


IBM TSM User Forum

IBM Tivoli Storage Manager
Trends und Kundenreferenzen



Neues zu Tape, Disk und SVC


Manuel Schweiger
IT Specialist





Agenda (Full)


- **Disk**
 - DS8870
 - XIV GEN3 SSD Upgrade
- **Storwize Family**
 - V3700
 - SVC/V7000 v6.4 Code
- **Tape**
 - LTO6





Agenda


- **Storwize Family**
 - V3700
 - SVC/N7000 v6.4 Code
- **Tape**
 - LTO6





DS8870

DS8000 Innovation and Continuity


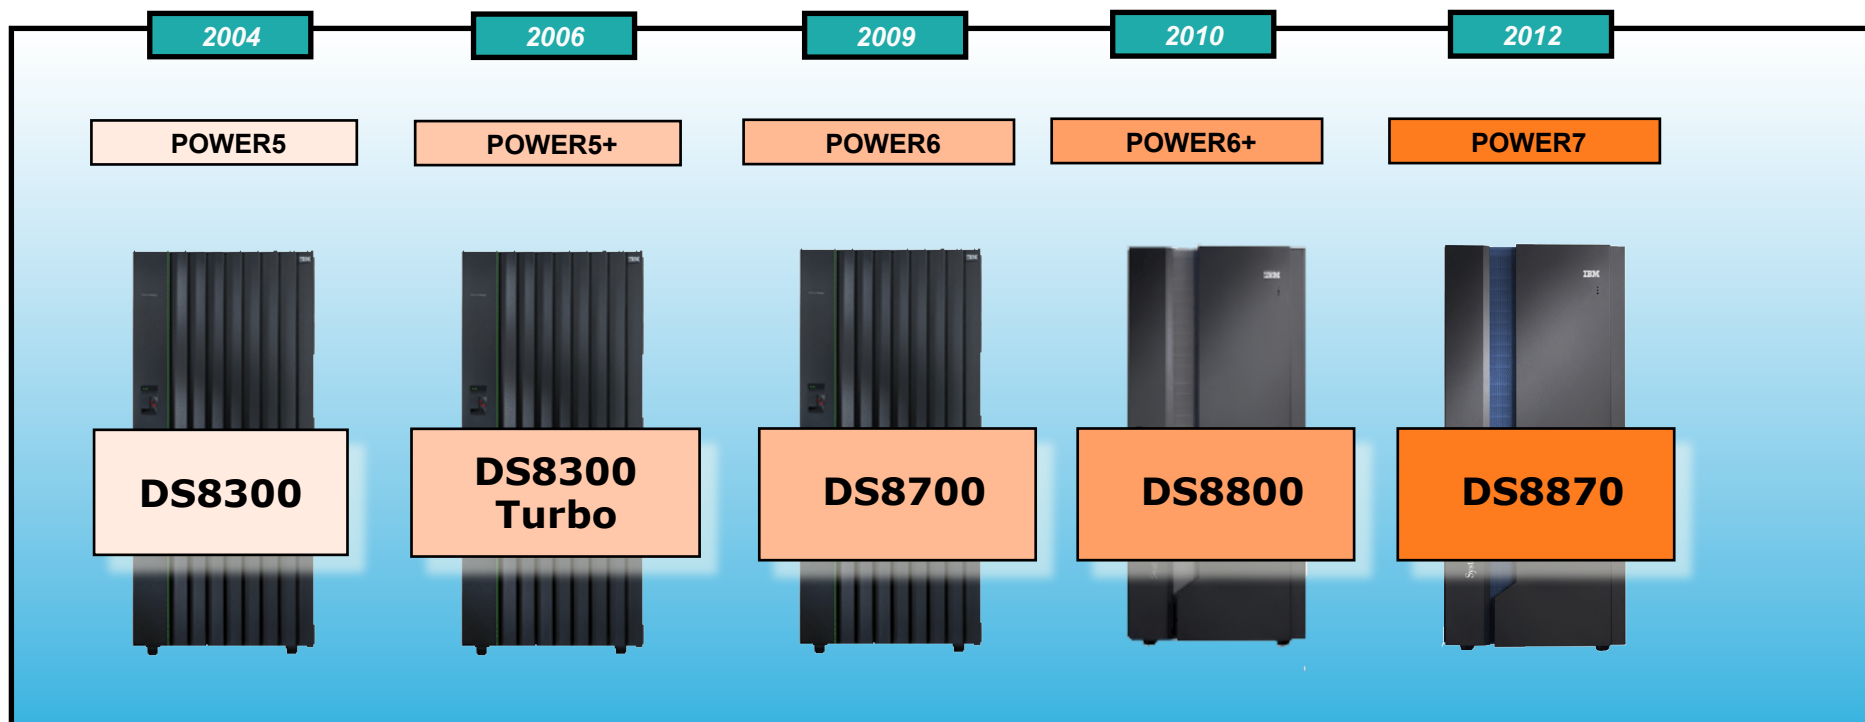


IBM TSM User Forum



5th-generation DS8000 enterprise disk system

- Building on a market-proven, reliable code base!
- 94% of the same proven microcode



- Designed for over 5-9's availability natively
- Designed for over 6-9's availability when DS8000 with Metro Mirror is combined with GDPS/PPRC HyperSwap

IBM TSM User Forum

New hardware delivers up to 3x performance boost



- IBM POWER7-based controllers
- Processor cores that scale from 2-16 per controller
 - Cores aligned with cache configuration
- Cache that scales from 16-1024 GB
- Redesigned power supply for higher energy efficiency
- Dedicated space for new Ultra SSD I/O Drawer (SoD in

June)

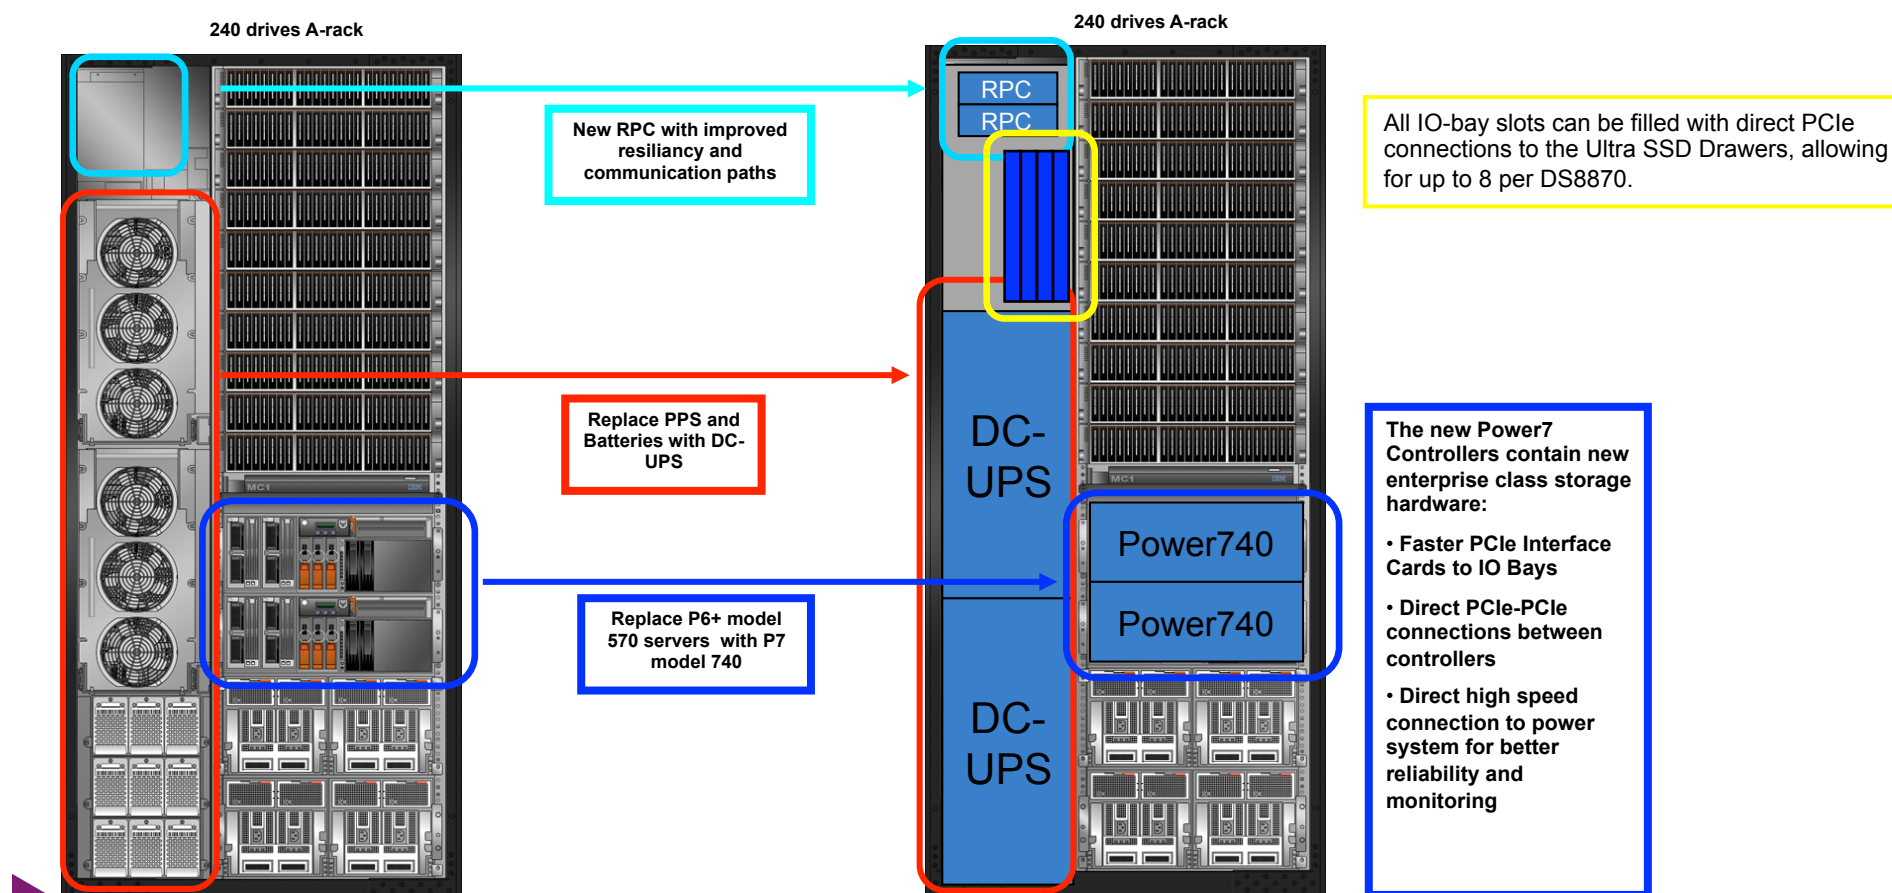


IBM TSM User Forum




Rack changes DS8800 to DS8870

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





Extraordinary performance for enterprise applications



	DS8300 (R4.3) P5+ 4-way	DS8700 (R5) P6 4-way	DS8800 (R6) P6+ 4-way	DS8870 (R7) P7 16-way	<i>Increase 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

	Increase vs. DS8300	Increase vs. DS8700
Seq. Read (GB/s)	5.4x	2.2x
Seq. Write (GB/s)	5.0x	2.3x
DB z/OS (K IOPS)	3.9x	3.2x
DB Open (K IOPS)	3.3x	2.9x




Official Storage Performance Council (SPC) results



- **SPC-1 throughput of 451,082 IOPS**
 - #1 result for single, enterprise-class all-HDD system
 - 67% faster than HDS VSP

- **SPC-2 throughput of 15,424 MB/s**
 - #1 result overall
 - 6% faster than prior #1 result
 - 17% faster than HDS VSP
 - 59% faster than DS8800

Note: DS8870 results will be published on Oct. 3

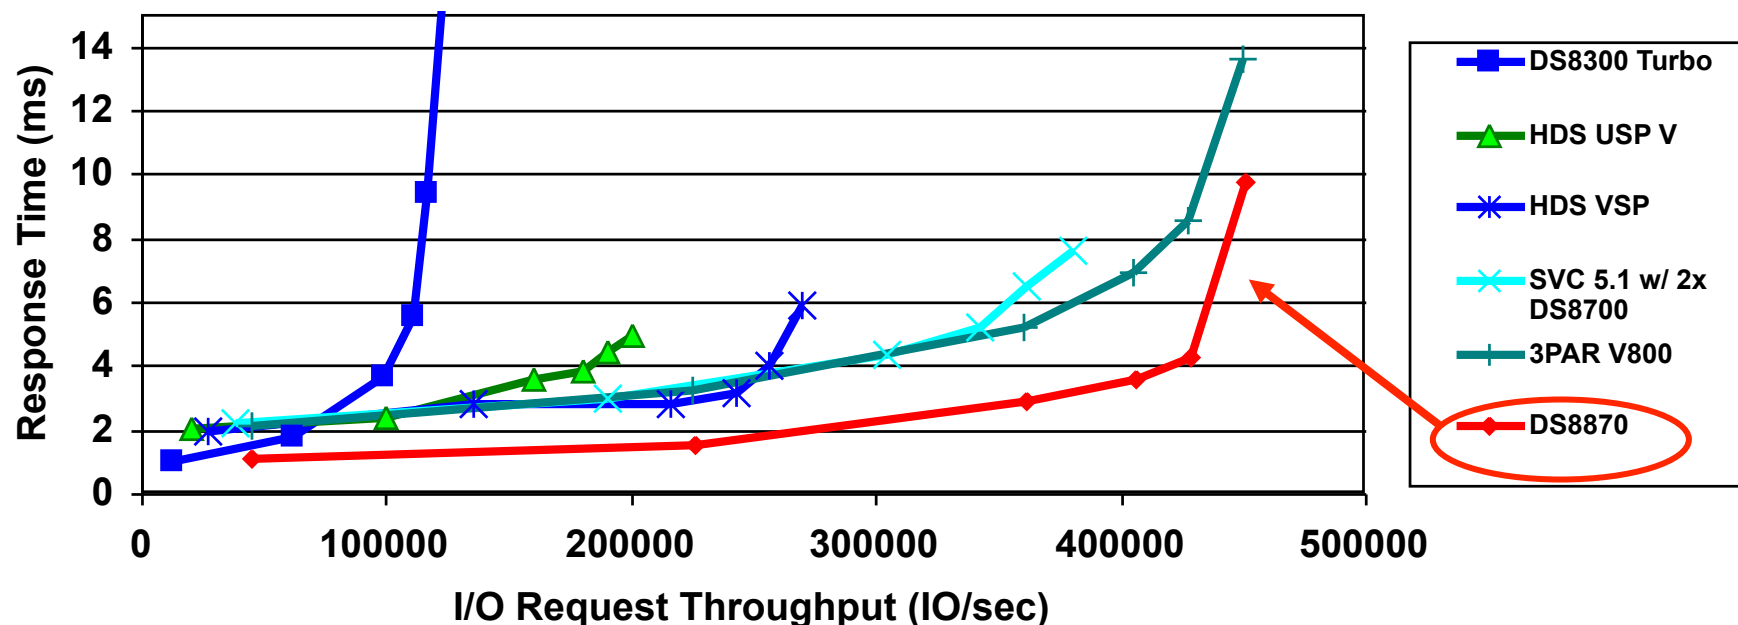




SPC-1 comparison vs competition*

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 x 73GB 15K hdds, RAID10, 256 GB Cache
- HDS USP V: 1024 x 146GB 15K hdds, RAID1, 256 GB Cache
- HDS VSP: 1152 x 146GB 15K hdds, RAID1, 512GB Cache+512GB Cache Flash
- 3PAR V800: 1920 x 300GB 15K hdds, RAID1, 8nodes, 768GB Cache
- SVC+2xDS8700: each DS8700 : 1024 x 146GB 15K hdds, RAID10, 384GB Cache
- DS8870: 1536 x 146GB 15K hdds, RAID10, 32x 8Gb FCP, 1024GB Cache

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

Energy consumption comparison



DS8700 with 1024 drives

- Base frame: 6.8kW
- Exp frame: 7.1kW
- Exp frame: 6.1kW
- Exp frame: 6.1kW
- Exp frame: 3.1kW

TOTAL: 29.2kW

DS8800 with 1536 drives

- Base frame: 7.5kW
- Exp frame: 6.2kW
- Exp frame: 6.3k W
- Exp frame: 6.3kW

TOTAL: 26.3kW

DS8870 with 1536 drives


- Base frame: 6.0kW
- Exp frame: 5.6kW
- Exp frame 5.8kW
- Exp frame: 5.8kW

TOTAL:

23.2 kW

Base frame is 20% more energy efficient

Higher efficiency rating positioned to meet emerging ENERGY STAR requirements





Full Disk Encryption is now standard

- FDE options across all drive tiers
- Same performance as standard drives
- Standards-based key management software supports key management interoperability protocol (KMIP)
- Key manager supports both disk and tape
- New European Union Privacy Directive makes breaches very costly
- Encryption is the least expensive data disposal technique
- Supports Easy Tier environments
- Over a thousand IBM encryption disk and tape solutions deployed worldwide

“Do not wait for an event-driven reaction to secure your data. Proactively securing your data will help ensure against a worst-case scenario and a financial impact that is likely to far surpass that of the data security purchase itself.”

John Monroe
Gartner Hype Cycle for Storage, July 2012



“Within five years, all HDDs and SSDs will be shipped preloaded with some kind of industry-standard FDE technology” – Gartner Hype Cycle for Storage, July

* Requires deployment of Tivoli Key Lifecycle Manager or IBM Security Key Lifecycle Manager



IBM TSM User Forum

Business Class and Enterprise Class configuration options



- List price for Business Class configuration option is roughly the same as DS8800
 - 30% price premium for Business Class upgrade to Enterprise Class
 - Small servers can be mounted in the Business Class frame to optimize space efficiency
- Enterprise Class offers better price/performance than previous models
- Non-disruptive upgrade from smallest to largest configuration



Model	Processor	Physical Capacity (max.)	Disk Drives (max.)	Memory	Host Adapters (max.)	9xE Attach
Business Class						
961	2-core	216 TB	144	16/32	4	0
Enterprise Class						
961	4-core	360 TB	240	64	8	0
961	8-core	2,304 TB	1536	128/256	16	0-3
961	16-core	2,304 TB	1536	512/1024	16	0-3
First Expansion Frame						
96E	N/A	504 TB	336	N/A	8	N/A
Second/Third Expansion Frame						
96E	N/A	720 TB	480	N/A	N/A	N/A





XIV SSD Cache Acceleration

Evolution of the Revolution



IBM TSM User Forum



XIV SSD Caching: Design Goals

Invisible


- Allow wide deployment of SSDs without increasing IT complexity
- Allow customers to focus on capacity and service management, not disk technologies

Reliable

- At least 99.999% availability, beyond the reliability of today's off-the-shelf SSD technology

Affordable




- Wide SSD deployment is currently limited by cost
- For maximum business impact SSD acceleration should be feasible for the entire data center



XIV SSD Solution



- **High capacity SSDs used as secondary cache**
 - 400GB SSD device per module
 - 6TB of total cache per rack
 - Housed in PCI caddy in rear of module
 - Added to all installed modules
- **No tuning necessary**
 - Immediate gain for all applications
 - Optional selective volume activation
- **Available to all Gen3 systems**
 - Order with SSD from factory
 - Field non-disruptive upgrade






XIV Caching Technology Highlights

Consistent, high hit ratio

Through efficient cache management


- Highly granular 4KB slots and aggressive 64KB pre-fetch
- Automatically bypass sequential reads larger than 64KB
- Massive parallelism of cache extension



Reliable use of commodity SSDs

Through patented SSD write-handling

- SSD are written to using log structured writes, tuned for caching
- Minimized wear-leveling, garbage collection and write amplification
- Ensure consistent, predictable performance and longevity of SSDs


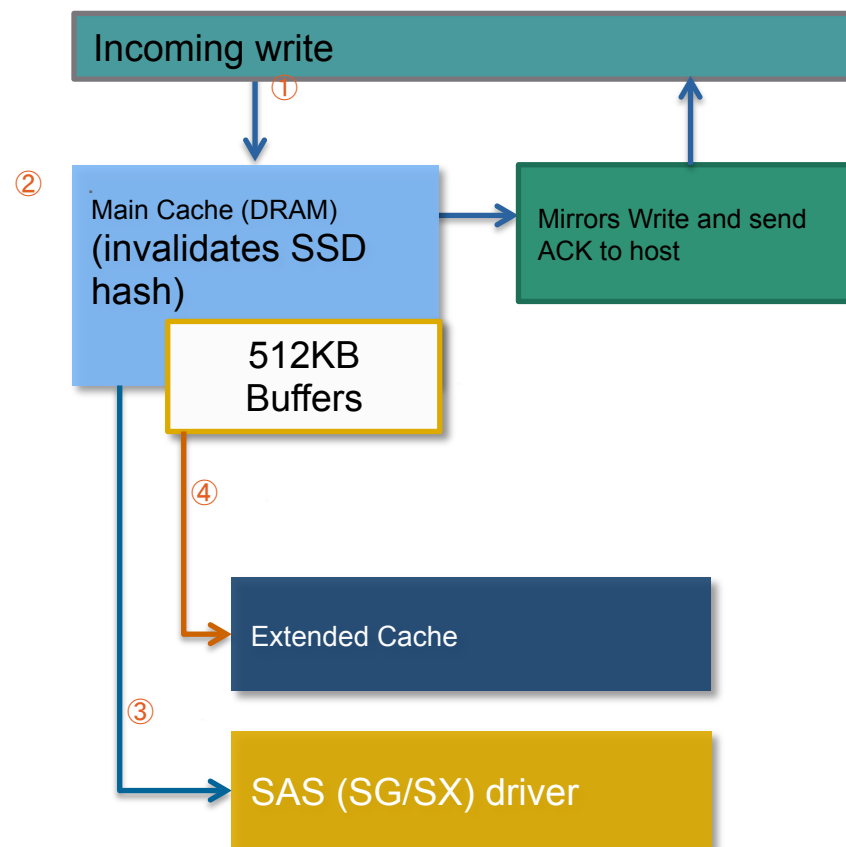




XIV SSD, How it works: Disk Writes With SSD Cache enabled

Write Operation:

- ① All writes go to DRAM (Main Cache)
- ② If SSD has a copy of the data being modified, then the corresponding hash table entry on SSD is invalidated to prevent subsequent hits from the extended cache
- ③ As writes are de-staged from Main Cache they are asynchronously staged on to the SSD buffer.
- ④ Once buffer is full, data is written to SSD and hash table is again valid.





XIV SSD, How it works: Random Read

Read Miss < 64K


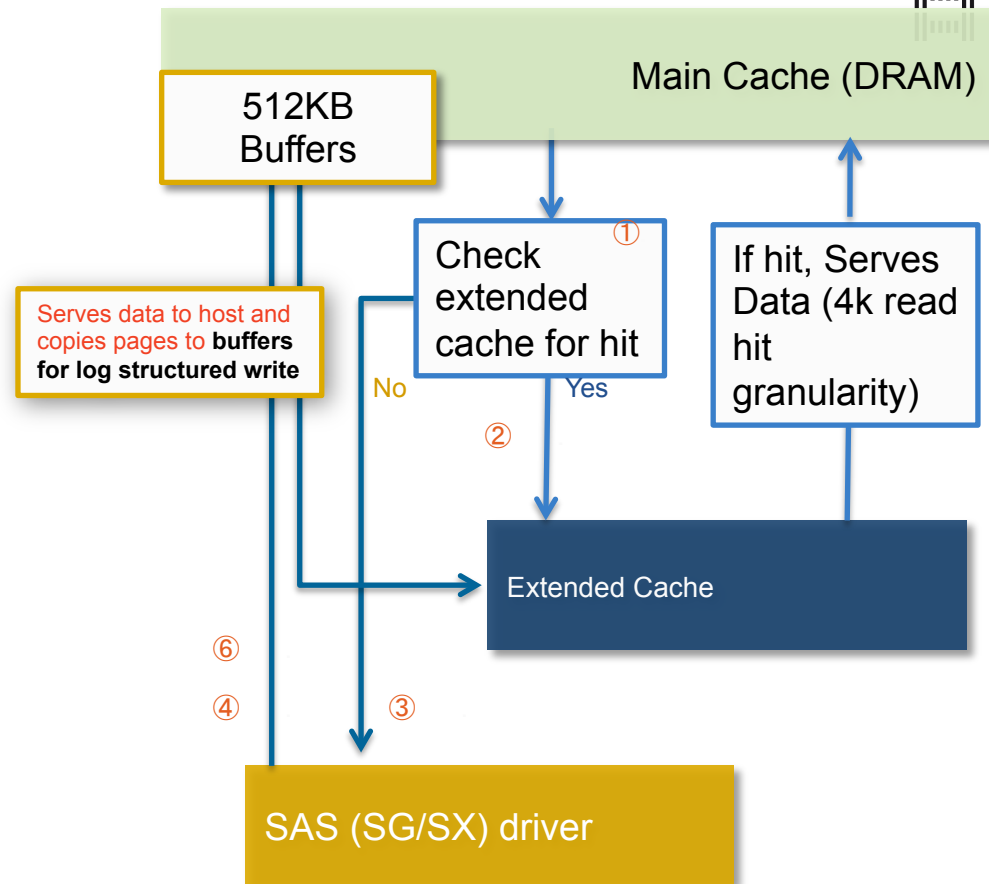
Read Miss Operation:

- ① Checks for hit on extended cache
- ② If hit, re-routes to the SSD
- ③ Else, *unmodified* request to HDD
- ④ Data copied onto 512KB buffer page
- ⑤ When buffer filled up, it is **sequentially** destaged to the SSD (*log-structured*).
- ⑥ Sequential IO detection bypasses SSD to avoid cache poisoning

Other cache operations:

– Sequential pre-fetch goes direct to SAS disks

– Reads > 64K bypass the Flash cache and are fetched directly from disk






Outstanding Performance Across Applications (Gen2 Vs Gen3)


Microsoft Exchange Mailboxes

- Requires latency under 20 ms
- ESRP-Storage test




Oracle Data Warehouse (IOPS)

- Oracle DHW Workload



SAS Business Analytics Workload


- Analytics reports created



Simulated via Swingbench load generator


Microsoft HyperV Simulation (IOPS)

- 200GB dataset
- 60% write activity




Transaction Processing (IOPS)

- Mixed read / write workload




Sequential Reads (MB/sec)

- System Bandwidth




File and Print Services (IOPS)

- Mixed file size workload
- XIV Gen3 also had 50% lower latency



Sequential Writes (MB/sec)

- System Bandwidth



IBM TSM User Forum

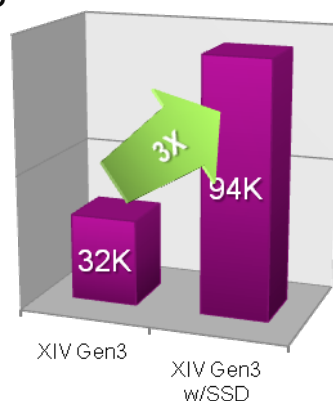
Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage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



Outstanding Applications Performance with SSD Caching


DB2 Brokerage (IOPS)

- Heavy Random Brokerage
- 90/10, Mixed block IO
- 84% Random Read Miss




Websphere Datastore (IOPS)

- Web 2.0 OLTP Workload
- 80/20/4k




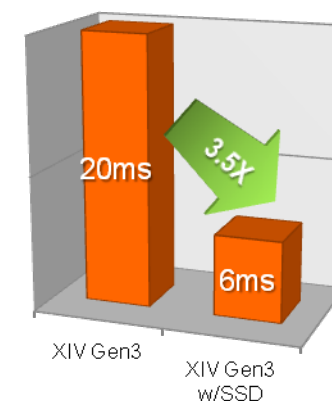
Core ERP (IOPS)

- CRM and Financial DB Workload
- 70/30/8k



Medical Record App Server (RT)

- Healthcare EMR Workload
- 100% random IO




IBM TSM User Forum

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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 improvements equivalent to the performance ratios stated here.



V3700

*The newest member
to the Storwize Family*



IBM TSM User Forum



Meet The Storwize Family

A comprehensive range of fully virtualized storage systems



Storwize V3500
(China only)



Storwize V3700



Storwize V7000

New!



Storwize V7000
Unified



Flex System
V7000

New!

Announcing Nov 13




SAN Volume
Controller

- One code base on all platforms
- One set of functions (selectively licensed)
- One user interface
- One CLI



Efficient Entry Storwize System

- **Enclosures** contain up to twelve LFF (3.5") or twenty-four SFF (2.5") drives in just 2U
- **Control enclosure:** dual active-active controllers and drives
 - 8GB Cache (4GB per controller) standard; optional upgrade to 16GB
 - 1Gb iSCSI host ports standard plus optional 8Gb FC or 10Gb iSCSI/FCoE or SAS¹ host interface upgrades
- **Expansion enclosure:**
 - Up to four expansion enclosures attach to one control enclosure
 - Up to 120 SFF or 60 LFF drives for up to 180TB raw capacity
 - Mix drive sizes and HDD/SSD in enclosure
 - Dual-active expansion canisters
- **RAID levels** 0, 1, 5, 6, and 10
- **Redundant, hot-swappable components** (i.e., power supplies and fans)
- **AC power** (110 to 240 V)
- **Optional Software¹:** Easy Tier, Remote Mirroring, FlashCopy upgrade (2040 targets)



IBM TSM User Forum

¹ Statement of Direction. IBM Plans are subject to change without notice



Storwize V3700: Graphical User Interface

Based on Common IBM Storage GUI

molly > Home > Overview

Suggested Tasks

System recommendations

6 Internal Drives → 1 MDisk → 1 Pool → 5 Volumes

0 Fibre Channel Hosts
0 iSCSI Hosts

Visual task flow guidance

Overview

Watch e-Learning: Overview

Integrated video instruction

Link to more information if needed

System status always available

Allocated: 5.00 GB / 1.10 TB (0%)

Running Tasks (0)

Health Status





Storwize V3700: High Performance

High performance in a midrange disk system

Almost 300,000 IOPS

Up to 21,000 IOPS for “database-like” 70/30 R/W workload

RAID-5 Performance Comparison

V7000 – 240x 15K V3700 – 120x 15K DS3500 – 96x 15K	Storwize V7000	Storwize V3700	DS3500
Cache Reads (IOPs)	900,000	280,000	140,000
Disk Reads (IOPs)	110,000	40,000	30,000
Disk Writes (IOPs)	24,000	9,000	7,500
Disk 70/30 (IOPs)	55,000	21,000	19,000
Cache Reads (MB/s)	5,500	3,300	2,500
Disk Reads (MB/s)	4,000	1,990	1,950
Disk Writes (MB/s) (cache mirrored)	2,200	650	500




Storwize Product Comparison



	Storwize V3700	Storwize V7000
Management SW	Storwize V3700 machine code	Storwize V7000 software
Control Canisters	Single ¹ or Dual	Dual
Standard Host Interface	6Gb SAS ¹ , 1Gb iSCSI	8Gb FC 1Gb iSCSI
Optional Host Interface	8Gb FC or 10Gb iSCSI / FCoE	10Gb iSCSI / FCoE
Cache (per controller)	4GB or 8GB	8GB
Expansion Unit Support	Up to 4 expansion units – up to 120 disks	Up to 9 expansion units – up to 240 disks per control enclosure; 960 disks per clustered system
Thin Provisioning	Yes, Standard	Yes, Standard
Data Migration	Yes, Standard (one-way)	Yes, Standard
FlashCopy (up to 64 targets)	Yes, Standard	N/A (see below)
FlashCopy upgrade	Yes, Optional ¹ (2040 images)	Standard (4,096 images)
Remote Mirroring (FC)	Yes, Optional ¹	Yes, Optional
Easy Tier	Yes, Optional ¹	Yes, Standard
System Clustering	NA	Yes, Standard
External Virtualization	NA	Yes, Optional
Compression	NA	Yes, Optional
NAS	NA	Yes, Optional


¹ Statement of Direction. IBM Plans are subject to change without notice





SVC / V7000 v6.4

Increase your space with RACE




IBM TSM User Forum



Compression Basics – Lempel Ziv

ASJKDFHAS **ABCDEF** QRIUFSDFWWEIRU
CMNXSDFKWIOEUZXCMZXNVSFJSDFL
SJCXCSKLRJHWEIOUOCZXCVMKMSNDK
SFSJZXM23NB33KJK1J1HJGJHHJ1VFH
JGFHJ1GHJG23GJ123 DHJKWE
IORUWOEIRIXCVLXVJLASDFSDFLSDG
RERMNJDFJKGDCJERTYUIRDJKGHDKJ
TEHTREUITYEUIDWOSIOSDFWEOIRU
 KDFHSDFJHWEIORWERYWEF
UWYEIRUWERYXDKJFHSWETR5DFGC
VBNA1SFSKLJFSKLDFJSLKDFJSLKDFJS
KLDFJSDLFKJSDFKLJSDFKLJSDLFJSDF
KLSJDFKLSJDEI4SDFDFDFDSDSDSDFSD
SDFSDFDSDSDFSSDF2834HKJH

- Detects repetitions in the data
- Replaces portions of the data with references to matching data






Compression Basics – Sliding Window

```
ASJKDFHASJHWRETORIUFSDFWWEIRU
CMNXSDFKWIOEUZXCMZXNVSFJSDFL
SJCXABCDEFHWEIOUOCZXCVMKSNDK
SFSJZXM23NB33KJK1J1HJGJHHJ1VFH
JGFHJ1GHJG23GJ123ABCDEFDHJKWE
IORUWGEIRIXCVLXVJLASDFSDFLSDG
RERMNJDFJKGDGJERTYUIRDJKGHDKJ
TEHTREUITYEUIDWOSIOSDFWEOIRU
KDFHSDFJHWEIORWERYWEFUWYEIR
UWERYXDKJFHSWETR5DFGCVBNA1SF
SKLJFSKLDJSLKDFJSLKDFJSLKDFJSD
LFKJSDFKLJSDFKLJSDABCDEF_LFJSDF
KLSJDFKLSJDEI4SDFDFDFDSDSDSDFSD
SDFSDFDSDSDFSSDF2834HKJH
```

- Repetitions can be detected only within the sliding window history
- Common sliding window size – 32K
- Repetitions outside the window can not be referenced

Window size limit

- Memory footprint required to hold history in searchable manner
- Processing power required for searching larger history window
- Size of pointer needed to reference small repetition






Compression Basics – Huffman Coding

IAIIIABIBIIDMIBBBMIIIIIBBBMADKLE
BBIIIBBBIIIIAJHJKJDAMMMIIIIIB
BBIIIIISDFDIOIIIIIIABBBBBMIIIMM
MIIIIIDD FMMMMIIGFMMAEERTGMM
DFMMIIIIIIAAABBBBBBIIIIUIIDIII
IIIDDGDBBBBBBMMMEERMBBIIBMI

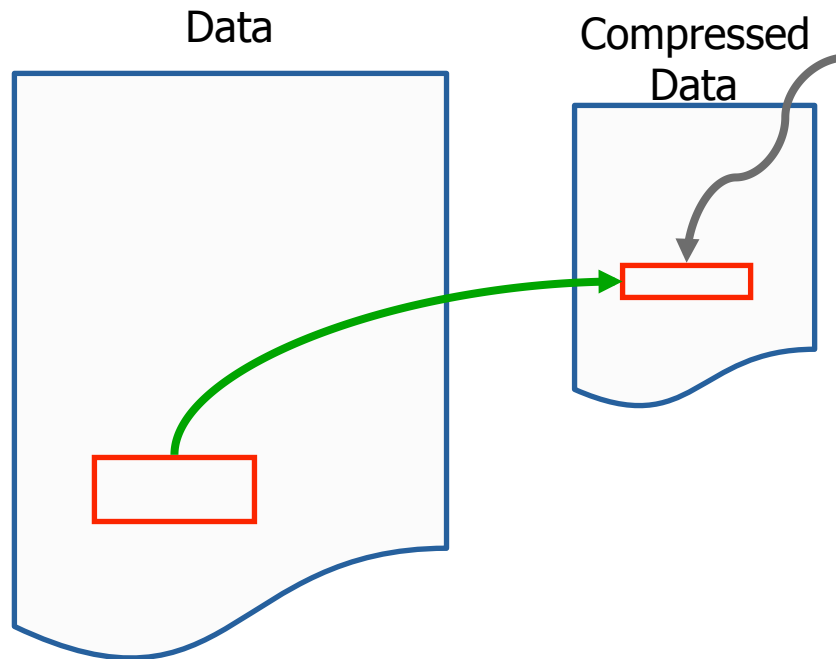
- Detects common characters in data
- Represents common character using less bits

Common Char	Bit Representation
I	0
B	10
M	110
Other	111 + 8 Bits






Compression – Random Access



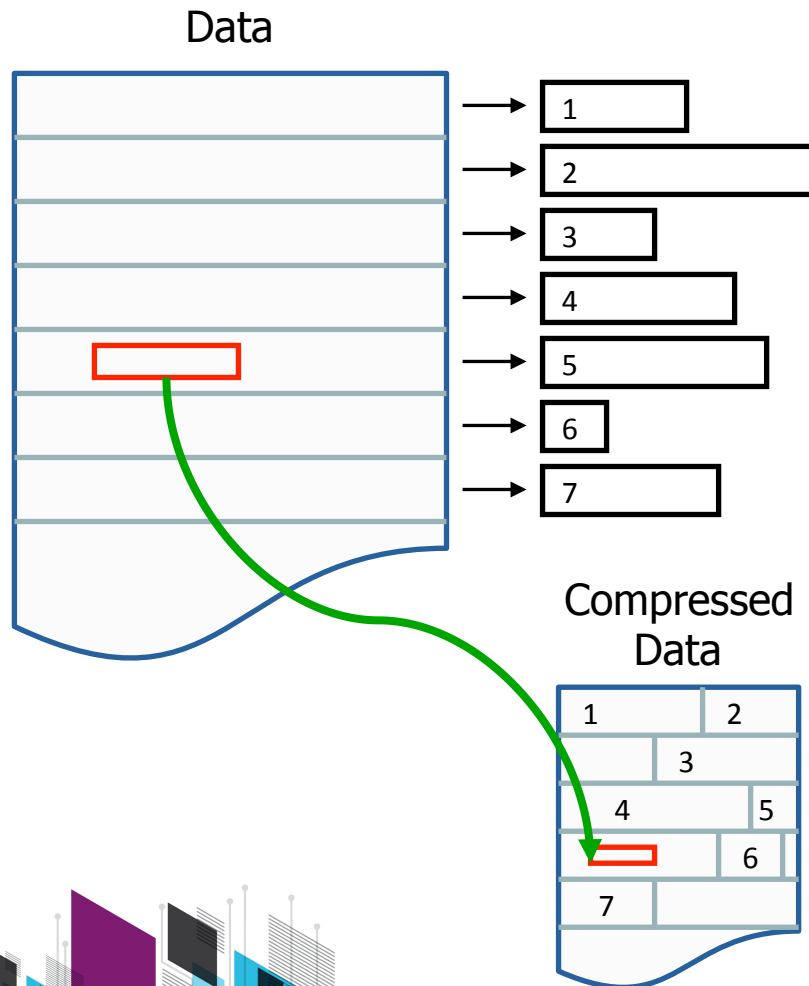
- Data is dependent on preceding data due to the nature of compression
- In order to read from a specific location, all data before it has to be decompressed
- To write to a specific location, all data after it has to be recompressed as well

- Not effective for large files or block devices
- Compression implementations do not support random access






Compression – Random Access Chunks




- Break original data to fixed chunks
- Each chunk is compressed and decompressed independently
- Enables some random access to the data (reads, not writes)
- Large chunks – Heavy I/O penalty
 - 4KB update = 1MB read + 1MB write
- Small chunks – Poor compression
- Variable output
 - Data fragmentation
 - Lower performance over time
 - Lower compression ratio over time





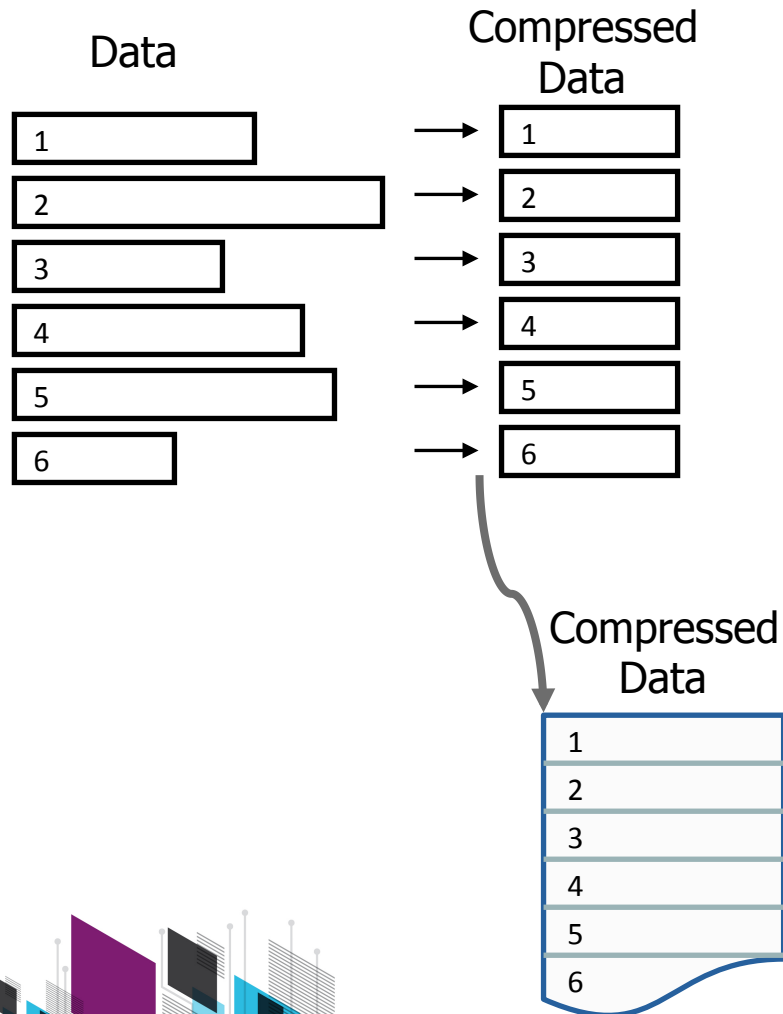
IBM RACE Technology



IBM TSM User Forum




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

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






Temporal Compression



- Applications make multiple updates to data
- Traditional and post-process compression uses fixed-sized chunks and compresses each update based on its location on a volume
- RACE compression acts on data that is written around the same time (“temporal locality”) not according to location
- Temporal locality is more related to real application operations
- RACE takes advantage of the structure of the data and its application level relations
- Better compression efficiency and performance




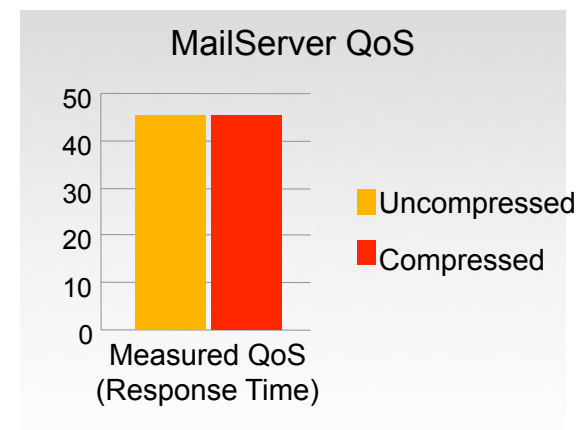





Compression Without Compromise – VMware

VMware VMmark Performance Benchmarks

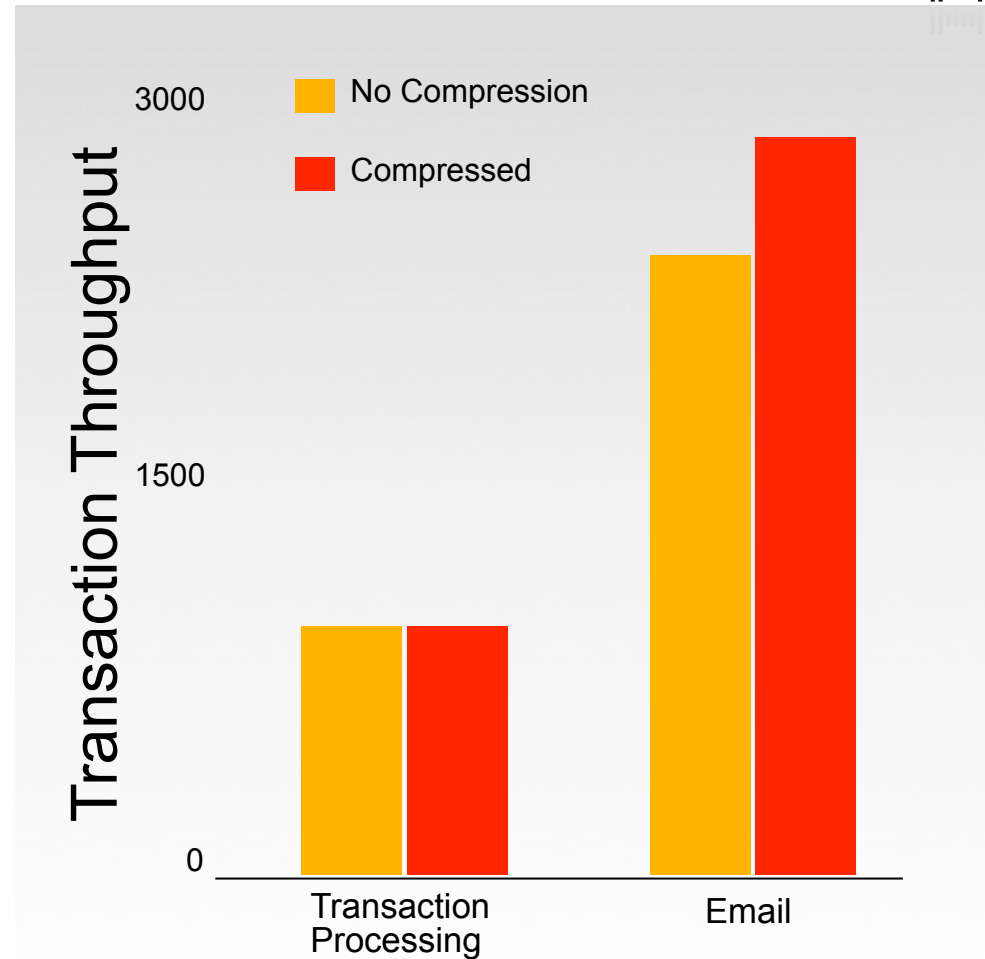
- Storwize V7000 with Real-time Compression delivers up to 4x compression while maintaining VMware and Application performance






Compression Without Compromise


- Storwize V7000 with Real-time Compression delivers up to 5x compression while maintaining or improving application business throughput
- Business throughput / Transactional IOPS – higher is better



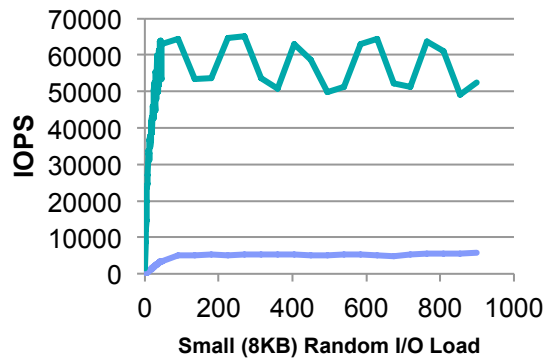
Source: IBM lab measurements, 96/48-drive configurations




Beta Customer database testing results SVC virtualizing 6-node XIV Gen2 configuration




Read IOPS




Read Latency




Read MBPS





Write IOPS



Write Latency



Write MBPS




❖ **Orion (Oracle I/O Calibration Tool)** is a standalone tool for calibrating the I/O performance for storage systems that are intended to be used for Oracle databases. The calibration results are useful for understanding the performance capabilities of a storage system, either to uncover issues that would impact the performance of an Oracle database or to size a new database installation



Expected Compression Rates

Databases	50-80%
Server Virtualization	45-70%
Collaboration Data	30-75%
Engineering Data	50-80%
E-mail	30-60%

- IBM *Comprestimator* tool should be used to evaluate expected compression benefits in existing environments





Does it work ?

- Just test for yourself (riskfree!)... only thing you will need is v6.4 installed
- You have 45 days with eval license – just activate the license.
- **Important**... only activate compression if you are below 25% CPU on V7000/CF8 or below 50% CPU on CG8 Nodes

SVC_IBMAT > Settings > General

General

- Date and Time
- Licensing**
- Upgrade Software
- GUI Preferences

Update License

Use this panel to set the licensing values for the IBM System Storage SAN Volume Controller. If you are sharing the total authorized capacities across multiple systems, enter only the capacities you wish to use on this system. The sum of the capacities across all systems must not exceed your authorized capacities.

Select Your License

- SAN Volume Controller Standard Edition (5639-VC)
- SAN Volume Controller Entry Edition (5639-VW)

Set Options

Option	Used	Set	
Virtualization Limit	1.65 TB	<input type="text" value="100"/>	TB
FlashCopy Limit	0.03 TB	<input type="text" value="100"/>	TB
Global and Metro Mirror Limit	0 TB	<input type="text" value="100"/>	TB
Real-time Compression Limit	0.52 TB	<input type="text" value="100"/>	TB





What compression ratio will I achieve?

- If you do not have a V7000 (new customer)
 - Then you can use the “Comprestimator” Tool if you use RHEL4+, ESX4+, Windows 2003+ or AIX 6.1+ hosts, HPUX, ...

1. Download Comprestimator from IBM website
2. Upload Comprestimator to Server
 - e.g. “scp Comprestimator root@IP_of_ESXServer:/
3. Change file to executable
 - “chmod +x Comprestimator”

4. Identify Volumes

```
esxcli storage core device list | grep Dev
Device Type: Direct-Access
Devfs Path: /vmfs/devices/disks/naa.600508e00000000005a02d8f29810870d
Device Type: Direct-Access
Devfs Path: /vmfs/devices/disks/naa.600507680183851a1000000000000004e
```

5. Run Comprestimator


```
~ # ./Comprestimator -d /vmfs/devices/disks/naa.600507680183851a1000000000000004e
Start time: 18/06/2012 07:27:55
Device name: /vmfs/devices/disks/naa.600507680183851a1000000000000004e
Device size: 250.0 GB
Number of processes: 10
Exhaustive: no
```

Sample#	Device Name	Size(GB)	Compressed Size(GB)	Total Savings(GB)	Total Savings(%)	With Provisioning Savings (%)	Compression Savings (%)	Compression Accuracy Range(%)
3627	/vmfs/devices/disks/naa.600507680183851a1000000000000004e	250.0	65.2	184.8	73.9%	1.3%	73.6%	4.6%



What compression ratio will I achieve?

- If you have a V7000 it is even more easy!
- You can use the internal mechanism “vdisk mirror” to check.
- No change to your original volume.
- Just an additional volume copy which is not used for production workload.
- Only writes are mirrored there... reads are still done from original volume!



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

Select a Pool

Filter...

Name	Status	Free Capa...	Capacity
DS3400_Quorum	Online	19.50 GB	19.50 GB
DS4700_EasyTier	Online	863.00 GB	1.63 TB

Volume	Status	Capacity	Usage
VMFS01	Online	250.00 GB	
Copy 0*	Online	250.00 GB	
Copy 1	Online	250.00 GB	89.21% (149.66 MB)




Will I have any performance impact?

- Once the vdisk mirror is created the new copy ("Copy 1") is used as write copy only – reads are still done from original volume.
- To test performance we just change "Copy 1" to primary.
- After that "Copy 1" will be used for read&write, "Copy 0" as write only.
- If there is any performance impact, just change back to "Copy 0" again.

VMFS01	Online	250.00 GB	
Copy 0*	Online	250.00 GB	
Copy 1	Online	250.00 GB	75.60% (186.10 GB)

VMFS00	Compressed	250.00 GB
Copy 0	Make Primary	250.00 GB
Copy 1*	Split into New Volume	250.00 GB
VMFS01	Validate Volume Copies	250.00 GB
Copy 0*	Delete this Copy	250.00 GB
Copy 1		250.00 GB

VMFS01	Online	250.00 GB	
Copy 0	Online	250.00 GB	
Copy 1*	Online	250.00 GB	75.60% (186.10 GB)



* Copy numbering is based on the assumption that no vdisk mirror was used. If already used then check for correct numbering!



After the Test

- If you are not satisfied with the compression feature
 - Just delete "Copy 1" (new volume) and remove the license
- If you are satisfied and want to use compression from now on
 - Officially license the compression feature and delete "Copy 0" (old volume)

Volume Name	Status	Size
VMFS00	Online	250.00 GB
Copy 0*	Online	250.00 GB
Copy 1	Online	250.00 GB
VMFS01		250.00 GB
Copy 0		250.00 GB
Copy 1*		250.00 GB
VMFS_Test		50.00 GB

Showing 17 volumes |

Warning


Do you want to remove copy 1 from volume VMFS00?

OK Cancel



LTO6

Newest updates with LTO




IBM TSM User Forum



Announcing new IBM LTO Generation 6 Tape Drives and Library support

- New IBM LTO Generation 6 integration into IBM Tape Portfolio
 - Automation TS2900, TS3100, TS3200, TS3310, TS3500
 - Preview of IBM Ultrium 6 Media
 - External Tape Drives (TS2260, TS2360)
- Features
 - **2.5 TB Native Physical Capacity** (6.25 TB with 2.5:1 compression) on LTO 6 media
 - Up to **160 MB/s** native data transfer rate
 - Supports **4 partitions**
 - Full High and Half High drive form factor
 - Encryption capable
 - Media partitioning
 - SAS at 6 Gbit/sec
 - FC at 8 GB/sec
 - Capability to Read/Write Ultrium 5 and read Ultrium 4 cartridges
 - Ethernet based statistics and drive status HTML page
- IBM exclusives:
 - IBM Linear Tape File System
 - Available on IBM tape family





The LTO Roadmap



	Generation 2	Generation 3	Generation 4	Generation 5	Next Generation 6
Native Capacity	200 GB	400 GB	800 GB	1.5 TB	2.5 TB
Native Transfer Rate	20 - 40 MB/sec	40 - 80 MB/sec	80-120 MB/sec	140MB/s	160MB/s
Attachment	FC-2 Ultra-160	FC-4 Ultra-160	FC-4 Ultra-160 SAS 3 Gbit SCSI 320	FC8 SAS 6 Gbit	FC8 SAS 6 Gbit
Form Factor	Full High	Half High Full High	Half High Full High	Half High Full High	Half High Full High


TS3500 LTO6 Drives
 Announced: Wed Oct 3, 2012
 Generally Available: Fri Nov 9, 2012

Other External LTO6 Drives
 Announced: Tues Nov 6, 2012
 Generally Available: Fri Dec 7, 2012

Integrated Drives & Storage Enclosure Drives
 2013

LTO6 Media
 Announced: tbd
 Generally Available: tbd

Plans are subject to change without notice





LTO Cartridge native (uncompressed) capacity, Interchangeability and Maximum tape drive native data transfer speeds

Tape Drive	Operation	Generation 1 cartridge  100GB Native	Generation 2 cartridge  200GB Native	Generation 3 cartridge  400GB Native	Generation 4 cartridge  800GB Native	Generation 5 cartridge  1.5 TB Native	Generation 6 cartridge  2.5 TB Native
Generation 1	Read	15 MB/sec					
	Write						
Generation 2	Read	15 MB/sec+	35 MB/sec				
	Write						
Generation 3 Full high	Read	15 MB/sec+	35 MB/sec+	80 MB/sec			
	Write						
Generation 3 Half High	Read	15 MB/sec+	35 MB/sec+	60 MB/sec			
	Write						
Generation 4	Read		35 MB/sec+	80 MB/sec+	120 MB/sec		
	Write						
Generation 5	Read			80 MB/sec+	120 MB/sec	140 MB/sec	
	Write						
Generation 6	Read				120 MB/sec	140 MB/sec	160 MB/sec
	Write						

Note: LTO6 media is mentioned in the anno letters as a "specification" but the media is not announced yet





Preview: Not announcing until December



New IBM Ultrium 6 2.5 TB Tape Cartridge

- Statement of Direction
 - **IBM intends to introduce tape cartridge media meeting the LTO Generation 6 media specifications.**
 - Testing and Certification is planned to time out at end of November
- IBM Ultrium 6 Media
 - Greater capacity
 - 2.5 TB native physical capacity (6.25 TB with 2.5:1 compression)
 - Barium Ferrite substrate
 - Data Cartridges and WORM Data Cartridges
- IBM Machine Type 3589 Ultrium Tape Cartridge
 - **Model 550** - IBM Ultrium 6 2.5 TB Data Cartridge, labeled with starting volume serial information and optionally packed in individual jewel cases
 - **Model 650** - IBM Ultrium 6 2.5 TB Data Cartridge, packed in individual jewel cases with two nonattached blank labels
 - **Model 570** - IBM Ultrium 6 2.5 TB Data Cartridge, labeled with starting volume serial information and optionally packed in jewel cases, are designed to provide WORM storage on a 2.5 TB capacity cartridge
 - **Model 670** - IBM Ultrium 6 2.5 TB Data Cartridge, packed in jewel cases with two nonattached blank labels, are also designed to provide WORM storage on a 2.5 TB capacity cartridge
- Feature Code Ultrium 6 Media in MT 3580, 3572, 3573
 - **FC 8605 (SEO # 39P1902) Ultrium 6 Data Cartridge (5-Pack)**
 - **FC 8601 Ultrium 6 Data Cartridge (Single) (3580 only)**



IBM TSM User Forum

Planned announcement: December 4th

Planned availability: December 7th

IBM LTFS Library Edition

- System support for IBM Tape Automation



Announcing on November 6, 2012...




New!

LTO 6 generation tape drives support and Windows Server 2008

- TS2900 tape autoloader
- TS3100 tape library
- TS3200 tape library
- TS3310 tape library
- TS3500 tape library




**Preview:
LTO 6 Support in LTFS SM
Not Until December**





LTFS Platform Support


	Single Drive Edition	Library Edition
Supported Tape Systems	TS2250, TS2260, TS2350, TS2360	TS2900, TS3100, TS3200, TS3310, TS3500, TS1140w/TS3500
Supported Operating Systems	<ul style="list-style-type: none">▪RHEL 5.5, 5.6, 5.7, 5.8, 6.1 and 6.2▪SLES 11 SP1 and SP2▪Windows 7 SP1▪Windows Server2008 R2 SP1▪Mac OS X 10.6, 10.7	<ul style="list-style-type: none">▪RHEL 5.5, 5.6, 5.7, 5.8, 6.1 and 6.2▪SLES 11 SP1 and SP2▪Windows Server 2008 R2 SP1
Other components	For details on supported hardware configurations please refer to the System Storage Interoperation Center (SSIC)	For details on supported hardware configurations please refer to the System Storage Interoperation Center (SSIC)






LTFS Use Cases

- ✓ Film digitization and production (Hollywood, etc)
 - Both digitization of existing libraries and production
 - Require archival of all digital masters for 50-100 years
- Archiving of Cultural and Historical data
 - University archives of books, movies, etc.
- ✓ Broadcast
 - Large broadcast archive initiatives at major TV networks (NBC, CBS, ZDF,..)
- Oil & Gas exploration
 - Seismic shots
- ✓ Digital Video Surveillance
 - Opportunity exists for long term tape storage
- ✓ Electronic Discovery
 - Archiving information for Computer Forensics purposes
- ✓ Medical
 - Hi-resolution files like Medical X-Rays or 3D/4D Ultrasound videos
- ✓ IT-Cloud Back-up / Archive



Look for
Lots of data
(Long files)
“Unstructured data”
Data doesn't change
Retention period
Need to exchange
Portability helpful



2012 Tape Technology Roadmap Details from INSIC*



*) Source: <http://www.insic.org/news/2012Roadmap/12roadmap.html>

Parameter/Year	2012	2014	2016	2018	2020	2022		
1. Capacity (TB)	4	8	16	32	64	128	41.42%	per year
2. Data rate per channel (MB/sec)	11.2	13.8	17.0	20.9	25.8	31.8	11.00%	per year
3. Total data rate (MB/sec)	270.0	405.0	607.4	911.0	1366.5	2049.6	22.47%	per year
4. FC Speed Roadmap (MB/sec)*	3200	6400	12800	12800	25600	25600		
5. Number of channels	24	29	36	43	53	64	10.33%	
6. Tape thickness (μm)	6.00	5.53	5.10	4.70	4.33	3.99	-4.00%	per year
7. Data capacity reserve	5.0%	5.0%	5.0%	5.0%	5.0%	5.0%		defect reserve
8. Tape length that is recordable (meters)	867	941	1,021	1,107	1,202	1,304	3.90%	winding reserve
9. Tape length total (meters)	902	979	1,062	1,152	1,250	1,357	4.17%	
10. Track density (TPI)	6,506	9,773	14,787	22,498	34,393	52,791	23.29%	per year
11. Linear bit density (KFCI)**	467	545	635	741	864	1,008	8.00%	per year
12. Areal density (Gbits/inch ²)	3.04	5.32	9.39	16.67	29.73	53.22	33.15%	per year
13. Tape speed (m/sec)	6.4	6.6	6.8	7.0	7.3	7.5	1.63%	
14. Tape width (mm)	12.65	12.65	12.65	12.65	12.65	12.65		
15. ECC and formatting overhead	24.00%	21.84%	19.88%	18.09%	16.47%	14.99%	-4.60%	per year
16. Servo track and layout overhead***	18.46%	16.59%	15.02%	13.67%	12.50%	11.49%	-7.00%	per year
17. Number of passes to write a tape	110	138	175	222	283	361		
18. Time to fill a tape (minutes)	247	329	439	585	781	1,041	15.47%	
19. Number of passes to end of media life	30,300	33,406	36,830	40,605	44,767	49,356	5.00%	per year
20. Number of data tracks	2,642	4,060	6,258	9,673	14,988	23,272	24.31%	per year
21. Bit Aspect Ratio (BAR)	88	67	51	38	29	22	-13.12%	per year

* See Reference [6]

** Defined as the 1T KFCI where T is the data cell length

*** On non-capacity reserve overhead only

Disclaimer: This technology roadmap is not a product roadmap and does not imply that any tape drive from any vendor may adhere to this list or timeframe.




IBM TSM User Forum

IBM Tivoli Storage Manager
Trends und Kundenreferenzen



Questions?



IBM TSM User Forum

IBM Tivoli Storage Manager
Trends und Kundenreferenzen



Thank You!

