sqladria

SQLADRIA SEMINAR

Zagreb, 23 September 2011

Sponsored by

bmcsoftware

SQLAdria Seminar

10.00 Opening

10.15 Rev up your SQL Applications with Enhanced Statistics

Ramon Menendez
BMC Software

11.15 Coffee Break

11.30 Squeeze the Most out of your Indexes in DB2 9!

Peter Plevka
BMC Software

12.30 Coffee Break

12.45 DSNZPARMs for DBA's

Ramon Menendez
BMC Software

13.45 Lunch

14.45 A reconstruction of 'AS IS' data model

Mateja Jankovič
Informatika d.d.

15.30 Process of Change Migration to Production Environment

Tomaž Dobnikar
NLB d.d.

16.00 Open forum

Sponsored by:

Zagreb, 23rd September 2011
Hotel Palace

Participants of the SQLAdria Seminar
Zagreb 23rd September 2011

	Name	Surname	Company	E-mail
1	Željko	Blagajić	HZMO	zeljko.blagajic@mirovinsko.hr
2	Anita	Lukačić	HZMO	anita.lukacic@mirovinsko.hr
3	Iva	Mesarić	HZMO	iva.mesaric@mirovinsko.hr
4	Goran	Petrinjac	HZMO	goran.petrinjac@mirovinsko.hr
5	Jelena	Šimić	HZMO	jelena.simic@mirovinsko.hr
6	Matija	Plavša	Informatika d.d.	matija.plavsa@informatika.si
7	Mateja	Janković	Informatika d.d.	mateja.jankovic@informatika.si
8	Boris	Brence	ZZZS Slovenija	boris.brence@zzzs.si
9	Sonja	Klančnik	ZZZS Slovenija	sonja.klancnik@zzzs.si
10	Tatjana	Šumec	ZZZS Slovenija	tatjana.sumec@zzzs.si
11	Lidija	Penezić	Apis IT	lidija.penezic@apis-it.hr
12	Dragica	Šmintić	Apis IT	dragica.smintic@apis-it.hr
13	Gordana	Kablar	Apis IT	gordana.kablar@apis-it.hr
14	Roman	Jurić	Apis IT	roman.juric@apis-it.hr
15	Goran	Bavčar	NLB d.d.	goran.bavcar@nlb.si
16	Tomaž	Dobnikar	NLB d.d.	tomaz.dobnikar@nlb.si
17	Peter	Plevka	BMC Software	peter_plevka@bmc.com
18	Ramon	Menendez	BMC Software	ramon_menendez@bmc.com
19	Marija	Musa	FINA	marija.musa@fina.hr
20	Ivan	Pinoza	FINA	ivan.pinoza@fina.hr
21	Branimir	Putniković	FINA	branimir.putnikovic@fina.hr
22	Darina	Giačić-Mitić	FINA	darina.giacic@fina.hr
23	Zvjezdana	Horvat	FINA	zvjezdana.horvat@fina.hr
24	Mirna	Kos	Zagrebačka banka d.d.	mirna.kos@unicreditgroup.zaba.hr
25	Ana	Ramljak	Zagrebačka banka d.d.	ana.ramljak@unicreditgroup.zaba.hr
26	Ivana	Žunac	Zagrebačka banka d.d.	ivana.zunac@unicreditgroup.zaba.hr
27	Sven	Lovrenčić	Zagrebačka banka d.d.	sven.lovrencic@unicreditgroup.zaba.hr
28	Nenad	Vidak	Zagrebačka banka d.d.	nenad.vidak@unicreditgroup.zaba.hr
29	Ivan	Jurinjak	Zagrebačka banka d.d.	ivan.jurinjak@unicreditgroup.zaba.hr
30	Igor	Mihevc	IMAVES d.o.o.	igor.mihevc@imaves.hr
31	Vedran	Vesel	IMAVES d.o.o.	vedran.vesel@imaves.hr
32	Milan	Vesel	IMAVES d.o.o.	milan.vesel@imaves.hr
33	Željen	Stanić	CA	zeljen.stanic@ca.com
34	Olivera	Stanić	CTK SQLAdria	ostanic@ctk-rijeka.hr

SQLADRIA SEMINAR
Zagreb, 23 September 2011

Rev up your SQL Applications with Enhanced Statistics

Ramon Menendez
BMC Software

Sponsored by

Rev Up Your SQL Application with Enhanced Statistics!

Ramón Menéndez, BMC Software
Zagreb, 23rd September 2011

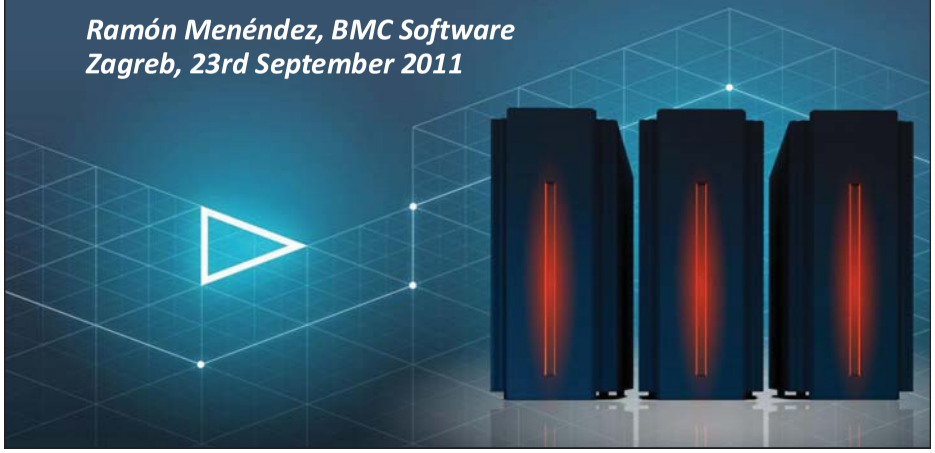

What Does "Enhanced" Mean?

en·hance [en-hans, -hahns]
–verb (used with object),-hanced, -hanc·ing.
1. to raise to a higher degree; intensify; magnify: *The candlelight enhanced her beauty.*
2. to raise the value or price of: *Rarity enhances the worth of old coins.*

© Copyright 9/15/2011 BMC Software, Inc

Introduction

- Changes made to optimizer statistics in DB2 9
- Columns in SYSINDEX: CLUSTERRATIOF and DATAREPEATFACTORF
- Enhancing these index statistics is optional, but recommended
- Point is to improve performance of key range extractions in SQL
  - No help for single row retrievals

© Copyright 9/15/2011 BMC Software, Inc

DB2 Index Structure

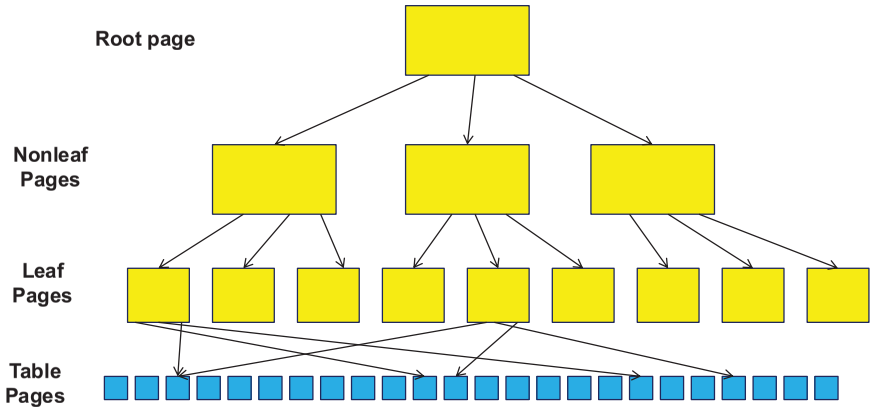

© Copyright 9/15/2011 BMC Software, Inc

Clustered Retrieval

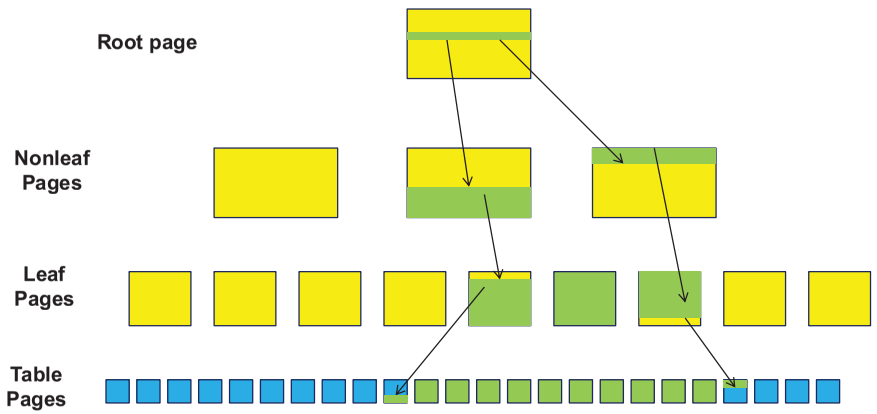

© Copyright 9/15/2011 BMC Software, Inc

Not So Clustered Retrieval

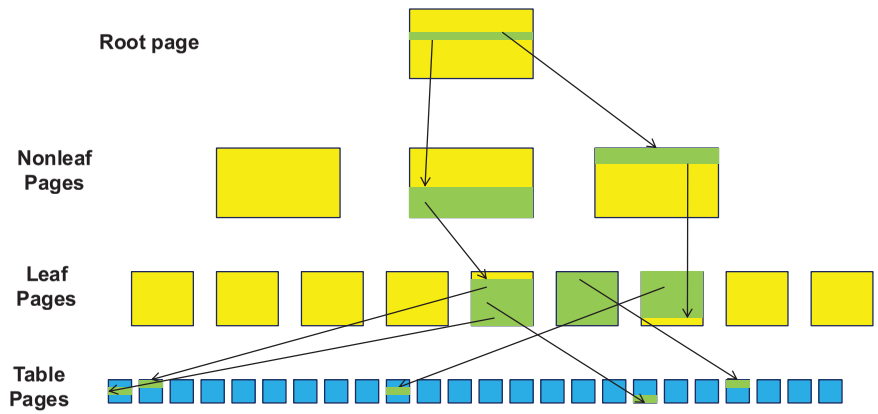

© Copyright 9/15/2011 BMC Software, Inc

More Clustering Examples

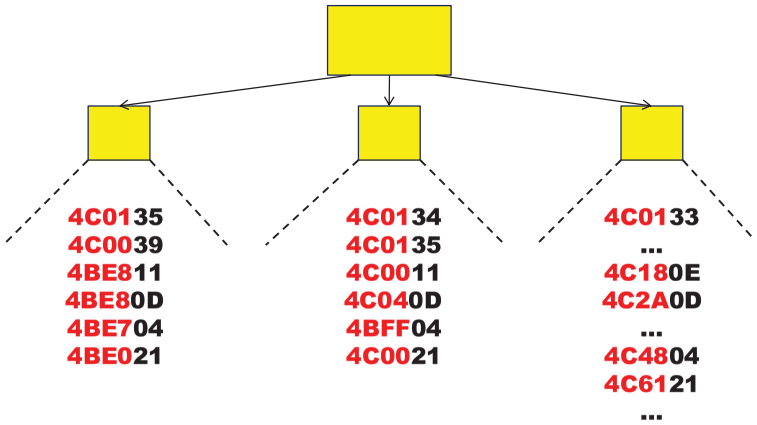

...page numbers descending

...up and down in a narrow range

...gradual drift in one direction

© Copyright 9/15/2011 BMC Software, Inc

What is DATAREPEATFACTORF?

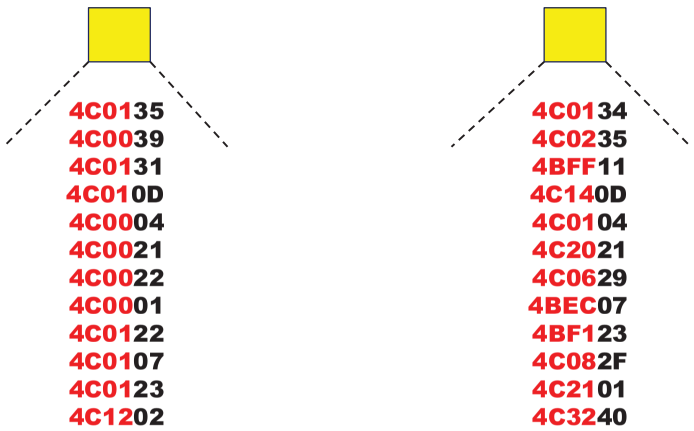

12 rids...3 page numbers

12 rids...11 page numbers

© Copyright 9/15/2011 BMC Software, Inc

STATCLUS

- Line 9 on DSNTIP6 panel ("STATISTICS CLUSTERING") in DB2 9
- "ENHANCED" is default
- You can specify STANDARD
- STATCLUS in DSN6SPRM macro
- Affects RUNSTATS operation, not the optimizer
- Moved to the list of opaque parameters: "Subsystem parameters that are not on installation panels" in DB2 10 and deprecated

© Copyright 9/15/2011 BMC Software, Inc

STATCLUS Migration

- What happens to existing stats and packages when STATCLUS changes from STANDARD to ENHANCED?

  - Nothing! Must run RUNSTATS and BIND
  - Dynamic SQL affected as soon as RUNSTATS runs

© Copyright 9/15/2011 BMC Software, Inc

CLUSTERRATIO Calculation Overview

- Based on analysis of tablespace page referred to by each RID (no need to access tablespace!)
- Consider index as list of tablespace page numbers
- Want one number that gives average expectation that synchronous I/O will be avoided

© Copyright 9/15/2011 BMC Software, Inc

CLUSTERRATIO Calculation Improvements

- Looks at every RID
- Considers declining values
- Considers size of prefetch window
- Point is to improve optimization of index scans

© Copyright 9/15/2011 BMC Software, Inc

CLUSTERRATIO Calculation

- Assume 32K tablespace
- Prefetch size is 4 pages (for buffer pool between 3.2M and 160M)
- Sliding window up to 64 pages

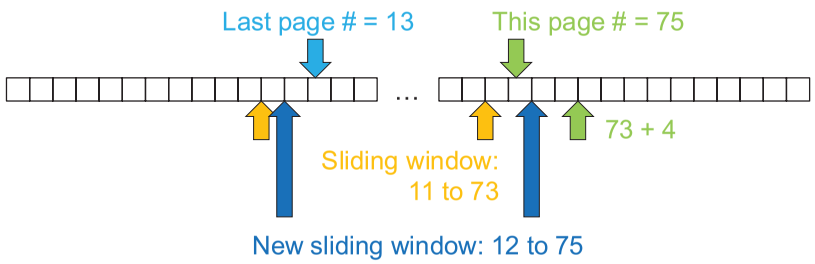

© Copyright 9/15/2011 BMC Software, Inc

CLUSTERRATIO Calculation Details

- Example calculation, based on prefetch amount of 8
  - Could be 16K tablespace in buffer pool with VPSIZE=8000
- Page numbers are decimal

Page #	13	35	40	…	96	29	104	3	7
Window after	13-13	35-35	35-40		35-96	33-96	41-104	3-3	3-7
# clustered	1	1	2		9	10	11	11	12
#unclustered	0	1	1		1	1	1	2	2

  - CLUSTERRATIOF = 12 / (12 + 2) = .8571

© Copyright 9/15/2011 BMC Software, Inc

Does This Work?

- All rids considered
  - More accurately simulates access to table pages
- Descending rids "count"
  - Supports dynamic prefetch as well as list prefetch
- "Oscillating" rids in small range
  - Within prefetch amount of range "counts"
- Gradual trend in page numbers
  - Will "move" sliding window

© Copyright 9/15/2011 BMC Software, Inc

What About Partitioned Spaces?

- Each partition of index calculated separately
- Results averaged for index, weighted by table partition cardinality
- NPI's?
  - They can include rids from many table partitions
  - "Switch" in page number from one partition to another is always counted as unclustered
  - We know the second page won't be part of the same prefetch!

© Copyright 9/15/2011 BMC Software, Inc

DATAREPEATFACTORF Calculation Overview

- Adds consideration of row density in tablespace
- For clustered data, want number to reflect relative probability of buffer hits after asynchronous I/O
- It is an absolute value
- Lower is better

© Copyright 9/15/2011 BMC Software, Inc

DATAREPEATFACTORF Calculation

- Assume 32K tablespace
- Prefetch size is 4 pages (for buffer pool between 3.2M and 160M)

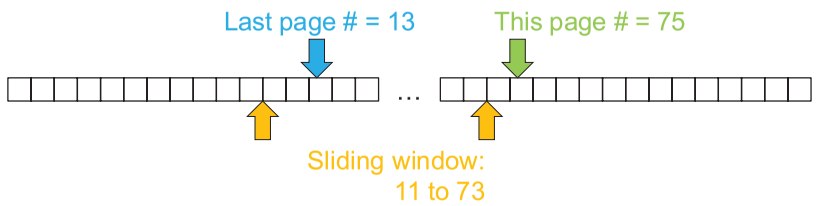

© Copyright 9/15/2011 BMC Software, Inc

DATAREPEATFACTORF Calculation Details

- Example calculation, based on prefetch amount of 8
  - Could be 16K tablespace in buffer pool with VPSIZE=8000
- Page numbers are decimal

Page #	13	35	35	2	10	18	26	3	15
Window after	13-13	35-35	35-35	2-2	2-10	2-18	2-26	2-26	2-26
DRF	1	2	2	3	4	5	6	6	6

- DATAREPEATFACTORF = 6

© Copyright 9/15/2011 BMC Software, Inc

When Do Enhanced Statistics Matter?

ACCESSTYPE	MATCHCOLS	ACCESSNAME	INDEXONLY
I	1	ITEMSHIP	N

- Access must be to the table, through a scan of the index
- Most of the time, the significant CLUSTERRATIOF threshold is 80%
- Change from 1% to 79% doesn't matter
- Change from 79% to 81% does matter

© Copyright 9/15/2011 BMC Software, Inc

Hybrid Join Example

METHOD	ACCESSTYPE	ACCESSNAME	SORTN_JOIN	PREFETCH
4	I	CUSTIX	N	L

- METHOD value of 4 indicates a hybrid join
- List prefetch usually selected to access the table
- SORTN_JOIN value of 'N' indicates the rid list is not sorted…CLUSTERRATIOF for CUSTIX > 80%
- Value is 'Y' if CLUSTERRATIOF <= 80%

© Copyright 9/15/2011 BMC Software, Inc

DATAREPEATFACTORF Example

METHOD	ACCESSTYPE	ACCESSNAME	INDEXONLY	PREFETCH
0	I	ORDERCUST	N	L

- Index scan to access a table
- Threshold between LIST PREFETCH /SEQ PREFETCH is at 80%
- A high DATAREPEATFACTORF (close to number of rows in table) causes threshold to move up to 90%

© Copyright 9/15/2011 BMC Software, Inc

When Will ENHANCED Help?

- Low CLUSTERRATIOF increases
  - All rids now counted
  - Descending rids
  - Page numbers going up and down within a range
  - Avoid sorts for hybrid joins, avoid prefetch in nested loop joins
- Clustered index with less dense data
  - Avoid prefetch in index scans

© Copyright 9/15/2011 BMC Software, Inc

When Will ENHANCED Hurt?

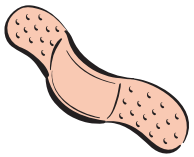

- When DB2 was doing the right thing for the wrong reasons
- Possible to decrease CLUSTERRATIOF which was artificially high
- Where CLUSTERRATIOF increases, maybe list prefetch was a good idea

© Copyright 9/15/2011 BMC Software, Inc

How To Get There?

- Do nothing
  - Accept the default
  - Wait and see
- Clone data to ENHANCED DB2?
  - Might work for limited amount of data
- RUNSTATS … REPORT YES UPDATE NONE
  - Compare reported CLUSTERRATIOF with catalog
    - Changes above or below threshold
    - EXPLAIN SQL identified as using indexes of interest
    - Change SQL if possible

© Copyright 9/15/2011 BMC Software, Inc

Alternatives

- Use Plan Stability!
- REBIND…PLANMGMT(EXTENDED)
  - Creates new "Original" package
  - If problems, REBIND…SWITCH(ORIGINAL)
- Copy SYSINDEXES to user table
- If problems, revert CLUSTERRATIOF and set DATAREPEATFACTORF to -1
  - Columns are updatable
- Use tools available

© Copyright 9/15/2011 BMC Software, Inc

Bibliography

- "DB2 9 for z/OS Performance Topics", IBM Redbook
- "DB2 9 for z/OS: Using the Utilities Suite". IBM Redbook
- "DB2 Version 9.1 for z/OS SQL Reference"
- "DB2 9 for z/OS Performance Monitoring and Tuning Guide"
- "DB2 Version 9.1 for z/OS Utility Guide and Reference"
- "DB2 Version 9.1 for z/OS Application Programming and SQL Guide"
- "DB2 Version 9.1 for z/OS Installation Guide",

© Copyright 9/15/2011 BMC Software, Inc

My own testing…

- Test table with about 1600 rows and 1600 pages
- Four indexes
  - IX1 always ascending
  - IX2 always descending
  - IX3 ascending/descending by tranches
  - IX4 random
- Reorg plus Runstats with STATCLUS=STANDARD

INDEX	CLUSTERED	CLUSTERRATIO	DATAREPEATFACTOR
IX1	Y	100	-1
IX2	N	12	-1
IX3	N	69	-1
IX4	N	10	-1

- SQL SELECT a range of 15% of the table accessing by the four indexes

© Copyright 9/15/2011 BMC Software, Inc

EXPLAIN

- With STATCLUS=STANDARD

QUERY	TIPO	MATCHING	PREFETCH	COST
NO WHERE	R	0	S	554
WHERE COL1	I	1	S	83
WHERE COL2	I	1	L	93
WHERE COL3	I	1	L	105
WHERE COL4	I	1	L	104

© Copyright 9/15/2011 BMC Software, Inc

Enhanced Statistics

- Set STATCLUS=ENHANCED
- Execute Runstats again

INDEX	CLUSTERED	CLUSTERRATIO	CLUSTERED	CLUSTERRATIO	DRF
IX1	Y	100	Y	100	1584
IX2	N	12	Y*	98*	354
IX3	N	69	Y*	99*	1119
IX4	N	10	N	28*	1584

  - Note: ROWS=1584, PAGES=1584

© Copyright 9/15/2011 BMC Software, Inc

EXPLAIN Compare

- Execute Workload again with STATCLUS=ENHANCED
- Compare

QUERY	TIPO	MATCHING	PREFETCH	COST	PREFETCH	COST
NO WHERE	R	0	S	554	S	554
WHERE COL1	I	1	S	83	S	83
WHERE COL2	I	1	L	93	S*	84*
WHERE COL3	I	1	L	105	S*	94*
WHERE COL4	I	1	L	104	L	106*

© Copyright 9/15/2011 BMC Software, Inc

Workload Compare Report Output

```
-- SQL EXPLORER 6.2.00         COPYRIGHT 1996 - 2009 BMC SOFTWARE, INC.

WORKLOAD COMPARE REPORT

STATEMENT SUMMARY IN DESCENDING COST DIFFERENCE ORDER
 PLAN      PROGRAM   STMT    STMT    COST      WEIGHTED  DIFFERENCE
 NAME      NAME      NUMBER,NUMBER   DIFF      COSTDIFF  IN
 --------  --------  -------------   --------  --------  ----------
 ACT930DM  ACT8SELE   87556, 87556       -11       -11   PATH STATS
 ACT930DM  ACT8SELE   87555, 87555        -9        -9   PATH STATS
 ACT930DM  ACT8SELE   87552, 87552         0         0   STATS
 ACT930DM  ACT8SELE   87554, 87554         0         0   STATS
 ACT930DM  ACT8SELE   87557, 87557         2         2   STATS
```

Nota: Weighted cost = (cost * executions)

© Copyright 9/15/2011 BMC Software, Inc

Differences in access path

```
SELECT   COL1, COL2, COL3, FILLER1
 FROM     ISPRXM.TEST_STATCLUS
 WHERE    COL3 BETWEEN '15' AND '30'

     QB PL M TB AC MT                          I  NEW  COMP P C MX LCK J
LABL BL NO E NO TY CL TABLE        INDEX        O  UJOG UJOG F F SQ MOD C
---- -- -- - -- -- -- ------------ ------------ -- ---- ---- - - -- --- -
XD00  1  1 0  1 I   1 STATCLUS     STATCLUS_IX3 N  NNNN NNNN*L    0  IS 0

XD01  1  1 0  1 I   1 STATCLUS     STATCLUS_IX3 N  NNNN NNNN S    0  IS 0

       SERVICE EXEC    WEIGHTED WEIGHTED DIFF
STMTNO UNITS   COUNT   COST     DIFF     IN
------ ------- ------- -------- -------- -----------------
 87556*    105       1*     105      -11 PATH STATS
 87556      94       1       94
```

© Copyright 9/15/2011 BMC Software, Inc

Additional testing…

- Table with about 1600 rows and 400 pages
- The same four indexes
  - IX1 always ascending
  - IX2 always descending
  - IX3 ascending/descending by tranches
  - IX4 Random
- SELECT SQL by range of 15% and join of both tables
- Repeat all process
  - Execute SQL Workload and Explain
  - Change ZPARM and Runstats
  - Explain
  - Workload Compare

© Copyright 9/15/2011 BMC Software, Inc

Enhanced Statistics

- Runstats with STATCLUS=ENHANCED

ÍNDICE	CLUSTERED	CLUSTERRATIO	CLUSTERED	CLUSTERRATIO	DRF
IX1	Y	100	Y	100	396
IX2	N	12	Y*	99*	114
IX3	N	69	Y*	99*	291
IX4	N	10	N	70*	1558

  - Note: ROWS=1584, PAGES=396
  - Note: IX4 benefits from higher row density

© Copyright 9/15/2011 BMC Software, Inc

Report Output

```
-- SQL EXPLORER 6.2.00       (C) COPYRIGHT 1996 - 2009 BMC SOFTWARE, INC.

WORKLOAD COMPARE REPORT

STATEMENT SUMMARY IN DESCENDING COST DIFFERENCE ORDER
 PLAN      PROGRAM   STMT   STMT    COST      WEIGHTED  DIFFERENCE
 NAME      NAME      NUMBER,NUMBER  DIFF      COSTDIFF  IN
 --------  --------  -------------  --------  --------  ----------
 ACT930DM  ACT8SELE   12096, 12096      -299      -299  PATH STATS

 EXPL SRC                                SERVICE
 LABL SSID NAME     COLLID        STMTNO UNITS
 ---- ---- -------- ------------ ------ -------
 XD00 DIA1 ACT8SELE ACT930_D_MAI  12096*   1314
 XD01 DIA1 ACT8SELE ACT930_D_MAI  12096    1015
```

Note: Weighted cost = (cost * executions)

© Copyright 9/15/2011 BMC Software, Inc

Report Output

```
SELECT   A.COL1, A.COL2, A.COL3, A.FILLER1, B.COL1, B.COL2, B.COL3, B.FILLER1
FROM     ISPRXM.TEST_STATCLUS A, ISPRXM.TEST_STATCLUS2 B
WHERE    A.COL1 BETWEEN '15' AND '30'
  AND    B.COL2 BETWEEN '15' AND '30'
  AND    A.COL1 = B.COL2

     QB PL M TB AC MT                          I  NEW  COMP P C MX LCK J
LABL BL NO E NO TY CL TABLE        INDEX       O  UJOG UJOG F F SQ MOD C
---- -- -- - -- -- -- ------------ ------------ -- ---- ---- - - -- --- -
XD00  1  1 0* 2 I   1*STATCLUS2   *STATCLUS2_IX2N  NNNN NNNN*L    0  IS 0
XD00  1  2 1* 1 I   1*STATCLUS    *STATCLUS_IX1 N  NNNN*NYNN*S    0  IS 0

XD01  1  1 0  1 I   1 STATCLUS     STATCLUS_IX1 N  NNNN NNNN S    0  IS 0
XD01  1  2 1  2 I   1 STATCLUS2    STATCLUS2_IX2N  NNNN NNNN      0  IS 0
```

© Copyright 9/15/2011 BMC Software, Inc

Conclusions and Questions

SQLADRIA SEMINAR
Zagreb, 23 September 2011

Squeeze the Most out of your Indexes in DB2 9!

Peter Plevka
BMC Software

Sponsored by

Squeeze the Most Out of Your Indexes in DB2 9!

Peter Plevka
BMC Software Inc.

Objectives

- Learn how index compression works.
- Learn how index compression differs from tablespace compression.
- Learn the steps necessary to implement compression for an index.
- Learn how to identify the indexes which are the best candidates for compression.
- Learn how to most effectively exploit index compression to save disk storage and meet your SQL performance goals.

© Copyright 9/16/2011 BMC Software, Inc

Index Compression – How Does It Work?

© Copyright 9/16/2011 BMC Software, Inc

Tablespace Compression

- Has been with us a long time
  - Since V3 of DB2!
- Compresses at the row level
- Uses Liv-Zempel algorithm
  - Needs dictionary
  - Hardware instructions – fast!
- Dictionary is built from data by Load or Reorg
- 85% is typical data reduction
- DSN1COMP utility can help plan

© Copyright 9/16/2011 BMC Software, Inc

Tablespace Compression

- Row is compressed and expanded by SQL
- Copies and logs are compressed
  - Log and data analysis tools can be slowed down
- Index data is not compressed
  - Can slow down utilities that extract keys
  - Index processing is not slowed down
  - Some indexes are bigger than their tablespaces!
- Effective technique to reduce disk usage without undue effect on SQL

© Copyright 9/16/2011 BMC Software, Inc

DB2 Index Structure

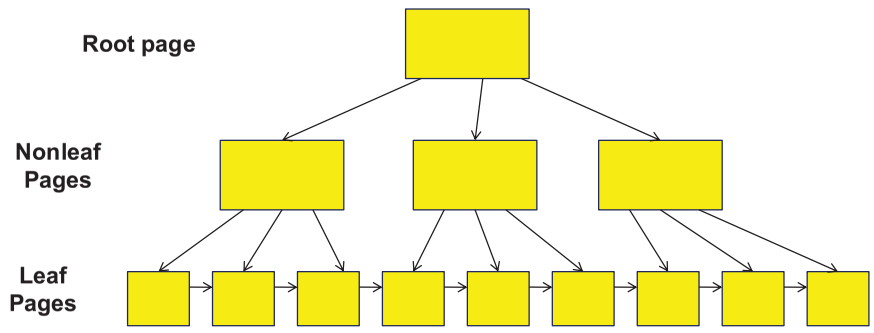

© Copyright 9/16/2011 BMC Software, Inc

DB2 Index Leaf Page

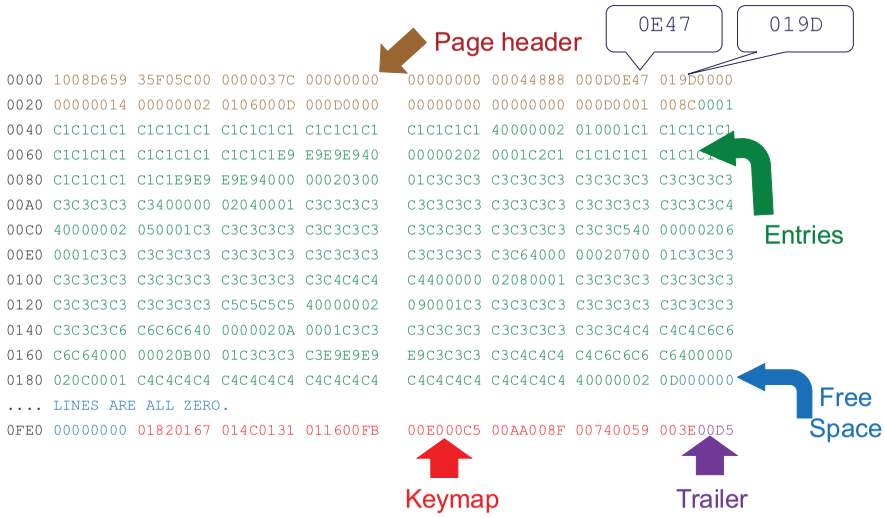

© Copyright 9/16/2011 BMC Software, Inc

DB2 Index Leaf Entries

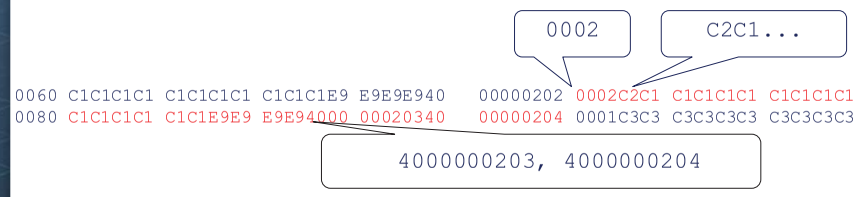

© Copyright 9/16/2011 BMC Software, Inc

Index Compression Algorithm

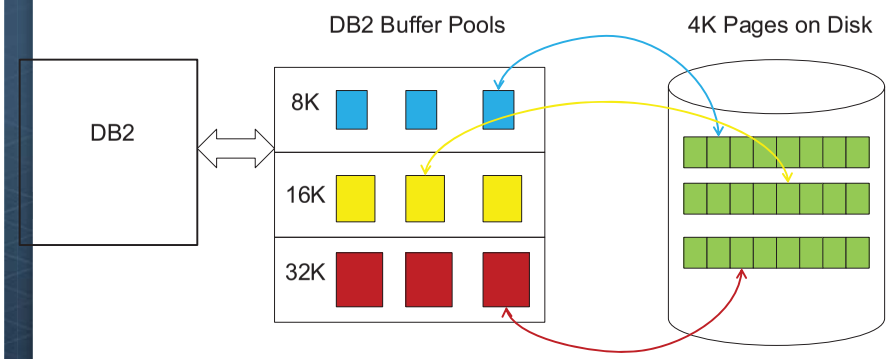

- Compression/expansion at time of I/O
- 4K on disk, 8/16/32K in buffer pool

© Copyright 9/16/2011 BMC Software, Inc

How Does DB2 Compress A Page?

Izzard	
Jace	
James	
Jameson	
Jamestown	
Jane	
Janes	
Janest	

© Copyright 9/16/2011 BMC Software, Inc

Index Entry Compression

Key value

4000000201,
4000000202

0002

```
0020 00000014 00000002 0106000D 000D0000   00000000 00000000 000D0001 008C0002
0040 C1C1C1C1 C1C1C1C1 C1C1C1C1 C1C1C1C1   C1C1C1C1 40000002 01400000 02020001
0060 C1C1C1C1 C1C1C1C1 C1C1C1C1 C1C1C1C1   E9E9E9E9 40000002 03
```

The two key values are:

AAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAZZZZ

02 | 00 | Key value

420200, 40

```
0040 0EFF0002 00C1C1C1 C1C1C1C1 C1C1C1C1   C1C1C1C1 C1C1C1C1 C1420200 40110E9
0060 E9E9E942 0202
```

420202

01 | 10

'ZZZZ'

© Copyright 9/16/2011 BMC Software, Inc

DB2 Index Leaf Page (Compressed)

0EFF | Page header | 3E43 | 01A1

```
0000 1008D659 35F0D800 0000037C 00000000   00000000 00064888 000D3E43 01A10000
0020 00000014 00000002 010A000F 000D0000   00000000 00000000 000D0001 808C0001
0040 0EFF0001 00C1C1C1 C1C1C1C1 C1C1C1C1   C1C1C1C1 C1C1C1C1 C1420200 0110E9E9
0060 E9E94202 010100C2 C1C1C1C1 C1C1C1C1   C1C1C1C1 C1C1C1E9 E9E9E942 02020100
0080 C3C3C3C3 C3C3C3C3 C3C3C3C3 C3C3C3C3   C3C3C3C3 42020301 13C44202 040113C5
00A0 42020501 13C64202 060110C4 C4C4C442   02070110 C5C5C5C5 42020801 10C6C6C6
00C0 C6420209 010CC4C4 C4C4C6C6 C6C64202   0A0104E9 E9E9E9C3 C3C3C3C4 C4C4C4C6
00E0 C6C6C642 020B0100 C4C4C4C4 C4C4C4C4   C4C4C4C4 C4C4C4C4 C4C4C4C4 42020C00
.... LINES ARE ALL ZERO.
0FE0 00000000 00000000 00000000 00000000   00000000 00000000 00000000 00000000
```

Entries

Free Space

Keymap ?

Trailer

© Copyright 9/16/2011 BMC Software, Inc

Implementation of Index Compression

- You can create a new index to be compressed
  - Specify 8K, 16K, or 32K BUFFERPOOL, and COMPRESS YES
- What about existing indexes?

© Copyright 9/16/2011 BMC Software, Inc

Implementation of Index Compression

- ALTER INDEX A.B BUFFERPOOL BP16K3; (for example)
  - Index must be stopped in data sharing
- ALTER INDEX A.B COMPRESS YES;
- Do at same time – avoid two REBUILDS
- Changing page size puts index in RBDP
- Changing to COMPRESS YES puts index in AREOR

© Copyright 9/16/2011 BMC Software, Inc

Effective Compression Ratio

- Depends on data in keys
- Consider a unique, padded index with a 40 byte key
- Can get 85 keys in a 4K page (uncompressed)
- Two extreme cases:
  - Consecutive keys with 38 matching bytes
  - Consecutive keys with 2 matching bytes

© Copyright 9/16/2011 BMC Software, Inc

Effective Compression Ratio

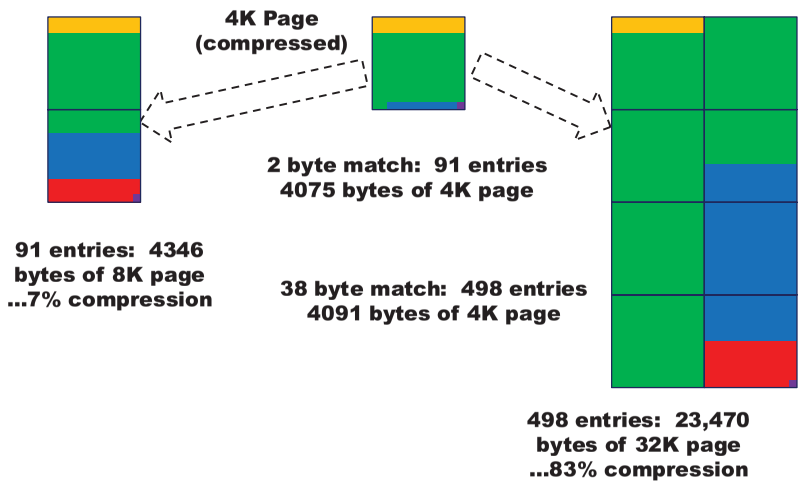

© Copyright 9/16/2011 BMC Software, Inc

DSN1COMP

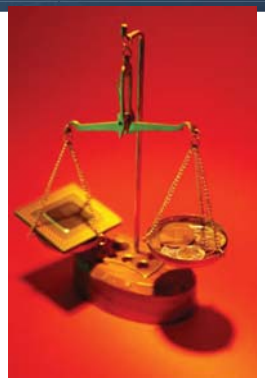

```
//MYJOB     JOB …
//DSN1      EXEC PGM=DSN1COMP,
//          PARM='LEAFLIM(10000)'
//STEPLIB   DD   DSN=DSN.LOAD,DISP=SHR
//SYSUT1    DD   DSN=VCAT.DSNDBC.MYDB.IX1.I0001.A001,DISP=SHR
//SYSPRINT  DD   SYSOUT=*
```

© Copyright 9/16/2011 BMC Software, Inc

DSN1COMP

```
DSN1999I START OF DSN1COMP FOR JOB SJDCX05  COMP1
DSN1998I INPUT DSNAME = VCAT.DSNDBC.MYDB.IX1.I0001.A001          , VSAM
                                  …
DSN1940I DSN1COMP COMPRESSION REPORT

              118  Index Leaf Pages Processed
           10,000  Keys Processed
           10,000  Rids Processed
              459  KB of Key Data Processed
              430  KB of Compressed Keys Produced
```

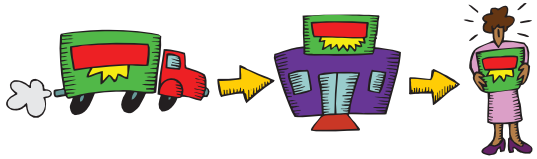

© Copyright 9/16/2011 BMC Software, Inc

DSN1COMP

```
EVALUATION OF COMPRESSION WITH DIFFERENT INDEX PAGE SIZES:

                 ---------------------------------------------
                8  K Page Buffer Size yields a
                6  % Reduction in Index Leaf Page Space
                   The Resulting Index would have approximately
               94  % of the original index's Leaf Page Space
               46  % of Bufferpool Space would be unused to
                   ensure keys fit into compressed buffers
                 ---------------------------------------------
                 ---------------------------------------------
               16  K Page Buffer Size yields a
                7  % Reduction in Index Leaf Page Space
                   The Resulting Index would have approximately
               93  % of the original index's Leaf Page Space
               73  % of Bufferpool Space would be unused to
                   ensure keys fit into compressed buffers
                 ---------------------------------------------
```

© Copyright 9/16/2011 BMC Software, Inc

DSN1COMP – another index

```
       8  K Page Buffer Size yields a
      51  % Reduction in Index Leaf Page Space
          The Resulting Index would have approximately
      49  % of the original index's Leaf Page Space
          No Bufferpool Space would be unused

      16  K Page Buffer Size yields a
      76  % Reduction in Index Leaf Page Space
          The Resulting Index would have approximately
      24  % of the original index's Leaf Page Space
          No Bufferpool Space would be unused

      32  K Page Buffer Size yields a
      82  % Reduction in Index Leaf Page Space
          The Resulting Index would have approximately
      18  % of the original index's Leaf Page Space
      30  % of Bufferpool Space would be unused to
          ensure keys fit into compressed buffers
```

© Copyright 9/16/2011 BMC Software, Inc

Which Indexes Should I Compress?

```
SELECT DBNAME, INDEXSPACE, NLEAF, PGSIZE, SPACEF
       FROM SYSIBM.SYSINDEXES
       WHERE SPACEF > 100000 AND COMPRESS <> 'Y'
       ORDER BY SPACEF DESC;
```

© Copyright 9/16/2011 BMC Software, Inc

Which Indexes Should I Compress?

```
SELECT I.DBNAME, I.INDEXSPACE, I.TBNAME, I.TBCREATOR,
       I.FIRSTKEYCARDF, I.FULLKEYCARDF, I.SPACEF, I.AVGKEYLEN,
       C.LENGTH,
       (I.SPACEF *
        (1.0 - I.FIRSTKEYCARDF/I.FULLKEYCARDF ) * C.LENGTH) /
       (I.AVGKLEN * I.FULLKEYCARDF) AS SAVING
       FROM SYSIBM.SYSINDEXES I, SYSIBM.SYSKEYS K,
            SYSIBM.SYSCOLUMNS C
       WHERE I.SPACEF > 100000 AND I.COMPRESS <> 'Y'
         AND I.NAME = K.IXNAME AND I.CREATOR = K.IXCREATOR
         AND I.TBNAME = C.TBNAME AND I.TBCREATOR = C.TBCREATOR
         AND I.FIRSTKEYCARDF > 0 AND I.FULLKEYCARDF > 0
         AND I.DBNAME <> 'DSNDB06'
         AND K.COLSEQ = 1 AND K.COLNO = C.COLNO
         AND C.COLTYPE <> 'VARCHAR' AND C.COLTYPE <> 'VARBIN'
       ORDER BY 10 DESC;
```

© Copyright 9/16/2011 BMC Software, Inc

Which Indexes Should I Compress?

- What did we miss with the SQL on the preceding slide?

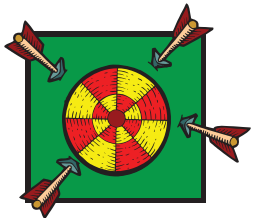

```
SELECT I.DBNAME, I.INDEXSPACE, I.TBNAME, I.TBCREATOR,
      I.FULLKEYCARDF, I.SPACEF, I.AVGKEYLEN, T.CARDF,
      T.CARDF / I.FULLKEYCARDF AS AVGRIDS
       FROM SYSIBM.SYSINDEXES I, SYSIBM.SYSTABLES T
       WHERE I.SPACEF > 100000 AND I.COMPRESS <> 'Y'
         AND I.TBNAME = T.NAME AND I.TBCREATOR = T.CREATOR
         AND I.FULLKEYCARDF > 0 AND I.DBNAME <> 'DSNDB06'
         AND I.CLUSTERRATIOF > .80
       ORDER BY I.SPACEF DESC, 9 DESC, I.AVGKEYLEN ASC;
```

© Copyright 9/16/2011 BMC Software, Inc

You Can Now Version Compressed Indexes!

- APAR PK79312, PTF UK52416 (F912 level)
- AREO* now set instead of RBDP
- Read the APAR – there are preconditioning PTF's

© Copyright 9/16/2011 BMC Software, Inc

Performance Considerations

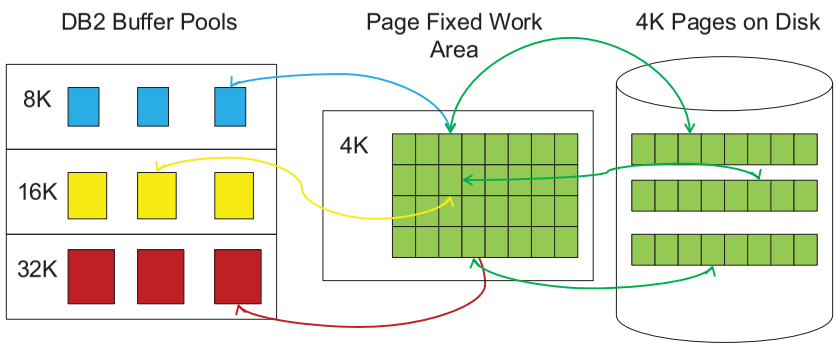

- 4K I/O done into/from page fixed work area
- Expansion/compression done at time of I/O
  - Appears as class 2 CPU (synch) or DBM1 SRB (asynch)

© Copyright 9/16/2011 BMC Software, Inc

Performance Considerations

- Performance characteristics are very different from tablespace compression:
- Random access works well with a compressed tablespace
  - No overhead sustained until row is fetched
- Not so well with a compressed index
  - Entire leaf page must be expanded to get one entry
  - Higher compression ratio increases hit rate but also increases CPU overhead
  - Elapsed time increase is probably negligible
  - Compare average cost of synchronous I/O

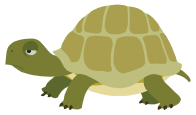

© Copyright 9/16/2011 BMC Software, Inc

Performance Considerations

- Performance characteristics are very different from tablespace compression:
- Sequential access does not work well with a compressed tablespace
  - Every row must be expanded
- Better performance with compressed index
  - One expansion call per page, not per entry
  - Dynamic prefetch can be triggered ➔ asynchronous I/O➔ no impact to elapsed time

© Copyright 9/16/2011 BMC Software, Inc

Performance Considerations

- So index scans are what we want to see
- Dynamic prefetch if index is well organized
  - Low elapsed time, no class 2 CPU
- No dynamic prefetch for disorganized index
  - Synch I/O and class 2 CPU ➔ elapsed time
  - Might record fewer I/O's (one per page)
- Consider REORG INDEX more often
- Elapsed time reduction will be higher with faster processor (less CPU bound)
- "Other Read I/O time" class 3 suspend time for I/O bound index scan

© Copyright 9/16/2011 BMC Software, Inc

Performance Considerations

- Dynamic prefetch changes when you compress
  - Still reserves 128K in buffers
  - 64K, 32K, or 16K in work area
  - Traces show more prefetch reads
  - CPU is DBM1 SRB
- Most writes should be asynch (deferred write engines)
  - No elapsed time hit
  - CPU is DBM1 SRB time

© Copyright 9/16/2011 BMC Software, Inc

Performance Considerations

- What is gross CPU overhead?
  - i.e., can you afford it?
- CPU depends primarily on number of keys
- ROT is 4 µs per 8K page on z990, 3 on z9, 1.9 on z10
  - Double for 16K, quad for 32K
- Multiply by I/O rate per second
  - 15,000 per sec, 8K, z10 ➔ 28,500 µs/sec ➔ 2.8% of processor

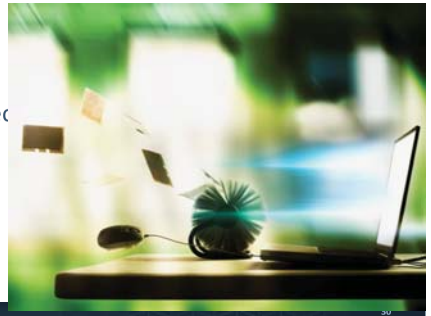

© Copyright 9/16/2011 BMC Software, Inc

Conclusion

- You can save space without destroying performance
- Select candidates
- Vet them for access

© Copyright 9/16/2011 BMC Software, Inc

Bibliography

- "DB2 9.1 for z/OS Diagnosis Guide and Reference", LY37-3218-03
- "Index Compression with DB2 9 for z/OS", Redpaper, Jeff Berger and Paolo Bruni
- "DB2 9 for z/OS Performance Topics", SG23-7473-00
- "DB2 9.1 for z/OS SQL Reference", SC18-9854-02
- "DB2 9.1 for z/OS Utility Guide and Reference", SC18-9855-03

© Copyright 9/16/2011 BMC Software, Inc

Questions

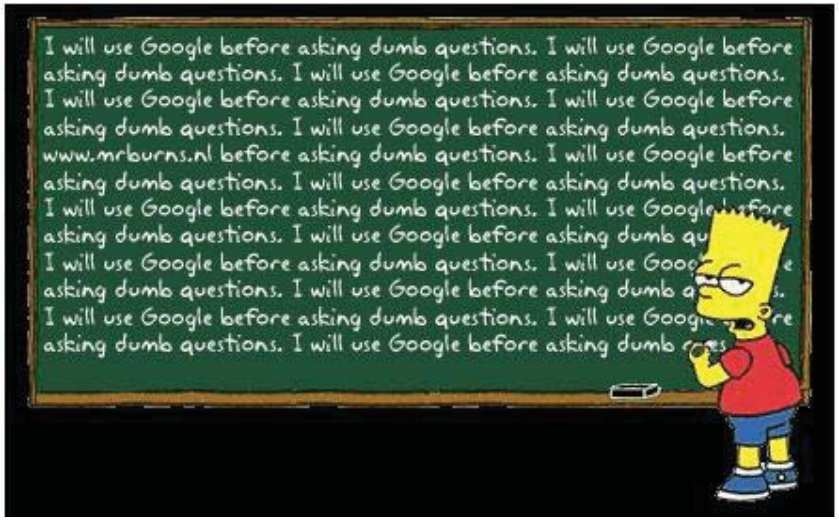

© Copyright 9/16/2011 BMC Software, Inc

SQLADRIA SEMINAR
Zagreb, 23 September 2011

DSNZPARMs for DBA's

Ramon Menendez
BMC Software

Sponsored by

The "you have to know the DSNZPARM" bar has been raised

Ramón Menéndez, BMC Software
Zagreb, 23rd September 2011

Preface

Not intended to be complete

Once you know one parameter, you see another one related to previous one

Ever increasing number of controls
- DB2 v8 +/- 257
- DB2 9 +/- 278
- DB2 10 +/- 289
  - new, removed, deprecated, opaque, externalized, changed maximum, changed default

Though looks like Systems Programmers task,

it affects Database and Application program design

© Copyright 9/15/2011 BMC Software, Inc

Recap DSNZPARM – Function

- A data-only subsystem parameter load module that contains DB2's execution time parameters
  - ZPARM is load module from assemble/link of DSN6xxx macro's
- They include the following functions:
  - Identifying information for the DB2 subsystem
  - Defining various internal memory areas
  - Timing and behavior of internal processes, for example, for storage management and optimization choices
  - Timers and controls that monitor thresholds, like locks
  - Default choices for such items as index structures
  - Environmental factors, such as language support

© Copyright 9/15/2011 BMC Software, Inc

Recap DSNZPARM – Controls

- Most visible parameters You all know about
  - maximum locks per user — NUMLKUS
  - maximum locks per user per table space — NUMLKTS
  - maximum users — CTHREAD
  - maximum TSO, batch — IDFORE, IDBACK
  - maximum remote connections, active — CONDBAT, MAXDBAT
  - maximum number of open data sets — DSMAX

Sample Display of parms

```
DSN6SYSP
    AUDITST=0,                  AUDIT TRACE
    CONDBAT=10000,              MAX REMOTE CONNECTED
    CTHREAD=200,                MAX USERS
    DLDFREQ=5,                  LEVELID UPDATE FREQUENCY
    PCLOSEN=5,                  RO SWITCH CHKPTS
    IDBACK=70,                  MAX BATCH CONNECT
    IDFORE=50,                  MAX TSO CONNECT
    CHKFREQ=500000,             CHECKPOINT FREQ
    MON=0,                      MONITOR TRACE
    MONSIZE=262144,             MONITOR SIZE
    SYNCVAL=NO,                 STATISTICS SYNC
    RLFAUTH=SYSIBM,             RESOURCE AUTHID
    RLF=NO,                     RLF AUTO START
    RLFERR=NOLIMIT,             RLST ACCESS ERROR
    RLFTBL=1,                   RLST NAME SUFFIX
```

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM - Considerations

- Apart from DB2 system configuration and performance management, other considerations include:
  - Which parameters are related to others and what are the limits?
  - Which parameters can be changed online, while DB2 is up, and which cannot.
  - And do you document what parameters gets changed when, and why?
  - How to perform change management on system parameters.
  - If, when, and how to go about modifying hidden DSNZPARM values.
- DBA's are glad there are people that worry about this?
- There are people worrying, right?

© Copyright 9/15/2011 BMC Software, Inc

Remember DSNZPARM – How to define

- Run DSNTINST CLIST:
  - INSTALL
  - UPDATE
  - MIGRATE
- CLIST will gen ALL jobs needed for INSTALL / MIGRATE
- CLIST will gen only DSNZPARM / DSNHDECP for update

```
                    INSTALL DB2 - THREAD MANAGEMENT
===>

Check numbers and reenter to change:

 1  DATABASES              ===> 100        Concurrently in use
 2  MAX USERS              ===> 80         Concurrently running in DB2
 3  MAX REMOTE ACTIVE      ===> 60         Maximum number of active
                                             database access threads
 4  MAX REMOTE CONNECTED   ===> 60         Maximum number of remote DDF
                                             connections that are supported
 5  MAX TSO CONNECT        ===> 100        Users on QMF or in DSN command
 6  MAX BATCH CONNECT      ===> 15         Users in DSN command or utilities
 7  SEQUENTIAL CACHE       ===> BYPASS     3990 storage for sequential IO.
                                            Values are SEQ or BYPASS.
 8  UTILITY CACHE OPTION   ===> NO         3990 storage for DB2 utility IO
 9  MAX KEPT DYN STMTS     ===> 5000       Maximum number of prepared dynamic
                                            statements saved past commit points
10  CONTRACT THREAD STG    ===> NO         Periodically free unused thread stg
11  MANAGE THREAD STORAGE  ===> NO         Manage thread stg to minimize size

PRESS:  ENTER to continue   RETURN to exit   HELP for more information
```

- BUT:
  - You lose values not on panels (not in DSNTIDxx member)
    - "Subsystem parameters that are not on installation panels"
  - You lose hidden ZPARM changes
    - Just to name a few: NPGTHRSH, SARGSWRP, SMSDCFL, SMSDCIX
  - You have NO history
  - There is NO relationship between panel-names and Macro's generated

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM - Maintaining

- Via CLIST
  - Migrate possible
  - Lose values
  - Lose history
- Via SOURCE
  - Migrate is INSTALL (can be very healthy to think about every value)
  - Can change EVERYTHING
  - Can add comment (assembler comment = asterisk in col 1)
    Stuff like: Old value, Who, when, why (change request number) etc.
  - Easier to create several copies
  - Put into source system
- DSNTXAZP tool
  - Updates a CLIST defaults input member (DSNTIDxx) with the current subsystem parameter and buffer pool settings.
  - Provides a report of all parameters.

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM – Dynamic Change

- Changing your DSNZPARM
- It is not Anymore
  - Edit Parameter List, assemble, link edit and Restart DB2
- It is not Really
  - Change DSNZPARM keyword values online
- It is not Exactly
  - Change Some DSNZPARM online
- What is it then?

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM – Dynamic Change

- Version 7 onwards will allow a subset of the most popular DSNZPARMS to be dynamically changed
- Advantages …
  - Tailor parameters to the current workload
  - No outages
- The ZPARM member is changed dynamically in its entirety by activating a different ZPARM member
- New command for changing parameter module

-SET SYSPARM LOAD (modname)
Loads the named parameter module-default (if not specified)=DSNZPARM

-SET SYSPARM RELOAD
Last named subsystem parameter module is loaded into storage

-SET SYSPARM STARTUP
Loads the initial parameters from DB2 startup

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM – Dynamic Change

- For most parameters
  - The online change will be transparent
  - The online change will be immediate
- However others react differently due to impact
  - Some only affect new threads for example, not existing threads
- Each parameter must be examined to determine impact
- Some when increased take effect immediately
- Some when decreased will require other events to occur

IFCID 0001 – has a new field to record the number of -SET SYSPARM commands executed

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM – Dynamic Change

- Parameters, when updated online, result in a change in system behavior.
  - CACHEDYN
  - CACHEDYN_FREELOCAL
  - MAXKEEPD
  - PARTKEYU
  - SYSADM/SYSADM2
  - XLKUPDLT
- Parameters updateable indicated in Installation Guide

Table 133. Directory of subsystem parameters and DSNHDECP values (continued)

Parameter	Macro	Panel	Update Online
ARCPFX2	DSN6ARVP	DSNTIPH	Yes
ARCRETN	DSN6ARVP	DSNTIPA	Yes
ARCWRTC	DSN6ARVP	DSNTIPA	Yes
ARCWTOR	DSN6ARVP	DSNTIPA	Yes
ARC2FRST	DSN6LOGP	DSNTIPO	Yes
ASCCSID	DSNHDECP	DSNTIPF	—
ASSIST	DSN6GRP	DSNTIPK	No
AUDITST	DSN6SYSP	DSNTIPN	No
AUTH	DSN6SPRM	DSNTIPP	No
AUTHCACH	DSN6SPRM	DSNTIPP	Yes
BACKODUR	DSN6SYSP	DSNTIPL	No
BINDNV	DSN6SPRM	DSNTIPP	Yes

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM – Short Summary

- DSNZPARM has IMPORTANT Tuning parameters
  - If you run for years without changing anything then maybe something is very wrong…

DSNT105I / DSNT106I THE LOADED DSN6SPRM CSECT IS OBSOLETE.
PLEASE REASSEMBLE THE DSN6SPRM CSECT
AND RELINKEDIT THE DSNZPARM LOAD MODULE.

- Who defines the DSNZPARM and why
- Differences between test and production
- How do you maintain your DSNZPARM afterwards

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM Changes in DB2 9

Not a complete list!

DB2 9 DSNZPARM Change Summary - Removed

- The option to hold a lock after commit is eliminated. DB2 always releases the page or row lock for cursor hold at commit, which allows for greater concurrency.
  - DSN6SPRM.RELCURHL:
    - Release at COMMIT any page or row locks on which a cursor WITH HOLD is positioned.
- Default protocol (DRDA or PRIVATE) when DBPROTOCOL is not explicitly specified at BIND
  - DSN6SYSP.DBPROTCL:
- Minimum number of table joins in a query to cause DB2 to monitor the resources consumed when determining the optimum access path.
  - DSN6SPRM.TABLES_JOINED_THRESHOLD:
- Maximum amount of elapsed time(seconds) consumed by the DB2 Optimizer.
  - DSN6SPRM.MAX_OPT_ELAP:

© Copyright 9/15/2011 BMC Software, Inc

DB2 9 DSNZPARM Change Summary - Modified

- DSN6SPRM.MGEXTSZ:
  - Use sliding secondary quantity for DB2-managed data sets.
  - Default changed from NO to YES.
  -
- DSN6SYSP.IDXBPOOL:
  - Default bufferpool for CREATE INDEX.
  - BP0, but now accepts 8 KB, 16 KB, and 32 KB page size BPs in addition to 4 KB.
- DSN6SYSP.WLMENV:
  - Name of the default WLM environment for DB2.
  - Length has increased from 18 to 32 characters.

© Copyright 9/15/2011 BMC Software, Inc

DB2 9 DSNZPARM Change - DSN6SPRM Added

- CACHEDYN_FREELOCAL
  - This parameter specifies the level at which DB2 can free cached dynamic statements to relieve DBM1 below-the-bar storage.
    - 0 does not free
    - 1 frees some memory
- EDM_SKELETON_POOL
  - This value specifies the minimum size in KB of the above-the-bar EDM skeleton pool for skeleton package tables (SKPTs) and skeleton cursor tables (SKCTs).
  - The default is 5Mb in V9 and 10Mb in V10.
- MAX_CONCURRENT_PKG_OPS
  - This parameter specifies the maximum number of automatic bind requests that can be processed simultaneously.
  - The default is 10.

© Copyright 9/15/2011 BMC Software, Inc

DB2 9 DSNZPARM Change - DSN6SPRM Added

- MXDTCACH
  - This value specifies the maximum size, in MB, of memory for data caching that DB2 allocates from above the bar.
  - Default is 20 MB.
- REOPTEXT
  - This values specifies whether to automatically reoptimize dynamic SQL.
  - Later removed by APAR (!)
- STATCLUS
  - default RUNSTATS clustering statistics type.
  - Default is ENHANCED (Alternative value is STANDARD)
  - This parameter should improve access path selection for duplicate and reverse-clustered indexes.
  - For indexes that contain either many duplicate key values or key values that are highly clustered in reverse order, cost estimation that is based purely on CLUSTERRATIOF can lead to repetitive index scans.
  - A new cost estimation formula based on the DATAREPEATFACTORF statistic to choose indexes avoids this performance problem.

© Copyright 9/15/2011 BMC Software, Inc

DB2 9 DSNZPARM Change - DSN6SYSP Added

- IMPDSDEF
  - This value specifies whether to define the data set when creating an implicit base table space or implicit index space.
  - Default is YES.
- IMPTSCMP
  - This value specifies whether to enable data compression on implicit base table space.
  - Default is NO.
- MAXOFILR
  - This specifies the maximum number of data sets that can be open concurrently for processing of LOB file references.
  - The default is 100
- MAXTEMPS
  - This value specifies the maximum MB of temporary storage in the workfile database that can be used by a single agent at any given time for all temporary tables.
  - The default is 0 (meaning no limit, the same as previous releases).

© Copyright 9/15/2011 BMC Software, Inc

DB2 9 DSNZPARM Change - DSN6SYSP Added

- TBSBPLOB
  - Name of the Bufferpool to be used by LOB table spaces that are implicitly created.
- TBSBPXML
  - Name of the Bufferpool to be used by XML table spaces that are implicitly created. Any valid 16K BP name is accepted. The default is BP16K0.
- TBSBPOOL - TBSBP8K -TBSBP16K - TBSBP32K
  - Default BPs for CREATE TABLESPACE depending on page size
- XMLVALA
  - in KB an upper limit for the amount of storage that each agent can have for storing XML values
  - Default is204800 KB.
- XMLVALS
  - in MB an upper limit for the amount of storage that each system can use for storing XML values
  - Default is 10240 MB.

© Copyright 9/15/2011 BMC Software, Inc

DB2 9 DSNZPARM Change Summary - Added

- DSN6SPRM.HONOR_KEEPDICTIONARY:
  - honor or ignore the KEEPDICTIONARY parameter on LOAD REPLACE or REORG request that converts a table space from basic row format to reordered row format.
  - Default is NO, meaning that dictionary is rebuilt.

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM
DB2's Application Environment

EDM pool changes for dynamic SQL statements

- Limit the number of Dynamic SQL Statements held in the Cache
- DSNZPARM MAXKEEPD
- Dynamic SQL execution storage is environment dependent and can create storage constraints
- DB2 9
  - Split below/above the bar pools
  - IFCID 225 (below)
    - QW0225LC - # stmts in cache
    - QW0225HC – High water mark
  - IFCID 217 (above)
- -START TRACE(GLOBAL) IFCID(217)

© Copyright 9/15/2011 BMC Software, Inc

Favor matching index access

- Special access path selection for tables under a given size.
- Not on installation panels
- DB2 9 Changed acceptable values for subsystem parameter NPGTHRSH
- DSNZPARM NPGTHRSH
  - -1: DB2 favors matching index access for all tables. **Not Longer valid**
  - 0: DB2 selects the access path based on cost, and no tables qualify for special handling. This is the default.
  - >=1: DB2 favors matching index access for tables for which the total number of pages on which rows of the table appear (NPAGES) is less than *n*.

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM
Objects

DSVCI -Support for VSAM Control Interval (CI) Greater than 4K

- Support for CI sizes of 8, 16, and 32 K table spaces.
- CI size DB2 Managed Datasets corresponds to the buffer pool of table space
- Implicitly redefined by LOAD or REORG (or RECOVER) of the table space.
- New CI sizes
  - Potentially reduce elapsed time for table space scans
  - I/O bound executions benefit with using larger VSAM CI sizes, including COPY and REORG
- DSNZPARM DSVCI (VARY DS CONTROL INTERVAL)
  - YES/NO

© Copyright 9/15/2011 BMC Software, Inc

Automatic Creation of Table Space

- If CREATE TABLE does not include explicit table space specification, a table space is automatically (implicitly) created with the following attributes
  - **Partition-by-growth universal table space**
    - **SEGSIZE=4**
    - **DSSIZE=4G**
    - **MAXPARTITIONS=256**
  - **Compression is controlled by DSNZPARM IMPTSCMP**
    - **Default is COMPRESS=NO**
  - **Dataset creation is controlled by DSNZPARM IMPDSDEF**
    - **Default is DEFINE=YES**
  - **Sliding scale for optimizing secondary extent allocation is used by default**
  - **The default buffer pool is determined based on the DSNZPARM settings.**
    - **Buffer pool page size is derived automatically using larger page size if max record size reaches 90%**

© Copyright 9/15/2011 BMC Software, Inc

Sliding Secondary Allocation Size

- From DB2 V8 onwards allocate secondary extents automatically
  - Set SECQTY to -1. Note all implicit tablespaces use this
  - Default PRIQTY
    - Controlled by ZPARM TSQTY and IXQTY
    - 1 cylinder for non-LOB tablespaces and indexes / 10 cylinders for LOB
- Extent sizes allocated gradually increase
  - Maximum Extent size is based on DSSIZE, LARGE and PIECESIZE and defaults
- Advantages:
  - Minimizes the potential for wasted space by increasing the size of secondary extents slowly at first
  - It prevents very large allocations for the remaining extents, which would likely cause fragmentation.
  - It does not require users to specify SECQTY values when creating and altering table spaces and index spaces.

DSNZPARM MGEXTSZ (OPTIMIZE EXTENT SIZING)
Default is NO for V8
Changed to YES automatically when upgrade to DB2 9

- Existing pagesets: Execute SQL to ALTER PRIQTY/SECQTY values to -1 plus schedule a REORG

© Copyright 9/15/2011 BMC Software, Inc

Reordered Row Format

- Basic Row Format
  - Variable-length columns are stored in the row as defined.
  - DB2 needed to do offset calculations to access each column.
- Reordered Row Format
  - Variable-length columns are stored at the end of the row
  - All the columns are directly addressable.
  - no change in any SQL syntax
    - Application Transparent
  - Exception: EDITPROC or VALIDPROC

```
CREATE TABLE TB1 (C1 INTEGER NOT NULL, C2 VARCHAR(10) NOT
NULL, C3 CHAR(10) NOT NULL, C4 VARCHAR(20)
```

Basic Row Format

C1	C2	C3	C4
8000000A	0006 WILSON	ANDREW	0008 SAN JOSE

2-byte length

Reordered Row Format

C1	C3	O2	O4	C2	C4
8000000A	ANDREW	12	18	WILSON	SAN JOSE

offset to C2 offset to C4

© Copyright 9/15/2011 BMC Software, Inc

Reordered Row Format

- Impact only for varying-length columns
- At the beginning enabled in NFM, then hidden DSNZPARM SPRMRRF and then converted to an opaque, online changeable DSNZPARM. Default value will remain SPRMRRF=ENABLE
  - Externalized as RRF in DB2 10
- REORG / LOAD REPLACE utility
  - From Basic Row Format to Reordered Row Format
  - KEEPDICTIONARY may result in poor compression
  - DB2 9 ignores KEEPDICTIONARY
  - UNLESS:
- DSNZPARM: HONOR_KEEPDICTIONARY
  - parameter defined by job DSNTIJUZ and does not appear on any installation panels

© Copyright 9/15/2011 BMC Software, Inc

Resource Timeout

- Force Utility to wait longer for a resource than SQL application
  - DSNZPARM UTIMOUT (Utility Timeout)
  - Default = 6
  - Default value for RETRY parameter of DB2 Utilities such as CHECK INDEX, and online REBUILD INDEX, and REORG.
- Trapping Long running units of Recovery
  - DSNZPARM URLGWTH
  - number of log records that are to be written by an uncommitted unit of recovery (UR) before DB2 issues a warning message to the console.
  - Default 0: No warning in V9, changed in V10
  - Activate IFCID313 (Statistics Class 3) for historic monitoring

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM Packages

Plan Stability

- REBINDs can cause access path changes
- No easy way to undo the REBIND
- Existing “solutions” inadequate
  - Preventing REBINDS altogether
  - REBINDing into alternate collections
  - Using hints
- Good Household practice to rebind on the new release to get the best performance
- Almost all Optimizer enhancements available in DB2 9 (CM)

© Copyright 9/15/2011 BMC Software, Inc

Plan Stability – The solution

- At REBIND, DB2 will save old copies of packages
- Two flavors
  - BASIC and EXTENDED
  - Controlled by new ZPARM PLANMGMT
  - Default was OFF for V9, now EXTENDED for V10 (!)
  - Also supported as REBIND options
- Catalog support
  - SYSPACKAGE reflects active copy
  - SYSPACKDEP reflects dependencies of all copies
  - Other catalogs (SYSPKSYSTEM, …) reflect metadata for all copies
- Invalidation and Auto Bind
  - Each copy invalidated separately
  - Auto bind replaces only the current copy – previous and original are not affected

DSNZPARM PLANMGMT

OFF – default	- only the current package in SPT01.
BASIC	- two copies in SPT01.
EXTENDED	- up to three copies of each package in SPT01 and SYSPACKDEP

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARMs – Packages

- ABIND - AUTO BIND
  - allow auto rebind of packages or plans at runtime
  - YES - DB2 automatically rebinds a package or plan one time at execution time
    - conditions
      - » If it is marked as non VALID
      - » Was last bound in a later DB2 version
      - » Was last autobound in a previous version of DB2
  - NO - Automatic rebind operations are not allowed under any circumstances.
    - If you attempt to run any plan or package on DB2 9 one of the situations described before, you receive SQLCODE -908 SQLSTATE 23510.
  - COEXIST - Automatic rebind operations to be performed in a DB2 data sharing environment
    - Conditions
      - » SYSPLAN or SYSPACKAGE is marked invalid
      - » Was last bound in a later DB2 version
    - Do not autobind if was last autobound in a previous version of DB2 to allow coexistence and reduce the rate of auto rebinds

SYSPACKAGE.RELBOUND identifies the DB2 version a package was last bound
- » **'L'** Version 8 dependency indicator
- » **'M'** Version 9 dependency indicator
- » **'O'** Version 10 dependency indicator

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM's - Packages

- ABEXP - EXPLAIN processing during AUTO BIND plan or package
  - Condition:
    - Plan/package bound EXPLAIN(YES) and OWNER Plan Table exists
- BINDNV - BIND NEW PACKAGE VERSION AUTHORITY
  - Valid Values
    - BIND: a user with BIND authority on a package can create a new VERSION
    - BINDADD (default): only users with BINDADD authority can add new VERSIONs of a package.
- CDSSRDEF - CURRENT DEGREE by DEFAULT
  - Default degree of parallelism for plans and packages without CURRENT DEGREE specified.
  - Valid values 1 or ANY

Set to 1 and use CURRENT DEGREE per package level to selectively implement parallelism.

- CMTSTAT - DDF THREADS
  - ACTIVE / INACTIVE after it successfully commits or rolls back and holds no cursors
- DESCSTAT - DESCRIBE FOR STATIC
  - YES/ NO build a DESCRIBE SQLDA when binding static SQL statements.

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM - Options Affecting SQL

- DB2 precompiler options, DB2 subsystem parameters (set through DSNZPARM), bind options, affect how SQL statements can be composed or determine how SQL statements are processed.
- DYNRULS - USE FOR DYNAMICRULES
  - use the application programming defaults that are specified on installation
  - use the values of the DB2 precompiler options for dynamic SQL statements that are bound by using DYNAMICRULES bind, define, or invoke behavior.
  - YES to use the application programming defaults for these fields regardless of the DYNAMICRULES BIND option:
    - DECIMAL POINT IS
    - STRING DELIMITER
    - SQL STRING DELIMITER
    - MIXED DATA
    - DECIMAL ARITHMETIC

SYSIBM.SYSPLAN

Column name	Data type	Description	Use
DYNAMICRULES	CHAR(1) NOT NULL WITH DEFAULT	The DYNAMICRULES option used when the plan was last bound: B BIND. Dynamic SQL statements are executed with DYNAMICRULES bind behavior. blank RUN. Dynamic SQL statements in the plan are executed with DYNAMICRULES run behavior.	G

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM's - Packages

- REFSHAGE - CURRENT REFRESH AGE default
  - Enables query rewrite for REFRESH DEFERRED materialized query tables (MQTs).
  - Default = 0 - disable query rewrite for REFRESH DEFERRED materialized query tables (MQT)
  - Use ANY if you are using MQTs
- MAINTYPE - CURRENT MAINT TYPES default
  - default value for the CURRENT MAINTAINED TABLE TYPES FOR OPTIMIZATION when CURRENT REFRESH AGE is set to ANY
    - SYSTEM: DB2 to use system maintained MQTs for query rewrite.
    - USER: DB2 to use user maintained MQTs for query rewrite.
    - ALL: DB2 to use both system and user maintained MQTs for query rewrite
    - NONE does not allow DB2 to use any MQTs for query rewrite.
- MAXKEEPD - MAX KEPT DYN STMTS
  - number of statements kept in Dynamic Statement Cache
- MAX_CONCURRENT_ PKG_OPS
  - number of automatic bind requests that can be processed concurrently.
  - Default value = 10.
  - AUTO BIND = NO → MAX_CONCURRENT_PKG_OPS has no effect.
- OPTHINTS - OPTIMIZATION HINTS
  - Guide the optimizer in choosing specific access paths.
  - Default = NO.
  - Optimization Hints changed in DB2 10 (SYSIBM.SYSQUERY tables)

© Copyright 9/15/2011 BMC Software, Inc

Others

- PARAMDEG - MAX DEGREE
  - Maximum number of parallel tasks DB2 subsystem can use
  - Default =0, → DB2 chooses a maximum degree based upon your environment.
- STORMXAB - MAX ABEND COUNT
  - Number of times stored procedure abend before requests to the stored procedure are rejected.
  - Zero (0) → After the first abend from a stored procedure, all subsequent calls are rejected until the procedure or UDF is started again.

© Copyright 9/15/2011 BMC Software, Inc

Others

- APPENSCH
  - system default application encoding scheme
  - Default EBCDIC
- AUTHCACH - PLAN AUTH CACHE default
  - authorization cache to reduce the amount of time DB2 spends checking authorizations for frequently executed plans.
  - 0 no cache
- CACHEDYN - CACHE DYNAMIC SQL
  - Cache prepared, dynamic SQL statements.
  - Impact on the size of your EDM POOL..
- CACHEPAC - PACKAGE AUTH CACHE
  - Package authorization caching
  - To reduce the amount of time DB2 spends checking authorizations for frequently executed packages.
- CACHERAC - ROUTINE AUTH CACHE
  - Stored procedures and user defined function caching
  - To reduce the amount of time DB2 spends checking authorizations for frequently executed stored procedures and UDF

© Copyright 9/15/2011 BMC Software, Inc

DSNZPARM Connectivity

REMOTE QUERY PERFORMANCE

- Large number of rows to fetch→ request multiple query blocks per transmission
  - Limited by query-block limit set for the server (DSNZPARM EXTRASRV)
  or
  - the query-block limit set by the client (DSNZPARM EXTRAREQ)
- Default is 100 query blocks per transmission
  - Default=Maximum
  - Could negatively affect the performance of your applications

As always your mileage may vary, depending on your workload mix.

© Copyright 9/15/2011 BMC Software, Inc

WAS/DB2 - Important DSNZPARM Values

- CONDBAT – Max. # of distributed connection into DB2 system
  - includes inactive and active connections, may be large
  - DB2 queues DBAT requests to become active up to CONDBAT
- MAXDBAT - Max # database access threads (DBATs) that can be active concurrently.
  - In many installations, max. value determined by available storage in DBM1 (check IFCID 225)
  - Set this value conservatively
- CMTSTAT INACTIVE – make a thread inactive after it successfully commits or rolls back and thread does not hold resources
  - prerequisite for sysplex workload balancing
  - inactive connections use less storage and free up DBM1 resources
- IDTHTOIN – time in sec an active server thread remain idle before it is canceled
  - inactive connections are not subject to idle thread timeout
  - Strongly recommended to not set to 0 – disable, default works well

© Copyright 9/15/2011 BMC Software, Inc

DDF-Related ZPARM Review

- CMTSTAT
  - DDF Threads Mode: Active/Inactive
- IDTHTOIN – Idle Thread Timeout
  - Idle thread timeout (DRDA), to clean up "active" DBATs doing no work
  - not recommended for ACTIVE mode
- TCPKPALV – TCP/IP Keepalive
  - Coordinate with IDTHTOIN to help detect network failures
- POOLINAC – Pool Thread Timeout
  - to reduce the number of pooled DBATs to match the workload (INACTIVE mode)
- ACCUMACC and ACCUMUID
  - to specify DDF accounting rollup to reduce record volume
- MAXTYPE1 (PP) – Max Inactive DBATs
  - to allow type 1 inactive processing for private protocol
- KEEPDYNAMIC(YES) / MAXKEEPD
  - to control storage usage in DBM1(driven by applications like SAP)
- EXTRAREQ / SRV – Extra Blocks REQ / SRV
- And of course:
  - MAXDBAT – Max Remote Active
  - CONDBAT – Max Remote Connected

© Copyright 9/15/2011 BMC Software, Inc

DB2 10 ZPARM changes

- PARMs removed or deprecated
  - Take into account skip-level migration
- New PARMS to support FlashCopy explotation (NFM)
- New PARMS to support new separate security
- New PARMS for SMS allocation of catalog and directory objects
- New PARMS for supporting Inline LOBS
- RRF Enable/Disable
- Default partition segsize
- EDMPOOL (limit below the bar) defaults to 0
  - If you are migrating to Version 10, override the default setting with the EDMPOOL value that you currently use in the previous release. (!)

© Copyright 9/15/2011 BMC Software, Inc

Summary

- As Systems Administrator you are
  - dependent upon database, application & subsystem design & practices
- As Systems Administrator you need to
  - Understand and avoid capacity issues
  - Understand & exploit capabilities
    - Substantial improvements at release level and also between releases
- As Database Administrator you are
  - Dependent upon subsystem installation and maintenance
- As Database Administrator you need to
  - Be involved in installation parameters
  - Know installation parameters can change
- Continuous Investment Required

© Copyright 9/15/2011 BMC Software, Inc

Summary

"Damn, everytime I think I got things under control they raise the bar on me."

© Copyright 9/15/2011 BMC Software, Inc

Conclusions and Questions

SQLADRIA SEMINAR
Zagreb, 23 September 2011

A reconstruction of 'AS IS' data model

Mateja Jankovič
Informatika d.d.

Sponsored by

The Reconstruction of ‚AS IS' Data Model

Mateja JANKOVIČ
INFORMATIKA d.d.

Agenda

- Data Architecture
- The Reconstruction of ‚AS IS' Data Model
- Informatika Case
- Summary

Data Architecture

• • •

Data Architecture

- Data architecture is **one of the pillars** of an Enterprise Architecture
  - which describes the structure of an organization's logical and physical data assets and the associated data management resources. (TOGAF)
- Designing a data architecture is **a complex process**
- **All organizations have data architecture**
  - Some are **better documented or used** than the others
  - Documents, which are not used or understood, is **not usefull for the organisation**

Data Architecture

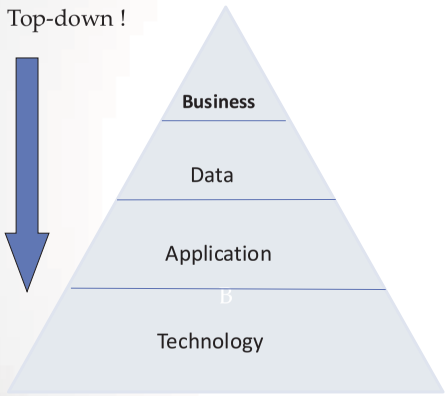

Enterprise level
- Business glossary
- Conceptual data model
- Logical data model

Business area level
- Business glossary
- Conceptual sub-model
- Logical sub-model

Aplication level
- Subject area in a logical data model (e.g. Data service, ETL data model, sync data model. Migration data model...)

Technology level
- Physical data model, DDL-i, DB2, Oracle...

Initiate a Data Governance

- A lot is written about data governance:
  - what it entails
  - who should be involved
  - and how it should be administered
- For many who are trying to initiate data governance in their organization this can be a confusing topic

Data Governance

- Is data governance about data?
  - It's about **how companies view their data**
  - A **data model IS NOT a database**, much like a blueprint is not a house.
- View of Metadata
  - **Without data**, there is no governance
  - Dynamic world of business need to better **define data**
  - An accurate, transparent and accessible **view of their metadata** should be created
- People must know how the program of **data governance will benefit** them specifically in their job
  - reduce the risk of project failures
  - improve communication
  - increase the value of strategic information

The Value of Data Modeling is Expanding

- Business demand for data and analytics
  - Data explosion
  - Need information and analysis
- Business efforts constrained by
  - Navigating data silos
  - Not knowing what data they have
  - Data quality and inconsistency
- Data modeling capability expands
- Data modeling role expands

Real Life

- ‚Do more with less'
  - More databases to manage
  - More platforms to support
  - Fewer people to do the work
- Technology is changing
- Business is involved
  - Visibility is needed
  - Provide information to users in their ‚language'
    - Different tools (BI, ETL, DDL,...)
    - Business users like Excell, Word, Web tools
  - Business doesn't know the mechanics, but they know the definitions
  - IT folks are enablers

Data Modeling Capability Expands

- In the past: IT oriented development
  - System-programming point of view
- Now: critical need for data architecture and blueprint
  - But ‚old school' is still valuable today: ‚Let's design before we bulid!'
  - Data modeling is usefull also in new situations
  - Not a technical challenge, but people challenge
- Data modeling capability expands
  - Team development
  - Data to information workflow
  - ETL, DW and BI
  - Cross project

Using Data Models can Reduce the Cost and Risk

Some examples:

- **Staff turnover**
  - Training
  - Importancy of accurate documentation
- **New development**
  - Communication
  - Sources for BI, migration, synchronization
- **Data integration**
  - Trend – to centralize data from variety of sources across the organization
  - Data reuse
- **Improving data quality**
- **Reduce redundant data and processes**

Mistakes we Make

- Forgetting that an enterprise architecture is **a living framework**
- Thinking that data models are **only about databases**
  - Logical and physical data models can support more than databases
    - Logical models are independent of technology and data requirements are documented better than in any text-based documentation
    - Physical models allow team members to see how business concepts are implemented without having to provide access to production systems.
- Thinking of them as **"your" models**
  - Models should be managed by a partnership of modelers and the business

A Data Model can be a Filter

- A way to communicate

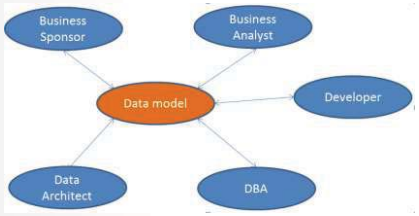

- A data model can add
  - Focus (subject area)
  - Visualization
    - different views for different audience
  - Translation (different DBMS, Excel, XML…)

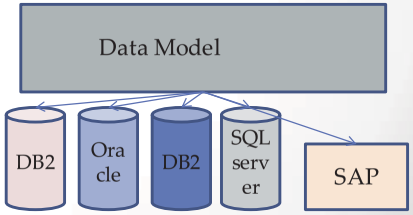

Using a Tool for Data Governance

- Single repository to store all of your data model assets
- A collaborative environment for multiple modeling teams
- Metadata storage for: multiple models, multiple DBMS, multiple tools, multiple audiences

Multiple models

Multiple DBMS

Multiple tools

Multiple audiences

The Reconstruction of 'AS IS' Data Model

…

Improve your Data Governance

- We simply follow a pattern of behavior that was established when **our business was significantly different**
  - Reasses what our business is today
- Identify existing assets
  - Hard to begin a new initiative
  - Profit by identifying existing assets
- You are already managing data
  - You are doing it informally

Modeling ‚AS IS'

- There is **no need to start from scratch** with a blank sheet of paper
- **Modeling the ‚AS IS' data architecture** can be extremely useful, and it can certainly show where things are sub-optimal
- If you want to **make future improvements**, you will need more than just good models
- **Most of the issues around improving data collection, usage and governance are non-technical!**

Data Architecture

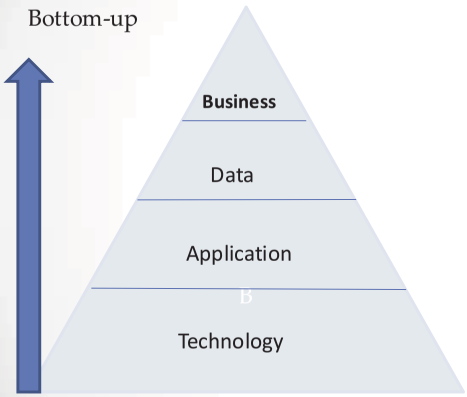

Enterprise level (EDM)

Business area level

Aplication level

Technology level

- Physical data model, DDL-i, DB2, Oracle...

Technology Support - Data Modeling Tools

- Forward Engineering (FE)
- Reverse Engineering (RE)
- Complete Compare (CC)
- Reporting
- Model Management

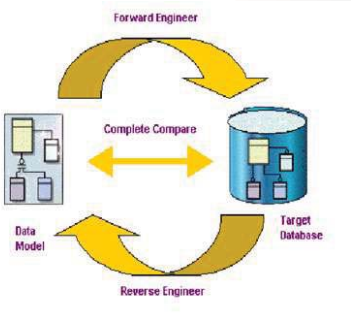

Data Architecture is Complex

- **A real enterprise has a very complicated data architecture**
- Most of the data are held in **large legacy or package systems**
- **You are already managing data**
- The work involved in **defining and implementing data governance** in an organization is not trivial, and it **requires time and money** to be successful
- **Estimating the return on investment** for improving data quality has always been hard

Informatika Case

...

Data Model Baseline

- Own development on basis of users requirements 89-92
- Initial data model(s) in Knowledgeware ADW
- For some business areas CA ERwin tool is used (DW,...)
- Data modeling procedure and data model change management is informal
- There was no migration done from ADW to ERwin, so most of the models exist only in a paper form, some of them exist in older versions in/of CA ERwin

A History of Data Modeling

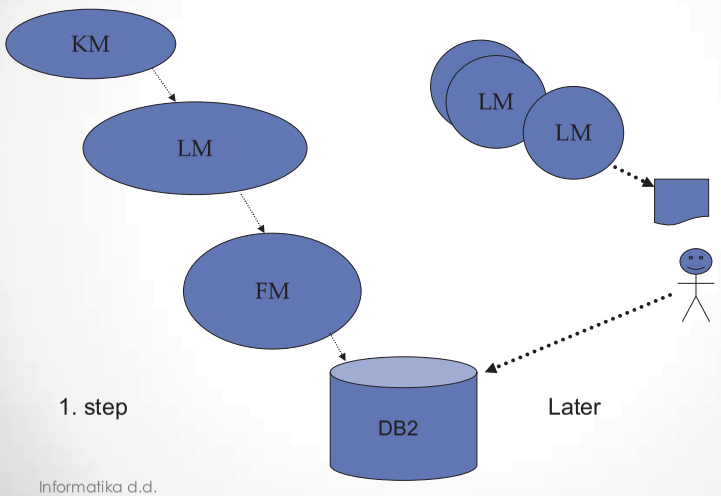

Some Dilemmas before Defining a Project

- How tools can help
- Provide the right amount of information
  - Don't use/expect too much information
- No impact on existing system!
  - Documentation, analysis and recomendations
- How to draw line
  - Roles and responsibilities

A Project: The Reconstruction of Data Model ‚AS IS'

Main objectives of the project:

- **Building an infrastructure** for data modeling (tool, training)
- Collect and analyse **existing documentation**
- **Physical data model reconstruction**
- Transformation from physical data model to **logical data model**

A Reverse Engineering with a Tool

- Tools used:
  - CA ERwin Data Modeler (V7)
  - QMF
  - Excell
    - Templates to be filled by analysts
- An **initial** data model
  - Filters
  - Grouping
  - Useful information

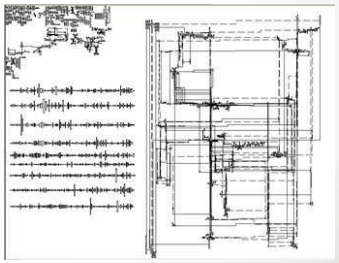

Data Model can be Improved by Additional Information

- Subject area
- Description, property
- Primary key
- (logical) Relation
- ....

BUT find an optimal way what to use/collect

Filtering and Grouping

- Filtering and grouping is important to organise further activities
- Initially naming standards can be used for mass filtering and grouping (naming standards for subsystem, creator, database, tablespace...)
- Additional activities are needed to organise subject areas

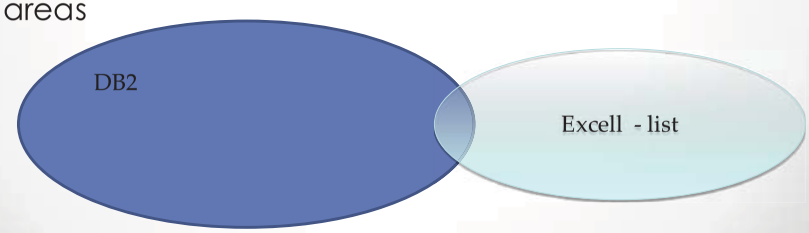

Creating Primary Key

- PK are not defined for all tables
- Usualy (unique) index is defined
  - Missing PK can be added manualy to the model (and database)
- Inability of defining PK

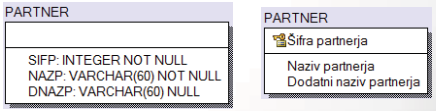

Creating Logical Relations

- Relations (RI) in DB2 are not always defined
- Missing (logical) relations can be added **manualy** to the data model
  - easy to say!
- Some technical reasons for inability of defining logical relations

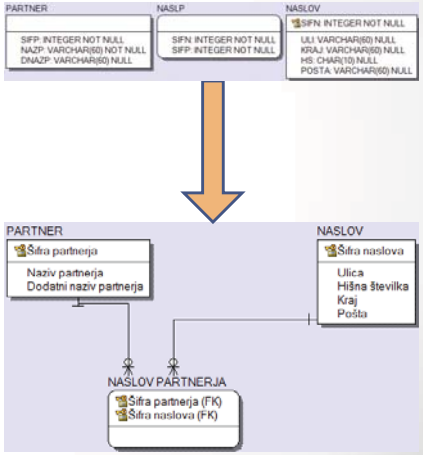

The Project Flow

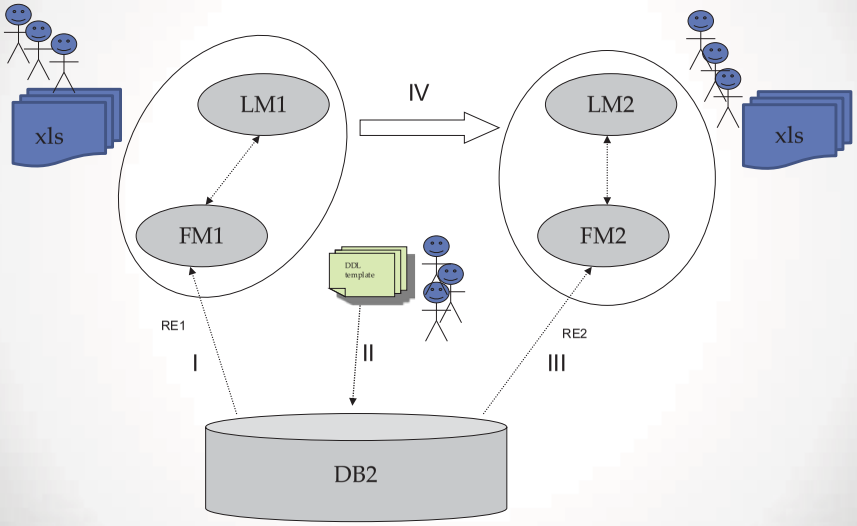

Important for the Project

- Support of the management
- Project team
  - Core team
  - Cooperation of collegues
- Adjust project plan – it's time consuming
- Provide the right amount of information
  - Find an optimal way what to use/collect
  - Step by step

At the End of the Project

- Data architecture in practise doesn't always proceeded from the "**top down**" or **"bottom up"**, the data architecture is usually developed from the "**middle-out**"
- It may be appropriate to undertake some work to "**fill in the gaps**"
- The data architecture may **never be "complete"** for the whole enterprise
- **Consistent approach and context for modeling activities**
- A project established a new baseline
- It should become a process

Summary

...

Summary

- Dynamic world of business needs to better define data. There's a critical need for data architecture and blueprint. We often forget that it is a living framework. An accurate, transparent and accessible view of metadata should be created. The value of data modeling is expanding. Both logical and physical data models can support more than databases.
- Usually an enterprise has a very complicated data architecture. Most of the data are held in large legacy or package systems.
- The work involved in defining and implementing data governance in an organization is not trivial. It's hard to begin a new initiative, but **we can profit by identifying existing assets.**
- There is no need to start from scratch with a blank sheet of paper. Often a simple redesign of a current process can deliver a desirable result. Modeling the ‚AS IS' data architecture can be extremely useful.
- For Informatika d.d. the reconstruction of existing data models was/is a challenge

References

- www.erwin.com
- www.information-management.com
- www.linkedin.com/groups

THANK YOU!

Mateja JANKOVIČ
mateja.jankovic@informatika.si

SQLADRIA SEMINAR
Zagreb, 23 September 2011

Process of Change Migration to Production Environment

Tomaž Dobnikar
NLB d.d.

Sponsored by

Postopek prenosa sprememb v produkcijsko okolje

Tomaž Dobnikar
svetovalec področja
NLB d.d.
Sektor za strateško načrtovanje in upravljanje informacijskega sistema banke
Šmartinska 130, SI-1520 Ljubljana
T:+386 1 520 73 92, M:+386 367 531, F:+386 1 541 44 60, E: tomaz.dobnikar@nlb.si
www.nlb.si

Vsebina

1. Upravljanje IT v NLB
2. Dokumentiranost delovnih postopkov
3. Izvajanje posegov v produkcijsko okolje

Upravljavski center za informacijsko tehnologijo v NLB

Interni dobavitelj strojne in programske opreme za podporo bančnim poslovnim procesom v NLB

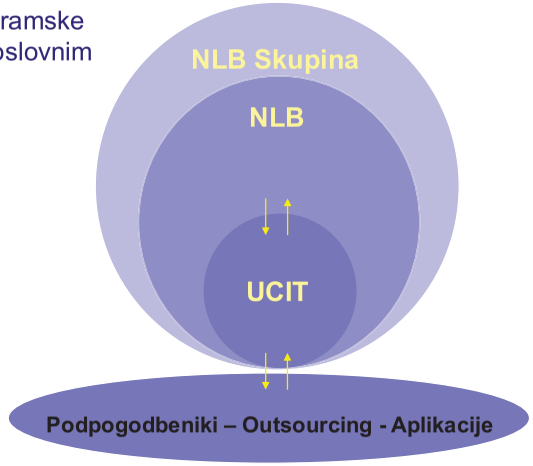

Organizacija

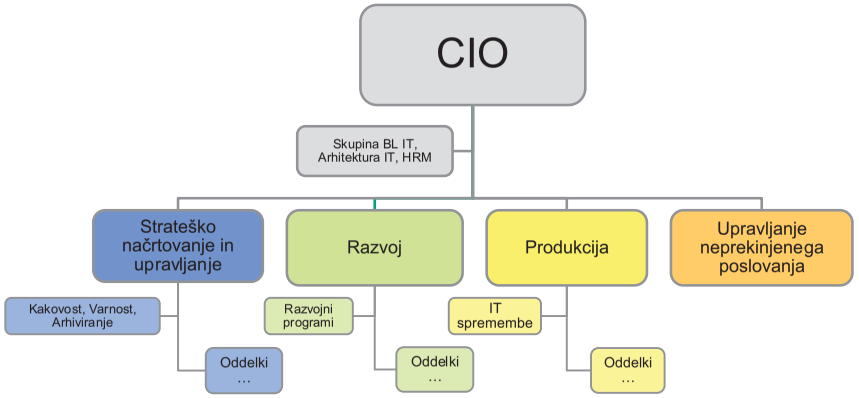

Procesni model

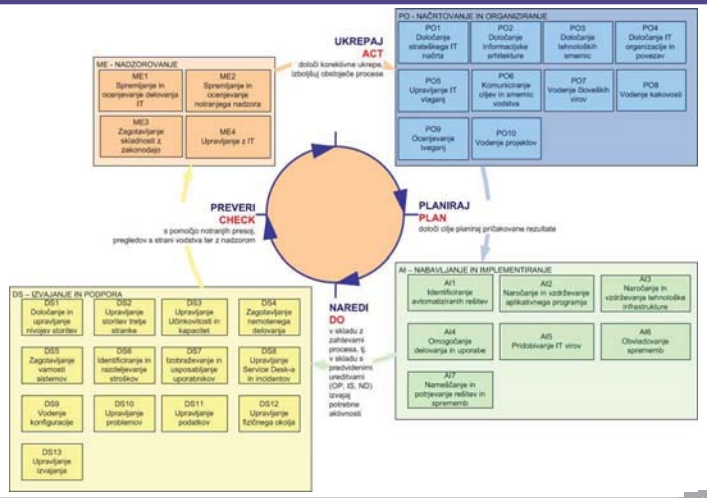

Sistem uravnoteženih kazalnikov (BSC)

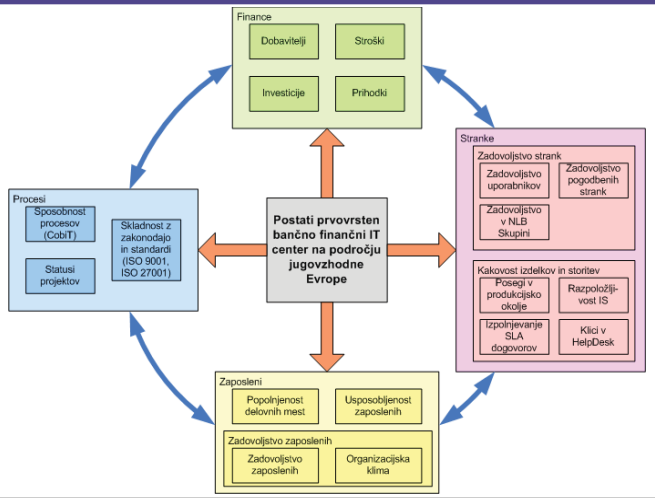

Skladnost z ISO 9001

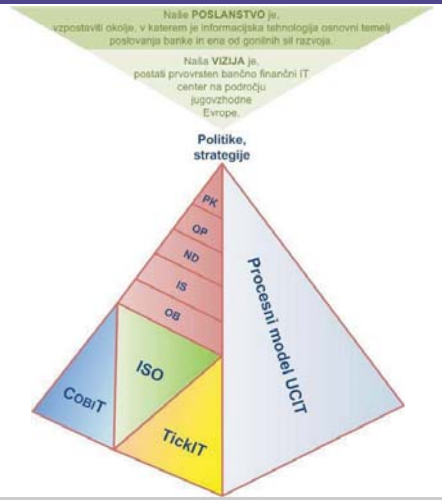

Življenjski cikel razvoja (in sprememb)

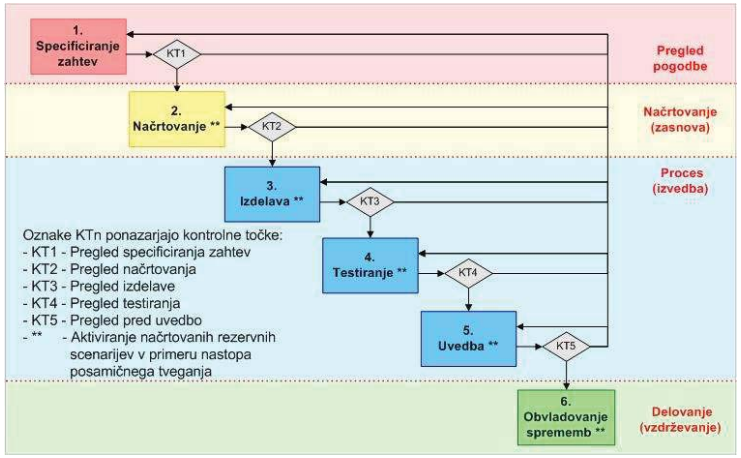

Upravljanje IT sprememb

1. Predpisan postopek uvajanja IT sprememb v produkcijsko okolje
2. Določena skupina za IT spremembe
3. Pripravljen letni koledar IT sprememb
4. Opredeljeni tipi IT sprememb
   a) Redni posegi
   b) Izredni posegi
   c) Standardni posegi
   d) Intervencijski posegi
5. Določeni koraki za izvedbo IT spremembe
6. Umeščenost izvajanja IT sprememb v druge postopke
7. Matrika odgovornosti

Koledar IT sprememb

KOLEDAR 2011 - 1. izdaja z dne 24.11.2010

Januar	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
POSLOVNO KRITIČNI DNEVI					P					DP				P			PP	DB		PP								DB			P
APL																											A		A	A	
SYS																															
IPL									IPL																						
Varnostni popravki SRV																				V	V			V	V	V	V	V			

Februar	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
POSLOVNO KRITIČNI DNEVI	PP			P					DB					PP	P	PP		DB			PP							DP
APL																												
SYS										S		S	S															
IPL													IPL															
Varnostni popravki SRV																	V	V			V	V	V	V	V			

Marec	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
POSLOVNO KRITIČNI DNEVI				P				DB			PP			PP	P			DB			PP							DB			P
APL																	A		A	A											
SYS			S		S	S																									
IPL																										BCP	TRA				
Varnostni popravki SRV																	V	V			V	V	V	V	V						

Legenda:

	Vikend ali praznik
A	APL - aplikativni posegi
S	SYS - sistemski posegi
I	IPL - ponoven zagon sistema / premik ure
B	BCP test
V	Varnostni popravki strežniki

Poslovno kritični dnevi

DB	Trajniki/direktne obremenitve (DB): 8., 18. in 28. v mesecu in prvi naslednji delovni dan, če je ta dan nedelovnik
P	plače zaposlenih v državnih ustanovah (do 5. v mesecu), plače 10. in 15. v mesecu in pokojnine
DP	kombinacija obdelav DB in izplačila plač ali pokojnin ali plačilnega prometa
PP	povečan obseg plačilnega prometa: plačila davkov (pravne osebe) - okoli 5. in 20. v mesecu, 15. v mesecu (oz. dan prej/potem, če je 15. nedelovni dan) in dnevi pred prazniki

Tipi IT sprememb

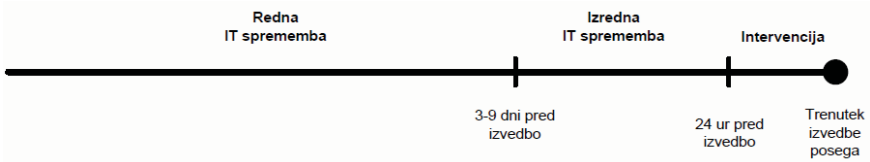

Standardna IT sprememba je IT sprememba, ki se redno oz. pogosto izvaja in praviloma zanjo obstaja utečen in dokumentiran postopek npr. v obliki operativnih navodil in je izvedba v splošnem dobro obvladovana. Standardne IT spremembe so praviloma nizko tvegane in sicer predvsem zaradi utečenega in večkrat preizkušenega postopka.

Koraki uvajanja IT sprememb

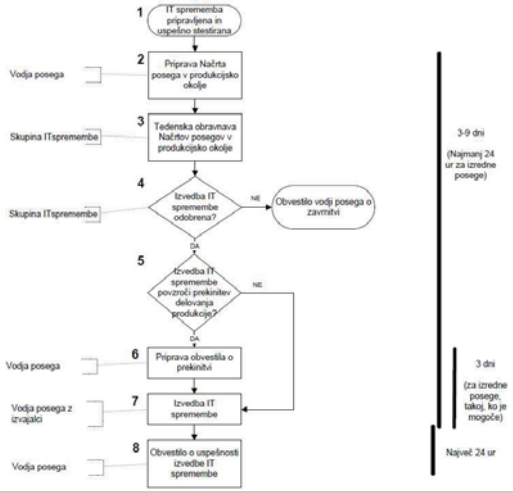

Umeščenost izvajanja IT sprememb v druge postopke

IZVOR IT SPREMEMBE ⇒ AKTIVNOST ⇓	Razvojni projekt	Vzdrževalni projekt	Aplikativna sprememba (zahtevek)	Sistemska sprememba	Nedelovanje produkcije
1. Izvedba kontrolnega pregleda	X	X	X		
2. Priprava in potrditev Dovoljenja za uvedbo dobavljenega sistema v produkcijsko okolje	X				
3. Priprava in potrditev Načrta posega	X	X	X	X	
4. Obveščanje končnih uporabnikov o prekinitvah	X*	X*	X*	X*	
5. Izvedba posega	X	X	X	X	X
6. Pošiljanje obvestila o uspešnosti izvedbe posega	X	X	X	X	
7. Pošiljanje Poročila o začetnem delovanju uvedenega sistema	X	X	X		
8. Pošiljanje Poročila o intervenciji					X

* v primeru, da poseg povzroči prekinitev v produkciji

Matrika odgovornosti

NOSILEC AKTIVNOSTI ⇒ AKTIVNOST ⇓	Izvajalec	Vodja posega	Vodja oddelka	Programski vodja	Direktor razvoja	Skupina IT spremembe	Direktor produkcije
1. Priprava Načrta posega, uskladitev z izvajalci	S	P					
2. Verifikacija Načrta posega			P	P*	P**		
3. Obravnava Načrtov posegov						P	
4. Odobritev izvedbe IT spremembe						S	P
5. Pošiljanje obvestila o prekinitvi na Help Desk		P					
6. Izvedba IT spremembe	S	P					
7. Evidentiranje statusa izvedbe IT spremembe		P					
8. Vodenje evidence standardnih sprememb			P				

*za posege, ki izhajajo iz projektov
**za aplikativne posege

P – primarna odgovornost (odgovornost za izvedbo); S – sekundarna odgovornost (odgovornost za sodelovanje)

Hvala!

KIT

KIT@nlb.si