IBM Software Group

DB2 and DB2 Tools: Version 8 and Beyond

Dan Wardman, IBM Information Management

DB2 Information Management Software

e business on demand software

Agenda

- Mainframe database direction
- DB2 V8 UPDATE
- DB2 and application development
- DB2 Futures
- DB2 Tools directions
- Summary

Mainframe Database Themes

- Provide the most reliable, scalable, and impregnable data base solutions in the industry
- Protect the investment of thousands of companies in IMS and DB2
- Enable applications of all types to exploit the information assets stored in IMS and DB2
- Provide a cost competitive solution
- Enable future generations of users

Greatest Hits: DB2 UDB for z/OS V8

- ✓ High availability
- ✓ Scalability or very large database
- ✓ Java and the web
- ✓ Queries and data warehouses
- ✓ Migrating or porting applications
- ✓ Application packages

SAP PEOPLESOFT SIEBEL

Continuous Availability

- **Online Schema Evolution: database changes with ALTER instead of DROP / CREATE e.g. ADD partition**
- **System-Level Log Point Recovery**
- **Data Partitioned Secondary Indexes**
- **Improved LPL Recovery**
- **Additional online zparms**

Scalability and Very Large Database

- Add partitions
- Separate partitioning & clustering
- Index improvements
- 4096 Partitions
- Rotate partitions
- Extend columns
- Optimization improvements
- Memory and scale increased

64 bit evolution

- zSeries, z/OS, z/Architecture & large real storage
- Real storage for performance
  - V6 data space advantages
    - zSeries & 64 bit addressing
- Version 8 64 bit virtual
  - z/OS 64-bit Virtual Storage Roadmap
  - Improve scalability, availability, ease
  - Move above the bar
  - Hiperpool/dataspace no longer needed, no ECSA for locks

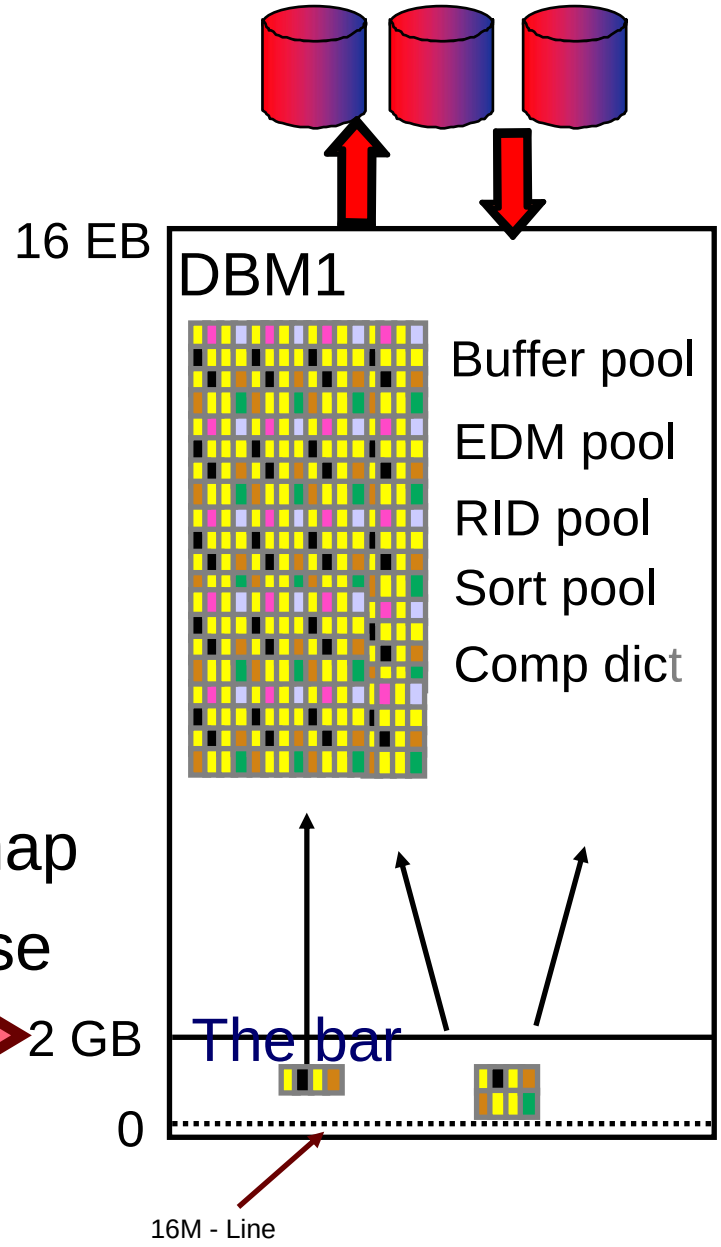

Java and the web

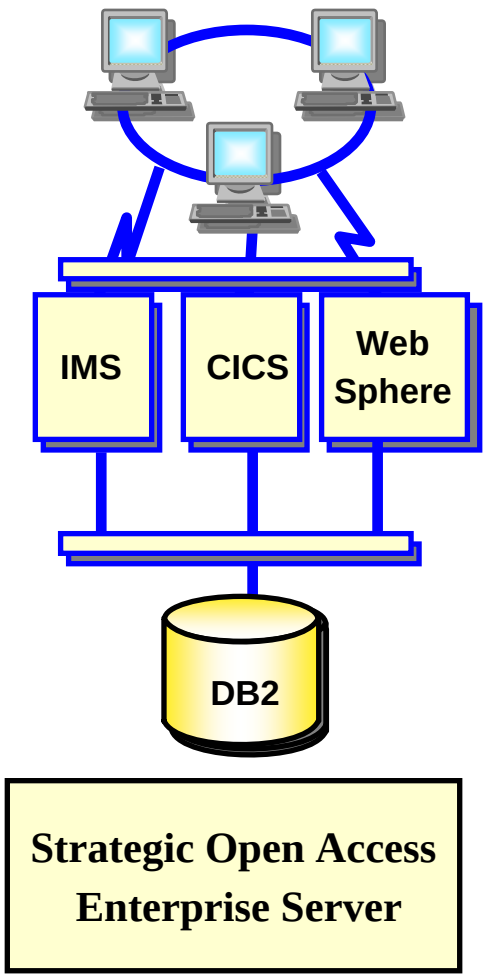

- Application Connectivity for JAVA
- Improve data sharing member routing
- LOCATION aliases at requester & server
- Extended function, standards
- Reduced processing
- Stored Procedure Enhancements
- Java Universal Driver

Key Java, XML and WebSphere Improvements

- Improved function & consistency
  - JDBC / SQLJ 3.0 standard
  - Java Universal Driver used across family
- Unicode improvements
- Complemented by more consistent SQL
- Increased integration with WebSphere
- XML Publishing

z/OS Application Connectivity to DB2 for z/OS & OS/390

- Pure Java type 4 JDBC driver
- From Java application on z/OS
- To DB2 for z/OS server or Linux, Unix & Windows
- JDBC 2.0 and 3.0, JDK V1.4
- Connect to DB2 for z/OS and WebSphere Application Server for z/OS

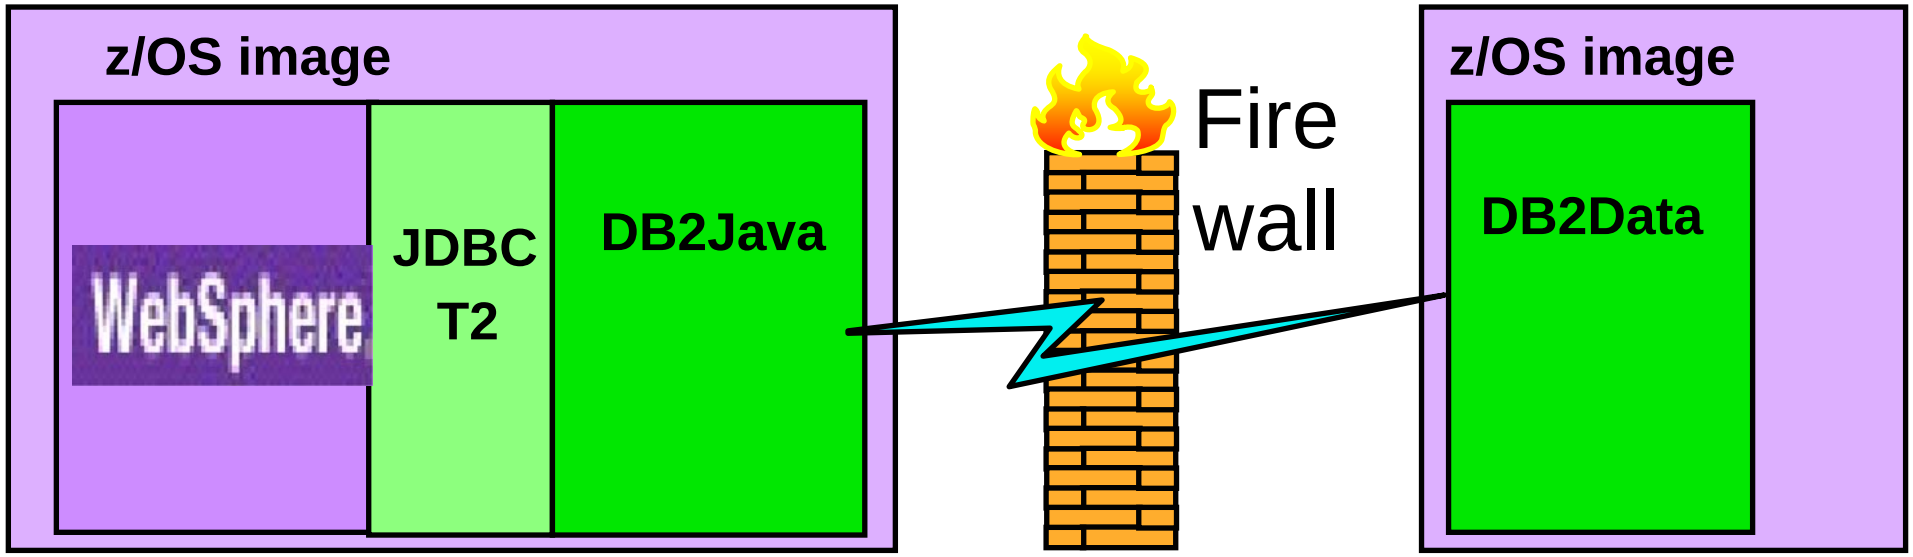

Current configuration

IBM z/OS Application Connectivity to DB2 for z/OS & OS/390

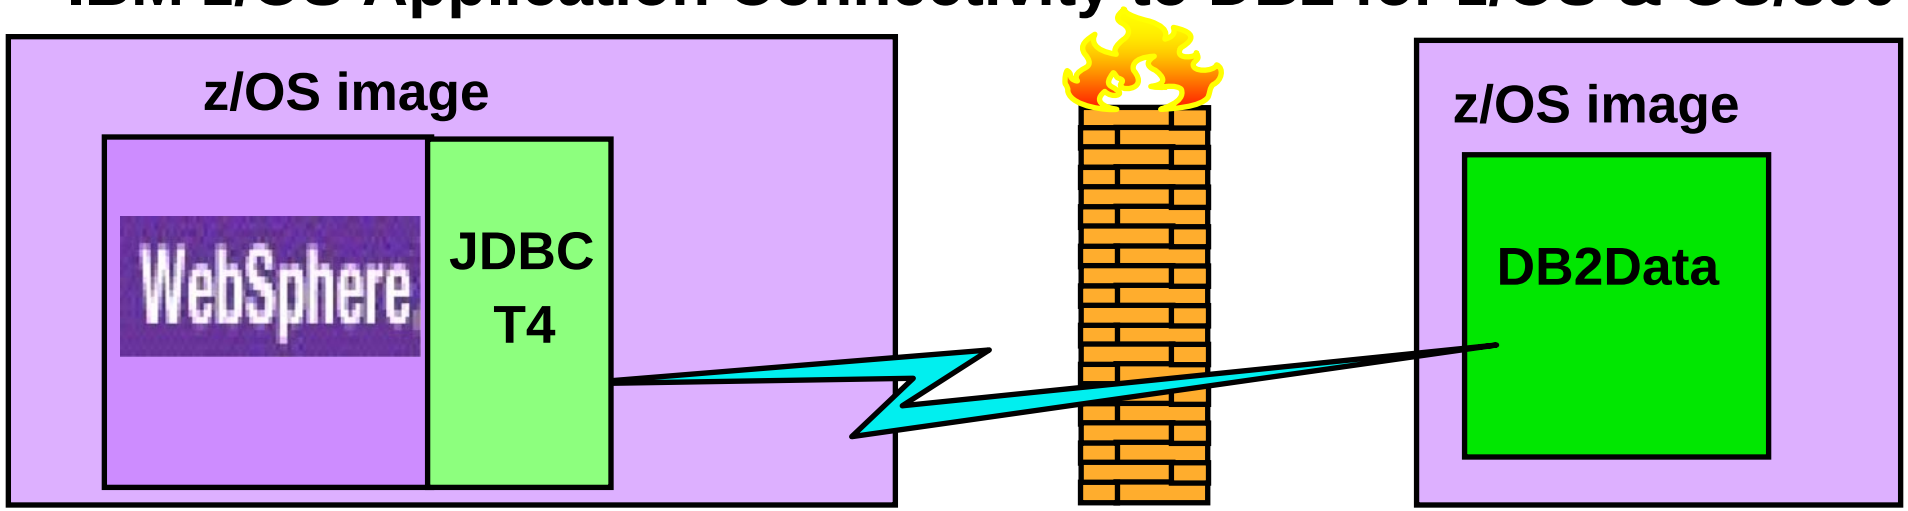

Save CPU, memory, license fees

Key Java Improvements

- Java Universal Driver
  - Updated to support JDBC/SQLJ 3.0 standard
    - savepoint support
    - connection pooling improvements
      - reusing PreparedStatements
      - standard connection pool properties
      - reset of server connection at getConnection
    - new metadata for PreparedStatements
    - return autogenerated keys
    - multiple open ResultSets for a single stored procedure
    - WITH HOLD cursors
    - improved BLOB/CLOB support

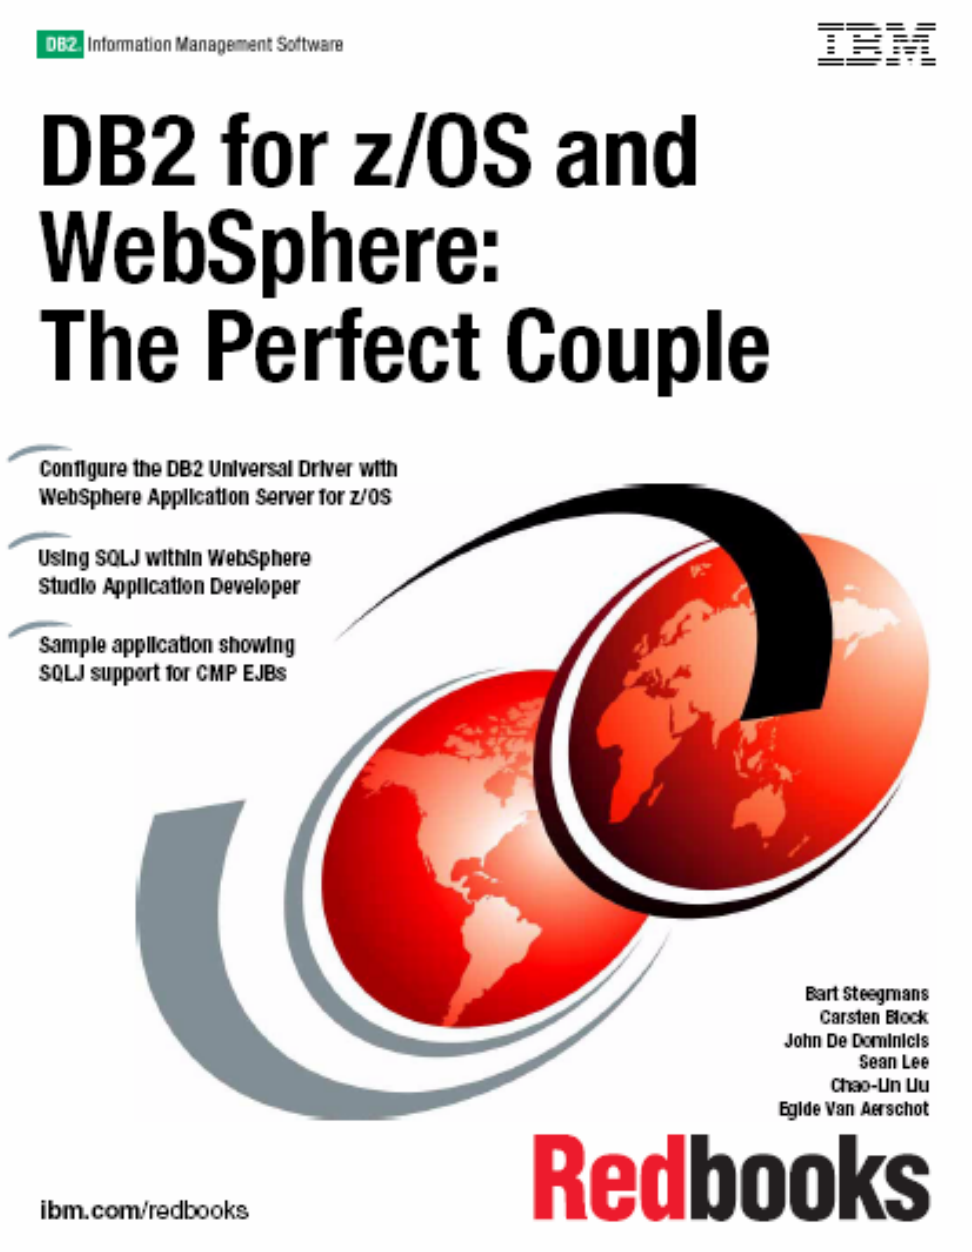

SG24-6319

Environment

Architecture

Application development

Security

Transaction management

Performance

Queries and data warehouses

- Optimization Improvements
  - Improved techniques
  - Enhanced data
  - Visual Explain
- Enhanced index options
- Materialized Query Tables
- New Partitioning options
- QMF improvements
- SQL enhancements

Migrating and porting applications

- Multi-row INSERT, FETCH & UPDATE
- GET DIAGNOSTICS
- INSERT within SELECT
- IDENTITY Column enhancements
- SEQUENCES
- CURRENT PACKAGE PATH
- SQL Procedure Language
- Dynamic Scrollable Cursor, Common Table Expression, Recursion
- Scalar Fullselect
- Materialized Query Table
- UNICODE SQL, Multiple CCSIDs
- XML Publishing
- Long names, long statements...

Enterprise Applications & : DB2 UDB for z/OS

4.6 certified

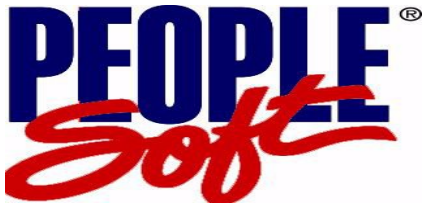

8.45 certified

7.8 certified

- **64 bit virtual storage**
- **Unicode**
- **Schema evolution**
- **System-level backup and recovery**
- **Multi-row fetch & insert**
- **Multiple DISTINCT Clauses**
- **Lock Contention on volatile Tables**
- **Fast Retrieval of Most Recent Value**
- **Longer Table Names & Column Names**
- **Additional statistics**
- **Convert Column Type**
- **Altering CLUSTER Option**
- **Adding Columns to Index**
- **Index-only Access Path for VARCHAR**
- **Adding New Partitions**
- **Separate Clustering from Partitioning**
- **...**

DB2 UDB for z/OS Version 8 is

- **SQL OLTP Leadership**: name lengths, statements, scrolling, expressions, predicates, diagnostics, …
- **Break through limitations**: storage, partitions, log
- **Performance enhancements**: index, materialized query tables, more efficient IO, multi-row
- **Database changes without an outage**: add partition, rotate partition, backup / restore
- **Integration**
  - zSeries, z/OS & ESS platform
  - Middleware stack

Rational.

WebSphere.

DB2®

Lotus®

Tivoli®

Reengineered for e-business on demand

DB2 Technology Themes

- **Enable high-volume transaction processing for next wave of Web applications**
  - **Extend the lead in transaction processing availability, scalability and performance**
  - **Reduce cost of ownership and zSeries-specific skill needs**
  - **Improve data warehousing and OLTP reporting**

DB2 Family SQL

z z/OS V7
common
LUW Linux, Unix & Windows V8.2

z

common

Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, Call from trigger, statement isolation

LUW

Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE, DELETE & MERGE, multi-site join, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT

DB2 Family SQL

z z/OS V8
common
LUW Linux, Unix & Windows V8.2

z Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables

common Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT

LUW Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE, DELETE & MERGE, multi-site join

DB2 Family SQL

z z/OS V8 Vnext
common
LUW Linux, Unix & Windows V8.2

z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE

common

Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE, DELETE & MERGE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect and fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables

LUW

Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE

DB2 for z/OS Vnext SQL, DB2 family & porting

- XML
- MERGE
- SELECT FROM UPDATE, DELETE, MERGE
- TRUNCATE
- INSTEAD OF TRIGGER
- BIGINT, VARBINARY, DECIMAL FLOAT
- Native SQL Procedure Language
- Optimistic locking
- LOB File reference variable & FETCH CONTINUE
- FETCH FIRST & ORDER BY in subselect and fullselect
- INTERSECT & EXCEPT
- ROLE & trusted context
- Many new built-in functions, caseless comparisons
- Index on expression
- Improved DDL consistency
- CURRENT SCHEMA

Native SQL Procedural Language

- Eliminates generated C code and compilation
- Fully integrated into the DB2 engine
- Extensive support for versioning:
  - VERSION keyword on CREATE PROCEDURE
  - CURRENT ROUTINE VERSION special register
  - ALTER ADD VERSION
  - ALTER REPLACE VERSION
  - ALTER ACTIVATE VERSION
- BIND PACKAGE with new DEPLOY keyword

Optimistic Locking Support

- Built-in timestamp for each row or page
  - Automatically updated by DB2
  - Allows simple timestamp predicate to validate that row has not changed since last access
- Eliminates need for complex predicates on WebSphere CMP updates, improves performance

XML Processing Paradigms

XML has become the "data interchange" format between B2B/B2C, inter- and intra-enterprise environments.

XML View Of Relational Data

- SQL data viewed and updated as XML
  - Done via document shredding and composition
- DTD and Schema Validation

XML Documents As Monolithic Entities

- Atomic Storage And Retrieval
- Search Capabilities

XML As A Rich Data Type

- Full Storage and Indexing
- Powerful Query Capabilities

DB2 vNext Themes

- **Enable high-volume transaction processing for next wave of Web applications**
- **Extend the lead in transaction processing availability, scalability and performance**
- **Reduce cost of ownership and zSeries-specific skill needs**
- **Improve reporting**

Schema Evolution – Database Definition On Demand

- Fast replacement of one table with another
- Rename column and index
- Alter index to compress or remove compression, change page size
- Table space that can add partitions, for growth
- Improve ability to rebuild an index online
- Online reorganization with no BUILD2 phase
- Modify early code without requiring an IPL
- Alter table space and index logging

CLONE Tables

- Allows fast replacing production data without renames and rebinds
  - A capability to support online load replace
- CREATE TABLE to create a Clone Table
  - All indexes are also cloned
  - Table and Index data are not copied
  - Base and Clone tables share the same table space and index names
  - Underlying data sets are differentiated by a data set instance number

Partition by Growth

- New partitioning scheme:
  - Single table tablespace, where each partition contains a segmented pageset (allows segmented to increase from 64GB to 16TB or 128 TB with 32K pages)
  - Eliminates need to define partitioning key and assign key ranges
  - A new partition is created when a given partition reaches DSSIZE (defaults to 64G)
  - Retains benefits of Utilities and SQL parallelism optimizations for partitioned tables

LOB Performance/Scalability

- Elimination of LOB locks – LRSN and page latching is used instead for consistency checks
- New network flows for delivering LOBs
  - JDBC, SQLJ, and CLI will let server determine whether to flow LOB values or LOCATORs based on size thresholds
  - Significant reduction in network traffic
  - Greatly reduces frequency of FREE LOCATOR statements

DB2 vNext Themes

- **Enable high-volume transaction processing for next wave of Web applications**
- **Extend the lead in transaction processing availability, scalability and performance**
- **Reduce cost of ownership and zSeries-specific skill needs**
- **Improve reporting**

Cost Of Ownership Trends

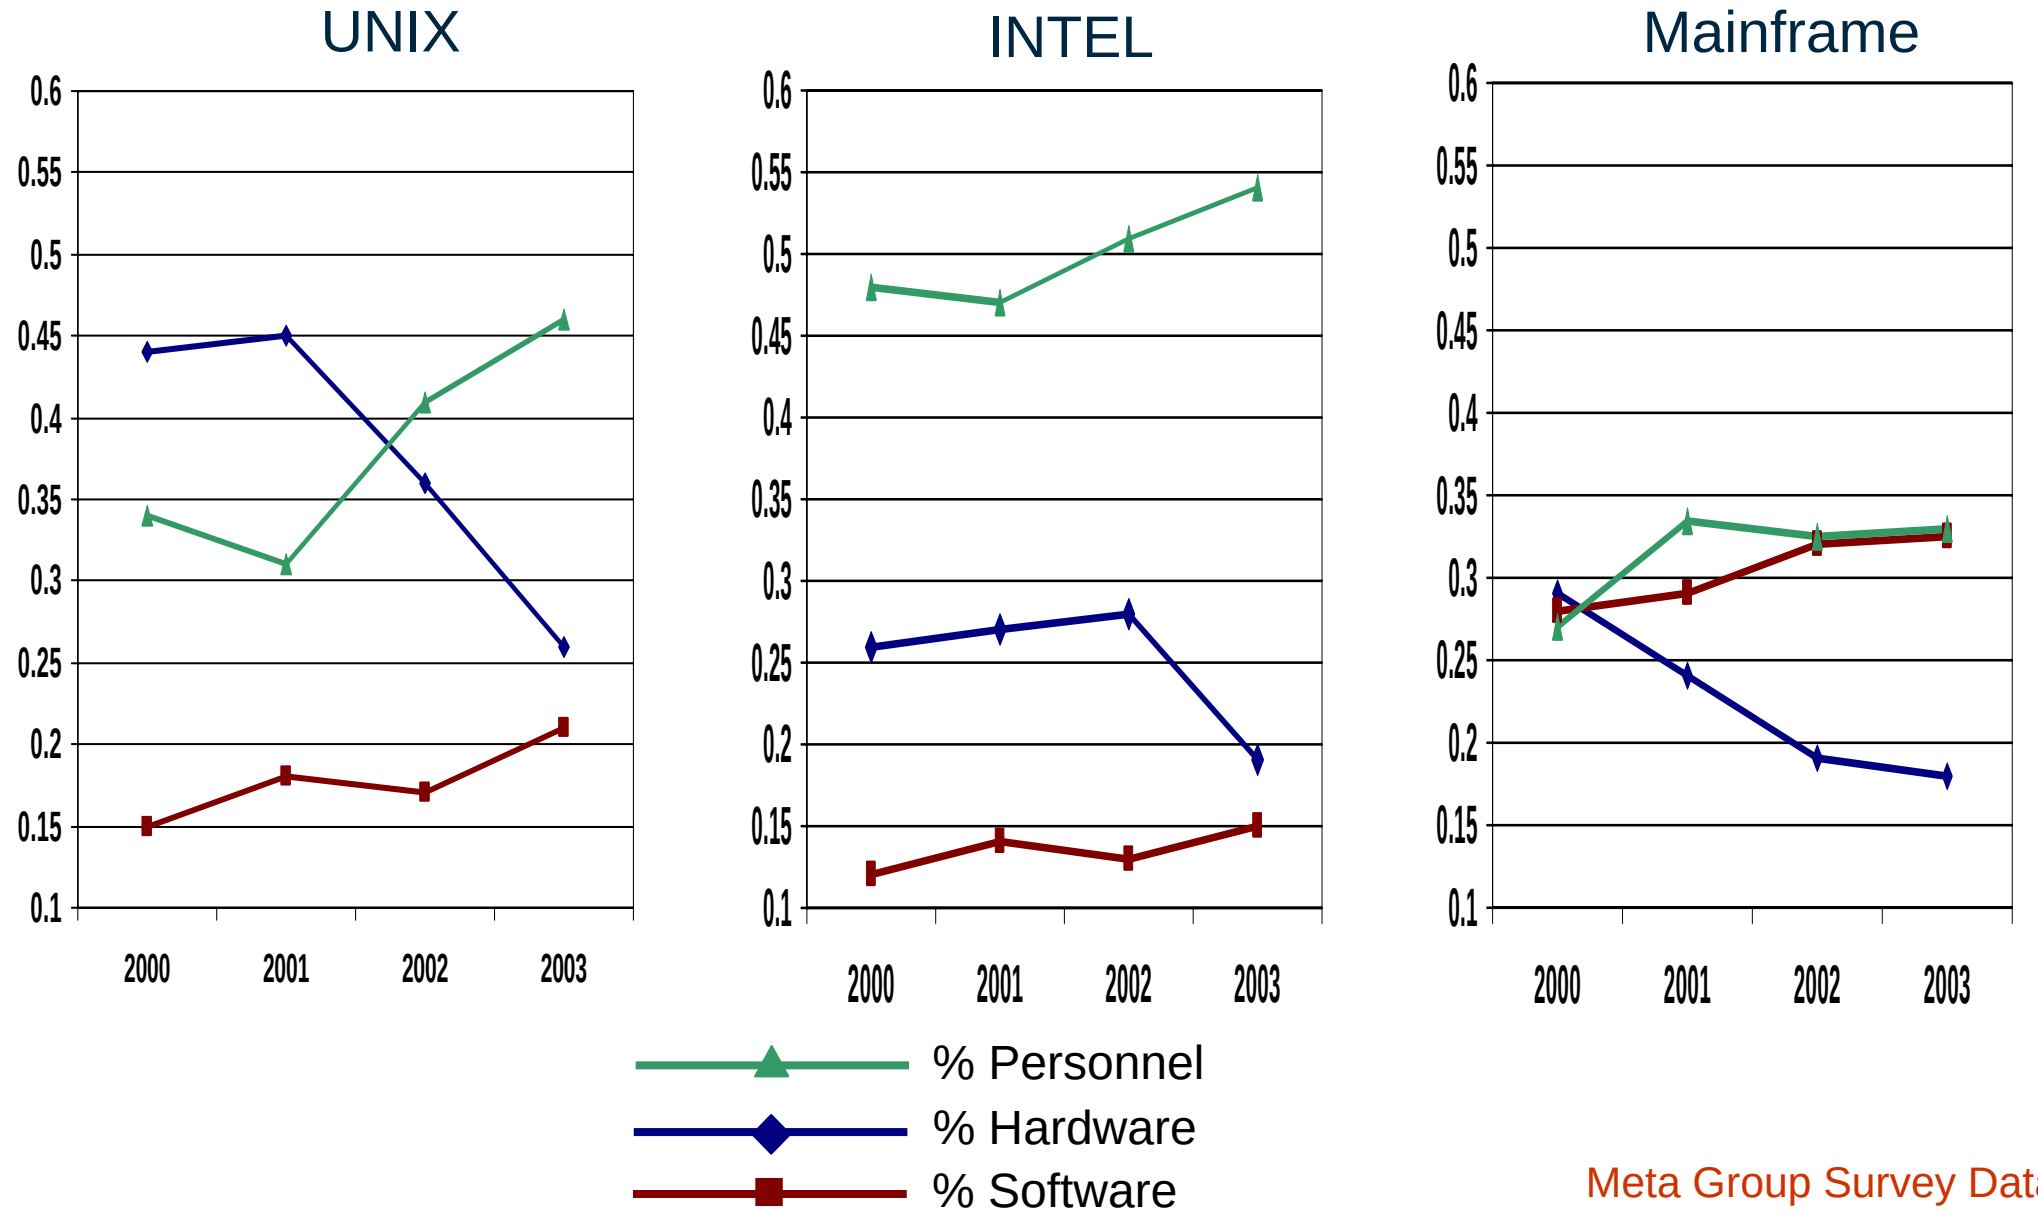

Meta Group Survey Data

Compliance/Auditing Pressure

- Regulatory compliance initiatives are impacting IT organizations in most countries/industries, and changing fast
  - Sarbanes-Oxley
  - Basel II
  - FDA: Food and Drug Administration 21 DFR Part 11
  - COPPA: Children's Online Privacy Protection Act of 2000
  - DPA: Data Protection Act (UK)
  - HIPAA: Health Insurance Portability and Accountability Act of 1996
  - PIPEDA: Personal Information Protection and Electronic Documents Act (Canada)
  - SEC Rule 17a-4: Records to be preserved by certain exchange members, brokers, dealers
  - USA Patriot Act: Uniting and Strengthening America by Providing Tools Required to Intercept and Obstruct Terrorism of 2001
- Focus is on both external threats (hackers) and internal employees

Security in DB2 for z/OS Vnext

Some key implementations

- Roles
- Network Trusted Contexts
- Instead of Triggers
- Improved auditing
- Secure Socket Layer
- Data Encryption

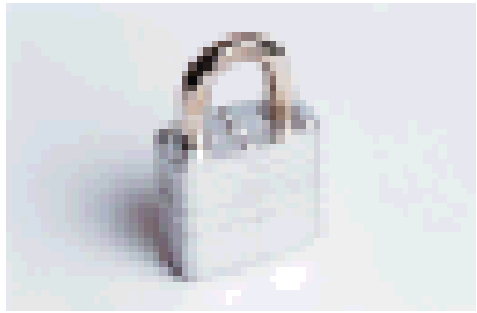

Protecting data on disk

- We will allow encryption for the key disk resources used by DB2:
  - Tables
  - LOBs
  - Indexes
  - Image copies
  - Logs
  - Archive logs

Database ROLEs

- ROLE is a “virtual authid”
  - Assigned via TRUSTED CONTEXT
  - Provides additional privileges only when in a trusted environment using existing primary AUTHID.
  - Can optionally be the OWNER of DB2 objects

```
CREATE ROLE PROD_DBA;
GRANT DBADM … TO PROD_DBA;

CREATE TRUSTED CONTEXT DBA1 …
   DEFAULT ROLE PROD_DBA OWNER(ROLE);
```

Database ROLEs Examples

- Dynamic SQL access to DB2 tables using JDBC or CLI, but only when running on a specific server.
- DBA can be temporarily assigned a DBA ROLE for weekend production table admin work – no table access at other times.
- DBA uses a ROLE for CREATE statements, so that the ROLE owns the objects he or she creates.
- Project librarian assigned a BIND ROLE only when running on the production code library server – can't BIND from any other server.

Trusted Security Context / ROLE WebSphere example

- WebSphere connection pool can be created with one DB2 AUTHID
- WebSphere can reuse pooled connections to DB2 with different AUTHIDs
- DB2 AUTHIDs can be given privileges that are only available when executing in WebSphere:
  - ✓ e.g. dynamic SQL access for JDBC only when using WebSphere

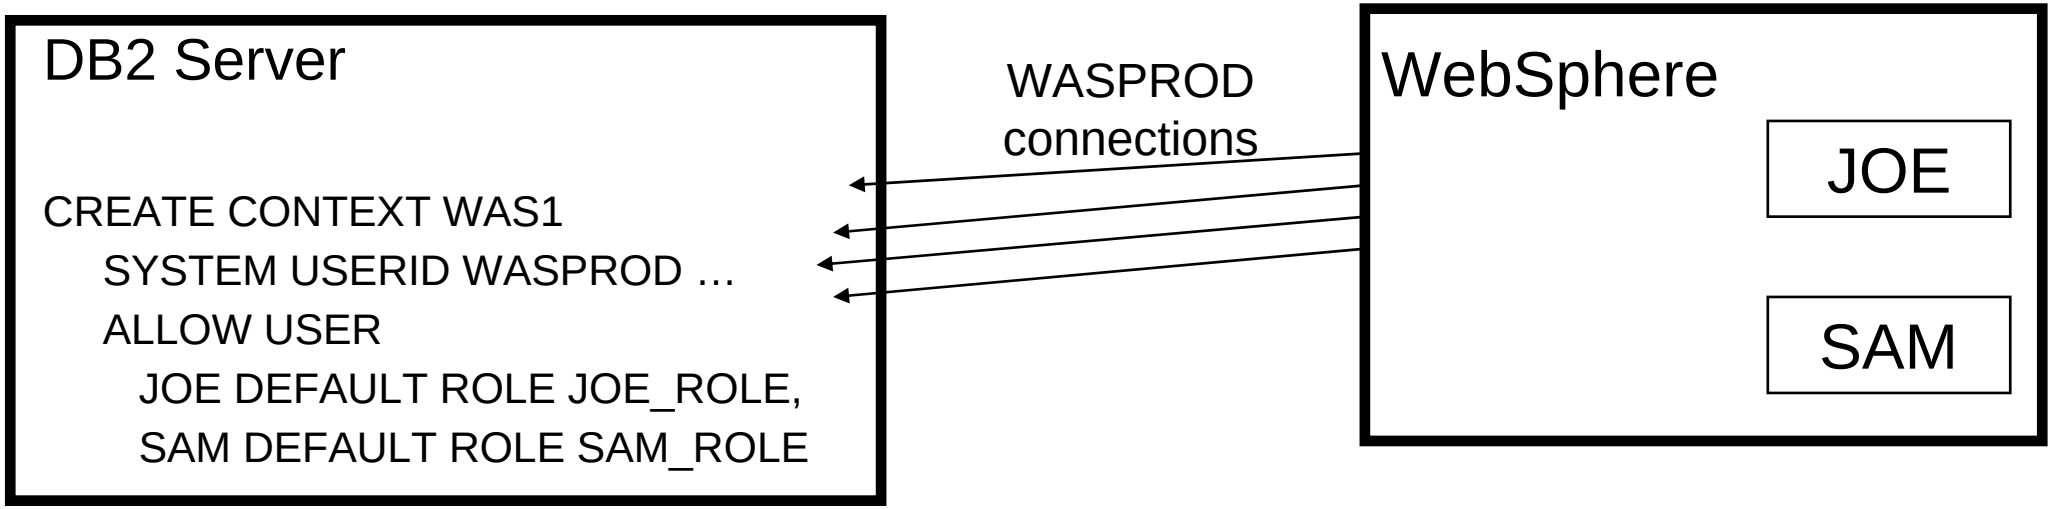

Volume-based COPY/RECOVER

- FlashCopy technology used to capture entire content of disk volumes
- RECOVER modified to enable object-level recovery from volume FlashCopy
- Eliminates labor associated with setting up COPY jobs for each database / table space

INDEX Compression

- INDEX pages will now be 4K, 8K, 16K or 32K
- Compression can be used to compress to 4K pages on disk
  - Reduces disk space requirements by up to 4X
- Compression/decompression takes place during I/O (pages in buffer pool are not compressed -- optimized for large REAL memory)
- No dictionaries used for index compression
  - New algorithm that compresses “on the fly”

DB2 vNext Themes

- **Enable high-volume transaction processing for next wave of Web applications**
- **Extend the lead in transaction processing availability, scalability and performance**
- **Reduce cost of ownership and zSeries-specific skill needs**
- **Improve data warehousing and OLTP reporting**

Data Warehousing, Reporting and Optimizer Improvements

- Cross query block optimization
- Histogram statistics exploitation
- Generalize sparse index and in-memory data cache method
- Dynamic Index ANDing for Star Schema
- EXCEPT and INTERSECT
- Many other SQL improvements

DB2 UDB for z/OS Vnext

- Integration
- Availability
- Scalability
- Productivity
- Total cost of ownership

- Data Definition On Demand
- XML, Unicode, LOBs

- SQL for DB2 family

- Utility enhancements

DB2 Tools Update

IBM DB2 and IMS Tools – how it started…

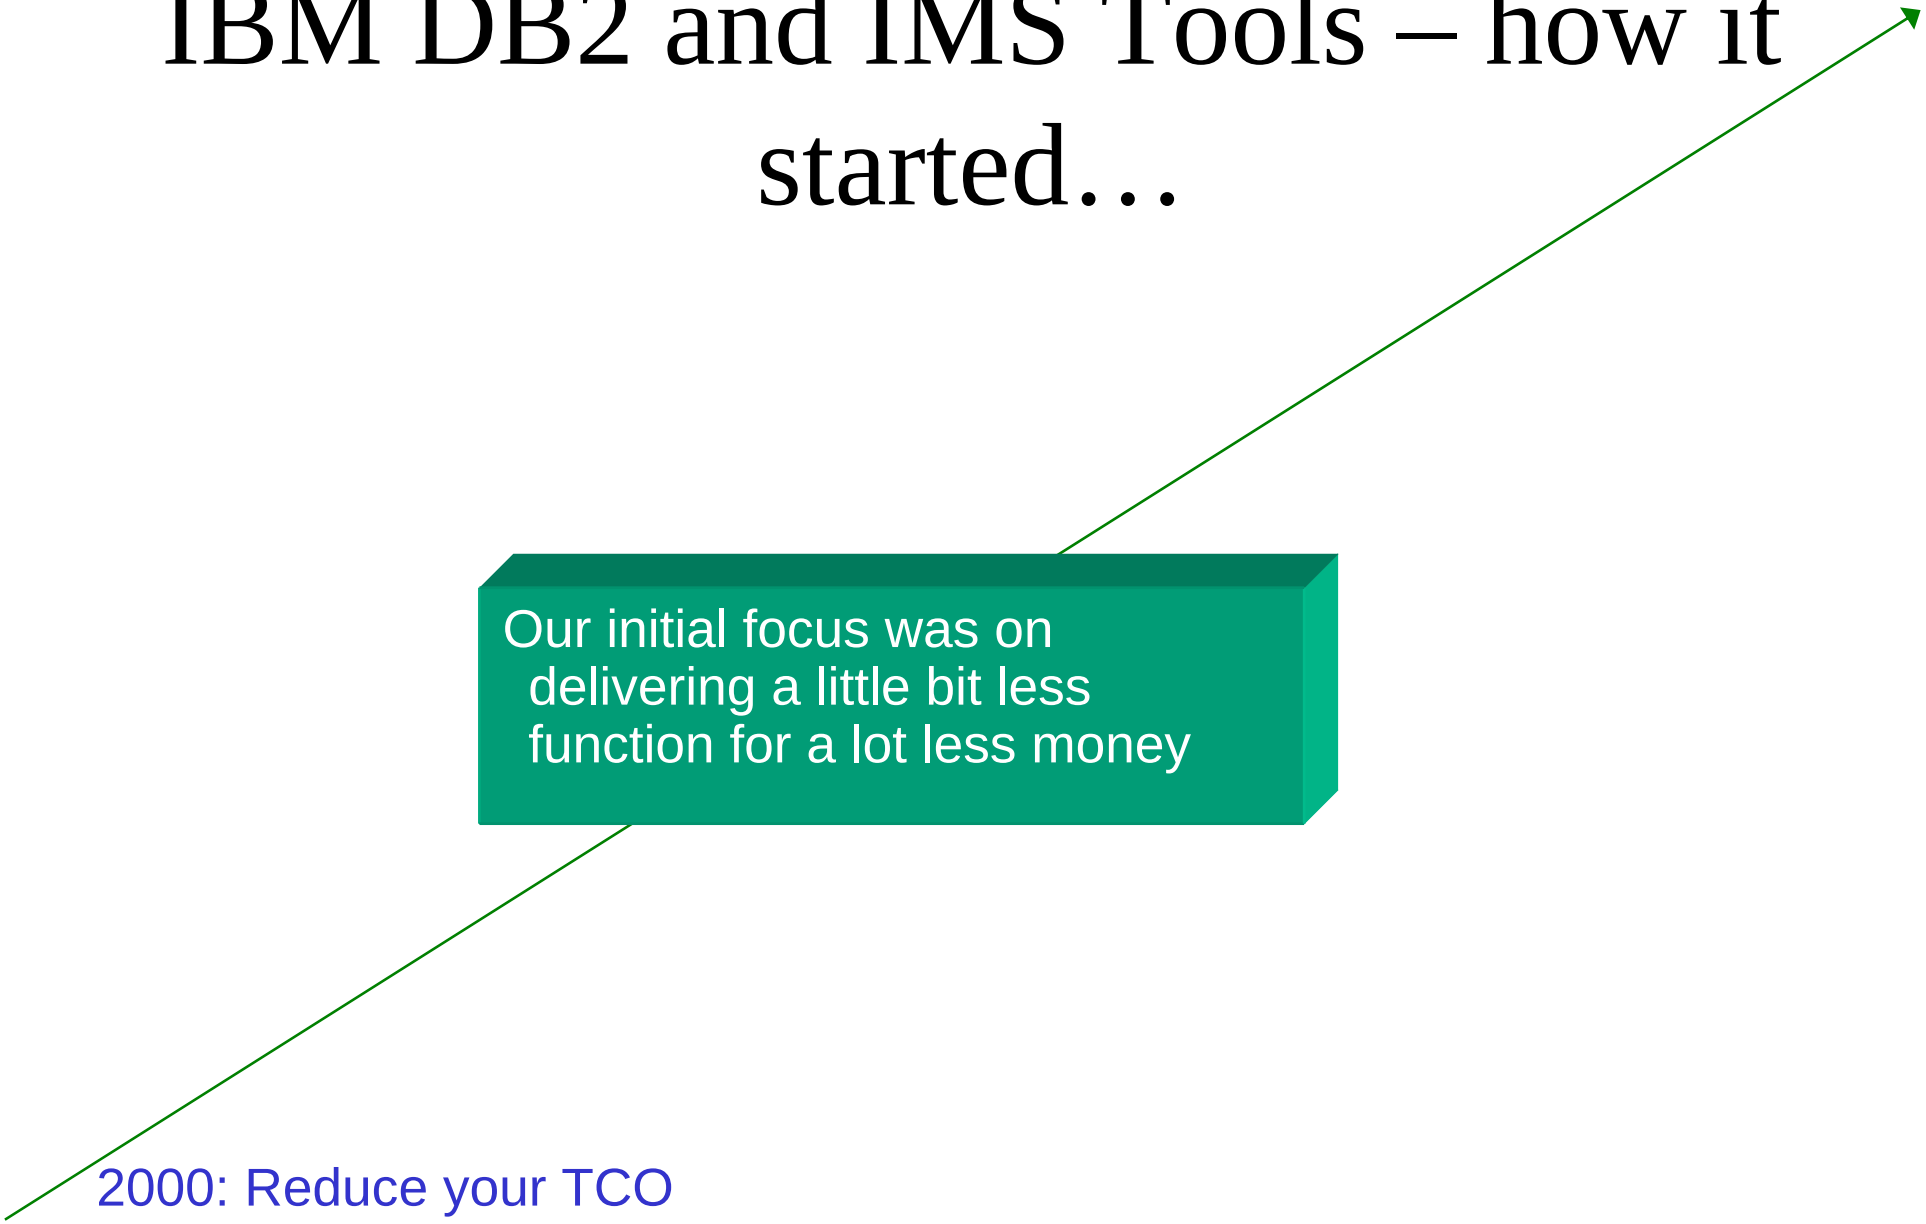

IBM DB2 Tools – stepping forward…

- All DB2 tools exploit V8 on z/OS and v8.2 on LUW today; All IMS tools support V9 today

We quickly produced a full set of products

2001-2002 Product replacements
We had to fill in the holes and provide tools to replace your existing tools, with Day-1 support

2000: Reduce your TCO

Today's IBM DB2 Tools

Database Administration

- DB2 Administration Tool
- DB2 Object Comparison Tool
- Data Encryption for DB2 and IMS
- DB2 Thread Expert
- DB2 Administration Toolkit SAP Edition

Utilities Management

- DB2 Utilities Suite
- DB2 Automation Tool
- **DB2 HP Unload**
- DB2 Automation Toolkit SAP Edition

Performance Management

- **DB2 Performance Expert**
- DB2 Performance Monitor
- OMEGAMON XE for DB2 on z/OS
- DB2 Buffer Pool Analyzer
- DB2 Query Monitor
- DB2 SQL Performance Analyzer
- DB2 PerformanceToolkit SAP Edition

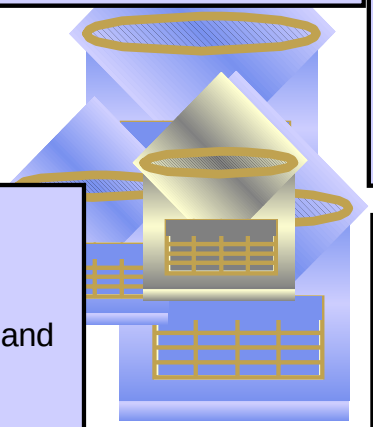

Recovery Management

- Application Recovery Tool for IMS and DB2 Databases
- DB2 Archive Log Accelerator
- DB2 Change Accumulation Tool
- DB2 Log Analysis Tool
- DB2 Object Restore Tool
- **DB2 Recovery Expert**

Application Management

- DB2 Data Export Facility
- DB2 Path Checker
- DB2 Bind Manager
- DB2 Data Archive Expert
- **DB2 Table Editor**
- **DB2 Web Query Tool**
- **DB2 Test Database Generator**

Information Integration

- DB2 DataPropagator
- SQData
- WebSphere II Replication
- WebSphere II Classic Federation
- WebSphere II Event Publisher for z/OS

Available on zSeries

Available on zSeries and multiplatforms

Currently only available on multiplatforms

Today's IBM IMS Tools Product Portfolio

Database Administration

- Data Encryption for IMS and DB2 Databases
- IMS Database Repair Facility
- IMS HALDB Conversion and Maintenance Aid
- IMS Hardware Data Compression Extended
- IMS Library Integrity Utilities
- IMS Parameter Manager
- IMS Sequential Randomizer Generator

Information Integration

- IMS DataPropagator
- WebSphere II Classic Event Publisher for IMS
- WebSphere II Classic Event Publisher for VSAM
- WebSphere II Classic Event Publisher for CA-IDMS
- WebSphere II Classic Federation for z/OS

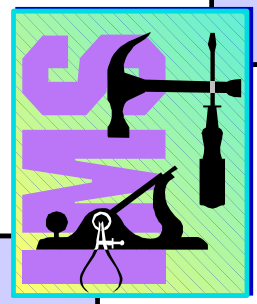

Performance Management

- IMS Buffer Pool Analyzer
- IMS Performance Analyzer
- IMS Problem Investigator
- IMS Network Compression Facility
- Tivoli OMEGAMON XE for IMS

Recovery Management

- Application Recovery Tool for IMS and DB2
- IMS DEDB Fast Recovery
- IMS Database Recovery Facility
- IMS High Performance Change Accumulation
- IMS High Performance Image Copy

Application Management

- IMS Batch Backout Manager
- IMS Batch Terminal Simulator
- IMS Connect
- IMS Connect Extensions
- IMS MFS Reversal Utilities
- IMS Program Restart Facility

Utilities Management

- IMS Database Control Suite
- IMS High Performance Fast Path Utilities
- IMS High Performance Load
- IMS High Performance Pointer Checker
- IMS High Performance Prefix Resolution
- IMS High Performance Unload
- IMS Index Builder
- IMS Parallel Reorganization

TM Management

- IMS Command Control Facility
- IMS ETO Support
- IMS High Performance System Generation Tools
- IMS Queue Control Facility

IBM DB2 Tools – looking ahead…

We invested, developed, and delivered in V8 – we continue to invest for deliveries in 2005 and beyond

Done, and continuing with common components, like DB2 Grouper

2003: Integration of products, new capabilities, better performance

2001-2002 Product replacements

2000: Reduce your TCO

zSeries Shipped MIPS History and Inventory Growth Trend

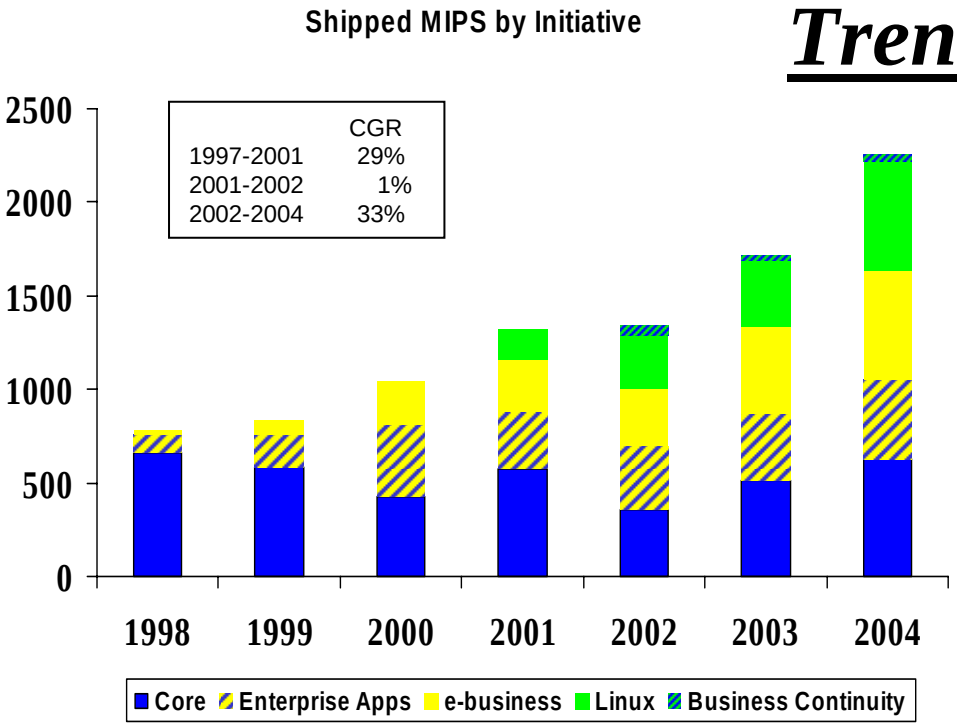

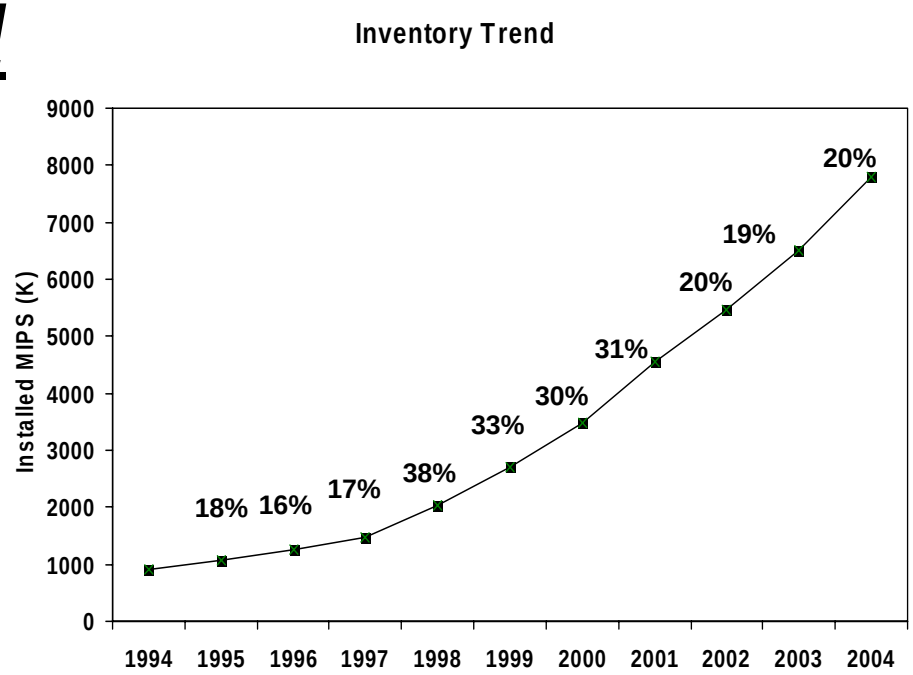

	1998	1999	2000	2001	2002	2003	2004
Shipped MIPS (K)	788	839	1049	1327	1338	1716	2258
YTY		+6%	+25%	+27%	+1%	+28%	+32%
Inventory (KMIPS)	2030	2693	3490	4556	5464	6500	7800
YTY	38%	33%	30%	31%	20%	19%	20%

Mainframe Database Tools Market
SW Costs flat as MIPs more than doubled

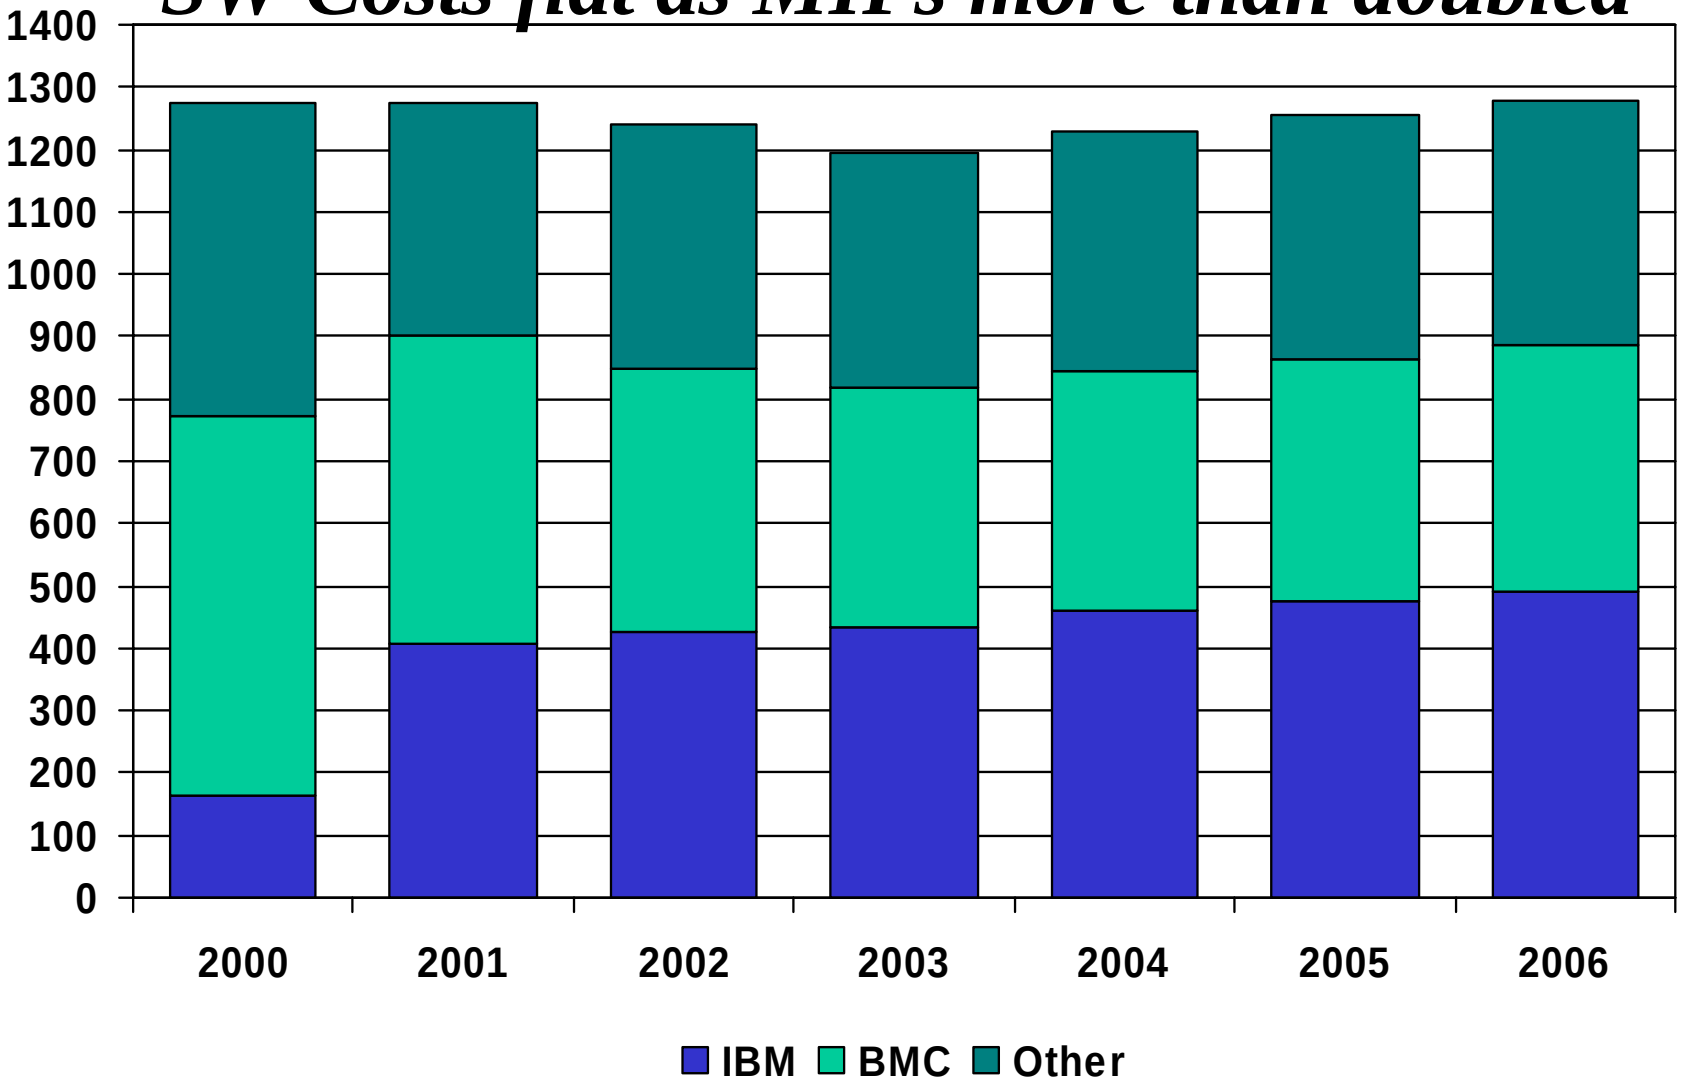

Sources: BMC Segment reports (actual rates). IBM FIW (actual rates).
Sources: Total/Other from IDC and GMV

Innovation

- Significant investment in Autonomic tooling
  - Test Database generator
  - Grouper
  - Performance Expert
  - Recovery Expert
  - Data Archive Expert
- More coming
  - Change management
  - Policy management
  - Audit Expert
  - Regulatory compliance
  - ….

"I think there is a world market for maybe five computers."

Thomas Watson, chairman of IBM, 1943

"Computers in the future may weigh no more than 1.5 tons. "

Popular Mechanics, 1949

"There is no reason anyone would want a computer in their home. "

Ken Olsen, founder of DEC, 1977

"640K ought to be enough for anybody. "

Bill Gates, 1981

"Prediction is difficult, especially about the future"

Yogi Berra

IBM DB2 Tools – "still" focusing on TCO…

Our first focus is still our current one: Reducing your Total Cost of Ownership

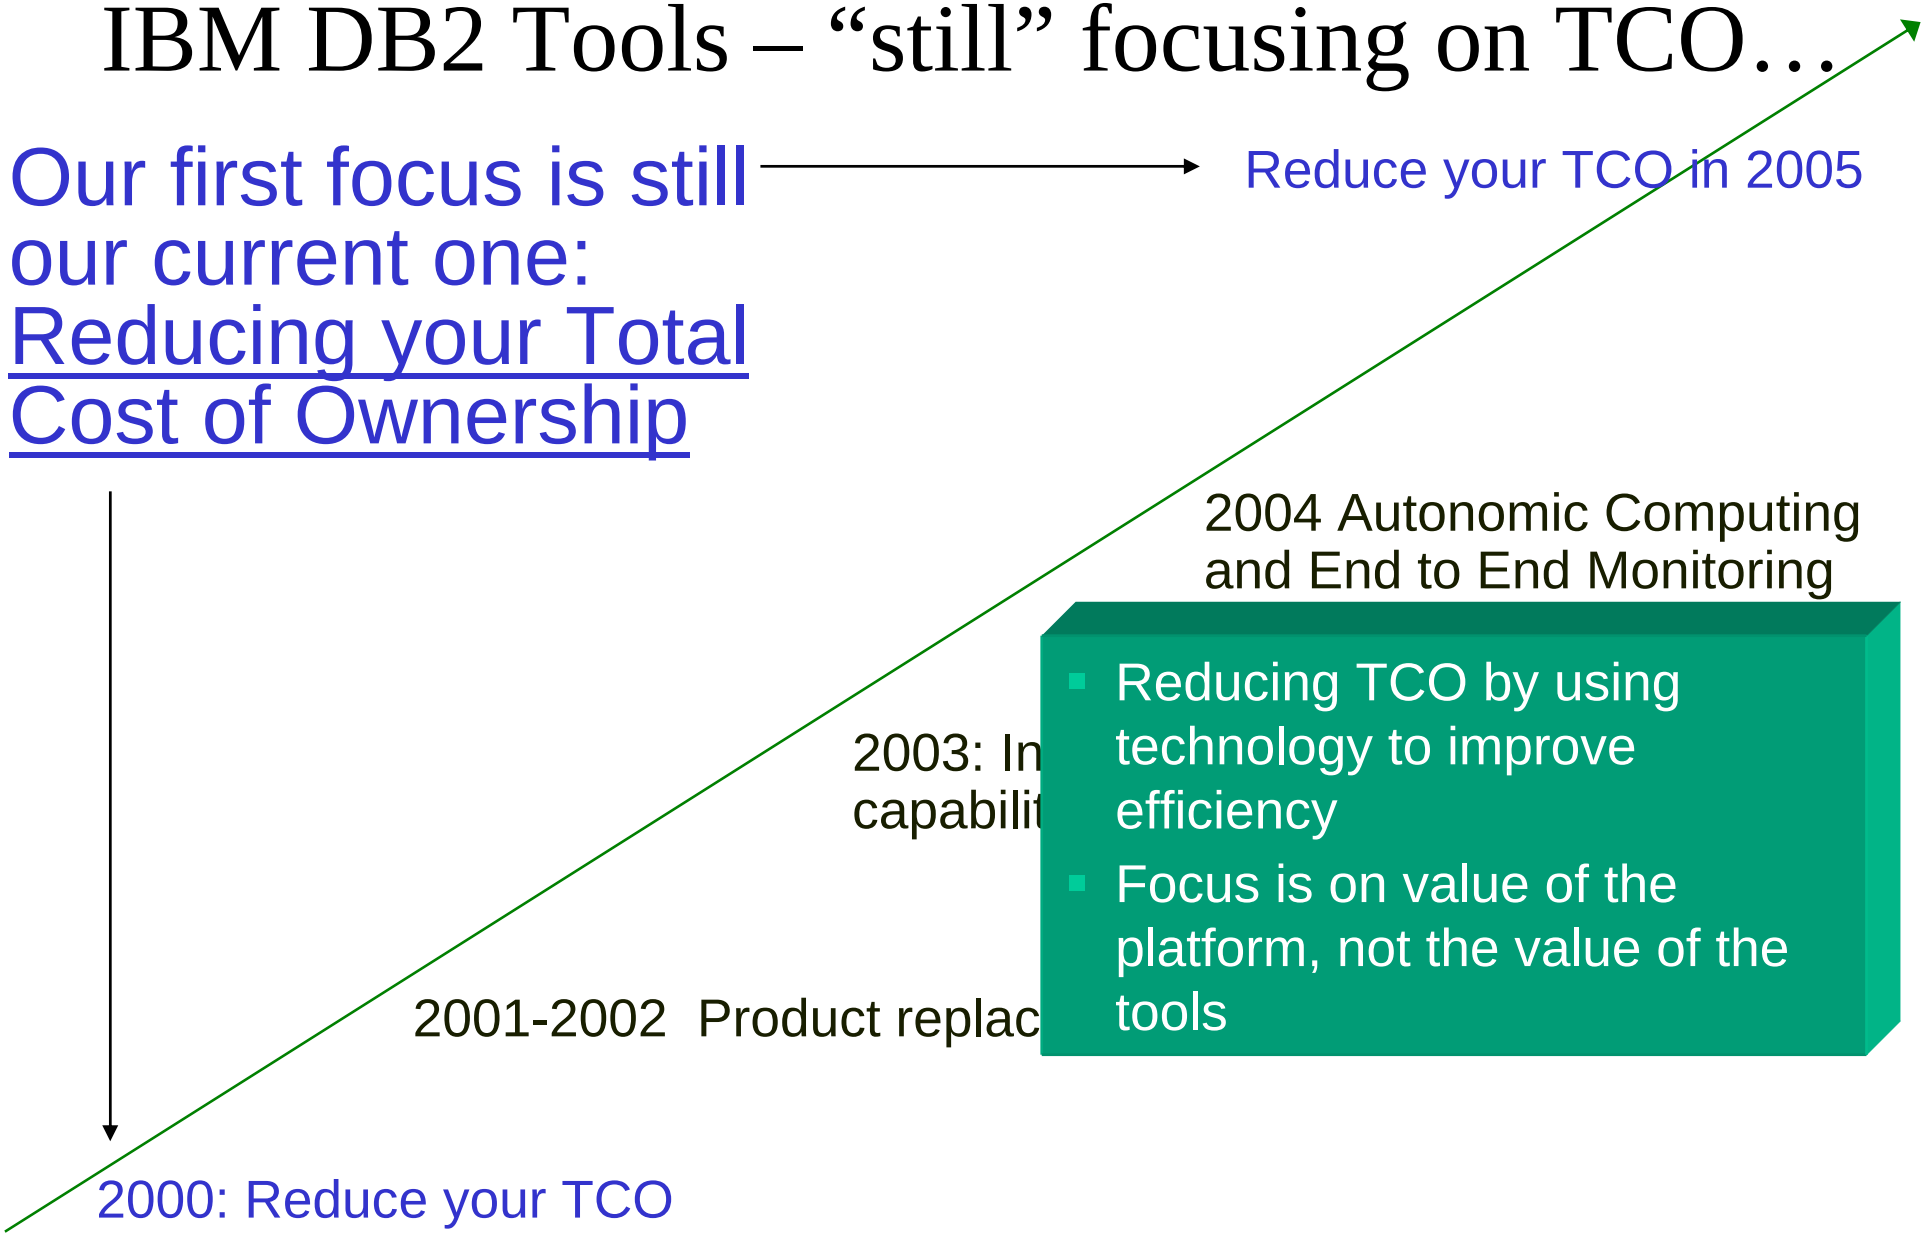

Advances in wheel and tire technology

In 1908, Hayes Wheels, manufactured its first product — a wooden spoked wheel for the Ford Model T

In 1904, mountable rims were introduced that allowed drivers to fix their own flats.

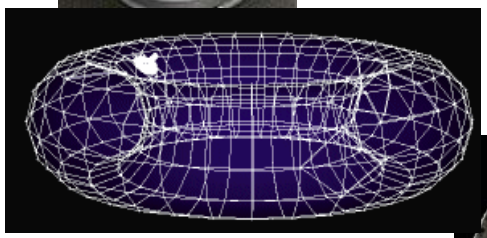

In 1911, Philip Strauss invented the first successful tire, which was a combination tire and air filled inner tube

In 1954, the Packard exploited P.W. Litchfield's 1903 patent for the tubeless tire belonging to the Goodyear Tire Company

In 1994, the first "run-flat" tire (by Goodyear) available on a regular production vehicle was offered as an option on the Chevrolet Corvette

With run-flat tires, there is no need for a spare wheel or a car jack, for that matter. This frees up space for luggage, lowers cost and saves some weight, which makes for a small gain in fuel economy.

http://inventors.about.com/library/inventors/bltires.htm

Change Tire

Topics

- Reorganization
- Recovery
- Change Management
- Regulatory compliance

OMEGAMON for DB2 and DB2 Performance Monitor/Expert - Convergence

Best of breed DB2 monitoring, analysis, and tuning solution

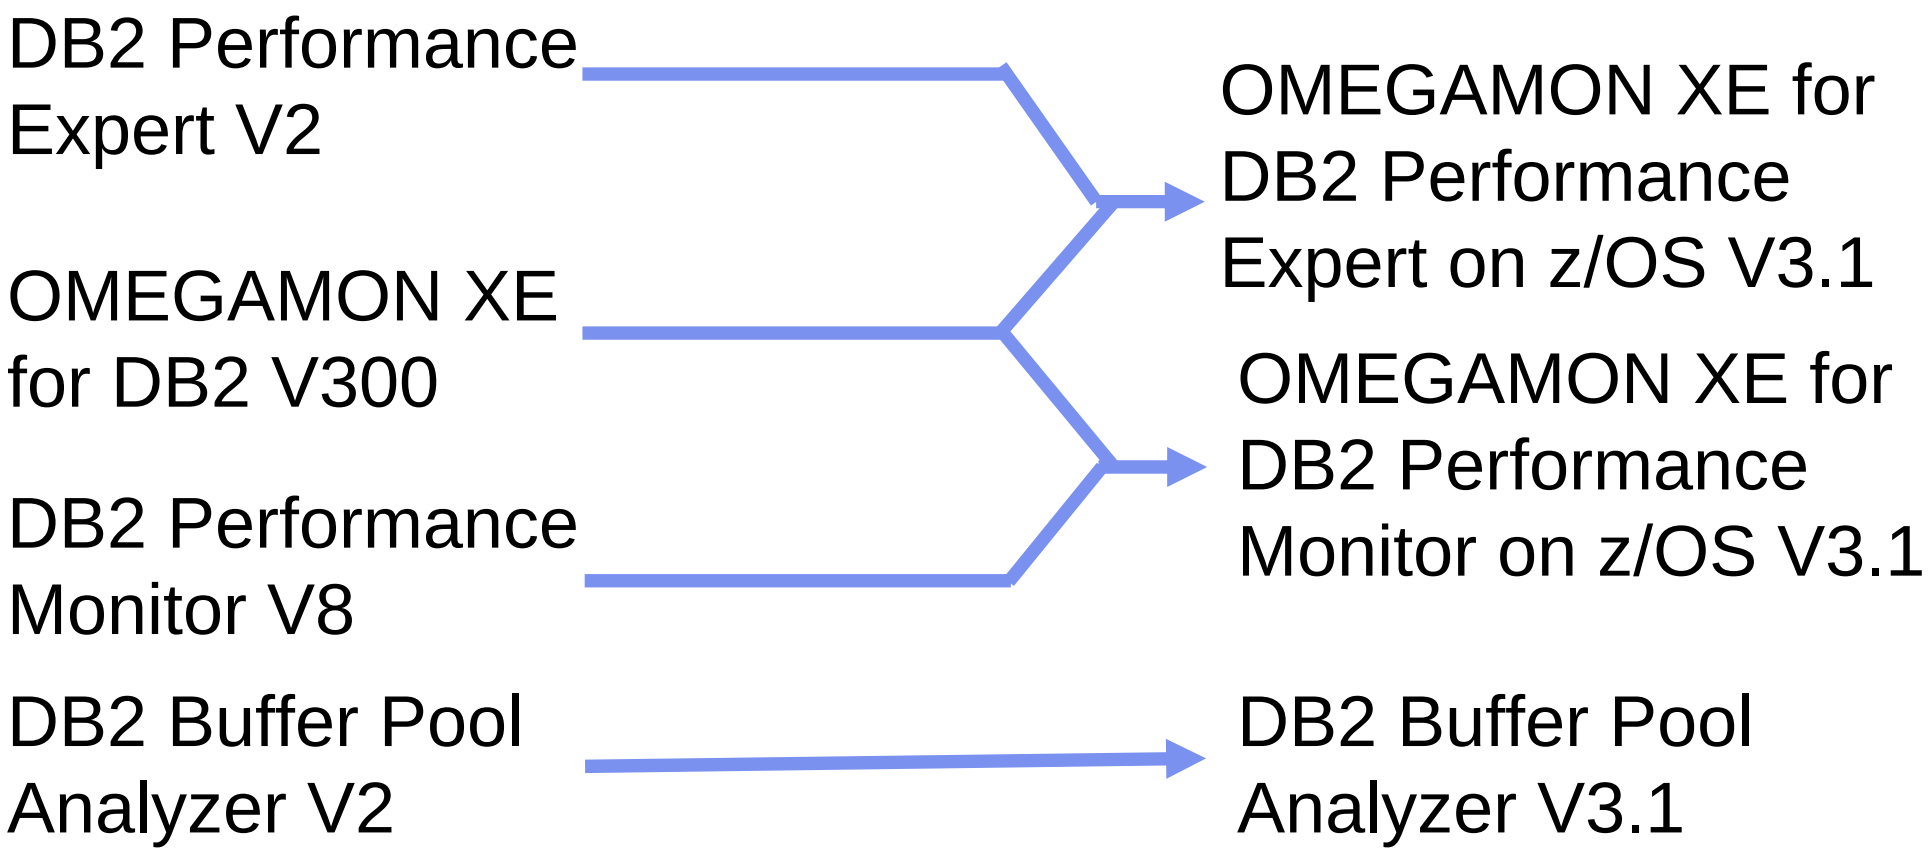

Determining the need to REORG --- without tools

- DBA reads the DB2 Administration Guide and Utility Guide and Reference
  - Find the recommended algorithms on when to REORG
  - Build SQL statements to query to catalog to find the values for the algorithms and extracts the information
  - Determine the REORG options desired
  - Build JCL to run REORG if the table space needs to be reorganized
  - Submit Job
- DBA repeats the above for all table spaces
- DBA repeats periodically to maintain table space reorganization

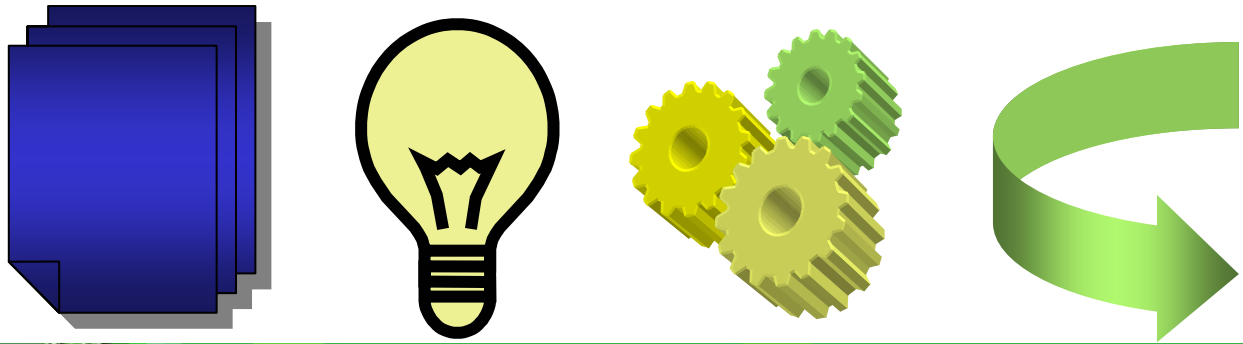

Determining the need to REORG … with DB2 Automation Tool

- DBA initiates a dialog with the tool
  - Define an object profile with ALL the table spaces
  - Define a utility profile for REORG with the proper options
  - Define an exception profile with checks for the proper statistics
  - Tie the three profiles together in a job profile
  - Set up Automation Monitor in job scheduler to run at a desired frequency.
- That's it !
  - The "monitor" runs the job profile, statistics for each table space in the utility profile are retrieved, and compared to the criteria in the exception profile.
  - Jobs are generated using the REORG utility profile for table spaces that meet the exception profile criteria

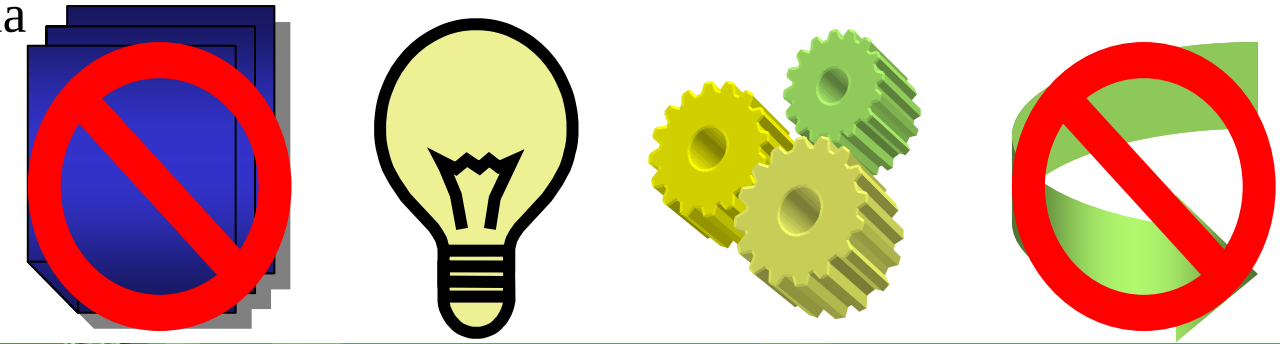

Determining the need to REORG … with DB2 Automation Tool in the future

- DBA initiates a dialog with DB2 Automation Tool
  - Define a utility profile, an object profile, and an exception profile.
- That's it !
  - No job profile necessary, nothing to set up and run on a periodic basis
  - Background "daemon" checks exception profile criteria against the statistics chosen
  - Automation Tool determines need to REORG a table space, generates the job, and runs it or schedules it to run during an acceptable batch windo

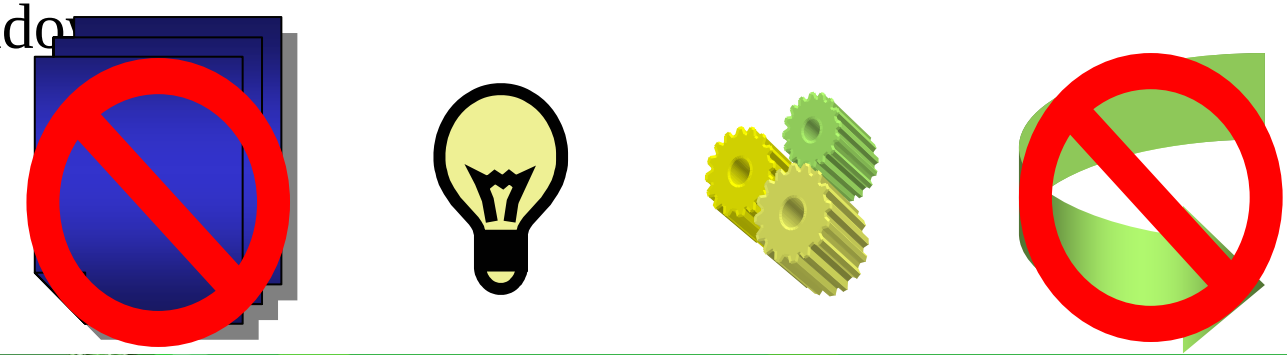

Application Recovery --- without tools

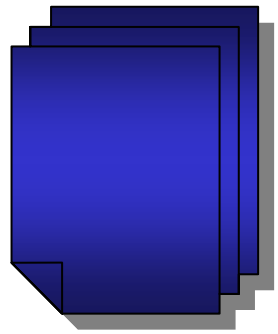

Recovery of a complete business application to three weeks ago, prior to a schema change.

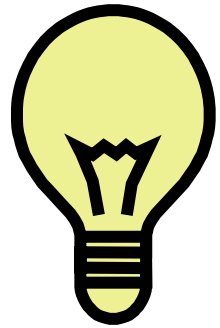

- DBA manual inventory of application parts if they exist
  - DDL definitions, RI, programs, data, indexes, views, triggers, stored procedures, security settings, binds, catalog statistics, etc,…
- Manually build SQL statements to query catalog to find definitions and extract current information to compare and see what needs to be reset and restored
  - Restore application parts
  - Determine methods to do point in time recovery on data and Indexes
  - Restore backups and roll forward to selected point in time
  - Test all application facets to be sure proper function and data are correct
- Complex, labor intensive, time consuming process

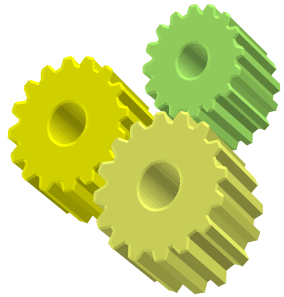

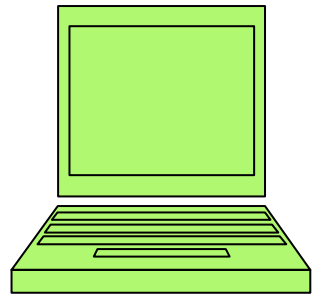

Application Recovery --- with current tools

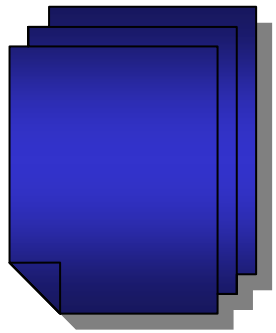

- DBA manual inventory of application parts if they exist
  - DDL definitions, RI, programs, data, indexes, views, triggers, stored procedures, security settings, binds, catalog statistics, etc,…
- Tools help build SQL statements to query catalog to find definitions and extract current information to compare and see what needs to be reset and restored
  - DB2 Object Restore can restore specified objects
  - DB2 Change Accumulation can reduce backup image copy processing
  - DB2 Log Analysis can display DB2 activity; create undo/redo SQL
    - These provide point in time recovery
  - Test all application facets to be sure proper function and data are correct
- Multi-step, labor intensive, time consuming process

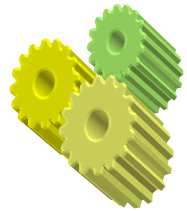

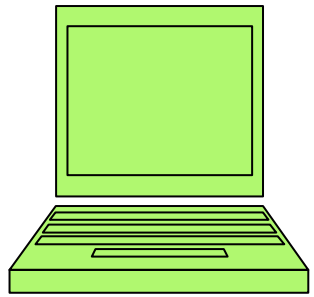

Application Recovery --- with future tools

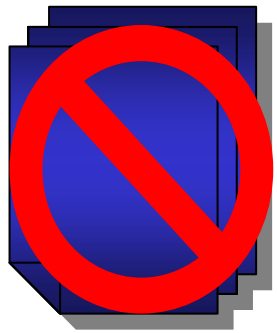

- Use DB2 Recovery Expert versioning repository to inspects version levels available for restoration including related dependent objects.
  - DDL definitions, RI, programs, data, indexes, views, triggers, stored procedures, security settings, binds, catalog statistics, etc,… available at different versions
  - Even finds tables that are used together for Dynamic SQL
- Use DB2 Recovery Expert to select point in time recovery specifications
  - Recovery Expert submits jobs to restore the application and return results
  - Test all application facets to be sure proper function and data are correct
- Simple, direct, autonomic process to recover a business application

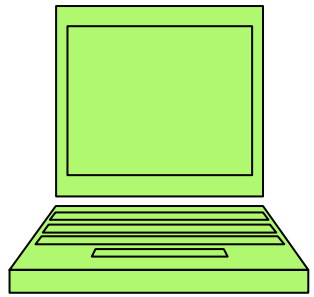

Database Change Management – without tools

- DBA determines some portion of an objects definition needs modification
  - Using SQL interrogates the catalog to find all dependencies this object has.
  - Using SQL extracts the definition of this object and all its dependents
  - As the extraction is rows from the catalog, DBA must construct DDL from the catalog information
  - Modifies the object for the new structure and any of the dependents that are affected by this change
  - Unloads the data
  - Drops the object and all its dependents
  - Recreates the object and all its dependents
  - Reloads the data
  - (Re)binds the application
  - Reissues all the authorization commands
- Without a tool, this is a very labor intensive, error prone operation

Database Change Management --- current tools

- Using the Tool, DBA indicates a change is needed to an object structure
  - Tool determines if another change for this object, or its dependents, is in process.
    - If yes, notifies the user of this and asks how to proceed (e.g., join the two changes, supersede it, ignore it, etc.)
    - If no, establishes a change ID for tracking purposes
  - Tool then automates the steps from the previous slide
    - Provides an easy interface to modify the existing object
    - Automatically propagates the change to dependent objects
    - Builds job streams to enact the change
  - User submits the job(s)
- Once the user is satisfied with this change, the tool provides an interface to promote this change to any desired target system.

Database Change Management – the future

- While the basic steps the tool must perform to enact a change are the same, the user experience changes as more graphical visualization is used.
  - A model of what is exists today is extracted from the catalog
  - Automatically finds all related schema change elements including “undocumented” relationships via GROUPER.
  - A filtered view of dependencies, the source schema and the change are displayed
  - DBA modifies the model to describe what the end result should be
  - Impact of applying the “current state of the change” is visible at all times during development
  - When the DBA is satisfied with the change, the tool generates change steps to any target system.
  - Inherently supports NLS, DBCS, Unicode, Translation, and usability requirements.....
  - Synergy with Rational modeling tools
  - Can insure the database definitions stay in line with company standards and needs.
  - Value proposition is productivity, faster change development and reduced errors

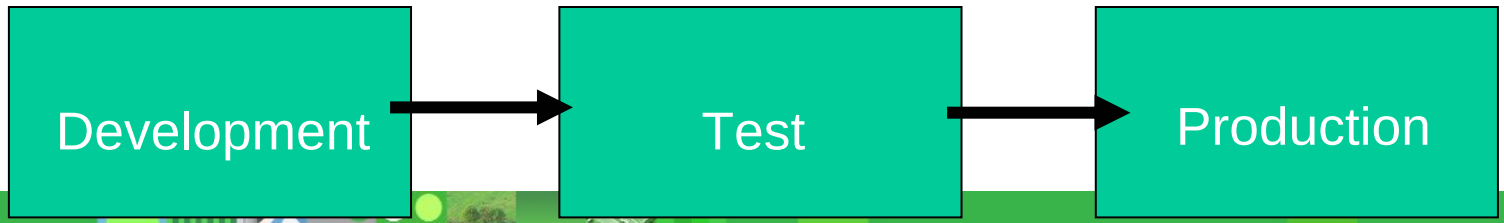

Change management encompasses the following distinct tasks:

- ✓ Comparing two environments (sets of objects) to determine where they differ.
- ✓ Analyzing the impact of a proposed change on a database.
- ✓ Migrating a set of objects or redefining the target objects to be like the source.
- ✓ Creating a historical baseline for future reference.
- ✓ Creating, viewing, adding, deleting, and modifying Change Managing the life cycle of structural changes to databases.
- ✓ Commands
  (DDL, DML, DCL, utilities, DB2 commands, system commands etc.).
- ✓ Managing the deployment of changes on the target database.
- ✓ Load, unload, and movement of data, such as running DB2 utilities bind/rebind, reorg, etc.

Regulatory compliance

- Protect sensitive data
- Save data for future audits and to comply with retention rules
- Discover responsible persons (who, what, where and when)
- For:
  - Sarbannes Oxley
  - HIPAA
  - Japanese Protecting Personal Freedom Act
  - …

Regulatory Compliance --- current & future tools

- Full support of data lifecycle – the elements of a suite:
  - TDBG - Control replication of sensitive data – cleanse test copies

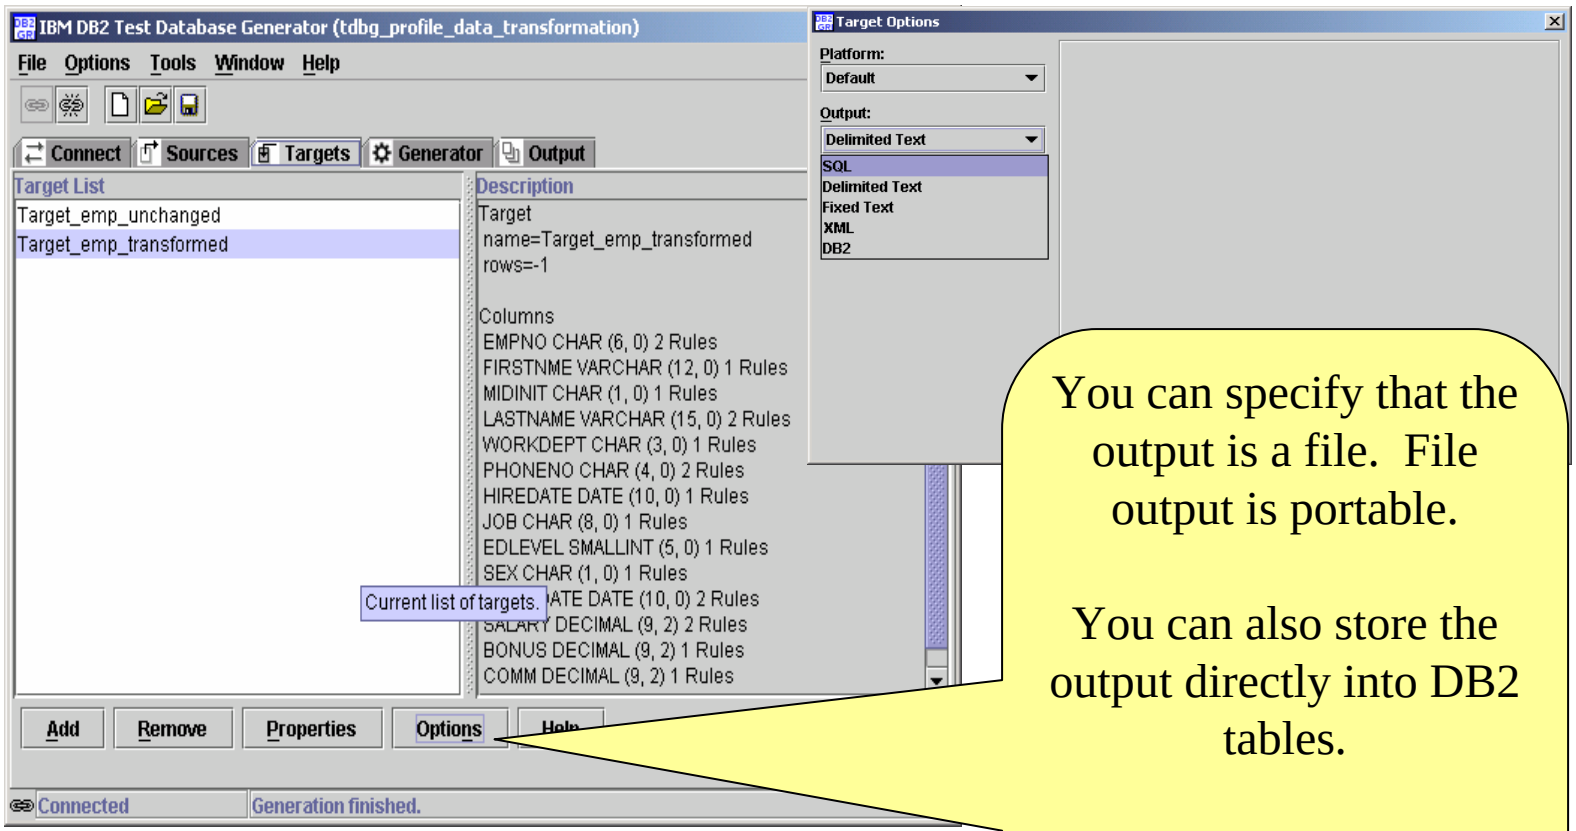

TDBG Target Transformations examples

- Examples
  - Social Security Number Char(9))
    - “Mask”, replace positions 1-6 with random numbers.
  - EMPNO (char(6))
    - “Static”, replace with 999999
  - Lastname (varchar(15))
    - “pattern”, replace with 5 random letters between A-Z
  - Phoneno Char(4))
    - “Mask”, replace positions 3-4 with random numbers.
  - Salary (Decimal)
    - “Random”, replace with random decimal between 1.04 and 112,000.05

- Archive data indefinitely
  - Sets of Related Rows - level of granularity
  - Schema evolution support
  - Durable Media support (worm devices)
  - Staging data to other tables and to files

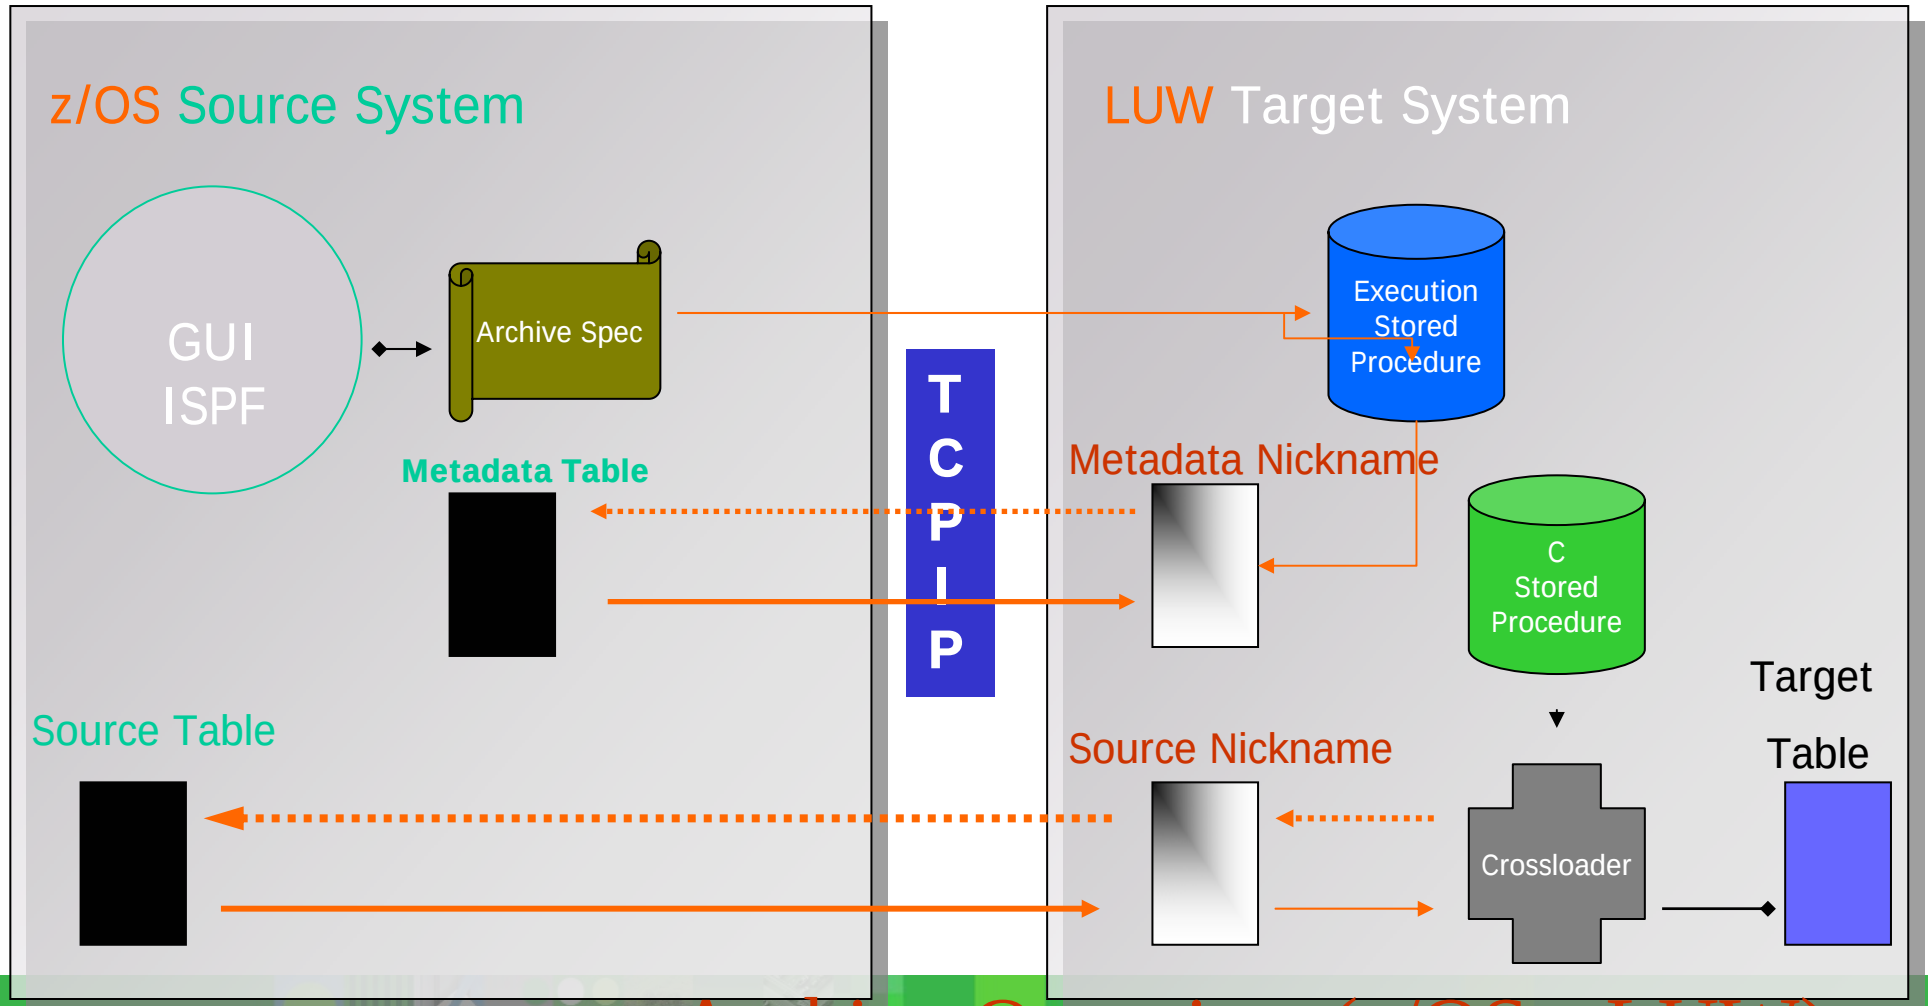

Archive Overview (z/OS->LUW)

Regulatory Compliance --- using future tools

- Support Auditing
  - Gather all the necessary data for auditing
    - People initiated and alert initiated
  - Automatic alerts of “out of policy” activities (authorization refusals)
  - Efficient and selective reporting to support authorized investigation

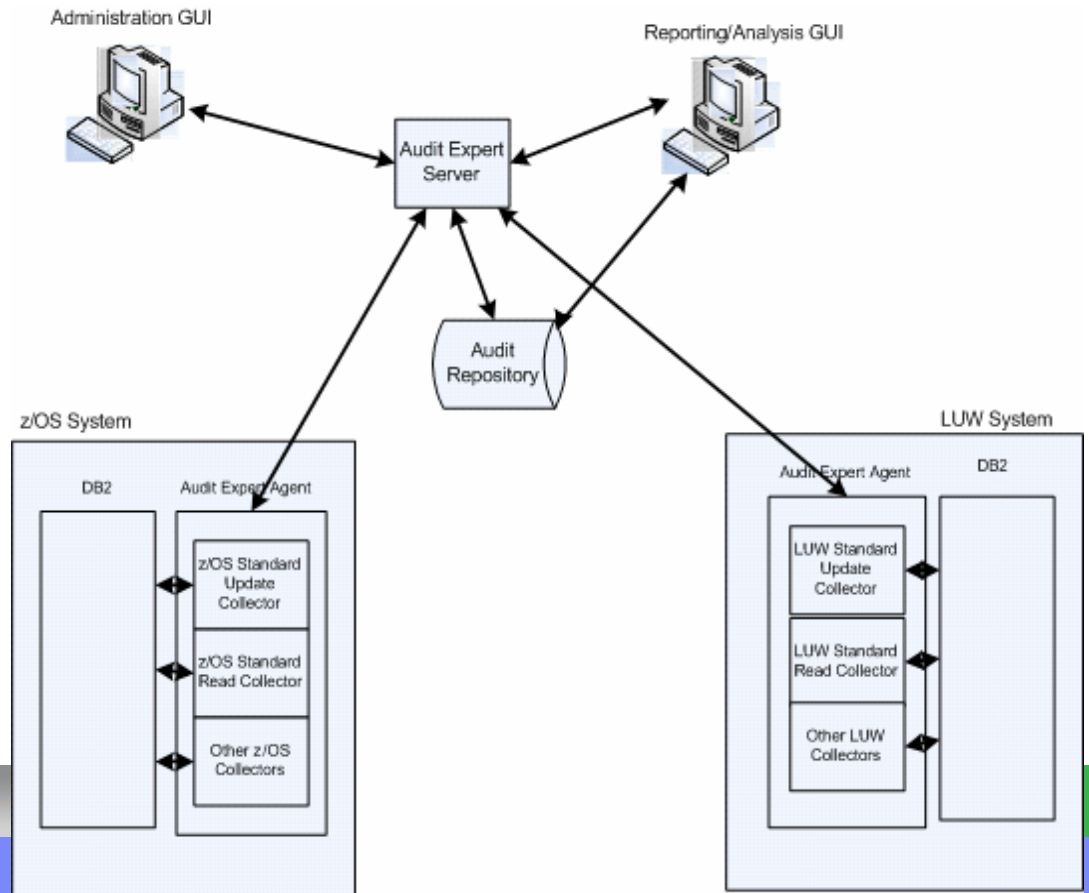

Summary

- Initial focus on Value to improve TCO
- Comprehensive solutions developed and deployed quickly
- Focus on autonomic features to add capability and simplify operations
- Now improving TCO by providing functions that let you
  - Manage bigger environments without adding staff
  - Focus on business value rather than maintenance
  - Avoid tedious tasks
  - Reduce errors
  - Enable people with less experience to be productive - fast
- Preserve your investment in DB2 applications and databases

Performance Monitoring --- using current tools

- DB2 Performance Expert is available today

- **Provides a CENTRALIZED, CONSOLIDATED performance facility for parallel monitoring of several DB2 systems**
  - Monitors for System Statistics - Applications - End2End - Bottlenecks …
  - Displays important performance data in graphical views
  - Shows current activities and history data
  - Provides exception and "expert" reports
  - Stores and manages performance data in a ***Performance Warehouse***

- Comprehensive tool, requires DB2 expertise, and additional analysis

Performance Monitoring --- using future tools

- DB2 Performance Expert extended

- Provides advanced autonomic capabilities
  - Emphasis on exception processing
  - Harvests warehouse data
    - Performance tuning, in the context of the OS
    - Trend analysis & capacity planning with automatic workload characterization
    - Problem diagnosis

- **Following examples are based on code running in the field today**

Performance Monitoring --- using future tools

- Autonomic emphasis on exception processing
  - Monitor online Event Exceptions (deadlocks)
    - Alert via audio alarm or message window pop-up - customizable
    - Display in the System Overview Panel
    - Drill down into details
  - Monitor user defined thresholds for periodic checking

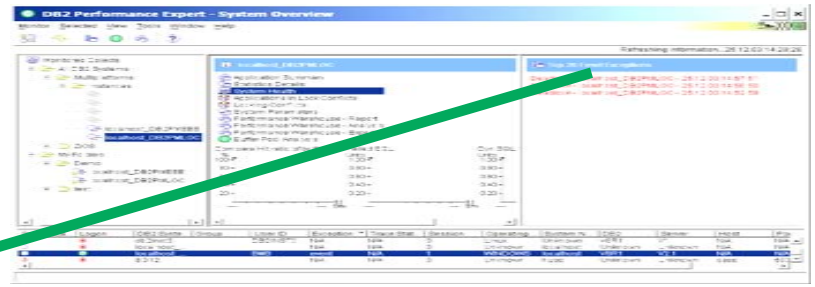

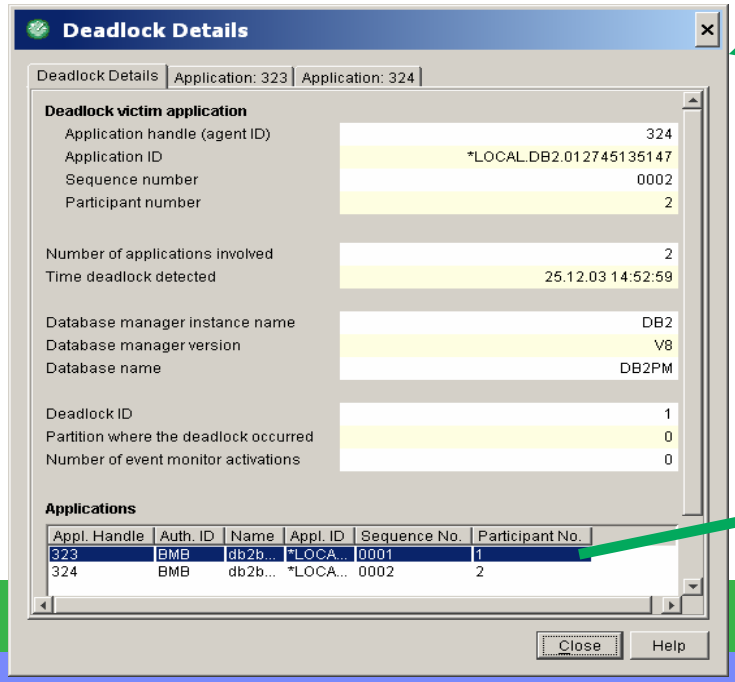

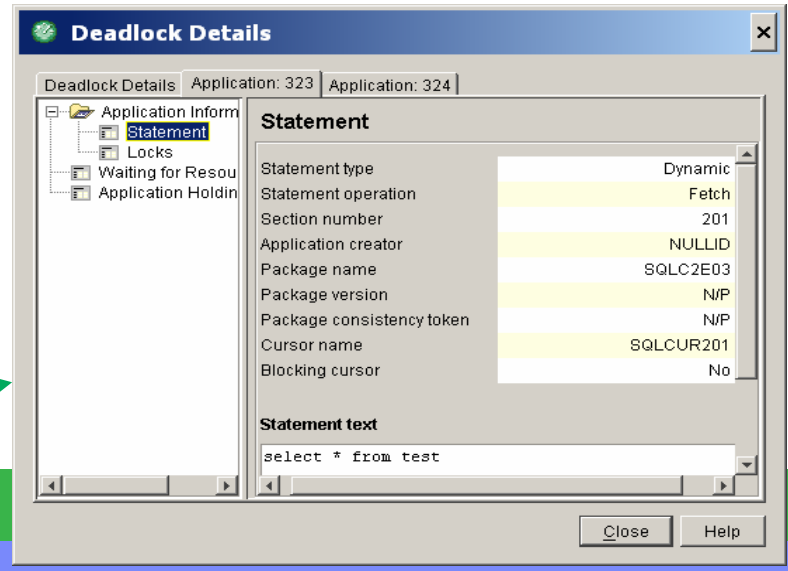

Performance Monitoring --- using future tools

- Autonomic performance tuning in the context of the OS
- DB2 system resources change (e.g. via dynamic LPARs on z/OS)
  - **OS system parameters collected in the warehouse on distributed DB**
    - Display of filesystems, processes, CPU/Memory usage and other
    - User defined exception thresholds on OS system parameters
  - **Providing easier problem analysis and problem prediction**
  - **Complete and detailed analysis of system performance and bottlenecks over time**

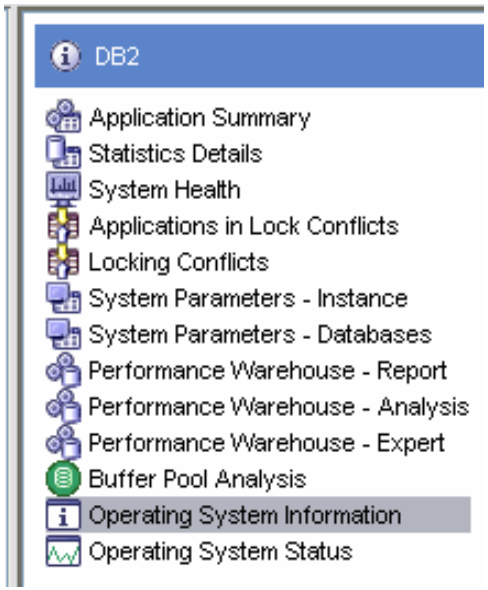

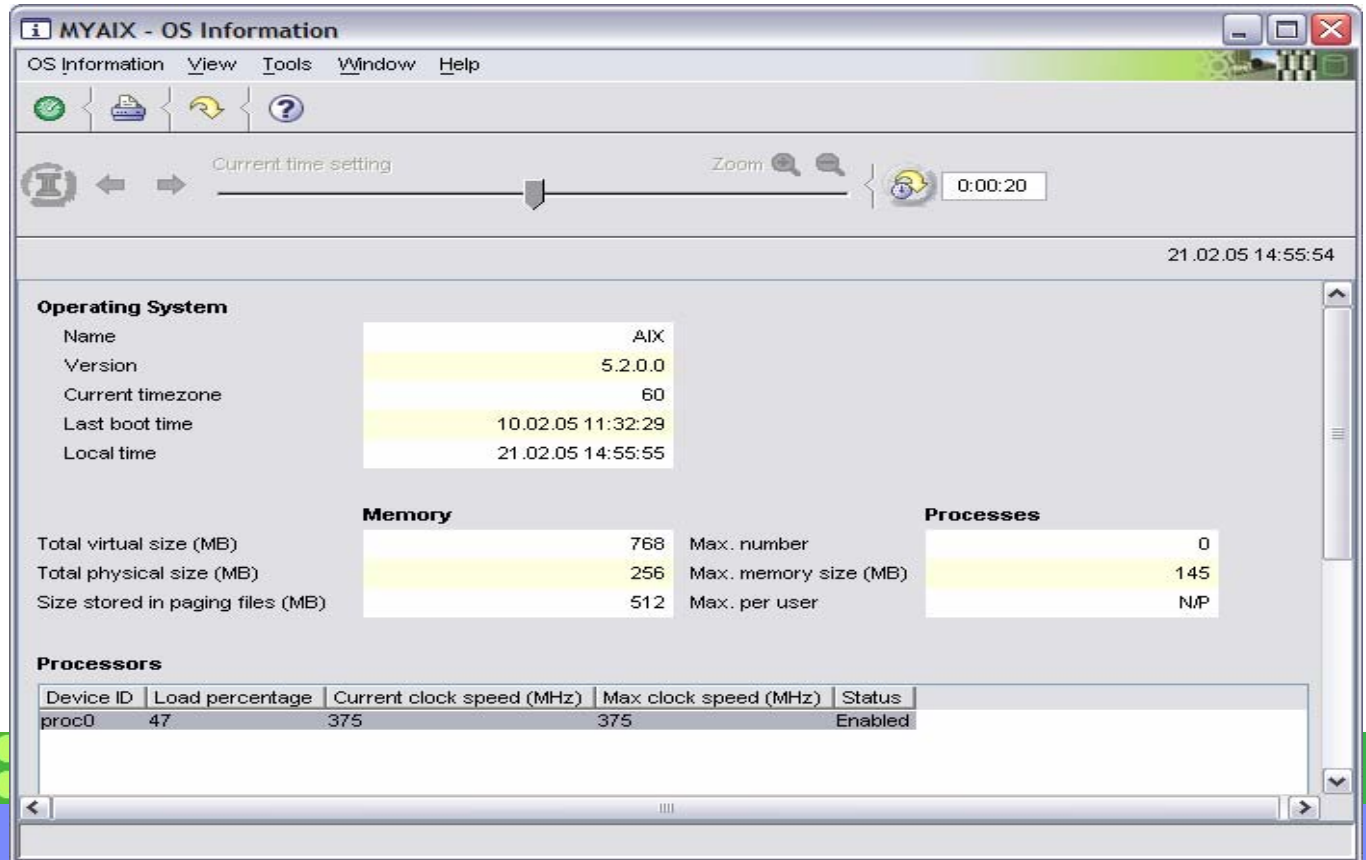

Performance Monitoring --- using future tools

- **Autonomic evaluation of trend data, and capacity planning -- automatic workload characterization**

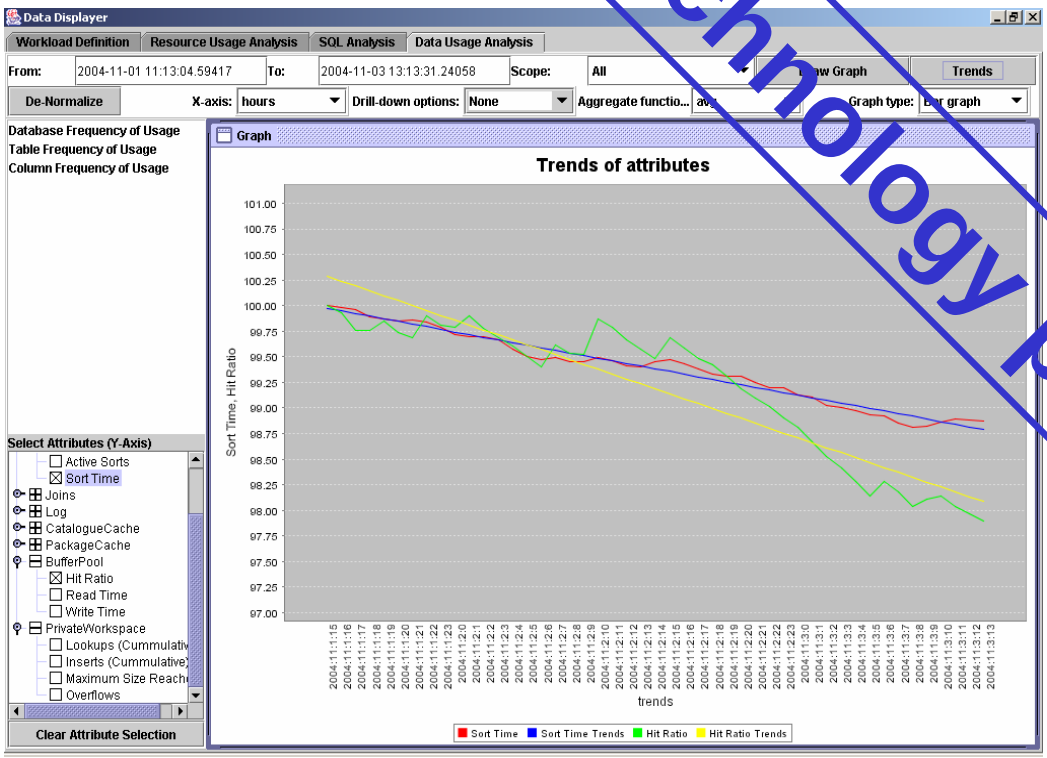

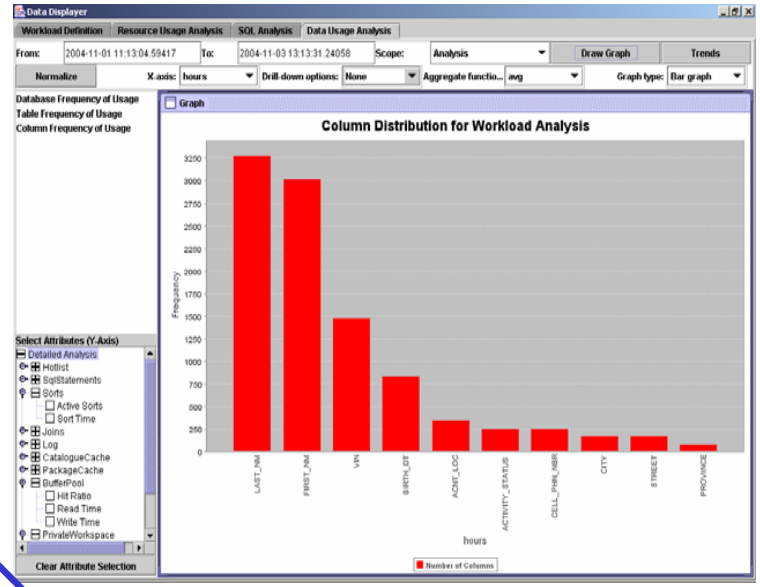

Technology preview

Performance Monitoring --- using future tools

- DB2 Performance Expert extended

- Provides advanced autonomic capabilities
  - Emphasis on exception processing
  - Harvests warehouse data
    - Performance tuning, in the context of the OS
    - Trend analysis & capacity planning with automatic workload characterization
    - Problem diagnosis

  - PE is available today on both DB2 for z/OS and LUW