

May 6-10, 2007  
San Jose Convention Center  
San Jose, California, USA

Session: X01

# DB2 9 for z/OS and Beyond

IDUG® 2007  
North America

Jeff Josten  
*IBM DB2 for z/OS Development*

May 10, 2007 12:40 p.m. – 2:10 p.m.

Platform: DB2 for z/OS

GoFurther




This session will provide insight into DB2 9 and new technologies that are being developed for future releases of the DB2 for z/OS product. Special focus will be given to technical advances that will assist customers in managing data privacy, security, regulatory compliance, porting SQL across multiple platforms, and the scalability/availability/performance improvements that customers are looking for in DB2.

## 5 Bullet Points

- DB2 9 for z/OS SQL enhancements
- Performance, availability, scalability enhancements
- Ease-of-use enhancements
- Improvements for data warehousing or reporting
- Vnext quick preview

## DB2 9 for z/OS – Addressing corporate data goals

- 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
  - Default databases and table spaces
- Data Warehousing
  - Dynamic index ANDing for star schema
  - EXCEPT and INTERSECT
- Decrease Complexity and Cost
  - Partition by growth
  - Performance improvements
  - Volume-based COPY/RECOVER
  - Optimization Service Center
  - System z Synergy
- Evolve Your Environment & SOA
  - Integrated pureXML®
  - WebSphere® integration
- Continuous Availability
  - Schema evolution enhancements
  - Fast table replacement



## DB2 9 for z/OS Innovation: Regulatory Compliance

- Trusted Security Context**
- Database Roles**
- Improved Auditing**
- Encryption improvements**
- Secure Socket Layer**
- Instead of triggers**



GoFurther

4

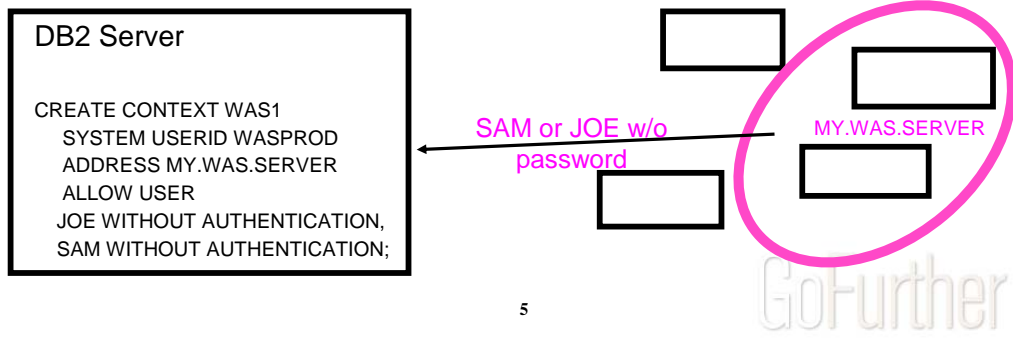
**The second category of innovation is**

### **COST SAVINGS THROUGH OPTIMIZATION**

- o Increased security and regulatory compliance through implementation of roles, network-trusted contexts, and enhanced auditing
- o Performance-boosting innovations such as load and reorg CPU reductions, improved performance for varying length data, and improved logging and insert performance
- o Synergy with IBM System z and z/OS in areas that include XML parsing, zIIP, MDAW channel improvements, encryption, IPv6 and Secure Socket Layer (SSL)
- o Query management enhancements to make accessing your data even faster and more accurate with indexing improvements that include index on expression, randomization, and larger index page sizes and optimization improvements that provide better data for the optimizer, improved optimization techniques, and better management with optimization services

## 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 with ability to switch users
  - more visibility/auditability of which user is current running
  - enables mixed security capabilities from a single app server



**Trusted security context:** Today, you have the option to set a system parameter which indicates to DB2 that all connections are to be trusted. It is unlikely that all connection types, such as DRDA, RRS, TSO, and batch, from all sources will fit into this category. It is likely that only a subset of connection requests for any type and source may be trusted or that you want to restrict trusted connections to a specific server. More granular flexibility will allow for the definition of *trusted connection objects*.

Once defined, connections from specific users via defined attachments and source servers will allow trusted connections to DB2. The users defined in this context can also be defined to obtain a *database role*.

## 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);
```

6

GoFurther

**Database role:** A database role is a virtual authorization ID that is assigned to the user via the context mentioned next. DB2 privileges are assigned to the defined role.

The role exists as an object independent of its creator, so creation of the role does not produce a dependency on its creator.

This capability can allow a DBA to have privileges to create objects and manage them for a time, even though ownership is to be another id.

The role can be assigned and removed from individuals via the trusted authorization context as needed. This allows a DBA to perform object maintenance during a change control window on a Saturday night, for example. But when Monday arrives, they do not have the authority to do this same work.

Auditing trails of the work completed during the maintenance window are available for verification by a security administrator or auditor.

## 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

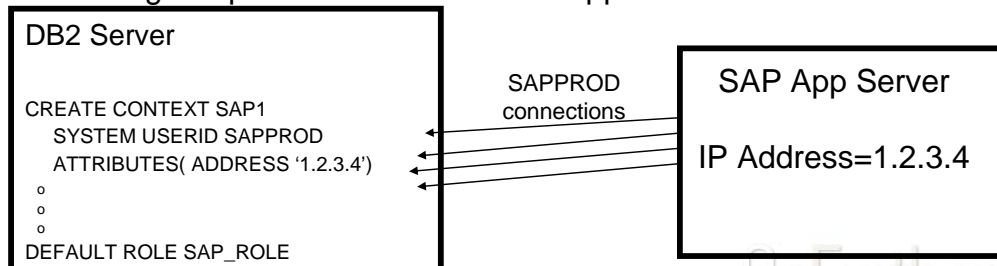
7

GoFurther

Improved trace filtering makes the jobs of auditing and of performance management easier. Many more options can be used to minimize the amount of data collected, so the overhead is reduced and the extraneous data does not need to be processed. Being able to include or exclude, based upon more qualifications reduces performance concerns and the amount of data to store and process.

## Example 1: ROLES and Trusted Context used to Secure App Servers

- Most existing application servers connect to DB2 using userid/password pairs:
  - Significant exposure if someone steals the userid/password!!!
- Trusted Context and ROLES can be used to limit exposure:
  - GRANTs to SAP\_ROLE can be restricted so that they are only valid when used by a valid SAP app server IP address
- No change required to the code in the application server




ROLES and Trusted Context can be used to provide added security for your network-attached application servers. These new capabilities allow the DBA to make GRANT statements conditional, so that they can only be used from a specified list of IP addresses. If someone steals the application server's userid/password, they won't be able to access the database unless they are also able to execute the SQL statement on one of the approved application servers.



## Example 2: ROLES and Trusted Context for Dynamic SQL Auditing

- Better auditing controls:
  - GRANT dynamic SQL privileges to a ROLE
  - End user identity can be delegated directly to DB2 without granting dynamic SQL privileges directly to the end user
  - End user passwords can be optional.
  - No added complexity for administration of GRANTS, while retaining the ability to audit the end user's identity!!!



ROLES and Trusted Context also enable customers to improve DB2 system auditing. Today, many customers use a “system” userid to access DB2 so that they don’t have to grant dynamic SQL privileges to their end users. A second benefit to the system userid is connection pooling in the application server. With V9, customers will be able to grant dynamic SQL table privileges to a ROLE, and specify that the end user can only use that ROLE when the end user is running on an approved application server. This has several benefits:

- A ROLE can be used as a single database authid that can be used to simplify administration of dynamic SQL privileges.
- The end user’s authid can be used to run database transactions, so that the DB2 audit is able to identify the end users individually (important capability for meeting some regulatory compliance requirements).
- The Trusted Context retains many of the performance benefits of connection pooling, while eliminating the restriction that a single authid (the system authid) must be used for all the uses of the connections in the pool.

## 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. **Tape TS1120**



**Enterprise-wide  
Key Management**

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



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.

The first change comes in the TS1120 tape drive, but the rest is beyond currently announced products, including DB2 V9.

The IBM System Storage TS1120 Tape Drive has been enhanced to provide the customer the option of using drive based data encryption. This encryption capability is now standard on all new TS1120 Tape Drives and is a chargeable upgrade feature for existing installed TS1120 Tape Drives. The encryption capability includes drive hardware as well as microcode additions and changes. Also being introduced is a new, separate IBM Encryption Key Manager component for the Java Platform(TM) program that supports the generation and communication of encryption keys for the tape drives across the enterprise.

The TS1120 based encryption and associated Encryption Key Manager component are supported in a wide variety of operating system environments including z/OS, i5/OS, AIX, HP, Sun, Linux and Windows. In addition, three different encryption management methods are supported: Application, System, or Library Managed. This encryption capability is supported when the TS1120 Tape Drive is integrated or attaches in the IBM System Storage TS3500 Tape Library, IBM System Storage TS1120 Tape Controller Model C06, IBM TotalStorage 3592 Tape Controller Model J70, IBM TotalStorage 3494 Tape Libraries, IBM TotalStorage C20 Silo Attach frame, and standalone environments. For more information on encryption please see:

[http://www.ibm.com/servers/eserver/zseries/zos/pdf/White\\_Paper\\_ESG\\_System\\_z\\_final.pdf](http://www.ibm.com/servers/eserver/zseries/zos/pdf/White_Paper_ESG_System_z_final.pdf)

## DB2 9 for z/OS Innovation: SQL

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



11

GoFurther

### **The second category of innovation is**

#### **COST SAVINGS THROUGH OPTIMIZATION**

- o Increased security and regulatory compliance through implementation of roles, network-trusted contexts, and enhanced auditing
- o Performance-boosting innovations such as load and reorg CPU reductions, improved performance for varying length data, and improved logging and insert performance
- o Synergy with IBM System z and z/OS in areas that include XML parsing, zIIP, MDAW channel improvements, encryption, IPv6 and Secure Socket Layer (SSL)
- o Query management enhancements to make accessing your data even faster and more accurate with indexing improvements that include index on expression, randomization, and larger index page sizes and optimization improvements that provide better data for the optimizer, improved optimization techniques, and better management with optimization services

## 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
- 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, 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 or DELETE, multi-site join, MDC

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS, comparing the z/OS Version 8 from March 2004 with the LUW version from October 2004.


There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. Sheryl Larsen provided the base for this information, but the mistakes are probably mine.

If you want to improve DB2 family consistency, then DB2 for z/OS Version 8 is a big step, changing the game from one of catch up to one of leapfrog.

IDUG® 2007 North America

## 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, **TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE, SELECT from 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, 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 or DELETE, 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, range partitioning, compression**

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

**u** {

**w** {

13

GoFurther

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 V9 for z/OS, (V9). V9 moves about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. At about this time we'll also have a new release of DB2 V9.1 for LUW, code named Viper. We are able to move more from the z list to the common list with Viper.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group.

Sheryl Larsen provided the base for this information, but the mistakes are probably mine.

## SQL: Productivity, DB2 family &amp; 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

As in Version 8, there are many improvements for SQL and for XML in V9. Improvements in the SQL have made migrating from other platforms, such as Unix and Windows much easier.

V9 continues the progress in SQL, with many new functions, statements and clauses. The biggest changes are in XML on the prior slide. There are new SQL data manipulation statements in MERGE and TRUNCATE. There are new data types with DECIMAL FLOAT, BIGINT, BINARY and VARBINARY types. Improvements in LOBs provides more consistent handling and improved performance. Intersect and Except set operations make some SQL operations simpler to specify. Security is improved with ROLES and network trusted context. Data definition consistency and usability are improved. V9 is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.

## 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



### Native SQL stored procedures

Stored procedures written in SQL procedure language enhance portability and ease of use when using DB2 for z/OS as your enterprise information source. This language is an ANSI standard language. It is similar to the proprietary stored procedure languages of several competitive databases, which assists in migrating and porting to DB2 for z/OS. This is very similar to the changes in DB2 for LUW V8.2.

SQL stored procedures are supported by the DB2 Development Center tooling, providing an environment to code, test, and debug modules from your connected workstation. This language is currently converted to C when the CREATE PROCEDURE statement is executed. The C program is then automatically prepared, compiled, linked, and bound. The developer does not need to work with the C code.

SQL stored procedures code will be natively integrated into the DB2 engine, eliminating the conversion to C. Additionally, extensions to the bind command will allow for the promotion of the program and access paths between environments without needing to recreate the stored procedure.

Native SQL stored procedures are eligible to have a portion run on zIIP processors when they are invoked from a remote client.

## 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 >
```

This statement provides a fast way to delete rows with SQL, with better application portability. Truncate Table provides for the rapid removal of rows from a table. You can use this function to delete the contents of a table before applying data via LOAD or INSERT or MERGE.



## Decimal Floating Point

- New datatype DECFLOAT
  - 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



New data type support is provided. Support will be added for:

- Decimal floating point numbers, similar to calculator mathematics and supporting the IEEE standard. These numbers can have more precision than current floating point.
- BIGINT support of double word (8 byte) integer values
- VARBINARY, providing better comparison support for binary strings

## 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

**Select from DELETE, UPDATE, and MERGE:** The object-relational capabilities of DB2 allow for the incorporation of business logic into the database. This extends the power of SQL. Sometimes the application needs to know the results of this logic, when applied to the SQL issued. A subsequent SELECT for the data adds complexity and execution time to the application.

The *insert within select* feature of DB2 for z/OS Version 8 has been expanded to include the retrieval of columns from rows that are modified via DELETE, UPDATE, and MERGE SQL. One SQL call to DB2 modifies the table contents and returns the resultant changes to the application program.

When used with DELETE, the application now has the option to code a destructive read from a table. This is particularly useful when a table is used as a data queue, as with many vendor packages.

**Subquery improvements:** Correlated and non-correlated subqueries will benefit from improved optimization. They will provide added flexibility with the support of ORDER BY and FETCH FIRST clauses.

## Text improvements in DB2 9

- **30 new & improved character functions**
- **Index on expression: e.g. UPPER, COLLATION\_KEY**
- **LOB improvements**
- **pureXML**
- **Text search server**

The text improvements in DB2 9 include many improvements and new functions which can provide fast search for text, including indexing, using the new index on expression with text functions like UPPER that is sensitive to locale, COLLATION\_KEY, SOUNDEX and DIFFERENCE functions. Improvements for large objects and XML also help with many text applications.

Specialized text search with a separate server is noted on the next pages.

### **Text function improvements in DB2 9**


- New built-in character functions: COLLATION\_KEY,
  - ASCII CHR, ASCII\_STR, COLLATION\_KEY
  - DIFFERENCE, EBCDIC CHR, EBCDIC\_STR
  - IS\_IDENTICAL\_TO, LEFT, LOCATE\_IN\_STRING
  - LPAD, NORMALIZE\_STRING, OVERLAY
  - RIGHT, RPAD, SOUNDEX, UNICODE
  - UNICODE\_STR, VARCHAR\_FORMAT
  - XMLATTRIBUTES, XMLCOMMENT,
  - XMLDOCUMENT, XMLPARSE, XMLPI
  - XMLQUERY, XMLSERIALIZE, XMLTEXT
- Index on expression: e.g. UPPER, LOWER

The text improvements in DB2 9 include many improvements and new functions which can provide fast search for text, including indexing, using the new index on expression with text functions like UPPER that is sensitive to locale, COLLATION\_KEY, SOUNDEX and DIFFERENCE functions.

Specialized text search is noted on the next page.

## Text Search Server

- Text search for CHAR, VARCHAR, CLOB & XML columns
- Provide a text index server
- Efficient communication interaction with DB2 for z/OS
- Text indexes are persisted into DB2 tables for backup & recovery purposes



21

GoFurther

A specialized text search engine on a separate server is expected to come in V9, but will deliver after general availability. While the server is separate, the text indexes are saved into DB2 tables to provide improved backup and recovery.

## DB2 9 Spatial Support

### **Enabling Open Geospatial Consortium (OGC) compliant geospatial applications**

- **Spatial data types**
- **Spatial functions and predicates**
- **Spatial indexes**
- **Spatial search**
- **OGC-compliant spatial catalog**

GIS = Geographic Information System. Aka Geospatial. E.g. Google Maps.

OGC = Open Geospatial Consortium

Spatial support will also need a separate server and may deliver after general availability for V9.



## 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


### **Additional family compatibility**


Additional compatibility will be added for:


- Default databases and table spaces: When porting from other DBMS that do not have the same use for database and table space, providing them automatically will reduce the effort for delivering good performance.
- Automatic unique indexes are created to support defined primary keys.
- The default for a table space will change to segmented, improving performance and manageability for many customers. Existing simple table spaces are supported, but new table spaces will be segmented or partitioned.



**Leverage Application Development Skills**



- Use new converged SQL
- Key Database Technologies
  - SQL, SQL Procedures
  - XML
  - SOA, Web Services
- Developer communities
  - COBOL, PL/I, REXX, C, C++, assembler, Fortran
  - Java (JDBC / SQLJ)
  - .NET (C#, VB .NET)
  - Open Source
    - PHP
    - Perl
    - Python
    - Ruby on Rails
    - TOAD for DB2





























24

Application programming is using a wider range of tools, environments and languages. The Eclipse framework is growing strongly. We need to connect the new languages and environments to the scale and value of the existing infrastructure. The Rational and WebSphere product lines provide part of the connection, with products like Rational Data Architect and WebSphere Information Integration. DB2 clients provide more support for new environments and new languages.



## DB2 9 for z/OS Innovation: Data Warehousing

- Dynamic index ANDing for star schema**
- INTERSECT, EXCEPT**
- Query optimization improvements**
- Improved query performance**
- Index compression**
- Plan stability**
- Optimization Service Center**



25

GoFurther

**The second category of innovation is**

### **COST SAVINGS THROUGH OPTIMIZATION**

- o Increased security and regulatory compliance through implementation of roles, network-trusted contexts, and enhanced auditing
- o Performance-boosting innovations such as load and reorg CPU reductions, improved performance for varying length data, and improved logging and insert performance
- o Synergy with IBM System z and z/OS in areas that include XML parsing, zIIP, MIDAW channel improvements, encryption, IPv6 and Secure Socket Layer (SSL)
- o Query management enhancements to make accessing your data even faster and more accurate with indexing improvements that include index on expression, randomization, and larger index page sizes and optimization improvements that provide better data for the optimizer, improved optimization techniques, and better management with optimization services

## Query Enhancements

- SQL enhancements: INTERSECT, EXCEPT, cultural sort, caseless comparisons, FETCH FIRST in fullselect, OLAP specifications: RANK, DENSE\_RANK, ROW\_NUMBER ...
- pureXML integration and text improvements
- Index improvements: Index on expression, Index compression, ...
- Improved Optimization statistics: Histogram
- Optimization techniques
  - Cross query block optimization and REOPT(AUTO)
  - Generalize sparse index & in-memory data cache method
  - Dynamic Index ANDing for Star Schema
- Analysis: instrumentation & Optimization Service Center

26

GoFurther

Improving data warehousing and reporting: Today's complex applications include both transactions and reporting, so performing both well is imperative. The key improvements for reporting are optimization enhancements to improve query and reporting performance and ease of use. Improved data is provided for the optimizer, with improved algorithms and a rewritten approach to handling performance exceptions.

More queries can be expressed in SQL with new SQL enhancements. The set operators INTERSECT and EXCEPT clauses make SQL easier to write. OLAP extensions for RANK, DENSE\_RANK and ROW\_NUMBER add new capabilities. Other SQL statements improve consistency with the DBMS industry. V9 continues the progress in SQL, with many new functions, statements and clauses. The biggest changes are in XML on a prior slide. New SQL data manipulation statements are MERGE and TRUNCATE. New data types with DECIMAL FLOAT, BIGINT, BINARY and VARBINARY. Improvements in LOBs provide new function, more consistent handling and improved performance. Data definition consistency and usability are improved. V9 is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.

Histogram statistics enable DB2 to improve access path selection by estimating predicate selectivity from value-distribution statistics that are collected over the entire range of values in a data set. RUNSTATS cannot collect histogram statistics on randomized key columns. DB2® chooses the best access path for a query based on predicate selectivity estimation, which in turn relies heavily on data distribution statistics. Histogram statistics summarize data distribution on an interval scale by dividing the entire range of

## Modified SAP/BW Workload

- **Data**
  - Populated with SAP benchmark BW 3.5 toolkits
  - Fact table size: 58.4M rows , 8 indexes
  - Dimension tables: 8 (2 ~ 99326 rows)
  - Snowflakes: 6 (added to increase query complexity)
- **Queries (100)**
  - Developed by DB2 development and performance
  - Based on V8 BW workload
  - New queries added to better reflect the customer scenarios learned from the V8 service stream
- **Represent customer workloads without adequate (multi-column) index support (this is the norm)**

## Performance Comparison

	DB2 V8	DB2 9	Improvement
Total Elapse Time (seconds)	<b>71660</b>	<b>8544</b>	<b>88%</b>
Total CPU time (seconds)	<b>7400</b>	<b>7514</b>	<b>-1.5 %</b>
CPU time eligible for zIIP off-load	<b>2924 (39.5%)</b>	<b>6775 (90%)</b>	

In V9, 71 out of 100 used PWJ

Key messages: V9 performs better than V8 for BW workloads with well-tuned index design (see V8 BW workload for details)

V9 outperforms V8 for BW workloads with inadequate multi-column index support (see new BW workload for details)

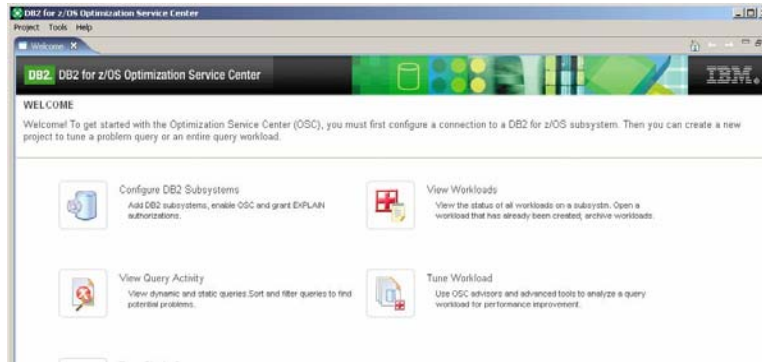
More room for zIIP off-load in V9 for BW workload

V9 is a great out-of-box solution which removes the burden of index design from users

Overall, a significant TCO reduction in V9 for BW customers

## Optimization Service Center

- Identify Problem Query
- Tune Problem Query
- Monitor & Capture Query Workload
- Tune Query Workload



## DB2 9 for z/OS Innovation: Cost

- ❑ **Cost Savings through Optimization**
  - ❑ **Security and Regulatory Compliance**
  - ❑ **Performance improvements**
  - ❑ **Synergy with System z**
  - ❑ **Query enhancements**
  - ❑ **Index Compression**
- ❑ **Reduced Complexity**



30

GoFurther


**The second category of innovation is**

### **COST SAVINGS THROUGH OPTIMIZATION**

- o Increased security and regulatory compliance through implementation of roles, network-trusted contexts, and enhanced auditing
- o Performance-boosting innovations such as load and reorg CPU reductions, improved performance for varying length data, and improved logging and insert performance
- o Synergy with IBM System z and z/OS in areas that include XML parsing, zIIP, MDAW channel improvements, encryption, IPv6 and Secure Socket Layer (SSL)
- o Query management enhancements to make accessing your data even faster and more accurate with indexing improvements that include index on expression, randomization, and larger index page sizes and optimization improvements that provide better data for the optimizer, improved optimization techniques, and better management with optimization services

## DB2 9 for z/OS Performance Improvements

- Synergy with new hardware: zIIP, MIDAW, DS8000
- Significant CPU time reduction in most utilities
- Performance/Scalability Enhancements
  - Especially Insert / Update / Delete
- Query/Access Path Performance Enhancements Other Performance Enhancements
  - Native SQL procedure, index compression, LOB, Varchar, ...
- DDF Improvements
- Improved virtual storage usage below bar



The key performance improvements in V9 are reduced cpu time in the utilities, improved LOB performance and scalability, improved optimization for SQL, the zIIP processing for remote native SQL procedures, reduced cpu time for data with varying length and better sequential access.

I'll discuss the optimization improvements on the query slide. V8 SQL procedures were not eligible to run on the zIIP, but changing to use the native SQL Procedure Language on V9 will make the work eligible for zIIP processing. Varying length data can improve substantially if there are large numbers of varying length columns. Several improvements in disk access can reduce the time for sequential disk access.

## 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

**Large object improvements:** Large objects (LOBs) were introduced in DB2 Version 6. Usage has increased substantially in the past few years, and major enhancements have been made in DB2 Version 8.

APARs on Version 8 deliver the ability to use utilities for loading and unloading large LOB data. File reference variables are used to let the large objects be accessed from data sets instead of from storage. The abilities to reorganize and to recover space are provided.

Future changes will help with improved function and usability, DB2 family compatibility, cost of ownership, performance, and scalability.



## Other Performance / Availability Items

- Insert performance APPEND INDEX LOG
  - INDEX on expression, larger page sizes, better split, ...
  - Log performance in data sharing, log archive striping
  - Not logged table space (scalability)
- CPU reductions in LOAD and REORG
- Online REBUILD INDEX REORG without BUILD2
- Improved varying length performance
- FETCH FIRST n ROWS improvements
  - Can now be specified in a subquery or fullselect
  - ORDER BY now exploits FETCH FIRST n ROWS, so that work files are not created (less I/O)

Many other improvements help with performance, with scalability and with availability. Index on an expression can be combined with caseless comparisons to improve text search. Improved insert rates can result from improved log performance. Data sharing logging rates should increase significantly. Archive logs can use striping, providing a performance boost when needed. Significant reductions in cpu usage are provided with the new utilities.

Improved cpu and elapsed times can be achieved with the FETCH FIRST clause specified on a subquery.

## 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

### **Table Append option**

The Table Append option offers increased performance for inserting data into the end of a table. It reduces the instructions used to target the locations for new rows. Index maintenance techniques are also made more efficient when DB2 detects that entries are added to the beginning or ending of the index.

## 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

35

GoFurther

Sequential insert performance is improved by avoiding page splits with larger index page sizes and the ability to split a page more effectively. Other changes improve logging rates.

Larger page sizes: up to 8x fewer splits

Asymmetric index split: up to 2x fewer splits, cpu/elapsed time savings. Up to -20% cl2 cpu, -31% elapsed

## Utility CPU time reduction up to –

\* primarily from index processing

- 10% to 20% in Copy, Recover table space / index\*
- 5% to 30% in Load\*, Reorg\*, Rebuild Index\*
- 20% to 60% in Check Index\*
- 35% in Load Partition\*
- 30% to 40% in Runstats Index\*
- 40% to 50% in Reorg Index\*
- 70% in Load Replace Partition with dummy input\*

DB2 9 preliminary figures for utility cpu use are very encouraging. There is substantial variation in the figures, depending upon the exact workload, but the improvements are large and across most of the key utilities: copy, reorg, load, rebuild, recover, runstats, and check index.

## System z Synergy & DB2 9

- ✓ **System z9 Integrated Information Processor (zIIP) Enterprise Class & Business Class**
- ✓ **Enhanced Cryptography**
- ✓ **Channels (4 Gb & MIDAW)**
- ✓ **Faster Processors**
- ✓ **Up to 54 Processors EC**
- ✓ **More memory, better value; 64 bit virtual storage**
- ✓ **z/Architecture new instructions**
- ✓ **Parallel Sysplex**



37

- ✓ IPv6
- ✓ SSL
- ✓ Java
- ✓ Decimal float
- ✓ Backup & restore
- ✓ Security
- ✓ Unicode collation
- ✓ Compression
- ✓ System z Application Assist Processor (zAAP)
- ✓ WLM enhanced ...



GoFurther

The latest System z9 processor improvements for DB2 are the zIIP and the new Business Class and Enterprise Class processors. DB2 V9 remote native SQL procedures are enabled for zIIP processing. V9 adds IPv6, SSL and decimal float and BIGINT data types, with enhancements for Parallel Sysplex, backup and restore, added security and encryption, more Unicode with collation, and uses the WLM in new ways. Channel enhancements (MIDAW) and improved DS8000 performance were included with the System z9 announcements. DB2 uses the latest improvements in hardware and operating system to provide better performance, improved value, more resilience and better function.

DB2 benefits from large real memory, faster processors, specialty engines, and better hardware compression. DB2 uses Parallel Access Volume and Multiple Allegiance features of the IBM DS8000 and Enterprise Storage Server™. FlashCopy® can be used for DB2 backup and restore. DB2 makes unique use of the z/Architecture™ instruction set, and recent instructions provide improvements in reliability, performance and availability. DB2 continues to deliver synergy with hardware data compression, FICON™ (fiber connector) channels, disk storage, advanced networking function, and Workload Manager (WLM).

## Synergy with new I/O hardware

DS8000 with Ficon Express and MIDAW

(Modified Indirect Data Address Word)

- MIDAW requires z9 (2094) and z/OS 1.6 or later
- Sequential read throughput
  - 40MB/sec on ESS 800      69MB/sec with DS8000
  - 109MB/sec with DS8000 and MIDAW
  - 138MB/sec with 2 stripes
- Bigger read, write, preformat quantity
  - 183MB/sec in sequential read with 2 stripes
- Similar for write
- Performance gap between EF (Extended Format) and non EF datasets or 4K and bigger page practically gone

38

GoFurther

MIDAW requires z9 (2094) and z/OS 1.6 with APARs OA10984 8/05, PK13324/13384 9/05 or later. Refer to earlier chart for performance figures with throughput and response times.

Then add

Faster channels (4 Gb)

Faster disk DS8000 Turbo

Sequential improvements in V9 – read and insert

Logging improvements in V9

## DDF Improvements

- 64-bit addressing by DDF
  - Special “shared private” with xxxDBM1 to eliminate many 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

39


GoFurther

Distributed connections to DB2 for z/OS will benefit from z/OS V1R7 changes that DB2 will exploit. DB2's distributed communication processes (the *distributed address space*) will access data directly from the database manager address space, instead of moving the data. The distributed address space will also exploit 64-bit addressing, as the database manager and lock manager address spaces do today with Version 8.

This internal change will benefit new and existing workloads, where distributed communications are configured with another logical partition (LPAR) or to an application running on the zSeries platform.

## Index Improvements

- INDEX on expression
- Page sizes 8K, 16K, 32K
- Improved page split
- Index compression
- Online REBUILD INDEX
- REORG without BUILD2 – not just for DPSI
- Randomized index key
- Not logged index space
- XML index



40

GoFurther

Many other improvements help with performance, with scalability and with availability. Index on an expression can be combined with caseless comparisons to improve text search. Improved insert rates can result from improved log performance, larger index page sizes, and improved index page splits. See the next slide for index compression. Data sharing logging rates should increase significantly. Archive logs can use striping, providing a performance boost when needed. A randomized index key will optimize for inserting, like the APPEND option for data.


Significant reductions in cpu usage are provided with the new utilities. REBUILD INDEX has an online option for the first time, while REORG becomes much more online by removing the BUILD2 phase.

The not logged table and index space can help with scaling where the log is a bottleneck, but won't increase performance significantly. The XML index is very different from any other 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



41

GoFurther

Index compression relies upon page level compression instead of row-level compression (technique used for tables). Indexes with 32K page sizes can save up to 8x on disk space with the compression feature turned on. The solution has to guarantee that a given page will compress to 4K on disk, so DB2 will sometimes force pages in memory to be “partial” so that the compression will indeed result in no more than a 4K byte page on disk.

## 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)

### DSN1COMP utility to simulate compression ratio without real index compression

#### Notes

- No compression or decompression in each Insert or Fetch; instead at I/O time
- Load or Reorg not required for compression
- Based on a very limited survey so far
  - Higher for relatively unique indexes with long keys
- CPU time impact under study

## 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

43

GoFurther

### **Temporary storage architecture and use**

Many recent improvements have enhanced the use of temporary storage when it is needed by the DB2 engine, for example, for interim materialized result sets.

The different temporary pools are converged to a single source. Use of temporary storage is eliminated when processing small amounts of data.

## 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

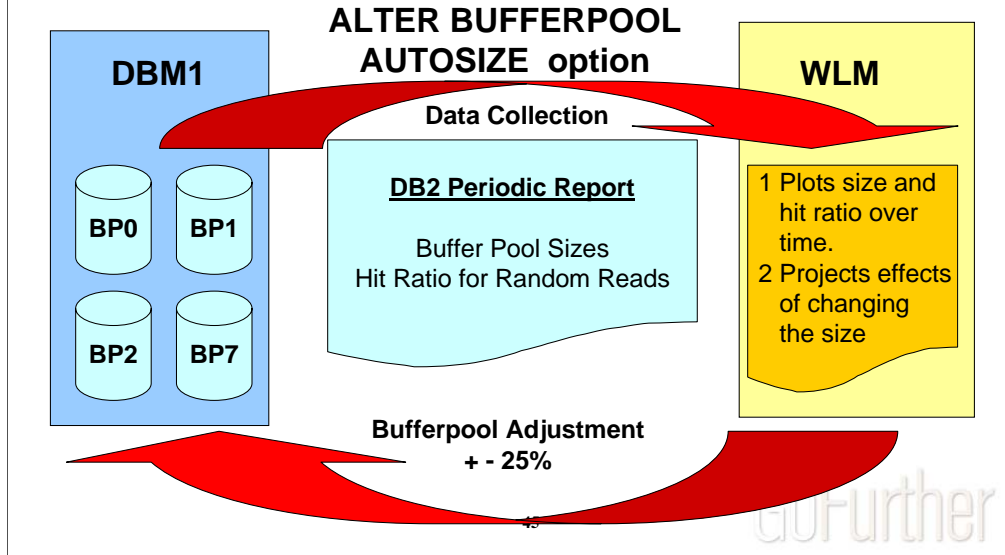
44

GoFurther

V8 provided a new BACKUP utility using FlashCopy technology to take very fast backups of the entire subsystem without any disruption. RECOVER is only for the whole subsystem.

With Vnext we expect to be able to recover an object, rather than the whole subsystem. This makes the job of backup and recovery simpler and easier.

# WLM assisted buffer pool management



DB2 periodically reports the buffer pool size and the hit ratio for random reads to WLM via a "data collection" exit. WLM plots the size and hit ratio over time to allow it to project the effects of changing the size, then adjusts the buffer pools, if WLM goals will improve.

## 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.

Version 9 builds upon the structure from V8, with the same modes, but clarifying the situation if you drop back after moving to ENFM or NFM.

## 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 or more of savings expected
  - Mainly from EDM related storage (static SQL) and dynamic statement cache (dynamic SQL)

47

GoFurther

Skeletons above the bar: For customers that use heavy package and plan activity such as banks, this is the most significant DBM1 below the bar storage relief in V9. For customers that use very few or small packages, such as SAP environments, the savings is smaller.

LI702 – move spaceblk (SPA) above the bar. SPA to be split into 2, 1 above, 1 below. Only a few, non-complex RTs are being considered for V9. Simple insert, delete Expected results will vary by SQL mix. (-5 to 30%)?

## 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)
- Command line processor
- Optimization Service Center

48

GoFurther

RLF is extended to support the Set Client Information APIs that are exposed in CLI and JDBC. This allows the customer to control CPU cost of dynamic queries for packaged applications and user-written applications that exploit these APIs.

CREATE | ALTER STOGROUP

```
DATACLAS          dcname
MGMTCLAS          mcname
STORCLAS scname
```

The existing VOLUME clause is now optional: it can be omitted if any of the DFSMS classes is specified. If explicitly specified to DB2, the new attributes are recorded in SYSIBM.SYSSTOGROUP



## DB2 9 for z/OS Innovation: SOA and XML

- ❑ Integration with WebSphere
- ❑ Native XML data type, hybrid data base server



GoFurther

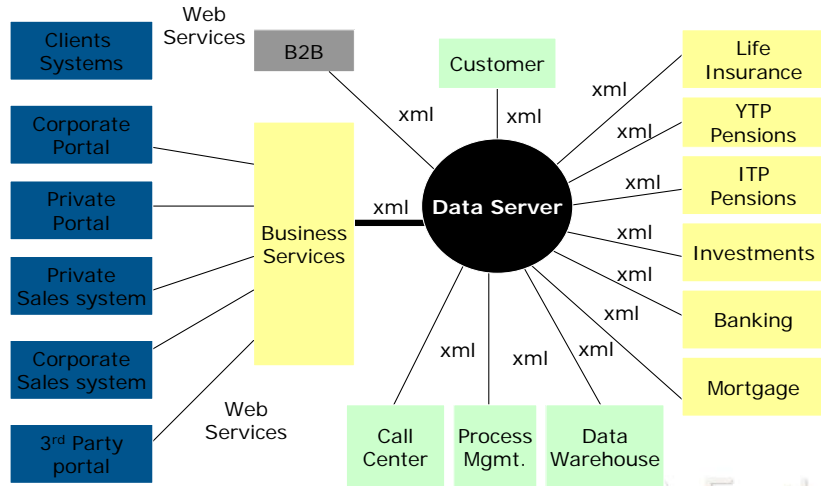
## Optimistic Locking Support

- 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

This change will help with migration of applications which use application consistency techniques. Application programmers will not need to be responsible for updating the timestamps. A simple timestamp predicate can be used to make sure that other updates are not lost.

# DB2 9 for z/OS ... Powering SOA Solutions

*XML is the Key Link*




## Contrasting the Models *XML and Relational*

- **Relational**
- Strength: Static data
  - Strict schema ensures data integrity
  - High performance indexing on fixed data
- Strength: 'Set-based' data
  - Multiple results returned
  - Retrieving rows
- **XML**
- Strength: Semi-structured, frequently changing data
  - Self-describing, flexible schema
  - Easily modified format
- Strength: Retrieving sequences
  - Documents, subdocuments, related documents



*Over \$20B Annual Customer Technology Investment in RDB Alone...*



*XML database investments growing twice as fast as total database investment...*

## XML Data Needs Relational Maturity

### Complementing XML Processing

- **XML Data Needs Protection**

- Backup and recovery features to ensure continuity
- Data is protected using database security



- **Simplified XML Data Access**

- Centrally store and access difficult to retrieve data
- SQL or XPath can be used to retrieve data
- Join XML data with it's related relational data



- **Search Speed**

- Search documents quickly and efficiently using proven search optimization engine of mature database



- **Optimize Existing Investments**

- Use existing technology infrastructure and skills to store and manage both relational and XML



## pureXML

- 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

GoFurther

Now let's discuss the wide range of new function being provided inside the DB2 for z/OS engine. This work is being done in parallel with similar changes in DB2 for Linux, Unix and Windows.

DB2 for z/OS extensibility has been implemented via *extenders*. Extenders for text, image, audio, video, and XML are delivered with DB2 and are optional installation components. They provide the tooling, user defined data types (UDT), user defined functions (UDF), and stored procedures to managed non-relational data. The XML extender provides for the storage and indexing of an XML document as a character large object (CLOB), or for the shredding of the XML document into relational columns for storage and query.

DB2 Version 8 expanded on XML support by implementing several XML publishing operations as built-in DB2 functions. This allows you to perform XML document composition from relational data with improved performance without the XML Extender.

There will be expanded support of XML in DB2 by integrating more features into the engine. This includes an XML data type, native storage of XML documents, integration of the XPath language, and catalog extensions to support definitions of XML schemas. Utilities will support creation and maintenance of XML data.

## pureXML -- 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")
```

Support for pureXML is provided in DB2 by integrating more features into the engine. This includes an XML data type, native storage of XML documents, integration of the XPath language, and extensions to support definitions of XML schemas. Utilities support creation and maintenance of XML data.

The SQL/XML publishing functions provided in DB2 V8 are enhanced, and many new functions and predicates are added.

## pureXML

- 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

The utilities are extended to support XML. You can LOAD, UNLOAD, CHECK DATA or INDEX, COPY, REBUILD, RECOVER and REORG the XML data.

The XML schema support handles a schema repository and uses tables to store XML schema information. Stored procedures are provided to register XML schemas.

IBM Systems Journal issue on XML:

<http://www.research.ibm.com/journal/sj45-2.html>



## DB2 9 for z/OS Innovation: Continuous Availability

- ❑ Online schema evolution
- ❑ More online utilities
- ❑ Data sharing enhancements



## Schema Evolution – Database Definition On Demand

- Fast replacement of one table with another
- Rename column and index
- Rename SCHEMA and VCAT
- 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

58

GoFurther

One of the key initiatives of V8 was online schema evolution, and that theme is expanding and changing to be data definition on demand.

One of the important changes is to be able to replace one table quickly with another. Another is to be able to rename a column or an index. Index page size can be altered. A new type of table space combines the attributes of segmented and partitioned, without a partitioning key. Rebuild index can be run with much less disruption. Online table space reorganization for a few partitions is improved a lot, removing the BUILD2 phase for all types of secondary indexes. Table space and index logging can be altered. A new ability to change the DB2 early code does not require an IPL.

## 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

### **Fast replacement of a table with another**

This function will allow you to generate a copy of a current table with the same attributes and same data, in the same table space. It appears to an application or user as a very fast replacement of data within a table. Web-based applications, striving for maximum availability will benefit from the option of implementing a pair of tables that are clones of each others' structure. Copies for application testing and auditing can be easily created. These *clone* tables will have the unique ability to change names quickly. Applications can quickly and almost transparently switch between dual mirror tables.

## CLONE Tables...

- A clone table can only be created
  - On a single table in a table space (partitioned or non-partitioned)
  - 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

The tables can be partitioned or non-partitioned, and are created with the CREATE TABLE syntax. The primary table's structure, including indexes, large objects (LOBs) and before triggers, will be copied. Information can then be inserted or loaded into the copy table, and the copy can have its own image copies. When the data in the copy table needs to become active to an application, an ALTER statement will switch the name, providing fast replacement of the original data

## 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

61

GoFurther

### Partitioned by growth

Partitioned tables have required key ranges to determine the target partition for row placement. When a table is partitioned, you gain the benefits of scaling objects to hold more data. You also benefit from more granular locking and parallel operations by spreading the data over more data sets.

The option to partition by growth will allow segmented tables to be partitioned as they grow, without needing key ranges. These segmented tables will gain increased table space limits and the SQL and utility parallelism, afforded for partitioned tables.

CREATE TABLESPACE ... (explicit specification)

MAXPARTITIONS integer

CREATE TABLE ... (implicit specification)

PARTITIONED BY SIZE EVERY integer G

PBR can also be UTS via `CREATE TABLESPACE ... SEGSIZE integer`  
`NUMPARTS integer`

## DB2 9 Utilities

- Support for all new functions in DB2 Version 9 for z/OS product (universal table spaces, XML, not logged, etc.)
- 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
- Recover to consistent PIT without need for a quiesce

62

GoFurther

More online utilities are coming! Online rebuild build an index in place after setting RBDP. Dynamic SQL will avoid the index until available. This works well for non-unique indexes.

Cloned tables are somewhat like Online Load Replace. There is the live table, and then there is the shadow or clone. One can load the clone, and then when ready for the clone to become live, a “switch” can be trigger with DDL.

Parallel log apply for SHRLEVEL REFERENCE and CHANGE  
 Multiple concurrent jobs no longer needed/supported – DSNU180I  
 NPSIs are also shadowed (higher disk requirement compared to when just logical part was reorganized), thus resulting in the NPIs also being reorganized -> therefore, don't follow with REORG of NPSIs.  
 Higher availability results in increased resource consumption  
 REORG SHRLEVEL REFERENCE now has a log phase since entire NPSI is shadowed and reorg'd

Online reorg solves the outage during the BUILD2 phase when reorging a partition at a time.

Online Check Data and Check Lob are the same as existing utilities, only with SHRLEVEL CHANGE.

Volume-based utilities allow recovery of individual objects after a volume based backup. It also extends system based backups and restores to and from tape.

Broken page recovery is a serviceability improvements when data has become broken because log apply fails.

Template switching allows using tape vs dasd if a size or limit on the ds is exceeded.

Modify Recovery is extended to allow n number of backups vs an age.

Check || is the SHRLEVEL CHANGE technique extended to SHRLEVEL REFERENCE

CPU reduction is to improve where we don't match up well against other vendors on a cost basis.

Histogram stats uses quantiles to yet give the optimizer more stats.

## DB2 9 Utilities

- More online utilities
  - Check data, LOB and repair locate ... SHRLEVEL CHANGE
  - Check index SHRLEVEL REFERENCE supports parallel for > 1 index
  - Load replace (shrlevel change) with CLONE TABLE function
- Always perform CHECKPAGE on the COPY utility
  - Prior to V9, CHECKPAGE was optional, with about ~5% CPU overhead, and if a broken page was encountered (DSNU441I for space maps or DSNU518I for others, both RC8), then copy-pending was set
  - Now, COPY always performs these checks (with reduced overall CPU!) and no longer sets copy-pending, so.... Check those RCs!
  - A new SYSCOPY record type is written if a broken page is detected to force a full image next since dirty bits may have already been flipped off in the space map pages

63

GoFurther

More online utilities are coming! Online rebuild build an index in place after setting RBDP. Dynamic SQL will avoid the index until available. This works well for non-unique indexes.

Cloned tables are somewhat like Online Load Replace. There is the live table, and then there is the shadow or clone. One can load the clone, and then when ready for the clone to become live, a “switch” can be trigger with DDL.

Online reorg solves the outage during the BUILD2 phase when reorganizing a partition at a time.

Online Check Data and Check Lob are the same as existing utilities, only with SHRLEVEL CHANGE.

Volume-based utilities allow recovery of individual objects after a volume based backup. It also extends system based backups and restores to and from tape.

Broken page recovery is a serviceability improvements when data has become broken because log apply fails.

Template switching allows using tape vs dasd if a size or limit on the ds is exceeded.

Modify Recovery is extended to allow n number of backups vs an age.

Check || is the SHRLEVEL CHANGE technique extended to SHRLEVEL REFERENCE

CPU reduction is to improve where we don't match up well against other vendors on a cost basis.

Histogram stats uses quantiles to yet give the optimizer more stats.

## 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

GoFurther



## 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<sub>65</sub> (V7, V8 usermod)

## DB2 9 for z/OS: Beta Summary

### Program Profile

- 10 core beta participants
- 15+ expanded beta participants
- 40+ external vendor program participants

### Core Customer Statistics

- 40% of core customers have committed DB2 9 deployment plans that start within 6 months, some as early as the day of GA
- 7 customers measured ~50% space savings with index compression
- 20-30% cpu savings in LOAD and REORG observed
- “Smooth” installation/migration reported

66

GoFurther

Core participants: code in hand for 9 months

Expanded: code in hand for 3-6 months, more being added

Vendors: specs in hand for 2 yrs. Code for 15 mo.

(~2500 for PVT, rest is split about 65/35 NST/CCE)

## DB2 9 for z/OS Beta Highlights

**Customer1** – industry: finance, geo: Europe

- Extensive regression & new function testing in a full clone of production environment
- Every DB2 test system successfully migrated to DB2 9 at least twice
- Measurement tests on rebuild index and index compression show 50% space savings; reorg utility elapsed times improved 33-147%, as measured during beta
- Eager to leverage new release capabilities: ***“We plan to order DB2 9 on the day of availability”***

**Customer2** – industry: finance, geo: Americas

- ***“DB2 is one of the core technologies we rely on to store business data & process time critical transactions”***
- Key interests: security & compliance, zIIP enhancement (native SQL procedures), LOB improvements, XML, QMF
- **pureXML testing successful**; POC started to explore redesign of internal system to exploit this exciting technology


## DB2 9 for z/OS Beta Highlights







**Customer3** – industry: insurance, geo: Americas

- Key interests: XML, native SQL stored procedures, IBM Developer Workbench (DWB), utilities, LOB enhancements
- Eclipse-based DWB interface simplifies development & debugging.
- **DB2 9 GA project starts one month after GA:** *“Need for new functionality will drive DB2 9 across our enterprise”*

**Customer4** – industry: mfg, geo: Americas

- Key interests: roles, enhanced backup & restore, clones tables, networking w/DB2 lab & other customers
- Nightly outage window for 1 table cut from 15-30 mins. to under 1 min. using cloned tables
- VCAT SWITCHes done for 60,000 table spaces, indexes, and stogroups in 9 mins. (vs. 2 hrs under prior process)
- **“Our input matters:** *DB2 COPY template switching added at our request; further refinement planned for incremental FlashCopy after GA; changes to DSN3@ATH and DSN3@SGN made due to our testing”*


**Beyond DB2 9**

- 
 TCO (improved performance, reduced people cost)
 
- 
 Availability improvements (fewer planned outages)
- 
 Data warehousing improvements
- Application development and application portability (SQL, XML)
 


GoFurther

69

**What is beyond DB2 9? Work for the future** provides the most comprehensive database functionality (OLTP and BI), the best performance and scalability, the highest RAS characteristics and, most importantly, much reduced administration cost. Today's database servers consist of independent components that require significant database and systems administration cost. Using tight integration of DBMS, operating system and hardware we can dramatically reduce the cost by:

- eliminating need for OS skills and System z-unique skills
- reducing DBA to absolute minimum
- use open, Web interfaces when interaction necessary

**What are targeted ecosystems?**

DB2 can be used for all applications, new and existing ones, home-grown or created by ISVs, transactions, queires, web serving or a mixture. One key target will be existing and new ISV installations such as SAP, Siebel, ...

**Why are these changes attractive to customers and ISVs?**

Customer and ISV objective is to drive down the customers total cost of ownership(TCO). Their game plan is to achieve this through non-ISV (such as IBM's) delivered hardware and software stack commoditization and people cost (systems administration) reduction. IBM's goal is also lowering the customers TCO. The people cost (systems administration) improvements through better usability, improved productivity and delivery of autonomic function complement hardware and software improvements in value.

## Some Vnext focus areas

- Performance
  - Dynamic, static
  - OLTP, batch, query
- Single-system scaling, Increased thread limits
- BIND/DDDL concurrency
- More online schema changes
- Plan stability (v9?)
- XML, SQL improvements
- Autonomics (stats, Reorg, compression, query tuning)
- Query analytics and aggregation functions

## 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>

V9 beta announcement: [http://www.ibm.com/common/ssi/rep\\_ca/8/897/ENUS206-098/ENUS206-098.PDF](http://www.ibm.com/common/ssi/rep_ca/8/897/ENUS206-098/ENUS206-098.PDF)

<http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&apiname=iSource&supplier=897&letternum=ENUS206-098>

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 <http://www.redbooks.ibm.com/abstracts/SG247259.html?Open>

LOBs book <http://www.redbooks.ibm.com/abstracts/SG247270.html?Open>

Security book coming SG24-6480

V9 Technical Overview coming SG24-7330

V9 Performance Topics coming



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, ... See notes below for more detail.

Main DB2 for z/OS web page: pointers to most of the following: <http://www.ibm.com/software/data/db2/zos/index.html>

V9 beta announcement: <http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&apiname=iSource&supplier=897&letternum=ENUS206-098>  
[http://www.ibm.com/common/ssi/rep\\_ca/8/897/ENUS206-098/ENUS206-098.PDF](http://www.ibm.com/common/ssi/rep_ca/8/897/ENUS206-098/ENUS206-098.PDF)

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:

SOA book includes overview of V9 XML <http://www.redbooks.ibm.com/abstracts/SG247259.html?Open>

LOBs book includes V9 <http://www.redbooks.ibm.com/abstracts/SG247270.html?Open>

Coming: security, V9 overview, V9 Performance Topics

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. Some require registration.

[http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008769&loc=en\\_US&cs=utf-8&lang=en](http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008769&loc=en_US&cs=utf-8&lang=en)

[http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008767&loc=en\\_US&cs=utf-8&lang=en](http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008767&loc=en_US&cs=utf-8&lang=en)

[http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008766&loc=en\\_US&cs=utf-8&lang=en](http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008766&loc=en_US&cs=utf-8&lang=en)

[http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008743&loc=en\\_US&cs=utf-8&lang=en](http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&dc=DA400&q1=presentation&uid=swg27008743&loc=en_US&cs=utf-8&lang=en)

Access the ftp site directly: <ftp://ftp.software.ibm.com/software/data/db2zos/>

|                                         |                                       |
|-----------------------------------------|---------------------------------------|
| IOD2559DB29ClonesLyle.pdf               | V9 clone tables                       |
| IOD1851DB2v9onlineUtilitiesHartmann.pdf | V9 utilities                          |
| IOD1855DB2v9designHartmann.pdf          | V9 applicationdesign                  |
| IOD1869DB2zOSv9performPreShibamiya.pdf  | V9 performance preview                |
| IOD1819PurcellV9.pdf                    | V9 optimization                       |
| IOD1730DB2v9xmlZhang.pdf                | V9 XML                                |
| IOD1729aDB2v9backuprecoveryTeng.pdf     | V9 backup and recovery                |
| IOD1641DB2V9autoQueryTuningFuh.pdf      | V9 Optimization Service Center part 1 |
| IOD1642DB2v9autoQueryTuningFuh.pdf      | V9 Optimization Service Center part 2 |
| IOD1450aDB2v9SAPHrle.pdf                | V9 for ERP and SAP                    |
| IOD1438db2v9LOBsWeihrach.pdf            | V9 LOBs                               |
| IOD1345DB2zOSmigrationg.pdf             | V8 & V9 migration                     |
| IOD1166DB2v9zOSbeyondCotner.pdf         | V9 overview                           |
| DB2V9forzOS.pdf                         | V9 overview                           |
| DB2V9zOS.ppt                            | V9 overview PowerPoint                |

IOD conference. There were about 20 presentations there, and most are on the conference web site and on the recently shipped CD.

Sessions from list above are 1166, 1345, 1438, 1450, 1641, 1642, 1729, 1730, 1819, 1851, 1855, 1869, and 2559.

1388 QMF V8 V9

1439 DB2 and SOA

Session: X01  
DB2 9 for z/OS and Beyond

**Jeff Josten**

IBM

[josten@us.ibm.com](mailto:josten@us.ibm.com)

