

DB2 10 for z/OS

Technical Overview

Jeff Josten
Distinguished Engineer
DB2 for z/OS Development

Session Code: S01
Tuesday 3rd May 2011 at 10:15am | Platform: z/OS





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.

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.

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

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

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

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

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

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

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

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



DB2 for z/OS Technical Strategy

- Continuous availability
- Performance and scalability
- Ease of management / autonomies
- Advanced application features



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

- Up to 20% utility CPU savings
- Compress indexes, save 50% disk
- Native SQL procedures
- More CPU on specialty engines

- Save up to 5-10% CPU batch & transactions out-of-the-box (rebind)
- 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

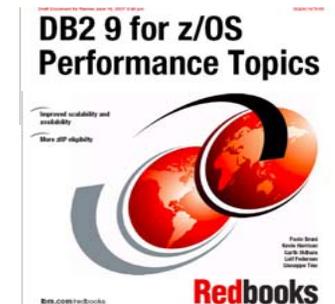
V8 out of service April 2012



DB2 Deep Synergy With System -

Key integration points include:

- Data sharing (availability and scale out)
- zIIP and other specialty engines
- Unicode conversion
- Encrypted communication & data
- Hardware data compression & encryption
- Cross-memory, memory protection keys
- Sorting
- Multi-core, large N-way
- 64-bit addressing and large memory
- z/OS Workload Manager
- z/OS Security Server (RACF)
- z/OS RRS integrated commit coordinator
- System z10 1 MB page size, decimal float
- Solid state disks
- zEnterprise z196, zBX, z10, ...



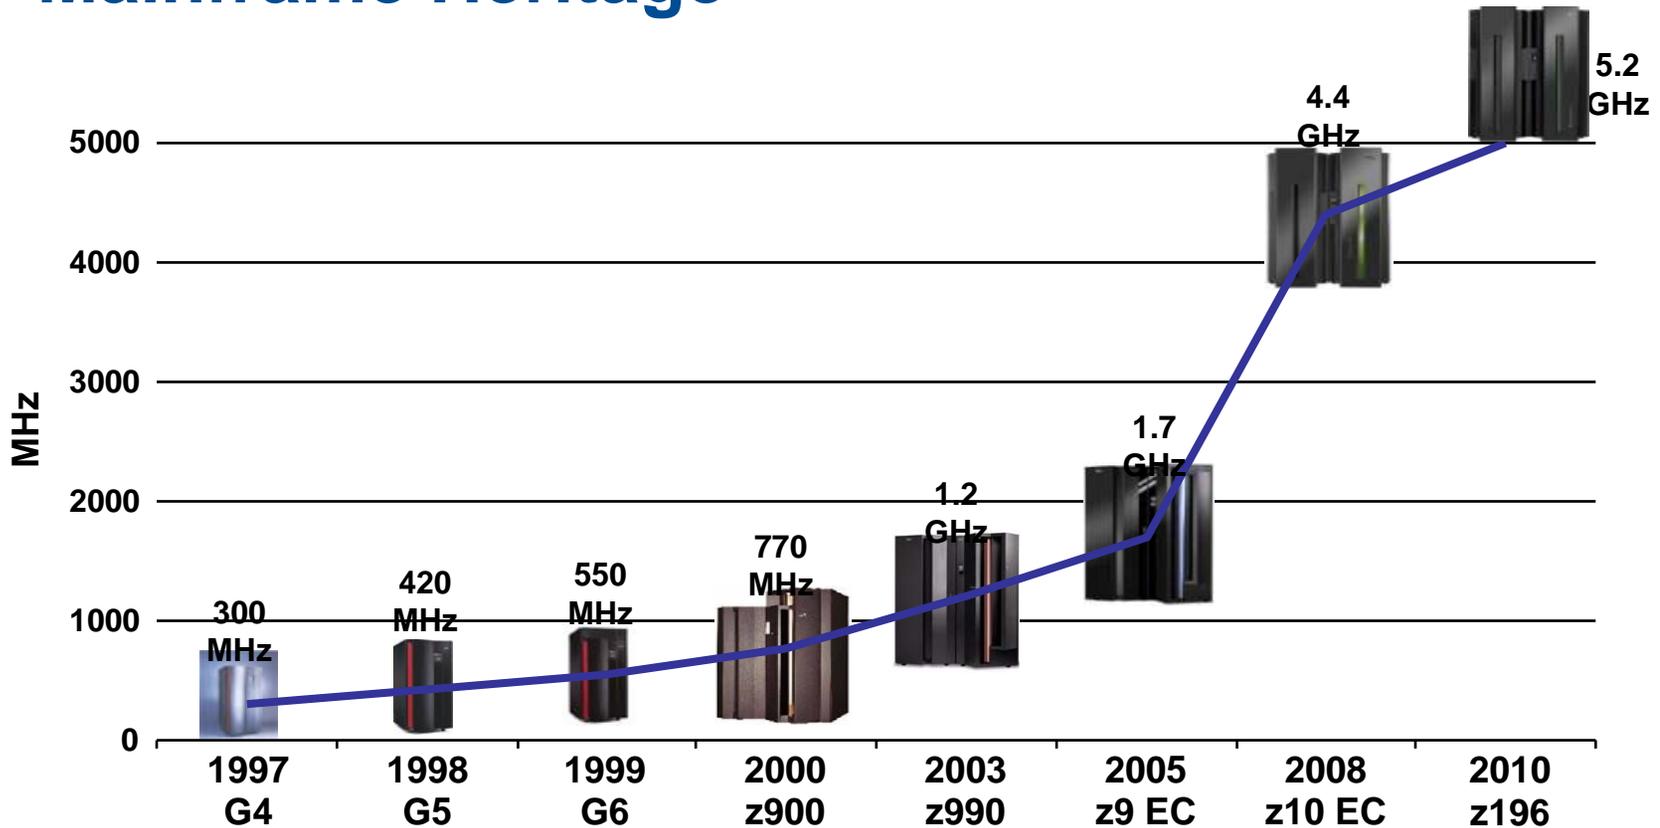


Hardware Trends Impacting DB2

- Drive towards multi core, slowing growth in processor frequency
 - 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
 - New programming paradigms (e.g. OpenCL)
- 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 over HDD



z196 Continues the CMOS Mainframe Heritage



- **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
- **z196** – Additional Architectural extensions & new cache structure



System zEnterprise Benefits for DB2

Taking System z synergy to the next level

- Faster CPUs, more CPUs, more memory → better DB2 performance, scalability
- Compression hardware expected to increase DB2 data compression performance
- Cache optimization, 192M L4 Cache expected to benefit DB2 work
- Hybrid architecture query performance acceleration with **IBM Smart Analytics Optimizer**
- Excellent synergy with DB2 10 → significant CPU reduction and scalability increase
 - CPU reductions
 - Remove key single system scaling inhibitors: virtual storage, latching, catalog, utilities, ...
 - Translation Lookaside Buffer changes expected to improve performance for 1MB page sizes
 - Buffer pool management





IBM zEnterprise System: The integration of Superior technologies



IBM zEnterprise 196 (z196)

- The Next generation of mainframe technology, more performance, more scale, more efficient
- Optimized to host large scale database, transaction, and mission critical applications

zEnterprise Unified Resource Manager

- Centralized management of heterogeneous resources for simplification and resiliency
- Unifies management of resources, extending IBM System z qualities of service end-to-end across workloads

zEnterprise BladeCenter Extension (zBX)

- Integrated IBM POWER7 blades and IBM x86 Blades
- High performance optimizers and appliances to accelerate time to insight and reduce cost



IBM zEnterprise Provides a New Dimension in Cloud Computing

The world's fastest and most scalable system:
IBM zEnterprise™ 196 (z196)

Unified management for a smarter system:
IBM zEnterprise Unified Resource Manager (zManager)

Scale out to a trillion instructions per second:
IBM zEnterprise BladeCenter® Extension (zBX)

Broad Network Access

Very large number of end user access from multiple sources including mobile devices

Rapid Elasticity

A new dimension of Scale. Most efficient platform for Large-scale Linux consolidation



Resource Pooling

1000s of virtualized systems across a heterogeneous resource pool

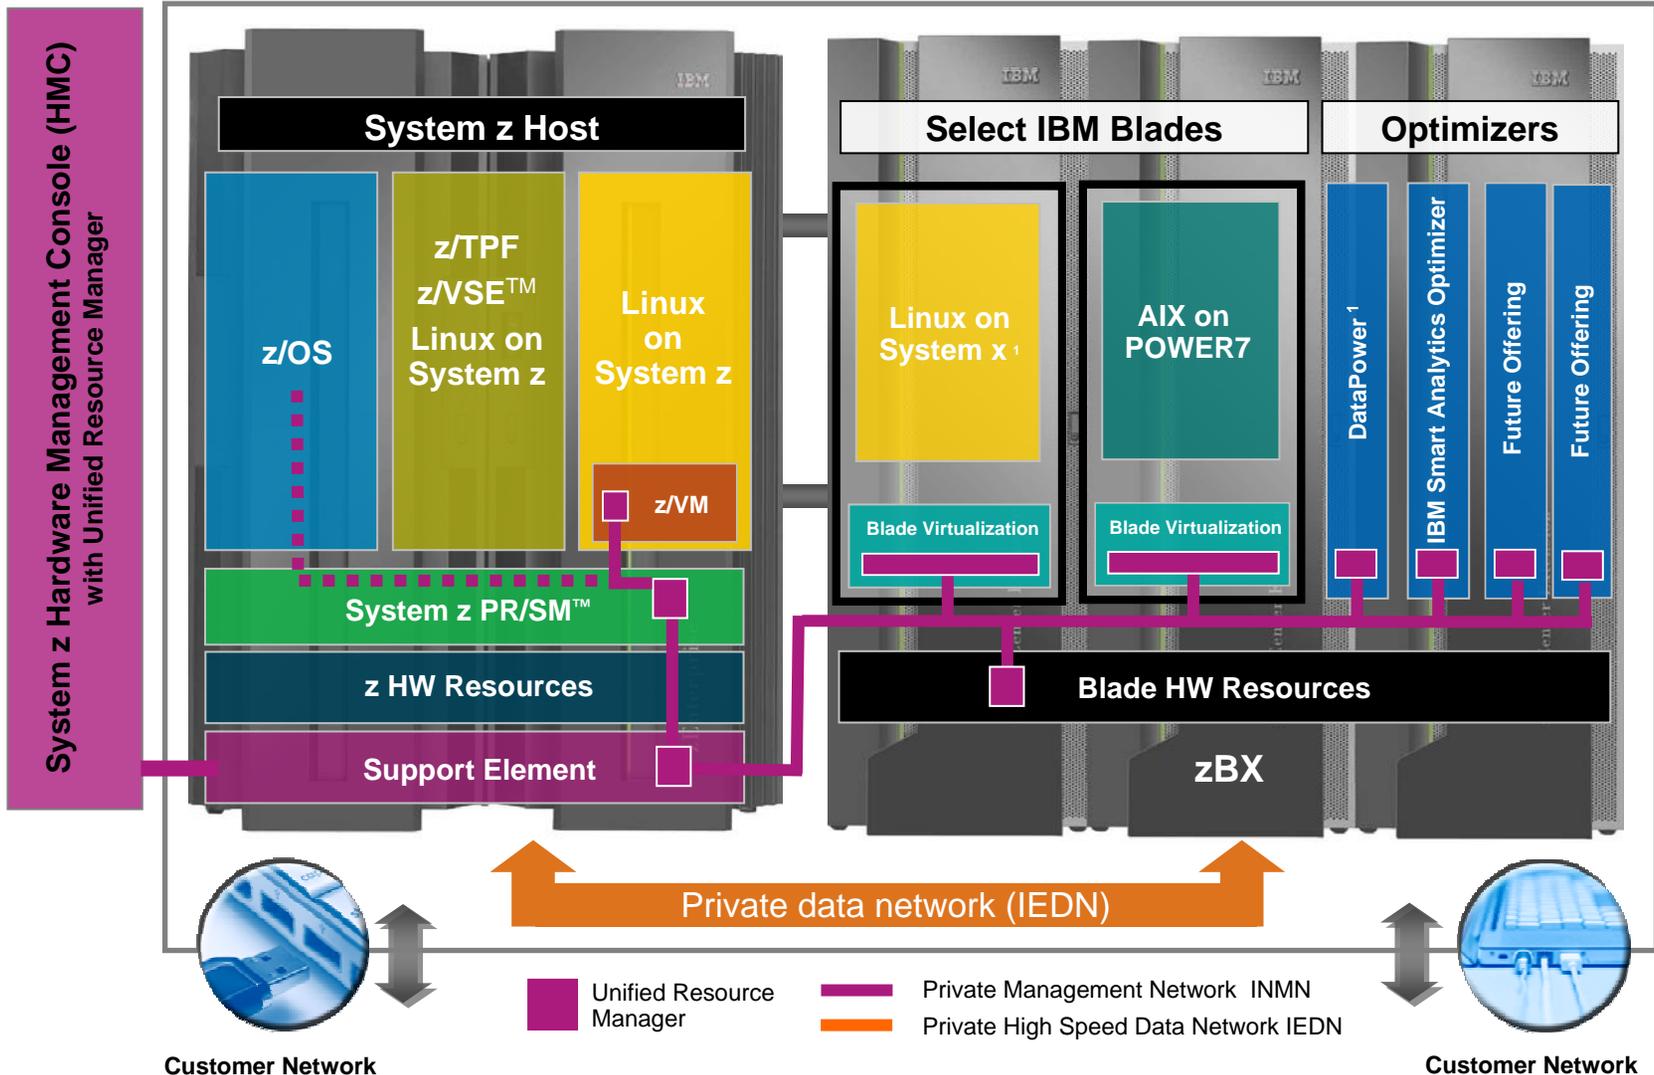
Measured Service

Meter, monitor, and track workloads for chargebacks and capacity expectations

On Demand Self-Service

Automate provisioning and service requests reducing provisioning cycles from weeks to minutes

A Look Inside the zEnterprise System



¹ All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

IBM Smart Analytics Optimizer V2

- *Next generation Smart Analytics Optimizer*
 - *Capitalizing on Netezza technology*
 - *Access through and highly integrated with DB2 for z/OS*

- *Beta begins in 3Q 2011*

- *Investment protection*
 - *The full value of ISAO V1 will be applied to ISAO V2*

Taking it to the Next Level

- *Next generation Smart Analytics Optimizer*
 - Enhances expandability
 - More robust query execution
 - Lowers cost
 - Removes past restrictions



IBM Smart Analytics Optimizer

Capitalizing on the best of both worlds – System z and Netezza

What is it?

The IBM Smart Analytics Optimizer is a workload optimized, appliance add-on, that enables the integration of business insights into operational processes to drive winning strategies. It accelerates select queries, with unprecedented response times.



How is it different

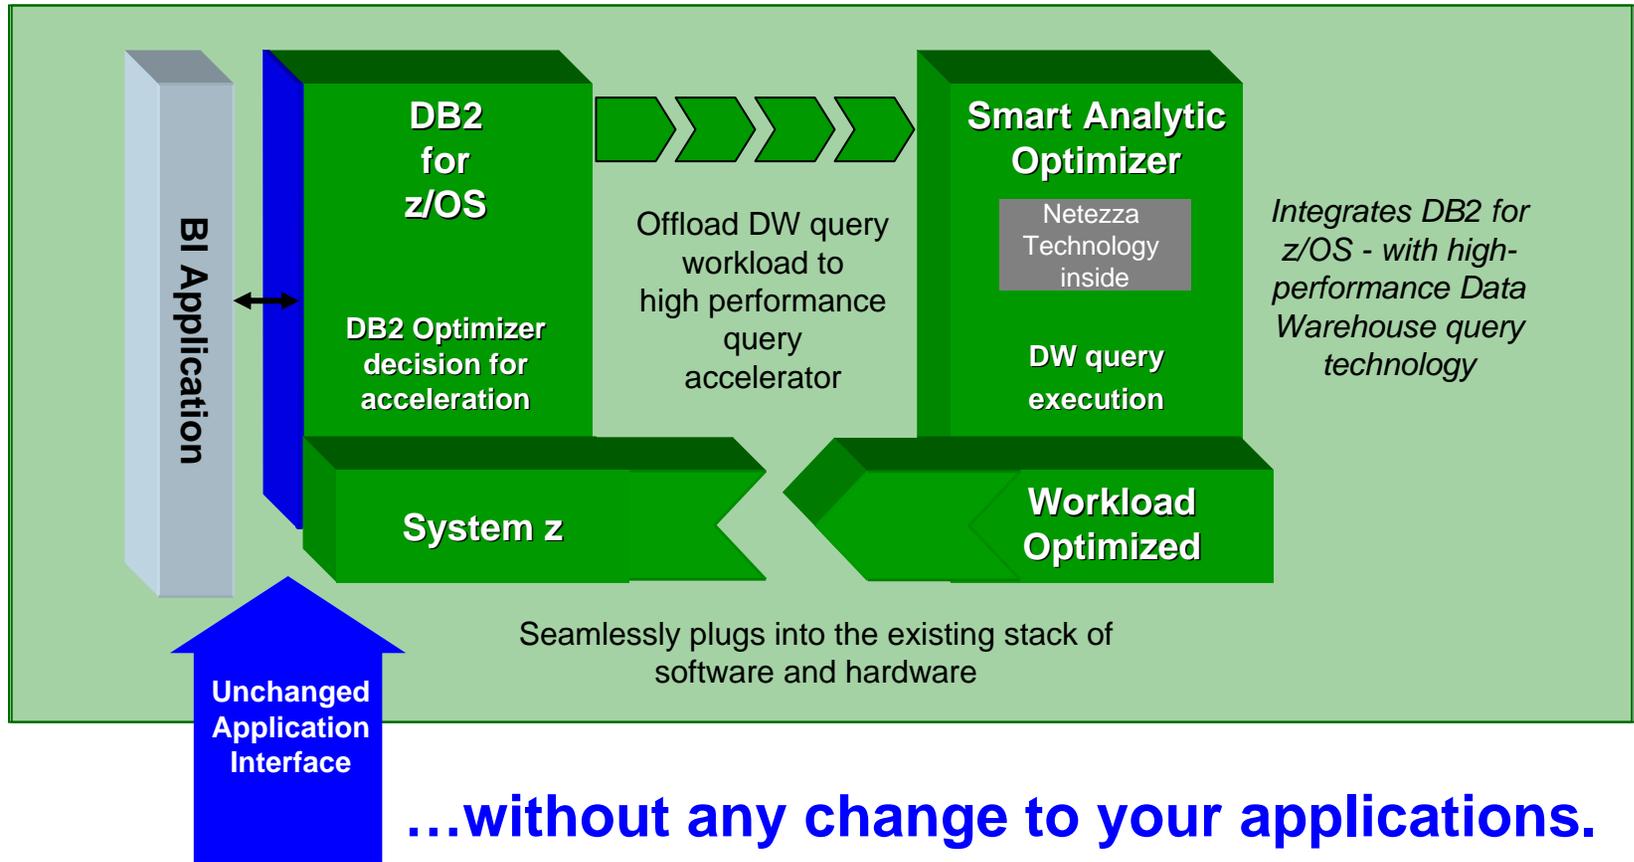
- **Performance:** Unprecedented response times to enable 'train of thought' analyses frequently blocked by poor query performance.
- **Integration:** Connects to DB2 through deep integration providing transparency to all applications.
- **Self-managed workloads:** queries are executed in the most efficient way
- **Transparency:** applications connected to DB2 are entirely unaware of the Optimizer
- **Simplified administration:** appliance hands-free operations, eliminating many database tuning tasks

Breakthrough Technology Enabling New Opportunities

Optimizing to the Workload

Marrying the best of each environment

Total solution remains centrally managed by System z...



...without any change to your applications.

Key Enhancements in V2 Release

- Extending acceleration to significantly larger number of queries
- Expanded size of the data to be accelerated
- Improved concurrent query execution
- Incremental update by partition
- DB2 9 and DB2 10 for z/OS support
- Supported only connected to z196/z114 (zEnterprise) forward



Mainframe Innovation: Specialty Engines



**Internal Coupling
Facility (ICF) 1997**



**Integrated Facility
for Linux® (IFL)
2000**



**System z Application
Assist Processor (zAAP)
2004**

Eligible for zAAP:

- Java execution environment
- z/OS XML System Services



**IBM System z Integrated
Information Processor and
(2006)**

Eligible for zIIP:

- DB2 remote access, XML, large parallel queries, utilities (index, sort, stats)
- ISVs
- IPsec encryption
- XML System Services
- Global Mirror (XRC)
- HiperSockets for large messages (e.g. DRDA)
- IBM GBS Scalable Architecture for Financial Reporting
- z/OS CIM Server
- zAAP on zIIP

* Statements represent the current intention of IBM. IBM development plans are subject to change or withdrawal without further notice.



DB2 & IBM zIIP Add Value to Database Work

Portions of the following DB2 for z/OS V8, DB2 9 and 10 workloads may benefit from zIIP or zAAP for XML (DB2 9 in blue, DB2 10 in green)*:

1 – DRDA over TCP/IP connections

- **DB2 9 for z/OS Remote native SQL procedures**
- **DB2 9 XML parsing, XML schema validation**
- **Increased portion of DRDA redirected to zIIPs to 60%**

Improved performance via reduced processor switching

2 - Requests that use parallel queries

- **DB2 9 higher percentage of parallel queries zIIP eligible**
- **DB2 10 more queries eligible, more parallelism**

3 - DB2 Utilities LOAD, REORG & REBUILD functions used to maintain index structures and sort

- **DB2 10 RUNSTATS – options other than column group, inline**

4 - DB2 10 buffer pool prefetch and deferred write

DB2 10 for z/OS

GA'ed October 2010

Completed Largest Beta Ever

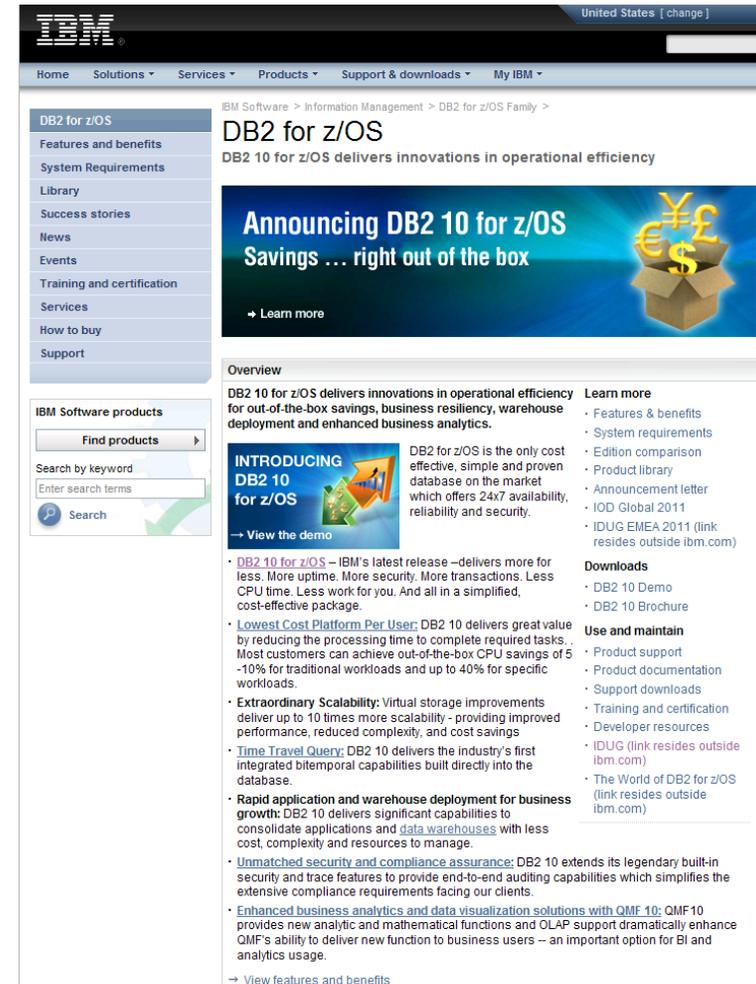
- 23 WW customers
- +10 Extended Beta
- Over 80 vendors

Fastest uptake out of the gate

- As of 1Q11, almost 100 customer orders
- More than 2x the number of licenses vs. V9
- About 25% are migrating from V8
- Every core beta customer is continuing with migration plans

First customer already in production

- Migrated from V8 to V10
- Quality/stability looking good



The screenshot shows the IBM website for DB2 10 for z/OS. The page features a navigation bar with links for Home, Solutions, Services, Products, Support & downloads, and My IBM. The main content area includes a sidebar with a table of contents for DB2 for z/OS (Features and benefits, System Requirements, Library, Success stories, News, Events, Training and certification, Services, How to buy, Support). The main content area has a header for 'DB2 for z/OS' with a sub-header 'DB2 10 for z/OS delivers innovations in operational efficiency'. Below this is a large banner with the text 'Announcing DB2 10 for z/OS Savings ... right out of the box' and a 'Learn more' link. The banner includes an image of a cardboard box with currency symbols (¥, €, \$) emerging from it. Below the banner is an 'Overview' section with a detailed description of the product's benefits, such as operational efficiency, business resiliency, and enhanced business analytics. There are also sections for 'Learn more' (with links to Features & benefits, System requirements, Edition comparison, Product library, Announcement letter, IOD Global 2011, IDUG EMEA 2011), 'Downloads' (with links to DB2 10 Demo, DB2 10 Brochure), and 'Use and maintain' (with links to Product support, Product documentation, Support downloads, Training and certification, Developer resources, IDUG, and The World of DB2 for z/OS). A search box is visible in the sidebar.



What do you want from DB2?

- ✓ Reduce CPU time 5% - 10%
- ✓ Increased ability to scale up
- ✓ Easier security compliance and audit
- ✓ Improved productivity
 - ✓ Temporal and more enhanced SQL & XML
 - ✓ Administration, scaling and performance
 - ✓ Move from DB2 9 or V8
- ✓ Ready for production, stable and available
- ✓ Customer references

Information Management

Draft Document for Review February 25, 2011 11:01 am

IBM
SG24-7942-00

DB2 10 for z/OS Performance Topics

Triton
CONSULTING

DB2 10 for z/OS
A Smarter Database for a Smarter Planet



DB2 10 has it!



DB2 10 Customers seeing reduced costs, simplified workloads through proven technology



Reduced Costs	Simplified Workloads	Proven Technology
<p>“Based on the performance metrics from our controlled test environment, we see a significant amount of CPU and Elapsed time savings. This release has many features that will help bring down our operating costs.”</p> <p>Morgan Stanley DB2 Team</p> 	<p>“With DB2 10 able to handle 5-10 times as many threads as the previous version, the upgrade will immediately give the bank some much-needed room for future workload growth while simultaneously reducing their data sharing overhead.”</p>  <p>Paulo Sahadi - Senior Product Information Management Division at Banco do Brasil</p>	<p>“Every single SQL statement we have tested has been better or the same as our current optimal paths – we have yet to see any significant access path regression. We had to spend a lot of time tuning SQL with DB2 9, but we expect that to disappear when we upgrade to DB2 10.”</p> <p>Philipp Nowak, BMW DB2 Product Manager</p>
<p>“We are particularly interested in the performance improvements due to the potential CPU reductions that we realized during our DB2 10 Beta testing. Our early testing has shown out-of-the-box processing cost reductions of between 5% - 10% and for some workloads as high as 30%. Potential cost savings of this magnitude cannot be ignored given today’s business climate.”</p>  <p>Large Global Bank</p>	<p>“We are really thrilled about “Temporal Data” feature – this feature has the potential to significantly reduce overheads. We have estimated that 80% of our existing temporal applications could have used “the DB2 10 temporal features” instead of application code - this feature will drastically save developer time, testing time – and even more importantly make applications easier to understand so improve business efficiency and effectiveness.”</p> <p>Frank Petersen – System Programmer Bankdata</p> 	<p>The new audit capabilities in DB2 10 will allow tables to be audited as soon as they are created, which is an obvious benefit for the business and will reduce costs and simplify our processes”</p> <p>Guenter Schinkel -Postbank Systems AG</p> 

For more customer references visit

<http://www.ibm.com/software/data/db2/zos/testimonials.html>



DB2 10 for z/OS: Cost Savings

CPU reductions for transactions, queries, and batch

- CPU reductions of 5-10% for traditional workloads
- CPU reductions of up to 20% for new workloads
- Up to additional 10% CPU savings using new functions
- For static SQL, REBIND typically required

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

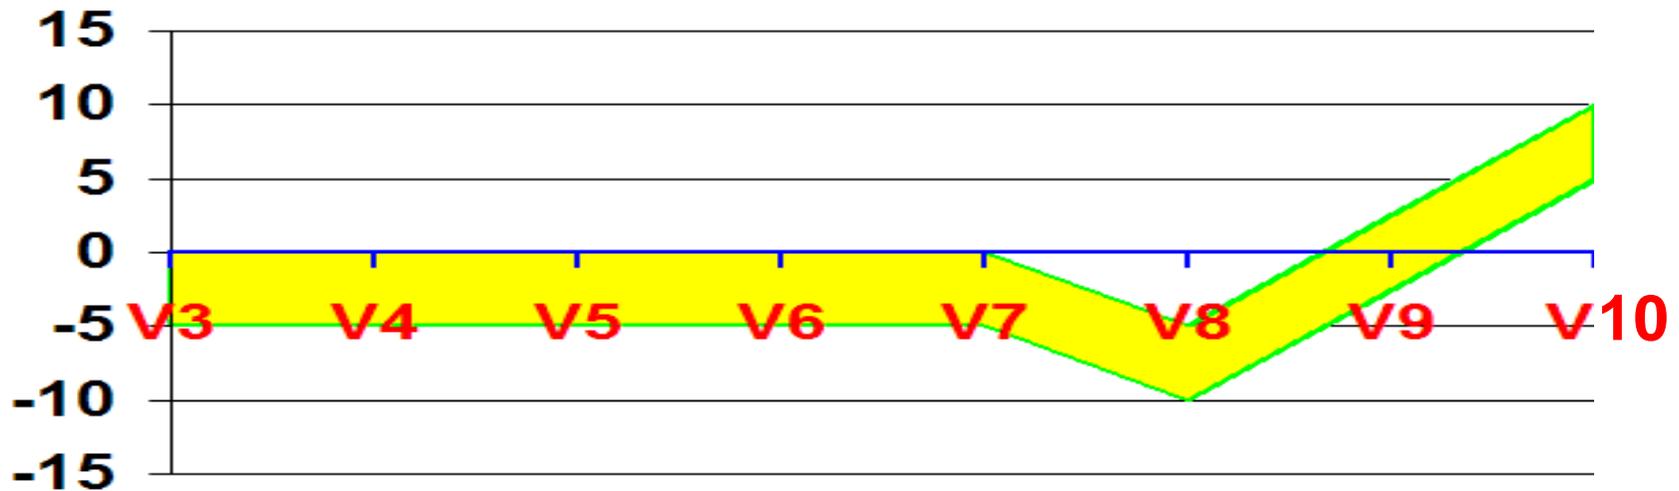




DB2 10 Performance

- Most customers 5% - 10% CPU reduction out of the box after rebind
- Some workloads and customer situations can reduce CPU time up to 20%

**Average %CPU improvements
version to version**





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
- JCC Type2 and ODBC for z/OS performance improvements
- Parallel index I/O at insert
- Workfile in-memory enhancements
- Index list prefetch
- Solid State Disk tracking in Real Time Statistics
- Buffer pool enhancements
 - Utilize 1MB page frames on z10 or z196
 - “Fully in memory” option (ALTER BUFFERPOOL)



Performance Enhancements requiring REBIND (CM)

- Most access path enhancements
- Further SQL runtime improvements
- Use of RELEASE(DEALLOCATE)
- 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
- Safe query optimization
- 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
- If migrating from V8, get new RUNSTATS before mass rebind



V10 Virtual Storage Relief Enables other performance options

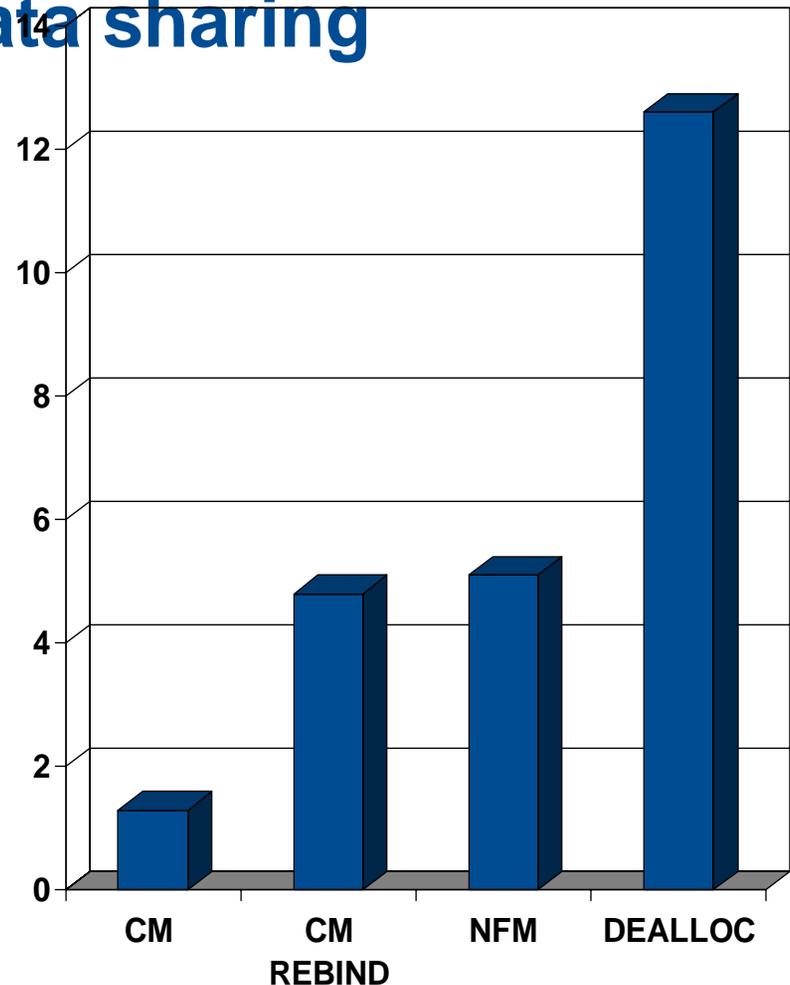
- After REBIND in v10, most thread storage moves above the bar (CM or NFM)
- Now you can consider using the RELEASE(DEALLOCATE) BIND option in more cases
 - IBM measurements show up to 10-20% cpu savings
 - Effective for high volume transactions with thread reuse
 - DBM1 vstor was an inhibitor to using this option in the past
- New feature called High Performance DBATs for DDF threads
 - Prior to v10, DDF enforced RELEASE(COMMIT) for all DDF threads
 - This was done because of DBM1 vstor concerns, which are now gone in v10
 - With v10, this enforcement is removed, and RELEASE(DEALLOCATE) will now be honored
 - KEEP DYNAMIC=YES will not get benefit from this enhancement
- With v10 vstor relief, MAXKEEPD can be increased which could improve dynamic SQL performance
- Tradeoff some increase in memory usage for reduced cpu



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



CPU Percent reduced from DB2 9



Performance Enhancements requiring NFM

- DB2 catalog concurrency and productivity
- Compress on insert
- Most utility enhancements
- LOB streaming between DDF and rest of DB2
- SQL Procedure Language performance improvements
- Workfile spanned records, partition by growth
- Access to currently committed data
- Insert improvement for universal table spaces
- LRSN spin avoidance for inserts to the same page
- Efficient caching of dynamic SQL statements with literals



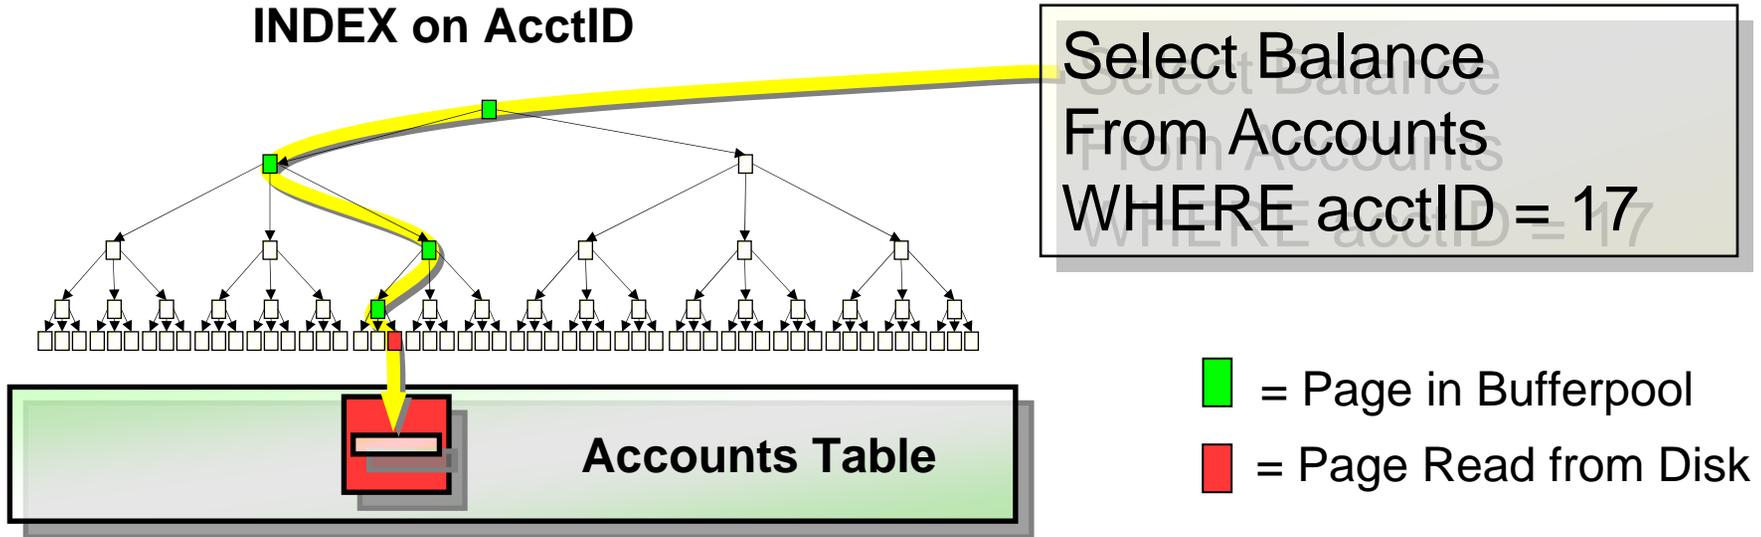
Performance Enhancements need NFM + DBA work

- Hash access path
- Index include columns
- Inline LOBs
 - Index on expression now possible for LOB columns
 - Important for spatial performance
 - LOAD/UNLOAD performance improvements
 - LOB compression for inline portion
- DEFINE NO for LOB and XML columns
- MEMBER CLUSTER for universal table space
- Alter to universal table space, page size, data set size, segment size
- Online reorg for all catalog and directory table spaces



Index to Data Access Path

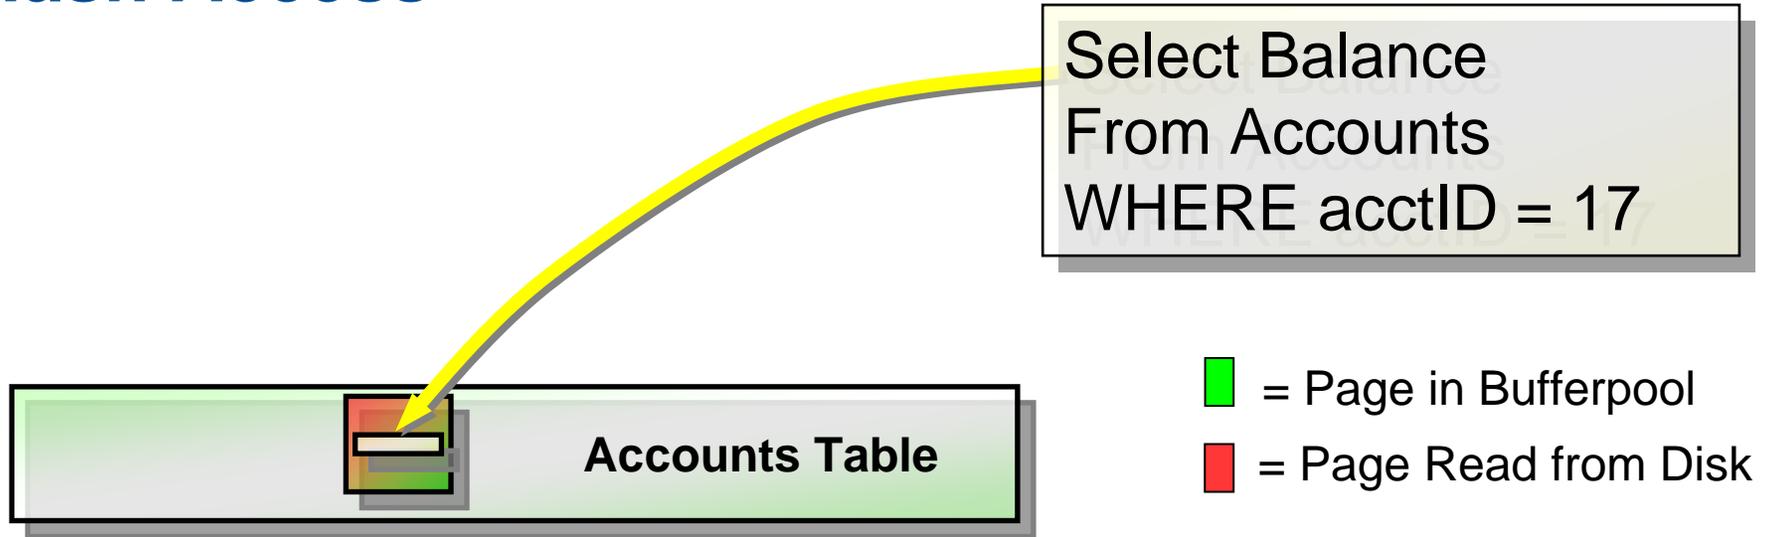
INDEX on AcctID



- Traverse down Index Tree
 - Typically non-leaf portion of tree in the bufferpool
 - Leaf Portion of the tree requires I/O
 - Requires searching pages at each level of the index
- Access the Data Page
 - Typically requires another I/O
- For a 5 Level Index
 - 6 GETP/RELPS, 2 I/O's, and 5 index page searches



Hash Access



- Hash Access provides the ability to directly locate a row in a table without having to use an index
- Single GETP/RELP in most cases
- 1 Synch I/Os in common case
 - 0 If In Memory Table
- Greatly reduced Search CPU expense

Hashing Summary



- Provides fast, direct location of most rows
 - Reduces I/O and CPU in most cases
 - Can replace an existing Primary or Unique Key Index
 - Faster Insertion/Deletion
- Size of Fixed Size Hash Area is important
 - Too small and performance degrades, too large and space is wasted
- DB2 helps you manage the size
 - REORG AUTOESTSPACE YES
 - RTS tracks the number of overflowed entries
- If clustering is important for query performance, then be aware that Hash will eliminate these benefits
- LOAD performance is slower with hash



Some Beta Customer Performance Feedback

Workload	Results
Customer1: Distributed Concurrent Insert	50% DB2 elapsed time reduction 15% chargeable CPU reduction after enabling high perf DBAT
Customer2: CICS online transactions	Approx. 7% CPU reduction in DB2 10 CM after REBIND, Another 4% reduction with 1MB page usage
Customer3: CICS online transactions	Approx 5% CPU reduction
Customer4: Data sharing heavy concurrent insert	38% CPU reduction
Customer5: Queries	Average CPU reduction 28% from V8 to DB2 10 NFM
Customer6: Batch	Overall 28% CPU reduction after rebind packages
Customer7: DDF OLTP	40% CPU reduction for JDBC stored procedures workload, 15% CPU reduction for securities trading app



Beta Customer Feedback on Selected New Functions

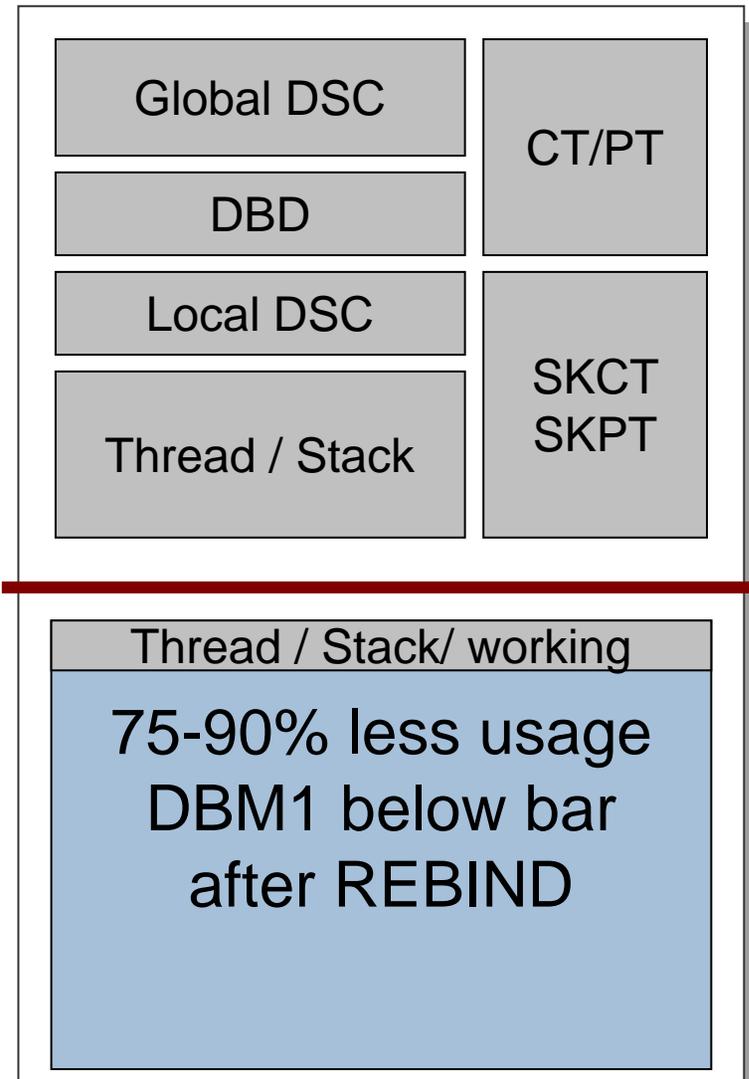
Workload	Results
Multi row insert (data sharing)	33% CPU reduction from V9, 4x improvement from V8 due to LRSN spin reduction
Parallel Index Update	30-40% Elapsed time improvement with class 2 CPU time reduction
Inline LOB	SELECT LOB shows 80% CPU reduction
Include Index	17% CPU reduction in insert after using INCLUDE INDEX
Hash Access	<p>20-30% CPU reduction in random access</p> <p>16% CPU reduction comparing Hash Access and Index-data access. 5% CPU reduction comparing Hash against Index only access Further improvements delivered late in the beta program.</p>



Virtual storage improvements

- 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 KEEP DYNAMIC=YES

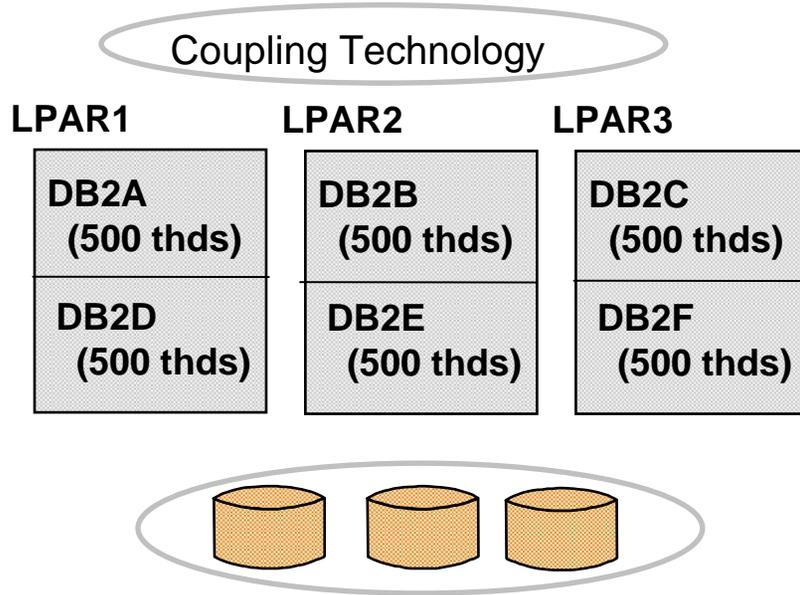
DB2 10





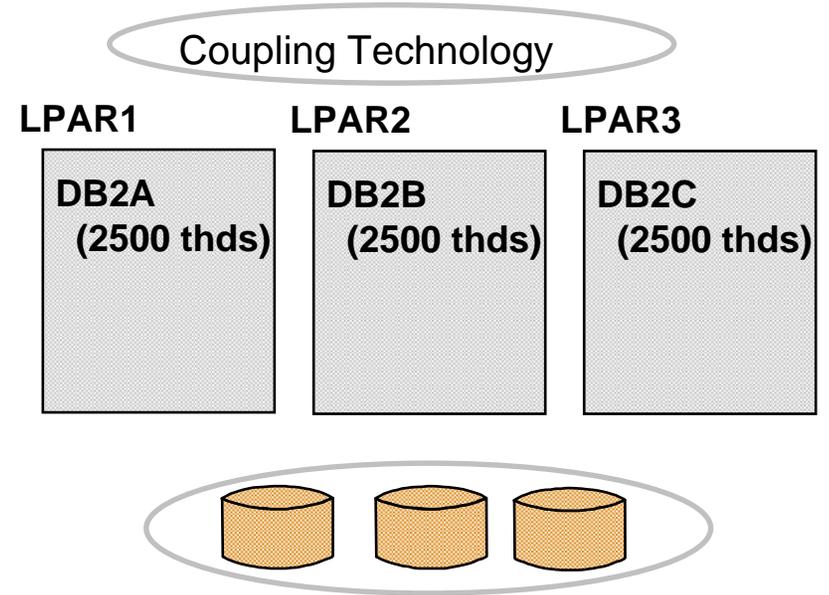
Running Many Active Threads

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



DB2 10 and Real Storage Usage

- DB2 10 will typically use somewhat more real storage than V8 or V9
- For one JDBC Type4 workload, V10 real storage usage increased 3-12% over V9 for the same number of threads
- If your V8 or V9 has little/no paging then additional memory is probably not needed when going to V10
 - If paging rates increase, then add more memory if possible
- Increasing threads, or converting to RELEASE(DEALLOCATE) will further increase memory usage
- APAR PM24723 adds real storage monitoring and management enhancements
 - IFCID 225 extensions
 - z/OS APAR OA35885 is required (to provide stats for 64-bit memory objects)
 - New REALSTORAGE_MANAGEMENT zparm to control when DB2 frees real frames back to z/OS
 - Default=AUTO: DB2 monitors paging rates to decide when to discard frames
 - New messages (DSNV516I, 517I) for when paging rate thresholds cause DB2 to free real frames



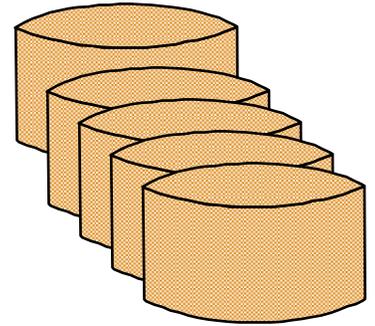
Other System Scaling Improvements

- Other bottlenecks can emerge in extremely heavy workloads
 - Reduced latch contention
 - Improved efficiency for latch suspend/resume
 - 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
- DB2 10 NFM catalog restructure improves BIND / Prepare concurrency
- SPT01 64GB constraint
- Improved accounting rollup, compress SMF option
 - Reduced SMF data volume
- Lower overhead for very large buffer pools



Major changes in DB2 10 catalog & directory

- Improve availability and productivity
- Increase maximum size substantially
- Reduce contention: BIND, Prepare, utilities
- Catalog changes: Remove links
 - 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





Data Sharing Improvements

- ACCESS DATABASE command wildcarding support - V9 PK80925
- Sub-group attach (v9 usermod)
- BP scan avoidance
- Delete data sharing member (planned for post GA)
 - Offline utility for “deactivate”, “reactivate”, “destroy”
- MEMBER CLUSTER support for UTS
- DDF Restart Light enhancements: Handle DDF indoubt URs
- Online DDF changes (planned for post GA)
- Auto rebuild CF lock structure on long IRLM waits during restart
 - Can avoid group-wide shutdowns
- LRSN spin avoidance for inserts to the same page (e.g. Multi Row Insert)
- IFCID 359 for index split
- New zparm to force deletion of CF structures on group start (e.g. DR testing)
- Expedited GBP DELETE_NAME processing
 - Avoid sending XI signals by deleting data only
 - Avoid potential lock timeout conditions when there are lots of directory entries for an object



Continuous Availability: More Online Schema Changes

- ALTER TABLESPACE

- Page size (not XML) (BUFFERPOOL)
- DSSIZE
- SEGSIZE
- Table space type
 - Single table simple -> PBG (inherit MC)
 - Single table segmented -> PBG
 - Classic partitioned -> PBR (inherit MC)
 - PBG -> Hash
- MEMBER CLUSTER

ALTER TABLESPACE ... MAXPARTITIONS m

ALTER TABLESPACE ... SEGSIZE s

ALTER TABLE ... ADD ORGANIZE BY HASH

- ALTER INDEX

- Page size (BUFFERPOOL)
 - In DB2 9 this was immediate with RBDP set

- Other schema change enhancements

- Table space no longer needs to be stopped to alter MAXROWS
- Object no longer needs to be stopped to alter BPOOL in data sharing



Other Availability Improvements

- Access currently committed data
- Change DDF location alias names online
 - New MODIFY DDF ALIAS command
 - New “dynamic alias” concept allows you to dynamically switch connections to an alias to different members
 - E.g. “penalty box” concept for poorly behaved applications
- Online DDF CDB changes
 - LOCATIONS, IPNAMES, IPLIST
- Dynamic add of active logs
 - New –SET LOG NEWLOG option
- Pre-emptable backout



DB2 10 REORG – improved availability & removed restrictions

- Reduced need for REORG INDEX
 - List prefetch of index leaf pages based on non-leaf information for range scans
- Improved performance for part-level REORG with NPIs & REORG INDEX
 - Index list prefetch results in up to 60% elapsed time reduction
- Reduced need for REORG with compress on insert
- New REORGCLUSTSENS RTS column
 - If no clustering-sensitive queries then avoid REORG to restore clustering
- REORG SHRLEVEL CHANGE for all cat/dir pagesets
- REORG SHRLEVEL CHANGE with REBALANCE (planned for post GA)
- REORG SHRLEVEL CHANGE for LOBs
- REORG SHRLEVEL REFERENCE|CHANGE to remove REORP



DB2 10 REORG – improved availability & removed restrictions

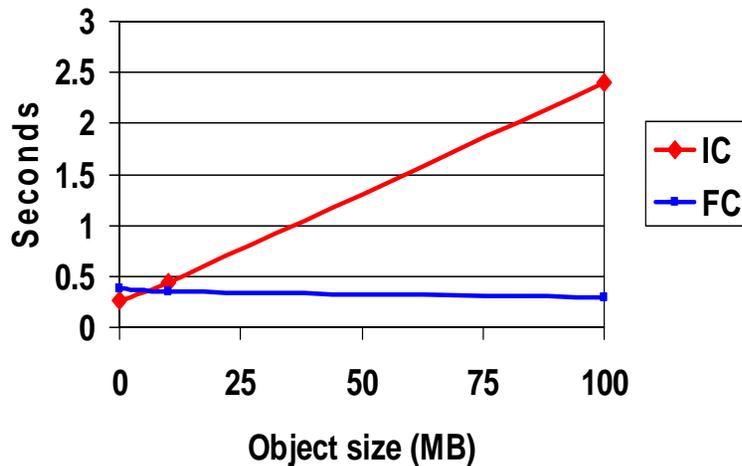
- REORG FORCE option to cancel blocking threads
 - Same process as –CANCEL THREAD so requires thread to be active in DB2 for it to be cancelled
- Reduce application outage on REORG with inline stats
- REORG of multiple part ranges
 - Retrofitted to V9
 - LISTDEF support is not retrofitted
- New AUX keyword on REORG for improved LOB handling
 - Permit rows to flow between partitions
 - Allows REORG REBALANCE with LOB columns
 - Allows ALTER of LIMITKEY with LOB columns



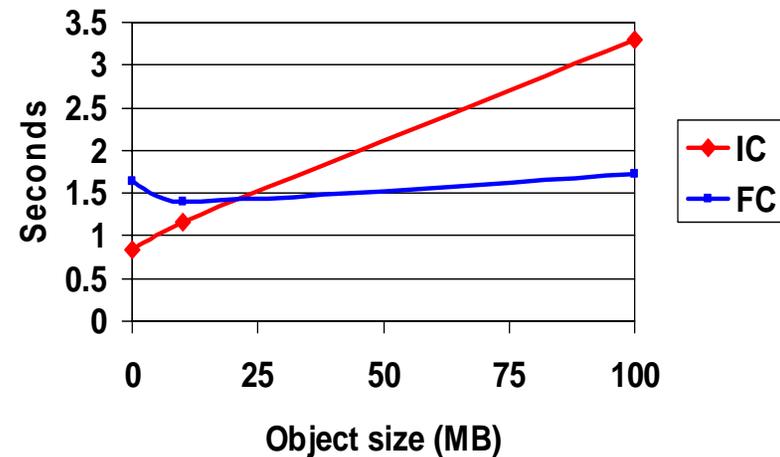
V10: COPY

- Dataset-level Flashcopy support
 - COPY, RECOVER, REORG, LOAD, REBUILD INDEX, REORG INDEX
 - New zparms & utility parms to govern
 - Dramatic CPU & elapsed time reduction
 - Create consistent imagecopies from SHRLEVEL CHANGE

CPU time per object (z10)



Elapsed time per object (z10)





V10: RECOVER

- New BACKOUT YES option for point in time recovery
 - True rollback, not run of generated SQL undo statements
 - Requires COPY YES for indexes
- VERIFYSET option to fail PIT recovery if entire set not included
 - Base, LOB, XML, history – not RI
- ENFORCE option to avoid CHKP/ACHKP when subset of set recovered
 - Improved performance due to avoidance of set checking (RI, aux)

Business Security & Compliance

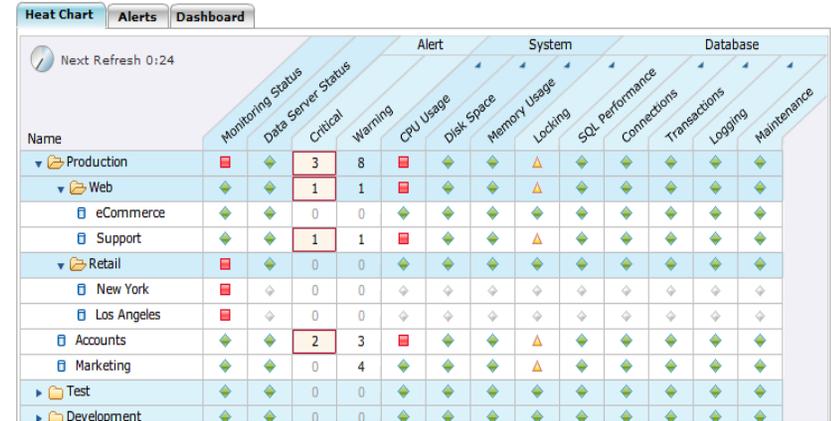
- Protect sensitive data from privileged users & improve productivity
 - SECADM & DBADM without data access
 - Usability: DBADM for all DB
 - Revoke without cascade
- Separate authorities to perform security related tasks, e.g. security administrator, EXPLAIN, performance monitoring and management
-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





DB2 10: Productivity – Doing More with Less!

- Easier scaling, simpler memory management
- Reduce contention, more online processing
- Access path stability: Post GA
- Reduced need for REORG
 - Build compression dictionary on the fly
 - Index list prefetch enhancements
 - Row-level sequential detection
- Auto statistics collection
- Configure IBM UDFs and stored procedures
- Allow one SDSNEXIT data set for many subsystems
- Statement level monitoring
- New DSNTIJXZ job to update migration input datasets with current zparm values
- DDF thread management enhancements



Heat Chart Alerts Dashboard
Next Refresh 0:24

Name	Monitoring Status		Data Server Status		Alert		System			Database		
	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	█	█	█	█	█	█	█	█	█	

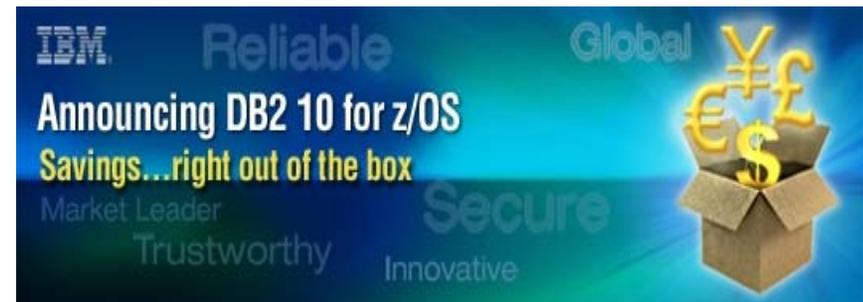
Manual invocation of
•RUNSTATS
•REORG



Query Processing Enhancements

- **Performance Improvements**

- Improved caching of dynamic SQL with literals
- Safe Query Optimization
- Aggressive View Merge
- IN List Processing
- SQL Pagination
- Parallelism Enhancements
- Index include columns



- **Access Path Stability**

- Relief from package REBIND regression



More Query Enhancements

- Optimization techniques
 - Remove parallelism restrictions; more even parallel distribution
 - Scalability: memory and latching relief allow more parallel
 - Optimization validation with Real Time Statistics
 - In-memory techniques for faster query performance
 - Multiple IN-List matching
 - IN-List predicate transitive closure
- RID overflow to workfile
 - Mitigate increased workfile usage by increasing RID pool size (default increased in DB2 10).
 - MAXTEMPS_RID zparm for maximum WF usage for each RID list
- Sort performance enhancements
 - Avoid workfile usage for small sorts
 - Hash support for large sorts



Query Parallelism enhancements

- Removing parallelism restrictions
 - Allow parallelism if a parallel group contains a work file
 - Support parallelism with multi-row fetch
- Parallelism effectiveness
 - Dynamic Record Range partitioning
 - Introduce sort composite to allow even distribution at runtime
 - Straw Model parallelism
 - Fine grained work elements to distribute work



DB2 10 Application Enablement and Portability

- Bitemporal data (data versioning)
- 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
- SQL Enhancements
 - Currently committed locking semantics
 - Implicit casting or loose typing
 - Timestamp with time zone
 - Variable timestamp precision – seconds to picoseconds
 - Moving Sum, Moving Average
 - SQLPL performance improvements



Versioned data or Temporal Data

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





DB2 SQL 2010



z z/OS 10
common

luw Linux, Unix & Windows 9.8

- 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 cast, date/time changes, currently committed, moving sum & average, index include columns, PureScale**
- 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, temp table compression
- u**
- w**



pureXML improved performance & 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 XML document nodes
- Stored procedure, UDF, Trigger enhanced support
- XML index matching with date/timestamp
- CHECK DATA utility checks XML



Expanding DB2 for z/OS ISV community





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....

Migration, fallback and data sharing coexistence fully supported

Mix of DB2 9 and 10 or DB2 V8 and 10

Key considerations:

- 
- Risk/reward analysis
 - What's your 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?
 - Do you have best practice service and test processes?
 - Migration cost savings is not 2X versus two migrations
 - Migration considerations for two versions still apply
 - Larger migration project, longer migration timeline
 - Applications and ISVs need to be ready
 - Timing: V8 end of service April 2012, other software, service & test process



DB2 10 News

Items that are planned for post GA delivery, based on customer input

- APREUSE, APCOMPARE – PM33767
- Delete data sharing member - PM31009
- Enhancements for new DBA authorities - PM28296
 - Prevent privileged users from stopping audit traces
 - No implicit system privileges for DBADM
- Inline LOBs for SPT01 - PM27811
 - Compression, BIND performance
- Online REORG concurrency for materializing deferred ALTERs – PM25648
- Temporal enhancements
 - TIMESTAMP WITH TIMEZONE support (PM31314)
 - Enhancement for data replication (PM31315)
 - ALTER ADD COLUMN, propagate to history table (PM31313)
- New system profile filters based on “client info” fields - PM28500
 - 3 new columns for userid, appname, and workstation
 - Wildcard support: if column is “*” then all threads pass that qualification
- New zparm force deletion of CF structures on group restart - PM28925
- Relief for incompatible change in CHAR of decimal data - PM29124
- Real storage monitoring enhancements – PM24723
- Hash LOAD performance – PM31214
- DSSIZE > 64GB – PM32429
- REORG REBALANCE SHRLEVEL CHANGE – Apar forthcoming (v9, v10)



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 function
- Premigration check DSNTIJPA PM04968
- Info APARs II14477 (DB2 9) II14474 (V8)

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

- Private protocol → DRDA
- Old plans and packages V5 or before → REBIND
- Plans containing DBRMs → packages
- ACQUIRE(ALLOCATE) → ACQUIRE(USE)



Sample Improvements for Guesstimate

- Run time CPU reductions 5% - 10%
- 1 MB page size 0% - 5% z10, z196
- Page fix buffers 0% - 8% V8 & high IO, in use?
- Release deallocate 0% - 15% short trans, batch
- Virtual storage constraints 0% - 5% memory, latches
- Data sharing fewer members 1% for each 2 members
- Improved dynamic SQL cache 0% - 20% literals
- Insert 0% - 40% high volume insert
- Predicate evaluation 0% - 60% complex predicates
- Access: hash, index include 0% - 5% access improved
- Increased use of zIIP 0% - 3% IO, RUNSTATS, parallel
- Utilities (from V8) 3% - 20% about same for 9 → 10
- Productivity: memory, temporal, security, admin ... priceless



Are you ready for DB2 10?

- Check prerequisites
- Contact vendors
- Migration planning workshop
- Plan gains, testing, memory, and performance
- Build detailed migration plan
- Check information APARs
- Apply required service
- Run premigration checks DSNTIJPA (or M) early and often
- Resolve incompatible changes
- Get rid of private protocol
- Convert to packages from DBRMs in plans
- Upgrade plan table formats to Unicode V8 or DB2 9 level
- Get ready for SMS
- Save performance and access path information
- Get all the parts out of the box



Free Migration Planning Workshops DB2 10 9

- Understand breadth of features in DB2 for z/OS
- Bring together a toolbox of resources for your migration planning
- Explain the current migration process
- Bring a project focus to migration
- Remain relevant through GA life of the product
 - Updated with field experiences





Improvements to installation & migration information based on customer feedback

- Pre-migration and migration checklists were added to help plan for and keep track of the migration process. The checklists include links to the pre-migration and migration steps. Access the checklists from the following locations:
 - Introduction to migration from DB2 Version 8:
http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migfromv8.htm
 - Introduction to migration from DB2 Version 9.1:
http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migfromv9.htm
- Installation and migration steps clearly define tasks to complete. Related concept and reference information are included as links from the tasks.
- Several new jobs simplify the setup and installation of DB2-supplied stored procedures and user-defined functions. The process is documented in [Installation step 20: Set up DB2-supplied routines \(optional\)](#) and [Migration step 25: Set up DB2-supplied routines \(optional\)](#).



DB2 10 for z/OS At a Glance

Performance, Scalability

- CPU reductions out-of-the-box
- Hash access to data, index include columns
- Ten times more threads per DB2 image

Availability Security

Productivity

- 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

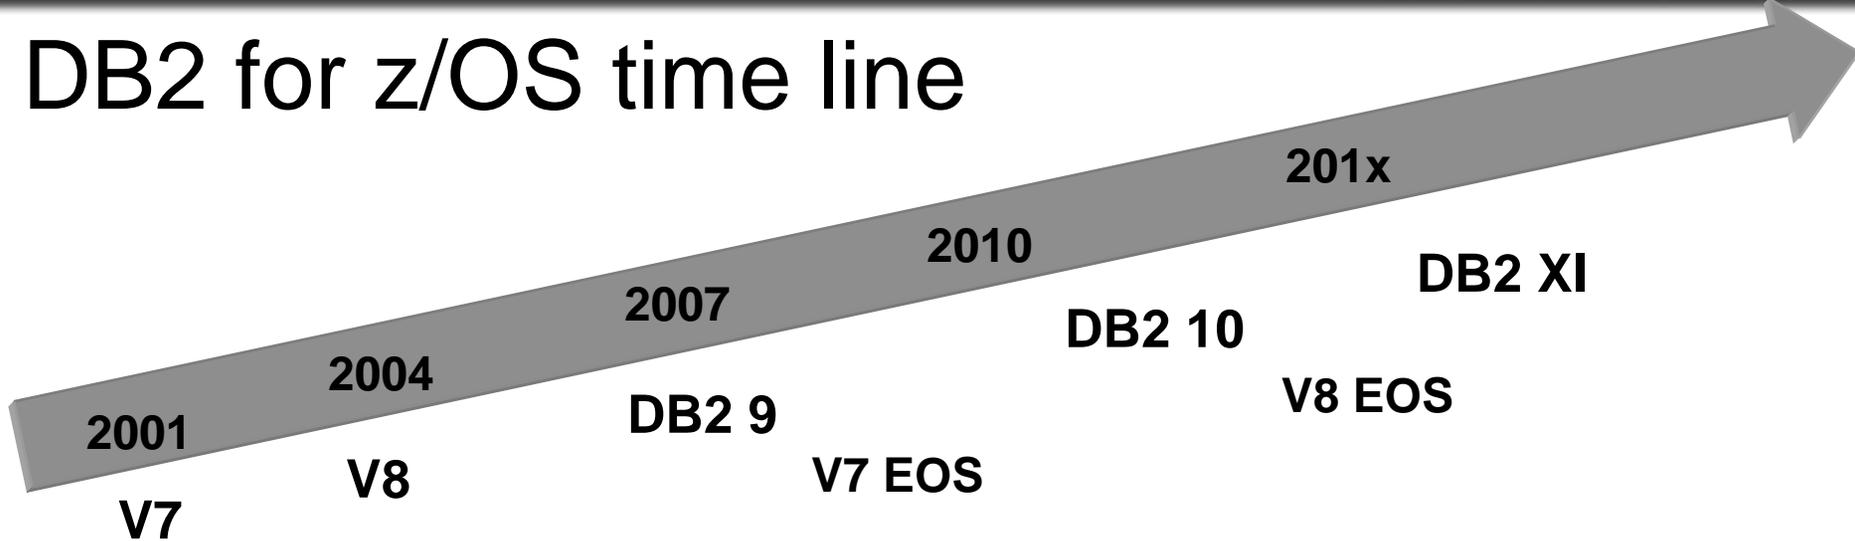
- Versioned data or temporal queries
- pureXML enhancements
- SQL improvements that simplify porting

Dynamic Warehousing

- Moving sum, moving average
- Many query optimization improvements
- Query parallelism restrictions removed



DB2 for z/OS time line



DB2 XI Focus Areas

- Continuous availability
- Performance and scalability
- Ease of management, autonomies
- SQL new function for better application performance, simplified application logic, and improved analytics

- Fast and easy migration planning and execution
- No application or DBA changes required to get DB2 XI benefits



Questions?



DB2 9 and 10 **IBM Redbooks** Publications

1. DB2 10 Technical Overview SG24-7892 updated
2. Extremely pureXML DB2 10 & 9 SG24-7915
3. DB2 10 Performance Topics SG24-7942 coming soon
4. DB2 10 Security coming
5. DB2 9 Technical Overview SG24-7330
6. DB2 9 Performance Topics SG24-7473
7. DB2 9 Stored Procedures SG24-7604
8. Serialization and Concurrency SG24-4725-01
9. Distributed Functions SG24-6952
10. Utilities SG24-6289-01
11. DB2 and Storage Management, SG24-7823
12. Index Compression with DB2 9 for z/OS redp4345
13. SQL Reference for Cross-Platform Development
14. Enterprise Database Warehouse, SG24-7637
15. 50 TB Data Warehouse on System z, SG24-7674
16. LOBs with DB2 for z/OS SG24-7270
17. Deploying SOA Solutions SG24-7663
18. Enhancing SAP - DB2 9 SG24-7239
19. Best practices SAP BI - DB2 9 SG24-6489-01
20. Data Sharing in a Nutshell, SG24-7322
21. Securing DB2 & MLS z/OS SG24-6480-01
22. Data Sharing: Dist Load Balancing & config. redp4449
23. Packages Revisited, SG24-7688
24. Ready to Access Solid-State Drives redp4537
25. Buffer Pool Monitoring & Tuning redp4604
26. Securing & Auditing Data SG24-7720



More information and resources

- **DB2 main web page**
<http://www.ibm.com/software/data/db2/zos/>
- **DB2 10 web page**
<http://www.ibm.com/software/data/db2/zos/db2-10/>
- **DB2 books, Information Center**
 - <http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656>
 - <http://publib.boulder.ibm.com/infocenter/imzic>
- **DB2 best practices web page**
 - <https://www.ibm.com/developerworks/data/bestpractices/db2zos/>
- **DB2 for z/OS IBM Redbooks publications**
 - <http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=db2&SearchOrder=4&SearchFuzzy=>
- **DB2 presentations**
<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/>

Information Management **IBM**
DB2 10 for z/OS
Proven, simplified and cost effective

**DB2 10 for z/OS
Technical Overview**

Explore the new system and application functions
Obtain information about expected performance improvements
Decide how and when to migrate to DB2 10

Information Management **IBM**
Draft Document for Review January 9, 2011 1:25 pm
SG24-7915-00

**Extremely pureXML
in DB2 10 for z/OS**

**Triton
CONSULTING**

Information Management **IBM**
Draft Document for Review February 25, 2011 11:01 am
SG24-7942-00

**DB2 10 for z/OS
Performance Topics**

base for a Smarter Planet

Discover the functions that provide reduced CPU time in CM and NFM



(cont) DB2 10 Resources and Contacts

SAP Whitepapers DB2 10 for z/OS is certified for SAP NetWeaver 7.30 and SAP R/3 4.6

- **SAP article on DB2 10 (published by SAP)** <http://www.sdn.sap.com/irj/sdn/db2>
- **SAP Best Practice Guide for Migrating to DB2 10 for z/OS (published by SAP)** <https://websmp207.sap-ag.de/~sapidb/011000358700001414122010E>
- **(Updated) Business Continuity Guide for Running SAP on System z – based on DB2 10 for z/OS, DB2 Connect 9.7 FP3a, SAP NetWeaver 7.10 and Tivoli Automation for z/OS V3.3** <http://publibfp.dhe.ibm.com/epubs/pdf/iapacs03.pdf>
- **DB2 10 for z/OS with SAP on IBM System z Performance Report – new techdocs white paper** <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101845>
- **DB2 10 for z/OS – Optimized for SAP –** <http://cattail.boulder.ibm.com/cattail/?source=s#view=andreas.r.mueller@de.ibm.com/files/3198290001883DDBA202FBE4093F23B6>
- **SAP on DB2 10 for z/OS - Being More Productive, Reducing Costs and Improving Performance –** <http://www.sdn.sap.com/irj/sdn/db2?rid=/library/uuid/005c6b33-aaf0-2d10-fcbb-b42e89ac5791>