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 remain at the sole discretion of Oracle.

ORACLE®

Advanced Tape Technologies for Future Archive Storage Systems

MSST - Media II (Tape Media and Libraries) - Wednesday May 8th, 2013

Dr. Mark L Watson
Oracle Corporation

with due acknowledgements to Dr. Robert M Raymond

Program Agenda

- Storage Growth / Requirements / Device Trends
- Recent Disk Trends
- Tape Storage Developments

Customer Challenge: “Big Data” is Exploding!!!

50%/yr Growth Rate

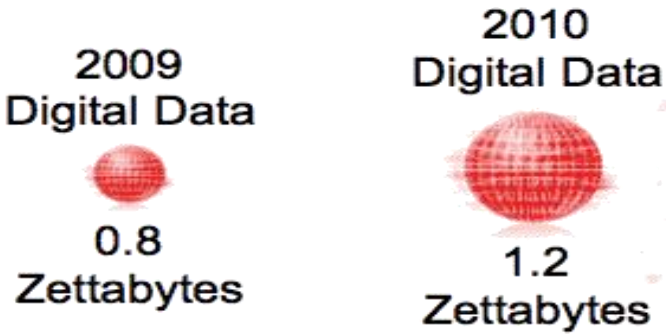

- IT budgets and headcounts can’t keep up
  - Budget are projected to grow 2-3%/yr*
- Disk prices are not declining at this rate
- **Customers can’t afford to just “put everything on disk”**

**Source: SearchStorage Magazine, July 2011*

Source: IDC, White Paper Sponsored by EMC, The Digital Universe Decade - Are You Ready? Doc.# IDC_925, May 2010

Storage Technologies: economics!

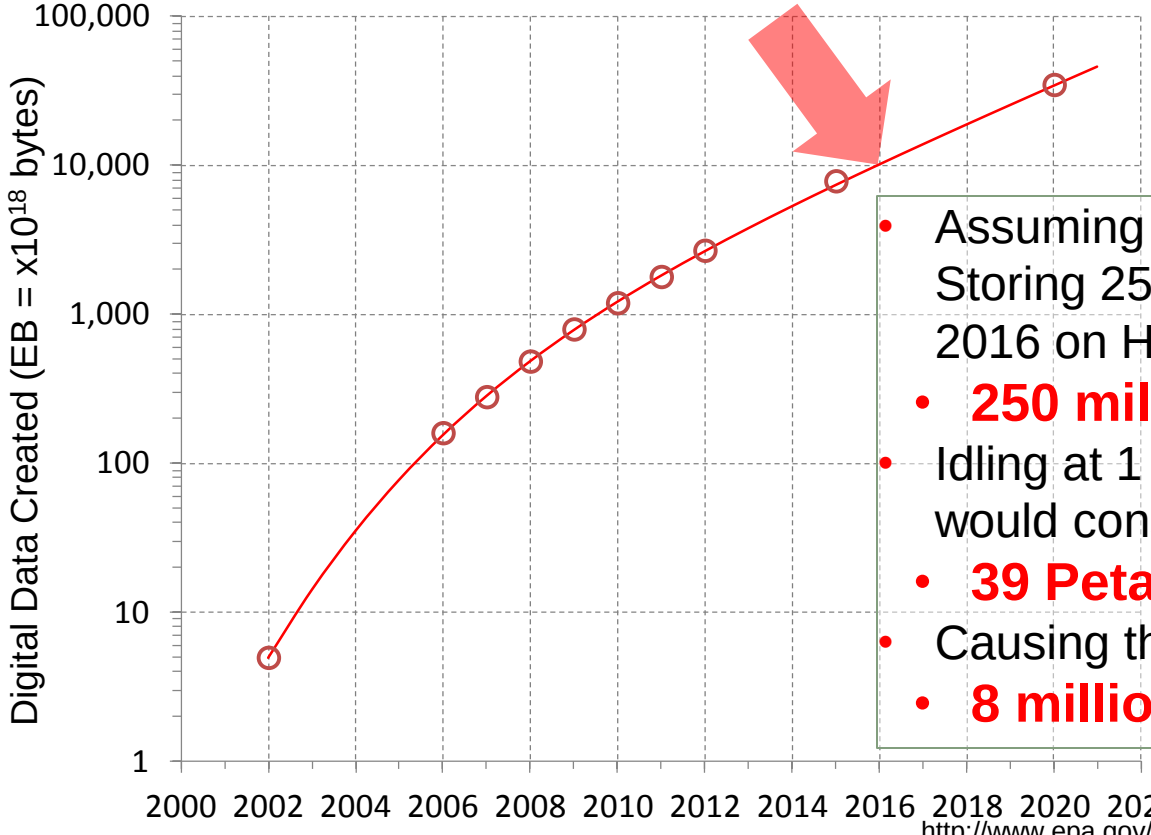

- Assuming a 10 TB HDD is available Storing 25% of the data created in 2016 on HDD would need:
  - **250 million hard drives**
- Idling at 1 watt over 5 years these would consume
  - **39 Peta-Joules of energy**
- Causing the emission of
  - **8 million tons of CO_2**

http://www.epa.gov/cleanenergy/energy-resources/calculator.html#results

Economics of Tiered Storage

Tape is the Foundation: Most of the Data Stored at the Lowest Cost

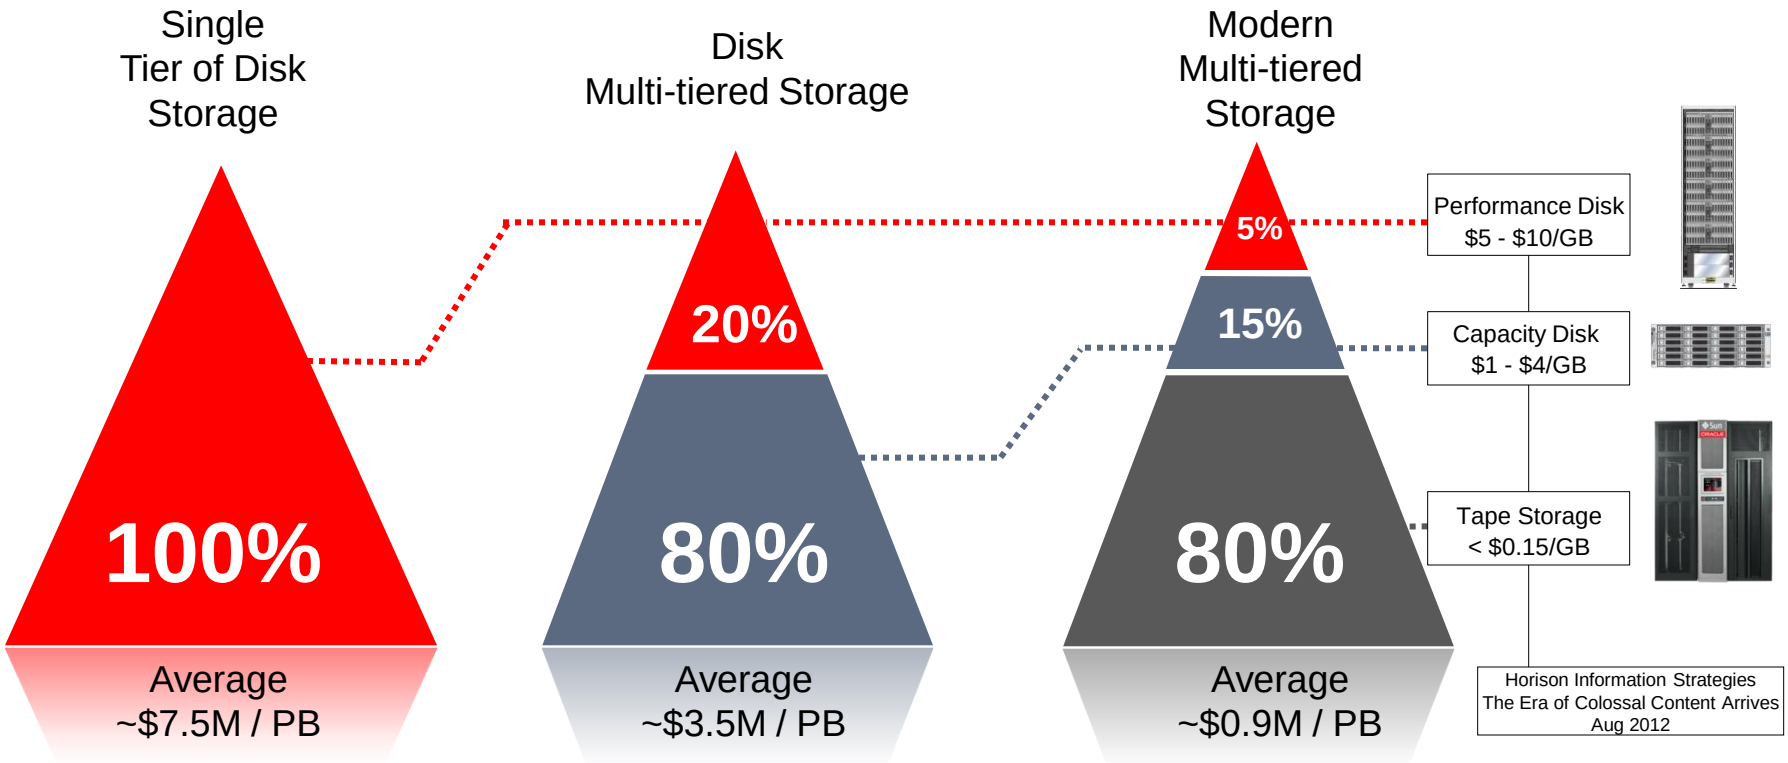

Storage Technologies Areal Density Trends

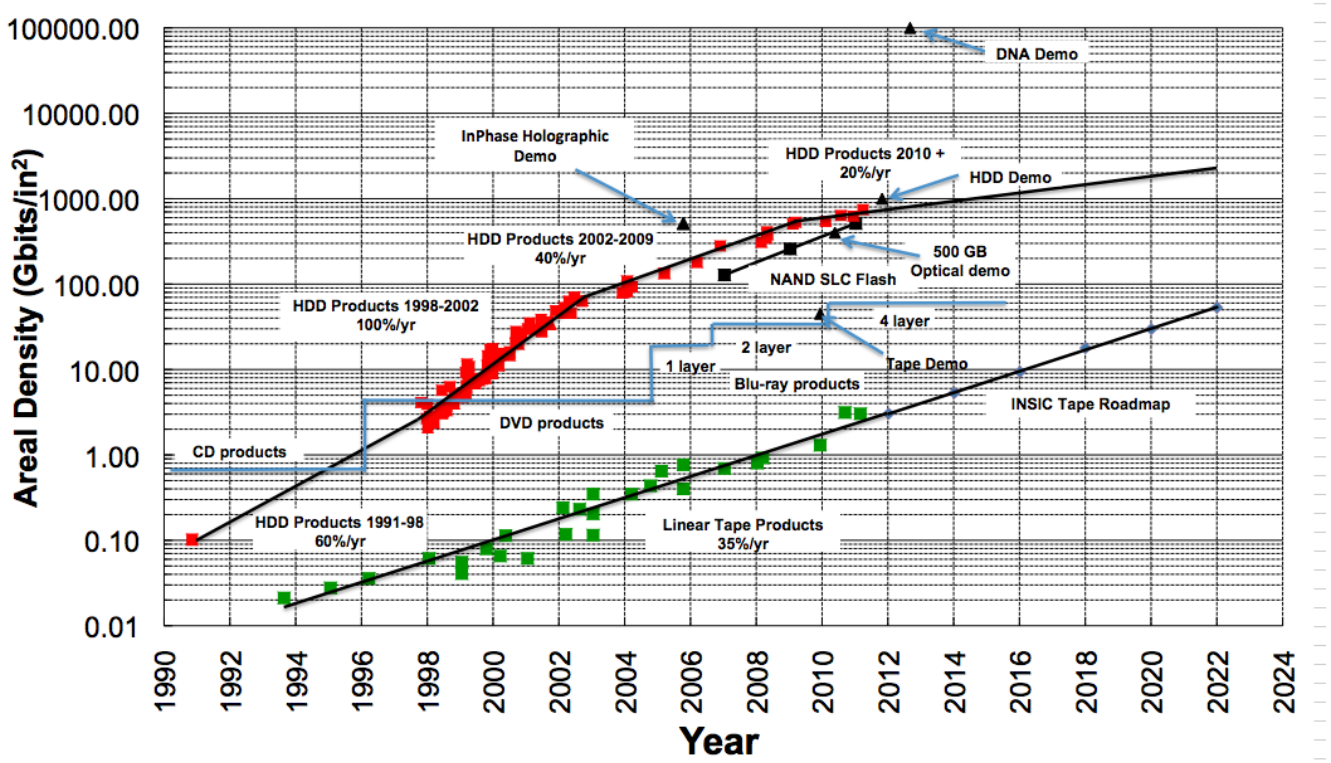

Source: disk areal density growth: http://www.forbes.com/sites/tomcoughlin/2012/10/03/have-hard-disk-drives-peaked/

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

Technology Price per GB Predictions

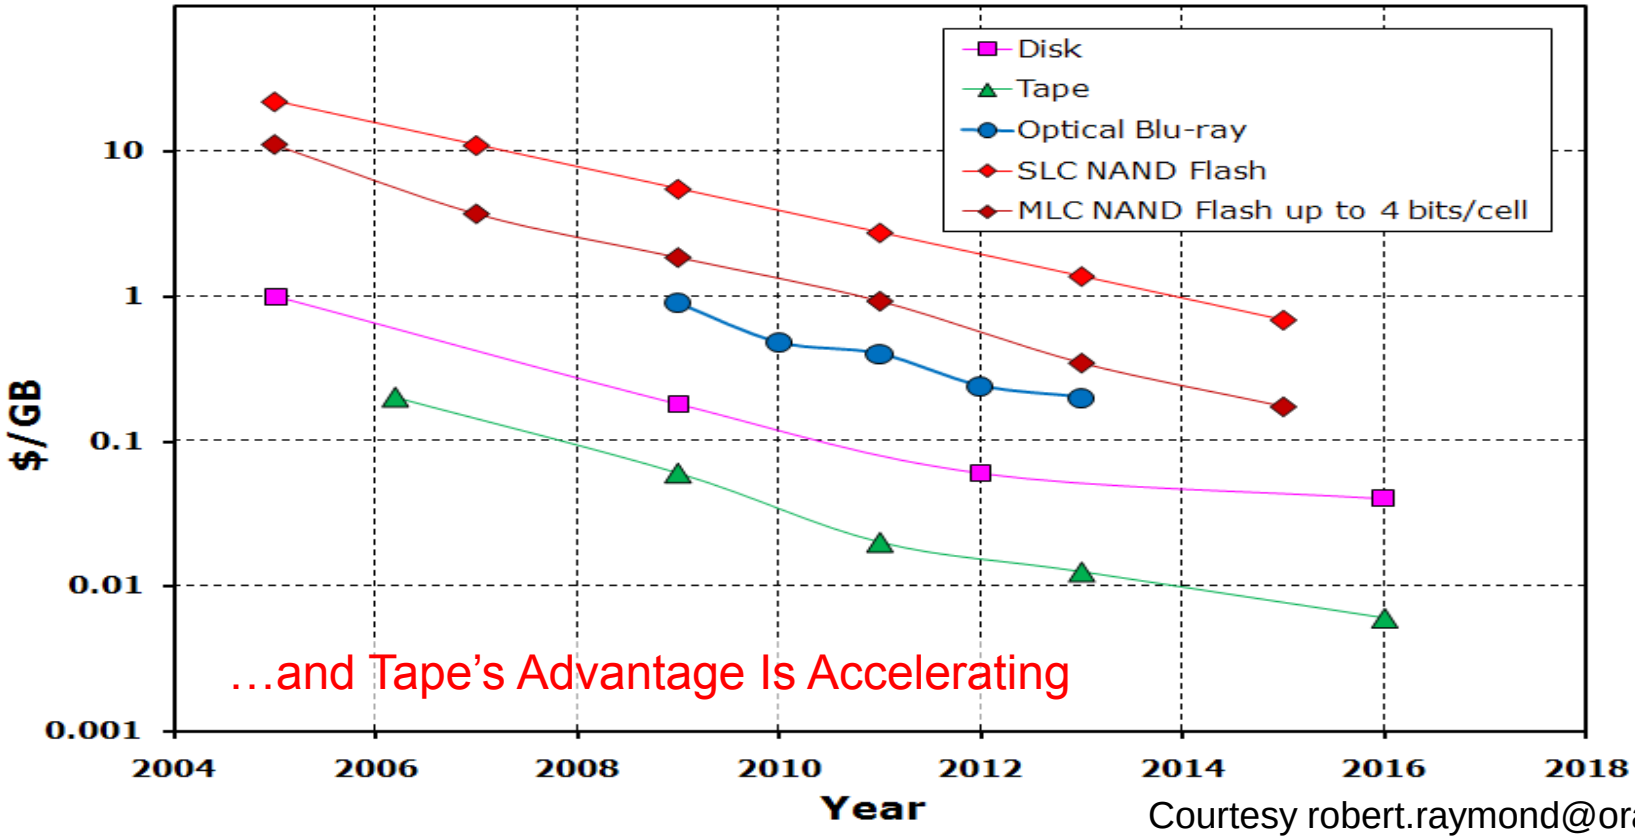

Courtesy robert.raymond@oracle.com

Magnetic Recording Definitions

Track Pitch = Track Width + distance between neighboring tracks

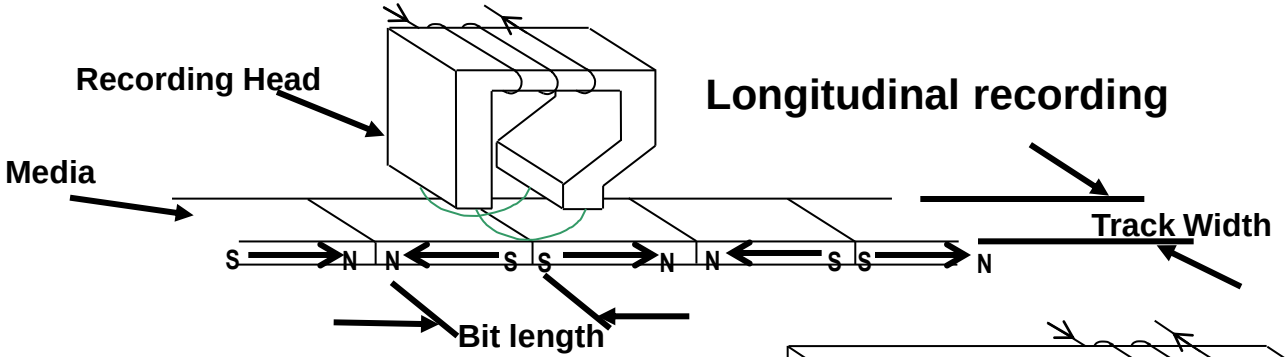

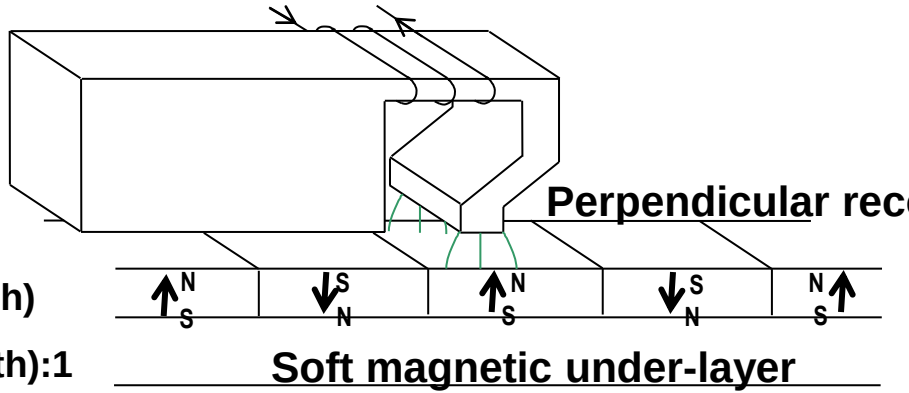

Tracks per inch (tpi) = (Track Pitch)$^{-1}$

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

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

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

Data Bit Size Comparison

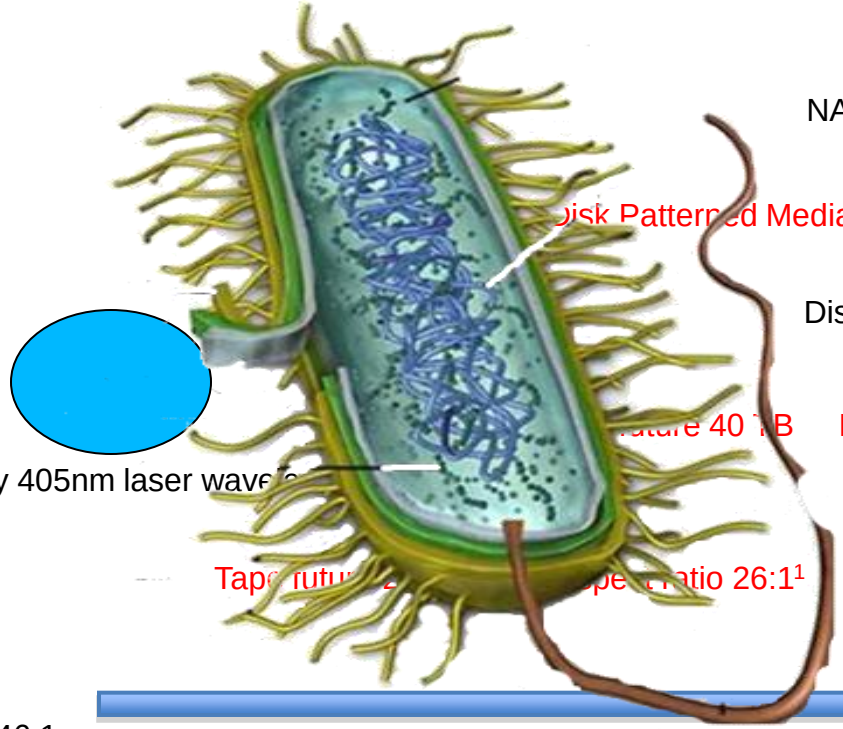

Bacterium
(2 um)

Blu-ray 405nm laser wave

NAND 250 Gb/inch^2 50nm X 50nm

isk Patterned Media future 1Tb/inch^2 Bit aspect ratio 1:1

Disk 500Gb/inch^2 Bit aspect ratio 5:1

ature 40 TB Bit aspect ratio 15:1[1]

Tape future 2 pect Ratio 26:1[1]

T10KC 5.0 TB
Bit aspect ratio 46:1

Bit length direction

Track width direction

1. Extrapolated from INSIC roadmap

Written Data Tracks on Tape (T10KC)

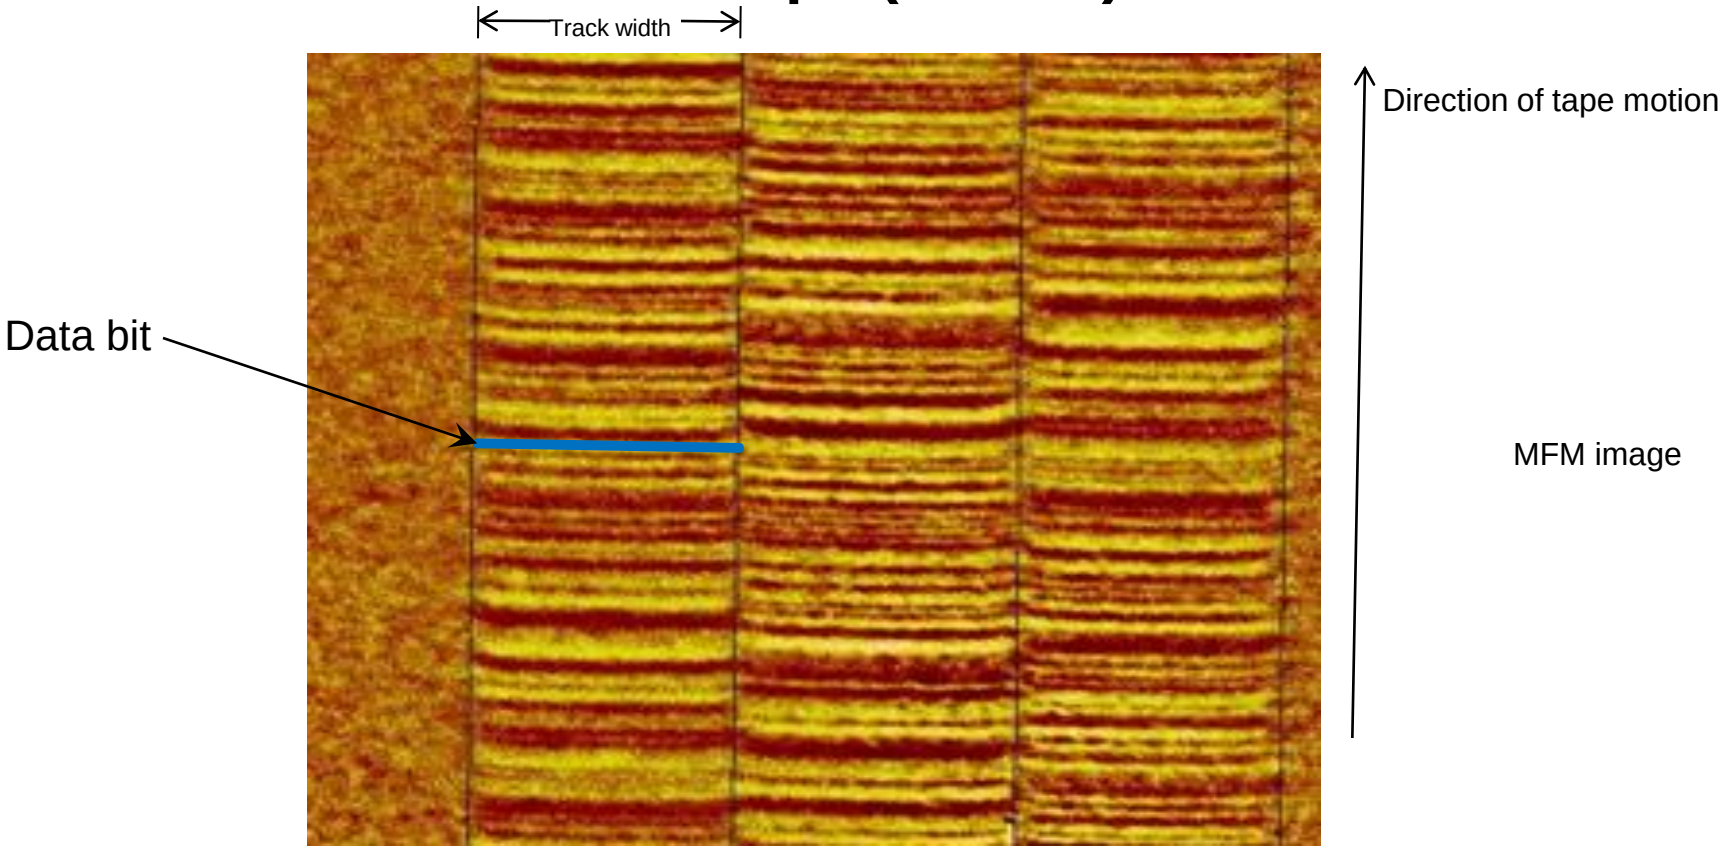

Tape Bit Size Roadmap

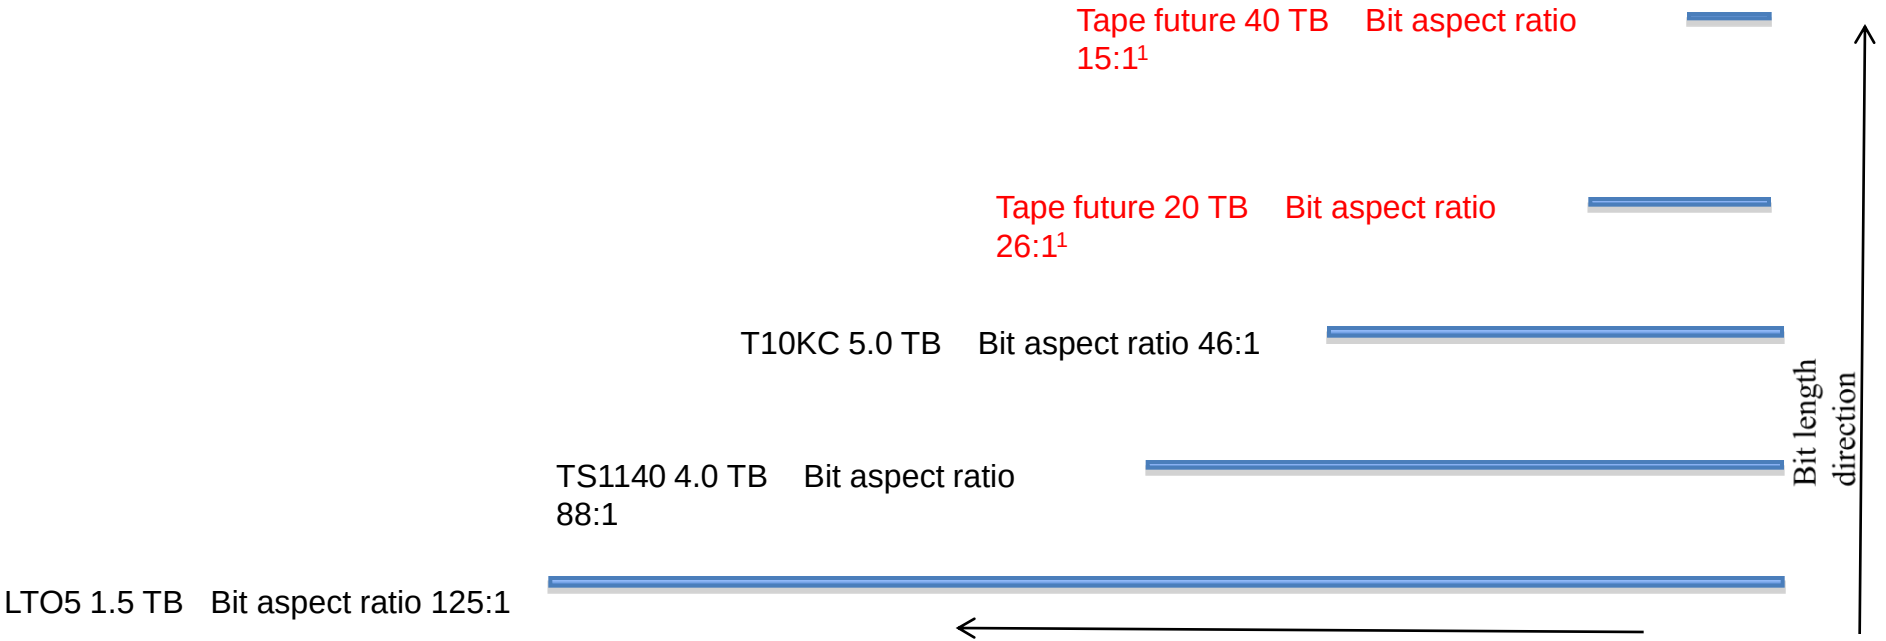

1. Extrapolated from INSIC roadmap

Recent Disk Trends

Disk Magnetic Recording Tri-Lemma

- 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

Problem: The Super-Paramagnetic Effect

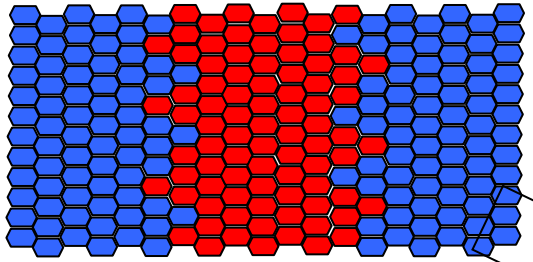

- In order to maintain the necessary signal to noise ration, smaller bits require smaller grains
- Small grains are easier to de-magnetize

Over time…

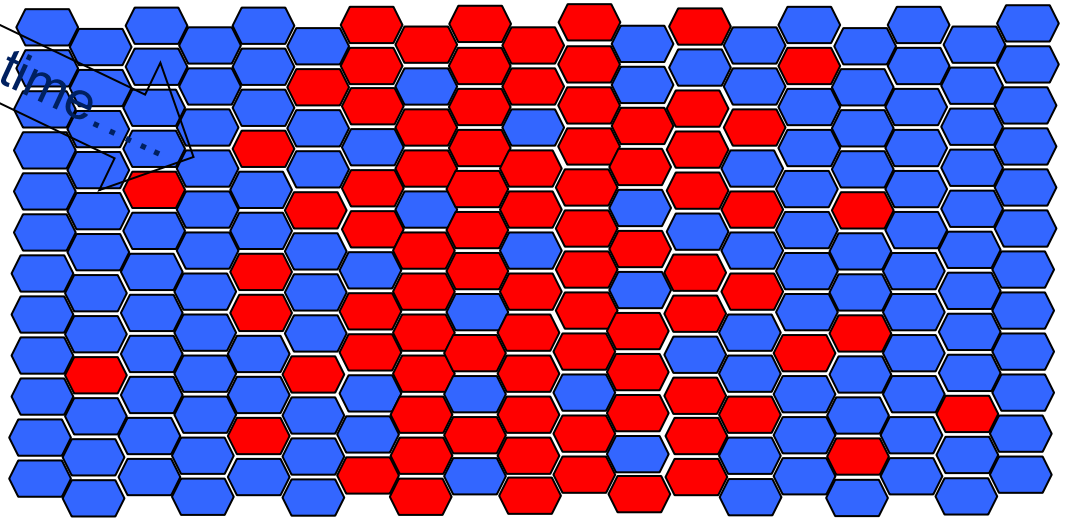

- Can result in loss of information
- Overcoming this requires
- Higher anisotropy recording materials
- Which need new methods of writing data

Possible New Disk Technologies

- Areal Density Options[1]
  - Discrete Track Recording (DTR) (~2 Tb/inch^2)
  - Shingled Recording (SWR) (~5 Tb/inch^2)
  - Heat Assisted Magnetic Recording (HAMR) (~5 Tb/inch^2)
  - Microwave Assisted Magnetic Recording (MAMR) (~5 Tb/inch^2)
  - Bit Patterned Media (BPMR) (~5 Tb/inch^2)
  - 2-D Recording (TDMR)
- Capacity Options[2]
  - Helium drive: more platters, 23% less power, 40% more capacity

1 Source: “Future Options For HDD Storage” Y. Shiroishi, et al, IEEE Trans. On Mag. Vol. 45, NO 10, Oct 2009
2 http://www.pcworld.com/article/262274/helium_filled_wd_drives_promise_huge_boost_in_capacity.html

Disk Observations

- Toshiba areal density (2 1/2): 744 Gb/inch^2 (1) or 500GB per platter
- Seagate areal density (3 1/2): 625 Gb/inch^2 (2) or 1TB per platter
  - New disk technology areal density demo by Seagate(3)
  - HAMR with granular media: >1 Tb/inch^2
- New disk technology demo by IMRE (Institute of Materials Research and Engineering)(4,5)
  - BPM process only not a recording demo ~ 3.3 Tb/inch^2
- New disk technology areal density demo by Toshiba(6)
  - TDK realized an areal density of 1.5 Tbits/(inch)^2 with a magnetic head technology that is based on a thermal assist recording method and uses near-field light. The bit error rate (BER) is 10^{-2}.
- More and more consolidation

1. http://www.techpowerup.com/149967/Toshiba-Boosts-Performance-with-New-High-Areal-Density-1-TB-2.5-Inch-Hard-Drive.html
2. http://www.theinquirer.net/inquirer/news/2047313/seagate-unveils-3tb-external-hard-drive
3. http://arstechnica.com/business/news/2012/03/hamr-time-seagate-demos-terabyte-per-inch-hard-disk-technology.ars
4. http://arstechnica.com/gadgets/news/2011/10/researchers-increase-hard-drive-density-sixfold-with-salt.ars
5 http://iopscience.iop.org/0957-4484/22/38/385301/pdf/0957-4484_22_38_385301.pdf
6 http://techon.nikkeibp.co.jp/english/NEWS_EN/20121002/243229/

Disk Drive Manufacturers Family Tree

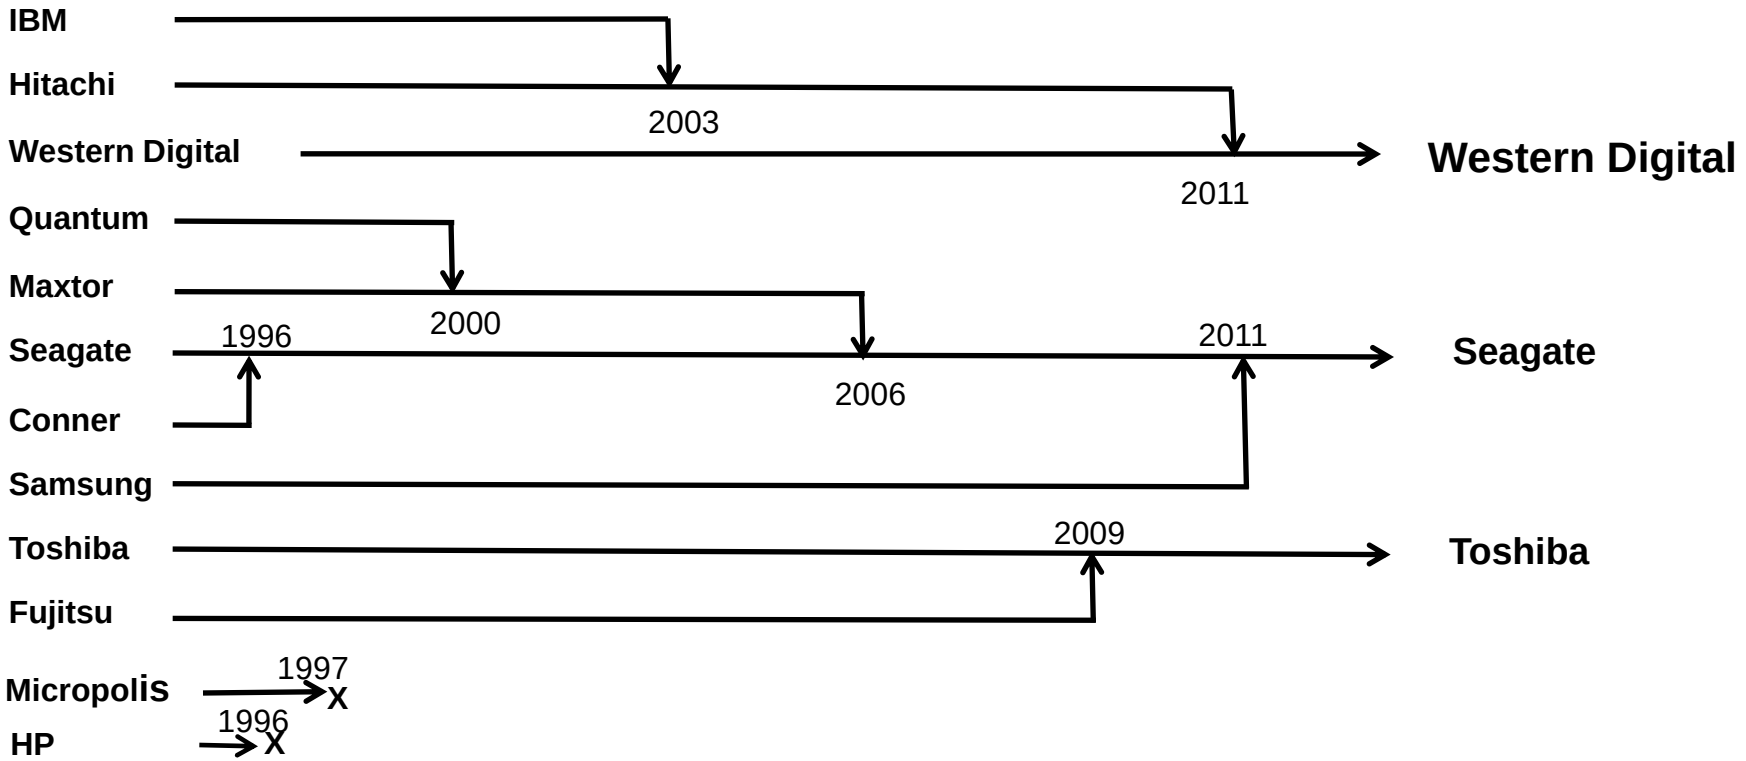

Source: http://en.wikipedia.org/wiki/History_of_hard_disk_drives

Tape Storage Developments

Technology Marches Forward

10 TB Example Over 15 Years

It is good to upgrade technology!

1996
- 6000 carts
- Timberline 9490 – 1.6 GB
- 357 sq ft
- 8200 lbs

2011
- 2 carts
- T10000C – 5.0 TB
- 0.3 sq ft
- 1.2 lbs

Large Libraries Provide Ideal Data Archiving Solutions

- The hardware costs are amortize over many tape cartridges providing the lowest TCO for large data storage archive repositories.
  - **15X less TCO and 238X less energy over 12 years than disk[1]**

Pictured: Ten string StorageTek SL8500 tape library, capable of holding 100,000 media cartridges

1. "In Search of the Long-Term Archiving Solution —Tape Delivers Significant TCO Advantage over Disk", Clipper Notes, December 23 2010, Report #TCG2010054LO

It’s Not Only About Cost/TB

	Disk	Tape
Max shelf life (bit rot)	10 years	30 years
Best practices for data migration to new technology	3-5 years	8-12 years
Uncorrected Bit Error Rate, Probability (avg 1 error in x TB)	10^{-14} (~10’s of TB)	10^{-19} (~1 million TB)
Power and cooling	238X	X

Each technology refresh or migration has a cost associated with it

Components of a Tape Drive

Recording heads

Take-up reel

Loader — The motor to engage the cartridge is beneath the loader

Tape cartridge

Red line shows tape

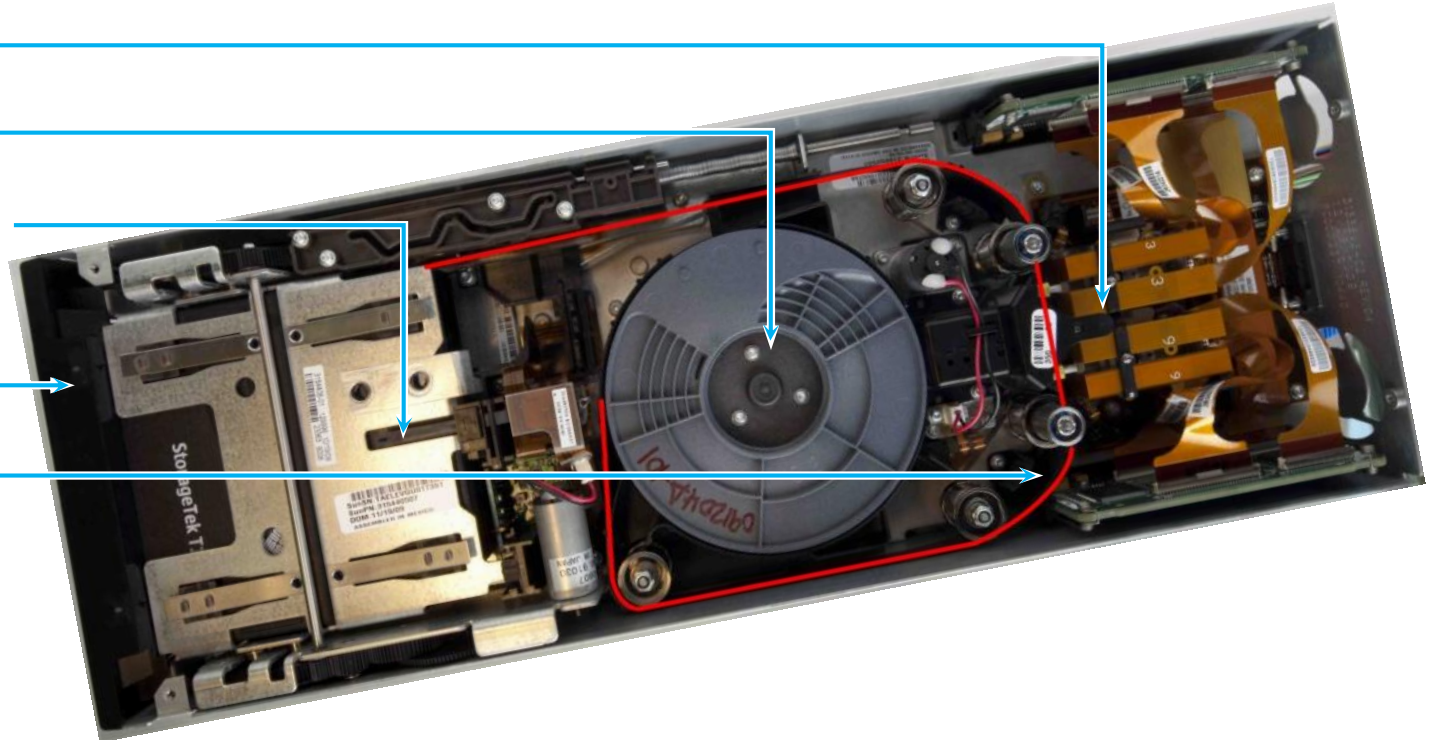

New Tape Technology Areas

- New Media (BaFe)
  - More stable than ever before
    - Magnetic (thermal), chemical, dimensional
  - Smaller particles or grains
  - Lower noise media structures
- New Recording Channels
  - New LDPC and other codes
- New Heads
  - MR (magneto resistive) >>>> GMR (giant magneto resistive)
  - 32 channels
  - Read older tapes back at least 2 drive generations
- Improved tape path / guiding

Media details – The Latest Technology

Top layer is magnetic layer (< 100 nm)
Bottom layer is non-magnetic iron-oxide layer

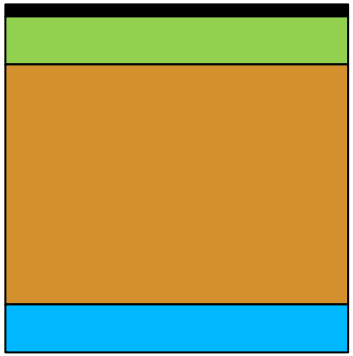

Magnetic layer
(MP or BaFe)

Underlayer
(non-magnetic)

Basefilm/Substrate
(PEN, PET, or Aramid)

Backcoat

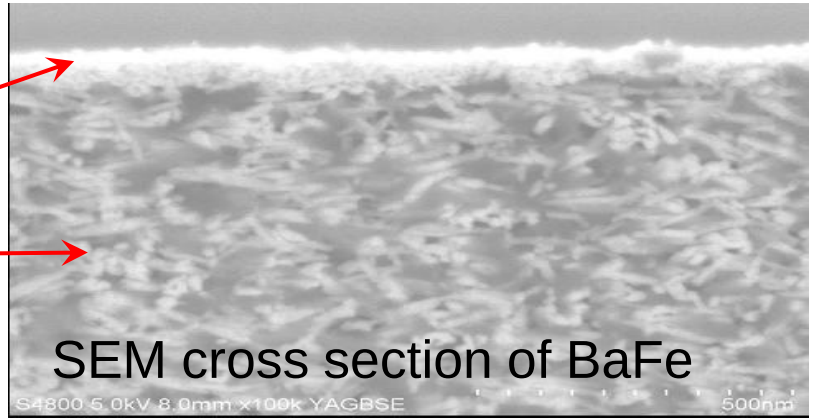

- Barium Ferrite ($BaFe_{12}O_{19}$)
- Hexagonal, platelet
- Naturally stable oxide
  - No corrosion
  - Chemically un-reactive

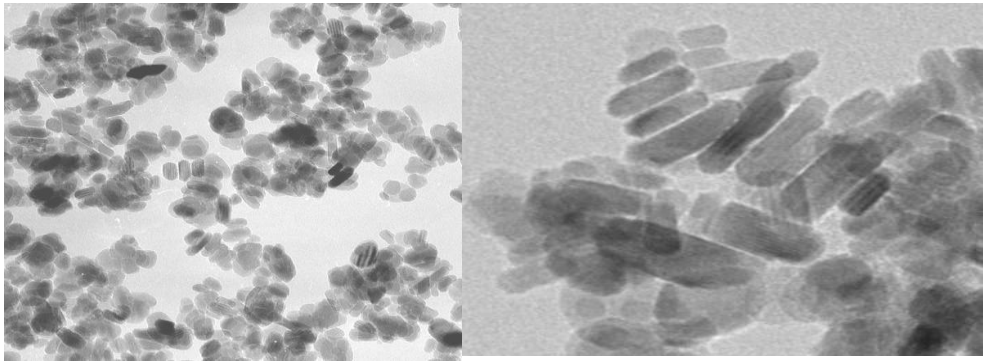

Dispersion of BaFe (TEM images)

The Future of Tape Media

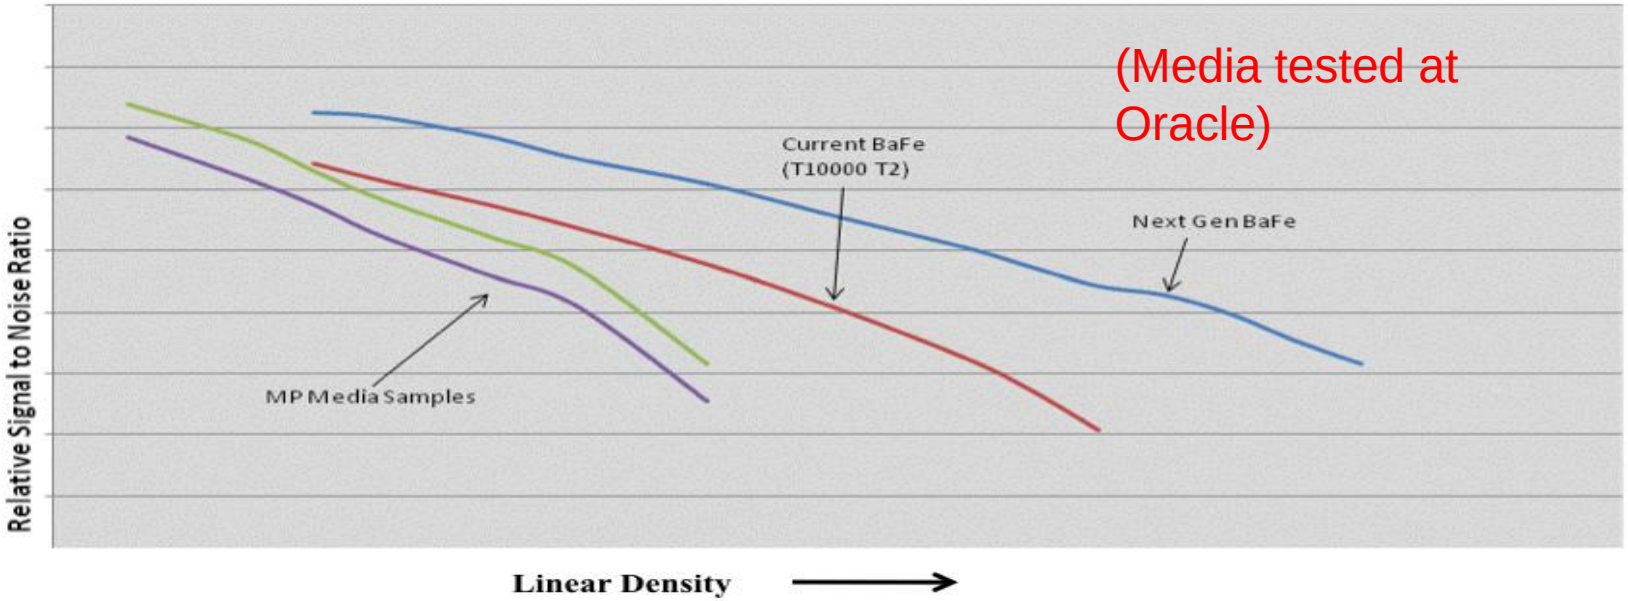

- Further improvements in recording properties of barium ferrite particles / media are possible
- Challenge will be developing drive technology to exploit this media development > new heads, guiding, R/W channels, servo positioning

Magnetic Degradation

Measuring Changes in Demagnetization

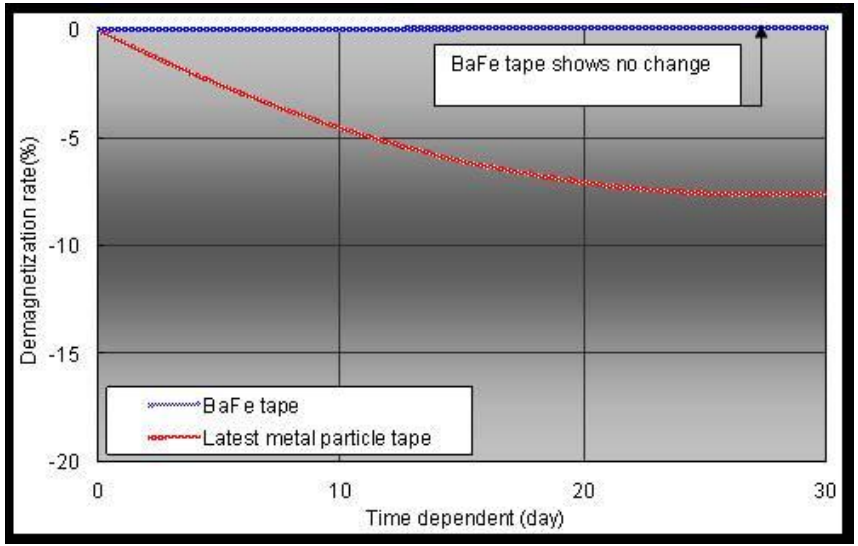

http://www.fujifilm.com/news/n100910.html

- 60C/90% RH for 30 days (equivalent to 30 years ambient storage)
- No change in BaFe magnetic properties
- BaFe ($BaFe_{12}O_{19}$) already is fully oxidized
  - Non-reactive
  - Extremely stable
- No impact to read or write performance for BaFe or MP

Thermal Stability of Magnetic Particles

No super-paramagnetic issues with tape!

Neel-Arrhenius law gives: Mean time to randomly flip the magnetization of a 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

$\frac{KV}{k_B T}$ > 90 for today's tape[1] ≡ Over a 30 year life!

1. 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 Areal Density Trends

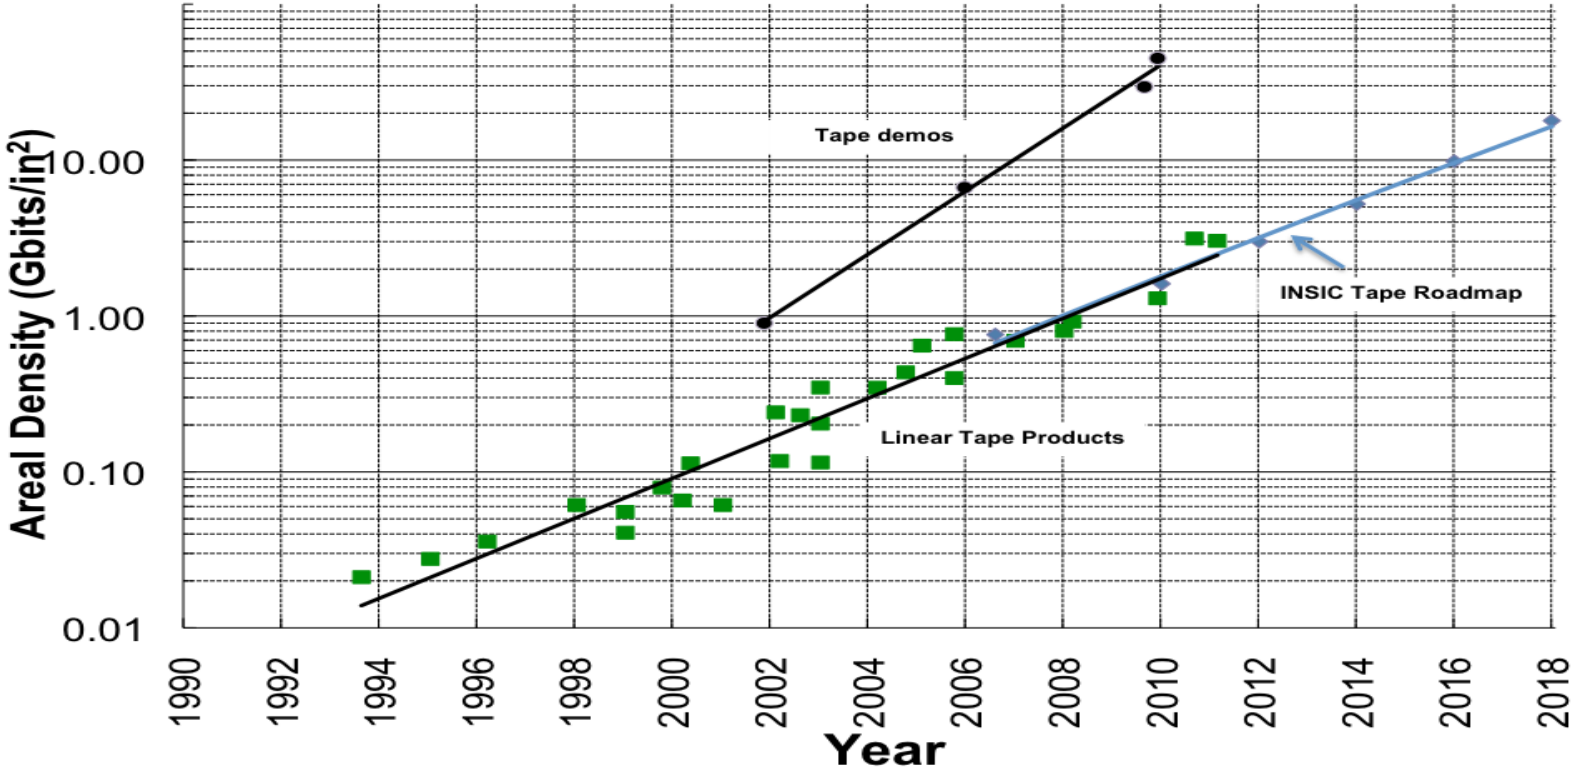

32

Tape Cartridge Capacity Trend

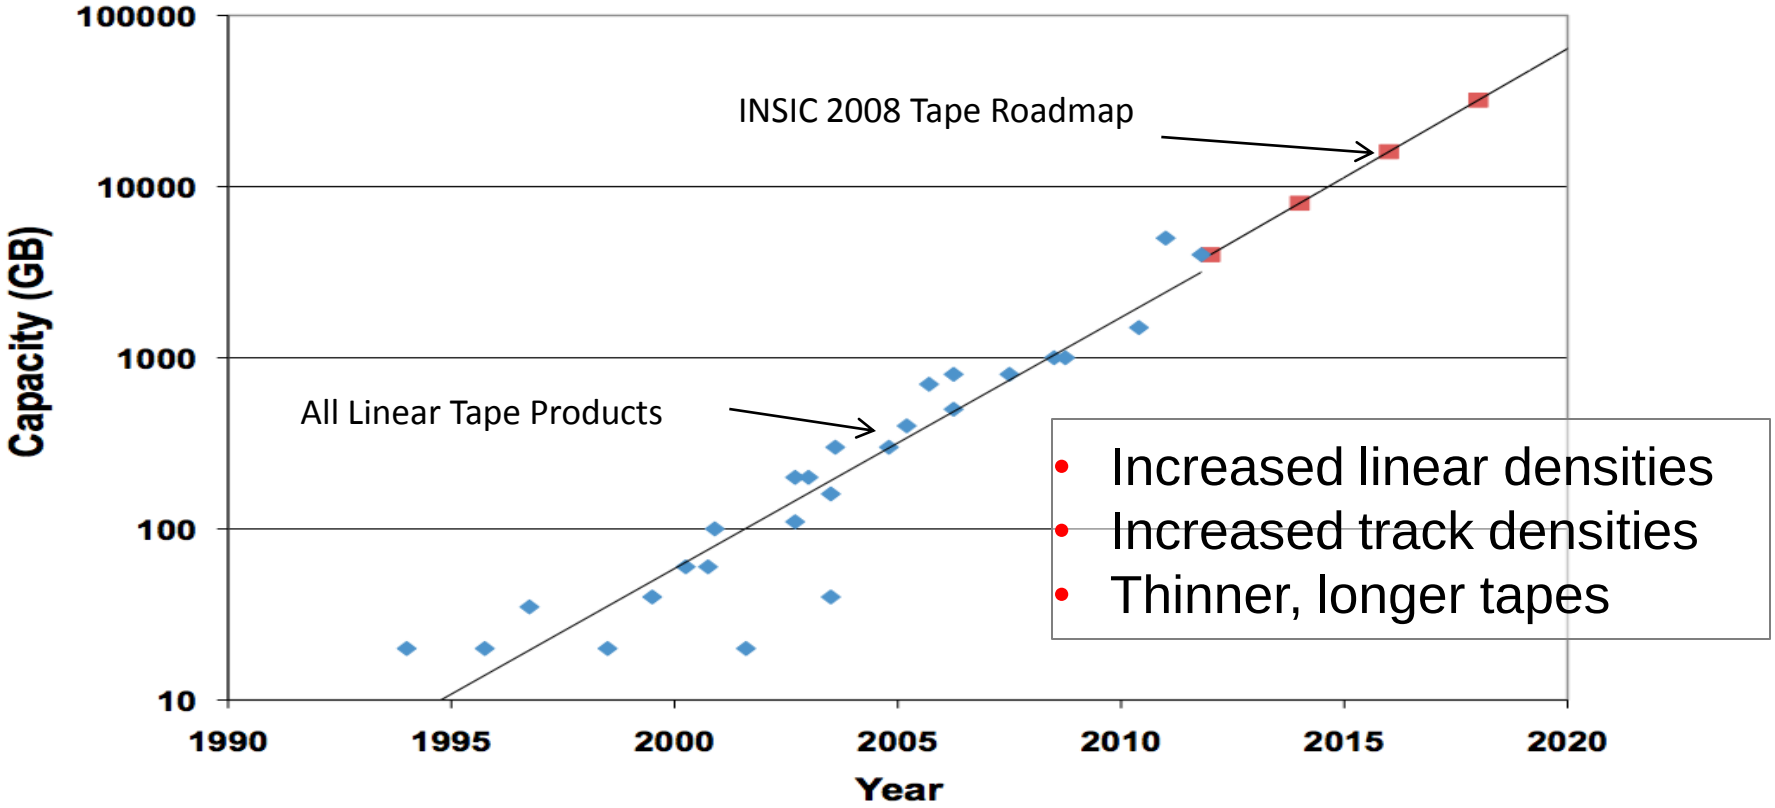

33

Tape Data Rate Trend

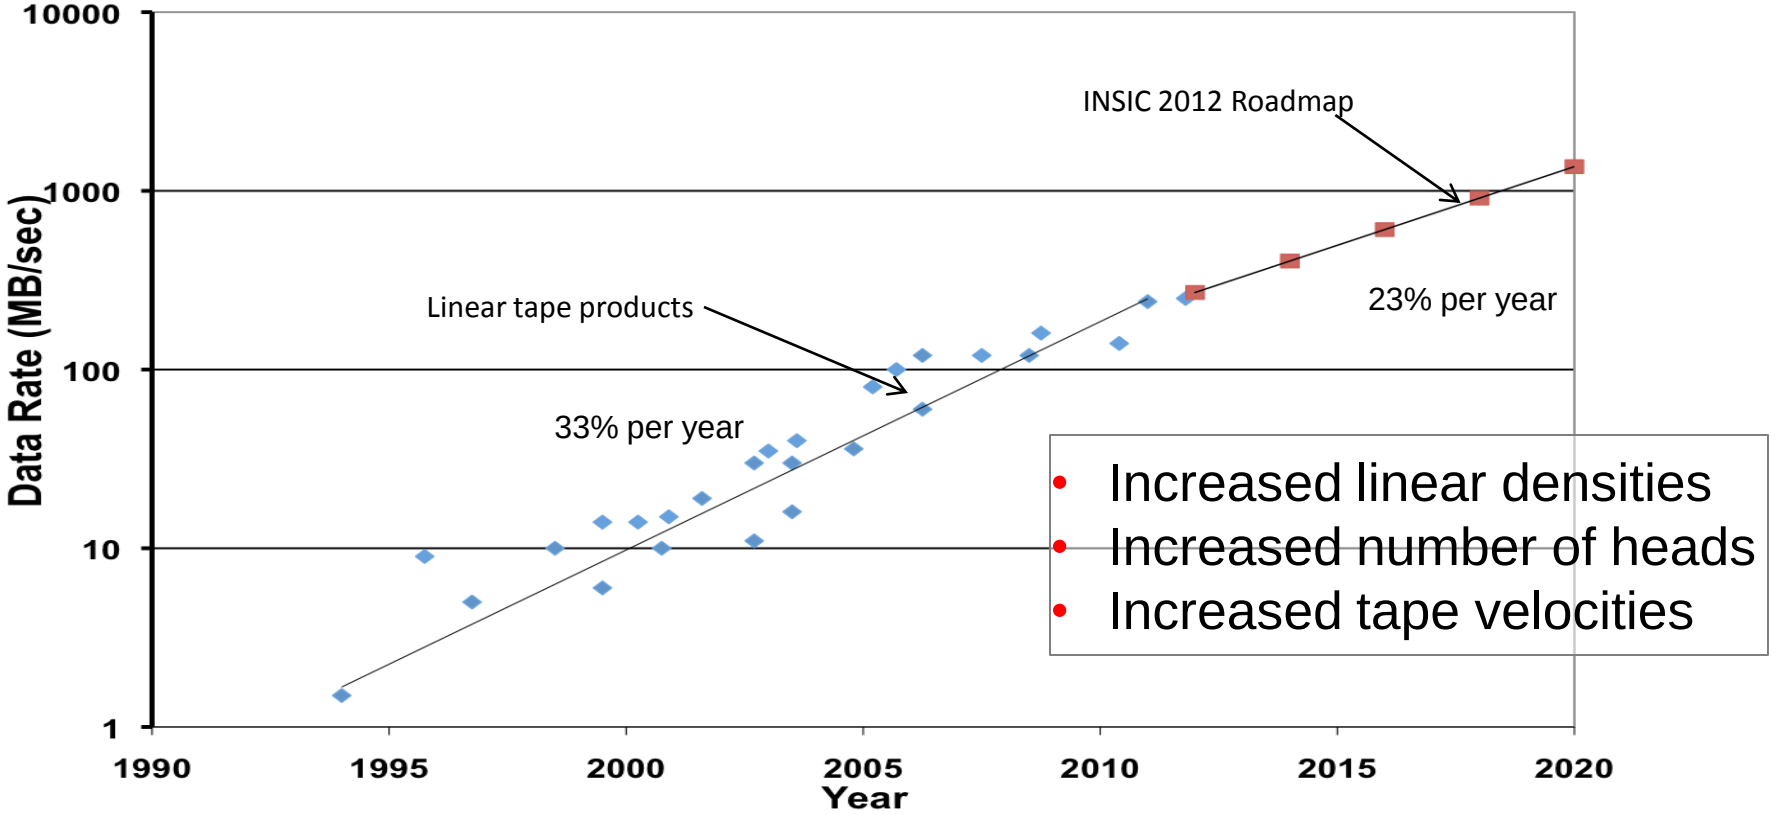

Summary

- Truly massive amounts of digital data are being created every year – and the rate of creation is increasing
- Managing / storing this data presents both problems and opportunities
- Magnetic Tape is the ideal large volume / extended period storage solution
  - Magnetic Tape continues to grow in both capacity and data rate
  - Tape has the lowest TCO over (any) other technologies
  - Tape has at least a 30 year archive life
  - New file systems (e.g. LTFS) are making tape much easier to use


ORACLE®