

**DB2 9 for z/OS**  
Secure Information Services You Can Trust



# DB2 9 for z/OS Overview

*Bart Steegmans - DB2 for z/OS L2 Performance*

GA  
March 16  
2007





*DB2 for z/OS 9.1 GSE Meeting*

*March – 22, 2007*

# Disclaimer and Trademarks

Information contained in this material has not been submitted to any formal IBM review and is distributed on "as is" basis without any warranty either expressed or implied. Measurements data have been obtained in laboratory environment. Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. The use of this information is a customer responsibility.

*IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.*

*TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/390, OS/400, ES/9000, MVS/ESA, Netfinity, RISC, RISC SYSTEM/6000, iSeries, pSeries, xSeries, SYSTEM/390, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, System z, System p*  

*The FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, WINDOWS NT, ODBC, WINDOWS 95*

***For additional information see [ibm.com/legal/copytrade.phtml](http://ibm.com/legal/copytrade.phtml)***

***Many thanks to Jeff Josten for allowing me to use his material***



---

# DB2 9 for z/OS

- ❑ **DB2 9 General Availability March 2007**
- ❑ **Requires Migration to V8 NFM**
- ❑ **IBM Tools V9 Ready at GA**
- ❑ **V7 End of Service Date is June 2008**
- ❑ **Still On V7? Migrate to V8 now**
- ❑ **Plan for V9**



---

# DB2 9 for z/OS 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 System z-specific skill needs**
- Improve data warehousing and OLTP reporting**





# DB2 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, range partitioning
- C** { 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
- L** { 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
- U**
- W**



# DB2 SQL

Z z/OS V9

Common

LUW Linux, Unix & Windows V9



- 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, range partitioning, **TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE**
- C** { 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, RANK, not logged tables**
- L** {
- U** { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET
- W** { CURRENT ISOLATION, multi-site join, MERGE





## DB2 9 for z/OS 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



# 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*

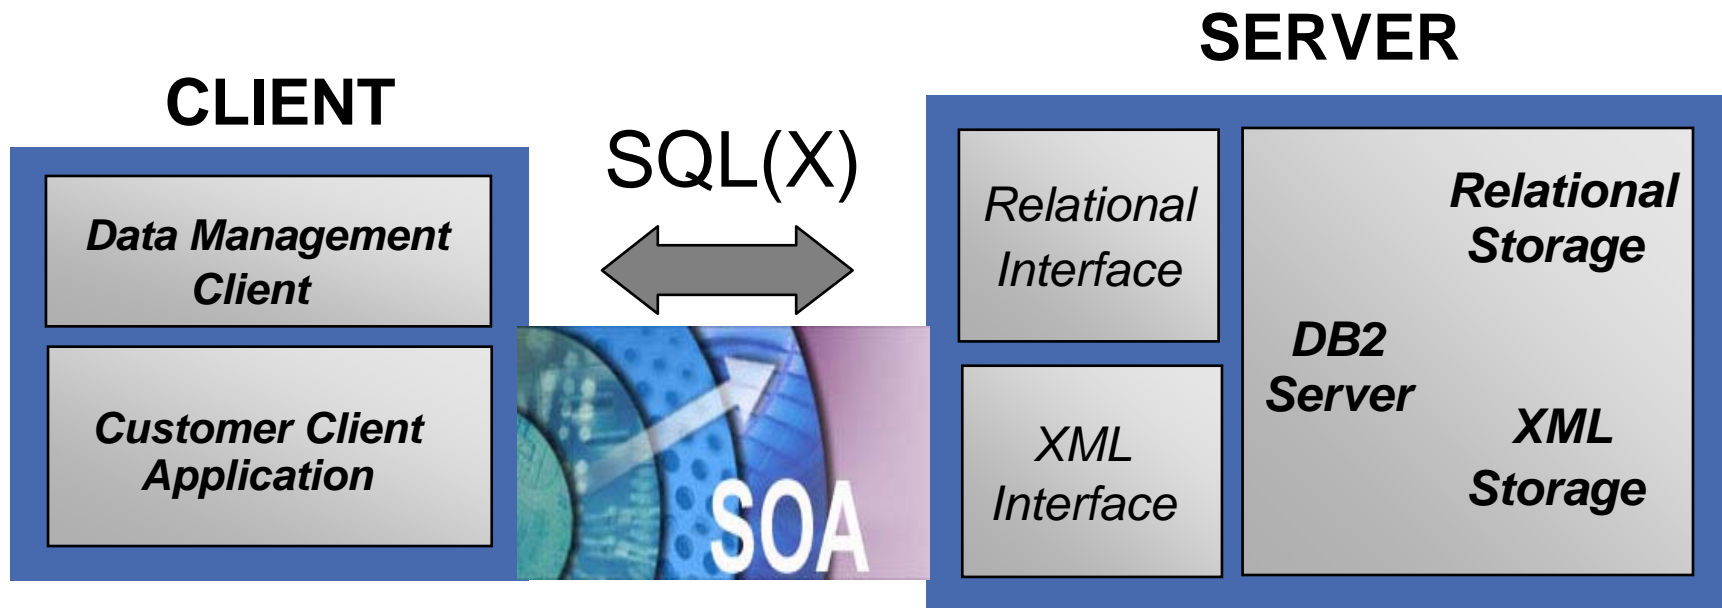
- Full Storage and Indexing
- Powerful Query Capabilities





# Capabilities Inside the Engine

**Performance, Performance, Performance**



Native storage Schema Index Functions Utilities



# pureXML<sup>tm</sup> Support

- Support XML data type
  - Store the XML document natively
  - DDL
    - CREATE/ALTER Table with XML type column
      - Implicitly create XML Auxiliary objects (tablespace/table/index) - one per XML column
    - Index support
      - Created by users
      - uses XPath to determine which nodes in the XML document to index.
- ```
CREATE INDEX dependentName ON deptTable(deptDocs)
GENERATE KEY USING XMLPATTERN
'/department/empl/dependent/name' ATOMIC AS SQL
VARCHAR(20);
```
- INSERT/UPDATE/DELETE
    - INSERT with VALUES and SUBSELECT
    - No Subdocument update



# pureXML Support -- Query

- Enhanced V8 XML Constructors (XML Publishing Functions)
- SQL/XML Functions and Predicates
  - XMLParse - Convert a XML text to XML value
  - XMLSerialize - Converts XML to character type
  - XMLQuery - executes an XPath expression against an XML value.  

```
SELECT XMLQUERY ( '//item[USPrice = $price] ' PASSING  
PO.POrder, T.price AS "price")  
  
FROM PurchaseOrders PO, T;
```
  - XMLCast - Cast XML to other types or other types to XML
  - XMLEXists - a predicate, which returns TRUE if the XPath expression evaluates to a non-empty sequence  

```
SELECT PO.pid  
FROM PurchaseOrders PO, T  
WHERE XMLEXISTS( '//item[USPrice = $price] '  
PASSING PO.POrder, T.price AS "price")
```



# pureXML Support (continued...)

- XPATH supported features from XPath 2.0:
- Utility Support
  - LOAD/UNLOAD, CHECK DATA/INDEX, COPY, REBUILD, RECOVER, REORG, etc.
- XML Schema Support
  - XSR – XML Schema Repository
  - Tables to store XML schemas
  - Stored procedures to register XML schemas
- DSN\_XMLVALIDATE() SQL/XML function
  - Test XML values for validity against XML schema
  - Obtain default values and schema normalized values from XML schema
- XML decomposition using annotated XML schema



# MERGE

- Array MERGE operation
- Targets OLTP applications like SAP

```
MERGE INTO account AS T
USING VALUES (:hv_id, :hv_amt) FOR 5 ROWS AS S(id,amt)
ON T.id = S.id
WHEN MATCHED THEN
    UPDATE SET balance = T.balance + S.amt
WHEN NOT MATCHED THEN
    INSERT (id, balance) VALUES (S.id, S.amt)
NOT ATOMIC CONTINUE ON SQLEXCEPTION
```





# TRUNCATE Statement

- Allows fast delete of all rows in a given table (simple, segmented, or partitioned)
- Very useful for nightly refresh of summary tables, warehouses, etc.

```
TRUNCATE TABLE TABLE-NAME
```

```
< DROP STORAGE | REUSE STORAGE >
```

```
< RESTRICT WHEN DELETE TRIGGERS |  
  IGNORE DELETE TRIGGERS >
```

```
< IMMEDIATE >
```



# 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

▶ **ROW CHANGE** [ **TIMESTAMP** ] **FOR** *table-designator* ◀  
                  [ **TOKEN** ]



---

# SQL Improvements – Family Compatibility

- INSTEAD OF triggers
- SELECT FROM UPDATE
- SELECT FROM DELETE
- SELECT FROM MERGE
- ORDER BY and FETCH FIRST in subselect
- BIGINT and VARBINARY data types



# Decimal Floating Point

- New datatype NUMBER
  - Well suited to typical customer financial calculations
  - Similar to “calculator” mathematics
    - Eliminates rounding errors by using base 10 math
    - Has up to 34 digits of precision
    - Floating point convenience with fixed point precision!!!
  - Hardware support will be provided in the next System z processor generation (new IEEE standard)
    - Software emulation provided for other models



# DDF Improvements

- 64-bit exploitation by DDF
  - Special “shared private” with xxxDBM1 to eliminate many of the data moves on SQL operations
- Support for IPv6 and SSL
- VTAM definition is now optional
- Prepare for elimination of PRIVATE protocol requester
  - Includes tools for identifying which packages need to be bound at remote servers



## DB2 9 Spatial Support

- Enabling OGC (Open Geospatial Consortium) compliant GIS (geospatial) applications
  - Spatial data types
  - Spatial functions and predicates
  - Spatial indexes
  - Spatial search
  - OGC-compliant spatial catalog



---

# DB2 9 for z/OS 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 System z-specific skill needs**
- ❑ **Improve reporting**





---

## Schema Evolution – Database Definition On Demand

- Fast replacement of one table with another
- Rename column and index
- 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
- Create & alter STOGROUP SMS constructs



# CLONE Tables

- Allows fast replacing production data without renames and rebinds
  - A capability to support online load replace
- ALTER 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



# CLONE Tables...

- Switch between base and clone
  - EXCHANGE DATA BETWEEN TABLE base-tab AND clone-tab
- A clone table can only be created
  - On a single table in a table space (partitioned or non-partitioned) – uses UTS
  - No RI or Trigger on the base table
- Use insert or load to populate clone tables
- Utilities (except RUNSTATS) can operate on clone tables with a new CLONE keyword



# 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



---

# CREATE TABLE ... APPEND(YES)

- New APPEND option:
  - Maximizes performance for “INSERT at end”
  - Avoids overhead of attempting to preserve clustering sequence
  - REORG after bulk insert to institute clustering



# Relief for Sequential Key INSERT

- New page sizes: 8K, 16K, 32K for INDEX pages
  - Fewer page splits for long keys
  - More key values per page
- INSERT at the end of the key range used to result in 50% free space in each index page
  - Enhanced support dynamically adapts page split boundary to minimize wasted space in index pages
- Index key randomization



# 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
  - FETCH CONTINUE
- Other LOB Improvements
  - REORG LOB reclaim space, SHRLEVEL(REFERENCE), allow LOG NO
  - Online CHECK LOB and DATA





# V9 Improvements for Vstor Constraint Relief

- DDF address space runs in 64-bit addressing mode
  - Shared 64-bit memory object avoids xmem moves between DBM1 and DDF and improves performance
  - Constraint relief
- DBM1, the following are moved above the bar in V9
  - Parse trees
    - peak below-the-bar storage for full prepare reduced 10%
  - EDM fixed pools
    - V8 customer dumps show as much as 50m will be moved.
    - Allows larger above the bar EDM pools
  - SKPTs / SKCTs (primarily static SQL). Also part of the CTs/PTs
    - New EDM pool for skeltons
    - Savings in below the bar 10m to 300m
  - Pageset blocks, RTS blocks: up to 10's of MB savings
  - Local SQL statement cache: rough ROT: about 60% moves above bar
  - Thread-related storage:
    - Certain RTs, space block, DMTR
    - 10's of MB or more in savings



# Other Performance / Availability Items

- Insert performance (APPEND, INDEX, LOG)
- INDEX on expression, INDEX compression
- Log contention relief for data sharing, archive log striping
- Health monitor task
- Utility TEMPLATE switching
- Enhancements for Point In Time recovery
- MODIFY Recovery enhancements
- Cancel in-progress DATABASE commands
- Data sharing restart availability enhancements
- CPU reductions in LOAD and REORG
- Online REBUILD INDEX
- Online REORG BUILD2 phase elimination
- Intra-REORG parallelism for UNLOAD, RELOAD, LOG phases
- DPSIs unique within partition
- FETCH FIRST n ROWS improvements
  - Can now be specified in a subquery or fullselect
  - ORDER BY now exploits FETCH FIRST n ROWS



---


# DB2 9 for z/OS 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 System z-specific skill needs**
- ❑ **Improve reporting**




# Cost Of Ownership Trends


## UNIX






## INTEL



## Mainframe



 % Personnel  
 % Hardware  
 % Software

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 CFR 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 9 for z/OS

Some key implementations

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



# Protecting data on disk

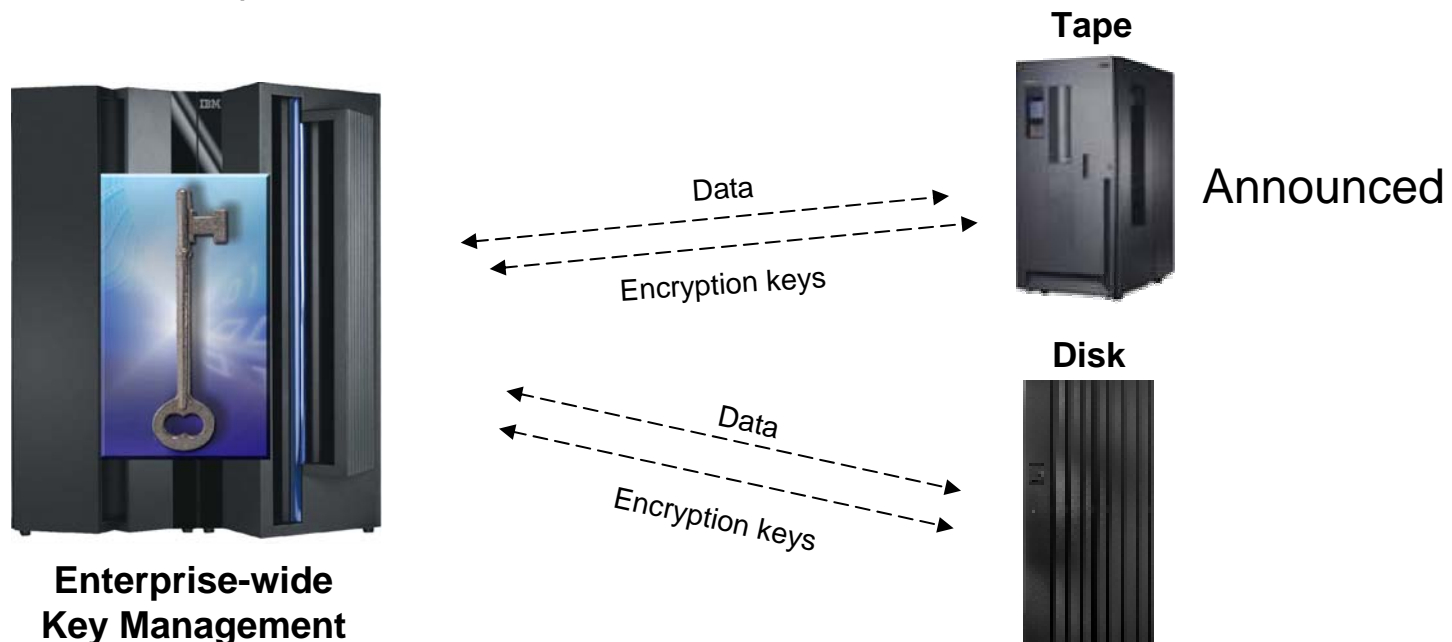
- We will allow encryption for the key disk resources used by DB2:
  - Tables
  - LOBs
  - Indexes
  - Image copies
  - Logs
  - Archive logs





# Future Directions – Extending Encryption to IBM TotalStorage

- Statement of Direction: To address customers' growing concern with data security, IBM is announcing a statement of direction for the development, enhancement and support of encryption capabilities within storage environments such that the capability does not require the use of host server resources.
- This includes the intent to offer, among other things, capabilities for products within the IBM TotalStorage® portfolio to support outboard encryption and to leverage the centralized key management functions planned for z/OS ICSF.

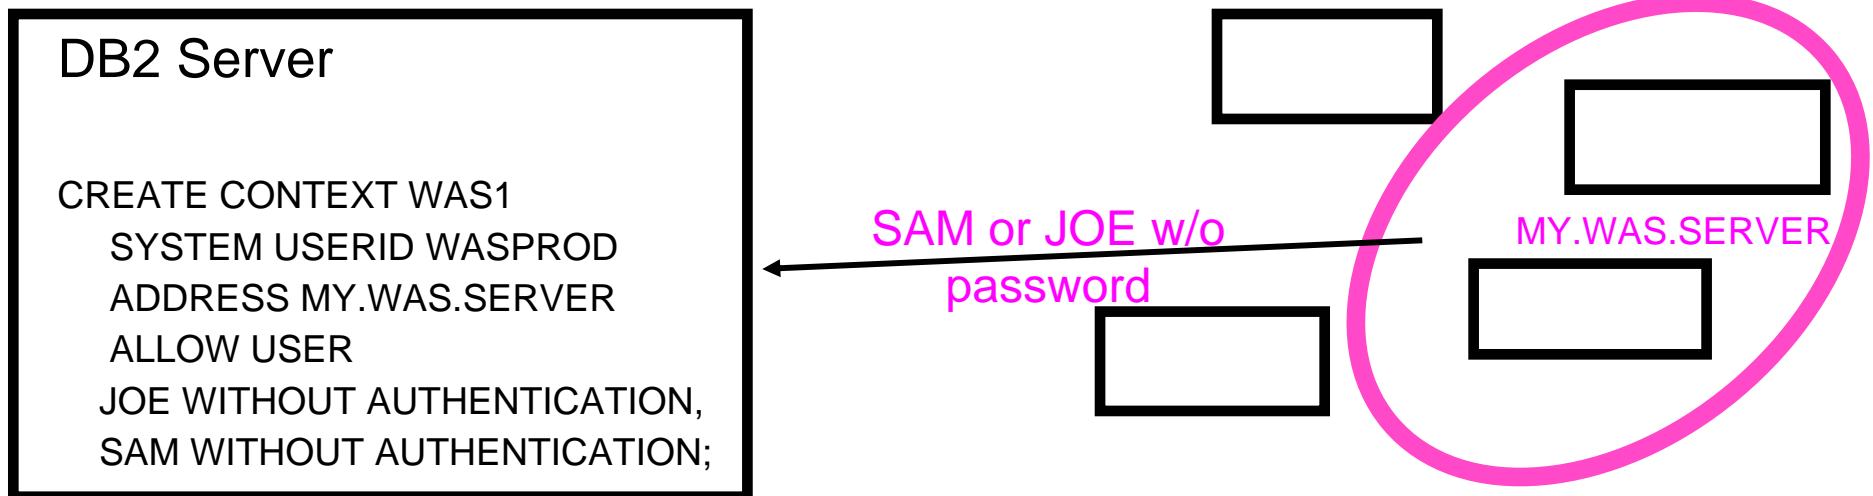


Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only



# Trusted Security Context

- Identifies “trusted” DDF, RRS Attach, or DSN application servers
- Allows selected DB2 authids on connections without passwords
  - Reduces complexity of password management
  - Reduces need for an all-inclusive “system authid” in app servers
  - More visibility/auditability of which user is current running
  - Enables mixed security capabilities from a single app server



# 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);
```



# Auditing: DB2 Trace Filtering

- New filtering capabilities for `–START TRACE` that `INCLUDE` or `EXCLUDE` based on these keywords:
  - `USERID` -- client userid
  - `WRKSTN` -- client workstation name
  - `APPNAME` -- client application name
  - `PKGLOC` -- package `LOCATION` name
  - `PKGCOL` -- package `COLLECTION` name
  - `PKGPROG` -- `PACKAGE` name
  - `CONNID` -- connection ID
  - `CORRID` -- correlation ID
  - `ROLE` – end user's database `ROLE`



# Volume-based COPY/RECOVER

- FlashCopy technology used to capture entire content of disk volumes
- RECOVER modified to enable object-level recovery from volume FlashCopy
  - Restore assumes that the object has not moved volumes
- Eliminates labor associated with setting up COPY jobs for each database / table space
- Full integration of tape into BACKUP/RESTORE SYSTEM utilities



# Converged TEMP Space

- Single source for all temporary space in DB2, replacing: DSNDB07, temp databases, workfile database
- Access is virtualized for small amounts of data, eliminating cost of work file creation (reduced CPU and I/O)
- Supports 4K and 32K page sizes, with automatic selection of the appropriate page size
- New Zparm for preventing workfile monopolization



# DDL Porting Improvements

- Automatic selection of DATABASE and TABLESPACE when DDL omits these keywords
- Automatic CREATE of UNIQUE index for PRIMARY KEY
- Deprecated simple table space, default to partition by growth (segmented structure, partitioned without a partitioning key)



# Other cost of ownership improvements

- Resource Limit Facility
  - RLF will be enhanced to allow CPU cost to be controlled based on:
    - Client workstation name
    - Client application name
    - Client userid
    - IP address
  - Allows fine-grain RLF limits for apps that exploit the Set Client Information APIs (SAP, PeopleSoft, Siebel, WebSphere)
- DBM1 and DDF virtual storage relief
- SMS integration
- RENAME SCHEMA, VCAT
- Utilities template switching support
- Optimization Service Center





---

# DB2 9 for z/OS 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 System z-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
- REOPT AUTO
- EXCEPT and INTERSECT
- RANK, DENSERANK, ROWNUMBER
- Many other SQL improvements



# Index Compression

- Compression of indexes for BI workloads
  - Indexes are often larger than tables in BI
- Solution provides page-level compression
  - Data is compressed to 4K pages on disk
  - 32K/16K/8K pages results in 8x/4x/2x disk savings
  - No compression dictionaries – compression on the fly



# TCO Improvements – DBA tools

- Autonomic Policy-based SQL query management/monitoring:
  - Automatic collection of performance data for long running queries
  - Automated query monitoring for the most frequent/expensive queries
  - REOPT(AUTO)
- Optimization Service Center (Web-based DBA admin – no 3270 screens)
  - DBA tool suite for tuning/managing SQL queries (Stats Advisor, Index Advisor, Query Rewrite Advisor, Query Workload Monitor, Resource Estimator, Query Formatter, Visual Explain, Visual Plan Hint, IBM Service Doc Generator, Partitioning/Clustering Advisor)
- Query Performance Warehouse
  - Execution history of queries
  - Identification of query patterns
  - Identification of usage patterns for tables/indexes



# V9 Customer Beta Feedback Highlights


- “Ready for GA”
- Install/migration smooth, no problems
  - Shorter, less intimidating than V8, CM\* is good
- Index compression: 50%+ space savings
- LOAD/REORG cpu reduction
- LOB performance / mgmt improvements are key
- Clones worked well
- Vstor savings will be important
- Not Logged: “be careful!!”
- High interest in XML
- Some are deploying as soon as GA



# Beyond DB2 9



❑ TCO (improved performance, reduced people cost)



❑ Availability improvements (fewer planned outages)



❑ Data warehousing improvements

❑ Application development and application portability (SQL, XML)



# Redbook coming soon

❑ **SG24-7330**

Information Management software

Draft Document for Review March 19, 2007 11:29 am

**IBM**


SG24-7330-00

## **DB2 9 for z/OS Technical Overview**

Read about the new features and functions of DB2 Version 9.1 for z/OS

Understand prerequisites and migration process

Evaluate applicability and major returns



Paolo Bruni  
Roy Cornford  
Rafael Garcia  
Sabine Kaschta  
Ravi Kumar

[ibm.com/redbooks](http://ibm.com/redbooks)

**Redbooks**

