DB2 11 for z/OS Overview

Charles Lewis
DB2 Advisor,
IBM Client Technical Professional
lewisc@us.ibm.com

December 10, 2013

IBM

Agenda

- DB2 11 Early Support Program Overview

- DB2 11 Technical Overview

- DB2 Development Outlook

DB2 11
ESP Clients

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

Overall Client Feedback

- Excellent quality and stability

- Good performance and CPU savings

- Full menu of functions, including

    - Utility improvements

    - Transparent archiving

    - RBA/LRSN

    - Optimizer and migration improvements

ESP Performance Summary

13 customers have provided performance data

- Summary of the data analysis with comparable workloads and after rebind for static workloads

  ✓ DRDA workload 0 to 20% CPU reduction

  ✓ CICS workload 3 to 18% CPU reduction

  ✓ 18% includes 4% CPU reduction with NFM extended RBA/LRSN

  ✓ Batch workload 3 to 20% CPU reduction

0 major performance issues

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

6   SIRDUG 12/9/13

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 Brasil 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

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 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**

The next step in big data starts with IBM.
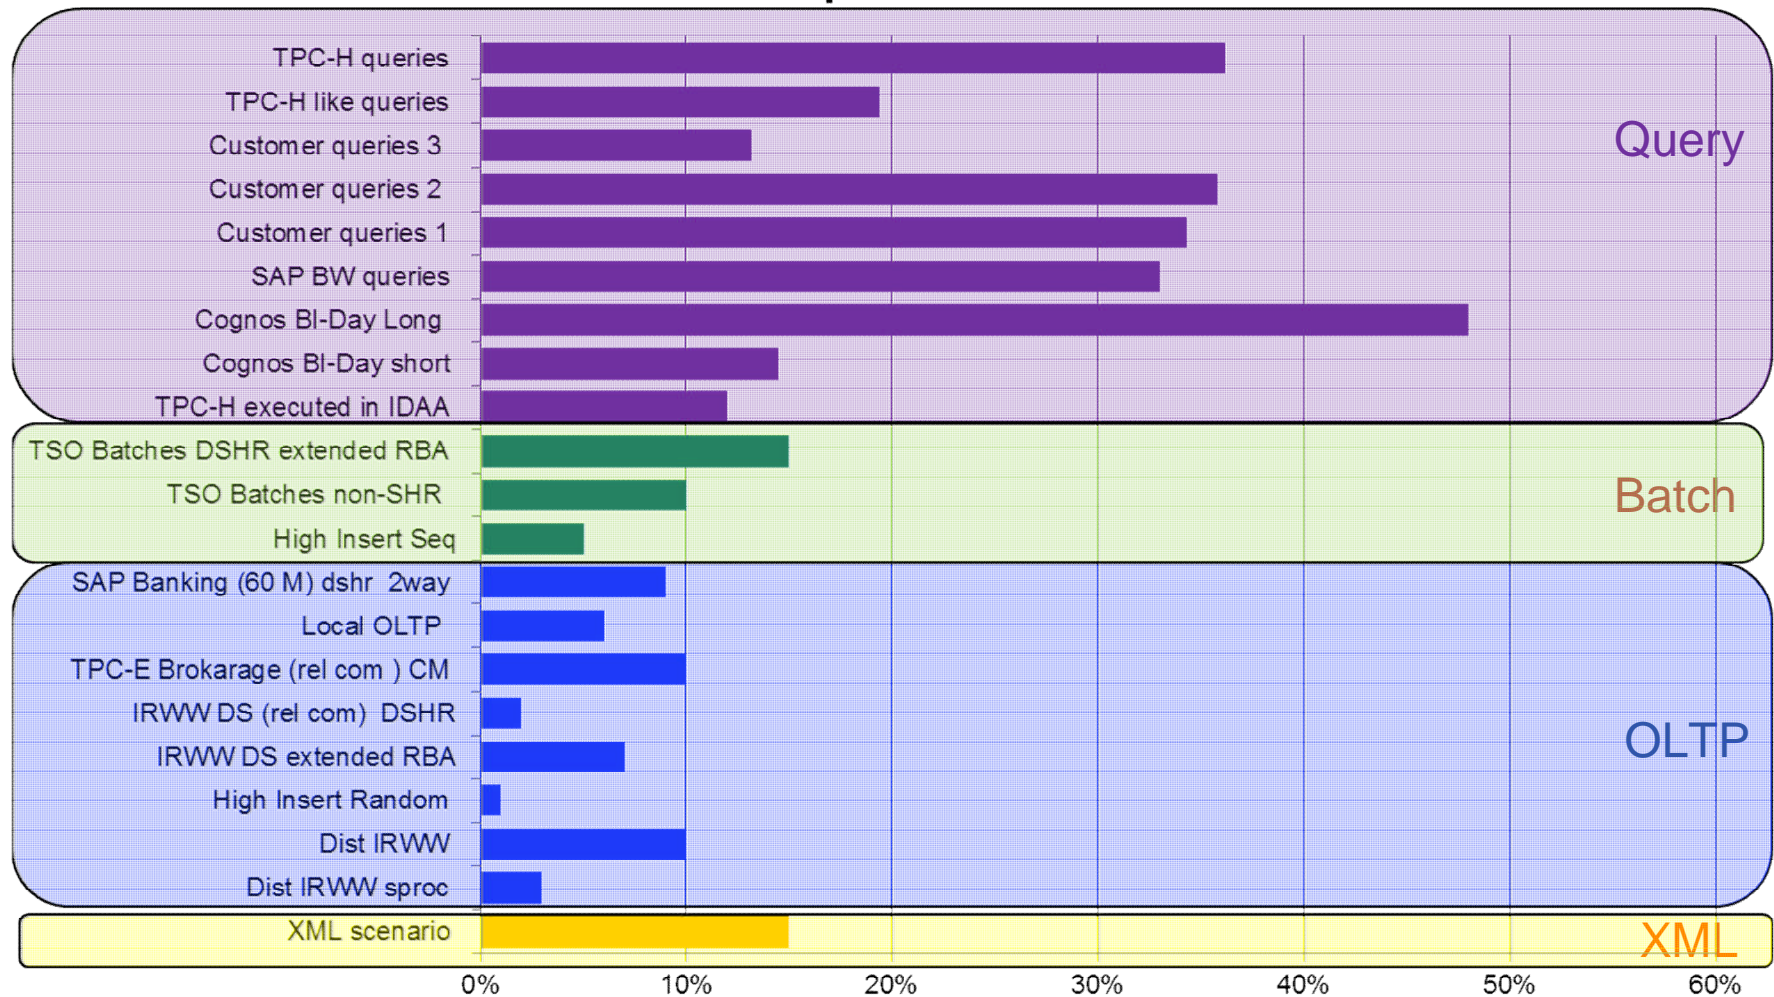
DB2 11 Major Themes

- **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
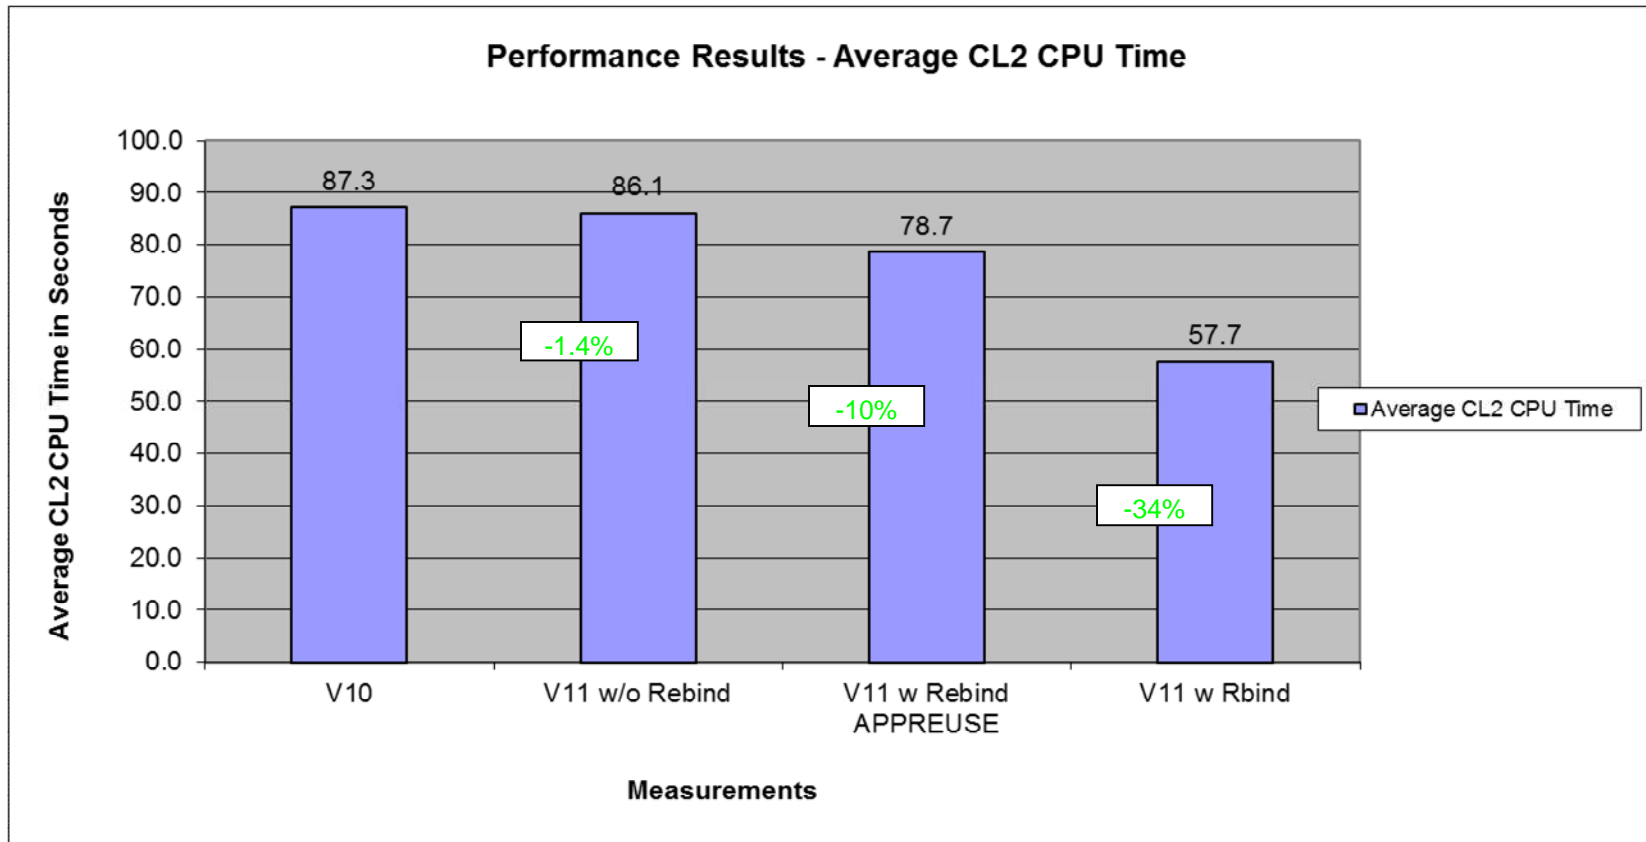 - Product quality/stability – raised the bar

Announce Oct. 1

GA Oct. 25

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

The next step in big data starts with IBM.

IBM

IBM Studies - Performance Quick Look

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

The next step in big data starts with IBM.

IBM

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

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
- DDF enhancement: DRDA package based continuous block fetch

The next step in big data starts with IBM.

Performance Improvements
Sysprog, DBA, or application 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

The next step in big data starts with IBM.

DB2 11 and zEC12 Synergy

- **Faster CPU – 1.25x compared to z196**
  - 5.5GHz processors, bigger/faster cache
  - 25% reduction measured with DB2 workloads
- **50% More System Capacity to help consolidation**
  - Up to 3TB real memory per server
  - Excellent synergy with DB2 10 and 11 scalability
- **New Features that DB2 11 Exploits**
  - FLASH Express and pageable 1MB frames, used for:
    - Buffer pool control blocks
    - DB2 executable code
  - 2GB frame support for buffer pools
    - Performance improvement expected for extremely large memory sizes
- **New zEC12 GA2 features that benefit DB2**
  - zEDC Express for enhanced DB2 SMF data compression
  - RoCE Express for faster, cheaper z/OS to z/OS DRDA communication
    - Preliminary measurements show up to 2x DRDA transaction throughput increase

The next step in big data starts with IBM.

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

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

- 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

The next step in big data starts with IBM.

IBM

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 using RACF exit
  - Dynamic SQL authorization checking improvements

- 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

The next step in big data starts with IBM.

IBM

Summary of Utilities Improvements

- Over 40 new enhancements!
- 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

The next step in big data starts with IBM.

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

The next step in big data starts with IBM.

IBM.

Integrating Big Data Analytics with DB2 for z/OS

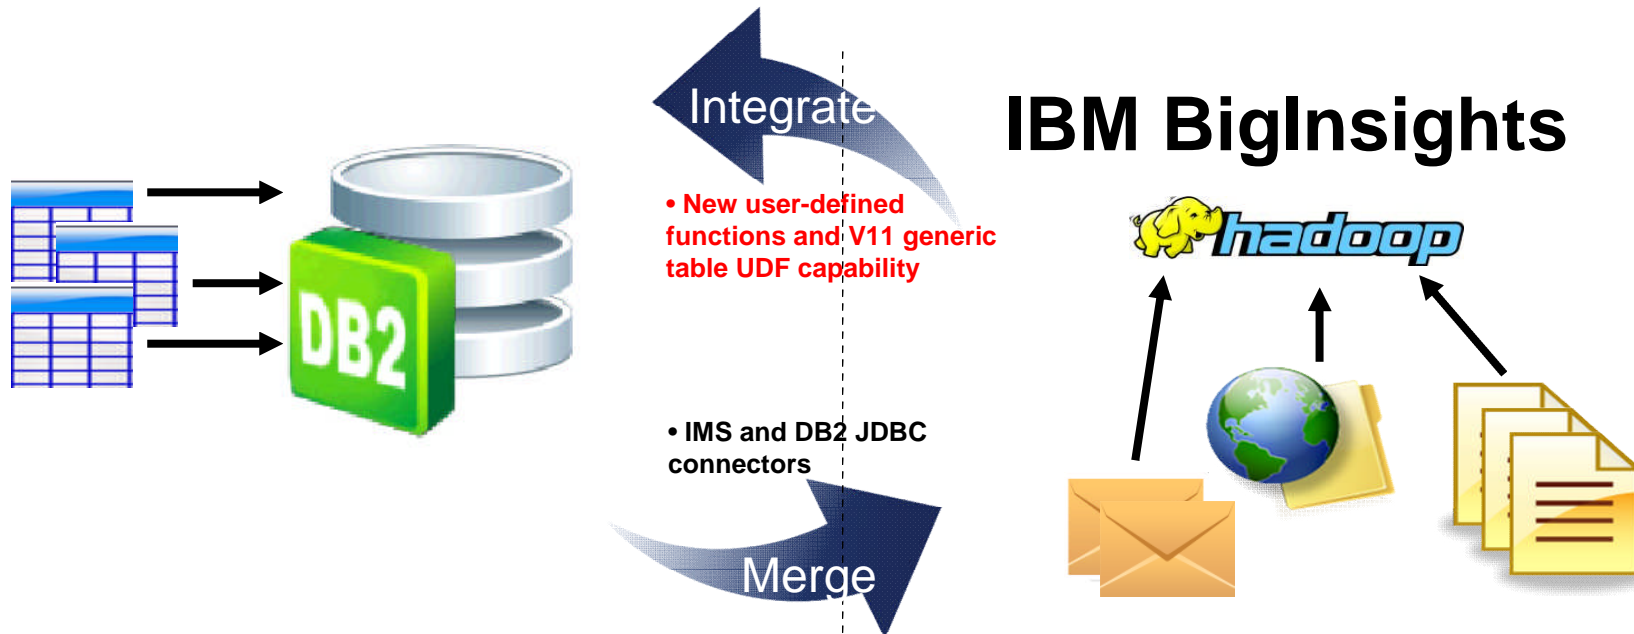| • **Much of the world's operational data resides on z/OS**	• **Unstructured data sources are growing fast**

▪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
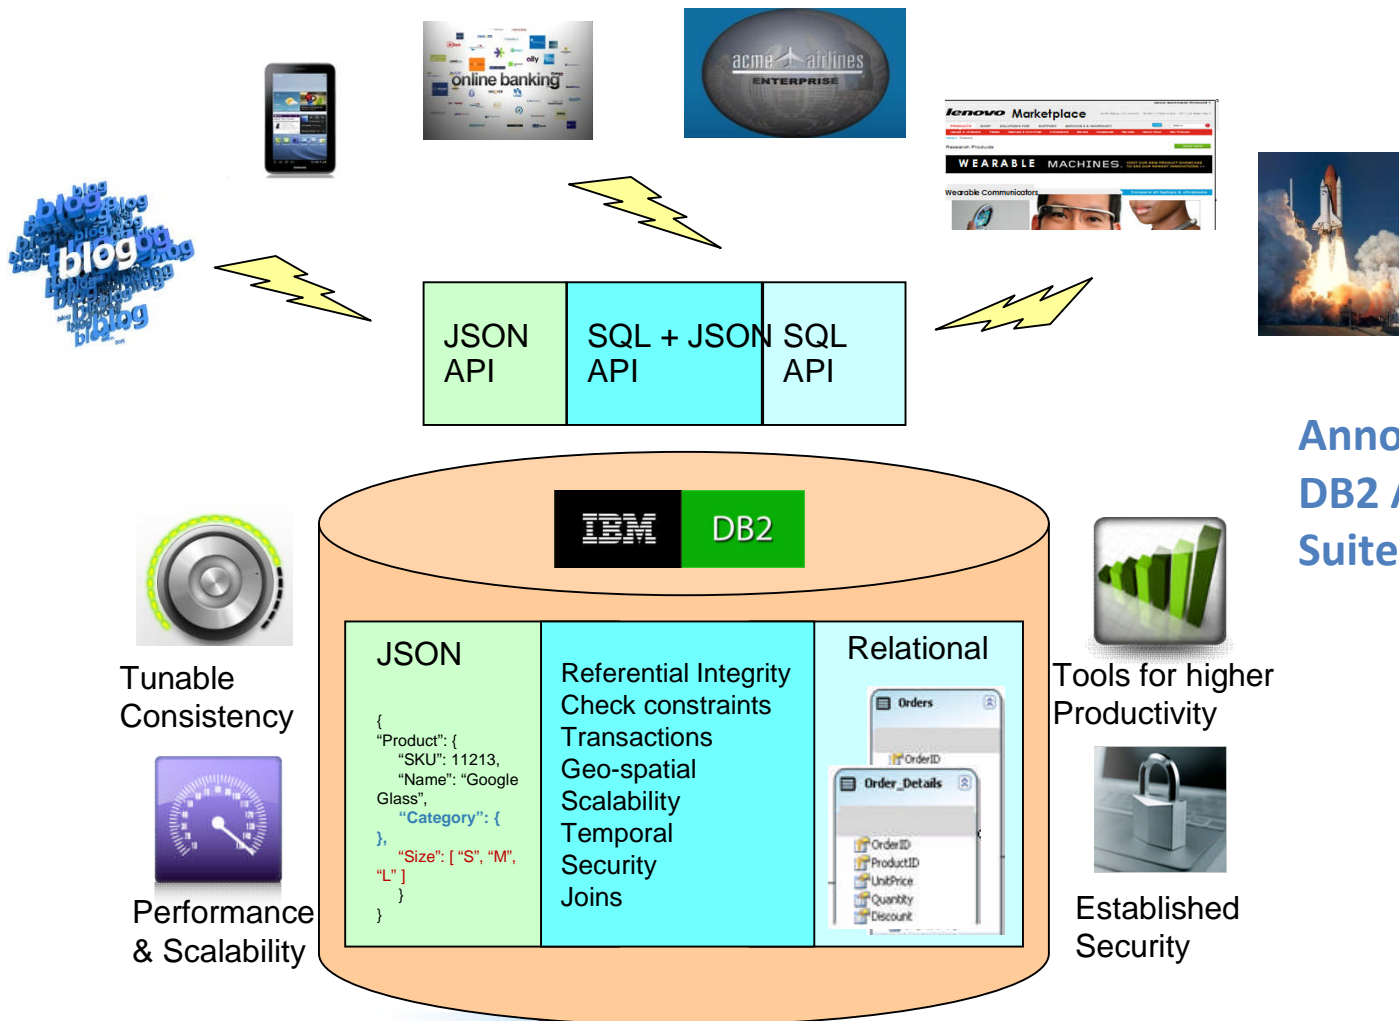
Integrate

IBM BigInsights

• **New user-defined functions and V11 generic table UDF capability**

• **IMS and DB2 JDBC connectors**

Merge

The next step in big data starts with IBM.

JSON Database Technology Preview
Providing the best of both worlds

JSON
API

SQL + JSON
API

SQL
API

Announced in DB2 Accessories Suite for z/OS

IBM DB2

JSON

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

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

Relational

Tunable
Consistency

Performance
& Scalability

Tools for higher
Productivity

Established
Security

The next step in big data starts with IBM.

IBM

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

The next step in big data starts with IBM.

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

"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 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)
  - No charge, 1-day education
  - DB2 11 MPW Community on DeveloperWorks

golden
living

"DB2 11 will provide enormous cost saving benefits across all DB2 shops"
Conrad Wolf

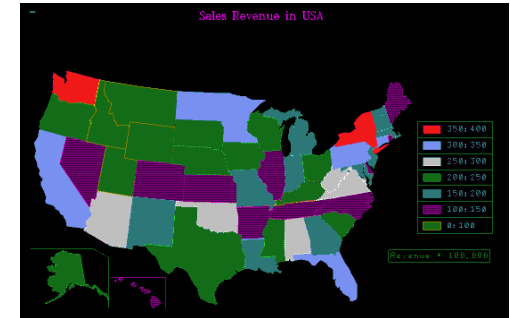The next step in big data starts with IBM.

IBM.

# 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
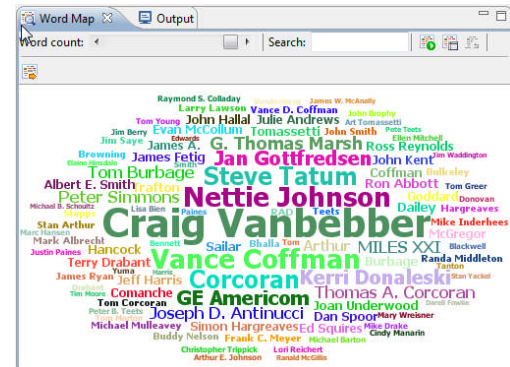
Analytics on unstructured data sources

- 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

The next step in big data starts with IBM.

IBM

DB2 11 Resources

- Information Center

- DB2 11 Announcement

  - http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS213-376&appname=USN

- DB2 11 Technical Overview Redbook (SG24-8180)

- DB2 11 links: https://www.ibm.com/software/data/db2/zos/family/db211/

  - Case study: BMW Group

  - Case study: JN Data

  - 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"

- eBook available at upcoming events

DB2 Cypress: Early Thoughts

- 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

- Self-optimizing system
  - More transparent SQL optimization
  - Temporal catalog for powerful problem diagnosis capabilities
  - Easier management of large tables

- Extend System z leadership for continuous availability

Join "The World of DB2 for z/OS" in 2013 (Nov or Dec) and Win Kindle Fire

Join   http://bit.ly/WorldofDB2

- **To celebrate 30th Anniversary of DB2 we are giving one lucky winner a Kindle Fire – profile picture must be added to qualify.**

The next step in big data starts with IBM.

IBM.