

A New Era in Storage Efficiency & Innovation

100 Jahre IBM Storage Vergangenheit – Gegenwart - Zukunft

Thomas Frey
Highend Storage Sales Leader
thfrey@de.ibm.com

Optimized systems.
Optimal innovation.



100 Years IBM Storage



100 Years IBM



1911



1924



1949



1956



NOW



Tape - Disk - CMOS


3

1952: IBM Modell 726 erster Bandspeicher

- 18.000 Lochkarten
- 1.440.000 Characters
- 1.44 MB
- Eisenoxydbeschichtung
- 7-Spur Technik



Photo 1951 Prototyp



720 Meter Bandlänge

Photo 1952 IBM 726 in Betrieb



100 BPI

4

IBM EDPM System 705,
die erste Installation in Europa bei

Kunde **Farbwerke Höchst AG**

Im Bild rechts erläutert Herr Horst Übner, Vertriebsbeauftragter der IBM GS Frankfurt/M in 1957 die Speicherkapazität einer **Magnetbandspule** im Vergleich zu Lochkarten. Er sagt: sehen Sie, diese 12 Kartons -mit je 2000 Lochkarten- die ich auf diesen Tisch gestapelt habe, beinhalten 24.000 Lochkarten, deren Daten können alle auf dieser einen Magnetbandspule -die ich in der linken Hand halte- gespeichert und in einem Bruchteil der Zeit von Lochkarten gelesen werden. Die Spule enthält etwa 750 m Magnetband und kostet runde 400 DM.



- 1953: IBM Modell 727 (728)
- 24.000 Lochkarten
 - 1.920.000 Characters
 - 1.92 MB
 - 7 Spur-Technik

IBM

© IBM Corporation, 2012

Enterprise Tape Entwicklung (nach Magstar)

IBM



- 2003 IBM 3592 (Jaguar Gen.1)
 - 40 MB/s,
 - 60 GB ,300 GB Kassetten
- 2005 IBM TS1120 (Jaguar Gen.2)
 - 104 MB/s,
 - 100 GB , 500 GB Reuse, 700 GB Kasette
- 2008 IBM TS1130 (Jaguar Gen.3)
 - GMR erste Generation
 - 160 MB/s,
 - 128 GB, 640 GB und 1 TB

Media Reuse



300-640GB 700-1600GB 4000GB *
*native ohne Kompression, Kapazität abhängig von Laufwerksgeneration

- 2011 IBM TS1140 (Jaguar Gen.4)
 - GMR dritte Generation
 - 250 MB/s,
 - 500GB , 4 TB

5



© IBM Corporation, 2012

Tape

29.5 Gbit/in² Demonstration mit GMR – 35 TB Kassette


Press release on Jan 22, 2010: <http://www.zurich.ibm.com/news/10/storage.html>

Demo Hardware Platform



Track density = 57 ktpi (track width = 0.446 µm)

- Perpendicular Recording (senkrechte Bitanordnung)
- Ultrafeine Barium-Ferrit-Medien ohne Metallaufdampfen
- Reduktion der Spurbreitung auf unter 25 nm
- Spurbreite auf unter 0.45 µm (Faktor 25)
- Erhöhung S/L-Geschwindigkeit um Faktor 38



Target BER: 1e-4
0.2 µm reader width

FPS4 Linear Density w/ 16 DDPML: 518 kbp

- Archivierung und Langzeitdatenspeicherung
- riesige Kassettenkapazitäten
- stromloser Datenträger
- Linear Tape File System LTFS

Reynold B. Johnson



Bild von 1971


Zielsetzung Anfang der 50-er Jahre:

Die Möglichkeit, jeden Geschäftsvorfall dann zu bearbeiten, wenn er anfällt mit einer Leistung von etwa 10.000 Fällen pro Tag!
(Rechenanlage ohne Stapelverarbeitung)



- take the best people
- go to San Jose
- make something cool
- find a new way to store data

1952 - 1961 die Anfangsepoche der elektromagnetischen Speicherung



1961

IBM 1301
28 und 56 MB per
Plattenmodul

Air Bearing Sliders
mit Landebahnen

Zugriffskamm mit
Köpfen für jede
Plattenoberfläche

Einführung der
Zylinderarchitektur

1975 - 1994 Epoche der fest eingebauten Platten mit Kontrolleinheiten



1975 IBM 3350 (2.5 GB)



1979 IBM 3370/75 (3 GB)



1981-87 IBM 3380
(60 GB)



1989-94 IBM 3390
(90 GB)

2004 - 2012 Epoche der Serverbasierenden Speichersysteme und der Speichervirtualisierung


2004
IBM DS8000 Series



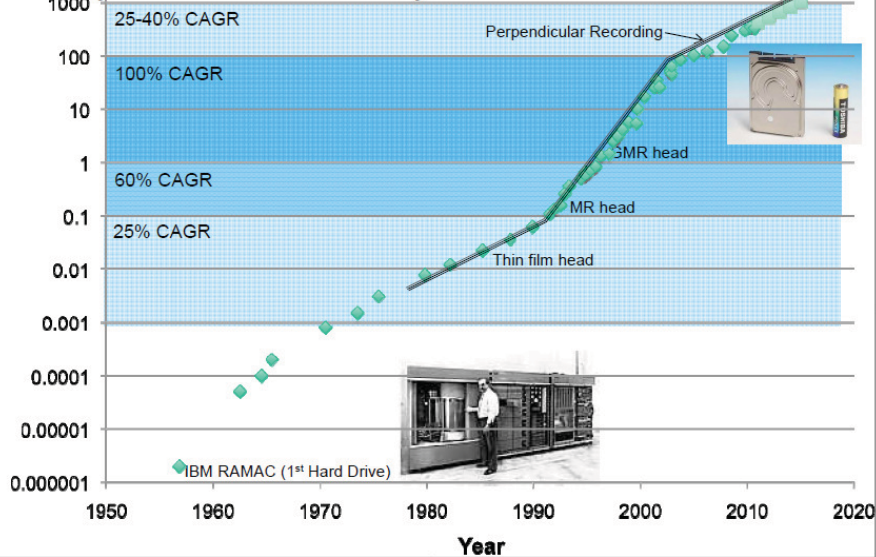
DS8800
1,38 PB SAS Disk



2008
IBM XIV Storage System 240TB




History of HDD Areal Density Growth



Explosion in Data Growth


- Lower storage costs allows..
 - Storing more types of data
 - Keeping data longer
 - Using disk in new ways
- Smarter planet
 - Instrumented and intelligent
- Business intelligence growth
 - From “C” level to end-users
- Results
 - Storage growing 20-40% per year
 - Storage budgets growing 1- 5% per year

Year	Tier 1 Disk Cost / TB
1980*	\$32,914,000
2011	\$7,000




* Source: Price for 3380 disk in 1980, wikipedia.org

Effizienz-Herausforderung:



Technologie-Ansätze helfen die Ideal-Linie zu finden



Technologien für Storage-Effizienz



▪ Store more with what's on the floor

- Storage Virtualization
- Thin Provisioning



▪ Move data to the right place

- Easy Tiering
- Automated Data Migration
- ILM/HSM



▪ Store less data

- Data Compression
- Data Deduplication



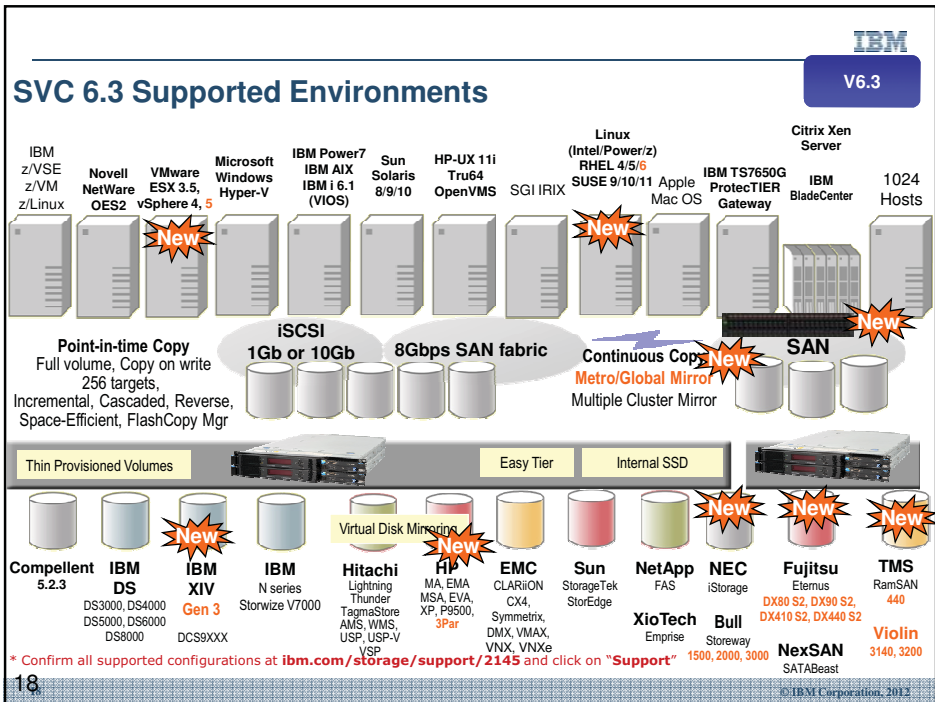
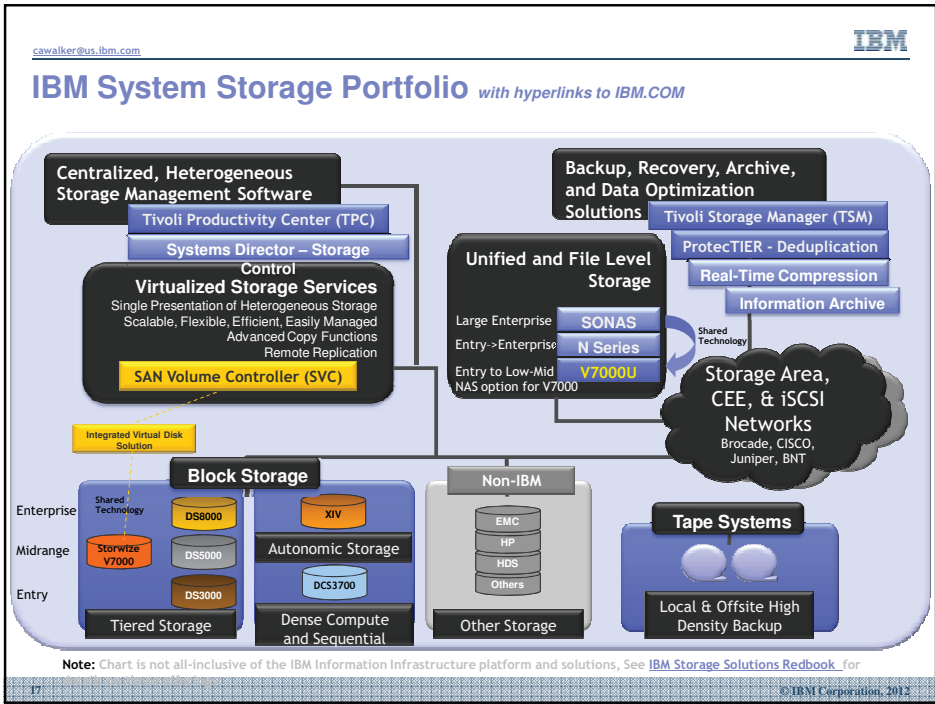
▪ Improve people productivity

- More intuitive user experience (GUI)
- Integration, Automation, System Director,

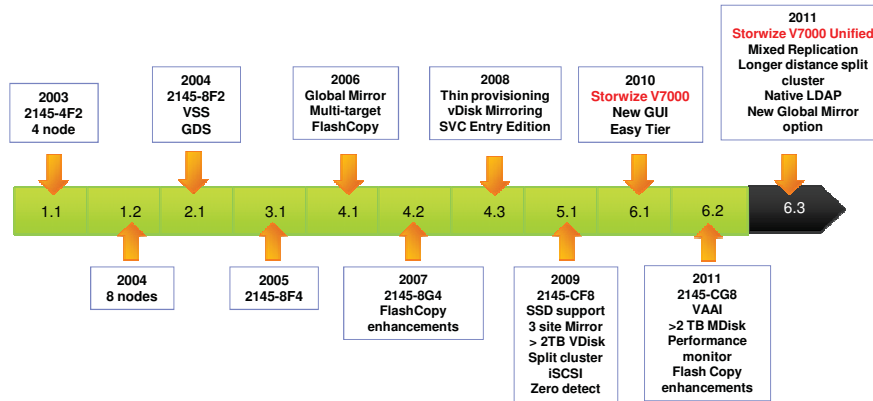


▪ Data protection

- Snapshot, Mirroring, HA, FCM
- Archive, Tape, VTL



SVC release history is a continuous evolution



Storwize V7000 – Modularity and Advanced Software Functions

Modular Hardware Building Blocks in 2U

- Enclosures contain up to twelve 3.5" or twenty-four 2.5" drives in just 2U
- Control enclosure: dual active-active controllers and drives; Expansion enclosure: drives only
- Up to nine expansion enclosures attach to one control enclosure
- Mix drive sizes and HDD/SSD in enclosure
- Eight 8Gbps FC ports plus four 1Gbps iSCSI ports per controller pair; 16GB cache per controller pair



Software inherited from prior offerings plus enhancements

New advanced software functions

- New GUI (*easy-to-use, web based*)
- RAS services and diagnostics
- Additional host, controller and ISV interoperability
- Integration with IBM Systems Director
- Enhancements to TPC, FCM and TSM support

Proven IBM software functionalities

- Easy Tier (*dynamic HDD/SSD management*)
- RAID 0, 1, 5, 6, 10
- Storage virtualization (*internal and external disks*)
- Non-disruptive data migration
- Global & Metro Mirror
- FlashCopy up to 256 copies of each volume
- Thin provisioning

Clustered system up to two Storwize V7000 control enclosures

Base Storwize V7000		Clustered System	
Maximums		Maximums	
Hosts (FC): 256	FC ports: 8x8Gb	Hosts (FC): 512	FC ports: 16x8Gb
Hosts (iSCSI): 64	Ethernet ports: 4x1Gb	Hosts (iSCSI): 128	Ethernet ports: 8x1Gb
Capacity: 24TB (optional)	4x10Gb	Capacity: 48TB (optional)	8x10Gb
Cache: 16GB	Volumes: 2048	Cache: 32GB	Volumes: 4096
Drives: 12 LFF or 24 SFF		Drives: 24 LFF or 48 SFF or mix	
Disk types: SAS, NL-SAS, SSD		Disk types: SAS, NL-SAS, SSD	

- Delivers ability to expand Storwize V7000 system capacity two-fold
- Enables distributing workloads and processing to another control enclosure with a new I/O group



+ External Virtualization for more capacity

Clustered System		System with External Virtualization	
Maximums		Maximums	
Hosts (FC): 256	FC ports: 8x8Gb	Hosts (FC): 512	FC ports: 16x8Gb
Hosts (iSCSI): 64	Ethernet ports: 4x1Gb	Hosts (iSCSI): 128	Ethernet ports: 8x1Gb
Capacity: 240TB (optional)	4x10Gb	Capacity: 480TB (optional)	8x10Gb
Cache: 16GB	Volumes: 2048	Cache: 32GB	Volumes: 4096
Drives: 120 LFF or 240 SFF or mix		Drives: 240 LFF or 480 SFF or mix	
Disk types: SAS, NL-SAS, SSD		Disk types: SAS, NL-SAS, SSD	

- Provides up to **2X capacity** and up to **2X throughput** for business growth
- **Fully upgradable without disruption** from smallest to largest configurations: minimizes need to purchase hardware in advance of need
- **Enables more consolidation**, including external virtualized storage, for greater efficiency



XIV Generation Comparison



	XIV 2810/2812-A14	XIV Gen3 2810/2812-114
Drives	72-180	72-180
Interconnect	Ethernet	InfiniBand
Disk Drives (7200 RPM)	SATA (1 TB or 2 TB)	SAS (2 TB or 3 TB)
Number of disk drives (min/max)	72/180	72/180
SSD capacity per module	N/A	400 GB
Max Capacity w/1 TB drives	79 TB	N/A
Max Capacity w/2 TB drives	161 TB	161 TB
Max Capacity w/3 TB drives	N/A	243 TB
Max FC ports	24	24
Max iSCSI ports	6	22
Max iSCSI ports in 6-module	N/A	6
Max number of CPU cores	84	60
Max Memory	120 GB (8 GB per module) 240 (16 GB per module) DDR2	360 GB (24 GB per module) DDR3
Max cache-to-disk bandwidth	240 Gb/sec	480 Gb/sec
Processor	E5410 Intel Quad Core XEON 2.33 GHz	E5620 Intel Quad Core XEON 2.4 GHz
Host FC Adapters	4 Gb/sec	8 Gb/sec
Host iSCSI Adapters	1 Gb/sec	1 Gb/sec

XIV Gen3 Technology Highlights

Smarter Storage, Built for Performance

- 20X more internal bandwidth
 - Using InfiniBand
- Over 2X more external bandwidth
 - With 8 Gb/sec FC ports and over 3x more iSCSI ports (6-22)
- New motherboards and processors
 - 2x disk bandwidth, 60 cores, 120 hyper-threads per rack
- 50% more cache capacity
 - Up to 360GB/system, (24GB per module)
- SSD ready
 - Optional cache upgrade of up to 6.0TB
- Announced July 12, 2011
- Generally available September 8, 2011



23

© IBM Corporation, 2012

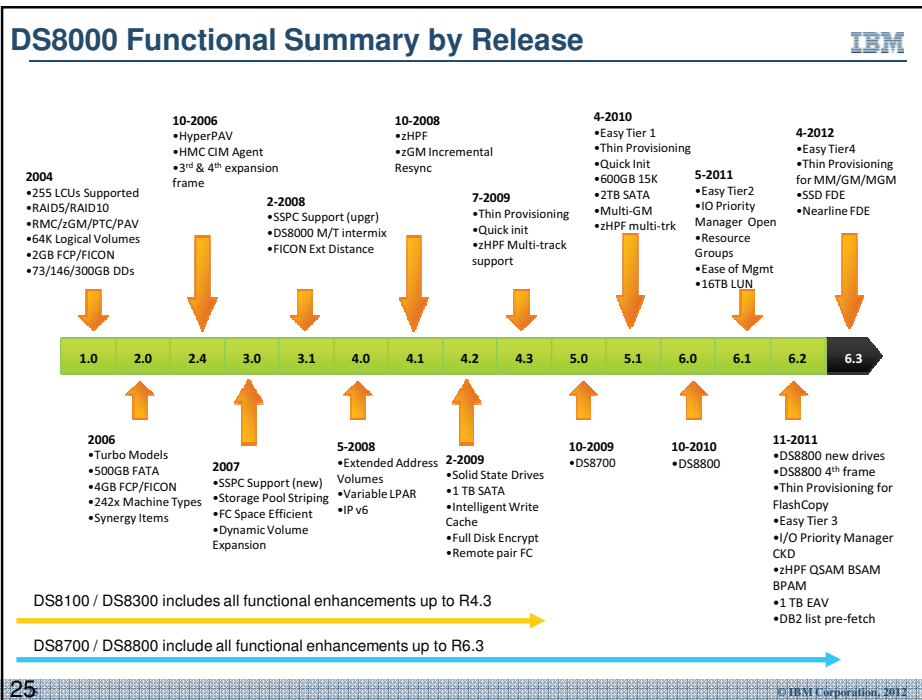
DS8000 – IBM's Premier Storage Array

- DS8000
 - POWER6+ based controllers
 - Up to 1,536 drives
 - Solid state disk, SAS, and NL-SAS drives
 - Up to 1,38 TB capacity
 - 24x7 warranty
- Advanced functions including
 - 3-level easy-tier and automatic rank balancing
 - Three-site replication
 - Hyperswap
 - Native System z and IBM I attachment
 - Capacity on demand



24

© IBM Corporation, 2012



Release 6.3 content

IBM

- **New 400GB SSDs**
Increases SSD capacity by 33% and improves performance
- **New self-encrypting drive options**
FDE now available for all drive tiers at virtually the same price as standard drives
- **Updated Easy Tier** (4th generation) support for encryption
The best automated tiering and the best security
- **Enhanced Thin Provisioning**
Advanced Metro Mirror and Global Mirror support for open systems

26
© IBM Corporation, 2012

R6.3 Support Matrix – DS8800

No changes to DS8700 drive support

Only drive w/out FDE option

	Solid State Disk	Enterprise Disk	Nearline Disk
Disk Type	2.5" SAS	2.5" SAS	3.5" SAS
DDMs per Enclosure Pair *	48	48	24
Disk Installation Group	8 or 16	16	8 or 16
RAID Types	5, 10 **	5, 6, 10	6
Size/RPM	300 400	146/15K 300/15K 450/10K 600/10K 900/10K	3TB/7.2K
Encryption-Capable Size/RPM	400	146/15K 300/15K 450/10K 600/10K 900/10K	3000/7.2K




* Cannot intermix 2.5" and 3.5" disks in same enclosure pair.

** RPO only based on customer request. Can increase long term drive failure rates and chances for double drive failures. © IBM Corporation, 2012

Open – 2.5" 400GB SSD Drive Performance

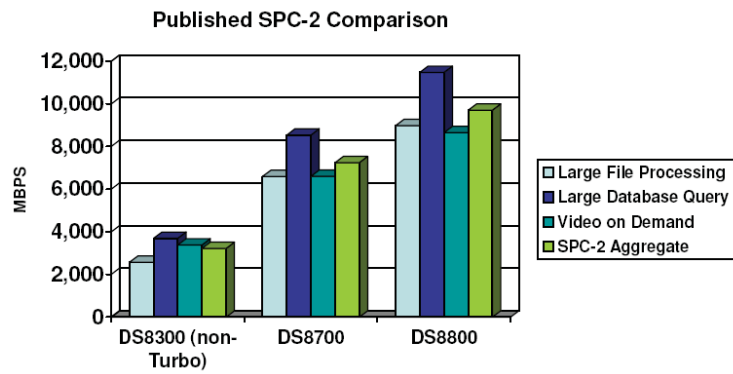
The new 2.5" 400GB SSDs on the DS8800 has equal or slightly better performance for 4KB random read/write and sequential read comparing with 2.5" 300GB SSDs, however, for sequential write 400GB SSDs perform more than 2X better.

Single Rank – RAID5



SPC-2 Benchmark

IBM



DS8800 host adapter performance measured using 8 Gb/second host adapter
 DS8300 and DS8700 host adapter performance measured using 4 Gb/second host adapter

29

© IBM Corporation, 2012

Technologien für Storage-Effizienz

IBM



Store more with what's on the floor

- Storage Virtualization
- Thin Provisioning



Move data to the right place

- Easy Tiering
- Automated Data Migration
- ILM/HSM



Store less data

- Data Compression
- Data Deduplication



Improve people productivity

- More intuitive user experience (GUI)
- Integration, Automation, System Director,



Data protection

- Snapshot, Mirroring, HA, FCM
- Archive, Tape, VTL

30

© IBM Corporation, 2012

Store More with What's on the Floor

31 © IBM Corporation, 2012

IBM

Virtualization Improves Utilization and Performance

- Virtualized storage pools
 - Stripe data over multiple RAID arrays
 - Increases available disk IOPs
 - Automatic configuration or administrator controlled


- Benefits
 - Allows increased drive utilization
 - Reduces hot spots
 - Reduces tuning effort

Virtualized Storage Pools

33 © IBM Corporation, 2012


XIV - Extreme Virtualization

- Every volume uses every drive
- Virtually eliminate hot spots and orphaned capacity
- Reduces management effort
 - No data placement
 - No snapshot layout
 - No spare drive management
 - No hot spot management
 - No RAID configuration tasks
 - No drive rebalancing
- Short rebuild time



Multi-Array Storage Virtualization

- Storage array virtualization
 - Scalable highly-available cluster design
 - Supports wide range of storage arrays
- Provides
 - Common LUN management
 - Common multi-path drivers
 - Common local and remote copy services
 - Transparent data migration between arrays
 - Easy Tier
- Benefits
 - IBM cut internal block storage costs by 50%



Move Data to the Right Place

37 © IBM Corporation, 2012


IBM

Storage Tiering

No Tiering	Transparent	Policy Driven
<ul style="list-style-type: none"> • Small or very simple storage footprints • Desire for uniformity • Tier-less architecture such as XIV 	<ul style="list-style-type: none"> • Tiering is desirable • Data may be difficult to manage with rules • Desire minimal administrator effort 	<ul style="list-style-type: none"> • Data can be managed with rules • Applications with special needs • Desire greater control

38 © IBM Corporation, 2012

**Evolution: storage media have gained capacity...
but no significant performance**




Performance × 50
Capacity × 50.000

- Hard Disk**
- ATA, SATA, FATA, SAS (2.0), FC-Platten
 - 2 ½ Zoll vs. 3 ½ Zoll Formfaktor
 - GMR (Giant Magneto-Resistance)



- Performance Entwicklung



IO Rate per GB und
Hard Disk Drive Generation

Seit 1990 um **Faktor 100**
schlechter

- Performance heute

	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

Solid State Disk (SSD)

Preisentwicklung

SSDs mit DRAM

- seit 1976 für Hochleistung
- 100x höhere Kosten als Disk

SSDs mit Flash

- SLC = Single Level Cell
 - 1 Bit (0/1) per Zelle
 - höhere Lebensdauer, schneller, teurer
- MLC = Multi Level Cell
 - Multiple States per Zelle (00/01/10/11)
 - geringere Lebensdauer, günstiger
- eMLC (Enterprise MLC)
 - Zellenabsicherung durch CU
 - Lebensdauer SLC vergleichbar

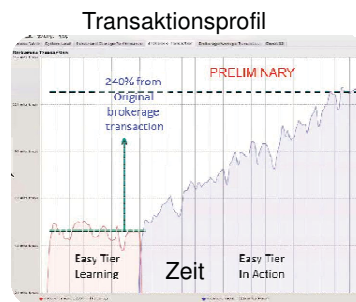
Zeit	SSD	Preis (\$)
2003	21 GB 2.5" SLC	11000.-
2006	32 GB 2.5" MLC	1000.-
2010	32 GB 2.5" MLC	100.-
2012	160 GB 2.5" MLC	80.-



Tiering oder Cache:

- Database DS8000: 20ms to 2ms responding improvement (Oracle/ IHK-GfK)

http://www.youtube.com/watch?v=zjg99ulzISl&feature=results_main&playnext=1&list=PL687BD0ECB092D739




XIV Gen3
Infiniband Interconnect

- Max Capacity: 161TB
- Max Memory: 360GB
- Max FC ports: 24 x 8Gb/s
- Max iSCSI ports: 22 x 1Gb/s
- iSCSI ports with 6 modules
- Disk Type: SAS, SSD


6TB SSD

- Cache SSD in XIV
- ERP/CRM 70/30/50/8k
- 20TB
- von 60.000 auf 123.000 IOs



SSDs Help to Improve XIV Performance

- SSDs is a secondary cache between the disks and the DRAM cache
- Used to improve random read hits
- Each module has SSD drive sized for optimum performance
- No administrative effort required
- Entire SSD is used as data is already on disk



Easy Tier for optimized tiering

- Monitors performance of each extent to determine the data 'temperature'
- Creates extent migration plan for optimal data placement every 24 hrs based on performance statistics
- Migrates extents across heterogeneous tiers **and** within a homogeneous tier according to migration plan(s) over 24 hour period
- A limited number of extents are chosen for migration every 5 minutes to avoid performance impact



Easy Tier generational enhancements (to date)

Self-tuning storage manages growth and complexity while lowering costs



1st Generation

- **Objective:** Optimizing use of expensive SSDs
- **Benefit:** Relocating just 5% of the data from HDDs to SSDs, reduced average I/O response time from 9ms to 2ms

2nd Generation

- **Objective:** Intra-tier rebalancing and support for any 2 tiers (no SSDs)
- **Benefit:** Automates performance optimization *within* a tier; Relocating less active data to slower nearline drives maintains performance with lower \$/GB as data grows

3rd Generation



- **Objective:** Full support across all 3 tiers and support for Thin Provisioned volumes
- **Benefit:** Faster performance when and where it's needed with SSDs; Cost savings (reduced footprint and \$/GB) for cold data; flexibility to support standard and Thin Provisioned volumes

4th Generation

- **Objective:** Support for Full Encryption
- **Benefit:** Combines advanced tiering with superior security for the ultimate in efficiency and data protection

Easy Tier V1 Benefits

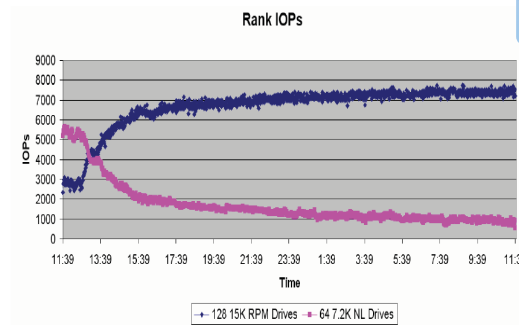

- Relocating just 5% of the data from HDDs to SSDs, reduced average I/O response time from 9ms to 2ms!
- Optimizing use of expensive SSDs



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

Easy Tier V2 Benefits

- “Any two tier” offered support for combining Enterprise Class and Nearline Class to maintain performance with lower \$/GB as data grows




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

Easy Tier v3 Benefits

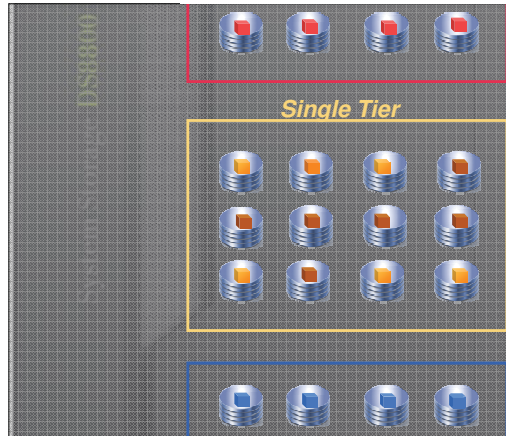
Combining the benefits across tiers and within a tier

- 3-tier support combines the benefits
 - Faster performance when and where it's needed with SSDs
 - Cost savings (reduced footprint and \$/GB) for cold data
- Plus, auto rebalancing supports single-tier (homogenous) as well as mixed tier pools
 - Rebalance after new resources are added or removed from the storage pool
 - Ensures balanced rank utilization to distribute natural performance skew
- Easy Tier support for thin provisioned volumes



Easy Tier intra-tier rebalancing

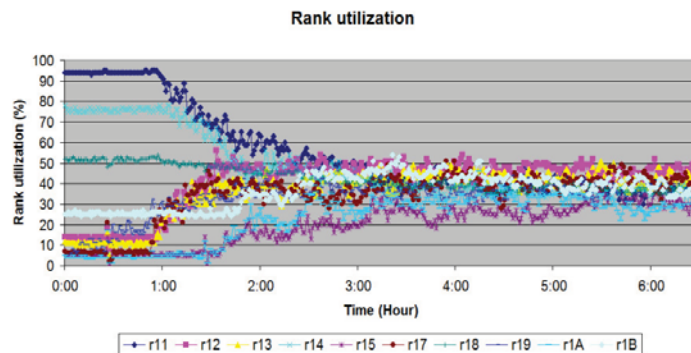
- Continuous rebalancing *within* each tier maintains peak performance across all drives
- Easy Tier complements other performance optimization features
 - Advanced caching algorithms
 - Storage Pool Striping (wide striping)
 - I/O Priority Manager for QoS



Easy Tier intra-tier Rebalancing – Balanced Utilization

Balanced Rank utilization on single-tier pool

- The re-balanced system shows flattened rank utilization across all the ranks, all without any operator activity!




All performance data contained here were obtained in the specific operating environment and under the conditions and/or configuration described above and is presented as an illustration. Details are available in the "IBM System Storage DS8800 and DS8700 Performance with Easy Tier 3rd Generation" whitepaper. Performance obtained in other operating environments may vary and customers should conduct their own testing.

Easy Tier intra-tier Rebalancing – Improved Throughput

Supports single-tier and multi-tier systems

Performance gain on single-tier pool

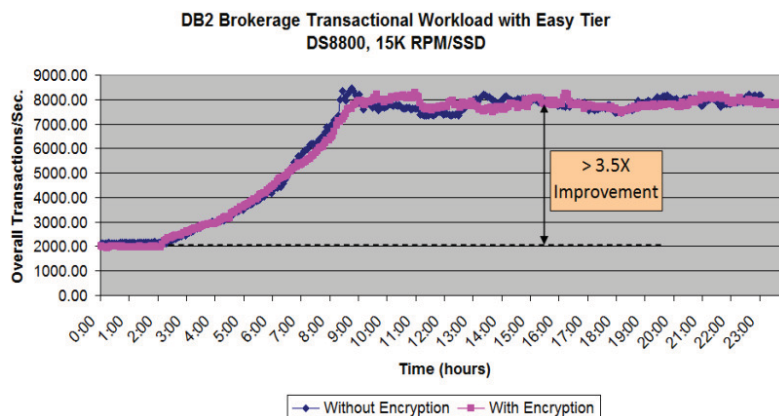
– The re-balanced system showed a **3x** throughput improvement and equal or better response times without any operator activity!



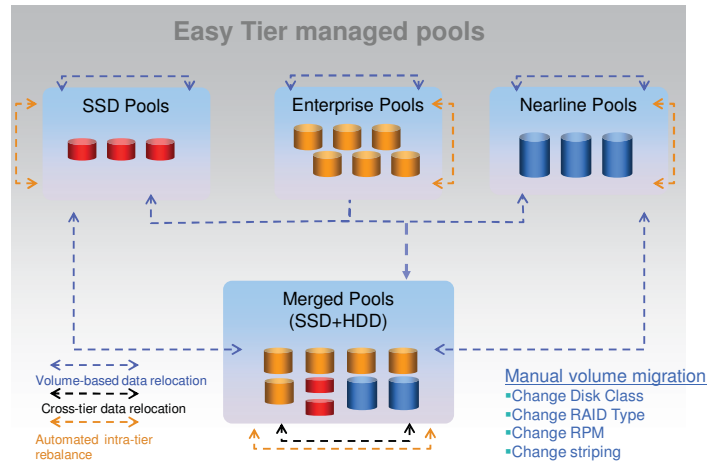
All performance data contained here were obtained in the specific operating environment and under the conditions and/or configuration described above and is presented as an illustration. Details are available in the "IBM System Storage DS8800 and DS8700 Performance with Easy Tier 3rd Generation" whitepaper. Performance obtained in other operating environments may vary and customers should conduct their own testing

Easy Tier4 Application Performance

- Equivalent performance with or without encryption drives in the DS8800



Summary of Easy Tier Migration Capabilities




The Coming Revolution for Data Center Efficiency Will Be Driven by SSD Technology

Published: 22 March 2012

Analyst(s): Gene Ruth, David J. Cappuccio

IT staff should look to SSD technology to help mitigate the avalanche of data flooding into data centers. While not inexpensive, SSDs have instigated a rethinking of storage array design, resulting in dramatic equipment efficiencies measured in performance, floor space and power consumption.

Will Flash Memory replace disks? \$ per GB prediction

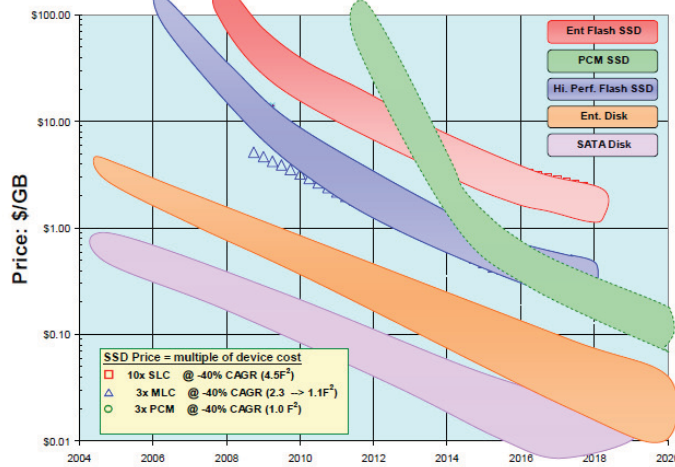


Source: IBM Almaden Research, Steven R. Hetzler, Sep 2009

Cost Evolution

- Device cost reductions will not keep up with increased use
- Lower cost devices do not improve management costs
- System level efficiency will be a key

Price Trends: Magnetic disks and Solid State Disks



First drift-tolerant 2bit/dot PCRAM from IBM Rueschlikon



	Dynamic Random Access Memory (DRAM)	Flash	Phase-Change Memory (PCM)	Multi-Level Phase-Change Memory
Timeframe	Invented in 1966	Invented in 1980s	Available in limited number of architectures - wide adoption expected by 2016	2016
Speed	Green	Red	Yellow	< × 100
Density	Yellow	Green	Yellow	Green
Endurance	Green	Red	Yellow	10 million write cycles
Retention	Red	Green	Green	Green
Scaling	Orange	Yellow	Green	Green

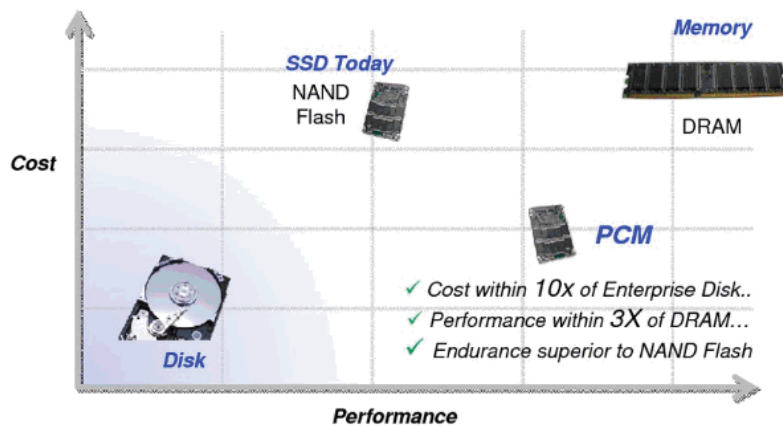
CREDIT: IBM RESEARCH

- Best in class
 - Average / Inadequate
 - Good / Adequate
 - Bad / Worse in class
- ▶ The qualification is relative and depends on the application
- ▶ RaceTrack Memory is not included in this time horizon

2 Mbit

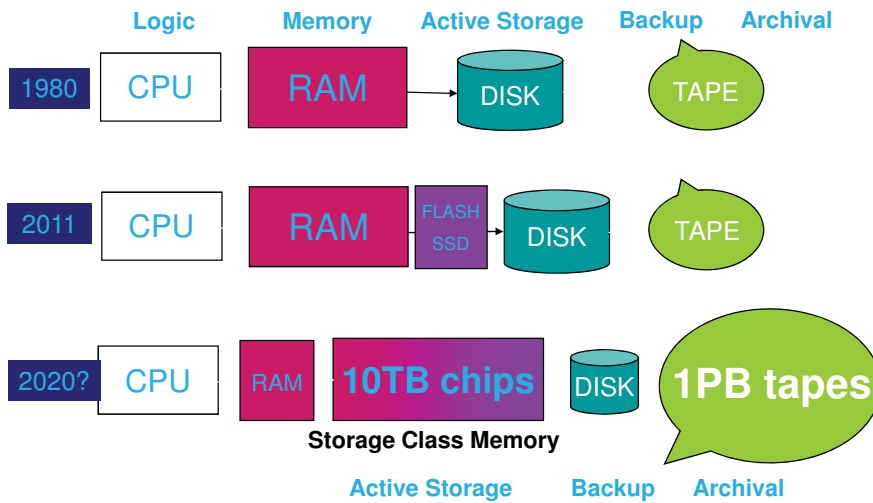
<http://www.zurich.ibm.com/news/11/pcm.html>
<http://www-043.ibm.com/ibm100/us/en/pcme/dram/>

Next Generation SSDs Storage Class Memory (SCM)



Ausblick: Evolution of Memory/Storage Stack

IBM



60

© IBM Corporation, 2012


Solid State Information Systems

IBM

- By 2020, Storage Class Memory should revolutionize data centers

Bandwidth Driven 2020 Storage System: 300 TB/s


Floor Space



322 Racks

3,650 Square Feet

Power



6


© IBM Corporation, 2012

Solid State Information Systems

- By 2020, Storage Class Memory should revolutionize data centers


Bandwidth Driven 2020 Storage System: 300 TB/s

Floor Space



6 Rack 85 Square Feet

Power




Solid State Information Systems

- By 2020, Storage Class Memory should revolutionize data centers


Transaction Driven 2020 Storage System: 500 MOP/s

Floor Space



1,250 Racks 13,996 Square Feet

Power




Solid State Information Systems

- By 2020, Storage Class Memory should revolutionize data centers


Transaction Driven 2020 Storage System: 500 MOP/s

Floor Space



1 Rack 11 Square Feet

Power



IBM Journal of Research and Development on Storage Technologies and Systems
 Volume 52, Combined Issues 4 and 5
 Online: <http://www.research.ibm.com/journal/rds2-45.html>
 Hardcopy Order number: G322-0255-0



- An architecture for storage-hosted application extensions
- An overview of candidate device technologies for Storage-Class Memory
- Archive storage system design for long-term storage of massive amounts of data
- Flash Memories: Successes and Challenges
- Galapagos: Model-Driven Discovery of End-to-End Application-Storage Relationships in Distributed Systems
- Glamor: an architecture for file system federation
- Intelligent Planners for Storage Provisioning and Disaster Recovery
- Phase-change random access memory --- a scalable technology
- Preservation DataStores: New Storage Paradigm for Preservation Environments
- Probe-based ultra-high density storage technology
- Scale Out File Services -- the new Area of NAS
- Scaling tape-recording areal densities to 100 Gbit/in²
- Storage-class memory: the next storage system technology
- The ANSI T10 object-based storage standard and current implementations
- The Evolution of Storage Management: Transforming raw data into information
- Transition-metal oxide based resistance-change memories
- Undetected Disk Errors in RAID Arrays

Improve People Productivity

67 © IBM Corporation, 2012

IBM

Simplified Storage Management

Simplified Administration

- Reduces number of tasks
- Allows for advanced options

Common User Interface

- Based on XIV GUI

New Volume

Select a Preset

Generic Thin Provision Mirror Thin Mirror

Select a Pool

Primary Pool: mdiskgrp0 Edit

Select Names and Sizes

Volume Name: SQL_Database Size: 100 GB +

Summary: 1 volume, 100.0 GB, 4.8 TB free in pool

Advanced... Create Create and Map to Host Cancel

68 © IBM Corporation, 2012

Storwize V7000 Unified - Integrates Block and File Storage

- Integrates file and block applications in a single storage system
- Adds supports for file protocols - NFS, CIFS, FTP, HTTPS, SCP
- Active Cloud Engine
 - Policy based file management
 - File replication and snapshots for business continuity



XIV Mobile Dashboard – iPhone app

- Login
- System view
- For Demo Mode
 - ▶ IP addr: demo
 - ▶ user: demo
 - ▶ password: demo
- Download
 - ▶ <http://itunes.apple.com/us/app/ibm-xiv-mobile-dashboard-for/id503500546?mt=8>



XIV Mobile Dashboard



71

XIV Mobile Dashboard – iPad app

- Download
 - <http://itunes.apple.com/us/app/ibm-xiv-mobile-dashboard/id465595012?mt=8>
- Info
 - <http://aussiestorageblog.wordpress.com/2011/10/06/ibm-xiv-mobile-dashboard-is-in-the-apple-store/>
- Available now
 - Real-time performance statistics
 - Similar to XIVTOP
- Coming Next
 - Alerting & notification
 - Provisioning

Total IOPS: 2054		XIV MN003		Full Redundancy	
Hosts	IOPS	BW(Mbps)	Latency(ms)		
Windows NT6	882	46.2	0.0		
ERP Server1	638	39.0	3.0		
DC0	315	65.0	0.0		
Demo server	219	2.3	1.0		



Total IOPS: 1378		XIV MN003		Full Redundancy	
Hosts	IOPS	BW(Mbps)	Latency(ms)		
Windows NT6	814	46.0	1.2		
ERP Server1	433	43.2	1.2		
DC0	274	39.4	1.2		
Demo server	158	16.4	0.0		

Total IOPS: 2361		XIV MN003		Full Redundancy	
Volumes	IOPS	BW(Mbps)	Latency(ms)		
Home folders	863	103.0	0.0		
Home files	791	138.0	0.0		
Backup	641	27.7	3.0		
Demo volume1	276	13.0	0.0		

72



Store Less Data

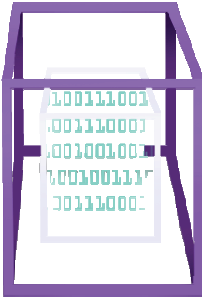
73

© IBM Corporation, 2012

IBM

IBM Real-time Compression Appliances Shrink Active Data

- Shrinks active data up to **80%** without degrading performance
- The **only** compression solution on the market that supports **active data**
- **Improve performance** of I/O-bound applications
- Helps keep more data online for **improved analytics and decision making**
- IBM Real-time Compression Appliances
 - 35 Patents granted or pending
 - Heterogeneous NAS target support




74

© IBM Corporation, 2012

ProtectTIER Deduplication Optimizes Backups and Restores

- Unmatched Efficiency:
Shrink backup data on disk up to **25 to 1**

- Unmatched Performance:
2000MB/sec Ingest Rate
2800MB/sec Restoration



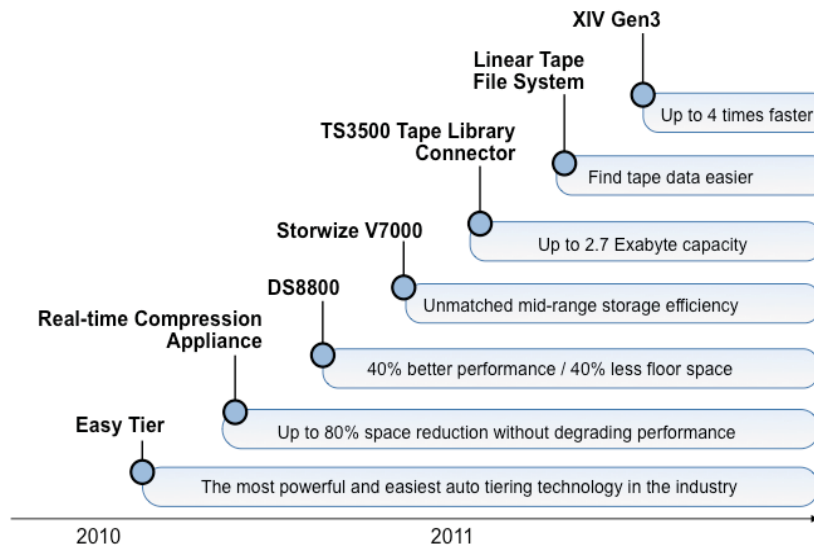
IBM ProtecTIER™ TS7610 Appliance
priced and packaged for mid-sized
environments

"Our backup window shrank immediately. Now, the typical restore request takes five minutes."

Karl Barth,
Director of Systems
and Operations



IBM Is Setting The Pace In Storage Efficiency



Thank You

Thank You