Partitioning…
and Relative Page Numbering
on Db2 12 for z/OS

David Simpson
dsimpson@themisinc.com

Themis

Themis Education

- Most complete DB2 Curriculum in the industry
- Offerings include a complete mainframe curriculum in addition to Oracle, Java, .NET, Linux, UNIX, etc.
- Training Venues:
  - At your facility for groups
  - Public enrollment at various locations in the USA
  - Distance learning public enrollment with live instructors.
- Webinars:
  - Visit www.themisinc.com/webinars for upcoming schedule and replays of past webinars and to download today's slides.
- This material comes from our Db2 12 migration course
  - View Outline and Register

Themis

David Simpson

David Simpson is currently the Vice President of Themis Inc. He teaches courses on SQL, Application Programming, Database Administration as well as optimization, performance and tuning. He also installs and maintains the database systems used for training at Themis and works with our network of instructors to deliver high quality training solutions to our customers worldwide.

Since 1993 David has worked as a developer and DBA in support of very large transactional and business intelligence systems. David is a certified DB2 DBA on both z/OS and LUW. David was voted Best User Speaker and Best Overall Speaker at IDUG North America 2006. He was also voted Best User Speaker at IDUG Europe 2006 and is a member of the IDUG Speakers Hall of Fame. David is also an IBM Gold Consultant.

dsimpson@themisinc.com
www.themisinc.com
www.linkedin.com/in/davidcsimpson/
@ThemisDave
@ThemisTraining

Find my work:
www.themisinc.com/webinars
www.idug.org/content

Themis

Partitioned Table Space Evolution

- **Index Controlled Range Partitioning**

- **Table Controlled Range Partitioning**

deprecated

- ***Universal Partition by Range (PBR)***
- ***Universal Partition by Growth (PBG)****

- ***Universal Partition by Range (PBR)*** **PAGENUM ABSOLUTE**
- ***Universal Partition by Range (PBR)*** **PAGENUM RELATIVE**

**For PBG fun see Julia Carter's article on the IDUG Beginner's Blog:*
https://www.idug.org/p/bl/et/blogaid=825

Themis

Possible Conversions

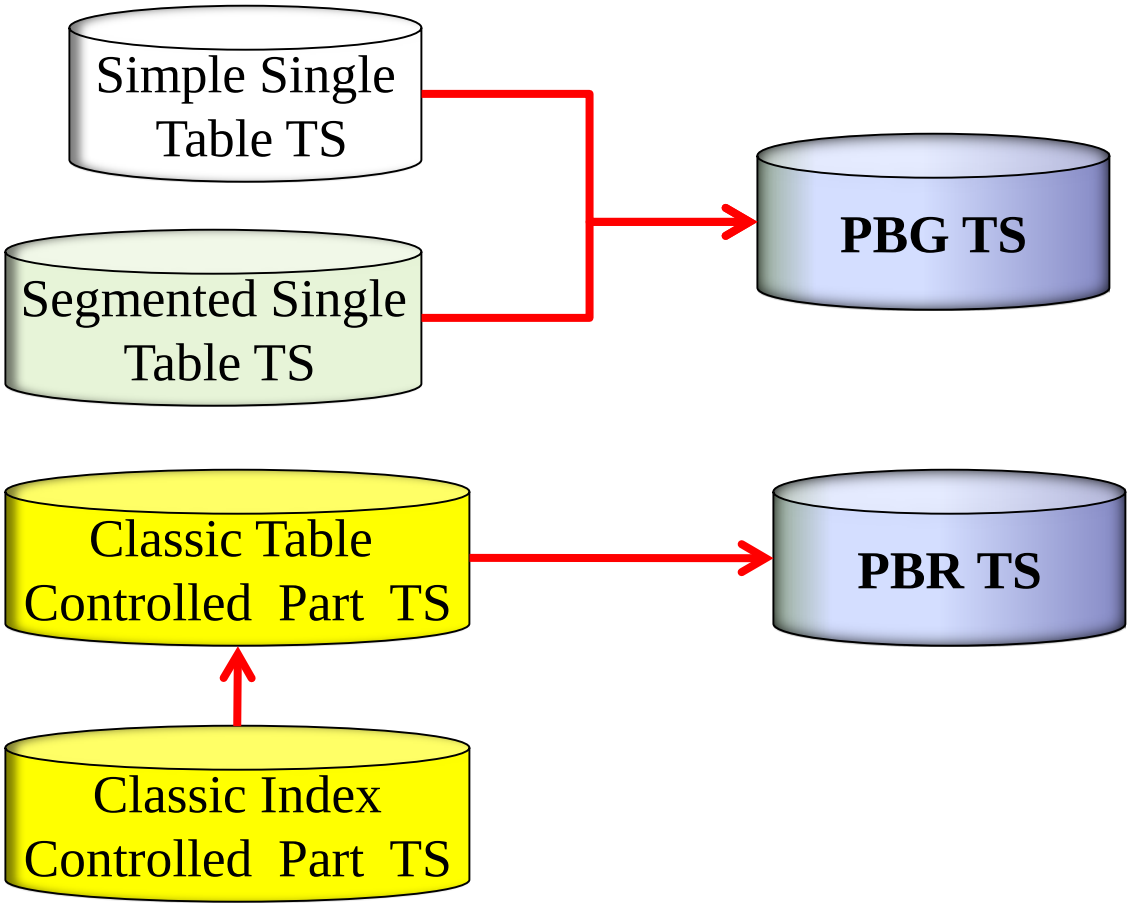

Themis

Partitioning… Last Limitkey

Last limitkey is enforced IF any of the following are true:

- The tablespace was created with the LARGE keyword
- The tablespace has more than 64 partitions
- The DSSIZE is 4G or greater
- Table controlled partitioning was used

Any of these things also lead to a 5 byte RID

Themis

Convert Index to Table Partitioning

```
      ALTER INDEX XEMPLS10
            NOT CLUSTER;
DSNT404I SQLCODE = 20272, WARNING:  TABLE SPACE CEMPLS10 HAS BEEN CONVERTED TO
         USE TABLE-CONTROLLED PARTITIONING INSTEAD OF INDEX-CONTROLLED
         PARTITIONING, ADDITIONAL INFORMATION: 2100000000
DSNT418I SQLSTATE   = 01666 SQLSTATE RETURN CODE
DSNT415I SQLERRP    = DSNXISB6 SQL PROCEDURE DETECTING
DSNT416I SQLERRD    = 42 0  0  -1  0  0 SQL DIAGNOST
DSNT416I SQLERRD    = X'0000002A'  X'00000000'  X'0
         X'00000000'  X'00000000' SQL DIAGNOSTIC IN
```

Clustering Index

Last limitkey was altered

Would be “*N” if alter was not necessary

Beware of converting any other way… especially add/rotate partition!!!

Themis

Convert Index to Table Partitioning

```
      ALTER INDEX XEMPLS10
            CLUSTER;
--------+---------+--------+---------+---------+---------+---------+--
DSNE616I STATEMENT EXECUTION                    LCODE IS 0
--------+---------+--------+---------+---------+---------+---------+--
--------+---------+--------+---------+---------+---------+---------+--
DSNE617I COMMIT PERFORMED, SQLCODE IS 0
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
--------+---------+--------+---------+---------+---------+---------+--
```

Put things back the way they were

Themis

Convert to PBR

```
      ALTER TABLESPACE DAVE.CEMPLS10
        SEGSIZE 32;
--------+---------+---------+---------+---------+---------+---------+---------+-------
DSNT404I SQLCODE = 610, WARNING:  A CREATE/ALTER ON OBJECT DAVE.CEMPLS10 HAS
         PLACED OBJECT IN ADVISORY REORG PENDING
DSNT418I SQLSTATE   = 01566 SQLSTATE RETURN
DSNT415I SQLERRP    = DSNXI14 SQL PROCEDU
DSNT416I SQLERRD    = 250 0  0  -1  0  0 SQL                    TION
DSNT416I SQLERRD    = X'000000FA'  X'0000000                    FFFFFFFF'
         X'00000000'  X'00000000' SQL DIAGNO
--------+---------+---------+---------+----                    -------+-------
```

Tablespace level reorg needed to convert

Themis

Table Controlled Partitioning (applies to PBR too!)

```
   CREATE TABLESPACE TSUPBR
    NUMPARTS 10
    SEGSIZE 32
     IN  THEMISDB
      USING STOGROUP WORKSHOP PRIQTY 720 SECQTY -1
     BUFFERPOOL BP2
     DEFINE NO
     LOCKSIZE PAGE LOCKMAX 0;
---------+---------+---------+---------+---------+---------+---------
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
---------+---------+---------+---------+---------+---------+---------
```

Themis

Table Controlled Partitioning (applies to PBR too!)

```
CREATE TABLE THEMIS.PEOPLE
(ID  INTEGER NOT NULL
,LAST_NAME  CHAR(20) NOT NULL
,FIRST_NAME CHAR(20) NOT NULL
,ZIP_CODE   CHAR(5) NOT NULL
,BIRTH_DTE  DATE NOT NULL
,PRIMARY KEY(ID))
PARTITION BY (ZIP_CODE)
(PARTITION 1 ENDING AT ('09999'), PARTITION  2 ENDING AT ('19999')
,PARTITION 3 ENDING AT ('29999'), PARTITION  4 ENDING AT ('39999')
,PARTITION 5 ENDING AT ('49999'), PARTITION  6 ENDING AT ('59999')
,PARTITION 7 ENDING AT ('69999'), PARTITION  8 ENDING AT ('79999')
,PARTITION 9 ENDING AT ('89999'), PARTITION 10 ENDING AT ('99999')
) IN THEMISDB.TSUPBR;
---------+---------+---------+---------+---------+---------+---------+--
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
---------+---------+---------+---------+---------+---------+---------+--
```

Limitkeys at the table level

Themis

Partitioning Index

```
CREATE INDEX XPART1 ON THEMIS.PEOPLE(ZIP_CODE,ID)
  PARTITIONED
  CLUSTER
  USING STOGROUP WORKSHOP
    PRIQTY 48 SECQTY -1
  PCTFREE 20 FREEPAGE 0
  ;
---------+---------+---------+---------+---------+---------+---------
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
---------+---------+---------+---------+---------+---------+---------
```

Partitioned index leads with the partitioning key

Themis

Another Partitioning Index

```
CREATE INDEX XPART2 ON THEMIS.PEOPLE(ZIP_CODE,LAST_NAME)
  PARTITIONED
  USING STOGROUP WORKSHOP
    PRIQTY 48 SECQTY -1
  PCTFREE 20 FREEPAGE 0
  ;
---------+---------+---------+-----+---------+---------
DSNE616I STATEMENT EXECUTION WAS SU      CODE IS 0
---------+---------+---------+---------+---------+---------
```

Partitioned index leads with the partitioning key

Themis

Non Partitioned Secondary Index (NPSI)

```
CREATE UNIQUE INDEX XPART3 ON THEMIS.PEOPLE(ID)
   USING STOGROUP WORKSHOP
     PRIQTY 48 SECQTY -1
   PCTFREE 20 FREEPAGE 0
   ;
--------+---------+---------+---------+---------+---------
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
--------+---------+---------+---------+---------+---------+---------
```

No keyword PARTITIONED So… NPSI

Themis

Data Partitioned Secondary Index (DPSI)

```
CREATE INDEX XPART4 ON THEMIS.PEOPLE(FIRST_NAME)
  PARTITIONED
  USING STOGROUP WORKSHOP
    PRIQTY 48 SECQTY -1
  PCTFREE 20 FREEPAGE 0
  ;
---------+---------+---------+---------+---------+---------+--------
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
---------+---------+---------+---------+---------+---------+--------
```

PARTITIONED
But does not lead with the partitioning key

Themis

Db2 12 PBR RPN Tablespaces

- Partition-by-range relative-page-number
- RPN represents the relative page number without embedded partition numbers
- 7-byte RID (2-byte part #, 4-byte page #1-byte row)
- DSSIZE per partition now 1TB with a max table size of 4PB for up to 256 trillion rows
- Index structure changed to support PBR RPN
- Utilities and Log modified to support PBR RPN
- V12R1M500 & above

Themis

CREATE TABLESPACE Changes

CREATE TABLESPACE

PAGENUM ABSOLUTE | RELATIVE –PBR only

DSSIZE *1g – 1024g*

PARTITION *integer* **DSSIZE *1g – 1024g***

Yes, DSSIZE can be at the partition level IF you use RPN!

Themis

NUMPARTS Limitations for Absolute

DSSIZE value	4K page size	8K page size	16K page size	32K page size
1G - 4G (1 GB to 4 GB)	4096	4096	4096	4096
8G (8 GB)	2048	4096	4096	4096
16G (16 GB)	1024	2048	4096	4096
32G (32 GB)	512	1024	2048	4096
64G (64 GB)	254	512	1024	2048
128G (128 GB)	128	256	512	1024
256G (256 GB)	64	128	256	512

Source: Db2 12 SQL Reference Guide

Assumes 5 byte RID using VSAM EA

Themis

Example: Absolute Page Numbering

```
CREATE TABLESPACE TSEMP
      IN THMC01EM
      NUMPARTS 4
      DSSIZE 1G
      SEGSIZE 32
      PAGENUM ABSOLUTE
      USING STOGROUP WORKSHOP
            PRIQTY 720 SECQTY -1
      PCTFREE 10
      FREEPAGE 0
      LOCKSIZE PAGE
      BUFFERPOOL BP2
      TRACKMOD NO
      CLOSE YES
      COMPRESS NO
;
```

Default controlled by zParm PAGESET_PAGENUM

Themis

Example: Absolute Page Numbering

```
//PRTEMP2  EXEC PGM=DSN1PRNT,
//         PARM='FORMAT,PRINT'
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSUT1   DD DSN=DC1B.DSNDBC.THMC01EM.TSEMP.I0001.A002,DISP=SHR
/*
```

```
PARTITION: # 0002
PAGE: # 00040000 ------------------------------
HEADER PAGE:   PGCOMB='10'X  PGBIGRBA='00000000
               HPGOBID='01010004'X  HPGHPREF='0
               HPGCATV='00'X  HPGTORBA='00000000
               HPGSSNM='DC1B'  HPGFOID='0003'X
               HPGZ3PNO='000000'X  HPGZNUMP='00
```

Themis

Example: Absolute Page Numbering

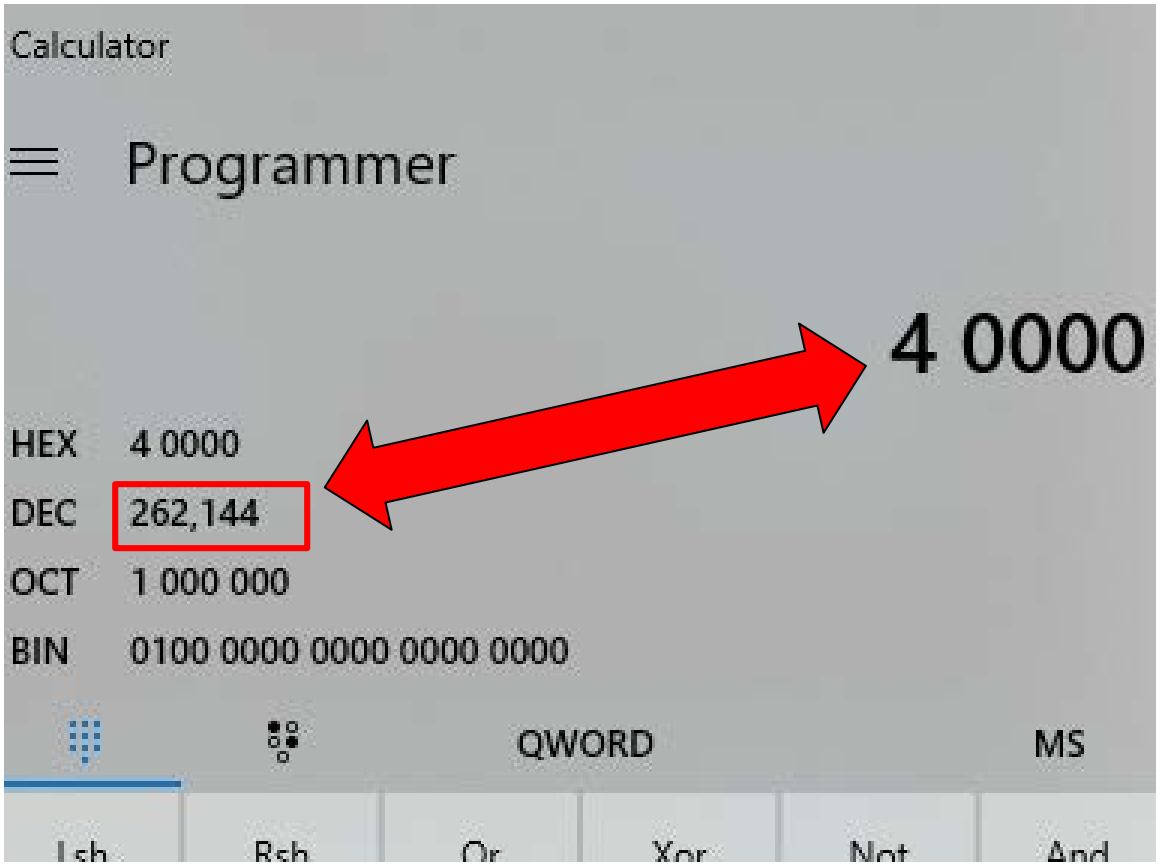

262,144 * 4 = 1,048,576k

In other words….

1G

Themis

Absolute Page Numbering Facts

- The DSSIZE will cause Db2 to reserve the appropriate number of page numbers in each partition
- Given the DSSIZE, Db2 can figure out which partition each page lives in by doing math
- Changing the DSSIZE results in renumbering every page after the first partition

Converting to RPN

```
---------+---------+---------+---------+---------+---------+---------+---------+---
      ALTER TABLESPACE THMC01EM.TSEMP PAGENUM RELATIVE                          0(
---------+---------+---------+---------+---------+---------+---------+---------+---
DSNT404I SQLCODE = 610, WARNING:  A CREATE/ALTER ON OBJECT THMC01EM.TSEMP
         PLACED OBJECT IN ADVISORY REORG PENDING
DSNT418I SQLSTATE   = 01566 SQLSTATE RETURN CODE
```

```
PARTITION: # 0002
PAGE: # 00000000 ------------------------------
HEADER PAGE:  PGCOMB='00'X  PGBIGRBA='00000000
              HPGOBID='01010004'X  HPGHPREF='0
              HPGCATV='00'X  HPGTORBA='0000000
              HPGSSNM='DC1B'  HPGFOID='0003'X
              HPGZ3PNO='000000'X  HPGZNUMP='00
```

Themis

PBR RPN Facts

- Altering DSSIZE becomes an immediate change (but packages are still invalidated!)
- DSSIZE may be altered at the partition level
- Maximums
  - Max DSSIZE is 1Tb
  - Max Tablespace is 4Pb / 256 trillion rows
- Eliminates the limitations shown on slide 9

Themis

PBR RPN Indexes

- Partitioned Index (PI)
  - Max partition size 1TB, DSSIZE keyword supported
  - Index partition size is independent from the tablespace
  - RID for index entry 7 bytes (2-byte part#, 4-byte page#, 1-byte record-id
  - Index header records index part#, index page # does not contain part#
  - SYSINDEXES & SYSINDEXPART PAGENUM setting reflects “R” for relative page numbering
- Non-Partitioned Index (NPI)
  - RID for index entry 7 bytes (2-byte part#, 4-byte page#, 1-byte record-id
  - SYSINDEXES & SYSINDEXPART PAGENUM setting reflects “A” for absolute page numbering

Themis

PBR RPN Indexes

```
NON-UNIQUE KEYS FOLLOW:
KEY ENTRY:  IPKMAP(XI)='003E'X  IPNRIDS=9
  KEY:
  00C4F1F1                                                                .D11
  RIDS:
  C0000010000002201 0000010000002202 0000010000002203 0000010000002204 000001000
  0000010000002207 0000010000002208 0000010000002209
```

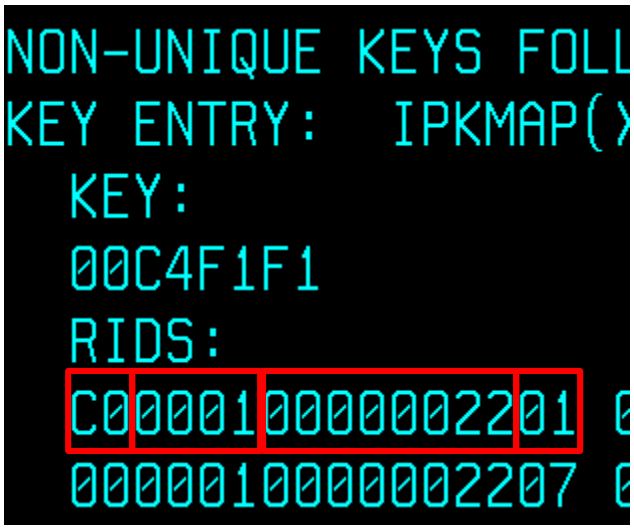

1 byte for pseudo-delete flag

2 byte partition number

4 byte page number

1 byte row number

Themis

zParms to Consider

- IX_TB_PART_CONV_EXCLUDE
- PREVENT_ALTERTB_LIMITKEY
- PREVENT_NEW_IXCTRL_PART
- DDL_MATERIALIZATION

Themis

SQL to determine partition type

```
 SELECT TS.DBNAME,
           TS.NAME,
           CASE WHEN TS.TYPE = 'R' THEN 'PBR ENFORCED'
                WHEN TS.TYPE = 'G' THEN 'PBG'
                WHEN TP.IXNAME = '' THEN 'TABLE CONTROLLED ENFORCED'
                WHEN TS.TYPE = 'L' THEN 'IX CONTROLLED ENFORCED'
                ELSE 'IX CONTROLLED NOT ENFORCED'
           END AS TYPE,
           TS.PARTITIONS
    FROM SYSIBM.SYSTABLESPACE TS JOIN
         (SELECT DISTINCT DBNAME, TSNAME, IXNAME
            FROM SYSIBM.SYSTABLEPART) TP
      ON TS.DBNAME = TP.DBNAME
     AND TS.NAME = TP.TSNAME
    WHERE TS.DBNAME NOT LIKE 'DSN%'    -- NO CATALOG / DIRECTORY
                                       -- MAY ADD OTHER EXCLUSIONS HERE
                                       -- (E.G. TOOLS, 3RD PARTY DBS)
      AND TS.PARTITIONS > 0            -- ONLY PARTITIONED THINGS
      AND TS.TYPE IN ('','R','G','L') -- NO WEIRD TYPES
ORDER BY TS.DBNAME, TS.NAME
```

Use at your own risk ☺

Themis

APAR discussed with REORG conversion issues…
PH02571

Themis

David Simpson

Themis Training

dsimpson@themisinc.com

@ThemisDave

@ThemisTraining

Themis