Marianne Hammer
IBM Corporation
Poughkeepsie, New York
hammerm@us.ibm.com

Session 8861: What's new in z/OS Performance
Share 116
Anaheim, CA
02/28/2011

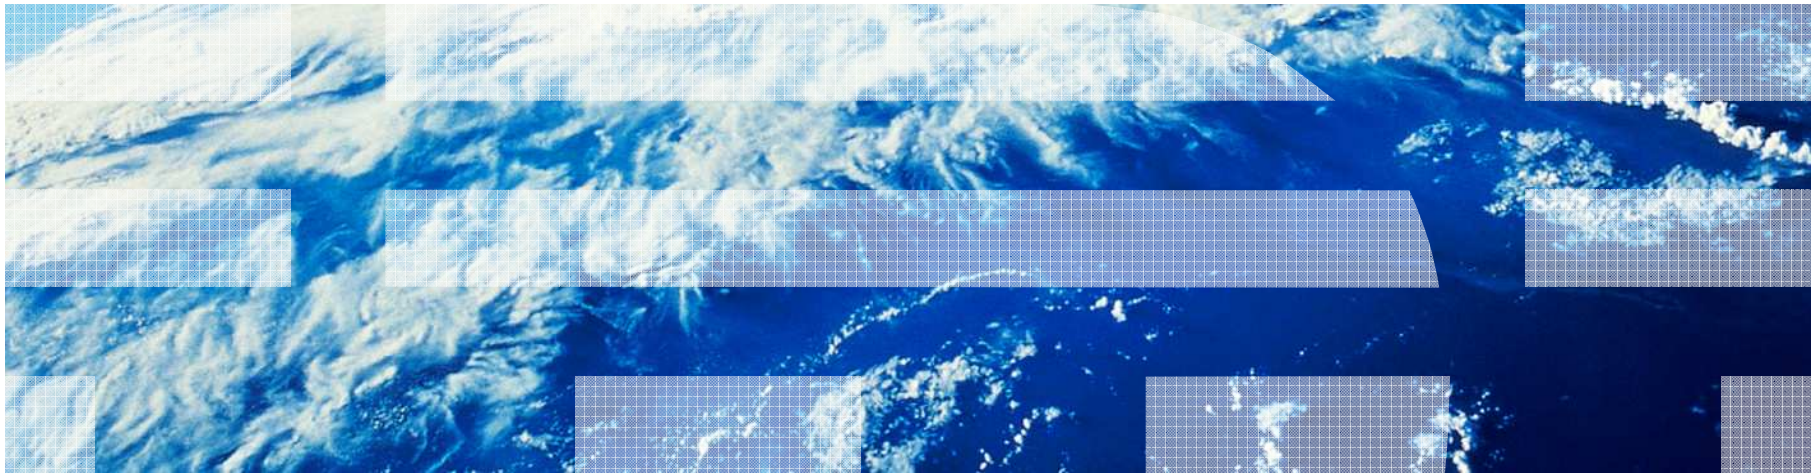

Trademarks

IBM Corporation 2009

- IBM, the IBM logo and ibm.com are registered trademarks, and other company, product or service names may be trademarks or service marks of International Business Machines Corporation in the United States, other countries, or both. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml
- Adobe, the Adobe logo, PostScript, the PostScript logo, Cell Broadband Engine, Intel, the Intel logo, Intel Inside, the Intel Inside logo, Intel Centrino, the Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, IT Infrastructure Library, ITIL, Java and all Java-based trademarks, Linux, Microsoft, Windows, Windows NT, the Windows logo, and UNIX are trademarks or service marks of others as described under "Special attributions" at: http://www.ibm.com/legal/copytrade.shtml#section-special
- Other company, product and service names may be trademarks or service marks of others.
- References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

Notes

- The performance observations are based on measurements and projections using standard IBM benchmarks in a controlled environment.
- The actual throughput that a 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.

What's new in z/OS Performance - Agenda

- z/OS V1.11 Performance:
  - Release performance content – BCP focus
  - Results from LSPR workloads
  - SADMP performance
  - zAAP on zIIP
  - zFS Sysplex enhancements

- z/OS V1.12 Performance:
  - Release performance content – BCP focus
  - Results from LSPR workloads
  - Dispatcher timeslices option
  - MTTR – Open Data sets
  - SVC dump performance

z/OS Performance Metrics and Tools

- Some of the performance metrics used in this presentation:
  - **ETR:** External Throughput rate – Number of transactions ended per second (IMS trans, CICS trans, WAS trans, TSO trans, batch jobs)
  - **ITR:** Internal Throughput Rate – Number of transactions per CPU busy second
  - **PL:** Pathlength – Instructions per transaction
- Performance data from RMF
- Internal version of HIS tool (customer Instrumentation)
- Other internal performance tools
- Release-to-release performance numbers in:
  - WSC Softcap tool
  - WSC zSoftcap tool – coming spring 2011. New user interface.

Methodology for testing and comparing two z/OS releases

- Compare new release to previous release
  - Previous release (base for comparison): GA-level of z/OS
  - When using performance workloads: measure systems at high CPU utilization – 90%
- Both releases tested on:
  - Same processor HW configuration
  - Same I/O configuration
  - Same workload setup and tuning parameters
  - Same number of simulated users
  - Same database layout and size
  - Same SW stack level (DB2, IMS, CICS, WAS, Java)
- Everything is the same except for the SYSRES that is used for IPL, and any PARMLIB and catalog changes required for new release.

z/OS Release Performance Goals

- Performance improvement goals for z/OS V1.11 and V1.12:
  - 5% performance improvements for the z/OS component measured as the average result of the four LSPR workloads on a 32w single image.
    z/OS component represents on average 25% of total CPU usage for the LSPR workloads.
    This translates into an average of 1.25% overall system improvement compared to previous release.

  - VSCR goal: Reduce 31-bit common storage by 4 MB for the release (compared to previous release). The 4 MB goal is based on a 10% reduction of the perceived 40 MB z/OS contribution to 31-bit common (based on the LSPR workloads), and is measured against customer use of common virtual storage.

CPU usage LSPR workloads

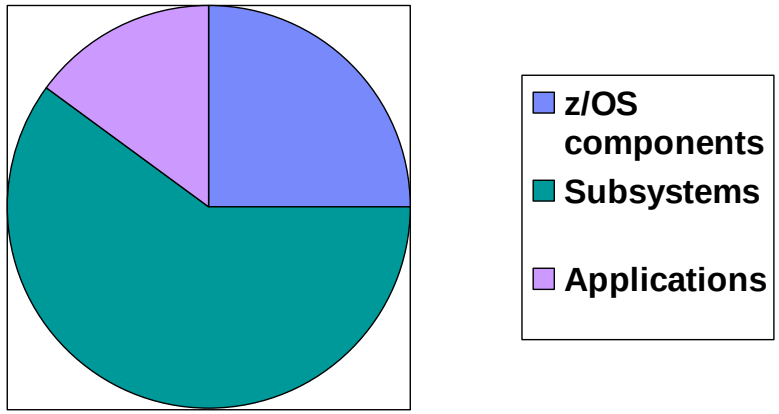

Module CPU usage based on instrumentation data (CPU samples)

z/OS V1.11 Performance Content - BCP

- VLF retrieve performance – prefetch:
  - STCMH instruction (STore Characters under Mask High)
  - Avoid cache misses by preloading data
  - Workload with 4-7% of CPU usage in VLF retrieve: 50-60% reduction in VLF CPU usage
- LE performance – mutex locking
  - Optimized code to provide faster locking and reduce CPU usage
    - 70%-80% reduction in CPU usage in the mutex locking module (LSPR workloads)
    - 90% reduction in ptd_mutex_lock elapsed time
    - 87% reduction in ptd_mutex_ulck elapsed time
- Hiperdispatch, other misc. BCP performance enhancements and recompile of top 200 csects
- NFS: change to use pause/release for serialization
  - 40% increase in throughput (NFS ops/sec) and 28% improvement in ITR
- Unicode performance improvements:
  - character conversion:
    - 31-bit: +47% for 8 byte data size, and +4% for 4K data size.
  - case conversion:
    - 31-bit: +57% for 8 byte data size, and +15% for 4K data size.
    - 64-bit: +50% for 8 byte data size, and +12% for 4K data size.

z/OS V1.11 Performance Results – LSPR Workloads

Workload	1w on z10	32w on z10 with HD=YES	64w on z10 with HD=YES	Comments
WASDB (WAS and DB2 OLTP)	0.7%	0.8%	9.7%	LE mutex. Large effect from hiperdispatch changes on 64w
OLTP-W (WAS, CICS, DB2 OLTP)	0.2%	4.4%	n/a	Mostly from LE mutex and hiperdispatch changes
OLTP-T (IMS OLTP)	6.0%	1.6%	1.5%	Mostly from VLF prefetch
CB-L (batch)	0.2%	0%	n/a	
Average	**1.8%**	**1.7%**	**5.7%**	

z/OS V1.11 SADMP Performance

- SADMP performance: MIDAWS support, pathlength reductions, data prefetch:
  - Up to 4x increase in dump data rate compared to V1.9 tests on z9 and same amount of real storage dumped (about 35GB). Improvements coming from the z/OS changes and z10 speed (vs z9).
  - V1.9 tests with 1 and 4 SADMP volumes dumped about 20-22GB of storage .

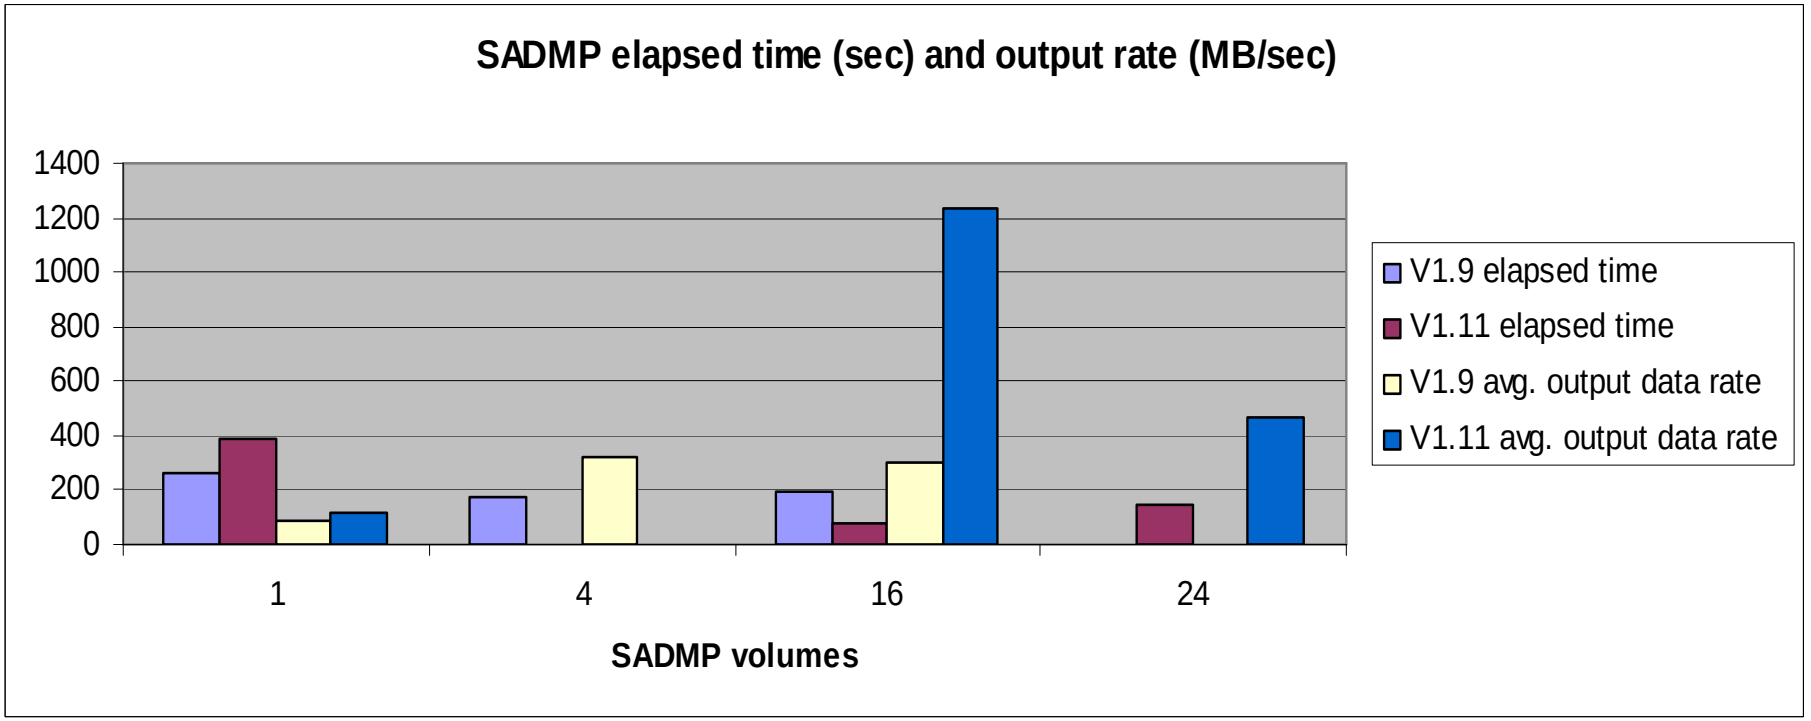

z/OS V1.11 zAAP on zIIP

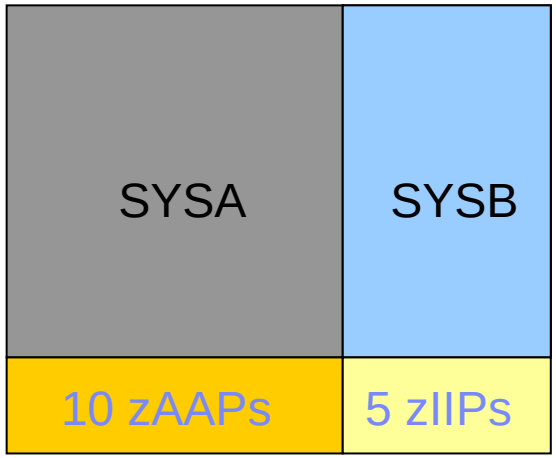

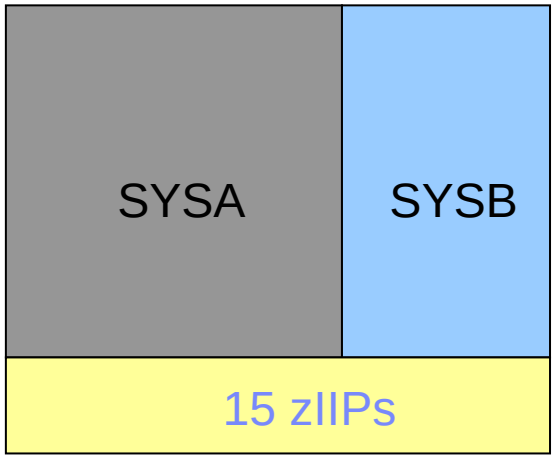

- Specified in IEASYSxx: ZAAPZIIP=NO|YES (default=YES since V1.11)
- No performance impact when running zAAPs on zIIPs
- Allows for more configuration flexibility
- Check RMF Workload Activity report for zAAP and zIIP usage:
  - AAPCP: percent of CPU for zAAP work running on a CP
  - IIPCP: percent of CPU for zIIP work running on a CP
  - AAP: percent of CPU for zAAP work
  - IIP: percent of CPU for zIIP work

```
      ---APPL %---
CP     201.62
AAPCP    0.00
IIPCP    0.44

AAP       N/A
IIP    314.42
```

zFS File System Sharing Read-only

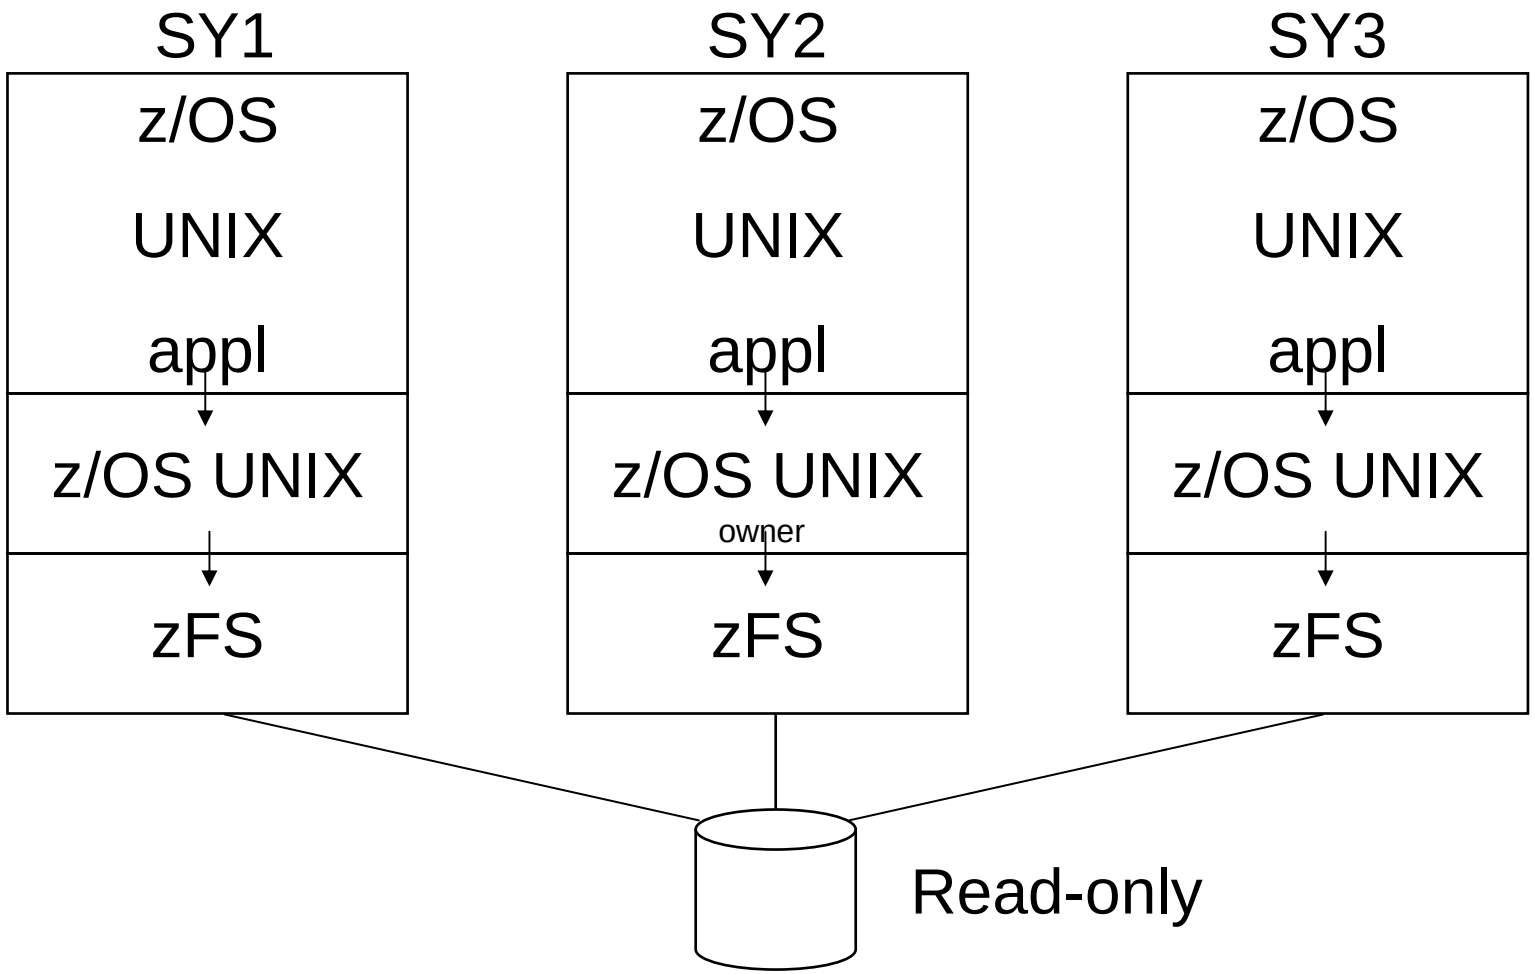

Sharing a zFS Filesystem Read-Write prior to z/OS V1.11

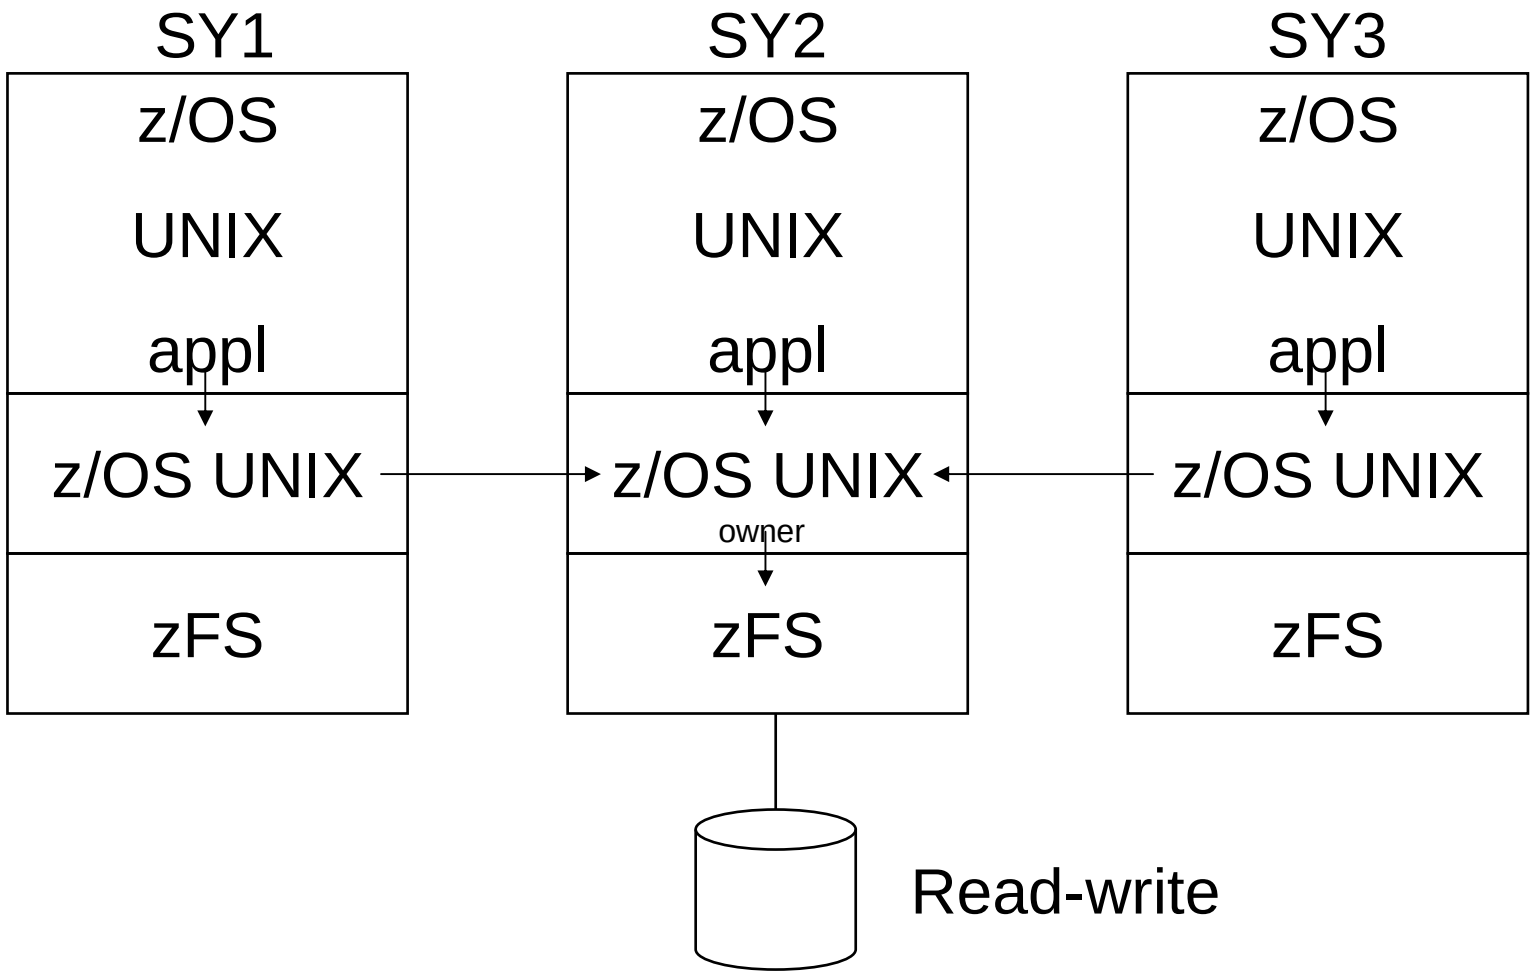

Sharing a zFS Filesystem Read-Write in V.11 with apar OA29619

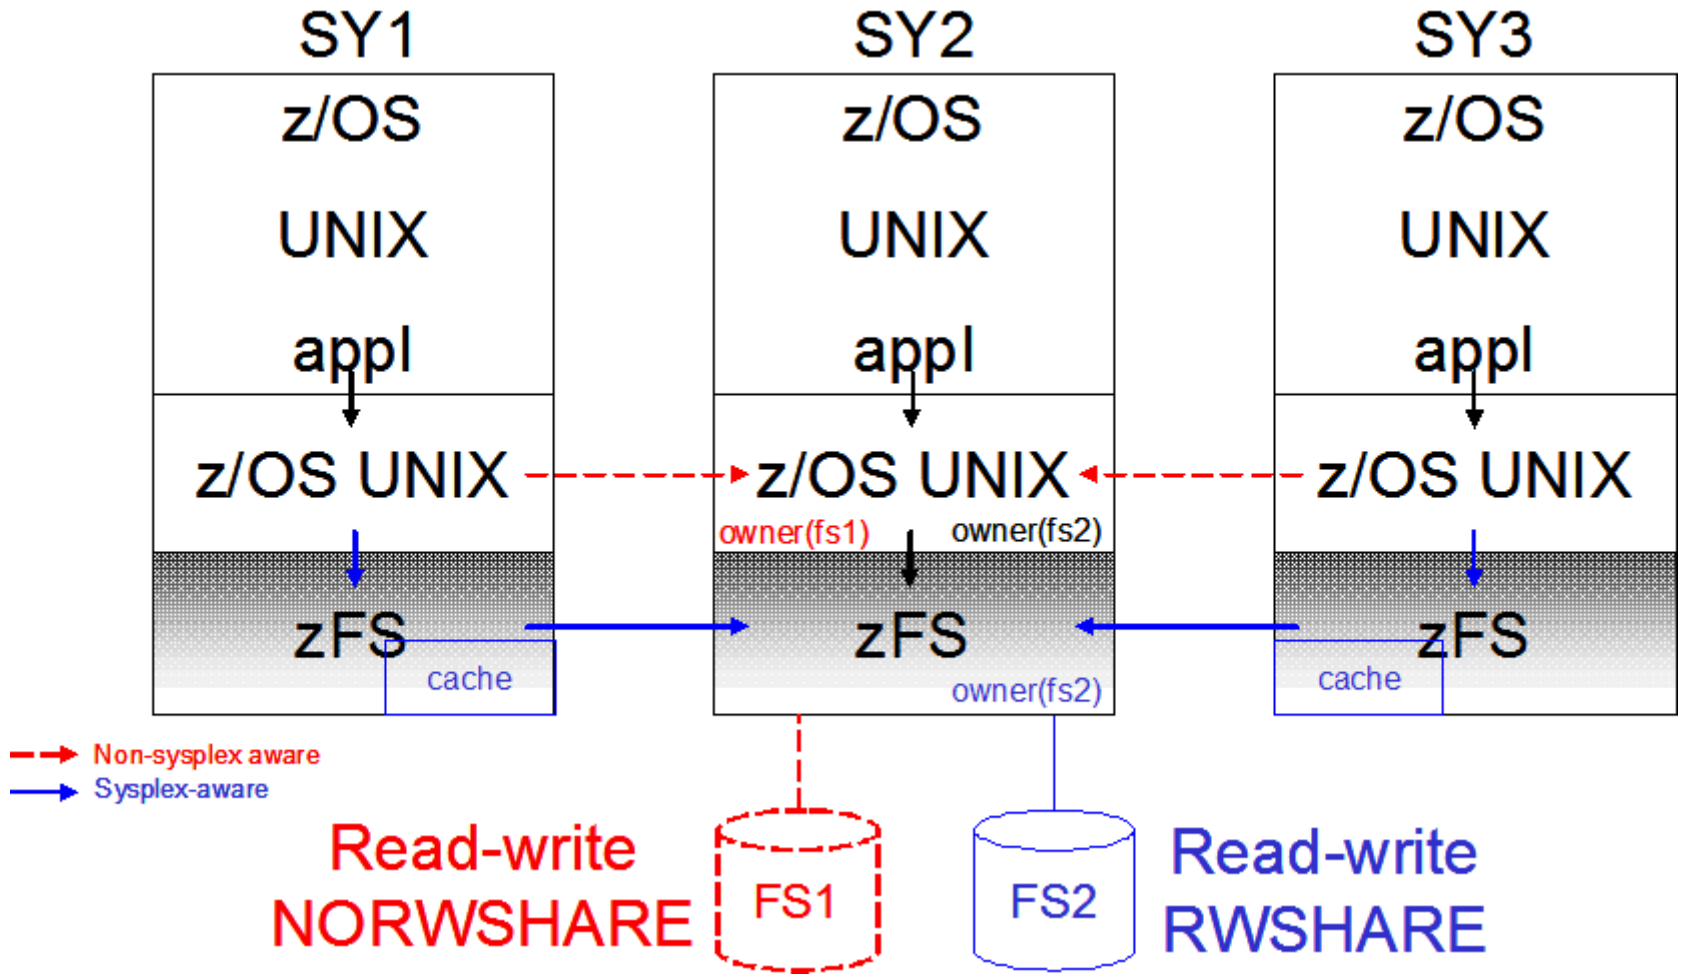

IOEPRMxx: SYSPLEX=FILESYS
BPXPRMxx: SYSPLEX(YES)

z/OS V1.12 Performance Content - BCP

- RRS zero interest commit optimization for WAS transactions
- LE code optimization
  - STCMH for heap pools
  - realloc() optimization
    - String manipulation intensive applications that use the realloc() function to grow the storage for the strings. Most benefits for applications that are heavy users of this function. Measured 20% improvement on z10.
- Overall system performance impact:
  - Changes to supervisor pause/release serialization – benefits to heavy users of pause/release (e.g. DB2)
  - Nucleus large page support
  - Misc hot cache items, and component re-compiles with higher optimization levels for top 200 csects
- Timeslices option for discretionary CPU-intenstive work
- MTTR Allocation improvements in a DB2 environment
- SVC dump performance improvements

z/OS V1.12 performance Results – LSPR Workloads

Workload	1w on z10	12w on z10 with HD=YES	32w on z10 with HD=YES	64w on z10 with HD=YES	Comments
WASDB (WAS and DB2 OLTP)	1.9%	1.5%	1.2%	2.5%	Mostly RRS and misc BCP
OLTP-W (WAS, CICS, DB2 OLTP)	-0.2%	0.1%	4.0%	3.4%	Pause-rel optimization and misc BCP
OLTP-T (IMS OLTP)	0.2%	1.7%	4.6%	-	Mostly VTAM OA33084 and misc BCP
CB-L (batch)	0.2%	-	1.2%	-	Misc BCP
Average	**0.5%**	**1.1%**	**2.7%**	**3.0%**	

z/OS V1.12 Dispatcher Timeslices Option

- The data shown are the result of a benchmark with more threads than CPs and very low pause times
  - only affects discretionary work that is CPU-intensive as determined by significant mean time to wait (MTTW)
- Tests done with default value and IEAOPTxx TIMESLICES=64 option
- 64 Bit multi-JVM had a 33 percent improvement at 7 Threads in each JVM
- TIMESLICES=255 at best had only a 2 percent improvement over TIMESLICES=64 in the Multi-JVM case

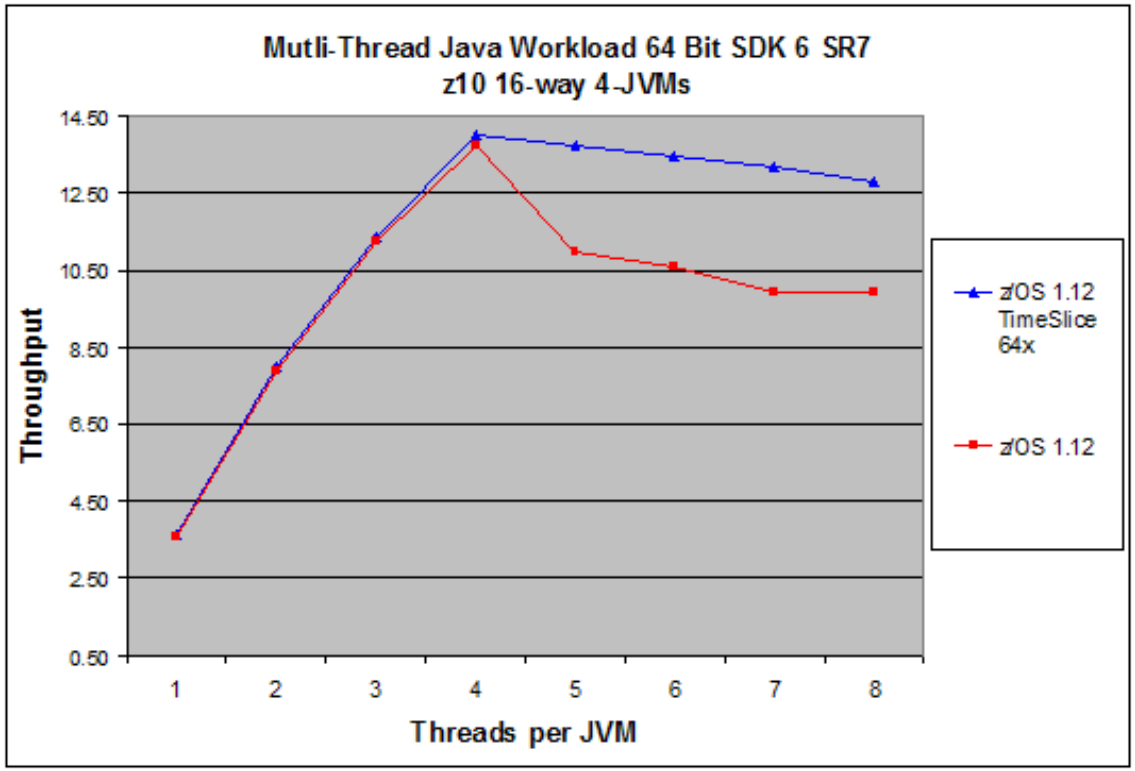

z/OS V1.12 Open DB2 Data Sets Performance

- DB2 PTF UK58205 – support up to 100000 open data sets in the DB2 DBM1 address space

- Allocation enhancement – MEMDSENQMGMT (Memory-based data set ENQ management)
  - ENQs managed in private storage instead of SWA blocks. Job will be non-restartable.
  - Enable via parmlib member or MVS command:
    - ALLOCxx: SYSTEM MEMDSENQMGMT(ENABLE)
    - SETALLOC SYSTEM,MEMDSENQMGMT=ENABLE
  - Exploit with DB2 APAR PM17542 (closed September 2010)

- GRS enhancement – special case for DB2 allocations:
  - Allow shared SYSZTIOT ENQs ahead of exclusive waiters to promote parallelism
    - Controlled by bit in JSCB (JSCBTIOD) which is set by program
    - Up to 50 shared request may jump ahead of an exclusive qname=SYSZTIOT
  - V1.12 APAR OA33633 (closed August 2010)
  - DB2 exploitation APAR PM18557 (closed August 2010)

z/OS V1.12 Open DB2 Data Sets – non-SMS Results

- Single member sysplex with GRS STAR on a z10 with 4 CPs (Pok tests)
- 20 concurrent batch jobs to open data sets
- DSNMAX=100000 in DSNZPARM

	V1.11	**V1.12 w/GRS exploitation**	**V1.12 w/MEMDSENQMGMT**	**Delta V1.11 to V.12 w/GRS and allocation**
Elapsed time	480.7	321.4	239.5	-50%
LPAR CPU%	59.5	66.5	75.7	+27%
ETR (number of open DS per sec)	199.7	298.7	400.8	+100%
ITR	335.6	449.1	529.3	+58%

z/OS V1.12 DB2 Open Data Sets – SMS-managed Results

- Tests using SMS-managed data sets done in DB2 lab (San Jose)
- z10 with 8 CPs, sysplex, DB2 9 for z/OS, DS8300 disk
- Table shows difference between V1.11 and V1.12 w/MEMDSENQMGT and GRS exploitation:

	Delta V1.11 to V.12 with GRS and MEMDSENQMGT
Elapsed time in sec	-76%
Avg. open time per data set	-69%
DBM1 TCB time	-85%
ETR	+4.3x

SVC Dump Exit Data Capture Problem

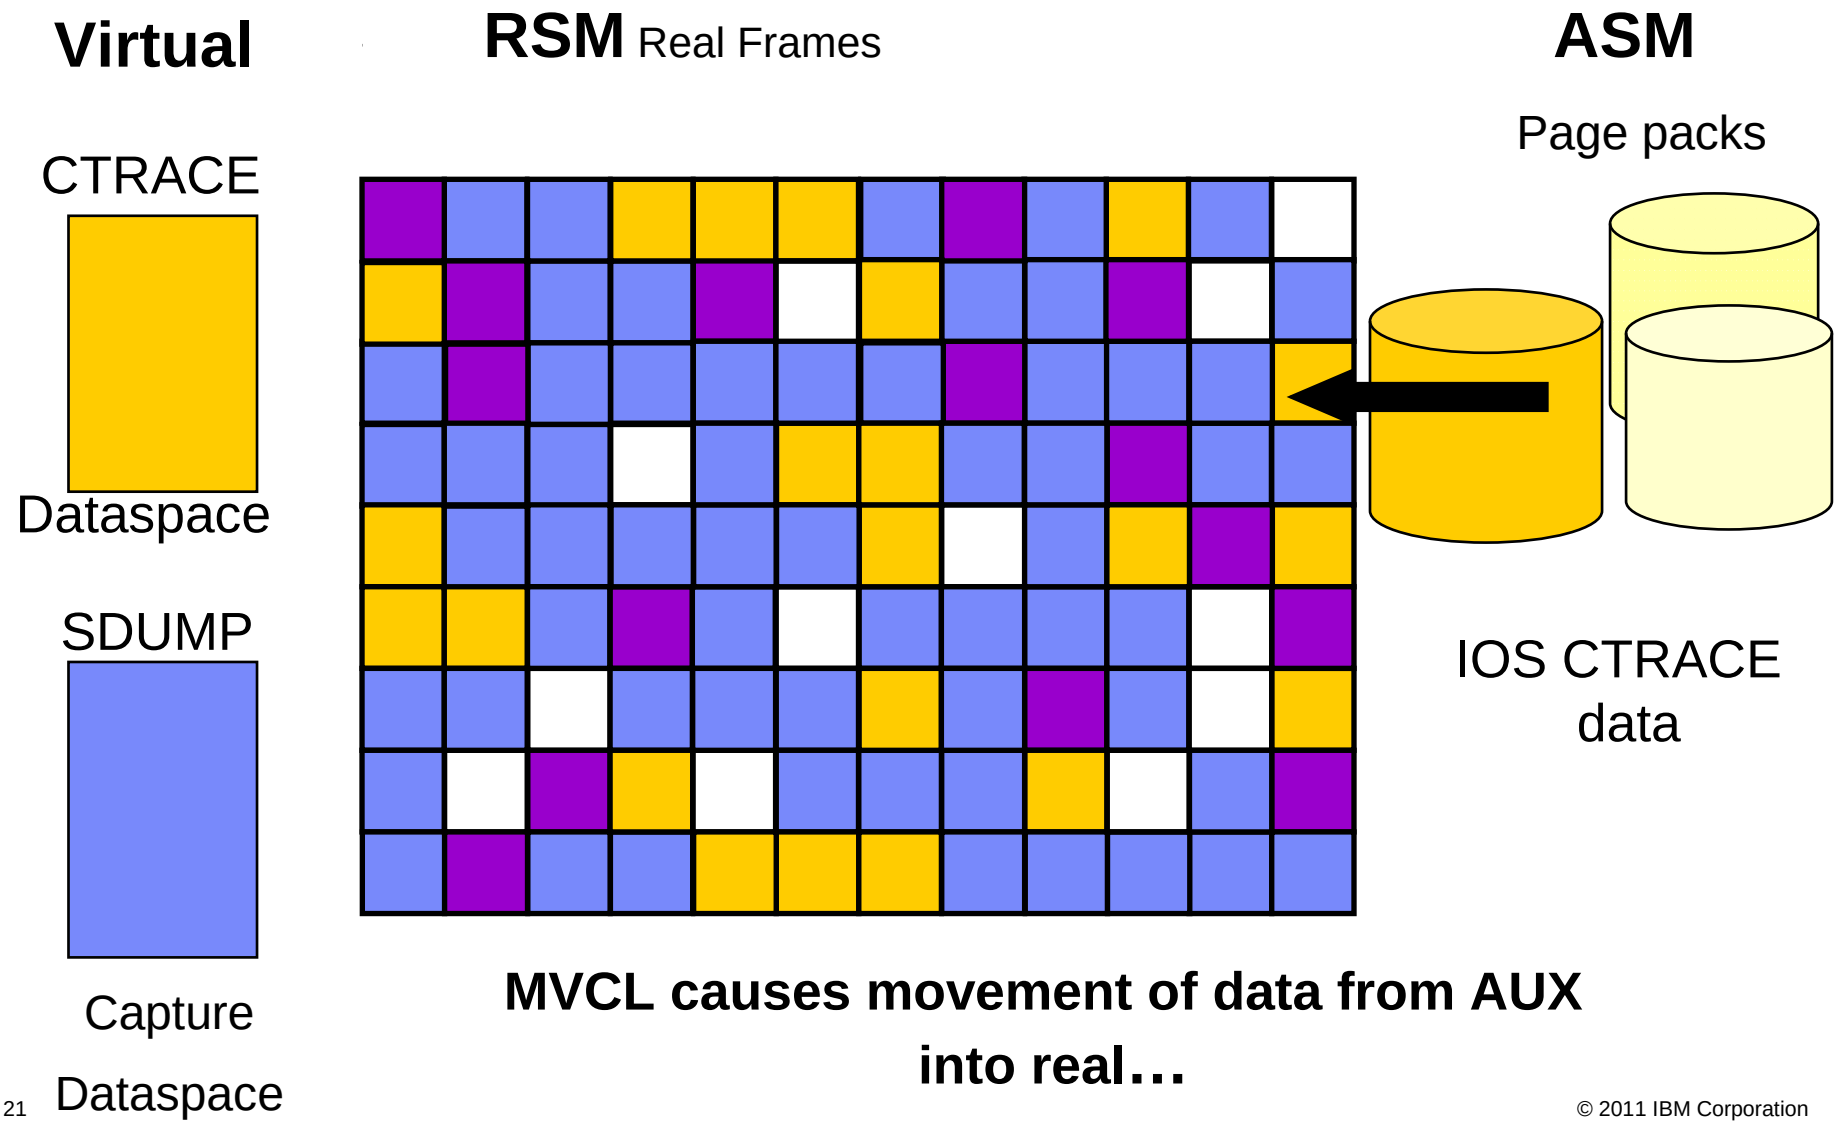

MVCL causes movement of data from AUX into real…

SVC Dump Exit Data Capture Problem

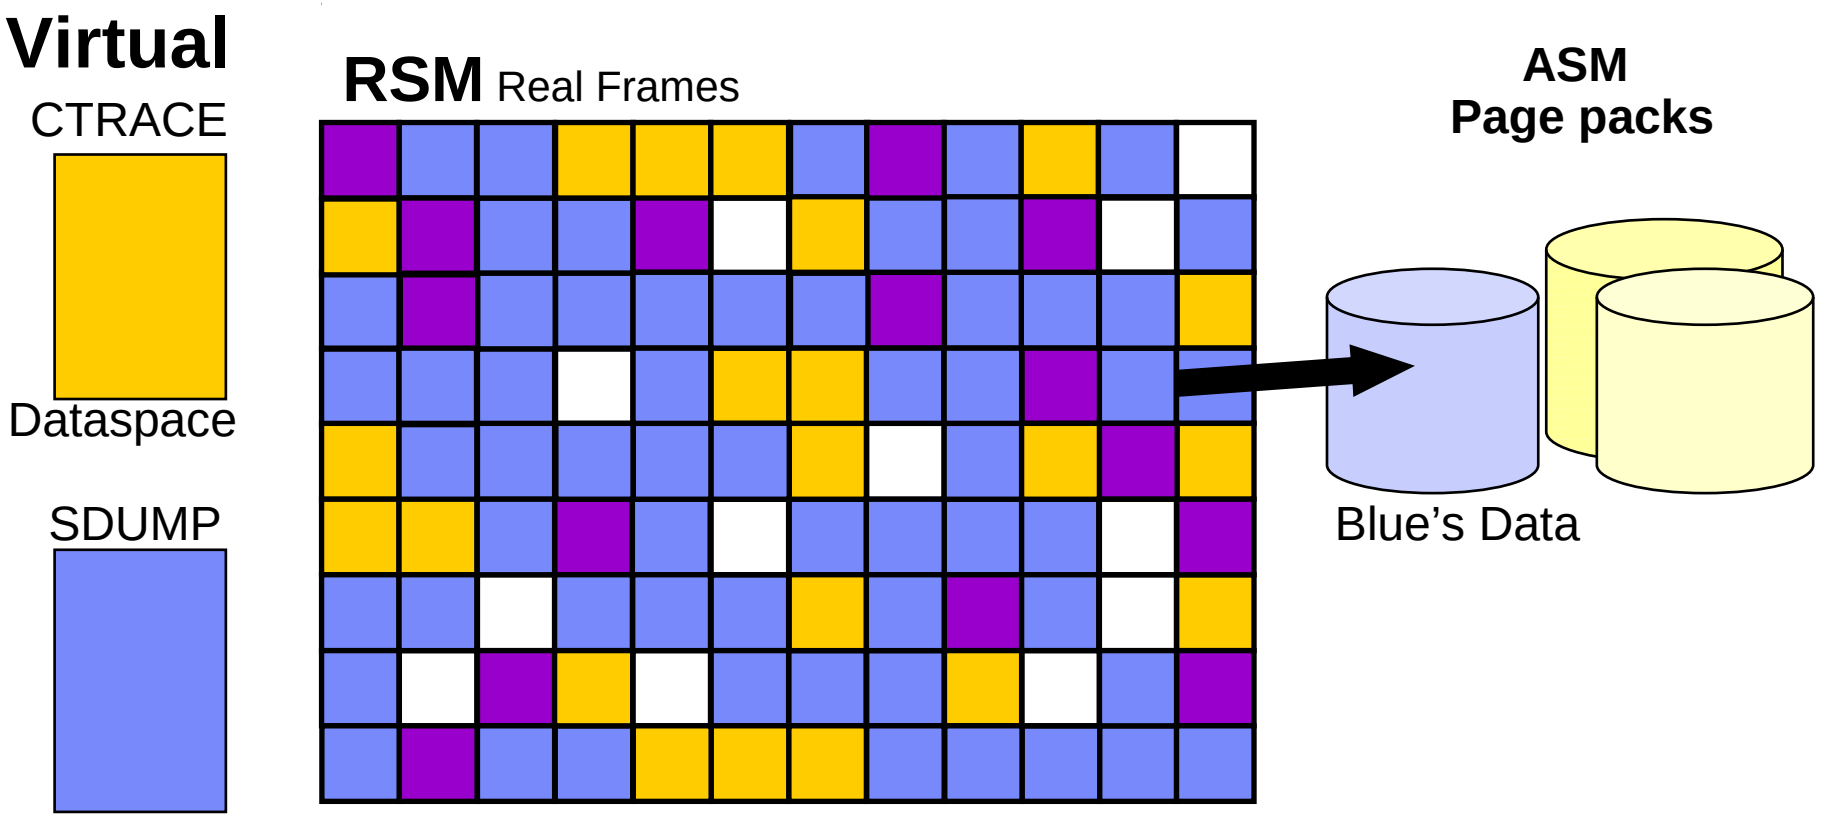

…which in turn forces page-out of other more important data

z/OS V1.12 SVC Dump Solution

- Reduce memory pressure when capturing exit data
- Do a smart copy such that if the source data is out on AUX we do an I/O directly into the SDUMP buffer space to capture the data while leaving the source data out on AUX
- In the previous example this would mean that the IOS component trace data will not be brought into real and an I/O will be done to copy the data from AUX directly into the SDUMP capture dataspace
- Avoid changing the reference pattern of the source data due to capture
- Copy the data via a special RSM service such that if the source data did not appear referenced before the capture it still remains unreferenced after the capture
- Data that has been captured will not cause other perhaps more important data to be paged out
- Data in the SDUMP capture space will be made to look old so that this data will be paged out before any important workload data is paged out

z/OS V1.12 SVC Dump Exit Data Capture Solution

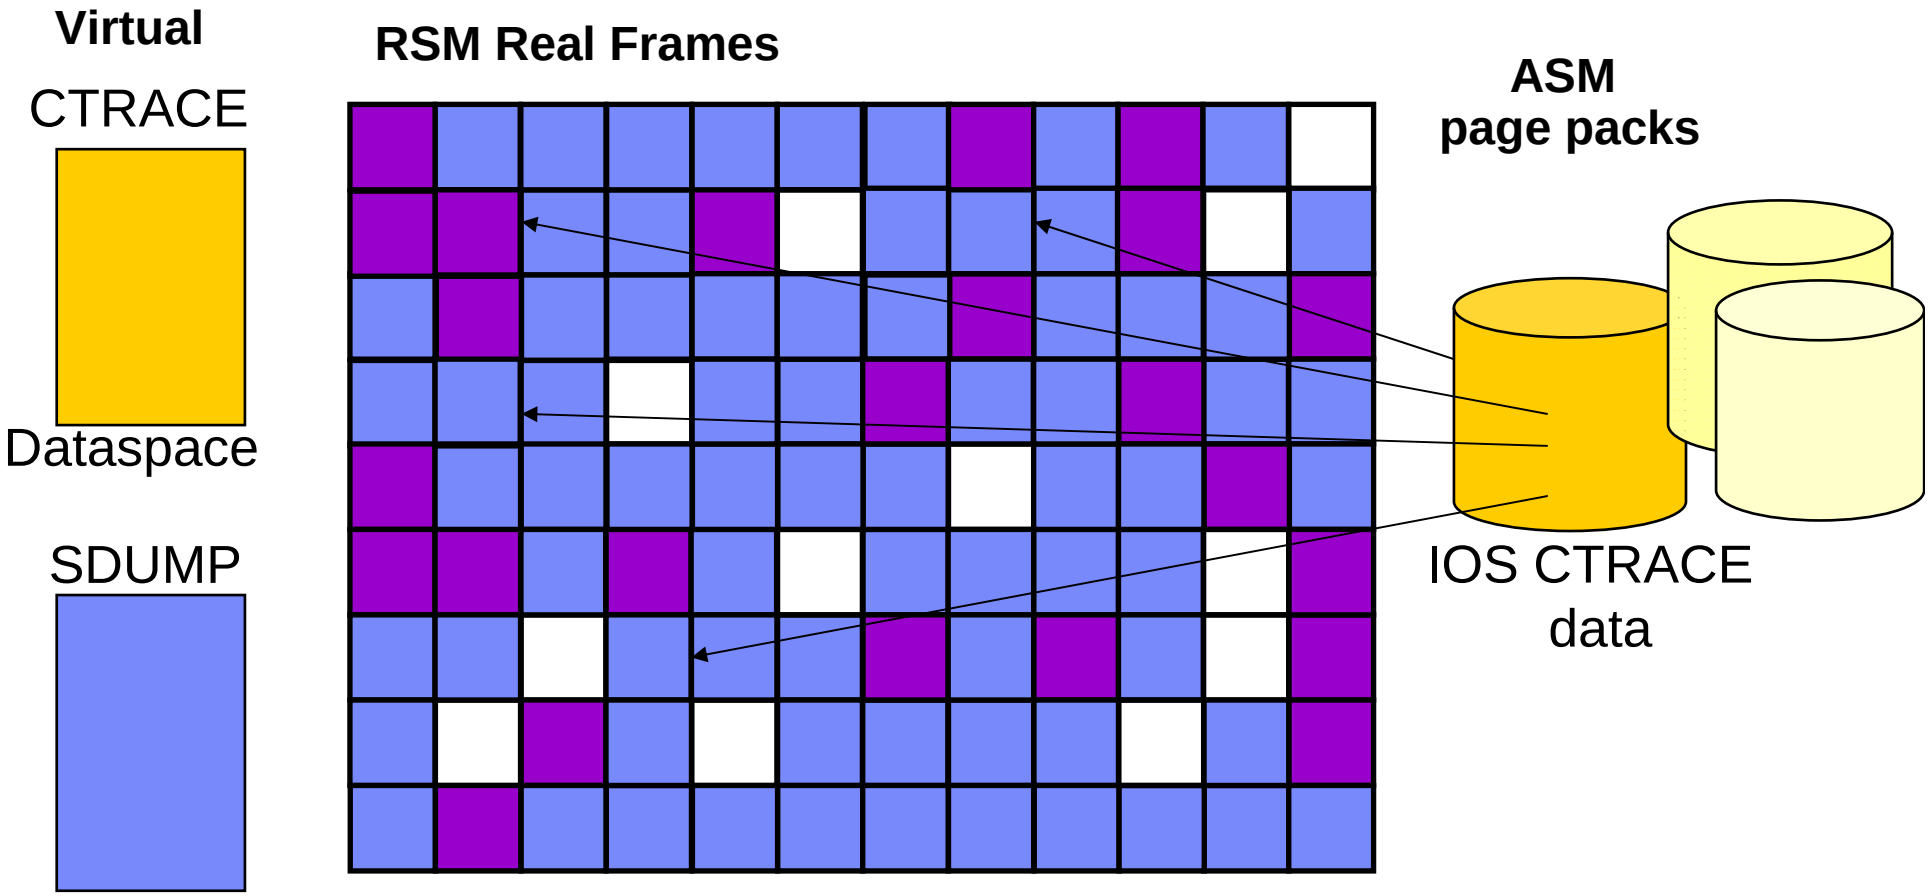

Smart Copy direct to SDUMP dataspace

z/OS V1.12 SVC Dump Test Configuration and Scenarios

- Three variations of amount of storage on page data sets: 0%, 5%, and 60%
- Tested with real storage constrained and non-constrained environments.
- Tests on a z10 with 4 CPs

Dump Size	Real storage	Page data sets	Dump data sets
10 GB	25 GB	15 x 4 GB page data sets on 2105-800	2 x 54 GB dump data set on 2107-922 w/116.7 GB cache
20GB	45 GB	15 x 4 GB page data sets on 2105-800	2 x 54 GB dump data set on 2107-922 w/116.7 GB cache
40GB	85 GB	15 x 6.74 GB page data sets on 2105-800	2 x 54 GB dump data set on 2107-922 w/116.7 GB cache

z/OS V1.12 with 60% of Data on Aux and No Real Storage Constraints

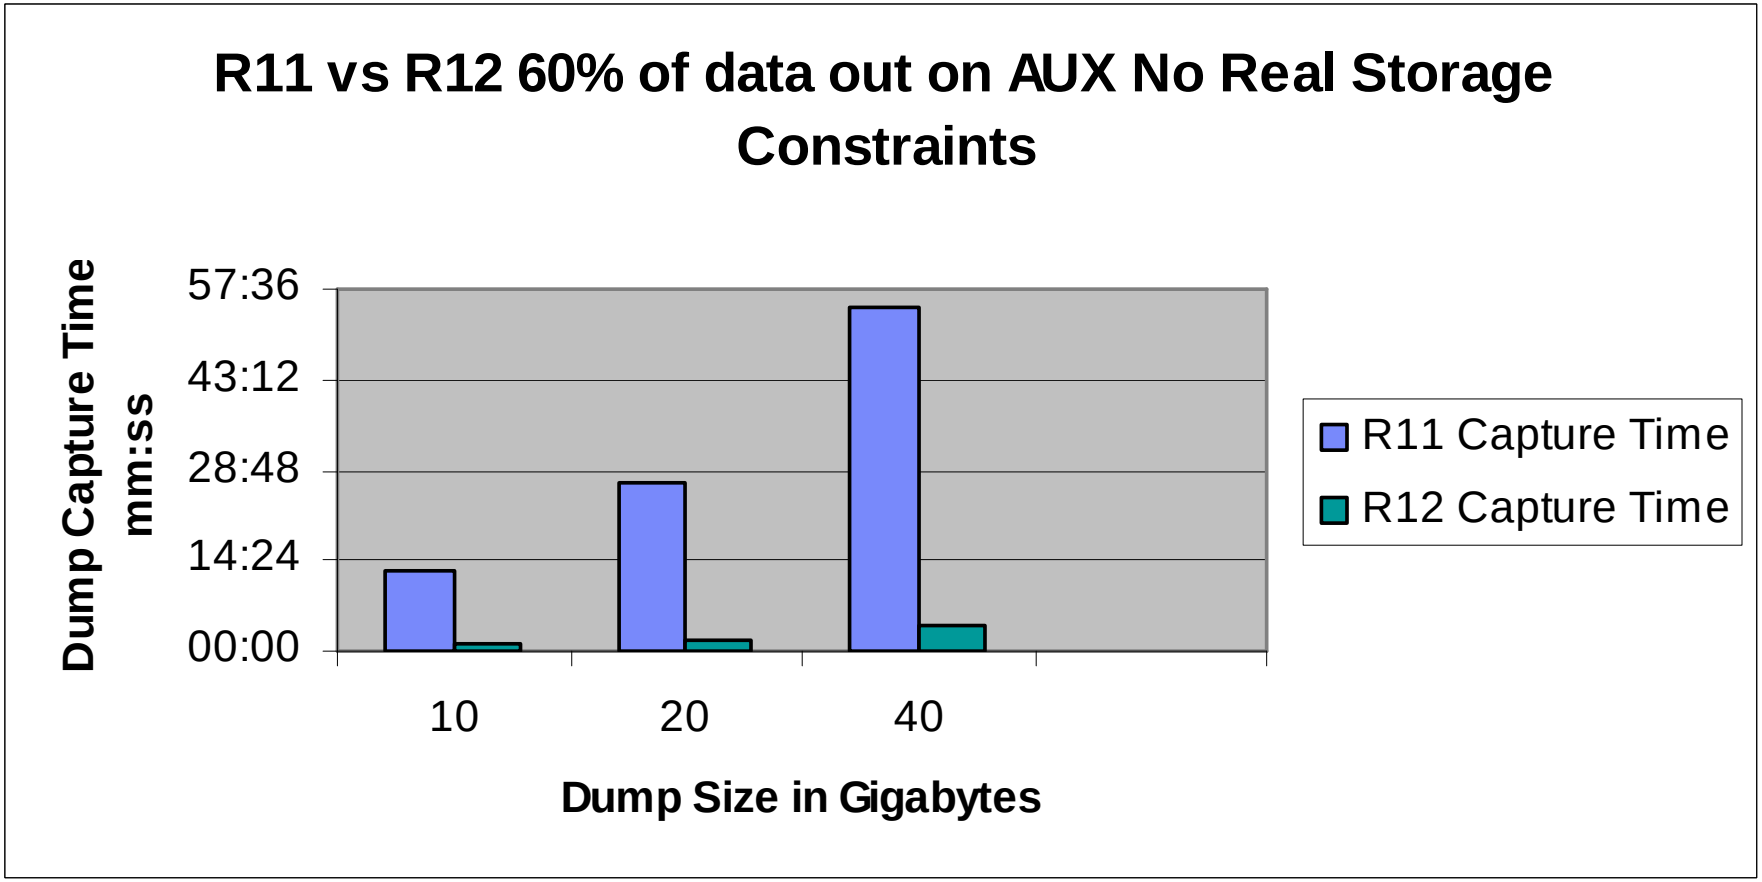

z/OS V1.12 with 60% of Data on Aux and Real Storage Constraints

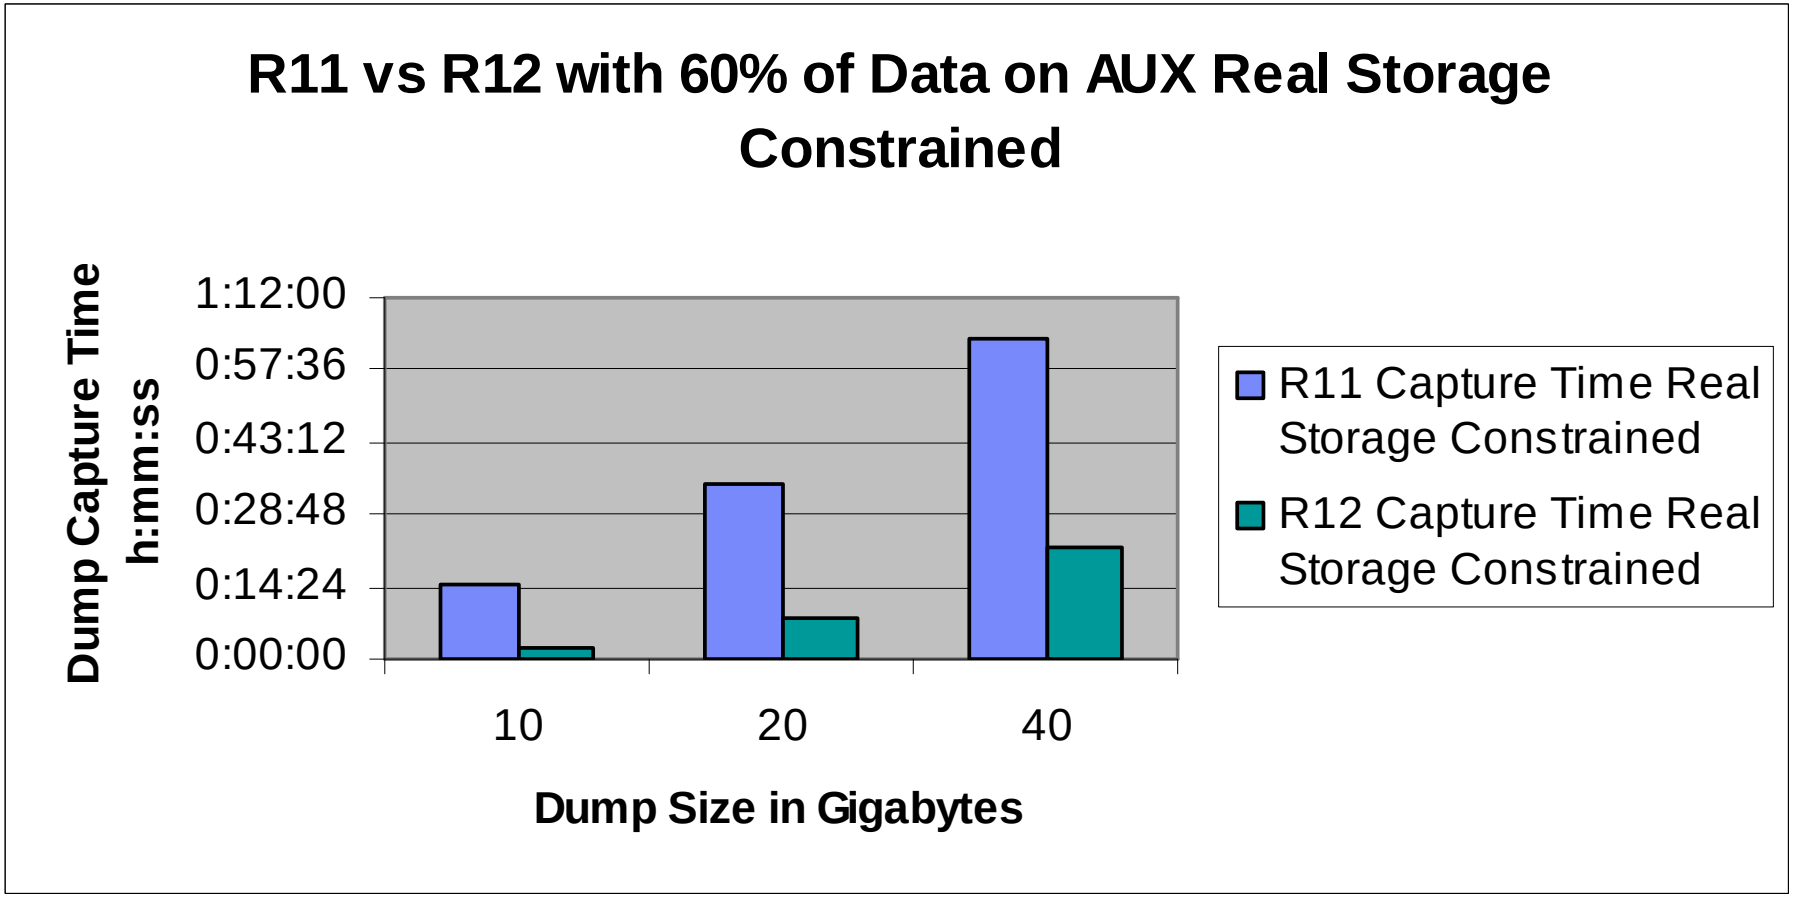

z/OS V1.12 10 GB Dump and No Real Storage Constraints

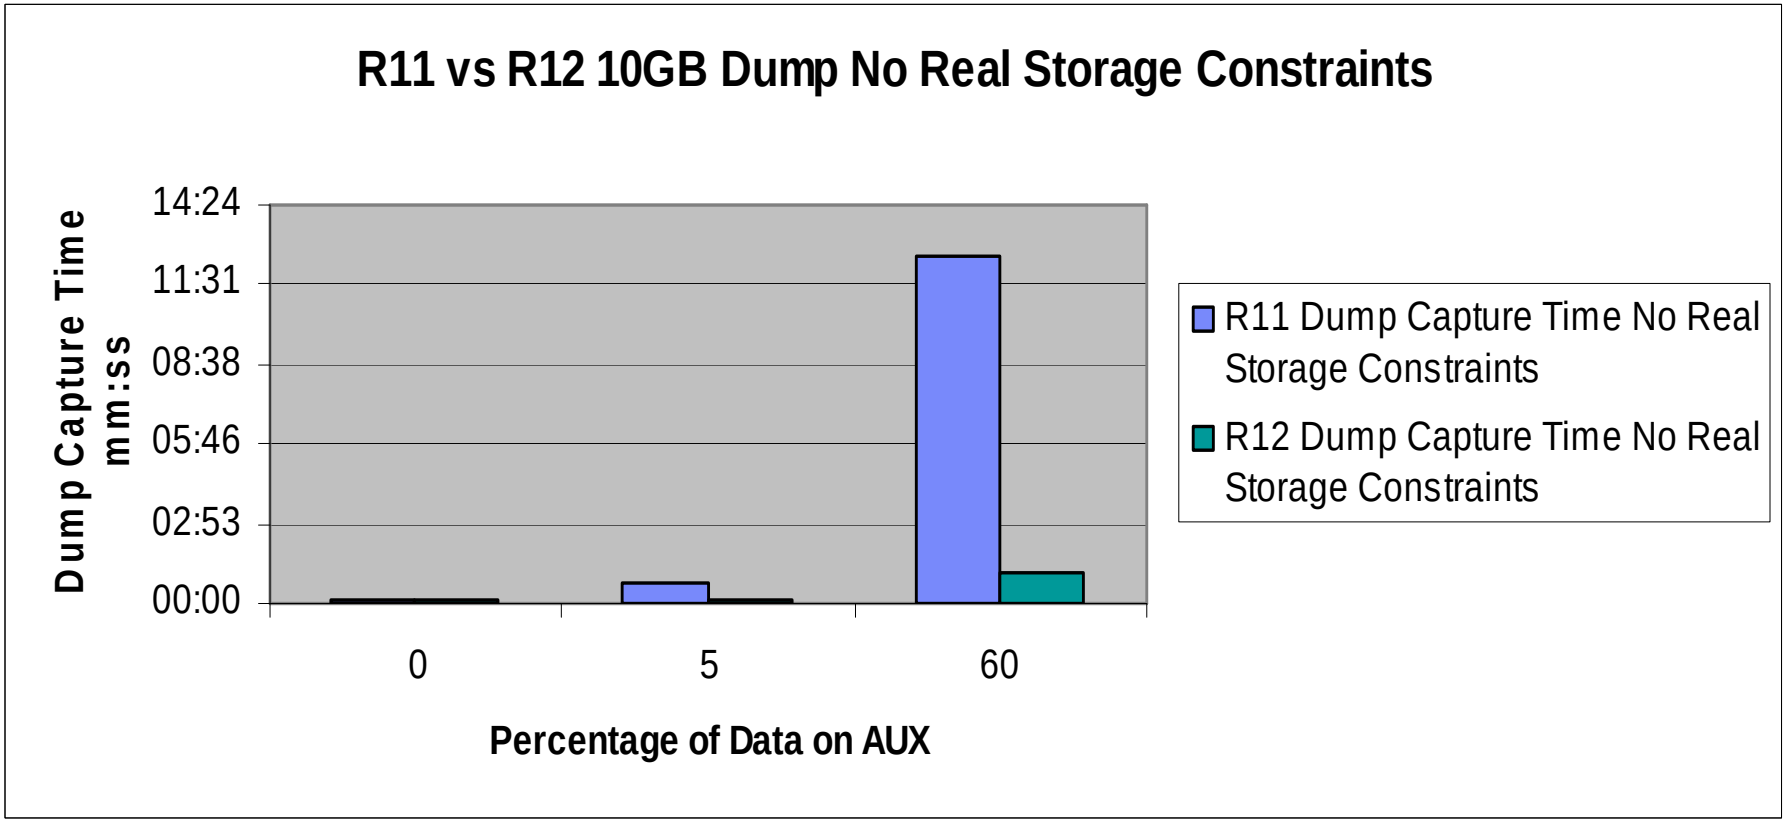

z/OS V1.12 10 GB Dump and Real Storage Cosntraints

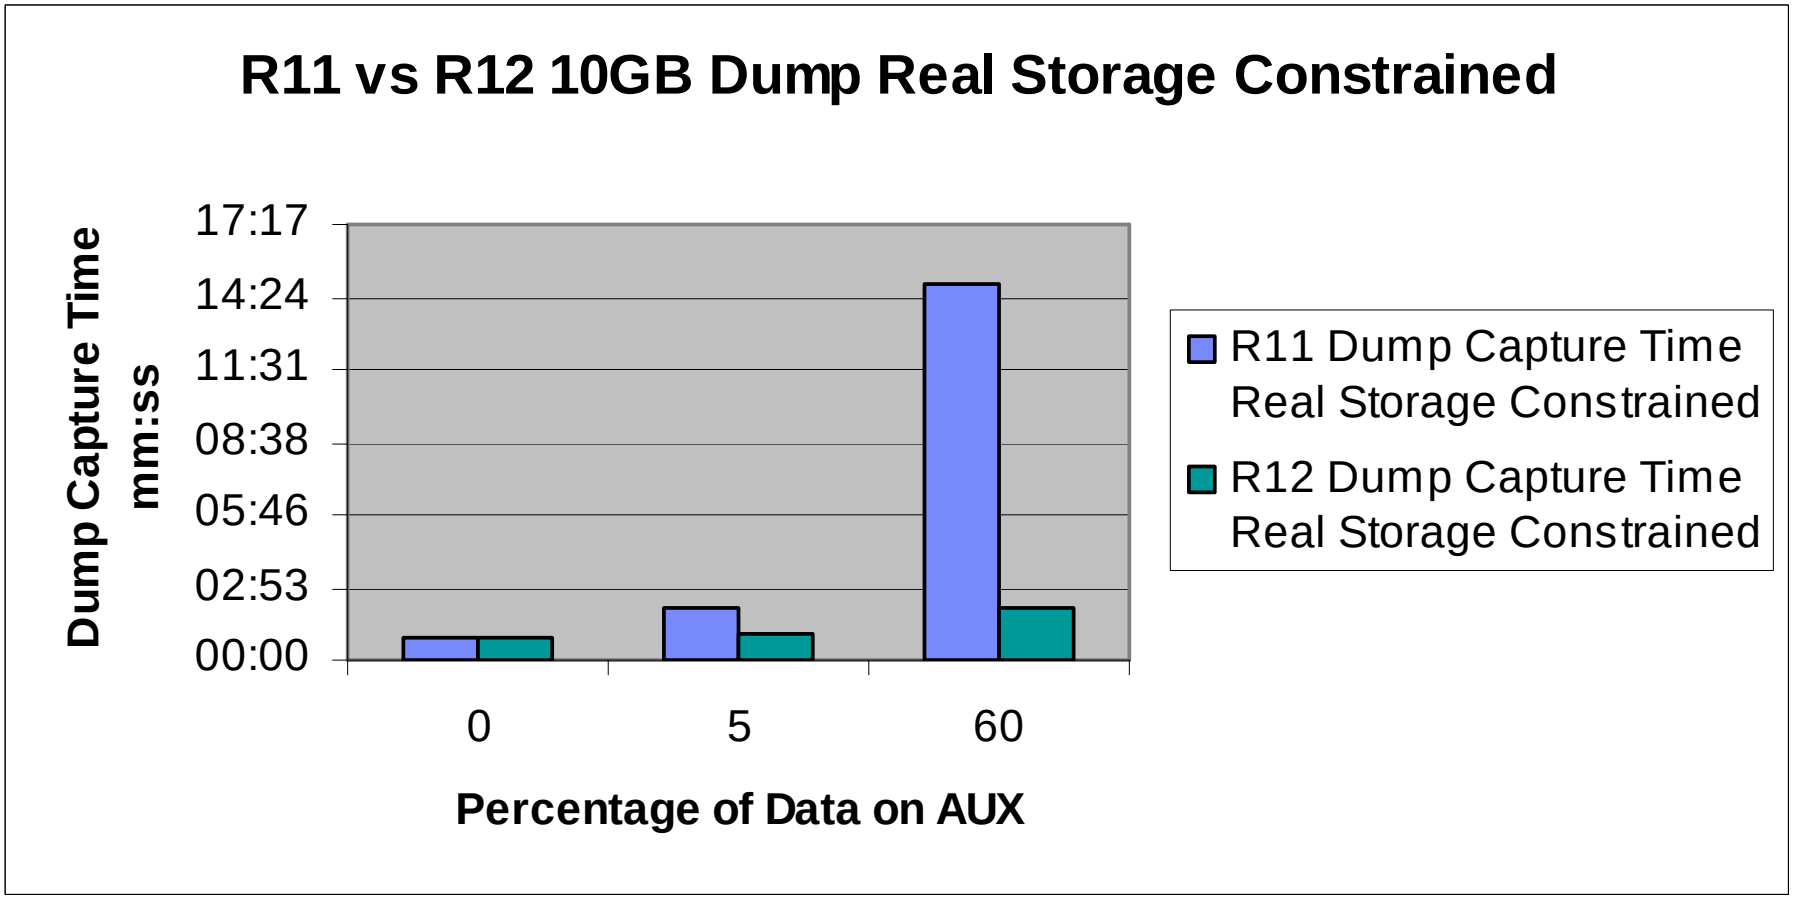

z/OS V1.12 SVC Dump Improvements with 60% of Data on Aux

Dump Size	No real storage constraint	Real storage constrained
10 GB	91%	86%
20 GB	92%	76%
40 GB	93%	65%

- Over 90% performance improvement measured in systems without real storage constraints, and only slightly increasing benefits as size of dump increases.

- In real-storage-constrained systems, the performance benefits are significant, although less pronounced (65%-86%). As the size of the dump increases, the observed performance benefits decrease from 86% to 65%.

z/OS V1.12 SVC Dump Improvements with 5% of Data on Aux

Dump Size	No real storage constraint	Real storage constrained
10 GB	78%	50%
20 GB	79%	52%
40 GB	82%	50%

- About 80% performance improvement measured in systems without real storage constraints, and only slightly increasing benefits as size of dump increases.

- In real-storage-constrained systems, the performance benefits are significant, although less pronounced (around 50%).

- Performance improvements are roughly 60% better in non-real storage constrained environments as compared to real-storage constrained environments.

z/OS V1.12 SVC Dump Improvements Summary

- Dramatic performance improvements observed in capture time for address spaces with high percentage on AUX, especially in environments with no-real-storage constraints.
  - 40 GB dump, 60% on AUX, 55min-R11 vs. 4 min-R12
- Significant performance improvements also observed in capture time for address spaces with as little as 5% on AUX, in constrained and non-constrained environments
  - 40 GB dump, 5% on AUX, 50-80% improvement
- Future performance runs will focus on the improvements in capture time for common storage (whole system non-dispatchable).

z/OS V1.11 and V1.12 Results Summary

- V1.11:
  - CPU performance goal: Overachieved: avg. 1.7% vs goal of 1.25%
  - VSCR performance goal:
    - EDT table compressed. The total reduction is about 2 MB based on LSPR workloads.
    - Telnet VSCR – not measured in LSPR workloads
- V1.12:
  - CPU performance goal: Overachieved: avg. 2.7% vs goal of 1.25%
  - VSCR results:
    - Across workloads : 8-12 MB savings (target 4 MB)
    - RMF CPools to 64-bit : ~10 MB
    - MMSB to 64-bit : ~ 1 MB

z/OS Release-to-Release Performance Improvements

- Focus on performance improvements:
  - Future z/OS releases
  - z/OS SW stack

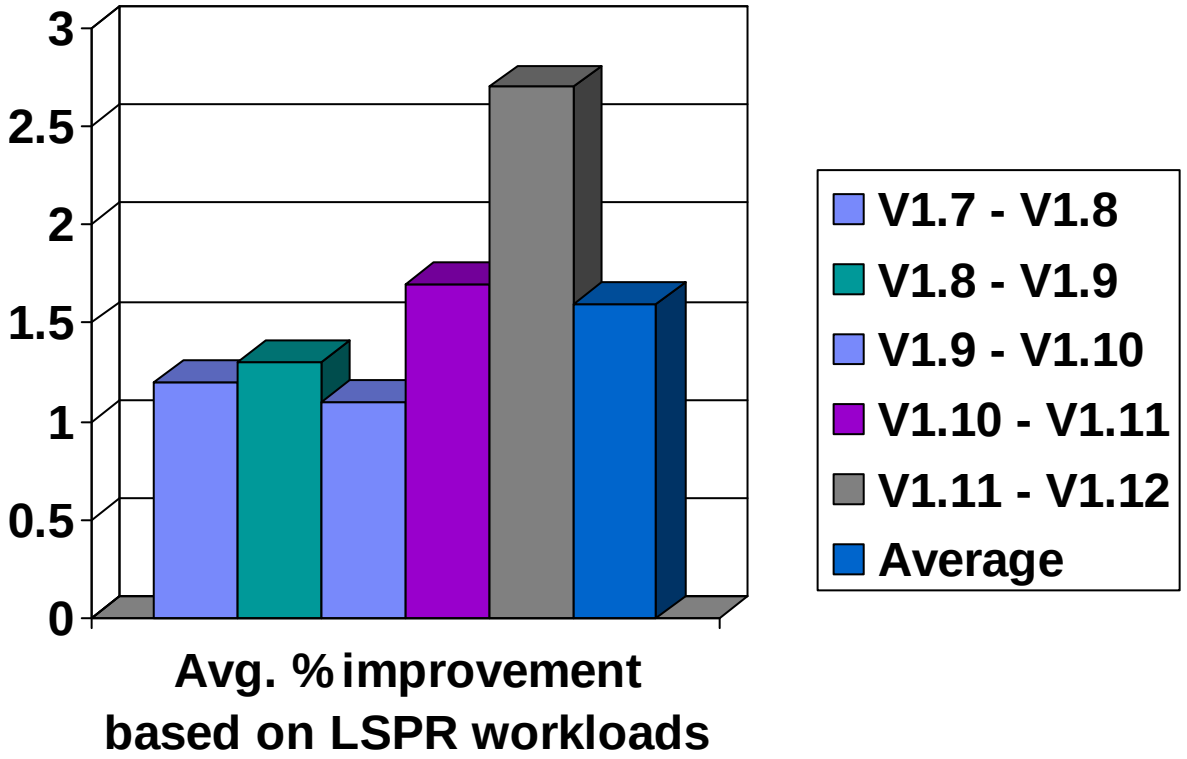

z/OS Performance References

- WSC SoftCap tool:
  - Customer version: http://www.ibm.com/support/techdocs
  - IBM internal version: http://w3.ibm.com/support/americas/wsc/cpsproducts.html
  - IBM business partners: http://partners.boulder.ibm.com/src/atsmastr.nsf/WebIndex/PRS1762
- LSPR web site (LSPR workloads):
  - https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex

SHARE sessions:

- 9024: z/OS Basics: The z/OS UNIX Shared File System Environment
  - Tuesday March 1: 9:30-10:30 in 201C (Jim Showalter)
- 9035: Smarter SVC Dump Processing for Improved z/OS Resiliency
  - Wednesday March 2: 8:00-9:00 in 201D (Elpida Tzortzatos)