DB2 X for z/OS
Technical Update

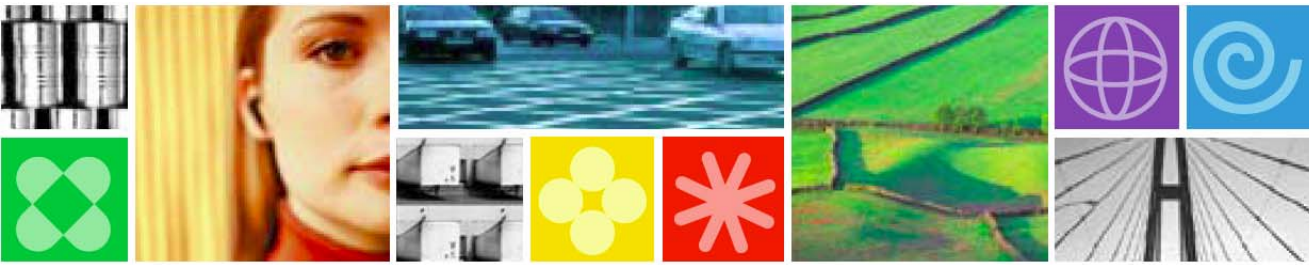

Standard disclaimer

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.

DB2 X for z/OS

- The next release of DB2 for z/OS
- Balanced release, satisfies major technical requirement across all major technology roadmap themes

DB2 for z/OS Technical Strategy

- Application enablement
  - Apps 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 engines
- Improved data warehousing capabilities

DB2 X Major Focus Areas

- Performance, reduced MIPS consumption
- Scalability, more concurrent active threads
- Continuous availability
- Catalog contention reduction
- DBA/sysprog productivity
- Advanced application functionality
- Security advancements

DB2 for z/OS Into the Future

Delivering Customer Value

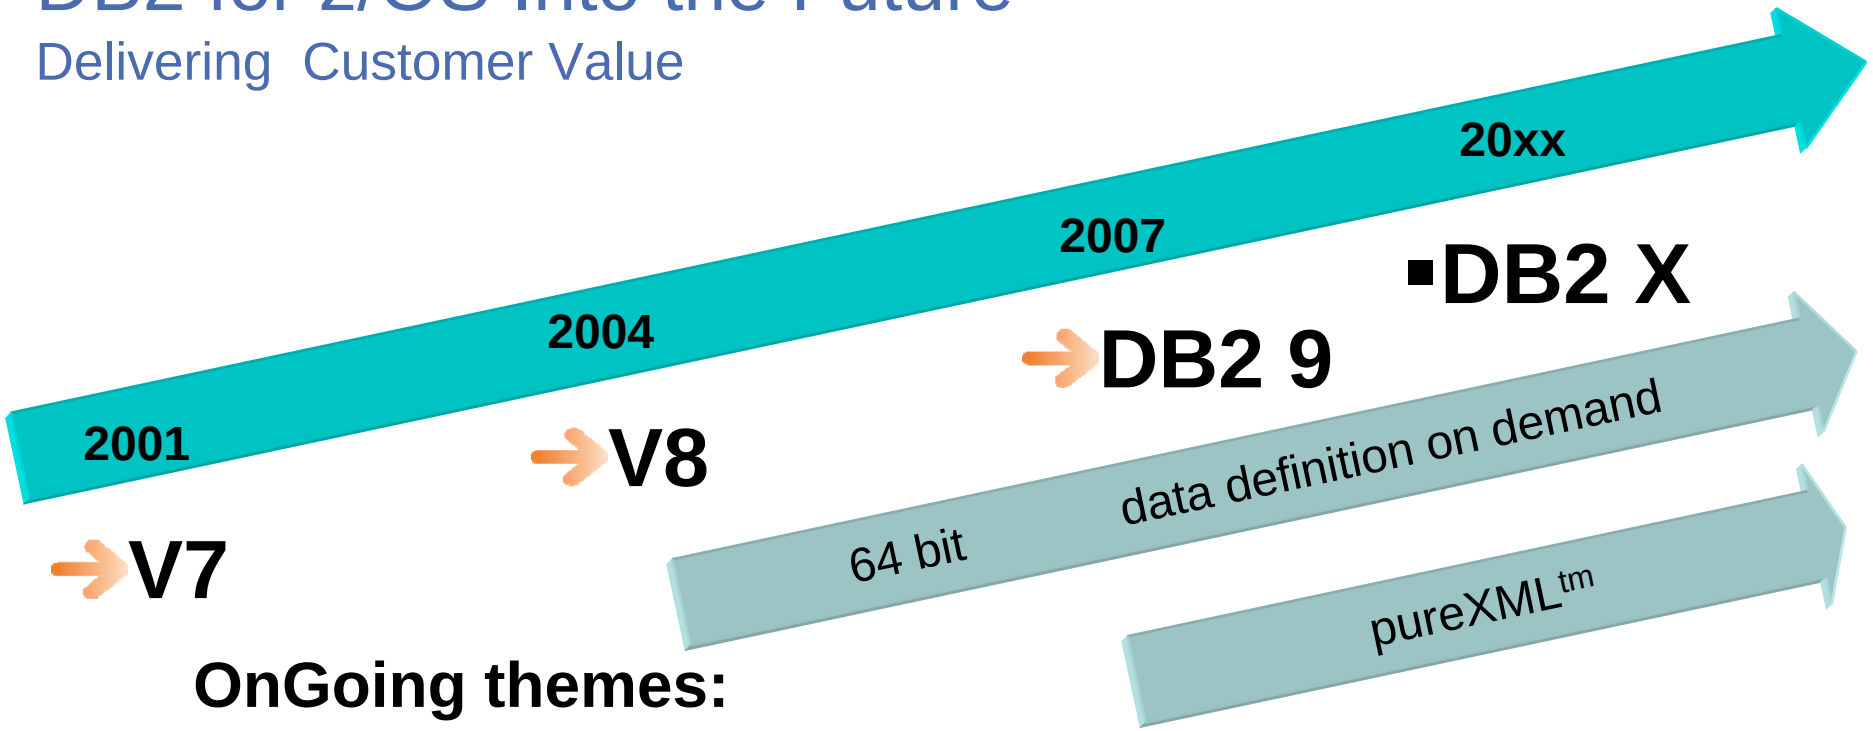

OnGoing themes:

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

DB2 X for z/OS Status

- The following slides represent DB2 Development's current thinking on some of the items that are candidates for DB2 X
- DB2 X is still in the development process, so details will change
- The intention is to give you some information on DB2's future technical directions
- DB2 Development values customer feedback

DB2 X 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 • Fine grained access control • Hash access to data • New DBA privileges with finer granularity
Simplification, Reduced TCO	• Full 64-bit SQL runtime • Auto stats • Data compression on the fly • Query stability enhancements • Reduced need for REORG • Utilities enhancements
Data Warehousing	• Moving sum, moving average • Many query optimization improvements • Query parallelism improvements • Advanced query acceleration

Performance

DB2 X Performance, Scalability Objectives

- Significant scalability and performance improvements
  - 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

DB2 X Performance Objectives

- Historical goal of <5 % version-to-version performance regression
- Goal of 5% -10% performance improvement for DB2 X
- More improvements will be possible through exploitation of new features

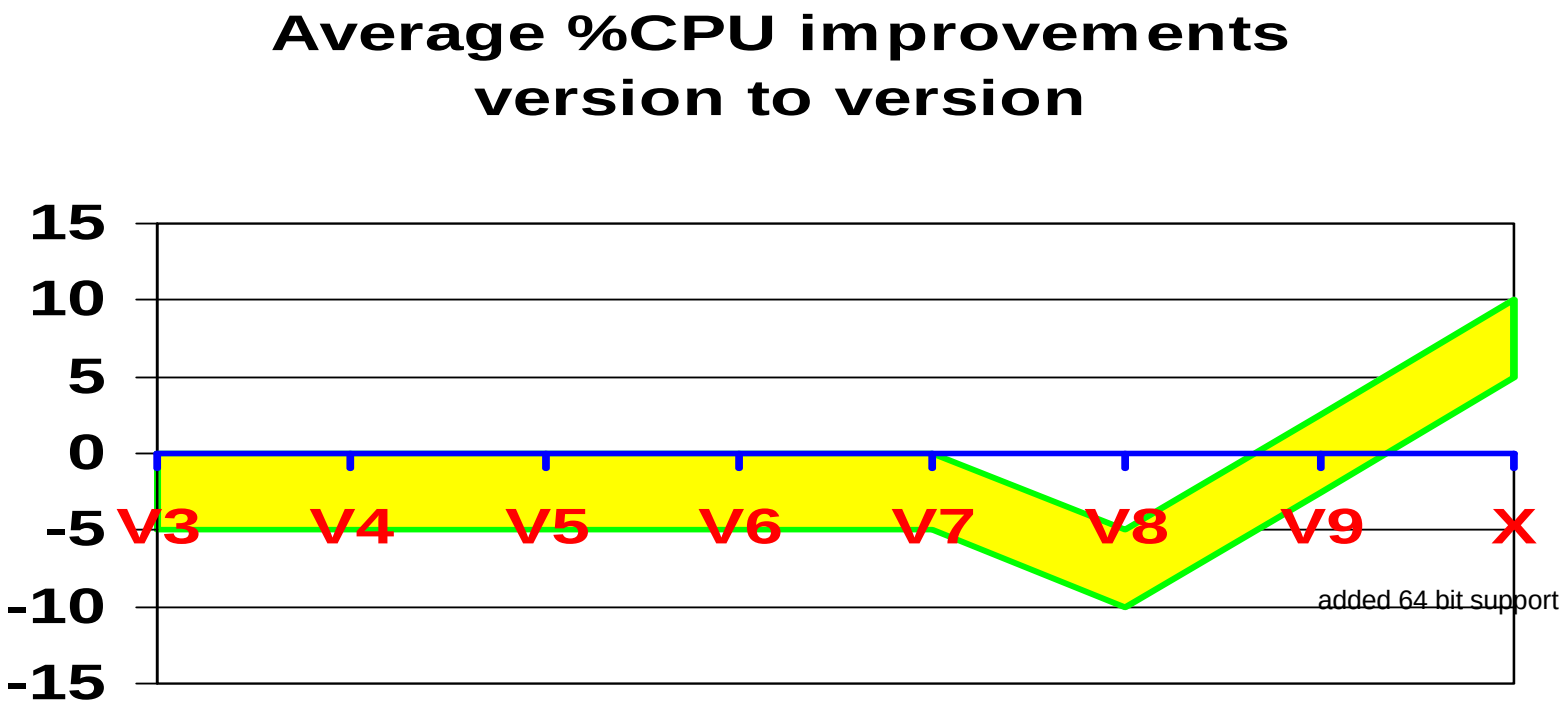

Performance Enhancements Requiring No Changes (“free”)

- SQL runtime improved efficiency
- Parallel index update at insert
- Faster single row retrievals
  - Open/fetch/close chaining
- INSERT improvements for UTS
- LOB streaming between DDF and rest of DB2
  - Faster fetch and insert, lower virtual storage consumption
- Workfile spanned records, PBG support, and in-memory enhancements
- Index list prefetch
- SQLPL performance improvements
- High Performance DBATs
- Exploitation of SSD

Performance Enhancements requiring DDL, BIND, etc. changes (but no app changes)

- Inline LOBs
- DEFINE NO for LOBs (and XML)
- MEMBER CLUSTER for UTS
- Efficient caching of dynamic SQL statements with literals
- Buffer pool enhancements
  - Utilize z10 1MB page size
  - “Fully in memory” option
- Hash access path
  - Alter + Reorg + rebind to activate
- Index include columns
  - Alter + Rebuild + rebind to activate

Performance Enhancements requiring REBIND (but no app changes)

- SQL paging performance enhancements
  - Single index access for complex OR predicates:
  - New EXPLAIN access method
- IN list performance
  - Optimized Stage1 processing (single or multiple IN lists)
  - Matching index scan on multiple IN lists
    - New EXPLAIN access method
- Query parallelism improvements
- More stage 2 predicates can be pushed down to stage 1
  - New DSN_FILTER_TABLE info to indicate stage1 or screening
- More aggressive merge of views and table expressions
  - Avoid materialization of views
- REBIND enables further SQL runtime improvements

Scalability

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

- DB2 X expects to move 80-90% of thread storage above the bar
  - More concurrent work
  - Reduce need to monitor
  - Consolidate LPARs
  - Reduced cost
  - Easier to manage
  - Easier to grow

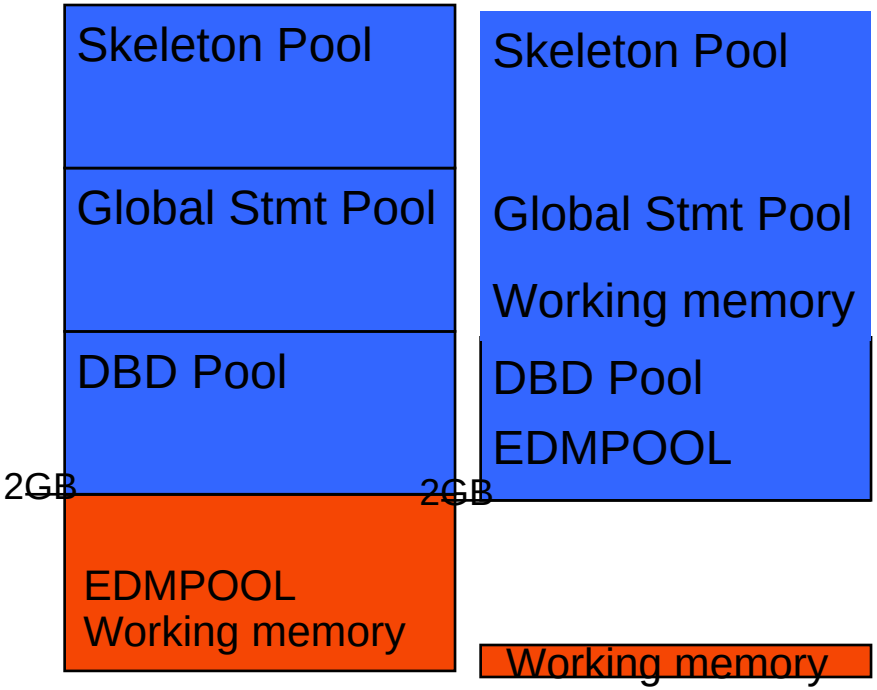

Running a Large Number of Threads

Today

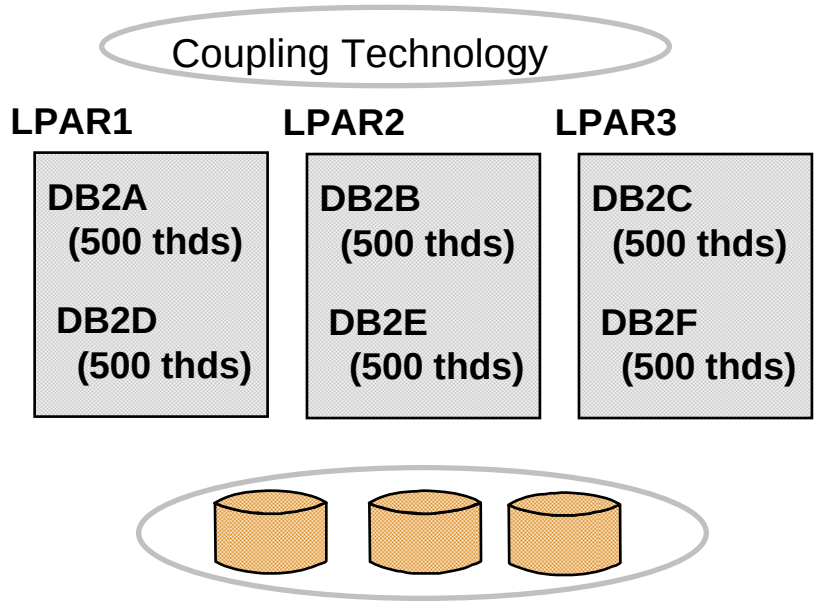

- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s / LPAR

DB2 X

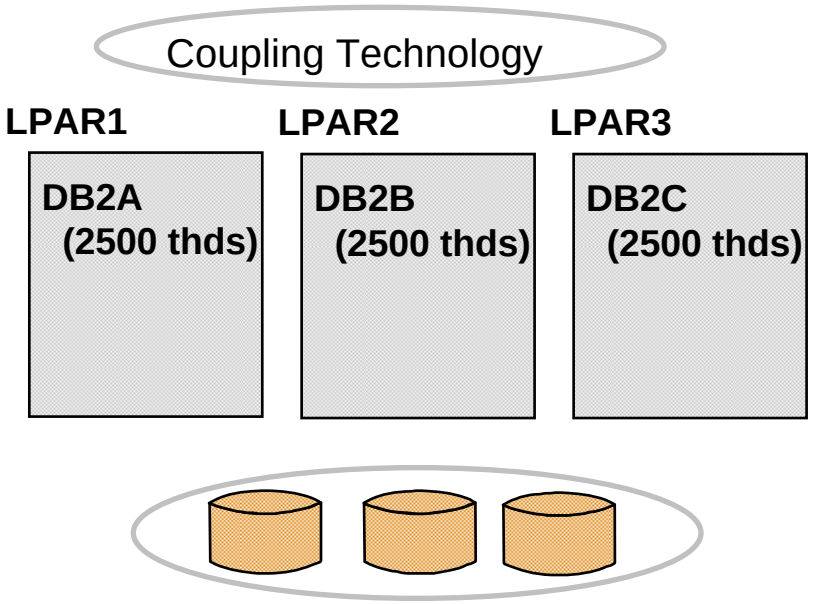

- More threads per DB2 image
- Potential for fewer members / LPARs
- More efficient use of large n-ways
- SSI constraints are relieved
- Easier growth, lower costs, easier management
- Data sharing required for continuous availability and XL scale

Other System Scaling Improvements

- Remove other potential bottlenecks
  - reduce latching and other system serialization contention
  - new option for readers to avoid waiting for updaters
  - eliminate UTSERIAL lock contention for utilities
  - Exploitation of 64-bit common storage to avoid ECSA constraints
- Improved DDL/BIND/Prepare concurrency
  - restructure parts of the DB2 catalog to avoid the contention
- Remove SPT01 64GB limit
- Lower overhead for very large buffer pools

Continuous Availability

Availability

- More online schema changes for table spaces, 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 table space 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

DB2 X 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

DB2 X more utilities enhancements

- Support of spanned records for UNLOAD of LOB data
  - Currently unload of LOBs >32K must use FRVs
  - This allows inlining of LOBs with base row in unload data set
  - Provides portability of data
- Performance enhancement for FRV processing with PDS data sets
  - UNLOAD 33% elapsed time reduction
  - LOAD 84% elapsed time reduction
- Extend support for UTF-16
  - Date, time & timestamp fields currently unloaded in UTF-8
  - Cannot specify a char value for a graphic column in WHEN clause

DB2 X: 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)
  - 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

Reduced Catalog Contention

Catalog Restructure for improved concurrency

- Remove links from the catalog and directory
  - DSN1CHKR no longer needed
- Row level locking enabled for the catalog and directory
- Done during ENFM for migrated systems
- BIND, PREPARE,and DDL will run with better concurrency, fewer timeouts/deadlocks

Other Catalog Changes

- Partition-by-growth (PBG) catalog/directory table spaces
  - Allow SPT01 to grow beyond 64GB
  - V8/V9 APAR PK80375 adds zparm for SPT01 compression
- DB2 managed catalog and directory data sets
  - DFSMS required
  - Eases admin and management burden
- New CLOB columns for storing SQL statements
  - Today, SQL statements can be split into several records with sequence numbers
  - CLOBs will make it easier to query SQL statements
- Convert SYSCOPY from EBCDIC to UNICODE.
- Online REORG enabled for all catalog/directory objects

Advanced Application Functionality

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
- Moving Sum, Moving Average

Application Enablement, Portability …

- SQLPL in Scalar and Table UDFs
- 64-bit ODBC Support
- Special null indicator to indicate value not supplied or default
- Allow caching of dynamic SQL statements with literals

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, Trigger enhanced support for XML
- XML index matching with date/timestamp
- CHECK XML utility

Temporal Data - Summary of Proposal

- Business Time (Effective Dates, Valid Time)
  - Every row has a pair of time stamps set by App
    - 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

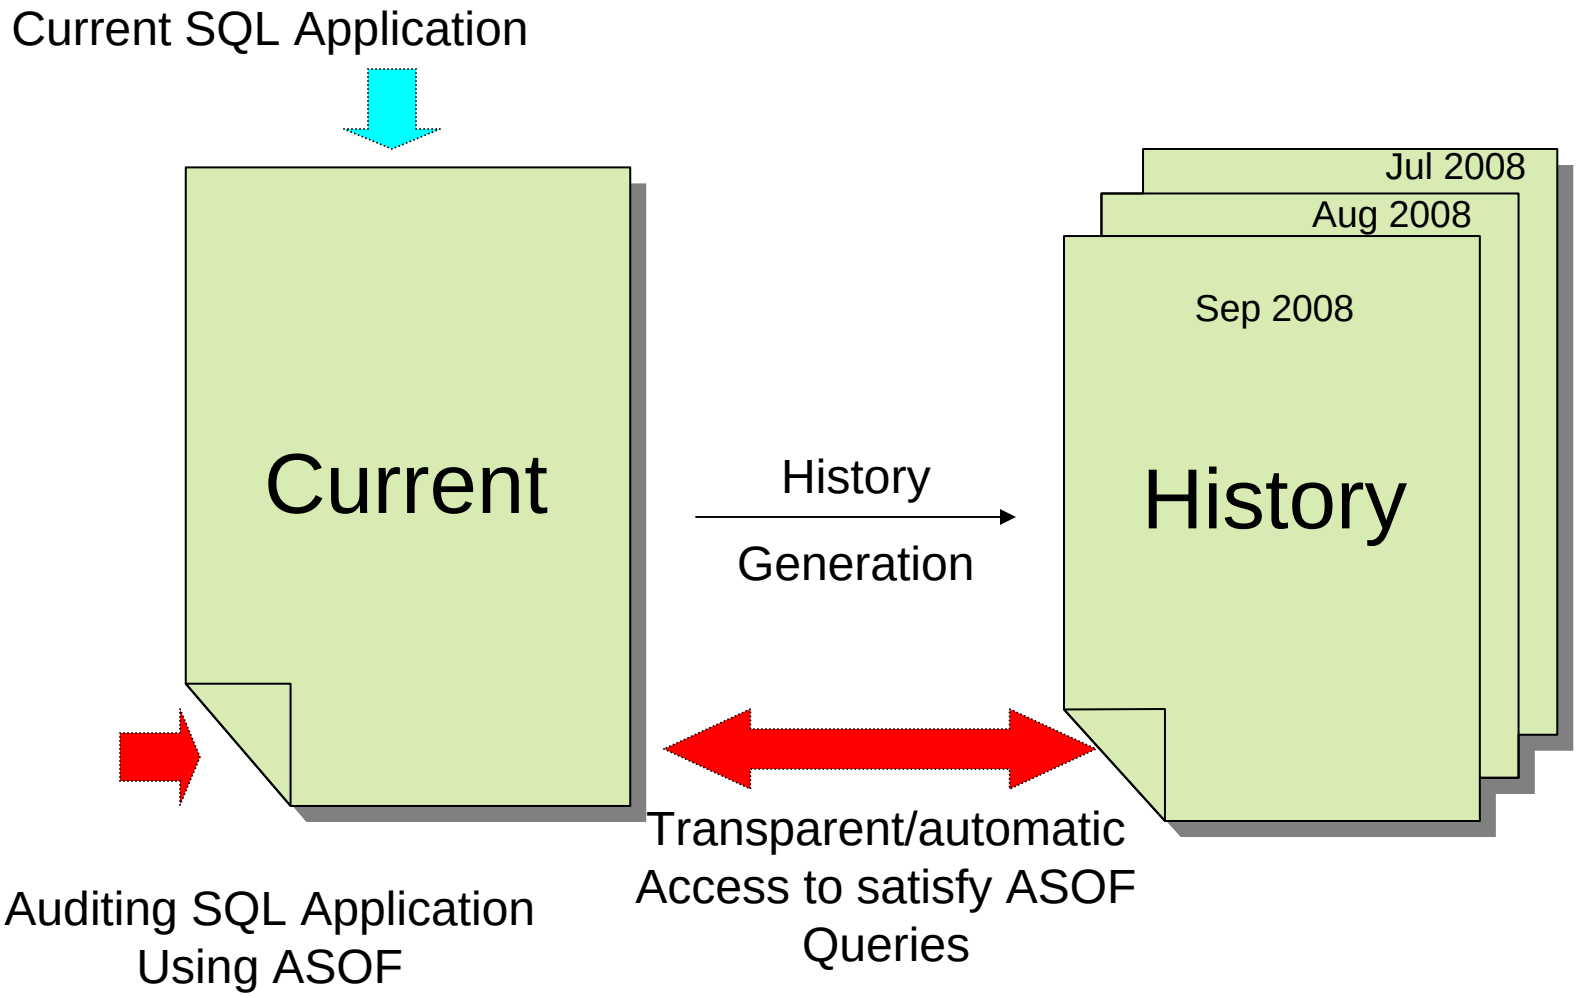

Temporal UPDATE example (business time)

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)

DBA/Sysprog Productivity

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 apps
- Per-statement BIND options

- Safe query optimization: assess “reliability” of access path choices
- RID pool overflow to workfiles

DB2 X: Productivity – Doing More with Less!

- 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

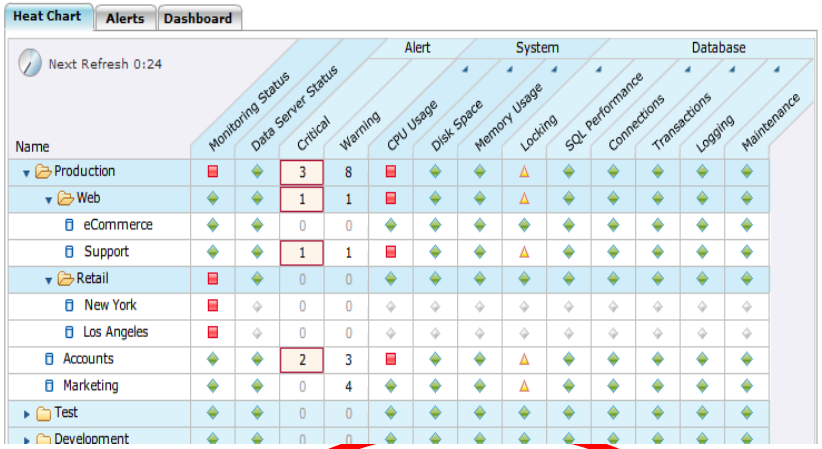

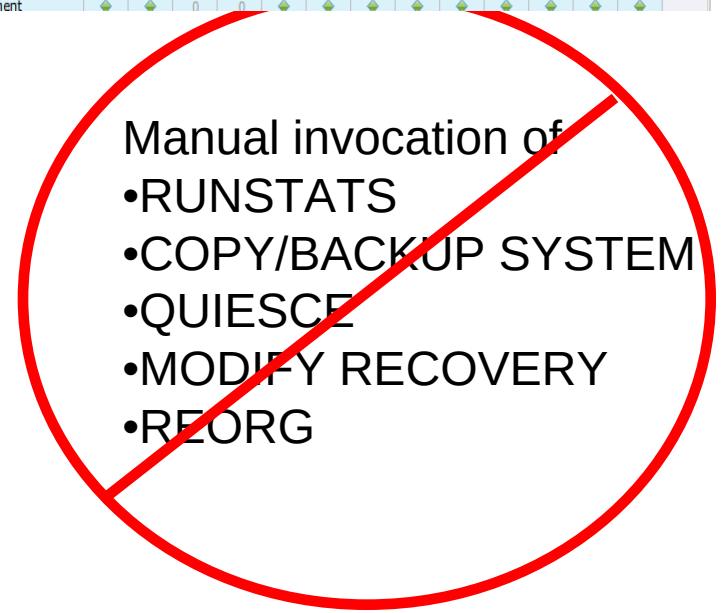

Autonomics and DBA Productivity…

- Checkpoint intervals based on both time and # log records
- Run ‘must complete’ backout under pre-emptable SRB
- Identify unused packages
- 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
- Manage max threads, connections, idle thread timeout on an application basis
  - Warning or exceptions issued when threshold is hit

Security Enhancements

DB2 X: 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

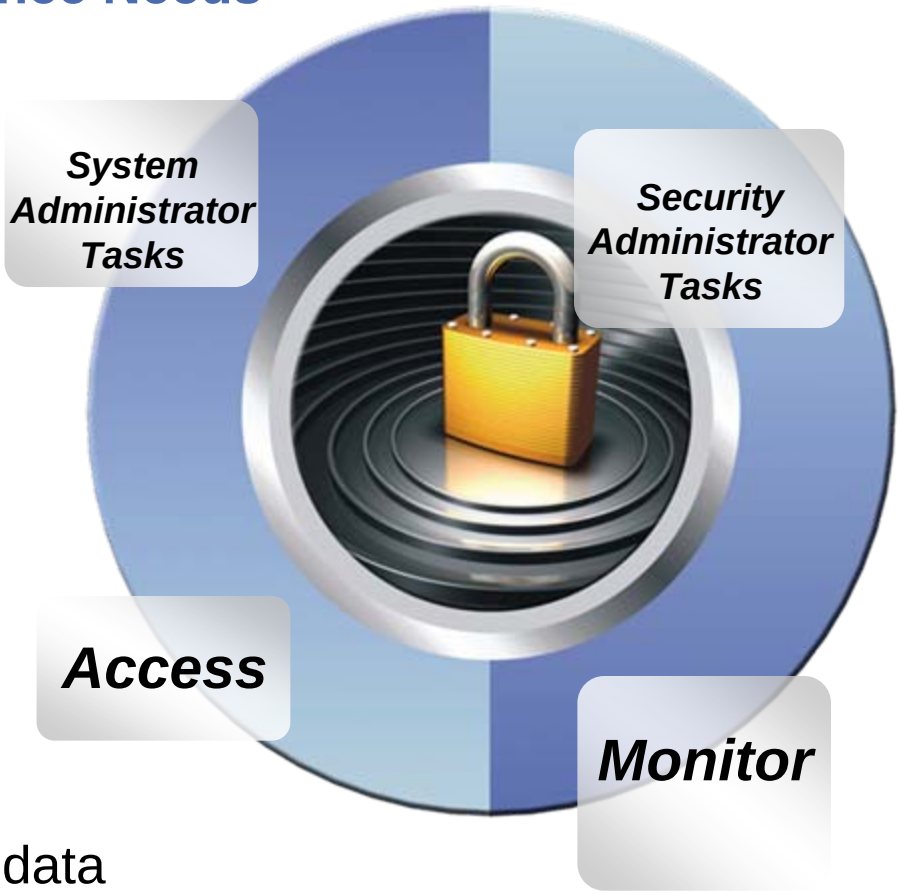

Use disk encryption

Key details about DB2 X

- CM, ENFM, NFM is planned
- Probable Prerequisites
  - z/OS V1.10
  - DFSMS required for DB2 catalog
  - DB2 9 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)
  - Workload capture through profile monitor
  - 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