

# IBM System z Technology Summit



## DB2 10 – The Secrets of Scalability

Lee Siegmund

Consulting IT Specialist



# Agenda

- **Introduction**
- **The Need for Scalability**
- **DB2 10 Scalability Enhancements**
  - Virtual storage constraint relief
  - Latch contention reduction
  - Catalog concurrency enhancements
  - SMF compression
  - Other scalability enhancements
- **Summary & Questions**

# The Need for Scalability

- **IT volumes continue to increase**
  - More applications
  - More data
  - More transactions
- **Performance is ever more important**
  - Customers need to support workload growth without a drop-off in performance
- **Availability is ever more important**
  - Pressure to reduce both planned and unplanned outages
- **End result: each DB2 environment is being asked to work harder, with less downtime**
- **Every DB2 release attempts to push back these boundaries, but major progress has been made in DB2 10**

## DB2 10 for z/OS

- **Extensive beta program running throughout 2009/10, with customers from all around the world**
- **Generally available since October 2010**
- **Support for skip migration from V8 as well as DB2 9**
  - Will make cost case for upgrade even more compelling
- **Excellent start**
  - First customers now running DB2 10 in production
  - Compared to DB2 9 at 12 months after GA
    - 3 x number of customers running DB2 10
    - 4 x number of DB2 10 licences
    - 3 x total number of MSUs
  - Many customers are planning their DB2 10 upgrades now, with most intending to begin real work in the next 6-18 months

# Top New Features

- **CPU/Performance Improvements**
- **Virtual Storage Enhancements**
- **Security Extensions**
- **Improved Catalog Concurrency**
- **Temporal Data**
- **Access Path Management**
- **pureXML enhancements**
- **Currently Committed semantics**
- **Autonomic statistics**
- **Dynamic schema change enhancements**
- **In-memory object support**
- **Optimiser enhancements**
- **MEMBER CLUSTER for UTS**
- **Backup and recovery enhancements**
- **Enhanced audit**
- **Include additional index columns**
- **Enhanced SQL OLAP functions**
- **Skip Migration (see later)**
- **And many more....**

# DB2 10 Scalability Enhancements

## Savings ... right out of the box

IBM DB2 10 for z/OS delivers faster queries and reduced cost with optimized technology

Simple

Cuts costs

Innovative

Proven

Secure




# Overview

- **Virtual storage constraint relief**
- **Latch contention reduction**
- **Catalog concurrency enhancements**
- **SMF compression**
- **Other scalability enhancements**

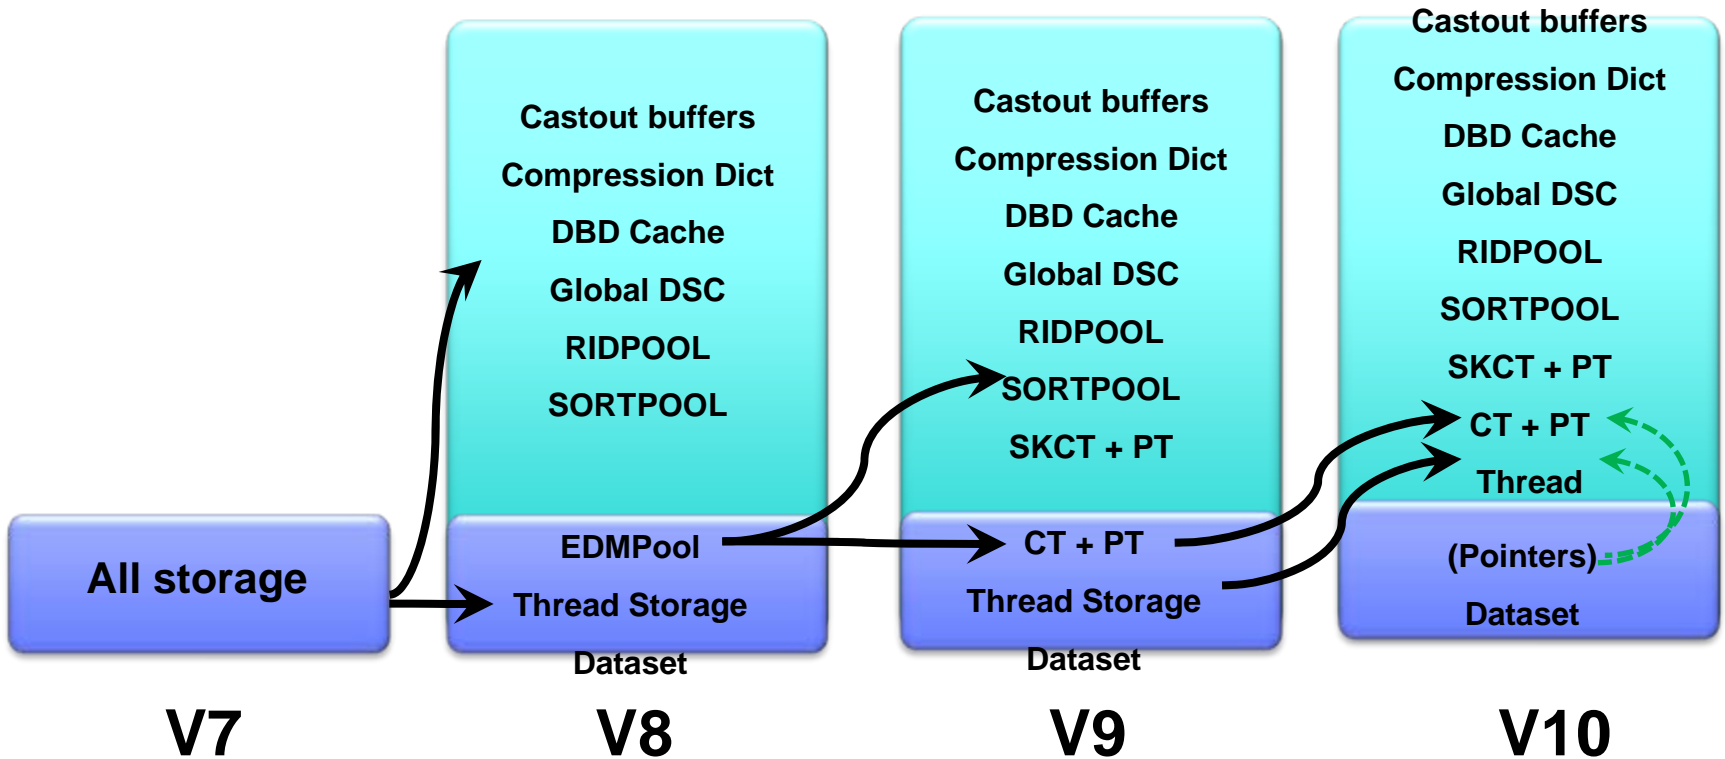
# Virtual Storage Enhancements

- **V8 began a major project to transform DB2 into a 64-bit RDBMS**
  - Laid the groundwork and provided some scalability improvements but a lot of DBM1 objects remained below the 2GB bar
- **DB2 9 improved things a little, but only by another 10-15% for most customers**
  - Practical limit of 300-500 threads per DB2 subsystem
- **DB2 10 moves 80-90% of the remaining objects above the bar, resulting in 5-10x improvement in threads per subsystem (CM)**



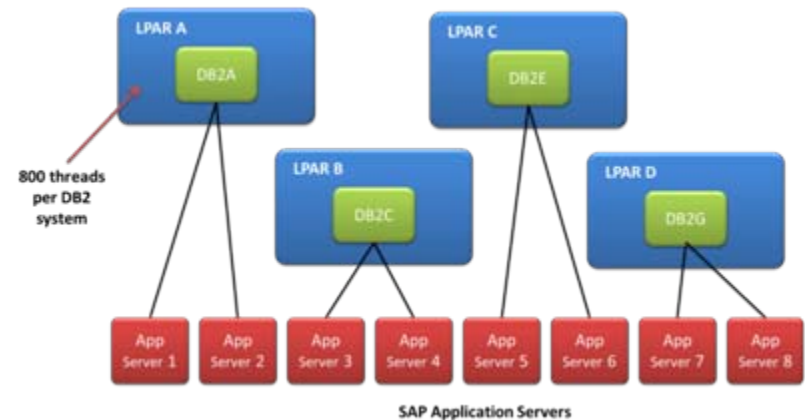



# Virtual Storage Enhancements



# Virtual Storage Enhancements

- **Possibility for less DB2 subsystems (and possibly less LPARs) in a data sharing environment**
  - Lower data sharing overhead
  - Less systems to manage / maintain
  - Minimum of 4 members still recommended for true continuous availability




# Virtual Storage Enhancements

- **More space for performance critical storage objects such as dynamic statement cache**
- **Potential to reduce legacy OLTP CPU cost through**
  - More use of CICS protected entry threads
  - More use of RELEASE(DEALLOCATE) with persistent threads (with trade-off on concurrency)
  - DB2 10 High-Performance DBATs
- **Other limiting factors on vertical scalability still remain**
  - Real storage
  - ESQA/ECSA (31-bit) storage
  - Active log write and SMF volumes


# Real Storage Enhancements

- For prior releases, z/OS always managed DB2 bufferpool pages as 4K frames
- Move to 64-bit architecture made much larger buffer pools viable
  - Bufferpools can use many millions of pages
  - Increased z/OS overheads for page management



# Real Storage Enhancements

- **DB2 10 introduces support for 1MB pages to reduce z/OS page management overheads**
  - Needs z10 or newer z196 server
  - Needs bufferpool to be defined with PGFIX=YES
  - z/OS sysprogs must partition real storage between 1K and 1MB frames (IESYSnn in PARMLIB)
- **Customer testing during beta program showed CPU reductions of 0-6% with this feature enabled**



# Storage Enhancements

- **Remains critical to ensure that there is no paging in DB2 address spaces**
  - Plan on additional 10-30% real memory following migration
- **Focus changes from virtual memory constraints & monitoring to real memory constraints & monitoring**
  - See (PM24723) for real storage monitoring and contraction enhancements – advised not to go into production without this!
- **Ensure use of PGFIX=YES to exploit 1MB real storage frames**
  - Many customers still haven't exploited this feature in their DB2 8 and DB2 9 systems – significant CPU savings!
  - Support for 1MB non page-fixed bufferpools in future release
- **Ensure you are up to date on z/OS maintenance before using 1MB pages**

# Latch contention reduction

- **Latch: DB2 mechanism for controlling concurrent events or the use of system resources**
  - Reported in accounting and statistics traces
  - Latch wait time can be significant for high-volume environments
- **DB2 10 reduces latch contention for a large number of situations, including:**
  - LC12: Global Transaction ID serialization
  - LC14: Buffer Manager serialization
  - LC19: Log write in both data sharing and non data sharing
  - LC24: EDM thread storage serialization
  - LC24: Buffer Manager serialization
  - LC27: WLM serialization latch for stored procedures and UDF
  - LC32: Storage Manager serialization
  - IRLM: IRLM hash contention
  - CML: z/OS Cross Memory Local suspend lock
  - UTSERIAL: Utility serialization lock for SYSLGRNG (removed in NFM)

# Catalog Concurrency

- **Contention on DB2 catalog is a major ongoing pain for most large DB2 customers**
- **DB2 10 introduces UTS PBG format for catalog tablespaces in NFM**
  - Internal hashes and links are removed during ENFM processing
  - Use of row-level locking and reordered row format
  - Use of new currently committed semantics and other lock avoidance techniques
  - No changes to utility jobs are necessary, but some SMS pre-reqs for migration
- **Greatly improves access to catalog/directory**
  - REORG SHRLEVEL(CHANGE) for complete catalog/directory
  - BIND concurrency much improved, but more work required in future releases – especially with heavy parallel DDL against different databases



# Catalog Contention Issues

- **Be prepared for some short-term degradation on entry to CM for single-thread BIND/REBIND processes, until you get to NFM**
  - PLANMGMT=EXTENDED the default, so multiple copies of access plan kept in catalog
  - New indexes are defined, in preparation for hash links to be removed in NFM
  - No concurrency improvement until catalog restructure in ENFM
  - Redbook testing showed worst-case elapsed time increases of 100-200% and CPU increases of 50-70%

# SMF compression

- **High transaction volume usually means high SMF volume, which can become a limiting factor**
- **Some customers forced to switch off useful accounting data, or resort to SMF rollup (via ACCUMACC ZPARM)**
- **New SMF compression feature can provide increased throughput due to I/O efficiency improvement**
  - Uses z/OS compression service to deliver approx. 60%-90% compression for 1% CPU cost

## SMF compression

- Enabled via new SMFCOMP DSNZPARM (member scope)
- All data after SMF header is compressed
- Needs vendor support to allow compressed SMF records to be processed
- New sample DSNTSMFD application to uncompress SMF data (via PM27872)
- Can be used in conjunction with accounting rollup to achieve up to 99% reduction

## Other scalability enhancements

### ▪ Others enhancements in DB2 10

- SPT01 restructured: split into several pieces with LOBs used for larger package sections
- Workfile enhancements: support for spanned records to increase maximum record length, better use of in-memory workfiles, use of PBG tablespaces
- Support for Extended Address Volumes: EAVs theoretically allow up to 221TB per volume (223GB in z/OS 1.10)
- Decrease dataset allocation/deallocation times: using new function in z/OS 1.12 DB2 startup/shutdown times can be improved (can be retrofitted to V8 and DB2 9 via APAR)

# Summary & Questions



## Summary

- **DB2 10 delivers some very significant enhancements for increasing throughput, supporting more users and reducing planned downtime**
  - Many of these enhancements available in Conversion Mode (CM)
- **Remember that sufficient real storage is needed to back any increase in virtual**
- **If you are still on DB2 V8, remember that support ends in April 2012**

## Useful URLs

- **DB2 10 Launch Website** <http://bit.ly/DB210Launch>
- **DB2 for z/OS Website** <http://www-01.ibm.com/software/data/db2/zos/>
- **DB2 Product Library** <http://www.ibm.com/software/data/db2/zos/library.html>
- **DB2 Newsletter** <http://www.ibm.com/vrm/newsletter/11065>
- **Latest Whitepapers**
  - [Gaining the Financial Benefits of DB2 10 for z/OS](#)
  - [Business Value of DB2 10 – Julian Stuhler](#)
  - [A Matter of Time: Temporal Data Management](#)
  - [Why DB2 for z/OS is BETTER than Oracle RAC ?](#)
- **DB2 for z/OS e-Kit** <http://bit.ly/DB210e-Kit>

# Top DB2 for z/OS Communities

- **World of DB2 for z/OS** <http://db2forzos.ning.com/>
- **DB2 10 LinkedIn** <http://linkd.in/IBMDB210>
- **DB2 for z/OS What's On LinkedIn** <http://linkd.in/kd05LH>
- **DB2 for z/OS YouTube** <http://www.youtube.com/user/IBMDB2forzOS>
- **WW IDUG LinkedIn Group** <http://linkd.in/IDUGLinkedIn>
- **IDUG.ORG** <http://www.idug.org>
- **IBM developerWorks** <http://www.ibm.com/developerworks/data/community/>



# DB2 for z/OS Best Practices

The screenshot shows the IBM developerWorks website interface. At the top, there's a navigation bar with 'English' and 'Sign in (or register)'. Below that, the 'developerWorks' logo is followed by links for 'Technical topics', 'Evaluation software', 'Community', and 'Events'. A search bar is also present.

The main content area is titled 'DB2 for z/OS best practices'. It features a 'Best practices' section with a graphic of a book and stars, and a brief description: 'Make your job easier, achieve your business goals more efficiently, learn techniques that work with other customers, and gain valuable knowledge directly from the experts.'

Below this, there are two columns of links:
 

- Introduction
- Planning
- Migrating
- Configuring
- Administering
- Developing
- DB2 stored procedures
- Security
- Monitoring and tuning

The 'Introduction' section contains text: 'These best practices present advice on the optimal way to use IBM® DB2® for z/OS® to satisfy key business data processing needs. These presentations and articles are authored by leading experts in IBM's development and consulting teams. Please send any feedback or suggestions for future DB2 for z/OS best practice topics to [db2zinfo@us.ibm.com](mailto:db2zinfo@us.ibm.com).'

The 'Planning' section includes a link: 'DB2 10 for z/OS: Helping you improve operational efficiencies and gain competitive advantage (IBM, April 2010)' by Roger Miller and David Beulke (David Beulke & Associates). It also mentions: 'Competitive pressures are common across most enterprise shops. These pressures give rise to challenges such as lowering operating costs through CPU cycle reductions while still building a strong foundation for SOA and XML initiatives, productivity improvements, and

On the right side, there are several widgets:
 

- 'Update My dW interests (Log in | What's this?)'
- 'Resources' with links to 'DB2 for z/OS family page' and 'DB2 for z/OS resources on developerWorks'.
- 'Special offers' with three items: 'On demand demos: An easy way to watch and learn', 'Get recognized! dW Author Program', and 'Cloud computing resources for IT professionals'.
- 'Trial software offers'.

<https://www.ibm.com/developerworks/data/bestpractices/db2zos/>

# Enjoy your complimentary DB2 10 Book The Smarter Faster Way to Upgrade !

## DB2 10 FOR z/OS

### The Smarter, Faster Way to Upgrade

In the current economic climate, businesses are under significant pressure to control costs and increase efficiency to improve their bottom line. DB2 for z/OS customers around the world are still trying to gain competitive advantage by doing more with less. More business insight, more performance, more operational efficiency, more functionality, more productivity with less cost, quicker time to market, and a lower TCO. With support for DB2 Version 11 scheduled to end in April 2012, there has never been a better time to start planning your DB2 10 upgrade. Here are the top 10 reasons to start planning today:

1. Improved performance, with reduced software license costs
2. Increased number of concurrent users, by a factor of 10
3. Reduced contention in database administration
4. More administrative capabilities with database in online
5. Improved security and auditing
6. Ability to maintain "temporality" of changing data - temporal data
7. Improved portability via enhanced SQL
8. Improved parallel performance and usability
9. Improved productivity for database systems administrators and application programmers
10. Better online transaction processing performance - Real Access


**John Campbell** is an IBMistinguished Engineer leading the team for z/OS development at the IBM Silicon Valley Lab.

**Cristian Molano** is an international IBM specialist and an IBM Red Consultant focused on DB2 for z/OS administration and performance.

**Surekha Parekh** is IBM world-wide Database Program Director - DB2 for z/OS.


IBM  
1980-2010  
© 2011 IBM Corp. All rights reserved.  
IBM, the IBM logo, and the DB2 logo are trademarks of International Business Machines Corporation in the United States, Canada, and other countries.

Price: \$14.95 (USA and CA)



## DB2 10 FOR z/OS

### The Smarter, Faster Way to Upgrade



John Campbell, Cristian Molano,  
and Surekha Parekh

The Smarter, Faster Way to Upgrade

THANK  
YOU

The image features the words "THANK YOU" in large, 3D, light blue letters. Each letter is filled with a different photograph of a diverse group of people. The 'T' shows a man in a suit and tie. The 'H' shows a woman with dark hair. The first 'A' shows a man with a beard and glasses. The 'N' shows a man with a beard. The 'K' shows a man with glasses. The 'Y' shows a man in a white lab coat. The 'O' shows a man in a white lab coat. The 'U' shows a woman with dark hair.