



# The Storage Lay<sup>1</sup> of the Land

**Ron Thielen – Winter 2007 CSG**

<sup>1</sup>Meaning “The way, position, or direction in which something is laid or lies” and not “A short lyric or narrative poem intended to be sung”

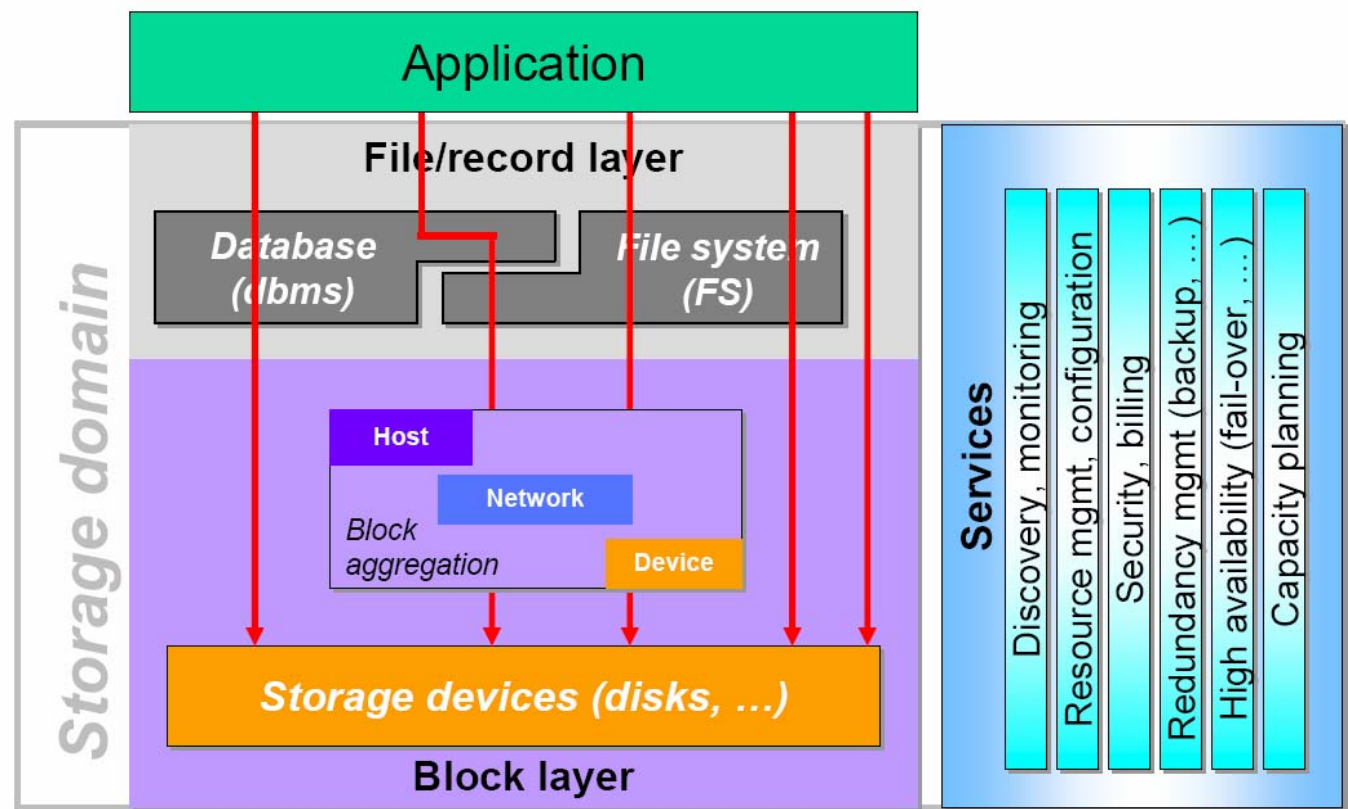
# Definitions

- Storage



# What are we talking about?

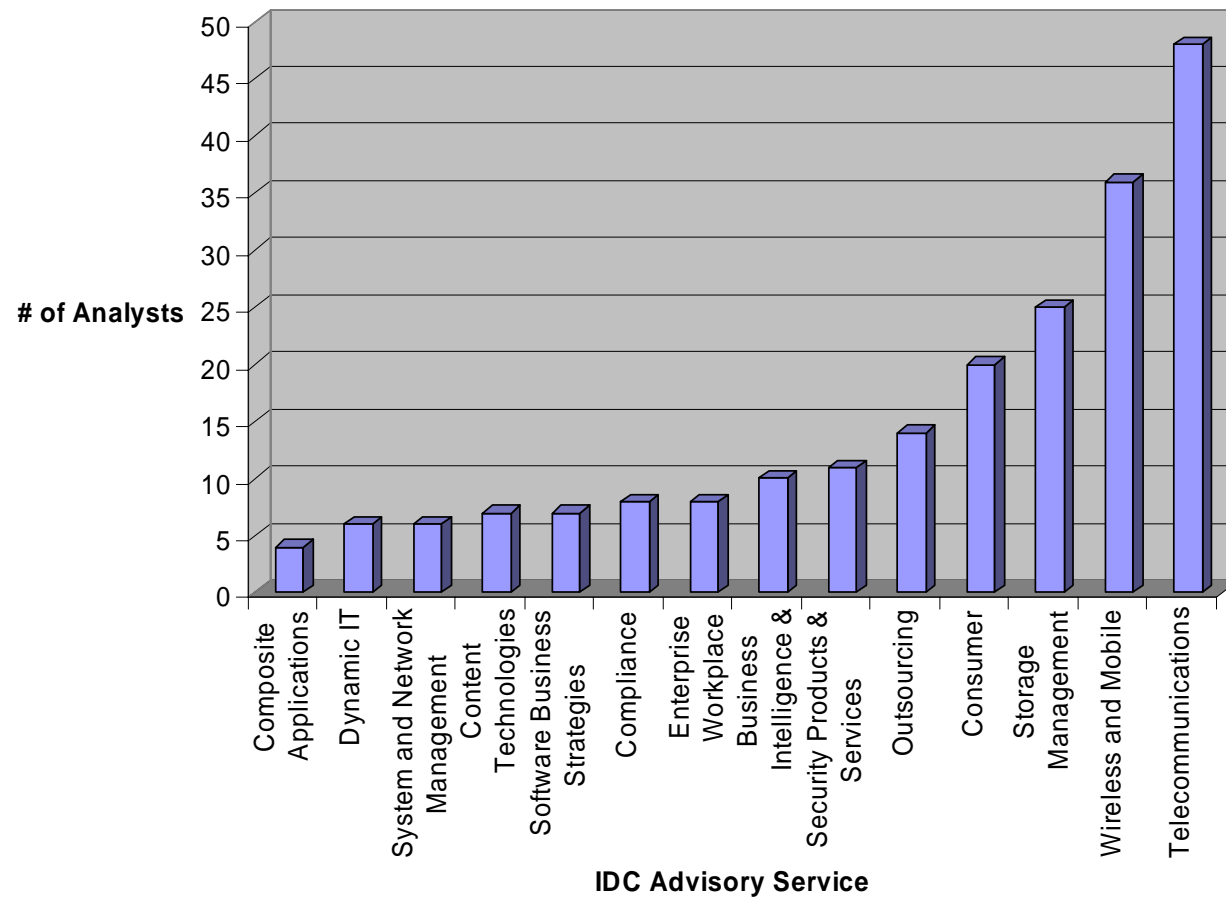
## The SNIA shared storage model



# Definitions

- Storage
- SAN or NAS
  - APIs – files, blocks, objects
  - Protocols – TCP/IP, FCP
  - Hardware – Fibre Channel, Ethernet, Infiniband
- File System – a metadata repository and related APIs
- <http://www.snia.org/education/dictionary>

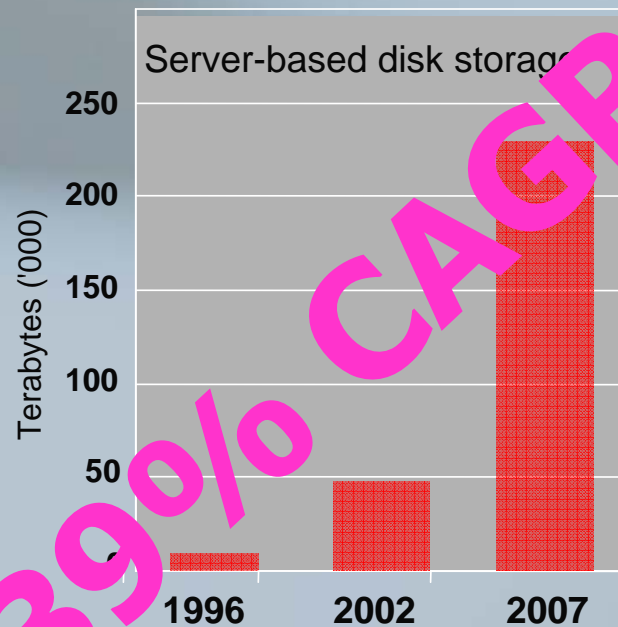
# IDC Analyst Index



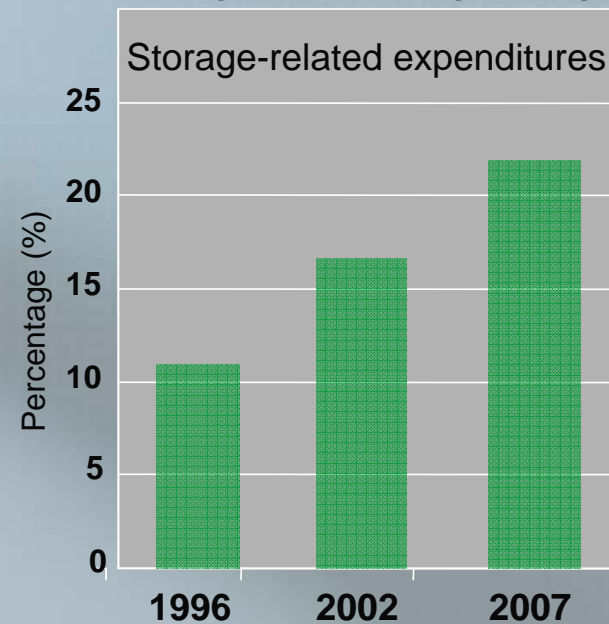
# A Tale of Two Growth Rates

In a typical Fortune 500 corporation...

Disk storage is growing rapidly



Storage-related expenditures\*, as a % of IT budgets, is also growing rapidly



Source: International Technology Group, Sept 2003

\* hardware, software, storage networking, personnel, backup operations, recovery, security

39% CAGR!

# Some Widely Reported CAGRs

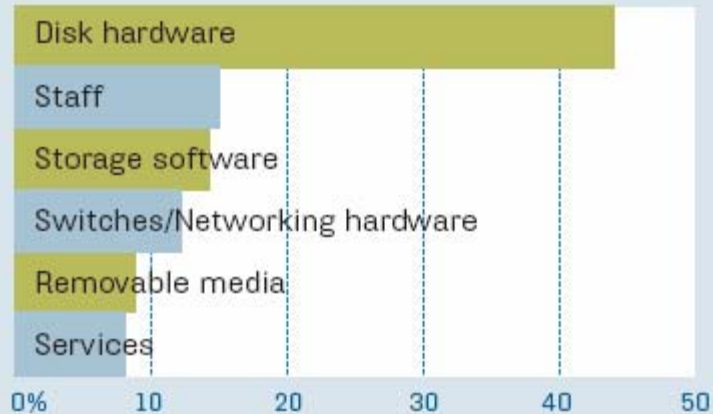
- >40% Storage Capacity (DMReview)
- 92% - Unstructured Reference Data (ESG)
- 64% - Data Subject to Retention for Compliance (ESG)
- 172% - Compliant Data on Disk
- -8% Direct Attached Storage
- 30% Disk Price Erosion

# Two Years of Growth on the Chicago Data Center SAN

	December 2004 TB	December 2006 TB	CAGR
Virtualized SATA Disk	14.0	72.8	128%
Virtualized 10K RPM Fibre Disk	8.0	32.9	103%
Virtualized 15K RPM Fibre Disk	7.2	11.2	25%
non-virtual 10K RPM Fibre Disk	2.4	4.8	41%
	31.6	121.7	96%
Staff	2 Partial	1 Partial	-29%

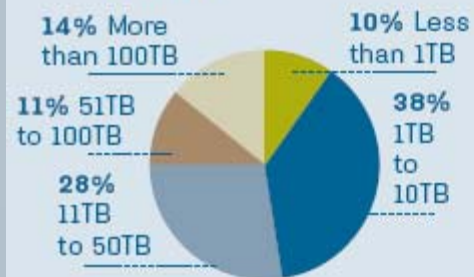
# Storage Magazine 2006 Purchasing Survey

## 1. WHERE YOUR BUDGET DOLLARS GO

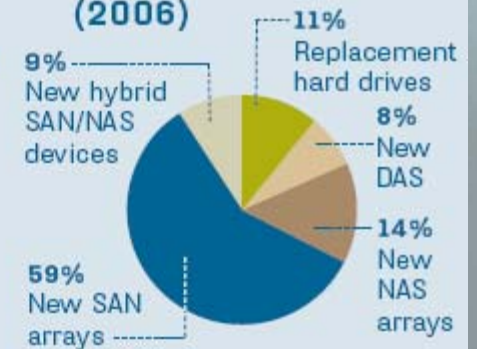


\* May not total 100% due to rounding

## 2. PLANNED DISK STORAGE PURCHASES\* (2006)



## 3. PRIMARY DISK STORAGE EXPENDITURES\* (2006)



# How Disks Work



# The Evolution of Disk Hardware

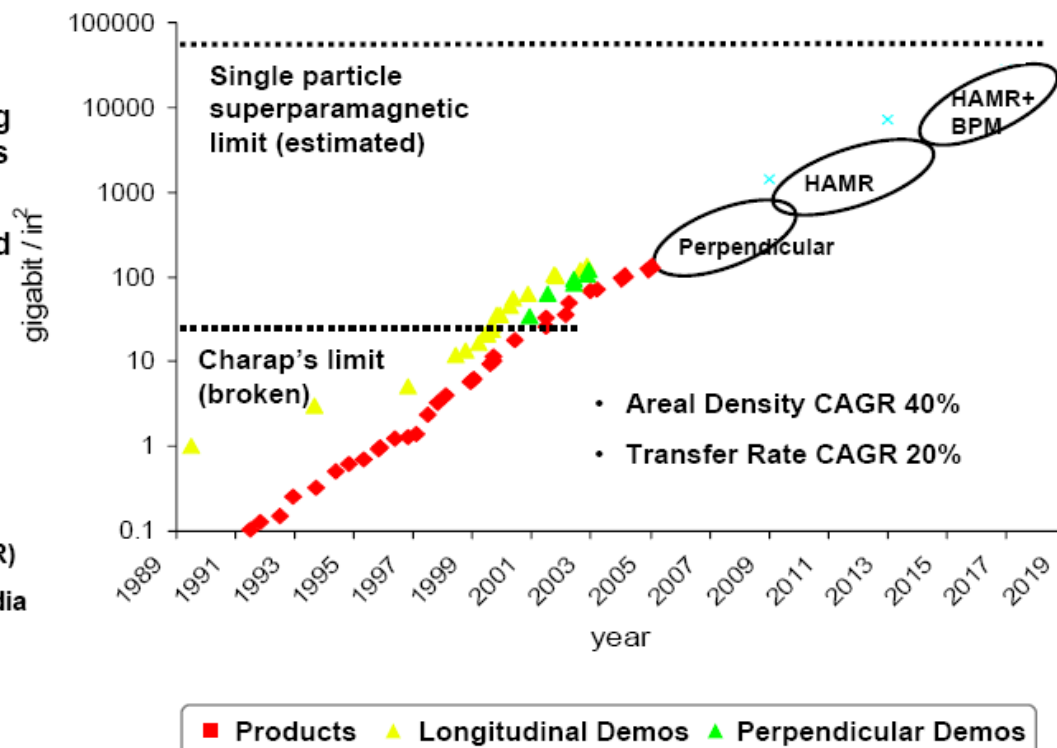
## Technical Specifications – Then and Now

	IBM RAMAC (1956)	Seagate Momentus (2006)	Delta
Capacity	~ 5 MB	160 GB	15,000 x
Areal Density	2000 bits/in <sup>2</sup>	130 Gb/in <sup>2</sup>	65,000,000 x
Disks	50 @ 24" dia	2 @ 2.5" dia	240 x
Price	\$50,000	<\$500	x / 100
Price/MB	\$1,000	<\$0.01	x / 320,000
Spindle Speed	1,200 rpm	5,400 rpm	5 x
Seek Time	600 ms	10 ms	x / 60
Data Rate (burst)	10 KB/s	44 MB/s	4,400 x
Power	5000 W	2 W	x / 2,500
Weight	~ 1 ton	4 oz	x / 9,000

# Where's Areal Density Going?

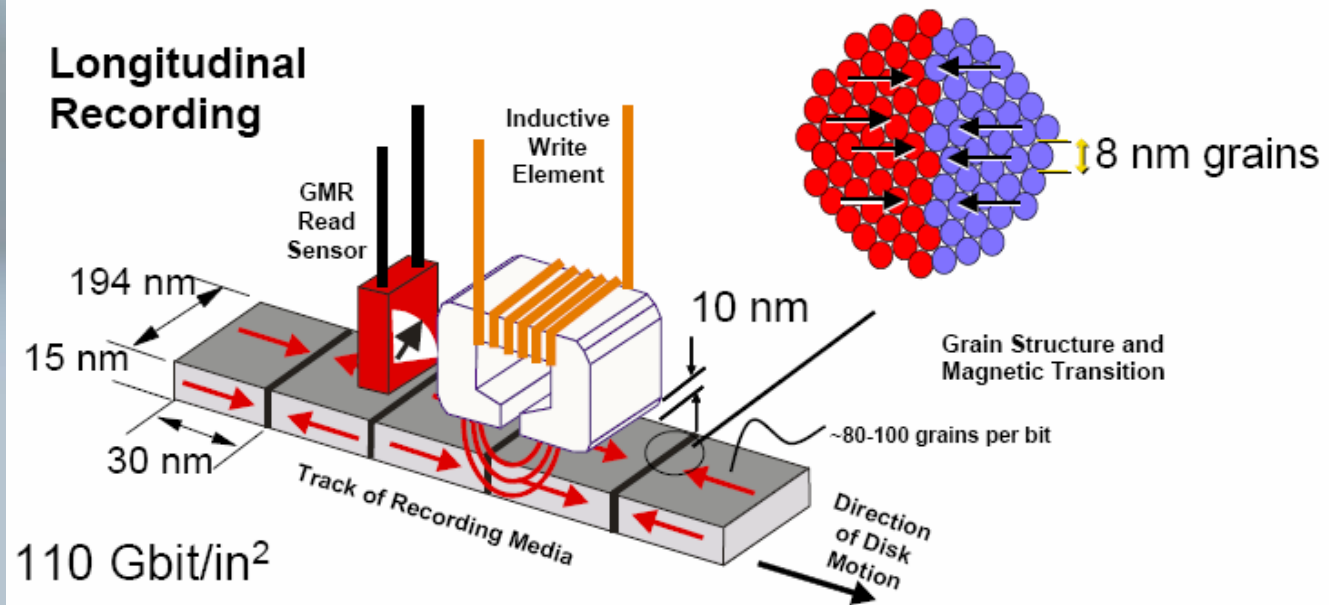
## Areal Density Growth

- Late 1990s – super paramagnetic limit demonstrated through modeling
- Longitudinal recording reaching areal density limits
- Perpendicular expected to extend to 0.5-1 Tb/in<sup>2</sup>
- Additional innovations required at that point
  - Heat-Assisted Magnetic Recording (HAMR)
  - Bit Patterned Media (BPM) recording



# What's a bit really look like?

## Recording Basics

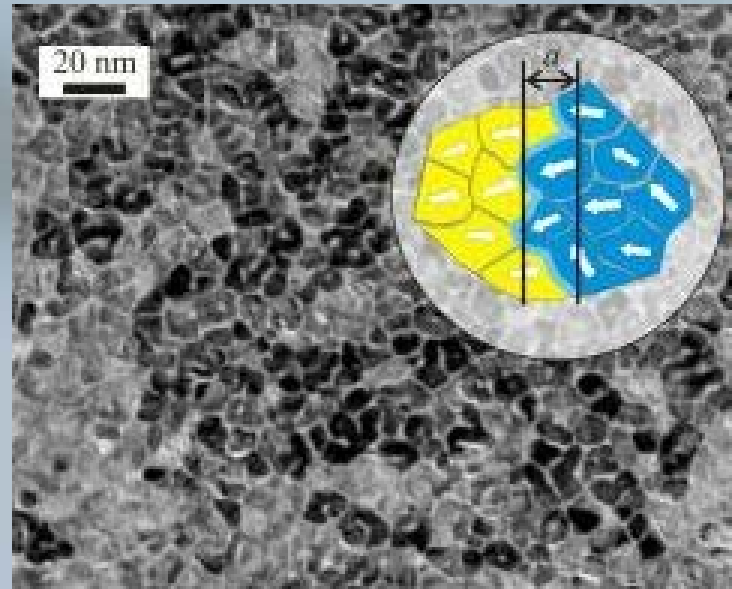


110 Gbit/in<sup>2</sup>

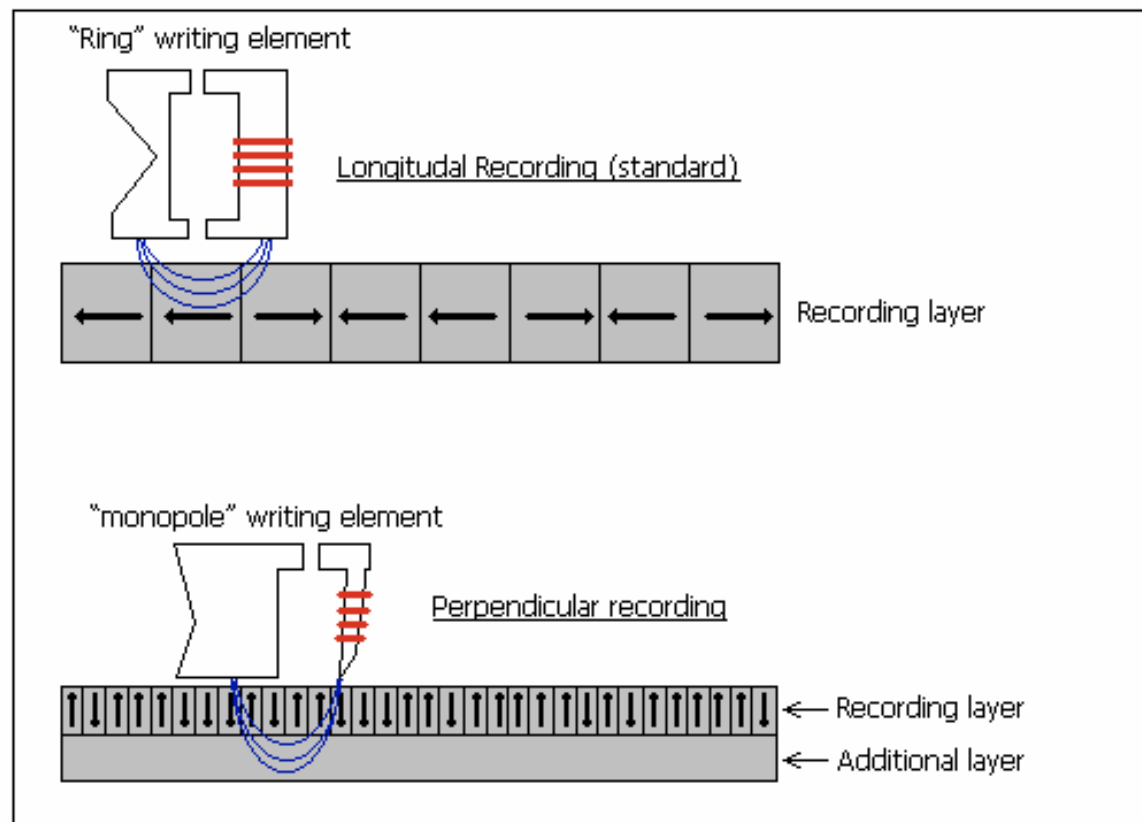
(Gigabits per square inch)

- Bit dimensions 30 x 194 x 15 nanometers
- Head-media spacing of 10 nanometers
- For 1 Tbit/in<sup>2</sup>, bits will need to be 3x smaller in width and length

# What's a bit really look like?

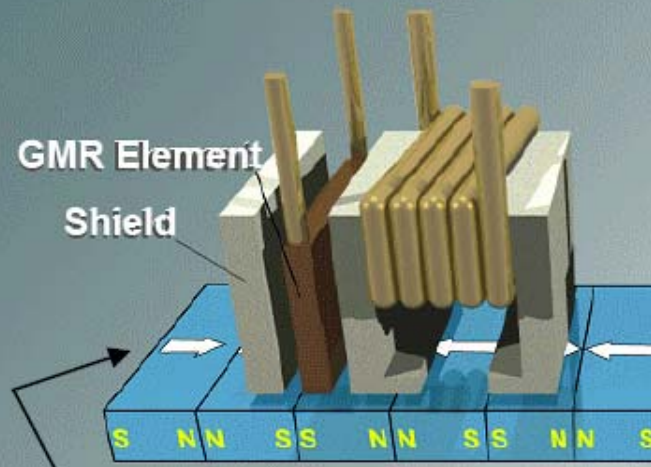


# Longitudinal vs. Perpendicular



# Longitudinal vs. Perpendicular

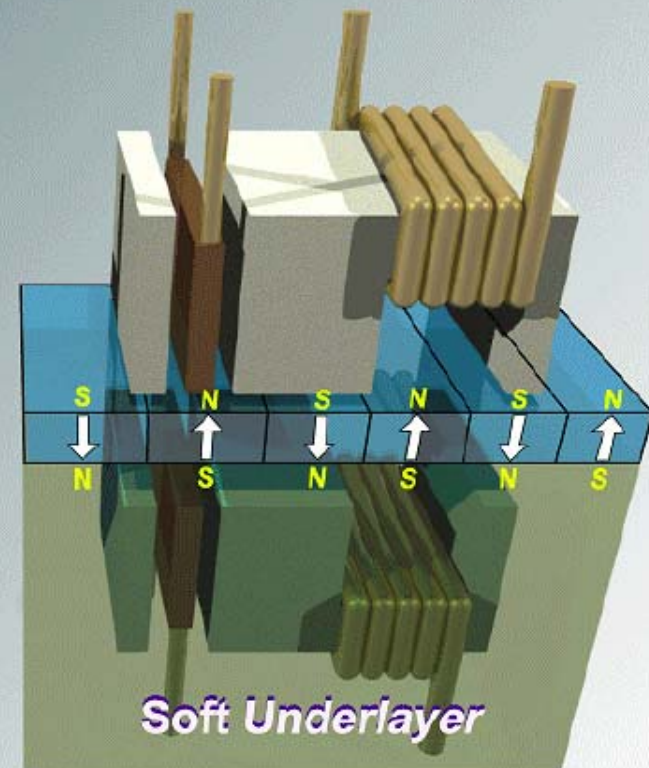
## Longitudinal Recording



Magnetic domains oriented in the direction of travel of the head.

Soft underlayer "mirrors" write head and makes it possible to write domains much closer together.

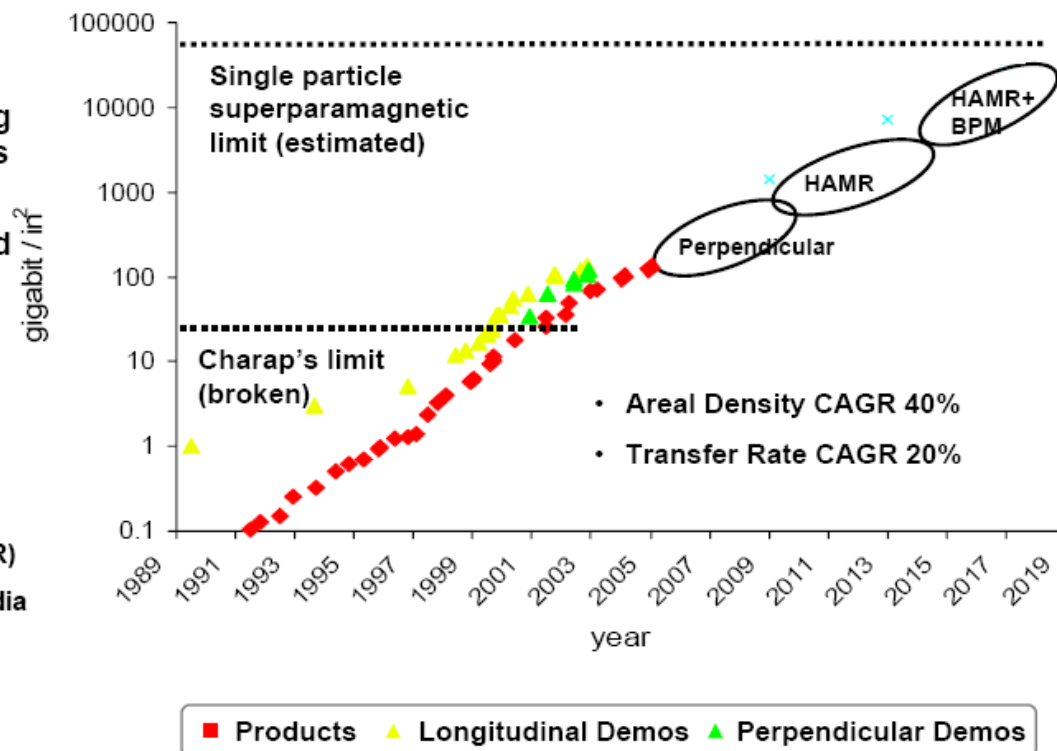
## Perpendicular Recording



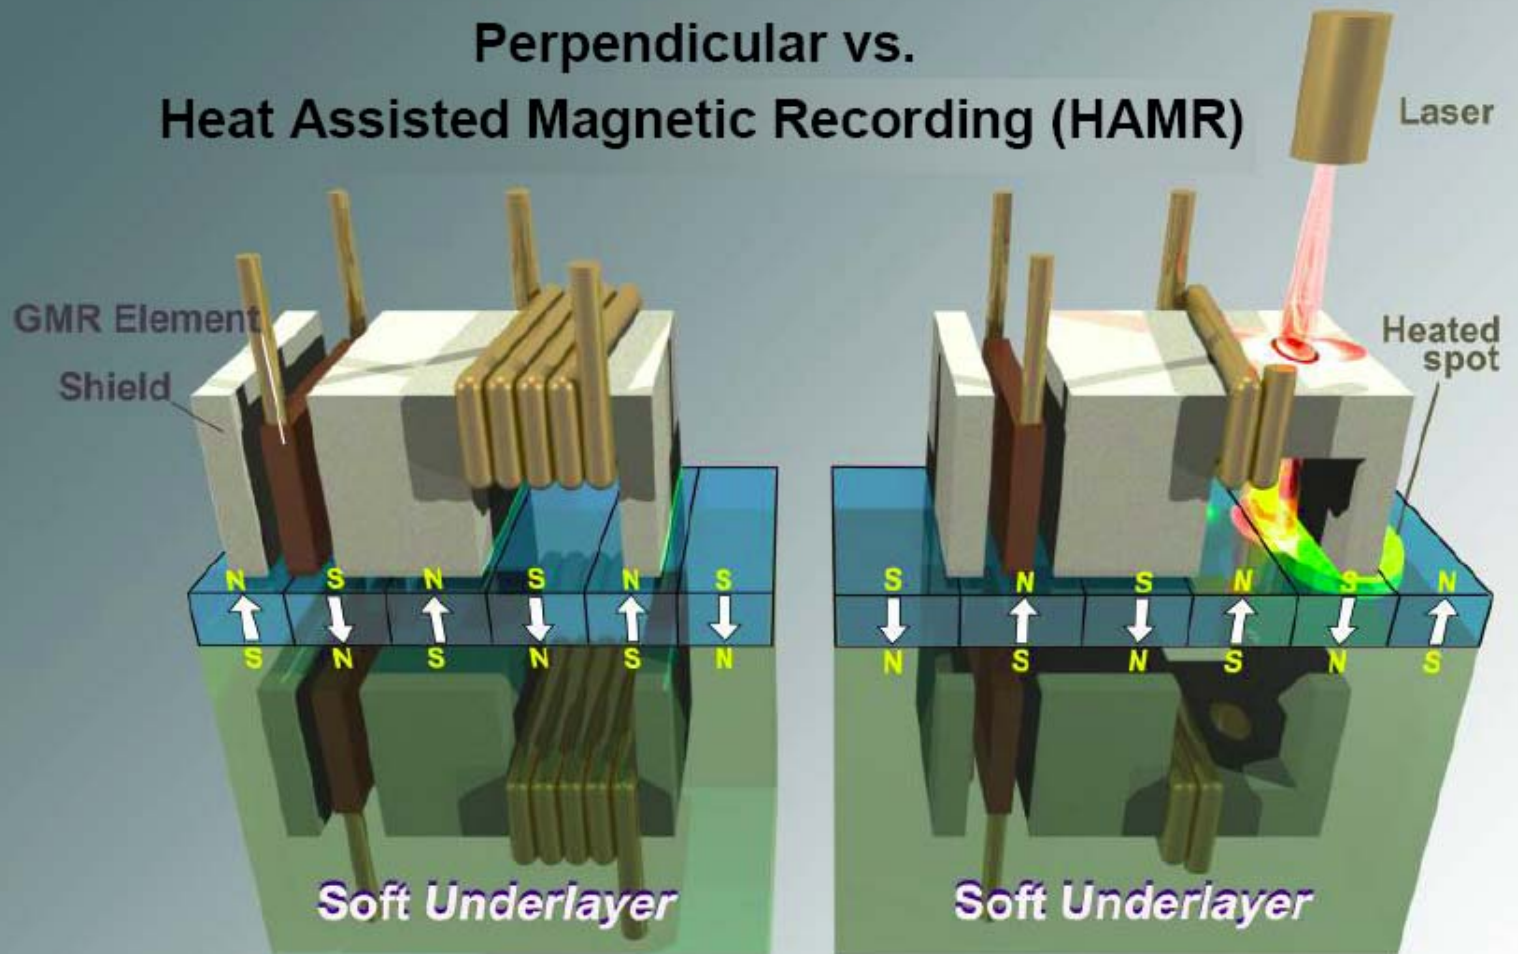
# Where's Areal Density Going?

## Areal Density Growth

- Late 1990s – super paramagnetic limit demonstrated through modeling
- Longitudinal recording reaching areal density limits
- Perpendicular expected to extend to 0.5-1 Tb/in<sup>2</sup>
- Additional innovations required at that point
  - Heat-Assisted Magnetic Recording (HAMR)
  - Bit Patterned Media (BPM) recording



## Perpendicular vs. Heat Assisted Magnetic Recording (HAMR)



HAMR can theoretically extend areal density beyond 5 Tbps

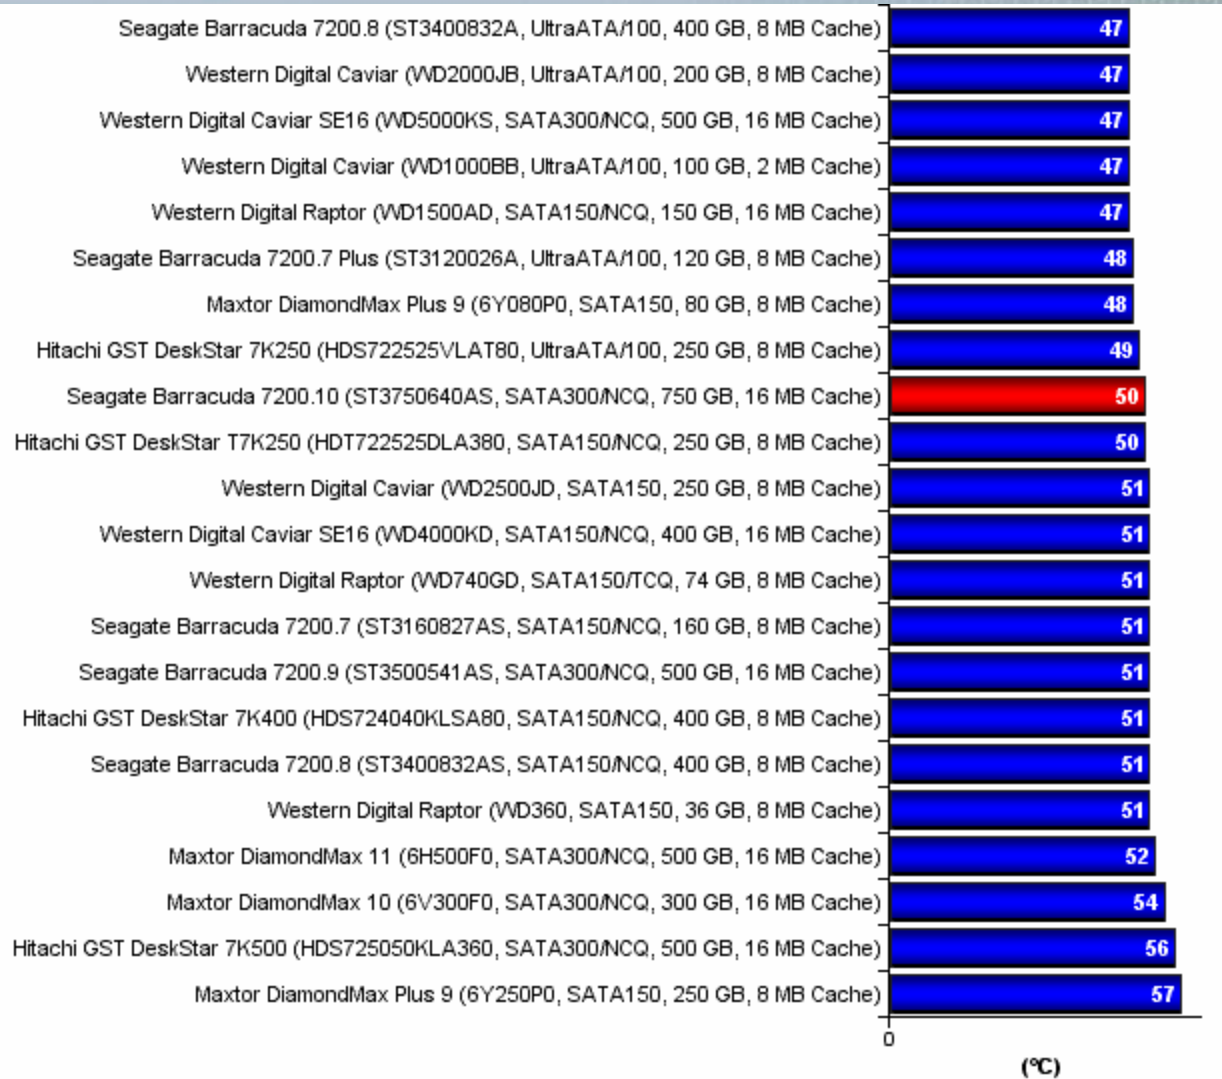
# So what does this all mean for disks in the near term?

3.5 inch Nearline	2005 (Longitudinal)	2009 (Perpendicular)	2013 (HAMR)
Drive Capacity (GB)	500	2,000	8,000
Number of Discs	3	3	3
Capacity (GB/disc)	168	670	2,670
Product Areal Density (Gbps)	120	500	1,800
Transfer Rate (Mb/sec)	995	2,000	5,000
RPM	7,200	7,200	10,000
Read Seek Time (ms)	8	7.2	6.5
2.5 inch Enterprise	2005	2009	2013
Drive Capacity (GB)	75	300	1,000
Number of Discs	2	2	2
Capacity (GB/disc)	40	150	500
Product Areal Density (Gbps)	70	300	1,000
Transfer Rate (Mb/sec)	750	2,000	4,000
RPM	10,000	15,000	15,000
Read Seek Time (ms)	4.7	3.8	3.1
1.0 inch Handheld	2005	2009	2013
Drive Capacity (GB)	8	30	100
Number of Discs	1	1	1
Capacity (GB/disc)	8	30	100
Product Areal Density (Gbps)	80	300	1,200
Transfer Rate (Mb/sec)	110	200	600
RPM	3,600	4,200	5,400
Read Seek Time (ms)	16	14.5	13.1

# Unanticipated Consequences

- Gartner Predicts – “By 2008, nearly 50% of data centers worldwide will lack the necessary power and cooling capacity to support high-density equipment.”
- Have you walked past your SAN storage arrays lately?

# Typical Modern 3.5" Disk Drive Temperatures



Source – Tom's Hardware @ <http://www.tomshardware.com/storage.html>

# What's the storage buzz

- SMI-S 1.2 & Aperi
- CDP – Continuous Data Protection
- Virtualization – Don't Laugh!
  - Heterogeneous virtualization
  - Switch based virtualization
  - HBA virtualization
- Global Name Spaces (File Virtualization)
- Clustered File Systems / Clustered Storage
- Archival File Systems (Archivas, Permabit)
- Database Archiving
- Wide Area File Systems

# What's the storage buzz

- OSD – Object-based Storage Devices
- TPM (Trusted Platform Module) in Storage Devices
- Solid State & Hybrid
  - Flash based (44 SSD manufacturers today)
  - MRAM (Magnetoresistive RAM)
  - Change Resistive, e.g. polymer
  - Nano-scale Mechanical (“milipede”)
- Intelligent Storage Grids & Storage Autonomics

# Regulatory Effects on Storage

- New Federal Rules for Civil Procedures (FRCP) causing much FUD

“U.S. companies will need to keep track of all the e-mails, instant messages and other electronic documents generated by their employees thanks to new federal rules that go into effect Friday, legal experts say.” – AP

“The rules also mean that colleges that are in litigation **or that suspect they may soon be in litigation** cannot destroy electronic evidence they know would be relevant to a lawsuit.” – The Chronicle of Higher Education

# The Case for an Architectural Approach to Storage

- It's all about the data!
- Storage continues to represent a significant (and probably growing) portion of your budgets.
- The technology landscape is perhaps a bit more complex than many of us think.
- The regulatory environment requires monitoring and strategic thought.
- Storage is not just about technology. There are policy and process issues to consider.
- The cost of getting it wrong may be significant.

# The Case for an Architectural Approach to Storage

Use case – VoIP and Unified Messaging

- Talk about unstructured data! What about all those little WAV files now flying around our networks, being stored in our e-mail systems?
  - What the impact of the new FRCP?
  - What's the impact on our backup systems, retention policies, DR & BC?
  - How will user behavior change as the result of introducing technology, and how will those new behaviors affect your storage plan?

# In Summary

- Storage Technology still rapidly evolving, and not just faster, better, cheaper but new and different
- Storage Requirements also rapidly evolving
- Need to have a storage strategy