

TSB-3274A

DB2 10 for SAP Benefits for SAP Customers

Johannes Schuetzner

SAP on DB2 for z/OS Development

IBM Boeblingen Lab

Housekeeping

- We value your feedback - don't forget to complete your evaluation for each session you attend and hand it to the room monitors at the end of each session
- Overall Conference Evaluation will be provided at the General Session on Friday
- Visit the Expo Solutions Centre
- Please remember this is a 'non-smoking' venue!
- Please switch off your mobile phones
- Please remember to wear your badge at all times

IBM Disclaimer

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.

Agenda

- Overview: DB2 10
- DB2 10 features for SAP – Scalability
- DB2 10 features for SAP – Performance
- DB2 10 features for SAP – Availability
- DB2 10 features for SAP – Autonomics and DBA Productivity
- DB2 10 features for SAP – Portability and Security



DB2 10 Beta Program

- Announced on February 9
- Beta program started on March 12
- Check www.ibm.com/software/data/db2/zos/db2-10
- Prereqs z/OS 1.10

Agenda

- Overview: DB2 10
- **DB2 10 features for SAP – Scalability**
- DB2 10 features for SAP – Performance
- DB2 10 features for SAP – Availability
- DB2 10 features for SAP – Autonomics and DBA Productivity
- DB2 10 features for SAP – Portability and Security



DB2 10 for SAP – Scalability

Overview and Values

- **Huge scalability improvement for concurrent workload within single DB2 subsystem**
 - Thousands of SAP work processes served by single DB2 subsystem
 - Better workload consolidation
 - More options for SAP landscape configurations
- **Partition-by-growth tablespaces used for DB2 Catalog tables**
 - Simplified setup and operations

Full 64-bit Runtime – 10x more DB2 threads

Challenge

Virtual storage for different DB2 storage areas in DBM1 address space restricted to 2 GB
→ Restricts maximum number of SAP work processes served by single DB2 subsystem

Solution

64-bit support for all relevant storage areas in DBM1

- Particularly, DB2 threads live above the 2 GB bar
- Goal: 5 to 10 times more active threads per DB2 subsystem
- Stack storage space also 64-bit enabled

Reducing DB2 Catalog Contention

Challenge

- Some DB2 catalog table spaces use so-called links, which requires page-level locking
→ Favours deadlock and timeout issues with concurrent DDL, DML, utilities
(for example, SAP BW)
- DB2 catalog and directory is user-managed → Additional complexity

Solution

The DB2 catalog and directory table spaces that used to have links are converted to regular UTS PBG table spaces

- One UTS PBG table space per table
 - DSSIZE 64 GB, RRF, SEGSIZE 32, MAXPARTITION 1, row-level locking
- Links are not used anymore
- Row-level locking can be used across the board for all catalog table spaces
- Also, DSNDB01.SYSLGRNX switches to row-level locking

Reducing DB2 Catalog Contention (continued)

Solution Details

- DB2 catalog and directory become DB2-managed
- Catalog and directory data sets are SMS-managed
- Existing data sets are converted during ENFM processing
- Reduces user efforts to maintain data sets
 - No need to allocated data sets and extends
 - No need to allocated data sets for shadows for online REREG
 - No need to estimate sizes for data sets (some customers used huge sizes)
 - SMS to determine which data set goes to which volume
- Also, new CLOB columns to store SQL statement texts
- Also, new BLOB columns to store packages and DBDs
 - Relief for previous 64 GB limit for SPT01

Agenda

- Overview: DB2 10
- DB2 10 features for SAP – Scalability
- **DB2 10 features for SAP – Performance**
- DB2 10 features for SAP – Availability
- DB2 10 features for SAP – Autonomics and DBA Productivity
- DB2 10 features for SAP – Portability and Security



DB2 10 for SAP – Performance Overview and Values

- **SAP database interface layer optimized specifically for DB2 10**
 - More streamlined and more efficient processing
 - Huge investment
- **Significant CPU and elapsed time improvements out of the box**
 - Cost savings
- **Performance boost for typical SAP Banking workloads (heavy INSERT)**
 - Elapsed time reduction for periodic banking batch jobs
- **Additional performance improvements and disk space savings for LOBs**
 - Cost savings
- **Combining member clustering and universal tablespaces**
 - Optimizing DB2 data sharing performance

Inline LOBs

Challenge

- Typically, average LOB size with SAP is between 1500 and 3000 bytes
- Performance of small-sized LOBs is slower than its compatible non-LOB data type of the same size

Solution

Inline LOBs

- An inline LOB is a LOB that resides completely or partially in the base table space of a table
 - In latter case, the rest of the LOB resides in LOB table space
- How to specify inline LOB
 - A new ZPARM is introduced: LOB_INLINE_LENGTH
 - A new clause `INLINE LENGTH` is introduced on `CREATE TABLE`, `ALTER TABLE ADD` column and `CREATE DISTINCT DATA TYPE`
 - For existing LOB, immediate change with advisory REORG-pending

Inline LOBs (continued)

Solution Details

- Inline length of a LOB can be changed
 - Can be decremented and incremented
 - Cannot become shorter than default value
 - Immediate change
 - New rows are inserted with new inline length
 - Existing rows will be changed at next REORG (advisory REORG state)
- Impact of inline LOBs on CPU consumption and elapsed time
 - If all LOB data that is processed is in inline part, big savings as no additional LOB table space needs to be accessed
- Impact of inline LOBs on space usage
 - If the base table space is compressed, inline portion of LOBs is automatically compressed as well
 - Compression dictionary applies to both non-LOB and inline LOB data
 - In LOB table space, each page can hold data of a single LOB only

Parallel Index Update at INSERT

Challenge

- For heavy INSERT workload like in SAP Banking or Retail, the defining factor for elapsed time often is the number of indices on the table
 - If table too large, unrealistic to have table and all indices resident in buffer pools
 - For each index, leaf page for new record needs to be read in from disk (sync I/O)

Solution

Relevant index leaf pages are read in from disk in parallel

- If DB2 encounters I/O wait time on GETPAGE operation for index leaf page, it inserts into next index
- After first round of inserting into each index, it checks if these succeeded for all indices
 - Very likely that the I/O operations have completed
- If not, DB2 will wait for the I/O operations for the remaining indices
- This processing applies to tables with at least three indices

Enabling MEMBER CLUSTER for UTS

Challenge

- Universal tablespace (UTS) is standard type for new SAP tables
 - For implicitly created tablespaces
 - Introduced with DB2 9
 - Better space reuse and delete processing
- UTS incompatible with MEMBER CLUSTER tablespace attribute
 - MEMBER CLUSTER provides best SQL INSERT performance in data sharing

Solution

- DB2 10 enables MEMBER CLUSTER for UTS
 - Including special INSERT processing activated by the combination of the following tablespace attributes:
 - FREEPAGE=0 AND PCTFREE = 0 and MEMBER CLUSTER
 - SAP Banking relies on this for crucial tables like BCA_PAYMITEM
- This enables exploitation of UTS across the board

Optimizer Enhancements

Challenge

Some scenarios not completely considered by optimizer

Solution

- Safe query optimization
 - If there are similar indices, index with more matching columns is preferred
 - RID list access never fails back to table space scan
- Predicate transitive closure on IN-lists
- Matching multiple IN-LIST predicates
 - DB2 translate INLIST into in-memory table
- List prefetch with IN-LIST
 - With IN-LIST matching predicate being transformed into in-memory
- Aggressively merge query blocks
 - Consistently merge views and table expressions

Optimizer Enhancements (continued)

Solution Details

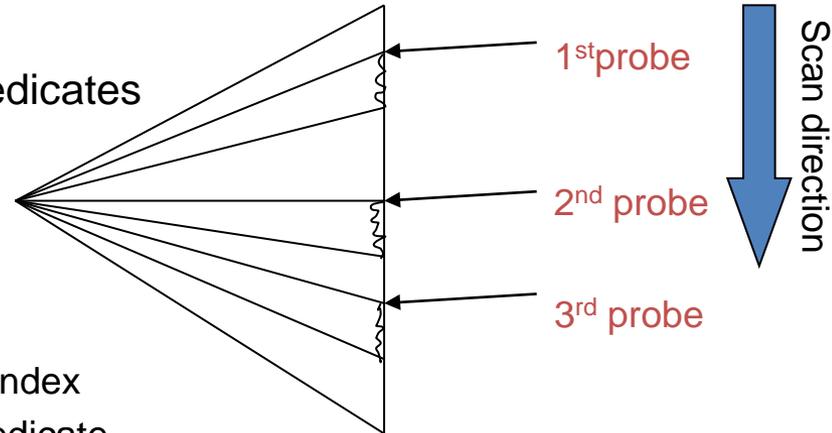
- SQL pagination

- Introduces new method to process OR predicates

- (C1 op ? AND C2 op ? AND C3 op ?) OR
 - (C1 op ? AND C2 op ?) OR
 - (C1 op ?)

- New method will be considered if:

- OR predicate is only referring single table
 - Each disjunct can be mapped to the same index
 - Each disjunct has at least one matching predicate



- Moving sum, moving average

- AVG(sh.sales) OVER (....)

- Three elements

- PARTITION BY sh.territory --- row should be assigned to partition according to territory
 - ORDER BY sh.month --- row sorted in the order of month within each partition
 - ROW 2 PRECEDING -- define the aggregation group. It's consists of the current row and the two preceding rows of the partition

Additional Improvements in Performance

- DEFINE NO for LOBs and XML
- LOB streaming between DDF and rest of DB2
 - Faster fetch and insert, lower virtual storage consumption
- Improve NOT PADDED index processing
 - Provide performance close to that of padded indexes
- Internal performance optimizations
 - Improved CPU cache performance
 - Exploit new hardware instructions
 - Streamlined DDF, RDS, DM, Index Mgr. performance-critical paths

Agenda

- Overview: DB2 10
- DB2 10 features for SAP – Scalability
- DB2 10 features for SAP – Performance
- **DB2 10 features for SAP – Availability**
- DB2 10 features for SAP – Autonomics and DBA Productivity
- DB2 10 features for SAP – Portability and Security



DB2 10 for SAP – Availability Overview and Values

- **Comprehensive FlashCopy exploitation**
 - Faster backup and recovery
- **Backward recovery**
 - Less downtime
- **Less downtime when changing table layout**
 - Better business continuity
- **Online reorganization of LOBs**
 - Better business continuity
- **More flexible online rotation of table partitions**
 - Faster reuse of existing partitions for new data like with SAP BW

More Online Schema Changes

Challenge

- During SAP lifecycle management tasks, the layout of table may change
 - Particularly during SAP upgrade
- To change different technical table settings, table needs to be converted (unloaded, dropped, recreated & loaded)
 - Table not available during this timespan

Solution

DB2 10 introduces deferred alterations for the following changes:

- Page size and buffer pool for table spaces and indices
- DSSIZE attribute for table spaces
- SEGSIZE attribute
- MEMBER CLUSTER attribute
- Convert single table segmented into UTS PBG
- Convert single table simple into UTS PBG

Deferred ALTER:

- Statement is put on to-do list in new DB2 catalog table
- Next REORG materializes pending ALTER

Online Reorganization for LOB table spaces

Challenge

DB2 9 introduced REORG support for LOB table spaces to allow reclaiming space. However, SHRLEVEL REFERENCE only → No concurrent changes allowed

Solution

Online REORG for LOBs

- SHRLEVEL CHANGE
- Builds on DB2 9 enhancement
- Base table space must be LOG YES

Agenda

- Overview: DB2 10
- DB2 10 features for SAP – Scalability
- DB2 10 features for SAP – Performance
- DB2 10 features for SAP – Availability
- **DB2 10 features for SAP – Autonomics and DBA Productivity**
- DB2 10 features for SAP – Portability and Security



DB2 10 for SAP – Autonomics and DBA Productivity

Overview and Values

- **DB2 default settings for SAP and automated setup of WLM environments**
 - Facilitated setup
- **Immediate compression**
 - Facilitated operations
- **Auto-stats**
 - Facilitated operations
- **Safe query optimization and stabilizing access paths**
 - Better out of the box performance

DB2 default system parameters for SAP

Challenge

- DB2 default values for system parameters (ZPARMs) not always adequate for SAP
- Detailed SAP recommendations for DB2 ZPARMs given that need to be implemented

Solution

- DB2 10 introduces system parameters defaults optimized for SAP systems
- Provided as new CLIST input member `DSNTIDXB`
 - Shipped as part of `SDSNSAMP`
- Besides DB2 ZPARMs, DB2 buffer pools are also configured according to SAP recommendations
 - BP attributes `VPSIZE`, `DWQT` and `VDWQT` are affected
- For DB2 subsystems migrated to DB2 10:
 - Existing bufferpool settings are honored
 - Existing DB2 ZPARM values are honored except if existing value is below new minimum
 - New DB2 ZPARM values are picked up from `DSNTIDXB`

Automated installation of WLM application environments

Challenge

- DB2 ships a growing number of stored procedures for different purposes like ease of administration
- SAP ships SAPCL stored procedure to collect DB2 performance data from IFI interface
- However, a prereq to run stored procedures is that suitable application environment exist in z/OS WLM
 - This is a manual step outside the control of DB2 (no batch or API interface)

Solution

- DB2 10 establishes WLM application environments for DB2 stored procedures as part of installation or DB2 release migration
- New single DB2 job to create all DB2-supplied stored procedures
- New validation program for DB2-supplied stored procedures

Compress on the fly during INSERT

Challenge

- Tables need to be reorganized in order to be compressed → Operational complexity
- For example, after SAP install or BW data load

Solution

Tables with COMPRESS YES attribute are directly compressed by DB2

- INSERT statement is instrumented to trigger creation of compression dictionary
 - Also applies to MERGE and LOAD SHRLEVEL CHANGE
- Conditions:
 - Table space defined with COMPRESS YES
 - No compression dictionary available
 - Table space contains enough data
- Compression dictionary is built asynchronously

Extended FlashCopy Exploitation

Challenge

General direction with SAP is to rely on FlashCopy technology for DB2 backups

- With DB2 9, some non-FlashCopy image copies still need to be taken

Solution

- DB2 10 `REPORT RECOVERY` includes `BACKUP SYSTEM` backups
- DB2 10 automatically deletes CF structures during first DB2 restart after recovery
- Data-set level FlashCopy for DB2 utilities
 - `COPY`
 - `REORG inline copy`
 - `LOAD inline copy`
 - Copy of indexes in `LOAD`, `REORG`, `REORG INDEX`, `REBUILD INDEX`
- FlashCopy backups can be used as input to:
 - `RECOVER`
 - `COPYTOCOPY`
 - Create sequential copies from FlashCopy
 - To exploit data set-level FlashCopy for `UNLOAD`:
 - `COPY` → `COPYTOCOPY` → `UNLOAD`

Backward Recovery

Challenge

Single SAP table needs to be recovered to a prior point in time and last backup is fairly old

- Single table recovery is exception in SAP context but may happen
- Recovery can be slow since DB2 always uses last backup as base and then applies log records

Solution

- Allow PIT recovery by rolling back on the log rather than restoring recovery base and rolling forward
 - New option for RECOVER utility: BACKOUT

Reduce Need for Reorganization of Indices

Challenge

If an index is disorganized:

- Sequential prefetch mechanisms may work poorly due to more gaps between leaf pages
→ More sync I/O
- Tends to have more index tree levels → More efforts to traverse the tree

Solution

- DB2 10 can do list prefetch for index leaf pages during query processing based on index non-leaf page information
 - Decision for list prefetch is made at execution time of query
- Reduces CPU time spent on index root page GETPAGE requests
- Rules adapted in stored procedure `DSNACCOX` for `REORG INDEX` recommendations

Dynamic SQL statements: More granular timers

Challenge

- For some wait times of SQL statements, DB2 does not provide dedicated timers
- Difficult to find all relevant statements on a table or for SAP transaction, end user, or report

Solution

- New IFCID 316 fields added:
 - QW0316W8 - "Wait for page latch"
 - QW0316W9 - "Wait for Drain Lock"
 - QW0316WA - "Wait for Drain Lock for claims to be released"
 - QW0316WB - "Wait for log writer"
- New filters for IFCID 316:
 - Filter for first table in a statement: WQALRTRL
 - Filter for client end user ID field of a statement: WQALEUID
 - Filter for client transaction name field of a statement: WQALEUTX
 - Filter for client workstation name field of a statement: WQALEUWN

Auto-Stats

Challenge

Customers need to ensure that the right statistics for the DB2 optimizer are collected at the right time

Solution

The new Auto-Stats feature provides the following:

- Profile support for RUNSTATS utility:
 - `RUNSTATS TABLESPACE.... SET / USE / UPDATE / DELETE PROFILE`
- New stored procedures
 - `ADMIN_UTL_MONITOR` checks need for new/better statistics based on RTS
 - `ADMIN_UTL_EXECUTE` collects new statistics using `ADMIN_UTIL_SCHEDULE`
- Control tables:
 - `SYSIBM.SYSTABLES_PROFILES` contains the RUNSTATS options at table level
 - `SYSIBM.SYSAUTOTIMEWINDOWS` defines when autonomic procedures can be run
 - `SYSIBM.SYSAUTOALERTS` is populated when `ADMIN_UTL_MONITOR` detects that an action needs to be scheduled for execution
- DB2 Admin Scheduler can be used to automatically run the stored procedures

Agenda

- Overview: DB2 10
- DB2 10 features for SAP – Scalability
- DB2 10 features for SAP – Performance
- DB2 10 features for SAP – Availability
- DB2 10 features for SAP – Autonomics and DBA Productivity
- **DB2 10 features for SAP – Portability and Security**



DB2 10 for SAP – Portability and Security Overview and Values

- **Completely overhauled new DBA privileges**
 - Based on business security policies, specific authorities can be granted
 - Different authorities dedicated to certain roles
 - Better complying with security regulations
 - Even SAP install does not require SYSADM privilege anymore

DB2 10 for SAP – Portability and Security

Overview and Values

- Specifying physical tablespace attributes at CREATE TABLE
 - Prereq for completely streamlined new SAP DBSL optimized for DB2 10
- More granular DBA privileges
 - Better complying with security regulations

New DBA privileges with finer granularity

Challenge

- Protect sensitive data from privileged users
- Separate authority to perform security related tasks
- **Reduce use of SYSADM authority**
- Allow EXPLAIN without execute privilege or ability to access data
- Audit the use of privileged users

Solution

Based on business security policies, new authorities can be granted

- New system administrative authorities:
 - SECADM
 - System DBADM
 - DATAACCESS
 - ACCESSCTRL
 - SQLADM
 - EXPLAIN
- **Fore REVOKE options:** For example, revoke SYSADM from a user using the new `NOT INCLUDING DEPENDENT PRIVILEGES` clause
 - Any grants made by the users do not get revoked

Q & A



schuetzner@de.ibm.com