#SHAREorg

SHARE
Technology · Connections · Results

DB2 10 for z/OS Technical Update

James Teng, Ph.D.
Distinguished Engineer
IBM Silicon Valley Laboratory
March 12, 2012

Disclaimers & Trademarks*

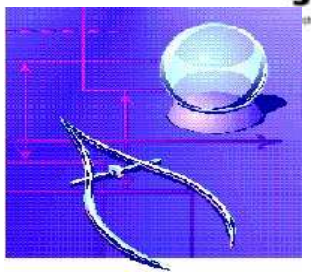

Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. Any discussion of OEM products is based upon information which has been publicly available and is subject to change. The opinions expressed are those of the presenter at the time, not necessarily the current opinion and certainly not that of the company.

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States and/or other countries: AIX, AS/400, DATABASE 2, DB2*, Enterprise Storage Server, ESCON*, IBM, iSeries, Lotus, NOTES, OS/400, pSeries, RISC, WebSphere, xSeries, z/Architecture, z/OS, zSeries, System p, System I, System z

The following terms are trademarks or registered trademarks of the Microsoft Corporation in the United States and/or other countries: MICROSOFT, WINDOWS, ODBC

For more copyright & trademark information see ibm.com/legal/copytrade.phtml

DB2 for z/OS The most robust and cost effective data server

	DB2	DB2 9	DB2 10
	▪ Deep synergy with System z ▪ HW Compression ▪ Consolidation	• Up to 20% utility CPU savings • Compress indexes, save 50% disk • Native SQL procedures • More CPU on specialty engines	• Save up to 5-10% CPU batch & transactions out-of-the-box (rebind) • On-the-fly data Compression • Temporal data support • Skip-level migration
	▪ Unmatched availability ▪ Unparalleled security ▪ Industry leading reliability	• Flexible context and role security • Expanded online schema changes • Volume level backup & recovery	• Ten times more concurrent users • More online schema changes • More granular access control
	▪ Near-linear scalability ▪ Optimized for SOA ▪ Flexible development ▪ Warehousing capabilities	• Seamless integration of XML and relational • Improved SQL • Partition by growth • OLAP expressions	• Enhanced query parallelism • More SQL compatibility • Improved pureXML and SQL PL

V8 out of service April 2012

DB2 Deep Synergy With System z

Key integration points include:

- Data sharing (availability and scale out)
- zIIP and other specialty engines
- Unicode conversion
- Encrypted communication & data
- Hardware data compression & encryption
- Cross-memory, memory protection keys
- Sorting
- Multi-core, large N-way
- 64-bit addressing and large memory
- z/OS Workload Manager
- z/OS Security Server (RACF)
- z/OS RRS integrated commit coordinator
- System z10 1 MB page size, decimal float
- Solid state disks
- zEnterprise z196, zBX, z10, …

zEnterprise 196 Benefits for DB2

Taking System z to the next level

- Faster CPUs, more CPUs, more memory means better DB2 performance, scalability
  - Excellent synergy with DB2 10, which will remove many single system scaling inhibitors
- Large cache to benefit DB2 workloads
- TLB Changes to improve DB2 10 performance for 1MB page sizes
- Hybrid architecture to open up new opportunities
  for DB2 query performance acceleration

DB2 and zEnterprise 196

- CPU reduction in all types of DB2 workloads
  - Larger processor cache (1.5MB L2 per core, 24MB L3 per chip, 192MB L4)
  - Various types of DB2 9 and 10 workloads show 20% to 40% DB2 CPU reduction compared to z10 processors.

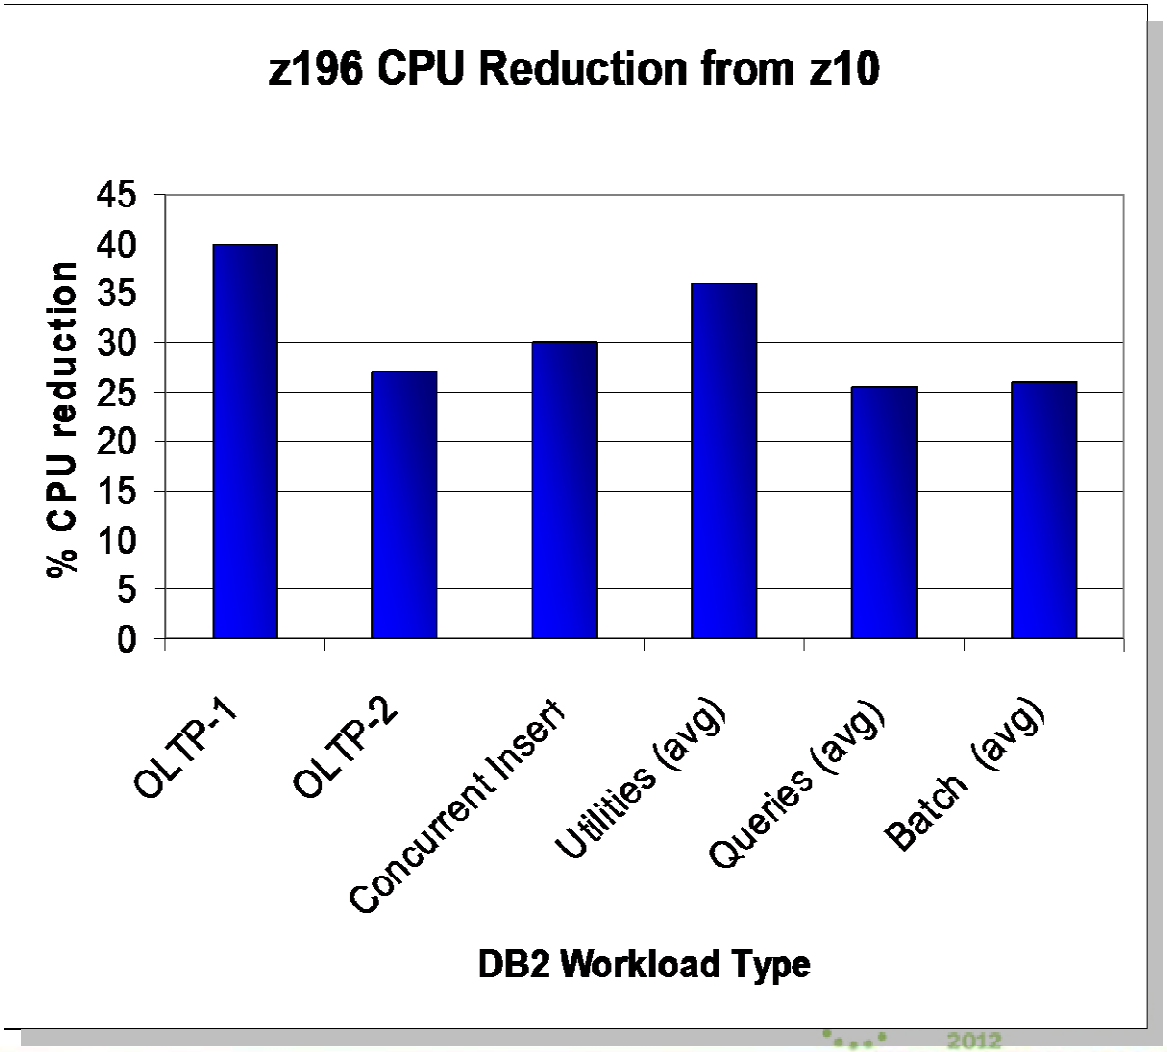

DB2 & IBM zIIP Add Value to Database Work

1 – **DRDA over TCP/IP connections**

- **DB2 9 for z/OS Remote native SQL procedures**
- **DB2 9 XML parsing via DRDA to fully utilize zIIP**

- **Increased portion of DRDA redirected to zIIPs to 60%**
  Improved performance via reduced processor switching

2 - Requests that use parallel queries

- ***DB2 9 higher percentage of parallel queries zIIP eligible***
- **DB2 10 more queries eligible, more parallelism**

3 - DB2 Utilities LOAD, REORG & REBUILD

- **DB2 utility functions used to maintain index structure and sort**

- **DB2 10 RUNSTATS - most options**

4 - DB2 10 buffer pool prefetch and deferred write

DB2 10 for z/OS

- **Fastest uptake**
  - **+2X customers**
  - **+3X licenses**
  - **25% coming from V8**
- **More customers in production**
  - **SAP, data warehouse and OLTP workloads**
  - **Skip-level and V2V**
- **Quality/stability looking good**
  - **Half TFSA per customer**
  - **Flat over the last 6 months**

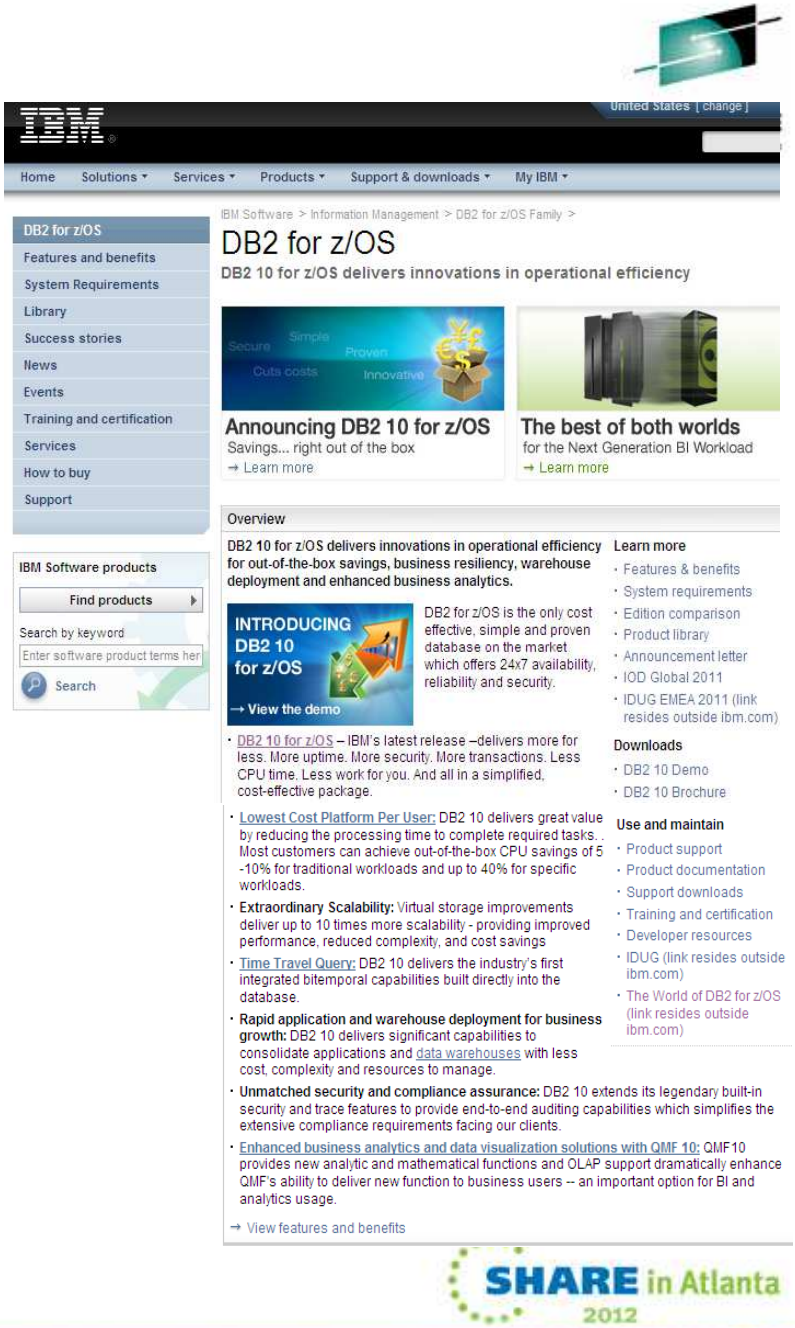

DB2 10 for z/OS: Cost Savings

CPU reductions for transactions, queries, and batch

- CPU reductions of 5-10% for traditional workloads
- CPU reductions of up to 20% for new workloads
- Up to additional 10% CPU savings using new functions
- For static SQL, REBIND typically required

Scales with less complexity and cost

- 5-10x more concurrent users – up to 20,000 per subsystem
- Significant scale-up capabilities in addition to existing scale-out support
- Consolidate to fewer LPARs and subsystems

Improved operational efficiencies and lower administration cost

- Automatic diagnostics, tuning, and compression

Even better performance

- Elapsed time improvement for small LOBS and Complex Queries

DB2 10 Danske Bank Experiences

Why DB2 10 Now?

- Reduced cost
- Improved scalability
- [illegible]

Benefits

- **90%** Virtual storage savings
- **10%** CPU savings on CICS transactions
- **30%** CPU savings on test batch workload

"Our DB2 10 experience has given us confidence about the virtual storage relief and CPU savings. I am looking forward to continuing our rollout and reaping the benefits."

Niels Simanis
Senior Technology Manager
Danske Bank

Migration Tips

- Plan well, including good maintenance practices
- Rebind can get you the highest CPU savings
- Expect increase in real storage consumption to support and exploit DB2 10

DB2 10 Dillard's Experiences

Why DB2 10 Now?

- Reduced cost
- Improved performance
- [illegible]

Benefits

- **20-30%** CPU savings out-of-the-box
- **5-15%** Performance improvements for batch, CICS, and DDF

Actual results may vary for other customers

> ***"We are pleasantly surprised with the out-of-the-box CPU savings we have seen during testing and early production phases."***
>
> **Terry Glover –Director IT Infrastructure Dillard's**

Migration Tips

- Thorough preparation and planning
- Good maintenance practices

Virtual storage improvements

- DBM1 below 2GB
  - 75-90% less usage in DB2 10 compared to DB2 9
  - Some of working storage (stack, xproc storage) stays below 2GB
- Larger number of threads
  - Possible data sharing member consolidation
- Improve CPU with storage
  - More release deallocate
  - Larger MAXKEEPD values for KEEPDYNAMIC=YES

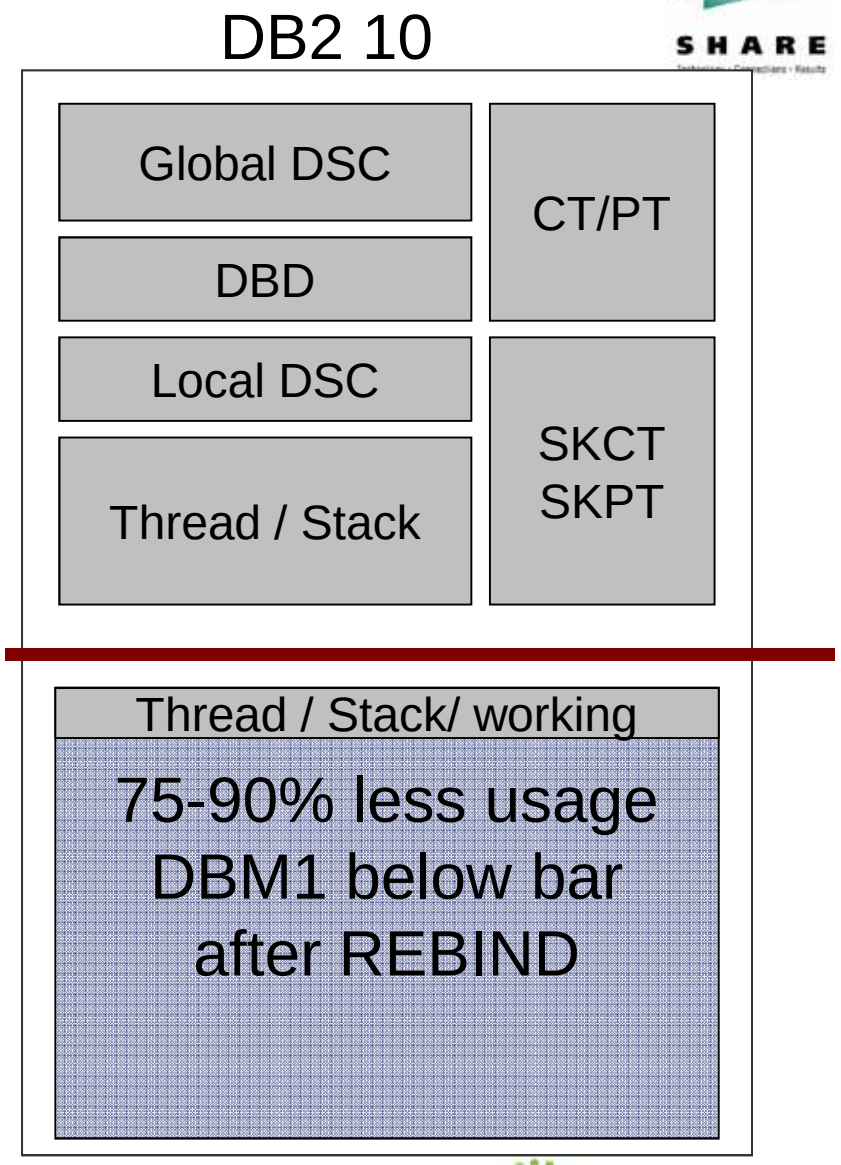

Running Many Active Threads

Today

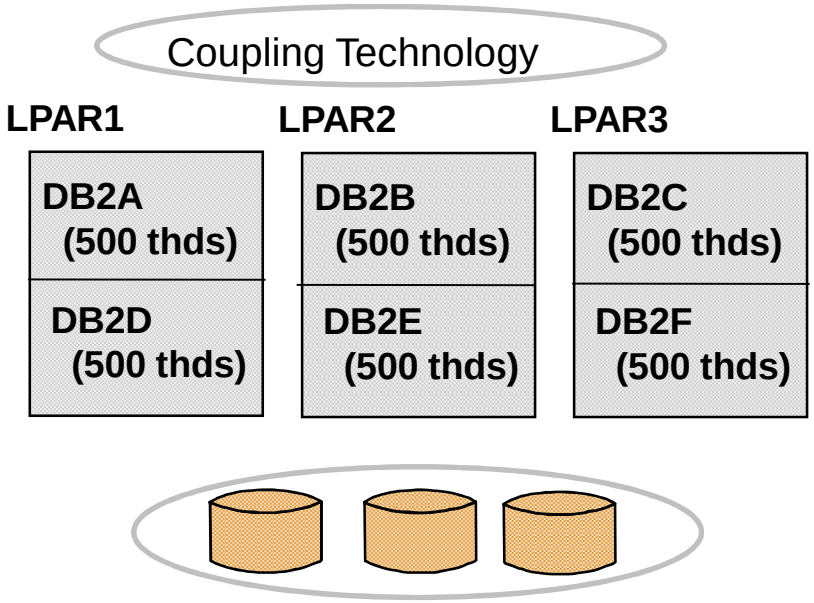

- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s per LPAR

DB2 10

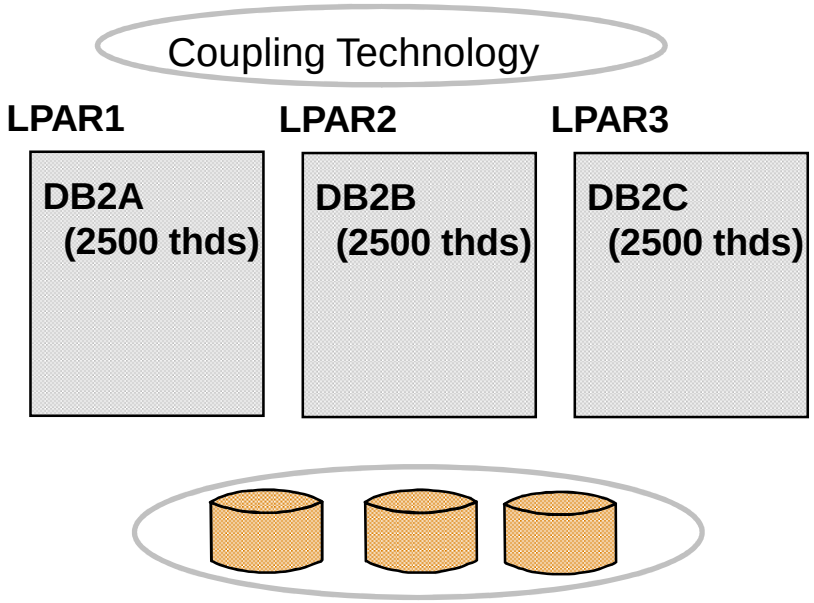

- More threads per DB2 image
- More efficient use of large n-ways
- Easier growth, lower costs, easier management
- Data sharing and Parallel Sysplex still required for very high availability and scale
- Rule of thumb: save ½% CPU for each member reduced, more on memory

Major changes in DB2 10 catalog & directory

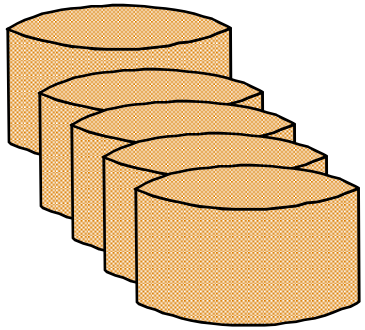

- Improve availability and productivity
- Increase maximum size substantially
- Reduce contention: BIND, Prepare, utilities
  - DDL concurrency also improved from removal of DBD01 hash anchor locks
- Catalog changes: Remove links, hashes
  - Many more table spaces, partition by growth
  - Row level locking, reordered row format
  - CLOB and BLOB columns for long strings
    - Inline for performance
  - Online reorganization and check
  - More automatic: DB2-managed SMS-controlled

Data Sharing Improvements

- ACCESS DATABASE command wildcarding support - V9 PK80925
- Sub-group attach
- BP scan avoidance
- Delete data sharing member
- MEMBER CLUSTER support for UTS
- DDF Restart Light enhancements: Handle DDF indoubt URs
- Online DDF changes
- Auto rebuild CF lock structure on long IRLM waits during restart
  - Can avoid group-wide shutdowns
- LRSN spin avoidance for inserts to the same page (e.g. Multi Row Insert)
- New zparm to force deletion of CF structures on group start
- Expedited GBP DELETE_NAME processing
  - Avoid sending XI signals by deleting data only
  - Avoid potential lock timeout conditions when there are lots of directory entries for an object

Online Schema Enhancements

- Table Space Changes - deferred
  - Page Size – UTS only
  - DSSIZE – UTS only
  - SEGSIZE – UTS only
  - MAXPARTITIONS
  - Convert single table simple/segmented into PBG
  - Convert classic partitioned table into PBR
  - Convert PBR to PBG
  - Convert Classic Partitioned/PBR/PBG to Member Cluster
  - Table Space is put into Advisory REORG pending
- Index Changes - deferred
  - Page Size
  - Index is put into Advisory REORG pending

Online Schema Enhancements …

- Deferred Alters are materialized at the next table/index space level REORG
  - Before REORG, pending changes can be dropped via ALTER … DROP …
  - Quiesce applications using the DBD, plan/package locks during the SWITCH phase
  - Invalidate plans/packages/dynamic statement cache

- Alter Buffer Pool with the same page size – immediate
  - No longer need to STOP TS/IS in data sharing
  - Use DRAIN(ALL) to quiesce applications
  - New BP will be used after commit

Index Improvements

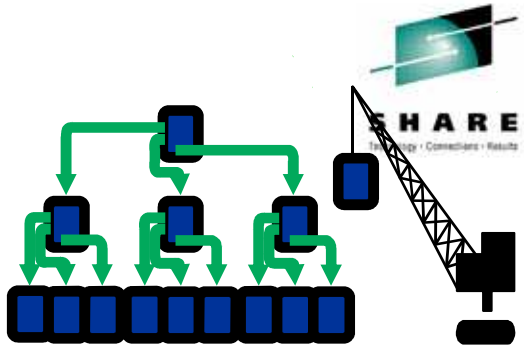

- Parallel index update at insert
  - For tables with more than 2 indexes
  - I/O parallelism
- Index with included columns
  - Add non-key columns
- Cache the index root page in buffer pool during open
- List prefetch for disorganized indexes
- CPU reduction for index access especially with NOT PADDED indexes on VARCHAR
- New IFCID 359 to monitor index split

Buffer Pool Enhancements

- Use 1MB real storage frame, instead of 4K, for PGFIX = YES Buffer Pools
  - Reduce cpu overhead when accessing pages in BP
  - Available only for z10 and zEnterprise
- Prefetch/Deferred Writes are running under zIIP
- Avoid BP scans in data sharing
  - Switch from non-Shared to Shared and vice versa
  - Data set close
  - During STOP DB2
  - Significantly improve performance for large BPs

Buffer Pool Enhancements …

- In-memory tables/Indexes
  - A new PGSTEAL = NONE on ALTER BUFFERPOOL
  - Load data into the buffer pool at data set open
    - ✓Done by DB2 Prefetch Engines under zIIP
  - Ensure BP size is large enough to cache all data/index pages
  - BP is managed in FIFO (First In First Out) order
    - ✓Not managed in LRU
    - ✓Disable all prefetch requests
  - Query optimizer uses zero I/O cost for tables and indexes using this Buffer Pool
  - Recommend to use AUTOSIZE(NO) which is the default
  - Recommend to set CLOSE = NO for tables/indexes

Hashed Table

SELECT * …WHERE
ITEMNO = 'W0133-1662996'

Locate the data row
by hashing the key value.

-The hash Key must be unique
-Reduced page visits
-Reduced CPU & elapsed time
-Possibly eliminate an index

-Tradeoff: extra space used

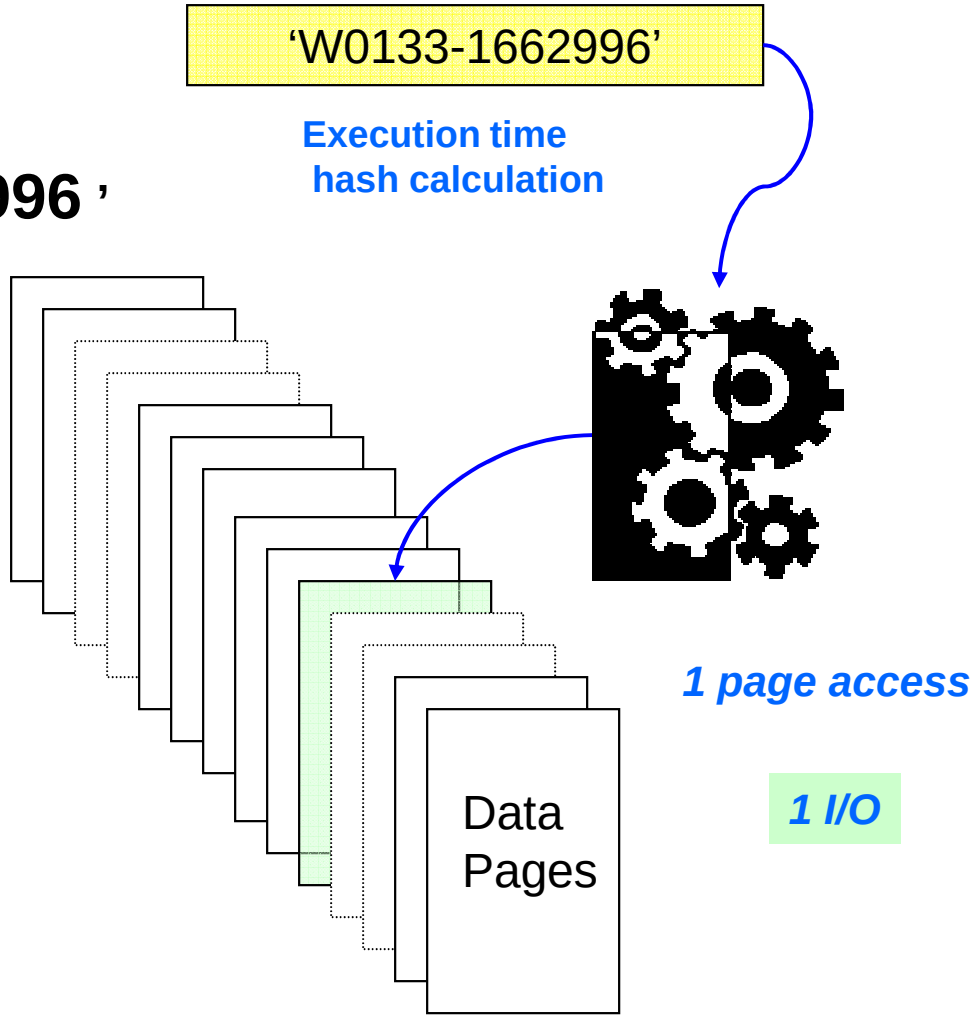

Versioned data or Temporal Data

- Table-level specification to control data management based upon time
- Two notions of time:
  - System time: notes the occurrence of a data base change
    - “row xyz was deleted at 10:05 pm”
    - Query at current or any prior period of time
    - Useful for auditing, compliance
  - Business time: notes the occurrence of a business event
    - “customer xyz’s service contract was modified on March 23”
    - Query at current or any prior/future period of time
    - Useful for tracking of business events over time, application logic greatly simplified
- New syntax in FROM clause to specify a time criteria for selecting historical data

Logging Enhancements

- Reduce Log Write Latch (LC19) contention
  - Eliminate CPU spin loop to generate an unique LRSN within a data sharing member – V9
  - LRSN will only need to be unique for updates against the same data/index page – V9
  - Allow multiple rows inserted into the same data page with the same LRSN – V10
- Parallel log force for dual active logs
- Option to take system checkpoint based on time or logs, whichever occurs first
- Online to add new active log data sets via
  - SET LOG NEWLOG(dsn) COPY(1 or 2)
  - Recommend to format the data set with DSNJLOGF

Utility Enhancements

- Removed UTSERIAL lock for greater utility concurrency
- Support data set level FlashCopy for COPY, RECOVER, LOAD, and REORG
- New BACKOUT YES option for point in time recovery
- zIIP-enablement for RUNSTATS
- Auto sampling rates & page sampling instead of row sampling
- REORG enhancements for LOB
  - Support SHRLEVEL CHANGE option
  - Permit rows to flow between partitions
  - Allows REORG REBALANCE with LOB columns
  - Allows ALTER of LIMITKEY with LOB columns
- LOAD/UNLOAD – spanned record support for LOB and XML

DB2 10 Summary

- Many opportunities for price/performance (cost) improvements
  - Major theme of this release
  - Most welcome to our customers
- Significant DBM1 31-bit VSCR after rebind
  - Opportunity for scale up and LPAR/DB2 consolidation
- REBIND required to obtain most performance and VSCR improvements
- Plan, provision, and monitor real storage consumption


Thank
You