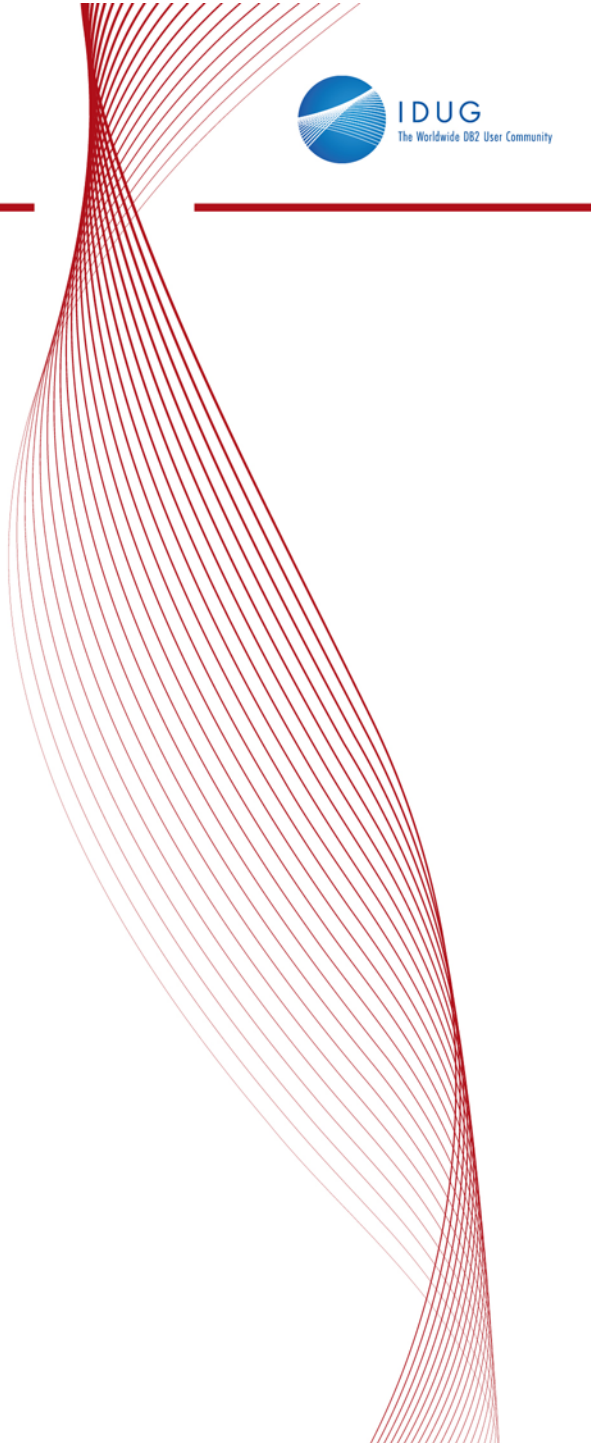

DB2 10 for z/OS Overview

Jeff Josten
IBM
josten@us.ibm.com

Session Code: S01

Tues. May 11 10:00 – 11:30
Platform: z/OS

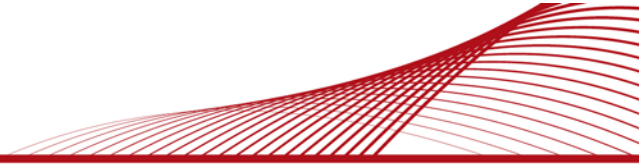

Disclaimer/Trademarks

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion. *

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AS/400, DataJoiner, DataPropagator, DB2, DB2 Connect, DB2 Extenders, DB2 OLAP Server, DB2 Universal Database, Distributed Relational Database Architecture, DRDA, eServer, IBM, IMS, iSeries, MVS, Net.Data, OS/390, OS/400, PowerPC, pSeries, RS/6000, SQL/400, SQL/DS, Tivoli, VisualAge, VM/ESA, VSE/ESA, WebSphere, z/OS, zSeries

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

DB2 for z/OS robust and cost effective

	DB2	DB2 9	DB2 10
Efficiency	• Deep synergy with System z • HW Compression • Mixed workloads at high utilization	▪ 20%-30% Utility CPU savings ▪ Compress indexes, save 50% disk ▪ More CPU on specialty engines	▪ Save up to 20% CPU batch & transactions ▪ Query performance ▪ Hashed data access ▪ Skip-level migration
Resilience	• Unmatched availability • Unparalleled security • Industry leading reliability	▪ Flexible context and role security ▪ Expanded online schema changes ▪ Volume level backup & recovery improvements	▪ 10x more concurrent users ▪ More online schema changes ▪ More granular access control ▪ Access path stability
Growth	• Near-linear scalability • Flexible development • Warehousing capabilities	▪ Seamless integration of XML and relational ▪ Improved SQL ▪ Partition by growth ▪ OLAP expressions	▪ Temporal data ▪ More SQL features ▪ Improved pureXML and SQL PL

Beta Announced: Feb 9, 2010

DB2 for z/OS Technical Strategy

- Extend the lead in availability, scalability and performance.
  - Parallel Sysplex: the best scale-out solution in the industry
  - Tight integration between DB2 and the System z hardware and z/OS operating system
  - Advanced solutions for compliance with data security and privacy regulations
  - Workload consolidation: System z is the ultimate consolidation platform
  - Eliminate all causes of outages
- Reduce cost of ownership
  - DB technology that can handle large workloads with fewer people
  - Advanced autonomics to make the system more self-managing and self-tuning
  - Storage and CPU optimization, including specialty engines
- Application enablement
  - Apps can easily connect to DB2 from anywhere
  - Application portability and DB2 family compatibility
  - Advanced SQL, XML capability. Easy application development
  - Expand data warehousing and analytics capabilities

DB2 for z/OS Into the Future

Delivering Customer Value

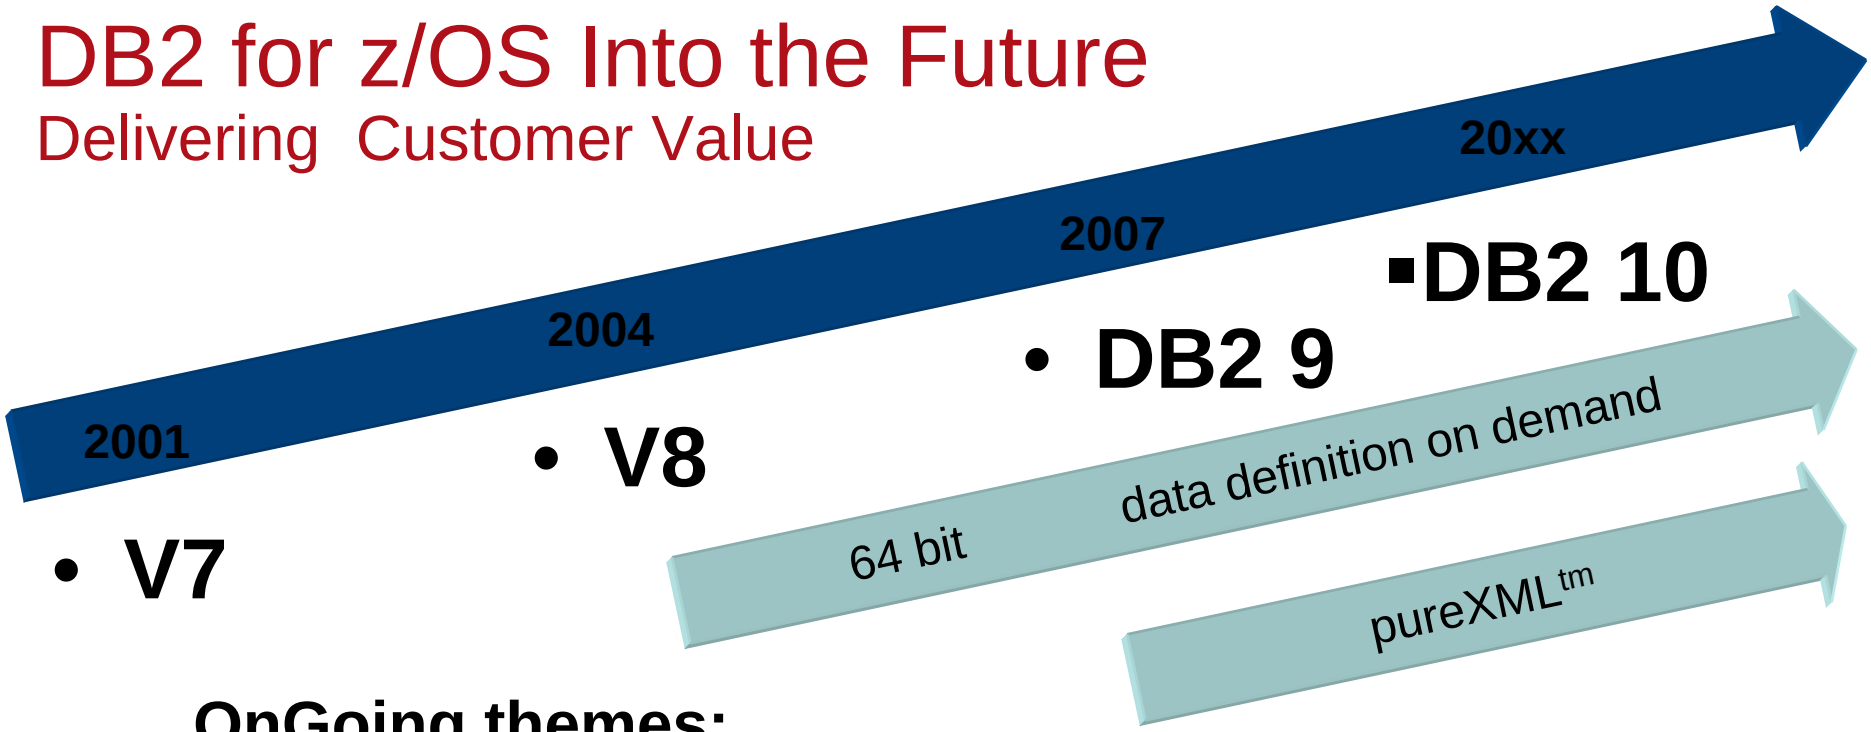

OnGoing themes:

Performance Scalability
Reliability Availability Serviceability
Security Productivity
Application Development
SQL XML

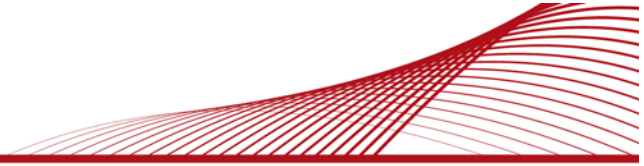

DB2 10 for z/OS At a Glance

Addressing Corporate Data Goals

RAS, Performance, Scalability, Security	• Wide range of performance improvements • More online schema changes • Catalog restructure for improved concurrency • Fine grained access control • Hash access to data • New DBA privileges with finer granularity • Query parallelism advancements
Simplification, Reduced TCO	• Full 64-bit SQL runtime • Auto stats • Data compression on the fly • Query stability and management enhancements • Reduced need for REORG • Utilities enhancements • More granular DDF thread management • Statement level diagnostics
Application Enablement	• pureXML enhancements • SQLPL enhancements • Temporal data • Last Committed reads • Timestamp with timezone • Many new SQL features • Moving sum, moving average

DB2 10 for z/OS: Savings Out-of-the-Box

Up to 20% CPU reductions for transactions, queries, and batch

- Out-of-the-box CPU reductions of 5-10% for traditional workloads
- Out-of-the box CPU reductions of up to 20% for new workloads
- Up to additional 10% CPU savings using new functions

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 Performance Objective

- Historical goal under 5% performance regression
- Goal 5% -10% initial performance improvement
- Many customers reduce CPU time 10% - 20%

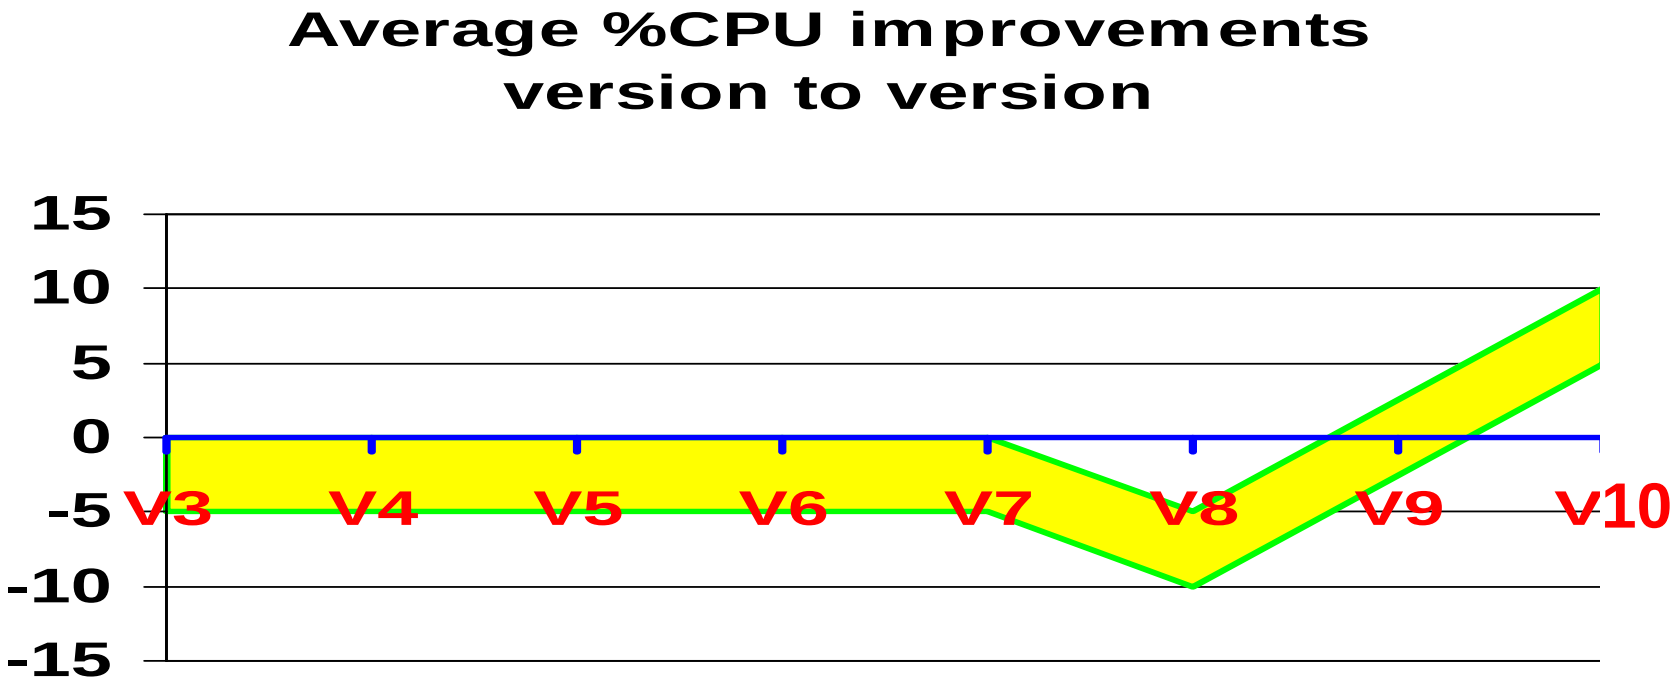

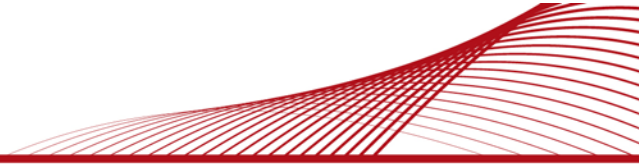

Performance Enhancements with Few Changes (CM)

- SQL runtime improved efficiency
- Address space, memory changes to 64 bit, some REBINDs
- Faster single row retrievals via open / fetch / close chaining
- Distributed thread reuse High Performance DBATs
  - RELEASE(DEALLOCATE) can be more aggressively used with v10 vstor relief
    - IBM measurements show up to 10-20% cpu savings
  - DDF enforcement of RELEASE(COMMIT) is removed
  - KEEPDYNAMIC=YES will not get benefit from this enhancement
- DB2 9 utility enhancements in CM8
- Parallel index update at insert
- Workfile in-memory enhancements
- Index list prefetch
- Solid State Disk use
- Buffer pool enhancements
  - Utilize z10 1MB page size
  - “Fully in memory” option (ALTER BUFFERPOOL)

Performance Enhancements requiring REBIND (CM)

- Most access path enhancements
- SQL paging performance enhancements
  - Single index access for complex OR predicates
  - New Range-list Index scan access path
- IN list performance
  - Optimized Stage1 processing (single or multiple IN lists)
  - Matching index scan on multiple IN lists
- Query parallelism improvements
- More stage 2 predicates can be pushed down to stage 1
- More aggressive merge of views and table expressions
  - Avoid materialization of views
- REBIND enables further SQL runtime improvements
- If migrate from V8, get new RUNSTATS before mass rebind

Performance Enhancements requiring NFM

- Efficient caching of dynamic SQL statements with literals
- Most utility enhancements
- LOB streaming between DDF and rest of DB2
- Faster fetch and insert, lower virtual storage consumption
- SQL Procedure Language performance improvements
- Workfile spanned records, PBG
- Insert improvement for UTS
- Local JDBC (Type2) and ODBC application performance
  - Limited block fetch, LOB progressive streaming, progressive CLOSE now available for JCC type2 and ODBC z/OS drivers
- Solid State Disk monitoring and exploitation

Performance Enhancements requiring NFM + DBA work

- Hash access path — Alter + Reorg + rebind to activate
- Index include columns — Alter + Rebuild + rebind to activate
- Inline LOBs — Alter (need UTS and RRF)
  - Index on expression now possible for LOB columns
  - Important for spatial performance
  - LOAD/UNLOAD perfromance improvements
  - LOB compression for inline portion
- MEMBER CLUSTER for UTS
- DEFINE NO for LOB and XML columns

Index to Data Access Path

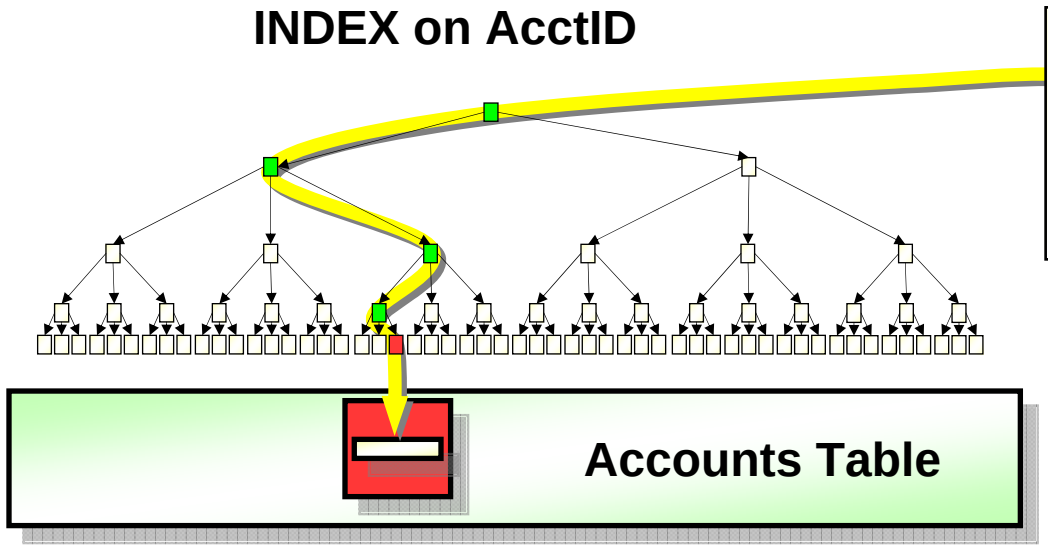


```
Select Balance
From Accounts
WHERE acctID = 17
```

= Page in Bufferpool

= Page Read from Disk

- Traverse down Index Tree
  - Typically non-leaf portion of tree in the bufferpool
  - Leaf Portion of the tree requires I/O
  - Requires searching pages at each level of the index
- Access the Data Page
  - Typically requires another I/O
- For a 5 Level Index
  - 6 GETP/RELPs, 2 I/O's, and 5 index page searches

Hash Access

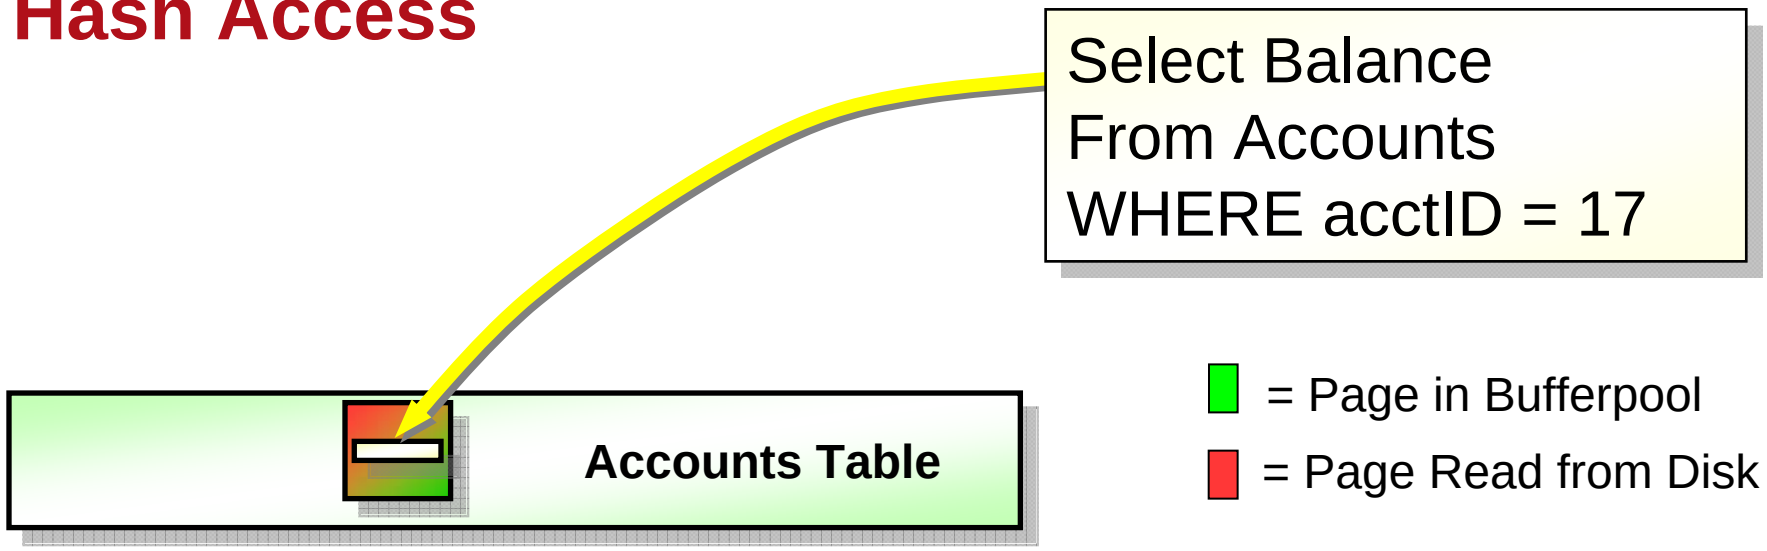

- Hash Access provides the ability to directly locate a row in a table without having to use an index
- Single GETP/RELP in most cases
- 1 Synch I/Os in common case
  - 0 If In Memory Table
- Greatly reduced Search CPU expense

Queries against Hashed Tables

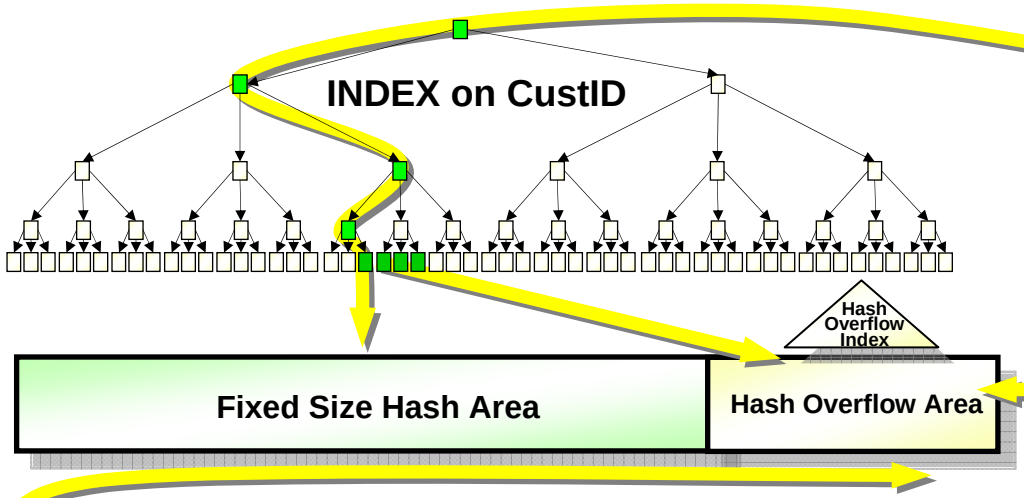

Index Lookup/Scan

```
Select Balance
From Accounts
WHERE CustID = 1717
```

Hash Access

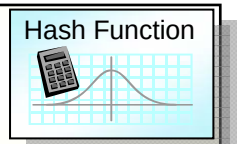


```
Select Balance
From Accounts
WHERE AcctID = 17
```

Table Space Scan

```
Select CustID
From Accounts
WHERE Balance > 100000
```

- Hash Access Path
  - Great for Equality and IN predicates
  - Can't do range scans
- Secondary indexes can be defined for Range Scans
- Table Space Scans still supported
- Hash Access can be used to enforce Primary Key and Unique constraints

Sizing a Hash Space

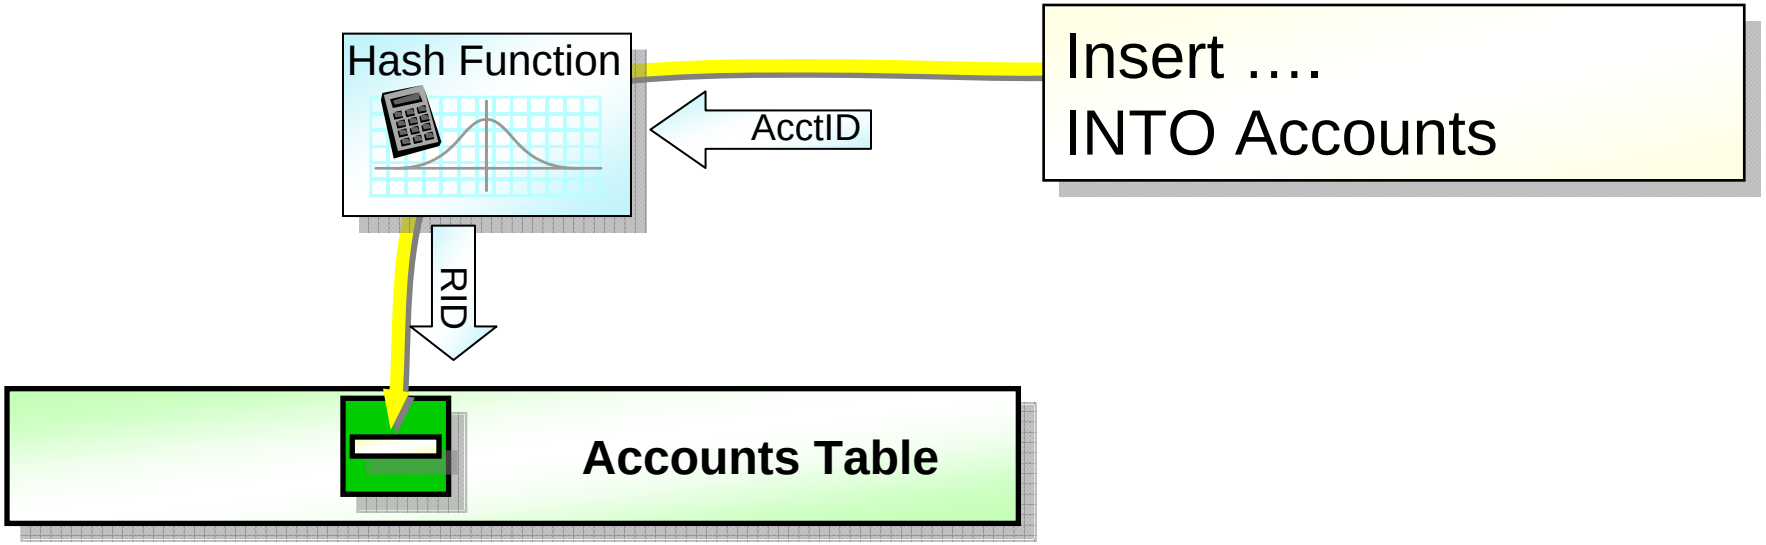

- What happens at insert time
  - We compute a ‘Random’ Page to insert on
    - If the Page is Full, the new entry becomes an overflow
    - The Page is ‘Randomly’ chosen, so there will be statistical variation
  - Suppose 20 Rows fit in a Page
  - What if we hash 20 Million rows into 1 Million pages
    - On Average, 20 rows will hash to each page

Hashing Summary

- Provides fast, direct location of most rows
  - Reduces I/O and CPU in most cases
  - Can replace an existing Primary or Unique Key Index
    - Faster Insertion/Deletion
- Size of Fixed Size Hash Area is important
  - Too small and performance degrades
  - Too large and space is wasted
- DB2 helps you manage the size
  - REORG AUTOESTSPACE YES
  - RTS tracks the number of overflowed entries

DB2 10: 64 bit Evolution (Virtual Storage Relief)

Scalability: Virtual storage constraint is still an important issue for many DB2 customers, until DB2 10

- DB2 10 supports 5-10x more active threads, up to 20,000 per member
  - 80-90% of thread storage moved above the bar
  - More concurrent work
  - Reduce need to monitor
  - Consolidate members and LPARs
  - Reduced cost
  - Easier to manage
  - Easier to grow

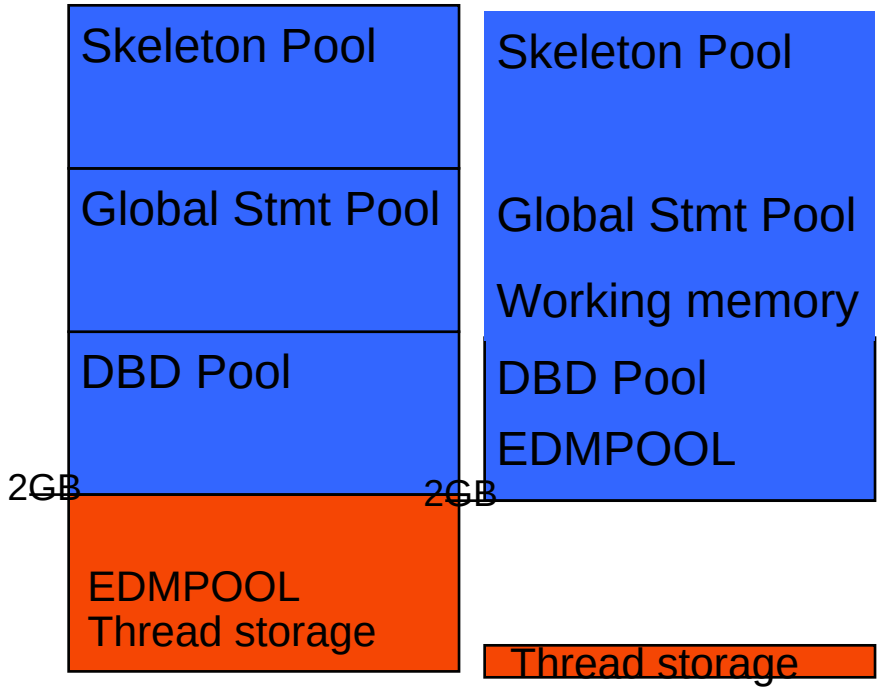

Running a Large Number of Threads

Today

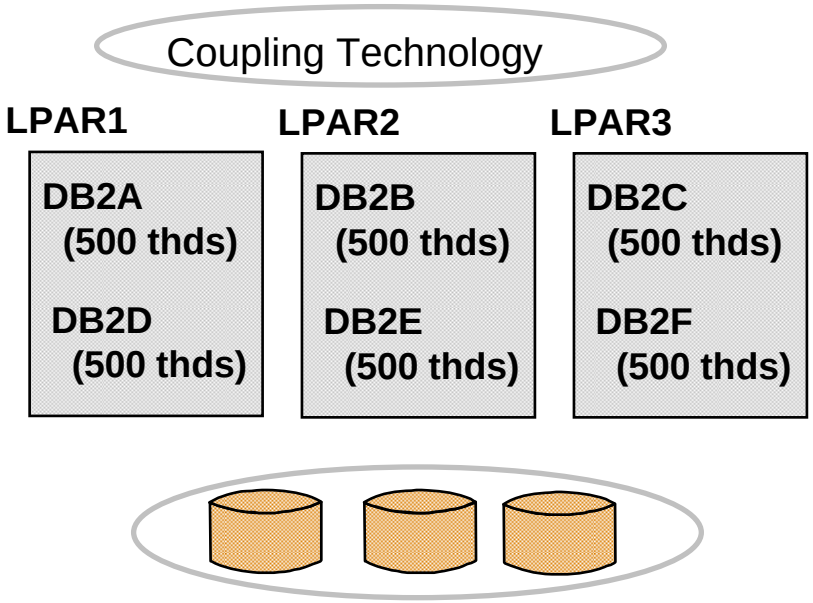

- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s / LPAR

DB2 10

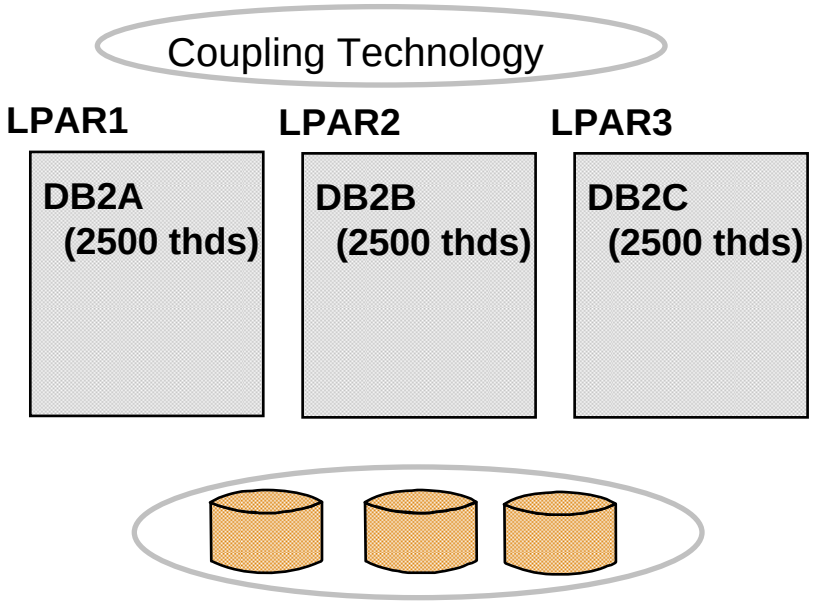

- Easier growth, lower costs, easier management
- More threads per DB2 image
- More efficient use of large n-ways
- Potential for fewer members / LPARs
- Rule of thumb: save ½% CPU for each member reduced, more on memory
- Data sharing and Parallel Sysplex still required for very high availability and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

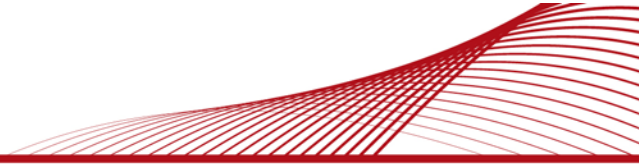

Other System Scaling Improvements

- Other bottlenecks can emerge in extremely heavy workloads
  - several improvements planned to reduce latching and other system serialization contention
  - new option to for readers to avoid waiting for inserters
  - eliminate UTSERIAL lock contention for utilities
  - Use 64-bit common storage to avoid ECSA constraints
- Concurrent DDL/BIND/Prepare processes can contend with one another
  - restructure parts of DB2 catalog to avoid the contention
- SPT01 64GB limit can be a constraint, especially if package stability is enabled
  - Allow many more packages by using LOBs
- Improved accounting rollup, compress SMF option
- Lower overhead for very large buffer pools

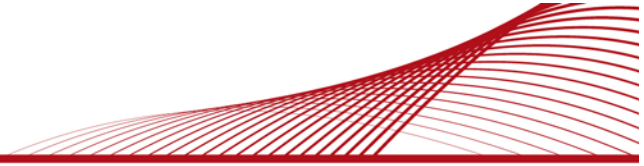

DB2 10 Catalog and Directory Improvements

- Remove links and enable row level locking on key catalog tables
  - Improved concurrency for BIND, PREPARE, and DDL
  - DSN1CHKR no longer needed in DB2 10 NFM
- Move long strings such as SQL statements and package binaries to LOBs
  - Removes max size constraint for SPT01 and others
  - Easier to query SQL statements from catalog
- Online REORG for all catalog and directory table spaces
- Allow SELECT from SYSLGRNX
- Easier management: DB2 managed and SMS controlled

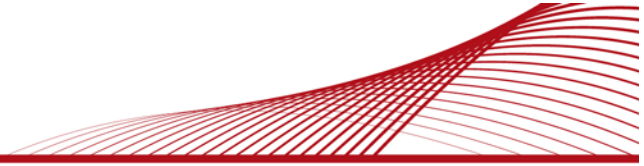

Data Sharing Improvements

- ACCESS DATABASE command wildcarding support - V9 PK80925
- Sub-group attach (v9 usermod)
- BP scan avoidance
- Delete data sharing member
  - Offline utility for “deactivate”, “reactivate”, “destroy”
  - Online utility will come later
- 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)
- IFCID 359 for index split
- New zparm to force deletion of CF structures on group start (e.g. DR testing)

Availability – More Online Schema Changes

- Online schema changes for table spaces, tables and indexes – PENDING with ALTER and Online REORG instead of DROP/CREATE or REBUILD INDEX
  Alterations occur with REORG, unless noted otherwise
  - Page size for table spaces and indexes
  - DSSIZE for table spaces
  - SEGSIZE
  - Convert single table segmented into UTS PBG
  - Convert single table simple into UTS PBG
  - Convert classic partitioned table space into UTS PBR
  - Convert UTS PBR to UTS PBG
  - Convert PBG to hash (immediate, but RBDP index)
  - MEMBER CLUSTER
  - Ability to drop pending changes

Availability – More Online Schema Changes…

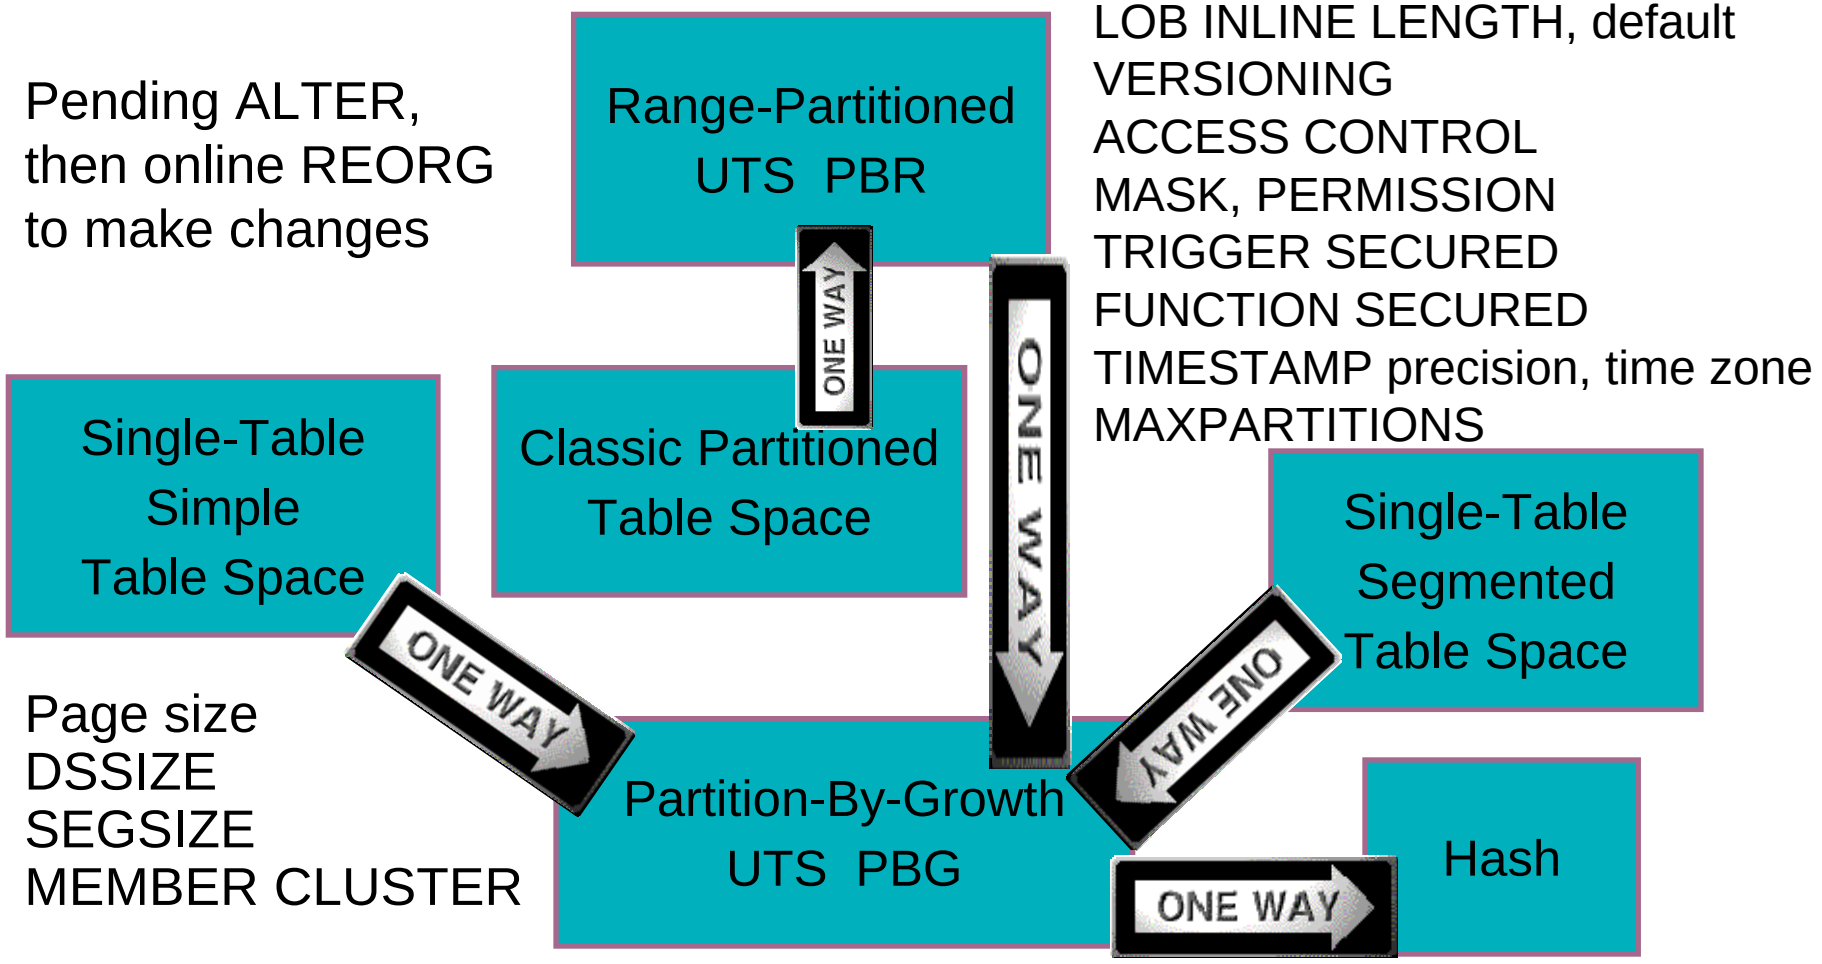

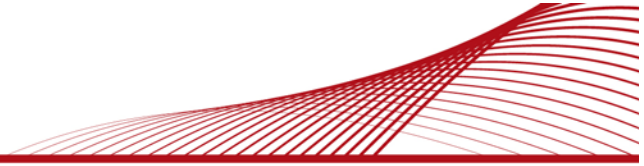

Other Availability Improvements

- Access currently committed data
- Change DDF location alias names online
- Online add active log

DB2 10 Utilities Enhancements – Online REORG

- REORG SHRLEVEL(CHANGE) for LOBs
- Online REORG enhancements
  - SHRLEVEL(CHANGE) support for all catalog/directory objects
  - Option to cancel blocking threads
  - Improved availability
    - Update stats after de-drain in UTILTERM phase
  - Allow disjoint partition ranges
  - Permit movement of rows between partitions when LOB columns exist
    - Allows REBALANCE and ALTER LIMITKEY even when LOB columns exist
    - Allows DISCARD to delete associated LOB values
  - Messages to estimate length of REORG phases and time to completion

DB2 10 more utilities enhancements

- Support of spanned records for UNLOAD of LOB data
  - Currently unload of LOBs >32K must use FRVs
  - Allow inline LOBs with base row in unload data set
  - Provides portability of data
  - Performance enhancement for FRV processing with PDS data sets, also in DB2 9
    - UNLOAD 33% elapsed time reduction
    - LOAD 84% elapsed time reduction
- Autonomic RUNSTATS & table profile

DB2 10: More Utility Improvements

- Improved COPY CHANGELIMIT performance
  - Use RTS instead of SM page scans
- Data set level FlashCopy option
- FlashCopy backups with consistency and no application outage
- FlashCopy backups as input to:
  - RECOVER (fast restore phase)
  - COPYTOCOPY, DSN1COPY
- RECOVER “back to” log point
- REPORT RECOVERY support for system level backups

DB2 10 Productivity – Doing More With Less!

- Auto statistics collection
- Easier scaling, simpler memory management
- Reduce contention, more online processing
- Access path stability
- Reduced need for REORG
  - Build compression dictionary on the fly
  - Index list prefetch enhancements
- Configure IBM UDFs and stored procedures
- Allow one SDSNEXIT data set for many subsystems
- Monitoring enhanced
  - Timeout / deadlock diagnostics
  - Identify SQL statements

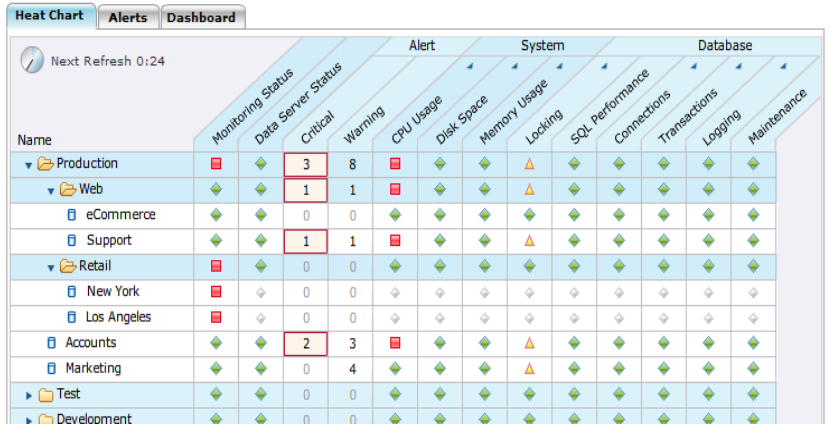

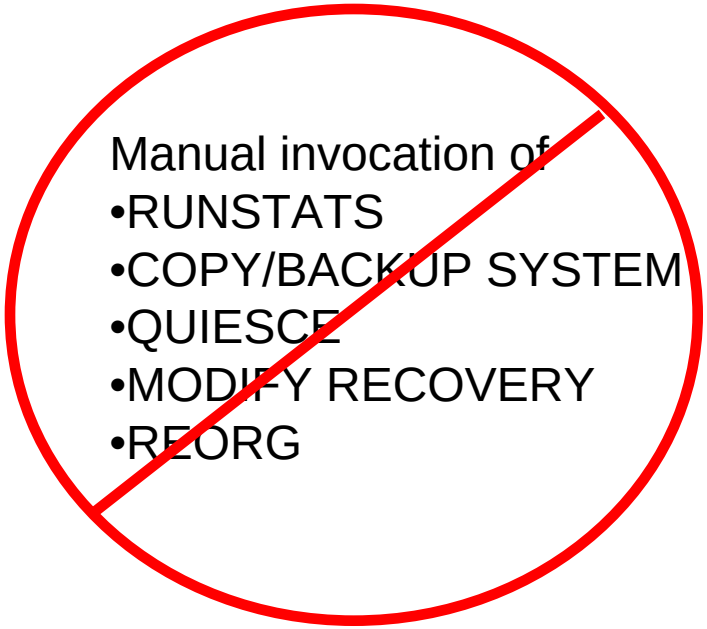

Autonomics and DBA Productivity…

- Checkpoint intervals based on both time and # log records
- Run ‘must complete’ backout under pre-emptable SRB
- Identify unused packages
- SQL Statement level monitoring
  - Statement ID introduced
  - Trace records & messages extended to include statement ID
  - New trace class for statement detail
    - GetPages, Locks, I/Os, cpu/elapsed time, etc. at statement level
    - Available via IFI for online monitors or via SMF/GTF
- Exploit z/OS 1.10 WLM service to temporarily boost priority of blocking lock holders
  - Complements V9 health monitor task which handles latches
- Manage max threads, connections, idle thread timeout on an application basis
  - Warning or exceptions issued when threshold is hit
  - Profiles can be set based on userids, packages, IP addresses, member names, …

Optimization Stability and Control

Provide unprecedented level of stability for query performance by stabilizing access paths for

- Static SQL - Relief from REBIND regressions
- Dynamic SQL
  - Remove the unpredictability of PREPARE
  - Extend Static SQL benefits to Dynamic SQL

Support:

- Access path repository
- Versioning
- "Fallback"
- "Lockdown"
- Manual overrides. Hints: easily influence access paths without changing apps
- Per-statement BIND options

- Safe query optimization: assess "reliability" of access path choices
- Adaptive algorithms, e.g. RID pool overflow to workfiles

DB2 10: Business Security and Compliance

- Protect sensitive data from privileged users & improve productivity
  - SECADM & DBADM without data access
  - Usability: DBADM for all DB
  - Revoke without cascade
- Separate authorities to perform security related tasks, e.g. security administrator, EXPLAIN, performance monitoring and management
- Audit privileged users
- Row and column access control
  - Allow masking of value
  - Restrict user access to individual cells

Use disk encryption

DB2 10 Security Benefits

- More flexible authorization
- Separation of duties
- Do job without access to data
- Policies for audit
- Simpler control
- Tighter security
- Avoid cascade delete
- Avoid views and application security logic
- Allow more tools
- Evolve security policies
- Easier to manage security policy

→ Improved productivity & tighter security

Use disk encryption

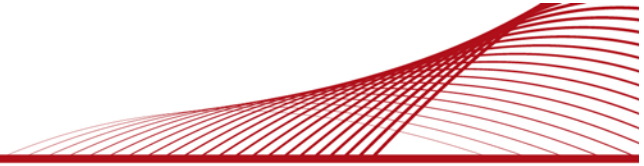

DB2 10 New Application Features

- Data versioning, temporal data
- pureXML enhancements
- Large object improvements
  - Allow non-NULL default values for inline LOBs
  - Loading and unloading tables with LOBs
    - LOBs in input/output files with other non-LOB data
- Currently committed locking semantics
- Implicit casting or loose typing
- Timestamp with timezone
- Greater timestamp precision – to picoseconds
- Moving Sum, Moving Average

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

Current and History

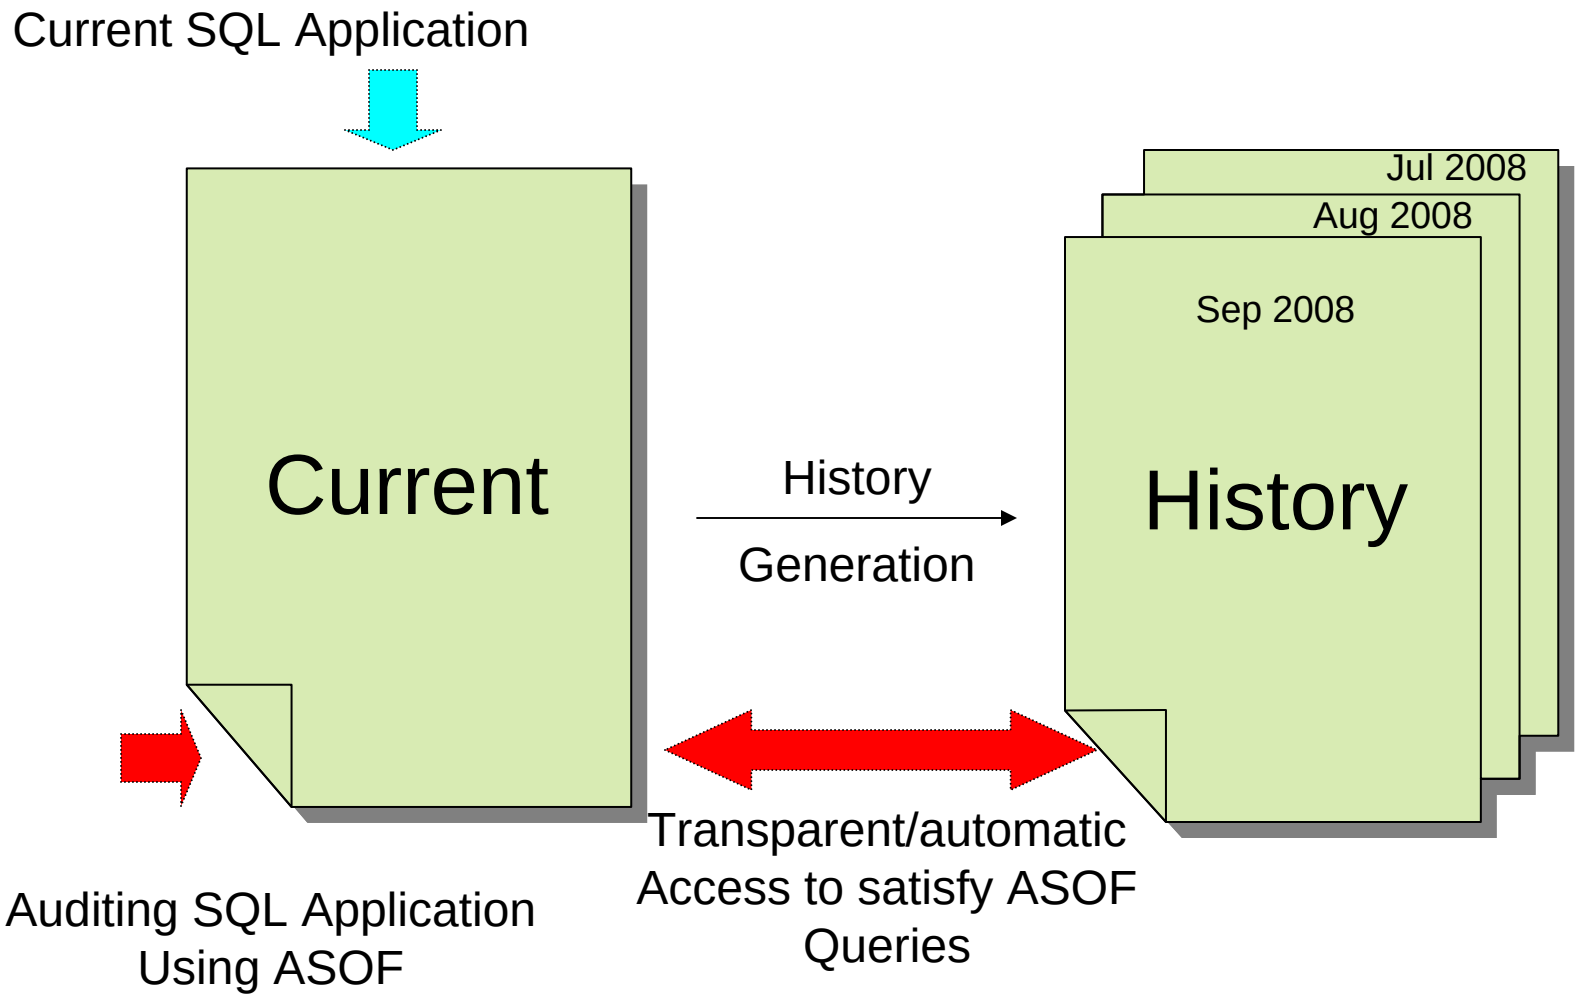

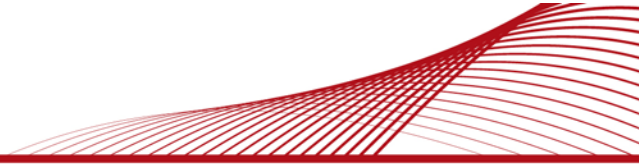

Temporal UPDATE example (business time)

Simple table definition (Policy#, start, end, coverage)

Table has 1 row of (123,'01/01/2001', '12/31/2001', 1000)

```
UPDATE policy p
   FOR BUSINESS_TIME FROM DATE('03/01/2001') TO DATE('03/31/2001')
   SET coverage = 2000;
```

Result of the update statement is 3 rows:

(123,'01/01/2001','03/01/2001',1000)
(123,'03/01/2001','03/31/2001',2000)
(123,'03/31/2001','12/31/2001',1000)

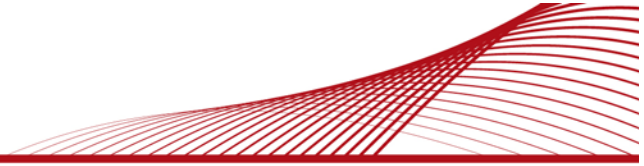

DB2 10 More New Application Features

- SQL stored procedure enhancements
  - SQL PL in Scalar UDFs & XML support
- 64-bit ODBC – also DB2 9 PK83072
- EXTENDED INDICATOR VARIABLES to indicate value not supplied or default
- DRDA support of Unicode for system code points
- CURRENT EXPLAIN MODE special register
- Allow caching of dynamic SQL statements with literals
- Data-dependent paging
  - When only a specific part of the result set is needed
  - Efficient access to desired portions of result set, based upon current position

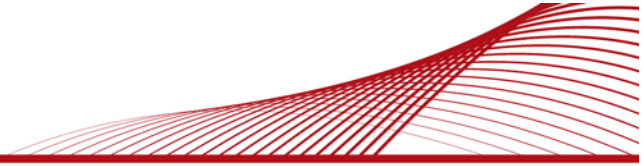

pureXML Performance and Usability Improvements

- XML schema validation in the engine for improved usability and performance
- Binary XML exchange format improves performance
- XML multi-versioning for more robust XML queries
- Allow easy update of sub-parts of XML document
- Stored proc, UDF, Trigger enhanced support
- XML index matching with date/timestamp
- CHECK utility checks XML

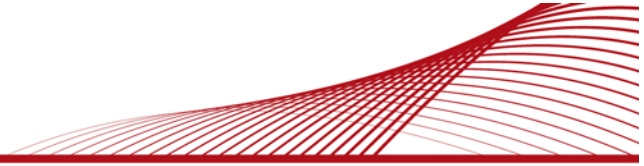

DB2 10 Query Enhancements

- CPU time reductions for queries, batch, & transactions
- SQL enhancements: Moving Sum, Moving Average, temporal, timestamp, implicit cast, SQL PL, …
- pureXML improvements
- Access improvements: Index include columns, hash, index list prefetch, workfile spanned records, …
- Optimization techniques
  - Remove parallelism restrictions and more even parallel distribution
    - increased zIIP use
  - In-memory techniques for faster query performance
  - Access path stability and control
- Analysis: instrumentation, Data Studio & Optim Query Tuner
- Advanced query acceleration techniques
  - IBM Smart Analytics Optimizer

Technology Preview: IBM Smart Analytics Optimizer

What is it?

- A special purpose, network-attached appliance that is an add-on to a DB2 for z/OS system
- Offloads typical DW/BI queries resulting in predictable and orders-of-magnitude faster query response times while reducing overall TCO

Business Value

- Dramatically lowers the cost for query and reporting on System z
- Advanced in-memory scale-out cluster technologies that keep the complete system centrally managed without having to change any requirements for BI applications
- Complements the many new Data Warehousing features in DB2 9 for z/OS
- Leverages the many new warehousing and business intelligence solutions now available on System z

Targeted Uses for DB2 for z/OS customers:

- Requirements to accelerate a subset of their warehouse or reporting queries
- Looking for more insight and business intelligence from operational data
- Needs to consolidate datamarts or data stores into one enterprise warehouse

DB2 10 for z/OS: Skip-Level Migration

May move from V8 to DB2 10,

but just because you can, doesn't mean you always should….

Data sharing mixed release coexistence fully supported
V8/10 or V9/10

Key considerations:

- Risk/reward analysis
  - What's the risk? Tolerance level?
  - How will you do it? What's your mitigation plan? Are ISVs ready?
  - What workloads do you need to test and can you test them properly?
  - Am I missing out on DB2 9 value in the meantime?
- Migration cost savings is not 2x versus two migrations
  - Migration considerations for two releases still apply
  - Larger migration project, longer migration timeline
  - Applications and ISVs may not be ready

Key Details About DB2 10: Getting Ready

Prerequisites: migrate from DB2 9 for z/OS or DB2 for z/OS V8

- z/OS V1.10 SMS-managed DB2-managed DB2 catalog
- System z10, z9, z890, z990, and above (no z800, z900)
- DB2 Connect 9 FP1, 9.7 FP3 for many 10 functions, FP2 beta
- IMS 10 & 11 (not 9) CICS compilers (See announcement)
- Info APARs for migration II14477 (9), II14474 (8)
- SPE PK56922 PK69411 PK61766 PK85956 PM04680 PK87280 PK87281 PM08102 PM08105
- Premigration check DSNTIJPA PM04968

Items deprecated in earlier versions eliminated: more for V8 mig.

- Private protocol → DRDA (DSNTP2DP, PK92339, PK64045)
- Old plans and packages V5 or before → REBIND
- Plans containing DBRMs → packages PK62876 PK79925 (V8)
- ACQUIRE(ALLOCATE) → ACQUIRE(USE)
- Old plan table formats → DB2 V8 or 9, Unicode, 59 cols PK85068
- BookManager use for DB2 publications → Info Center, pdf

DB2 Sort for z/OS v1.0 – Announced May 11th

- Provides high speed utility sort processing for DB2 for z/OS data
- Using DB2 Sort V1.0* for utility sort processing:
  - Up to 30% in reduction of elapsed time
  - Up to 50% reduction of CPU
- Relief from constraints for applications that have:
  - Large volumes of data:
  - High-transaction workloads
  - Lots of insert, update and delete operations
- Continued commitment from IBM to deliver DB2 solutions to provide the highest level of ROI

****Customer results may vary. Results based on analysis done at SVL lab***

DB2 10 for z/OS on the Web

- http://www.ibm.com/software/data/db2/zos/db2-10/

Jeff Josten
josten@us.ibm.com