Hints and Tips to get the most out of DB2 9 for z/OS

John J. Campbell - IBM Distinguished Engineer

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

Agenda

Economic downturn – changing the game

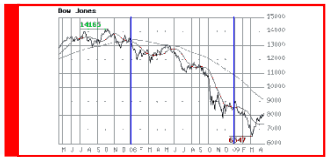

DB2 9 for z/OS – Business Value & Benefits

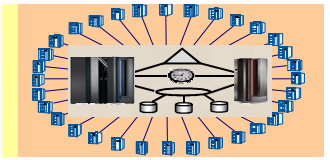

- Price/performance (reduced cost)
- Improved availability
- Reduced total cost of ownership
- Application enablement and support for new workloads

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)

- **Larger index page size**
  - Available in V9 NFM
  - Can be used with or without index compression
  - Index must be defined in 8 KB, 16 KB, or 32 KB bufferpool
  - Potential to reduce the number of index leaf page splits, which are painful especially for GBP-dependent index (data sharing)
    - Reduce index tree lotch contention
    - Reduce index tree p-lock contention
    - Reduce active log write I/O
  - Potential to reduce the number of index levels
    - Reduce the number of getpages for index traversal
    - Reduce CPU resource consumption
  - Possibility that large index page size may aggravate index bufferpool hit ratio for random access

Large Index Page Size Examples

Rows In Table	1,000,000,000								
Key Length	**4**	**8**	**16**	**32**	**64**	**128**	**256**	**512**	**1024**
Page Size									
4096									
Entries/Leaf	336	252	168	100	56	29	15	7	3
Leafs	2,976,191	3,968,254	5,952,381	10,000,000	17,857,143	34,482,759	66,666,667	142,857,143	333,333,334
Non-Leaf fanout	331	242	158	93	51	26	13	7	3
Index Levels	**4**	**4**	**5**	**5**	**6**	**7**	**9**	**11**	**19**
8192									
Entries/Leaf	677	508	338	203	112	59	30	15	7
Leafs	1,477,105	1,968,504	2,958,580	4,926,109	8,928,572	16,949,153	33,333,334	66,666,667	142,857,143
Non-Leaf fanout	666	488	318	187	103	54	27	14	7
Index Levels	**4**	**4**	**4**	**4**	**5**	**6**	**7**	**8**	**11**
16,384									
Entries/Leaf	1360	1020	680	408	226	120	61	31	15
Leafs	735,295	980,393	1,470,589	2,450,981	4,424,779	8,333,334	16,393,443	32,258,065	66,666,667
Non-Leaf fanout	1,336	980	639	376	207	108	55	28	14
Index Levels	**3**	**4**	**4**	**4**	**4**	**5**	**6**	**7**	**8**
32,768									
Entries/Leaf	2725	2044	1362	817	454	240	123	62	31
Leafs	366,973	489,237	734,215	1,223,991	2,202,644	4,166,667	8,130,082	16,129,033	32,258,065
Non-Leaf fanout	2,676	1,963	1,280	755	414	218	111	56	28
Index Levels	**3**	**3**	**3**	**4**	**4**	**4**	**5**	**6**	**7**

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
  - Potential for significant reduction in CPU resource consumption by avoiding
    - Overhead of stored procedure invocation overhead
    - Overhead of roundtrip between WLM and DBM1 address spaces for each SQL call
  - Short running SQL procedure could achieve up to an 40% ITR improvement
  - But little or no improvement for long-running SQL procedure
  - When invoked from DRDA connection over TCP/IP
    - zIIP eligible
    - As it runs in DBM1 address space under DDF enclave SRB
  - Easy to code, develop and manage
  - Selective application re-engineering may be required when migrating existing SQL procedures to Native SQL Procedures

Improved Price/Performance (Reduced Cost)

- **Stored Procedures - Performance of different languages**
  - Environment Configuration
    - z/OS 1.9
    - DB2 9 for z/OS
    - Universal Driver 3.52.76
    - JDK 1.4.2 (SQLJ/JDBC stored procedures)
    - 3 CP's
    - 2 zIIP's
    - 2 zAAP's
    - IRWW OLTP workload

More Quick Hits …

- **Stored Procedures - Performance of different languages …**

Language/API	Base CPU/tran Cost	Billable CPU/tran Cost after zIIP and/or zAAP redirect
COBOL stored proc	1X (BASE)	.88x
C stored proc	.95x	.83x
SQLJ stored proc	1.7x	1.15x (zIIP + zAAP)
JDBC stored proc	2.95x	1.76x (zIIP + zAAP)
External SQL stored proc	1.62x	1.49x
Native SQL stored proc	1.14x	.65x
Remote SQLJ	1.78x	1.06x

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
  - Use with MEMBER CLUSTER if GBP-dependent
- **Data sharing logging improvement**
  - Now only necessary to generate unique LRSNs when log records are for same index or data page
  - Reduced LRSN spin when holding log latch saves CPU
- **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 > = 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 KEEPDYNAMIC(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**
  - Simplication and safety

Reduced Total Cost of Ownership …

- **Exploitation of volume level backups**
  - Tape support/control for BACKUP and RESTORE SYSTEM utilities
  - Recovery of individual tables paces 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

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

- **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**