

19. April 2010



# Storage Innovations

## Smarter Storage Systems for a Smarter Planet

**Thomas Frey**

[thfrey@de.ibm.com](mailto:thfrey@de.ibm.com)

**IBM Deutschland Enterprise Storage Sales Leader**



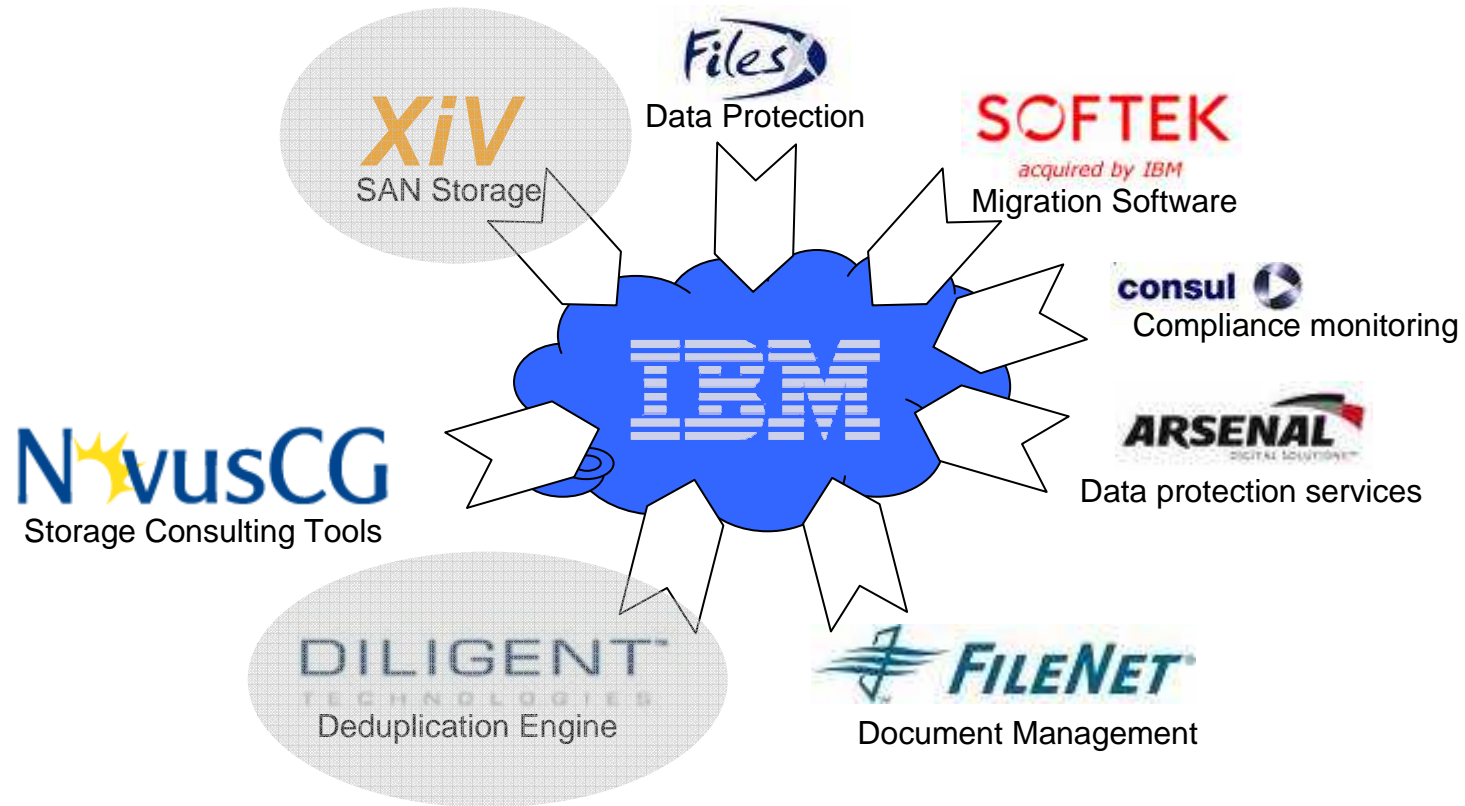
## **IBM Storage Leadership**

- #1** Virtualization
- #1** Tape
- #2** External Disk
- #1** Storage Mgmt
- #1** Storage Services
- #1** Storage Solutions
- #1** Total Storage

Source: IDC, Gartner, IBM MI




## Important STG Acquisitions in Recent Years



# The explosion of information meets budget realities

## *Top issues and initiatives for storage managers*



Source: TheInfoPro Storage Study (12/14/09), n=186.

Source: The InfoPro Storage Study (12/14/09): F1000 Sample. Wave 13, n=183.

\*Note that due to multiple responses per interview, total exceeds 100%. This is a partial list of responses received, showing only the top issues and initiatives.

## Announcing a new smarter systems portfolio

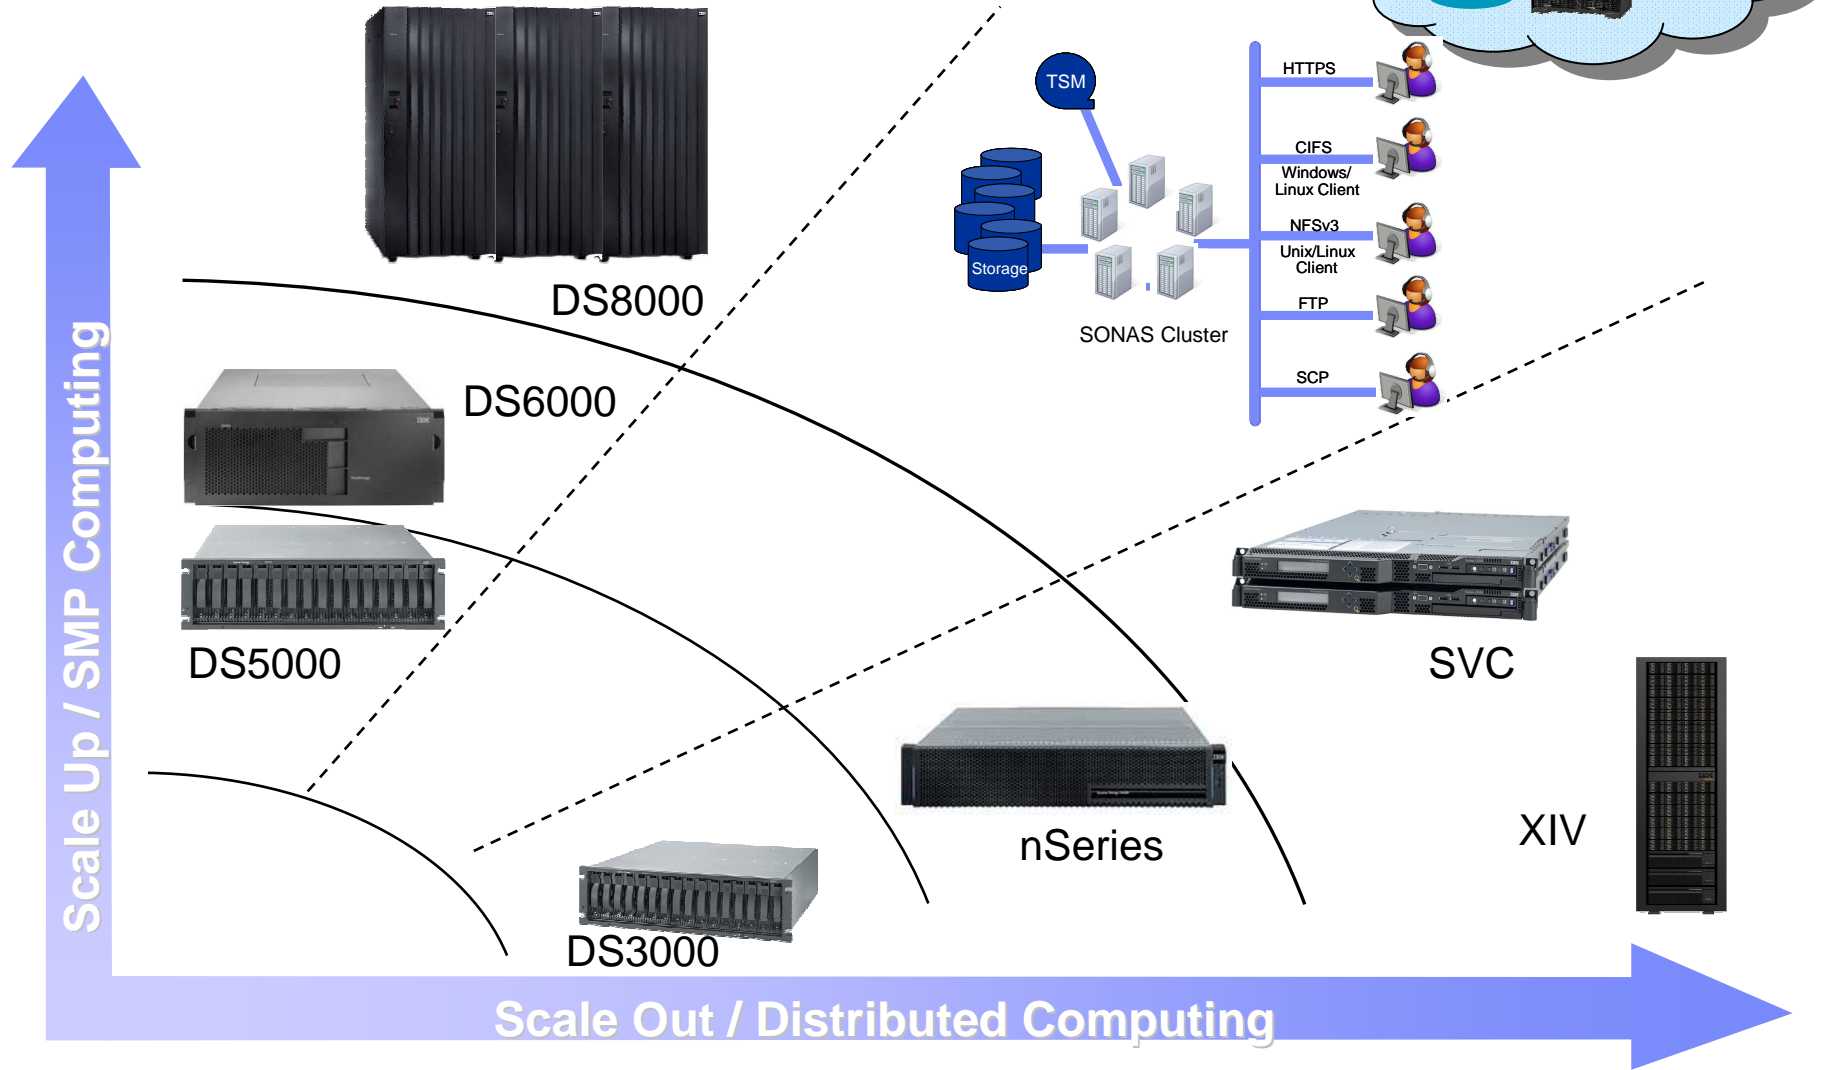
### Smarter Systems for a Smarter Planet

February 9
POWER7 <i>(midrange)</i>
Storage leadership <i>(SONAS, data deduplication)</i>
March 2
System x eX5 <i>(introduce portfolio)</i>

April 13
<b>Smarter Systems for a Smarter Planet</b>
Increase emphasis on leadership Storage offerings. DS8700 Easy Tier LTO5 XIV Erweiterungen

July	October
Storage leadership	Storage leadership

# Storage architectures (Scale up and Scale out)



## Challenge: Reduce the cost and complexity of storing vast amounts of data

***Announcing: Storage offerings that enable smart movement and management of information and capacity growth without added complexity.***

### **DS8700 Easy Tier scalable storage for vast amounts of data with less management complexity.**

- Relocating 10% of data to SSDs can increase system throughput by 300%.

### **ProtecTIER many-to-one replication lessens duplication while storing backup data.**

- Reduce WAN bandwidth by up to 95% or more saving cost and time.

### **XIV 2TB high-density, scalable drives to manage vast amounts of data efficiently.**

- Double storage capacity in the same footprint plus optimized performance.

### **LTO 5 midrange industry standard tape drives and libraries to rapidly retrieve data.**

- 88% more capacity and 17% better performance than previous generation.

## IBM XIV – Next generation storage - Scale Out

***Next-generation storage with exceptionally low Total Cost of Ownership***

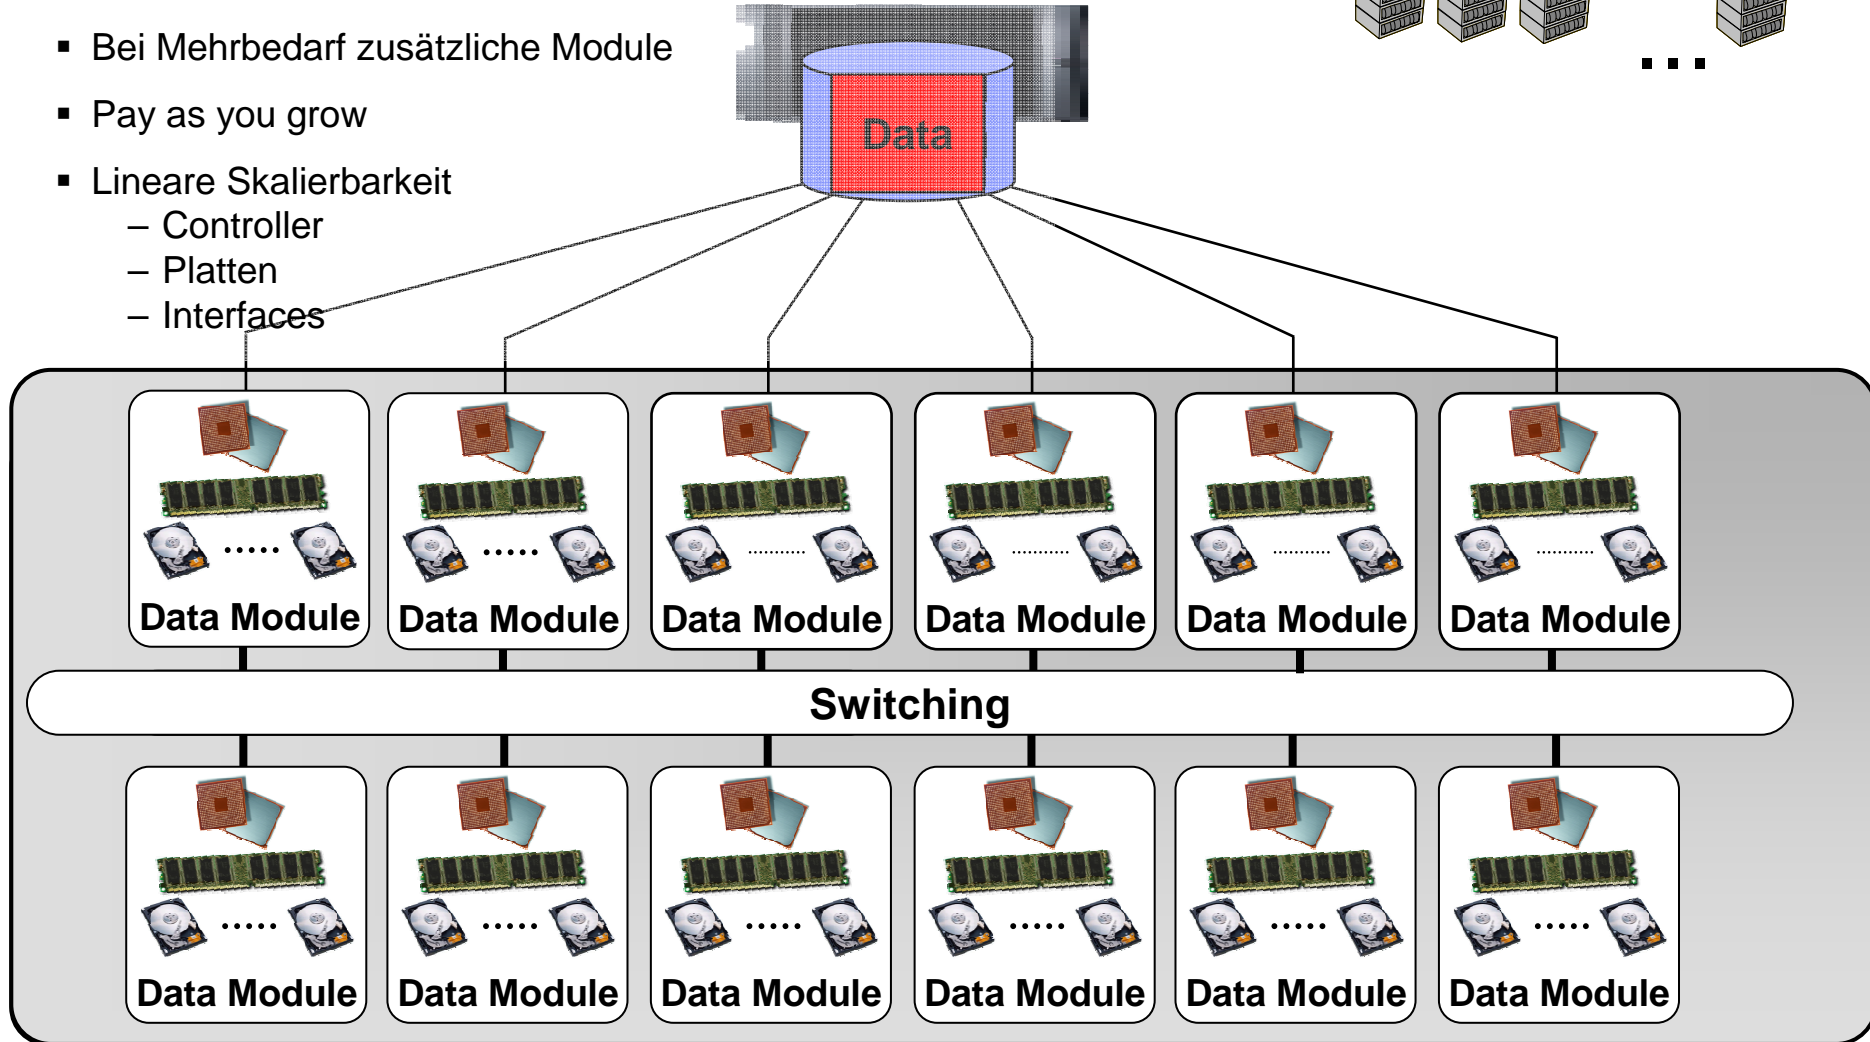

- **Capacity without complexity – Now double capacity to 161TB.**
  - Remarkably easy installation. Set it up and it starts migrating data.
  - Automatically tunes itself. Never a hotspot. Always fast.
  - Performance gets better as you fill it up.
- **Advanced grid architecture.**
  - XIV delivers superior performance with high capacity drives
- **Enterprise class storage that's ideal for white space.**
  - Almost 2,000 new installations since we acquired XIV.
  - In 2009 there were over 200 installations into accounts that had not done business with IBM before.





### Scale-Out

- Daten sind über mehrere Controller verteilt
- Bei Mehrbedarf zusätzliche Module
- Pay as you grow
- Lineare Skalierbarkeit
  - Controller
  - Platten
  - Interfaces



**Current 79TB Useable with 1TB**



**April 13, 2010**



**161TB Useable with 2TB drives**



**79TB Useable with 1TB drives**



IBM XIV Configurations											
# Modules	# Disks	Usable capacity* 1 TB/2 TB	Connectivity		Power usage 1 TB/2 TB			Internal (grid) switching (GBps)	# CPUs		Memory (GB)
			FC ports	iSCSI ports	kVA (peak)	kVA (idle)	kBTU		Single	Dual**	
6	72	27/55	8	0	3.4/3.1	2.9/2.8	11.0	30	6	9	48
9	108	43/87	16	4	5.0/4.5	4.2/4.0	16.4	48	9	15	72
10	120	50/102	16	4	5.5/4.9	4.7/4.4	17.8	52	10	16	80
11	132	54/111	16	4	6.1/5.4	5.2/4.7	19.3	56	11	17	88
12	144	61/125	20	6	6.7/5.8	5.7/5.1	20.7	60	12	18	96
13	156	66/134	20	6	7.2/6.2	6.2/5.5	22.2	64	13	19	104
14	168	73/149	24	6	7.8/6.6	6.7/5.9	23.6	68	14	20	112
15	180	79/161	24	6	8.4/7.1	7.2/6.2	25.1	72	15	21	120

# Linux on System z – XIV Support Statement

April 30, 2009

## IBM is announcing qualification and general availability of support for Linux on System z (SLES 10) with the IBM XIV Storage System.

- IBM eServer™ zSeries® 890, 990 (z890, z990), all IBM System z9® and all IBM System z10™ servers
- IBM XIV Storage System (2810-A14)
- Environment:
  - Native LPAR mode: Linux on System z SLES 10 SP2
  - Guest OS mode: Linux on System z SLES 10 SP2 z/VM® is supported as a Hypervisor only. VM System volumes must reside on non XIV storage. z/VM release 5.4 and 5.3 are supported.
- SLES 10 2.6.16.60-0.34-default (or higher) is required

### Linux on IBM System z – IBM XIV Storage System Support Statement

IBM now supports Linux® on IBM System z® (SLES 10 SP2) with the IBM XIV® Storage System!

Linux on System z combines the advantages of the IBM mainframes with the flexibility and open standards of the Linux operating systems. Linux can help simplify business integration through the use of open industry standards, and it can also support deployment of new solutions more quickly.

Now the benefits of Linux on System z can be combined with the phenomenal capabilities of XIV – Storage Reinvented to support today's fast growing, dynamic environments. The IBM XIV Storage System is a revolutionary open disk system that represents the next generation of high-end disk storage, offering self-tuning and self-healing for consistently high performance and reliability as well as management simplicity and low total costs.

---






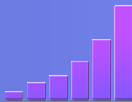
IBM is announcing qualification and general availability of support for Linux on System z (SLES 10) with the IBM XIV Storage System. This includes the integration into the IBM enterprise support mechanisms as well as all needed qualification items (hardware and software). Support qualification is as follows:

<b>System z Host Type:</b>	IBM eServer™ zSeries® 890, 990 (z890, z990), all IBM System z9® and all IBM System z10™ servers
<b>Storage hardware:</b>	IBM XIV Storage System (2810-A14)
<b>Environment:</b>	1. Native LPAR mode: Linux on System z SLES 10 SP2 2. Guest OS mode: Linux on System z SLES 10 SP2 z/VM® is supported as a Hypervisor only. VM System volumes must reside on non XIV storage. z/VM release 5.4 and 5.3 are supported.
<b>Linux code level:</b>	SLES 10 2.6.16.60-0.34-default (or higher) is required
<b>XIV code release:</b>	IBM XIV Storage System Software release 10.0.1.b (or higher) is required
<b>Known restrictions:</b>	255 WWPns in a zone with an XIV FC port 128 WWPns per single Host connected to an XIV FC port
<b>Date:</b>	April 30, 2009
<b>URL:</b>	<a href="http://www-03.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes">http://www-03.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes</a> Under Product Family, you would select IBM System Storage Enterprise Disk Under Product Model, you would select IBM XIV Storage System You would then see IBM System z and S/390 listed under Host Platform select that and you see SUSE SLES 10 under OS

IBM, IBM logo, IBM eServer, System z, System z9, System z10, XIV, zSeries and z/VM are trademarks of IBM Corporation in the United States, other countries or both. Linux is a registered trademark of Linux Torvalds in the United States, other countries, or both.

IBM Linux on System z / XIV – Support Statement

## Scalability: XIV Storage System

50 to 70% lower total cost of storage (no added charge for XIV software features – mirroring, snapshot, data migration, management)		<b>Capital Cost</b>
80/160 TB useable on one floor tile		<b>Space</b>
XIV uses 4 to 9 times less power for the same (or better for same capacity) performance and reliability levels		<b>Energy</b>
10-20% of traditional systems space is orphaned and will never be reclaimed. With XIV space is never lost		<b>Waste</b>
Using differential snapshots yields 15-30% saving in infrastructure cost		<b>Snapshots</b>
Thin provisioning yields 20-50% saving in infrastructure cost over a period of time		<b>Stretch Your TB</b>

# Cloud-onomics...

## CLOUD COMPUTING



...leverages virtualization, standardization and automation to free up operational budget for new investment



## New DS8700 Easy Tier Capability – Scale Up


### *Breakthrough Performance with Smart Management of Data*

- **Increase performance by up to 300%!**
  - Solid state drives (SSD) provides breakthrough performance.
  
- **Easy Tier maximizes performance gains while minimizing costs.**
  - Relocating just 5% of data to SSDs with Easy Tier can reduce response time 78%!
  
- **Easy Tier is easy.**
  - Smart data placement is completely automated.



## Scale-Up

- Daten sind in der Box gehalten
- Bei Mehrbedarf Aufrüsten der Box
- Hoher Einstiegsinvest
- Skalierbarkeits-Limit ist Box-Limit
  - In Anzahl Controller/Performance
  - In Anzahl Platten/ Kapazität
  - In Anzahl der Interfaces

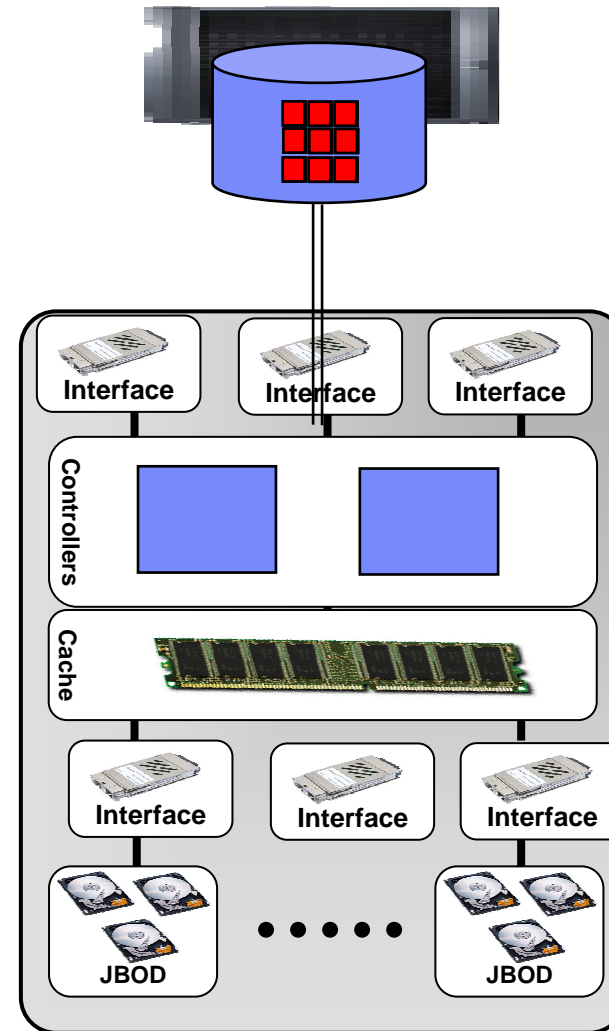


### Shared Prozessoren

- starke redundante SMP Prozessoren (2-way, 4-way, .....
- hoch effiziente Cache Algorithmen (4KB Segmente, Cache Verwaltung, Pre-Staging, Write efficient)

### Dedizierte Prozessoren

- Spezielle Host und Disk Interfaces

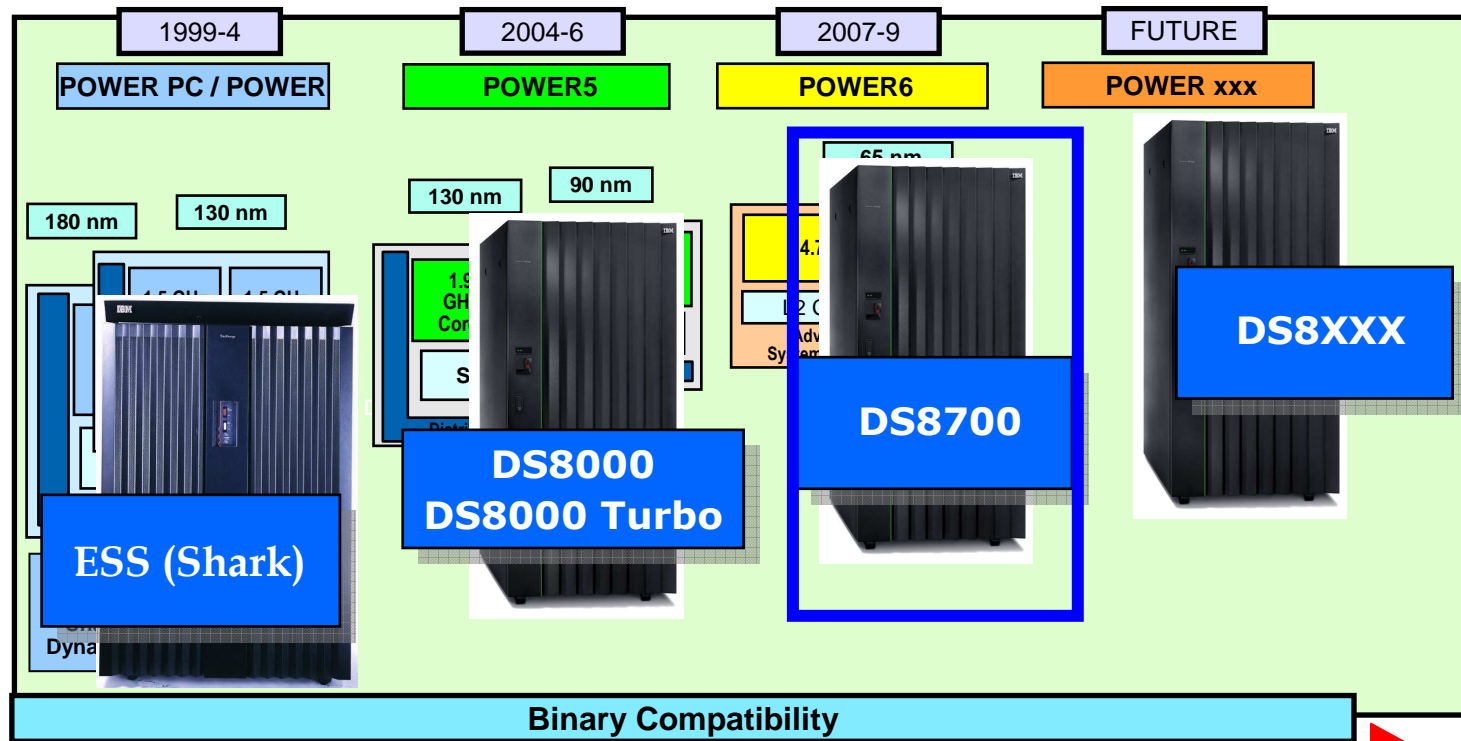


# IBM POWER and DS8700 Reliability and Resiliency

**Over 12,000 systems sold worldwide!!!**

The IBM POWER processor has been behind the success of IBM enterprise storage beginning with the Enterprise Storage Server (Shark) in 1999

***Benefit: a steady, 10-year lineage of RAS improvements!***



**DS8700 leverages the DS8000's highly reliable code base!**



# Introducing the New DS8700

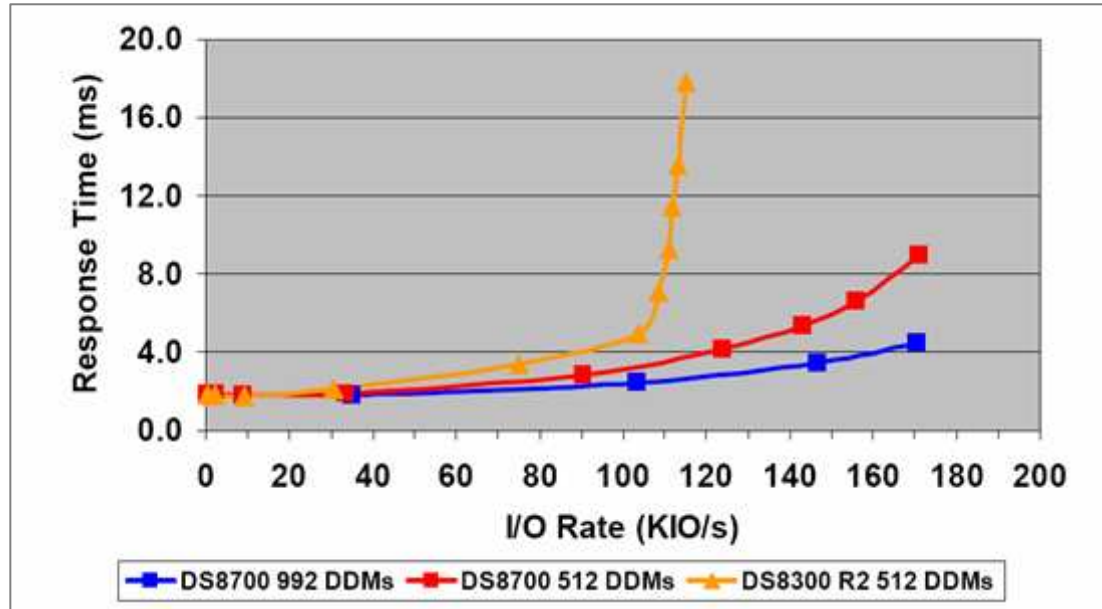
## *The Next Chapter in IBM's Flagship Disk Platform*

### What's New – Why Move to DS8700

- **Performance**
  - Up to over 150% performance boost with new IBM POWER6-based controllers
  - New, faster PCI Express (PCI-E) internal fabric enables much higher performance and scalability
  - Almost 70% faster ASIC on the device adapters
  - Increased FlashCopy performance in every metric
  - Increased SSD performance on sequential reads
- **Availability**
  - Single model, scalable via concurrent upgrade of all components
  - Shorter service windows with faster concurrent microcode updates
  - Better than 99.999% availability
- **Investment Protection and Scalability**
  - DS8700 architecture allows for future expansion of additional controllers and future 8Gb Host Adapters
  - Future features, such as SSD optimization and M/GM Multiple Session coming soon
- **Management**
  - Simplified management and application-aware FlashCopy
- **Security**
  - Full Disk Encryption enhancements address PCI-DSS compliance



# Performance improvements: OLTP FBA



FBA

(4KB 50R/50W/50HR)

## IBM System Storage DS8700 Release 5.1


- Release highlights
  - **IBM System Storage Easy Tier**
    - Sub-volume automated relocation
    - Full volume manual relocation
  - Storage Tier Advisor tool
  - New drive options double capacity
  - Thin Provisioning
  
- Announcement is April 13 (GA date is May 21)



**DS8700**

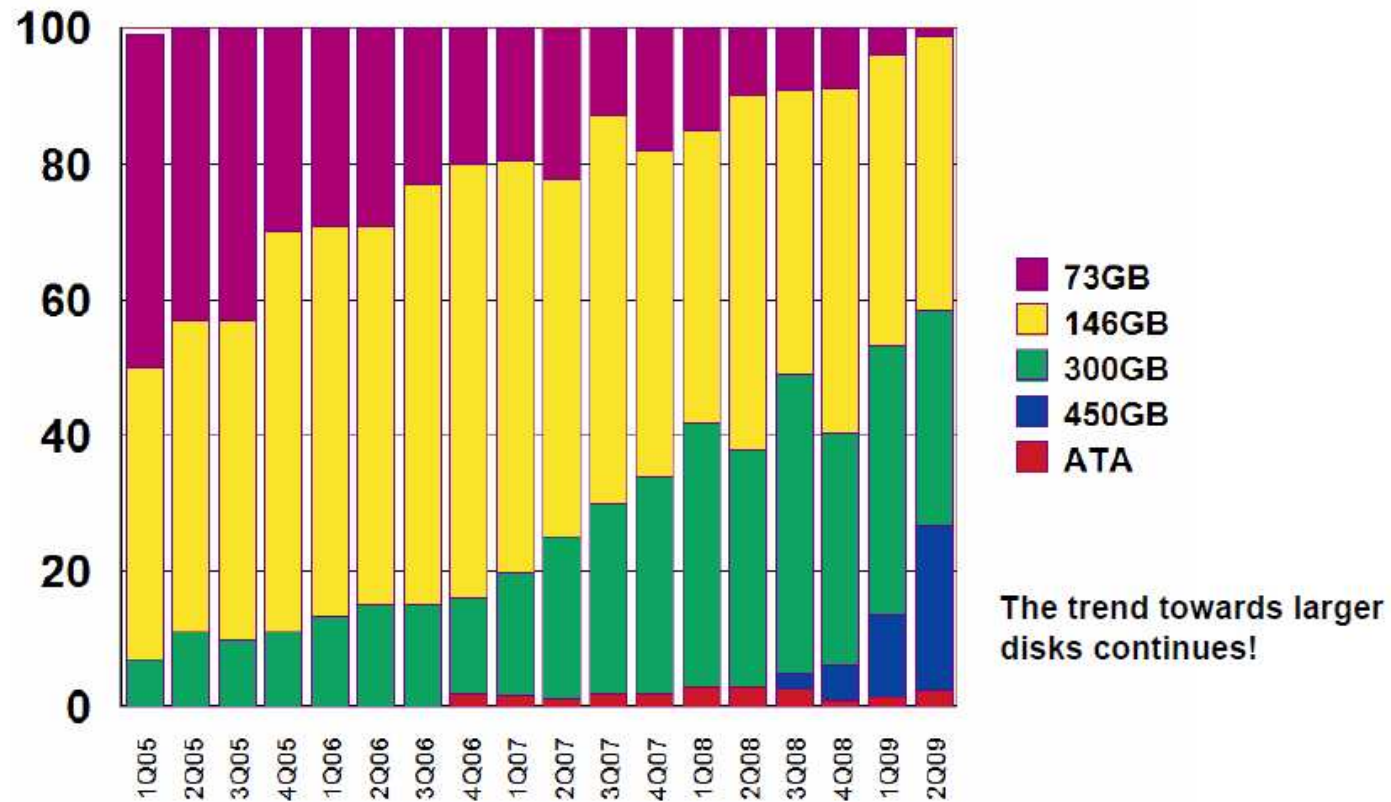
## Performance constrained by current drive limitations

- Processor capabilities are out-stripping disk drive and RAID controller performance (rotational speed and IOPS)
- As a result, servers and storage systems become more unbalanced between CPU/controller capability and storage performance
- Clients add more drive spindles to improve performance




Performance gains through HDDs has become ineffective and wasteful

## DS8000 Drive Sizes Shipped by Quarter (percent)



## Why is solid-state storage a business opportunity?

- Database Example – Use of Rotating Disk Drives
  - Reducing I/O wait time can allow for higher server utilization

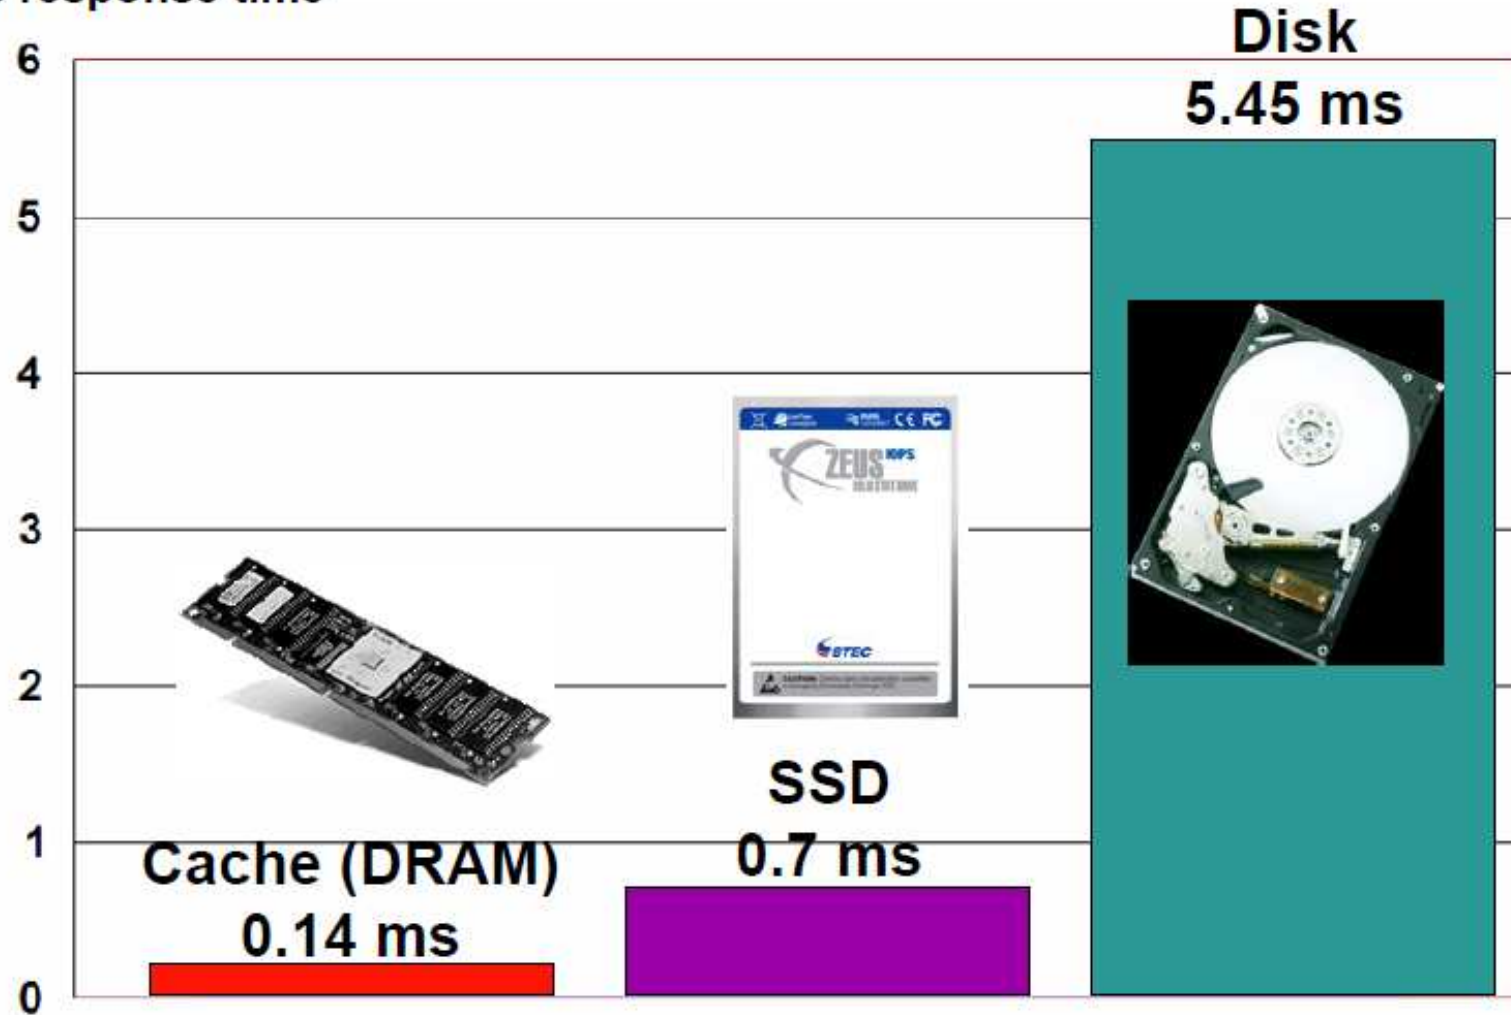


Even well-tuned databases have the opportunity to improve performance and reduce hardware resources

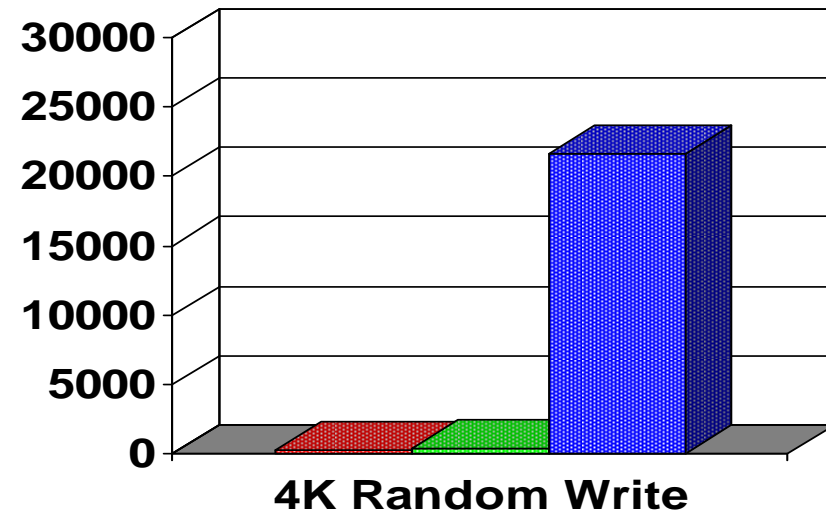
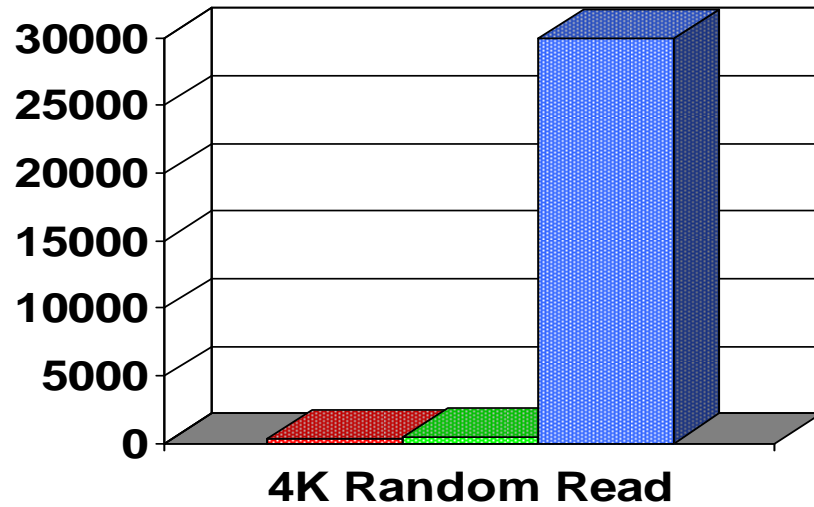
Source: If applicable, describe source origin

# Latency – Read Operation

ms response time



## Random Performance (Iometer)



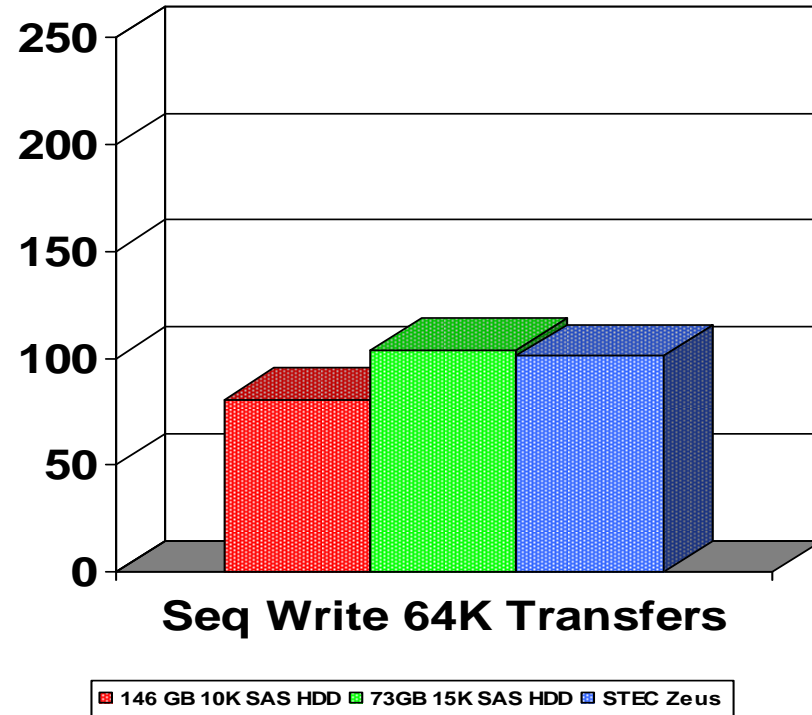

■ 146GB 10K SAS HDD   ■ 73GB15K SAS HDD   ■ STEC Zeus

■ 146GB 10K SAS HDD   ■ 73GB15K SAS HDD   ■ STEC Zeus

•STEC Zeus SSDs have a large advantage over HDD drives in random 4K IOP workloads.




# Sequential Write Performance (Iometer)



## Potential Systems Usage of Flash based SSDs


- Database OLTP & Data Warehouse
- File systems
- Collaboration
- SOA and messaging infrastructure
- Content Management, metadata, ingest
- Stream analytics
- High-event rate messaging systems
- System availability: fast dump, checkpoint, mirror, replication
- Operating System fast paging store for real memory partitions
- Hypervisor fast paging store for shared memory partitioning
- Ultrascale internet applications
- High Performance computing



## Client Challenges with SSD

*Inefficient use of a very expensive asset is difficult to justify*

- SSDs are considerably more expensive than traditional disks
- Without optimization tools, clients have been over-provisioning them




**Result: Many clients feel they can't afford solid-state storage yet**

## Solid-state drives (SSDs) positioned to address this

*New Tier-0 drives for high priority, time-sensitive applications*

### ▪ Potential client benefits

- **Increase revenue opportunities**
  - More transactions in less time
- **Reduce storage infrastructure costs**
  - Reduce acquisition and operating costs
- **Reduce server infrastructure costs**
  - Smaller servers, DRAM memory capacity, cost and power
- **Improve availability**
  - Lower component failure rates and faster error discovery
- **Enable new capabilities**
  - New functions and applications become feasible



### Solid-state drives

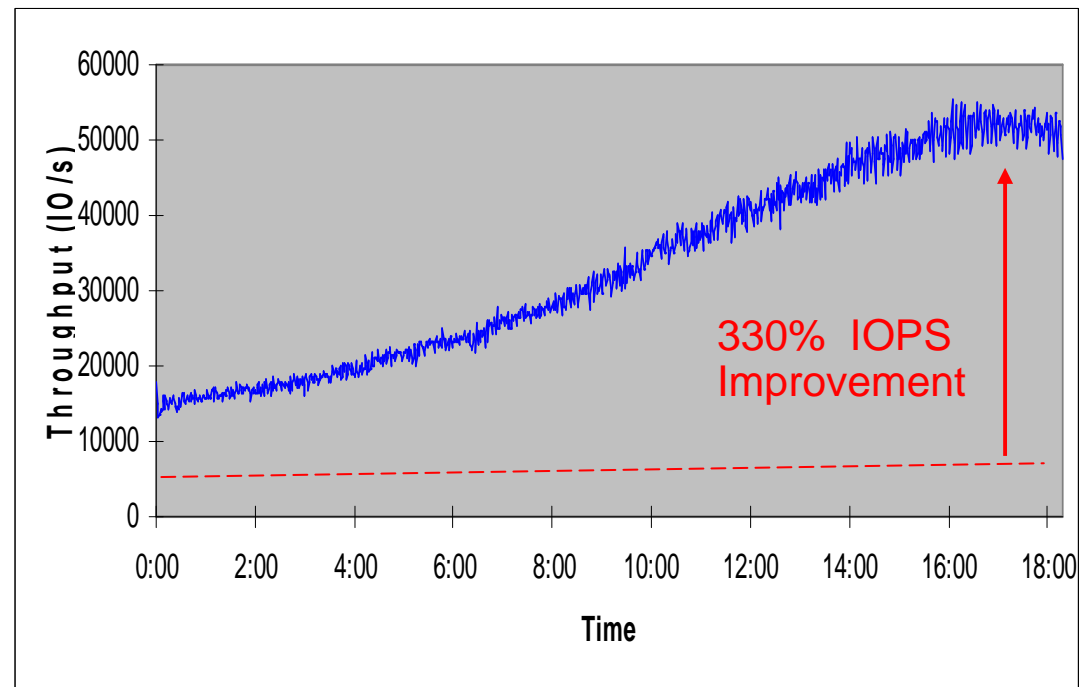

- Random access storage
- Non-volatile, semiconductor (NAND flash)
- No mechanical parts
- No rotating parts
- Same form factor as traditional HDDs

## Smart data placement with Easy Tier

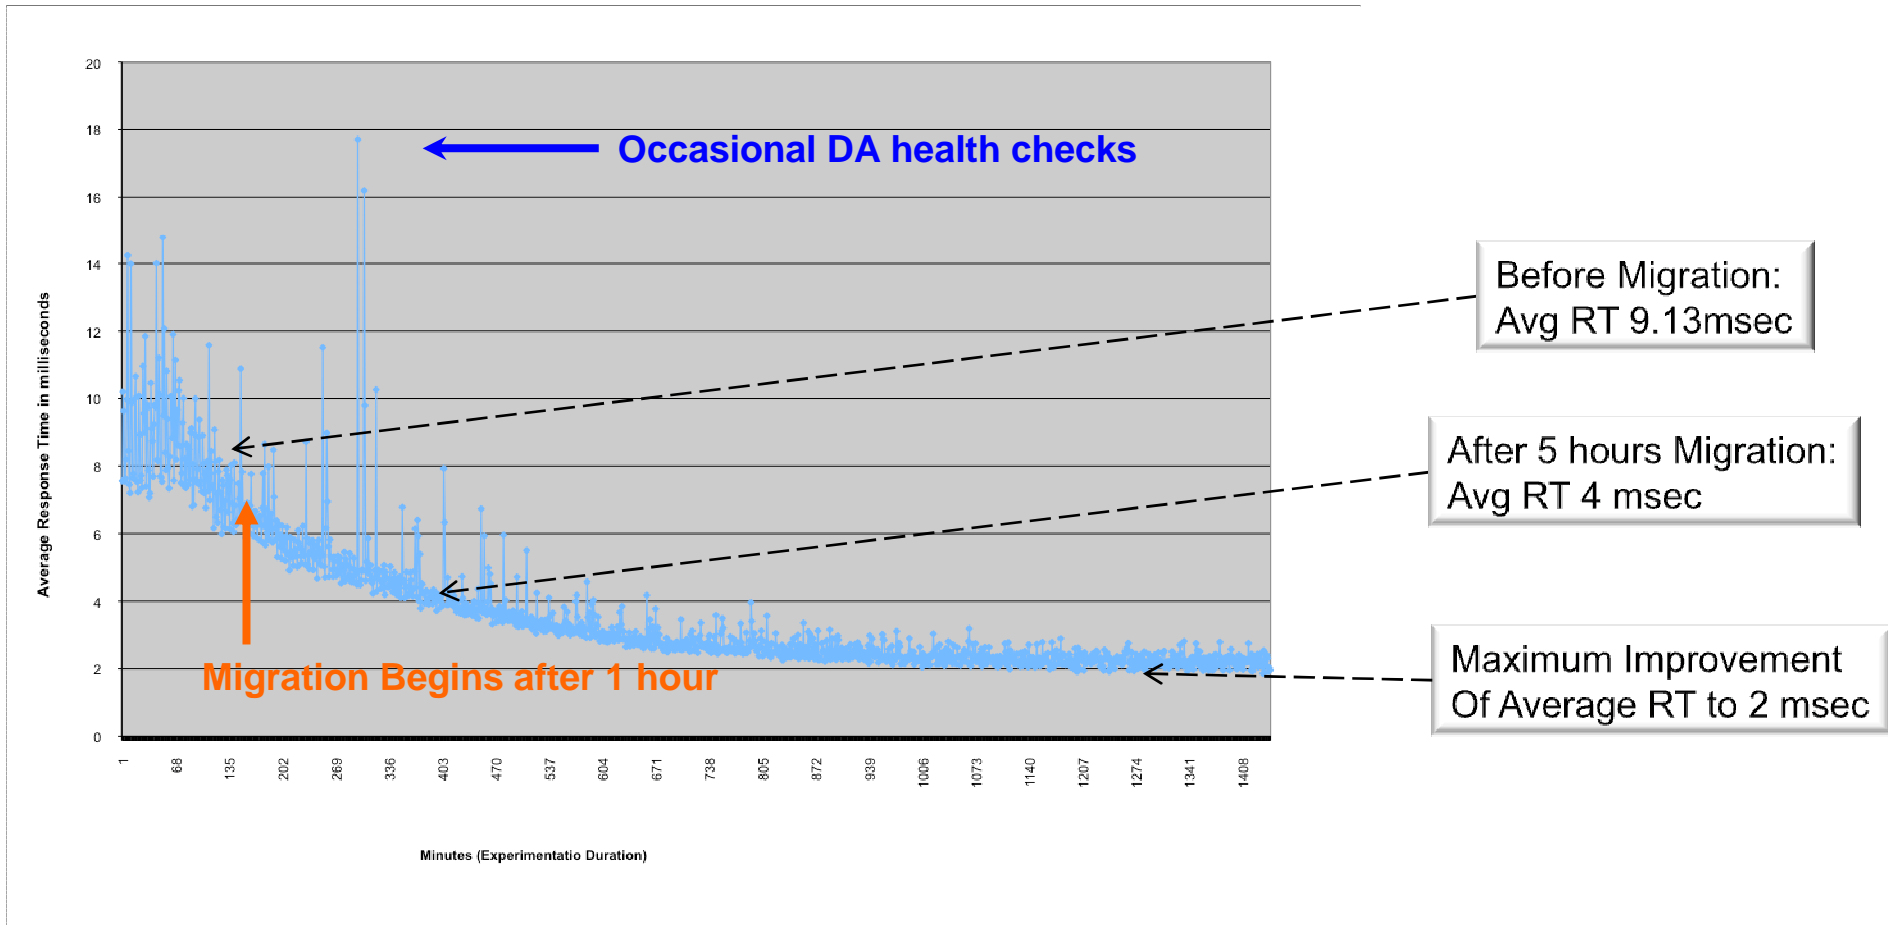
- First ever Storage Performance Council (SPC-1) benchmark submission with SATA and SSD technology
  - System configuration: 2.3 TB SSD + 96 TB SATA

Increase of  
**330%!**

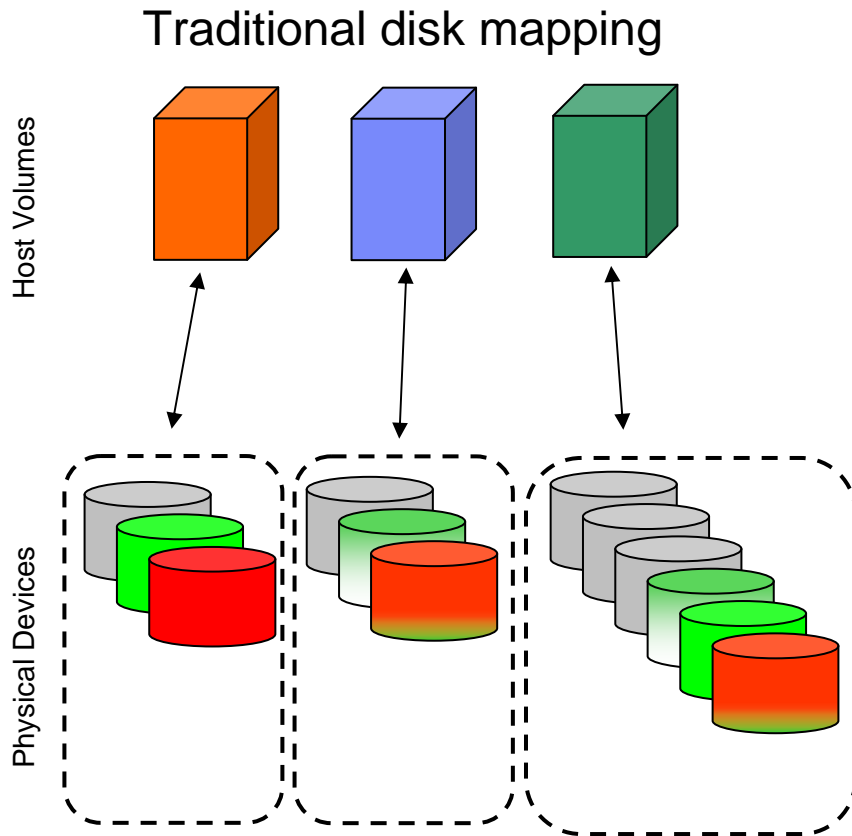
### Easy Tier




# Significant Improvement on Average Application Response Time Concurrently with Live Workload



# Traditional Disk Mapping vs. *Dynamic Virtual* Storage Mapping

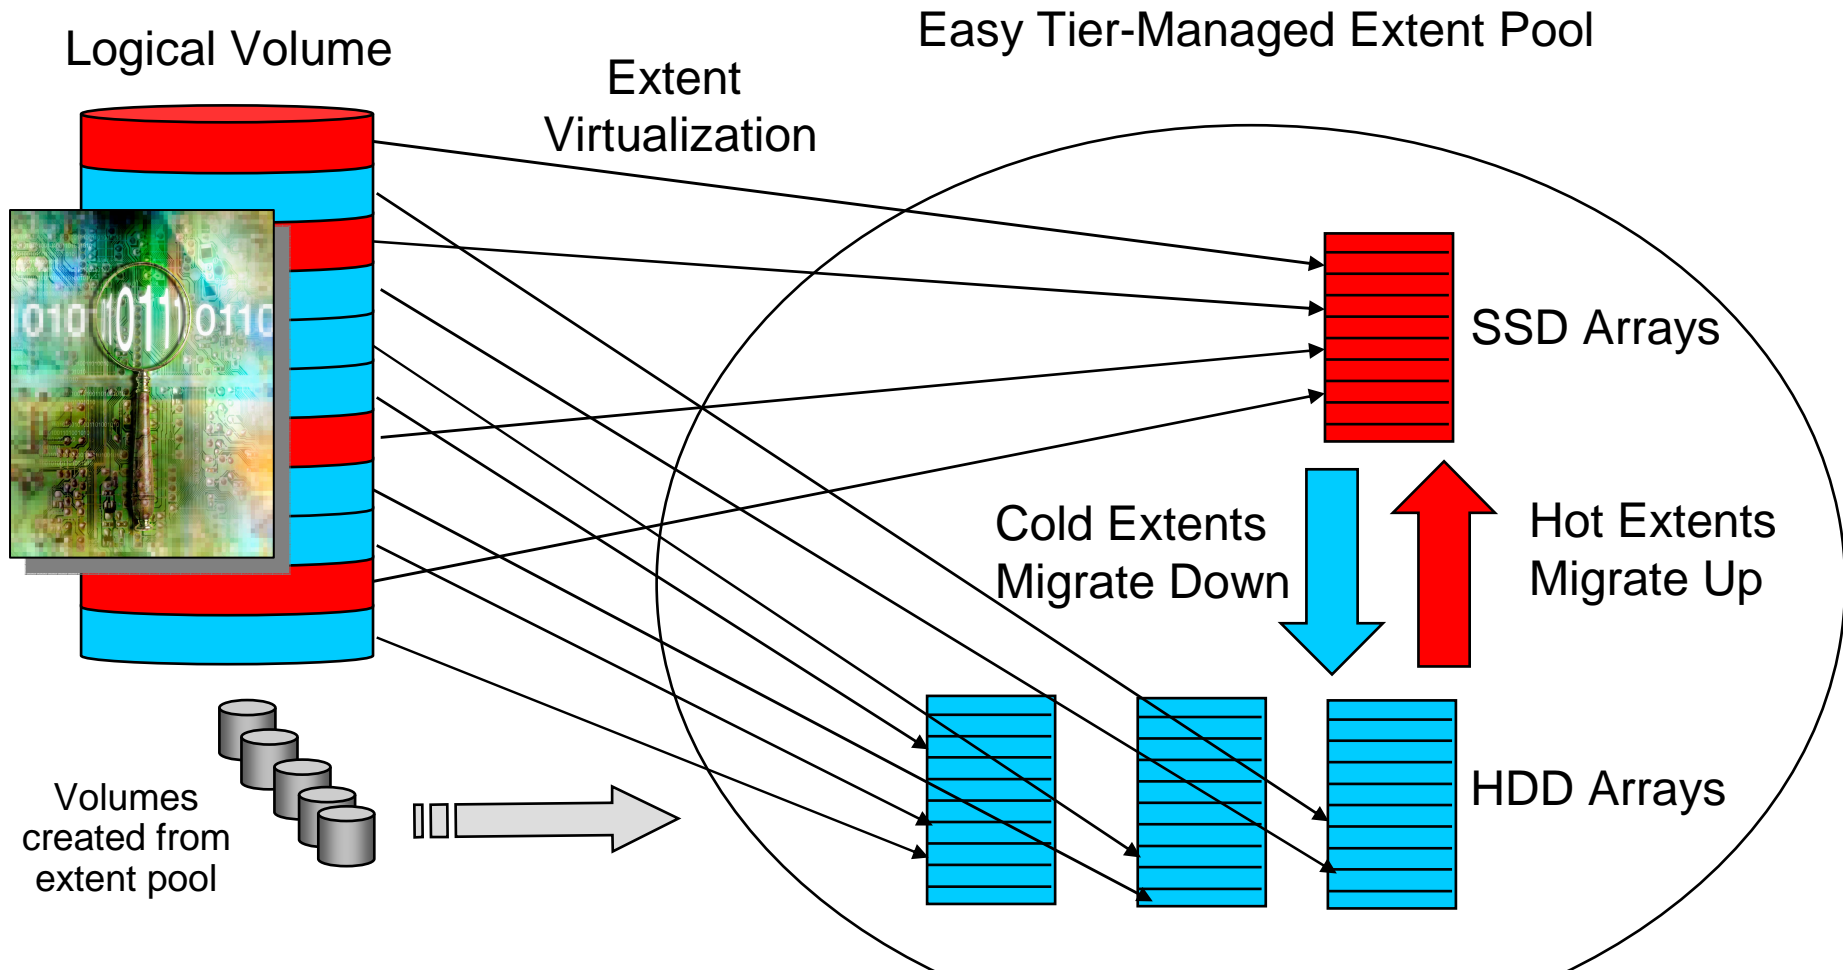


Volumes have different characteristics. Applications need to place them on correct tiers of storage based on usage.



All volumes appear to be “logically” homogenous to apps. But data is placed at the right tier of storage based on its usage through smart data placement and migration.

# Easy Tier automates movement of extents based on performance




1. Easy Tier monitors performance of each extent to determine the data 'temperature'
2. Easy Tier creates extent migration plan for optimal data placement every 24 hrs based on performance statistics
3. Easy Tier migrates extents within an extent pool according to plan over 24 hour period (limited number of extents are chosen for migration every 5 minutes)



# IBM Storage Tier Advisor tool – providing guidance with current workloads

IBM System Storage™ DS8000



Summary Report

[Volume Heat Distribution on LPAR SF75FK840ESS01](#)

### System Summary

The data is collected from Tue Dec 08 01:20:55 PST 2009 to Tue Dec 08 09:23:55 PST 2009.

Total Volumes Monitored	4
Total Capacity Monitored	416G
Hot Data Capacity (% of Total)	61G (14.6%)
Capacity Allocated on SSD/Total SSD Capacity	0G/146G
Estimated Migration Time	2 hours
Random I/O Percentage (Random of Total)	100% (3162056 of 3162056)
Data Validity	Valid
System State	Latest Warmstart: No Warmstart Latest Failback: No Failback

### Recommended SSD Configuration\*1

SSD Configuration	Predicted Performance Improvement*2
Take Advantage of Existing SSD Spare Capacity(RAID 5)*3	35%~ 55%
Add 1 New 73G SSD Rank Pair(s) (RAID 5)*3	54%~ 74%

\*1 The recommended SSD configuration is only the suggested SSD capacity to add or to take advantage of the existing S resource, for detailed physical configuration, please consult IBM service team.  
 \*2 The predicted performance improvement is the possible response time reduction at the backend in a balanced system configuration, and it may vary with different system workload and configuration.  
 \*3 Assume the rank pair will be configured as RAID5 (6+P+S), and the equivalency capacity is 876G.

- Calculates amount of Hot Data Capacity and estimates migration time to move hot data to SSDs

- Provides guidance on how current workload can benefit from existing and additional SSDs
- ROI can be determined by rate of performance change
- Recommendation can change over time

# IBM Storage Tier Advisor tool – providing guidance with current workloads

IBM System Storage™ DS8000

## Volume Heat Distribution

The data is collected from Tue Dec 08 01:20:55 PST 2009 to Tue Dec 08 09:23:55 PST 2009

Volume ID <sup>*1</sup>	Extent Pool ID	Configured Size <sup>*2</sup>	Capacity on SSD <sup>*3</sup>	Heat Distribution <sup>*4</sup>
0x0001	0x0000	100G	0G	90G <span style="float: right;">10G</span>
0x0002	0x0000	100G	0G	92G <span style="float: right;">8G</span>
0x0003	0x0000	100G	10G	70G <span style="float: right;">30G</span>
0x0004	0x0000	116G	0G	93G <span style="float: right;">13G</span>

20 Entries Per Page

GO
| < >> > |
Displaying Page 1 of 1

\*1. Volume ID is the id to represent the volume inside DS8000, which is specified during the volume creation.  
 \*2. The configured capacity of the volume.  
 \*3. The Capacity on SSD means the capacity of data already be migrated to SSD.  
 \*4. The Heat Distribution shows the heat distribution of the data in this volume, and blue bar is the capacity of the cold data, and the red bar is the capacity of the hot data.

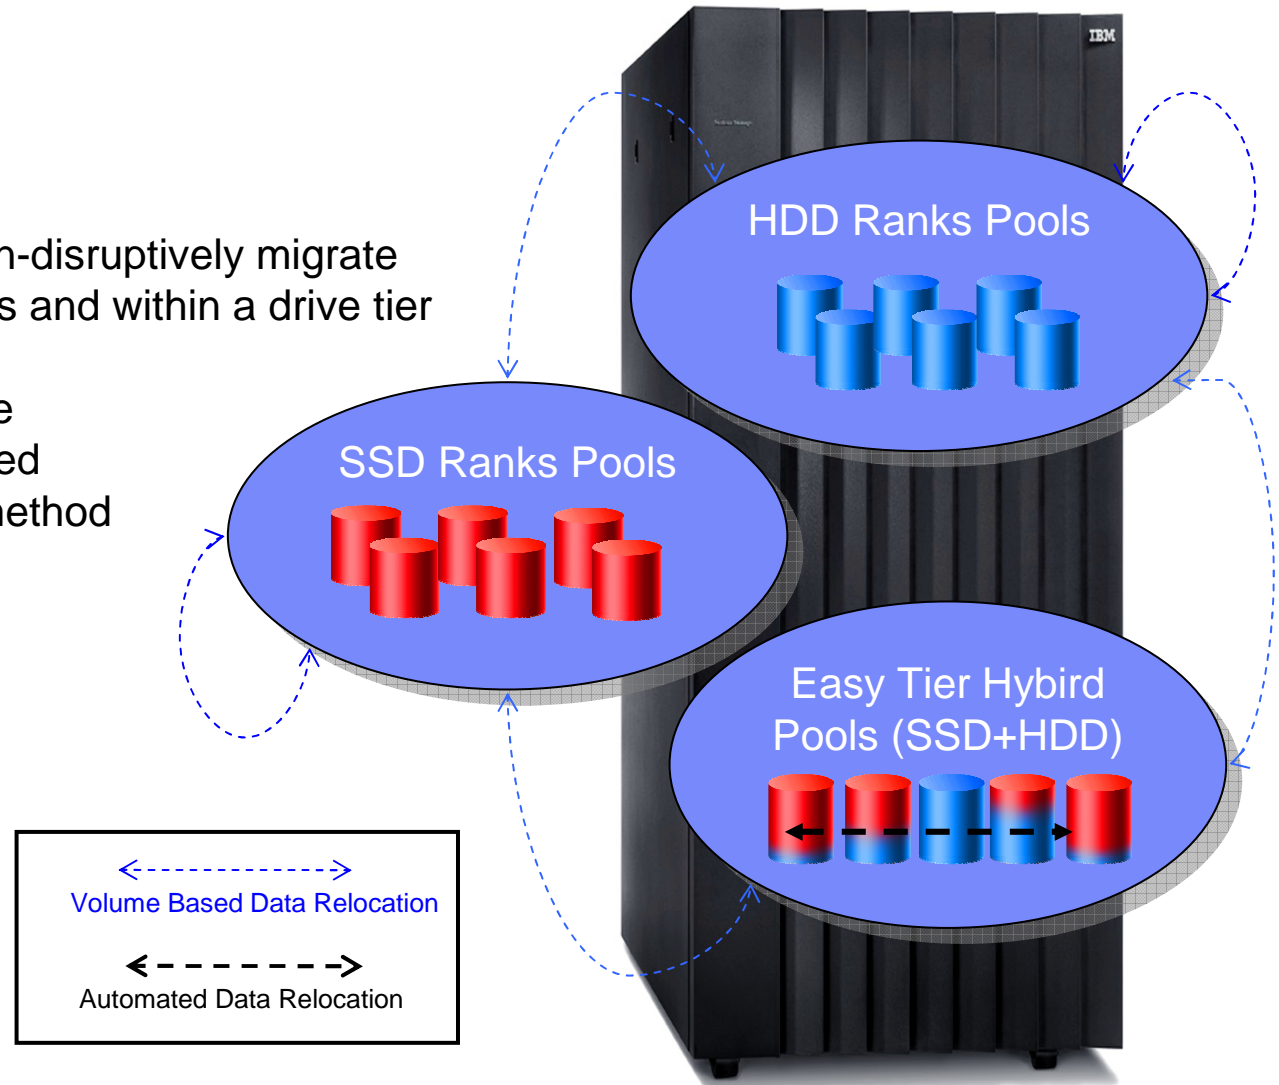
**Summary Report**

Volume Heat Distribution on LPAR SF75FK840ESS01

- Shows amount of hot data in each volume
- Shows capacity in each volume that already has SSD capacity
- 30GB is the hot data in volume 3 and 10GB is already in SSD ranks

# Dynamic volume relocation with Easy Tier

- Clients can manually and non-disruptively migrate full volumes across drive tiers and within a drive tier
  - Migrate to new disk type
  - Migrate to new RAID type
  - Migrate to new drive speed
  - Migrate to new striping method



<http://www.ibm.com/systems/storage>



## IBM announces IBM System Storage™ Easy Tier solid-state storage


The IBM System Storage™ DS8700 is the most advanced enterprise storage system in the industry's most extensive storage portfolio and is designed to support the most demanding business applications with its superior performance, resiliency, scalability, security and value.



### New IBM System Storage Easy Tier

The latest DS8700 release offers the new IBM System Storage Easy Tier feature that enables clients to deploy solid-state storage confidently, effectively and economically by automatically and dynamically moving only the appropriate data to the SSDs in the system, based on ongoing performance monitoring. Such effective storage tiering will help clients enjoy the performance benefits of SSDs without requiring administrators to create and manage storage tier policies and without the excessive costs associated with placing too much of the wrong data on these relatively expensive drives. Easy Tier also includes the ability to manually and non-disruptively relocate full logical volumes around the system, which provides additional flexibility and control



**We're here to help**

 Easy ways to get the answers you need.


-  Chat now
-  Request a quote


Or call us at:  
1-866-883-8901


**Highlights**



# Podcast

IBM System Storage Easy Tier - Solid-State Storage Made Simple 




 Allen Marin  
mehr

Inhalt


- 1. IBM System Stora... 0:28
- 2. Maximizing the bu... 0:59
- 3. Performance cons... 2:33
- 4. Why is solid-state... 1:17
- 5. Solid-state drives ... 0:43
- 6. Client challenges ... 0:45
- 7. The solution: IBM... 6:24
- 8. Other DS8700 Rel... 2:47
- 9. Summary
- 10. Untitled

Gesamtdauer: 15:54 / 15:56

13 April 2010 

## IBM System Storage Easy Tier Solid-State Storage Made Simple

Allen Marin  
IBM System Storage Marketing



**IBM System Storage**  
Smart movement and management of  
information and capacity growth  
without complexity

→ Announcing next generation storage technologies

© 2010 IBM Corporation

## New drive options

- **New 600GB 15k rpm Fibre Channel Drives**
  - Enterprise performance with greater capacity
  - Feature Codes #2716, 2717 (CoD)
  
- **New 2TB 7.2k rpm SATA Drives**
  - Doubles raw capacity
  - Feature Codes #2916, 2917 (CoD)
  
- **8-drive minipack option for SSDs**
  - Aligns with Easy Tier SSD granularity
  - Limitations apply (see backup charts)
  - Feature code
    - 73GB – #6014; no CoD option
    - 146GB - #6114; no CoD option



## Other DS8700 Release 5.1 enhancements

- **Faster concurrent, non-disruptively code load times for faster microcode updates (down to two hours)**
  
- High Performance FICON Multi-track Extended Distance Support
  - Enables high-speed write performance at longer distances for clients running long distance, multi-site (10-100KM) environments.
  - High Performance FICON is the protocol that enables faster performance for z/OS workloads
  - Pricing is a one-time charge of \$10,000

## Notice on withdrawals

- **WfM (Announce Date 2/23, 2010 and Effective Date 6/11, 2010):**

- DS6000 Model 522 and some features
  - EX2's still available for current customers
- DS8100/DS8300 242x HW M/Ts Models 931 and 932 in EMEA (RoHAS)



- **WfM (Announce Date 5/11, 2010) and effective Sept-Oct 2010**


- 146 GB Fibre Channel Drive
- 146 GB Encryption FC Drive
- 1TB SATA
- DS8100/DS8300 242x HW M/Ts Models 931 and 932 in US, AP, LA, and CAN





## z/VSE V4.2 Enhancement: Parallel Access Volume (PAV)

- Allows a z/VSE V4.2 host to access a single ECKD disk volume with multiple concurrent requests
  - multiple addresses (alias) to a single logical device
  - enables more than one I/O operation to a single logical device
  - may reduce device queue delays
  - volume sharing – not file sharing
- PAV is an optional, licensed feature of IBM DS8000 and DS6000
  - no changes needed for application programs
- Examples of PAV candidates
  - VSAM catalogs, shared clusters, libraries
  - spool files, work files, log files
- Potential benefits include possibility of improved performance/throughput
  - multiple jobs, multiple partitions, CICS
  - gains are *highly dependent on workload*




# DS8000 PAV for z/OS Parallel I/O Processing

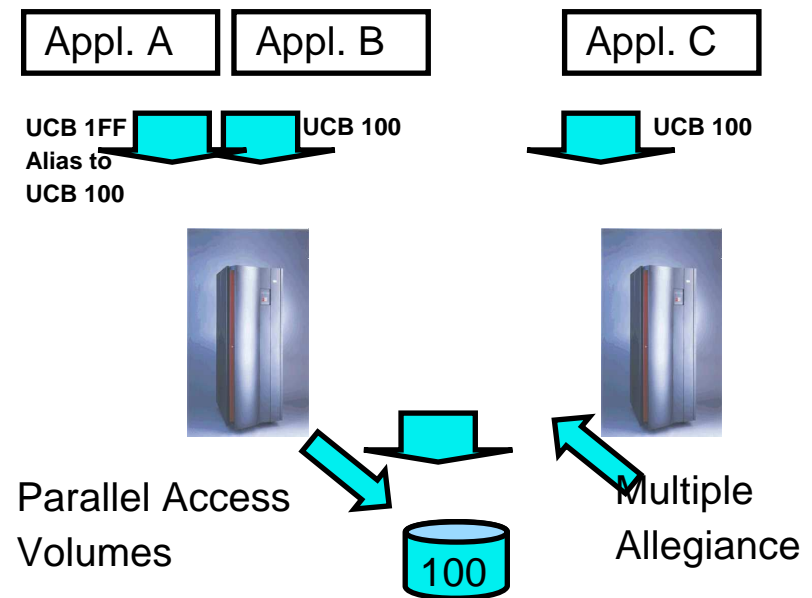
## *Throughput beyond raw performance*

- IBM PAV enables System z to process multiple I/Os to same logical volume at once
- IBM Multiple Allegiance expands capability across multiple System z servers
- I/O Priority enables use of info from z/OS Workload Manager to manage i/o processing order

### Traditional z/OS Operations



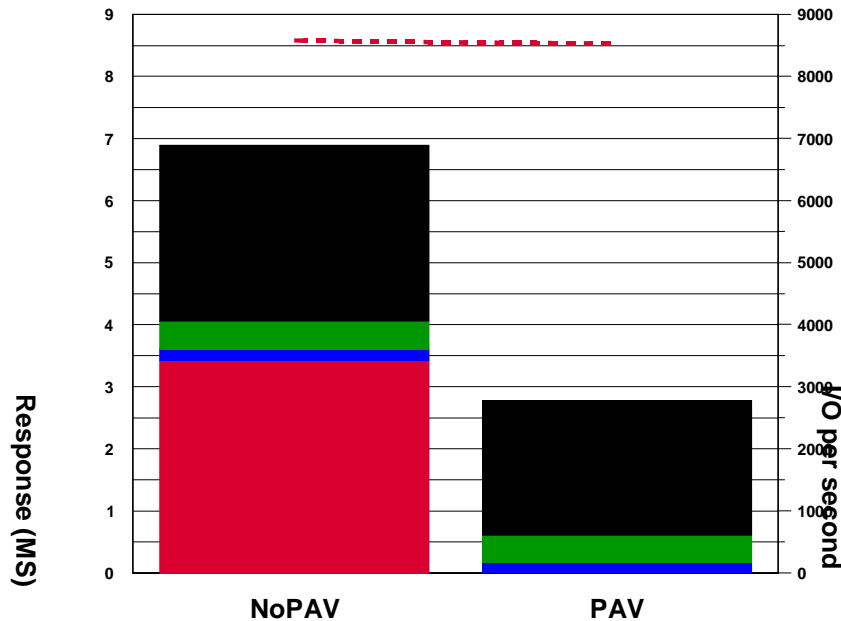
### Parallel Access Volumes/MA



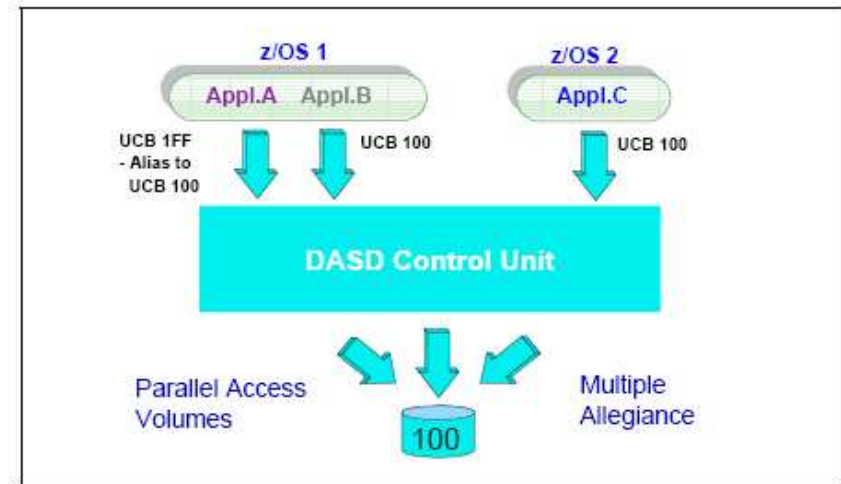
# IBM PAV and IBM Multiple Allegiance

*Enable different hosts to simultaneously process multiple I/O operations to the same logical volume*

- **Multiple Allegiance and PAV functions allow multiple I/Os to be executed concurrently against the same volume in a z/OS environment**
  - With Multiple Allegiance, the I/O are coming from different LPAR of z/OS systems
  - With Parallel Access Volumes, the I/O are coming from the same LPAR of z/OS systems
    - **Static PAV:** Aliases are always associated with the same Base Address
    - **Dynamic PAV:** Aliases are assigned up front but can be reassigned to any base address as need dictates: WLM function call Dynamic Alias Management - reactive alias assignment
    - **HyperPAV:** On demand/ Proactive alias assignment



<Queue  
<Pend  
<Connect  
<Disc  
<I/O Rate>



## IBM Information Infrastructure for Storage Virtualization – Scale Out *IBM System Storage SAN Volume Controller*


- **Value**
  - **Improves storage utilization** (+30%) and reduces storage growth
  - **Boosts performance** and simplifies storage management for IBM and non-IBM disk
  - **Improve storage administration productivity by up to 2x**
  - **Redundant architecture supports enterprise-class availability**
    - Non-disruptive upgrades of both hardware and software
  - **Supports non-disruptive data movement**
  
- **Powerful data management capabilities**
  - Space-Efficient Virtual Disks support on demand provisioning
  - Space-Efficient FlashCopy dramatically reduces storage needed for backup copies by as much as 75% or more
  - Virtual Disk Mirroring helps improve availability for critical applications

**Over 14,000 storage engines shipped running in more than 4,600 SVC systems**

***Outstanding performance,  
flexibility, and high availability  
while controlling storage TCO***

***Information Availability***

# Sweet Spot Architektur





## Appliance System

Concurrent Code und Hardware Upgrade  
 Concurrent Performance and Capacity Scaling  
 Skaliert bis 8 Knoten je Cluster



Any to Any Architektur  
 Jeder Speicher für jeden Knoten erreichbar  
 → Exzellente Ressource Utilization  
 → Durchgehendes Mirror Konzept

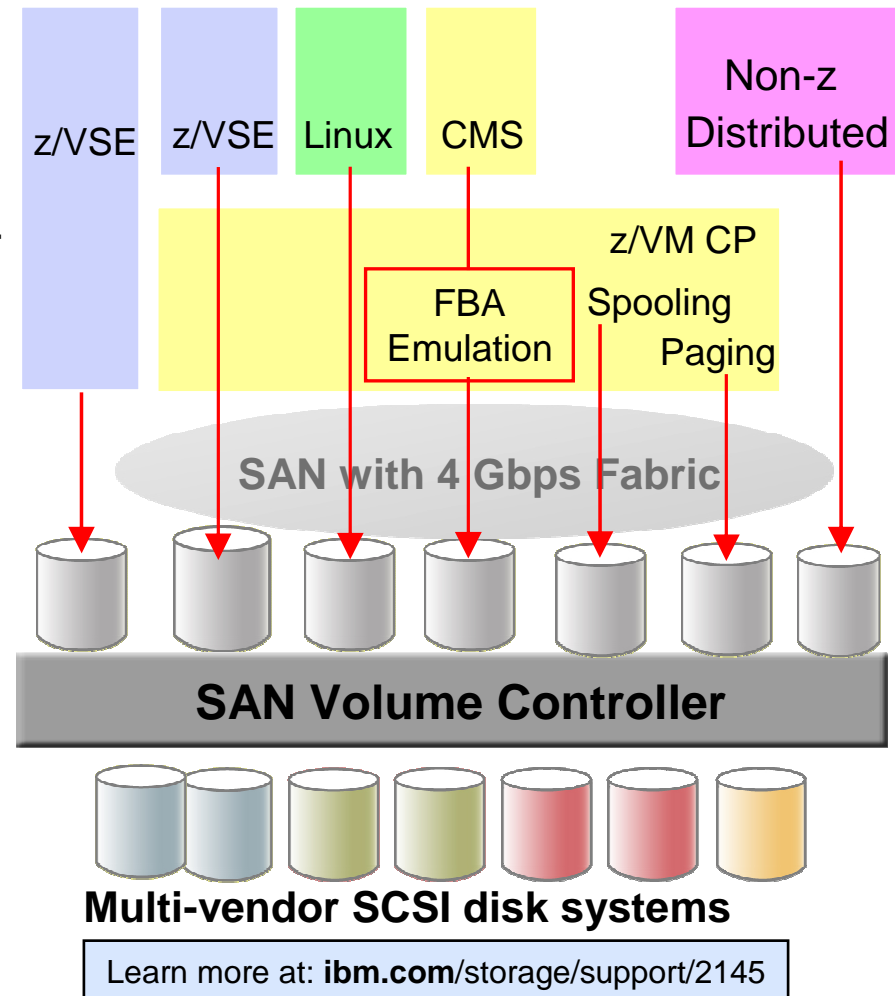


Enorme Entwicklung der Leistungsfähigkeit

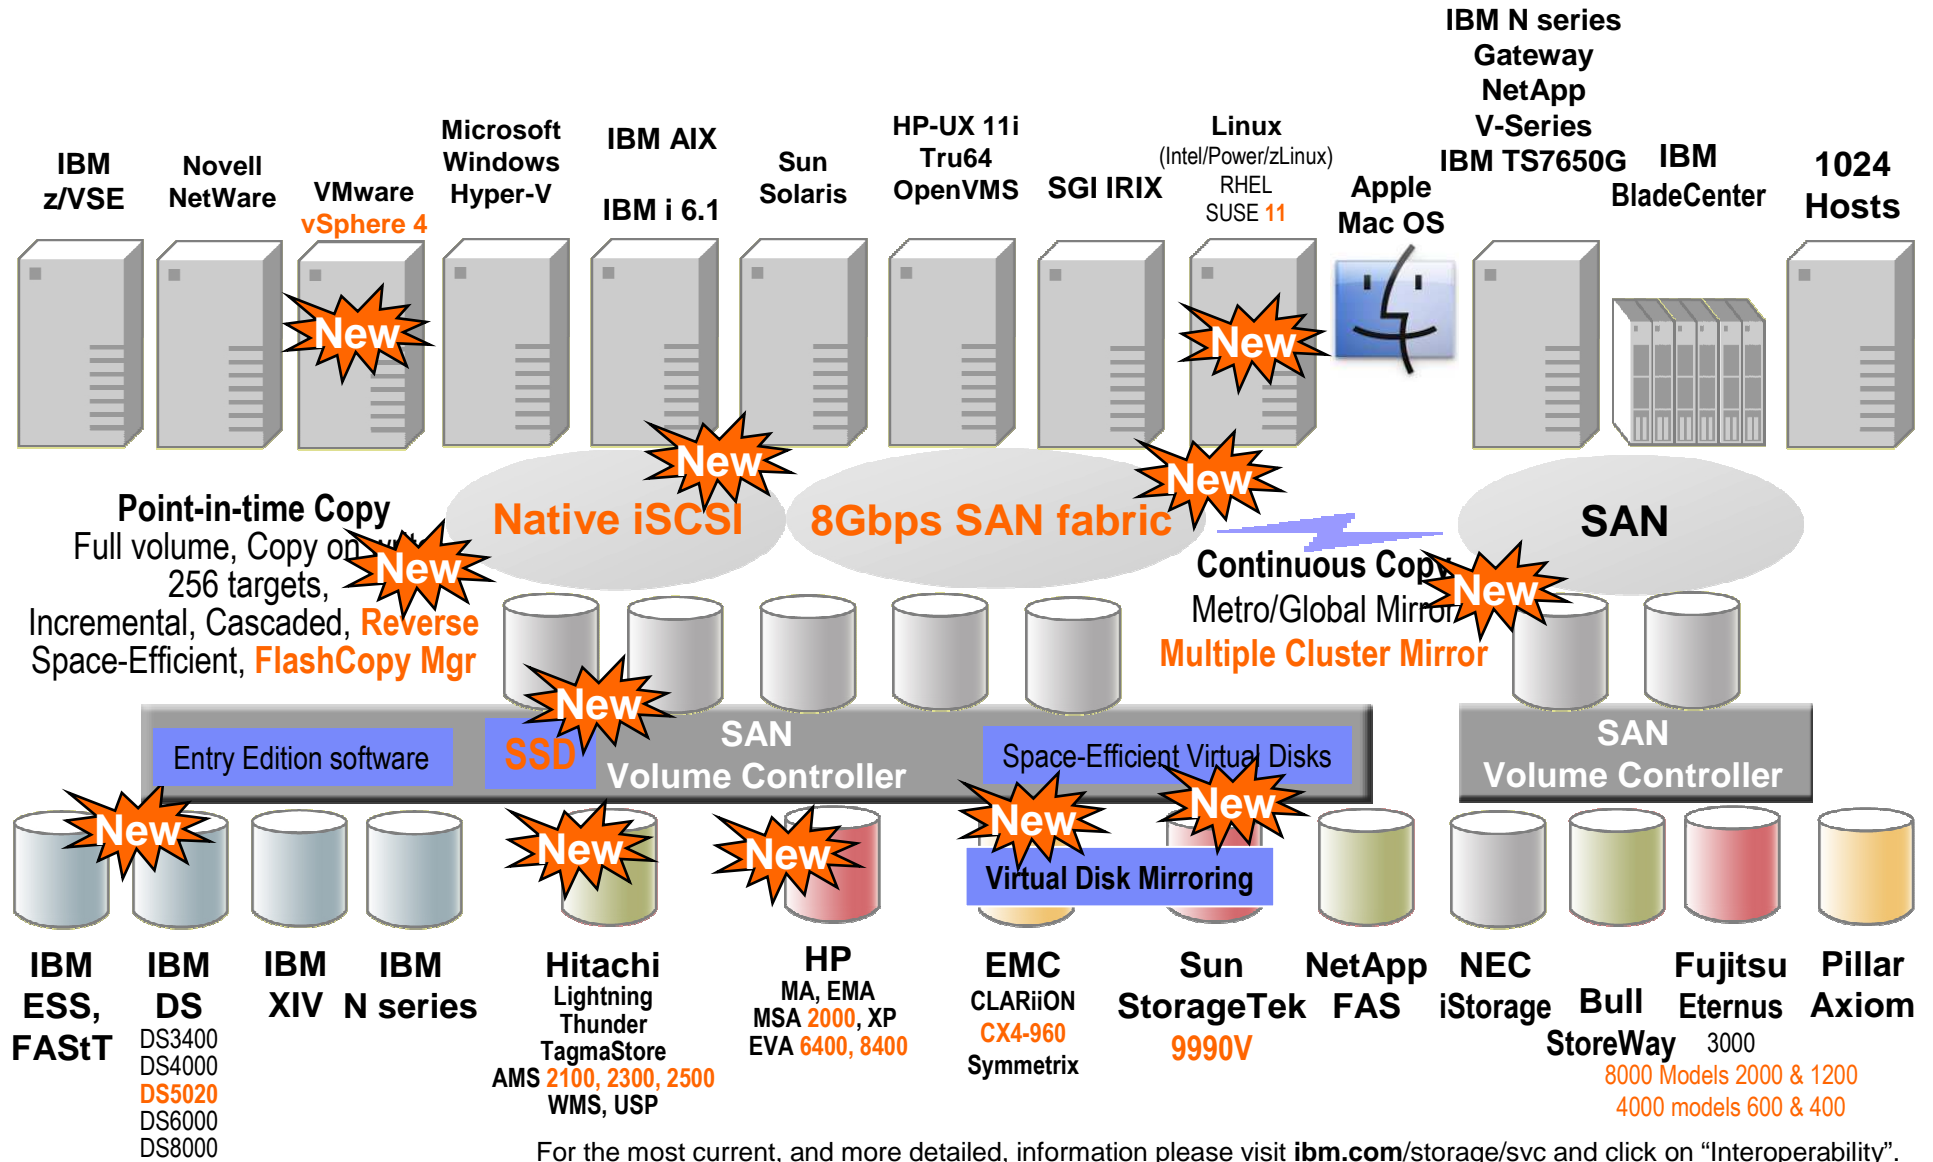
Nachgewiesen durch SPC1 Tests

## z/VSE V4.2 Enhancement: SAN Volume Controller (SVC)

- SAN Volume Controller (SVC) creates a single pool of SCSI disk capacity
- Disk storage options include IBM DS8000, DS6000, ESS, DS4000, etc. plus qualified systems from various non-IBM vendors
- *SVC platform* includes both hardware and software components:
  - SVC ‘nodes’ provide redundant components plus cache
  - Systems Storage Productivity Center (SSPC) software provides administrative and copy services
- Also supported in z/VM V5.3 and later, as well as Linux on System z



# Asset Utilization: Virtualizing Existing IT with SVC



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

# DS8700 & SVC-CF8

Component	Features	Quantity
Storage Engines	CF8	4 and 6
DS8700 Storage Systems	RAID 10	2
Disks	146 GB, 15K RPM	2048
Switches	IBM 2498-B24	4
Max Storage Capacity	RAID 10	138 TB
Max Storage Capacity	RAID 5	245 TB

Figure 4 Overview of tested DS8700 configurations.

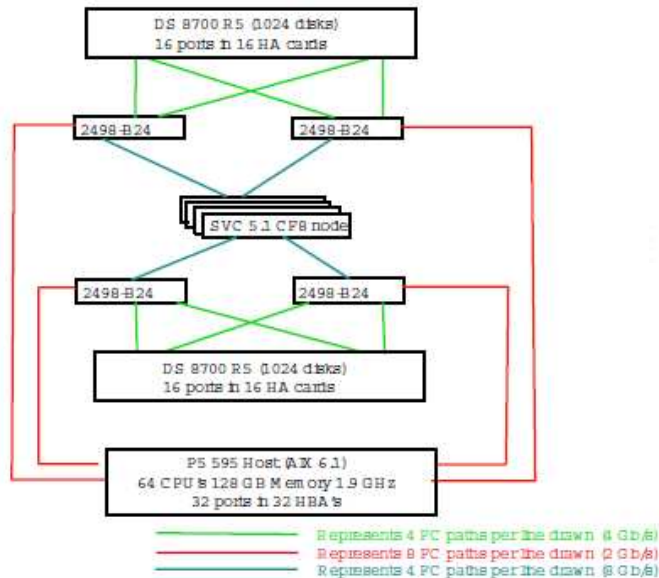


Figure 5 Tested SVC 5.1 configuration (4 node cluster with DS8700).

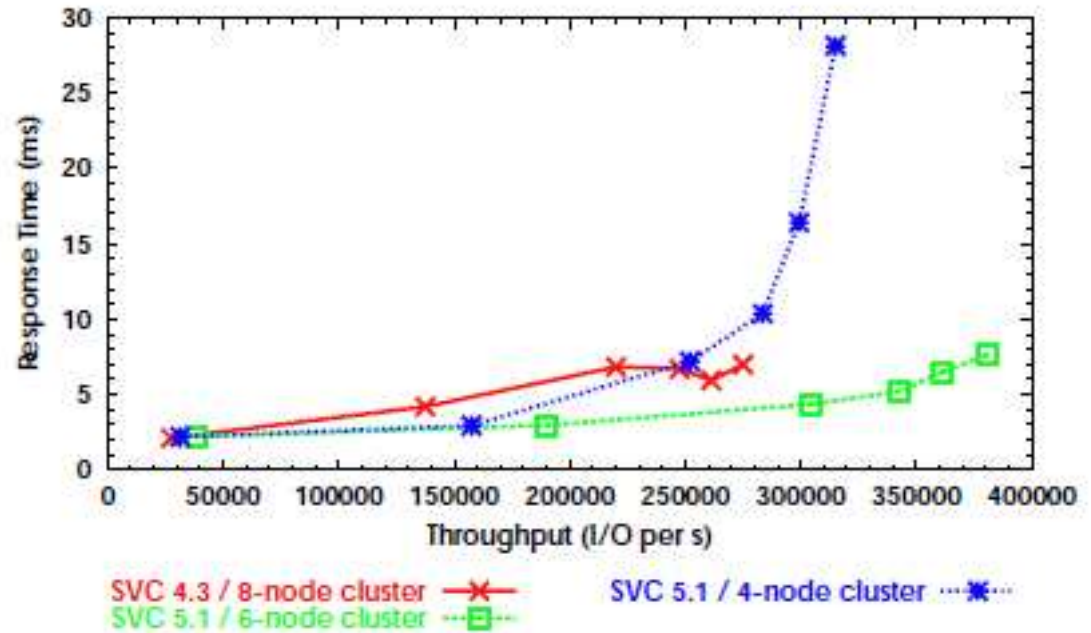


Figure 6 SPC-1 performance comparison of SVC 5.1 versus SVC 4.3.

## IBM® System Storage™ SAN Volume Controller Version 5.1 Enterprise Performance

Section 3 uses the SPC-1 workload to illustrate a large OLTP environment built using SVC and DS8700 technology. The results reported in that section are also published on the Storage Performance Council web site [1, 2]. The publication of these results yielded a new record for SPC-1 throughput, breaking the previous record by more than 25 percent. Like many of the results of Section 2, those of Section 3 demonstrate a better than factor of two increase in the throughput per node achieved with SVC 5.1.



## IBM Scale Out Network Attached Storage (SONAS)

### *The “C:” drive of the enterprise*

- **Just like you, enterprises have to manage files on drives and in directories. It’s a pain.**
  - High administration costs from managing multiple file servers.
  - Hard to know where a specific file is when you need it.
- **SONAS makes it easy.**
  - Single namespace – like one big “C:” drive but everyone can share it.
  - Up to 14 Petabytes of storage.
    - Starting point is for customers needing 100s of TBs of storage.
  - 180 million 4-drawer filing cabinets full of text.
    - Every MRI done in the world this year.
- **This is a game-changer. No other product can scale like this.**
- **Perfect storage solution for cloud implementations.**
- **Announced on February 9, 2010; GA on March 12, 2010.**



## Ist die Infrastruktur auf künftige Anforderungen ausgelegt?

### 30 Milliarden

Geschätzte Anzahl RFID-Tags, die 2010 produziert und in Produkten, Gebäuden, Fuhrparks, Pässen und sogar Tieren eingesetzt werden.

### 1 Billion

Die Anzahl der an das Internet angeschlossenen Geräte steigt von 500 Mio. im Jahr 2006 bis zum Jahr 2011 sprunghaft auf 1 Billion an.

### 10x

10-facher Anstieg bei den digitalen Daten zwischen 2006 und 2011 – 80 Prozent davon in unstrukturierten Formaten

### 1 von 3

Geschäftsführer treffen wichtige Geschäftsentscheidungen häufig anhand unvollständiger Informationen oder anhand von Informationen, denen sie nicht trauen.

### 70 Cent pro 1 €


Zur Wartung derzeitiger IT-Infrastrukturen werden durchschnittlich 70% der Budgets verwendet, nur 30% für neue Lösungen.

### 33 Prozent beenden Geschäftskontakt

33 Prozent der Kunden, die von einem Sicherheitsleck erfahren, brechen den Geschäftskontakt zu dem Unternehmen ab, das sie dafür verantwortlich halten.


Globally, file-based unstructured data is 80% of new data and growing

Worldwide Storage Capacity Shipped by Segment, 2008–2013



Source: IDC, State of File-Based Storage Use in Organizations: Results from IDC's 2009 Trends in File-Based Storage Survey: Dec 2009: Doc # 221138

## Unstrukturierte Daten- neben dem Wachstum ein weiteres Problem:

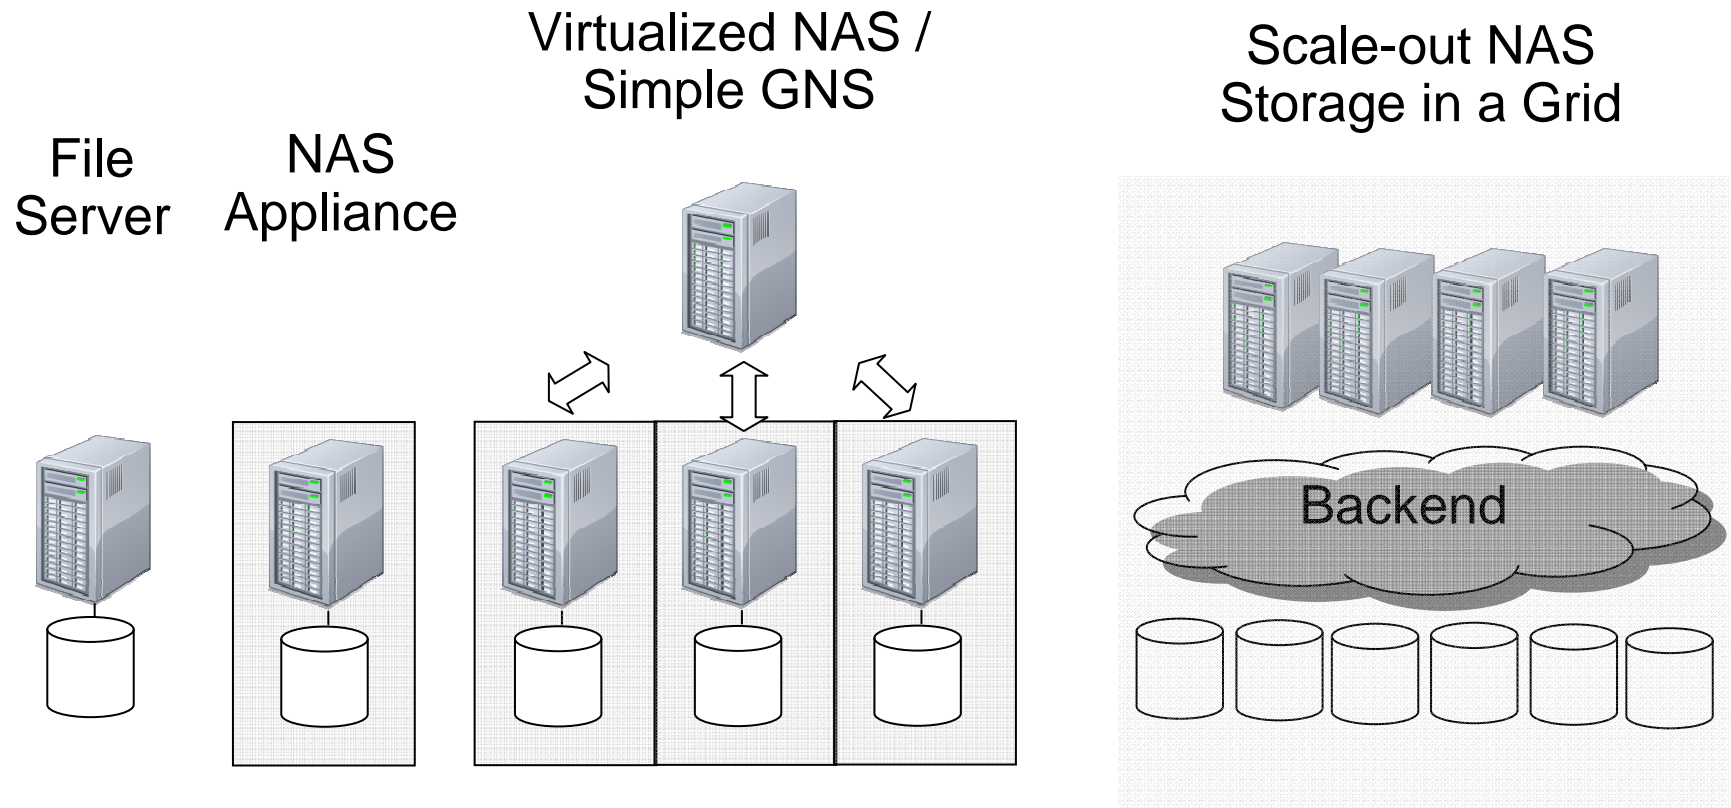


### Inactive Files verbrauchen Ressourcen wie:

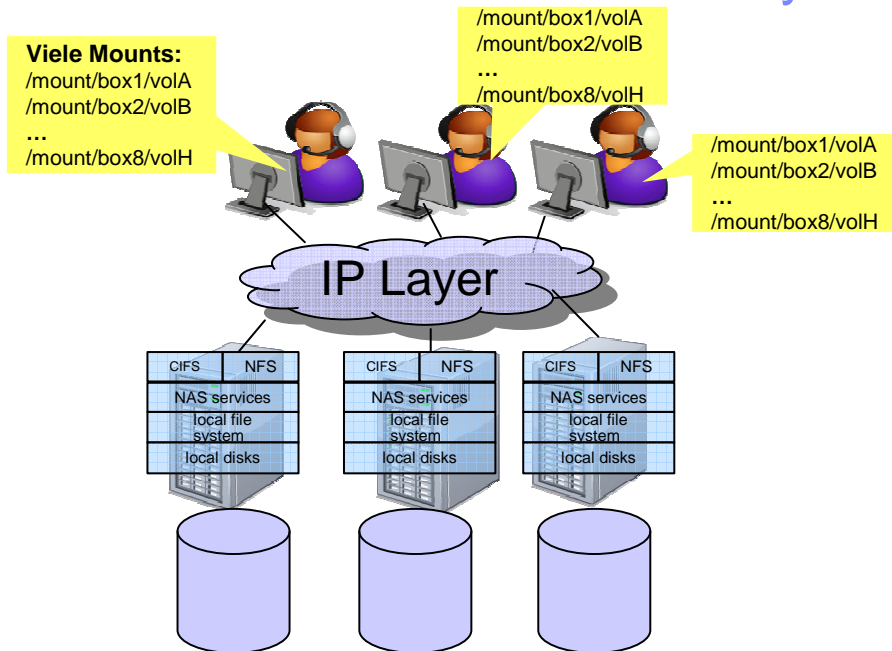
- Virens Scanner
- Aufwendigere Datenorganisation (store / find )- Content?
- Backup Volumen: „Rebuild Time einer 1TB Drive mit 1.5 Mio Files“
- Datenspiegelung, Archivierung

### Können SLA's noch eingehalten werden?

# Speicher für Unstructured Data - vom File Server zum Scale-out-NAS

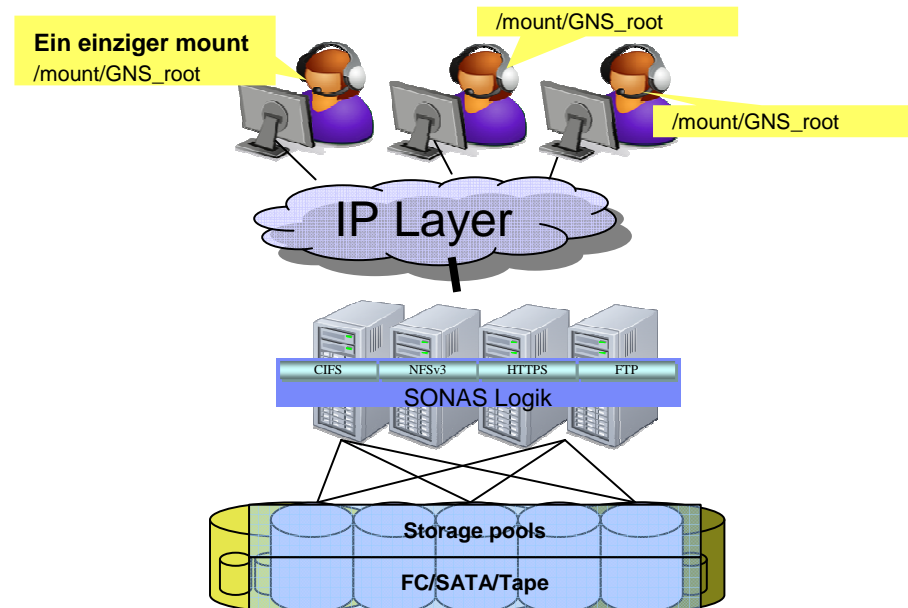


## Fileer oder skalierbare NAS-Systeme?



### Klassischer Fileer

- ~70% inaktive Daten
- Auslastung < 50%
- Anpassung der Datenverzeichnisse bei Änderung der Datenlokation
- Aufwendiges Backup je Fileer
- Skalierbarkeit nur in Fileergrenzen
- 54 ▪ Keine ILM Policy



### SONAS

- ~0% inaktive Daten (wg. ILM)
- Auslastung > 80%
- Global Namespace, auch bei Datenumzug
- Backup SW-Feature im Standard
- Skalierbarkeit über physikalische Grenzen
- ILM bis hin zum Tape

# SONAS Architektur



## Parallele Grid Architektur:

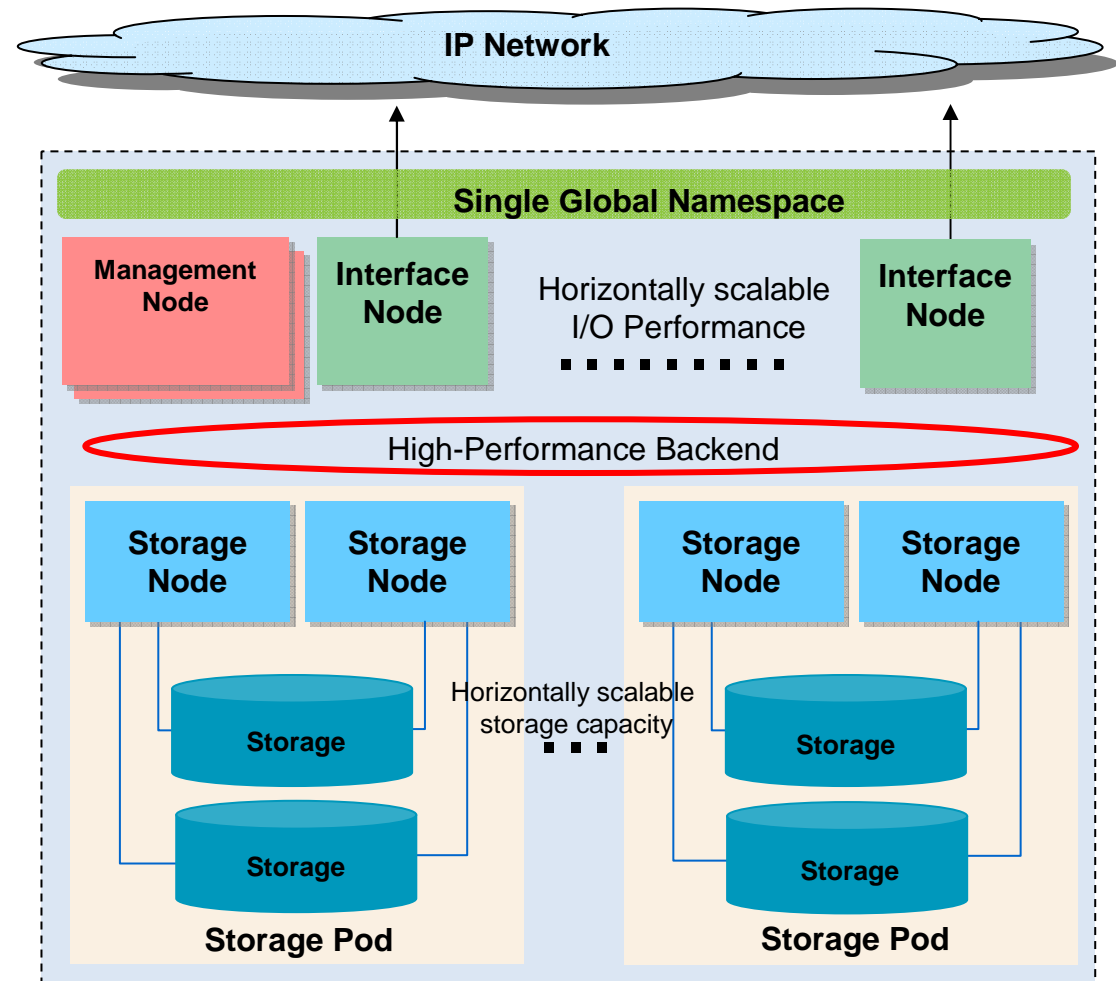
- Massive Skalierbarkeit
- Extreme Performance
- Automatische Speicherhierarchie
- Globaler virtueller FileServer


## SONAS Appliance:

- Management Node- bis zu 2
- Interface Node- bis zu 30
- Storage Pods- bis zu 30

das heißt:

- Alles Nodes sind clustered
- SONAS SW auf allen Nodes
- Einsatz „off-the-shelf“ Produkte





**SONAS**




**Cloud**






# Schreiben von Files

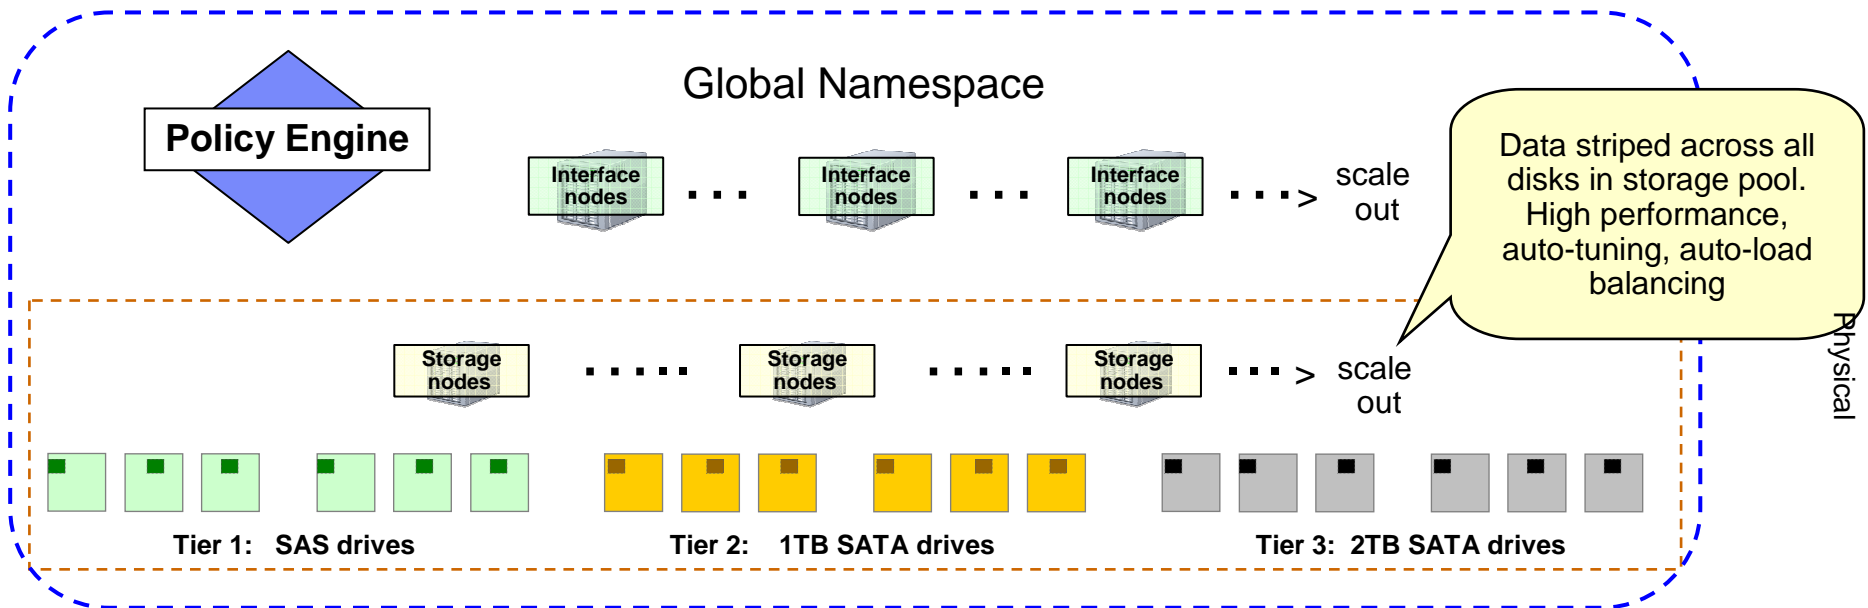


/home/appl/data/web/important\_big\_spreadsheet.xls  
 /home/appl/data/web/big\_architecture\_drawing.ppt  
 /home/appl/data/web/unstructured\_big\_video.mpg




**Note:** all three files, in same directory, but each allocated to *different* physical storage pool

## IBM Scale Out NAS



# Lesen von Files

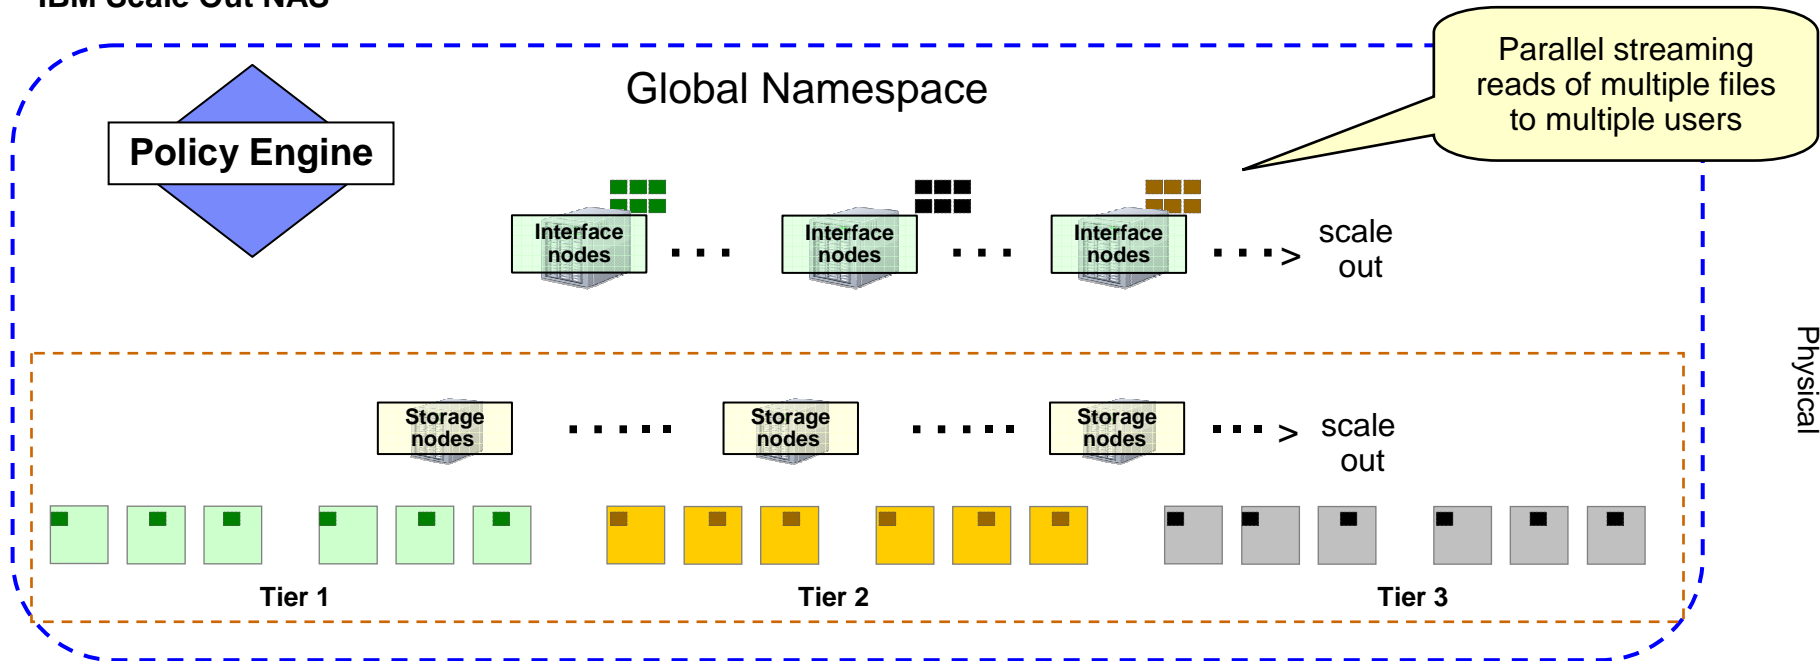


/home/appl/data/web/important\_big\_spreadsheet.xls  
/home/appl/data/web/big\_architecture\_design.ppt  
/home/appl/data/web/unstructured\_big\_video.mpg




Logical

## IBM Scale Out NAS



Physical

# Scan Engine



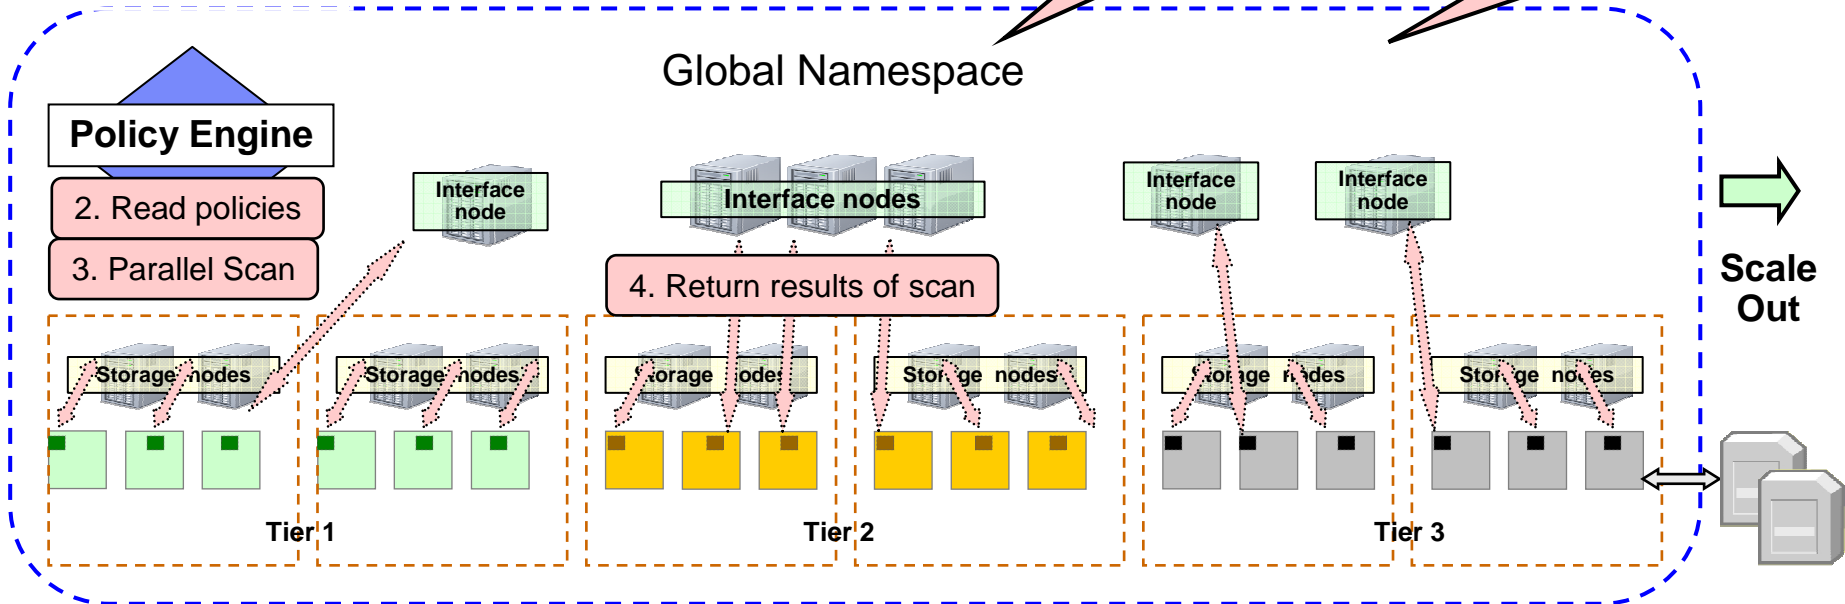
- Scan Engine reads internal SONAS file system metadata
- Does not need to read the file or directory tree
- All nodes can participate in scan of file system

1. Start scan

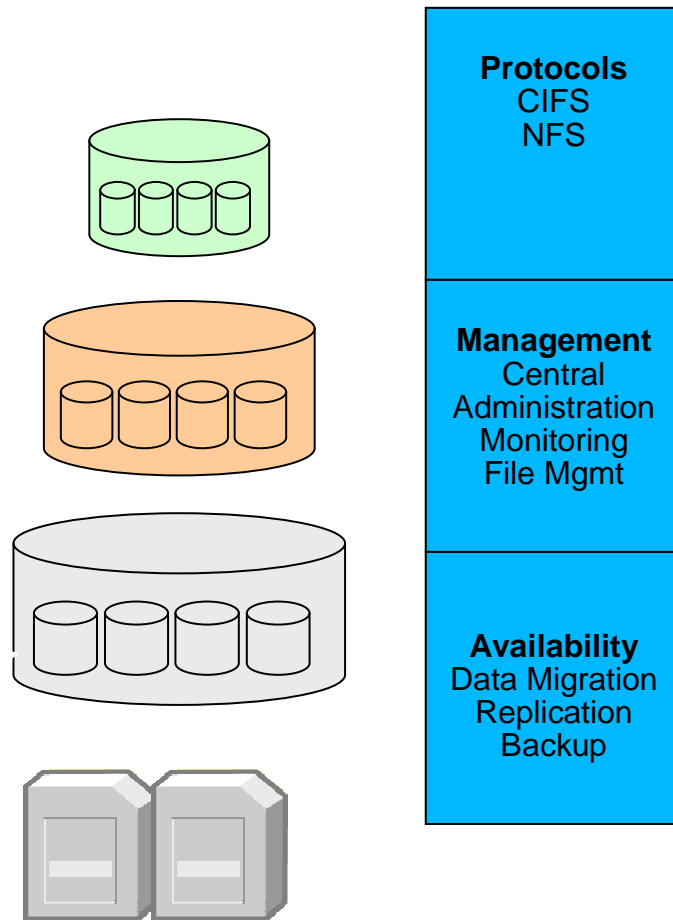
Metadata scan  
Scan > 15 million files/minute

Some or all nodes (both storage and interface) can participate in *parallel scan engine*

IBM Scale Out NAS




## ILM/HSM- Leben einer Datei von der Wiege bis in den Tod



- Automatic Tiered storage
- Liefert:
  - Einen globalen File System Name Space über unabhängige logische Storagepools
  - Dateien in gleicher Directory können in unterschiedlichen Pools liegen
  - Dateien werden über Policies zum Zeitpunkt der Anlage in Storagepools angelegt (SQL like statements)
  - Dateien werden automatisch (policy-driven) zwischen den Storage-Tiers verschoben (inklusive Platte und Tape)
  - Erlaubt klassifizierung der Daten analog der SLAs

# File Verschiebung -> ILM / HSM



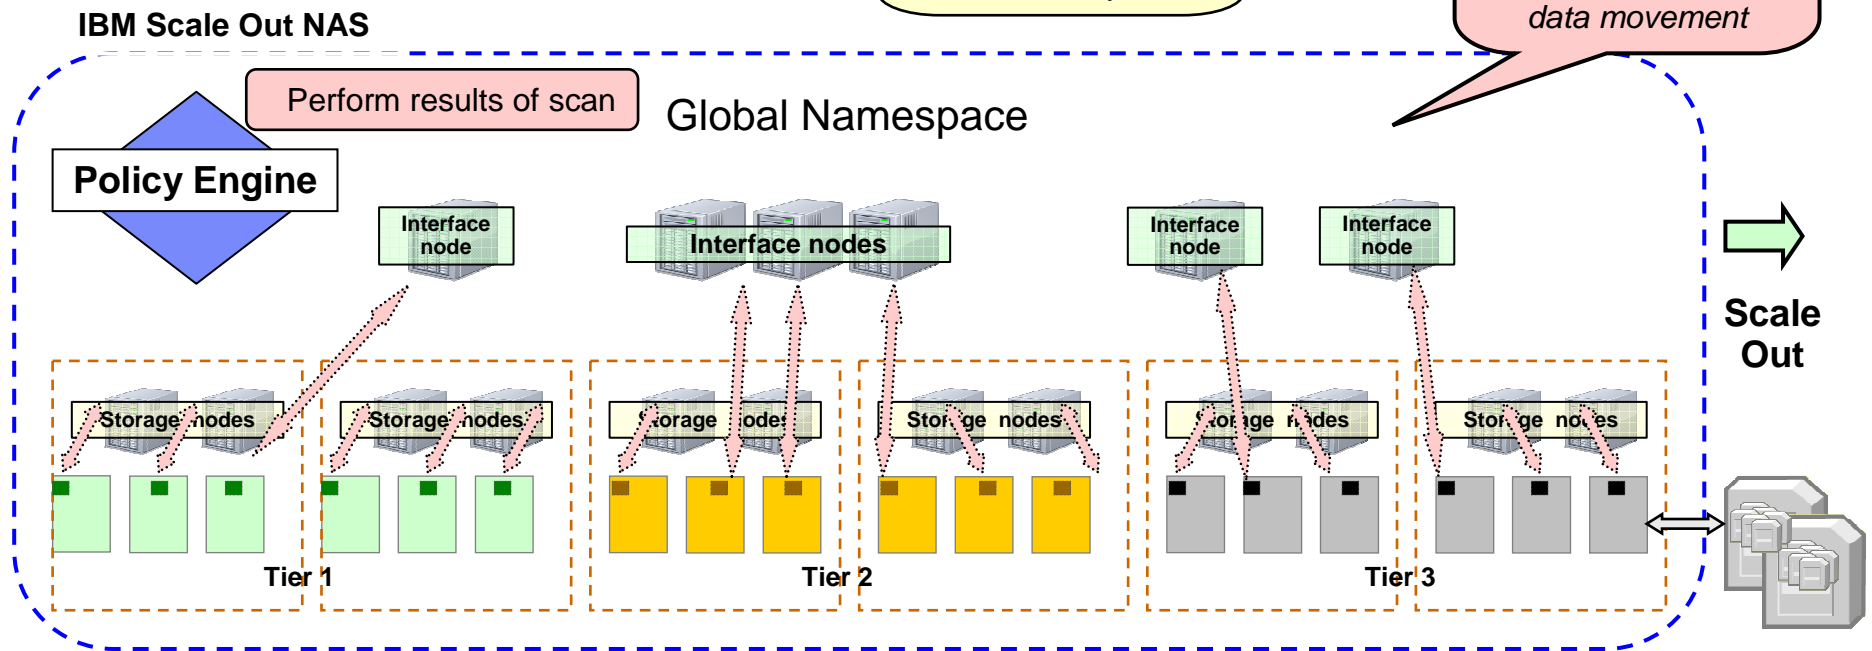
/home/appl/data/web/important\_big\_spreadsheet.xls

/home/appl/data/web/big\_architecture\_drawing.ppt

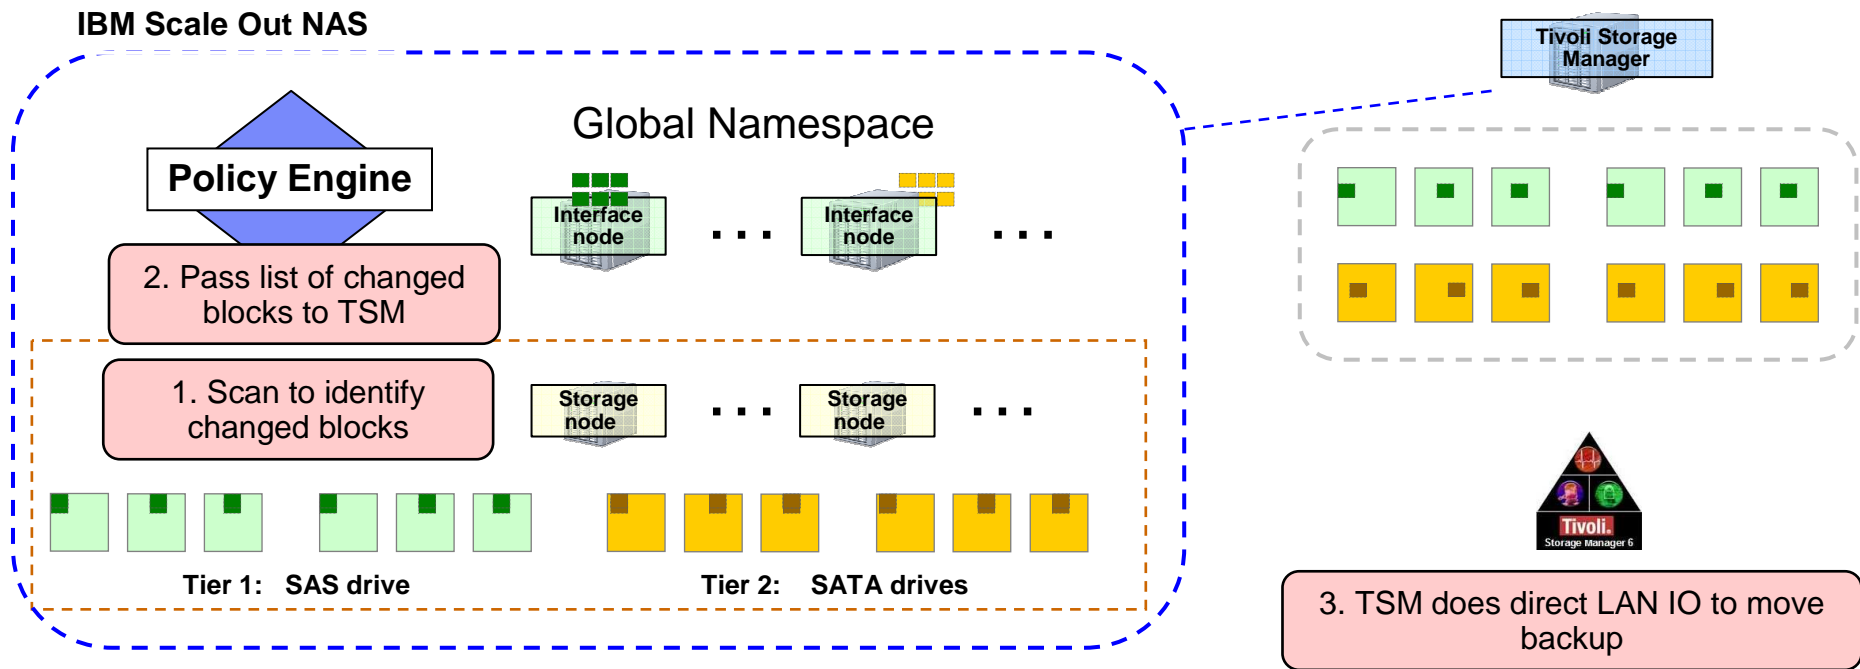
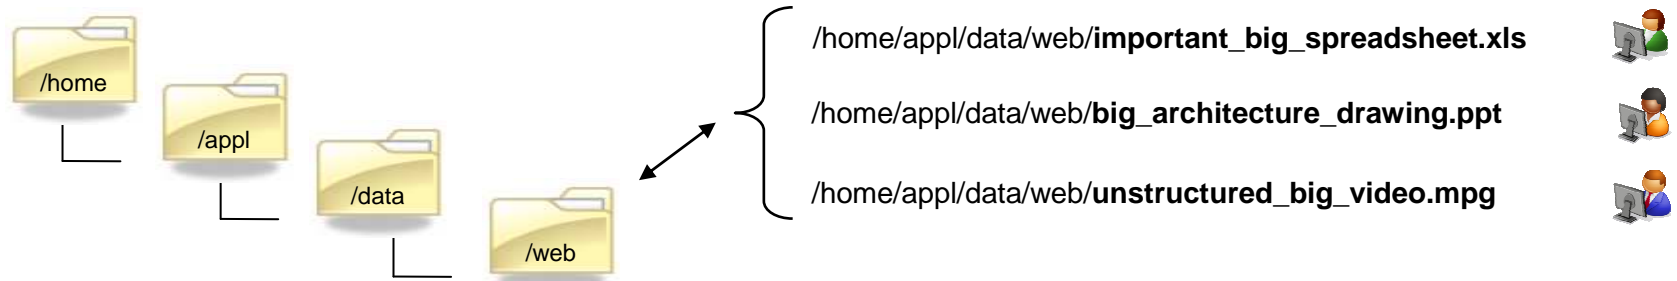
/home/appl/data/web/unstructured\_big\_video.mpg

**Note:** all three files,  
no change to logical  
directory

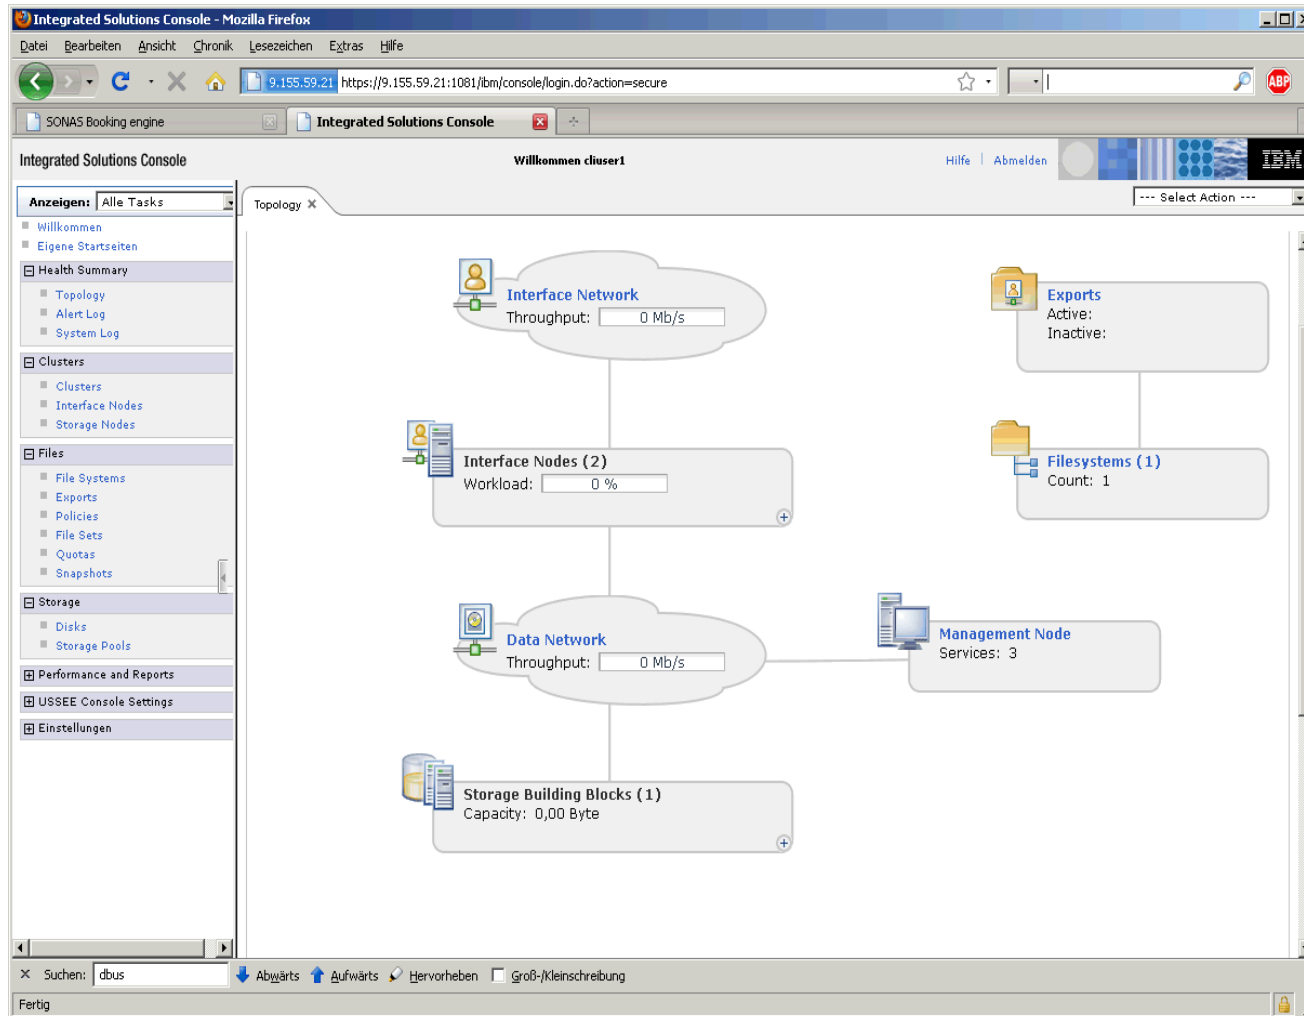
All nodes (both  
storage and  
interface) can  
participate in *parallel*  
data movement



# Backup / Restore mit Tivoli Storage Manager



# SONAS - Integrierte Oberfläche



## Einfache Administration

- Eine Oberfläche
- Alle Komponenten
- Alle Funktionen
  
- Automatisierbar
- CLI fähig
  
- Integrierbar in LDAP/AD

# Cloud-onomics...

## CLOUD COMPUTING



...leverages virtualization, standardization and automation to free up operational budget for new investment





# TAPE LTO5



## LTO 5

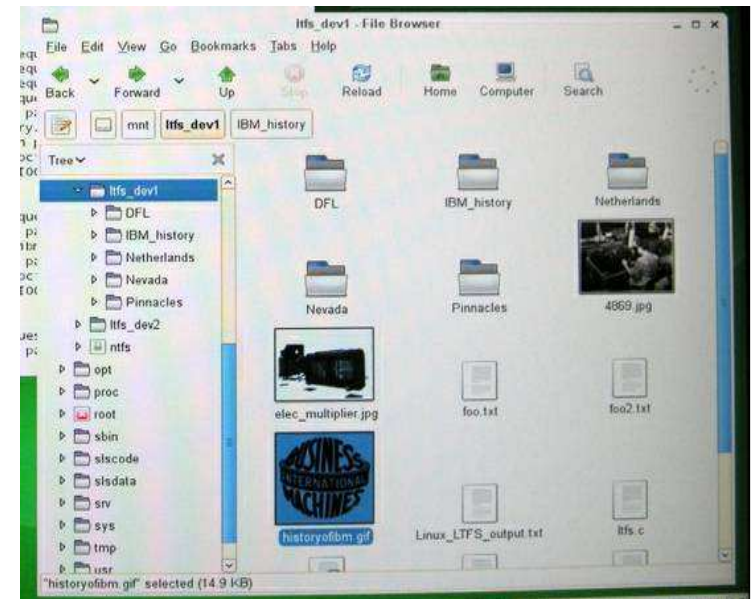
- 1.5TB native Kassettenkapazität
- 3TB compressed (2:1)
- 140MB/s Datenrate
- 6Gbps SAS oder 8Gbps FC
- Drive Encryption
- IBM LTFS (Long Term File System)
- Media Partitioning
  
- TS1050 für Libraries TS3100, TS3200, TS350
- Preview LTO 5 für TS3310
- Standalone Drives TS2250, TS2350



**LTO5: 17% schneller und 88% mehr Daten  
auf einer Kassette im Vergleich zu LTO4!**

## Media Partitioning und IBM Long Term File System

- **Media Partitioning**
  - Cartridge hat zwei Media Partitionen für schnelleren Zugriff
  - Eine Partition speichert Content, die andere den Index
- **IBM Long Term File System**
  - Filesystem Zugriff auf Betriebssystem Ebene
  - User sieht und hat Zugriff auf Daten des Tapes-Content Darstellung wie bei einer Disk!
  - Initial support on Linux and Apple
- Supported by External drives only (not in libraries)



Tape die bessere Disk ? :-)

## IBM Storage on Twitter

- Twitter Name: **ibmstorage** - <http://twitter.com/ibmstorage>
- Official Tweeter: **Raj Subramaniam is the official Community manager for storage. He does this role in addition to his web technical lead role. Raj was key in setting up IBM developerworks forum, blogs and spaces for IBM Storage brand. Raj has a clear purpose for IBM's tweets, which includes, " promoting and hashtagging events (e.g. #ibmpulse), promoting storage solutions, news distribution and engaging stakeholders (where appropriate)."**
- Stats: **137 following/140 followers**
- Why Storage tweets: **"As Twitter adoption rose among customers and partners, we wanted to engage them, keep them up to date, and promote our brand."**



Thank You.

