

Storage Efficiency 2013



IBM Smarter Storage for Smarter Computing

Dr. Roland Wolf
IBM Technical Sales Support
rolwolf@de.ibm.com



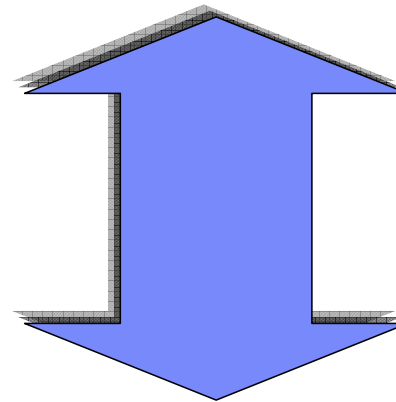
Storage Effizienz hat viele Gesichter

- Schnellster „Flash“-Speicher
 - Beschleunigt die Datenverarbeitung
 - Schnellere Business-Entscheidungen
- Compression
 - Mehr Daten pro GB speichern
- Tiering
 - Daten dort halten, wo sie nach Zugriffsverhalten hin gehören
 - Schnelle vs. Hochkapazitiver Speicher
- Virtualisierung
 - Flexibilität
 - Einfaches Management
- Daten-Cloud
 - „Selbstbedienung“
- Datensicherheit
 - Datenspiegelung
 - Backup Verfahren, FlashCopy



z/VSE
z/VM
z/Linux

FICON



DS8870

mit BC Modell

Fibre Channel



SVC

XIV




V7000 / V3700

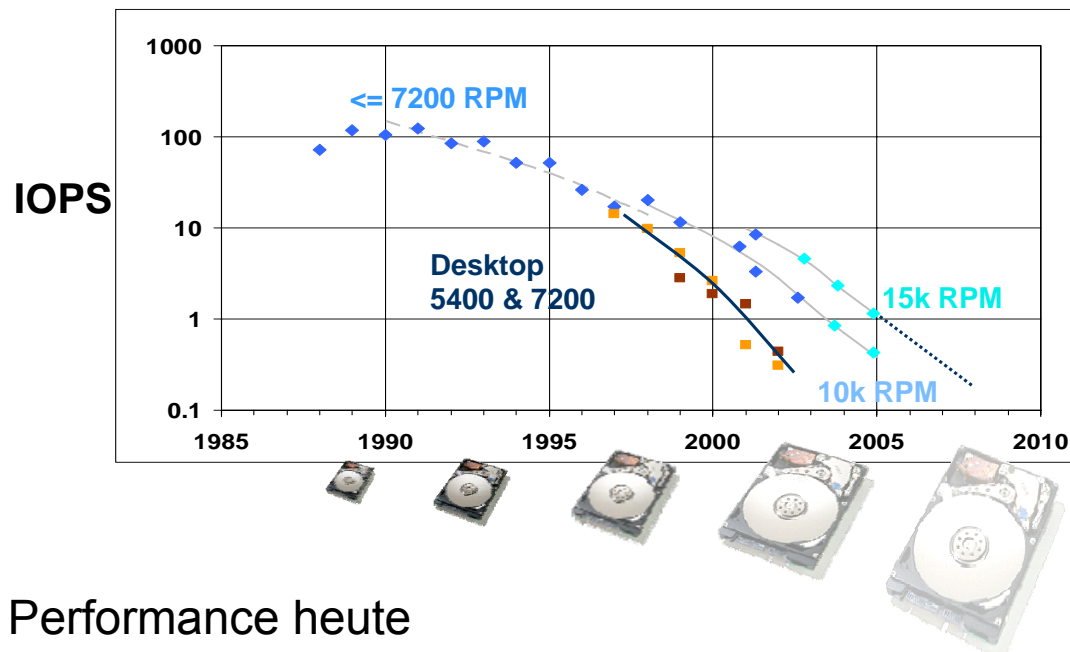


RamSan

Hard Disk



- Performance Entwicklung



IO Rate per GB und
Hard Disk Drive Generation

Seit 1990 um **Faktor 100**
schlechter

- Performance heute

* 100% Random

	RPM	Latency (ms)	Seek (ms)	IO/s *
FC/SAS	15000	2	4	167
FC/SAS	10000	3	5	125
SATA/SAS	7200	4.2	9	76

Flash Technologie

Warten Sie noch auf die
25K RPM Disk?




Drehst du noch
oder „blitzt“ du schon?

- Performance-Lücke steigt dramatisch
- Anwendungen werden zunehmend behindert
- SSD als Technologie zur **Performance-Steigerung**

- Bessere Verfahren zur Identifizierung der inaktiven Daten
- SSD als **Technologie zur Kostensenkung**

Flash Storage Technologie


Flash Preisentwicklung per GB



IBM Disk- und Virtualisierungssysteme



Pure-Familie
V7000 ITE



TMS-RamSan



XIV



DS8870


FICON



SVC



V7000
V7000Unified




V3700

Alle IBM Storage-Systeme sind „Flash-enabled“!


I/O Serviced by Disk

- I/O request ~ 100 μ s, Wait for I/O to be serviced ~ 5,000 μ s, Process I/O ~ 100 μ s
- CPU Utilization = Processing time / (Wait time+ Processing time) = 200 / 5,200 = ~4%



I/O Serviced by RamSan Flash

- Issue I/O request ~ 100 μ s, Wait for I/O to be serviced ~ 200 μ s, Process I/O ~ 100 μ s
- CPU Utilization = Processing time / (Wait time+ Processing time) = 200 / 400 = ~50%



5 reasons why Flash Storage is a “must” in every data center

What if you could...

1. Reduce DB and “per core” licensing
2. Improve application efficiency
3. Improve server efficiency
4. Increase storage operations efficiency
- 5. Lowest Latency!**




Application Sweet Spots: *Do More, Do it Faster!*



OLTP Databases

- Financial, gaming, real-time billing, trading, real-time monitoring, query acceleration (DB2/Oracle), etc.




Analytical applications (OLAP)

- Business intelligence, batch processing, ERP systems, reporting, massive data feeds, etc.



Virtual Infrastructures

- VDI, Consolidated virtual infrastructures, user profiles, etc.



HPC/Computational Applications

- Simulation, modeling, rendering, FS metadata, scratch space, video on demand, thread efficiency, etc.



Cloud-scale Infrastructures

- On-demand computing, content distribution, web, caching, metadata, GPFS, active file management, etc.

Financial

Government

E-Commerce

HPC

Telecom

SSD Hersteller



mehr
als 200
Anbieter



Western Digital Solid State Storage



Warum TMS?



**Solid State Storage
Leader**

- **Höchste Performance, höchste Reliabilität**, geringste Latency, geringster Energie Verbrauch in der Branche
Systeme sind 200% schneller als andere

**Deep Domain
Expertise**

- **34 Jahre** Erfahrung im Design von SSDs
30+ Patente
20 Generationen von Speicher Systemen

34
years

**Global Enterprise
Customers**

- Wachsende Kunden Basis in über 34 Ländern


**Strong Financial
Performance**

- Kein Venture Capital / Keine langfristigen Schulden

**World Class
Team**

- **Über 400 Mann-Jahre** Erfahrungen mit SSD

Hardware Accelerated with Extreme Parallelism





- No CPU in data-path but FPGAs
- Hardware instead of “Lines of Code”
- The Lookup Tables and Write Buffer is RAM
- Extreme parallelism
- A RamSan-810 has 40 controllers, so it can do 1440 4KB operations simultaneously

“You cannot increase performance by adding lines of code.”

*DMA = Direct Memory Access

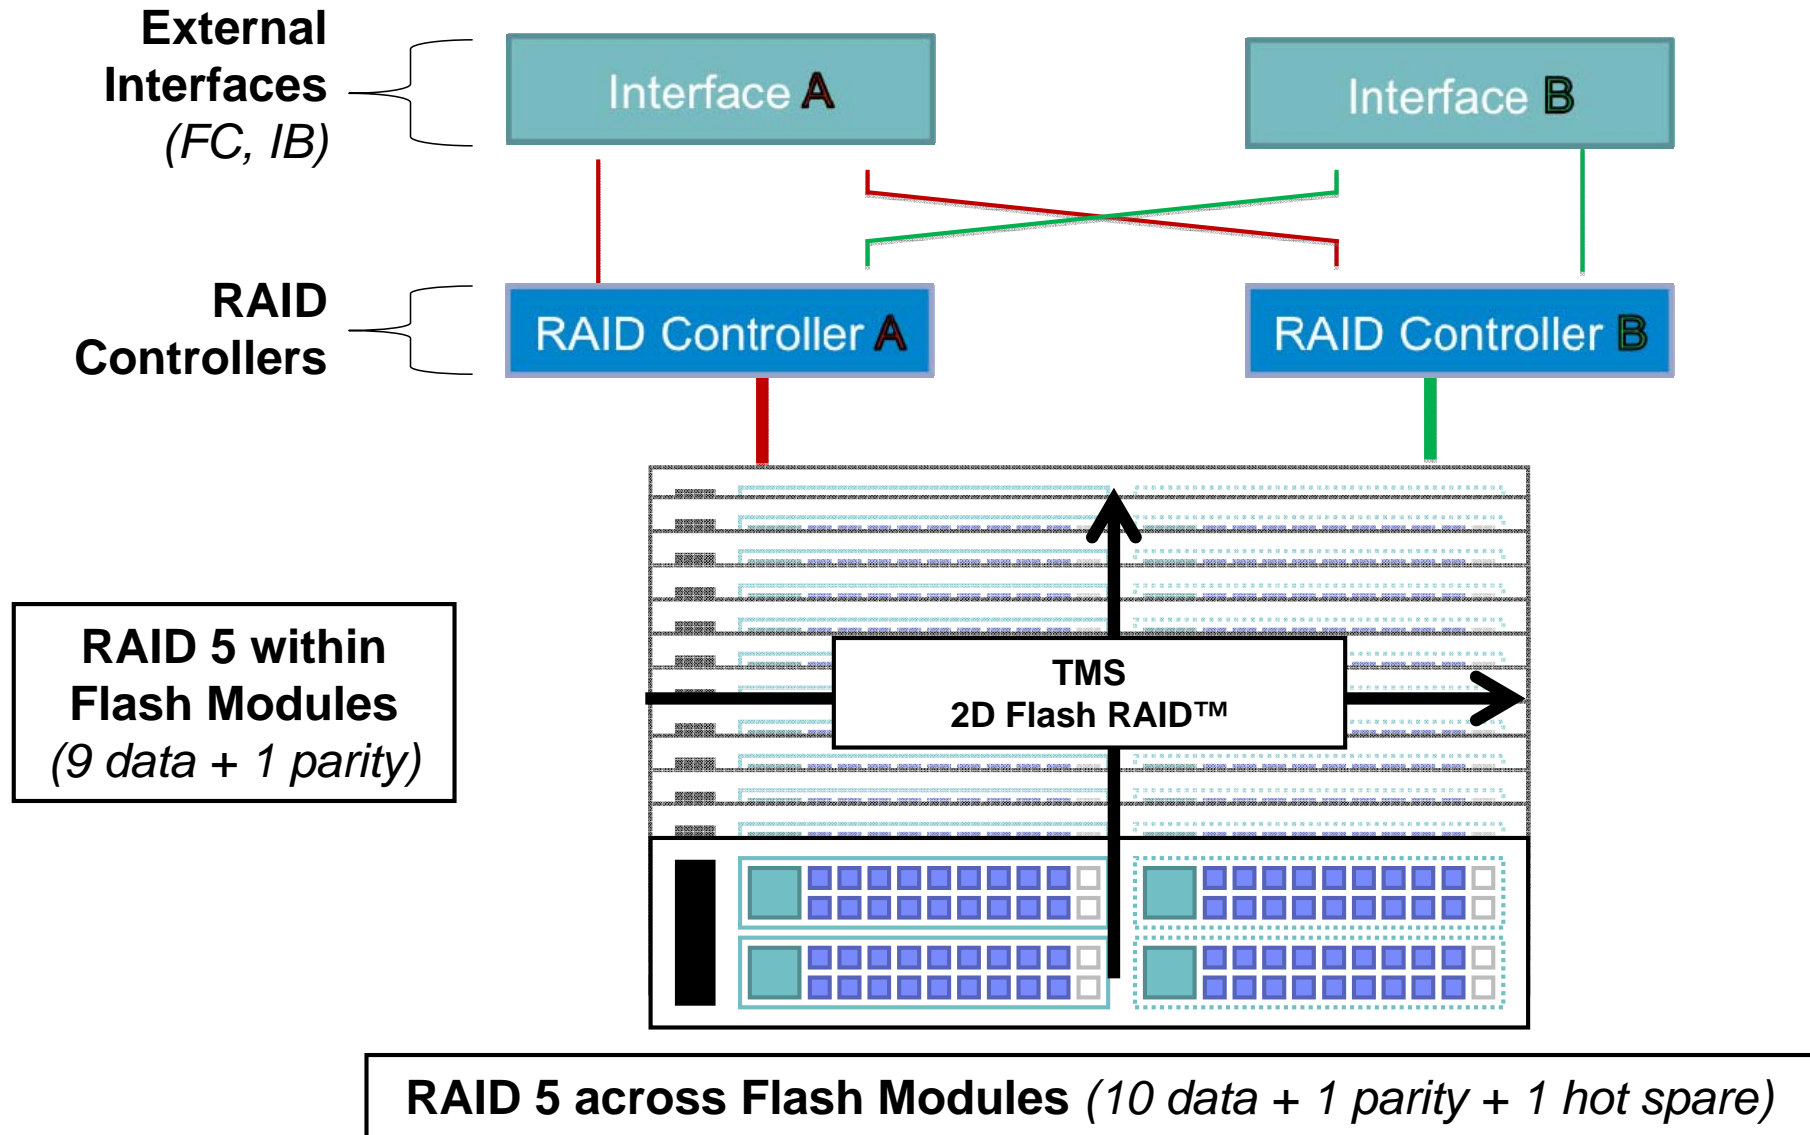
**FPGA = Field Programmable Gate Array

RamSan Choices



RamSan-70	RamSan-710/810	RamSan-720/820
SLC Flash	SLC (7) / eMLC (8) Flash	SLC (7) / eMLC (8) Flash
90/30 us R/W Latency	100/25 us R/W Latency	100/25 us R/W Latency
450 or 900 GB	1-5 or 2-10 TB	(5/10) or (10/20) TB w/HA
1.5M IOPS (512B)	450K/400K IOPS (4K)	500K/450K IOPS (4K)
2.5 GB/s	5/4 GB/s	5/4 GB/s
Full-height, half-length PCIe x8 2.0	1U rackmount, 4x 8Gb FC ports, 4x 40Gb QDR InfiniBand	1U rackmount, 4x 8Gb FC ports, 4x 40Gb QDR InfiniBand

2D Flash RAID™ (RamSan 720 / 820)



Performance Scenario: Oracle RAC, 4 Nodes

ORACLE

Enterprise Array, No Flash

2 million queries
12.25 minutes to complete

16K Total IOPS
4K per RAC Node

```
[oracle]$ time ./spawn_50.sh
```

```
real    12m15.434s
user    0m5.464s
sys     0m4.031s
```

RamSan

2 million queries
1.3 minutes to complete

160K Total IOPS
40K per RAC Node

```
[oracle]$ time ./spawn_50.sh
```

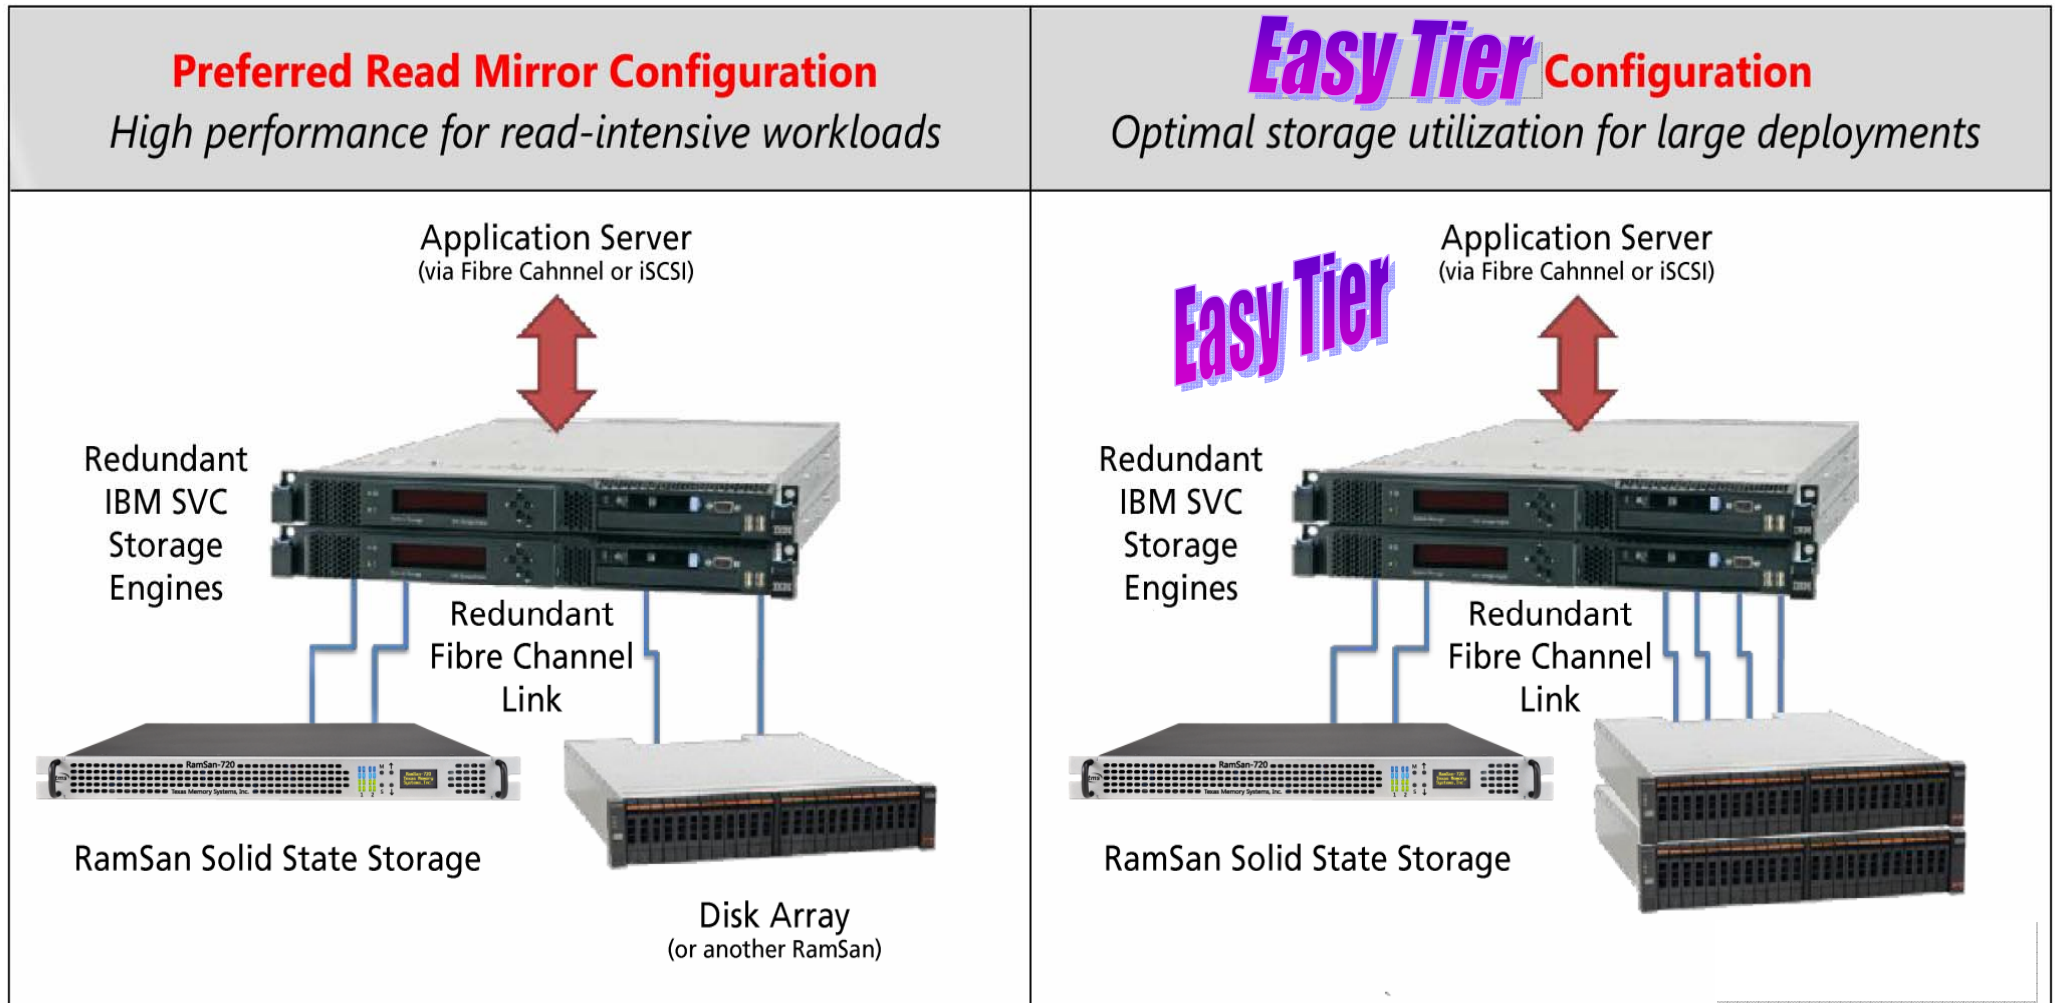
```
real    1m19.838s
user    0m4.439s
sys     0m3.215s
```

A factor of about 10x improvement!

SVC and RamSan – Combines Performance with Functionality

Basic RamSan SVC Configurations


In most cases, a basic RamSan SVC configuration includes one RamSan and one or more disk arrays connected via Fibre Channel to two IBM SVC Storage Engines, which use Fibre Channel to connect to application servers.



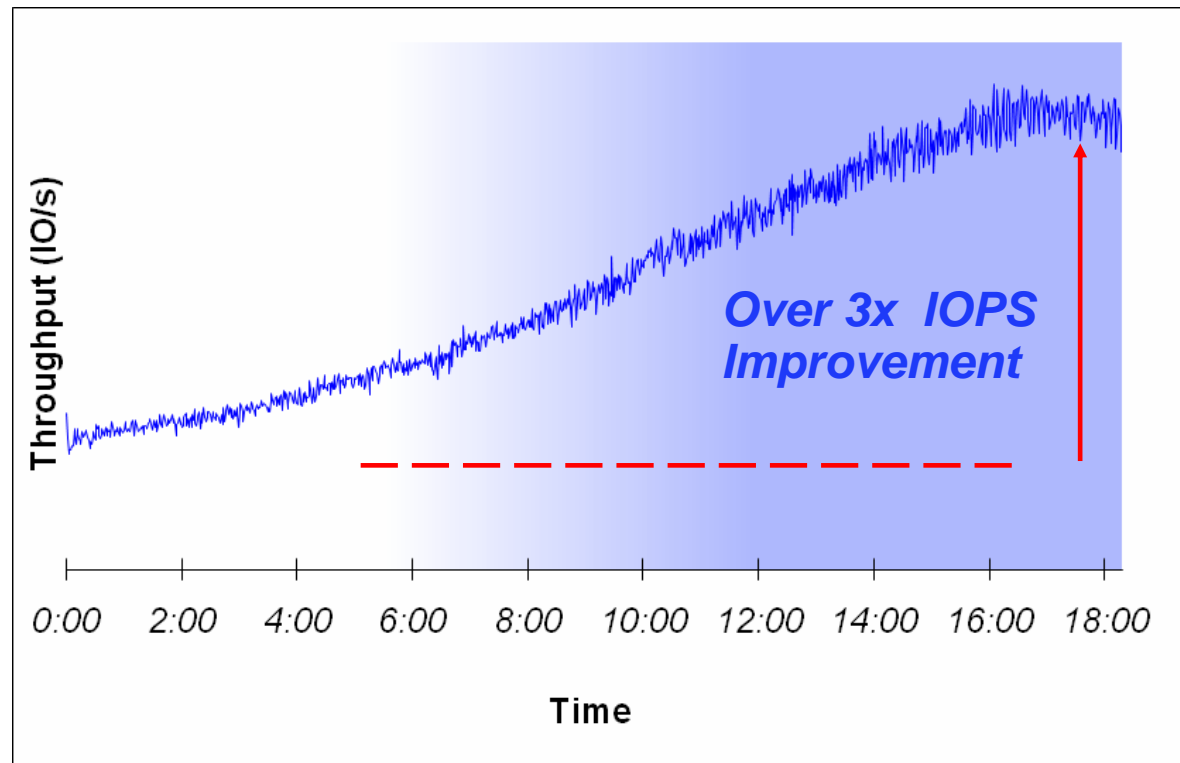
Storage Easy Tier technology

System learns workload, does automated relocation, removes user complexity

*Easy Tier for DS8000,
SVC V7000*




System configuration:
2.3 TB SSD + 96 TB SATA



Source: Storage Performance Council, April 2010: http://www.storageperformance.org/results/benchmark_results_spc1#a00092

Easy Tier provides better performance in **1/2** the floor space & with **40%** less energy consumed
96 SATA HDDs + 16 SSDs (5.7 KW) vs. 192 FC HDDs (9.5 KW)

GB Kunde aus Stuttgart




- SAP, VMware
- 2x 200TB

- IBM SVC
- IBM Disk, EMC
- RamSan (2x)

- Tiering

Batch-Jobs (SAP):
von 7h auf 2min!

Einsatz von SSD/Flash im Rechenzentrum – heute und morgen



Storage Performance Council (SPC) results

SPC Benchmark 1 (SPC-1) results page – http://www.storageperformance.org/results/benchmark_results_spc1

SPC Benchmark 2 (SPC-2) results page – http://www.storageperformance.org/results/benchmark_results_spc2



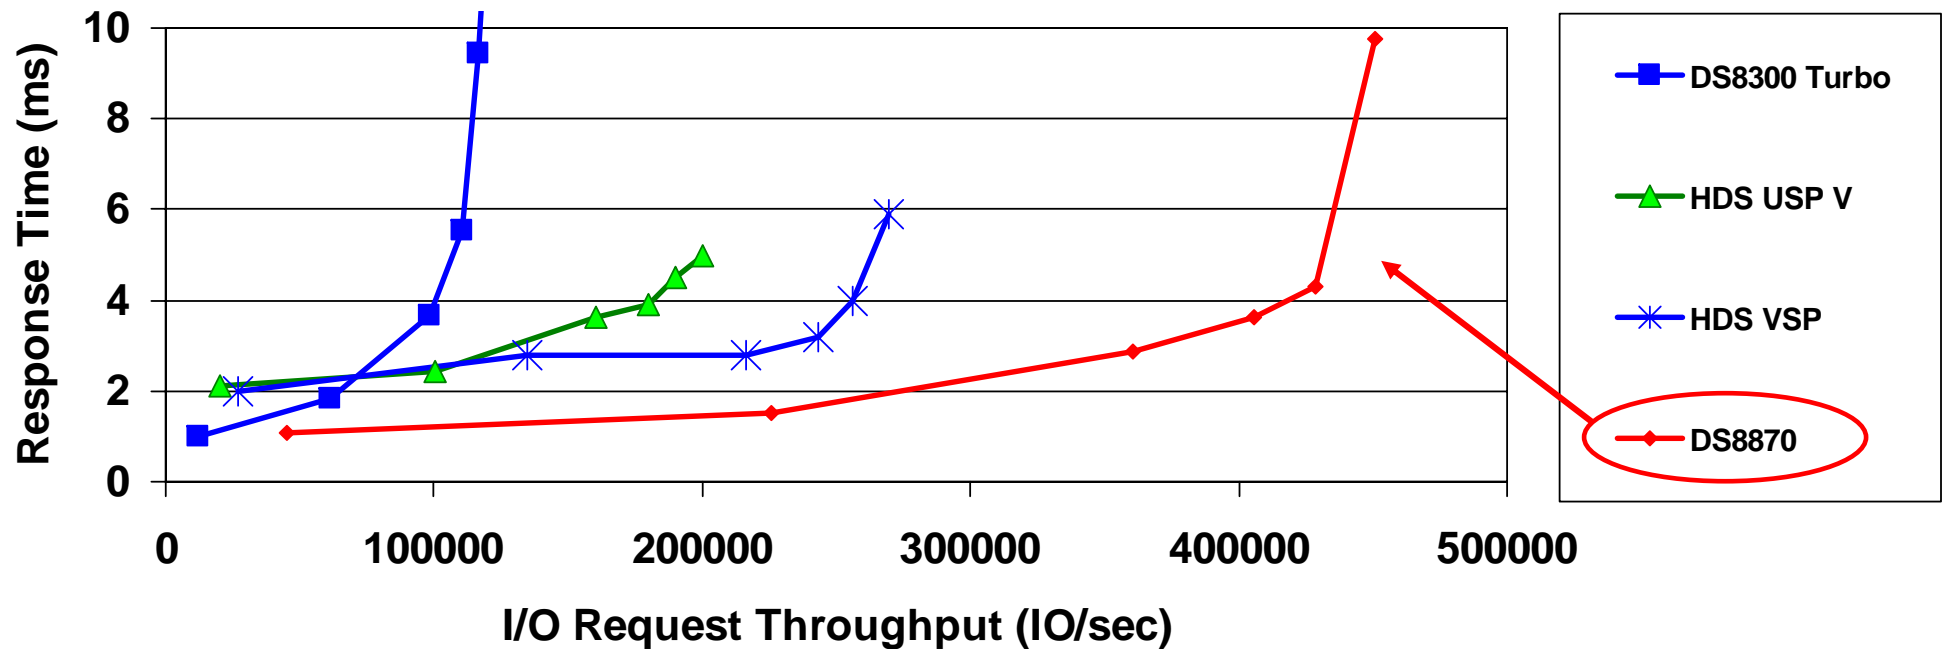
- SPC-1 throughput 451.082 IOPS
 - #1 result for, enterprise-class **all-HDD systems**
- SPC-2 throughput 15.424 MB/s
 - #1 result
- Internal benchmarks:

	DS8300 (R4.3)	DS8700 (R5)	DS8800 (R6)	DS8870 (R7)	<i>Increase</i>
	P5+ 4-way	P6 4-way	P6+ 4-way	P7 16-way	<i>vs. DS8800</i>
Seq. Read (GB/s)	3.9	9.7	11.8	21.0	1.8x
Seq. Write (GB/s)	2.2	4.7	6.7	11.0	1.6x
DB z/OS (K IOPS)	165	201	204	640	3.1x
DB Open (K IOPS)	165	191	198	550	2.8x

SPC-1 comparison vs. competition^{*)}

Higher throughput with 451K IOPS and lower response time!

SPC-1 Published Results



^{*)} Source: Published SPC-1 results: http://www.storageperformance.org/results/benchmark_results_spc1

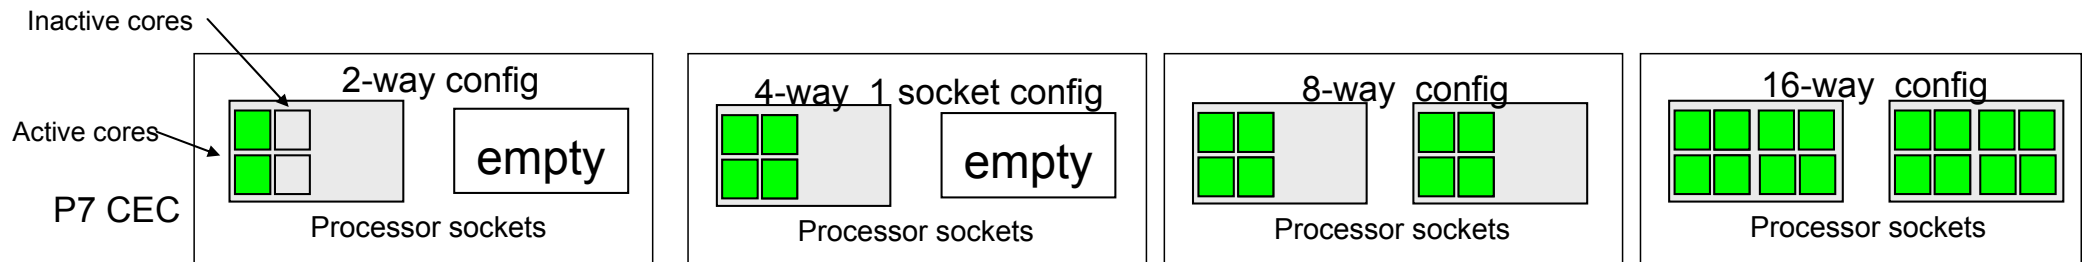
- DS8300 Turbo: 480 × 73 GB 15K HDDs, RAID-10, 256 GB Cache
- HDS USP V: 1024 × 146 GB 15K HDDs, RAID-1, 256 GB Cache
- HDS VSP: 1152 × 146 GB 15K HDDs, RAID-1, 512 GB Cache + 512 GB Cache Flash
- DS8870: 1536 × 146 GB 15K HDDs, RAID-10, 32 × 8Gb FCP, 1024 GB Cache

EMC refuses to publish results, so we can assume less-than-impressive performance

DS8870 - R7 configuration table

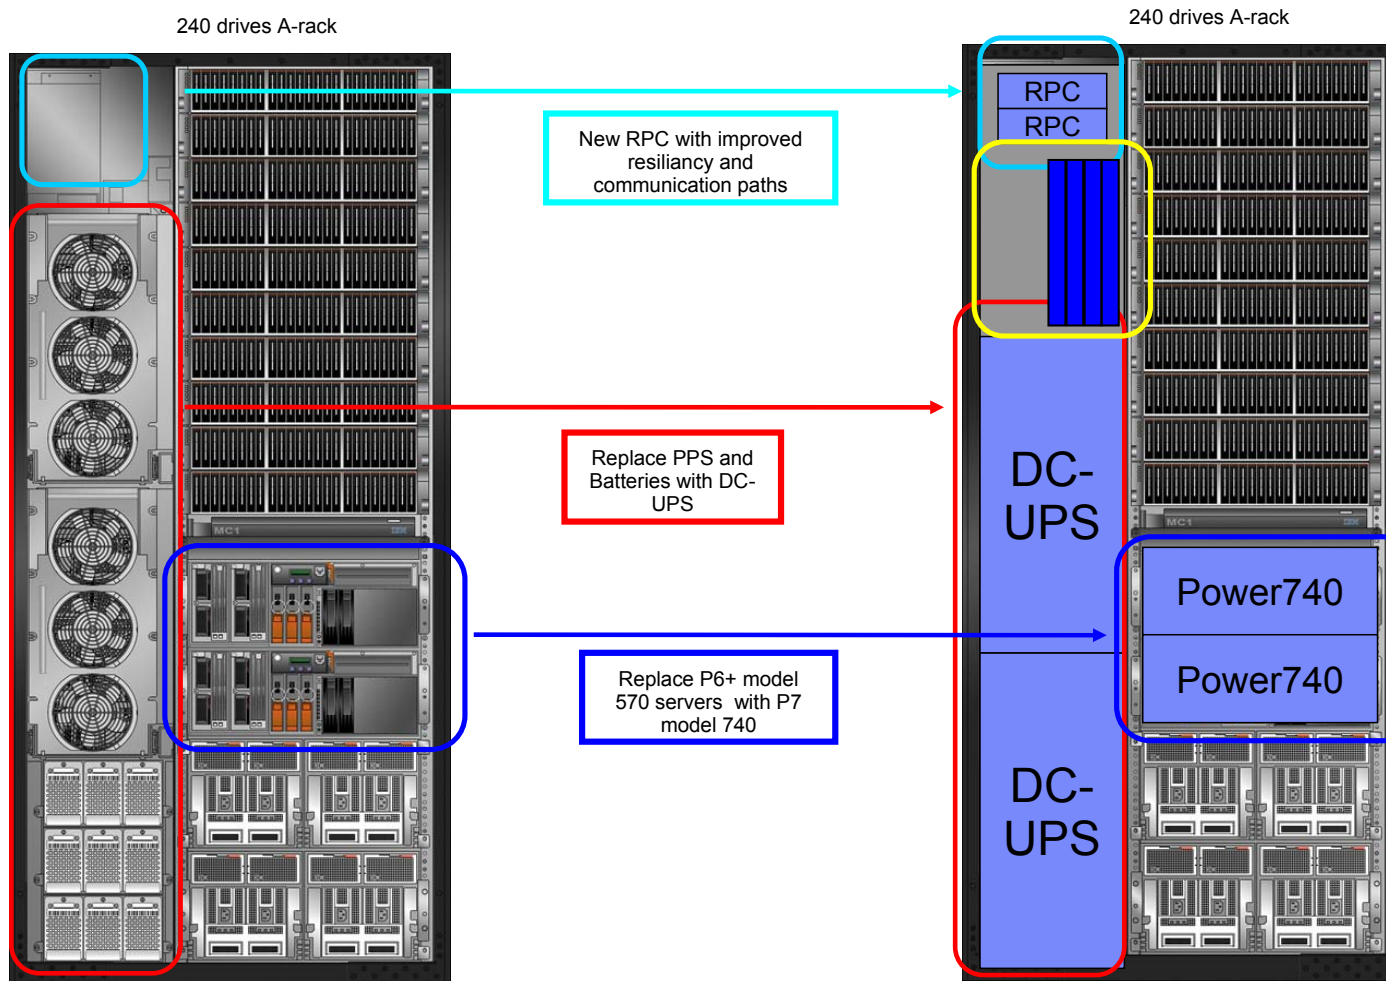
System Class	Active Processor Config per CEC	Processor Cards per CEC	Total System Memory Sizes (GB)	NVS Sizes (GB)	CEC DIMM Size (GB) / # of DIMM's per CEC / # of Memory Controlers	Expansion Frame	IOA's Max DA Pairs/Max HA's	Max DDM's
Business Class***	2- core	1x4 core	16	1	4(GB)x2x1	No	2/4(32 ports)	144
			32		4(GB)x4x1			
Enterprise Class	4-core	1x4 core	64	2	4(GB)x8x1	No	4/8(64 ports)	240
	8-core	2x4 core	128	4	4(GB)x16x2	0, 1 or 2	8/16(128 ports)	1056
			256	8	4(GB)x32x4	0, 1, 2 or 3	8/16(128 ports)	
	16-core	2x8 core	512	16	8(GB)x32x4	0, 1, 2 or 3	8/16(128 ports)	1536
			1024	32*	16(GB)*x32x4			

* - Business Class upgradeable to Enterprise Class, however requires additional FO/FB cycle to adjust NVS size



Rack changes DS8800 to DS8870

- System rack dimensions unchanged from previous models
- Location of drives, adapters, host cables carry forward.



All IO-bay slots can be filled with direct PCIe connections to the Ultra SSD Drawers, allowing for up to 8 per DS8870.

The new Power7 Controllers contain new enterprise class storage hardware:

- Faster PCIe Interface Cards to IO Bays
- Direct PCIe-PCIe connections between controllers
- Direct high speed connection to power system for better reliability and monitoring

DS8000 integration of EXP30 Ultra SSD storage

SoD: Annc. Letter 116-106

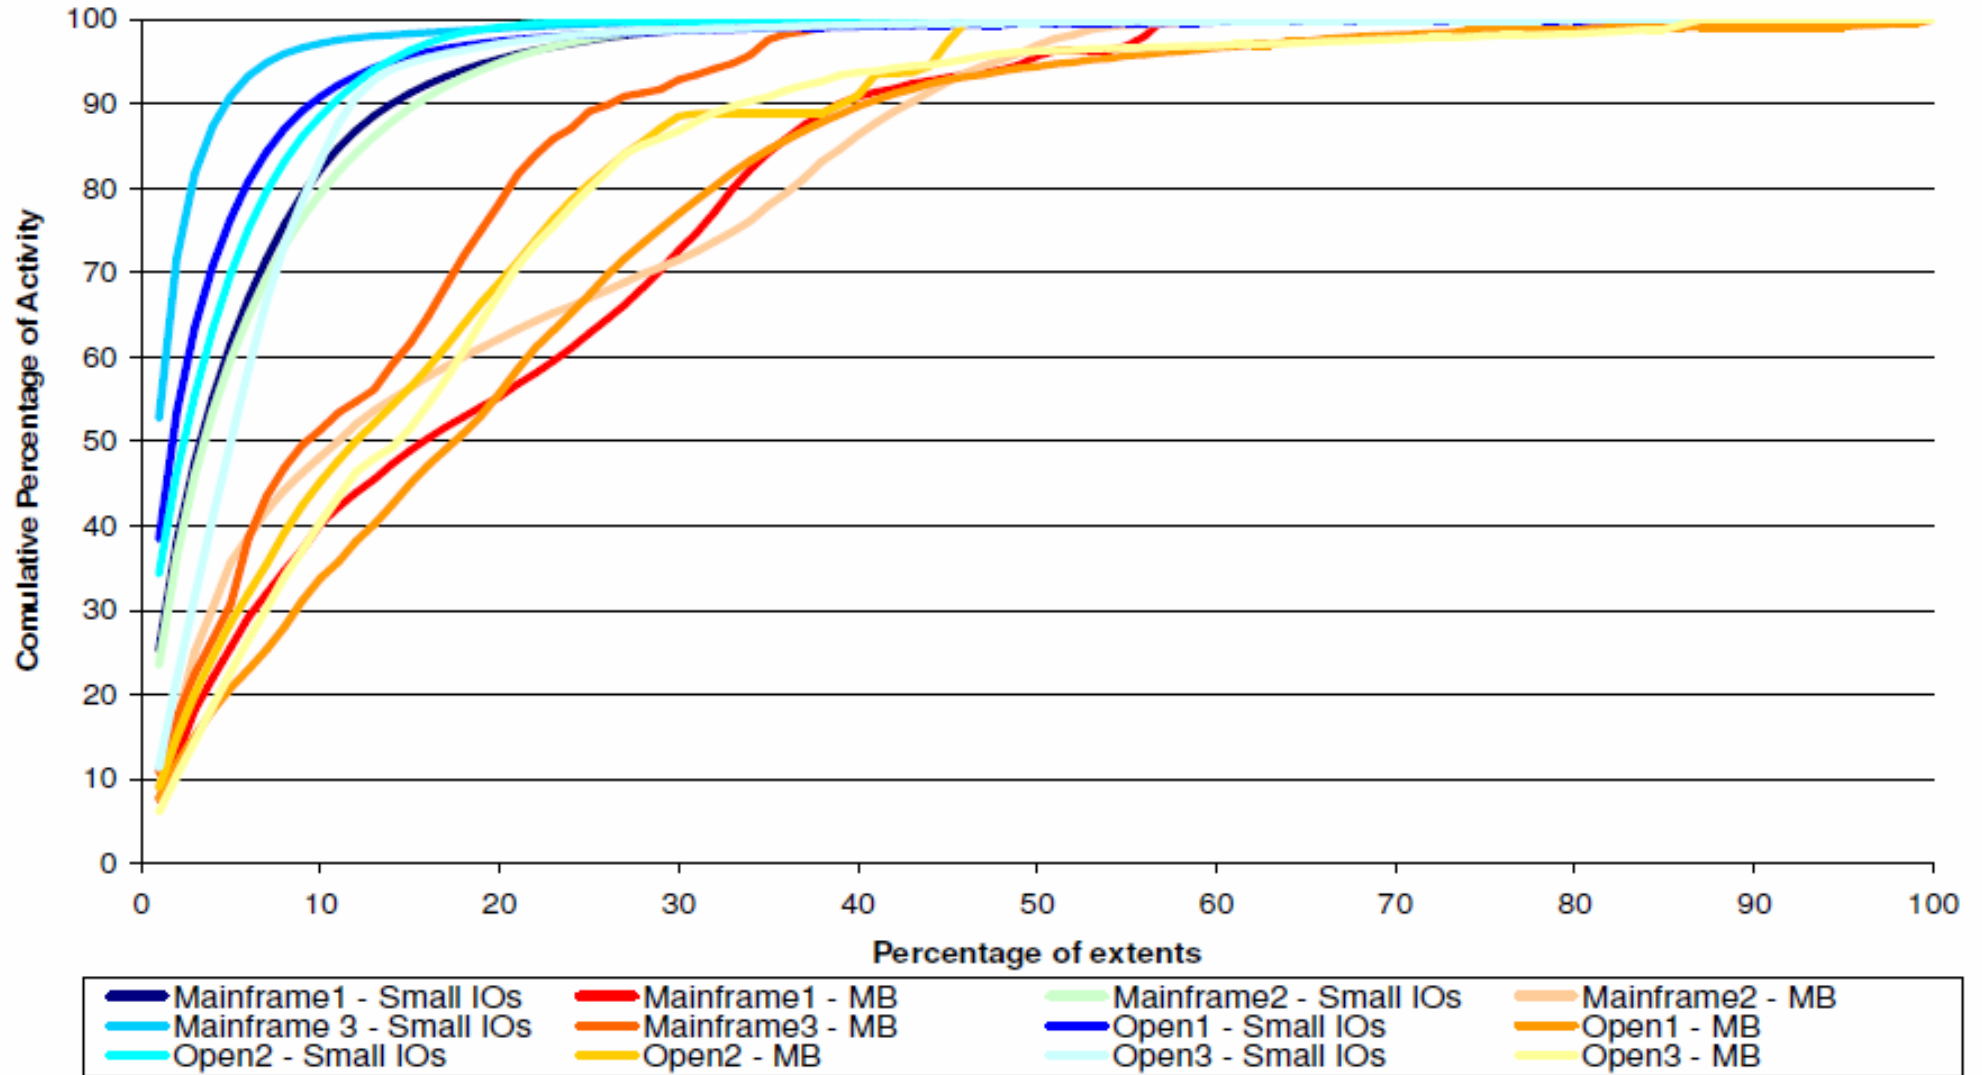
Coming



- 1U drawer . . . Up to 30 SSD
- Optimised for SSD performance
- Two large-cache SAS adapters especially designed for SSD
 - With 387 GB drives = up to 11.6 TB
 - Directly connected to DS8000 internal PCIe fabric
- Enterprise class RAS
 - High-Availability 2W Cluster
 - RAID 5 and 6 and mirrored/protected write cache
- Easy Tier Enabled

*Statements of IBM's future plans and direction are provided for informational purposes only. Plans and directions are subject to change without notice

Workload Skew for Different Environments



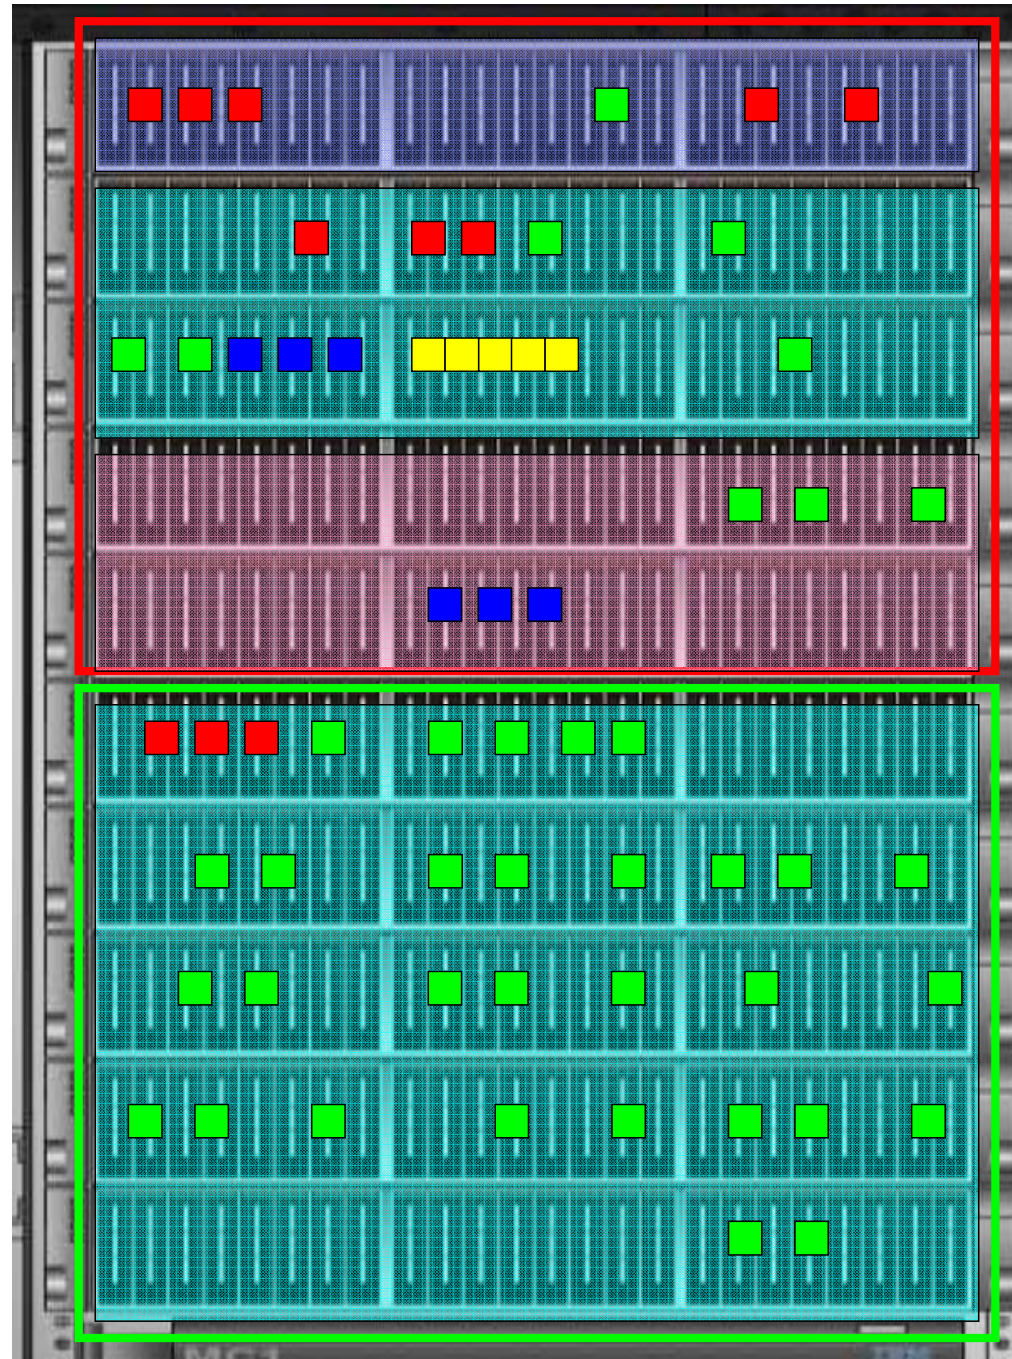
Easy Tiering

Extent Pool P0

Easy Tier across 3 Tiers

Extent Pool P1

Easy Tier within homogenous pool



SSD / Flash

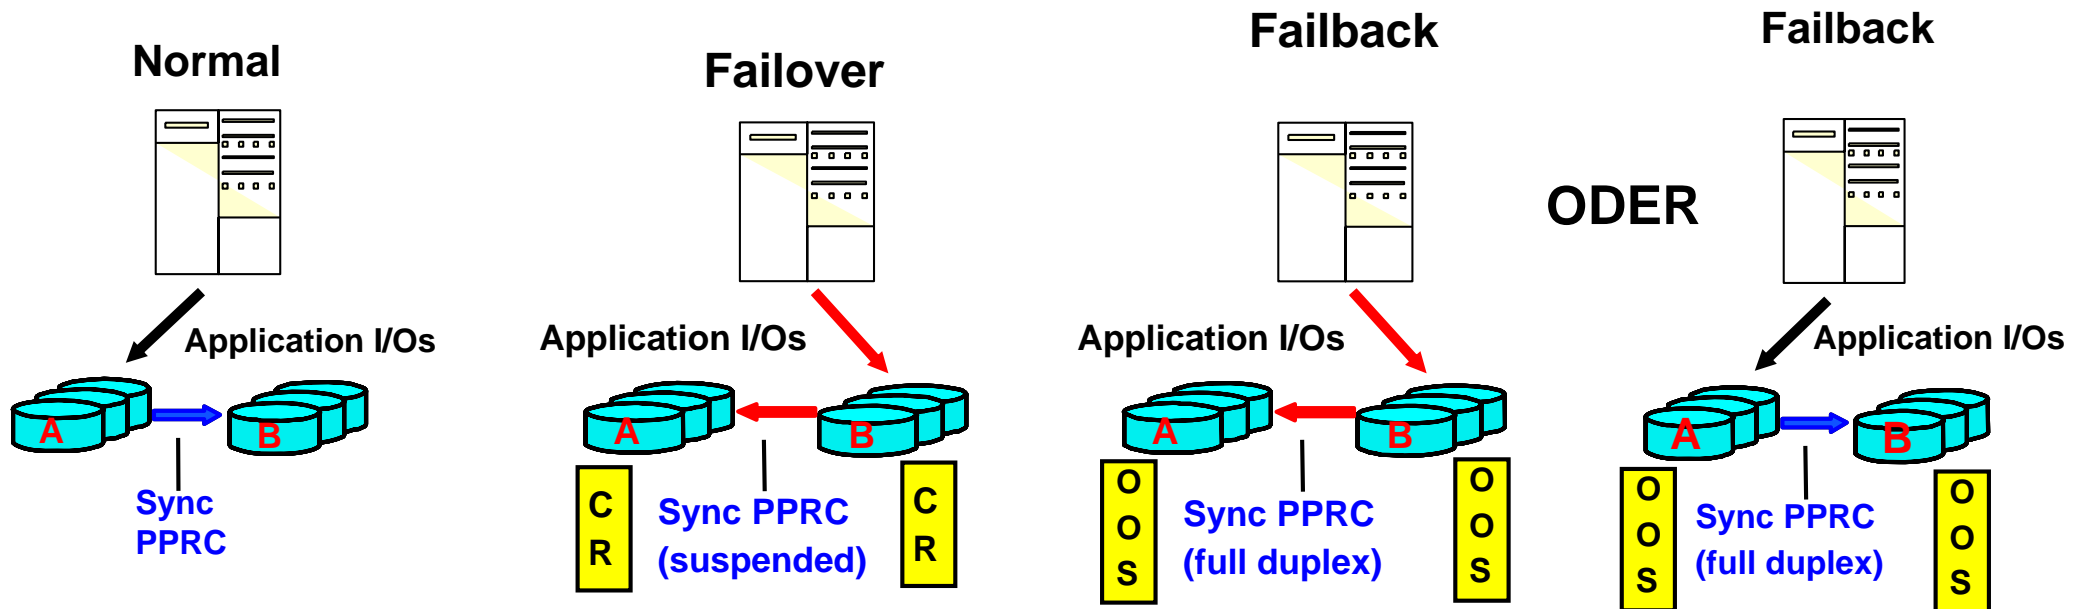
FC / SAS

Nearline

FC / SAS

Metro Mirror (PPRC) Failover / Failback

- Methode zur Umkehr der Spiegelrichtung und danach Wiederherstellung des ursprünglichen Zustandes
- Geeignet für Disaster Recovery Tests
- Die Mitführung von Change Bitmaps erlaubt eine sehr schnelle Resynchronisierung



**Faster Resynchronization
Less Resource Consuming**

IBM Disk- and virtualization Systems

- Virtualization as independent layer in a SAN – separation of HW and logic
- Standardized storage services
- Ease of administrations by virtualising all LUNs
- Increase performance by memory cache of SVC nodes
- Striping across many drives, arrays, subsystems
- Optimum usage / utilization of SAN storage (no boundaries)
- Storage Pools




SVC



V7000Unified



V7000



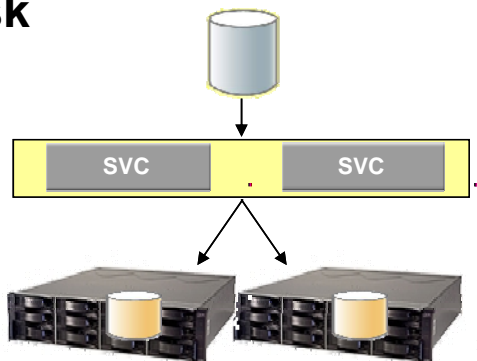
V3700

- EasyTier / automatic tiering
- High availability solution with SVC stretched clusters
- Remote Mirroring solutions independent of storage subsystem
- Online volume migration
- Snapshots for all managed disks
- Thin-Provisioning
- Dynamic volume expansion or shrink
- Quality of service with I/O governing
- RACE compression


SVC Funktion interner Datenspiegel (vdisk Mirror)

Erweiterung zur hochverfügbaren Stretched Cluster Lösung

Virtual Disk Mirror



Stretched Cluster Anwendung



Transparentes Verschieben
(Motion) auf allen Ebenen

Interner Datenspiegel

- Schützt vor Ausfall eines Speichersystems
- Transparenter Failover für Server
- Keine Lizenz erforderlich


No Single Point of Failure Konzept Stretched Cluster

- Transparenter Failover beim Ausfall einer Site
- Auch für zLinux und FC attached zVM, zVSE

IBM SVC und Storwize V7000 with Integrated Real-time Compression

What this is

- Easy-to-use **real-time compression** fully integrated into SVC/Storwize V7000 for active primary data such as production databases and e-mail applications
- Compression technology that operates immediately as data is written to disk so no space is wasted storing uncompressed data awaiting post-processing



Why It Matters

- Store **up to five times** more active data in the same physical disk space
- Compression helps **reduce storage purchase costs**, rack space, power and cooling, and software costs for additional functions
- Deferral of storage costs throughout the lifetime of the system
- Immediate extra free space on primary storage
- Compression can help freeze storage growth or delay need for additional purchases
- Savings on the back-end storage including Disaster Recovery, allowing for extra backup retention
- Transparency to users and support staff

SVC Compression ohne Komplexität und Performanceverlust!

- Lizenziert auf TB Basis (Größe der LUNs)
- Real-time Compression als Volume in die GUI integriert
- Compression der Daten beim Initial Write
 - Weniger Daten werden gespeichert
 - Weniger I/O's im Backend
- Keine Performanceverlust
- Einfache Konvertierung von bestehenden LUNs in Compressed-LUNs
 - Volume mirroring Basis Funktion
 - Auch externe Volumes können konvertiert werden
 - Kein "process scheduling" oder "post processing"!
 - Sizingtools ("Comprestimator")




Add Volume Copy - vdisk0 (mdg0)

Mirroring creates two identical copies of the volume and, if each copy belongs to a different storage pool, still allows data access even if the physical storage in one pool goes offline.

Volume Type

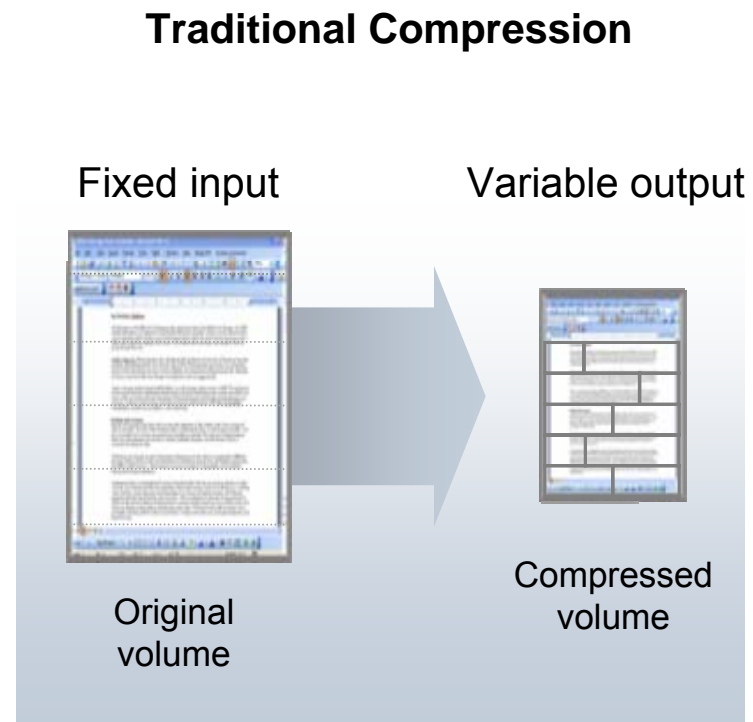
Generic
 Thin Provisioned
 Compressed



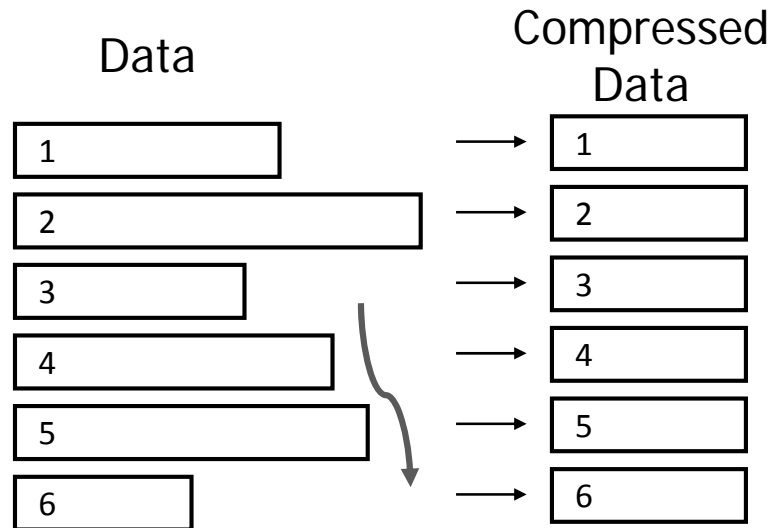
IBM Data-compression: einfach, schnell, effizient

Traditional Compression

- Traditional compression engines break volumes up into fixed-sized chunks and compress them
- Some or all data written is stored uncompressed and only compressed later
 - Reduces efficiency because additional storage needed for data waiting to be compressed
- Compression ratios depend on chunk size used
 - Big chunks can have poor I/O performance
 - Small chunks offer lower compression ratios
 - Compression typically degenerates over time creating fragmentation requiring more garbage collection
 - Poorer I/O performance




Variable Input Fixed Output



- RACE flips this approach, taking a variable data stream size and producing “fixed” output units
 - Compressed volumes have a consistent layout
 - Temporal locality: data that’s accessed together is compressed together
 - Variable sized input chunks get better compression
 - Requires fewer disk I/Os
 - Delivers better performance

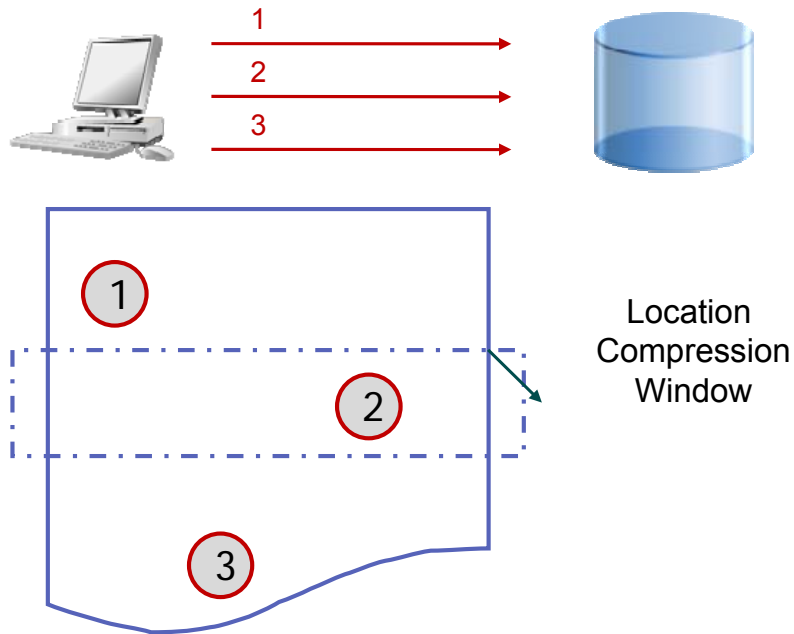
Compressed Data



- No Fragmentation
- Consistent performance over time
- Consistent compression ratio over time

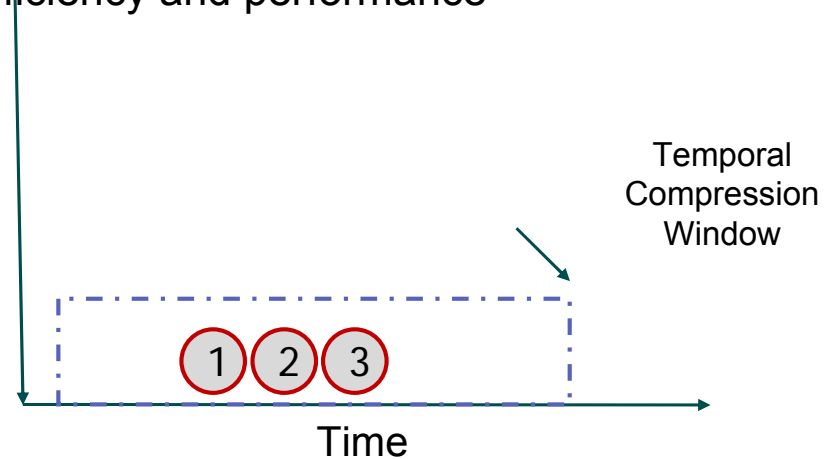
Compression Without Compromise

Designed for Real-Time Random Access to Active Data



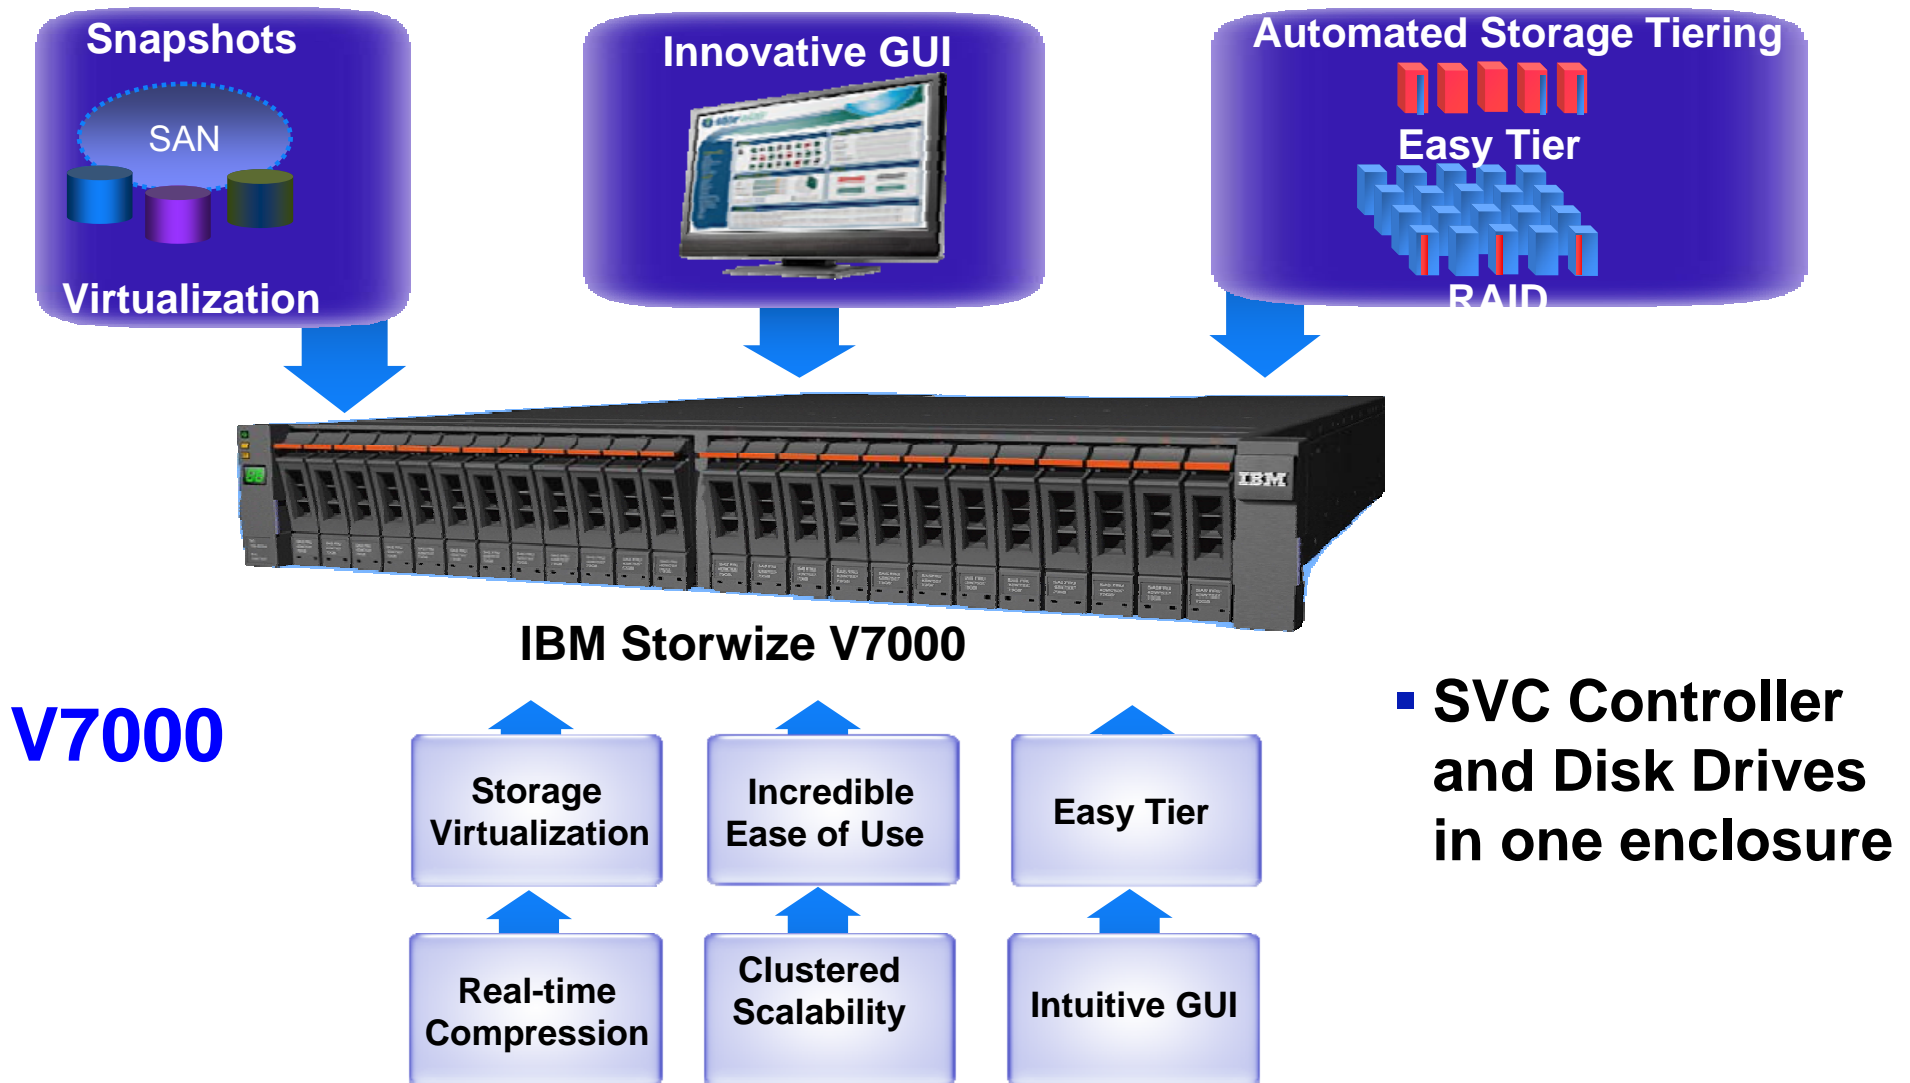
- Applications make multiple updates to data
- Traditional compression uses fixed-sized chunks and compresses each update based on its physical location on a volume
- In this example, **three** separate compression actions

- RACE compression acts on data that is written around the same time (“temporal locality”) not according to physical location
- Temporal locality is more related to real system operation
 - Applications may make related updates to different parts of a volume
- RACE takes advantage of the structure of the data and its relations
- In this example, only **one** compression action
- Better compression and decompression efficiency and performance



Technology Integration Continues, Storage Efficiency Continues

...IBM brings you IBM Storwize V7000 Storage System with integrated Real-time Compression




IBM Storwize V7000 Clustered System

Clients can **Scale UP** and/or **Scale OUT** as business needs demand

- **SCALE UP** to add capacity
- **SCALE OUT** to add performance with up to 960 disk drives
- Cluster four systems to provide smarter scalability
- Buy only as much capability as you need, when you need it
- Increase performance, connectivity and cache
- Clustered system operates and is managed as a single system

Start Small

One 24 bay control enclosure




Cluster
Up to
4X
V7000

Note: IBM Storwize V7000 Unified does not currently support clustered systems. A maximum of 240 drives are currently supported.

Das “Smarter Storage” Konzept: V3700, der kleine Bruder der V7000

IBM bringt das Beste aus den IBM Speichertechnologien in ein Entry System



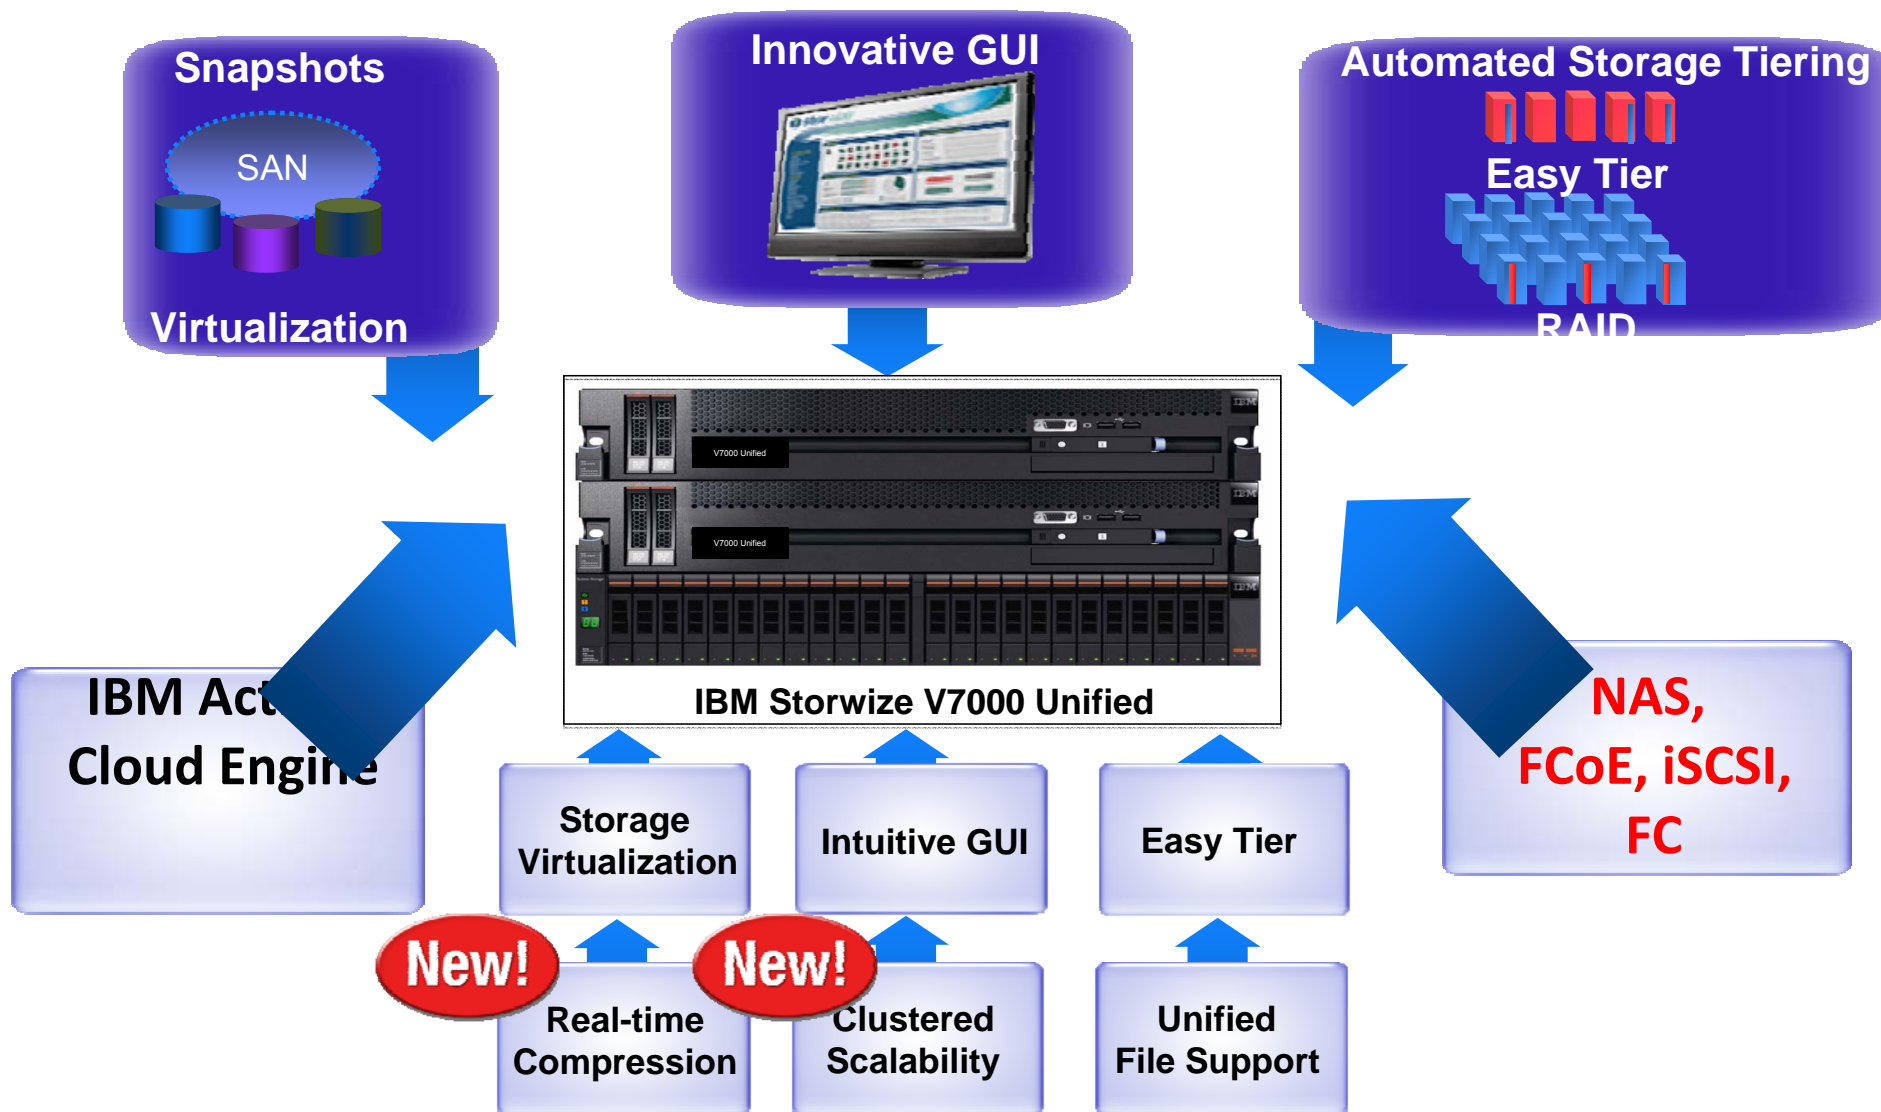
**Platzsparende
Daten-
sicherung**

**Einfache
Bedienbarkeit
verbessert
Produktivität**

**Easy Tier
verbessert
Antwortzeiten**

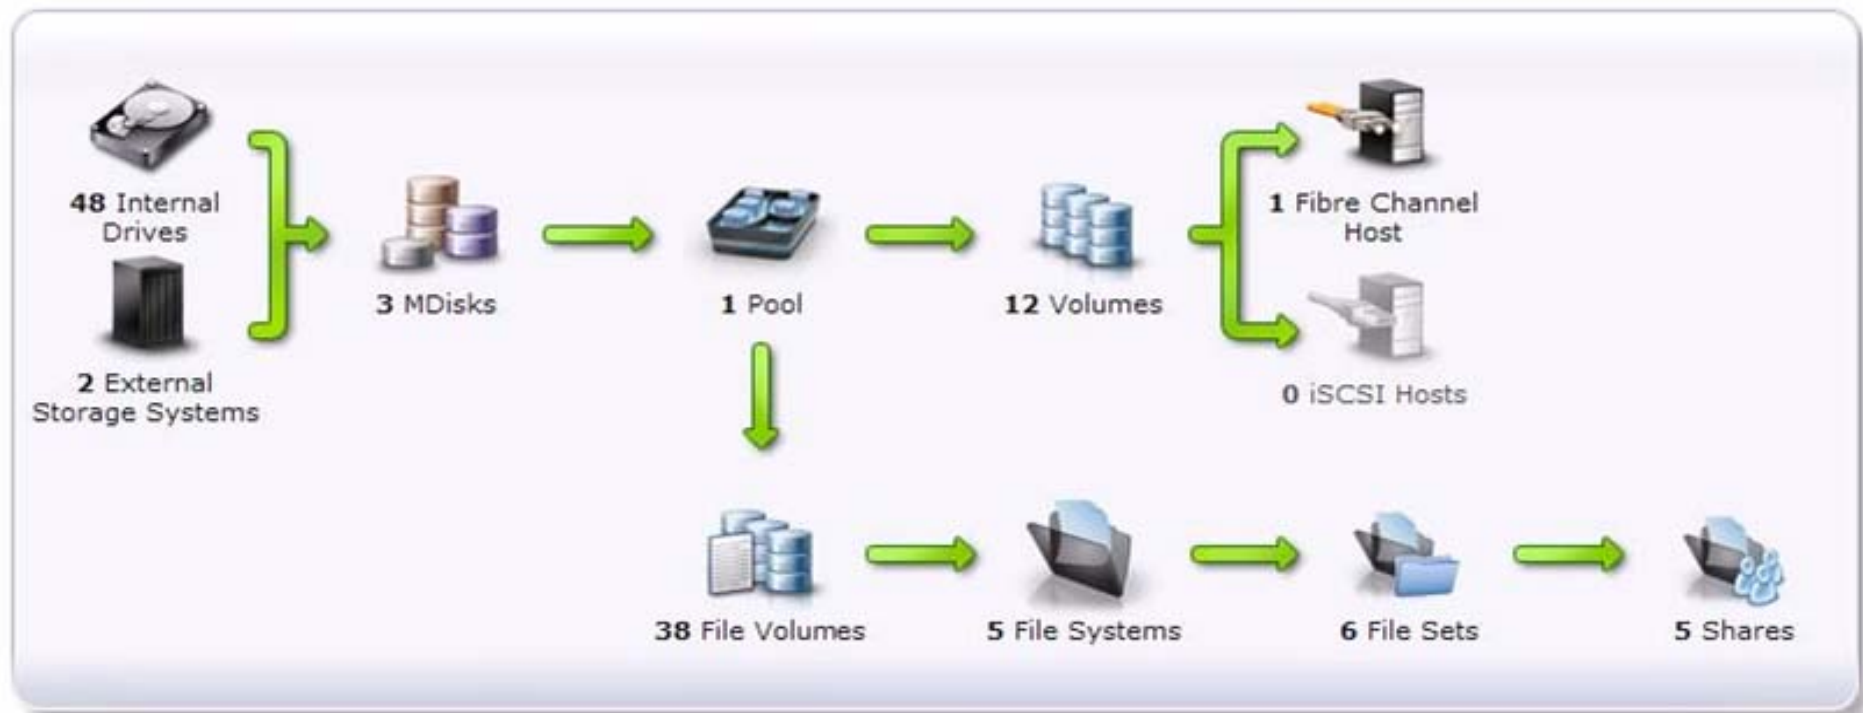
V7000 Unified, Storage Efficiency Continues

...IBM brings you IBM Storwize V7000 Unified Storage System with integrated Real-time Compression and 4-Way Clustering



Intuitive GUI is Truly Integrated

- One administration interface for block and file data
 - Not a launcher for two different interfaces
- Integrated data protection
 - Built-in NDMP and IBM Tivoli Storage Manager client



To deliver on its promised value, cloud storage must provide a few key capabilities.



- **Self-Service** ability to manage their own IT environment in their own way



- **Quick Access** getting the right data to the right people at the right time



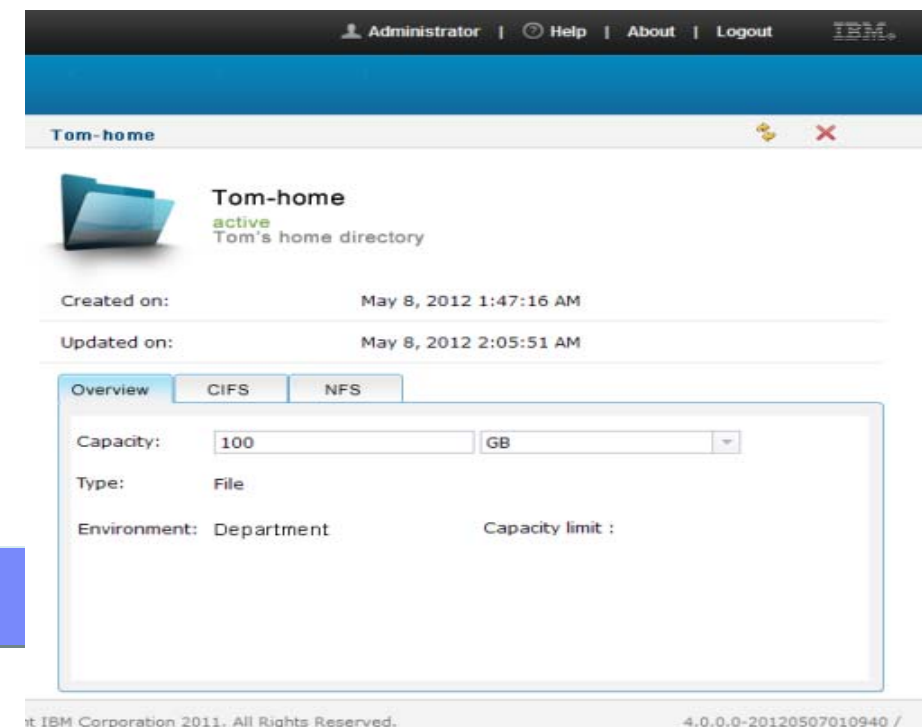
- **Pay per use** pay only for what is consumed



- **Elasticity** capacity growth without constraints, release resources when not needed

Introducing IBM SmartCloud Storage Access

- **Self-service**
 - Capacity provisioning, monitoring, reporting
- **Web-based access**
 - Simple GUI via Internet
- **Capacity elasticity**
 - Automated or policy-driven
- **Monitoring and reporting capability**
 - End-user and administrator
- **Service Class**
 - Gold, Silver, Bronze
- **Secure**
 - Access controls
- **Approval**
 - Complete approval control process

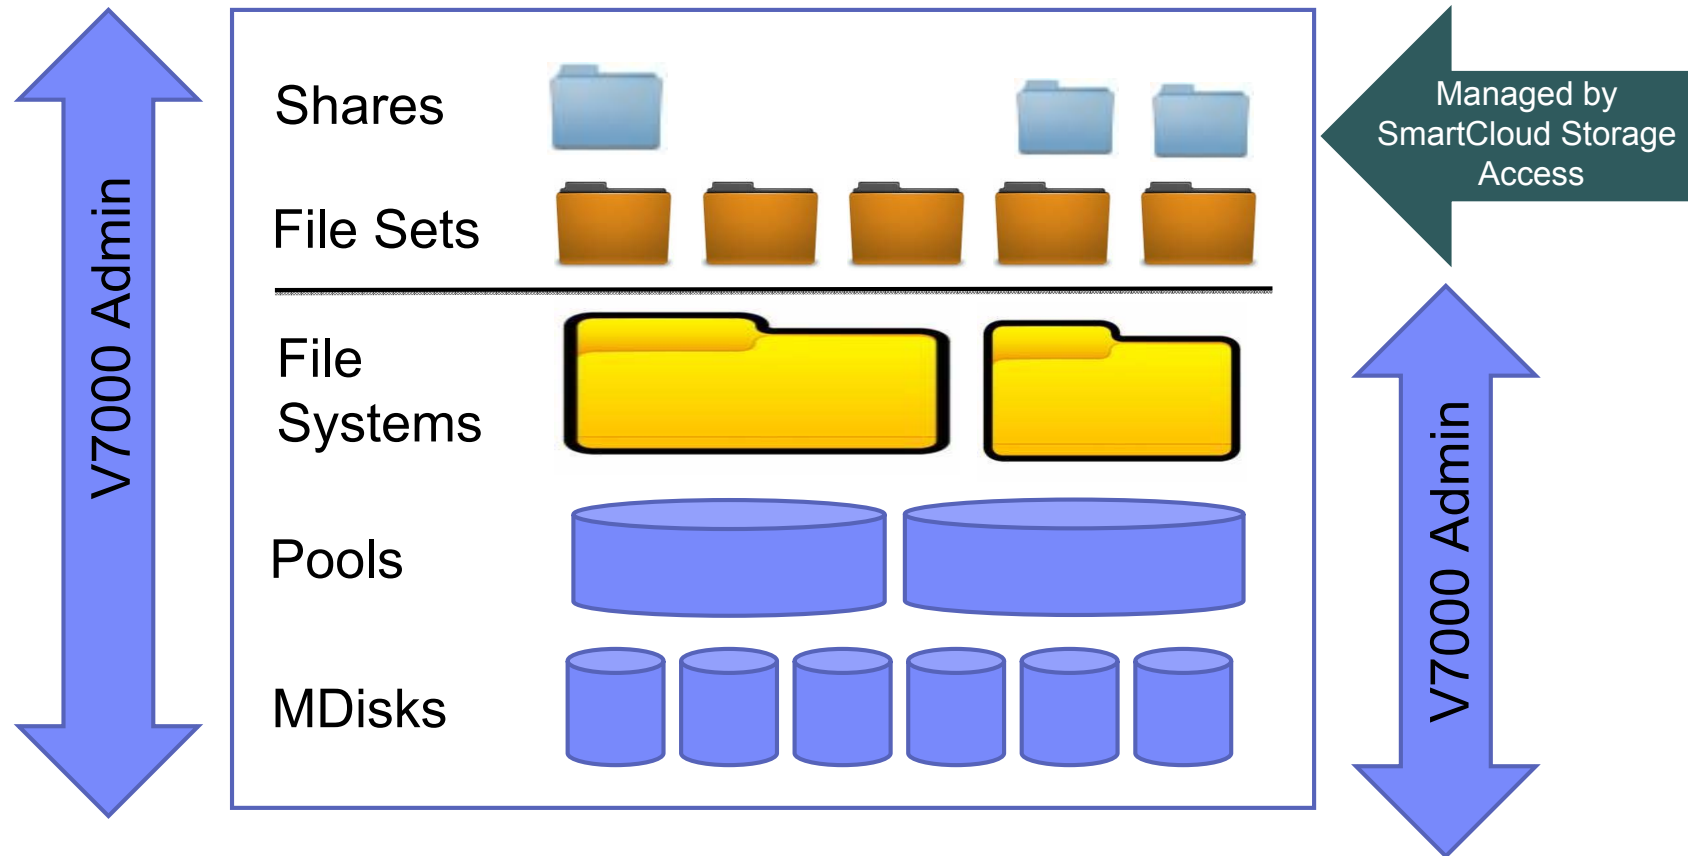


Enables private cloud storage service

V7000 Unified SmartCloud Storage Access management

Without SmartCloud Storage Access

With SmartCloud Storage Access



V7000 Unified structure

XIV Announcement Highlights

Announce: February 5, 2013; General Availability: March 7, 2013

Price-performance efficiencies for cloud and virtualization

<p>Cloud Agile</p>	<p>Up to 5X 10Gb iSCSI throughput Windows Server 2012 support</p>	<p>➔</p> <p>Bullet-proof cost-effective speed for IP-based SANs Value for Microsoft Hyper-V, VSS, and MSCS</p>
--------------------	--	---

Higher performance and availability, ALWAYS tuning-free

<p>Self-optimizing</p>	<p>Up to 4.5X boost in DB performance with enhanced SSD caching Up to 13.7 GB/s sequential read per rack Dramatically reduced drive rebuild times</p>	<p>➔</p> <p>OLTP Tier-1 applications High throughput for Analytics, Streaming Media, and Backup XIV provides the lowest rebuild times of a 2TB drive in 26 minutes at 100% utilization</p>
------------------------	---	--

Reduced TCO through energy efficiency and capacity optimization

<p>Efficient by Design</p>	<p>Up to 16% reduced power consumption Extended Space Reclamation/UNMAP</p>	<p>➔</p> <p>Reduce \$ and carbon footprint Stretch your TBs; defer CapEx Efficient thin provisioning</p>
----------------------------	--	--


XIV Model 214 - New CPU

■ New 6-core CPU

- Intel Xeon E5645
- 2.4 GHz
 - Sized to support module I/O processing
 - Key benefit of XIV architecture
- XIV uses Intel Hyper-Threading
 - Delivers 2 logical processing threads per physical core
 - ✓ 6 physical CPUs per module
 - ✓ 12 logical CPUs per module
 - 15 Module system
 - ✓ 90 physical cores
 - ✓ 180 logical cores (hyper-threading)

■ Support for:


- New high-speed iSCSI interfaces
 - Utilizes Light Weight IP Stack (LWIP)
- Platform for future hardware/software enhancements



Multiprocessing
2-core processor

Multiprocessing
2-core processor with Hyper-Threading

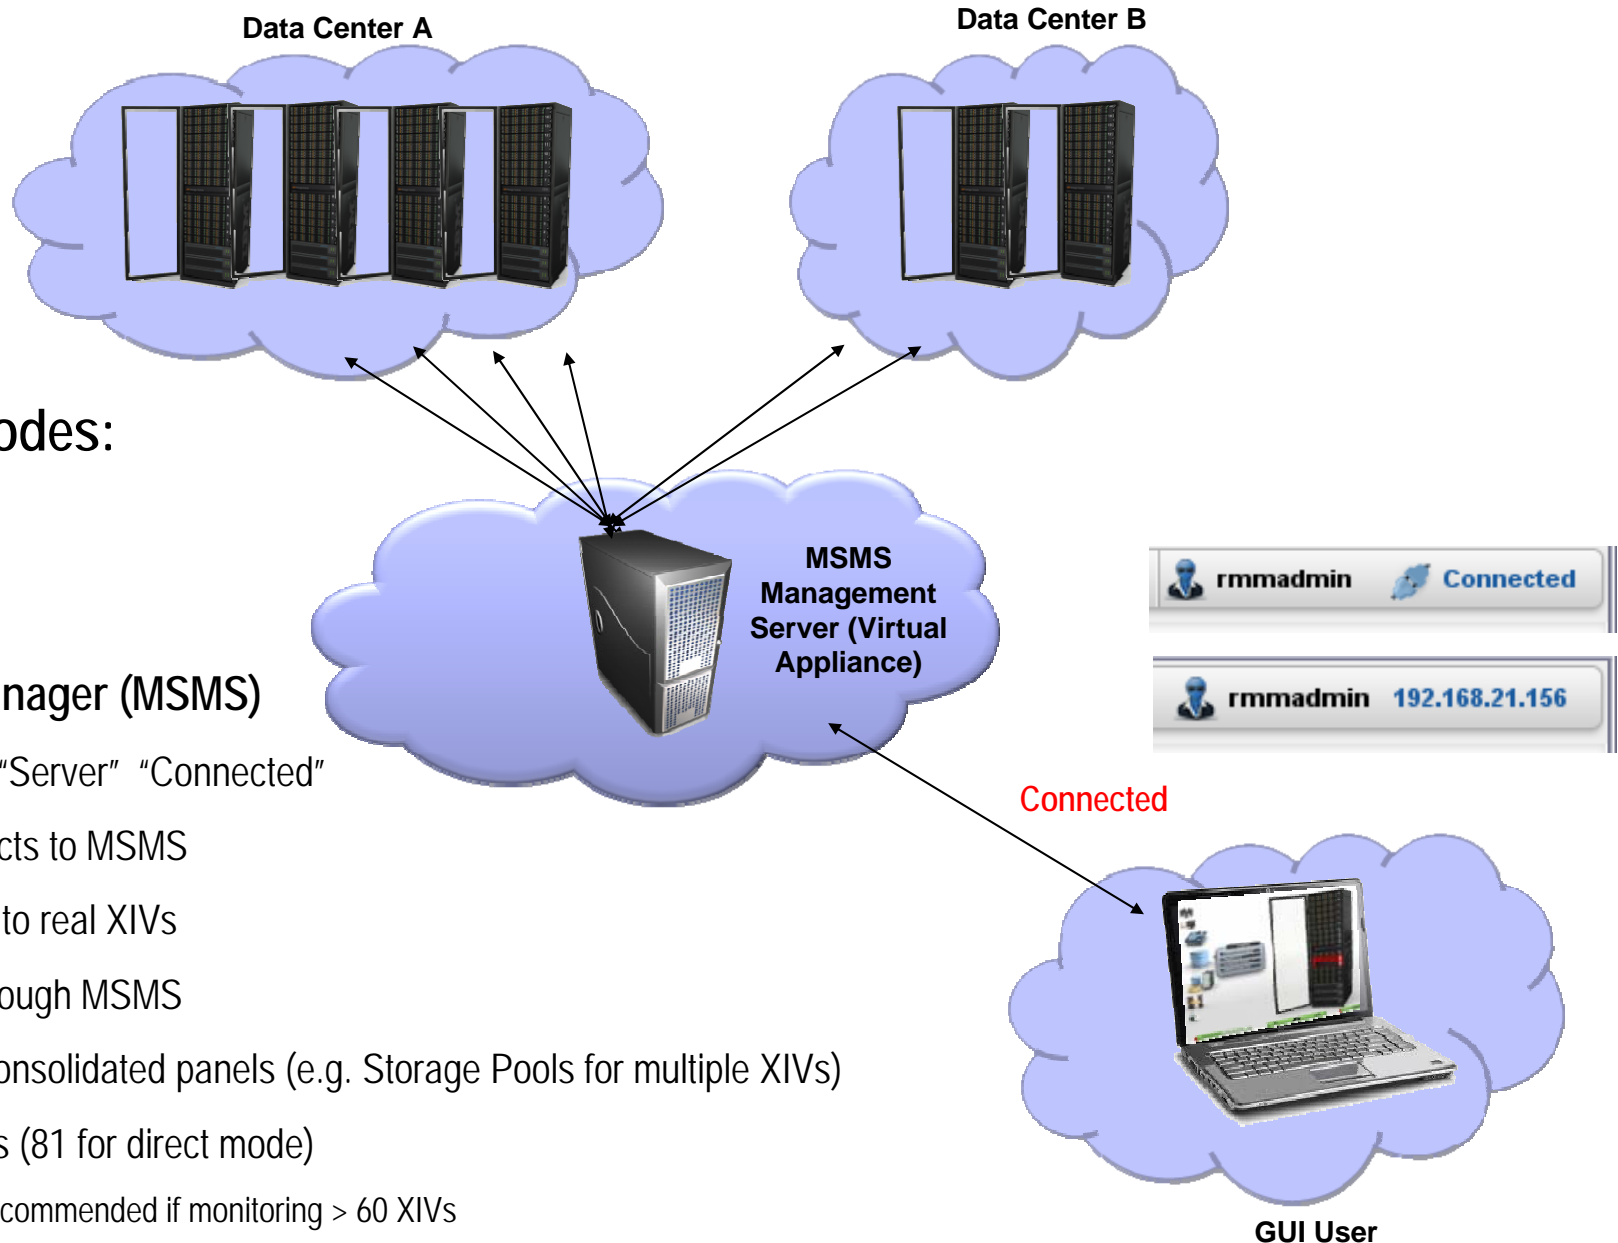
http://cache-www.intel.com/cd/00/00/01/77/17705_htt_user_guide.pdf



XIV 6-core CPU
with hyper-threading

Cloud agile **Self-optimizing** Efficient by design

XIV GUI 4.0 Multi-System Manager (MSM) Mode



■ 3 XIV GUI 4.x modes:

- ▶ Demomode
- ▶ Direct

★ ▶ Multi-System Manager (MSMS)

- Also "Manager" "Server" "Connected"
- GUI client connects to MSMS
- MSMS connects to real XIVs
- All actions go through MSMS
- Allows viewing consolidated panels (e.g. Storage Pools for multiple XIVs)
- Max 144 systems (81 for direct mode)
 - MSMS mode recommended if monitoring > 60 XIVs

MSMS 1.1 - Consolidated Views (All Systems or Group of Systems)

- Storage Pools
- Volumes by Pools
- Volumes and Snapshots
- Consistency Groups
- Hosts and Clusters
- Host Connectivity
- QoS Performance Class
- Volumes by Hosts
- iSCSI Connectivity
- Mirroring
- Migration
- Users
- **Usage**
 - Check server connectivity to multiple XIVs
 - Check mirror master and secondary configuration or status

The screenshot displays the XIV Storage Management console. The breadcrumb trail is 'All Systems (4) > ATS_Lab (2) > Consistency Groups'. The table below shows various storage volumes with their names, sizes, and redundancy status.

Name	Size	Master	Pool	Created (GU...)	Cre...
Unassigned Volumes ATXIV-1310115					
Unassigned Volumes ATXIV-7803015					
MY_ESX_CG_DR ATXIV-7803015					
MY_SRM_GM ATXIV-7803015					
RMM_PROD_CG_ASYNC ATXIV-1310115		RMM_PROD_ASYNC	0%	10,015 GB	
RMM_PROD_CG_ASYNC ATXIV-7803015		RMM_PROD_ASYNC	6%	1,013 GB	
RMM_PROD_CG_SYNC ATXIV-1310115		RMM_PROD_SYNC	2%	1,015 GB	
RMM_PROD_CG_SYNC ATXIV-7803015		RMM_PROD_SYNC	3%	1,013 GB	
TEAM00_VMWARE_DR_CG ATXIV-7803015		TEAM00_VMWARE_P...	20%	515 GB	
TEAM01_VMWARE_DR_CG ATXIV-7803015		TEAM01_VMWARE_P...	20%	1,013 GB	
TEAM02_VMWARE_DR_CG ATXIV-7803015		TEAM02_VMWARE_P...	20%	515 GB	
TEAM03_VMWARE_DR_CG ATXIV-7803015		TEAM03_VMWARE_P...	23%	515 GB	
TEAM04_VMWARE_DR_CG ATXIV-7803015		TEAM04_VMWARE_P...	20%	515 GB	

Bottom status bar: **Hard: 92,587 of 215,440 GB (43%)** | **IOPS 307,646 Total** | **Full Redundancy (2)**

- Change password (or edit or delete user) for same user across multiple systems

Trademarks and Disclaimers

IBM Corporation 1994–2012 All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at <http://www.ibm.com/legal/copytrade.shtml>.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered

trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided “AS IS” without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.