DB2 10 for z/OS and Beyond

Session Number 1110

Jeff Josten, IBM
Michael Perera, IBM

IBM Software
Information On Demand 2011

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.

Agenda

- DB2 10 status update
- DB2 10 technical overview
- Customer experiences
- A look into the future

DB2 for z/OS Technical Strategy

- Continuous availability, RAS leadership
- Performance and scalability
- Ease of management / autonomics
- Advanced application features

- Low cost of computing
- OLTP, batch, query, mixed workloads

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 end of service April 2012

DB2 10 for z/OS

- **Fastest uptake**
  - **3X** more customers
  - **4X** more licenses
  - **25%** migrating from V8
- **First customers in production**
  - Banking leading the way
    - SAP, data warehouse and OLTP workloads
    - Skip-level and V2V
  - Quality/stability looking good
    - **Next migration wave in 2H2011**
  - **22%** surveyed plan production by YE2011

Who has DB2 10?

Slide intentionally left blank

…. and many more

What Customers are Saying

- **Mike, it's important that you ….**
  - Maintain service stream stability
    - Appropriately flag HIPERs
    - Minimize PEs
  - Simplify migration
    - Provide new features without forcing application changes
  - Maintain leadership in RAS

zEnterprise 196 Benefits for DB2

Taking System z to the next level

- Faster CPUs, more CPUs, more memory means better DB2 performance, scalability
  - Excellent synergy with DB2 10 removes many single system scaling inhibitors
- Large cache benefits DB2 workloads
- TLB changes improve DB2 10 performance for 1MB page sizes
- Hybrid architecture opens up new opportunities for DB2 query performance acceleration

DB2 10 and zEnterprise 196

- Scalability with DB2 10 and z196
  - Higher DB2 CPU reduction can be achieved as the number of processors per LPAR increases
    - Best fit with DB2 10 scalability
  - More than 20% improvement with DB2 10 compared to DB2 9 on z196 64 way

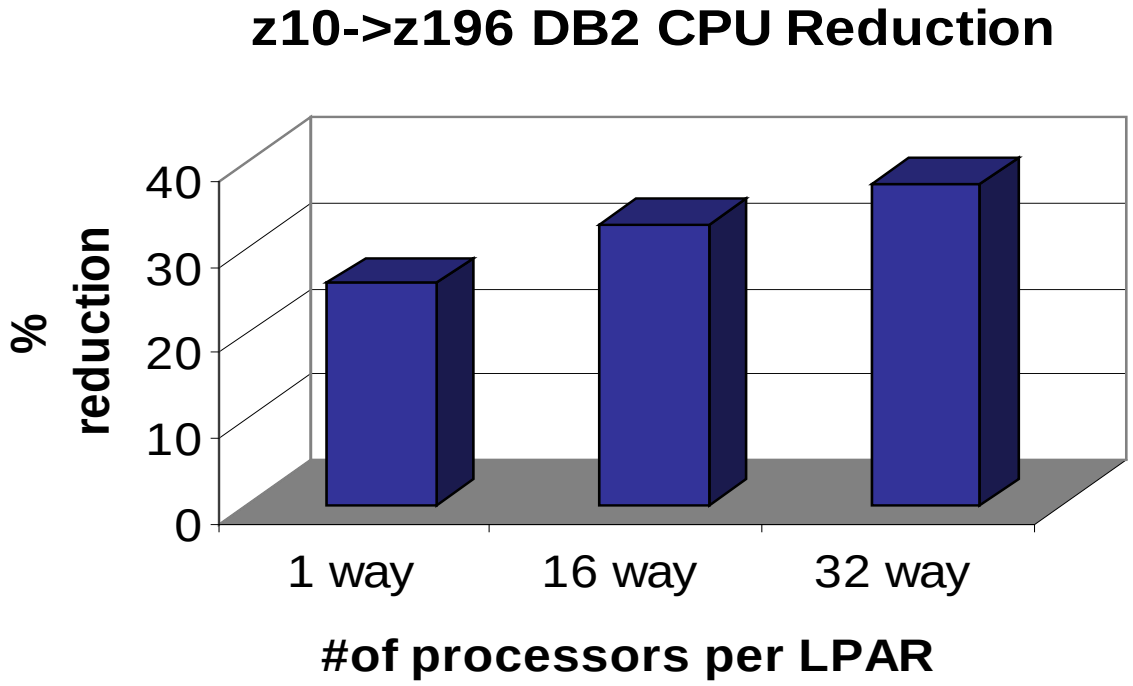

SAP Sales and Distribution Benchmark

- SAP SD Benchmark
  - z10 and z196 with 4 processors per LPAR
  - Enabled 1 MB large frames
  - 76% improvement in ITR when comparing the results of running the SAP SD workload with DB2 9 on z196 (almost a factor of 1.8x)

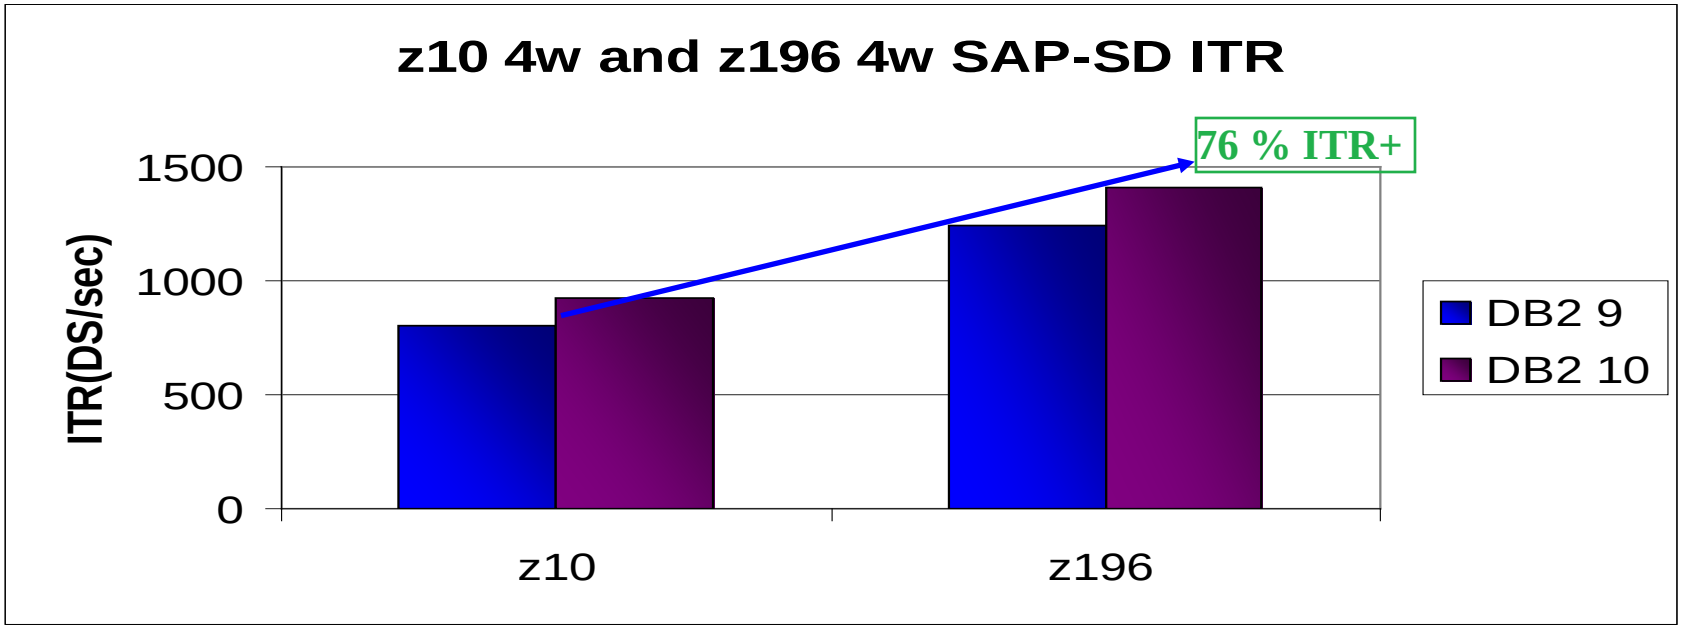

Mainframe Innovation
Specialty Engines

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

System z Application Assist Processor (zAAP) 2004

IBM System z Integrated Information Processor and (2006)

Internal Coupling Facility (ICF) 1997

Integrated Facility for Linux® (IFL) 2000

Eligible for zAAP:

- Java execution environment
- z/OS XML System Services

* 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**, **DB2 10**)*:

1. **DRDA over TCP/IP connections**
   - DB2 for z/OS Remote native SQL procedures
   - 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**
   - Higher percentage of parallel queries zIIP eligible
   - More queries eligible, more parallelism
3. **DB2 Utilities LOAD, REORG & REBUILD functions used to maintain index structures and sort**
   - RUNSTATS options other than column group, inline
4. **Buffer pool prefetch and deferred write**

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

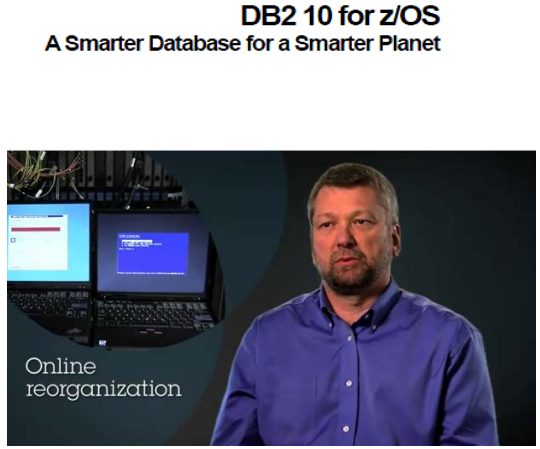

DB2 10 has it!

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

Reduced Costs	Simplified Workloads	Proven Technology
"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." Morgan Stanley DB2 Team	**"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."** Paulo Sahadi - Senior Production Manager, Information Management Division at Banco do Brasil	**"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."** Philipp Nowak, BMW DB2 Product Manager
"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."  Large Global Bank	*"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."*** Frank Petersen – System Programmer Bankdata	**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"** Guenter Schinkel -Postbank Systems AG

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
  - Up to 5-10% for traditional workloads
  - Up to 20% for new workloads
  - Up to an additional 10% 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 added 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 for small LOBS and Complex Queries

DB2 10 Performance

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

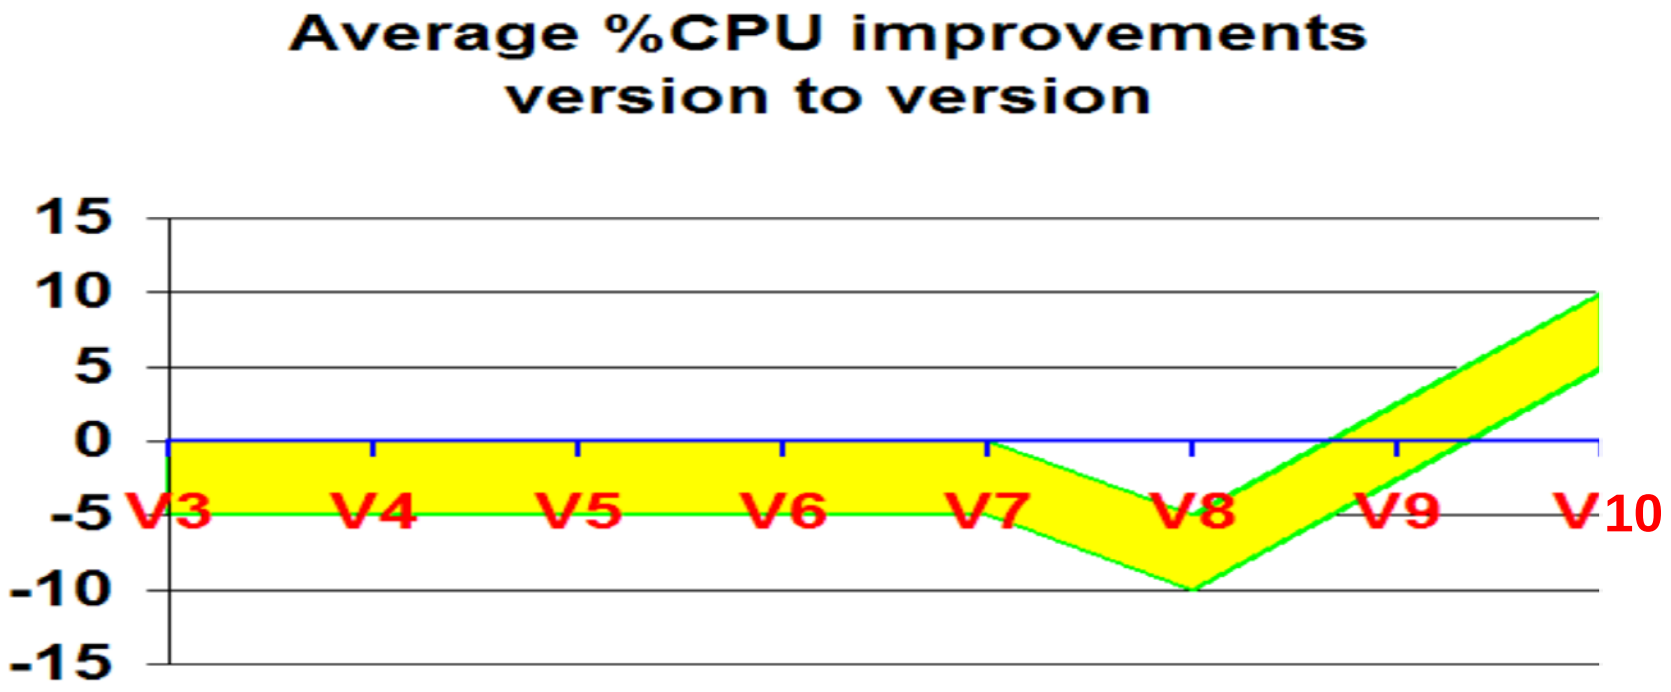

Performance Enhancements Requiring Few Changes (CM)

- 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 – requires PGFIX(YES)
  - "Fully in memory" option (ALTER BUFFERPOOL): PGSTEAL(NONE)

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

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 perfromance 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

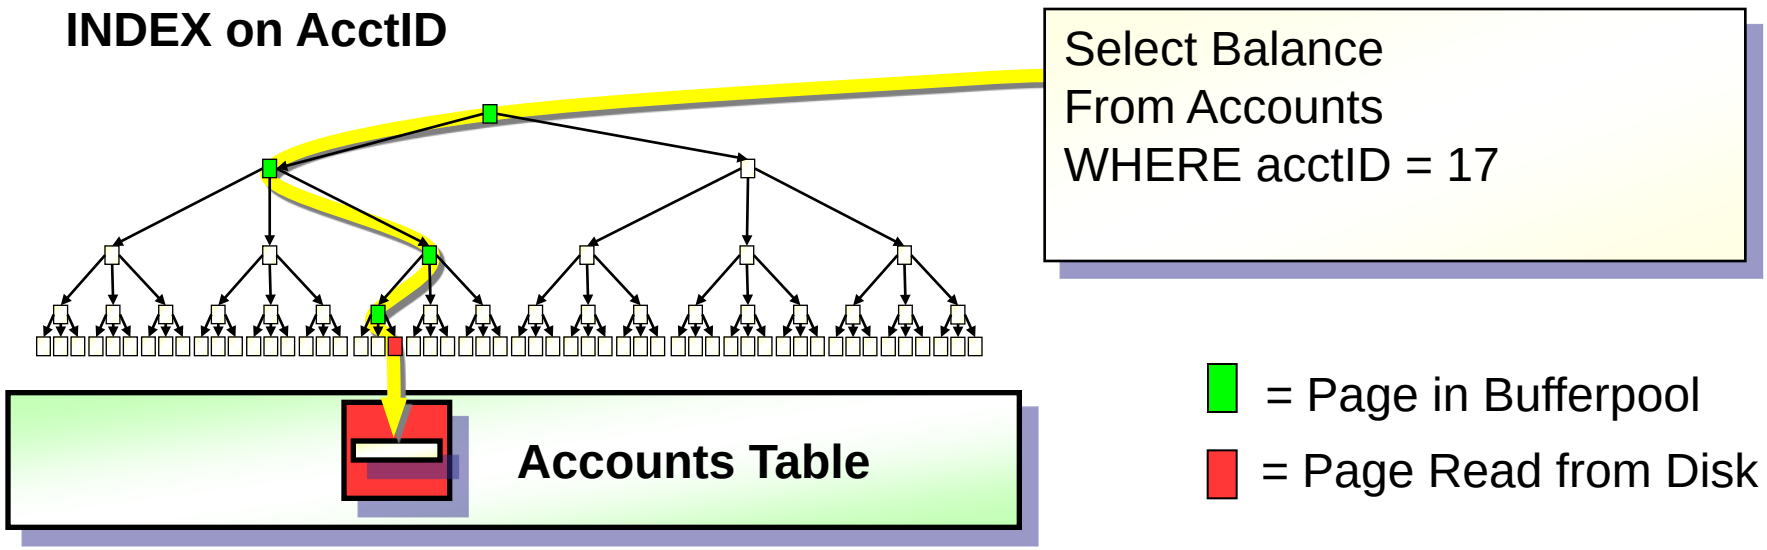

- **Traverse down Index Tree**
  - Typically non-leaf portion of tree in the bufferpool
  - Leaf Portion of the tree requires I/O
  - Requires searching pages at each level of the index
- **Access the Data Page**
  - Typically requires another I/O
- **For a 5 Level Index**
  - 6 GETP/RELPs, 2 I/O's,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

Hash Access

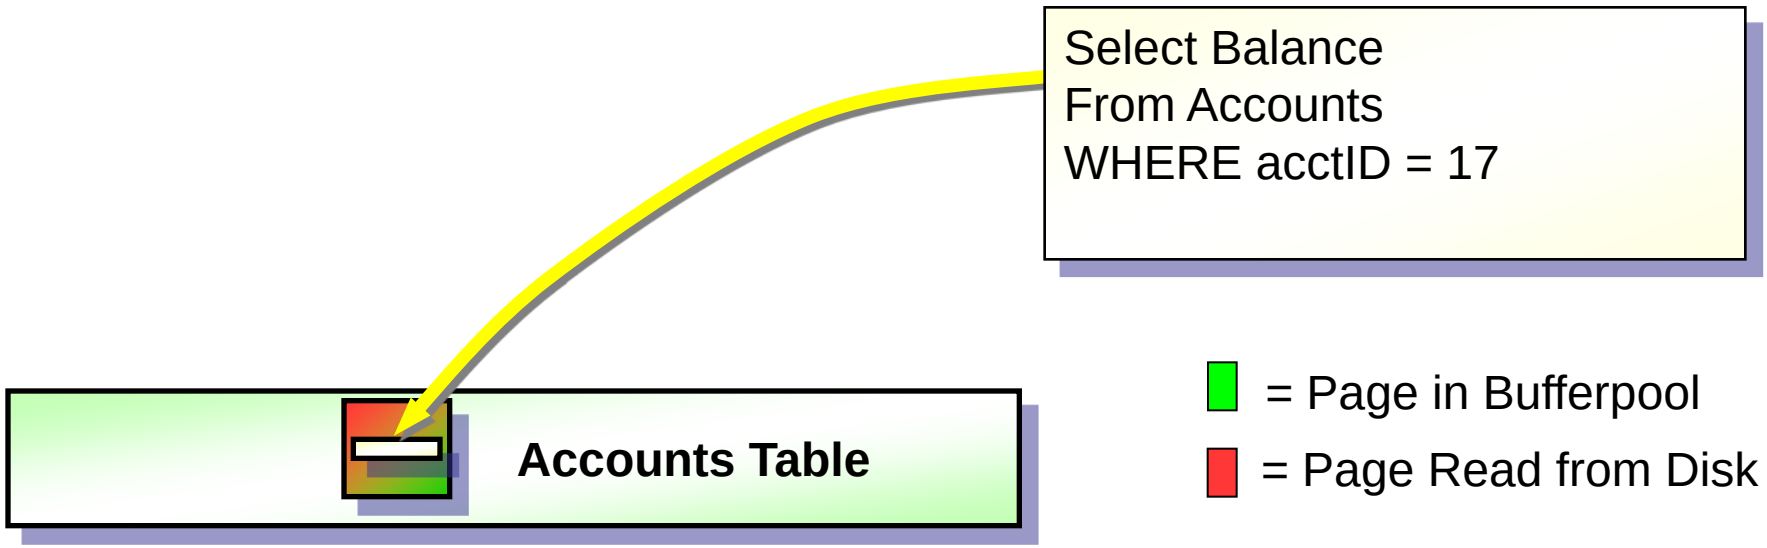

- Hash Access provides the ability to directly locate a row in a table without having to use an index
- Single GETP/RELP in most cases
- 1 Synch I/Os in common case
  - 0 If In Memory Table
- Greatly reduced Search CPU expense

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

Beta Customer Performance Feedback

Workload	Customer Results
CICS online transactions	Approx. 7% CPU reduction in DB2 10 CM after REBIND, additional 4% reduction when 1MB page frames are used for selective buffer pools
CICS online transactions	Approx 10% CPU reduction from DB2 9
CICS online transactions	Approx 5% CPU reduction from DB2 V8
CICS online transactions	10+% CPU increase (APAR was shipped to address)
Distributed Concurrent Insert	50% DB2 elapsed time reduction, 15% chargeable CPU reduction after enabling high performance DBAT
Data sharing heavy concurrent insert	38% CPU reduction
Queries	Average CPU reduction 28% from V8 to DB2 10 NFM
Batch	Overall 20-25% CPU reduction after rebind packages

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	20-30% CPU reduction in random access 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.

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 KEEPDYNAMIC=YES

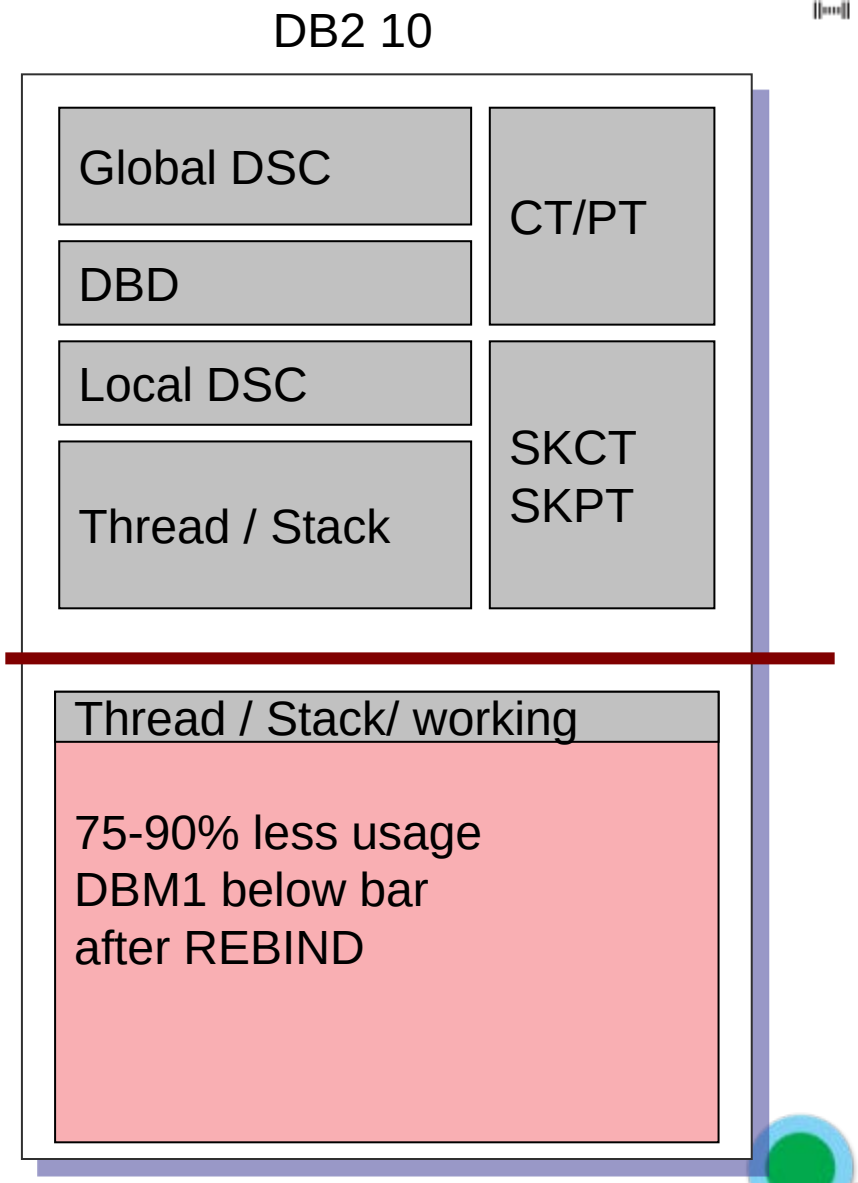

Running Many Active Threads

Today

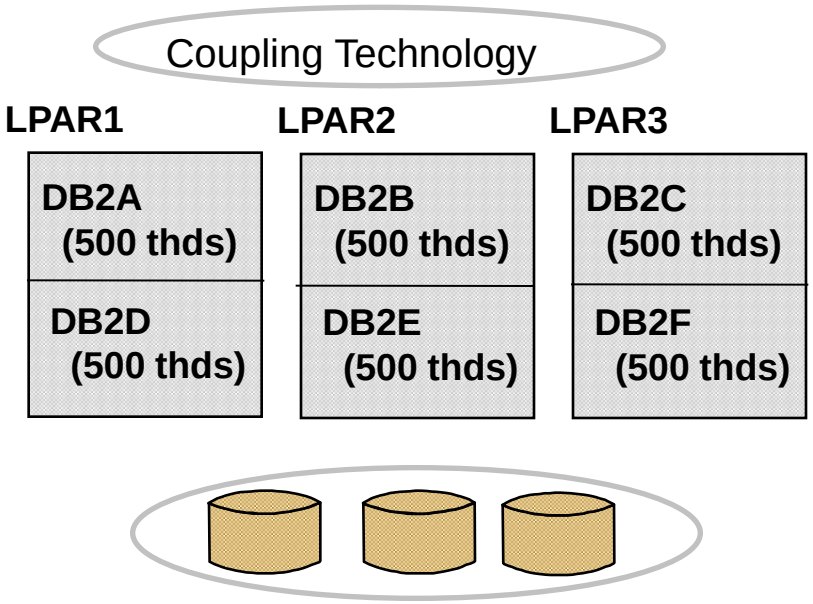

- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s per LPAR

DB2 10

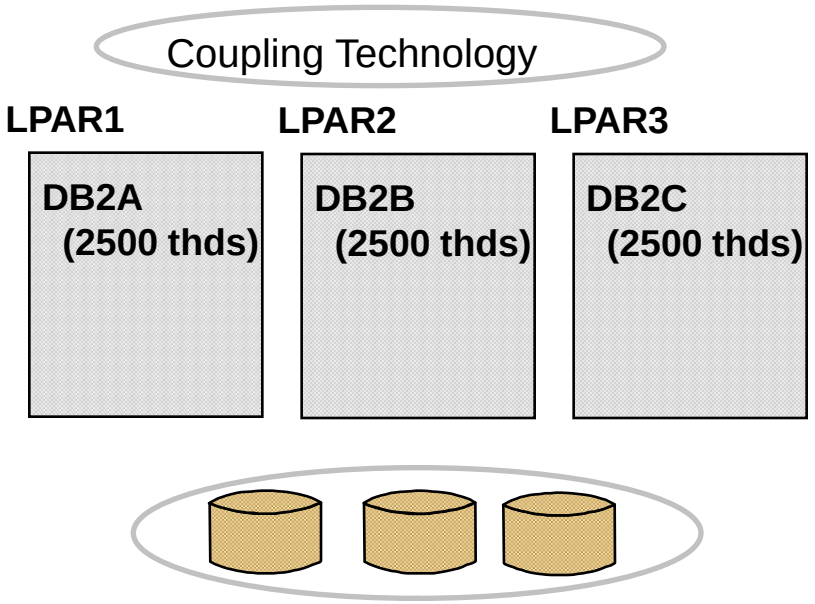

- More threads per DB2 image
- More efficient use of large n-ways
- Easier growth, lower costs, easier management
- Data sharing & Parallel Sysplex still required for very high availability and scale
- Rule of thumb: save ½% CPU for each member reduced, more on memory

Performance and Scalability Summary

- Many traditional OLTP workloads see 5-10% CPU reduction in CM mode after (some more, some less)
- Prerequisites
  - REBIND packages to
    - Generate new SQL run time and get maximum VSCR
    - Gain performance benefits that are conditional on REBIND
  - Use PGFIX=YES on buffer pools to exploit 1MB real storage frames available on z10 and z196 (100% backed)
- There were some exceptions < 5% CPU savings for OLTP with very light transaction, skinny packages with few simple SQL
  - Package allocation cost overrides benefit from SQL optimizations
  - APAR PM31614 helps to solve this by improving package allocation performance
  - Use of persistent threads with RELEASE(DEALLOCATE) will compensate

Performance and Scalability Summary…

- DBM1 Virtual Storage Constraint Relief with 64-bit SQL run time:
  - Available in CM
  - Requirement to REBIND static SQL packages to accrue maximum benefit
  - Very good results achieved
  - Have high degree of confidence that problem addressed
    - Real world proposition: 500 -> 2500-3000+ threads
  - Limiting factors now on vertical scalability (# number of threads per DB2 image)
    - Amount of real storage provisioned
    - ESQA/ECSA (31-bit) storage
    - Active log write

DB2 10 Virtual Storage Relief Enables Other Performance Options

- After REBIND in DB2 10, most thread storage moves above the bar (CM or NFM)
- Major customer opportunities for 31-bit VSCR and improved price/performance
  - Potential to reduce legacy OLTP transaction CPU cost through use of
    - More CICS protected ENTRY (persistent) threads
    - More use of RELEASE(DEALLOCATE) with persistent threads
    - Must consider additional real storage to back the requirement
  - Potential to reduce CPU for DRDA transactions by using High Performance DBAT
    - Must be using CMTSTAT=INACTIVE so that threads can be pooled and reused
    - Packages must be bound with RELEASE(DEALLOCATE) to get reuse for same connection
    - MODIFY DDF PKGREL(BNDOPT) must also be in effect
    - Do not to overuse RELEASE(DEALLOCATE) on packages
      - Will drive up the MAXDBAT requirement
      - Will need additional real storage to support increased number of threads
  - Some tradeoff with BIND/REBIND and DDL concurrency

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

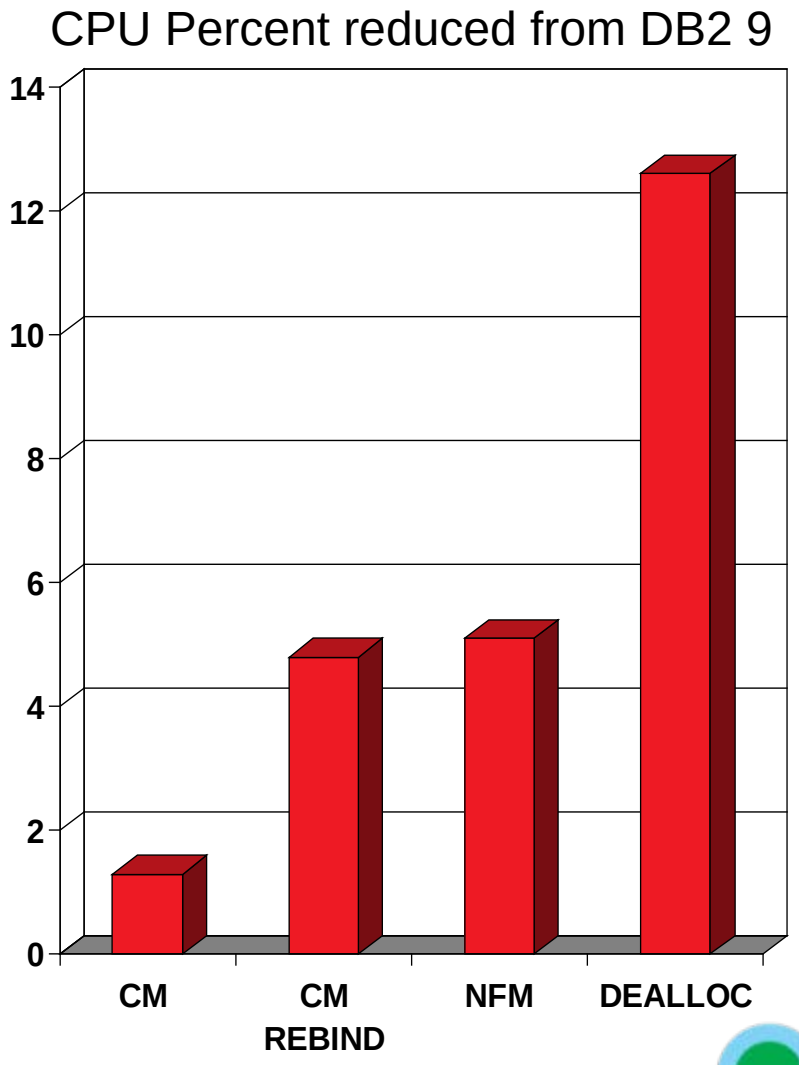

DB2 10 and Real Storage Usage

- DB2 10 will likely use incrementally more real memory
  - Lab measurements showed about a 10% increase
- 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

Use of z/OS 1MB Page Frames on z10 or z196

- Potential for reduced for CPU through less TLB misses
  - Early customer experience: 0-6% reduced CPU
- Buffer pools must be defined as PGFIX=YES
- Buffer pool page fix introduced in V8 to reduce CPU
- Must partition real storage between 4K frames and 1MB frames
  - Specified by LFAREA xx% in IESYSnn parmlib member and only changeable by IPL
  - 1MB frames are non-pageable
  - If 1MB page frames are overcommitted, will use 4K page frames
  - Recommendation: to add 20% in size to allow for growth and tuning

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
  - 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

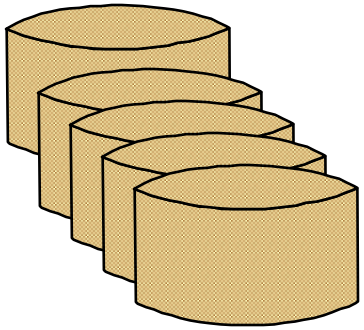

Continuous Availability: More Online Schema Changes

- Universal Table Space (UTS) type was introduced in DB2 9
  - Two types: Partition By Growth (PBG) and Partition By Range (PBR)
  - Combined advantages of the old segmented and classic partitioned table spaces
  - UTS is the strategic table space type going forward, many new functions are enabled for UTS only
  - But DB2 9 gives no easy way to convert existing TS's to UTS – unload/drop/recreate/reload required
  - DB2 10 allows for ALTER + Online REORG – easy, non disruptive
  - DB2 10 allows MEMBER CLUSTER for UTS
- DB2 10: Convert to Universal Table Space (UTS) with no outage
  - Single table simple -> UTS/PBG
  - Single table segmented -> UTS/PBG
  - Classic partitioned -> UTS/PBR
- DB2 10: Online ALTER of other UTS attributes
  - DSSIZE (also for XML or LOB table spaces)
  - page size (also for LOB table spaces)
  - SEGSIZE (also for segmented or XML table spaces)
- DB2 10: Alter index page size for indexes on UTS's
- DB2 10: Convert PBG/PBR -> Hash
- DB2 10: 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 and 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 DB2 9
    - 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

Business Security and 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

Row and Column Access Control

- A table level authorization function fully implemented in the database that provides:
  - Row level access control based on customer-supplied rules
    - A doctor can see rows representing his patients only
    - A manager can see rows representing his employees only
  - Column level access control based on customer-supplied rules
    - This is also called *data masking*
    - A teller can see only the last 4 digits of the credit card number column
- Defined and managed by SECADM

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
  - New RTS columns: cluster sensitivity for queries
- 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

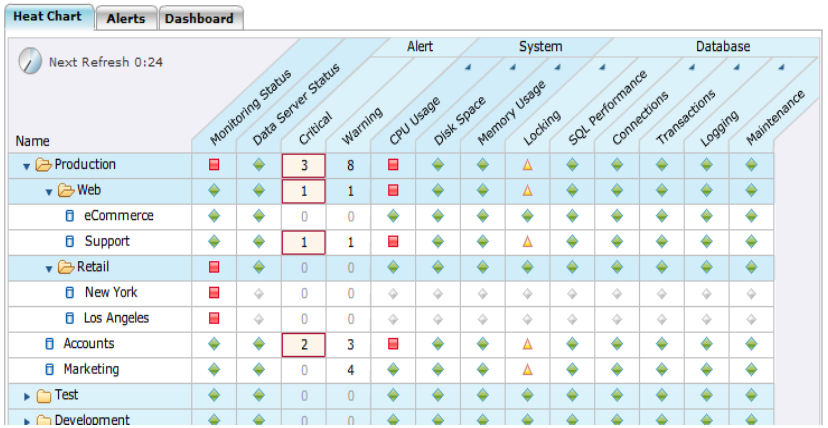

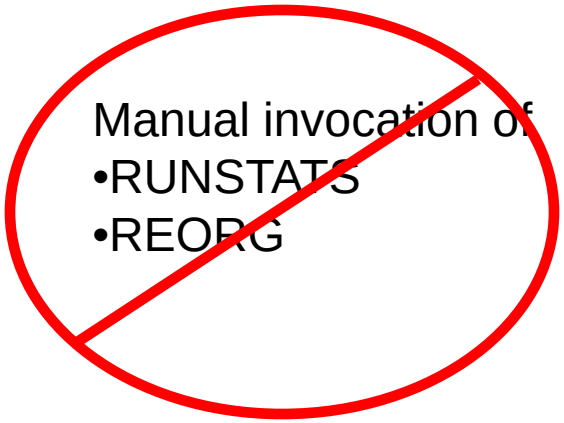

DB2 10 Connection Profile

- Problem: For distributed workloads, low priority or poorly behaving client applications may monopolize DB2 resources and prevent high-priority applications from executing.
- Solution: Increased granularity of monitoring for system level activities
  - Number of connections
  - Number of threads
  - Idle thread timeout
- Profiles specified in SYSIBM.DSN_PROFILE_TABLE
- DB2 adds support for filtering and threshold monitoring of system related activities via new keywords
  - Number of threads
  - Number of connections
  - Idle thread timeout
- New scope filters
  - ROLE (available through Trusted Context support)
  - Product-specific identifier

DW and OLTP Co-residency in a Parallel Sysplex

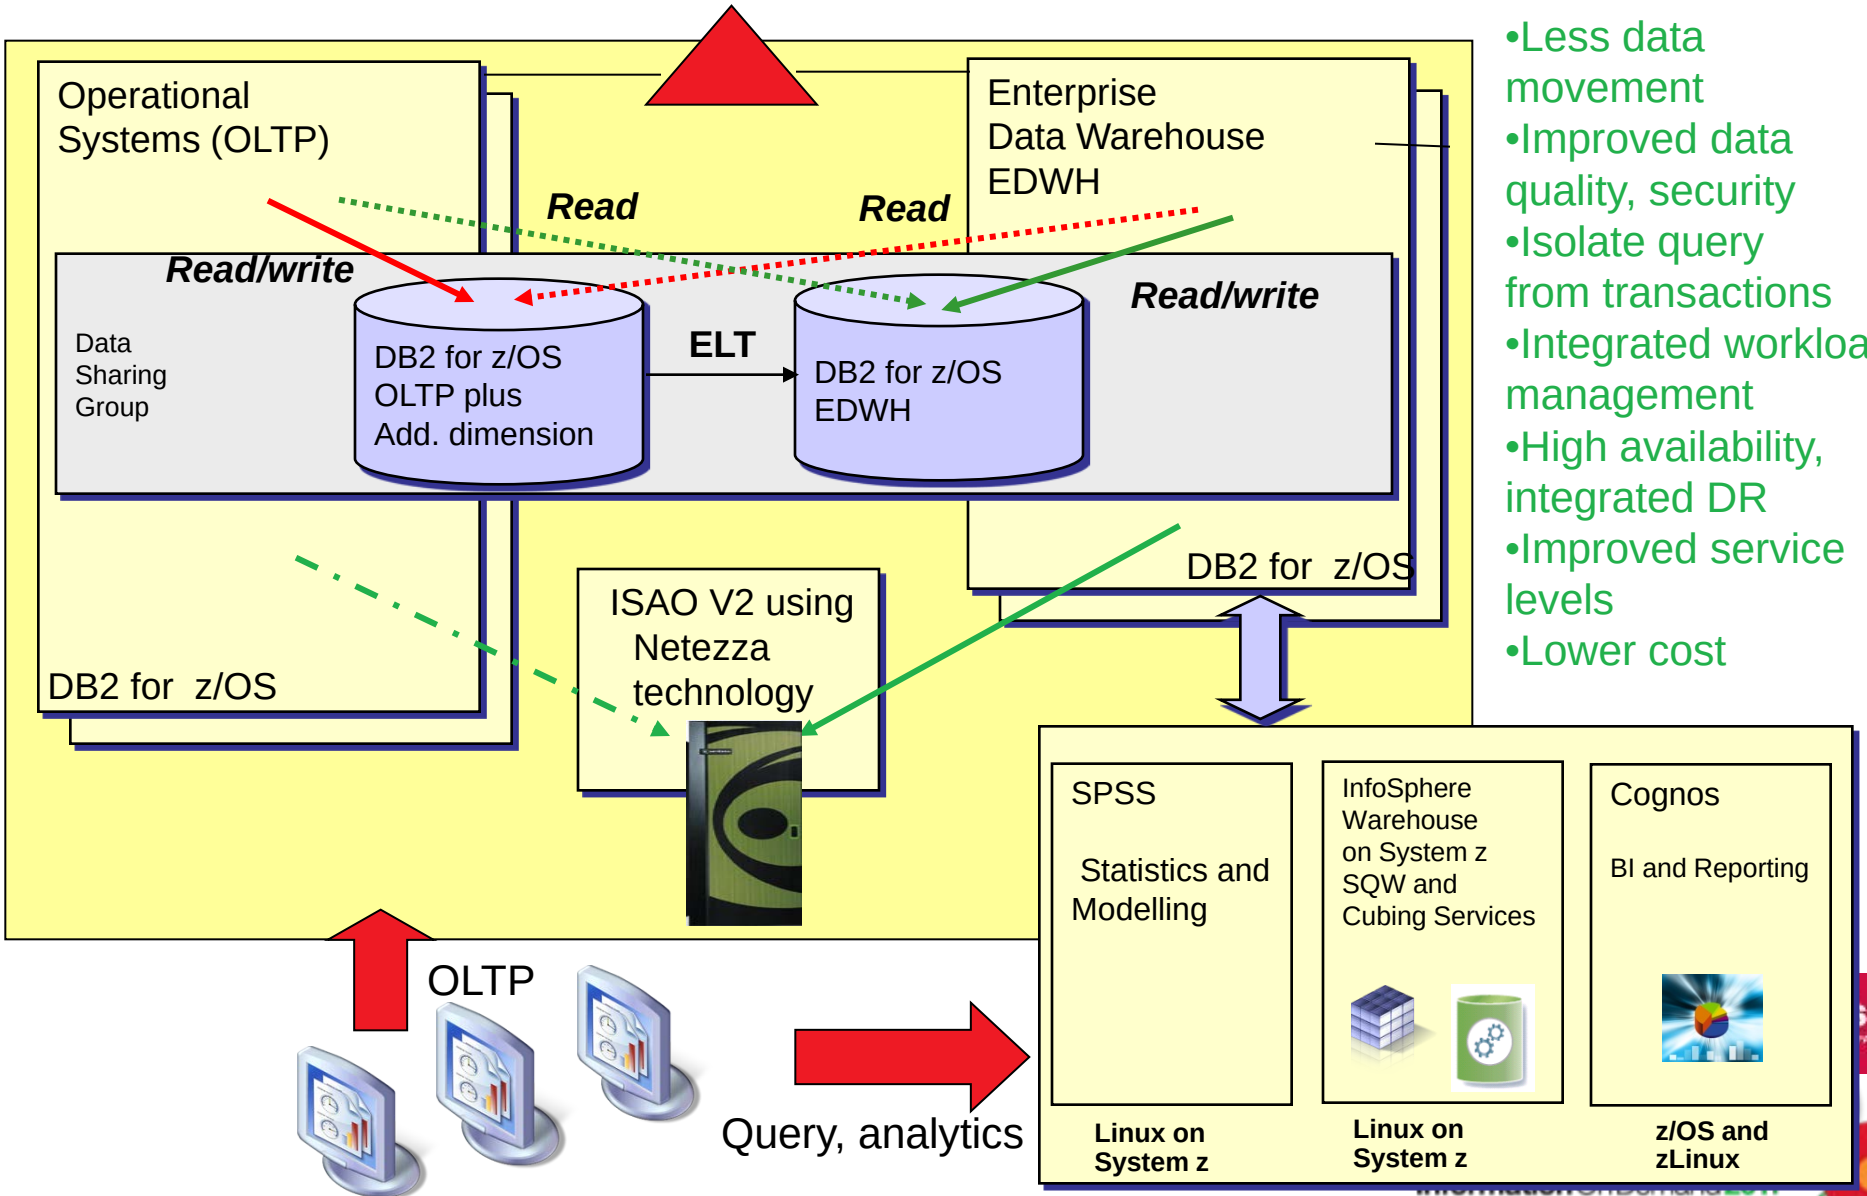

- Less data movement
- Improved data quality, security
- Isolate query from transactions
- Integrated workload management
- High availability, integrated DR
- Improved service levels
- Lower cost

IBM Smart Analytics Optimizer

Capitalizing on the best of relational and columnar databases

Data Warehousing

What is it?

The IBM Smart Analytics Optimizer is a workload optimized, appliance-like 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 accelerator
- Simplified administration: appliance-like hands-free operations, eliminating many database tuning tasks

Breakthrough technology enabling new opportunities

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

Optimizing to the Workload

Total solution remains centrally managed by System z…

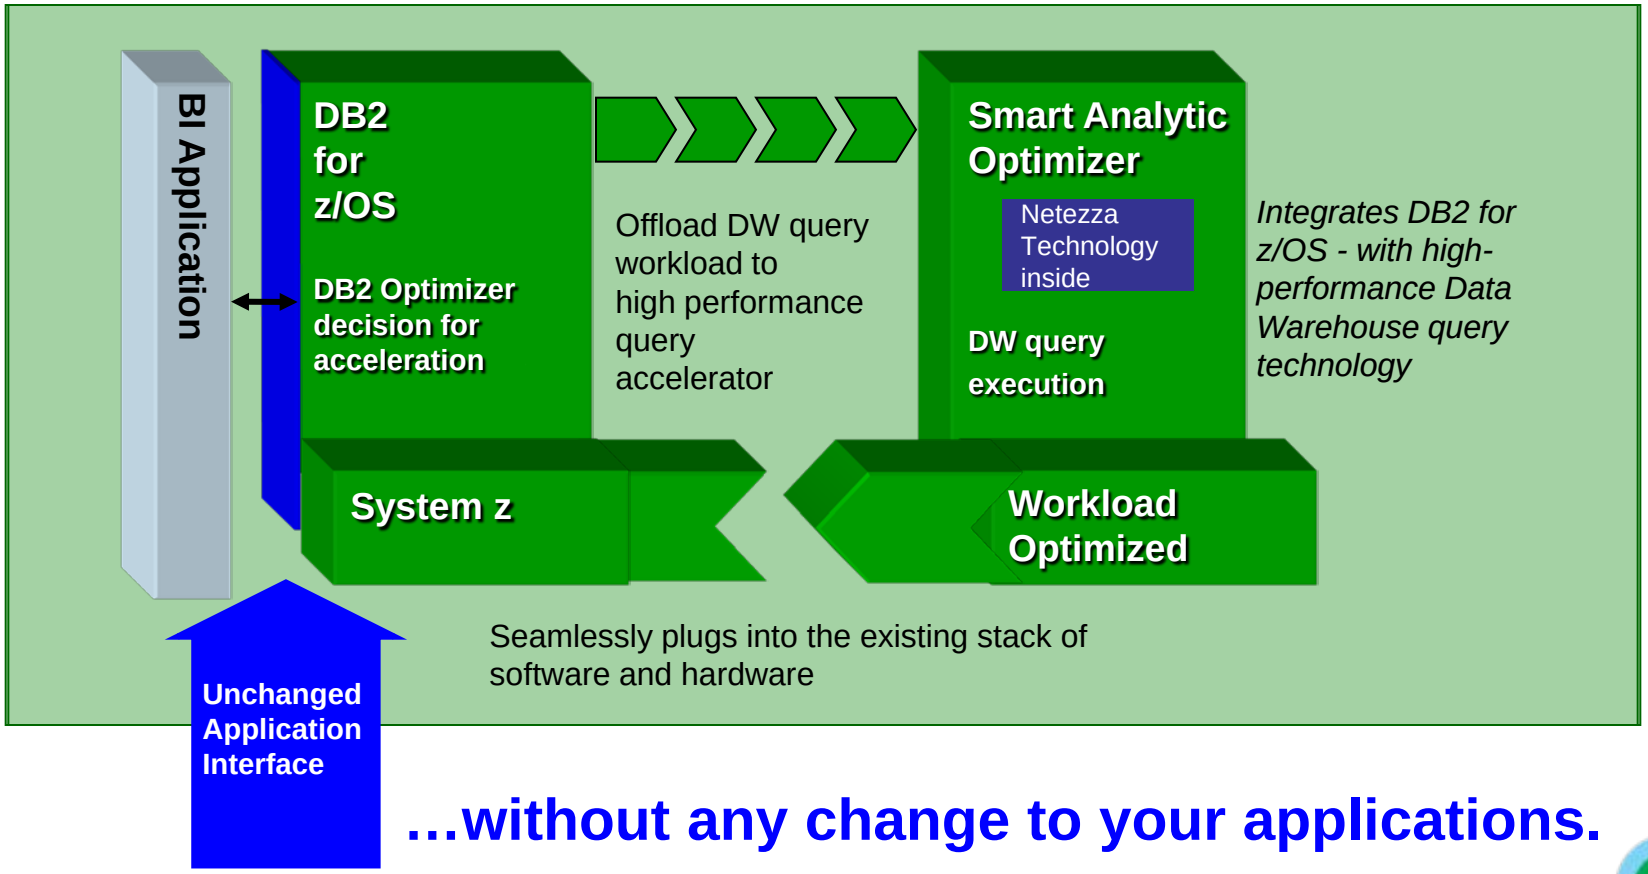

…without any change to your applications.

Key Enhancements in ISAO V2

- 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

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

DB2 10 Application Enablement & 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
  - Extended indicator variables

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

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

DB2 10 for z/OS Skip-Level Migration

- Save migration cost & effort
  - Leverage the benefits of both DB2 9 & 10 with a single migration
  - Migration, fallback and data sharing coexistence fully supported
  - Mix of DB2 9 and 10 or DB2 V8 and 10
  - Migration considerations for both versions still apply
- Migration best practices still apply
  - Are your critical apps ready?
  - Do you have a good test plan?
  - Are you following maintenance best practices?
  - Are your ISVs ready?
- Migration Planning Workshops & DB2 Advisors can help you build a successful migration plan

DB2 10 News
Items planned for post GA delivery, based on customer input

- APREUSE, APCOMPARE – PM25679
  - V9 or V10 bound package required as a base starting point
- Delete data sharing member - PM42528
- 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
- Relief for incompatible change in CHAR of decimal data - PM29124
- Real storage monitoring enhancements – PM24723
- DSSIZE > 64GB – PM42175
- Authorization changes when there is no private protocol – PM37300

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 II 14477 (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)

Free Migration Planning Workshops (MPW) DB2 9 & 10

- 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

DB2 10 Summary

- Many opportunities for price/performance (cost) improvements
  - Major theme of this release
  - Most welcome by our customers
- Significant DBM1 31-bit VSCR after rebind
  - Opportunity for scale up and LPAR/DB2 consolidation
- REBIND required to obtain most performance and VSCR improvements
- Plan, provision, and monitor real storage consumption

DB2 for z/OS Timeline

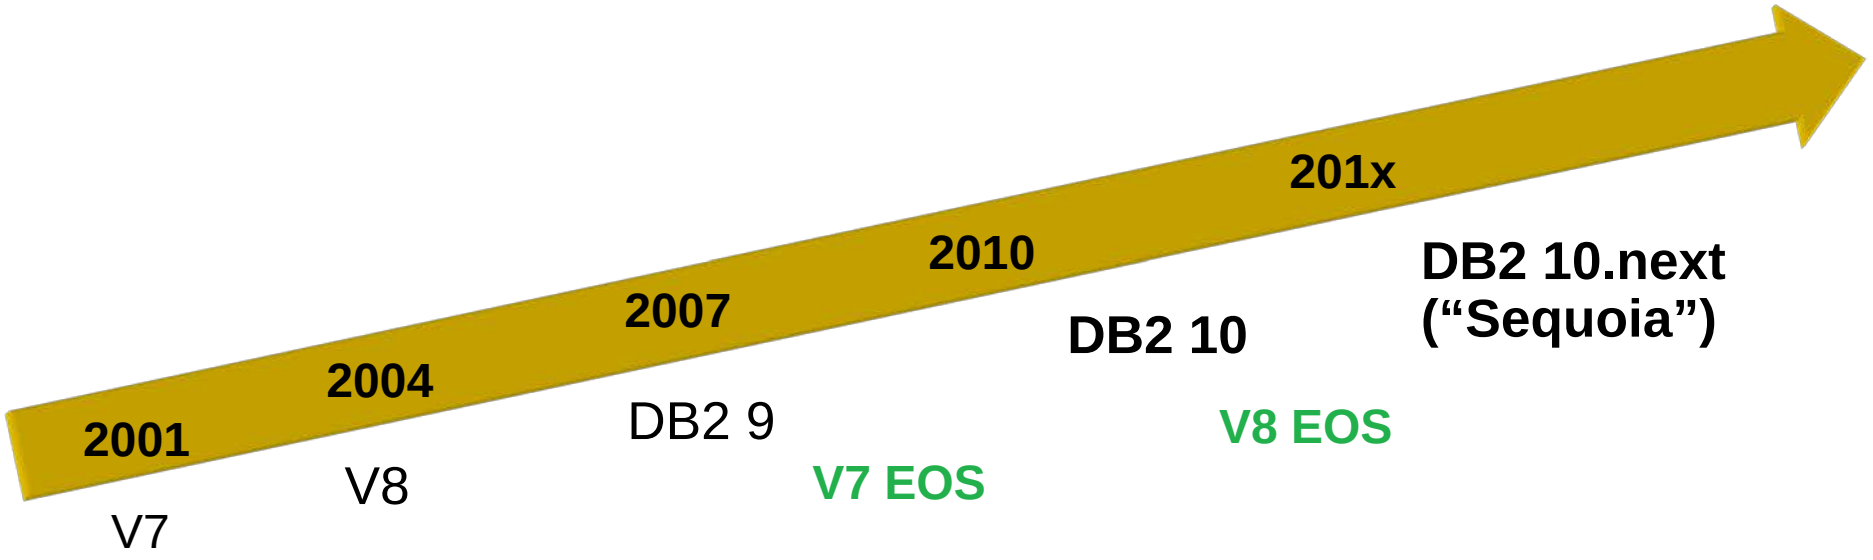

DB2 V8 End of Service effective April 30, 2011

DB2 10.next – “Sequoia”
Some early thoughts

- **Unmatched reliability, scalability, and availability**
  - Improved Data Sharing performance and efficiency
  - Even less downtime by removing growth limitations
  - Simplify management, reduce planned outages with more online schema changes and utilities improvements
  - Expanded logging capacity and scalability
- **Simpler, faster migration**
  - No application changes needed to account for SQL interface differences
  - Access path stability improvements
  - Better application performance with SQL and XML enhancements
  - Simplify application with expanded temporal and SQL PL
- **Enhanced business analytics**
  - Faster, more efficient performance for query workloads
  - Transparent archiving
  - More optimization for ISAO.next
  - More efficient inline database scoring enabling predictive analytics
- **Save money, Save time**
  - Reduce CPU usage
    - OLTP and heavy INSERT workloads
    - Key query workloads
    - SAP day posting in data sharing
    - Additional utilities performance and CPU improvements
  - Save time and resources with new autonomic and application development capabilities

QUESTIONS?

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_intro2migrfromv8.htm
  - Introduction to migration from DB2 Version 9.1:

    http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migrfromv9.htm

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_intro2migrfromv8.htm
  - Introduction to migration from DB2 Version 9.1:

    http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migrfromv9.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).

Expanding DB2 for z/OS ISV Community

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

DB2 10 Resources

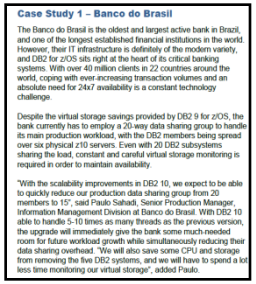

- **Website** http://www.ibm.com/software/data/db2/zos/db2-10/
  - Case Studies, Customer statements
  - Demos: DB2 10 for z/OS, QMF 10
  - Brochures: DB2 10 for z/OS Highlights, QMF 10 What's New

- **Presentations**
  - DB2 10's new functions – ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/
  - Overviews - ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/overview
  - Migration - ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/migration

- **Books**
  - DB2 10 for z/OS Technical Overview http://www.redbooks.ibm.com/abstracts/sg247892.html
  - DB2 10 for z/OS Performance Topics – *coming soon* http://www.redbooks.ibm.com/abstracts/sg247942.html
  - Extremely pureXML in DB2 10 for z/OS http://www.redbooks.ibm.com/abstracts/sg247915.html
  - DB2 10 for z/OS Book ftp://public.dhe.ibm.com/common/ssi/ecm/en/imm14075usen/IMM14075USEN.PDF
  - DB2 10 Security – *coming*
  - DB2 9 Technical Overview http://www.redbooks.ibm.com/abstracts/sg247330.html
  - DB2 9 Performance Topics http://www.redbooks.ibm.com/abstracts/sg247473.html
  - DB2 9 Stored Procedures http://www.redbooks.ibm.com/abstracts/sg247604.html
  - DB2 9 Resource Serialization and Concurrency Control http://www.redbooks.ibm.com/abstracts/sg244725.html
  - DB2 9 Distributed Functions http://www.redbooks.ibm.com/abstracts/sg246952.html
  - DB2 9 Utilities http://www.redbooks.ibm.com/abstracts/sg246289.html
  - DB2 and Storage Management http://www.redbooks.ibm.com/abstracts/sg247823.html
  - Index Compression with DB2 9 for z/OS http://www.redbooks.ibm.com/abstracts/redp4345.html
  - Enterprise Data Warehousing with DB2 9 for z/OS http://www.redbooks.ibm.com/abstracts/sg247637.html
  - 50 TB Data Warehouse Benchmark on IBM System z http://www.redbooks.ibm.com/abstracts/sg247674.html
  - LOBs with DB2 for z/OS http://www.redbooks.ibm.com/abstracts/sg247270.html
  - Deploying SOA Solutions http://www.redbooks.ibm.com/abstracts/sg247663.html
  - Data Sharing in a Nutshell http://www.redbooks.ibm.com/abstracts/sg247322.html

(cont) DB2 10 Resources

- **(cont) Books**
  - Securing DB2 and Implementing MLS on z/OS http://www.redbooks.ibm.com/abstracts/sg246480.html
  - DB2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant http://www.redbooks.ibm.com/abstracts/redp4449.html
  - DB2 9 for z/OS Packages Revisited http://www.redbooks.ibm.com/abstracts/sg247688.html
  - Ready to Access DB2 for z/OS Data on Solid-State Drives http://www.redbooks.ibm.com/abstracts/redp4537.html
  - DB2 9 for z/OS: Buffer Pool Monitoring and Tuning http://www.redbooks.ibm.com/abstracts/redp4604.html
  - Securing and Auditing Data on DB2 for z/OS http://www.redbooks.ibm.com/abstracts/sg247720.html

- **Whitepapers**
  - Business Value Whitepaper – Julian Stuhler, Triton Consulting: "DB2 10 for z/OS: A Smarter Database for a Smarter Planet" *http://public.dhe.ibm.com/software/data/sw-library/db2/analystreports/tritonconsulting-db210forzos-smarterdatabase.pdf*
  - A Matter of Time: Temporal Data Management *http://public.dhe.ibm.com/software/data/sw-library/db2/papers/A_Matter_of_Time_-_DB2_zOS_Temporal_Tables_-_White_Paper_v1.4.1.pdf*
  - Why DB2 for z/OS is Better than Oracle RAC https://www14.software.ibm.com/webapp/iwm/web/signup.do?lang=en_US&source=sw-infomgt&S_PKG=db2z-better-thank-oracle-rac-wp
  - zJournal article by Willy Favero http://www.mainframezone.com/z-journal

(cont) DB2 10 Resources

- SAP Whitepapers
  - 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/atsmastr.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
  - Enhancing SAP by Using Db2 9 for z/oS http://www.redbooks.ibm.com/abstracts/sg247239.html
  - Best Practices for SAP BI using DB2 9 for z/OS http://www.redbooks.ibm.com/abstracts/sg246489.html

Some Certifications so far...

- 1Q2011: SAP NetWeaver 7.30 and SAP R/3 4.6 is certified for DB2 10
- 2Q2011: All SAP products based on NetWeaver 7.00/7.01 is certified for DB2 10
- 2Q2011: PeopleSoft PeopleTools 8.50 and 8.51 is certified for DB2 10

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
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

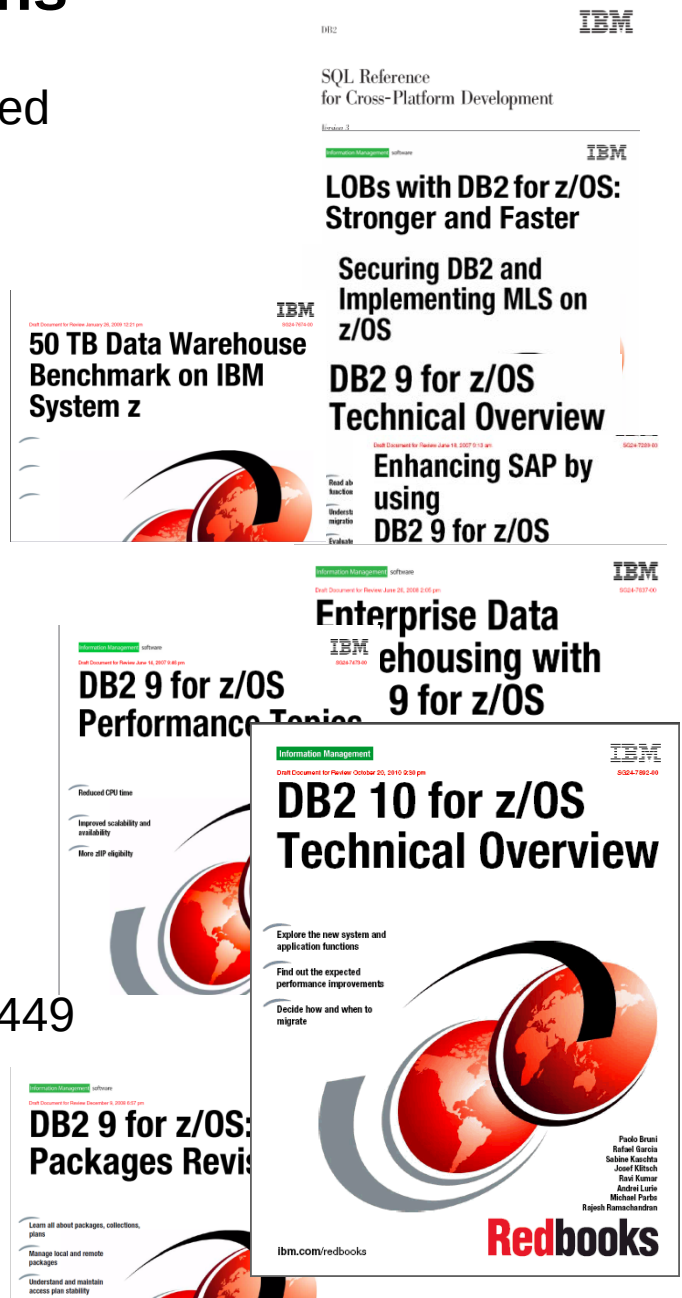

DB2 10 for z/OS At a Glance

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 and zEnterprise 196

- CPU reduction in all types of DB2 workloads
  - Larger processor cache (1.5MB L2 per core, 24MB L3 per chip, 129MB L4)
  - Various types of DB2 9 and 10 workloads show 20% to 40% DB2 CPU reduction compared to z10 processors.

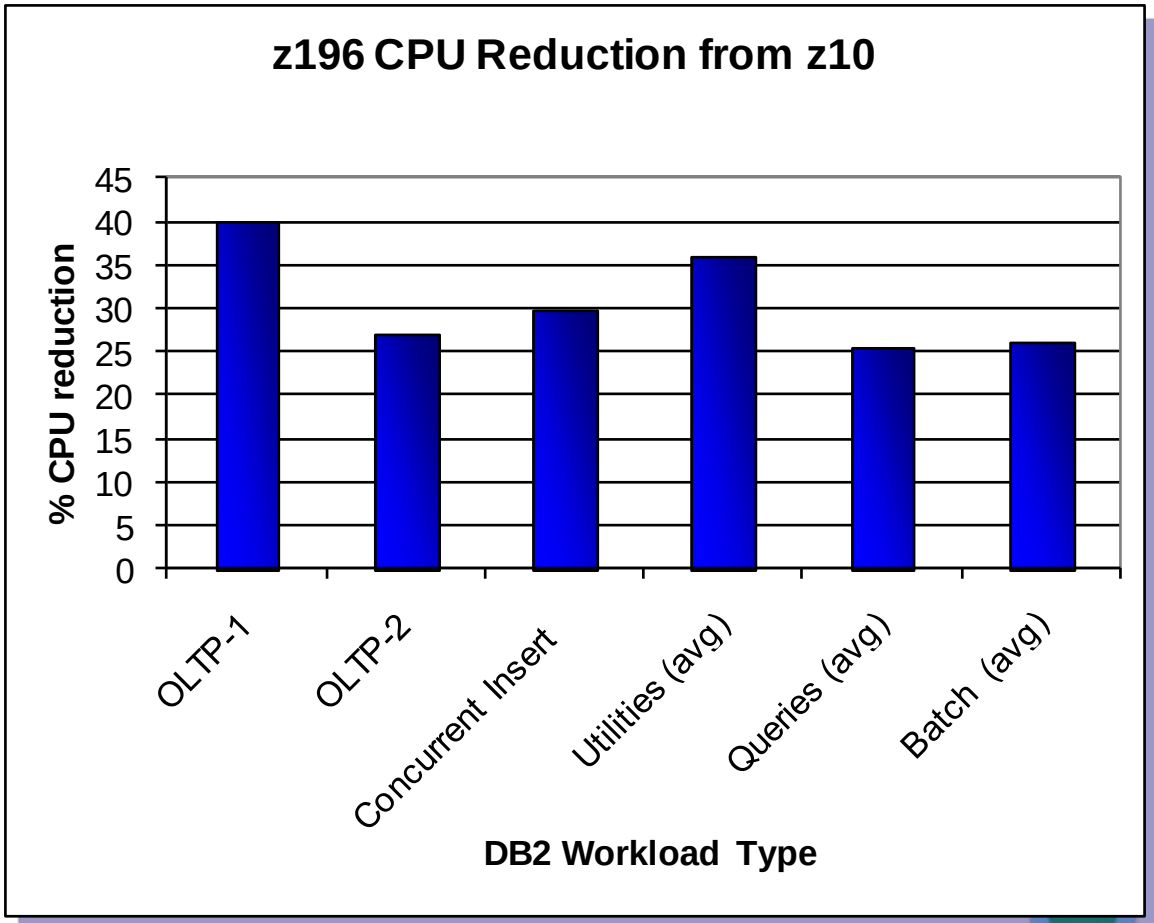

NOTE: zHPF performance summary
Hardware environment: DS8800, FICON Express 8, z196

- Format write throughput increases up to 52% (4K pages)
  - Applies to Load, Reorg writes to shadow, Restores
  - DB2 V8 increases only 30%. Higher gains requires VPSIZE x VPSEQT >= 400MB.
- Preformatting throughput increases up to100%
  - Insert preformatting is asynchronous, except for when allocating a new extent
- Synch I/O cache hit response time decreases by up to 30%
- Sequential prefetch throughput increases up to 19%
- Dynamic prefetch throughput increases up to 23% (40% with SSD)
- Disorganized index scans
  - DB2 10 throughput increases up to 111% (more with 8K pages)
  - Together DB2 10 and zHPF is up to 11 times faster
- Skip sequential index-to-data access
  - Cache misses are 3 to 4 times faster

Read 4K pages from cache
FICON Express 8

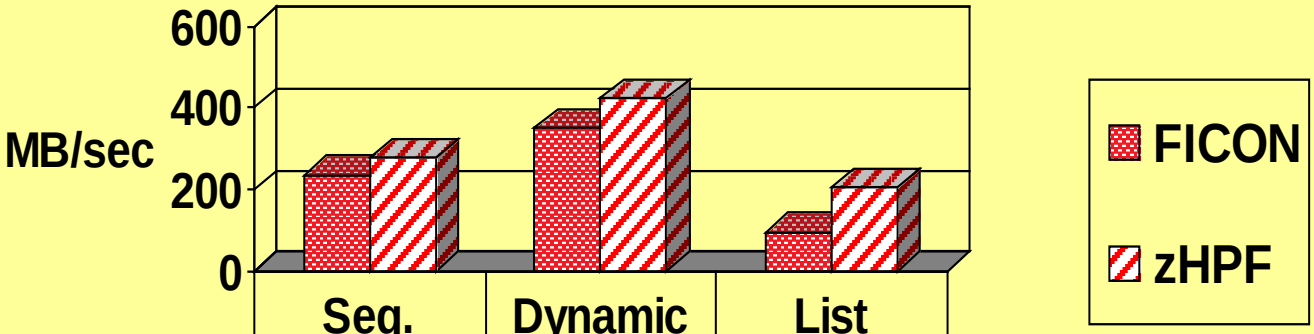

	Seq.	Dynamic	List
FICON	234.3	351	96.9
zHPF	278	423	204.3

- zHPF increases sequential prefetch by 19%
- zHPF increases dynamic prefetch by 23%
- zHPF increases list prefetch by 111%
- Dynamic prefetch still has 110% more throughput than list prefetch

Disorganized Index Scan, hot cache 4K pages

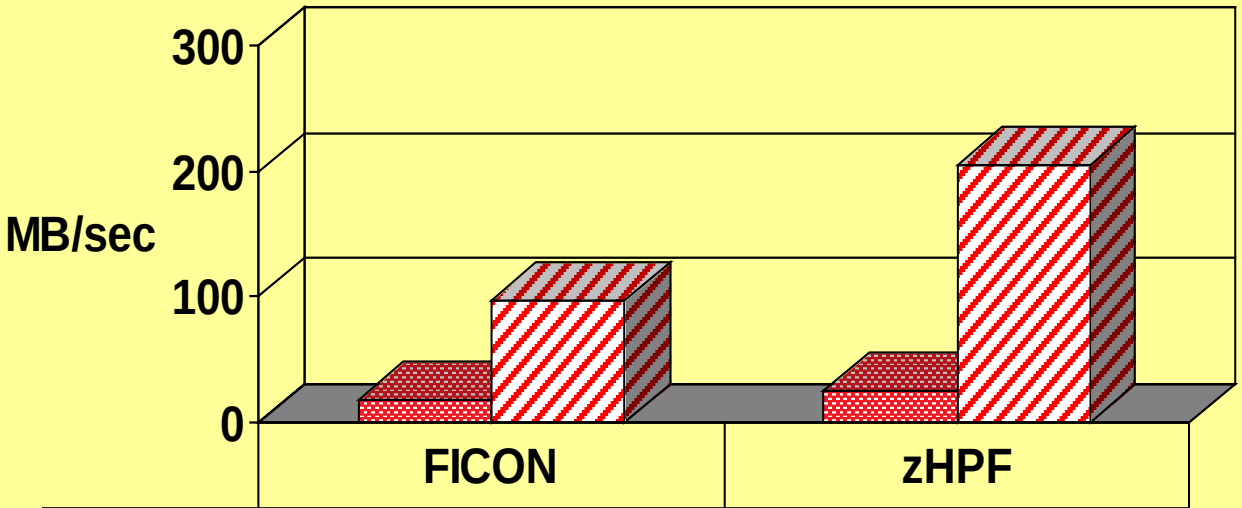

	FICON	zHPF
DB2 9	17.5	25
DB2 10	96.9	204.3

- DB2 9 throughput increases list prefetch by 43%
- DB2 10 throughput increases list prefetch by 111%
- Together DB2 10 zHPF is 11 times faster than DB2 9 FICON

Directions for presentation template – IMPORTANT!!!

- The title slide of your presentation must include the conference title slide with the requested information for Session Title (this should match the title in our system and included in your session confirmation note), speaker information, and session number (9999 format), also in your confirmation note.
- We encourage you to use the entire template for your presentation, but only require the title page be in the Information On Demand template. DO NOT USE TEMPLATE FROM ANOTHER CONFERENCE!

- **To change the cover sheet of existing Powerpoint presentation**
  - Open the conference presentation template
  - Open your existing presentation
  - Highlight cover page of the presentation template - click Edit on menu bar, then Copy in the drop down menu
  - Go to your existing presentation and click Edit on menu bar, then Paste in the drop down menu
  - Immediately look to the bottom right corner of your new paste and you'll see a little folder
  - Click on the folder and choose "Keep Source Formatting"
  - You should now have the new cover slide and none of your existing presentation should be affected
  - When saving, use the following naming convention: “IOD 2011 xxxx speaker last name.ppt” (replace xxxx with your 4-digit session number and add speaker last name)

- **To create new presentation**
  - Open the conference presentation template and click on File on menu bar and Save As in the drop down menu
  - Use the following naming convention: “IOD 2011 xxxx speaker last name.ppt” (replace xxxx with your 4-digit session number and add speaker last name)
  - Add slides as needed
  - Delete any of the master template slides that you don't need

➔ IMPORTANT Directions for presentation templates

- **IBM speakers**, **IF** you make any reference to product/services content that won't be announced until after IOD 2011 concludes, you must include the disclaimer language (on slide 3 here) verbatim, the disclaimer language must be on the second slide in your presentation and is in addition to the Acknowledgements and Disclaimers page which is to be included in all presentations (slide 4 here).
- **IBM speakers**, you must include the Acknowledgements and Disclaimers page (slide 3 here) at the end of your presentation. Please see the notes in red text and update appropriately for your presentation. IF you are speaking with a customer or Business Partner, you MUST use the disclaimer slide.
- Include the survey feedback slide (slide 6) and ask your audience to fill out the session survey on-line. Also include the Communities slide (5)
- Delete any template slides not used. Save and convert to PDF format. Please see below for naming convention for presentation files.
- If you don't have access to PDF converter via Acrobat, there are many available as open source, following is one:
  - http://www.artofsolving.com/opensource/jodconverter

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.

Acknowledgements and Disclaimers:

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS-IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Communities

- **On-line communities, User Groups, Technical Forums, Blogs, Social networks, and more**
  - Find the community that interests you …
    - **Information Management** **ibm.com**/software/data/community
    - **Business Analytics** **ibm.com**/software/analytics/community
    - **Enterprise Content Management** **ibm.com**/software/data/content-management/usernet.html

- **IBM Champions**
  - Recognizing individuals who have made the most outstanding contributions to Information Management, Business Analytics, and Enterprise Content Management communities
    - **ibm.com**/champion

Thank You!
Your Feedback is Important to Us

- Access your personal session survey list and complete via SmartSite
  - Your smart phone or web browser at: iodsmartsite.com
  - Any SmartSite kiosk onsite
  - Each completed session survey increases your chance to win an Apple iPod Touch with daily drawing sponsored by Alliance Tech

1 Line Title Text – 24 Points

- The first level bullet is 20 points
  - The second level bullet is 20 points
    - The third level bullet is 18 points
      - The fourth level bullet is 18 points

2 Line Title Text – 24 Points

This is the look for a sub-title – 20 points

- For a sub-title look, turn off bullet and make bold, italic & choose the green from the color palette

Divider Chart

Use this chart for paragraph statements or cut and paste this box into other charts to provide additional text fields.

Customized these graphics based on speaker content.

Adjust font size as needed.

*** Please note this chart is in addition to the master layouts**

Arrows

- Please use these shapes to highlight important information

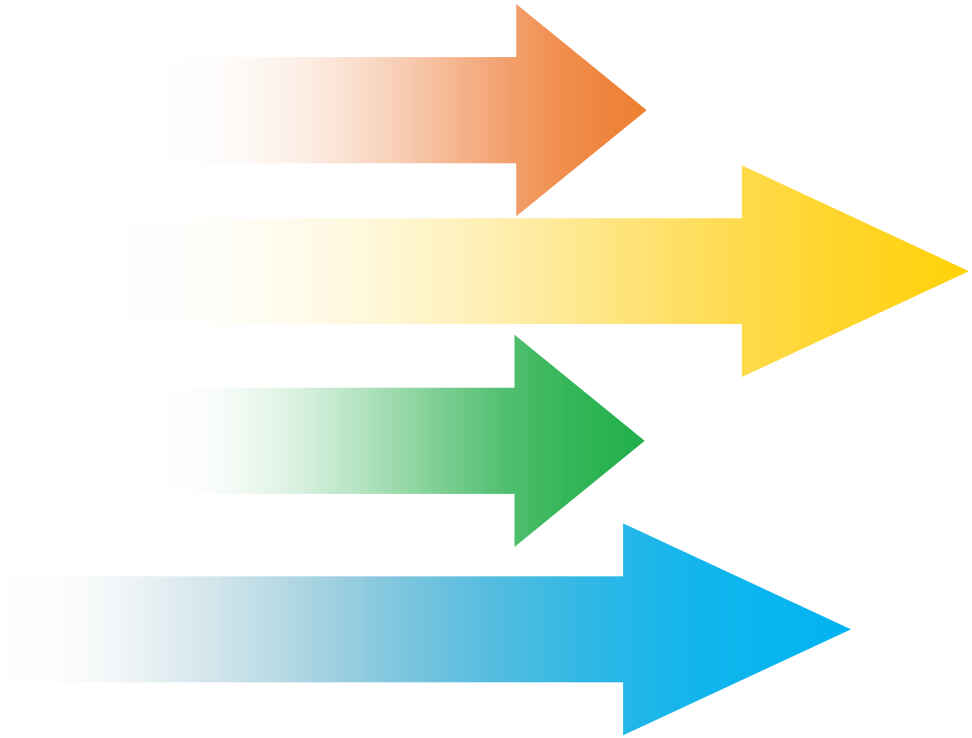

Accent Colors for Diagrams

- Please use these shapes to highlight important information
- Or to break up long lists of bulleted text

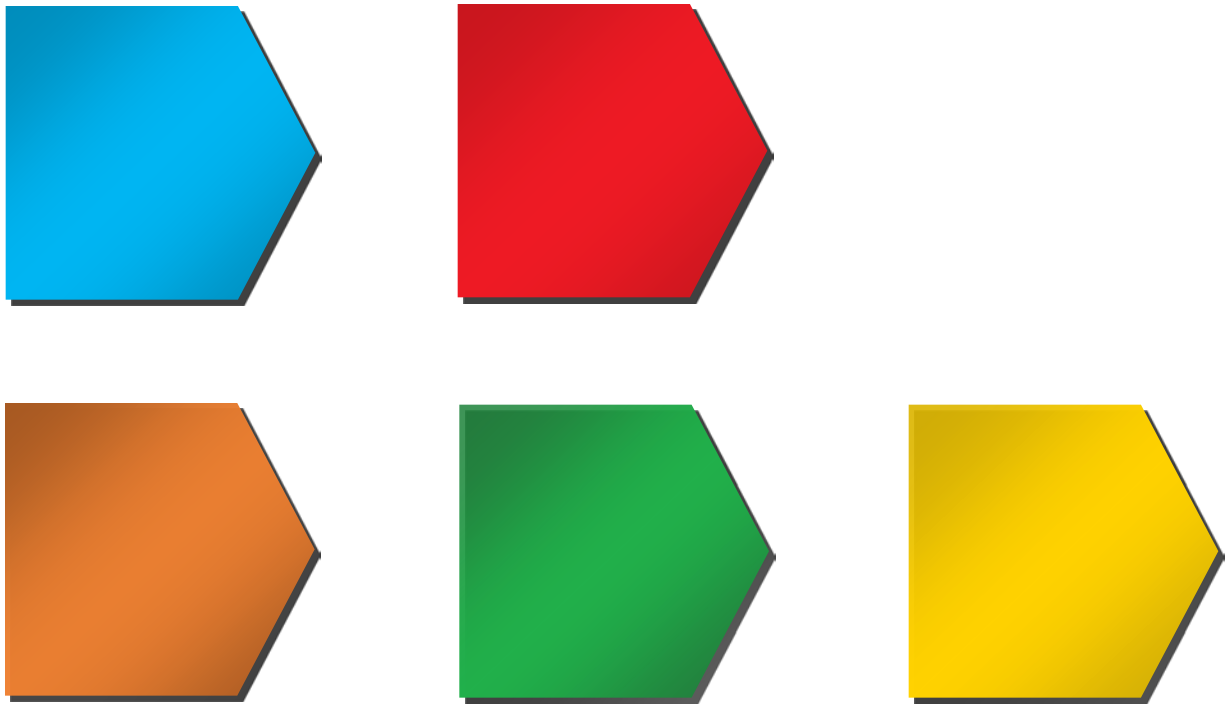

Accent Colors for Diagrams

- Please use these shapes to highlight important information
- Or to break up long lists of bulleted text

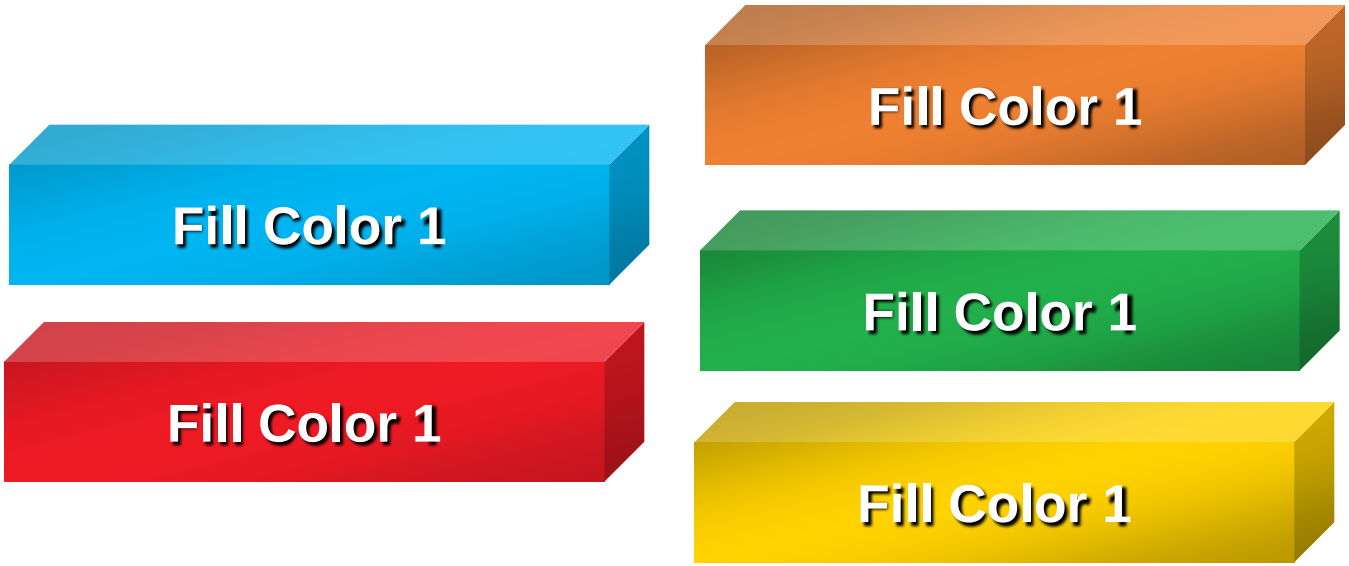

Bar Chart Sample

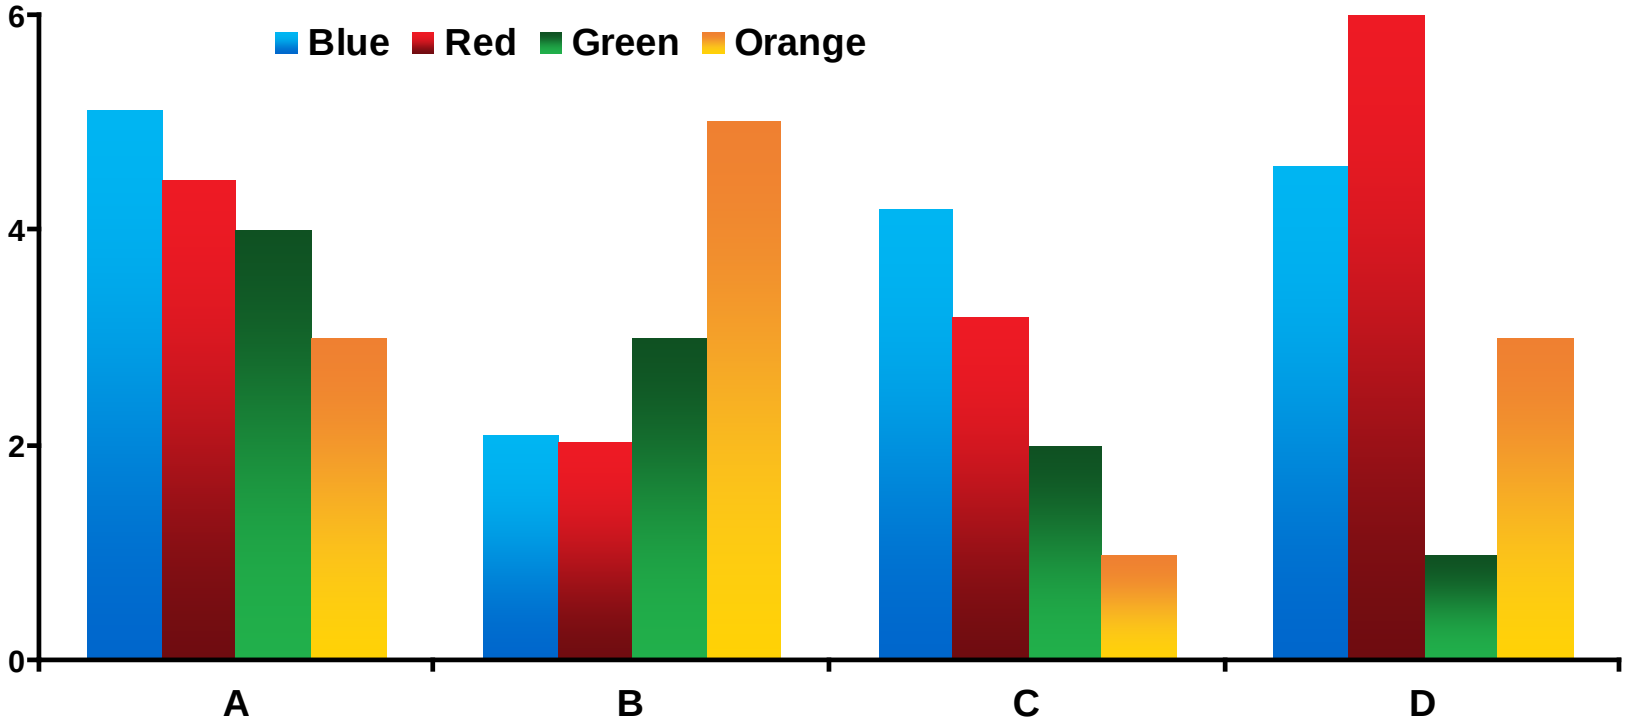

* Legend text here

Pie Chart Sample

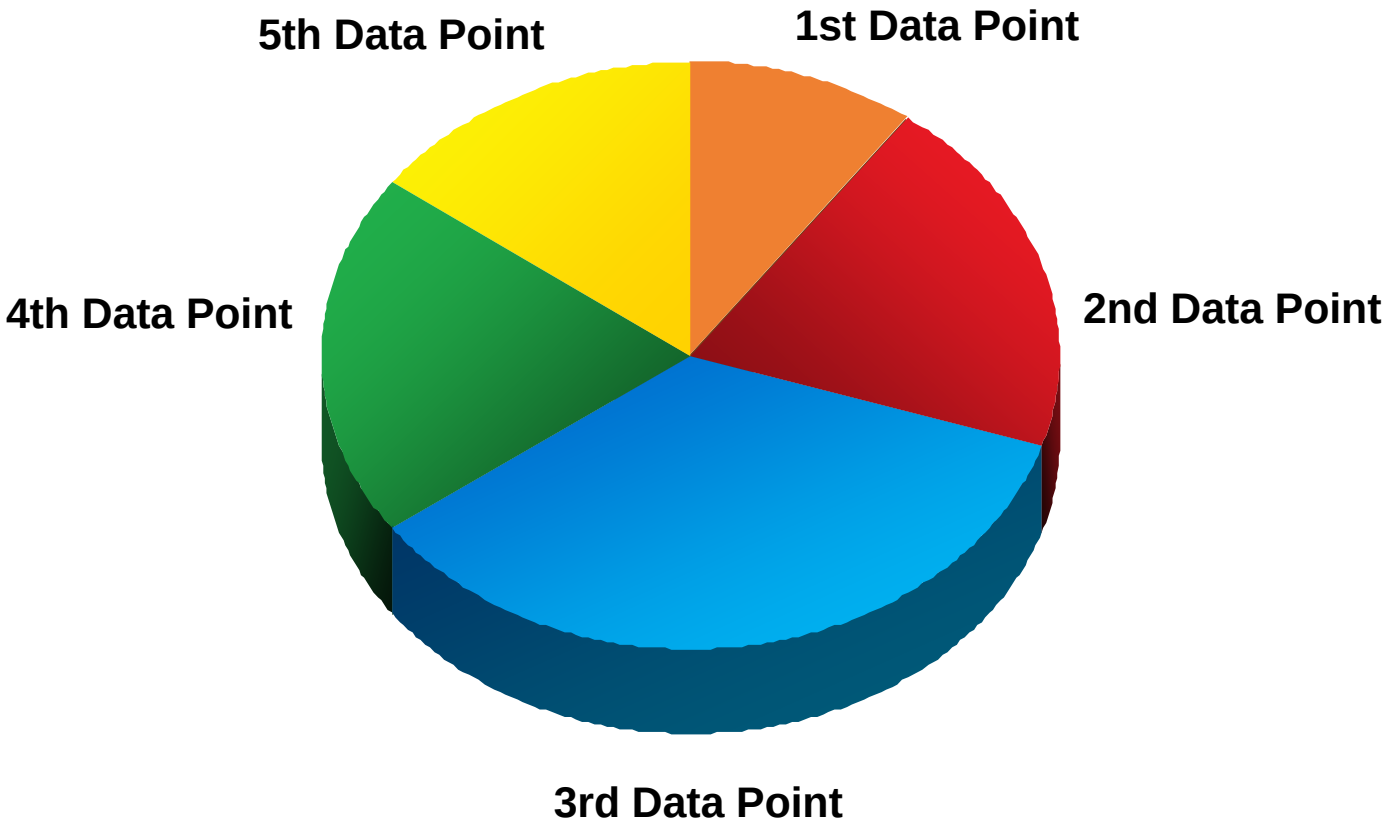

Bar Graph

Graph Title (turn bullets off)

Subtitle

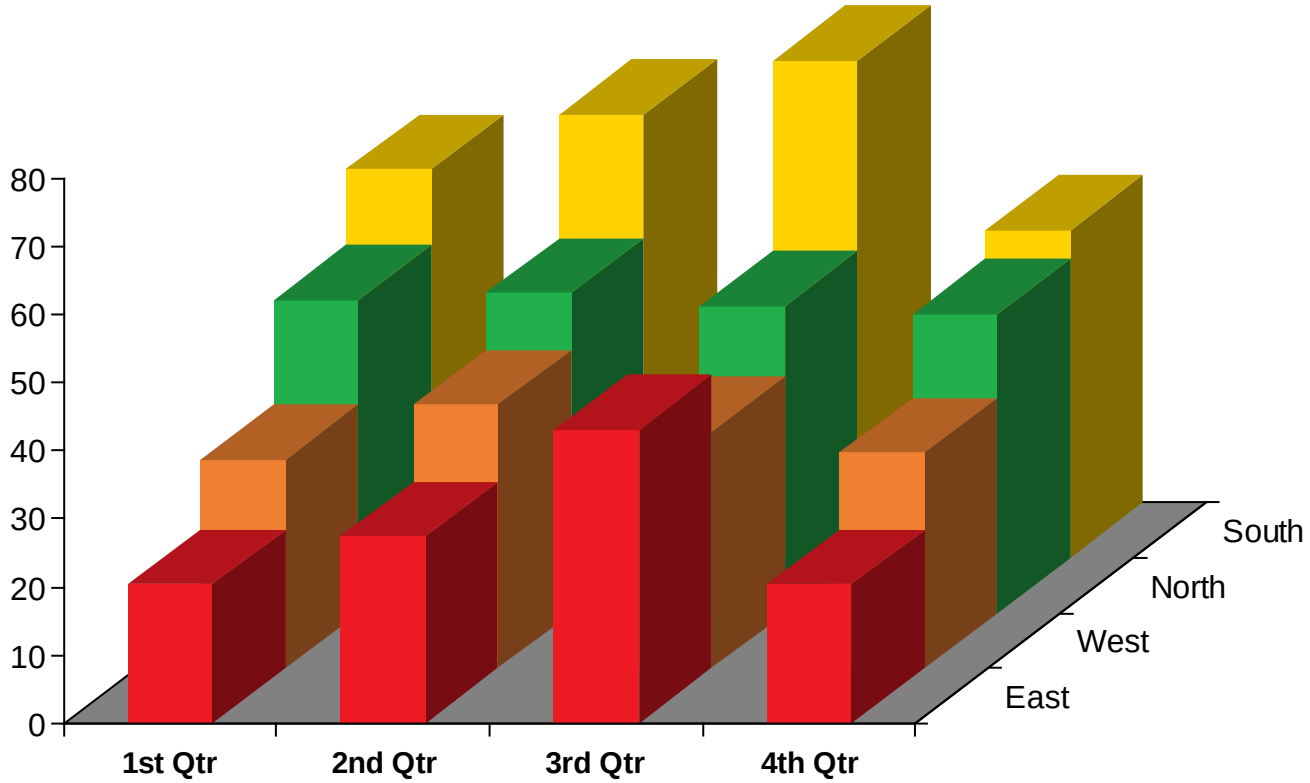

Table

	Header	Header	Header
Text Goes Here	0.00	0.00	0.00
Text Goes Here	0.00	0.00	0.00
Text Goes Here	0.00	0.00	0.00
Text Goes Here	0.00	0.00	0.00