IBM Systems & Technology Group

Smarter SVC Dump Processing for Improved z/OS Resiliency

SHARE in Anaheim

Session 9035

Elpida Tzortzatos elpida@us.ibm.com

Agenda

SVC Dump Enhancements

- Overview of today's SVC Dump
- New/Enhanced External Control Knobs for the SysProg
  - CHNGDUMP SET,SDUMP,AUXMGMT=ON/OFF
  - CHNGDUMP MAXSNDSP=sss
- Internal Algorithm Improvements (z/OS V1R12)
  - SVC Dump "SmartCopy"
  - ***Performance Data/Results***
- Aux Storage Configuration Strategies

SVC Dump Enhancements
Part 1:
Overview of today's SVCDUMP

OVERVIEW: Objectives of SVC Dump

- Never cause an outage taking a dump
- Capture diagnostic data before it is overwritten
  - This translates to capture it fast enough
- Cause minimal performance disruption
  - Due to their memory intensive nature, dumps cannot be processed transparently, but their impact should be mitigated to be just what is essential

OVERVIEW: Current Controls

- The Sysprog has some responsibility in making sure SVC Dump meets its objectives
  - Provide sufficient auxiliary storage for normal system operation plus the dump capture phase (ideally for multiple dumps)
  - Limit the virtual used by SVC Dump to protect the system from an out-of-aux wait state (MAXSPACE= parameter on CHNGDUMP command)
  - Provide some "reserved" real storage to expedite the capture of common storage (BUFFERS= parameter on CHNGDUMP command)

OVERVIEW: Current Controls

Prior to z/OS V1R11:

- MAXSPACE defines the maximum amount of virtual storage for DUMPSRV to use
  - CD SET,SDUMP,MAXSPACE=xxxxxxxxM
- Default is 500M
- SysProg must estimate the impact upon auxiliary (paging) storage
- SVCDUMP processing truncates the dump when MAXSPACE is reached or SRM detects that 85% of paging space is used up

OVERVIEW: Limitations of Current Design

- **Problem Statements:**
  - Too late to react when critical auxiliary storage shortage (85% utilization) is detected by SRM
  - Large exposure that dumping could cause a WAIT state 03C RSN01 (paging space exhausted)
  - How large should MAXSPACE be to prevent WAIT03C and to capture the largest dump?
    - First Failure Data Capture (FFDC)

    -vs-

    - System Availability/RAS

SVC Dump Enhancements

Part 2:

New/Enhanced External Control Knobs

New External Control Knob: CD SET..AUXMGMT..

- New keyword AUXMGMT=On/Off on CHNGDUMP
  - SDUMP will continually monitor AUX utilization during the dumping process.

- Benefits:
  - Aux Monitoring is enhanced to detect AUX storage utilization changes more rapidly.
  - Improve the management of virtual storage when an SVC DUMP is taken.
  - Allows a dump to complete if the customer has provided sufficient AUX storage.

New External Control Knob: CD SET..AUXMGMT..

Usage & Invocation

- AUXMGMT=ON
  - New keyword AUXMGMT=On/Off is added on CHNGDUMP
  - Prior to AUXMGMT, only MAXSPACE restricted DUMPSRV's use of virtual storage
  - With AUXMGMT, the installation's auxiliary storage resource restricts the behavior of DUMPSRV
  - This is the default which makes availability a higher priority over first-failure data capture

CHNGDUMP SET,SDUMP,AUXMGMT=ON/OFF

New External Control Knobs: CD SET..AUXMGMT…

- AUXMGMT=ON
  - No new dumps are allowed when AUX storage utilization reaches 50%
  - Current dump data capture stops when AUX storage utilization reaches 68%
  - Once the limit is exceeded, new dumps will not be processed until the AUX storage utilization drops below 35%
  - Always honor MAXSPACE when it is more restrictive than AUXMGMT. (i.e. When MAXSPACE=35Meg, stop SVC dumps when MAXSPACE is exceeded even if AUX utilization is only 3%,)

New External Control Knob: CD SET..AUXMGMT..

AUXMGMT=OFF

- SDUMP virtual storage management reverts to control via MAXSPACE
- Dump in progress is stopped and made as a partial dump when a critical AUX storage shortage (85% utilization) is detected or MAXSPACE is exceeded.
- After critical AUX storage shortage, AUX storage utilization must be 35% or less before dump capture can resume
- Installation must turn AUXMGMT off to have previous behavior

System non-dispatchable during global capture

- Taking down or inhibiting the customer's system functions in order to take an SVC dump is certainly not desirable
  - Slowly progressing global data capture may leave the system non-dispatchable long enough for it to be partitioned from the Sysplex
- Add another factor – maximum system non-dispatchability to the criteria of determining when to reset system to dispatchable in z/OS V1R11
  - Default MAXSNDSP is set to 15 seconds and can be modified via the CHNGDUMP command

SVC Dump Enhancements
Part 3:
Internal Algorithm Improvements
“Smart Copy”

SVC Dump Exit Data Capture Problem

- **An Example:**
  - The IOS component trace data resides on AUX and it is now being dumped by the IOS SDUMP exit
  - As part of the SVC dump capture processing for the IOS component trace data, all the data will be brought into real and copied into the SDUMP capture dataspace
  - Now, even though the IOS component trace data will not be referenced again in the near future, it is all in real as recently referenced data
  - Having the IOS component trace data in real may put pressure on real memory availability, forcing page-out of other (more likely to be referenced) data.
    - The IOS component trace data will stay in real since it is recently referenced and may cause other more important data to be paged out.

SVC Dump Exit Data Capture Problem

Virtual

CTRACE

Dataspace

SDUMP

CAPTURE Dataspace

RSM Real Frames

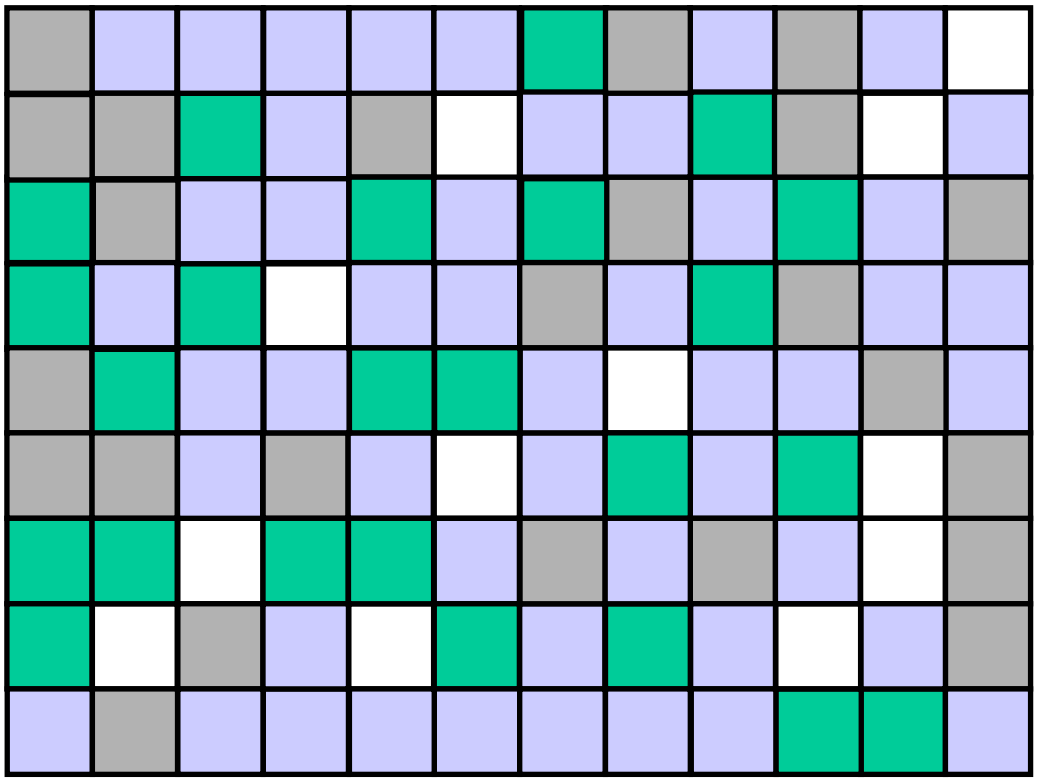

ASM

Page Packs

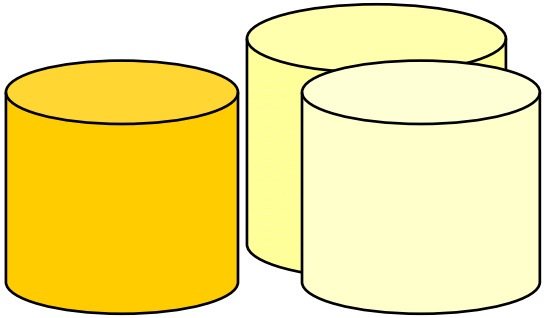

IOS CTRACE data

With IOS CTRACE data paged out…

SVC Dump Exit Data Capture Problem

Virtual

CTRACE

Dataspace

SDUMP

CAPTURE Dataspace

RSM Real Frames

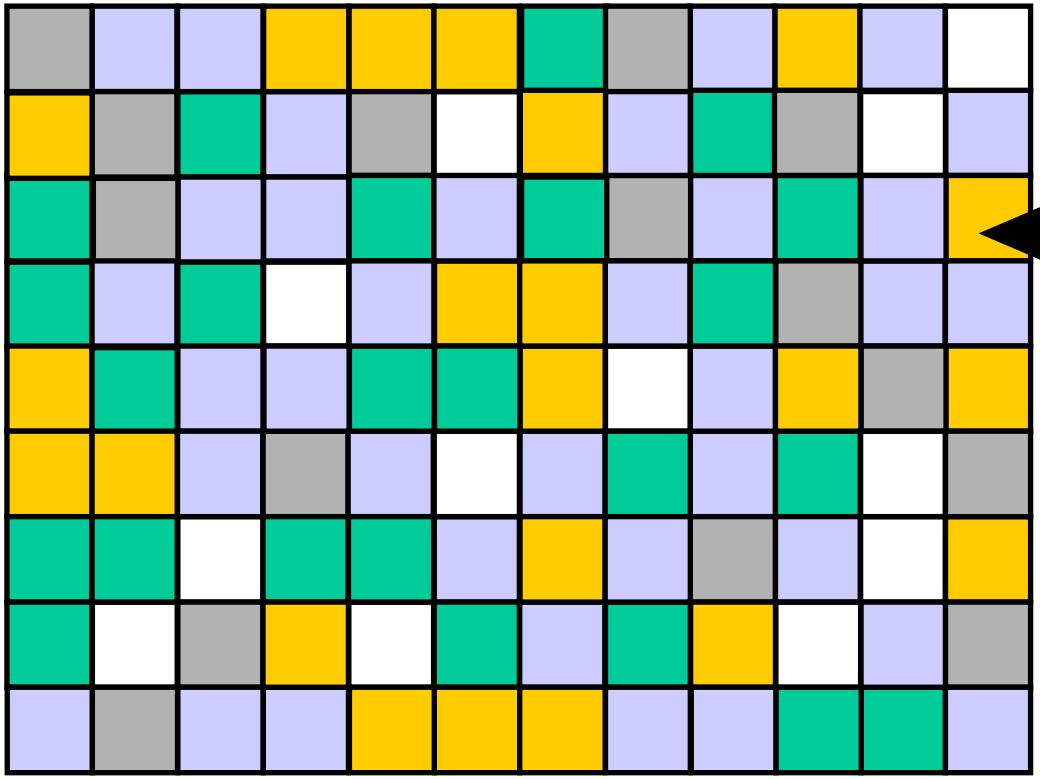

ASM

Page Packs

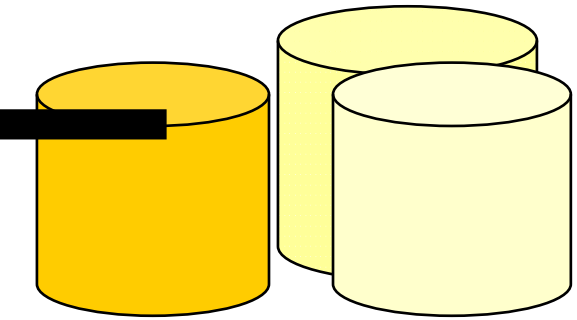

IOS CTRACE data

MVCL causes movement of data from AUX into real…

SVC Dump Exit Data Capture Problem

Virtual

CTRACE

Dataspace

SDUMP

CAPTURE Dataspace

RSM Real Frames

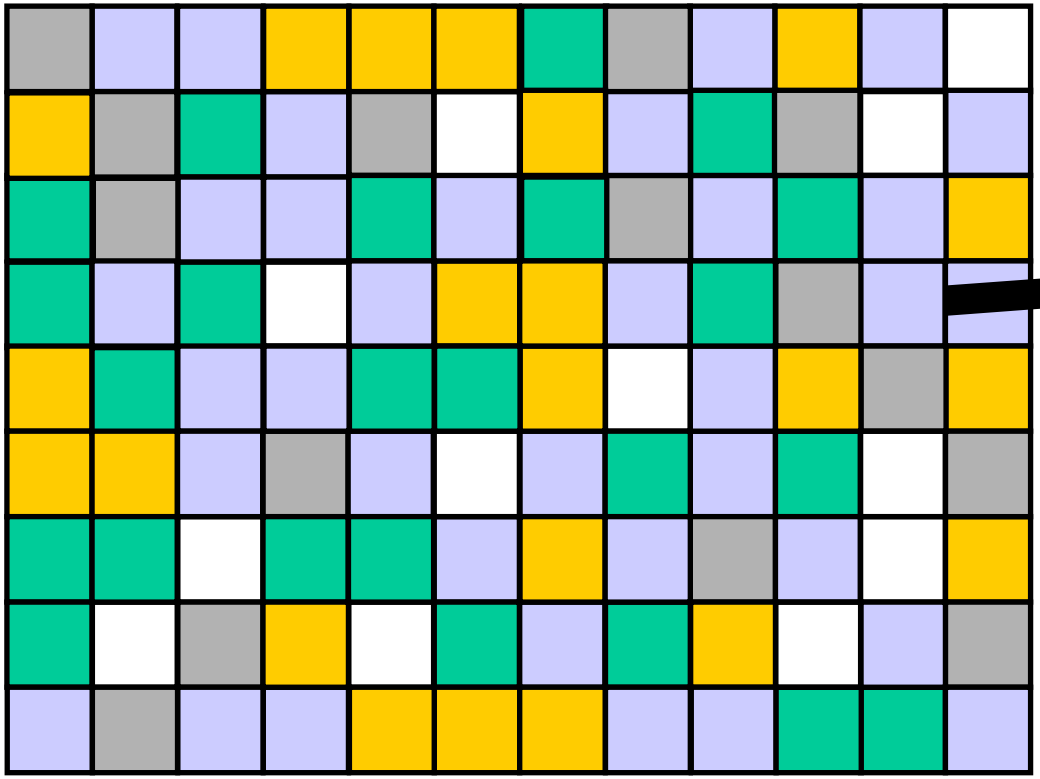

ASM

Page Packs

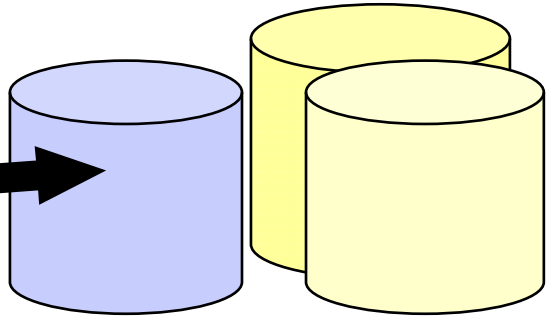

Blue's Data

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

z/OS R12 - SVC Dump Capture 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 R12 - SVC Dump Exit Data Capture Solution

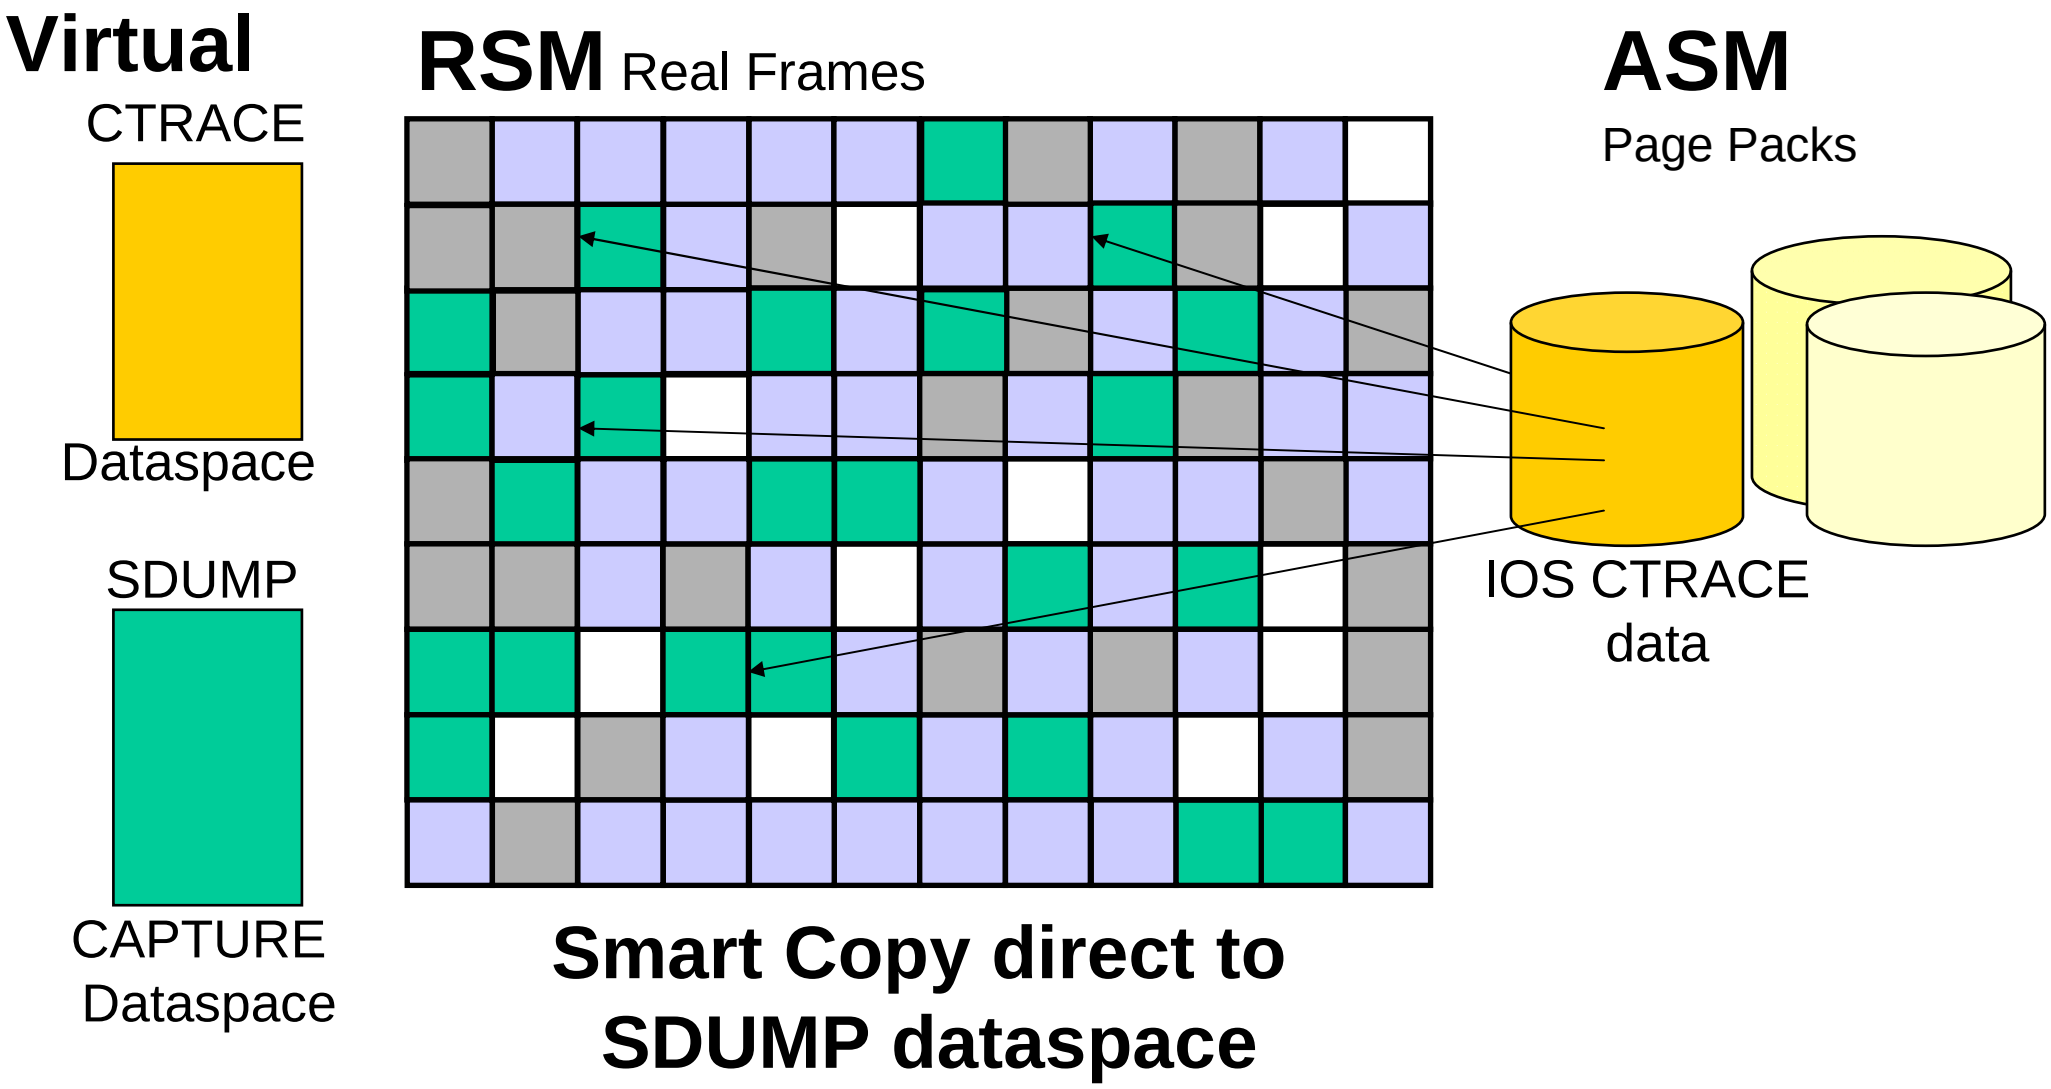

z/OS R12 - SVC Dump Exit Data Capture Solution

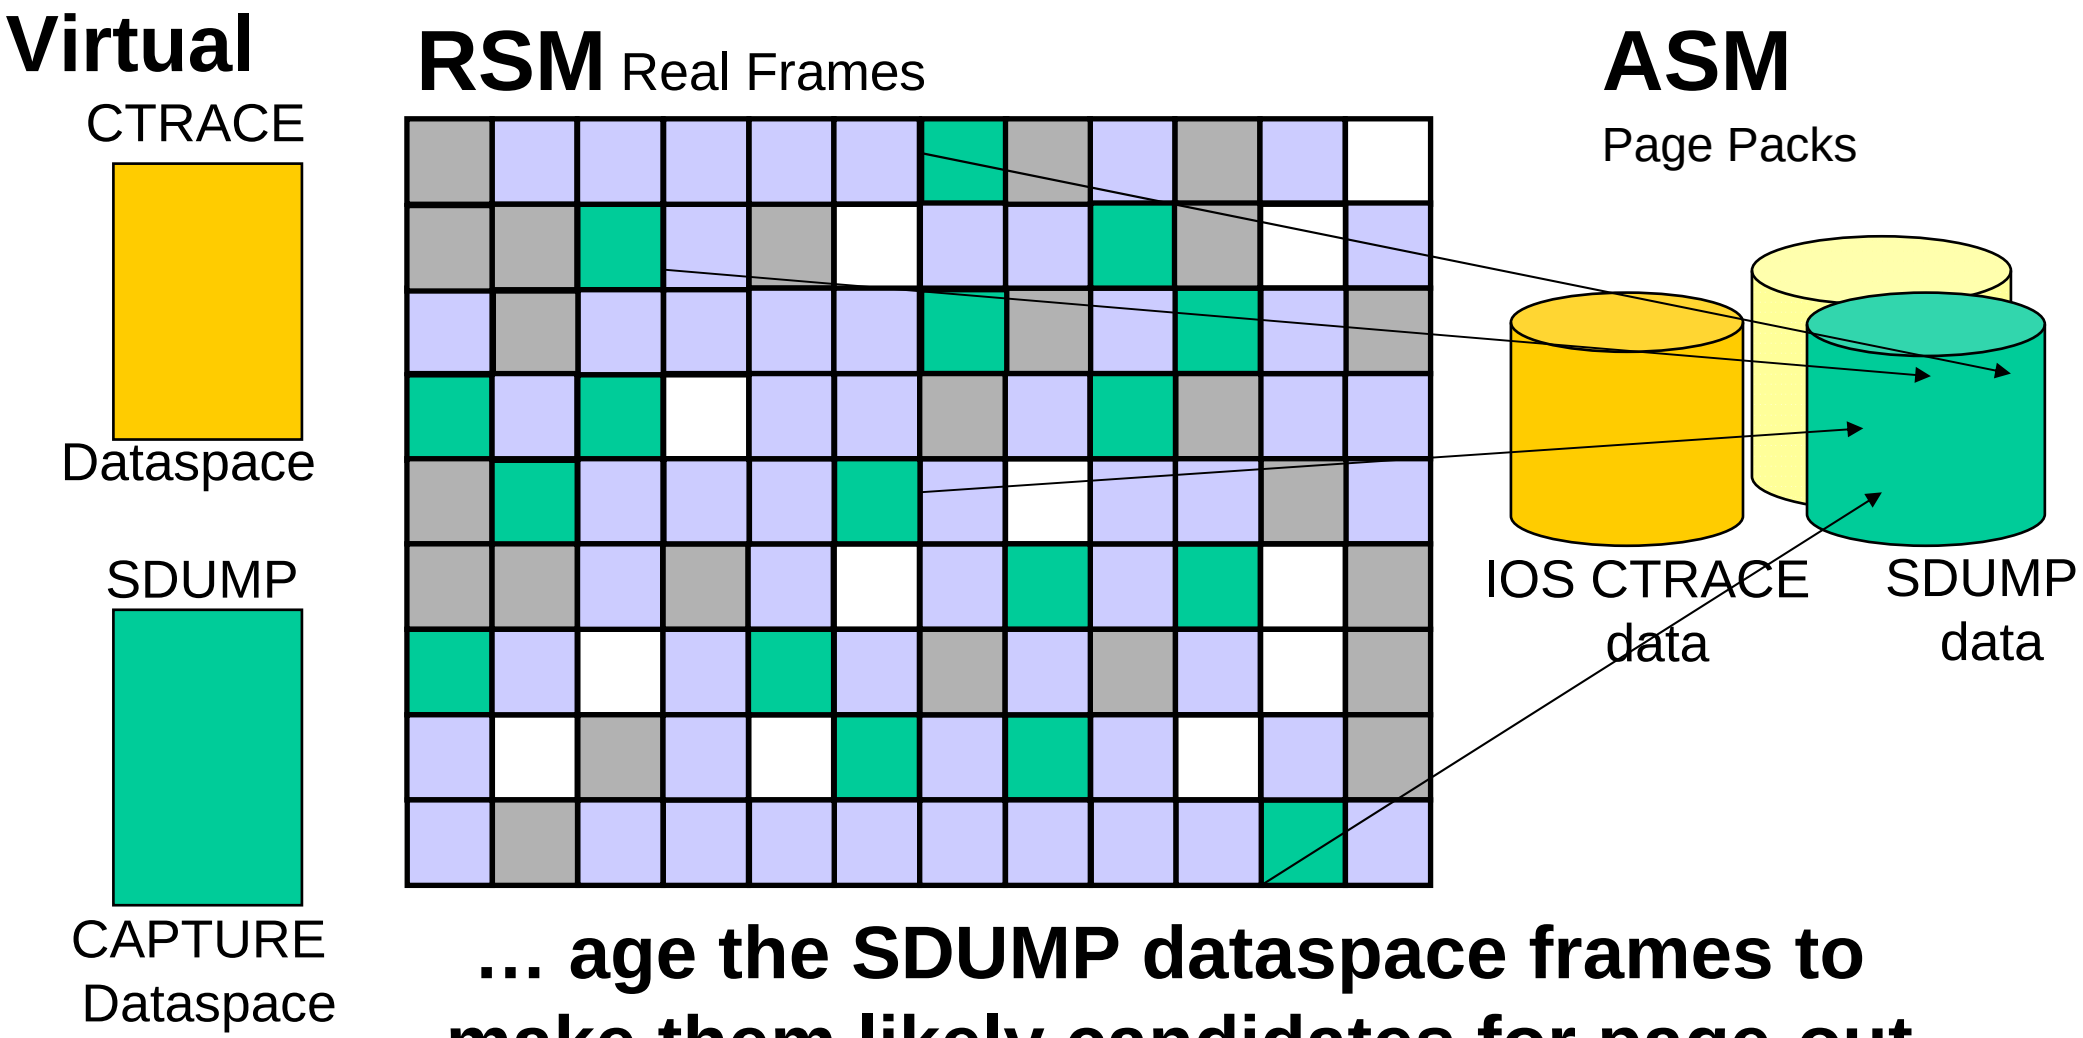

… age the SDUMP dataspace frames to make them likely candidates for page-out

SVC Dump Enhancements

Performance Test Results
with Internal Algorithm Improvements

Configuration set up 10GB and 20GB dumps

ASM Paging Packs**
(with 4GB datasets)

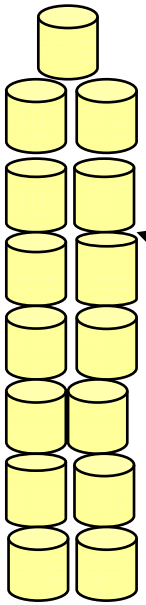

** 2105 model 800 (Silvertip Shark) with 32G cache (3390-9)

SDUMP Volumes*

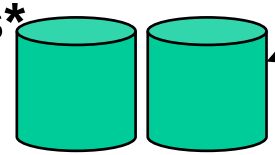

* 2107 model 922 (MegaMouth Shark) with 116.7G cache (3390-54)

4-Way z10
(25GB real for 10GB dump
45 GB real for 20 GB dump)

Configuration set up for 40GB dump

ASM Paging Packs**
(with 6.74GB datasets)

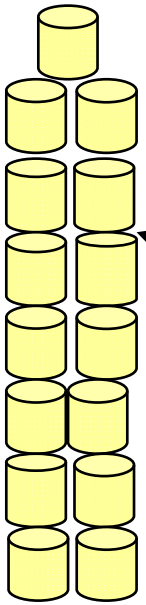

** 2105 model 800 (Silvertip Shark) with 32G cache (3390-9)

SDUMP Volumes*

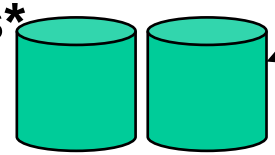

* 2107 model 922 (MegaMouth Shark) with 116.7G cache (3390-54)

4-Way z10
(85 GB real)

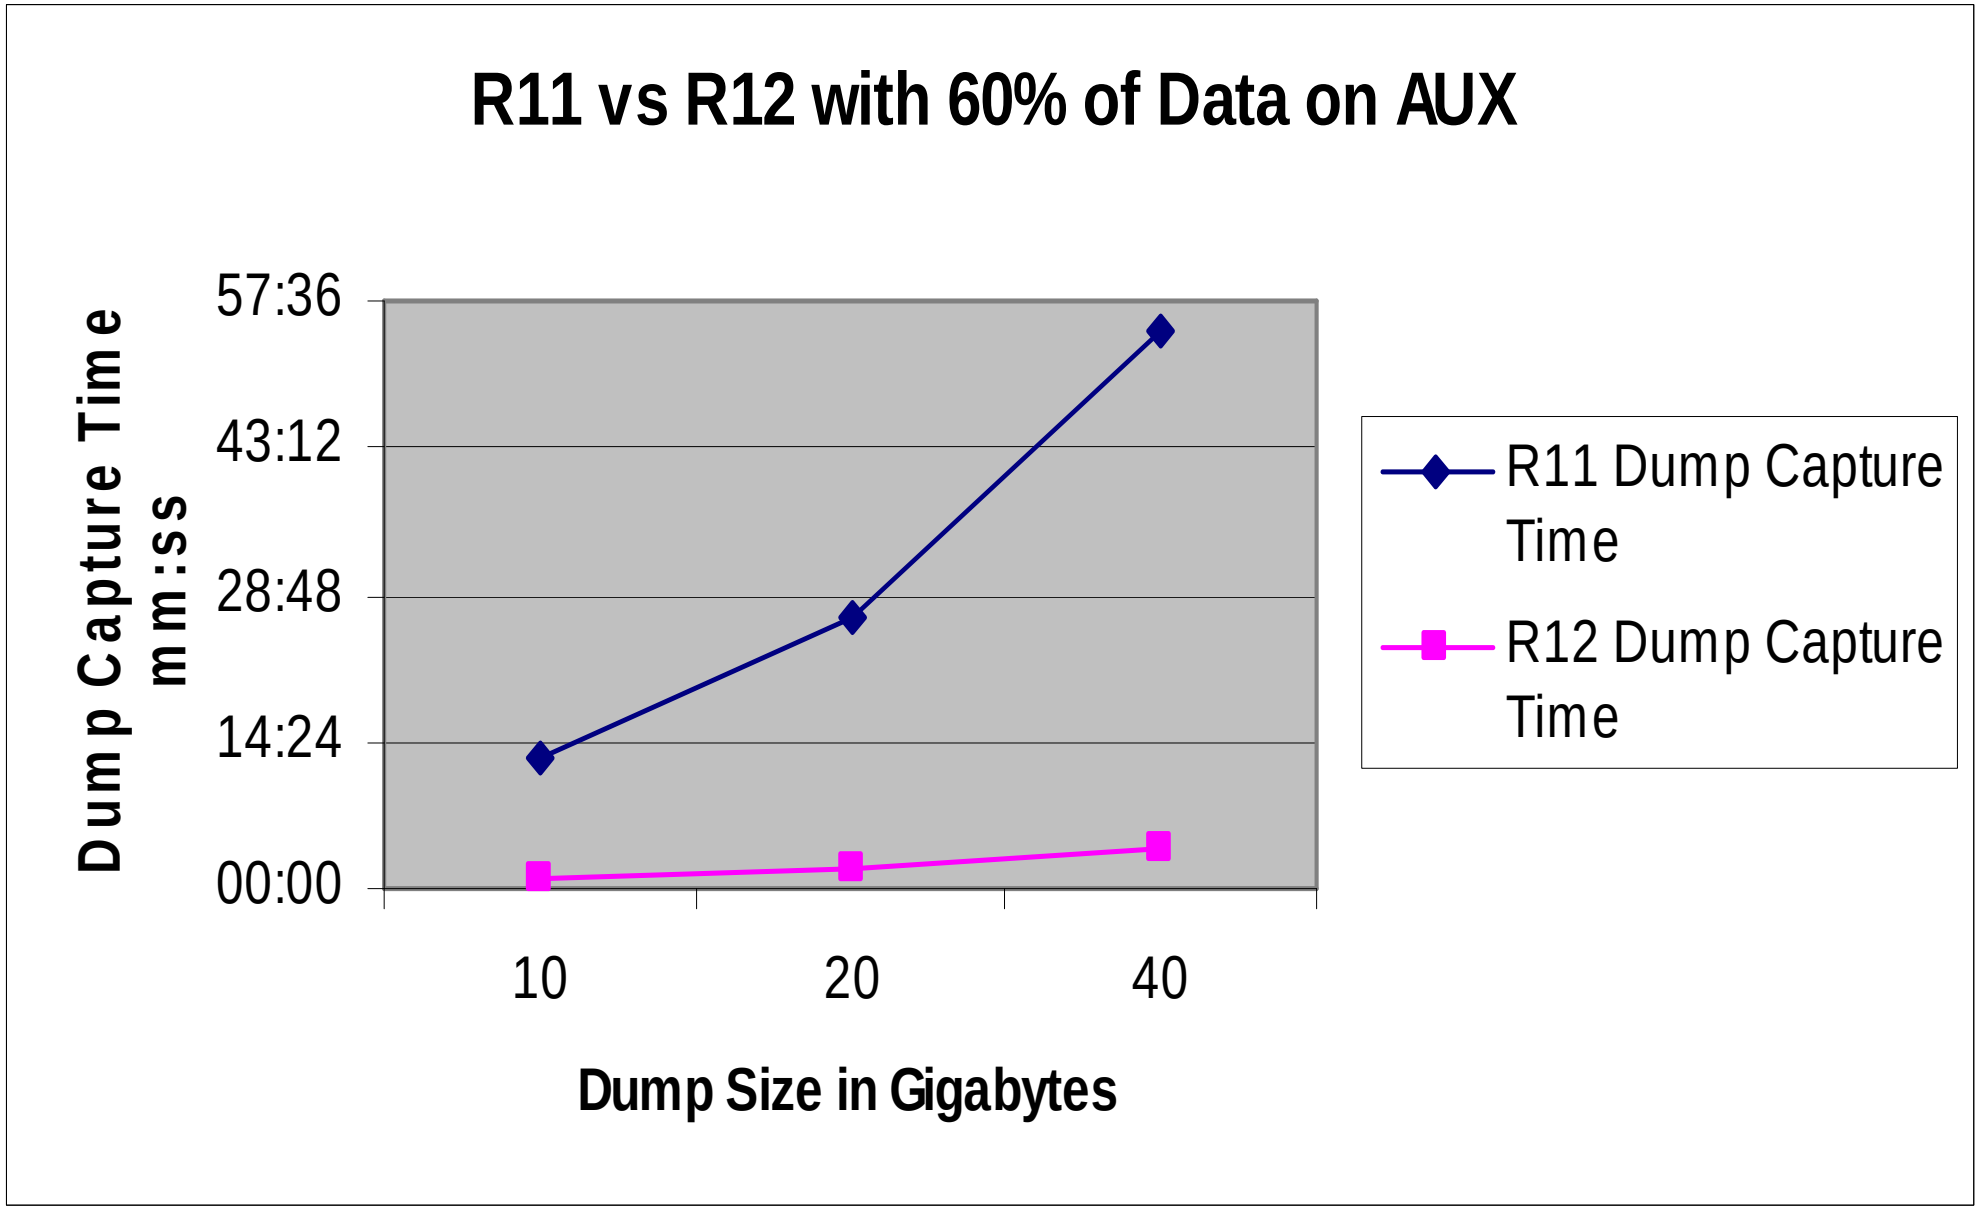

R11 vs R12 60% of data out on AUX No Real Storage Constraints

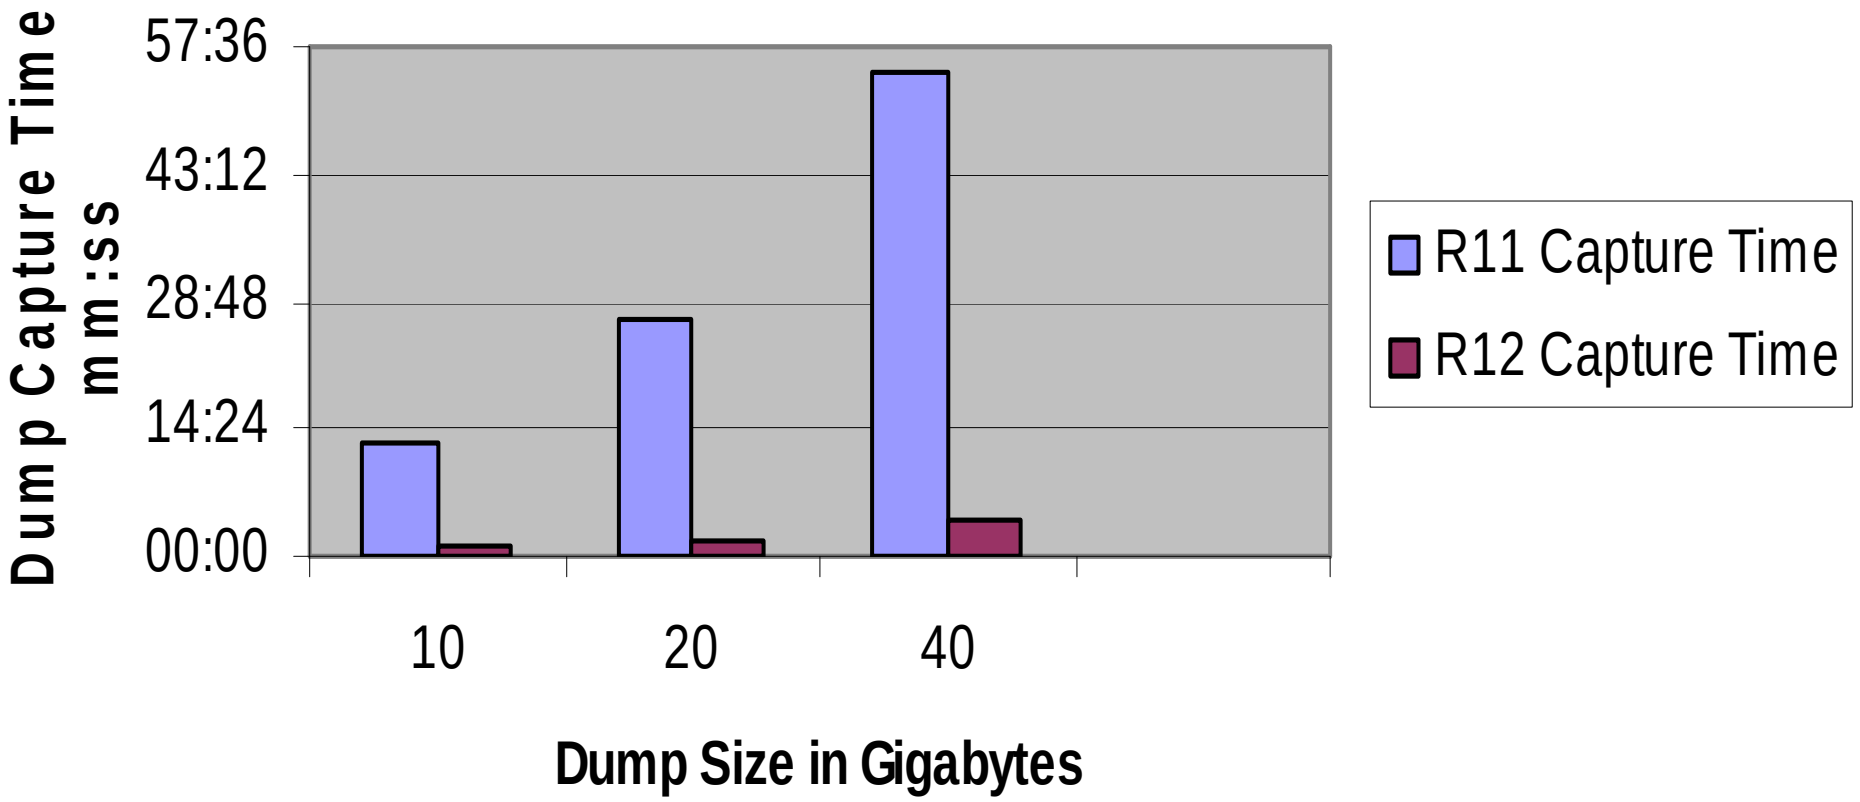

R11 vs R12 with 60% of Data on AUX Real Storage Constrained

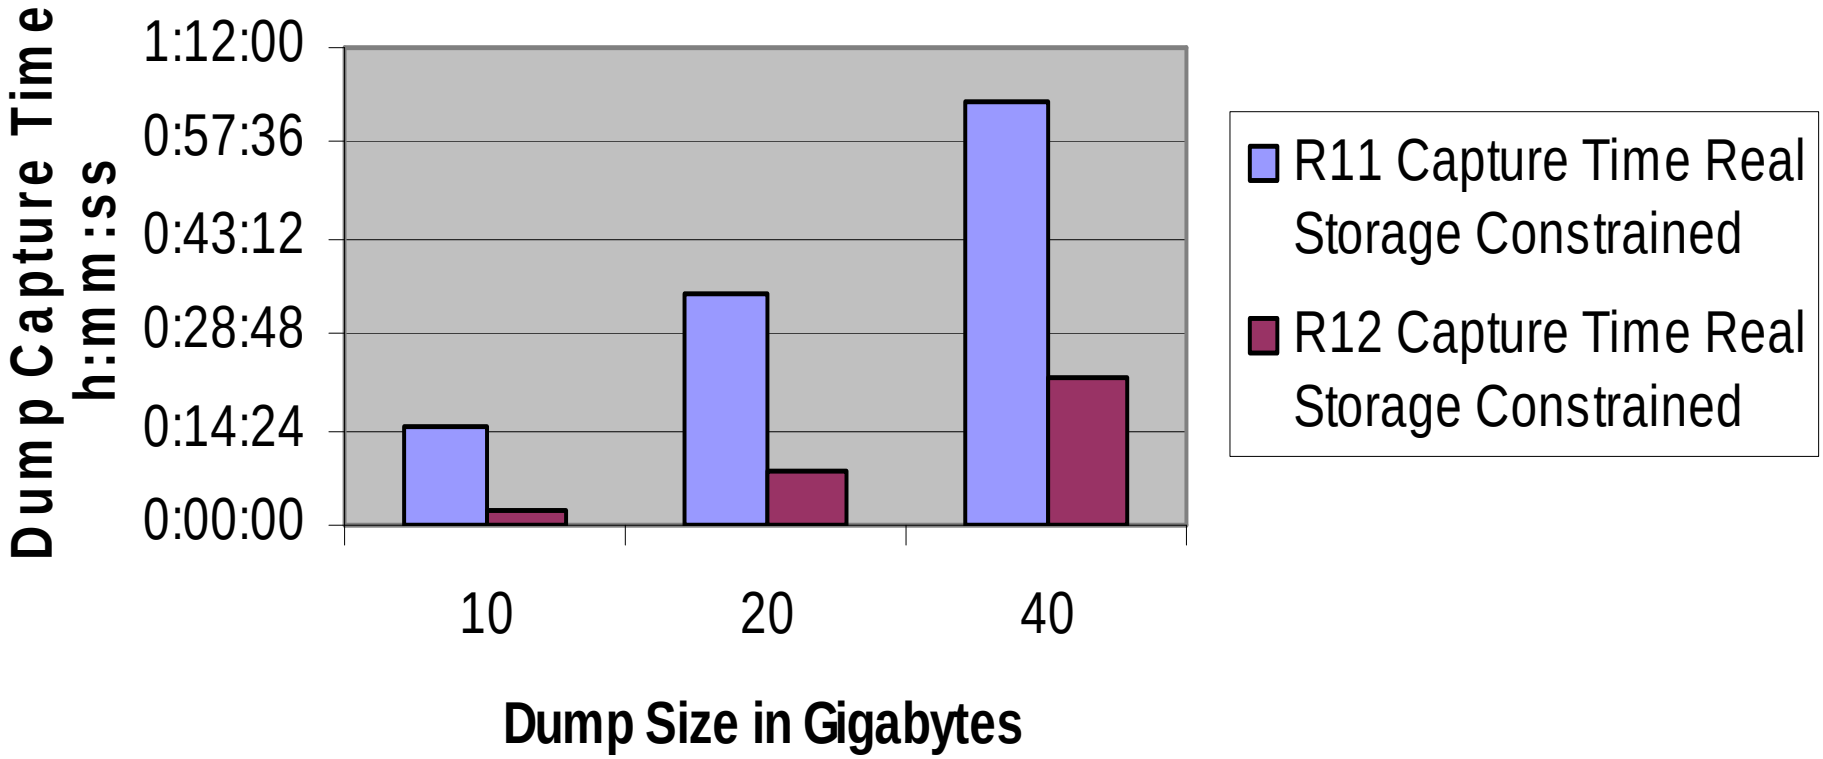

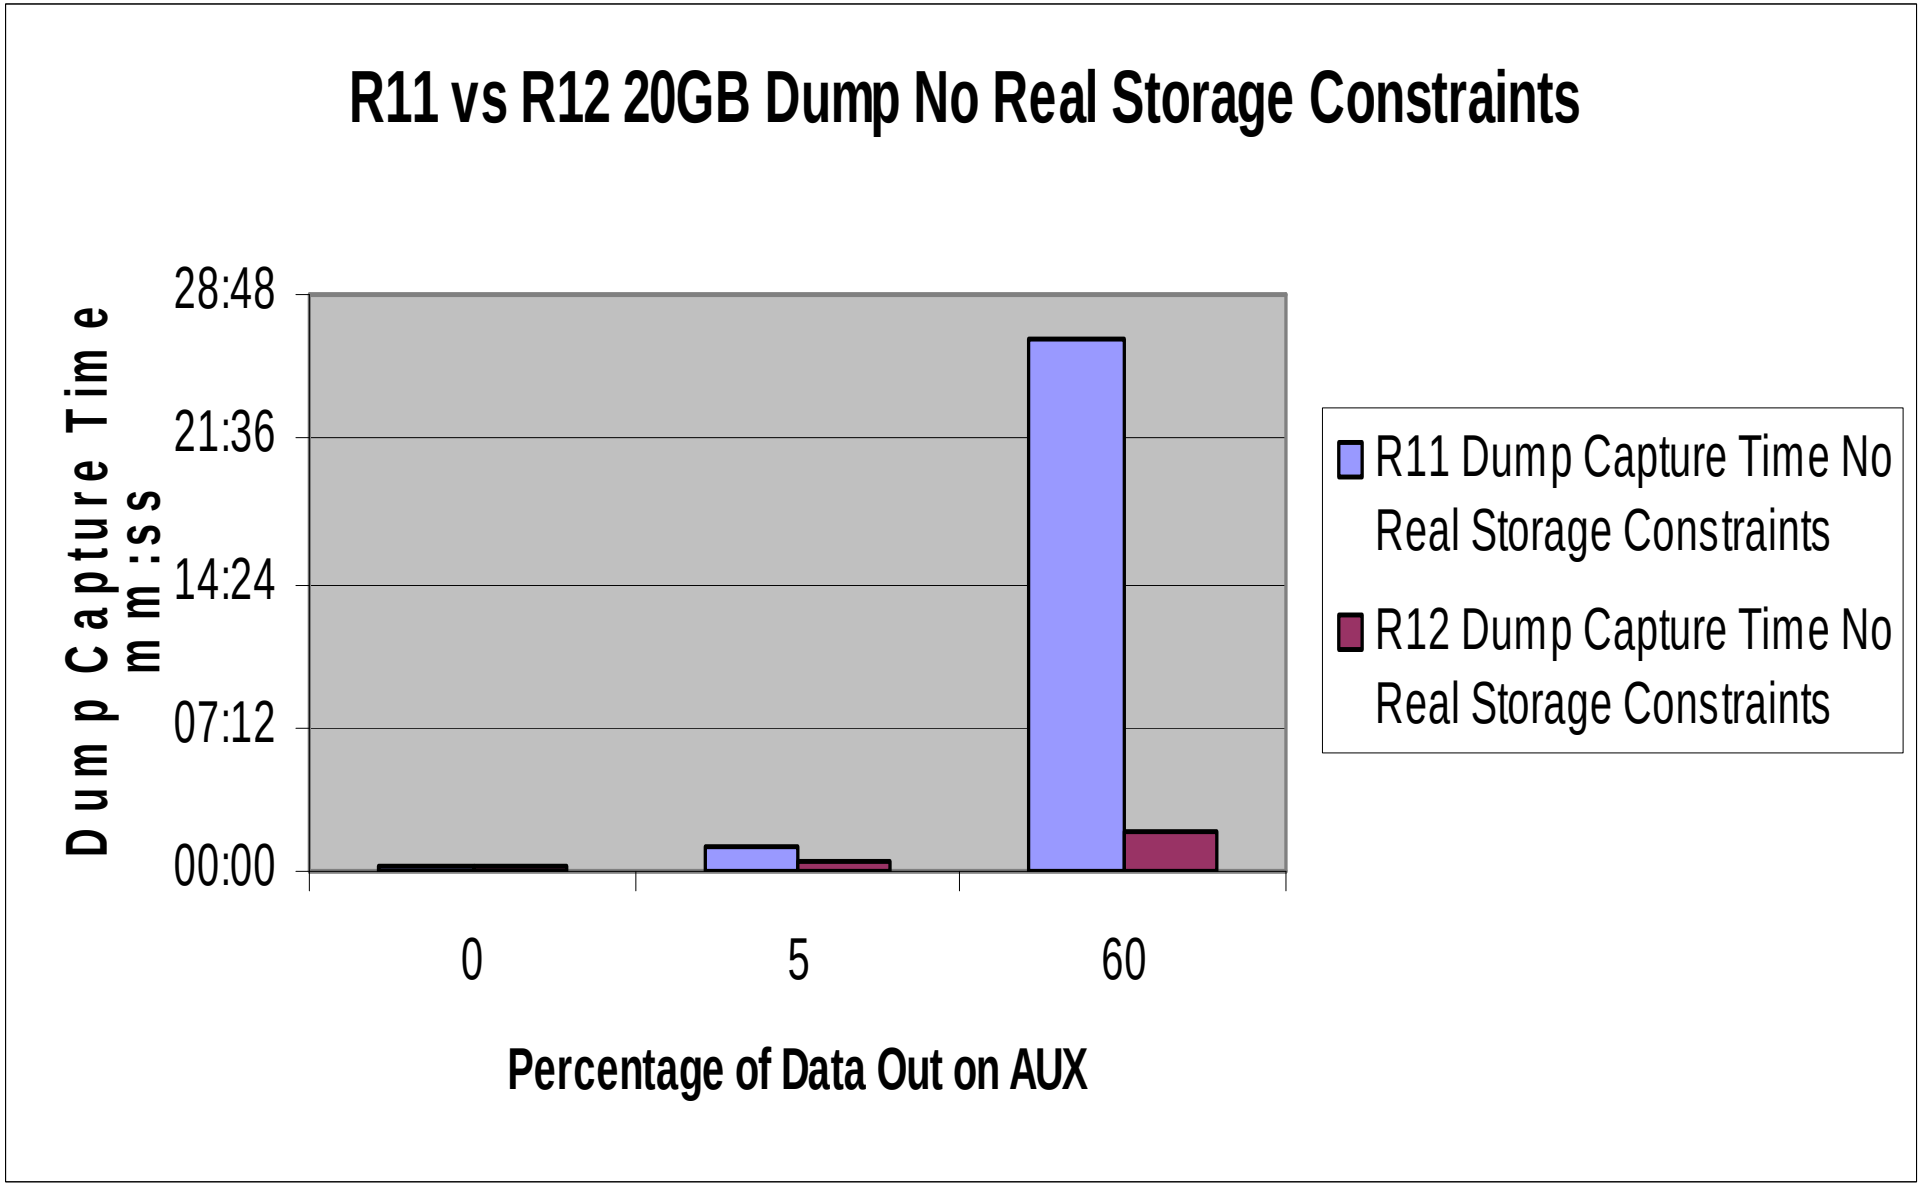

R11 vs R12 10GB Dump Real Storage Constrained

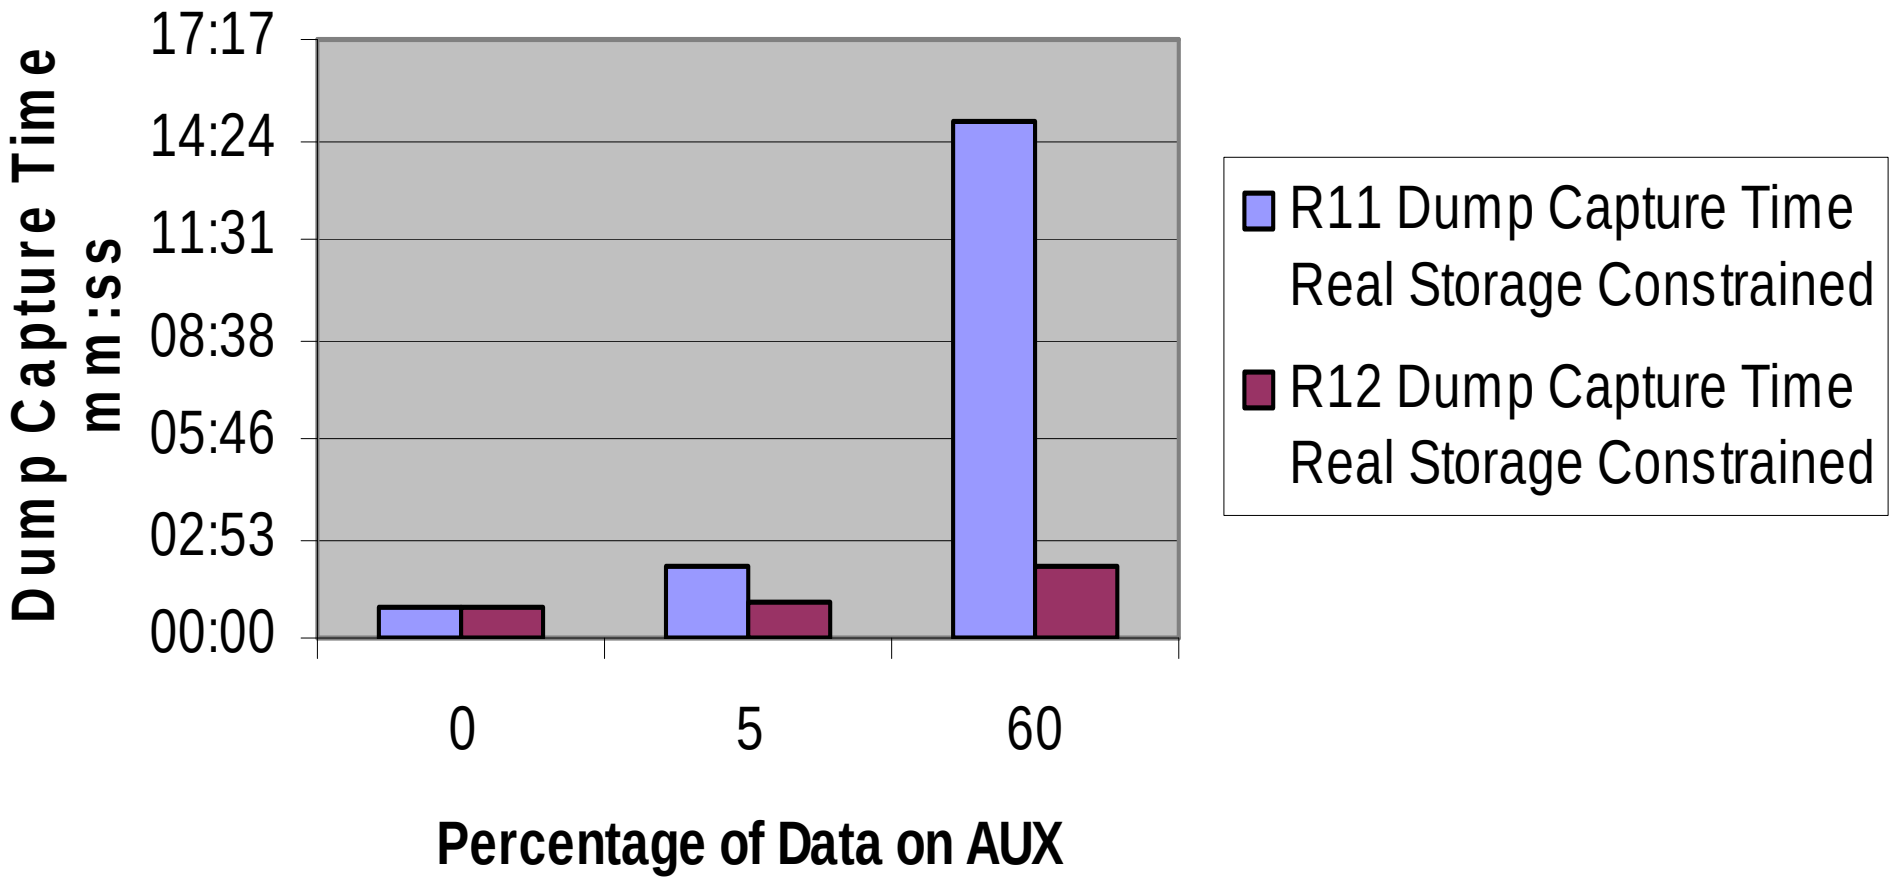

R11 vs R12 10GB Dump No Real Storage Constraints

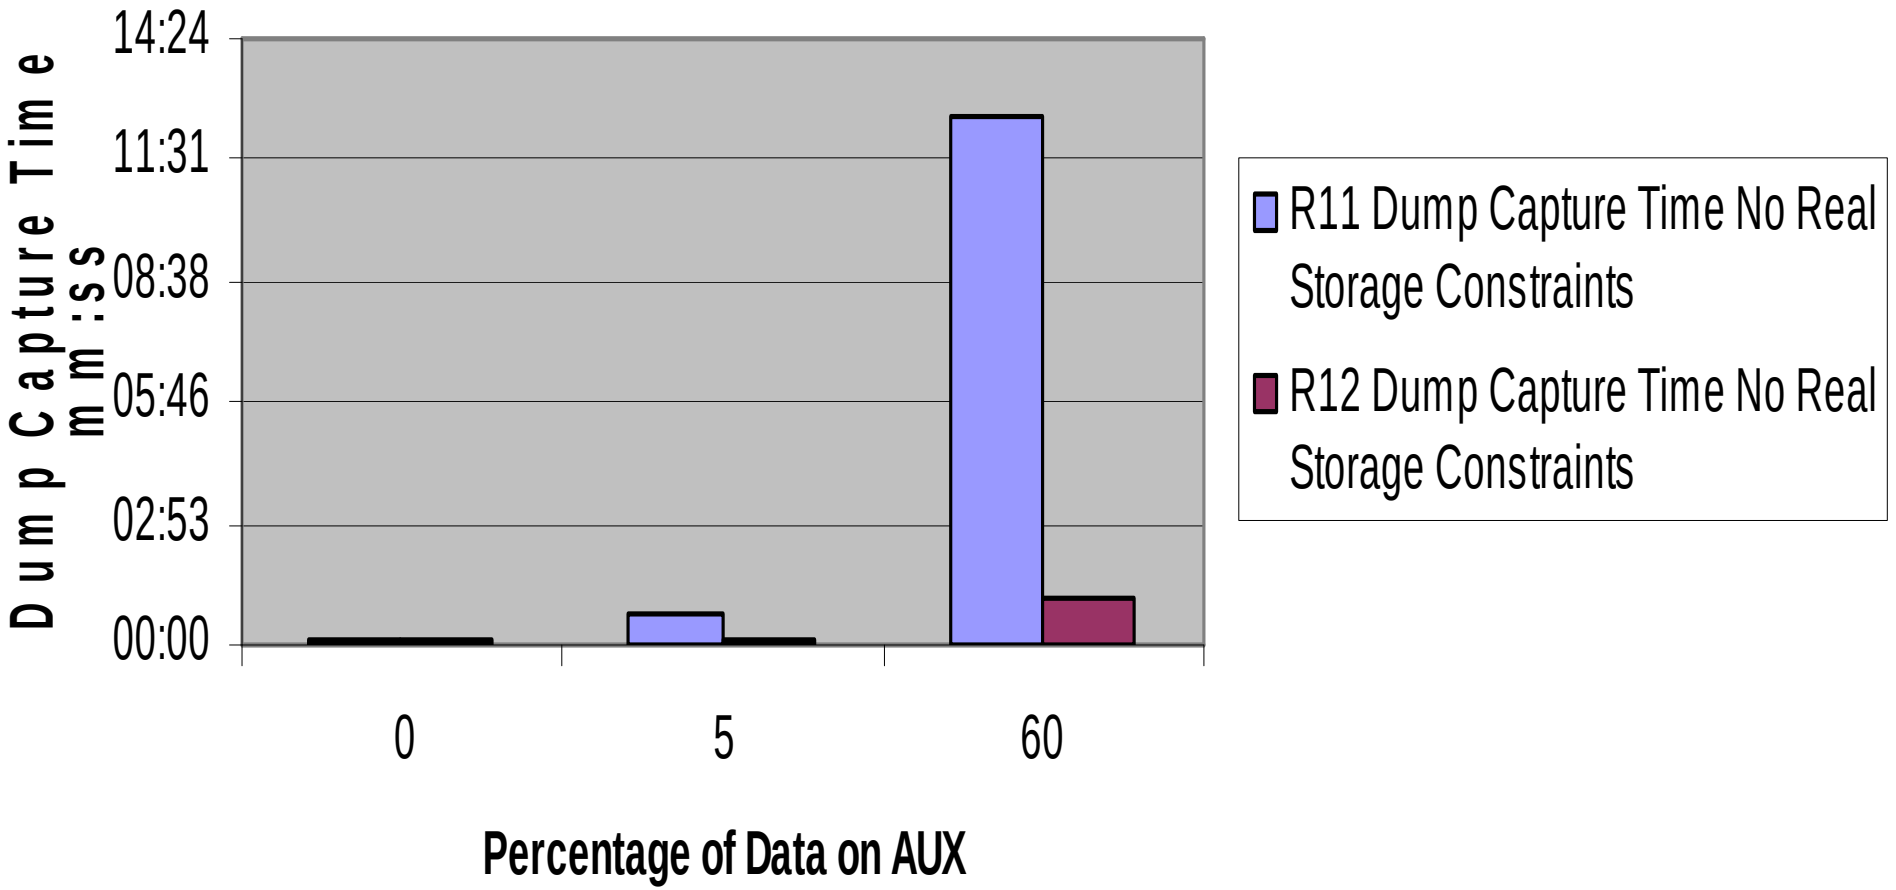

Performance Results: R12 vs R11 with 60% of data out on AUX

System: / Size Of Dump	No Real Storage Constraints	Real Storage Constrained	Capture Time (No constraints)		Capture Time (Constrained)	
	% Performance Improvement		R11 VS R12		R11 VS R12	
10 GB	91%	86%	12:35	1:03	14:52	2:00
20 GB	92%	76%	26:37	2:00	35:14	8:13
40 GB	93%	65%	54:36	4:05	1:03:56	22:34

- **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%.**

Performance Results: R12 vs R11 with 5% of data out on AUX

System: / Size Of Dump	No Real Storage Constraints	Real Storage Constrained
	% Performance Improvement	
10 GB	78%	50%
20 GB	79 %	52%
40 GB	82%	50.4%

- 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.

Performance Results: R12 vs R11 (No Real Storage Constraints)

Size of Dump: / % of data on Aux:	10 GB	20 GB	40 GB
	% Performance Improvement		
0%	4% (0:08 vs 0:07)	6% (0:16 vs 0:15)	no data
5%	78% (0:43 vs 0:09)	79% (1:08 vs 0:23)	82%
60%	91% (12:35 vs 1:03)	92% (26:37 vs 2:00)	93% (54:36 vs 4:05)

- **As the % of data on Aux increases, the observed performance benefit increases *dramatically*, from 4% up to 91%.**
- **For no real storage constraint environments, the best performance data observed (93%) was with the larger dump (40GB) and where more data (60%) was on Aux.**

Performance Results: R12 vs R11 (Real Storage Constrained)

Size of Dump: / % of data on Aux:	10 GB	20 GB	40 GB
	% Performance Improvement		
0%	3.8%	Data not available	Data not available
5%	50%	50%	48%
60%	86.5%	76%	65%

- As the % of data on Aux increases, the observed performance benefit increases *dramatically*, from less than 4% up to almost 90%.
- For real storage constrained environments, the best performance data observed was with the smaller dump (10GB) since the real memory needed to capture the dump was less than the real memory needed for the bigger dumps.

Performance Results: R12 vs R11 with SSDs

Percent of Data on AUX	No SSDs Capture Time		SSDs Capture Time	
	R11 VS	R12	R11 VS	R12
5%	11:23	5:39	7:50	3:12
60%	1:03:56	22:47	43:01	8:57

40GB Dump with 60% of data out on AUX real storage constrained

Performance Results: R12 vs R11 with varying the AFC

Available Frame Queue Count	Capture Time	
	R11 VS R12	
	mm:ss	mm:ss
8GB	01:25	00:15
3GB	02:16	00:38
1.5GB	02:27	00:56

6GB dump with 20% of data out on AUX with varying the Available Frame Queue Size

SVC DUMP 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).

Configuring AUX Storage Needs for SVC Dump

Aux Strategies

- Exhausting auxiliary storage during SVC Dump processing will result in an outage
- System Resources Manager (SRM) identifies a shortage of auxiliary storage when 70% of the slots are in use. So, it is recommended that you plan for a maximum of 60% of slot utilization by your workloads and dump processing.
- The monitoring of the maximum number of slots used should be continuous rather than an observation.
- Examine the “RMF Paging Activity Report” to determine the “high water mark” of auxiliary slots in use

Overview ... Postprocessor Paging Activity Report

FRAME AND SLOT COUNTS

	CENTRAL STORAGE		
(91 SAMPLES)	MIN	MAX	AVG
AVAILABLE	2068166	2074202	2071567
SQA	19,220	19,295	19,263
LPA	4,751	4,759	4,757
CSA	16,588	16,609	16,601
LSQA	49,851	50,763	50,165
REGIONS+SWA	454,309	459,554	456,700
TOTAL FRAMES	2621440	2621440	2621440

	FIXED FRAMES		
NUCLEUS	2,385	2,385	2,385
SQA	15,622	15,697	15,665
LPA	67	67	67
CSA	8,583	8,599	8,594
LSQA	14,466	14,569	14,516
REGIONS+SWA	36,129	36,693	36,538
BELOW 16 MEG	97	110	102
BETWEEN 16M-2G	24,058	24,630	24,467
TOTAL FRAMES	77,388	77,928	77,767

LOCAL PAGE DATA SET SLOT COUNTS			
	MIN	MAX	AVG
AVAILABLE SLOTS	3,346,554	4,183,193	3,585,594
VIO SLOTS	0	0	0
NON-VIO SLOTS	597,599	1,434,238	1,195,198
BAD SLOTS	0	0	0
TOTAL SLOTS	4,780,790	4,780,790	4,780,790
Percent Used	13%	30%	25%

SHARED FRAMES AND SLOTS			
CENTRAL STORAGE	9,119	9,185	9,136
FIXED TOTAL	38	39	38
FIXED BELOW 16 M	0	0	0
AUXILIARY SLOTS	0	0	0
TOTAL	9,119	9,185	9,136

MEMORY OBJECTS AND FRAMES			
OBJECTS COMMON	**6**	**6**	**6**
SHARED	**0**	**0**	**0**
LARGE	**1**	**1**	**1**
FRAMES COMMON	**200**	**200**	**200**
COMMON FIXED	**0**	**0**	**0**
SHARED	**0**	**0**	**0**
1 MB	**2**	**2**	**2**

Aux Strategies…

- In the example configuration above the page datasets support 18.2GB of space
  - There 256 pages/slots per MB
  - Total Auxiliary slots available = 4,780,790/256=18,675MB
  - 18,675MB/1024=18.2GB
- Plan for a maximum of 60% of slot utilization by your workloads and dump processing.
- The configuration used in this example could accommodate
  - (60%-30%)=30% of slots left for dump use
  - 30%*18.2=5.5GB of slots for dumps
- These values should be reviewed frequently

Aux Strategies…

- For a rough approximation of the largest dump that may be taken sum up the following:
  - Max amount of CSA allocated and backed in real plus aux+ max amount of real and auxiliary storage in use by your 6 largest address spaces
    - Use the RMF STORF report
  - Multiply the above number by the number of dumps taken before they can be written to dump datasets (perhaps 5-6)
- A better approximation may be available from dumps already taken

Overview ... Monitor III STORF report

```
                        RMF V1R9   Storage Frames                    Line 1 of
219
 Command ===>                                                     Scroll ===>
PAGE

 Samples: 120     System: SCLM  Date: 07/12/06  Time: 11.19.00  Range: 120
Sec

            Service      -- Frame Occup.--  - Active Frames -  AUX   PGIN Large
Jobname  C Class     Cr Total  ACTV  Idle   WSET Fixed   Div Slots Rate Page

OMVS     S SYSTEM       228K  228K     0   228K  1569     0  216K    0    0
SMSPDSE1 S SYSTEM       108K  108K     0   108K  1457     0  196K    0    9
SMSVSAM  S SYSTEM     27448 27448     0 27448   647     0 12126    0    0
RMF      S SYSSTC     21713 21713     0 21713   157     0 11023    0    0
SEL55530 B BATCHMED   20463 20463     0 20463   208   655     0    0   15
DFSZFS   S SYSSTC     17622 17622     0 17622   380     0 24218    0    0
PEGBLD7  O OE         12077  2077     0 12077   141     0     0    0    0
GRS      S SYSTEM       8836  8836     0  8836   567     0 11406    0    0
XCFAS    S STCSYS       8122  8122     0  8122  2331     0  1211    0    0
IXGLOGR  S SYSTEM       7943  7943     0  7943   140     0   620    0    0
SPASCM   S STCCMD       5907     0  5907     0   136     0     1    0    0
RMFGAT   S SYSSTC       5871  5871     0  5871   102     0  1238    0    0
*MASTER* S SYSTEM       5818  5818     0  5818  5473     0   643    0    0
VLF      S SYSSTC       5253  5253     0  5253   141     0  1334    0    0
NET      S SYSSTC       4862  4862     0  4862   182     0  2397    0    0
```

Questions ?

SVC Dump Enhancements

Performance Test Results with Internal Algorithm Improvements (Back-Up Slides)