



DB2 10 for z/OS and Beyond

John Iczkovits
IBM

August 13, 2013
Session Number 14295



Please note



SHARE
Software Hardware Architecture Research Evaluation

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 and the new zEnterprise EC12 machine
- DB2 11 preview

DB2 for z/OS

The most robust and cost effective data server

DB2

DB2 9

DB2 10



Efficiency

- Deep synergy with System z
- HW Compression
- Consolidation

- Up to 20% utility CPU savings
- Compress indexes, save 50% disk
- 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



Resilience

- 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



Growth

- Near-linear scalability
- Flexible development
- Warehousing capabilities

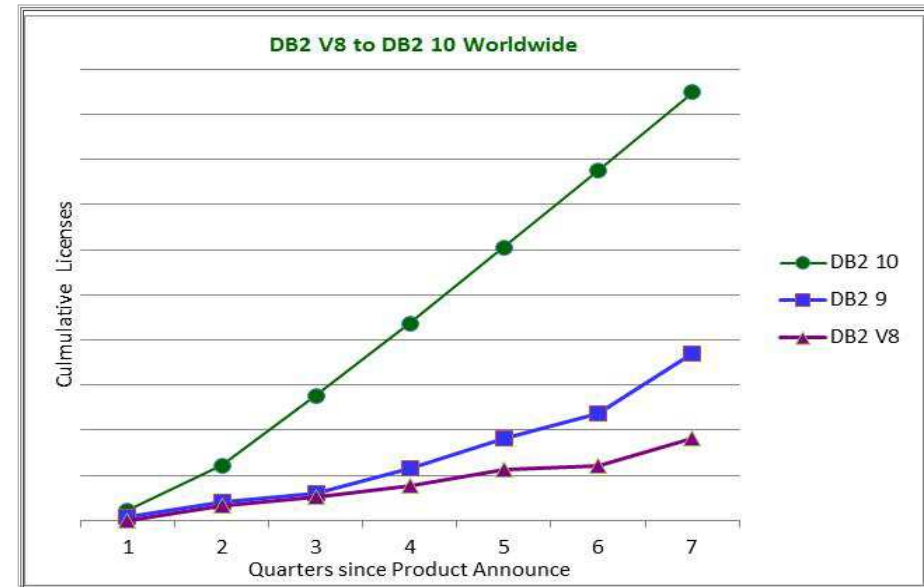
- Seamless integration of XML and relational
- Improved SQL
- Native SQL PL
- Partition by growth
- OLAP expressions

- Enhanced query parallelism
- Advanced query acceleration
- Query optimizations
- More SQL compatibility
- Improved pureXML and SQL PL

DB2 10 for z/OS Snapshot

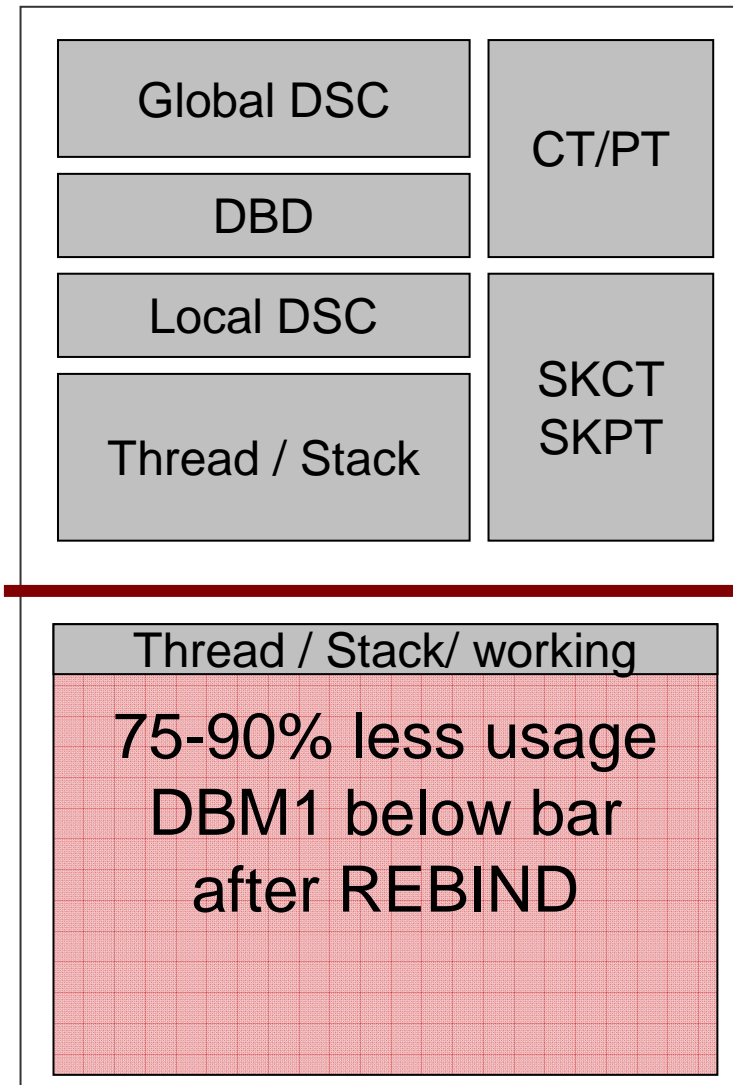
- **Fastest uptake**
 - **+2x** customers vs. V9
 - **+2.5x** licenses vs. V9
 - **25%** coming from DB2 V8

- **Adoption Driven by:**
 - **Performance improvements without application changes**
 - **Virtual Storage Constraint relief for more threads**
 - **Security, RAS improvements**
 - **Bitemporal data**



DB2 10 Virtual Storage Constraint Relief DB2 10

- DBM1 below 2GB
 - 75-90% less usage in DB2 10 compared to DB2 9
 - Some of working storage (stack, xproc storage) stays below 2GB
- Larger number of threads
 - Possible data sharing member consolidation
- Improve CPU with storage
 - More release deallocate
 - Larger MAXKEEPD values for KEEP DYNAMIC=YES


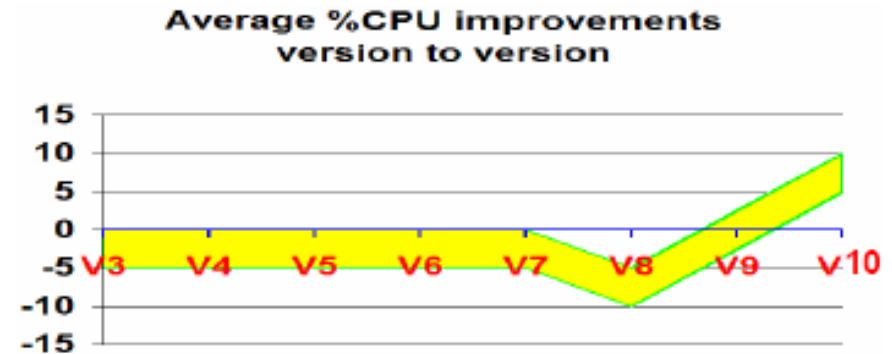


DB2 10 Performance

- 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




CPU Percent reduced from DB2 9


DB2 High Performance DBAT

- **High Performance DBATs reduce CPU consumption by**
 - Supporting RELEASE(DEALLOCATE) to avoid repeated package allocation/deallocation
 - Avoids processing to go inactive and then back to active
- **Enabling High Performance DBAT, e.g. for WebSphere**
 - BIND client packages into different collection coll2 with RELEASE(DEALLOCATE)
 - BIND other frequently executed packages with RELEASE(DEALLOCATE)
 - Set -MODIFY DDF PKGREL(BNDOPT) to enable
 - In WAS datasource property definition point to new collection
 - E.g. jdbcCollection=coll2

Index to Data Access Path vs. Hash Access



■ = Page in Bufferpool
■ = Page Read from Disk




- Index->Data access
 - Traverse down Index Tree
 - For a 5 Level Index
 - 6 GETP
 - 2 I/O's
 - 5 index page searches

- Hash Access
 - Locate a row without having to use an index
 - Single GETP in most cases
 - 1 Synch I/O in common case
 - Potential for significantly improved performance

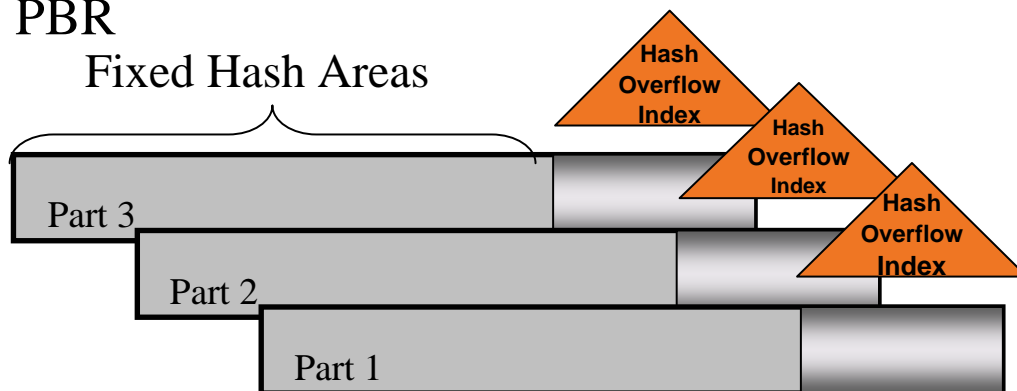
Hash Access and Hash Space

PBG



- Optimal to get from fixed area
 - 1 getpage, 1 I/O
- Overflow
 - 3 getpages, 2-3 I/Os

PBR



- Use REORG with **AUTOESTSPACE YES** unless you know better
- Real Time Statistics (RTS)
 - # of overflow
 - TOTALENTRIES**
 - **TOTALENTRIES / TOTALROWS < 10%**
- FREEPAGE is not valid for HASH space but PCTFREE is honored

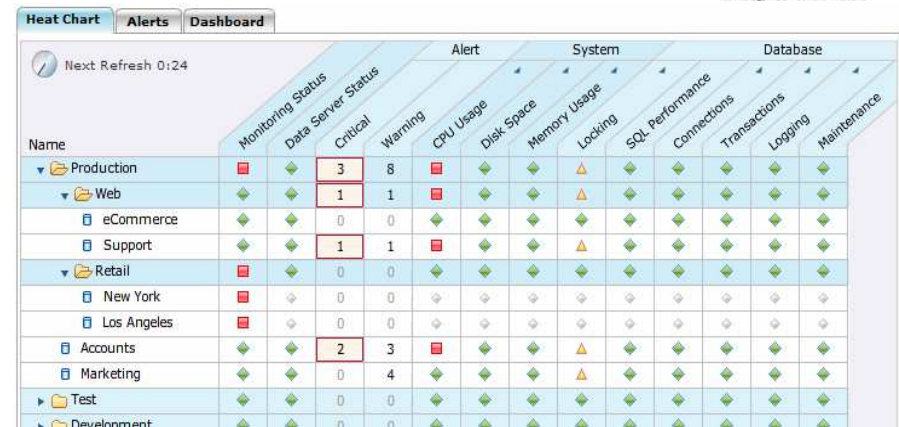
Hashing Summary



- 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

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



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

Manual invocation of

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

DB2 10: Protecting DB2 from bad applications



- **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
- 12 different types of filtering criteria including
 - IP address, domain name, location name, workstation name, app name, userid, product ID, package name, collection ID, role name, authid
- Can add default filtering: an asterisk is accepted on any filtering type to match all connections/threads
- Dynamic Location Alias
 - Create and manage subsets of member in a data sharing group
 - Dynamically change which members are used to balance client workloads
 - Without stopping and restarting DDF or DB2
- Use InfoSphere Optim Configuration manager to manage profiles

DB2 10 Access Path Comparison (APCOMPARE)



- APCOMPARE
 - “Tell me if static SQL statements had changes in access paths”
 - Optionally, stop the BIND/REBIND if there were changes
- New option on
 - BIND PACKAGE
 - REBIND PACKAGE and REBIND TRIGGER PACKAGE
- For all statements in the package ...
 - Load old access path from the “EDB”
 - Optimizer generates new access path, as usual
 - Compare old access path with new access path
 - Report results via messages / PLAN_TABLE output
 - Determine REBIND success vs failure

DB2 10 APCOMPARE option values



- APCOMPARE(NONE/NO)
 - No comparison performed
 - This is the default
- APCOMPARE(WARN)
 - RC = 4
 - DB2 will continue processing the package
- APCOMPARE(ERROR)
 - RC = 8
 - DB2 will terminate the processing of the package

DB2 10 Access Path Reuse (APREUSE)



- APREUSE - “Do the BIND/REBIND but try to avoid access path changes”
 - At migration from DB2 9 to 10
 - After service fixes that require the regeneration of runtime structures (++HOLDS directives)
 - To bring invalid packages back to life
 - Due to application changes (BIND PACKAGE)
- New option on
 - BIND PACKAGE
 - REBIND PACKAGE and REBIND TRIGGER PACKAGE
- For all statements in the package ...
 - Load old access path (“EDB”)
 - Feed “EDB” as hint to the optimizer
 - As a final check, compare old access path with new
 - Report results via messages / PLAN_TABLE output
 - Determine REBIND success vs failure
- APREUSE implicitly turns on APCOMPARE



APREUSE option values



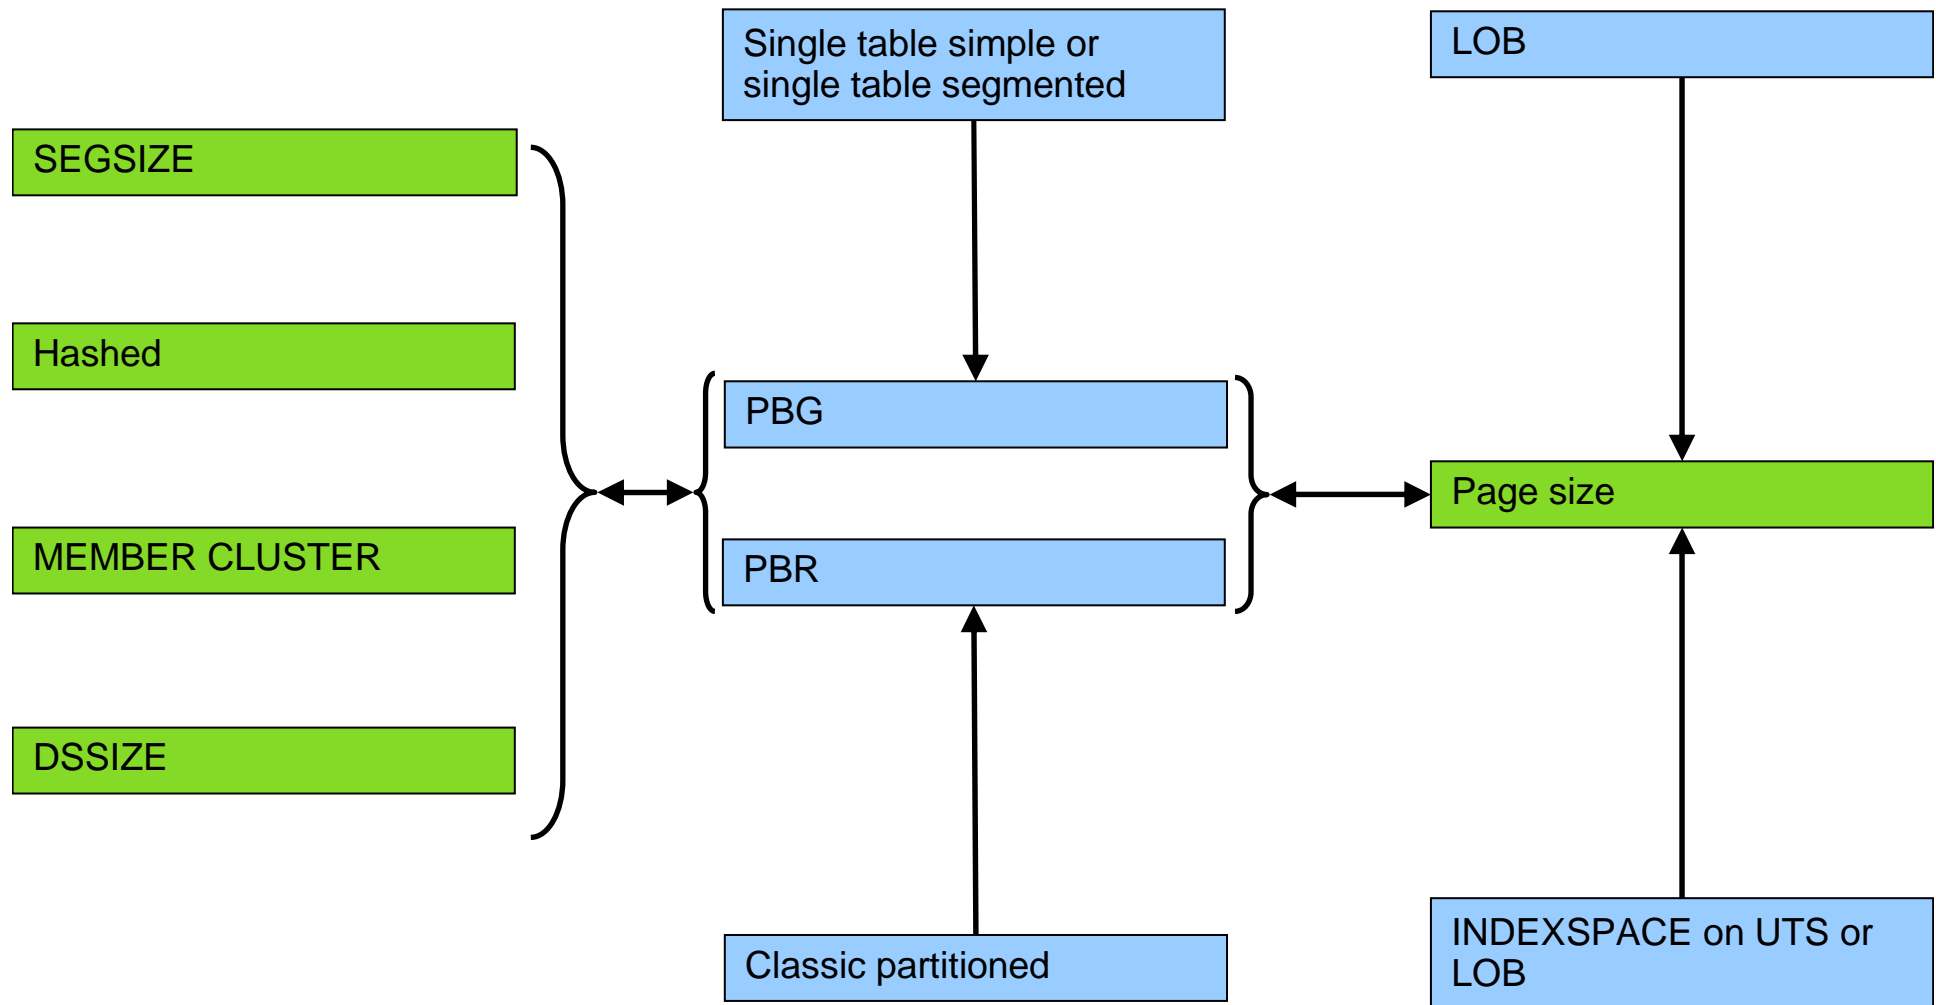
- APREUSE(NONE/NO)
 - No reuse performed
 - This is the default
- APREUSE(ERROR)
 - RC = 8
 - DB2 will terminate the processing of the package

DB2 10 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
 - 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

Online Schema – DB2 10



DB2 10 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
- 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 to run REORG
 - List prefetch of index leaf pages based on non-leaf information for range scans
 - Row level sequential detection
 - Auto compress on insert
 - New REORGCLUSTSENS RTS column: If no clustering-sensitive queries then avoid REORG to restore clustering
- Improved performance for part-level REORG with NPIs & REORG INDEX
 - Index list prefetch results in up to 60% elapsed time reduction
- REORG SHRLEVEL CHANGE for all cat/dir pagesets
- 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


DB2 10: Backup & Recovery

- Dataset-level Flashcopy support
 - COPY, RECOVER, REORG, LOAD, REBUILD INDEX, REORG INDEX
 - New zparms & utility parms to govern
 - Virtually eliminate CPU & elapsed time for large pagesets
 - Create transaction-consistent image copies from COPY SHRLEVEL CHANGE
 - Create partition-level inline image copies from REORG

CPU time per object (z10)




Elapsed time per object (z10)




DB2 10: Backup & Recovery

- VERIFYSET option to fail PIT recovery if entire set not included
 - Base, LOB, XML & history objects
- ENFORCE NO option to avoid CHKP/ACHKP on PIT recovery of subset of set
 - Improved performance due to avoidance of set checking (RI, aux)
- Fast recovery to point in time through new BACKOUT option
 - Include indexes in RECOVER list to avoid the need to rebuild them
 - Indexes must be COPY YES
 - No imagecopy required though
 - LOBs requires APAR PM45650
 - PIT recovery always with consistency since V9



Changes in DB2 10 catalog & directory

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



DB2 10 Business Security & Compliance

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



Temporal Based Analysis

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



System Time Query

employees

EmpID	Dept	System_start	System_end
12345	M15	05/31/2000	12/31/9999

employees_history

EmpID	Dept	System_start	System_end
12345	J13	11/15/1995	01/31/1998
12345	M24	01/31/1998	05/31/2000
67890	K25	11/15/1995	03/31/2000

1. Which department is employee 12345 in?

```
SELECT Dept
FROM employees
WHERE EmpID=12345
```

M15

Note that without the FOR SYSTEM_TIME clause, it defaults to the current table only.

2. Which department was employee 12345 in on 12/01/1997?

```
SELECT Dept
FROM employees FOR SYSTEM_TIME AS OF '12/01/1997'
WHERE EmpID=12345
```

J13

3. How many departments has employee 12345 worked in?

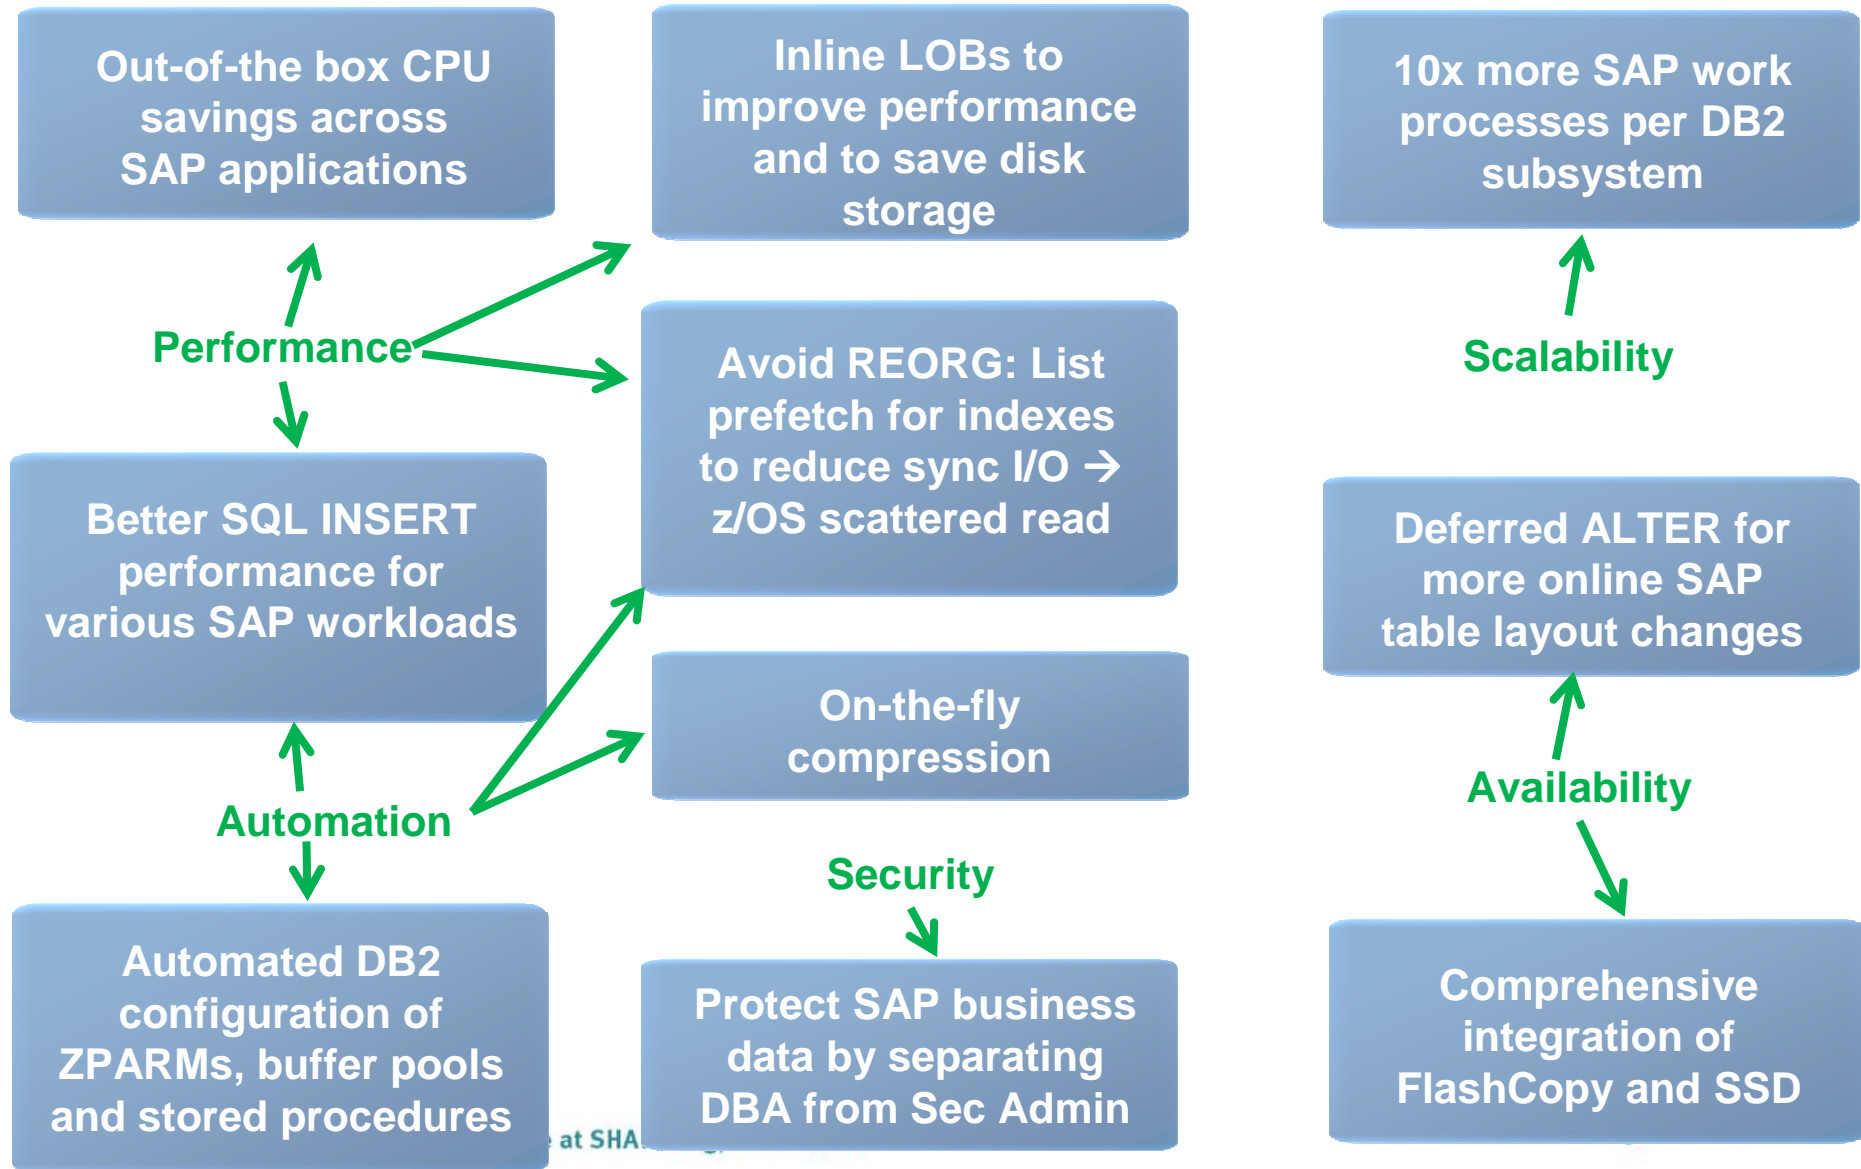
```
SELECT count(distinct Dept)
FROM employees FOR SYSTEM_TIME FROM MINVALUE TO MAXVALUE
WHERE EmpID=12345
```

3

DB2 Family 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 - Top 10 Features for SAP



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)

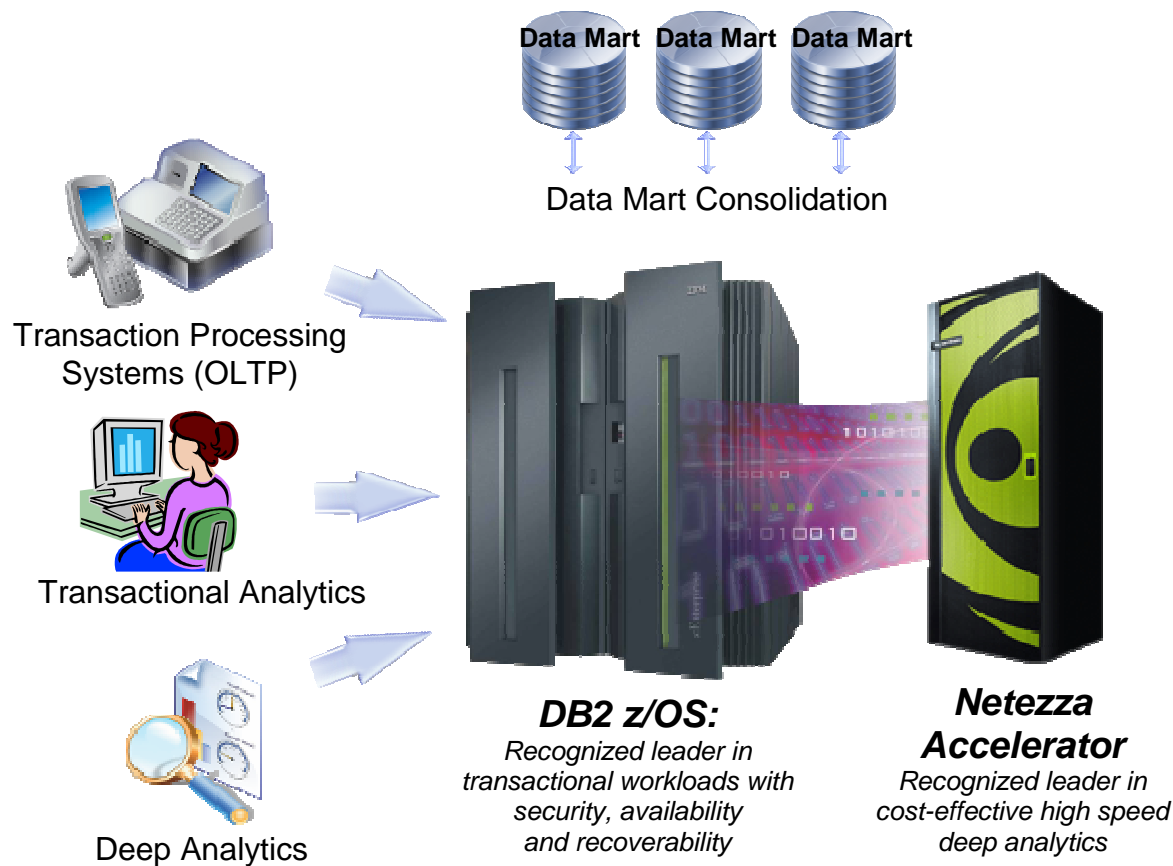
DB2 and zEnterprise EC12

- **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 perform better than pageable 4k for non-fixed BPs
 - Recommendation for best performance is still 1M fixed with sufficient real memory
 - 2GB frame support
 - Additional CPU savings, especially for very large memory
- **Transactional Memory provides further possibilities for performance gains**



Creating the Hybrid Data Server - Netezza and Z

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



Best in OLTP and Transactional Analytics

Industry recognized leader in mission critical transaction systems

Best in Deep Analytics

Proven appliance leader in high speed analytic systems

Best in Consolidation

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

Together:

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

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)

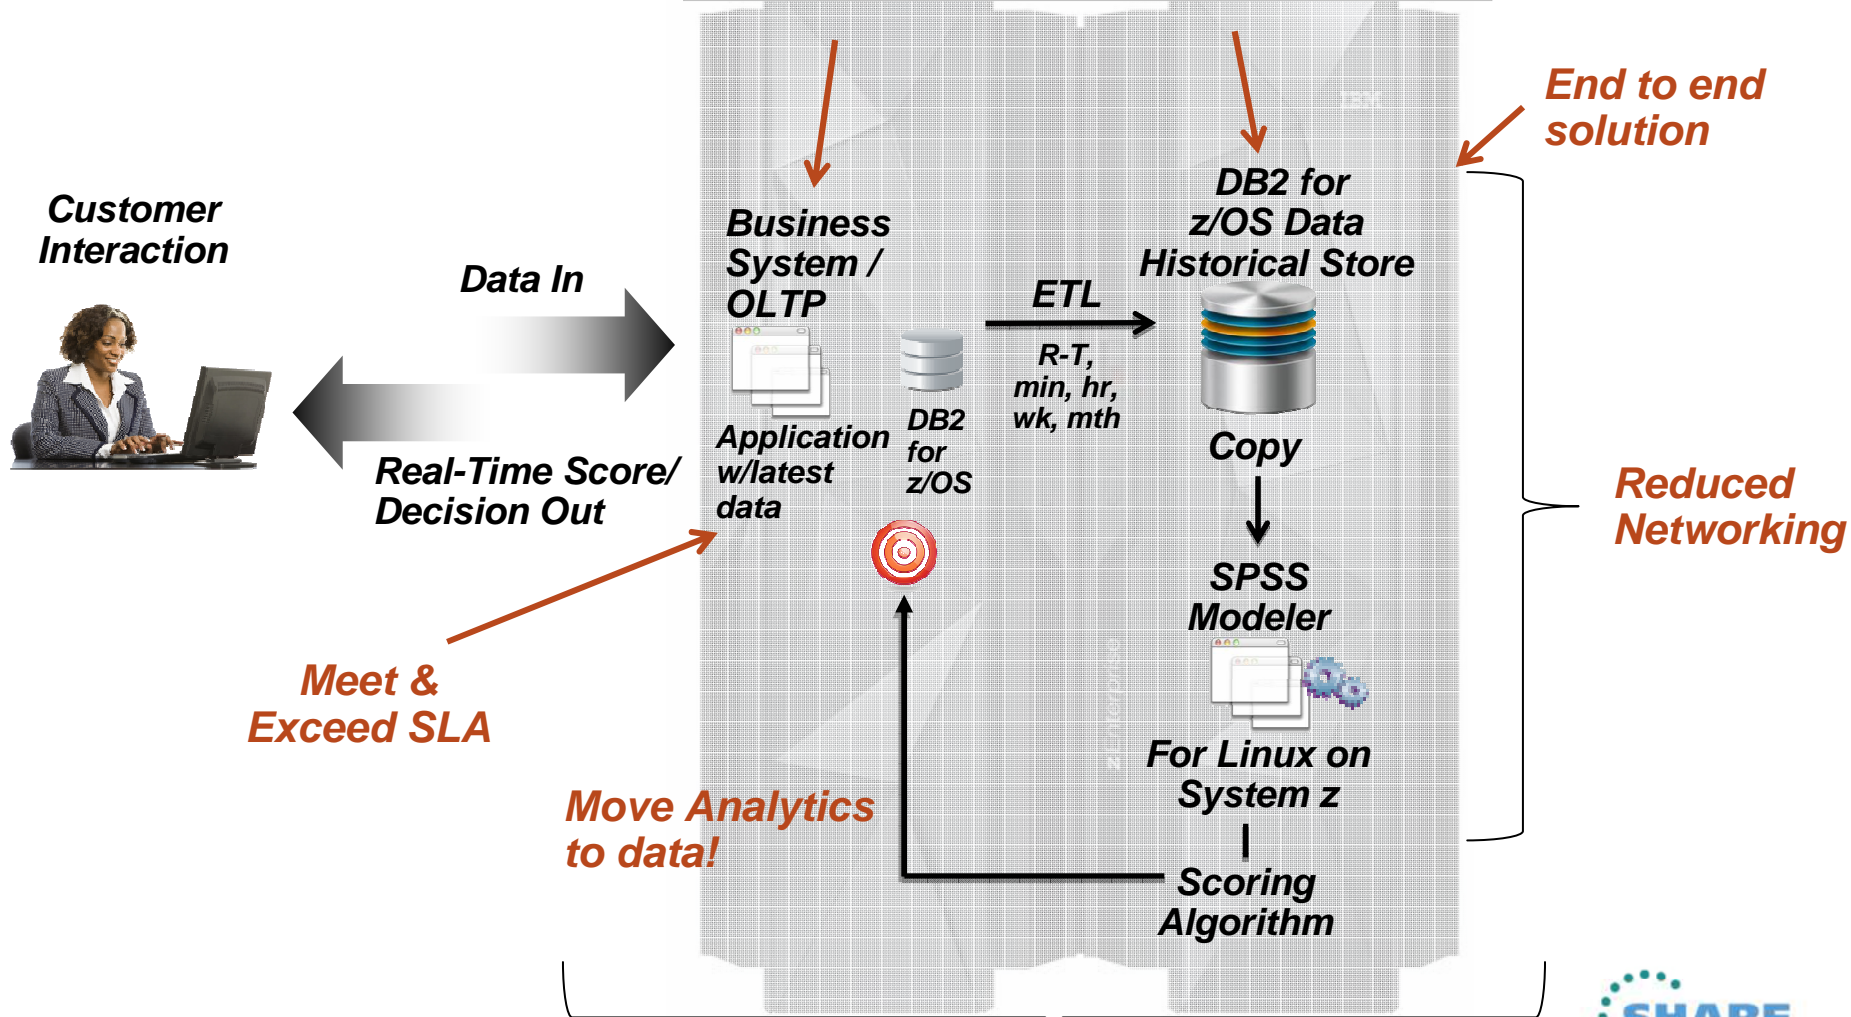
• UNLOAD Lite

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

Modeler 15 Real-time Scoring with DB2 z/OS



Support for both in-transaction and in-database scoring on the same platform



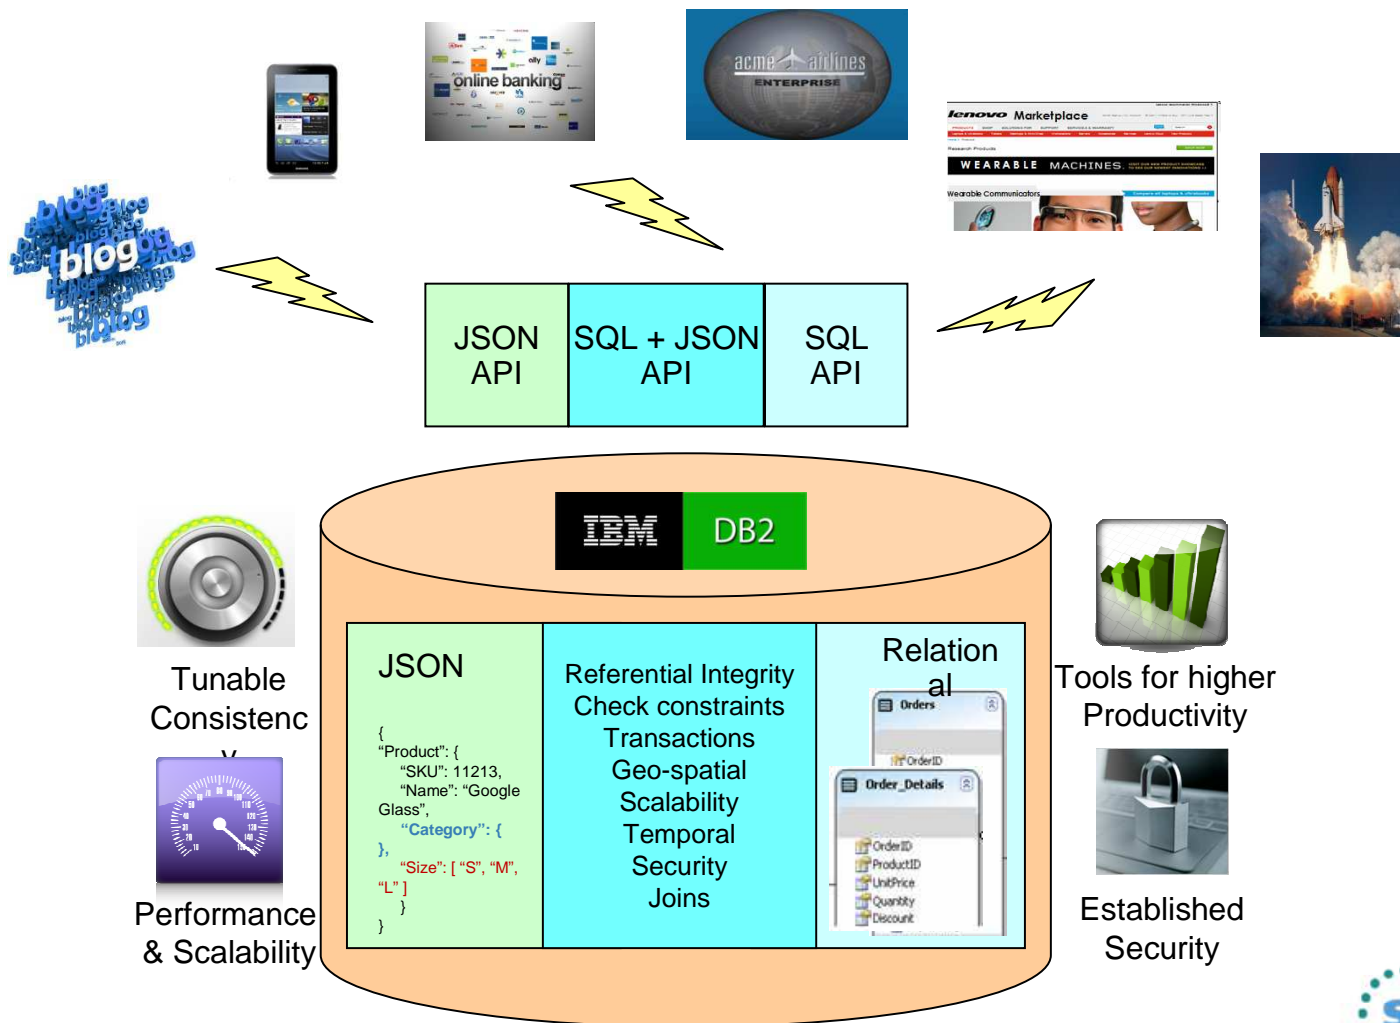
Complete your sessions evaluation online at SHARE.org/BostonEval



Consolidates Resources

JSON Database Technology Preview

Providing the best of both worlds



Complete your sessions evaluation online at SHARE.org/BostonEval

IBM Confidential until April 3, 2013



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 for z/OS Customer Trends



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

DB2 11 Major Themes



- **Performance Improvements**
 - Improving efficiency, reducing costs, no application changes
 - 0-5% for OLTP, 5-15% for update intensive batch
 - 5-20% for query workloads
 - Less overhead for data de-compression
 - Exploitation of new zEC12 hardware features
- **Continuous Availability Features**
 - Improved autonomies which reduces costs and improves availability
 - Making online changes without affecting applications
 - Online REORG improvements, less disruption
 - DROP COLUMN, online change of partition limit keys
 - Extended log record addressing capacity (1 yottabyte)
 - BIND/REBIND, DDL break into persistent threads
- **Enhanced business analytics**
 - Faster, more efficient performance for query workloads
 - Temporal and SQLPL enhancements
 - Transparent archiving
 - SQL improvements and IDAA enhancements
- **Simpler, faster DB2 version upgrades**
 - No application changes required for DB2 upgrade
 - Access path stability improvements
 - Product stability: support pre GA customer production



Extended RBA Problem Statement



- DB2's Relative Byte Address (RBA) for logging is 6 bytes
 - Gives 256TB of log record addressing capacity per DB2 subsystem/member
- With heavy sustained logging rates, DB2 can exhaust the 6-byte RBA
 - DSNJ032I and DSNJ033E warning messages
 - Alert-level = 'WARNING' when RBA reaches x'F00000000000'
 - Alert-level = 'CRITICAL' when RBA reaches x'FFFF00000000'
 - Manual recovery actions are needed
 - Data Sharing: shut down the affected member and start a new member in its place
 - Non Data Sharing: reset all PGLOGRBA values back to zero (extended outage)
 - Documented in the DB2 Administration Guide
 - If alert-level reaches 'CRITICAL' then DB2 terminates to protect data integrity and force recovery actions
 - Reason code 00D10251
 - ACCESS(MAINT) restart allowed to prepare for recovery actions

Extended LRSN Problem Statement



- The data sharing Log Record Sequence Number (LRSN) is derived from the 8-byte time-of-day clock which hits end of range in 2042
- However, some data sharing groups have a non-zero LRSN “delta” which gets added to the TOD clock
 - If a non-zero “delta” exists, then the LRSN will hit end of range prior to 2042
 - Use DSNJU004 to determine if you have a non-zero LRSN delta value
 - A “delta” value could be set when data sharing is enabled or re-enabled
 - Whenever the end-of-log RBA of the enabling member is past the TOD clock
- Some non data sharing customers have enabled data sharing to circumvent RBA nearing end-of-range
 - This would cause a non-zero LRSN delta, so LRSN hits end of range before 2042
- 6-byte LRSN value has precision to only 16 microseconds
 - Can cause LRSN ‘spinning’ which burns extra cpu and aggravated log latch contention
 - V9 NFM addresses most LRSN spin situations, and V10 NFM enhances further. But some spins still exist due to the 16 usec granularity (log latch not held, page latches are)

DB2 11 Planned Solution for Extended RBA/LRSN



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

Some Planned DB2 11 RAS Improvements



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

DB2 11 Some Planned New Application Features

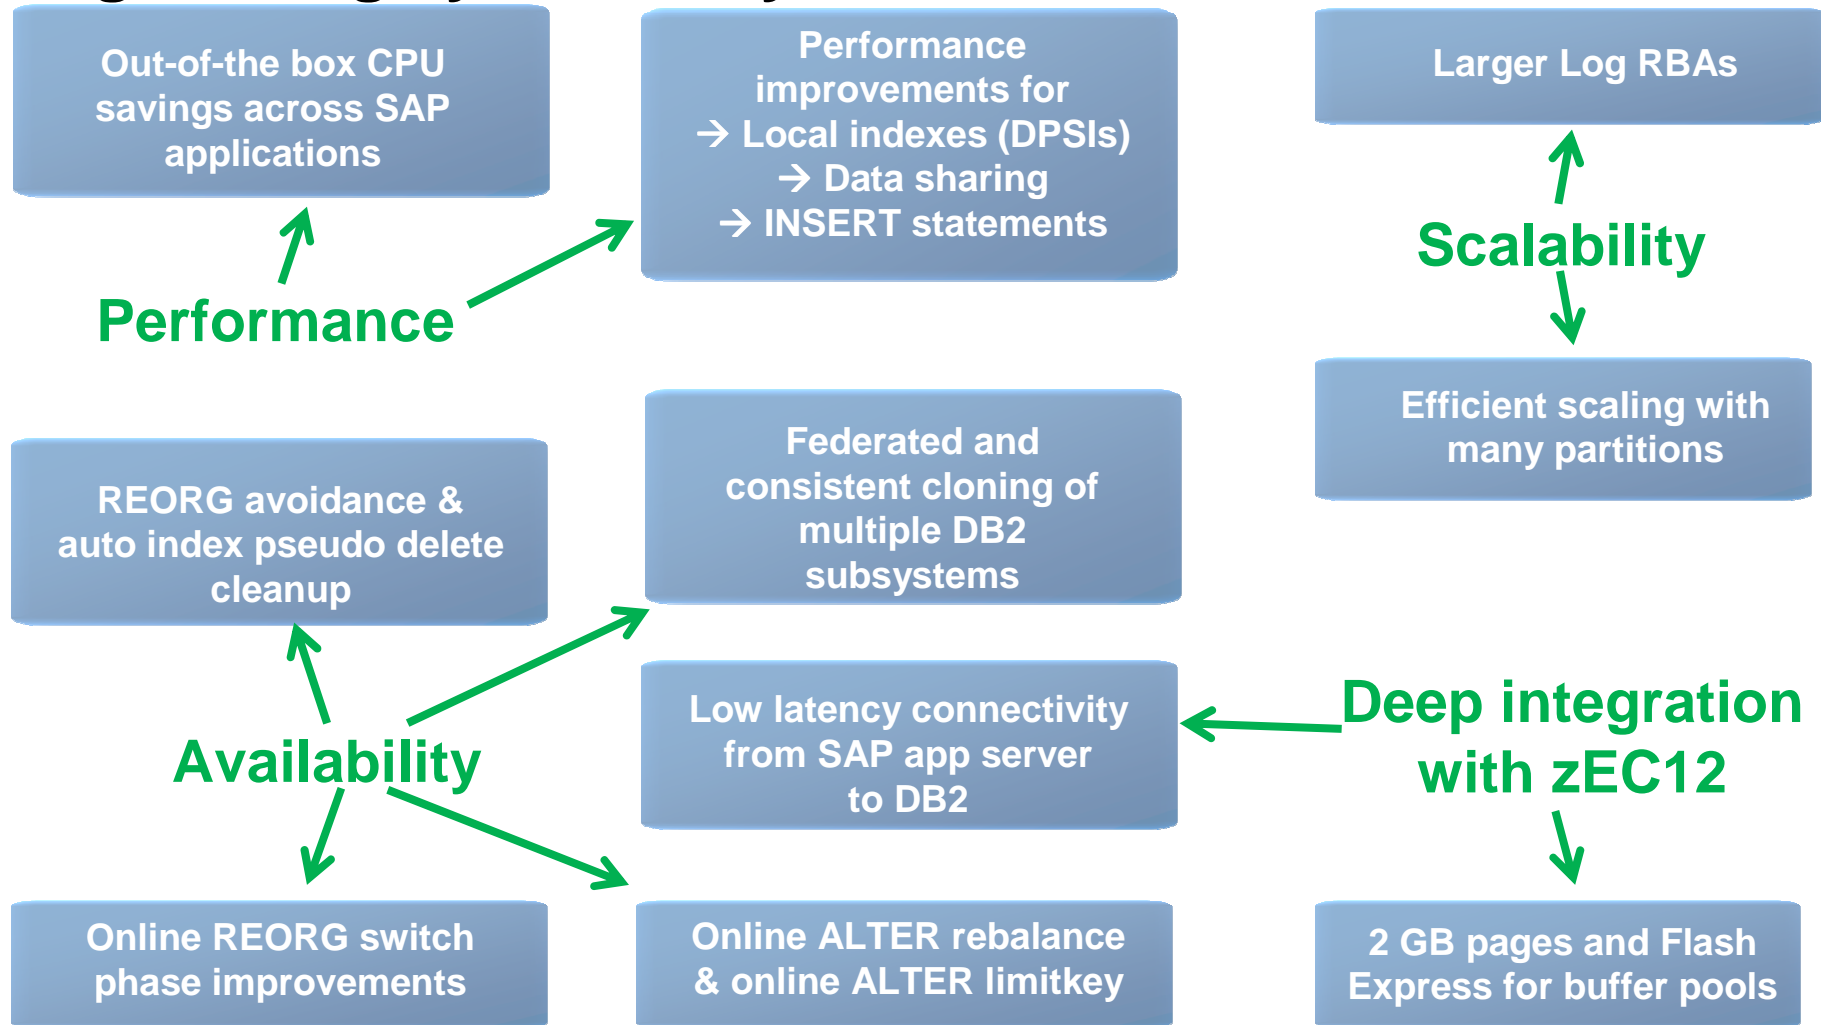


- Global variables
- SQLPL improvements (performance, manageability, function)
 - Autonomous transactions
 - Array data type support
- Alias/synonym support for Sequence Objects
- Grouping sets
- DPSI performance improvements
- DGTT performance improvements
- Temporal data enhancements
 - Support for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

- Enhancements for DB2 and Big Data integration

March 2013 : Start of DB2 11 Early Support Program

Again, largely driven by features for SAP

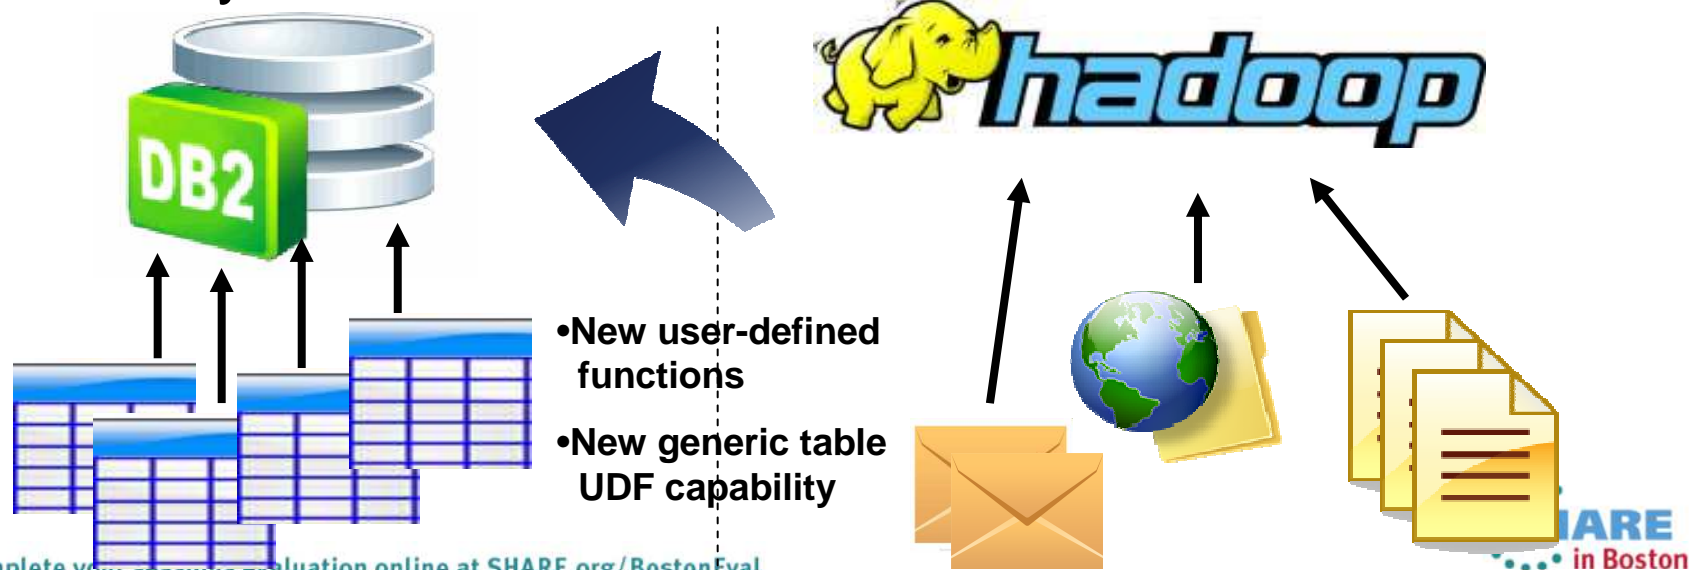


Enhancing DB2 Analytics on “z” with Big Data



Use of
unstructured
raw data is
GROWING

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



Easier DB2 Version Upgrade



- SQL Compatibility – new option for enforcement
 - Limit SQL incompatibilities when possible
 - Provide mechanism to identify applications affected by SQL changes
 - Provide seamless mechanism to make changes at an application (package) level
 - This mechanism will enable support for up to two back level releases (N-2)
 - DB2 11 will be the initial deployment of this capability
 - DB2 10 will be the lowest level of compatibility supported
- Access path stability improvements




DB2 11 Planning



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

DB2 for z/OS Request for Enhancements (RFE)



More Effectively Communicating with Customer on Requirements

SHARE

- **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



Complete your sessions evaluation online at SHARE.org/BostonEval



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

