



DB2 9 for z/OS Is Here: Part 2 – A Deeper Dive



*Gareth Jones – jonesgth@uk.ibm.com
Florence Dubois – fldubois@uk.ibm.com*

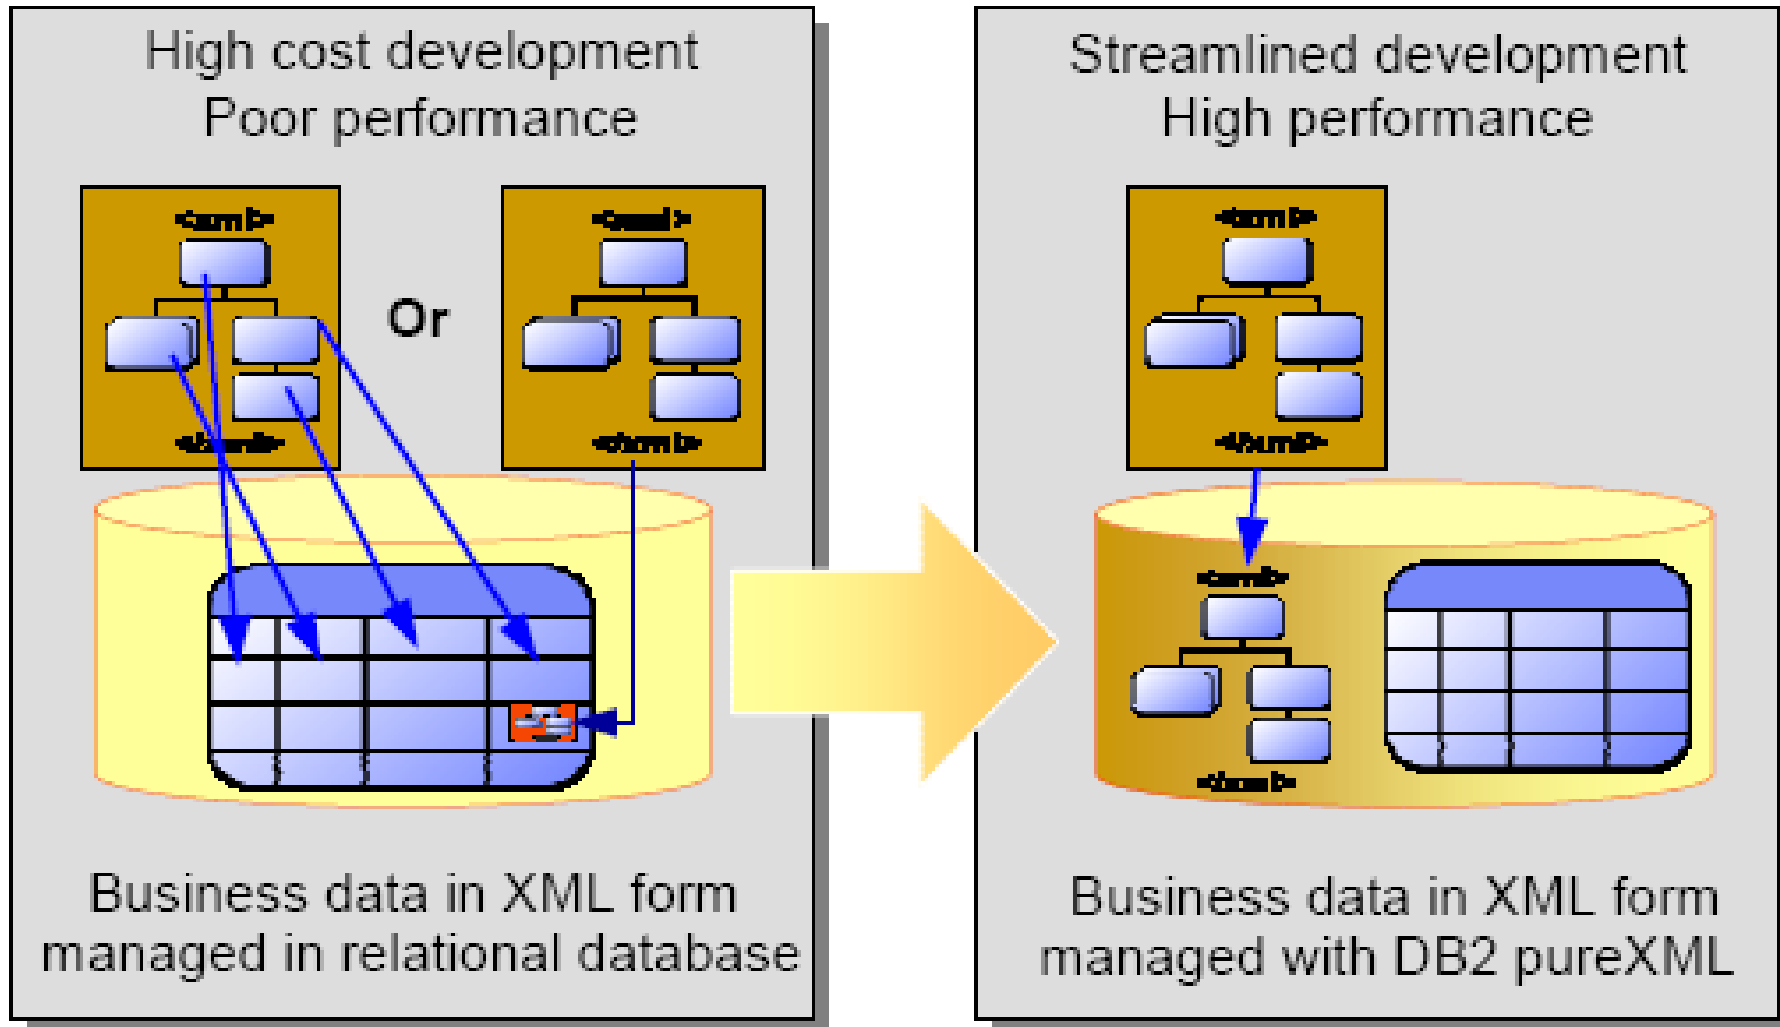
TAKE BACK CONTROL

DB2 9 for z/OS – Addressing corporate goals

- **Evolve Your Environment & SOA**
 - Integrated pureXML®
 - WebSphere® integration
- **Improved IT Infrastructure for Compliance Efforts**
 - Trusted security context
 - Database roles
 - Auditing, encryption improved
- **Simplify development and porting**
 - Many SQL improvements
 - Native SQL stored procedures
 - DDL porting improvements
- **Data Warehousing**
 - SQL enhancements
 - Query optimization improvements
- **Decrease Complexity and Cost**
 - Performance improvements
 - Volume-based backup/recovery
 - Index compression
 - Optimization Service Center
- **Continuous Availability**
 - Schema evolution enhancements
 - Partition by growth
 - Fast table replacement



DB2 9: A New Generation Hybrid Data Server



DB2 9 – Summary of pureXMLTM Support

- XML as a native data type
- Pure XML storage and indexing
- SQL/XML and XPath support
- Integration with traditional relational data
- XML Schema Repository
- Schema validation
- Application Support (Java, C/C++, .NET, PHP, COBOL, PL/1 etc.)
- Visual Tooling, Control Center Enhancements
- DB2 Utilities: Load, Unload, Reorg, etc.
- ...and more

DB2 9

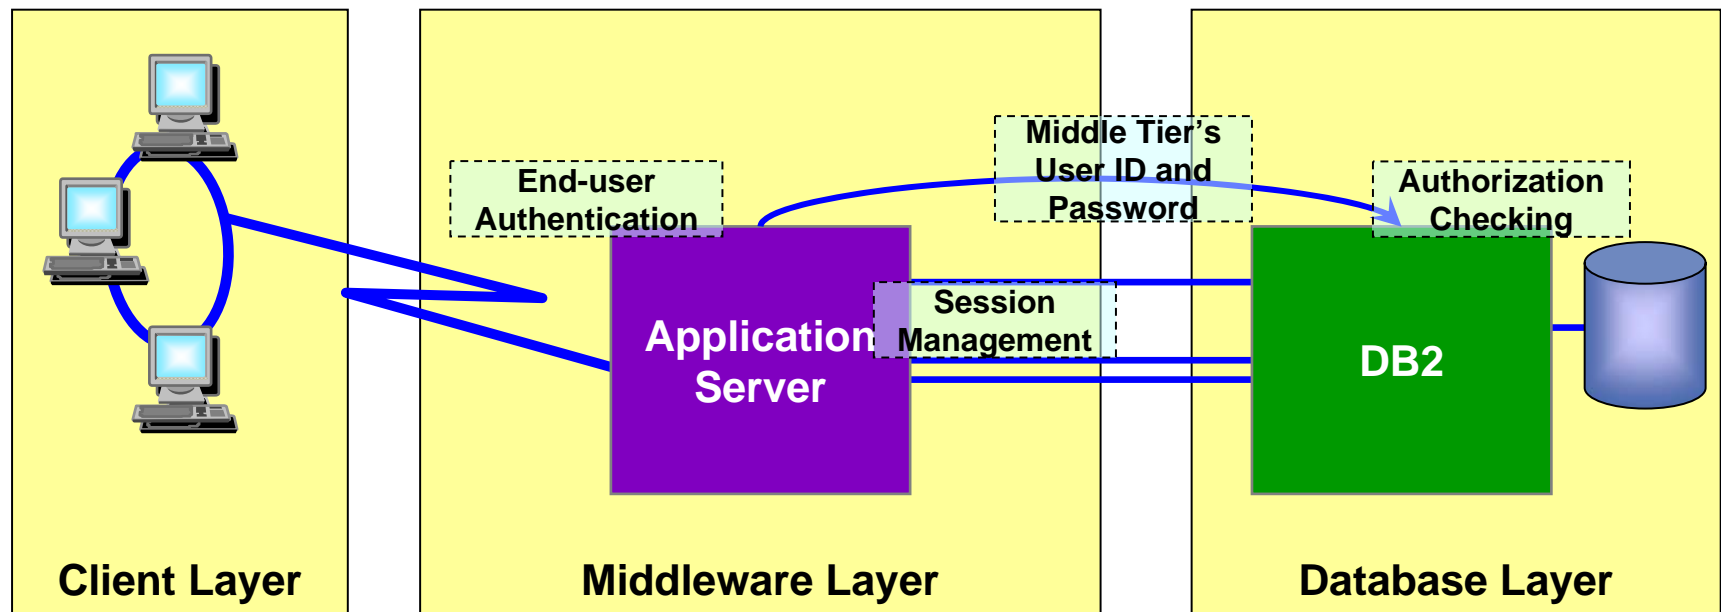
Secure and
Resilient
Infrastructure
for a New
Breed of Agile
Applications



*"This is not a bolt-on or band-aid approach,
DB2 9 for z/OS is XML without compromise"*

Kevin Campbell, Application Architect, Univar USA

Authentication in a Three-Tier Architecture



- A three-tiered application model with DB2 as the database server:
 - The middle layer authenticates users running client applications.
 - It also manages interactions with the database server.
 - The middle layer's user ID and password are used for database authentication.
 - The privileges of the associated authorization id are checked when accessing the database, including all access on behalf of all end-users.



Trusted Contexts

- A ***TRUSTED CONTEXT*** establishes a trusted relationship between DB2 and an external entity such as a middleware server. For example:
 - WebSphere Application Server
 - Lotus Domino
 - SAP NetWeaver
 - PeopleSoft V7
- A set of *trust attributes* is evaluated to determine if a specific context is to be trusted.
- A trusted context allows the external entity to use a database connection under a different user ID without the database server authenticating that ID.
- It also allows an AUTHID to acquire database privileges associated with that trusted context, and not available outside it, via a ***ROLE***.



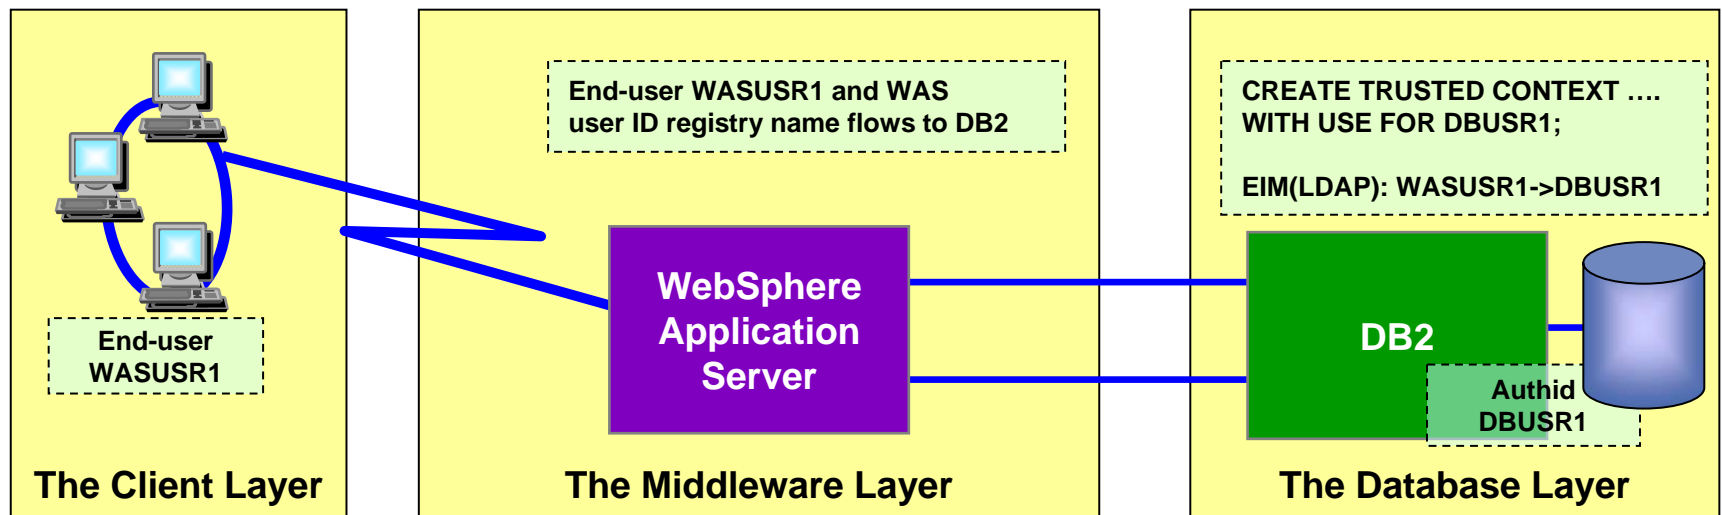
Roles and Context-specific Privileges

- **Roles** provide the flexibility to grant privileges to an AUTHID only when the user is connected via a trusted context.
- They greatly simplify management of authorization.
- An individual **role** can be defined for any AUTHID using the trusted connection, in which case the user inherits the privileges granted to the individual **role**.
- Where there is no individual **role**, any AUTHID using a trusted context inherits the privileges of the trusted context's default **role**, if defined.

```
CREATE TRUSTED CONTEXT CTX1  
BASED UPON CONNECTION USING SYSTEM AUTHID WASADM1  
DEFAULT ROLE CTXROLE  
ATTRIBUTES (ADDRESS '9.67.40.219')  
ENABLE  
WITH USE FOR JOE ROLE JROLE;
```



WebSphere Integration – Propagating Client Identity



- Configure WebSphere app server to create trusted connection and send the registry name to the server
 - Database property 'propagateClientIdentityUsingTrustedContext' is set to 'true'
 - Application parameter 'TargetRealmName' is set to the registry name at the server
- DB2 maps the registry name and client user ID to obtain DB2 auth ID
- DB2 checks if the DB2 auth ID is allowed to use the trusted connection
- Client user ID, WAS security token & DB2 auth ID are recorded in DB2 audit logs
- Client user ID and DB2 auth ID are recorded in RACF audit logs



DB2 9 for z/OS Innovation: SQL

- ❑ Numerous new SQL capabilities
- ❑ Easier application porting
- ❑ Simplified application development



SQL: Productivity, DB2 family & porting




- XML
- MERGE & TRUNCATE
- SELECT FROM UPDATE, DELETE, MERGE
- INSTEAD OF TRIGGER
- BIGINT, VARBINARY, BINARY, DECIMAL FLOAT
- Native SQL Procedure Language
- Nested compound
- 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



TRUNCATE Statement

- Allows fast delete of all rows in a given table (segmented, partitioned or simple)
- 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 >
```



MERGE

- Multi-row MERGE operation, using arrays
- 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
```



SQL Improvements – Family Compatibility

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



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 segmented structure, partition by growth




Query Enhancements


- SQL enhancements
 - INTERSECT, EXCEPT, cultural sort, caseless comparisons, RANK, DENSE_RANK, ROW_NUMBER ...
- Index improvements
 - Index on expression, Index compression, ...
 - Ability to identify unused indexes
- Optimization techniques
 - Histogram Statistics
 - Cross query block optimization and REOPT(AUTO)
 - Generalize sparse index & in-memory data cache method
 - Dynamic Index ANDing for Star Schema




Intersect/Except/Union



INTERSECT



EXCEPT



UNION /
UNION ALL


** There are some variations and restrictions*

** UNION ALL includes duplicates*




Index Changes

- INDEX on expression
- Ability to identify unused indexes
- Page sizes 8K, 16K, 32K
- Improved page split
- Index compression
- Online REBUILD INDEX
- REORG without BUILD2 – not just for DPSI
- DPSI can be unique when superset of partitioning key
- Randomized index key
- Not logged index space
- XML index



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
 - 8K, 16K or 32K pages results in 2x, 4x or 8x disk savings
 - No compression dictionaries – compression on the fly



Index Compression: Differences between data and index compression

	Data	Index
Level	Row	Page (1)
Comp on disk	Yes	Yes
Comp in Buffer Pool	Yes	No
Comp in Log	Yes	No
Comp Dictionary	Yes	No (2)
Average Comp Ratio	10% - 90%	25% - 75% (3)



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




CREATE TABLE ... APPEND(YES)

- New APPEND option:
 - Maximizes performance for “INSERT at end”
 - Avoids overhead of attempting to preserve clustering sequence
 - CREATE or ALTER table



Varchar Performance Improvement

- Old tuning recommendation for rows with many columns with any varchar present



- V9 DB2 internally executes this recommendation and more
- 2 times or more improvement observed when many rows with many varchars are scanned and/or fetched using many predicates
- No difference if no varchar , Under 5% improvement for a typical online transaction
- Reorg with rebuild compression dictionary if varchar columns when migrating to V9



LOB Performance/Scalability

- LOB lock avoidance – 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



Usability Improvements

- Converged temp space
- WLM-assisted buffer pool sizing
- Migration/fallback
- Improved virtual storage usage
- Others...

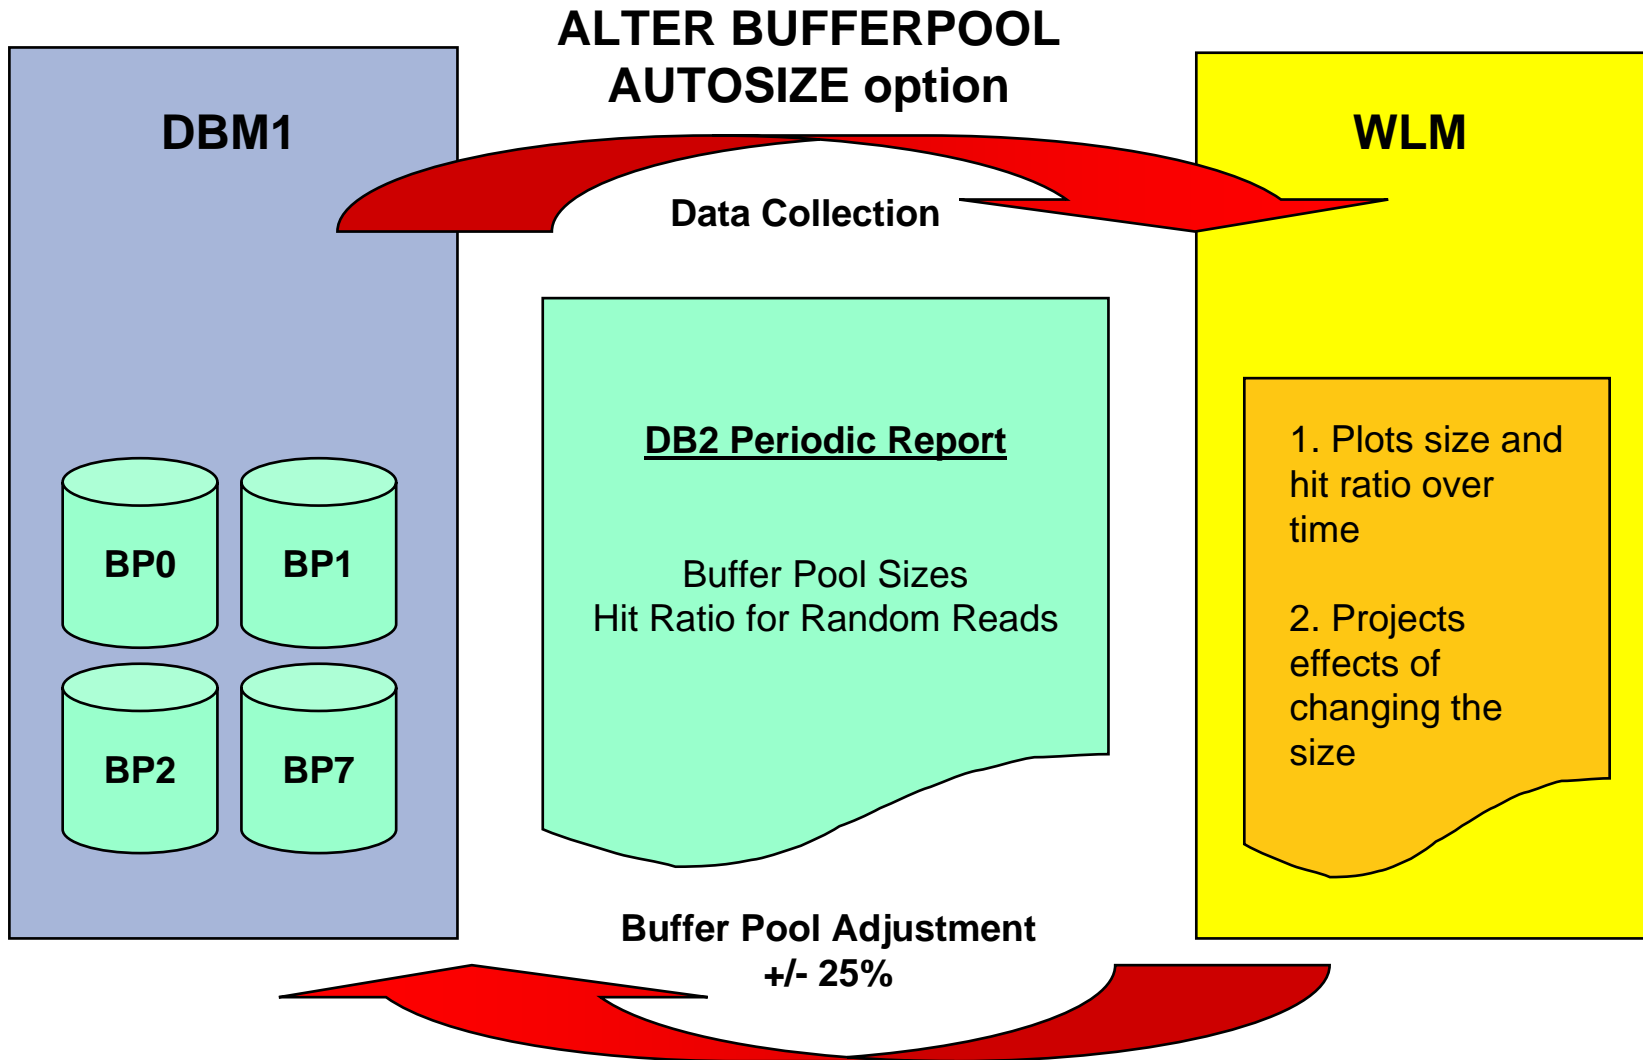


Converged TEMP Space

- 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



WLM assisted buffer pool management



V9 Modes – An Overview

CM Compatibility Mode - This is the mode DB2 is in when V9 is started for the first time from V8. It will still be in CM when migration job DSNTIJTC has completed. No new function can be executed in CM. Data sharing systems can have V8 and V9 members in this mode. DB2 can only migrate to CM from V8 NFM.

ENFM Enabling New Function Mode - This mode is entered when CATENFM START is executed (the first step of job DSNTIJEN). DB2 remains in this mode until all the enabling functions are completed. Data sharing systems can only have V9 members in this mode.

NFM New Function Mode - This mode is entered when CATENFM COMPLETE is executed (the only step of job DSNTIJNF). This mode indicates that all catalog changes are complete and new function can be used.

ENFM* This is the same as ENFM but the * indicates that at one time DB2 was at NFM. Objects that were created when the system was at NFM can still be accessed but no new objects can be created. When the system is in ENFM* it can not fallback to V8 or coexist with a V8 system.

CM* This is the same as CM but the * indicates that at one time DB2 was at a higher level. Objects that were created at the higher level can still be accessed. When DB2 is in CM* it can not fallback to V8 or coexist with a V8 system.



DB2 9 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
 - EDM fixed pools
 - SKPTs / SKCTs (primarily static SQL). Also part of CTs/PTs
 - Pageset blocks, RTS blocks
 - Local SQL statement cache
 - Some thread-related storage
- For installations that are constrained on DBM1 VSTOR:
 - 200 to 300MB of savings expected with rebinds, but mileage will vary
 - Mainly from EDM related storage (static SQL) and dynamic statement cache (dynamic SQL)



Other cost of ownership improvements

- Resource Limit Facility enhanced to allow CPU cost to be controlled based on:
 - Client workstation / app name
 - Client userid, IP address
- SMS integration
- Utilities template switching
- RENAME SCHEMA, VCAT
- REOPT(AUTO)



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
- ALTER TABLE EXCHANGE to switch base to clone
 - Provides load replace (shrlevel change) functionality



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



DB2 9 Utilities

- Always perform CHECKPAGE on the COPY utility
- More online utilities
 - Rebuild Index SHRLEVEL CHANGE
 - Great for building new non-unique indexes
 - Reorg enhancements
 - Reorg LOB now supports SHRLEVEL REFERENCE
 - LOB space reclamation
 - Partition-level capabilities (not available with REBALANCE)
 - Partition parallelism (UNLOAD/RELOAD) in a single utility statement
 - Elimination of the BUILD2 phase outage
 - Check data, LOB and repair locate ... SHRLEVEL CHANGE
 - Check index SHRLEVEL REFERENCE supports parallel for > 1 index



DB2 9 Data Sharing Enhancements

- Restart performance enhancements
 - Reduced impact of retained locks
 - Open datasets ahead of log apply
 - Avoid locks that have potential of ‘hanging’ restart
- Auto-recover GRECP/LPL objects on group restart
 - Useful in Disaster Recovery or GDPS scenarios
- Index performance improvements
 - Option to randomize key
 - Sequential key insert performance improvement
 - Larger page size, index compression
- Use of DPSI to avoid data sharing overhead
 - DPSI can be unique within partition
- Detect use of unused indexes so they can be dropped
 - Reduces data sharing overhead



DB2 9 Data Sharing Enhancements...

- Command to remove GBP-dependency at object level
 - ACCESS DB MODE(NGBPDEP)
 - Typical usage would be before batch run
 - Issue on the member on which you plan to run batch
- Command to “prime” open dataset
 - START DB MODE(OPEN) [PART]
- Log latch contention relief (avoid LRSN “spinning”)
- Improved performance for GBP writes
 - Avoid copying pages for batched writes
- DB2 overall health taken into account for WLM routing
- Balance group attach connections across multiple members on same LPAR (V7, V8 usermod)



References



TAKE BACK CONTROL

DB2 9 for z/OS References

Main DB2 for z/OS web page: pointers to most of the following:

<http://www.ibm.com/software/data/db2/zos/index.html>

DB2 9 for z/OS main page:

<http://www.ibm.com/software/data/db2/zos/db2zosv91.html>

Overview presentation, webcast and foils with notes:

<http://www.ibm.com/software/os/zseries/webcast/18may/>

<ftp://ftp.software.ibm.com/software/data/db2zos/DB2V9zOS.pdf>

Redbooks including V9:

- SOA book - overview of V9 XML SG24-7259
<http://www.redbooks.ibm.com/abstracts/SG247259.html?Open>
- LOBs book SG24-7270 <http://www.redbooks.ibm.com/abstracts/SG247270.html?Open>
- V9 Technical Overview SG24-7330 <http://www.redbooks.ibm.com/redbooks/pdfs/sg247330.pdf>
- V9 Performance Topics <http://www.redbooks.ibm.com/redbooks/pdfs/sg247473.pdf>

DB2 9 What's New

<http://publib.boulder.ibm.com/epubs/pdf/dsnwnk10.pdf>

Detailed presentations:

- Start on the Events page <http://www.ibm.com/software/data/db2/zos/events.html>
- Click on Presentations from previous conferences. Sort results by date - newest first.
- Access the ftp site directly: <ftp://ftp.software.ibm.com/software/data/db2zos/>
- About 15 of more than 200 presentations that address V9.

Presentations from IOD conference, IDUG, Share, ...



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, zSeries, System z, pureXML

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

