



DB2 9 for z/OS and Beyond

The future runs on System z



Kevin Harrison

Senior IT architect/Software Engineer

IBM Silicon Valley Lab

© 2008 IBM Corporation

What's new in DB2 for z/OS?

- Synergy with DB2 for z/OS and System z
- What does DB2 9 have for me?
- What are DB2 9 performance characteristics?
- How is database administration improved?
- What are the improvements for application programming? Is XML for me?
- What is the vision beyond DB2 9? What early planning can I do for DB2 9 and beyond?

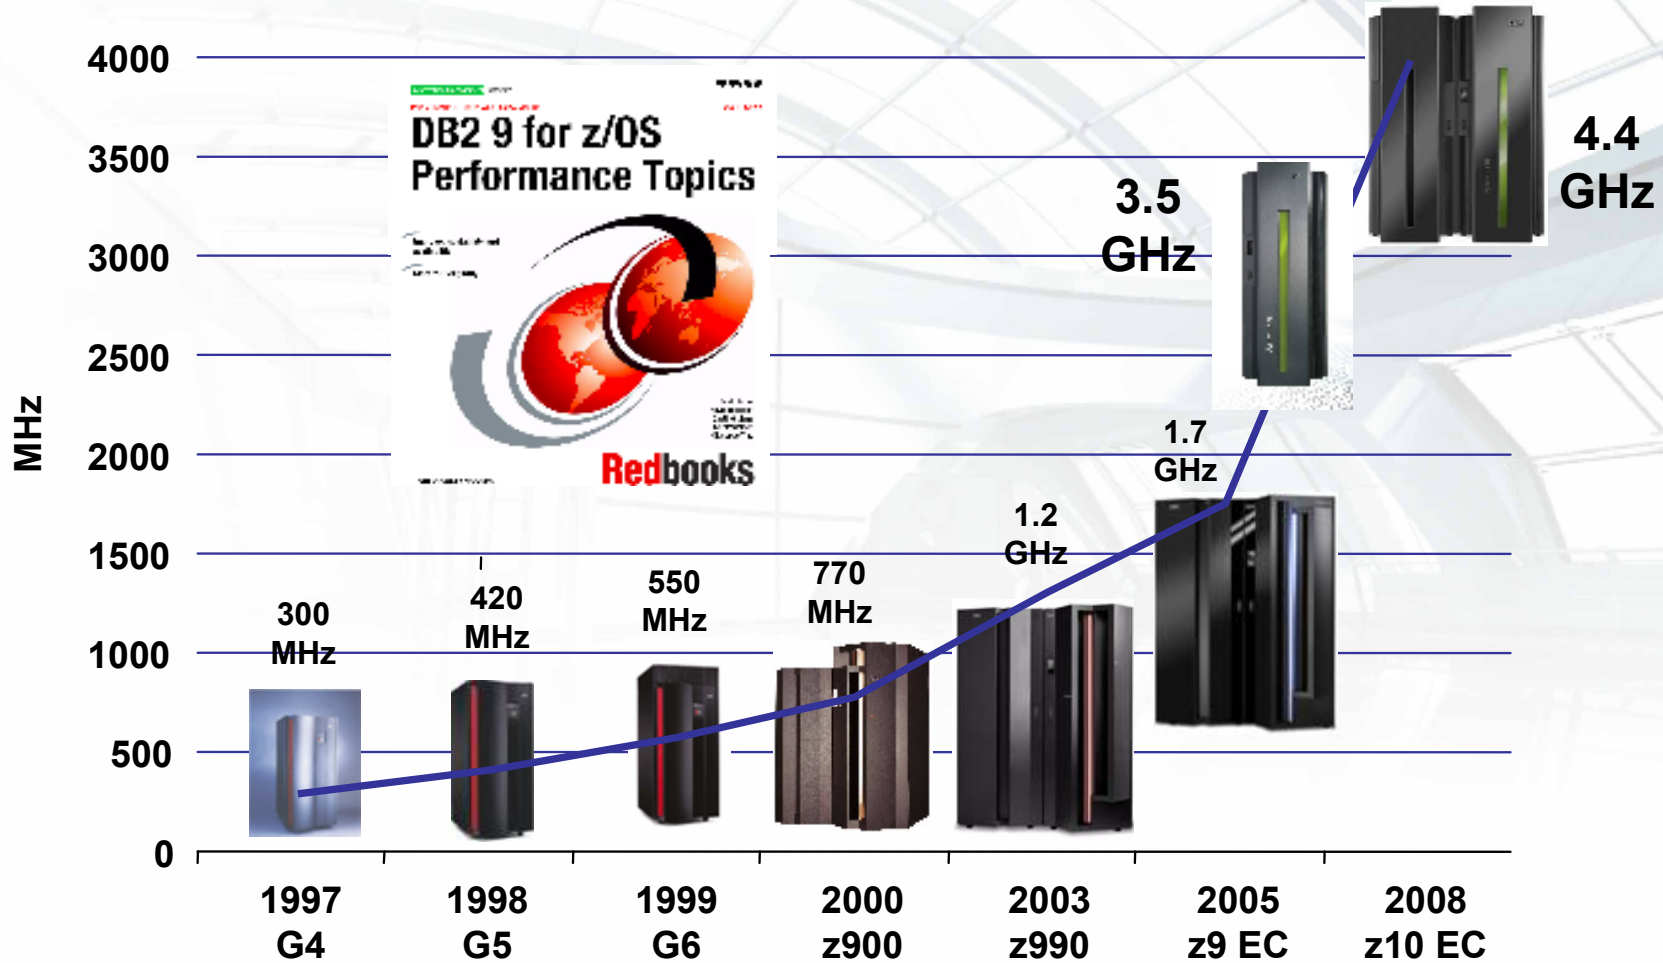
DB2 Deep Synergy With System z

Key integration points include:

- ▶ Data sharing (availability and scale out)
- ▶ Hardware data compression & encryption
- ▶ zIIP specialty engines
- ▶ Unicode conversion
- ▶ Encrypted TCP/IP communication (SSL), encrypted data
- ▶ Cross-memory, memory protection keys
- ▶ Sorting
- ▶ Multi-core, large N-way
- ▶ 1 MB page size (z10)
- ▶ Decimal float arithmetic (z10)
- ▶ 64-bit addressing and large memory
- ▶ z/OS Workload Manager
- ▶ z/OS Security Server (RACF)
- ▶ z/OS RRS integrated commit coordinator
- ▶ Solid state disks



IBM z10 Extends Scale and Value



- G4 - 1st full-custom CMOS S/390®
- G5 - IEEE-standard BFP; branch target prediction
- G6 – Copper Technology (Cu BEOL)
- z900 - Full 64-bit z/Architecture®
- z990 - Superscalar CISC pipeline
- z9 EC - System level scaling
- z10 EC – Architectural extensions

DB2 for z/OS

Lowering TCO

- **Maximum value for dollar investment**

- ▶ Hardware pricing
 - CPU saving specialty engines (zIIP, zAAP..)
 - Compression of disk space (data, index)
- ▶ Software pricing
 - Reduction for tiers
 - Parallel Sysplex aggregation
 - z990, z9, z10 technology dividend
 - 10% reduction in charge units for each step
 - zNALC, Value Unit Edition, Subcapacity pricing



- **CPU + Memory + I/O and disk + Software + Energy and floor space + People = Improved Total Cost of Ownership (TCO)**

Helping to drive down the cost of IT *Now even more workloads can benefit from zIIP*

- Integrate data across the enterprise, optimize resources and lower the cost of ownership
 - ▶ Centralized data serving
 - ▶ Serving XML data
 - ▶ Use by ISVs
 - ▶ New HiperSockets™ for large messages
 - ▶ New IBM GBS Scalable Architecture for Financial Reporting™
- zIIPs offer economics to help you
 - ▶ **PLUS** zIIP price same for z10 EC as z9 EC



***IBM System z10 Integrated Information Processor and
IBM System z9 Integrated Information Processor***

DB2 9 for z/OS

Addressing Corporate Data Goals

<p>Application Enablement</p>	<ul style="list-style-type: none"> • pureXML • Optimistic locking for WebSphere • LOB performance, usability • Native SQL procedure language • SQL improvements that simplify porting
<p>RAS, Performance, Scalability, Security</p>	<ul style="list-style-type: none"> • More online schema changes • Online REBUILD INDEX, Online REORG improvements, Clone tables • Trusted context and ROLES • Parallel Sysplex clustering improvements • 64-bit virtual storage improvements
<p>Simplification, Reduced TCO</p>	<ul style="list-style-type: none"> • Index compression • Partition By Growth tables • Package stability • Volume based backup / recovery • Automatic object creation
<p>Dynamic Warehousing</p>	<ul style="list-style-type: none"> • Many SQL improvements • Dynamic index ANDing • Histogram statistics • New built-in OLAP expressions • Optimization Service Center

DB2 SQL 2004

z z/OS V8

common

luw Linux, Unix & Windows V8.2



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, range partitioning

c
o
m
m
o
n

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
u
w

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 2007

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 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**
o
m
m
o
n { 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, range partitioning, compression**
- l**
u
w { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, **XQuery**

DB2 SQL new cross platform reference

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 for SQL, join across encoding schemes, IS NOT DISTINCT FROM, TRUNCATE, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, index compression
- 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, range partitions, data compression, **session variables, DECIMAL FLOAT, optimistic locking, ROLE**
- 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, vendor syntax**
- u**
- w**

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 for 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, 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**
- u**
- w**

What You Can Do with pureXML

- Create tables with XML columns or alter table add XML columns
- Insert XML data, optionally validated against schemas
- Create indexes on XML data
- Efficiently search XML data
- Extract XML data
- Decompose XML data into relational data
- Construct XML documents from relational and XML data
- All the utilities and tools support for XML



Universal Table Space

- ▶ Combination of segmented with partitioning options
 - Better space management
 - Support of mass deletes / TRUNCATE
- ▶ If partitioned
 - Still must be one table per table space
 - Can choose [Range Based partitioning](#) (as before: PBR)
 - Can choose [Partitioned By Growth](#) (PBG)
- ▶ DROP / CREATE to migrate existing page sets
- ▶ Simple table spaces can not be created
 - Default table space is now Segmented (CM) or PGB (NFM)


LOB Improvements

- Progressive Streaming for LOB Locator Values
 - ▶ DB2 uses LOB size to determine whether to send LOB data to Java or DB2 CLI clients in one go (<32KB), in chunks (<1MB) or as LOB locator (>=1MB)
 - Transparent to application using LOB locators
- FETCH CONTINUE
 - ▶ Allows applications to retrieve LOB/XML data in pieces without the use of locators
- File reference variables
 - ▶ A file reference variable allows direct transfer of LOB data between DB2 and the file named in the variable
- Utility Changes
 - ▶ REORG LOB reclaim space
 - ▶ LOAD / Cross load LOB column lengths > 32KB in V7, V8 APARs
 - ▶ Logging for > 1GB LOBs
 - ▶ Online CHECK LOB and DATA
- Elimination of LOB locks for improved availability and performance



Indexing Enhancements

- Larger index pages allow for more efficient use of storage
 - ▶ Fewer page splits for long keys
 - ▶ More key values per page
- Index compression provides page-level compression
 - ▶ Data is compressed to 4K pages on disk
 - ▶ 32K/16K/8K pages results in up to 8x/4x/2x disk savings
 - ▶ No compression dictionaries
 - Compression on the fly
 - No LOAD or REORG required
- Rebuild Index SHRLEVEL CHANGE
- Define RANDOM index keys to avoid hot spots with multiple processes inserting sequential keys



Data Sharing DB2 9 Enhancements

- Log latch contention relief
- Restart performance enhancements
 - ▶ Reduced impact of retained locks – released as rollbacks are completed
 - ▶ Open data sets ahead of log apply
- Command to remove GBP-dependency at object level
 - ▶ ACCESS DB MODE(NGBPDEP)
 - ▶ Typical usage would be before batch run
 - ▶ Command to “prime” open data set
 - ▶ ACCESS DB MODE(OPEN) [PART]
- Auto-recover GRECP/LPL objects on group restart
 - ▶ Useful in Disaster Recovery or GDPS scenarios
- DB2 overall health taken into account for WLM routing
- Balance group attach connections across multiple members on same LPAR (V7 & V8 usermod)
- Group wide outage no longer needed for new LOB locking protocol (apar)

Utilities Highlights


- More online utilities
 - ▶ Rebuild Index SHRLEVEL CHANGE
 - ▶ Reorg LOB now supports SHRLEVEL REFERENCE (space reclamation)
 - ▶ Check data, LOB and repair locate ... SHRLEVEL CHANGE
 - ▶ Check index SHRLEVEL REFERENCE supports parallel for > 1 index
 - ▶ Clones for “online LOAD REPLACE”
- Online REORG BUILD2 phase elimination
- REORG parallelism for UNLOAD, RELOAD, LOG phases
- Utility TEMPLATE switching
- UNLOAD SKIP LOCKED DATA option

Utilities Highlights...

- RECOVER to any point-in-time with consistency
- MODIFY Recovery enhancements
 - ▶ “*Retain*” keyword added to improve management of copies
 - LAST(n), LOGLIMIT, GDGLIMIT
- Volume-based COPY/RECOVER (BACKUP SYSTEM/RESTORE SYSTEM)
 - ▶ RECOVER modified to enable object-level recovery from volume FlashCopy
 - ▶ Full integration of tape into BACKUP/RESTORE SYSTEM utilities
 - ▶ Incremental FlashCopy, APAR PK41001
- Truncate log based on timestamp
- RECOVER RESTOREBEFORE to use an earlier image copy
- Display progress of RECOVER during log apply
- COPY CHECKPAGE option always active
 - ▶ “Copy Pending” avoided if broken page encountered
- COPY SCOPE PENDING to copy only objects in “Copy Pending”

Why Migrate to DB2 9 for z/OS?

- Business needs
 - 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




Why is migration easier to DB2 9 for z/OS?

- Migration process enhancements: ENFM speed, CM*
- Much less performance regression:
 - ▶ Earlier improvements
 - ▶ Package stability & tools for avoiding access path issues
- CCSIDs and old product issues resolved in V8
- Simpler virtual storage considerations
- Less impact from incompatible changes
- Earlier deliveries from vendors



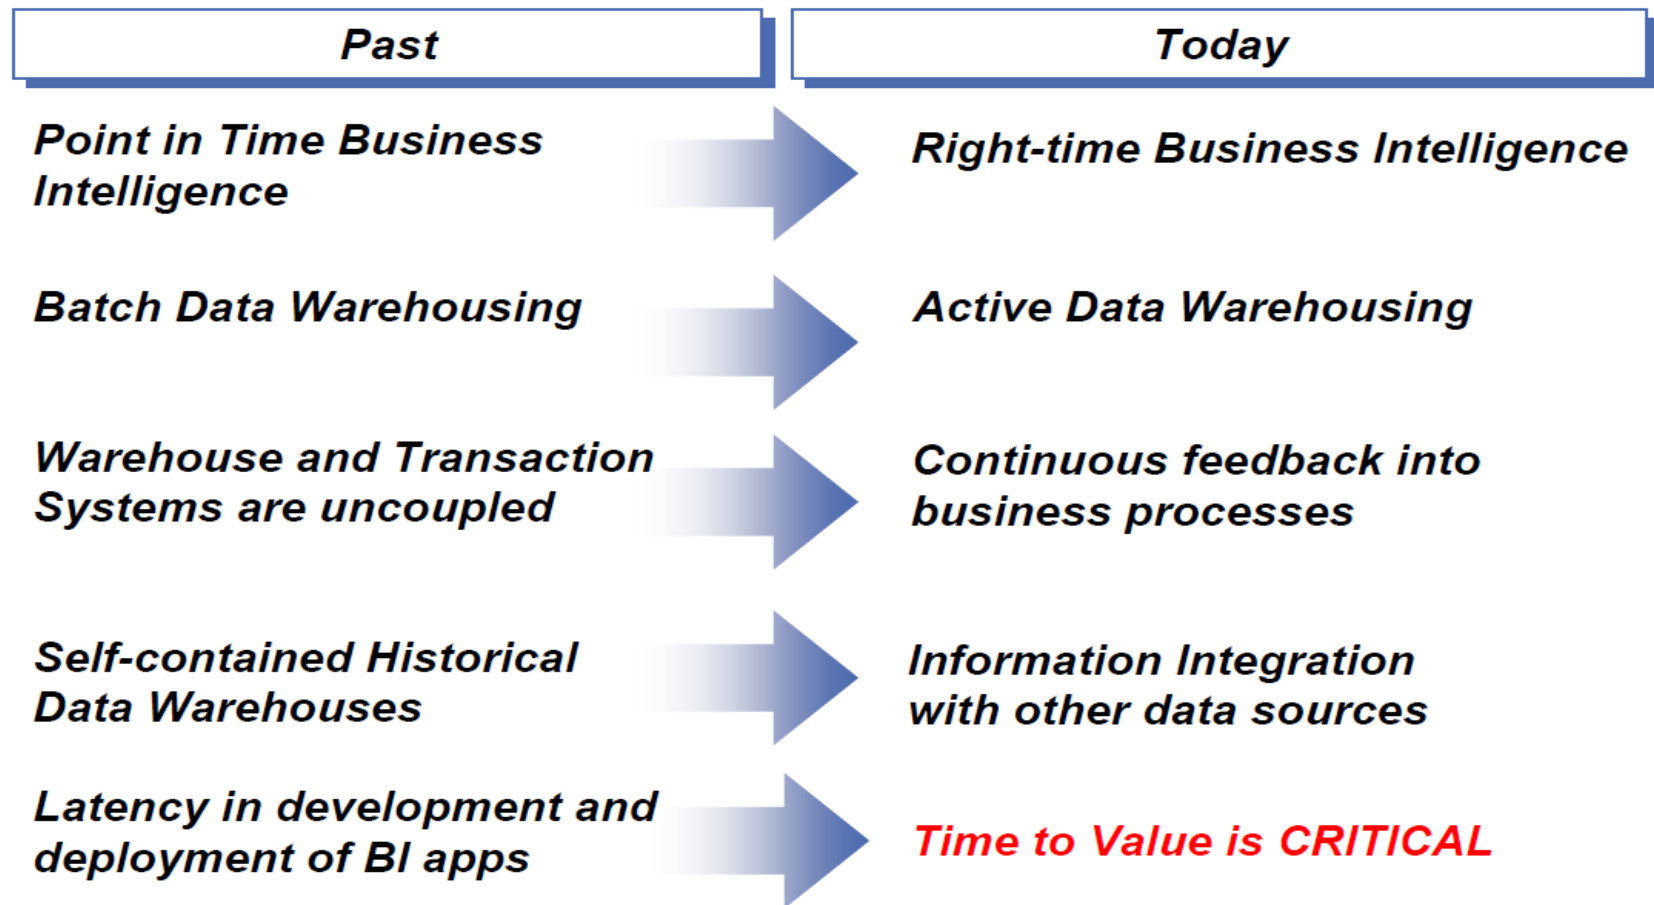
Most consumable DB2 9 improvements

- CM very little to no action:
 - Utility CPU reductions
 - Logging improvements
 - Improved index page split
 - Larger prefetch, write & preformat quantities
 - LOB performance
 - DDF VSCR
 - Optimization Service Center, Opt. Expert, & Data Studio
 - Changed online REORG
 - Improved RUNSTATS
 - Optimization improvements, EDMPOOL VSCR
- NFM
- LOB lock avoidance
 - Reordered row format
 - Index: larger page sizes, compression, index on expression



Data Warehouse & Business Intelligence Trends

Business Intelligence Becoming Mission-Critical



Data Warehousing on z/OS – What is driving this?

- Customer commitment to the z platform
 - ▶ Customers want to protect their significant investment in System z
 - ▶ TCO can be reduced through the utilization of existing processors, people, practices
 - ▶ TCO may also be achieved through a consolidation approach
- New BI trends are changing the DBMS landscape
 - ▶ The distinction is blurring between warehouse and OLTP databases based on new trends such as Dynamic Warehouse and Operational BI, driving:
 - The need for increased reliability, availability, security, and compliance in a DWH DBMS
 - The need for very current warehouse data, where proximity to the source provides an advantage
- Many z customers already have a warehouse on DB2 for z/OS
 - ▶ This drives requirements into hardware and software, which in turn drives a trend
 - ▶ DB2 has responded with increased functionality and performance; hardware changes are driving down costs
- Specialty processors provide new ways to optimize TCO
 - ▶ zIIPs and IFLs are driving down hardware and software costs; DWH/BI can make excellent use of these processors, ultimately driving TCO advantages

InfoSphere Warehouse on System z

New Information On Demand Software for System z

Better business decisions,
faster and with a lower overall TCO

InfoSphere Warehouse on System z adds to :

- DB2 for z/OS
- InfoSphere Information Server for System z
- InfoSphere MDM Server for System z
- Cognos 8 BI for System z

DB2 Developer Workbench → Data Studio → Optim Studio tools

before

now

IBM DB2 Developer Workbench V9.1

- SQL Query Editor
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

IBM Data Studio

- Integrated Query Editor – SQL + XQuery
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

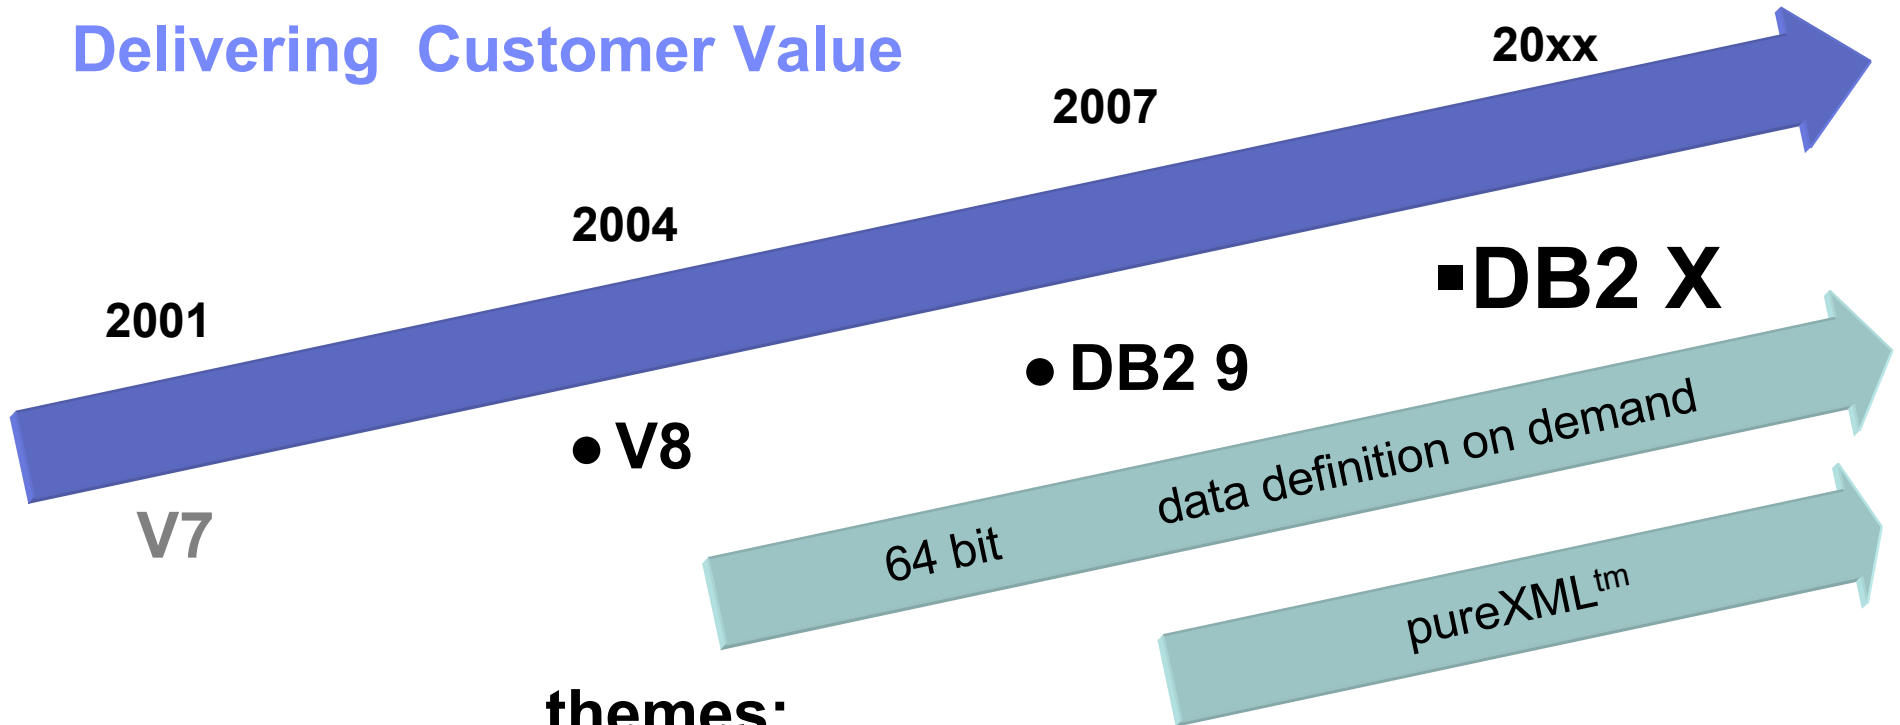
***Data Studio is a full replacement of
DB2 Developer Workbench
plus much more***

- ***DB2 for Linux, Unix, Windows v9.1, v9.5, v9.7***
- ***DB2 for z/OS v7, v8, 9***
- ***DB2 for i v5r2, v5r3, v5r4, v6***
- ***Informix Dynamic Server (IDS) v9, v10, v11***

- ER Diagramming
- Data Distribution Viewer
- Object Management
- Browse & Update Statistics
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Server Support
- Health Monitoring DB2 for LUW 9.5 and DB2 9 for z/OS

DB2 for z/OS Into the Future


Delivering Customer Value



themes:

Performance Scalability
 Reliability Availability Serviceability
 Security Productivity
 Application Development
 SQL XML SOA

DB2 X: Preliminary Performance Plan → Significant CPU Reductions: Best on z10



Transactions	DB design changes
Batch	Hash access
REBIND	application changes
	SQL adjustments

Your situation will vary. Less CPU is better.

Processors z890, z990, z9 and later z/OS 1.10 and later

DB2 X Performance, Scalability Objectives

- Provide significant Scalability and Performance improvements
 - ▶ Important objective for DB2 X
 - ▶ Synergy with latest System z and follow-on machines
 - High n-way scalability
 - Large real memory use
 - Hardware level optimization (cache & instructions)
 - ▶ Synergy with z/OS 1.10 software and beyond
 - ▶ Improve transaction times
 - ▶ Lower CPU usage for large & small DB2 subsystems

Performance

- Hash access path
- Parallel index update at insert
- Faster single row retrievals
- Inline LOBs
- LOB streaming between DDF and rest of DB2
 - ▶ Faster fetch and insert, reduced virtual storage
- DEFINE NO for LOBs (and XML)
- MEMBER CLUSTER for UTS
- Efficient dynamic SQL statement caching with literals

Performance

- Index include columns
- Workfile spanned records, PBG support, and in-memory enhancements
- Buffer pool enhancements
 - ▶ Utilize z10 1MB page size
 - ▶ “Fully in memory” option
- Internal performance optimizations
 - ▶ Improved cpu cache performance
 - ▶ Use new hardware instructions
 - ▶ Streamline performance-critical paths
- Solid state disk

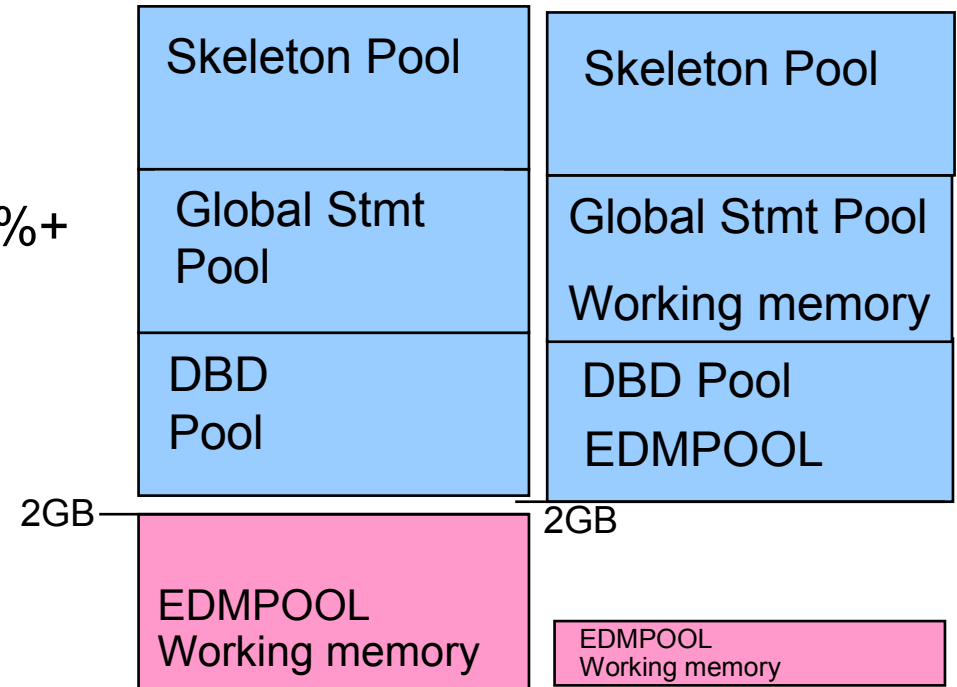
Query Performance

- Index include columns
- Hash access
- Safe query optimization
- More effective optimizer query rewrite
- Minimize materialization and size of intermediate results
- IN list performance
- UNION ALL performance
- RID pool usage enhancements
- Query parallelism enhancements: lifting restrictions, improving efficiency
- Workfile spanned records, PBG support, and in-memory enhancements
- Use of solid state disk

DB2 X: 64 bit Evolution (Virtual Storage Relief)


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

- DB2 9 helped (~ 10% – 15%)
- DB2 X expect to move 80% - 90%+
 - ▶ More concurrent work
 - ▶ Reduce need to monitor
 - ▶ Able to consolidate LPARs
 - ▶ Reduced cost
 - ▶ Easier to manage
 - ▶ Easier to grow




Running a Large Number of Active Threads

Today



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

DB2 X



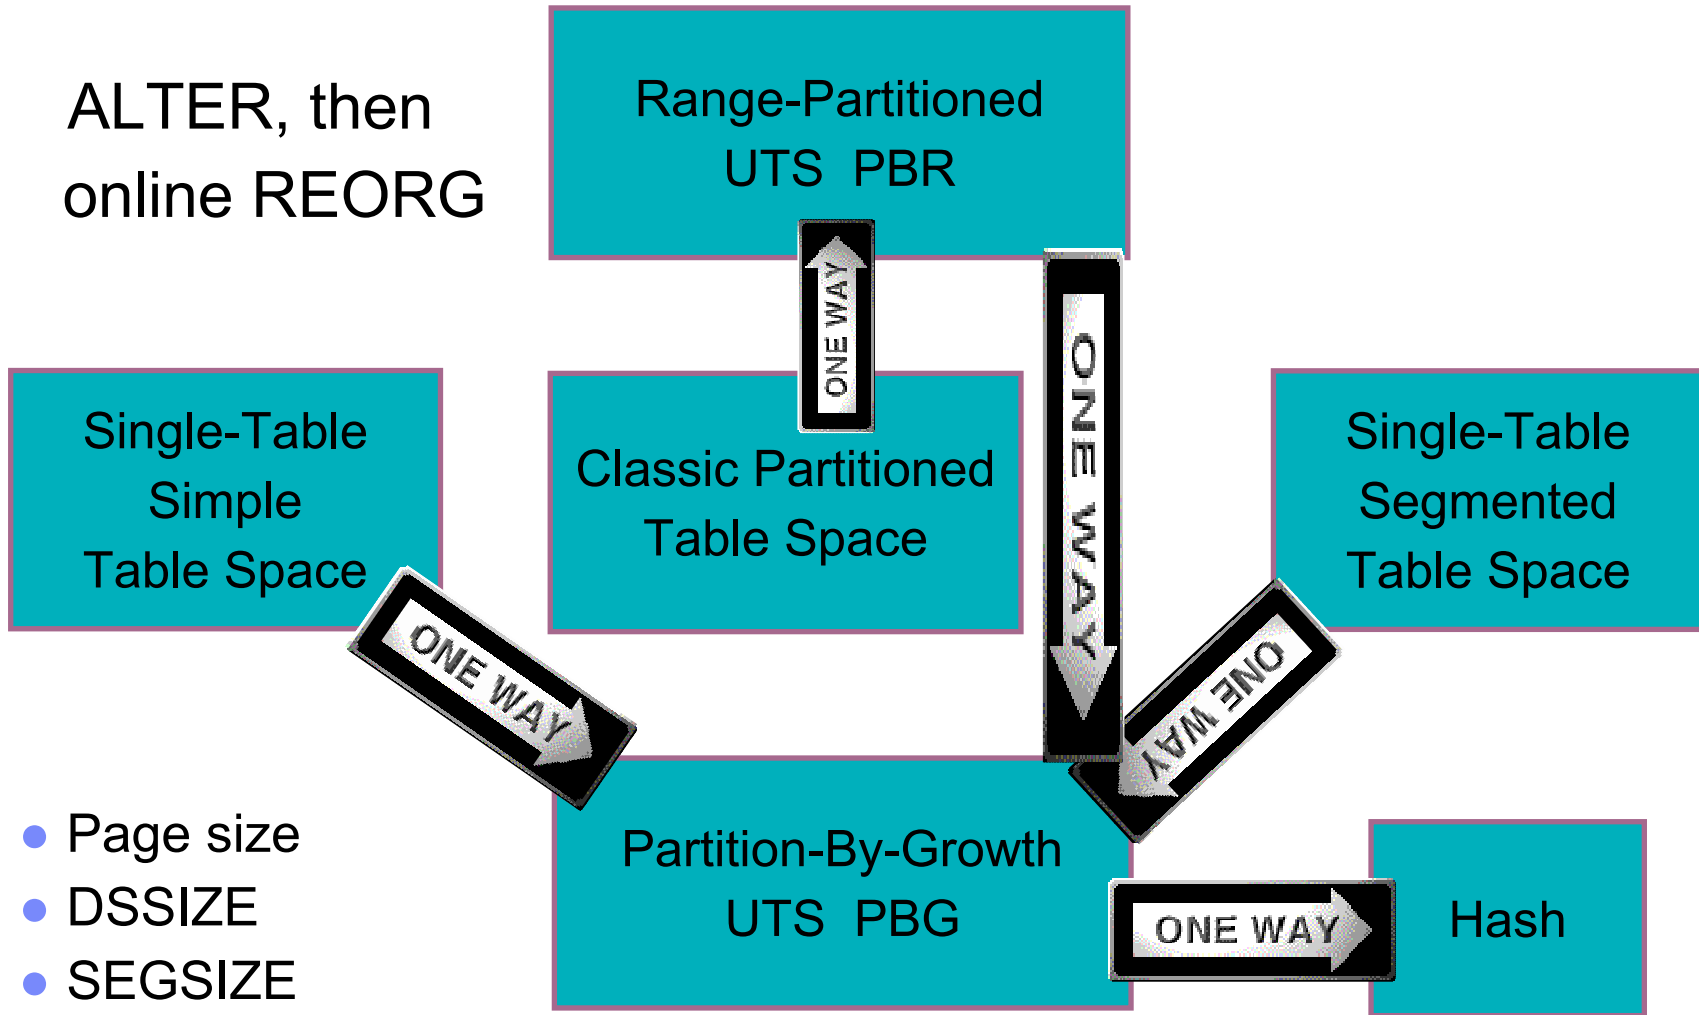
- More threads per DB2 image
- More efficient use of large n-ways
- Easier growth, lower costs, easier management
- Data sharing and Parallel Sysplex still required for HA and XXL scale

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” support is enabled
 - ▶ relieve 64GB limit for SPT01

DB2 X Availability: ALTER table spaces

ALTER, then
online REORG



- Page size
- DSSIZE
- SEGSIZE
- MEMBER CLUSTER

Availability

- More online schema changes for tablespaces, tables and indexes
 - Online REORG instead of DROP/CREATE or REBUILD INDEX
 - Alterations are manifested with REORG, unless noted otherwise
 - ▶ Page size for table spaces and indexes
 - ▶ DSSIZE for table spaces
 - ▶ SEGSIZE
 - ▶ MEMBER CLUSTER
 - ▶ Convert single table segmented into UTS PBG
 - ▶ Convert single table simple into UTS PBG
 - ▶ Convert classic partitioned tablespace into UTS PBR
 - ▶ Convert UTS PBR to UTS PBG
 - ▶ Convert PBG to hash (immediate, but RBDP index)
 - ▶ Ability to drop pending changes
- Online REORG for LOBs, other Online REORG / utility improvements
- Online add active log

DB2 X: Business Security & Compliance

- Protect sensitive data from privileged users
 - ▶ SYSADM without data access, DBADM for all DB
- Separate authority to perform security related tasks
- Allow EXPLAIN without execute privilege or ability to access data
- Audit privileged users
- More granular administrative authorities
- Row and column access control
 - Allow masking of value
 - Restrict user access to individual cells



Use disk encryption

DB2 X: Productivity – Doing More with Less!

- Auto statistics collection
- Compress 'on the fly'
 - ▶ Avoid need to run utility
- Reduce contention, more online processing
- Automatic config of IBM supplied UDFs and SPs
- Enhancements for Data Studio

Name	Availability Status	Cache Access Status	Crash	Swapping	CPU Usage	Cap. Status	Memory Usage	Loading	SQL Performance	Compression	Transactions	Logging	Standards
Production	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Adm	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Commerc	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Support	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Adm2	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
New York	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Los Angeles	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Assurance	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Marketing	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Test	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK
Development	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	OK

Manual invocation of

- RUNSTATS
- COPY/BACKUP SYSTEM
- QUIESCE
- MODIFY RECOVERY
- REORG

DB2 X Utilities Enhancements

- ▶ REORG SHRLEVEL(CHANGE) for LOBs
- ▶ Online REORG enhancements
 - SHRLEVEL(CHANGE) support for all catalog/directory objects
 - Support shrinking of PBG partitions
 - Option to cancel blocking threads
 - Faster SWITCH phase
 - Allow disjoint partition ranges
 - Permit movement of rows between partitions when LOB columns exist
 - Allows REBALANCE or shrinking of PBG even though LOB columns exist
 - Allows DISCARD to delete associated LOB values
 - Messages to estimate length of REORG phases and time to completion

Application Enablement, Portability

- Versioned data or Temporal
- pureXML enhancements
- Allow non-NULL default values for inline LOBs
- Loading and unloading tables with LOBs
 - LOBs in input/output files with other non-LOB data
- Last committed' locking semantics
- Implicit casting
- Timestamp with time zone
- Greater timestamp precision

Versioned data or Temporal Data

- Table-level specification to control the management of data 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, app logic greatly simplified
- New syntax in FROM clause to specify a time criteria for selecting historical data

Application Enablement, Portability ...

- SQL stored procedure enhancements
 - ▶ SQL PL in Scalar UDFs
- 64-bit ODBC
- Prepare/Describe enhancements for optimistic locking
- Special null indicator to indicate value not supplied or default
- DRDA support of Unicode for system code points
- Instance based statement hints
- Allow caching of dynamic SQL statements with literals
- Data-dependent paging
 - When only a specific part of the result set is needed
 - Efficient access to desired portions of result set, based on size or on data values

DB2 SQL

z z/OS X
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**

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 Enhancements

- XML schema validation in the engine for improved usability, performance
- Binary XML exchange format for improved performance
- XML multi-versioning for more robust XML queries
- Allow easy update of sub-parts of an XML document
- Stored proc, UDF, Trigger enhanced support for XML
- Composite index support for XML data

Data Warehousing

- Moving Sum, Moving Average
- Enhanced query parallelism technology for improved performance
 - ▶ Remove query parallelism restrictions
- In-memory techniques for faster query performance
- Advanced query acceleration techniques

Key details about DB2 X

- CM, ENFM, NFM modes
- Prerequisites
 - ▶ z/OS V1.10
 - ▶ DB2 9 for z/OS in NFM
 - ▶ System z10, z9, z890, z990, and above (no z800, z900)
- Items deprecated in earlier versions eliminated:
 - ▶ Private protocol → DRDA (new help in DSNTDP2DP)
 - ▶ Old plans and packages V5 or before → REBIND
 - ▶ Plans containing DBRMs → packages PK62876
 - ▶ ACQUIRE(ALLOCATE) → ACQUIRE(USE)
 - ▶ Old Plan table formats → DB2 V8 or 9 format (59 columns)
 - ▶ XML Extender → XML type
 - ▶ DB2 MQ XML user-defined functions and stored procedures → XML functions
 - ▶ DB2 Management Clients feature (DB2 Administration Server, Control Center, & Development Center) → IBM Data Studio application & administration services
 - ▶ msys for Setup DB2 Customization Center → install panels
 - ▶ BookManager use for DB2 publications → Info Center, pdf


DB2 X for z/OS At a Glance

<p>Application Enablement</p>	<ul style="list-style-type: none"> • Versioned data or Temporal • pureXML enhancements • Last Committed reads • SQL improvements that simplify porting
<p>RAS, Performance, Scalability, Security</p>	<ul style="list-style-type: none"> • Wide range of performance improvements • Hash access to data • More online schema changes • Catalog restructure for improved concurrency • Row and column access control • Administrator privileges with finer granularity
<p>Simplification, Reduced TCO</p>	<ul style="list-style-type: none"> • 5 – 10 times more threads per DB2 image • Auto statistics • Data compression on the fly • Query stability enhancements • Reduced need for REORG • Utilities enhancements
<p>Dynamic Warehousing</p>	<ul style="list-style-type: none"> • Moving sum, moving average • Many query optimization improvements • Query parallelism improvements • Advanced query acceleration

Get updated books **latest June 26, 2009**

<http://publib.boulder.ibm.com/infocenter/imzic/>

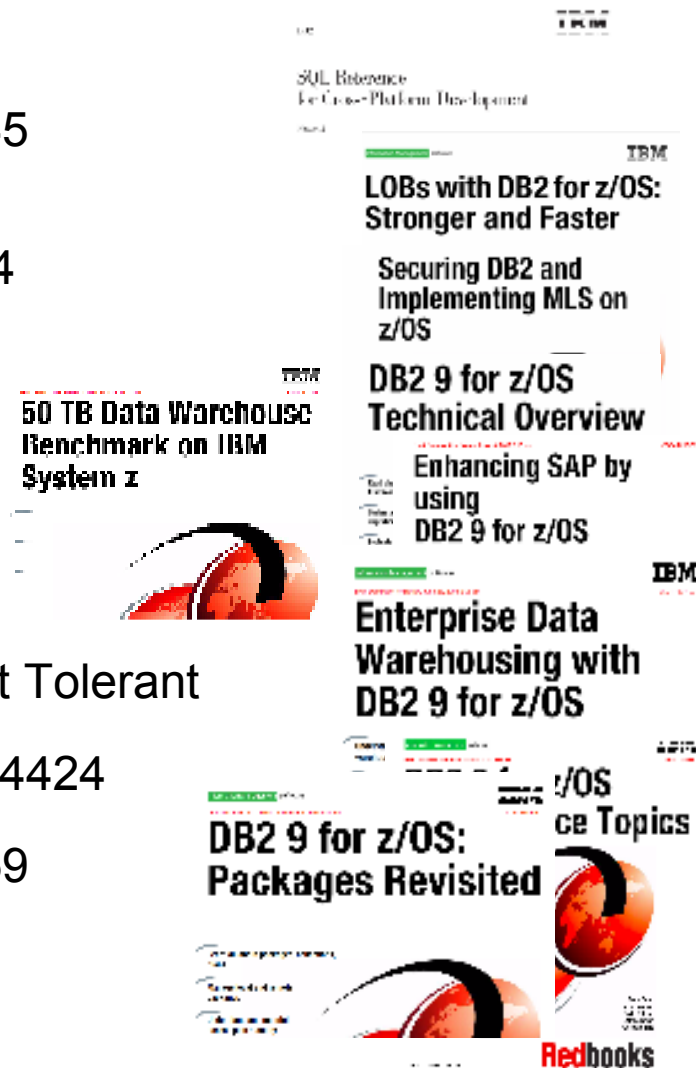
- ✓ Administration Guide
- ✓ Data Sharing: Planning and Administration
- ✓ Performance Guide
- ✓ Utility Guide and Reference
- ✓ Application Programming Guide & Reference
- ✓ Application Programming Guide & Reference for JAVA™
- ✓ ODBC, Spatial, ...
- ✓ Redbooks
- ✓ Installation Guide
- ✓ RACF Access Control Module Guide
- ✓ Messages
- ✓ Codes



...

DB2 9 in IBM Redbooks Publications

1. DB2 9 Technical Overview SG24-7330
2. DB2 9 Performance Topics SG24-7473
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. DB2 9 Optimization Service Center 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 Architecture, SG24-6952



Important Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.

WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- **CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR**
- **ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.**

DB2 9 for z/OS and Beyond

Thank
You

