




IBM Systems and Technology Group

# **Green Shoots:** **What is emerging in Storage Technology**


**Rick Terry**  
**Storage Solutions Specialist**  
**IBM Systems Storage - UKI**

1698



© Thomas Savery

1712




© Thomas Newcomen

2009?



© Hitachi Rail Systems

1814







© George Stephenson

1923



© H. N. Gresley

# Everything digital – from the cradle to beyond the grave



Page 5

The information shown on this sheet within the boundaries of the


No. of	Serial No. of	Serial No. of	Serial No. of	Serial No. of	Serial No. of	Serial No. of	Serial No. of	Serial No. of	Serial No. of	Serial No. of
1	1	1	1	1	1	1	1	1	1	1
2	2	2	2	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	4	4	4	4
5	5	5	5	5	5	5	5	5	5	5
6	6	6	6	6	6	6	6	6	6	6
7	7	7	7	7	7	7	7	7	7	7
8	8	8	8	8	8	8	8	8	8	8
9	9	9	9	9	9	9	9	9	9	9
10	10	10	10	10	10	10	10	10	10	10
11	11	11	11	11	11	11	11	11	11	11
12	12	12	12	12	12	12	12	12	12	12
13	13	13	13	13	13	13	13	13	13	13
14	14	14	14	14	14	14	14	14	14	14
15	15	15	15	15	15	15	15	15	15	15
16	16	16	16	16	16	16	16	16	16	16
17	17	17	17	17	17	17	17	17	17	17
18	18	18	18	18	18	18	18	18	18	18
19	19	19	19	19	19	19	19	19	19	19
20	20	20	20	20	20	20	20	20	20	20
21	21	21	21	21	21	21	21	21	21	21
22	22	22	22	22	22	22	22	22	22	22
23	23	23	23	23	23	23	23	23	23	23
24	24	24	24	24	24	24	24	24	24	24
25	25	25	25	25	25	25	25	25	25	25
26	26	26	26	26	26	26	26	26	26	26
27	27	27	27	27	27	27	27	27	27	27
28	28	28	28	28	28	28	28	28	28	28
29	29	29	29	29	29	29	29	29	29	29
30	30	30	30	30	30	30	30	30	30	30
31	31	31	31	31	31	31	31	31	31	31
32	32	32	32	32	32	32	32	32	32	32
33	33	33	33	33	33	33	33	33	33	33
34	34	34	34	34	34	34	34	34	34	34
35	35	35	35	35	35	35	35	35	35	35
36	36	36	36	36	36	36	36	36	36	36
37	37	37	37	37	37	37	37	37	37	37
38	38	38	38	38	38	38	38	38	38	38
39	39	39	39	39	39	39	39	39	39	39
40	40	40	40	40	40	40	40	40	40	40
41	41	41	41	41	41	41	41	41	41	41
42	42	42	42	42	42	42	42	42	42	42
43	43	43	43	43	43	43	43	43	43	43
44	44	44	44	44	44	44	44	44	44	44
45	45	45	45	45	45	45	45	45	45	45
46	46	46	46	46	46	46	46	46	46	46
47	47	47	47	47	47	47	47	47	47	47
48	48	48	48	48	48	48	48	48	48	48
49	49	49	49	49	49	49	49	49	49	49
50	50	50	50	50	50	50	50	50	50	50

Total of Slides: 50

Serial of Slides and Frames: 10 / 20

Notes: - Check the page number and the serial of the slides as you register.

RCH TM




# New Enterprise Data Center Strategy

*Traditional  
Data Centers*



300,000 trades per trader per day

***The Globally Integrated  
Enterprise***



*Web 2.0  
Data Centers*



**You Tube**

10 hours upload per minute

***Business-to-consumer  
uber-portals***



**eBay**

**amazon.com**

**Google**

**Dynamic  
Infrastructure**



# Facing up to the Demands of the new Business World



**Optimize IT**  
Save money



**Leverage Information**  
Make money



**Manage Risk**  
Stay out of jail while protecting business



**Business Agility**  
plan to keep going in a changing world

*Information On Demand*



**Information Infrastructure**  
*Availability, Security, Retention, Compliance.*



**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 – it's a three way relationship





## Compliance – Whose job is it anyway?

### **Chief information security officer (CISO)**

The chief information security officer often reports to the **chief information officer** or even directly to the **chief executive officer**.

### **An Anti-money laundering officer (AMLO)**

Anti-money laundering guidelines came into prominence globally after the September 11, 2001 attacks and the subsequent enactment of the USA PATRIOT Act.

Ultimately the Buck stops here:

- **Chief Executive**
- **Chief Finance Officer**
- **Managing Director**
- **Company Secretary**

**Board members can go to jail when things go wrong!**

## Information Availability – Operational Definitions

- 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*







# Legal Availability



**28 Days for Information Notices**

DATA PROTECTION  
ACT 1998

**20 Days for General Access**



**1 working Day**



**2 working Days  
or  
48 hours**







# 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, GWYNNEETH 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.

*Gwynneeth G. Powell*  
Registrar of Births and Deaths.



**Verifiable**

**Immutable**

**Authentic**



## Retention of e-records

- **How does an organisation prove authenticity with *electronic* records?**
  - Policies
  - Technologies





- How long do you have to retain records?

5 years for Anti-Money Laundering purposes

Data<sup>3</sup> Protection Act  
Engineering Drawings and Plans  
as long as necessary  
50 years for ionization (inc X-Rays)



## Breach notification

- **40 US States have implemented breach notification laws, but what is the UK position – Does the Data Protection Act specifically require a breach notification strategy?**
  - Data Protection Act implies a requirement for a breach handling strategy by extension, through the Human Rights Act and other regulation
- **By April 2008 the Information Commissioner said he had been notified of 94 data breaches since the loss of the child benefit database – two thirds were committed by the Government and other public sector bodies**



It's not just about fines and jail sentences ...but also reputation



## ISO 27001 – The ONLY defence

- **There is now a significant “soft law” preference for ISO 27001 methodologies:**
  - ICO “our approach to encryption”:
    - “Personal information, which is stored, transmitted or processed in information, communication and technical infrastructures, should also be managed and protected in accordance with the organisation’s security policy and using best practice methodologies such as using the International Standard 27001.”

## Encryption – the only defence against Enforcement








IBM Systems and Technology Group

**So, what's happening to YOUR business?**

# The Information Explosion

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



**NOT GOING TO STOP**



Only 20% of data is structured. The rest is unstructured. Unstructured data includes text documents, spreadsheets, e-mail, photos, audio, video, and images. Unstructured data is growing rapidly and will continue to do so. **IT WILL GET WORSE**

Source: Various external consultant reports

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




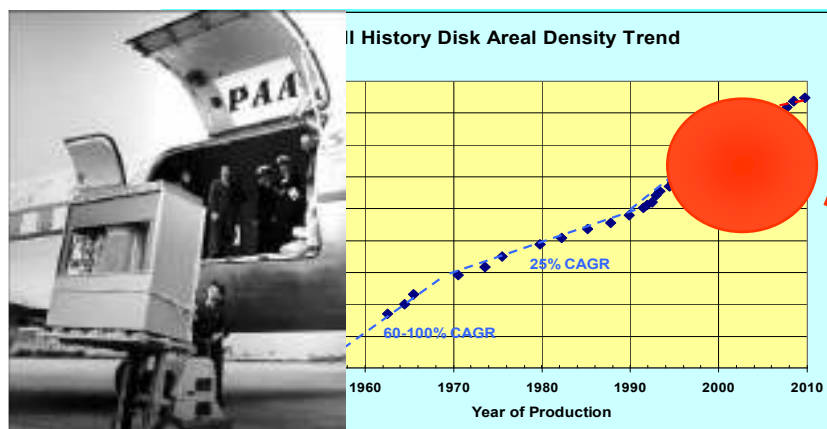
# Change in Disk Areal Density Trend

## Historic trend

- From 1957 until today CAGR of Areal Density has averaged 30-35%/yr
- Between 1957 and 1970 and again between ~1990 and ~2004 it averaged 60-100%/yr

## Current Decade

- Around 2004 the CAGR dropped to 25-35%/yr as we approach the *superparamagnetic* limit
- Disk vendors indicate that it will continue at 25-35% CAGR for the foreseeable future
- This will have a significant impact on the price and performance of disk drives

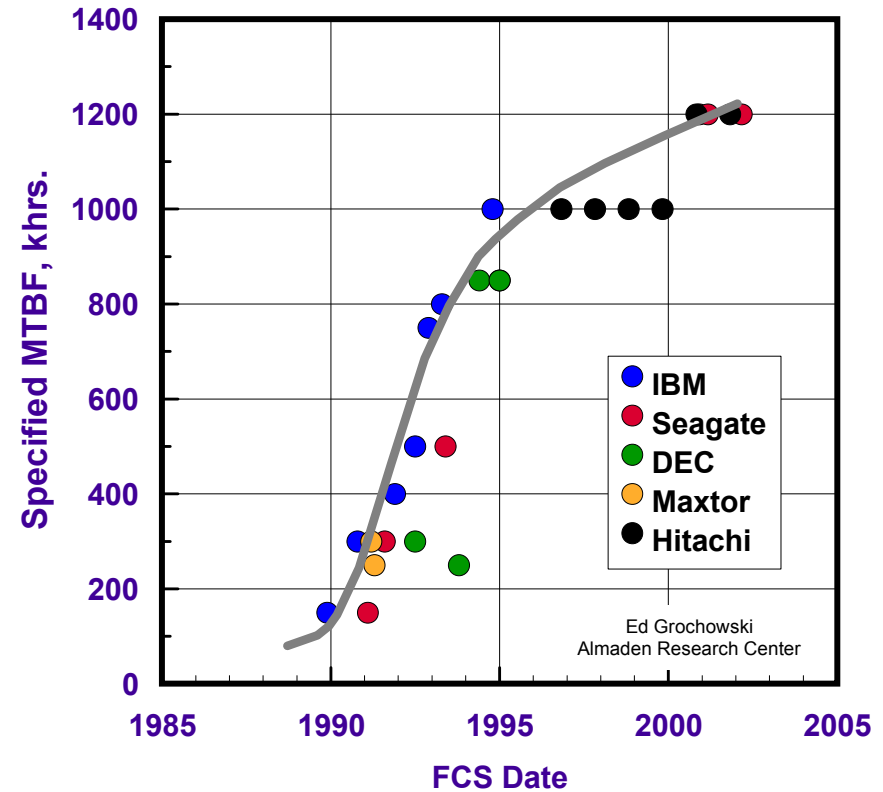
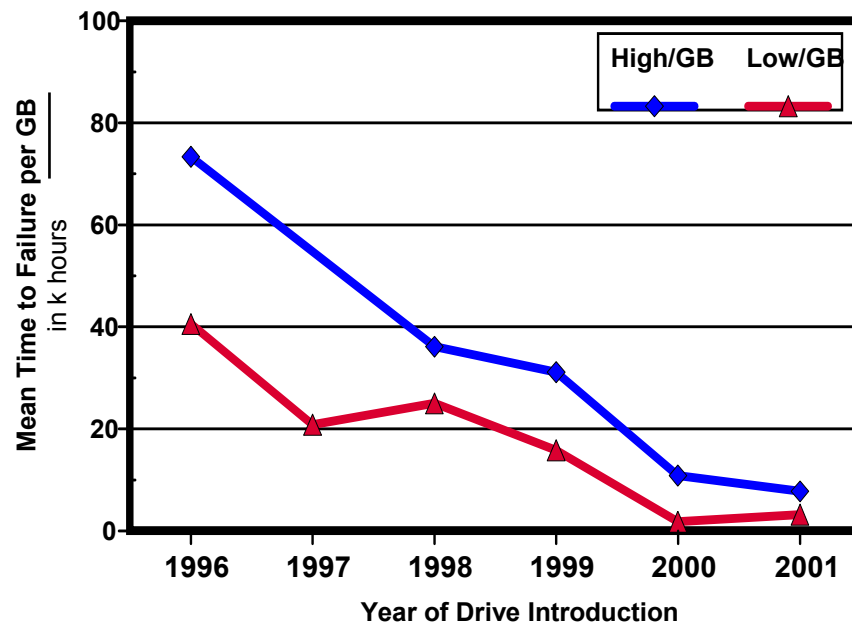


Sources—Historic: IBM & Hitachi, Current: disk vendors through IBM procurement



# Trends in Drive Reliability

- Drive reliability (MTBF and hard error rates) not keeping up with capacity growth.
  - Implies need for multi-parity RAID designs



HDD MTBF Manufacturer Specifications

## Emerging Trends

- **£ per GB or TB**
- **£ per IOPS**
- **kWatt per TB**
- **kgCO<sub>2</sub> per TB**
- **FTE per TB**
- **New style arrays**

# Emerging Disk Array Types

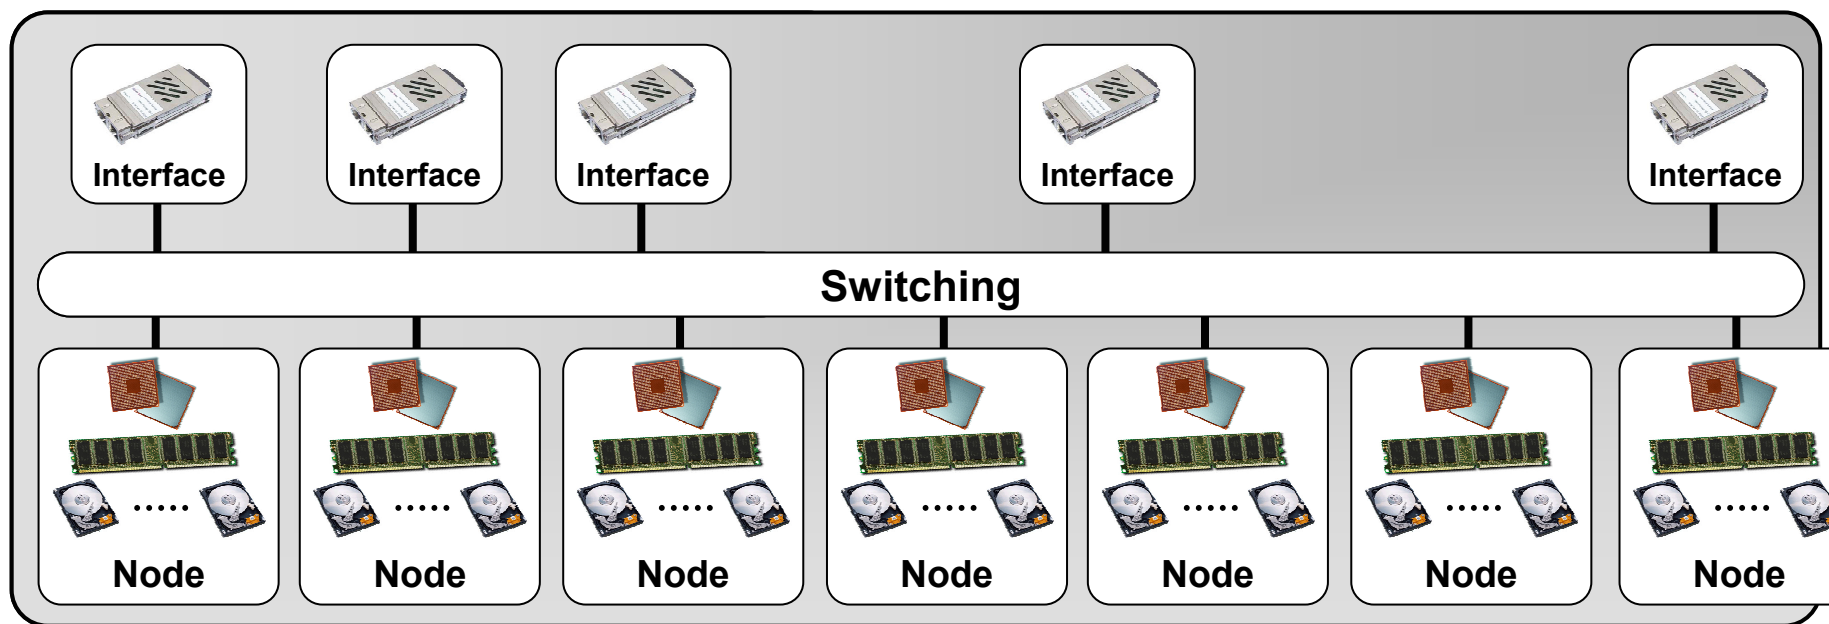
**Scale Out**

▪ **Clustered Modular**

- Clusters of standard/modular components, with single storage system image
- Usual goal – build mid/large/very large scale arrays with Entry or Modular controllers and drive packages

**Design principles:**

- Massive parallelism
- Off-the-shelf components
- Coupled disk, RAM and CPU
- User simplicity




## Emerging Trends

- **£ per GB or TB**
- **£ per IOPS**
- **New Style Arrays**
- **Solid State Disks**

# SSD Flash Storage Design 1

## Place SSD At Same Spot In Hierarchy As HDD

- Combine DRAM and Flash
  - Mask slow writes
  - Spread writes to minimize cell wear
- Use existing Storage Controller
  - Not optimized for SSDs – a limited number of SSDs will match controller performance
- Can provide very high IOP rate to a small storage capacity
  - e.g. 100K IOPS/sec to 1TB of data
  - Good match for very high performance database applications
    - Compare to spreading data across many HDDs to achieve equivalent Ops/sec



• Very limited applications currently

• Only Enterprise Class where “speed counts”



## SSD Flash Storage Design 2

### SSD In Front Of The Storage Controller

- **Combine DRAM and Flash**

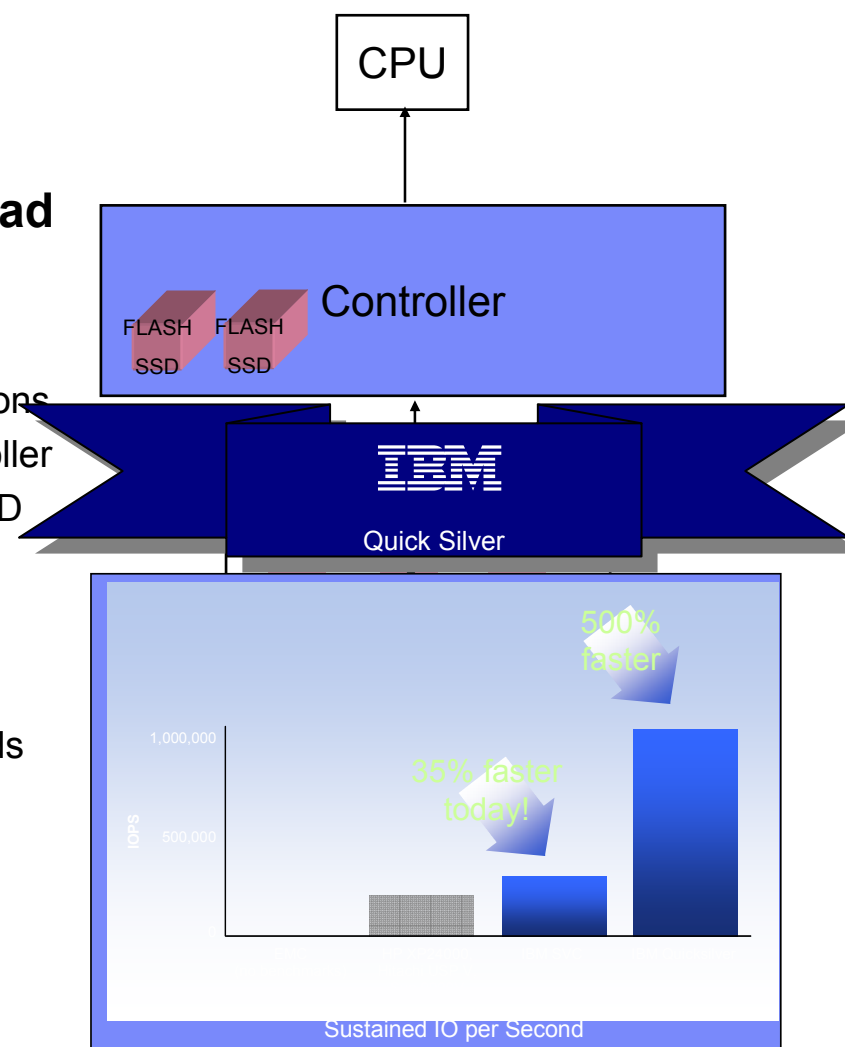
- Mask slow writes
- Spread writes to minimize cell wear

- **Minimize Storage Controller Overhead**

- Improve Performance
  - IOPS/sec
  - Response time
- Use SSD's Controller Cache and Copy functions
- Mirroring or RAID of SSDs provided by Controller
- ILM – migration of data between SSD and HDD


- **Can provide very high IOP rate to a small storage capacity**

- Impacts Open Applications
  - Largest benefit to cache unfriendly workloads
  - Good match for very high performance database applications



# SSD Flash Storage Design 3 - SSD Based Systems

- Include Non-Volatile Memory as part of Server Memory System
  - Potential large benefit to certain applications
  - Requires significant modifications to Middleware or OS to achieve benefit
  - Use in Database, Search, Indexing applications



## Emerging Trends


- **£ per GB or TB**
- **£ per IOPS**
- **kWatt per TB**
- **New Style Arrays**
- **Solid State Disks**
- **Low Energy Options**

# 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
- Data De-Duplication



# Average Storage Cost Trends



Sources: Disk - Industry Consultants, SATA Estimate based on Market data, Tape - IBM



## Emerging Trends

- **£ per GB or TB**
- **£ per IOPS**
- **kWatt per TB**
- **kgCO<sub>2</sub> per TB**
- **New style arrays**
- **Solid State Disks**
- **Low Energy Options**
- **Footprinting**



# Are we solving Symptoms or Problems?



**Data centre  
disruption  
'to hit UK in  
five years'**

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



**Is it a  
GREEN  
issue?**



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

November 24<sup>th</sup> 2008

# The output from a SPACE Footprint energy efficiency engagement

## Power and Environmental Comparison

<b>One Year Savings</b>	Power	£4399	21.38%	Floor Space	£-222	-19.65%	Maint 1 Yr	£2000	Calculated: 14/04/2008 Version: 1.81
	Cooling(BTU)	145,012	26.22%	CO <sub>2</sub> (kg)	21,019	21.38%	Maint 5 Yr	£10000	

Power	Base kVA		1 Year Usage		One Year Cost		Three Year Cost		Five Year Cost	
	Current	Proposed	Current	Proposed	Current	Proposed	Current	Proposed	Current	Proposed
Fabric	0.0	0.2	0	1,577	0	142	0	426	0	710
Storage	17.4	13.5	152,424	118,260	13,718	10,643	41,154	31,930	68,591	53,217
<b>Totals</b>	<b>17.4</b>	<b>13.7</b>	<b>152,424</b>	<b>119,837</b>	<b>£13718</b>	<b>£10785</b>	<b>£41154</b>	<b>£32356</b>	<b>£68591</b>	<b>£53927</b>

Floor Space	Base Floor Space		One Year Cost		Three Year Cost		Five Year Cost	
	Current	Proposed	Current	Proposed	Current	Proposed	Current	Proposed
Fabric	0.0	0.0	0	0	0	0	0	0
Storage	2.3	2.0	1,130	1,002	3,390	3,006	5,650	5,010
19" Rack	0.0	0.7	0	350	0	1,050	0	1,750
<b>Totals</b>	<b>2.3</b>	<b>2.7</b>	<b>£1130</b>	<b>£1352</b>	<b>£3390</b>	<b>£4056</b>	<b>£5650</b>	<b>£6760</b>

Cooling Requirements	Base kBTU		One Year kBTU		Three Year kBTU		Five Year kBTU	
	Current	Proposed	Current	Proposed	Current	Proposed	Current	Proposed
Fabric	0.00	0.51	0	4,485	0	13,454	0	22,423
Storage	63.14	46.07	553,106	403,610	1,659,319	1,210,829	2,765,532	2,018,048
<b>Totals</b>	<b>63.14</b>	<b>46.59</b>	<b>553,106</b>	<b>408,094</b>	<b>1,659,319</b>	<b>1,224,282</b>	<b>2,765,532</b>	<b>2,040,471</b>

Carbon Footprint	Base (kW/year)		One Year Output (kg)		Three Year Output (kg)		Five Year Output (kg)	
	Current	Proposed	Current	Proposed	Current	Proposed	Current	Proposed
<b>Totals</b>	<b>152,424</b>	<b>119,837</b>	<b>65,542</b>	<b>51,530</b>	<b>196,627</b>	<b>154,586</b>	<b>327,712</b>	<b>257,649</b>

Energy Efficient Architectural Building Block		Offset Calculation - U Value conversion			
		Today		Proposed	
		Fabric	Storage	Fabric	Storage
Rack 'U's		0	0	2	0
Racks		0.00	0.00	0.05	0.00
Combined			0.00		0.05
			0		1

■ **Focuses on reducing waste and improving reclamations**  
 Virtualisation, Rationalisation and Right Tiering  
 Improved Capacity Management and Information Management  
 — Alignment of storage services to business requirements

## Emerging Trends

- **£ per GB or TB**
- **£ per IOPS**
- **kWatt per TB**
- **kgCO<sub>2</sub> per TB**
- **FTE per TB**
- **New style arrays**
- **Solid State Disks**
- **Low energy options**
- **Footprinting**
- **Virtualisation**

# Universal Application of the Virtualized Information Infrastructure



Users

**Desktop Infrastructure Virtualization Applications**

*physical server*

virtual server   virtual server   virtual server

**Server Virtualization**

*physical server*

virtual server   virtual server   virtual server

**Server Virtualization**

**File Virtualization**

Storage network(s)

**Network Virtualization**

**Disk Storage Virtualization**

**File System Virtualization**

**Tape Storage Virtualization**

Drives

Key Principles

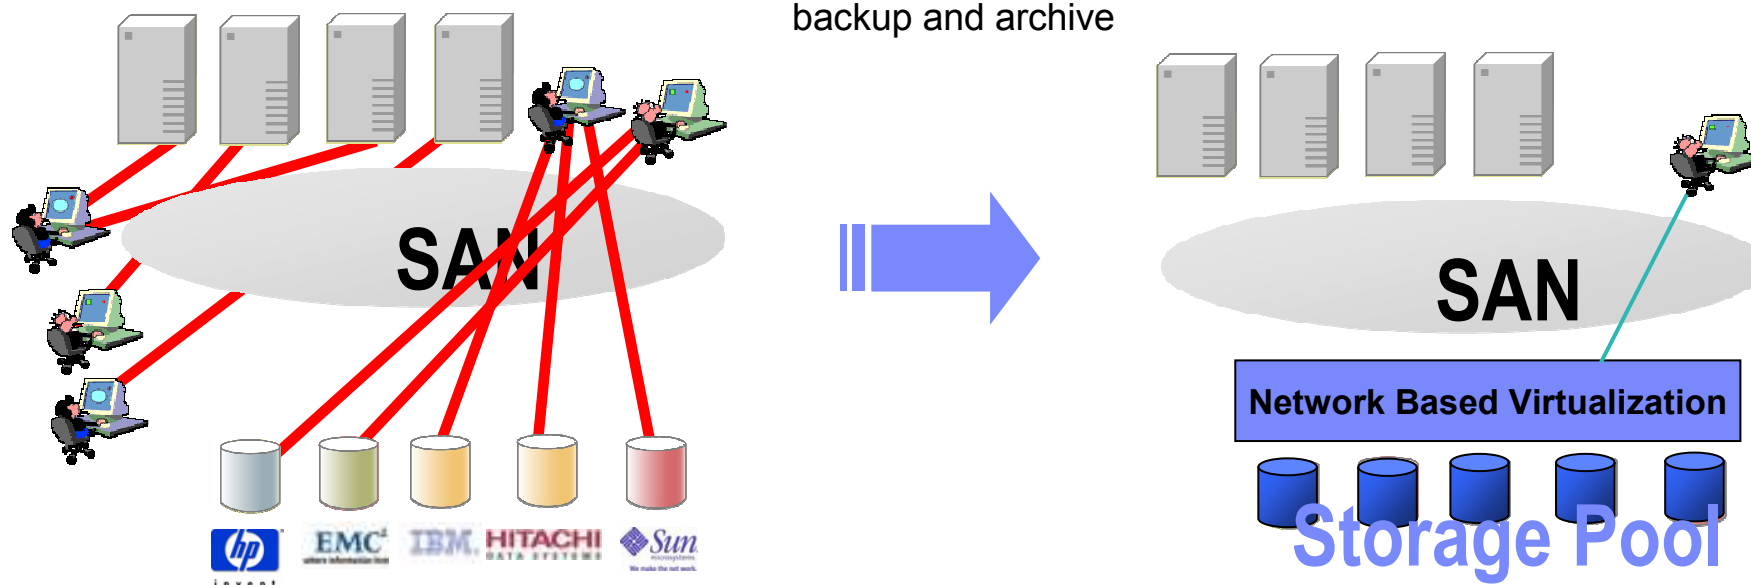
- Comprehensive
- Open
- Heterogeneous
- Common skills

Virtual/Physical Management

## The SAN Challenges tackled by Virtualisation

- **Complexity**
- **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



## Why is SVC Important?

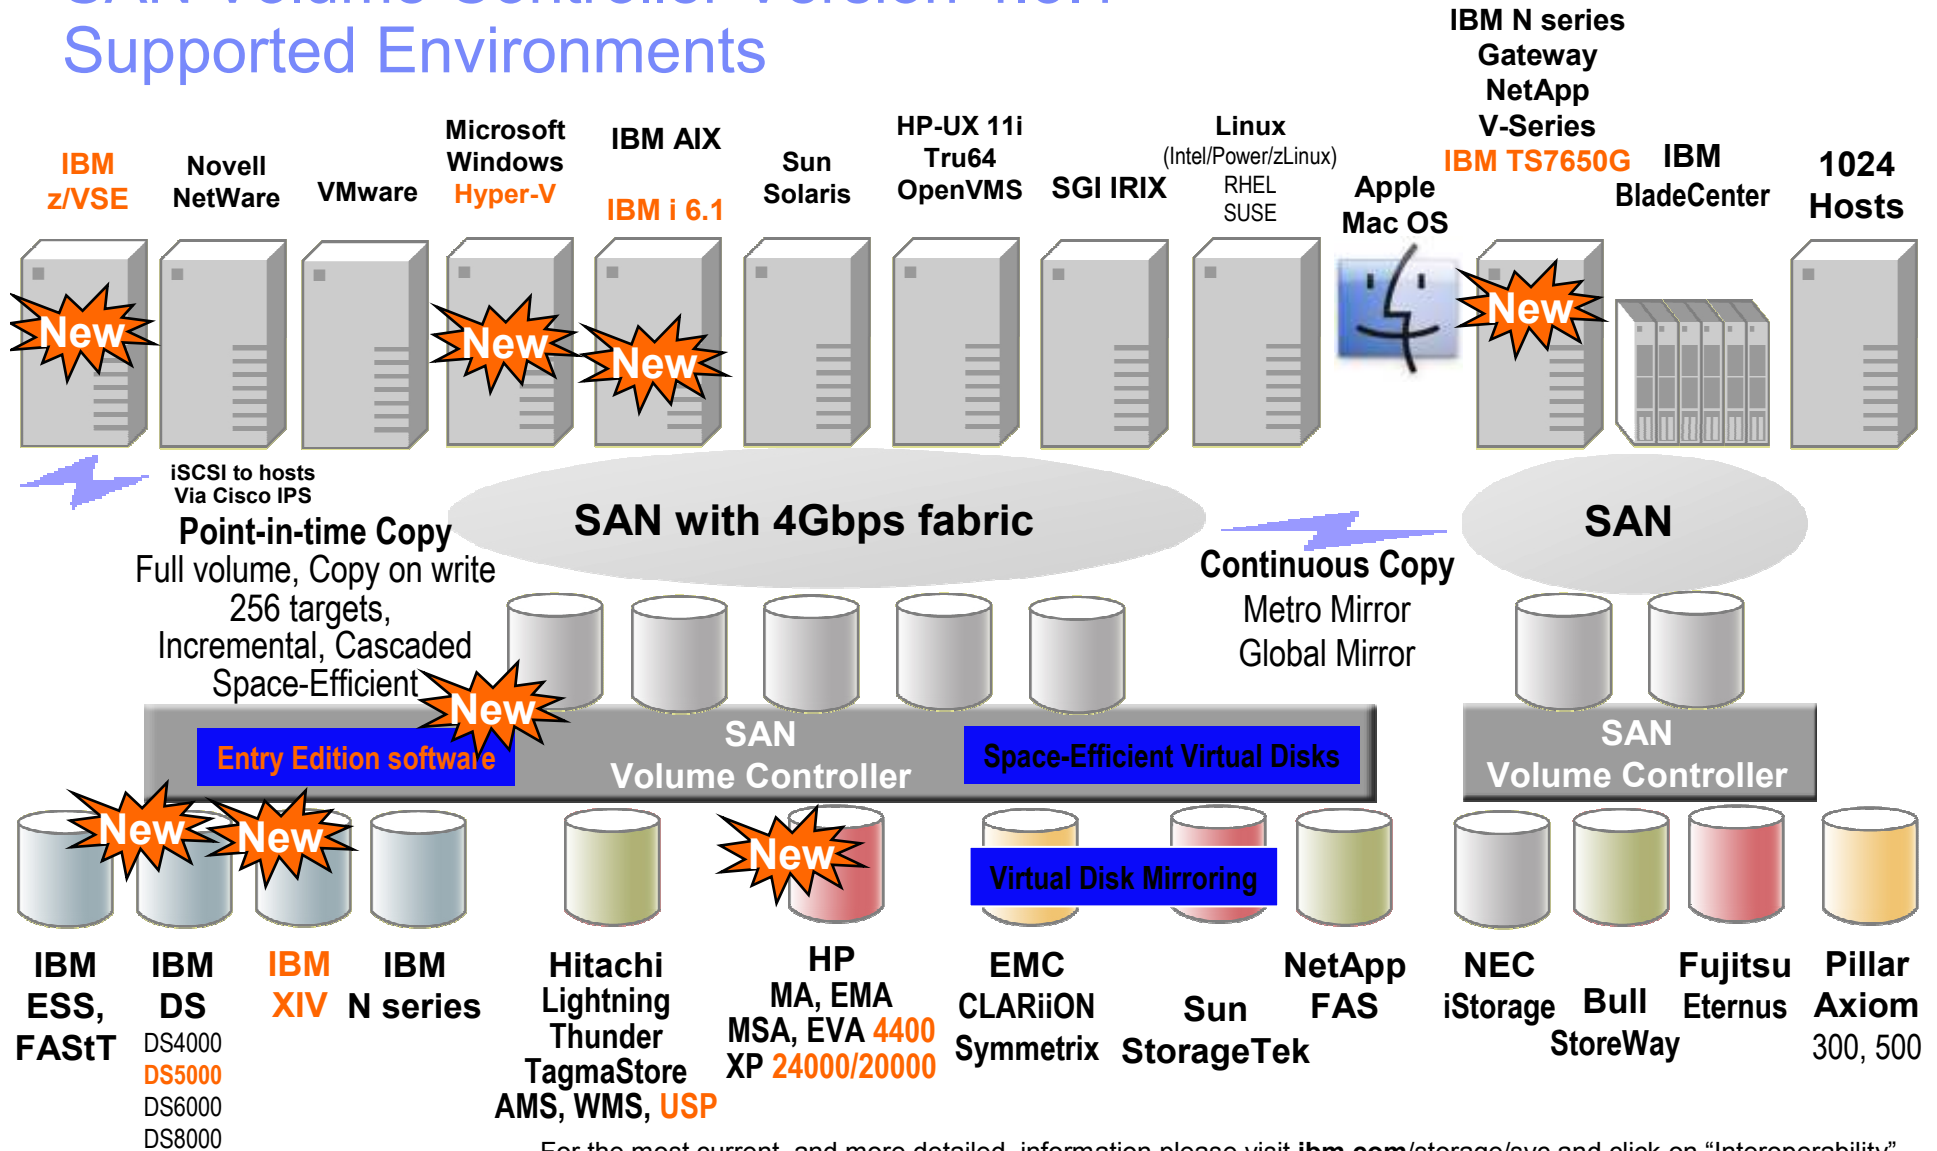


### ***Overall, SVC helps reduce storage cost***

- **Helps improve storage utilization**
  - Make better use of existing storage and control growth
- **Designed to improve application availability**
  - Make changes to storage and move data without taking applications down
- **Helps simplify management**
  - Greater efficiency and productivity for storage management staff
- **Offers network-based replication**
  - Helps enable greater choice when buying storage



# SAN Volume Controller Version 4.3.1 Supported Environments




For the most current, and more detailed, information please visit [ibm.com/storage/svc](http://ibm.com/storage/svc) and click on "Interoperability".

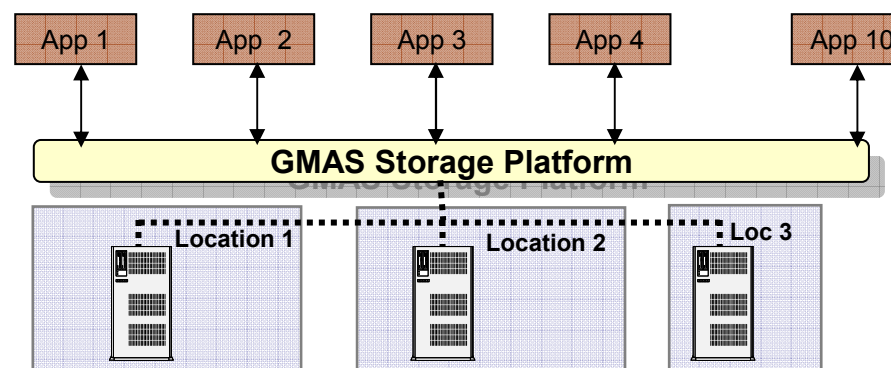
## Virtualized Infrastructure

GMAS is a storage solution that intelligently manages the interaction between applications and storage resources via client business rules

### Before GMAS



### After GMAS



#### “Siloed” infrastructure:


- Difficulty sharing resources across applications
- Requires application downtime for maintenance
- Manual administration, upgrades & conversions
- Inherently vulnerable to storage failures

#### “Virtualized” infrastructure:

- Collapses silos into a single shareable storage pool across applications
- Enables maintenance, support & data conversion without application downtime
- Automates upgrades, conversions & administration
- No single point of failure


## Cloud Computing – the ultimate virtualized destination

- **Everyone will go there... eventually**
- **Trains took off when they were used by Consumers**
  - No one user owns the train, the coaches, rail track or even the drivers
  - This will offer true “Utility” computing
  - Generational change required....it’s a matter of thinking
    - Hotmail
    - Google
    - Amazon
    - Ebay
    - Facebook






# Solid state storage – Emerging technologies



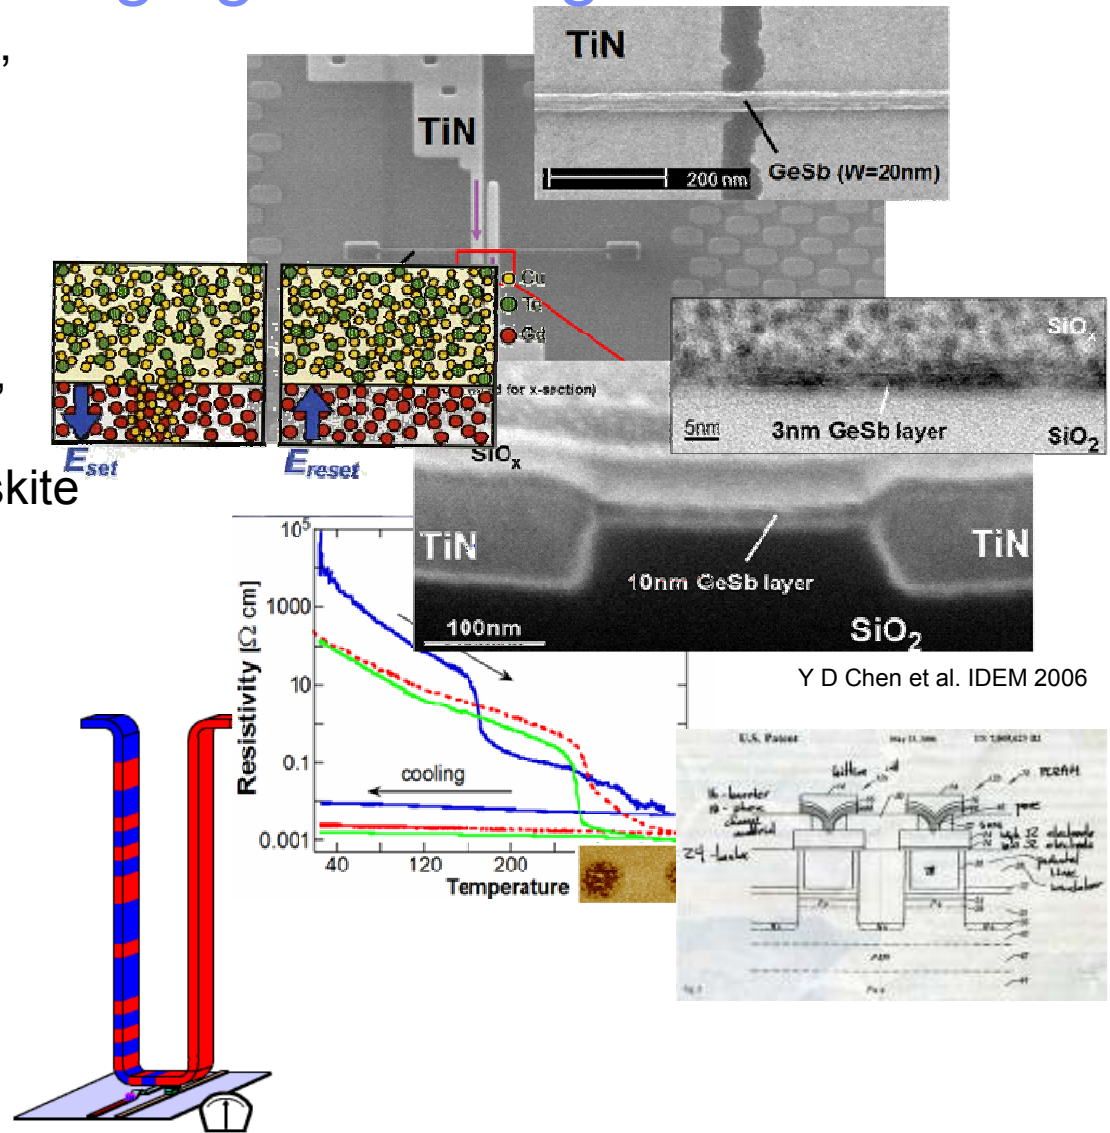
Quantity: FRAM, PFRAM, SiC Bipolar, Molecular, NanoCrystal



Flow: PCRAM, PMC, Polymer, RRAM, Perovskite




Spin: MRAM, Magnetic Shift Register



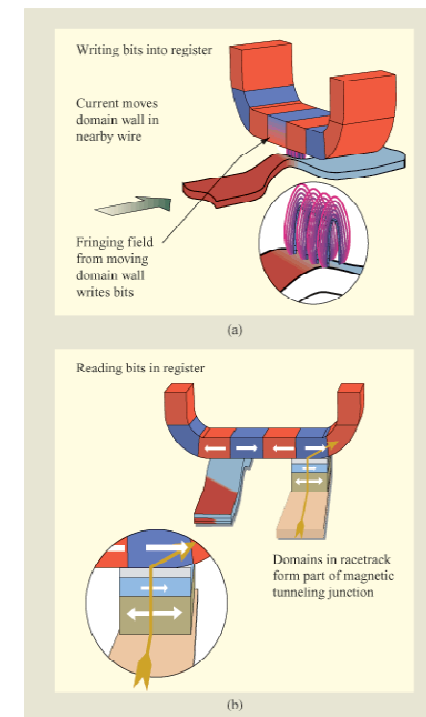



# Latest in Storage Research : "RaceTrack" Memory




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


GMR = giant magneto-resistivity  
Nobel prize for Peter Gruenberg (D) and Albert Fert (F), 2007







# Why choose IBM?



Technology as a  
**SHARED  
SERVICE**

### Service means:

- Use, not buy
- Virtualized
- Consumption based costing
- On demand
- Low cost of entry
- Shared infrastructure

### Safe Harboring means:

- IT Hedging
- “Good enough” technology
- Cost predictability
- Fewer relationships
- Simpler answers
- Reduced Risk



## **SAFE HARBOURING**

Drive unit costs down thru  
trusted partnering



Massive  
**EFFICIENCY**  
programmes

### Efficiency programmes means:

- Search for gold
- Re-use
- Reclamation
- De-duplicate
- Eradicate
- Share
- Re-purpose



IBM Systems and Technology Group

# Thank You

**Rick Terry**  
Storage Solutions Sales Specialist  
UKI S P A C E Programme  
Contact Details +44 77 25 70 64 81  
[ricterry@uk.ibm.com](mailto:ricterry@uk.ibm.com)

