

May 6-10, 2007

San Jose Convention Center

San Jose, California, USA

DB2 9 for z/OS and Beyond

IDUG® 2007

North America




Bryan F. Smith bfsmith@us.ibm.com
IBM DB2 for z/OS Development



GoFurther

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



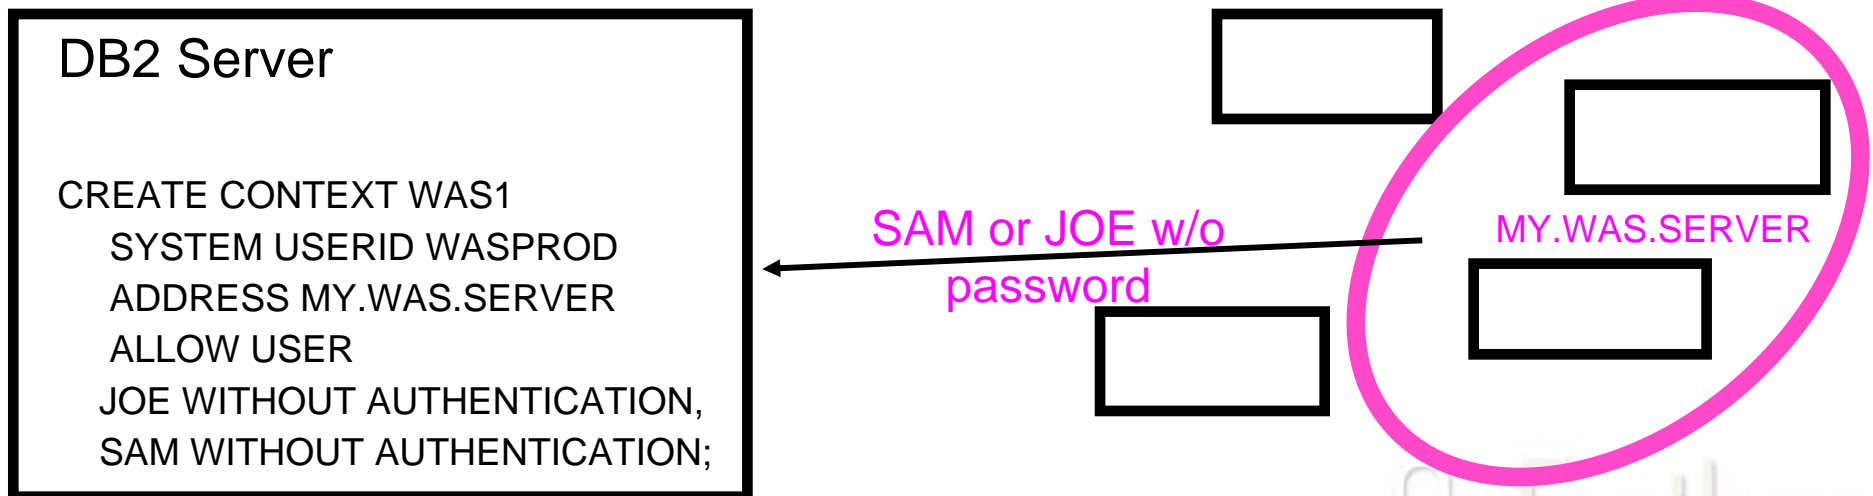
DB2 9 for z/OS: Regulatory Compliance

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



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



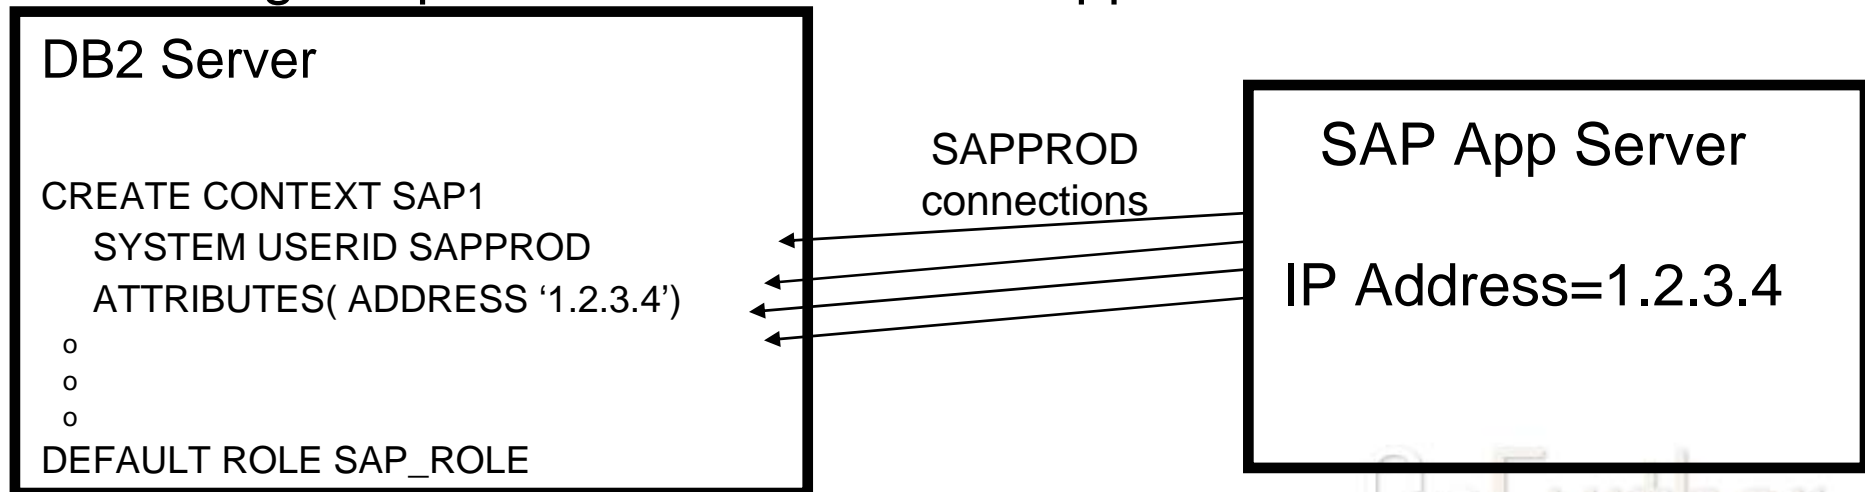
Database ROLES

- 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);
```


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



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!!!



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

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.



**Enterprise-wide
Key Management**



Tape TS1120



Disk



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

DB2 9 for z/OS Innovation: SQL

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



DB2 SQL

Z: z/OS V8

common

Iuw: 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
- Iuw** { 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

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**


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 >
```


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



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

Text improvements in DB2 9


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

Text function improvements in DB2 9

- New built-in character functions: COLLATION_KEY,
 - ASCII_CHAR, ASCII_STR, COLLATION_KEY
 - DIFFERENCE, EBCDIC_CHAR, 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

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



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**

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

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



GoFurther

DB2 9 for z/OS: Data Warehousing

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



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

Modified SAP/BW Workload

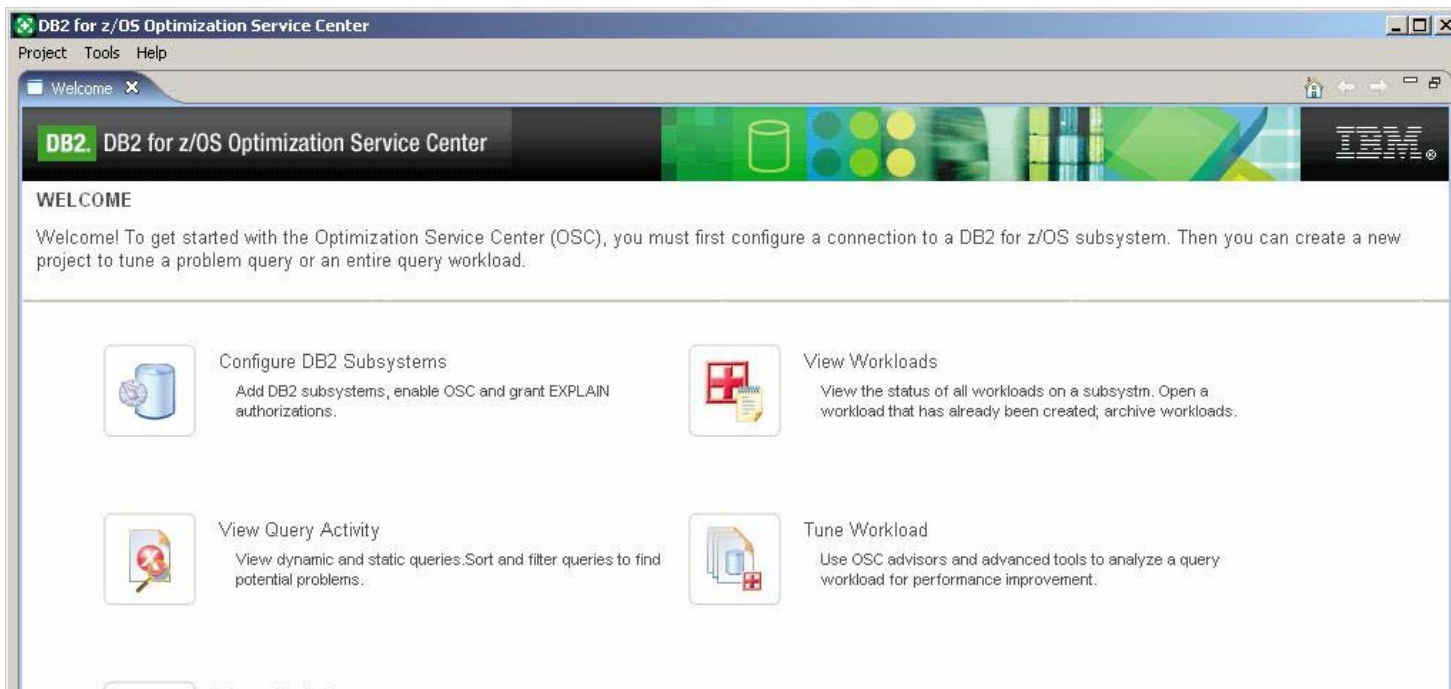
- **Data**
 - Populated with SAP benchmark BW 3.5 toolkits
 - Fact table size: 58.4M rows , 8 indexes
 - Dimension tables: 8 (2 ~ 99326 rows)
 - Snowflakes: 6 (added to increase query complexity)
- **Queries (100)**
 - Developed by DB2 development and performance
 - Based on V8 BW workload
 - New queries 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)	71660	8544	88%
Total CPU time (seconds)	7400	7514	-1.5 %
CPU time eligible for zIIP off-load	2924 (39.5%)	6775 (90%)	

Optimization Service Center

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



DB2 9 for z/OS: Cost

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



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



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

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)

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

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

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***

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**



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

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

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


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



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)

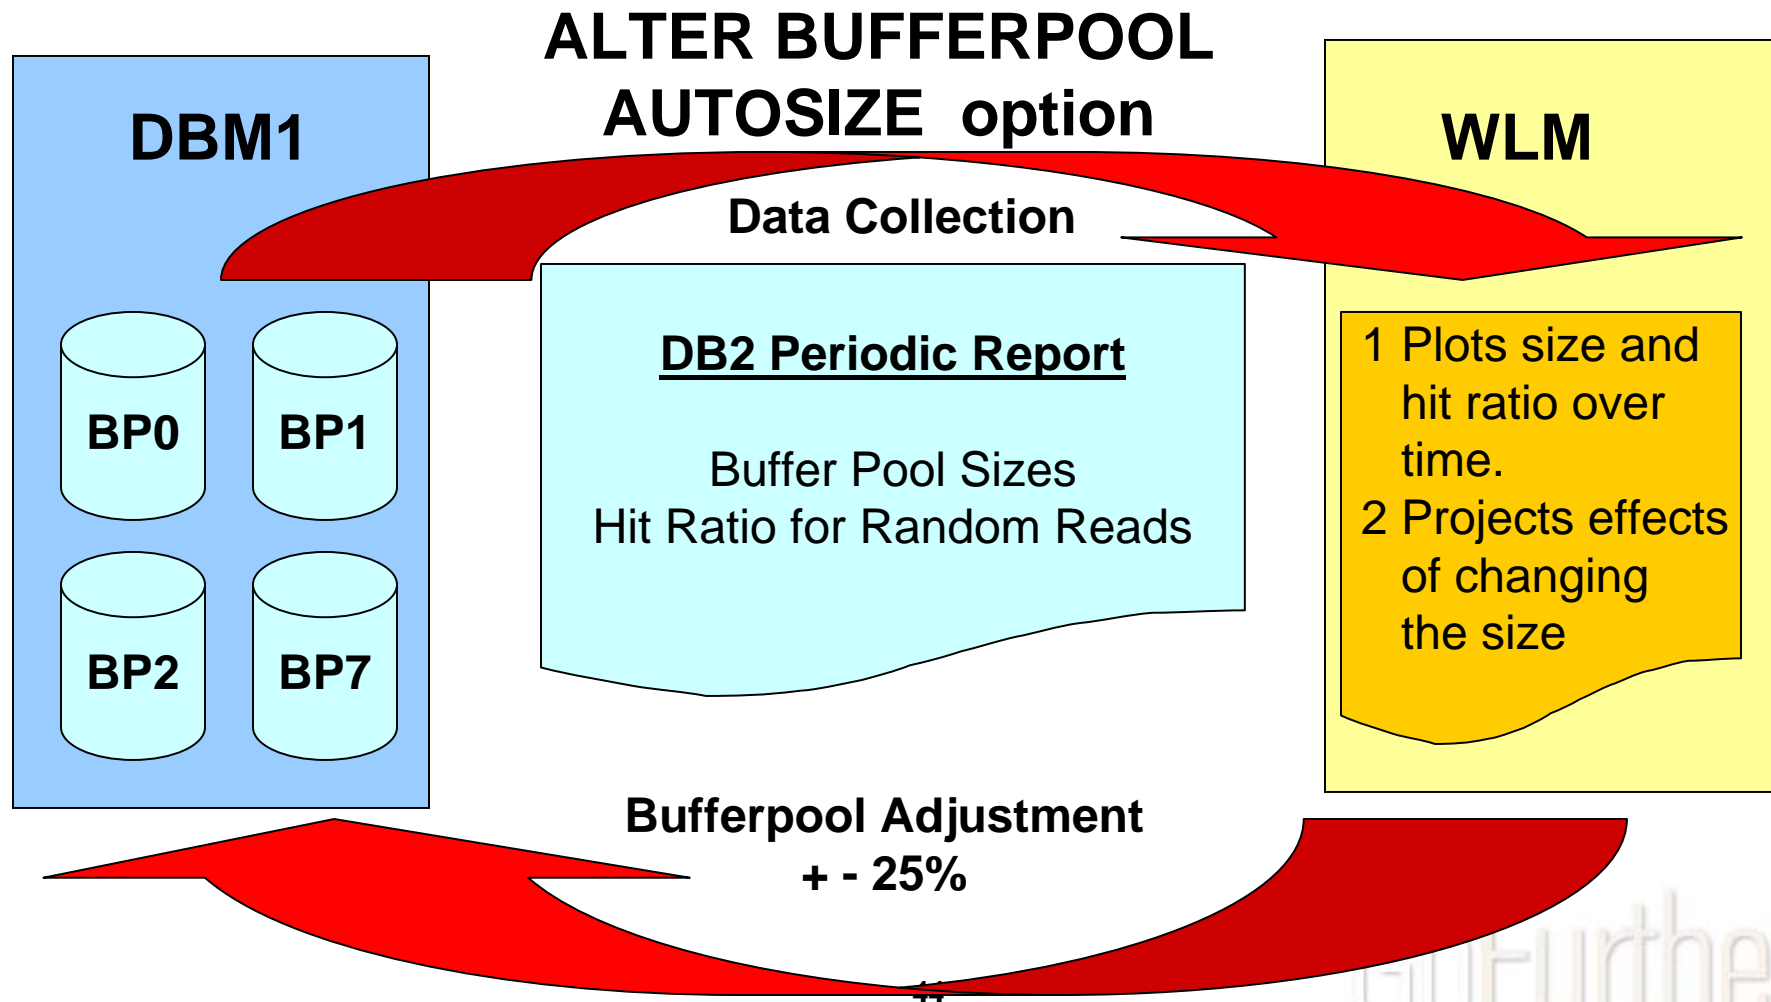
Converged TEMP Space

- 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

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

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 or more of savings expected
 - 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)
- Command line processor
- Optimization Service Center

DB2 9 for z/OS Innovation: SOA and XML

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

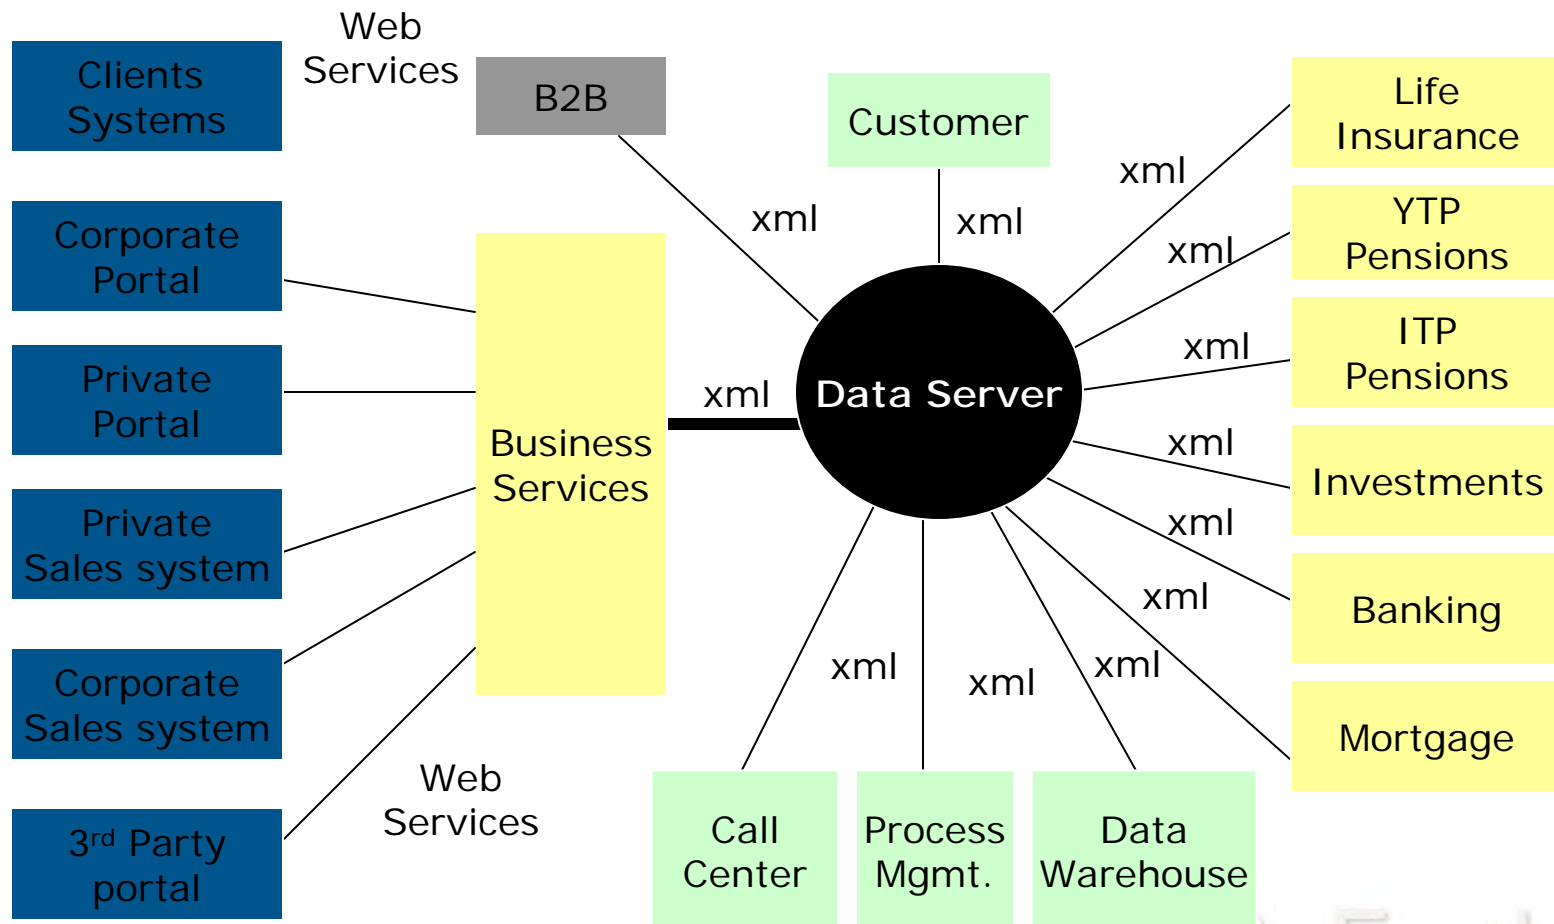


Optimistic Locking Support

- 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

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

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")
```

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

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

CLONE Tables

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

CLONE Tables...

- 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

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

- 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

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

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 LPAR₆₄ (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

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

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/fcgi-bin/ssialias?subtype=ca&infotype=an&appname=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.

