HITACHI
Inspire the Next

Is DAS really cheaper than SAN?

Storage Economic Methods help to differentiate Price and Cost

June 2009

David Merrill
Chief Economist
Hitachi Data Systems

Hitachi Data Systems
WebTech Educational Seminar Series

HITACHI
Inspire the Next

SAN versus DAS – A TCO Perspective

Recent reports suggest that DAS may have a lower cost than SAN architectures. David Merrill, Hitachi Data Systems Chief Economist and Global Business Consultant will present several dimensions of total cost of ownership that goes beyond the cost of acquisition. If the full capability of virtualized Storage and SAN are not employed and used to your advantage, then in some cases DAS can be preferable. Economically superior storage and SAN solutions continue to evolve, make sure you are realizing the full benefit of SAN (and NAS) storage architectures. At this session, you'll learn:

- Differentiate between cost of acquisition and cost of ownership when evaluating DAS and SAN Architectures
- Identify and measure the crossover point (economically) of SAN and DAS systems with regard to operations, backup, disaster recovery protection
- Compare basic SAN costs and benefits with virtualized and thin provisioned storage within the SAN

Agenda

HITACHI
Inspire the Next

- Intro on Storage Economics
  - 33 types of money
  - Price ≠Cost
- Separate Connection Protocol and Costs from Storage
  - Connection and storage are both intertwined in the architecture
  - Asking the right questions on DAS and SAN costs
- A short technical and econ history on SAN
- How to economically exploit SAN architectures
  - Virtualization, thin provisioning, dynamic tiers
- Separate the discussion on TCA and TCO
  - Econometrics
- Conclusions

Introduction to Storage Economics

HITACHI
Inspire the Next

- Price does not equal ownership costs – Purchase is 20-25% of TCO
- 33 types of cost
  - Total lifecycle costs
  - Procurement, maintenance
  - Cost of waste, cost of copies and growth
  - Labor, administration, migration effort
  - Risk of loss, recovery, litigation, compliance
  - Power, space, cooling
  - Download paper @ http://www.hds.com/assets/pdf/33-types-of-costs-for-storage-tco.pdf
- ROI, ROA and TCO
  - Business justification
  - Comparative economics
- Economically superior storage architectures exist and continue to evolve
- Econometrics
  - Cannot improve what you cannot measure
- Read more at www.StorageEconomics.com

Reality Check and Provocative Questions

HITACHI
Inspire the Next

- It is cheaper to waste disk capacity than it is to manage it better
- Price of disk is approaching zero, what is happening to the TCO?
- Total cost of storage or total cost of data?
  - Container economics
  - Data vs. storage economics (TCDO)
    - Don't divide costs into usable or raw
    - Divide costs into the amount of original instance data
    - TCDO exposes other 'hidden costs'
      - Waste, power, copies, RAID overhead
    - TCDO can be 4 to 10x the TCO

- How do I justify higher costs storage architectures to management?
- At what point is SAN, pooled storage architectures cheaper to own?
- What type of money is relevant in this discussion

What are the real questions here?

HITACHI
Inspire the Next

- **Fact 1**: DAS disk architecture is cheaper to buy
  - SAN costs, overhead can be high with switches, HBA, directors, ISL ports
  - Controllers, RAID overhead, Cache, FC ports add to the $/TB rate
- **Fact 2**: Price does not equal cost
  - Acquisition cost is only 20% of the TCO
  - Hitachi Data Systems has documented/measured 33 different types of ownership costs
- Perception problems
  - Connection type (DAS and FC SAN) and the full impact on cost
  - Pooled disk versus dedicated disk – both tend to have poor utilization…
  - Local cheap disk versus enterprise-class storage systems
  - Lets don't confuse the connection protocol with the storage architecture and storage TCO
- Cloud Computing is causing us to revisit this ten year old question
  - DAS is connection method being promoted for some cloud architectures
    - Application clouds, software as a service cloud
  - Traditional DC clouds (virtualized) still relying on SAN and pooled storage

Looking beyond Cost of Acquisition

HITACHI
Inspire the Next

- Data is not at rest on the disk
  - Protected from unplanned loss (backup)
  - Protected and recovered in the case of catastrophe (DR)
  - Kept around, auditable, traceable for long periods of time (archive, indexed)
  - Copies for various business support (DR, regression testing, dev)
  - Stateful and stateless
  - Poor performance can cost the company money
  - Data availability needed for competitive advantage, business need
- Nature and age of data
  - Structured, unstructured
  - Copies of data
  - Depots of data
- All of the above functions add cost and complexity to a seemingly simple task of storing data
  - DAS and SAN both have overhead, some might not be obvious

Capacity Definitions – Use the Proper Denominator!

HITACHI
Inspire the Next

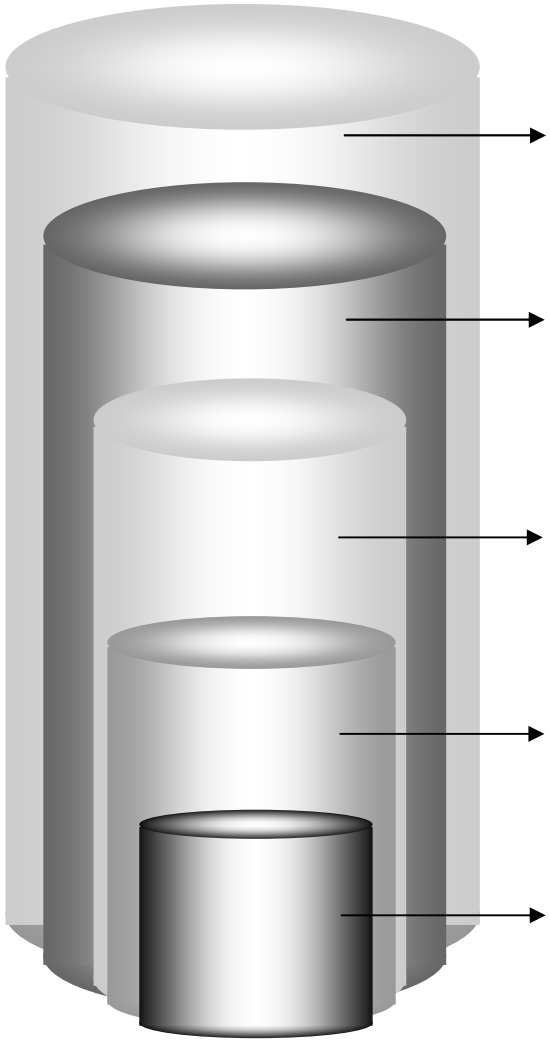

Raw Capacity
Disk volumes sitting in the box or frame

Usable Capacity
After RAID protection, formatted Ready for host allocation

Allocated Volumes
Disk volumes allocated to servers

Utilized Capacity
Disk volumes used by the operating system of the server, e.g. data bases, operating system

Application Content
Data content of the application

Capacity Definitions – Use the Proper Denominator!

HITACHI
Inspire the Next

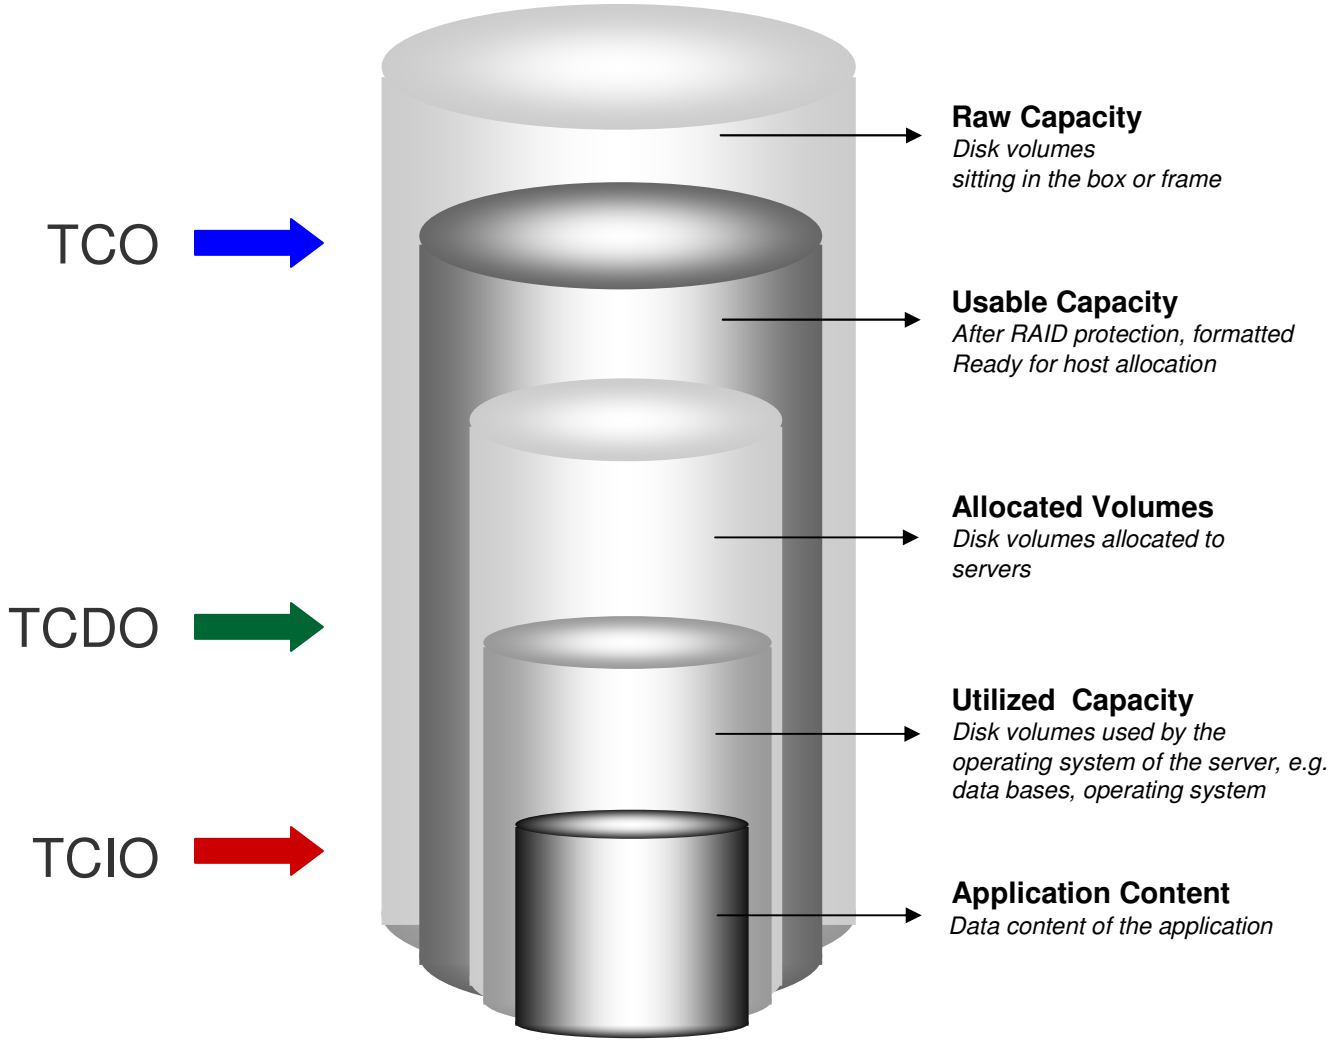

Hitachi Data Systems

HITACHI
Inspire the Next

History (and Economics) of Storage Architectures

Direct Attached Storage before SAN (circa 1999)

HITACHI
Inspire the Next

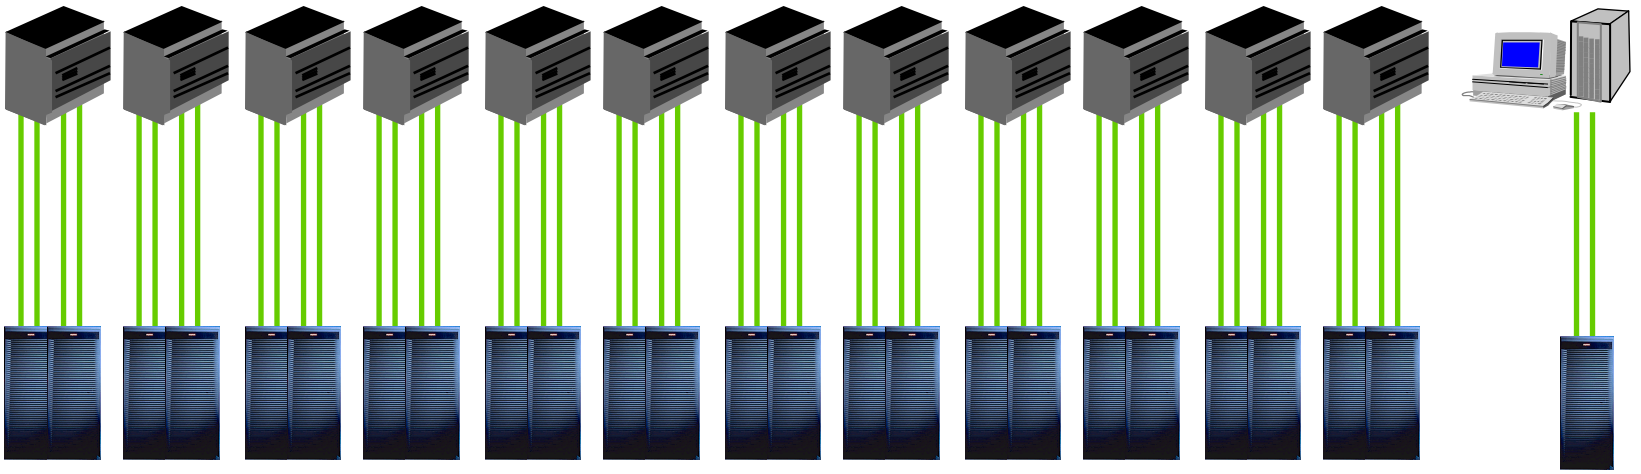

Storage and Server Islands

Maximum 1 to 2 TB per storage frame

20 to 30% Utilization of storage

Bulky SCSI cables at 32 Mb/sec speeds

Storage foot prints were more than half the data center floor space

Storage After SAN

HITACHI
Inspire the Next

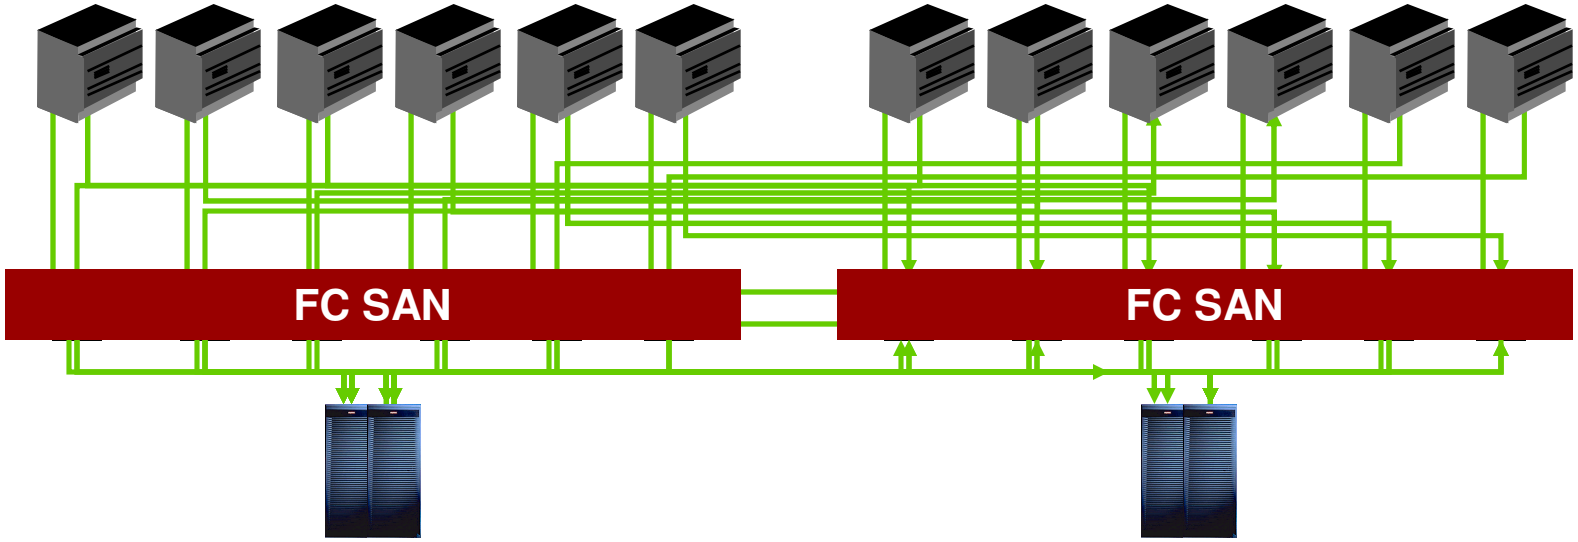

Servers are now networked to storage to access stranded storage

Capacity is consolidated to16 - 32TB storage frames

Storage foot prints were reduced by 6 to 8 times with denser storage frames

Thin Fibre Channel cables with 1000 Mb/sec performance and greater flexibility in configuration

Higher availability with switched paths

Hitachi Data Systems

SAN Implications for Storage

HITACHI
Inspire the Next

- Storage consolidation puts more demands on the Storage Subsystem
  - Internal bandwidths must be increased
  - Higher capacity, scalable subsystems are required
- Sharing a common network requires zoning of the Servers to storage
  - Prevent servers from accessing other servers storage on the network
- Consolidation puts more users at risk
  - Demands higher availability subsystem

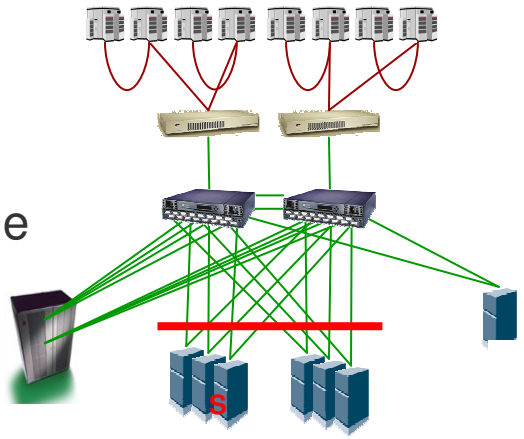

- SANs require cooperation and collaboration between vendors
  - SNIA was established by storage vendors to sort out integration of interfaces and set up cooperative support agreements
- SAN's create another level of management complexity
  - SAN's become the third leading cause of application failure

What Happened to the Promise of SAN

HITACHI
Inspire the Next

- SAN promised consolidation through the elimination of islands of storage
- SAN would increase storage utilization and provide centralized management
- Ten Years later the utilization of storage is still about 20% to 30%
- **What happened?**
  - SAN Networked servers to storage but left the islands of storage
  - Storage systems are not networked to work together
  - Storage systems are still provisioned separately
  - There is no networking, no data mobility between storage systems

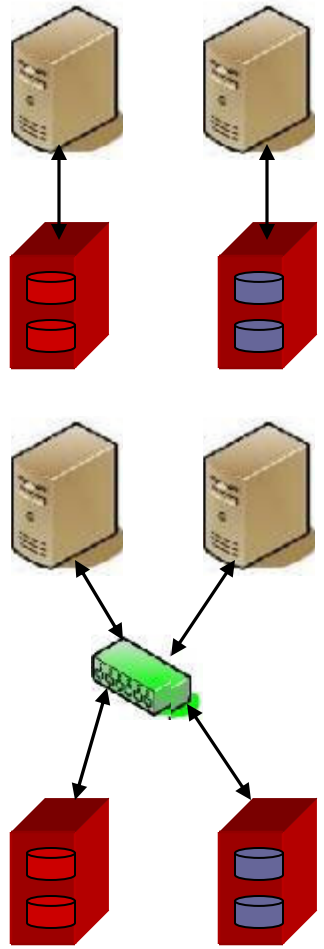

SAN Summary

HITACHI
Inspire the Next

- SAN connectivity is required for consolidation and data mobility
  - Connect multiple servers to one physical storage port that supports multiple virtual storage ports
  - Connect multiple role based servers to a common pool of virtual storage capacity
  - Connect different content services platform to a common archive platform
  - Connect Virtual Server Nodes for Application mobility (VMotion) and Site Recovery
  - SAN is required with storage virtualization to reduce device migration costs

- SAN requires storage virtualization to increased storage utilization
  - Requires storage virtualization to provide mobility between storage arrays
  - Requires Capacity virtualization to provide a seamless pool of storage capacity

Economically Superior Storage Architectures

HITACHI
Inspire the Next

- In order to drive down unit costs and meet increasing growth requirements, new storage architectures have to be employed
  - Current island SAN and DAS solutions tend to be unsustainable
- An Economic Triumvirate
  - Storage Virtualization – in the controller
    - Reclaimed space, better aggregate utilization, better mgmt
    - Significant time and cost savings with data/system-based migration
  - Dynamic Tiering
    - Data mobility within the tiers, policy-based
    - Move and tier the copies!
  - Dynamic Provisioning (thin provisioning)
    - Space reclamation, wide striping and very fast provisioning
- Individually these 3 provide economic benefit, but there is a compounded impact when bundled
  - Better utilization, reduced power and cooling per TB, big labor savings
  - Fewer arrays serving up the storage capacity, lower cost of growth

Services versus the Standalone Approach

HITACHI
Inspire the Next

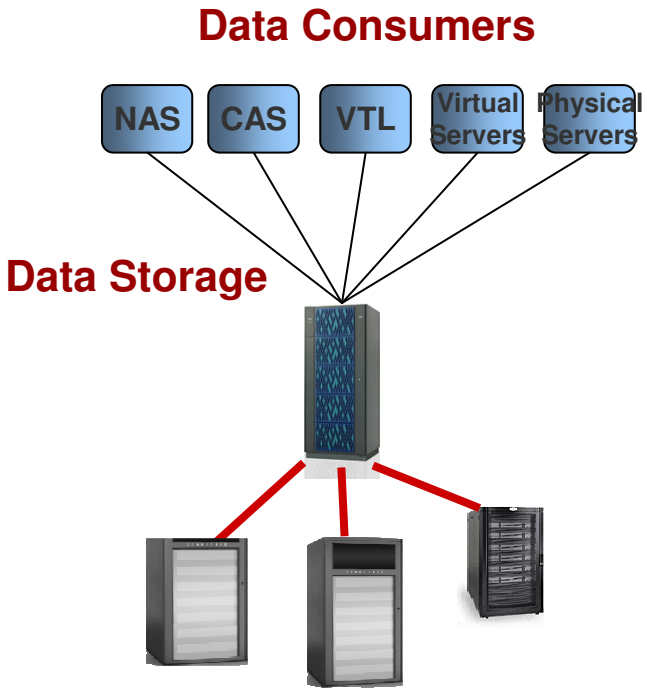

"I need 10TB of disk space"

- Where is the best place to put it? (enabled by Tiered Storage)
- I have 30TB free across all platforms
- If the platforms are virtualized in a storage pool, no need to buy

Data Consumers and Data Storage

"I need 10TB of disk space"

- What type storage do you need it for?
  - Virtual Tape?
  - Content archival?
  - Normal disk?
- I have 30TB free –, 9TB on NAS, 8TB on CAS, 7TB on VTL, 6TB on DAS
- Buy additional 10 TB even though 30TB is available

Planning for Different Data Types

HITACHI
Inspire the Next

Consumption of Enterprise Disk Capacity by Type

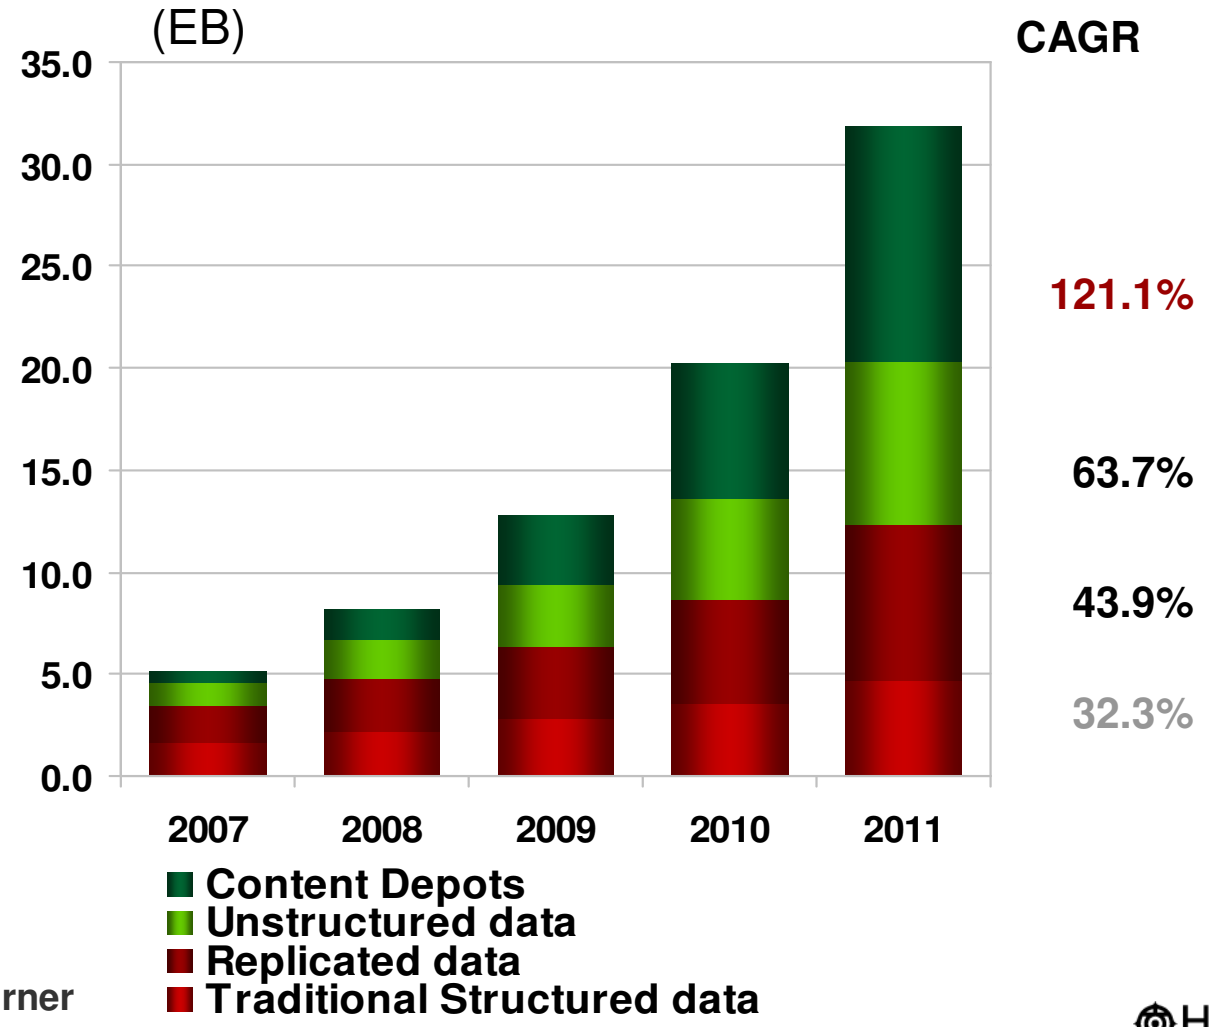

IDC 2008
Vernon Turner

Hitachi Data Systems

Hitachi Can Consolidate Different Storage Requirements with a Common Services Platform

HITACHI
Inspire the Next

RISS

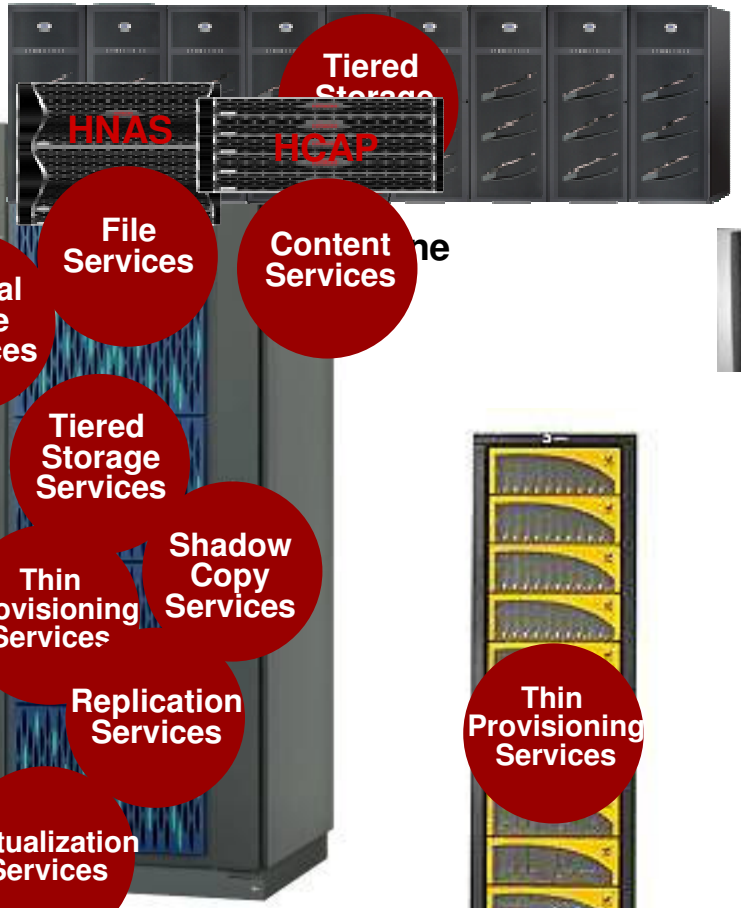

NAS

VTL

InServ S800

AXIOM

Hitachi Data Systems

HITACHI
Inspire the Next

Two Assessment Methods

Total Cost of Acquisition

Total Cost of Ownership

Premise on DAS and TCA

HITACHI
Inspire the Next

- Some basic assumptions around total cost of acquisition (TCA)
  - Low cost SAS disks local to the CPU do present an attractive price point
  - Capacity, performance, RAID protection may have to be factored
  - RAID overhead
  - Performance, short-stroking, wide stripes for better access
- Direct attached to the CPU, the host has to do all the management
  - Backup, Disaster Recovery, replication, data mining
  - De-duplication, encryption, data mining, etc.
- Where are storage management functions deployed?
  - Central console or array
  - Within the host
  - Bundled and controlled by the application
    - Standards, interoperability
- At what capacity does the TCA cross-over between DAS and SAN?
  - Buying more spindles to handle workload
  - Increasing business risk with increased capacity and value of data
  - File servers, cluster nodes to be added for capacity only

TCA Analytics for a Cloud Architecture

HITACHI
Inspire the Next

	Innput Parameters	Unit		Case 1	Case 2	Case 3	Case 4	Case 5
General	Number of Nodes, CPU	Per cloud		12	60	100	200	400
	Workload Type	for cloud						
	Averages Disks per Node	Per node		3	3	3	3	3
DAS Parameters	RAW Capacity	per disk		1024	1024	1024	1024	1024
	Unit Cost	Per disk		$68	$68	$68	$68	$68
	RAID Overhead	array group		0%	0%	0%	0%	0%
	Average DAS Utilization	per disk		20%	20%	20%	20%	20%
	Blade/Server HW Cost	per node		$2,800	$2,800	$2,800	$2,800	$2,800
	Copy factor	per TB		3	3	3	3	3
	Capacity Overhead	per cloud		5%	5%	5%	5%	5%
	failure ratio	per cloud		20%	20%	20%	20%	20%
	Performance overhead	per cloud		30%	30%	30%	30%	30%

DAS Calculations			Case 1	Case 2	Case 3	Case 4	Case 5
# of total drives			36	180	300	600	1,200
TCA of drives			$2,448	$12,240	$20,400	$40,800	$81,600
failure (overhead)	extra nodes		2.40	12.00	20.00	40.00	80.00
performance overhead	extra nodes		3.60	18.00	30.00	60.00	120.00
capacity growth overhead	extra nodes		0.60	3.00	5.00	10.00	20.00
Other HW overhead			$18,480	$92,400	$154,000	$308,000	$616,000
Other SW overhead							
Total RAW Capacity	TB		36.9	184.3	307.2	614.4	1,228.8
Total Usable Capacity	TB		36.9	184.3	307.2	614.4	1,228.8
Total Used Capacity			2.46	12.29	20.48	40.96	81.92
TCA	All cloud Disk		$20,935	$104,673	$174,455	$348,910	$697,820
TCA / Usable TB	Usable		$568	$568	$568	$568	$568
TCA / Data Written TB	Used TB		$8,518	$8,518	$8,518	$8,518	$8,518

TCA Comparisons for DAS and SAN in a Cloud

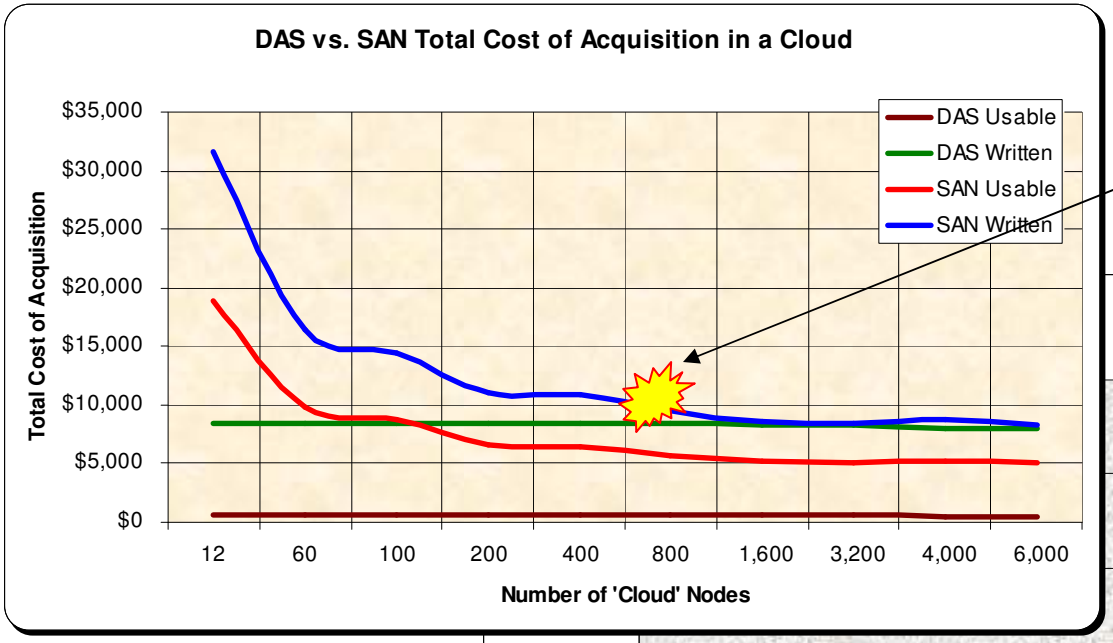

TCA Parity at around 1,000 Nodes

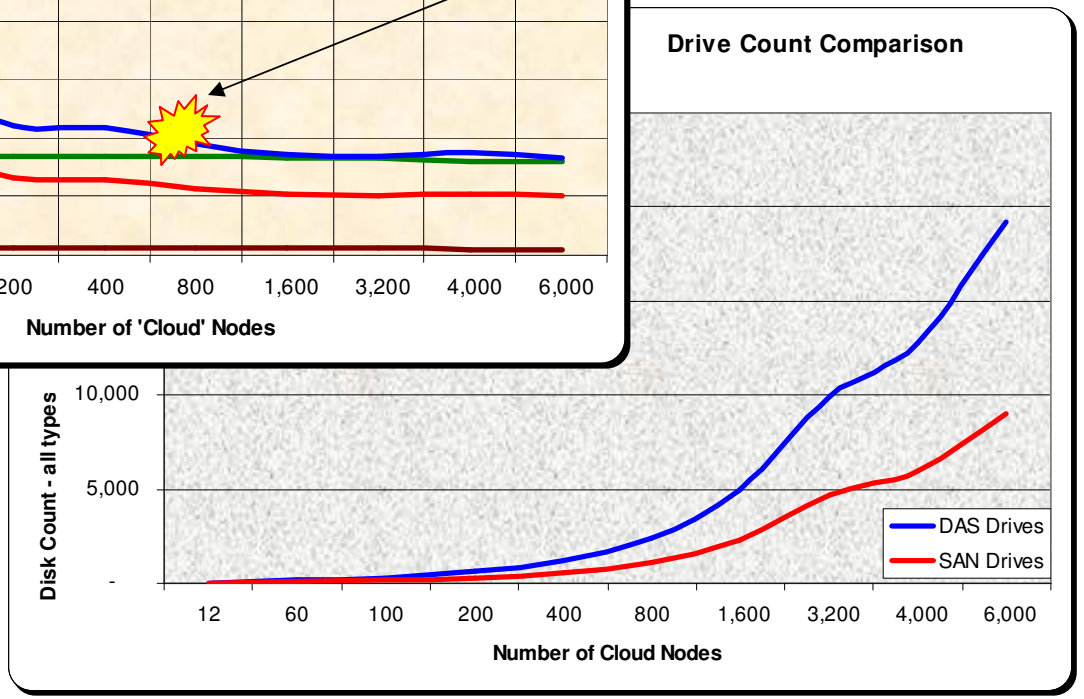

Total Cost of Ownership (TCO/TB/Year)

HITACHI
Inspire the Next

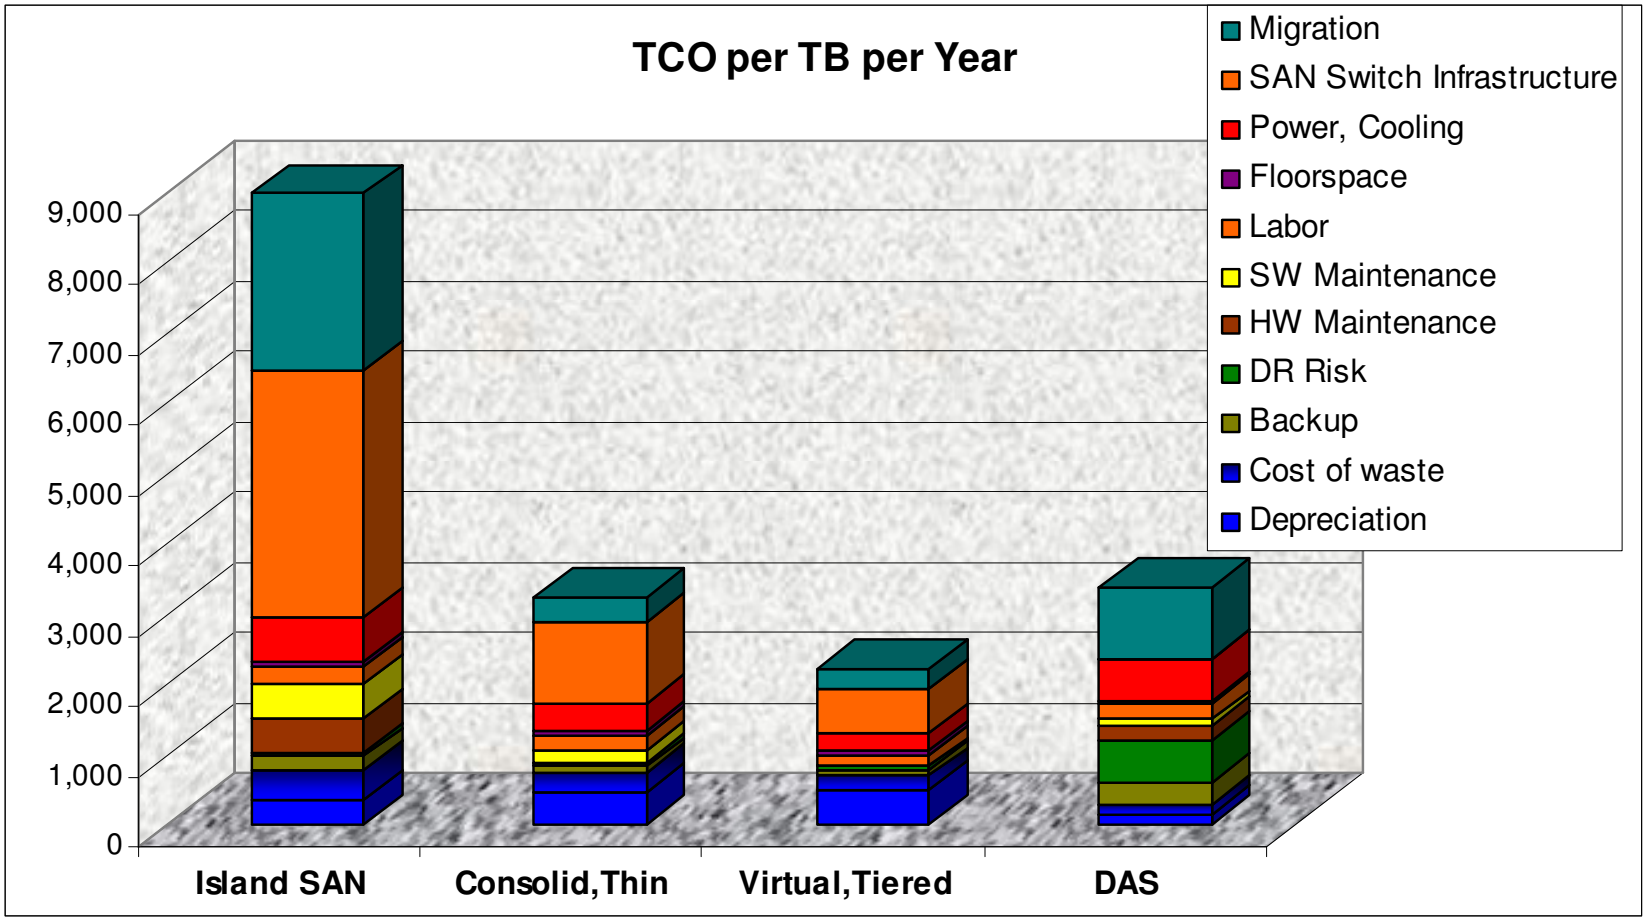

One site analysis, 320 TB usable capacity

Hitachi Data Systems

Cost Impact Areas that Differ with Architectures

HITACHI
Inspire the Next

- **SAN Connection Tariff**
  - Without virtual ports, large cache and virtualization, stand-alone SAN architectures can present a relatively high per port and per GB Cost
- ***Migration Cost* will vary greatly between architectures**
  - Sometimes a hidden cost since it occurs in the distant future
- **Labor and Management**
  - Server admin and storage experts, management $ woven into all storage solutions
- ***Cost of waste*, cost of growth and ability to scale**
  - DAS and SAN Islands tend to have an unsustainable physical growth pattern that will impact DC environmentals, Capital costs etc.
  - Presenting virtual, thin and tiered volumes to the host (through FC, NAS, FCoE) can reduce the physical growth or appetite, and do more with less
- **Backup, copies, DR *protection* present added services that have different cost factors for different storage implementations/architectures**
  - CPU and app impact on doing other boring things related to data and storage
  - Don't burden apps and CPU to some storage functions

Overhead and Overhead

HITACHI
Inspire the Next

- DAS does have real and economically impacting overhead costs
  - Backup servers
  - Disaster protection infrastructure
  - CPU workload that could be moved to a specialized processor or controller
  - Many extra spindles for performance, growth, utilization, copies etc.
- SANs have an overhead
  - HBA, Switches, higher-end storage systems and various software
  - SAN based storage can be prone to poor utilization and higher-than-needed costs of acquisition and ownership
- SAN with advanced storage can present new TCO and price points that are easy to overlook
  - Thin volumes
  - Capacity on demand provisioning
  - Advanced protection, encryption, data management, retrieval

Summary

HITACHI
Inspire the Next

- SAN connectivity is required for consolidation and data mobility
- SAN requires storage virtualization to increased storage utilization
  - Requires storage virtualization to provide mobility between storage systems
  - Requires Capacity virtualization to provide a seamless pool of storage capacity
- SAN is required to integrate server and storage virtualization
- SAN Connectivity is a key factor for increasing ROA through Virtualization
- SAN and DAS are connection protocols, but also implicate the underlying storage architecture and overall storage efficiencies
- Differences of TCA and TCO
  - DAS TCA does see a cross over point when comparing to SAN and pooled storage (virtualized, thin)
  - SAN and Pooled storage tend to have better TCO as more cost categories come into play

HITACHI
Inspire the Next

Questions and Discussion

Upcoming WebTech Seminars

HITACHI
Inspire the Next

- **SAN Series Webcasts**
  - SAN Consolidation – The Next Step, July 15, 2009, 9am PT
- **Upcoming Webcasts**
  - Thin Provisioning – Lessons Learned, June 17, 9am PT
  - Storage Reclamation: A Case Study*, "How to Live off your Body Fat in a Down Economy",* July 22, 2009, 9am PT
  - What's Next for Sustainable IT?, July 29, 2009, 9am PT
- **Please check www.hds.com/webtech for:**
  - Link to the recording, the presentation, and Q&As (available next week)
  - Schedule and registration for upcoming WebTech sessions

HITACHI
Inspire the Next

Thank You