




# DB2 10 for z/OS

How can it help you?

<http://www.ibm.com/support/docview.wss?uid=swg27017960>





# Disclaimer/Trademarks

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

**The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion. \***

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AS/400, DataJoiner, DataPropagator, DB2, DB2 Connect, DB2 Extenders, DB2 OLAP Server, DB2 Universal Database, Distributed Relational Database Architecture, DRDA, eServer, IBM, IMS, iSeries, MVS, Net.Data, OS/390, OS/400, PowerPC, pSeries, RS/6000, SQL/400, SQL/DS, Tivoli, VisualAge, VM/ESA, VSE/ESA, WebSphere, z/OS, zSeries

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.



# DB2 for z/OS

The most robust and cost effective data server

## DB2

## DB2 9

## DB2 10



- Deep synergy with System z
- HW Compression
- Consolidation

- 20%-30% Utility CPU savings
- Compress indexes, save 50% disk
- More CPU on specialty engines

- Save up to 20% CPU batch & transactions
- On-the-fly data Compression
- Temporal data support
- Skip-level migration



- Unmatched availability
- Unparalleled security
- Industry leading reliability

- Flexible context and role security
- Expanded online schema changes
- Volume level backup & recovery

- Ten times more concurrent users
- More online schema changes
- More granular access control



- Near-linear scalability
- Optimized for SOA
- Flexible development
- Warehousing capabilities

- Seamless integration of XML and relational
- Improved SQL
- Partition by growth
- OLAP expressions

- Enhanced query parallelism
- More SQL compatibility
- Improved pureXML and SQL PL

**Beta Announced:  
Feb 9, 2010**



## Top 10 in DB2 10 for z/OS

1. CPU reductions for transactions, queries, & batch
2. Ten times more users by avoiding memory constraints
3. More concurrency for catalog, utilities, and SQL
4. More online changes for data definition, utilities and subsystems
5. Improved security with more granularity
6. Temporal or versioned data
7. SQL enhancements improve portability
8. pureXML performance and usability enhancements
9. Hash, index include columns, access path stability, skip migration
10. Productivity improved for database & systems administrators, and application programmers





# DB2 10 for z/OS At a Glance


<b>Performance, Scalability</b>	<ul style="list-style-type: none"><li>• CPU reductions out-of-the-box</li><li>• Hash access to data, index include columns</li><li>• Ten times more threads per DB2 image</li></ul>
<b>Availability Security Productivity</b>	<ul style="list-style-type: none"><li>• More online schema changes</li><li>• Improved concurrency: catalog, data, &amp; utilities</li><li>• Row and column access control, masking</li><li>• Administrator privileges with finer granularity</li><li>• Administration productivity enhancements</li></ul>
<b>Application Enablement</b>	<ul style="list-style-type: none"><li>• Versioned data or temporal queries</li><li>• pureXML enhancements</li><li>• Last committed reads</li><li>• SQL improvements that simplify porting</li></ul>
<b>Dynamic Warehousing</b>	<ul style="list-style-type: none"><li>• Moving sum, moving average</li><li>• Many query optimization improvements</li><li>• Query parallelism improvements</li><li>• IBM Smart Analytics Optimizer</li></ul>



## DB2 10 Performance Objective

- Historical goal under 5% performance regression
- Goal 5% -10% initial performance improvement
- Many customers reduce CPU time 10% - 20%

**Average %CPU improvements  
version to version**





## DB2 10 Performance, Scalability Objectives

- Significant scalability and performance improvements
  - Synergy with latest System z hardware & software
    - High n-way scalability
    - Large real memory exploitation
    - Hardware level optimization
  - Improve transaction times
  - Lower CPU usage for large & small DB2 subsystems
- Virtual storage is most common vertical scale constraint for large customers
  - Limited number of concurrent threads for a single member / subsystem
  - Address next tier of constraints: latches, concurrency



## DB2 10 for z/OS: Out-of-the-Box Savings

### Up to 20% CPU reductions for transactions, queries, and batch

- Out-of-the-box CPU reductions of 5-10% for traditional workloads
- Out-of-the box CPU reductions of up to 20% for new workloads
- Up to additional 10% CPU savings using new functions

### Scales with less complexity and cost

- 5-10x more concurrent users – up to 20,000 per subsystem
- Significant scale-up capabilities in addition to existing scale-out support
- Consolidate to fewer LPARs and subsystems

### Improved operational efficiencies and lower administration cost

- Automatic diagnostics, tuning, and compression

### Even better performance

- Elapsed time improvement for small LOBS and Complex Queries







## Performance Enhancements Requiring Few Changes (CM)

- SQL runtime improved efficiency
- Address space, memory changes to 64 bit, some REBINDs
- Faster single row retrievals via open / fetch / close chaining
- Distributed thread reuse High Performance DBATs
- DB2 9 utility enhancements in CM8
- Parallel index update at insert
- Workfile in-memory enhancements
- Index list prefetch
- Solid State Disk use
- Buffer pool enhancements (Utilize z10 1MB page size)



## Performance Enhancements requiring REBIND (CM)

- Most access path enhancements
- SQL paging performance enhancements
  - Single index access for complex OR predicates:
- IN list performance
  - Optimized Stage1 processing (single or multiple IN lists)
  - Matching index scan on multiple IN lists
- Query parallelism improvements
- More stage 2 predicates can be pushed down to stage 1
- More aggressive merge of views and table expressions
  - Avoid materialization of views
- REBIND enables further SQL runtime improvements
- If migrate from V8, get new RUNSTATS before mass rebind



## Performance Enhancements requiring NFM

- Efficient caching of dynamic SQL statements with literals
- Most utility enhancements
- LOB streaming between DDF and rest of DB2
- Faster fetch and insert, lower virtual storage consumption
- SQL Procedure Language performance improvements
- Insert improvement for UTS



## Performance Enhancements requiring NFM + DBA work

- Hash access path      Alter + Reorg + rebind to activate
- Index include columns    Alter + Rebuild + rebind to activate
- Inline LOBs              Alter (need UTS and RRF)
- MEMBER CLUSTER for UTS

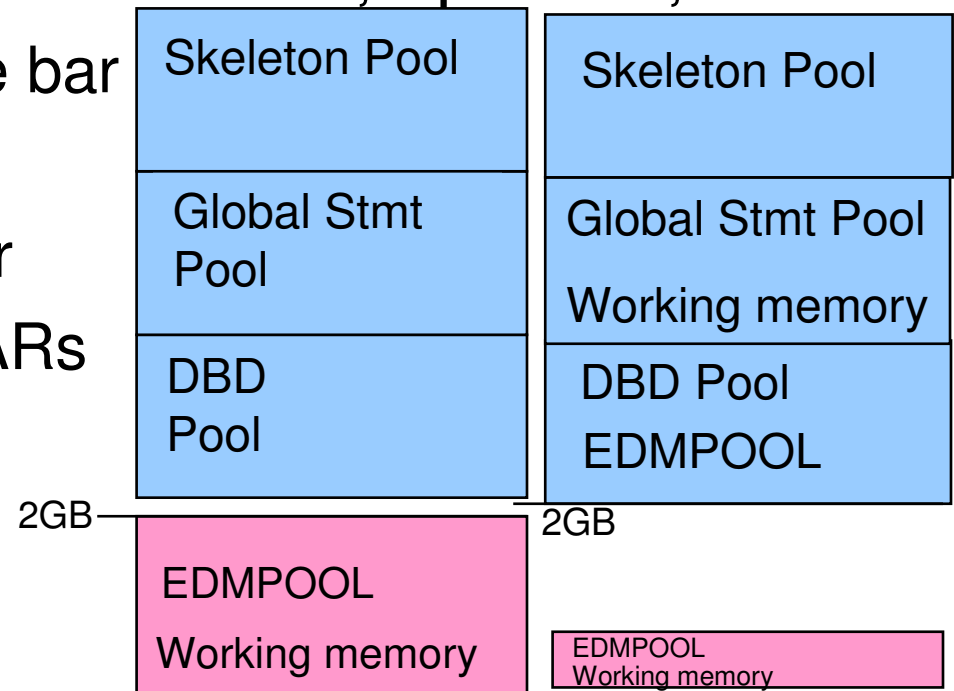


# DB2 10: 64 bit Evolution Virtual Storage Relief

DB2 9 helped (~ 10% – 15%)

DB2 10: 5 to 10 times more threads, up to 20,000

- Move 80% - 90% above bar
- More concurrent work
- Reduce need to monitor
- Able to consolidate LPARs
- Reduced cost
- Easier to manage
- Easier to grow




**Scalability: Virtual storage constraint is still an important issue for many DB2 customers.**




# Running a Large Number of Active Threads

## Today



- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s per LPAR

## DB2 10



- More threads per DB2 image
- More efficient use of large n-ways
- Easier growth, lower costs, easier management
- Data sharing and Parallel Sysplex still required for very high availability and scale
- Rule of thumb: save 1/2% CPU for each member reduced, more on memory



## Other System Scaling Improvements

- Other bottlenecks can emerge in extremely heavy workloads
  - several improvements planned to reduce latching and other system serialization contention
  - new option to for readers to avoid waiting for updaters
  - eliminate UTSERIAL lock contention for utilities
  - Use 64-bit common storage to avoid ECSA constraints
- Concurrent DDL/BIND/Prepare processes can contend with one another
  - restructure parts of DB2 catalog to avoid the contention
- SPT01 64GB limit can be a constraint, especially if package stability is enabled
  - Allow many more packages by using LOBs



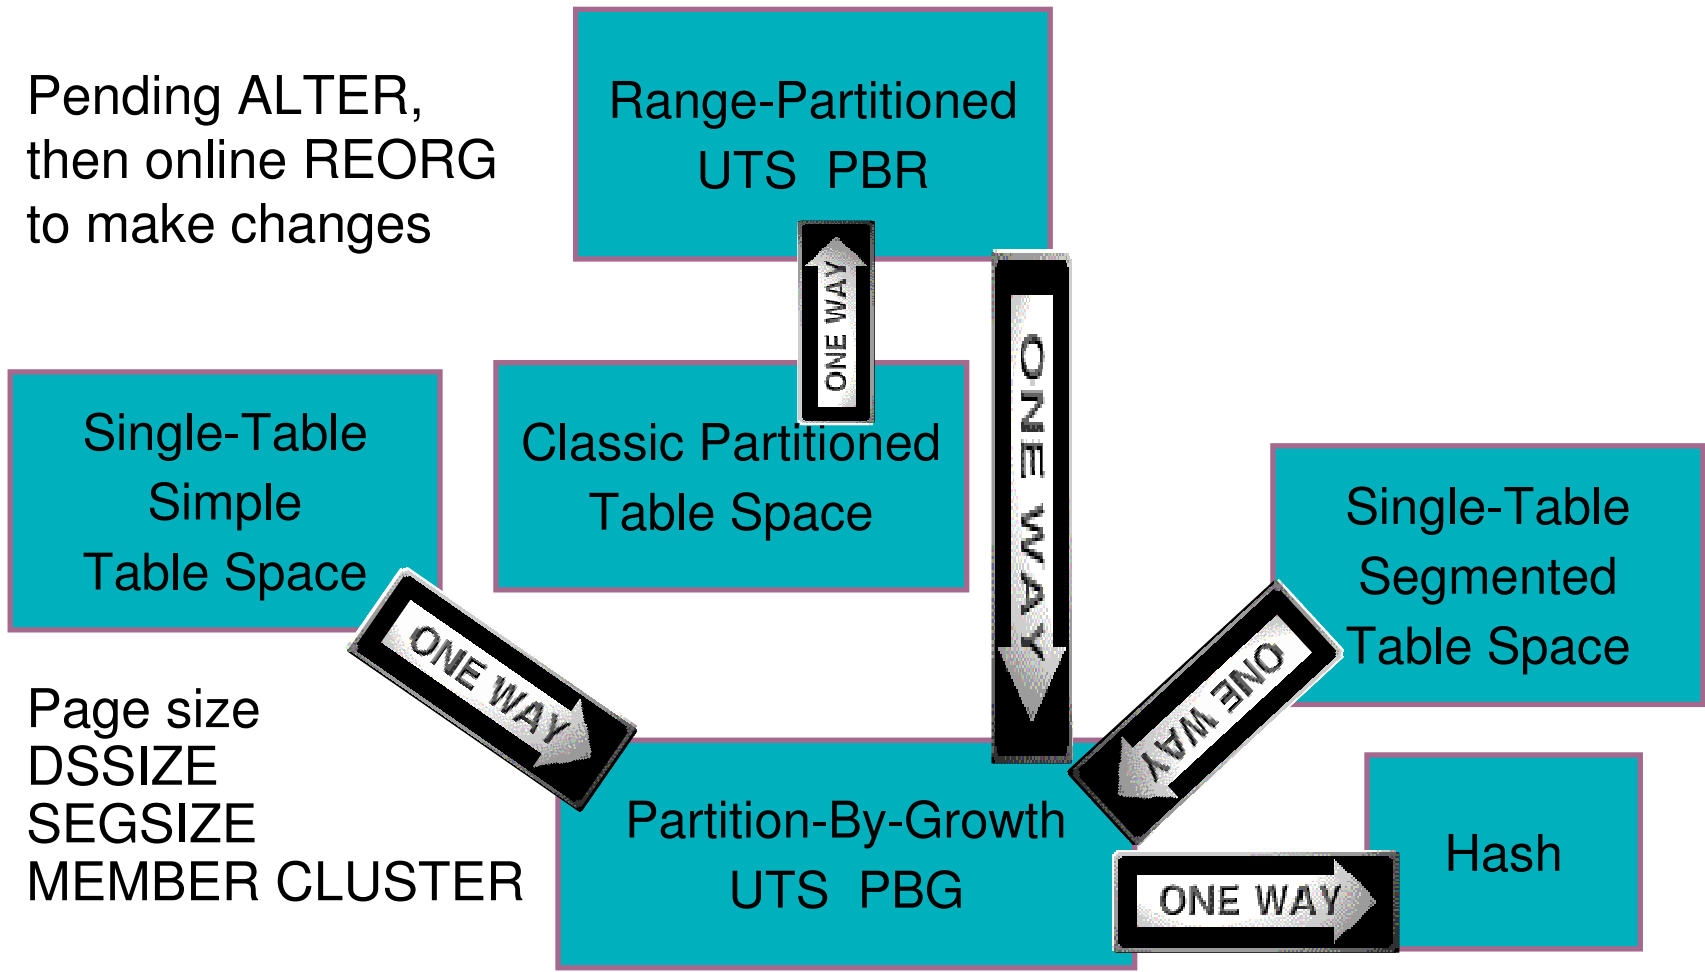
# DB2 10 for z/OS At a Glance

<b>Performance, Scalability</b>	<ul style="list-style-type: none"><li>• CPU reductions out-of-the-box</li><li>• Hash access to data, index include columns</li><li>• Ten times more threads per DB2 image</li></ul>
<b>Availability Security Productivity</b>	<ul style="list-style-type: none"><li>• More online schema changes</li><li>• Improved concurrency: catalog, data, &amp; utilities</li><li>• Row and column access control, masking</li><li>• Administrator privileges with finer granularity</li><li>• Administration productivity enhancements</li></ul>
<b>Application Enablement</b>	<ul style="list-style-type: none"><li>• Versioned data or temporal queries</li><li>• pureXML enhancements</li><li>• Last committed reads</li><li>• SQL improvements that simplify porting</li></ul>
<b>Dynamic Warehousing</b>	<ul style="list-style-type: none"><li>• Moving sum, moving average</li><li>• Many query optimization improvements</li><li>• Query parallelism improvements</li><li>• IBM Smart Analytics Optimizer</li></ul>





# Improved availability ALTER REORG





## Major changes in DB2 10 catalog & directory

- Improve availability and productivity
- Increase maximum size substantially
- Reduce contention: BIND, DDL, utilities
- Allow SELECT from SYSLGRNX



## DB2 10: Business Security & Compliance

- Protect sensitive data from privileged users & improve productivity
  - SYSADM & DBADM without data access
  - Usability: DBADM for all DB
  - Revoke without cascade
- Separate authorities to perform security related tasks, e.g. security administrator, EXPLAIN, performance monitoring and management
- Audit privileged users
- Row and column access control
  - Allow masking of value
  - Restrict user access to individual cells



***Use disk encryption***



# Optimization Stability and Control

Provide unprecedented level of stability for query performance by stabilizing access paths for

- Static SQL - Relief from REBIND regressions
- Dynamic SQL
  - Remove the unpredictability of PREPARE
  - Extend Static SQL benefits to Dynamic SQL



## DB2 10 for z/OS At a Glance

<b>Performance, Scalability</b>	<ul style="list-style-type: none"><li>• CPU reductions out-of-the-box</li><li>• Hash access to data, index include columns</li><li>• Ten times more threads per DB2 image</li></ul>
<b>Availability Security Productivity</b>	<ul style="list-style-type: none"><li>• More online schema changes</li><li>• Improved concurrency: catalog, data, &amp; utilities</li><li>• Row and column access control, masking</li><li>• Administrator privileges with finer granularity</li><li>• Administration productivity enhancements</li></ul>
<b>Application Enablement</b>	<ul style="list-style-type: none"><li>• <b>Versioned data or temporal queries</b></li><li>• <b>pureXML enhancements</b></li><li>• <b>Last committed reads</b></li><li>• <b>SQL improvements that simplify porting</b></li></ul>
<b>Dynamic Warehousing</b>	<ul style="list-style-type: none"><li>• Moving sum, moving average</li><li>• Many query optimization improvements</li><li>• Query parallelism improvements</li><li>• IBM Smart Analytics Optimizer</li></ul>



## DB2 10 Application Enablement and Portability

- Data versioning by date
- pureXML enhancements
- Large object improvements
  - Inline LOBs
  - Loading and unloading tables with LOBs
    - LOBs in input/output files with other non-LOB data
- Improved portability and SQL consistency
  - Currently committed locking semantics
  - Implicit casting or loose typing
  - Timestamp with time zone
  - Variable timestamp precision



## Versioned data or Temporal Data

- Table-level specification to control data management based upon time
- Two notions of time:
  - System time: notes the occurrence of a data base change
    - “row xyz was deleted at 10:05 pm”
    - Query at current or any prior period of time
    - Useful for auditing, compliance
  - Business time: notes the occurrence of a business event
    - “customer xyz’s service contract was modified on March 23”
    - Query at current or any prior/future period of time
    - Useful for tracking of business events over time, application logic greatly simplified
- New syntax in FROM clause to specify a time criteria for selecting historical data



# 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 SQL, join across encoding schemes, IS NOT DISTINCT FROM, session variables, range partitioning
- c** { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT
- l** { Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE or DELETE, multi-site join, MDC
- u**
- w**





## cross-platform SQL book V3

# DB2 SQL

z z/OS 9

common

luw Linux, Unix & Windows 9



- z** { Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE SQL, join across encoding schemes, IS NOT DISTINCT FROM, session variables, **TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE, SELECT from MERGE, index & XML compression**
- common** { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, **UPDATE 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, OmniFind, Spatial, range partitioning, compression**
- luw** { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, **XQuery**



# DB2 SQL

## cross-platform SQL book V3.1

z z/OS 9

common

luw Linux, Unix & Windows 9.5



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE SQL, join across encoding schemes, IS NOT DISTINCT FROM, TRUNCATE, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, index & XML compression

c  
o  
m  
m  
o  
n  
l  
u  
w

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 & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, OmniFind, spatial, range partitions, data compression, **session variables, DECIMAL FLOAT, optimistic locking, ROLE**

Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, **more** Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, **XML enhancements, array data type, global variables, more vendor syntax**



# DB2 SQL

z z/OS 9

common

luw Linux, Unix & Windows 9.7



- z { Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE SQL, join across encoding schemes, IS NOT DISTINCT FROM, VARBINARY, FETCH CONTINUE, 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 & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, OmniFind, spatial, range partitions, data compression, session variables, DECIMAL FLOAT, optimistic locking, ROLE, **TRUNCATE, index & XML compression, created temps**
- l { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, **more** Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, XML enhancements, array data type, global variables, **even more vendor syntax, LOB & temp table compression, inline LOB, administrative privileges, implicit casting, date/time changes, currently committed**



# DB2 SQL

z z/OS 10

common

luw Linux, Unix & Windows 9.7



- z { Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE SQL, join across encoding schemes, IS NOT DISTINCT FROM, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, **data versioning, access controls**
- 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 & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, OmniFind, spatial, range partitions, data compression, session variables, DECIMAL FLOAT, optimistic locking, ROLE, TRUNCATE, index & XML compression, created temps, **inline LOB, administrative privileges, implicit casting, date/time changes, currently committed, moving sum & avg.**
- l { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, more Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, XML enhancements, array data type, global variables, even more vendor syntax, LOB & temp table compression,
- u
- w



# pureXML improved performance and usability

## Integrated XML Support

- **XML schema validation in the engine for improved usability, performance**
  - XML schema association with XML columns
  - Using z/OS XML System Services, 100% zIIP / zAAP eligible
- **Native XML Date and Time in business processing**
  - xs:date, xs:dateTime, and xs:dateTime support and XML index support
- **Allow easy update of sub-parts of an XML document**
- **XML support in SQL PL stored procedures & user defined functions**
- **Performance enhancements**



# DB2 10 for z/OS At a Glance

<b>Performance, Scalability</b>	<ul style="list-style-type: none"><li>• CPU reductions out-of-the-box</li><li>• Hash access to data, index include columns</li><li>• Ten times more threads per DB2 image</li></ul>
<b>Availability Security Productivity</b>	<ul style="list-style-type: none"><li>• More online schema changes</li><li>• Improved concurrency: catalog, data, &amp; utilities</li><li>• Row and column access control, masking</li><li>• Administrator privileges with finer granularity</li><li>• Administration productivity enhancements</li></ul>
<b>Application Enablement</b>	<ul style="list-style-type: none"><li>• Versioned data or temporal queries</li><li>• pureXML enhancements</li><li>• Last committed reads</li><li>• SQL improvements that simplify porting</li></ul>
<b>Dynamic Warehousing</b>	<ul style="list-style-type: none"><li>• Moving sum, moving average</li><li>• Many query optimization improvements</li><li>• Query parallelism improvements</li><li>• IBM Smart Analytics Optimizer</li></ul>




## DB2 10 Query Enhancements

- CPU time reductions for queries, batch, & transactions
- SQL enhancements: Moving Sum, Moving Average
- Optimization techniques
  - Remove parallelism restrictions and more even parallel distribution
    - increased zIIP use
  - In-memory techniques for faster query performance
  - Access path stability and control
- Advanced query acceleration techniques
  - IBM Smart Analytics Optimizer



# DB2 10 for z/OS

Planning to migrate







## DB2 10 for z/OS: Skip-Level Migration

May move from V8 to DB2 10,  
but just because you can, doesn't mean you always should....

Key considerations:

- Risk/reward analysis
  - What's the risk? Tolerance level?
  - How will you do it? What's your mitigation plan? Are ISVs ready?
  - What workloads do you need to test and can you test them properly?
  - Am I missing out on DB2 9 value in the meantime?
- May not see large migration cost savings
  - Expect 20% to 25% cost savings versus two migrations
  - Larger migration project, longer migration timeline, more risk
  - Applications and ISVs may not be ready



If you are on V7 or earlier, go to V8


If you plan to migrate in 2010 or 2011, go to DB2 9

If you are on V8 for over 2 years, go to DB2 9



## Why Migrate to DB2 9 for z/OS?

- Business needs to save money
  - Reduce CPU time & disk space
  - Improve business agility
  - Service Oriented Architecture
- Application developers need
  - PureXML for a powerful SQL and XML interface to XML data
  - Powerful new SQL enhancements
  - Portability with SQL and data definition compatibility
- Database Administrators need
  - Improve availability and performance
  - More flexible security and easier regulatory compliance
  - Better web application & data warehouse function and performance
  - LOB function, performance, usability





## Key details about DB2 10: getting ready

Prerequisites: migrate from DB2 9 for z/OS or DB2 for z/OS V8

- z/OS V1.10 SMS-managed DB2-managed DB2 catalog
- System z10, z9, z890, z990, and above (no z800, z900)
- DB2 Connect 9 FP1, 9.7 FP3 for many 10 functions, FP2 beta
- IMS 10 & 11 (not 9) CICS compilers (See announcement)
- SPE PK56922 PK69411 PK61766 PK85956 PM04680 PK87280 PK87281
- Premigration check DSNTIJPA PM04968

Items deprecated in earlier versions eliminated: more for V8 mig.

- Private protocol → DRDA (DSNTP2DP, PK92339, PK64045)
- Old plans and packages V5 or before → REBIND
- Plans containing DBRMs → packages PK62876 PK79925 (V8)
- ACQUIRE(ALLOCATE) → ACQUIRE(USE)
- Old plan table formats → DB2 V8 or 9, Unicode, 59 cols PK85068
- BookManager use for DB2 publications → Info Center, pdf




## No longer supported in DB2 10

- Private protocol: Convert to DRDA.
- EXPLAIN tables prior to Version 8 format: Alter add cols.
- Plans containing DBRMs. Acquire allocate. Old packages.
- DB2 catalog tables are DB2-managed & SMS-managed. No links. More LOBs & table spaces. No SPT01 compress.
- REORG TABLESPACE SHRLEVEL NONE on LOB table spaces. Use SHRLEVEL CHANGE or REFERENCE.
- DB2 MQ XML functions: Use pureXML functions.
- DB2 XML Extender: Use pureXML.
- Some subsystem parameters removed, many changed
- Accessories Suite parts Optimization Service Center, Developer Workbench, Visual Explain for DB2 → Data Studio



# Questions?





## DB2 9 and 10 in **IBM Redbooks Publications**

1. DB2 9 Technical Overview SG24-7330
2. DB2 9 Performance Topics SG24-7473 updated Dec. 2009
3. DB2 9 Stored Procedures SG24-7604
4. Index Compression with DB2 9 for z/OS redp4345
5. SQL Reference for Cross-Platform Development
6. Enterprise Database Warehouse, SG24-7637
7. 50 TB Data Warehouse on System z, SG24-7674
8. New Tools for Query Optimization SG24-7421
9. LOBs with DB2 for z/OS SG24-7270
10. Deploying SOA Solutions SG24-7663
11. Enhancing SAP - DB2 9 SG24-7239
12. SAP Application on Linux z SG24-6847
13. Best practices SAP BI - DB2 9 SG24-6489-01
14. Data Sharing in a Nutshell, SG24-7322
15. Securing DB2 & MLS z/OS SG24-6480-01
16. Data Sharing: Distributed Load Balancing & Fault Tolerant Configuration redp4449
17. Considerations on Small & Large Packages redp4424
18. Backup and Recovery Considerations redp4452
19. Powering SOA with IBM Data Servers SG24-7259
20. Packages Revisited, SG24-7688
21. Data Studio V2.1 Web Services redp4510
22. Ready to Access Solid-State Drives redp4537
23. Distributed Functions SG24-6952
24. Buffer Pool Monitoring & Tuning redp4604
25. Securing & Auditing Data SG24-7720
26. Serialization and Concurrency SG24-4725-01 new
27. Utilities SG24-6289-01 draft

