

**SHARE**

Technology • Connections • Results

# ATS (Advanced Technical Skills) DB2 Health Check - What Will the Specialist Prescribe for You?

John Iczkovits [iczkovit@us.ibm.com](mailto:iczkovit@us.ibm.com)  
IBM

March 4, 2011  
Session #8387




# Disclaimer

- © Copyright IBM Corporation 2011. All rights reserved. U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
- *THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.*
- **IBM, the IBM logo, ibm.com, DB2 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)**

## Abstract

- Maybe you are concerned about your DB2 environment or maybe just want to validate your environment, so you are thinking about an ATS (Advanced Technical Skills) DB2 Health Check. You are thinking, what really happens in a DB2 Health Check? What will I walk away with? Will my management approve of the final analysis? How will my company and I benefit from the analysis? If you've seen the required documentation list, after reviving yourself you ask, is this all really necessary? Why does ATS really need this detailed information and what will they do with it? Find out the answers to these questions and more by attending this presentation. You will leave understanding why the ATS DB2 Health Check is truly one of the best in the industry and how it will benefit you and your company.




**SHARE**  
Technology • Connections • Results

# Agenda

- Some basics
- OMPE reports
- RMF reports
- Common ZPARM issues
- Remaining reports
- Final analysis


# NOTICE



**SHARE**  
Efficiency • Connections • Results

- Note – many customer reports look similar. The reports and information in this presentation were changed from their original format. If the information you are reviewing seems familiar, it is purely coincidental and is not from a specific customer.
- DB2 Health Checks can encompass many different variations of requests. This presentation illustrates some common areas we concentrate on, however Health Checks are not limited to these areas.
- Although we cover some areas outside of DB2 at a minimal level, this does not replace the need for such things as an z/OS Health Check. Keep in mind, our focus and expertise is in DB2. Detailed questions regarding z/OS or other products should be handled by those products experts.

**SHARE**  
in Anaheim  
2011




**SHARE**  
Technology • Connections • Results

# Some Basics

## Why ask for an ATS DB2 Health Check?

- Generally an ATS DB2 Health Check provides expert advice on your overall environment. Some reasons customers use our service:
  - Many customers do not work in an environment where the total view of performance can be analyzed. If there is a problem, what is causing it?
  - Availability is another key concern. Are there any single points of failure or major issues?
  - Skills development. It is not unusual for an ATS DB2 Health Check report to have 50 or more pages. Each finding is outlined with detailed information regarding its significance, cause, and solutions. Customers who wish to learn more about our analysis can work alongside the ATS specialist to gain a better understanding of their environment. Detailed analysis of top 10 problem SQL statements will provide the customer with a framework to understand how to address different SQL issues.
  - An independent report provides customers with trusted expert and detailed analysis without political pressure from different and at times opposing groups.
- 7 • The final result is an extensive and detailed report with a follow up meeting to discuss the findings.



**SHARE**  
Technology • Connections • Results

# What data do we ask for?

- SMF data - 30,70-79,100,101,102. One hour for online peak and another hour for batch peak.
- Copy of ZPARMs in the format of DSNTIJUZ– if this is a Data Sharing Health Check and member ZPARMs are the same, send one member's ZPARM only.
- Copy of DSNHDECP
- Output from DSNJU004 as well as LISTCAT output for all active log data sets and BSDS. For Data Sharing Health Checks send the output for each member.
- SELECT \* output for the RTS (Real Time Statistics) tables:
  - Prior to V9: TABLESPACESTATS and INDEXSPACESTATS
  - V9: SYSTABLESPACESTATS and SYSINDEXSPACESTATS
- BIND options for major applications
- Output from command DISPLAY GROUP
- Output from command DISPLAY BUFFERPOOL(ACTIVE) DETAIL(\*). For Data Sharing Health Checks, run this command for each member.
- For Data Sharing Health Checks, provide the output from one member only for command: DIS GBPOOL(\*) GDETAIL(\*) TYPE(GCONN)
- Explanation of buffer pool designations
- DIAGNOSE DISPLAY MEPL output
- Output of DB2 started tasks, such as MSTR, DBM1, IRLM, DIST, ADMT, and WLM stored procedure address spaces for DB2.
- OSC/OE or VE output for the top 10 problem SQL statements
- Copy of the CFRM policy
- Copy of the WLM policy
- Disk layout diagram
- Data Sharing member/subsystem LPAR layout diagram
- <sup>8</sup> We review backup and recovery issues when requested

**SHARE**  
in Anaheim  
2011



## FAQ- WOW! You have asked for a huge amount of data. Do you really look at all of the data provided?



- YES! We actually do. When performance and/or availability is an issue for customers we need very detailed information.
- Running EXPLAINS will not help you diagnose such problems as incorrect WLM assignments or disk problems. Is it the LPAR that is saturated or the CEC? In depth analysis must be performed in order to understand why performance goals are not being met.
- Running EXPLAINS will not tell you when you have availability issues.
  - Are there single points of failure?
  - Are there more optimal ways of structuring your environment?

## What are the SMF types used for and why ask for all of them?

- We ask for specific SMF record types so that we do not need to go back and ask the customer for even more data if further analysis is required.
- SMF record types (subtypes have additional information not outlined below):
  - 30 - Common Address Space Work
  - 70 - RMF Processor Activity
  - 71 - RMF Paging Activity
  - 72 - RMF Workload Activity and Storage Data
  - 73 - RMF Channel Path Activity
  - 74 - RMF Activity of several resources, such as coupling facility
  - 75 - RMF Page Data Set Activity
  - 76 - RMF Trace Activity
  - 77 - RMF Enqueue Activity
  - 78 - RMF Virtual Storage and I/O Queuing Activity
  - 79 - RMF Monitor II Activity
  - 100 - DB2 Statistics
  - 101 - DB2 Accounting
  - 102 - DB2 Statistics and Performance

# Output from DISPLAY GROUP

DSN7100I -DBP1 DSN7GCMD

\*\*\* BEGIN DISPLAY OF GROUP(AAADBP0 ) GROUP LEVEL(910) MODE(N )  
PROTOCOL LEVEL(3) GROUP ATTACH NAME(DBP0)

```
-----  
DB2                                DB2 SYSTEM      IRLM  
MEMBER  ID  SUBSYS  CMDPREF  STATUS  LVL  NAME      SUBSYS  IRLMPROC  
-----  
DBP1      1  DBP1   -DBP1   ACTIVE  910  AAAA      IRP1   DBP1IRLM  
DBP2      2  DBP2   -DBP2   ACTIVE  910  AAAB      IRP2   DBP2IRLM  
DBP3      3  DBP3   -DBP3   ACTIVE  910  AAAC      IRP3   DBP3IRLM  
DBP4      4  DBP4   -DBP4   ACTIVE  910  AAAD      IRP4   DBP4IRLM  
-----
```

SCA STRUCTURE SIZE: 8192 KB, STATUS= AC, SCA IN USE: 2 %

LOCK1 STRUCTURE SIZE: 32256 KB

NUMBER LOCK ENTRIES: 8388608


NUMBER LIST ENTRIES: 53801, LIST ENTRIES IN USE: 20

\*\*\* END DISPLAY OF GROUP(AAADBP0 )

DSN9022I -DBP1 DSN7GCMD 'DISPLAY GROUP ' NORMAL COMPLETION



- We start out with the DISPLAY GROUP output. This gives us a good overview of the customer's environment.
- From this report we find:
  - DB2 members/ subsystem names
  - LPARs they run on
  - Version of DB2
  - Mode of DB2
  - IRLM names
  - DB2 command prefix
  - SCA and lock sizes.



**SHARE**  
Technology • Connections • Results

# OMPE Reports


# What OMPE reports are produced from the SMF data?

- OMPE (Omegamon XE Performance Expert) reports:
  - Accounting Short reports
  - Accounting Long reports
  - Accounting reports by connection types
  - Statistics reports
  - Statistics traces
  - More detailed OMPE reports when required

# What RMF reports are produced from the SMF data?

- RMF reports:
  - Summary
  - CPU Activity
  - Partition Data
  - Paging Activity
  - Workload Activity (SCPER)
  - Coupling Facility Activity
  - Additional RMF reports when required

# GMT vs. Local FAQ



- SMF records for RMF are recorded in local time, however DB2 SMF type 100-102 records are in GMT. Why the difference?
- Answer: It has to do with DB2 logging and recovery:
  - GMT is always constant and never changes an hour forward or back.
  - DB2 requires that there is no overlap in log information that could cause corruption problems because the same timestamp is used after switching between daylight savings time and standard time.
  - RMF reports do not need to deal with logging issues, there are no data corruption issues and therefore the requirements are relaxed.
- For RMF reports we do not use GMT offsets as they are local time, however for DB2 data we need to add the GMT offset. We need to be very careful when dealing with offsets. The customer may be from the west coast, however the CPUs are running with the time in central where they are located. We must always verify this information before proceeding.
- Physical SMF type 100-102 as well as GTF records are created as local time plus the offset. Reporting products require the GMT offset.

E  
results

SHARE  
in Anaheim  
2011

```

IFA010I SMF DUMP PARAMETERS
IFA010I END(2400) -- DEFAULT
IFA010I START(0000) -- DEFAULT
IFA010I DATE(1900000,2099366) -- DEFAULT
IFA010I OUTDD(DUMPOUT,TYPE(100:102)) -- SYSIN
IFA010I INDD(DUMPIN,OPTIONS(ALL)) -- SYSIN
IFA020I DUMPOUT -- NULLFILE
IFA013I 'CLEAR' OPTION IS VALID ONLY FOR VSAM DATA SETS. OPTION
        IGNORED FOR DATASET DUMPIN
IFA020I DUMPIN -- JOHNICZ.SMF.UNTERSED

```

is the SMF data for the right period?



SUMMARY ACTIVITY REPORT

START DATE-TIME 03/21/2008-02:00:00 END DATE-TIME 03/21/2008-15:59:59


RECORD TYPE	RECORDS READ	PERCENT OF TOTAL	AVG. RECORD LENGTH	MIN. RECORD LENGTH	MAX. RECORD LENGTH	RECORDS WRITTEN
2	1	.00 %	18.00	18	18	1
3	1	.00 %	18.00	18	18	1
30	39,790	9.80 %	1,363.03	398	32,738	0
70	72	.02 %	2,140.50	356	4,040	0
71	36	.01 %	1,714.00	1,676	1,752	0
72	12,240	3.01 %	1,142.02	1,068	14,604	0
73	36	.01 %	19,704.00	19,704	19,704	0
74	1,368	.34 %	28,204.78	364	32,720	0
75	258	.06 %	264.00	264	264	0
77	36	.01 %	9,688.88	2,560	18,880	0
78	72	.02 %	8,803.94	1,888	18,840	0
100	96	.02 %	2,559.66	556	7,774	96
101	342,512	84.36 %	2,039.56	1,210	2,876	342,512
102	9,509	2.34 %	408.07	200	2,586	9,509
TOTAL	406,027	100 %	1,998.57	18	32,738	352,119

do we have all of the SMF types requested?

- NUMBER OF RECORDS IN ERROR 0

output from dump SMF program IFASMFDP





# First things first – did we get the right data?

are we looking at the right environment?

LOCATION: DB2  
GROUP: DB2GRP  
MEMBER: DB2  
SUBSYSTEM: DB2  
DB2 VERSION: V8

OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V4)

STATISTICS REPORT - LONG

SCOPE: MEMBER

PAGE: 1-1  
REQUESTED FROM: 12/19/08 00:00:00.00  
TO: 12/19/08 23:59:59.99  
INTERVAL FROM: 12/19/08 02:13:22.77  
TO: 12/19/08 15:43:33.84

----- HIGHLIGHTS -----

INTERVAL START	: 12/19/08 02:13:22.77	SAMPLING START:	12/19/08 02:13:22.77	TOTAL THREADS	:	692.2K
INTERVAL END	: 12/19/08 15:43:33.84	SAMPLING END	: 12/19/08 15:43:33.84	TOTAL COMMITS	:	1058.7K
INTERVAL ELAPSED:	13:30:11.077457	OUTAGE ELAPSED:	0.000000	DATA SHARING MEMBER:	:	N/A

we did not specify a time in our input, are the dates and time for the period requested correct?  
We really need to watch those time zones!

is the interval long enough?

OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V4)


PAGE: 1

JOB SUMMARY LOG

RUN DATE: 02/24/09 09:51:21.51

MSG.ID.	LOCATION	GROUP	SSID	MEMBER	TIMESTAMP	DESCRIPTION
0FPEC4065I	DB2P	AAADB2P	DBP1	DBP1	02/19/09 16:00:00.40	DB2 STOP TRACE NUMBER 03 DB2 SUBSYSTEM ID = DBP1 TEXT = -STOP TRACE (ACCTG )CLASS (* )RMID (* )PLAN (* )AUTHID (* )TNO (* )
0FPEC4005I	DB2P	AAADB2P	DBP1	DBP1		NUMBER OF RECORDS PROCESSED WITHOUT A CPU HEADER WAS 658
0FPEC4010I	DB2P	AAADB2P	DBP1	DBP1		NUMBER OF RECORDS PROCESSED WITHOUT A CORRELATION HEADER WAS 579

Any traces change during the interval of interest?  
Are we missing any records?



# Are the right IFCIDs produced?

LOCATION: DB2

OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V4)

PAGE: 4

GROUP: N/P

IFCID FREQUENCY DISTRIBUTION LOG

RUN DATE: 02/04/09 10:22:02.75

MEMBER: N/P

SUBSYSTEM: DB2


ACTUAL FROM: 12/19/08 02:00:00.18

DB2 VERSION: V8

TO: 12/19/08 15:59:59.72

IFCID	INPUT COUNT	INPUT PCT OF TOTAL	PROCESSED COUNT	PROCESSED PCT OF TOTAL	IFCID	INPUT COUNT	INPUT PCT OF TOTAL	PROCESSED COUNT	PROCESSED PCT OF TOTAL
1	4	0.00%	0	0.00%	202	4	0.00%	0	0.00%
2	4	0.00%	0	0.00%	225	4	0.00%	0	0.00%
3	40,476	99.76%	0	0.00%	230	4	0.00%	0	0.00%
105	28	0.06%	0	0.00%	254	40	0.09%	0	0.00%
106	4	0.00%	0	0.00%	337	3	0.00%	0	0.00%
TOTAL INPUT TRACE RECORDS =			40,571						
TOTAL PROCESSED TRACE RECORDS =			0						

# What are we looking for in the OMPE Statistics Reports?



**SHARE**  
S H A R E  
S H A R E  
S H A R E

- Lots of different things, but the following pages illustrate some of the key details we look for.
- Keep in mind, some values are accumulated, some are the values at the time the statistics record is cut.
- We must be very careful regarding values in relation to the interval period. For some items a wide interval period is very helpful for a birds eye view of the environment, however for other items a wide interval period is worthless.
- We must also be very careful when looking at High Water Marks (HWM). When was the HWM set? Was it a second ago? A week ago? Several months ago? We cannot tell just by looking at the report.
  - HWM FAQ – What are some circumstances where the HWM is irrelevant because it was not set recently? The environment could have changed. For example, did you dynamically add or decrease CPUs, memory, disk, etc.? Or maybe dynamically alter some ZPARM values? We can even move disk or data sets without taking a DB2 outage. Investigating HWM issues takes time and effort.

**SHARE**  
in Anaheim  
2011


## Statistics FAQ

Why at times do you run the statistics with REPORT and at other times TRACE?

ANSWER: It depends on the situation. Some key questions are:

- What is the SMF interval?
- Are we looking for specific problems? If so, are the related fields accumulated, or are they a snapshot view when the record was cut? If the latter, a field probably does not contain information from the entire interval.

REPORT can work very well for the big picture, however we may need to run a TRACE for specific time periods.




**SHARE**

Technology • Connections • Results

SQL DDL	QUANTITY	/SECOND	/THREAD	/COMMIT	SQL DDL	CONTINUED	QUANTITY	/SECOND	/THREAD	/COMMIT
CREATE TABLE	1069.00	0.28	0.00	0.00	DROP TABLE		928.00	0.27	0.00	0.00
CREATE GLOBAL TEMP TABLE	0.00	0.00	0.00	0.00	DROP INDEX		82.00	0.00	0.00	0.00
DECLARE GLOBAL TEMP TABLE	74.00	0.00	0.00	0.00	DROP VIEW		0.00	0.00	0.00	0.00
CREATE AUXILIARY TABLE	0.00	0.00	0.00	0.00	DROP SYNONYM		0.00	0.00	0.00	0.00
CREATE INDEX	249.00	0.01	0.00	0.00	DROP TABLESPACE		92.00	0.00	0.00	0.00
. . . .										

Are there a large number of DDL executions for a short interval?

For example, a large number of CREATEs and DROPs in a short interval deserves further investigation.



**SHARE**  
Technology • Connections • Results

EDM POOL	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
PAGES IN RDS POOL (BELOW)	30740.00	N/A	N/A	N/A
...				
<b>FAILS DUE TO POOL FULL</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>
PAGES IN RDS POOL (ABOVE)	524.3K	N/A	N/A	N/A
...				
<b>FAILS DUE TO RDS POOL FULL</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>
PAGES IN DBD POOL (ABOVE)	75000.00	N/A	N/A	N/A
...				
<b>FAILS DUE TO DBD POOL FULL</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>
PAGES IN STMT POOL (ABOVE)	25000.00	N/A	N/A	N/A
...				
<b>FAILS DUE TO STMT POOL FULL</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>
PAGES IN SKEL POOL (ABOVE)	25600.00	N/A	N/A	N/A
...				
<b>FAILS DUE TO SKEL POOL FULL</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>
DBD REQUESTS	30822.83	45.91	1426.03	3.39
DBD NOT FOUND	0.00	0.00	0.00	0.00
<b>DBD HIT RATIO (%)</b>	<b>100.00</b>	<b>N/A</b>	<b>N/A</b>	<b>N/A</b>
CT REQUESTS	21.61	0.03	1.00	0.00
CT NOT FOUND	0.00	0.00	0.00	0.00
<b>CT HIT RATIO (%)</b>	<b>100.00</b>	<b>N/A</b>	<b>N/A</b>	<b>N/A</b>
PT REQUESTS	58196.59	86.69	2692.48	6.41
PT NOT FOUND	0.00	0.00	0.00	0.00
<b>PT HIT RATIO (%)</b>	<b>100.00</b>	<b>N/A</b>	<b>N/A</b>	<b>N/A</b>

any pool full conditions?

any hit ratio problems?

**SHARE**  
in Anaheim  
2011

DYNAMIC SQL STMT	QUANTITY	/SECOND	/THREAD	/COMMIT	SUBSYSTEM SERVICES	QUANTITY	/SECOND	/THREAD	/COMMIT
PREPARE REQUESTS	24509.41	36.36	1129.31	2.69	IDENTIFY	24.85	0.04	1.15	0.00
FULL PREPARES	66.19	0.10	3.06	0.01	CREATE THREAD	21.61	0.03	1.00	0.00
SHORT PREPARES	24332.51	36.25	1125.75	2.68	SIGNON	14.90	0.02	0.69	0.00
GLOBAL CACHE HIT RATIO (%)	99.73	N/A	N/A	N/A	TERMINATE	73.13	0.11	3.38	0.01
					ROLLBACK	7.19	0.01	0.33	0.00
IMPLICIT PREPARES	0.00	0.00	0.00	0.00					
PREPARES AVOIDED	0.00	0.00	0.00	0.00	COMMIT PHASE 1	4195.78	6.25	194.12	0.46
CACHE LIMIT EXCEEDED	0.00	0.00	0.00	0.00	COMMIT PHASE 2	8430.60	12.56	390.04	0.93
PREP STMT PURGED	0.00	0.00	0.00	0.00	READ ONLY COMMIT	127.50	0.19	5.90	0.01
LOCAL CACHE HIT RATIO (%)	N/C	N/A	N/A	N/A	UNITS OF RECOVERY INDOUBT	0.00	0.00	0.00	0.00
					UNITS OF REC.INDBT RESOLVED	0.00	0.00	0.00	0.00
					SYNCHS(SINGLE PHASE COMMIT)	368.28	0.55	17.04	0.04
					QUEUED AT CREATE THREAD	0.00	0.00	0.00	0.00
					SUBSYSTEM ALLIED MEMORY EOT	0.00	0.00	0.00	0.00
					SUBSYSTEM ALLIED MEMORY EOM	0.00	0.00	0.00	0.00
					SYSTEM EVENT CHECKPOINT	2.24	0.00	0.10	0.00
					HIGH WATER MARK IDBACK	21.00	0.03	0.97	0.00
					HIGH WATER MARK IDFORE	1.00	0.00	0.05	0.00
					HIGH WATER MARK CTHREAD	20.00	0.03	0.93	0.00



any hit ratio problems?


for the interval, were there too many:

Rollback?

Queued at create thread (MAXDBAT)?

Memory EOT and EOM?

Number of checkpoints?



how high did the water marks get?






OPEN/CLOSE ACTIVITY	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
OPEN DATASETS - HWM	2404.00	N/A	N/A	N/A
OPEN DATASETS	2255.00	N/A	N/A	N/A
DS NOT IN USE,NOT CLOSE-HWM	2178.00	N/A	N/A	N/A
DS NOT IN USE,NOT CLOSED	1882.00	N/A	N/A	N/A
IN USE DATA SETS	373.00	N/A	N/A	N/A
DSETS CLOSED-THRESH.REACHED	0.00	0.00	0.00	0.00
DSETS CONVERTED R/W -> R/O	353.00	0.20	0.01	0.00

What are the ZPARM values for DSMAX, PCLOSEN and PCLOSET?

Are the ZPARM values set correctly based on the high water marks?

Were data sets closed due to threshold reached?

How many data sets converted from R/W to R/O?



LOG ACTIVITY	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
READS SATISFIED-OUTPUT BUFF	195.00	0.11	0.00	0.00
READS SATISFIED-OUTP.BUF(%)	100.00			
READS SATISFIED-ACTIVE LOG	0.00	0.00	0.00	0.00
READS SATISFIED-ACTV.LOG(%)	0.00			
READS SATISFIED-ARCHIVE LOG	0.00	0.00	0.00	0.00
READS SATISFIED-ARCH.LOG(%)	0.00			
TAPE VOLUME CONTENTION WAIT	0.00	0.00	0.00	0.00
READ DELAYED-UNAVAIL.RESOUR	0.00	0.00	0.00	0.00
ARCHIVE LOG READ ALLOCATION	0.00	0.00	0.00	0.00
ARCHIVE LOG WRITE ALLOCAT.	4.00	0.00	0.00	0.00
CONTR.INTERV.OFFLOADED-ARCH	198.0K	109.97	2.87	2.07
LOOK-AHEAD MOUNT ATTEMPTED	0.00	0.00	0.00	0.00
LOOK-AHEAD MOUNT SUCCESSFUL	0.00	0.00	0.00	0.00
UNAVAILABLE OUTPUT LOG BUFF	0.00	0.00	0.00	0.00
OUTPUT LOG BUFFER PAGED IN	0.00	0.00	0.00	0.00

If there are many reads from the archive log – why? Are the active logs too small and switching too often? Are there COMMIT or rollback issues?

Are the log output buffers too small?

Were there enough tape drives for the archive logs?


DB2 COMMANDS	QUANTITY	/SECOND	DB2 COMMANDS	CONTINUED	QUANTITY	/SECOND
-----	-----	-----	-----	-----	-----	-----
DISPLAY DATABASE	0.00	0.00	MODIFY TRACE		0.00	0.00
DISPLAY THREAD	0.00	0.00	CANCEL THREAD		0.00	0.00
DISPLAY UTILITY	0.00	0.00	TERM UTILITY		0.00	0.00
DISPLAY TRACE	0.00	0.00				
DISPLAY RLIMIT	0.00	0.00	RECOVER BSDS		0.00	0.00
. . .						

Are any of the DB2 command values unusually high for the interval?  
 If so, investigate the cause further. Is there a problem with automation?  
 Was there a specific problem during this time frame?

RID LIST PROCESSING	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
MAX RID BLOCKS ALLOCATED	5499.00	N/A	N/A	N/A
CURRENT RID BLOCKS ALLOCAT.	2.00	N/A	N/A	N/A
TERMINATED-NO STORAGE	0.00	0.00	0.00	0.00
TERMINATED-EXCEED RDS LIMIT	0.00	0.00	0.00	0.00
TERMINATED-EXCEED DM LIMIT	0.00	0.00	0.00	0.00
TERMINATED-EXCEED PROC.LIM.	0.00	0.00	0.00	0.00

Any RID related problems? If so, what is the cause?

# Notes regarding RID related problems



What is the ZPARM value for MAXRBLK? Review the termination reasons:

NO STORAGE – DBM1 storage exhausted. Review memory section.

EXCEED RDS LIMIT – RIDs exceeded 25% of the table size. Most common reason is stale RUNSTATS. Other reason is an optimizer error. In this instance, you could disable RID list processing by adding the clause OPTIMIZE FOR 1 ROW to the SQL statement, or force the access path to index only by adding the necessary columns to the index. Another is multi-index access based on complex OR statement in SQL.

EXCEED DM LIMIT - number of RID entries was greater than the physical limit of approximately 16 million RIDs. Follow the same instructions as in EXCEED RDS LIMIT.

EXCEED PROC. LIM. - The number of times the maximum RID pool storage was exceeded. Increase MAXRBLK if possible.

These are much less of an issue in DB2 10 where processes will continue using a workfile instead of switching throwing away the index reads and then doing a table space scan!

# Authorization Management



AUTHORIZATION MANAGEMENT	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
TOTAL AUTH ATTEMPTS	18859.98	28.09	872.42	2.08
. . .				
PKG CACHE OVERWRT - AUTH ID	0.00	0.00	0.00	0.00
PKG CACHE OVERWRT - ENTRY	0.00	0.00	0.00	0.00
. . .				
RTN CACHE OVERWRT - AUTH ID	0.00	0.00	0.00	0.00
RTN CACHE OVERWRT - ENTRY	0.00	0.00	0.00	0.00
RTN CACHE - ENTRY NOT ADDED	0.00	0.00	0.00	0.00

Review ZPARM value for CACHEPAC. Increase if necessary and memory is available.

Review ZPARM value for CACHERAC. Increase if necessary and memory is available.

LOCKING ACTIVITY	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
SUSPENSIONS (ALL)	2572.81	3.83	119.03	0.28
SUSPENSIONS (LOCK ONLY)	1891.15	2.82	87.49	0.21
SUSPENSIONS (IRLM LATCH)	596.57	0.89	27.60	0.07
SUSPENSIONS (OTHER)	85.10	0.13	3.94	0.01
TIMEOUTS	0.00	0.00	0.00	0.00
DEADLOCKS	0.00	0.00	0.00	0.00
LOCK REQUESTS	268.8K	398.85	12.4K	29.48
UNLOCK REQUESTS	98374.70	146.54	4551.33	10.83
QUERY REQUESTS	0.00	0.00	0.00	0.00
CHANGE REQUESTS	13155.82	19.60	608.66	1.45
OTHER REQUESTS	0.00	0.00	0.00	0.00
LOCK ESCALATION (SHARED)	0.00	0.00	0.00	0.00
LOCK ESCALATION (EXCLUSIVE)	0.00	0.00	0.00	0.00
DRAIN REQUESTS	13.66	0.02	0.63	0.00
DRAIN REQUESTS FAILED	0.00	0.00	0.00	0.00
CLAIM REQUESTS	102.3K	152.36	4732.13	11.26
CLAIM REQUESTS FAILED	0.00	0.00	0.00	0.00

→ Investigate high number of  
→ deadlocks and timeouts

→ Are there many more  
→ locks than unlocks? Is  
lock avoidance occurring?

→ Are there many lock  
→ escalations?

→ Investigate high number of  
→ drain or claim failures.

DATA SHARING LOCKING	QUANTITY	/SECOND	/THREAD	/COMMI
-----	-----	-----	-----	-----
GLOBAL CONTENTION RATE (%)	0.23			
P/L-LOCKS XES RATE (%)	16.65			
LOCK REQUESTS (P-LOCKS)	25474.80	87.26	1.39	0.61
UNLOCK REQUESTS (P-LOCKS)	22541.22	77.21	1.23	0.54
CHANGE REQUESTS (P-LOCKS)	73.73	0.25	0.00	0.00
SYNCH.XES - LOCK REQUESTS	223.6K	762.48	12.12	5.33
SYNCH.XES - CHANGE REQUESTS	10351.93	35.46	0.56	0.25
SYNCH.XES - UNLOCK REQUESTS	203.4K	696.86	11.08	4.87
ASYNCH.XES - RESOURCES	8.73	0.03	0.00	0.00
SUSPENDS - IRLM GLOBAL CONT	760.56	2.61	0.04	0.02
SUSPENDS - XES GLOBAL CONT.	0.00	0.00	0.00	0.00
SUSPENDS - FALSE CONTENTION	239.61	0.82	0.01	0.01
INCOMPATIBLE RETAINED LOCK	0.00	0.00	0.00	0.00
NOTIFY MESSAGES SENT	113.50	0.39	0.01	0.00
NOTIFY MESSAGES RECEIVED	49.48	0.17	0.00	0.00
P-LOCK/NOTIFY EXITS ENGINES	500.00	N/A	N/A	N/A
P-LCK/NFY EX.ENGINE UNAVAIL	0.00	0.00	0.00	0.00
PSET/PART P-LCK NEGOTIATION	138.72	0.48	0.01	0.00
PAGE P-LOCK NEGOTIATION	385.13	1.32	0.02	0.01
OTHER P-LOCK NEGOTIATION	41.71	0.14	0.00	0.00
P-LOCK CHANGE DURING NEG.	446.25	1.53	0.02	0.01

What is the global contention rate?

Should be under 5%

Should be very low for DB2 V8

NFM or beyond (locking protocol 2)

Investigate false contention using

RMF CF Activity Report






GLOBAL DDF ACTIVITY	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
DBAT QUEUED-MAXIMUM ACTIVE	0.00	0.00	0.00	N/A
CONV.DEALLOC-MAX.CONNECTED	0.00	0.00	0.00	N/A
COLD START CONNECTIONS	0.00	0.00	0.00	0.00
WARM START CONNECTIONS	0.00	0.00	0.00	0.00
RESYNCHRONIZATION ATTEMPTED	0.00	0.00	0.00	0.00
RESYNCHRONIZATION SUCCEEDED	0.00	0.00	0.00	0.00
CUR TYPE 1 INACTIVE DBATS	0.00	N/A	N/A	N/A
TYPE 1 INACTIVE DBATS HWM	1.00	N/A	N/A	N/A
TYPE 1 CONNECTIONS TERMINAT	0.00	0.00	N/A	N/A
CUR TYPE 2 INACTIVE DBATS	14.00	N/A	N/A	N/A
TYPE 2 INACTIVE DBATS HWM	24.00	N/A	N/A	N/A
ACC QUEUED TYPE 2 INACT THR	927.00	0.51	N/A	N/A
CUR QUEUED TYPE 2 INACT THR	0.00	N/A	N/A	N/A
QUEUED TYPE 2 INACT THR HWM	4.00	N/A	N/A	N/A
CURRENT ACTIVE DBATS	2.00	N/A	N/A	N/A
ACTIVE DBATS HWM	10.00	N/A	N/A	N/A
TOTAL DBATS HWM	24.00	N/A	N/A	N/A
CURRENT DBATS NOT IN USE	2.00	N/A	N/A	N/A
DBATS NOT IN USE HWM	10.00	N/A	N/A	N/A
DBATS CREATED	4.00	N/A	N/A	N/A
POOL DBATS REUSED	1853.00	N/A	N/A	N/A

→ Exceed MAXDBAT?

→ Exceed CONDBAT?


→ Any Type 1 inactive?

→ How close to MAXDBAT?




CPU TIMES	TCB TIME	PREEMPT SRB	NONPREEMPT SRB	TOTAL TIME	PREEMPT IIP SRB	/COMMIT
SYSTEM SERVICES ADDRESS SPACE	7.364756	0.000000	13.895361	21.260107	N/A	0.000206
DATABASE SERVICES ADDRESS SPACE	0.623223	6.820931	23.002857	30.447011	0.023340	0.000295
IRLM	0.000016	0.000000	0.040580	0.040596	N/A	0.000000
DDF ADDRESS SPACE	1.645657	3:57.015778	2.445729	4:01.107164	0.422653	0.002339
TOTAL	9.633642	4:03.836708	39.384527	4:52.854878	0.445993	0.002841

- What is the relationship between SRB and TCB? If TCB is higher than SRB then we need to investigate further.
- Do MSTR and IRLM have higher values then DBM1?




**SHARE**  
Technology • Connections • Results

QUERY PARALLELISM	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
MAX.DEGREE OF PARALLELISM	1.00	N/A	N/A	N/A
PARALLEL GROUPS EXECUTED	171.00	0.00	0.00	0.00
RAN AS PLANNED	0.00	0.00	0.00	0.00
RAN REDUCED	0.00	0.00	0.00	0.00
SEQUENTIAL-CURSOR	0.00	0.00	0.00	0.00
SEQUENTIAL-NO BUFFER	50.00	0.00	0.00	0.00
SEQUENTIAL-ENCLAVE SER.	0.00	0.00	0.00	0.00
ONE DB2 - COORDINATOR = NO	0.00	0.00	0.00	0.00
ONE DB2 - ISOLATION LEVEL	0.00	0.00	0.00	0.00
ONE DB2 - DCL TTABLE	0.00	0.00	0.00	0.00
MEMBER SKIPPED (%)	N/C			
REFORM PARAL-CONFIG CHANGED	0.00	0.00	0.00	0.00
REFORM PARAL-NO BUFFER	0.00	0.00	0.00	0.00



- Any parallelism related problems? Not enough storage or other issues?




**SHARE**

Technology • Connections • Results

IFC DEST.	WRITTEN	NOT WRTN	BUF.OVER	NOT ACCP	WRT.FAIL	IFC RECORD COUNTS	WRITTEN	NOT WRTN
-----	-----	-----	-----	-----	-----	-----	-----	-----
SMF	55342.0K	0.00	0.00	0.00	0.00	SYSTEM RELATED	20375.00	0.00
GTF	0.00	0.00	N/A	0.00	0.00	DATABASE RELATED	16306.00	0.00
OP1	526.9M	162.2K	N/A	31119.00	N/A	ACCOUNTING	62876.8K	591.00
OP2	48798.2K	0.00	N/A	0.00	N/A	START TRACE	5.00	0.00
OP3	2.00	0.00	N/A	0.00	N/A	STOP TRACE	5.00	0.00
OP4	0.00	0.00	N/A	0.00	N/A	SYSTEM PARAMETERS	5801.00	1.00
OP5	0.00	0.00	N/A	0.00	N/A	SYS.PARMS-BPOOLS	1723.00	0.00
OP6	0.00	0.00	N/A	0.00	N/A	AUDIT	8758.2K	0.00
OP7	0.00	0.00	N/A	0.00	N/A			
OP8	0.00	0.00	N/A	0.00	N/A	TOTAL	71679.2K	592.00
RES	0.00	N/A	N/A	N/A	N/A			
TOTAL	631.0M	162.2K		31119.00	0.00			

Review the “NOT” columns. Any write recording problems? Any SMF data missing that will throw off results? In our scenario the problem was OP1, which means the monitor the customer was using is missing a considerable amount of data. In this scenario we start by reviewing the ZPARM value for MONSIZE.



LATCH CNT	/SECOND	/SECOND	/SECOND	/SECOND
-----	-----	-----	-----	-----
LC01-LC04	0.00	0.00	2.16	0.00
LC05-LC08	0.00	0.00	0.00	0.00
LC09-LC12	0.00	0.20	0.00	0.06
LC13-LC16	0.32	219.02	0.00	0.03
LC17-LC20	0.00	0.00	8.97	0.00
LC21-LC24	0.03	0.00	22.13	5169.81
LC25-LC28	2.96	0.00	1.82	0.18
LC29-LC32	0.01	0.57	4.96	1.37

How many latches per second? If we are in the thousands or very high in a relatively short period we need to review why the specific latch count is excessive.



DBM1 AND MVS STORAGE BELOW 2 GB			QUANTITY	DBM1 AND MVS STORAGE BELOW 2 GB CONTINUED			QUANTITY
-----			-----	-----			-----
TOTAL DBM1 STORAGE BELOW 2 GB	(MB)	402.04	24 BIT LOW PRIVATE	(MB)	0.21		
TOTAL GETMAINED STORAGE	(MB)	91.36	24 BIT HIGH PRIVATE	(MB)	0.33		
VIRTUAL BUFFER POOLS	(MB)	N/A	31 BIT EXTENDED LOW PRIVATE	(MB)	47.74		
VIRTUAL POOL CONTROL BLOCKS	(MB)	N/A	31 BIT EXTENDED HIGH PRIVATE	(MB)	442.33		
EDM POOL	(MB)	87.89	EXTENDED REGION SIZE (MAX)	(MB)	1557.00		
COMPRESSION DICTIONARY	(MB)	N/A	EXTENDED CSA SIZE	(MB)	400.80		
CASTOUT BUFFERS	(MB)	N/A					
DATA SPACE LOOKASIDE BUFFER	(MB)	N/A	AVERAGE THREAD FOOTPRINT	(MB)	2.98		
HIPERPOOL CONTROL BLOCKS	(MB)	N/A	MAX NUMBER OF POSSIBLE THREADS		390		
DATA SPACE BP CONTROL BLOCKS	(MB)	N/A					
TOTAL VARIABLE STORAGE	(MB)	256.29	AVERAGE THREAD FOOTPRINT (TYPE II)	(MB)	1.01		
TOTAL AGENT LOCAL STORAGE	(MB)	143.38	MAX NUMBER OF POSSIBLE THREADS (TYPE II)		1037		
. . .							
DBM1 STORAGE ABOVE 2 GB			QUANTITY	REAL AND AUXILIARY STORAGE			QUANTITY
-----			-----	-----			-----
FIXED STORAGE	(MB)	7.56	REAL STORAGE IN USE	(MB)	3826.16		
GETMAINED STORAGE	(MB)	1508.60	AUXILIARY STORAGE IN USE	(MB)	164.98		
COMPRESSION DICTIONARY	(MB)	194.48					
IN USE EDM DBD POOL	(MB)	33.15					


Keep in mind that this is cut at the end of the statistics interval – how large is the interval? This is true for much of the report, however how useful is it in this case? If the interval is small enough and therefore useful, review the total virtual and real memory used. Also review “MAX NUMBER OF POSSIBLE THREADS”. Once trended over time, determine increase or decrease of workloads and memory.

- OMPE can also be used to download this information to a worksheet similar to MEMU2

BP18	GENERAL	QUANTITY	/SECOND	/THREAD	/COMMIT	BP18	READ OPERATIONS	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
CURRENT ACTIVE BUFFERS		1908.46	N/A	N/A	N/A	BPOOL HIT RATIO (%)		-10.13			
UNAVAIL.BUFFER-VPOOL FULL		0.00	0.00	0.00	0.00	GETPAGE REQUEST	4799.3M	27.9K	622.22	93.42	
NUMBER OF DATASET OPENS		15724.00	0.09	0.00	0.00	GETPAGE REQUEST-SEQUENTIAL	4479.3M	26.1K	580.74	87.19	
BUFFERS ALLOCATED - VPOOL		190.9K	N/A	N/A	N/A	GETPAGE REQUEST-RANDOM	320.0M	1861.51	41.49	6.23	
DFHSM MIGRATED DATASET		0.00	0.00	0.00	0.00	SYNCHRONOUS READS	580.5M	3377.01	75.26	11.30	
DFHSM RECALL TIMEOUTS		0.00	0.00	0.00	0.00	SYNCHRON. READS-SEQUENTIAL	36532.8K	212.52	4.74	0.71	
VPOOL EXPANS. OR CONTRACT.		0.00	0.00	0.00	0.00	SYNCHRON. READS-RANDOM	544.0M	3164.49	70.53	10.59	
VPOOL OR HPOOL EXP.FAILURE		0.00	0.00	0.00	0.00	GETPAGE PER SYN.READ-RANDOM		0.59			
CONCUR.PREF.I/O STREAMS-HWM		0.00	N/A	N/A	N/A	SEQUENTIAL PREFETCH REQUEST	106.3M	618.62	13.79	2.07	
PREF.I/O STREAMS REDUCTION		16582.00	0.10	0.00	0.00	SEQUENTIAL PREFETCH READS	91721.7K	533.57	11.89	1.79	
PARALLEL QUERY REQUESTS		7594.00	0.04	0.00	0.00	PAGES READ VIA SEQ.PREFETCH	3445.5M	20.0K	446.70	67.07	
PARALL.QUERY REQ.REDUCTION		7594.00	0.04	0.00	0.00	S.PRF.PAGES READ/S.PRF.READ		37.56			
PREF.QUANT.REDUCED TO 1/2		0.00	0.00	0.00	0.00	LIST PREFETCH REQUESTS	15554.8K	90.49	2.02	0.30	
PREF.QUANT.REDUCED TO 1/4		0.00	0.00	0.00	0.00	LIST PREFETCH READS	8657.5K	50.36	1.12	0.17	
						PAGES READ VIA LIST PREFETCH	130.4M	758.84	16.91	2.54	
						L.PRF.PAGES READ/L.PRF.READ		15.07			
						DYNAMIC PREFETCH REQUESTED	181.9M	1058.33	23.59	3.54	
						DYNAMIC PREFETCH READS	42651.4K	248.11	5.53	0.83	
						PAGES READ VIA DYN.PREFETCH	1128.8M	6566.37	146.35	21.97	
						D.PRF.PAGES READ/D.PRF.READ		26.47			
						PREF.DISABLED-NO BUFFER	0.00	0.00	0.00	0.00	
						PREF.DISABLED-NO READ ENG	0.00	0.00	0.00	0.00	
						PAGE-INS REQUIRED FOR READ	339.9K	1.98	0.04	0.01	

## Some things to consider:

- BP hit ratio, sometimes it matters, some times it does not. Watch for negative.
- How much is sequential vs. random?
- Any disabled, reduction, or failure issues?
- Review “PAGE-INS” carefully? Any MVS memory issues?




BP18	GENERAL	QUANTITY	/SECOND	/THREAD	/COMMIT	BP18	READ OPERATIONS	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
	BUFFER UPDATES	1610.4M	9367.97	208.78	41.39		MAX WORKFILES CONCURR. USED	6.11	N/A	N/A	N/A
	PAGES WRITTEN	329.2M	1915.33	42.69	8.46		MERGE PASSES REQUESTED	18178.4K	105.75	2.36	0.47
	BUFF.UPDATES/PAGES WRITTEN	4.89					MERGE PASS DEGRADED-LOW BUF	0.00	0.00	0.00	0.00
							WORKFILE REQ.REJCTD-LOW BUF	0.00	0.00	0.00	0.00
	SYNCHRONOUS WRITES	5738.2K	33.38	0.74	0.15		WORKFILE REQ-ALL MERGE PASS	36377.8K	211.62	4.72	0.94
	ASYNCHRONOUS WRITES	127.9M	744.29	16.59	3.29		WORKFILE NOT CREATED-NO BUF	0.00	0.00	0.00	0.00
							WORKFILE PRF NOT SCHEDULED	0.00	0.00	0.00	0.00
	PAGES WRITTEN PER WRITE I/O	2.46									
	HORIZ.DEF.WRITE THRESHOLD	8589.9M	50.0K	1113.68	220.80						
	VERTI.DEF.WRITE THRESHOLD	2145.5K	12.48	0.28	0.06						
	DM THRESHOLD	0.00	0.00	0.00	0.00						
	WRITE ENGINE NOT AVAILABLE	85938.00	0.50	0.01	0.00						
	PAGE-INS REQUIRED FOR WRITE	0.00	0.00	0.00	0.00						

Review values for DM THRESHOLD, WRITE ENGINE NOT AVAILABLE, PAGE-INS REQUIRED FOR WRITE – determine what is causing these problems.

Review the WORKFILE section, any problems?

Review the TOTAL section after all of the buffer pools sections. Actually start at the TOTAL section and work your way backwards. This will provide a birds eye view for potential specific problems to review. In the TOTAL section, ignore such things as the BP hit ratio.





# Newer versions of OMPE will break down the hit ratio

BP2	GENERAL	QUANTITY	/SECOND	/THREAD	/COMMIT	BP2	READ OPERATIONS	QUANTITY	/SECOND	/THREAD	/COMMIT
	CURRENT ACTIVE BUFFERS	1642.68	N/A	N/A	N/A		BPOOL HIT RATIO (%)	75.91			
	UNAVAIL.BUFFER-VPOOL FULL	0.00	0.00	0.00	0.00		BPOOL HIT RATIO (%) SEQU	-78.35			
							BPOOL HIT RATIO (%) RANDOM	97.87			
	NUMBER OF DATASET OPENS	120.00	0.01	0.00	0.00		GETPAGE REQUEST	40236.7K	2794.02	1095.89	134.97
	BUFFERS ALLOCATED - VPOOL	30000.00	N/A	N/A	N/A		GETPAGE REQUEST-SEQUENTIAL	5015.0K	348.24	136.59	16.82
							GETPAGE REQUEST-RANDOM	35221.7K	2445.78	959.30	118.14
	DFHSM MIGRATED DATASET	0.00	0.00	0.00	0.00		SYNCHRONOUS READS	749.3K	52.03	20.41	2.51
	DFHSM RECALL TIMEOUTS	0.00	0.00	0.00	0.00		SYNCHRON. READS-SEQUENTIAL	284.7K	19.77	7.76	0.96
	VPOOL EXPANS. OR CONTRACT.	0.00	0.00	0.00	0.00		SYNCHRON. READS-RANDOM	464.6K	32.26	12.65	1.56
	VPOOL OR HPOOL EXP.FAILURE	0.00	0.00	0.00	0.00		GETPAGE PER SYN.READ-RANDOM	76.81			

• • • •


- What is really dragging down your buffer pool hit ratio?
- What is going on between random and sequential?

GROUP BP1	QUANTITY	/SECOND	/THREAD	/COMMIT	GROUP BP1	CONTINUED	QUANTITY	/SECOND	/THREAD	/COMMIT
GROUP BP R/W RATIO (%)	23.30	N/A	N/A	N/A	WRITE AND REGISTER		1867.44	6.40	0.10	0.04
GBP-DEPENDENT GETPAGES	1696.5K	5811.13	92.39	40.63	WRITE AND REGISTER MULT		3576.76	12.25	0.19	0.09
SYN.READ(XI)-DATA RETURNED	350.21	1.20	0.02	0.01	CHANGED PGS SYNC.WRTN		23245.52	79.63	1.27	0.56
SYN.READ(XI)-NO DATA RETURN	1163.15	3.98	0.06	0.03	CHANGED PGS ASYNC.WRTN		179.47	0.61	0.01	0.00
SYN.READ(NF)-DATA RETURNED	4509.99	15.45	0.25	0.11	PAGES WRITE & REG MULT		21557.54	73.84	1.17	0.52
SYN.READ(NF)-NO DATA RETURN	91635.56	313.89	4.99	2.19	READ FOR CASTOUT		71.79	0.25	0.00	0.00
UNREGISTER PAGE	70632.92	241.95	3.85	1.69	READ FOR CASTOUT MULT		1900.42	6.51	0.10	0.05
CLEAN PAGES SYNC.WRITTEN	0.00	0.00	0.00	0.00	WRITE TO SEC-GBP		N/A	N/A	N/A	N/A
REG.PAGE LIST (RPL) REQUEST	24469.78	83.82	1.33	0.59	CLEAN PAGES ASYNC.WRTN		N/A	N/A	N/A	N/A
NUMBER OF PAGES RETR.FROM GBP	2004.22	6.87	0.11	0.05	CLEAN PGS READ AFT.RPL		N/A	N/A	N/A	N/A
PGS READ FRM DASD AFTER RPL	N/A	N/A	N/A	N/A	PARTICIPAT.GBP REBUILD					
ASYNC.READ-DATA RETURNED	N/A	N/A	N/A	N/A						
PAGES CASTOUT	14204.19	48.65	0.77	0.34	PAGE P-LOCK LOCK REQ		4939.74	16.92	0.27	0.12
UNLOCK CASTOUT	697.50	2.39	0.04	0.02	SPACE MAP PAGES		4182.10	14.33	0.23	0.10
					DATA PAGES		757.65	2.60	0.04	0.02
READ CASTOUT CLASS	231.85	0.79	0.01	0.01	INDEX LEAF PAGES		0.00	0.00	0.00	0.00
READ DIRECTORY INFO	0.00	0.00	0.00	0.00						
READ STORAGE STATISTICS	45.59	0.16	0.00	0.00	PAGE P-LOCK UNLOCK REQ		4403.28	15.08	0.24	0.11
REGISTER PAGE	4798.11	16.44	0.26	0.11						
DELETE NAME	8.73	0.03	0.00	0.00	PAGE P-LOCK LOCK SUSP		25.22	0.09	0.00	0.00
ASYNCH GBP REQUESTS	26148.05	89.57	1.42	0.63	SPACE MAP PAGES		22.31	0.08	0.00	0.00
EXPLICIT X-INVALIDATIONS	0.00	0.00	0.00	0.00	DATA PAGES		2.91	0.01	0.00	0.00
CASTOUT CLASS THRESHOLD	0.97	0.00	0.00	0.00	INDEX LEAF PAGES		0.00	0.00	0.00	0.00
GROUP BP CASTOUT THRESHOLD	0.00	0.00	0.00	0.00						
GBP CHECKPOINTS TRIGGERED	1.94	0.01	0.00	0.00	PAGE P-LOCK LOCK NEG		26.19	0.09	0.00	0.00
CASTOUT ENGINE NOT AVAIL.	N/A	N/A	N/A	N/A	SPACE MAP PAGES		26.19	0.09	0.00	0.00
WRITE ENGINE NOT AVAILABLE	N/A	N/A	N/A	N/A	DATA PAGES		0.00	0.00	0.00	0.00
READ FAILED-NO STORAGE	N/A	N/A	N/A	N/A	INDEX LEAF PAGES		0.00	0.00	0.00	0.00
WRITE FAILED-NO STORAGE	0.00	0.00	0.00	0.00						
WRITE TO SEC-GBP FAILED	0.00	0.00	0.00	0.00						
DELETE NAME LIST SEC-GBP	698.47	2.39	0.04	0.02						
DELETE NAME FROM SEC-GBP	2.91	0.01	0.00	0.00						
UNLOCK CASTOUT STATS SEC-GBP	0.00	0.00	0.00	0.00						
ASYNCH SEC-GBP REQUESTS	2.91	0.01	0.00	0.00						

How is CF  
batching  
working?

Enough storage in GBP structures?

Tie this view of GBP activity with what is in the  
RMF CF Activity Report.



**SHARE**

Technology • Connections • Results

DRDA REMOTE LOCS	SENT	RECEIVED	KIHDB2A	SENT	RECEIVED
-----	-----	-----	-----	-----	-----
TRANSACTIONS	0.00	214.00	TRANSACTIONS	0.00	2.00
CONVERSATIONS	0.00	214.00	CONVERSATIONS	0.00	2.00
CONVERSATIONS QUEUED	0.00		CONVERSATIONS QUEUED	0.00	
SQL STATEMENTS	0.00	5153.0K	SQL STATEMENTS	0.00	24.00
SINGLE PHASE COMMITS	0.00	993.9K	SINGLE PHASE COMMITS	0.00	0.00

Are there different DRDA locations? If there are, we know that we are using private protocol. Look at using DRDA instead of private protocol.

## What are we looking for in the OMPE Accounting Reports?

- Lots of different things, but the following pages illustrate some of the key details we look for.
- Keep in mind, some values are averages while others are totals.

# FAQ: Why run different types of Accounting reports?



- To see the big picture, we run an Accounting Short Report by connection type. At times we will further investigate different categories within a connection type. This can be viewed as a 20,000 foot approach.
- For a 10,000 foot approach, we run an Accounting Short Report. We can restrict the report by categories depending on the amount of data and/or looking for specific problems.
- To expand on the 20,000 foot approach, we will run the same report by connection type, but this time as a Long Report.
- For specific issues, we run Accounting Long Reports.
- As in the Statistics runs, some fields are averages, some are totals. We must always be careful to understand what we are looking for.

# Accounting Short Report by Connection Type




CONNTYPE	#OCCURS	#COMMIT	INSERTS	OPENS	PREPARE	CLASS2	EL.TIME	BUF.UPDT	LOCK	SUS
	#DISTR	SELECTS	UPDATES	CLOSES	CLASS1	EL.TIME	CLASS2	CPUTIME	SYN.READ	#LOCKOUT
	#ROLLBK	FETCHES	MERGES	DELETES	CLASS1	CPUTIME	GETPAGES	TOT.PREF		
CICS	1436	1434	0.09	11.23		0.00	0.029457	64.69		0.08
	0	43.69	0.05	11.23		0.064630	0.007892	3.52		0
	1	30.76	N/A	0.06		0.011044	142.23	0.17		
. . .										
DRDA	70898	71914	0.41	0.31		0.13	0.061132	7.18		0.02
	70866	0.19	0.01	0.12		0.085829	0.006714	11.69		0
	1078	0.57	N/A	0.02		0.007213	336.45	11.36		
. . .										
*** GRAND TOTAL ***	570426	572475	0.18	23.48		0.02	0.051351	73.65		0.14
	71013	26.03	0.06	22.92		0.482827	0.014834	6.86		0
	1167	145.68	N/A	0.10		0.038221	434.77	10.97		

- Which connection type is doing the most amount of work? First, we really need to understand what we are trying to find. Is the issue CPU, elapsed time, getpages, etc.?
- Does anything look out of proportion in the big picture?




**First, we get a birds eye view of the elapsed time and class 2 time by reviewing the graph for distribution.**

ELAPSED TIME DISTRIBUTION



CLASS 2 TIME DISTRIBUTION



- What percentage of time was spent for what major reason?


AVERAGE	APPL(CL.1)	DB2 (CL.2)	IFI (CL.5)	CLASS 3 SUSPENSIONS	AVERAGE TIME	AV.EVENT	HIGHLIGHTS
ELAPSED TIME	1.042238	0.889962	N/P	LOCK/LATCH(DB2+IRLM)	0.014614	9.40	#OCCURRENCES : 4083
NONNESTED	1.042238	0.889962	N/A	SYNCHRON. I/O	0.108989	50.88	#ALLIEDS : 3
STORED PROC	0.000000	0.000000	N/A	DATABASE I/O	0.106553	50.46	#ALLIEDS DISTRIB: 0
UDF	0.000000	0.000000	N/A	LOG WRITE I/O	0.002437	0.42	#DBATS : 4080
TRIGGER	0.000000	0.000000	N/A	OTHER READ I/O	0.089436	47.55	#DBATS DISTRIB. : 0
				OTHER WRTE I/O	0.000008	0.00	#NO PROGRAM DATA: N/P
CP CPU TIME	0.243579	0.242119	N/P	SER.TASK SWTCH	0.002463	0.08	#NORMAL TERMINAT: 0
AGENT	0.243579	0.242119	N/A	UPDATE COMMIT	0.000165	0.02	#DDFRRSF ROLLUP: 409
NONNESTED	0.243579	0.242119	N/P	OPEN/CLOSE	0.000221	0.00	#ABNORMAL TERMIN: 0
STORED PRC	0.000000	0.000000	N/A	SYSLGRNG REC	0.001881	0.05	#CP/X PARALLEL. : 0
UDF	0.000000	0.000000	N/A	EXT/DEL/DEF	0.000154	0.00	#IO PARALLELISM : 0
TRIGGER	0.000000	0.000000	N/A	OTHER SERVICE	0.000041	0.00	#INCREMENT. BIND: 0
PAR.TASKS	0.000000	0.000000	N/A	ARC.LOG(QUIES)	0.000000	0.00	#COMMITTS : 4011
				LOG READ	0.000000	0.00	#ROLLBACKS : 80
SECP CPU	0.133938	N/A	N/A	DRAIN LOCK	0.000000	0.00	#SVPT REQUESTS : 0
				CLAIM RELEASE	0.000000	0.00	#SVPT RELEASE : 0
SE CPU TIME	0.000000	0.000000	N/A	PAGE LATCH	0.000000	0.00	#SVPT ROLLBACK : 0
NONNESTED	N/A	N/A	N/A	NOTIFY MSGS	0.000000	0.00	MAX SQL CASC LVL: 0
STORED PROC	N/A	N/A	N/A	GLOBAL CONTENTION	0.001430	0.60	UPDATE/COMMIT : 0.98
UDF	N/A	N/A	N/A	COMMIT PH1 WRITE I/O	0.000000	0.00	SYNCH I/O AVG. : 0.002142
TRIGGER	N/A	N/A	N/A	ASYNCH CF REQUESTS	0.000054	0.49	
				TCP/IP LOB	N/A	N/A	
PAR.TASKS	N/A	N/A	N/A	TOTAL CLASS 3	0.216994	109.00	
. . .							
NOT ACCOUNT.		430849	N/A				

- Where are we spending time?
- Are we using zIIPs/zAAPs? Efficiently?
- Not accounted for time high?

- If suspension time was a primary concern, what should we focus on?

- How many occurrences?
- Disk problems?
- Review all fields





GLOBAL	CONTENTION	L-LOCKS	AVERAGE TIME	AV.EVENT	GLOBAL	CONTENTION	P-LOCKS	AVERAGE TIME	AV.EVENT
L-LOCKS			0.001023	0.25	P-LOCKS			0.000408	0.36
PARENT (DB, TS, TAB, PART)			0.000028	0.12	PAGESET/PARTITION			0.000000	0.00
CHILD (PAGE, ROW)			0.000004	0.03	PAGE			0.000004	0.00
OTHER			0.000991	0.10	OTHER			0.000404	0.35


- If locking time was high, review this section closely.

SQL DML	AVERAGE	TOTAL	SQL DCL	TOTAL	SQL DDL	CREATE	DROP	ALTER	LOCKING	AVERAGE	TOTAL
SELECT	0.00	0	LOCK TABLE	0	TABLE	0	0	0	TIMEOUTS	0.00	0
INSERT	0.57	2320	GRANT	0	CRT TTABLE	0	N/A	N/A	DEADLOCKS	0.00	0
UPDATE	0.40	1633	REVOKE	0	DCL TTABLE	0	N/A	N/A	ESCAL.(SHARED)	0.00	0
MERGE	N/A	N/A	SET CURR.SQLID	0	AUX TABLE	0	N/A	N/A	ESCAL.(EXCLUS)	0.00	0
DELETE	0.02	69	SET HOST VAR.	0	INDEX	0	0	0	MAX PG/ROW LOCKS HELD	85.88	14954
			SET CUR.DEGREE	0	TABLESPACE	0	0	0	LOCK REQUEST	239.72	978793
DESCRIBE	3.60	14713	SET RULES	0	DATABASE	0	0	0	UNLOCK REQUEST	4.00	16316
DESC.TBL	0.00	0	SET CURR.PATH	0	STOGROUP	0	0	0	QUERY REQUEST	0.00	0
PREPARE	3.81	15550	SET CURR.PREC.	0	SYNONYM	0	0	N/A	CHANGE REQUEST	3.17	12937
OPEN	3.90	15932	CONNECT TYPE 1	0	VIEW	0	0	N/A	OTHER REQUEST	0.00	0
FETCH	10.19	41624	CONNECT TYPE 2	0	ALIAS	0	0	N/A	TOTAL SUSPENSIONS	0.43	1750
CLOSE	0.00	18	SET CONNECTION	0	PACKAGE	N/A	0	N/A	LOCK SUSPENSIONS	0.00	18
			RELEASE	0	PROCEDURE	0	0	0	IRLM LATCH SUSPENS.	0.42	1732
DML-ALL	22.50	91859	CALL	0	FUNCTION	0	0	0	OTHER SUSPENS.	0.00	0
			ASSOC LOCATORS	0	TRIGGER	0	0	N/A			
			ALLOC CURSOR	0	DIST TYPE	0	0	N/A			
			HOLD LOCATOR	0	SEQUENCE	0	0	0			
			. . .								

•How much DML in relation to the entire process? Compare with Information from 2 pages ago.

•Do any DCL and DDL values seem unusual?

•Any timeouts, deadlocks, or lock escalations?  
•Review other locking information as well.



**SHARE**  
Technology • Connections • Results


NORMAL TERM.	AVERAGE	TOTAL	ABNORMAL TERM.	TOTAL	IN DOUBT	TOTAL	DRAIN/CLAIM	AVERAGE	TOTAL
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
NEW USER	0.00	0	APPL.PROGR. ABEND	0	APPL.PGM ABEND	0	DRAIN REQUESTS	0.00	0
DEALLOCATION	0.00	0	END OF MEMORY	0	END OF MEMORY	0	DRAIN FAILED	0.00	0
APPL.PROGR. END	0.00	0	RESOL.IN DOUBT	0	END OF TASK	0	CLAIM REQUESTS	16.00	65347
RESIGNON	0.00	0	CANCEL FORCE	0	CANCEL FORCE	0	CLAIM FAILED	0.00	0
DBAT INACTIVE	0.00	0							
TYPE2 INACTIVE	0.00	0							
RRS COMMIT	0.00	0							
END USER THRESH	0.10	409							
BLOCK STOR THR	0.00	0							
STALENESS THR	0.00	1							

- Any failures reported?
- Any ABENDs or unusual results?

DATA SHARING	AVERAGE	TOTAL	QUERY PARALLELISM	AVERAGE	TOTAL
P/L-LOCKS XES(%)	3.67	N/A	MAXIMUM MEMBERS USED	N/A	1
LOCK REQ - PLOCKS	0.45	1843	MAXIMUM DEGREE	N/A	0
UNLOCK REQ - PLOCKS	0.07	282	GROUPS EXECUTED	0.00	0
CHANGE REQ - PLOCKS	0.08	338	RAN AS PLANNED	0.00	0
LOCK REQ - XES	8.77	35804	RAN REDUCED	0.00	0
UNLOCK REQ - XES	0.30	1228	ONE DB2-COORDINATOR = NO	0.00	0
CHANGE REQ - XES	1.27	5183	ONE DB2-ISOLATION LEVEL	0.00	0
SUSPENDS - IRLM	0.30	1216	ONE DB2-DCL TEMPORARY TABLE	0.00	0
SUSPENDS - XES	0.00	3	SEQUENTIAL-CURSOR	0.00	0
SUSPENDS - FALSE	N/A	N/A	SEQUENTIAL-NO ESA SORT	0.00	0
INCOMPATIBLE LOCKS	0.00	0	SEQUENTIAL-NO BUFFER	0.00	0
NOTIFY MSGS SENT	0.00	0	SEQUENTIAL-ENCLAVE SERVICES	0.00	0
			MEMBER SKIPPED (%)	N/C	N/A
			DISABLED BY RLF	N/P	N/P
			REFORM PARAL-CONFIG	0.00	0
			REFORM PARAL-NO BUF	0.00	0

•Any Data Sharing lock problems?

•Any query parallelism problems?



**SHARE**


Technology • Connections • Results

STORED PROCEDURES	AVERAGE	TOTAL	UDF	AVERAGE	TOTAL	TRIGGERS	AVERAGE	TOTAL
CALL STATEMENTS	0.00	0	EXECUTED	0.00	0	STATEMENT TRIGGER	0.00	0
ABENDED	0.00	0	ABENDED	0.00	0	ROW TRIGGER	0.00	0
TIMED OUT	0.00	0	TIMED OUT	0.00	0	SQL ERROR OCCUR	0.00	0
REJECTED	0.00	0	REJECTED	0.00	0			

- How many SPs, UDFs, and triggers are execute? Does it seem to be a high number?
- Any errors or problems?

RID LIST	AVERAGE	TOTAL
USED	N/P	N/P
FAIL-NO STORAGE	N/P	N/P
FAIL-LIMIT EXCEEDED	N/P	N/P

- Any RID failures?




**SHARE**  
Technology • Connections • Results

BP5	BPOOL ACTIVITY	AVERAGE	TOTAL
-----			
	BPOOL HIT RATIO (%)	78.64	N/A
	GETPAGES	2.51	10260
	BUFFER UPDATES	0.15	597
	SYNCHRONOUS WRITE	0.00	0
	SYNCHRONOUS READ	0.45	1845
	SEQ. PREFETCH REQS	0.00	17
	LIST PREFETCH REQS	0.00	0
	DYN. PREFETCH REQS	0.00	2
	PAGES READ ASYNCHR.	0.09	348

- Review the buffer pool information.
- Are there many synchronous reads, etc.?
- What was the buffer pool hit ratio?


GROUP BP5	AVERAGE	TOTAL
-----	-----	-----
GBP-DEPEND GETPAGES	1.04	4274
READ(XI)-DATA RETUR	0.00	0
READ(XI)-NO DATA RT	0.00	0
READ(NF)-DATA RETUR	0.00	0
READ(NF)-NO DATA RT	0.00	9
PREFETCH PAGES READ	0.00	0
CLEAN PAGES WRITTEN	0.00	0
UNREGISTER PAGE	0.09	355
ASYNCH GBP REQUESTS	0.02	77
EXPLICIT X-INVALID	0.00	0
ASYNCH SEC-GBP REQ	0.00	0
PG P-LOCK LOCK REQ	0.02	66
SPACE MAP PAGES	0.02	66
DATA PAGES	0.00	0
INDEX LEAF PAGES	0.00	0
PG P-LOCK UNLOCK REQ	0.00	1
PG P-LOCK LOCK SUSP	0.00	3
SPACE MAP PAGES	0.00	3
DATA PAGES	0.00	0
INDEX LEAF PAGES	0.00	0
WRITE AND REGISTER	0.08	336
WRITE & REGISTER MULT	0.02	86
CHANGED PAGES WRITTEN	0.14	565
WRITE TO SEC-GBP	N/A	N/A



**SHARE**  
Technology • Connections • Results

Review the group buffer pool information.  
Since this data is limited to the plan, there are fewer fields here, but specifics can be helpful.

**SHARE**  
in Anaheim  
2011

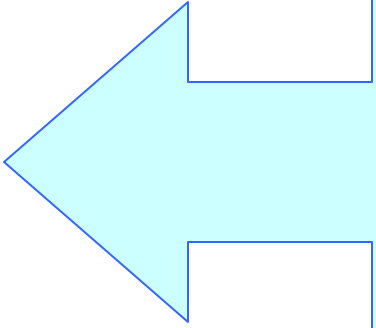


**SHARE**  
Technology • Connections • Results

# RMF Reports




```
ERB103I PPS: OPTIONS IN EFFECT
ERB103I PPS:  SUMMARY(TOT)  -- DEFAULT
ERB103I PPS:  SUMMARY(INT)  -- DEFAULT
ERB103I PPS:  NODELTA      -- DEFAULT
ERB103I PPS:  NOEXITS      -- DEFAULT
ERB103I PPS:  MAXPLEN(50)  -- DEFAULT
ERB103I PPS:  ETOD(0000,2400)  -- DEFAULT
ERB103I PPS:  STOD(0000,2400)  -- DEFAULT
ERB103I PPS:  PTOD(0000,2400)  -- DEFAULT
ERB103I PPS:  RTOD(0000,2400)  -- DEFAULT
ERB103I PPS:  SYSOUT(X)    -- SYSIN
ERB103I PPS:  REPORTS(CPU)  -- SYSIN
ERB103I PPS:  DATE(01011959,12312058)  -- DEFAULT
```

- 
- Were the correct options chosen?
  - Were the correct dates and times chosen?
  - Were the correct reports chosen?

```
z/OS V1R7                SYSTEM ID AAAA                START 01/21/2009-08.59.00  INTERVAL 00.14.59
                        CONVERTED TO z/OS V1R10 RMF      END   01/21/2009-10.29.00  CYCLE 1.000 SECONDS
NUMBER OF INTERVALS 6                TOTAL LENGTH OF INTERVALS 01.29.56
```

- Verify the correct dates and times are present. Keep in mind, RMF reports time as local.
- How long is the interval counter? Is it what we are looking for?
- How many cycles are recorded?
- Notice that we report the level of z/OS from both the source system as well as the system we created the reports on in IBM.



# RMF Summary Report

## R M F S U M M A R Y R E P O R T

z/OS V1R7

SYSTEM ID AAAA

START 01/21/2009-08.59.00 INTERVAL 00.14.59

CONVERTED TO z/OS V1R10 RMF

END 01/21/2009-10.29.00 CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 6

TOTAL LENGTH OF INTERVALS 01.29.56

-DATE	TIME	INT	CPU	DASD	DASD	TAPE	JOB	JOB	TSO	TSO	STC	STC	ASCH	ASCH	OMVS	OMVS	SWAP	DEMAND
MM/DD	HH.MM.SS	MM.SS	BUSY	RESP	RATE	RATE	MAX	AVE	MAX	AVE	MAX	AVE	MAX	AVE	MAX	AVE	RATE	PAGING
01/21	08.59.00	14.59	76.4	1.1	9938	356.0	91	86	39	35	254	245	0	0	9	7	0.00	37.12
01/21	09.14.00	15.00	77.0	1.3	9518	497.1	85	78	43	40	256	253	0	0	7	7	0.00	22.57
01/21	09.29.00	14.59	77.3	1.2	9866	358.2	82	80	43	39	254	251	0	0	7	7	0.00	51.27
01/21	09.44.00	14.59	77.7	1.0	10188	144.7	85	83	47	42	249	247	0	0	7	7	0.00	18.96
01/21	09.59.00	15.00	71.2	1.2	9033	531.4	94	91	48	44	247	238	0	0	7	7	0.00	180.69
01/21	10.14.00	14.59	46.4	1.6	5653	449.4	90	76	49	45	255	229	0	0	7	6	0.00	77.32
-TOTAL/AVERAGE			71.0	1.2	9033	389.5	94	82	49	41	256	244	0	0	9	7	0.00	64.66

- Review the interval dates and times, match them up to the OMPE reports we are reviewing.
- How are we doing with CPU busy? Looks really good – right? We need to look further, there is actually a big problem, but we cannot tell from the Summary Report. At times this report will show more realistic numbers, but not in this case.
- Any paging problems?

# RMF CPU Activity Report



## C P U A C T I V I T Y

z/OS V1R7

SYSTEM ID AAAA

DATE 01/21/2009

INTERVAL 14.59.968

CONVERTED TO z/OS V1R10 RMF

TIME 08.59.00

CYCLE 1.000 SECONDS

```
-CPU 2097  MODEL 705  H/W MODEL E12  SEQUENCE CODE 00000000000A486C  HIPERDISPATCH=NO
---CPU---  ----- TIME % -----  LOG PROC  --I/O INTERRUPTS--
NUM  TYPE  ONLINE  LPAR BUSY  MVS BUSY  PARKED  SHARE %  RATE  % VIA TPI
 0   CP   100.00  76.40     99.99     -----  74.3     0.39  88.51
 1   CP   100.00  76.40     99.99     -----  74.3     0.67  89.11
 2   CP   100.00  76.40     99.99     -----  74.3     0.50  90.18
 3   CP   100.00  76.39     99.99     -----  74.3    11685  12.20
TOTAL/AVERAGE                76.40     99.99                297.2    11686  12.21
 4   IIP   100.00  49.46     51.51     -----  55.0
TOTAL/AVERAGE                49.46     51.51                55.0
```

- How many CPs and specially engines are available?
- Was each processor online 100% of the time?
- Compare the LPAR BUSY time to the MVS BUSY time. The summary report showed a good LPAR BUSY rate, however we now see that MVS BUSY is almost 100%. We must investigate further, however we know that the CEC is saturated, even though the CPs do not seem to be.




# RMF Partition Data Report

MVS PARTITION NAME	AAAA	NUMBER OF PHYSICAL PROCESSORS	6	GROUP NAME	N/A
IMAGE CAPACITY	488	CP	5	LIMIT	N/A
NUMBER OF CONFIGURED PARTITIONS	11	AAP	0		
WAIT COMPLETION	NO	IFL	0		
DISPATCH INTERVAL	DYNAMIC	ICF	0		
		IIP	1		

- We need to review the big picture for the CEC, such as the number and type of processors.

# RMF Partition Data Report - continued



**SHARE**  
Technology • Connections • Results

----- PARTITION DATA -----							-- LOGICAL PARTITION PROCESSOR DATA --				-- AVERAGE PROCESSOR UTILIZATION PERCENTAGES					
-----MSU----- -CAPPING--							PROCESSOR-	----DISPATCH TIME DATA----		LOGICAL PROCESSORS		--- PHYSICAL PROCESSORS ---				
NAME	S	WGT	DEF	ACT	DEF	WLM%	NUM	TYPE	EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR	MGMT	EFFECTIVE	TOTAL
AAAA	A	600	0	298	NO	0.0	4	CP	00.45.48.269	00.45.50.298	76.34	76.40	0.05		61.07	61.12
AAAB	A	280	0	139	NO	0.0	2	CP	00.21.17.532	00.21.18.048	70.98	71.01	0.01		28.39	28.40
AAAC	A	50	0	25	NO	0.0	2	CP	00.03.54.790	00.03.54.967	13.04	13.05	0.00		5.22	5.22
AAAD	A	50	0	25	NO	0.0	2	CP	00.03.54.176	00.03.54.568	13.01	13.03	0.01		5.20	5.21
AAAF	A	20	0	0	NO	0.0	1	CP	00.00.00.000	00.00.00.000	0.00	0.00	0.00		0.00	0.00
DUMMY	A	10	0	0	NO	0.0	1	CP	00.00.00.000	00.00.00.000	0.00	0.00	0.00		0.00	0.00
*PHYSICAL*												0.04				0.04
TOTAL									01.14.54.768	01.14.59.806			0.11		99.89	100.0
-----																
AAAA	A	550					1	IIP	00.07.24.570	00.07.25.115	49.40	49.46	0.06		49.40	49.46
AAAB	A	250					1	IIP	00.00.06.017	00.00.06.068	0.67	0.67	0.01		0.67	0.67
AAAC	A	100					1	IIP	00.00.00.081	00.00.00.092	0.01	0.01	0.00		0.01	0.01
AAAD	A	100					1	IIP	00.00.29.992	00.00.30.106	3.33	3.35	0.01		3.33	3.35
*PHYSICAL*												0.14				0.14
TOTAL									00.08.00.661	00.08.02.628			0.22		53.41	53.63

- We review each LPAR, CP or specialty engine type, weighing factor, etc.
- Now we see the real problem. Because of the weighing factor, although the LPARs are running below 80%, the CEC itself is saturated at about 100% busy.
- Keep in mind, as we saw in the OMPE Statistics reports, RMF reports show you a picture from the end of the interval, before the record is cut.

**IARE**  
in Anaheim  
2011

# RMF Paging Activity Report


OPT = IEAOPTSA      MODE = ESAME      CENTRAL STORAGE MOVEMENT RATES - IN PAGES PER SECOND

```

-----
HIGH UIC (AVG) = 1180      (MAX) = 2540      (MIN) = 560
-
          WRITTEN TO          READ FROM          *--- CENTRAL STORAGE FRAME COUNTS ----*
          CENTRAL STOR      CENTRAL STOR          MIN            MAX            AVG
HIPERSPACE RT            0.19            0.00            92,308        92,478        92,371
PAGES
VIO          RT            0.17            0.12            0              6              2
PAGES
  
```

- We have to review the UIC value. Are we having paging problems during this interval

# RMF Paging Activity Report - continued




**SHARE**  
Technology • Connections • Results

## FRAME AND SLOT COUNTS

CENTRAL STORAGE				LOCAL PAGE DATA SET SLOT COUNTS			
	MIN	MAX	AVG		MIN	MAX	AVG
(91 SAMPLES)							
AVAILABLE	11,234	142,283	50,507	AVAILABLE SLOTS	4,090,162	4,194,608	4,120,823
SQA	14,864	15,223	15,031	VIO SLOTS	339	351	348
LPA	5,777	5,778	5,777				
CSA	25,586	26,890	26,337	NON-VIO SLOTS	870,236	974,670	944,012
LSQA	65,233	67,092	66,372				
REGIONS+SWA	2626907	2757614	2717875	BAD SLOTS	0	0	0
TOTAL FRAMES	2883584	2883584	2883584	TOTAL SLOTS	5,065,183	5,065,183	5,065,183
FIXED FRAMES				SHARED FRAMES AND SLOTS			
NUCLEUS	2,436	2,435	2,435	CENTRAL STORAGE	1,287	1,468	1,348
SQA	13,032	13,390	13,185				
LPA	143	144	143	FIXED TOTAL	28	55	33
CSA	2,706	2,718	2,713	FIXED BELOW 16 M	0	0	0
LSQA	43,478	44,990	44,383	AUXILIARY SLOTS	3,384	3,528	3,487
REGIONS+SWA	49,451	52,178	50,050	TOTAL	4,932	5,094	5,036
BELOW 16 MEG	723	760	727				
BETWEEN 16M-2G	69,652	70,593	70,308				
TOTAL FRAMES	111,774	115,365	112,911				

- How many frames do we have available? Can we afford to alter DB2 that will take up more memory and take away from storage? Want to add 100,000 32K buffer pool pages and make sure they are 100% memory backed? Check this type of report first.





**SHARE**  
Technology • Connections • Results


# RMF Workload Activity report - SCPER

REPORT BY: POLICY=AAAAPLEX    WORKLOAD=DDF    SERVICE CLASS=DISJ12A    RESOURCE GROUP=\*NONE    PERIOD=1    IMPORTANCE=2  
CRITICAL    =NONE

-TRANSACTIONS-	TRANS-TIME	HHH.MM.SS.TTT	--DASD I/O--	---SERVICE---	SERVICE TIME	---APPL %---	--PROMOTED--	---STORAGE---				
AVG	11.54	ACTUAL	135	SSCHRT 113.4	IOC	0	CPU 301.368	CP 16.78	BLK 0.000	AVG	0.00	
MPL	11.54	EXECUTION	132	RESP 3.0	CPU 12557K	SRB 0.000	AAPCP 0.00	ENQ 0.000	TOTAL	0.00		
ENDED	78139	QUEUED	3	CONN 0.3	MSO 0	RCT 0.000	IIPCP 1.96	CRM 0.000	SHARED	0.00		
END/S	86.82	R/S AFFIN	0	DISC 2.5	SRB 0	IIT 0.000						
#SWAPS	0	INELIGIBLE	0	Q+PEND 0.2	TOT 12557K	HST 0.000	AAP N/A				-PAGE-IN RATES-	
EXCTD	0	CONVERSION	0	IOSQ 0.0	/SEC 13952	AAP N/A	IIP 16.70				SINGLE	0.0
AVG ENC	11.55	STD DEV	605			IIP 150.314					BLOCK	0.0
REM ENC	0.00					ABSRPTN 1209					SHARED	0.0
MS ENC	0.00					TRX SERV 1209					HSP	0.0

GOAL: EXECUTION VELOCITY 40.0%    VELOCITY MIGRATION:    I/O MGMT 22.9%    INIT MGMT 22.9%

SYSTEM	RESPONSE TIME	EX	PERF	AVG	--EXEC USING%--	EXEC DELAYS %	-USING%-	--- DELAY % ---	%										
	VEL%	INDX	ADRSP	CPU	AAP	IIP	I/O	TOT	CPU	IIP	I/O	CRY	CNT	UNK	IDL	CRY	CNT	QUI	
AAAA	--N/A--	22.9	1.7	11.5	1.1	N/A	1.1	0.2	8.2	5.8	2.2	0.2	0.0	0.0	89	0.0	0.0	0.0	0.0



•We need to check the Performance index. If it is above 1, we need to Investigate further.


•Understand what all of the DB2 related services classes are and verify the performance index and other values.  
•Does the period, importance, and velocity make sense?:  
•Verify all DB2 related service classes are properly assigned.

# RMF Coupling Facility report

- Coupling Facility Usage Summary
  - Structure Summary
    - LOCK1 – size, requests per second
    - GBPs – size, directory reclaims, XIs due to directory reclaims
      - *Reclaims and XIs should be zero*
  - Storage Summary
    - Total used, total available
    - Compare values on both CF LPARs
      - *Sufficient for structures to rebuild to other CF*
  - Processor Summary
    - CF machine – model, CFCC level
    - Average CF utilization, number of processors
      - *Higher utilization ok with more processors*

## RMF Coupling Facility report, cont.

- Coupling Facility Structure Activity
  - LOCK1
    - Service times, async conversions, delays
      - *Acceptable values dependent on Parallel Sysplex configuration*
    - Contention
      - *Total contention under 5%, preferably under 2%*
      - *False contention under 1%*
  - GBPs
    - Service times, async conversions, delays
      - *Acceptable values dependent on Parallel Sysplex configuration*
    - Requests per second: if under 100, focus on other GBPs
    - GBP duplexing – focus on primary structure
    - XIs under ‘Totals’ generally expected



**SHARE**  
Technology • Connections • Results

# Common ZPARM Issues

## Common ZPARM Issues


- As you know, there are many ZPARM parameters. We have already seen in the OMPE Statistics and Accounting reports the effects of some ZPARM values set incorrectly or to an inefficient level. For example, setting DSMAX too low may cause problems with data set open and close activity.
- The ATS DB2 Health Check report covers these types of problems and provides recommendations for appropriate ZPARM values based on the customer environment. The following slides list some of the problem ZPARM values we encounter most often.

# ZPARM parameters commonly not set correctly



- ADMTPROC=,
- LRDRTHLD=0
- MAXKEEPD=5000. Are we using dynamic statement cache efficiently? If not, and we cannot improve the usage, we are taking up memory with little or no return.
- NPGTHRSH=0
- PARAMDEG, when the value is too high
- RECALL=YES and RECALLD=120 when using manual tape mounting
- SEQCACH=BYPASS
- SEQPRES=NO
- STATHIST=NONE
- STATSINT=30. Keep in mind what happens when DB2<sup>70</sup> crashes before it can destage data to disk.


## ZPARM parameters commonly not set correctly - continued



**E**  
results

- SYSADM, SYSADM2, SYSOPR1, and SYSOPR2. Use secondary auth ids when possible instead of specific userids.
- SYSTEM\_LEVEL\_BACKUPS. Set this to YES when using the BACKUP utility and the data is available.
- ARCRETN. Make sure you really keep the archive log data sets as long as you need them for recovery.
- BLKSIZE=28672. When using parallelism for mass recoveries on disk – reduce to 24K.
- UNIT=tape and UNIT2=tape. Both duplexed copies on the same output tape?
- DEALLCT=(1440) – depends on archive log media.
- MAXARCH – when less than 10000.
- CHKFREQ=5 or above
- <sup>71</sup> CONDBAT vs. MAXDBAT

## ZPARM parameters commonly not set correctly - continued




- CTHREAD in relation to IDBACK, IDFORE, and other non DRDA connections.
- CTHREAD+MAXDBAT – Can we really handle the workload?
- DLDFREQ=0
- IDXBPOOL and TBSBPOOL. Same buffer pool as other data?
- LOGAPSTG=0
- PCLOSEN and PCLOSET
- SMFACCT=1. Preferably (1,2,3,7,8)
- SMFSTAT=1. Preferably YES
- STATIME=5
- SYNCVAL=NO
- TBSBP8K, TBSBP16K, and TBSBP32K – watch out for buffer pools already allocated for specific purposes.
- TBSBPLOB=BP0



## ZPARM parameters commonly not set correctly - continued

- TBSBPXML=BP16K0
- URCHKTH=0
- URLGWTH=0
- WLMENV=name ← set to some not existent WLM environment.



**SHARE**  
Technology • Connections • Results

# Remaining Reports

## Output from DSNJU004

- From the output of DSNJU004, can we keep up with ZPARM requirements for ARCRETN and MAXARCH?
- How many of the active log data sets are reusable? Are there problems with keeping up with the tape writes when archiving?
- Review the archive log history. How often do you explicitly archive logs? With what MODE, etc.?

## LISTCAT output for the active log data sets and BSDS

- How large are the active log data sets?
- Are there problems with keeping up with the tape writes when archiving?
- How many active log data sets exist?
- What volumes are the data sets allocated to? Are the active log data sets and BSDS split by software duplexed data sets?
- From the previous bullet, if you have more than one disk controller, are the software duplexed data sets split across different controllers?

## What do we find in the RTS (Real Time Statistics) output?

- Most customers we deal with do not have RTS enabled (starting in DB2 9 RTS is always enabled). The ones that have the RTS enabled typically do not review the RTS data nor are they aware of the jewels that the RTS offers.
- RTS is what it sounds like - Real Time Statistics as opposed to the DB2 catalog which can have very stale data. We highly encourage customers to use the RTS to its full extent.
- Most customers do have RTS Stored Procedures DSNACCOR and in DB2 9 DSNACCOX installed, however they typically do not execute it.
- We use the RTS information for such information as the number of pseudo deletes, extents, the last time utilities were run, etc.

## BIND output

- What BIND parameters are used?
  - What is the isolation level?
  - What is CURRENTDATA set to?
- How often are BIND and REBIND executed? Based on what criteria?
- How often is EXPLAIN executed? Any changes after BIND or REBIND?
- How often are triggers rebound?
- When in DB2 9, which flavor is used for plan stability – OFF, BASIC, or EXTENDED?

## DISPLAY BUFFERPOOL(ACTIVE) DETAIL(\*) output

- Much of the information for this output is also found in the OMPE Statistics Reports. Here are some items that are not:
  - Is the buffer pool page fixed?
  - Is AUTOSIZE turned on?
  - What is the stealing method?
  - What are the different thresholds?
- We may want to look at buffer pool information cumulatively or in total.

## **DIS GBPOOL(\*) GDETAIL(\*) TYPE(GCONN) output**

- Much of the information for this output is also found in the OMPE Statistics Reports plus the CFRM policy. Here are some items that are not (or current vs. pending):
  - What is the directory to data ratio?
  - What is the GBPCACHE attribute?
  - What are the thresholds and intervals?
  - What is the recovery status and is auto recovery on?
  - What is the CF level?
  - Are there directory reclaims? XIs due to directory reclaims?
- We may want to look at group buffer pool information cumulatively or in total.



## Explanation of buffer pool designations

- Which buffer pools have memory allocated to them and for what use?
- How does the customer decide to segregate buffer pools?
  - Is it efficient?
  - What characteristics are used?
- How often are the buffer pool assignments and statistics reviewed?
- What resides in BP0? Does anything need to be moved to a different buffer pool?


## DIAGNOSE DISPLAY MEPL output

- Part of the Health Check is to understand your maintenance philosophy. Do you apply normal maintenance and HIPERs in a timely fashion?
- We take the output from DIAGNOSE DISPLAY MEPL and feed it into an IBM program that tells us:
  - A summary of missing HIPERs, number of PEs, and the number of inconsistencies.
  - Specifics behind the summary numbers.
  - We also see the percentage of PTFs applied by quarter over a period of time.

## **Output of DB2 started tasks, such as MSTR, DBM1, IRLM, DIST, ADMT, and WLM stored procedure address spaces for DB2.**

- We take a look at all of the STC. This provides a good overall view of what is happening in DB2.
- Does anything look unusual in any of the STC outputs? We can get a feel for a number of problems by reviewing the STCs. These problems may not occur during the interval covered by the SMF data provided. In that case the STC output is a good way of seeing what problems may have occurred on other dates and times.


# OSC/OE or VE output for the top 10 problem SQL statements



**E**  
results

- Optimization Service Center (OSC) and Optimization Expert (OE) were announced with DB2 9 for z/OS and are available for DB2 V8 for z/OS. Visual Explain (VE) is available for DB2 V8 and prior versions. OSC and VE are products available at no additional cost.
- These tools provide many helpful hints, tips, and solutions. For example, you have heard about RUNSTATS and use of quantiles and histograms, but how do you use them? These products will provide the exact RUNSTATS utility input required.
- OE is available for an additional cost. It includes all of the functions of OSC plus powerful advisors. For example, Index advisor (IA) will not only tell you what indexes you are missing, it even creates the DDL for you.
- The functions in OSC and OE can be applied to a single query or an application workload.

# OSC/OE or VE output for the top 10 problem SQL statements



**SHARE**  
Technology • Connections • Results

- Most customers we work with do not have OSC/OE or VE installed. We request that they install VE at a minimum to produce the output we need for our analysis.
- There is a learning curve, but once you understand how to take advantage of these powerful tools, they will be among your favorites.
- We do not try to solve all of your application problems with these tools. We ask for a small sample of application data to analyze so you can understand our analysis and findings. We take the output you provide and apply the tools to that output, then we show you what we have learned and what recommendations the tools produce. You can then apply a similar process to other application and SQL problems you may have.


**SHARE**  
in Anaheim  
2011

## CFRM Policy

- Are SIZE, INITSIZE, and MINSIZE set correctly?
- Is Auto Alter used?
- Are the group buffer pools duplexed?
- Should the SCA and lock structures duplexed?
- Are the structures in a specific preference order?
- What is the physical location of the structures? Are they internal coupling facilities or external? We need to make sure you are covered when a failure does occur.

## WLM Policy

- Are the DB2 related tasks defined in a proper hierarchy?
- Are the DRDA related tasks setup as a hierarchy or does everyone get the same priority? Do the CEO and mail room operators have the same priority?
- What is the step down approach when required to leave a period? (Period to period migration)
- What happens when the CPU or CEC are at or close to 100% busy? What work suffers at that point? Is that what you and the business expect?



**SHARE**  
Technology • Connections • Results

# Final Analysis



## What to expect once the analysis is complete

- The entire process from the time we start working on your project is typically four weeks.
  - Week one is for customer interviews and receiving the required documentation. This is also a perfect time for customers to work along side us and gain additional skills and expertise from knowledge transfer. Part of this week is typically onsite at the customer location. Typically we have an initial conference call as well as the data FTPed to us prior to week one.
  - Weeks two and three are offsite. We further analyze the documentation provided and review our interview notes. We may work with customer personnel over the phone as we go through this process, which provides additional knowledge transfer.
  - Week four we finalize the report.
- We deliver the detailed report. This report includes an executive summary and typically requires dozens of pages to present our findings and analysis.

## What to expect once the analysis is complete - continued

- After we deliver the final report, the customer typically distributes the report and discusses it internally.
- We schedule a conference call with the customer to discuss the report and answer any questions the customer may have.

# Interested in an ATS DB2 Health Check?



- Contact John Iczkovits at [iczkovit@us.ibm.com](mailto:iczkovit@us.ibm.com) (757) 564-2205