IBM Information On Demand 2010

Create Insight. Transform. Go Beyond.

DB2 10 for z/OS Technical Overview

John J. Campbell

Distinguished Engineer

DB2 for z/OS Development

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.

DB2 10 for z/OS

- The next release of DB2 for z/OS
- Satisfies major technical requirements across all of the themes
- Major focus areas include:
  - Price/performance
  - Scalability
  - Catalog contention reduction
  - DBA productivity
  - New SQL functionality
  - Query performance and manageability
  - Skip release migration

DB2 for z/OS Technical Strategy

- Application Development
  - Applications can easily connect to DB2 from anywhere
  - Advanced SQL, XML capability, application portability
- 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 engine
- Improved data warehousing capabilities

DB2 for z/OS Into the Future

Delivering Customer Value

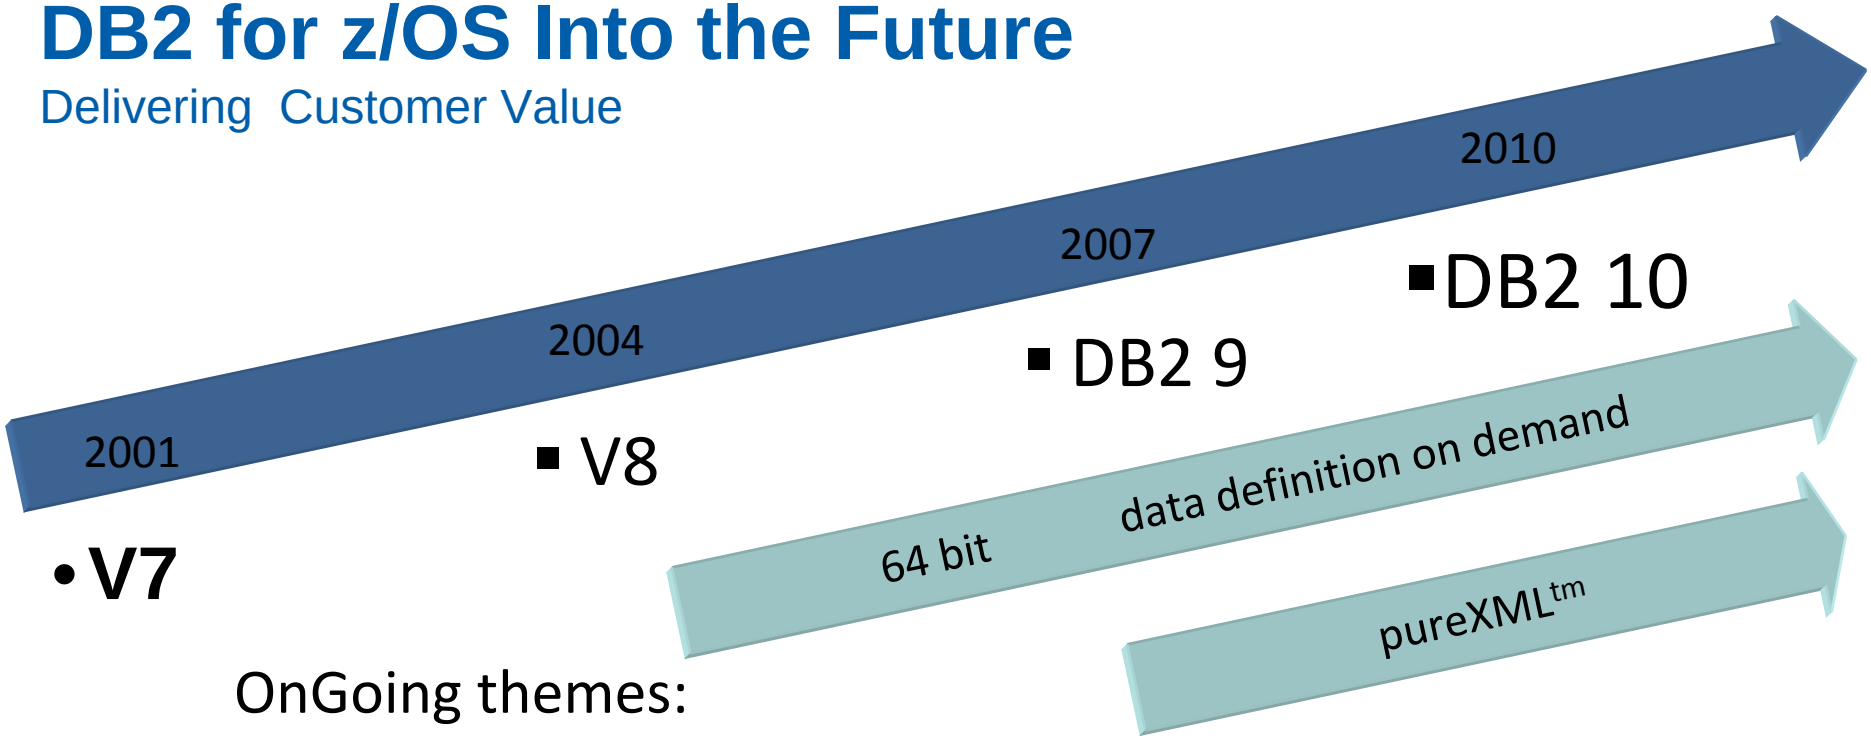

OnGoing themes:

Performance Scalability
Reliability Availability Serviceability
Security Productivity
Application Development
SQL XML SOA

DB2 10 for z/OS At a Glance

Addressing Corporate Data Goals

Application Enablement	• pureXML enhancements • Temporal queries • Last Committed reads • Timestamp with timezone • SQL improvements that simplify porting
RAS, Performance, Scalability, Security	• Wide range of performance improvements • More online schema changes • Catalog restructure for improved concurrency • Row and Column access control • Hash access to data • New DBA privileges with finer granularity
Simplification, Reduced TCO	• Full 64-bit SQL runtime (5x – 10x more threads) • Auto stats • Data compression on the fly • Query stability enhancements • Reduced need for REORG • Utilities enhancements
Dynamic Warehousing	• Moving sum, moving average • Many query optimization improvements • Query parallelism improvements • Advanced query acceleration

Application Enablement, Portability

- Allow non-NULL default values for inline LOBs
- Loading and unloading tables with LOBs
  - LOBs in input/output files with other non-LOB data
- ‘Last committed’ locking semantics
- Implicit casting
- Timestamp with timezone
- Greater timestamp precision

Application Enablement, Portability …

- SQLPL in Scalar UDFs
- 64-bit ODBC Support (APAR PK83072 for DB2 9)
- Special null indicator to indicate value not supplied or default
- DRDA support of Unicode for system code points
- Instance based statement hints
- Allow caching of dynamic SQL statements with literals
- Improved efficient access for “SQL paging”

pureXML Enhancements

- XML schema validation in the engine for improved usability, performance
- Binary XML exchange format for improved performance
- XML multi-versioning for more robust XML queries
- Allow easy update of sub-parts of an XML document
- Stored proc, UDF enhanced support for XML

Temporal Data - Summary

- Business Time (Effective Dates, Valid Time)
  - Every row has a pair of time stamps set by Application
    - Start time: when the business deems the row valid
    - End Time: when the business deems row validity ends
  - Query over current, any prior, present or future period in business time
  - Useful for tracking of business events over time, app logic greatly simplified
- System Time (Assertion Dates, Knowledge Dates, Transaction Time)
  - Every row has another pair of time stamps set by DBMS
    - Start time: when the row was inserted in the DBMS
    - End Time: when the row was modified/deleted
    - Modified rows start time is the modification time
  - Query at current or any prior period in system time
  - Useful for auditing, compliance
- Bi-temporal
  - Inclusion of both System Time and Business Time in row

Current and History

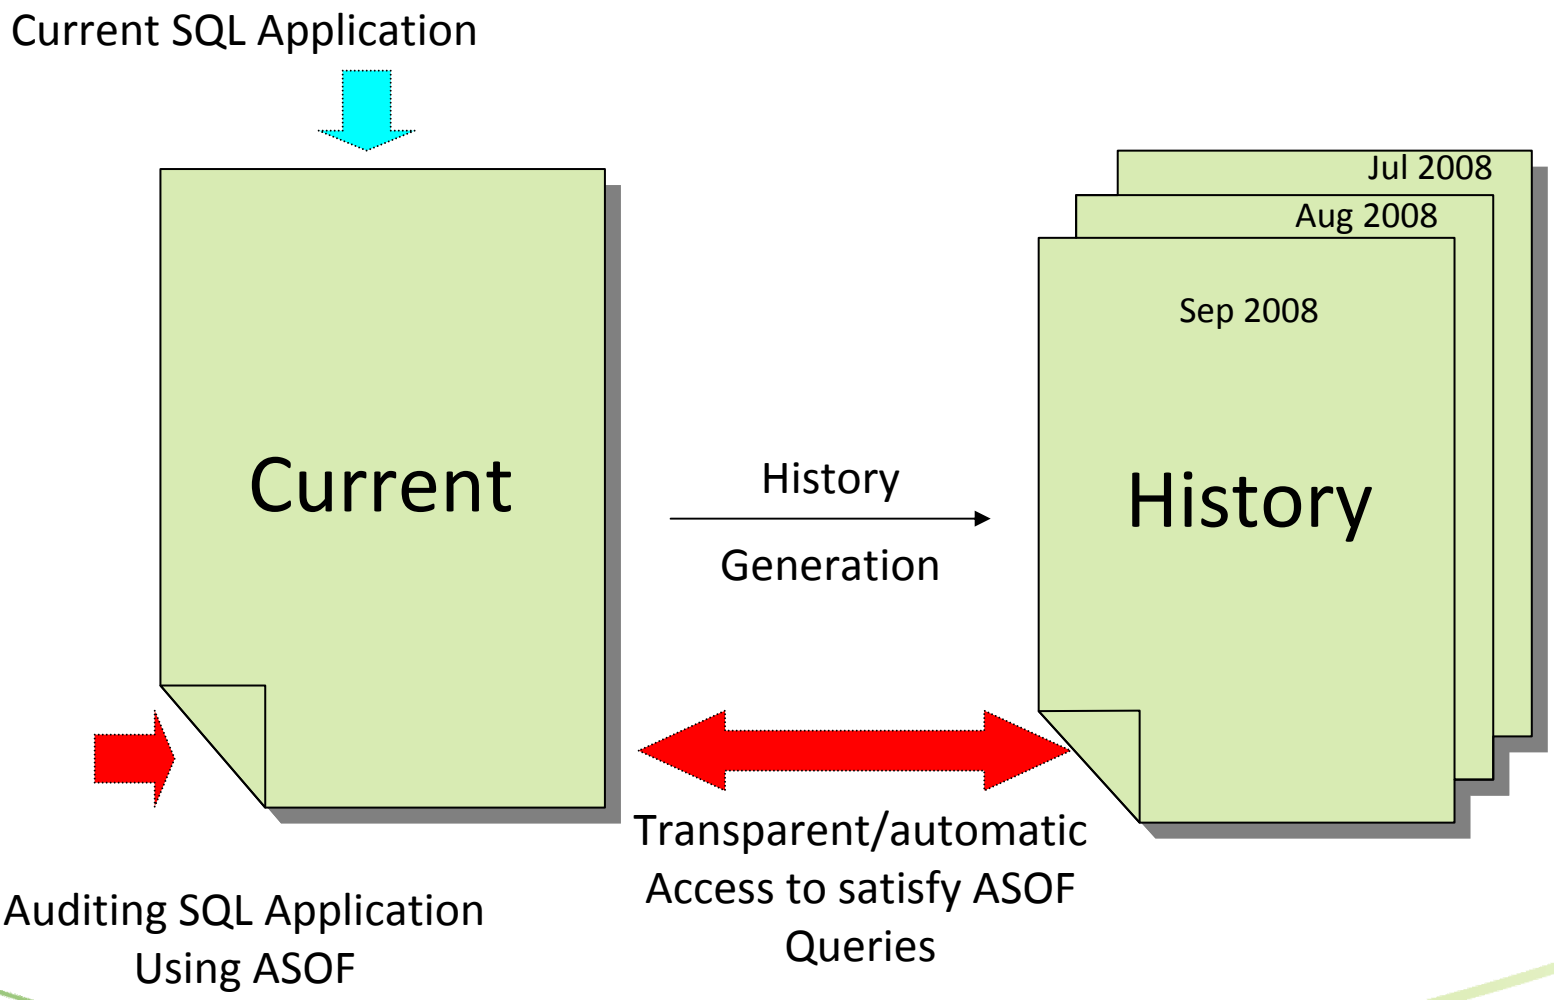

Temporal UPDATE example (business time)

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)

Availability

- More online schema changes for tablespaces, tables and indexes
  Online REORG instead of DROP/CREATE or REBUILD INDEX
  Alterations are manifested with REORG, unless noted otherwise
  - Page size for table spaces and indexes
  - DSSIZE for table spaces
  - SEGSIZE
  - MEMBER CLUSTER
  - Convert single table segmented into UTS PBG
  - Convert single table simple into UTS PBG
  - Convert classic partitioned tablespace into UTS PBR
  - Convert UTS PBR to UTS PBG
  - Convert PBG to hash (immediate, but RBDP index)
  - Ability to drop pending changes
- Online REORG for LOBs, other Online REORG / utility improvements
- Online add active log

Performance, Scalability Objectives

- Provide significant Scalability and Performance improvements
  - Will be an important “feature” for DB2 10
  - Synergistic operation with latest System z hardware
    - High n-way scalability
    - Large real memory exploitation
    - Hardware level optimization
  - Improve transaction times
  - Lower CPU usage for both large and small DB2 subsystems
- Virtual storage is most common constraint for large customers
  - Can limit the number of concurrent threads for a single member/subsystem
- Increasing the number of concurrent threads will expose the next tier of constraints, which should also be addressed

64 bit Evolution (VSCR)

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

- DB2 9 helped (~ 10% – 15%)
- DB2 10 expects to move 90%
  - More concurrent work
  - Reduce need to monitor
  - Consolidate LPARs
  - Reduced cost
  - Easier to manage
  - Easier to grow

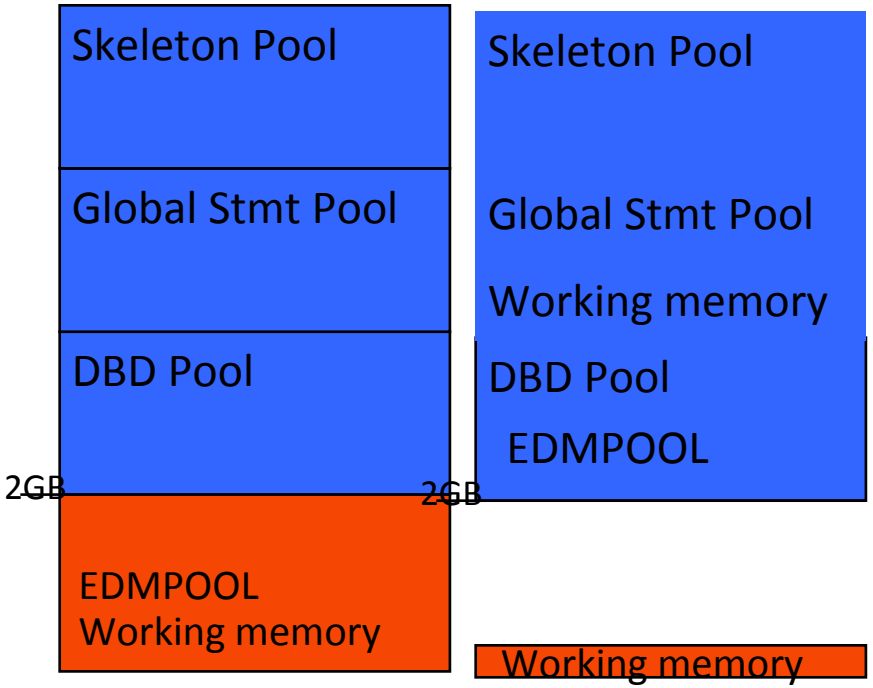

Running a Large Number of Threads

Today

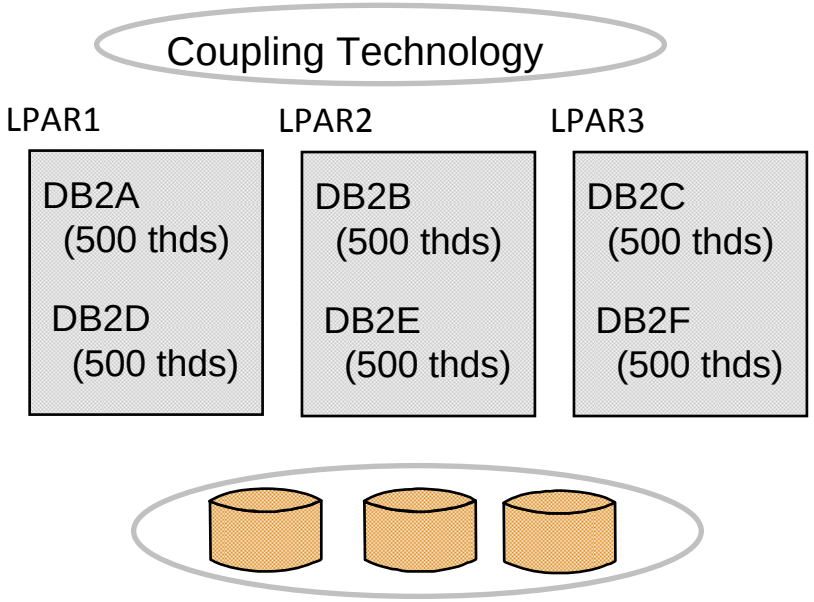

- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s / LPAR

DB2 10

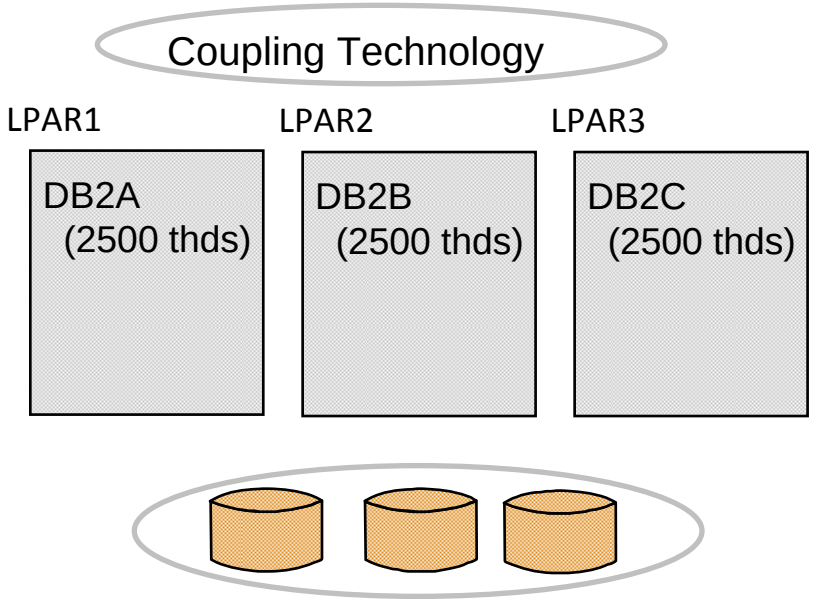

- More threads per DB2 image
- More efficient use of large n-ways
- SSI constraints are relieved
- Easier growth, lower costs, easier management
- Data sharing required for continuous availability and XXXL scale

Other System Scaling Improvements

- Other bottlenecks can emerge in extremely heavy workloads
  - Several improvements planned to reduce latching and other ystem serialization contention
  - New option to for readers to avoid waiting for updaters
  - Eliminate UTSERIAL lock contention for utilities
  - Exploitation of 64-bit common storage to avoid ECSA constraints
- Concurrent DDL/BIND/Prepare processes can hit contention with one another
  - Restructure parts of the DB2 catalog to avoid the contention
- SPT01 64GB limit can be a constraint, especially if “plan stability” support is enabled
  - Relieve 64GB limit for SPT01

Performance

- Hash access path
- Parallel index update at insert
- Faster single row retrievals
- Inline LOBs
- LOB streaming between DDF and rest of DB2
  - Faster fetch and insert, lower virtual storage consumption
- DEFINE NO for LOBs (and XML)
- Enabling MEMBER CLUSTER for UTS
- Efficient caching of dynamic SQL statements with literals

Performance …

- Buffer pool enhancements
  - Utilize z10 1MB page size
  - “Fully in memory” option
- Internal performance optimizations
  - Improved cpu cache performance
  - Exploit new h/w instructions
  - Streamlined DDF, RDS, DM, Index Mgr. performance-critical paths
- Exploitation of SSD

Query Performance and Manageability

- Safe query optimization: assess “reliability” of access path choices
- More Access path stability
- IN list performance
- RID pool overflow to workfiles
- Index include columns
- Workfile spanned records, PBG support, and in-memory enhancements
- Auto Stats
- Instance based statement hints
- Single index access for complex OR predicates*
  - commonly used for cursor scrolling
- Query parallelism improvements*
- Index list prefetch to reduce need for index REORG

Optimization Stability and Control

- Provide an unprecedented level of stability of query performance achieved by stabilizing access paths:
  - 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 applications
  - Per-statement BIND options

Business Security & Compliance Needs

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

- "As of" query, temporal or versioned data
- Fine grained access control
  - Allow masking of value
  - Restrict user access to individual cells

Use disk encryption

Autonomics and DBA Productivity

- Auto statistics collection
- Compress ‘on the fly’
  - Avoid need to run utility
- Timeout / deadlock diagnostics:
  - Identify SQL statements
- Automatic config of IBM supplied UDFs and SPs
- Access path stability
- Reduced need for REORG
  - Build compression dictionary on the fly
  - Index list prefetch enhancements
- Allow tailored names for DSNHDECP

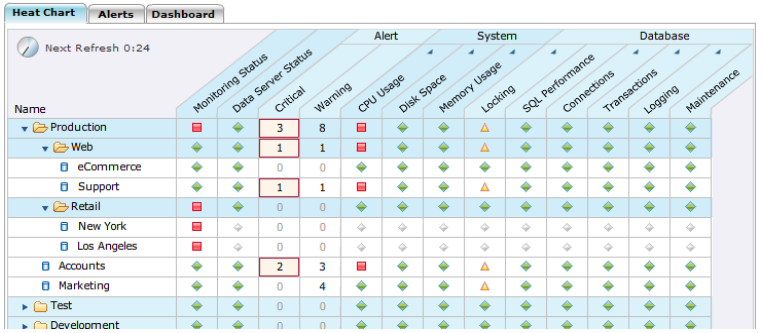

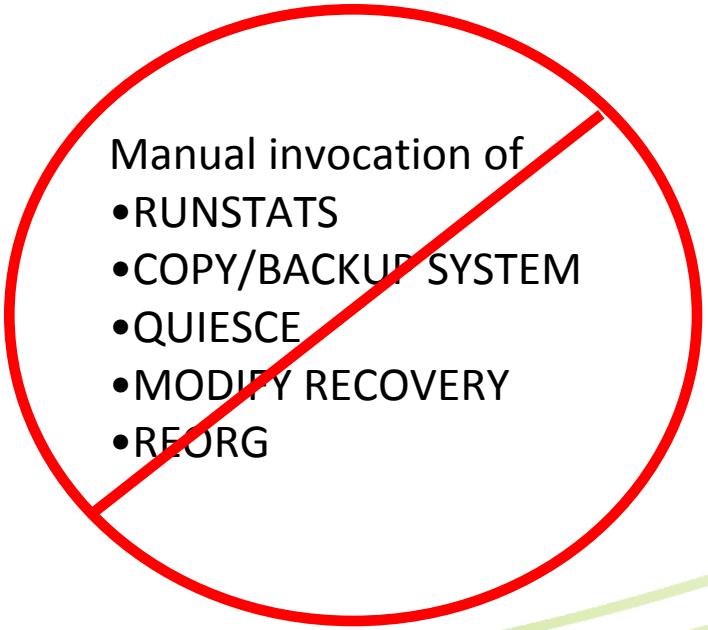

Autonomics and DBA Productivity …

- Checkpoint intervals based on both time and # log records
- Run ‘must complete’ backout under pre-emptable SRB
- Identify unused packages

Utilities Enhancements

- REORG SHRLEVEL(CHANGE) for LOBs
- Online REORG enhancements
  - SHRLEVEL(CHANGE) support for all catalog/directory objects
  - Option to cancel blocking threads
  - Faster SWITCH phase
  - Allow disjoint partition ranges
  - Permit movement of rows between partitions when LOB columns exist
    - Allows REBALANCE or shrinking of PBG even though LOB columns exist
    - Allows DISCARD to delete associated LOB values
  - Messages to estimate length of REORG phases and time to completion

Utilities Enhancements …

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

Data Warehousing

- Moving Sum, Moving Average
- Enhanced query parallelism technology for improved performance
  - Remove query parallelism restrictions
- In-memory techniques for faster query performance
- Advanced query acceleration techniques

Key details about DB2 10

- CM, ENFM, NFM modes lanned
- Prerequisites
  - z/OS V1.10
  - SMS managed, DB2 managed for DB2 catalog
  - DB2 9 for z/OS in NFM or DB2 V8 for z/OS in NFM
  - z890, z990, z9 and above (no z800, z900)
- Eliminated:
  - Private protocol → DRDA (new help in DSNTP2DP)
  - Old plans and packages V5 or before → REBIND
  - Plans containing DBRMs → packages
  - ACQUIRE(ALLOCATE) → ACQUIRE(USE)
  - XML Extender → XML type
  - DB2 MQ XML user-defined functions and stored procedures → XML functions
  - DB2 Management Clients feature (DB2 Administration Server, Control Center, & Development Center) → IBM Data Studio application & administration services
  - msys for Setup DB2 Customization Center → install panels
  - BookManager use for DB2 publications → Info Center, pdf

DB2 10 for z/OS Summary

- Major new release of DB2 for z/OS
- Satisfies major technical requirements
- Improved price/performance
- 64-bit run time for scalability
- Query performance and manageability
- Skip release migration opportunity