
ORACLE®

Storage Technology Futures and Trends

iCAS2015

Dr. Robert M. Raymond
Director, Tape Systems Development
September 17, 2015

ORACLE

Program Agenda

1. Data storage trends
2. Storage Device trends
3. Tape
4. Disk
5. Flash

Data Growth

Driven by Unstructured Data

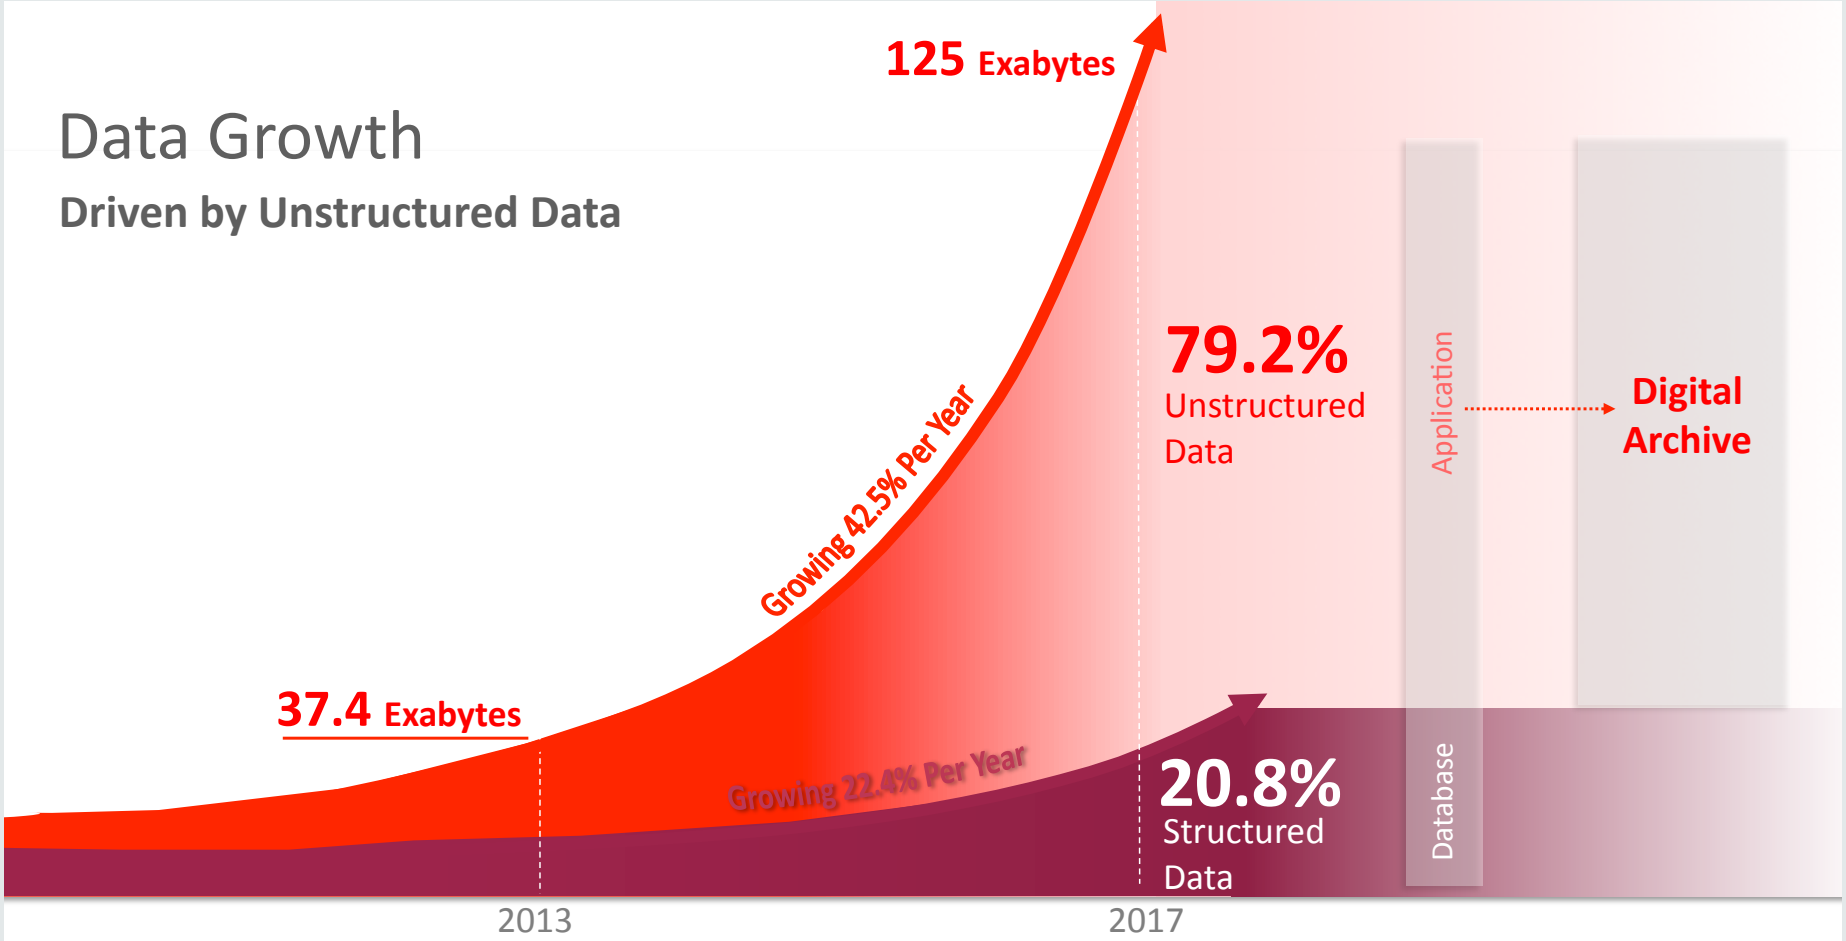

Source: IDC - 2014, Structured Data vs. Unstructured Data: The Balance of Power Continues to Shift

ORACLE

80%

Of Data Rarely Used After 90 Days

The Efficiency of Tiered Storage

Analyst Study: 1 PB Growing at 45% for 9 Years

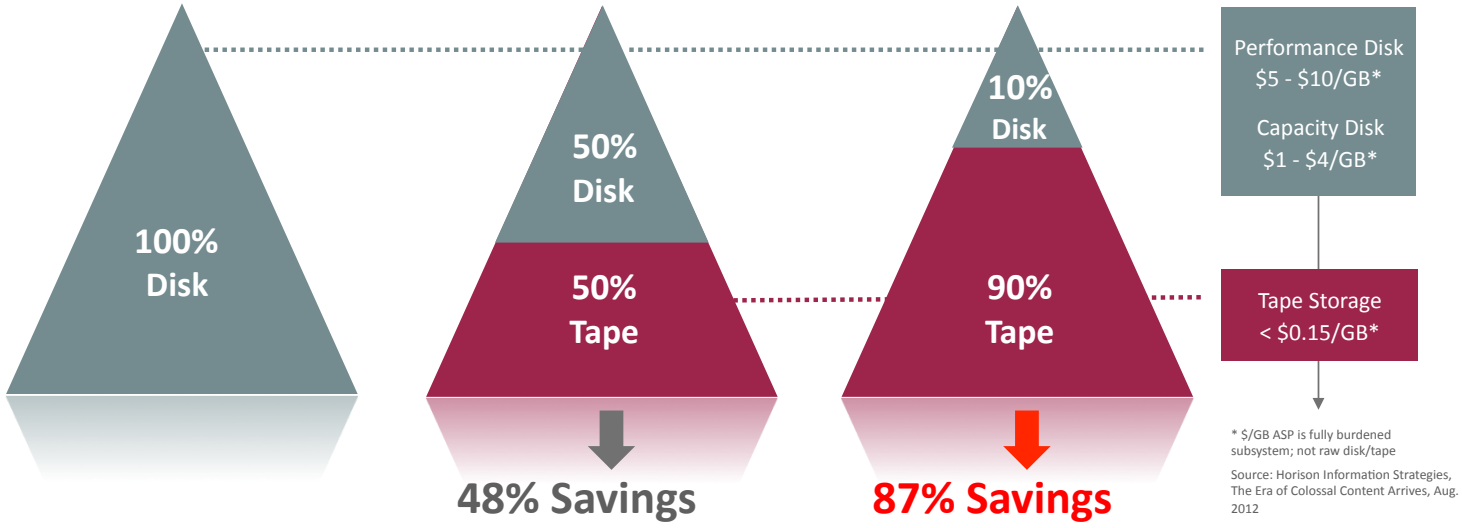

Source: The Clipper Group, Revisiting the Search for Long-Term Storage — A TCO Analysis of Tape and Disk, May 13, 2013

ORACLE

Storage Trends

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Storage Technologies Areal Density Trends

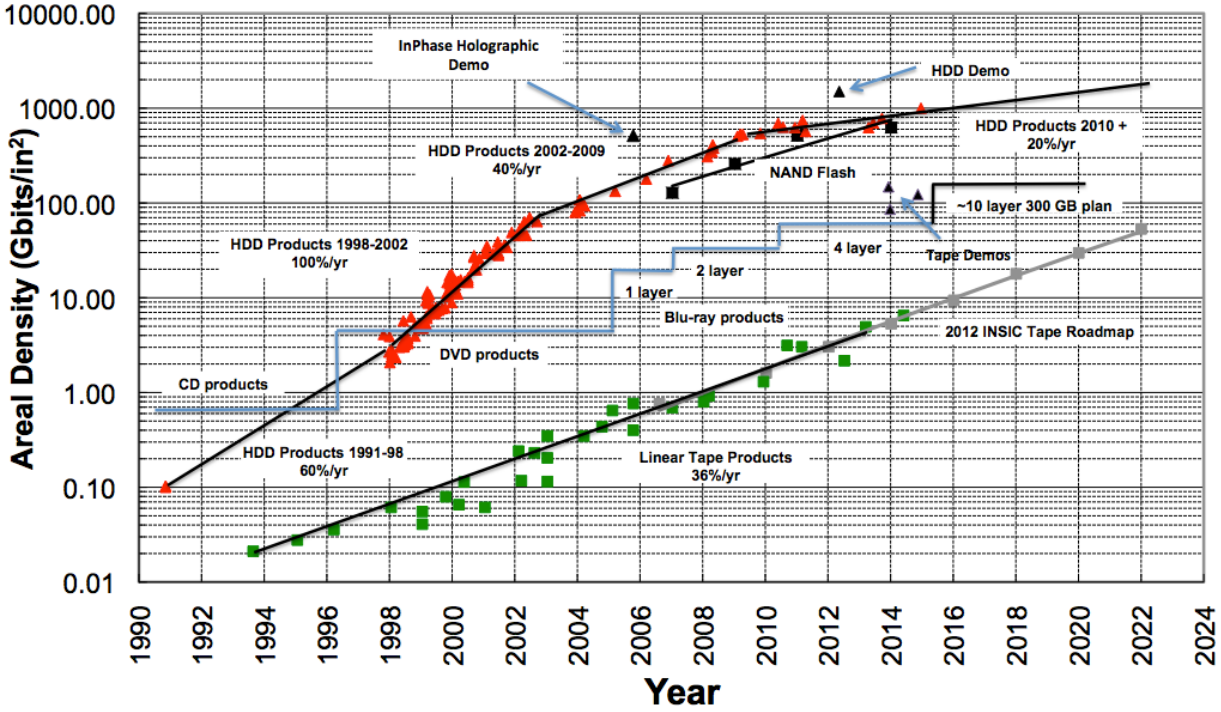

Tape gets its capacity by having 1000X the recording surface area comparing a 1/2 inch cartridge to a 3 1/2 inch disk.

Tape and disk data courtesy of INSIC

Magnetic Recording

Storage Technologies Areal Density Trends

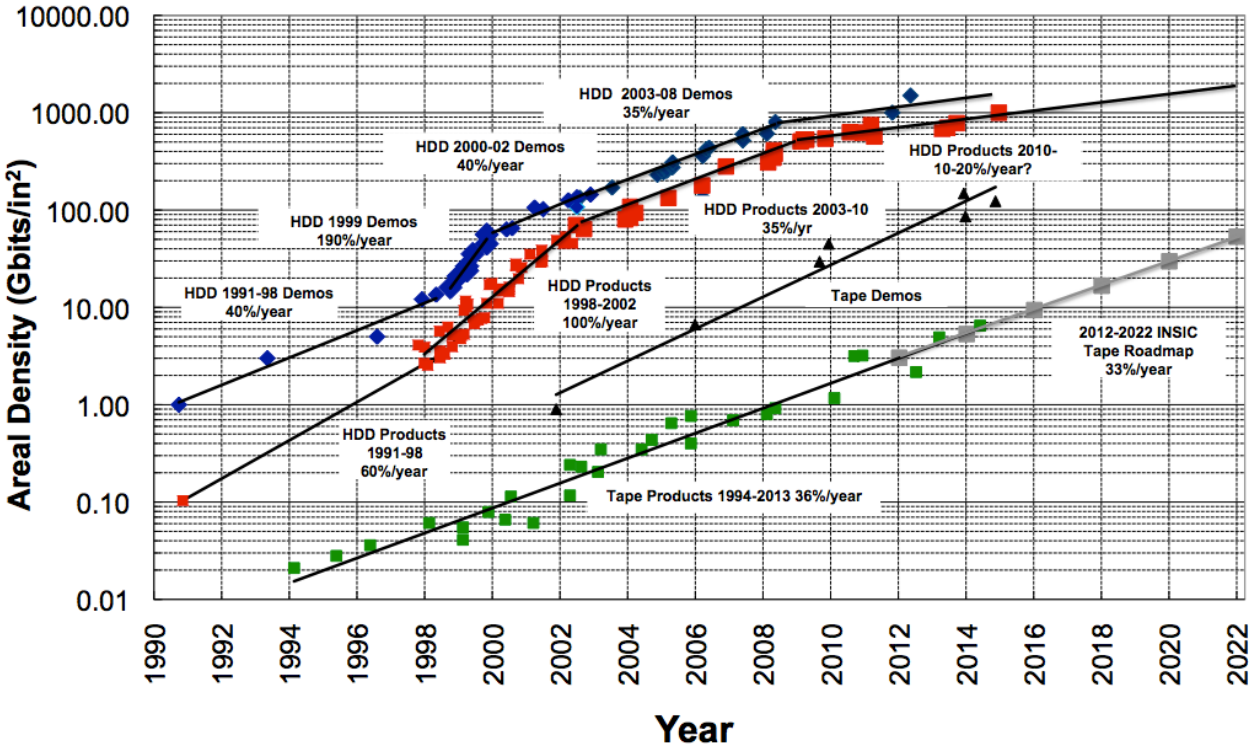

Chart courtesy of INSIC

Tape gets its capacity by having 1000X the recording surface area comparing a 1/2 inch cartridge to a 3 1/2 inch disk.

Magnetic Recording Definitions

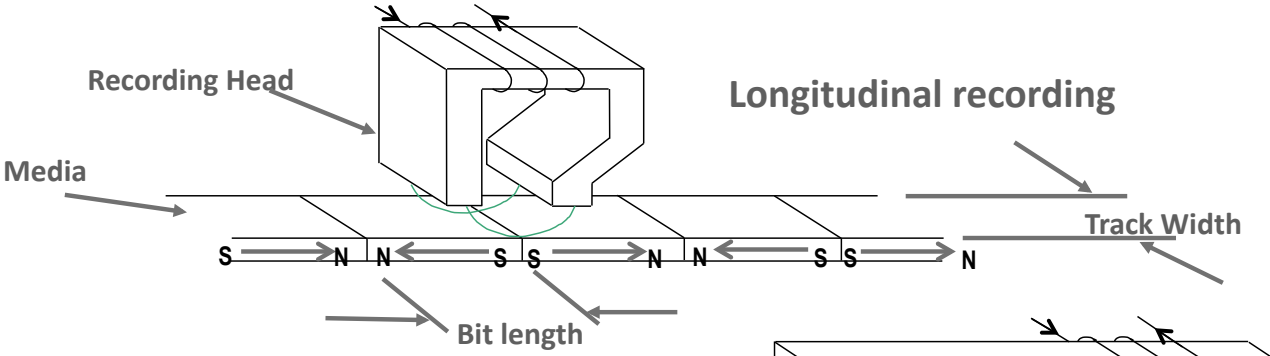

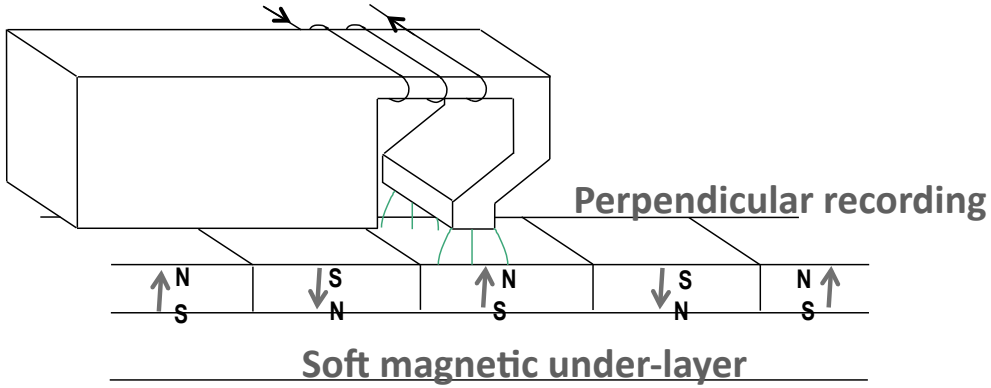

Tracks per inch (tpi) = $(\text{Track Pitch})^{-1}$

Bits per inch (bpi) = $(\text{Bit Length})^{-1}$

Areal density = (Tracks per inch) X (Bits per inch)

Bit aspect ratio (BAR) = (Track width) /(Bit length):1

A Closer Look at the Magnetic Layer (grains or particles)

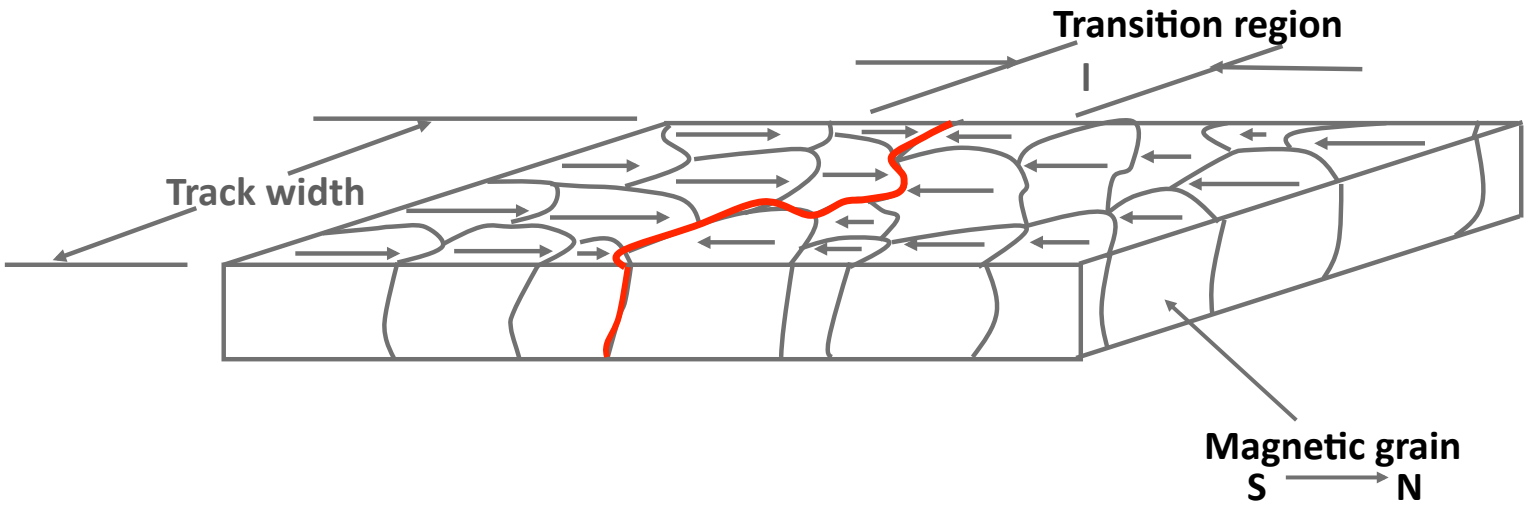

Magnetic Force Microscope (MFM) Track Images

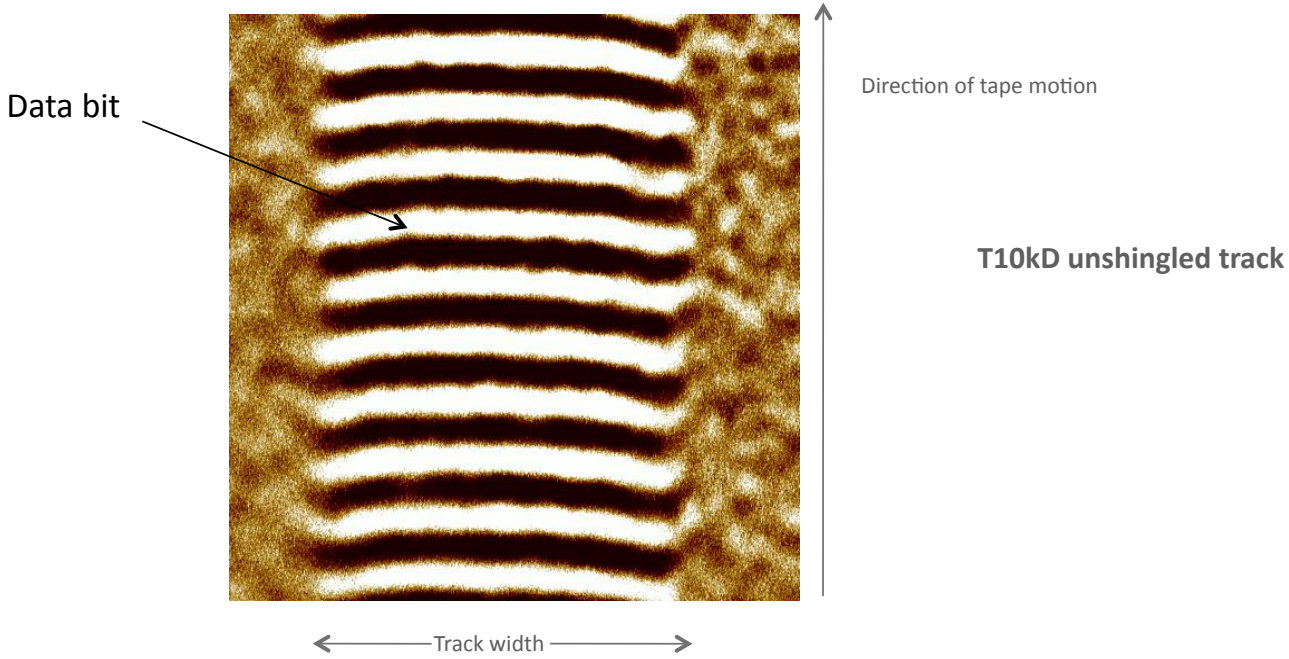

T10kD unshingled track

Data Bit Size Comparison

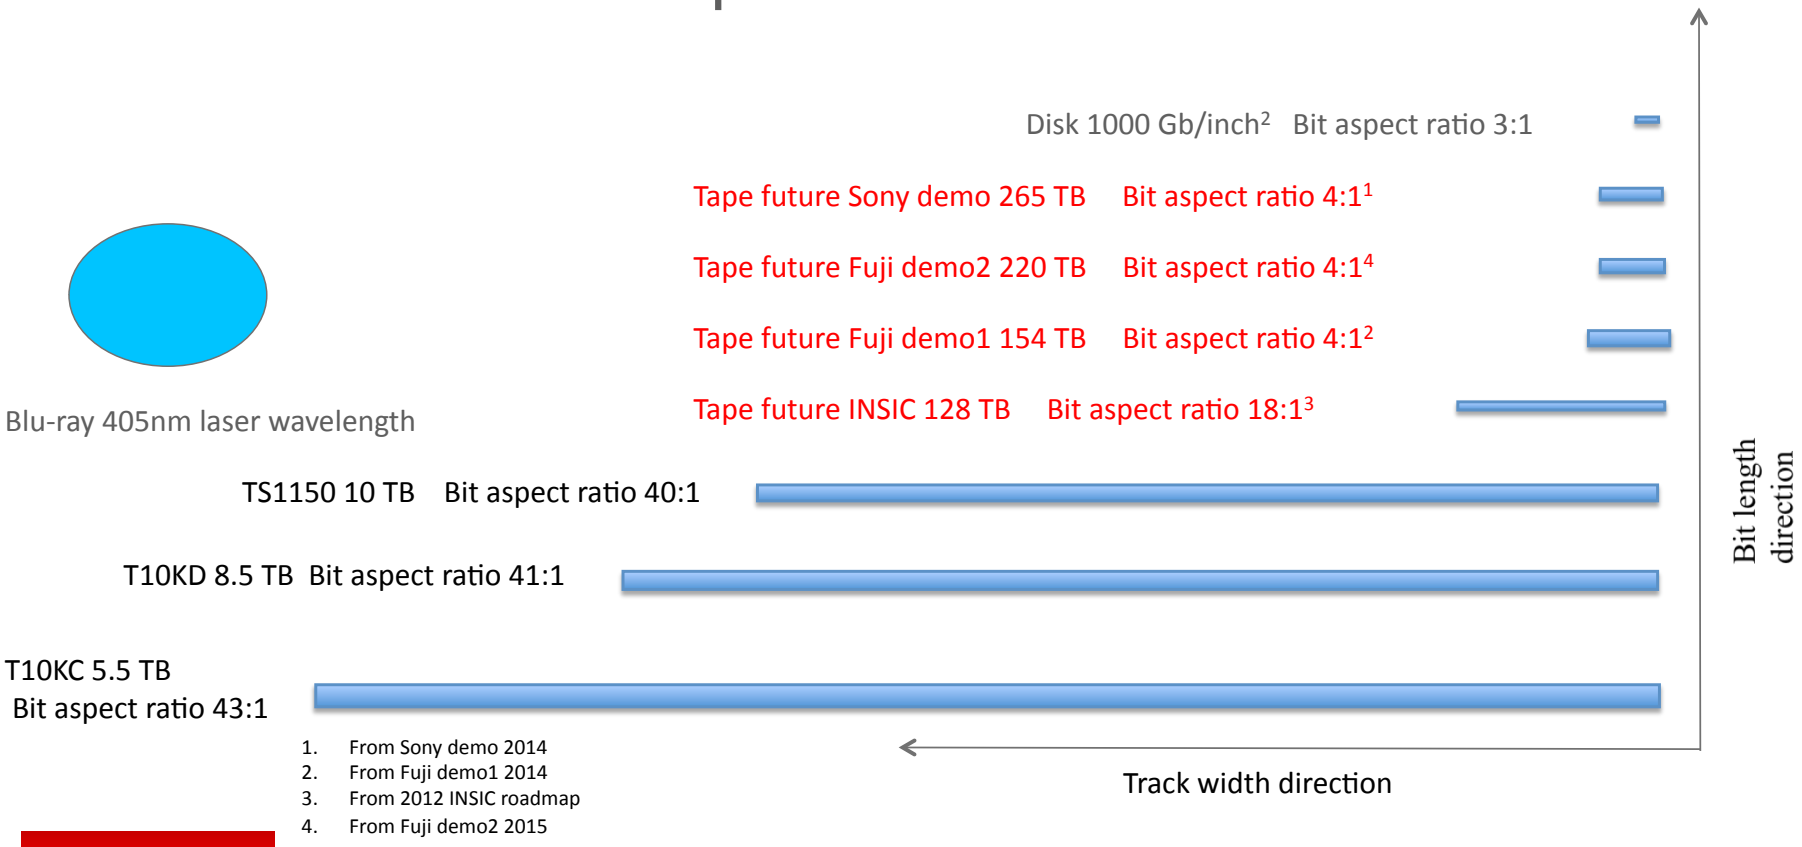

1. From Sony demo 2014
2. From Fuji demo1 2014
3. From 2012 INSIC roadmap
4. From Fuji demo2 2015

ORACLE

As Bits Get Small: Bit rot (super-paramagnetic effect ,thermal instability)

This is what is slowing down disk capacity growth!

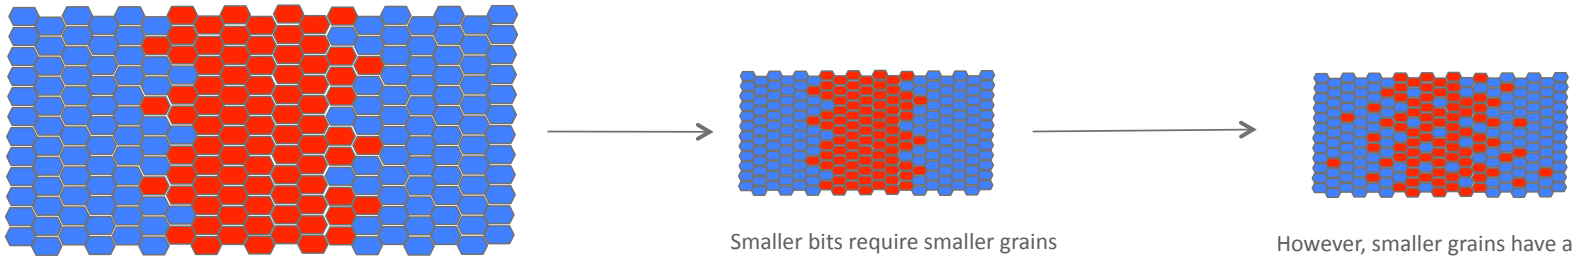

Smaller bits require smaller grains for required SNR

However, smaller grains have a higher probability of reversing over time

Neel-Arrhenius law gives: Mean time to randomly flip grain due to thermal fluctuations $\tau_N = \tau_0 \exp\left(\frac{KV}{k_B T}\right)$

V is the volume of the grain, T is the temperature and K is the grain's magnetic anisotropy energy

$\frac{KV}{k_B T}$ > 60 for good thermal stability, 10 year data life [1,2]

$\frac{KV}{k_B T}$ > 90 for today's tape [3], 30 year data life

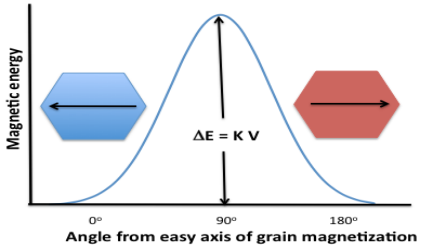

1. Dobisz et al. Patterned Media: Nanofabrication Challenges of Future Disk Drives, Proceedings of the IEEE, Vol. 96, No. 11, November 2008
2. Weller et al. Thermal Effect Limits in Ultrahigh-Density Magnetic Recording, IEEE Transactions on Magnetics, VOL. 35, NO. 6,November1999
3. Watson et al. Investigation of Thermal Demagnetization Effects in Data Recorded on Advanced Barium Ferrite Recording Media, IEEE Transactions on Magnetics, Vol. 44, No. 11, November. 2008

Tape Storage Trends

Oracle StorageTek Tape – A Look Back 16 years

10 TB in 1998

~ 6000 carts
TimberLine 9490EE – 1.6 GB ea
357 sq ft
8200 lbs

10 TB in 2014

< 2 carts
T10000D – 8.5 TB ea
0.3 sq ft
1.2 lbs

ORACLE

Oracle StorageTek Tape – A Look Back 16 years

10 PB in 1998

x10

~ 6,000,000 carts
~ 8 acres
~ 4,100 tons

10 PB in 2014

1,177 carts
StorageTek SL3000 with T10000D
37 sq. feet
~1.5 tons

Storage Technologies Areal Density Trends

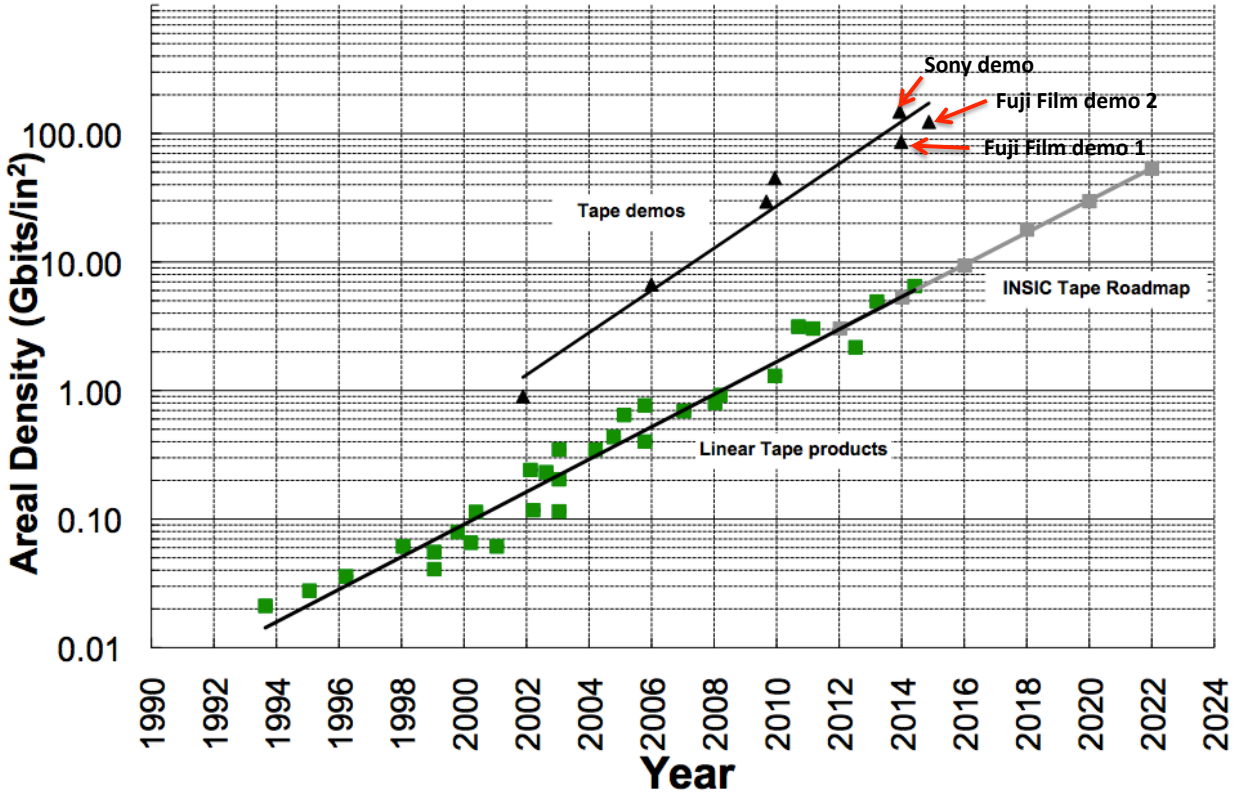

Tape Storage Projections - Recent Technology Demos

Demos show we've got solid technology to achieve roadmap goals

- INSIC tape roadmap shows technology path to 128 TB on a cartridge
- Sony sputtered media demo (4/14)
  - Areal density of 148 GB/in^2
  - '265 TB' cartridge
  - http://www.sony.net/SonyInfo/News/Press/201404/14-044E/index.html
- Fujifilm advanced BaFe demo1 (5/14)
  - Areal density of 85.9 GB/in^2
  - '154 TB' cartridge
  - http://www.fujifilmusa.com/press/news/display_news?newsID=880613
- Fujifilm advanced BaFe demo2 (4/15)
  - Areal density of 123 GB/in^2
  - '220 TB' cartridge
  - http://www.research.ibm.com/labs/zurich/sto/tape/arealdensity.html

ORACLE

Disk Storage

Disk Magnetic Recording Tri-Lemma Review

- Smaller bits => Smaller grains for required SNR
- Smaller grains => Higher Hc[1] for thermal stability
- Higher Hc => Can not write on the media

1. Hc is the media Coercivity, which is the strength of the magnetic field required to flip the magnetization in the media

New Disk Technologies Required

- Helium drives
- Shingled recording
- Energy assisted recording
- Bit pattern recording

Shingled Recording Concept

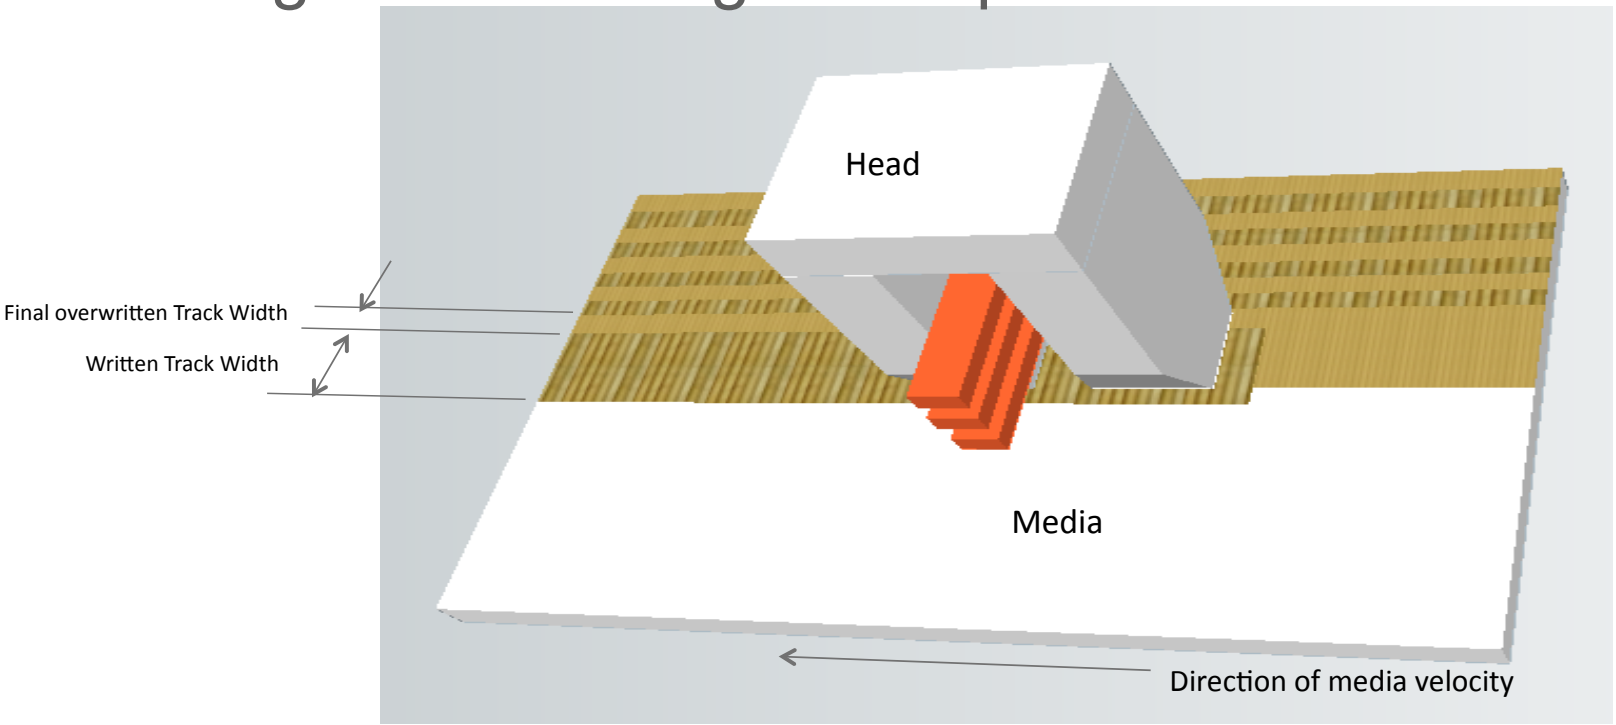

Wide tracks are partially overwritten to get narrower tracks

HAMR Concept

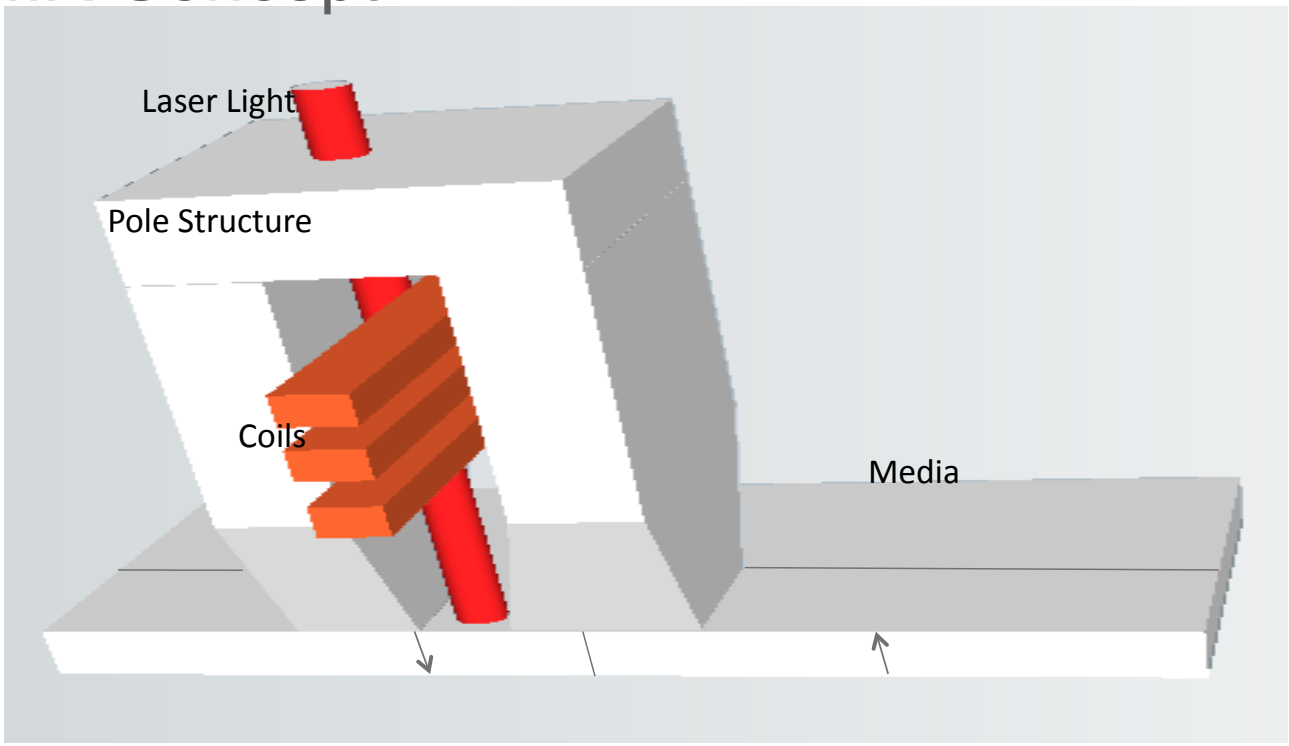

Laser heats media reducing media Hc so head magnetic field can write media

Flash Storage Trends

FLASH Challenges

- Reduction in cell size and more bits per cell results in degradation of retention time and endurance
  - 10 year retention dropping to 1year at end of endurance due to write cycles[1]
  - State detection level is determined by a small number of electrons
    - ~ 8 electrons per level for 16 nm TLC device
- As cell size shrink interference between cells increases[2]
- Basic performance has not improved (read, write and erase latencies) over the last decade[1]

1 International Technology Roadmap For Semiconductors, 2011 Edition Emerging Research Devices page 18
2 http://www.forbes.com/sites/michaelkanellos/2013/08/14/with-3d-chips-samsung-leaves-moores-law-behind/

3D NAND

- 3D stacking cells on top of each other enabling significant density increases
- Eliminate the need to reduce dimensions
  - no new lithographic technology needed, just add more layers to increase capacity
- Compared to latest 2-D NAND[1]
  - 2X the number of cells/inch[2]
  - ½ the power,
  - 2X as fast
  - 10X the endurance

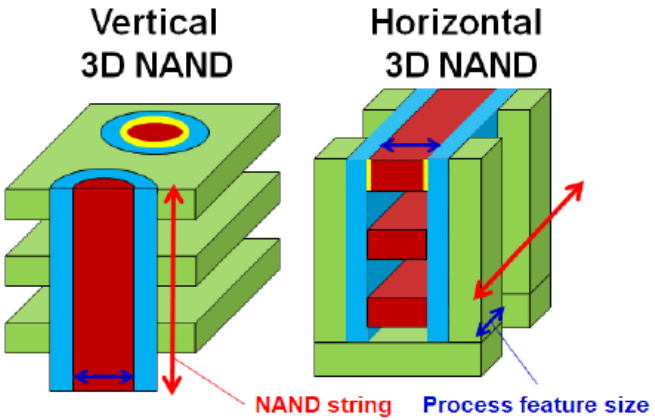

3D NAND architectures.

1. http://www.forbes.com/sites/michaelkanellos/2013/08/14/with-3d-chips-samsung-leaves-moores-law-behind/

3D NAND FLASH: Moving forward now with 5 manufacturers

- 1st to market: Samsung in 2013
  - 128 Gbit chip[1]
  - 24 layers of Flash cells
  - > 2.9 billion cells
  - 32 layer version released 5/14
- Intel and Micro announce 3D[2]
  - Could see 10TB in SSD drive format
- Toshiba and Sandisk announce 3D[3]

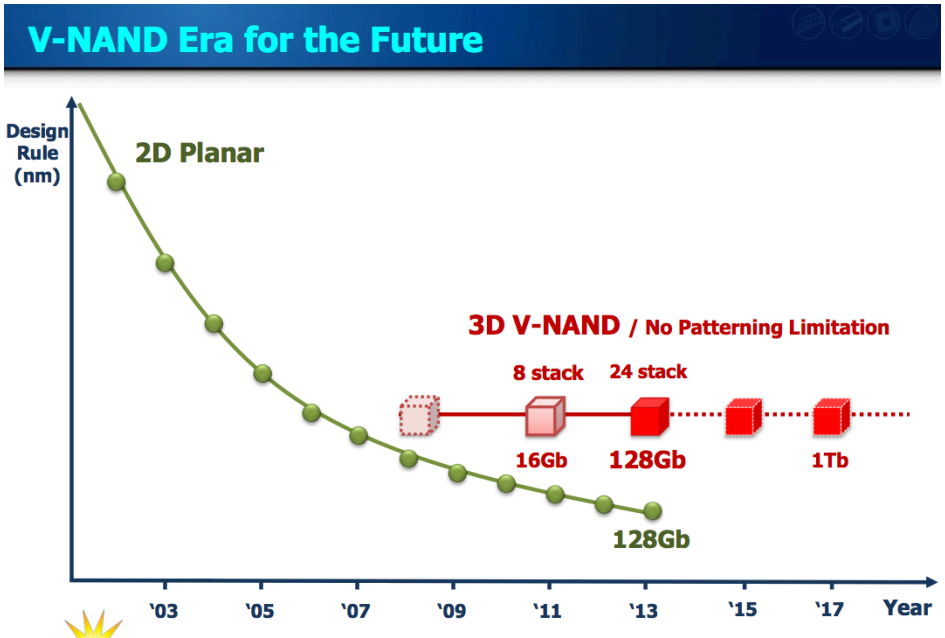

1. http://www.anandtech.com/show/7237/samsungs-vnand-hitting-the-reset-button-on-nand-scaling
2. http://www.digitaltrends.com/computing/need-space-ssd-intels-3d-nand-may-answer/
3. http://www.cbronline.com/news/tech/hardware/storage/toshiba-and-sandisk-partner-to-produce-high-capacity-3d-memory-chips-4268156

Hardware and Software
Engineered to Work Together


ORACLE®