

DB2 for z/OS Universal Table Spaces: The DB2 10 Story

Willie Favero

IBM Silicon Valley Lab
Data Warehousing on System z Swat Team

Friday, March 16, 2012
11:00 AM-12:30 PM
Session Number: 10508



Please Note:



IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Acknowledgements and Disclaimers:



Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS-IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

© **Copyright IBM Corporation 2012. All rights reserved.**

- ***U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.***

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

Table Space History

In the Beginning...

- There were simple table spaces
 - Multiple tables in the same table space
 - Multiple tables could occupy the same page
 - Simplistic space map
 - 64GB size limitation
 - Deprecated in DB2 9 for z/OS
 - DB2 Catalog still uses them
 - DB2 10 discontinues use of simple table spaces for DB2 Catalog

Note

- Customer really shouldn't be creating simple table spaces anymore in any DB2 Version

And There Were Also...

- Partitioned table spaces
 - Only one table allowed per table space
 - Table space is divided into multiple partitions, data sets
 - Requires a partitioning column
 - SQL and utilities have partition independence
 - Could be up to 128Tb
 - Customer must pick number of partitions

Note

- Deprecated in DB2 10 for z/OS

Then DB2 V2.1 Added...

- Segmented table spaces
 - Multiple tables per table space
 - Pages are organized into segments
 - Only one table per segment
 - Still 64GB size limitation
 - Better space maps, better DELETES
 - And better INSERT processing
 - Customer choice, either segmented or partitioned
 - Of the three types, best performance in most cases
 - Default in DB2 9 when SEGSIZE, NUMPARTS, or MAXPARTITIONS options are not specified

Note

- No longer default in DB2 10 – partition-by-growth becomes default

Later, DB2 V6 Introduced...



- Large Object (LOB) table spaces

And Introduced with DB2 9...



- Universal Table Space (UTS)
- XML table spaces (discussion for another time)

Universal Table Space

What Was Needed



- A table space needs both partitioned and segmented organization:
 - A table's growth is unpredictable
 - How do you handle large without a convenient key for range partitioning
 - It needs to be larger than 64GB
 - Inter-partition parallelism or independent processing is necessary
 - Partition scope operations (ADD, ROTATE) apply
 - Rows are variable in length and a fast insert is required
 - Mass delete operations should be fast

What Was Needed



- Partitioning by a ROWID column introduces additional table space administration overhead:
 - estimating optimal number of partitions
 - ADDing partitions if necessary
 - less than optimal space utilization

The Solution...



- Universal table space
 - The very best of segmented and partitioned table spaces delivered in one object

What is a Universal Table Space?

- All the best features from
 - Segmented table spaces and partitioned table spaces
 - Hybrid
 - Extra space maps and space map information
 - Multiple data sets (partitions)
 - Segmentation
- Plus a bunch of really cool new stuff
 - Better space management means less REORG
 - SQL TRUNCATE supported
 - ALTER TABLE ROTATE PARTITION supported
 - CLONE table supported (UTS required)
 - Improved insert performance

Things to Remember

- Only available ~~AFTER~~ upgrading to DB2 9 new function mode (NFM)
- Only one table per table space allowed
- Reordered Row Format (RRF) only
- Partition independence
- No longer has a 64GB limitation
 - Depending on DSSIZE and the number of partitions, the table space could grow up to 128 TB
- Incompatible with MEMBER CLUSTER

Note

DB2 10 allows MEMBER CLUSTER for universal table spaces

Two Flavors are Available

- Universal table spaces are available in two flavors
 - Range-partitioned (PBR)
 - All the features of classic partitioning
 - Table controlled partitioning only
 - Using partition column
 - Partitioned and segmented
 - Partition-by-growth (PBG)
 - Partitions added as space is needed
 - No partitioning key
 - Partitioned and segmented

Common UTS Function limitations



- Cannot be used for the WORKFILE database

Note

- Version 10 supports partition-by-growth table spaces in the WORKFILE database

- No LOCKSIZE TABLE (uses partitioned table space locking scheme)

Common UTS Function limitations



- No easy way to convert current type of table space to UTS
 - Required to use DROP/CREATE
- DB2 10 ALTER TABLESPACE abc MAXPARTITIONS
 - Single table simple table space to partition-by-growth UTS
 - Segmented table space to partition-by-growth UTS
 - Plans and packages are invalidated
- DB2 10 ALTER TABLESPACE abc SEGSIZE
 - Partition table space to range-partitioned UTS

Note

Common UTS Function limitations

- No MEMBER CLUSTER (not supported for segmented)

- DB2 10 allows MEMBER CLUSTER only for partition-by-growth or range-partitioned universal table space
 - ALTER MEMBER CLUSTER
 - *Still not allowed for segmented, LOB, work file, or XML table spaces*
 - *If table space is a partitioned table space, partitioned table space is converted to range-partitioned universal table space*
 - *Altering MEMBER CLUSTER places table space in advisory REORG pending state (AREOR)*
 - SYSIBM.SYSTABLESPACE.MEMBER_CLUSTER

Note

- No ALTER SEGSIZE/DSSIZE

- So get SEGSIZE/DSSIZE right or it is a DROP/CREATE
- DB2 10 allows ALTER SEGSIZE
 - If SEGSIZE is specified on ALTER, no other clause is allowed
 - ALTER SEGSIZE can only be specified for universal table space or partitioned table space that uses table-controlled partitioning
- DB2 10 has ALTER DSSIZE also

Note

DSNZPARM for SEGSIZE Default

- When SEGSIZE is NOT specified
- DB2 10 – The picture changes considerably
 - If ZPARM DPSEGSZ = 0
 - If MAXPARTITIONS is not specified
 - *If NUMPARTS is not specified*
 - *SEGSIZE 4 for segmented table space*
 - *If NUMPARTS is specified*
 - *Classic partitioned table space*
 - If MAXPARTITIONS is specified
 - *With or without NUMPARTS being specified*
 - *partition-by-growth table space w/ SEGSIZE = 32*
 - If ZPARM DPSEGSZ > 0 (a greater than zero value)
 - If MAXPARTITIONS is not specified
 - *If NUMPARTS is not specified*
 - *SEGSIZE 4 for segmented table space*
 - *If NUMPARTS is specified*
 - *Partitioned by range-partitioned table space w/ SEGSIZE = DPSEGSZ*
 - If MAXPARTITIONS is specified
 - *With or without NUMPARTS being specified*
 - *partition-by-growth table space w/ SEGSIZE = DPSEGSZ*

DB2 9
Default SEGSIZE 4

SYSIBM.SYSCOPY



- STYPE
 - *CHAR(1) NOT NULL WITH DEFAULT*
- When ICTYPE=A (ALTER) , the new/changed values are:
 - B The MEMBER CLUSTER value was changed
 - D The DSSIZE attribute of the table space was altered
 - M The MAXPARTITIONS attribute of the table space was altered
 - S The SEGSIZE attribute of the table space was altered

Other values were added/changed in SYSCOPY but are not affected by universal table spaces

Catalog Table SYSIBM.SYSCOPY



- TTYPE
 - *CHAR(8) NOT NULL WITH DEFAULT*
- When ICTYPE=A (ALTER) and STYPE=B
 - This column indicates if the previous value for the MEMBER CLUSTER attribute is being used:
 - Y
 - The previous member cluster attribute of the table space is being used
 - N
 - The previous member cluster attribute of the table space is not being used
- When ICTYPE=A (ALTER) and STYPE=D
 - This column indicates the previous DSSIZE attribute value for the table space in units of G, M, or K

Catalog Table SYSIBM.SYSCOPY



- TTYPE
 - *CHAR(8) NOT NULL WITH DEFAULT*
- When ICTYPE=A (ALTER) and STYPE=M
 - This column indicates either the previous value of the MAXPARTITIONS attribute for the table space or the type of table space conversion that was performed on the table space
 - I
 - The table space was converted from a single-table simple table space to a partition-by-growth universal table space
 - n
 - The previous value of the MAXPARTITIONS attribute for the table space
 - S
 - The table space was converted from single-table segmented table space to a partition-by-growth universal table space

Catalog Table SYSIBM.SYSCOPY



- TTYPE
 - *CHAR(8) NOT NULL WITH DEFAULT*
- When ICTYPE=A (ALTER) and STYPE=S
 - This column indicates either the previous value of the SEGSIZE attribute for the table space or the type of table space conversion that was performed on the table space
 - n
 - The previous value of the SEGSIZE attribute for the table space
 - P
 - The table space was converted from a partitioned table space to a range-partitioned universal table space

Pause for *Questions*

Partition-by-growth Universal Table Space

Partition-by-growth Table Space

explicit specification

```
CREATE TABLESPACE ...  
MAXPARTITIONS integer
```

implicit specification


```
CREATE TABLE ...  
PARTITIONED BY SIZE EVERY  
integer G
```

- ✓ Associated SYSTABLESPACES columns
 - MAXPARTITIONS =max number of partitions
 - PARTITIONS =actual number of partitions
 - TYPE =G
- ✓ Only single-table table space
- ✓ Universal table space organization: although the table space is partitioned, the data within each partition is organized according to segmented architecture
- ✓ Incompatible with MEMBER CLUSTER, ADD PARTITION, ROTATE PARTITION

How Partition-By-Growth Works



- ✓ The table space starts with one partition, additional partitions will be added on demand until the maximum partition is reached.



Partitioned Table Space (parts added on demand)

Partition-By-Growth CREATE

- SQL CREATE TABLESPACE statement for PBG
 - CREATE TABLESPACE TS1 IN DB1
 - **MAXPARTITIONS 55** ← **Makes PBG**
 - **SEGSIZE 64**
 - **DSSIZE 2G** ← **Partition size**
 - LOCKSIZE ANY;
- A new key word MAXPARTITIONS - specifies the maximum # of partition for a table space.
- Maxpartitions can be changed by ALTER TABLESPACE
 - Keep in mind that ALTER MAXPARTITIONS may require down time because it needs to physically close the datasets

Partition-By-Growth Create

- SQL CREATE TABLE statement for PBG

- CREATE TABLE Mytable
- **PARTITION BY SIZE EVERY integer G;**
where integer ≤ 64

- Only available when you don't specify a table space name on the CREATE TABLE
- Table space is implicitly created
- mG specifies DSSIZE of the table space

More on Implicitly Created PBG



- Implicitly created table space defaults to PBG
- It defaults to row locking
- The LOCKMAX defaults to SYSTEM
- Default value for MAXPARTITIONS = 256
- Default SEGSIZE = 4 if not specified on DDL

Note


- In DB2 Version 10, the default SEGSIZE value for universal table spaces has changed from 4 to 32
 - New DSNZPARAM – DPSEGSZ (default 32) on DSN6SYSP macro
 - DPSEGSZ affects the SEGSIZE default chosen
 - DPSEGSZ becomes available in DB2 10 new function mode (NFM)
- Default DSSIZE = 4G if not specified on DDL
 - Note: DSSIZE and SEGSIZE require a DROP to change, no ALTER option

Note

- DB2 10 has ALTER DSSIZE/SEGSIZE

Partition-By-Growth Space Search

- ✓ No more space in the partition...
 - Search forward to next partition if there is one
 - Search backward to previous partitions



Note: If there is any restricted DBET state of any part during the backward space search. New part will not be added.

PBG UTS and Catalog Table

- Catalog table SYSIBM.SYSTABLESPACE
- TYPE column value
 - “R” - Range-partitioned universal table space
 - “P” - Implicit table space created for pureXML columns
 - “O” - Table space is a LOB table space
 - “G” - Partitioned-by-growth table space
- MAXPARTITIONS (new column)
 - Maximum number of partitions
 - 0 (zero) if table space NOT partition-by-growth
- PARTITIONS
 - Column contains the number of physical partition (dataset) that currently exist

PBG UTS and Catalog Table



- Catalog table SYSIBM.SYSTABLEPART
 - When table space created, one partition created and one row inserted to SYSTABLEPART (assuming created with DEFINE YES)
 - Additional row added to SYSTABLEPART for each new partition required

Additional Characteristics of PBG



- PBG is partitioned according to space requirements
 - A partition is allocated when one is needed due to growth
- Each partition has a one-to-one correspondence to a VSAM data sets and **MUST** be DB2-managed
- No partitioning key to bound the data within a table space, so no PI index
- Only non-partitioned indexes can be created
 - No data-partitioned secondary index (DPSI)
- Only single table allowed per table space
 - can not totally replace segmented table space

Additional Characteristics of PBG



- When a partition fills and MAXPARTITIONS has not been reached
 - New partition created and catalog is updated
 - Even if unit of work adding a partition issues a rollback, new partition remains
 - Compression dictionary will be copied from previous partition to the new partition
 - Freespace, caching, define, logging and trackmod attributes are same for each partition
 - Drains and Claims of new partition are inherited from prior partition

Additional Characteristics of PBG



- Some DBET states are also inherited by the new partition from the previous partition
 - RO*, UTUT, UTRO for table space, PSRBD, ICOPY for NPI
- CLONE table can be created
 - Both CLONE and base table grow at the same time
- All utilities can operate at the partition level except LOAD utility

PBG – Additional Function Limitations



- No partition key range can be defined
- ~~No~~ **ALTER ADD PART**
- No ALTER ROTATE PART
- No ALTER Stogroup
- No LOAD PART
- No user-directed define partition
 - Required to use UNLOAD/LOAD instead of DSN1COPY for copying data between table space if source table space has more than 1 partition

Note

DB2 10 allows a partition to be added up to the value of MAXPARTITIONS

Practical Applications for PBG

- When no obvious partitioning column exists
- When a table requiring > 64G
 - Lift 64G size limitation of segmented table space
 - Increase overall size of table space on demand
- Space on Demand
- Large table space and manage utilities at a data subset is needed
 - Partition level utility
- There's a need for CLONE table
- Planned hash table use (added in DB2 10)

Note

Partition-By-Growth and REORG



- Reorganization of data could result in more or less partitions
 - If n # of parts to start with will be n or more # of parts at the end of the REORG
 - No delete of existing partitions
 - If REORG is at the table space level, could result in empty partitions at the end of table space
 - Tables cannot contain LOB or XML columns
- If MAXPARTITIONS is reached, REORG will fail
- If new partition is added, dictionary pages are copied from the previous partition into the new partition

Partition-by-growth and REORG



- REORG SHRLEVEL CHANGE – new partition is added to both shadow and base (I and J data set)

Partition-By-Growth and REORG PART



- Data must fit back to original part or REORG will fail
 - To prevent failure, adjust space attributes (PCTFREE or PGFREE)
- If partition range level REORG, the data on one part can overflow to the other
- To avoid failure, run REORG for entire partitioned table space

Partition-by-growth and REORG



- Table has LOB column,
 - Holes within each partition will be eliminated
 - REORG does not move the data from one partition to another

Partition-By-Growth - REORG limitations



- No parallelism to ensure data reduced to minimum number of partitions
- No REBALANCE
- No shrinking of partitions even if there are only empty partitions at the end of table space
 - The empty partitions could have header, space map page, dictionary page and system pages

Partition-By-Growth and COPY



- Copies can be made at the part level or the table space level
- Will copy empty partition
- Will also pick up new partition added during COPY for COPY SHRLEVEL CHANGE at the table space level
 - Remember this is a fuzzy copy. It is not recommend to be used for RECOVER TOCOPY

Partition-By-Growth and RECOVER



- RECOVER to currency with image copy
 - Pick up new added parts since last copy via log apply.
- RECOVER to image copy, PIT or NOT LOGGED table space
 - The excess partitions (in base, LOB or XML) will be empty (header/space map/system pages).

Partition-By-Growth and LOAD



- Only support table space level operation
 - No partition level load
- No parallelism for Load Utility
- Can accommodate growth of table space
- Copy dictionary from previous partition to the new partition
- Excess partitions remain empty
- `LOAD ... COPYDICTIONARY` not available
 - Keyword made available in DB2 9 by PK63324 and PK63325

Partition-By-Growth & Other Utilities



- CHECK INDEX SHRLEVEL CHANGE
 - Partition added during the course of the CHECK INDEX utility is NOT checked

- REBUILD INDEX SHRLEVEL CHANGE
 - Index for record inserted into new added partition during the course of the REBUILD is reflected in the index page set via log apply

Partition-By-Growth and DSN1COPY



- Partition number may be inconsistent between DSN1COPY and the target table space
 - If partition number of TARGET table space is greater than partition number of SOURCE table space:
 - Use TRUNCATE TABLE on the target table before DSN1COPY to make sure the target table is empty
 - If partition number of TARGET table space is less than partition number of SOURCE table space:
 - DSN1COPY cannot be used
 - Unload/Load may be used
- Use Numparts and Segment keywords
 - Numparts = Maxpartitions for PbG

Pause *simply for effect*

Range-partitioned Universal Table Space

Create Range-partitioned UTS

- SQL CREATE statement
 - CREATE TABLESPACE PRB_TS1 IN UTS_DB1
 - Numparts 3
 - **SEGSIZE 64** ← Makes it PBR
 - LOCKSIZE ANY;
- Create a partitioned table space and just add the SEGSIZE clause = Range-partitioned table space

Note

- Range-partitioned table space is now **DEFAULT in DB2 10**
 - Classic partitioned table spaces still supported
 - Create classic by specifying SEGSIZE 0 on CREATE

Create Table in Range-partitioned UTS



- CREATE TABLE MyTable
 - (C1 CHAR(4),
 - C2 VARCHAR(20),
 - C3 INTEGER)
 - PARTITION BY (C1)
 - (PARTITION 1 ENDING AT ('DDDD'),
 - PARTITION 2 ENDING AT ('HHHH'),
 - PARTITION 3 ENDING AT ('ZZZZ'))
 - IN UTS_DB1.PRB_TS1 ;
- Must use table-controlled partitioning

Range-partitioned UTS and Catalog Table



- Catalog table SYSIBM.SYSTABLESPACE
- TYPE column value
 - “R” - Range-partitioned universal table space
 - “P” - Implicit table space created for XML columns.
 - “O” - Table space is a LOB table space
 - “G” - Partitioned-by-growth table space

Range-partitioned Aux Table Space



- LOB table space
 - Can be user defined via SQLRULES
- XML AUX table space
 - Its table space type is also Range-partitioned UTS
 - XML rows are in the same part as base row

PBR – Additional Function Limitations



- No index-controlled partitioning definition

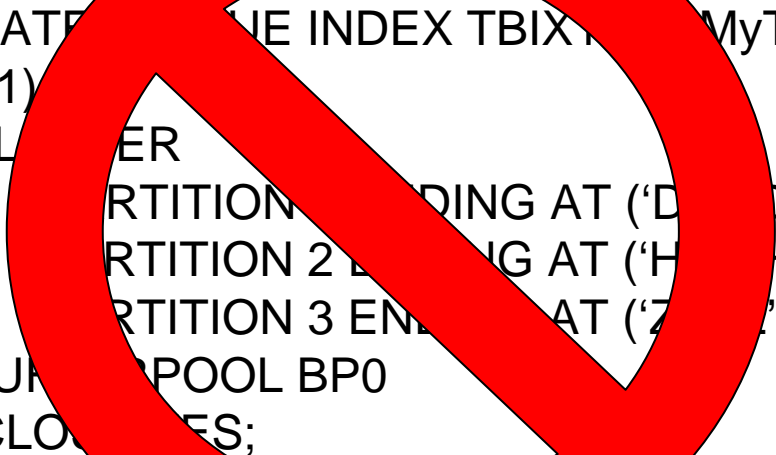
Example of invalid way to create partition range:

```
CREATE UNIQUE INDEX TBIX1 ON MyTable
(C1)
  CLUSTER
    (PARTITION 1 ENDING AT ('DDDD'),
     PARTITION 2 ENDING AT ('HHHH'),
     PARTITION 3 ENDING AT ('ZZZZ') )
  BUFFERPOOL BP0
  CLOSE YES;
```


PBR – Additional Function Limitations

- No index-controlled partitioning definition

Example of invalid way to create partition range:



```

CREATE INDEX TBIX ON MyTable
(C1)
CLUSTER
PARTITION 1 ENDING AT ('D'),
PARTITION 2 ENDING AT ('H'),
PARTITION 3 ENDING AT ('Z')
BUFFER POOL BP0
CLOSE;
  
```

SQLCODE = -662

A PARTITIONED INDEX
CANNOT BE CREATED ON A
NON-PARTITIONED,
PARTITION-BY-GROWTH, OR
RANGE-PARTITIONED
UNIVERSAL TABLE SPACE

Range-partitioned Practical Applications



- When a partitioned table space and a partitioning key is required
- When better performance than classic partitioned table space is required
- Parallelism and partition-independence capabilities
- When a CLONE table is required

Note

- Hash table use (added in DB2 10)

Conclusion

What Can You Have?

- ☑ Partition-by-growth universal table space
 - MAXPARTITIONS
- ☑ Range-partitioned universal table space
 - SEGSIZE and Numparts
- ☑ Classic partitioned table space
 - Numparts, no SEGSIZE
- ☑ Segmented table space
 - SEGSIZE, no Numparts, no MAXPARTITIONS
- ☑ Simple table space
 - Not allowed

Universal Table Spaces are Very Cool

- ☑ Only one table per table space
 - Possibly more VSAM data sets
 - DSMAX may have to be increased
 - ☑ No member cluster in DB2 9
 - Could be a concern if you are using data sharing
 - **Allowed in DB2 10**
 - ☑ No ALTER migration in DB2 9
 - Must DROP and re-CREATE
 - **ALTER in DB2 10**
- Remember...**

Pause
simply for effect

Questions

Additional Information

Shameless Self promotion

My DB2 for z/OS blog...

<http://blogs.ittoolbox.com/database/db2zos>

References

Presentations

“DB2 V9 Universal Table Spaces for z/OS”,
Frances Villafuerte,
IBM DB2 Development,
IDUG NA 2007

Article

“[Structure and Format Enhancements in DB2 9 for z/OS](#)”,
William Favero,
z/Journal Magazine (<http://www.zJournal.com>),
August/September 2007

Redbooks -- <http://www.ibm.com/redbooks>

[SG24-7330 - DB2 9 for z/OS Technical Overview](#) (8.2MB)

[SG24-7473 - DB2 9 for z/OS Performance Topics](#) (6.5 MB)

[SG24-7892 - DB2 10 for z/OS Technical Overview](#) (11 MB)

Useful DB2 for z/OS URLs



- DB2 10 Launch Website
- DB2 for z/OS Website
- DB2 Product Library
- DB2 Newsletter

<http://bit.ly/DB210Launch>

<http://www-01.ibm.com/software/data/db2/zos/>

<http://www.ibm.com/software/data/db2/zos/library.html>

<http://www.ibm.com/vrm/newsletter/11065>

- Latest Whitepapers

- [Business Value of DB2 10 – Julian Stuhler](#)
- [A Matter of Time: Temporal Data Management](#)
- [Why DB2 for z/OS is BETTER than Oracle RAC ?](#)

Top DB2 for z/OS e-Communities



- **World of DB2 for z/OS - 1700+ members**

<http://db2forzos.ning.com/>

- **DB2 10 LinkedIn - 1000+ members**

<http://linkd.in/IBMDB210>

- **DB2 for z/OS What's On LinkedIn – 2000+ members**

<http://linkd.in/kd05LH>

- **DB2 for z/OS YouTube**

<http://www.youtube.com/user/IBMDB2forzOS>

- **WW IDUG LinkedIn Group - 2000 +members**

<http://linkd.in/IDUGLinkedIn>

- **IBM DeveloperWorks**

<http://www.ibm.com/developerworks/data/community/>

**Thank You
for Attending!
Willie**

Willie Favero

DB2 SME

Data Warehousing for System z Swat Team

IBM Silicon Valley Laboratory

My DB2 Blog

www.it.toolbox.com/blogs/db2zos/

<http://www.WillieFavero.com>

