



High Availability in DB2 10 for z/OS

Session Number 1590

Haakon Roberts
IBM DB2 Development

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 2011. 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, ibm.com and DB2 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



Information Management Communities



- **On-line communities, User Groups, Technical Forums, Blogs, Social networks, and more**
 - Find the community that interests you...
 - **World of DB2 for z/OS** <http://db2forzos.ning.com/>
 - **Information Management** ibm.com/software/data/community
 - **Business Analytics** ibm.com/software/analytics/community
 - **International DB2 User Group** www.idug.org
- **IBM Champions**
 - Recognizing individuals who have made the most outstanding contributions to Information Management, Business Analytics, and Enterprise Content Management communities
 - ibm.com/champion





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 ?

● **DB2 for z/OS e-Kit**



<http://bit.ly/DB210e-Kit>

● **Upcoming Conferences/Events**

- IDUG DB2 Tech Conference EMEA Prague - 13th -18th Nov 2011
- DB2 10 Migration Planning Workshop Prague - 13th November 2011
- IDUG DB2 Tech Conference AG Denver - 14th -16th May 2012
- IDUG 10 Migration Planning Workshop Denver - 13th May 2012



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/>

DB2 10 is Here !

Customers seeing reduced costs, simplified workloads through proven technology



Reduced Costs	Simplified Workloads	Proven Technology
<p>“Based on the performance metrics from our controlled test environment, we see a significant amount of CPU and Elapsed time savings. This release has many features that will help bring down our operating costs.”</p>  <p>Morgan Stanley DB2 Team</p>	<p>“With DB2 10 able to handle 5-10 times as many threads as the previous version, the upgrade will immediately give the bank some much-needed room for future workload growth while simultaneously reducing their data sharing overhead.”</p>  <p>Paulo Sahadi - Senior Production Manager, Information Management Division at Banco do Brasil</p>	<p>“Every single SQL statement we have tested has been better or the same as our current optimal paths – we have yet to see any significant access path regression. We had to spend a lot of time tuning SQL with DB2 9, but we expect that to disappear when we upgrade to DB2 10.”</p> <p>Philipp Nowak, BMW DB2 Product Manager</p>
<p>“We are particularly interested in the performance improvements due to the potential CPU reductions that we realized during our DB2 10 Beta testing. Our early testing has shown out-of-the-box processing cost reductions of between 5% - 10% and for some workloads as high as 30%. Potential cost savings of this magnitude cannot be ignored given today’s business climate.”</p>  <p>Large Global Bank</p>	<p>“The new temporal functionality in DB2 10 for z/OS will allow us to drastically simplify our date-related queries. In addition, we’ll be able to reduce our storage costs by using cheaper storage for inactive rows and reduce our processing cost by having DB2 handle data movement more efficiently than the custom code we’ve written to do the same work in the past”</p> <p>Large Insurance Company</p>	<p>The new audit capabilities in DB2 10 will allow tables to be audited as soon as they are created, which is an obvious benefit for the business and will reduce costs and simplify our processes”</p> <p>Guenter Schinkel -Postbank Systems AG</p> 

For more customer references visit : <http://www.ibm.com/software/data/db2/zos/testimonials.html>

Agenda



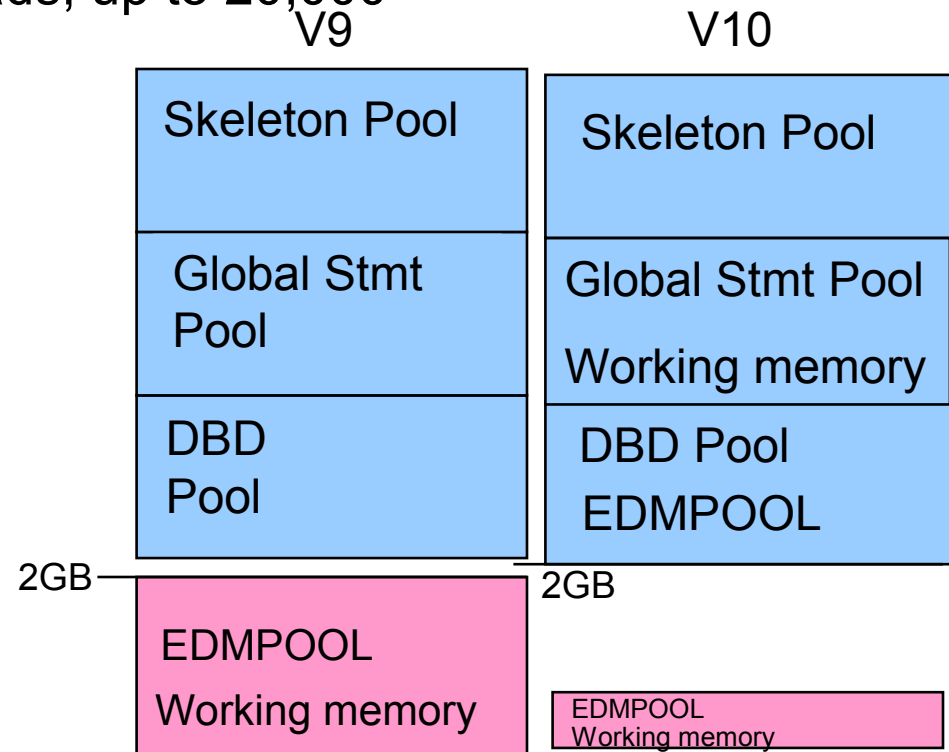
- VSCR
- Other Scalability & Contention Relief
- Logging & Restart
- Online Schema
- Utilities
- Recent Updates
- Summary



DB2 10: 64 bit Evolution Virtual Storage Relief



- DB2 9 helped (~ 10% – 15%)
- DB2 10: 5 to 10 times more threads, up to 20,000
 - Move 80% - 90% above bar
 - More concurrent work
 - Reduce need to monitor
 - Able to consolidate LPARs
 - Reduced cost
 - Easier to manage
 - Easier to grow



Scalability: Virtual storage constraint is still an important issue for many DB2 customers in 9





DB2 10: Other Scalability & Contention Improvements

- Restructure catalog & directory
 - Improved concurrency for BIND, PREPARE and DDL
 - Increased use of LOBs
 - Eliminate 64Gb limit on SPT01 and other pagesets
 - Eliminate hash chains & contention on DBD01 hash anchors
 - No need for DSN1CHKR any more
 - Enable row level locking on key catalog tables
 - Enable REORG SHRLEVEL CHANGE on entire catalog & directory
 - Eliminate UTSERIAL lock for improved utility concurrency
 - Easier management: DB2 managed & SMS controlled pagesets



DB2 10: Other Scalability & Contention Improvements



- Access currently committed data for application lock avoidance
 - New USECURRENTLYCOMMITTED BIND/PREPARE option
 - May return committed data without waiting for update locks
 - Available for UTS only
 - Supports retrieval of uncommitted deletes (pseudo-deletes)
 - Skips uncommitted inserts
 - No support for uncommitted updates
 - New instrumentation in statistics IFCID 2
 - QISTRCCI – uncommitted inserted rows skipped
 - QISTRCCD – uncommitted deleted rows read



Logging & restart



- Ability to checkpoint based on both time and number of log records
 - Meaning of CHKFREQ is unchanged
 - Minimum # of log records raised from 200 to 1000
 - New ZPARAMs to control new behavior
 - CHKLOGR – number of log records between checkpoints
 - 1000 – 99,999,999
 - CHKMINS – number of minutes between checkpoints
 - 1-1439
 - CHKTYPE SINGLE|BOTH – govern old/new
 - Set by dynamic ZPARAM or –SET LOG command
 - -SET LOG change does not persist across restart
 - -DIS LOG command indicates settings and if mode is SINGLE or BOTH



Logging & restart



- Dynamic add of active logs
 - New –SET LOG NEWLOG option
 - New active log must be IDCAMS defined & preformatted by DSNJLOGF
 - Only a single log dataset at a time
 - Issue command twice for dual logging
 - Limit is still 93 active log pairs
 - No dynamic delete of active logs
- Pre-emptable backout
 - Pre-V10, abort/backout schedules non-preemptable SRB
 - On single CPU system may give impression of DB2 hang
 - V10: Create enclave at restart for preemptable SRB backout processing



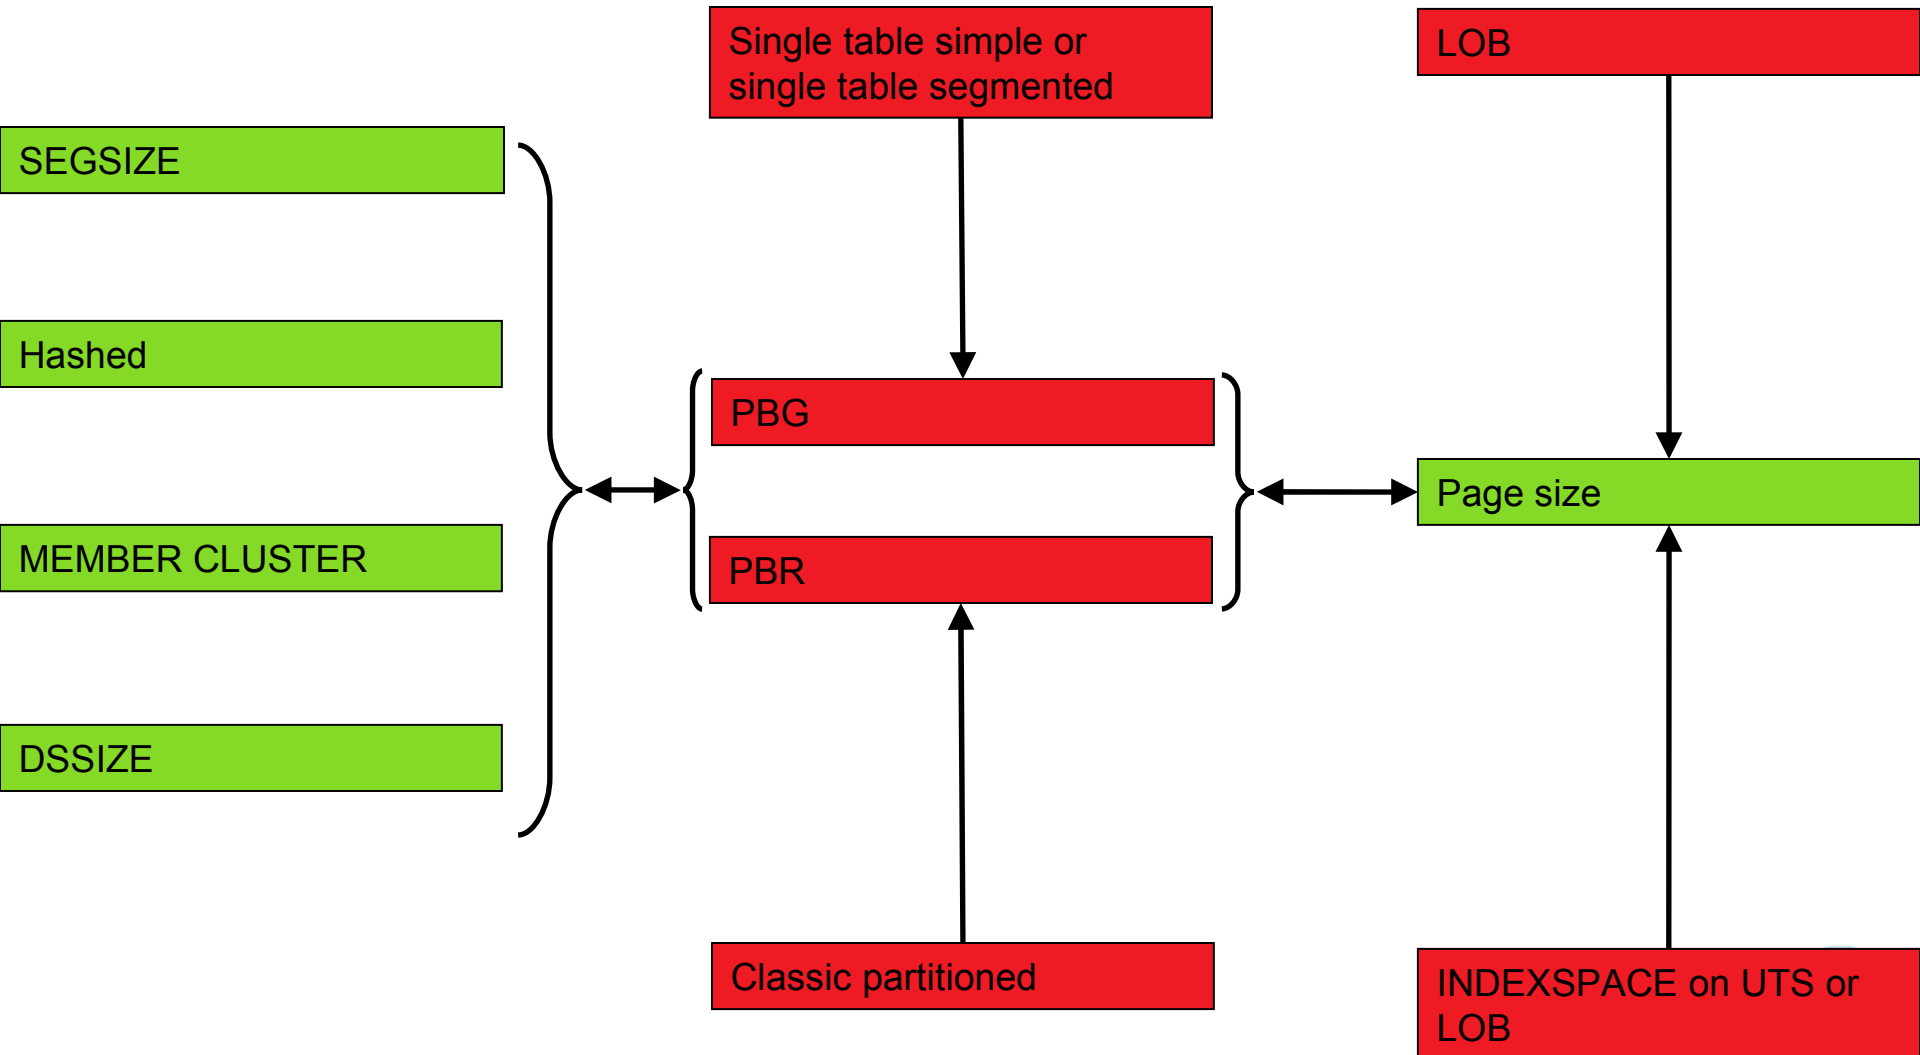
Online Schema



- V9:
 - Change of table - or index space attributes can require an outage
 - Unload data
 - Drop table space
 - Recreate table space, tables, indexes, views
 - Re-establish authorization & RI
 - Reload data
 - Undo of changes requires same process
- V10:
 - Execute ALTER statement
 - Affected pagesets placed in AREOR state
 - Changes are cached and materialized on next table/index space-level REORG
 - Pending changes can be dropped
 - Some restrictions exist
 - Mixing immediate & deferred DDL
 - PIT recovery to prior to a materializing REORG



Online Schema – DB2 10



DB2 10: REORG



- **REORG avoidance**

- Reduced need for REORG INDEX with index leaf page list prefetch
- Reduced need for REORG with compress on insert
- Easily identify whether REORG to regain clustering is important

- **Improved availability & removed restrictions - do things you could not do before**


- Reduced application outage for REORG with inline stats
- REORG support for multiple part ranges
- REORG SHRLEVEL CHANGE for all cat/dir page sets
- REORG SHRLEVEL REFERENCE|CHANGE to remove REORP
- REORG SHRLEVEL CHANGE for LOBs
 - SHRLEVEL NONE deprecated in 10 NFM
- REORG FORCE option to cancel blocking threads
- REORG REBALANCE with LOB columns
- -ALTER LIMITKEY permitted with LOB columns



DB2 10: LOAD & UNLOAD



- Remove MAX_UTIL_PARTS zparm
 - Restriction removed for REORG in V9
- Improved performance for LOAD REPLACE with LOB data
 - Up to 50% elapsed time reduction
- Spanned record support for LOB/XML data
 - LOBs & XML documents inlined in SYSREC with base data
 - Option in addition to FRVs
 - Performance & portability




DB2 10: Backup & Recovery




- Dataset-level Flashcopy support
 - COPY, RECOVER, REORG, LOAD, REBUILD INDEX, REORG INDEX
 - New zparms & utility parms to govern
 - Virtually eliminate CPU & elapsed time for large pagesets
 - Create transaction-consistent image copies from COPY SHRLEVEL CHANGE
 - Create partition-level inline image copies from REORG

CPU time per object (z10)




Elapsed time per object (z10)



DB2 10: Backup & Recovery



- VERIFYSET option to fail PIT recovery if entire set not included
 - Base, LOB, XML & history objects
- ENFORCE NO option to avoid CHKP/ACHKP on PIT recovery of subset of set
 - Improved performance due to avoidance of set checking (RI, aux)
- Fast recovery to point in time through new BACKOUT option
 - Include indexes in RECOVER list to avoid the need to rebuild them
 - Indexes must be COPY YES
 - No imagecopy required though
 - LOBs requires APAR PM45650
 - PIT recovery always with consistency since V9



DB2 10: Other Utilities



- CHECK DATA
 - CHECK utilities will no longer set CHKP/ACHKP
 - CHECK DATA enhanced for XML support
 - Document validation
 - Schema validation
 - Automated exception table processing for XML documents
- Spatial index support for REORG, REBUILD INDEX & CHECK INDEX
 - Avoid DROP/CREATE of spatial indexes
- Removed UTSERIAL lock for greater utility concurrency



Recent Availability APARs



- New DEL_CFSTRUCTS_ON_RESTART ZPARM for auto-delete of CF structures on restart if no active connections exist
 - Useful for DR environments
 - APAR PM28925 (V10)
 - APAR PM31807 (IRLM support)
- New LBACKOUT options LIGHT and LIGHTAUTO to allow restart light with postponed abort
 - Retained locks will be kept for PA URs
 - APAR PM34474 (V9)
 - APAR PM37675 (V10)
 - Previously, LBACKOUT would be ignored and all URs would be backed out if LIGHT specified on restart
- New DDLTOX zparm to provide a timeout multiplier for DDL & DCL
 - PM32921 (V9) & PM37660 (V10)
- ALTER MAXROWS to set AREO* rather than AREOR
 - PM43597 (V10)
 - Allows MAXROWS to take effect at partition level





Recent Availability APARs

- Retrofit of support of multiple part ranges for REORG
 - More efficient, improved availability, exploit parallelism
 - PK87762 & PM13259 (V9)
 - PM25525 (V9) – new PARALLEL keyword on REORG
 - PM37293 (V9) – new REORG_LIST_PROCESSING zparm to control PARALLEL REORG processing
- Online REORG materialisation of inline LOBs
 - PM29037 (V10)
- REORG to correct BRF/RRF mismatch after DSN1COPY
 - PM40646 (V9)
- Faster handling of zero length LOBs in LOAD/UNLOAD
 - PM12286 (V9)





Recent Availability APARs

- LOAD/UNLOAD FORMAT INTERNAL
 - PM19584
 - Unload and load data in true internal format
 - 85% CPU & elapsed time reduction on UNLOAD
 - 77% elapsed time, 56% CPU reduction on LOAD
 - Supported by High Performance Unload
- LOAD PRESORTED
 - PM19584
 - Avoid sort overhead when data already sorted in clustering order
 - Up to 25% CPU reduction, 33% ET reduction depending on no. of indexes
 - Works well with Utility Enhancement Tool PRESORT option
- Fast LOAD through index avoidance
 - PM27962 (V9)
 - New INDEXDEFER option to skip index key insert
 - Leaves indexes or logical partitions in RBDP
 - For LOAD RESUME or partition-level LOAD REPLACE with NPIs
 - Can even skip unique indexes
 - LOAD single part (5% of data) with 5 NPIs: Save 64% ET



Recent Availability APARs



- LOB pageset support for RECOVER BACKOUT YES
 - PM45650 (V10)
 - Lost LOBs will be marked invalid if necessary
 - Tablespace will be placed in AUXW
- Improved MODIFY RECOVERY support for SLBs
 - PM24237 (V9)
 - MODIFY will check for existence of SLB before setting copy-pending
 - REPORT RECOVERY support for SLBs delivered in base V10
- Faster image copy to tape
 - PM23786 (V9)
 - Improved tape mark handling when copying multiple pagesets to same tape
 - One customer measured 40% elapsed time improvement
- SELECT from SYSLGRNX
 - PM35190 & PM42331 (V10)
 - ISO(UR) enforced



Recent Availability APARs



- Retrofit utility spatial index support to V9
 - PM35200 (V9)
- Prevent read-only outage from CHECK SHRLEVEL CHANGE utility when Flashcopy support not available
 - PM19034 (V9)
 - New CHECK_FASTREPLICATION zparm
 - Recommend setting to REQUIRED
- Empty lists generated by LISTDEF changed to return RC4
 - PM27099 (V10)
 - Previously returned RC8
 - RC4 desirable, and necessary since default changed to DEFINED YES
- Improve DSN1COPY usability with REPAIR VERSIONS
 - PM27940 (V9)
 - For DSN1COPY of partitioned tablespaces



Recent Availability APARs



- Allow utility restart when template uses &UNIQ or &UQ
 - PM19063 (V9)
- Extend DFSORT variable length record sort limit from 2bn to 4bn
 - PM43006 (V9)
 - 2bn limit affects single REORG data sort subtask
 - Prevent ICE121A FILE SIZE IS TOO LARGE error



Summary



- Consistent & continuing focus
- Remove scalability inhibitors
- Reduce or eliminate contention
- Eliminate planned & unplanned outages





Thank You!

Your Feedback is Important to Us

- Access your personal session survey list and complete via SmartSite
 - Your smart phone or web browser at: iodsmartsite.com
 - Any SmartSite kiosk onsite
 - Each completed session survey increases your chance to win an Apple iPod Touch with daily drawing sponsored by Alliance Tech
- Visit Us at

www.ibm.com/software/data/db2/db210

A graphic with a blue and green gradient background. On the right, an open cardboard box contains gold currency symbols (Euro, Yen, Pound, Dollar). The text 'Savings ... right out of the box' is in white. Below it, 'IBM DB2 10 for z/OS delivers faster queries and reduced cost with optimized technology' is in yellow. Other text includes 'Proven', 'Secure', 'Simple', 'Cuts costs', and 'Innovative' in various shades of blue and green.

Savings ... right out of the box

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

Proven

Secure

Simple

Cuts costs

Innovative

