

DB2 Performance in Practice

Steve McDougall
Independent Database Consultant



BMC DB2 Performance Seminars
Europe, September 16 - 25 2003



“Tuning DB2 is like flying a Jumbo Jet with a thousand dials and switches”



DB2 Performance in Practice

Agenda

- The tuning challenge
- Memory resource tuning
- Gone fishing
- Tuning SQL
- Index design issues




All screen shots courtesy of
T-Systems, Bielefeld, Germany



Copyright 2003 Steve McDougall - DB2 Performance in Practice

3

Increased Complexity



Slide courtesy of BMC Software



Copyright 2003 Steve McDougall - DB2 Performance in Practice

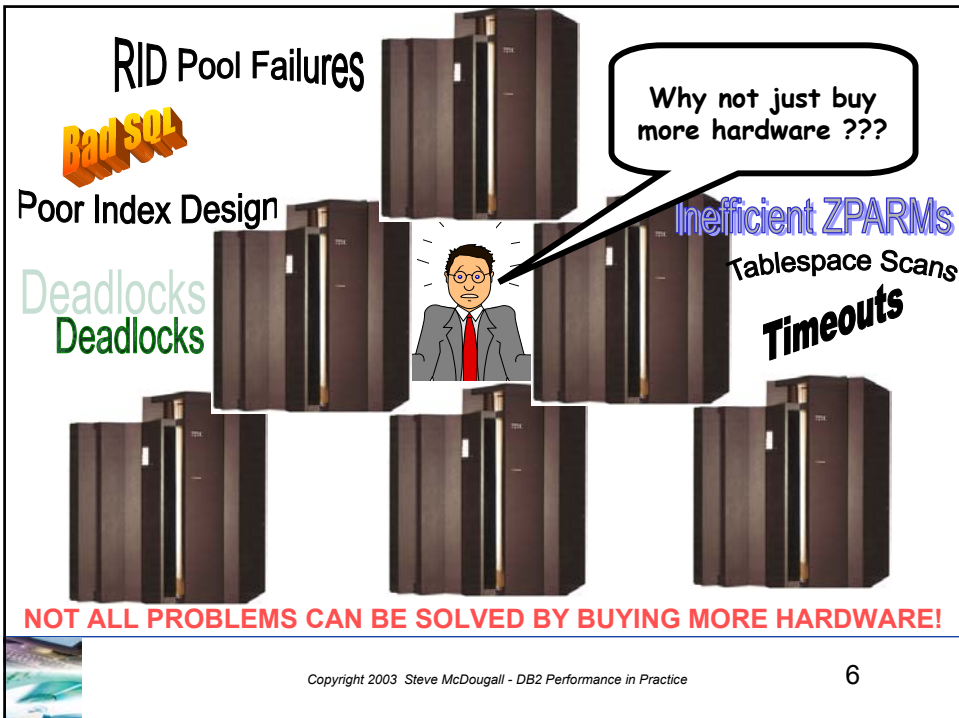
4

Consequences for DB2

- Added functionality
 - stored procedures, UDF, UDT, LOBs, BLOBs, CLOBs, Extenders
- Connectivity
 - ICLI, DB2 Connect, TCP/IP
- High availability
 - Web Sphere, e-business
- More data
 - Terabyte DBs not unusual
- Related systems
 - CICS, MQ-Series
 - Oracle, SQL-Server
- New applications coming
 - Federated databases
 - XML, XQUERY



**Increased demands on DB2
and the DBA!**



RID Pool Failures

Bad SQL

Poor Index Design

Deadlocks

Deadlocks

Inefficient ZPARMs

Tablespace Scans

Timeouts


Why not just buy more hardware ???

NOT ALL PROBLEMS CAN BE SOLVED BY BUYING MORE HARDWARE!



The Tuning Challenge

- More hardware may **mask** the problem without **removing** the cause
- Pilots **can fly** Jumbo Jets, DBAs **can tune** DB2 (and often related systems too!)
- Know how to identify the problems
- Know how to fix them
- Where do we begin?



Copyright 2003 Steve McDougall - DB2 Performance in Practice

7

Traditional Memory Resource Tuning

- Monitor the system
- Size the various pools (not just the BPs!) and set the ZPARMs
- Tune for the observed workloads, perhaps two sets of ZPARMs
- Monitor the system
- When an exception occurs, fix it
- Monitor the system --- etc.



TIMES THE NUMBER OF DB2 SYSTEMS!

Can we automate this somehow?



Copyright 2003 Steve McDougall - DB2 Performance in Practice

8


Tuning Memory Resources with BMC's Pool Advisor (*Cool Pool Tool* !)

- Statistics collected every five minutes
- Analysis performed every five minutes
 - based on a fifteen minute rolling average
- Recommends actions for the various Pools
 - BP *and* EDM, DSC, RID, Sort, GBP
- Text Advisor explains recommendations
- Recommendations can often be carried out *automatically*, if we wish
- Operates within parameters that *we* set
 - MIN, MAX, INC, DEC



Modes of Operation

- Increase and decrease pool sizes for best performance within overall storage usage considering MVS paging rates
 - **NORMAL mode** - all increase recommendations honored
 - **BALANCE mode** -equal trade-offs between increase and decrease recommendations
 - **DECREASE mode** - decrease recommendations honored (virtual storage constraints exist)
 - **FORCE mode** - attempts to reduce storage for non-critical components (virtual storage constraint is serious)



PMDEQRPN/P Pool Advisor Report Viewer LINE 1 OF 4
 Command =====> Scroll =====> CSR_

BMCSftwr.PMDMAIN -- DB2 POOLS STATUS MONITOR -- 02/21 18:12:30
 --< [HELP] >----- [WARNING] -----< [CONFIG] >-----
 Actions: T - Text A - Analysis H - History C - Configuration advisor

| DB2 | Rel | Health | Region | PageRt | BPGpRt | Efficiency | | | | | |
|--------|-----|--------|--------|--------|--------|------------|------|------|------|------|------|
| | | | | | | BP | DSC | EDM | RID | SORT | GBP |
| + DBM1 | 6.1 | GOOD | 158M | 0 | 0 | 100% | 100% | 100% | 100% | 100% | 100% |
| + DB2G | 6.1 | GOOD | 619M | 42 | 58 | 100% | 100% | 100% | 100% | 100% | N/A% |
| + DB2H | 6.1 | GOOD | 358M | 7 | 0 | 100% | 100% | 100% | 100% | 100% | N/A% |
| + DB2L | 6.1 | GOOD | 540M | 0 | 4 | 100% | 100% | 100% | 100% | 100% | N/A% |

18:06:02 I PMDZSYSA - ***** Background advisor cycle starting *****
 18:06:02 I PMDZSYSA - VS PSR: 20030221 18:06:01 R PMDZVSC0 DB2G 2
 18:06:02 I PMDZSYSA - VS AVAIL: 20030221 18:06:01 R PMDZVSC0 DB2G 0
 18:06:02 I PMDZSYSA - DB2G - VSC Storage mode changed from NORMAL to DECREASE
 18:06:02 I PMDZSYSA - DB2G - PSR_STATE = 2; AVL_STATE = 0; CUR_VS = 619MB; MAX_V

18:31:02 I PMDZSYSA - ***** Background advisor cycle starting *****
 18:31:02 I PMDZSYSA - VS PSR: 20030221 18:31:02 R PMDZVSC0 DB2G 1
 18:31:02 I PMDZSYSA - VS AVAIL: 20030221 18:31:02 R PMDZVSC0 DB2G 0
 18:31:02 I PMDZSYSA - DB2G - VSC Storage mode changed from DECREASE to BALANCE
 18:31:02 I PMDZSYSA - DB2G - PSR_STATE = 1; AVL_STATE = 0; CUR_VS = 620MB; MAX_V

18:36:04 I PMDZSYSA - ***** Background advisor cycle starting *****
 18:36:04 I PMDZSYSA - VS PSR: 20030221 18:36:04 R PMDZVSC0 DB2G 0
 18:36:04 I PMDZSYSA - VS AVAIL: 20030221 18:36:04 R PMDZVSC0 DB2G 0
 18:36:04 I PMDZSYSA - DB2G - VSC Storage mode changed from BALANCE to NORMAL
 18:36:04 I PMDZSYSA - DB2G - PSR_STATE = 0; AVL_STATE = 0; CUR_VS = 620MB; MAX_V

Copyright 2003 Steve McDougall - DB2 Performance in Practice

11

DB2 Pools Status Monitor

PMDEQRPN/P Pool Advisor Report Viewer LINE 1 OF 12
 Command =====> Scroll =====> CSR_

BMCSftwr.PMDMAIN -- DB2 POOLS STATUS MONITOR -- 08/05 15:39:09
 --< [HELP] >----- [ALERT] -----< [CONFIG] >-----
 Actions: T - Text A - Analysis H - History C - Configuration advisor

| DB2 | Rel | Health | Region | PageRt | BPGpRt | Efficiency | | | | | |
|--------|-----|--------|--------|--------|--------|------------|------|------|------|------|------|
| | | | | | | BP | DSC | EDM | RID | SORT | GBP |
| + DBF1 | 6.1 | GOOD | 464M | 0 | 14 | 100% | 100% | 100% | 100% | 100% | 100% |
| + DBF2 | 6.1 | GOOD | 150M | 0 | 0 | 100% | 100% | 100% | 100% | 100% | 100% |
| + DBM1 | 6.1 | GOOD | 246M | 0 | 948 | 91% | 100% | 100% | 100% | 100% | 100% |
| + DBM2 | 6.1 | GOOD | 157M | 0 | 0 | 100% | 100% | 100% | 100% | 100% | 100% |
| + DB2G | 6.1 | GOOD | 741M | 0 | 1059 | 97% | 100% | 100% | 100% | 100% | N/A% |
| + DB2H | 6.1 | GOOD | 403M | 0 | 55 | 100% | 100% | 100% | 100% | 100% | N/A% |
| + DB2I | 6.1 | GOOD | 195M | 0 | 0 | 100% | 100% | 100% | 100% | 100% | N/A% |
| + DB2L | 6.1 | GOOD | 711M | 1 | 665 | 99% | 100% | 100% | 100% | 100% | N/A% |
| + DB2Q | 6.1 | GOOD | 391M | 0 | 5 | 100% | 100% | 100% | 100% | 100% | N/A% |
| + DB2U | 6.1 | GOOD | 444M | 2 | 20 | 100% | 100% | 100% | 100% | 100% | N/A% |
| + DB2W | 6.1 | GOOD | 359M | 0 | 2 | 100% | 100% | 100% | 100% | 100% | N/A% |
| + DB2Z | 6.1 | GOOD | 597M | 1 | 410 | 99% | 100% | 100% | 100% | 100% | N/A% |

Copyright 2003 Steve McDougall - DB2 Performance in Practice

12

Dynamic Buffer Pool Tuning

```
DOMEADV0/P                               Advisor Response                               LINE 1 OF 26
Command =====> ----- Scroll =====> CSR_

ADVISOR:PMDZBPM1                          -- ----- 08/05/02 15:42:01
                                   Buffer Pools Status Monitor Analysis

DB2:
Buffer Pool: DBM1
Interval start: 2002-08-05-15.25.59
Current duration: 15 minutes.

-----
The efficiency rating for this buffer pool is 62%
For most busy DB2 systems, this number should be quite high. The efficiency
for this interval was poor.

Based on a recent analysis of all relevant performance indicators, an INCREASE
in the size of this buffer pool should be considered.

We recommend NO CHANGE to the Sequential Steal Threshold.
We recommend the Deferred Write Threshold be INCREASED by 10%.
No adverse performance event counts were detected.
To see additional information click here.
To see the list of Action Items, click here.
-----
```



Dynamic Buffer Pool Tuning

```
DOMEADV0/P                               Advisor Response                               LINE 1 OF 19
Command =====> ----- Scroll =====> CSR_

ADVISOR:PMDZSYSYSA                      -- ----- 08/05/02 15:42:38
                                   Pending Action List

Following is a list of all the pending action items for the following DB2 and
interval, with the highest priority items listed first:

DB2:
Interval start: DBM1
Current duration: 2002-08-05-15.25.59
Storage Mode: NORMAL

-----
Status: Auto      Priority: ?  Advisor: BPM  Recmd: INCREASE BP5 100 PGS
Action: RULE      EXEC
Status: Auto      Priority: ?  Advisor: BPM  Recmd: SET BP5 DWT/VDWT 50%
Action: RULE      EXEC
-----
```



Dynamic Buffer Pool Tuning

```
DOMEADVO/P                               Advisor Response                               LINE 1 OF 20
Command ==>> ----- Scroll ==> CSR_

ADVISOR:PMDZRULE      -- ----- --      08/05/02 15:43:11
                               Rule Details


Rule Identifier:  BR2           Rule Version:  2.0           DB2:  DBM1
Rule Description:
A significant IO rate,
Sequential Access is low,
Virtual Pool Residency Time is less than 30 seconds and
System Hit Ratio is less than 90%.

Execution Mode:  Automatic click here to set manual mode.
Times Triggered: 477         click here to reset this counter.
Times Executed:  4          click here to reset this counter.

-----
```



Buffer Pool Configuration



Strategy: Place page-sets with similar characteristics in the same bufferpools.



Configuration Advisor

```

PMDEBPCA/P          Buffer Pool Configuration Summary          LINE 1 OF 14
Command ==>>>          Scroll ==>> CSR_

DB2: DBS1 Rel: 6.1 Data sharing: YES Configuration: 0
Compatibility: 98

Actions: < ANALYZE >

Type one or more of the following action codes. Then press Enter.
S - Details/set attributes L - Lock pool U - Unlock pool
P - List page sets M - Lock page sets U - Unlock page sets

--Attributes-- --UPool-- --Hpool-- Thresholds--
PoolID T P S A U Z PgSets Cmp Size Typ Size C/O Seq DWT VDW PSA Chg
-----
BP0 S M L L L S 203 91 2500 P 0 Y 10 20 5 LRU Y
BP1 T M H L L L 13 85 2500 D 0 Y 50 20 5 LRU Y
BP2 T M L L L S 6244 96 100000 D 0 Y 10 20 5 LRU Y
BP3 I M L L L S 28468 99 80000 D 0 Y 10 20 5 LRU Y

```

- We can examine the current configuration and change it
- Pool Advisor generates the necessary commands to move objects to pools in the new configuration



EDM Pool

```

PMDEQRPM/P          Pool Advisor Report Viewer          LINE 1 OF 13
Command ==>>>          Scroll ==>> CSR_

BMCSFtwr.PMDEDMP1 -- EDM POOL STATUS MONITOR -- 08/11 07:41:55
--< HELP >-----< NOTES >-----< CONFIG >-----
Actions: T - Text A - Analysis H - History

DB2: DBS1 6.1 Interval start: 2003-08-11-07.30.59 Duration: 00:10:56

+ EDM pool efficiency : 100%
EDM pool size (pages) : 30000 DSC dataspace size . . . : 14623
EDM pool full failures: 0 DSC dataspace full failures: 0

Number Number Hit Request Loads Pages Pages
requests loads ratio per min per min alloc alloc
-----
DBD 114914 0 100% 1915 0 10138
CT 0 0 N/P% 0 0 627
PT 2344 0 100% 39 0 9076
DynC 108080 208 100% 1801 3 0 14319
SKCT
SKPT 635

```



Dynamic Statement Cache

```

PMDEQRPN/P                               Pool Advisor Report Viewer                               LINE 1 OF 6
Command ==>>>                               Scroll ==> CSR_

BMCSftwr.PMDDSC1  --  DYNAMIC STATEMENT CACHE  --  08/11 07:39:56
---< HELP >-----< NOTES >-----< CONFIG >-----
  Actions: T - Text    A - Analysis    H - History
           P - Plan    R - Program    N - ConnID    C - CorRID

DB2: DBS1 6.1      Interval start: 2003-08-10-23.59.59  Duration: 07:39:57

+ Prepares : 907 Avg Prepare CPU: 00.011616 Avg Elapsed: 00.031322
  Duplicates: 200 Avg Duplicate CPU: 00.011957 Avg Elapsed: 00.027638
  Recalled : 534744 Avg Recalled CPU: 00.000530 Avg Elapsed: 00.000932
  
```



RID Pool

```

PMDEQRPN/P                               Pool Advisor Report Viewer                               LINE 1 OF 3
Command ==>>>                               Scroll ==> CSR_

BMCSftwr.PMDRIDP1  --  RID POOL STATUS MONITOR  --  08/11 07:43:56
---< HELP >-----< NOTES >-----< CONFIG >-----
  Actions: T - Text    A - Analysis    H - History

DB2: DBS1 6.1      Interval start: 2003-08-11-07.30.59  Duration: 00:12:57

      RID Pool      RID Pool
      Size (Blks)    HWM (Blks)
      (bytes)        (bytes)
      Times          Times
      Used           Failed
      Aborts -      Aborts -
      No Storage    Int. Limit
      -----
+ 100%      3125      4309      64029      1      0      1
      51M      71M
  
```

- Pool Advisor reacts on two critical values
 - High Water Mark
 - Aborts caused by no storage



Sort Pool

```

PMDEQRPN/P                               Pool Advisor Report Viewer                               LINE 1 OF 17
Command =====>                               Scroll ==> CSR_
BMCSftwr.PMDSRTP1  --  SORT POOL STATUS MONITOR  --  More: +
--< [HELP] >-----< [NOTES] >-----< [CONFIG] >-----
  Actions: T - Text      A - Analysis      H - History

  DB2: DBS1 6.1      Interval start: 2003-08-11-07:30:59  Duration: 00:14:02

+ Sort pool merge/pass efficiency          100%
Sort work file data base name             DBS1WRK

Maximum sort pool size (per thread)       29M

Number of sort merge passes . . . . . 283
Number of workfiles requested for merge passes . . . . . 573
Average number of workfiles requested per merge pass . . 2

Number of inefficient sort merge passes . . . . . 0
Number of workfiles denied for merge passes . . . . . 0
Average number of workfiles denied per merge pass . . . . N/P

Maximum number of workfiles concurrently allocated . . . 15
Number of times the work buffer pool did not have enough
  
```



Copyright 2003 Steve McDougall - DB2 Performance in Practice

21

Group Buffer Pools

```

PMDEQRPN/P                               Pool Advisor Report Viewer                               LINE 1 OF 14
Command =====>                               Scroll ==> CSR_
BMCSftwr.PMDGBP1  --  GROUP BPOL STATUS MONITOR  --  08/11 07:45:43
--< [HELP] >-----< [NOTES] >-----< [CONFIG] >-----
  Actions: T - Text      A - Analysis      H - History
           G - Group Details      M - Member Details

  Interval start: 07:30:59      Duration: 00:14:40
  Data sharing group: DSNDBS0      Member: DBS1      SSID: DBS1

  GBP ID  Eff  Read  Writ  Size  # Alloc  # Alloc  Castouts  XI Dir
         hit% hit% hit%          DIR Entr  DAT Entr
-----
* GBP0    80% 100% 100% 15872K    626    439    139      0
* GBP2    34%  8% 100% 791040K 150721 28935 10064    0
* GBP3    31%  3% 100% 632832K 123898 28026 6834    0
  
```

- Pool Advisor will make recommendations to change
 - Directory to Data Entries Ratio
 - GBP Size
 - Checkpoint interval



Copyright 2003 Steve McDougall - DB2 Performance in Practice


22

Memory Savings

| DB2 Subsystem | Storage allocation with SAP recommendations | Possible Storage allocation with Pool Advisor |
|----------------------|---|---|
| DB2G | 536 MB | 323 MB |
| DB2H | 357 MB | * 317 MB |
| DB2L | 635 MB | * 533 MB |
| DBM1 | 201 MB | * 100 MB |
| DBM2 | 164 MB | * 100 MB |
| DB2W | 371 MB | * 200 MB |
| DBF1 | 513 MB | * 250 MB |
| DBF2 | 162 MB | * 100 MB |
| DB2Q | 423 MB | * 250 MB |
| DB2U | 451 MB | * 250 MB |
| DB2Z | 430 MB | * 250 MB |
| DBS1 | 1103 MB | ** 920 MB |
| DBS2 | 1103 MB | ** 950 MB |
| DBX1 | 536 MB | 391 MB |
| Total Storage | 6985 MB | 4934 MB |



SQL Everywhere




Static
Dynamic
Internet

ERP/CRM
SQL
IDBC

Not always able to check SQL with EXPLAIN



Gone Fishing



Why is this inadequate?



Fishing for SQL



- Traditional SQL fishing
 - Accounting report plus Explain can miss SQL which executes thousands of times - but poorly
 - Brings up too many SQLs that are OK
- Better SQL fishing
 - Tune the **heavy hitters**
 - Tune SQL which accesses the **heaviest used objects**



Tune the Heavy Hitters



- Find the big fish but consider
 - Many small fish = One big fish
- Saving 1/100's of a second on frequently executed SQL helps
 - cumulative effect
- Deals with SQL that gets through the traditional net
 - i.e. those with *low* elapsed time, *low* getpages and *low* CPU time - but *high* execution rate
- Identification problem



Using BMC's APPTUNE

ASQEQRPW/P View a Report LINE 1 OF 2
 Command =====> _____ Scroll ==> CSR_

BMCsftwr.SQMCCTS -- SUBSYSTEM ANALYSIS (DATA) -- 08/02 10:11:47
 Source : DC01-ACTIVE Intv : 07/31 00:00 - 07/31 23:59 More: >

Actions: G-AppGrps R-Programs C-ConnID T-Detail A-Averages O-Object
 P-Plans U-Users S-Stmts E-SQL Errors L-CorrID

Actions for SAP systems: V-Server N-Work Process Number I-Work Process ID

| AppGrp: | User: | ConnID: | Plan: |
|----------|-----------------|-----------|-------------|
| CorrID: | Program: | Stmt No: | |
| DB2 SSID | MVS ID | SQL Calls | Elapsed |
| + | DB2G (6.1) I211 | 4595K | 32:23.54293 |
| + | DB2L (6.1) I211 | 11816K | 27:23.39748 |

| SQL Calls | Elapsed | Total IN-SQL Time | % CPU | % Getpage | + S Numb |
|-----------|-------------|-------------------|-------------|-----------|----------|
| 4595K | 32:23.54293 | 54.2% | 08:28.61695 | 32.2% | 1932327 |
| 11816K | 27:23.39748 | 45.8% | 17:49.40214 | 67.8% | 4515860 |



A stylized illustration of a screech owl. The owl has a large, round head with prominent yellow eyes and a red face. Its body is covered in brown and white mottled feathers. It has large, dark wings and a long, pointed beak. The owl is shown in a frontal, slightly hunched position.

- At first ignored (low elapsed / IN-SQL Time)
- Executes millions of times
- Worth tuning

| | Before | After |
|-----------------|---------|---------|
| Avg IN-SQL | 0,074 | 0,0307 |
| per row fetched | 0,00045 | 0,00023 |
| Avg CPU | 0,0378 | 0,0227 |
| per row fetched | 0,00023 | 0,00017 |
| Avg GETPAGES | 965 | 189 |
| per row fetched | 6 | 1 |



A cartoon illustration of an owl, facing forward. It has large, yellow eyes with black pupils, a small beak, and large, brown feet. Its body is brown with darker brown spots on its wings and back.

- Find the top N busiest objects in the system
 - number of getpages in a given time period
- Look at SQLs which access these objects
 - elapsed time, getpages, CPU
- Fewer SQLs to look at
 - better chance of finding tuning candidates



Heaviest Used Objects



| Database | Page Set | Type | BPool | BPool HRatio | Number | % | Getpage | Tot Time | Avg Time | Number |
|----------|----------|------|-------|--------------|--------|-------|---------|----------|----------|--------|
| + POOL17 | SPOOL17 | T | BP2 | 99.9 % | 989095 | 21.9% | 14.4377 | 00.0000 | | 161 |
| + STAB11 | SSTAB11 | T | BP2 | 99.8 % | 753841 | 16.7% | 17.7640 | 00.0000 | | 418 |
| + STAB07 | SSTAB07 | T | BP2 | 99.9 % | 475402 | 10.5% | 09.6107 | 00.0000 | | 131 |
| + STAB07 | HRP11RIR | T | BP3 | 100 % | 144154 | 3.2% | 02.2857 | 00.0000 | | 47 |
| + POOL00 | SPOOL00 | T | BP2 | 98.6 % | 142922 | 3.2% | 04.7682 | 00.0000 | | 11 |

| Creator | Object | Type | BPool | BPool HRatio | Number | % | Getpage | Tot Time | Avg Ti |
|---------|---------|------|-------|--------------|--------|-------|---------|----------|--------|
| S SAPR3 | HRP1001 | T | BP2 | 99.9 % | 464017 | 97.6% | 09.2189 | 00.00 | |
| + SAPR3 | HRP1000 | T | BP2 | 99.4 % | 11379 | 2.4% | 00.3709 | 00.00 | |
| + SAPR3 | HRP1007 | T | BP2 | 83.3 % | 6 | 0.0% | 00.0208 | 00.00 | |

| Program | Stmt Type | Call Type | Sect No. | BPool HRatio | Number | % | Getpage | Tot Time | Avg Time | Number |
|------------|-----------|-----------|----------|--------------|--------|-------|---------|----------|----------|--------|
| + U8PW0E02 | DYNAMIC | CURSOR | 53 | 77.7 % | 143880 | 22.5% | 04,1149 | 00,0000 | | 969 |
| + U8PW0E02 | DYNAMIC | CURSOR | 55 | 89.4 % | 128585 | 20.1% | 02,6156 | 00,0000 | | 845 |
| + U8PW0E02 | DYNAMIC | CURSOR | 57 | 94.7 % | 123051 | 19.2% | 01,4923 | 00,0000 | | 73 |



Copyright 2003 Steve McDougall - DB2 Performance in Practice

31

Explain

DOMEADVO/P

Command ==>>>

Advisor Response

Scroll ==> CSR_

LINE 109 OF 296

08/02/02 11:13:28

ADVISOR:EXPLAINX

STEP METHOD

ACCESS TYPE

ACCESS NAME

TABLE NAME

1 0-FIRST TABLE

N -INDEX SCAN

HRP1001B4

HRP1001

TAB NO.

IX ONLY

MTCH COLS

LOCK MODE

MIAP

PRE FTCH

SORTN UJOG

SORTC UJOG

COL EVAL

ACCESS DEGREE

1 1

N

5

N

0

NNNN

NNNN

0

SQL TEXT FOR STATEMENT: 0

SELECT *

FROM "HRP1001"

WHERE "MANDT" = ?

AND "PLVAR" = ?

AND "OTJID" IN(?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?)

AND "SUBTY" = ?

AND "ISTAT" IN(?, ?)

AND "BEGDA" <= ?

AND "ENDDA" >= ?

AND "SCLAS" = ?

FOR FETCH ONLY

WITH UR

Easy to read !

Easy to read !



Copyright 2003 Steve McDougall - DB2 Performance in Practice

32

Statistics Recommendations

High2Key

Sample

BMC184240S-RUNSTATS have not been run on this object. The Optimizer will use defaults to determine the access paths.

BMC184264W-The HIGH2KEY and LOW2KEY values are not available. This can adversely impact performance. Run RUNSTATS and rebind the object.


BMC184280W-This index is the Clustering index but it is no longer clustered. Consider running a REORG and RUNSTATS, followed by a REBIND.

BMC184290W-The average Index Leaf Distribution exceeds the installation standard. The Leaf Distribution/100 is ...

BMC184260W-The ratio of Full Key Cardinality to First Key Cardinality is almost equal. This index is not a good candidate for the Clustering index. Consider changing the CLUSTERING option to another index that will provide some degree of grouping.



Access Path Recommendations



Sample

BMC184594I-DB2 will be using a tablespace scan to satisfy this query which has been designated for ONLINE Only usage. The table does not currently have an index. The number of rows in the table is

BMC184596W-DB2 will be using multiple indexes to satisfy this query which has been designated for BATCH Only usage. The use of multiple indexes to satisfy this query may cause excessive overhead. The number of rows in the table is ...

BMC184600S-DB2 will be using a Hybrid Join to satisfy this query which has been designated for ONLINE Only usage. Hybrid Joins are complex to resolve and may cause excessive overhead. The number of rows in the table is ...



SQL Recommendations



Sample

BMC184000S-SELECT * on TABLE is not recommended since the answer set may change in the future.

BMC184326W-Tip: Consider using OPTIMIZE FOR N ROWS if the desired answer set is small and you want to avoid the cost of List Prefetch.

BMC184044I-A NOT BETWEEN Value1 and Value2 predicate is Non-indexable. Consider redesigning the predicate to use OR instead of NOT BETWEEN.

BMC184314W-The ORDER BY contains a column(s) which is not indexed. Drop the column from the ORDER BY or add it to the index to eliminate sorting.



Locking Recommendations



Sample

BMC184186W-The object was bound with the ISOLATION option of Repeatable Read. Review with DBA on this usage.

BMC184376W-An INTENT EXCLUSIVE lock will be acquired on a table accessed via a Table Space Scan. Consider using LOCK TABLE to minimize locks and DB2/IRLM Lock management overhead. This can reduce CPU usage and improve elapse time.

BMC184582W-To improve concurrency and throughput consider modifying the SQL text to use WITH UR or bind the object with a Cursor Isolation=UR (Uncommitted Read).



Index Design Issues

- Performance indexes
 - to avoid TS scans on large tables
 - to avoid sorts
 - index only access (not with VARCHAR till V8)
 - help joins
- Common Problems
 - Clustering on primary key
 - Low cardinality column first
 - Column ordering
 - unused and redundant indexes



Index Design Example

```
SELECT * FROM "TEVEN"  
WHERE "MANDT" = ?  
AND "PERNR" = ?  
AND "LDATE" >= ?  
AND "LDATE" <= ?  
AND "STOKZ" = ?  
ORDER BY "MANDT" , "PERNR" , "LDATE" , "LTIME" , "ERDAT" , "ERTIM" ,  
"PDSNR" FOR FETCH ONLY WITH UR
```


- | | | | | |
|--------------|---|------------|---|--|
| • INDEX COLS | ➔ | • ADD COLS | ➔ | • Make IX clustering |
| – MANDT | | – ERDAT | | • Original clustering was on the primary key |
| – PERNR | | – ERTIM | | |
| – LDATE | | – PDSNR | | |
| – LTIME | | | | |

Elapsed 0,269

Elapsed 0,039



Index Advisor Component



- *NEW*
- Tables by getpage volume
- Table CRUD matrix
- Statement column usage
- What if analysis



DB2 Performance in Practice Summary

- Increasingly complex applications are driving tuning requirements
 - more hardware is not always the best answer
- Automation of DB2 memory resource tuning
 - smart memory tuning - now
- Superior strategies to find SQL to tune
 - heavy hitters
 - heaviest used objects
- Recommendations to improve SQL / access path
- Index design issues and common problems





Even Jumbo Jets have auto-pilots!



41

DB2 Performance in Practice

Steve McDougall
Independent Database Consultant
SteveMcDou@aol.com

