Experience the DB2 10 experience

Session Number 1409

Frank Petersen,
Bankdata/JN Data

IBM Software
Information On Demand 2011

Agenda :

- Few words about the Beta process
- Plan early – Do the prereq's
- Migration to DB2 10 – New stuff
- Migration experiences
- Performance experiences
- Conclusions

Activities during Beta DB2 10

- Bankdata participated in the beta running from late 2009 to GA
  - Some 30 areas of the product were evaluated
  - Including code – quality and readiness
  - Including documentation – readability and completeness
  - Including business value where applicable
  - Including performance measuring and saving potential where applicable
- Huge effort : +2000 hours used
- Extremely committed people
  - IBM organization
  - Documentation team
  - Fellow beta-testers
  - Beta group

Need of early planning

- Is DB2 10 a change in paradigm ?
  - Is migration more planning and prereq than a technical issue ?
  - Is the code quality better ?
  - If so – can we migrate earlier ?
    - Gain the savings earlier ?

- Areas that need your attention :
  - Time to get out of V8/V9 (1 year/double license)
  - Level set performance expectations (*)
  - Identify business cases for new development features (*)
  - Re-consider multi row inserts (*)
  - Real storage (1 Mb page size)
  - Elimination of Private Protocol (*)
  - Subsystem consolidation (*)
  - Monitors and tooling
  - SMS on the catalog/directory (*)
  - Catalog restructure (*)

Lines marked with * indicates areas that will be touched in this presentation

Catalog/Directory changes

- SMS on catalog and directory :
  - In DB2 10 this is mandatory.
  - Plan to set in your ACS routines/Zparm deck
    - Must use the Extended addressability.
  - During migration you do not define new pageset anymore
  - DB2 will allocate additional datasets automatically
  - You must erase old shadow datasets for cat/dir before going to 10
  - Review your automation and monitoring of these datasets

- Catalog restructure (NFM).
  - Pointers and links removed
  - Row level locking and inline Lobs used where applicable
  - Most of known concurrency problems removed
  - Review serialization in your flow of utilities, binds and DDL.

Level set performance expectations

- Investigate how much CPU in your peak hour is for DB2 work (SQL)
- Establish a way to reevaluate after upgrading
  - after CM
  - after rebind
  - after NFM
- Understand your ability to do a mass rebind
- Identify key workloads
  - Get an expectation of savings potential
  - Look through list of DB2 10 "sweet spots"
  - Judge real saving ASAP in test ?
  - In QA ?
- Document saving so you know what DB2 10 gave you
- Produce list of things that can give additional savings
  - Implement those in an intelligent priority
  - Document savings

Frank's unofficial Performance chart

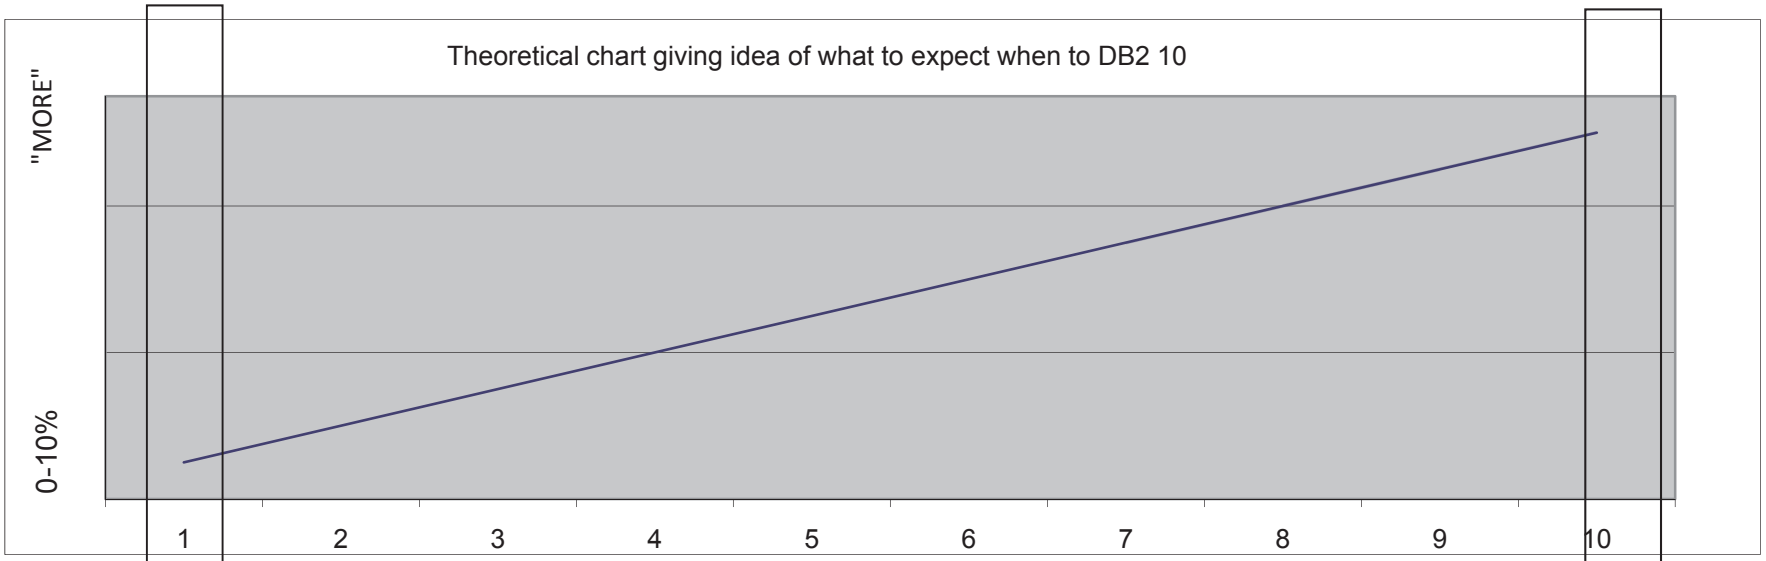

<<

- Concurrent inserts in sequential (or skip sequential) order (CM)
- Concurrent inserts into non-partitioned by taking an advantage of PBG Member Cluster
- Queries with multi range predicates and/or IN predicates (stage1) & evaluating many rows (CM)
- Queries with stage 2 predicates and evaluating many rows (CM)
- Native SQL procedures after regenerated in DB2 10 (CM)
- Workloads prev. tuned to save virtual storage and can take advantage of DB2 10 VSCR - using protected threads, RELEASE DEALLOCATE bind option (CM)
- Distributed workload by turning on high performance DBATs in DB2 10 (CM)
- Dyn.SQL with KEEPDYNAMIC YES with limited MAXKEEPD values by increasing this (CM)
- Concurrent insert suffering from log latch contentions - LC19 (CM/NFM)
- Online transactions requiring a large number of package allocation - LC24 (CM)
- Queries with CPU parallelism - LC24 (CM)
- Dyn.SQL using literals by enabling CONCENTRATE STATEMENTS WITH LITERALS (NFM)
- JCC type 2 (e.g., WAS on z/OS) or ODBC on z/OS applications executing queries returning multiple rows by using limited block fetch (CM)
- Workload accessing small sized LOB by taking an advantage of INLINE LOB (NFM)
- Workload updating small nodes on large XML documents by using XML MODIFY (NFM)
- Applications accessing data in random orders by taking an advantage of HASH ACCESS (NFM)
- Applications accessing the table with multiple indexes by using INCLUDE index (NFM)
- By replacing application logic (triggers, stored procedures) with Bi-Temporal features (NFM)

>>

Eliminate private protocol

- Private Protocol has to be removed before DB2 10
  - Check/rebind packages bound with DBPROT(P)
  - Install APAR PK64045 on V8/V9 systems
    - new ZPARM DRDA_RESOLVE_ALIAS
  - Look at DSNTP2DP (DSNTIJPD) REXX tool to help convert
  - Look at PM04968
    - delivers DSNTIJPM – check if your are ready for DB2 10
  - Install APAR PK92339 on V8/V9 systems (today)
    - Introduces ZPARM to turn off PP.
  - Revoke “AT ALL LOCATIONS” (“public*”) before entering DB2 10 !
  - Hopefully you are not still using SNA between z/OS subsystems ?

Storage considerations

- DB2 10 has significant reduction of storage below the bar
  - Increases the number of threads possible
  - Ease up on monitoring of storage
  - Re-evaluate decisions from the "JC storage diet"
  - Do your homework
    - Monitor your storage before and after 10
    - Not least your real storage consumption
  - Consider consolidating subsystems in datasharing
    - Only if these were build because of # of threads etc
  - Do you have real storage enough ?
  - Consider page fixing bufferpools
  - Consider going for large page support in 10
  - Consider candidates for PGSTEAL(NONE)
  - Notice that DB2 10 allocates bufferpools in chucks on demand

Going to DB2 10 ? When and why ?

- Reducing running costs :
  - CPU savings out of the box
  - HASH access
  - Buffer pool enhancements
  - Reduction in # of latches
  - Reduction in CAT/DIR contention
  - SMS/DB2 managed CAT/DIR
  - UTS improvements
  - DB2 subsys consolidation
  - Inline LOBS
  - Dynamic statement Cache enhancements
  - Additional non-key columns to index
  - I/O parallelism on index update

- Reducing developing costs:
  - Timestamps with TZ and greater precision
  - Bi-temporal tables
  - Support for SQL table functions
  - Enhanced XML features
  - SQL enhancements
  - Addition of extended indicator variables
  - Access to currently committed data
  - Statement level monitoring

- Reducing administ
  - Online schema evolution
  - Online reorg enhancements
  - Skip level migration
  - Adding log DS on the fly
  - Statement level Opt. hints
  - Rotating partitions
  - Security and Auditing
  - Compress on INSERT
  - Automated statistics

Hash Access To Data

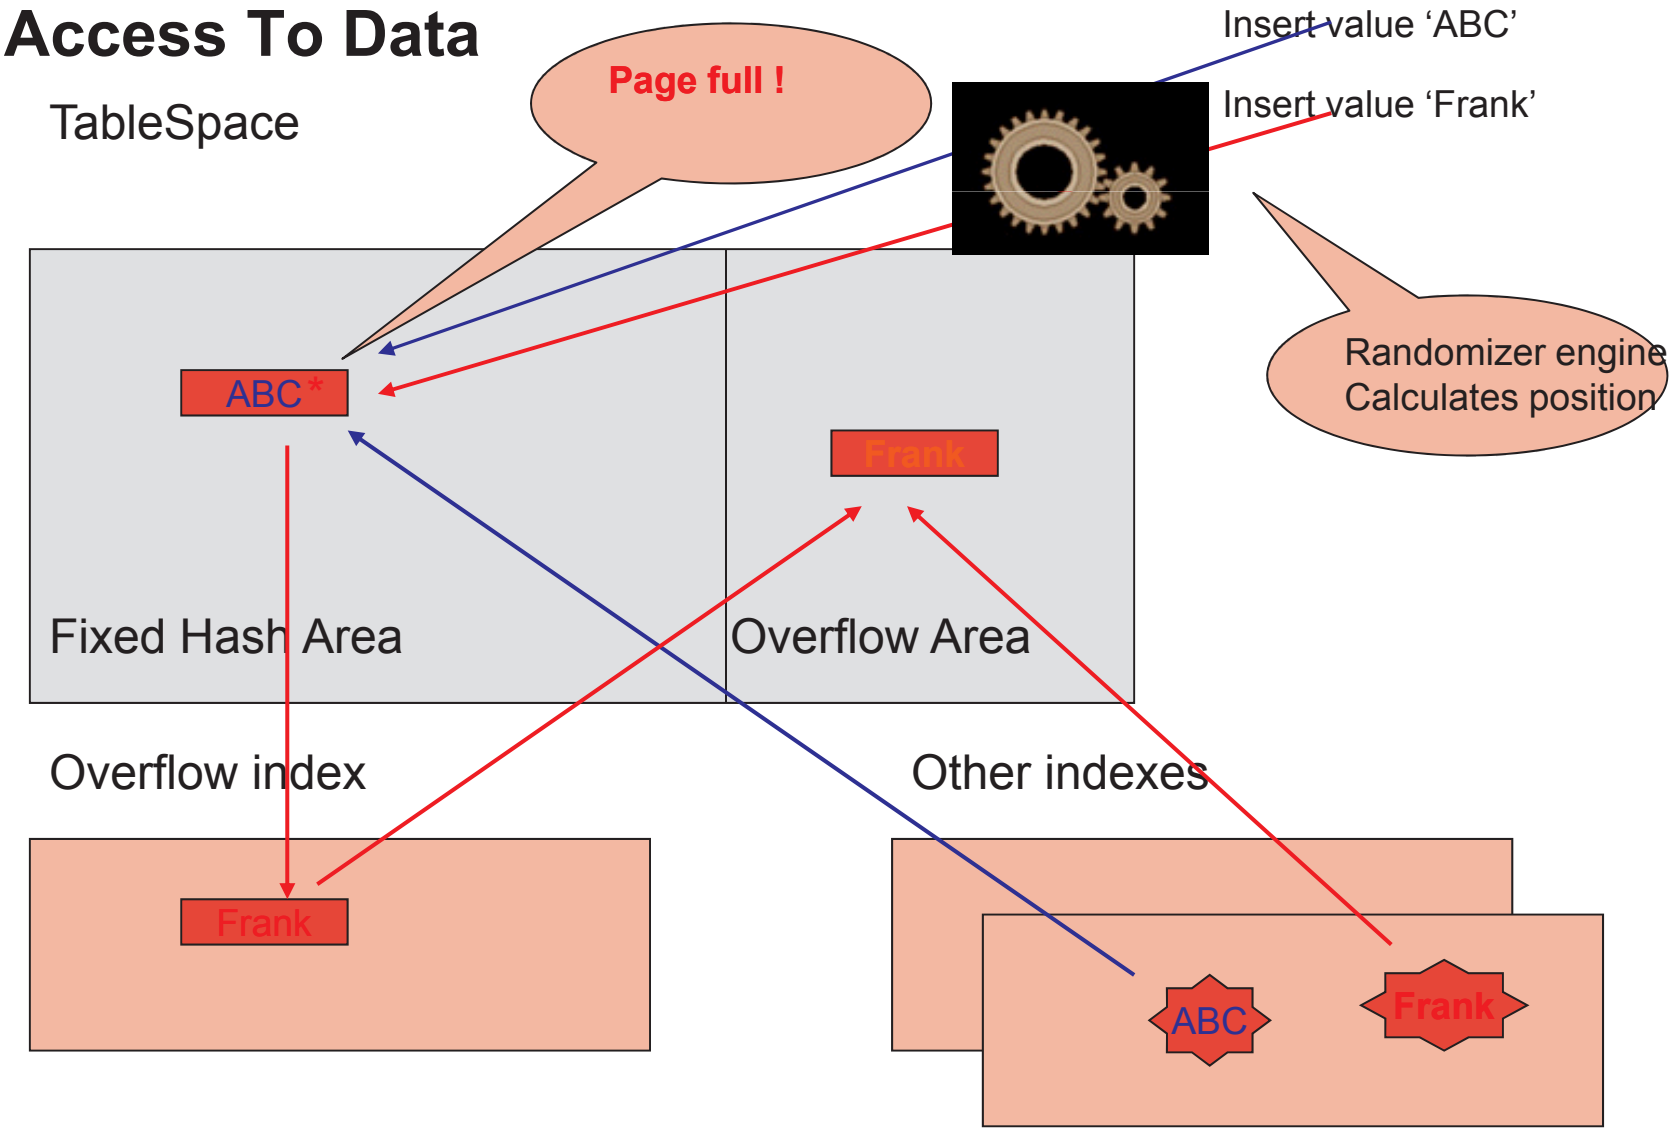

Hash Access To Data

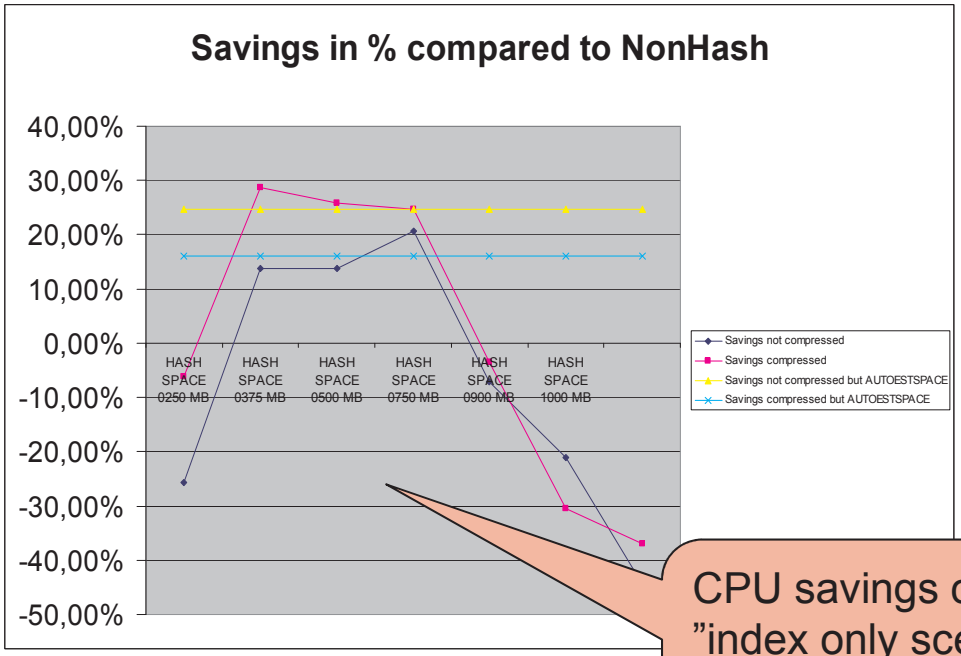

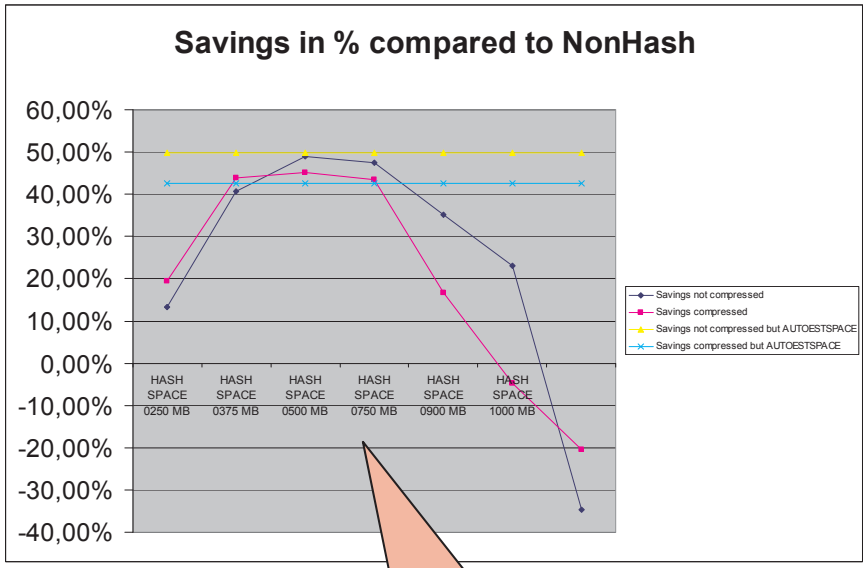

CPU savings compared to ”index only scenario”
Flat lines are ”autoestspace”
The rest is varying HashArea

CPU savings compared to ”index and data scenario”
Flat lines are ”autoestspace”
The rest is varying HashArea

XML – The new things

- DB2 10 XML support contains various new ‘goodies’ :
  - Engine Schema Validation
  - Multi Versioning
  - Update of parts of the document
  - Binary support
  - Date and Time support
  - Check XML

- Delivery that makes the DB2 for z/OS XML support more complete

XML Engine Schema Validation

- Example :

```
INSERT INTO Vehicles (VehicleNumber, Status, VehicleData)
values(18,'Running',
XMLPARSE(DOCUMENT SYSFUN.DSN_XMLVALIDATE('
<Root><VehicleNumber  ID=18> <FirstReg Date=1966-06-01/>
<IDData>
<Mark>Triumph</Mark>
<Location>
<line>Coventry</line>
<line>UK</line>
</Location></IDData>
<repairs>
<repair>
<partnumber>T14821</partnumber>
<name>Piston</name>
<quantity>2</quantity>
<price>112.25</price>
</repair>
</repairs>
</VehicleNumber></Root>','SYSXSR.Vehicles')));
```

CREATE TABLE Vehicles (
VehicleNumber BIGINT NOT NULL,

STATUS VARCHAR(10) NOT NULL WITH DEFAULT 'UnRestored',

VehicleData XML (XMLSCHEMA ID SYSXSR.VEHICLES), …

If XML document does not adhere to the schema the insert will be rejected by DB2

```
DSNT408I SQLCODE = -20399, ERROR:  ERROR ENCOUNTERED DURING XML PARSING
      VALIDATION AT LOCATION 364 Unknown reason RC=0018,RSN=8604,  OR XML
      PARSING OR VALIDATION AGAINST XML SCHEMA XSRID 26.
```

XML Engine Schema Validation

- Description and outcome XML CHECK
  - Now we alter the schema
    - Change the physical schema and do a new registration
    - Point the table to another schema (or add a schema)

```
ALTER TABLE Vehicles
  ALTER INFO SET DATA TYPE
  XML(XMLSCHEMA ID SYSXSR.Vehicles1);

---------+---------+---------+---------+---------+
DSNT404I SQLCODE = 162, WARNING:  TABLE SPACE
JSKDBTST.XMYC0001 HAS BEEN  PLACED IN CHECK
PENDING
```

  - Existing XML are now doubtful..
  - TableSpace goes into CHECK PENDING
  - Do a CHECK UTILITY :
    - SCOPE PENDING/REFONLY/AUXONLY
      - REFONLY check base table
      - AUXONLY checks the LOB/XML objects
    - XMLERROR REPORT/INVALIDATE
    - AUXERROR REPORT/INVALIDATE
  - If AUXERROR INVALIDATE and invalid XML doc found :
    - Document MOVED to dynamic shadow XML space
    - XML column in invalid base row marked invalid
    - Space goes info AUX WARNING-You have to clean up !

XML UPDATE

- Description and outcome
  - Before DB2 10 you could not update part of a XML document.
  - Now, DB2 comes with real SQL update syntax for XML.
  - Delivers subset of XQUERY update facility
    - Insert
    - Delete
    - Replace
  - Implemented using the scalar function XMLMODIFY
  - Easy to implement in SQL syntax.

- ***Let's see how XML Update works:***

XML UPDATE

```
....
<Location>
<line>Coventry</line>
<line>UK</line>
</Location>
....
```

```
UPDATE Vehicles SET
VehicleData = xmlmodify(
'replace value of node
/Vehiclenumber/iddata/location/line
[2]
with "United Kingdom"")
WHERE VehicleNumber=18;
```

Replace node

```
....
<Location>
<line>Coventry</line>
<line>United Kingdom</line>
</Location>
....
```

```
....
<Location>
<line>Coventry</line>
<line>United
Kingdom</line>
</Location>
....
```

```
UPDATE Vehicles SET VehicleData =
xmlmodify(
'insert node $in as last into
/Vehiclenumber/location',
xmlparse(document
'<region>Britain</region>') as "in")
WHERE VehicleNumber=18;
```

```
....
<Location>
<line>Coventry</line>
<line>United
Kingdom</line>
<Region>Britain</Region>
</Location>
....
```

```
....
<Location>
<line>Coventry</line>
<line>United
Kingdom</line>
<Region>Britain</Region>
</Location>
....
```

```
UPDATE Vehicles SET VehicleData =
xmlmodify(
'delete node
/Vehiclenumber/location/region') WHERE
VehicleNumber=18;
```

```
.....
<Location>
<line>Coventry</line>
<line>United
Kingdom</line>
</Location>
....
```

Bi Temporal Tables

- Description System Time
  - Timing is everything !
  - DB2 can now automatically archive a row when data is changed

```
CREATE TABLE Vehicles (
VehicleNumber INTEGER,Color CHAR(10),
SYS_START TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW BEGIN,
SYS_END TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW END,
CREATE_ID TIMESTAMP(12) GENERATED ALWAYS AS TRANSACTION START ID,
PERIOD SYSTEM_TIME(SYS_START, SYS_END));

CREATE TABLE Vehicles_HIST (
VehicleNumber INTEGER,Color CHAR(10),
SYS_START TIMESTAMP(12) NOT NULL,
SYS_END TIMESTAMP(12) NOT NULL,
CREATE_ID TIMESTAMP(12));

ALTER TABLE Vehicles ADD VERSIONING USE HISTORY TABLE Vehicles_HIST;
```


Bi Temporal Tables

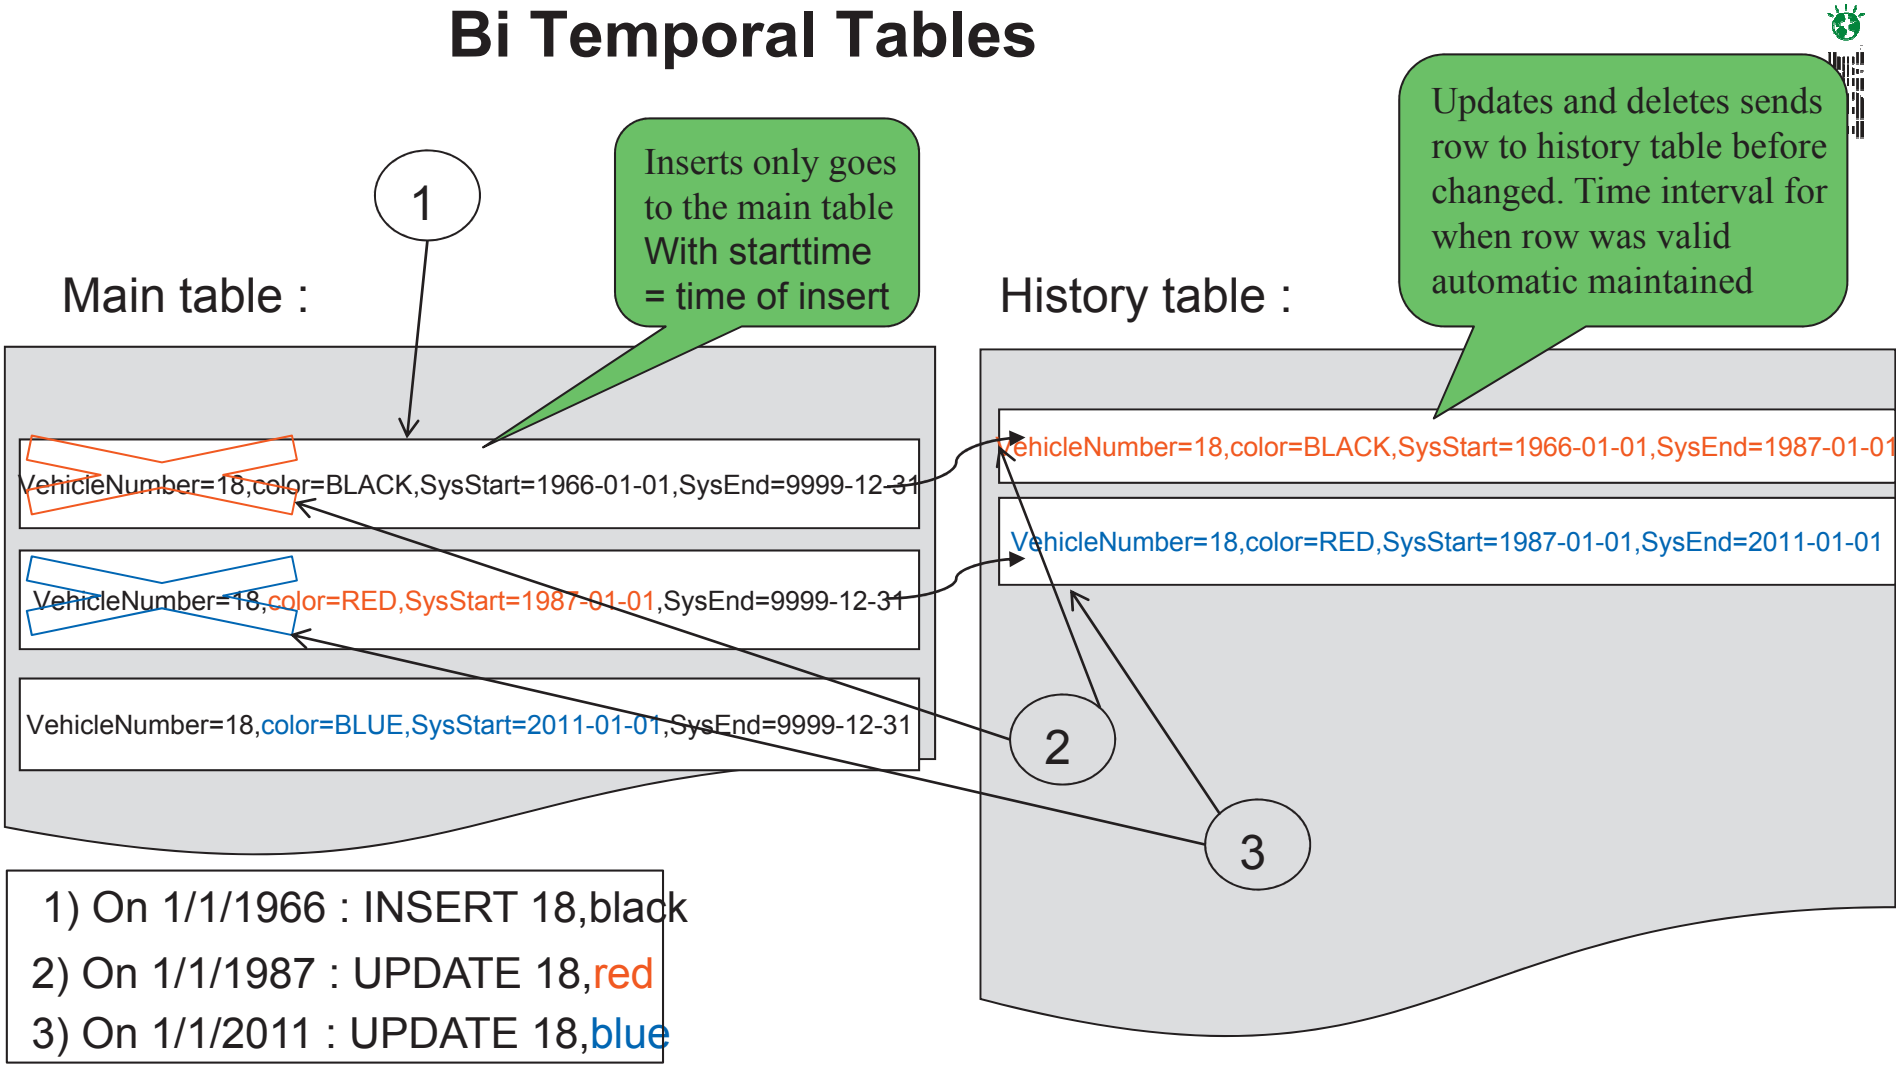

1) On 1/1/1966 : INSERT 18,black
2) On 1/1/1987 : UPDATE 18,red
3) On 1/1/2011 : UPDATE 18,blue

How to retrieve a historical row :

SELECT * FROM Vehicles FOR SYSTEM_TIME AS OF '2010-01-01' WHERE VehicleNumber=18-Result =”RED”
SELECT * FROM Vehicles FOR SYSTEM_TIME AS OF ’1965-01-01' WHERE VehicleNumber=18-Result =SQLCode +100
SELECT * FROM Vehicles FOR SYSTEM_TIME AS OF '2020-01-01' WHERE VehicleNumber=18-Result =”BLUE”

Bi Temporal Tables

- Description System Time Cont'd
  - A few things to observe :
    - BeginTime inclusive, EndTime exclusive
      - NOT customizable !!
    - Can be DROP'ped
    - Security at the history table is separate
    - SELECT from history table directly is indeed possible, though not advised !
    - INSERT into the history table also possible, even more strongly not advised
    - Will naturally require that you calculate on the space demand for history table.
    - Will only carve 1 history row per row per UOW. Can be a problem if a batch program is updating rows repeatedly per UOW.
    - Restriction to have the same columns in the tables
    - NO row retention time and auto purge… ☹

Bi Temporal Tables

- Description business Time
  - DB2 can now maintain multiple copies of ”same” row on a date/timestamp level

Main table :

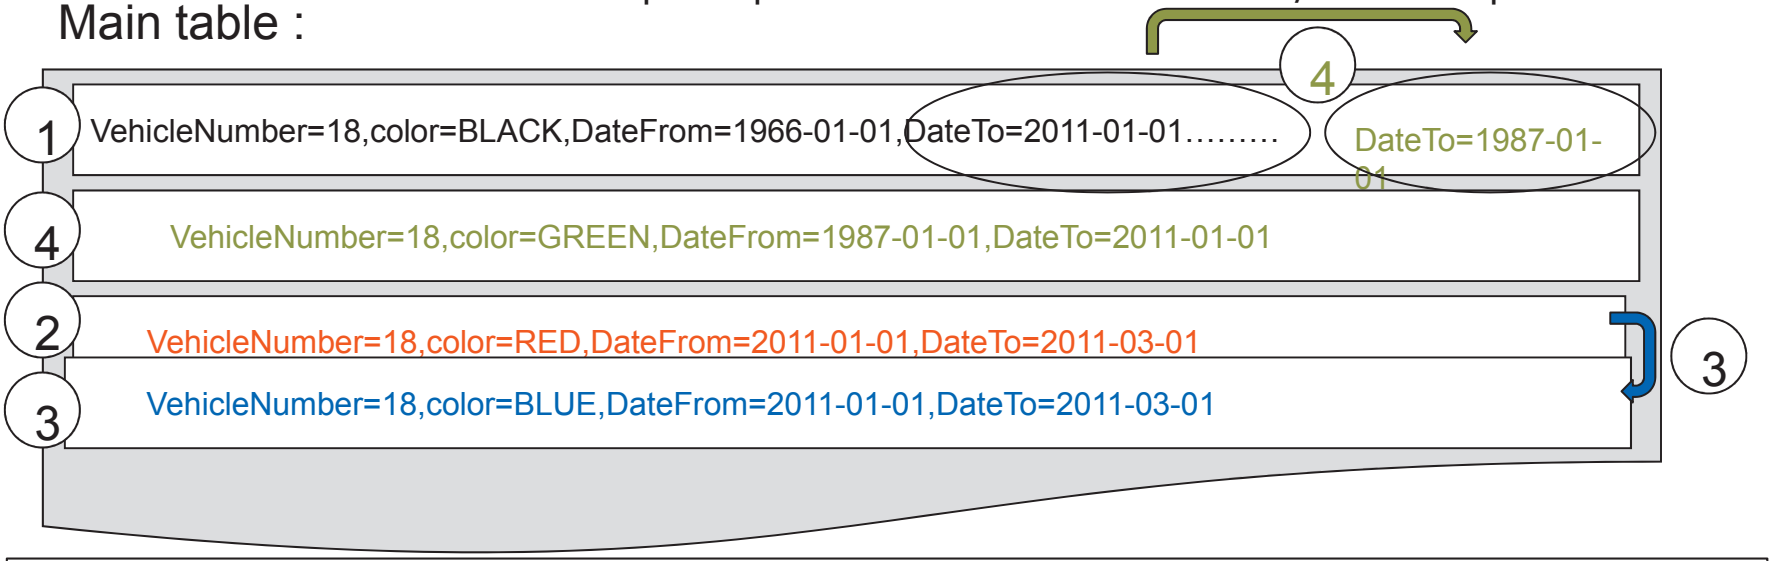

1) INSERT VehicleNumber=18,Color=black,DateFrom=1966-01-01,DateTo=2011-01-01
2) INSERT VehicleNumber=18,Color=red,DateFrom=2011-01-01,DateTo=2011-03-01
3) UPDATE FOR PORTION OF BUSINESS_TIME FROM '2011-01-01' TO '2011-03-01‘Color=BLUE WHERE VehicleNumber=18
4) UPDATE FOR PORTION OF BUSINESS_TIME FROM ‘1987-01-01' TO '2011-01-01‘Color=GREEN WHERE VehicleNumber=18

How to retrieve a time specific row :

SELECT * FROM Vehicles FOR BUSINESS_TIME AS OF ‘1988-01-03' WHERE VehicleNumber=18
Result =”GREEN”

- Bi-temporal tables and the future……
  - During the Beta Bankdata estimated that :
    - this feature could have coped with 80% of our temporal needs
    - The last is mostly because of a financial need for a third dimension
      - "Visibility" time – the need to set a value for a time in the future
        - But covering backwards in time.
        - Interest rate to be set 14 days from now but valid from yesterday !
    - Bankdata took discussions with a very committed IBM
      - There will be a "demo" at this conference by IBM
        - Using SP, triggers and UDF's
        - Can be implemented on top of DB2 10 NFM
        - Future is an open question (for me)

Timestamp changes

- Greater precision
  - In DB2 10 timestamp precision can go from micro to pico (6->12)
  - In fact YOU decide how many fractions you want (0->12)
  - Even at the biggest precision no guaranty for uniqueness (in theory)
  - From static or dynamic SQL.
- With timezone
  - Designed for data entered from different timezones to enable an easy calculation of duration and conversions.
  - Can it be used to keep systems running across change in DST ?
    - Evaluate how it works with related scalar functions (like HOUR())

Inline LOBS

- Description :
  - New Zparm LOB_INLINE_LENGTH sets the installation standard
    - Can also be defined in the DDL
  - If a given LOB is smaller than this LOB goes to base table.
  - If a given LOB is greater it partly goes to base and partly to aux.
  - Expect a massive performance benefit
    - On both SELECT and INSERT
    - Index on expression can be defined for inline portion
  - Testing was done by a purpose build copy of an existing Bankdata program.
    - INSERT XML-documents of varying length and SELECT those again.
    - The SELECT could be limited to selecting LOBs with a given length.
    - We have then performed executions with following combinations :
      - At DB2 V9
      - At DB2 10 with no INLINE length
      - At DB2 10 with INLINE length of 2100 bytes, 2800 bytes, 16000 bytes
  - The idea is that as we move the INLINE LENGTH up more and more of the documents will be kept in the base table.

Inline LOBS

- Description (cont'd)
  - A few measurements

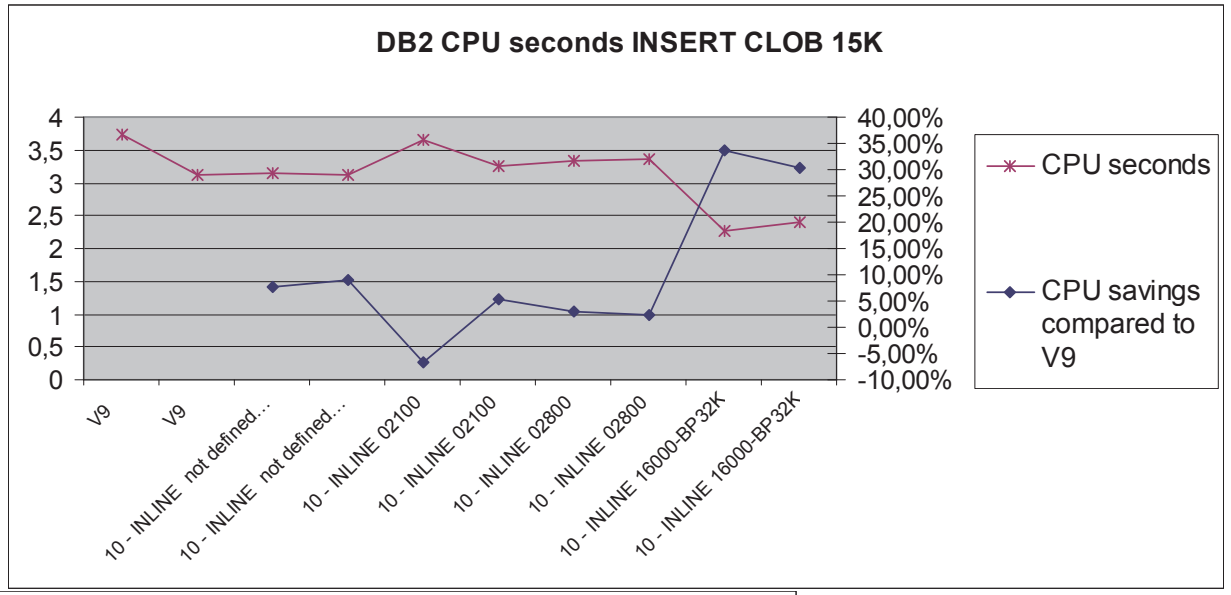

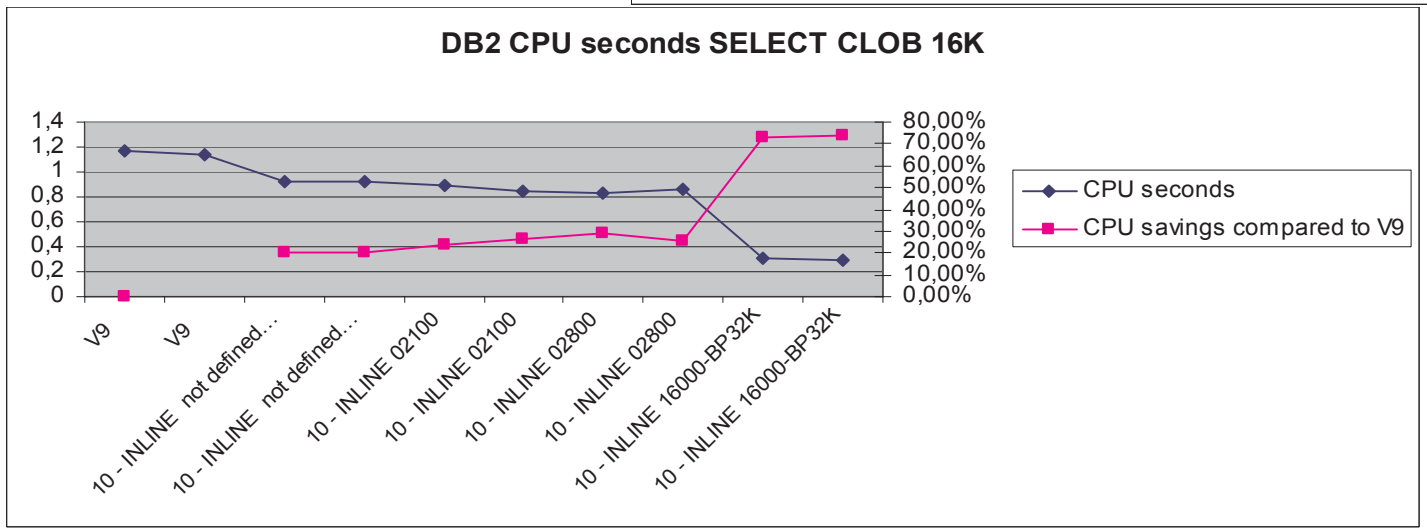

Reduce Log-Latch Contention

- Description :
  - A problem with unknown impact for heavy inserts in Data Sharing.
  - In version 8 spin loop in log record carving
    - In version 9 partly solved
  - A benchmark comparing V8->V9->V10 both datasharing/non-datasharing in all thinkable combinations of multirow and singlerow.
  - ***The conclusion is that DB2 10 DS can now run as fast as non-DS and finally take full advantage of MultiRow Insert.***
  - ***We are extremely pleased with this result !***

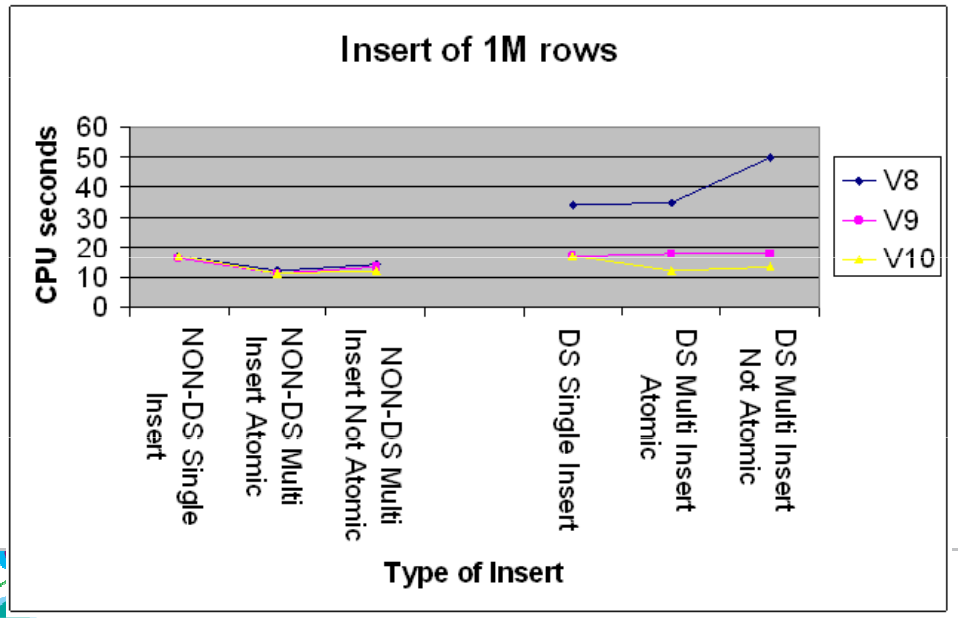

Statement level monitoring

- Statement level monitoring for static and dynamic SQL
- Produces IFCID records with timing info per statement.
- Can remove much of the need for doing tracing
- YOU can do IFI programming or wait for boring support in your tooling
- In assembler with IFI-programming issue a :
  - -STA TRA(MON) DEST(OP8) CLASS(29) IFCID(400)
- Empty out the monitor-data with READS-calls :

```
CALL  DSNWLI,(READS,IFCA,(R9),IFCIDS,QUALS),VL
CLC   IFCARC2,=XL4'00E60803'          NO TRACE DATA PRESENT
BE    NO_DATA_NOW
CLC   IFCARC2,=XL4'00E60827'          TRACE STOPPED
BE    TRACE_STOPPED
```

- The 401-record is mapped in macro DSNDQW05 and contains
  - Statement identifier as found in SYSPACKSTMT
  - All you desire : Getpages, # of rows examined, CPU time, WAIT time, ELAPS time etc etc
- Convert STCK value for timings to readable values in assembler:
  - STCKCONV STCKVAL=QW0401CP,CONVVAL=OUTAREA,TIMETYPE=DEC, DATETYPE=DDMMYYYY
- So from my assembler program I write (could be an insert to Performance DW) :
  - DATA CAPTURED : STMT_ID **161176** COLL: **DWS** PROG:**MY_PROG** CPU TIME FOR STMT uSEC **8400**

•Migration of a DS-group

- As example we will use a 2 way DS QA environment at Bankdata.
  - Migration jobs build in advance
    - Not using jobs generated by panels but for the first system
    - Manually modifying the jobs is easier as good naming conventions
    - Doc-member maintained through out all migrations
  - •Pre-req
    - Fall back/toleration SPE
    - Including the PP-fix
    - Early code activated
    - SMS definitions enabled
    - DBA's alerted concerning REBIND's
    - Ensuring performance datawarehouse is fit for fight !
    - Set MAXSPACE to at least 15G to ensure a full dump in case…..
    - Make sure your SYSDUMP-datasets are VSAM EA-enables (dumps are BIG)

•Migration of a DS-group

- Execute the migration steps that can be done in advance – in advance
  - "way before"
  - The pre-migration queries (DSNTIJPM)
    - You need time to deal with problems
  - Run the link checker – try with DISP=SHR
  - Run the CHECK INDEX on catalog
  - Run the SQL that checks for packages that will be invalidated.
  - Run the SQL that checks DB2 CATALOG consistency
  - Build and try your backup job
    - Cat/Dir image copies
    - DFDSS backup of logs and BSDS and selected tablespaces ?
    - Mirror copy of selected parts of the system ?
    - Execution datasets (SDSNEXIT, SDSNLOAD …. ZFS)
  - Build and assemble the new ZPARM-deck
    - Compare the source with your V9 version
  - Remove shadow copies of old Cat/Dir datasets

•Zparm changes-deceased parameters

- •MVS – z/OS architecture
- •EDMBFIT - Better fit. Allocate storage adjacent to existing slot in EDM-pool
- •LOGAPSTG - LOG fast apply storage- Now build into the product
- •HOPAUTH - DDF is removed
- •EDMPOOL - EDMPOOL split into numerous. Old EDMPOOL is only for packages bound before 10.
- •MAX_UTIL_PARTS - There is no longer a limit on the number of partitions for LOAD
- •PRIVATE_PROTOTAL - Disable PP. Should be done in the last weeks running DB2 V9
- •DRDA_RESOLVE_ALIAS - Set to YES if you use 3-part names in DDF....
- •PARTKEYU - controlled whether values in columns in partitioning keys can be updated.
- •OPTJBPR - Index costing when BP size was increased (PK69079).
- •OPTHYBCST - Enhanced cost model for Hybrid Join enabled.
- •OPTIXOPREF - index-only access for queries is favored over index plus data access when the index-only access has equal or better filtering.
- •OPTOIRCPF - more accurately estimate the cost of accessing inner table indexes when outer composite is ordered in inner index key ordering.
- •PREVALKEEP - Controlled SQL sequence references NEXTVAL or PREVAL,
- •REORG_IGNORE_FREESPACE - Controlled whether DB2 used the PCTFREE and FREEPAGE values that are defined for the table space in certain situations.
- •SJMISSKY - Star join optimization
- •XMLTABJPD - XMLTABLE join optimization. Now incorporated into the new pureXML implementation.

•Zparm changes-new born

•DB2SORT- Set if you have license for DB2 sort

•CHECK_SETCHKP - specifies whether the CHECK DATA and CHECK LOB utilities are to place inconsistent objects in CHECK PENDING status

–CHECK_FASTREPLICATION- specifies the type of replication that DSS COPY uses to copy objects to shadow data sets.

•CATXSTCL,CATXMGCL,CATXDACL,CATDSTCL,CATDMGCL,CATDDACL -SMS defs for CAT

•SMFCOMP- Compression on SMF-records.

•SIGNON_MODULE - Name for signon exit

–LOB_INLINE_LENGTH - Length of inline portion of LOBS

•IMPTKMOD - DB2 is to track modifications to pages of implicitly created table spaces.

•IMPDSSIZE - the maximum data set size (DSSIZE) in gigabytes that DB2 should use for creating an implicit base table space.

•DSN6SYSP IDAUTH_MODULE - Name of auth exit

•DPSEGSZ - Default segment size that is to be used for partitioned table spaces when the CREATE TABLESPACE statement does not include the SEGSIZE parameter.

•CHKTYPE- defines whether the interval between log checkpoints is to be based on the number of written log records, the time between checkpoints, or both

•CHKMINS- # of minutes between checkpoints

–CHKLOGR - # of log records between checkpoints

–ACCESS_CNTL_MODULE- Name of access control exit

•Zparm changes-new born cont'd

•SEPARATE_SECURITY - specifies whether DB2 security administrator duties are to be separated from system administrator duties for this subsystem.

•SECADM2 ,SECADM2_TYPE,SECADM1,SECADM1_TYPE

•REVOKE_DEP_PRIVILEGES - controls whether revoking a privilege from a user is to cause dependent privileges to be revoked. If dependent privileges are to be revoked, revoking a privilege from a user also revokes the privilege from anyone that the user has granted that privilege to.

•REC_FASTREPLICATION- specifies whether the RECOVER utility should use FlashCopy to recover from a FlashCopy image copy.

•PLANMGMTSCOPE- specifies the default plan management scope to use when the PLANMGMTSCOPE option is not explicitly specified for the bind or rebind of a package.

•PARA_EFF - controls the efficiency that DB2 assumes for parallelism when DB2 chooses an access path. The integer value that is used for this parameter represents a percentage efficiency.

•MAXTEMPS_RIDS - determines the maximum amount of temporary storage in the work file database that a single RID list can use at a time

•INDEX_IO_PARALLELISM - enable I/O parallelism for INSERT operations to improve insert workload performance on tables with 3 or more indexes

•FLASHCOPY_COPY,FLASHCOPY_LOAD,FLASHCOPY_PPRC,FLASHCOPY_REORG_TS,FLASHCOPY_REBUILD_INDEX,FLASHCOPY_REORG_INDEX - Controls the usage of flashcopy

•DSN6SPRM FCCOPYDDN HLQ.&&DB..&&SN.. - defines the default value that is to be used for the FCCOPYDDN parameter of the FLASHCOPY option of DB2 utilities control statements

–COMPRESS_SPT01 - specifies whether the SPT01 directory table space is to be compressed

–BIF_COMPATIBILITY – controls format of return value of built-in CHAR function

•CATMAINT

- Do you suffer of CATMAINTFOBIA ?
  - –Take a DFDSS dump of entire production system
    - CAT/DIR, sort, LOGs, BSDS's, exits etc
    - No user data
  - Restore to test or sandbox system
  - Define SMS catalog etc, etc
  - Adjust RACF (on sandbox system you can do it)
  - Very limited security exposure
    - No user data
    - Besides a little RUNSTATS and RTS values
    - System Auditors should not have any objections
  - – Start up the system on V9
  - – Run the migration steps
  - – And enjoy CATMAINT
    - Cancel the job, restart it and feel the CATMAINTFOBIA vanish……

•CATMAINT

- Additional comments regarding CATMAINT
  - Restarting need the RESTART keyword
  - For next Christmas I wish :
    - Better status in CATMAINT output during the run (“Index X of Y nn% completed”)
  - – Looking in JES-output from DBM1 will give you an indication :

DBD01.I0001.A001 RETAINED
DSNADH02.I0001.A001 RETAINED
DSNSSH02.I0001.A001 RETAINED
DSNDTX05.I0001.A001 RETAINED
DSNDYX02.I0001.A001 RETAINED
DSNDFX01.I0001.A001 RETAINED
DSNACX02.I0001.A001 RETAINED
DSNACX03.I0001.A001 RETAINED
DSNACX04.I0001.A001 RETAINED
DSNDBX01.I0001.A001 RETAINED
DSNDBX02.I0001.A001 RETAINED
DSNAPX02.I0001.A001 RETAINED
DSNGGX05.I0001.A001 RETAINED

DSNPSX01.I0001.A001 RETAINED
DSNPSX02.I0001.A001 RETAINED
DSNGGX06.I0001.A001 RETAINED
DSNDLX04.I0001.A001 RETAINED
DSNATX05.I0001.A001 RETAINED
DSNDPX05.I0001.A001 RETAINED
DSNDCX05.I0001.A001 RETAINED
DSNDXX07.I0001.A001 RETAINED
DSNDRX03.I0001.A001 RETAINED
DSNDKX02.I0001.A001 RETAINED
DSNDKX03.I0001.A001 RETAINED
SYSTSUNI.I0001.A001 RETAINED
SYSTSASC.I0001.A001 RETAINED

•Migration of a DS-group

- Time map for a real migration (11000 tables, 100000 package versions,SPT01 700 Cyls) :
  - 00.00 – 00.07 Run the pre-migration queries
  - 00.09 – 00.17 Run the link checker
  - 00.19 – 00.23 Run the SQL to find plans to be invalidated
  - 00.25 – 00.34 Run the SQL to check catalog inconsistency
  - 00.36 – 01.07 Image copies and backup
  - 01.09 – 01.15 DIS THREAD INDOUBT etc and STOP DB2 V9
  - 01.17 – 01.20 Activate DB2 10 code (copy SDSNLOAD etc etc)
  - 01.22 – 01.25 START DB2 10
  - 01.25 – 01.35 CATMAINT
    - Workload of CICS, CLI, UTILITIES and distributed on V9 member
    - CATMAINT crashes with lock UTILITIES and DRDA. These forced out !
  - 01.35 – 01.50 CATMAINT restarted and runs successful
  - 01.52 – 02.20 Check indexes – problem getting sort space
  - 02.22 – 02.24 Rebuild indexes
  - 02.26 – 02.28 Bind REXX and define sort space
  - 02.30 – 02.35 Bind SPUFI and DCLGEN. Define Query Opt. DB.
  - 02.37 – 02.39 Check for VIEWs to redefine
  - 02.41 – 02.45 Bind of CLI
  - 02.47 – 02.52 DB2 supplied Stored Procedures.

•Migration of a DS-group

- Issues found :
  - Rexx using DRDA to a DB2 V9 WITHOUT fallback/toleration SPE fails (DRDA/PP)
  - Index inconsistent (00C90101) – an old V9 problem ? Rebuild helped
  - 00c90101 erqual 53B6 – Under investigation. DB2 10 problem but rebind solves !!!!
  - DB2 CONNECT call in REXX will give RC > 0 if multiple CONNECTs – V9 gave RC = 0
  - DB2 V9 accepts timestamps as “YYYY-MM-DD-HH:mm:ss.tttttt” - DB2 10 rejects
  - Problem when REXX comes in from a DB2 10 system into a V9 member on a group where one member runs DB2 10
    - The DSNREXX package from DB2 10 must be bound on the receiving system (DRDA)
    - The first member in the receiving group upgraded to DB2 10 will need that package
    - If a 3 part request from a DB2 10 system ‘hits’ a member running DB2 V9
      - » An auto rebind is triggered
      - » If users on the DB2 10 member is running REXX the package is held and a LOCK TIMEOUT is due
      - »The user hangs for a minute or more and the bind fails. And again and again and again
    - SOLUTION : SET ZPARM ABIND to COEXIST
  - Problem with Stored Procedures if WLM setting is SYSPLEX-wide
    - Work from one member can be routed to another LPAR’s WLM address space
    - If that other member is running V9 a DSNX983I is received (code level mismatch)
    - You have to change WLM to use Address Spaces on the same LPAR as the caller
    - –Only one Address Space will be active (though remember NUMTCB=x)

•Issues found

- 8 h after migrating the main development system fallback was done-Not because of DB2 but
  - Parts of the DB2 Optim product supporting DB2 10 did not support DB2 10
  - Because of license limitations we had not been able to test in advance
- But fallback worked as expected
  - We had tested this very carefully during the Beta – payback time !!
- Delayed our DB2 10 migration with around 1 month
  - ((which is 1/12 of the available time in the legal agreement))
- Problem in COBOL and CLOB's :

V9 Precompiler :

```
10  CLOB-DATA.
 11  CLOB-DATA-LENGTH    PIC S9(9) COMP.
 11  CLOB-DATA-DATA.
     49 FILLER        PIC X(32767).
     49 FILLER        PIC X(32767).
     .... etc, etc
```

DB2 10 Precompiler :

```
10  CLOB-DATA.
 49 CLOB-DATA-LENGTH     PIC S9(9) COMP-5.
  49 CLOB-DATA-DATA      PIC X(52428800).
```

```
RECL er PIC S9(3)V USAGE COMP-3.
…
PERFORM VARYING TALLY FROM 1 BY RECL
  UNTIL TALLY > (ARKIV-DATA-LENGTH)
```

  - COBOL compile option TRUNC(OPT) sends application into endless loop !!!!!!!
- GROUP BY with no ORDER BY can give different order on the result set
  - Really an application problem but how do we track SQL with problems ???
- In SYSTABAUTH we had GRANTEE = PUBLIC* which was from old GRANT….PUBLIC AT ALL LOCATIONS
  - But revoke gives -104 - Solved when going to NFM
- SPT01 explodes in size :
  - from 330K tracks to 600 K tracks – thanks for catalogue not being user managed anymore ☺

•Migration of a DS-group

- Time map for going to NFM (11000 tables, 100000 package versions,SPT01 700 Cyls) :
  - 00.00 – 00.23 Run the DSNTIJEN
  - 00.30 – 00.30 Run the DSNTIJNF
- How long will you stay in ENFM before DSNTIJNF
  - Minutes ? Hours ? one day ?
  - But not days
- After NFM :
  - Rebinds should not be needed to gain performance
  - Look at parallelism in your BIND-flow
  - Look at removing serialization in your utility flow
  - If you experience locking problems in the catalog
    - Examine – it could justify a PMR
  - Examine and document savings from V9 – DB2 CM – DB2 CM after rebind – DB2 NFM
  - Look at taking advantage of DB2 10 features
    - INLINE LOBS
    - CONCENTRATE STATEMENTS WITH LITERALS
    - INCLUDE columns in indexes with drop of redundant indexes
  - Examine and document savings from V9 – DB2 CM – DB2 CM after rebind – DB2 NFM – DB2 NFM new features
  - Notice that some rows in RTS will change as some cat/dir now are standard partitioned spaces
  - Clean up old rubbish from V9 and before : SYSIBMTA and SYSIBMTB databases ??

Is it an urban myth that DB2 10 consumes more storage ?

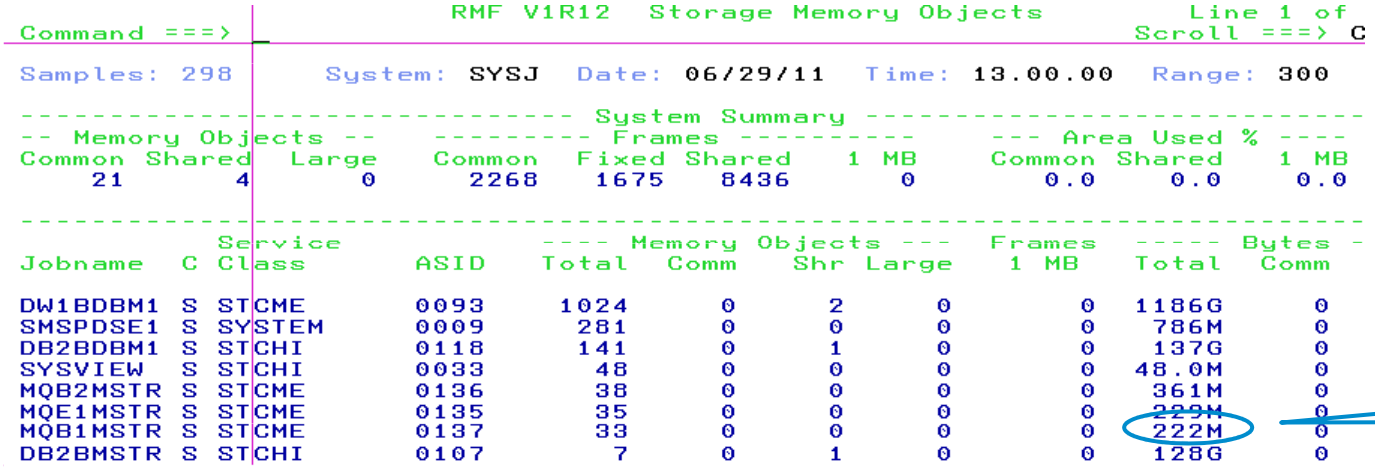

```
                RMF V1R12  Storage Memory Objects          Line 1 of
Command ===>                                              Scroll ===> C

Samples: 298      System: SYSJ   Date: 06/29/11   Time: 13.00.00   Range: 300

------------------------------ System Summary -----------------------------
-- Memory Objects --    --------- Frames ----------     --- Area Used % ----
Common Shared   Large    Common   Fixed Shared    1 MB     Common Shared    1 MB
    21      4       0      2268    1675   8436       0        0.0    0.0     0.0

-------------------------------------------------------------------------------
            Service                ---- Memory Objects ---   Frames   ----- Bytes -
Jobname  C  Class       ASID     Total   Comm    Shr Large    1 MB    Total   Comm

DW1BDBM1 S  STCME       0093      1024      0      2     0       0   1186G      0
SMSPDSE1 S  SYSTEM      0009       281      0      0     0       0    786M      0
DB2BDBM1 S  STCHI       0118       141      0      1     0       0    137G      0
SYSVIEW  S  STCHI       0033        48      0      0     0       0   48.0M      0
MQB2MSTR S  STCME       0136        38      0      0     0       0    361M      0
MQE1MSTR S  STCME       0135        35      0      0     0       0    229M      0
MQB1MSTR S  STCME       0137        33      0      0     0       0    222M      0
DB2BMSTR S  STCHI       0107         7      0      1     0       0    128G      0

Samples: 298      System: SYSJ   Date: 06/30/11   Time: 13.00.00   Range: 300

------------------------------ System Summary -----------------------------
-- Memory Objects --    --------- Frames ----------     --- Area Used % ----
Common Shared   Large    Common   Fixed Shared    1 MB     Common Shared    1 MB
    22      5       0      2977    1675  25679       0        0.0    0.0     0.0

-------------------------------------------------------------------------------
            Service                ---- Memory Objects ---   Frames   ----- Bytes
Jobname  C  Class       ASID     Total   Comm    Shr Large    1 MB    Total   Comm

DW1BDBM1 S  STCME       0093      1082      0      2     0       0   1186G      0
DB2BDBM1 S  STCHI       0266       939      0      2     0       0   1186G      0
SMSPDSE1 S  SYSTEM      0009       281      0      0     0       0    786M      0
SYSVIEW  S  STCHI       0033        48      0      0     0       0   48.0M      0
MQB2MSTR S  STCME       0136        40      0      0     0       0    363M      0
MQE1MSTR S  STCME       0135        35      0      0     0       0    229M      0
MQB1MSTR S  STCME       0137        35      0      0     0       0    226M      0
DW1BMSTR S  STCME       0085        10      1      2     0       0  1190G 6144M
DB2BMSTR S  STCHI       0180         8      1      2     0       0  1190G 6144M
```

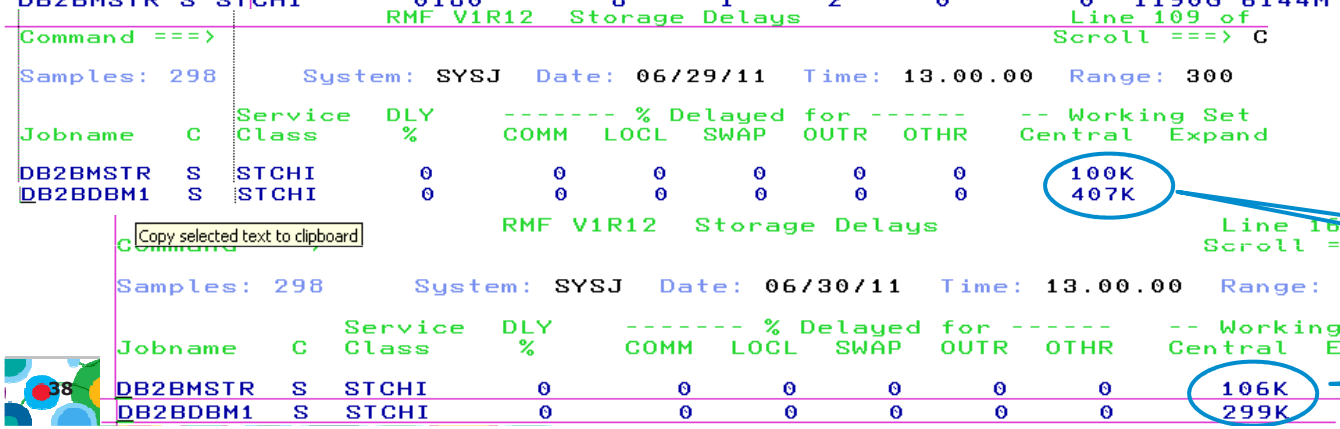

```
                RMF V1R12  Storage Delays            Line 109 of
Command ===>                                        Scroll ===> C

Samples: 298      System: SYSJ   Date: 06/29/11   Time: 13.00.00   Range: 300

            Service  DLY     ------- % Delayed for ------   -- Working Set
Jobname  C  Class     %      COMM   LOCL   SWAP   OUTR   OTHR   Central  Expand

DB2BMSTR S  STCHI       0       0      0      0      0      0      100K
DB2BDBM1 S  STCHI       0       0      0      0      0      0      407K
```

Copy selected text to clipboard

```
                RMF V1R12  Storage Delays            Line 165 of
Command ===>                                        Scroll ===> C

Samples: 298      System: SYSJ   Date: 06/30/11   Time: 13.00.00   Range: 300

            Service  DLY     ------- % Delayed for ------   -- Working Set
Jobname  C  Class     %      COMM   LOCL   SWAP   OUTR   OTHR   Central  Expand

DB2BMSTR S  STCHI       0       0      0      0      0      0      106K
DB2BDBM1 S  STCHI       0       0      0      0      0      0      299K
```

- DB2 DBM1 LMO (above the bar) :
  - V9 =137 Gb
  - DB2 10 = 1186 Gb

DB2 V9

DB2 10

- DB2 Working set :
  - V9 =0.5 Gb
  - DB2 10 = 0.4 Gb

- RECENT MEASUREMENTS ON PRODUCTION LIKE SYSTEM:
- +15% REAL STORAGE

DB2 V9

DB2 10

Performance measurements

- Numbers here are for a development system
  - Many of applications show no gain
  - Few gives additional overhead
    - Many gives overhead before rebind
  - Some show a CPU reduction :

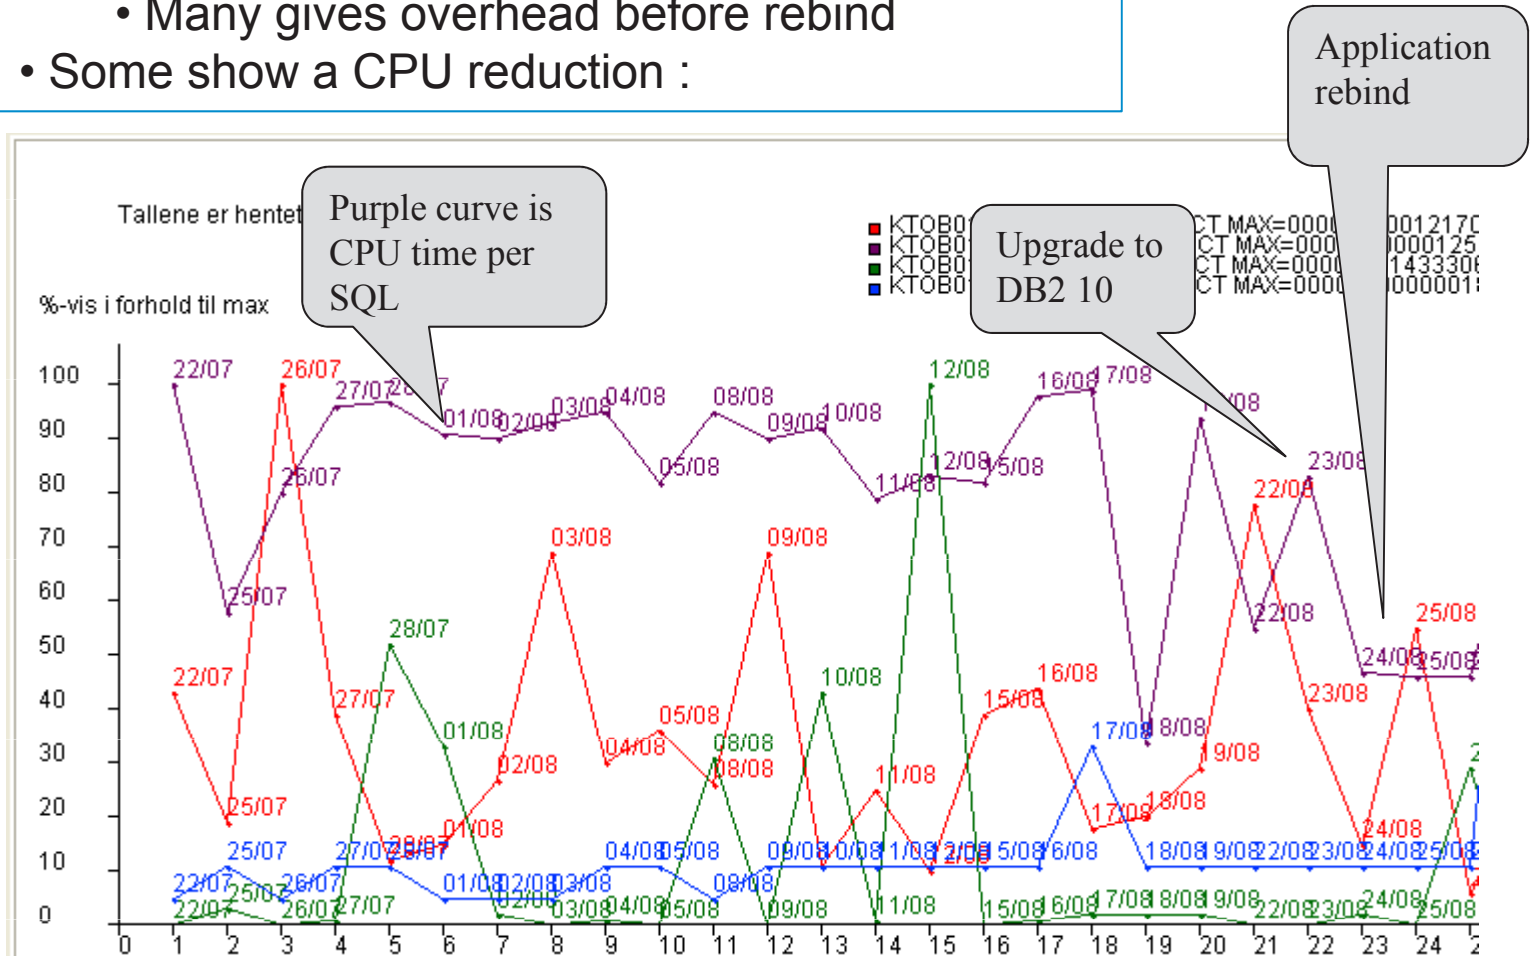

Performance measurements

- One very static SELECT application with same SELECT
  - On same data again and again
  - No overhead after upgrade before rebind

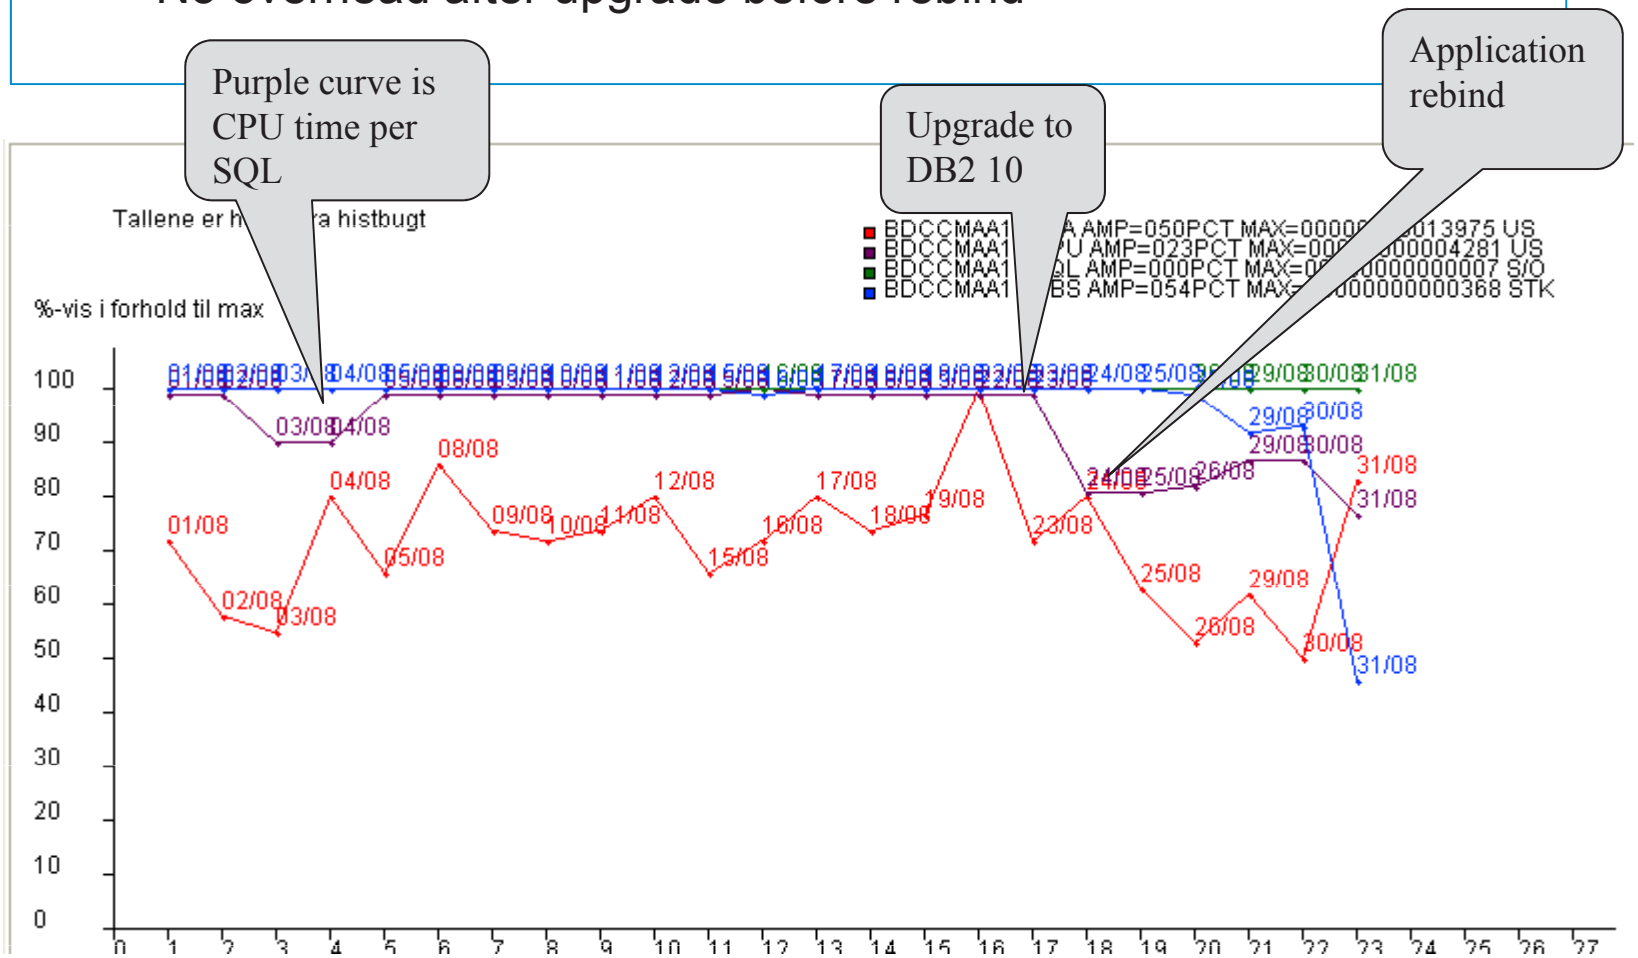

Performance measurements

- Here I show the GrandTotal (all applications).
  - Difficult to see any clear improvement
  - improvements of 5-10% can easily drown in the variations.
  - Huge variation as development/QA system
  - Hopefully I will have production data at time of conference

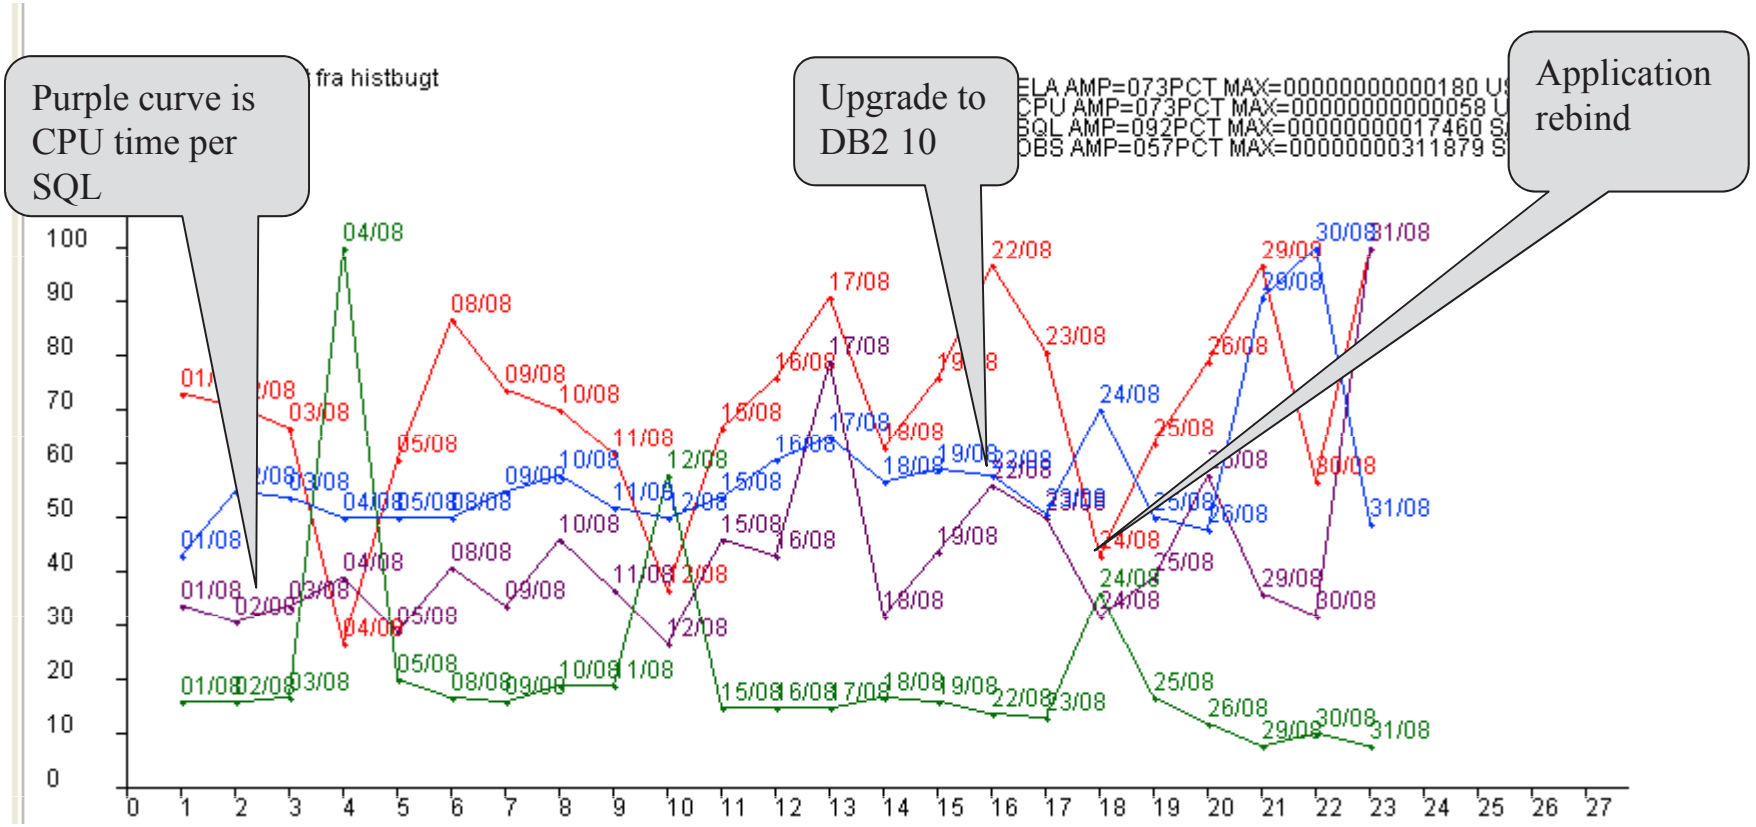

•Conclusion DB2 10

- DB2 10 is an important release
- Will give significant CPU saving in many applications
- Delivers extremely useful application features :
  - •Temporal tables
  - • XML
  - •Hashing
- Sets the cornerstone to new important features :
  - • Read last committed version of data
  - • Autonomic maintenance (hopefully)
- Finally solves some scalability and administrative issues :
  - •VSCR – Reduce # of sub-systems
  - • Online schema evolution – next step taken
  - • New DBA privilege
- But don’t under-estimate the smaller issues :
  - • Catalog restructure
  - • Statement level monitoring

Thank you !

Frank Petersen
JN Data
fap@jndata.dk

- Session 1409
- Experience the DB2 10 experience

Thank You!
Your Feedback is Important to Us

- Access your personal session survey list and complete via SmartSite
  - Your smart phone or web browser at: iodsmartsite.com
  - Any SmartSite kiosk onsite
  - Each completed session survey increases your chance to win an Apple iPod Touch with daily drawing sponsored by Alliance Tech