

DB2 10 and Beyond – Celebrating 30 Years of Superior Technology

Jeff Josten Distinguished Engineer

DB2 for z/OS Development

IBM Silicon Valley Lab

© 2013 IBM Corporation

Please note

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.



Agenda

- 30 years a quick retrospective
- DB2 10 news and highlights
- DB2 and the new zEnterprise EC12 machine
- DB2 11 preview



A Peek at Some DB2 Museum Artifacts





V1R1 announced June 7, 1983 & GA April 2, 1985











DB2 for z/OS Delivering Customer Value 2010

64 bit

2004

V8

2007 DB2 10 DB2 9

pureXMLtm

2001

V7

Ongoing themes: Performance Scalability Reliability Availability Serviceability Security Productivity Application Development SQL XML

The DB2 Family

1983: IBM gives DATABASE 2 (DB2) for MVS (once known internally as Eagle) its wings, announcing V1

> **1988:** IBM announces and ships SQL/400 on the new AS/400 server with integrated RDBMS. The International DB2 Users Group (IDUG) is formed.



DB2。



1993: IBM announces and ships DB2 for OS/2 V1 and DB2 for AIX V1





DB2 10 for z/OS Snapshot

- Fastest uptake
 - +2x customers vs. V9
 - +2.5x licenses vs. V9
 - 25% coming from DB2 V8
- Adoption Driven by:
 - Performance improvements without application changes
 - Virtual Storage Constraint relief for more threads
 - Security, RAS improvements
 - Bitemporal data



DB2 10 Virtual Storage Constraint Relief DB2 10

DBM1 below 2GB

- 75-90% less usage in DB2 10 compared to DB2 9
- Some of working storage (stack, xproc storage) stays below 2GB

Larger number of threads

 Possible data sharing member consolidation

Improve CPU with storage

- More release deallocate
- Larger MAXKEEPD values for
- KEEPDYNAMIC=YES

10





DB2 10 Performance

- Most customers can see a 5% 10% out-of-the-box CPU reduction (transactions and batch) after rebind
- Some workloads and customer situations can see a CPU reduction of up to 20%
- Synergistic operation with latest System z hardware
- <u>Sample</u>: Preliminary Measurements of IBM Relational Warehouse Workload (IRWW) with Data Sharing
- Base: DB2 9 NFM REBIND with PLANMGMT EXTENDED
- DB2 9 NFM → DB2 10 CM without REBIND showed 1.3% CPU reduction
- DB2 10 CM REBIND with same access path showed 4.8% CPU reduction
- DB2 10 NFM brought 5.1% CPU reduction
- DB2 10 CM or NFM with RELEASE DEALLOCATE 12.6% CPU reduction from DB2 9



Average %CPU improvements version to version

© 2012 IBM Corporation



DB2 10 Productivity – Doing More with Less!

- Easier scaling, simpler memory management
- Reduce contention, more online processing
- Reduced need for REORG
 - Build compression dictionary on the fly
 - Index list prefetch enhancements
 - Row-level sequential detection
- Configure IBM UDFs and stored procedures
- Statement level monitoring
- Access path stability, APREUSE & APCOMPARE
- DDF thread management enhancements

					Alert			System			Database		
// Next Refresh 0:24	Monit	oring stat	server sta	ous al	ng pu'	Jsage Disk	5pace	or Usage	n9 01	erforman	ections	actions	9 Mainte
C Production	v.	~	-		-		N ⁻	~				V A	N ⁻
	-		3	0	-	•	•	<u> </u>	•		•	•	—
🔻 🗁 Web	+	•	1	1		+	+	Δ	+	•	+	+	+
eCommerce	•	•	0	0	•	-	•	•	•	•	•		
Support	•	٠	1	1		•	•	Δ	٠	٠	٠	٠	•
🔻 🗁 Retail		-	0	0	٠	\	\	\	\	\	\	٠	\
New York			0	0	<i>\</i>	<i>\</i>	<i>\</i>	<i>\</i>	<i>\</i>	<i>\</i>	<i>\</i>	<i>\</i>	\$
Los Angeles		- \$	0	0	<i>\</i>		<i>\</i>	<i>\</i>		<i>\</i>	<i>\</i>	- \$	\$
Accounts	•	•	2	3		•	•	Δ	•	٠	٠	٠	•
Marketing	+	٠	0	4	٠		•	Δ		٠	٠	٠	٠
🕨 🧀 Test	•	-	0	0	\	\	•	\	\	\	\	\	\
 Condense 			0	0									







Online Schema – DB2 10





Changes in DB2 10 catalog & directory

- Improve availability and productivity
- Increase maximum size substantially
- Reduce contention: BIND, Prepare, Utilities
 - DDL concurrency also improved from removal of DBD01 hash anchor locks
- Catalog changes: Remove links, hashes
 - Many more table spaces, partition by growth
 - Row level locking, reordered row format
 - CLOB and BLOB columns for long strings
 - Inline for performance
 - Online reorganization and check
 - More automatic: DB2-managed SMS-controlled
 - Allow query of SYSLGRNX
 - Allow SQL statements in catalog to be queried with normal SQL





DB2 10 Business Security & Compliance

- Protect sensitive data from privileged users and improve productivity
 - Separation of duties
 - More granular control
 - New SECADM authority for security administration
 - System level DB admin with/without access control and with/without data access privileges
 - Usability: DBADM for all DB
 - New Explain privilege
- Audit policy for improved data auditing
 - Audit admin authority usage
 - Dynamic table auditing
 - All access, statement id level
- Support distributed identities introduced in z/OS 1.11
- Support client certificate authentication in z/OS 1.10
- Support password phrases in z/OS 1.10
- Connection level security enforcement using strong authentication
- Row and column access control
 - Allow masking of value
 - Restrict user access to individual cells





Temporal Based Analysis

- System-maintained temporal tables
 - DB2 generated history
 - AS OF query
- User-maintained temporal tables
 - User provide time period
 - Automatic business time key enforcement.
 - Query over any current, any prior, future point/period in business time.
 - New time range update/delete statements support automatic row splitting, exploited by the merge statements.
- Bi-temporal, combination of the above two





Key details about DB2 10: getting ready

Prerequisites: migrate from DB2 9 NFM or DB2 V8 NFM

- z/OS V1.10 SMS-controlled DB2-managed DB2 catalog
- System z z196, z10, z9, z890, z990, and above (no z800, z900)
- DB2 Connect 9 FP1, 9.7 FP3a for 10 new function
- Premigration check DSNTIJPA PM04968
- Info APARs II14477 (DB2 9) II14474 (V8)

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

- Private protocol \rightarrow DRDA
- Old plans and packages V5 or before \rightarrow REBIND
- Plans containing DBRMs \rightarrow packages
- ACQUIRE(ALLÕCATE) \rightarrow ACQUIRE(USE)



Large Main Memory Exploitation

- Latest z servers can support TB main memory sizes
- A single DB2 member can currently support up to 1 TB of BP space
- >DB2 10 has improvements for large BP scalability/performance
 - ➤ 1MB page frame support
 - ➤ "In Memory" BP attribute
 - > Avoid BP scans for data sharing TS/part/index inter-DB2 interest level changes
 - Internal BP hashing and latching improvements
 - VSCR improvements for 5-10x more active threads per DB2 member/subsys
 - > Opportunities to tradeoff increased memory usage for reduced CPU
- ➤DB2 10 new RTS column to indicate SSD usage
- >New zEC12 machine introduces flash memory
- Future: continue to leverage large main memory and automatic storage hierarchy management for performance advantages



Hardware Trends Leveraged by DB2 for z/OS

• Multi core, future slowing growth in single thread performance

- Higher n-ways, more parallelism bring potential latching bottlenecks, memory cache thrashing, ...
- S/W techniques for single threaded performance growth
- Clustered systems for massive scale out and continuous availability
- Specialty engines (price / performance)
- Hybrid systems, accelerators
 - Use cores that are more specialized to their purpose
 - New performance opportunities

Memory hierarchy design

- Higher CPU frequencies, n-ways make cache utilization a critical factor
- Translation lookaside buffer design, large System z page sizes

Solid state disk (and other disk related improvements)

- Performance, energy consumption, reliability benefits of HDD
- Large and rapidly growing main memory sizes along with falling prices
 - Performance improvements through more memory use
 - DB2 10 enables more persistent threads with RELEASE(DEALLOCATE)
 - Larger buffer pools, in-memory workfiles, I/O avoidance

DB2 Deep Synergy with System z and DS8000

- Hardware-based Coupling Facility for efficient DB2 data sharing (availability and scale out)
- zIIP engines to offload certain DB2 workloads
- Hardware instructions for Unicode conversion
- Hardware data compression & encryption
- Cross-memory, memory protection keys
- Sorting
- z/OS Workload Manager (WLM)
- z/OS Security Server (RACF)
- I MB page size
- Volume-level and object-level FlashCopy
- Solid state disks (SSDs)
- System z High Performance FICON (zHPF) with multi-track data transfer and FICON Express 8





DB2 and zEnterprise EC12 Team to Bolster Capacity and Performance

- Faster CPU 1.25X compared to z196
 - 20-28% CPU reduction measured with DB2 OLTP workloads
 - 25% reduction measured with DB2 query and utilities workloads
 - Less compression overhead with DB2 data (1-15%)
- 50% More System Capacity
- New Features DB2 plans to exploit
 - FLASH memory and 1MB Pageable Large Pages
 - Larger DB2 Buffer Pools
 - 2GB frame support
 - Larger frames should provide additional CPU savings
 - DB2 code backed by large frames for CPU reductions
- Flash memory improves SVC dump and z/OS paging performance
- Initial support planned for DB2 10 with APARS
- Additional 1 M PLP being considered for DB2 11





DB2 transaction throughput improvement stems from reduced CPU needed for pageable buffer pool management blocks

Initial results: Pageable Large Pages for DB2 helps DB2 achieve up to a 3% transaction throughput improvement from CPU savings. The savings are due to reduced buffer pool management with 1M PLP.

STATE IDIVI COLPORATION





Creating the Hybrid Data Server - Netezza and Z

Combine DB2 for z/OS with Netezza to provide an industry exclusive



Best in OLTP and Transactional Analytics

Industry recognized leader in mission critical transaction systems

Best in Deep Analytics

Proven appliance leader in high speed analytic systems

Best in Consolidation

Unprecedented mixed workload flexibility and virtualization providing the most options for cost effective consolidation

ven appliance leader in



DB2 Analytics Accelerator V3 *Lowering the costs of trusted analytics*



What's New?

- High Performance Storage Saver
 - Store a DB2 table or partition of data solely on the Accelerator. Removes the requirement for the data to be replicated on both DB2 and the Accelerator

Incremental Update

• Enables tables within the Accelerator to be continually updated throughout the day.

zEnterprise EC12 Support

Version 3 will support the zEnterprise
 EC12, z196 and z114 System z platforms

Query Prioritization

- Brings System z workload management down to the individual query being routed to the Accelerator
- High Capacity
 - Support has been extended to include the entire Netezza 1000 line (1.28 PB)
- UNLOAD Lite
 - Reduces z/OS MIPS consumption, by moving the preparation off System z.

Smarter**Analytics**

Challenges associated with traditional approach to real-time scoring



Multiple infrastructures required

Predictive Analytics for Linux on System z

SPSS Modeler for Linux on System z



- Version: Modeler v15
- Data mining tool used for generating hypotheses and scoring
- Text analysis for unstructured data to model consumer behavior
- In-Transaction Scoring with DB2 z/OS: Embeds the Scoring Algorithm Directly within the Transactional Application





Smarter Analytics



Smarter**Analytics**





DB2 for z/OS Customer Trends

- Proliferation of mobile and other network-connected devices is driving increases in:
 - -transaction workloads
 - -data volumes
 - -24x7 requirements
- Continued focus on cost containment and resource efficiency
- Competitive pressures continue to drive an increasing need for innovation, analytics, and data integration
- DB2 for z/OS has leading edge capabilities to support these requirements and DB2 11 makes important improvements

DB2 11 Major Themes

Performance Improvements

- Improving efficiency, reducing costs, no application changes
- 0-5% for OLTP
- 20-30% for query workloads
- Less overhead for data de-compression
- Exploitation of new zEC12 hardware features

Continuous Availability Features

- Improved autonomics which reduces costs and improves availability
- Making online changes without affecting applications
- Online REORG improvements, less disruption
- DROP COLUMN, online change of partition limit keys
- Extended log record addressing capacity (1 yottabyte)
- BIND/REBIND, DDL break into persistent threads

Enhanced business analytics

- Faster, more efficient performance for query workloads
- Temporal and SQLPL enhancements
- Transparent archiving
- SQL improvements and IDAA enhancements

Simpler, faster DB2 version upgrades

- No application changes required for DB2 upgrade
- Access path stability improvements
- Product stability: support pre GA customer production





IBM announced DB2 11 Early Support Program → Learn more

DB2 11 Planned Solution for Extended RBA/LRSN

Expand the RBA and LRSN to 10 bytes

- RBA addressing capacity of 1 yottabyte (2**80)
- LRSN extended on left by 1 byte, on the right by 3 bytes
 - >30,000 years and 16Mx more precision
- 8 bytes is not sufficient to solve LRSN issues and may not give sufficient capacity for the longer term
- >NFM only (6 byte RBA/LRSN continues to be used in CM)
- > Once in NFM, DB2 continues to use 6-byte values until you take action to convert
- > Two conversion tasks:
 - Convert BSDSes to new format to enable logging with larger RBAs/LRSNs

a smarter

- Convert pagesets to new page format
- These tasks are optional
 - If you don't care about larger RBAs/LRSNs then you don't have to convert
 - But performance will be better if you convert BSDSes (avoid internal conversion overhead on log write)
- BSDSes can be converted without converting pagesets
- Pagesets can be converted in a piecemeal fashion
 - Expectation is that most customers will roll the conversion over a period of days/weeks/months

IBM

Some Planned DB2 11 RAS Improvements

BIND / DDL concurrency with persistent threads

- Use of persistent threads likely will increase in V10 with vstor relief
- More online schema changes
 - Alter partitioning limit keys
 - DROP column
 - Point in time recovery support for deferred schema changes
- Cancel DDF Threads new FORCE option
- Open data set limit raised to 200K
- REORG avoidance
 - Automatic cleanup of index pseudo deleted entries
- > Online REORG improvements if a REORG is needed, then the goal is non-disruptive
 - SWITCH phase performance improvements
 - Drain improvements to improve concurrency of log apply and switch phases
 - REORG REBALANCE SHRLEVEL(CHANGE)
- Easier query performance management
 - Optimizer input to Runstats
- DB2/RACF authorization control enhancements
- Buffer pool management improvements
 - New FRAMESIZE BP attribute for direct control of z/OS large page frame usage
 - Max/min size for WLM system-managed BPs

DB2 11 Some Planned New Application Features

- Global variables
- > SQLPL improvements (performance, manageability, function)
 - Autonomous transactions
 - Array data type support
- Alias/synonym support for Sequence Objects
- Grouping sets
- DPSI performance improvements
- DGTT performance improvements
- > Temporal data enhancements
 - Support for views
 - Special register support
- Transparent archive query
 - New DDL to relate the current table to the archive
 - Applications can query current + archive with no SQL changes
- Enhancements for DB2 and Big Data integration

Information Management

Enhancing DB2 Analytics on "z" with Big Data

- Much of the world's operational data resides on z/OS
- Unstructured data sources are growing fast
- There is a need to <u>integrate</u> this data so that insights from BigData sources can drive business actions
- DB2 is providing the connectors and the DB capability to allow DB2 apps to easily and efficiently access these data sources



Use of unstructured raw data is GROWING



DB2 11 Planning

- Dual mode migration (CM, ENFM, NFM)
- Migration from DB2 10 only (no skip)
- z/OS 1.13 or above. z10 or above.
- No pre-V9 bound packages
- Sysplex query parallelism support is removed

DB2 for z/OS Request for Enhancements (RFE)

More Effectively Communicating with Customer on Requirements

- Historically... internal IBM tool (FITS). No direct customer access to requirements
- Going forward with <u>DB2 for z/OS RFE</u> you can:
 - Directly manage/track your requirements greater accessibility
 - Access, vote, comment, and watch other public requirements
 - Directly interact with DB2 development
- DB2 for z/OS RFE Community link: http://www.ibm.com/developerworks/rfe/infomgmt/
- Online help and tutorials available from the RFE Community including YouTube videos

YouTube videos: http://www.ibm.com/developerworks/rfe/execute?use_case=tutorials

Watching



Customers

RFE Community

- •Submit and Edit
- •Search and View
- •Attachments
-
- •Voting
- •Commenting

Notifications (email and RSS)
Groups
Escalation process

- . .
- •And more...



Brand Development





Typical Utilization for Servers Windows: 5-10% Unix: 10-20% System z: 85-100%

System z can help **reduce** your floor space up to **75%-85%** in the data center







System z can lower your total cost of ownership, requiring as little as 30% of the power of a distributed server farm running equivalent workloads

The cost of storage is typically three times more in distributed environments





Thank You for Joining Us today!

If you would take a moment to fill out the feedback form which will display on the next slide, it would be greatly appreciated. Your comments are very important to us.

Go to www.ibm.com/software/systemz/events/calendar to:

- Replay this webcast
- View previously broadcast webcasts
- Register for upcoming webcasts