

May 6-10, 2007  
San Jose Convention Center  
San Jose, California, USA

Session: C06



*What's New in DB2 9 for z/OS for Administrators*

IDUG® 2007  
North America


Dr. Jim Teng  
IBM Silicon Valley Laboratory

May 8, 2007 10:40 a.m. – 11:40 a.m.

Platform: DB2 for z/OS


 

GoFurther



Do you want to learn more about DB2 9 for z/OS Administration functions? This presentation covers V9 functions: Universal Table Spaces, Clone Table, Append Table option, Converged Temporary Space, Automatic Creation of Databases, Automatic Buffer Pool Management, Online Utilities, CopyPool enhancements.

## Disclaimers & Trademarks\*



Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. Any discussion of OEM products is based upon information which has been publicly available and is subject to change. The opinions expressed are those of the presenter at the time, not necessarily the current opinion and certainly not that of the company.

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States and/or other countries: AIX, AS/400, DATABASE 2, DB2\*, Enterprise Storage Server, ESCON\*, IBM, iSeries, Lotus, NOTES, OS/400, pSeries, RISC, WebSphere, xSeries, z/Architecture, z/OS, zSeries, System p, System I, System z

The following terms are trademarks or registered trademarks of the Microsoft Corporation in the United States and/or other countries: MICROSOFT, WINDOWS, ODBC

For more copyright & trademark information see [ibm.com/legal/copytrade.phtml](http://ibm.com/legal/copytrade.phtml)

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in the operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only. This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

## Agenda

- *Universal Table Space*
- *Index Enhancements*
- *Implicit table space and database*
- *New format for variable length rows*
- *Automatic Buffer Pool Management*
- *Converged TEMP space*
- *Data Sharing Enhancements*
- *Clone Table*
- *Utility Enhancements*
- *Copy Pool Enhancements*

## Table spaces up to and including V8

- Simple
  - f* Multiple tables per table space, different tables share pages.
    - *Deprecated* as of this release.
- Segmented
  - f* Multiple tables per table space, different tables don't share pages.
  - f* Handles space usage and mass deletes well.
  - f* Limited to 64Gb (32 x 2Gb datasets).
- Partitioned
  - f* One table per table space.
  - f* Needs a partitioning column.
  - f* But doesn't handle space usage or mass deletes as well as segmented.

## Universal Table Space

- *Two Types: Partition By Range and Partition By Growth*
- *Partition By Range*
  - f* A partitioned segmented table space.
  - f* Partitioning column required.
  - f* Just like a partitioned table space today with segmented space map page format.
- *Partition By Growth*
  - f* A large single table segmented table space (> 64GB).
  - f* No partitioning column required.
  - f* On demand to add a partition due to space needs.
- Must be SMS managed EA data set if DSSIZE > 4GB

## *Universal Table Space - Benefits*

- *Improved table scan performance*
  - f I/O eliminated for empty pages*
- *Improved mass delete performance*
  - f Only space map pages are updated and logged*
  - f Improved ALTER TABLE ROTATE performance*
- Copy will not copy empty data pages caused by mass deletes
- Improved Insert performance for variable length records
  - f Better space management = Fewer REORGs*

## Partition By Range - DDL

**CREATE TABLESPACE TS1 IN DB1  
NUMPARTS 32 SEGSIZE 64;**

*Basically creates a partitioned table space but just add the **SEGSIZE** clause.*

- *Only table-controlled partitioning allowed (i.e. key limits need to specify on CREATE TABLE).*
- *Key limit on the last part is always enforced.*
- *SYSTABLESPACE has TYPE column value "R".*
- *No MEMBER CLUSTER.*
- *No ALTER SEGSIZE so get SEGSIZE right or it is a DROP/CREATE.*
- *All index RIDs will be 5 bytes.*

## Partition By Growth – DDL

- *CREATE TABLESPACE .... MAXPARTITIONS n*
  - f You can set DSSIZE and SEGSIZE.*
  - f Default 4G and SEGSIZE 4.*
- *CREATE TABLE .... PARTITION BY SIZE EVERY mG*
  - f Only available when you don't specify a table space on the CREATE TABLE.*
  - f MAXPARTITIONS 256, DSSIZE mGB and SEGSIZE 4.*
  - f You can alter MAXPARTITIONS.*
- *PBG is used for implicitly created table spaces in NFM*
  - f MAXPARTITIONS 256, DSSIZE 4GB, and SEGSIZE 4*



## Partition By Growth – Characteristics

- *CREATE TABLESPACE* allocates the first data partition
  - f* Except for *DEFINE NO*
- Add a new partition when space is needed until the *MAXPARTITIONS* limit is reached
  - f* Inherent compression dictionary if exists on the previous partition
  - f* Once a new partition is added, it will stay until the table is dropped
- *REORG TABLESPACE* will not drop partitions even if they are empty
  - f* For a table that has *LOB* columns), base rows will not be moved across partition boundary
- Utilities can operate at the partition level except for the *LOAD* utility

## *CREATE TABLE ... APPEND(YES)*

- *New APPEND option:*
  - f Maximizes performance for “INSERT at end”*
  - f Avoids overhead of attempting to preserve clustering sequence*
  - f CREATE or ALTER table*

## *Index Enhancements*

- *Support 8K, 16K, and 32K index page size*
- *Intelligent index split for sequential key inserts*
- *Option to randomized index key*
- *Index Compression*
- *Index on expression*

## *Implicit table space and database*

- *In CM will create a segmented table space with SEGSIZE 4, LOCKSIZE ROW in DSNDB04*
- *In NFM:*
  - f Partition by Growth Table Space - DSSIZE 4G, MAXPARTITIONS 256, LOCKSIZE ROW, SEGSIZE 4*
  - f Implicitly create up to 60,000 databases (DSNxxxxx), creator will be SYSIBM*
  - f ZPARM IMPSDEF = YES or NO - DEFINE YES or NO for the dataset*
  - f ZPARM IMPTSCMP = YES or NO - COMPRESS YES or NO*
- *Can not explicitly create a table space or table in an implicitly created database*

## Synergy between DB2 and SMS

- *CREATE STOGROUP* is enhanced to include SMS data class, management class, and storage class as optional parameters.

```
CREATE STOGROUP stogroup-name  
  < VOLUMES (volume-id,...) >      Now OPTIONAL  
  VCAT catalog-name  
  < DATACLAS dc-name >  
  < MGMTCLAS mc-name >  
  < STORCLAS sc-name >      } NEW Parameters
```

## *Reordered Row Format (RRF)*

- *Improved performance for tables that have variable length columns*
- *RRF is only available in V9 NFM*
- *REORG/LOAD REPLACE to convert existing tables to RRF*
  - f Can be done at PART level*
  - f Don't specify KEEPDICTIONARY*
- *Not applicable for table spaces that contain tables with variable length columns with EDITPROC or VALIDPROC*
- *Not applicable for DB2 catalog/directory tables*

## Basic Row Format and Reordered Row Format

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

### Basic Row Format

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

2-byte length

### Reordered Row Format

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

offset to C2

offset to C4

- The offset is a hexadecimal value, indicating the offset for the column from the start of the data in the row.
- Note: in BRF, for a varying length field, DB2 logs from the first changed byte to the end of the row. In RRF, the first changed byte may be the offset field for the first varying length field following the one being updated.

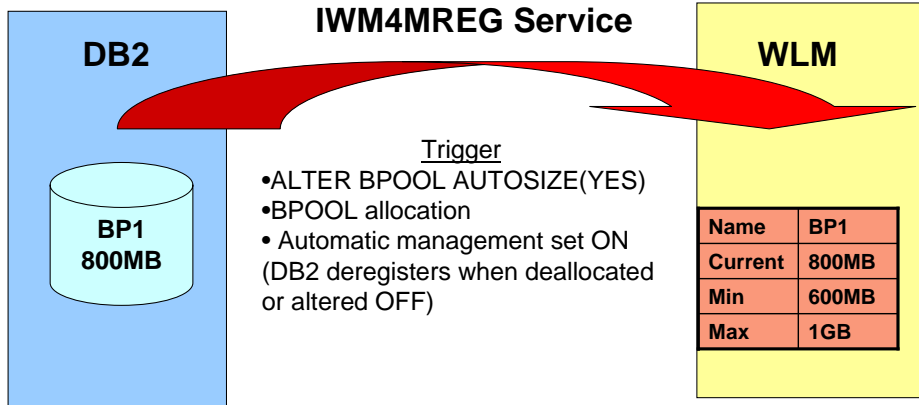
## *Automatic buffer pool management*

- *Only the size attribute of the buffer pool.*
- *Can be enabled or disabled at the individual buffer pool level.*
- *Automatic management entails the following:*
  - f DB2 Registers the BPOOL with WLM*
  - f DB2 provides sizing information to WLM*
  - f DB2 communicates to WLM each time allied agents encounter delays*
  - f DB2 periodically reports BPOOL size and random read hit ratios to WLM*

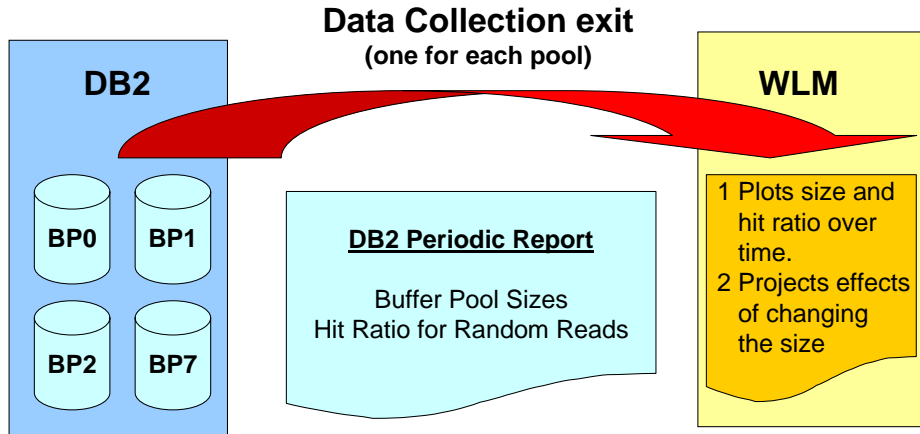
GoFurther



## DB2 Registers BPOOL to WLM



## Periodic reporting



### *Buffer Pool adjusting*

- ***If the buffer pool is adjusted, the result will be just as though an ALTER BUFFERPOOL VPSIZE command had been issued***

*f The new size is stored by DB2 in the BSDS*

- ***If the buffer pool is deallocated (e.g. because DB2 is being stopped) it will subsequently be reallocated at its most recently allocated size.***

#### ***Example***

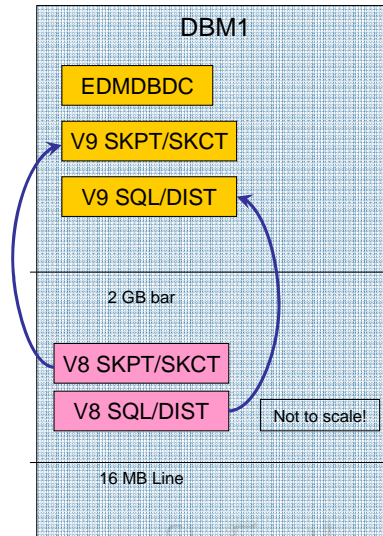
*f If BPOOL is adjusted from 800 MB to 900 MB*

*f Then DB2 is stopped and restarted*

*f BPOOL will be subsequently allocated at 900 MB*

## Virtual Storage Constraint Relief

- Virtual Storage Constraint is still an important issue for many DB2 customers. The following changes provide some relief:
- EDMPOOL Changes:
  - f V8 – DBD storage moved above 2GB bar.
  - f V9 – SKCT, SKPT storage moved above 2GB bar.
- Other changes:
  - f Some storage acquired for distributed applications moved above 2GB bar.
  - f Some storage acquired for dynamic SQL statement execution moved above 2GB bar.



## Converged WORKFILE and TEMP databases

- Use the WORKFILE database for all temporary tables
- Workfile Table Spaces are Segmented
  - f* Existing workfile table spaces will be reformatted
  - f* Default SEGSIZE = 16
  - f* Only 4K and 32K page sizes are supported
- Remove the 32K system limit for DGTTs
- Option to limit the maximum amount of TEMP storage used at the thread level
- Provide instrumentation data to monitor storage utilization of the WORKFILE database

## Data Sharing Restart Enhancements



**Improve performance,  
usability and availability  
During DB2 restart.**

### *Performance*

- f* Avoid acquiring certain locks for GBP dependent objects
- f* Open the objects involved in restart as early as possible.

### • *Usability*

- f* Automatically initiating the GBP Recovery (GRECP) of the GRECP objects at the end of restart.

### • *Availability*

- f* Supporting table level granularity of the retain locks for postponed abort Unit of Recovery

22

GoFurther

## Clone Table

- ➔ Allows fast replacing production data without renames and rebinds
- ➔ A capability to support online load replace
- ➔ Use ALTER TABLE to create a Clone Table
  - All indexes are also cloned
  - Table and Index data are not copied
  - Base and Clone tables share the same tablespace and index names
  - Underlying data sets are differentiated by a data set instance number

## Clone Table ...

- ➔ A clone table can only be created
  - On a single table Universal Table Space
  - Must be DB2 Managed
  - No RI or After Trigger on the base table
- ➔ New indexes created on the base table will set Rebuild Pending for the clone
- ➔ Use insert or load to populate clone tables
- ➔ Utilities (except RUNSTATS) can operate on clone tables with a new CLONE keyword



## Clone Table ...

- ➔ EXCHANGE DATA to switch data between the base and the clone tables
- ➔ No FASTSWICH for online REORG
- ➔ No ALTERs (i.e. online schema not allowed)
- ➔ Image copies taken on clone can continue to be used for base table recovery
- ➔ Use ALTER TABLE to drop a clone
- ➔ No clones allowed on MQT
- ➔ Catalog tables can't be cloned

## *Utility Enhancements*

- *Online Rebuild Index*
- *Online CHECK DATA and LOB*
- *Online REORG enhancements*
  - f Eliminate BUILD2 phase*
  - f Unload/Reload/Log Apply are done in Parallel*
  - f No longer support concurrent REORG PART jobs*
- *Support SHRLEVEL REFERENCE for REORG LOB*
- *CHECK INDEX Parallelism*
- *MODIFY RECOVERY enhancements*

## *Copy Pool Enhancements in Version 9*

- **BACKUP SYSTEM** and **RESTORE SYSTEM** utilities will manage Copy Pool backups to/from tapes
  - f* Up to five tape copies can be created
- Support Incremental FlashCopy
- The **RECOVER** utility can use Copy Pool backups as the source for recovery of DB2 tables and indexes
  - f* Backups can be on DASDs or Tapes
- Prereq. DFSMSHsm and z/OS 1.8

GoFurther

## *RECOVER to PIT with consistency*

*Enhance DB2 RECOVER utility to:*

- f Automatically detect the uncommitted transactions running at the recover PIT*
  - f Roll back their changes on the recovered objects.*
  - f Thus ensuring data consistency after PIT recoveries.*
  - f Recovered objects left in a transaction consistent state.*
- 
- Avoid the need to regularly run the QUIESCE utility*
    - f Reduces disruption to DB2 users and applications*

## Summary

- *Universal Table Space*
  - *f Partition by Growth and Partition by Range*
- *Index Enhancements – compression, large page size,...*
- *Automatic creation of Table Spaces and Databases*
- *Goal Oriented Buffer Pool Management*
- *Improved Row Format for Variable Length Rows*
- *Converged TEMP space*
- *More online utilities – Rebuild Index, Check Data/Lob,...*
- *Clone Table to support Online LOAD Replace*
- *Automatic GRECP Recovery for Disaster Recovery*
- *CopyPool Enhancements*

Session: C06

Session Title : *What's New in DB2 9 for z/OS for Administrators*

*Dr. Jim Teng*

*IBM Silicon Valley Laboratory*

*jteng@us.ibm.com*

