



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 10 for z/OS

How can it help you?





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**



Software > Information Management > DB2 Product Family > DB2 for z/OS >

Announcing DB2 10 for z/OS Beta

The undisputed leader in total system availability, scalability, security and reliability

- DB2 for z/OS Family
- Library
- Technical resources
- Success stories
- News
- How to buy
- Events
- Training and certification
- Services
- Support



DB2 10 for z/OS Highlights

DB2 10 for z/OS Beta
Cut Costs & Improve Performance

- [Announcement Letter](#)
- [DB2 for z/OS and Data Warehousing Nordic Seminars](#)
- [Hear what people are saying about DB2](#)

DB2 for z/OS Roadshow - including DB2 10 for z/OS Sneak Preview



Coming to a city near you!

Feb 9 St. Louis, Feb 11 NYC, Feb 18 Toronto, Mar 4 Chicago and many more

→ [Register Now - Track 2](#)

DB2 10 for z/OS TECHNICAL PREVIEW

Related links

- [DB2 for z/OS](#)
- [DB2 9 for z/OS](#)
- [Data Warehousing and Business Intelligence on System z](#)
- [DB2 Tools for z/OS](#)
- [DB2 pureXML](#)

In today's business and economic environment, the challenge for IT is clear: improve operational efficiencies, reduce costs, and adapt quickly to support business growth -- all without sacrificing the resiliency required for today's demanding business requirements. DB2 for z/OS is the undisputed leader in total system availability, scalability, security, and reliability at the lowest cost per transaction. DB2 10 builds on the formidable capabilities of [DB2 9 for z/OS](#) and continues to set the standard, delivering key innovations and resource savings, including:

Out-of-the-box Savings by improving operational efficiencies

IBM continues to invest in new features to support your efforts to make your business more efficient, and DB2 10 delivers great value in this area. Compared to previous DB2 versions, some customers can achieve a 5% to 10% out-of-the-box CPU savings for traditional workloads and up to 20% out-of-the-box CPU savings for non-traditional workloads. Productivity improvements in DB2 10 for database and system administrators can drive additional operational efficiencies and cost savings. Synergy with other IBM System z platform components reduces CPU use by leveraging the latest processor improvements, larger amounts of memory, solid-state disk and z/OS enhancements.



DB2 10 for z/OS What's exciting?

Efficiency	<ul style="list-style-type: none">• CPU reduced: transactions & queries• Ten times more concurrent users
Resiliency	<ul style="list-style-type: none">• More online schema changes• Concurrency for catalog & utilities• Improved security controls and audit
Applications	<ul style="list-style-type: none">• Versioned data or temporal queries• pureXML and SQL enhancements

**→ Productivity improved for DBAs,
application programmers, & systems**



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, ... Insert your favorite.
10. Productivity improved for database & systems administrators, and application programmers





Why Migrate to DB2 10 for z/OS?

- Business needs to save money
 - **Reduce CPU time**
 - Service Oriented Architecture
- Application developers need improved productivity and integration
 - pureXML for a faster, more capable interface to XML data
 - Powerful new SQL temporal enhancements & portability
- Database Administrators need
 - Improved performance
 - Availability, scalability & memory management
 - Simpler security and regulatory compliance
 - More productive database administration






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 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





DB2 10: Performance Plan → Very Significant CPU Reductions: Best on z10



Transactions	DB design changes
Batch	Hash access, LOBs
REBIND	Index include cols.
	application changes
	SQL adjustments

Your situation will vary. Less CPU is better.

Processors z10, z9, z890, z990 and later z/OS 1.10 and later



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



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
 - “Fully in memory” option (ALTER BUFFERPOOL)



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
- Workfile spanned records, PBG
- 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
- DEFINE NO for LOB and XML columns



Converting to Hash Access

Evaluate potential candidate tables:

*Check if they have **unique keys***

used for single fetch queries

*Tables that are fairly **static** in size*

Evaluate existing index that can be replaced with the hash access

If the index is used for range queries, it cannot be dropped

Estimate number of rows, average row size and space needed

ALTER ADD *organization-clause*

REORG AUTOESTSPACE YES

Rebind applications with fully qualified equal predicates on hash key

To pick up Hash access.

Followup:

Check to see if hash access was chosen

Check that RTS for appropriate space specifications

Monitor index last-used RTS info to see if index can be dropped



DB2 10 for z/OS Resiliency

Scalability	<ul style="list-style-type: none">• Ten times more concurrent users• More concurrent utilities
Availability	<ul style="list-style-type: none">• More online schema changes• Improved concurrency on user & catalog data
Audit and Security	<ul style="list-style-type: none">• Improved audit and control• Increase administrative authority granularity• Masking for data in database




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 ½% 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 inserters
 - 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



Major changes in DB2 10 catalog & directory

- Improve availability and productivity
- Increase maximum size substantially
- Reduce contention: BIND, DDL, utilities
- Allow SELECT from SYSLGRNX
- Catalog changes: Remove links
 - Many more table spaces, partition by growth
 - Row level locking
 - CLOB and BLOB columns for long strings
 - Online reorganization and check
 - More automatic: DB2-managed SMS-controlled



Catalog Restructure for improved concurrency

- Remove links from the catalog and directory
 - DSN1CHKR no longer needed
- Sixty new table spaces; 7 old ones dropped
- Row level locking enabled for catalog and directory
- Conversion during ENFM for migrated systems
- BIND, PREPARE, and DDL will run with better concurrency, fewer timeouts/deadlocks
- Allow online REORG for all catalog and directory table spaces.



Other Catalog Changes

- Partition-by-growth (PBG) catalog/directory table spaces
 - Allow packages to grow beyond 64GB using LOBs
 - V8/DB2 9 APAR PK80375 adds zparm for SPT01 compression; no compression in DB2 10
- DB2 managed catalog and directory data sets
 - SMS-managed, DB2-managed catalog required
 - Eases admin and management burden
- New CLOB & BLOB columns for storing SQL statements
 - Today, SQL statements can be split into several records with sequence numbers
 - CLOBs will make it easier to query SQL statements
- Online REORG enabled for all catalog/directory objects



Improved availability ALTER REORG

Pending ALTER,
then online REORG
to make changes

Range-Partitioned
UTS PBR

LOB INLINE LENGTH, default
VERSIONING
ACCESS CONTROL
MASK, PERMISSION
TRIGGER SECURED
FUNCTION SECURED
TIMESTAMP precision, time zone
MAXPARTITIONS

Single-Table
Simple
Table Space

Classic Partitioned
Table Space

Single-Table
Segmented
Table Space

Page size
DSSIZE
SEGSIZE
MEMBER CLUSTER

Partition-By-Growth
UTS PBG

Hash

INDEX page size
INCLUDE cols

ADD active log
BUFFERPOOL PGSTEAL NONE





Availability

- Online schema changes for table spaces, tables and indexes – PENDING with ALTER and Online REORG instead of DROP/CREATE or REBUILD INDEX

Alterations occur with REORG, unless noted otherwise

- Page size for table spaces and indexes BUFFERPOOL
- DSSIZE for table spaces
- SEGSIZE
- Convert single table segmented into UTS PBG
- Convert single table simple into UTS PBG
- Convert classic partitioned table space into UTS PBR
- Convert UTS PBR to UTS PBG
- Convert PBG to hash (immediate, but RBDP index)
- MEMBER CLUSTER
- Ability to drop pending changes



Availability ...

- More ALTERs (not pending)
 - LOB INLINE LENGTH, default
 - VERSIONING
 - ACCESS CONTROL
 - MASK, PERMISSION
 - TRIGGER SECURED
 - FUNCTION SECURED
 - TIMESTAMP precision, time zone
 - INDEX INCLUDE cols
 - BUFFERPOOL PGSTEAL NONE
 - MAXPARTITIONS
- Online REORG for LOBs, other Online REORG / utility improvements & restriction removal
- Online add active log



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



DB2 10 Security Benefits

- More flexible authorization
- Separation of duties
- Do job without access to data
- Policies for audit
- Simpler control
- Tighter security
- Avoid cascade delete
- Avoid views and application security logic
- Allow more tools
- Evolve security policies
- Easier to manage security policy
- Improved productivity & tighter security





DB2 10: Productivity – Doing More with Less!

- Auto statistics collection
- Easier scaling, simpler memory management
- Reduce contention, more online processing
- Access path stability
- Reduced need for REORG
 - Build compression dictionary on the fly
 - Index list prefetch enhancements
- Configure IBM UDFs and stored procedures
- Allow one SDSNEXIT data set for many subsystems
- Monitoring enhanced
 - Timeout / deadlock diagnostics
 - Identify SQL statements

Name	Monitoring Status		Alert		System		Database					
	Data Server Status	Critical	Warning	CPU Usage	Disk Space	Memory Usage	Locking	SQL Performance	Connections	Transactions	Logging	Maintenance
Production	3	8										
Web	1	1										
eCommerce	0	0										
Support	1	1										
Retail	0	0										
New York	0	0										
Los Angeles	0	0										
Accounts	2	3										
Marketing	0	4										
Test	0	0										
Development	0	0										

Manual invocation of

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



DB2 10 Utilities Enhancements

- REORG SHRLEVEL(CHANGE) for LOBs
- Online REORG enhancements
 - SHRLEVEL(CHANGE) support for all catalog/directory objects
 - Option to cancel blocking threads
 - Improved availability
 - Allow disjoint partition ranges
 - Permit movement of rows between partitions when LOB columns exist
 - Allows REBALANCE and ALTER LIMITKEY even when LOB columns exist
 - Allows DISCARD to delete associated LOB values
- Messages to estimate length of REORG phases and time to completion



DB2 10 more utilities enhancements

- Support of spanned records for UNLOAD of LOB data
 - Currently unload of LOBs >32K must use FRVs
 - Allow inline LOBs with base row in unload data set
 - Provides portability of data
 - Performance enhancement for FRV processing with PDS data sets, also in DB2 9
 - UNLOAD 33% elapsed time reduction
 - LOAD 84% elapsed time reduction
- Autonomic RUNSTATS & table profile



DB2 10: More Utility Improvements

- Improved COPY CHANGELIMIT performance
 - Use RTS instead of SM page scans
- Data set level FlashCopy option
- FlashCopy backups with consistency and no application outage
- FlashCopy backups as input to:
 - RECOVER (fast restore phase)
 - COPYTOCOPY, DSN1COPY
- RECOVER “back to” log point
- REPORT RECOVERY support for system level backups



Autonomics and DBA Productivity...

- Checkpoint intervals based on both time and log records
- Run 'must complete' backout under pre-emptable SRB
- Identify unused packages
- SQL Statement level monitoring
 - Statement ID introduced
 - Trace records & messages extended to include statement ID
 - New trace class for statement detail
 - GetPages, Locks, I/Os, cpu/elapsed time, etc. at statement level
- Manage max threads, connections, idle thread timeout on an application basis
 - Warning or exceptions issued when threshold is hit



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

➤ Access path repository

➤ Versioning

➤ “Fallback”

➤ “Lockdown”

➤ Manual overrides. Hints: easily influence access paths without changing apps

➤ Per-statement BIND options

- Safe query optimization: assess “reliability” of access path choices
- RID pool overflow to workfiles



Many improvements for SAP & web applications

- **Autonomics**
- Compress on the fly on INSERT
- Auto-statistics
- Access path stability and hints enhancements
- Access path lock-in and fallback for dynamic SQL
- Automatic checkpoint interval
- Automated installation, configuration & activation of DB2 supplied stored procedures & UDFs
- Data set FlashCopy in COPY & inline copy
- Inline image copies for COPY YES indexes
- UNLOAD from FlashCopy backup
- REORG enhancements
- Reduce need for reorganizations for indices
- **Performance**
- CPU reductions
- Hash access path
- Numerous optimizer enhancements, paging through result sets
- Parallel index update at insert
- Faster single row retrievals
- Inline LOBs
- LOB streaming between DDF and rest of DB2
- Faster fetch and insert, lower virtual storage consumption
- DEFINE NO for LOBs and XML
- MEMBER CLUSTER for UTS
- Query parallelism enhancements: lifting restrictions
- Dynamic Index ANDing Enhancements
- Option to avoid index entry creation for NULL value
- Index include columns
- Buffer pool enhancements
- **Scalability**
- Many more threads
- Reducing latch contention
- Workfile spanned records, PBG support, and in-memory enhancements
- **Availability**
- More online schema changes for table spaces, tables and indexes via online REORG
- Online REORG for LOBs
- Online add log
- Automatically delete CF structures before/during first DB2 restart
- **Portability**
- Allow non-NULL default values for inline LOBs
- Loading and unloading tables with LOBs in stream
- Currently committed locking semantics
- Default SAP settings for DB2
- **Security**
- More granular DBA privileges



DB2 10 Application Enablement and Portability

- Data versioning by date
- pureXML enhancements
- Large object improvements
 - Allow non-NULL default values for 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 – seconds to picoseconds
 - Moving Sum, Moving Average



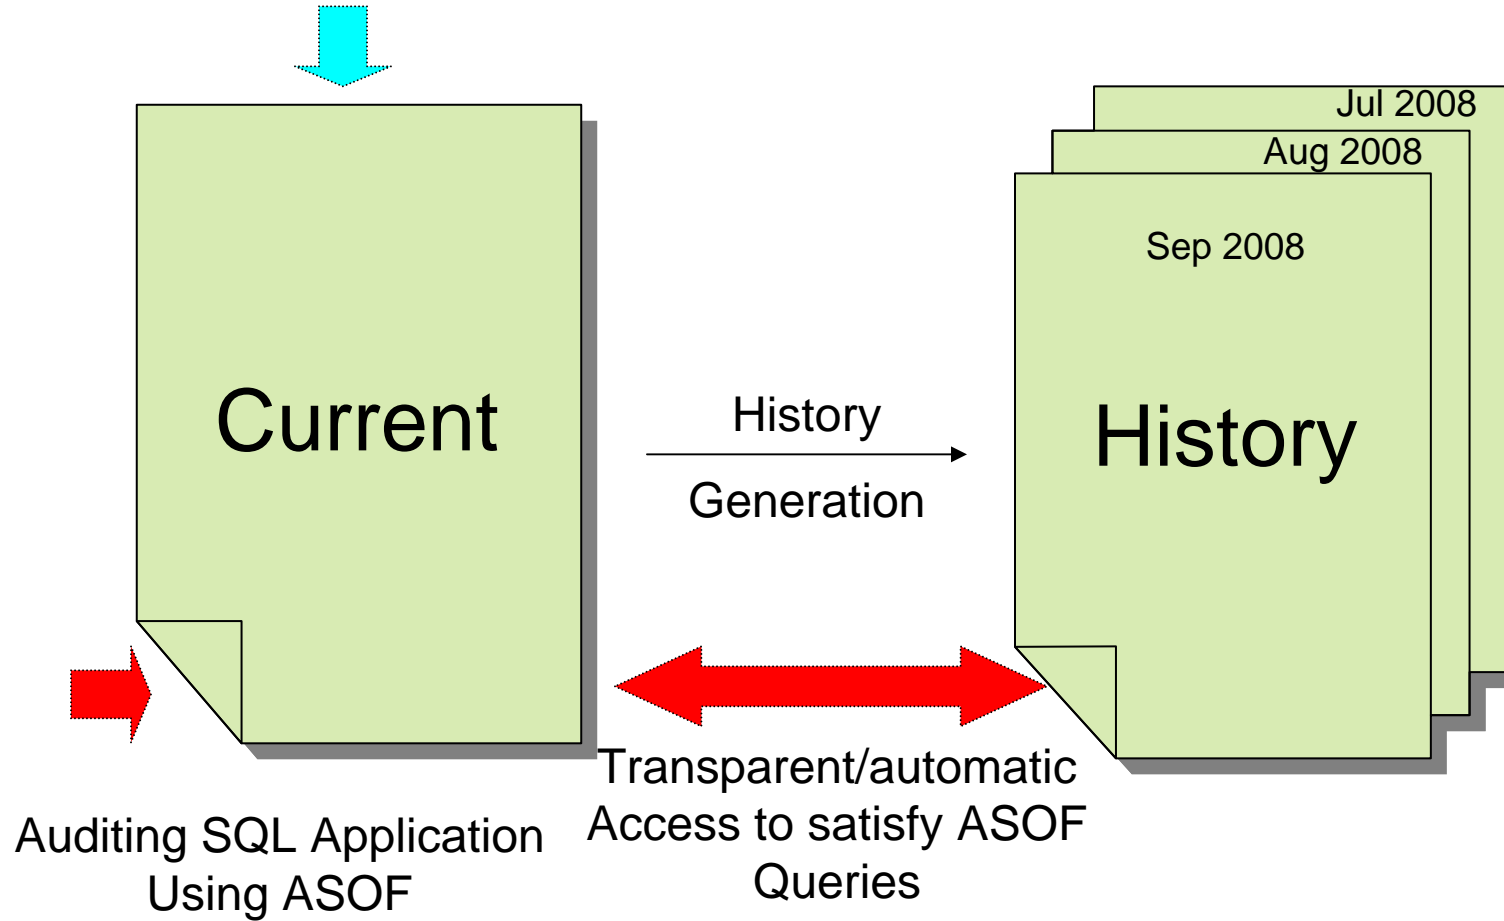
Versioned data or Temporal Data

- 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



Current and History

Current SQL Application





DB2 10 Application Enablement, Portability ...

- SQL stored procedure enhancements
 - SQL PL in Scalar UDFs & XML support
- 64-bit ODBC – also DB2 9 PK83072
- 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 upon current position



DB2 SQL

z z/OS V7

common

luw Linux, Unix & Windows V8.2



- z { Range partitioning
- c { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions, 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, Call from trigger, statement isolation
- o {
- m {
- n { 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 INSERT, UPDATE, or DELETE, multi-site join, 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, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT
- l {
- u {
- w {



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

commonluw

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

- XML schema validation in the engine for improved usability and performance
- Binary XML exchange format improves performance
- XML multi-versioning for more robust XML queries
- Allow easy update of sub-parts of XML document
- Stored proc, UDF, Trigger enhanced support
- XML index matching with date/timestamp
- CHECK DATA utility checks XML



DB2 10 Query Enhancements

- CPU time reductions for queries, batch, & transactions
- SQL enhancements: Moving Sum, Moving Average, temporal, timestamp, implicit cast, SQL PL, ...
- pureXML improvements
- Access improvements: Index include columns, Hash
- 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
- Analysis: instrumentation, Data Studio & Optim Query Tuner
- Advanced query acceleration techniques
 - IBM Smart Analytics Optimizer



Technology Preview: IBM Smart Analytics Optimizer

What is it?

- A special purpose, network-attached appliance that is an add-on to a DB2 for z/OS system
- Offloads typical DW/BI queries resulting in predictable and orders-of-magnitude faster query response times while reducing overall TCO



Business Value

- Dramatically lowers the cost for query and reporting on System z
- Advanced in-memory scale-out cluster technologies that keep the complete system centrally managed without having to change any requirements for BI applications
- Complements the many new Data Warehousing features in DB2 9 for z/OS
- Leverages the many new warehousing and business intelligence solutions now available on System z

Targeted Uses for DB2 for z/OS customers:

- Requirements to accelerate a subset of their warehouse or reporting queries
- Looking for more insight and business intelligence from operational data
- Needs to consolidate datamarts or data stores into one enterprise warehouse



DB2 10 for z/OS At a Glance

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



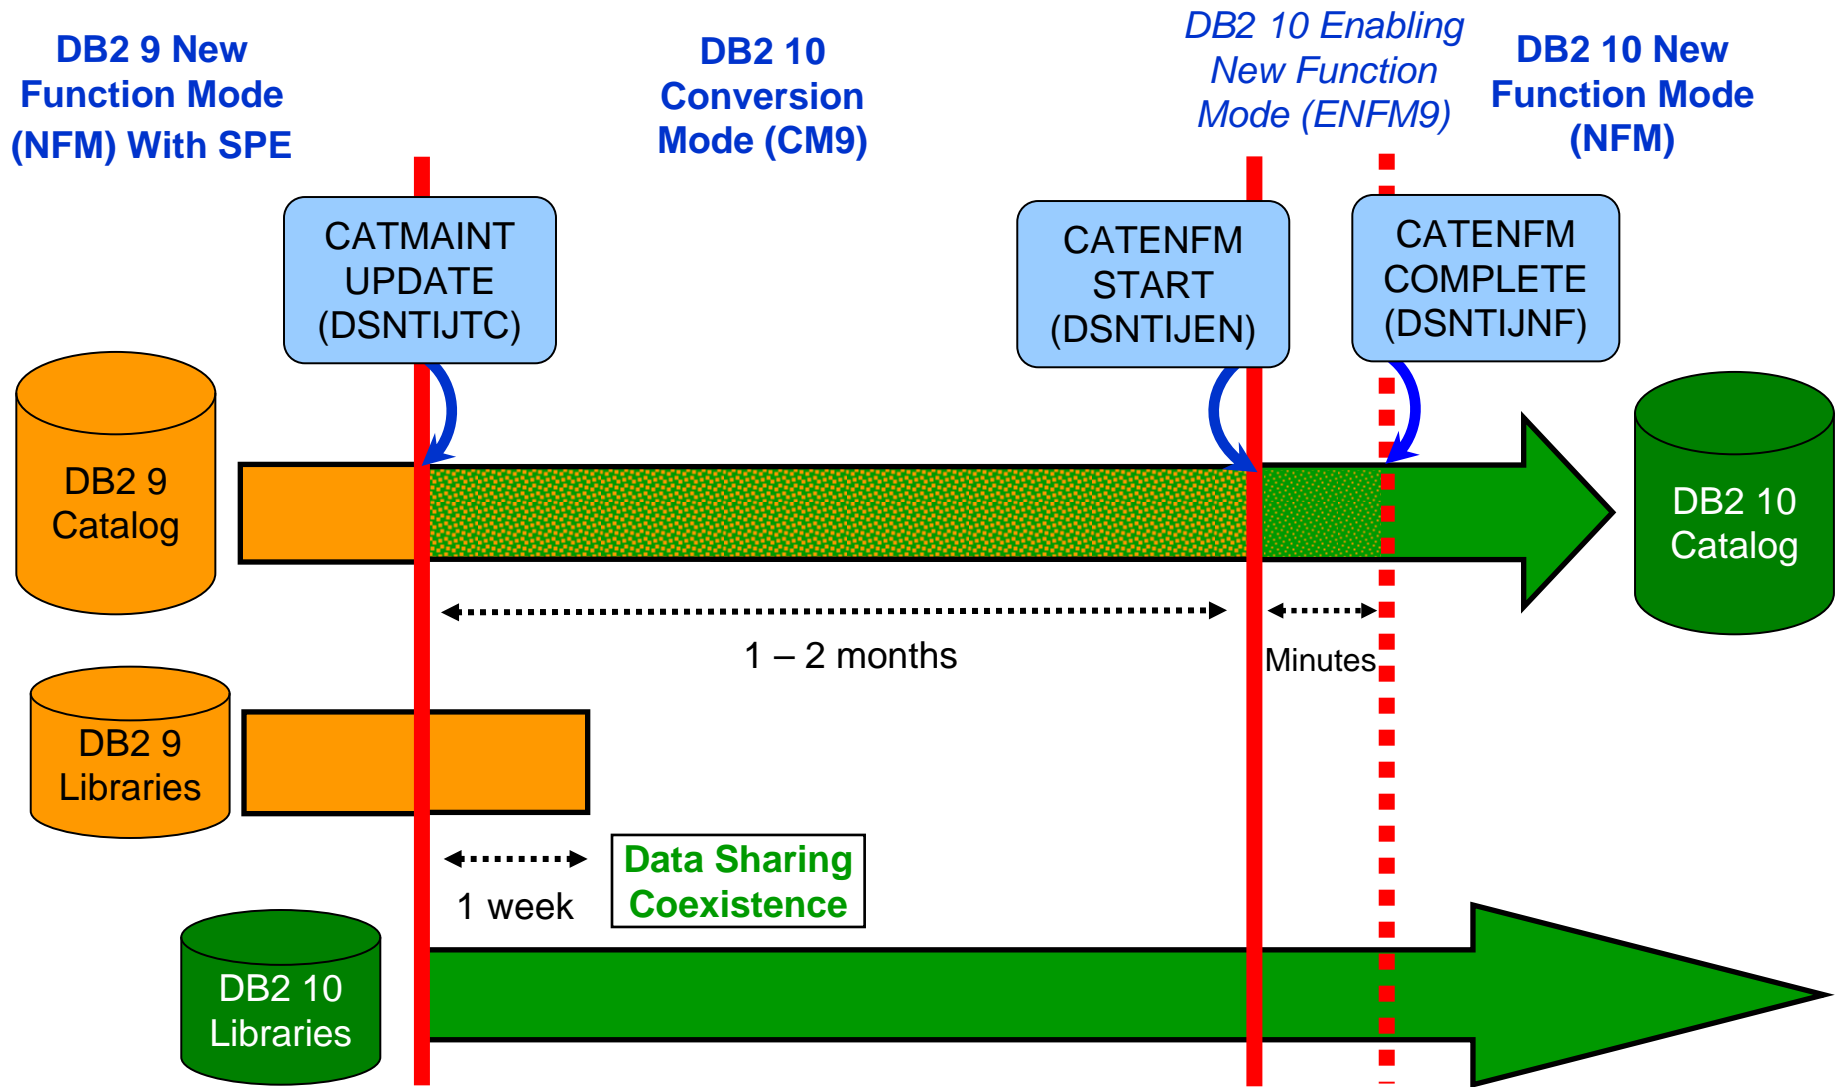
If you are on V7 or earlier, go to V8

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



Normal Migration Overview DB2 9 → DB2 10





Overview of Modes when migrating 9 → 10

CM9 Conversion Mode – The mode DB2 is in when DB2 10 is started for the first time after migrating direct from DB2 9. It will still be in CM9 when migration job DSNTIJTC has completed. Very little new function can be executed in CM9 Data sharing systems can have DB2 9 and DB2 10 members in this mode. DB2 can only migrate to CM9 from DB2 9 NFM.

ENFM9 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 DB2 10 members in this mode.

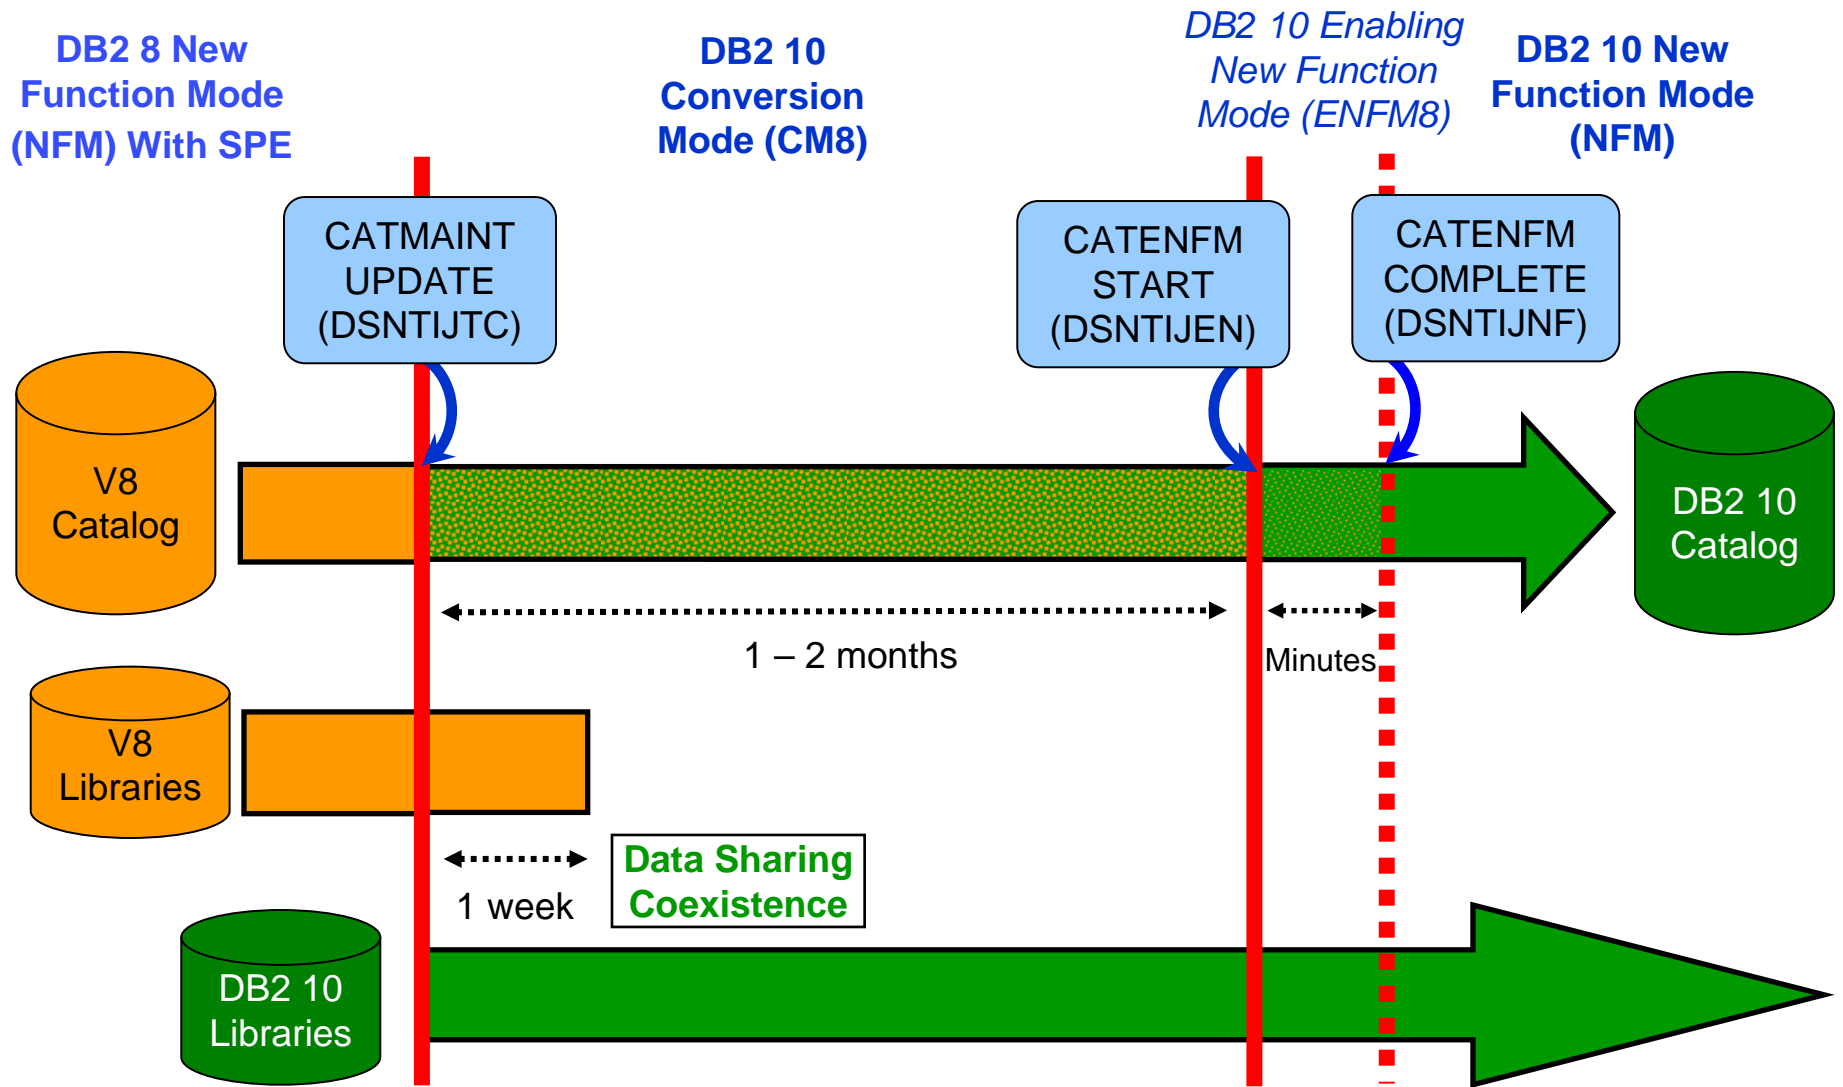
NFM New Function Mode - This mode is entered when CATENFM COMPLETE is executed (the only step of job DSNTIJNF). This mode indicates that all catalog changes are complete and new function can be used.

ENFM9* This is the same as ENFM9 but the * indicates that at one time DB2 was at DB2 10 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 ENFM9* it can not fallback to DB2 9 or coexist with a DB2 9 system.

CM9* This is the same as CM9 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 CM9* it can not fallback to DB2 9 or coexist with a DB2 9 system.



Skip Migration Overview V8 → 10





Overview of Modes when migrating V8 → 10

CM8 Conversion Mode - This is the mode DB2 is in when DB2 10 is started for the first time after migrating direct from DB2 V8. It will still be in CM8 when migration job DSNTIJTC has completed. Very little new function can be executed in CM8. Data sharing systems can have DB2 V8 and DB2 10 members in this mode. DB2 can only migrate to CM8 from DB2 V8 NFM.

ENFM8 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 DB2 10 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.

ENFM8* This is the same as ENFM8 but the * indicates that at one time DB2 was at DB2 10 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 ENFM8* it can not fallback to DB2 V8 or coexist with a DB2 V8 system.

CM8* This is the same as CM8 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 CM8* it can not fallback to DB2 V8 or coexist with a DB2 V8 system.



Customer migration paths:



Skip DB2 9 is possible, but most customers will go to DB2 9

Estimate save 20% to 25% vs 2 steps

Reasonable timing to skip


Value in DB2 9 deferred until then

When will applications add support for 10? Drop 8?

Skip versions V5 → 7, V8 → 10, not every version



DB2 for z/OS time line



V7

V8

V7 EOS

8 → 10 Skip zone 2012 →

No new value zone
If skipping 8 → 10
Application support timing?




If you are not on DB2 9, look what you are missing.

<i>Availability, Scale & Resiliency</i>	More Online Schema Changes		Business Flexibility
	Volume Level Backup & Recovery		Faster, Cheaper, Granular Recovery
<i>Compliance & Security</i>	Database Roles & Trusted Context		Efficient Auditing & Compliance
<i>Easier Application Development</i>	PureXML		Streamlined Data Integration
<i>OLTP & Warehousing</i>	Query Optimization		Highly Available, Secure Data
<i>Reduce Cost of Ownership & Skills</i>	Index Compression		50% Disk Savings



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 from DB2 9:

- 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.
- DB2 Management Clients feature, Control Center
- msys for Setup DB2 Customization Center
- Some subsystem parameters removed, many changed
- XDBDECOMPXML and XDBDECOMPXML100MB
- Accessories Suite parts Optimization Service Center, Developer Workbench, Visual Explain for DB2 → Data Studio



No longer supported in DB2 10 from V8:

- Private protocol: Convert to DRDA.
- EXPLAIN tables prior to Version 8 format: Alter add columns.
- 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.
- DB2 Management Clients feature, Control Center
- msys for Setup DB2 Customization Center
- Some subsystem parameters removed, many changed
- ...



No longer supported in DB2 10 from V8: ...

- Net.Data
- DB2-established stored procedures
- Old JDBC driver
- Pascal L string data type from VAX
- Creating simple table spaces
- QMF Visionary Studio
- DB2 Estimator
- BookManager help
- DB2 Extenders: AIV, text, Net Search
- Java stored procedures in resettable JVMs
- ...



Deprecated, may be removed from later versions:

- Simple table spaces
- Partitioned table spaces other than universal table spaces
- DSNHDECP load
- NEWFUN(YES) and NEWFUN(NO)
- DSNHPC7 precompiler
- EXPLAIN stored procedures
- ...

All of these lists are much longer, see book for details.



Questions?





DB2 for z/OS Technical Strategy

- Application enablement
 - Apps can easily connect to DB2 from anywhere
 - Advanced SQL, XML capability, application portability
- Extend the lead in availability, scalability and performance.
 - Parallel Sysplex: the best scale-out solution in the industry
 - Tight integration between DB2 and the System z hardware and z/OS operating system
 - Advanced solutions for compliance with data security and privacy regulations
 - Workload consolidation: System z is the ultimate consolidation platform
 - Eliminate all causes of outages
- Reduce cost of ownership
 - DB technology that can handle large workloads with fewer people
 - Advanced autonomics to make the system more self-managing and self-tuning
 - Storage and CPU optimization, including specialty engines
- Improved data warehousing capabilities



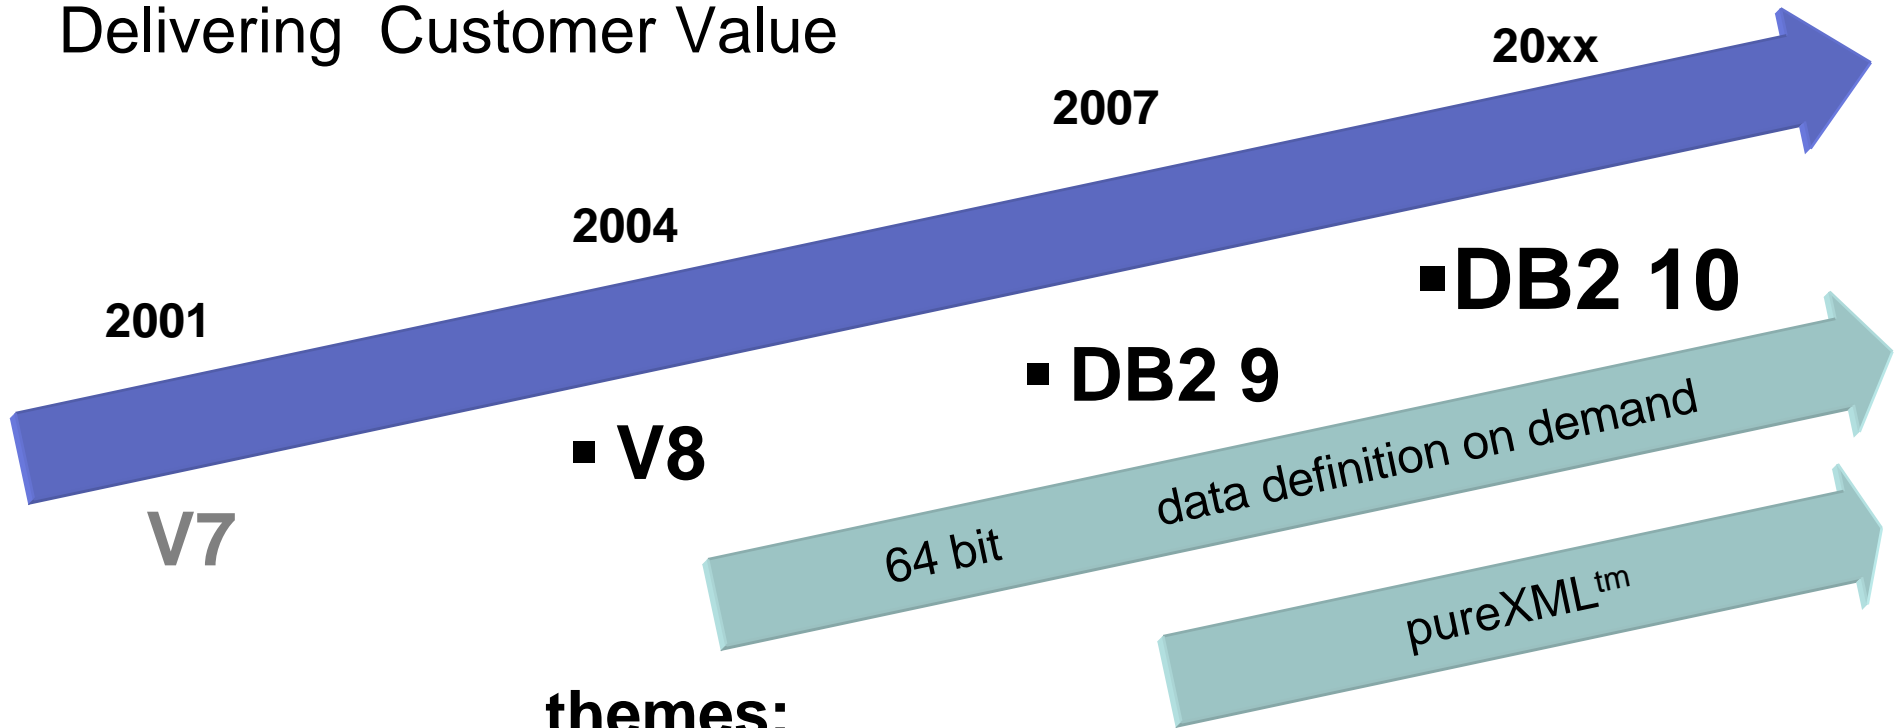
DB2 10 Major Focus Areas

- Performance, reduced CPU consumption
- Scalability, more concurrent active threads
- Continuous availability
- Contention reduction
- DBA / system administration productivity
- Advanced application functionality
- Security advancements



DB2 for z/OS Into the Future

Delivering Customer Value



themes:

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



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

