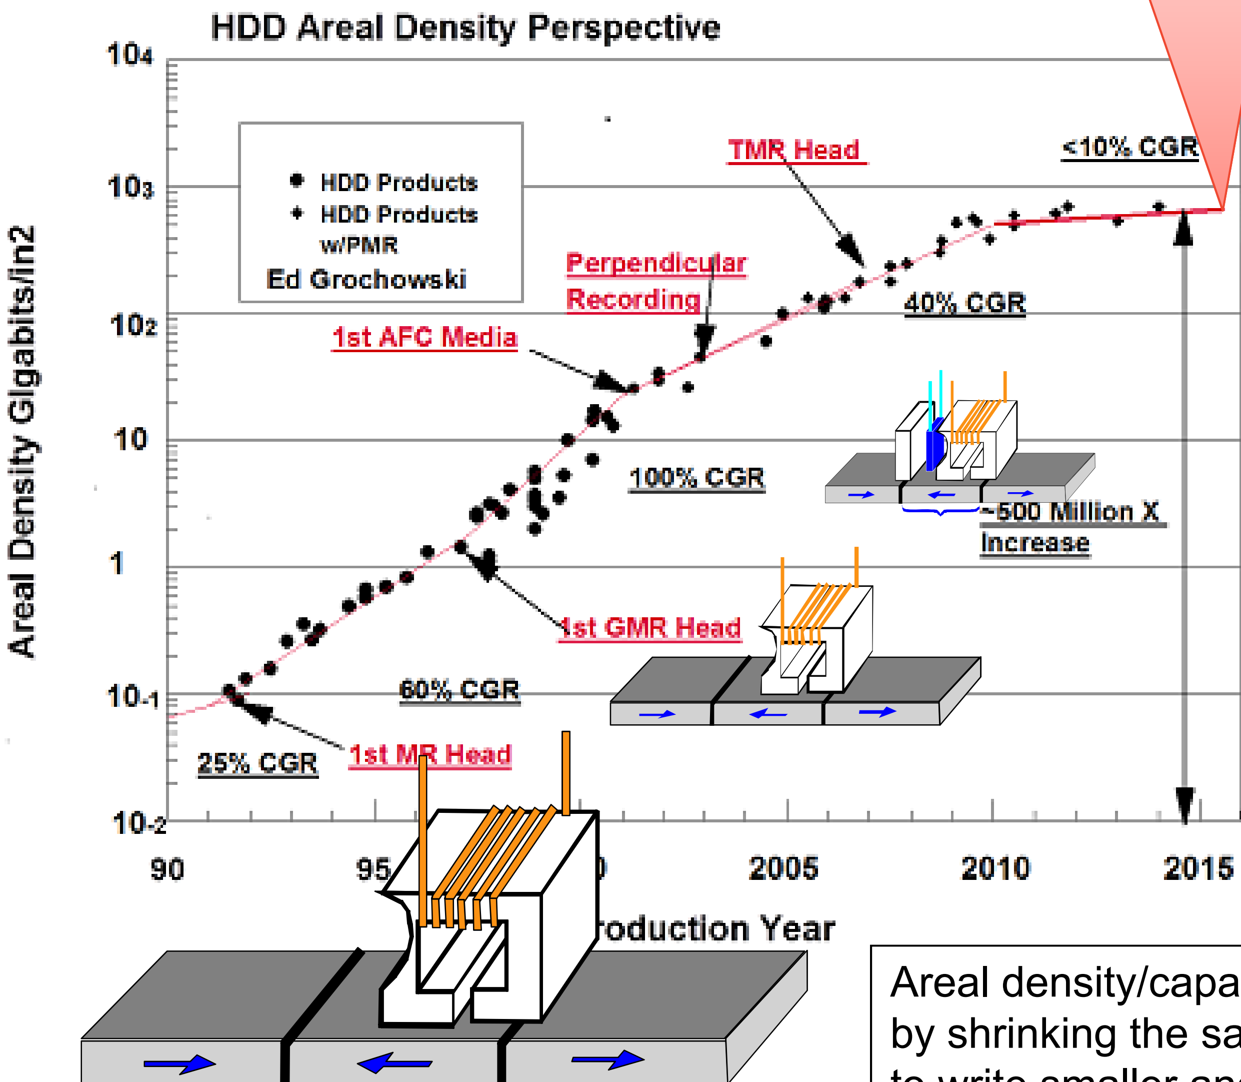

HAMR or patterned media required to advance. Both difficult technology transitions

→ HDD scaling will continue to be slow until at least mid 2018

Areal density/capacity scaling achieved by shrinking the same basic technology to write smaller and smaller bits on disk

From Mark Lantz – Future of Tape presentation

Ref: http://www.storageacceleration.com/author.asp?section_id=3670&doc_id=274482

Edge2015
#ibmedge

Optical Roadmap – A Squeeze Play

		CD-ROM	DVD-ROM	DVD+R	BluRay	BluRay	BluRay	BluRay	BluRay
		1982	1995	2004	2006	2014	2015		
		ECMA-130	ECMA-267	ECMA-364	BD-ROM	BD-R	BD-R	BD-R	BD-R
GB		0.682	4.70	8.55	25	100	300	500	1000
Layers		1	1	2	1	3	6	6	6
Capacity / Layer		0.682	4.70	4.27	25.00	33.33	50.00	83.33	166.67
Areal Density (User)	Gb/in^2	0.41	2.77	2.52	14.73	19.64	29.46	49.11	98.21
Areal Density (Raw)	Gb/in^2	1.44	6.54	5.95	26.72	25.23	37.85	37.85	37.85
			Red		Blue				
Capacity (per layer)	factor		6.89	0.91	5.32	1.33	1.50	1.67	2.00
Optical	factor		2.35	1.00	5.17	1.00	1.00	1.00	1.00
Track Squeeze	factor		1.41	1.00	1.02	1.00	1.42	1.00	1.00
Bit Squeeze	factor		1.37	0.91	0.78	0.94	1.05	1.00	1.00
Channel Eff	factor		1.09	1.00	1.33	1.41	1.00	1.67	2.00
Format Eff	factor		1.37	1.00	0.98	1.00	1.00	1.00	1.00
Disk Area Eff	factor		1.02	1.00	1.00	1.00	1.00	1.00	1.00
Total	factor		6.89	0.91	5.36	1.33	1.50	1.67	2.00

Smaller Wavelength Laser

Can not be scaled continually

Spot squeeze (ie. areal density increase not due to smaller spot size or improved format efficiency

Less channel overhead

PRML channel & modulation, 6% loss in LD due to layers

Elimination of guard band on write, 19% of spot overwriting adjacent trk, cancellation technology

Improved Channel?

Multi-state Recording

Storage Technologies Areal Density Trends

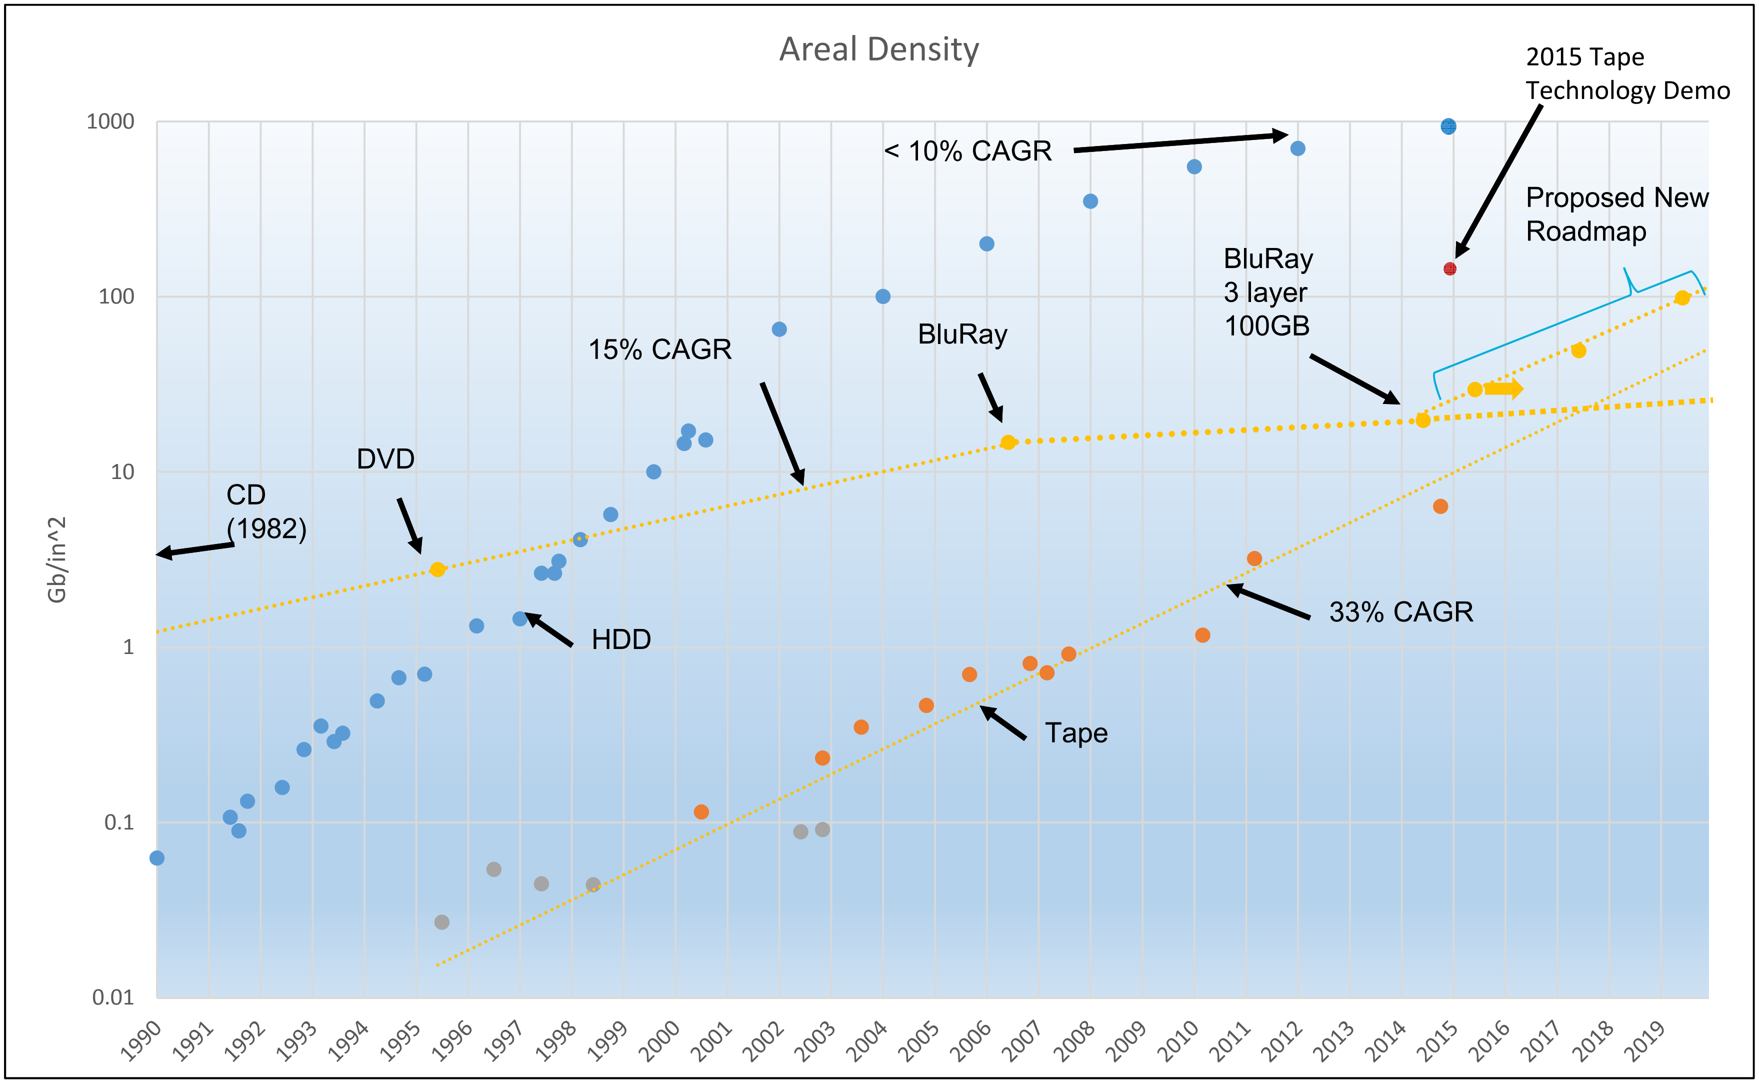

Engineers Solve Technology Problems but…

Market forces drive business models

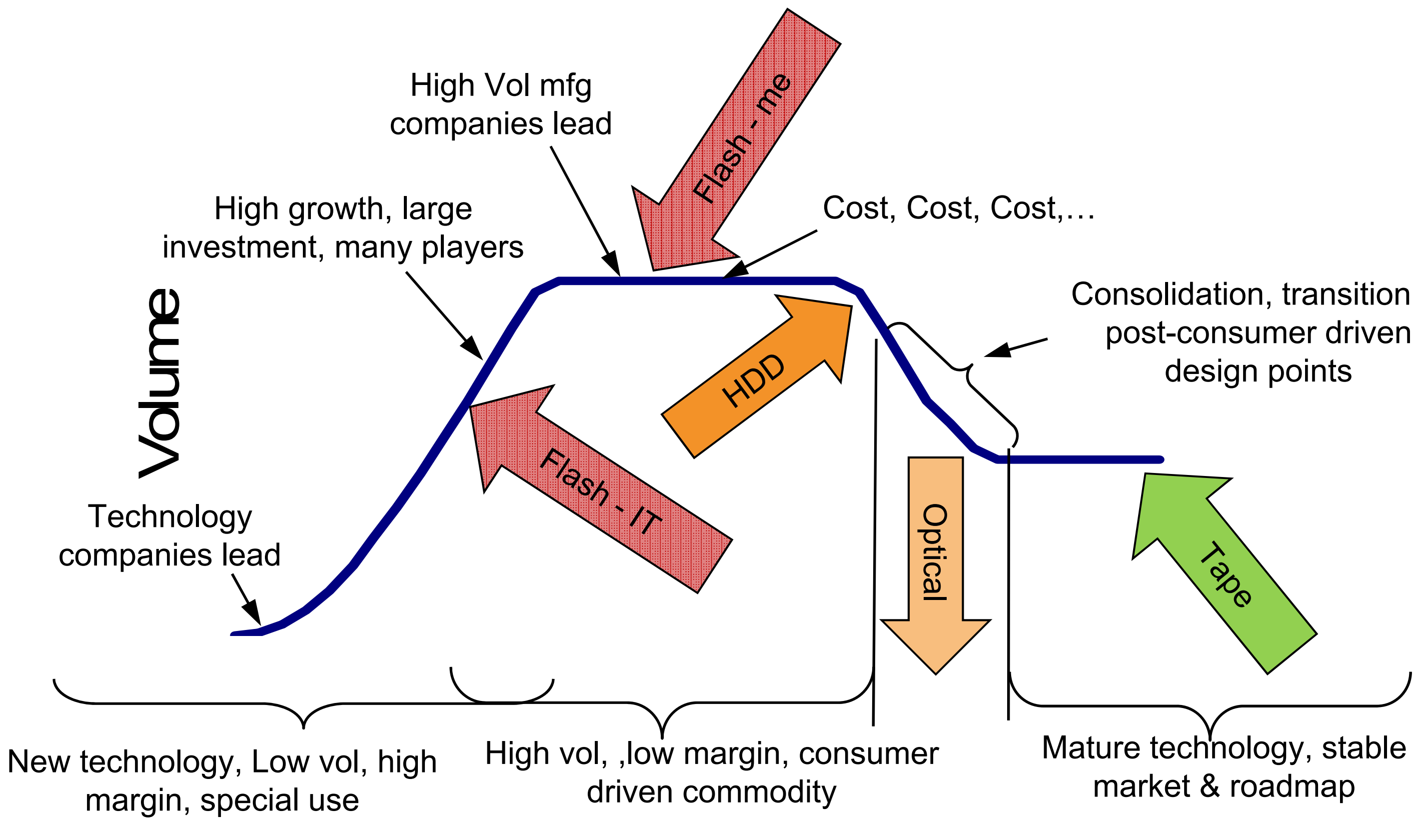

Fujifilm's 7th Annual Global IT Executive Summit

October 7 – 10, 2015

Investing in the Future

Ed Childers
STSM, Manager Tape and LTFS
IBM Storage Development, Tucson, AZ

Excerpts from IBM EDGE May 11-15, 2015

How do we capitalize on the technology?

- The Problem
  - (ie. where to aim)
- What's Required?
  - (in what do we need invest?)
- What's the return?
  - (it's about $)

Session objectives

- The Problem
- Data use shift
  - Data is an asset
    - It needs to be accessible to have value
- Storage Market Transition / Disruption
  - Traditional Scaling breaking down
    - Consumer Volume Shifts

- What's Required?

- The Bottom Line

Data was traditionally a tool of the business

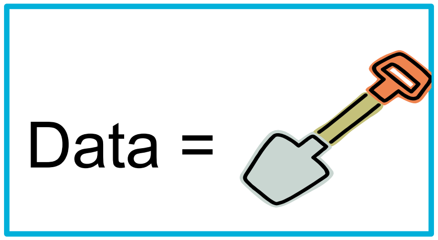

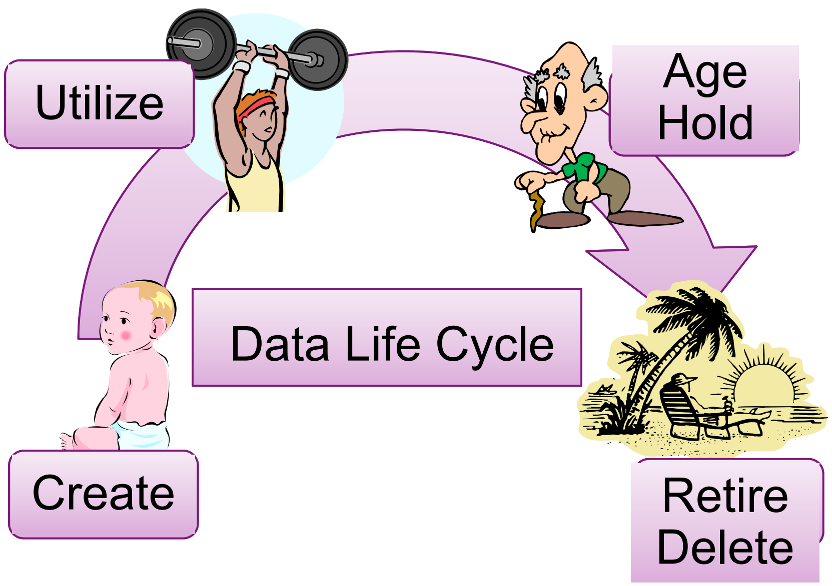

Now, Data *is* the business

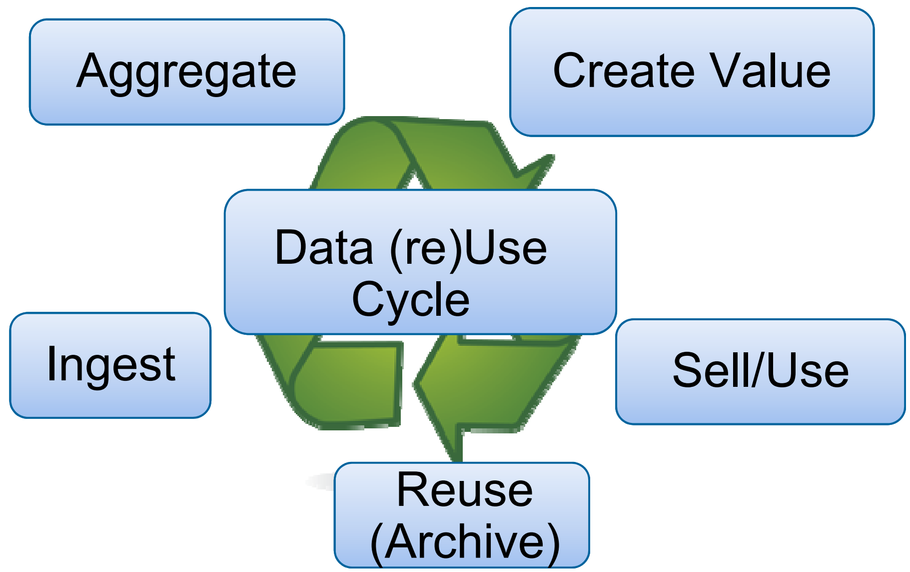

Demands on IT, Storage, and Shift to Cloud

Underneath, more to the point

Fundamentally – It's a Storage Scaling Problem

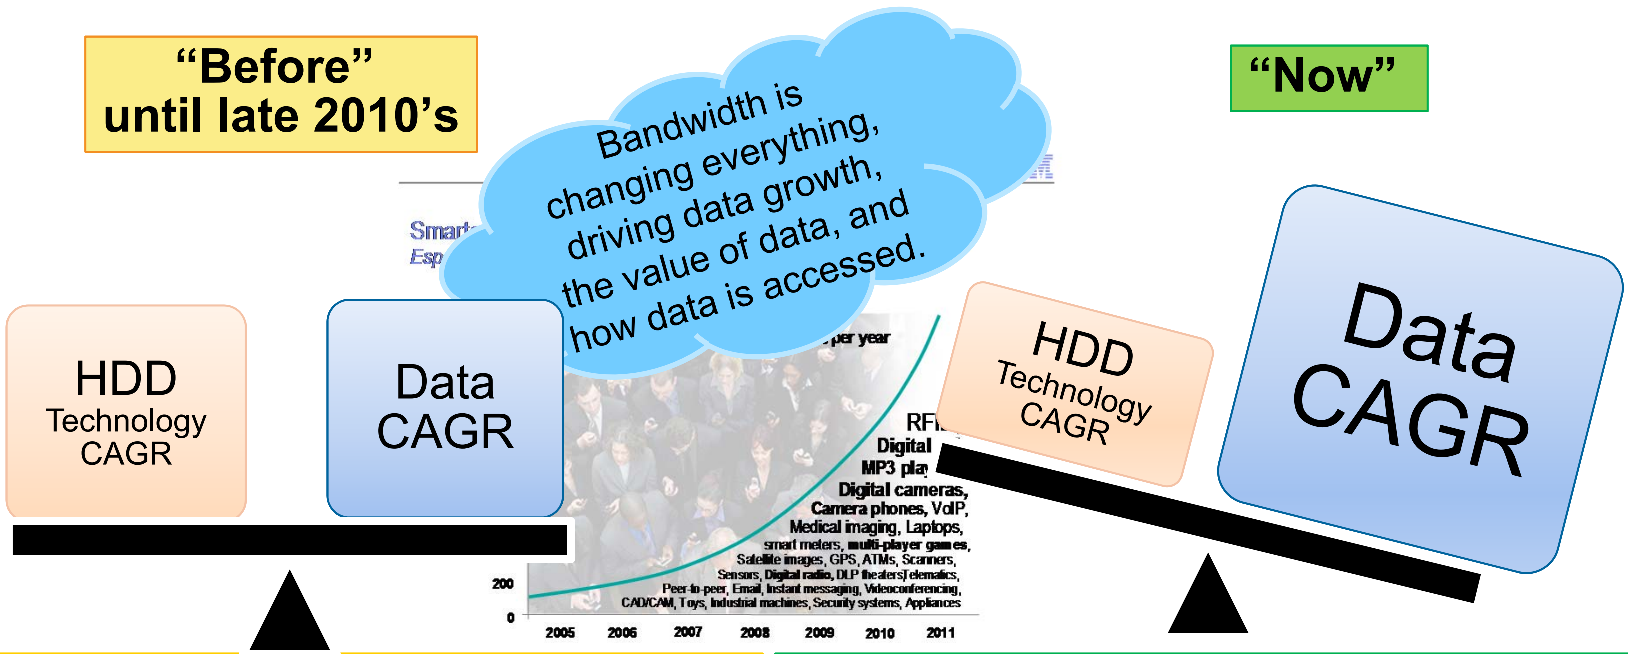

The rate of data growth approx the same as:
- ➢ advances in HDD technology areal density (ie. double every 2 years = 40% CAGR)

The rate of data growth is greater than:
- ➢ advances in HDD technology areal density

(No increases in IT spend to offset imbalance)

Session objectives

- The Problem

- What's Required?
  - Investments in:
    - ✓ Tape Technology
    - ✓ Tape Product Delivery
    - Software Stack
      - LTFS

- The Bottom Line

IBM Tape Investment Strategy

Edge2015
#ibmedge

- **Tape Technology Pipeline - IBM Development & IBM Research**
  - Large research investment Zurich and Almaden laboratories
  - Technology demonstrations – 220 TB in 2015
    - ➢Substantiates roadmap thru 2025
  - Drive technology value into Enterprise, leverage into LTO when industry is capable
- **TS1150 Enterprise Tape product line**
  - Reliability, Performance and Function differentiation
  - Enterprise media cartridge ***with reuse***
  - Enterprise Automation compatible
- **LTO Midrange product line**
  - Open Tape Streaming product family
  - Full Automation Product support – 1U to HD Frames
  - TPC Consortium driven development/function
- **Software - LTFS**
  - IBM invented, open source, open standard
  - Provides file system support, integration of tape into Spectrum Scale

Announcing Availability of LTO Gen 7
Oct 6 2015

IBM – Tape Product Innovations

124 Patents Issued
2014

Edge2015
#ibmedge

- **Timing-Based Servo and Surface control guiding**
  - TBS revolutionized tape track following with 10X precision of analog patterns
  - Surface control guiding reduces frequency of tape guiding transients for robust track following
- **Compact tape path with flat head geometry contour and flangeless rollers**
  - Flat head design reduces friction – head/media wear improved
  - Flangeless tape path eliminated edge contact with the tape – improved media life
  - This enables short tape path – reduced complexity, friction, and improved access
- **GMR head technology with protective overcoat**
  - 3X signal amplitude enabling track density and longer head life
  - TMR on way in future
- **Multi-level ECC with NPML(LTO) and DD-NPML(TS1140) channels**
  - Multi-level Reed-Solomon Error Correction codes with End-to-End CRC
  - Tape format interleaves data both vertically and longitudinally on the media
  - Advanced self-optimizing channel technology for maximum readback reliability
  - Full on-the-fly decryption/decompression
- **Barium Ferrite Tape Media Collaboration**
  - Collaboration with Fujifilm on media technology advancements

Edge2015 #ibmedge

Spectrum Storage – Unlocking Data, Enabling Lower Costs

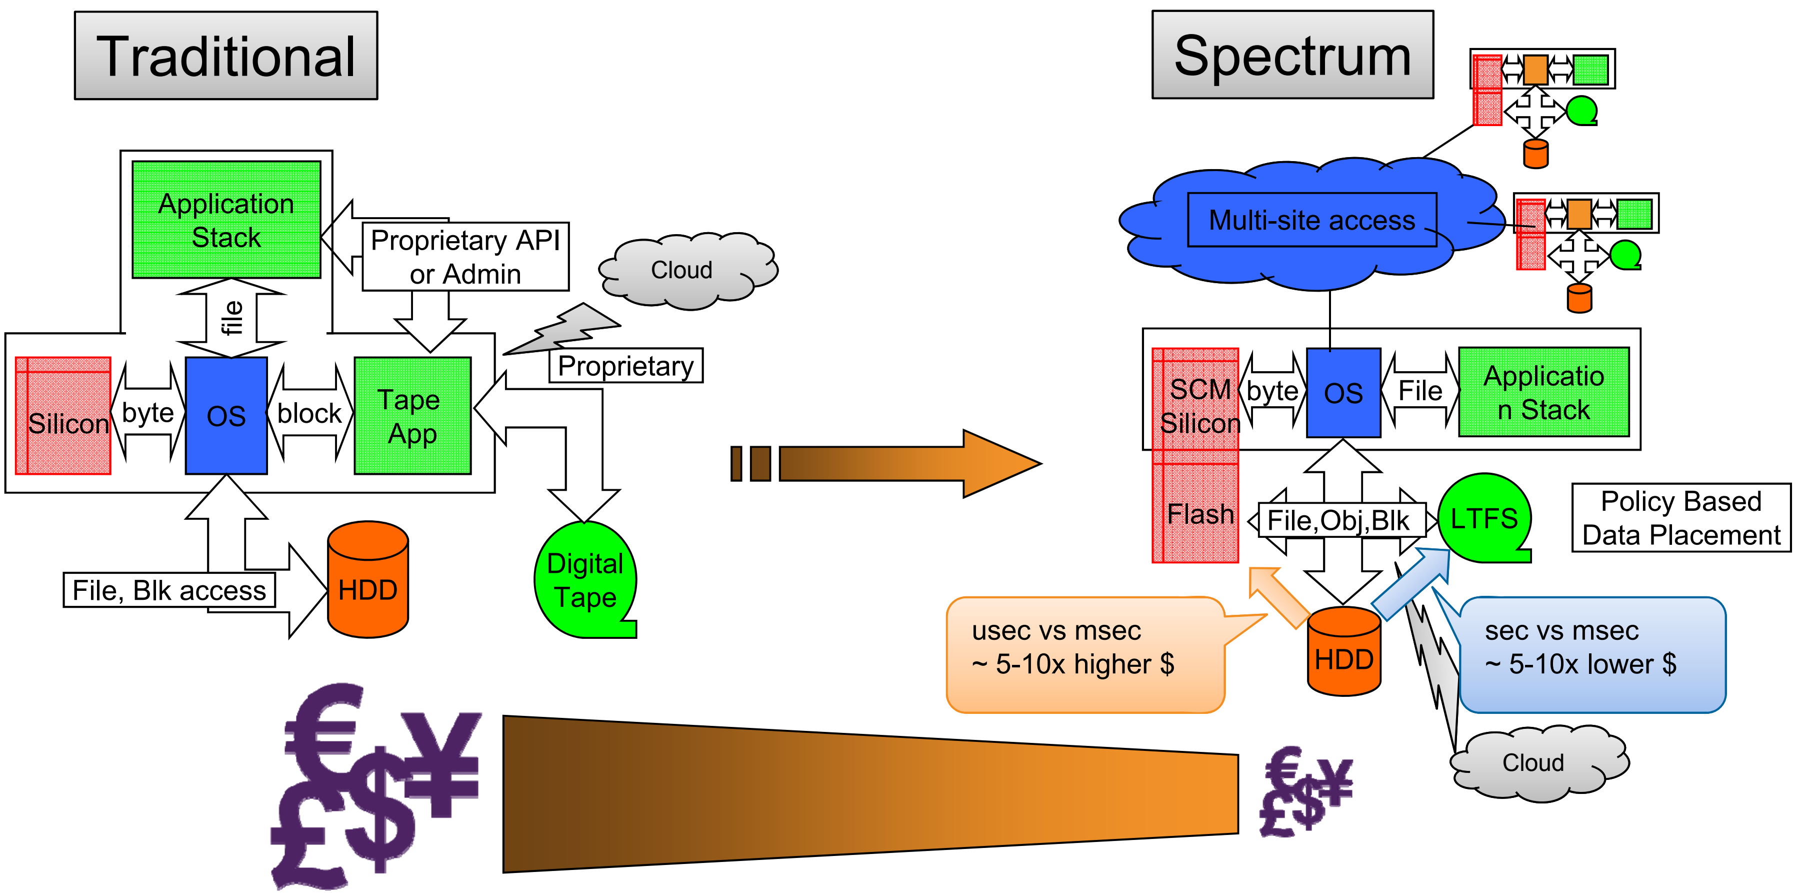

Spectrum Storage

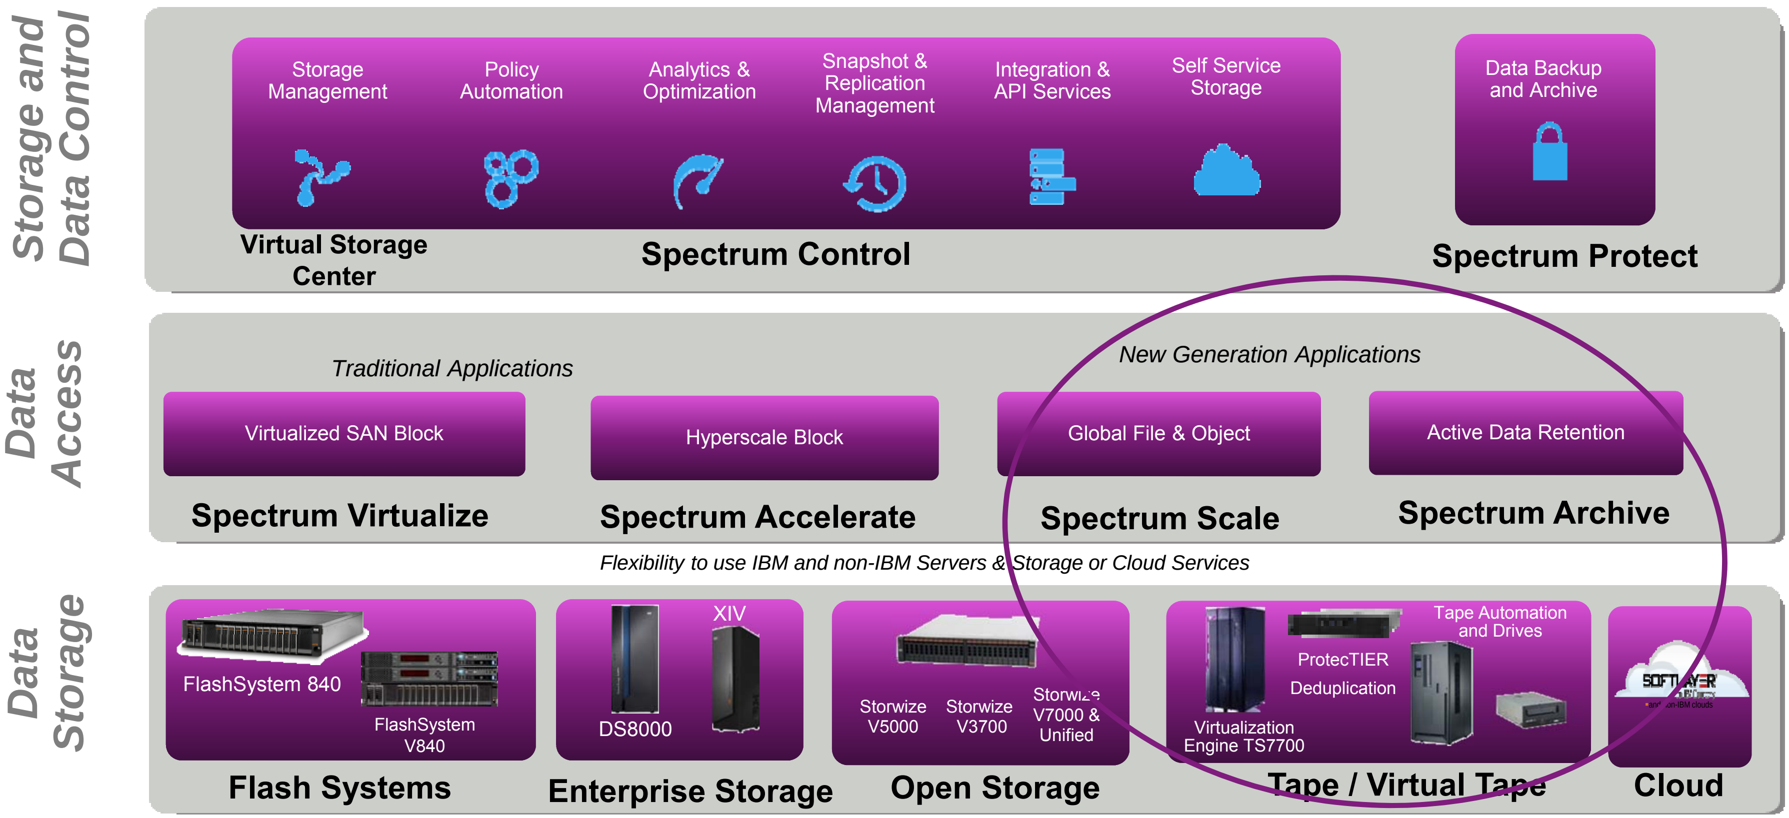

Basic LTFS

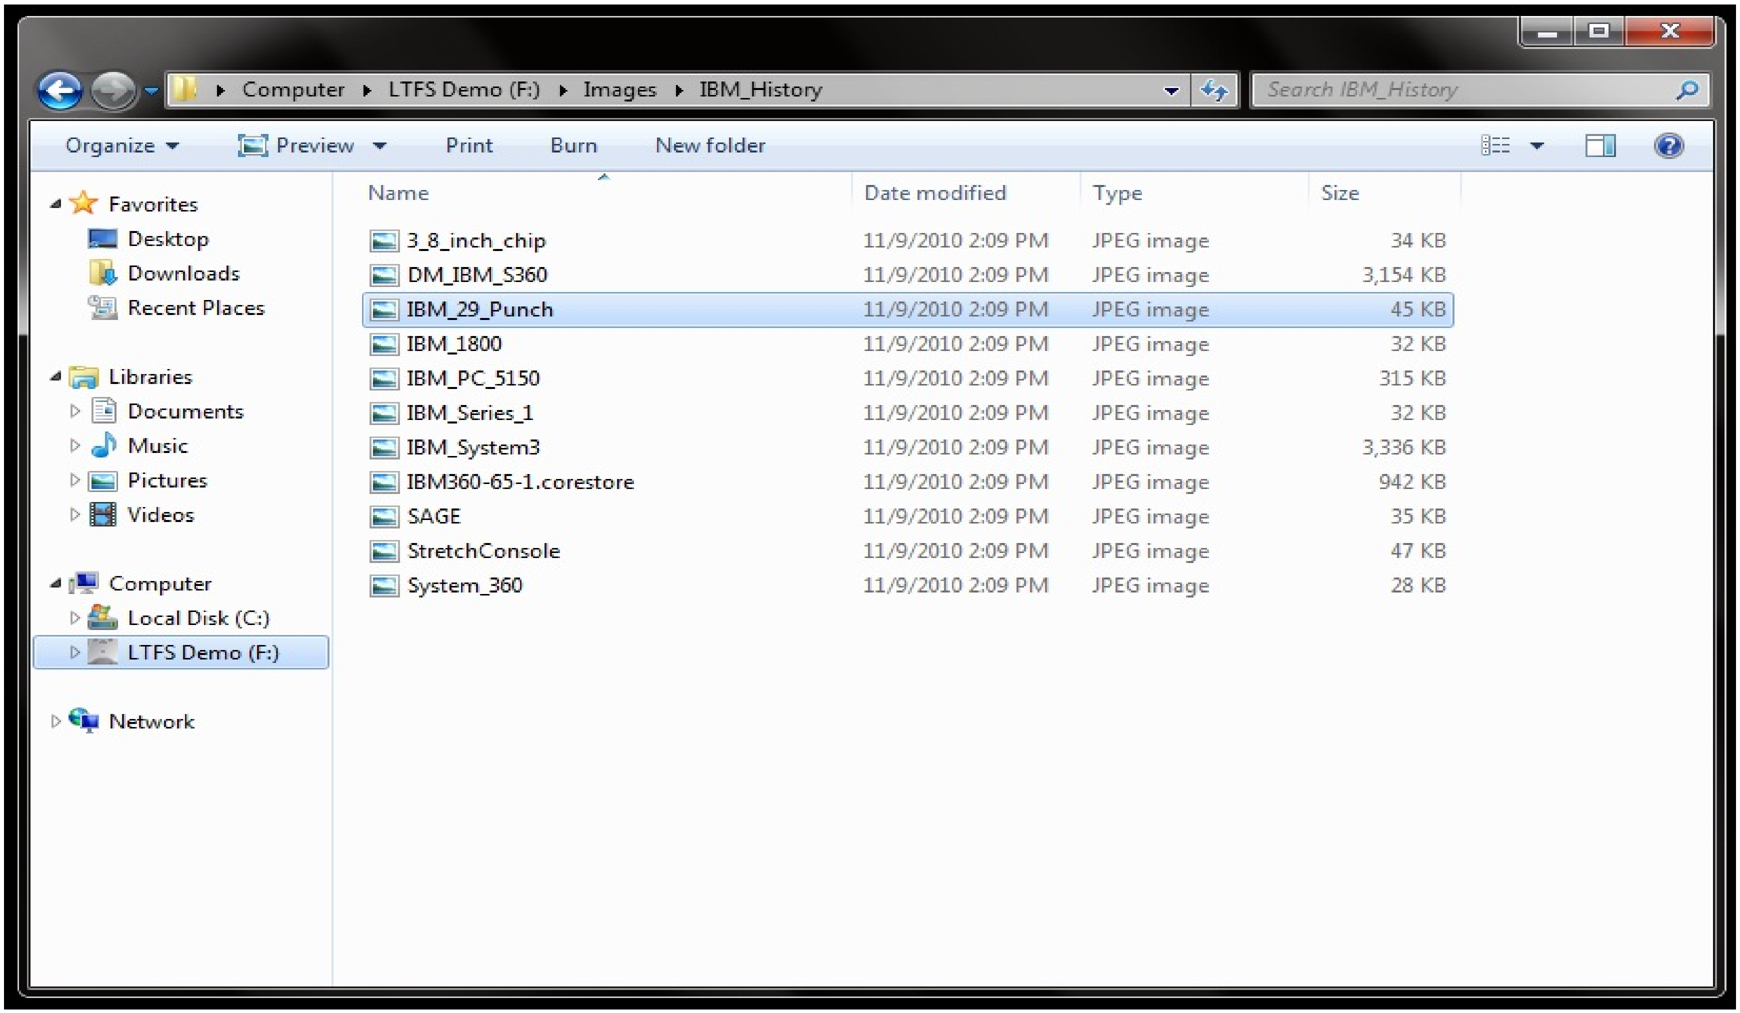
Computer ▸ LTFS Demo (F:) ▸ Images ▸ IBM_History
Search IBM_History
Organize
Preview
Print
Burn
New folder
Favorites
Desktop
Downloads
Recent Places
Libraries
Documents
Music
Pictures
Videos
Computer
Local Disk (C:)
LTFS Demo (F:)
Network
Name
Date modified
Type
Size
3_8_inch_chip 11/9/2010 2:09 PM JPEG image 34 KB
DM_IBM_S360 11/9/2010 2:09 PM JPEG image 3,154 KB
IBM_29_Punch 11/9/2010 2:09 PM JPEG image 45 KB
IBM_1800 11/9/2010 2:09 PM JPEG image 32 KB
IBM_PC_5150 11/9/2010 2:09 PM JPEG image 315 KB
IBM_Series_1 11/9/2010 2:09 PM JPEG image 32 KB
IBM_System3 11/9/2010 2:09 PM JPEG image 3,336 KB
IBM360-65-1.corestore 11/9/2010 2:09 PM JPEG image 942 KB
SAGE 11/9/2010 2:09 PM JPEG image 35 KB
StretchConsole 11/9/2010 2:09 PM JPEG image 47 KB
System_360 11/9/2010 2:09 PM JPEG image 28 KB

Spectrum Archive Tiered Operational Storage

Edge2015 #ibmedge

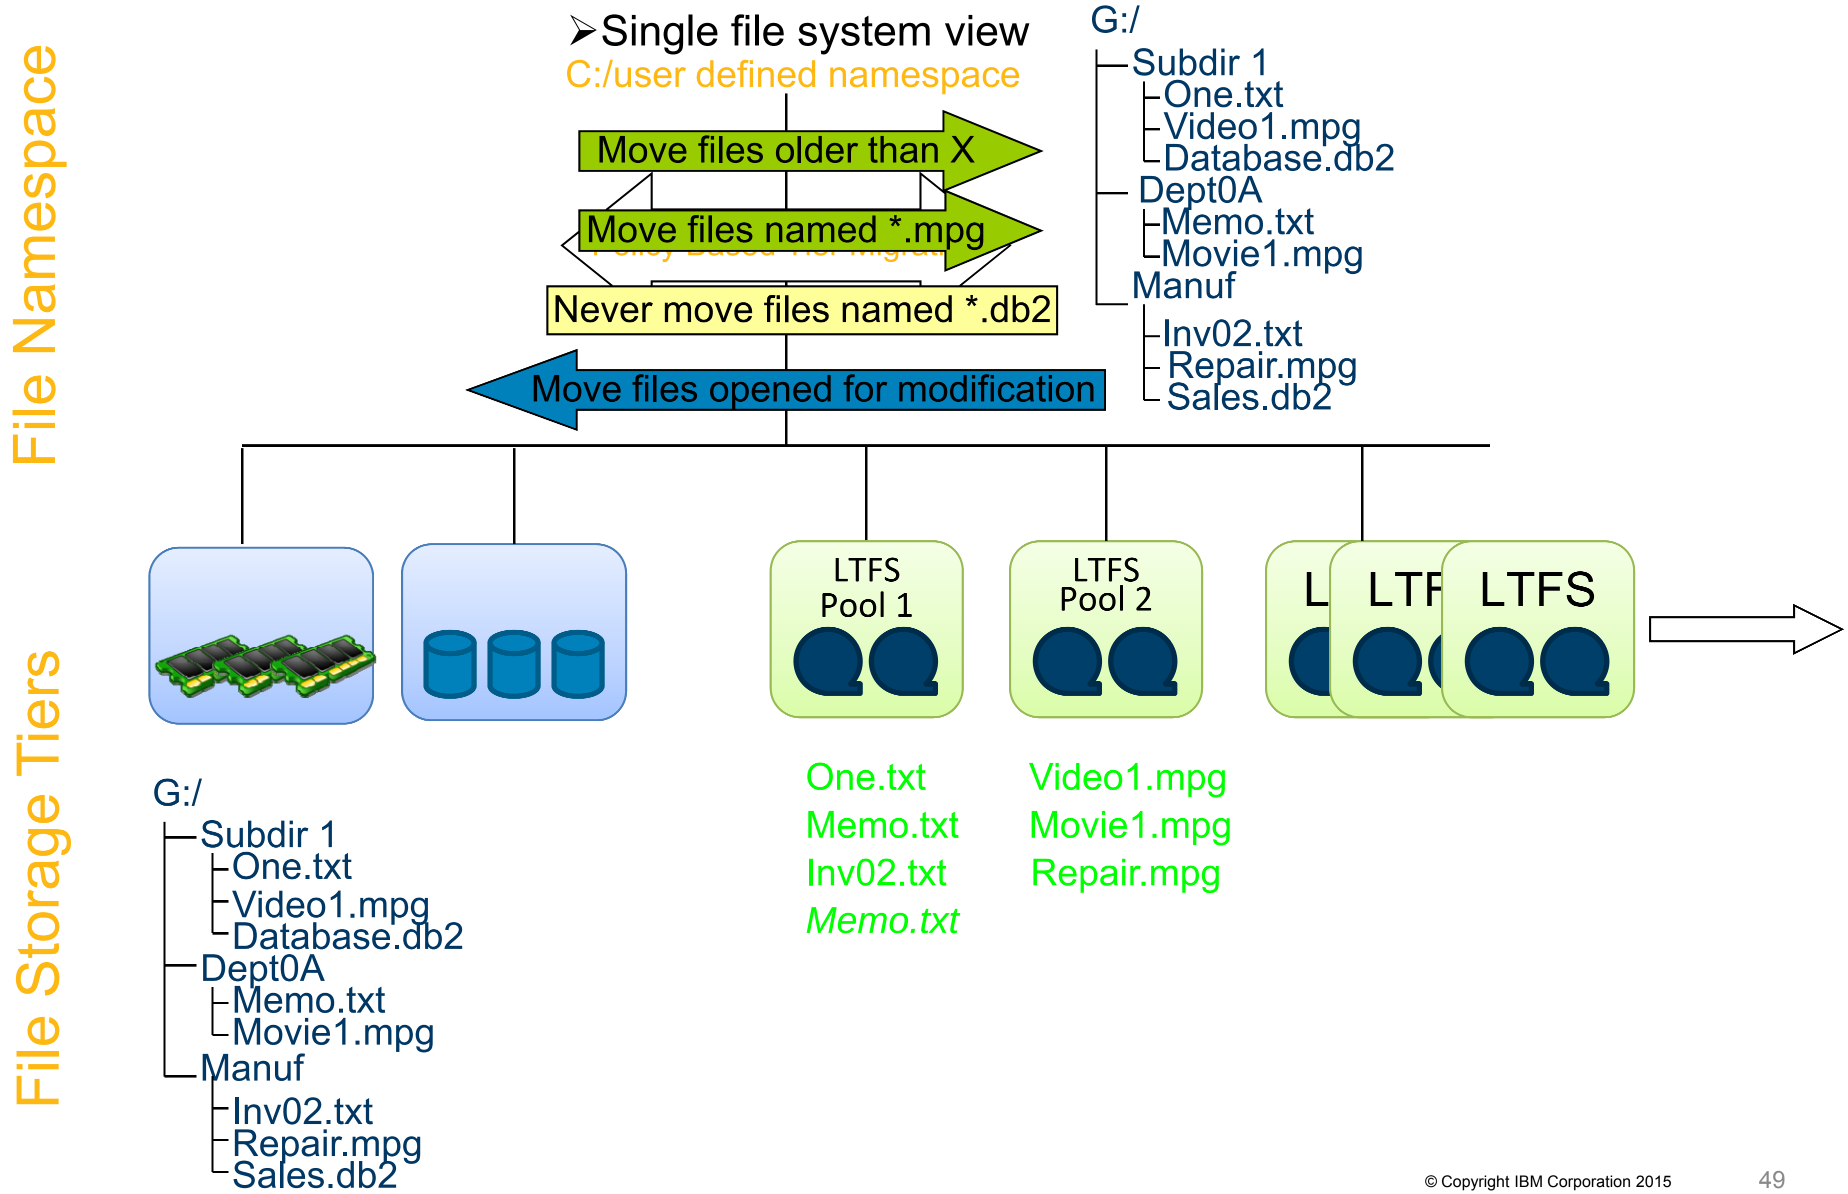

Session objectives

- The Problem

- Why Spectrum Archive?
- Breaking Down Storage Islands

- The Bottom Line
  - Lower costs for cold data while maintaining access

Tape & HDD side by side
Setting up the comparison

Component
$/TB (raw)

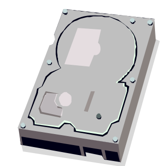

Hardware
$/TB (raw)

System
$/TB/mo (raw)

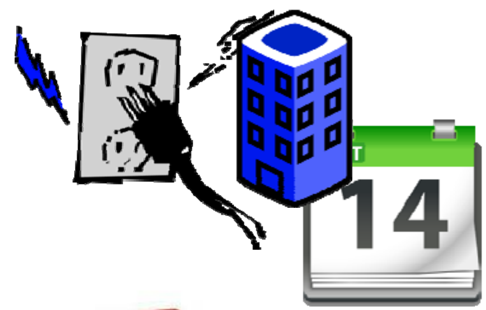

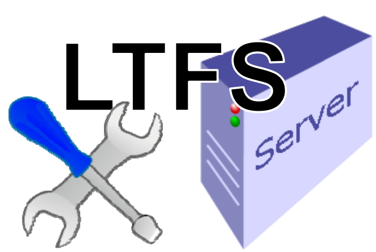

Solution
$/TB/mo (user)

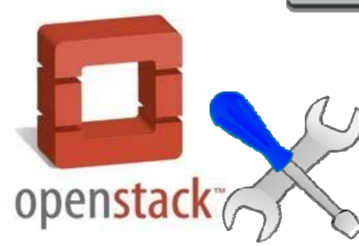

20PB Static Archive

Edge2015
#ibmedge

2.5TB LTO 6 @ $35

6TB HDD @ $235

Component Hardware (raw)

	LTFS	HDD
$/TB	14.00	39.17
$/TB	21.26	63.95

7 frame TS4500
24 drive, 8K carts

74 45 6TB
Backblaze like pods

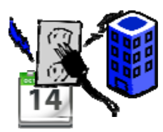

$0.14 /KWH
$200 / sq ft/yr

System (raw)

	LTFS	HDD
3 years $/TB/mo	0.71	2.72
5 years $/TB/mo	0.48	2.00
8 years $/TB/mo	0.35	NA

284 TB/sq ft

66TB/sq ft

+LTFS $ @ 2x redundancy
+0$ SWIFT @ 3x redundancy
$2600/mo admin/maint

Solution (user)

	LTFS	Glacier	HDD
3 years $/TB/mo	1.79	10.00	8.29
5 years $/TB/mo	1.23	10.00	6.13
8 years $/TB/mo	0.91	10.00	NA

Storage Only – no R/W
No connectivity $

20PB Static Archive – Looking Forward

Edge2015 #ibmedge

6.0TB LTO 7 @ $35

12TB HDD @ $235

Component Hardware (raw)

	LTFS	HDD
$/TB	5.83	19.58
$/TB	10.11	32.00

4 frame TS4500
24 drive, 8K carts

37 45 12TB
Backblaze like pods

$0.14 /KWH
$200 / sq ft/yr

System (raw)

	LTFS	HDD
3 years $/TB/mo	0.34	1.38
5 years $/TB/mo	0.22	1.01
8 years $/TB/mo	0.16	NA

660 TB/sq ft

221TB/sq ft

Solution (user)

+LTFS $ @ 2x redundancy
+0$ SWIFT @ 3x redundancy
$2600/mo admin/maint

	LTFS	Glacier	HDD
3 years $/TB/mo	0.92	10.00	4.13
5 years $/TB/mo	0.63	10.00	3.03
8 years $/TB/mo	0.47	10.00	NA

Storage Only – no R/W
No connectivity $

Summary:

- **The era of big data is creating demand for cost effective storage solutions**
- **Tape remains the most cost-efficient and greenest technology for archival storage and active archive applications**
- **Tape has a sustainable roadmap for at least another decade**
  - 123 Gbit/in^2 areal density demo shows feasibility of multiple future tape generations
  - Potential exists for the continued of scaling of tape beyond 123 Gbit/in^2
- **The cost advantage of tape over HDD and optical disk will continue to grow**