Development in a Virtualized Production Environment

–

True DB2 performance simulation covering

- Environment changes
- Schema changes
- Application changes

Ulf Heinrich
SEGUS, Inc
u.heinrich@segus.com

Agenda

- What the DB2 Optimizer needs to make a choice
- How to model a production environment
- Compare access paths resulting from
  - Buffer pool changes
  - RID expansion
  - CPU changes
- Test “what if” you drop an index, or change the sort order of an index

What the optimizer needs to make a choice

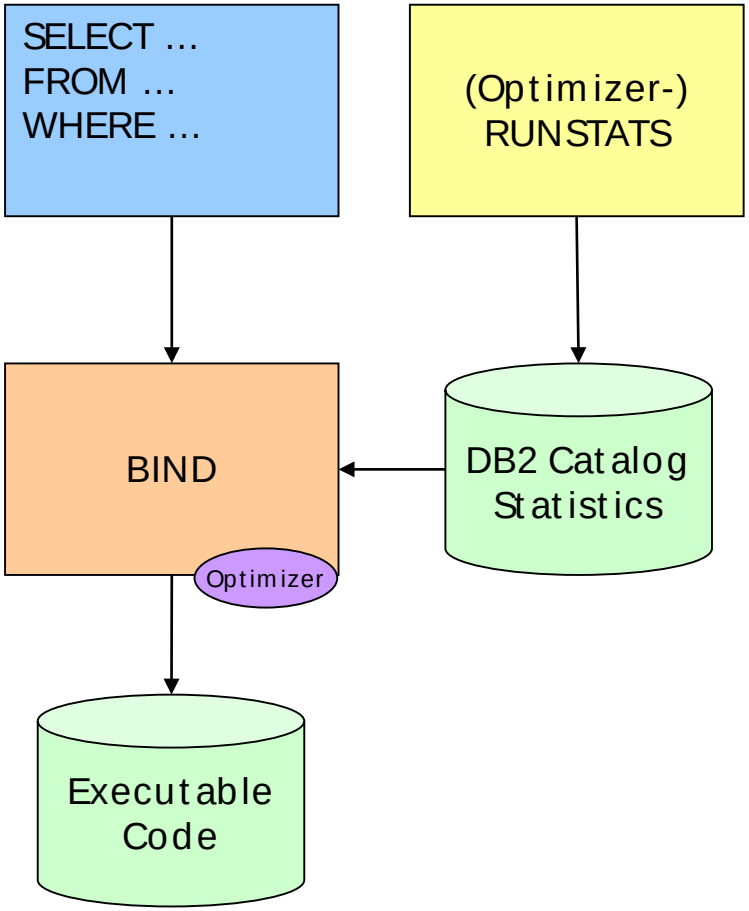

What the optimizer needs to make a choice

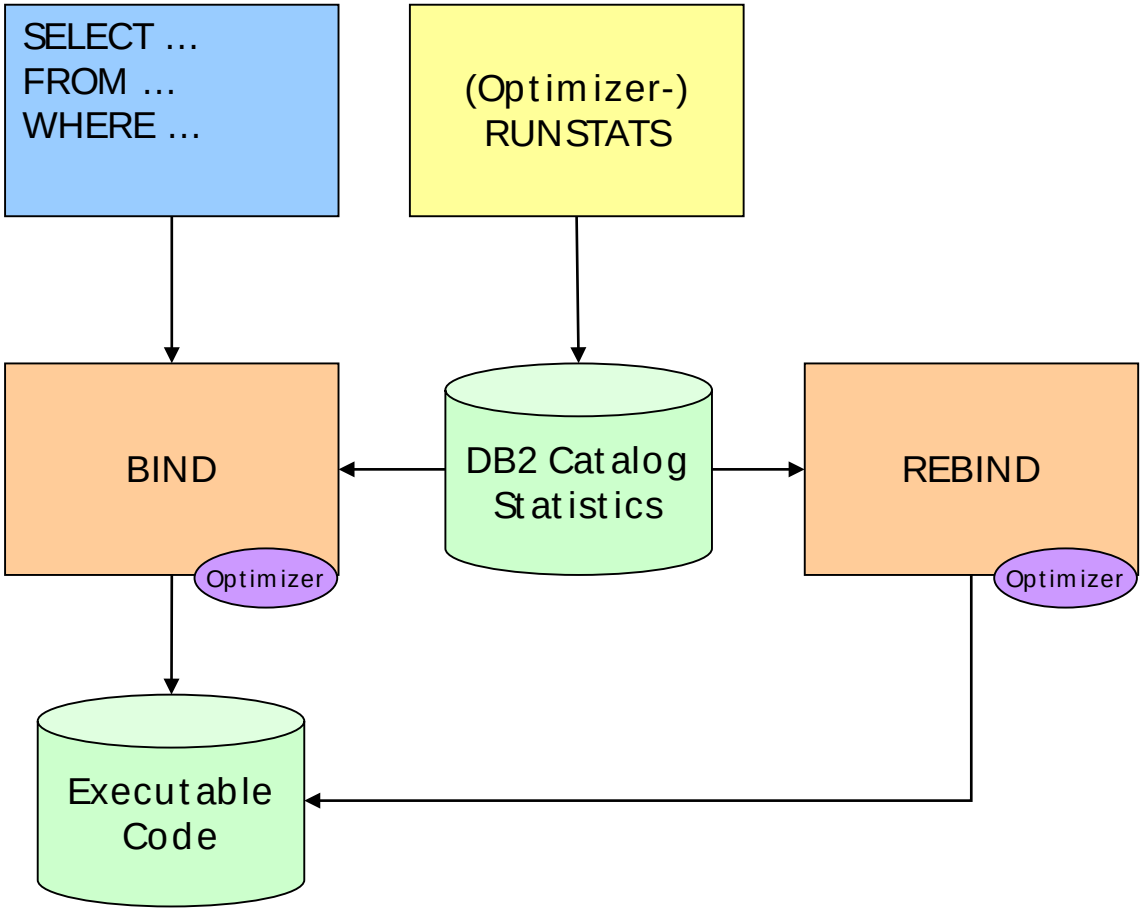

What the optimizer needs to make a choice

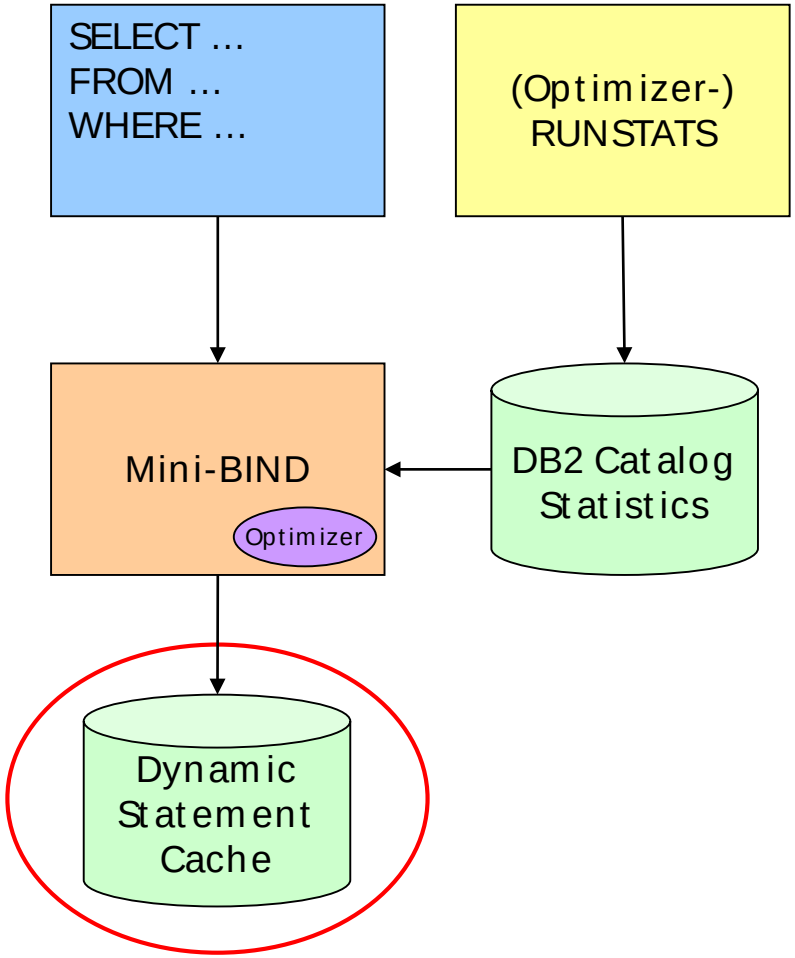

What the optimizer needs to make a choice

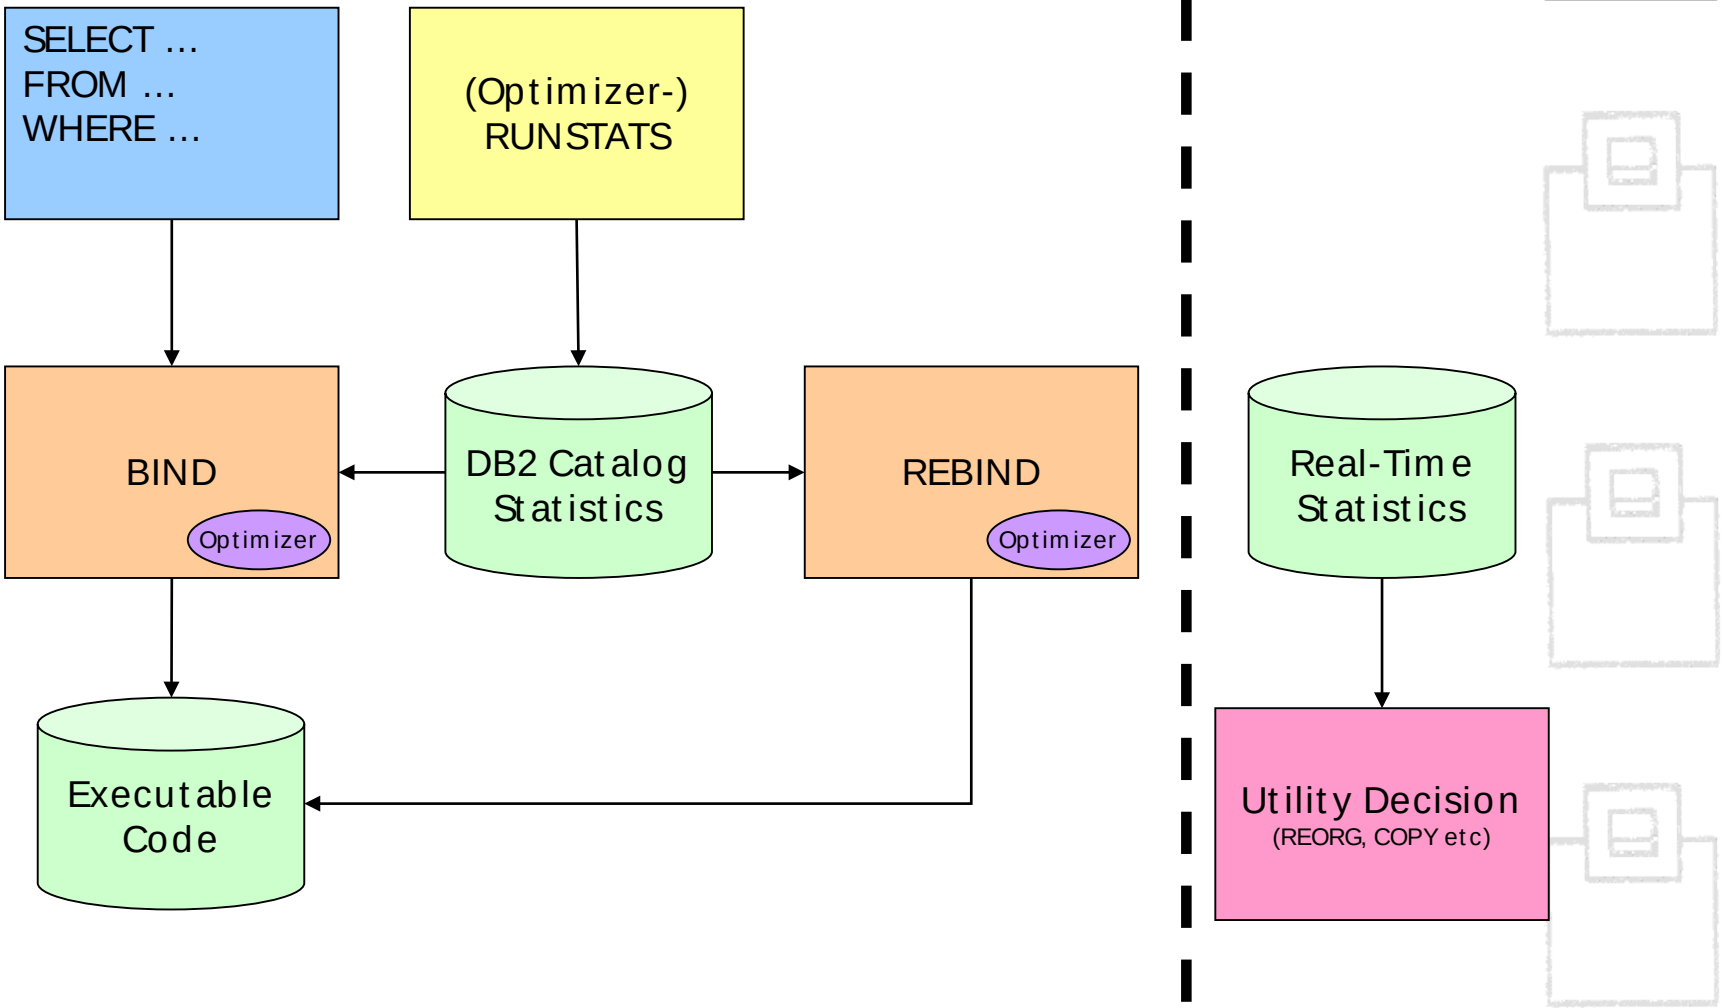

What the optimizer needs to make a choice

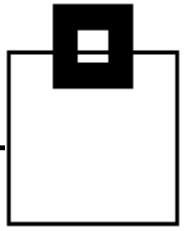

1st DDL:

```
SELECT …
FROM …
WHERE …
```


- Which object are referenced in the SQL
  - SELECT <columns> FROM <table> or <view> or …
  - WHERE <local predicates>
  - ORDER BY or GROUP BY or UNION or … <columns>

- Which objects are defined and how
  - INDEX
  - PARTITIONING

What the optimizer needs to make a choice

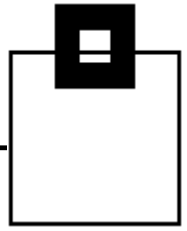

2nd STATISTICS:

SYSCOLDIST / SYSKEYTGTDIST
- CARDF
- COLGROUPCOLNO / KEYGROUPKEYNO
- COLVALUE / KEYVALUE
- FREQUENCYF
- HIGHVALUE
- LOWVALUE
- NUMCOLUMNS / NUMKEYS
- QUANTILENO
- STATSTIME

SYSCOLUMNS / SYSKEYTARGETS
- COLCARDF / CARDF
- HIGH2KEY
- LOW2KEY
- n/a / STATS_FORMAT

SYSCOLSTATS
- COLCARD
- HIGHKEY
- LOWKEY

SYSINDEXES
- CLUSTERING*
- CLUSTERRATIO
- CLUSTERRATIOF
- DATAREPEATFACTORF
- FIRSTKEYCARDF
- FULLKEYCARDF
- NLEAF
- NLEVELS

SYSINDEXPART
- LIMITKEY*

* Columns are not updated by RUNSTATS

_ Columns are not updatable

SYSROUTINES
- CARDINALITY*
- INITIAL_INSTS*
- INITIAL_IOS*
- INSTS_PER_INVOC*
- IOS_PER_INVOC*

SYSTABLES
- CARDF
- EDPROC*
- NPAGES
- NPAGESF
- PCTROWCOMP

SYSTABLESPACE
- NACTIVE
- NACTIVEF

SYSTABSTATS
- CARD
- CARDF
- NPAGES

What the optimizer needs to make a choice

3rd ENVIRONMENT:

- CP speed
- # of CPs
- BPs
- RID pool
- Sort pool

How to model a production environment

Create a Baseline…

- DDL
- STATISTICS
- ENVIRONMENTAL data

… allow to select any existing SQL along with the current access path details for comparison

- Production PLAN_TABLE
- Dynamic Statement Cache (DSC)

… and consider workload data

- DSC
- Monitor

How to model a production environment

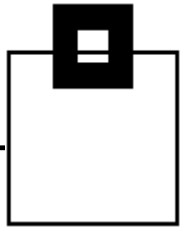

- A Baseline represents a given system, or subset

- Do multiple Baselines to allow covering
  - Different systems
  - Different amounts of data
  - Different distribution
  - Different workload, or time frames

- Apply naming conversions if you need to allow
  - Multiplications of Baselines (e.g. for different users)
  - Merging multiple production environments

How to model a production environment

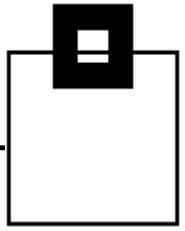

→ A test/simulation environment that behaves exactly like the production environment

→ This allows to pre-check the results from a
- New application version
- New DB2 version (or APARs affecting performance)
- New statistics
  - RUNSTATS (especially with dynamic SQL)

→ Keep control of your environment!
- Quickly identify and understand performance changes (degradation)
- Protect your production environment

How to model a production environment

Gathering data for the simulation environment

- DB2P: Production
  - Extract all/subset DDL
  - Extract the relevant catalog statistics
  - Extract CP speed, # of CPs, BPs, RID & Sort pool
  - Static SQL:
    - Extract Packages and PLAN_TABLE data
  - Dynamic SQL:
    - Take a snapshot of the Dynamic Statement Cache
    - Explain all captured statements to a temporary PLAN_TABLE

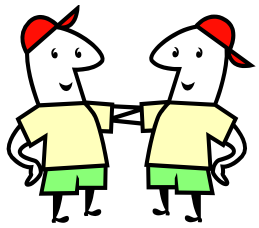

How to model a production environment

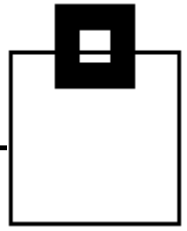

Applying data to the simulation environment

... Import the flat files:

- Apply naming changes if desired

- DB2Q: Quality Assurance
  - Create DDL DEFINE NO
  - Update the relevant catalog statistics
  - Apply environment modeling
  - Static SQL:
    - Import Packages and PLAN_TABLE
  - Dynamic SQL:
    - Import captured SQL and PLAN_TABLE

→ A true simulated environment, ready for any desired analysis/ precheck

How to model a production environment

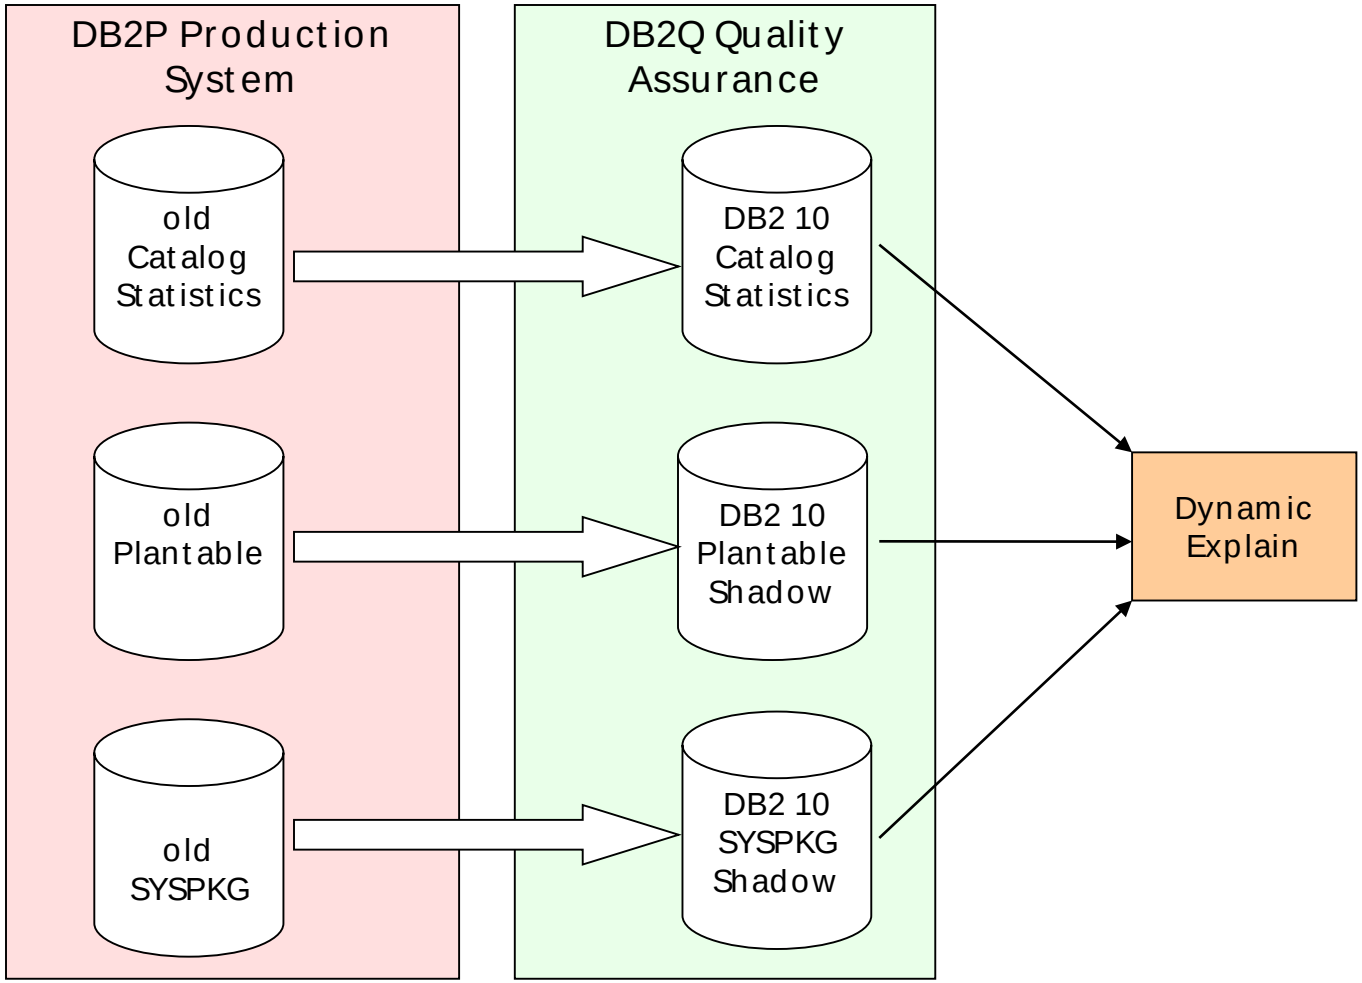

How to model a production environment

Dynamic SQL management and protection:

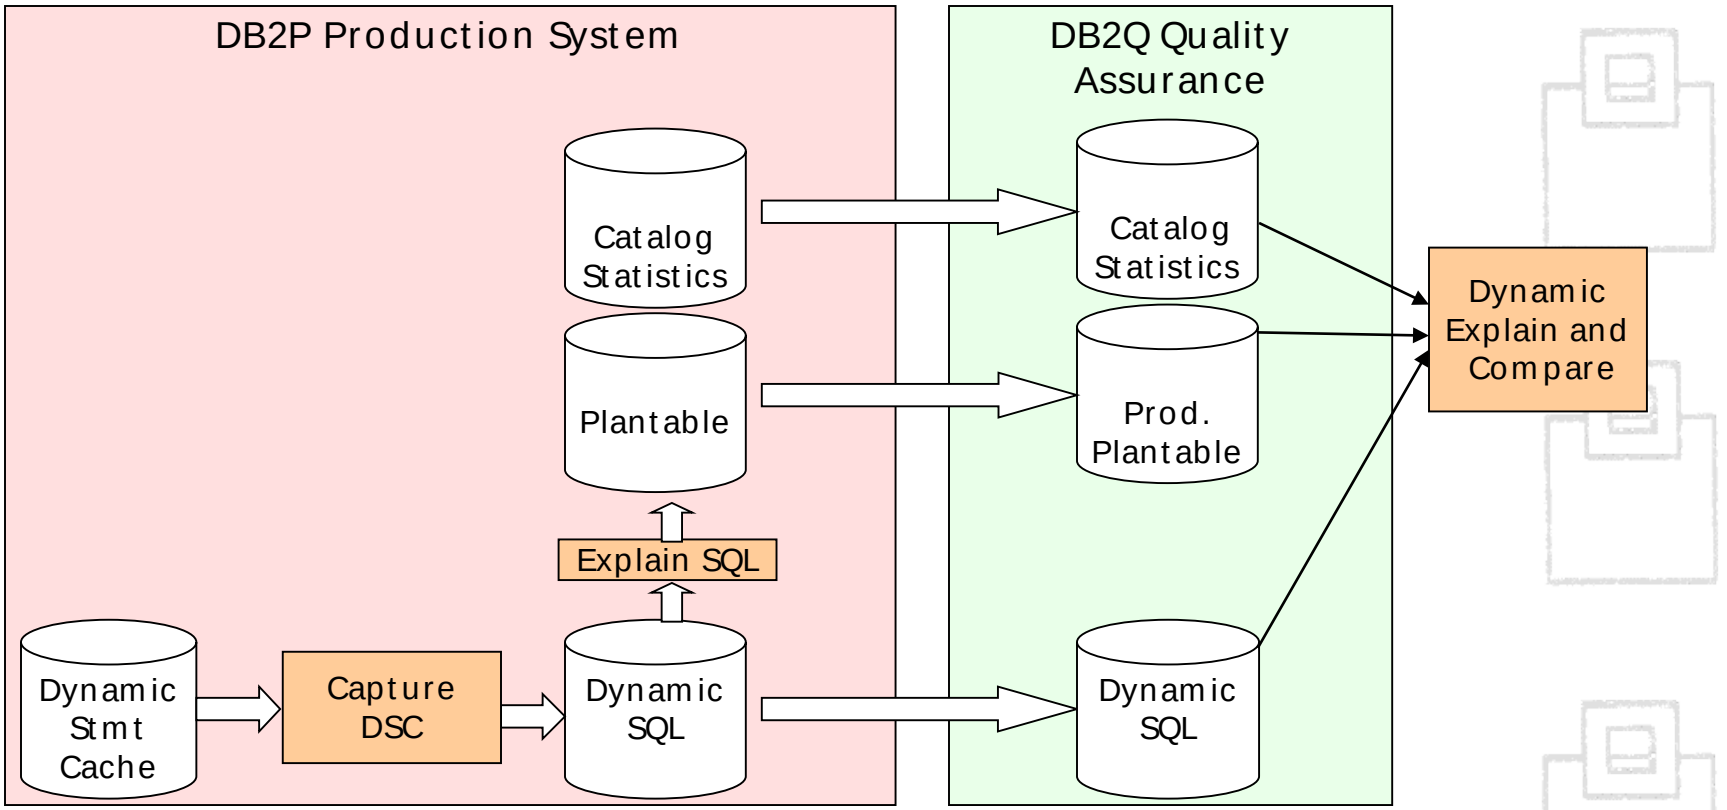

Compare Access Paths – Environment changes

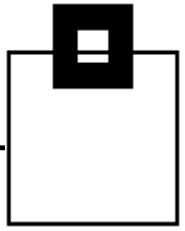

- CPU simulation
  - Copy production to test
  - Check a faster newer machine (Upsize)
  - Check a slower older machine (Downsize)

- ZPARM simulation
  - Change size of SRTPOOL
  - Change size of RID Pool
  - Change size of data cache or Star Join Pool

- BUFFERPOOL
  - Change size of any BUFFERPOOL

Compare Access Paths – Environment changes

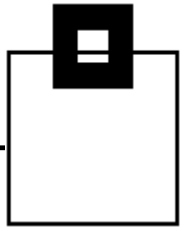

Production Modelling

- Supports optimizer overrides for optimizer relevant system settings
  - New zparms
    - SIMULATED_CPU_SPEED
    - SIMULATED_COUNT
  - New SYSIBM.DSN_PROFILE_ATTRIBUTES*
    - SORT_POOL_SIZE
    - MAX_RIDBLOCKS
    - For bufferpools

*Find DDL in member DSNTIJOS of your SDSNSAMP

Compare Access Paths – Environment changes

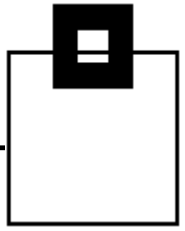

Production Modelling

→ ZPARM SIMULATED_CPU_SPEED

- µs of task or SRB execution time per SU
- OFF or an integer from 1 – 2147483647
- Gather from production by

```
EXPLAIN ALL SET QUERYNO=12345°
FOR SELECT * FROM SYSIBM.SYSDUMMY1;

SELECT HEX(SUBSTR(IBM_SERVICE_DATA,69,4))
           AS CPU_SPEED,
  FROM PLAN_TABLE
  WHERE QUERYNO=12345°;
```

- Apply the value to DSN6SPRM on test

° must be unique

Compare Access Paths – Environment changes

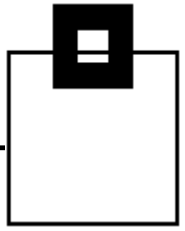

Production Modelling

→ ZPARM SIMULATED_CPU_COUNT*

- OFF or an integer from 1 – 127
- Gather from production by

```
SET CURRENT DEGREE='ANY';
EXPLAIN ALL SET QUERYNO=12345°
FOR SELECT * FROM SYSIBM.SYSDUMMY1;

SELECT HEX(SUBSTR(IBM_SERVICE_DATA,25,2))
           AS CPU_COUNT,
  FROM PLAN_TABLE
  WHERE QUERYNO=12345°;
```

- Apply the value to DSN6SPRM on test

*only if DEGREE='ANY'

° must be unique

Compare Access Paths – Environment changes

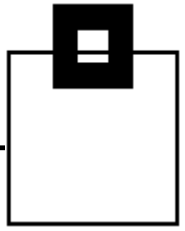

Production Modelling

- → DSN_PROFILE_TABLE
- → DSN_PROFILE_ATTRIBUTES SORT_POOL_SIZE
  - Gather from production by

```
EXPLAIN ALL SET QUERYNO=12345°
FOR SELECT * FROM SYSIBM.SYSDUMMY1;

SELECT HEX(SUBSTR(IBM_SERVICE_DATA,9,4))
           AS SORT_POOL_SIZE,
  FROM PLAN_TABLE
  WHERE QUERYNO=12345°;
```

° must be unique

Compare Access Paths – Environment changes

Production Modelling

- → DSN_PROFILE_TABLE
- → DSN_PROFILE_ATTRIBUTES SORT_POOL_SIZE
  - Apply on test by

```
INSERT INTO SYSIBM.DSN_PROFILE_TABLE
  VALUES (1234);

INSERT INTO SYSIBM.DSN_PROFILE_ATTRIBUTES
  VALUES (1234, 'SORT_POOL_SIZE',NULL,
  307200);
```

profile ID°

Desired size to simulate

° must be unique

Compare Access Paths – Environment changes

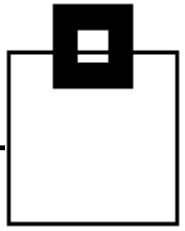

Production Modelling

- → DSN_PROFILE_TABLE
- → DSN_PROFILE_ATTRIBUTES MAX_RID_BLOCKS
  - Gather from production by

```
EXPLAIN ALL SET QUERYNO=12345°
FOR SELECT * FROM SYSIBM.SYSDUMMY1;

SELECT HEX(SUBSTR(IBM_SERVICE_DATA,13,4))
           AS MAX_RID_BLOCKS,
  FROM PLAN_TABLE
  WHERE QUERYNO=12345°;
```


° must be unique

Compare Access Paths – Environment changes

Production Modelling

- → DSN_PROFILE_TABLE
- → DSN_PROFILE_ATTRIBUTES MAX_RID_BLOCKS
  - Apply on test by

```
INSERT INTO SYSIBM.DSN_PROFILE_TABLE
  VALUES (1234);

INSERT INTO SYSIBM.DSN_PROFILE_ATTRIBUTES
  VALUES (1234, 'MAX_RIDBLOCKS',NULL,
  307200);
```

profile ID°

Desired size to simulate

° must be unique

Compare Access Paths – Environment changes

Production Modelling

- → DSN_PROFILE_TABLE
- → DSN_PROFILE_ATTRIBUTES bufferpools
  - Gather from production from DSNTIP1 panel
  - Apply on test by

```
INSERT INTO SYSIBM.DSN_PROFILE_TABLE
  VALUES (1234);

INSERT INTO SYSIBM.DSN_PROFILE_ATTRIBUTES
VALUES(1234, 'BP0',NULL, 25000);
```

profile ID°

Desired BP to simulate

Desired size to simulate

° must be unique

Compare Access Paths – Environment changes

All changes can be done „on the fly“ with no restart of DB2

Before rushing off and changing everything think!

- Will my change affect anyone apart from me?
- Am I sure about that?
- Am I really sure?
- Ask just in case!

(e.g. We don‘t want to set BP0 to 350.000 in test etc.)

Compare Access Paths – Environment changes

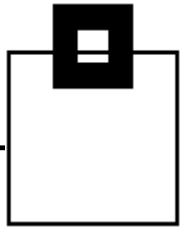

Production Modelling

- All simulation values are assigned to a **global** profile for a single DB2 subsystem
- Activation of a **global** profile by `-START PROFILE`
- Execute a dynamic EXPLAIN for the SQL queries you'd like to assess
- Refer to column REASON of the DSN_STATEMNT_TABLE
  → 'PROFILEID 1234,

Your assigned profile ID

- Refer to SYSIBM.DSN_PROFILE_TABLE
  → PROFILE_ENABLED = 'Y'

Compare Access Paths – Environment changes

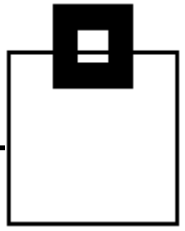

```
ImpactExpert for DB2 z/OS ----- Statements Summary -------
Command ===>
Mode: EARLY PRECHECK DYNAMIC
Primary cmd: END, R(efresh), Z(oom)
Line    cmd: S(elect)

AUTHID      : *

    Status                                      Statements
    ----------------------------------------    ----------
_   PROCESSED                                         8489
_   - IMPROVED                                          50
_   - CHANGED                                           44
_   - UNCHANGED                                       8027
_   - WORSENED                                         360
_   - PREFETCH CHANGE                                    8
_   NON-DETERMINABLE                                     0
_   PLAN_TABLE ISSUES                                    0
_   ERROR                                               11
_   - EXPLAIN ERROR                                     11
    ----------------------------------------    ----------
```

Compare Access Paths – Environment changes

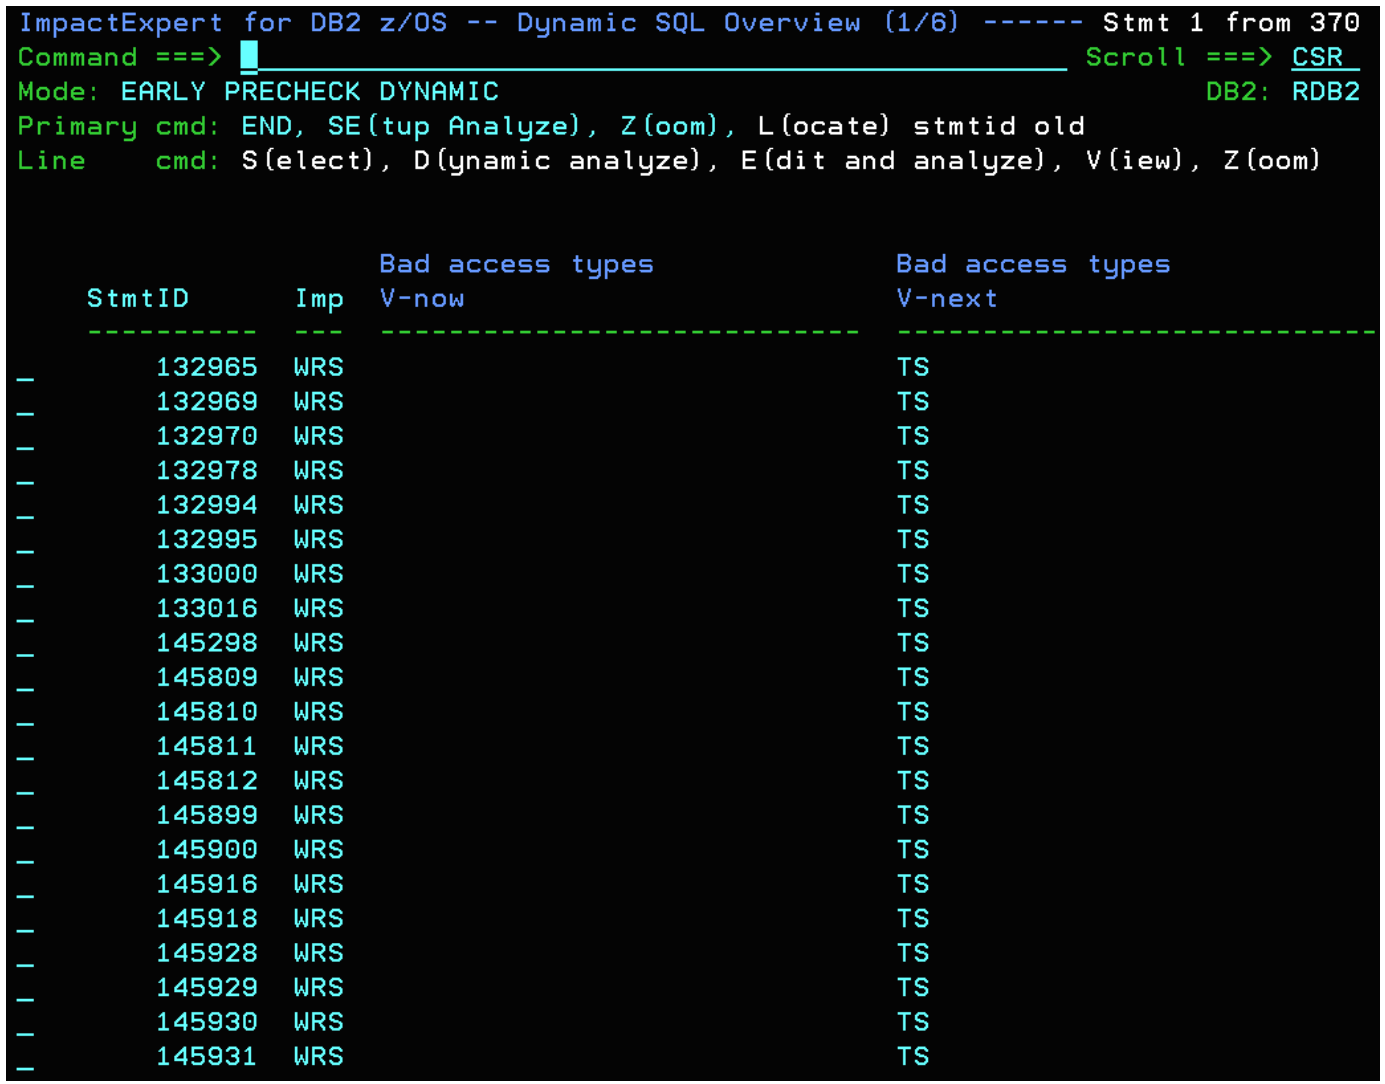

Compare Access Paths – Environment changes

```
ImpactExpert for DB2 z/OS ----- Access Paths ------- LINE 00000001 COL 001 080
Command ===> _____________________________________________ Scroll ===> CSR
Mode: EARLY PRECHECK DYNAMIC                                  DB2: RDB2
Primary cmd: END, C(atalog data), D(etails on/off), S(tatement text)

RunID. . . . VOX                          ExplainID V-now . . 1
Created TS . 2011-03-22-08.52.35.118244   ExplainID V-next. . 5
StmtID . . . 150185

Access path OLD ----------------------! Access path NEW ---------------------

TABLE              QB PN AC MA ME IX PR ! TABLE              QB PN AC MA ME IX PR
  INDEX                 TY CO TH ON FT !   INDEX                 TY CO TH ON FT
- - - - - - - - - - - - - - - - - - - - ! - - - - - - - - - - - - - - - - - - - -
T5DF4               1  1 I   1      N  ! T5DF4               1  1 R   0      N  S
  T5DF4~0                              !
                                       !
Milliseconds:              2           ! Milliseconds:              1
Serviceunits:             48           ! Serviceunits:             57
-------------------------------------------------------------------------------
******************************** Bottom of Data ********************************
```

- What happens with REALLY fast machines is that a TS scan is, in elapsed time, quicker than Index access!

Compare Access Paths – Schema changes

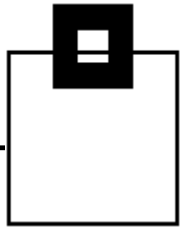

How to reliably simulate index changes?

E.g.

- Create index
- Drop Index
- Alter Index

→ DB2 9 for z/OS

→ DB2 8 (retro fitted via APAR PK46687
- add PK67554 and PK64951 for NFM)

"Recommendation*: Do not manually insert data into or delete data from this table, it is intended to be used only by optimization tools.“

Compare Access Paths – Schema changes

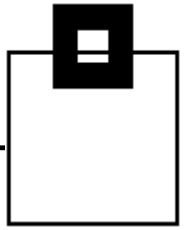

How to build an entry yourself:

DSN_VIRTUAL_INDEXES*	SYSINDEXES	Description
TBCREATOR	TBCREATOR	Auth. ID of owner/schema of table on which the entry is simul.
TBNAME	TBNAME	Name of the table on which the entry is being simulated
IXCREATOR	IXCREATOR	Auth. ID/schema of the owner of the index
IXNAME	IXNAME	Name of the index to simulate
ENABLE		Whether entry will be considered ('Y') or not ('N')
MODE		Whether the index is being created ('C') or dropped ('D')
UNIQUERULE	UNIQUERULE	Index is uniqueness: D for No (duplicates are allowed); U for Yes
COLCOUNT	COLCOUNT	The number of columns in the key
CLUSTERING	CLUSTERING	Whether the index is clustered ('Y' or 'N')
NLEAF	NLEAF	# of active leaf pages in the index, or -1 if unknown
NLEVELS	NLEVELS	# of levels in the index tree, or -1 if unknown
INDEXTYPE	INDEXTYPE	The index type: '2' - NPSI; 'D' – DPSI
PGSIZE	PGSIZE	Size, in bytes, of the leaf pages in the index: 4K, 8K, 16K, 32K
FIRSTKEYCARDF	FIRSTKEYCARDF	# of distinct values of the first key column, or -1 if unknown
FULLKEYCARDF	FULLKEYCARDF	# of distinct values of the key, or -1 if unknown
CLUSTERRATIOF	CLUSTERRATIONF	Clustering ratio, or -1 if unknown
PADDED	PADDED	Index keys padded for varying-length column data ('Y' or 'N')
COLNO1		Column # of the first column in the index key
ORDERING1		Ordering ('A' or 'D') of the first column in the index key
COLNOn		Column # repeated up to 64
ORDERINGn		Ordering ('A' or 'D') repeated up to 64

*Find DDL in member DSNTIJOS of your SDSNSAMP

! needs to have the same schema name (authid) as the PLAN_TABLE !

Compare Access Paths – Schema changes

How to build an entry yourself:

The easiest way for drop simulation is to insert from sysindexes:

```
INSERT INTO <qualifier>.DSN_VIRTUAL_INDEXES
   (TBCREATOR, TBNAME, IXCREATOR, IXNAME, ENABLE, MODE,
   UNIQUERULE, COLCOUNT, CLUSTERING, NLEAF, NLEVELS, INDEXTYPE,
   PGSIZE, FIRSTKEYCARDF, FULLKEYCARDF, CLUSTERRATIOF, PADDED
   , COLNO1, ORDERING1
   , COLNO2, ORDERING2)
   SELECT TBCREATOR, TBNAME, CREATOR, NAME, 'Y', 'D',
   UNIQUERULE, COLCOUNT, CLUSTERING, NLEAF, NLEVELS, '2', 4,
   FIRSTKEYCARDF, FULLKEYCARDF, CLUSTERRATIOF, PADDED
   , 30, 'A'
   , 32, 'A'
   FROM SYSIBM.SYSINDEXES
   WHERE CREATOR = '<ixcreator>'
   AND NAME = '<ixname>'
   AND TBCREATOR = '<tbcreator>'
   AND TBNAME = '<tbname>'
```

active

drop

pgsize

type 2 IX

colno [1-64], ordering

Your desired index for drop simulation

Compare Access Paths – Schema changes

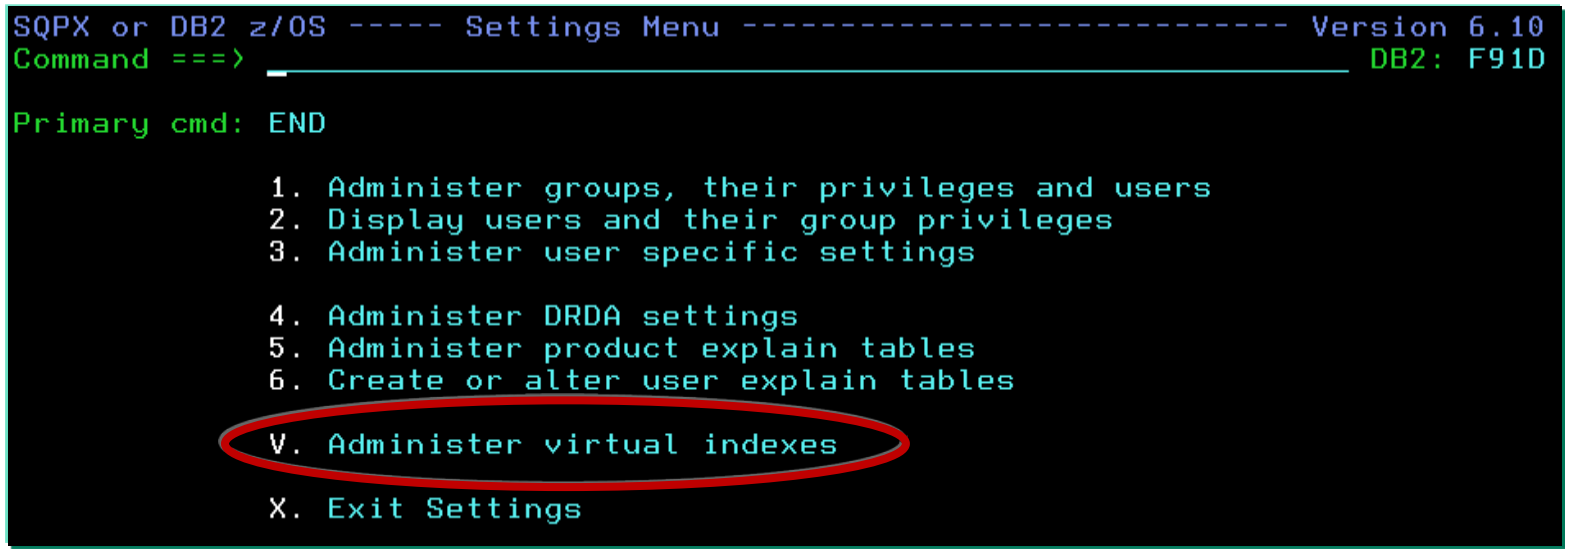

Compare Access Paths – Schema changes

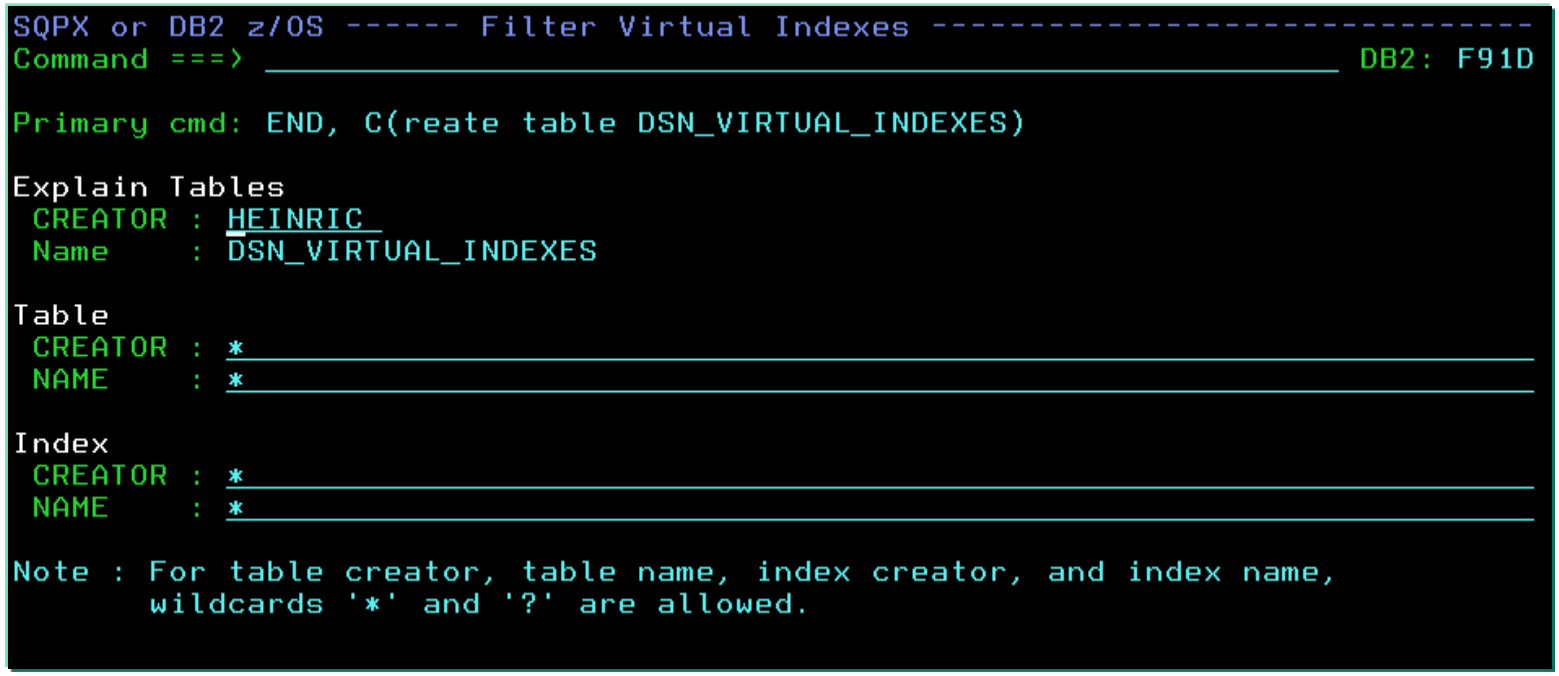

Compare Access Paths – Schema changes

```
SQPX or DB2 z/OS ----- Virtual Index Overview ----------------- Index 1 from 2
Command ===> _____________________________________________ Scroll ===> CSR
                                                            DB2: F91D

Primary cmd: END, N(ew), C(opy), L(ocate) index creator
Line    cmd: C(olumns), D(isable), E(nable), R(emove), U(pdate), Z(oom)

    Explain table used : HEINRIC.DSN_VIRTUAL_INDEXES

    Index     Index                Table     Table
    Creator   Name                 Creator   Name                 Enable  Mode
    --------  -------------------  --------  -------------------  ------  ------
          +                    +         +                    +
_   HEINRIC   VRT_IX_20110809212   SYSIBM    SYSDATABASE          YES     CREATE
_   SYSIBM    DSNDDH01             SYSIBM    SYSDATABASE          YES     DROP
    --------  -------------------  --------  -------------------  ------  ------
```

Compare Access Paths – Schema changes

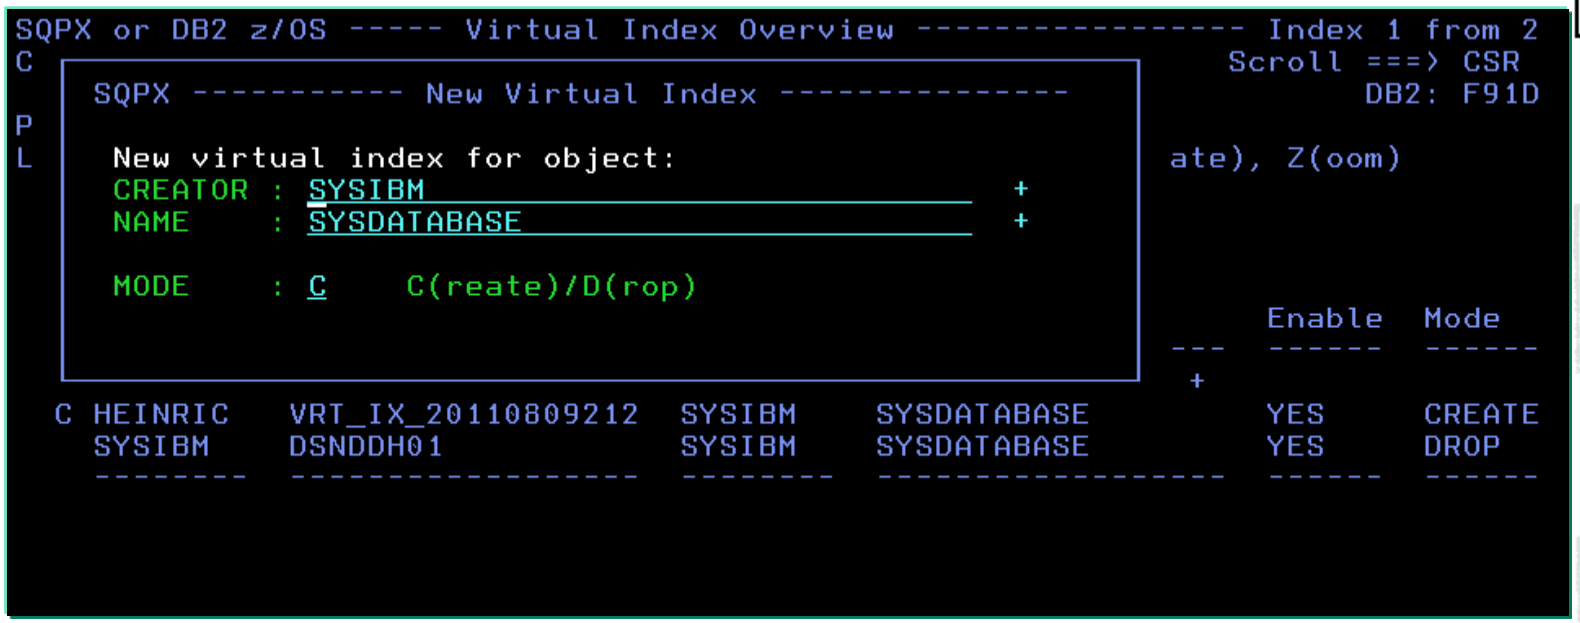

Compare Access Paths – Schema changes

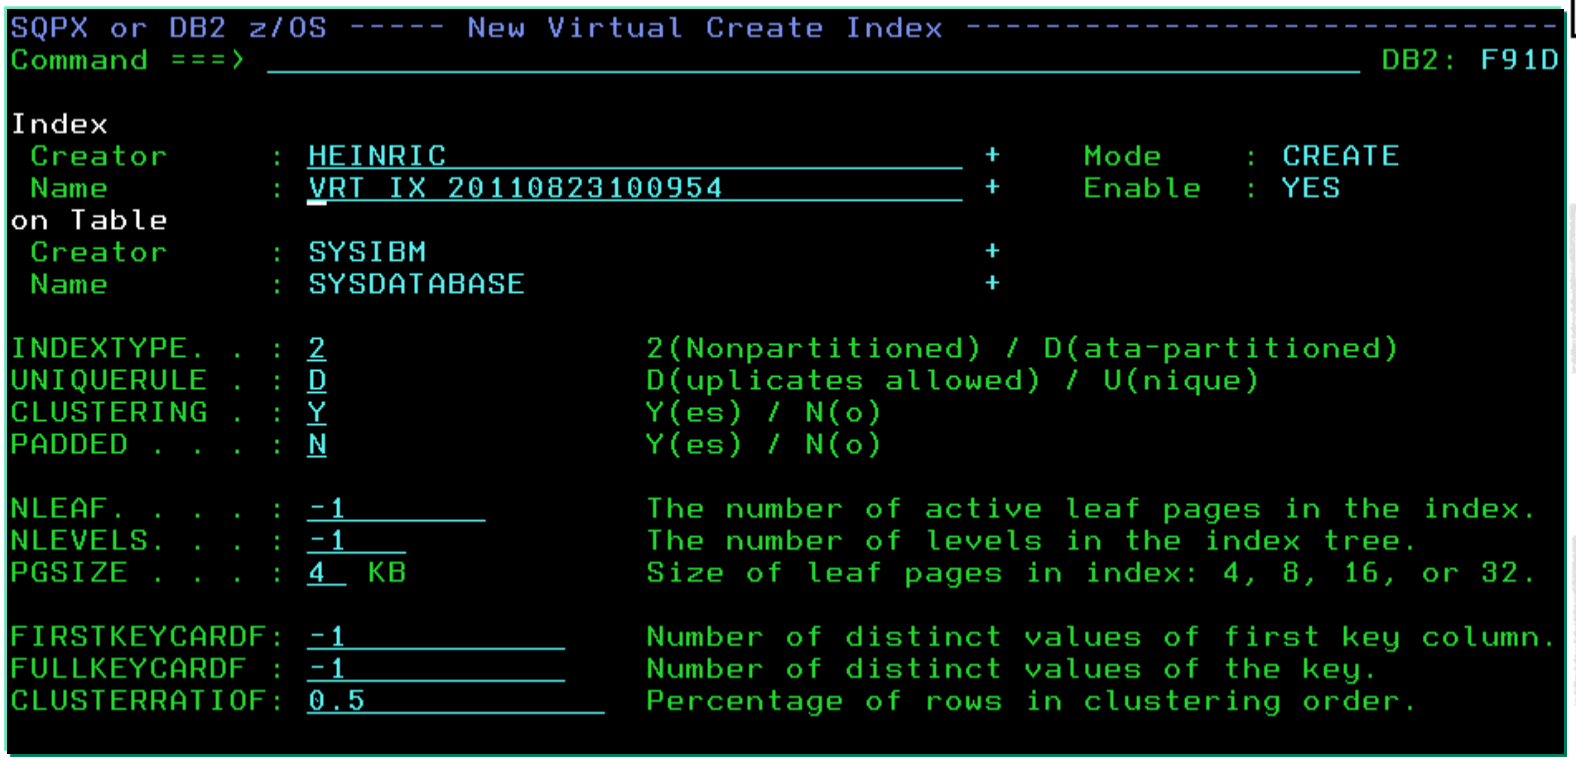

Compare Access Paths – Schema changes

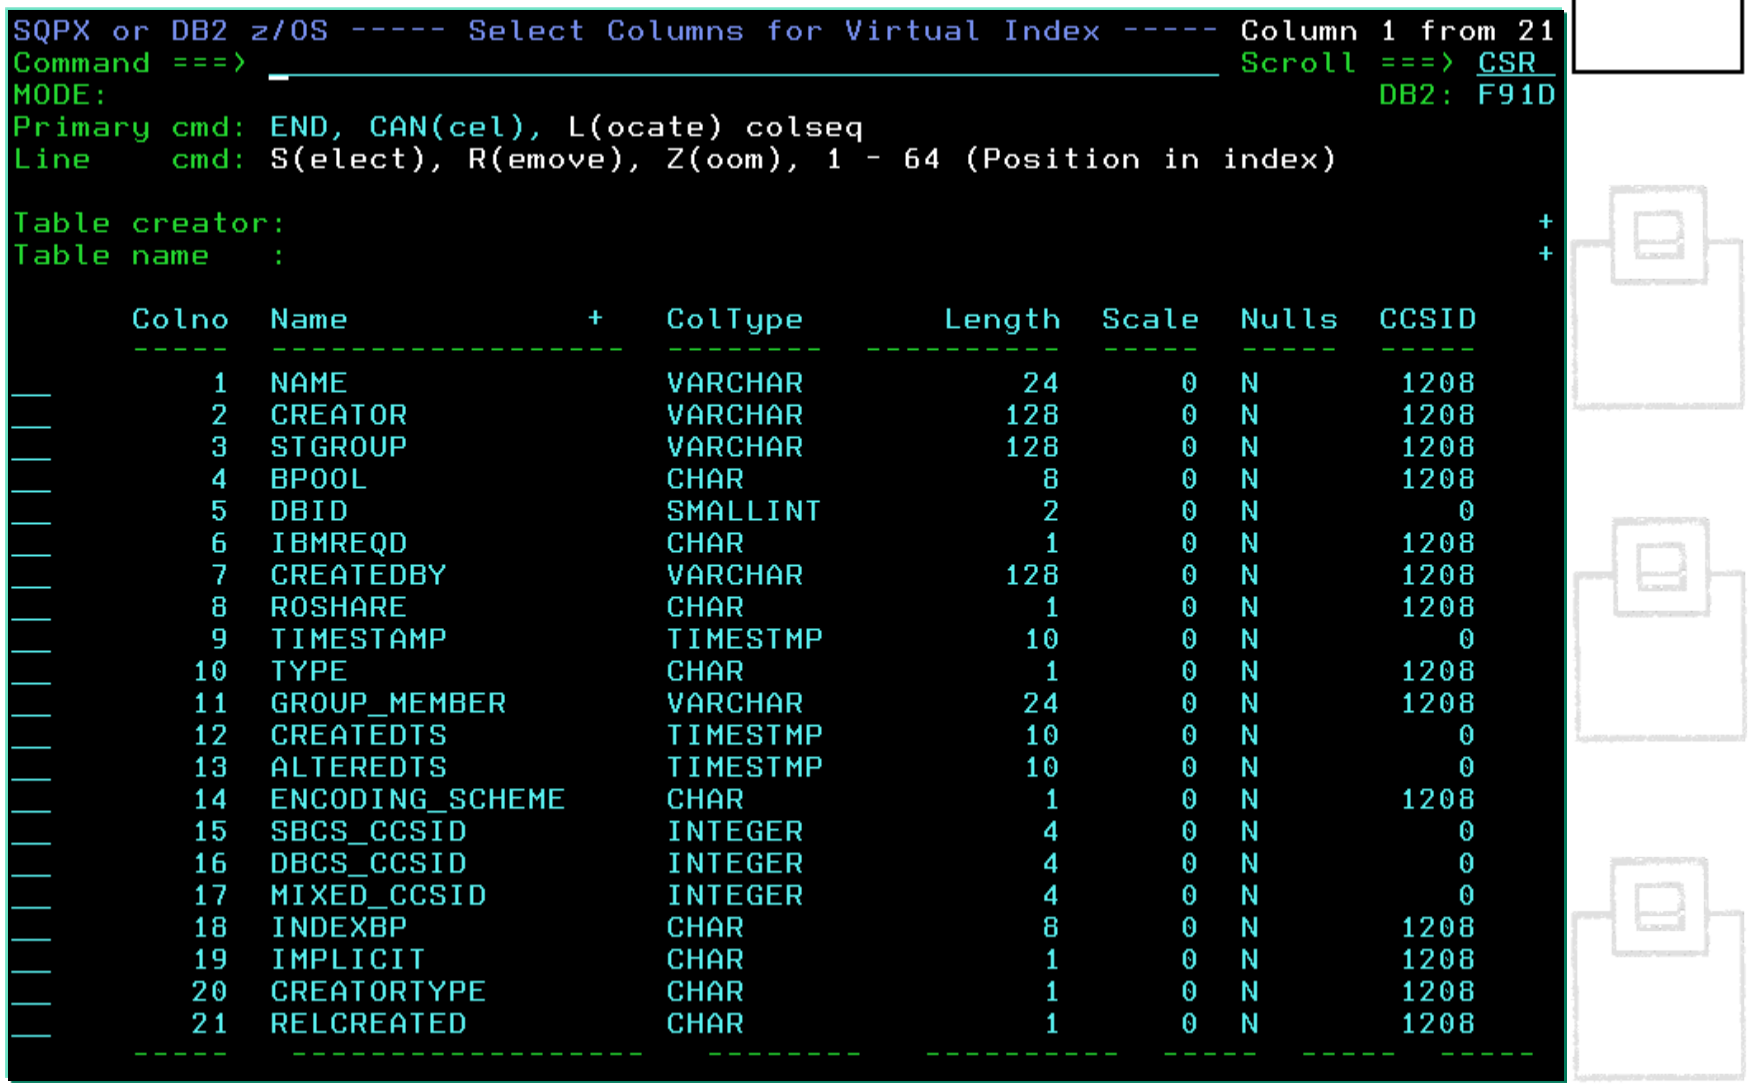

```
SQPX or DB2 z/OS ----- Select Columns for Virtual Index ----- Column 1 from 21
Command ===> _                                               Scroll ===> CSR
MODE:                                                        DB2: F91D
Primary cmd: END, CAN(cel), L(ocate) colseq
Line    cmd: S(elect), R(emove), Z(oom), 1 - 64 (Position in index)

Table creator:                                                              +
Table name   :                                                              +

      Colno  Name              +   ColType       Length  Scale  Nulls  CCSID
      -----  ------------------  --------  ----------  -----  -----  -----
__        1  NAME                VARCHAR           24      0  N       1208
__        2  CREATOR             VARCHAR          128      0  N       1208
__        3  STGROUP             VARCHAR          128      0  N       1208
__        4  BPOOL               CHAR               8      0  N       1208
__        5  DBID                SMALLINT           2      0  N          0
__        6  IBMREQD             CHAR               1      0  N       1208
__        7  CREATEDBY           VARCHAR          128      0  N       1208
__        8  ROSHARE             CHAR               1      0  N       1208
__        9  TIMESTAMP           TIMESTMP          10      0  N          0
__       10  TYPE                CHAR               1      0  N       1208
__       11  GROUP_MEMBER        VARCHAR           24      0  N       1208
__       12  CREATEDTS           TIMESTMP          10      0  N          0
__       13  ALTEREDTS           TIMESTMP          10      0  N          0
__       14  ENCODING_SCHEME     CHAR               1      0  N       1208
__       15  SBCS_CCSID          INTEGER            4      0  N          0
__       16  DBCS_CCSID          INTEGER            4      0  N          0
__       17  MIXED_CCSID         INTEGER            4      0  N          0
__       18  INDEXBP             CHAR               8      0  N       1208
__       19  IMPLICIT            CHAR               1      0  N       1208
__       20  CREATORTYPE         CHAR               1      0  N       1208
__       21  RELCREATED          CHAR               1      0  N       1208
      -----  ------------------  --------  ----------  -----  -----  -----
```

Compare Access Paths – Schema changes

```
 ImpactExpert for DB2 z/OS ----- Access Paths ------- LINE 00000001 COL 001 080
Command ===> _____________________________________________ Scroll ===> CSR
Mode: REBIND ANALYSIS WITH/WITHOUT MIGRATION RULES                  DB2: Q91A
Primary cmd: END, C(atalog data), D(etails on/off), S(tatement text), T(ables),
             I(All Indexes), X(Used indeXes)

Collection . DDLGEN                Timestamp. . 2011-02-25-14.09.29.890000
Package. . . SQLDDLD               Contoken . . 18E9104F0C187370
StmtNo old . 2262                  Bindtime . . 2011-02-25-14.08.59.956412
StmtNo new .

Access path before REBIND ----------- ! Access path with REBIND -------------
                                      !
TABLE             QB PN AC MA ME IX PR ! TABLE             QB PN AC MA ME IX PR
  INDEX                TY CO TH ON FT !   INDEX                TY CO TH ON FT
---------------- -- -- -- -- -- -- -- ! ---------------- -- -- -- -- -- -- --
SYSTABLES         1  1 I   1  0 N  S  ! SYSTABLES         1  1 I   3  0 N
  DSNDTX02                            !   VRT_IX_20110218115 VIRTUALLY CREATED
                  1  2     0  3 N     !
                                      !
Milliseconds:           5             ! Milliseconds:           1
Serviceunits:          10             ! Serviceunits:           2
--------------------------------------------------------------------------------
****************************** Bottom of Data ********************************
```

Compare Access Paths Metadata

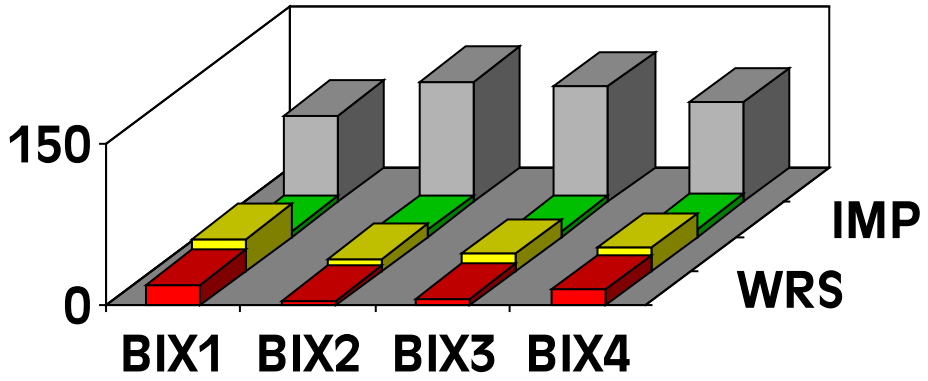

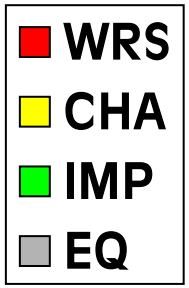

- Here are four batch runs with normal data spread
- The question is: Where are the data deltas and especially

Which combination gives the BEST results

- The answer is: Metadata analysis

Compare Access Paths Metadata

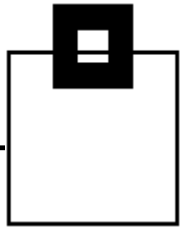

```
BIX6MENU ert for DB2 z/OS ------ Main Menu ---------------------- Version 6.10
Command ===> _____________________________________________________ DB2: Q91A

Primary cmd: END, S(ettings), C(leanup), A(bout), F(ilter Jobs), H(istory)
Line    cmd: S(elect), I(nfo), F(ilter Jobs)

                                                    Dyn  Migr. Convert
  Scenario                  Base         Recent     Expl Rules Qual. DRDA VOX
  ------------------------  -----------  ---------- ---- ----- ----- ---- ---
_ REBIND Analysis           Catalog      -          YES  N     -     -    -
_ Pre-BIND  Local           Catalog      DBRM       YES  -     -     -    -
_ Post-BIND Local           History      Catalog    NO   -     -     -    -
_ Pre-BIND  Prod-Baseline   Export (*)   DBRM       YES  -     Y     Y    N
_ Post-BIND Prod-Baseline   Export (*)   Catalog    NO   -     Y     Y    -
_ Early Precheck Static     Export (*)   -          YES  Y     N     Y    N
_ Early Precheck Dynamic    Export (*)   -          YES  Y     Y     -    N
_ DSC Protection            Export (*)   -          YES  -     N     -    -
_ Dynamic SQL               DynStmtCache -          YES  -     -     -    -
_ Static and dynamic SQL    Trace        -          YES  -     -     -    -

_ Plan_table compare        Plan_table 1 Plan_table 2
_ DBRM reconstruct          Catalog

NOTE (*): Use export/import function to update BIX export tables.
```

Compare Access Paths Metadata

```
 ImpactExpert for DB2 z/OS ----- Job Overview (5/5) ----------- Job 16 from 113
Command ===> _____________________________________________ Scroll ===> CSR
Mode: REBIND ANALYSIS WITH/WITHOUT MIGRATION RULES                 DB2: Q91A
Primary cmd: END, A(ll), N(ew), R(efresh), Z(oom), L(ocate) submit time
Line    cmd: P(rograms), C(ompare), D(elete), G(raphic), L(PLans),
             S(tatements), Z(oom)

     Submit time         Submitter Jobname  Stepname PLTB owner  PTF Level Mode
     ------------------- --------- -------- -------- ----------  --------- ----
C    2011-02-25-14.13.14 BOXWELL   RBNDPACK CHKRBIND BOXWELL     PTF_10179 RC P
C    2011-02-25-14.09.26 BOXWELL   RBNDPACK CHKRBIND BOXWELL     PTF_10179 RC P
_    2011-02-07-16.08.26 DUDEK     RBNDPALE CHKRBIND IMPALE      PTF_10179 RC P
_    2011-02-07-16.08.26 DUDEK     RBNDPAIR CHKRBIND IMPAIR      PTF_10179 RC P
_    2011-02-07-14.19.40 DUDEK     RBNDPAC2 CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-07-12.56.03 DUDEK     RBNDPAXX CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-07-12.56.03 DUDEK     RBNDPACK CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-07-11.08.49 DUDEK     RBNDPAC2 CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-07-11.08.49 DUDEK     RBNDPAC1 CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-03-14.32.42 DUDEK     RBNDPAC1 CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-03-14.04.36 DUDEK     RBNDPACK CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-02-16.43.14 DUDEK     RBNDPACK CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-02-16.19.42 DUDEK     RBNDPACK CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-02-15.57.27 DUDEK     RBNDPACK CHKRBIND IMPACT      PTF_10179 RC P
_    2011-02-02-14.04.59 DUDEK     RBNDPACK CHKRBIND IMPACT      PTF_10179 RC P
```

Compare Access Paths Metadata

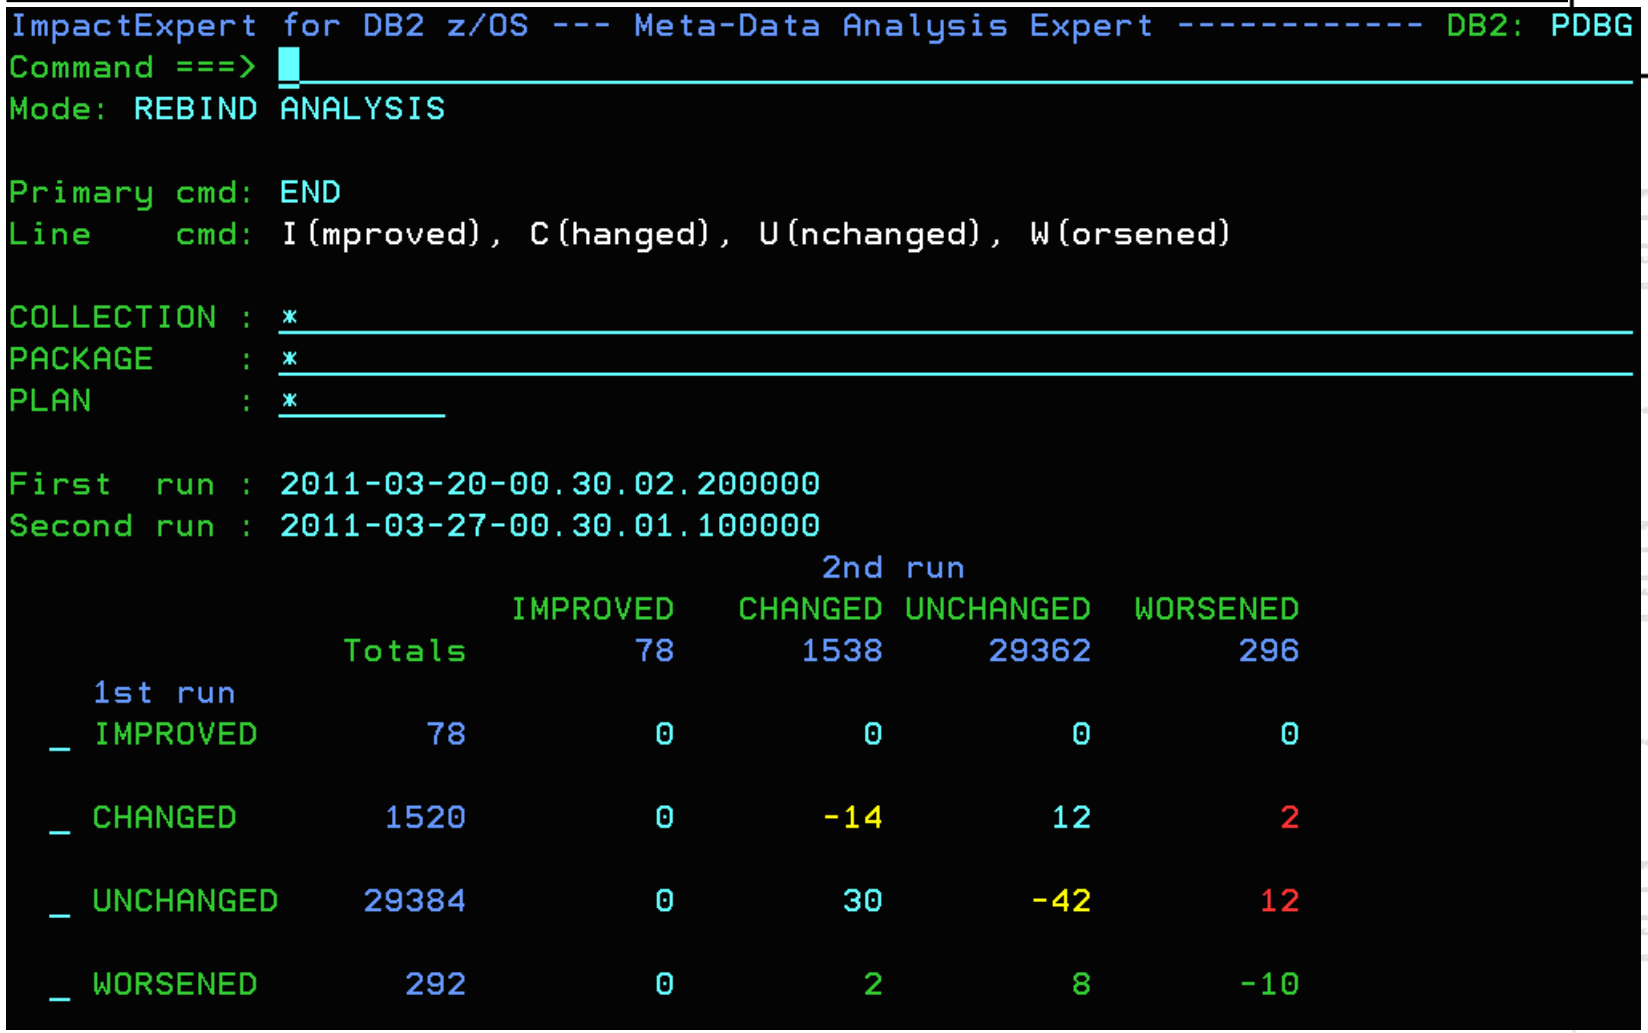

```
ImpactExpert for DB2 z/OS --- Meta-Data Analysis Expert ------------ DB2: PDBG
Command ===>
Mode: REBIND ANALYSIS

Primary cmd: END
Line    cmd: I(mproved), C(hanged), U(nchanged), W(orsened)

COLLECTION : *
PACKAGE    : *
PLAN       : *

First  run : 2011-03-20-00.30.02.200000
Second run : 2011-03-27-00.30.01.100000
                                    2nd run
                      IMPROVED    CHANGED UNCHANGED  WORSENED
                Totals       78       1538     29362       296
    1st run
  _ IMPROVED         78          0          0          0          0

  _ CHANGED        1520          0        -14         12          2

  _ UNCHANGED     29384          0         30        -42         12

  _ WORSENED        292          0          2          8        -10
```

Compare Access Paths Metadata

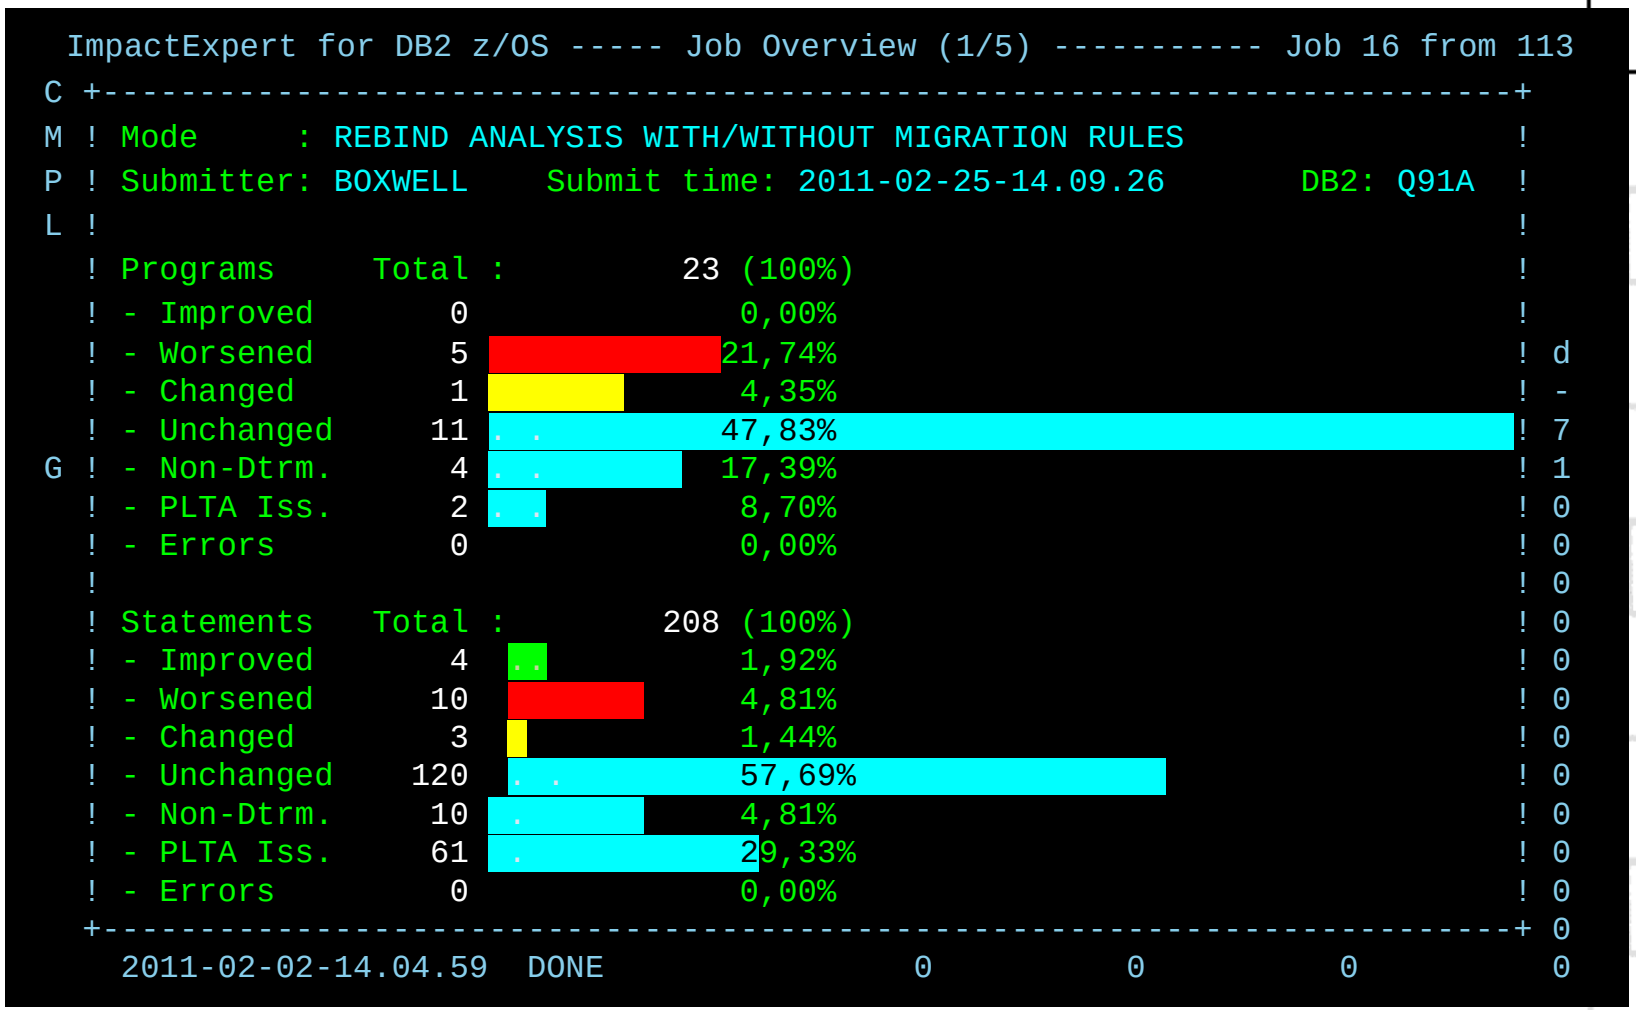

Compare Access Paths Metadata

```
 ImpactExpert for DB2 z/OS --- Meta-Data Analysis Expert ------------ DB2: Q91A
Command ===> ____________________________________________________________________
Mode: REBIND ANALYSIS WITH/WITHOUT MIGRATION RULES

Primary cmd: END
Line    cmd: I(mproved), C(hanged), U(nchanged), W(orsened)

COLLECTION : *___________________________________________________________________
PACKAGE    : *___________________________________________________________________
PLAN       : *________
First  run : 2011-02-25-14.09.26    Second run : 2011-02-25-14.13.14
                                        2nd run
                       IMPROVED   CHANGED UNCHANGED  WORSENED
               Totals         0         0       137         0
    1st run
  U IMPROVED         4        -4         0         4         0

  _ CHANGED          3         0        -3         3         0

  _ UNCHANGED      120         0         0         0         0

  _ WORSENED        10         0         0        10       -10
```

Compare Access Paths Metadata

```
 ImpactExpert for DB2 z/OS -- Statement Overview ---------------- Stmt 1 from 4
Command ===> _____________________________________________ Scroll ===> CSR
Mode: REBIND ANALYSIS WITH/WITHOUT MIGRATION RULES                 DB2: Q91A
Primary cmd: END, Z(oom), L(ocate) collid/plan
Line    cmd: S(tatement text), Z(oom), 1(First run), 2(Second run)

First  run : 2011-02-25-14.09.26    Second run : 2011-02-25-14.13.14

    Collection/Plan +   Program +   Version/Contoken     +     StmtNo    SectNo
    ------------------  ----------  ------------------------  --------  --------
2 1 DDLGEN              SQLDDLD     CONTOKEN=18E9104F0C1873      2262         4
_   DDLGEN              SQLDDLD     CONTOKEN=18E9104F0C1873      4642        23
_   DDLGEN              SQLDDLS     CONTOKEN=18E33349101C91      1862        21
_   DDLGEN              SQLDDLS     CONTOKEN=18E33349101C91      2007        23
    ------------------  ----------  ------------------------  --------  --------
```

Compare Access Paths Metadata

```
 ImpactExpert for DB2 z/OS ----- Access Paths ------- LINE 00000001 COL 001 080
Command ===> _____________________________________________ Scroll ===> CSR
Mode: REBIND ANALYSIS WITH/WITHOUT MIGRATION RULES                 DB2: Q91A
Primary cmd: END, C(atalog data), D(etails on/off), S(tatement text), T(ables),
             I(All Indexes), X(Used indeXes)

Collection . DDLGEN                Timestamp. . 2011-02-25-14.13.17.060000
Package. . . SQLDDLD               Contoken . . 18E9104F0C187370
StmtNo old . 2262                  Bindtime . . 2011-02-25-14.08.59.956412
StmtNo new .

Access path before REBIND ----------- ! Access path with REBIND -------------
                                      !
TABLE             QB PN AC MA ME IX PR ! TABLE             QB PN AC MA ME IX PR
  INDEX                TY CO TH ON FT !   INDEX                TY CO TH ON FT
----------------- -- -- -- -- -- -- -- ! ----------------- -- -- -- -- -- -- --
SYSTABLES          1  1 I   1  0 N  S  ! SYSTABLES          1  1 I   1  0 N  S
  DSNDTX02                            !   DSNDTX02
                   1  2     0  3 N     !                    1  2     0  3 N
                                      !
Milliseconds:            5            ! Milliseconds:            5
Serviceunits:           10            ! Serviceunits:           10
-------------------------------------------------------------------------------
******************************* Bottom of Data ********************************
```

Conclusion

- Precheck any changes
  - Environment
  - Schema
  - Application
- Allow flexibility in developing and running your applications
- Insight out of the box into the effect of your desired changes before you apply them
- Cost efficiency and cost avoidance in development and operations

Assure Best Practices and precheck!

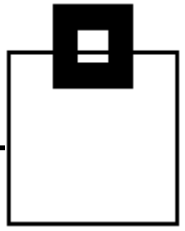

Before you apply

- Hardware modifications
- Environment setup changes
- New application release
- New DB2 code
  - Version
  - APAR
- Index modifications
  - Create
  - Drop
  - Change

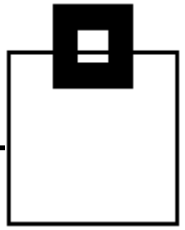

Ulf Heinrich

SEGUS, Inc

u.heinrich@segus.com