



## DB2 9 for z/OS Reduce Costs & Increase Business Value

John J. Campbell - IBM Distinguished Engineer  
Kevin Campbell – Univar Application Architect

## Introduction




- **DB2 9 for z /OS is a major release of DB2 for the IBM System Z platform**
- **Many features that provide technical and business benefit across the following areas:-**
  - Price/performance (reduced cost)
  - Improved availability
  - Reduced total cost of ownership
  - Application enablement and support for new workloads

**John Campbell - IBM Distinguished Engineer**


**Email: John [Campbell1/UK/IBM@IBMGB](mailto:John_Campbell1/UK/IBM@IBMGB)**

## Agenda

### Economic downturn – changing the game



### DB2 9 for z/OS – Business Value & Benefits



- Price/performance (reduced cost)
- Improved availability
- Reduced total cost of ownership
- Application enablement and support for new workloads

### Univar – Customer Reference



### Next Steps



## Succeeding in Turbulent times Top 3 Challenges for CIOs



- Reducing Costs of Information Technology
- Boost Business Resilience and Reduce Risks
- Demand to drive more innovative solutions

## Improved Price/Performance (Reduced Cost)



- **Native SQL Procedures**

- Standard SQL type language that programmers on all platforms will be able to easily understand
  - No need for COBOL skills
  - No need for external C program which must be prepared and executed
- External SQL procedures are easily converted over
- Runs entirely inside DBM1 address space and avoids SQL API trips
- Eligible for zIIP offload when invoked by DRDA over TCP/IP connection

- **Larger index page size**

- Index must be defined in 8 KB, 16 KB, or 32 KB bufferpool
- Reduce number of index levels and reduce index retrieval costs
- Reduce number of painful index leaf page splits

## Improved Price/Performance (Reduced Cost)



- **Autonomic asymmetric leaf page split**
  - Performance relief for sequential key insert with better space utilization
  - Reduce number of painful index leaf page splits
- **Fast table APPEND ('insert at the end')**
  - Reduced space searching
- **Data sharing logging improvement**
  - Now only necessary to generate unique LRSNs when log records are for same index or data page
  - Reduced LRSN spin saves CPU and reduces log latch contention
- **Identify and remove unused indexes**
  - Difficult to determine in a dynamic SQL environment
  - SYSINDEXSPACESTATS.LASTUSED (RTS) records last used date
  - Improved insert and delete performance

## Improved Price/Performance (Reduced Cost)



- **Improved sort avoidance and performance**
  - Use of in-memory workfile if number of rows can fit into one page
  - Use of 32KB workfile if row size  $\geq$  100 bytes to reduce IO
  - New GROUP BY sort – group collapsing during sort input phase
  - Sort avoidance for DISTINCT on non-unique index
- **SELECT, INSERT, UPDATE, DELETE for LOBs**
  - Improved performance
  - Significant reduction in locking and holding locks for shorter duration
  - No lock escalation
  - New dynamic data format (progressive streaming) for JCC T4 applications
- **Reduced CPU for LOAD and REORG for charge back**
  - Improvements related to reduced index manager costs, use of shared memory objects to avoid data movement, improved index key generation

## Improved Availability



- **Online REORG**

- Eliminates the BUILD2 phase for REORG PART operation
- NPIs also shadowed and implicitly reorganised
- Partition level unload/reload/log apply parallelism
- Removes prime cause of outage

- **Online REBUILD INDEX**

- Good for CREATE INDEX DEFER YES

- **CLONE Table**

- Fast replacement of one table with another (flip-flop)
- Addresses requirement to replace the entire contents of a table while maintaining access to the old data until the new dataset has been loaded
  - Aka 'Online LOAD REPLACE SHRLEVEL(CHANGE)'
- Reduce or even eliminate service outage caused by batch processes



## Improved Availability ...



- **REORG of LOB table space**
  - Complete REORG of LOB data to reclaim space
- **Modify EARLY code with no IPL needed**
  - New command to refresh early code and then recycle DB2
- **Consistent RECOVER**
  - Automatically detects uncommitted transactions that are running at the PIT recovery point
  - Rolls back changes on the object to be recovered to ensure data consistency after the PIT recovery
  - URs that are INFLIGHT, INABORT, POSTPONED ABORT are rolled back
  - Leaves the recovered objects in a consistent state from a transaction point of view
  - Reduces even eliminates the need for taking successful QUIESCE points

## Improved Availability ...



- **Data sharing restart availability enhancements**
  - Initiating automatic GRECP recovery at the end of restart
  - Deferring the updates of SYSLGRNX beyond end of restart
  - Opening data sets earlier in restart processing
  - Removing need for conversion locks during special open
  - Allowing table-level retained locks to support postponed abort unit of recovery
- **Cancel in progress database commands**
- **Online schema change**
  - RENAME COLUMN and RENAME INDEX
  - Eliminate destructive changes

## Reduced Total Cost of Ownership



- **Plan stability**
  - Ability to backup your static SQL packages
  - Save old copies of packages in Catalog/Directory
  - Can switch back to previous or original version when bad access path change
  - Removes the fear of REBIND
- **Histogram statistics**
  - Represents pockets of data
  - Improved filter factor estimation when gaps in the range
- **Universal Table Space Partition-by-Growth (PBG)**
  - New partition added automatically when more space needed
  - Max size controlled by MAXPARTITIONS, DSSIZE, and page size
  - Help deal with potentially large unpredictable data volumes

## Reduced Total Cost of Ownership ...



DB2 9 for z/OS  
is here

- **Trusted network context and SQL ROLE**
  - Addresses security/audit issue in 3-tier architectures where a 'surrogate' user id (or function id) is used to access DB2
  - Provides better control access to applications
  - Provides better audit ability both when making database change and when a user executes 'transactions'
- **Selective tracing**
  - New trace filters available to help minimize trace overhead
  - Filters include the ability to include or exclude data with wild card capability
  - Use filters to target detailed trace classes selectively and reduce CPU overhead

## Reduced Total Cost of Ownership ...



- **Incremental DBM1 31-bit storage VSCR (5-10%)**
  - Reduced EDM Pool requirement for static SQL
  - Reduced Local Dynamic Statement Cache when using KEEP DYNAMIC(YES)
- **Index compression for informational systems**
  - Save DASD space
  - Requires large index page size
  - Target large indexes e.g., NPIs
- **Utility TEMPLATE switching**
  - Extends the capability of the template command to allow different output locations to be specified based on the size of the dataset
  - Reduces the ongoing effort required to monitor and maintain backup jobs by automatically selecting the correct output location as DB2 tables grow over time
- **MODIFY RECOVERY**
  - Simplification and safety

## Reduced Total Cost of Ownership ...



DB2 9 for z/OS  
is here

- **Exploitation of volume level backups**
  - Tape support/control for BACKUP and RESTORE SYSTEM utilities
  - Recovery of individual tables spaces and indexes from volume-level backups
  - Exploitation of Incremental FlashCopy
- **Automatic object creation**
  - Implicit creation of
    - Database
    - Primary key index
    - Unique key index
    - ROWID index
    - LOB table space, table & auxiliary index

## Application enablement and support for new workloads

- **Integrated XML support**
  - Declarative language, reduce complexity, dramatically improve application development productivity
  - Directly store and query XML in inherent hierarchical format
    - No decomposition/composition
    - No normalize/de-normalize
  - Native processing with good XML index design = high performance
  - Ideally suited
    - Versatile schemas that are diverse and evolve, and end-user customizable applications
    - Sparsely populated attribute values (null vs. absence)
  - Manage XML data with ACID properties, auditing and regulatory compliance, together with relational data

## Application enablement and support for new workloads

- **INSTEAD OF triggers**

- Usability feature provides an extension to the updatability of views
- Trigger logic performs the operation against the table on behalf of the view
- Transparent to the application

- **Index on expression**

- General application for multi key column browsing
- Eliminate non-matching index scans
- Remove column concatenation
- Reduce number of destructive index changes



## Application enablement and support for new workloads

- **FETCH FIRST and ORDER BY in subselect**
  - Can perform mass insert/update/delete in increments
- **Optimistic locking control**
  - Positioned updates and deletes performed with optimistic concurrency control method
  - Uses RID and a row change token to test whether data has been changed by another application since the last read operation
  - Ensures data integrity while limiting the time that locks are held
  - Faster and more scalable than database locking for concurrent data access

## Univar Background



- **Largest National distributor of Industrial Chemical and Food products and related services**
- **Data center in Redmond, WA services 4,500 US and 755 Canadian employees**
- **Core business system is CICS/DB2 procurement, inventory management and sales order processing ERP**
- **Approximately 14 DB2 subsystems on 3 LPARs**

## Introduction

**DB2 9 for z/OS  
is here**



**Kevin Campbell**

**Manager, Application and Data  
Architecture – Univar**

**Recently migrated to DB2 9 for z/OS**



## DB2 9 Migration



- **ESP Participant, tested migration several times during beta**
- **No significant application issues during migration, DB2 connect upgraded in some cases**
- **All sub-systems migrated within a 5 month timeframe**

## DB2 9 Feature Exploitation



- **Partition By Growth table spaces**
  - Very useful for packaged applications where key structure is unknown and no obvious candidate for range partitioning
  - Univar has used this feature extensively with DB2 Content Manager

## DB2 9 Feature Exploitation



DB2 9 for z/OS  
is here

- **Native SQL Procedures**
  - Easier deployment
  - Improved performance
  - Great development/debug support with Data Studio Developer, very accessible to more workstation oriented developers
  - Better portability, richer language support
  - zIIP!

## DB2 9 Feature Exploitation



DB2 9 for z/OS  
is here

- **OmniFind Text Search Server**
  - No charge feature ordered as part of accessories suite
  - Used to provide sophisticated search function for users with minimal z overhead - index processing occurs on Intel platform
  - Univar creates indexes based on user defined functions which concatenate several character fields on customer and product records

## DB2 9 Feature Exploitation



DB2 9 for z/OS  
is here

- **Pure XML**
  - Developed message warehousing application to store XML documents generated by internal and external integration
  - Greatly helps auditing and issue resolution with trading partners
  - Single XML column used to store multiple document types



## DB2 9 Feature Exploitation



DB2 9 for z/OS  
is here

- **“Instead of” triggers**
  - Often use views to mask underlying data model
  - Instead of triggers allow these to be used for update and insert
  - Maintain tight control of data integrity while simplifying access for developers

## DB2 9 Feature Exploitation



- **Index on expression**
  - Create an index on the result of a scalar function:  
UPPER, LEFT, MONTH
  - Avoids the need to have a separate, indexable, column representing the function result
  - We've used this to create an index on the UPPER() value of mixed case fields

## DB2 9 Feature Exploitation



DB2 9 for z/OS  
is here

- **BI Improvements**

- Migrated Data Warehouse from an Other product
- Small (10 MSU) LPAR able to meet or beat internal query benchmarks
- Saw improvements when enabling index ANDing for star schema
- Histogram statistics help with skewed data distribution
- Housekeeping/utilities much, much faster than Other product

## Next Steps & More Information !



- **Are you ready to Migrate to DB2 9 for z/OS ?**

Contact your local IBM representative or email

WW DB2 for z/OS Market Manager [Surekha21@uk.ibm.com](mailto:Surekha21@uk.ibm.com)

- **Need More Information**

- [DB2 for z/OS Landing page](#)

- **Whitepaper**

[DB2 9 for z/OS Data On Demand](#)

- **IBM Redbooks**

[Latest Redbooks](#)



- **Univar Case Study**

[Univar uses DB2 9 for z/OS with pureXML to speed development and reduce cost](#)

- **Join [“The World of DB2 for z/OS !”](#)**

- **Twitter DB2: <http://twitter.com/IBMDB2>**

