




Industry Trends and Perspectives: *State of Storage V2.009*

Greg Schulz, Founder & Sr. Analyst -The StorageIO Group
Author "The Green and Virtual Data Center" (Auerbach)
And "Resilient Storage Networks" (Elsevier)



Introduction

Who is Greg Schulz and StorageIO

- Former IT customer/user
 - Electrical power G&T utility, transportation, financial services
 - Servers, storage, networks, hardware and software (Cap. plan/BC/DR)
- Recovering Ex-vendor
 - Manufacturer and VAR, hardware, software and services
 - Storage (SSD, RAID, NAS, VTLs, tape, backup & archiving & HSM)
 - Storage networking (mainframe and open, LAN, MAN, WAN, SAN)
- Industry analyst and involvement with trade groups
 - Involved with SNIA and other trade and vendor groups
 - Worked for an analyst firm before founding StorageIO
 - Real-world analysis based on experience talking with people like you
- Author, presenter and blogger (www.storageio.com & www.storageioblog.com)
 - Two solo books, contributed to several others
 - Thousands of articles, tips, columns, papers and reports
 - Present and keynote speaking on a global basis

Industry Trends: Myths and Realities

These are constantly changing and evolving


- Myth: Storage efficiency is all about capacity optimization
 - Reality – It is also about boosting performance and productivity
- Myth: Tape is dead, tape is root of data protection problems
 - Reality – Tape is alive co-existing with disk to disk backup
- Myth: BC/DR are only issues for large enterprises
 - Reality – Applicable to organizations of all sizes
- Myth: Energy efficiency is energy avoidance (e.g. turn things off)
 - Reality – Also about boosting effective performance/productivity
- Myth: Virtualization is only for consolidation
 - Reality – Also for emulation (e.g. VTLs) and management

Industry Trends: Increased Threat Risks

More data, more risks, more reliance and more awareness

Threat Risks

- Headline news
- Non-headline news
- Local and Regional
- Acts of Nature
- Acts of man
 - Accidental
 - Intentional
- Technology
 - Failure
 - Miss-configuration



Issues, challenges and drivers


- Issues no longer constrained to rich and famous
- Continued dependence on information services
- Reliance on Information and services availability
- Increased visibility and awareness of issues
- Regulation and compliance, data loss prevention



Industry Trends: I/O Performance Gap

Disk I/O lags server performance and storage capacity

Sustaining IT and Business Growth, Maintaining QoS



Action Items

- ✓ Address I/O issues/problems
- ✓ Storage & I/O optimization
- ✓ User fewer, yet fast devices
- ✓ Balance SSD (RAM & FLASH)
- ✓ Fast 15.5K SAS & FC disks
- ✓ Data footprint reduction

- Server processor performance curve
- Disk storage capacity curve
- - - Disk storage performance curve (IOPS)


See "Data Center Performance Bottlenecks" www.storageio.com



Industry Trends: Expanding Data Footprint

Management: Counter impact of expanding data footprint

Sparse, Duplicate
Files and Content



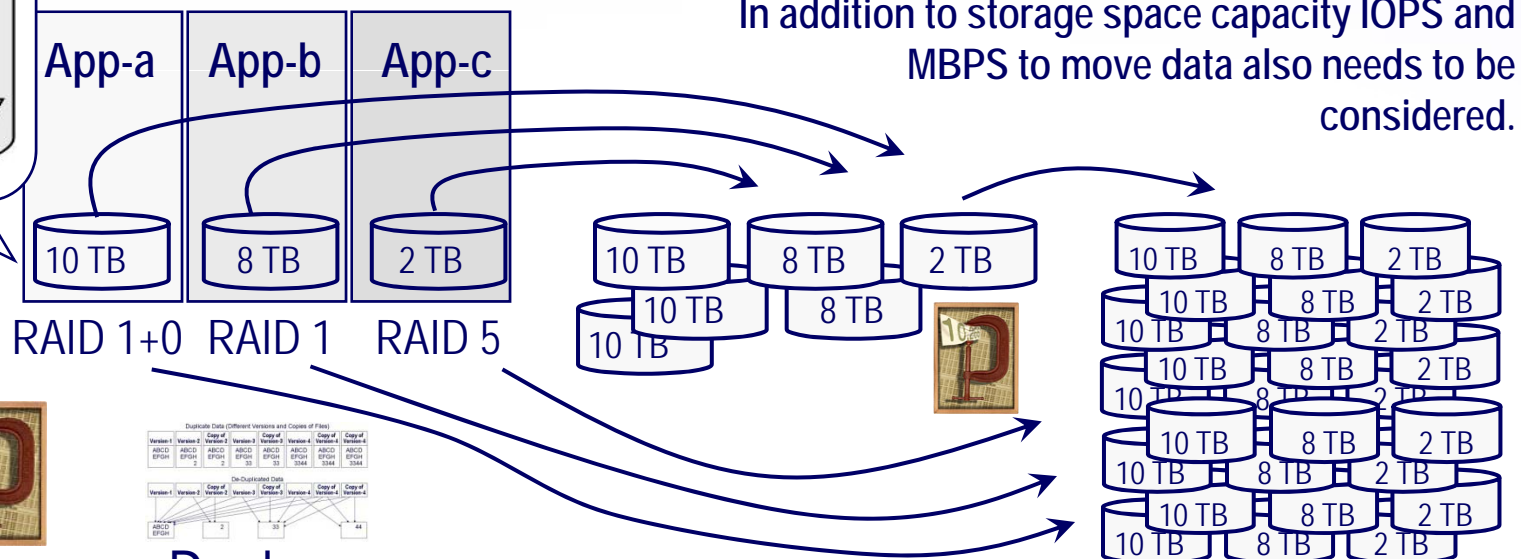
Archive Compress



De-dupe

Original Data
Primary Database,
Email, File serving

Copies – Data Proliferation and expanding data footprint
DSS, Training, Test, Dev
QA, Operational, Needs
Backup, BC, DR, HA,
Archive, Compliance




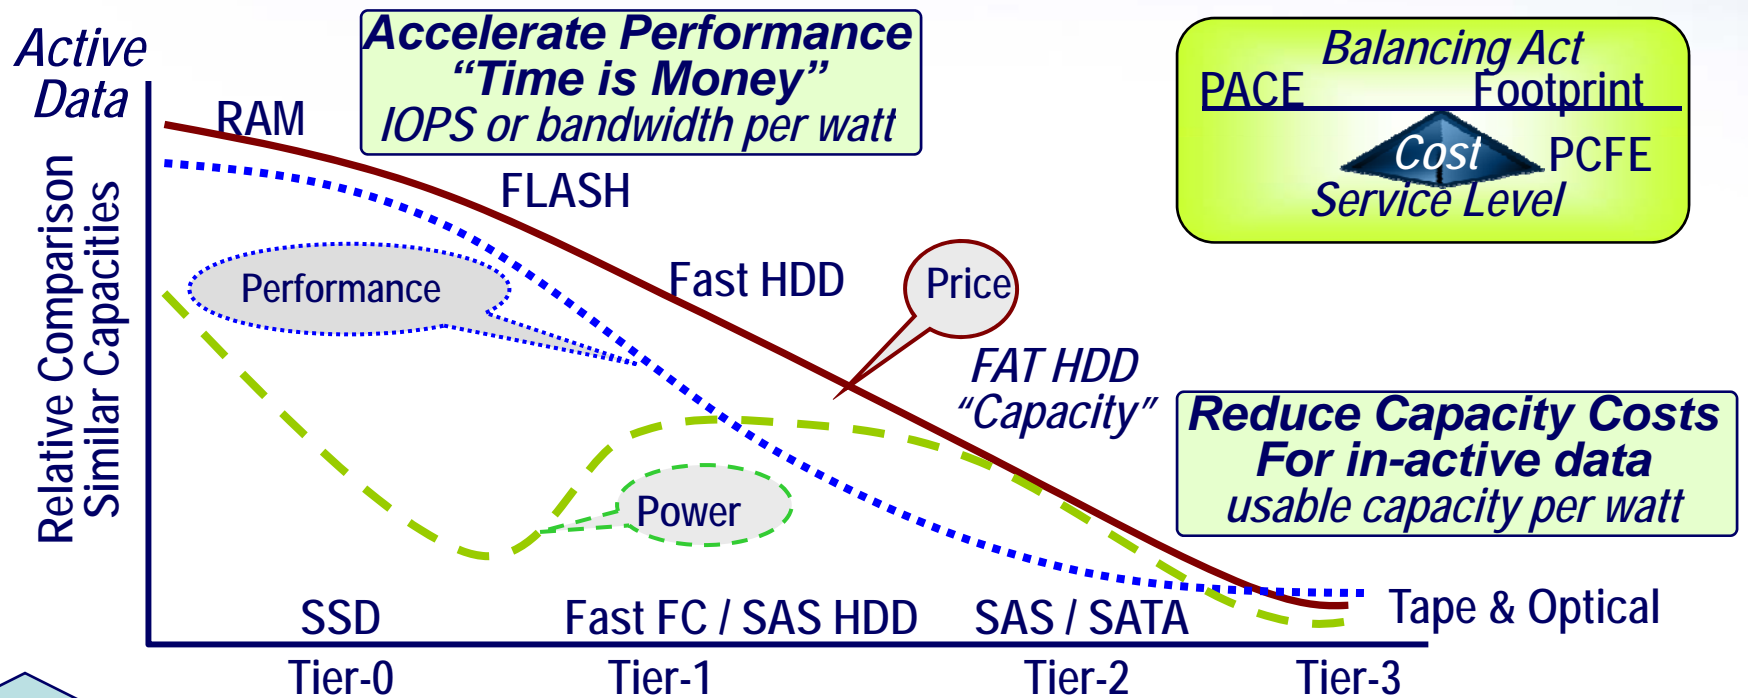
Challenge: More data to backup, protect and manage

Action: Reduce footprint impact: Archive, Compress, De-dupe, Tiered Storage

See “Business Benefits of Data Footprint Reduction” www.storageio.com

Industry Trend: Tiered Resources

Balance Performance, Availability, Capacity, Energy (PACE)



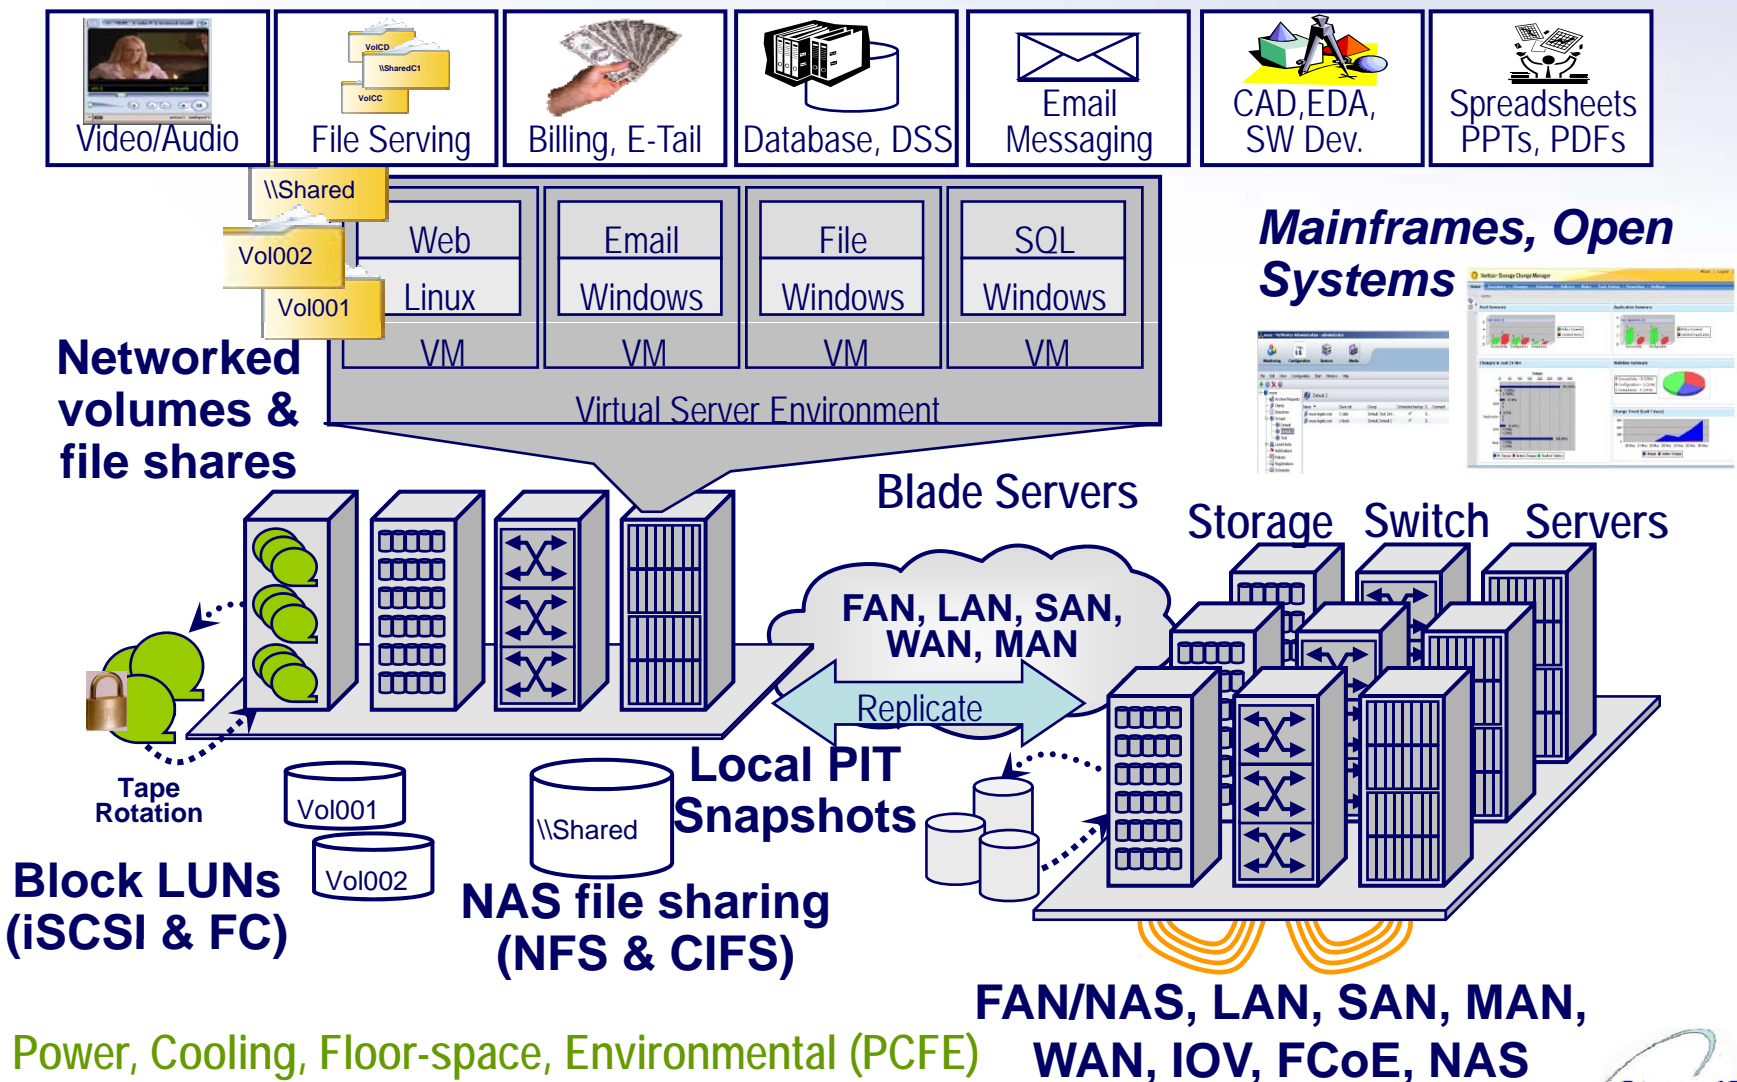
See "Achieving Efficiency with SSD" reports and articles at www.storageio.com

Source: "The Green and Virtual Data Center" (Auerbach)



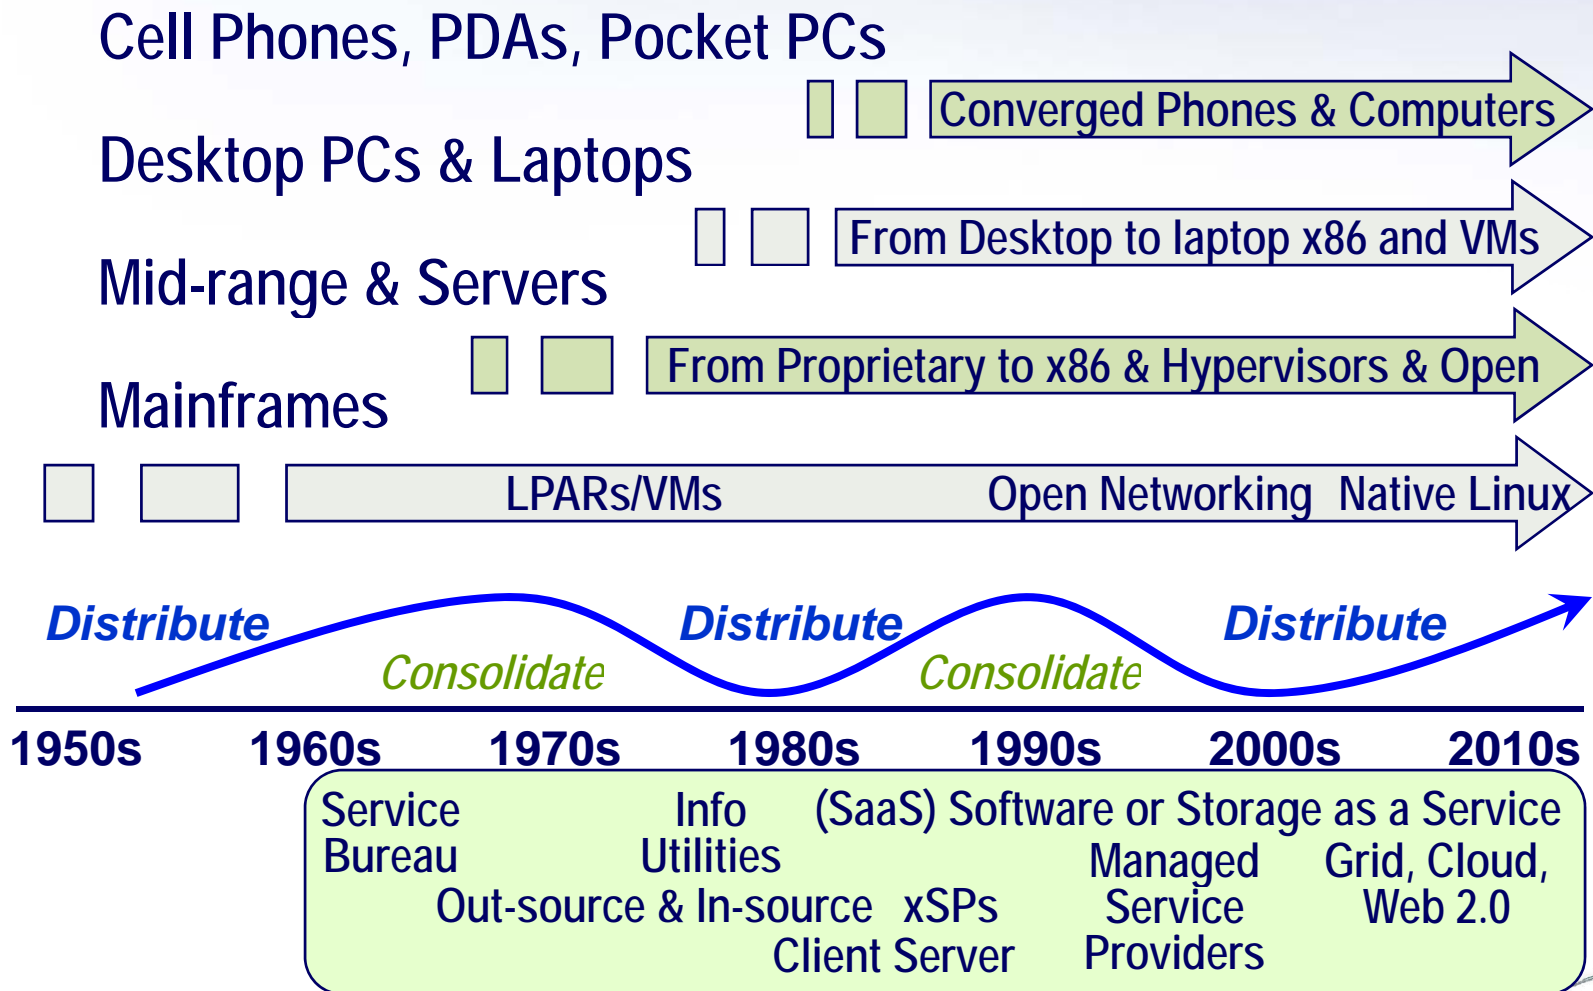
Industry Trends: Virtualization

Current focus is consolidation of physical resources



Industry Trends: Compute Continuum

We are in a consolidation phase (**Again!**)



What's The “Buzz” Out There?

What I commonly hear from IT customers and others

- **Interoperability:** With scalability and stability, hardware and software
- **Cost avoidance:** Doing more with less, supporting growth, tactical vs. strategic
- **Tape is not dead (yet):** Being used for performance and capacity (larger sites)
- **Dedupe scalability concerns:** Compliance, RAS & performance (large sites)
- **Skepticism or confusion:** Green, FCoE, SSD, Clouds, Clusters & Grid
- **Limited “broad” adoption:** Thin provision, MAID, Storage Encryption
- **Expanding data footprint:** Impact on storage and management costs
- **Confusion of price band, tiers & categories:** SNIA taxonomy opportunity!
- **Concern over degrading service levels due to consolidation**
- **Reducing software footprint impact:** Software and application consolidation
- **Data protection and DLP:** For virtual and physical environments
- **Who will manage:** Virtual & converged networks, converged servers & storage

Industry Trends: Lack of Innovation?

That depends on what your definition of innovation is!

What determines innovation?

- Ideas and technology?
- Number of existing startups?
- Level of marketing and awareness?
- Press/media and analyst coverage?
- Number of customer beta "seed" sites?
- Amount of paying Customer deployment?



Do these represent Innovation, you be the judge!
Proverbial "turn the crank", FCoE/CEE/DCE, MR-IOV/SR-IOV,
Open source, server/storage convergence, green tech,
transactional flash, tape, disk, grid, cluster, cloud, RAID,
BC/DR/DLP/Data protection and other IRM software storage
systems and architectures.

Next Generation Data Centers

What the “Information Factory” of the future will look like

New and Emerging



Time Tested and Field Proven



Balance of new and old technologies

SSD (RAM & FLASH)
Clusters & Grids
Clouds, policy based
automation and
Virtualization, tiered
access (FCoE, Object
Access), DLP

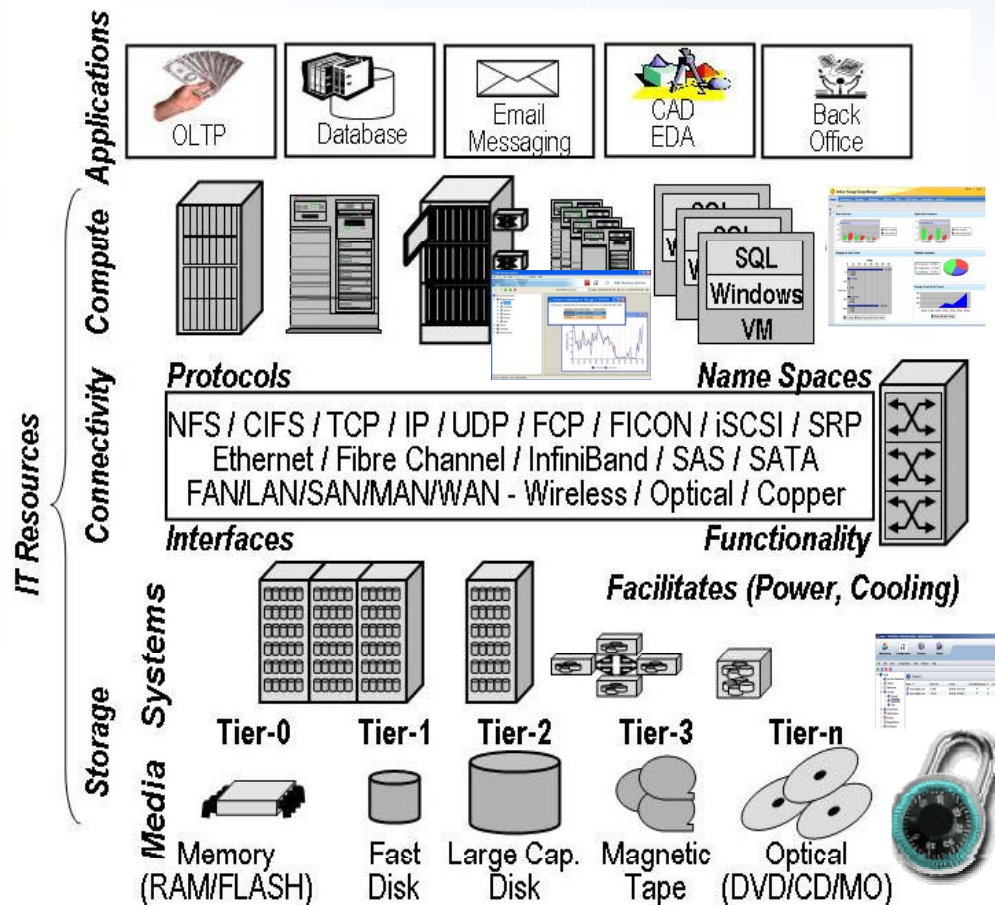
Magnetic disks and tape,
clusters and tiered
storage, RAID, tiered
data protection,
encryption, tiered access
(FC, iSCSI, SAS, NAS)

- Highly optimized, secure, resilient and flexible, dense resources, (beyond just consolidation), automated, efficient enabling more work and productivity to be done
- Clouds are complimentary to traditional IT resources as another tier of resource for IT service delivery
- Keep in mind that virtual, cloud, grid and SOA solutions require physical resources, software, tools, facilities and people!

Cross Technology Domain Management

Server, storage, networking and software convergence

IRM = Software, tools, procedures, policies, best practices



Infrastructure Resource Management (IRM) functions and activities **Processes and Tools**

- Namespace and virtualization
- Measurements and metrics
- Monitoring and reporting
- Modeling, analysis, planning
- Resource usage and allocation
- Performance and capacity plan
- Thin provisioning and purposing
- Diagnostic and resolution
- Change & configuration validation
- Data protection and footprint reduction
- Policy management and service levels
- Facilities and asset management
- Logical and physical security
- Procurement and disposition

Source: "The Green and Virtual Data Center" (Auerbach)

See "Data Protection Options for Virtualized Servers" www.storageio.com

Next Generation Data Centers

Zombie “Declared Dead” technologies still being bought & used

RIP – Dead or dying

Floppy disk drive, parallel SCSI, Y2K and DotCom SSPs, Britton Lee database machine, punch cards, etc...

Endangered List

Green Hype, Optical & CDs, FANs and SMIS among others

“Zombies” - Declared Dead (vendors sell, customers buy)

Magnetic tape, magnetic disk, backup, IBM “z” mainframe, HDDs, DEC/Compaq/HP VMS, PC, Fibre Channel & FICON, RAID, Windows, storage arrays, printers, etc...



Prediction: Tape will be “Trendy” in 2009


“Zombie Technologies” are profitable for vendors, productive for customers, boring for marketing (thus limited or no coverage)

Opportunity: *Make some news to remain relevant!*

Industry Trends: Shifting Landscapes

Changing data lifecycle and access patterns

Legacy and Transactional Data




Profile = Data is created, worked with and then goes dormant after some period of time with probability of little to no future access or use

Examples = Database, Email, Transactional, general file serving, project oriented data

Action = Ideal candidate for archiving off of primary or on-line storage too off-line and removable media or MAID 2.0 & IPM based storage combined with purging or deletion of data no longer needed to meet compliance or other commitments

Web 2.0 and On-line Data



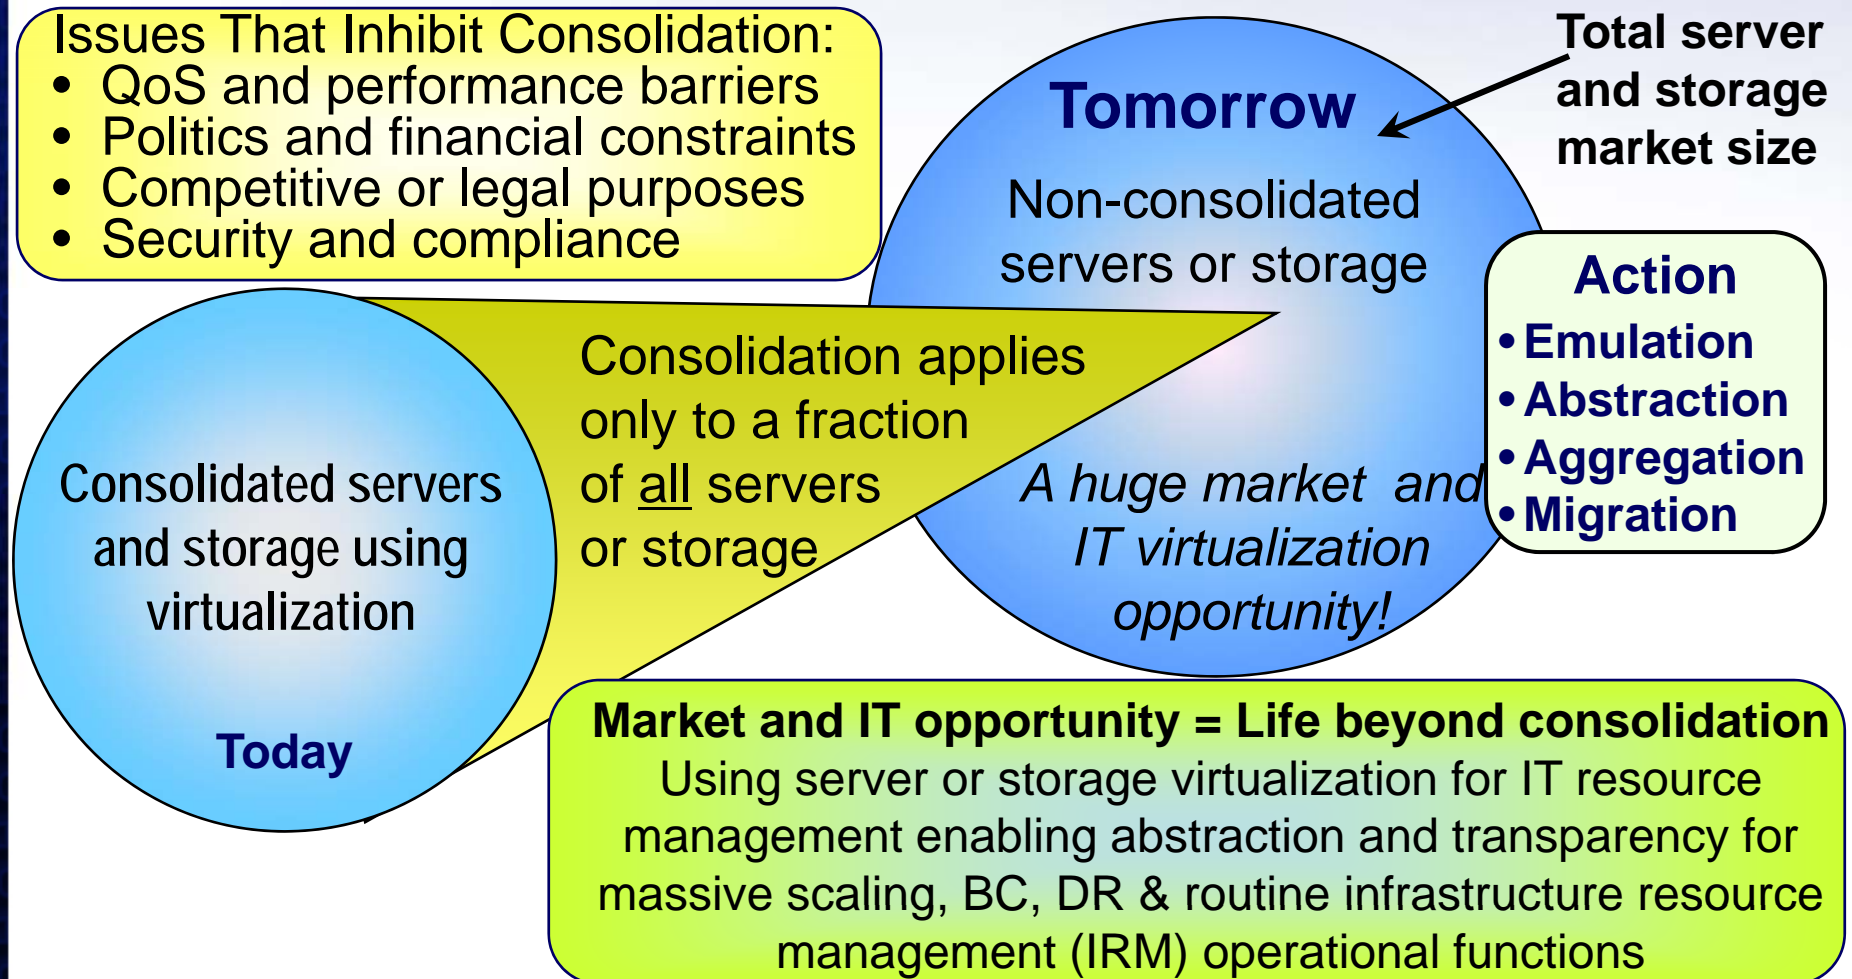
Profile = Data is created, worked with and then may go idle briefly, then accessed, then idle, then active, then idle, then active...

Examples = Web, Reference and lookup, fixed content, Web 2.0 and social networking, media and entertainment, some email, search, seasonal or event and research based data

Action = On-line storage with variable performance to meet changing workload demands, bulk and clustered storage, IPM enabled storage and storage caching

Industry Trends: Virtualization (Next Phase)

Life beyond consolidation: Abstract, Emulate and Management



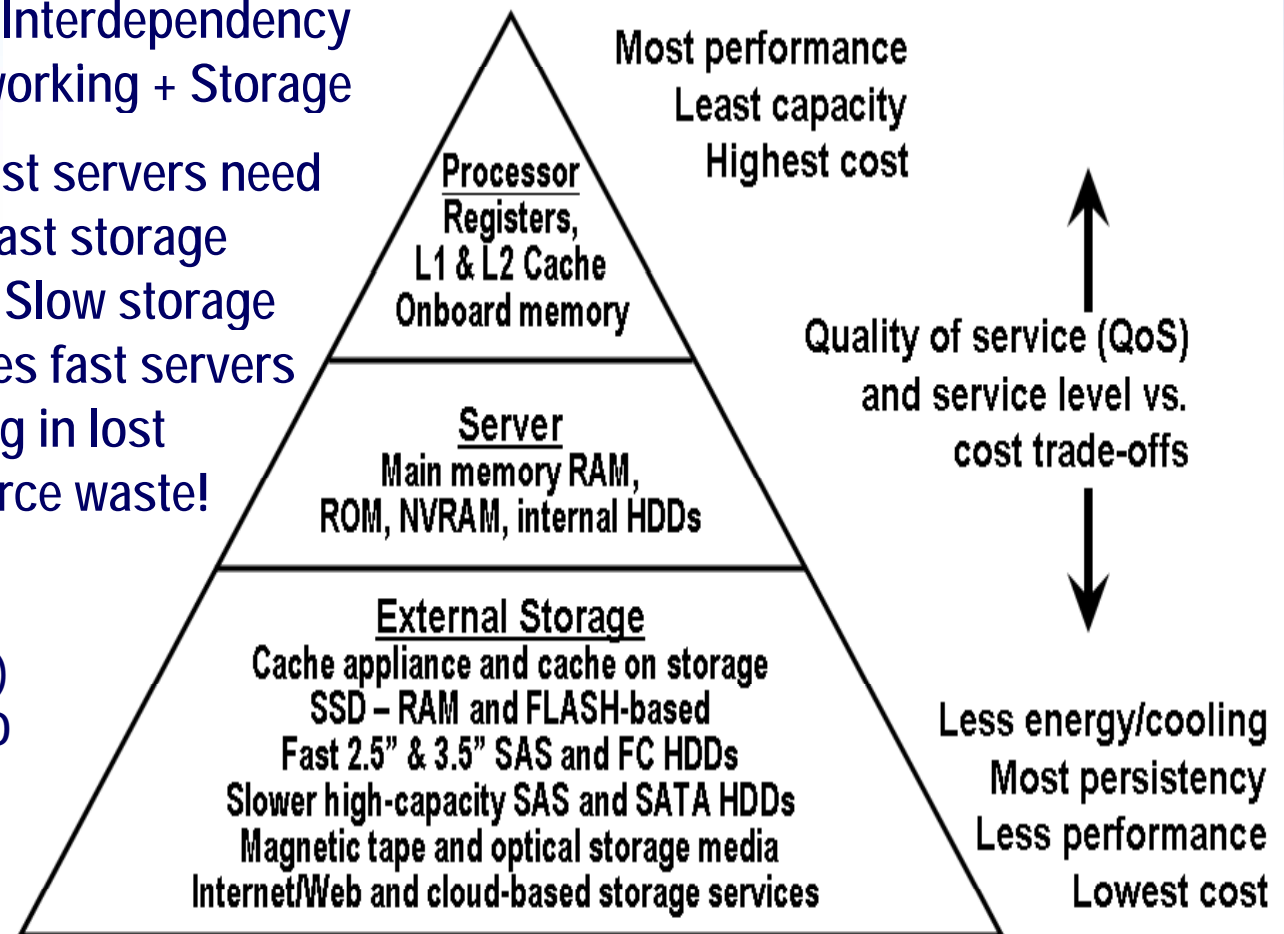
Industry Trends: Technology Alignment

Tiered IT Resources: Fix Performance Bottlenecks

Convergence and Interdependency
Servers + I/O Networking + Storage

The need for speed. Fast servers need fast I/O networks and fast storage mediums and devices. Slow storage and I/O networks causes fast servers to have to wait resulting in lost productivity and resource waste!

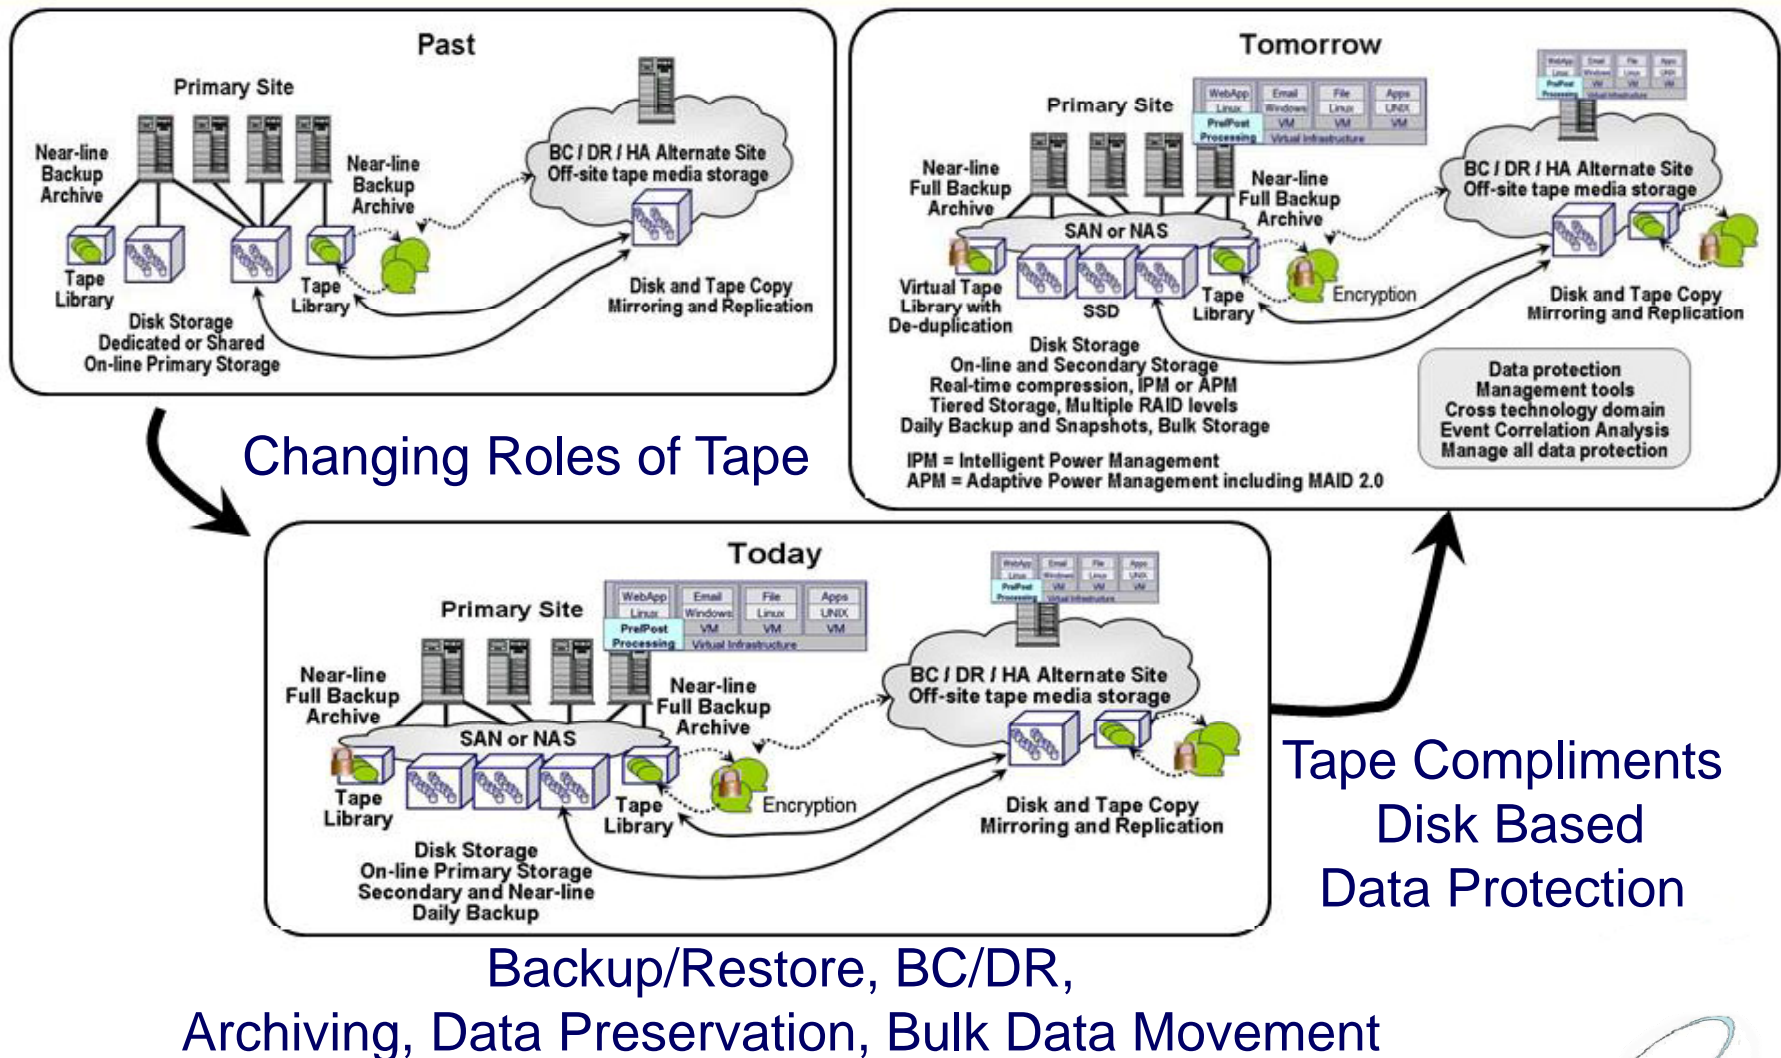
T0 = SSD (FLASH & RAM)
T1 = Fast 15.5K RPM HDD
T2 = High Capacity HDD
T3 = Tape, Optical, other



Source: "The Green and Virtual Data Center" (Auerbach)


Industry Trends: Technology Alignment

Tiered IT Resources: Tape, Still Relevant, Changing Role



Industry Trends: Cloud Computing

Cloudy weather & confusion vs. clear skies & opportunity



"Say . . . What's a mountain goat doing way up here in a cloud bank?"

Many types of clouds

- Scientific vs. commercial
- Services vs. solutions
- Architecture vs. products
- Public vs. private
- Protocols and personalities

Confusion and skepticism

- Lingering reminder of SSPs
- Replacement vs. complimentary
- BC/DR and availability

Opportunities & call to action

- Clarify - Avoid flying blind
- Position as a IT tiered resource
- Complimentary vs. competitive
- Protect & preserve, part of BC/DR


Industry Trends: Converged Networking

I/O Virtualization (IOV) and Converged I/O Networking

I/O, I/O, Its Off To Virtual Work We Go...


Traditional Approaches

Separate networks & interconnects
(Fibre Channel, GbE, IBA, Etc.)



Evolving Approaches

Unified & converged interconnects
(Virtualized FC, GbE, FCoE, Etc.
Physical Data Center Ethernet or IBA)



Need for speed: Consolidation, Reduce Response Time (Productivity), Support Growth

See “FCoE Overview” and “I/O, I/O, its Off To Virtual Work we Go” www.storageio.com

Green Computing: Shifting Focus

From green-wash to action, from energy avoidance to energy efficiency

“A little less conversation, a little more action please...” - Elvis



RIP: Green-wash and Green hype!

- **Green-hype**
On endangered species list
Less on perception
More on substance
- **Closing the green gap**
Economics vs. CO₂ only focus
Sustain growth and service
From energy avoidance to energy efficiency

Energy Avoidance

Power Down, Over Consolidate
Decrease Amount of Useful Work
Decrease Energy Used

Some Energy Efficiency


Faster Components, Same Power
Increase Amount of Useful Work
Same Amount Energy Used

Some Energy Efficiency

Lower Power Draw Components
Same Useful Work Done
Decrease Energy Used

More Energy Efficiency

Faster Components/Less Power
Increase Amount Useful Work
Decrease Energy Used




Source: "The Green and Virtual Data Center" (Auerbach)




Green Computing: Shifting Focus - Continued

Boosting IT efficiency & productivity, addressing PCFE and other IT issues

Wheel of opportunity: PCFE optimization and IT productivity



Policies, metrics,
monitoring and
best practices
efficiency rebates



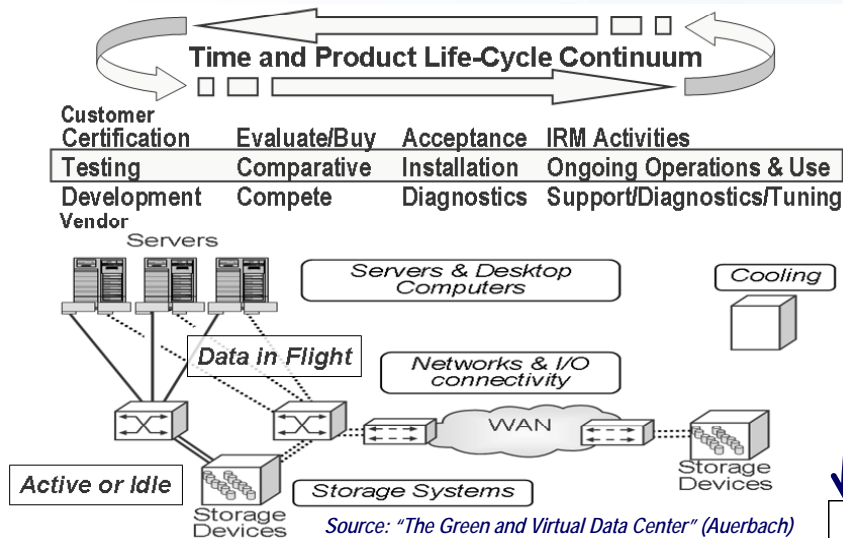
Also see SNIA and
other venues for
various tips

PCFE: Power, Cooling, Footprint, EH&S

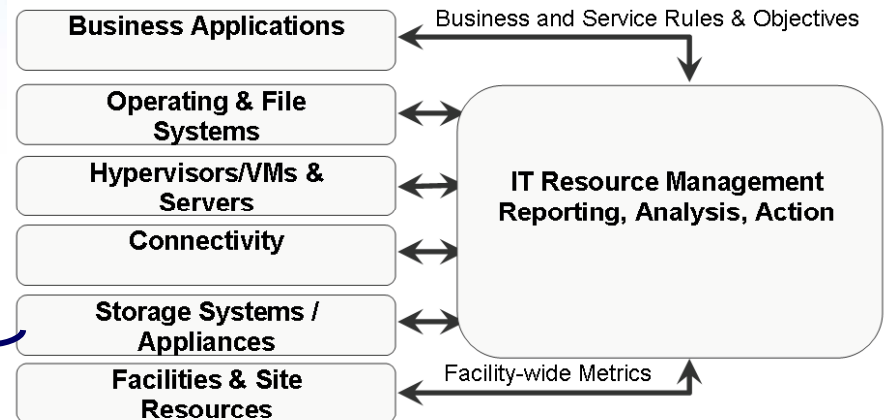
Metrics and Measurements

Industry is busy as are EPA and others!

Different usage scenarios



Different areas of focus



Source: "The Green and Virtual Data Center" (Auerbach)

SNIA metrics and taxonomy

	Online Storage	Near Online Storage	Removable Media Libraries	Virtual Media Libraries	Infrastructure Appliances	Infrastructure Interconnects
Storage Taxonomy Summary	Prime storage, able to serve random as well as sequential workloads with minimal delay	Intended as second tier storage behind Online Storage. Able to service random and sequential workloads, but perhaps with noticeable delay in time to 10" data-cs="6" data-kind="parent">				
	Max Storage Devices	Max Storage Devices	Max Tape Drives	Max Storage Devices Supported	Max Port Count	
Group 1) SoHo & Consumer	Up to 4 Devices		Stand Alone Drive (No Backup)			
Group 2) Entry, DAS, or JBOD	More than 4 Devices	Up to 4 Devices	Up to 4 Drives			Up to 32
Group 3) Entry / Midrange	More than 20 Devices	More than 4 Devices	More than 4 Drives	Up to 100 Devices	Support for up to 20 Devices	Up to 128
Group 4) Midrange / Enterprise	More than 100 Devices	More than 100 Devices	More than 24 Drives	More than 100 Devices	Support for more than 20 Devices	More than 128
Group 5) Enterprise / Mainframe	More than 1000 Devices		More than 11 Drives	More than 100 Devices	Support for more than 100 Devices	

© SNIA 2009

80 Plus, EPA, Greengrid, JDEC, SNIA, SPC and others

What's Next?

Some predictions, opinions and perspectives

- ✓ Archiving escapes from the realm of compliance
- ✓ FCoE prime time “main event” in 2010/2011
- ✓ NAS gains DB popularity when using DIO/CIO
- ✓ OSD and Holographic appear at “some day isle”
- ✓ Object storage systems beginning to get mindshare
- ✓ SAS gets respect for HDD to storage, storage to server
- ✓ RAID remains relevant (Anyone remember DTDS+ ?)
- ✓ Coarser and integrated CDP gaining adoption
- ✓ FLASH and SSD make HDDs relevant, HDDs make tape relevant
- ✓ Small improvements on large basis have big benefits:
 - 2:1 compression on 100TB = 50TB vs. 20:1 dedupe on 5TB = 1TB
 - 4 watts savings per HDD x 1,000 HDDs = 4kW
- ✓ Networks are faster, however more data to move in less time
- ✓ Virtualization is driving need for cross technology domain IRM
- ✓ DLP and security integration across IRM domains
- ✓ More compliance / regulations including ETS and “Green” schemes



General Comments

Basic Premises – Gain management insight and control

- IT data centers are information factories
 - Efficient equipment, management tools and resources needed
- Many approaches depending upon your customers issues
 - Do more with less: Shift from energy avoidance to efficiency
 - Small improvements on large scale have big benefits!
 - More IOPS or Transactions or Files processed per watt of energy
 - More usable capacity in given footprint per watt of energy
- Avoid simply moving IT problems around
 - Solve problems and issues to enable IT and business growth
 - Instead of race to replace tape, revamp data protection architecture
- Balance between futures and what works today
 - Leverage virtualization to bridge from the past to the future
 - Virtualization and cloud computing require real resources & people

Next to Closing Comments



So what's my point in all of this anyway?

- In general
 - More of the same: "Meet the new boss, same as the old boss" - The who
- 2009 will be a challenging year, there are opportunities
 - Execution and delivery on prior hype and announcements
 - Prepare for future growth and scaling requirements
 - Flexibility, adaptability, leveraging opportunities and challenges
 - Preparation for the next technology and economic phase
 - Optimization to support and sustain growth now and in the future
- What will be important this year
 - Flexibility, adaptability, affordability, value vs. low cost
 - Articulate vision and strategy for near term and long term
 - Educate, generate awareness, add value, be of assistance

Closing Comments

Look at 2009 as a challenging opportunity

- Where you can learn more: 
 - www.storageio.com (Book info, white papers, articles, tips, videos)
 - www.storageioblog.com & www.twitter.com/storageio
 - www.thegreenandvirtualdatacenter.com



**Order your copy of “The Green and Virtual Data Center” (Auerbach)
At Amazon and other global venues**

Thank You

And, Let's Stay In Touch!



Greg@storageio.com
Twitter.com/storageio