

May 6-10, 2007  
San Jose Convention Center  
San Jose, California, USA

# DB2 for z/OS System Performance Basics

## A DB2 Statistics Review

Session code: A06



IDUG® 2007  
North America

Randy Ebersole IBM

May 8, 2007 10:40 a.m. – 11:40 a.m.

Platform: DB2 for z/OS

GoFurther

There is a wealth of information available from DB2 statistics. In this presentation we will discuss DB2 statistics and how to interpret these statistics based on various production scenarios.

We'll focus on what the many statistics mean and steps to take on how to address the issues to improve DB2 performance.

Our main objective in this session is to improve the understanding of the DB2 Statistics information and allow the attendee to increase the ability to exploit this information in their specific DB2 environments.

## Presentation Objectives

- How DB2 provides DB2 Statistics
- Five primary areas to focus on when reviewing DB2 Statistics information
- How to learn about their specific environment and take steps to improve
- DB2 system and application performance
- Interpretation of DB2 workload characteristics
- What new information is available in DB2 Statistics

2

GoFurther

### Abstract Outline:

- General Statistics Information
- Pools - Buffer, RID, EDM
- DB2 Storage Information
- Thread Management
- Locking Analysis

## Who Are We?

North American DB2 for z/OS Lab Services  
DB2 Information Management, IBM Software Group

---

The DB2 for z/OS Software Services Team delivers specialized services to accelerate your implementation and to fully leverage and exploit your investment in DB2 for z/OS technology. Our consultants are highly skilled in DB2 systems administration, DB2 application performance, DB2 scalability and DB2 high availability. We also specialize in Version to Version migration assistance for DB2 for z/OS.

We also have peer groups who specialize in DB2 on other platforms.

---

**For additional information, contact:**

Randy Ebersole  
ebersole@us.ibm.com

GoFurther

## Agenda

- Statistics – Getting Started
- Statistics – Time – It's Important
- Statistics – A review and things I look at.....
- Statistics – New Things....

Focus is on the Statistics "Long" Report

## DB2 Statistics Information – Page 1

- SMFSTAT = YES
  - Statistics trace classes 1, 3, 4, 5, and 6
- Class 1 - Information on system services and database statistics
  - Includes system parameters in effect
- Class 3 - Information on deadlocks and timeouts
- Class 4 - Information on exceptional conditions
- Class 5 – Information on data sharing
- Class 6 – Information on storage statistics for the DBM1 address space
- SMF Statistics data in both SMF type 100 and 102 records
  - IFCIDs 0001, 0002, 0202, and 0230 are of SMF type 100
  - All other IFCIDs in statistics trace classes are of SMF type 102
- Statistics collection interval(STATIME) – Interval of time statistics trace is written

5

GoFurther

SMF (System management facilities) collects and records system and job-related information

SMF collects data, writes to SMF datasets, with a periodic Dump to Dump datasets for ongoing use

For DB2:

Record Type 100 – Statistics - DB2 writes record type 100 to record data collected at event monitoring points

Record Type 101 – Accounting - DB2 writes record type 101 to account for resources during a transaction

Record Type 102 - DB2 writes record type 102 to record performance information (Audit also)

**Accounting and statistics traces** Enabling accounting class 2 along with accounting classes 1 and 3 provides additional detail relating directly to the accounting record IFCID 0003, as well as recording thread level entry into and exit from DB2. This allows you to separate DB2 times from application times. Running accounting class 2 does add to the cost of processing. How much overhead occurs depends on how much SQL the application issues. Typically, an online transaction incurs an additional 2.5% when running with accounting class 2. A typical batch query application, which accesses DB2 more often, incurs about 10% overhead when running with accounting class 2. If most of your work is through CICS, you most likely do not need to run with class 2, because the class 1 and class 2 times are very close.

**Exception:** If you are using CICS Transaction Server for z/OS 2.2 with the Open Transaction Environment (OTE), activate and run class 2.

## DB2 Statistics Information – Page 2

- DB2 statistics counters contain DB2 activity between when the record is externalized and when the DB2 system started. There are three types of counters:
  - An accumulated value since the DB2 system was last started. The total number of INSERTS statements executed is an example
  - A snapshot value would be statistics like the number of open data sets at the time the DB2 statistics record was externalized
  - A high water mark value is a statistic that would represent the highest value since the time the system was last started. since the system was last started. An example of this would the high water mark for CREATE THREAD

## Time and Time.....

---- HIGHLIGHTS -----  
 INTERVAL START : 02/12/06 07:01:20.81  
 INTERVAL END : 02/12/06 07:16:21.54  
 INTERVAL ELAPSED : 15:00.732524

LOG ACTIVITY	QUANTITY	/SECOND
LOG RECORDS CREATED	1875.6K	1043.20
LOG CI CREATED	65552.24	36.46
LOG WRITE I/O REQ (LOG1&2)	125.1K	69.59
LOG CI WRITTEN (LOG1&2)	177.6K	98.76
<b>LOG RATE FOR 1 LOG (MB)</b>	<b>N/A</b>	<b>0.19</b>
LOG WRITE SUSPENDED	23940.49	13.32

- Per Second statistics provides a good workload comparison
- Also available, Per Thread, Per Commit

GoFurther

## DB2 Time

CPU TIMES	TCB TIME	PREEMPT SRB	NONPREEMPT SRB
SYSTEM SERVICES ADDRESS SPACE	2:24.976187	2:00.000001	7:56.915621
DATABASE SERVICES ADDRESS SPACE	3:42.009408	0:36:00.000001	2:00:06.017125
IRLM	0.051992	0:38.995937	9:00.047929
DDF ADDRESS SPACE	14.253401	14.397188	1:25:31.65058

CPU TIMES	TOTAL TIME	PREEMPT IIP SRB
SYSTEM SERVICES ADDRESS SPACE	12:21.891809	0:00.000000
DATABASE SERVICES ADDRESS SPACE	2:39:48.02653	0:00.000000
IRLM	9:39.047929	0:00.000000
DDF ADDRESS SPACE	1:25:31.65058	0:00.000000
TOTAL	4:27:20.61686	0:00.000000

8

GoFurther

TCB Example - Time under a DB2 address space TCB running asynchronously associated With DB2 logging and BP writes

SRB Example: Scheduled under the DB2 address spaces to manage the asynchronous read engine for sequential prefetch.

DDF Workload - DDF DIST AS SRB CPU - All the CPU used by DB2 to process Non-nested distributed connect work will get charged to the DIST AS (SRB mainly)  
 In a CICS DB2 environment, the processor time from the DB2 accounting records is typically much greater than the processor time reported in the DB2 statistical records, because most of the processor time used is in the thread TCB itself and in the DB2 address spaces using cross memory services.



## Open Datasets

- DS MAX should be larger than the maximum number of data sets that are open/in use at one time
- OPEN/CLOSE ACTIVITY – 1/ sec

OPEN/CLOSE ACTIVITY	QUANTITY	/SECOND
OPEN DATASETS – HWM	7001.00	N/A
OPEN DATASETS	6811.00	N/A
DS NOT IN USE,NOT CLOSE-HWM	6144.00	N/A
DS NOT IN USE,NOT CLOSED	4634.00	N/A
IN USE DATA SETS	2177.00	N/A
DSETS CLOSED-THRESH.REACHED	140.80	<b>2.88</b>
DSETS CONVERTED R/W -> R/O	109.85	0.06

9

GoFurther

Make sure the statistics trace was run at a peak period, so that you can obtain the most accurate maximum figure.

## EDM POOL – PAGE 1

- Quantity and /SECOND statistics available for the following details for EDM POOL Management

PAGES IN RDS POOL (BELOW)  
HELD BY CT  
HELD BY PT  
FREE PAGES  
FAILS DUE TO POOL FULL

PAGES IN STMT POOL (ABOVE)  
HELD BY STATEMENTS  
FREE PAGES  
FAILS DUE TO STMT POOL FULL

PAGES IN RDS POOL (ABOVE)  
HELD BY CT  
HELD BY PT  
FREE PAGES  
FAILS DUE TO RDS POOL FULL

PAGES IN SKEL POOL (ABOVE)  
HELD BY SKCT  
HELD BY SKPT  
FREE PAGES  
FAILS DUE TO SKEL POOL FULL

PAGES IN DBD POOL (ABOVE)  
HELD BY DBD  
FREE PAGES  
FAILS DUE TO DBD POOL FULL

## EDM POOL – PAGE 2

- Quantity and /SECOND statistics available for the following details for EDM POOL Management

DBD REQUESTS  
DBD NOT FOUND  
**DBD HIT RATIO (%)**

PKG SEARCH NOT FOUND  
PKG SEARCH NOT FOUND INSERT  
PKG SEARCH NOT FOUND DELETE

CT REQUESTS  
CT NOT FOUND  
CT HIT RATIO (%)

STATEMENTS IN GLOBAL CACHE

PT REQUESTS  
PT NOT FOUND  
PT HIT RATIO (%)

## Dynamic SQL Cache

DYNAMIC SQL STMT	QUANTITY	/SECOND
PREPARE REQUESTS	68108.00	12.61
FULL PREPARES	30.00	0.01
SHORT PREPARES	68075.00	12.60
GLOBAL CACHE HIT RATIO (%)	99.96	N/A
IMPLICIT PREPARES	0.00	0.00
PREPARES AVOIDED	0.00	0.00
CACHE LIMIT EXCEEDED	0.00	0.00
PREP STMT PURGED	19.00	0.00
LOCAL CACHE HIT RATIO (%)	N/C	N/A

## Logging Activity – Page 1

LOG ACTIVITY	QUANTITY	/SECOND
READS SATISFIED-OUTPUT BUFF	1815.3K	1009.66
READS SATISFIED-OUTP.BUF(%)	96.91	
READS SATISFIED-ACTIVE LOG	57869.99	32.19
READS SATISFIED-ACTV.LOG(%)	3.09 <i>DISK, NOT TAPE</i>	
READS SATISFIED-ARCHIVE LOG	0.00	0.00
READS SATISFIED-ARCH.LOG(%)	0.00	
TAPE VOLUME CONTENTION WAIT	0.00	0.00
READ DELAYED-UNAVAIL.RESOUR	0.00	0.00
ARCHIVE LOG READ ALLOCATION	0.00	0.00
ARCHIVE LOG WRITE ALLOCAT.	2.00	0.00
CONTR.INTERV.OFFLOADED-ARCH	107.8K	59.99
LOOK-AHEAD MOUNT ATTEMPTED	0.00	0.00
LOOK-AHEAD MOUNT SUCCESSFUL	0.00	0.00

13

GoFurther

Archive to disk: With Archive log data sets on disk, it can be shared by many log readers.

Every rollback that accesses the archive log must wait for any previous rollback work that accesses the same archive tape volume to complete.

## Logging Activity – Page 2

LOG ACTIVITY	QUANTITY	/SECOND
UNAVAILABLE OUTPUT LOG BUFF	0.00	0.00
OUTPUT LOG BUFFER PAGED IN	0.00	0.00
LOG RECORDS CREATED	1875.6K	1043.20
LOG CI CREATED	65552.24	36.46
LOG WRITE I/O REQ (LOG1&2)	125.1K	69.59
LOG CI WRITTEN (LOG1&2)	177.6K	98.76
LOG RATE FOR 1 LOG (MB)	N/A	0.19
LOG WRITE SUSPENDED	23940.49	13.32

200 Bytes - "Average" size of log record

Log CI CREATED x 4 KB / LOG RECORDS CREATED = Calculated Log Record  
 $65552.24 \times 4\text{KB} / 1,875,600 = 143 \text{ bytes}$

A non-zero value for UNAVAILABLE OUTPUT LOG BUFF is an indicator that your output buffer is too small. Make sure Output Log Buffer is allocated to a size that at the size you choose is backed up by real storage. A non-zero value for B in Figure 72 on page 676 is an indicator that your output buffer is too large for the amount of available real storage.

For fixed-length non-compressed rows

changes are logged only from the beginning of the first updated column to the end of the last updated column

For varying-length rows (contains one/or more varying-length columns)

data is logged from the first changed byte to the end of the last updated column

You can have 93 Active Logs

Logging less than 5 MB per second, the workload is read-intensive

Logging more than 5 MB per second, the workload is update-intensive

## Subsystem Activity

SUBSYSTEM SERVICES	QUANTITY	/SECOND
IDENTIFY	7253.00	0.14
CREATE THREAD	20109.00	0.40
SIGNON	1.00	0.00
TERMINATE	27366.00	0.54
ROLLBACK	13407.00	0.27
COMMIT PHASE 1	232.0K	4.62
COMMIT PHASE 2	232.3K	4.62
READ ONLY COMMIT	1692.00	0.03
UNITS OF RECOVERY INDOUBT	0.00	0.00
UNITS OF REC.INDBT RESOLVED	0.00	0.00
SYNCHS(SINGLE PHASE COMMIT)	74175.00	1.48
QUEUED AT CREATE THREAD	0.00	0.00
SUBSYSTEM ALLIED MEMORY EOT	560.00	0.01
SUBSYSTEM ALLIED MEMORY EOM	0.00	0.00
SYSTEM EVENT CHECKPOINT	328.00	0.01
HIGH WATER MARK IDBACK	10.00	0.00
HIGH WATER MARK IDFORE	3.00	0.00
HIGH WATER MARK CTHREAD	15 12.00	0.00

GoFurther

## Latching

LATCH CNT	/SECOND	/SECOND	/SECOND	SECOND
LC01-LC04	0.00	0.00	0.00	0.01
LC05-LC08	0.00	0.23	0.00	0.00
LC09-LC12	0.00	0.24	0.00	1.56
LC13-LC16	0.51	48.78	0.00	0.08
LC17-LC20	0.00	0.00	835.32	0.00
LC21-LC24	0.02	0.00	168.41	176.78
LC25-LC28	3.73	0.00	0.00	0.00
LC29-LC32	0.00	0.20	11.54	3.81

DB2 Latch Management controls DB2's external access to internal DB2 resources. Latching is used for short-term serialization of internal DB2 resources such as storage or control block chains. It is performed by the DB2 agent services manager subcomponent and the Buffer Manager.



## Global DDF Activity – Page 1

GLOBAL DDF ACTIVITY	QUANTITY	/SECOND
DBAT QUEUED-MAXIMUM ACTIVE	0.00	0.00
CONV.DEALLOC-MAX.CONNECTED	0.00	0.00
COLD START CONNECTIONS	0.00	0.00
WARM START CONNECTIONS	1.00	0.00
RESYNCHRONIZATION ATTEMPTED	0.00	0.00
RESYNCHRONIZATION SUCCEEDED	0.00	0.00
CUR TYPE 1 INACTIVE DBATS	0.00	N/A
TYPE 1 INACTIVE DBATS <b>HWM</b>	1.00	N/A
TYPE 1 CONNECTIONS TERMINAT	0.00	0.00

DBAT QUEUED-MAXIMUM ACTIVE shows the number of DBATs that had to wait because the maximum number of threads allowed in the system had been reached. This value should be zero, or close to zero. Increase the value of MAXDBAT to raise the number of current active DBATs allowed

## Global DDF Activity – Page 2

GLOBAL DDF ACTIVITY	QUANTITY	/SECOND
CUR TYPE 2 INACTIVE DBATS	96.16	N/A
TYPE 2 INACTIVE DBATS <b>HWM</b>	628.00	N/A
ACC QUEUED TYPE 2 INACT THR	659.2K	11.71
CUR QUEUED TYPE 2 INACT THR	0.00	N/A
QUEUED TYPE 2 INACT THR <b>HWM</b>	25.00	N/A
CURRENT ACTIVE DBATS	8.83	N/A
ACTIVE DBATS <b>HWM</b>	90.00	N/A
TOTAL DBATS <b>HWM</b>	628.00	N/A
CURRENT DBATS NOT IN USE	5.03	N/A
DBATS NOT IN USE <b>HWM</b>	74.00	N/A
DBATS CREATED	356.00	N/A
POOL DBATS REUSED	1318.1K	N/A

If the inactive option is specified for DDF threads on the installation panel DSNTIPR, it is possible that this field and the currently active DBATs exceed the MAX REMOTE ACTIVE value specified on installation panel DSNTIPE. This is because DDF allows database access threads to be initialized as long as the MAX REMOTE CONNECTED value specified on installation panel DSNTIPE is not exceeded. The resulting database access threads are regarded as active during database access thread initialization and are thus accounted for in both this field and the Active DBATs currently counts. If the current database access thread number is still greater than the MAXIMUM REMOTE ACTIVE value after the completion of the database access thread initialization, the extra database access threads are made inactive. The Active DBATs currently value is then decreased and the value of this field is increased by the number of threads that were made inactive.

## DBM1 Storage Stats – Page 1

DBM1 AND MVS STORAGE <b>BELOW</b> 2 GB		QUANTITY
TOTAL DBM1 STORAGE BELOW 2 GB	(MB)	923.40
TOTAL GETMAINED STORAGE	(MB)	221.31
VIRTUAL BUFFER POOLS	(MB)	N/A
VIRTUAL POOL CONTROL BLOCKS	(MB)	N/A
EDM POOL	(MB)	195.31
COMPRESSION DICTIONARY	(MB)	N/A
CASTOUT BUFFERS	(MB)	N/A
DATA SPACE LOOKASIDE BUFFER	(MB)	N/A
HIPERPOOL CONTROL BLOCKS	(MB)	N/A
DATA SPACE BP CONTROL BLOCKS	(MB)	N/A

**SEE DB2 PM EXPERT RELEASE 2 for more details**

## DBM1 Storage Stats – Page 2

DBM1 AND MVS STORAGE <b>BELOW</b> 2 GB		QUANTITY
TOTAL VARIABLE STORAGE	(MB)	434.09
TOTAL AGENT LOCAL STORAGE	(MB)	384.59
TOTAL AGENT SYSTEM STORAGE	(MB)	284.91
NUMBER OF PREFETCH ENGINES		290.00
NUMBER OF DEFERRED WRITE ENGINES		254.00
NUMBER OF CASTOUT ENGINES		600.00
NUMBER OF GBP WRITE ENGINES		524.00
NUMBER OF P-LOCK/NOTIFY EXIT ENGINES		342.00
TOTAL AGENT NON-SYSTEM STORAGE	(MB)	99.68
TOTAL NUMBER OF ACTIVE USER THREADS		108.00
RDS OP POOL	(MB)	N/A
RID POOL	(MB)	0.97
PIPE MANAGER SUB POOL	(MB)	0.00
LOCAL DYNAMIC STMT CACHE CNTL BLKS	(MB)	0.99

## DBM1 Storage Stats – Page 3

DBM1 AND MVS STORAGE <b>BELOW 2 GB</b>		QUANTITY
THREAD COPIES OF CACHED SQL STMTS	(MB)	0.00
IN USE STORAGE	(MB)	0.00
STATEMENTS COUNT		0.00
HWM FOR ALLOCATED STATEMENTS	(MB)	0.00
STATEMENT COUNT AT HWM		0.00
DATE AT HWM		07/29/06
TIME AT HWM		23:58:43.22
BUFFER & DATA MANAGER TRACE TBL	(MB)	N/A
TOTAL FIXED STORAGE	(MB)	4.31
TOTAL GETMAINED STACK STORAGE	(MB)	263.68
TOTAL STACK STORAGE IN USE	(MB)	222.22u
STORAGE CUSHION	(MB)	108.60

## DBM1 Storage Stats – Page 4

DBM1 AND MVS STORAGE <b>BELOW 2 GB</b> CONTINUED		QUANTITY
24 BIT LOW PRIVATE	(MB)	0.22
24 BIT HIGH PRIVATE	(MB)	0.36
31 BIT EXTENDED LOW PRIVATE	(MB)	56.55
31 BIT EXTENDED HIGH PRIVATE	(MB)	982.51
EXTENDED REGION SIZE (MAX)	(MB)	1679.00
EXTENDED CSA SIZE	(MB)	200.23
AVERAGE THREAD FOOTPRINT	(MB)	<b>1.38</b>
MAX NUMBER OF POSSIBLE THREADS		675.56

## DBM1 Storage Stats – Page 5

DBM1 STORAGE ABOVE 2 GB		QUANTITY
FIXED STORAGE	(MB)	31.10
GETMAINED STORAGE	(MB)	2296.37
IN USE EDM DBD POOL	(MB)	274.19
IN USE EDM STATEMENT POOL	(MB)	4.88
IN USE EDM RDS POOL	(MB)	126.95
IN USE EDM SKELETON POOL	(MB)	36.95
VIRTUAL BUFFER POOLS	(MB)	3394.53
VIRTUAL POOL CONTROL BLOCKS	(MB)	0.14
CASTOUT BUFFERS	(MB)	75.00
VARIABLE STORAGE	(MB)	1860.02
THREAD COPIES OF CACHED SQL STMTS	(MB)	33.44
IN USE STORAGE	(MB)	12.44
HWM FOR ALLOCATED STATEMENTS	(MB)	14.33
REAL AND AUXILIARY STORAGE		QUANTITY
REAL STORAGE IN USE	(MB)	4775.47
AUXILIARY STORAGE IN USE	(MB)	<b>4393.71</b>

23

Go further

Also captured are additional storage stats

SHARED MEMORY STORAGE	(MB)	9.06
TOTAL FIXED VIRTUAL 64BIT SHARED	(MB)	1.14
TOTAL GETMAINED VIRTUAL 64BIT SHARED	(MB)	3.11
TOTAL VARIABLE VIRTUAL 64BIT SHARED	(MB)	4.81

## BUFFERPOOL – Page 1

BP19 GENERAL	QUANTITY	/SECOND
CURRENT ACTIVE BUFFERS	4117.69	N/A
UNAVAIL.BUFFER-VPOOL FULL	0.00	0.00
NUMBER OF DATASET OPENS	716.00	0.01
BUFFERS ALLOCATED - VPOOL	743.0K	N/A
DFHSM MIGRATED DATASET	0.00	0.00
DFHSM RECALL TIMEOUTS	0.00	0.00
VPOOL EXPANS. OR CONTRACT.	0.00	0.00
VPOOL OR HPOOL EXP.FAILURE	0.00	0.00
CONCUR.PREF.I/O STREAMS-HWM	0.00	N/A
PREF.I/O STREAMS REDUCTION	0.00	0.00
PARALLEL QUERY REQUESTS	0.00	0.00
PARALL.QUERY REQ.REDUCTION	0.00	0.00
PREF.QUANT.REDUCED TO 1/2	0.00	0.00
PREF.QUANT.REDUCED TO 1/4	0.00	0.00



## BUFFERPOOL – Page 2

BP19 READ OPERATIONS	QUANTITY	/SECOND
-----		
BPOOL HIT RATIO (%)	86.62	
GETPAGE REQUEST	5897.5M	68.9K
GETPAGE REQUEST-SEQUENTIAL	875.2M	10.2K
GETPAGE REQUEST-RANDOM	5022.4M	58.7K
SYNCHRONOUS READS	265.8M	3106.24
SYNCHRON. READS-SEQUENTIAL	5861.8K	68.50
SYNCHRON. READS-RANDOM	259.9M	3037.73
GETPAGE PER SYN.READ-RANDOM	19.32	

12 hours of processing

## BUFFERPOOL – Page 3

BP19 READ OPERATIONS CONTINUED)	QUANTITY	/SECOND
SEQUENTIAL PREFETCH REQUEST	75716.7K	884.86
SEQUENTIAL PREFETCH READS	3588.2K	41.93
PAGES READ VIA SEQ.PREFETCH	88130.3K	1029.93
SEQ.PRF.PAGES READ/S.PRF.READ	24.56	
LIST PREFETCH REQUESTS	10792.1K	126.12
LIST PREFETCH READS	2515.2K	29.39
PAGES READ VIA LIST PREFETCH	52470.1K	613.19
L.PRF.PAGES READ/L.PRF.READ	20.86	
DYNAMIC PREFETCH REQUESTED	61668.3K	720.68
DYNAMIC PREFETCH READS	15127.4K	176.78
PAGES READ VIA DYN.PREFETCH	382.7M	4472.40
DRF.PAGES READ/D.PRF.READ	25.30	
PREF.DISABLED-NO BUFFER	0.00	0.00
PREF.DISABLED-NO READ ENG	0.00	0.00
PAGE-INS REQUIRED FOR READ	8.00	0.00

## BUFFERPOOL – Page 4

BP19 WRITE OPERATIONS	QUANTITY	/SECOND
-----	-----	-----
BUFFER UPDATES	559.6M	6540.25
PAGES WRITTEN	119.5M	1396.01
BUFF.UPDATES/PAGES WRITTEN	4.68	
SYNCHRONOUS WRITES	5727.00	0.07
ASYNCHRONOUS WRITES	5036.1K	58.85
PAGES WRITTEN PER WRITE I/O	23.69	
HORIZ.DEF.WRITE THRESHOLD	1717.00	0.02
VERTI.DEF.WRITE THRESHOLD	1419.9K	16.59
DM THRESHOLD	0.00	0.00
WRITE ENGINE NOT AVAILABLE	0.00	0.00
PAGE-INS REQUIRED FOR WRITE	0.00	0.00

## BUFFERPOOL – Page 5

BP7 SORT/MERGE	QUANTITY	/SECOND
MAX WORKFILES CONCURR. USED	286.12	N/A
MERGE PASSES REQUESTED	2016.3K	23.56
MERGE PASS DEGRADED-LOW BUF	0.00	0.00
WORKFILE REQ.REJCTD-LOW BUF	0.00	0.00
WORKFILE REQ-ALL MERGE PASS	4036.4K	47.17
WORKFILE NOT CREATED-NO BUF	0.00	0.00
WORKFILE PRF NOT SCHEDULED	0.00	0.00

## RID Processing

RID LIST PROCESSING	QUANTITY	/SECOND
MAX RID BLOCKS ALLOCATED	739.00	N/A
CURRENT RID BLOCKS ALLOCAT.	208.00	N/A
TERMINATED-NO STORAGE	0.00	0.00
TERMINATED-EXCEED RDS LIMIT	327.00	0.36
TERMINATED-EXCEED DM LIMIT	0.00	0.00
TERMINATED-EXCEED PROC.LIM.	0.00	0.00

**IFCID 125 from Performance Class 8 provides the REAL detail of RID pool process**

**(TERMINATED-NO STORAGE)** is not zero, reduce overall storage use.

RDS or Data Manager limits are exceeded

**(TERMINATED-EXCEED RDS LIMIT and TERMINATED-EXCEED DM LIMIT)**

Inaccurate or incomplete RUNSTATS statistics

Optimizer error

Consider the OPTIMIZE FOR 1 ROW clause to avoid RID list processing

, or add the necessary columns to the index to make access INDEX ONLY.

I **(TERMINATED-EXCEED PROC LIMIT)** is not zero, or close to zero, increase RID pool size.

## LOCKING – PAGE 1

LOCKING ACTIVITY	QUANTITY	/SECOND
SUSPENSIONS (ALL)	246.0K	4.89
SUSPENSIONS (LOCK ONLY)	89050	1.77
SUSPENSIONS (IRLM LATCH)	121.4K	2.41
SUSPENSIONS (OTHER)	96110.13	28.64
TIMEOUTS	0.00	0.00
DEADLOCKS	0.00	0.00

## LOCKING – PAGE 2

LOCKING ACTIVITY	QUANTITY	/SECOND
LOCK REQUESTS	214.6M	4270.50
UNLOCK REQUESTS	38131.7K	758.73
QUERY REQUESTS	0.00	0.00
CHANGE REQUESTS	1266.5K	25.20
OTHER REQUESTS	0.00	0.00
LOCK ESCALATION (SHARED)	0.00	0.00
LOCK ESCALATION (EXCLUSIVE)	0.00	0.00
DRAIN REQUESTS	11499	0.23
DRAIN REQUESTS FAILED	3286.00	0.07
CLAIM REQUESTS	12713.7K	252.97
CLAIM REQUESTS FAILED	0.00	0.00

## What's New.....

- Stats for new features coming with DB2 9 will deliver new information
  - XML Locking
  - TRUNCATE
  - NATIVE SQL PROCEDURES
  - MERGE
  - SKIPPED LOCK DATA
  - IP address - now may require 39 character colon hexadecimal representation of an IPv6 address
  - DEBUGSESSION
  - Security – Trusted Connection
    - context name, role name, original application user, security token, etc....



## More What's New.....

Stats for new features coming with DB2 9 will deliver new information

- The IFCID 106 mapping macro is updated:
  - For the new TBSBP8K, TBSBP16K, TBPS32K, IMPDB, IMPTSSEG, IMPDSDEF, and IMPTSCMP system parameters
  - SYSTEM\_LEVEL\_BACKUPS, RESTORE\_RECOVER\_FROMDUMP, UTILS\_DUMP\_CLASS\_NAME and RESTORE\_TAPEUNITS
- FETCH CURRENT CONTINUE for LOB/ XML
- CLONE TABLE SUPPORT
- CHANGE TO BUILD2 – UTILITY CHANGE
- RENAME INDEX

## And a few more....

- Stats for new features coming with DB2 9 will deliver new information
  - New counters for DM Stats to track storage use of WORKFILE database
  - PLAN TABLE CHANGES
  - ALTER BUFFERPOOL

## Things to Remember....

- Consider benefit of the view of Long Report for the system
  - Create a baseline report for the system
- Understand Statistics pattern for your environment
  - What is Normal (or not !)
- Be aware of:
  - what Statistics classes are normally started by subsystem
  - How long is SMF data stored
  - Exploitation of storing DB2 performance information in DB2 tables

DB2 for z/OS System Performance Basics

A DB2 Statistics Review

Session code: A06

# Questions

Randy Ebersole

IBM

[ebersole@us.ibm.com](mailto:ebersole@us.ibm.com)

Session code: A06

GoFurther