DB2 10, 11 and Beyond

Jay Yothers

DB2 for z/OS Development

IBM Silicon Valley Lab

Please note

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

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

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

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

Agenda

- DB2 10 news and highlights
- DB2 11 preview

InformationOnDemand**2012**
#ibmiod

DB2 10 for z/OS Snapshot

- **Fastest uptake**

    - **+2x** customers vs. V9

    - **+2.5x** licenses vs. V9

    - **25%** coming from DB2 V8

- **Adoption Driven by:**

    - **Performance improvements without application changes**

    - **Virtual Storage Constraint relief for more threads**

    - **Security, RAS improvements**

    - **Bitemporal data**

DB2 V8 to DB2 10 Worldwide

Culmulative Licenses

Quarters since Product Announce
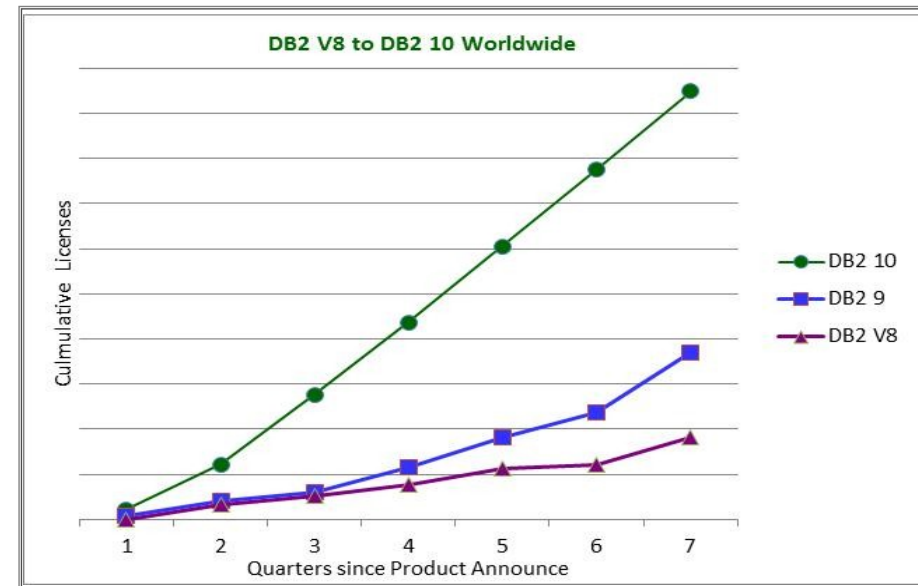
DB2 10
DB2 9
DB2 V8

DB2 10 Performance

- Most customers can see a 5% - 10% out-of-the-box CPU reduction (transactions and batch) after rebind

- Some workloads and customer situations can see a CPU reduction of up to 20%

- Synergistic operation with latest System z hardware

Sample: Preliminary Measurements of IBM Relational Warehouse Workload (IRWW) with Data Sharing

- Base: DB2 9 NFM REBIND with PLANMGMT EXTENDED

- DB2 9 NFM → DB2 10 CM without REBIND showed 1.3% CPU reduction

- DB2 10 CM REBIND with same access path showed 4.8% CPU reduction

- DB2 10 NFM brought 5.1% CPU reduction

- DB2 10 CM or NFM with RELEASE DEALLOCATE 12.6% CPU reduction from DB2 9

Average %CPU improvements version to version
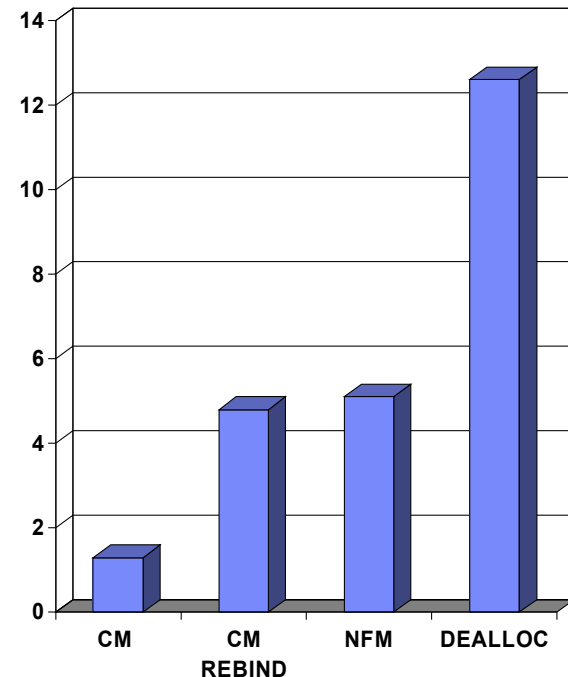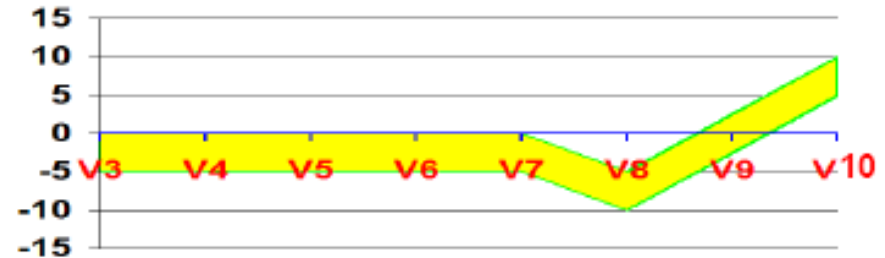
CPU Percent reduced from DB2 9

DB2 10: A Few Interesting Performance Features

- **CM:**

  - High performance DBATs

  - Parallel index I/O at insert

  - Index list prefetch

  - SQLPL performance improvements

  - SQL pagination performance enhancement

  - 1M page frames for buffer pools

- **NFM:**

  - Hash access to data

  - Index include columns

  - Inline LOBs

DB2 10 Virtual Storage Constraint Relief

DB2 10
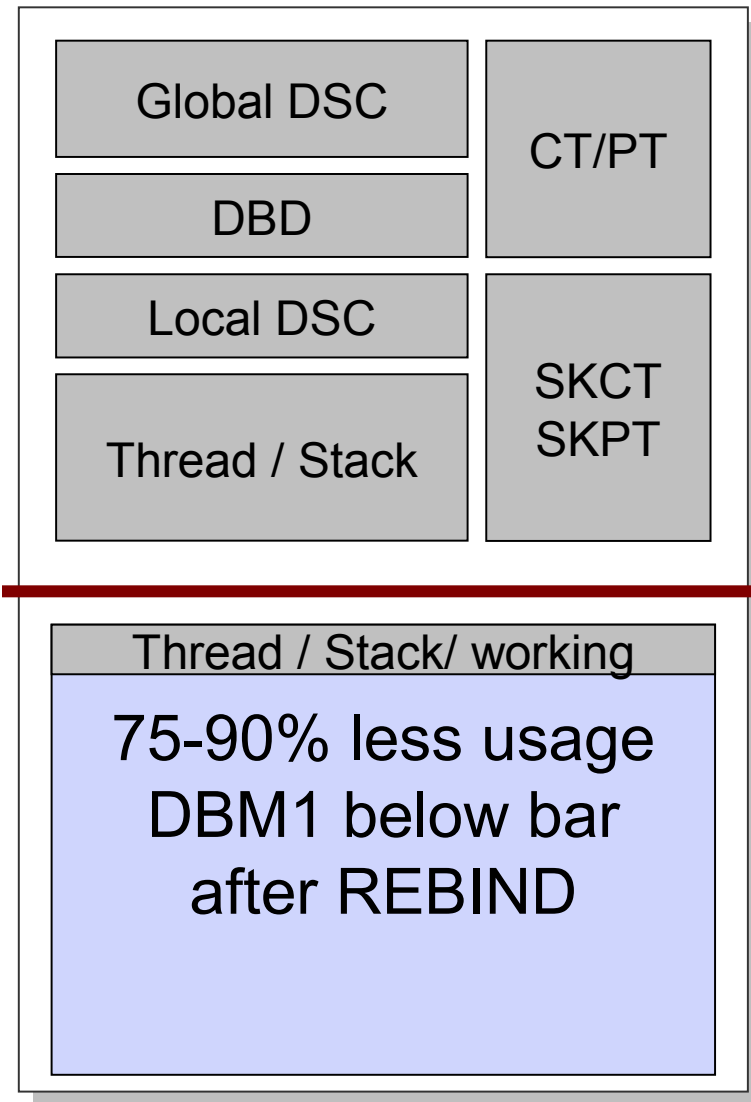
- **DBM1 below 2GB**

  – 75-90% less usage in DB2 10 compared to DB2 9

  – Some of working storage (stack, xproc storage) stays below 2GB

- **Larger number of threads**

  – Possible data sharing member consolidation

- **Improve CPU with storage**

  – More release deallocate

  – Larger MAXKEEPD values for KEEPDYNAMIC=YES

Global DSC	CT/PT
DBD	
Local DSC	SKCT SKPT
Thread / Stack	

Thread / Stack/ working

75-90% less usage DBM1 below bar after REBIND

DB2 10 Productivity – Doing More with Less!

- **Easier scaling, simpler memory management**

- **Reduce contention, more online processing**

- **Reduced need for REORG**
  - Build compression dictionary on the fly
  - Index list prefetch enhancements
  - Row-level sequential detection

- **Configure IBM UDFs and stored procedures**

- **Statement level monitoring**

- **Access path stability, APREUSE & APCOMPARE**

- **DDF thread management enhancements**

Manual invocation of
- RUNSTATS
- COPY/BACKUP SYSTEM
- QUIESCE
- MODIFY RECOVERY
- REORG

Changes in DB2 10 catalog & directory

- Improve availability and productivity

- Increase maximum size substantially

- Reduce contention: BIND, Prepare, Utilities

  - DDL concurrency also improved from removal of DBD01 hash anchor locks

- Catalog changes: Remove links, hashes

  - Many more table spaces, partition by growth

  - Row level locking, reordered row format

  - CLOB and BLOB columns for long strings

    - Inline for performance

  - Online reorganization and check

  - More automatic: DB2-managed SMS-controlled

  - Allow query of SYSLGRNX

  - Allow SQL statements in catalog to be queried with normal SQL
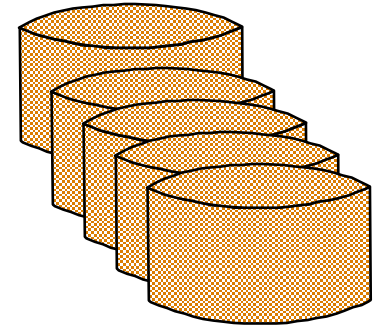
Key details about DB2 10: getting ready

Prerequisites: migrate from DB2 9 NFM or DB2 V8 NFM

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

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

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

Creating the Hybrid Data Server - Netezza and Z
Combine DB2 for z/OS with Netezza to provide an industry exclusive

Data Mart **Data Mart** **Data Mart**

Data Mart Consolidation

Transaction Processing
Systems (OLTP)

Transactional Analytics

Deep Analytics

DB2 z/OS:
*Recognized leader in
transactional workloads with
security, availability
and recoverability*

**Netezza
Accelerator**
*Recognized leader in
cost-effective high speed
deep analytics*

Together:
*Destroying the myth that transactional and
decision support workloads have to be on
separate platforms*

**Best in OLTP and
Transactional Analytics**

*Industry recognized leader in
mission critical transaction
systems*

Best in Deep Analytics

*Proven appliance leader in high
speed analytic systems*

Best in Consolidation

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

DB2 Analytics Accelerator V3
Lowering the costs of trusted analytics

What's New?

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

- **Incremental Update**
  - Enables tables within the Accelerator to be continually updated throughout the day.

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

- **Query Prioritization**
  - Brings System z workload management down to the individual query being routed to the Accelerator

- **High Capacity**
  - Support has been extended to include the entire Netezza 1000 line *(1.28 PB)*

- **UNLOAD Lite**
  - *Reduces z/OS MIPS consumption, by moving the preparation off System z.*

New Technology Emerges

NoSQL Databases

Column-store Databases

"Impedance mismatch" (storing aggregates)

In-memory Databases

XML Databases

Object Databases

Optimized storage format for query & analytics

Leverage larger, cheaper memories for I/O avoidance

1990s   **2000s**   **2010s**
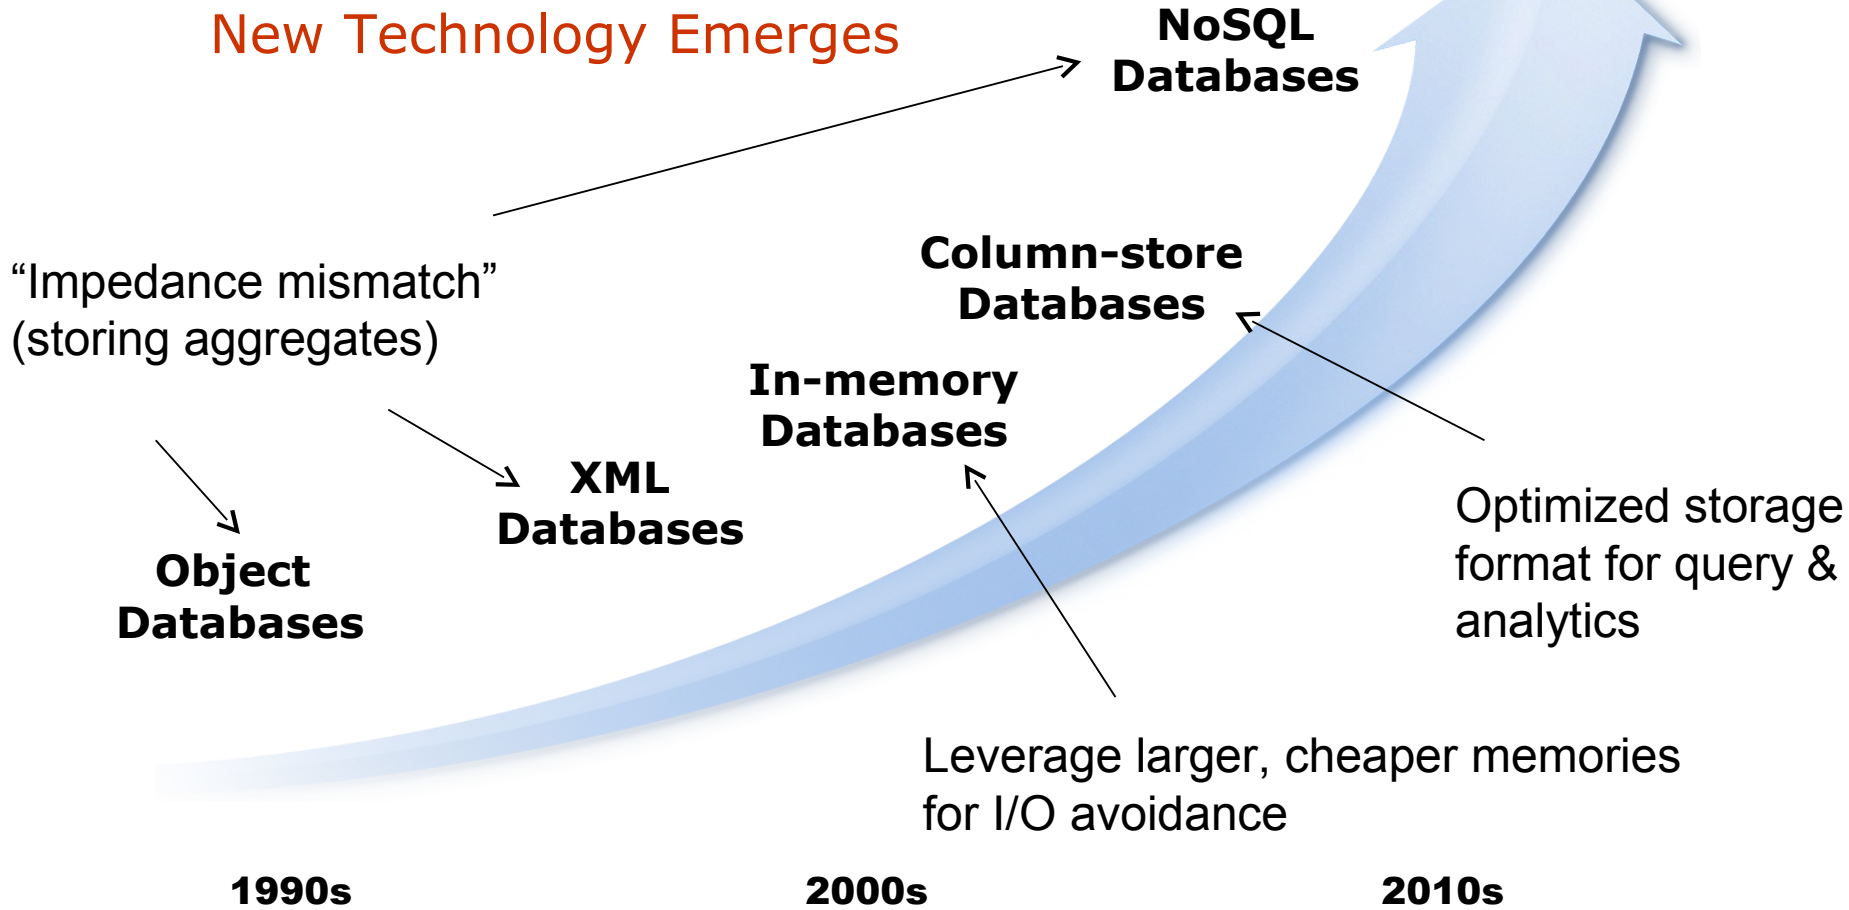
Eliminate mapping and transformation through the tiers
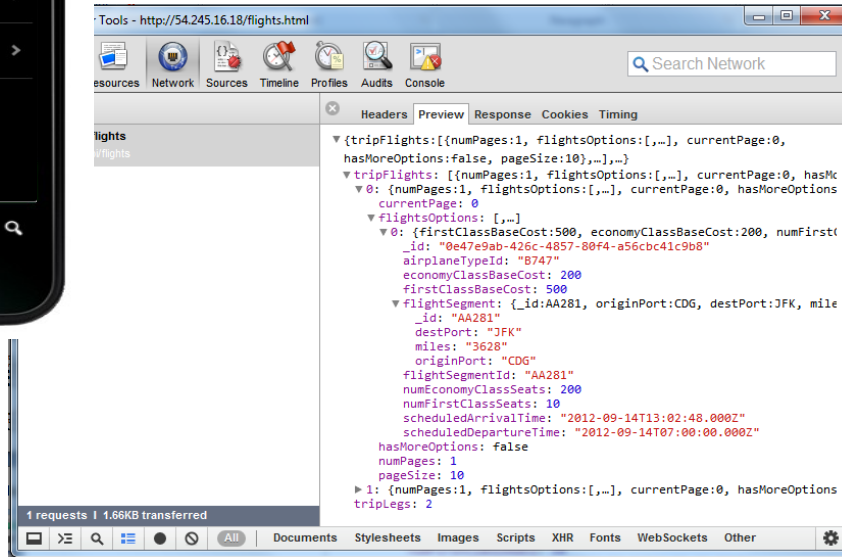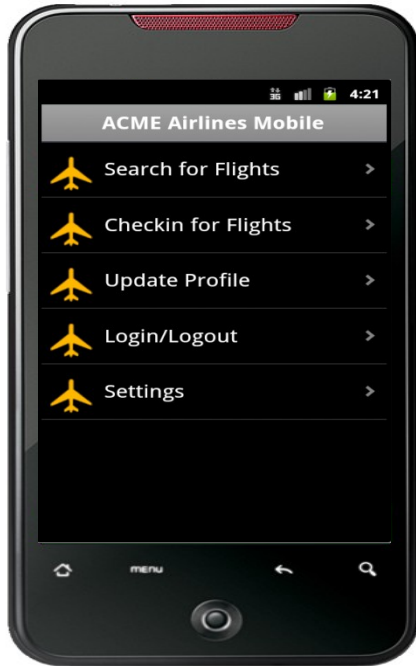
Store business objects (e.g. itinerary) as a single entity in DB2 instead of rows that must be pieced together with SQL
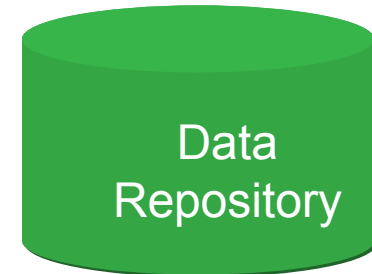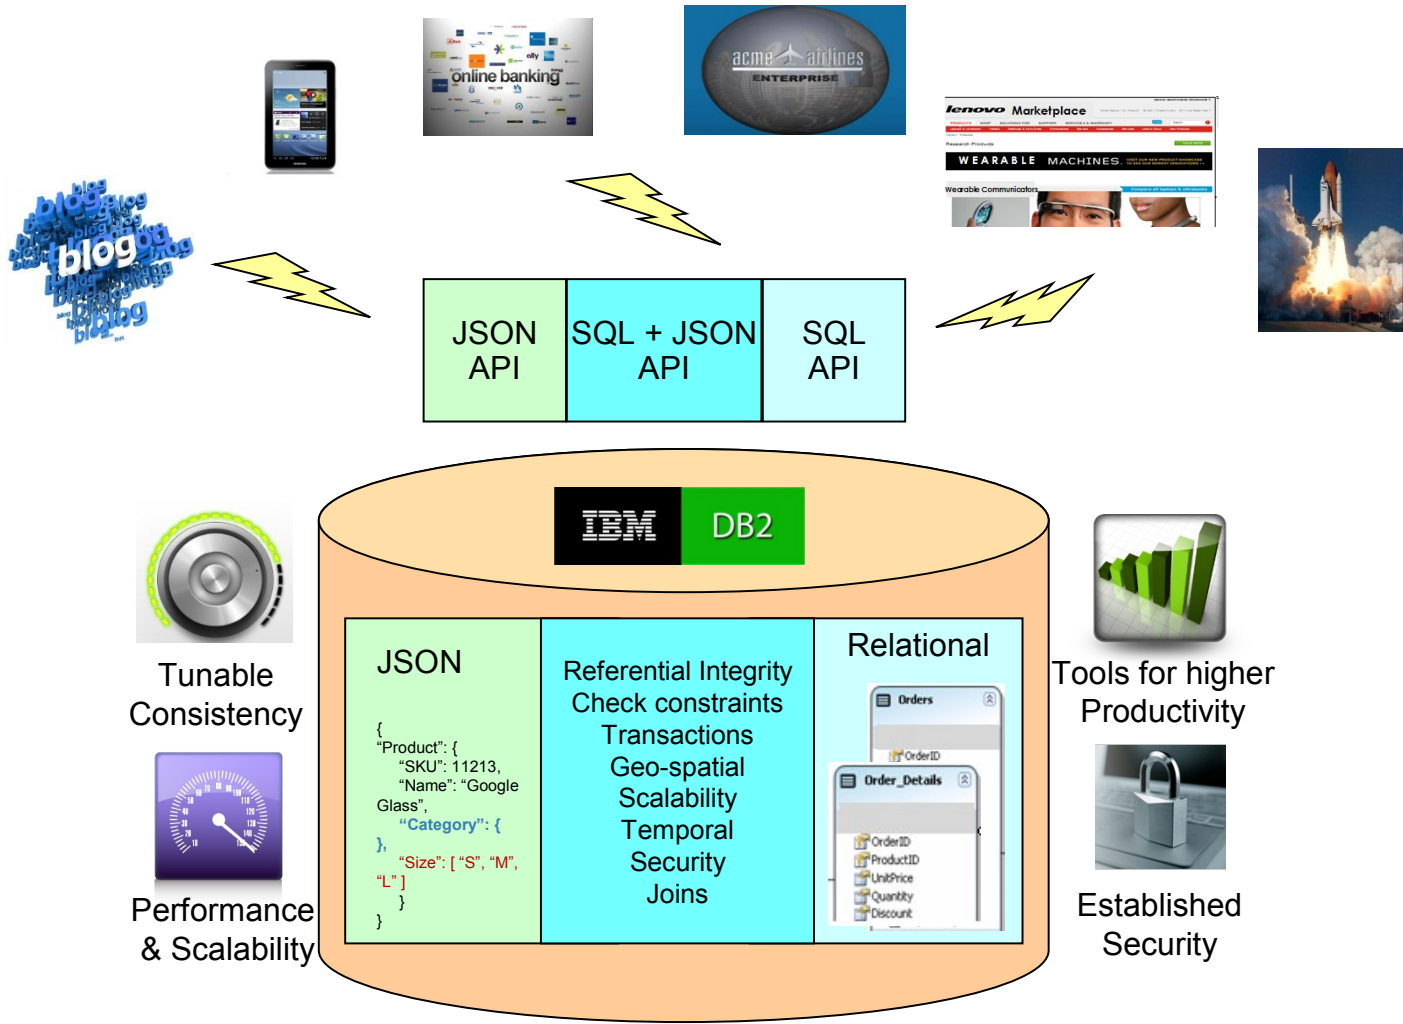
(XML can also be used to represent business objects in DB2)

Data Repository

{tripFlights:[{numPages:1, flightsOptions:[,…], currentPage:0, hasMoreOptions:false, pageSize:10},…],…}
▼tripFlights: [{numPages:1, flightsOptions:[,…], currentPage:0, hasMc
  ▼0: {numPages:1, flightsOptions:[,…], currentPage:0, hasMoreOptions
      currentPage: 0
    ▼flightsOptions: [,…]
      ▼0: {firstClassBaseCost:500, economyClassBaseCost:200, numFirstC
          _id: "0e47e9ab-426c-4857-80f4-a56cbc41c9b8"
          airplaneTypeId: "B747"
          economyClassBaseCost: 200
          firstClassBaseCost: 500
        ▼flightSegment: {_id:AA281, originPort:CDG, destPort:JFK, mile
            _id: "AA281"
            destPort: "JFK"
            miles: "3628"
            originPort: "CDG"
          flightSegmentId: "AA281"
          numEconomyClassSeats: 200
          numFirstClassSeats: 10
          scheduledArrivalTime: "2012-09-14T13:02:48.000Z"
          scheduledDepartureTime: "2012-09-14T07:00:00.000Z"
      hasMoreOptions: false
      numPages: 1
      pageSize: 10
  ▶1: {numPages:1, flightsOptions:[,…], currentPage:0, hasMoreOptions
  tripLegs: 2

JSON Database Technology Preview
Providing the best of both worlds

| JSON API | SQL + JSON API | SQL API |

IBM DB2

| JSON | Referential Integrity
Check constraints
Transactions
Geo-spatial
Scalability
Temporal
Security
Joins | Relational |

```
{
"Product": {
  "SKU": 11213,
  "Name": "Google Glass",
  "Category": {
},
  "Size": [ "S", "M", "L" ]
  }
}
```

Orders
OrderID

Order_Details
OrderID
ProductID
UnitPrice
Quantity
Discount

Tunable Consistency

Performance & Scalability

Tools for higher Productivity

Established Security
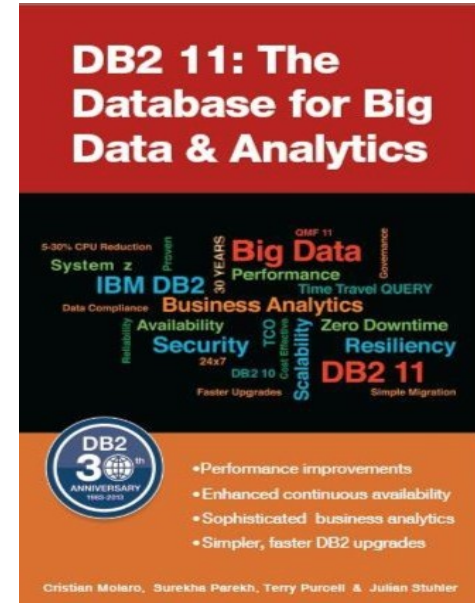
DB2 for z/OS Technical Strategy

➢ Continuous availability, RAS leadership

➢ Performance and scalability

➢ Autonomics and simplification

➢ Advanced application features

➢ Grow real time analytics capabilities

DB2 11 Major Themes

- **Out-of-the-box CPU Savings***
  - Improving efficiency, reducing costs, no application changes
  - Up to 10% for complex OLTP
  - Up to 10% for update intensive batch
  - Up to 40% for queries
  - Additional performance improvements through use of new DB2 11 features

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

- **Enhanced business analytics**
  - Expanded SQL, XML, and analytics capabilities
  - Temporal and SQLPL enhancements
  - Transparent archiving

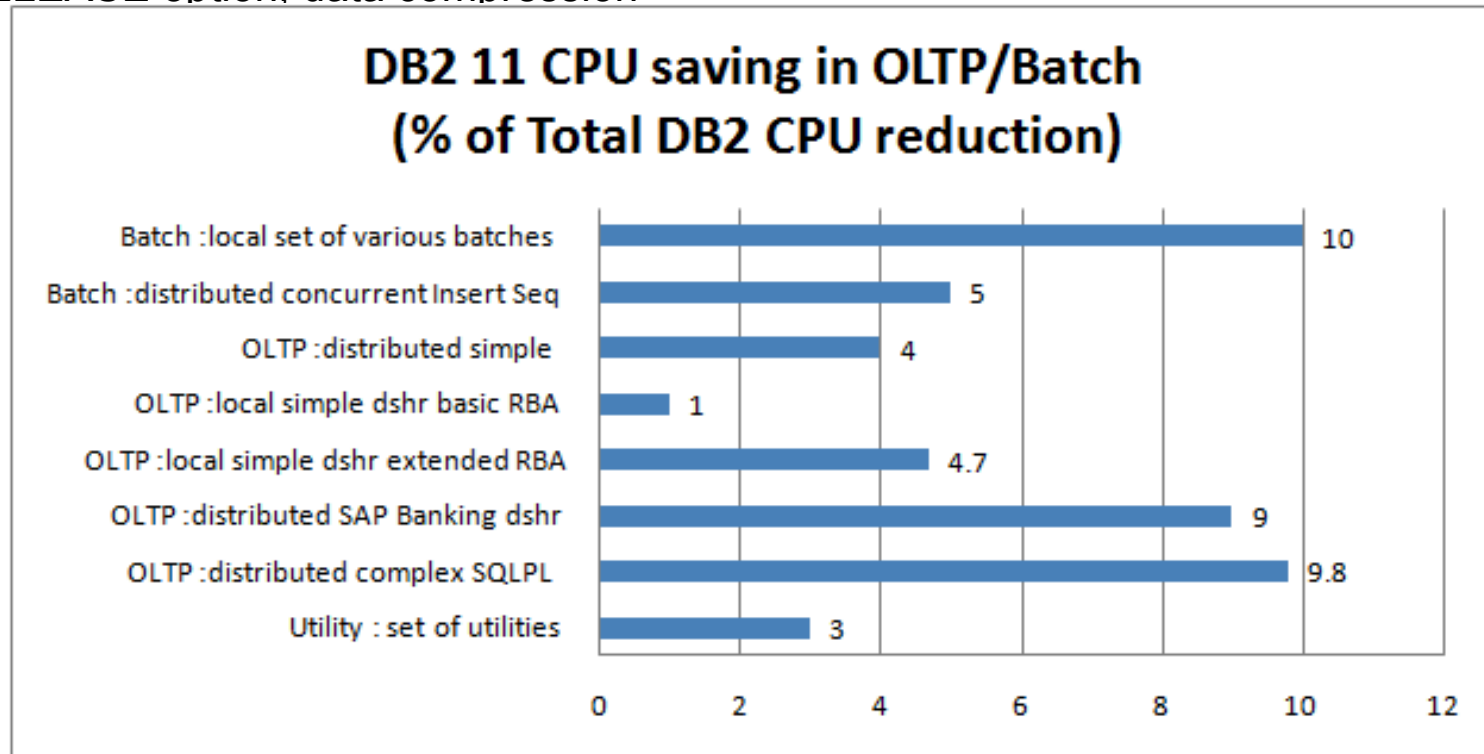- **Simpler, faster DB2 version upgrades**
  - No application changes required for DB2 upgrade
  - Access path stability improvements
  - Product quality/stability: support pre GA customer production

*REBIND may be required for best results

**IBM announced DB2 11
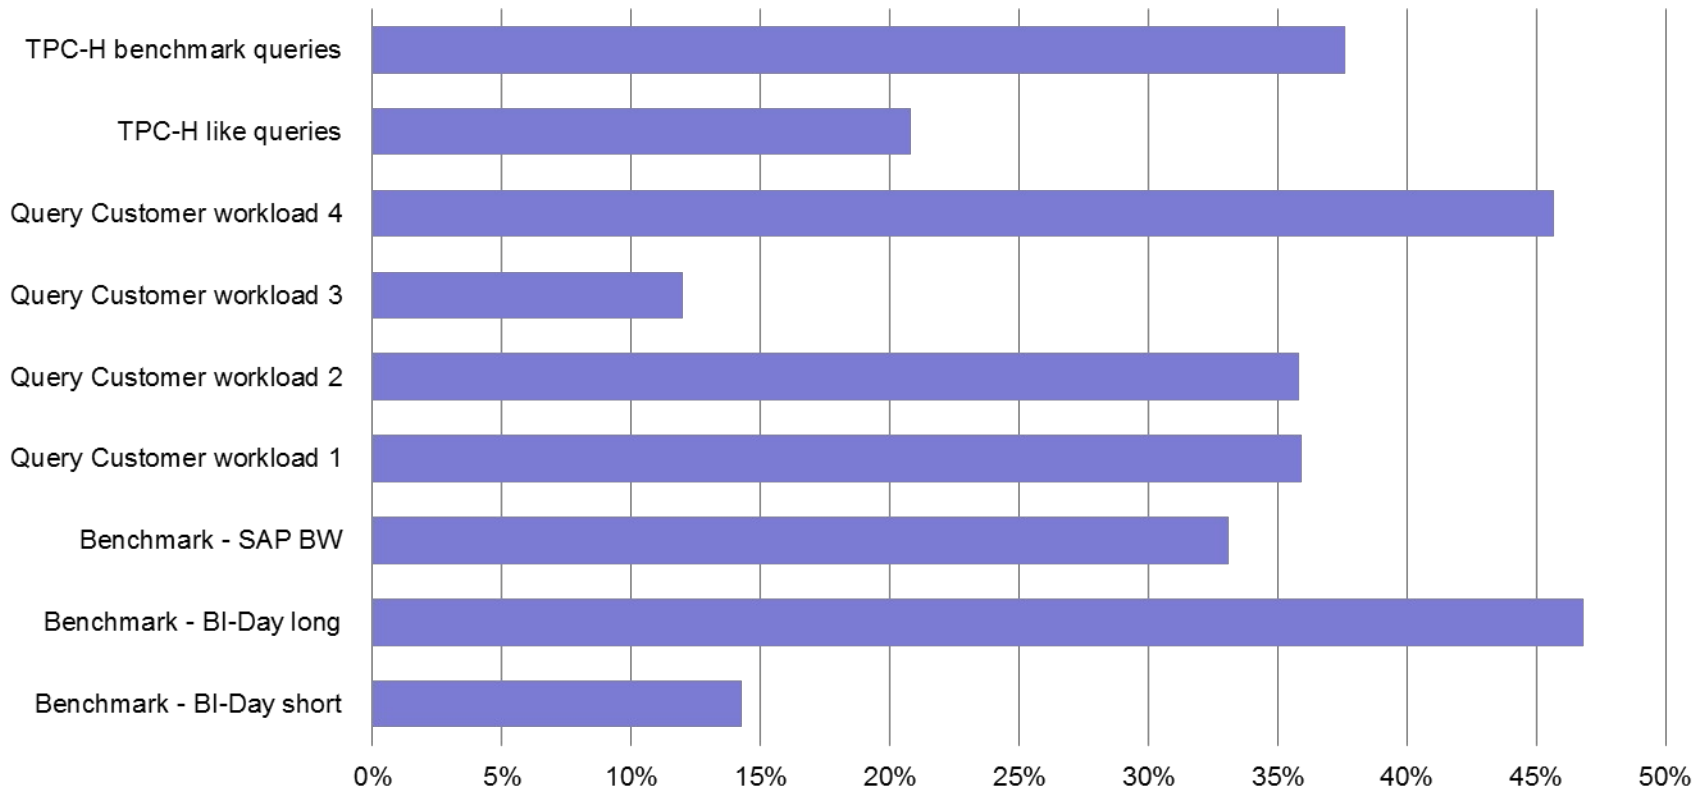Early Support Program**
→ Learn more

DB2 11 OLTP/Batch Performance Expectations

- These are preliminary results from IBM testing
- Performance expectations vary depending on many factors, including
    - Access path selection, Read/Write ratio, Number of rows returned
    - Number and type of columns returned, Number of partitions touched
    - Schema - Number of partitions defined, DPSI, etc
    - RELEASE option, data compression

**DB2 11 CPU saving in OLTP/Batch
(% of Total DB2 CPU reduction)**

Category	Value
Batch :local set of various batches	10
Batch :distributed concurrent Insert Seq	5
OLTP :distributed simple	4
OLTP :local simple dshr basic RBA	1
OLTP :local simple dshr extended RBA	4.7
OLTP :distributed SAP Banking dshr	9
OLTP :distributed complex SQLPL	9.8
Utility : set of utilities	3

Performance Improvements – no REBIND needed

- DDF performance improvements
  - Reduced SRB scheduling on TCP/IP receive using new CommServer capabilities
  - Improved autocommit OLTP performance
  - DRDA package based continuous block fetch
- xProcs above the bar
  - 31-bit Vstor relief enabled by RMODE 64 support in z/OS 1.13 and above
  - Enables other internal performance improvements
- zIIP enablement for all SRB-mode DB2 system agents that are not response time critical
- Avoid cross-memory overhead for writing log records
- Data decompression performance improvement
- INSERT performance
  - Latch contention reduction for classes 6, 14, 19
  - CPU reduction for Insert column processing and log record creation
  - Data sharing LRSN spin avoidance
  - Page fix/free avoidance in GBP write

Performance Improvements – no REBIND needed…

- Automatic index pseudo delete cleanup
  - For fine-tuning, DBA work would be required
- ODBC/JDBC type2 performance improvements
  - Stored procedure invocation
- Java stored procedure multi-threading improvements
- Sort performance improvements
- DPSI performance improvements for merge
- Performance improvements with large number of partitions
- XML performance improvements
- Optimize RELEASE(DEALLOCATE) execution so that it is consistently better performing than RELEASE(COMMIT)
  - Monitor # parent locks and cleanup internal structures when threshold is hit
- IFI 306 filtering capabilities to improve Replication capture performance
- Utilities performance improvements

Performance Improvements – no REBIND needed…

- ACCESS DATABASE command performance
- DGTT performance improvements
  - Avoid incremental binds for reduced CPU overhead
- P-procs for LIKE predicates against Unicode tables
- Improved performance for ROLLBACK TO SAVEPOINT
- zEC12 exploitation:
  - Pageable 1M frames for buffer pool control structures
  - 2G page frame size
  - 1M page frames for DB2 code.
    - Requires z/OS 2.1 or above and zEC12 Flash Express
- Latch contention reduction and other high n-way scalability improvements
- Data sharing performance improvements
  - LRSN spin reduction with extended LRSNs
  - Castout performance
  - GBP write-around
  - Index split performance

Performance Improvements – REBIND required (with or without APRESUE)

- Query transformation improvements – less expertise required to write performant SQL
  - Enhanced query rewrite to improve predicate indexability
    - new situations where non-indexable predicates can be rewritten by Optimizer to be indexable
    - Convert some common stage 2 predicates to indexable (YEAR(), DATE(), SUBSTR(col,1,x), value BETWEEN COL1 AND COL2)
    - Improved indexability for OR COL IS NULL predicates
    - Push complex predicates inside materialized views/table expressions
  - Enhanced pruning of "always true" and "always false" predicates
- Enhanced duplicate removal
  - Lots of queries require duplicate removal: e.g. DISTINCT, GROUP BY, etc.
  - Dup elimination via sorting can be expensive
  - New techniques: Index duplicate removal, early out
  - Will not show in Explain table, need to look at IXSCAN_SKIP_DUPS column in DSN_DETCOST_TABLE to determine if sort avoided

Performance Improvements – REBIND required (with or without APRESUE)…

- **In-memory techniques**
  - In-memory, reusable workfile
  - Sparse index (limited hash join support)
  - Non-correlated subquery using MXDTCACH
  - Correlated subquery caching
- **Non correlated subquery with mismatched length**
- **Select list do-once**
  - Non column expressions in the select list can be executed once rather than per-row
- **Column processing improvements**
  - Xproc (generated machine code) for output column processing
  - Optimized machine instructions for input/output column processing

Performance Improvements – REBIND required (with or without APRESUE)…

- ## RID overflow to workfile handled for Data Manager set functions
  - DB2 10 added RID overflow to workfile
  - DB2 11 adds support for set functions (COUNT, MAX, MIN etc) which was excluded in DB2 10
- ## Performance improvements for common operators
  - MOVE, CAST, output hostvar processing, CASE, SUBSTR, DATE, others
- ## DECFLOAT data type performance improvements
  - Up to 23% CPU reduction for conversion to/from decfloat
  - Approx. 50% cpu reduction in INSERT, FETCH for decfloat columns
  - Helped further by zEC12 hw improvements for decimal floating point

Performance Improvements – REBIND required (without APRESUE)

- DPSI and page range performance improvements
  - Page range screening for join/correlation predicates
  - Parallelism optimization for DPSI access

- Optimizer CPU and I/O cost balancing improvements
  - Measured results: 3% to >30% performance improvement for query workloads

Performance Improvements – DBA or application effort required

- **Suppress-null indexes**
  - Index entries not created when all values for indexed columns are NULL
  - Reduced index size, improved insert/update/delete performance, compatibility with other DBMSes
  - Improved utility CREATE INDEX performance
- **New PCTFREE FOR UPDATE attribute to reduce indirect references**
- **DGTT performance improvements**
  - Non logged DGTTs
- **Global variables**
  - Easier, more efficient sharing of data between SQL statements

Performance Improvements – DBA or application effort required

- Optimizer externalization of missing/conflicting statistics
  - Identify missing statistics during bind/prepare/explain
  - DBA or tooling to convert output to RUNSTATS input
- Extended optimization - selectivity overrides (filter factor hints)
  - Improve optimizer's ability to find the cheapest access path
  - Collect filter factors for predicates in a Selectivity Profile
  - Selectivity Profile is populated via BIND QUERY
- Open data set limit raised to 200K

DB2 and zEnterprise EC12

- **Faster CPU – 1.25x compared to z196**
  - 20-28% CPU reduction measured with DB2 OLTP workloads
  - 25% reduction measured with DB2 query and utilities workloads
  - Less compression overhead with DB2 data (1-15%)
- **50% More System Capacity to help consolidation**
  - Excellent synergy with DB2 10 scalability
- **New Features DB2 plans to exploit**
  - FLASH memory and pageable 1MB frames
    - Improved DUMP and z/OS paging performance
    - Pageable 1M frames for buffer pool control blocks for CPU savings (retrofit to V10 via APAR)
  - 2GB frame support (V11 only)
    - Additional CPU savings, especially for very large memory
  - Optimizer CPU & I/O cost balancing improvements improves query performance on zEC12 (V11 only)
- **Transactional Memory provides further possibilities for performance gains**

DB2 and IBM zIIP Add Value to Database Work

Portions of the following DB2 workloads in enclave SRB mode are eligible for zIIP*

DB2 9 in blue   DB2 10 in green   DB2 11 in orange

1. DRDA over TCP/IP connections: up to 60% of the processing
   - DB2 9 for zOS remote native SQL procedures
   - DB2 9 XML parsing, schema validation

2. Requests that use parallel queries: up to 80% of the processing after reaching a CPU usage threshold
   - DB2 9 and DB2 10 remove restrictions for query parallelism enabling more queries to run with parallelism and therefore to potentially increase zIIP eligibility

3. DB2 utilities: up to 100% of the processing
   - LOAD, REORG and REBUILD functions used to maintain index structures and sort
   - DB2 10 RUNSTATS – options other than column group, inline
   - DB2 11 RUNSTATS column group and inline

4. Asynchronous processing that is charged to a DB2 address space (introduced in DB2 10, expanded in DB2 11): up to 100% of the processing
   - DB2 10 buffer pool prefetch and deferred write
   - All other such asynchronous processing, except for P-lock negotiation

DB2 Related enhancements in z/OS 2.1

- Enhancing DB2 BACKUP SYSTEM solution (DB2 11 only)
  - ❏ **Enable recovery of single page set from DB2 system-level backup even if original volume does not have sufficient space**
  - ❏ **Enable exploitation of FlashCopy consistency group for DB2 BACKUP SYSTEM**
  - ❏ **Enable restore of a page set to a different name**

- z/OS DFSMS RLS for Catalog support (DB2 9 and above)
  - ❏ **Improved DB2 data set open/close performance**

- z/OS DFSMS StorageTiers (DB2 9 and above)
  - ❏ **Optimizes disk placement on SSD and HDD**

- Enhanced WLM managed buffer pools (DB2 9 and above)
  - ❏ **WLM now allows buffer pool sizes to shrink**

- XES GBP write-around support (DB2 11 only)
  - ❏ **CFLEVEL 17 or above (retrofit to z/OS 1.13 with OA37550)**

DB2 Related enhancements in z/OS 2.1...

- Reducde DRDA message latency and DB2 CPU reduction (DB2 11 only)
- **TCP/IP synchronous receive (retrofit to z/OS 1.13 with OA39810)**

- Reduce DRDA message latency and DB2 CPU reduction (DB2 9 and above)
- **CommServer SMC-R support for RDMA over ethernet (RoCE),**
- **Requires zEC12 or zBC12, z/OS to z/OS only for now**

- 2G page frame support (DB2 11 only)
- **Requires zEC12 or zBC12 (retrofit to z/OS 1.13)**
- **Retrofit to z/OS 1.13 – OA40967**

- Pageable 1M frame support (DB2 10 and 11)
- **Requires zEC12 or zBC12 (retrofit to z/OS 1.13)**
- **DB2 V10 APAR PM85944**

- Ability to execute code loaded into 1M frames (DB2 11 only)
- **SVL measured 1.8% CPU reduction for an OLTP workload**

RAS and Usability Improvements

- Expanded RBA/LRSN. Expand to 10 bytes
- Increase 2G limit for a single internal DB2 storage pool
- BIND / DDL / Online REORG concurrency with persistent threads
  - Use of persistent threads will increase in V10 with vstor relief
  - Examples: IMS PWFI, CICS protected entry
  - Avoid having to shut down these apps to get a REBIND through
- DEFER DEFINE improved concurrency for first insert
- More online schema changes
  - Alter partitioning limit keys
  - DROP column
  - Alter Drop Pending Changes: AREOR status is now removed
  - Point in time recovery support for deferred schema changes

DB2 11 Planned RBA/LRSN Solution

- Expand the RBA and LRSN to 10 bytes
  - RBA addressing capacity of 1 yottabyte (2**80)
  - LRSN extended on left by 1 byte, on the right by 3 bytes
    - >30,000 years and 16Mx more precision
  - 8 bytes is not sufficient to solve LRSN issues and may not give sufficient capacity for the longer term
- NFM only (6 byte RBA/LRSN continues to be used in CM)
- Once in NFM, DB2 continues to use 6-byte values until you take action to convert
- Two conversion tasks:
  - Convert BSDSes to new format to enable logging with larger RBAs/LRSNs
  - Convert page sets to new page format
- These tasks are optional
  - If you don't care about larger RBAs/LRSNs then you don't have to convert
  - But performance will be better if you convert BSDSes (avoid internal conversion overhead on log write)
- BSDSes can be converted without converting page sets
- Page sets can be converted in a piecemeal fashion
  - Expectation is that most customers will roll the conversion over a period of days/weeks/months

Bind / DDL / Online REORG break in to persistent threads

- DB2 11 delivers a break-in mechanism for persistent RELEASE(DEALLOCATE) threads
  - Persistent thread automatically detects operations that would like to break in, this detection happens at commit time
  - If detected, then RELEASE(DEALLOCATE) will behave like RELEASE(COMMIT)
  - zPARM PKGREL_COMMIT=YES must be set (parameter is online changeable)
    - Default is YES
  - Idle threads are also handled
  - Not designed to handle long running threads that don't commit or transactions with held cursors
- Packages resume normal RELEASE(DEALLOCATE) behavior after the break-in operation completes

RAS and Usability Improvements…

- Cancel DDF Threads – new FORCE option
  - Prior command without FORCE must be issued first
  - Only DDF threads
  - z/OS 1.13 APAR OA39392 required
- DRDA SQLCancel() improvements
  - Interrupt even when waiting on locks, executing SPs, or statement forwarded to another DB2
- Open data set limit raised to 200K (retrofit to V10)
- Workfile space shortage warning new instrumentation and messages
- Restrict hybrid Join to 80% of the total RID pool
- Query parallelism dynamic adjustment to available system resources
- Virtual storage scalability improvements
  - Shared memory object increased from 128G to 1T
  - Internal max for single storage pool lifted from 2G to 4G
  - Query parallelism more robust virtual storage allocation to avoid overruns

RAS Improvements…

- Query management improvements
- Autonomics improvements
  - Automatic index pseudo delete cleanup
  - Overflow row reduction
  - Optimizer externalizes missing stats to enable automated RUNSTATS
- DDF enhanced client info fields for improved granularity
- New command option to externalize RTS stats (ACCESS DB)
- Performance monitoring improvements
  - zIIP time added to CPU trace header
  - Package detail for rollup accounting
  - Reduction in 'not accounted for' time for query parallelism
  - Accumulated transaction summary data by connection type (new IFCID 369)
  - More granular stored procedure and UDF monitoring
- Support setting of APPLCOMPAT and other special registers through profile table
  - Plan to deliver APAR in 4Q13

Query Management Improvements

- Optimizer externalization of missing statistics
- Plan management improvements - APREUSE(WARN) support
  - BIND succeeds even if access path cannot be reused for one or more statements
  - Makes mass REBIND operations more feasible with APREUSE
  - Better Explain information:
    - PLAN_TABLE describes new access path even in case of APREUSE failure
    - PLAN_TABLE.REMARKS reports APRRUSE failures
- EXPLAIN and virtual index improvements
- New zparm to control max storage allocation for sort
  - (1-128M), default=1M (same as V10)

RAS and Usability Improvements…

- Data Sharing Enhancements
- Index mgr avoid RBLDP during group restart
- Restart performance: fast log apply enabled
- SELECT from SPT01 & DBD01
- DESCSTAT BIND option
- New admin stored procedure to issue z/OS commands
- Compression dictionary handling for IFCID 306

Data Sharing Improvements

➤ Group buffer pool write-around

➤Castout enhancements

➤ CF DELETE_NAME enhancements

➤ Restart light enhancements

➤Index split performance and other indexing improvements

➤Auto LPL recovery improvements

➤Full LRSN spin avoidance

➤Avoid child 'U' lock propagation for single-member R/O case

DB2 11 Security Enhancements

- **DB2/RACF authorization control enhancements**
  - Invalidate cached authorization info when RACF changes are made. Also invalidate static SQL packages
    - New AUTHEXT_CACHEREFRESH Zparm to activate the new behavior
  - AUTOBIND, BIND, REBIND present PKG-owner ACEE to RACF
  - Dynamic SQL authorization checking:
    - When DYNAMICRULES not equal to RUN, DB2 presents AUTHID to RACF
    - DYNAMICRULES defines whether AUTHID is
      - PKG owner
      - ID that defined the routine
      - ID that invokes the routine
  - New AUTHEXIT_CHECK zparm to activate the new behavior
- **Bind plan option to ensure the program is authorized to use the plan**
  - New PROGAUTH bind option
- **Remove column masking restrictions for GROUP BY and DISTINCT**

Summary of DB2 11 Utilities Improvements

- **Availability**
  - Online data repartitioning
    - REORG REBALANCE SHRLEVEL(CHANGE)
    - Online ALTER of limit keys
  - Online REORG availability improvements
    - SWITCH phase reduction
    - Improved drain processing
  - Part level inline image copies for REORG

- **Usability**
  - Online REORG automated mapping tables
  - Improved utility parallelism and control
  - DISPLAY UTILITY enhancements

- **CPU reduction**
  - More zIIP offload for LOAD and RUNSTATS

- **Performance**
  - Faster LOAD processing
  - Inline statistics improvements, reduced need for RUNSTATS
  - Optimizer input to statistics collection
  - Reduced system resources for utilities (MRU buffer management)
  - DSNACCOX performance

Expanded Analytics Capabilities

- Query performance improvements
- Temporal data enhancements
  - Support for views
  - Special register support
  - Integrated auditing support (planned)
- Transparent archive query

Transparent Archive Query
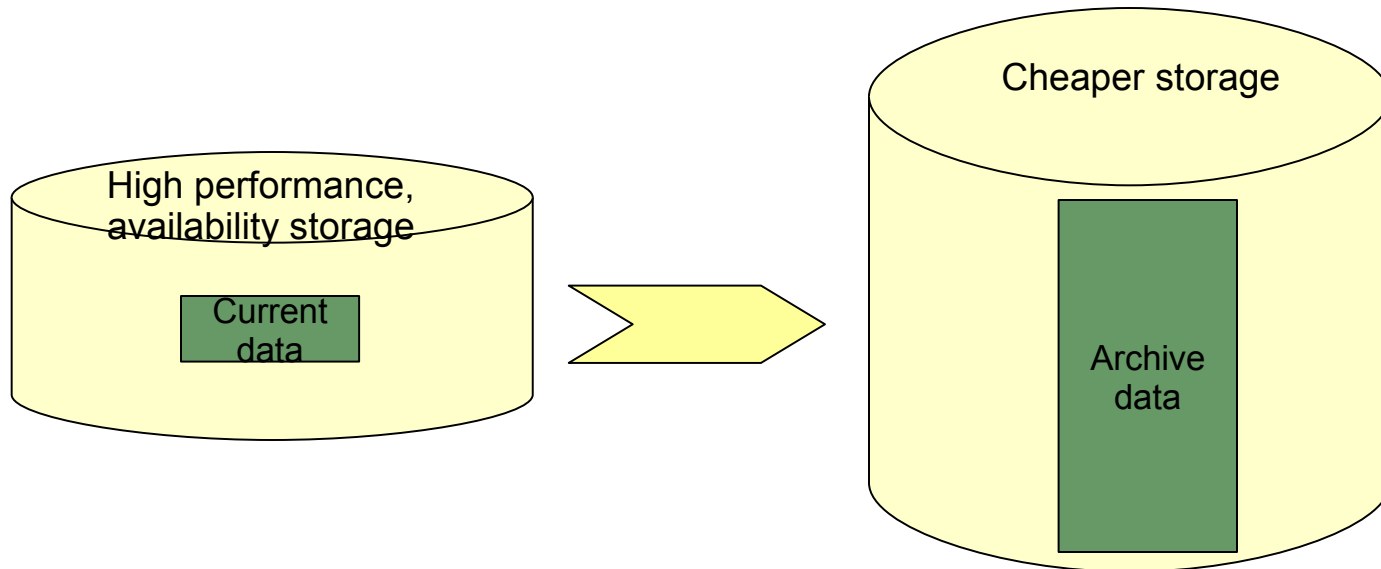
- Applications can query current + archive with no SQL changes
  - By default, data is retrieved from base table only, as usual
  - Set a new global variable when archive data is desired
  - DB2 automatically converts SQL to UNION ALL via dynamic plan switching technique (high performance)
- Archiving process is user-controlled
- Move_To_Archive global variable allows DELETEs to be automatically archived
- Leverages DB2 10 temporal constructs for archiving use cases
- Future potential for more IDAA synergy

Expanded Analytics Capabilities...

- SQL Grouping Sets, including Rollup, Cube
  - Rollup is helpful in providing subtotaling along a hierarchical dimension such as time or geography
  - CUBE is helpful in queries that aggregate based on columns from multiple dimensions
- DB2 support for IDAA V3 and V4 (rolled back to V10)
  - Support for static SQL
  - Propagating DB2 changes to the accelerator as they happen – V11 improved CDC capture performance with new IFI 306 filtering capabilities
  - Detect staleness of data via RTS
  - Reducing disk storage cost by archiving data in the accelerator and maintaining the excellent performance for analytical queries: *High Performance Storage Saver*
  - Workload Manager integration and better monitoring capabilities
  - Increasing the query off-load scope via new special register CURRENT QUERY ACCELERATION
- High performance SPSS in-database scoring via PACK/UNPACK (rolled back to v10)
- Hadoop access via table UDF
  - UDFs shipped with BigInsights
  - Uses new V11 generic table UDF capability
- JSON support (planned)

Integrating Big Data Analytics with DB2 for z/OS

■ **Much of the world's operational data resides on z/OS**	■ **Unstructured data sources are growing fast**

■ Two significant needs:

1. Merge this data with trusted OLTP data from zEnterprise data sources
2. **Integrate this data so that insights from Big Data sources can drive business actions**

New V11 features enable this

■ Connectors to allow BigInsights to easily & efficiently access DB2 data

■ DB2 is providing the connectors & the DB capability to allow DB2 apps to easily and efficiently access hadoop data sources

Integrate

IBM BigInsights

• New user-defined functions and generic table UDF capability

DB2

hadoop

Merge
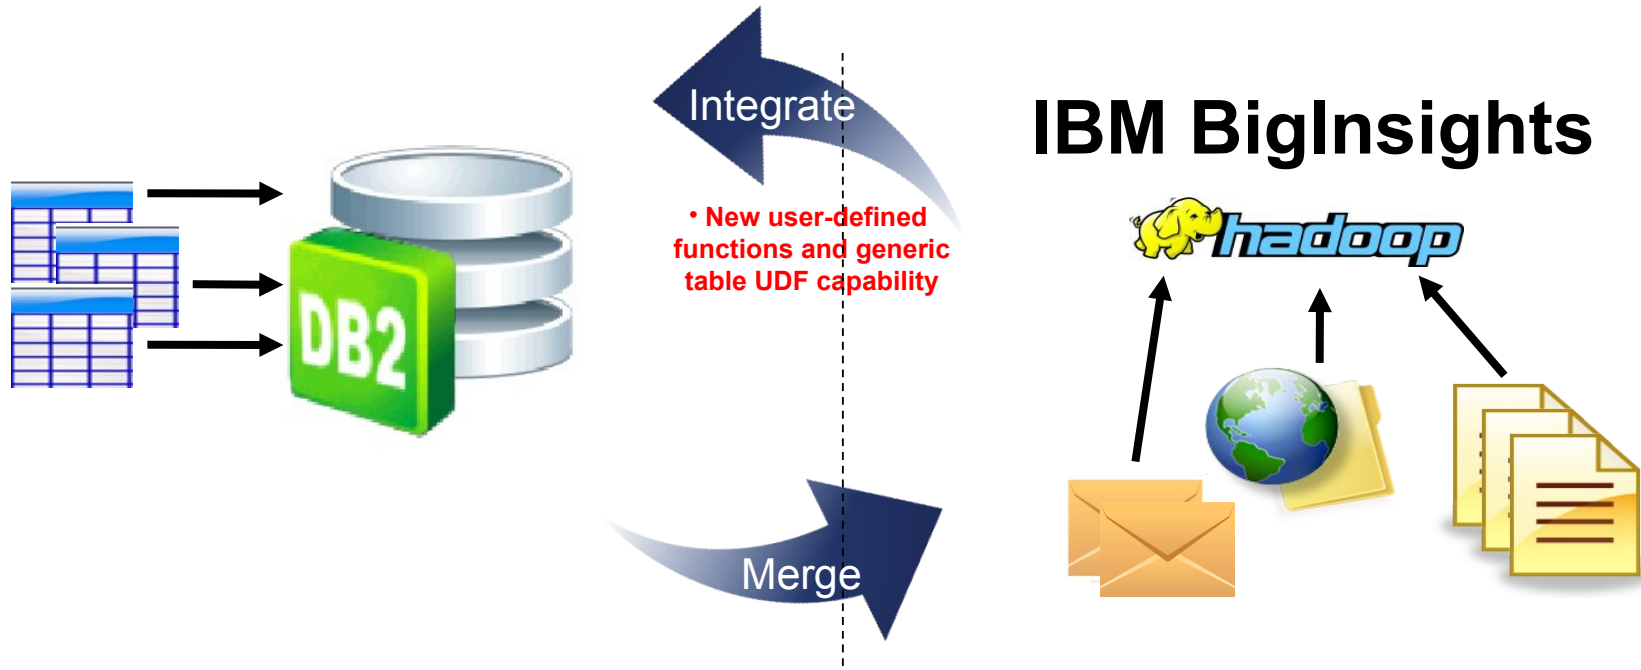
The Big Data Starting Point

Where are organizations getting the most return on Big Data projects?

"What types of data/records are you planning to analyze using big data technologies?"

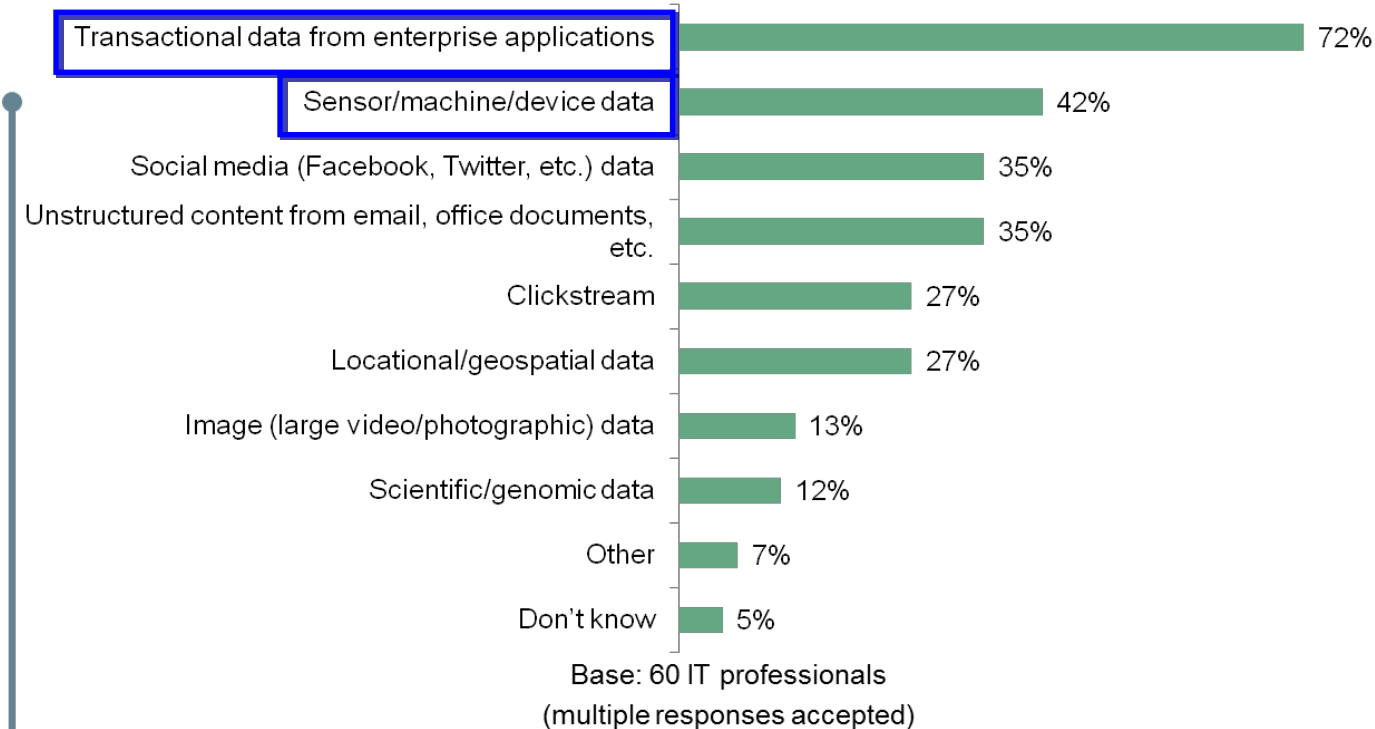| Data Type	Percentage
Transactional data from enterprise applications	72%
Sensor/machine/device data	42%
Social media (Facebook, Twitter, etc.) data	35%
Unstructured content from email, office documents, etc.	35%
Clickstream	27%
Locational/geospatial data	27%
Image (large video/photographic) data	13%
Scientific/genomic data	12%
Other	7%
Don't know	5%

Base: 60 IT professionals
(multiple responses accepted)

> (!) Most big data use cases hype its application for analysis of new, raw data from social media, sensors, and web traffic, but we found that firms are being very practical, with early adopters using it to operate on enterprise data they already have.

Source: 2012 IBM Global Big Data Online Survey

New Application Features...

- **Global variables**
  - Named memory variables that you can access and modify through SQL
  - Share relational data between SQL statements
    - Without the need for application logic to support the data transfer
- **SQLPL improvements (performance, manageability, function)**
  - Autonomous transactions
  - Array data type support
- **Alias support for Sequence objects**
  - Private alias, as currently supported for tables/views
  - Or new public alias support, enabled only for sequence objects
    - Implicit SYSPUBLIC qualifier
- **Row/Column Access Control UNION/UNION ALL support**
- **Unicode column support for an EBCDIC table (planned)**
- **BIND support for DBRMs with long & mixed cased names in zFS input files**
- **Provide REST UDFs as DB2 samples on DeveloperWorks**

XML Enhancements

- New Features
  - Basic xQuery (retrofit to v10, PM47617, PM47618)
  - COBOL samples for XML (published on Developerworks website)
- Features Enhancements
  - Implicitly add doc node during insert/update
  - Crossloader support
  - Fix error reporting position predicate
  - Support xquery constructor as the source expression of insert and replace
- Performance Enhancements
  - Binary XML validation *(retrofit to DB2 V10)*
  - Partial validation after update
  - Date/Time Predicate Pushdown
  - XQuery(FLWOR) and XMLQUERY enhancement
  - Optimize Index Search Keys
  - XML Operator Improvements, use less storage and CPU
  - XQuery deferred construction
  - XMLTABLE pushdown cast
  - Avoid validation of validated binary XML data during LOAD

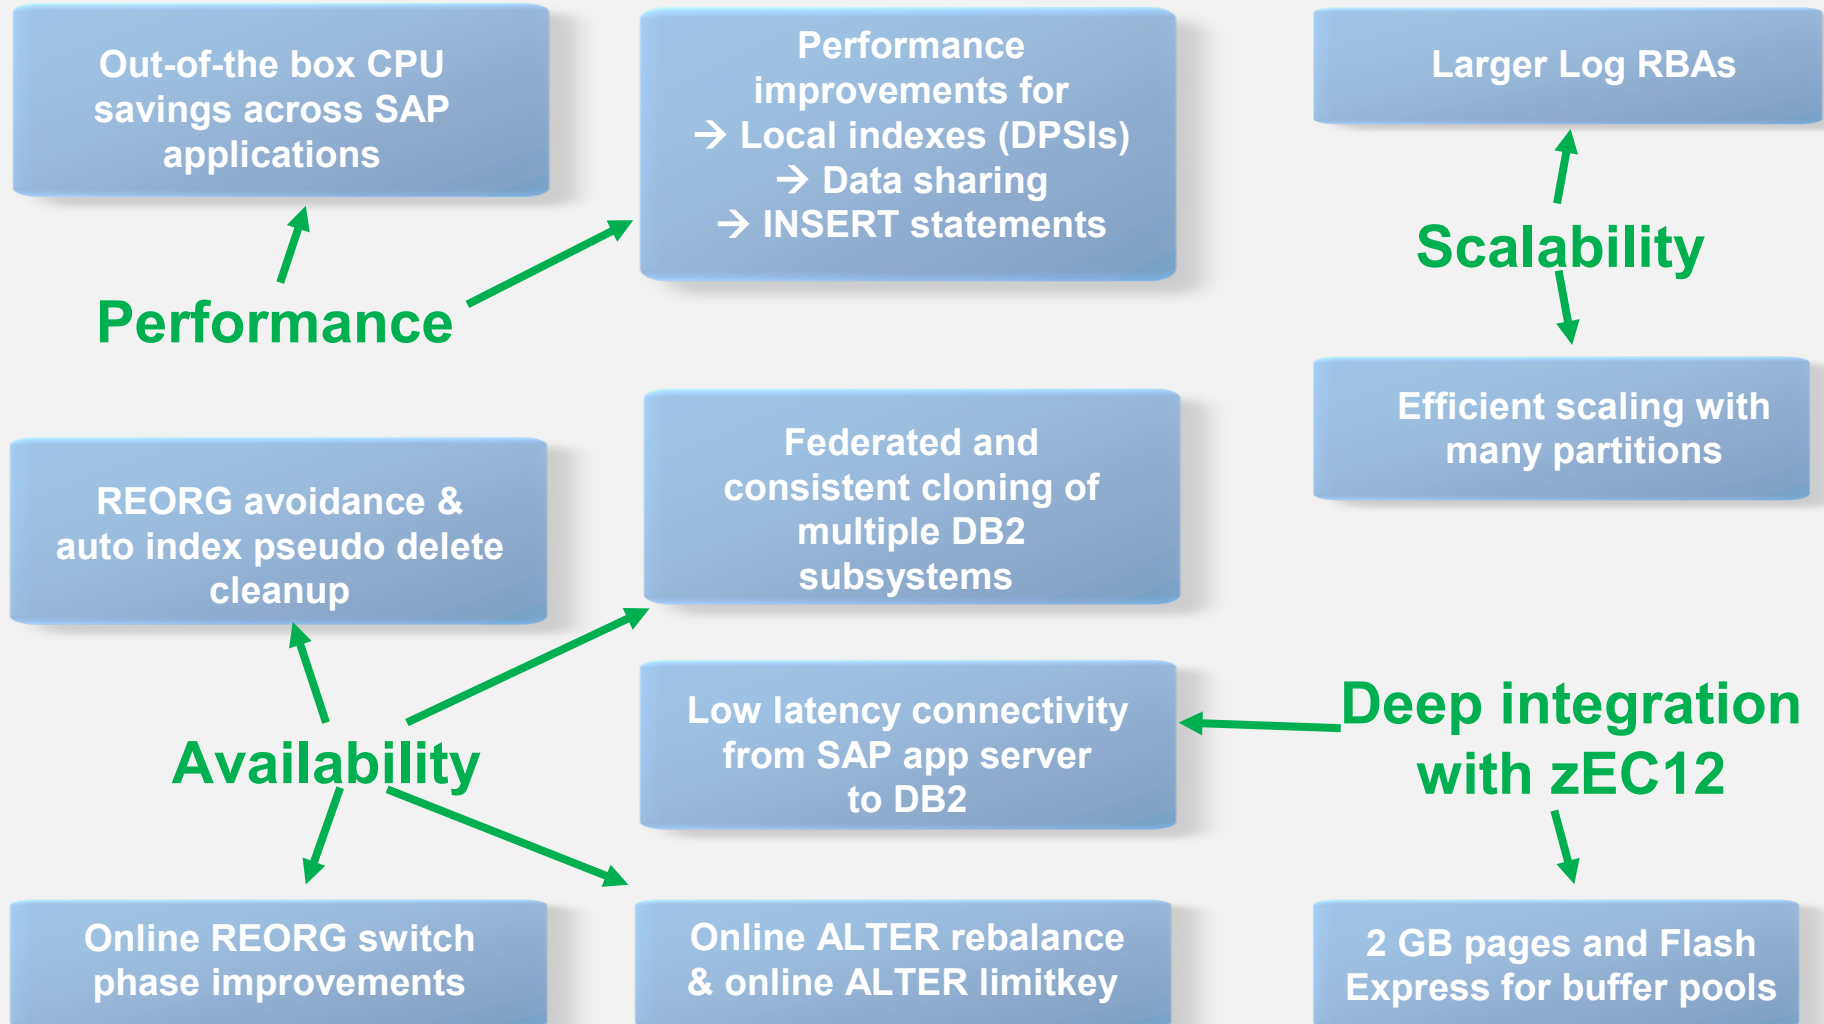# Easier DB2 Version Upgrade – application compatibility

- New DB2 releases can introduce SQL behavior changes which can break existing applications
  - For example, changes for SQL standards compliance
  - Example: DB2 10 CHAR function with decimal input no longer returns leading zeros when there is a decmial point
- Application Compatibility (APPLCOMPAT) – new option for enforcement
  - Provide mechanism to identify applications affected by SQL changes
  - Provide seamless mechanism to make changes at an application (package) level or at a system level
    - This mechanism will enable support for up to two back level releases (N-2)
    - The release after DB2 10 will be the initial deployment of this capability
    - DB2 10 will be the lowest level of compatibility supported

Easier DB2 Version Upgrade...

- Faster ENFM processing
  - Lab measurement showed 18x faster in V11 vs. V10 using a large customer catalog
  - Note: V11 ENFM performance is sensitive to size of SYSLGRNX. Consider running MODIFY RECOVER to clean up old entries if SYSLGRNX is very large
- Access path stability improvements
- Higher code quality stability levels
- New SQL Capture/Replay tooling can help testing of DB2 version upgrades

March 2013 : Start of DB2 11 Early Support Program
Again, largely driven by features for SAP

Out-of-the box CPU savings across SAP applications

Performance improvements for
→ Local indexes (DPSIs)
→ Data sharing
→ INSERT statements

Larger Log RBAs

Performance

Scalability

Efficient scaling with many partitions

REORG avoidance & auto index pseudo delete cleanup

Federated and consistent cloning of multiple DB2 subsystems

Availability

Low latency connectivity from SAP app server to DB2

Deep integration with zEC12

Online REORG switch phase improvements

Online ALTER rebalance & online ALTER limitkey

2 GB pages and Flash Express for buffer pools

52

DB2 11 Planning

- Dual mode migration (CM, ENFM, NFM)

- DB2 10 is the platform for migration

- z/OS 1.13 or above. z10 or above.

- No pre-V9 bound packages

- DB2 Connect V10.5 FP2 is the recommended level for V11
  - This level is required to exploit most new V11 features
  - Any in-service level DB2 Connect supports V11
  - Seamless migration DB2 Connect V9.7 FP6 or V10.1 FP2

- Sysplex query parallelism support is removed

DB2 Cypress: Early Thoughts

- **Out-of-the-box performance improvements**
  - No application or DBA changes needed
  - In-memory optimizations to reduce CPU
  - HW/SW integration

- **Ease of use improvements**
  - Application developers: more transparent SQL performance optimization, SQL, XML enhancements
  - DBAs: easier SQL tuning, large table management improvements, improved system autonomics

- **RAS Improvements**
  - More online schema change capabilities, enhanced Parallel Sysplex capabilities, utilities improvements

- **Expanded SQL and analytics capabilities**

DB2 for z/OS Request for Enhancements (RFE)
More Effectively Communicating with Customer on Requirements

- **Historically…** internal IBM tool (FITS). No direct customer access to requirements

- **Going forward** with **DB2 for z/OS RFE** you can:
  - Directly manage/track your requirements – greater accessibility
  - Access, vote, comment, and watch other public requirements
  - Directly interact with DB2 development

- DB2 for z/OS RFE Community link: http://www.ibm.com/developerworks/rfe/infomgmt/

- Online help and tutorials available from the RFE Community – including YouTube videos

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

Customers

RFE Community

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

- Watching
- Notifications (email and RSS)
- Groups
- Escalation process
- And more

Brand Development

Typical Utilization for Servers

Windows: 5-10%   Unix: 10-20%   **System z: 85-100%**

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

Thank You

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

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