Pulse
Comes to You 2009

IBM
Managing the World's Infrastructure

Green Shoots:
What is emerging in Storage Technology

Rick Terry, Information Infrastructure Evangelist
Steve Legg, Chief Technology Officer Systems & Technology Group

INDUSTRIAL REVOLUTION

IRON AGE

AGE OF THE MICRO-PROCESSOR

INTERNET REVOLUTION

This is “Smarter Planet”

Everyone, Everything, Everywhere becoming connected in a joined up technological world

Everything digital – from before the cradle to beyond the grave

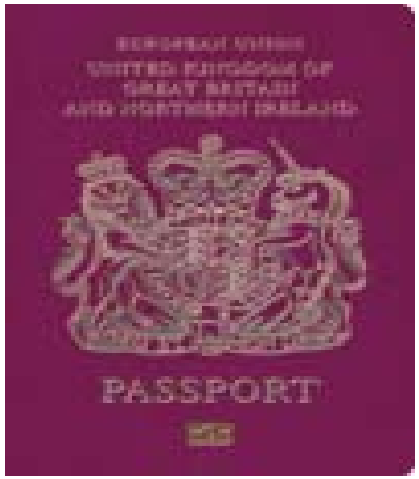

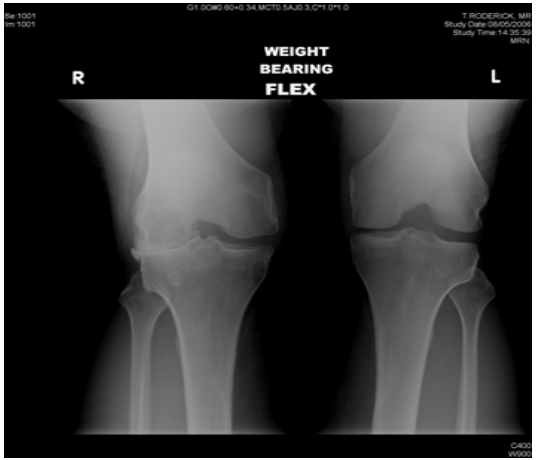

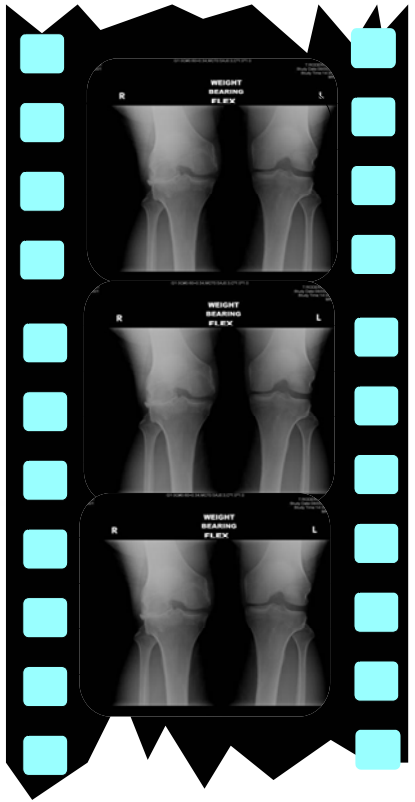

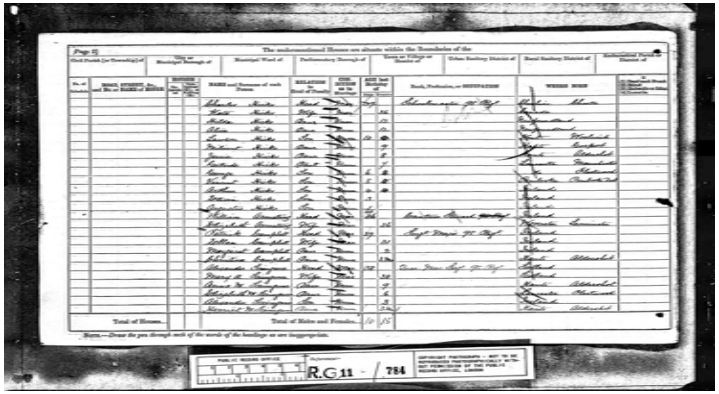

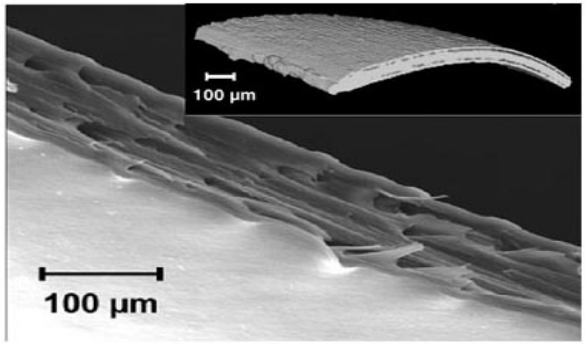

The Need for a Dynamic Infrastructure

Traditional Data Centers

Web 2.0 Data Centers

The Globally Integrated Enterprise

Business-to-consumer uber-portals

ebay

amazon.com

Google

Building the infrastructure for a Smarter Planet

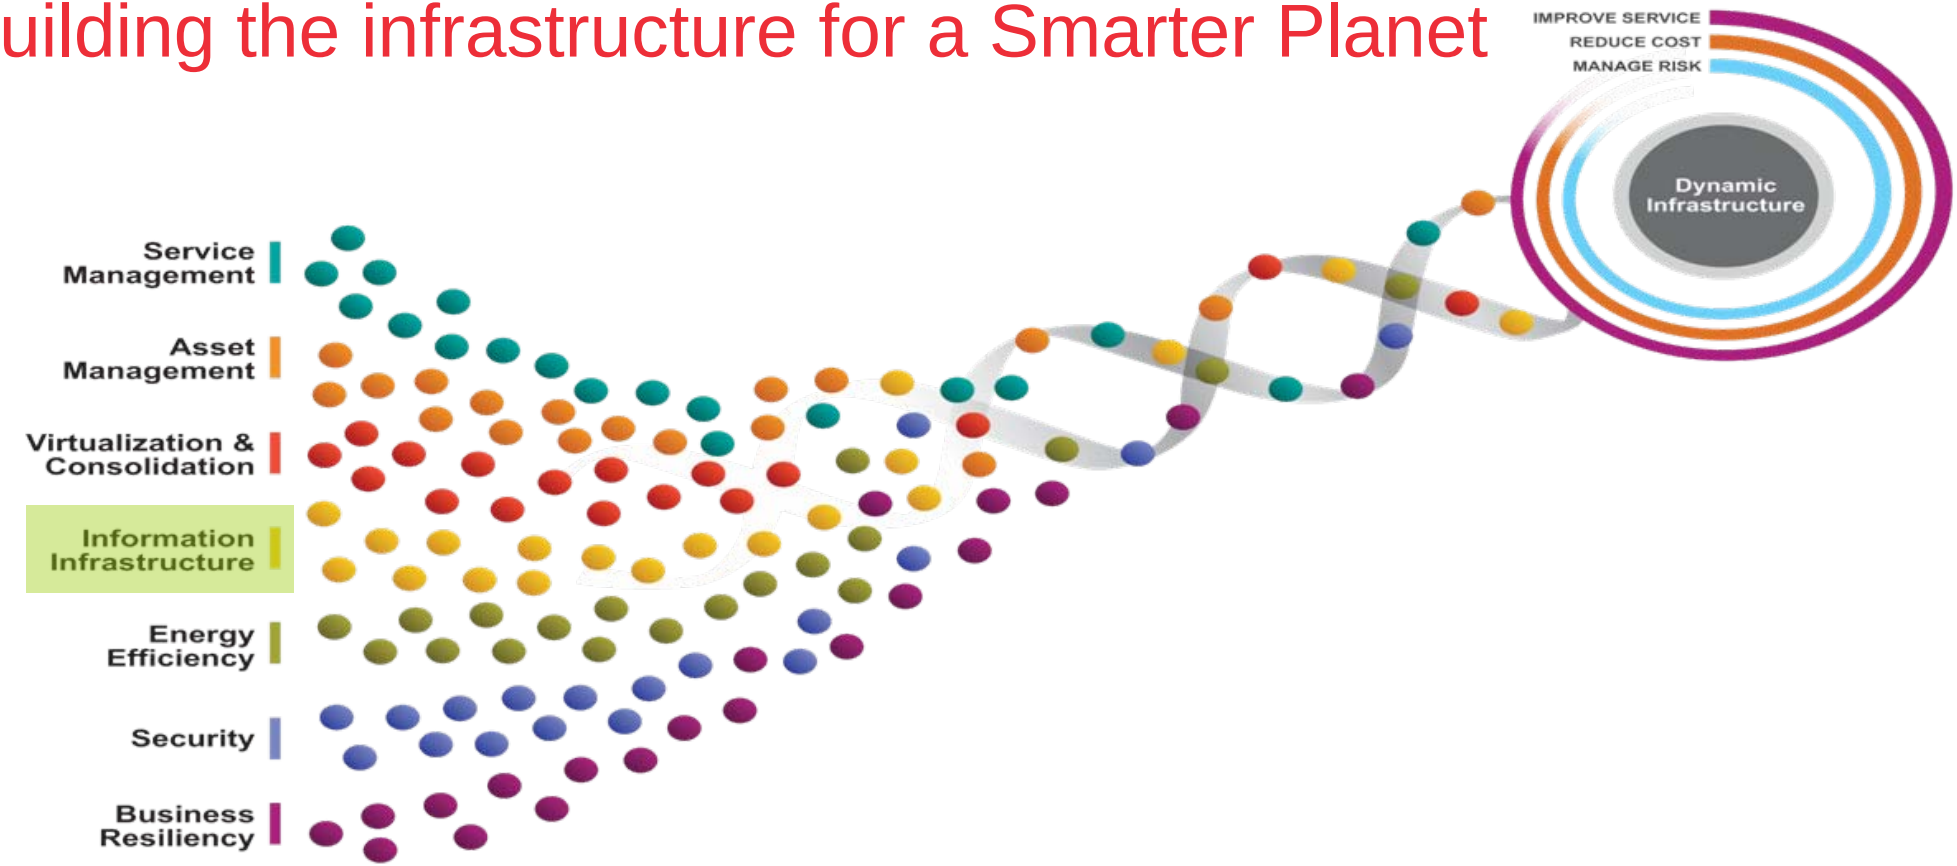

Why Information Infrastructure

Reduce Cost
Optimise IT to Save/Make Money

Mitigate Risk
stay out of jail while staying in business

Improve Service
plan to keep going in a changing world

Availability

Deliver continuous and reliable access to information

Security

Protect and enable secure sharing of information

Retention

Support your information retention policies

Compliance

Reduce reputation risks and audit deficiencies

Compliance

- Law Courts
- Parliament
- Regulatory Bodies

Availability

- 1990's – Have you got the data *somewhere else*?

  Disaster Recovery

- After 9/11 – Have you got the *same data* somewhere else?

  Business Continuity

- In an "On Demand World" – I want the data *where I am*, NOW!

  True Availability

Availability plus Compliance =

Availability plus Compliance = Retention

Availability plus Compliance = Retention

Up to 4 yrs for Phone Traffic Data

Availability plus Compliance = Retention

5 years for Anti Money Laundering purposes

3-6 years for Tax, VAT, PAYE, etc

Availability plus Compliance = Retention

Up to 40 years for exposure to hazardous materials

Availability plus Compliance = Retention

50 years for ionizing radiation

Availability plus Compliance = Retention

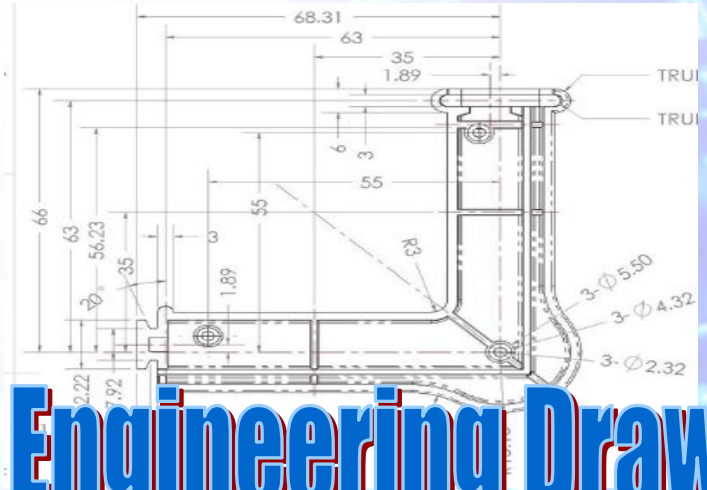

Engineering Drawings /Plans~effectively Forever

Availability plus Compliance = Retention

- Balancing business benefit AND risk
- Ongoing Environmental Cost
- Potential Fines &/or Loss of Reputation or customers

DATA PROTECTION ACT 1998

"as long as necessary"

Availability plus Compliance = Disclosure

DATA PROTECTION ACT 1998

40 Days for Subject access
28 Days for Information Notices

FREEDOM OF INFORMATION

20 Days for General Access

2 working Days
or
48 hours

1 working Day

- No discernible business benefit
- Fines &/or Loss of Reputation
- Ultimate sanctions

What is a record?

YH 796971

B. Cert. S.
R.B.D.

1 & 2 ELIZ. 2 CH. 20

CERTIFICATE OF BIRTH

Name and Surname.....

Sex.....

Date of Birth.....

Place of Birth { Registration District..... Sub-district.....

I, GWYNNETH G. POWELL Registrar of Births and Deaths for the Sub-district of.....do hereby certify that the above particulars have been compiled from an entry in a register in my custody.

Witness my hand this.....day of.....19.....

CAUTION :—*Any person who (1) falsifies any of the particulars on this certificate, or (2) uses a falsified certificate as true, knowing it to be false, is liable to prosecution.*

Gwynneth G. Powell

Registrar of Births and Deaths.

Verifiable

Immutable

Authentic

Security

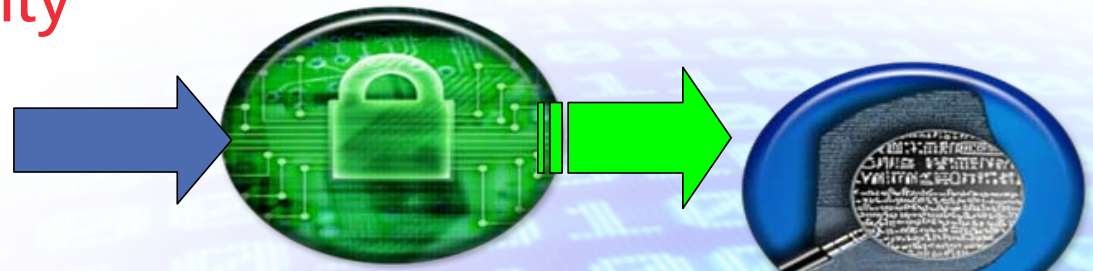

- How does an organisation prove authenticity with electronic records?
  - Policies
  - Technologies
- How does an organisation provide protection for the media?
  - Encryption

- ~100 losses in 6 mths after HMRC

- ISO 27001 - the Information Security Management System (ISMS) certification standard.
- It's not just about fines and jail sentences …but also reputation

So, what's happening to YOUR business?

Hyper Data Growth

Data created and copied expected to grow at
57% CAGR through 2010

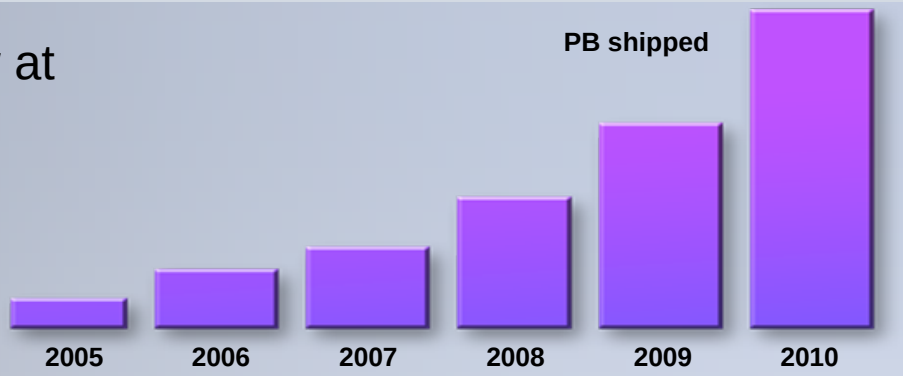

Hyper Data Growth

Data created and copied expected to grow at
57% CAGR through 2010

NOT GOING TO STOP

Today, only 20% of data is structured, with unstructured content coming in the form of email, other documents, images, and video

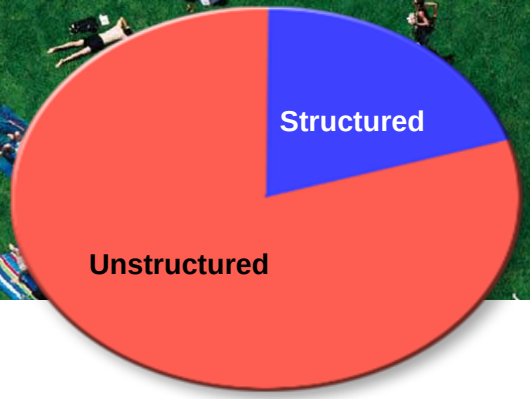

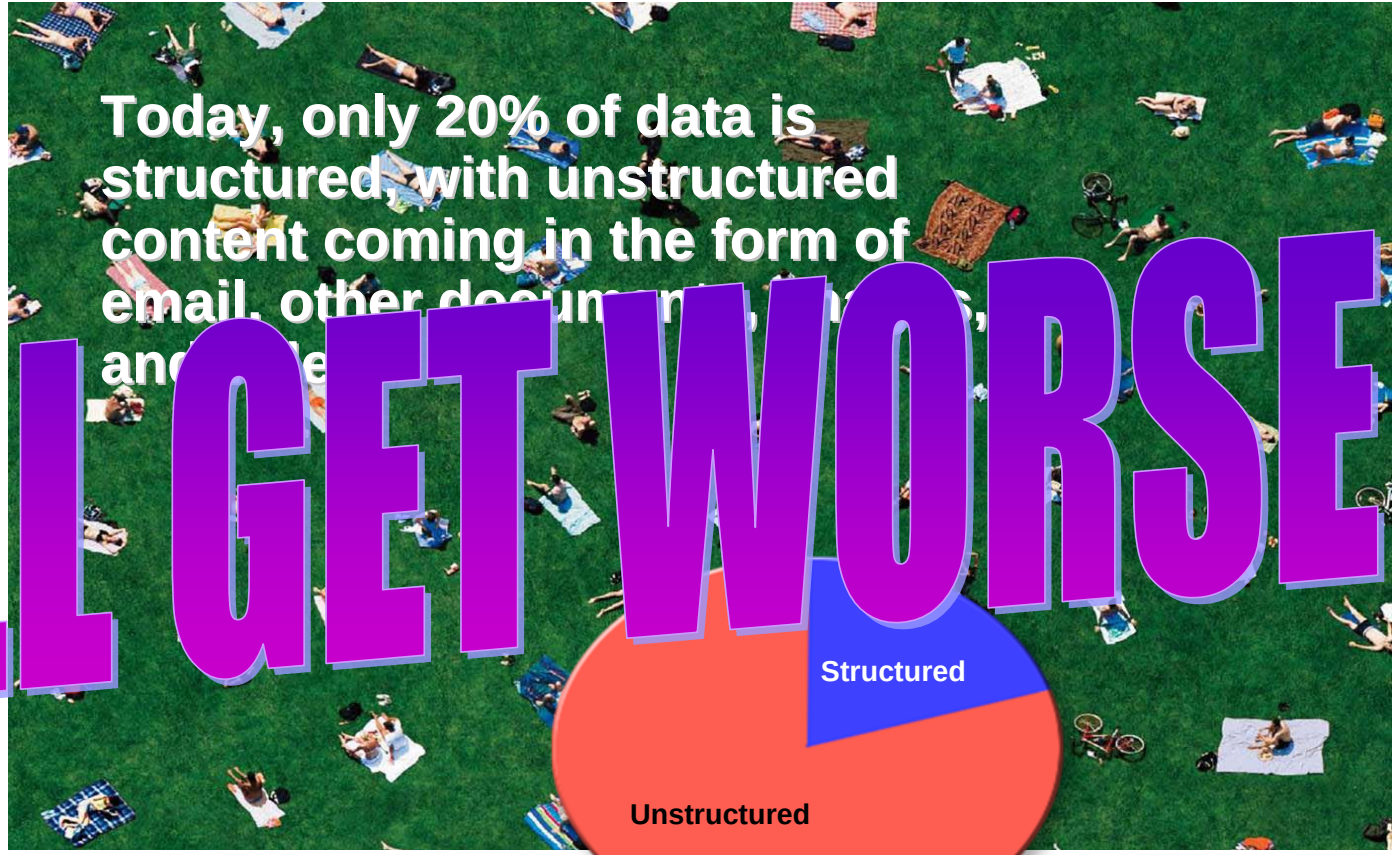

IT WILL GET WORSE

“Ye canna change the laws of physics Cap’n”

Trends: Magnetic Disk Technology

Trends: Magnetic Disk Technology

4,700,000 bytes

Trends: Magnetic Disk Technology

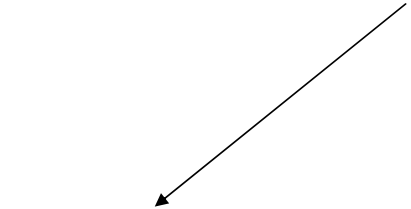

1,000,000,000,000 bytes

Trends: Magnetic Disk Technology

Trends in the physics of magnetic recording

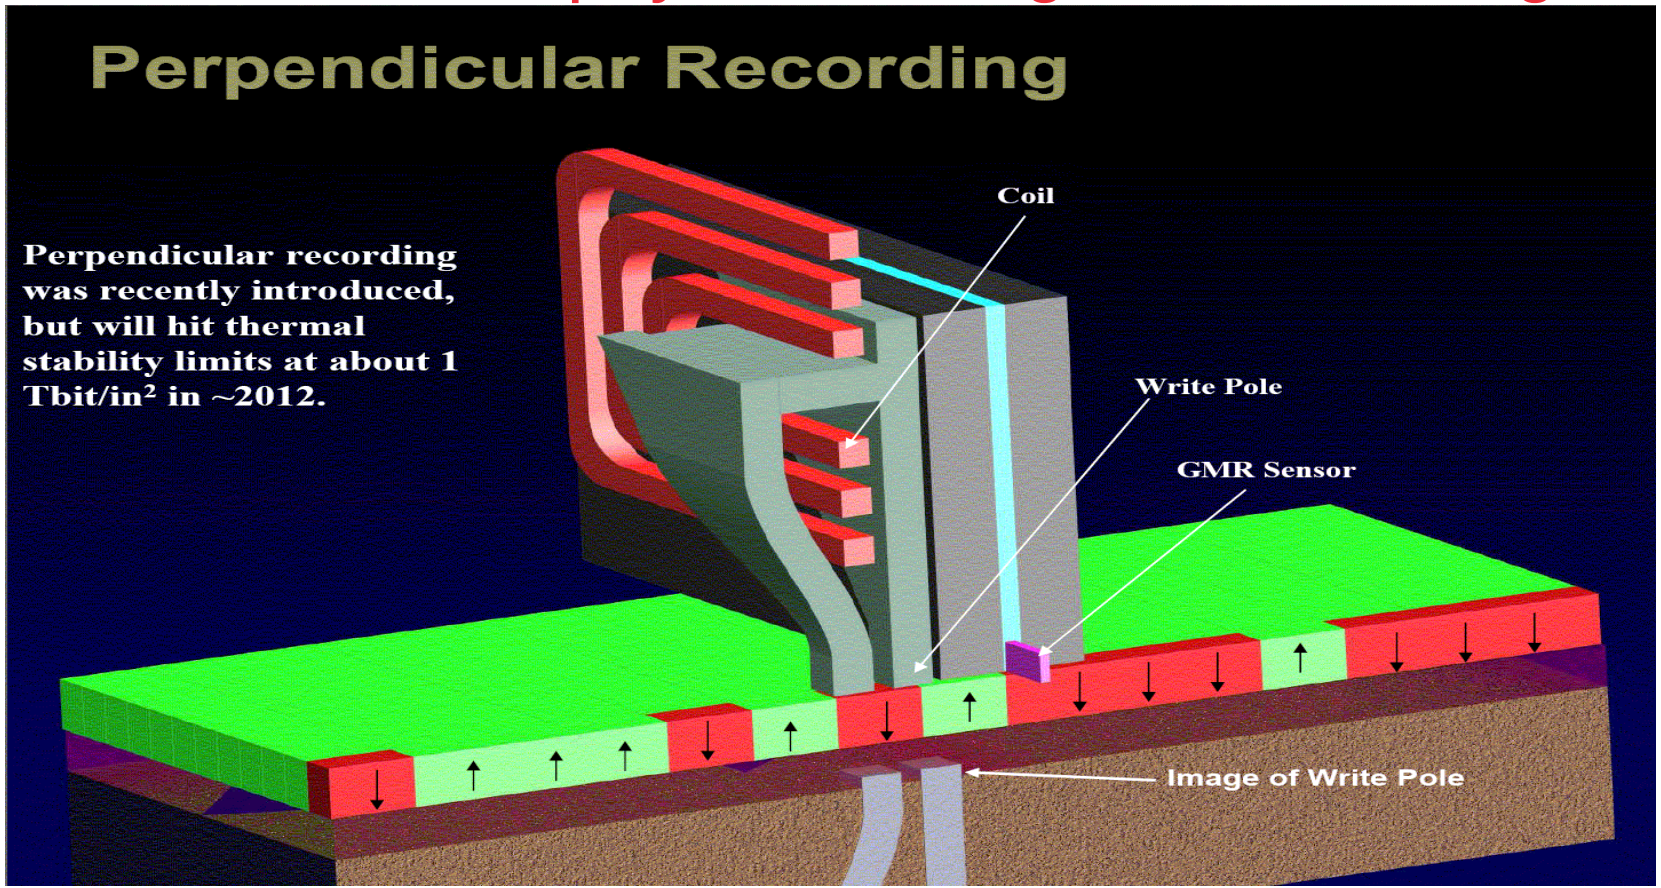

Trends: Magnetic Disk Technology history

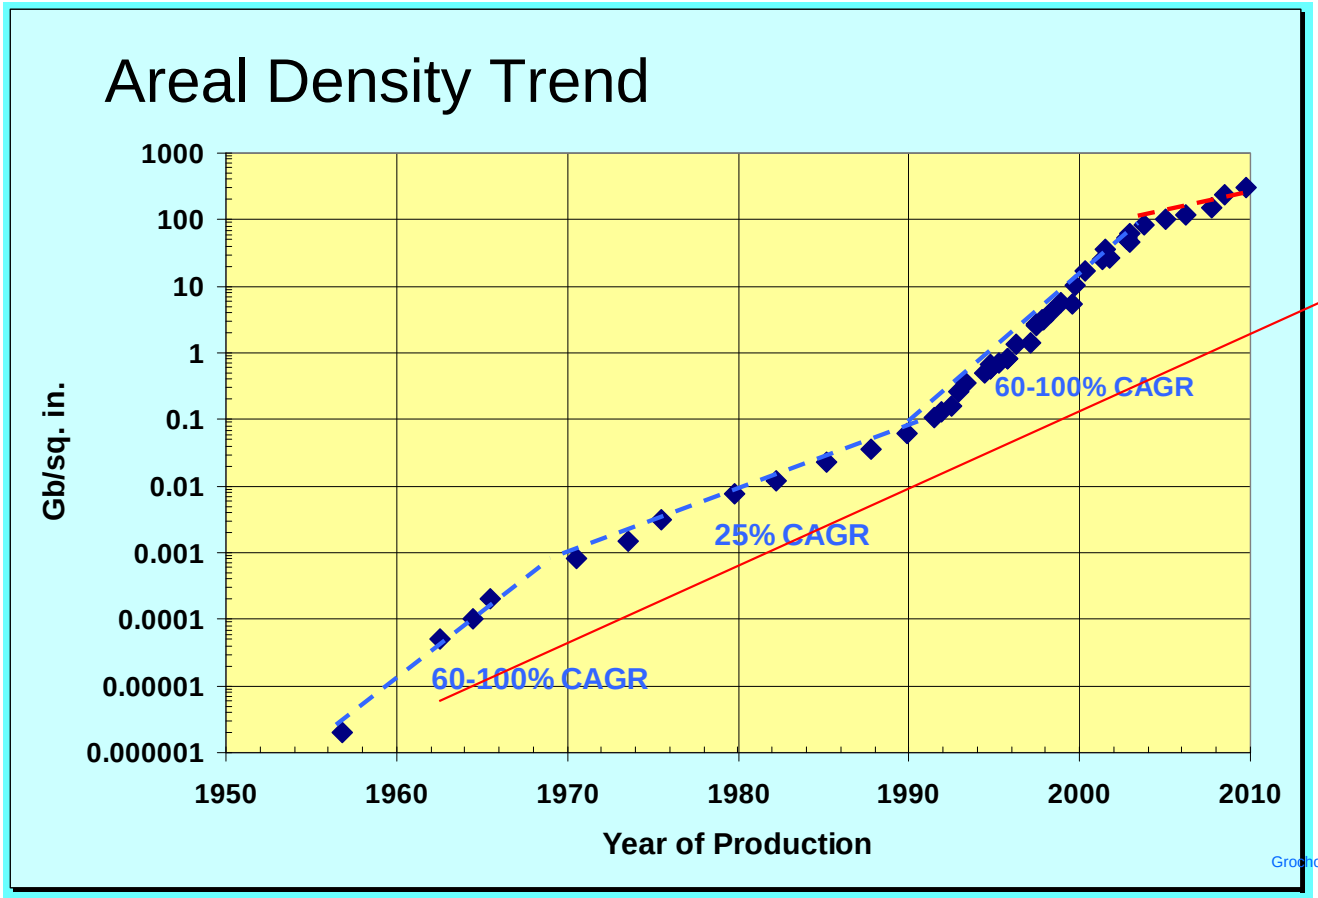

…looking a bit further forward

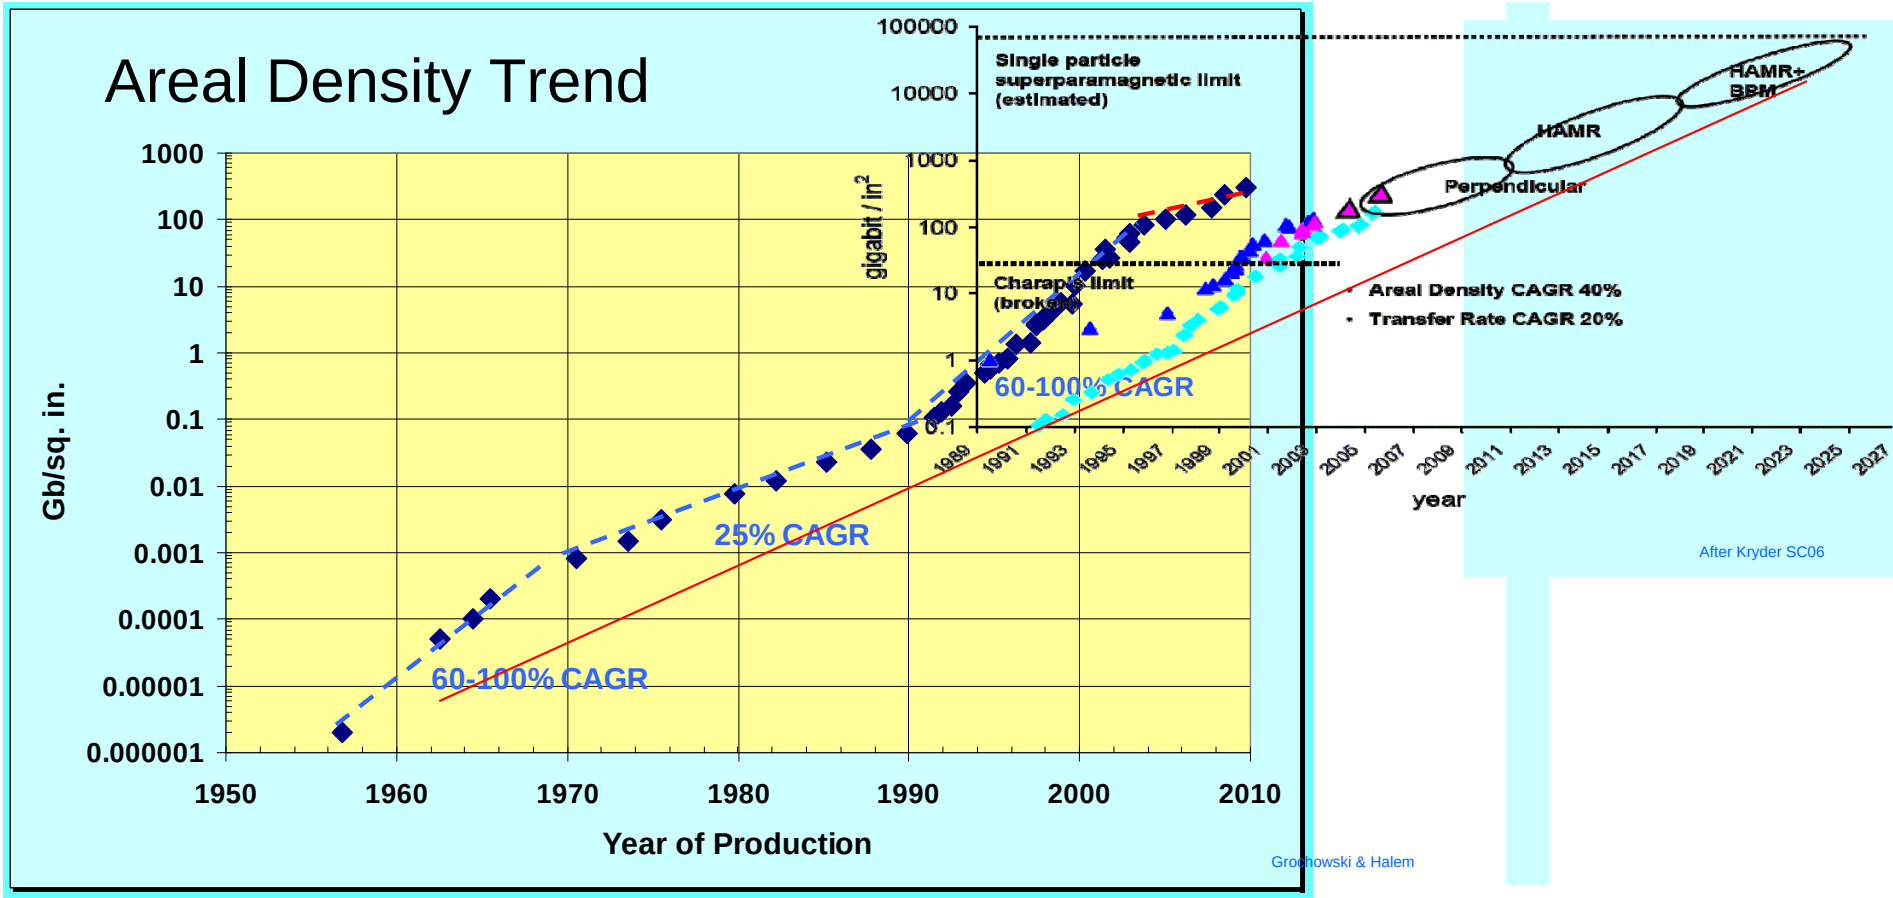

Trends: Technology Commoditisation

- **Functionality**
- **Availability**
- **Disaster tolerance**
- **Scalability**
- **Manageability**

Emerging Workloads

Searchable Archives

Grid File Serving

Web 2.0

IPTV

Medical Imaging

Digital Video Surveillance

Video Serving

File Serving

Next Gen. Networks

SMB

- **Low cost/commodity**
- **Modular**
- **New data types**
- **Innovation/business value**

Data Warehousing

Business Intelligence

Data Backup & Protection

Automated Management

Storage Security

Transaction Processing

Disaster Recovery

HPC

"Classic" Workloads

Trends: Storage Power Consumption

Despite improved efficiency, energy consumption is growing…

Storage Power Consumption/GB

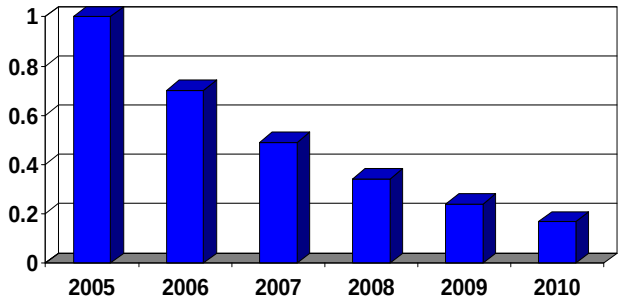

Source: IBM

Data Center Storage - Petabytes

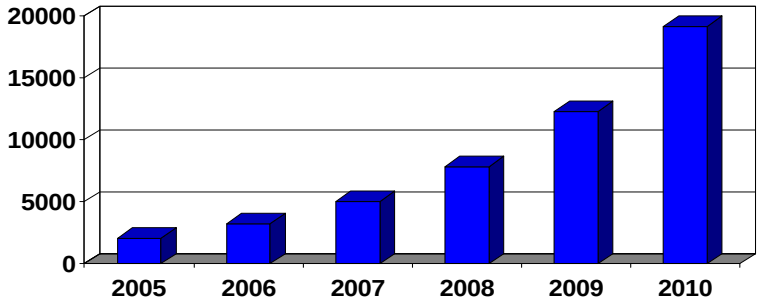

Data Center Storage Power Growth

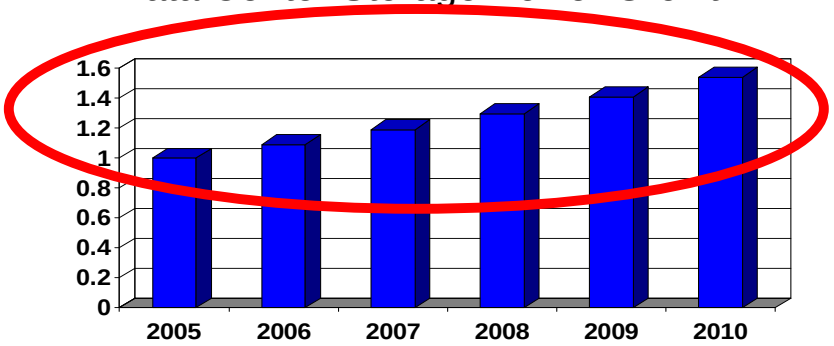

Source: IBM

Storage Power Consumption – in the Data Centre

Is it a GREEN issue?

Data centre disruption 'to hit UK in five years'

SOURCE: Rakesh Kumar, Gartner VP of Research, October 2007

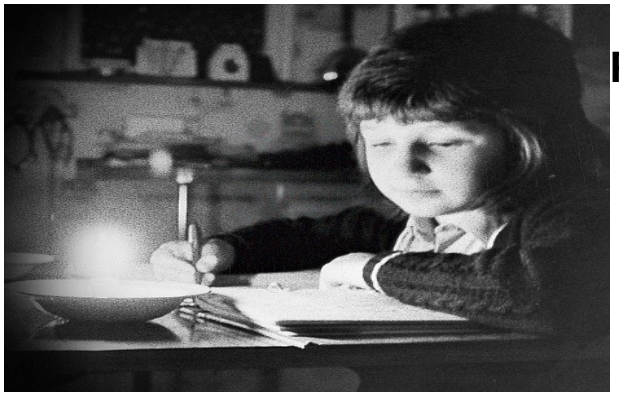

Regular blackouts to hit Britain within three years because there is a shortage of new power stations, CapGemini study claims

November 24th 2008

Trends in energy efficiency

STOP PRESS…“abc disk uses *less* electricity than brand xyz”…

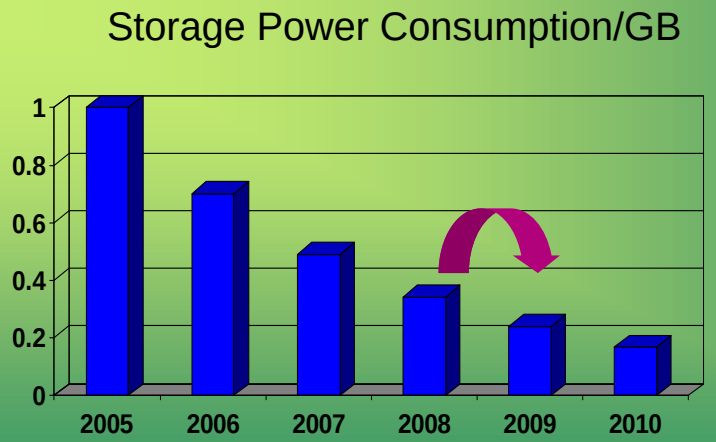

Average Storage Cost Trends

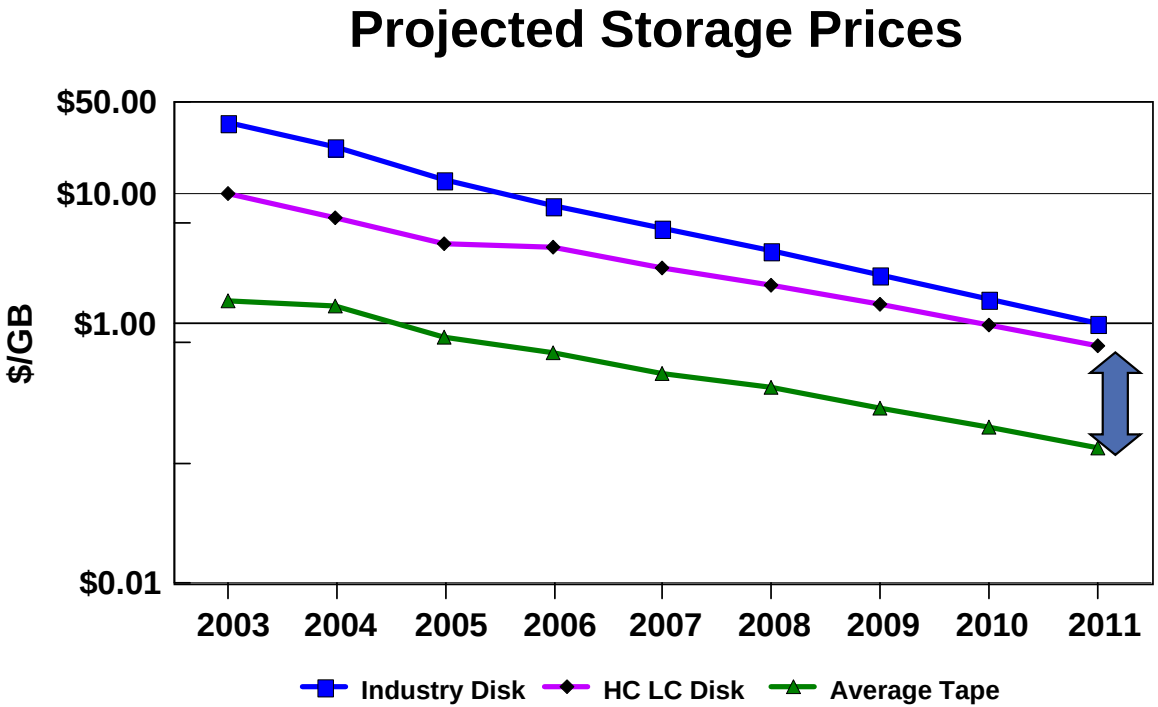

What are we trying to optimise?

- £ per GB or TB
- £ per IOPS
- kW per TB
- $kgCO_2$ per TB
- FTE per TB

- New style arrays
- Solid State Disks
- Low energy options
- Footprinting
- Virtualisation

Pulse

Comes to You 2009

Pulse
Comes to You 2009

IBM
Managing the World's Infrastructure

Responsible Reactions
IBM's Information Infrastructure Strategy

Rick Terry, Information Infrastructure Evangelist
Steve Legg, Chief Technology Officer Systems & Technology Group

How do you control that raging teenager in us all?

How do you control that raging teenager in us all?

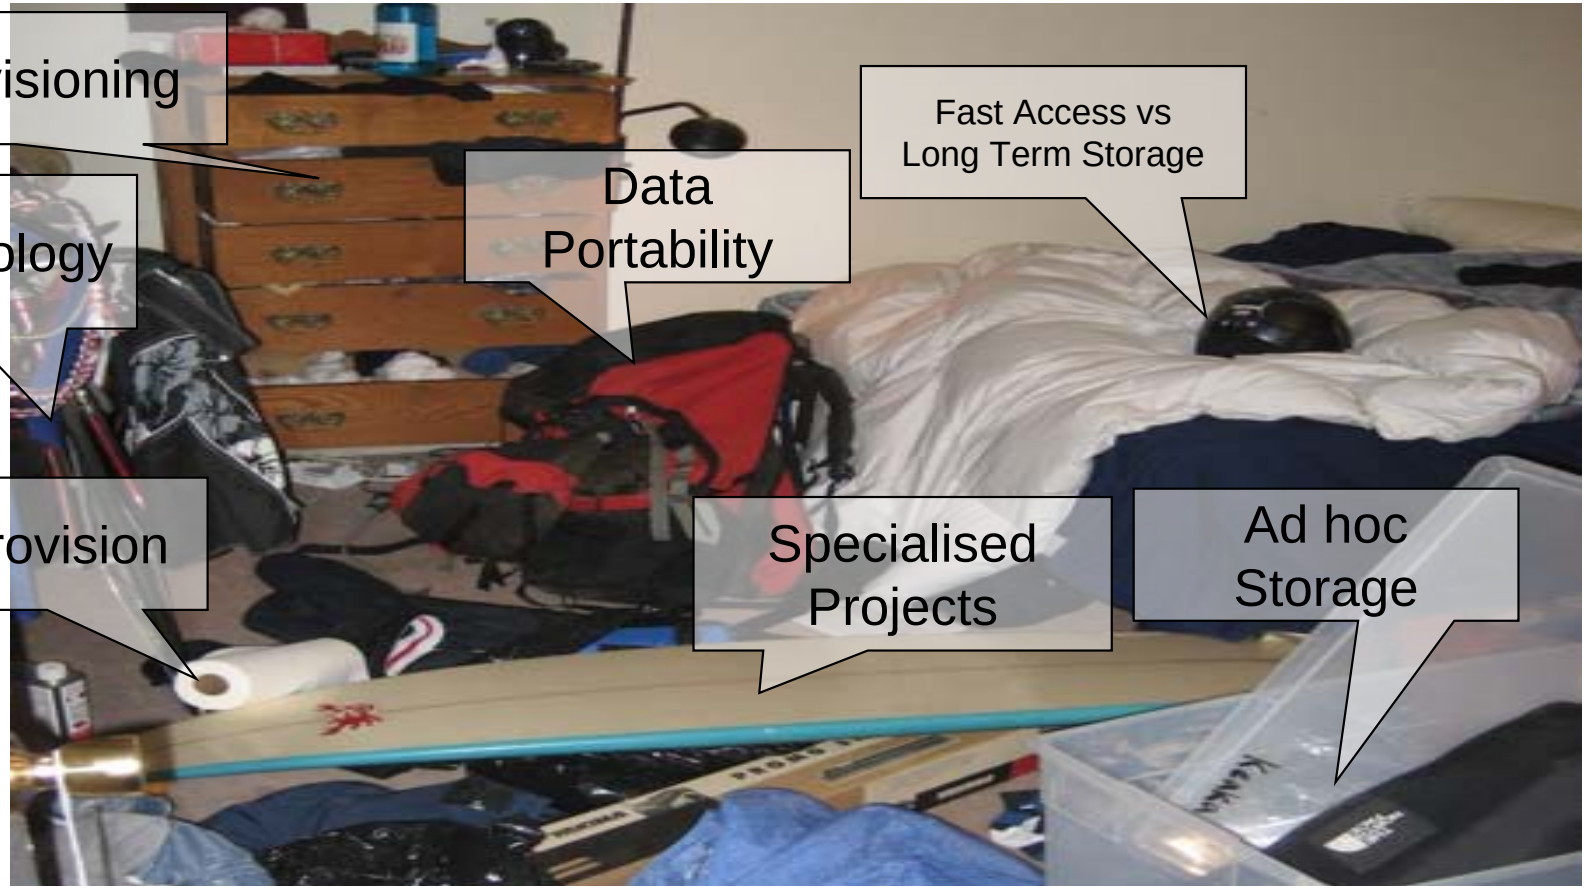

Persistent Data, Ephemeral Value

Ephemeral: an adjective to refer to a fast deteriorating importance or temporary nature of an object to a person.

The majority of the data sitting on a "production" array is non-transactional, or post-transactional.

Even worse, at least half the data is duplicated

Steve Duplessie
Enterprise Strategy Group

Persistent Data, Ephemeral Value

6th May 1840

worth 1 penny

7th May 1840

worthless

19th May 2009

worth up to £2,275

2d Blue….

£9,000!

Persistent Data, Ephemeral Value

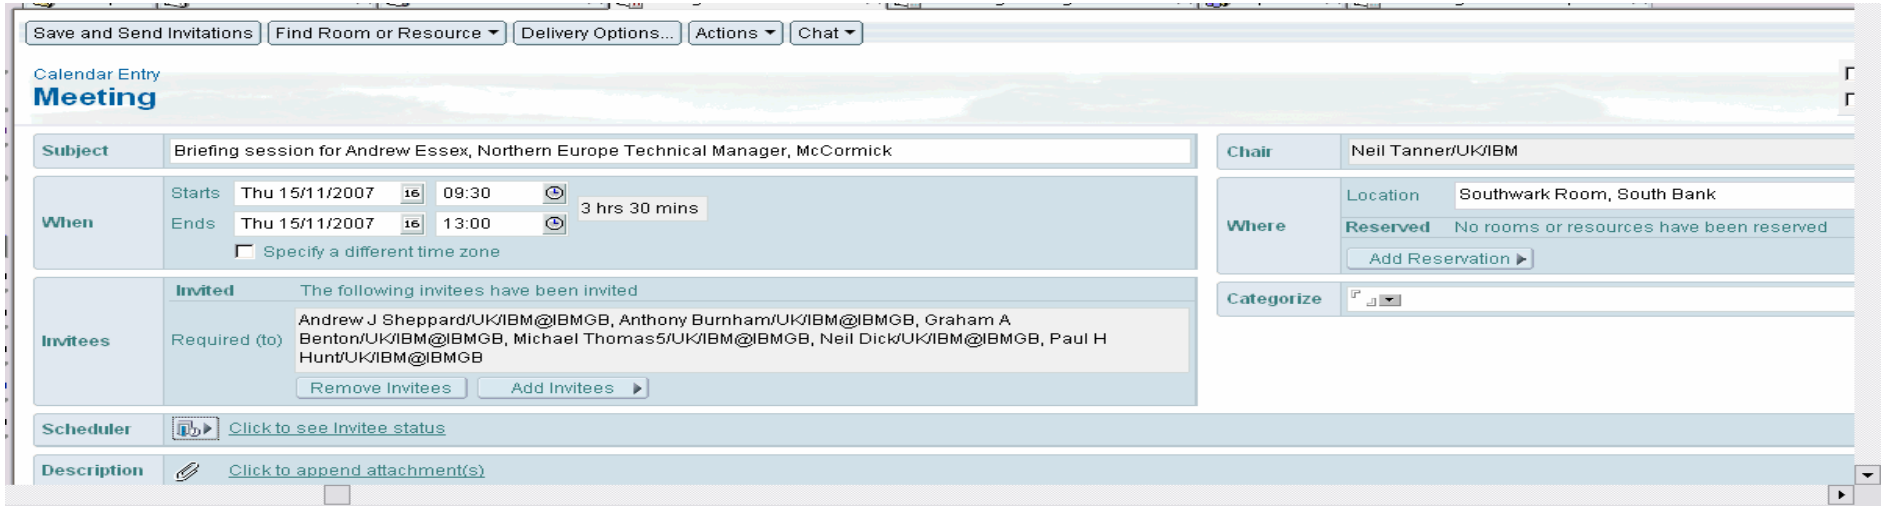

Persistent Data, Ephemeral Value

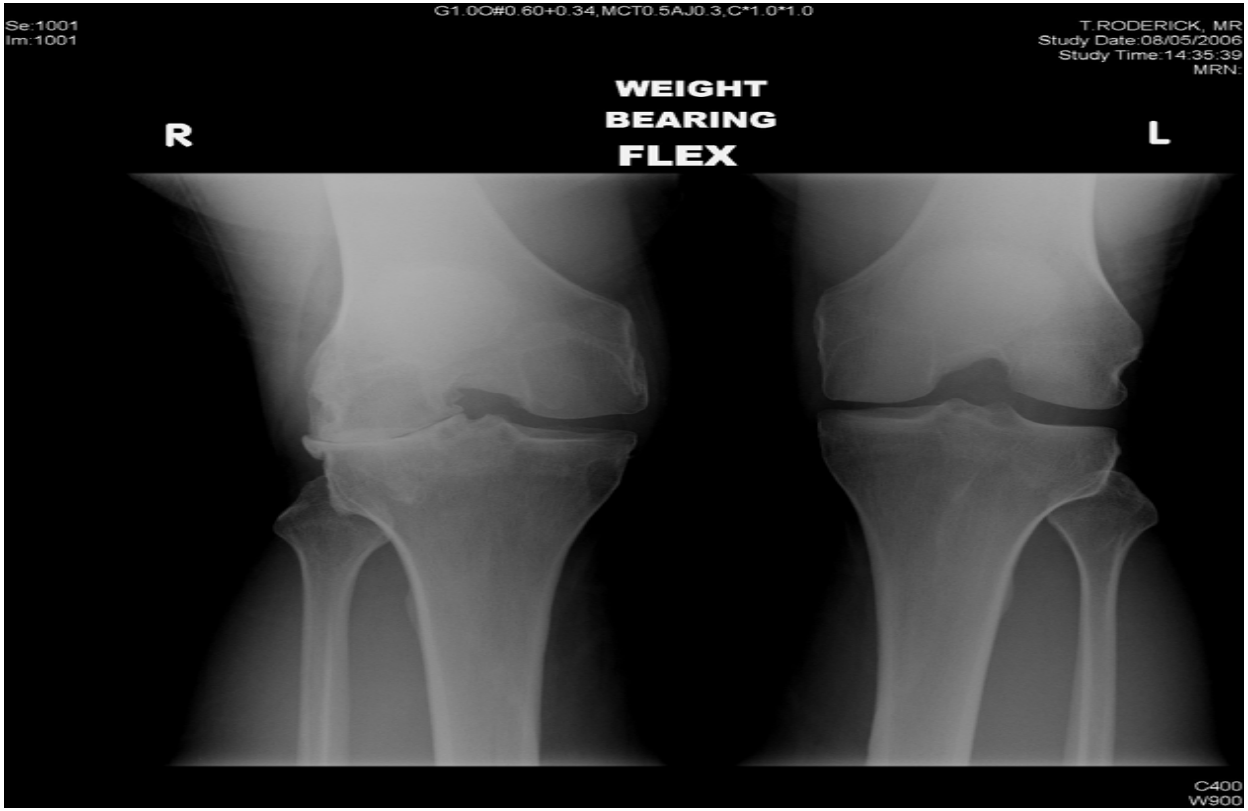

Are we solving Symptoms or Problems?

3

Are we solving Symptoms or Problems?

Are we solving Symptoms or Problems?

Are we solving Symptoms or Problems?

Are we solving Symptoms or Problems?

Are we solving Symptoms or Problems?

Are we solving Symptoms or Problems?

Are we solving Symptoms or Problems?

Remember this?

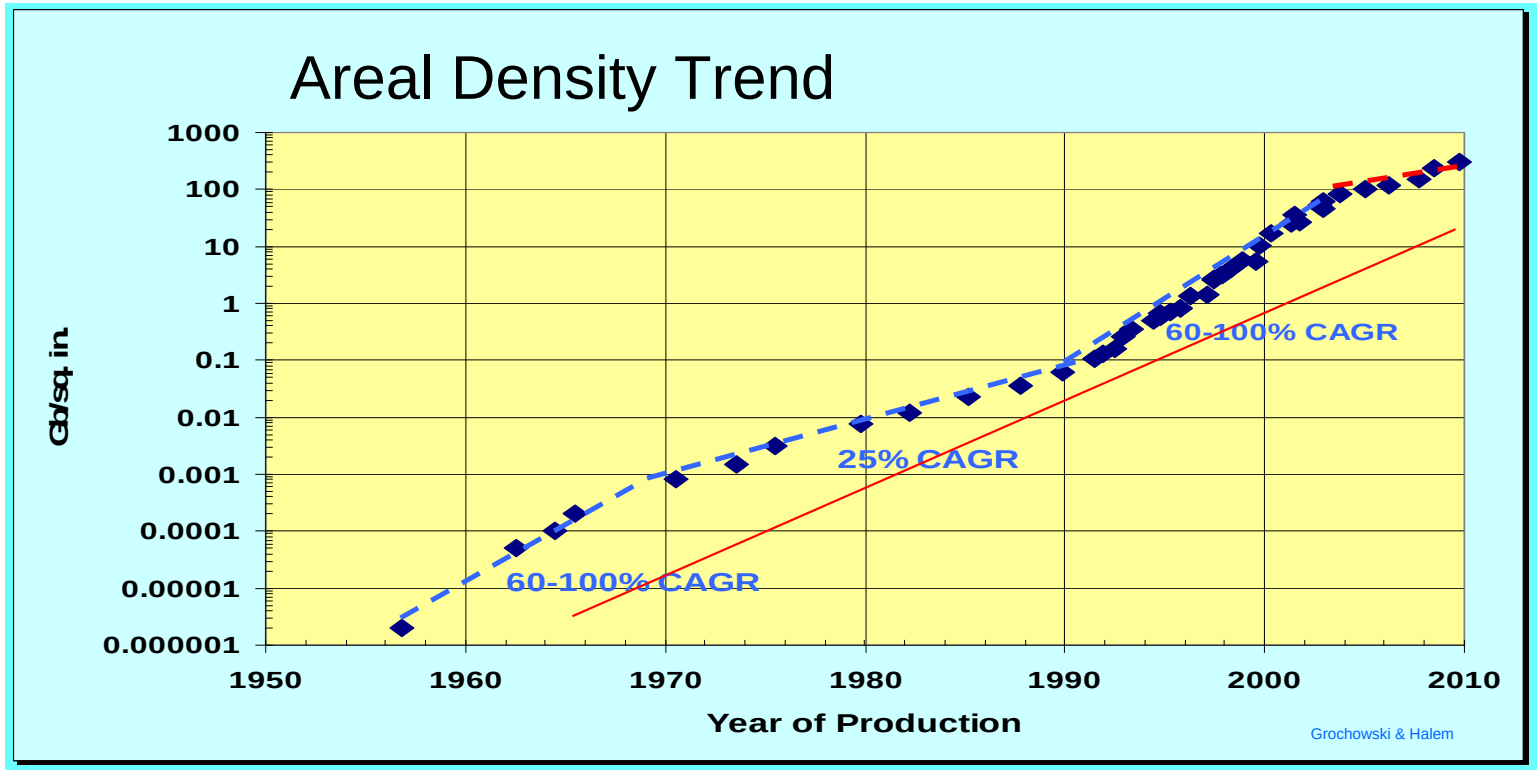

Remember this?

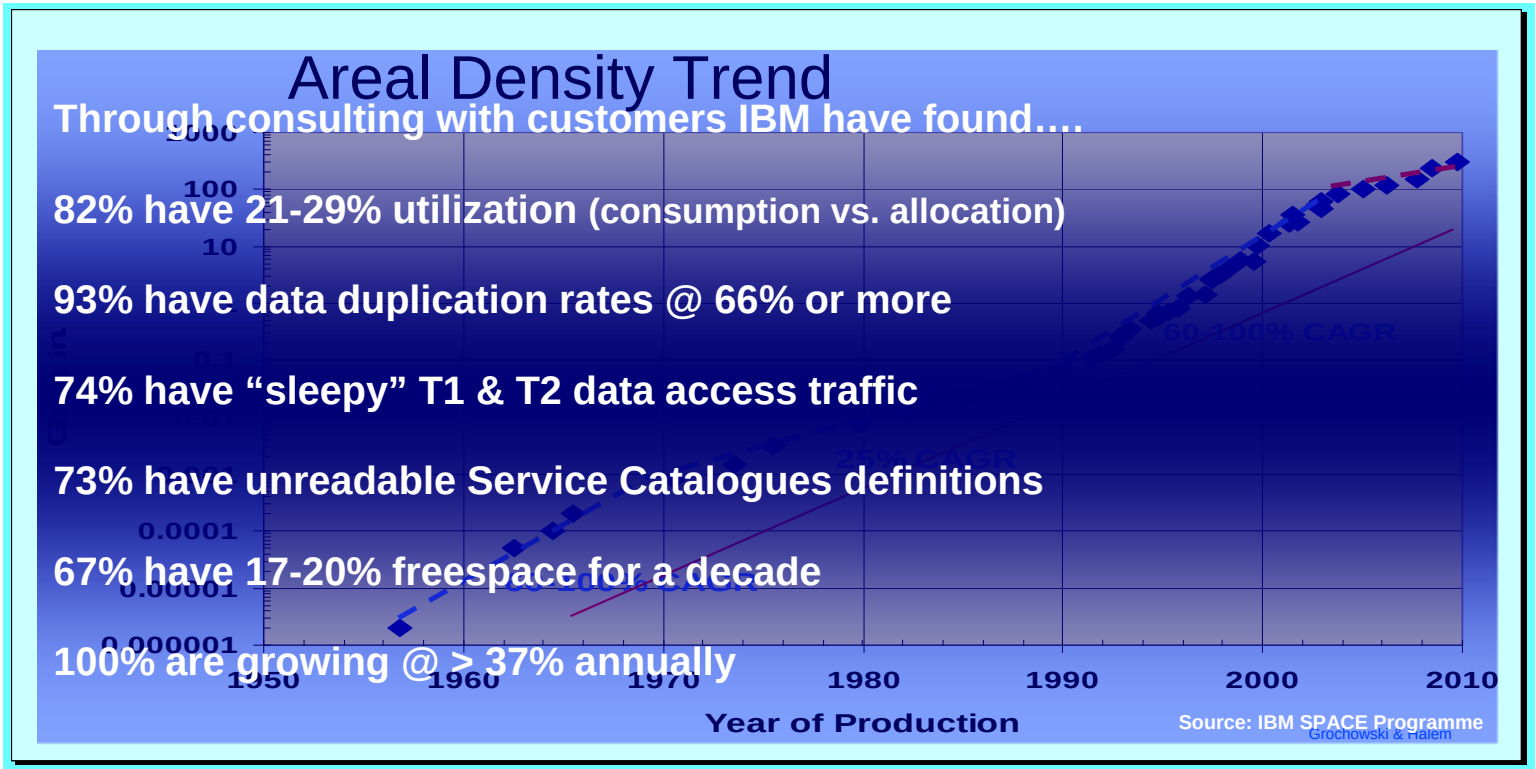

So what does that really mean?

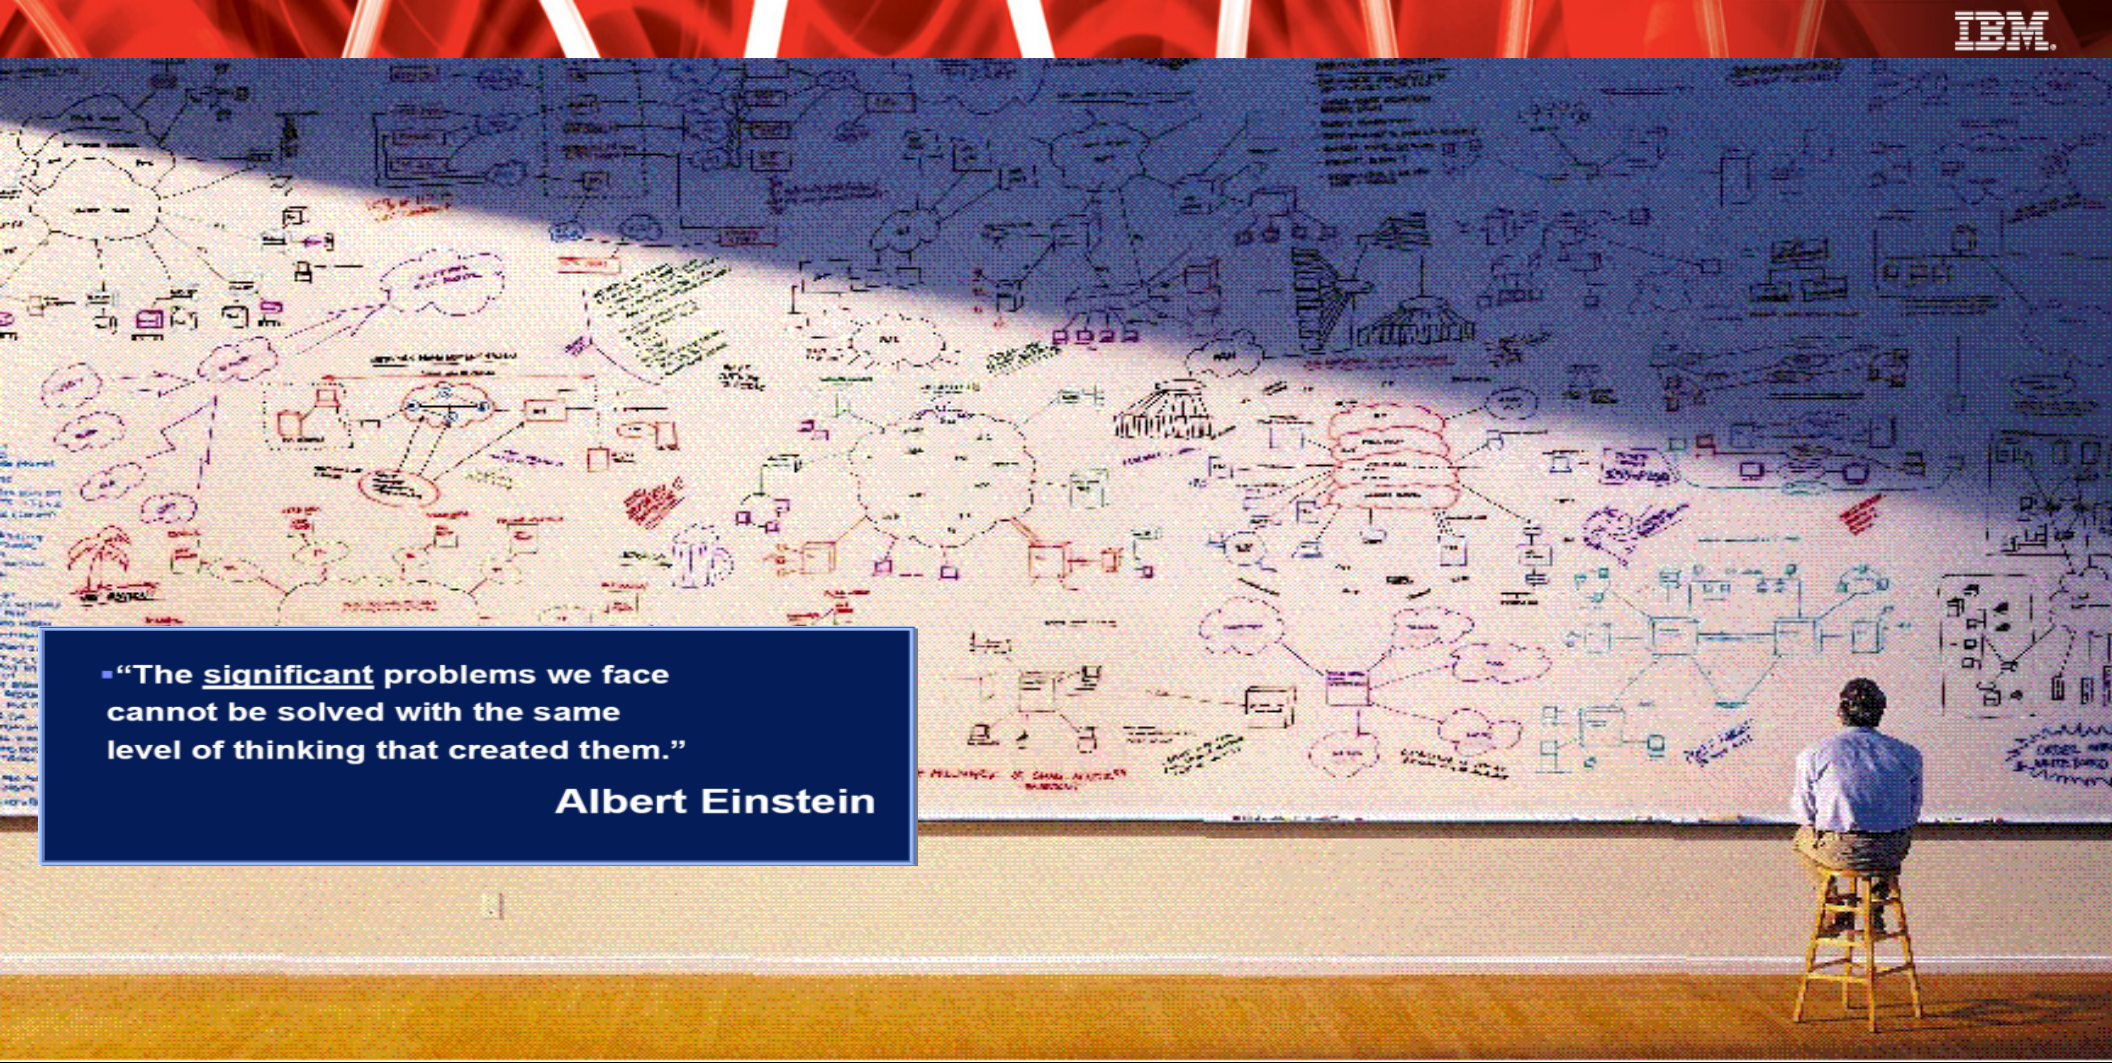
"The significant problems we face cannot be solved with the same level of thinking that created them."
Albert Einstein

What is the IBM SPACE Programme?

SPACE Probe

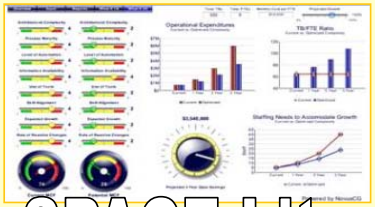

SPACE Lite

SPACE FOUNDATION / RAPID MCF

MCF Dashboard (all topics)

1. Capacity Management
2. Process Maturity
3. Skill Alignment
4. Technology Challenges
   - Architectural complexity
   - Level of automation
   - Use of tools
5. Business Challenges
   - Information availability
   - Expected Growth
   - Rate of re-active change (stability)

Prioritised roadmap with some costing

12 DAYS
(of which 5 are SPACE LITE / MCF Lite)

Report & Presentation

SPACE Delta

SPACE Foundation

SPACE Modules

- **Detailed Cost Discovery & Business Case**
  10 days
- **Storage Metrics (SRA/SERP)**
  3 days + license
- **Backup Metrics (Prism)**
  3 days
- **Data Profile (TekTools)**
  20 servers
  10 days
- **Architecture Assessment (ARC WP)**
  5 days
- **Application Mapping (SLRQ / A2I)**
  5 days
  (10 apps)
- **Environmental Maturity Assessment (MCF/Footprint)**
  4 days

RECOMMENDATIONS AND NEXT ACTIONS

SPACE Probe

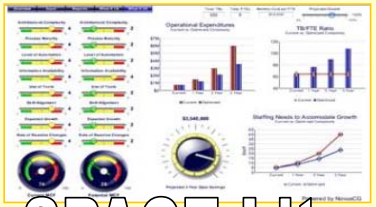

SPACE Lite

SPACE Delta

SPACE FOUNDATION / RAPID MCF

MCF Dashboard (all topics)

1. Capacity Management
2. Process Maturity
3. Skill Alignment
4. Technology Challenges
   - Architectural complexity
   - Level of automation
   - Use of tools
5. Business Challenges
   - Information availability
   - Expected Growth
   - Rate of re-active change (stability)

Prioritised roadmap with some costing

12 DAYS
(of which 5 are SPACE LITE / MCF Lite)

Report & Presentation

SPACE Foundation

SPACE Modules

- Detailed Cost Discovery & Business Case — 10 days
- Storage Metrics (SRA/SERP) — 3 days + license
- Backup Metrics (Prism) — 3 days
- Data Profile (TekTools) — 20 servers, 10 days
- Architecture Assessment (ARC WP) — 5 days
- Application Mapping (SLRQ / A2I) — 5 days (10 apps)
- Environmental Maturity Assessment (MCF/Footprint) — 4 days

RECOMMENDATIONS AND NEXT ACTIONS

SPACE Modules

SPACE Probe

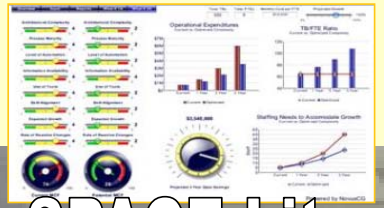

SPACE Lite

SPACE Delta

SPACE FOUNDATION / RAPID MCF

MCF Dashboard (all topics)

1. Capacity Management
2. Process Maturity
3. Skill Alignment
4. Technology Challenges
   - Architectural complexity
   - Level of automation
   - Use of tools
5. Business Challenges
   - Information availability
   - Expected Growth
   - Rate of re-active change (stability)

Prioritised roadmap with some costing

12 DAYS
(of which 5 are SPACE LITE / MCF Lite)

Report & Presentation

SPACE Foundation

- Detailed Cost Discovery & Business Case — 10 days
- Storage Metrics (SRA/SERP) — 3 days + license
- Backup Metrics (Prism) — 3 days
- Data Profile (TekTools) — 20 servers, 10 days
- Architecture Assessment (ARC WP) — 5 days
- Application Mapping (SLRQ / A2I) — 5 days (10 apps)
- Environmental Maturity Assessment (MCF/Footprint) — 4 days

RECOMMENDATIONS AND NEXT ACTIONS

Total TBs	Total FTEs	Monthly Cost per FTE	Projected Growth
320	5	$12,500	100% (0% – 150%)

Current		Potential	
Architectural Complexity	4	Architectural Complexity	2
Process Maturity	4	Process Maturity	2
Level of Automation	4	Level of Automation	2
Information Availability	4	Information Availability	4
Use of Tools	4	Use of Tools	2
Skill Alignment	4	Skill Alignment	2
Expected Growth	4	Expected Growth	2
Rate of Reactive Changes	4	Rate of Reactive Changes	2

Current MCF

Potential MCF

Operational Expenditures
Current vs. Optimized Complexity

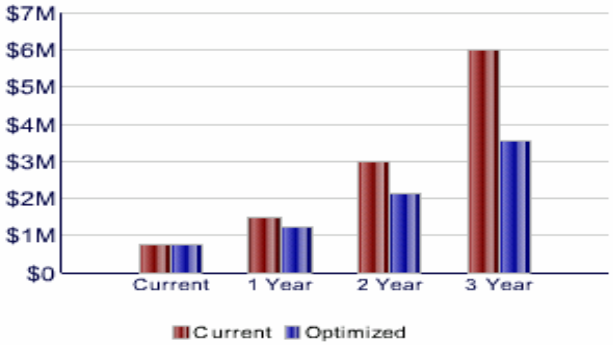

TB/FTE Ratio
Current vs. Optimized Complexity

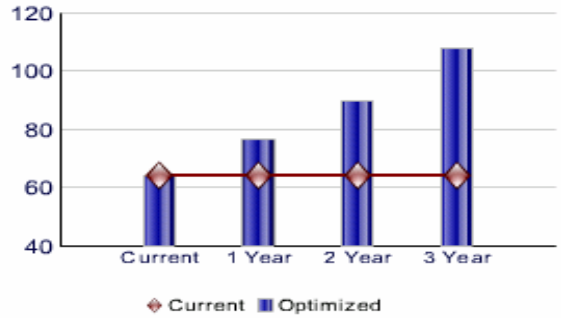

$3,540,000

Projected 3 Year Opex Savings

Staffing Needs to Accomodate Growth
Current vs. Optimized Complexity

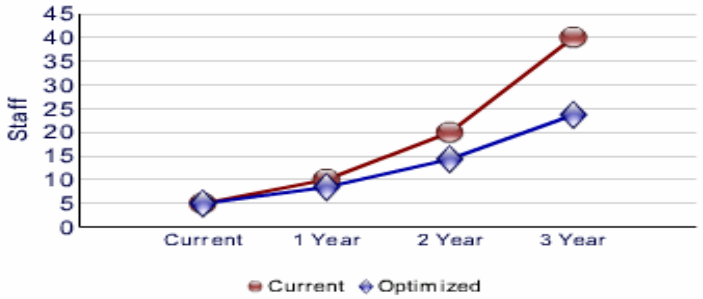

Powered by NovusCG

Example Reclamation Method & Capacity Released

Storage Targets and complexity to recover

Method	Recoverable TB	Complexity
Orph GT50	39	2
Orph LT50	42	3
Orph UI	64	6
Reclm GT50	12	2
Reclm LT50	67	3
Mapped	117	2
Low Utilization Reclaim	65	8

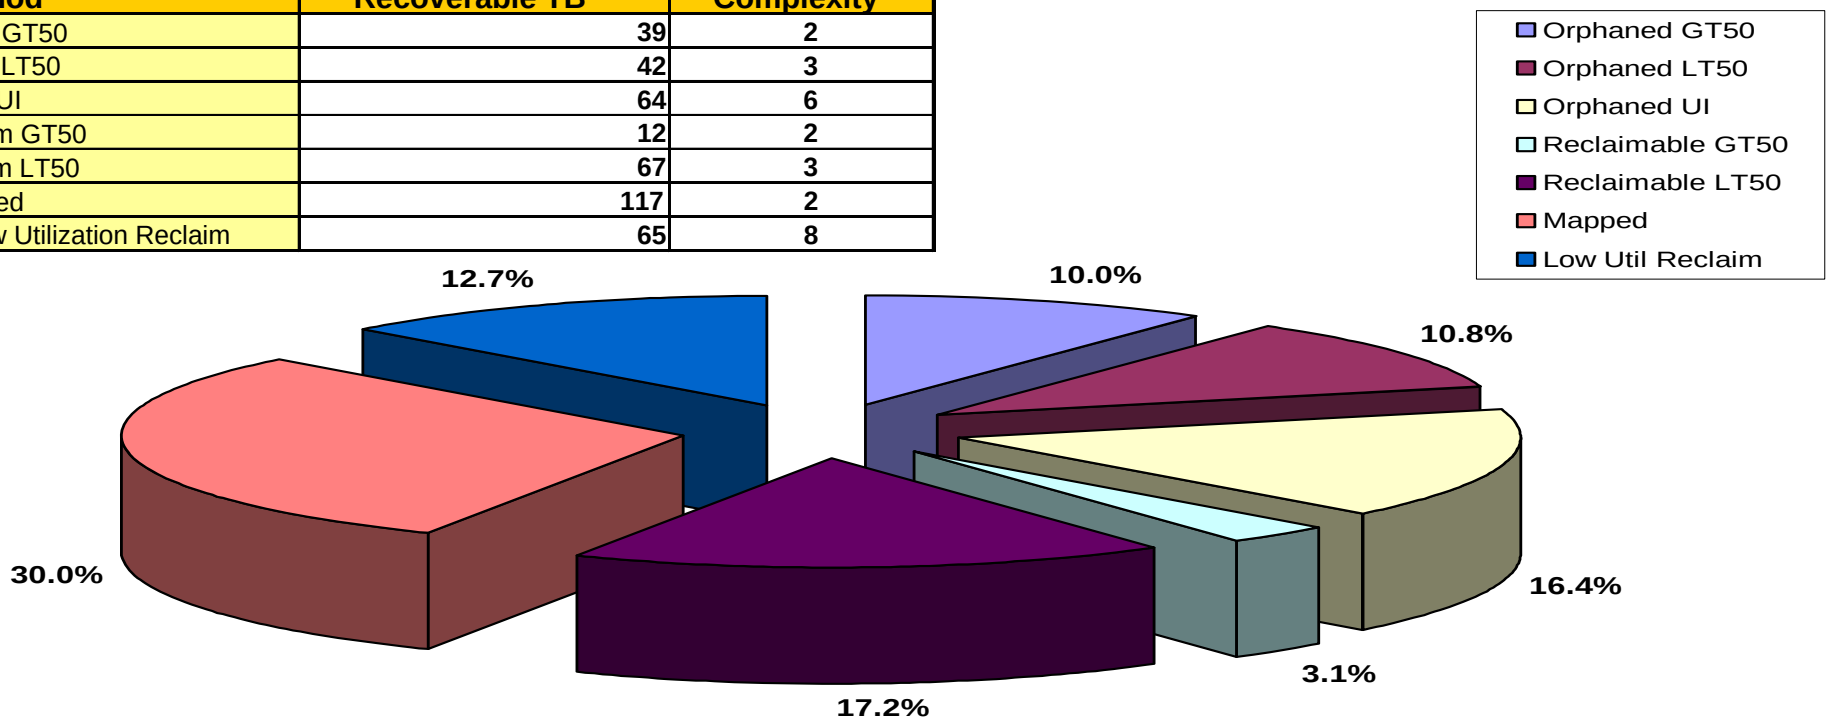

Reclamation Categories

Trapped Cost - Categories	Category Description	Recoverable TB's	Time to Free Up	Complexity (1-10)	Recovery Actions
Orphaned Storage	**Storage configured to servers but not picked up on the host side. Could be a server that was removed, but with volumes that weren't unmasked or a server on which the storage hasn't been picked upon scanning the bus.**	**81**	**Days**	**3**	**Identify, unmask, rezone and mask the storage for immediate use.**
Orphaned Uninstrumented	**Storage that may or may not be orphaned. This is caused by SRM agents not being properly rolled out or discovery not being completed.**	**64**	**Days/ Weeks**	**5**	**Requires validation on the host side to ensure that the volumes are truly not being used and remediation.**
Reclaimable Storage	**Storage presented to a server but not in a volume group**	**79**	**Days**	**3**	**Storage can be pulled into a volume group and carved into a file system.**
Mapped	**Storage mapped down an FA/front end channel port, but not masked to a host**	**117**	**Days**	**2**	**Storage can be masked to a HBA and made available to a host.**
Storage Utilisation less than 10% - Server view	**Storage configured but with marginal usage at the host/file system level.**	**65**	**Weeks/ Months**	**8**	**Defer future 'capex' upgrades and drive up utilisation.**

When you mix the ingredients well

Information Infrastructure supports the Dynamic Infrastructure Imperatives

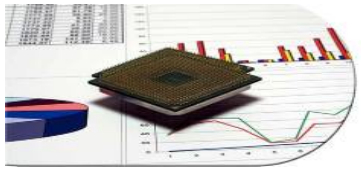

- *Provision new applications resources in minutes*
- *Triple asset utilization without most upgrades*
- *Shrink physical disk storage up to 50%*
- *Begin savings in year 1*
- *System outages down 58%*
- *Reduce floor space by 80%*
- *Power & cooling reduced by 60%*

JAL

ING

FDA

CITY AND COUNTY OF HONOLULU STATE OF HAWAII

Nationwide On Your Side

DHL

NATURAL HISTORY MUSEUM

How do I participate in the IBM SPACE Programme?

Remember, *there are no “secret pills”*

Remember, *there are no “secret pills”*

Remember, *there are no “secret pills”*

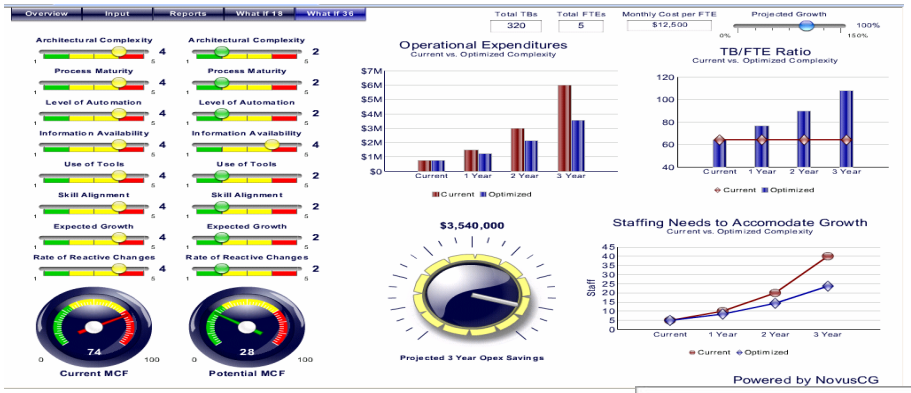

Data Deduplication
Data Deduplication
Data Deduplication

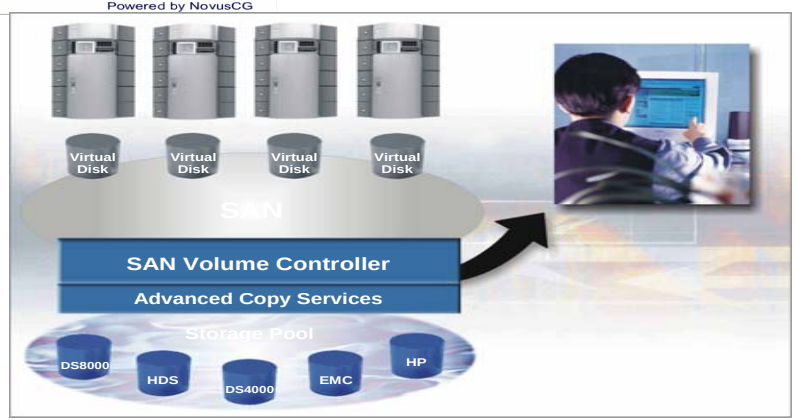

Remember, *there are no "secret pills"*

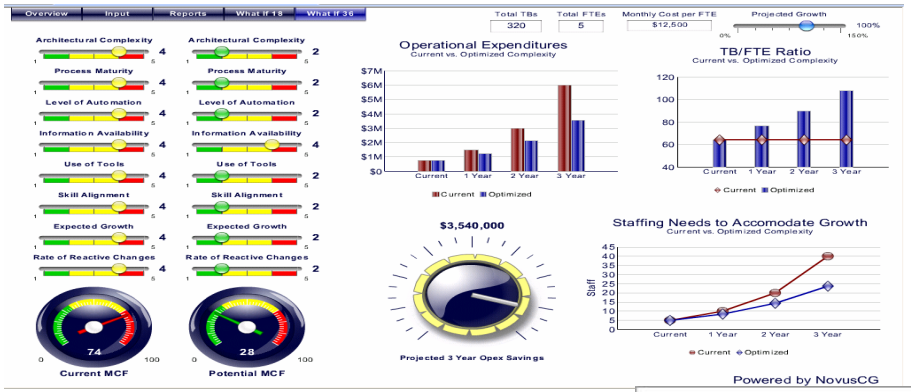

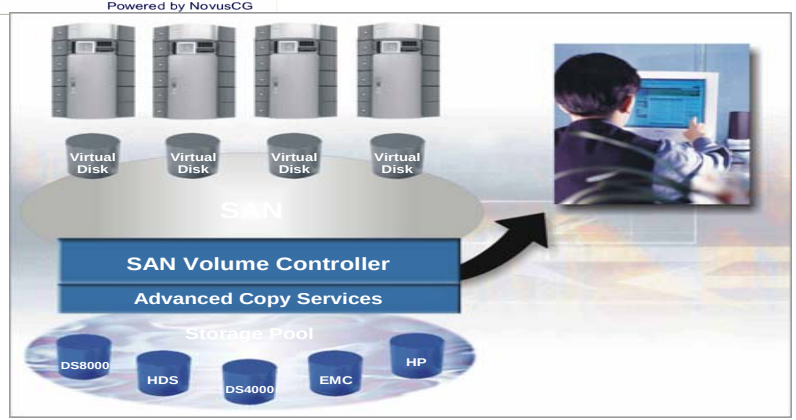

Data Deduplication

Current events mean a change in IT Prioritization

And the leaders are already planning…

Old Thinking

IT *maintains* IT **resources** that support the business

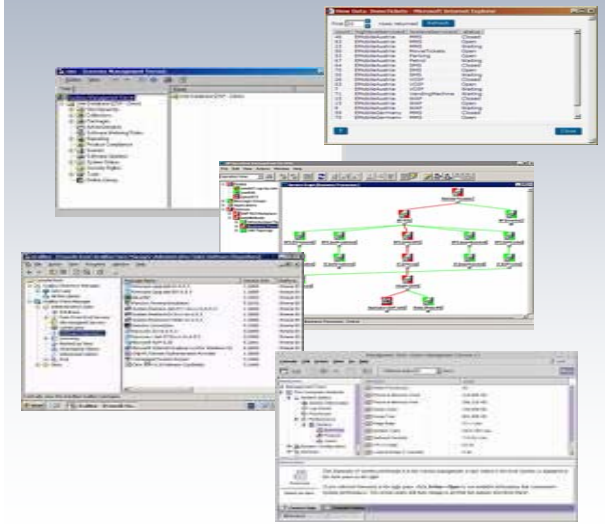

From Silos …

New Thinking

IT *delivers* **services** designed to meet business **goals**

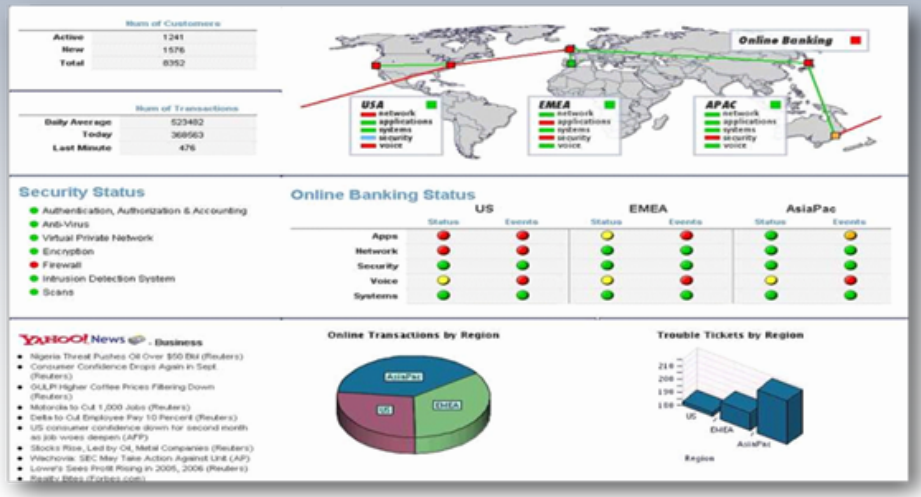

… to Services

IBM System Storage – Offering Matrix

Disk Systems

- SAN Volume Controller
- DS family
- XIV
- N series

Availability Management

- TotalStorage Productivity Center, SSPC
- SAN Fabric Management software
- Tivoli Provisioning Manager
- Tivoli Storage Process Manager
- IBM Systems Director family

Tape Systems

- TS family of drives, libraries
- VTL
- ProtectTIER

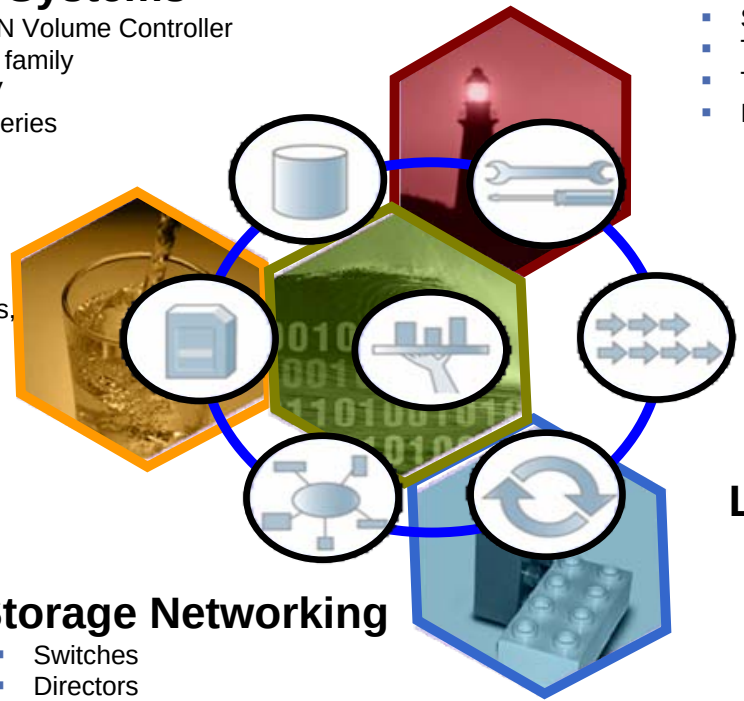

Business Continuity

- Productivity Center for Replication
- Advanced copy services
- Tivoli Storage Manager (TSM) family
- Tivoli Continuous Data Protection (CDP)
- Tape cluster grids and Peer-to-Peer
- GDOC, GDPS

Services

- Consulting
- Assessments
- Design
- Migration
- Deployment
- Outsourcing
- Hosting

Storage Networking

- Switches
- Directors
- Routers

Lifecycle and Retention

- DR550, DR550 Express, FS gateway
- Grid Archive Manager, GMAS
- TSM Space Management for Unix/Windows
- GPFS, DFSMS
- N series with SnapLock™
- WORM tape support

Responsible reactions - Technologies

- Virtualise
- Commoditise
- De-Duplicate
- Match data value to storage cost
- Future thoughts

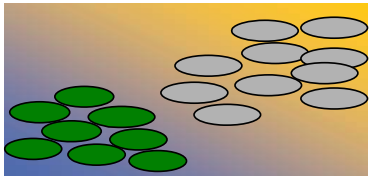

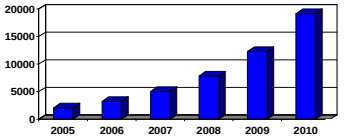

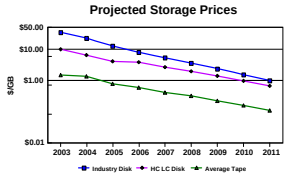

Universal Application of the Virtualized Information Infrastructure

Key Principles

- Comprehensive
- Open
- Heterogeneous
- Common skills

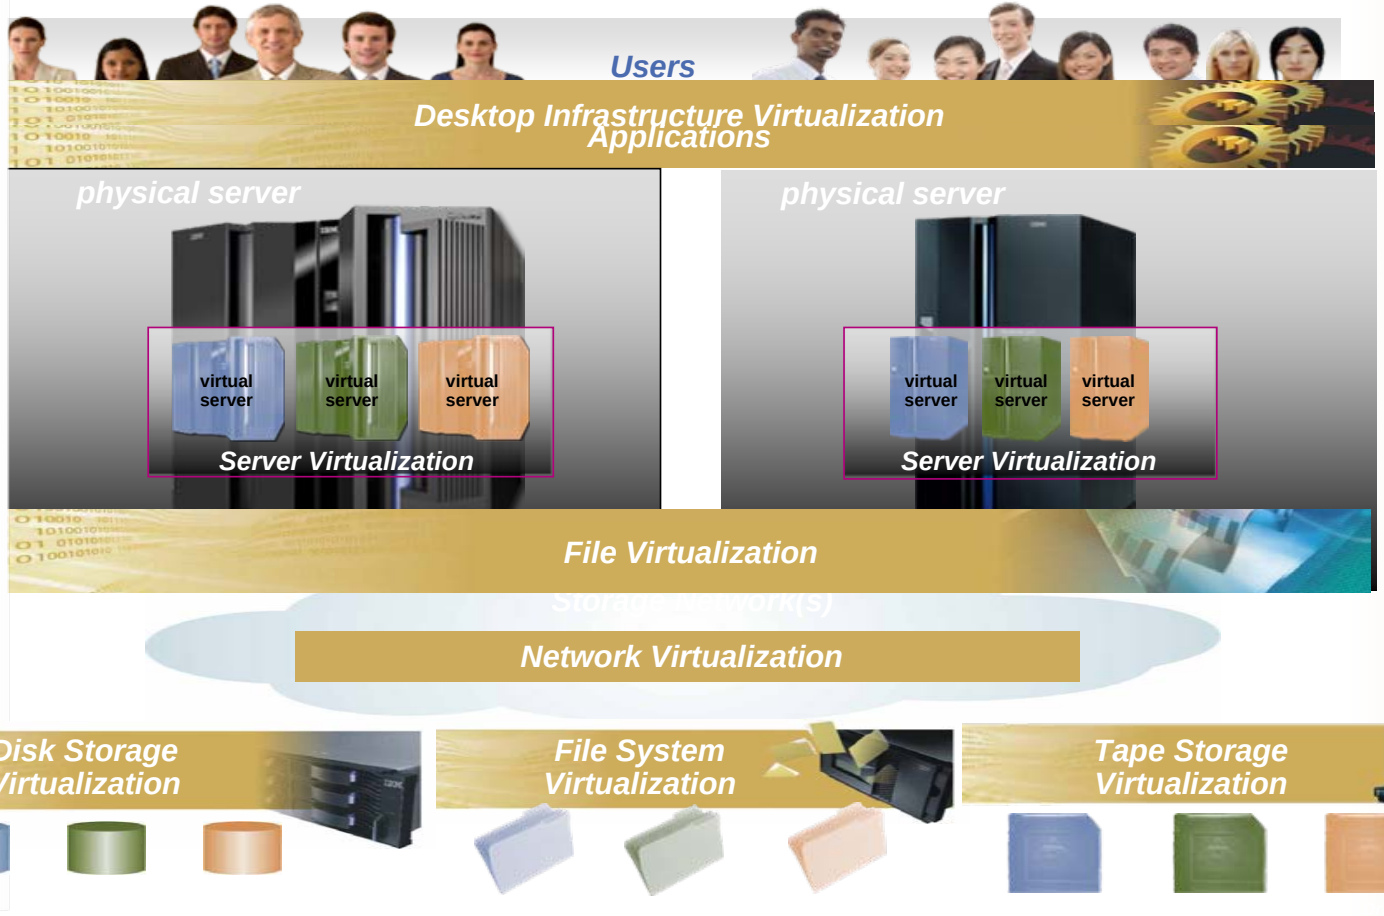

Virtual/Physical Management

The Storage Challenges tackled by Virtualisation

- Compexity
- Migration
- Interoperability
- Vendor lock-in

- Benefits of SAN Virtualization (according to IDC)
- >50% boost in effective capacity use
  - – share space across file servers and NAS systems
- >70% decrease in migration/consolidation time
  - – nondisruptive migration during business hours
- >40% reduction in spending on new file-based storage
  - – leverage automated movement and tiered storage for backup and archive

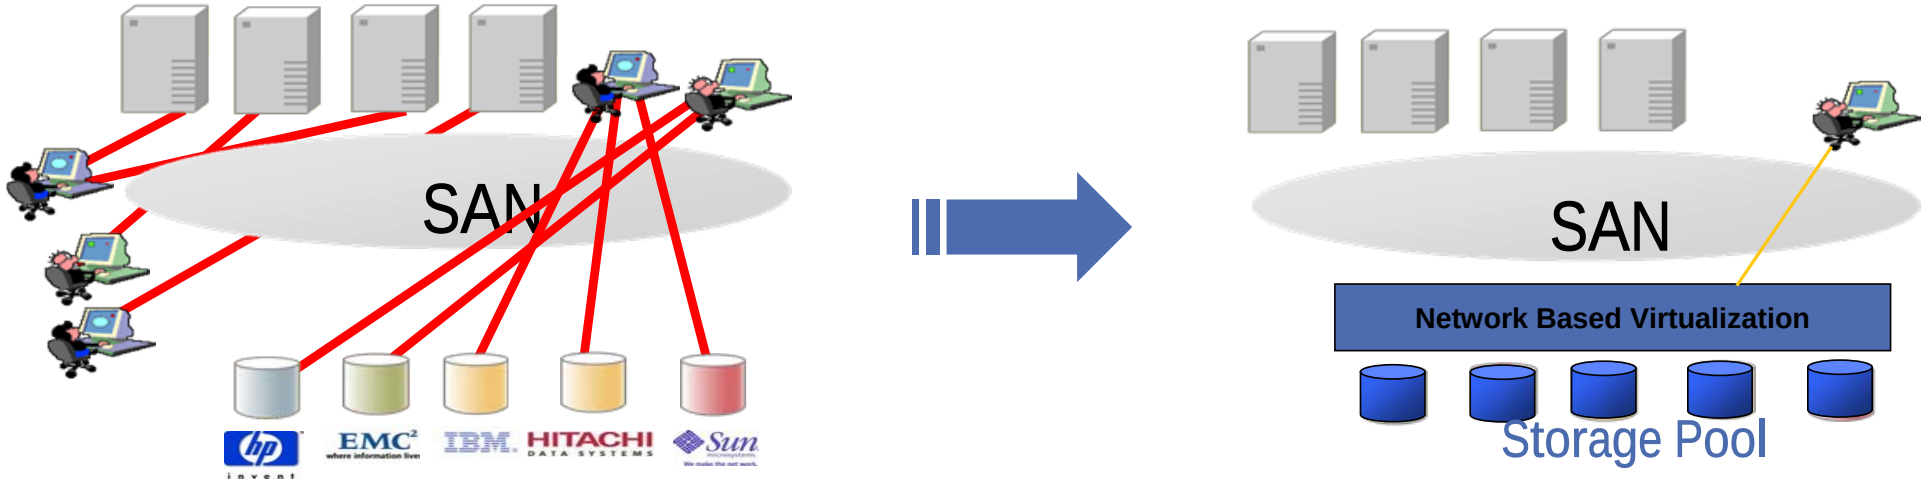

IBM Disk Storage Portfolio

	Mainframe	Open Systems	Heterogeneous	NAS
		SVC		SoFS For customers with massive IO / backup / restores

High-end	DS8000
	For clients requiring:
	• One solution for mainframe & open systems
	• Disaster Recovery
	  • across 3 sites
	  • across 2 sites > 100 km apart
	• Secure encryption
	• Continuous availability, no downtime for upgrades
	• Best-in-class response time
		XIV
		For clients requiring:
		• Open systems environment support
		• Ease of management
		• Web 2.0 workloads
		• Self tuning
			For clients requiring:
			• Virtualization of multiple vendor environments, including IBM, EMC, HP and others
			• Open systems environment support, including IBM system i support
				N series
				For clients with:
				• Combined NAS (file) and SAN requirements
				• Simple two-site high availability

Mid-range	DS6000 For clients requiring mid-range mainframe support. DS8000 compatible	DS4000/5000 For customers requiring: • Open systems environments including IBM i support • Cost efficient storage for capacity • Modular storage		
Entry		DS3000 • For very price-sensitive customers • Modular storage • Affordable disk shelves behind SVC EE		
Special Purpose		DCS9900 (HPC, Digital Media, IPCCTV, Archive)		

Storage Differentiation

- New uses for information, and growth in scale of information will strain existing infrastructure designs and operational practices for data management
- The major opportunities for differentiation are:
  - Scale
  - Management
  - Archival function
  - Security
- *Progress in basic recording technology, computing, and networking will not resolve the storage issues*
- Investments in new infrastructure and processes will be needed to prepare for the coming uses for digital information

XIV Architecture - Scale out

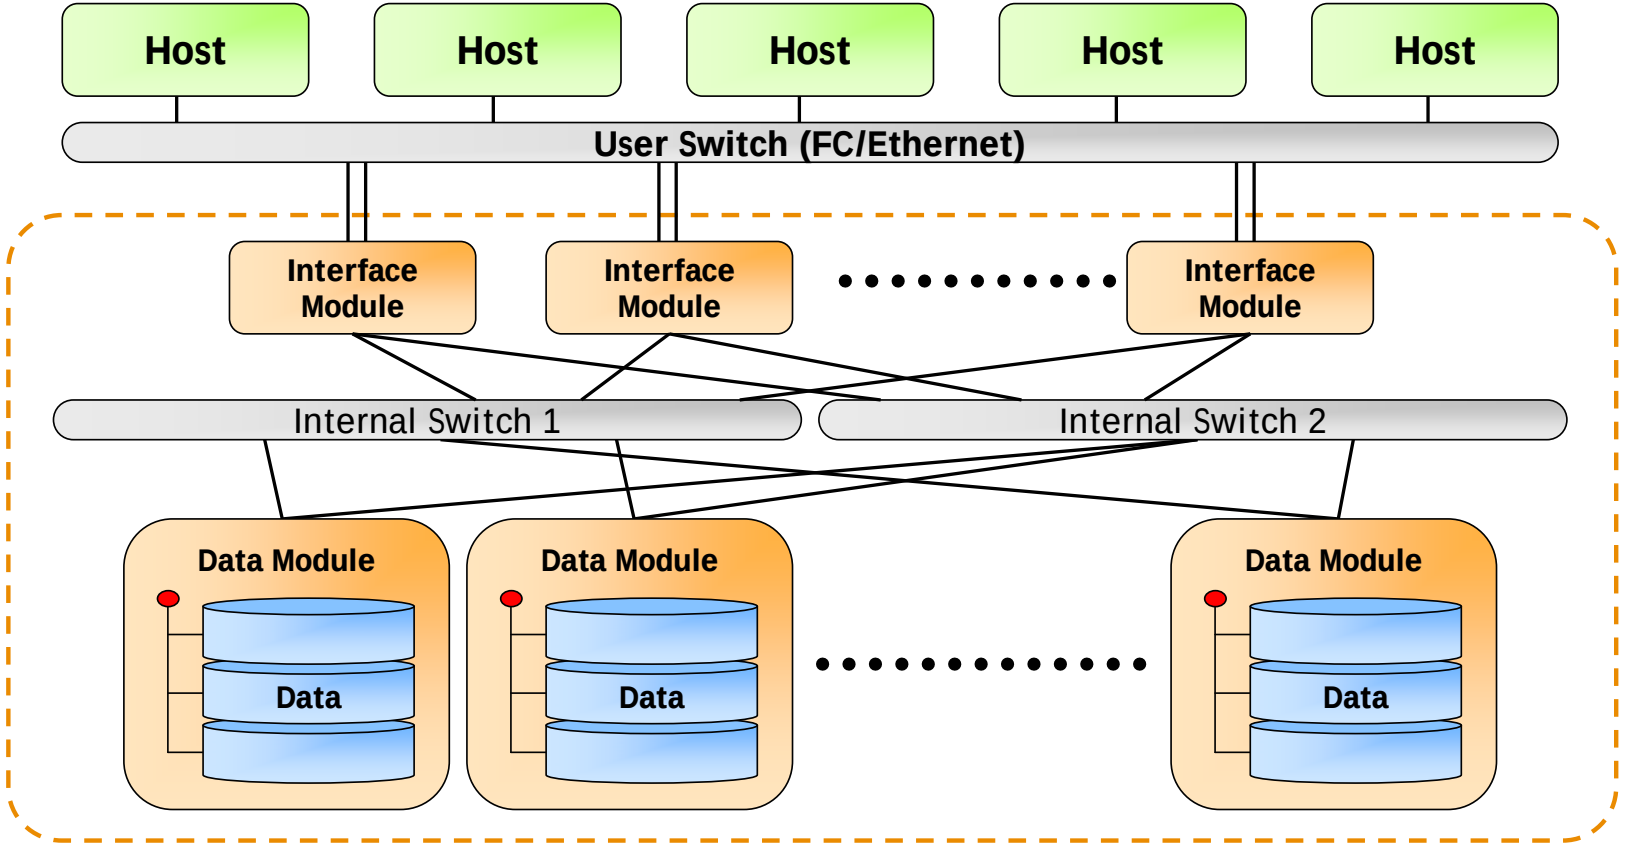

Hardware implementation

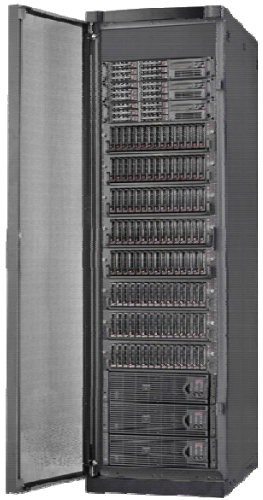

- **Commodity gigabit Ethernet switches act as a backplane**
- **Interface modules: commodity servers act as ‘unintelligent’ routers**
- **Data modules: commodity servers with SATA drives act as ‘smart’ storage sub systems:**
  - **Caching**
  - **Advanced Replication functionality**
  - **Virtualization**
- **UPSs**

Use Case Scenarios for the Scale-out Architecture

- **Digital Archive**
- **Web 2.0**
- **Service Providers**
- **Digital Media**
- **Development / Test**
- **Clustered Computing**

What is Data De-duplication?

- **Data de-duplication** (often called "intelligent compression") is a method of reducing storage needs by eliminating redundant data. Only one unique instance of the data is actually retained on storage media. Redundant data is replaced with a reference or pointer to the unique data copy.

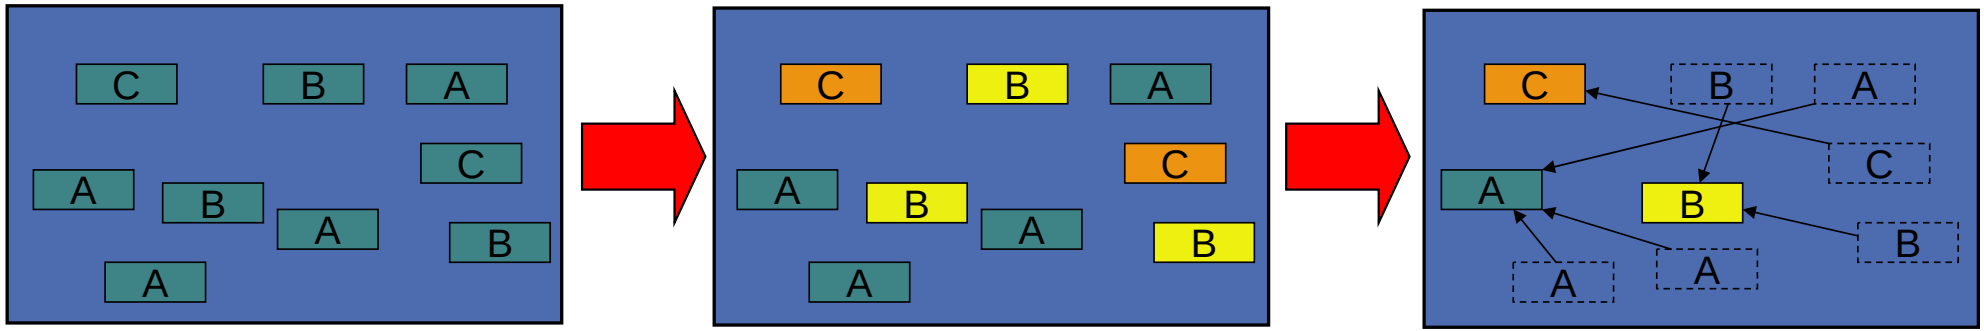

1. Data elements are evaluated to determine a unique signature for each

2. Signature values are compared to identify all duplicates

3. Duplicate data elements are eliminated and are replaced with pointers to the existing reference element

ProtecTIER Differentiators

Performance

Sustainable 450 MB/s per node (900 MB/s two node cluster), performing ***inline*** de-duplication

Capacity

Up to 1 PB physical capacity per node

Enterprise-Class Data Integrity

Binary diff process during de-dupe designed for the highest data integrity

Non-Disruption

Daily Operations

Inline de-duplication eliminates need for significant secondary processing

Implementation

Integrates well with existing backup environment and infrastructure

Reducing Total Costs through Energy Efficiency

IBM is the Leader in Green Data Center Initiatives

Cut TCO 50% with Blended Tape and Disk*

10 year TCO example. Assumes 250TB of storage, 25% growth per year.

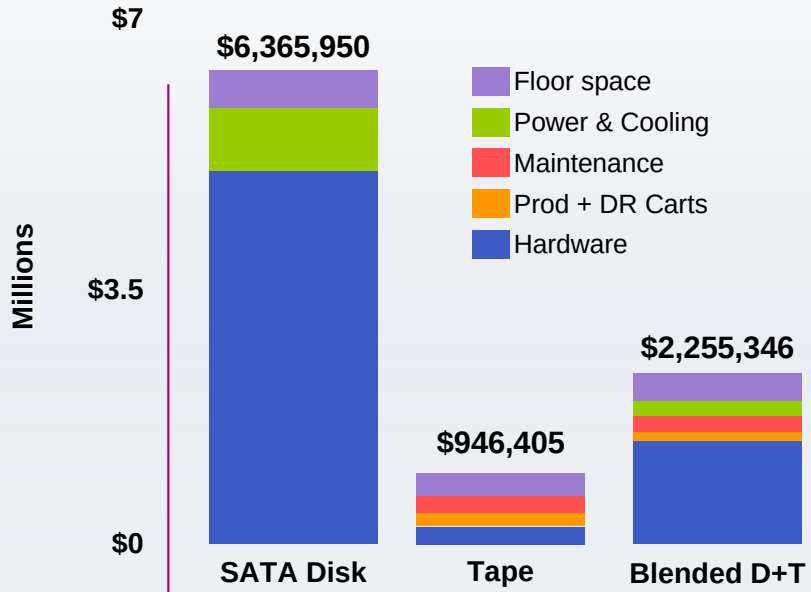

"... tape is still less expensive to acquire and maintain than disk over the long term and it can offer power and cooling efficiencies over disk."

Source: Addressing Archiving and Retention Challenges In the Government Sector, Heidi Biggar and Brian Babineau, Enterprise Strategy Group, March, 2008

Green Projects for Information Infrastructure Reduce TCO

- Virtualization
- Best Practices
- Technology Refresh
- Tiered Storage
- Storage Consolidation

* TCO estimates based on IBM internal studies.

So what's next?

GLOBAL
INNOVATION
OUTLOOK

Why Solid State Storage? Time Scale

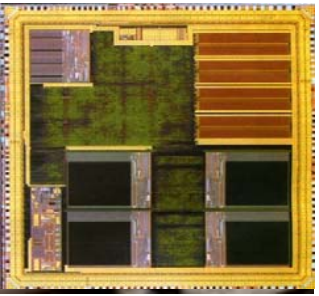

Time ((s))

	Human scale	IT scale		
10^{10}	century	3s	Get data from TAPE (40s)	
10^{9}	decade	0.3s		
10^{8}	year	31ms		
10^{7}	month	2.5ms	Get data from DISK (5ms)	*Magnetic Storage*
10^{6}	week	0.6ms	Write to FLASH, random (1 ms)	
10^{5}	day	86us		
10^{4}			Read a FLASH device (20 us)	*SCM*
10^{3}	hour	3.6us		
10^{2}	minute	60ns	Get data from DRAM (60ns)	
10			Get data from L2 cache (10ns)	*Memory*
1	second	1ns	CPU operations (1ns)	

Storage Class Memory compared to other technologies

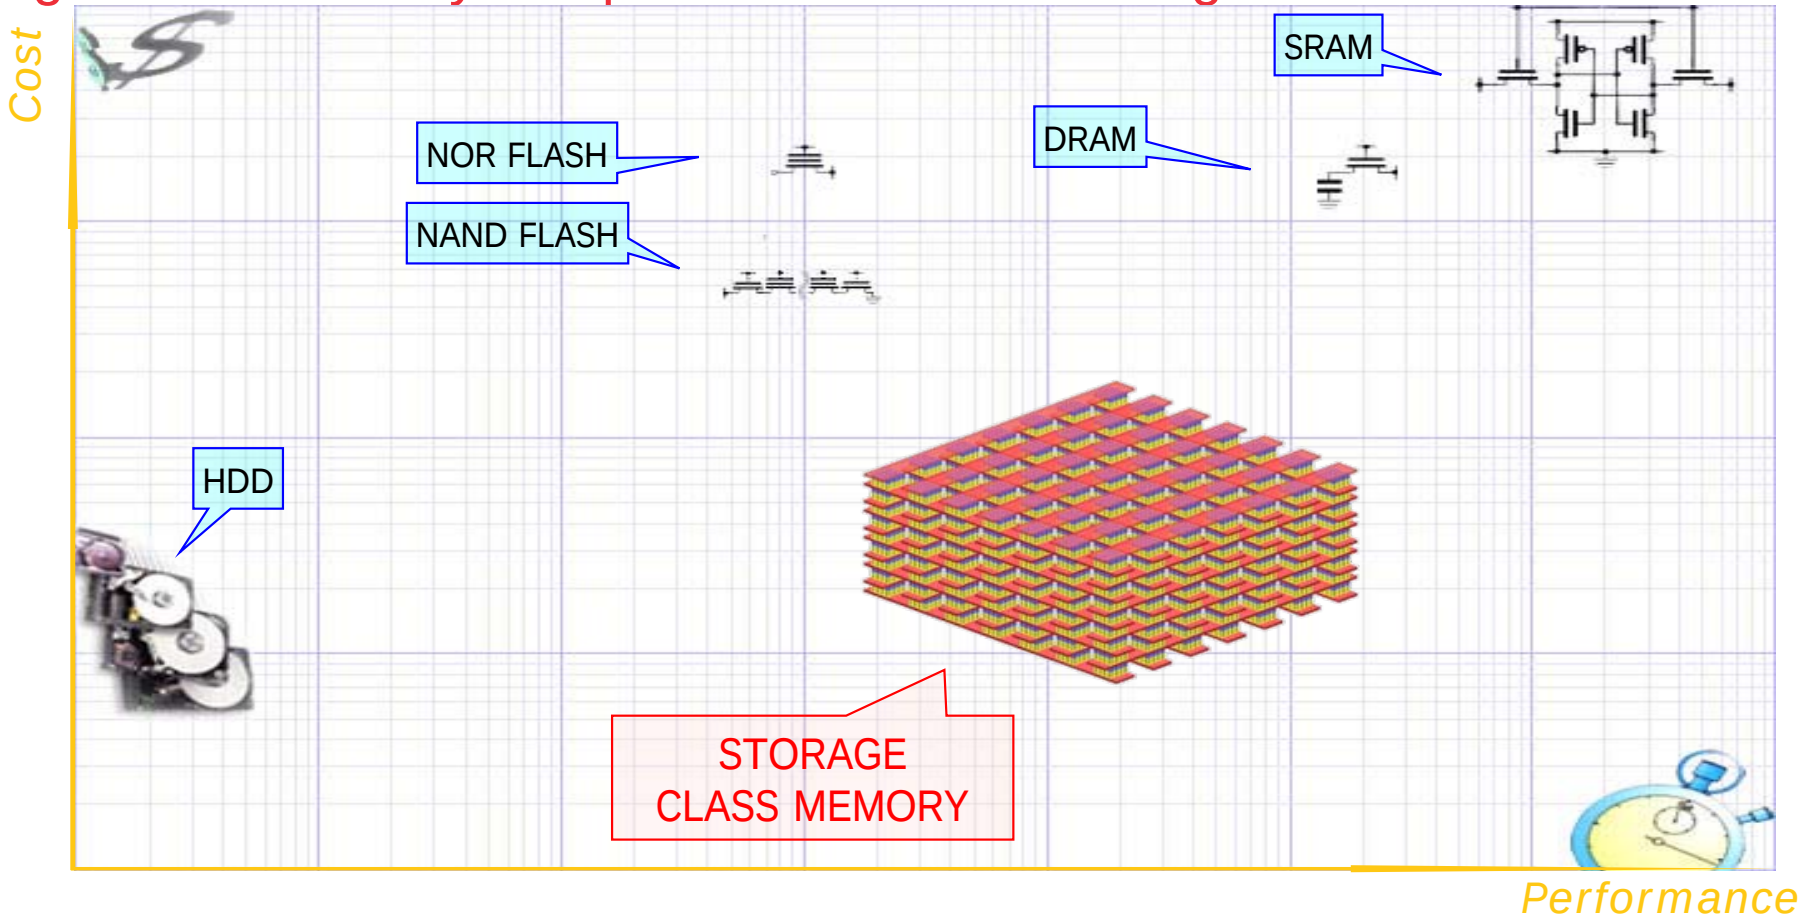

Solid state storage – Emerging technologies

Quantity: FRAM, PFRAM, SiC Bipolar, Molecular, NanoCrystal

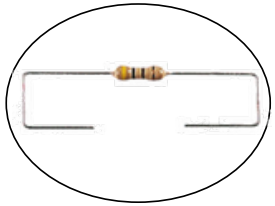

Flow: PCRAM, PMC, Polymer, RRAM, Perovskite

Spin: MRAM, Magnetic Shift Register

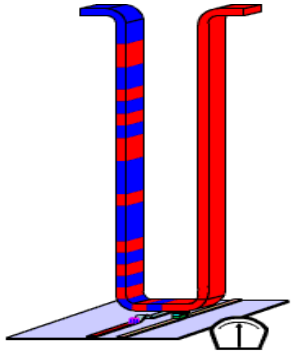

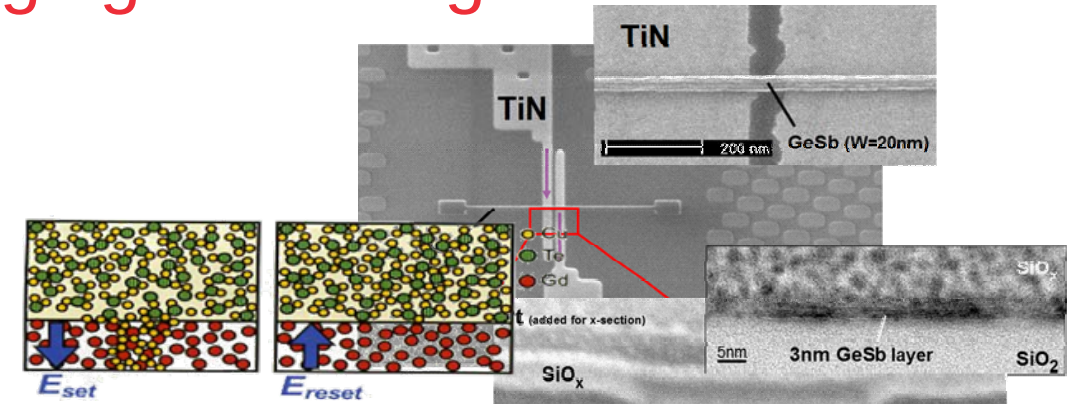

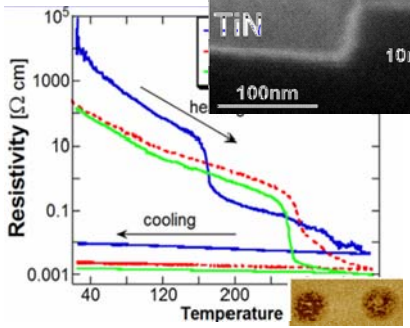

Y D Chen et al. IDEM 2006

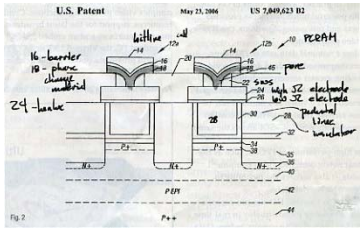

Latest in Storage Research : "RaceTrack" Memory

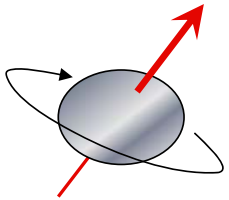

Spintronics

Stuart Parkin, IBM Fellow, inventor of GMR read heads, is developing “racetrack memory”, basis for unprecedented storage density, in IBM Almaden Research labs

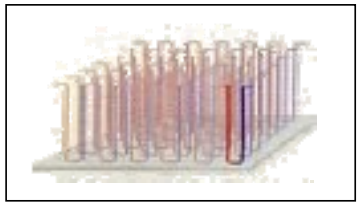

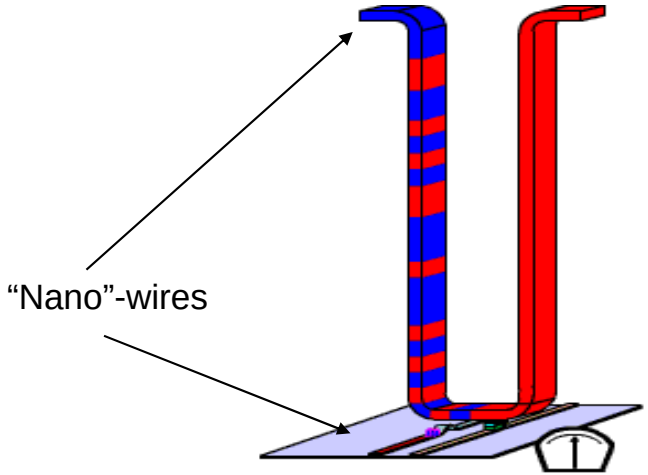

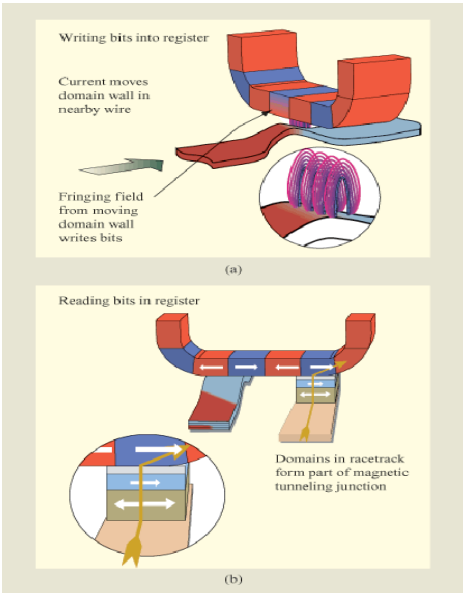

Remember this?

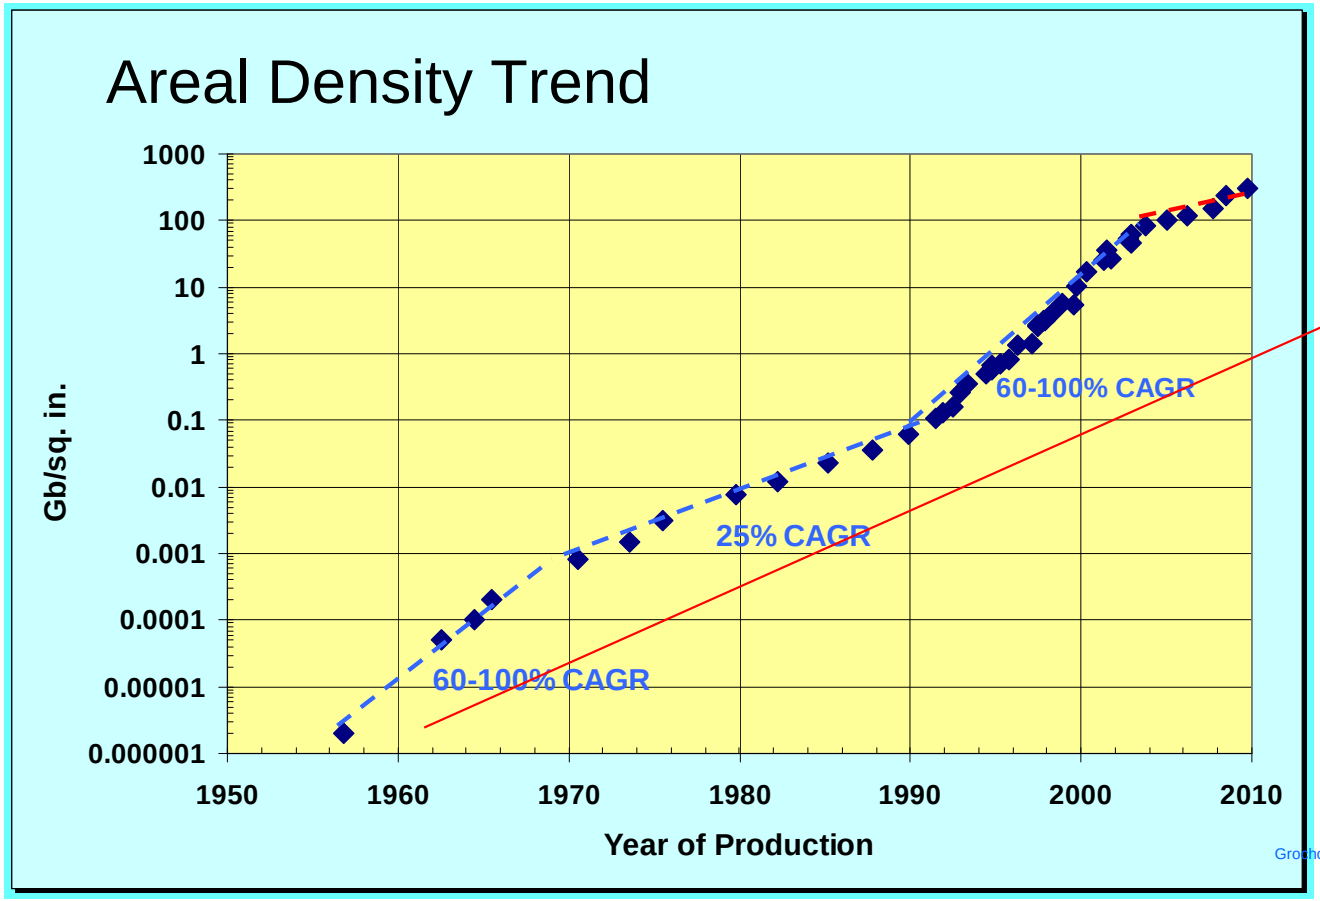

Racetrack?

Factor of 100x

Grochowski & Halem

Solid State Disk Summary

- Disk density (& cost take-down rate) slowing
- Solid state memory effects are being very actively researched
- Mainstream first half of next decade, probability for some niche applications sooner
- Abrupt “phase change” from magnetic disk to solid state (compare with flat screen introduction)

- Positioning for flexibility of storage using virtualisation

Why IBM Information Infrastructure

Reduce Cost
Optimise IT to Save/Make Money

Mitigate Risk
stay out of jail while staying in business

Improve Service
plan to keep going in a changing world

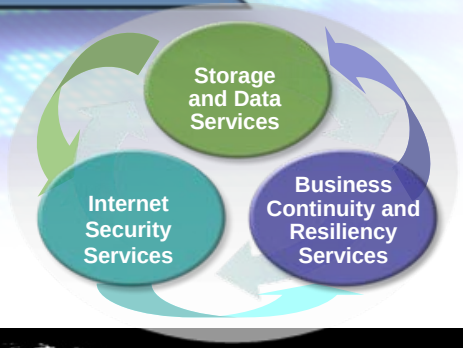

If not now, when…

Thank You

Steve Legg
Chief Technology Officer – Storage UKI
Contact Details +44 79 67 27 54 06
splegg@uk.ibm.com

Rick Terry
Information Infrastructure Evangelist
Contact Details +44 77 25 70 64 81
ricterry@uk.ibm.com