



V9 Migration Experiences @ KBC

Ronny Vandegehuchte



KBC Configuration

- 50 subsystems (15 in production)
- Datasharing (3 way)
- 24X7
- sandbox, development, acceptance, production

..

Environment	DB2 V9 CM	DB2 V9 NFM
Sandbox	01/02/2009	15/04/2009
Development	17/03/2009	20/05/2009
Acceptance	31/03/2009	29/05/2009
Production	16/05/2009	13/06/2009

V9 is using more 32K for SORTing

- Development subsystem(s) :
 - 4K sort files isolated in BP3
 - 3 files, each size of 3000 cyls each
 - 3000 buffers
 - BP Thresholds :
 - VP SEQUENTIAL=100
 - DEFERRED WRITE = 90
 - VERTICAL DEFERRED WRT = 85, 0
 - 32K sort files isolated BP32K3
 - **1 file, 100 cyls**
 - **50 buffers**
 - BP thresholds
 - **Defaults**
 - Rarely used.....

V9 is using more 32K for SORTing

- After migrating development systems to CM

- SQLcodes -904 32K resource

```
DSNP007I -SDTC DSNPXTN0 - EXTEND FAILED FOR 995  
SSYSDTC.DSNDBD.SWDSBTC.DSN32K01.I0001.A001.  
RC=00D70027  
CONNECTION-ID=BATCH, CORRELATION-ID=PD0Y1ODO,  
LUW-ID=BEBBCM01.A2SDTC.C3E6CF9D5628=4763
```

- Increased BP32K3 size (50->**500->2500**)
- Increased to **3 workdatasets, 3000cyls each**
- Changed BP32K3 thresholds to BP3 thresholds
- Got problem....
 - -904 on 32K sort tables, not accessable, no locks, no use
 - Stop/start didn't work, need to recycle DB2 system

..


V9 is using more 32K for SORTing

```
DSNP008I -SDTC DSNPXTN0 - REQUIRED USER-DEFINED DATA 931  
SET NOT AVAILABLE-  
SSYSDTC.DSNDBC.SWDSBTC.DSN32K02.I0001.A002.  
RC=00D70002  
CONNECTION-ID=DB2CALL, CORRELATION-ID=CIIT00D,  
LUW-ID=BEBBCM01.A2SDTC.C4120A6B41FE=107468  
DSNP007I -SDTC DSNPXTN0 - EXTEND FAILED FOR 932  
SSYSDTC.DSNDBD.SWDSBTC.DSN32K02.I0001.A002.  
RC=00D70002  
CONNECTION-ID=DB2CALL, CORRELATION-ID=CIIT00D,  
LUW-ID=BEBBCM01.A2SDTC.C4120A6B41FE=107468
```

- Defined the sort table spaces, `cyls(3000,0)` which is more than 2 Gb for primary
- When we filled 2 Gb (usually 2913 cyl), we tried to extend to the second pageset which hadn't been defined.
- Defined sortfiles as 2900cyls !
- Due to create of 32K3 spaces, they become segmented !

V9 is using more 32K for SORTing

DPC1



Unicode improvements

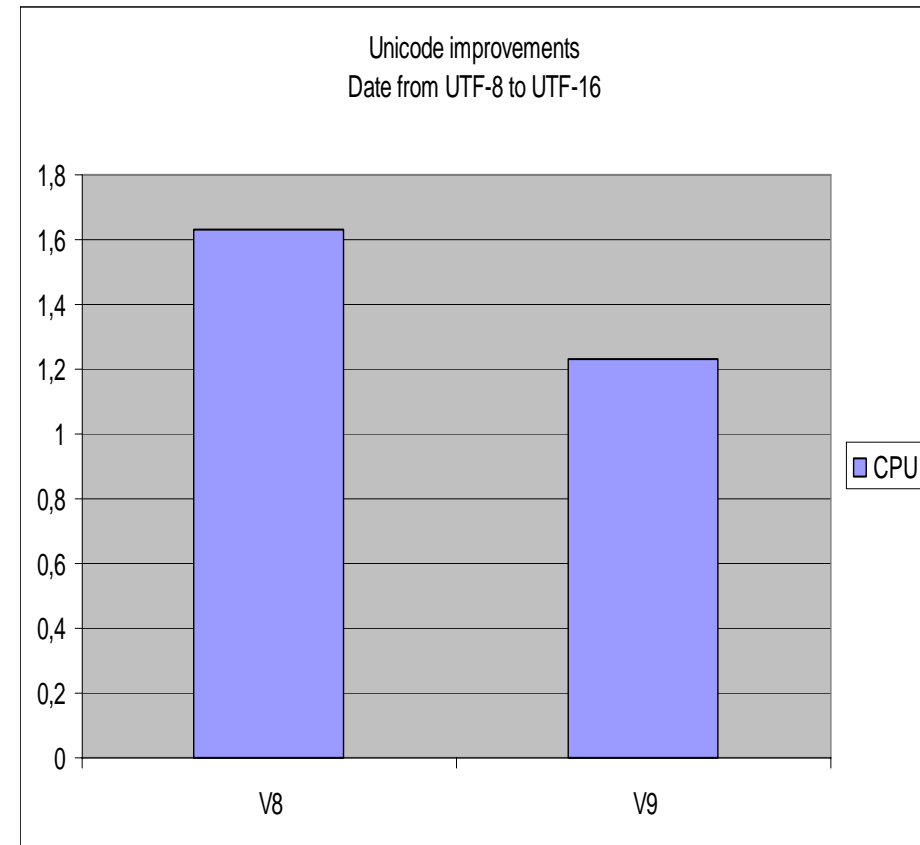
- Date/time/timestamp are stored as UTF-8 in a unicode table
 - Pre V9 :
 - Conversions to UTF-16 are done by calls to z/OS conversions services,....Conversions cost
 - V9 :
 - Conversions to UTF-16 are done by DB2 it self

..

Unicode improvements

- Test case:
 - Table CBTUNIDD has 39 Date Cols, and 40K rows. Select all rows and all columns into PIC N(10) , via an COBOL pgm.
 - 33% CPU decrease
 - Conversions are done in DSNXRTIM instead of CUNMUNI

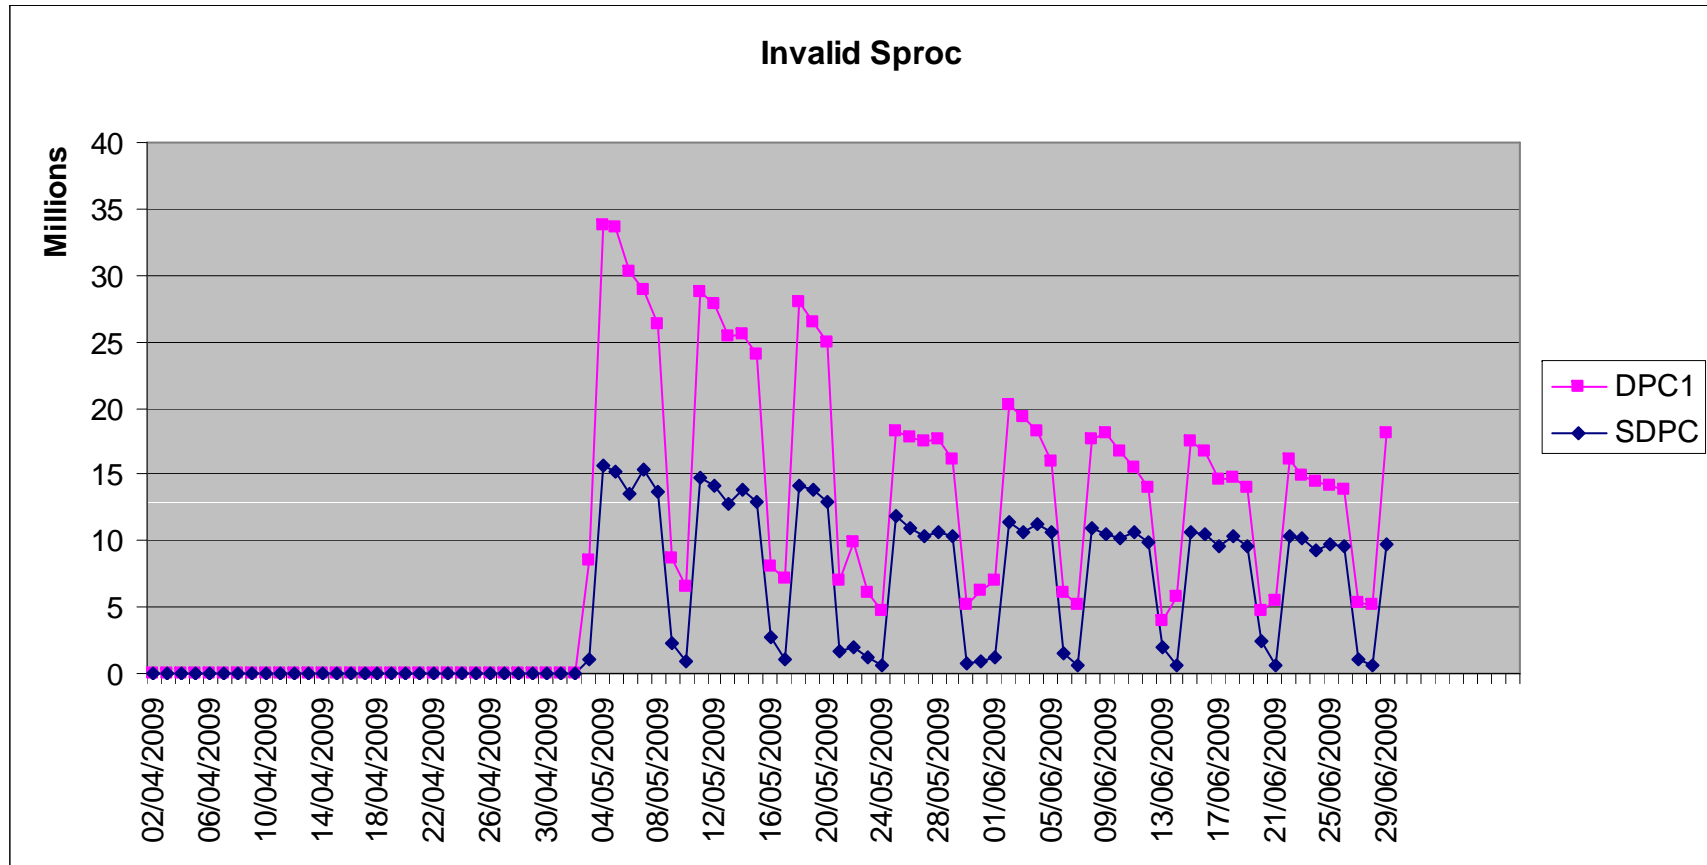
..



Rebind BEFORE migration


- We did get Abend0c4 in program's, on *some* acceptance & production systems
 - in packages that were bound on Pre V8 and even V8 (UK44259)
 - rebind solved all cases
- Wrong query results!!
 - in packages that were bound on Pre V8 and even V8
 - rebind solved all cases
- To avoid problem's after migration :
 - rebind **before** migrating to V9
- ..

Rebind AFTER migration Invalid SPROC's

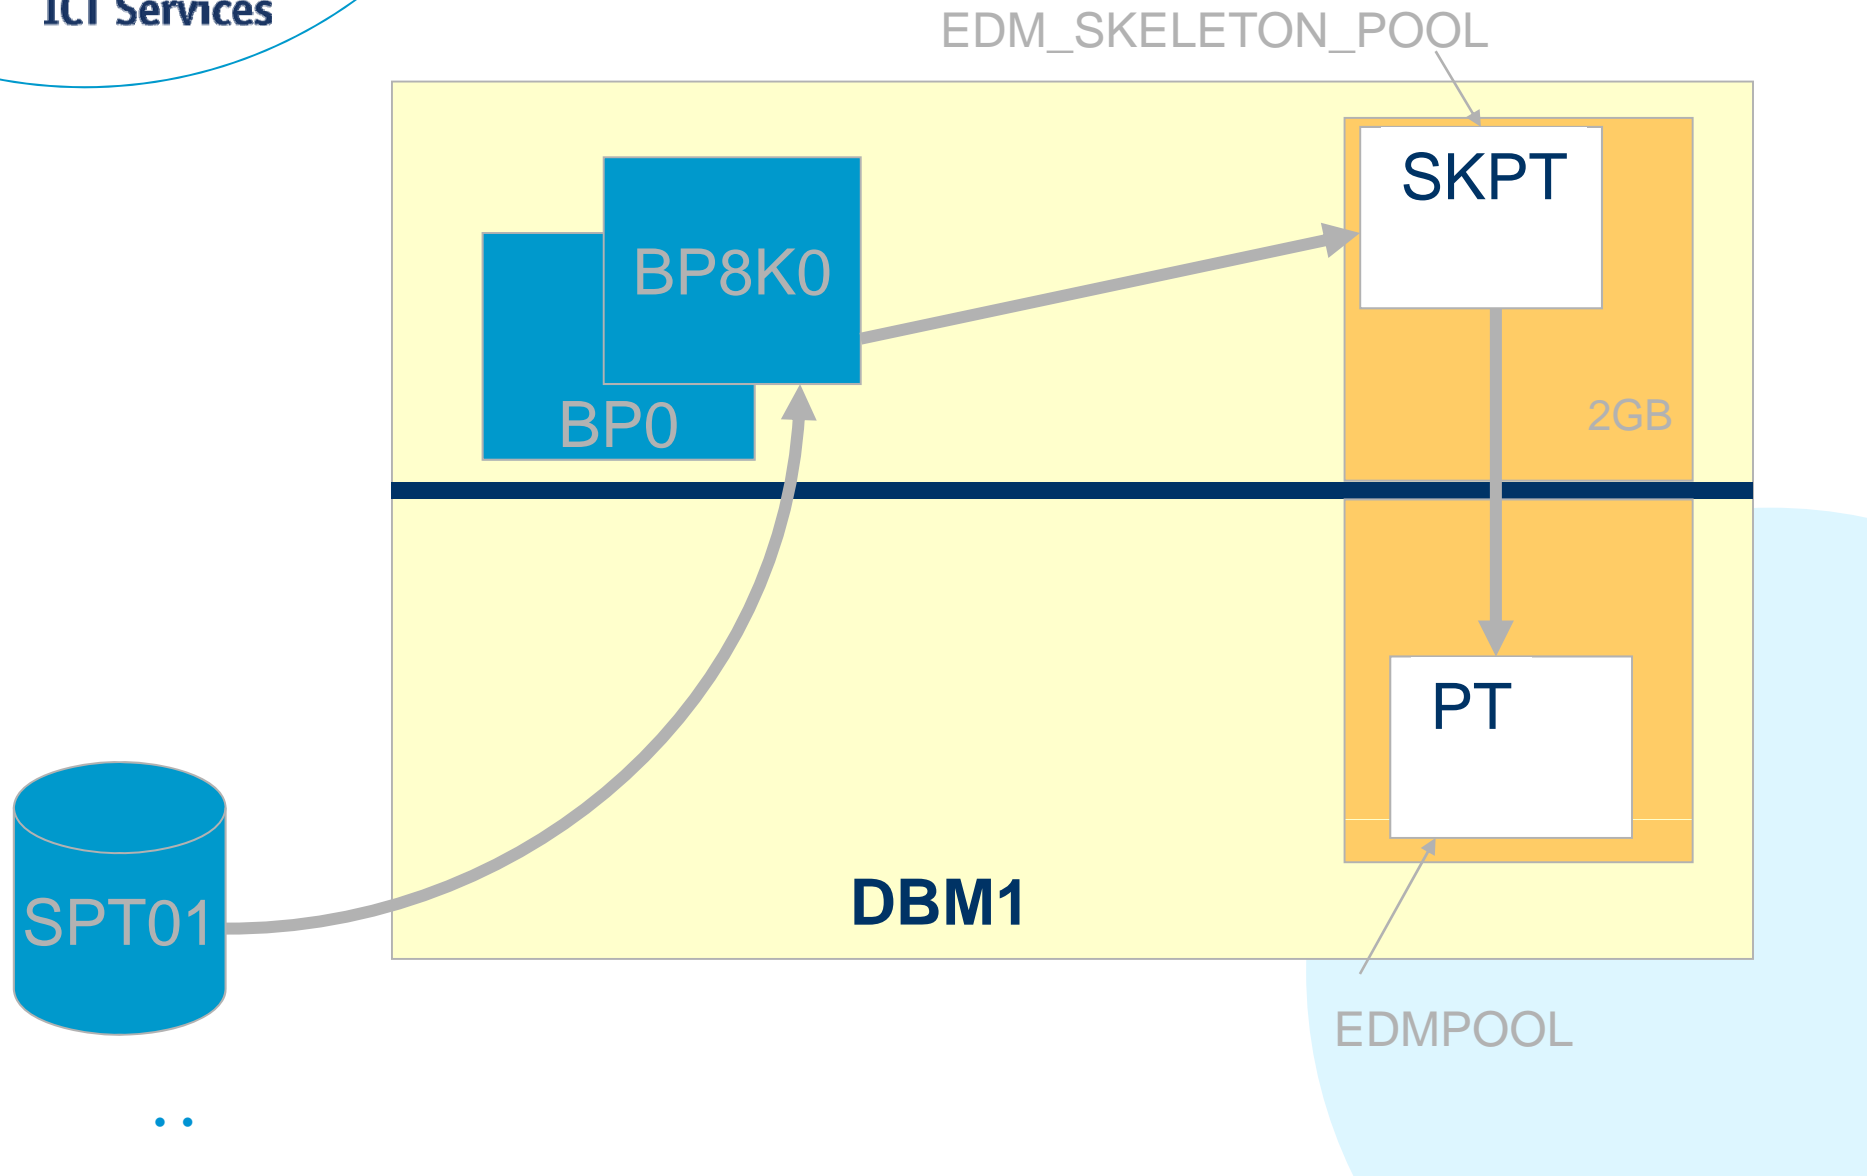


Rebind AFTER migration Invalid SPROC's

- A pgm uses 19% more CPU



EDM Pool Structure V9 (no Rebind)



EDM Pool Structure V9 (no Rebind)

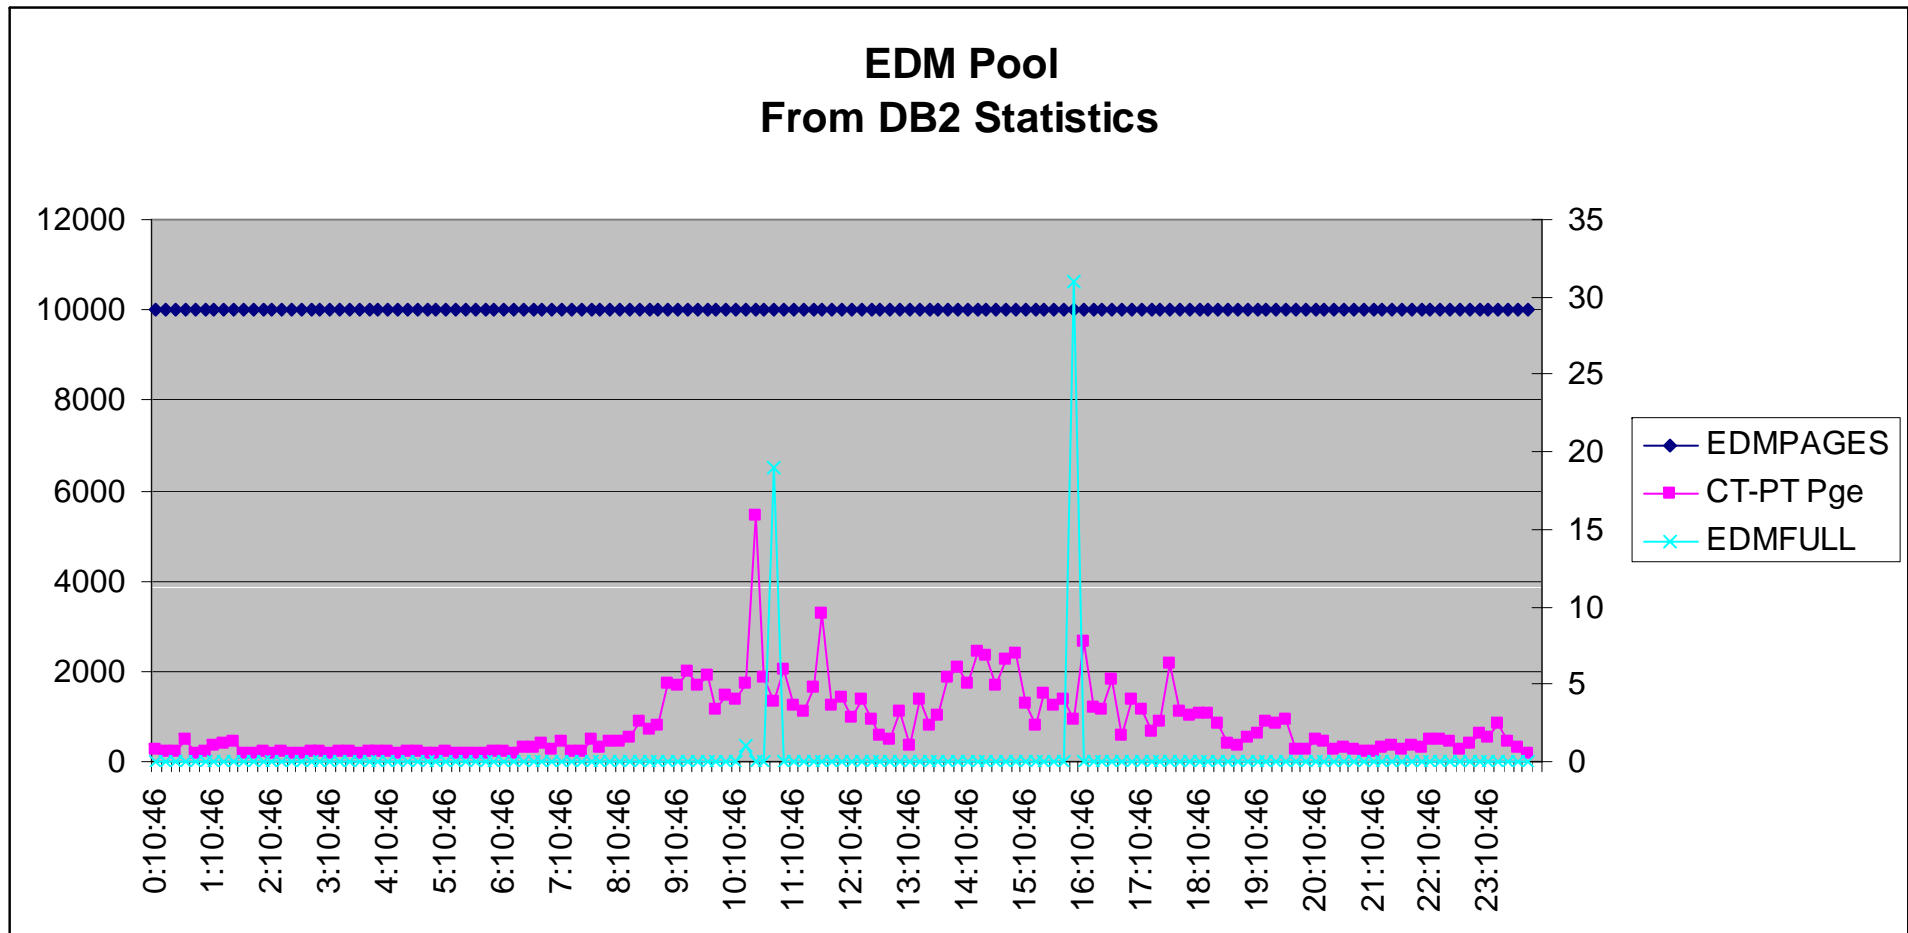
- EDM_SKELETON_POOL :
 - DSNZPARM based on calculations,
 - 6389KB
 - From DB2 stats : $\text{Max}(\text{QISESKPT} + \text{QISESKCT})$
 - 100000KB
- EDMPOOL = 125000KB down to 40000KB
- WHAT IS/WAS THE IMPACT ?
 - Nothing..... Or
 - HIGHer CPU usage in DBM1
 - HIGHer CPU usage in the application
 - HIGHer I/O rates – LOWER BP8K0 hitratios

..

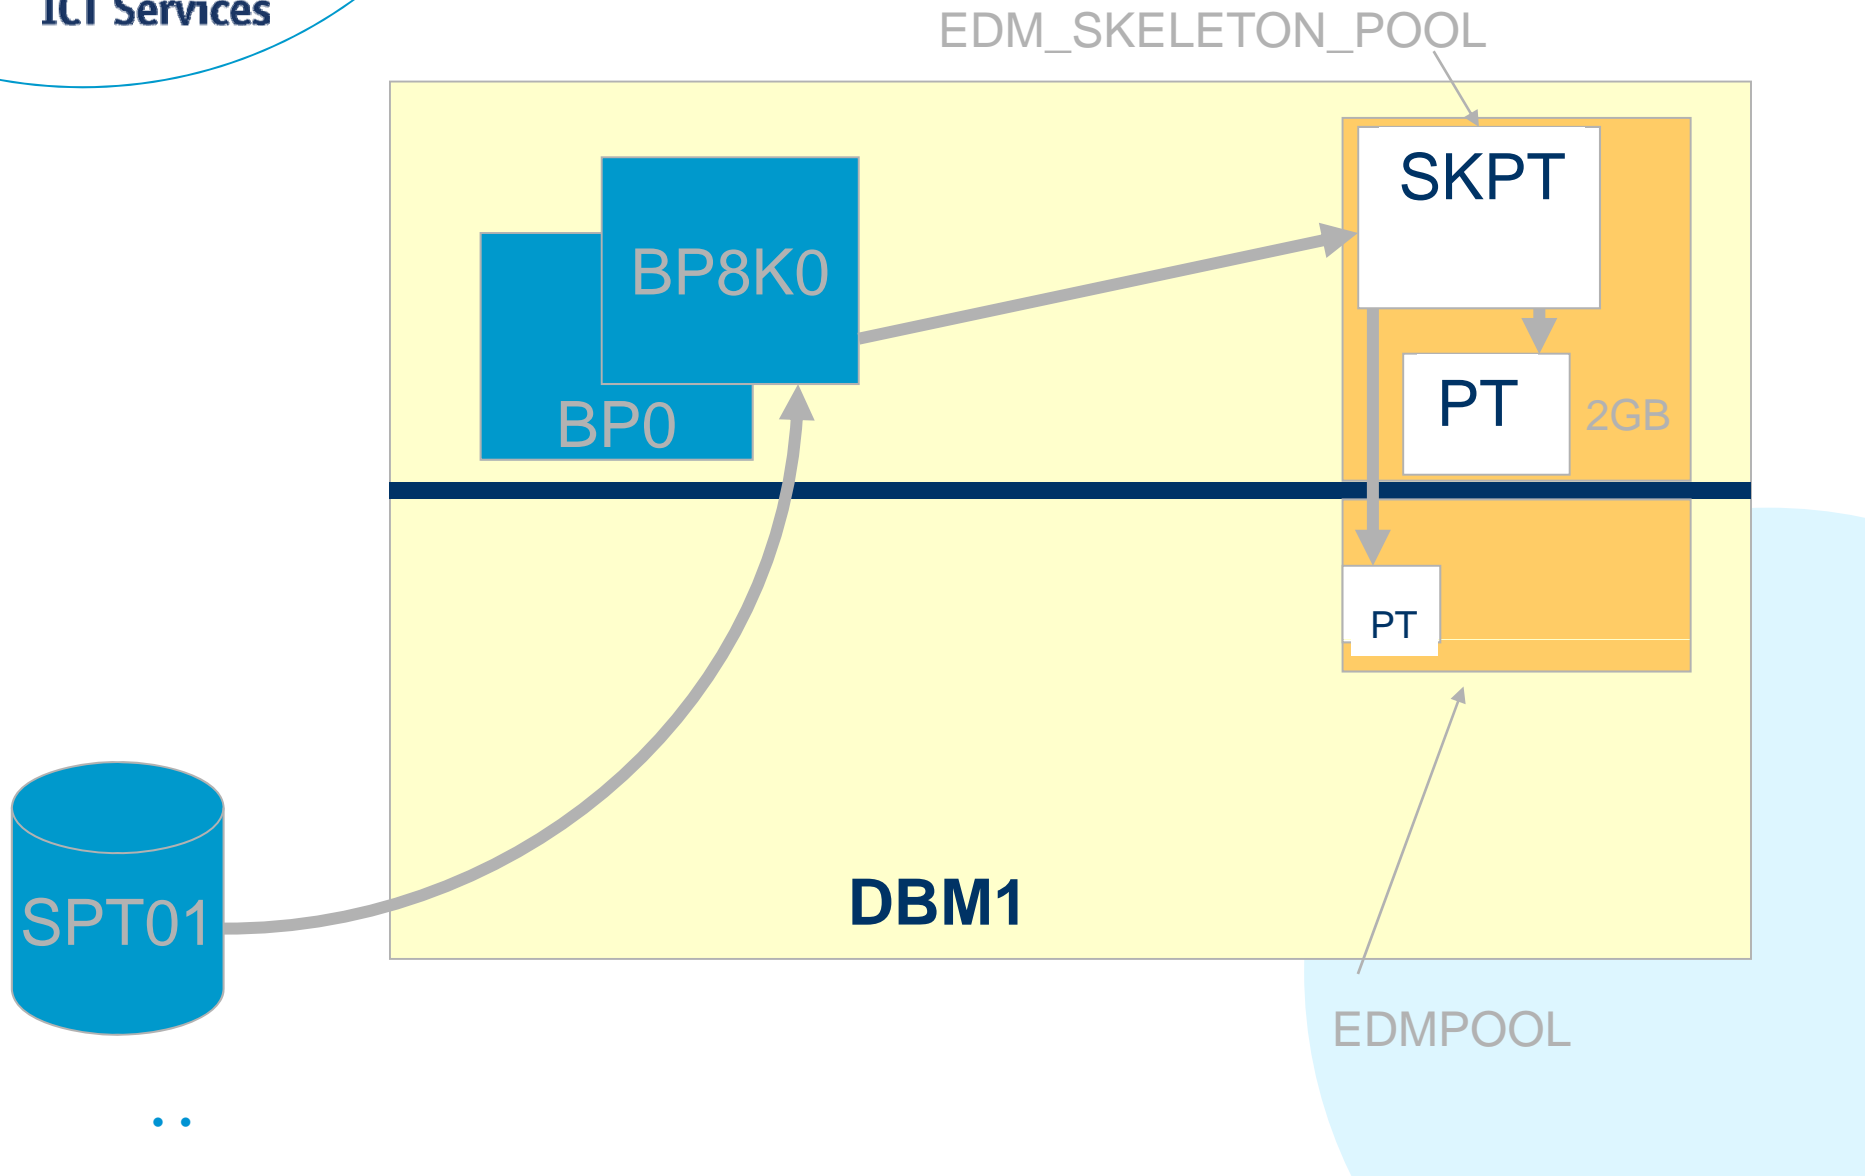


EDM Pool Structure V9 (no Rebind)

- Statistics interval 10min is to big !!



EDM Pool Structure V9 (Rebind)



Increased Logging Activity

- Load table Resume No Replace Log No Part 1
 - Table has 5 NPI's, logging occurs for these NPI's
 - Same # of rows (no changes of the datamodel)
 - V8 : 0.041 GB -> V9 : **8.25 GB**
 - **PMR to IBM**
 - This is caused by the structure modification of the NPI (i.e. delete leaf page). This in terms is a secondary effect of a new algorithm we use to 'build' the NPI during load part replace operation in DB2 9.
 - **APAR PK83683**
UK50265

..

Access Path Changes

- NO mass rebind, risk to high
- Look at what happens
- Solve it case by case

Access Path Changes Static SQL

- Bad access path after first bind in DB2 9
 - Runstats + REBIND : NOT ok
 - Advisory runstats OSC + REBIND : Problem solved

```
RUNSTATS TABLESPACE DF2BR01.SF23060  
TABLE(PIBBRP.TF2BR3060) SAMPLE 5  
COLUMN(CMED,RVERS)  
COLGROUP(CMED) FREQVAL COUNT 10  
SORTDEVT SYSDA SORTNUM 4  
INDEX(PIBBRP.XF2BR30601 HISTOGRAM NUMCOLS 1 NUMQUANTILES 20,  
PIBBRP.XF2BR30600 KEYCARD)
```

Access Path Changes Static SQL

- In most cases better performing SQL
- Less Sorting
- Nested Loop + sort -> Hybrid Join (no sort)
- In some cases use of other IXs
- Other access sequence
- More Multiple IX Access :
 - Sometimes simple rebind solves the problem
 - In other cases Advisory runstats OSC + REBIND
- ..

Access Path Changes Dynamic SQL

- DB2 V8 and Tivoli Directory Server (LDAP):
 - Use of dynamic statement cache
 - Reduce “dynamic” overhead
 - => Trying to guarantee stable performance
 - No regular collection of DB2 statistics
 - Unique collection of DB2 statistics with specific colgroup stats
 - => Trying to avoid accesspath changes

..

Access Path Changes Dynamic SQL


- DB2 V9 CM and LDAP
 - Bad LDAP performance
 - => new collection of DB2 stats with colgroup stats (cfr V8) did not help
 - => collection of stats advised by OSC Statistics Advisor helped in certain cases, but due to data skewing still some problems
 - we bound DSNCLINF with reopt(always)
 - => initial CPU overhead 10%
 - no use of dynamic cache anymore
 - every stmt is reoptimized
 - ➔ but stable performance!
 - ➔ but worse than in V8

Access Path Changes Dynamic SQL

- DB2 V9 NF and LDAP
 - We tried to bind DSNCLINF with reopt(NONE)
=> again bad accesspaths
 - We rebound DSNCLINF with reopt(always)
=> performance OK but...
initial CPU overhead 10%
+ day after day increase of LDAP DB2 CPU (open PMR)
 - Ultimately DSNCLINF was bound with reopt(once)
=> performance OK
=> **LDAP CPU usage down to the CPU levels of DB2 V8**

..

Access Path Changes Dynamic SQL



Access Path Changes Dynamic SQL

- Support IBM
 - Difficult to address Access Path problems/changes
- Other access path changes ?
 - YES (data warehousing)
 - Use OSC statistics advisor !

..



Remove BUILD2 Phase in Reorg

- DB2 V9 removes the need for a BUILD2 phase, by reorganizing the **whole** of the NPI.
 - The reorganized NPIs are then switched in the SWITCH phase along with the other data sets. The **only outage** is during the relatively short SWITCH phase
 - Relieve of massive GBP writes, short on storage
- Our benefit None... got worse...
 - We run reorg part x jobs in parallel.....
 - Less Concurrency between individual concurrent REORG Part job's (DSNU180I)
 - 'Critical' reorg's jobs do not run
 - ..

Remove BUILD2 Phase in Reorg

- Users should be aware of the following implications (V9 Technical Overview SG247330)
 - DB2 will need additional temporary disk storage for the shadow data sets for each NPI.
 - The cost of the REORG will increase, as the build of the entire NPI will require more CPU than the updates that occurred in the BUILD2 phase.
 - **As the entire NPI will be built and then switched, it is no longer possible to run reorganizations of different parts in separate jobs in parallel. Jobs that ran in parallel **will need to be changed.****
 - We REORG on a Partition Level
 - Need to change to REORG flow

Other problems

- OSC problem:
 - LOB-tablespaces – FULL
 - First tried to drop those tablespaces
 - But EXPLAIN failed
 - Problem solved by shedding clean-up procedures
 - UK54574


..




System Monitoring


- CPU
- I/O
- Memory

..






I/O (Sync – ASync)




I/O (Prefetching)





ICT Services

Memory



- Note that IFCID 225 is moved from record type 102 to 100

Lessons learned

- Rebind BEFORE migrating to V9
- Rebind AFTER migrating to V9
- Optimizer needs more detailed statistics
- Size of EDM SKELETON pool!
- Adjust the 32K workdatasets
- Be aware of changed REORG NPI behaviour

..