DB2 11 for z/OS Technical Overview

Session 17206

August 10, 2015

John Iczkovits
iczkovit@us.ibm.com

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 for z/OS

- Strong uptake out of the gate
- Over 150 customers* (* As of Dec 2014)
  - Faster migration success
  - 2x faster adoption
- Out-of-the-box quality/stability
  - 68% fewer PMRs
  - 35% fewer APARs
- DB2 10
  - Withdraw from Marketing: July 6, 2015
  - End of Service: Sept 30, 2017

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
- Sysplex query parallelism support is removed
- DB2 11 Migration Planning Workshop (MPW)
  - Free, 1-day education
  - DB2 11 MPW Community on DeveloperWorks

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 (or 1B petabytes)
  - BIND/REBIND, DDL break into persistent threads
- **Enhanced business analytics**
  - Expanded SQL, XML, and analytics capabilities
  - Temporal and SQLPL enhancements
  - Transparent archiving
  - Hadoop integration, NoSQL and JSON support
- **Simpler, faster DB2 version upgrades**
  - No application changes required for DB2 upgrade
  - Access path stability improvements
  - Product quality/stability – raised the bar

DB2 11 for z/OS
Save money, save time, reduce cost

SHARE
in Orlando 2015

Impressive DB2 11 Performance Results!

DB2 11 % CPU Improvement From DB2 10
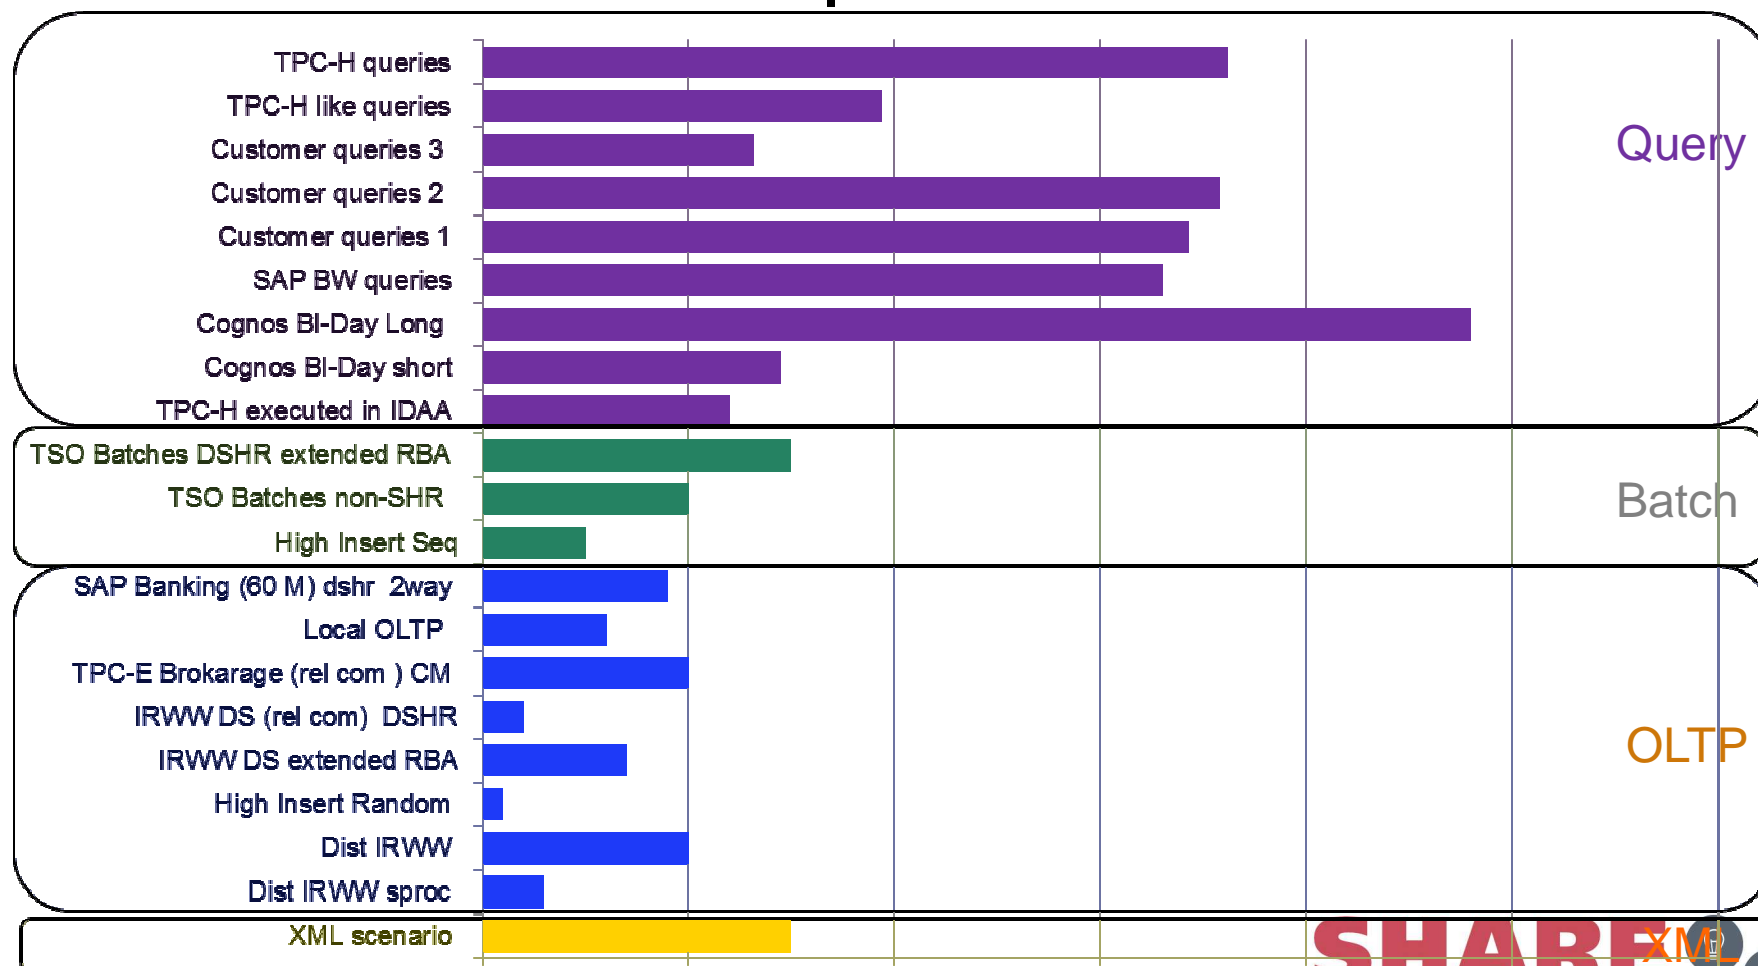
Query
- TPC-H queries
- TPC-H like queries
- Customer queries 3
- Customer queries 2
- Customer queries 1
- SAP BW queries
- Cognos BI-Day Long
- Cognos BI-Day short
- TPC-H executed in IDAA

Batch
- TSO Batches DSHR extended RBA
- TSO Batches non-SHR
- High Insert Seq

OLTP
- SAP Banking (60 M) dshr 2way
- Local OLTP
- TPC-E Brokarage (rel com) CM
- IRWW DS (rel com) DSHR
- IRWW DS extended RBA
- High Insert Random
- Dist IRWW
- Dist IRWW sproc

XML
- XML scenario

0% 10% 20% 30% 40% 50% 60%

TPC-H Using Static SQLPL

Performance Results - Average CL2 CPU Time
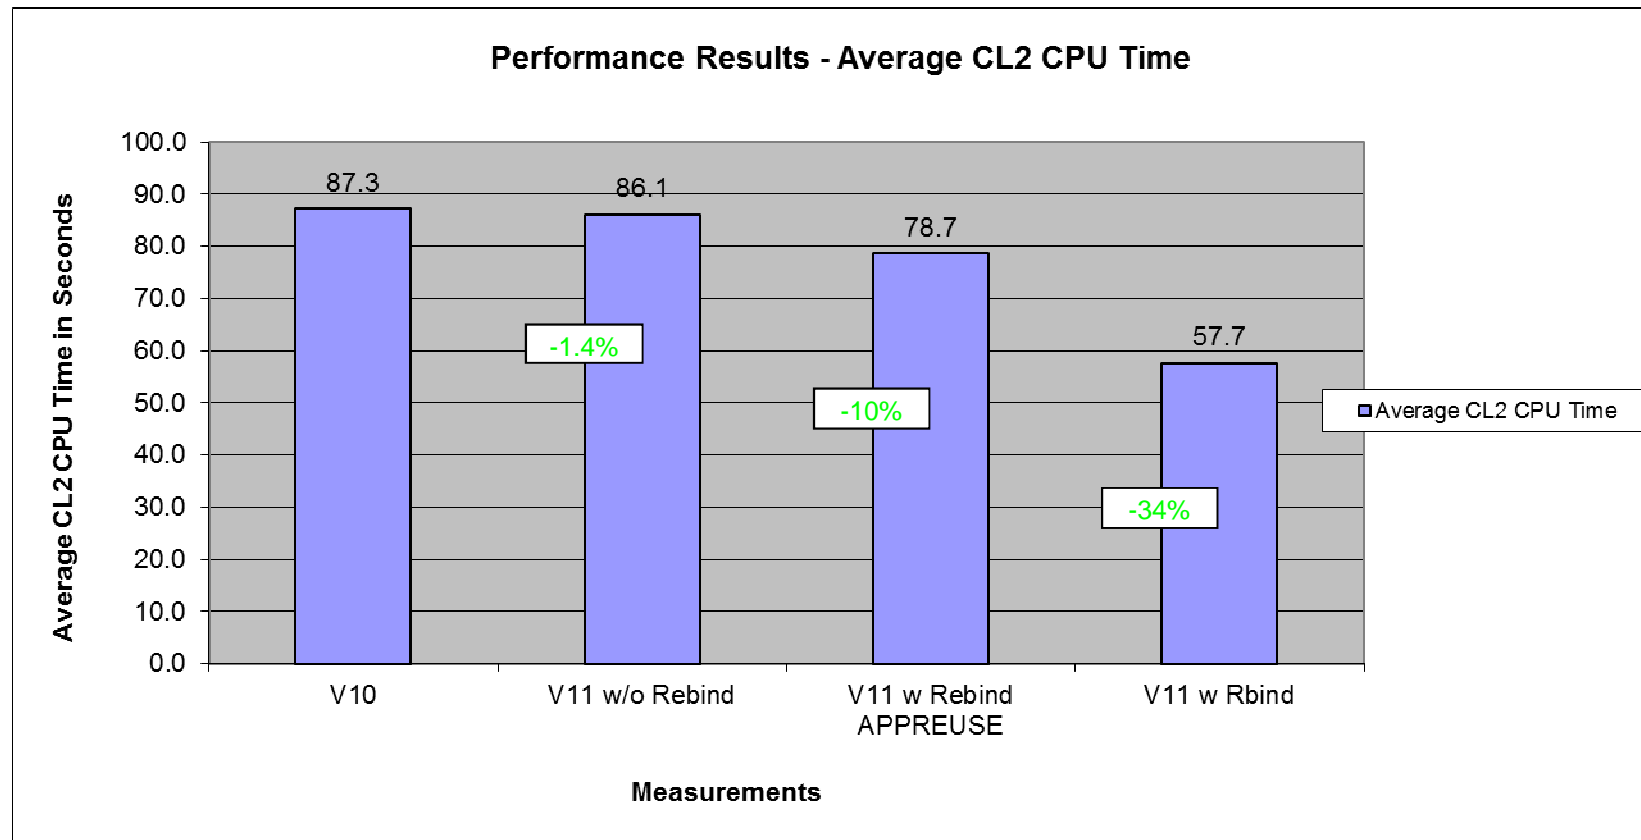
-10% out-of-box improvement with DB2 11 when rebinding with APREUSE
- 34% improvement in DB2 11 when rebinding to obtain
DB2 11 AP

Performance Improvements
no REBIND needed – Partial List

- DDF performance improvements
  - Reduced SRB scheduling on tcp/ip receive using new CommServer capabilities
  - Improved autocommit OLTP performance
- INSERT performance
  - Latch contention reduction
  - CPU reduction for Insert column processing and log record creation
  - Data sharing LRSN spin avoidance
  - Page fix/free avoidance in GBP write
- Automatic index pseudo delete cleanup
- IFI 306 filtering capabilities to improve Replication capture performance
- DGTT performance improvements
  - Avoid incremental binds for reduced cpu overhead *
- Utilities performance improvements
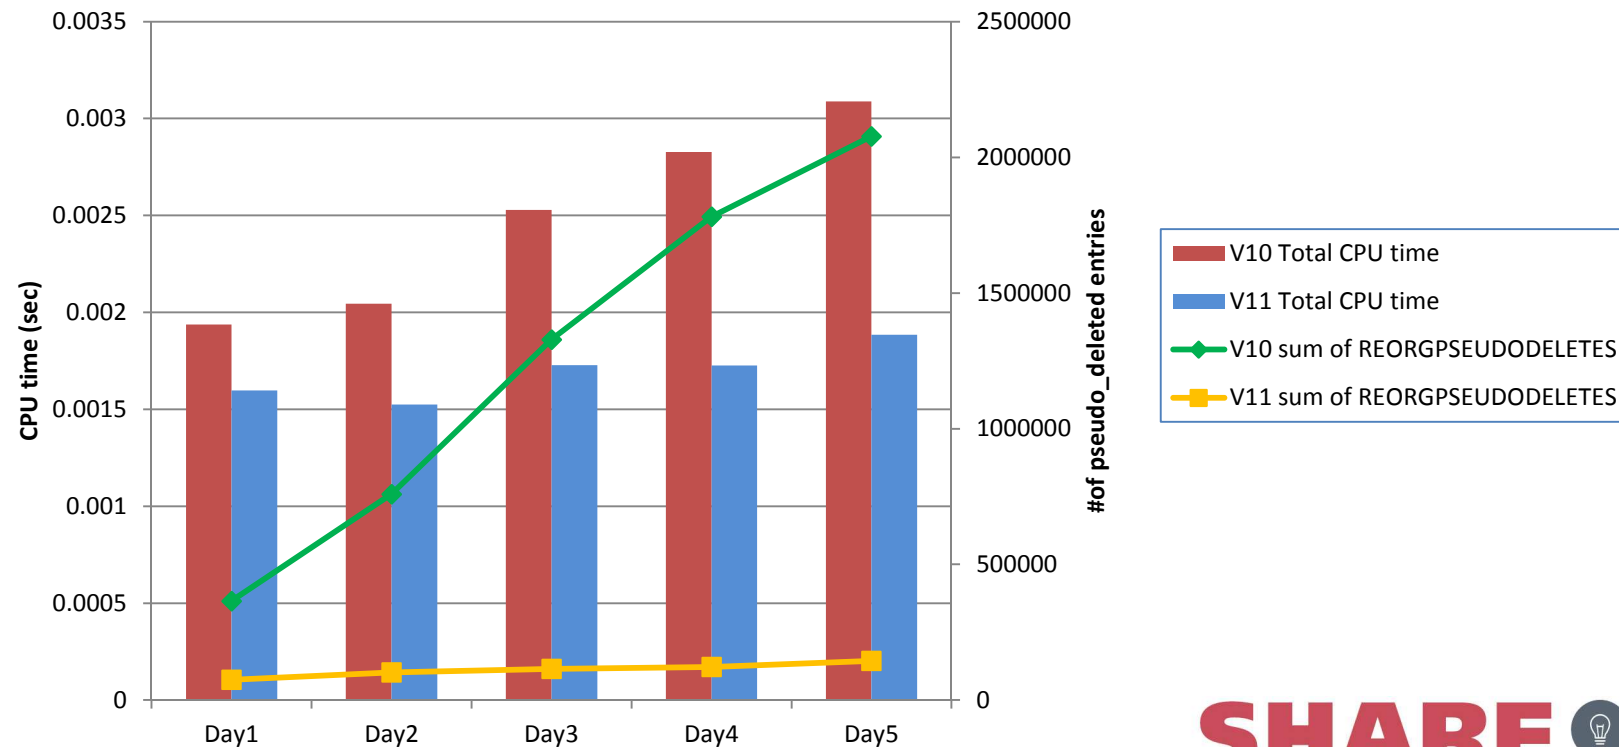- Java stored procedures: multi threaded JVMs, 64-bit JVM – more efficient

DB2 11 Auto Pseudo Delete Cleanup

- Up to 39% DB2 CPU reduction per transaction in DB2 11 compared to DB2 10
- Up to 93% reduction in Pseudo deleted entries in DB2 11
- Consistent performance and less need of REORG in DB2 11

WAS Portal Workload 5 Days Performance

SHARE in Orlando 2015

Performance Improvements
REBIND required – Partial List

- Query transformation improvements – less expertise required for performant SQL
- 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
- In-memory techniques
  - In-memory, reusable workfile
  - Sparse index (limited hash join support)
  - Non-correlated subquery using MXDTCACH
  - Correlated subquery caching
- 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 column processing
- DPSI performance improvements
- Data de-compression optimizations
- Optimizer CPU and I/O cost balancing improvements
- DRDA package based continuous block fetch

Performance Improvements
Sysprog, DBA, or appl effort required – Partial List

- 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 and CREATE INDEX performance
- New PCTFREE FOR UPDATE attribute to reduce indirect references
- DGTT performance improvements
  - Non logged DGTTs
- 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
- Open dataset limit raised to 200K

DB2 and z13 Synergy

- Increased Uniprocessor capacity
  - Up to 10% more capacity per core vs. zEC12
  - Up to 38% more capacity per core vs. z196
- Bigger/faster cache
- Increased System Capacity
  - Up to 40% more z/OS system capacity
  - Max LPARs increased from 60 to 85
  - Max usuableCPs increased from 101 to 141
  - System I/O bandwidth up to 832 GB/sec.
- Increased memory sizes
  - Up to 10TB per server
  - Up to 4 TB per z/OS LPAR
- Significant increase in zIIPprocessor capacity
  - zIIPs(and IFLs) are SMT enabled
  - Up to 1.4x capacity per zIIPcore vs. zEC12
- CFCC Level 20 allows for much larger DB2 CF structures

Buffer Pool Enhancements

- New frame size option for fixed pages
  - FRAMESIZE(2G)
    - Requires new z/OS and new hardware support for 2G frame size
- Pageable large pages
  - Allows FRAMESIZE(1M) PGFIX(NO)
  - Improves performance by reducing TLB misses
    - Requires new z/OS and new hardware support for pageable1M large pages
- New keywords to control the growth of buffer pool
  - With AUTOSIZE enabled
    - VPSIZEMIN specifies minimum number of buffers to allocate
    - VPSIZEMAX specifies maximum number of buffers to allocate

Buffer Pool Simulation

- Provides accurate estimation of I/O savings of simulated larger BP size
  - Intended for production environments as well as test/dev
- For more sophisticated BP analysis, a tool would still be required
- ALTER BUFFERPOOL command will support
- DB2 11 APAR PI22091
- Refer to DB2 11 for z/OS Managing Performance publication, SC19-4060

RAS and Usability Improvement Highlights

- Logging capacity and performance: RBA/LRSN optionally expands to 10 bytes
- BIND / DDL / Online REORG concurrency with persistent threads
  - Avoid having to shut down apps to get a REBIND through, e.g. for application upgrades
- More online schema changes
  - Alter partitioning limit keys
  - DROP column
  - Point in time recovery support for deferred schema changes
- Autonomics improvements
  - Automatic index pseudo delete cleanup
  - Overflow row reduction
  - Optimizer externalizes missing stats to enable automated RUNSTATS
- Data sharing improvements
  - Group buffer pool write-around
  - Restart light enhancements
  - Index split performance and other indexing improvements
  - Full LRSN spin avoidance
- Plan management improvements - APREUSE(WARN) support
- -ACCESS DATABASE … MODE(STATS) option to externalize RTS statistics

99.999% availability because your business never stops.[1]

Security Enhancements

- Remove inconsistencies between DB2 and RACF access controls
  - Automatic DB2 cache refresh when RACF changes are made
    - Package auth cache, dynamic statement cache, user authentication cache
  - Support BIND OWNER when using RACF exit
  - Support auto REBIND using owner's authid when usin RACF exit
  - Dynamic SQL authorisation checking improvements
- Bind plan option to ensure the program is authori use the plan
  - New PROGAUTH bind option
- Remove column masking restrictions for GROUP BY and DISTINCT

Summary of 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
  - REORG delete unused PBG datasets
  - System cloning improvements

- 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
  - REORG option to avoid sorting data for clustering
  - DSNACCOX performance

Over 40 new enhancements!

ROI

Key utilities performance numbers

- Up to 81% zIIP-eligible CPU with RUNSTATS COLGROUP
- Up to 40% zIIP-eligible CPU in REORG & LOAD with inline distribution stats
- REORG SWITCH phase outage reduced by up to 91%
- Up to 71% elapsed time reduction for REORG of subset of partitions
  - SORTNPSI option retrofitted to V9 & V10
- RECOVER from part-level image copies reduced CPU by up to 50%, elapsed by up to 40%
- LOAD from single input dataset elapsed time reduced by up to 70%
- Crossloader support for FETCH CONTINUE for LOB & XML data
  - 28% CPU reduction

SWITCH phase impact relief – reduced application impact

- Easier drain acquisition
- Prevent new claims on all target partitions whilst waiting for drains
  - Faster drain acquisition for part-level REORG
- New DRAIN_ALLPARTS option to momentarily drain all data parts
  - Eliminates claim-drain "deadlocks" for part-level REORG with NPSIs
- Restructure SWITCH phase processing for outage reduction
  - SWITCH phase ET reduction of 91% measured when reorging 20 parts
- New SWITCHTIME parameter to determine earliest point at which drain processing will be attempted
  - Govern timing of drain without the need to schedule separate –ALTER UTILITY command

SHARE in Orlando **2015**

Expanded SQL and Analytics Capabilities

- Global variables
- SQLPL improvements: array data type, autonomous transactions
- Alias support for sequence objects
- Temporal data enhancements
  - Support for views
  - Special register support
  - Integrated auditing support (planned)
- Transparent archive query
- SQL Grouping Sets, including Rollup, Cube
- Unicode column support for EBCDIC tables
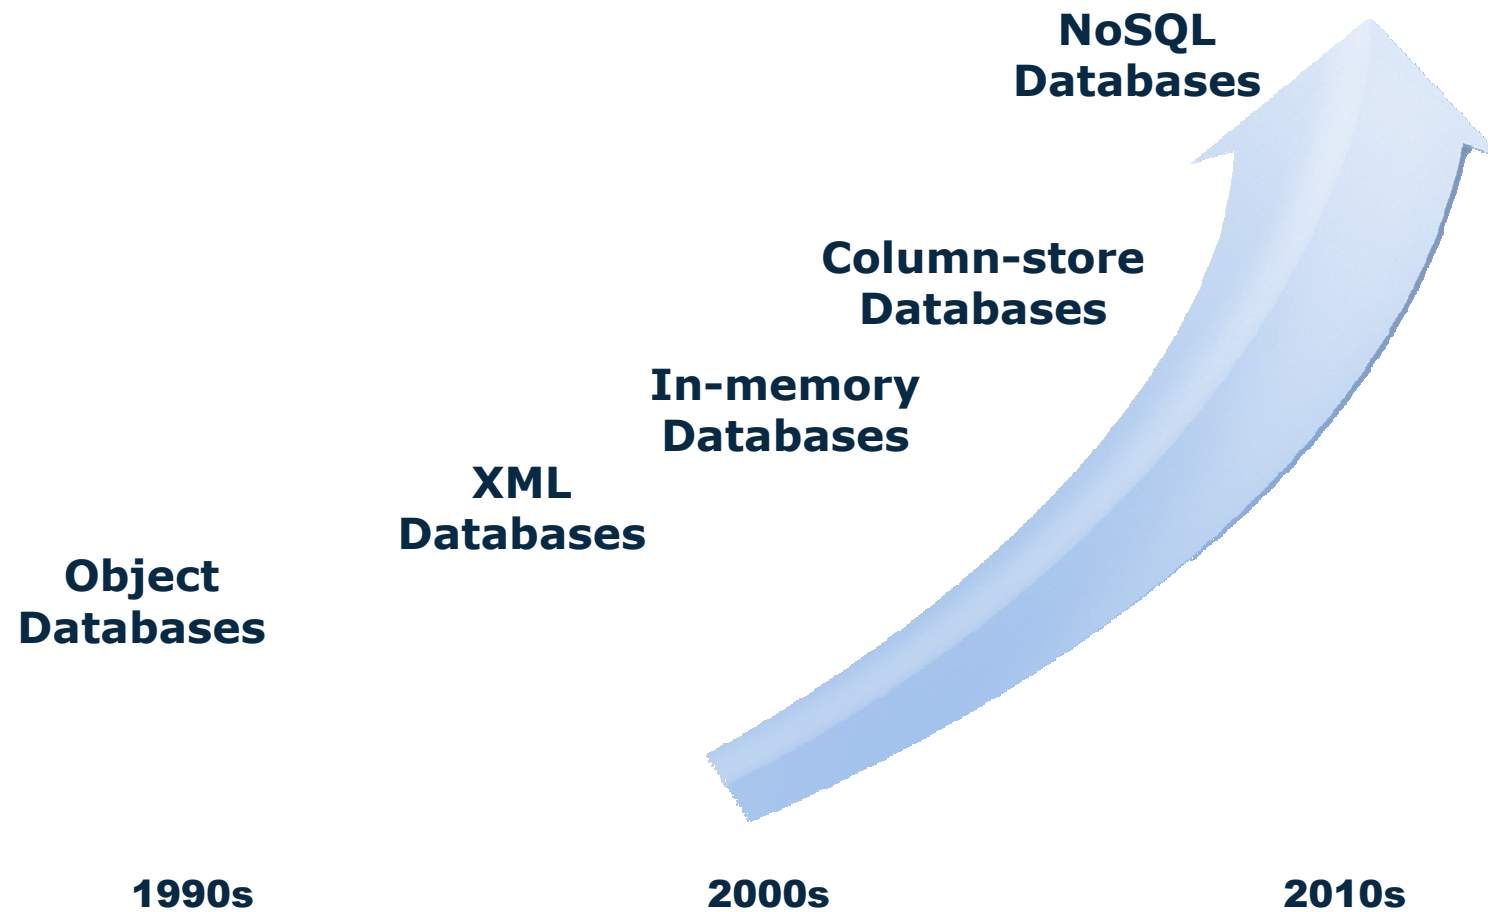- Hadoop access via table UDF
- JSON support

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
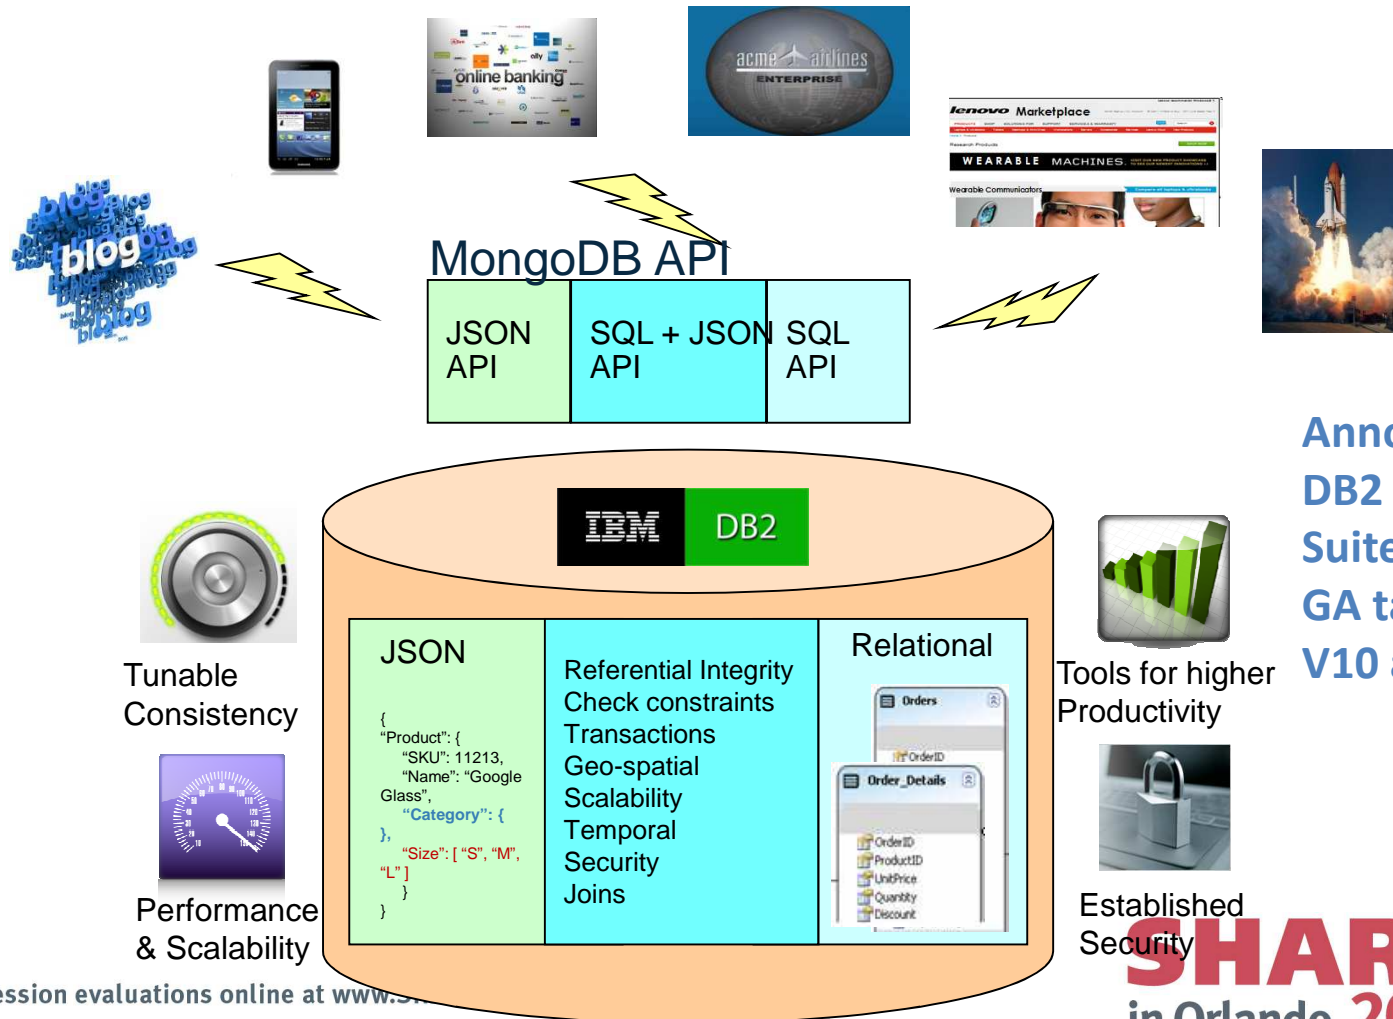- Move_To_Archive global variable allows DELETEs to be automatically archived

8/5/2015

New Technology Emerges

NoSQL
Databases

Column-store
Databases

In-memory
Databases

XML
Databases

Object
Databases

1990s   2000s   2010s

JSON Database Technology Preview
Providing the best of both worlds

MongoDB API

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" ]
    }
 }
```

Tunable Consistency

Performance & Scalability

Tools for higher Productivity

Established Security

Announced in DB2 Accessories Suite for z/OS, GA target Dec. 6 V10 and V11

Complete your session evaluations online at www.S

Why is System z Important for Big Data and Analytics?

Because the world's largest and most successful companies store their operational data on z

- Data that originates and/or resides on zEnterprise
    - 2/3 of business transactions for U.S. retail banks
    - 80% of world's corporate data
- Businesses that run on zEnterprise
    - 66 of the top 66 worldwide banks
    - 24 of the top 25 U.S. retailers
    - 10 of the top 10 global life/ health insurance providers
- The downtime of an application running on zEnterprise = apprx 5 minutes per yr
- 1,300+ ISVs run zEnterprise today
    - More than 275 of these selling over 800 applications on Linux

Complete your session evaluations online at www.SHARE.org/Orlando-Eval
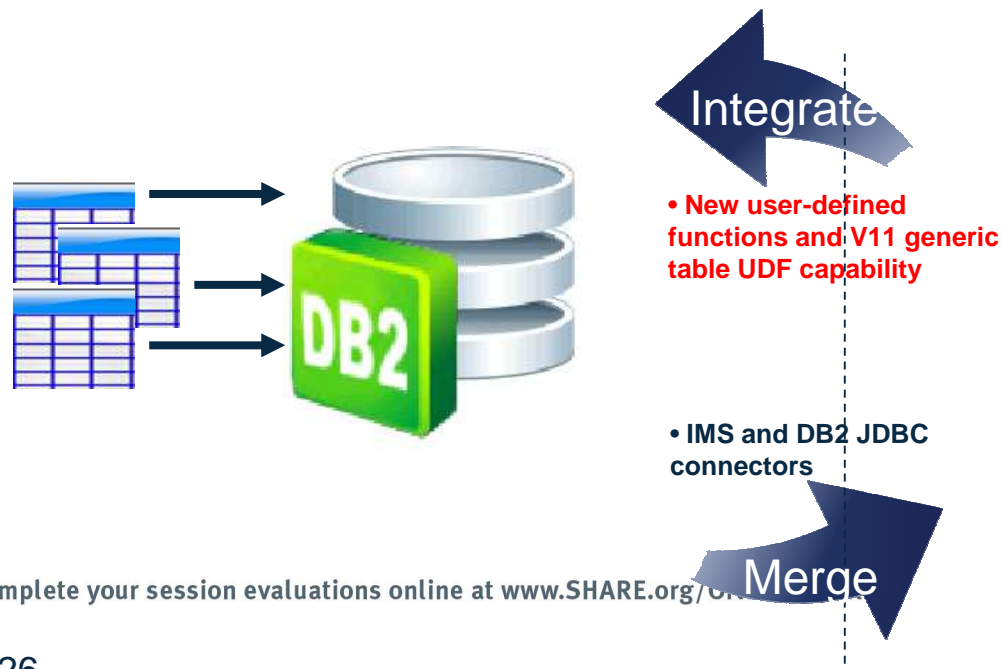
*Source IBM CEO Study 2011

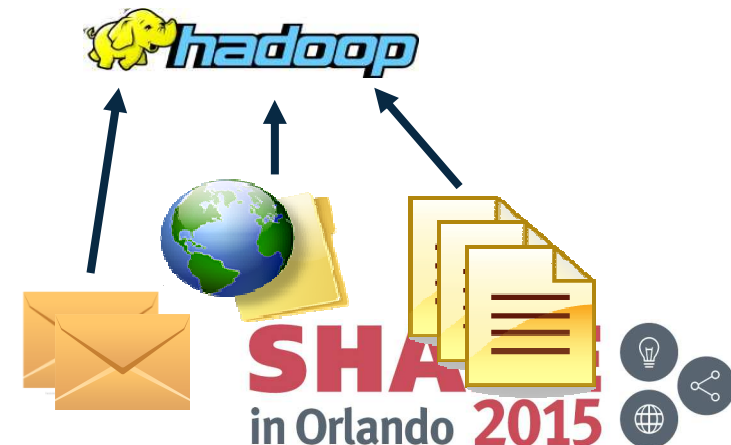# Integrating Big Data Analytics with DB2 for z/OS

New V11 features enable this

- 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
- 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 V11 generic table UDF capability

• IMS and DB2 JDBC connectors

Merge

XML Enhancements

- New Features
  - Basic xQuery (retrofit to v10)
  - COBOL samples for XML (published on Developerworks website)
- Feature 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 (APPLCOMPAT)
  - New feature to ease DB2 version upgrades – avoid impact to applications
  - New mechanism to identify applications affected by SQL changes in the new release
  - Seamless mechanism to make changes at an application (package) level or system level
- Faster ENFM processing
  - Lab measurement showed 18x faster in V11 vs. V10 using a large customer catalog
- Access path stability improvements
- Higher code quality stability levels
- SQL Capture/Replay tooling can help testing of DB2 version upgrades
- Migration Planning Workshops (MPW)
  - See the DB2 11 MPW community in DeveloperWorks for latest info

"We have seen some really good results regarding CPU savings - we have been so impressed with the product stability and have already moved an internal production system to DB2 11"
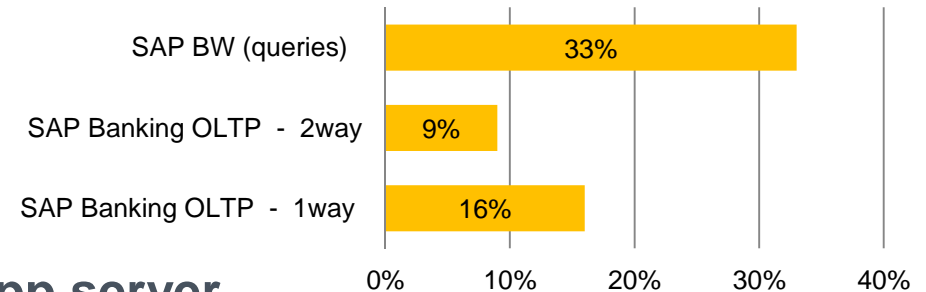
Stefan Korte   GAD

DB2 11 Optimized for SAP

Sample CPU reductions from DB2 10

- **Immediate SAP certification for DB2 11 at GA!**
  - See SAP Note 1850403
  - Easy migration from DB2 10:
    No new SAP service packs required
    Facilitated online DB2 migration
- **Save with CPU reductions**
- **Low latency connectivity from SAP app server**
- **Federated and consistent cloning of SAP business processes spanning multiple SAP/DB2 systems**
- **Online data maintenance**
  - Better online REORG, online repartitioning
- **Better scaling**
  - Larger log RBAs, larger statement cache

	CPU reduction
SAP BW (queries)	33%
SAP Banking OLTP - 2way	9%
SAP Banking OLTP - 1way	16%

0%   10%   20%   30%   40%

DB2 V8: 40+ features

DB2 9: 50+ features

DB2 10: 40+ features

DB2 11: 40+ features

DB2 11 SAP Certified at GA - the fastest ever certification for any DB2 for z/OS release in history

DB2 11 ESP Highlights

ESP Start	**February 2013**
First Code Drop	**March 2013**
"Regular" service process	**July 2013**
GA	**October 25, 2013**

Core - 21 WW Customers

Geography
- 11 EMEA
- 9 NA
- 1 SA

Industry
- 7 Banking
- 5 Insurance
- 3 Healthcare
- 2 Financial Markets
- 1 Automotive

Extended – 6 WW Customers

Geography
- 3 EMEA
- 2 NA
- 1 SA

Industry
- 3 Banking
- 2 Computer Services
- 1 Professional Services

in Orlando 2015

DB2 11 ESP Client Feedback

- Excellent quality and stability

- Good performance and CPU savings
  - ✓ DRDA workload up to 20% CPU reduction
  - ✓ CICS workload up to 18% CPU reduction
  - ✓ Batch workload up to 20% CPU reduction

- Full menu of functions, including
  - Utility improvements
  - Transparent archiving
  - Large RBA/LSRN
  - Optimizer and migration improvements
  - Big Data Integration
  - JSON Support for modern workloads

DB2 11 for z/OS - over 30 quotes

DB2 11 - SPEED & COST

"The Archive Transparency feature addresses an issue we have needed to resolve for a long time at the Bank and will reduce."

Paulo Sahadi, IT Executive Banco do Brasil

bankdata

BANCO DO BRASIL

"We have seen some really good results regarding CPU savings while running IMS-driven batch workload in our ESP test environment with DB2 11 CM/NFM - we have been so impressed with the **product stability** and have already moved an internal **production system to DB2 11**"

Stefan Korte   GAD

DB2 11 Early Support Program (ESP)
CPU savings, very high quality, production level stability

*"We have been involved in several DB2 for z/OS ESP's. This one will **rank as one of, if not the smoothest one yet**." – Large NA retailer*

*"Overall they are **very satisfied and astonished about the system stability** of DB2 V11. In V10 they experienced this in another way."*
– European Insurance

*"We have seen very few problems in [Installation, Migration, and Performance]. Overall, it has been a very pleasant experience!!…The **quality of the code is clearly much higher** than for the ESP for DB2 10…" - European Banking/FSS*

*"**Good code stability, no outages, no main failures, only a few PMRs….**" – European Banking*

SHARE in Orlando 2015

DB2 11 Early Support Program (ESP)
CPU savings, very high quality, production level stability

" *Higher availability, performance, lower CPU consumption amongst other new features were the benefits perceived by Banco do Brazil with DB2 11 for z/OS. During our testing with DB2 11 we noticed improved performance, along with stability.* " - Paulo Sahadi, IT Executive, Banco do Brasil

"We have seen some incredible performance results with DB2 11, a major reduction of CPU time, 3.5% before REBIND and nearly 5% after REBIND. This will significantly bring down our operating costs"
– Conrad Wolf, Golden Living

"I saw a significant performance improvement in recovery of catalog and directory. (V10 5:53 minutes, V11 2:50 minutes) That rocks! … **DB2 11 is the best version I have ever seen***."*
- European Gov't

*"**Overall, we have been impressed** with the new version of DB2."*
– NA Manufacturer

ESP Customer Experiences

- ## Stadtwerke Bielefeld GmbH
  - Major business benefits: Performance and SAP feature exploitation. Expecting to move to DB2 11 as soon as SAP certification complete
  - "The SAP IS-U unbilled revenue batch workload showed an elapsed time reduction of about 20% in Conversion Mode" -- Bernd Klawa, Stadtwerke Bielefeld DB2 DBA

- ## JN Data
  - Major business benefits: Operational enhancements, extended log addressing, DBA productivity improvements
  - "We love autonomics. DB2 11 has some really nice features for reducing the burden on the DBA" -- Frank Petersen, JN Data DB2 Systems Programmer

- ## BMW Group
  - Major business benefits: Forthcoming zEC12 upgrade will allow use of 2GB page frames, ability to break into persistent threads and undertake more dynamic schema change will help business agility
  - "Virtual storage isn't a big limitation for us any more, but we expect the CPU savings in DB2 11 to provide the major business benefit for us" -- BMW Group DB2 for z/OS Product Manager

DB2 11 Resources

- Information Center
- DB2 11 Technical Overview Redbook (SG24-8180)
  - Draft version available, final version coming soon.
- DB2 11 links: https://www.ibm.com/software/data/db2/zos/family/db211/
  - Link to DB2 11 Announcement Letter
  - Links to webcasts
  - Customer case studies
  - Whitepaper: "DB2 11 for z/OS: Unmatched Efficiency for Big Data and Analytics"
  - Whitepaper: "How DB2 11 for z/OS Can Help Reduce Total Cost of Ownership"
- DB2 11 Migration Planning Workshop
  - http://ibm.co/IIJxw8
- Free eBook available for download
  - http://ibm.co/160vQgM
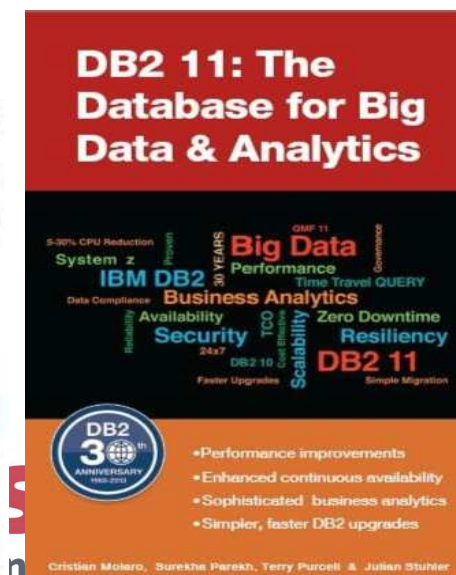
DB2 Cypress Themes

- In-memory processing
  - HW/SW integration into the future on z
  - Out-of-the-box performance improvement
- "Mobile-scale" data bases
  - More schema flexibility
  - Extreme scale tables, indexes
  - Higher data ingest rates
- Cloud enablement
  - Developer self-service, cloud-based provisioning, deployment
  - Self-optimizing system
    - More transparent SQL optimization
    - Temporal catalog for powerful problem diagnosis capabilities
    - Easier management of large tables
- Analytics and Big Data
- Extend System z leadership for continuous availability

DB2 Cypress Planning

- DB2 11 NFM is the prereq for migration
- Single phase catalog migration
- z/OS 2.1 or above. z196 hw or above
- No pre-V10 bound packages
- More memory = more performance
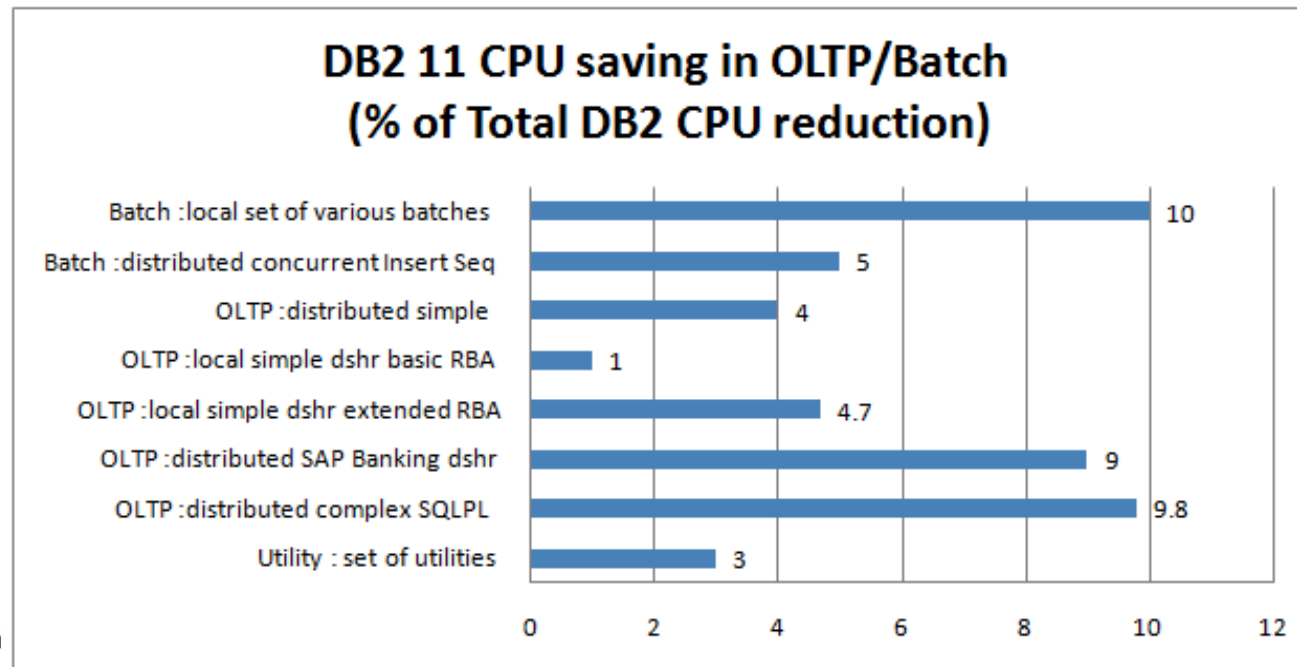- BRF deprecated
- Require BSDS conversion to new format

Thank You

SHARE

in Orlando 2015

DB2 11 OLTP/Batch Performance Expectations

- These are 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

Significant CPU Reduction in Query Workloads

DB2 11 Query Workloads - After REBIND w/o APREUSE
% of DB2 Class 2 CPU Reduction from DB2 10
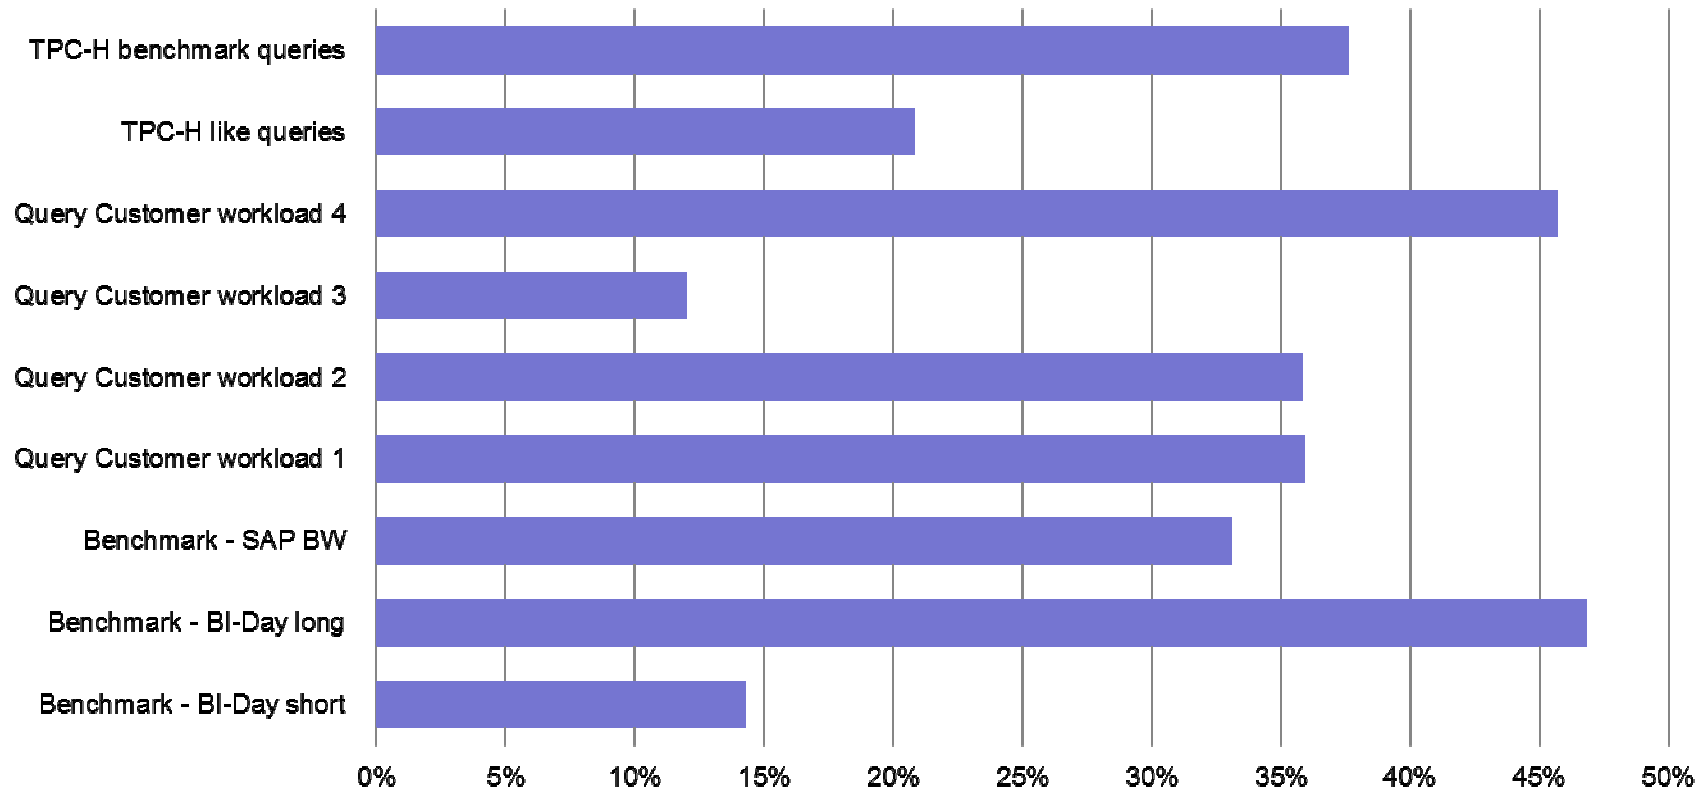
- Most performance improvements are also available with APREUSE
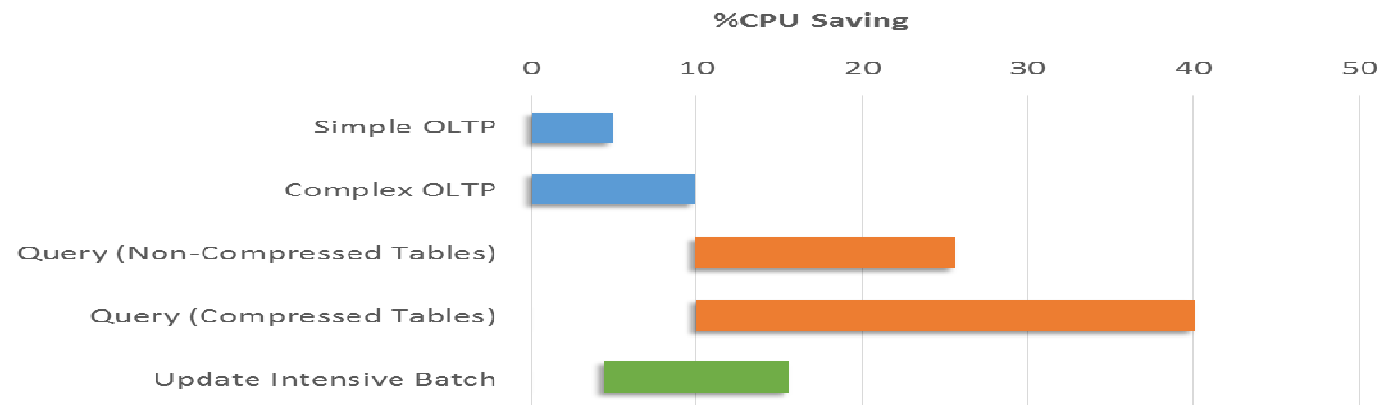- New and improved access path choices may be available without APREUSE

DB2 11 Affordable for Every Type of Workload

- **Out-of-the-box CPU Savings**

Expected Savings by Workload Type
DB2 11 Conversion Mode v DB2 10 New Function Mode

%CPU Saving

Workload Type	
Simple OLTP	~5
Complex OLTP	~10
Query (Non-Compressed Tables)	10–25
Query (Compressed Tables)	10–40
Update Intensive Batch	5–15

Axis: 0 10 20 30 40 50

REBIND may be needed for best performance

- **DB2 base LOAD and REORG inline statistics collection** now executed under enclave SRBs, so are **now zIIP eligible**

- **More potential savings with application or system changes**
  - Log replication capture
  - Data sharing using extended log record format
  - Up to 20-90% CPU savings from pureXML performance enhancements

c

DB2 11 – Foundation for Business Critical Analytics

DB2 for z/OS and
IBM DB2 Analytics Accelerator

LTP Transactions

eal time data ingestion

gh concurrency

DB2 Native
Processing

perational analytics

andard reports

LAP

mplex queries
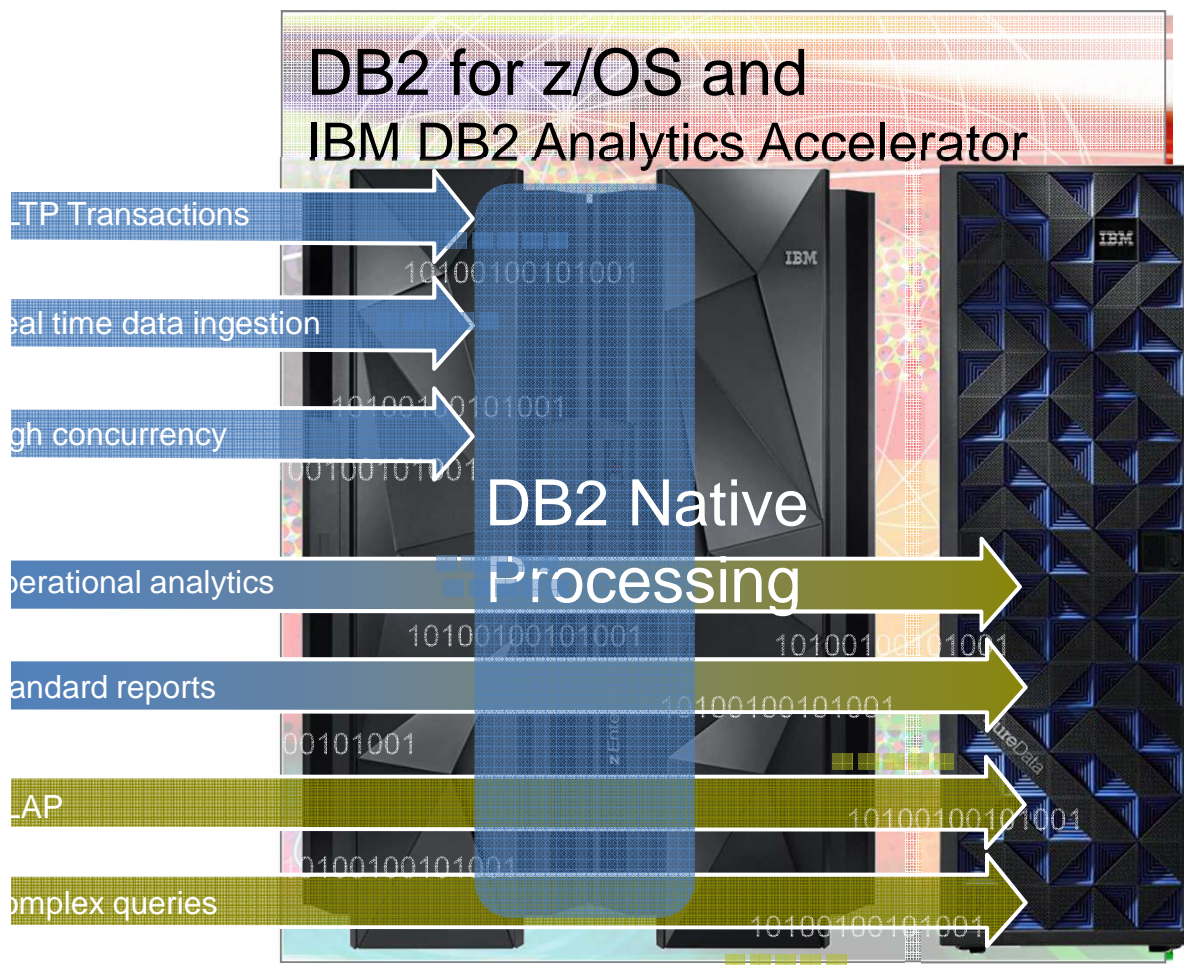
- DB2 11 CPU savings benefit query workloads with or without IDAA

Query Management Improvement Highlights

- 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)

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)
    - DB2 11 is the initial deployment of this capability
    - DB2 10 will be the lowest level of compatibility supported

QMF 11: Business Analytics for the System z Enterprise

QMF Analytics for TSO

- Brand new component available in QMF Enterprise Edition 11
- Delivers unprecedented charting and statistical analysis capabilities directly to the mainframe
- Completely menu driven

Faster up and running with QMF reporting

- Adhoc Reports and Quick Reports
- Allows users to quickly and easily create their own sophisticated reporting objects using an open canvas
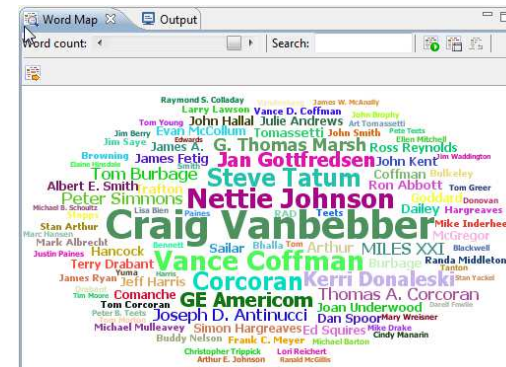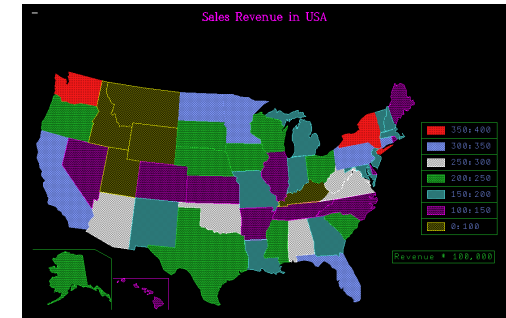
Analytics on unstructured data sources

- Text Analytics allows users to extract entities from unstructured data sources (either file-based or database-based) and display the results graphically

Increased support for the business user

- Dynamarts allow users to save their result sets with their query objects for offline use
- Mobile device support for iPad and Android tablets

iOS Orlando 2015