

SMF Buffers: How NOT to Loose Your Precious Assets

Richard S. Ralston
Humana, Inc.

August 4, 2010
Session Number 6836



SHARE in Boston

SMF Buffers

- Interesting word – Assume
 - SMF processing had been handled by many in the past
 - Assumed everything was OK
- Several occurrences of lost SMF data in 2009
- Guess who got to fix the problem?
- Initially had a small SMF Buffer space and small Clsize
 - 4k Clsize
 - 256m buffer space
- Upgraded to max
 - 26K CI size
 - 1 gig buffer space

SMF Buffers — Do You Have Enough?

- Watch and track data from your type 23 records.
- 2 types of records
 - SMF dataset switch
 - 15 minute SMF status records ****
- Make sure you have current SMF maintenance
- Use either 1/3 or 1/2 track Csize
- For high activity LPARs use 1 gig buffer space
 - BUFFSIZEMAX(1G) in SMFPRM00
- For other LPARs the type 23s will provide clues to determine buffer space
- <https://www.ibm.com/developerworks/wikis/display/zosperfinstr/Home> - Martin Packer

Console Messages

- *IEE986E SMF HAS USED 50% OF AVAILABLE BUFFER SPACE
- Messages start at 50% buffer space used
 - Actually BUFUSEWARN in SMFPRM00, Default=25%
- 1 message for every 1% increase in buffer space, or 8m whichever is larger
- Messages stop when decreasing and reaches BUFUSEWARN
 - Fewer messages as space decreases
- If NOBUFFS(MSG) in SMFPRM00 (default?) then IEE979W message and buffers are lost until buffer space becomes available


SMF Type 23 Status Record

- 15 minute interval record
- Contains:
 - SMF23BFA Number of Buffer Allocation Requests
 - SMF23BFH High Water Mark of Storage Allocation
 - SMF23BFW Number of SMF Buffers Written
 - SMF23BFT Total Buffer Storage Allocated
 - SMF23SUS Number of SMF Buffer Suspensions
 - SMF23BFL Percent Buffer Usage Warning Level Reqs
 - SMF23RCW SMF Records Written
 - Amongst other things.
- MXG & MICS can process these records


Issue - Situation

- SMF gets bent out of shape if a big bunch (>2000?) of big records hit it within 1 second followed by continued high SMF activity.
- Normally (98+% of the time) SMF handles all records in 8 meg of buffers, no matter what you have specified.
- When 'bent', SMF gets/adds buffers in 8 meg increments but may only use 1 buffer from each increment.
- Possible to 'use up' all buffer space in a matter of seconds
- Handful of APARs on this, all improve the situation, but it is not completely eliminated.
- Use SMF23BFH to determine how big to make SMF Buffer Space, max is 1 gig


Hit 92% of Available Buffer Space




Buffers vs. SMF Records



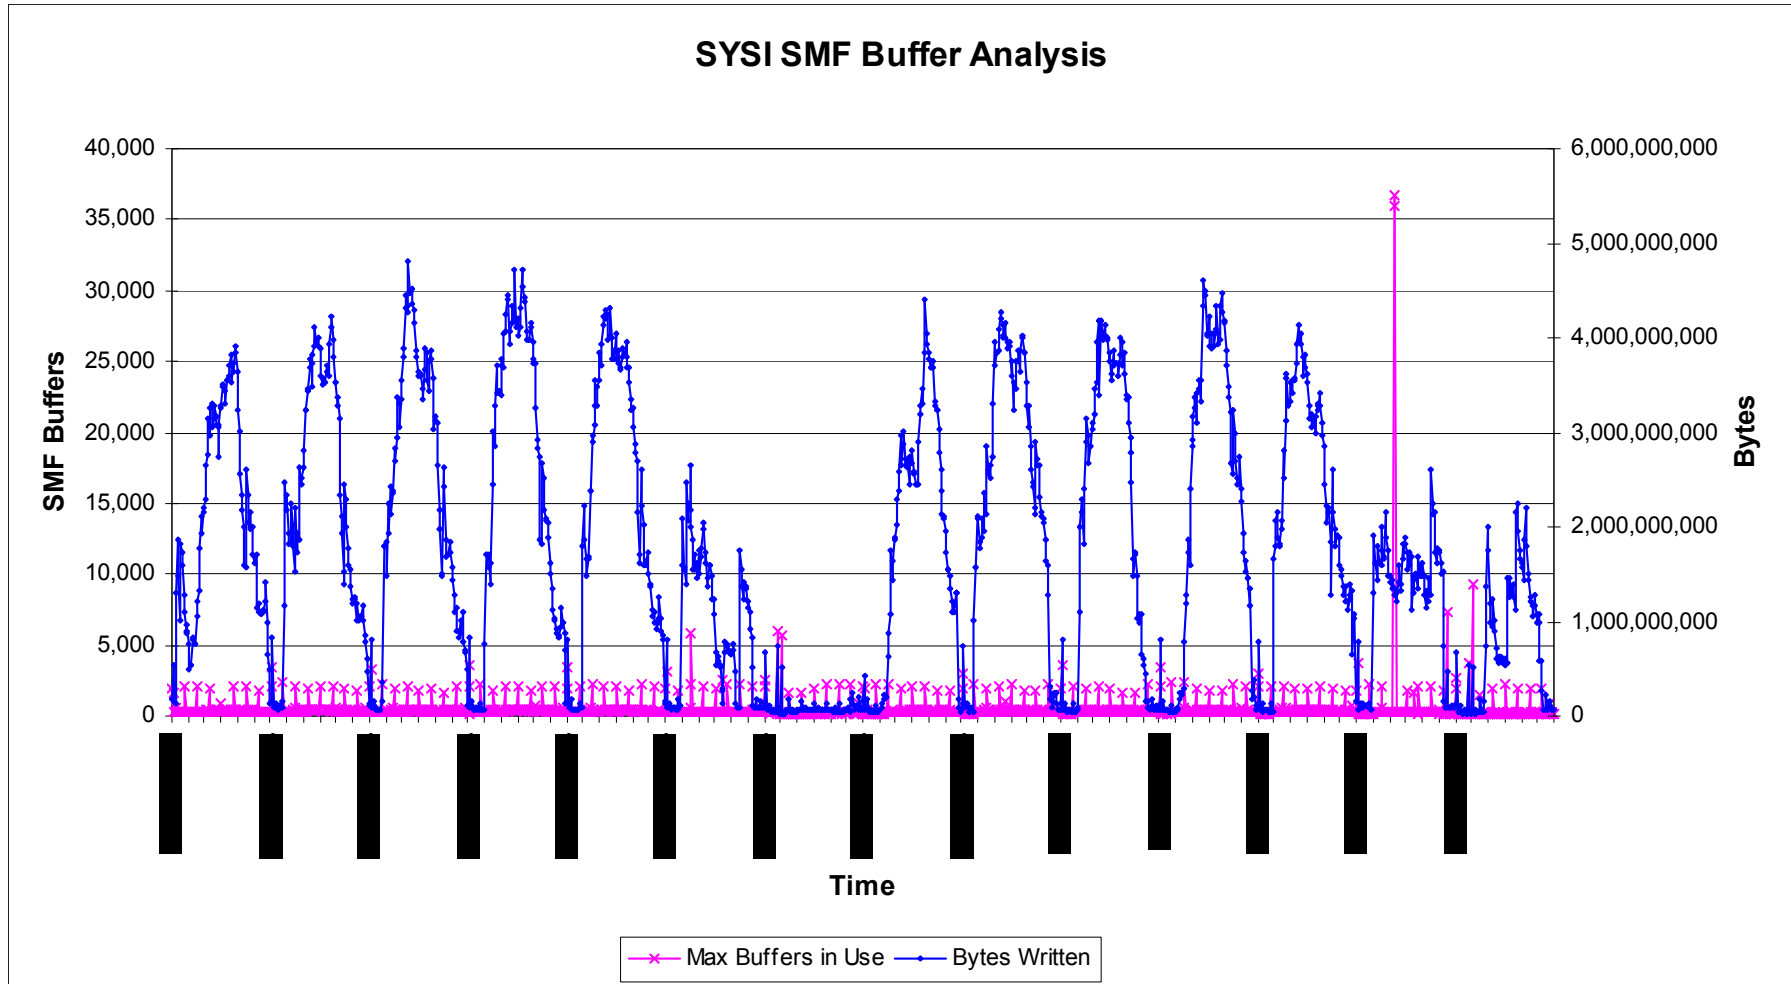
Buffer Turnover




SMF/Buffer



Bytes Written



Buffer Memory Used

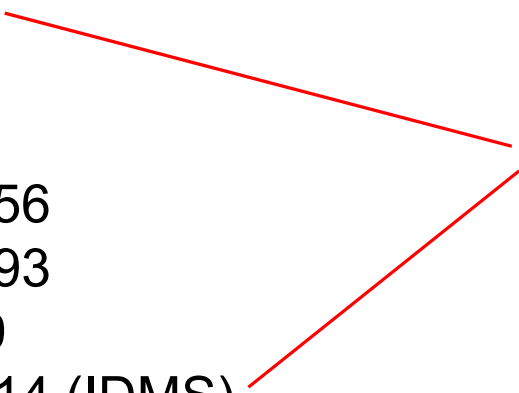


Analysis

- Counted every SMF record type & subtype written for 15 minutes before and after the event in 1 second increments.
 - CSV file and loaded into Excel
- Attempt to determine exactly what type of records were written, when and how many.
- Eyeballed the data on a large screen
- Used SAS & MXG code – homegrown

- MXG module ANALSMF which can provide useful information

Ultimate Cause?

- 3521 SMF records written at 9:00:01
 - 3230 were 110-2
 - Event Duration: 13:30 – very long
 - 50% buffer full: 9:07:49
 - 92% buffer full: 9:13:31
 - 70916 total records
 - Type 14: 9509
 - Type 30: 2369
 - Type 42: 4930
 - Type 101: 12,856
 - Type 110: 12,793
 - Type 115: 5029
 - Type 254: 12,714 (IDMS)
 - Plenty of CPU, CPC: 40%, LPAR 20%
 - Root Cause: Still Unknown – I/O Issues?
- 60,200 records, 84.9% of total
- 

Different Time, Different LPAR


- December 27 - 04:06:51 – Test LPAR
- Hit 82% buffer full – 1 gig buffer space
- Similar SMF analysis by second
- Cause: 10 Test DB2 regions shutdown at the same time - automation
- Solution: stagger DB2 shutdowns by 1 minute

Total recs
45203

t30-4	t30-5	total	%
17598	17598	35196	77.9%

Start	End	Duration	recs/sec
4:06:22	4:06:51	0:00:29	1213.655

DB2 Shutdowns SMF Impact

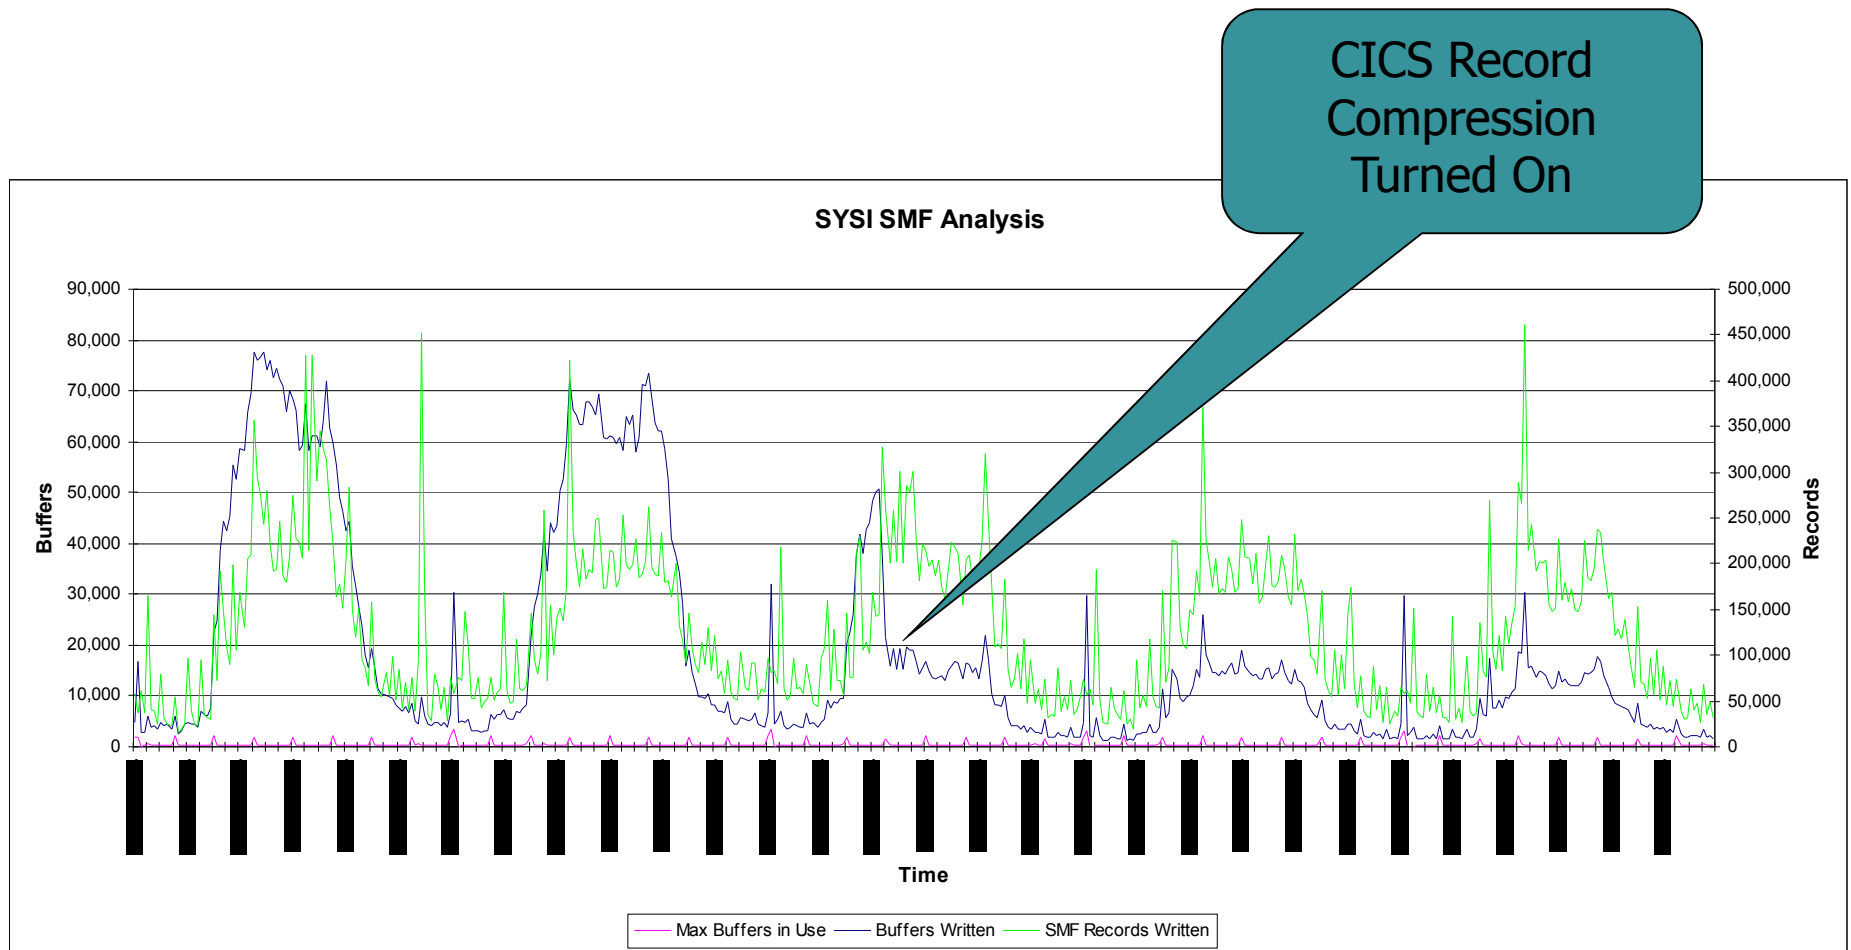


It looks like if the shutdowns occur too fast and use over 30,000 buffers we impact the SMF process.

Assistance: CICS Record Compression

- Compresses 110 records as they are written
- 110 records are normally 32K
- After compression they are about 6K
- Does not reduce the number records
- Does reduce the number of SMF buffers written
- Can be turned on dynamically
- Default in CICS V4, optional in CICS V3

CICS Record Compression



Assistance: DB2 ACCUMAC

- DB2 accounting data for DDF and RRSF threads is accumulated by end user
- Default: ACCUMAC=10
 - Up to 10 SQL calls are accumulated/summarized into 1 record.
- Was set to NO on ASG's recommendation (TMONDB2)
 - SQL Explain does not work if not set to NO
- Currently testing
 - DBAs use TMONDB2 and Detector
 - We will probably not use this

Assistance: **EMPTYEXCPSEC(SUPRESS)**

- SMFPRMxx option
 - Available with z/OS 1.10
- Suppress empty execute channel program (EXCP) entries in the SMF type 30 record
- Specify EMPTYEXCPSEC(SUPPRESS) to ensure that the system generates no empty EXCP sections for non-allocated candidate volumes in the SMS storage group.
- Default: EMPTYEXCPSEC(NOSUPPRESS)
- Currently testing to ensure no issues.

CPExpert

- CPExpert V20.1 has added SMF & type 23 analysis
 - WLM760 - *Large percent of SMF buffer space allocated to active buffers*
 - WLM761 - *IEE986E BUFUSEWARN message was produced*
 - WLM762 - *SMF data was lost because no buffer space was available*
 - WLM763 - *Control Interval (CI) size might be too small for SMF data sets*
 - WLM764 - *SMF recording was switched to more than two data sets*
 - WLM765 - *SMF records were written to the SMF data set with recording spikes*
 - WLM767 - *NOBUFFS(HALT) was specified in SMFPRMxx*
 - WLM768 - *LASTDS(HALT) was specified in SMFPRMxx*

SMF Logger

- Intended to fix SMF data loss problems
- Gets better with each z/OS release
- Move to it when you think its ready

- We will review and probably implement SMF Logger in 4Q10 after we have finished our migration to z/OS 1.11

More Information

- Cheryl Watson's Tuning Letter has discussed SMF buffer issues in a few recent issues
- MXG Newsletters (www.mxg.com)

Contact me

- rralston@humana.com
- 502-476-1121
- cyclotourist@bellsouth.net
 - Facebook & LinkedIn
- rsrcyclotourist@gmail.com